

Calhoun: The NPS Institutional Archive DSpace Repository

## ALGORITHM 686, FORTRAN Subroutines for Updating the QR Decomposition

Reichel, L.; Bragg, W.B.
https://hdl.handle.net/10945/38295

This publication is a work of the U.S. Government as defined in Title 17, United States Code, Section 101. Copyright protection is not available for this work in the United States.

Downloaded from NPS Archive: Calhoun


# ALGORITHM 686 FORTRAN Subroutines for Updating the QR Decomposition 

L. REICHEL<br>Bergen Scientific Center<br>and<br>W. B. GRAGG<br>Naval Postgraduate School


#### Abstract

Let the matrix $A \in \mathbf{R}^{m \times n}, m \geq n$, have a QR decomposition $A=\mathrm{QR}$, where $\mathrm{Q} \in \mathbf{R}^{m \times n}$ has orthonormal columns, and $R \in \mathbf{R}^{n \times n}$ is upper triangular. Assume that $Q$ and $R$ are explicitly known. We present FORTRAN subroutines that update the QR decomposition in a numerically stable manner when $A$ is modified by a matrix of rank one, or when a row or a column is inserted or deleted. These subroutines are modifications of the Algol procedures in Daniel et al. [5]. We also present a subroutine that permutes the columns of $A$ and updates the QR decomposition so that the elements in the lower right corner of R will generally be small if the columns of $A$ are nearly linearly dependent. This subroutine is an implementation of the rank-revealing $Q R$ decomposition scheme recently proposed by Chan [3]. Categories and Subject Descriptors: G.1.3 [Numerical Analysis]: Numerical Linear Algebra; G. 4 [Mathematics of Computing]: Mathematical Software

General Terms: Algorithms


Additional Key Words and Phrases: QR decomposition, subset selection, updating

## 1. INTRODUCTION

We present several FORTRAN subroutines for updating the QR decomposition of a matrix. Let $A \in \mathbf{R}^{m \times n}, m \geq n$, have a QR decomposition $A=\mathrm{QR}$, where $\mathrm{Q} \in \mathbf{R}^{m \times n}$ has orthonormal columns, and $\mathrm{R} \in \mathbf{R}^{n \times n}$ is upper triangular. Assume that the elements of Q and R are explicitly known. Let $\bar{A} \in \mathbf{R}^{p \times q}, p \geq q$, be obtained from $A$ by inserting or deleting a row or a column, or let $\bar{A}$ be a rankone modification of $A$ (i.e., $\bar{A}=A+v u^{T}$, where $u \in \mathbf{R}^{n}, v \in \mathbf{R}^{m}$ ). Then, a QR -decomposition of $\bar{A}, \bar{A}=\overline{\mathrm{Q}} \overline{\mathrm{R}}$, where $\overline{\mathrm{Q}} \in \mathbf{R}^{p \times q}$ has orthonormal columns and

[^0](C) 1990 ACM 0098-3500/90/1200-0369 $\$ 01.50$
$\overline{\mathrm{R}} \in \mathbf{R}^{q \times q}$ is upper triangular, can be computed in $O(m n)$ arithmetic operations by updating Q and R ; see Daniel et al. [5]. The updating is done by applying Givens reflectors. The operation count for updating $Q$ and $R$ compares favorably with the $O\left(m n^{2}\right)$ arithmetic operations necessary to compute a QR decomposition of a general $m \times n$ matrix.

Algol procedures for computing $\overline{\mathrm{Q}}$ and $\overline{\mathrm{R}}$ from Q and R are presented by Daniel et al. [5]. Buckley [2] translated these procedures into FORTRAN. Our FORTRAN subroutines implement modifications of the algorithms in [5]. These modifications reduce the operation count. Also, the structure of our code differs from [2]. We use subroutines such as the Basic Linear Algebra Subprograms (BLAS), described in [12], for simple matrix and vector operations. This makes it fairly easy to adapt the frequently executed parts of the code in order to achieve faster execution. The subroutines run quickly on an IBM 3090 VF vector computer.

