

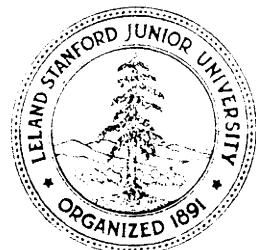
ON COMPUTING THE SINGULAR VALUE DECOMPOSITION

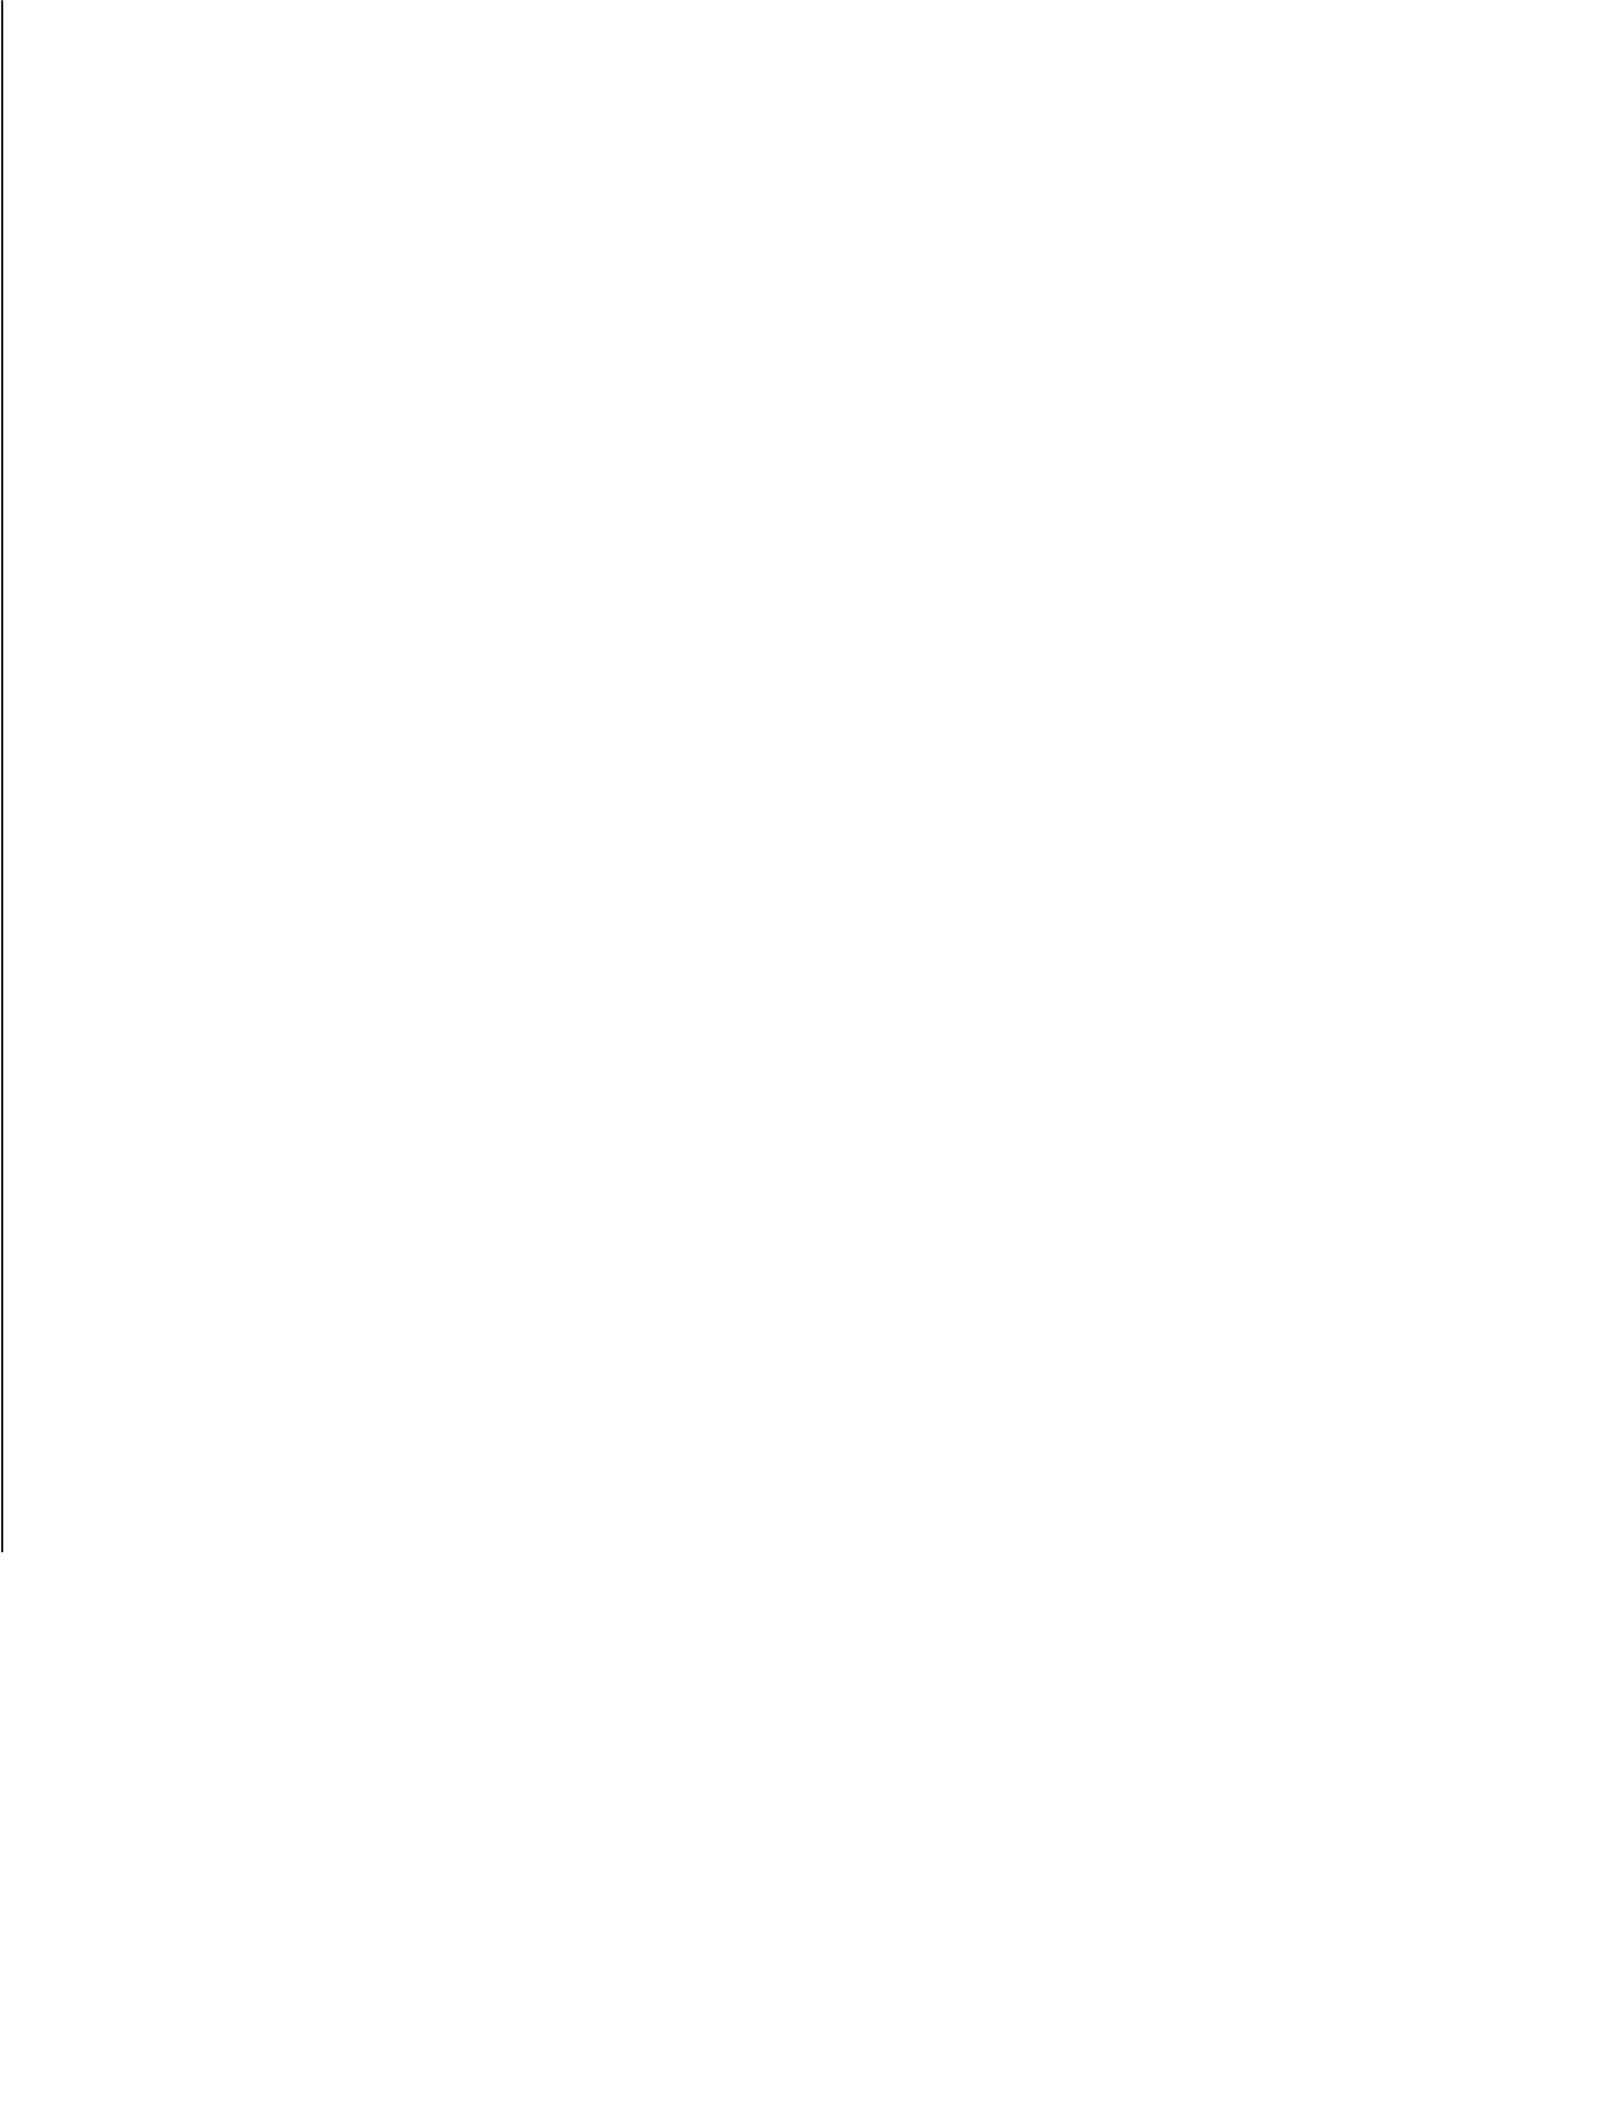
by

Tony Fan C. Chan

STAN-CS-77-588  
FEBRUARY 1977

COMPUTER SCIENCE DEPARTMENT  
School of Humanities and Sciences  
STANFORD UNIVERSITY





