

Notes 5: **Basis Updates**

## 1 Many Solves with $B$ and $B^T$

Each iteration, Simplex and Reduced-Gradient methods need to solve systems of the form  $Bx = v$  and  $B^T y = w$  (at least one of each), where  $B$  is the current square, nonsingular basis. Large problems may require tens of thousands of iterations. It is therefore normal to use *direct* methods rather than *iterative* methods as linear-system solvers. (Typically this means sparse LU factorization methods.)

## 2 Basis Updating

In general it is too expensive to compute brand new LU factors for every  $B$ . Instead, we take advantage of the fact that each  $B$  is almost the same as the one before. (One column of  $B$  has been replaced by a column of  $A$ .) In some way we *update* the current factors of  $B$ . Sometimes we keep the existing factors (unaltered) and accumulate updating information on the side. Other methods modify the existing factors themselves.

Suppose a column  $\bar{a}$  replaces the  $p$ th column of  $B$  to give the next basis  $\bar{B}$ . Then

$$\bar{B} = B + (\bar{a} - a)e_p^T, \quad (1)$$

where  $e_p$  is the  $p$ th column of the identity. Given some kind of factorization of  $B$ , we need to obtain a factorization of  $\bar{B}$ . We study several methods that have varying degrees of stability, efficiency, and ease of implementation.

## 3 PFI and EFI

For many years, large-scale implementations of the Simplex Method used what was called “the product form of inverse”, in which the solution of  $Bx = b$  was thought of (and implemented) as a product of the form

$$x = E_l \dots E_2 E_1 b,$$

where each matrix  $E_j$  was the identity except for one (somewhat sparse) column  $v_j$ . After a basis change, the solution of  $\bar{B}\bar{x} = \bar{b}$  would be computed as

$$\bar{x} = E_{l+1} E_l \dots E_2 E_1 \bar{b},$$

and similarly for systems involving  $B^T$  and  $\bar{B}^T$ . The method allowed efficient use of auxiliary storage, with the growing list of vectors  $\{v_1, v_2, \dots\}$  being accessed sequentially (either forwards or backwards).

Eventually the terms arising from an initial basis  $B_0$  were constructed from LU factors of  $B_0$ , leading to the name “elimination form of inverse” [11]. However, it is wiser to eliminate the word “inverse” from all discussions of Simplex. We describe the updating methods in terms of *factorizations* of  $B_0$ ,  $B$ , and  $\bar{B}$ .

## 4 The Product-Form Update

We may factorize  $\bar{B}$  in (1) as  $\bar{B} = BT$ , where

$$Bv = \bar{a}, \quad T = I + (v - e_p)e_p^T.$$

Note that  $T$  is the identity matrix with the  $p$ th column replaced by  $v$ . Most importantly,  $T$  is a *permuted triangle* and it is easy to solve a system  $Ty = d$ . After  $k$  updates we have

$$B_k = B_0T_1T_2 \dots T_k,$$

where  $B_0$  may be treated as a “black box”, and the factors  $T_j$  may be stored sequentially as a sequence of sparse vectors  $\{v_j\}$ .

Since  $B_k$  contains factorizations of all previous bases, it is clear that if any previous basis is ill-conditioned then the updated factorization involves ill-conditioned factors, even if  $B_k$  itself becomes well-conditioned. The PF update does not have the ability to “recover” if Simplex passes through one or more unfriendly bases.

A practical safeguard is to test each update matrix  $T$ . If the diagonal element  $v_p$  is small, the update should be abandoned because  $T$  will be ill-conditioned. For example,

$$\text{if } |v_p|/\|v\|_\infty \geq 10^{-5} \text{ then Update else Factorize } \bar{B}.$$

Such a test is “prudent”, but we cannot say that it makes the PF update a stable method. A tolerance nearer  $10^{-2}$  or  $10^{-1}$  would help, but then the update would be rejected rather often.


## 5 The Bartels-Golub-Reid Update

The need for a stable update was emphasized by Bartels and Golub [1], and a more efficient implementation was proposed in [2]. The key idea was to update LU factors of  $B$  by replacing a column of  $U$  and restoring  $U$  to triangular form using stabilized elimination matrices. The updates to  $L$  may be kept in product form, but  $U$  must be updated explicitly. A satisfactory sparse implementation, LA05, was developed by Reid [12], and a similar implementation is included in LUSOL [7, 14]. Reid later updated LA05 as LA15 [13].

