

GENERALIZED FINITE ALGORITHMS FOR CONSTRUCTING HERMITIAN MATRICES WITH PRESCRIBED DIAGONAL AND SPECTRUM*

INDERJIT S. DHILLON^{†¶}, ROBERT W. HEATH JR.[‡], MÁTYÁS A. SUSTIK[†], AND
JOEL A. TROPP[§]

Abstract. In this paper, we present new algorithms that can replace the diagonal entries of a Hermitian matrix by any set of diagonal entries that majorize the original set without altering the eigenvalues of the matrix. They perform this feat by applying a sequence of $(N - 1)$ or fewer plane rotations, where N is the dimension of the matrix. Both the Bendel–Mickey and the Chan–Li algorithms are special cases of the proposed procedures. Using the fact that a positive semidefinite matrix can always be factored as X^*X , we also provide more efficient versions of the algorithms that can directly construct factors with specified singular values and column norms. We conclude with some open problems related to the construction of Hermitian matrices with joint diagonal and spectral properties.

Key words. Hermitian matrices, majorization, Schur–Horn theorem, inverse eigenvalue problem, plane rotations, finite-step algorithms

AMS subject classifications. 15A18, 15A57, 65F18

DOI. 10.1137/S0895479803438183

1. Introduction. It is sometimes of interest to construct a collection of Hermitian matrices that have specified diagonal elements and eigenvalues. When all the eigenvalues are nonnegative, the problem is essentially equivalent to constructing a collection of rectangular matrices with specified column norms and singular values. In particular, if a rectangular matrix X has requirements on its singular values and squared column norms, the Hermitian matrix X^*X has a corresponding requirement on its eigenvalues and diagonal entries.

A specific example of this problem is to construct Hermitian matrices with unit diagonal and prescribed nonnegative eigenvalues. Such matrices are called correlation matrices—Davies and Higham [6] discuss several applications that require such matrices, ranging from the generation of test matrices for eigenvalue solvers to the design of statistical experiments. A related matrix construction problem has also arisen in connection with wireless communications. It turns out that $d \times N$ matrices, $d < N$, with d identical nonzero singular values and with prescribed column norms satisfy a certain “sum capacity” bound and “minimum squared correlation” property that is important in wireless applications. These matrices only exist if a majorization condition holds, as discussed in section 2.1. For a quick introduction to squared correlation; see the appendix in [7]. Refer to [20, 22] for details on how the problem arises in wireless applications.

*Received by the editors November 21, 2003; accepted for publication (in revised form) by U. Helmke September 24, 2004; published electronically June 22, 2005.

<http://www.siam.org/journals/simax/27-1/43818.html>

[†]Department of Computer Sciences, The University of Texas, Austin, TX 78712 (inderjit@cs.utexas.edu, sustik@cs.utexas.edu).

[‡]Department of Electrical and Computer Engineering, The University of Texas, Austin, TX 78712 (rheath@ece.utexas.edu).

[§]Institute for Computational Engineering and Sciences (ICES), The University of Texas, Austin, TX 78712 (jtropp@ices.utexas.edu).

[¶]The research of this author was supported by NSF CAREER Award ACI-0093404.

Two finite step techniques, the Bendel–Mickey [2] and Chan–Li [4] algorithms, are available for special cases. Both algorithms apply a sequence of plane rotations to an initial matrix that change its diagonal entries while preserving its spectrum. The Chan–Li algorithm starts with the diagonal matrix of eigenvalues and can reach a real, symmetric matrix with a specified majorizing diagonal. On the other hand, the Bendel–Mickey algorithm can start with an arbitrary Hermitian matrix and transform it to a Hermitian matrix with *equal* diagonal entries.

In this paper, we present new algorithms that generalize the Chan–Li and Bendel–Mickey procedures so that we can start with an arbitrary Hermitian matrix and change its diagonal entries to specified values while retaining its original spectrum. The only requirement is that the new diagonal elements majorize (in essence, average) the original ones. Thus our generalized algorithms permit us to construct a class of Hermitian matrices satisfying spectral and diagonal constraints that is richer than the collection yielded by earlier algorithms.

We now give a brief outline of the paper. In section 2, we provide the necessary background and summarize previous results. We present our generalized algorithms in section 3, and section 4 contains some numerical examples. We conclude in section 5 with some open problems.

2. Background and related work.

2.1. Majorization. Majorization is a relation between vectors that appears in a striking number of apparently unrelated contexts. Lorenz originally developed the ordering for econometrics, where he used it to compare the equitability of income distributions [1]. An intuitive definition is that one vector majorizes another if the former has “more average” entries than the latter. Let us make this notion precise.

