SIAM J. MATRIX ANAL. APPL.
Vol. 20, No. 3, pp. 720–755

A SUPERNODAL APPROACH TO SPARSE PARTIAL PIVOTING*

JAMES W. DEMMEL†, STANLEY C. EISENSTAT‡, JOHN R. GILBERT§, XIAOYE S. LI¶, AND JOSEPH W. H. LIU‖

Abstract. We investigate several ways to improve the performance of sparse LU factorization with partial pivoting, as used to solve unsymmetric linear systems. We introduce the notion of unsymmetric supernodes to perform most of the numerical computation in dense matrix kernels. We introduce unsymmetric supernode-panel updates and two-dimensional data partitioning to better exploit the memory hierarchy. We use Gilbert and Peierls's depth-first search with Eisenstat and Liu's symmetric structural reductions to speed up symbolic factorization.

We have developed a sparse LU code using all these ideas. We present experiments demonstrating that it is significantly faster than earlier partial pivoting codes. We also compare its performance with UMFPACK, which uses a multifrontal approach; our code is very competitive in time and storage requirements, especially for large problems.

Key words. sparse matrix algorithms, unsymmetric linear systems, supernodes, column elimination tree, partial pivoting

AMS subject classifications. 65F05, 65F50

PII. S0895479895291765

1. Introduction.

The problem of solving sparse symmetric positive definite systems of linear equations on sequential and vector processors is fairly well understood. Normally, the solution process is performed in two phases:

- symbolic determination of the nonzero structure of the Cholesky factor;
- numeric factorization and solution.

Elimination trees [35] and compressed subscripts [41] reduce the time and space for symbolic factorization. Supernodal [5] and multifrontal [15] elimination allow the use of dense vector operations for nearly all of the floating-point computation, thus reducing the symbolic overhead in numeric factorization. Overall, the Megaflop rates of modern sparse Cholesky codes are nearly comparable to those of dense solvers [37, 38, 39].

*Received by the editors September 11, 1995; accepted for publication (in revised form) by A. Greenbaum May 15, 1998; published electronically April 14, 1999.

http://www.siam.org/journals/simax/20-3/29176.html

†Computer Science Division, University of California, Berkeley, CA 94720 (demmel@cs.berkeley.edu). The research of the first and fourth authors was supported in part by NSF grant ASC–9313958, DOE grant DE–FG03–94ER25219, UT Subcontract ORA4466 from ARPA Contract DAAL03–91–C0047, DOE grant DE–FG03–94ER25206, and NSF Infrastructure grants CDA–8722788 and CDA–9401156.

‡Department of Computer Science, Yale University, P.O. Box 208285, New Haven, CT 06520-8285 (sce@cs.yale.edu). The research of this author was supported in part by NSF grant CCR-9400921.

§Xerox Palo Alto Research Center, 3333 Coyote Hill Road, Palo Alto, CA 94304 (gilbert@xerox.com). The research of this author was supported in part by the Institute for Mathematics and Its Applications, University of Minnesota, and in part by DARPA Contract DABT63-95-C0087.

¶National Energy Research Scientific Computing (NERSC) Center, Lawrence Berkeley National Lab, 1 Cyclotron Rd., Berkeley, CA 94720 (xiaoye@nersc.gov). The research of this author was done at UC Berkeley and as a summer intern and consultant at Xerox PARC.

‖Department of Computer Science, York University, North York, ON, M3J 1P3, Canada, (joseph@cs.yorku.ca). The research of this author was supported in part by the Natural Sciences and Engineering Research Council of Canada under grant A5509.

```
for column j = 1 to n do
    f = A(:, j);
    Symbolic factor: determine which columns of L will update f;
    for each updating column r < j in topological order do
        Col-col update: f = f − f(r) · L(r + 1:n, r);
    end for;
    Pivot: interchange f(j) and f(m), where |f(m)| = max |f(j:n)|;
    Separate L and U: U(1:j, j) = f(1:j);    L(j:n, j) = f(j:n);
    Scale: L(:, j) = L(:, j)/L(j, j);
    Prune symbolic structure based on column j;
end for;
```

FIG. 1.1. *Left-looking LU factorization with column-column updates.*

For unsymmetric systems, where pivoting is required to maintain numerical stability, progress has been less satisfactory. Recent research has concentrated on two basic approaches: submatrix-based methods and column-based (or row-based) methods. Submatrix methods typically use some form of Markowitz ordering with threshold pivoting, in which each stage's pivot element is chosen from the uneliminated submatrix by criteria that attempt to balance numerical quality and preservation of sparsity. Recent submatrix codes include Amestoy and Duff's symmetric pattern multifrontal code MUPS [2] and Davis and Duff's unsymmetric multifrontal code UMFPACK [7].

Column methods, by contrast, typically use ordinary partial pivoting.[1] The pivot is chosen from the current column according to numerical considerations alone; the columns may be preordered before factorization to preserve sparsity. Figure 1.1 sketches a generic left-looking column LU factorization.[2] Notice that the bulk of the numeric computation occurs in column-column (col-col) updates, or, to use BLAS terminology [13, 14], in sparse AXPYs.

In column methods, the preordering for sparsity is completely separate from the factorization, just as in the symmetric case. This is an advantage when several matrices with the same nonzero structure but different numerical values must be factored. However, symbolic factorization cannot be separated from numeric factorization, because the nonzero structures of the factors depend on the numerical pivoting choices. Thus column codes must do some symbolic factorization at each stage; typically this amounts to predicting the structure of each column of the factors immediately before computing it. (George and Ng [22, 23] described ways to obtain upper bounds on the structure of the factors based only on the nonzero structure of the original matrix.) A disadvantage of the column methods is that they do not reorder the columns dynamically, so there may be more fill.

An early example of such a code is Sherman's NSPIV [42] (which is actually a row code). Gilbert and Peierls [29] showed how to use depth-first search and topological ordering to get the structure of each factor column. This gives a column code that runs in total time proportional to the number of nonzero floating-point operations, unlike the other partial pivoting codes. Eisenstat and Liu [21] designed a pruning

[1] Row methods are exactly analogous to column methods, and codes of both sorts exist. We will use column terminology; those who prefer rows may interchange the terms throughout the paper.

[2] We use Matlab notation for integer ranges and submatrices: $r{:}s$ or $(r{:}s)$ is the range of integers $(r, r+1, \ldots, s)$. If I and J are sets of integers, then $A(I, J)$ is the submatrix of A with rows whose indices are from I and with columns whose indices are from J. $A(:, J)$ abbreviates $A(1:n, J)$. $nnz(A)$ denotes the number of nonzeros in A.

technique to reduce the amount of structural information required for the symbolic factorization, as we describe further in section 4. The result is that the time and space for symbolic factorization are typically reduced to a small fraction of the entire factorization.

In view of the success of supernodal techniques for symmetric matrices, it is natural to consider the use of supernodes to enhance the performance of unsymmetric solvers. Like the nonzero structure of the factors, the boundaries of the supernodes cannot be determined in advance; rather, they emerge depending on pivoting choices during the factorization.

In this paper, we generalize supernodes to unsymmetric matrices, and we give efficient algorithms for locating and using unsymmetric supernodes during a column-based LU factorization. We describe a new code called SuperLU that uses depth-first search and symmetric pruning to speed up symbolic factorization, and uses unsymmetric supernodes to speed up numeric factorization. The rest of the paper is organized as follows. Section 2 introduces the tools we use: unsymmetric supernodes, panels, and the column elimination tree. Section 3 describes the supernodal numeric factorization. Section 4 describes the supernodal symbolic factorization. In section 5, we present experimental results: we benchmark our code on several test matrices, we compare its performance to other column and submatrix codes, and we investigate its cache behavior in some detail. Finally, section 6 presents conclusions and open questions.

2. Unsymmetric supernodes. The idea of a supernode is to group together columns with the same nonzero structure, so they can be treated as a dense matrix for storage and computation. Supernodes were originally used for (symmetric) sparse Cholesky factorization [5, 15]. In the factorization $A = LL^T$ (or $A = LDL^T$), a supernode is a range $(r{:}s)$ of columns of L with the same nonzero structure below the diagonal; that is, $L(r{:}s, r{:}s)$ is full lower triangular and every row of $L(s{:}n, r{:}s)$ is either full or zero.

Rothberg and Gupta [38, 39] and Ng and Peyton [37] analyzed the effect of supernodes in Cholesky factorization on modern uniprocessor machines with memory hierarchies and vector or superscalar hardware. All the updates from columns of a supernode are summed into a dense vector before the sparse update is performed. This reduces indirect addressing and allows the inner loops to be unrolled. In effect, a sequence of col-col updates is replaced by a supernode-column (sup-col) update. The sup-col update can be implemented using a call to a standard dense Level 2 BLAS matrix-vector multiplication kernel. This idea can be further extended to supernode-supernode (sup-sup) updates, which can be implemented using a Level 3 BLAS dense matrix-matrix kernel. This can reduce memory traffic by an order of magnitude, because a supernode in the cache can participate in multiple column updates. Ng and Peyton reported that a sparse Cholesky algorithm based on sup-sup updates typically runs 2.5 to 4.5 times as fast as a col-col algorithm. Indeed, supernodes have become a standard tool in sparse Cholesky factorization [5, 37, 38, 43].

To sum up, supernodes as the source of updates help because of the following:

1. The inner loop (over rows) has no indirect addressing. (Sparse Level 1 BLAS is replaced by dense Level 1 BLAS.)
2. The outer loop (over columns in the supernode) can be unrolled to save memory references. (Level 1 BLAS is replaced by Level 2 BLAS.)

Supernodes as the destination of updates help because of the following:

3. Elements of the source supernode can be reused in multiple columns of the destination supernode to reduce cache misses. (Level 2 BLAS is replaced by Level 3 BLAS.)

Supernodes in sparse Cholesky can be determined during symbolic factorization, before the numeric factorization begins. However, in sparse LU, the nonzero structure cannot be predicted before numeric factorization, so we must identify supernodes on the fly. Furthermore, since the factors L and U are no longer transposes of each other, we must generalize the definition of a supernode.

2.1. Definition of unsymmetric supernodes. We considered several possible ways to generalize the symmetric definition of supernodes to unsymmetric factorization. We define $F = L + U - I$ to be the *filled matrix* containing both L and U.

T1. Same row and column structures: A supernode is a range $(r{:}s)$ of columns of L and rows of U, such that the diagonal block $F(r{:}s, r{:}s)$ is full, and outside that block all the columns of L in the range have the same structure and all the rows of U in the range have the same structure. T1 supernodes make it possible to do sup-sup updates, realizing all three benefits.

T2. Same column structure in L: A supernode is a range $(r{:}s)$ of columns of L with triangular diagonal block full and the same structure below the diagonal block. T2 supernodes allow sup-col updates, realizing the first two benefits.

T3. Same column structure in L, full diagonal block in U: A supernode is a range $(r{:}s)$ of columns of L and U, such that the diagonal block $F(r{:}s, r{:}s)$ is full, and below the diagonal block the columns of L have the same structure. T3 supernodes allow sup-col updates, like T2. In addition, if the storage for a supernode is organized as for a two-dimensional (2-D) array (for Level 2 or 3 BLAS calls), T3 supernodes do not waste any space in the diagonal block of U.

T4. Same column structure in L and U: A supernode is a range $(r{:}s)$ of columns of L and U with identical structure. (Since the diagonal is nonzero, the diagonal block must be full.) T4 supernodes allow sup-col updates, and also simplify storage of L and U.

T5. Supernodes of A^TA: A supernode is a range $(r{:}s)$ of columns of L corresponding to a Cholesky supernode of the symmetric matrix A^TA. T5 supernodes are motivated by George and Ng's observation [22] that (with suitable representations) the structures of L and U in the unsymmetric factorization $PA = LU$ are contained in the structure of the Cholesky factor of A^TA. In unsymmetric LU, these supernodes themselves are sparse, so we would waste time and space operating on them. Thus we do not consider them further.

Figure 2.1 is a schematic of definitions T1 through T4.

Supernodes are only useful if they actually occur in practice. The occurrence of symmetric supernodes is related to the clique structure of the chordal graph of the Cholesky factor, which arises because of fill during the factorization. Unsymmetric supernodes seem harder to characterize, but they also are related to dense submatrices arising from fill. We measured the supernodes according to each definition for 126 unsymmetric matrices from the Harwell–Boeing sparse matrix test collection [17] under various column orderings. Table 2.1 shows, for each definition, the fraction of nonzeros of L that are not in the first column of a supernode; this measures how much row index storage is saved by using supernodes. Corresponding values for symmetric supernodes for the symmetric Harwell–Boeing structural analysis problems usually range from about 0.5 to 0.9. Larger numbers are better, indicating larger supernodes.

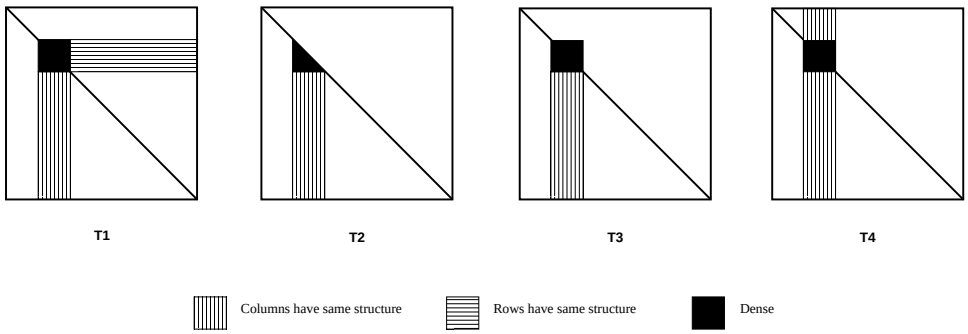

FIG. 2.1. *Four possible types of unsymmetric supernodes.*

TABLE 2.1
Fraction of nonzeros not in the first column of supernode.

	T1	T2	T3	T4
median	0.236	0.345	0.326	0.006
mean	0.284	0.365	0.342	0.052

We reject T4 supernodes as being too rare to make up for the simplicity of their storage scheme. T1 supernodes allow Level 3 BLAS updates, but as we will see in section 3.2 we can get most of their cache advantage with the more common T2 or T3 supernodes by using supernode-panel updates. Thus we conclude that either T2 or T3 is best by our criteria. Our code uses T2, which gives slightly larger supernodes than T3 at a small extra cost in storage (see section 2.2).

Figure 2.2 shows a sample matrix and the nonzero structure of its factors with no pivoting. Using definition T2, this matrix has four supernodes: $\{1,2\}$, $\{3\}$, $\{4,5,6\}$, and $\{7,8,9,10\}$. For example, in columns 4, 5, and 6 the diagonal blocks of L and U are full, and the columns of L all have nonzeros in rows 8 and 9. By definition T3, the matrix has five supernodes: $\{1,2\}$, $\{3\}$, $\{4,5,6\}$, $\{7\}$, and $\{8,9,10\}$. Column 7 fails to join $\{8,9,10\}$ as a T3 supernode because u_{78} is zero.