Several program libraries, such as LINPACK [6] and NAG [15], provide subroutines for updating R, but contain no routines for updating the complete QR decomposition. Two advantages of updating both Q and R are that QR decomposition can be modified stably when a row or a column is deleted from $A$, and that the individual elements of projections are easily accessible; see LINPACK [6, p. 10.23], Daniel et al. [5], and Stewart [18].

The first comprehensive survey of updating algorithms was presented by Gill et al. [8]; and a recent discussion with references to applications can be found in Golub and Van Loan [10, Ch. 12.6]. The applications include linear least squares problems, regression analysis, and the solution of nonlinear systems of equations. (See Allen [1], Goldfarb [9], Gragg and Stewart [11], Moré and Sorensen [14].) The algorithms also appear to be applicable to recursive least squares problems of signal processing; see Ling et al. [13].

We also present a subroutine that implements the rank-revealing $Q R$ decomposition method recently proposed by Chan [3]. In this method the QR decomposition $A=\mathrm{QR}$ is updated to yield the QR decomposition $\bar{A}=\overline{\mathrm{Q}} \overline{\mathrm{R}}$, where $\bar{A}$ is obtained from $A$ by column permutation. This permutation is selected so that, in general, the element(s) in the lower right corner of $\overline{\mathrm{R}}$ are small if $A$ has nearly linearly dependent columns. The subroutine can be used to solve the subset selection problem; see Golub and Van Loan [10]. Table I lists the FORTRAN subroutines for updating the QR decomposition. All subroutines use double precision arithmetic and are written in FORTRAN 77. Section 2 contains programming details for the subroutines of Table I and for certain auxiliary subprograms. For all subroutines of Table I, except DRRPM, the numerical method as well as Algol procedures have been presented in [5]. For these subroutines we discuss differences between our FORTRAN subroutines and the Algol procedures only.

Section 3 discusses programming details for the subroutines for simple matrix and vector operations. These subroutines have been written to vectorize well on an IBM 3090 VF vector computer when the vectorizing compiler VS FORTRAN 2.3.0 is used. The last section, Section 4, contains some timing results for our subroutines and comparisons with the code [2].

Table I. Subroutines for Updating a QR Decomposition
$A=\mathrm{QR}$ to Yield a QR Decomposition $\bar{A}=\overline{\mathrm{Q}} \overline{\mathrm{R}}$
Subroutine $\quad$ Purpose

| DDECL | Computes $\bar{Q}, \overline{\mathrm{R}}$ from $\mathrm{Q}, \mathrm{R}$ when $\bar{A}$ is obtained from $A$ by deleting a column; see [5]. |
| :--- | :--- |
| DDELR | Computes $\overline{\mathrm{Q}}, \overline{\mathrm{R}}$ from $\mathrm{Q}, \mathrm{R}$ when $\bar{A}$ is obtained from $A$ by deleting a row; see $[5]$. |
| DINSC | Computes $\overline{\mathrm{Q}}, \overline{\mathrm{R}}$ from $\mathrm{Q}, \mathrm{R}$ when $\bar{A}$ is obtained from $A$ by inserting a column; see [5]. |
| DINSR | Computes $\overline{\mathrm{Q}}, \overline{\mathrm{R}}$ from $\mathrm{Q}, \mathrm{R}$ when $\bar{A}$ is obtained from $A$ by inserting a row; see $[5]$. |
| DRNK1 | Computes $\overline{\mathrm{Q}}, \overline{\mathrm{R}}$ from $\mathrm{Q}, \mathrm{R}$ when $\bar{A}$ is a rank-one modification of $A ;$ see [5]. |
| DRRPM | Computes $\overline{\mathrm{Q}}, \overline{\mathrm{R}}$ from $\mathrm{Q}, \mathrm{R}$ when $\bar{A}$ is obtained by permuting the columns of $A$ in a <br> manner that generally reveals if columns of $A$ are nearly linearly dependent; see [3]. |

## 2. THE UPDATING SUBROUTINES

We consider the subroutines of Table I in order. These subroutines use auxiliary subroutines, which we need to introduce first. They are listed in Table II.