DEFINITION 2.1. *Let \mathbf{a} be a real, N -dimensional vector, and denote its k th smallest component by $a_{(k)}$. This number is called the k th order statistic of \mathbf{a} .*

DEFINITION 2.2. *Let \mathbf{a} and \mathbf{z} be real N -dimensional vectors, and suppose that their order statistics satisfy the following relationships:*

$$\begin{aligned} a_{(1)} &\leq z_{(1)}, \\ a_{(1)} + a_{(2)} &\leq z_{(1)} + z_{(2)}, \\ &\vdots \\ a_{(1)} + a_{(2)} + \cdots + a_{(N-1)} &\leq z_{(1)} + z_{(2)} + \cdots + z_{(N-1)}, \quad \text{and also} \\ a_{(1)} + a_{(2)} + \cdots + a_{(N)} &= z_{(1)} + z_{(2)} + \cdots + z_{(N)}. \end{aligned}$$

Then we say that \mathbf{z} majorizes \mathbf{a} , and we write $\mathbf{z} \succcurlyeq \mathbf{a}$. If each of the inequalities is strict, then \mathbf{z} strictly majorizes \mathbf{a} , and we write $\mathbf{z} \succ \mathbf{a}$.

The inequality in the above definition is reversed in some treatments, but we follow the convention in Horn and Johnson’s book [13, p. 192]. An equivalent definition is that $\mathbf{z} \succcurlyeq \mathbf{a}$ if and only if $\mathbf{z} = M\mathbf{a}$ for some doubly-stochastic matrix M . Birkhoff’s theorem states that the collection of doubly-stochastic matrices of size N is identical with the convex hull of the permutation matrices having size N . It follows that those vectors which majorize a fixed vector form a compact, convex set; see [13, p. 197 and p. 527] and [18] for more details.

It is easy to verify that majorization is antisymmetric and transitive on \mathbb{R}^N . But the majorization relation is reflexive, and hence a partial ordering, only when we restrict our attention to nondecreasing vectors in \mathbb{R}^N [3].

Majorization plays a role on our stage because it defines the precise relationship between the diagonal entries and eigenvalues of a Hermitian matrix.

THEOREM 2.3 (Schur–Horn [13, pp. 193–196]). *The diagonal entries of a Hermitian matrix majorize its eigenvalues. Conversely, if $\mathbf{a} \succcurlyeq \boldsymbol{\lambda}$, then there exists a Hermitian matrix with diagonal entries listed by \mathbf{a} and eigenvalues listed by $\boldsymbol{\lambda}$.*

I. Schur demonstrated the necessity of the majorization condition in 1923, and A. Horn proved the converse some thirty years later [12]. Horn’s original proof is quite complicated and a small cottage industry has grown up to produce simpler, more constructive arguments; see, for example, [4, 5, 16]. A comprehensive reference on majorization is Marshall and Olkin’s monograph [18].

2.2. Some posets. First, we define some concepts related to partial orderings, and then we develop some new partial orderings on Hermitian matrices that are closely related to the matrix construction problem.

DEFINITION 2.4. *A set S equipped with a partial ordering \succcurlyeq is called a poset. It is denoted as (S, \succcurlyeq) . Two elements $a, b \in S$ are comparable if and only if $a \succcurlyeq b$ or $b \succcurlyeq a$. Any totally ordered subset of a poset is called a chain. Every pair of elements in a chain is comparable.*

We may equip any poset with the ordering topology, where each basic open set is given by $\{a \neq b : b \succ a\}$ for some point b . This is the minimal topology in which the ordering is continuous [15, p. 57].

Observe that vectors are comparable by majorization only when their entries have the same sum. Let \mathcal{C}_α denote the set of N -dimensional nondecreasing vectors whose entries sum to α . Every \mathcal{C}_α has a unique majorization-maximal element: the constant vector with entries α/N . On the other hand, there are no minimal vectors under the majorization relation; every \mathbf{z} succeeds an infinite number of other vectors.

DEFINITION 2.5. *We say that two Hermitian matrices A and B are Schur–Horn equivalent if and only if they have identical spectra and identical diagonals (up to permutation). We write $A \equiv B$ and we use $[A]$ to denote the equivalence classes induced by this relation.*

We indicate the collection of Schur–Horn equivalence classes by \mathcal{H} . Notice that the members of \mathcal{H} vary significantly. For example, the Schur–Horn equivalence class of a diagonal matrix is the set of diagonal matrices with the same entries in permuted order. Meanwhile, the equivalence class of a matrix with unit diagonal and nonnegative eigenvalues $\boldsymbol{\lambda}$ is the set of “correlation matrices” that have spectrum $\boldsymbol{\lambda}$ [6, 17]. Even though similarity transformations preserve the eigenvalues of a Hermitian matrix, very few simultaneously preserve the diagonal. Therefore, Schur–Horn equivalence classes are not stable under most transformations. Exceptions include symmetric permutations and diagonal similarity transforms.

