

# ROW MODIFICATIONS OF A SPARSE CHOLESKY FACTORIZATION \*

TIMOTHY A. DAVIS<sup>†</sup> AND WILLIAM W. HAGER<sup>‡</sup>

**Abstract.** Given a sparse, symmetric positive definite matrix  $\mathbf{C}$  and an associated sparse Cholesky factorization  $\mathbf{LDL}^T$ , we develop sparse techniques for updating the factorization after a symmetric modification of a row and column of  $\mathbf{C}$ . We show how the modification in the Cholesky factorization associated with this rank two modification of  $\mathbf{C}$  can be computed efficiently using a sparse rank one technique developed in an earlier paper [*SIAM J. Matrix Anal. Appl.*, 20 (1999), pp. 606-627]. We also determine how the solution of a linear system  $\mathbf{Lx} = \mathbf{b}$  changes after changing a row and column of  $\mathbf{C}$  or after a rank- $r$  change in  $\mathbf{C}$ .

**Key words.** numerical linear algebra, direct methods, Cholesky factorization, sparse matrices, mathematical software, matrix updates.

**AMS subject classifications.** 65F05, 65F50, 65Y20, 65-04.

**1. Introduction.** The problem of updating a Cholesky factorization after a small rank change in the matrix is a fundamental problem with many applications including optimization algorithms, least-squares problems in statistics, the analysis of electrical circuits and power systems, structural mechanics, boundary condition changes in partial differential equations, domain decomposition methods, and boundary element methods [12]. A variety of techniques for modifying a dense Cholesky factorization are given in the classic reference [11]. Recently [3, 4], we consider a sparse Cholesky factorization  $\mathbf{LDL}^T$  of a symmetric, positive definite matrix  $\mathbf{C}$ , and the modification associated with a rank- $r$  change of the form  $\bar{\mathbf{C}} = \mathbf{C} \pm \mathbf{W}\mathbf{W}^T$ , where  $\mathbf{W}$  is  $n$ -by- $r$  with  $r$  typically much less than  $n$ . In this paper, we consider a special, but important, rank-2 change corresponding to a symmetric modification of a row and column of  $\mathbf{C}$ . Although we could, in principle, use our previous methodology to update the factorization, we observe that this rank-2 approach is much less efficient than the streamlined approach we develop here. In fact, the rank-2 approach could result in a completely dense modification of the factorization, where nonzero entries are first introduced and then canceled out. With the new approach, the work connected with *removing* a nonzero row and column is comparable to the work associated with a sparse rank-1 *update*, while the work associated with *adding* a new nonzero row and column is comparable to the work associated with a sparse rank-1 *downdate*. This connection between the modification of the matrix and the modification of the factorization is nonintuitive: When we remove elements from the matrix, we update the factorization; when we add elements to the matrix, we downdate the factorization. As a by-product of our analysis, we show how the solution to a triangular system  $\mathbf{Lx} = \mathbf{b}$  changes when both  $\mathbf{L}$  and  $\mathbf{b}$  change as a result of the row and column modification

---

\* October 21, 2003. This material is based upon work supported by the National Science Foundation under Grant No. CCR-0203270. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the National Science Foundation.

<sup>†</sup>davis@cise.ufl.edu, <http://www.cise.ufl.edu/~davis>, PO Box 116120, Department of Computer and Information Science and Engineering, University of Florida, Gainesville, FL 32611-6120. Phone (352) 392-1481. Fax (352) 392-1220.

<sup>‡</sup>hager@math.ufl.edu, <http://www.math.ufl.edu/~hager>, PO Box 118105, Department of Mathematics, University of Florida, Gainesville, FL 32611-8105. Phone (352) 392-0281. Fax (352) 392-8357.

problem discussed in this paper, or as a result of a rank- $r$  change to  $\mathbf{C}$  [3, 4].

One specific application for the techniques developed in this paper is the Linear Programming Dual Active Set Algorithm (LPDASA) [5, 13, 14, 15, 16, 17]. In this dual approach to linear programming, inactive inequalities are identified dynamically, during the solution process; dropping these inactive inequalities amounts to removing a row and a column from a Cholesky factored matrix. In the same way, when a dropped inequality later becomes active, a new row and column must be inserted in the matrix, and the resulting modification in the Cholesky factorization evaluated. In general, the techniques developed in this paper are useful in any setting where a system of equations is solved repeatedly, with equations added or dropped before each solve.

A brief overview of our paper follows. In Section 2 we consider the *row addition* problem, in which a row and column of a dense matrix  $\mathbf{C}$ , originally zero except for the diagonal element, are modified. Section 3 extends this method to a sparse matrix  $\mathbf{C}$ . The *row deletion* problem is the opposite of row addition, and is discussed in Sections 4 and 5 (the dense and sparse cases, respectively). Sections 6 and 7 describe how to make arbitrary modifications to a row and column of  $\mathbf{C}$ . The efficient methods presented in Sections 2 through 7 are contrasted with performing the modifications as a rank-2 outer product modification in Section 8, which is shown to be costly, particularly in the sparse case. Section 9 shows how to efficiently modify the solution to  $\mathbf{L}\mathbf{x} = \mathbf{b}$  when  $\mathbf{L}$  and  $\mathbf{b}$  change. A brief presentation of the experimental performance of these methods in the context of matrices arising in linear programming is given in Section 10.

We use the notation  $\overline{\mathbf{C}}$  to denote the matrix  $\mathbf{C}$  after it has been modified. Bold uppercase  $\mathbf{A}$  refers to a matrix. Bold lowercase  $\mathbf{a}$  is a column vector, thus,  $\mathbf{a}^\top$  always refers to a row vector. Plain lowercase letters (such as  $a$  and  $\alpha$ ) are scalars. We use  $|\mathbf{A}|$  to denote the number of nonzero entries in the sparse matrix  $\mathbf{A}$ . Without parentheses, the notation  $\mathbf{A}_i$  or  $\mathbf{A}_{ij}$  refer to submatrices of a matrix  $\mathbf{A}$  (sometimes 1-by-1 submatrices). We use parenthesis  $(\mathbf{A})_{ij}$  to refer to the entry in row  $i$  and column  $j$  of the matrix  $\mathbf{A}$ ,  $(\mathbf{A})_{*j}$  to refer to all of column  $j$  of  $\mathbf{A}$ , and  $(\mathbf{A})_{i*}$  to refer to all of row  $i$  of  $\mathbf{A}$ . When counting floating-point operations, we count one flop for any arithmetic operation including  $*$ ,  $/$ ,  $+$ ,  $-$ , and  $\sqrt{\phantom{x}}$ .

**2. Adding a row and column to  $\mathbf{C}$ .** Modifying the  $k$ th row and column of a matrix  $\mathbf{C}$  can be considered in two different contexts. If we have a rectangular  $n$ -by- $m$  matrix  $\mathbf{A}$ , and  $\mathbf{C} = \alpha\mathbf{I} + \mathbf{A}\mathbf{A}^\top$ , then modifying row  $k$  of  $\mathbf{A}$  leads to changes in the  $k$ th row and column of  $\mathbf{C}$ . In the more general context, we consider arbitrary changes to the  $k$ th row and column of a symmetric positive definite matrix  $\mathbf{C}$ . We will first consider the special case where row  $k$  is initially zero and becomes nonzero in  $\overline{\mathbf{A}}$  (the *row addition* case). Alternatively, the  $k$ th row and column of  $\mathbf{C}$  is initially a multiple of the  $k$ th row and column of the identity matrix, and changes to some other value.

Let  $\overline{\mathbf{a}}_2^\top$  be the new nonzero  $k$ th row of  $\overline{\mathbf{A}}$ ,

$$\mathbf{A} = \begin{bmatrix} \mathbf{A}_1 \\ \mathbf{0}^\top \\ \mathbf{A}_3 \end{bmatrix}, \quad \overline{\mathbf{A}} = \begin{bmatrix} \mathbf{A}_1 \\ \overline{\mathbf{a}}_2^\top \\ \mathbf{A}_3 \end{bmatrix},$$

where  $\mathbf{0}^\top$  is a row vector whose entries are all 0. This leads to a modification to row

and column  $k$  of  $\mathbf{C}$ ,

$$\mathbf{C} = \begin{bmatrix} \alpha\mathbf{I} + \mathbf{A}_1\mathbf{A}_1^\top & \mathbf{0} & \mathbf{A}_1\mathbf{A}_3^\top \\ \mathbf{0}^\top & \alpha & \mathbf{0}^\top \\ \mathbf{A}_3\mathbf{A}_1^\top & \mathbf{0} & \alpha\mathbf{I} + \mathbf{A}_3\mathbf{A}_3^\top \end{bmatrix} = \begin{bmatrix} \mathbf{C}_{11} & \mathbf{0} & \mathbf{C}_{31}^\top \\ \mathbf{0}^\top & c_{22} & \mathbf{0}^\top \\ \mathbf{C}_{31} & \mathbf{0} & \mathbf{C}_{33} \end{bmatrix},$$