### 2.1 Subroutines DORTHO and DORTHX

Given a matrix $Q \in \mathbf{R}^{m \times n}, m \geq n$, with orthonormal columns and a vector $w \in \mathbf{R}^{m}$, the subroutine DORTHO computes the Fourier coefficients $s:=\mathbf{Q}^{T} w$ and the orthogonal projection of $w$ into the null space of $\mathbf{Q}^{T}, v:=\left(I-\mathbf{Q Q}^{T}\right) w$. At most one reorthogonalization is carried out. Since the subroutine DORTHO differs from the corresponding Algol procedure "orthogonalize" in [5], we discuss DORTHO and its use in some detail.

Subroutine DORTHO is called by routine DINSC, which updates the QR factorization of a matrix $A=\mathrm{QR} \in \mathbf{R}^{m \times n}, m>n$, when a column $w$ is inserted into $A$. Updating may not be meaningful if $w$ is a nearly linear combination of the columns of $Q$. Therefore, DOR'IHO computes the condition number of the matrix $\tilde{\mathrm{Q}}:=[\mathrm{Q}, w /\|w\|] \in \mathbf{R}^{m \times(n+1)}$, where $\|\|$ is the Euclidean norm. Using $Q^{T} \mathrm{Q}=I$, we obtain the following expressions for the singular values $\sigma_{1} \geq \sigma_{2} \geq \cdots \geq \sigma_{n+1}$ of $\tilde{\mathrm{Q}}$ :

$$
\begin{align*}
\sigma_{1} & =\left(1+\left\|\mathbf{Q}^{T} w\right\| /\|w\|\right)^{1 / 2}  \tag{2.1a}\\
\sigma_{j} & =1, \quad 2 \leq j \leq n,  \tag{2.1b}\\
\sigma_{n+1} & =\left(1-\left\|\mathbf{Q}^{T} w\right\| /\|w\|\right)^{1 / 2} \tag{2.1c}
\end{align*}
$$

Further, for $v:=\left(I-\mathrm{QQ}^{t}\right) w /\|w\|$,

$$
\begin{equation*}
\|v\|=\sigma_{1} \sigma_{n+1} \tag{2.2}
\end{equation*}
$$

Since $1 \leq \sigma_{1} \leq \sqrt{2}, \sigma_{n+1}$ is also an accurate estimate of the length of the orthogonal projection of $w /\|w\|$ into the null space of $\mathbf{Q}^{T}$. In order to avoid severe cancellation of significant digits in (2.1c), we first determine $\sigma_{1}$ from (2.1a) and then $\sigma_{n+1}$ from (2.2).

Subroutines DINSC and DORTHO have an input parameter RCOND which is a lower bound for the reciprocal condition number. The computations are discontinued and an error flag is set if RCOND $>\sigma_{n+1} / \sigma_{1}$. On exit, RCOND : $=\sigma_{n+1} / \sigma_{1}$.

Table II. Auxiliary Subroutines

| $\begin{array}{c}\text { Auxiliary } \\ \text { subroutine }\end{array}$ | $\begin{array}{c}\text { Called by } \\ \text { subroutine }\end{array}$ | Purpose |
| :--- | :---: | :--- |$]$| DORTHO | DINSC, DRNK1 | Compute $s:=\mathrm{Q}^{T} w, v:=\left(I-\mathrm{QQ}^{T}\right) w$ with reorthogonalization <br> for arbitrary vector $w$. |
| :--- | :--- | :--- |
| DORTHX | DDELR | Compute $s:=\mathrm{Q}^{T} e_{j}, v:=\left(I-\mathrm{QQ}^{T}\right) e_{j}$, with reorthogonalization <br> for axis vector $e_{j}$. |
| DINVIT | DRRPM | Compute approximation of a right singular vector <br> corresponding to a least singular value of R. A first <br> approximation is obtained from the LINPACK condition <br> number estimator DTRCO, and is improved by inverse |
| DTRLSL | DINVIT | iteration. <br> Solve lower triangular system of equations with frequent <br> rescalings in order to avoid overflow. Similar to part of <br> DTRCO. |
|  | DINVIT | Solve upper triangular linear system of equations with <br> frequent rescalings in order to avoid overflow. Similar to <br> part of DTRCO. |

