

2023-02-17

1 Tricky tridiagonals

Consider the tridiagonal matrix

$$T = \begin{bmatrix} a_1 & b_1 & & & & & \\ c_2 & a_2 & b_2 & & & & \\ & c_3 & a_3 & b_3 & & & \\ & & \ddots & \ddots & \ddots & & \\ & & & c_{n-1} & a_{n-1} & b_{n-1} & \\ & & & & c_n & a_n & \end{bmatrix}$$

What happens if we factor using Gaussian elimination? We'll consider first the case of no pivoting. At the first step of elimination, we subtract a multiple of the first row from the second row, introducing a zero in the $(2, 1)$ position. Then we subtract a multiple of the second row from the third row, and so on. At each step, we only work with two rows of the matrix. Let's write the first two steps in pictures, assuming an implementation where we are systematically overwriting the original matrix with lower and upper triangular factors. At each step, we color the entries of the matrix that are transformed. In the first step, we have

$$\begin{bmatrix} a_1 & b_1 & & & & & \\ c_2 & a_2 & b_2 & & & & \\ & c_3 & a_3 & b_3 & & & \\ & & \ddots & \ddots & \ddots & & \\ & & & c_{n-1} & a_{n-1} & b_{n-1} & \\ & & & & c_n & a_n & \end{bmatrix} \rightarrow \begin{bmatrix} a_1 & b_1 & & & & & \\ \mathbf{l}_2 & \hat{\mathbf{a}}_2 & b_2 & & & & \\ & c_3 & a_3 & b_3 & & & \\ & & \ddots & \ddots & \ddots & & \\ & & & c_{n-1} & a_{n-1} & b_{n-1} & \\ & & & & c_n & a_n & \end{bmatrix}$$

At the next step, we have

$$\begin{bmatrix} a_1 & b_1 & & & & & \\ c_2 & a_2 & b_2 & & & & \\ & c_3 & a_3 & b_3 & & & \\ & & \ddots & \ddots & \ddots & & \\ & & & c_{n-1} & a_{n-1} & b_{n-1} & \\ & & & & c_n & a_n & \end{bmatrix} \rightarrow \begin{bmatrix} a_1 & b_1 & & & & & \\ l_2 & \hat{a}_2 & b_2 & & & & \\ & \mathbf{l}_3 & \hat{\mathbf{a}}_3 & b_3 & & & \\ & & \ddots & \ddots & \ddots & & \\ & & & c_{n-1} & a_{n-1} & b_{n-1} & \\ & & & & c_n & a_n & \end{bmatrix}$$

At the end of the iteration, we have

$$\begin{bmatrix} a_1 & b_1 & & & & \\ c_2 & a_2 & b_2 & & & \\ & \ddots & \ddots & \ddots & & \\ & & c_{n-1} & a_{n-1} & b_{n-1} & \\ & & & c_n & a_n & \end{bmatrix} = \begin{bmatrix} 1 & & & & & \\ l_2 & 1 & & & & \\ & \ddots & \ddots & & & \\ & & l_{n-1} & 1 & & \\ & & & l_n & 1 & \end{bmatrix} \begin{bmatrix} a_1 & b_1 & & & & \\ & \hat{a}_2 & b_2 & & & \\ & & \ddots & \ddots & & \\ & & & \hat{a}_{n-1} & b_{n-1} & \\ & & & & \hat{a}_n & \end{bmatrix}$$

In code, we go from the standard Gaussian elimination loop

```

1 # Overwrite A with L and U
2 for j = 1:n-1
3   A[j+1:n,j] /= A(j,j);
4   A[j+1:n,j+1:n] -= A[j+1:n,j]*A[j,j+1:n]'
5 end

```

to the simplified loop

```

1 # Overwrite T with L and U
2 for j = 1:n-1
3   T[j+1,j] /= T[j,j]
4   T[j+1,j+1] -= T[j+1,j]*T[j,j+1]
5 end

```

The former loop does $O(n^2)$ work per step, while the latter does $O(1)$ work per step; that is, factoring a tridiagonal costs $O(n)$ work per time, not the $O(n^3)$ cost of factoring an ordinary matrix.

