

Maintaining a Sparse Inverse in the Simplex Method

Abstract: Improved methods are discussed for handling sparse matrices in practical linear programming. An analytical comparison is made of four methods for updating the inverse in the iterations following a reinversion. Of these, one technique using the elimination form of inverse is selected for some computational experiments and its advantages in terms of speed and sparseness demonstrated.

Introduction

Sparse matrix methods have been important in practical linear programming from the very earliest implementations of the revised simplex method. These methods have become even more important as the size of problems presented for solution has continued to grow until, at present, linear programs with a few thousand constraints are quite commonplace. Fortunately, and not surprisingly, the densities of the constraint matrices decline with the increase in dimension, thus keeping within manageable proportions the amount of data storage and arithmetic.

Techniques for handling sparse matrices are important in two major areas in linear programming:

1. Storage and retrieval of the problem data. There are several techniques for storing sparse matrices in use at present [1], [2].
2. Maintaining and applying the basis inverse or substitute inverse.

This paper is concerned with the second area. Readers will be assumed to be familiar with the simplex method using the product form of inverse [3] and the Gaussian elimination for solving systems of equations [4].

In recent years it has been realized that the sparsity of the inverse representation produced by the linear programming inversion routine can be drastically improved

by using the Gaussian or "elimination" form of inverse (E.F.I.) rather than using the Gauss-Jordan or product form of inverse (P.F.I.). This scheme was first advocated by Markowitz [5] and subsequently, for the special case of staircase matrices, by Dantzig [6]. Later still Dantzig, Harvey, McKnight and Smith [7] experimentally demonstrated the superiority of the E.F.I. over the P.F.I. for a number of linear programming problems and Brayton, Gustavson and Willoughby [8] were able to prove this superiority for general sparse matrices.

The elimination form of inverse is now known to be implemented in at least two commercial linear programming codes; Standard Oil's M5 code [7] and Scientific Control Systems' UMPIRE code (see Beale [9]). Other codes are in the course of being modified and may have been already altered. Some aspects of the P.F.I. and E.F.I. inversion techniques will be examined in the next section.

The success of the Gaussian or elimination form of inverse naturally leads one to ask whether the resulting triangular factors can be used to advantage in updating the inverse in the iterations following a reinversion. The answer would certainly appear to be yes. Dantzig [6] advocated updating the triangular factors of the basis for his staircase algorithm, since this preserves the staircase form of the factors. It seems extremely likely that in this

special case updating the factors, rather than using the ordinary product form of updating, would lead to increased efficiency. Bennett and Green [10] pursued this technique for general sparse matrices and Bartels [11] and Bartels and Golub [12] have advocated a form of Dantzig's method on the grounds of numerical stability without, however, commenting on the sparsity implications. Brayton et al. [8] have proposed two other updating techniques, although their concern is not primarily with efficiency in terms of the simplex method.

The three new proposals will be reviewed and compared in the third section together with their implications for the simplex method. One of these methods is then selected for some computational experiments and results are given. A discussion follows subsequently on how this method might be incorporated efficiently into a production code.

Product and elimination forms of inverse

• P.F.I.

A popular method of carrying out reinversion [13] is to order the rows (implicitly) and the columns (explicitly) of B and then to partition as follows:

$$B = \begin{bmatrix} \Lambda_1 & & \\ A & \tilde{B} & \\ D & E & \Lambda_2 \end{bmatrix},$$

where Λ_1 and Λ_2 are lower triangular. The submatrix $\tilde{B}$ is sometimes referred to as the "bump" and those vectors with elements in Λ_1 are "above the bump", those in Λ_2 being "below the bump". It seems to be characteristic of most linear programs that the submatrices Λ_1 and, particularly, Λ_2 are quite substantial portions of B .

The above partitioned form of B may be factorized thus:

$$B = \begin{bmatrix} \Lambda_1 & & \\ A & I & \\ D & & I \end{bmatrix} \begin{bmatrix} I & & \\ & \tilde{B} & \\ & & I \end{bmatrix} \begin{bmatrix} I & & \\ & I & \\ & E & I \end{bmatrix} \begin{bmatrix} I & & \\ & I & \\ & & \Lambda_2 \end{bmatrix}.$$

Since Λ_1 and Λ_2 are triangular, the calculation of the product form of inverse is nontrivial only for the submatrix $\tilde{B}$. Various "merit" schemes have been proposed for selecting the order of pivots in $\tilde{B}$ (see, e.g., Tewarson [14]) but they will not be discussed here.