2.2. Storage of supernodes. A standard way to organize storage for a sparse matrix is a one-dimensional (1-D) array of nonzero values in column-major order, plus integer arrays giving row numbers and column starting positions. This is called *compressed column storage* and is also the scheme used in the Harwell–Boeing collection. We use this layout for both L and U, but with a slight modification: we store the entire square diagonal block of each supernode as part of L, including both the strict lower triangle of values from L and the upper triangle of values from U. We store this square block as if it were completely full (it is full in T3 supernodes, but its upper triangle may contain zeros in T2 supernodes). This allows us to address each supernode as a 2-D array in calls to BLAS routines. In other words, if columns $(r{:}s)$ form a supernode, then all the nonzeros in $F(r{:}n, r{:}s)$ are stored as a single dense 2-D array. This also lets us save some storage for row indices: only the indices of nonzero rows outside the diagonal block need be stored, and the structures of all columns within a supernode can be described by one set of row indices. This is similar to the effect of compressed subscripts in the symmetric case [41].

We represent the part of U outside the supernodal blocks with compressed column

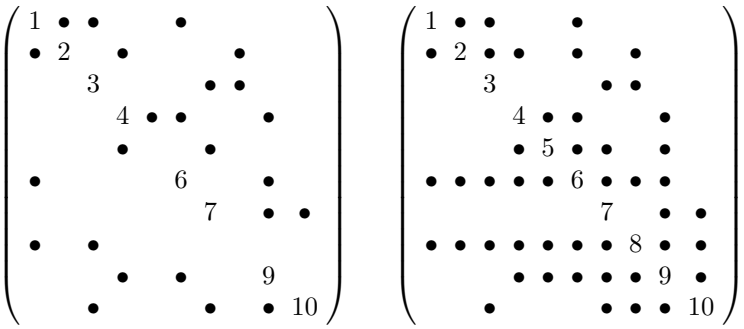

FIG. 2.2. *A sample matrix and its LU factors. Diagonal elements a_{55} and a_{88} are zero.*

$$
\begin{array}{c}
1 \\ 2 \\ 3 \\ 4 \\ 5 \\ 6 \\ 7 \\ 8 \\ 9 \\ 10
\end{array}
\begin{pmatrix}
s_1 & s_1 & u_3 & & & u_6 & & & & \\
s_1 & s_1 & u_3 & u_4 & & u_6 & & u_8 & & \\
 & & s_2 & & & & u_7 & u_8 & & \\
 & & & s_3 & s_3 & s_3 & & & u_9 & \\
 & & & s_3 & s_3 & s_3 & u_7 & & u_9 & \\
s_1 & s_1 & s_2 & s_3 & s_3 & s_3 & u_7 & u_8 & u_9 & \\
 & & & & & & s_4 & & s_4 & s_4 \\
s_1 & s_1 & s_2 & s_3 & s_3 & s_3 & s_4 & s_4 & s_4 & s_4 \\
 & & & s_3 & s_3 & s_3 & s_4 & s_4 & s_4 & s_4 \\
 & & s_2 & & & & s_4 & s_4 & s_4 & s_4
\end{pmatrix}
$$

FIG. 2.3. *Supernodal structure (by definition* T2*) of the factors of the sample matrix.*

storage: the values are stored by columns, with a companion integer array the same size to store row indices; another array of n integers indicates the start of each column.

Figure 2.3 shows the structure of the factors in the example from Figure 2.2, with s_k denoting a nonzero in the kth supernode and u_k denoting a nonzero in the kth column of U outside the supernodal block. Figure 2.4 shows the storage layout. (We omit the indexing vectors that point to the beginning of each supernode and the beginning of each column of U.)

2.3. The column elimination tree. Since our definition requires the columns of a supernode to be contiguous, we should get larger supernodes if we bring together columns of L with the same nonzero structure. But the column ordering is fixed, for sparsity, before numeric factorization; what can we do?

In symmetric Cholesky factorization, one type of supernodes—the "fundamental" supernodes—can be made contiguous by permuting the matrix (symmetrically) according to a postorder on its elimination tree [4]. This postorder is an example of what Liu calls an equivalent reordering [35], which does not change the sparsity of the factor. The postordered elimination tree can also be used to locate the supernodes before the numeric factorization.

We proceed similarly for the unsymmetric case. Here the appropriate analogue of the symmetric elimination tree is the *column elimination tree*, or column etree

$$\begin{array}{c} \\ 6 \\ 8 \end{array}\!\!\begin{array}{c} \begin{array}{cc} 1 & 2 \end{array} \\ \begin{pmatrix} s_1 & s_1 \\ s_1 & s_1 \\ s_1 & s_1 \\ s_1 & s_1 \end{pmatrix} \end{array} \qquad \begin{array}{c} \\ 6 \\ 8 \\ 10 \end{array}\!\!\begin{array}{c} \begin{array}{c} 3 \end{array} \\ \begin{pmatrix} s_2 \\ s_2 \\ s_2 \\ s_2 \end{pmatrix} \end{array} \qquad \begin{array}{c} \\ 8 \\ 9 \end{array}\!\!\begin{array}{c} \begin{array}{ccc} 4 & 5 & 6 \end{array} \\ \begin{pmatrix} s_3 & s_3 & s_3 \\ s_3 & s_3 & s_3 \\ s_3 & s_3 & s_3 \\ s_3 & s_3 & s_3 \\ s_3 & s_3 & s_3 \end{pmatrix} \end{array} \qquad \begin{array}{c} \begin{array}{cccc} 7 & 8 & 9 & 10 \end{array} \\ \begin{pmatrix} s_4 & 0 & s_4 & s_4 \\ s_4 & s_4 & s_4 & s_4 \\ s_4 & s_4 & s_4 & s_4 \\ s_4 & s_4 & s_4 & s_4 \end{pmatrix} \end{array}$$

Supernodal blocks (stored in column-major order)

row	1 2	2	1 2	3 5 6	2 3 6	4 5 6
	u_3 u_3	u_4	u_6 u_6	u_7 u_7 u_7	u_8 u_8 u_8	u_9 u_9 u_9

Nonzeros in columns of U outside supernodes

FIG. 2.4. *Storage layout for factors of the sample matrix, using* T2 *supernodes.*

for short. The vertices of this tree are the integers 1 through n, representing the columns of A. The column etree of A is the (symmetric) elimination tree of the column intersection graph of A, or equivalently the elimination tree of A^TA provided there is no cancellation in computing A^TA. See Gilbert and Ng [27] for complete definitions. The column etree can be computed from A in time almost linear in the number of nonzeros of A by a variation of an algorithm of Liu [35].

The following theorem states that the column etree represents potential dependencies among columns in LU factorization, and that (for strong Hall matrices) no stronger information is obtainable from the nonzero structure of A. Note that column i updates column j in LU factorization if and only if $u_{ij} \neq 0$.

THEOREM 2.1 (column etree [27]). *Let A be a square, nonsingular, possibly unsymmetric matrix, and let $PA = LU$ be any factorization of A with pivoting by row interchanges. Let T be the column etree of A.*

1. *If vertex i is an ancestor of vertex j in T, then $i \geq j$.*
2. *If $l_{ij} \neq 0$, then vertex i is an ancestor of vertex j in T.*
3. *If $u_{ij} \neq 0$, then vertex j is an ancestor of vertex i in T.*
4. *Suppose in addition that A is strong Hall (that is, A cannot be permuted to a nontrivial block triangular form). If vertex j is the parent of vertex i in T, then there is some choice of values for the nonzeros of A that makes $u_{ij} \neq 0$ when the factorization $PA = LU$ is computed with partial pivoting.*

Just as a postorder on the symmetric elimination tree brings together symmetric supernodes, we expect a postorder on the column etree to bring together unsymmetric supernodes. Thus, before we factor the matrix, we compute its column etree and permute the matrix columns according to a postorder on the tree. We now show that this does not change the factorization in any essential way.

THEOREM 2.2. *Let A be a matrix with column etree T. Let π be a permutation such that whenever $\pi(i)$ is an ancestor of $\pi(j)$ in T, we have $i \geq j$. Let P be the permutation matrix such that $\pi = P \cdot (1\colon n)^T$. Let $\bar{A} = PAP^T$.*

1. *$\bar{A} = A(\pi, \pi)$.*
2. *The column etree $\bar{T}$ of $\bar{A}$ is isomorphic to T; in particular, relabeling each node i of $\bar{T}$ as $\pi(i)$ yields T.*

3. *Suppose in addition that $\bar{A}$ has an LU factorization without pivoting, $\bar{A} = \bar{L}\bar{U}$. Then $P^T\bar{L}P$ and $P^T\bar{U}P$ are, respectively, unit lower triangular and upper triangular, so $A = (P^T\bar{L}P)(P^T\bar{U}P)$ is also an LU factorization.*

Remark 2.3. Liu [35] attributes to F. Peters a result similar to part 3 for the symmetric positive definite case, concerning the Cholesky factor and the (usual, symmetric) elimination tree.

Proof. Part 1 is immediate from the definition of P. Part 2 follows from Corollary 6.2 in Liu [35], with the symmetric structure of the column intersection graph of our matrix A taking the place of Liu's symmetric matrix A. (Liu exhibits the isomorphism explicitly in the proof of his Theorem 6.1.)

Now we prove part 3. We have $a_{\pi(i)\pi(j)} = \bar{a}_{ij}$ for all i and j. Write $L = P^T\bar{L}P$ and $U = P^T\bar{U}P$, so that $l_{\pi(i)\pi(j)} = \bar{l}_{ij}$ and $u_{\pi(i)\pi(j)} = \bar{u}_{ij}$. Then $A = LU$; we need show only that L and U are triangular.

Consider a nonzero $u_{\pi(i)\pi(j)}$ of U. In the triangular factorization $\bar{A} = \bar{L}\bar{U}$, element $\bar{u}_{ij}$ is equal to $u_{\pi(i)\pi(j)}$ and is therefore nonzero. By part 3 of Theorem 2.1, then, j is an ancestor of i in $\bar{T}$. By the isomorphism between $\bar{T}$ and T, this implies that $\pi(j)$ is an ancestor of $\pi(i)$ in T. Then it follows from part 1 of Theorem 2.1 that $\pi(j) \geq \pi(i)$. Thus every nonzero of U is on or above the diagonal, so U is upper triangular. A similar argument shows that every nonzero of L is on or below the diagonal, so L is lower triangular. The diagonal elements of L are a permutation of those of $\bar{L}$, so they are all equal to 1. □

Since the triangular factors of A are just permutations of the triangular factors of PAP^T, they have the same sparsity. Indeed, they require the same arithmetic to compute; the only possible difference is the order of updates. If addition for updates is commutative and associative, this implies that with partial pivoting (i, j) is a legal pivot in $\bar{A}$ if and only if $(\pi(i), \pi(j))$ is a legal pivot in A. In floating-point arithmetic, the different order of updates could conceivably change the pivot sequence. Thus we have the following corollary.

COROLLARY 2.4. *Let π be a postorder on the column etree of A, let P_1 be any permutation matrix, and let P_2 be the permutation matrix with $\pi = P_2 \cdot (1{:}n)^T$. If $P_1AP_2^T = LU$ is an LU factorization, then so is $(P_2^TP_1)A = (P_2^TLP_2)(P_2^TUP_2)$. In exact arithmetic, the former is an LU factorization with partial pivoting of AP_2^T if and only if the latter is an LU factorization with partial pivoting of A.*

This corollary states that an LU code can permute the columns of its input matrix by postorder on the column etree, and then fold the column permutation into the row permutation on output. Thus our SuperLU code has the option of returning either four matrices P_1, P_2, L, and U (with $P_1AP_2^T = LU$), or just the three matrices $P_2^TP_1$, $P_2^TLP_2$, and $P_2^TUP_2$, which are a row permutation and two triangular matrices. The advantage of returning all four matrices is that the columns of each supernode are contiguous in L, which permits the use of a Level 2 BLAS supernodal triangular solve for the forward-substitution phase of a linear system solver. The supernodes are not contiguous in $P_2^TLP_2$.

2.4. Relaxed supernodes. We observe that, for most matrices, the average size of a supernode is only about 2 to 3 columns (though a few supernodes are much larger). A large percentage of supernodes consists of only a single column, many of which are leaves of the column etree. Therefore we have devised a scheme to merge groups of columns at the fringe of the etree into *artificial supernodes* regardless of their row structures. A parameter r controls the granularity of the merge. Our merge rule is: node i is merged with its parent node j when the subtree rooted at j has at

```
1.  for column j = 1 to n do
2.      f = A(:, j);
3.      Symbolic factorization: determine which supernodes of L will update f;
4.      Determine whether j belongs to the same supernode as j − 1;
5.      for each updating supernode (r:s) < j in topological order do
6.          Apply supernode-column update to column j:
7.              f(r:s) = L(r:s, r:s)^{-1} · f(r:s);  /* Now f(r:s) = U(r:s, j) */
8.              f(s + 1:n) = f(s + 1:n) − L(s + 1:n, r:s) · f(r:s);
9.      end for;
10.     Pivot: interchange f(j) and f(m), where |f(m)| = max |f(j:n)|;
11.     Separate L and U: U(1:j, j) = f(1:j);    L(j:n, j) = f(j:n);
12.     Scale: L(j:n, j) = L(j:n, j)/L(j, j);
13.     Prune symbolic structure based on column j;
14. end for;
```

FIG. 3.1. *LU factorization with supernode-column updates.*

most r nodes. In practice, the best values of r are generally between 4 and 8 and yield improvements in running time of 5% to 15%.

Artificial supernodes are a special case of *relaxed supernodes*, which Duff and Reid [15] and Ashcraft and Grimes [4] have used in the context of multifrontal methods for systems with symmetric nonzero structure. They allow a small number of zeros in the structure of any supernode, thus relaxing the condition that the columns must have strictly nested structures. It would be possible to use this idea in the unsymmetric code as well, though we have not experimented with it. Relaxed supernodes could be constructed either on the fly (by relaxing the nonzero count condition described in section 4.3 below), or by preprocessing the column etree to identify small subtrees that we would merge into supernodes.

3. Supernodal numeric factorization. Now we show how to modify the col-col algorithm to use sup-col updates and supernode-panel updates. This section describes the numerical computation involved in the updates. Section 4 describes the symbolic factorization that determines which supernodes update which columns and also the boundaries between supernodes.

3.1. Sup-col updates. Figure 3.1 sketches the sup-col algorithm. The only difference from the col-col algorithm is that all the updates to a column from a single supernode are done together. Consider a supernode $(r{:}s)$ that updates column j. The coefficients of the updates are the values from a segment of column j of U, namely $U(r{:}s, j)$. The nonzero structure of such a segment is particularly simple: all the nonzeros are contiguous, and follow all the zeros (as proved in Corollary 4.2, which appears in section 4.1). Thus, if k is the index of the first nonzero row in $U(r{:}s, j)$, the updates to column j from supernode $(r{:}s)$ come from columns k through s. Since the supernode is stored as a dense matrix, these updates can be performed by a dense lower triangular solve (with the matrix $L(k{:}s, k{:}s)$) and a dense matrix-vector multiplication (with the matrix $L(s + 1{:}n, k{:}s)$). As described in section 4, the symbolic phase determines the value of k, that is, the position of the first nonzero in the segment $U(r{:}s, j)$.