where  $c_{22} = \alpha$ . We can let  $\alpha$  be zero; even though  $\mathbf{C}$  would no longer be positive definite, the submatrix excluding row and column  $k$  would still be positive definite. The linear system  $\mathbf{C}\mathbf{x} = \mathbf{b}$  is well-defined in this case, except for  $\mathbf{x}_k$ . The new matrix  $\bar{\mathbf{C}}$  is given as

$$\bar{\mathbf{C}} = \begin{bmatrix} \alpha\mathbf{I} + \mathbf{A}_1\mathbf{A}_1^\top & \mathbf{A}_1\bar{\mathbf{a}}_2 & \mathbf{A}_1\mathbf{A}_3^\top \\ \bar{\mathbf{a}}_2^\top\mathbf{A}_1^\top & \alpha + \bar{\mathbf{a}}_2^\top\bar{\mathbf{a}}_2 & \bar{\mathbf{a}}_2^\top\mathbf{A}_3^\top \\ \mathbf{A}_3\mathbf{A}_1^\top & \mathbf{A}_3\bar{\mathbf{a}}_2 & \alpha\mathbf{I} + \mathbf{A}_3\mathbf{A}_3^\top \end{bmatrix} = \begin{bmatrix} \mathbf{C}_{11} & \bar{\mathbf{c}}_{12} & \mathbf{C}_{31}^\top \\ \bar{\mathbf{c}}_{12}^\top & \bar{c}_{22} & \bar{\mathbf{c}}_{32}^\top \\ \mathbf{C}_{31} & \bar{\mathbf{c}}_{32} & \mathbf{C}_{33} \end{bmatrix}.$$

Thus, adding a row  $k$  to  $\mathbf{A}$  is equivalent to adding a row and column  $k$  to  $\mathbf{C}$ . Note that changing row and column  $k$  of  $\mathbf{C}$  from zero (except for the diagonal entry) to a nonzero value can also be viewed as increasing the dimension of the  $(n-1)$ -by- $(n-1)$  matrix,

$$\begin{bmatrix} \mathbf{C}_{11} & \mathbf{C}_{31}^\top \\ \mathbf{C}_{31} & \mathbf{C}_{33} \end{bmatrix}.$$

The prior factorization of the  $n$ -by- $n$  matrix  $\mathbf{C}$  is given as

$$\begin{aligned} \mathbf{LDL}^\top &= \begin{bmatrix} \mathbf{L}_{11} & & \\ \mathbf{0}^\top & 1 & \\ \mathbf{L}_{31} & \mathbf{0} & \mathbf{L}_{33} \end{bmatrix} \begin{bmatrix} \mathbf{D}_1 & & \\ & d_2 & \\ & & \mathbf{D}_3 \end{bmatrix} \begin{bmatrix} \mathbf{L}_{11}^\top & \mathbf{0} & \mathbf{L}_{31}^\top \\ & 1 & \mathbf{0}^\top \\ & & \mathbf{L}_{33}^\top \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{C}_{11} & \mathbf{0} & \mathbf{C}_{31}^\top \\ \mathbf{0}^\top & \alpha & \mathbf{0}^\top \\ \mathbf{C}_{31} & \mathbf{0} & \mathbf{C}_{33} \end{bmatrix}. \end{aligned}$$

which leads to the four equations,

$$\begin{aligned} \mathbf{L}_{11}\mathbf{D}_1\mathbf{L}_{11}^\top &= \mathbf{C}_{11}, \\ d_2 &= \alpha, \\ \mathbf{L}_{31}\mathbf{D}_1\mathbf{L}_{11}^\top &= \mathbf{C}_{31}, \\ \mathbf{L}_{31}\mathbf{D}_1\mathbf{L}_{31}^\top + \mathbf{L}_{33}\mathbf{D}_3\mathbf{L}_{33}^\top &= \mathbf{C}_{33}. \end{aligned} \tag{2.1}$$

After adding row and column  $k$  to obtain  $\bar{\mathbf{C}}$ , we have the factorization

$$\begin{aligned} \bar{\mathbf{LDL}}^\top &= \begin{bmatrix} \mathbf{L}_{11} & & \\ \bar{\mathbf{l}}_{12}^\top & 1 & \\ \mathbf{L}_{31} & \bar{\mathbf{l}}_{32} & \bar{\mathbf{L}}_{33} \end{bmatrix} \begin{bmatrix} \mathbf{D}_1 & & \\ & \bar{d}_2 & \\ & & \bar{\mathbf{D}}_3 \end{bmatrix} \begin{bmatrix} \mathbf{L}_{11}^\top & \bar{\mathbf{l}}_{12} & \mathbf{L}_{31}^\top \\ & 1 & \bar{\mathbf{l}}_{32}^\top \\ & & \bar{\mathbf{L}}_{33}^\top \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{C}_{11} & \bar{\mathbf{c}}_{12} & \mathbf{C}_{31}^\top \\ \bar{\mathbf{c}}_{12}^\top & \bar{c}_{22} & \bar{\mathbf{c}}_{32}^\top \\ \mathbf{C}_{31} & \bar{\mathbf{c}}_{32} & \mathbf{C}_{33} \end{bmatrix}. \end{aligned} \tag{2.2}$$

Note that  $\mathbf{L}_{11}$  and  $\mathbf{L}_{31}$  do not change as a result of modifying row and column  $k$  of  $\mathbf{C}$ . From (2.2), the relevant equations are

$$\mathbf{L}_{11}\mathbf{D}_1\bar{\mathbf{l}}_{12} = \bar{\mathbf{c}}_{12}, \tag{2.3}$$

$$\begin{aligned}
& \bar{\mathbf{I}}_{12}^{\top} \mathbf{D}_1 \bar{\mathbf{I}}_{12} + \bar{d}_2 = \bar{c}_{22}, \\
& \mathbf{L}_{31} \mathbf{D}_1 \bar{\mathbf{I}}_{12} + \bar{\mathbf{I}}_{32} \bar{d}_2 = \bar{c}_{32}, \\
(2.4) \quad & \mathbf{L}_{31} \mathbf{D}_1 \mathbf{L}_{31}^{\top} + \bar{\mathbf{I}}_{32} \bar{d}_2 \bar{\mathbf{I}}_{32}^{\top} + \bar{\mathbf{L}}_{33} \bar{\mathbf{D}}_3 \bar{\mathbf{L}}_{33}^{\top} = \mathbf{C}_{33}.
\end{aligned}$$

Let  $\mathbf{w} = \bar{\mathbf{I}}_{32} \sqrt{\bar{d}_2}$ . Combining (2.4) with the original equation (2.1), we obtain

$$\bar{\mathbf{L}}_{33} \bar{\mathbf{D}}_3 \bar{\mathbf{L}}_{33}^{\top} = (\mathbf{C}_{33} - \mathbf{L}_{31} \mathbf{D}_1 \mathbf{L}_{31}^{\top}) - \bar{\mathbf{I}}_{32} \bar{d}_2 \bar{\mathbf{I}}_{32}^{\top} = \mathbf{L}_{33} \mathbf{D}_3 \mathbf{L}_{33}^{\top} - \mathbf{w} \mathbf{w}^{\top}.$$

The factorization of  $\mathbf{L}_{33} \mathbf{D}_3 \mathbf{L}_{33}^{\top} - \mathbf{w} \mathbf{w}^{\top}$  can be computed as a rank-1 downdate of the original factorization  $\mathbf{L}_{33} \mathbf{D}_3 \mathbf{L}_{33}^{\top}$  [3]. This derivation leads to Algorithm 1 for computing the modified factorization  $\bar{\mathbf{L}} \bar{\mathbf{D}} \bar{\mathbf{L}}^{\top}$ :

**Algorithm 1 (Row addition)**

1. Solve the lower triangular system  $\mathbf{L}_{11} \mathbf{D}_1 \bar{\mathbf{I}}_{12} = \bar{c}_{12}$  for  $\bar{\mathbf{I}}_{12}$ .
2.  $\bar{d}_2 = \bar{c}_{22} - \bar{\mathbf{I}}_{12}^{\top} \mathbf{D}_1 \bar{\mathbf{I}}_{12}$
3.  $\bar{\mathbf{I}}_{32} = (\bar{c}_{32} - \mathbf{L}_{31} \mathbf{D}_1 \bar{\mathbf{I}}_{12}) / \bar{d}_2$
4.  $\mathbf{w} = \bar{\mathbf{I}}_{32} \sqrt{\bar{d}_2}$
5. Perform the rank-1 downdate  $\bar{\mathbf{L}}_{33} \bar{\mathbf{D}}_3 \bar{\mathbf{L}}_{33}^{\top} = \mathbf{L}_{33} \mathbf{D}_3 \mathbf{L}_{33}^{\top} - \mathbf{w} \mathbf{w}^{\top}$ .