Using the notation of Ref. 8 we may then express B as a product of elementary transformations

$$B = T_1 T_2 \cdots T_l,$$

or equivalently

$$B^{-1} = T_l^{-1} \cdots T_2^{-1} T_1^{-1},$$

where T_k is of the form assuming B is $m \times m$,

$$T_k = \begin{bmatrix} 1 & & & & & \\ & \ddots & & & & \\ & & 1 & & & \\ & & & \ddots & & \\ & & & & 1 & \\ & & & & & \ddots \\ & & & & & & 1 \end{bmatrix} \begin{matrix} t_{lk} \\ \\ \\ t_{kk} \\ \\ t_{mk} \end{matrix}$$

$$= I + (t_k - e_k) e_k',$$

where e_k is the k th unit m vector, e_k' its transpose, and

$$t_k = T_{k-1}^{-1} \cdots T_2^{-1} T_1^{-1} b_k,$$

where b_k is the k th column of B .

Note, however, that two transformations are made for each column in the "bump"; hence the elements of the T_k for such columns are zero on the rows corresponding to Λ_2 and in addition a second transformation is also made, with a unit pivot and the k th column of E on the rows of Λ_2 . Note also that $T_k^{-1} = I - t_{kk}^{-1} (t_k - e_k) e_k'$ and that it is essentially necessary to keep only the vector t_k to store T_k , and t_k and t_{kk}^{-1} to store T_k^{-1} . (See [8]).

• E.F.I.

The P.F.I. partitioning scheme may be used with only a slight modification in the elimination form of inverse. By rearranging the rows and columns of the submatrices Λ_2 , E and D , we may repartition B to give [9]

$$B = \begin{bmatrix} \Lambda_1 & & \\ \tilde{D} & \tilde{\Lambda}_2 & \tilde{E} \\ A & & \tilde{B} \end{bmatrix},$$

where $\tilde{\Lambda}_2$ is now upper triangular. This procedure obviously facilitates factorization into triangular factors by Gaussian elimination, since again only the factorization of $\tilde{B}$ is nontrivial.

Suppose B is now factorized into LU , where L is lower and U is upper triangular. Let l_k and u_k be the k th columns of L and U ; then defining

$$L_k = I + (l_k - e_k) e_k';$$

$$U_k = I + (u_k - e_k) e_k',$$

we may factorize

$$L = L_1 L_2 \cdots L_m,$$

$$U = U_m U_{m-1} \cdots U_1.$$

Hence

$$B^{-1} = U_1^{-1} \cdots U_m^{-1} L_m^{-1} \cdots L_1^{-1}.$$

There is some latitude in the choice of values of diagonal elements of L and U . The most efficient choice would appear to be setting u_{kk} to 1 for k in partitions Λ_1 , $\bar{B}$ and l_{kk} to 1 in partition Λ_2 . If this is done the number of transformations (disregarding unit transformations) is exactly the same as for the P.F.I., since two transformations are calculated only for the columns in $\bar{B}$.

The order of choice of pivot elements in $\bar{B}$ may be decided here, as for the P.F.I., by merit schemes. These are exhaustively discussed by Dantzig et al. [7].

Methods of updating inverse

In this section we briefly review, without rigorous proof, and compare four methods of updating the inverse. All except the first assume the elimination form of inverse discussed in the previous section. With the exception of Method II all are described at greater length by Brayton et al. [8].

• Method I [3]

Method I is the standard procedure used in the product form of the simplex method.

Suppose column b_k of the basis is to be replaced by $\bar{b}_k$. Let B be the original basis (factorized in either E.F.I. or P.F.I. form) and let $\bar{B}$ be the new basis.

$$\text{Let } t_{l+1} = B^{-1} \bar{b}_k,$$

$$\text{then } \bar{B}^{-1} = T_{l+1}^{-1} B^{-1},$$

$$\text{where } T_{l+1} = I + (t_{l+1} - e_k) e_k',$$

$$\text{Hence if } B^{-1} = T_l^{-1} \cdots T_2^{-1} T_1^{-1},$$

$$\text{then } \bar{B}^{-1} = T_{l+1}^{-1} T_l^{-1} \cdots T_1^{-1}.$$

Further changes are made by repeating the above process.

• Method II [11]

This is Bartels's version of Dantzig's algorithm [6]. Suppose we have $B = LU$, i.e., $B^{-1} = U^{-1} L_m^{-1} \cdots L_1^{-1}$, and we again change b_k to $\bar{b}_k$. Then $L^{-1} \bar{B}$ will be identical to U except in column k . Permute columns $k+1, \dots, m$ one place to the left and place column k in position m to give a Hessenberg matrix,

$$H = \begin{array}{|c|} \hline \text{[Diagram of a Hessenberg matrix with shaded upper triangular region]} \\ \hline \end{array} = L^{-1} \bar{B} Q,$$

where Q is a permutation matrix.