The advantages of using sup-col updates are similar to those in the symmetric case [37]. Efficient Level 2 BLAS matrix-vector kernels can be used for the triangular solve and matrix-vector multiply. Furthermore, all the updates from the supernodal

1. **for** column $j = 1$ **to** n **step** w **do**
2. Symbolic factor: determine which supernodes will update any of $L(:, j{:}j+w-1)$;
3. **for** each updating supernode $(r{:}s) < j$ in topological order **do**
4. **for** column $jj = j$ **to** $j+w-1$ **do**
5. Apply supernode-column update to column jj;
6. **end for**;
7. **end for**;
8. Inner factorization:
   Apply the sup-col algorithm on columns and supernodes within the panel;
9. **end for**;

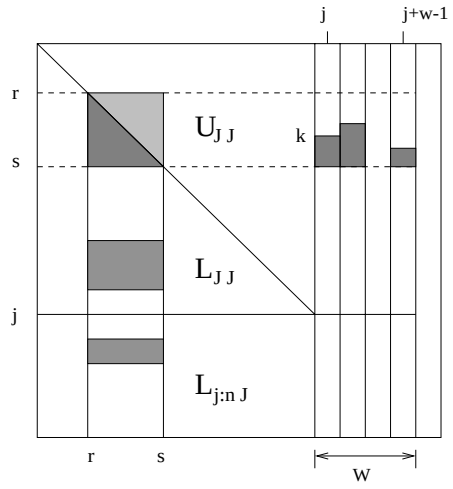

FIG. 3.2. *The supernode-panel algorithm, with columnwise blocking.* $J = 1{:}j-1$.

columns can be collected in a dense vector before doing a single scatter into the target vector. This reduces the amount of indirect addressing.

3.2. Supernode-panel updates. We can improve the sup-col algorithm further on machines with a memory hierarchy by changing the data access pattern. The data we are accessing in the inner loop (lines 5–9 of Figure 3.1) include the destination column j and all the updating supernodes $(r{:}s)$ to the left of column j. Column j is accessed many times, while each supernode $(r{:}s)$ is used only once. In practice, the number of nonzero elements in column j is much less than that in the updating supernodes. Therefore, the access pattern given by this loop provides little opportunity to reuse cached data. In particular, the same supernode $(r{:}s)$ may be needed to update both columns j and $j+1$. But when we factor the $(j+1)$st column (in the next iteration of the outer loop), we will have to fetch supernode $(r{:}s)$ again from memory, instead of from cache (unless the supernodes are small compared to the cache).

3.2.1. Panels. To exploit memory locality, we factor several columns (say w of them) at a time in the outer loop, so that one updating supernode $(r{:}s)$ can be used to update as many of the w columns as possible. We refer to these w consecutive columns as a *panel* to differentiate them from a supernode; the row structures of these columns may not be correlated in any fashion, and the boundaries between panels may be different from those between supernodes. The new method requires rewriting

the doubly nested loop as the triple loop shown in Figure 3.2. This is analogous to loop tiling techniques used in optimizing compilers to improve cache behavior for 2-D arrays with regular stride. It is also somewhat analogous to the sup-sup updates that Ng and Peyton [37] and Rothberg and Gupta [38] have used in symmetric Cholesky factorization.

The structure of each sup-col update is the same as in the sup-col algorithm. For each supernode $(r{:}s)$ to the left of column j, if $u_{kj} \neq 0$ for some $r \leq k \leq s$, then $u_{ij} \neq 0$ for all $k \leq i \leq s$. Therefore, the nonzero structure of the panel of U consists of dense column segments that are rowwise separated by supernodal boundaries, as in Figure 3.2. Thus, it is sufficient for the symbolic factorization algorithm to record only the first nonzero position of each column segment. As detailed in section 4.4, symbolic factorization is applied to all the columns in a panel at once, over all the updating supernodes, before the numeric factorization step.

In dense factorization, the entire supernode-panel update in lines 3–7 of Figure 3.2 would be implemented as two Level 3 BLAS calls: a dense triangular solve with w right-hand sides, followed by a dense matrix-matrix multiply. In the sparse case, this is not possible, because the different sup-col updates begin at different positions k within the supernode, and the submatrix $U(r{:}s, j{:}j + w - 1)$ is not dense. Thus the sparse supernode-panel algorithm still calls the Level 2 BLAS. However, we get similar cache benefits to those from the Level 3 BLAS, at the cost of doing the loop reorganization ourselves. Thus we sometimes call the kernel of this algorithm a "BLAS-$2\frac{1}{2}$" method.

In the doubly nested loop (lines 3–7 of Figure 3.2), the ideal circumstance is that all w columns in the panel require updates from supernode $(r{:}s)$. Then this supernode will be used w times before it is forced out of the cache. There is a trade-off between the value of w and the size of the cache. For this scheme to work efficiently, we need to ensure that the nonzeros in the w columns do not cause cache thrashing. That is, we must keep w small enough so that all the data accessed in this doubly nested loop fit in cache. Otherwise, the cache interference between the source supernode and the destination panel can offset the benefit of the new algorithm.

3.2.2. Outer and inner factorization. At the end of the supernode-panel update (line 7), columns j through $j + w - 1$ of L and U have received all their updates from columns to the left of j. We call this the *outer factorization*. What remains is to apply updates that come from columns within the panel. This amounts to forming the LU factorization of the panel itself (in columns $(j{:}j + w - 1)$ and rows $(j{:}n)$). This *inner factorization* is performed by the sup-col algorithm, almost exactly as shown in Figure 3.1. The inner factorization includes a columnwise symbolic factorization just as in the sup-col algorithm. The inner factorization also includes the supernode identification, partial pivoting, and symmetric structure reduction for the entire algorithm. Section 4 contains details of the inner factorization.

3.2.3. Reducing cache misses by rowwise blocking. Our first experiments with the supernode-panel algorithm showed speedups for some medium-sized problems of around 20–30%. However, the improvement for large matrices was often only a few percentage points. We now study the reasons and remedies for this.

To implement loops (lines 3–7 of Figure 3.2), we first expand the nonzeros of the panel columns of A into an n by w full working array, called the *sparse accumulator* [26] or SPA. This allows random access to the entries of the active panel. A temporary array stores the results of the BLAS operations, and the updates are scattered into the SPA. At the end of panel factorization, the data in the SPA are copied into storage for L and U. Although increasing the panel size w gives more opportunity for data

```
1.  for j = 1 to n step w do
2.      ...
3.      for each updating supernode (r:s) < j in topological order do
4.          Apply triangular solves to A(r:s, j:j + w − 1) using L(r:s, r:s);
5.          for each row block B in L(s + 1:n, r:s) do
6.              for jj = j to j + w − 1 do
7.                  Multiply B · U(r:s, jj), and scatter into SPA(:, jj);
8.              end for;
9.          end for;
10.     end for;
11.     ...
12  end for;
```

FIG. 3.3. *The supernode-panel algorithm, with 2-D blocking.*

reuse, it also increases the size of the active data set that must fit into cache. The supernode-panel update loop accesses the following data:

- the nonzeros in the updating supernode $L(r{:}n, r{:}s)$;
- the SPA data structure, consisting of an n by w full array and a temporary store of size n.

By instrumenting the code, we found that the working sets of large matrices are much larger than the cache size. Hence, cache thrashing limits performance.

We experimented with a scheme suggested by Rothberg [39], in which the SPA has only as many rows as the number of nonzero rows in the panel (as predicted by symbolic factorization), and an extra indirection array of size n is used to address the SPA. Unfortunately, the cost incurred by double indirection is not negligible, and this scheme was not as effective as the 2-D blocking method we now describe.

We implemented a rowwise blocking scheme on top of the columnwise blocking in the supernode-panel update. The 2-D blocking adds another level of looping between the two loops in lines 3 and 4 of Figure 3.2. This partitions the supernodes (and the SPA structure) into block rows. Then each block row of the updating supernode is used for up to w partial matrix-vector multiplies, which are pushed all the way through into the SPA before the next block row of the supernode is accessed. The active data set accessed in the inner loops is thus much smaller than in the 1-D scheme. The 2-D blocking algorithm is organized as in Figure 3.3. The key performance gains come from the loops (lines 5–9), where each row block is reused as much as possible before the next row block is brought into the cache. The innermost loop is still a dense matrix-vector multiply, performed by a Level 2 BLAS kernel.

3.2.4. Combining 1-D and 2-D blocking. The 2-D blocking works well when the rectangular supernodal matrix $L(r{:}n, r{:}s)$ is large in both dimensions. If all of $L(r{:}n, r{:}s)$ fits in cache, then the rowwise blocking gives no benefit, but still incurs overhead for setting up loop variables, skipping the empty loop body, and so on. This overhead can be nearly 10% for some of the sparser problems. Thus we have devised a hybrid update algorithm that uses either the 1-D or 2-D partitioning scheme, depending on the size of each updating supernode. This decision is made at runtime, with the overhead limited to a one-line test. It turns out that this hybrid scheme works better than either 1-D or 2-D code for many problems. Therefore, this is the algorithm that we used in our code.

4. Symbolic factorization. Symbolic factorization is the process that determines the nonzero structure of the triangular factors L and U from the nonzero structure of the matrix A. This in turn determines which columns of L update each column j of the factors (namely, those columns r for which $u_{rj} \neq 0$), and also determines which columns of L can be combined into supernodes.

Without numeric pivoting, the complete symbolic factorization can be performed before any numeric factorization. Partial pivoting, however, requires that the numeric and symbolic factorizations be interleaved. The sup-col algorithm performs symbolic factorization for each column just before it is computed, as described in section 4.1. The supernode-panel algorithm performs symbolic factorization for an entire panel at once, as described in section 4.4.

4.1. Column depth-first search. From the numeric factorization algorithm, it is clear that the structure of column $F(:, j)$ depends on the structure of column $A(:, j)$ of the original matrix and on the structure of $L(:, J)$, where $J = 1{:}j-1$. Indeed, $F(:, j)$ has the same structure as the solution vector for the following triangular system [29]:

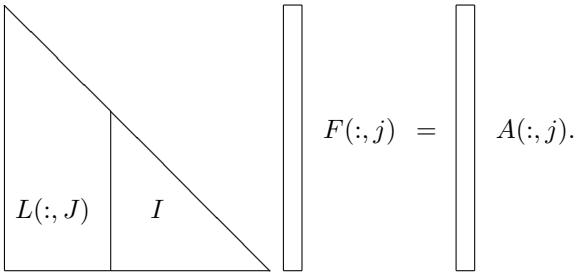

A straightforward way to compute the structure of $F(:, j)$ from the structures of $L(:, J)$ and $A(:, j)$ is to simulate the numerical computation. A less expensive way is to use the following characterization in terms of paths in the directed graph of $L(:, J)$.

For any matrix M, the notation $i \xrightarrow{M} j$ means that there is an edge from i to j in the directed graph of M, that is, $m_{ij} \neq 0$. Edges in the directed graph of M are directed from rows to columns. The notation $i \overset{M}{\Longrightarrow} j$ means that there is a directed path from i to j in the directed graph of M. Such a path may have length zero; that is, $i \overset{M}{\Longrightarrow} i$ always holds.

THEOREM 4.1 (see [24]). *f_{ij} is nonzero (equivalently, $i \xrightarrow{F} j$) if and only if $i \overset{L(:,J)}{\Longrightarrow} k \xrightarrow{A} j$ for some $k \leq i$.*

This result implies that the symbolic factorization of column j can be obtained as follows. Consider the nonzeros in $A(:, j)$ as a subset of the vertices of the directed graph $G = G(L(:, J)^T)$, which is the reverse of the directed graph of $L(:, J)$. The nonzero positions of $F(:, j)$ are then given by the vertices reachable by paths from this set in G. We use the graph of L^T here because of the convention that edges are directed from rows to columns. Since L is actually stored by columns, our data structure gives precisely the adjacency information for G. Therefore, we can determine the structure of column j of L and U by traversing G from the set of starting nodes given by the structure of $A(:, j)$.

The traversal of G determines the structure of $U(:, j)$, which in turn determines the columns of L that will participate in updates to column j in the numerical fac-

torization. These updates must be applied in an order consistent with a topological ordering of G. We use depth-first search to perform the traversal, which makes it possible to generate a topological order (specifically, reverse postorder) on the nonzeros of $U(:, j)$ as they are located [29].

Another consequence of the path theorem is the following corollary. It states that if we divide each column of U into *segments*, one per supernode, then within each segment the column of U consists of a consecutive sequence of nonzeros. Thus we need only keep track of the position of the first nonzero in each segment.

COROLLARY 4.2. *Let $(r\!:\!s)$ be a supernode (of either type T2 or T3) in the factorization $PA = LU$. Suppose u_{kj} is nonzero for some j with $r \le k \le s$. Then $u_{ij} \neq 0$ for all i with $k \le i \le s$.*

Proof. Let $k \le i \le s$. Since $u_{kj} \neq 0$, we have $k \overset{L(:,J)}{\Longrightarrow} m \overset{A}{\longrightarrow} j$ for some $m \le k$ by Theorem 4.1. Now l_{ik} is in the diagonal block of the supernode, and hence is nonzero. Thus $i \overset{L(:,J)}{\longrightarrow} k$, so $i \overset{L(:,J)}{\Longrightarrow} m \overset{A}{\longrightarrow} j$, whence u_{ij} is nonzero by Theorem 4.1. ∎

4.2. Pruning the symbolic structure. We can speed up the depth-first search traversals by using a reduced graph in place of G, the reverse of the graph of $L(:, J)$. We have explored this idea in a series of papers [20, 21, 25]. Any graph H can be substituted for G, provided that $i \overset{H}{\Longrightarrow} j$ if and only if $i \overset{G}{\Longrightarrow} j$. The traversals are more efficient if H has fewer edges; but any gain in efficiency must be traded off against the cost of computing H.

An extreme choice of H is the elimination directed acyclic graph (*elimination dag*) [25], which is the transitive reduction of G, or the minimal subgraph of G that preserves paths. However, the elimination dag is expensive to compute. The *symmetric reduction* [20] is a subgraph that has (in general) fewer edges than G but more edges than the elimination dag, and that is much less expensive to compute. The symmetric reduction takes advantage of symmetry in the structure of the filled matrix F; if F is completely symmetric, it is just the symmetric elimination tree. The symmetric reduction of $L(:, J)$ is obtained by removing all nonzeros l_{rs} for which $l_{ts}u_{st} \neq 0$ for some $t < \min(r, j)$. Eisenstat and Liu [21] give an efficient method to compute the symmetric reduction during symbolic factorization and demonstrate experimentally that it significantly reduces the total factorization time when used in an algorithm that does col-col updates.

Our supernodal code uses symmetric reduction to speed up its symbolic factorization. Using the sample matrix in Figure 2.2, Figures 4.1 and 4.2 illustrate symmetric reduction in the presence of supernodes. We use S to represent the supernodal structure of $L(:, J)^T$ and R to represent the symmetric reduction of S. It is this R that we use in our code. Note that the edges of the graph of R are directed from columns of L to rows of L.

