

Exploiting Fast Matrix Multiplication Within the Level 3 BLAS

NICHOLAS J. HIGHAM
Cornell University

The Level 3 BLAS (BLAS3) are a set of specifications of FORTRAN 77 subprograms for carrying out matrix multiplications and the solution of triangular systems with multiple right-hand sides. They are intended to provide efficient and portable building blocks for linear algebra algorithms on high-performance computers. We describe algorithms for the BLAS3 operations that are asymptotically faster than the conventional ones. These algorithms are based on Strassen's method for fast matrix multiplication, which is now recognized to be a practically useful technique once matrix dimensions exceed about 100. We pay particular attention to the numerical stability of these "fast BLAS3." Error bounds are given and their significance is explained and illustrated with the aid of numerical experiments. Our conclusion is that the fast BLAS3, although not as strongly stable as conventional implementations, are stable enough to merit careful consideration in many applications.

Categories and Subject Descriptors: G.1.3 [Numerical Analysis]: Numerical Linear Algebra

General Terms: Algorithms

Additional Key Words and Phrases: Error analysis, level 3 BLAS, matrix multiplication, numerical stability, Strassen's algorithm, triangular systems

1. INTRODUCTION

In 1969, Strassen [24] showed how to multiply two $n \times n$ matrices with less than $4.7n^{\log_2 7}$ arithmetic operations. Since $\log_2 7 \approx 2.807 < 3$, his method improves asymptotically on the standard algorithm for matrix multiplication, which requires $O(n^3)$ operations.

Some have regarded Strassen's algorithm as being of theoretical interest only (see, for example, [21, p. 76; 23, p. 533]). However, in 1970 Brent [5] implemented Strassen's algorithm in Algol-W on an IBM 360/67 and concluded that in this environment Strassen's method (with just one level of recursion) runs faster than the conventional method for $n \geq 110$. Furthermore, recently, Bailey [2] compared his FORTRAN implementation of Strassen's algorithm for the Cray-2 with the Cray library routine for matrix multiplication and observed speed-up factors

Author's current address: Department of Mathematics, University of Manchester, Manchester M13 9PL, England.

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.

© 1990 ACM 0098-3500/90/1200-0352 \$01.50

ACM Transactions on Mathematical Software, Vol. 16, No. 4, December 1990, Pages 352-368.

ranging from 1.45 for $n = 128$ to 2.01 for $n = 2048$ (although 35 percent of these speed-ups are due to Cray-specific techniques). These empirical results of Brent and Bailey show that Strassen's algorithm is indeed of practical interest when n is in the hundreds. Indicative of this interest is the inclusion in IBM's ESSL library [18] of routines for real and complex matrix multiplication using a variant of Strassen's method due to Winograd.

The exponent for matrix multiplication has been reduced several times to the current record value of 2.376 [7], but as far as we know none of these asymptotically faster algorithms is quicker than Strassen's method for values of n for which dense matrix multiplication is currently performed in practice ($n \leq 10,000$, say).

In this work we show how Strassen's algorithm can be exploited in all the Level 3 Basic Linear Algebra Subprograms (BLAS3). The BLAS3 [10, 11] are a set of specifications of FORTRAN 77 subprograms for carrying out matrix-matrix operations. They are intended to provide efficient and portable building blocks for linear algebra algorithms on high-performance computers. One reason for the importance of the BLAS3 is that organizing algorithms in a block structure (treating matrices as arrays of smaller matrices) and using calls to the high-level BLAS3 primitives is an effective way to achieve high performance on machines with a hierarchy of memory (such as cache memory, global memory, or vector registers); see, for example, [8, 12, 13, 14, 22].

In Section 2 we present Strassen's algorithm, in its most general form for evaluating products of rectangular matrices, and we discuss practical issues concerning its implementation. In Section 3 we summarize the BLAS3 primitives and describe fast algorithms for the BLAS3 operations involving symmetry and triangularity; these algorithms are recursive and make use of Strassen's method. In addition to having asymptotically smaller operation counts than conventional BLAS3 implementations, the ones we propose have much scope for parallelization, by virtue of their divide and conquer nature.

As explained in [10]: "Although it is intended that the Level 3 BLAS be implemented as efficiently as possible, it is essential that efficiency should not be achieved at the cost of sacrificing numerical stability." We therefore pay particular attention to the stability properties of the algorithms discussed here. Rounding error bounds are given and analyzed in Section 4, and Section 5 contains experiments designed to give further insight into the stability properties of these "fast BLAS3." Our conclusion is that while fast BLAS3 are not as strongly numerically stable as conventional implementations of the BLAS3, they are stable enough to warrant careful consideration in many applications.

2. FAST MULTIPLICATION OF RECTANGULAR MATRICES

Strassen's method is usually presented as a way to multiply square matrices. This is true of the original paper [24] as well as most subsequent descriptions. An exception is the unpublished report of Brent [5], which treats the rectangular case, and which we follow here.

To develop the general version of Strassen's method, consider the product $C = AB$, where A and B are matrices of dimensions $m \times n$ and $n \times p$, respectively.

Assume, for the moment, that $m = 2^i$, $n = 2^j$, $p = 2^k$. Partitioning each of C , A , and B into four equally sized blocks, the product

$$\begin{bmatrix} C_{11} & C_{12} \\ C_{21} & C_{22} \end{bmatrix} = \begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix} \begin{bmatrix} B_{11} & B_{12} \\ B_{21} & B_{22} \end{bmatrix} \quad (2.1)$$

can be accomplished using the following formulas:

$$\begin{aligned} P_1 &= (A_{11} + A_{22})(B_{11} + B_{22}), \\ P_2 &= (A_{21} + A_{22})B_{11}, \\ P_3 &= A_{11}(B_{12} - B_{22}), \\ P_4 &= A_{22}(B_{21} - B_{11}), \\ P_5 &= (A_{11} + A_{12})B_{22}, \\ P_6 &= (A_{21} - A_{11})(B_{11} + B_{12}), \\ P_7 &= (A_{12} - A_{22})(B_{21} + B_{22}), \\ C_{11} &= P_1 + P_4 - P_5 + P_7, \\ C_{12} &= P_3 + P_5, \\ C_{21} &= P_2 + P_4, \\ C_{22} &= P_1 + P_3 - P_2 + P_6. \end{aligned} \quad (2.2)$$

These equations are easily confirmed by substitution. Counting the additions (A) and multiplications (M) we find that while conventional multiplication requires

$$mnpM + m(n-1)pA, \quad (2.3)$$

Strassen's algorithm, using conventional multiplication at the block level, requires

$$\frac{7}{8} mnpM + \left(\frac{7}{8} m(n-2)p + \frac{5}{4} mn + \frac{5}{4} np + \frac{8}{4} mp \right) A. \quad (2.4)$$

