BULL. AUSTRAL. MATH. SOC. 65F05
VOL. 26 (1982), 107-119.

STABLE ALGORITHMS FOR SOLVING SYMMETRIC AND SKEW-SYMMETRIC SYSTEMS*

JAMES R. BUNCH

Communicated by James M. Hill

Algorithms for decomposing symmetric and skew-symmetric matrices in order to solve systems of linear equations will be discussed. The algorithms are numerically stable, yet take advantage of the symmetry or skew-symmetry to halve the work and storage.

1. Introduction

We shall consider solving $n \times n$ systems of linear equations when A is symmetric $\left(A = A^T\right)$ or skew-symmetric $\left(A = -A^T\right)$ - or Hermitian $\left(A = \bar{A}^T\right)$ or skew-Hermitian $\left(A = -\bar{A}^T\right)$. We shall, in general, only discuss the case when A is real, pointing out any differences when A is complex.

In practice, most symmetric systems are also positive definite, that is, $x^TAx > 0$ for all $x \neq 0$. This is the easiest of the three cases to solve and will be discussed in §3. If A is symmetric indefinite, that is, there exist $x, y \neq 0$ such that $x^TAx > 0$ and $y^TAy < 0$, then this is the hardest of the three cases and will be discussed in §4. Skew-

Received 8 March 1982. Support for this research was provided under NSF grant MCS 79-20491.

* This paper is based on an invited lecture given at the Australian Mathematical Society Applied Mathematics Conference held in Bundanoon, February 7-11, 1982. Other papers delivered at this Conference appear in Volumes 25 and 26.

https://doi.org/10.1017/S0004972700005621 Published online by Cambridge University Press

symmetric systems lie intermediate in difficulty between definite and indefinite systems and will be discussed in §5.

If A is (real) symmetric, then all its eigenvalues are real. We define the *inertia* of A to be the triple (π, ν, ζ), where π, ν, ζ are the number of positive, negative, and zero eigenvalues of A. If A is nonsingular then $\zeta = 0$; if A is positive definite then $\pi = n$, $\nu = 0$, and $\zeta = 0$. By Sylvester's Inertia Theorem [9], the inertia of a symmetric matrix is preserved under (nonsingular) congruence transformations, that is, A and $B = CAC^T$ have the same inertia where C is nonsingular.

If A is (real) skew-symmetric, then all its eigenvalues are purely imaginary. Hence, here we define the *inertia* of A to be the triple (π, ν, ζ), where π, ν, ζ are the number of positive, negative, and zero imaginary parts of the eigenvalues. But, since A is real, its nonzero eigenvalues occur in complex conjugate pairs, that is, $\pm i\mu_j$ where μ_j are positive. Hence the inertia of any real skew-symmetric matrix is $\left((n-\zeta)/2, (n-\zeta)/2, \zeta\right)$. If A is also nonsingular, its inertia is $(n/2, n/2, 0)$. This fixed inertia property makes skew-symmetric matrices easier to decompose stably than symmetric indefinite matrices. If A is skew-symmetric then $B = CAC^T$ is skew-symmetric and has the same inertia as A, where C is nonsingular.

2. Lagrange's method

The classical method [9] for calculating the inertia of a symmetric matrix is Lagrange's method (1759): a (real) quadratic form

$$\varphi(x) = x^T A x = \sum_{i=1}^{n} \sum_{j=1}^{n} a_{ij} x_i x_j = \varphi\left(x_1, \ldots, x_n\right),$$

where $A = A^T$, is reduced to a diagonal form

$$\sum_{k=1}^{n} d_k z_k^2$$

by linear congruence transformations. Hence the inertia of A is the same as the number of positive, negative, and zero d_k's.

https://doi.org/10.1017/S0004972700005621 Published online by Cambridge University Press