The elements below the diagonal in columns k to $m-1$ of H must now be eliminated by a set of elementary transformations. An important feature of Bartels' method is that, in carrying out the elimination, rows may be interchanged or permuted to place the maximum of the elements $h_{il}, h_{l+1,l}$ on the diagonal (Wilkinson's "partial pivoting" strategy [15]) to ensure numerical stability.

The result is a new upper triangular matrix

$$\bar{U} = G_{m-1} P_{m-1} \cdots G_{k+1} P_{k+1} G_k P_k H,$$

where the P_l are either unit matrices or the permutation matrices exchanging rows l and $l+1$, and G_l are the elementary transformation matrices eliminating the element below the diagonal in column l . Hence

$$\bar{B}^{-1} = Q \bar{U}^{-1} G_{m-1} P_{m-1} \cdots G_k P_k L^{-1}.$$

Dantzig [6] and Bennett and Green [10] have shown that the "near commutativity" of the matrices G_l may be used to calculate a new lower triangular matrix $\bar{L}$ such that

$$\bar{L}^{-1} = G_{m-1} P_{m-1} \cdots G_k P_k L^{-1},$$

thus maintaining a true triangular decomposition. This is not, however, an essential feature of the method.

Of course the permutation matrices P_l , Q need not appear explicitly. All that is necessary is to record the order of pivoting. The process is repeated for the next iteration using the new $\bar{U}$ and (implicitly) the new $\bar{L}$.

• Method III [8]

Neither this method nor the next maintain a true triangular decomposition. However, both require initially the elimination form of inverse.

Suppose we have $B = LU$ or, equivalently,

$$B^{-1} = U_1^{-1} \cdots U_m^{-1} L_m^{-1} \cdots L_1^{-1}$$

and again change column b_k to $\bar{b}_k$ and B to $\bar{B}$.

The method proceeds as follows: Let

$$v_k = L^{-1} \bar{b}_k.$$

Then $L^{-1} \bar{B}$ is identical to U except for column k which is now v_k instead of u_k .

To reduce $L^{-1} \bar{B}$ back to upper triangular form the first step carried out is to reduce row k to zero in columns $k+1, \dots, m$.

$$\begin{array}{|c|} \hline \text{[Diagram of a matrix with shaded upper triangular region and a vertical line at column k]} \\ \hline \end{array}$$

To accomplish this we form a row transformation using the row vector

$$\begin{aligned} w_k' &= (0, \dots, 0, w_{k,k+1}, \dots, w_{km}) \\ &= (0, \dots, 0, u_{k,k+1}, \dots, u_{km})U^{-1}. \end{aligned}$$

The row transformation is then the elementary matrix $W_k = I + e_k w_k'$ and the desired elimination is performed by forming $W_k^{-1}L^{-1}\bar{B}$, since the columns $k+1, \dots, m$ of $L^{-1}\bar{B}$ are identical to U , $W_k^{-1}U = (I - e_k w_k')U$, and $w_k'U$ is simply the row $(0, \dots, 0, u_{k,k+1}, \dots, u_{km})$. Furthermore the first k columns of W_k are unit vectors and hence the first k columns of $L^{-1}\bar{B}$ are unchanged.

Now let $t_k = W_k^{-1}L^{-1}\bar{b}_k$ and form the elementary column transformation

$$T_k = I + (t_k - e_k)e_k'.$$

Then forming $T_k^{-1}W_k^{-1}L^{-1}\bar{B}$ clearly reduces column k to the unit vector e_k and columns $k+1, \dots, m$ are unchanged since they have a zero on the pivot row. The result then is a new upper triangular matrix, say $U^{(k)}$, obtained from U by replacing the k th row and column by the unit vector e_k ,

i.e.,

$$U^{(k)} = T_k^{-1}W_k^{-1}L^{-1}\bar{B},$$

or

$$\bar{B}^{-1} = U^{(k)-1}T_k^{-1}W_k^{-1}L^{-1}.$$

Using product form notation, if we let $U_l^{(k)}$ denote the elementary matrix U_l with u_{kl} set to zero we may rewrite $\bar{B}^{-1}$ as

$$\begin{aligned} \bar{B}^{-1} &= U_1^{-1} \dots U_{k-1}^{-1} U_{k+1}^{(k)-1} \\ &\dots U_m^{(k)-1} T_k^{-1} W_k^{-1} L_m^{-1} \dots L_1^{-1}. \end{aligned}$$

Note that the row transformation W_k can be represented, if desired, by a product of two-element elementary column transformations. Also note that, although we have introduced two new transformations, if U was previously in product form, $U = U_m \dots U_1$, the new form $U^{(k)}$ is obtained simply by deleting U_k and a single element at most from each $U_{k+1} \dots U_m$ for a net gain of one transformation in general, unless U_k happens to be the unit matrix.