ON COMPUTING THE SINGULAR VALUE DECOMPOSITION

by

Tony Fan C. **CHAN**<sup>\*</sup>

**August**, 1976.

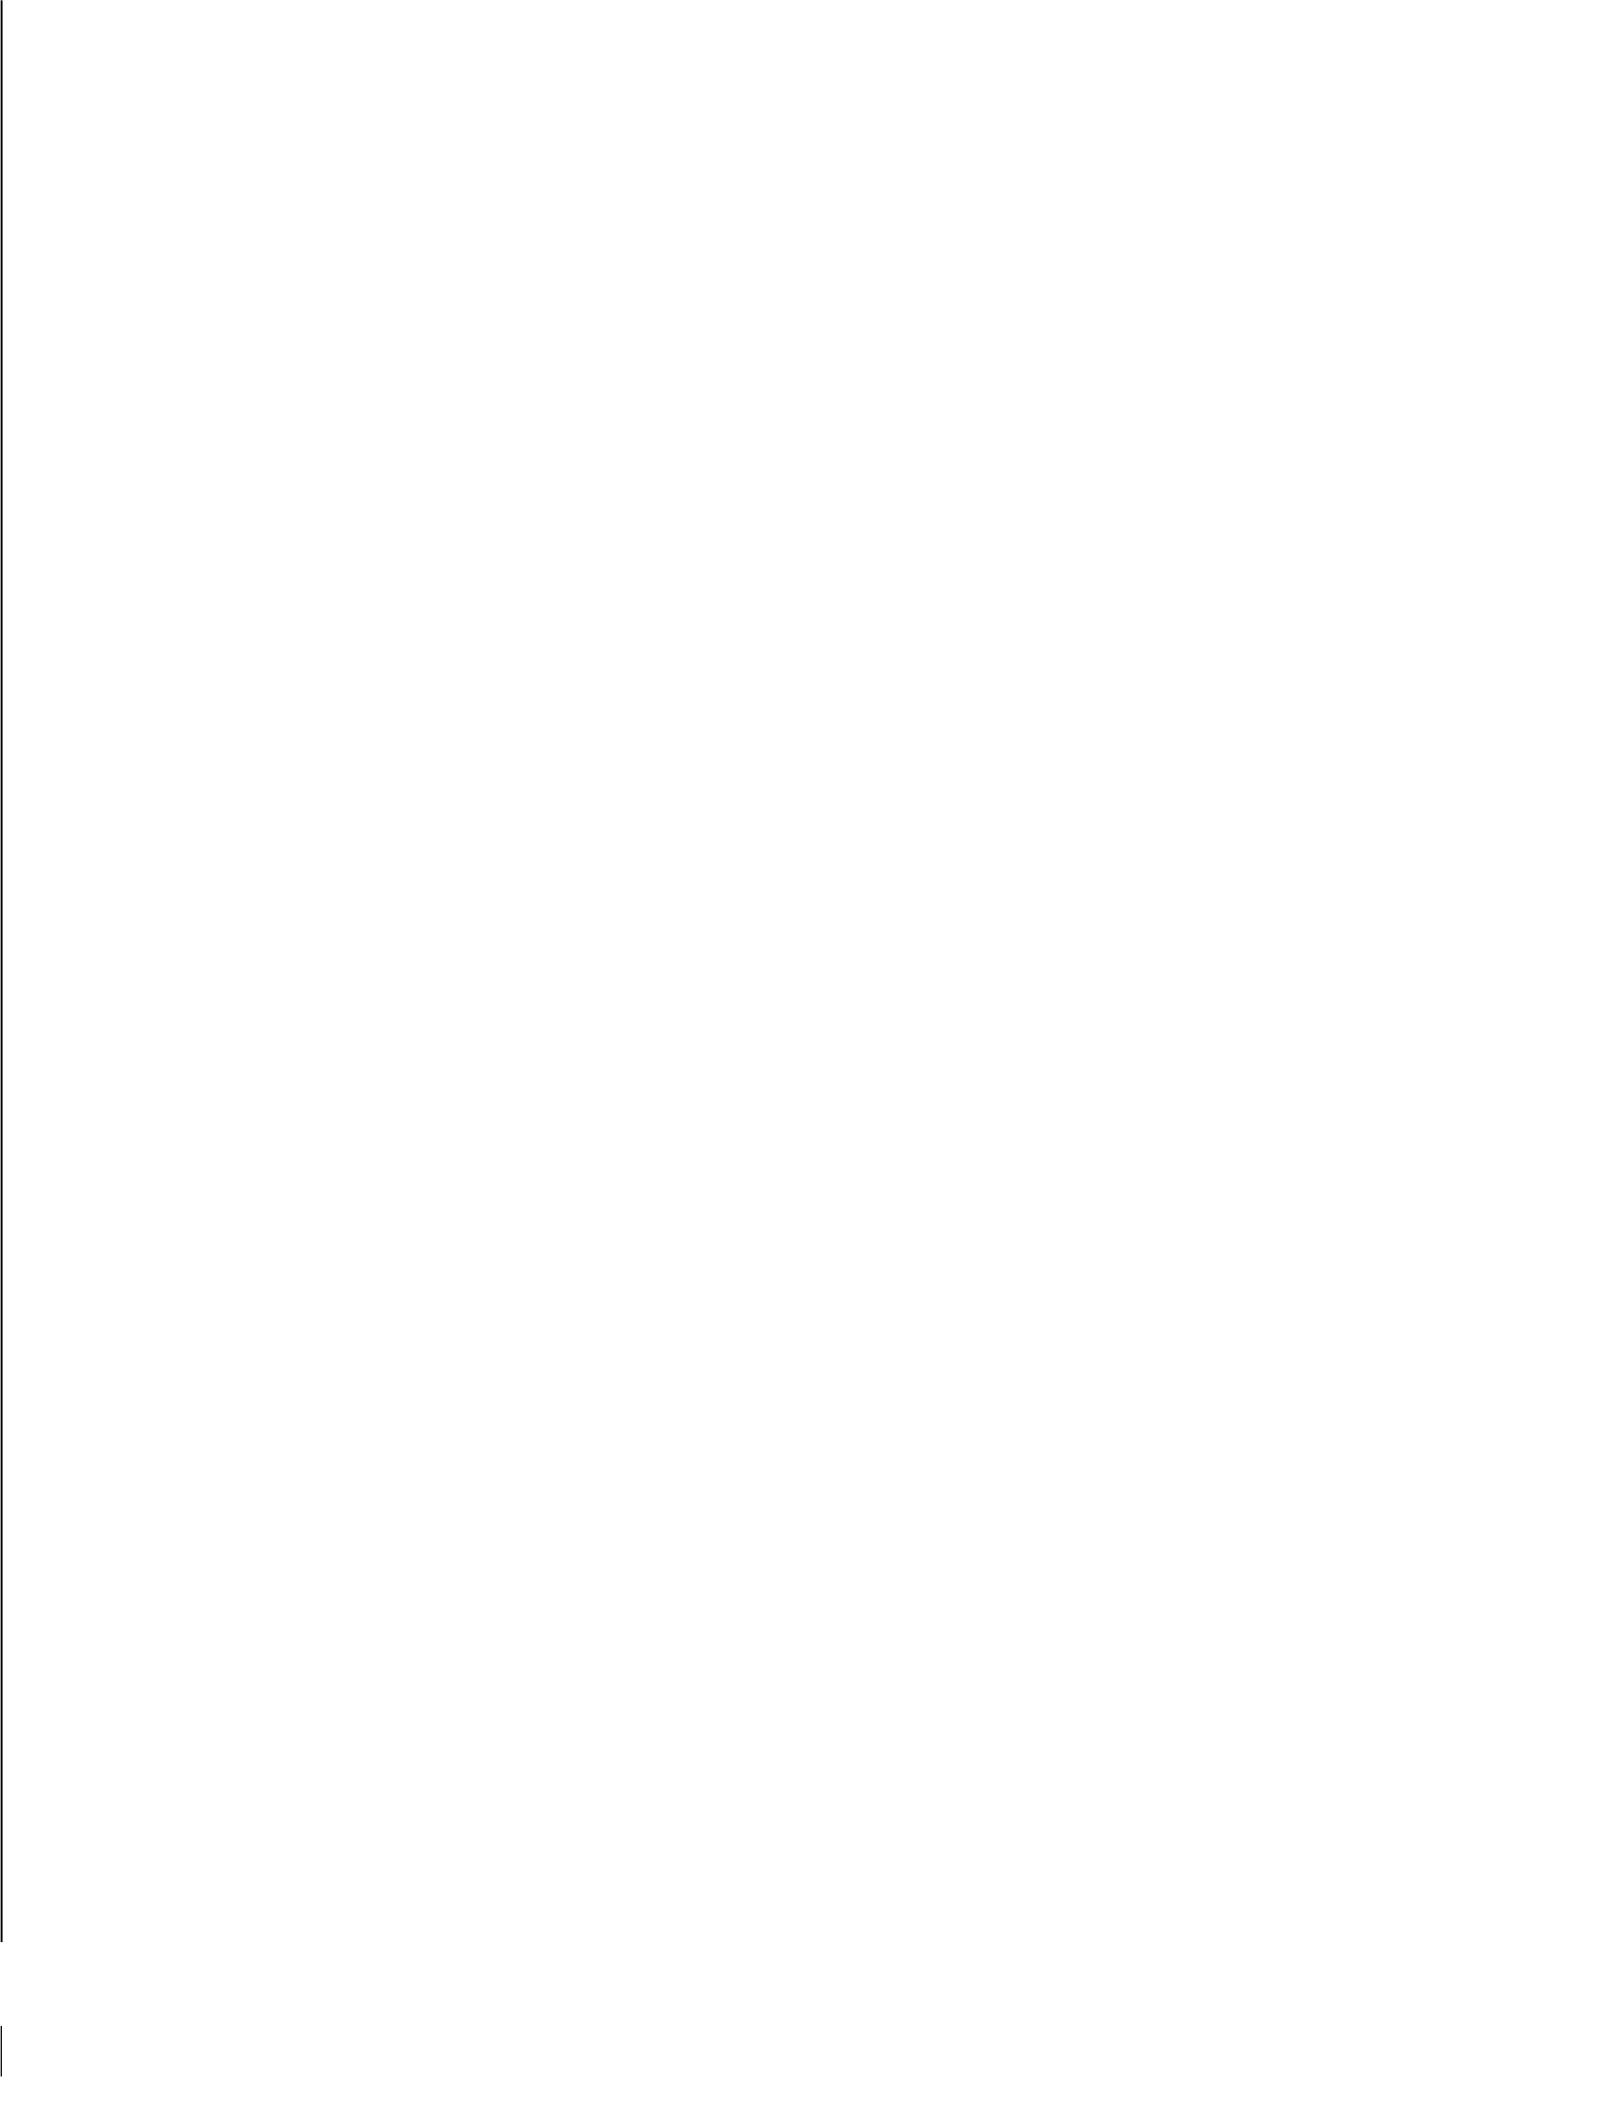
\*Computer Science Dept., Stanford Univ., Ca94305.

This work was supported by NSF Grant **DCR75-13497** and NASA Ames Contract **NCA2-0R745-520**. The computing time was provided by the Stanford Linear Accelerator Center (**SLAC**).



## ABSTRACT

The most well-known and widely-used algorithm for computing the Singular Value Decomposition (SVD) of an  $m \times n$  rectangular matrix  $A$  nowadays is the Golub-Reinsch algorithm [1]. In this **paper**, it is shown that by (1) first triangularizing the matrix  $A$  by Householder transformations before bidiagonalizing it, and (2) accumulating some left transformations on an  $n \times n$  array instead of on an  $m \times n$  array, the resulting algorithm is often more efficient than the Golub-Reinsch algorithm, especially for matrices with considerably more rows than columns ( $m \gg n$ ), such as in least squares applications. The two algorithms are compared in terms of operation counts, and computational experiments that have been carried out verify the theoretical comparisons. The modified algorithm is more efficient even when  $m$  is only slightly greater than  $n$ , and in some cases can achieve as much as 50% savings when  $m \gg n$ . If accumulation of left transformations is desired, then  $n^2$  extra storage locations are required (relatively small if  $m \gg n$ ), but otherwise no extra storage is required. The modified algorithm uses only orthogonal transformations and is therefore numerically stable. In the Appendix, we give the FORTRAN code of a hybrid method which automatically selects the more **efficient** of the two algorithms to use depending upon the input values for  $m$  and  $n$ .



(0) INTRODUCTION

Let  $A$  be a real  $m \times n$  matrix, with  $m \gg n$ . It is well-known [1,2] that the following decomposition of  $A$  always exists :

$$A = U \sum v^T \quad (0.1)$$

where  $U$  is a  $m \times n$  matrix and consists of  $n$  orthonormalized eigenvectors associated with the  $n$  largest eigenvalues of  $AA^T$ ,  $V$  is a  $n \times n$  matrix and consists of the orthonormalized eigenvectors of  $A^T A$ , and  $\Sigma$  is a diagonal matrix consisting of the "singular values" of  $A$ , which are the non-negative square roots of the eigenvalues of  $A^T A$ .

Thus,

$$U^T U = V^T V = V V^T = I_n$$

and  $\Sigma = \text{diag}(\sigma_1, \dots, \sigma_n)$ . (0.2)

It is usually assumed for convenience that

$$\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_n > 0.$$

The decomposition (0.1) is called the Singular Value Decomposition (SVD) of  $A$ .

Remarks:

(1) If  $\text{rank}(A) = r$ , then  $\sigma_{r+1} = \sigma_{r+2} = \dots = \sigma_n = 0$ .

(2) There is no loss of generality in assuming that  $m \geq n$ , for if  $m < n$ , then we can instead compute the SVD of  $A^T$ . If the SVD of  $A^T$  is equal to  $U \sum v^T$ , then the SVD of  $A$  is equal to  $V \sum U^T$ .

The SVD plays a very important role in linear algebra. It has applications in such areas as least squares problems [1,2,3], in computing the pseudo-inverse [2], in computing the Jordan Canonical form [4], in solving integral equations [5], in digital image processing [6], and in optimization [7]. Many of the applications often involve large matrices. It is therefore important that the computational procedures for obtaining the SVD be as efficient as possible.

It is perhaps difficult to find an algorithm that has optimal efficiency for all matrices, so it would be desirable to know for what kind of matrices a given algorithm is best suited. It is in this spirit that we were first motivated to look for improvements of the Golub-Reinsch algorithm when the matrix  $A$  has considerably more rows than columns, i.e.  $m \gg n$ . It turns out that such an improvement is indeed possible, with only slight modifications to the Golub-Reinsch algorithm, even when  $m$  is only slightly greater than  $n$ , and can sometimes achieve as much as 50% savings in execution time when  $m \gg n$ .

In section (1) we will briefly describe the Golub-Reinsch algorithm. We will then present the modified algorithm in section (2), with some computational details deferred to section (3). Operation counts for the two algorithms will be given in section (4) and some computational results in section (5). We will make some conclusions in section (6). In the Appendix,

we will give the FORTRAN implementation of a hybrid method which automatically selects the more efficient of the two algorithms to use depending upon the input values for m and n.

(1) THE GOLUB-REINSCH ALGORITHM (GR-SVD)

We will use the same notation as in [1].

This algorithm consists of two phases. In the first phase one constructs two finite sequences of Householder transformations

$$p^{(k)} \quad (k=1, 2, \dots, n)$$

and  $q^{(k)} \quad (k=1, 2, \dots, n-2)$

such that

$$p^{(n)} \dots p^{(1)} A q^{(1)} \dots q^{(n-2)} = \begin{bmatrix} * & * & \dots & * & 0 \\ 0 & * & \dots & * & * \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & 0 & * \\ 0 & 0 & \dots & 0 & 0 \end{bmatrix} = J^{(0)},$$

$(m-n) \times n$

an upper bidiagonal matrix. Specifically,  $p^{(i)}$  zeros out the subdiagonal elements in column  $i$  and  $q^{(j)}$  zeros out the appropriate elements in row  $j$ .

The singular values of  $J^{(0)}$  are the same as those of  $A$ .

Thus,

$$\begin{aligned} \text{if } J &= G \sum H^T \text{ is the SVD of } J, \\ \text{then } A &= P G \sum H^T Q^T \\ \text{so that } U &= P G, \quad V = Q H \quad (1.2) \\ \text{with } P &= p^{(1)} \dots p^{(n)}, \quad Q = q^{(1)} \dots q^{(n-2)}. \end{aligned}$$

The second phase is to iteratively diagonalize  $J^{(0)}$  by the QR method so that

$$J^{(0)} \rightarrow J^{(1)} \rightarrow \dots \rightarrow \Sigma$$

where  $J^{(i+1)} = S^{(i)}^T J^{(i)} T^{(i)}$ , (1.3)

where  $S^{(i)}$  and  $T^{(i)}$  are products of Givens transformations and are therefore orthogonal.

The matrices  $T^{(i)}$  are chosen so that the sequence  $M^{(i)} = J^{(i)}^T J^{(i)}$  converges to a diagonal matrix while the matrices  $S^{(i)}$  are chosen so that all  $J^{(i)}$  are of bidiagonal form.

The products of the  $T^{(i)}$ 's and the  $S^{(i)}$ 's are exactly the matrices  $H^T$  and  $G^T$  respectively in Eqn (1.2). For more details, see [1].

It has been reported in [1] that the average number of iterations on  $J^{(i)}$  in (1.3) is usually less than  $2n$ . In other words,  $J^{(2n)}$  in Eqn (1.3) is usually a good approximation to a diagonal matrix.

We will briefly describe how the computation is usually implemented. Assume for simplicity, that we can destroy  $A$  and return  $U$  in the storage for  $A$ . In the first phase, the  $P^{(i)}$  are stored in the lower part of  $A$ , and the  $Q^{(i)}$  are stored in the upper triangular part of  $A$ . After the bidiagonalization, the  $Q^{(i)}$  are accumulated in the storage provided for  $V$ , the two diagonals of  $J^{(0)}$  are copied to two other linear arrays, and the  $P^{(i)}$  are accumulated in  $A$ .

In the second phase, for each  $i$ ,

$s^{(i)}$  is applied to  $P$  from the right and

$T^{(i)T}$  is applied to  $Q^T$  from the left

in order to accumulate the transformations.

## (2) THE MODIFIED ALGORITHM (MOD-SVD)

Our original motivation for this algorithm is to find an improvement of GR-SVD when  $m \gg n$ . In that case, two improvements are possible:

(1) In Eqn (1.1), each of the transformations  $P^{(1)}$  and  $Q^{(1)}$  has to be applied to a submatrix of size  $(m-i+1) \times (n-i+1)$ .

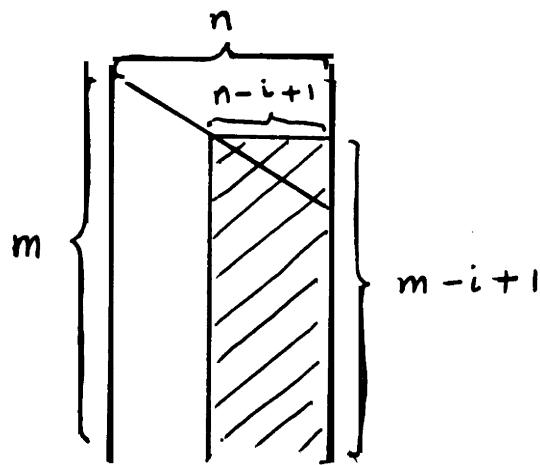


Fig. 2.1  $P^{(1)}$  and  $Q^{(1)}$  affects the shaded portion of the matrix

Now, since most entries of this submatrix are ultimately going to be zeros, it is intuitive that if it can somehow be arranged that the  $Q^{(1)}$  does not have to be applied to the subdiagonal part of this submatrix, then we will be saving a great amount of work when  $m \gg n$ .

This can indeed be done by first transforming  $A$  into upper triangular form by Householder transformations on the left.

$$L^T \begin{bmatrix} A \end{bmatrix} \rightarrow \begin{bmatrix} \text{triangular} \\ 0 \end{bmatrix} \equiv \begin{bmatrix} R \\ 0 \end{bmatrix}$$

where  $R$  is  $n \times n$  upper triangular and  $L$  is orthogonal, and then proceed to bidiagonalize  $R$ . The important difference is that this time we will be working with a much smaller matrix  $R$  than  $A$  (if  $n^2 \ll mn$ ), and so it is conceivable that the work required to bidiagonalize  $R$  is much smaller than that originally done by the right transformations when  $m \gg n$ .

The question still remains as to how to bidiagonalize  $R$ . An obvious way is to treat  $R$  as an input matrix to GR-SVD, using alternating left and right Householder transformations. In fact, it can be easily verified that if the SVD of  $R$  is equal to  $X \Sigma Y^T$ , then the SVD of  $A$  is given by

$$A = L \begin{bmatrix} X \\ 0 \end{bmatrix} \Sigma Y^T \quad (2.1)$$

We can identify  $U$  with  $L \begin{bmatrix} X \\ 0 \end{bmatrix}$  and  $V$  with  $Y$ . Notice that in order to obtain  $U$ , we have to form the extra product  $L \begin{bmatrix} X \\ 0 \end{bmatrix}$ . If  $U$  is not needed (e.g. in least squares), then we do not have to accumulate any left transformations and in that case, for  $m \gg n$ , it seems likely that we will make a substantial saving.

It is also possible to take advantage of the structure of  $R$  to bidiagonalize it. This will be discussed in section (3).

(ii) The second improvement over GR-SVD that can be made is the following. In GR-SVD, each of the  $s^{(i)}$  is applied to the  $m \times n$  matrix  $P$  from the right to accumulate  $U$ . If  $m \gg n$ , then this accumulation involves a large amount of work because a single Givens transformation affects two columns of  $P$  (of length  $m$ ) and each  $s^{(i)}$  is the product of on the average  $n/2$  Givens transformations. Therefore, in such cases, it would seem more efficient to first accumulate all  $s^{(i)}$  on a  $n \times n$  array  $Z$  and later form the matrix product  $PZ$  after  $J^{(i)}$  has converged to  $\sum$ .

In essence, improvement (i) works best when  $U$  is not needed, improvement (ii) works best when  $U$  is needed and both work best when  $m \gg n$ .