DEFINITION 2.6. *For any two elements of \mathcal{H} , $[A]$ and $[Z]$, we say that $[Z] \succcurlyeq [A]$ if and only if the two matrices have the same spectrum and $\text{diag } Z \succcurlyeq \text{diag } A$.*

It is not hard to check that this construction yields a well-defined partial ordering on \mathcal{H} . Clearly, two Schur–Horn equivalence classes are comparable only if their members have the same spectrum. Suppose that the entries of $\boldsymbol{\lambda} \in \mathbb{R}^N$ already occur in nondecreasing order, viz. $\lambda_k = \lambda_{(k)}$ for each k . Then we may write $\mathcal{H}_\boldsymbol{\lambda}$ to denote the elements of \mathcal{H} with spectrum $\boldsymbol{\lambda}$. Each $\mathcal{H}_\boldsymbol{\lambda}$ forms an isolated component of the poset $(\mathcal{H}, \succcurlyeq)$, and it has a unique maximal element: the equivalence class of matrices with eigenvalues $\boldsymbol{\lambda}$ and with a constant diagonal. A significant difference between majorization and the matrix ordering is that every chain under the matrix ordering has a minimal element: $[\text{diag } \boldsymbol{\lambda}]$, where $\boldsymbol{\lambda}$ lists the (common) eigenvalues of the members

of the chain.

2.3. Algorithms. Now we discuss two algorithms which have been proposed for constructing Hermitian matrices with diagonal and spectral properties. In what follows, we use \mathbb{M}_N to denote the set of complex $N \times N$ matrices and $\mathbb{M}_{d,N}$ to denote the set of complex $d \times N$ matrices.

The Bendel–Mickey algorithm produces random (Hermitian) correlation matrices with given spectrum [2]. Suppose that $A \in \mathbb{M}_N$ is a Hermitian matrix with $\text{Tr } A = N$. If A does not have a unit diagonal, we can locate two diagonal elements so that $A_{jj} < 1 < A_{kk}$; otherwise, the trace condition would be violated. It is then possible to construct a real rotation Q in the jk -plane for which $(Q^*AQ)_{jj} = 1$. The transformation $A \mapsto Q^*AQ$ preserves the conjugate symmetry and the spectrum of A , but it reduces the number of nonunit diagonal entries by at least one. Therefore, at most $(N - 1)$ rotations are required before the resulting matrix has a unit diagonal. If the output matrix is Z , it follows that $[Z] \succcurlyeq [A]$. Indeed, $[Z]$ is the unique \succcurlyeq -maximal element in every chain that contains $[A]$.

The Chan–Li algorithm, on the other hand, was developed as a constructive proof of the Schur–Horn theorem [4]. Suppose that $\mathbf{a} \succcurlyeq \boldsymbol{\lambda}$. The Chan–Li algorithm begins with the diagonal matrix $\Lambda \stackrel{\text{def}}{=} \text{diag } \boldsymbol{\lambda}$. Then it applies a sequence of $(N - 1)$ cleverly chosen (real) plane rotations to generate a real, symmetric matrix A with the same eigenvalues as Λ but with diagonal entries listed by \mathbf{a} . Once again, the output and input satisfy the relationship $[A] \succcurlyeq [\Lambda]$. While the Bendel–Mickey algorithm starts from any element of a chain and moves to the top; the Chan–Li algorithm starts at the bottom of a chain and moves upward.

The Bendel–Mickey algorithm is a surjective map from the set of Hermitian matrices with spectrum $\boldsymbol{\lambda}$ onto the set of correlation matrices with spectrum $\boldsymbol{\lambda}$. If the initial matrix is chosen uniformly at random (which may be accomplished with standard techniques [19]), the result may be construed as a random correlation matrix. The distribution of the output, however, is unknown [11]. On the other hand, due to the special form of the initial matrix and the rigid choice of rotations, the Chan–Li algorithm cannot construct very many distinct matrices with a specified diagonal. It would be interesting to develop a procedure which can calculate every member of a given equivalence class; see section 5 for a brief discussion.

A brief discussion of how to use plane rotations to equalize the diagonal entries of a Hermitian matrix appears on page 77 of Horn and Johnson [13]. Problems 8.4.1 and 8.4.2 of Golub and Van Loan outline the Bendel–Mickey algorithm [10].

Davies and Higham present a numerically stable version of the Bendel–Mickey algorithm in their article [6]. Other references on this topic include [14, 23].

3. Generalized algorithms. We propose methods that generalize the Bendel–Mickey and Chan–Li algorithms. Like them, our techniques use a sequence of $(N - 1)$ or fewer plane rotations to move upward between two points in a chain. The crux of the matter is the strategy for selecting the planes of rotation. The two methods we present can be viewed, respectively, as direct generalizations of the Chan–Li strategy and the Bendel–Mickey strategy. Unlike the earlier algorithms, these new techniques do not require ending at the top of a chain like Bendel–Mickey nor starting at the bottom like Chan–Li. Therefore, our techniques allow the construction of a much larger set of matrices than the Chan–Li algorithm, while retaining its ability to select the final diagonal entries.