Let us look more closely at Lagrange's method. If $a_{11} \neq 0$, then

$$\begin{aligned}\varphi(x) &= a_{11}x_1^2 + 2a_{12}x_1x_2 + \ldots + 2a_{1n}x_1x_n + \sum_{i=2}^{n}\sum_{j=2}^{n} a_{ij}x_ix_j \\ &= a_{11}\left(x_1^2 + 2\,\frac{a_{12}}{a_{11}}\,x_1x_2 + \ldots + 2\,\frac{a_{1n}}{a_{11}}\,x_1x_n\right) + \sum_{i=2}^{n}\sum_{j=2}^{n} a_{ij}x_ix_j \\ &= a_{11}\left(x_1 + \frac{a_{12}}{a_{11}}\,x_2 + \ldots + \frac{a_{1n}}{a_{11}}\,x_n\right)^2 + \sum_{i=2}^{n}\sum_{j=2}^{n}\left(a_{ij} - \frac{a_{1i}a_{1j}}{a_{11}}\right)x_ix_j \\ &= d_1z_1^2 + \varphi(x_2, \ldots, x_n) ,\end{aligned}$$

where

$$d_1 \equiv a_{11} ,$$

$$z_1 \equiv x_1 + \frac{a_{12}}{a_{11}}\,x_2 + \ldots \frac{a_{1n}}{a_{11}}\,x_n ,$$

and

$$\varphi(x_2, \ldots, x_n) \equiv \sum_{i=2}^{n}\sum_{j=2}^{n}\left(a_{ij} - \frac{a_{1i}a_{1j}}{a_{11}}\right)x_ix_j$$

is a quadratic form in the $n - 1$ variables $x_2, \ldots, x_n$. If $a_{22} - a_{12}^2/a_{11} \neq 0$, we can continue as above to eliminate x_2 .

Let us write this first part of Lagrange's method in matrix form. If $a_{11} \neq 0$, let

$$L_1 = \begin{bmatrix} 1 & & & \\ l_{21} & 1 & 0 & \\ \vdots & 0 & \ddots & \\ l_{n1} & & & 1 \end{bmatrix} ,$$

where $l_{j1} = a_{j1}/a_{11}$; let $z_1 \equiv L_1x$; let

$$D_1 = \left[\begin{array}{c|c} d_1 & 0 \\ \hline 0 & A^{(n-1)} \end{array}\right]$$

https://doi.org/10.1017/S0004972700005621 Published online by Cambridge University Press

where $d_1 = a_{11}$ and

$$(D_1)_{ij} = A^{(n-1)}_{i-1,j-1} = a_{ij} - a_{1i}a_{1j}/a_{11}$$

for $2 \le i, j \le n$. Then

$$A = L_1 D_1 L_1^T$$

and

$$\varphi(x) = x^T A x = d_1 z_1^2 + \varphi(x_2, \ldots, x_n) ,$$

where

$$\varphi(x_2, \ldots, x_n) = [x_2, \ldots, x_n] A^{(n-1)} \begin{bmatrix} x_2 \\ \vdots \\ x_n \end{bmatrix} .$$

But $A = L_1 D_1 L_1^T$ is exactly the matrix form of the first step of symmetric Gaussian elimination performed on A , where the l_{j1} are the multipliers and $A^{(n-1)}$ is the reduced matrix. Thus the first part of Lagrange's method is just symmetric Gaussian elimination.

If $a_{11} = 0$ but $a_{kk} \ne 0$ for some $k > 1$, then let P be the permutation matrix obtained by interchanging the kth and first row and column of the identity matrix. Then $P = P^T = P^{-1}$ and

$$\varphi(x) = x^T A x = \tilde{x}^T (P^T A P) \tilde{x} ,$$

where $\tilde{x} = Px$. Now $(P^T A P)_{11} = a_{kk} \ne 0$, so we may eliminate $\tilde{x}_1 = x_k$. In matrix form, we obtain, as before,

$$P^T A P = L_1 D_1 L_1^T .$$

However, if the diagonal of A was null (or if at some stage during the process the diagonal was null), we could not do this. If $A \equiv 0$, we would be finished. Otherwise, there exists $a_{rs} \ne 0$ with $r \ne s$. For simplicity, assume $a_{12} \ne 0$ (otherwise, interchange the rth and first

https://doi.org/10.1017/S0004972700005621 Published online by Cambridge University Press

variables and sth and second variables). In this case, Lagrange suggested applying the transformation:

$$x_1 = y_1 + y_2 \ , \quad x_2 = y_1 - y_2 \ , \quad x_3 = y_3, \ \ldots, \ x_n = y_n \ .$$

