

 Open access • Journal Article • DOI:10.1016/0167-8191(94)90079-5

## Updating URV decompositions in parallel — [Source link](#)

G. W. Stewart

**Institutions:** University of Maryland, College Park

**Published on:** 01 Apr 1992 - Parallel Computing

**Topics:** Matrix decomposition, Symmetric matrix, Block matrix, LU decomposition and QR decomposition

Related papers:

- [An updating algorithm for subspace tracking](#)
- [Updating a Generalized URV Decomposition](#)
- [Computing low-dimensional signal subspaces](#)
- [Updating a rank-revealing ULV decomposition](#)
- [Direction-of-arrival estimation using the rank-revealing URV decomposition](#)

Share this paper:    

View more about this paper here: <https://typeset.io/papers/updating-urv-decompositions-in-parallel-8liw22pwp5>

**UPDATING URV DECOMPOSITIONS IN PARALLEL**

By

**G.W. Stewart**

**IMA Preprint Series # 986**

May 1992

# Updating URV Decompositions in Parallel\*

G. W. Stewart<sup>†</sup>

April 1992

## ABSTRACT

A URV decomposition of a matrix is a factorization of the matrix into the product of a unitary matrix ( $U$ ), an upper triangular matrix ( $R$ ), and another unitary matrix ( $V$ ). In [8] it was shown how to update a URV decomposition in such a way that it reveals the effective rank of the matrix. It was also argued that the updating procedure could be implemented in parallel on a linear array of processors; however, no specific algorithms were given. This paper gives a detailed implementation of the updating procedure.

---

\*This report is available by anonymous ftp from [thales.cs.umd.edu](http://thales.cs.umd.edu) in the directory `pub/reports`.

<sup>†</sup>Department of Computer Science and Institute for Advanced Computer Studies, University of Maryland, College Park, MD 20742. This work was supported in part by the Air Force Office of Scientific Research under Contract AFOSR-87-0188. Part of it was done while the author was visiting the Institute for Mathematics and Its Applications, University of Minnesota, Minneapolis, MN 55455.

Updating URV Decompositions  
in Parallel  
G. W. Stewart

**1. Introduction**

A matrix is of effective rank  $k$  if a small perturbation is sufficient to make it of rank  $k$  but a large perturbation is required to make it of rank  $k-1$ . Matrices that are effectively rank degenerate occur in a number of applications, in which it is required to determine the effective rank and to compute an orthonormal basis for its effective null space of the matrix (e.g., see [9]). Of course, the notions of “large” and “small”, which define the effective rank, depend on the application.

If the  $n \times p$  matrix  $X$  is of effective rank  $k$ , there are orthogonal matrices  $U$  and  $V$  such that

$$U^T X V = \begin{pmatrix} R & F \\ 0 & G \\ 0 & 0 \end{pmatrix}, \quad (1.1)$$

where  $R$  and  $G$  are upper triangular matrices of orders  $k$  and  $p-k$  and  $F$  and  $G$  are small. Such a decomposition is called a *rank-revealing URV decomposition*. If  $V = (V_1 \ V_2)$  is partitioned conformally with (1.1), then

$$U^T X V_2 = \begin{pmatrix} F \\ G \end{pmatrix}$$

is small, so that the columns of  $V_2$  furnish an orthonormal basis for the effective null space of  $X$ .

An example of a rank-revealing URV decomposition is the singular value decomposition [1], in which  $R$  and  $G$  are diagonal and  $F$  is zero. However, the singular value decomposition is expensive to compute and often furnishes more information than is needed to solve the problem at hand. By stepping back from diagonality we obtain a triangular decomposition that is cheaper to compute and easy to update.

The problem of updating occurs when rows are added to  $X$ . If  $z^T$  is a new row, then the updating problem is to determine a rank revealing URV decomposition of

$$\begin{pmatrix} X \\ z^T \end{pmatrix}$$

from that of  $X$ . Since

$$\begin{pmatrix} U^T & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} X \\ z^T \end{pmatrix} V = \begin{pmatrix} R & F \\ 0 & G \\ 0 & 0 \\ x^T & y \end{pmatrix},$$

where  $(x^T \ y^T) = z^T(V_1 \ V_2)$ , the problem reduces to one of computing a rank-revealing URV factorization of

$$\begin{pmatrix} R & F \\ 0 & G \\ x^T & y \end{pmatrix} \quad (1.2)$$

from the matrix

$$\begin{pmatrix} R & F \\ 0 & G \end{pmatrix}.$$

In [8] the author described an algorithm requiring  $O(p^2)$  time for updating (1.2). The algorithm alternates between two steps:

Update: Reduce (1.2) to upper triangular form, preserving as far as possible the small elements of  $F$  and  $G$ .

Deflate: If the resulting matrix  $R$  is defective in rank, find its rank-revealing URV decomposition.

In the same paper it was shown that the algorithm could be made to run on a linear array of  $p$  processors in  $O(p)$  time; however, no specific algorithms were given. One purpose of this paper is to provide the missing algorithms.

Another purpose is to illustrate a style of programming. Our model of computation will be fine-grain MIMD; that is, short messages interleaved with short computations without global control. Fine-grain SIMD algorithms have an extensive literature, as do coarse-grain MIMD algorithms (e.g., see [5,6]). Less attention has been paid to fine-grain MIMD algorithms, chiefly because machines to run them have not been widely available. Now that such machines as the IWARP are in production, it is appropriate to illustrate coding techniques with a completely new and fairly complex set of algorithms.

The paper is organized as follows. In the next section we will describe the model of computation we will use. In Section 3 we show how to use precedence diagrams to pass from sequential to parallel algorithms. The parallel updating algorithm will be described piece by piece in the remaining sections. Although we will give brief sketches of the sequential algorithms, the reader should look to [8] for background and details.

## 2. Architectural Details

The way a parallel algorithm is implemented depends on the architecture of the system for which it is intended. In this paper we will work with a linear MIMD array that is capable of fine-grained communication. The number of processors,  $p$ , will be equal to the order of the matrix in question, and the matrices will be stored row-wise, one row to each processor. These are not the only possible choices, but they are well suited to matrices of moderate order, say thirty, such as arise in applications like signal processing.

The programming style will be SPMD; that is, each processor will run the same program with different data. A processor is identified by its position in the array, or equivalently by the row of the matrix it contains. This identifier is denoted by  $I$  in the programs to follow.