To illustrate, suppose  $B = LU$  and note that  $\bar{B} = LU'$ , where  $U'$  is identical to  $U$  except for its  $p$ th column  $\bar{u}$ , which solves the system  $L\bar{u} = \bar{a}$ . If  $U'$  is already triangular, we are lucky—the new factors are on hand. More generally, the last nonzero element of  $\bar{u}$  will be in row  $\ell$  with  $\ell > p$ , and to restore triangularity we need to find an LU factorization of  $U'$ .

As a first step, let  $P$  be a cyclic permutation that moves the  $p$ th column and row of  $U'$  to position  $\ell$  and shifts the intervening columns and rows forward. We see that  $U'$  is upper triangular except for one *column*, while the permuted matrix

$U''$  is upper triangular except for one row:



We refer to the exceptional column and row as *spikes*. The circles remind us that the column and row spikes are *sparse*. We may be lucky again and find that the row spike has no subdiagonal nonzeros ( $U''$  will be already triangular), but in general we need to find a factorization  $U'' = \tilde{L}\tilde{U}$ .

Ideally the existing diagonals in rows  $p:l-1$  will be large enough to serve as pivots, but to ensure stability we must allow *row interchanges*. The factor  $\tilde{L}$  is constructed as a product of *stabilized elementary transformations* whose essential part is one of the  $2 \times 2$  matrices

$$M = \begin{pmatrix} 1 & \\ \mu & 1 \end{pmatrix} \quad \text{or} \quad \tilde{M} = \begin{pmatrix} & 1 \\ 1 & \mu \end{pmatrix},$$

where the choice is made to keep  $\mu$  bounded ( $|\mu| \leq \text{UpdateTo1} = 10, 5, \text{ or } 2$ , say) while optimizing sparsity to some extent. The updated factorization is then  $\bar{B} = \tilde{L}\tilde{U}$  with  $\bar{L} = \tilde{L}\tilde{L}$  (in product form) and  $\bar{U} = \tilde{U}$  (explicit).

The row operations involved in factorizing  $U''$  require a fairly complex data structure. In LA05 and LUSOL, the nonzeros in  $U$  are stored *row-wise* and initially in natural order (so that solving a system with  $U$  or  $U^T$  involves a sweep through contiguous memory). During updates, several rows may be modified when the column vector  $\bar{u}$  is inserted, and further rows may be modified during the factorization of  $U''$  (depending how often  $\tilde{M}$  is chosen). Modifications are made in place when possible. If too many nonzeros are generated in a particular row, the row must be moved to the beginning of free storage. Periodic garbage collection reclaims the unused storage.

In the context of Bartels-Golub-Reid (BGR) updates, the initial LU factorization should ensure that  $L$  is well-conditioned. With `UpdateTo1` in the range  $[1, 10]$ , the modified  $L$  factors tend to be well-conditioned also. Although the storage of  $U$  becomes increasingly fragmented, the benefit is that hundreds of BGR updates can be carried out safely. The decision to refactorize rather than update may be based on storage concerns rather than loss of precision.

## 6 The Forrest-Tomlin Update

The update proposed by Forrest and Tomlin in 1972 [5] followed Bartels and Golub in updating  $L$  in product form and  $U$  explicitly, but with emphasis on dealing efficiently with *sparse* LU factors. Algebraically, the FT update is equivalent to the BGR update with the restriction that  $\ell = m$  and row interchanges are not allowed in the factorization  $U'' = \tilde{L}\tilde{U}$ .

Note that  $\ell = m$  means  $P$  moves the column spike to the end of  $U''$ , where it is no longer a spike. Also, “no row interchanges” means the row spike is simply deleted from the existing  $U$ . Hence, the FT update can be implemented efficiently with a *column-oriented* data structure for  $U$ .

To derive the FT update directly, let  $\delta$  be the  $p$ th diagonal of  $U$  and consider the system  $U^T r = \delta e_p$ . Then  $U$  and  $r$  have the following structure:

$$U = \begin{pmatrix} U_1 & u & U_2 \\ & \delta & v^T \\ & & U_3 \end{pmatrix}, \quad r = \begin{pmatrix} 0 \\ 1 \\ s \end{pmatrix},$$

where  $U_3^T s = -v$ . We can now obtain a triangular factorization  $U = R\tilde{U}$  in which most of the  $p$ th row of  $U$  is eliminated:

$$R = \begin{pmatrix} I & & \\ & 1 & -s^T \\ & & I \end{pmatrix}, \quad \tilde{U} = \begin{pmatrix} U_1 & u & U_2 \\ & \delta & \\ & & U_3 \end{pmatrix}.$$