**end Algorithm 1**

Step (1) of Algorithm 1 requires the solution of a unit lower triangular system  $\mathbf{L}_{11} \mathbf{y} = \bar{c}_{12}$  of order  $k-1$ . The computation of  $\mathbf{y}$  takes  $(k-1)^2 - (k-1)$  flops, and computing  $\bar{\mathbf{I}}_{12} = \mathbf{D}_1^{-1} \mathbf{y}$  takes another  $k-1$  flops. Step (2) requires  $2(k-1)$  work, using  $\mathbf{y}$ . Step (3) is the matrix-vector multiply  $\mathbf{L}_{31} \mathbf{y}$ , where  $\mathbf{L}_{31}$  is  $(n-k)$ -by- $(k-1)$ , and thus takes  $2(n-k)(k-1)$  operations, and  $k-1$  more to divide by  $\bar{d}_2$ . Step (4) requires one square root operation and  $n-k$  multiplications. Finally, the rank-1 downdate of step (5) takes  $2(n-k)^2 + 4(n-k)$  operations using method C1 of [11] (see [4]). For the dense case, the total number of floating-point operations performed by Algorithm 1 is  $2n^2 + 3n + k^2 - (2nk + 2k + 1)$ . This is roughly  $2n^2$  when  $k=1$ ,  $n^2$  when  $k=n$ , and  $(5/4)n^2$  when  $k=n/2$ .

**3. Sparse row addition.** If  $\mathbf{C}$  is sparse, each step of Algorithm 1 must operate on sparse matrices. The graph algorithms and data structures must efficiently support each step. We will assume that  $\mathbf{L}$  is stored in a compressed column vector form, where the row indices in each column are sorted in ascending order. This is the same data structure used in [3, 4], except that the algorithms presented there do not require sorted row indices, but they do require the integer *multiplicity* of each nonzero entry of  $\mathbf{L}$  to support an efficient symbolic downdate operation. The algorithm discussed below will not require the multiplicities.

Maintaining the row indices in sorted order requires a merge operation for the set union computation to determine the new nonzero patterns of the columns of  $\mathbf{L}$ , rather than a simpler unsorted set union used in [3, 4]. It has no effect on asymptotic complexity, and little effect on the run time. Although more work is required to maintain the row indices in sorted order, time is gained elsewhere in the algorithm. Operating on columns in sorted order in the forward solve of  $\mathbf{L} \mathbf{x} = \mathbf{b}$ , for example, is faster than operating on a matrix with jumbled columns.

Step (1) of Algorithm 1 solves the lower triangular system  $\mathbf{L}_{11} \mathbf{y} = \bar{c}_{12}$ , where all three terms in this system are sparse. Gilbert and Peierls have shown how to solve

this system optimally, in time proportional to the number of floating-point operations required [9, 10]. We review their method here.

Consider the  $n$ -by- $n$  system  $\mathbf{L}\mathbf{x} = \mathbf{b}$ , where  $\mathbf{L}$  is lower triangular and both  $\mathbf{L}$  and  $\mathbf{b}$  are sparse. The solution will be sparse, and the total work may be less than  $O(n)$ . We cannot use a conventional algorithm that iterates over each column of  $\mathbf{L}$  and skips those for which  $\mathbf{x}$  is zero, since the work involved will then be at least  $n$ . Instead, the nonzero pattern of  $\mathbf{x}$  must first be computed, and then the corresponding columns of  $\mathbf{L}$  can be used to compute  $\mathbf{x}$  in time proportional to the floating-point work required.

Let  $G_L$  be a graph with  $n$  nodes and with a directed edge from node  $j$  to node  $i$  if and only if  $l_{ij}$  is nonzero. Gilbert and Peierls show that  $x_j$  is nonzero (ignoring numerical cancellation) if and only if there is a path of length zero or more from some node  $i$ , where  $b_i \neq 0$ , to node  $j$  in the graph  $G_L$ . Computing the pattern of  $\mathbf{x}$  requires a graph traversal, starting from the nodes corresponding to the nonzero pattern of  $\mathbf{b}$ . It can be done in time proportional to the number of edges traversed. Each of these edges is a nonzero in  $\mathbf{L}$  that takes part in the subsequent numerical computation.

The above result holds for any lower triangular matrix  $\mathbf{L}$ . In our case,  $\mathbf{L}$  arises from a Cholesky factorization, and has an *elimination tree* [18, 19]. The elimination tree of  $\mathbf{L}$  has  $n$  nodes. The parent of node  $j$  in the elimination tree is the smallest index  $i > j$  for which  $l_{ij} \neq 0$ ; node  $j$  is a root if there is no such  $i$ . Since the nonzero pattern of  $(\mathbf{L})_{*j}$  is a subset of its path to the root of the elimination tree [20], all the nodes in  $G_L$  that can be reached from node  $j$  correspond to the path from node  $j$  to the root of the elimination tree. Traversing the paths in the tree, starting at nodes corresponding to nonzero entries in  $\mathbf{b}$ , takes time proportional to the number of nonzero entries in  $\mathbf{x}$ . A general graph traversal of  $G_L$  is not required. Step (1) of Algorithm 1 takes

$$O\left(\sum_{(\bar{\mathbf{I}}_{12})_j \neq 0} |(\mathbf{L}_{11})_{*j}|\right)$$

time to compute the nonzero pattern and numerical values of both  $\mathbf{y}$  and  $\bar{\mathbf{I}}_{12}$ . The insertion of the nonzero entries of  $\bar{\mathbf{I}}_{12}$  into the data structure of  $\mathbf{L}$  is performed in conjunction with step (3).

Step (2) is a scaled dot product operation, and can be computed in time proportional to the number of nonzero entries in  $\bar{\mathbf{I}}_{12}$ .

Step (3) is a matrix-vector multiply operation. It accesses the same columns of  $\mathbf{L}$  used by the sparse lower triangular solve, namely, each column  $j$  for which the  $j$ th entry in  $\bar{\mathbf{I}}_{12}$  is nonzero. These same columns need to be modified, by shifting entries in  $\mathbf{L}_{31}$  down by one, and inserting the new entries in  $\bar{\mathbf{I}}_{12}$ , the  $k$ th row of  $\bar{\mathbf{L}}$ . No other columns in the range 1 to  $k - 1$  need to be accessed or modified by steps (1) through (3). When step (3) completes, the new column  $k$  of  $\bar{\mathbf{L}}$  needs to be inserted into the data structure. This can be done in one of two ways. In the general case, we can store the columns themselves in a non-contiguous manner, and simply allocate new space for this column. A similar strategy can be used for any columns 1 through  $k - 1$  of  $\bar{\mathbf{L}}$  that outgrow their originally allocated space, with no increase in asymptotic run time. Alternatively, we may know an a priori upper bound on the size of each column of  $\mathbf{L}$  after all row additions have been performed. In this case, a simpler static allocation strategy is possible. This latter case occurs in our use of the row addition algorithm in LPDASA, our target application [5]. In either case, the time to insert  $\bar{\mathbf{I}}_{12}$  into the

data structure and to compute  $\bar{\mathbf{I}}_{32}$  in step (3) is

$$O\left(\sum_{(\bar{\mathbf{I}}_{12})_j \neq 0} |(\mathbf{L}_{31})_{*j}| \right).$$

Step (4) is a simple scalar-times-vector operation. The total time for steps (1) through (4) is

$$O\left(\sum_{(\bar{\mathbf{I}}_{12})_j \neq 0} |(\mathbf{L})_{*j}| \right).$$

Step (5) almost fits the specifications of the sparse rank-1 modification in [3], but with one interesting twist. The original  $k$ th row and column of  $\mathbf{C}$  are zero, except for the placeholder diagonal entry,  $\alpha$ . The new row and column only adds entries to  $\mathbf{C}$ , and thus the nonzero pattern of the original factor  $\mathbf{L}$  is a subset of the nonzero pattern of  $\bar{\mathbf{L}}$  (ignoring numerical cancellation). The rank-1 modification in Algorithm 1 is a symbolic update (new nonzero entries are added, not removed) and a numeric downdate  $\mathbf{L}_{33}\mathbf{D}_3\mathbf{L}_{33}^\top - \mathbf{w}\mathbf{w}^\top$ . Since the multiplicities used in [3] are needed only for a subsequent symbolic downdate, they are not required by the row addition algorithm. They would be required by a row deletion algorithm that maintains a strict nonzero pattern of  $\mathbf{L}$ ; this issue is addressed in Section 5.

The rank-1 modification to obtain the factorization  $\bar{\mathbf{L}}_{33}\bar{\mathbf{D}}_3\bar{\mathbf{L}}_{33}^\top$  takes time proportional to the number of nonzero entries in  $\bar{\mathbf{L}}_{33}$  that change. The columns that change correspond to the path from node  $k$  to the root of the elimination tree of  $\bar{\mathbf{L}}$ . This path is denoted  $\bar{\mathcal{P}}$  in [3]. At each node  $j$  along the path  $\bar{\mathcal{P}}$ , at most four floating-point operations are performed for each nonzero entry in  $(\bar{\mathbf{L}})_{*j}$ .