We now present the modified algorithm:

MOD-SVD:

(1)  $L^T \begin{bmatrix} A \end{bmatrix} \rightarrow \begin{bmatrix} R \\ 0 \end{bmatrix}$  where  $R$  is  $n \times n$  upper triangular,

(2) Find the SVD of  $R$  by GR-SVD,  $R = X \sum Y^T$ ,

(3) Form  $A = L \begin{bmatrix} X \\ 0 \end{bmatrix} \sum Y^T$ , the SVD of  $A$ .

The idea of transforming A to upper triangular form when  $m \gg n$  and then calculating the SVD of R is mentioned in Lawson & Hanson [3,pp.119,122] in the context of least squares problems where U is not explicitly required.

In the next section we will discuss some computational details of this modified algorithm, and in section (4) we will compare the operation counts of the two algorithms.

### (3) SOME COMPUTATIONAL DETAILS

(i) It should be obvious that when  $U$  is not needed then MOD-SVD does not require any extra storage. When  $U$  is needed, we can store  $L^T$  in the lower part of  $A$ , copy  $R$  into another  $n \times n$  array  $W$  and ask GR-SVD to return  $X$  in  $W$ . Therefore we need at most  $n^2$  extra storage locations which is relatively small when  $m \gg n$ .

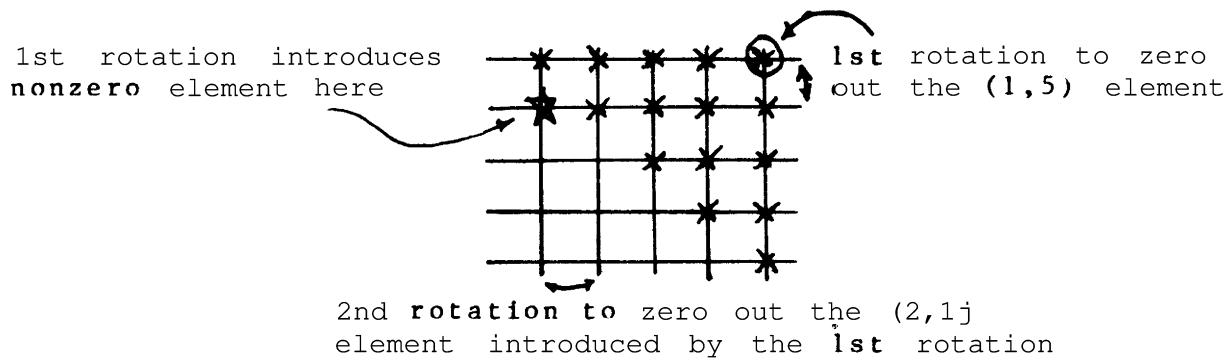
(ii) The next question is how to form  $L \begin{bmatrix} X \\ 0 \end{bmatrix}$  without using extra storage. This can be done by noting that

$$L \begin{bmatrix} X \\ 0 \end{bmatrix} = L \begin{bmatrix} I \\ 0 \end{bmatrix} X$$

so we can first accumulate  $L \begin{bmatrix} I \\ 0 \end{bmatrix}$  in the space provided for  $U$  and then do a matrix multiplication by  $X$ .

In the experiments that we have carried out, we actually accumulate the Householder transformations  $L$  on  $\begin{bmatrix} X \\ 0 \end{bmatrix}$ . We do not recommend doing this in practice because it requires  $mn$  instead of  $n^2$  extra storage locations. But one can show that both methods take about the same amount of work and so it will not affect the comparisons.

(iii) The question arises whether it is possible to bidiagonalize  $R$  in a way that takes advantage of the zeros that are already in  $R$ . One way is to use **Givens** transformations to zero out the elements at the upper right hand corner of  $R$ , one column or one row at a time. **Pictorially**, (for  $n=5$ ) to zero out the  $(1,5)$  element, we do two Givens transformations as follows:



It turns out however, by simple counting, that this method takes about the same operations ( $4n^3/3$  multiplications) as the previous method to bidiagonalize  $R$ , provided that we do not have to accumulate transformations. If we do need to accumulate either the left or the right transformations, then this method will require more work ( $4n^3$  versus  $4n^3/3$  mult.) mainly because it requires two rotations to zero out each element and these rotations have to be accumulated. So it seems that taking advantage of the zero structure of  $R$  in this fashion actually makes the method less efficient.

We have to note, however, that Givens transformations involve fewer additions and array accesses than Householder transformations per multiplication (see section 4.1). Therefore this method tends to be more competitive on modern computers where the time taken for floating point additions and multi-dimensional array indexings are not negligible compared to that for multiplications.

There may be other ways to bidiagonalize  $R$  using orthogonal transformations, but we shall not pursue this subject further.

(4) OPERATION COUNTS

In section (2), we indicated that MOD-SVD should be more efficient than CR-SVD when  $m \gg n$ . In this section, we study the relative efficiency between CR-SVD and MOD-SVD as a function of  $m$  and  $n$ . We do this by computing the asymptotic operation counts for each algorithm.

In the operation counts given below, we only keep the highest order terms in  $m$  and  $n$ , and so the results are correct for relatively large  $m$  and  $n$ .

CR-SVD:

(1) Bidiagonalization (using Householder transformations)

$$J = P(n) \dots P(1) A Q(1) \dots Q(n-2) \quad 2(mn^2 - n^3/3) \text{ mult.}$$

$$\text{accumulate } P = P(1) \dots P(n) \quad mn^2 - n^3/3 \text{ mult.}$$

$$\text{accumulate } Q = Q(1) \dots Q(n-2) \quad 2n^3 \text{ mult.}$$

(2) Diagonalization (using Givens transformations)

$$\text{accumulate } S^{(i)} \text{ on } P \quad Cmn^2 (C=4) \text{ mult.}$$

$$\text{accumulate } T^{(i)} \text{ on } Q \quad Cn^3 (C=4) \text{ mult.}$$

MOD-SVD:

(1) Triangularization (using Householder transformations)

$$L^T \begin{bmatrix} A \end{bmatrix} \rightarrow \begin{bmatrix} R \\ 0 \end{bmatrix} \quad mn^2 - n^3/3 \text{ mult.}$$

(2) CR-SVD of  $R$ ,  $R = X \sum Y^T$  depends on whether accumulations are needed.

(3) Form  $L \begin{bmatrix} X \\ 0 \end{bmatrix}$  (using Householder transf.)  $2mn^2 - n^3$  mult.

Some comments are in order:

(0) The entries  $C_{mn}^2$  and  $C_n^3$  with C-4 in the diagonalization phase of CR-SVD are obtained by assuming that the iterative phase of the SVD takes on the average two complete QR iterations per singular value [1], [3, p122]. We have checked this experimentally and found it to be quite accurate.

**It** is assumed that slow Givens is used throughout the calculation. If fast Givens [8] had been used, then the entries would become approximately  $2mn^2$  and  $2n^3$  instead (viz C-2).

(1) For the Householder transformations, each multiplication also invokes 1 addition and approximately 2 array addressings. For the **Givens** transformations, each multiplication invokes **1/2** an addition and 1 array addressing. On many large computers today, a floating point multiplication is not much slower than a floating point addition. Also, array indexing is usually quite expensive. In such cases, a Householder multiplication actually involves more work than a Givens multiplication because of the extra additions and array indexings. Therefore, the operation counts given for the diagonalization phase of GR-SVD may be misleading because it may actually involve relatively less work. The total effect, however, can be accounted for by using a smaller value for C. For example, if 1 Givens "multiplication" takes half the work needed by a Householder "multiplication", then the effect on the relative efficiency can be accounted for by

setting C-2 instead of C-4. On older or **non-scientific** machines where multiplications take much more time than additions and array addressings, the operation count based on multiplications alone is usually a good measure of relative efficiency.

(2) The application of  $S^{(1)T}$  and  $T^{(1)}$  on  $J^{(1)}$  is actually of order  $O(n^2)$  and is therefore not included in the above counts.

(3) We have to distinguish between 4 cases in the comparison:

Case a: both U and V are required explicitly,

Case b: only U is required explicitly,

case c: only V is required explicitly,

Case d: only  $\sum$  is required explicitly.

These four cases do arise in applications. We will mention a few here:

Case a arises in the computation of pseudo-inverses [1].

Case b is Case c for  $A^T$ .

Case c arises in least squares applications [1,3] and in the solution of homogeneous linear equations [1].

Case d arises in the estimation of the condition number of a matrix and in the determination of the rank of a matrix [10].

The total operation count for each case is given in Table 4.1 .

Table 4.1

Total operation counts of GR-SVD and MOD-SVD for each of the cases a, b, c, and d.

Case	GR-SVD	MOD-SVD
a	$(3+C)mn^2 + (C-1/3)*3$	$3mn^2 + (2C+4/3)n^3$
b	$(3+C)mn^2 - n^3$	$3mn^2 + (C+2/3)n^3$
c	$2mn^2 + Cn^3$	$mn^2 + (C+5/3)n^3$
d	$2mn^2 - 2n^3/3$	$mn^2 + n^3$

Using Table 4.1 , we can compute the ratio of the operation counts of MOD-SVD to that of GR-SVD for each of the four cases. This is given in Table 4.2 where the ratio is expressed as a function of  $r = m/n$ .

Table 4.2

Ratio of operation count of MOD-SVD to that of GR-SVD.

$$r = m/n$$

Case	Ratio	Cross-over point
a	$[3r + (2C+4/3)] / [(3+C)r + (C-1/3)]$	$(C+5/3)/C$
b	$[3r + (C+2/3)] / [(3+C)r - 1]$	$(C+5/3)/C$
c	$[r + (C+5/3)] / [2r + C]$	5/3
d	$[r + 1] / [2r - 2/3]$	5/3

These ratios are plotted in Fig. 4.1 to Fig. 4.4 for  $C=2, 3, 4$ .

The cross-over point  $r^*$  is the value of  $r$  which makes the ratio equal to 1. If  $r > r^*$ , then MOD-SVD is more efficient than CR-SVD.

From Figures 4.1 - 4.4, we see that, in all 4 cases a,b,c and d, MOD-SVD becomes more efficient than CR-SVD when  $r$  starts to get bigger than 2 approximately, and the savings can be as much as 50% when  $r$  is about 10. On the other hand, when  $r$  is about 1, CR-SVD is more efficient. This agrees with our earlier conjectures. However, the important