3.1. Plane rotations. First, we discuss how to use a plane rotation to modify the diagonal entries while preserving the spectrum. Let us suppose that A is a 2×2 matrix with diagonal \mathbf{a} that we wish to transform to \mathbf{z} , where $\mathbf{z} \succcurlyeq \mathbf{a}$. Without loss of generality, we can assume that $a_1 \leq z_1 \leq z_2 \leq a_2$. We will explicitly construct a real plane rotation Q so that the diagonal of Q^*AQ equals \mathbf{z} . Recall that a two-dimensional plane rotation is an orthogonal matrix of the form

$$Q = \begin{bmatrix} c & s \\ -s & c \end{bmatrix},$$

where $c^2 + s^2 = 1$ [10]. The desired plane rotation yields the matrix equation

$$\begin{bmatrix} c & s \\ -s & c \end{bmatrix}^* \begin{bmatrix} a_1 & a_{21}^* \\ a_{21} & a_2 \end{bmatrix} \begin{bmatrix} c & s \\ -s & c \end{bmatrix} = \begin{bmatrix} z_1 & z_{21}^* \\ z_{21} & \tilde{z}_2 \end{bmatrix}.$$

The equality of the upper-left entries can be stated as

$$c^2 a_1 - 2sc \operatorname{Re} a_{21} + s^2 a_2 = z_1.$$

This equation is quadratic in $t \stackrel{\text{def}}{=} s/c$,

$$(3.1) \quad (a_2 - z_1)t^2 - 2t \operatorname{Re} a_{21} + (a_1 - z_1) = 0,$$

whence

$$(3.2) \quad t = \frac{\operatorname{Re} a_{21} \pm \sqrt{(\operatorname{Re} a_{21})^2 - (a_1 - z_1)(a_2 - z_1)}}{a_2 - z_1}.$$

Notice that the discriminant is nonnegative due to the majorization condition. The \pm sign in (3.2) is taken to avoid sign cancellations with $\operatorname{Re} a_{21}$. If necessary, we can extract the other root of (3.1) using the fact that the product of its roots equals $(a_1 - z_1)/(a_2 - z_1)$. Finally, determine the parameters of the rotation using

$$(3.3) \quad c = \frac{1}{\sqrt{1+t^2}} \quad \text{and} \quad s = ct.$$

Floating-point arithmetic is inexact so the rotation may not yield $(Q^*AQ)_{11} = z_1$. A better implementation sets this entry to z_1 explicitly. Davies and Higham have shown that this method of computing rotations is numerically stable [6]. Since Q is orthogonal, Q^*AQ preserves the spectrum of A but replaces its diagonal with \mathbf{z} .

3.2. Generalized Bendel–Mickey. Let \mathbf{z} and \mathbf{a} be N -dimensional vectors for which $\mathbf{z} \succcurlyeq \mathbf{a}$. We will show how to transform a Hermitian matrix A with diagonal \mathbf{a} and spectrum $\boldsymbol{\lambda}$ into one with diagonal \mathbf{z} and spectrum $\boldsymbol{\lambda}$ using a sequence of plane rotations. It is enough to prove the result when the components of \mathbf{a} and \mathbf{z} are sorted in ascending order, so we place that restriction in the sequel.

Suppose that $\mathbf{a} \neq \mathbf{z}$. On account of the majorization relationship, it is possible to select indices $i < j$ that satisfy two properties: $a_i < z_i \leq z_j < a_j$ and $a_k = z_k$ for all k strictly between i and j . If $z_i - a_i \leq a_j - z_j$, then we construct a plane rotation Q in the (i, j) -plane such that $(Q^*AQ)_{ii} = z_i$. Otherwise, we find Q such that $(Q^*AQ)_{jj} = z_j$. Either rotation can be calculated using appropriate versions of (3.2) and (3.3). To see that this strategy can be repeated, we just need to check that \mathbf{z} majorizes the diagonal of Q^*AQ . In the first case, the plane rotation transforms a_i

to z_i and a_j to $a_i + a_j - z_i$, while the remaining diagonal entries do not change. Since $a_i < z_i \leq z_j \leq a_i + a_j - z_i < a_j$ the diagonal entries of Q^*AQ remain in ascending order. The first $(i-1)$ majorization conditions are clearly unaffected. Notice that

$$\sum_{\ell=1}^{i-1} a_\ell + z_i \leq \sum_{\ell=1}^{i-1} z_\ell + z_i,$$

which proves the i th majorization condition. The next $(j-i-1)$ majorization inequalities follow in consequence of a_k being equal to z_k whenever $i < k < j$. The rest of the majorization conditions hold since

$$\sum_{\ell=1}^{i-1} a_\ell + z_i + \sum_{k=i+1}^{j-1} a_k + (a_i + a_j - z_i) = \sum_{\ell=1}^j a_\ell \leq \sum_{\ell=1}^j z_\ell.$$

The argument in the case when $z_i - a_i > a_j - z_j$ is similar. It follows that our rotation strategy may be applied until $\text{diag } A = \mathbf{z}$. This proof leads to the following algorithm.