Communications is effected by bidirectional communication channels between processors. In the programs we will use the following conventions. We will imagine the processors as being arranged in a vertical line, north to south. The function call (all code is in C)

```
fget(NORTH, &x);
```

assigns the variable  $x$  of type `float` the value in the north input register. If no input is available, the program is blocked until input appears. Similarly the statement

```
fget(SOUTH, &x);
```

gets input from the channel to the south. The statements

```
fput(NORTH, x);  
fput(SOUTH, x);
```

send output to the north and south respectively. If there is already output in the channel, the program is blocked until the channel becomes free. We make no explicit assumptions about the speed of communication; but if the programs are to remain balanced, the time taken for two processors to exchange a floating-point number should be commensurate with the time required for a floating-point operation. This style of computing has the flavor of computations on a systolic array; however, it differs from it in the fact that no global synchronization of the processors is required. To give in a name we will call it *quasi-systolic*.

Instead of presenting entire programs, we will break the code into more manageable fragments. The fragments are from programs that have been debugged on a simulator that forks a process for each processor in the array. The communications functions are implemented by manipulating queues in shared memory. It has been shown [4] that if a program works on such a simulator it will work on a truly parallel system.

### 3. Precedence Diagrams and Parallel Code

It is not a trivial matter to pass from sequential code for a matrix algorithm to the corresponding quasi-systolic code. One obvious reason is that a quasi-systolic program must take on the additional responsibility of passing data between processors. However, a more subtle reason is that the two kinds of code look at the computations from different points of view. The sequential program follows the computations as they pass through the matrix, a global view analogous to the Lagrangian approach to fluid flow. The




Figure 3.1: Precedence Diagram for a Triangular System

quasi-systolic program takes the Eulerian approach; it sits at one part in the matrix and manipulates data as it flows by.

Precedence diagrams are a useful intermediary between sequential and parallel code. Consider, for example, the following sequential code for solving the  $p \times p$  upper triangular system  $Rw = b$ .

**Fragment 3.1.** Sequential solution of  $Rw = b$ .

```

for (i=p; i>=1; i--){
    w[i] = 0.
    for (j=p; j>i; j--)
        w[i] = w[i] + r[i][j]*w[j];
    w[i] = (b[i] - w[i])/r[i][i];
}

```

A transition diagram for this code is given in Figure 3.1.

The matrix is stored row-wise on the linear array. The nodes of the diagram are associated with matrix elements and they represent a computation involving nearby elements. The computation is specified by a symbol, whose meaning is given in the legend below the diagram. Thus the nodes labeled  $*$  are part of the computation of the inner product in the sequential program Fragment 3.1, while the nodes labeled  $\div$  produce the components of the solution.

The arrows specify the order in which the calculation must occur. Thus the  $\div$  operations in the diagram must proceed the  $*$  operations above them. The numbers represent times at which the calculation can occur, and they must be distinct for each node on a processor. The vertical arrows are doubled to indicate that they involve communication across processors—in this case the passing of the components  $w[j]$  of the solution. To keep the numbers in the diagram to a reasonable size, a computation shares a number with its input; but the output from a computation has a higher number.

The easy part of the code implementing a precedence diagram like the one in Figure 3.1 is a loop that follows the arrows across a row—gathering data, computing, and communicating the results. The hard part is handling edge conditions, which invariably have special communication requirements. For example, if the direction of communication is generally north, then the first processor must not send output, and the last must not request input. Again, each row in a diagram will usually have a special node—the  $\div$  node in Figure 3.1—that must set things up for the next processor in line.

In the following code the  $I$ th processor contains the  $I$ th row of the matrix  $R$  in a singly subscripted array named  $r$ .

**Fragment 3.2.** Quasi-systolic solution of  $Rw = b$ .

```
w[I] = 0.;
for (j=p; j>I; j--){
    fget(SOUTH, &w[j]);
    w[I] = w[I] + r[j]*w[j];
    if (I != 1)
        fput(NORTH, w[j]);
}
w[I] = (b - w[I])/r[I];
if (I != 1)
    fput(NORTH, w[I]);
```

The  $I$ th right-hand side  $b$  is contained in  $b$ , while the components of the solution are contained in an array  $w$ . At the end of the program the entire solution is on the first processor, ready for distribution to the other processors, if it is needed.

It is worth noting that by replacing the horizontal arrows with double arrows, we obtain a precedence diagram suitable for a square array of processors, in which each processor is responsible for one matrix element. This will be true of many of the precedence diagrams to follow. The additional processors do not result in a corresponding savings in time since the longest path through the precedence graph is  $O(p)$ . However, on a square array it is often possible to pipeline such computations, something that is not possible on a linear array.



BEFORE							
1	2	3	4	5	6	7	8
X	X	X	0	0	X	E	0
X	$\checkmark$	X	X	0	E	E	E

AFTER							
1	2	3	4	5	6	7	8
X	X	X	X	0	X	E	E
X	0	X	X	0	X	E	E

Figure 4.1: Application of a Plane Rotation



#### 4. Plane Rotations and Simple Updating<sup>1</sup>

Most of the computation in URV updating involves the use of plane rotations to introduce zeros into the matrix  $R$ . Since treatments of plane rotations are widely available (e.g., see [1]), we will not go into the numerical details here. Instead we will sketch the few basic facts needed to understand our algorithms and introduce some conventions for describing reductions based on plane rotations.

Figure 4.1 shows two rows of a matrix before and after the application of a plane rotation. The X's represent nonzero elements, the 0's represent zero elements, and the E's represent small elements. The plane rotation has been chosen to introduce a zero into the position occupied by the checked X in column 2. When the rotation is applied the following rules hold.

1. A pair of X's remains a pair of X's (columns 1 and 3).
2. An X and an 0 are replaced by a pair of X's (column 4).
3. A pair of 0's remains a pair of 0's (column 5).
4. An X and an E are replaced by a pair of X's (column 6).
5. A pair of E's remains a pair of E's (column 7).
6. An E and an 0 are replaced by a pair of E's (column 8).

The fact that a pair of small elements remains small (columns 7 and 8) follows from the fact that a plane rotation is orthogonal and cannot change the norm of any vector to which it is applied. This is a key observations, since the point of the updating algorithm is to keep small elements small.

