

Lawrence Berkeley National Laboratory

LBL Publications

Title

Improving the Numerical Stability of Fast Matrix Multiplication

Permalink

<https://escholarship.org/uc/item/26v4646n>

Journal

SIAM Journal on Matrix Analysis and Applications, 37(4)

ISSN

0895-4798

Authors

Ballard, Grey
Benson, Austin R
Druinsky, Alex
[et al.](#)

Publication Date

2016

DOI

10.1137/15m1032168

Peer reviewed

IMPROVING THE NUMERICAL STABILITY OF FAST MATRIX MULTIPLICATION*

GREY BALLARD[†], AUSTIN R. BENSON[‡], ALEX DRUINSKY[§], BENJAMIN LIPSHITZ[¶],
AND ODED SCHWARTZ^{||}

Abstract. Fast algorithms for matrix multiplication, namely those that perform asymptotically fewer scalar operations than the classical algorithm, have been considered primarily of theoretical interest. Apart from Strassen’s original algorithm, few fast algorithms have been efficiently implemented or used in practical applications. However, there exist many practical alternatives to Strassen’s algorithm with varying performance and numerical properties. Fast algorithms are known to be numerically stable, but because their error bounds are slightly weaker than the classical algorithm, they are not used even in cases where they provide a performance benefit. We argue in this paper that the numerical sacrifice of fast algorithms, particularly for the typical use cases of practical algorithms, is not prohibitive, and we explore ways to improve the accuracy both theoretically and empirically. The numerical accuracy of fast matrix multiplication depends on properties of the algorithm and of the input matrices, and we consider both contributions independently. We generalize and tighten previous error analyses of fast algorithms and compare their properties. We discuss algorithmic techniques for improving the error guarantees from two perspectives: manipulating the algorithms, and reducing input anomalies by various forms of diagonal scaling. Finally, we benchmark performance and demonstrate our improved numerical accuracy.

Key words. practical fast matrix multiplication, error bounds, diagonal scaling

AMS subject classifications. 65F35, 65G50

DOI. 10.1137/15M1032168

*Received by the editors July 23, 2015; accepted for publication (in revised form) by M. L. Overton July 19, 2016; published electronically October 4, 2016. This work was performed by an employee of the U.S. Government or under U.S. Government contract. The U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or allow others to do so, for U.S. Government purposes. Copyright is owned by SIAM to the extent not limited by these rights.

<http://www.siam.org/journals/simax/37-4/M103216.html>

Funding: The first author was supported by an appointment to the Sandia National Laboratories Truman Fellowship in National Security Science and Engineering, sponsored by Sandia Corporation (a wholly owned subsidiary of Lockheed Martin Corporation) as operator of Sandia National Laboratories under its U.S. Department of Energy contract DE-AC04-94AL85000. The second author was supported by an Office of Technology Licensing Stanford Graduate Fellowship. This material is based upon work of the third author supported by the U.S. Department of Energy (DOE), Office of Science, Office of Advanced Scientific Computing Research (ASCR), Applied Mathematics program under contract DE-AC02-05CH11231. The fifth author was supported by grants 1878/14 and 1901/14 from the Israel Science Foundation (founded by the Israel Academy of Sciences and Humanities) and grant 3-10891 from the Ministry of Science and Technology, Israel; the Einstein Foundation and the Minerva Foundation; the Intel Collaborative Research Institute for Computational Intelligence (ICRI-CI); a grant from the United States-Israel Binational Science Foundation (BSF), Jerusalem, Israel; and the HUJI Cyber Security Research Center in conjunction with the Israel National Cyber Bureau in the Prime Minister’s Office.

[†]Sandia National Laboratories, Livermore, CA 94550. Current address: Computer Science Department, Wake Forest University, Winston-Salem, NC 27109 (ballard@wfu.edu).

[‡]Institute for Computational and Mathematical Engineering, Stanford University, Stanford, CA 94305 (arbenon@stanford.edu).

[§]Lawrence Berkeley National Laboratory, Berkeley, CA 94720 (adruinsky@lbl.gov).

[¶]San Bruno, CA 94066 (benjamin.lipshitz@gmail.com).

^{||}School of Computer Science and Engineering, The Hebrew University of Jerusalem, Jerusalem 9190416, Israel (odedsc@cs.huji.ac.il).

1. Introduction. After Strassen’s discovery in 1969 of an algorithm for dense matrix-matrix multiplication [25] that reduced the computational complexity from the classical $O(N^3)$ (for multiplying two $N \times N$ matrices) to $O(N^{\log_2 7})$, extensive effort has been made to understand fast matrix multiplication, based on algorithms with computational complexity exponent less than 3. From a theoretical perspective, there remains a gap between the best known lower bound [21] and best known upper bound [14] on the exponent. From a practical perspective, it is unlikely that the techniques for obtaining the best upper bounds on the exponent can be translated into practical algorithms that will execute faster than the classical one for reasonably sized matrices. In this paper, we are interested in the numerical stability of practical algorithms that have been demonstrated to outperform the classical algorithm (as well as Strassen’s in some instances) on modern hardware [3].