With our choice of data structures, exploitation of the elimination tree, and the rank-1 modification from [3], the total time taken by Algorithm 1 is proportional to the total number of nonzero entries in columns corresponding to nonzero entries in  $\bar{\mathbf{I}}_{12}$ , to compute steps (1) through (4), plus the time required for the rank-1 modification in step (5). The total time is

$$O\left(\sum_{(\bar{\mathbf{I}}_{12})_j \neq 0} |(\mathbf{L})_{*j}| + \sum_{j \in \bar{\mathcal{P}}} |(\bar{\mathbf{L}})_{*j}| \right)$$

This time includes all data structure manipulations, sparsity pattern computation, and graph algorithms required to implement the algorithm. It is identical to the total number of floating-point operations required, and thus Algorithm 1 is optimal. In the sparse case, if every column  $j$  takes part in the computation, the time is  $O(|\bar{\mathbf{L}}|)$ . Normally, not all columns will be affected by the sparse row addition. If the new row and column of  $\bar{\mathbf{L}}$  are very sparse, only a few columns take part in the computation.

**4. Row deletion.** By *deleting* a row and column  $k$  from the matrix  $\mathbf{C}$ , we mean setting the entire row and column to zero, except for the diagonal entry  $(\mathbf{C})_{kk}$  which is set to  $\alpha$ . Prior to deleting row and column  $k$ , we have the original factorization

$$\mathbf{L}\mathbf{D}\mathbf{L}^\top = \begin{bmatrix} \mathbf{L}_{11} & & & \\ \mathbf{l}_{12}^\top & 1 & & \\ \mathbf{L}_{31} & \mathbf{l}_{32} & \mathbf{L}_{33} & \end{bmatrix} \begin{bmatrix} \mathbf{D}_1 & & & \\ & d_2 & & \\ & & \mathbf{D}_3 & \end{bmatrix} \begin{bmatrix} \mathbf{L}_{11}^\top & \mathbf{l}_{12} & \mathbf{L}_{31}^\top \\ & 1 & \mathbf{l}_{32}^\top \\ & & \mathbf{L}_{33}^\top \end{bmatrix}$$

$$= \begin{bmatrix} \mathbf{C}_{11} & \mathbf{c}_{12} & \mathbf{C}_{31}^\top \\ \mathbf{c}_{12}^\top & c_{22} & \mathbf{c}_{32}^\top \\ \mathbf{C}_{31} & \mathbf{c}_{32} & \mathbf{C}_{33} \end{bmatrix}$$

After deleting row and column  $k$ , we have

$$\begin{aligned} \overline{\mathbf{LDL}}^\top &= \begin{bmatrix} \mathbf{L}_{11} & & \\ \mathbf{0}^\top & 1 & \\ \mathbf{L}_{31} & \mathbf{0} & \overline{\mathbf{L}}_{33} \end{bmatrix} \begin{bmatrix} \mathbf{D}_1 & & \\ & \alpha & \\ & & \overline{\mathbf{D}}_3 \end{bmatrix} \begin{bmatrix} \mathbf{L}_{11}^\top & \mathbf{0} & \mathbf{L}_{31}^\top \\ & 1 & \mathbf{0}^\top \\ & & \overline{\mathbf{L}}_{33}^\top \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{C}_{11} & \mathbf{0} & \mathbf{C}_{31}^\top \\ \mathbf{0}^\top & \alpha & \mathbf{0}^\top \\ \mathbf{C}_{31} & \mathbf{0} & \mathbf{C}_{33} \end{bmatrix} \end{aligned}$$

Thus we only need to set row and column  $k$  of  $\overline{\mathbf{L}}$  to zero, set the diagonal entry to  $\alpha$ , and compute  $\overline{\mathbf{L}}_{33}$  and  $\overline{\mathbf{D}}_3$ . The original factorization is

$$\mathbf{L}_{33}\mathbf{D}_3\mathbf{L}_{33}^\top = \mathbf{C}_{33} - \mathbf{L}_{31}\mathbf{D}_1\mathbf{L}_{31}^\top - \mathbf{l}_{32}d_2\mathbf{l}_{32}^\top,$$

while the new factorization is given as

$$\overline{\mathbf{L}}_{33}\overline{\mathbf{D}}_3\overline{\mathbf{L}}_{33}^\top = \mathbf{C}_{33} - \mathbf{L}_{31}\mathbf{D}_1\mathbf{L}_{31}^\top.$$

Combining these two equations, we have a numeric rank-1 update,

$$(4.1) \quad \overline{\mathbf{L}}_{33}\overline{\mathbf{D}}_3\overline{\mathbf{L}}_{33}^\top = \mathbf{L}_{33}\mathbf{D}_3\mathbf{L}_{33}^\top + \mathbf{w}\mathbf{w}^\top,$$

where  $\mathbf{w} = \mathbf{l}_{32}\sqrt{d_2}$ . Algorithm 2 gives the complete row deletion algorithm:

**Algorithm 2 (Row deletion)**

1.  $\overline{\mathbf{l}}_{12} = \mathbf{0}$
2.  $\overline{d}_2 = \alpha$
3.  $\overline{\mathbf{l}}_{32} = \mathbf{0}$
4.  $\mathbf{w} = \mathbf{l}_{32}\sqrt{d_2}$
5. Perform the rank-1 update  $\overline{\mathbf{L}}_{33}\overline{\mathbf{D}}_3\overline{\mathbf{L}}_{33}^\top = \mathbf{L}_{33}\mathbf{D}_3\mathbf{L}_{33}^\top + \mathbf{w}\mathbf{w}^\top$ .

**end Algorithm 2**

When  $\mathbf{C}$  is dense, the number of floating-point operations performed by Algorithm 2 is  $2(n-k)^2 + 5(n-k) + 1$  (the same as steps (4) and (5) of Algorithm 1). This is roughly  $2n^2$  when  $k = 1$ , and  $(1/2)n^2$  when  $k = n/2$ . No work is required when  $k = n$ .

**5. Sparse row deletion.** When row and column  $k$  of  $\mathbf{C}$  are “deleted” (set to zero) to obtain  $\overline{\mathbf{C}}$ , no new nonzero terms will appear in  $\overline{\mathbf{L}}$ , and some nonzero entries in  $\mathbf{L}$  may become zero. We refer to the deletion of entries in  $\mathbf{L}$  as a symbolic downdate [3]. The symbolic downdate is combined with a numeric rank-1 update, because of the addition of  $\mathbf{w}\mathbf{w}^\top$  in (4.1).

We cannot simply delete entries from  $\overline{\mathbf{L}}$  that become numerically zero. An entry in  $\overline{\mathbf{L}}$  can only be removed if it becomes *symbolically zero* (that is, its value is zero regardless of the assignment of numerical values to the nonzero pattern of  $\mathbf{C}$ ). If entries are zero because of exact numerical cancellation, and are dropped from the data structure of  $\mathbf{L}$ , then the elimination tree no longer characterizes the structure

of the matrix. If the elimination tree is no longer valid, subsequent updates and downdates will not be able to determine which columns of  $\mathbf{L}$  must be modified.

Steps (1) through (3) of Algorithm 2 require no numerical work, but do require some data structure modifications. All of the entries in row and column  $k$  of  $\bar{\mathbf{L}}$  become symbolically zero and can be removed. If  $\mathbf{L}$  is stored by columns, it is trivial in step (3) to immediately delete all entries in column  $\mathbf{l}_{32}$ . On the other hand, the statement  $\bar{\mathbf{l}}_{12} = \mathbf{0}$  in step (1) is less obvious. Each nonzero element in row  $k$  lies in a different column. We must either set these values to zero, delete them from the data structure, or flag row  $k$  as zero and require any subsequent algorithm that accesses the matrix  $\mathbf{L}$  to ignore flagged rows. The latter option would lead to a sparse row deletion algorithm with optimal run time, but complicates all other algorithms in our application and increases their run time. We choose to search for them in each column that they appear and delete them from the data structure.

To set  $\mathbf{l}_{12}$  to zero and delete the entries from the data structure for  $\mathbf{L}$  requires a scan of all the columns  $j$  of  $\mathbf{L}$  for which the  $j$ th entry of  $\mathbf{l}_{12}$  is nonzero. The time taken for this operation is asymptotically bounded by the time taken for steps (1) through (3) of sparse row addition (Algorithm 1), but the bound is not tight. The nonzero pattern of row  $k$  of  $\mathbf{L}$  can easily be found from the elimination tree. Finding the row index  $k$  in the columns takes less time than step (1) of Algorithm 1 since a binary search can be used. Deleting the entries takes time equivalent to step (3) of Algorithm 1, since we maintain each column with sorted row indices.