Thus, if m , n , and p are large, Strassen's algorithm reduces the arithmetic by a factor of about $\frac{7}{8}$. Since all the blocks in (2.1) have even dimensions we can use the idea recursively on the multiplications associated with the P_i . A total of $q = \min(i, j, k)$ recursions are possible, after which we have matrix multiplications of size $(2^{i-q} \times 2^{j-q}) \times (2^{j-q} \times 2^{k-q})$, in which one of the dimensions is 1 and conventional multiplication must be used. Overall, the number of scalar multiplications is

$$7^q 2^{i+j+k-3q} = \left(\frac{7}{8} \right)^q 2^{i+j+k} = \min(m, n, p)^{\log_2(7/8)} mnp. \quad (2.5)$$

In the case $m = n = p$, this reduces to $n^{\log_2 7} = n^{2.807\dots}$. There does not seem to be any simple formula for the number of additions in the general case, but it is of the same order as the number of multiplications.

There are several ways of modifying the algorithm to handle odd dimensions. One technique is to pad A and B with zeros to achieve even dimensions, compute

the extended product recursively, and then extract the desired product by “unpadding.” Two ways to pad are as follows. One can extend any odd dimension to the next even one by padding with a single row and/or column of zeros; in this case padding may have to be done on each step of the recursion (e.g., if $m = n = p = 2^r + 1$ for some r). Or, as suggested in [24], one can pad once and for all to make each dimension a power of 2 times t , for some small t , and then recur until some dimension is t , at which point conventional multiplication is used. Finally, there is the “chopping” approach suggested in [5]: here one temporarily drops last rows and/or columns to achieve even dimensions, computes the product, and then reinstates the lost information via a rank 1, 2, or 3 correction. For example, if m , n and p are all odd, we can write

$$AB = \begin{bmatrix} A_1 & a_c \\ a_r^T & \alpha \end{bmatrix} \begin{bmatrix} B_1 & b_c \\ b_r^T & \beta \end{bmatrix} = \begin{bmatrix} A_1 B_1 & 0 \\ 0 & 0 \end{bmatrix} + \Delta,$$

where

$$\Delta = \begin{bmatrix} a_c \\ \alpha \end{bmatrix} \begin{bmatrix} b_r^T & \beta \end{bmatrix} + \begin{bmatrix} 0 & A_1 b_c \\ a_r^T B_1 & a_r^T b_c \end{bmatrix}.$$

This last approach is the most attractive for practical computation, because it has the lowest storage requirement and is the easiest to implement in a language such as FORTRAN that does not allow dynamic expansion of arrays. (We mention that in the Matlab language [20], which supports arbitrary redimensioning of arrays, the padding approach is trivial to implement.)

In practice, it is expedient not to recur to the level of scalars or vectors, but to use conventional multiplication once the dimensions are so small that any further reduction in the number of arithmetic operations is offset by an increase in bookkeeping costs. For insight into how the cut-off level should be determined, it is helpful to compare the number of operations in Strassen’s algorithm with one level of recursion with that for conventional multiplication. Assuming the dimensions are even, we have, on subtracting (2.4) from (2.3), $\frac{1}{8}mnpM + (\frac{1}{8}mnp - \frac{5}{4}mn - \frac{5}{4}np - \frac{5}{4}mp)A$. If we assume “ $M \approx A$ ”, this reduces to $\frac{1}{4}(mnp - 5(mn + np + mp))M$, from which we can conclude that Strassen’s method with one level of recursion requires less arithmetic than conventional multiplication if $mnp \geq 5(mn + np + mp)$. This suggests that a criterion of the form “if $mnp \leq n_0(mn + np + mp)/3$ ” is appropriate to terminate the recursions, where n_0 is a machine and compiler-dependent value that must be chosen empirically (we divide by 3 to make the test reduce to “if $n \leq n_0$ ” when $m = n = p$). In his FORTRAN implementation for square matrices on the Cray-2, Bailey [2] found that $n_0 = 127$ minimized the execution time. The Strassen routines in IBM’s ESSL library allow up to four levels of recursion and take $n_0 = 184$ or $n_6 = 260$ for real matrices (depending on the architecture) and $n_0 = 35$ for complex matrices [18].

It is interesting to ask what value of n_0 minimizes the number of arithmetic operations. We can answer this question in the case $m = n = p = 2^k$. Let $n_0 = 2^r$. The number of multiplications and additions can be shown to be

$$M(k) = 7^{k-r}8^r, \quad A(k) = 4^r(2^r + 5)7^{k-r} - 6 \cdot 4^k.$$

The sum $M(k) + A(k)$ is minimized over all integers r by $r = 3$; interestingly, this value is independent of k . Thus, for all $n = 2^k \geq 8$, the choice $n_0 = 8$ minimizes a reasonable measure of the computational work. For $n_0 = 8$, there are about 1.6 times more additions than multiplications when n is large; for $n_0 = 1$, this ratio increases to around 6. Moreover, the total number of arithmetic operations for $n_0 = 1$ is about 1.8 times that for $n_0 = 8$ when n is large; this emphasizes that the choice of n_0 can have a significant impact on the efficiency of Strassen's method. Finally, we note the interesting statistics that the total number of arithmetic operations for Strassen's method with $n_0 = 8$ is smaller than that for conventional multiplication by a factor 0.76 for $n = 128$ and 0.52 for $n = 1024$.

3. FAST ALGORITHMS FOR THE BLAS3

The BLAS3 cover four basic matrix-matrix operations.

(a) Matrix-matrix products:

$$C \leftarrow \alpha AB + \beta C, \quad A \in \mathbf{R}^{m \times n}, \quad B \in \mathbf{R}^{n \times p}, \quad C \in \mathbf{R}^{m \times p}.$$

(b) Rank- r and rank- $2r$ updates of a symmetric matrix $C \in \mathbf{R}^{n \times n}$:

$$\begin{aligned} C &\leftarrow \alpha A^T A + \beta C, & A &\in \mathbf{R}^{r \times n}, \\ C &\leftarrow \alpha A^T B + \alpha B^T A + \beta C, & A, B &\in \mathbf{R}^{r \times n}. \end{aligned}$$

(c) Multiplication of a matrix by a triangular matrix:

$$B \leftarrow \alpha TB, \quad T \in \mathbf{R}^{m \times m} \text{ triangular}, \quad B \in \mathbf{R}^{m \times p}.$$

(d) Solving a triangular system of equations with multiple right-hand sides:

$$B \leftarrow \alpha T^{-1}B, \quad T \in \mathbf{R}^{m \times m} \text{ nonsingular and triangular}, \quad B \in \mathbf{R}^{m \times p}.$$

This is a simplified description. The BLAS3 include variations such as $C \leftarrow \alpha A^T B + \beta C$ in (a), and $B \leftarrow \alpha B T^{-T}$ in (c). For our purposes it is sufficient to give algorithms for (a)–(d), as the variations are handled with obvious modifications.