In the figures, the symbol "$\oslash$" indicates an entry in S that was pruned from R by symmetric reduction. The $(8, 2)$ entry was pruned due to the symmetric nonzero pair $(6, 2)$ and $(2, 6)$. The figure shows the current state of the reduced structure based on the first seven columns of the filled matrix.

It is instructive to follow this example through one more column to see how symbolic factorization is carried out in the reduced supernodal structure. Consider the symbolic step for column 8. Suppose a_{28} and a_{38} are nonzero. The other nonzeros in column 8 of the factor are generated by paths in the reduced supernodal structure (we show just one possible path for each nonzero):

$$8 \overset{A^T}{\rightarrow} 2 \overset{R}{\rightarrow} 6,$$

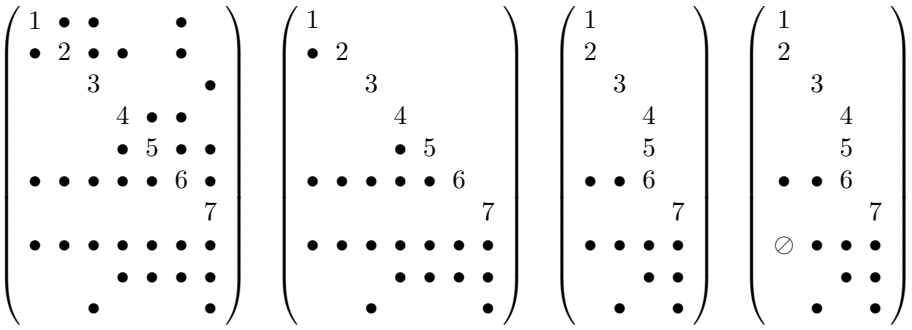

FIG. 4.1. *Supernodal and symmetrically reduced structures.*

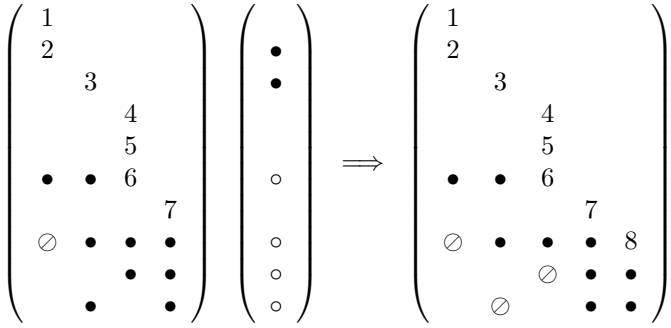

FIG. 4.2. *One step of symbolic factorization in the reduced structure.*

$$8 \xrightarrow{A^T} 3 \xrightarrow{R} 8,$$
$$8 \xrightarrow{A^T} 2 \xrightarrow{R} 6 \xrightarrow{R} 9,$$
$$8 \xrightarrow{A^T} 3 \xrightarrow{R} 10.$$

Figure 4.2 shows the reduced supernodal structure before and after column 8. In column 8 of A, the original nonzeros are shown as "•" and the fill nonzeros are shown as "○". Once the structure of column 8 of U is known, more symmetric reduction is possible. The entry $l_{10,3}$ is pruned due to the symmetric nonzeros in l_{83} and u_{38}. Also, l_{96} is pruned by l_{86} and u_{68}. Figure 4.2 shows the new structure.

The supernodal symbolic factorization relies on the path characterization in Theorem 4.1 and on the path-preserving property of symmetric reduction. In effect, we use the modified path condition

$$i \xrightarrow{A^T} m \xRightarrow{R} j$$

on the symmetrically reduced supernodal structure R of $L(:, J)^T$.

4.3. Detecting supernodes. Since supernodes consist of contiguous columns of L, we can decide at the end of each symbolic factorization step whether the new column j belongs to the same supernode as column $j-1$.

For T2 supernodes, the test is straightforward. During symbolic factorization, we test whether $L(:,j) \subseteq L(:,j-1)$ (where the containment applies to the set of nonzero indices). At the end of the symbolic factorization step, we test whether $nnz(L(:,j)) = nnz(L(:,j-1)) - 1$. Column j joins column $(j-1)$'s supernode if and only if both tests are passed.

T3 supernodes also require the diagonal block of U to be full. To check this, it suffices to check whether the single element u_{rj} is nonzero, where r is the first column index of the supernode. This follows from Corollary 4.2: $u_{rj} \neq 0$ implies that $u_{ij} \neq 0$ for all $r \leq i \leq j$. Indeed, we can even omit the test $L(:,j) \subseteq L(:,j-1)$ for T3 supernodes. If $u_{rj} \neq 0$, then $u_{j-1,j} \neq 0$, which means that column $j-1$ updates column j, which implies $L(:,j) \subseteq L(:,j-1)$. Thus we have proved the following theorem.

THEOREM 4.3. *Suppose a T3 supernode begins with column r and extends at least through column $j-1$. Column j belongs to this supernode if and only if $u_{rj} \neq 0$ and $nnz(L(:,j)) = nnz(L(:,j-1)) - 1$.*

For either T2 or T3 supernodes, it is straightforward to implement the relaxed versions discussed in section 2.4. Also, since the main benefits of supernodes come when they fit in the cache, we impose a maximum size for a supernode.

4.4. Panel depth-first search. The supernode-panel algorithm consists of an outer factorization (applying updates from supernodes to the active panel) and an inner factorization (applying sup-col updates within the active panel). Each has its own symbolic factorization. The outer symbolic factorization happens once per panel and determines two things: (1) a single column structure, which is the union of the structures of the panel columns, and (2) which supernodes update each column of the panel, and in what order. This is the information that the supernode-panel update loop in Figure 3.2 needs.

The inner symbolic factorization happens once for each column of the panel, interleaved column by column with the inner numeric factorization. In addition to determining the nonzero structure of the active column and which supernodes within the panel will update the active column, the inner symbolic factorization is also responsible for forming supernodes (that is, for deciding whether the active column will start a new supernode) and for symmetric structural pruning. The inner symbolic factorization is, therefore, exactly the sup-col symbolic factorization described above.

The outer symbolic factorization must determine the structures of columns j to $j+w-1$, i.e., the structure of the whole panel, and also a topological order for $U(1{:}j, j{:}j+w-1)$ en masse. To this end, we developed an efficient panel depth-first search scheme, which is a slight modification of the column depth-first search. The panel depth-first search algorithm maintains a single postorder depth-first search list for all w columns of the panel. Let us call this the *PO* list, which is initially empty. The algorithm invokes the column depth-search procedure for each column from j to $j+w-1$. In the column depth-first search, each time the search backtracks from a vertex, that vertex is appended to the *PO* list. In the panel depth-first search, however, the vertex is appended to the *PO* list *only if it is not already on the list.* This gives a single list that includes every position that is nonzero in any panel column, just once; and the entire list is in (reverse) topological order. Thus the order of updates specified by the list is acceptable for each of the w individual panel columns.

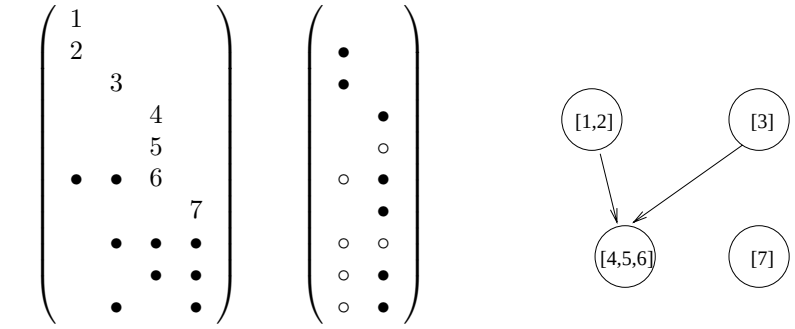

FIG. 4.3. *The supernodal directed graph corresponding to* $L(1{:}7, 1{:}7)^T$.

We illustrate the idea in Figure 4.3, using the sample matrix from Figures 4.1 and 4.2, and a panel of width two. The first seven columns have been factored, and the current panel consists of columns 8 and 9. In the panel, nonzeros of A are shown as "•" and fill in F is shown as "○". The depth-first search for column 8 starts from vertices 2 and 3. After that search is finished, the panel postorder list is $PO = (6, 2, 3)$. Now the depth-first search for column 9 starts from vertices 6 and 7 (not 4, since 6 is the representative vertex for the supernode containing column 4). This depth-first search only appends 7 to the PO list, because 6 is already on the list. Thus, the final list for this panel is $PO = (6, 2, 3, 7)$. The postorder list of column 8 is $(6, 2, 3)$ and the postorder list of column 9 is $(6, 7)$, which are simply two sublists of the panel PO list. The topological order is the reverse of PO, or $(7, 3, 2, 6)$. In the loop of line 3 of Figure 3.2, we follow this topological order to schedule the updating supernodes and perform numeric updates to columns of the panel.

5. Evaluation. In this section, we evaluate our algorithms using matrices from several applications and several sources. We compare the performance of SuperLU, our supernode-panel code, with its predecessors and with one other approach to sparse LU factorization.

5.1. Experimental setup. Table 5.1 lists 23 matrices with some characteristics of their nonzero structures. Some of the matrices are from the Harwell–Boeing collection [17]. Many of the larger matrices are from the ftp site maintained by Tim Davis of the University of Florida.[3] Those matrices are as follows: MEMPLUS is a circuit simulation matrix from Steve Hamm of Motorola. RDIST1 is a reactive distillation problem in chemical process separation calculations, provided by Stephen Zitney at Cray Research, Inc. SHYY161 is derived from a direct, fully coupled method for solving the Navier–Stokes equations for viscous flow calculations, provided by Wei Shyy of the University of Florida. GOODWIN is a finite element matrix in a nonlinear solver for a fluid mechanics problem, provided by Ralph Goodwin of the University of Illinois at Urbana–Champaign. VENKAT01, INACCURA, and RAEFSKY3/4 were provided by Horst Simon, then of NASA. VENKAT01 comes from an implicit 2-D Euler solver for an unstructured grid in a flow simulation. RAEFSKY3 is from a fluid structure interaction turbulence problem. RAEFSKY4 is from a buckling problem for a container model.

[3]URL: http://www.cis.ufl.edu/∼davis

TABLE 5.1
Characteristics of the test matrices. Structural symmetry s is the fraction of the nonzeros matched by nonzeros in symmetric locations. None of the matrices are numerically symmetric.

	Matrix	s	n	$nnz(A)$	$nnz(A)/n$
1	MEMPLUS	.983	17758	99147	5.6
2	GEMAT11	.002	4929	33185	6.7
3	RDIST1	.062	4134	9408	2.3
4	ORANI678	.073	2529	90158	35.6
5	MCFE	.709	765	24382	31.8
6	LNSP3937	.869	3937	25407	6.5
7	LNS3937	.869	3937	25407	6.5
8	SHERMAN5	.780	3312	20793	6.3
9	JPWH991	.947	991	6027	6.1
10	SHERMAN3	1.000	5005	20033	4.0
11	ORSREG1	1.000	2205	14133	6.4
12	SAYLR4	1.000	3564	22316	6.3
13	SHYY161	.769	76480	329762	4.3
14	GOODWIN	.642	7320	324772	44.4
15	VENKAT01	1.000	62424	1717792	27.5
16	INACCURA	1.000	16146	1015156	62.9
17	AF23560	.947	23560	460598	19.6
18	DENSE1000	1.000	1000	1000000	1000
19	RAEFSKY3	1.000	21200	1488768	70.2
20	EX11	1.000	16614	1096948	66.0
21	WANG3	1.000	26064	177168	6.8
22	RAEFSKY4	1.000	19779	1316789	66.6
23	VAVASIS3	.001	41092	1683902	41.0

TABLE 5.2
Machines used to compare various column LU codes. Column "#" is maximum number of instruction issues per clock cycle.

	Clock MHz	On-chip cache	External cache	#	Peak Mflops	DGEMM Mflops	DGEMV Mflops
IBM RS/6000-590	66.5	256 KB		6	266	250	235
SGI MIPS R8000	90	16 KB	4 MB	4	360	340	210
DEC Alpha 21164	300	8 KB-L1 96 KB-L2	4 MB	4	600	350	135
SUN UltraSparc-I	143	16 KB	512 KB	4	286	227	–

AF23560 is from solving an unsymmetric eigenvalue problem, provided by Zhaojun Bai of the University of Kentucky. EX11 is from a three-dimensional (3-D) steady flow calculation in the SPARSKIT collection maintained by Youcef Saad at the University of Minnesota. WANG3 is from solving a coupled nonlinear PDE system in a 3-D ($30\times30\times30$ uniform mesh) semiconductor device simulation, as provided by Song Wang of the University of New South Wales, Sydney. VAVASIS3 is an unsymmetric augmented matrix for a 2-D PDE with highly varying coefficients [44]. DENSE1000 is a dense 1000×1000 random matrix.

The matrices are sorted in increasing order of $flops/nnz(F)$, the ratio of the number of floating-point operations to the number of nonzeros $nnz(F)$ in the factored matrix $F = U + L - I$. The reason for this order will be described in more detail in section 5.4.

This paper does not address the performance of column preordering for sparsity. We simply use the existing ordering algorithms provided by Matlab [26]. For all

matrices, except 1, 15, and 21, the columns were permuted by Matlab's minimum degree ordering of A^TA, also known as "column minimum degree." However, this ordering produces excessive fill for matrices 1, 15, and 21, because it attempts only to minimize the upper bound on the actual fill, and the upper bounds are too loose in these cases. When these three matrices are symmetrically permuted by Matlab's symmetric minimum degree ordering on $A + A^T$, the amount of fill is much smaller than using column minimum degree ordering.

We conducted performance analysis on high-end workstations from four vendors (IBM, SGI, DEC, and SUN). Some characteristics of these machines are tabulated in Table 5.2. The instruction caches, if separate from the data cache, are not listed in the table. In most cases, the on-chip L1 caches are fairly small, so we use either the L2 cache or the off-chip cache as a reference. The DGEMM and DGEMV Mflop rates were measured using vendor-supplied BLAS libraries. (Exception: SUN does not supply a BLAS library, so we report the DGEMM speed from PHiPAC [6]. PHiPAC does not include DGEMV.) Our UltraSparc-I has less physical memory than the other machines, so some large problems could not be tested on this machine.

5.2. Performance of SuperLU on an IBM RS/6000-590. Table 5.3 presents the performance of SuperLU on this system. The CPU clock rate is 66.5 MHz. The processor has two branch units, two fixed-point units, and two floating-point units, which can all operate in parallel if there are no dependencies. Each FPU can perform two operations (a multiply and an add or subtract) in each cycle. Thus, the peak floating-point performance is 266 Mflops. The size of the main memory is 768 MB. SuperLU is implemented in C; we used the AIX xlc compiler with -O3 optimization. All floating-point computations are performed in double precision.