Assume now that the input value of $\mathrm{RCOND} \leq \sigma_{n+1} / \sigma_{1}$. Then DORTHO computes $s:=\mathrm{Q}^{T} w$ and $v:=\left(I-\mathrm{QQ}^{T}\right) w$ by a scheme analogous to the method described by Parlett [15, p. 107] for orthogonalizing a vector against another vector. For definiteness, we present the orthogonalization scheme. References to $\sigma_{1}, \sigma_{n+1}$, and RCOND are neglected for simplicity.

```
Orthogonalization Algorithm. Input \(\mathrm{Q} \in \mathbf{R}^{m \times n}\) ( Q has orthonormal
columns), \(m, n(m>n), w \in \mathbf{R}^{m}(w \neq 0)\); output \(v\left(v=\left(I-\mathbf{Q Q}^{T}\right) w\right)\),
\(s\left(s=\mathbf{Q}^{T} w\right) ;\)
    \(\check{w}:=w /\|w\| ;\)
    \(s:=\mathbf{Q}^{T} \tilde{w} ; v:=\tilde{w}-\mathbf{Q} s ;\)
    if \(\|v\| \geq 0.707\) then
        \(v:=v /\|v\| ; \quad s:=s\|w\| ; \quad\) exit; \(\quad *\|v\|=1, \quad \mathbf{Q}^{T} v=0 *\)
    \(s^{\prime}:=\mathbf{Q}^{T} v ; \quad v^{\prime}:=v-\mathrm{Q}^{\prime} ;\)
    if \(\left\|v^{\prime}\right\| \leq 0.707\|v\|\) then
        * \(w\) lies in \(\operatorname{span}\{\mathrm{Q}\}\) numerically *
    \(v:=0 ; s:=\left(s+s^{\prime}\right)\|w\| ;\) set flag; exit;
    \(v:=\left(v+v^{\prime}\right) /\left\|v+v^{\prime}\right\| ; \quad s:=\left(s+s^{\prime}\right)\|w\| ; \quad\) exit; \(\quad *\|v\|=1, \quad \mathrm{Q}^{T} v=0 *\)
```

The proof in Parlett [16, pp. 107-108], that one reorthogonalization suffices, carries over the present algorithm, using $\mathrm{Q}^{T} \mathrm{Q}=I$.

We note that there are other ways to carry out the computations on lines (2.3)-(2.4). In [5], $v$ and $v^{\prime}$ are updated immediately after a component of $s$ is computed. Our scheme has the advantage of being faster on vector computers because it allows matrix vector operations; it is also, generally, more accurate, since there is less accumulation of rounding errors. The latter advantage can be shown easily, but we omit the details.

We turn to subroutine DORTHX. This is a faster version of subroutine DORTHO. DORTHX assumes that $w$ in the orthogonalization algorithm is an axis vector. This simplifies the computations in (2.3). DORTHX may perform nearly twice as fast as DORTHO.

### 2.2 Subroutines DINVIT, DTRLSL, and DTRUSL

Given a nonsingular upper triangular matrix $U=\left[\mu_{j k}\right] \in \mathbf{R}^{n \times n}$ and a vector $b=\left[\beta_{j}\right] \in \mathbf{R}^{n}$, DTRUSL solves $U x=b \rho$, where $|\rho| \leq 1$ is a scaling factor such that $\left|\beta_{j} \rho / \mu_{j j}\right| \leq 1$ for all $j$. The scaling factor is introduced in order to avoid overflow when solving very ill-conditioned linear systems of equations. DTRLSL is an analogous subroutine for lower triangular systems.

DTRLSL and DTRUSL are called by DINVIT, a subroutine for computing an approximation of a right singular vector belonging to a least singular value of a right triangular matrix $R$. If $R$ is singular, then such a singular vector is computed by solving a triangular linear system of equations. Otherwise an initial approximate right singular vector $a^{(o)}=\left\{\alpha_{j}^{(o)}\right\}_{j=1}^{n}$ is obtained by the LINPACK condition number estimator DTRCO, and inverse iteration with $\mathrm{R}^{T} \mathrm{R}$ is used to obtain improved approximations $a^{(j)}, j=1,2, \ldots$, NMBI'I', where NMBIT is an input parameter to DINVIT and DRRPM. On exit from DINVIT and DRRPM, $\operatorname{IPOS}(j)$ contains the least index $k$ such that $\left|\alpha_{k}^{(j)}\right| \geq\left|\alpha_{l}^{(j)}\right|, 1 \leq l \leq n$, $0 \leq j \leq$ NMBIT. On return from DINVIT and DRRPM, the parameter DELTA is given by DELTA $:=\left\|\mathrm{R}^{T} \mathrm{R} a^{(\mathrm{NMBIT})}\right\| /\left\|a^{(\mathrm{NMBIT})}\right\|$. Hence, DELTA is an upper bound for the least singular value of R .

