# ARGONNE NATIONAL LABORATORY 

9700 South Cass Avenue
Argonne, IL 60439
$\qquad$

```
ANL/MCS-TM-175
```


# A Summary of Block Schemes for Reducing a General Matrix to Hessenberg Form 

by<br>Christian Bischof

Mathematics and Computer Science Division
Technical Memorandum No. 175

February 1993

This work was supported by the Office of Scientific Computing, U.S. Department of Energy, under Contract W-31-109-Eng-38.

## Contents

Abstract ..... 1
1 Introduction ..... 1
2 Block Representations for Representing Products of Householder Matrices ..... 1
3 Block Schemes for Reducing a General Matrix to Hessenberg Form ..... 2
4 Discussion of the Various Update Schemes ..... 4
References ..... 6

# A Summary of Block Schemes for Reducing a General Matrix to Hessenberg Form 

by

Christian H. Bischof


#### Abstract

Various strategies have been proposed for arriving at block algorithms for reducing a general matrix to Hessenberg form by means of orthogonal similarity transformations. This paper reviews and systematically categorizes the various strategies and discusses their computational characteristics.


## 1 Introduction

Let $A$ be an $n \times n$ symmetric matrix. Our goal is to compute an orthogonal matrix $Q, Q^{T} Q=I$ such that $Q^{T} A Q=H$ where $H$ is of upper Hessenberg form. The standard algorithm [10] reduces $A$ one column at a time through Householder transformation at a cost of $O\left(4 / 3 n^{3}\right)$ flops. It mainly employs matrix-vector multiplications and symmetric rank-one updates, which require more memory references than matrix-matrix operations [5,4,7].

On the other hand, so-called block reduction methods allow the formulation of most of the computation in terms of matrix-matrix operations [8], which, at a small increase in the number of floating-point operations, dramatically reduce the number of memory accesses. As a result, the block approach is preferable for machines employing a memory hierarchy, as do (in one form or another) most high-performance computers, from workstations to massively parallel machines. Thus, block algorithms play a prominent role in the LAPACK library of portable linear algebra codes for highperformance architectures [1].

The orthogonal reduction of a matrix to Hessenberg form involves a sequence of $m-2$ Householder reductions

$$
P_{i}=I-2 u_{i} u_{i}^{T}, u_{i}^{T} u_{i}=2 .
$$

The key to finding a block formulation is the expression of

$$
B=P_{p} \cdots P_{1} A P_{1} \cdots P_{p}
$$

in a form that is easy to compute and requires as little access to $A$ as possible. Dongarra, Hammarling, and Sorensen [8] suggested $B=A-U V^{T}-W U^{T}$. Dubrulle [9], on the other hand, uses a block formulation of $Q_{p}=P_{1} \cdots P_{p}$ based on the WY representation [6] to express $B=\left(I-U Y^{T}\right)^{T}\left(A-Z Y^{T}\right) . U, V, W, Y$, and $Z$ are all $n \times p$ matrices.

The paper is structured as follows: In $\S 2$ we review the possibilities for expressing a series $P_{1} \cdots P_{p}$ of Householder transformations in block form. In $\S 3$ we then discuss the various schemes that has been proposed and suggest some new ones. The relative merits of the various approaches are discussed in $\S 4$.

## 2 Block Representations for Representing Products of Householder Matrices

There are several ways for representing a product

$$
Q_{i}=P_{1} \cdots P_{i}
$$

of Householder transformations.

WY1: [6]

$$
\begin{equation*}
Q_{i}=I-U_{i} Y_{i}^{T} \tag{1}
\end{equation*}
$$

WY2: [6]

$$
\begin{equation*}
Q_{i}=I-W_{i} U_{i}^{T} \tag{2}
\end{equation*}
$$

compactWY: [12]

$$
\begin{equation*}
Q_{i}=I-U_{i} S_{i} U_{i}^{T} \tag{3}
\end{equation*}
$$

In all those schemes, $U_{i}$ is the matrix of Householder vectors, and extra storage is required for the other matrices. In the above formulas, $U_{i}, Y_{i}$, and $W_{i}$ are $m \times i$ matrices. $U_{i}$ and $Y_{i}$ are lower trapezoidal and $S_{i}$ is an upper triangular $i \times i$ matrix. Since typically $n \gg i$, the "compact" scheme requires less work space, at the expense of slightly more work in applying $Q_{i}$. These matrices are accumulated as follows $(k=1, \ldots, i-1)$. For all schemes, we have