The immediate removal of entries in  $\bar{\mathbf{L}}_{33}$  can be done using the symbolic rank-1 downdate presented in [3]. However, this requires an additional array of multiplicities, which is one additional integer value for each nonzero in the matrix. Instead, we can allow these entries to become numerically zero (or very small values due to numerical roundoff) and not remove them immediately. Since they become numerically zero (or tiny), they can simply remain in the data structure for the matrix, and have no effect on subsequent operations that use the matrix  $\mathbf{L}$ . The entries can be pruned later on by a complete symbolic factorization, taking  $O(|\mathbf{L}|)$  time [6, 7, 8]. If this is done rarely, the overall run time of the application that uses the sparse row deletion algorithm will not be affected adversely.

The asymptotic run time of our sparse row deletion algorithm is the same as sparse row addition (or less, because of the binary search), even though sparse row addition requires more numerical work. This is non-optimal, but no worse than sparse row addition, whose run time is optimal.

**6. Arbitrary row modification.** It is possible to generalize our row deletion and addition algorithms to handle the case where the  $k$ th row and column of  $\mathbf{C}$  is neither originally zero (the row addition case) nor set to zero (the row deletion) case, but is changed arbitrarily. Any change of this form can be handled as a row deletion followed by a row addition, but the question may remain as to whether or not it can be done faster as a single step. In this section we show that no floating-point operations are saved in a single-pass row modification algorithm, as compared to the row deletion + row addition approach, if the change in the  $k$ th row and column is arbitrary.

The original matrix  $\mathbf{C}$  and the new matrix  $\bar{\mathbf{C}}$  are

$$(6.1) \quad \mathbf{C} = \begin{bmatrix} \mathbf{C}_{11} & \mathbf{c}_{12} & \mathbf{C}_{31}^T \\ \mathbf{c}_{12}^T & c_{22} & \mathbf{c}_{32}^T \\ \mathbf{C}_{31} & \mathbf{c}_{32} & \mathbf{C}_{33} \end{bmatrix} \quad \text{and} \quad \bar{\mathbf{C}} = \begin{bmatrix} \mathbf{C}_{11} & \bar{\mathbf{c}}_{12} & \mathbf{C}_{31}^T \\ \bar{\mathbf{c}}_{12}^T & \bar{c}_{22} & \bar{\mathbf{c}}_{32}^T \\ \mathbf{C}_{31} & \bar{\mathbf{c}}_{32} & \mathbf{C}_{33} \end{bmatrix}.$$



Computing  $\bar{\mathbf{l}}_{12}$ ,  $\bar{d}_2$ , and  $\bar{\mathbf{l}}_{32}$  is the same as the row addition algorithm, and takes exactly the same number of floating-point operations. The original factorization of the (33)-block is

$$\mathbf{L}_{33}\mathbf{D}_3\mathbf{L}_{33}^\top = \mathbf{C}_{33} - \mathbf{L}_{31}\mathbf{D}_1\mathbf{L}_{31}^\top - \mathbf{l}_{32}d_2\mathbf{l}_{32}^\top,$$

while the new factorization is

$$\bar{\mathbf{L}}_{33}\bar{\mathbf{D}}_3\bar{\mathbf{L}}_{33}^\top = \mathbf{C}_{33} - \mathbf{L}_{31}\mathbf{D}_1\mathbf{L}_{31}^\top - \bar{\mathbf{l}}_{32}\bar{d}_2\bar{\mathbf{l}}_{32}^\top.$$

These can be combined into a rank-1 update and rank-1 downdate

$$(6.2) \quad \bar{\mathbf{L}}_{33}\bar{\mathbf{D}}_3\bar{\mathbf{L}}_{33}^\top = \mathbf{L}_{33}\mathbf{D}_3\mathbf{L}_{33}^\top + \mathbf{w}_1\mathbf{w}_1^\top - \mathbf{w}_2\mathbf{w}_2^\top,$$

where  $\mathbf{w}_1 = \mathbf{l}_{32}\sqrt{d_2}$  and  $\mathbf{w}_2 = \bar{\mathbf{l}}_{32}\sqrt{\bar{d}_2}$ . The multiple rank update/downdate presented in [4] cannot perform a simultaneous update and downdate, but if the sparse downdate (removal of entries that become symbolically zero) is not performed, it would be possible to compute the simultaneous update and downdate (6.2) in a single pass. This may result in some time savings, since the data structure for  $\mathbf{L}$  is scanned once, not twice. No floating-point work would be saved, however. The total floating-point operation count of the row modification algorithm is identical to a row deletion followed by a row addition.

**7. Making sparse changes to a sparse row.** Suppose we constrain the modifications to the  $k$ th row and column of  $\mathbf{C}$ . This assumption can reduce the total amount of floating-point work required to compute the modified factorization  $\bar{\mathbf{L}}\bar{\mathbf{D}}\bar{\mathbf{L}}^\top$ , as shown below.

More precisely, consider the case where only a few entries in row and column  $k$  of  $\mathbf{C}$  are changed. Let

$$\begin{aligned} \Delta\mathbf{c}_{12} &= \bar{\mathbf{c}}_{12} - \mathbf{c}_{12}, \\ \Delta c_{22} &= \bar{c}_{22} - c_{22}, \\ \Delta\mathbf{c}_{32} &= \bar{\mathbf{c}}_{32} - \mathbf{c}_{32}, \end{aligned}$$

and assume that the change in the  $k$ th row and column is much sparser than the original  $k$ th row and column of  $\mathbf{C}$  (for example,  $|\Delta\mathbf{c}_{12}| \ll |\bar{\mathbf{c}}_{12}|$ ).

If we consider (2.3) and its analog for the original matrix  $\mathbf{C}$ , we have

$$(7.1) \quad \mathbf{L}_{11}\mathbf{D}_1\mathbf{l}_{12} = \mathbf{c}_{12},$$

$$(7.2) \quad \mathbf{L}_{11}\mathbf{D}_1\bar{\mathbf{l}}_{12} = \bar{\mathbf{c}}_{12}.$$

Combining these two equations together, we have

$$\mathbf{L}_{11}\mathbf{D}_1\Delta\mathbf{l}_{12} = \Delta\mathbf{c}_{12},$$

where  $\Delta\mathbf{l}_{12} = \bar{\mathbf{l}}_{12} - \mathbf{l}_{12}$ . Since  $\Delta\mathbf{c}_{12}$  is sparse, the solution to this lower triangular system will also likely be sparse. It can be solved in time proportional to the time required to multiply  $\mathbf{L}_{11}$  times  $\Delta\mathbf{l}_{12}$ , or

$$O\left(\sum_{(\Delta\mathbf{l}_{12})_j \neq 0} |(\mathbf{L}_{11})_{*j}|\right)$$

[9, 10]. We can then compute  $\bar{\mathbf{l}}_{12} = \mathbf{l}_{12} + \Delta\mathbf{l}_{12}$ . This approach for computing  $\bar{\mathbf{l}}_{12}$  can be much faster than solving (7.2) directly, which would take

$$O\left(\sum_{(\bar{\mathbf{l}}_{12})_j \neq 0} |(\mathbf{L}_{11})_{*j}|\right)$$

time. Computing  $\bar{d}_2$  can be done in time proportional to  $|\Delta\mathbf{l}_{12}|$ , using the following formula that modifies the dot product computation in Algorithm 1,

$$\bar{d}_2 = d_2 + \Delta d_2 = d_2 + \Delta c_{22} - \sum_{(\Delta\mathbf{l}_{12})_j \neq 0} (\Delta\mathbf{l}_{12})_j ((\mathbf{l}_{12})_j + (\bar{\mathbf{l}}_{12})_j) (\mathbf{D}_1)_{jj}.$$

Similarly, the  $k$ th column of  $\bar{\mathbf{L}}$  can be computed as

$$\bar{\mathbf{l}}_{32} = (\Delta\mathbf{C}_{32} + \mathbf{l}_{32}d_2 - \mathbf{L}_{31}\mathbf{D}_1\Delta\mathbf{l}_{12})/\bar{d}_2$$

The key component in this computation is the sparse-matrix-vector multiplication  $\mathbf{L}_{31}\mathbf{D}_1\Delta\mathbf{l}_{12}$ , which takes less time to compute than the corresponding computation  $\mathbf{L}_{31}\mathbf{D}_1\mathbf{l}_{12}$  in step (2) of Algorithm 1.

The remaining work for the rank-1 update/downdate of  $\mathbf{L}_{33}\mathbf{D}_{33}\mathbf{L}_{33}^\top$  is identical to the arbitrary row modification (6.2). If  $k$  is small, it is likely that this update/downdate computation will take much more time than the computation of  $\bar{\mathbf{l}}_{12}$ ,  $d_2$ , and  $\bar{\mathbf{l}}_{32}$ . If  $k$  is large, however, a significant reduction in the total amount of work can be obtained by exploiting sparsity in the change in the row and column of  $\mathbf{C}$ .