Nearly all fast matrix multiplication algorithms are based on recursion, using a recursive rule that defines a method for multiplying matrices of fixed dimension $M_0 \times K_0$ by $K_0 \times N_0$ $M_0 K_0 N_0$ scalar multiplications. In this work, we use the notation $\langle M_0, K_0, N_0 \rangle$ for these algorithms. For practical algorithms, these fixed dimensions need to be very small, typically $M_0, K_0, N_0 < 10$, as they define the factors by which the dimensions of subproblems are reduced within the recursion. Many such algorithms have been recently discovered [3, 24]. Most fast algorithms share a common bilinear structure and can be compactly represented by three matrices that we denote by $\llbracket \mathbf{U}, \mathbf{V}, \mathbf{W} \rrbracket$, following the notation of Bini and Lotti [4]. Many key properties of the practicality of an algorithm, including its numerical stability, can be derived quickly from its $\llbracket \mathbf{U}, \mathbf{V}, \mathbf{W} \rrbracket$ representation. We also note that, because recursive subproblems are again matrix multiplications, different recursive rules can be combined arbitrarily. Following the terminology of Ballard et al. [2] and Demmel et al. [12], we refer to algorithms that vary recursive rules across different recursive levels and within each level as *nonuniform, nonstationary* algorithms. If an algorithm uses the same rule for every subproblem in each recursive level but varies the rule across levels, we call it a *uniform, nonstationary* algorithm; those defined by only one rule are called *stationary* algorithms.

Fast matrix multiplication is known to yield larger numerical errors than those of the classical algorithm. The forward error guarantee for the classical algorithm is componentwise: the error bound for each entry in the output matrix depends only on the dot product between the corresponding row and column of the input matrices. Fast algorithms perform computations involving other input matrix entries that do not appear in a given dot product (their contributions eventually cancel out), and therefore the error bounds for these algorithms depend on more global properties of the input matrices. Thus, fast algorithms with no modification are known to exhibit so-called normwise stability [4] (sometimes referred to as Brent stability [23]), while the classical algorithm exhibits the stronger componentwise stability, which is unattainable for fast algorithms [23].

Our main goals in this paper are to explore means for improving the theoretical error bounds of fast matrix multiplication algorithms and to test the improvements with numerical experiments, focusing particularly on those algorithms that yield performance benefits in practice. For computing $\mathbf{C} = \mathbf{A} \cdot \mathbf{B}$, where \mathbf{A} is $M \times K$ and \mathbf{B} is $K \times N$, normwise stability bounds for full recursion take the following form:

$$\|\hat{\mathbf{C}} - \mathbf{C}\| \leq f_{\text{alg}}(K) \|\mathbf{A}\| \|\mathbf{B}\| \epsilon + O(\epsilon^2),$$

where $\|\cdot\|$ is the max-norm, ϵ is the machine precision, and f_{alg} is a polynomial

function that depends on the algorithm [4, 12, 16].¹ For example, $f_{\text{alg}}(K) = K^2$ for the classical algorithm, with no assumption on the ordering of dot product computations. We note that f_{alg} is independent of the input matrices, and $\|\mathbf{A}\|\|\mathbf{B}\|$ is independent of the algorithm. In this paper, we explore ways of improving each factor separately. Our main contributions include

1. generalizing and tightening previous error analysis of stationary fast algorithms to bound f_{alg} in terms of the number of recursive steps used and two principal quantities derived from $\llbracket \mathbf{U}, \mathbf{V}, \mathbf{W} \rrbracket$;
2. presenting and comparing the stability quantities of recently discovered practical algorithms;
3. exploring means of improving algorithmic stability through algorithm selection and nonuniform, nonstationary combination of algorithms;
4. presenting diagonal scaling techniques to improve accuracy for inputs with entries of widely varying magnitudes; and
5. showing empirical results of the effects of the various improvement techniques on both error and performance.

The structure of the remainder of the paper is as follows. We describe related work in section 2 and introduce our notation for fast matrix multiplication algorithms in section 3. Section 4 presents the error analysis for bounding f_{alg} for general fast algorithms, and section 5 discusses the implications of the bounds on known practical algorithms. We present diagonal scaling techniques in section 6, showing how to reduce the contribution of the input matrices to the error bound. Finally, we discuss our results in section 7 and provide directions for improving implementations of fast matrix multiplication algorithms.