---

<sup>1</sup>This section draws heavily on a similar section in [8].

$$\begin{array}{cccc}
\rightarrow & r & r & r & r \\
& 0 & r & r & r \\
& 0 & 0 & r & r \\
& 0 & 0 & 0 & r \\
\rightarrow & \check{x} & x & x & x
\end{array}
\Rightarrow
\begin{array}{cccc}
& r & r & r & r \\
\rightarrow & 0 & r & r & r \\
& 0 & 0 & r & r \\
& 0 & 0 & 0 & r \\
\rightarrow & 0 & \check{x} & x & x
\end{array}
\Rightarrow
\begin{array}{cccc}
& r & r & r & r \\
& 0 & r & r & r \\
& 0 & 0 & r & r \\
& 0 & 0 & 0 & r \\
\rightarrow & 0 & 0 & \check{x} & x
\end{array}$$
  

$$\begin{array}{cccc}
& r & r & r & r \\
& 0 & r & r & r \\
& 0 & 0 & r & r \\
\rightarrow & 0 & 0 & 0 & r \\
\rightarrow & 0 & 0 & 0 & \check{x}
\end{array}
\Rightarrow
\begin{array}{cccc}
& r & r & r & r \\
& 0 & r & r & r \\
& 0 & 0 & r & r \\
& 0 & 0 & 0 & r \\
& 0 & 0 & 0 & 0
\end{array}$$

Figure 4.2: Simple Updating

Premultiplication by a plane rotation operates on the rows of the matrix. We will call such rotations *left rotations*. Postmultiplication by *right rotations* operates on the columns. Analogous rules hold for the application of a right rotation to two columns of a matrix.

When rotations are used to update a URV decomposition, the right rotations must be multiplied into  $V$ . To get a complete update of the decomposition, we must also multiply the left rotations into  $U$ . However, in many applications  $U$  is not needed, and this step can be omitted. In this paper we will not give code for updating  $U$ .

Algorithms that use plane rotations are best described by pictures. To fix our conventions, we will show how the matrix

$$\begin{pmatrix} R \\ x^T \end{pmatrix},$$

where  $R$  is upper triangular, can be reduced to upper triangular form by left rotations. This updating procedure occurs twice in the following algorithms and will be called *simple updating*.

The reduction is illustrated in Figure 4.2. The elements of  $R$  and  $x^T$  are represented generically by  $r$ 's and  $x$ 's. The first step in the reduction is to eliminate the first element of  $x^T$  by a rotation that acts on  $x^T$  and the first row of  $R$ . The element to be eliminated has a check over it and the two rows that are being combined are indicated by the arrows to the left of the array.

According to this notation, the second step combines the second row of  $R$  with  $x^T$  to eliminate the second element of the latter. Note that  $r_{21}$ , which is zero, forms a pair

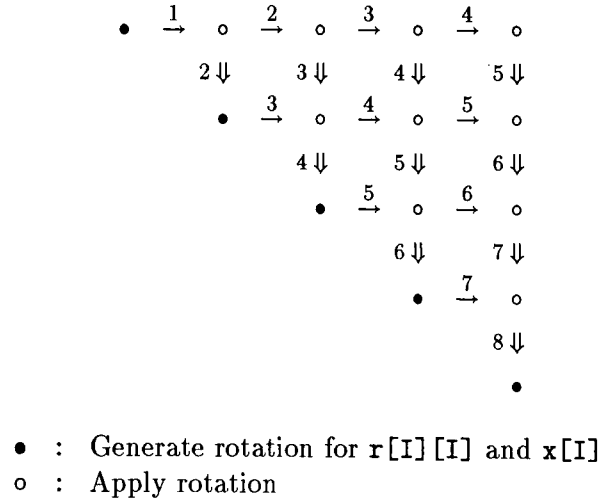


Figure 5.1: Precedence Diagram for Simple Updating

of zeros with the first component of  $x^T$ , so that the zero we introduced in the first step is not destroyed in the second step. The third and fourth steps of the reduction are similar.

For column operations with right rotations we will use an analogous notation. The main difference is that the arrows will point down to the columns being combined.

## 5. Simple Updating in Parallel

Figure 5.1 contains a precedence diagram for simple updating. The diagram assumes that the vector  $x^T$  lies in the first processor. After a component of  $x^T$  has been combined with the row on one processor, it drops to the next to be combined with the next row.

The following fragment implements the algorithm.

**Fragment 5.1.** Simple updating.

```

for (j=1; j<=p; j++){
  if (I == 1)
    xx = x[j];
  else
    fget(NORTH, &xx);
  if (j == I)

```

```

        rotgen(&r[j], &xx, &c, &s);
    else{
        rotapp(&r[j], &xx, &c, &s);
        if (I != p)
            fput(SOUTH, xx);
    }
}

```

The function `rotgen(&r, &x, &c, &s)` generates a rotation from  $\mathbf{r}$  and  $\mathbf{x}$ , returning  $\sqrt{r^2 + x^2}$  in  $\mathbf{r}$  and zero in  $\mathbf{x}$ . The function `rotapp(&r, &x, &c, &s)` applies the rotation whose cosine and sine are  $c$  and  $s$  to  $\mathbf{r}$  and  $\mathbf{s}$ , returning the altered values of  $\mathbf{r}$  and  $\mathbf{x}$ .

## 6. Updating

We now turn to the problem of updating the matrix (1.2). The first step is to compute  $z^T V$ . We will assume that  $z^T$  is situated in the topmost processor. We begin by passing the  $I$ th component down the array until it reaches processor  $I$ , where it is stored in  $z_i$ . This passing of data is simple enough not to require a precedence diagram, and we give only the code.

**Fragment 6.1.** Distribution of  $z^T$ .

```

for (j=p; j>I; j--){
    if (I == 1)
        temp = z[j];
    else
        fget(NORTH, &temp);
    fput(SOUTH, temp);
}
if (I == 1)
    zi = z[1];
else
    fget(NORTH, &zi);

```


Once the components of  $z^T$  are in place, they are multiplied into the components of the rows of  $V$ , the products being passed down and summed on the way. The precedence diagram is given in Figure 6.1. The corresponding fragment is

**Fragment 6.2.** Formation of  $z^T V$ .

```

for (j=1; j<=p; j++){
    if (I == 1)
        sum = 0;
    else

```

Figure 6.1: Precedence Diagram for  $z^T V$ 

```

      fget(NORTH, &sum);
sum = sum + zi*v[j];
if (I != p)
  fput(SOUTH, sum);
else
  x[j] = sum;
}

```

Finally the vector  $x^T$ , which is now on the  $p$ th processor, is distributed to all the processors:

**Fragment 6.3.** Distribution of  $x^T$

```

for (j=1; j<=p; j++){
  if (I != p)
    fget(SOUTH, &x[j]);
  if (I != 1)
    fput(NORTH, x[j]);
}

