

2019-09-20

1 Packed storage

In LU factorization, the locations where we write the multipliers in L are exactly the same locations where we introduce zeros in A as we transform to U . Thus, we can re-use the storage space for A to store both L (except for the diagonal ones, which are implicit) and U . Using this strategy, we have the following code:

```

1 %
2 % Overwrite A with L and U factors
3 %
4 function [A] = mylu(A)
5     n = length(A);
6     for j=1:n-1
7         A(j+1:n,j) = A(j+1:n,j)/A(j,j);
8         A(j+1:n,j+1:n) = A(j+1:n,j+1:n) - A(j+1:n,j)*A(j,j+1:n);
9     end

```

If we wanted to extract the L and U factors explicitly, we could then do

```

1 LU = mylu(A);
2 L = eye(length(A)) + tril(A,-1);
3 U = triu(A);

```

The bulk of the work at step j of the elimination algorithm is in the computation of a rank-one update to the trailing submatrix. How much work is there in total? In eliminating column j , we do $(n-j)^2$ multiplies and the same number of subtractions; so in all, the number of multiplies (and adds) is

$$\sum_{j=1}^{n-1} (n-j)^2 = \sum_{k=1}^{n-1} k^2 = \frac{1}{3}n^3 + O(n^2)$$

We also perform $O(n^2)$ divisions. Thus, Gaussian elimination, like matrix multiplication, is an $O(n^3)$ algorithm operating on $O(n^2)$ data.

2 Schur complements

The idea of expressing a step of Gaussian elimination as a low-rank submatrix update turns out to be sufficiently useful that we give it a name. At

any given step of Gaussian elimination, the trailing submatrix is called a *Schur complement*. We investigate the structure of the Schur complements by looking at an LU factorization in block 2-by-2 form:

$$\begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix} = \begin{bmatrix} L_{11} & 0 \\ L_{21} & L_{22} \end{bmatrix} \begin{bmatrix} U_{11} & U_{12} \\ 0 & U_{22} \end{bmatrix} = \begin{bmatrix} L_{11}U_{11} & L_{11}U_{12} \\ L_{21}U_{11} & L_{22}U_{22} + L_{21}U_{12} \end{bmatrix}.$$

We can compute L_{11} and U_{11} as LU factors of the leading sub-block A_{11} , and

$$\begin{aligned} U_{12} &= L_{11}^{-1}A_{12} \\ L_{21} &= A_{21}U_{11}^{-1}. \end{aligned}$$

What about L_{22} and U_{22} ? We have

$$\begin{aligned} L_{22}U_{22} &= A_{22} - L_{21}U_{12} \\ &= A_{22} - A_{21}U_{11}^{-1}L_{11}^{-1}A_{12} \\ &= A_{22} - A_{21}A_{11}^{-1}A_{12}. \end{aligned}$$

This matrix $S = A_{22} - A_{21}A_{11}^{-1}A_{12}$ is the block analogue of the rank-1 update computed in the first step of the standard Gaussian elimination algorithm.

For our purposes, the idea of a Schur complement is important because it will allow us to write blocked variants of Gaussian elimination — an idea we will take up in more detail shortly. But the Schur complement actually has meaning beyond being a matrix that mysteriously appears as a by-product of Gaussian elimination. In particular, note that if A and A_{11} are both invertible, then

$$\begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix} \begin{bmatrix} X \\ S^{-1} \end{bmatrix} = \begin{bmatrix} 0 \\ I \end{bmatrix},$$

i.e. S^{-1} is the $(2, 2)$ submatrix of A^{-1} .

3 Blocked Gaussian elimination

Just as we could rewrite matrix multiplication in block form, we can also rewrite Gaussian elimination in block form. For example, if we want

$$\begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix} = \begin{bmatrix} L_{11} & 0 \\ L_{21} & L_{22} \end{bmatrix} \begin{bmatrix} U_{11} & U_{12} \\ 0 & U_{22} \end{bmatrix}$$

then we can write Gaussian elimination as:

1. Factor $A_{11} = L_{11}U_{11}$.
2. Compute $L_{21} = A_{21}U_{11}^{-1}$ and $U_{12} = L_{11}^{-1}A_{12}$.
3. Form the Schur complement $S = A_{22} - L_{21}U_{12}$ and factor $L_{22}U_{22} = S$.