Of course, since most of the entries in a tridiagonal are zero, we waste space by storing every entry explicitly, just as we waste time if we use the general-purpose Gaussian elimination procedure. In general, it is better to store only the nonzero entries of the matrix. For example, in the case of the tridiagonal matrix, we might store the a , b , and c coefficients as vectors, in which case Gaussian elimination looks like

```

1 # Overwrite a with the L entries, and b and c with the U entries
2 for j = 1:n-1
3   a[j+1] /= b[j]
4   b[j+1] -= a[j+1]*c[j]
5 end

```

We can similarly write $O(n)$ time routines for forward and backward substitution with the (bidiagonal) triangular factors of T .

2 Consider sparsity

Tridiagonal matrices (and more generally banded matrices) belong to a more general family of *sparse* matrices, or matrices in which the vast majority of the entries are zero.

Sparse matrices are a bit different from dense matrices. In a dense matrix, we store the elements of the matrix in a fixed order. There is no need to specify that a value at a specific location in memory corresponds to a particular matrix element; that information is implicit. For special sparse matrices, such as banded matrices, one can again come up with storage formats that have this property. But in a general sparse matrix format, though we only store the nonzero entries, we also have to store the location of those entries. This has several consequences. One is that memory access for sparse matrix computations is less regular than in the dense case, and so they run at a lower flop rate than dense computations (though they typically do far fewer flops than corresponding dense computations might, so it's still a win). Also, operations that are trivial in the dense case become more complicated in the sparse case. For example, adding two dense matrices is easy, but adding two sparse matrices can be rather complicated, depending on the data structure that is used to store them.

The one operation on sparse matrices that is really simple is matrix-vector multiplication. This costs $O(\text{nnz})$ time, where nnz is the number of nonzeros in the matrix. We will see that we can build iterative solvers that just use this operation, but that's a topic for another day. For now, the question is: can we do fast Gaussian elimination on sparse matrices? The answer is yes, but whether it's a good idea or not depends a lot on the matrix. We turn to why this is the case in the next section.

3 General sparse direct methods

Suppose A is a sparse matrix, and $PA = LU$. Will L and U also be sparse? The answer depends in a somewhat complicated way on the structure of the graph associated with the matrix A , the pivot order, and the order in which variables are eliminated. Except in very special circumstances, there will generally be more nonzeros in L and U than there are in A ; these extra nonzeros are referred to as *fill*. There are two standard ideas for minimizing fill:

1. Apply a *fill-reducing ordering* to the variables; that is, use a factorization

$$PAQ = LU,$$

where Q is a column permutation chosen to approximately minimize the fill in L and U , and P is the row permutation used for stability.

The problem of finding an elimination order that minimizes fill is NP-hard, so it is hard to say that any ordering strategy is really optimal. But there is canned software for some heuristic orderings that tend to work well in practice. From a practical perspective, then, the important thing is to remember that a fill-reducing elimination order tends to be critical to using sparse Gaussian elimination in practice.

2. Relax the standard partial pivoting condition, choosing the row permutation P to balance the desire for numerical stability against the desire to minimize fill.

For the rest of this lecture, we will consider the simplified case of *structurally* symmetric matrices and factorization without pivoting (which is stable for diagonally dominant systems and positive definite systems).

