Linear Algebra Appl. 452 (2014) 40–45

Contents lists available at ScienceDirect

Linear Algebra and its Applications

www.elsevier.com/locate/laa

Symmetric elimination without pivoting

Paul Van Dooren [a,*], Gilbert Strang [b]

[a] *Department of Mathematical Engineering, Université catholique de Louvain, B-1348 Louvain-la-Neuve, Belgium*
[b] *Department of Mathematics, MIT, Cambridge, MA 02139, USA*

ARTICLE INFO

Article history:
Received 12 September 2013
Accepted 19 March 2014
Available online xxxx
Submitted by Z. Strakoš

MSC:
65F05
15A23

Keywords:
Inertia
Symmetric factorization

ABSTRACT

Positive definite matrices factor into $A = LL^{\mathrm{T}}$ (Cholesky). Symmetric indefinite matrices need a symmetric middle factor in $A = LPL^{\mathrm{T}}$. Then A and P have the same inertia (eigenvalues of the same sign). We construct P through elimination, so the inertias agree for all leading minors of A and P. When restricting P to be a variant of a symmetric permutation in which diagonal 1's can be replaced by 0's or -1's, it is unique.

1. Main result

This paper is about matrix algebra, not matrix analysis. The matrices are symmetric and generally indefinite. Pivoting (row exchange) may be numerically necessary in Gaussian elimination, but we don't do it. If the matrix were positive definite, we would create its Cholesky factorization $A = LL^{\mathrm{T}}$. In the general case, L can remain lower triangular and invertible, but the factorization then needs a symmetric matrix P between L and L^{T}:

$$A = LPL^{\mathrm{T}}.$$

* Corresponding author.
E-mail addresses: paul.vandooren@uclouvain.be (P. Van Dooren), gs@math.mit.edu (G. Strang).

http://dx.doi.org/10.1016/j.laa.2014.03.026

There are many such factorizations, unless we impose P to have a special structure. We require that P is almost a symmetric permutation matrix, except that elements 1 on its main diagonal may be replaced by 0 or -1. We will call this *Property A* and define it formally as follows

Property A: P is symmetric and has at most one nonzero entry in each row and column and $P(i,i) \in \{0,1,-1\}$, $P(i,j) = P(j,i) \in \{0,1\}$.

Then P^2 is diagonal with 1's or 0's on its diagonal. It is the orthogonal projector on the range of P. Therefore P is its own Moore–Penrose inverse (pseudoinverse): $P = P^\dagger$. This explains the qualification "almost a symmetric permutation matrix". We will call P the **inertia matrix**[1] **of A**.

$A = LPL^{\mathrm{T}}$ says that P is "congruent" to A—perhaps an old-fashioned word. The Inertia Theorem [4,6] asserts that P and A have the same inertia:

$$\mathrm{In}(P) = \{n, z, p\} = \mathrm{In}(A).$$

P and A have the same number of negative eigenvalues, zero eigenvalues, and positive eigenvalues. More than this, the inertias of all the upper left square submatrices P_k and A_k also agree:

Property I:

$$\mathrm{In}(P_k) = \{n_k, z_k, p_k\} = \mathrm{In}(A_k).$$

The point is that each $A_k = L_k P_k L_k^{\mathrm{T}}$ because L is upper triangular. Notice that P_k also inherits Property **A** from the matrix P, from the normal descending order of elimination. Because we impose Property **A** on P, we show it is unique.

Theorem. *Every real symmetric matrix A factors into $A = LPL^{\mathrm{T}}$ with a lower triangular invertible matrix L and a unique inertia matrix P that has Property **A**. It then follows that P has also Property **I**.* □

There are ten 2×2 inertia matrices: nine diagonal matrices of 1's or 0's or -1's and the row exchange matrix.