To carry on the process, one simply treats $T_k^{-1}W_k^{-1}L^{-1}$ as the lower triangular factor as before, although in fact if multiplied out the resulting matrix would not, in general, be triangular.

• Method IV [8]

Method IV, which also assumes the elimination form of inverse, proceeds by inserting matrices T_k into the list $U_m \dots U_1$ and sometimes deleting a U_k .

To change b_k to $\bar{b}_k$ we form the vector

$$t_k = U_{k+1}^{-1} \dots U_m^{-1} L_m^{-1} \dots L_1^{-1} \bar{b}_k$$

and replace U_k by the transformation

$$T_k = I + (t_k - e_k)e_k'.$$

Then

$$\begin{aligned} \bar{B}^{-1} &= U_1^{-1} \dots U_{k-1}^{-1} T_k^{-1} U_{k+1}^{-1} \dots U_m^{-1} L_m^{-1} \\ &\dots L_1^{-1}. \end{aligned}$$

Continuation of this process is rather more complicated than in the preceding methods. Suppose now we have changed b_k to $\bar{b}_k$ and wish to change b_l to $\bar{b}_l$. There are two cases:

1. If $l \leq k$, then let $t_l = U_{l+1}^{-1} \dots U_{k-1}^{-1} T_k^{-1} U_{k+1}^{-1} \dots U_m^{-1} L_m^{-1} \dots L_1^{-1} \bar{b}_l$ form $T_l = I + (t_l - e_l)e_l'$ and replace U_l by T_l as before.
2. If $l > k$ form $t_l = T_k^{-1} U_{k+1}^{-1} \dots U_m^{-1} L_m^{-1} \dots L_1^{-1} \bar{b}_l$ and the transformation T_l as above. However, U_l is not replaced and T_l^{-1} is placed directly before T_k^{-1} in the inverse transformation list, i.e.,

$$\begin{aligned} \bar{B}^{-1} &= U_1^{-1} \dots U_{k-1}^{-1} T_l^{-1} T_k^{-1} U_{k+1}^{-1} \dots U_m^{-1} L_m^{-1} \\ &\dots L_1^{-1}. \end{aligned}$$

Continuing further, if columns $k_1, \dots, k_p$ have been changed and column l is to be altered, let $k = \min(k_1, \dots, k_p)$ and carry out the above procedure.

Having completed our summary of methods for updating a basis inverse or substitute inverse, we must consider their relative merits in the context of the simplex method in terms of preserving sparsity, work required per iteration, and storage and data handling considerations.

The great advantages of Method I are, of course, its simplicity and the minimum of computational steps and housekeeping involved. The essential steps in changing the basis in the simplex method, given the constraints $Ax = b$ and the new column a_j to enter the basis, are the calculation of the updated column $\alpha = B^{-1}a_j$ and the ratio test with the updated right-hand side, $\beta = B^{-1}b$, to find, assuming $\beta \geq 0$, the pivot row k from

$$\theta = \min_{\alpha_i > 0} \frac{\beta_i}{\alpha_i}.$$

The vector α , which we cannot avoid calculating, is precisely the vector t_{l+1} used in Method I to update the basis inverse (and the right-hand side).

Methods II through IV all suffer from the disadvantage that they require not only the vector α for determining the pivot row but also another vector t_k , which is only a partially updated form of a_j . Methods II and III can more or less overcome this disadvantage by putting the vector $L^{-1}a_j$ in temporary storage in the course of cal-

culating α , at the expense either of using more of the computer's core storage, thus effectively limiting the number of vectors which can be saved for multiple pricing, or of using backing store. This burden is, however, comparatively light compared to the immediate and acute difficulty encountered in Method IV, where the vector t_k to be calculated for updating purposes cannot in general be known until *after* the ratio test operation to determine the pivot row. This implies that t_k must be calculated from scratch or obtained by backward transformations on the vector α .