$$
\begin{equation*}
U_{k+1}=\left[U_{k}, u_{k+1}\right] \tag{4}
\end{equation*}
$$

that is, $U$ just collects the Householder vectors used in the transformation. The other matrices get updated as follows:

$$
\begin{align*}
Y_{k+1} & =\left[P_{k} Y_{k}, u_{k+1}\right]  \tag{5}\\
W_{k+1} & =\left[W_{k}, Q_{k} u_{k+1}\right]  \tag{6}\\
S_{k+1} & =\left(\begin{array}{cc}
S_{k} & -S_{k} U_{k}^{T} u_{k+1} \\
0 & 1
\end{array}\right) \tag{7}
\end{align*}
$$

With these block formulations, the one-sided application of $Q_{i}$ to a matrix now involves two or three matrix-matrix multiplications instead of a series of $i$ rank-one updates and matrix-vector multiplications. The resulting decrease in memory traffic usually more than makes up for the additional $O\left(i^{2} n\right)$ floating-point operations. Examples of the use of block orthogonal transformations in the computation of the QR factorization $A=Q R$, where $Q$ is orthogonal and $R$ upper triangular, can be found in $[6,5,3]$.

We also mention that Walker [13] and Puglisi [11] suggested a variant of the compact WY representation which represents

$$
\begin{equation*}
Q_{i}=I-U_{i}\left(T_{i}\right)^{-1} U_{i}^{T} \tag{8}
\end{equation*}
$$

Computationally, the formulation (8) behaves very much like (3) and can be easily substituted in formulae based on (3). Hence, for the sake of brevity, we will not consider it further.

## 3 Block Schemes for Reducing a General Matrix to Hessenberg Form

The two main steps in the block algorithm are the generation of a block transformation and its application to the remaining part of the matrix. As a shorthand, let $\mathrm{HH}(x)$ be the Householder vector that reduces $x$ to a multiple of $e_{1}$. Let

$$
A_{i}=P_{i} \cdots P_{1} A P_{1} \cdots P_{i}
$$

be the matrix obtained after the $i$ th symmetric Householder update. Further, let $A^{(i: n, j: n)}$ be the $(n-i+1) \times(n-j+1)$ submatrix beginning at location $(i, j)$ of $A$ with analogous notation for vectors.

Having partitioned $A$ into $M$ block columns $\tilde{A}_{j}$ of width $p, m=M p, j=1, \ldots, M$, the "blueprint" for a block algorithm is shown in Figure 1. In the "rightlooking" variant we immediately apply $Q_{j}$ to all remaining block columns, whereas in the "leftlooking" algorithm we apply

```
\(\tilde{Q} \leftarrow I ;\)
for \(j=1\) to \(M\) do
    Find an orthogonal matrix \(\tilde{Q}_{p}\) such that
        \(\left(\tilde{Q}_{p}^{T} A(p(j-1)+1: p M, p(j-1)+1: p M) \tilde{Q}_{p}\right)(p(j-1)+1: p M, p(j-1)+1: p j)\)
    is upper Hessenberg.
    Either rightlooking algorithm:
            Update block columns \(j+1, \ldots, M\) with \(\tilde{Q}_{p}\).
    or leftlooking algorithm:
    Update block column \(j+1\) with \(\tilde{Q}_{1}, \ldots, \tilde{Q}_{j}\).
end for
```

Figure 1:
all previous block orthogonal transformations to the next block column only. In particular, the leftlooking algorithm requires storage of the block transformations $Q_{1}, \ldots, Q_{j}$. Hence the rightlooking algorithm is usually preferred (see [2] for a thorough discussion of the leftlooking vs. rightlooking issue). In order to develop an efficient block algorithm, we must be to find $\tilde{Q}_{j}$ and update $\tilde{A}_{j}$ in a fashion that requires as little as possible access to block columns $j+1, \ldots, M$. We consider the following alternatives:
implicit_Q1: ([8])

$$
\begin{equation*}
A_{p}=A-U_{p} V_{p}^{T}-W_{p} U_{p}^{T} \tag{9}
\end{equation*}
$$

implicit_Q2: ([9])

$$
\begin{equation*}
A_{p}=\left(I-U_{p} Y_{p}^{T}\right)^{T} A\left(I-U_{p} Y_{p}^{T}\right)=A-U_{p} \tilde{V}_{p}^{T}-\tilde{W}_{p} U_{p}^{T} \tag{10}
\end{equation*}
$$