### 2.3 Updating Subroutines

We are now in a position to consider the subroutines of Table I. The vectorization is done mainly in the subroutines for simple matrix and vector operations, discussed in Section 3, but some loops of the subroutines of Table I vectorize as well. Comments in the source code reveal which loops vectorize or are eligible for vectorization on an IBM 3090 VF computer with compiler VS FORTRAN 2.3.0 when the default vectorization directives are used. For applications to particular problem classes, changing the default vectorization by compiler directives may decrease execution time.

We list the differences between the subroutines of Table I and the corresponding Algol procedures of [5]. Some of these modifications are suggested in [5], but not implemented in the Algol procedures. In subroutine DDELC, the column deleted in $A:=Q R$ is determined optionally. Not computing this column saves $O(m n)$ arithmetic operations. In subroutine DDELR, the auxiliary subroutine DORTHX is used instead of DORTHO. As indicated in Section 2.1, the former subroutine may perform nearly twice as fast. In subroutine DINSC, a column $w$ is inserted into $A:=\mathrm{QR}$ only if the reciprocal condition number of the matrix $[\mathrm{Q}, w /\|w\|]$ is larger than a bound given by the parameter RCOND on entry. The parameter RCOND can be used to prevent updating when $w /\|w\|$ is nearly in the range of Q. Finally, DRNK1 performs slightly faster if the updated matrix $A+v u^{T}$ is such that $v$ lies numerically in the range of $A$.

The subroutine DRRPM implements an algorithm presented by Chan [3]. The computation of an approximate right singular vector corresponding to a least singular value is done by subroutine DINVIT, and has already been discussed. The position of the component of largest magnitude of this singular vector has to be determined, and we found, in agreement with Chan's suggestion [3], that two inverse iterations suffice. In fact, in all computed examples, one inverse
iteration was sufficient, even for problems with multiple or close least singular values. The subroutine permutes the order of columns 1 through $k$ of $A \Pi$ where $k$ is an input parameter, $A \in \mathbf{R}^{m \times n}, m \geq n$, and $\Pi$ is a permutation matrix. Typically, DRRPM is called with $k=n, n-1, n-2, \ldots$ until no further permutation is made or until the computer upper bound DELTA for the least singular value of the matrix consisting of the first $k$ columns of $A \Pi$ is not small.

The subroutines of Table I neither require nor produce a factorization with nonnegative diagonal elements of the upper triangular matrix.

## 3. SUBROUTINES FOR SIMPLE MATRIX AND VECTOR OPERATIONS

Much computational experience on a variety of computers led Dongarra and Sorensen [7] to conclude that nearly optimal performance of numerical linear algebra code can be achieved if the subroutines for simple matrix and vector operations, such as BLAS [12], are written to perform on the computer at hand. We have written these subroutines so that they run efficiently on an IBM 3090 VF vector computer when the compiler VS FORTRAN 2.3 .0 with parameter vlev $=2$ is used. No vectorization directives are required in the code. In particular, we note that when implementing the code for updating the QR decomposition of a matrix on other (vector) computers, it may be possible to gain execution speed by modifying the subroutines for simple linear algebra tasks. We illustrate this with an example that describes a partricular feature of the VS FORTRAN 2.3 .0 compiler. By introducing a temporary scalar variable, denoted ACC in the subroutine DAPX in Example 3.1 below, this compiler generates code that avoids unnecessary loads and stores.

During execution, ACC should be thought of as a vector variable stored in a vector register. Timings for DAPX and comparison with code with explicitly unrolled loops have been carried out by Robert and Squazzero [17]. These timings show subroutine DAPX to perform better than equivalent subroutines with explicitly unrolled loops. However, when implemented on other computers with other compilers, the subroutine DAPX might perform better if the order of the DO-loops is interchanged.