**8. Row modifications as a rank-2 outer-product.** Modifying row and column  $k$  of a symmetric matrix  $\mathbf{C}$  can be written as a rank-2 modification  $\bar{\mathbf{C}} = \mathbf{C} + \mathbf{w}_1\mathbf{w}_1^\top - \mathbf{w}_2\mathbf{w}_2^\top$ . Suppose  $\mathbf{C}$  and  $\bar{\mathbf{C}}$  are as given in (6.1). Let

$$\mathbf{d} = \begin{bmatrix} \bar{c}_{12} - \mathbf{c}_{12} \\ (\bar{c}_{22} - c_{22})/2 \\ \bar{c}_{32} - \mathbf{c}_{32} \end{bmatrix}.$$

Let  $\mathbf{e}_k$  be the  $k$ th column of the identity. Then  $\bar{\mathbf{C}} = \mathbf{C} + \mathbf{d}\mathbf{e}_k^\top + \mathbf{e}_k\mathbf{d}^\top$ . This can be put into the form  $\bar{\mathbf{C}} = \mathbf{C} + \mathbf{w}_1\mathbf{w}_1^\top - \mathbf{w}_2\mathbf{w}_2^\top$  using the following relationship (note that  $\mathbf{e}$ , below, is an arbitrary column vector, not necessarily  $\mathbf{e}_k$ ). Given  $\mathbf{d}$  and  $\mathbf{e} \in \mathbb{R}^n$ , define

$$\mathbf{p} = \frac{\mathbf{d}}{\|\mathbf{d}\|} + \frac{\mathbf{e}}{\|\mathbf{e}\|} \quad \text{and} \quad \mathbf{q} = \frac{\mathbf{d}}{\|\mathbf{d}\|} - \frac{\mathbf{e}}{\|\mathbf{e}\|}.$$

Then we have

$$(8.1) \quad \mathbf{d}\mathbf{e}^\top + \mathbf{e}\mathbf{d}^\top = \frac{\|\mathbf{d}\|\|\mathbf{e}\|}{2} (\mathbf{p}\mathbf{p}^\top - \mathbf{q}\mathbf{q}^\top).$$

In our case,  $\mathbf{e} = \mathbf{e}_k$  and  $\|\mathbf{e}\| = 1$ . Defining

$$\mathbf{w}_1 = \sqrt{\frac{\|\mathbf{d}\|}{2}} \left( \frac{\mathbf{d}}{\|\mathbf{d}\|} + \mathbf{e}_k \right) \quad \text{and} \quad \mathbf{w}_2 = \sqrt{\frac{\|\mathbf{d}\|}{2}} \left( \frac{\mathbf{d}}{\|\mathbf{d}\|} - \mathbf{e}_k \right),$$

it follows from (8.1)

$$(8.2) \quad \bar{\mathbf{C}} = \mathbf{C} + \mathbf{w}_1\mathbf{w}_1^\top - \mathbf{w}_2\mathbf{w}_2^\top.$$

In the dense case, computing (8.2) using a rank-1 update and rank-1 downdate takes  $4n^2 + 11n$  floating-point operations, independent of  $k$  (including the work to compute  $\mathbf{w}_1$  and  $\mathbf{w}_2$ ). If we use Algorithm 1 and 2 to make an arbitrary change in the  $k$ th row and column of  $\mathbf{C}$ , the total work is  $4n^2 + 8n + 3k^2 - (6nk + 7k)$ , which is roughly  $4n^2$  when  $k = 1$ ,  $n^2 + n$  when  $k = n$ , and  $(7/4)n^2$  when  $k = n/2$ . Using the rank-2 update/downdate method to modify row and column  $k$  of  $\mathbf{C}$  can thus take up to four times the work compared to the row addition/deletion presented here.

If the modification requires only the addition or deletion of row and column  $k$ , then only Algorithm 1 or 2 needs to be used, but the entire rank-2 update/downdate method presented in this section is still required (about  $4n^2$  work, independent of  $k$ ). Considering only the quadratic terms, Algorithm 1 performs  $2n^2 + k^2 - 2nk$  operations, and Algorithm 2 performs  $2(n - k)^2$  operations. The row modification methods require much less work, particularly for the row deletion case when  $k \approx n$ , in which the work drops from  $4n^2$  for the rank-2 update/downdate method to nearly no work at all.

In the sparse case, the differences in work and memory usage can be extreme. Both the rank-1 update and downdate could affect the entire matrix  $\mathbf{L}$ , and could cause catastrophic intermediate fill-in. Consider the row addition case when  $k \approx n$ , and the new row and column of  $\mathbf{C}$  is dense. The factorization of  $\bar{\mathbf{C}}$  will still be fairly sparse, since the large submatrix  $\mathbf{C}_{11}$  does not change, and remains very sparse. The factor  $\mathbf{L}_{11}$  can have as few as  $O(n)$  entries. However, after one rank-1 update, the  $(11)$ -block of  $\mathbf{C} + \mathbf{w}_1\mathbf{w}_1^T$  is completely dense, since  $\mathbf{w}_1$  is a dense column vector. After the rank-1 downdate (with  $-\mathbf{w}_2\mathbf{w}_2^T$ ), massive cancellation occurs, and the factorization of  $\mathbf{C}_{11}$  is restored to its original sparse nonzero pattern. But the damage has been done, since we require  $O(n^2)$  memory to hold the intermediate factorization. This is infeasible in a sparse matrix algorithm. The memory problem could be solved if a single-pass rank-2 update/downdate algorithm were to be used, but even then, the total work required would be  $O(n^2)$  which is much more than the  $O(|\bar{\mathbf{L}}|)$  time required for Algorithm 1 in this case. The same problem occurs if the downdate with  $-\mathbf{w}_2\mathbf{w}_2^T$  is applied first.

**9. Modifying a lower triangular system.** We now consider a related operation that can be performed efficiently at the same time that we modify the Cholesky factorization. Suppose we have a linear system  $\mathbf{C}\mathbf{x} = \mathbf{b}$ , and the system is modified, either by changing a row and column of  $\mathbf{C}$  as discussed above, or due to a low-rank change  $\bar{\mathbf{C}} = \mathbf{C} \pm \mathbf{W}\mathbf{W}^T$  as discussed in [3, 4]. After the factorization is modified, the new factorization  $\bar{\mathbf{L}}\bar{\mathbf{D}}\bar{\mathbf{L}}^T = \bar{\mathbf{C}}$  will normally be used to solve a modified linear system,  $\bar{\mathbf{C}}\bar{\mathbf{x}} = \bar{\mathbf{b}}$ . The complete solution  $\bar{\mathbf{x}}$  will likely be different in every component because of the backsolve, but a significant reduction of work can be obtained in the forward solve of the lower triangular system  $\bar{\mathbf{L}}\bar{\mathbf{x}} = \bar{\mathbf{b}}$ . We thus focus only on this lower triangular system, not the complete system.

First, consider the simpler case of a low-rank change of a dense matrix. As shown below, it takes double the work to modify the solution instead of computing it from scratch, but the dense matrix method can be used for submatrices in the sparse case, resulting in a significant reduction in work. Suppose we have the system  $\mathbf{L}\mathbf{x} = \mathbf{b}$ , including its solution  $\mathbf{x}$ , and the new linear system is  $\bar{\mathbf{L}}\bar{\mathbf{x}} = \bar{\mathbf{b}}$ . Combining these two equations gives

$$\bar{\mathbf{L}}\bar{\mathbf{x}} = \bar{\mathbf{b}} - \mathbf{b} + \mathbf{L}\mathbf{x} = \Delta\mathbf{b} + \mathbf{L}\mathbf{x}$$

where  $\Delta\mathbf{b} = \bar{\mathbf{b}} - \mathbf{b}$ . Suppose we are given  $\mathbf{L}$ ,  $\mathbf{x}$ , and  $\Delta\mathbf{b}$ . The matrix  $\bar{\mathbf{L}}$  is computed

column-by-column by either the low-rank update/downdate algorithm in [3, 4], or the row and column modification algorithm discussed in this paper. We can combine the modification of  $\mathbf{L}$  with the computation of  $\bar{\mathbf{x}}$ , as shown in Algorithm 3:

**Algorithm 3 (Dense modification of  $\mathbf{L}\mathbf{x} = \mathbf{b}$  to solve  $\bar{\mathbf{L}}\bar{\mathbf{x}} = \bar{\mathbf{b}}$ )**

```

 $\bar{\mathbf{x}} = \Delta\mathbf{b}$ 
for  $j = 1$  to  $n$  do
     $\bar{\mathbf{x}} = \bar{\mathbf{x}} + (\mathbf{L})_{*j}x_j$ 
    compute the new column  $(\bar{\mathbf{L}})_{*j}$ 
     $\bar{\mathbf{x}}_{j+1\dots n} = \bar{\mathbf{x}}_{j+1\dots n} - (\bar{\mathbf{L}})_{j+1\dots n,j}\bar{x}_j$ 
end for

```

**end Algorithm 3**

The total work to modify the solution to  $\bar{\mathbf{L}}\bar{\mathbf{x}} = \bar{\mathbf{b}}$  is roughly  $2n^2$ , as compared to  $n^2$  work to solve  $\bar{\mathbf{L}}\bar{\mathbf{x}} = \bar{\mathbf{b}}$  from scratch. One would never use this method if  $\mathbf{L}$  and  $\mathbf{b}$  are dense.