This maps $2a_{12}x_1x_2$ into $2a_{12}\left(y_1^2-y_2^2\right)$ and the coefficient of the y_1y_2 term is zero. Let

$$R = \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} \oplus I_{n-2} \ ,$$

$$x = \begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix} , \text{ and } \ y = \begin{bmatrix} y_1 \\ \vdots \\ y_n \end{bmatrix} .$$

Then

$$x = Ry$$

and

$$\varphi(x) = x^TAx = y^T\left(R^TAR\right)y \equiv \psi(y)$$

is a quadratic form in y with

$$\left(R^TAR\right)_{11} = 2a_{12} \ ,$$

$$\left(R^TAR\right)_{22} = -2a_{12} \ ,$$

and

$$\left(R^TAR\right)_{12} = 0 \ .$$

We may now eliminate y_1 . But, since $\left(R^TAR\right)_{12} = 0$, the coefficient of y_2^2 in the new quadratic form in $y_2, \ldots, y_n$ is still $-2a_{12}$, and is, hence, nonzero. So, we may also eliminate y_2 . Thus the change of variables above guarantees the elimination of two variables. Later in §4 we shall relate this process to symmetric Gaussian elimination.

3. Symmetric positive definite systems

If A is symmetric and positive definite $\left(x^T A x > 0 \text{ for all } x \neq 0\right)$, then $a_{11} > 0$ and the first part of Lagrange's method can be done: $A = L_1 D_1 L_1^T$ as in §2. Then the reduced matrix $A^{(n-1)}$ is once again symmetric positive definite; hence the first part of Lagrange's method is applicable at each step. So $A = LDL^T$, where L is unit lower triangular and D is diagonal with positive diagonal elements; this is exactly symmetric Gaussian elimination. In order to solve $Ax = b$, we solve $Ly = b$ for y and then $L^T x = D^{-1} y$ for x.

Another well-known method for solving symmetric positive definite systems is the Cholesky decomposition. Here A is decomposed as $A = \tilde{L}\tilde{L}^T$ where $\tilde{L}$ is lower triangular. The two methods are related mathematically by $\tilde{L} = LD^{\frac{1}{2}}$. The Cholesky decomposition is used in LINPACK [8] for solving symmetric positive definite systems.

Each method requires $\frac{1}{6}n^3$ multiplications, $\frac{1}{6}n^3$ additions, and no comparisons. Let $B(n)$ be the largest element (in modulus) that occurs in any reduced matrix during the decomposition process divided by the largest element (in modulus) in the original matrix A. Then $B(n) = 1$ for symmetric Gaussian elimination and $B(n) = \frac{1}{\sqrt{\max a_{jj}}}$ for Cholesky's method [11] $\left(B(n) = 1 \text{ if } \max_{i,j} |a_{ij}| = 1\right)$. Since $B(n)$ measures the stability of an algorithm [11], both symmetric Gaussian elimination and Cholesky's method are very stable for symmetric positive definite systems.

4. Symmetric indefinite systems

The two well-known algorithms for decomposing symmetric indefinite matrices are the tridiagonal method [1], [10] and the diagonal pivoting method [2, 3], [5], [6], [7], [8].

The tridiagonal method uses stabilized elementary congruence transformations to reduce a symmetric matrix A to a symmetric tridiagonal matrix T :

https://doi.org/10.1017/S0004972700005621 Published online by Cambridge University Press

$$A = P_2L_2 \ldots P_nL_nTL_n^TP_n \ldots L_2^TP_2 ,$$

where the P_j are elementary permutation matrices and the L_j are unit lower triangular. At each step the largest element in the pivot column is interchanged to the (2, 1) position by symmetric permutation.

This requires $\frac{1}{6}n^3$ multiplications, $\frac{1}{6}n^3$ additions, $\frac{1}{2}n^2$ comparisons, and $B(n) \le 4^{n-2}$ [3, p. 525].

(Then Gaussian elimination with partial pivoting is used to decompose T further.)

The diagonal pivoting method reduces a symmetric matrix A by stabilized congruence transformations to a symmetric block diagonal matrix D with blocks of order 1 or 2 :

$$A = P_1M_1P_2M_2 \ldots P_{n-1}M_{n-1}DM_{n-1}^TP_{n-1} \ldots M_2^TP_2M_1^TP_1 ,$$

where the P_j are elementary permutation matrices and the M_j are unit lower triangular matrices.