The proof of our main theorem will be constructive and elementary. Each upper left submatrix $A_{k-1} = L_{k-1} P_{k-1} L_{k-1}^{\mathrm{T}}$ is *bordered* by a row and column to create $A_k = L_k P_k L_k^{\mathrm{T}}$:

$$\begin{bmatrix} A_{k-1} & a_k \\ a_k^{\mathrm{T}} & b_k \end{bmatrix} = \begin{bmatrix} L_{k-1} & 0 \\ \ell_k^{\mathrm{T}} & d_k \end{bmatrix} \begin{bmatrix} P_k \end{bmatrix} \begin{bmatrix} L_{k-1}^{\mathrm{T}} & \ell_k \\ 0^{\mathrm{T}} & d_k \end{bmatrix}.$$

[1] The same term is used in mechanics but there it clearly has a different meaning

The chief pleasure lies in identifying the four possible changes of inertia n, z, p at each step of symmetric elimination. This is captured in the following lemma.

Bordering Lemma. *At each step from A_{k-1} and P_{k-1} to A_k and P_k, the count of negative-zero-positive eigenvalues satisfies*

$$n_{k-1} \leqslant n_k \leqslant n_{k-1} + 1,$$
$$z_{k-1} - 1 \leqslant z_k \leqslant z_{k-1} + 1,$$
$$p_{k-1} \leqslant p_k \leqslant p_{k-1} + 1.$$

The sum $n + z + p = k - 1$ increases by 1 to k. So there are four possible changes from the inertia $\mathrm{In}(A_{k-1}) = (n_{k-1}, z_{k-1}, p_{k-1})$ to $\mathrm{In}(A_k)$:

1. $\{n_k, z_k, p_k\} = \{n_{k-1}, z_{k-1} + 1, p_{k-1}\}$,
2. $\{n_k, z_k, p_k\} = \{n_{k-1}, z_{k-1}, p_{k-1} + 1\}$,
3. $\{n_k, z_k, p_k\} = \{n_{k-1} + 1, z_{k-1}, p_{k-1}\}$,
4. $\{n_k, z_k, p_k\} = \{n_{k-1} + 1, z_{k-1} - 1, p_{k-1} + 1\}$.

That fourth possibility is illustrated by the exchange matrix: $z_1 = 1$ but $z_2 = 0$.

$$A = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} = P \quad \textit{with } \mathrm{In}(A_1) = (0, 1, 0) \textit{ and } \mathrm{In}(A) = (1, 0, 1).$$

The proof of this lemma is based on the Cauchy interlacing property. Let the ordered eigenvalues of the symmetric matrices A_{k-1} and A_k be μ_i and λ_i, with inertias $\{n, z, p\}$ and $\{N, Z, P\}$:

$$\mu_1 \leqslant \cdots \leqslant \mu_n < \mu_{n+1} = \cdots = \mu_{n+z} < \mu_{n+z+1} \leqslant \cdots \leqslant \mu_{n+z+p}$$
$$\lambda_1 \leqslant \cdots \leqslant \lambda_N < \lambda_{N+1} = \cdots = \lambda_{N+Z} < \lambda_{N+Z+1} \leqslant \cdots \leqslant \mu_{N+Z+P}.$$

The Cauchy interlacing property [2] says that

$$\lambda_j \leqslant \mu_j \leqslant \lambda_{j+1} \quad \text{for } 1 \leqslant j \leqslant n + z + p.$$

This implies that A_k has at least n negative eigenvalues, $z - 1$ zero eigenvalues and p positive eigenvalues. The count $n + z + p + 1 = N + Z + P$ yields the desired bounds and therefore also the four cases shown above. □

In the recursive construction of the matrices L and P, suppose we have reached $A_{k-1} = L_{k-1} P_{k-1} L_{k-1}^{\mathrm{T}}$. Then we also have a first decomposition for the bordered matrix:

$$\begin{bmatrix} A_{k-1} & a_k \\ a_k^{\mathrm{T}} & b_k \end{bmatrix} = \begin{bmatrix} L_{k-1} & \\ 0 & 1 \end{bmatrix} \begin{bmatrix} P_{k-1} & c_k \\ c_k^{\mathrm{T}} & b_k \end{bmatrix} \begin{bmatrix} L_{k-1}^{\mathrm{T}} & 0 \\ & 1 \end{bmatrix} \tag{1}$$

where c_k is the unique solution of $L_{k-1}c_k = a_k$. We now eliminate as much as possible of the vector c_k by using the elimination vector $\ell_k = P_{k-1}c_k$:

