

## Week 6: Friday, Oct 2

### Logistics

1. I will be out of town next Tuesday, October 6, and so will not have office hours on that day. I will be around on Monday, except during the SCAN seminar (1:25-2:15); feel free to stop by my office with questions.
2. On HW2 p4: “negative elements” should be “nonzero elements”.

### Sparse matrices, graphs, and tree elimination

Recall the following example from last lecture. 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

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

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  $vw^T/\alpha$  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.

## 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*.

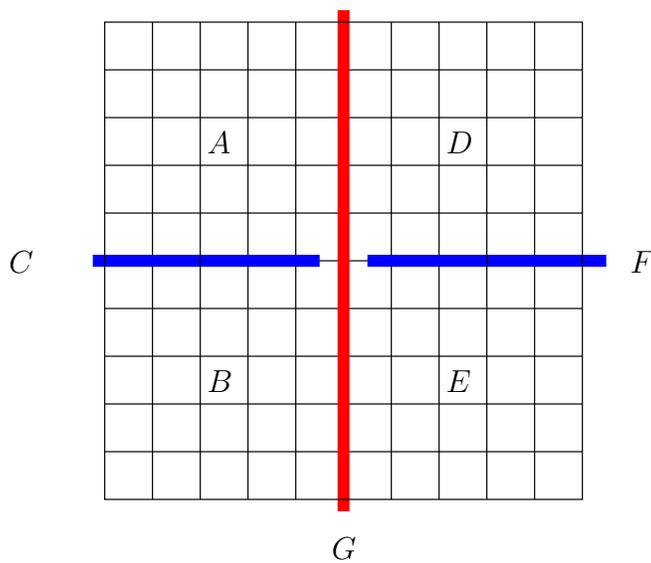
We can get a lower bound on the cost of the factorization by figuring out the cost of factoring the Schur complement associated with  $G$ ,  $C$ ,  $F$ , etc. After we eliminate everything except the nodes associated with  $G$ , we pay about  $2n^3/3$  flops to factor the remaining (dense)  $n$ -by- $n$  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 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  $O(N \log N)$ <sup>1</sup>.

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.

<sup>1</sup> The 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.



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

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  $F$  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.

## Other structured matrices

Band matrices with upper and lower bandwidth  $p$  and  $q$  are described by  $n(p+q+1)$  numbers, and the  $L$  and  $U$  factors can be described by a total of  $n(p+q+1)$  numbers.  $U$  and  $L$  have upper and lower bandwidth respectively of  $p$  and  $q$  in the absence of pivoting; with pivoting,  $U$  has upper bandwidth at most  $p+q$  and  $L$  is “essentially” lower banded with bandwidth  $q$ . The details are described in the book. Also described in the book are several other classes of problems for which the matrix depends on fewer than  $O(n^2)$  parameters, and there are ways in which those compact representations lead to fast factorization or solution algorithms. It’s worth knowing that these methods exist, but for the moment we will say little about them. We will, however, return to Toeplitz matrices and FFTs later in the course.