Q_WY: (see [9])

$$
\begin{equation*}
A_{p}=\left(I-U_{p} S_{p} U_{p}^{T}\right)^{T} A\left(I-U_{p} S_{p} U_{p}^{T}\right)=\left(I-U_{p} Y_{p}^{T}\right)^{T}\left(A-Z_{p} Y_{p}^{T}\right) \tag{11}
\end{equation*}
$$

## Q_compactWY:

$$
\begin{equation*}
A_{p}=\left(I-U_{p} S_{p} U_{p}^{T}\right)^{T}\left(A-Z_{p} S_{p} U_{p}^{T}\right) \tag{12}
\end{equation*}
$$

These matrices are accumulated as shown below for $i=1, \ldots, p-1$. For all cases we have

$$
U_{i+1}=\left[U_{i}, u_{i+1}\right]
$$

## implicit_Q1:

$$
\begin{gather*}
u_{i+1}=\operatorname{HH}\left(a_{i}^{(i: m)}-\left(U_{i} V_{i}^{T}+W_{i} U_{i}^{T}\right) e_{i}^{(i: m)}\right) \\
v=\left(A^{T}-V_{i} U_{i}^{T}-U_{i} W_{i}^{T}\right) u_{i+1}, x=\left(A-U_{i} V_{i}^{T}-W_{i} U_{i}^{T}\right) u_{i+1}  \tag{13}\\
V_{i+1}=\left[V_{i}, v-\frac{x^{T} u_{i+1}}{2} u_{i+1}\right], W_{i+1}=\left[W_{i}, x-\frac{v^{T} u_{i+1}}{2} u_{i+1}\right]
\end{gather*}
$$

implicit_Q2:

$$
\begin{gather*}
\left.u_{i+1}=\operatorname{HH}\left(a_{i}^{(i: m)}-U_{i} \tilde{V}_{i}^{T}-\tilde{W}_{i} U_{i}^{T}\right) e_{i}^{(i: m)}\right) \\
x=\left(A^{T}-U_{i} \tilde{V}_{i}^{T}-\tilde{W}_{i} U_{i}^{T}-u_{i+1} u_{i+1}^{T}\left(A^{T}-U_{i} \tilde{V}_{i}^{T}-\tilde{W}_{i} U_{i}^{T}\right)\right) u_{i+1}  \tag{14}\\
\tilde{V}_{i+1}=\left[\tilde{V}_{i}, x\right], \tilde{W}_{i+1}=\left[\tilde{W}_{i}, A u_{i+1}\right]
\end{gather*}
$$

Q_WY:

$$
\begin{gather*}
u_{i+1}=\mathrm{HH}\left(\left(I-U_{p} Y_{p}^{T}\right)^{T}\left(a_{i}^{(i: m)}-Z_{i} Y_{i}^{T} e_{i}^{(i: m)}\right)\right) \\
Z_{i+1}=\left[Z_{i}, A u_{i+1}\right]  \tag{15}\\
Y_{i+1} \text { is updated as in (5). }
\end{gather*}
$$

Q_compactWY;

$$
\begin{gather*}
u_{i+1}=\operatorname{HH}\left(\left(I-U_{i} S_{i} U_{i}^{T}\right)^{T}\left(a_{i}^{(i: m)}-Z_{i} S_{i} U_{i}^{T} e_{i}^{(i: m)}\right)\right) \\
Z_{i+1}=\left[Z_{i}, A u_{i+1}\right]  \tag{16}\\
S_{i+1} \text { is updated as in (7). }
\end{gather*}
$$

While the latter two schemes are reformulations of similarity transformations of $Q$ in block form, the first two schemes do not explicitly involve a block orthogonal transformation. Thus, for the 'Q_WY' and 'Q_compactWY' schemes, the orthogonal transformations $\tilde{Q}_{j}$ are readily available in a convenient block form, whereas there is no block form of $\tilde{Q}_{j}$ readily accessible in the "implicit" schemes. Hence, the "implicit" schemes allow only for a rightlooking algorithm, whereas the other schemes allow for both a leftlooking and rightlooking formulation.