Example 3.1 Subroutine for Matrix Vector Multiplication
SUBROUTINE DAPX(A,LDA,M,N,X,Y)
C
C DAPX COMPUTES Y:=A*X.
INTEGER LDA,M,N,I,J
REAL*8 A(LDA,N), X(N), Y(M), ACC
C
C OUTER LOOP VECTORIZES.
DO $10 \mathrm{I}=1, \mathrm{M}$
$\mathrm{ACC}=\mathrm{ODO}$ DO $20 \mathrm{~J}=1, \mathrm{~N}$
$\mathrm{ACC}=\mathrm{ACC}+\mathrm{A}(\mathrm{I}, \mathrm{J}) * \mathrm{X}(\mathrm{J})$
20 CONTINUE
$Y(I)=A C C$
10 CONTINUE
RETURN
END

Table III. Timings for DDELCO

|  |  | CPU time in seconds |  |  |
| ---: | :---: | :---: | :---: | :---: |
|  |  | Scalar | Vector <br> arithmetic | Scalar time <br> $n$ |
|  | $n$ | arithmetic | Vector time |  |
| 10 | 10 | $4 \cdot 10^{-4}$ | $4 \cdot 10^{-4}$ | 1 |
| 20 | 10 | $4 \cdot 10^{-4}$ | $4 \cdot 10^{-4}$ | 1 |
| 30 | 10 | $5 \cdot 10^{-4}$ | $4 \cdot 10^{-4}$ | 1.5 |
| 50 | 10 | $6 \cdot 10^{-4}$ | $4 \cdot 10^{-4}$ | 1.5 |
| 75 | 10 | $8 \cdot 10^{-4}$ | $4 \cdot 10^{-4}$ | 2 |
| 128 | 10 | $1 \cdot 10^{-3}$ | $5 \cdot 10^{-4}$ | 2 |
| 1024 | 10 | $7 \cdot 10^{-3}$ | $2 \cdot 10^{-3}$ | 3.5 |
| 1280 | 10 | $9 \cdot 10^{-3}$ | $3 \cdot 10^{-3}$ | 3.5 |

Table IV. Timings for DRNK1

|  |  | CPU time in seconds |  |  |
| ---: | :---: | :---: | :---: | :---: |
|  |  | Scalar <br> $m$ | Vector <br> arithmetic | Scalar time |
|  | $n$ | Vector time |  |  |
| 16 | 12 | $1 \cdot 10^{-3}$ | $1 \cdot 10^{-3}$ | $\mathbf{1}$ |
| 32 | 25 | $4 \cdot 10^{-3}$ | $3 \cdot 10^{-3}$ | $\mathbf{1 . 5}$ |
| 64 | 50 | $2 \cdot 10^{-2}$ | $7 \cdot 10^{-3}$ | 2 |
| 128 | 100 | $6 \cdot 10^{-2}$ | $2 \cdot 10^{-2}$ | 2.5 |
| 1024 | 100 | $4 \cdot 10^{-1}$ | $8 \cdot 10^{-2}$ | 4.5 |
| 1250 | 100 | $5 \cdot 10^{-1}$ | $9 \cdot 10^{-1}$ | 5 |

## 4. COMPUTED EXAMPLES

Example 4.1 Execution times for subroutines DDELCO and DRNK1 are compared for scalar and vector arithmetic. The measured CPU times differ somewhat between different executions of the same code. The times reported are therefore rounded to one significant digit and the quotient of measured CPU times rounded to the nearest multiple of $\frac{1}{2}$.

Table III shows the CPU times for DDELCO. This routine and its subroutines are compiled with the VS FORTRAN 2.3 .0 compiler. The times for vector arithmetic are obtained from code generated with compiler option vlev $=2$, which makes the compiler generate code that utilizes the vector registers and arithmetic. The times for scalar arithmetic are obtained from code generated with compiler option vlev $=0$, which makes the compiler generate code that does not use vector instructions. Given a QR decomposition of a matrix $A \in \mathbf{R}^{m \times n}$, Table III shows the CPU time required by DDELCO to compute the QR decomposition of $\bar{A} \in \mathbf{R}^{m \times(n-1)}$ obtained by deleting column one of $A$.