2. Related work. Bini and Lotti [4] provide the first general error bound for fast matrix multiplication algorithms, and their analysis provides the basis for our results. Demmel et al. [12] generalize Bini and Lotti's results and show that all fast algorithms are stable. A more complete summary of the numerical stability of fast algorithms, with a detailed discussion of Strassen's algorithm along with Winograd's variant, appears in Higham's textbook [16, Chapter 23]. We discuss these previous works in more detail and compare them to our error bounds in section 4.

Castrapel and Gustafson [8] and D'Alberto [9] discuss means of improving the numerical stability of Strassen's algorithm (and Winograd's variant) using the flexibility of nonuniform, nonstationary algorithms. Castrapel and Gustafson propose general approaches to such algorithms, and D'Alberto provides a specific improvement in the case of two or more levels of recursion.

Smirnov [24] describes strategies for discovering practical fast algorithms and presents several new algorithms, including a rank-23 algorithm for $\langle 3, 3, 3 \rangle$ with the fewest known nonzeros and an algorithm for $\langle 6, 3, 3 \rangle$ that yields a better exponent than that of Strassen's. Similar techniques are used by Benson and Ballard [3], who demonstrate performance improvements over the classical and Strassen's algorithms for both single-threaded and shared-memory multithreaded implementations. Laderman, Pan, and Sha [20], and later Kaporin [18, 19], considered another form of practical algorithm that can achieve fewer floating point operations than the Strassen–Winograd variant for certain matrix dimensions. Kaporin demonstrates better numerical stability than that of Strassen–Winograd and shows comparable performance. However, because the base case dimensions proposed are relatively large (e.g., 13 or 20), we

¹Here and elsewhere, the $O(\epsilon^2)$ term hides dependence on dimensions and norms of input matrices.

suspect that the performance will not be competitive on today’s hardware. Furthermore, because the $[[\mathbf{U}, \mathbf{V}, \mathbf{W}]]$ representations are not readily available, we do not consider these types of algorithms in this work.

Dumitrescu [13] proposes a form of diagonal scaling to improve the error bounds for Strassen’s algorithm. We refer to his approach as *outside scaling* and discuss it in more detail in section 6. Higham [16] points out that inside scaling can also affect the error bound, but he does not propose a technique for improving it. Demmel, Dumitriu, and Holtz [11] and Ballard et al. [1] state (without proof) improved error bounds using either inside or outside diagonal scaling.

3. Fast matrix multiplication algorithms. Fast algorithms for matrix multiplication are those that perform fewer arithmetic operations than the classical algorithm in an asymptotic sense, achieving a computational complexity exponent less than 3 for the square case. We consider such fast algorithms to be practical if it has been (or likely can be) demonstrated that they outperform the most efficient implementations of the classical algorithm on current hardware [3]. From a practical viewpoint, because matrices arising in current applications have limited size, we can consider a fast algorithm’s recursive rule being applied only a few times. In light of this viewpoint, we state our algorithms (and error bounds) in terms of the number of recursive levels rather than the dimension of the base case, where the number of recursive levels need not be considered a fixed quantity. In the rest of this section, we state the notation and terminology of the fast algorithms we consider in this paper.

3.1. Base case algorithms. A bilinear noncommutative algorithm that computes a product of an $M_0 \times K_0$ matrix and a $K_0 \times N_0$ matrix ($\mathbf{C} = \mathbf{AB}$) using R nonscalar (active) multiplications is determined by an $M_0K_0 \times R$ matrix \mathbf{U} , a $K_0N_0 \times R$ matrix \mathbf{V} , and an $M_0N_0 \times R$ matrix \mathbf{W} such that

$$(1) \quad c_k = \sum_{r=1}^R w_{kr} m_r \quad \text{where } m_r := s_r \cdot t_r, \quad s_r := \sum_{i=1}^{M_0K_0} u_{ir} a_i, \quad t_r := \sum_{j=1}^{K_0N_0} v_{jr} b_j,$$

for $k = 1, \dots, M_0N_0$. Here, the single indices of entries of \mathbf{A} and \mathbf{B} assume column-major order, the single indices of entries of \mathbf{C} assume row-major order, and (\cdot) signifies an active multiplication. We denote the dimensions of such an algorithm by $\langle M_0, K_0, N_0 \rangle$, the rank of the algorithm by R , and the set of coefficients that determine the algorithm by $[[\mathbf{U}, \mathbf{V}, \mathbf{W}]]$.