The general matrix product in (a) can be implemented using Strassen's method as described in Section 2. (Note that it may take fewer operations to compute $(\alpha A)B$ or $A(\alpha B)$ rather than $\alpha(AB)$, depending on the dimensions of A and B .)

In the rest of this section we develop recursive algorithms for operations (b)–(d). For ease of presentation, we assume that all matrix dimensions are a power of 2. For general dimensions, one can use the chopping technique discussed in Section 2, taking advantage of structure when forming the low-rank correction matrix. When stating operation counts, we assume full recursion (i.e., until some dimension is 1). In practice, one would have a cut-off threshold: once the problem size is sufficiently small, conventional multiplication, or in the case of (d), the substitution algorithm for triangular systems would be used.

In (b) the rank- r update requires the computation of $\text{CP}(A) := A^T A$, where $A \in \mathbf{R}^{r \times n}$ (“CP” stands for cross product). Partitioning A into four blocks of dimensions $r/2 \times n/2$, we have

$$A^T A = \begin{bmatrix} A_{11}^T & A_{21}^T \\ A_{12}^T & A_{22}^T \end{bmatrix} \begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix} = \begin{bmatrix} A_{11}^T A_{11} + A_{21}^T A_{21} & A_{11}^T A_{12} + A_{21}^T A_{22} \\ \text{symm.} & A_{12}^T A_{12} + A_{22}^T A_{22} \end{bmatrix}.$$

A recursive algorithm to compute $\text{CP}(A)$ is as follows:

$$\text{CP}(A) = \begin{bmatrix} \text{CP}(A_{11}) + \text{CP}(A_{21}) & \text{Strass}(A_{11}^T, A_{12}) + \text{Strass}(A_{21}^T, A_{22}) \\ \text{symm.} & \text{CP}(A_{12}) + \text{CP}(A_{22}) \end{bmatrix}, \quad (3.1)$$

where $\text{Strass}(A, B)$ denotes the use of Strassen’s method to compute AB . Thus, the idea is to apply the algorithm recursively to the “ $A^T A$ ” subproducts and use Strassen’s method on the others.

The number of multiplications required by this algorithm is

$$rn^2 \left(\frac{2}{3} \mu^{\log_2 7/8} - \frac{1}{6} \mu^{-1} \right) + \frac{1}{2} rn, \quad \mu = \min(r, n), \quad (3.2)$$

which reduces to $\frac{2}{3} n^{\log_2 7} + \frac{1}{3} n^2$ when $n = r$. In comparison with the count (2.5) for Strassen’s method, the dominant term in (3.2) has the same exponent but a smaller constant: $\frac{2}{3}$ instead of 1. Note that this improvement is not as good as for conventional multiplication, where symmetry of the product halves the work. We mention that S. A. Vavasis (private communication) has devised a “ $\frac{3}{5} n^{\log_2 7}$ ” method. His method employs three calls to Strassen’s method and two recursive calls to the method itself. Thus, compared to (3.1), it involves 5 rather than 6 recursive calls on each level, but the formulas defining the method are more complicated, and so it uses more additions and transpositions on each level.

The rank- $2r$ update in (b) can be handled by computing $D = A^T B$ using Strassen’s algorithm and then forming $C \leftarrow \alpha D + \alpha D^T + \beta C$.

Next we turn to the multiplication by a triangular matrix in (c), $B \leftarrow TB$. We can use a technique analogous to the one just discussed for computing $\text{CP}(A)$. Assuming T is upper triangular, we can partition T and B into four equally sized blocks and write

$$A = \begin{bmatrix} T_{11} & T_{12} \\ 0 & T_{22} \end{bmatrix} \begin{bmatrix} B_{11} & B_{12} \\ B_{21} & B_{22} \end{bmatrix} = \begin{bmatrix} T_{11} B_{11} + T_{12} B_{21} & T_{11} B_{12} + T_{12} B_{22} \\ T_{22} B_{21} & T_{22} B_{22} \end{bmatrix}.$$

The idea is to use Strassen’s method on the full matrix products $T_{12} B_{21}$ and $T_{12} B_{22}$, and the same algorithm recursively on the other products. The operation count is essentially the same as for the cross product algorithm above; just replace n by m and r by p in (3.2).

Finally, we consider the BLAS3 operation (d). Our task is to compute $X = T^{-1} B$, where $T \in \mathbf{R}^{m \times m}$ is triangular and $B \in \mathbf{R}^{m \times p}$. We partition X , T , and B

into four equally sized blocks, to obtain

$$\begin{aligned}
 \begin{bmatrix} T_{11} & T_{12} \\ 0 & T_{22} \end{bmatrix}^{-1} \begin{bmatrix} B_{11} & B_{12} \\ B_{21} & B_{22} \end{bmatrix} &= \begin{bmatrix} T_{11}^{-1} & -T_{11}^{-1}T_{12}T_{22}^{-1} \\ 0 & T_{22}^{-1} \end{bmatrix} \begin{bmatrix} B_{11} & B_{12} \\ B_{21} & B_{22} \end{bmatrix} \\
 &= \begin{bmatrix} T_{11}^{-1}(B_{11} - T_{12}(T_{22}^{-1}B_{21})) & T_{11}^{-1}(B_{12} - T_{12}(T_{22}^{-1}B_{22})) \\ T_{22}^{-1}B_{21} & T_{22}^{-1}B_{22} \end{bmatrix} \\
 &\equiv \begin{bmatrix} X_{11} & X_{12} \\ X_{21} & X_{22} \end{bmatrix}, \tag{3.3}
 \end{aligned}$$

where the X_{ij} are computed by recursive applications of the same algorithm, and the products $T_{12}X_{21}$ and $T_{12}X_{22}$ (involved in the computation of X_{11} and X_{12} , respectively) are computed using Strassen's algorithm. In the case $B = I$, this method reduces to the matrix inversion technique described in [24]. The operation count is exactly the same as for the method described for (b), assuming that divisions are counted as multiplications.

4. NUMERICAL STABILITY

Although Strassen's method is well known, its numerical stability, that is, its behavior in floating point arithmetic, is much less widely appreciated. Partly, this is because the early error analysis of the method in [5] was not published (Brent's paper [6] contains some material from [5], but not the error analysis of Strassen's method). Miller [19] states a stability result for Strassen's method in general terms. His result is presented in a more specific form by Bailey in [2], though, unfortunately, an error in the statement makes the result too strong. Bini and Lotti [3] give an error analysis of a class of fast matrix multiplication techniques that includes Strassen's; when specialized to Strassen's method, their quite general error bound is similar to the result given below, but it has an extra factor $\log_2 n$.