Table IV is similar to Table III, and contains execution times for DRNK1. The reduction in execution time obtained by using vector instructions is of the same order of magnitude for the other updating routines, too.

Example 4.2 Execution times for subroutines written by Buckley [2] and those of Table I are compared. The vectorized and scalar codes were generated as

Table V. The First Row of $A=\mathrm{QR}$ is Deleted. CPU Times for Vectorized Code for Updating Q and R are Given in Seconds; $n=10$.

|  | Time for | Time for | Time for DELROW [2] |
| ---: | :---: | :---: | :---: |
| $m$ | DDELR | DELROW [2] | Time for DDELR |
| 10 | $3 \cdot 10^{-5}$ | $9 \cdot 10^{-4}$ | $3.5 \cdot 10^{1}$ |
| 64 | $7 \cdot 10^{-5}$ | $2 \cdot 10^{-3}$ | $2 \cdot 10^{1}$ |
| 128 | $8 \cdot 10^{-4}$ | $3 \cdot 10^{-3}$ | 3 |
| 1024 | $4 \cdot 10^{-3}$ | $2 \cdot 10^{-2}$ | 4 |

Table VI. The Last Column of $A=Q R$ is deleted. CPU Times for Vectorized Code for Updating $Q$ and $R$ are Given in Seconds. DDELC Does Not Compute the Last Column of $A$. (IFLAG $=0$ on Entry).

| $m$ | $n$ | Time for DELC | $\begin{aligned} & \text { Time for } \\ & \text { DELCOL [2] } \end{aligned}$ | Time for DELCOL [2] |
| :---: | :---: | :---: | :---: | :---: |
|  |  |  |  | Time for DDELC |
| 1024 | 10 | 1. $10^{-4}$ | $2 \cdot 10^{-4}$ | 2 |
| 1024 | 100 | $1 \cdot 10^{-4}$ | $7 \cdot 10^{-3}$ | $7.5 \cdot 10^{1}$ |
| 1280 | 100 | $1 \cdot 10^{-4}$ | $9 \cdot 10^{-3}$ | $9.5 \cdot 10^{1}$ |

Table VII. A New First Column is Inserted into $A=\mathrm{QR}$. CPI Times for Vectorized Code for Updating Q and R are Given in Seconds; $n=10$.

| $m$ | Time for DINSC | Time for INSCOL [2] | Time for INSCOL [2] |
| :---: | :---: | :---: | :---: |
|  |  |  | Time for DINSC |
| 64 | $1 \cdot 10^{-3}$ | $2 \cdot 10^{-3}$ | 2 |
| 128 | 1. $10^{-3}$ | $3 \cdot 10^{-3}$ | 2 |
| 1024 | $7 \cdot 10^{-3}$ | $2 \cdot 10^{-2}$ | 2.5 |

Table VIII. The First Row of $A=Q R$ is Deleted. CPU Times for Scalar Code for Updating Q and R are Given in Seconds; $n=10$.

|  | Time for | Time for | Time for DELROW [2] |
| ---: | :---: | :---: | :---: |
| $m$ | DDELR | DELROW [2] | Time for DDELR |
| 10 | $2 \cdot 10^{-5}$ | $6 \cdot 10^{-4}$ | $3 \cdot 10^{1}$ |
| 64 | $1 \cdot 10^{-3}$ | $2 \cdot 10^{-3}$ | 1.5 |
| 128 | $2 \cdot 10^{-3}$ | $3 \cdot 10^{-3}$ | 1.5 |
| 1024 | $1 \cdot 10^{-2}$ | $2 \cdot 10^{-2}$ | 1.5 |

explained in Example 4.1. We found that vectorization of the subroutines in [2] did not change the execution times significantly, generally less than $20 \%$. In all computed examples the vectorized subroutines in [2] required at least twice as much execution time as the vectorized subroutines of Table I. For certain problems, our vectorized code executed up to 95 times faster than the vectorized code in [2]. For scalar code the differences in execution time often decreased with increasing matrix size. Tables V-VIII present some sample timings.