One step of the decomposition process looks as follows:

$$A \equiv \begin{bmatrix} S & C^T \\ C & B \end{bmatrix} = \underbrace{\begin{bmatrix} I & 0 \\ CS^{-1} & I \end{bmatrix}}_{\equiv M_1} \begin{bmatrix} S & 0 \\ 0 & A^{(n-k)} \end{bmatrix} \underbrace{\begin{bmatrix} I & S^{-1}C^T \\ 0 & I \end{bmatrix}}_{= M_1^T} ,$$

where S is $k \times k$, nonsingular, B is $(n-k) \times (n-k)$, C is $(n-k) \times k$, and $A^{(n-k)} \equiv B - CS^{-1}C^T$, $k = 1$ or 2 .

Let us now look at Lagrange's method when the diagonal of A is null. If

$$R = U \oplus I_{n-2} ,$$

where U is of order 2 , then

$$R^TAR = \left[\begin{array}{c|c} U^TSU & U^TC^T \\ \hline CU & B \end{array}\right] .$$

https://doi.org/10.1017/S0004972700005621 Published online by Cambridge University Press

If U is chosen so that $U^T S U \equiv D$ is diagonal $\left(U = \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}\right.$ is only one such choice$\left.\right)$ and the associated variables (that is, the first two variables) are eliminated, then the resulting reduced matrix is

$$B - CUD^{-1}U^T C^T .$$

But if we use S as a 2×2 pivot and performed block symmetric Gaussian elimination, the reduced matrix is

$$A^{(n-2)} = B - CS^{-1}C^T .$$

Since

$$CUD^{-1}U^T C^T = CS^{-1}C^T ,$$

the reduced matrices are identical.

Thus there is no need to find (as Lagrange did) a matrix U which diagonalizes S ; rather, we may reduce A by congruences to a symmetric block diagonal form with blocks of order 1 or 2 . Any block of order 2 is of the form

$$\begin{bmatrix} 0 & a_{12} \\ a_{12} & 0 \end{bmatrix} .$$

Since its determinant is negative, it has one positive and one negative eigenvalue.

In order to maintain stability, we must also employ 2×2 pivots whenever the diagonal of A is small. This can be done while preserving the property that the determinant of any 2×2 pivot block is negative [5], [7], [8] (it is nonzero in the complex symmetric case).

The method requires $\frac{1}{6}n^3$ multiplications, $\frac{1}{6}n^3$ additions, with more than or the same number of $\frac{1}{2}n^2$ but less than or the same number of n^2 comparisons and $B(n) < (2.57)^{n-1}$ for a partial pivoting strategy or with more than or the same number of $\frac{1}{12}n^3$ but less than or the same number of $\frac{1}{6}n^3$ comparisons and $B(n) < 3nf(n)$ for a complete pivoting strategy where

https://doi.org/10.1017/S0004972700005621 Published online by Cambridge University Press

$$f(n) = \left[\prod_{k=2} k^{1/(k-1)}\right]^{\frac{1}{2}} < 1.8n^{\frac{1}{4}\ln n} .$$

(Recall that Gaussian elimination requires $\frac{1}{3}n^3$ multiplications, $\frac{1}{3}n^3$ additions, with $\frac{1}{2}n^2$ comparisons and $B(n) \le 2^{n-1}$ for partial pivoting or with $\frac{1}{3}n^3$ comparisons and $B(n) < \sqrt{n}\, f(n)$ for complete pivoting.)

There are actually three cases here: real symmetric, complex symmetric, and complex Hermitian. Both algorithms cover all three cases. The diagonal pivoting algorithm with partial pivoting is used in LINPACK [8, Chapter 5].

5. Skew systems

If A is a skew-symmetric $(A = -A^T)$, then the diagonal of A is null, and if n is odd then $\det A = 0$. Thus, if A is skew-symmetric and nonsingular then n is even. The diagonal of A being null may seem to pose a difficulty at first glance, but we shall see that this property makes the skew-symmetric case easier than the symmetric indefinite case.