## 4 Discussion of the Various Update Schemes

To assess the work required for a block update, let us assume that $A$ is $m \times m$ and that we choose a block size of $p$. Then at step $i=1, \ldots, p$ the length of $u_{i}$ is $m-i+1$ and the work required to perform an update step can be expressed in the following units:
$T_{m v}$ : Work required to multiply an $m \times(i-1)$ matrix with an $(i-1)$ vector or an $(i-1) \times m$ matrix by an $m$-vector. This requires $2 m(i-1)$ flops.
$T_{s v}$ : Work required to multiply an $(i-1) \times(i-1)$ upper triangular matrix by an $(i-1)$-vector. This requires $(i-1)^{2}$ flops.
$T_{A u}:$ Work required to multiply an $(m-i+1) \times(m-i+1)$ matrix by an $(m-i+1)$-vector. This requires $2(m-i)^{2}$ flops.

After a block update of width $p$ is computed, this update will have to be applied to an $m \times k$ submatrix of $A$, where $k$ is determined primarily by our choice of a left- or rightlooking algorithm. If we choose a leftlooking algorithm, then $k=p$ (assuming fixed block sizes), whereas for the rightlooking algorithm $k=m-p$. This work can be expressed in terms of
$T_{r k_{-} p}:$ Work required to apply a rank-p update to an $m \times k$ matrix. This requires $2 m k p$ flops.
$T_{m m 1}$ : Work required to multiply a $p \times m$ by an $m \times k$ matrix. This requires $2 m k p$ flops.
$T_{m m 2}$ : Work required to multiply a $p \times p$ upper triangular matrix by a $p \times k$ matrix. This requires $k p^{2}$ flops.

Another point to consider is whether the block formulation of $A$ also produces a block formulation for $Q$. The latter may be important if $Q$ has to be applied to the eigenvectors of the Hessenberg matrix in a back-transform step. Using the above criteria, we summarize the differences between the different blocking schemes in Table 1. The first row (labeled "Gen. Step $i$ ") shows the work to be performed at the $i$ th step of generating the block transformation, the second row (labeled "Application") shows the work required in applying the block update, and the last row (labeled "Block Q ?") indicates whether this formulation also results in a block formulation for $Q$.

Table 1: Work required at Step $i$

|  | implicit_Q1 | implicit_Q2 | Q_WY | Q_compactWY |
| :---: | :---: | :---: | :---: | :---: |
| Gen. Step $i$ | $10 T_{m v}+2 T_{A u}$ | $6 T_{m v}+2 T_{A u}$ | $4 T_{m v}+T_{A u}$ | $4 T_{m v}+3 T_{s v}+T_{A u}$ |
| Application | $2 T_{r k_{-} p}$ | $2 T_{r k_{-} p}$ | $2 T_{r k_{-} p}+T_{m m 1}$ | $2 T_{r k_{-p}}+T_{m m 1}+2 T_{m m 2}$ |
| Block Q ? | no | no | yes | yes |

Table 2: Storage Requirements

|  | implicit_Q1 | implicit_Q2 | Q_WY | Q_compactWY |
| :--- | :---: | :---: | :---: | :---: |
| Representing blocks | $2 m p$ | $2 m p$ | $2 m p$ | $m p+\frac{p^{2}}{2}$ |
| Extra storage discarding blocks | none | none | none | $k p$ |
| Extra storage preserving blocks | none | none | $k p$ | $k p$ |

We compare in Table 2 the amount of storage required for the different schemes in a rightlooking algorithm. Assuming that $U$ is stored in the lower triangular part of $A$, we consider the amount of storage required to represent the blocking factors, and the extra amount of storage required to apply an $m \times p$ block update to an $m \times k$ matrix with and without discarding of block orthogonal transformations.

Dubrulle [9] compared a rightlooking algorithm using formulations "implicit_Q2" and "Q_WY" on the IBM 3090 /VF using straight Fortran. In these experiments he found the "Q_WY" formulation to be superior to the "implicit_Q2" formulation.