Tables V-VII present timings for vectorized code. The next table shows timings for scalar code for the same updatings as in Table V. Table VIII shows that, without vectorization, DELROW [2] requires $50 \%$ more CPU time then DDELR for moderately large problems.

## ACKNOWLEDGMENTS

One of the authors (L.R.) would like to thank Jack Dongarra and Pat Gaffney for valuable discussions, and Aladin Kamel for help in handling the computer.

## REFERENCES

1. Allen, D. M. Mean square error of prediction as a criterion for selecting variables. Technometrics 13 (1971), 469-475.
2. Buckley, A. Algorithm 580, QRUP: A set of FORTRAN routines for updating QR factorizations. ACM Trans. Math. Softw. 7 (1981), 548-549 and 8 (1982), 405.
3. Chan, T. F. Rank revealing QR factorizations. Lin. Alg. Appl. 88/89 (1987), 67-82.
4. Cline, A. K., Moler, C. B., Stewart, G. W., and Wilkinson, J. H. An estimate for the condition number of a matrix. SIAM J. Numer. Anal. I6 (1979), 368-375.
5. Daniel, J. W., Gragg, W. B., Kaufman, L., and Stewart, G. W. Reorthogonalization and stable algorithms for updating the Gram-Schmidt QR factorization. Math. Comput. 30 (1976), 772-795.
6. Dongarra, J. J., Bunch, J. R., Moler, C. B., and Stewart, G. W. Linpack Users' Guide. SIAM, Philadelphia, 1979.
7. Dongarra, J. J., and Sorensen, D. C. Linear algebra on high performance computers. In Applications of Supercomputers, D. F. Lockhart, and D. L. Hicks, Eds., Elsevier, Amsterdam, 1986, 5788.
8. Gill, P. E., Golub, G. H., Murray, W., and Saunders, M. A. Methods for modifying martrix factorizations. Math. Comput. 28 (1974), 505-535.
9. Goldfarb, D. Factorized variable metric methods for unconstrained optimization. Math. Comput. 30 (1976), 796-811.
10. Golub, G. H., and Van Loan, C. F. Matrix Computations. Johns Hopkins University Press, Baltimore, Md., 1983.
11. Gragg, W. B., and Stewart, G. W. A stable variant of the secant method for solving nonlinear equations. SIAM J. Numer. Anal. 13 (1976), 889-903.
12. Lawson, C., Hanson, R., Kincaid, D., and Krogh, F. Basic linear algebra subprograms for FORTRAN usage. ACM Trans. Math. Softw. 5 (1979), 308-323.
13. Ling, F., Manolakis, D., and Proakis, J. G. A recursive modified Gram-Schmidt algorithm for least-squares estimation. IEEE Trans. Acoustics, Speech and Signal Process. ASSP-34 (1986), 829-835.
14. Moré, J. J., and Sorensen, D. C. Newton methods. In Studies in Numerical Analysis, G. H. Golub, Ed., The Mathematical Association of America, Washington, D.C., 1984, 29-82.
15. NAG FORTRAN Library Manual, Mark 12. Numerical Algorithms Group, 1987.
16. Parlett, B. N. The Symmetric Eigenvalue Problem. Prentice-Hall, Englewood Cliffs, N.J., 1980.
17. Robert, Y., and Sguazzero, P. The LU decomposition algorithm and its efficient FORTRAN implementation on the IBM 3090 vector multiprocessor. Tech. Rep. ICE-0006, IBM European Center for Scientific and Engineering Computing, Rome, 1987.
18. Stewart, G. W. The effect of rounding error on an algorithm for downdating a Cholesky factorization. J. Inst. Math. Applic. 23 (1979), 203-213.

[^0]:    This research was supported by NSF grant DMS-8704196 and by the Foundation Research Program of the Naval Postgraduate School.
    Authors' addresses: L. Reichel, Bergen Scientific Center, Thormøhlensgaten 55, 36, N-5008, Bergen, Norway and Department of Mathematics, University of Kentucky, Lexington, KY 40506; W. B. Gragg, Department of Mathematics, Naval Postgraduate School, Monterey, CA 93943.

    Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