```

The next step in the updating algorithm depends on the size of the vector  $y^T$  in (1.2); that is, the vector whose components are  $x[k+1]$ ,  $x[k+2]$ ,  $\dots$ ,  $x[p]$ . If simple updating is performed on (1.2), the result will be a new matrix

$$\begin{pmatrix} \hat{R} & \hat{F} \\ 0 & \hat{G} \end{pmatrix},$$

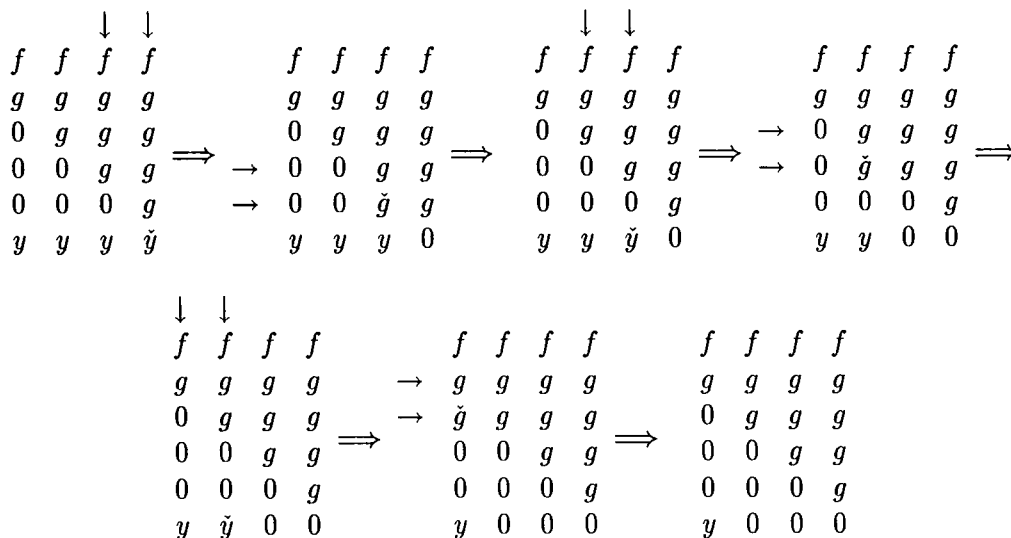


Figure 6.2: Reduction of  $y^T$


where  $\|\hat{F}\|_F^2 + \|\hat{G}\|_F^2 = \|F\|_F^2 + \|G\|_F^2 + \|y^T\|_F^2$ . Consequently, if  $y^T$  is small enough, the updated matrix will be of effective rank not greater than  $k$ . In this case, simple updating followed by the deflation procedure to be described later is what we want to do.

However, if  $y^T$  is sufficiently large, we must take into the account a possible increase in effective rank. Here simple updating will not do, since it will fold the large elements of  $y^T$  into  $F$  and  $G$  and destroy the rank-revealing character of the decomposition. The cure is to preprocess the matrix so that only the first component of  $y^T$  is nonzero while  $G$  remains upper triangular; for then simple updating will make at most the  $(k + 1)$ th column large.

The reduction is illustrated in Figure 6.2. Left rotations introduce zeros into the component of  $y$ . In the process each rotation introduces a nonzero element below the diagonal of  $G$ , which is removed by a right rotation.

The parallel implementation of this reduction is the trickiest part of the algorithm. The reason is that the left and right rotations must be interleaved. We will proceed in two stages.

Figure 6.3 gives a precedence diagram for part of the reduction. The array begins with the  $(k + 1, k + 1)$ -element. The right rotations come from the south and are applied to  $R$  and  $V$  before being passed north. The application of the rotations continues



- : 1. Generate left rotation for rows  $I$  and  $I+1$
- : 2. Generate right rotation for columns  $I-1$  and  $I$
- : 3. Apply right rotation to  $R$  and  $V$

Figure 6.3: Precedence Diagram for the Reduction of  $y^T$ , Part 1

through the matrix  $F$  up to row one (not shown in the diagram). Here is the code.

**Fragment 6.4.** Reduction of  $y^T$ , Part 1.

```

for (j=p; j>=(k+2>I+1?k+2:I+1); j--){
  fget(SOUTH, &cr[j]);
  fget(SOUTH, &sr[j]);
  rotapp(&r[j-1], &r[j], &cr[j], &sr[j]);
  rotapp(&v[j-1], &v[j], &cr[j], &sr[j]);
  if (I != 1){
    fput(NORTH, cr[j]);
    fput(NORTH, sr[j]);
  }
}

```

To generate the left rotation, processor  $I$  requires the element  $r_{I,I+1}$  from the south. To generate the right rotation, processor  $I$  requires the current value of  $y_I$  from the south. It is important that the corresponding information be sent north in the proper order.

**Fragment 6.5.** Reduction of  $y^T$ , Part 1 continued.

```

if (I!=p && I>k){
  fget(SOUTH, &rinf);
  rotgen(&r[I], &rinf, &cl, &sl);
}
if (I > k+1){
  if (I != p)
    fget(SOUTH, &x[I]);
  rotgen(&x[I-1], &x[I], &cr[I], &sr[I]);
  fput(NORTH, cr[I]);
  fput(NORTH, sr[I]);
  rotapp(&r[I-1], &r[I], &cr[I], &sr[I]);
  fput(NORTH, r[I-1]);
  r[I-1] = 0.;
  if (I != k+2)
    fput(NORTH, x[I-1]);
}

```

We now turn to the application of the left rotations. Since a rotation can be applied to a pair of elements only after rotations have been applied to all inferior pairs, some care must be taken to get an  $O(p)$  algorithm. The trick is to apply a rotation and allow the superior processor to apply the same rotation before going on to the next rotation. Figure 6.4 contains a precedence diagram for this process. The array begins at the  $(k+1, k+1)$ -element. The double arrows indicate that an element must be transmitted and then returned. Code implementing this diagram follows.

**Fragment 6.6.** Reduction of  $y^T$ , Part 2.

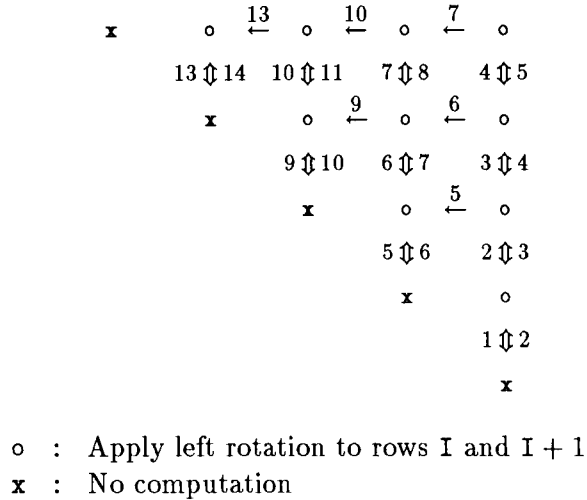
```

if (I > k){
  for (j=p; j>=I; j--){
    if (j != I){
      fget(SOUTH, &rinf);
      rotapp(&r[j], &rinf, &cl, &sl);
      fput(SOUTH, rinf);
    }
    if (I != k+1){
      fput(NORTH, r[j]);
      fget(NORTH, &r[j]);
    }
  }
}

```

It remains to apply the right rotations to the part of  $V$  lying on and below the diagonal. The code is simple and requires no particular explanation.