Now consider the sparse case. Suppose we have a low-rank sparse update of  $\mathbf{L}$ . The columns that change in  $\bar{\mathbf{L}}$  correspond to a single path  $\bar{\mathcal{P}}$  in the elimination tree for a rank-1 update. Every entry in these specific columns is modified. For a rank- $r$  update, where  $r > 1$ , the columns that change correspond to a set of paths in the elimination tree, which we will also refer to as  $\bar{\mathcal{P}}$  in this more general case. In both cases, the nonzero pattern of each column  $j$  in the path  $\bar{\mathcal{P}}$  is a subset of the path [20]. We can thus partition the matrices  $\mathbf{L}$  and  $\bar{\mathbf{L}}$  into two parts, according to the set  $\bar{\mathcal{P}}$ . The original system is

$$\begin{bmatrix} \mathbf{L}_{11} & \mathbf{0} \\ \mathbf{L}_{21} & \mathbf{L}_{22} \end{bmatrix} \begin{bmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \end{bmatrix} = \begin{bmatrix} \mathbf{b}_1 \\ \mathbf{b}_2 \end{bmatrix}$$

where the matrix  $\mathbf{L}_{22}$  consists of all the rows and columns corresponding to nodes in the path  $\bar{\mathcal{P}}$ . If the changes in the right-hand-side  $\bar{\mathbf{b}}$  are also constrained to the set  $\bar{\mathcal{P}}$ , the new linear system is

$$\begin{bmatrix} \mathbf{L}_{11} & \mathbf{0} \\ \mathbf{L}_{12} & \bar{\mathbf{L}}_{22} \end{bmatrix} \begin{bmatrix} \mathbf{x}_1 \\ \bar{\mathbf{x}}_2 \end{bmatrix} = \begin{bmatrix} \mathbf{b}_1 \\ \bar{\mathbf{b}}_2 \end{bmatrix}.$$

Note that  $\mathbf{L}_{11}$ ,  $\mathbf{L}_{12}$  and  $\mathbf{b}_1$  do not change, and thus  $\mathbf{x}_1$  does not change. We have made a sparse change to both the matrix  $\mathbf{L}$  and the right-hand-side. The solution to the subsystem  $\mathbf{L}_{11}\mathbf{x}_1 = \mathbf{b}_1$  does not change. We have

$$\begin{aligned} \mathbf{L}_{21}\mathbf{x}_1 + \mathbf{L}_{22}\mathbf{x}_2 &= \mathbf{b}_2 \\ \mathbf{L}_{21}\mathbf{x}_1 + \bar{\mathbf{L}}_{22}\bar{\mathbf{x}}_2 &= \bar{\mathbf{b}}_2. \end{aligned}$$

We can apply Algorithm 3 to the subsystem

$$\bar{\mathbf{L}}_{22}\bar{\mathbf{x}}_2 = \Delta\mathbf{b}_2 + \mathbf{L}_{22}\mathbf{x}_{22},$$

where  $\Delta\mathbf{b}_2 = \bar{\mathbf{b}}_2 - \mathbf{b}_2$ , to obtain the new solution  $\bar{\mathbf{x}}_2$ . Algorithm 3 takes  $4|\bar{\mathbf{L}}_{22}| + O(\bar{\mathcal{P}})$  floating-point operations to compute  $\bar{\mathbf{x}}_2$ . An additional  $4r|\bar{\mathbf{L}}_{22}| + O(\bar{\mathcal{P}})$  floating-point operations are required to update  $\bar{\mathbf{L}}_{22}$ . Solving  $\bar{\mathbf{L}}\bar{\mathbf{x}} = \bar{\mathbf{b}}$  after  $\bar{\mathbf{L}}$  is updated takes  $2|\bar{\mathbf{L}}|$  floating-point operations. If the path  $\bar{\mathcal{P}}$  is short, making a sparse modification to the old solution  $\mathbf{x}$  to obtain the new solution  $\bar{\mathbf{x}}$  during the update of  $\bar{\mathbf{L}}$  takes much less work than solving the new linear system after  $\bar{\mathbf{L}}$  is updated.

Finally, consider the case discussed in this paper of modifying row and column  $k$  of  $\mathbf{C}$ . If we add row  $k$ , the original lower triangular system is

$$(9.1) \quad \begin{bmatrix} \mathbf{L}_{11} & & \\ \mathbf{0}^\top & 1 & \\ \mathbf{L}_{31} & \mathbf{0} & \mathbf{L}_{33} \end{bmatrix} \begin{bmatrix} \mathbf{x}_1 \\ x_2 \\ \mathbf{x}_3 \end{bmatrix} = \begin{bmatrix} \mathbf{b}_1 \\ b_2 \\ \mathbf{b}_3 \end{bmatrix}.$$

The new lower triangular system is

$$(9.2) \quad \begin{bmatrix} \mathbf{L}_{11} & & \\ \bar{\mathbf{I}}_{12}^\top & 1 & \\ \mathbf{L}_{31} & \bar{\mathbf{I}}_{32} & \bar{\mathbf{L}}_{33} \end{bmatrix} \begin{bmatrix} \mathbf{x}_1 \\ \bar{x}_2 \\ \bar{\mathbf{x}}_3 \end{bmatrix} = \begin{bmatrix} \mathbf{b}_1 \\ \bar{b}_2 \\ \bar{\mathbf{b}}_3 \end{bmatrix},$$

where we assume  $\mathbf{b}_1$  does not change, and the entries that change in  $\mathbf{b}_3$  are a subset of the path  $\bar{\mathcal{P}}$ . Let  $\Delta\mathbf{b}_3 = \bar{\mathbf{b}}_3 - \mathbf{b}_3$ . The term  $\bar{x}_2$  can be computed as

$$\bar{x}_2 = \bar{b}_2 - \bar{\mathbf{I}}_{12}^\top \mathbf{x}_1.$$

The (33)-blocks of (9.1) and (9.2) give the equations

$$(9.3) \quad \begin{aligned} \mathbf{L}_{33}\mathbf{x}_3 &= \mathbf{b}_3 - \mathbf{L}_{31}\mathbf{x}_1 \\ \bar{\mathbf{L}}_{33}\bar{\mathbf{x}}_3 &= (\mathbf{b}_3 - \mathbf{L}_{31}\mathbf{x}_1) + (\Delta\mathbf{b}_3 - \bar{\mathbf{I}}_{32}\bar{x}_2). \end{aligned}$$

The change in the right-hand-side of this system is  $\Delta\mathbf{b}_3 - \bar{\mathbf{I}}_{32}\bar{x}_2$ . Since the nonzero pattern of  $\bar{\mathbf{I}}_{32}$  is a subset of  $\bar{\mathcal{P}}$ , this computation fits the same requirements for the sparse change in  $\mathbf{x}$  due to a sparse low-rank change in  $\mathbf{L}$  and  $\mathbf{b}$ , discussed above. The row deletion case is analogous.

The number of floating-point operations in Algorithm 3 to modify the solution to  $\bar{\mathbf{L}}\bar{\mathbf{x}} = \bar{\mathbf{b}}$  is roughly the same as a rank-1 update to compute the modified  $\bar{\mathbf{L}}$ . However, the run time is not doubled compared to a stand-alone rank-1 update. At step  $j$ , for each  $j$  in the path  $\bar{\mathcal{P}}$ , the rank-1 update must read column  $j$  of  $\mathbf{L}$ , modify it, and then write the  $j$ th column of  $\bar{\mathbf{L}}$  back into memory. No extra memory traffic to access  $\bar{\mathbf{L}}$  is required to compute the solution to the lower triangular system using (9.3). With current technology, memory traffic can consume more time than floating point operations. Thus, when modifying the  $k$ th row and column of  $\mathbf{C}$ , or performing a rank- $r$  update/downdate to  $\mathbf{C}$ , we can update the solution to the lower triangular system at almost no extra cost. In our target application [5], solving the linear system  $\mathbf{C}\mathbf{x} = \mathbf{b}$ , given its  $\mathbf{LDL}^\top$  factorization, can often be the dominant step. The method presented here can cut this time almost in half since the forward solve time is virtually eliminated, leaving us with the time for the upper triangular backsolve.

**10. Experimental results.** To illustrate the performance of the algorithms developed in this paper in a specific application, Table 10.1 gives the flops associated with the solution of the four largest problems in the Netlib LP test set using the Linear Programming Dual Active Set Algorithm (LPDASA [5]). The problems passed to the solver were first simplified using the ILOG CPLEX [2] version 7.0 presolve routine. An LP presolver [1] preprocesses the problem by removing redundant constraints and variables whose values can be easily determined. Hence, the number of rows  $n$  of the problems listed in Table 10.1 is much smaller than the number of rows in the original Netlib problems.