4 Sparse matrices, graphs, and tree elimination

Consider the following illustrative example of how factoring a sparse matrix can lead to more or less dense factors depending on the order of elimination. Putting in \times to indicate a nonzero element, we have

$$\begin{bmatrix} \times & \times & \times & \times & \times \\ \times & \times & & & \\ \times & & \times & & \\ \times & & & \times & \\ \times & & & & \times \end{bmatrix} = \begin{bmatrix} \times & & & & \\ \times & \times & & & \\ \times & \times & \times & & \\ \times & \times & \times & \times & \\ \times & \times & \times & \times & \times \end{bmatrix} \begin{bmatrix} \times & \times & \times & \times & \times \\ & \times & \times & \times & \times \\ & & \times & \times & \times \\ & & & \times & \times \\ & & & & \times \end{bmatrix}.$$


That is, L and U have many more nonzeros than A . These nonzero locations that appear in L and U and not in A are called *fill-in*. On the other hand,

if we cyclically permute the rows and columns of A , we have

$$\begin{bmatrix} \times & & & \times \\ & \times & & \times \\ & & \times & \times \\ & & & \times & \times \\ \times & \times & \times & \times & \times \end{bmatrix} = \begin{bmatrix} \times & & & & \\ & \times & & & \\ & & \times & & \\ & & & \times & \\ \times & \times & \times & \times & \times \end{bmatrix} \begin{bmatrix} \times & & & \times \\ & \times & & \times \\ & & \times & \times \\ & & & \times & \times \\ & & & & \times & \times \end{bmatrix}.$$

That is, the factorization of PAP^T has *no* fill-in.

A sparse matrix A can be viewed as an *adjacency matrices* for an associated graphs: make one node for each row, and connect node i to node j if $A_{ij} \neq 0$. The graphs for the two “arrow” matrices above are:



These graphs of both our example matrices are *trees*, and they differ only in how the nodes are labeled. In the original matrix, the root node is assigned the first label; in the second matrix, the root node is labeled after all the children. Clearly, the latter label order is superior for Gaussian elimination. This turns out to be a general fact: if the graph for a (structurally symmetric) sparse matrix S is a tree, and if the labels are ordered so that each node appears after any children it may have, then there is no fill-in: that is, L and U have nonzeros only where S has nonzeros.

Why should we have no fill when factoring a matrix for a tree ordered from the leaves up? To answer this, we think about what happens in the first step of Gaussian elimination. Our original matrix has the form

$$S = \begin{bmatrix} \alpha & w^T \\ v & S_{22} \end{bmatrix}$$

The first row of U is identical to the first row of S , and the first column of L has the same nonzero structure as the first column of A , so we are fine there.


The only question is about the nonzero structure of the Schur complement $S_{22} - vw^T/\alpha$. Note that the update vw^T/α has nonzeros only where v_i and w_j are both nonzero — that is, only when nodes i and j are both connected to node 1. But node 1 is a leaf node; the only thing it connects to is its parent! So if p is the index of the parent of node 1 in the tree, then we only change the (p, p) entry of the trailing submatrix during the update — and we assume that entry is already nonzero. Thus, the graph associated with the Schur complement is the same as the graph of the original matrix, but with one leaf trimmed off.

5 Nested dissection

Tree-structured matrices are marvelous because we can do everything in $O(n)$ time: we process the tree from the leaves to the root in order to compute L and U , then recurse from the root to the leaves in order to do back substitution with U , and then go back from the leaves to the root in order to do forward substitution with L . Sadly, many of the graphs we encounter in practice do not look like trees. However, we can often profitably think of clustering nodes so that we get a *block* structure associated with a tree.

For illustrative purposes, let us consider Gaussian elimination on a matrix whose graph is a regular $n \times n$ mesh. Such a matrix might arise, for example, if we were solving Poisson's equation using a standard five-point stencil to discretize the Laplacian operator. We then think of cutting the mesh in half by removing a set of separator nodes, cutting the halves in half, and so forth. This yields a block structure of a tree consisting of a root (the separator nodes) and two children (the blocks on either side of the separator). We can now dissect each of the sub-blocks with a smaller separator, and continue on in this fashion until we have cut the mesh into blocks containing only a few nodes each. Figure 1 illustrates the first two steps in this process of *nested dissection*.