Figure 6.4: Reduction of  $y^T$ , Part 2

**Fragment 6.7.** Reduction of  $y^T$ , Part 2 continued.

```

for (j=I; j>k+1; j--){
  if (j != I){
    fget(NORTH, &cr[j]);
    fget(NORTH, &sr[j]);
  }
  rotapp(&v[j-1], &v[j], &cr[j], &sr[j]);
  if (I != p){
    fput(SOUTH, cr[j]);
    fput(SOUTH, sr[j]);
  }
}

```

The updating process is completed with a simple update. Consequently, we must move the reduced vector to the topmost processor. Note that only nonzero (i.e., the first  $k + 1$ ) components have to be moved. The follow code accomplishes this task.

**Fragment 6.8.** Moving  $x^T$  to the top.

```

if (I==1)
  for (j=k+2; j<=p; j++)
    x[j] = 0;
if (I < k+2 && I!=p)

```

```

    fget(SOUTH, &x[k+1]);
    if (I<=k+2 && I!=1)
        fput(NORTH, x[k+1]);

```

## 7. Deflation

An increase in effective rank will show up in the updating process. A decrease in effective rank is also possible; however, it is not as easy to detect. This section is devoted to algorithms to detect a decrease in effective rank and make the URV decomposition reveal it, a process we will call *deflation*.

The procedure is as follows. First determine a vector  $w$  of norm one such that  $b = Rw$  is as small as possible, or at least nearly so. If  $\|b\|$  is greater than an (application dependent) tolerance, declare  $R$  to be of full rank and stop. Otherwise, determine orthogonal matrices  $P$  and  $Q$  such that  $Q^T w = \mathbf{e}_k$ , where  $\mathbf{e}_k$  is the  $k$ th unit vector, and  $P^T RQ$  is upper triangular. It then follows that

$$\|b\| = \|Rw\| = \|(P^T RQ)Q^T w\| = \|(P^T RQ)\mathbf{e}_k\|,$$

or equivalently that the norm of the last column of the upper triangular matrix  $P^T RQ$  is equal to  $\|b\|$ , which is small. Thus,  $P^T RQ$  is a rank-revealing URV decomposition of  $R$ .

The determination of the vector  $w$  falls under the rubric of condition estimation (for a survey see [3]). Here we give a very simple condition estimator [2]. The idea is to solve the system  $R\hat{w} = b$ , where the components of  $b$  are  $\pm 1$ , the sign being chosen to maximize the growth of  $\hat{w}$ . The vector  $w$  is given by  $\hat{w}/\|\hat{w}\|$ . The following code is adapted from Fragment 3.2.

**Fragment 7.1.** Condition estimation.

```

    if (I <= k){
        w[I] = 0;
        for (j=k; j>I; j--){
            fget(SOUTH, &w[j]);
            w[I] = w[I] + r[j]*w[j];
            if (I != 1)
                fput(NORTH, w[j]);
        }
        w[I] = ((w[I]<0 ? 1 : -1) - w[I])/r[I];
        if (I != 1)
            fput(NORTH, w[I]);
    }

```

The next step is to compute the reciprocal of the norm of  $\hat{w}$  and distribute it through the array, so that each processor can decide whether deflation is necessary. The function `nrm2` computes the norm.

$$\begin{array}{cccc}
\downarrow & \downarrow & & \\
r & r & r & r \left| \begin{array}{l} \rightarrow w \\ \rightarrow \check{w} \\ \rightarrow w \\ \rightarrow w \end{array} \right. \\
0 & r & r & r \\
0 & 0 & r & r \\
0 & 0 & 0 & r
\end{array}
\Rightarrow
\begin{array}{cccc}
\downarrow & \downarrow & & \\
r & r & r & r \left| \begin{array}{l} \rightarrow w \\ \rightarrow \check{w} \\ \rightarrow w \\ \rightarrow w \end{array} \right. \\
0 & 0 & r & r \\
0 & 0 & 0 & r
\end{array}
\Rightarrow
\begin{array}{cccc}
\downarrow & \downarrow & & \\
r & r & r & r \left| \begin{array}{l} \rightarrow w \\ \rightarrow \check{w} \\ \rightarrow w \\ \rightarrow w \end{array} \right. \\
0 & r & r & r \\
0 & 0 & r & r \\
0 & 0 & 0 & r
\end{array}$$
  

$$\begin{array}{cccc}
\downarrow & \downarrow & & \\
r & r & r & r \left| \begin{array}{l} \rightarrow 0 \\ \rightarrow 0 \\ \rightarrow w \\ \rightarrow \check{w} \end{array} \right. \\
r & r & r & r \\
0 & r & r & r \\
0 & 0 & 0 & r
\end{array}
\Rightarrow
\begin{array}{cccc}
\downarrow & \downarrow & & \\
r & r & r & r \left| \begin{array}{l} \rightarrow 0 \\ \rightarrow 0 \\ \rightarrow 0 \\ \rightarrow 1 \end{array} \right. \\
0 & r & r & r \\
0 & 0 & r & r \\
0 & 0 & 0 & r
\end{array}$$

Figure 7.1: Reduction of  $\hat{w}$ 

**Fragment 7.2.** Distribution of  $\|\hat{w}\|$ .

```

if (I == 1){
    *rnormw = sqrt((double) k)/nrm2(k, w);
} else
    fget(NORTH, rnormw);
if (I != p)
    fput(SOUTH, *rnormw);

```

Assuming a deflation is necessary, the next step is to reduce  $w$  to a multiple of  $e_k$  by plane rotations. Since  $w$  and  $\hat{w}$  produce the same sequence of rotations, we work with the latter. The rotations are post multiplied into  $R$  to produce  $RQ$ , which is upper Hessenberg. The process is illustrated in Figure 7.1. A precedence diagram is given in Figure 7.2.


We assume that  $\hat{w}$  is in the first row, where the rotations are generated. They are then passed south, where they are applied to  $R$  and  $V$ , as in the following code.