It is not difficult to do an error analysis of Strassen's method, at least in the case $m = n = p = 2^k$. We did an analysis before seeing [5], and arrived at almost exactly the same result by the same route. Since [5] is not readily accessible, we present the analysis in an appendix.

The result may be stated as follows. Let $n_0 = 2^r$ be the threshold beyond which conventional multiplication is used. If u denotes the unit round-off and \hat{C} denotes the computed product $C = AB$ from Strassen's method, then

$$\|\hat{C} - C\| \leq \left[\left(\frac{n}{n_0} \right)^{\log_2 12} (n_0^2 + 5n_0) - 5n \right] u \|A\| \|B\| + O(u^2), \tag{4.1}$$

where $\|A\| = \max_{i,j} |a_{ij}|$ (note that this is not a consistent matrix norm since $\|AB\| \leq \|A\| \|B\|$ is generally false). For comparison, if $C = AB$ is computed the usual way, then

$$|\hat{C} - C| \leq nu |A| |B| + O(u^2), \tag{4.2}$$

where $|\cdot|$ denotes the operation of replacing each matrix element by its absolute value, and (4.2) implies the (generally much weaker) bound

$$\|\hat{C} - C\| \leq n^2 u \|A\| \|B\| + O(u^2). \quad (4.3)$$

To interpret the above bounds, note first that all three fall short of the ideal bound

$$|\hat{C} - C| \leq u |C| + O(u^2), \quad (4.4)$$

which says that each component of C is computed with high relative accuracy. Nevertheless (4.2) is a strong bound—the best we can expect when we accept the possibility of numerical cancellation. It treats each element of the error matrix $E = \hat{C} - C$ individually, and is similar to (4.4) if $|A| |B| \approx |C|$.

The norm bounds (4.1) and (4.3) are weaker than (4.2), since they provide the same bound for each element of E . Note also that the scaling $AB \rightarrow (AD)(D^{-1}B)$, where D is diagonal, leaves (4.2) unchanged but alters (4.1) and (4.3).

The bounds (4.3) and (4.1) differ only in the constant term. For Strassen's method, the greater the depth of recursion the bigger the constant in (4.1): if we use just one level of recursion ($n_0 = n/2$), then the constant is $3n^2 + 25n$, whereas with full recursion ($n_0 = 1$), the constant is $6n^{\log_2 12} - 5n = 6n^{3.585\dots} - 5n$.

To summarize, Strassen's method has less favorable stability properties than conventional multiplication in two respects: it satisfies a weaker error bound (norm-wise rather than component-wise) and it has a larger constant in the bound (how much larger depending on n_0). The norm-wise bound is a consequence of the fact that Strassen's method adds together elements of A matrix-wide (and similarly for B); for example, in (2.2), A_{11} is added to A_{22} , A_{12} , and A_{21} . This intermingling of elements is particularly undesirable when A or B has elements of widely differing magnitudes because then large errors can contaminate small components of the product. This phenomenon is well illustrated by the example

$$C = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & \epsilon \\ \epsilon & \epsilon^2 \end{bmatrix},$$

which is evaluated exactly in floating point arithmetic if we use conventional multiplication. However, Strassen's method computes

$$c_{22} = 2(1 + \epsilon^2) + (\epsilon - \epsilon^2) - 1 - (1 + \epsilon).$$

Because c_{22} involves subterms of order unity, the error $\hat{c}_{22} - c_{22}$ will be of order u . Thus the relative error $|\hat{c}_{22} - c_{22}| / |c_{22}| = O(u/\epsilon^2)$, which is much larger than u if ϵ is small. This is an example where Strassen's method does not satisfy the bound (4.2).

Another interesting property of Strassen's method is that it always involves some genuine subtractions (assuming that all additions are of nonzero terms). This is easily deduced from the formulas (2.2). As noted in [14], this makes Strassen's method unattractive in applications where all the elements of A and B are nonnegative (for example, in Markov processes [15]). Here, conventional multiplication yields low relative error component-wise because in (4.2) $|A| |B| = |AB| = |C|$, yet comparable accuracy cannot be guaranteed for Strassen's method.

We mention that Winograd discovered a variant of Strassen's method that requires only 15 additions instead of 18 (see [1, p. 247; 4, p. 133]). However, for this variant the error bound corresponding to (4.1) contains a larger exponent: $\log_2 18 \approx 4.170$ in place of $\log_2 12 \approx 3.585$. This result is proved in [3] (in which both error bounds have an extra factor $\log_2 n$), where it is also shown that Strassen's method has the minimum exponent in its error bound over the set of all fast matrix multiplication methods that are based on computation of a 2×2 matrix product in seven multiplications and that employ integer constants of the form $\pm 2^i$, where i is an integer (this set breaks into 26 equivalence classes).

Finally, we comment on the stability properties of the methods of Section 3. For square matrices of dimension n , the recursive methods for the symmetric matrix update and the triangular matrix times matrix product both satisfy the same bound (4.1) as Strassen's method; this is not surprising, since in both methods nearly all the work is done by calls to Strassen's method. The method for solving triangular systems $TX = B$ satisfies the following bound, a proof of which is given in the appendix (we make the same assumptions as for (4.1)):

$$T\hat{X} = B + E, \quad \|E\| \leq c(n, n_0)u \|T\| \|\hat{X}\| + O(u^2),$$

$$c(n, n_0) = \left(\frac{n}{n_0}\right)^{\log_2 12} \left(\frac{n_0^2}{11} + \frac{23}{55}n_0\right) + \frac{10}{11}n_0^2 + \frac{35}{11}n_0 - \frac{143}{55}n. \quad (4.5)$$

For comparison, consider the computed solution \bar{X} obtained using back substitutions. The i th column \bar{x}_i of \bar{X} satisfies (see, for example, [16]) $(T + E_i)\bar{x}_i = b_i$, $|E_i| \leq (n + 1)u |T|$. It follows that

$$T\bar{X} = B + F, \quad |F| \leq (n + 1)u |T| |\bar{X}|, \quad (4.6)$$

and the latter bound implies $\|F\| \leq n(n + 1)u \|T\| \|\bar{X}\|$. Thus, the same comments apply as for Strassen's method: the error bound for the fast $TX = B$ solver has a weaker form than that for the conventional technique and has a larger constant.

5. NUMERICAL EXPERIMENTS

We have carried out numerical experiments to gain insight into the bounds of Section 4 and to explore the effect of using the fast matrix multiplication techniques in place of conventional multiplication within one particular algorithm.