3.2. Stationary algorithms. Now we consider multiplying an $M \times K$ matrix \mathbf{A} by a $K \times N$ matrix \mathbf{B} . We will assume that M , K , and N are powers of M_0 , K_0 , and N_0 ; otherwise, we can always pad the matrices with zeros, and the same analysis will hold. The fast algorithm proceeds recursively by first partitioning \mathbf{A} into $M_0 \times K_0$ submatrices of size $(M/M_0) \times (K/K_0)$ and \mathbf{B} into $K_0 \times N_0$ submatrices of size $(K/K_0) \times (N/N_0)$ and then following (1) by matrix blocks, i.e.,

$$(2) \quad \mathbf{C}_k = \sum_{r=1}^R w_{kr} \mathbf{M}_r \quad \text{where } \mathbf{M}_r := \mathbf{S}_r \cdot \mathbf{T}_r, \quad \mathbf{S}_r := \sum_{i=1}^{M_0K_0} u_{ir} \mathbf{A}_i, \quad \mathbf{T}_r := \sum_{j=1}^{K_0N_0} v_{jr} \mathbf{B}_j,$$

for $k = 1, \dots, M_0N_0$, where (\cdot) signifies a recursive call to the algorithm. Here, we are using single subscripts on matrices as an index for the column- or row-major ordering of the matrix blocks. The algorithms in this class of fast matrix multiplication are called *stationary algorithms* because they use a fixed base case algorithm at each

recursive step [12]. However, we do not assume that stationary algorithms recurse all the way to a base case of dimension 1; we assume only that the base case computation (of whatever dimension) is performed using the classical algorithm. Thus, a stationary algorithm is defined by the triplet of matrices $[\mathbf{U}, \mathbf{V}, \mathbf{W}]$ along with a number of recursive levels L used before switching to the classical algorithm.

3.3. Uniform, nonstationary algorithms. In contrast to the stationary algorithms discussed above, *uniform, nonstationary algorithms* employ a different fast algorithm, in the sense of (1) and (2), at each recursive level [2]. The fast algorithm is the same at a given recursive level. Specifically, we will consider uniform, nonstationary algorithms with L steps of recursion, so the algorithm is specified by matrices $\mathbf{U}^{[l]}$, $\mathbf{V}^{[l]}$, $\mathbf{W}^{[l]}$ of dimensions $M_0^{[l]} K_0^{[l]} \times R^{[l]}$, $K_0^{[l]} N_0^{[l]} \times R^{[l]}$, $M_0^{[l]} N_0^{[l]} \times R^{[l]}$ for $l = 1, \dots, L$.

Uniform, nonstationary algorithms are of particular interest for maximizing performance. The fastest algorithm for a particular triplet of dimensions M , K , and N may depend on many factors; the same algorithm may not be optimal for the recursive subproblems of smaller dimensions. Assuming performance is fixed for a given triplet of dimensions, the flexibility of nonstationary algorithms allows for performance optimization over a given set of fast algorithms. However, in parallel and more heterogeneous settings, better performance may be obtained by the greater generality of nonuniform, nonstationary algorithms described in the next section.

3.4. Nonuniform, nonstationary algorithms. The final class of matrix multiplication algorithms we consider contains *nonuniform, nonstationary algorithms*. In contrast to the previous case, nonuniform, nonstationary algorithms use different algorithms within a single recursive level as well across recursive levels [2], though we restrict the dimensions of the partition to be fixed across base case algorithms at a given recursive level. To define such algorithms, we must specify $[\mathbf{U}, \mathbf{V}, \mathbf{W}]$ for every node in the recursion tree, a total of $1 + R^{[1]} + R^{[1]}R^{[2]} + \dots + \prod_{l=1}^{L-1} R^{[l]}$ recursive rules. We use superscript notation $[l, r_1, r_2, \dots, r_{l-1}]$ to denote a recursive node at level l , in the top-level subtree r_1 , in the second level subtree r_2 , and so on.

We demonstrate in subsection 4.5 that the flexibility of these algorithms allows for an improvement in the numerical stability of multilevel recursive algorithms. We suspect that they also provide a performance benefit over stationary algorithms, though this has never been demonstrated empirically.

4. Error analysis. The work of Bini and Lotti [4] provides the basic framework for the forward error analysis of fast matrix multiplication algorithms. They provide general bounds for any square, stationary, bilinear algorithm with power-of-two coefficients (so that there is no error in scalar multiplications), assuming that full recursion is used (a base case of dimension 1). Demmel et al. [12] extend the work of Bini and Lotti by (1) accounting for errors induced by the scalar multiplications in bilinear algorithms, (2) analyzing uniform, nonstationary bilinear fast matrix multiplication algorithms (algorithms that use different fast matrix multiplication routines at different levels of recursion), and (3) analyzing group-theoretic fast matrix multiplication algorithms. The bounds provided by Demmel et al. also assume the use of square algorithms and that full recursion is used. Higham [16] provides bounds for Strassen's original algorithm as well as Winograd's variant in terms of the base case dimension n_0 , where the recursion switches to the classical algorithm. Higham's bounds are also slightly tighter (in the case of Strassen's and Winograd's algorithms) than the general bounds previously mentioned. We note that any matrix multiplication algorithm satisfying the componentwise error bound must perform at least N^3

arithmetic operations; that is, we cannot get the same componentwise error bounds even when using just one step of recursion of a fast algorithm [23].