If A is skew-Hermitian $(A = -\bar{A}^T)$, then the diagonal of A is purely imaginary but not necessarily null, for example,

$$A = \begin{bmatrix} i & -1+2i \\ 1+2i & 4i \end{bmatrix} .$$

If n is odd then $\mathrm{Re}(\det A) = 0$, and if n is even then $\mathrm{Im}(\det A) = 0$. If A is skew-Hermitian, then $B = iA$ is Hermitian, and we can apply any of the algorithms in §4 to B. (In order to solve $Ax = b$, we solve $Bx = ib$.)

Similarly, if A is (real or complex) skew-symmetric, then $B = iA$ is Hermitian, and we could apply the algorithms in §4. However, if A is *real* skew-symmetric, we would prefer to stay in *real* arithmetic. Is there a *stable* decomposition of real skew-symmetric matrices which would allow us to stay in *real* airthmetic? Such a decomposition should be based on congruence transformations, since they preserve the inertia and guarantee that each reduced matrix during the process remains skew-symmetric.

Let us find such congruence transformations. Let

$$A = \begin{bmatrix} S & -C^T \\ C & B \end{bmatrix}$$

where S is $k \times k$, C is $(n-k) \times k$, B is $(n-k) \times (n-k)$; S and B are skew-symmetric.

If S is nonsingular then

$$A = \underbrace{\begin{bmatrix} I & 0 \\ CS^{-1} & I \end{bmatrix}}_{\equiv M} \begin{bmatrix} S & 0 \\ 0 & A^{(n-k)} \end{bmatrix} \underbrace{\begin{bmatrix} I & -S^{-1}C^T \\ 0 & I \end{bmatrix}}_{\equiv \tilde{M} = M^T},$$

where $A^{(n-k)} \equiv B + CS^{-1}C^T$ is skew-symmetric, and $\tilde{M} = M^T$ since $S^{-T} = -S^{-1}$.

Thus we have performed a congruence transformation, and A and $\begin{bmatrix} S & 0 \\ 0 & A^{(n-k)} \end{bmatrix}$ are congruent (and have the same inertia).

Note that S being $k \times k$, skew-symmetric, and nonsingular implies that k is even. Since the diagonal of A is null, $k \neq 1$ unless $C \equiv 0$. Let $k = 2$ and

$$S = \begin{bmatrix} 0 & -a_{21} \\ a_{21} & 0 \end{bmatrix}.$$

If $a_{21} = 0$ but $a_{j1} \neq 0$, interchange the jth and second row and column. (If A is nonsingular, then n is even and $k = 2$ suffices at each step.) On conclusion,

$$A = \left(P_1M_1P_2M_2 \ldots P_{n-1}M_{n-1}\right)D\left(M_{n-1}^TP_{n-1} \ldots M_2^TP_2M_1^TP_1\right),$$

where the P_j are elementary permutation matrices (possibly $P_j = I$), the M_j are unit lower triangular matrices (possibly $M_j = I$), and D is a skew-symmetric block diagonal matrix with blocks of order 1 (zero blocks) or of order 2 (nonsingular blocks).

This gives existence of the decomposition and requires $\frac{1}{6}n^3$ multiplications, $\frac{1}{6}n^3$ additions, and $\frac{1}{2}n^2$ comparisons; the decomposition can be stored in the strictly upper (or lower) triangular part of A, plus one n-vector to store the permutation information.

Stability of the decomposition can be obtained by either a partial or complete pivoting strategy.

If $|a_{21}| = \max_{2 \le i \le n} \{|a_{i1}|, |a_{i2}|\}$, then

$$\max_{i,j} \left|\left(A^{(n-2)}\right)_{ij}\right| \le 3 \max_{r,s} |a_{rs}| .$$

If $|a_{m1}| = \max_{2 \le i \le n} \{|a_{i1}|, |a_{i2}|\}$, then interchange the mth and second row and column. If $|a_{m2}| = \max_{2 \le i \le n} \{|a_{i1}|, |a_{i2}|\}$, then interchange the first and second row and column and then the mth and second row and column. This provides a partial pivoting strategy with $B(n) \le (\sqrt{3})^{n-2} < (1.7321)^{n-2}$ at a cost of $\frac{1}{2}n^2$ comparisons.