All computations were performed in Matlab [20] on a Sun 3/50 workstation. The (double precision) unit round-off $u_d = 2^{-52} \approx 2.2 \times 10^{-16}$. Our Matlab codes use recursion and are quite short (under 50 lines of code for each fast BLAS3 routine). Because of the overhead of interpretation and recursion in Matlab, these "fast" routines are in fact quite slow, and so we do not report timings.

In the first experiment we looked at the error in Strassen's algorithm. Let $*$ denote conventional multiplication. For several A and B , we computed $C = A * B$ in double precision, and then $C_* = A * B$ and $C_S = \text{Strass}(A, B, n_0)$ in simulated single precision ($u_s = 2^{-23} \approx 1.2 \times 10^{-7}$). We

define the following quantities:

$$\rho_N(\hat{C}) = \frac{\|\hat{C} - C\|}{n^2 u_s \|A\| \|B\|} \quad (\text{norm-wise relative residual}),$$

$$\rho_C(\hat{C}) = \max_{i,j} \left\{ \frac{|\hat{c}_{ij} - c_{ij}|}{n u_s (|A| |B|)_{i,j}} \right\} \quad (\text{component-wise relative residual}),$$

$$e_N(\hat{C}) = \frac{\|\hat{C} - C\|}{u_s \|C\|} \quad (\text{norm-wise relative error}),$$

$$e_C(\hat{C}) = \max_{i,j} \left\{ \frac{|\hat{c}_{ij} - c_{ij}|}{u_s |c_{ij}|} \right\} \quad (\text{component-wise relative error}).$$

The quantities $\rho_N(C_*) \leq 1$ and $\rho_C(C_*) \leq 1$ (to within $O(u_s)$) measure the sharpness of the bounds (4.3) and (4.2), respectively, while $\rho_N(C_S)$ and $\rho_C(C_S)$ indicate whether the computed Strassen product C_S satisfies these same bounds. The quantity

$$\rho_S(C_S) = \frac{\|C_S - C\|}{d_n u_s \|A\| \|B\|} \leq 1, \quad d_n = \binom{n}{n_0}^{\log_2 12} (n_0^2 + 5n_0) - 5n,$$

measures the sharpness of the residual bound (4.1) for Strassen's method.

We present results for square A and B of dimension $n = 64$ in Table I. We tried both $n_0 = 32$ and $n_0 = 4$, to see how the depth of recursion affects the errors. In practical use of Strassen's method, n and n_0 would be somewhat larger than these values, but they are sufficient to give insight into the error behavior.

The matrices used are defined as follows: urand_i and nrnd_i are random matrices with elements from the uniform $[0, 1]$ and normal $(0, 1)$ distributions, respectively. Z_i is a random matrix with 2-norm condition number 10^4 . P is the Pascal matrix, made up from the numbers in Pascal's triangle; its (i, j) element is $(i + j - 2)! / [(i - 1)!(j - 1)!]$. In each case the same A and B were used for each of the two n_0 values. The C_* statistics are listed in the columns $n_0 = 32$, although they are independent of n_0 .

The results display several interesting features:

- The results confirm the error bounds (4.1–4.3), since $\rho_S(C_S) \leq 1$, $\rho_N(C_*) \leq 1$ and $\rho_C(C_*) \leq 1$ in all cases. The bounds are one or more orders of magnitude from being equalities, and the bound (4.1) for Strassen's method is particularly weak.
- $\rho_N(C_S) \leq 1$ in each case, that is, in these examples Strassen's method satisfies the norm bound (4.3) for conventional multiplication. The component-wise bound (4.2) is severely violated by Strassen's method in the Pascal matrix example! In this example A and B are nonnegative and A has elements of widely varying magnitude—both of these properties are unfavorable for Strassen's method, as noted in Section 4.
- The $e_N(C_S)$ and $e_N(C_*)$ values show that the two multiplication techniques gave products with similar norm-wise relative errors. The component-wise relative errors were also similar for the first three products, being large for the

Table I. Results for Strassen's Method; $m = n = 64$

$(A, B):$ $n_0:$	(urand ₁ , 32)	(urand ₂ , 4)	(nrand ₁ , 32)	(nrand ₂ , 4)	(Z ₁ , 32)	(Z ₂ , 4)	(P, 32)	(urand ₃ , 4)
$\rho_N(C_S)$	3.23e-2	8.17e-2	8.39e-3	2.94e-2	5.08e-3	1.83e-2	1.19e-2	5.14e-2
$\rho_C(C_S)$	2.00e-1	1.12e0	7.85e-2	2.81e-1	9.00e-2	7.08e-1	4.11e22	1.75e34
$e_N(C_S)$	2.51e0	6.33e0	4.06e0	1.42e1	3.87e0	1.39e1	1.00e0	4.35e0
$e_C(C_S)$	1.28e1	7.18e1	8.50e4	4.34e4	1.65e4	1.63e5	2.63e24	1.12e36
$\rho_S(C_S)$	1.46e-2	1.76e-3	7.98e-4	4.80e-5	8.45e-4	5.46e-5	2.53e-4	1.39e-4
$\rho_N(C_*)$	2.55e-2		7.01e-3		4.31e-3		6.89e-3	
$\rho_C(C_*)$	1.18e-1		5.35e-2		5.22e-2		9.37e-2	
$e_N(C_*)$	1.98e0		3.39e0		3.28e0		5.83e-1	
$e_C(C_*)$	7.57e0		3.27e4		7.98e3		6.00e0	

second and third products, in accord with the fact that neither method satisfies the bound (4.4).

- The error measures for Strassen's method are in most cases bigger for $n_0 = 4$ than for $n_0 = 32$, as the bound (4.1) "predicts."

Next, we consider the fast triangular system solver of Section 3. We chose $n = p = 64$ and solved four systems in double precision, first with the fast algorithm, obtaining X_S , and then with back substitutions, obtaining X_B . We computed the relative residuals

$$\rho_N(\hat{X}) = \frac{\|T\hat{X} - B\|_\infty}{(n+1)u_d\|T\|_\infty\|\hat{X}\|_\infty}, \quad \rho_C(X) = \max_{i,j} \left\{ \frac{(|T\hat{X} - B|)_{ij}}{(n+1)u_d(|T||\hat{X}|)_{ij}} \right\},$$

which are bounded by 1 for $\hat{X} = X_B$, in view of (4.6). A quantity $\rho_S(X_S)$ measures the weakness in the bound (4.5) and is defined analogously to $\rho_S(C_S)$ above. In each case B was of the type nrand described above. T_1, T_2, T_3 , and T_4 are the triangular factors from the QR factorizations of random matrices with 2-norm condition numbers 10, 10^5 , 10^{10} , and 10^{15} , respectively. The results are given in Table II.