$$\begin{bmatrix} P_{k-1} & c_k \\ c_k^{\mathrm{T}} & b_k \end{bmatrix} = \begin{bmatrix} I_{k-1} & \\ \ell_k^{\mathrm{T}} & 1 \end{bmatrix} \begin{bmatrix} P_{k-1} & e_k \\ e_k^{\mathrm{T}} & f_k \end{bmatrix} \begin{bmatrix} I_{k-1} & \ell_k \\ & 1 \end{bmatrix}, \tag{2}$$

where $e_k = c_k - P_{k-1}\ell_k = (I_{k-1} - P_{k-1}^2)c_k$. That vector is orthogonal to the range of P_{k-1} since it satisfies $P_{k-1}e_k = 0$. The new diagonal entry is $f_k = b_k - c_k^{\mathrm{T}} P_{k-1} c_k$. Combining (1) and (2) we have reached

$$\begin{bmatrix} A_{k-1} & a_k \\ a_k^{\mathrm{T}} & b_k \end{bmatrix} = \begin{bmatrix} L_{k-1} & \\ \ell_k^{\mathrm{T}} & 1 \end{bmatrix} \begin{bmatrix} P_{k-1} & e_k \\ e_k^{\mathrm{T}} & f_k \end{bmatrix} \begin{bmatrix} L_{k-1}^{\mathrm{T}} & \ell_k \\ & 1 \end{bmatrix}. \tag{3}$$

If $e_k = 0$ the construction is essentially complete. We are in cases **1**–**3**, where the diagonal entry f_k should become 0 or 1 or -1. For this we just choose the positive diagonal entry d_k of L:

Cases **1**, **2**, **3**. Choose $d_k = 1$ or $\sqrt{f_k}$ or $\sqrt{-f_k}$. Then $p_k = 0$ or 1 or -1.

$$\begin{bmatrix} A_{k-1} & a_k \\ a_k^{\mathrm{T}} & b_k \end{bmatrix} = \begin{bmatrix} L_{k-1} & \\ \ell_k^{\mathrm{T}} & d_k \end{bmatrix} \begin{bmatrix} P_{k-1} & 0 \\ 0^{\mathrm{T}} & p_k \end{bmatrix} \begin{bmatrix} L_{k-1}^{\mathrm{T}} & \ell_k \\ & d_k \end{bmatrix}. \tag{4}$$

If $e_k \neq 0$ in (3), we are in case **4**. The nullity drops to $z_k = z_{k-1} - 1$ and there is more work to do. Let i be the index of the *first nonzero entry* of e_k. If the unit column vector $u_i = (0, \ldots, 1, \ldots, 0)$ has 1 in position i, then *we want u_i to replace e_k in P_k*, so that every row and column of P_k has at most one nonzero element. This then ensures that Property **A** is preserved in P_k.

$P_{k-1}e_k = 0$ implies that $P_{k-1}u_i = 0$. The triangular update matrix L_{up} then gives the desired inertia matrix P_k starting from (3):

$$\begin{bmatrix} P_{k-1} & e_k \\ e_k^{\mathrm{T}} & f_k \end{bmatrix} = L_{up} \begin{bmatrix} P_{k-1} & u_i \\ u_i^{\mathrm{T}} & 0 \end{bmatrix} L_{up}^{\mathrm{T}} \quad \text{if } L_{up} = \begin{bmatrix} I_{k-1} + (e_k - u_i)u_i^{\mathrm{T}} & 0 \\ f_k u_i^{\mathrm{T}}/2 & 1 \end{bmatrix}. \tag{5}$$

Now the inertial factorization $A_k = L_k P_k L_k^{\mathrm{T}}$ is complete:

$$L_k := \begin{bmatrix} L_{k-1} & \\ \ell_k^{\mathrm{T}} & 1 \end{bmatrix} L_{up} \quad \text{and} \quad P_k := \begin{bmatrix} P_{k-1} & u_i \\ u_i^{\mathrm{T}} & 0 \end{bmatrix}. \tag{6}$$

This covers the four cases and completes the proof of our main theorem. □

P_k is clearly unique in cases **1**, **2**, **3**. Uniqueness in case **4** comes from the uniqueness of i (the index of the first nonzero entry).

L is not unique. But the recursive construction shows that the only degrees of freedom lie in the diagonal elements of L that correspond to nonzero diagonal elements in P (cases **2** and **3**), and in the columns of L that correspond to zero columns of P (case **1**).