Now let us consider the implications of Methods II to IV in terms of work per iteration expended in updating the basis inverse and of the necessary data manipulation. Current packing methods of storing sparse data, particularly strings of transformations [1], [2] make it rather difficult to modify an existing product form of L and U . Replacement of U_k by a T_k with more nonzeros in Method IV, for example, necessitates either a wholesale shifting of elements of U , maintenance of a second transformation file to be interwoven with the first, copying out a complete new basis file, or initial and wasteful leaving of space for such eventualities. A similar, in fact worse, situation occurs in Method II where on the average half of the U_i transformations must be modified at each iteration, increasing the number of nonzeros to be stored. Explicit storage of U can be an answer only for very small problems. Some kind of two-file system, each taking the new inverse in turn, would appear in general to be necessary, involving a great deal of read/write activity. Method III on the other hand requires only the elimination of previously nonzero elements in U —a comparatively easy task. This point will be returned to in the section "Implementation of Method III."

The amount of arithmetic work per iteration would appear to be greatest in Methods II and III. The bulk of the work in Method II comes from applying the G_i transformations to H to reduce it to triangular form. On the average we expect to have to deal with $m/2$ columns on the right-hand side of the matrix. This is a considerable amount of computation and becomes progressively worse as U fills in with each iteration. For Method III the main effort is in computing the row vector w_k' . It will be shown later that this calculation may be performed concurrently with the backward-transformation phase of the next simplex iteration. The most important point, however, is that U becomes progressively less dense in this method and the calculation progressively easier. There is admittedly the extra complication in Method III that row transformations are called for, but this is more of an inconvenience than a problem and, as already pointed out, could be circumvented.

Finally we consider the most important point—preservation of sparsity. In their discussion Brayton et al. [8]

concluded that Method III was superior to Method IV in this respect since both a row and column of U are eliminated in exchange for a vector w_k' , which only has elements in positions $k+1, \dots, m$, and a transformation vector t_k , which is the incoming column multiplied only by L^{-1} . On the other hand if in Method IV some columns with very small pivot row index k are introduced, then virtually every change thereafter will be of type 2 and we cannot expect much improvement over the ordinary product form method.

In comparing Methods II and III it again seems almost certain that Method III will win out. The number of new elements added in the new transformations would appear to be much the same, but as already mentioned in Method II, U becomes more and more dense while in Method III the density of U declines. In this connection it should be pointed out that the columns of U most frequently operated upon are the right-most. These columns, however, are just those columns which will be initially the most dense after inversion, since inversion schemes generally postpone operating on the most dense columns until last.

In the light of these considerations it appears that Method III is the most promising for incorporation into the simplex method, both on the grounds of work per updating iteration and preservation of sparsity. However, there is certainly some increase in updating work per iteration over the standard product form (Method I) and a significant improvement in growth of nonzeros in the basis representation is required to justify its use. Some experimental results are presented in the next section.

Although we have chosen Method III as the best for our purposes, this is not to say that the others may not be superior in other circumstances. Bartels's method (II) is of considerable importance in achieving highly accurate solutions to small, dense problems (see Bartels and Golub [12]) while Dantzig's original version [6] of Method II for specially structured matrices is as yet untried.

Computational results

Some computational experiments have been carried out to determine whether Method III gives a sufficient improvement over Method I, in terms of growth of nonzeros, to warrant further investigation.

To carry out these experiments two linear programming codes were written and run on the IBM 360/67 at Stanford University. Both are all-in-core FORTRAN codes using the elimination form of inverse. Neither employs multiple pricing. The first code updates the inverse in standard product form fashion (Method I); the second is a modified version using Method III to update the inverse. No attempt has been made to evaluate timings, dependent as they are on the machine used and the pro-

Table 1 Problem statistics. (See text).

	Problem number				
	1	2	3	4	5
Name	SHARE 2B	SHARE 1B	BRANDY	BANDM	MAR
Rows	99	118	221	306	325
Columns (structural)	79	225	249	472	452
% Density (of structural columns)	10.25	4.45	3.91	1.59	1.77

gramming technique. Only the number of nonzero elements in the inverse representation is considered.

Five problems ranging from small to moderate have been used. These problems are two of the well-known SHARE standard test problems, 2B and 1B [16] and three other test problems supplied by J. Cord of IBM. The relevant statistics on these problems are given in Table 1.