The goal of the error analysis provided in this section is to generalize the previous work in two main directions and to tighten the analysis particularly in the case when nonzeros of \mathbf{U} , \mathbf{V} , and \mathbf{W} are not all ± 1 . First, we will consider rectangular fast algorithms; that is, instead of considering recursive rules for multiplying two $K \times K$ matrices, we consider the more general set of rules for multiplying an $M \times K$ matrix by a $K \times N$ matrix. Second, we will state our general bounds in terms of the number of levels of recursion used. Motivated by the results of recently discovered practical algorithms [3, 24], we would like to understand the theoretical error guarantees of an algorithm in terms of its $\llbracket \mathbf{U}, \mathbf{V}, \mathbf{W} \rrbracket$ representation. The recent performance results show that rectangular algorithms have practical value (even for multiplying square matrices) and that, for performance reasons, typically only a small number of recursive steps is used in practice. Several recently discovered practical algorithms include fractional power-of-two coefficients (e.g., $1/2$, $1/4$, $1/8$), and we expect that other currently undiscovered, useful algorithms will include fractional coefficients that are not powers of two. Therefore, we make no assumptions on the entries of \mathbf{U} , \mathbf{V} , and \mathbf{W} , and we derive principal quantities that can be tighter than the analogous quantities in the previous works by Bini and Lotti [4] and Demmel et al. [12], particularly in the case of fractional coefficients. This sometimes leads to much sharper error bounds (see Example 4).

Finally, we point out that our representation of nonuniform, nonstationary algorithms is more convenient than previous work. Careful choices of nonuniform, nonstationary algorithms have been shown to improve the numerical stability over stationary approaches (see Example 6) [9]. Bini and Lotti's bounds [4] can be applied to such algorithms in terms of the global $\llbracket \mathbf{U}, \mathbf{V}, \mathbf{W} \rrbracket$ representation, but the size of that representation grows quickly with the number of recursive levels. Our representation (see subsection 3.4) and error bound (see subsection 4.5), given in terms of the base case rule used at each node in the recursion tree, allow for a more efficient search of combinations of rules and have led to automatic discovery of more stable algorithms (see Example 7).

After defining the principal quantities of interest and specifying our model of computation, the rest of this section provides forward error bounds for each of the types of fast algorithms defined in section 3. We warn the reader that there are notational similarities and (sometimes subtle) inconsistencies with previous work as a result of our tightening of the analysis.

4.1. Principal quantities. Following the approach of Bini and Lotti [4], we identify two principal quantities associated with a fast algorithm that, along with the dimensions of the algorithm and the number of levels of recursion, determine its theoretical error bounds. The two principal quantities can be easily computed from the $\llbracket \mathbf{U}, \mathbf{V}, \mathbf{W} \rrbracket$ representation, and we define them in terms of the following vectors:

$$(3) \quad \alpha_r := \sum_{i=1}^{M_0 K_0} \mathbb{I}(u_{ir} \neq 0), \quad \beta_r := \sum_{j=1}^{K_0 N_0} \mathbb{I}(v_{jr} \neq 0), \quad \gamma_k := \sum_{r=1}^R \mathbb{I}(w_{kr} \neq 0),$$

$$(4) \quad a_r := \sum_{i=1}^{M_0 K_0} |u_{ir}|, \quad b_r := \sum_{j=1}^{K_0 N_0} |v_{jr}|$$

for $r = 1, \dots, R$ and $k = 1, \dots, M_0 N_0$, where \mathbb{I} is the Boolean-valued indicator function with value 1 for true and 0 for false. That is, $\boldsymbol{\alpha}$ is the vector of numbers of

TABLE 1

Principal quantities for a variety of fast matrix multiplication algorithms. The rank of the algorithm (R) drives the asymptotic complexity, and the total number of nonzeros in the \mathbf{U} , \mathbf{V} , and \mathbf{W} (nnz) affects the constant in the complexity. Likewise, the E parameter drives the asymptotic behavior of the stability bound, and the Q parameter affects the constant (see Theorem 3). The stability exponent (*stab. exp.*) denotes the asymptotic stability of the algorithm assuming square matrix multiplication (see (12)), which allows for comparison of algorithms with different base case sizes.