These numbers suggest that the "Q_WY" scheme is superior to the "implicit_Q2" scheme even if we do not need the block factors defining $Q$ later on. The reason seems to be that in the "implicit_Q2" scheme we have to compute $A u$ and $A^{T} u$ at every step of accumulating a block transform, whereas "Q_WY" gets by with only computing $A u$ and with two less matrix-vector multiplications. On the other hand "Q_WY" requires an extra matrix-matrix product (the work needed is denoted by $T_{m m 1}$ in Table 1) when the block transformation is applied. Thus it seems as if "Q_WY" profits from the usual BLAS 3 to BLAS 2 tradeoff in that matrix-matrix operations execute more efficiently than a series of matrix-vector operations. Furthermore the "Q_WY" scheme requires less data traffic since we touch $A$ only once at every update step, whereas "implicit_Q2" touches $A$ twice.*

The scheme "implicit_Q2" scheme is similar to "implicit_Q1" which was implemented in the second test release of LAPACK, and it is reasonable to assume that it would perform similarly. The "Q_compactWY" scheme is similar to scheme "Q_WY" but requires less storage, in particular if a left-looking scheme is used (i.e., $k \ll m$ in Table 2) or when block transformations are stored for successive back transformations. Since the "implicit" schemes do not require less storage, do not give us any block formulation for $Q$, and do not seem to perform any better, it seems advantageous to use one of the WY-based schemes. In particular, the one based on the compact WY representation seems to be appropriate, since it allows economical storage of the block transforms for subsequent back transformations of the eigenvectors of the Hessenberg matrix. Hence, the "Q_compactWY" scheme was chosen for the final release of LAPACK [1].
${ }^{*}$ We assume here that $A u$ and $A^{T} u$ are computed separately using BLAS 2 calls. We can get by with touching $A$ only once, but then we cannot use a BLAS 2 call.

## Acknowledgments

I thank Jeremy Du Croz and Sven Hammarling for the many stimulating discussions that eventually prompted me to methodologically evaluate the various possibilities.

## References

[1] E. Anderson, Z. Bai, C. Bischof, J. Demmel, J. Dongarra, J. DuCroz, A. Greenbaum, S. Hammarling, A. McKenney, and D. Sorensen. LAPACK User's Guide. SIAM, Philadelphia, Penna., 1992.
[2] E. Anderson and J. Dongarra. Evaluating block algorithm variants in LAPACK. Technical Report CS-90-103, Computer Science Department, The University of Tennessee, April 1990.
[3] C. H. Bischof. A block QR factorization algorithm using restricted pivoting. In Proceedings SUPERCOMPUTING '89, pages 248-256, Baltimore, Md., 1989. ACM Press.
[4] C. H. Bischof. Fundamental linear algebra computations on high-performance computers, in volume 250 of Informatik Fachberichte, pages 167-182. Springer-Verlag, Berlin, 1990.
[5] C. H. Bischof and J. J. Dongarra. A project for developing a linear algebra library for highperformance computers. In Graham Carey, editor, Parallel and Vector Supercomputing: Methods and Algorithms, pages 45-56. John Wiley \& Sons, Somerset, N.J., 1989.
[6] C. H. Bischof and C. F. Van Loan. The WY representation for products of Householder matrices. SIAM Journal on Scientific and Statistical Computing, 8:s2-s13, 1987.
[7] J. Dongarra and S. Hammarling. Evolution of Numerical Software for Dense Linear Algebra, pages 297-327. Oxford University Press, Oxford, U.K., 1989.
[8] J. J. Dongarra, S. J. Hammarling, and D. C. Sorensen. Block reduction of matrices to condensed form for eigenvalue computations. Technical Report MCS-TM-99, Mathematics and Computer Science Division, Argonne National Laboratory, September 1987.
[9] A. A. Dubrulle. On block Householder algorithms for the reduction of a marix to Hessenberg form. In Joanne L. Martin and Stephen F. Lundstrom, editors, Supercomputing '88: Volume II, Science and Applications, Washington, DC, 1989. IEEE Computer Society Press.
[10] G. H. Golub and C. F. Van Loan. Matrix Computations. The Johns Hopkins University Press, Baltimore, 1983.
[11] C. Puglisi. Modification of the Householder method based on the compact WY representation. CERFACS Report TR/PA/90/29, 1990.
[12] R. Schreiber and C. Van Loan. A storage efficient WY representation for products of Householder transformations. SIAM Journal on Scientific and Statistical Computing, 10(1):53-57, 1989.
[13] H. F. Walker. Implementation of the GMRES method using Householder transformations. SIAM Journal on Scientific and Statistical Computing, 9(1):152-163, 1988.