In the inner loops of our sparse code, we call the two Level 2 BLAS routines DTRSV (triangular solve) and DGEMV (matrix-vector multiply) provided in the IBM ESSL library [32], whose BLAS-3 matrix-matrix multiply routine (DGEMM) achieves about 250 Mflops when the dimension of the matrix is larger than 60 [1]. In our sparse algorithm, we find that DGEMV typically accounts for more than 80% of the floating-point operations. As shown in the second to last column of Table 5.3, this percentage is 95% higher than for many matrices. Our measurements reveal that for typical dimensions arising from the benchmark matrices, DGEMV achieves roughly 235 Mflops if the data are from cache. If the data are fetched from main memory, this rate can drop by a factor of 2 or 3.

The BLAS speed is clearly an upper bound on the overall factorization rate. However, because symbolic manipulation usually takes a nontrivial amount of time, this bound could be much higher than the actual sparse code performance. The last column in Table 5.3 presents the percentage of the total execution time spent in numeric computation. For matrices 1 and 2, the program spent less than 35% of its time in the numeric part. Compared to the others, these two matrices are sparser, have less fill, and have smaller supernodes, so our supernodal techniques are less applicable. Matrix 2 is also highly unsymmetric, which makes the symmetric structural reduction less effective. However, it is important to note that the execution times for these two matrices are small.

For larger and denser matrices such as 18–23, the algorithm achieves between 110 and 125 Mflops, which is about half of the machine peak. These matrices take much longer to factor, which could be a serious bottleneck in an iterative simulation process. Our techniques are successful in reducing the solution times for this type of problem.

TABLE 5.3
*Performance of SuperLU on an IBM RS/*6000-590.

	Matrix	$nnz(F)$	$\frac{nnz(F)}{nnz(A)}$	#flops (10^6)	Time (sec)	Mflops	% flops DGEMV	% num time
1	MEMPLUS	140388	1.4	1.8	0.57	3.08	70	16
2	GEMAT11	93370	2.8	1.5	0.27	5.64	82	33
3	RDIST1	338624	3.6	12.9	0.96	13.47	85	48
4	ORANI678	280788	3.1	14.9	1.11	13.48	98	51
5	MCFE	69053	2.8	4.1	0.24	17.42	96	54
6	LNSP3937	427600	16.8	38.9	1.50	25.97	95	48
7	LNS3937	449346	17.7	44.8	1.65	27.16	96	48
8	SHERMAN5	249199	12.0	25.2	0.82	30.78	93	57
9	JPWH991	140746	23.4	18.0	0.52	34.57	94	58
10	SHERMAN3	433376	21.6	60.6	1.37	44.24	85	56
11	ORSREG1	402478	28.5	59.8	1.21	49.42	87	50
12	SAYLR4	654908	29.3	104.8	2.18	48.07	87	57
13	SHYY161	7634810	23.2	1571.6	25.42	61.83	88	57
14	GOODWIN	3109585	9.6	665.1	12.55	52.99	92	63
15	VENKAT01	12987004	7.6	3219.9	42.99	74.90	91	63
16	INACCURA	9941478	9.8	4118.7	67.73	60.81	96	64
17	AF23560	13986992	30.4	6363.7	75.91	83.83	92	73
18	DENSE1000	1000000	1.0	666.2	5.68	117.28	93	72
19	RAEFSKY3	17544134	11.8	12118.7	107.60	112.62	94	77
20	EX11	26207974	23.8	26814.5	247.05	108.54	95	81
21	WANG3	13287108	74.9	14557.5	116.58	124.86	96	81
22	RAEFSKY4	26678597	20.3	31283.4	263.13	118.89	97	83
23	VAVASIS3	49192880	29.2	89209.3	786.94	113.36	98	80

For a dense 1000×1000 matrix, our code achieves 117 Mflops. This may be compared to 168 Mflops reported in the LAPACK manual [3] on a matrix of this size, and 236 Mflops reported in the online Linpack benchmark files [36].

5.3. Comparison with previous column LU algorithms. In this section, we compare the performance of SuperLU with three of its predecessors, including the column code GP by Gilbert and Peierls [29] (Figure 1.1), GP-Mod by Eisenstat and Liu [21] (section 4.2), and SupCol by Eisenstat, Gilbert, and Liu [19] (Figure 3.1). GP and GP-Mod were written in Fortran. SupCol was first written in Fortran, and later translated literally into C; no changes in algorithms or data structures were made in this translation. SuperLU is written in C. (Matlab contains C implementations of GP and GP-Mod [26], which we did not test here.)

For the Fortran codes, we use Fortran 77 compilers; for the C codes, we use ANSI C compilers. In all cases, we use highest possible optimization provided by each compiler. Both SupCol and SuperLU call Level 2 BLAS routines. For the RS/6000-590, we use the BLAS routines from IBM's ESSL library. For the DEC Alpha, we use the BLAS routines from DEC's DXML library. There are no vendor-supplied BLAS libraries on the Sparc, so we use our own routines implemented in C.

Tables 5.4 through 5.7 present the results of comparisons on the four machines. The blocking parameters w, t, and b (Figure 5.6) for SuperLU are chosen according to the size of the data cache (Table 5.2) and are reported in each comparison table. In all these tables, the column labeled "GP" gives the raw factorization time in seconds of GP. The numbers in each successive column are speedups achieved by the corresponding enhancement over GP. Thus, for example, a speedup of 2 means that the running time was half that of GP. The numbers in the last two rows of each table

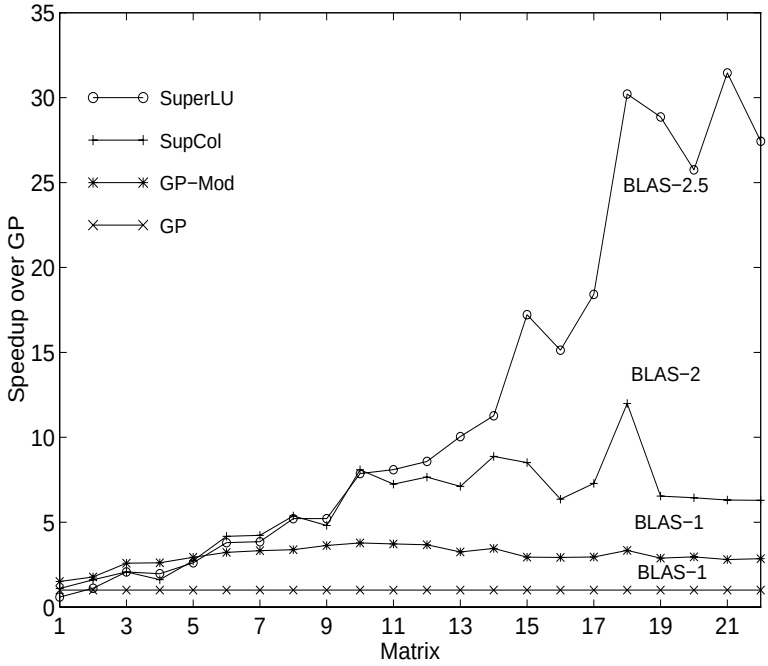

FIG. 5.1. *Speedups of each enhancement over GP on the MIPS R*8000.

show the average speedup and its standard deviation.

Figure 5.1 gives a visual look at the speedups for SuperLU over its predecessors on one of the workstations we experimented with (using data from Table 5.5). Of course, machines with different memory architectures would give different plots.

We make the following observations from these results:

- The symmetric structure pruning in GP-Mod is very effective in reducing the graph search time. This significantly decreases the symbolic factorization time in the GP code. It achieves speedup for all problems, on all machines. Its average speedup on the RS/6000-590 is 3.64, the highest among all the machines.
- Supernodes in SupCol restrict the search to the supernodal graph and allow the numeric kernels to employ dense BLAS-2 operations. The effects are not as dramatic as the pruning technique. For some matrices, such as 1–3, the runtimes are actually longer than GP-Mod. This is because supernodes are often small in the sparser matrices.
- Supernode-panel updates in SuperLU reduce the cache miss rate and exploit dense substructures in the factor F. For problems without much structure, the gain is often offset by various overheads. However, the advantage of SuperLU over SupCol becomes significant for larger or denser problems, or on machines with small cache, such as Alpha 21164, on which SuperLU achieves more than a factor of 2 speedup over SupCol for the six large matrices 18–23.

With more and more sophisticated techniques introduced, the added complications of the code increase the runtime overhead to some extent. This overhead can show up prominently in small or very sparse problems. The two supernodal codes are particularly sensitive to the characteristics of the problems. This can be seen from the

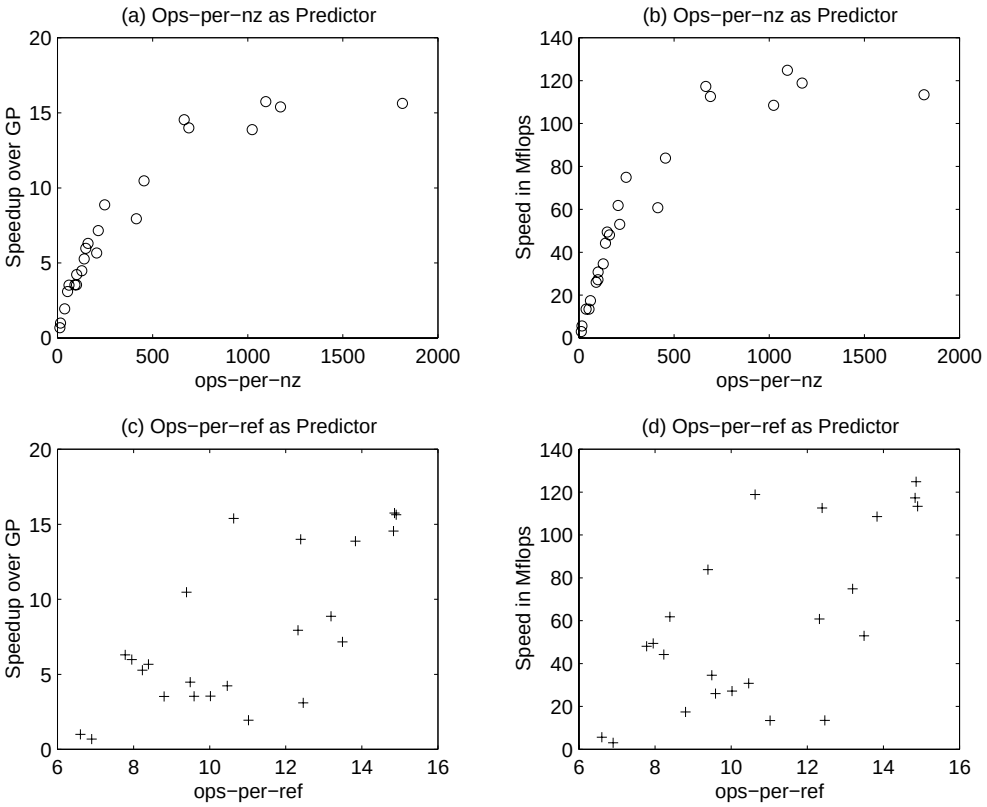

FIG. 5.2. *Matrix characteristics as predictors of performance.*

large standard deviations of their average speedups.

5.4. Understanding cache behavior and parameters. We now analyze the behavior of SuperLU in detail. We wish to understand when our algorithm is significantly faster than other algorithms. We would like performance-predicting variable(s) that depend on "intrinsic" properties of the problem, such as the sparsity structure, rather than algorithmic details and machine characteristics.

5.4.1. How much cache reuse can we expect? As discussed in section 3.2, the supernode-panel algorithm gets its primary gains from improved data locality by reusing a cached supernode several times. To understand how much cache reuse we can
hope for, we computed two statistics: *ops-per-nz* and *ops-per-ref*. After defining these statistics carefully, we discuss which more successfully measures reuse.

Ops-per-nz is simply calculated as $\#flops/nnz(F)$, the total number of floating-point operations per nonzero in the filled matrix F. If there were perfect cache behavior, i.e., one cache miss per data item (ignoring the effect of cache line size), then *ops-per-nz* would exactly measure the amount of work per cache miss. In reality, *ops-per-nz* is an upper bound on the reuse. Note that *ops-per-nz* depends only on the fact that we are performing Gaussian elimination with partial pivoting, not on implementation or machine details. *Ops-per-nz* is a natural measure of potential data reuse, because it has long been used to distinguish among the different levels of BLAS.

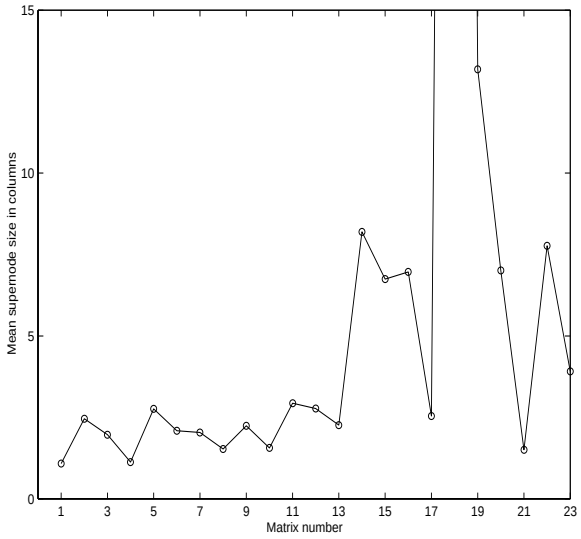

FIG. 5.3. *Mean number of columns per supernode.*

In contrast, *ops-per-ref* provides a lower bound on cache reuse and does depend on the details of the SuperLU algorithm. *Ops-per-ref* looks at each supernode-panel update separately and assumes that all the associated data are outside the cache before beginning the update. This pessimistic assumption limits the potential reuse to twice the panel size, $2\,w$.

Now we describe how we compute the average *ops-per-ref* for the entire factorization. Consider a single update from supernode $(r\colon s)$ to panel $(j\colon j+w-1)$. We assume that the supernode entry is brought into cache from main memory exactly once for the entire supernode-panel update, if it is used at all. Thus, during a single supernode-panel update, each entry accessed in the updating supernode accounts for between 2 and $2\,w$ *operations per reference*. Define *kmin* to be the number of the first row containing a nonzero in the panel,

$$kmin = \min_{j \le jj < j+w} \{k \mid k = \min_{r \le i \le s} \{i \mid A(i, jj) \neq 0\}\} .$$

Then $nnz(L(r\colon n, kmin\colon s))$ entries of the supernode are referenced in the supernode-panel update. The dense triangular solve in column jj of the update takes $(s-k+1)\cdot(s-k)$ flops. The matrix-vector multiply takes $2\cdot(s-k+1)\cdot nnz(L(s+1\colon n, s))$ flops. (We count both additions and multiplications.) For all panel updates, we sum the memory reference counts and the flop counts, then divide the second sum by the first to arrive at an average *ops-per-ref*. *Ops-per-ref* ranges from 2 to $2\,w$, with larger values indicating better cache use.

Figure 5.2 plots these two statistics against the speedup SuperLU achieved over the col-col code GP and against SuperLU's raw execution rate. It is clear that (perhaps surprisingly) *ops-per-nz* is superior to *ops-per-ref* as a predictor of either of these measures of performance. This is good news, because *ops-per-nz* measures the best case reuse, and *ops-per-ref* the worst case. But neither statistic captures all the variation in the performance.