$\langle M_0, K_0, N_0 \rangle$	Ref.	$M_0 K_0 N_0$	R	nnz	Q	E	Stab. exp.
$\langle 2, 2, 2 \rangle$	(classical)	8	8	24	4	2	1
$\langle 2, 2, 2 \rangle$	[25]	8	7	36	8	12	3.58
$\langle 3, 2, 2 \rangle$	[25]*	12	11	48	8	12	3.03
$\langle 2, 3, 2 \rangle$	[25]*	12	11	48	9	13	3.03
$\langle 4, 2, 2 \rangle$	[25]*	16	14	72	8	12	2.94
$\langle 2, 4, 2 \rangle$	[25]*	16	14	72	12	24	2.94
$\langle 3, 2, 3 \rangle$	Appendix B	18	15	94	10	20	3.21
$\langle 3, 3, 2 \rangle$	Appendix B	18	15	94	11	23	3.21
$\langle 3, 3, 3 \rangle$	[24]	27	23	139	15	29	3.07
$\langle 4, 2, 3 \rangle$	[3]	24	20	130	14	34	3.38
$\langle 3, 4, 2 \rangle$	[3]	24	20	130	14	30	3.38
$\langle 2, 3, 4 \rangle$	[3]	24	20	130	14	35	3.38
$\langle 4, 4, 2 \rangle$	Appendix C	32	26	257	22	89	3.90
$\langle 4, 2, 4 \rangle$	Appendix C	32	26	257	23	92	3.93
$\langle 3, 4, 3 \rangle$	[3]	36	29	234	23	100	3.66
$\langle 3, 3, 4 \rangle$	[3]	36	29	234	18	71	3.66
$\langle 3, 3, 6 \rangle$	[24]	54	40	960	39	428	4.69
$\langle 3, 6, 3 \rangle$	[24]	54	40	960	48	728.5	4.69

*These algorithms correspond to straightforward generalizations of Strassen's $\langle 2, 2, 2 \rangle$ algorithm, using either two copies of the algorithm or one copy of the algorithm combined with the classical algorithm.

nonzeros in the columns of \mathbf{U} , β is the vector of numbers of nonzeros in the columns of \mathbf{V} , γ is the vector of numbers of nonzeros in the rows of \mathbf{W} , \mathbf{a} is the vector of column 1-norms of \mathbf{U} , and \mathbf{b} is the vector of column 1-norms of \mathbf{V} . When \mathbf{U} and \mathbf{V} have ± 1 entries, $\alpha = \mathbf{a}$ and $\beta = \mathbf{b}$.

DEFINITION 1. The prefactor vector \mathbf{q} is defined entrywise by

$$(5) \quad q_k = \gamma_k + \max_r (\alpha_r + \beta_r) \mathbb{I}(w_{kr} \neq 0)$$

for $k = 1, \dots, M_0 N_0$, and the prefactor Q is defined as

$$Q = \max_k q_k.$$

DEFINITION 2. The stability vector \mathbf{e} is defined entrywise by

$$(6) \quad e_k = \sum_{r=1}^R a_r \cdot b_r \cdot |w_{kr}|$$

for $k = 1, \dots, M_0 N_0$, and the stability factor E is defined as

$$E = \max_k e_k.$$

The principal quantities for several fast algorithms are listed in Table 1.

Bini and Lotti [4] provide a definition of \mathbf{q} for two different summation algorithms: sequential summation and serialized divide-and-conquer (see subsection 4.2). We choose the looser of these two bounds (sequential summation) for generality and simpler notation. However, our results are easily converted to the tighter case. Demmel et al. use the serialized divide-and-conquer algorithm in their analysis. Bini and Lotti's analysis does not account for scalar (nonactive) multiplication by elements of \mathbf{U} , \mathbf{V} , and \mathbf{W} , so their E parameter depends only on the nonzero structure, rather than the magnitude of the elements in these matrices (cf. (4) and Definition 2). Demmel et al. do account for the multiplication by elements of \mathbf{U} , \mathbf{V} , and \mathbf{W} . However, their E parameter is identical to that of Bini and Lotti, and their bound includes an additional factor of $(\|U\|\|V\|\|W\|)^L$, where L is the number of recursive levels and $\|\cdot\|$ is the max-norm.

4.2. Model of arithmetic and notation. We follow the notation of Demmel et al. [12]. Let $\Theta = \{\theta \mid |\theta| < \epsilon\}$ be the set of all errors bounded by ϵ (machine precision), and let $\Delta = \{1 + \theta \mid \theta \in \Theta\}$. We assume the standard model of rounded arithmetic, where the computed value of $op(a, b)$ is $op(a, b)(1 + \theta)$ for some $\theta \in \Theta$. We use the set operation notation: $A + B := \{a + b \mid a \in A, b \in B\}$, $A - B := \{a - b \mid a \in A, b \in B\}$, and $A \cdot B := \{a \cdot b \mid a \in A, b \in B\}$.

We define $A^j = A \cdot A \cdot \dots \cdot A$ and note that $\Delta^j \subset \Delta^{j+1}$ as $1 \in \Delta$. Furthermore, we will not distinguish between singleton sets and an element when using this notation, e.g., $op(a, b)(1 + \theta) \in op(a, b)\Delta$. Finally, we will use the standard hat or $fl(\cdot)$ notation to denote a computed value, e.g., \hat{C} or $fl(op(a, b)) \in op(a, b)\Delta$.