ALGORITHM 1 (Generalized Bendel–Mickey). *Let A be an $N \times N$ Hermitian matrix with diagonal \mathbf{a} and, furthermore, let \mathbf{z} be a vector such that $\mathbf{z} \succcurlyeq \mathbf{a}$, where both \mathbf{a} and \mathbf{z} are arranged in ascending order. The following procedure computes a Hermitian matrix with diagonal entries \mathbf{z} and eigenvalues equal to that of A .*

1. Find $i < j$ for which $a_i < z_i$, $z_j < a_j$ and $a_k = z_k$ for $i < k < j$ (in our implementation we pick the smallest such i). If no such pair exists, we are either done ($\mathbf{z} = \mathbf{a}$) or the majorization condition is violated.

2. Construct a plane rotation Q in the (i, j) -plane using (3.2) and (3.3) with appropriate modifications to transform a_i to z_i in the case $z_i - a_i \leq a_j - z_j$ or transform a_j to z_j otherwise.

3. Replace A by Q^*AQ .

4. Repeat Steps 1–3 until the diagonal is transformed to \mathbf{z} .

This algorithm requires about $12N^2$ real floating-point operations if conjugate symmetry is exploited. The storage requirement is about $N(N+1)/2$ complex floating-point numbers.

3.3. Generalized Chan–Li. Distinct algorithms arise by changing the strategy for selecting the planes of rotation. Let \mathbf{z} and \mathbf{a} be N -dimensional vectors for which $\mathbf{z} \succcurlyeq \mathbf{a}$. As before, we assume that they are sorted in ascending order, and suppose that A is a Hermitian matrix with diagonal \mathbf{a} . We now exhibit a different method for transforming the diagonal of A to \mathbf{z} while preserving its eigenvalues. It can be viewed as a generalization of the Chan–Li algorithm [4].

We will use induction on the dimension, so grant us for a moment that we can perform the advertised feat on Hermitian matrices of size $(N-1)$. Now we consider N -dimensional vectors for which $\mathbf{z} \succcurlyeq \mathbf{a}$, and suppose that $\text{diag } A = \mathbf{a}$. The majorization condition implies that $a_1 \leq z_1 \leq z_N \leq a_N$, so it is always possible to select a least integer $j > 1$ so that $a_{j-1} \leq z_1 \leq a_j$. Let P_1 be a permutation matrix for which

$$\text{diag}(P_1^*AP_1) = (a_1, a_j, a_2, \dots, a_{j-1}, a_{j+1}, \dots, a_N).$$

Observe that $a_1 \leq z_1 \leq a_j$ and $a_1 \leq a_1 + a_j - z_1 \leq a_j$. Thus we modify (3.2) and (3.3) to construct a two-dimensional plane rotation Q_2 that sets the upper-left entry of

$$Q_2^* \begin{bmatrix} a_1 & a_{j1}^* \\ a_{j1} & a_j \end{bmatrix} Q_2$$

to z_1 . If we define the rotation

$$P_2 \stackrel{\text{def}}{=} \begin{bmatrix} Q_2 & \mathbf{0}^* \\ \mathbf{0} & I_{N-2} \end{bmatrix},$$

then

$$P_2^* P_1^* A P_1 P_2 = \begin{bmatrix} z_1 & \mathbf{v}^* \\ \mathbf{v} & A_{N-1} \end{bmatrix},$$

where \mathbf{v} is an appropriate vector and A_{N-1} is an appropriate submatrix with

$$\text{diag}(A_{N-1}) = (a_1 + a_j - z_1, a_2, \dots, a_{j-1}, a_{j+1}, \dots, a_N).$$

In order to apply the induction hypothesis, it remains to check that the vector (z_2, z_3, \dots, z_N) majorizes the diagonal of A_{N-1} . We accomplish this in three steps. First, recall that $a_k \leq z_1$ for $k = 2, \dots, j-1$. Therefore,

$$\sum_{k=2}^m z_k \geq (m-1)z_1 \geq \sum_{k=2}^m a_k$$

for each $m = 2, \dots, j-1$. The sum on the right-hand side obviously exceeds the sum of the smallest $(m-1)$ entries of the vector $\text{diag } A_{N-1}$, so the first $(j-2)$ majorization inequalities are in force. Second, we use the fact that $\mathbf{z} \succcurlyeq \mathbf{a}$ to calculate that for $m = j, \dots, N$,

$$\sum_{k=2}^m z_k = \sum_{k=1}^m z_k - z_1 \geq \sum_{k=1}^m a_k - z_1 = (a_1 + a_j - z_1) + \sum_{k=2}^{j-1} a_k + \sum_{k=j+1}^m a_k.$$

Once again, observe that the sum on the right-hand side exceeds the sum of the smallest $(m-1)$ entries of the vector $\text{diag } A_{N-1}$, so the remaining majorization inequalities are in force. Finally, rearranging the relation $\sum_{k=1}^N z_k = \sum_{k=1}^N a_k$ yields $\sum_{k=2}^N z_k = \text{Tr } A_{N-1}$. In consequence, the induction furnishes a rotation Q_{N-1} that sets the diagonal of A_{N-1} equal to the vector (z_2, \dots, z_N) . Defining

$$P_3 \stackrel{\text{def}}{=} \begin{bmatrix} 1 & \mathbf{0}^* \\ \mathbf{0} & Q_{N-1} \end{bmatrix},$$

we see that conjugating A by the orthogonal matrix $P = P_1 P_2 P_3$ transforms the diagonal entries of A to \mathbf{z} while retaining the spectrum $\boldsymbol{\lambda}$. This proof leads to the following algorithm.