**Fragment 7.3.** Reduction of  $\hat{w}$ .

```

for (j=2; j<=k; j++){
    if (I == 1)
        rotgen(&w[j], &w[j-1], &c[j], &s[j]);
    else{
        fget(NORTH, &c[j]);
        fget(NORTH, &s[j]);
    }
}
if (I <= k)

```



- : Generate and apply rotation, for  $w[j-1]$  and  $w[j]$
- o : Apply rotation to  $r$  and  $v$
- v : Apply rotation to  $v$
- x : No computation

Figure 7.2: Precedence Diagram for Reduction of  $\hat{w}$ 

```

    rotapp(&r[j], &r[j-1], &c[j], &s[j]);
rotapp(&v[j], &v[j-1], &c[j], &s[j]);
if (I != p){
    fput(SOUTH, c[j]);
    fput(SOUTH, s[j]);
}
}

```

The matrix  $R$  is now upper Hessenberg, and the final step of the deflation is to reduce the subdiagonal elements to zero. This process is shown in Figure 7.3.

A precedence diagram is given in Figure 7.4. As always with rotations that cross processor boundaries, it is important for a processor feed the next one before continuing to do its own thing.


**Fragment 7.4.** Reduction of Hessenberg  $R$ .

```

if (I > k) return;
for (j=I-1; j<=p; j++){
    if (I != 1){
        fget(NORTH, &rsup);
    }
}

```

$$\begin{array}{cccc}
\rightarrow & r & r & r & r \\
\rightarrow & \tilde{r} & r & r & r \\
& 0 & r & r & r \\
& 0 & 0 & r & r
\end{array}
\Rightarrow
\begin{array}{cccc}
r & r & r & r \\
0 & r & r & r \\
0 & \tilde{r} & r & r \\
0 & 0 & r & r
\end{array}
\Rightarrow
\begin{array}{cccc}
r & r & r & r \\
0 & r & r & r \\
0 & 0 & r & r \\
0 & 0 & \tilde{r} & r
\end{array}
\Rightarrow
\begin{array}{cccc}
r & r & r & r \\
0 & r & r & r \\
0 & 0 & r & r \\
0 & 0 & 0 & r
\end{array}$$

Figure 7.3: Reduction of Hessenberg  $R$ 

- : Generate rotation for rows  $I-1$  and  $I$
- : Apply rotation
- x : No computation

Figure 7.4: Precedence Diagram for Reduction of Hessenberg  $R$ 


```

if (j == I-1)
    rotgen(&rsup, &r[j], &c, &s);
else
    rotapp(&rsup, &r[j], &c, &s);
fput(NORTH, rsup);
}

if (j!=I-1 && I!=k){
    fput(SOUTH, r[j]);
}


```





- : Generate rotation for  $r[I][I]$ ,  $r[I][k]$ .
- : Apply rotation.
- x : Pass  $r[k][k]$  north.

Figure 8.2: Precedence Diagram for Reducing the Last Column



```

if (I < k){
  for (j=k-1; j>I; j--){
    fget(SOUTH, &c);
    fget(SOUTH, &s);
    fget(SOUTH, &rk[j]);
    if (I != 1){
      fput(NORTH, c);
      fput(NORTH, s);
      fput(NORTH, rk[j]);
    }
    rotapp(&r[j], &r[k], &c, &s);
    rotapp(&v[j], &v[k], &c, &s);
  }
  rotgen(&r[I], &r[k], &c, &s);
  rotapp(&v[I], &v[k], &c, &s);
  fget(SOUTH, &rk[k]);
  rk[I] = 0.;
  rotapp(&rk[I], &rk[k], &c, &s);
  if (I != 1){
    fput(NORTH, c);
    fput(NORTH, s);
    fput(NORTH, rk[I]);
    fput(NORTH, rk[k]);
  }
  for (j=k+1; j<=p; j++){
    fget(SOUTH, &rk[j]);
    if (I != 1)
      fput(NORTH, rk[j]);
  }
  if (I != p){
    fput (SOUTH, c);
    fput (SOUTH, s);
  }
}

```

After the reduction, the right rotations, which are in the first processor, are applied to  $V$  as in Fragment 6.7.

The last row of  $R$ , which is now nonzero is in the first processor. This row is reduced to zero by a variant of the simple updating algorithm Fragment 5.1. This restores  $R$  to upper triangular form, after which  $k$  is reduced by one and the new, smaller  $R$  is checked for effective deficiency in rank.



## 9. Observations and Conclusions

The code presented in this paper has a finality that belies the work involved in actually writing it down. In this section I would like to make some informal observations on coding and debugging quasi-systolic algorithms.

Unless the hardware and software on your target system are completely reliable, it is a good idea to do preliminary debugging on a simulator. A simulator is easy to write using standard system routines. Since it runs on a single machine, it is easy to insert debugging statements and direct output to wherever you wish. Moreover, if a put-get discipline, such as the one in this paper, is used to synchronize computations, successful execution on a simulator guarantees that the code will run in parallel.

In most cases the starting point for parallel code is a sequential algorithm. If you do not understand the latter there is little hope of your producing the former. Begin by coding and debugging the sequential version (unless of course you already understand it well).

I found precedence diagrams to be extremely useful in passing from the sequential algorithm to parallel code. However, there is a learning curve here, and toward the end of the project my diagrams degenerated into impressionistic sketches.

Unless you are careful, you will find you are spending more time debugging your test cases than your code. It is worth your while to put hard thought into devising tests whose correctness can be recognized without elaborate post-processing. In running test cases, all the usual caveats about testing boundary conditions apply. In particular, it is easy to write code that performs adequately when communication is well buffered but fails when the buffer queues have length one.

Most of the standard debugging tools do not work well with simulators, and you will probably be reduced to inserting explicit debugging output in your code. If you understand the sequential algorithm and your tests are well constructed, the main problems will be with the flow of your algorithm. The vast majority of my debugging output consisted of statements like

`<Proc id>: I am here <position id>`

Although the code given here assumes only one row is allocated to each processor, it would not be a difficult task to extend it to the case where several consecutive rows are assigned to each processor.

I will conclude by expressing my belief that quasi-systolic algorithms implemented on linear arrays will become increasingly important in the solution of medium size, dense matrix problems. Such problems arise regularly in real-time control and signal processing applications, where they must be solved with dispatch. A linear array insures that the hardware can be both manageably small and inexpensive. The MIMD implementation of quasi-systolic algorithms means that the same system can be adapted to

solve a wide variety of problems.