Under this arithmetic, the following fact for summation will be useful in our analysis:

$$(7) \quad fl\left(\sum_{i=1}^N fl(c_i \cdot a_i)\right) \in \left(\sum_{i=1}^N c_i \cdot a_i\right) \Delta^N,$$

where the algorithm for summation is simply to accumulate the terms a_i one at a time in sequential order. By using a serialized divide-and-conquer summation, we can also achieve

$$(8) \quad fl\left(\sum_{i=1}^N fl(c_i \cdot a_i)\right) \in \left(\sum_{i=1}^N c_i \cdot a_i\right) \Delta^{1+\lceil \log_2 N \rceil}.$$

For generality, we will assume the more pessimistic bound in (7). Our results can easily be modified for the error bounds in (8).

We will also use the following property:

$$(9) \quad fl\left(\sum_{i=1}^N c_i \Delta^{a_j}\right) \in \left(\sum_{i=1}^N c_i\right) \Delta^{N+\max_j a_j}.$$

4.3. Forward error analysis of stationary algorithms. The following theorem states the forward error bound for a stationary algorithm in terms of the principal quantities Q and E defined in subsection 4.1, which can be readily determined from its $[\mathbf{U}, \mathbf{V}, \mathbf{W}]$ representation. The sources of error are floating point error accumulation and possible growth in magnitude of intermediate quantities. The floating point error accumulation depends in part on Q and grows at worst linearly in L . The growth of intermediate quantities depends on E and grows exponentially in L , which typically dominates the bound. Table 1 shows the values of these quantities for a variety of algorithms.

THEOREM 3. *Suppose that $\mathbf{C} = \mathbf{A} \cdot \mathbf{B}$, where $\mathbf{A} \in \mathbb{R}^{M \times K}$ and $\mathbf{B} \in \mathbb{R}^{K \times N}$ is computed by using L recursive steps of the fast matrix multiplication in (2), with the classical algorithm used to multiply the $(M/M_0^L) \times (K/K_0^L)$ matrices by the $(K/K_0^L) \times (N/N_0^L)$ matrices at the base cases of the recursion. Then the computed matrix $\hat{\mathbf{C}}$ satisfies*

$$\|\hat{\mathbf{C}} - \mathbf{C}\| \leq (K/K_0^L + Q \cdot L) (K/K_0^L) \cdot E^L \|\mathbf{A}\| \|\mathbf{B}\| \epsilon + O(\epsilon^2),$$

where $\|\cdot\|$ is the max-norm.

Proof. We begin by analyzing how relative errors propagate as we form the \mathbf{S} and \mathbf{T} matrices. Let a superscript index in brackets denote a matrix formed at the specified level of recursion. Following (7), we have the following error at the first recursive level:

$$\hat{\mathbf{S}}_r^{[1]} \in \sum_{i=1}^{M_0 K_0} u_{ir} \mathbf{A}_i \Delta^{\alpha_r}, \quad \hat{\mathbf{T}}_r^{[1]} \in \sum_{j=1}^{K_0 N_0} v_{jr} \mathbf{B}_j \Delta^{\beta_r},$$

where α and β are defined in (3).

This error propagates as we recurse. At the l th level of recursion, the inputs to the fast algorithm are given as sums of matrices \mathbf{A}_ϕ and \mathbf{B}_ψ , each with a possible error of Δ^a and Δ^b , respectively, for some index sets ϕ and ψ and some integers a and b . Following (2) and (7), the algorithm simply accumulates an additional factor of $\Delta^{\alpha_{r_l}}$ and $\Delta^{\beta_{r_l}}$ before the matrices are passed to the subsequent level of recursion. Thus, at the L th level of recursion, we have

$$(10) \quad \hat{\mathbf{S}}_r^{[L]} \in \mathbf{S}_r^{[L]} \Delta^{\alpha_{r_1} + \dots + \alpha_{r_L}}, \quad \hat{\mathbf{T}}_r^{[L]} \in \mathbf{T}_r^{[L]} \Delta^{\beta_{r_1} + \dots + \beta_{r_L}},$$

with $r = r_1 + (r_2 - 1)R + \dots + (r_L - 1)R^{L-1}$. Note that in exact arithmetic,

$$(11) \quad \mathbf{S}_r^{[L]} = \sum_{i=1}^{M_0^L K_0^L} u_{i_1 r_1} \dots u_{i_L r_L} \mathbf{A}_i, \quad \mathbf{T}_r^{[L]} = \sum_{j=1}^{K_0^L N_0^L} v_{j_1 r_1} \dots v_{j_L r_L} \mathbf{B}_j,$$

where $i = i_1 + (i_2 - 1)M_0 K_0 + \dots + (i_L - 1)(M_0 K_0)^{L-1}$ and $j = j_1 + (j_2 - 1)K_0 N_0 + \dots + (j_L - 1)(K_0 N_0)^{L-1}$ represent recursive orderings of the subblocks of \mathbf{A} and \mathbf{B} .