We already have  $L\bar{u} = \bar{a}$ , and solving  $LR\tilde{u} = \bar{a}$  alters only the  $p$ th element of  $\bar{u}$  to give  $\bar{\delta} = r^T \tilde{u}$ . With suitable permutations,  $\bar{u}$  and the modified factors follow:

$$\bar{u} = \begin{pmatrix} \bar{u}_1 \\ \bar{u}_p \\ \bar{u}_3 \end{pmatrix}, \quad \bar{L} = LR, \quad \bar{U} = \begin{pmatrix} U_1 & U_2 & \bar{u}_1 \\ & U_3 & \bar{u}_3 \\ & & \bar{\delta} \end{pmatrix}.$$

In sequence, the steps are

- Solve  $L\bar{u} = \bar{a}$ .
- Find  $\delta = U_{pp}$  and solve  $U^T r = \delta e_p$ .
- Set  $\bar{\delta} = r^T \bar{u}$ .
- Delete the  $p$ th column and row of  $U$ .
- Define  $\bar{L} = LR$  and insert the new last column with  $\bar{\delta}$  to define  $\bar{U}$ .

Bearing in mind that this is a restricted form of the BGR update, we know that the numerical stability of the FT update depends on the size of the elements of  $s$ . (These are the multipliers  $\mu$  from the equivalent sequence of  $2 \times 2$  matrices  $M$ .) Thus, a practical safeguard could take the form

$$\text{if } \|s\|_\infty \leq 10^5 \text{ then Update else Factorize } \bar{B}.$$

As with the PF update, such a test is prudent but not completely reliable unless the tolerance is chosen much closer to 1. Nevertheless, the ease of implementation (and the existence of some sort of stability check) means that the FT update has been adopted in virtually *all* sparse implementations of Simplex. Efficiency wins!

## 7 The Block-LU Update

As before, suppose  $B_0$  is a given  $m \times m$  starting basis, and  $B_k$  is the same as  $B_{k-1}$  except for one column. The columns of each  $B_k$  are distinct columns of a larger matrix  $A$ . In general,  $B_k$  is a matrix in which the columns of  $V_k \equiv (v_1 \ v_2 \ \dots)$  have replaced columns  $p_1, p_2, \dots$  of  $B_0$ . If a sparse LU factorization

$$B_0 = L_0 U_0 \quad (2)$$

is available, the solution of  $B_k x = b$  may be obtained from a larger system

$$\begin{pmatrix} B_0 & V_k \\ E_k^T & \end{pmatrix} \begin{pmatrix} y \\ z \end{pmatrix} = \begin{pmatrix} b \\ 0 \end{pmatrix}, \quad (3)$$

where the columns of  $E_k$  are columns  $p_1, p_2, \dots$  of the  $m \times m$  identity. The larger system may be solved using the block-triangular factorization

$$\begin{pmatrix} B_0 & V_k \\ E_k^T & \end{pmatrix} = \begin{pmatrix} L_0 & \\ Z_k^T & I \end{pmatrix} \begin{pmatrix} U_0 & Y_k \\ & C_k \end{pmatrix}, \quad (4)$$

where we need

$$L_0 Y_k = V_k, \quad U_0^T Z_k = E_k, \quad C_k = -Z_k^T Y_k \quad (= -E_k^T B_0^{-1} V_k).$$

Thus,  $Y_k$  and  $Z_k$  are solutions of sparse systems and may be stored column-wise as  $m \times k$  sparse matrices, while  $C_k$  is  $k \times k$  and may be treated as a dense matrix (or sparse if we prefer—it often is). Note that  $C_k$  is the Schur complement of  $B_0$  in the matrix on the left of (4), but it is more important to regard (4) as a block-LU factorization.

### 7.1 Solving with $B_k$ and $B_k^T$

From (2)–(4), we can solve  $B_k x = b$  by the following sequence:

$$\begin{aligned} &\text{Solve } L_0 w = b. \\ &\text{Solve } C_k z = -Z_k^T w. \\ &\text{Solve } U_0 y = w - Y_k z. \\ &\text{Extract } x \text{ from the appropriate parts of } \begin{pmatrix} y \\ z \end{pmatrix}. \end{aligned} \quad (5)$$

Similarly, we can solve  $B_k^T y = c$  as follows:

$$\begin{aligned} &\text{Permute } \begin{pmatrix} c \\ 0 \end{pmatrix} \text{ to } \begin{pmatrix} c_1 \\ c_2 \end{pmatrix} \text{ to match rows of } B_k^T. \\ &\text{Solve } U_0^T w = c_1. \\ &\text{Solve } C_k^T z = c_2 - Y_k^T w. \\ &\text{Solve } L_0^T y = w - Z_k z. \end{aligned} \quad (6)$$