5.4.2. How large are the supernodes? The supernode size determines the

(a) Matrix 1: 17758 rows, 16378 supernodes

(b) Matrix 2: 4929 rows, 2002 supernodes

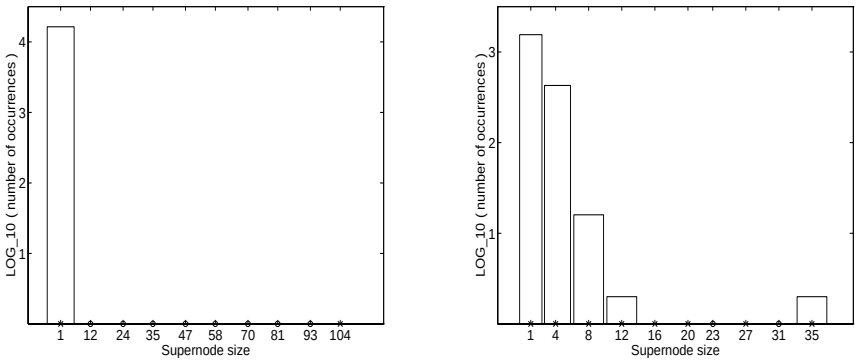

(c) Matrix 3: 4134 rows, 2099 supernodes

(d) Matrix 14: 7320 rows, 893 supernodes

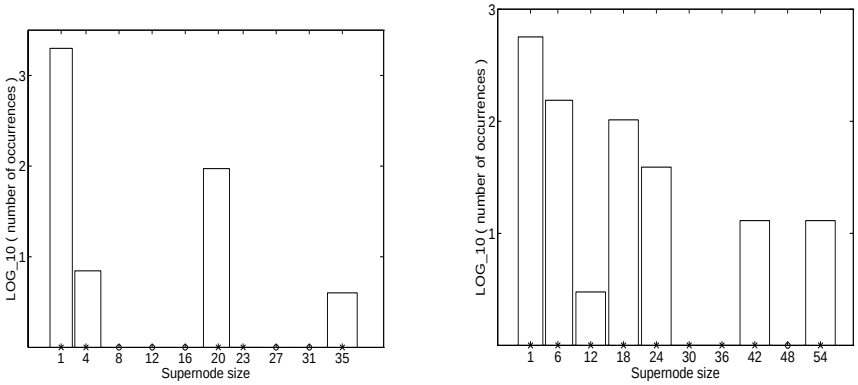

FIG. 5.4. *Distribution of supernode size for four matrices.*

size of the matrix passed to matrix-vector multiply and other Level 2 BLAS routines. Figure 5.3 shows the average number of columns in the supernodes of the matrices after amalgamating the relaxed supernodes at the bottom of the column etree (section 2.4). The average size is usually quite small.

More important than average size is the distribution of supernode sizes. In sparse Gaussian elimination, more fill tends to occur in the later stages. Usually there is a large percentage of small supernodes corresponding to the leaves of the column etree, even after amalgamation. Larger supernodes appear near the root. In Figure 5.4 we plot the histograms of the supernode size for four matrices chosen to exhibit a wide range of behavior. In the figure, the number at the bottom of each bar is the smallest supernode size in that bin. The mark "o" at the bottom of a bin indicates zero occurrences; otherwise, a "$*$" is put at the bottom of a bin. Relaxed supernodes of granularity $r = 4$ are used. Matrix 1 has 16378 supernodes, all but one of which have less than 12 columns; the single large supernode, with 115 columns, is the dense submatrix at the bottom right corner of F. Matrix 14 has more supernodes distributed over a wider spectrum; it has 13 supernodes with 54 to 59 columns. This matrix shows greater speedups over the nonsupernodal codes.

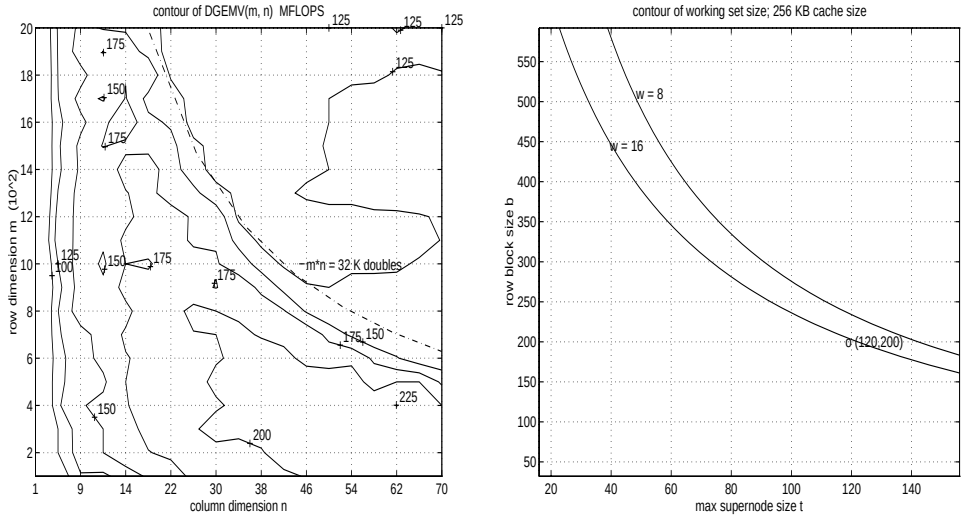

FIG. 5.5. (a) *Contour plot of* DGEMV *performance.* (b) *Contour plot of working set in* 2*-D algorithm.*

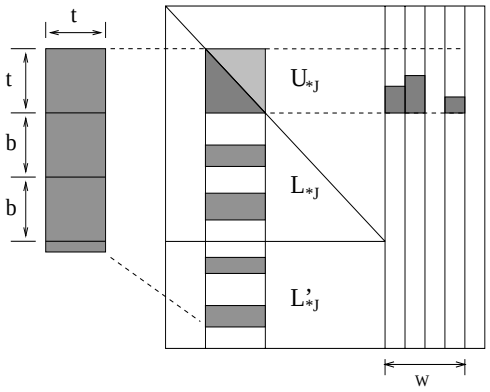

FIG. 5.6. *Parameters of the working set in the* 2*-D algorithm.*

5.4.3. Blocking parameters. In our hybrid algorithm (section 3.2.4), we need to select appropriate values for the parameters that describe the 2-D data blocking: panel width w, maximum supernode size t, and row block size b. The key considerations are that the active data we access in the inner loop (the *working set*) should fit into the cache, and that the matrices presented to the Level 2 BLAS routine DGEMV should be the sizes and shapes for which that routine is optimized. Here we describe in detail the methodology we used to choose parameters for the IBM RS/6000-590.

- **DGEMV optimization.** As indicated in the last column of Table 5.3, the majority of the floating-point operations are in the matrix-vector multiply. The dimensions (m, n) of the matrices in calls to DGEMV vary greatly depending on the supernode dimensions. Very often, the supernode is a tall and skinny matrix, that is, $m \gg n$. We measured the DGEMV Mflops rate for various m and n and present a contour plot in the (m, n) plane in Fig-

TABLE 5.4
Speedups achieved by each enhancement over GP on the RS/6000–590. The blocking parameters for SuperLU are $w = 8, t = 100,$ and $b = 200$.

	Matrix	GP (Seconds)	GP-Mod	SupCol	SuperLU
1	MEMPLUS	0.40	1.48	1.05	0.68
2	GEMAT11	0.27	1.69	1.29	1.00
3	RDIST1	1.90	2.75	2.24	1.94
4	ORANI678	13.86	3.55	2.98	3.10
5	MCFE	1.55	3.44	3.52	3.52
6	LNSP3937	7.11	3.39	3.86	3.54
7	LNS3937	7.77	3.39	3.85	3.55
8	SHERMAN5	3.98	3.43	4.57	4.23
9	JPWH991	2.78	3.61	4.21	4.48
10	SHERMAN3	7.43	3.54	5.99	5.27
11	ORSREG1	8.73	3.64	5.86	5.98
12	SAYLR4	17.51	3.67	5.99	6.30
13	SHYY161	163.14	3.65	6.46	5.67
14	GOODWIN	90.63	3.84	6.46	7.16
15	VENKAT01	355.50	3.86	8.33	8.87
16	INACCURA	544.91	4.17	7.24	7.94
17	AF23560	823.47	4.23	9.58	10.47
18	DENSE1000	83.48	4.21	10.22	14.54
19	RAEFSKY3	1571.63	4.30	11.54	14.00
20	EX11	3439.41	4.36	11.42	13.87
21	WANG3	1841.27	4.34	12.23	15.75
22	RAEFSKY4	3968.16	4.35	11.89	15.39
23	VAVASIS3	12342.97	4.79	13.11	15.63
Mean			3.64	6.67	7.52
Std			0.79	3.69	5.04

ure 5.5(a). Each contour represents a constant Mflops rate. The dashed curve represents $mn = 32$K double reals, or a cache capacity of 256 Kbytes. In the optimum region, we achieve more than 200 Mflops; outside this region, performance drops either because the matrices exceed the cache capacity or because the column dimension n is too small.

- **Working set.** By studying the data access pattern in the inner loop of the 2-D algorithm, lines 6–8 in Figure 3.3, we find that the working set size is the following function of w, t, and b, as shown in Figure 5.6:

$$WS = \underbrace{b \times t}_{\text{supernode rows}} + \underbrace{(t+b) \times w}_{\text{DGEMV vectors}} + \underbrace{b \times w}_{\text{SPA rows}} .$$

  In Figure 5.5(b), we fix two w values and plot the contour lines for $WS = 32$K in the (t, b) plane. If the point (t, b) is below the contour curve, then the working set can fit in a cache of 32K double reals, or 256 Kbytes.

Based on this analysis, we use the following rules to set the parameters.

First we choose w, the width of the panel in columns. Larger panels mean more reuse of cached data in the outer factorization, but also mean that the inner factorization (by the sup-col algorithm) must be applied to larger matrices. We find empirically that the best choice for w is between 8 and 16. Performance tends to degrade for larger w.

Next we choose b, the number of rows per block, and t, the maximum number of columns in a supernode. Recall that b and t are upper bounds on the row and column

TABLE 5.5
Speedups achieved by each enhancement over GP on the MIPS R8000. The blocking parameters for SuperLU are $w = 16, t = 100,$ and $b = 800$.

	Matrix	GP (Seconds)	GP-Mod	SupCol	SuperLU
1	MEMPLUS	0.42	1.51	1.10	0.59
2	GEMAT11	0.29	1.77	1.61	1.11
3	RDIST1	2.03	2.58	2.07	2.07
4	ORANI678	2.26	2.61	1.61	1.96
5	MCFE	0.60	2.93	2.73	2.61
6	LNSP3937	5.13	3.23	4.17	3.80
7	LNS3937	5.74	3.32	4.22	3.85
8	SHERMAN5	3.70	3.38	5.37	5.22
9	JPWH991	2.50	3.63	4.81	5.21
10	SHERMAN3	8.73	3.78	8.08	7.87
11	ORSREG1	8.18	3.72	7.24	8.10
12	SAYLR4	14.92	3.67	7.65	8.58
13	SHYY161	235.77	3.24	7.11	10.04
14	GOODWIN	103.66	3.45	8.87	11.27
15	VENKAT01	524.46	2.95	8.51	17.22
16	INACCURA	720.86	2.93	6.36	15.13
17	AF23560	1095.30	2.95	7.28	18.42
18	DENSE1000	113.28	3.34	11.99	30.21
19	RAEFSKY3	2263.80	2.88	6.54	28.87
20	EX11	5302.74	2.96	6.44	25.75
21	WANG3	2710.19	2.80	6.31	31.46
22	RAEFSKY4	6005.72	2.85	6.29	27.44
Mean			3.02	5.74	12.13
Std			0.57	2.75	10.48

dimensions of the call to DGEMV. We choose $t = 120$ and $b \approx 200$, which guarantees that the working set fits in cache (per Figure 5.5(b)), and that we can hope to be near the optimum region of DGEMV performance (per Figure 5.5(a)).

Recall that b is relevant only when we use rowwise blocking. This implies that the 2-D scheme adds overhead only if the updating supernode is small. In the actual code, the test for a large supernode is

if $ncol > 40$ **and** $nrow > b$ **then** the supernode is large,

where $nrow$ is the number of dense rows below the diagonal block of the supernode since $ncol$ is the number of dense columns of the supernode updating the panel. In practice, this choice usually gives the best performance.

The best choice of the parameters w, t, and b depends on the machine architecture and on the BLAS implementation, but it is largely independent of the matrix structure. Thus we do not expect each user of SuperLU to choose values for these parameters. Instead, our library code provides an inquiry function that returns the parameter values, much in the spirit of the LAPACK environment routine ILAENV. The machine-independent defaults will often give satisfactory performance. The methodology we have described here for the RS/6000 can serve as a guide for users who want to modify the inquiry function to give optimal performance for particular computer systems.

5.5. Comparison between SuperLU and multifrontal factorization. A number of codes for solving unsymmetric linear systems are available at the time of this writing, differing along several axes: emphasizing unsymmetric versus symmetric nonzero structure; using direct versus iterative methods; intended to be robust for

TABLE 5.6
Speedups achieved by each enhancement over GP on the Alpha 21164. *The blocking parameters for SuperLU are* $w = 16, t = 50,$ *and* $b = 100$.

	Matrix	GP (Seconds)	GP-Mod	SupCol	SuperLU
1	MEMPLUS	0.17	1.25	1.01	0.45
2	GEMAT11	0.13	1.54	1.26	0.84
3	RDIST1	0.80	1.76	1.77	1.45
4	ORANI678	0.92	1.74	1.47	1.45
5	MCFE	0.24	1.71	2.01	1.85
6	LNSP3937	2.09	1.93	2.61	2.27
7	LNS3937	2.33	1.94	2.59	2.27
8	SHERMAN5	1.50	1.92	3.13	3.00
9	JPWH991	1.06	2.14	3.20	3.20
10	SHERMAN3	3.65	2.10	4.06	3.93
11	ORSREG1	3.41	2.07	3.87	3.91
12	SAYLR4	6.73	2.05	4.01	4.34
13	SHYY161	102.19	1.81	3.97	4.58
14	GOODWIN	46.18	1.92	3.84	4.90
15	VENKAT01	235.01	1.71	4.08	7.00
16	INACCURA	333.24	1.72	3.48	6.07
17	AF23560	497.36	1.68	4.03	7.45
18	DENSE1000	49.29	1.82	4.82	10.38
19	RAEFSKY3	1065.88	1.68	4.00	10.02
20	EX11	1563.17	1.73	4.12	10.61
21	WANG3	1324.79	1.74	3.92	11.06
22	RAEFSKY4	2939.42	1.73	3.96	10.36
23	VAVASIS3	9477.62	1.83	4.51	11.48
Mean			1.80	3.29	5.34
Std			0.20	1.10	3.69

general problems versus efficient for specific applications; and in the public domain versus subject to commercial restrictions. A comprehensive comparison of all the codes against all possible metrics would be valuable but is not the purpose of the present paper. Rather, to locate the performance of SuperLU in the constellation of linear solvers, we compare it in detail with one alternative: UMFPACK version 2.1 [7, 8, 9]. This is a modern code that, like SuperLU, emphasizes unsymmetric structure and robustness for general problems. (A recent report [30] compares SuperLU and GP with some unsymmetric iterative algorithms.)