We can get a lower bound on the cost of the factorization by figuring out the cost of factoring the Schur complement associated with G , C , F , etc. After we eliminate everything except the nodes associated with G , we pay about $2n^3/3$ flops to factor the remaining (dense) n -by- n Schur complement matrix G . Similarly, we pay about $2(n/2)^3/3$ time to factor the dense $(n/2)$ -by- $(n/2)$ complements associated with the separators C and F . Eliminating all four separators then costs a total of $\approx 10n^3/12$ flops. Now, think of applying



$$S = \begin{bmatrix} S_{AA} & & S_{AC} & & & & & & & S_{AG} \\ & S_{BB} & S_{BC} & & & & & & & S_{BG} \\ S_{CA} & S_{CB} & S_{CC} & & & & & & & S_{CG} \\ & & & S_{DD} & & S_{DF} & & & & S_{DG} \\ & & & & S_{EE} & S_{EF} & & & & S_{EG} \\ & & & S_{FD} & S_{FE} & S_{FF} & & & & S_{FG} \\ S_{GA} & S_{GB} & S_{GC} & S_{GD} & S_{GE} & S_{GF} & S_{GG} & & & \end{bmatrix}$$

Figure 1: Nested dissection on a square mesh. We first cut the graph in half with the red separator G , then further dissect the halves with the blue separators C and F . Nodes in A , B , D , and E are only connected through these separator nodes, which is reflected in the sparsity pattern of the adjacency matrix S when it is ordered so that separators appear after the things they separate.

nested dissection to blocks A , B , D , and E ; eliminating the Schur complements associated with separators inside each of these blocks will take about $5(n/2)^3/6$ flops; all four together cost a total of $4(5(n/2)^3/6) = (1/2)(5n^3/6)$ flops to factor. If we keep recursing, we find that the cost of factoring Schur complements associated with all the separators looks like

$$\frac{5}{6}n^3 \left(1 + \frac{1}{2} + \frac{1}{4} + \dots \right) \approx \frac{5}{3}n^3.$$

It turns out that forming each Schur complement is asymptotically not more expensive than eliminating it, so that the overall cost of doing nested dissection on an $n \times n$ mesh with $N = n^2$ unknowns is also $O(n^3) = O(N^{1.5})$. It also turns out that the fill-in is $O(N \log N)$ ¹.

Now think about doing the same thing with a three-dimensional mesh. In this case, the top-level separators for an $n \times n \times n$ mesh with $N = n^3$ unknowns would involve n^2 unknowns, and we would take $O(n^6) = O(N^2)$ time to do the elimination, and $O(N^{4/3})$ fill. This relatively poor scaling explains why sparse direct methods are attractive for solving 2D PDEs, but are less popular for 3D problems.

6 E pluribus unum

So far, we have described a few ideas about how to perform Gaussian elimination for linear systems. But you might find yourself asking a reasonable — if a bit lazy — question: “why bother?” After all, the Julia backslash operator is a marvelous invention, and the simple expression

1 `x = A\b`

mostly “does the right thing” for a variety of types of matrices and factorizations, based partly on type information. If A is square and dense, this line causes Julia to factor the matrix A using Gaussian elimination, then carry out forward and backward substitution. If A is triangular, Julia does forward or backward substitution. If A is a sparse matrix, Julia uses the UMFPACK sparse LU package, including applying a reasonable column permutation. If

¹The explanation of why is not so hard, at least for regular 2D meshes, but it requires more drawing than I feel like at the moment. The paper “Nested Dissection of a Regular Finite Element Mesh” by Alan George (SIAM J. Numer. Anal. 10(2), April 1973) gives a fairly readable explanation for the curious.

Julia can do all this automatically for you, why do you need to know the details?

There's a deeper answer to this question than the superficial "because you're taking a numerical methods class." It even goes beyond needing to understand things like when a sparse system is best solved by a direct method vs. an iteration. One very important reason to understand the factorizations that are being computed behind the scenes is that those factorizations can be *reused* when you are solving more than one linear system at a go. And solving more than one problem at a time, as it turns out, is often what we want to do.