The key features of the results are:

- In these tests, X_S satisfies the norm-wise residual bound satisfied by X_B , but does not always satisfy the component-wise bound (4.6).
- Generally, the various bounds are far from being equalities, particularly the bound (4.5) for X_S .

We summarize a third experiment, in which we used the fast BLAS3 within the matrix multiplication-rich polar decomposition algorithm of [17]. This algorithm employs the iteration $X_{k+1} = X_k + \frac{1}{2}X_k(I - X_k^T X_k)$, $X_k \in \mathbf{R}^{n \times n}$, which we implemented as $X_{k+1} = X_k + \frac{1}{2}\text{Strass}(X_k, I - \text{CP}(X_k))$. The iteration converges to an orthogonal U , and the symmetric positive semidefinite polar factor is given by $H = U^T A$, which we implemented as $H = \text{Strass}(U^T, A)$. We computed the polar decompositions of various A of dimensions 32 and 64, with $n_0 = 4, 8$, or 16. A natural measure of the quality of the computed polar factors \hat{U}, \hat{H} is their backward error, $\|A - \hat{U}\hat{H}\|_\infty$. In all cases where the iteration converged, the backward error was of the same order of magnitude as when conventional

Table II. Results for the Fast $TX = B$ Solver; $m = p = 64$

T :	T_1		T_2		T_3		T_4	
$\kappa_2(T)$:	10		10^5		10^{10}		10^{15}	
n_0 :	32	4	32	4	32	4	32	4
$\rho_N(X_S)$	1.62e-3	4.39e-3	5.19e-4	2.42e-3	5.21e-4	2.28e-3	5.37e4	2.00e-3
$\rho_C(X_S)$	2.20e-2	6.38e-2	1.99e-2	1.72e0	1.52e-2	2.57e2	1.47e-2	5.80e4
$\rho_S(X_S)$	5.96e-4	5.12e-5	4.67e-4	7.44e-5	3.86e-4	9.56e-5	5.60e-4	7.00e-5
$\rho_N(X_B)$	2.12e-3		9.76e-4		8.73e-4		7.62e-4	
$\rho_C(X_B)$	2.20e-2		1.99e-2		1.52e-2		1.47e-2	

multiplication was used, namely, as small as could be expected. In several cases, with $n_0 = 4$ or 8 , the iteration failed to converge, although it had converged for the same matrices when using conventional multiplication. This behavior can be explained by the less accurate answers returned by the fast BLAS3. Relaxing the convergence tolerance slightly restored convergence (in the same number of iterations as for conventional multiplication) and still gave an acceptable backward error in all cases.

Taking into account the error analysis of Section 3 and the experiments of this section, the numerical stability properties of the fast BLAS3 may be summarized as follows. Here, by “conventional BLAS3” we mean the BLAS3 implemented using conventional multiplication and the substitution algorithm. The conventional BLAS3 satisfy strong component-wise bounds for the residuals of the computed matrices. The fast BLAS3 do not satisfy any such component-wise bounds, as simple examples show, although in specific cases component-wise small residuals may be obtained (e.g., in most of our numerical tests). The two BLAS3 implementations satisfy similar *norm-wise* bounds, but the fast BLAS3 have larger constant terms, which increase as the cut-off threshold n_0 decreases. In our tests with $n \leq 64$, the norm-wise residuals for the fast BLAS3 were never more than twice as large as those for the conventional BLAS3.

For applications in which the BLAS3 are employed as building blocks, an important consideration is whether it is crucial that component-wise small residuals be achieved. LAPACK [9] makes use of the BLAS3 in its block factorization algorithms, and it is desirable to know whether these algorithms remain backwards stable when the fast BLAS3 are used. Specifically, are the computed factors the true factors of a perturbed matrix where the perturbation satisfies a norm-wise bound commensurate with the error bounds for the fast BLAS3? In joint work with J. W. Demmel, we are investigating this question; preliminary results show that the answer is yes for block *LU* factorization, and we are currently examining other block factorizations, including those that use block application of Householder transformations.

The potentially faster growth of errors with dimension for the fast BLAS3 than for the conventional BLAS3 may necessitate some minor algorithm retuning, as in our polar decomposition example. It may also reduce the achievable accuracy, although in practice one is unlikely to have more than a few levels of recursion for reasons of efficiency, so the additional error growth may not be too

serious. For software developers, these issues lead to a possible dilemma: whether to assume that fast BLAS3 may be used instead of conventional BLAS3. Catering for both types of BLAS3 implementation would lead to difficulties in cases where algorithm parameters need to be different for each implementation, as may be the case when convergence tests are involved. In other cases, such as for block factorizations, the developer may merely need to add a section in the documentation identifying how statements about accuracy must be modified for the fast BLAS3.

Overall, we conclude that the fast BLAS3 are stable enough to deserve further investigation in a variety of applications.

Finally, we point out that while we have implemented and tested all the algorithms described here in Matlab on a workstation, we have not implemented the algorithms in a compiled language on a high-performance computer. This task we leave to further work, but Bailey's results [2] for Strassen's method assure us that the fast BLAS3 will yield useful speed-ups for n in the hundreds on appropriate machines.

APPENDIX

In this appendix we give proofs of the error bounds (4.1) for Strassen's method and (4.5) for the fast $TX = B$ solver. As our model for floating point arithmetic, we take

$$\begin{aligned} fl(x \text{ op } y) &= (x \text{ op } y)(1 + \delta), & |\delta| &\leq u, \quad \text{op} = *, /, \\ fl(x \pm y) &= x(1 + \alpha) \pm y(1 + \beta), & |\alpha|, |\beta| &\leq u, \end{aligned}$$

where u is the unit round-off. This model encompasses machines that do not use a guard digit in addition/subtraction.

Throughout the appendix, A , B , T , and X are all $n \times n$ matrices with $n = 2^k$, and the fast algorithms use the threshold $n_0 = 2^r$. Recall that $\|A\| = \max_{i,j} |a_{ij}|$. We use without comment the norm inequality $\|AB\| \leq n \|A\| \|B\|$.

First we consider Strassen's method. Assume that the computed product $\hat{C} \approx AB$ from Strassen's method satisfies

$$\hat{C} = AB + E, \quad \|E\| \leq c_k u \|A\| \|B\| + O(u^2), \quad (\text{A.1})$$

where c_k is a constant. In view of (4.3), (A.1) certainly holds for $n = n_0$, with $c_r = n_0^2$. Our aim is to verify (A.1) inductively and, at the same time, to derive a recurrence for the unknown constant c_k .