This same idea works for more than a block 2-by-2 matrix. Suppose `idx` is a MATLAB vector that indicates the first index in each block of variables, so that block A_{IJ} is extracted as

```

1  I = idx(i):idx(i+1)-1;
2  J = idx(j):idx(j+1)-1;
3  A_IJ = A(idx(i):idx(i+1)-1, idx(j):idx(j+1)-1);

```

Then we can write the following code for block LU factorization:

```

1  M = length(idx)-1;      % Number of blocks
2  for j = 1:M
3      J = idx(j):idx(j+1)-1; % Indices for block J
4      rest = idx(j+1):n;    % Indices after block J
5      A(J,J) = lu(A(J,J));  % Factor submatrix A_JJ
6
7      % Extract L and U (this could be implicit)
8      L_JJ = tril(A(J,J),-1) + eye(length(J));
9      U_JJ = triu(A(J,J));
10
11     % Compute block column of L and row of U, Schur complement
12     A(rest,J) = A(rest,J)/U_JJ;
13     A(J,rest) = L_JJ\A(J,rest);
14     A(rest,rest) = A(rest,rest)-A(rest,J)*A(J,rest);
15 end

```

As with matrix multiply, thinking about Gaussian elimination in this blocky form lets us derive variants that have better cache efficiency. Notice that all the operations in this blocked code involve matrix-matrix multiplies and multiple back solves with the same matrix. These routines can be written in a cache-efficient way, since they do many floating point operations relative to the total amount of data involved.

Though some of you might make use of cache blocking ideas in your own work, most of you will never try to write a cache-efficient Gaussian elimination routine of your own. The routines in LAPACK and MATLAB (really the same routines) are plenty efficient, so you would most likely turn to them. Still, it is worth knowing how to think about block Gaussian elimination, because sometimes the ideas can be specialized to build fast solvers for linear systems when there are fast solvers for sub-matrices

For example, consider the *bordered* matrix

$$A = \begin{bmatrix} B & W \\ V^T & C \end{bmatrix},$$

where B is an n -by- n matrix for which we have a fast solver and C is a p -by- p matrix, $p \ll n$. We can factor A into a product of *block* lower and upper triangular factors with a simple form:

$$\begin{bmatrix} B & W \\ V^T & C \end{bmatrix} = \begin{bmatrix} B & 0 \\ V^T & L_{22} \end{bmatrix} \begin{bmatrix} I & B^{-1}W \\ 0 & U_{22} \end{bmatrix}$$

where $L_{22}U_{22} = C - V^TB^{-1}W$ is an ordinary (small) factorization of the trailing Schur complement. To solve the linear system

$$\begin{bmatrix} B & W \\ V^T & C \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} b_1 \\ b_2 \end{bmatrix},$$

we would then run block forward and backward substitution:

$$\begin{aligned} y_1 &= B^{-1}b_1 \\ y_2 &= L_{22}^{-1}(b_2 - V^T y_1) \end{aligned}$$

$$\begin{aligned} x_2 &= U_{22}^{-1}y_2 \\ x_1 &= y_1 - B^{-1}(Wx_2) \end{aligned}$$

4 Backward error in Gaussian elimination

Solving $Ax = b$ in finite precision using Gaussian elimination followed by forward and backward substitution yields a computed solution \hat{x} *exactly* satisfies

$$(1) \quad (A + \delta A)\hat{x} = b,$$

where $|\delta A| \lesssim 3n\epsilon_{\text{mach}}|\hat{L}||\hat{U}|$, assuming \hat{L} and \hat{U} are the computed L and U factors.

I will now briefly sketch a part of the error analysis following Demmel's treatment (§2.4.2). Mostly, this is because I find the treatment in §3.3.1 of Van Loan less clear than I would like – but also, the bound in Demmel's book

is marginally tighter. Here is the idea. Suppose \hat{L} and \hat{U} are the computed L and U factors. We obtain \hat{u}_{jk} by repeatedly subtracting $l_{ji}u_{ik}$ from the original a_{jk} , i.e.

$$\hat{u}_{jk} = \text{fl} \left(a_{jk} - \sum_{i=1}^{j-1} \hat{l}_{ji} \hat{u}_{ik} \right).$$