## 7.2 Updating $Y_k, Z_k, C_k$

Suppose the basis matrices  $B_k$  arise when column  $a_s$  of  $A$  replaces column  $a_r$  for a sequence of sensible pairs  $(s, r)$ . There are four updates to deal with:

1. Usually  $a_s$  is a new column of  $A$  (not in  $B_0$  or  $V_k$ ) and  $a_r$  is column of  $B_0$ . Then  $Y_k, Z_k, C_k$  expand.
2. Sometimes  $a_s$  is a new column and  $a_r$  is a column of  $V_k$ . All dimensions remain the same.
3. Sometimes  $a_s$  is a column of  $B_0$  that was replaced earlier, and  $a_r$  is another column of  $B_0$ . All dimensions remain the same.
4. Sometimes  $a_s$  is a column of  $B_0$  that was replaced earlier, and  $a_r$  is a column of  $V_k$ . Then  $Y_k, Z_k, C_k$  shrink.

## 7.3 Origins

The block-LU (BLU) update above was first described in [6] (where it was called the Schur-complement update). It was motivated by the work of Bisschop and Meeraus [3], who used the related block factorization

$$\begin{pmatrix} B_0 & V_k \\ E_k^T & I \end{pmatrix} = \begin{pmatrix} B_0 & \\ E_k^T & I \end{pmatrix} \begin{pmatrix} I & B_0^{-1}V_k \\ & C_k \end{pmatrix} \quad (7)$$

without storing  $B_0^{-1}V_k$  (at the expense of additional solves with  $B_0$  and  $B_0^T$ ). The Schur complement  $C_k$  is the same as before. In [3],  $C_k^{-1}$  was updated explicitly.

A more stable alternative is to update a factorization of  $C_k$ . Our Fortran 77 code LUMOD [14] maintains a dense factorization  $L_k C_k = U_k$  when rows and columns of  $C_k$  are added, deleted or replaced, as needed in section 7.2. Note that  $L_k$  is *square* and well-conditioned, and  $U_k$  is upper triangular. LUMOD stores both factors row-wise in 1D arrays. It was designed for future implementations of the BLU update.

## 7.4 QPBLU

An initial factorization  $B_0 = L_0 U_0$  may be used to define “ $L_0$ ” and “ $U_0$ ” in (5)–(6) in several combinations:

	“ $L_0$ ”	“ $U_0$ ”
1	$L_0 U_0$	$I$
2	$I$	$L_0 U_0$
3	$L_0$	$U_0$

In [4], combination 1 was implemented on a Cray Y-MP and found to be faster than our LUSOL implementation of the Bartels-Golub-Reid update. This combination gives  $Z_k = E_k$  (very sparse) and  $B_0 Y_k = V_k$  as in (7). The vectors  $Y_k$  are rather dense (as for the Product-Form update), but it does save time to store them and avoid double solves with  $B_0$  and  $B_0^T$ .

Combination 3 is likely to keep both  $Y_k$  and  $Z_k$  sparse. A first hint was given in [9], and extensive results are given in Hanh Huynh's thesis [8] using her Fortran 95 QP solver QPBLU. Given a semidefinite diagonal matrix  $H$ , QPBLU is designed to solve convex QP problems of the form

$$\begin{aligned} \min_x \quad & c^T x + \frac{1}{2} x^T H x \\ & Ax = b, \quad l \leq x \leq u. \end{aligned}$$

The active-set method factorizes the current KKT matrix

$$K_0 = \begin{pmatrix} H_0 & A_0^T \\ A_0 & \end{pmatrix}$$

and then applies BLU updates as rows and columns are symmetrically added and deleted. Sparse factors  $K_0 = L_0 U_0$  are obtained from the linear system packages LUSOL, MA57, PARDISO, SuperLU, or UMFPACK (where all except PARDISO can treat  $L_0$  and  $U_0$  separately). To maintain  $C_k$  for the BLU updates, LUMOD has been reimplemented as `lumod.f95`. Further BLU updates are used to accommodate QP Hessians of the form

$$H_\ell = R_\ell^T H_0 R_\ell, \quad \text{where} \quad R_\ell = (I + u_1 v_1^T)(I + u_2 v_2^T) \dots (I + u_\ell v_\ell^T)$$

for a sequence of vector pairs  $(u_1, v_1), (u_2, v_2), \dots$ , as required by the nonlinear optimizer SNOPT for its sequence of convex QP subproblems. We expect QPBLU to complement SQOPT for subproblems involving thousands of degrees of freedom (where  $A_0$  has many more columns than rows).