Fig. 4.1 Case a

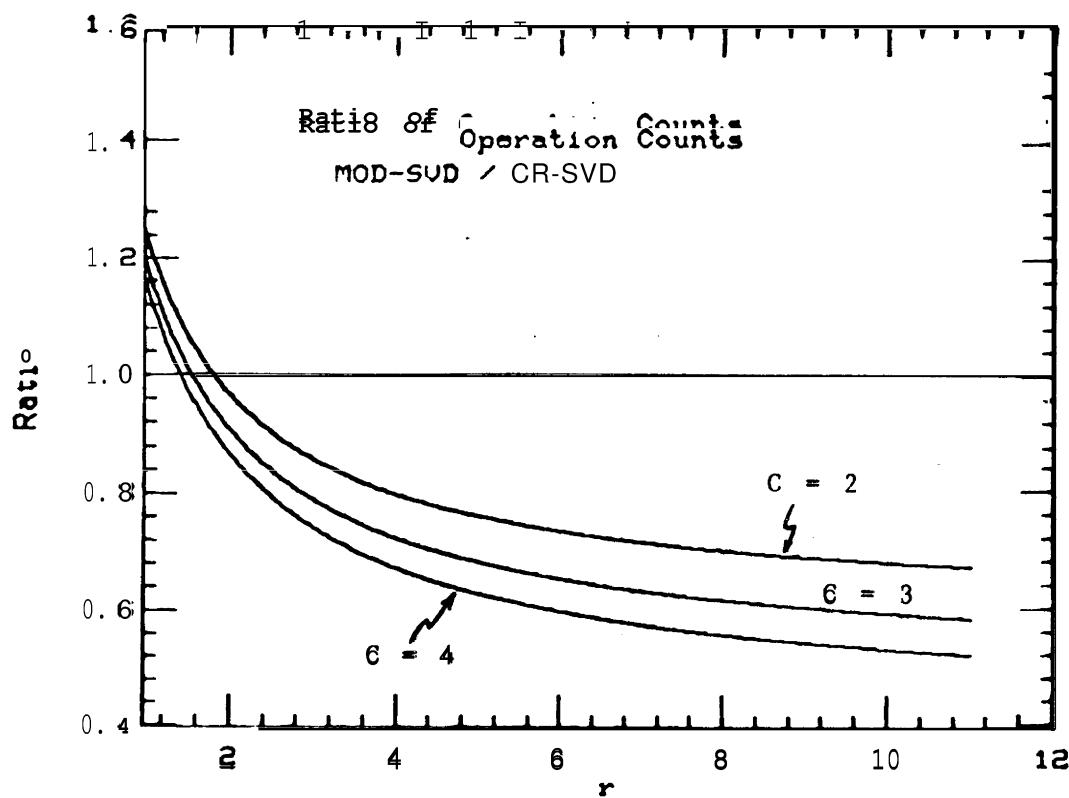


Fig. 4.2 Case b

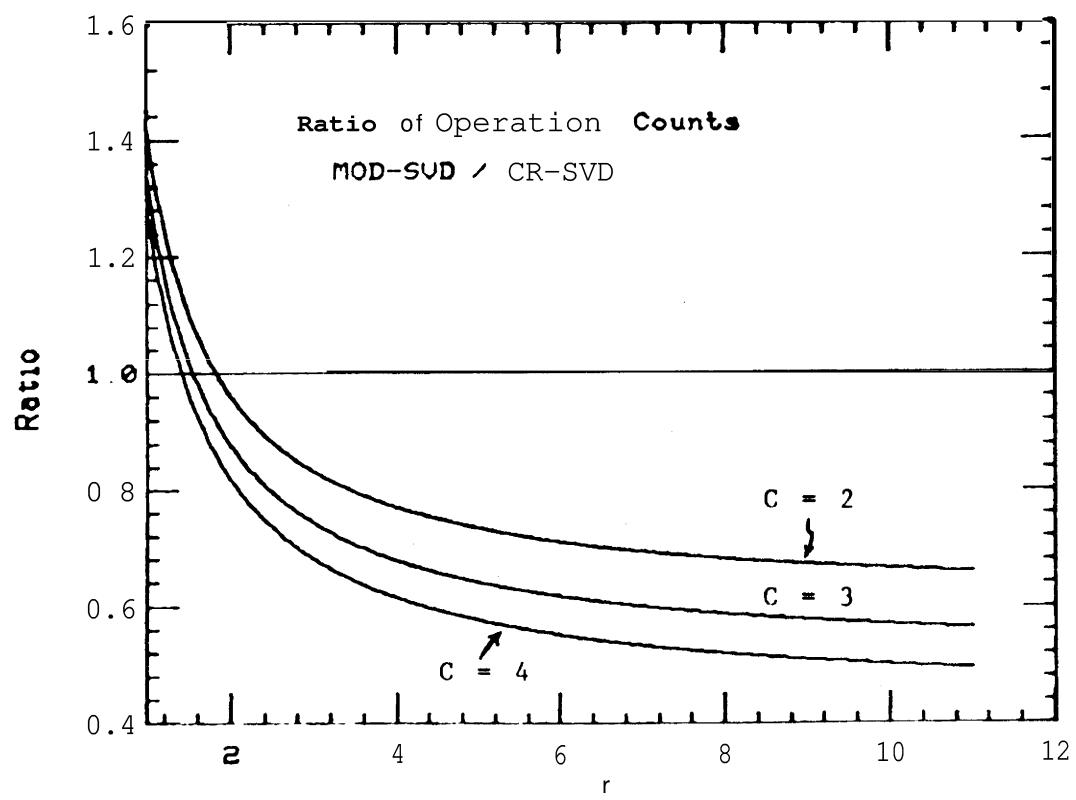


Fig. 4.3 Case c

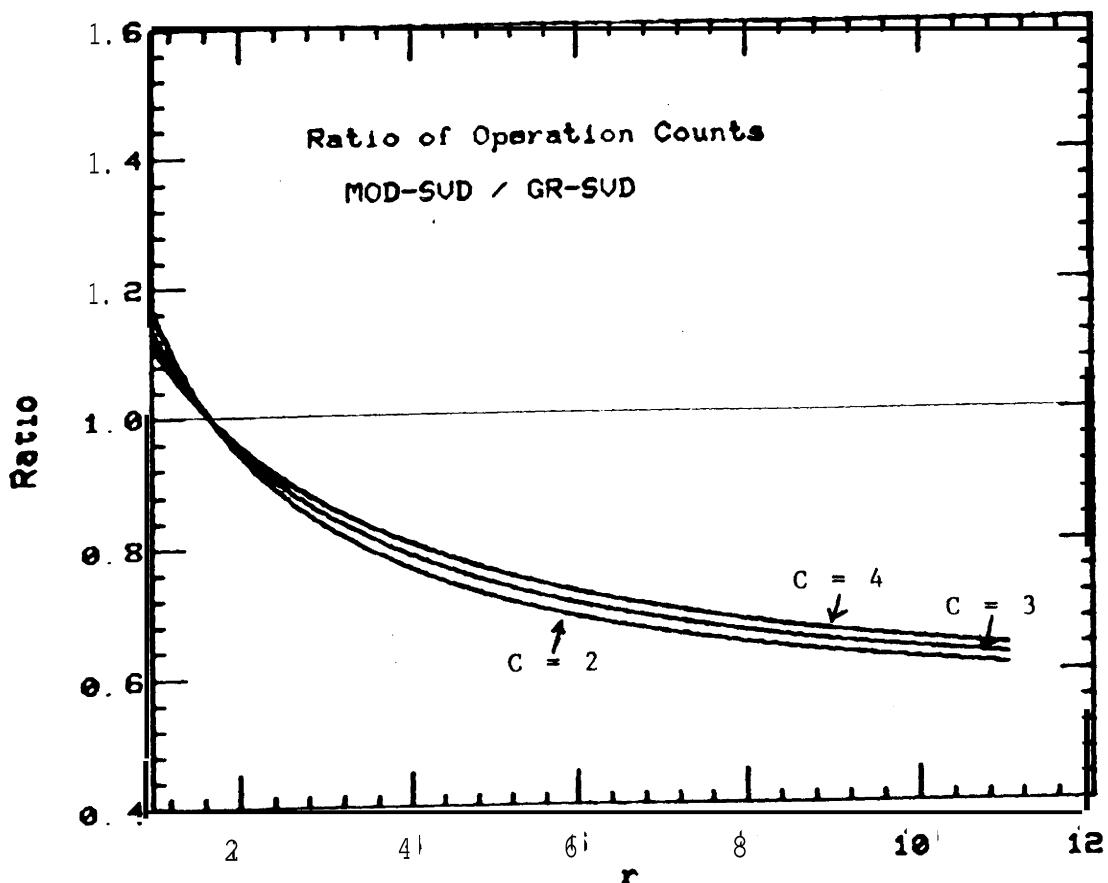
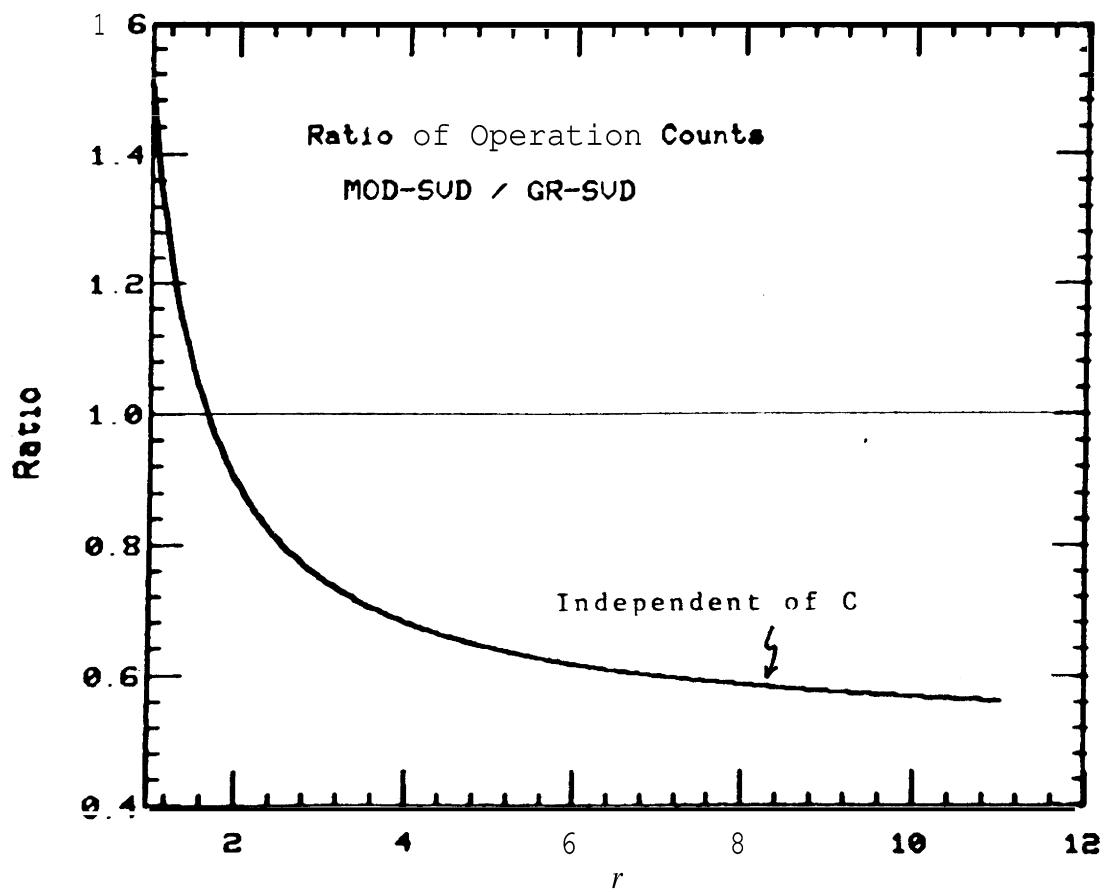


Fig. 4.4 Case d



thing is that all the curves decrease quite fast as  $r$  becomes large. If we assume that it is equally likely to encounter matrices with any value of  $r \geq 1$  (this is not an unreasonable assumption for designers of general mathematical software, for example), then MOD-SVD is obviously preferable. In any case, Fig. 4.1 - 4.4 give indications as to when one of the methods is more efficient, at least when  $m$  and  $n$  are large enough so that our operation counts apply.

In the context of least squares applications, we can also compare the operation counts of GR-SVD and MOD-SVD to that of the orthogonal triangularization methods [9] (OTLS) often used for such problems. This comparison is shown in Table 4.4 .

Table 4.4

Least squares using orthogonal triangularization versus  
**using** SVD

---

OTLS = orthogonal triangularization method  
 for least squares problems.

---

OTLS : GR-SVD =  $[r-1/3] / [2r+C]$

---

OTLS : MOD-SVD =  $[r-1/3] / [r+C+5/3]$

---

These ratios are plotted in Fig. 4.5 and Fig. 4.6 for  $C=2, 3, 4$  .

Fig. 4.5

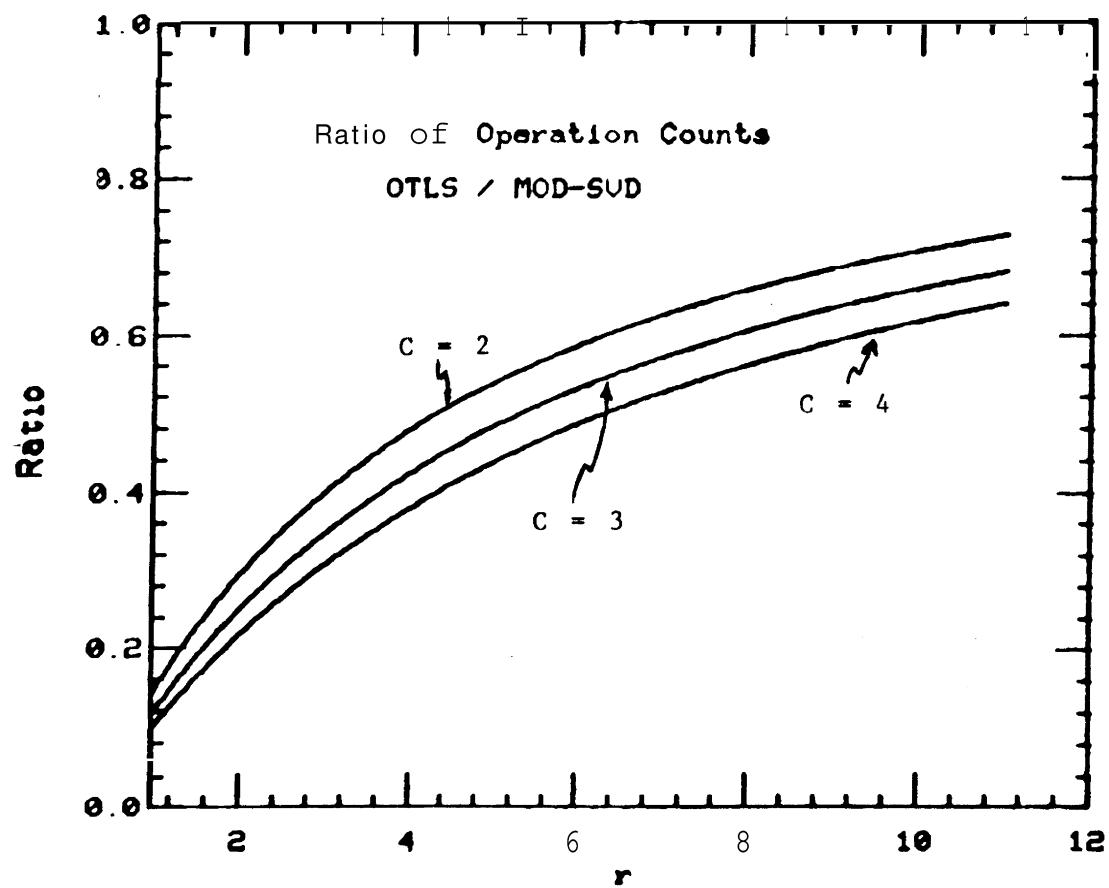
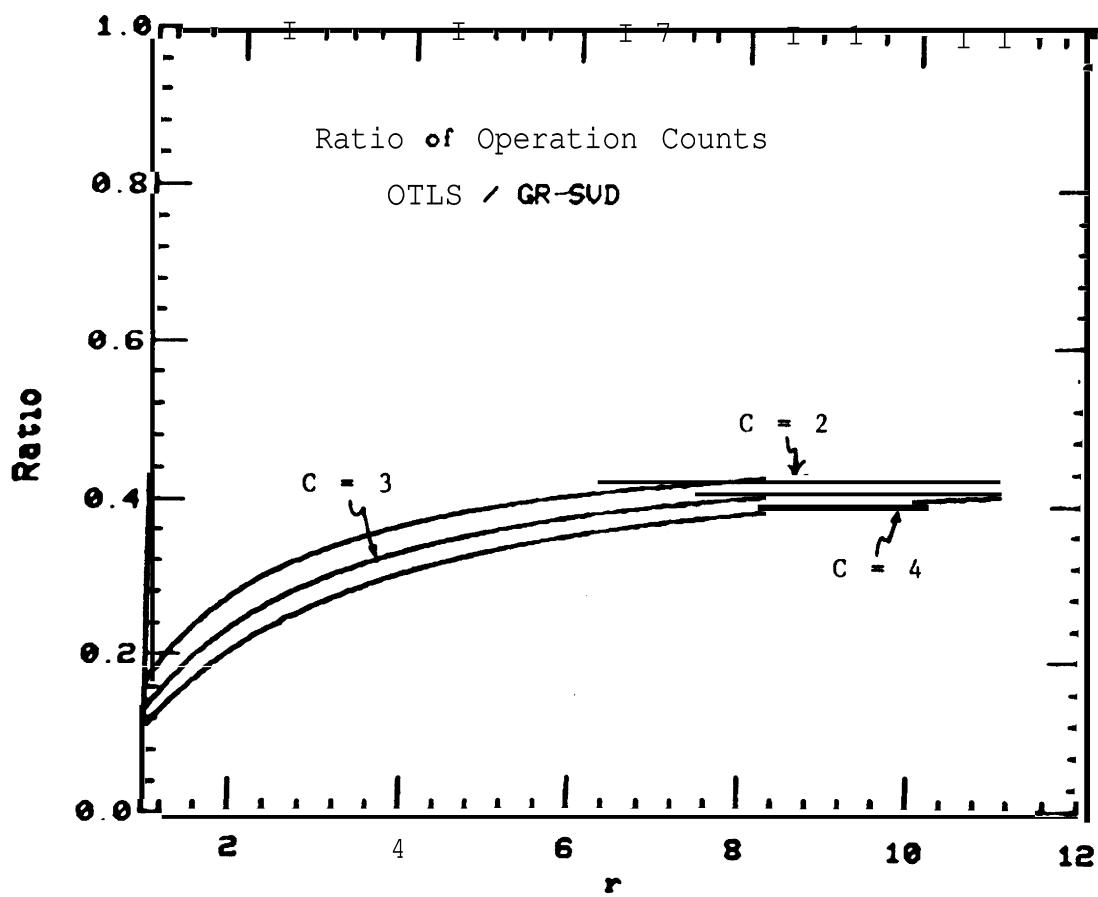


Fig. 4.6



One sees from these figures that for  $m$  nearly equal to  $n$ , the two SVD algorithms require much more work than OTLS. However, when  $r$  is bigger than about 3, MOD-SVD requires only about 3 times more work than OTLS. It may therefore become economically feasible to solve the least squares problems at hand by MOD-SVD instead of OTLS. The reward is that the SVD returns much more useful information about the problem than OTLS [3].

It is easy to see that as  $r$  becomes arbitrarily large, MOD-SVD is as efficient as OTLS since the bulk of the work is in the triangufarization of the data matrix  $A$ . However, GR-SVD can be at most half as efficient as OTLS.

## (5) COMPUTATIONAL RESULTS

The conclusions in the last section hold only if  $m$  and  $n$  are both large. In this section, some computational experiments are carried out to see if the conclusions are still valid for matrices with realistic sizes.

We computed the SVD of some randomly generated **matrices** using both GR-SVD and MOD-SVD. The version of GR-SVD that we used is a modified ALGOL W translation of the procedure that appeared in [1]. MOD-SVD is realized by writing a procedure to triangularize the input matrix by Householder transformations and then using the same above-mentioned GR-SVD procedure for computing the SVD of  $R$ .

All tests were run on the **IBM** 370/168's at the Stanford Linear Accelerator Center (SLAC). Long precision was used throughout the calculation. The mantissa of a floating point number is represented by 56 bits (approximately 16 decimal digits).

For each of the 4 cases, we fixed some values for  $n$  and computed the SVD of a sequence of randomly generated matrices with different values of  $r$ . The execution times taken by GR-SVD and MOD-SVD were then compared, together with the accuracies of the computed answers. Since we are working in a multi-programming environment, the execution times we measured cannot be taken as the

actual computing time taken. Moreover, the influence of the compiler on the relative efficiency of the two algorithms may be the deciding factor [11]. However, keeping these points in mind, we can still expect a qualitative agreement with the analysis based on operation counts.

On the IBM 370/168's at SLAC, a floating point multiplication takes only about 1.5 times the work taken for a floating point addition. Also, array indexing in ALGOL W is very expensive due to subscript checking (it actually can be more expensive than floating point multiplications). Therefore, as noted in section 4.1, we should use  $C$  approximately equal to 2 instead of 4 in Table 4.2 and Table 4.4, for the purpose of comparing the relative efficiency of the two algorithms based on the computational results.

The results of the computations are plotted in Fig. 5.1 - Fig. 5.6. In general, they agree very well qualitatively with the asymptotic results we obtained by operation counts (with  $C=2$ ). We observe that the larger  $n$  is the better the agreement, as it should be. However, even when  $n$  is small, the theoretical results based on asymptotic operation counts still describe very well the qualitative behavior of the computational results in many cases. The computational results also show that large savings in work are indeed realizable for reasonably-sized matrices (For example, see Fig. 5.3 and Fig. 5.4).