ALGORITHM 2 (Generalized Chan–Li). *Let A be an $N \times N$ Hermitian matrix with diagonal \mathbf{a} , and let \mathbf{z} be a vector such that $\mathbf{z} \succcurlyeq \mathbf{a}$, where both \mathbf{a} and \mathbf{z} are arranged in ascending order. The following procedure computes a Hermitian matrix with diagonal entries \mathbf{z} and eigenvalues equal to that of A .*

1. Set $i = 1$.
2. Find the least $j > i$ so that $a_{j-1} \leq z_i \leq a_j$.
3. Use a symmetric permutation to set a_{i+1} equal to a_j while shifting diagonal entries $i+1, \dots, j-1$ one place down the diagonal.
4. Construct a plane rotation Q in the $(i, i+1)$ -plane using (3.2) and (3.3) with appropriate modifications.

5. Replace A by Q^*AQ .
6. Use a symmetric permutation to re-sort the diagonal entries of A in ascending order.
7. Increment i and repeat Steps 2–7 while $i < N$.

This algorithm has the same functionality and complexity as Algorithm 1 but it is different in the plane rotations used. It is conceptually simpler to perform the permutations described in the algorithm, but an implementation can avoid them. The MATLAB code provided in the appendix of [7] demonstrates one such implementation.

3.4. One-sided algorithms. It is well known that any positive semidefinite matrix $A \in \mathbb{M}_N$ can be expressed as the product X^*X where $X \in \mathbb{M}_{d,N}$ and $d \geq \text{rank } A$. With this factorization, the two-sided transformation $A \mapsto Q^*AQ$ is equivalent to a one-sided transformation $X \mapsto XQ$. In consequence, the machinery of Algorithm 1 requires little adjustment to produce these factors. The following algorithm generalizes the one-sided version proposed by Davies and Higham in [6].

ALGORITHM 3 (One-sided generalized Bendel–Mickey). *Suppose that \mathbf{z} and \mathbf{a} are nonnegative vectors of length N with ascending entries. Assume, moreover, that $\mathbf{z} \succcurlyeq \mathbf{a}$. The following algorithm takes as input a $d \times N$ complex matrix X whose squared column norms are listed by \mathbf{a} and transforms it into a matrix with the same singular spectrum and with squared column norms listed by \mathbf{z} .*

1. Find $i < j$ for which $\|\mathbf{x}_{i-1}\|_2^2 < z_i$, $z_j < \|\mathbf{x}_j\|_2^2$ and $\|\mathbf{x}_k\|_2^2 = z_k$ for $i < k < j$. If no such pair exists, we are either done or the majorization condition is violated.
2. Form the quantities

$$a_i = \|\mathbf{x}_i\|_2^2, \quad a_{ji} = \langle \mathbf{x}_j, \mathbf{x}_i \rangle, \quad \text{and} \quad a_j = \|\mathbf{x}_j\|_2^2.$$

3. Construct a plane rotation Q in the (i, j) -plane using (3.2) and (3.3) with appropriate modifications to transform a_i to z_i in the case $z_i - a_i \leq a_j - z_j$, or transform a_j to z_j otherwise.
4. Replace X by XQ .
5. Repeat Steps 1–4 until all column norms are as desired.

The algorithm requires about $12dN$ real floating-point operations and storage of $N(d+2)$ complex floating-point numbers including the desired column norms and the current column norms. A similar modification of our generalized Chan–Li algorithm also leads to a one-sided version.

4. Numerical examples. The generalized algorithms can produce a richer set of matrices with prescribed diagonal entries and eigenvalues, making it possible to find solutions that satisfy additional properties or better suit the application.