The problems are displayed according to the average density of the lower triangular  $\mathbf{L}$  with the densest problems at the top of the table. The density measures the

TABLE 10.1  
*Experimental results*

Problem	Performance	Forward solve	Column updates	Row deletions
DFL001 <i>n</i> : 3881 density: .097	number: avg. mod. rank avg. mod. flops avg. solve flops		2629 1.2 $2.43 \times 10^6$ $2.03 \times 10^6$	87 1 $1.64 \times 10^6$ $1.64 \times 10^6$
PDS20 <i>n</i> : 10214 density: .011	number: avg. mod. rank avg. mod. flops avg. solve flops		1736 3.5 $3.37 \times 10^6$ $1.21 \times 10^6$	65 1 $2.13 \times 10^6$ $2.12 \times 10^6$
KEN18 <i>n</i> : 39856 density: .00012	number: avg. mod. rank avg. mod. flops avg. solve flops		2799 15.6 $61.7 \times 10^3$ 7155	23 1 2242 2075
OSA60 <i>n</i> : 10209 density: .00011	number: avg. mod. rank avg. mod. flops avg. solve flops		71 58.9 4415 270	277 1 62 37

fraction of the elements beneath the diagonal that are nonzero. The density varies from 0.097 for DFL001 down to 0.00011 OSA60. The column labeled “forward solve” gives the average number of flops that would have been required if forward solves were done in a conventional way, by a forward elimination process. As the LP is solved, the sparsity of  $\mathbf{L}$  changes slightly due to changes in the current basis. Hence, we give the average flops needed for a conventional solve of  $\mathbf{L}\mathbf{x} = \mathbf{b}$ , which can be compared with the average flops in modifying the solution using the technique developed in this paper. The columns labeled “column update” and “row deletion” partition the update flops into those associated with a change of the form  $\mathbf{C} + \mathbf{W}\mathbf{W}^T$  and those associated with the deletion of a row and column from  $\mathbf{C}$ .

In Table 10.1 we also give the average rank of the update. Since we have developed [4] a multiple rank approach for performing column updates, the average ranks listed in column 4 are all greater than 1; they are near 1 for the densest problem DFL001 and near 60 for the sparsest problem OSA60. Since the row deletion algorithm developed in this paper is a single rank process, the ranks listed in column 5 are all 1.

We also list the flops associated with modifying  $\mathbf{L}$ . For the column updates, this is the flop count for a typical rank- $r$  update, where  $r$  is the average rank of the updates as given in the table. Recall that the column update requires about  $4r$  flops per entry in  $\mathbf{L}$  that change. Modifying the forward solve takes 4 flops per entry in  $\mathbf{L}$  that change. The average flops associated with the modification of  $\mathbf{L}$  is thus always greater than the flops associated with the forward solve update, especially for multiple rank column updates. The conventional forward solve requires 2 flops for each nonzero entry in  $\mathbf{L}$ , whether they change or not. In the worst case, when all of the entries in  $\mathbf{L}$  change, modifying the forward solve takes no more than twice the work of the conventional forward solve, and a column update takes no more than  $2r$  times the work.

When the matrix  $\mathbf{L}$  is very sparse (such as OSA60), both the column update and row deletion methods modify only a small portion of the matrix. The elimination tree tends to be short and wide, with short paths from any node to the root. Thus, for very sparse problems the conventional forward solve (which accesses all of  $\mathbf{L}$ ) takes much more work than either the column update or the row deletion. For the OSA60

matrix, the conventional forward solve takes about 40 times the work compared to modifying the forward solve during a column update.

For a matrix  $\mathbf{L}$  that is fairly dense (such as DFL001), the elimination tree tends to be tall and thin, and the column updates or row deletions modify most of entries in  $\mathbf{L}$ . The work required to modify the forward solve is more on average than a full forward solve. However, since the update of the forward solve can be done at the same time as the update of the factorization, this single pass over  $\mathbf{L}$  is typically faster than the two pass algorithm, even though the flops are slightly higher on average.

**11. Summary.** We have presented a method for modifying the sparse Cholesky factorization of a symmetric, positive definite matrix  $\mathbf{C}$  after a row and column of  $\mathbf{C}$  have been modified. One algorithm, the sparse row addition, is optimal. The corresponding row deletion algorithm is not optimal, but takes no more time than the row addition. Although changing a row and column of  $\mathbf{C}$  can be cast as rank-2 change of the form  $\mathbf{C} \pm \mathbf{W}\mathbf{W}^T$ , the latter is impractical in a sparse context. We have shown how to modify the solution to a lower triangular system  $\mathbf{L}\mathbf{x} = \mathbf{b}$  when the matrix  $\mathbf{L}$  changes as a result of either the row addition/deletion operation discussed here, or an update/downdate of the form  $\mathbf{C} \pm \mathbf{W}\mathbf{W}^T$  described in our previous papers [3, 4]. By postponing the symbolic downdate, the memory usage has been reduced by 25% (assuming 8-byte floating-point values and 4-byte integers), compared with the column update/downdate methods described in [3, 4]. Together, the row addition/deletion algorithms and the column update/downdate algorithms form a useful suite of tools for modifying a sparse Cholesky factorization and for solving a sparse system of linear equations. Using these algorithms, the linear programming solver LPDASA is able to achieve an overall performance that rivals, and sometimes exceeds, the performance of current commercially available solvers [5].

## REFERENCES

- [1] E. D. ANDERSEN AND K. D. ANDERSEN, *Presolving in linear programming*, Math. Program., 71 (1995), pp. 221–245.
- [2] R. E. BIXBY, *Progress in linear programming*, ORSA J. on Computing, 6 (1994), pp. 15–22.
- [3] T. A. DAVIS AND W. W. HAGER, *Modifying a sparse Cholesky factorization*, SIAM J. Matrix Anal. Appl., 20 (1999), pp. 606–627.
- [4] ———, *Multiple-rank modifications of a sparse Cholesky factorization*, SIAM J. Matrix Anal. Appl., 22 (2001), pp. 997–1013.
- [5] ———, *A sparse multilevel implementation of the LP dual active set algorithm*, tech. report, Dept. of Mathematics, Univ. of Florida, Gainesville, FL 32611, Submitted to Math. Program. Available at <http://www.math.ufl.edu/~hager> and as CISE tech report TR-03-014 at [http://www.cise.ufl.edu/tech\\_reports](http://www.cise.ufl.edu/tech_reports), July 26, 2003.
- [6] S. C. EISENSTAT, M. C. GURSKY, M. H. SCHULTZ, AND A. H. SHERMAN, *Yale sparse matrix package, I: The symmetric codes*, Internat. J. Numer. Methods Engrg., 18 (1982), pp. 1145–1151.
- [7] A. GEORGE AND J. W. H. LIU, *An optimal algorithm for symbolic factorization of symmetric matrices*, SIAM J. Computing, 9 (1980), pp. 583–593.
- [8] ———, *Computer Solution of Large Sparse Positive Definite Systems*, Englewood Cliffs, New Jersey: Prentice-Hall, 1981.
- [9] J. R. GILBERT, *Predicting structure in sparse matrix computations*, SIAM J. Matrix Anal. Appl., 15 (1994), pp. 62–79.
- [10] 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.
- [11] P. E. GILL, G. H. GOLUB, W. MURRAY, AND M. A. SAUNDERS, *Methods for modifying matrix factorizations*, Math. Comp., 28 (1974), pp. 505–535.
- [12] W. W. HAGER, *Updating the inverse of a matrix*, SIAM Rev., 31 (1989), pp. 221–239.

- [13] ———, *The dual active set algorithm*, in *Advances in Optimization and Parallel Computing*, P. M. Pardalos, ed., North Holland, Amsterdam, 1992, pp. 137–142.
- [14] ———, *The LP dual active set algorithm*, in *High Performance Algorithms and Software in Nonlinear Optimization*, R. D. Leone, A. Murli, P. M. Pardalos, and G. Toraldo, eds., Dordrecht, 1998, Kluwer, pp. 243–254.
- [15] ———, *The dual active set algorithm and its application to linear programming*, *Comput. Optim. Appl.*, 21 (2002).
- [16] ———, *The dual active set algorithm and the iterative solution of linear programs*, in *Novel Approaches to Hard Discrete Optimization*, P. M. Pardalos and H. Wolkowicz, eds., Kluwer Academic Publishers, 2003, pp. 95–107.
- [17] W. W. HAGER AND D. W. HEARN, *Application of the dual active set algorithm to quadratic network optimization*, *Comput. Optim. Appl.*, 1 (1993), pp. 349–373.
- [18] J. W. H. LIU, *A compact row storage scheme for Cholesky factors using elimination trees*, *ACM Trans. Math. Software*, 12 (1986), pp. 127–148.
- [19] ———, *The role of elimination trees in sparse factorization*, *SIAM J. Matrix Anal. Appl.*, 11 (1990), pp. 134–172.
- [20] R. SCHREIBER, *A new implementation of sparse Gaussian elimination*, *ACM Trans. Math. Software*, 8 (1982), pp. 256–276.