Fig. 5.1 Case a

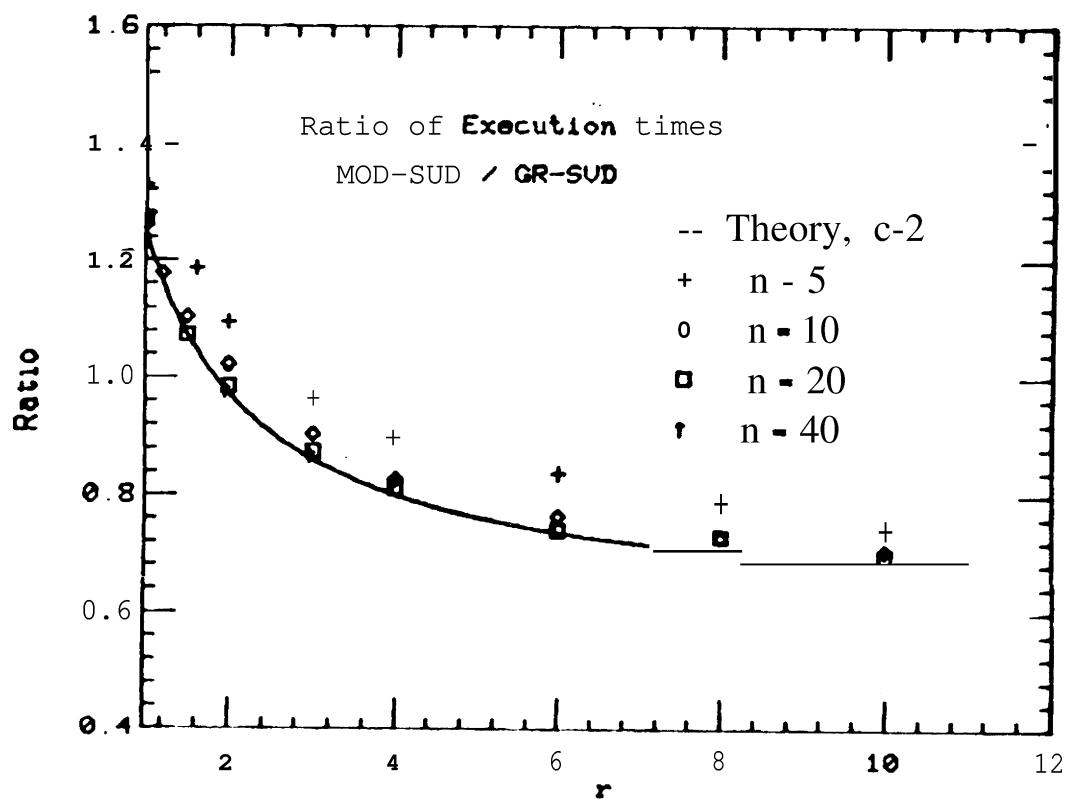


Fig. 5.2 Case b

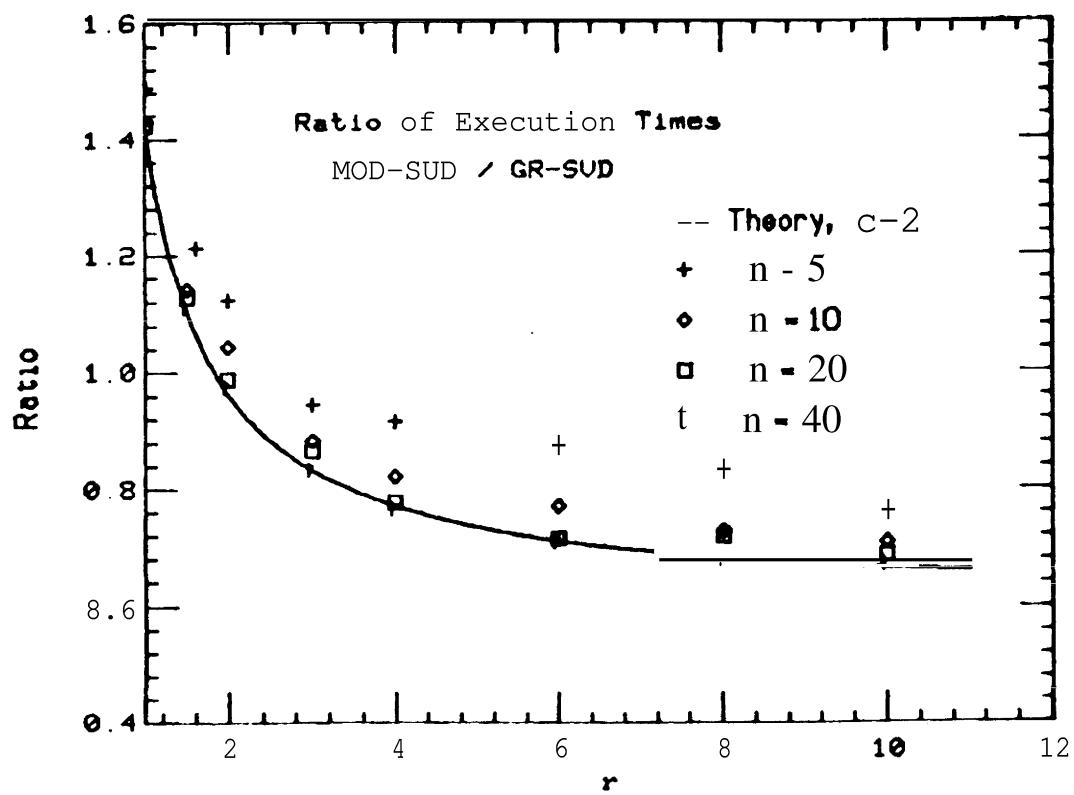


Fig. 5.3 Case c

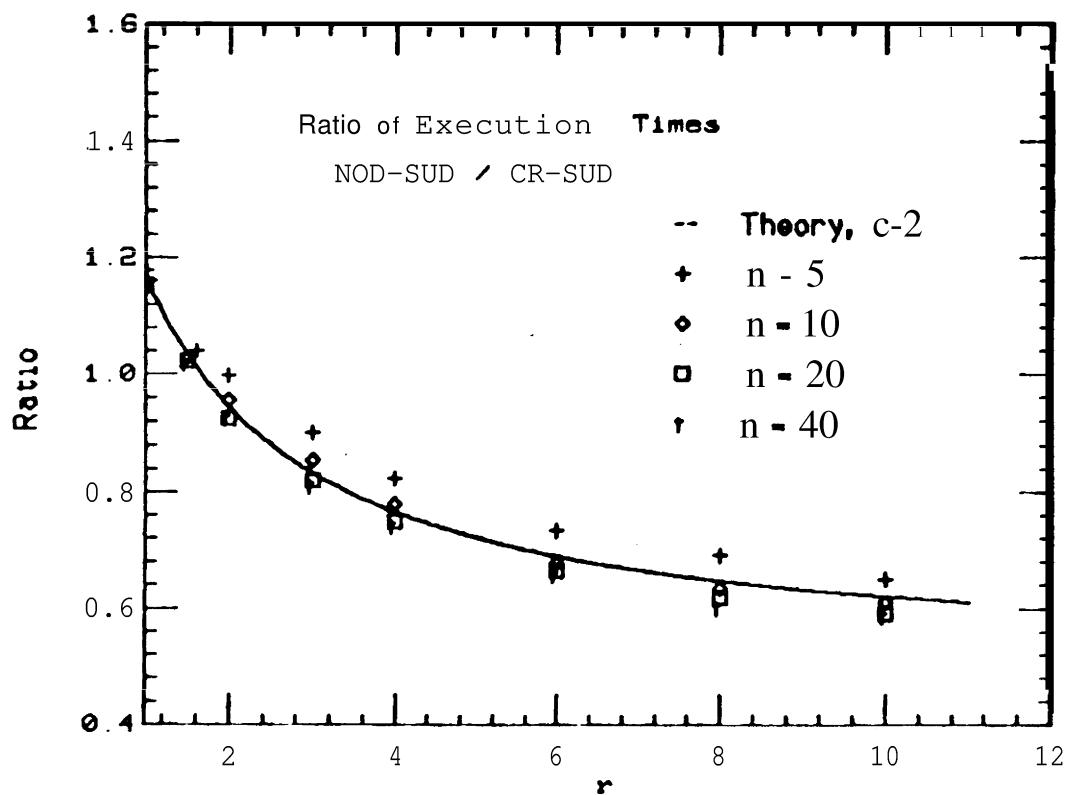


Fig. 5.4 Case d

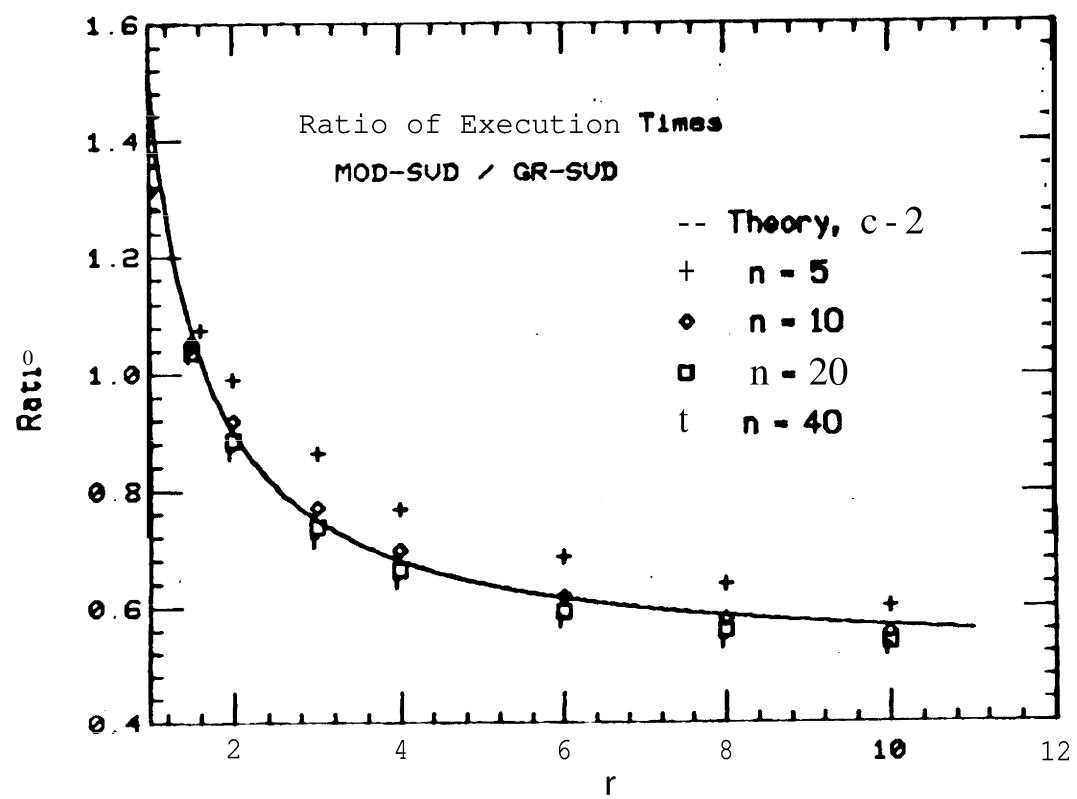


Fig. 5.5

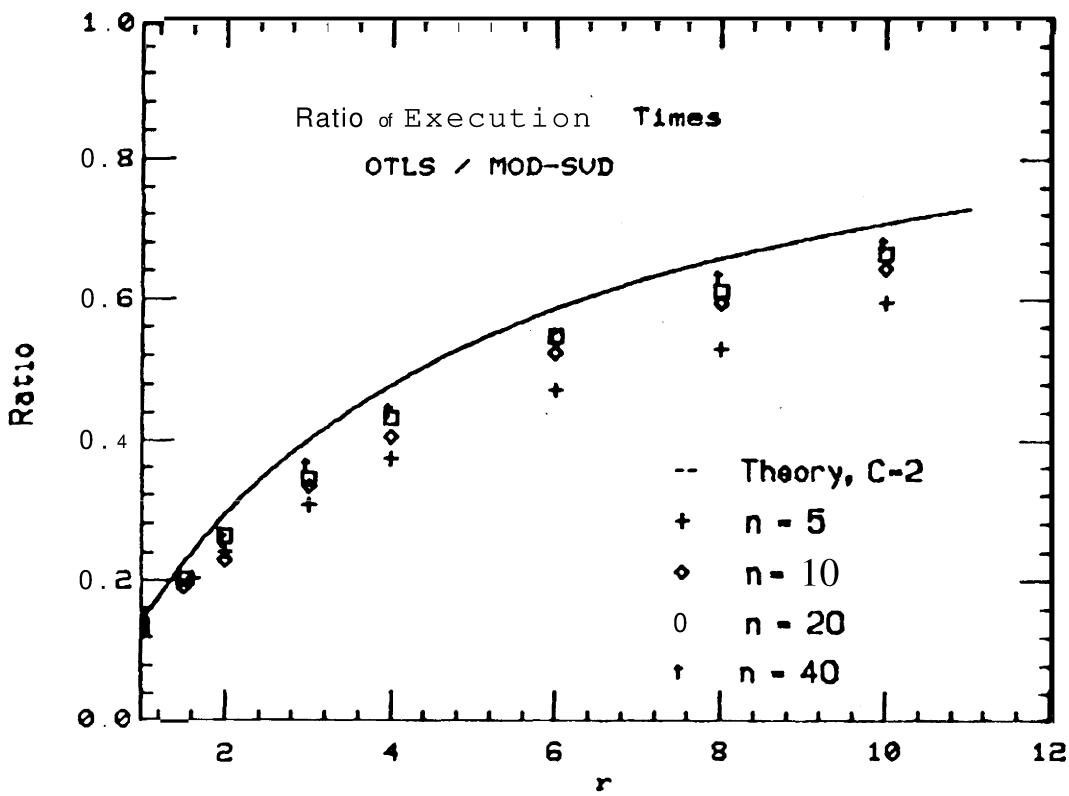
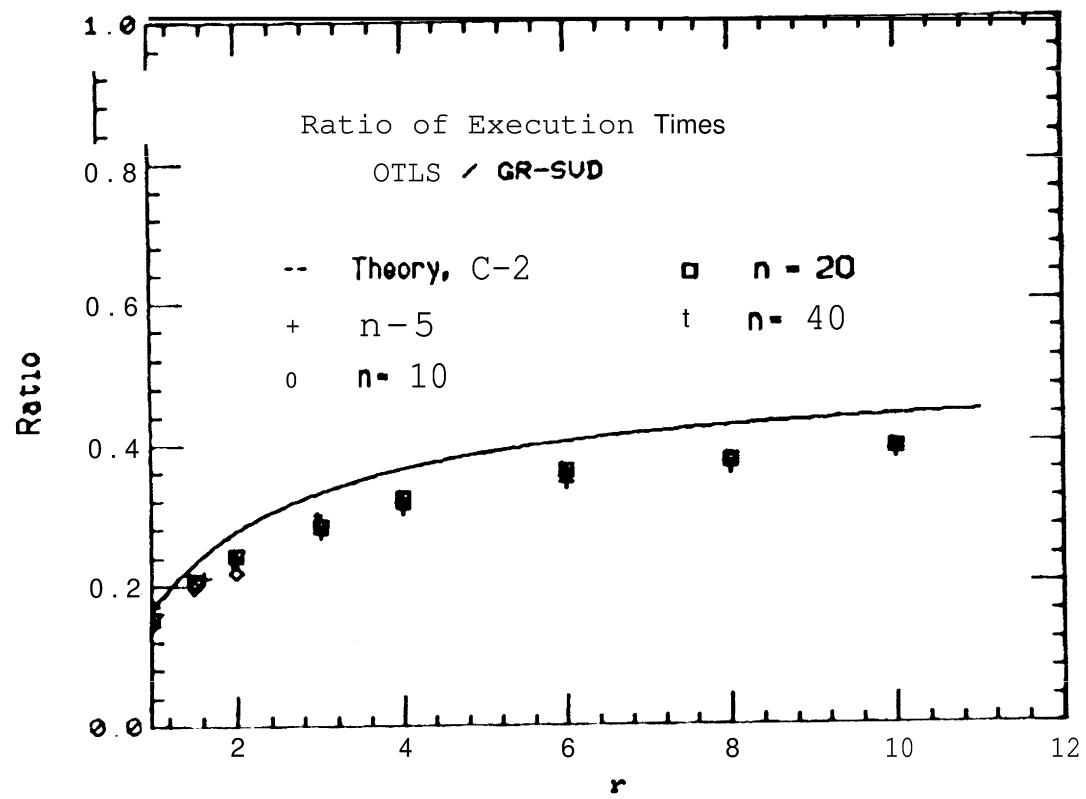


Fig. 5.6



We also checked the accuracies of the computed results, The singular values returned by both procedures GR-SVD and MOD-SVD agree to within a few units of the machine precision in almost all cases that we have tested. The matrices  $U$  and  $V$  also agree to the same precision but the signs of the corresponding columns may be reversed. However, the SVD is only unique to within such a sign change, so this is acceptable [10].

We also computed the singular values of the following  $30 \times 30$  matrix:

$$\begin{vmatrix} 1 & -1 & \dots & -1 \\ 1 & -1 & & \vdots \\ \vdots & \ddots & \ddots & \vdots \\ 0 & \ddots & \ddots & -1 \end{vmatrix}$$

This matrix is very ill-conditioned (with respect to computing its inverse) and is very close to being a matrix of rank 29 even though the determinant equals 1 for all values of  $n$ . The computed singular values from both GR-SVD and MOD-SVD agree exactly with those given in [1] to 15 significant digits (which are all the digit8 printed in ALGOL W).

## (6) CONCLUSIONS

**Firstly**, the theoretical results we obtained do seem to predict the actual computational efficiencies quite well, and they can therefore be used to indicate which algorithm to choose for a given matrix.