A corresponding decomposition for complex Hermitian matrices $H = LPL^*$ is straightforward. L is now a complex triangular (and nonsingular) matrix, and P still has Properties **A** and **I**.

Notice that P is completely defined by the inertias of the leading minors of A. As a consequence we have the following corollary.

Corollary. *Two symmetric matrices A and B are congruent with a lower triangular invertible matrix L if and only if they have the same inertia matrix P. It follows directly that the inertia matrix is a canonical form under the group of lower triangular congruence transformations.* □

2. Symmetric factorizations

This decomposition is of course similar to the symmetric factorization of Bunch–Parlett and its variants [3]. A is still symmetric and possibly indefinite, and is factorized as:

$$PAP^{\mathrm{T}} = LDL^{\mathrm{T}}$$

where the permutation P is crucial to make sure that L has bounded elements, and that D is made of 1×1 or 2×2 diagonal blocks. There exist pivoting strategies with reasonable complexity that ensure boundedness of elements of L, but the factorization says nothing about the inertias of the leading principal minors of A. We refer to [3, Section 11.1] for a discussion on complete, partial and rook pivoting. On the other hand, the inertias of the leading principal minors of PAP^{T} are easily derived from the small blocks in D.

Another related result is the anti-triangular factorization of Mastronardi and Van Dooren [5], which is a decomposition of the form

$$UAU^{\mathrm{T}} = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & Y^{\mathrm{T}} \\ 0 & 0 & X & Z^{\mathrm{T}} \\ 0 & Y & Z & W \end{bmatrix}.$$

U is an orthogonal transformation, X and W are symmetric, and X and Y are invertible. Here the dimensions of the blocks on the right hand side reveal the inertia of A.

Both these factorizations are backward stable, because much care is taken in constructing the transformations P and U. But they do not yield the information provided by the present factorization. Our decomposition, on the other hand, is unstable since we are not allowed to perform any pivoting. A backward stable alternative to our decomposition would be to use a stable decomposition to each of the leading submatrices A

but that would result in an $\mathcal{O}(n^4)$ algorithm. Notice that P is not a continuous function of the parameters of A when some of the leading submatrices are singular, since arbitrary small perturbations will result in a different matrix P. Using a backward stable algorithm does not affect the possible ill-condition of the problem.

The inertia of the leading principal minors is e.g. important when dealing with symmetric (block)-Hankel matrices. It was shown in [1] that the rank pattern of the growing block Hankel matrix

$$\mathcal{H} = \begin{bmatrix} H_1 & H_2 & H_3 & \dots \\ H_2 & H_3 & \ddots & \ddots \\ H_3 & \ddots & \ddots & \ddots \\ \vdots & \ddots & \ddots & \ddots \end{bmatrix},$$

decides the controllability and observability indices of the realization of the formal power series $H(z) = \sum_{i=1}^{\infty} H_i z^{-i}$. This is important since these indices describe how much time is needed to drive a particular state to zero.

Acknowledgements

The work of the first author is partly supported by the Belgian Network DYSCO (Dynamical Systems, Control, and Optimization), funded by the Interuniversity Attraction Poles Programme, initiated by the Belgian State, Science Policy Office (grant VII/19 (2012–2017)). The scientific responsibility rests with the authors. The second author was supported by the National Science Foundation (grant 1023152) and by a Mathworks Professorship.

References

[1] O. Bosgra, On parametrizations for the minimal partial realization problem, Systems Control Lett. 3 (Sept. 1983) 181–187.
[2] G. Golub, C. Van Loan, Matrix Computations, 4th ed., Johns Hopkins University Press, Baltimore, 2013.
[3] N. Higham, Accuracy and Stability of Numerical Algorithms, SIAM Publications, Philadelphia, 2013.
[4] P. Lancaster, M. Tismenetsky, The Theory of Matrices, 2nd ed., Academic Press, Orlando, FL, 1990.
[5] N. Mastronardi, P. Van Dooren, An algorithm for computing the anti-triangular factorization of symmetric matrices, SIAM J. Matrix Anal. Appl. 34 (1) (2013) 173–196.
[6] G. Strang, Linear Algebra and Its Applications, 4th ed., Cengage Learning, 2009.