UMFPACK uses a multifrontal algorithm. Where the outer loop of a left-looking algorithm like SuperLU is over columns (or panels of columns) of the factors being computed, the outer loop of a multifrontal algorithm is over pivots (or blocks of pivots) being eliminated. All the updates created when a block is eliminated are computed at once and stored as a dense *update matrix*. Before a block of pivots is eliminated, all the update matrices contributing to that block are summed into a *frontal matrix*. The elimination step can use Level 2 or Level 3 BLAS because the arithmetic is carried out on the dense frontal matrix. Some extra intermediate storage is needed to record update matrices that have not yet been assembled into frontal matrices, and some extra data movement is needed for the assembly. UMFPACK does not use a column preordering; rather, it chooses row and column pivots to balance considerations of stability and sparsity by using approximate Markowitz counts with a pivot threshold. In principle, the pivot threshold can lead to a less accurate solution than strict partial pivoting; in practice, the lost accuracy can usually be retrieved by iterative refinement of the solution. In principle, the freedom to choose both row and column pivots dy-

TABLE 5.7
Speedups achieved by each enhancement over GP on the UltraSparc-I. The blocking parameters for SuperLU are $w = 8, t = 100$, and $b = 400$.

	Matrix	GP (Seconds)	GP-Mod	SupCol	SuperLU
1	MEMPLUS	0.36	1.17	1.08	0.58
2	GEMAT11	0.23	1.27	1.16	0.93
3	RDIST1	1.53	1.69	1.56	1.46
4	ORANI678	1.86	1.64	1.25	1.33
5	MCFE	0.52	1.97	1.85	1.92
6	LNSP3937	4.26	1.86	2.16	2.24
7	LNS3937	4.89	1.94	2.11	2.33
8	SHERMAN5	3.15	1.94	2.28	3.03
9	JPWH991	2.32	2.18	2.47	3.09
10	SHERMAN3	7.73	2.01	2.84	3.59
11	ORSREG1	7.2	1.97	2.69	3.52
12	SAYLR4	13.88	1.96	2.52	3.84
13	SHYY161	188.72	1.91	3.01	3.43
14	GOODWIN	89.30	1.89	2.62	4.41
18	DENSE1000	94.77	2.05	3.33	4.25
Mean			1.83	2.19	2.66
Std			0.28	0.69	1.22

namically could lead to sparser factors than strict partial pivoting; in practice, some matrices have sparser factors by one method and some by the other.

We compared UMFPACK and SuperLU on a group of 45 structurally unsymmetric matrices from a variety of applications, as described in Table 5.8. (This is a more comprehensive test set than the one we used in the earlier experiments with other left-looking codes described above.) We performed the experiments on the IBM RS/6000-590 described earlier. UMFPACK is written in Fortran; we compiled it with the AIX xlf compiler with -O3 optimization and linked it with the IBM ESSL library for BLAS calls. We used the parameter settings recommended by UMFPACK's authors [7].

UMFPACK does not include an initial column ordering step. For the initial column ordering in SuperLU, we ran Liu's multiple minimum degree algorithm [34] on the structure of $A^T A$. We report times for ordering and factorization separately. In applications where many matrices with the same nonzero structure but different values are factored, the cost of column ordering can be amortized over all the factorizations; in applications where only a single matrix is to be factored, preordering is part of the solution cost.

Table 5.9 gives time requirements and Table 5.10 gives memory requirements for the two codes on the matrices from the test set. The memory requirement we report includes only the memory actually used for the factorization, including working storage. Figures 5.7 and 5.8 summarize the comparison; each figure plots the relative time requirements against the relative space requirements for the two codes. Column preordering time is omitted in Figure 5.7 and included in Figure 5.8.

Neither code always dominates the other in either storage cost or time. Somewhat surprisingly, for 24 of the 45 matrices, the dynamic fill-reducing approach in UMFPACK seems to be less effective than the static preordering. SuperLU uses less memory for 60% of the matrices. When ordering time is not counted, SuperLU takes less time for 77% of the matrices. When ordering time is included, SuperLU takes less time for 44% of the matrices. For some matrices, such as MEMPLUS and ORANI678, the

TABLE 5.8
Characteristics of the unsymmetric matrices. StrSym is the fraction of nonzeros matched by nonzeros in symmetric locations. NumSym is the fraction of nonzeros matched by equal values in symmetric locations.

Matrix	n	$nnz(A)$	StrSym	NumSym	Discipline
CRY10000	10000	49699	.9979	.2012	crystal growth simulation
MEMPLUS	17758	99147	.9830	.5864	circuit simulation
DW8192	8192	41746	.9699	.9320	square Dielectric waveguide
GARON2	13535	390607	.9542	.6528	2D FEM, Navier-Stokes
JPWH_991	991	6027	.9469	.9469	circuit physics
AF23560	23560	460598	.9465	.0511	eigenvalue problem
BRAMLEY2	17933	1021849	.9257	.0466	nonlinear CFD
BRAMLEY1	17933	1021849	.9254	.2806	nonlinear CFD
LNSP3937	3937	25407	.8686	.1272	fluid flow
LNS_3937	3937	25407	.8686	.1272	fluid flow
WATSON5	1853	7803	.8590	.6170	circuit simulation
ADD32	4960	23884	.8310	.3979	computer component design
SHERMAN5	3312	20793	.7802	.2882	petroleum engineering
MHD4800A	4800	102252	.7718	.2806	magnetohydrodynamics
SHYY161	76480	329762	.7685	.3085	fluid flow
SHYY41	4720	20042	.7664	.3113	fluid flow
OLM5000	5000	19996	.7500	.5000	hydrodynamics
FS_541_2	541	4285	.7227	.1262	chemical kinetics
MCFE	765	24382	.7088	.0313	astrophysics
PORES_2	1224	9613	.6613	.4689	petroleum engineering
GOODWIN	7320	324772	.6423	.0194	fluid mechanics
TOLS4000	4000	8784	.5935	.3642	aeroelasticity
UTM5940	5940	83842	.5624	.0708	plasmas nuclear physics
BBMAT	38744	1771722	.5398	.0224	structure engineering CFD
RW5151	5151	20199	.4902	.0000	Markov chain transition
PSMIGR_1	3140	543162	.4816	.0161	demography
GRE_1107	1107	5664	.1954	.1954	discrete simulation
ONETONE2	36057	227628	.1482	.1020	nonlinear circuit
RDIST3A	2398	61896	.1468	.0074	chemical engineering
ONETONE1	36057	341088	.0989	.0681	nonlinear circuit
ORANI678	2529	90158	.0728	.0023	economics
RDIST1	4134	94408	.0620	.0034	chemical engineering
RADFR1	1048	13299	.0600	.0066	chemical engineering
RDIST2	3198	56934	.0491	.0037	chemical engineering
MAHINDAS	1258	7682	.0302	.0166	economics
LHR04	4101	82682	.0159	.0010	chemical engineering
LHR01	1477	18592	.0085	.0013	chemical engineering
HYDR1	5308	23752	.0040	.0006	chemical engineering
EXTR1	2837	11407	.0040	.0005	chemical engineering
WEST2021	2021	7353	.0039	.0006	chemical engineering
VAVASIS1	4408	95752	.0033	.0003	2D PDE
VAVASIS2	11924	306842	.0025	.0001	2D PDE
GEMAT11	4929	33108	.0017	.0003	electrical power
LHR71	70304	1528092	.0014	.0002	chemical engineering
VAVASIS3	41092	1683902	.0009	.0000	2D PDE

MMD ordering takes significantly more time than factorization. It should be noted that our current approach to ordering can be improved. For example, the column minimum degree algorithm used in Matlab [26] implements the minimum degree algorithm on A^TA without explicitly forming the structure of A^TA. In recent work, Davis, Gilbert, and Ng [10, 45] are investigating better minimum degree algorithms for unsymmetric matrices that we expect to improve both fill and runtime.

For 9 of the 13 problems whose dimensions are at least 10000, SuperLU outper-

TABLE 5.9
*Performance of SuperLU and UMFPACK on an IBM RS/*6000-590 *with* 768 *MB of memory. Numbers in parenthesis are factorization rate in Mflops. A "+" before a number indicates SuperLU outperforms UMFPACK. UMFPACK ran out of memory on* BBMAT.

Matrix	Seconds (Mflops)				
	SuperLU			UMFPACK	
	order	factor			
CRY10000	+0.35	+1.73	(28)	2.30	(31)
MEMPLUS	89.58	+1.65	(22)	1.94	(1)
DW8192	+0.38	+3.69	(28)	5.64	(53)
GARON2	+3.86	+22.12	(58)	53.83	(97)
JPWH_991	0.13	0.53	(36)	0.25	(20)
AF23560	+9.22	+62.11	(80)	224.89	(112)
BRAMLEY2	+23.01	+57.97	(93)	279.42	(101)
BRAMLEY1	+23.00	+58.10	(93)	274.46	(106)
LNSP3937	+0.40	+1.28	(21)	3.11	(36)
LNS_3937	+0.40	+1.12	(24)	2.52	(31)
WATSON5	0.22	+0.13	(12)	0.20	(.4)
ADD32	0.24	+0.17	(.9)	0.36	(.3)
SHERMAN5	0.24	+0.77	(29)	0.88	(33)
MHD4800A	+0.74	+0.58	(8)	7.10	(65)
SHYY161	+2.33	+22.04	(47)	54.14	(63)
SHYY41	+0.14	+0.49	(17)	0.69	(13)
OLM5000	0.12	+0.14	(.6)	0.21	(.4)
FS_541_2	0.08	+0.05	(5)	0.12	(3)
MCFE	1.21	+0.22	(13)	0.25	(17)
PORES_2	+0.10	+0.17	(14)	0.28	(9)
GOODWIN	+5.34	+10.31	(49)	22.37	(64)
TOLS4000	+0.06	+0.06	(.6)	0.23	(.1)
UTM5940	+0.93	+3.02	(43)	4.23	(61)
BBMAT	+185.77	+821.49	(54)	failed	
RW5151	+0.16	+1.24	(27)	1.97	(28)
PSMIGR_1	290.08	187.86	(89)	93.93	(100)
GRE_1107	0.15	0.39	(23)	0.31	(20)
ONETONE2	3.00	+4.98	(26)	7.02	(26)
RDIST3A	0.77	0.48	(24)	0.44	(17)
ONETONE1	15.77	48.00	(53)	38.87	(79)
ORANI678	73.41	+1.32	(24)	1.63	(4)
RDIST1	1.02	0.76	(23)	0.74	(12)
RADFR1	0.09	+0.10	(11)	0.11	(4)
RDIST2	0.50	0.45	(16)	0.41	(8)
MAHINDAS	0.69	+0.09	(7)	0.17	(2)
LHR04	2.65	+1.10	(15)	1.71	(9)
LHR01	0.46	+0.21	(11)	0.29	(7)
HYDR1	+0.53	+0.26	(3)	0.81	(3)
EXTR1	+0.13	+0.13	(1)	0.32	(1)
WEST2021	+0.08	+0.08	(.9)	0.26	(.2)
VAVASIS1	12.49	7.50	(81)	4.40	(46)
VAVASIS2	76.51	38.06	(94)	37.09	(56)
GEMAT11	0.24	+0.26	(4)	0.39	(2)
LHR71	48.23	+23.12	(21)	40.74	(17)
VAVASIS3	1091.72	660.89	(98)	533.35	(122)

forms UMFPACK both in factorization time and in memory.

6. Remarks.

6.1. The rest of the SuperLU package. In addition to the LU factorization kernel described in this paper, we have developed a suite of supporting routines to

TABLE 5.10
*Performance of SuperLU and UMFPACK on an IBM RS/*6000-590 *with* 768 *MB of memory. A "+" before a number indicates SuperLU outperforms UMFPACK. UMFPACK ran out of memory on* BBMAT.

Matrix	$nnz(F)$		Memory (MB)	
	SuperLU	UMF	SuperLU	UMF
CRY10000	+651412	713789	+8.73	9.01
MEMPLUS	356842	157474	7.49	2.29
DW8192	+771299	1315910	+9.87	25.76
GARON2	+4744645	8298907	+51.83	144.82
JPWH_991	144499	67703	1.71	1.56
AF23560	+12950915	27979292	+136.35	392.83
BRAMLEY2	+12652411	29803064	+136.67	376.05
BRAMLEY1	+12684407	29355331	+137.04	445.65
LNSP3937	+367691	586976	+4.79	14.62
LNS_3937	+373477	494120	+4.76	12.36
WATSON5	41152	17471	0.48	0.23
ADD32	40444	35660	1.45	0.47
SHERMAN5	+238905	246358	+3.20	5.16
MHD4800A	+227469	1550890	+3.38	31.29
SHYY161	+6695357	9524169	+83.29	145.83
SHYY41	195685	167723	2.97	2.67
OLM5000	40002	35004	1.47	0.42
FS_541_2	+16157	18458	0.30	0.28
MCFE	+63940	80595	+0.87	1.52
PORES_2	+54028	54405	+0.83	1.01
GOODWIN	+2764580	5162683	+31.30	64.33
TOLS4000	8872	8784	0.97	0.16
UTM5940	+994050	1225186	+11.46	22.31
BBMAT	+49987183	failed	+538.76	failed
RW5151	+379538	440271	+5.12	9.78
PSMIGR_1	8709183	6369782	+88.75	355.98
GRE_1107	110883	79048	+1.40	1.74
ONETONE2	+1270569	1337258	+20.92	22.34
RDIST3A	254176	189852	+2.99	3.53
ONETONE1	+4676473	5145165	+56.47	103.99
ORANI678	804641	122079	7.99	6.49
RDIST1	398008	277947	4.76	2.97
RADFR1	50483	29699	0.72	0.35
RDIST2	230084	151860	2.95	1.72
MAHINDAS	23909	14126	0.54	0.31
LHR04	342084	313473	+4.62	4.65
LHR01	66863	60264	1.04	0.85
HYDR1	+81335	111919	2.03	1.61
EXTR1	+34461	36836	0.97	0.48
WEST2021	19179	14615	0.63	0.20
VAVASIS1	1495345	915869	+16.21	18.21
VAVASIS2	5523041	4249821	+59.07	89.19
GEMAT11	87054	67149	1.93	0.82
LHR71	+7508569	8740127	+96.14	124.55
VAVASIS3	+39798599	41746313	+418.65	470.75

solve linear systems, all of which are available from Netlib.[4] The complete SuperLU package [12] includes column preordering for sparsity, based on Liu's multiple minimum degree algorithm [34] applied to the structure of A^TA. (As noted above, we plan to replace this with a new code that does not form A^TA.) SuperLU also includes condition number estimation, iterative refinement of solutions, and componentwise error

[4]URL: http://www.netlib.org/scalapack/prototype/

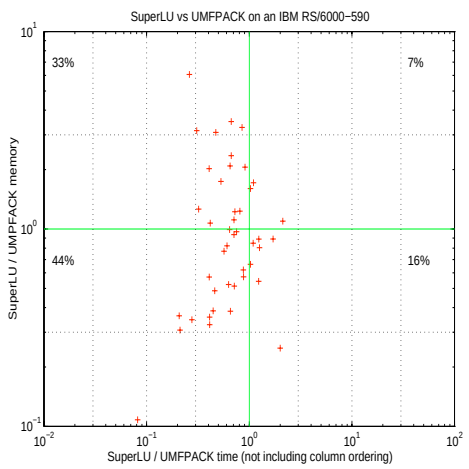

FIG. 5.7. *Compare SuperLU to UMFPACK, when MMD ordering time is not included.*

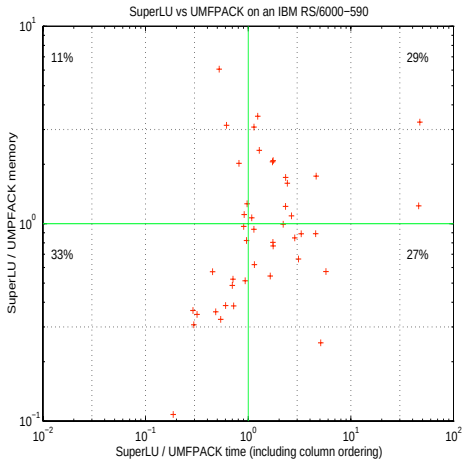

FIG. 5.8. *Compare SuperLU to UMFPACK, when MMD ordering time is included.*

bounds for the refined solutions. These are all based on the dense matrix routines in LAPACK [3]. In addition, SuperLU includes a Matlab mex-file interface, so that our factor and solve routines can be called as alternatives to those built into Matlab.