Regardless of the order of the sum, we get an error that looks like

$$\hat{u}_{jk} = a_{jk}(1 + \delta_0) - \sum_{i=1}^{j-1} \hat{l}_{ji} \hat{u}_{ik}(1 + \delta_i) + O(\epsilon_{\text{mach}}^2)$$

where $|\delta_i| \leq (j-1)\epsilon_{\text{mach}}$. Turning this around gives

$$\begin{aligned} a_{jk} &= \frac{1}{1 + \delta_0} \left(\hat{l}_{jj} \hat{u}_{jk} + \sum_{i=1}^{j-1} \hat{l}_{ji} \hat{u}_{ik}(1 + \delta_i) \right) + O(\epsilon_{\text{mach}}^2) \\ &= \hat{l}_{jj} \hat{u}_{jk}(1 - \delta_0) + \sum_{i=1}^{j-1} \hat{l}_{ji} \hat{u}_{ik}(1 + \delta_i - \delta_0) + O(\epsilon_{\text{mach}}^2) \\ &= \left(\hat{L}\hat{U} \right)_{jk} + e_{jk}, \end{aligned}$$

where

$$e_{jk} = -\hat{l}_{jj} \hat{u}_{jk} \delta_0 + \sum_{i=1}^{j-1} \hat{l}_{ji} \hat{u}_{ik} (\delta_i - \delta_0) + O(\epsilon_{\text{mach}}^2)$$

is bounded in magnitude by $(j-1)\epsilon_{\text{mach}}(|L||U|)_{jk} + O(\epsilon_{\text{mach}}^2)$ ¹. A similar argument for the components of \hat{L} yields

$$A = \hat{L}\hat{U} + E, \text{ where } |E| \leq n\epsilon_{\text{mach}}|\hat{L}||\hat{U}| + O(\epsilon_{\text{mach}}^2).$$

In addition to the backward error due to the computation of the LU factors, there is also backward error in the forward and backward substitution phases, which gives the overall bound (1).

¹ It's obvious that e_{jk} is bounded in magnitude by $2(j-1)\epsilon_{\text{mach}}(|L||U|)_{jk} + O(\epsilon_{\text{mach}}^2)$. We cut a factor of two if we go down to the level of looking at the individual rounding errors during the dot product, because some of those errors cancel.

5 Pivoting

The backward error analysis in the previous section is not completely satisfactory, since $|L||U|$ may be much larger than $|A|$, yielding a large backward error overall. For example, consider the matrix

$$A = \begin{bmatrix} \delta & 1 \\ 1 & 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ \delta^{-1} & 1 \end{bmatrix} \begin{bmatrix} \delta & 1 \\ 0 & 1 - \delta^{-1} \end{bmatrix}.$$

If $0 < \delta \ll 1$ then $\|L\|_\infty \|U\|_\infty \approx \delta^{-2}$, even though $\|A\|_\infty \approx 2$. The problem is that we ended up subtracting a huge multiple of the first row from the second row because δ is close to zero — that is, the leading principle minor is *nearly* singular. If δ were exactly zero, then the factorization would fall apart even in exact arithmetic. The solution to the woes of singular and near singular minors is pivoting; instead of solving a system with A , we re-order the equations to get

$$\hat{A} = \begin{bmatrix} 1 & 1 \\ \delta & 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ \delta & 1 \end{bmatrix} \begin{bmatrix} 1 & 1 \\ 0 & 1 - \delta \end{bmatrix}.$$

Now the triangular factors for the re-ordered system matrix \hat{A} have very modest norms, and so we are happy. If we think of the re-ordering as the effect of a permutation matrix P , we can write

$$A = \begin{bmatrix} \delta & 1 \\ 1 & 1 \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ \delta & 1 \end{bmatrix} \begin{bmatrix} 1 & 1 \\ 0 & 1 - \delta \end{bmatrix} = P^T LU.$$

Note that this is equivalent to writing $PA = LU$ where P is another permutation (which undoes the action of P^T).