Only the smallest problem was run to optimality, the remainder being cut off after 200, 250, 300 and 350 iterations. To compare the two methods the set of problems was run on both codes with a fixed inversion frequency: 40 iterations for the SHARE problems and 50 for the others. The initial comparison is arrived at by counting, at intervals of 10 iterations, the number of new nonzeros added to the inverse since the last reinversion. To further remove constant factors only nonpivot elements are counted. (Both methods give a net addition of one pivot element per iteration in general.) All problems were started from a slack basis and hence the iterations up to the first reinversion are ignored. The number of new nonpivot nonzero elements after multiples of 10 iterations from reinversion is then averaged for the remainder of the run. These figures are given in Table 2.

Table 2 Growth of new nonzeros in basis.

Problem		1		2		3		4		5	
Method		I	III	I	III	I	III	I	III	I	III
Iterations since invert	10	391	176	298	206	717	332	863	252	304	118
	20	891	492	638	497	1589	901	1866	693	628	354
	30	1112	847	993	830	2405	1545	2779	1238	1147	729
	40	1602	1515	1458	1272	3417	2342	3649	1882	1796	1135
	50	—	—	—	—	4254	3062	4646	2642	2261	1441

The raw data appear in more digestible form in Table 3, where Method I has been used as a base and the fraction of new elements produced by Method III is given. The improvement is obviously very satisfactory, with the possible exception of problem 2. As might be expected the improvement tails off as the number of iterations increases, particularly for the smaller, denser problems.

To obtain a better estimate of the effect of Method III on the simplex algorithm we should also consider the total number of nonpivot nonzeros in the inverse. To this end we list the inversion statistics in Table 4 and compute the fraction of total nonpivot nonzeros produced by Method III, again using Method I as a base. The inversion statistics do not include the original all-slack basis and are of some interest in themselves. The pivot-choosing procedure in decomposing $\tilde{B}$ (see the second section) is essentially algorithm 7 in Dantzig et al. [7].

The modified data for Method III appear in Table 5.

Although the problems are comparatively small and the sample is also small, the results in Table 5 are quite encouraging. It appears that for reasonably sparse problems with 200 to 300 constraints a reduction of 20 to 30 percent in the number of nonzeros in the substitute inverse may be expected with Method III, thus considerably reducing the time spent in forward and backward transformation in the simplex method. The results for problem 4 show that the reduction may be dramatic.

Implementation of Method III

The results of the preceding section suggest that further experiments using a modified production or commercial code would be worthwhile. It would also appear that the modifications could be made and the extra updating work per iteration accomplished at comparatively little cost.

The first essential is clearly to have a code that employs the elimination form of inverse. Given such a code it should be a simple matter to adjust the storage and buffering activities such that the $T_k^{-1}W_k^{-1}$ transforms can be adjoined to the left of the product form $L_m^{-1} \cdots L_1^{-1}$ of L^{-1} .

The reduction of the k th row and column of U to the unit vector e_k can probably be accomplished satisfac-

torily by simply setting indicator bits. Usually we will have U^{-1} represented as $U_1^{-1}U_2^{-1} \cdots U_m^{-1}$, or rather the packed vectors u_l corresponding to these elementary matrices. The reduction may then be accomplished by tagging U_k^{-1} with a bit to indicate that it is to be skipped. Furthermore by using, say, the blocked index scheme [1] for storing the vectors u_l , the row index k , if it appears, may be tagged to indicate that the element is zero for $l = k + 1, \dots, m$. Failing this the u_{kl} elements could be physically replaced by zeros if necessary.

This scheme is of course inefficient in the sense that in an out-of-core system progressively more and more valueless data must be buffered in and out of core. On the other hand if the system is already compute-bound rather than I/O-bound this is no restriction, and even if this is not the case the very frequent reinversions carried out for large problems should prevent this inefficiency from reaching really noticeable proportions.

Finally let us consider the process of computing w_k' . For simplicity let us assume we are carrying out the first iteration after a reinversion. We then have

$$B^{-1} = U_1^{-1} \cdots U_m^{-1} L_m^{-1} \cdots L_1^{-1}.$$

Now the calculation of

$$w_k' = (0, \dots, 0, u_{k,k+1}, \dots, u_{km})U^{-1},$$

and also the modification of U to $U^{(k)}$, may be carried out in a *single pass* from left to right through the transformations $U_1^{-1} \cdots U_m^{-1}$. To see this, note that in computing, say,

$$(0, \dots, 0, z_{k+1}, \dots, z_m)U_l^{-1} \quad \text{for } l > k,$$

only the elements $z_{k+1}, \dots, z_l$ play any part in the arithmetic, since U and U_l are upper triangular. The elements $z_{l+1}, \dots, z_m$ are unchanged. We therefore arrange the computation of w_k' and the modification of U to $U^{(k)}$ as follows.