### References

- [1] G. H. Golub and C. F. Van Loan. *Matrix Computations*. Johns Hopkins University Press, Baltimore, Maryland, 2nd edition, 1989.
- [2] W. B. Gragg and G. W. Stewart. A stable variant of the secant method for solving nonlinear equations. *SIAM Journal on Numerical Analysis*, 13:880–903, 1976.
- [3] N. J. Higham. A survey of condition number estimation for triangular matrices. *SIAM Review*, 29:575–596, 1987.
- [4] D. P. O’Leary and G. W. Stewart. From determinacy to systaltic arrays. *IEEE Transactions on Computers*, C-36:1355–1359, 1987.
- [5] J. M. Ortega, R. G. Voigt, and C. H. Romine. *A Bibliography on Parallel and Vector Numerical Algorithms*. Technical Report TM-10998, Oak Ridge National Laboratory, 1989.
- [6] Y. Robert. *The Impact of Vector and Parallel Architectures on the Gaussian Elimination Algorithm*. Halsted Press, New York, 1990.
- [7] G. W. Stewart. *On an Algorithm for Refining a Rank-Revealing URV Decomposition and a Perturbation Theorem for Singular Values*. Technical Report CS-TR 2626, Department of Computer Science, University of Maryland, 1991. To appear in *Linear Algebra and Its Applications*.
- [8] G. W. Stewart. *An Updating Algorithm for Subspace Tracking*. Technical Report CS-TR 2494, Department of Computer Science, University of Maryland, 1990. To appear in *IEEE Transactions on Signal Processing*.
- [9] R. J. Vaccaro, editor. *SVD and Signal Processing, II*, Elsevier Science Publishers, Amsterdam, 1990.

## Recent IMA Preprints

#	Author/s	Title
892	<b>E.G. Kalnins, Willard Miller, Jr. and Sanchita Mukherjee,</b>	Models of $q$ -algebra representations: the group of plane motions
893	<b>T.R. Hoffend Jr. and R.K. Kaul,</b>	Relativistic theory of superpotentials for a nonhomogeneous, spatially isotropic medium
894	<b>Reinhold von Schwerin,</b>	Two metal deposition on a microdisk electrode
895	<b>Vladimir I. Oliker and Nina N. Uraltseva,</b>	Evolution of nonparametric surfaces with speed depending on curvature, III. Some remarks on mean curvature and anisotropic flows
896	<b>Wayne Barrett, Charles R. Johnson, Raphael Loewy and Tamir Shalom,</b>	Rank incrementation via diagonal perturbations
898	<b>Mingxiang Chen, Xu-Yan Chen and Jack K. Hale,</b>	Structural stability for time-periodic one-dimensional parabolic equations
899	<b>Hong-Ming Yin,</b>	Global solutions of Maxwell's equations in an electromagnetic field with the temperature- dependent electrical conductivity
900	<b>Robert Grone, Russell Merris and William Watkins,</b>	Laplacian unimodular equivalence of graphs
901	<b>Miroslav Fiedler,</b>	Structure-ranks of matrices
902	<b>Miroslav Fiedler,</b>	An estimate for the nonstochastic eigenvalues of doubly stochastic matrices
903	<b>Miroslav Fiedler,</b>	Remarks on eigenvalues of Hankel matrices
904	<b>Charles R. Johnson, D.D. Olesky, Michael Tsatsomeros and P. van den Driessche,</b>	Spectra with positive elementary symmetric functions
905	<b>Pierre-Alain Gremaud,</b>	Thermal contraction as a free boundary problem
906	<b>K.L. Cooke, Janos Turi and Gregg Turner,</b>	Stabilization of hybrid systems in the presence of feedback delays
907	<b>Robert P. Gilbert and Yongzhi Xu,</b>	A numerical transmutation approach for underwater sound propagation
908	<b>LeRoy B. Beasley, Richard A. Brualdi and Bryan L. Shader,</b>	Combinatorial orthogonality
909	<b>Richard A. Brualdi and Bryan L. Shader,</b>	Strong hall matrices
910	<b>Håkan Wennerström and David M. Anderson,</b>	Difference versus Gaussian curvature energies; monolayer versus bilayer curvature energies applications to vesicle stability
911	<b>Shmuel Friedland,</b>	Eigenvalues of almost skew symmetric matrices and tournament matrices
912	<b>Avner Friedman, Bei Hu and J.L. Velazquez,</b>	A Free Boundary Problem Modeling Loop Dislocations in Crystals
913	<b>Ezio Venturino,</b>	The Influence of Diseases on Lotka-Volterra Systems
914	<b>Steve Kirkland and Bryan L. Shader,</b>	On Multipartite Tournament Matrices with Constant Team Size
915	<b>Richard A. Brualdi and Jennifer J.Q. Massey,</b>	More on Structure-Ranks of Matrices
916	<b>Douglas B. Meade,</b>	Qualitative Analysis of an Epidemic Model with Directed Dispersion
917	<b>Kazuo Murota,</b>	Mixed Matrices Irreducibility and Decomposition
918	<b>Richard A. Brualdi and Jennifer J.Q. Massey,</b>	Some Applications of Elementary Linear Algebra in Combinations
919	<b>Carl D. Meyer,</b>	Sensitivity of Markov Chains
920	<b>Hong-Ming Yin,</b>	Weak and Classical Solutions of Some Nonlinear Volterra Integrodifferential Equations
921	<b>B. Leinkuhler and A. Ruehli,</b>	Exploiting Symmetry and Regularity in Waveform Relaxation Convergence Estimation
922	<b>Xinfu Chen and Charles M. Elliott,</b>	Asymptotics for a Parabolic Double Obstacle Problem
923	<b>Yongzhi Xu and Yi Yan,</b>	An Approximate Boundary Integral Method for Acoustic Scattering in Shallow Oceans
924	<b>Yongzhi Xu and Yi Yan,</b>	Source Localization Processing in Perturbed Waveguides
925	<b>Kenneth L. Cooke and Janos Turi,</b>	Stability, Instability in Delay Equations Modeling Human Respiration
926	<b>F. Bethuel, H. Brezis, B.D. Coleman and F. Hélein,</b>	Bifurcation Analysis of Minimizing Harmonic Maps Describing the Equilibrium of Nematic Phases Between Cylinders
927	<b>Frank W. Elliott, Jr.,</b>	Signed Random Measures: Stochastic Order and Kolmogorov Consistency Conditions
928	<b>D.A. Gregory, S.J. Kirkland and B.L. Shader,</b>	Pick's Inequality and Tournaments
929	<b>J.W. Demmel, N.J. Higham and R.S. Schreiber,</b>	Block $LU$ Factorization
930	<b>Victor A. Galaktionov and Juan L. Vazquez,</b>	Regional Blow-Up in a Semilinear Heat Equation with Convergence to a Hamilton-Jacobi Equation
931	<b>Bryan L. Shader,</b>	Convertible, Nearly Decomposable and Nearly Reducible Matrices
932	<b>Dianne P. O'Leary,</b>	Iterative Methods for Finding the Stationary Vector for Markov Chains
933	<b>Nicholas J. Higham,</b>	Perturbation theory and backward error for $AX - XB = C$
934	<b>Z. Strakos and A. Greenbaum,</b>	Open questions in the convergence analysis of the lanczos process for the real symmetric eigenvalue problem
935	<b>Zhaojun Bai,</b>	Error analysis of the lanczos algorithm for the nonsymmetric eigenvalue problem
936	<b>Pierre-Alain Gremaud,</b>	On an elliptic-parabolic problem related to phase transitions in shape memory alloys
937	<b>Bojan Mohar and Neil Robertson,</b>	Disjoint essential circuits in toroidal maps