If we wish to control the multipliers, it's natural to choose the permutation P so that each of the multipliers is at most one. This standard choice leads to the following algorithm:

```

1  for j = 1:n-1
2
3      % Find ipiv >= j to maximize |A(i,j)|
4      [absAij, ipiv] = max(abs(A(j:n,j)));
5      ipiv = ipiv + j-1;
6
7      % Swap row ipiv and row j
8      Aj = A(j,j:n);
```

```
9     A(j,j:n) = A(ipiv,j:n);
10     A(ipiv,j:n) = Aj;
11
12     % Record the pivot row
13     piv(j) = ipiv;
14
15     % Update trailing submatrix
16     A(j+1:n,j+1:n) = A(j+1:n,j+1:n) - A(j+1:n,j)*A(j,j+1:n);
17
18     end
```

In practice, we would typically use a strategy of *deferred updating*: that is, rather than applying the pivot immediately across all columns, we would only apply the pivoting within a block of columns. At the end of the block, we would apply all the pivots simultaneously. As with other blocking strategies we have discussed, this has no impact on the total amount of work done in some abstract machine model, but it is much more friendly to the memory architecture of real machines.

By design, this algorithm produces an L factor in which all the elements are bounded by one. But what about the U factor? There exist pathological matrices for which the elements of U grow exponentially with n . But these examples are extremely uncommon in practice, and so, in general, Gaussian elimination with partial pivoting does indeed have a small backward error. Of course, the pivot growth is something that we can monitor, so in the unlikely event that it *does* look like things are blowing up, we can tell there is a problem and try something different. But when problems do occur, it is more frequently the result of ill-conditioning in the problem than of pivot growth during the factorization.

6 Beyond partial pivoting

Gaussian elimination with partial pivoting has been the mainstay of linear system solving for many years. But the partial pivoting strategy is far from the end of the story! GECP, or Gaussian elimination with *complete* pivoting (involving both rows and columns), is often held up as the next step beyond partial pivoting, but this is really a strawman — though complete pivoting fixes the aesthetically unsatisfactory lack of backward stability in the partial pivoted variant, the cost of the GECP pivot search is more expensive than is usually worthwhile in practice. We instead briefly describe two other pivoting

strategies that are generally useful: *rook pivoting* and *tournament pivoting*. Next week, we will also briefly mention *threshold pivoting*, which is relevant to sparse Gaussian elimination.

6.1 Rook pivoting

In Gaussian elimination with rook pivoting, we choose a pivot at each step by choosing the largest magnitude element in the first row *or column* of the current Schur complement. This eliminates the possibility of exponential pivot growth that occurs in the partial pivoting strategy, but does not involve the cost of searching the entire Schur complement for a pivot (as occurs in the GECP case).

For the problem of solving linear systems, it is unclear whether rook pivoting really holds a practical edge over partial pivoting. The complexity is not really worse than partial pivoting, but there is more overhead (both in runtime and in implementation cost) to handle deferred pivoting for performance. Where rook pivoting has a great deal of potential is in Gaussian elimination on (nearly) *singular* matrices. If $A \in \mathbb{R}^{m \times n}$ has a large gap between σ_k and σ_{k+1} for $k < \min(m, n)$, then GERP on A tends to yield the factorization

$$PAQ = LU, U = \begin{bmatrix} U_{11} & U_{12} \\ 0 & U_{22} \end{bmatrix}$$

where $U_{11} \in \mathbb{R}^{k \times k}$ and $\|U_{22}\|$ is very small (on the order of σ_{k+1}).

Rook pivoting and the closely-related threshold rook pivoting are particularly useful in constrained optimization problems in which constraints can become redundant. Apart from its speed, the rook pivoting strategy has the advantage over other rank-revealing factorizations that when A is sparse, as one can often control the fill (nonzeros in L and U that are not present in A). The LUSOL package of Michael Saunders is a particularly effective example.

6.2 Tournament pivoting

In parallel dense linear algebra libraries, a major disadvantage of partial pivoting is that the pivot search is a communication bottleneck, even with deferred pivoting. This is increasingly an issue, as communication between processors is far more expensive than arithmetic, and (depending on the matrix layout) GEPP requires communication each time a pivot is selected.

For this reason, a number of recent *communication-avoiding* LU variants use an alternate pivoting strategy called *tournament pivoting*.

The idea behind *tournament pivoting* is to choose b pivot rows in one go, rather than iterating between choosing a pivot row and performing elimination. The algorithm involves each processor proposing several candidate pivot rows for a hierarchical tournament. There are different methods for managing this tournament, with different levels of complexity. One intriguing variant, for example, is the remarkable (though awkwardly named) CALU_PRRP algorithm, which uses rank-revealing QR factorizations to choose the pivots in the tournament. The CALU_PRRP algorithm does a modest amount of work beyond what is done by partial pivoting, but has better behavior both in terms of communication complexity and in terms of numerical stability.