We illustrate the generalized Chan–Li algorithm by comparing it with the classical algorithm. Suppose we want to produce a Hermitian matrix with eigenvalues $(1, 4, 5, 7, 9)$ and diagonal entries $(2, 5, 6, 6, 7)$. This example was presented in [4]; our generalized algorithm (essentially) yields the same result

$$A_{\text{ChanLi}}^{(1)} = \begin{bmatrix} 2.0000 & 0 & 0.7071 & -0.9487 & 0.7746 \\ 0 & 5.0000 & 0 & 0 & 0 \\ 0.7071 & 0 & 6.0000 & 1.3416 & -1.0954 \\ -0.9487 & 0 & 1.3416 & 6.0000 & 2.4495 \\ 0.7746 & 0 & -1.0954 & 2.4495 & 7.0000 \end{bmatrix}.$$

Notice the sparsity structure in the above matrix. In applications such as designing matrices for testing eigenvalue solvers [9], it would be better to produce a more

random Hermitian matrix that satisfies the diagonal and eigenvalue constraints. The generalized algorithms can be used for this purpose. We first generate a sequence of (say) six vectors satisfying

$$(1, 4, 5, 7, 9) = \mathbf{z}_0 \preceq \mathbf{z}_1 \preceq \mathbf{z}_2 \preceq \mathbf{z}_3 \preceq \mathbf{z}_4 \preceq \mathbf{z}_5 = (2, 5, 6, 6, 7).$$

The vectors $\mathbf{z}_1, \mathbf{z}_2, \mathbf{z}_3, \mathbf{z}_4$ may be generated as random convex combinations of \mathbf{z}_0 and \mathbf{z}_5 . This construction ensures that the majorization conditions hold. Five steps of the generalized Chan–Li algorithm are then successively used to transform the diagonal matrix $\text{diag}(\mathbf{z}_0)$ to have diagonals $\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_5$. We arrive at the final matrix

$$A_{\text{GenChanLi}}^{(1)} = \begin{bmatrix} 2.0000 & 1.0400 & 1.4517 & -0.6294 & -0.3720 \\ 1.0400 & 5.0000 & 0.3620 & -0.2157 & 1.4731 \\ 1.4517 & 0.3620 & 6.0000 & 1.6901 & -0.6544 \\ -0.6294 & -0.2157 & 1.6901 & 6.0000 & -1.2822 \\ -0.3720 & 1.4731 & -0.6544 & -1.2822 & 7.0000 \end{bmatrix}.$$

For the wireless application mentioned in the introduction, the matrices in question must have all nonzero eigenvalues equal to one; see [20] for more details. The following example calls for the generation of matrices with eigenvalues $(0, 0, 1, 1, 1)$ and diagonal $(0.4, 0.6, 0.6, 0.6, 0.8)$. The Chan–Li algorithm produces

$$A_{\text{ChanLi}}^{(2)} = \begin{bmatrix} 0.4000 & -0.4899 & 0 & 0 & 0 \\ -0.4899 & 0.6000 & 0 & 0 & 0 \\ 0 & 0 & 0.6000 & 0.4000 & -0.2828 \\ 0 & 0 & 0.4000 & 0.6000 & 0.2828 \\ 0 & 0 & -0.2828 & 0.2828 & 0.8000 \end{bmatrix}.$$

In the wireless application, it is often desirable to have lower variance in the magnitudes of the off-diagonal entries, which are also known as “cross-correlations.”

We can use the generalized algorithms to obtain “random” solutions as follows. Generate random matrices with the given spectrum and select those which have a diagonal majorized by the target diagonal. Then apply either one of the generalized algorithms. Using such a strategy, the generalized Bendel–Mickey algorithm provides the following matrix:

$$A_{\text{GenBendelMickey}}^{(2)} = \begin{bmatrix} 0.4000 & 0.1803 & 0.2101 & -0.3558 & 0.1916 \\ 0.1803 & 0.6000 & -0.3991 & -0.1487 & 0.1615 \\ 0.2101 & -0.3991 & 0.6000 & -0.1237 & 0.1457 \\ -0.3558 & -0.1487 & -0.1237 & 0.6000 & 0.2756 \\ 0.1916 & 0.1615 & 0.1457 & 0.2756 & 0.8000 \end{bmatrix}.$$

5. Conclusions. We have shown that a sequence of $(N-1)$ rotations is sufficient to replace the original diagonal of $N \times N$ Hermitian matrix with any set of diagonal entries that majorizes the original set, all the while preserving the spectrum of the matrix. The algorithms we have presented can move up a chain in the poset of Schur–Horn equivalence classes as given in Definition 2.5.

An obvious question is whether it is possible to obtain an algorithm that moves down a chain instead. In other words, is it possible to construct a finite sequence of rotations to replace the diagonal with a set of entries that majorizes the eigenvalues but not necessarily the original diagonal? Since the diagonal matrix of eigenvalues lies at the bottom of every chain, it might seem at first glance that, as a special case, we

are attempting to calculate the eigenvalues in finite time. We avoid this paradox since we assume that the target diagonal is already known. In fact, to get to the bottom of the chain, $O(N^2)$ Givens rotations can be used to reduce the initial matrix to tridiagonal form and then to the desired diagonal matrix of eigenvalues (these include the application of perfect shifts to the tridiagonal matrix; see [8]). By combining these $O(N^2)$ rotations with the Chan–Li algorithm, we answer our question in the affirmative. However this strategy goes through the bottom of the chain, so a natural question is whether there is a more direct construction. A direct algorithm may also be computationally more efficient. However, it seems harder to transform a vector into one of its predecessors than into one of its successors. Entropy may provide a reasonable explanation: it is easier to make things more uniform than less uniform.