The MOD-SVD algorithm clearly works better than GR-SVD for matrices that have many more rows than columns. The price that MOD-SVD has to pay when  $m$  is nearly equal to  $n$  is not that big (usually **less** than 30%). We have also seen that the cost of solving a least squares problem by MOD-SVD can often be less than twice that of the usual orthogonal triangularization algorithms. It may therefore become economically feasible to solve many least squares problems by the SVD algorithms.

Some improvements can probably be made on the bidiagonalization of the upper triangular matrix  $R$  in MOD-SVD by taking advantage of the special structure of  $R$ . We also want to note again that MOD-SVD requires  $n^2$  extra storage locations if the left transformations have to be accumulated. This may be a disadvantage when storage is at a premium.

We have also seen that the usual practice of counting only multiplications in operation counts for numerical algorithms is no longer viable for many modern computers. Other properties, such as the amount of array accesses involved, may influence the efficiencies of algorithms decisively.

To be sure, there may be other ways to compute the SVD that will work better in some cases but not in others. It is perhaps impossible to find an "optimal" algorithm that works best for all matrices. Nevertheless, we hope this paper has shown that it may be worthwhile to look for improvements in the organizations of existing algorithms.

Appendix : Fortran Code of a Hybrid Algorithm

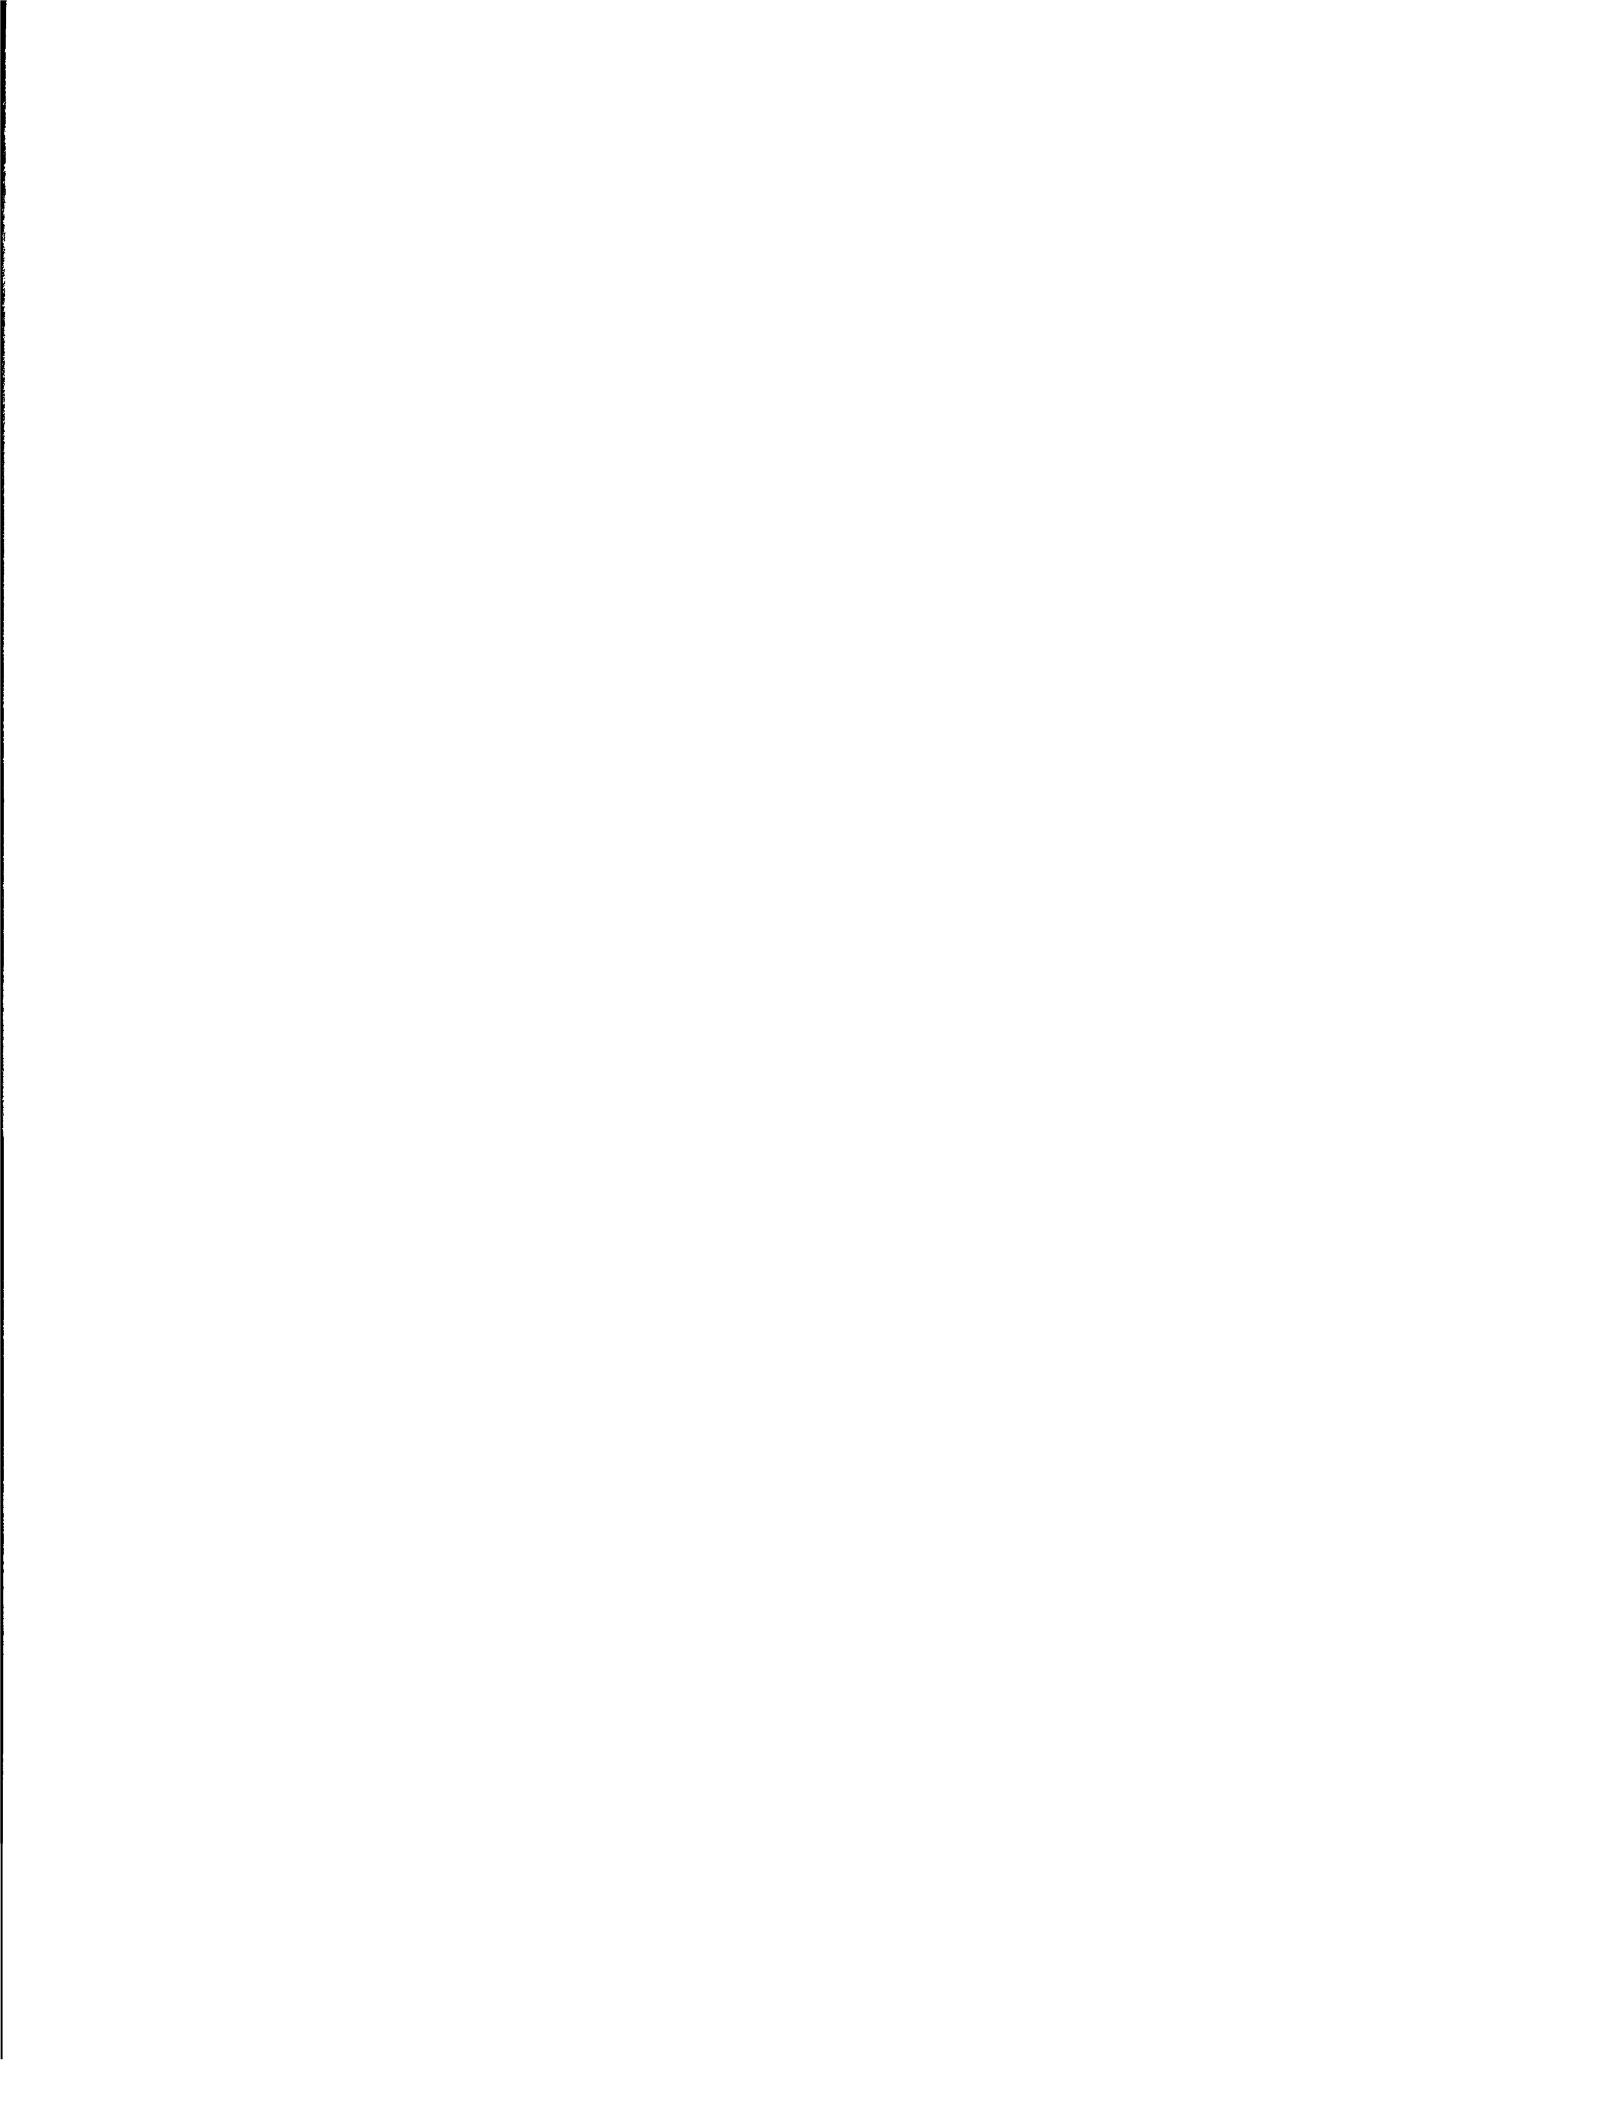
Based on the results of earlier sections, we can implement a hybrid method for computing the SVD of a rectangular matrix A which automatically chooses to use the more efficient algorithm between GR-SVD and MOD-SVD. For each of the four Cases a,b,c and d, if the input matrix A has a value of  $r (= m/n)$  which is less than the cross-over point  $r^*$  for that case, then we use GR-SVD, otherwise we use MOD-SVD. The cross-over points depend on the value of C used. As noted before, the value of C to be used depends on the relative efficiencies of floating point multiplications, floating point additions and array indexings on the particular machine concerned. However, C can be determined once for all for any particular machine and compiler combination. For example, if floating point multiplications take much more time than floating point additions and array indexings on the machine in question, then we should use C approximately equal to 4.

In this Appendix, we give the codes of a **Fortran** subroutine called HYBSVD which implements the above-mentioned hybrid algorithm. HYBSVD will need to call a standard Golub-Reinsch SVD subroutine during part of its computation and so we have included such a routine, called GRSVD, in the listing of the codes of HYBSVD.

The routine GRSVD is actually a slightly modified version of the subroutine SVD in the EISPACK [12] package. The main modification that we have made is to eliminate the requirement in subroutine SVD that the row dimension of V declared in the calling program be equal to that of A. This minimizes the storage requirements of GRSVD at the cost of one more argument in the argument list.

There is one additional feature implemented in HYBSVD (and also in GRSVD). In least squares applications, where we are looking for the minimal length least squares solution to the overdetermined linear system  $Ax = b$ , the left transformations  $U^T$  have to be accumulated on the right-hand side vectors  $b$  (there may be more than one  $b$ ). This can be done by putting the vectors  $b$  in the matrix argument B when calling HYBSVD and -setting IRHS to the number of  $b$ 's.

The calling sequences and usages of HYBSVD and GRSVD are explained in the comments in the beginning of the listings of the subroutines.



\*\*\*\*\* FIRST CARD OF HYBSVD\*\*\*\*\*

```
SUBROUTINE HYBSVD(NAU,NV,NZ,M,N,A,W,MATU,U,MATV,V,Z,B,IRHS,IER,  
,RV1)  
,  
INTEGER NAU,NV,NZ,M,N,IRHS,IER,IP1,I,J,K,IM1,IBACK  
DOUBLE PRECISION A(NAU,N),W(N),U(NAU,N),V(NV,N),Z(NZ,N),  
B(NAU,IRHS),RV1(N)  
DOUBLE PRECISION X0VRPT,C,R,G,SCALE,DSIGN,DABS,DSQRT,F,S,H  
REAL FLOAT  
LOGICAL MATU,MATV
```

THIS SUBROUTINE IS A MODIFICATION OF THE GOLUB-REINSCH PROCEDURE

(1) FOR COMPUTING THE SINGULAR VALUE DECOMPOSITION  $A = U W V^T$  OF A REAL  $M \times N$  RECTANGULAR MATRIX. THE ALGORITHM IMPLEMENTED IN THIS ROUTINE HAS A HYBRID NATURE. WHEN  $M$  IS APPROXIMATELY EQUAL TO  $N$ , THE GOLUB-REINSCH ALGORITHM IS USED, BUT WHEN  $M$  IS GREATER THAN APPROXIMATELY  $2N$ , A MODIFIED VERSION OF THE GOLUB-REINSCH ALGORITHM IS USED. THIS MODIFIED ALGORITHM FIRST TRANSFORMS  $A$  INTO UPPER TRIANGULAR FORM  $B$  BY HOUSEHOLDER TRANSFORMATIONS  $L$  AND THEN USES THE GOLUB-REINSCH ALGORITHM TO FIND THE SINGULAR VALUE DECOMPOSITION OF THE RESULTING UPPER TRIANGULAR MATRIX  $R$ . WHEN  $U$  IS NEEDED EXPLICITLY, AN EXTRA ARRAY  $Z$  (OF SIZE AT LEAST  $N \times N$ ) IS NEEDED, BUT OTHERWISE  $Z$  MAY COINCIDE WITH EITHER  $A$  OR  $V$  AND NO EXTRA STORAGE IS REQUIRED. THIS HYBRID METHOD SHOULD BE MUCH MORE EFFICIENT THAN THE GOLUB-REINSCH ALGORITHM WHEN  $M$  IS MUCH BIGGER THAN  $N$ . FOR DETAILS, SEE (2).

HYBSVD CAN ALSO BE USED TO COMPUTE THE MINIMAL LENGTH LEAST SQUARES SOLUTION TO THE OVERDETERMINED LINEAR SYSTEM  $A^*X = B$ .

NOTICE THAT THE SINGULAR VALUE DECOMPOSITION OF A MATRIX IS UNIQUE ONLY UP TO THE SIGN OF THE CORRESPONDING COLUMNS OF  $U$  AND  $V$ .

THIS ROUTINE HAS BEEN CHECKED BY THE PFORTRAN VERIFIER (3) FOR ADHERENCE TO A LARGE, CAREFULLY DEFINED, PORTABLE SUBSET OF AMERICAN NATIONAL STANDARD FORTRAN CALLED PFORTRAN.