Consider C_{11} in (2.2), and in particular, its subterm P_1 . Accounting for the errors in matrix addition, and invoking (A.1), we obtain $\hat{P}_1 = (A_{11} + A_{22} + \Delta_A)(B_{11} + B_{22} + \Delta_B) + E_1$, where

$$\begin{aligned} |\Delta_A| &\leq u(|A_{11}| + |A_{22}|), \\ |\Delta_B| &\leq u(|B_{11}| + |B_{22}|), \\ \|E_1\| &\leq c_{k-1} u \|A_{11} + A_{22} + \Delta_A\| \|B_{11} + B_{22} + \Delta_B\| + O(u^2) \\ &\leq 4c_{k-1} u \|A\| \|B\| + O(u^2). \end{aligned}$$

Hence,

$$\begin{aligned}\hat{P}_1 &= P_1 + F_1, \\ \|F_1\| &\leq (8 \cdot 2^{k-1} + 4c_{k-1})u \|A\| \|B\| + O(u^2).\end{aligned}$$

Similarly,

$$\hat{P}_4 = A_{22}(B_{21} - B_{11} + \Delta_B) + E_4,$$

where

$$\begin{aligned}|\Delta_B| &\leq u(|B_{21}| + |B_{11}|), \\ \|E_4\| &\leq c_{k-1}u \|A_{22}\| \|B_{21} - B_{11} + \Delta_B\| + O(u^2),\end{aligned}$$

which gives

$$\begin{aligned}\hat{P}_4 &= P_4 + F_4, \\ \|F_4\| &\leq (2 \cdot 2^{k-1} + 2c_{k-1})u \|A\| \|B\| + O(u^2).\end{aligned}$$

Now,

$$\hat{C}_{11} = fl(\hat{P}_1 + \hat{P}_4 - \hat{P}_5 + \hat{P}_7)$$

where $\hat{P}_5 \equiv P_5 + F_5$ and $\hat{P}_7 \equiv P_7 + F_7$ satisfy exactly the same error bounds as \hat{P}_4 and \hat{P}_1 , respectively. Assuming that these four matrices are added in the order indicated, we have

$$\begin{aligned}\hat{C}_{11} &= C_{11} + \Delta C_{11} \\ \|\Delta C_{11}\| &\leq u(3\|\hat{P}_1\| + 3\|\hat{P}_4\| + 2\|\hat{P}_5\| + \|\hat{P}_7\|) \\ &\quad + \|F_1 + F_4 - F_5 + F_7\| + O(u^2) \\ &\leq 26 \cdot 2^{k-1}u \|A\| \|B\| + 4(5 \cdot 2^{k-1} + 3c_{k-1})u \|A\| \|B\| + O(u^2) \\ &= (46 \cdot 2^{k-1} + 12c_{k-1})u \|A\| \|B\| + O(u^2).\end{aligned}$$

Clearly, the same bound holds for the other three $\|\Delta C_{ij}\|$ terms. Thus, overall,

$$\hat{C} = AB + E, \quad \|E\| \leq (46 \cdot 2^{k-1} + 12c_{k-1})u \|A\| \|B\| + O(u^2).$$

A comparison with (A.1) shows that we need to define the c_k by

$$c_k = 12c_{k-1} + 46 \cdot 2^{k-1}, \quad k \geq r, \quad c_r = 4^r, \quad (\text{A.2})$$

where the c_r value follows from (4.3). Solving this recurrence, we obtain

$$\begin{aligned}c_k &= 12^{k-r}4^r + 46 \cdot 2^{k-1} \left(\frac{6^{k-r} - 1}{5} \right) \\ &= \left(\frac{n}{n_0} \right)^{\log_2 12} n_0^2 + \frac{46}{5} \frac{n}{2} \left(\left(\frac{n}{n_0} \right)^{\log_2 6} - 1 \right) \\ &\leq \left(\frac{n}{n_0} \right)^{\log_2 12} (n_0^2 + 5n_0) - 5n,\end{aligned}$$

which gives (4.1).

Next, we turn to the fast $TX = B$ solver. Following the lines of the proof above, we assume that the computed solution \hat{X} satisfies

$$T\hat{X} = B + G, \quad \|G\| \leq d_k u \|T\| \|\hat{X}\| + O(u^2). \quad (\text{A.3})$$

From (3.3), we have $T_{22}X_{21} = B_{21}$, so by (A.3):

$$\begin{aligned} T_{22}\hat{X}_{21} &= B_{21} + G_{21}, \\ \|G_{21}\| &\leq d_{k-1}u \|T_{22}\| \|\hat{X}_{21}\| + O(u^2) \leq d_{k-1}u \|T\| \|\hat{X}\| + O(u^2). \end{aligned} \quad (\text{A.4})$$

Also, $T_{11}X_{11} = B_{11} - T_{12}X_{21}$, where the product is computed using Strassen's method. Thus, using (A.1) and (A.3),

$$T_{11}\hat{X}_{11} = B_{11} - T_{12}\hat{X}_{21} + H_{11} + E + \Delta, \quad (\text{A.5})$$

where

$$\begin{aligned} \|E\| &\leq c_{k-1}u \|T_{12}\| \|\hat{X}_{21}\| + O(u^2) && \text{(error in multiplication),} \\ \|\Delta\| &\leq u(\|B_{11}\| + \|T_{12}\hat{X}_{21} + E\|) && \text{(error in subtraction),} \\ \|H_{11}\| &\leq d_{k-1}u \|T_{11}\| \|\hat{X}_{11}\| + O(u^2) && \text{(error in triangular solve).} \end{aligned}$$

We can bound $\|B_{11}\|$ as follows, using (A.5):

$$\|B_{11}\| = \|T_{11}\hat{X}_{11} + T_{12}\hat{X}_{21}\| + O(u) \leq 2^{k-1} \|T\| \|\hat{X}\| + O(u).$$

Using this in the bound for $\|\Delta\|$, we obtain

$$\begin{aligned} T_{11}\hat{X}_{11} &= B_{11} - T_{12}\hat{X}_{21} + G_{11}, \\ \|G_{11}\| &\leq (2^k + c_{k-1} + d_{k-1})u \|T\| \|\hat{X}\| + O(u^2). \end{aligned} \quad (\text{A.6})$$

Results analogous to (A.4) and (A.6) hold for \hat{X}_{22} and \hat{X}_{12} . It is clear, then, that we need to define $d_k = d_{k-1} + c_{k-1} + 2^k$. Combining this with the recursion (A.2) for the c_k , we have $y_k = Wy_{k-1} + v_{k-1}$, $k \geq r$, where

$$y_k = \begin{bmatrix} d_k \\ c_k \end{bmatrix}, \quad W = \begin{bmatrix} 1 & 1 \\ 0 & 12 \end{bmatrix}, \quad v_{k-1} = \begin{bmatrix} 2^k \\ 46 \cdot 2^{k-1} \end{bmatrix}.$$