Other interesting questions arise. What is the structure of a general Schur–Horn equivalence class of Hermitian matrices? Is there a fast procedure that is capable of constructing every member of a given equivalence class? Is it possible to define a uniform probability measure on each class and to construct members from a class uniformly at random? In this paper, we have restricted our attention to finite step algorithms. Iterative algorithms are an alternative, especially for the case of producing Hermitian matrices that satisfy additional constraints; see [21] for such an approach. It would be useful to understand these problems better, and we hope that other researchers will take interest.

REFERENCES

- [1] B. C. ARNOLD, *Majorization and the Lorenz Order: A Brief Introduction*, Lecture Notes in Statistics, 43, Springer-Verlag, Berlin, 1987.
- [2] R. B. BENDEL AND M. R. MICKEY, *Population correlation matrices for sampling experiments*, Commun. Statist. Simul. Comp., B7 (1978), pp. 163–182.
- [3] R. BHATIA, *Matrix Analysis*, Springer-Verlag, New York, 1996.
- [4] N. N. CHAN AND K.-H. LI, *Diagonal elements and eigenvalues of a real symmetric matrix*, J. Math. Anal. Appl., 91 (1983), pp. 562–566.
- [5] M. T. CHU, *Constructing a Hermitian matrix from its diagonal entries and eigenvalues*, SIAM J. Matrix Anal. Appl., 16 (1995), pp. 207–217.
- [6] P. I. DAVIES AND N. J. HIGHAM, *Numerically stable generation of correlation matrices and their factors*, BIT, 40 (2000), pp. 640–651.
- [7] I. S. DHILLON, R. W. HEATH JR., M. A. SUSTIK, AND J. A. TROPP, *Generalized finite algorithms for constructing Hermitian matrices with prescribed diagonal and spectrum*, Technical Report TR-03-49, Department of Computer Sciences, The University of Texas at Austin, Austin, TX, 2003.
- [8] I. S. DHILLON AND A. N. MALYSHEV, *Inner deflation for symmetric tridiagonal matrices*, Linear Algebra Appl., 358 (2003), pp. 139–144.
- [9] Z. DRMAČ, *Accurate computation of the product-induced singular value decomposition with applications*, SIAM J. Numer. Anal., 35 (1998), pp. 1969–1994.
- [10] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations*, 3rd ed., Johns Hopkins University Press, Baltimore, MD, 1996.
- [11] R. B. HOLMES, *On random correlation matrices*, SIAM J. Matrix Anal. Ap., 12 (1991), pp. 239–272.
- [12] A. HORN, *Doubly stochastic matrices and the diagonal of a rotation matrix*, Amer. J. Math., 76 (1954), pp. 620–630.
- [13] R. A. HORN AND C. R. JOHNSON, *Matrix Analysis*, Cambridge University Press, Cambridge, UK, 1985.
- [14] K. D. IKRAMOV, *A remark on “A note on constructing a symmetric matrix with specified diagonal entries and eigenvalues,”* BIT, 38 (1998), p. 807.
- [15] J. L. KELLEY, *General Topology*, Van Nostrand, New York, 1955.
- [16] R. S. LEITE, R. W. RICHA, AND C. TOMEI, *Geometric proofs of some theorems of Schur–Horn type*, Linear Algebra Appl., 286 (1999), pp. 149–173.
- [17] G. MARSAGLIA AND I. OLKIN, *Generating correlation matrices*, SIAM J. Sci. Statist. Comput.,

- 5 (1984), pp. 470–475.
- [18] A. W. MARSHALL AND I. OLKIN, *Inequalities: Theory of Majorization and its Applications*, Academic Press, New York, 1979.
 - [19] G. W. STEWART, *The efficient generation of random orthogonal matrices with an application to condition estimation*, SIAM J. Numer. Anal., 17 (1980), pp. 403–409.
 - [20] J. A. TROPP, I. S. DHILLON, AND R. W. HEATH JR., *Finite-step algorithms for constructing optimal CMDA signature sequences*, IEEE Trans. Inform. Theory, 50 (2004), pp. 2916–2921.
 - [21] J. A. TROPP, I. S. DHILLON, R. W. HEATH JR., AND T. STROHMER, *Designing structured tight frames via an alternating projection method*, IEEE Trans. Inform. Theory, 51 (2004), pp. 188–209.
 - [22] P. VISWANATH AND V. ANANTHARAM, *Optimal sequences and sum capacity of synchronous CDMA systems*, IEEE Trans. Inform. Theory, 45 (1999), pp. 1984–1991.
 - [23] H. ZHA AND Z. ZHANG, *A note on constructing a symmetric matrix with specified diagonal entries and eigenvalues*, BIT, 35 (1995), pp. 448–451.