#### REFERENCES:

- (1) GOLUB, G.H. AND REINSCH, C. (1970) "SINGULAR VALUE DECOMPOSITION AND LEAST SQUARES SOLUTIONS," *NUMER. MATH.* 14, 403-420, 1970.
- (2) CHAN, T.F. (1976) "ON COMPUTING THE SINGULAR VALUE DECOMPOSITION," TO APPEAR AS A STANFORD COMPUTER SCIENCE REPORT.
- (3) FYDER, B.G. (1974) "THE PFORTRAN VERIFIER," *SOFTYARD PRACTICE AND EXPERIENCE*, VOL. 4, 359-377, 1974.

HYBSVD ASSUMES  $M \geq N$  AND  $N \neq 0$ . THEN COMPUTE THE SINGULAR VALUE DECOMPOSITION OF  $A$ . IF  $A = U W V^T$ , THEN  $A = V W U^T$ .

ON INPUT:

NAU MUST BE SET TO THE ROW DIMENSION OF THE TWO-DIMENSIONAL ARRAY PARAMETERS A, U AND V AS DECLARED IN THE CALLING PROGRAM DIMENSION STATEMENT. NOTE THAT NAU MUST BE AT LEAST AS LARGE AS M;

NV MUST BE SET TO THE ROW DIMENSION OF THE TWO-DIMENSIONAL ARRAY PARAMETER V AS DECLARED IN THE CALLING PROGRAM DIMENSION STATEMENT. NV MUST BE AT LEAST AS LARGE AS N;

NZ MUST BE SET TO THE ROW DIMENSION OF THE TWO-DIMENSIONAL ARRAY PARAMETER Z AS DECLARED IN THE CALLING PROGRAM DIMENSION STATEMENT. NOTE THAT NZ MUST BE AT LEAST AS LARGE AS N;

M IS THE NUMBER OF ROWS OF A (AND U);

N IS THE NUMBER OF COLUMNS OF A (AND U) AND THE ORDER OF V;

A CONTAINS THE RECTANGULAR INPUT MATRIX TO BE DECOMPOSED;

B CONTAINS THE IRHS RIGHT-HAND-SIDES OF THE OVERRDETERMINED LINEAR SYSTEM A\*X=B. IF IRHS.GT.0, THEN ON OUTPUT, THESE IRHS COLUMNS IN B

WILL CONTAIN U\*B. THUS, TO COMPUTE THE MINIMAL LENGTH LEAST SQUARES SOLUTION, ONE MUST COMPUTE V \* W<sup>+</sup> TIMES THE COLUMNS OF B, WHERE W<sup>+</sup> IS A DIAGONAL MATRIX, W<sup>+</sup>(I)=0 IF W(I) IS NEGLIGIBLE, OTHERWISE IS 1/W(I). IF IRHS=0, B M A Y COINCIDE WITH A OR U AND WILL NOT BE REFERENCED;

IRHS IS THE NUMBER OF RIGHT HAND-SIDES OF THE OVERRDETERMINED SYSTEM A\*X=B. IRHS SHOULD BE SET TO ZERO IF ONLY THE SINGULAR VALUE DECOMPOSITION OF A IS DESIRED;

MATU SHOULD BE SET TO .TRUE. IF THE U MATRIX IN THE DECOMPOSITION IS DESIRED, AND TO .FALSE. OTHERWISE;

MATV SHOULD BE SET TO .TRUE. IF THE V MATRIX IN THE DECOMPOSITION IS DESIRED, AND TO .FALSE. OTHERWISE.

WHEN HYBSVO IS USED TO COMPUTE THE MINIMAL LENGTH LEAST SQUARES SOLUTION TO AN OVERRDETERMINED SYSTEM, MATU SHOULD BE SET TO .FALSE., AND MATV SHOULD BE SET TO .TRUE..

ON OUTPUT:

4 IS UNALTERED (UNLESS OVERWRITTEN BY UORV);

W CONTAINS THE (NON NEGATIVE) SINGULAR VALUES OF A (THE DIAGONAL ELEMENTS OF W). THEY ARE UNORDERED. IF AN ERROR EXIT IS MADE, THE SINGULAR VALUES SHOULD BE CORRECT FOR INDICES IERR+1, IERR+2, . . . M, N;

U CONTAINS THE MATRIX U (ORTHOGONAL COLUMN VECTORS) OF THE DECOMPOSITION IF MATU IS BEEN SET TO .TRUE. OTHERWISE U IS USED AS A TEMPORARY ARRAY. U MAY COINCIDE WITH A. IF AN ERROR EXIT IS MADE, THE COLUMNS OF U CORRESPONDING TO INDICES OF CORRECT SINGULAR VALUES SHOULD BE CORRECT;

V CONTAINS THE MATRIX V (ORTHOGONAL) OF THE DECOMPOSITION IF  
MAT V HAS BEEN SET TO .TRUE. OTHERWISE V IS NOT REFERENCED.  
V MAY ALSO COINCIDE WITH A IF U IS NOT NEEDED. IF AN ERROR  
EXIT IS MADE, THE COLUMNS OF V CORRESPONDING TO INDICES OF  
CORRECTS SINGULAR VALUES SHOULD BE CORRECTED.

Z CONTAINS THE MATRIX X IN THE SINGULAR VALUE DECOMPOSITION  
OF R=XS<sup>T</sup>Y, IF THE MODIFIED ALGORITHM IS USED. IF THE  
GOLUB-FEINSCH PROCEDURE IS USED, THEN IT IS NOT REFERENCED.  
IF MAT U HAS BEEN SET TO .FALSE., Z MAY COINCIDE  
WITH A OR V AND IS NOT REFERENCED;

IERR IS SET TO  
ZERO                   FCR NORMAL RETURN.  
K                    IF THE K-TH SINGULAR VALUE HAS NOT BEEN  
DETERMINED AFTER 30 ITERATIONS:  
  1            IF IRHS .LT. 0 .  
  2            IF M .LT. N .  
  3            IF NAU .LT. M .  
  4            IF NV .LT. N .  
  5            IF NZ .LT. N .

RV1 IS A TEMPORARY STORAGE ARRAY.

PROGRAMMED BY TONY CHAN, COMP-SCI-DEPT.  
STANFORD UNI V. CA 94305  
LAST MODIFIED: 12 SEPTEMBER, 1976.

```
IERR=0
IF(IRHS.GE.0) GO TO 2
IERR=1
RETURN
2 IF(M.GE.N) GO TO 3
IERR=2
RETURN
3 IF(NAU.GE.M) GO TO 4
IERR=3
RETURN
4 IF(NV.GE.N) GO TO 5
IERR=4
RETURN
5 IF(NZ.GE.N) GO TO 6
IERR=5
RETURN
6 CONTINUE
```

SET VALUE FOR C. THE VALUE FOR C DEPENDS ON THE RELATIVE  
EFFICIENCY OF FLOATING POINT MULTIPLICATIONS, FLOATING POINT  
ADDITIONS AND TWO DIMENSIONAL ARRAY INDEXINGS ON THE  
COMPUTER WHERE THIS SUBROUTINE IS TO BE RUN. C SHOULD  
USUALLY BE BETWEEN 2 AND 4. FOR DETAILS ON CHOOSING C, SEE  
(2). THE ALGORITHM IS NOT SENSITIVE TO THE VALUE 0 = C  
ACTUALLY USED AS LONG AS C IS BETWEEN 2 AND 4.

C = 4.000

DETERMINING CROSS-OVER POINT

```

C
IF (MATU .AND. MATV) X O V R P T = (C+5.D0/3.D0)/C
I  F (MATU .AND. .NOT. MATV ) X O V R P T = (C+5.D0/3.D0)/C
I  F (.NOT. MATU . A N D . MATV) X O V R P T = 5.D0/3.D0
I  F (.NOT. MATU .AND..NOT. MATV) X O V R P T = 5.D0/3.D0

C
C DETERMINE WHETHER TO USE GOLUB-REINSCH OR THE MODIFIED
C ALGORITHM.
C
R = FLOAT(M)/FLCAT(N)
I F(R. G E . XCVFPT) GO TO 8
C
C USE GOLUBREINSCH PROCEDURE
C
CALL GRSVD(NAU,NV,M,N,A,W,MATU,U,MATV,V,B,IRHS,IERR,RV1)
FETURN

C
C USE MODIFIED ALGCRITHM

8  DO 10 I=1,M
    D C 10 J=1,N
10      U(I,J)=A(I,J)

C
C TRIANGULARIZE U BY HOUSEHOLDER TRANSFORMATIONS ,USING
C W AND RV1 A TEMPORARY STORAGE.
C
DO ? C I=1,N
G=0.0D0
S=0.0D0
SCALE=0.0D0

C
C PERFORM SCALING OF COLUMNS TO AVOID UNNECESSARY OVERFLOW
C OR UNDERFLOW

DO 3 0 K=I,M
    SCALE = SCALE + DABS(U(K,I))
3 0 IF (SCALE.EQ. 0.0D0) GO TO 20
0 0 4 0 K=I,M
    U(K,I) = U(K,I)/SCALE
    S = 3 + U(K,I)**2
4 0 CONTINUE

C
C THE VECTOR E OF THE HOUSEHOLDER TRANSFORMATION I + EE*/H
C WILL BE RESTORED IN COLUMN I OF U. THE TRANSFORMED ELEMENT
C U(I,I) WILL BE RESTORED IN W(I) AND THE SCALAR H IN
C RV1(I).

C
F = U(I,I)
G = -DSIGN(DSQRT(S),F)
H = F*G - S
U(I,I) = F * G
RV1(I) = H
W(I) = SCALE * G

I F(I.EQ. N) GO TO 85

C
C APPLY TRANSFCRMATIONS TO REMAINING COLUMNS OF A

IP1 = I + 1
0 0 5 0 J=IP1,N

```

```

S = 0.0DC
63   DO 60 K=I,M
      S = S + U(K,I)*U(K,J)
      F = S/H
      DO 70 K=I,M
      U(K,J) = U(K,J) + F*U(K,I)
70   CONTINUE
50   CONTINUE
C
C   APPLY TRANSFORMATIONS TO COLUMNS OF B IF IRHS .GT. 0
C
85   IF (IRHS.EQ.0) GO TO 20
D   O 80 J=1,IRHS
      S = 0.000
      DO 90 K=I,M
      S = S + U(K,I)*B(K,J)
90   F = S/H
      DO 100 K=I,M
      B(K,J) = B(K,J) + F*U(K,I)
100  CONTINUE
90  CONTINUE
20  CONTINUE
C
C   COPY R INTO Z IF MATU = .TRUE.
C
110  IF (.NOT.MATU) GO TO 300
DO 110 I=1,N
      D   O 110 J=1,N
      I   F(J.GE. I) GO TO 112
      Z(I,J) = 0.900
      GO TO 110
112  I   F(J.EQ. I) GO TO 114
      Z(I,J) = U(I,J)
      GO TO 110
114  Z(I,I) = W(I)
10  CONTINUE
C
C   ACCUMULATE HOUSEHOLDER TRANSFORMATIONS IN U
C
120  D   O 120 IBACK=1,N
      I = N - IBACK + 1
      IP1 = I + 1
      G = W(I)
      H = RV1(I)
      IF (I.EQ.N) GO TO 130
C
140  D   C 140 J=IP1,N
      U(I,J) = 0.300
C
130  I   F(H.EQ. C.000) GO TO 150
      I   F(I.EQ. N) GO TO 160
C
170  D   O 170 J=IP1,N
      S = 0.0DC
      DO 180 K=IP1,M
      S = S + U(K,I)*U(K,J)
      F = S/H
      DO 170 K=I,M
      U(K,J) = U(K,J) + F*U(K,I)
70   CONTINUE

```



C . . . . . FIRST CARD OF GRSVD . . . . .  
 C  
 C SUBROUTINE GRSVD(NAU,NV,M,N,A,W,MATU,U,MATV,V,B,IRHS,IERR,RV1)  
 C  
 C INTEGER I,J,K,L,M,N,II,II,KK,K1,LL,L1,MN,NAU,NV,ITS,IERR,IRHS  
 C DOUBLE PRECISION A(NAU,N),W(N),U(NAU,N),V(NV,N),B(NAU,IRHS),RV1(N)  
 C DOUBLE PRECISION C,F,G,H,S,X,Y,Z,EPS,SCALE,MACHEP  
 C DOUBLE PRECISION DSQRT,DMAX1,DABS,DSIGN  
 C LOGICAL MATU,MATV  
 C  
 C THIS SUBROUTINE IS A TRANSLATION OF THE ALGOL PROCEDURE SVD.  
 C NUM. MATH. 14, 403-420 (1970) BY GOLUB AND DREINSCH.  
 C HANDBOOK FOR AUTO. COMP., VOL I-F - LINEAR ALGEBRA, 134-151 (1971).  
 C  
 C THIS SUBROUTINE DETERMINES THE SINGULAR VALUE DECOMPOSITION  
 C  
 C A=UWV<sup>T</sup> OF A REAL M BY N RECTANGULAR MATRIX. HOUSEHOLDER  
 C BIDIAGONALIZATION AND A VARIANT OF THE QR ALGORITHM ARE USED.  
 C GRSVD ASSUMES M .GE. N. IF M .LT. N, THEN COMPUTE THE SINGULAR  
 C  
 C VALUE DECOMPOSITION OF A. IF A =UWV<sup>T</sup> THEN A=VWU<sup>T</sup>.  
 C  
 C GRSVD CAN ALSO BE USED TO COMPUTE THE MINIMAL LENGTH LEAST SQUARES  
 C SOLUTION TO THE OVERRDETERMINED LINEAR SYSTEM A\*X=B.  
 C  
 C ON INPUT:  
 C  
 C NAU MUST BE SET TO THE ROW DIMENSION OF THE TWO-DIMENSIONAL  
 C ARRAY PARAMETERS A, U AND B AS DECLARED IN THE CALLING PROGRAM  
 C DIMENSION STATEMENT. NOTE THAT NAU MUST BE AT LEAST  
 C AS LARGE AS M;  
 C  
 C NV MUST BE SET TO THE ROW DIMENSION OF THE TWO-DIMENSIONAL  
 C ARRAY PARAMETER V AS DECLARED IN THE CALLING PROGRAM  
 C DIMENSION STATEMENT. NV MUST BE AT LEAST AS LARGE AS N;  
 C  
 C M IS THE NUMBER OF ROWS OF A (AND U);  
 C  
 C N IS THE NUMBER OF COLUMNS OF A (AND U) AND THE ORDER OF V;  
 C  
 C A COULD CONTAIN THE RECTANGULAR INPUT MATRIX TO BE DECOMPOSED;  
 C  
 C B CONTAINS THE IRHS RIGHT-HAND-SIDES OF THE OVERRDETERMINED  
 C LINEAR SYSTEM A\*X=B. IF IRHS .GT. 0,  
 C THEN ON OUTPUT, THESE IRHS COLUMNS  
 C  
 C WILL CONTAIN U<sup>T</sup> B. THUS, TO COMPUTE THE MINIMAL LENGTH LEAST  
 C SQUARED SOLUTION. ONE MUST COMPUTE V\*W<sup>T</sup> TIMES THE COLUMNS OF  
 C  
 C B, WHERE W<sup>T</sup> IS A DIAGONAL MATRIX, W<sup>T</sup>(I)=0 IF W(I) IS  
 C NEGLIGIBLE. OTHERWISE IS 1/W(I). IF IRHS=0, B MAY COINCIDE  
 C WITH A OR U AND WILL NOT BE REFERENCED;  
 C  
 C IPHS IS THE NUMBER OF RIGHT-HAND-SIDES OF THE OVERRDETERMINED  
 C SYSTEM A\*X=B. IRHS SHOULD BE SET TO ZERO IF ONLY THE SINGULAR  
 C VALUE DECOMPOSITION OF A IS DESIRED;

**MATU** SHOULD BE SET TO •TRUE• IF THE U MATRIX IN THE DECOMPOSITION IS DESIRED, AND TO •FALSE• OTHERWISE:

MATV SHOULD BE SET TO .TRUE. IF THE V MATRIX IN THE DECOMPOSITION IS DESIRED, AND TO .FALSE. OTHERWISE.