## 7.5 QPBLUR

On some of the test problems reported in [8], QPBLU encountered KKT systems that were essentially singular. To date, no ‘‘KKT repair’’ procedure has been implemented in QPBLU. Instead, we have considered *primal and dual regularization* of the QP problem itself. A new solver QPBLUR is being developed by Chris Maes to solve strictly convex QP problems of the form

$$\begin{aligned} \min_{x,y} \quad & c^T x + \frac{1}{2} x^T H x + \frac{1}{2} \delta \|x\|_2^2 + \frac{1}{2} \mu \|y\|_2^2 \\ & Ax + \mu y = b, \quad l \leq x \leq u, \end{aligned}$$

where  $\delta$  and  $\mu$  are small scalars of order  $10^{-6}$  to  $10^{-12}$ . Given any  $x$  that satisfies the bounds  $l \leq x \leq u$ , we can choose  $y$  to satisfy the constraints  $Ax + \mu y = b$ . Thus, a single-phase active-set algorithm becomes possible. Search directions are computed from KKT systems of the form

$$\begin{pmatrix} H_{FR} + \delta I & A_{FR}^T \\ A_{FR} & -\mu I \end{pmatrix} \begin{pmatrix} \Delta x_{FR} \\ -\Delta y \end{pmatrix} = - \begin{pmatrix} g_{FR} - A_{FR}^T y \\ 0 \end{pmatrix}.$$

A major benefit is that the KKT system is nonsingular for *any active set*. The BLU updates from QPBLU can still be used. Preliminary results with a MATLAB implementation of QPBLUR are encouraging [10].

## References

- [1] R. H. Bartels and G. H. Golub. The simplex method of linear programming using the LU decomposition. *Commun. ACM*, 12:266–268, 1969.
- [2] R. H. Bartels, J. Stoer, and Ch. Zenger. A realization of the simplex method based on triangular decompositions. In J. H. Wilkinson and C. Reinsch, editors, *Handbook for Automatic Computation, Volume II*, pages 219–239. Springer Verlag, Berlin, Heidelberg, New York, 1971.
- [3] J. Bisschop and A. Meeraus. Matrix augmentation and structure preservation in linearly constrained control problems. *Math. Program.*, 18:7–15, 1980.
- [4] S. K. Eldersveld and M. A. Saunders. A block-LU update for large-scale linear programming. *SIAM J. Matrix Anal. Appl.*, 13:191–201, 1992.
- [5] J. J. H. Forrest and J. A. Tomlin. Updating triangular factors of the basis to maintain sparsity in the product form simplex method. *Math. Program.*, 2:263–278, 1972.
- [6] P. E. Gill, W. Murray, M. A. Saunders, and M. H. Wright. Sparse matrix methods in optimization. *SIAM J. Sci. and Statist. Comput.*, 5:562–589, 1984.
- [7] P. E. Gill, W. Murray, M. A. Saunders, and M. H. Wright. Maintaining LU factors of a general sparse matrix. *Linear Algebra and its Applications*, 88/89:239–270, 1987.
- [8] H. M. Huynh. *A Large-scale Quadratic Programming Solver Based On Block-LU Updates of the KKT System*. PhD thesis, SCCM, Stanford University, 2008.
- [9] H. M. Huynh and M. A. Saunders. A convex QP solver based on block-LU updates. Presented at SIAM Conference on Parallel Processing for Scientific Computing, San Francisco, CA, February 22–24, 2006.
- [10] C. Maes and M. A. Saunders. An active-set convex QP solver based on regularized KKT systems. Presented at BIRS Workshop 09w5101, Advances and Perspectives on Numerical Methods for Saddle Point Problems, Banff, Alberta, Canada, Apr 12-17, 2009. <http://www.stanford.edu/group/SOL/talks.html>.
- [11] H. M. Markowitz. The elimination form of the inverse and its application to linear programming. *Management Science*, 3(3):255–269, 1957.
- [12] J. K. Reid. A sparsity-exploiting variant of the Bartels-Golub decomposition for linear programming bases. *Math. Program.*, 24:55–69, 1982.
- [13] J. K. Reid. LA15 basis factorization package (thread-safe version of LA05). HSL 2004 Catalogue, <http://www.aspentech.com/hsl/hslnav/hsl2004.htm>, 2004.
- [14] SOL downloadable software. <http://www.stanford.edu/group/SOL/software.html>.