Expanding the recurrence and picking out the first component of y_k , we obtain

$$d_k = e_1^T y_k = e_1^T \left(W^{k-r} y_r + \sum_{j=0}^{k-r-1} W^j v_{k-1-j} \right). \quad (\text{A.7})$$

Noting that

$$W^j = \begin{bmatrix} 1 & \frac{12^j - 1}{11} \\ 0 & 12^j \end{bmatrix},$$

$d_r = n_0(n_0 + 1)$ (using (4.6)) and $c_r = n_0^2$, we can evaluate (A.7), to obtain

$$\begin{aligned} d_k &= n_0(n_0 + 1) + \left(\frac{12^{k-r} - 1}{11}\right)n_0^2 + (2^{k+1} - 2^{r+1}) \\ &\quad + \frac{46}{11}\left(2^{k-1}\left(\frac{6^{k-r} - 6}{5}\right) - (2^{k-1} - 2^r)\right) \\ &= \left(\frac{n}{n_0}\right)^{\log_2 12} \left(\frac{n_0^2}{11} + \frac{23}{55}n_0\right) + \frac{10}{11}n_0^2 + \frac{35}{11}n_0 - \frac{143}{55}n, \end{aligned}$$

which completes the proof of Eq. (4.5).

ACKNOWLEDGMENTS

I am grateful to Richard Manuck of the Stanford Mathematics and Computer Science Library for providing me with a copy of [5]. Part of this work was done during a visit to the Computer Science Department at Stanford University in January 1989. I thank Gene Golub, Walter Murray, and Michael Saunders for financial support and for making my stay so pleasant. Des Higham corrected an error in (2.4) in a draft manuscript and made many helpful suggestions. Steve Vavasis and the referees also offered helpful comments.

REFERENCES

1. AHO, A. V., HOPCROFT, J. E, AND ULLMAN, J. D. *The Design and Analysis of Computer Algorithms*. Addison-Wesley, Reading, Mass., 1974.
2. BAILEY, D. H. Extra high speed matrix multiplication on the Cray-2. *SIAM J. Sci. Stat. Comput.* 9 (1988), 603–607.
3. BINI, D., AND LOTTI, D. Stability of fast algorithms for matrix multiplication. *Numer. Math.* 36 (1980), 63–72.
4. BRASSARD, G., AND BRATLEY, P. *Algorithmics: Theory and Practice*. Prentice-Hall, Englewood Cliffs, N.J., 1988.
5. BRENT, R. P. Algorithms for matrix multiplication. Tech. Rep. CS 157, Computer Science Dept., Stanford Univ., Palo Alto, Calif., 1970.
6. BRENT, R. P. Error analysis of algorithms for matrix multiplication and triangular decomposition using Winograd's identity. *Numer. Math.* 16 (1970), 145–156.
7. COPPERSMITH, D., AND WINOGRAD, S. Matrix multiplication via arithmetic progression. In *Proceedings of the Nineteenth Annual ACM Symposium on Theory of Computing*, 1987, 1–6.
8. DAYDÉ, M. J., AND DUFF, I. S. Use of level 3 BLAS in LU factorization on the Cray-2, the ETA-10P, and the IBM 3090-200/VF. Tech. Rep. CSS 229, Computer Science and Systems Div., Harwell Lab., 1988.
9. DEMMEL, J. W., DONGARRA, J. J., DU CROZ, J. J., GREENBAUM, A., HAMMARLING, S. J., AND SORENSEN, D. C. Prospectus for the development of a linear algebra library for high-performance computers. Tech. Memor. 97. Mathematics and Computer Science Div., Argonne National Lab., Argonne, Ill., 1987.
10. DONGARRA, J. J., DU CROZ, J. J., DUFF, I. S., AND HAMMARLING, S. J. A set of Level 3 basic linear algebra subprograms. *ACM Trans. Math. Softw.* 16 (1990), 1–17.
11. DONGARRA, J. J., DU CROZ, J. J., DUFF, I. S., AND HAMMARLING, S. J. Algorithm 679: A set of Level 3 basic linear algebra subprograms. *ACM Trans. Math. Softw.* 16 (1990), 18–28.
12. GALLIVAN, K., JALBY, W., AND MEIER, U. The use of BLAS3 in linear algebra on a parallel processor with a hierarchical memory. *SIAM J. Sci. Stat. Comput.* 8 (1987), 1079–1084.
13. GALLIVAN, K., JALBY, W., MEIER, U., AND SAMEH, A. H. Impact of hierarchical memory systems on linear algebra algorithm design. *Int. J. Supercomput. Appl.* 2 (1988), 12–48.

14. GOLUB, G. H., AND VAN LOAN, C. F. *Matrix Computations*, 2nd ed. Johns Hopkins University Press, Baltimore, Md., 1989.
15. HEYMAN, D. P. Further comparisons of direct methods for computing stationary distributions of Markov chains. *SIAM J. Alg. Discrete Meth.* 8 (1987), 226–232.
16. HIGHAM, N. J. The accuracy of solutions to triangular systems. *SIAM J. Numer. Anal.* 26 (1989), 1252–1265.
17. HIGHAM, N. J., AND SCHREIBER, R. S. Fast polar decomposition of an arbitrary matrix. *SIAM J. Sci. Stat. Comput.* 11 (1990), 648–655.
18. IBM. *Engineering and Scientific Subroutine Library, Guide and Reference, Release 3*. 4th ed., Program 5668-863, 1988.
19. MILLER, W. Computational complexity and numerical stability. *SIAM J. Comput.* 4 (1975), 97–107.
20. MOLER, C. B., LITTLE, J. N., AND BANGERT, S. *Pro-Matlab User's Guide*. The MathWorks, Inc., South Natick, Mass., 1987.
21. PRESS, W. H., FLANNERY, B. P., TEUKOLSKY, S. A., AND VETTERLING, W. T. *Numerical Recipes: The Art of Scientific Computing*. Cambridge University Press, Cambridge, England, 1986.
22. SCHREIBER, R. S. Block algorithms for parallel machines. In *Numerical Algorithms for Modern Parallel Computer Architectures*, M. H. Schultz, Ed., IMA Volumes In Mathematics and Its Applications 13, Springer-Verlag, Berlin, 1988, 197–207.
23. SEDGEWICK, R. *Algorithms*. 2nd ed., Addison-Wesley, Reading, Mass., 1988.
24. STRASSEN, V. Gaussian elimination is not optimal. *Numer. Math.* 13 (1969), 354–356.

Received May 1989; revised October 1989; accepted November 1989