• Initial step

Create a zero row vector and skip transformations $U_1^{-1} \cdots U_{k-1}^{-1}$. Tag U_k^{-1} . Let $l = k + 1$.

• General step

For transformation U_l^{-1} , if u_{kl} is nonzero, extract it and add it in position l of the row vector, marking the element as now being zero in U_l^{-1} . Postmultiply the row vector by U_l^{-1} . Proceed for $l = k + 1, \dots, m$.

A very advantageous result of this procedure is that the updating of the inverse can be carried out concurrently with the backward transformation of the *following* simplex iteration. In calculating the pricing vector for the next iteration we must compute

$$\pi' = c' U^{(k)-1} T_k^{-1} W_k^{-1} L^{-1}$$

Table 3 Fraction of new nonzeros produced by Method III.

Iterations since invert	Problem number				
	1	2	3	4	5
10	0.45	0.69	0.46	0.34	0.39
20	0.55	0.78	0.57	0.37	0.56
30	0.76	0.84	0.64	0.45	0.63
40	0.95	0.87	0.69	0.53	0.63
50	—	—	0.72	0.58	0.64

Table 4 Inversion statistics (averages).

	Problem number				
	1	2	3	4	5
Original basic nonzeros	488	498	735	755	822
Nonpivot, nonzero transformation elements	417	398	595	608	561
Number of transformations	76	107	133	126	172

Table 5 Fraction of total nonzeros produced by Method III.

Iterations since invert	Problem number				
	1	2	3	4	5
10	0.73	0.87	0.71	0.58	0.78
20	0.70	0.86	0.68	0.53	0.77
30	0.83	0.88	0.71	0.54	0.75
40	0.96	0.90	0.73	0.59	0.72
50	—	—	0.76	0.62	0.71

for some row vector c' . Now the first $k - 1$ columns of $U^{(k)}$ are identical with those of U . The modification of the remaining columns of U may be carried out in the same pass through the transformation file as that for the calculation of π' and furthermore only one unpacking of the nonzeros from each u_l is necessary. Having obtained w_k' the new transformations T_k^{-1} , W_k^{-1} may be attached to L^{-1} and the remainder of the calculation of π' can proceed. Looking to the future, we observe that this procedure is admirably suited to parallel processing facilities.

Conclusion

Linear programming inversion routines have now reached a very high level of sophistication and efficiency, in terms of both speed and sparseness of the resulting inverse. Although there is doubtless still some room for improvement it seems likely that further improvements in efficiency must be sought in other areas. The attempt to maintain a sparse inverse is one such area and our computational results give considerable grounds for optimism. What is needed now is full-scale experimentation with some of the methods reviewed here on large problems of 1000 and more constraints to confirm and extend these results.

Acknowledgments

This research was carried out under an IBM Post-Doctoral Fellowship at Stanford University. The author thanks IBM for their support. He is also greatly indebted to G. B. Dantzig, G. H. Golub and S. Maier for many helpful discussions, and to R. Harvey, A. Colville and J. Cord for supplying data for the computational experiments.

The work was partially supported by the office of Naval Research, Contract N-00014-67-A0112-0011; the U.S. Atomic Energy Commission, Contract AT[04-3]326 PA #18; and National Science Foundation Grant GP 9329.

This paper is based on Stanford Operations Research Technical Report 70-15.

Postscript

Since the work described in this paper was completed, J. J. H. Forrest and the author have devised and implemented an improved algorithm, based on those of Bartels and Golub, and Brayton et al. In principle this new algorithm could be regarded as carrying out the first step of Method II, i.e., the formation of the upper Hessenberg matrix H , and then applying a row transformation, as in Method III, to reduce H back to permuted triangular form by eliminating the elements h_{kj} ($j = k, \dots, m-1$) without row interchanges. This procedure has the great advantage that column transformations T_k^{-1} are not adjoined to the left of L^{-1} , but instead become part of the product form of the new upper triangular factor. Hence the "partially updated vector" v_k of Method III is now updated only by L^{-1} and the row transformations are produced once per iteration. Since a row transformation can introduce at most one new nonzero element to a vector, in contrast to a column transformation T_k^{-1} the rate of growth of nonzero elements is drastically reduced.

The new algorithm is fully described in Refs. 17 and 19, together with computational experience. The basic computational result is that the growth rate of new nonzeros in the transformation files is reduced by a consis-

tent 90 percent, rather than the 40 to 50 percent obtained with Method III, and the overall reduction in number of nonzeros by 60 to 70 percent rather than 20 to 30 percent on average. At present we find that this reduction results in an improvement of about 40 percent in terms of the number of simplex iterations per unit time.