A complete pivoting strategy brings the largest element in the reduced matrix to the $(2, 1)$ position at each step. This yields $B(n) < \sqrt{n}\, f(n)$ at a cost of $\frac{1}{12}n^3$ comparisons.

One can similarly modify the tridiagonal method [1], [10] yielding

$$A = \left(P_2L_2 \ldots P_nL_n\right)T\left(L_n^TP_n \ldots L_2^TP_2\right) ,$$

where the P_j are elementary permutation matrices, the L_j are unit lower triangular matrices, and T is skew-symmetric and tridiagonal. Here $B(n) \le 3^{n-2}$ and $\frac{1}{6}n^3$ multiplications. $\frac{1}{6}n^3$ additions, and $\frac{1}{2}n^2$ comparisons are required.

In conclusion, we see that skew-symmetric systems may be solved stably using congruence transformations; they are intermediate in difficulty between symmetric positive definite systems and symmetric indefinite systems. The situation can be summarized in the table below.

Matrix		Sym. Pos. Def.	Skew-Sym. (n even)	Sym. Indef.	General
Method		LDL^T	Diagonal Pivoting		Gaussian Elimination
Pivot size		1	2	1 or 2	1
Number of multiplications or additions		$\frac{1}{6}n^3$	$\frac{1}{6}n^3$	$\frac{1}{6}n^3$	$\frac{1}{3}n^3$
Number of comparisons	partial	0	$\frac{1}{2}n^2$	$\frac{1}{2}n^2$ to n^2	$\frac{1}{2}n^2$
	complete		$\frac{1}{12}n^3$	$\frac{1}{12}n^3$ to $\frac{1}{6}n^3$	$\frac{1}{3}n^3$
$B(n)$	partial	1	$< (1.74)^{n-2}$	$< (2.57)^{n-1}$	$= 2^{n-1}$
	complete		$< \sqrt{n}\, f(n)$	$< 3nf(n)$	$< \sqrt{n}\, f(n)$

Method		Tridiagonal method	
Pivot size		1	1
Number of multiplications or additions		$\frac{1}{6}n^3$	$\frac{1}{6}n^3$
Number of comparisons		$\frac{1}{2}n^2$	$\frac{1}{2}n^2$
$B(n)$		$\leq 3^{n-2}$	$\leq 4^{n-2}$

$$f(n) = \left(\prod_{k=2}^{n} k^{1/(k-1)}\right)^{\frac{1}{2}} < 1.8n^{\frac{1}{4}\ln n} .$$

References

[1] Jan Ole Aasen, "On the reduction of a symmetric matrix to tridiagonal form", *BIT* 11 (1971), 233-242.

[2] J.R. Bunch, "Analysis of the diagonal pivoting method", *SIAM J. Numer. Anal.* 8 (1971), 656-680.

https://doi.org/10.1017/S0004972700005621 Published online by Cambridge University Press

[3] James R. Bunch, "Partial pivoting strategies for symmetric matrices", *SIAM J. Numer. Anal.* 11 (1974), 521-528.

[4] James R. Bunch, "Stable decomposition of skew-symmetric matrices", *Math. Comp.* (to appear).

[5] James R. Bunch and Linda Kaufman, "Stome stable methods for calculating inertia and solving symmetric linear systems", *Math. Comp.* 31 (1977), 163-179.

[6] James R. Bunch, Linda Kaufman and Beresford N. Parlett, "Decomposition of a symmetric matrix", *Numer. Math.* 27 (1976), 95-109.

[7] J.R. Bunch and B.N. Parlett, "Direct methods for solving symmetric indefinite systems of linear equations", *SIAM J. Numer. Anal.* 8 (1971), 639-655.

[8] J.J. Dongarra, J.R. Bunch, C.B. Moler and G.W. Stewart, *LINPACK users' guide* (Society for Industrial and Applied Mathematics, Philadelphia, 1979).

[9] L. Mirsky, *An introduction to linear algebra* (Clarendon, Oxford, 1955).

[10] B.N. Parlett and J.K. Reid, "On the solution of a system of linear equations whose matrix is symmetric but not definite", *BIT* 10 (1970), 386-397.

[11] J.H. Wilkinson, *The algebraic eigenvalue problem* (Clarendon, Oxford, 1965).

Department of Mathematics,
University of California, San Diego,
La Jolla,
California 92093,
USA.