6.2. Effect of the matrix on performance. The supernodal approach reduces both symbolic and numeric computation time. But unsymmetric supernodes tend to be smaller than supernodes in symmetric matrices. The supernode-panel method is most effective for large problems with enough dense submatrices to use dense block operations and exploit data locality. In this regard, the dense 1000×1000 example illustrates the largest possible gains. Dense blocks are necessary for top performance in all modern factorization algorithms, whether left-looking, right-looking, multifrontal, or any other style.

Our goal has been to develop sparse LU software that works well for problems with a wide range of characteristics. It is harder to achieve high flop rates on problems that are very sparse and have no structure to exploit; it is easier on problems that are denser or become so during elimination. Fortunately, the "hard" matrices by this definition generally take many fewer floating-point operations than the "easy" ones, and hence take much less time to factor. Our combination of 1-D and 2-D blocking techniques gives a good performance compromise for all the problems we have studied, and with particularly good performance on the largest problems.

6.3. Effect of the computer system on performance. We have studied several characteristics of the computing platform that can affect the overall performance, including the Level 2 BLAS speed and the cache size. We showed how to systematically make a good choice of the blocking parameters in the code so as to maximize the speed of the numeric kernel, using the IBM RS/6000-590 as an example. We expect this methodology to be applicable to other systems (and BLAS implementations) as well.

6.4. Possible enhancements. We are considering several possible enhancements to the SuperLU code. One is to switch to a dense LU code at a late stage of the factorization. It would be difficult to implement this in a sup-col code, because that code is strictly left-looking, and only one column of the matrix is factored at a

time. However, this would be much easier in the supernode-panel code. At the time we decide to switch, we simply treat the rest of the matrix columns (say, d of them) as one panel, and perform the panel update to $A(1{:}n, n-d+1{:}n)$. (One might want to split this panel up for better cache behavior.) Then the reduced matrix at the bottom right corner can be factored by calling an efficient dense code, for example, from LAPACK [3]. The dense code does not spend time on symbolic structure prediction and pruning, thus streamlining the numeric computation. We believe that, for large problems, the final dense submatrix will be big enough to make the switch beneficial. For example, for a 2-D $k \times k$ square grid problem ordered by nested dissection, the dimension of the final dense submatrix is $\frac{3}{2}k \times \frac{3}{2}k$; for a 3-D $k \times k \times k$ cubic grid, it is $\frac{3}{2}k^2 \times \frac{3}{2}k^2$, if pivots come from the diagonal. The Harwell library code MA48 [16, 18] employs such a switch to dense code, which has a significant and beneficial effect on performance.

To enhance SuperLU's performance on small and extremely sparse problems, it would be possible to make a choice at runtime whether to use supernode-panel, supcol, or col-col updates. The choice would depend on the size of the matrix A and the expected properties of its supernodes; it might be based on an efficient symbolic computation of the density and supernode distribution of the Cholesky factor of A^TA [28].

Could we make supernode-panel updates more effective by improving the similarity between the row structures of the columns in a panel? We believe this could be accomplished with a more sophisticated column permutation strategy. We could partition the nodes of the column etree into connected subtrees, grouping together nodes that have common descendants (and therefore the potential for updates from the same supernodes). Then the overall column order would be a two-level postorder, first within the subtrees (panels) and then among them. Again, it might be possible to use information about the Cholesky supernodes of A^TA to guide this grouping.

We are also developing a parallel sparse LU algorithm based on SuperLU [11, 33]. In this context, we target large problems, especially those too big to be solved on a uniprocessor system. Therefore, we plan to parallelize the 2-D blocked supernode-panel algorithm, which has very good asymptotic behavior for large problems. The 2-D block-oriented layout has been shown to scale well for parallel sparse Cholesky factorization [31, 40].

Acknowledgments. We thank Rob Schreiber and Ed Rothberg for very helpful discussions during the development of SuperLU. Rob pushed us to find a way to do the supernode-panel update; Ed suggested the *ops-per-ref* statistic and also provided access to the SGI MIPS R8000. We are grateful to Tim Davis, Steve Vavasis, and Randy Bramley for making their matrices available to us. We thank the referees for their suggestions to improve the presentation of the paper. The Institute for Mathematics and Its Applications at the University of Minnesota provided the fertile environment in which this work began.

REFERENCES

[1] R. Agarwal, F. Gustavson, P. Palkar, and M. Zubair, *A performance analysis of the subroutines in the ESSL/LAPACK call conversion interface (CCI)*, IBM T.J. Watson Research Center, Yorktown Heights, NY, 1994.

[2] P. R. Amestoy and I. Duff, *Vectorization of a multiprocessor multifrontal code*, Int. J. Supercomputer Appl., 7 (1993), pp. 64–82.

[3] E. Anderson et al., *LAPACK User's Guide,* 2*nd ed.*, SIAM, Philadelphia, PA, 1995.

[4] C. Ashcraft and R. Grimes, *The influence of relaxed supernode partitions on the multifrontal method*, ACM Trans. Math. Software, 15 (1989), pp. 291–309.
[5] C. Ashcraft, R. Grimes, J. Lewis, B. Peyton, and H. Simon, *Progress in sparse matrix methods for large sparse linear systems on vector supercomputers*, Internat. J. of Supercomputer Applications, 1 (1987), pp. 10–30.
[6] J. Bilmes, K. Asanovic, J. Demmel, D. Lam, and C.-W. Chin, *Optimizing Matrix Multiply Using PHiPAC: A Portable, High-Performance, ANSI C Coding Methodology*, Tech. Report CS-96-326, Computer Science Department., University of Tennessee, Knoxville, TN, May 1996. (LAPACK Working Note #111).
[7] T. A. Davis, *User's Guide for the Unsymmetric-Pattern MultiFrontal Package (UMFPACK)*, Tech. Report TR-95-004, Computer and Information Sciences Department, University of Florida, Gainesville, FL, January 1995.
[8] T. A. Davis and I. S. Duff, *A combined unifrontal/multifrontal method for unsymmetric sparse matrices*, Tech. Report TR-95-020, Computer and Information Sciences Department, University of Florida, Gainesville, FL, 1995.
[9] T. A. Davis and I. S. Duff, *An unsymmetric-pattern multifrontal method for sparse LU factorization*, SIAM J. Matrix Anal. Appl., 18 (1997), pp. 140–158.
[10] T. A. Davis, J. R. Gilbert, E. Ng, and B. W. Peyton, *A column approximate minimum degree ordering algorithm*, Presented at Sixth SIAM Symposium on Applied Linear Algebra, Snowbird, UT, 1997.
[11] J. W. Demmel, J. R. Gilbert, and X. S. Li, *An Asynchronous Parallel Supernodal Algorithm for Sparse Gaussian Elimination*, Tech. Report UCB//CSD-97-943, Computer Science Division, University of California, Berkeley, CA, 1997.
[12] J. W. Demmel, J. R. Gilbert, and X. S. Li, *SuperLU User's Guide*, Tech. Report UCB//CSD-97-944, Computer Science Division, University of California, Berkeley, CA, 1997.
[13] D. Dodson and J. Lewis, *Issues relating to extension of the basic linear algebra subprograms*, ACM SIGNUM Newsletter, 20 (1985), pp. 2–18.
[14] J. J. Dongarra, J. D. Croz, S. Hammarling, and R. J. Hanson, *An extended set of basic linear algebra subroutines*, ACM Trans. Math. Software, 14 (1988), pp. 1–17, 18–32.
[15] I. Duff and J. Reid, *The multifrontal solution of indefinite sparse symmetric linear equations*, ACM Trans. Math. Software, 9 (1983), pp. 302–325.
[16] I. Duff and J. K. Reid, *The design of MA*48*, a code for the direct solution of sparse unsymmetric linear systems of equations*, ACM Trans. Math. Software, 22 (1996), pp. 187–226.
[17] I. S. Duff, R. Grimes, and J. Lewis, *Sparse matrix test problems*, ACM Trans. Math. Software, 15 (1989), pp. 1–14.
[18] I. S. Duff and J. K. Reid, *MA*48*, a Fortran Code for Direct Solution of Sparse Unsymmetric Linear Systems of Equations*, Tech. Report RAL–93–072, Rutherford Appleton Laboratory, Oxon, UK, 1993.
[19] S. C. Eisenstat, J. R. Gilbert, and J. W. Liu, *A supernodal approach to a sparse partial pivoting code*, in Householder Symposium 12, Los Angeles, CA, 1993.
[20] S. C. Eisenstat and J. W. H. Liu, *Exploiting structural symmetry in sparse unsymmetric symbolic factorization*, SIAM J. Matrix Anal. Appl., 13 (1992), pp. 202–211.
[21] S. C. Eisenstat and J. W. H. Liu, *Exploiting structural symmetry in a sparse partial pivoting code*, SIAM J. Sci. Statist. Comput., 14 (1993), pp. 253–257.
[22] J. A. George and E. Ng, *An implementation of Gaussian elimination with partial pivoting for sparse systems*, SIAM J. Sci. Statist. Comput., 6 (1985), pp. 390–409.
[23] J. A. George and E. Ng, *Symbolic factorization for sparse Gaussian elimination with partial pivoting*, SIAM J. Sci. Statist. Comput., 8 (1987), pp. 877–898.
[24] J. R. Gilbert, *Predicting structure in sparse matrix computations*, SIAM J. Matrix Anal. Appl., 15 (1994), pp. 62–79.
[25] J. R. Gilbert and J. W. H. Liu, *Elimination structures for unsymmetric sparse LU factors*, SIAM J. Matrix Anal. Appl., 14 (1993), pp. 334–352.
[26] J. R. Gilbert, C. Moler, and R. Schreiber, *Sparse matrices in Matlab: Design and implementation*, SIAM J. Matrix Anal. Appl., 13 (1992), pp. 333–356.
[27] J. R. Gilbert and E. Ng, *Predicting structure in nonsymmetric sparse matrix factorizations*, in Graph Theory and Sparse Matrix Computation, A. George, J. R. Gilbert, and J. W. H. Liu, eds., Springer–Verlag, New York, Berlin, 1993.
[28] J. R. Gilbert, E. G. Ng, and B. W. Peyton, *An efficient algorithm to compute row and column counts for sparse Cholesky factorization*, SIAM J. Matrix Anal. Appl., 15 (1994), pp. 1075–1091.
[29] J. R. Gilbert and T. Peierls, *Sparse partial pivoting in time proportional to arithmetic operations*, SIAM J. Sci. Statist. Comput., 9 (1988), pp. 862–874.

[30] J. R. Gilbert and S. Toledo, *An assessment of incomplete LU preconditioners for nonsymmetric linear systems*, manuscript, 1997.
[31] A. Gupta and V. Kumar, *Optimally scalable parallel sparse Cholesky factorization*, in Proc. 7th SIAM Conference on Parallel Processing for Scientific Computing, D.H. Bailey et. al., eds., SIAM, Philadelphia, 1995, pp. 442–447.
[32] *International Business Machines Corporation Engineering and Scientific Subroutine Library, Guide and Reference*, Version 2 Release 2, Order No. SC23-0526-01, 1994.
[33] X. S. Li, *Sparse Gaussian Elimination on High Performance Computers*, Tech. Report UCB//CSD-96-919, Computer Science Division, Ph.D. dissertation, University of California, Berkeley, CA, 1996.
[34] J. W. Liu, *Modification of the minimum degree algorithm by multiple elimination*, ACM Trans. Math. Software, 11 (1985), pp. 141–153.
[35] J. W. H. Liu, *The role of elimination trees in sparse factorization*, SIAM J. Matrix Anal. Appl., 11 (1990), pp. 134–172.
[36] *PDS: The performance database server*, *http://performance.netlib.org/performance/*, May 1995.
[37] E. G. Ng and B. W. Peyton, *Block sparse Cholesky algorithms on advanced uniprocessor computers*, SIAM J. Sci. Statist. Comput., 14 (1993), pp. 1034–1056.
[38] E. Rothberg and A. Gupta, *Efficient sparse matrix factorization on high-performance workstations—exploiting the memory hierarchy*, ACM Trans. Math. Software, 17 (1991), pp. 313–334.
[39] E. Rothberg and A. Gupta, *An evaluation of left-looking, right-looking and multifrontal approaches to sparse Cholesky factorization on hierarchical-memory machines*, Internat. J. High Speed Comput., 5 (1993), pp. 537–593.
[40] E. E. Rothberg and A. Gupta, *An efficient block-oriented approach to parallel sparse Cholesky factorization*, in Supercomput., 1993, pp. 503–512.
[41] A. H. Sherman, *On the Efficient Solution of Sparse Systems of Linear and Nonlinear Equations*, Ph.D. thesis, Yale University, New Haven, CT, 1975.
[42] A. H. Sherman, *Algorithm* 533: *NSPIV, a FORTRAN subroutine for sparse Gaussian elimination with partial pivoting*, ACM Trans. Math. Software, 4 (1978), pp. 391–398.
[43] H. Simon, P. Vu, and C. Yang, *Performance of a Supernodal General Sparse Solver on the CRAY Y-MP:* 1.68 *GFLOPS with Autotasking*, Tech. Report TR SCA-TR-117, Boeing Computer Services, Seattle, WA, 1989.
[44] S. A. Vavasis, *Stable finite elements for problems with wild coefficients*, SIAM J. Numer. Anal., 33 (1996), pp. 890–916.
[45] S. I. Larimore, *An Approximate Minimum Degree Ordering Algorithm*, Tech. Report CISE-TR-98-016, Department of Computer and Information Science and Engineering, University of Florida, Gainesville, FL, 1998.