Some experiments with Method II are also given in Ref. 19, which indicate that the Bartels and Golub method also gives considerable improvement in sparsity terms over the standard product form method, provided that the I/O problem discussed in the third section can be overcome.

The new triangular updating scheme given in Refs. 17 and 19, combined with the extremely efficient new inversion technique recently advocated by Hellerman and Rarick [18], have greatly increased the power of the simplex method.

References

1. D. M. Smith, "Data Logistics for Matrix Inversion", in *Sparse Matrix Proceedings*, R. Willoughby, Ed., RA-1, IBM Thomas J. Watson Research Center, Yorktown Heights, N.Y., 1969, pp. 127-138.
2. J. de Buchet, "How to Take Into Account the Low Density of Matrices to Design a Mathematical Programming Package—Relevant Effects on Optimization and Inversion Algorithms," Paper presented to the Oxford Symposium on Sparse Systems of Linear Equations, April, 1970.
3. G. B. Dantzig, *Linear Programming and Extensions*, Princeton University Press, Princeton, 1963.
4. G. Forsythe and C. B. Moler, *Computer Solution of Linear Algebraic Equations*, Prentice-Hall, Englewood Cliffs, N.J., 1967.
5. H. M. Markowitz, "The Elimination Form of Inverse and its Application to Linear Programming", *Management Sci.* 3, 255-269 (1957).
6. G. B. Dantzig, "Compact Basis Triangularization for the Simplex Method," in *Recent Advances in Mathematical Programming*, R. L. Graves and P. Wolfe, Eds., McGraw-Hill Book Co., Inc., New York, 1963, pp. 125-132.
7. G. B. Dantzig, R. P. Harvey, R. D. McKnight, and S. S. Smith, "Sparse Matrix Techniques in Two Mathematical Programming Codes," in *Sparse Matrix Proceedings* (R. Willoughby, Ed.), RA-1, IBM Thomas J. Watson Research Center, Yorktown Heights, N.Y., 1969, pp. 85-99.
8. R. K. Brayton, F. G. Gustavson, and R. A. Willoughby, "Some Results on Sparse Matrices", *Report RC-2332*, IBM Thomas J. Watson Research Center, Yorktown Heights, N.Y., February 14, 1969.
9. E. M. L. Beale, "Sparseness in Linear Programming," Paper presented to the Oxford Symposium on Sparse Systems of Linear Equations, April, 1970.
10. J. M. Bennett and D. R. Green, "Updating the Inverse or the Triangular Factors of a Modified Matrix," *Basser Computing Department, University of Sydney, Technical Report No. 42*, April, 1966.
11. R. H. Bartels, "A Numerical Investigation of the Simplex Method," *Computer Science Department, Stanford University, Technical Report No. CS-104*, July 31, 1968.
12. R. H. Bartels and G. H. Golub, "The Simplex Method of Linear Programming Using LU Decomposition," *Comm. ACM*, 12, 266-268, 275-278 (1969).
13. W. Orchard-Hays, *Advanced Linear Programming Computing Techniques*, McGraw-Hill Book Co., Inc., New York, 1968.

14. R. P. Tewarson, "On the Product Form of Inverses of Sparse Matrices", *SIAM Rev.* 8, 336-342 (1966).
15. J. H. Wilkinson, *Rounding Errors in Algebraic Processes*, Prentice-Hall, Englewood Cliffs, N.J., 1963.
16. D. M. Smith and W. Orchard-Hays, "Computational Efficiency in Product Form LP Codes", in *Recent Advances in Mathematical Programming*, R. L. Graves and P. Wolfe, Eds., McGraw-Hill Book Co., Inc., New York, 1963, pp. 211-218.
17. J. J. H. Forrest and J. A. Tomlin, "Updating Triangular Factors of the Basis to Maintain Sparsity in the Product Form Simplex Method," paper presented to the NATO Conference on Large Scale Mathematical Programming, Elsinore, Denmark, July 1971. To appear in *Mathematical Programming*.
18. E. Hellerman and D. Rarick, "Reinversion with the Pre-assigned Pivot Procedure", *Math. Prog.* 1, 195-216 (1971).
19. J. A. Tomlin, "Modifying Triangular Factors of the Basis of the Simplex Method", to appear in *Sparse Matrices and Their Application*, Rose & Willoughby, Eds., Plenum Press, New York, 1972.

Received November 23, 1971

The author is located at the Operations Research Department, Stanford University, Stanford, California 94305.