- 939 **Bojan Mohar and Svatopluk Poljak** Eigenvalues in combinatorial optimization
- 940 **Richard A. Brualdi, Keith L. Chavey and Bryan L. Shader**, Conditional sign-solvability
- 941 **Roger Fosdick and Ying Zhang**, The torsion problem for a nonconvex stored energy function
- 942 **René Ferland and Gaston Giroux**, An unbounded mean-field intensity model:  
Propagation of the convergence of the empirical laws and compactness of the fluctuations
- 943 **Wei-Ming Ni and Izumi Takagi**, Spike-layers in semilinear elliptic singular Perturbation Problems
- 944 **Henk A. Van der Vorst and Gerard G.L. Sleijpen**, The effect of incomplete decomposition preconditioning  
on the convergence of conjugate gradients
- 945 **S.P. Hastings and L.A. Peletier**, On the decay of turbulent bursts
- 946 **Apostolos Hadjidimos and Robert J. Plemmons**, Analysis of  $p$ -cyclic iterations for Markov chains
- 947 **ÅBjörck, H. Park and L. Eldén**, Accurate downdating of least squares solutions
- 948 **E.G. Kalnins, Willard Miller, Jr. and G.C. Williams**, Recent advances in the use of separation of  
variables methods in general relativity
- 949 **G.W. Stewart**, On the perturbation of LU, Cholesky and QR factorizations
- 950 **G.W. Stewart**, Gaussian elimination, perturbation theory and Markov chains
- 951 **G.W. Stewart**, On a new way of solving the linear equations that arise in the method of least squares
- 952 **G.W. Stewart**, On the early history of the singular value decomposition
- 953 **G.W. Stewart**, On the perturbation of Markov chains with nearly transient states
- 954 **Umberto Mosco**, Composite media and asymptotic dirichlet forms
- 955 **Walter F. Mascarenhas**, The structure of the eigenvectors of sparse matrices
- 956 **Walter F. Mascarenhas**, A note on Jacobi being more accurate than QR
- 957 **Raymond H. Chan, James G. Nagy and Robert J. Plemmons**, FFT-based preconditioners for  
Toeplitz-Block least squares problems
- 958 **Zhaojun Bai**, The CSD, GSVD, their applications and computations
- 959 **D.A. Gregory, S.J. Kirkland and N.J. Pullman**, A bound on the exponent of a primitive matrix using  
Boolean rank
- 960 **Richard A. Brualdi, Shmuel Friedland and Alex Pothén**, Sparse bases, elementary vectors and nonzero  
minors of compound matrices
- 961 **J.W. Demmel**, Open problems in numerical linear algebra
- 962 **James W. Demmel and William Gragg**, On computing accurate singular values and eigenvalues of acyclic  
matrices
- 963 **James W. Demmel**, The inherent inaccuracy of implicit tridiagonal QR
- 964 **J.J.L. Velázquez**, Estimates on the  $(N - 1)$ -dimensional Hausdorff measure of the blow-up set  
for a semilinear heat equation
- 965 **David C. Dobson**, Optimal design of periodic antireflective structures for the Helmholtz equation
- 966 **C.J. van Duijn and Joseph D. Fehribach**, Analysis of planar model for the molten carbonate fuel cell
- 967 **Yongzhi Xu, T. Craig Poling and Trent Brundage**, Source localization in a waveguide with unknown  
large inclusions
- 968 **J.J.L. Velázquez**, Higher dimensional blow up for semilinear parabolic equations
- 969 **E.G. Kalnins and Willard Miller, Jr.**, Separable coordinates, integrability and the Niven equations
- 970 **John M. Chadam and Hong-Ming Yin**, A diffusion equation with localized chemical reactions
- 971 **A. Greenbaum and L. Gurvits**, Max-min properties of matrix factor norms
- 972 **Bei Hu**, A free boundary problem arising in smoulder combustion
- 973 **C.M. Elliott and A.M. Stuart**, The global dynamics of discrete semilinear parabolic equations
- 974 **Avner Friedman and Jianhua Zhang**, Swelling of a rubber ball in the presence of good solvent
- 975 **Avner Friedman and Juan J.L. Velázquez**, A time-dependence free boundary problem modeling  
the visual image in electrophotography
- 976 **Richard A. Brualdi, Hyung Chan Jung and William T. Trotter, Jr.**, On the poset of all posets on  
 $n$  elements
- 977 **Ricardo D. Fierro and James R. Bunch**, Multicollinearity and total least squares
- 978 **Adam W. Bojanczyk, James G. Nagy and Robert J. Plemmons**, Row householder transformations for  
rank- $k$  Cholesky inverse modifications
- 979 **Chaocheng Huang**, An age-dependent population model with nonlinear diffusion in  $R^n$
- 980 **Emad Fatemi and Faroukh Odeh**, Upwind finite difference solution of Boltzmann equation applied to  
electron transport in semiconductor devices
- 981 **Esmond G. Ng and Barry W. Peyton**, A tight and explicit representation of  $Q$  in sparse  $QR$   
factorization
- 982 **Robert J. Plemmons**, A proposal for  $FFT$ -based fast recursive least-squares
- 983 **Anne Greenbaum and Zdenek Strakos**, Matrices that generate the same Krylov residual spaces
- 984 **Alan Edelman and G.W. Stewart**, Scaling for orthogonality
- 985 **G.W. Stewart**, Note on a generalized sylvester equation
- 986 **G.W. Stewart**, Updating URV decompositions in parallel