## ON OUTPUT:

A IS UNALTERED (UNLESS OVERWRITTEN BY U OR R);

W CONTAINS THE N (NON-NEGATIVE) SINGULAR VALUES OF A (THE  
DIAGONAL ELEMENTS OF W). THEY ARE UNORDERED. IF AN  
ERROR EXIT IS MADE, THE SINGULAR VALUES SHOULD BE CORRECT  
FOR INDICES IERR+1, IERR+2, . . . , N;

U CONTAINS THE MATRIX U (ORTHOGONAL COLUMN VECTORS1 OF THE DECOMPOSITION IF MATU HAS BEEN SET TO .TRUE. OTHERWISE U KS USED AS A TEMPORARY ARRAY. U MAY COINCIDE WITH A. IF AN ERRO? EXIT IS MADE. THE COLUMNS OF U CORRESPONDING TO INDICES OF CORRECT SINGULAR VALUES SHOULD BE CORRECT;

V CONTAINS THE MATRIX V (ORTHOGONAL) OF THE DECOMPOSITION IF  
MATV HAS BEEN SET TO .TRUE.. OTHERWISE VIS NOT REFERENCED.  
V MAY ALSO COINCIDE WITH A IF U IS NOT NEEDED. IF AN ERROR  
EXIT IS MADE, THE COLUMNS OF V CORRESPONDING TO INDICES OF  
CORRECT SINGULAR VALUES SHOULD BE CORRECT;

## TERRISSSETTO

```

ZER0 FOR NDRLAL RETURN,
K IF THE K-TH SINGULAR VALUE HAS NOT BEEN
DETERMINED AFTER 30 ITERATIONS;
-1 IF IRHS .LT.. 0 .
-2 IF M .LT. N .
-3 IF NNU .LT. M .
4 IF NV .LT. N .

```

RV1 IS A TEMPORARY STORAGE ARRAY.

THIS SUBROUTINE HAS BEEN CHECKED BY THE PPORT VERIFIER  
(RYDER, B.G. "THE PPORT VERIFIER", SOFTWARE - PRACTICE AND  
EXPERIENCE, VOL. 4, 359-377, 1974) FOR ADHERENCE TO A LARGE,  
CAREFULLY DEFINED, PORTABLE SUBSET OF AMERICAN NATIONAL STANDARD  
FORTRAN CALLED PPORT.

ORIGINAL VERSION OF THIS CODE IS SUBROUTINE SVD IN RELEASE 2 OF EISPACK.

MODIFIED BY TONY CHAN. COMP. SCI. DEPT. STANFORD UNIV., CA94305.  
LAST MODIFIED: 2 SEPTEMBER, 1976

MACHEP IS A MACHINE DEPENDENT PARAMETER SPECIFYING  
THE RELATIVE PRECISION OF FLOATING POINT ARITHMETIC.  
MACHEP = 16.0D0\*(-13) FOR LONG FORM ARITHMETIC  
ON 5360

DATA MATCHUP/2-22D-15/

IERR = 3  
IF (IRHS.GE.0) GC TO 2

```

IERR=-1
RETURN
2 IF (M .GE. N ) GO TO 3
IERR=-2
RETURN
3 IF (NAU .GE. M) GO TO 4
IERR=-3
RETURN
4 IF (NV .GE. N) GO TO 5
IERR=-4
RETURN
5 CONTINUE
C
C      DO 100 I = 1, M
C
C          DO 190 J = 1, N
C              U(I,J) = A(I,J)
100 CONTINUE
C      :::::::::::::: HOUSEHOLDER REDUCTION TO BIDIAGONAL FORM ::::::::::::
C      G = 0.0D0
C      SCALE = 0.0D0
C      X = 0.0D0
C
C      DO 303 I = 1, N
C          L = I + 1
C          RV1(I) = SCALE * G
C          G = 0.0D0
C          S = 0.0D0
C          SCALE = 0.0D0
C
C      COMPUTE LEFT TRANSFORMATIONS THAT ZEROS THE SUBDIAGONAL ELEMENTS
C      OF THE IITH COLUMN.
C
C      DO 120 K = I, M
C          SCALE = SCALE + DABS(U(K,I))
C
C          IF (SCALE.EQ.0.0D0) GO TO 210
C
C          DO 130 K = I, M
C              U(K,I) = U(K,I) / SCALE
C              S = S + U(K,I)**2
130 CONTINUE
C
C          F = U(I,I)
C          G = -DSIGN(DSQRT(S),F)
C          H = F * G - S
C          U(I,I) = F - G
C          IF (I.EQ. N) GO TO 155
C
C      APPLY LEFT TRANSFOR YATTONS TO REMAINING COLUMNS OF A
C
C      DO 150 J = L, N
C          S = 0.0D0
C
C          DO 140 K = I, M
C              S = S + U(K,I) * U(K,J)
C
C              F = S / H
C
C              DO 150 K = I, M

```

```

      U(K,J) = U(K,J) + F * U(K,I)
150    CONTINUE

C      APPLY LEFT TRANSFORMATIONS TO THE COLUMNS OF B IF IRHS .GT. 0.
C
155      IF (IRHS.EQ.0) GO TO 190
      DO 160 J=1,IRHS
      S=0.0D0
      DO 170 K=I,M
      S = S + U(K,I)*B(K,J)
      F = S/H
      DO 180 K=I,M
      B(K,J) = B(K,J) + F*U(K,I)
160    CONTINUE

C      COMPUTE R I G H T TRANSFORMATIONS.
C
190      DO 200 K = I,M
200      U(K,I) = SCALE * U(K,I)
C
210      W(I) = SCALE * G
      G = 0.0D0
      S = 0.0D0
      SCALE = 0.0D0
      IF (I.GT.M .OR. I.EQ.N) GO TO 290
C
220      DO 220 K = L, N
      SCALE = SCALE + DABS(U(I,K))
C
      IF (SCALE.EQ. 0.0D0) GO TO 290
C
      DO 230 K = L, N
      U(I,K) = U(I,K)/SCALE
      S = S + U(I,K)**2
230    CONTINUE

C      F = U(I,L)
      G = -DSIGN(DSQR(T(S)),F)
      H = F*I*G - S
      U(I,L) = F - G
C
240      DO 240 K = L, N
      RV1(K) = U(I,K) / H
C
      IF (I.EQ.M) GO TO 270
C
      DO 260 J = L,M
      S = 0.0D0
C
      DO 250 K = L,N
      S = S + U(J,K)*U(I,K)
C
      DO 260 K = L,N
      U(J,K) = U(J,K) + S * RV1(K)
260    CONTINUE

C      DO 280 K = L, N
280      U(I,K) = SCALE * U(I,K)
C
290      X = DMAX1(X,DABS(W(I))+DABS(RV1(I)))

```



```

450      CONTINUE
C
460      DO 470 J = I, M
470      U(J,I) = U(J,I) / G
C
GOTO 490
C
475      DC 480 J = I, M
480      U(J,I) = 0.000
C
490      U(I,I) = U(I,I) + 1.000
500  CONTINUE
C      ::::::::::: DIAGONALIZATION OF THE BIDIAGONAL FORM . . . . . . . . .
510  EPS = MACHEP * X
C      ::::::::::: FOR K=N STEP -1 UNTIL 1 DO - - - - - - - - - - - - - - - -
DO 700  KK = 1, N
      K1 = N - KK
      K = K1 + 1
      ITS = 0
C      - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - -
      TEST FOR SPLITTING.
      FOR L=K STEP -1 UNTIL 1 DO - - - - - - - - - - - - - - - - - - - - - -
520      DO 530 LL = 1, K
      L1 = K - LL
      L = L1 + 1
      IF (DABS(RV1(L)) .LE. EPS) GO TO 565
C      ::::::::::: RV1(1) IS ALWAYS ZERO, SO THERE IS NO EXIT
C      THROOUGH THE BOTTOM OF THE LOOP :::::::::::
      IF (DABS(W(L1)) .LE. EPS) GO TO 540
530  CONTINUE
C      ::::::::::: CANCELLATION OF RV1(L) IF L GREATER THAN 1 :::::::::::
540      C = 0.000
      S = 1.000
C
DO 560  I = L, K
      F = S * RV1(I)
      RV1(I) = C * RV1(I)
      IF (DABS(F) .LE. EPS) GO TO 565
      G = W(I)
      Ti = DSQRT(F*F+G*G)
      W(I) = H
      C = G / H
      S = -F / H
C
C      APPLY LEFT TRANSFORMATIONS TO a IF IRHS .GT. 0.
C
      IF (IRHS .EQ. 0) GO TO 542
      DO 545 J=1, IRHS
          Y=B(L1,J)
          Z=B(I,J)
          B(L1,J) = Y*C + Z*S
          B(I,J) = -Y*S + Z*C
545  CONTINUE
54%  CONTINUE
C
      IF (.NOT. MATU) GO TO 560
C
DO 550  J = 1, M
      Y = U(J,L1)
      Z = U(J,I)
      U(J,L1) = Y * C + Z * S

```

```

      U(J,I) = -Y * S + Z * C
  550      CONTINUE
C
  560      CONTINUE
C      :::::::::: TEST FOR CONVERGENCE ::::::::::::
  565      Z = W(K)
      IF (L .EQ. K) GO TO 650
C      :::::::::: SHIFT FROM BOTTOM 2 BY 2 MINOR ::::::::::::
      IF (ITS .EQ. 30) GO TO 1000
      ITS = ITS + 1
      X = W(L)
      Y = W(K1)
      G = RV1(K1)
      H = RV1(K)
      F = ((Y - Z) * (Y + Z) + (G - H) * (G + H)) / (2.000 * H * Y)
      G = DSQRT(F*F+1.0D0)
      F = ((X - Z) * (X + Z) + H * (Y / (F + DSIGN(G,F)) - H)) / X
C      :::::::::: NEXT QR TRANSFORMATION ::::::::::::
      C = 1.0D0
      S = 1.0D0
C
      DO 500 I1 = L, K1
      I = I1 + 1
      G = RV1(I)
      Y = W(I)
      H = S * G
      G = C * G
      Z = DSQRT(F*F+H*H)
      RV1(I1) = Z
      C = F / Z
      S = H / Z
      F = X * C + G * S
      G = -X * S + G * C
      H = Y * S
      Y = Y * C
      IF (.NOT. MATV) GO TO 575
C
      DO 570 J = 1, N
      X = V(J,I1)
      Z = V(J,I)
      V(J,I1) = X * C + Z * S
      V(J,I) = -X * S + Z * C
  570      CONTINUE
C
  575      Z = DSQRT(F*F+H*H)
      W(I1) = Z
C      :::::::::::: ROTATION CAN BE ARBITRARY IF Z IS ZERO ::::::::::::
      IF (Z .EQ. 0.0D0) GO TO 580
      C = F / Z
      S = H / Z
  580      F = C * G + S * Y
      X = -S * G + C * Y
C
C      APPLY LEFT TRANSFORMATIONS TO B IF IRHS .GT. 0.
C
      IF (IRHS .EQ. 0) GO TO 582
      DO 585 J=1,IRHS
          Y = B(I1,J)
          Z = B(I,J)
          B(I1,J) = Y*C + Z*S

```

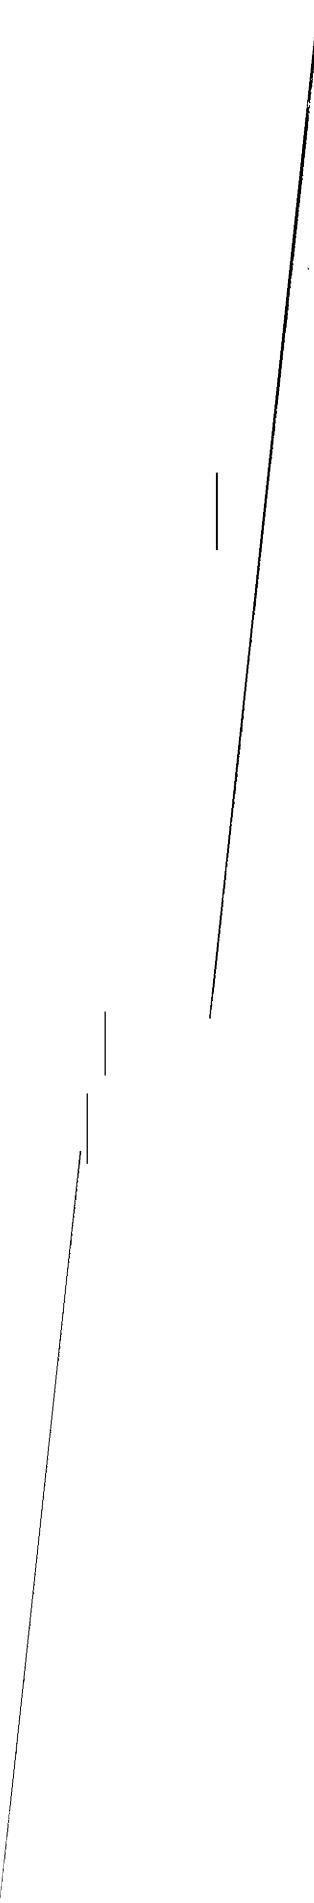
```

585      B(I,J) = -Y*S + Z*C
582      CONTINUE
C
C      IF (.NOT. MATU) GO TO 600
C
C      DO 590 J = 1, M
C          Y = U(J,I1)
C          Z = U(J,I)
C          U(J,I1) = Y * C + Z * S
C          U(J,I) = -Y * S + Z * C
590      CONTINUE
C
C      600  CONTINUE
C
C          RV1(L) = 0.000
C          RV1(K) = F
C          W(K) = X
C          GO TO 520
C      ::::::: CONVERGEYCE :::::::
C      650  IF (Z .GE. 0.000) GO TO 700
C      ::::::: W(K) IS MADE NON-NEGATIVE :::::::
C          W(K) = -Z
C          IF (.NOT. MATV) GO TO 700
C
C          DO 590 J = 1, N
690      V(J,<) = -V(J,K)
C
C      700 CONTINUE
C
C          GO TO 1001
C      ::::::: SET ERROR -- NO CONVERGENCE TO A
C      ::::::: SINGULAR VALUE AFTER 30 ITERATIONS :::::::
1000 IERR = K
100 1 RETURN
C      ::::::: LAST CARD OF GRSVD :::::::*
C      END

```

#### ACKNOWLEDGEMENT

The author would like to thank Prof. John Gregg Lewis for his initial interest in this work and for providing the ALGOL W SVD procedure that was used in the experiments. Special thanks are also due Prof. Gene Golub, Prof. Charles Van Loan and Mr. **W.Coughran** for their helpful discussions. The following persons also helped at one time or another: Dr. C.Lawson, Prof. R-Hanson, Prof. M.Gentleman, Prof. **J.Oliger**, Dr. **P.Gill**, Prof. J-Dennis, Mr. **J.Bolstad** and Dr. **J.S.Pang**. Finally, the author thanks SLAC for providing the computing time that was used.



REFERENCES

[1] **Golub, G.H.** and **Reinsch, C.** (1971) "Singular Value Decomposition and Least Squares Solutions" in *Handbook for Automatic Computation, II, Linear Algebra*, by **J.H. Wilkinson** and **C. Reinsch**, **Springer-Verlag**, New York.

[2] Golub, G.H. and **Kahan, W.** (1965) "Calculating the Singular Values and Pseudoinverse of a Matrix," *SIAM J. Numer. Anal.*, '2, No.3, 205-224.

[3] **Lawson, C.L.** and **Hanson, R.J.** (1974) "Solving Least Squares Problems," Prentice-Hall, New Jersey.

[4] **Golub, G.H.** and **Wilkinson, J.H.** (1975) "Ill-conditioned Eigensystems and the Computation of the Jordan Canonical Form," To appear in **SIAM** Review, 1976.

[5] **Hanson, R.J.** (1971) "A Numerical Method for Solving Fredholm Integral Equations of the First Kind Using Singular Values," *SIAM J. Numer. Anal.*, 8, No.3, 616-626.

[6] **Andrews, H.C.** and **Patterson, C.L.** (1976) "Singular Value Decompositions and Digital Image Processing," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, Vol. ASSP - 24, No.1, Feb. 1976.

[7] **Bartels, R.H., Golub, G.H.,** and **Saunders, M.A.** (1970)  
"Numerical Techniques in Mathematical Programming," in  
Nonlinear Programming. Academic Press, New York, 123-176.

[8] **Gentleman, W.M.** (1972) 'Least Squares Computations by Givens  
Transformations without Square Roots,' Univ. of Waterloo  
Report CSRR-2062, Waterloo, Ontario, Canada.

[9] **Golub, G.H.** and **Businger, P.A.** (1965) "Linear Least Squares Solution  
by Householder Transformations,' Numer. Math., 7, Handbook series  
**Linear Algebra**, 269-276.

[10] **Stewart, G.W.** (1973) 'Introduction to Matrix Computations,'  
Academic Press, New York.

[11] **Parlett, B.N.** and **Wang, Y.** (1975) "The influence of the Compiler  
on the Cost of Mathematical Software,' ACM TOMS, **Vol.1, No.1**,  
March 1975, **pp. 35-46.**

[12] **Smith, B.T.** et al (1976) "Matrix Eigensystem Routines -  
**EISPACK** Guide,' Second Edition, Springer **Verlag**, Lecture  
Notes in Computer Science Series.