We now use the classical algorithm to multiply the computed $\mathbf{S}^{[L]}$ and $\mathbf{T}^{[L]}$ matrices at the leaves of the recursion. Because the inner dimension of each leaf-level matrix multiplication is K/K_0^L , from (7) and (10) we accumulate another factor of Δ^{K/K_0^L} to obtain

$$\hat{\mathbf{M}}_r^{[L]} \in \mathbf{S}_r^{[L]} \mathbf{T}_r^{[L]} \Delta^{\chi_r + K/K_0^L},$$

where $\chi_r = \alpha_{r_1} + \beta_{r_1} + \dots + \alpha_{r_L} + \beta_{r_L}$ for $1 \leq r \leq R^L$.

Next, the computed matrices $\mathbf{M}^{[L]}$ are added to form \mathbf{C} following (2). At the l th level of recursion, sums of matrices $\mathbf{M}_\phi^{[L]}$, for appropriate index sets ϕ and including accumulated error Δ^a for some integers a , are added together to form the intermediate computed quantities $\mathbf{M}^{[l]}$. In the final step at the top of the recursion tree, we have

$$\hat{\mathbf{C}}_k \in \sum_{r=1}^R w_{kr} \hat{\mathbf{M}}_r^{[1]} \Delta^{\gamma_k},$$

where γ is as defined in (3). Following (9), if $\hat{\mathbf{M}}_r^{[1]} \in \mathbf{M}_r^{[1]} \Delta^{x_r}$ for some integers x_r , then

$$\hat{\mathbf{C}}_k \in \sum_{r=1}^R w_{k_r} \mathbf{M}_r^{[1]} \Delta^{\gamma_k + \max_r x_r \cdot \mathbb{I}(w_{k_r} \neq 0)}.$$

Likewise, a factor of $\Delta^{\gamma_{k_l}}$ is accumulated at every recursive step, and the contributed error from the $\mathbf{M}^{[L]}$ matrices comes from the leaf that is involved in the summation with maximum error. Leaf matrix $\mathbf{M}_r^{[L]}$ is involved in the summation for \mathbf{C}_k if $w_{k_1 r_1} \cdots w_{k_L r_L} \neq 0$, where $r = r_1 + (r_2 - 1)R + \cdots + (r_L - 1)R^{L-1}$ and $k = k_1 + (k_2 - 1)M_0 N_0 + \cdots + (k_L - 1)(M_0 N_0)^{L-1}$. Thus, we have

$$\hat{\mathbf{C}}_k \in \sum_{r=1}^{R^L} w_{k_1 r_1} \cdots w_{k_L r_L} \mathbf{M}_r^{[L]} \Delta^{\mu_k + \max_r \chi_r \cdot \mathbb{I}(w_{k_1 r_1} \cdots w_{k_L r_L} \neq 0) + K/K_0^L},$$

where $\mu_k = \gamma_{k_1} + \cdots + \gamma_{k_L}$.

Let $\delta_k = \mu_k + \max_r \chi_r \cdot \mathbb{I}(w_{k_1 r_1} \cdots w_{k_L r_L} \neq 0) + K/K_0^L$. In order to determine the largest accumulated error, we compute the maximum over all output blocks \mathbf{C}_k :

$$\begin{aligned} \max_k \delta_k &= K/K_0^L + \max_{k_1, \dots, k_L} \left\{ \mu_k + \max_{r_1, \dots, r_L} \chi_r \cdot \mathbb{I}(w_{k_1 r_1} \cdots w_{k_L r_L} \neq 0) \right\} \\ &= K/K_0^L + \max_{k_1} \left\{ \gamma_{k_1} + \max_{r_1} (\alpha_{r_1} + \beta_{r_1}) \mathbb{I}(w_{k_1 r_1} \neq 0) \right\} + \cdots \\ &\quad + \max_{k_L} \left\{ \gamma_{k_L} + \max_{r_L} (\alpha_{r_L} + \beta_{r_L}) \mathbb{I}(w_{k_L r_L} \neq 0) \right\} \\ &= K/K_0^L + \max_{k_1} \left\{ \gamma_{k_1} + \max_{r_1} (\alpha_{r_1} + \beta_{r_1}) \mathbb{I}(w_{k_1 r_1} \neq 0) \right\} \cdot L = K/K_0^L + Q \cdot L, \end{aligned}$$

where Q is given in Definition 1.

We now compute the forward error bound for each block of the output matrix. We have $\mathbf{E}_k = \mathbf{C}_k - \hat{\mathbf{C}}_k \in \sum_r w_{k_1 r_1} \cdots w_{k_L r_L} \mathbf{M}_r^{[L]} \Theta^{\delta_k}$, which implies (using (11))

$$\begin{aligned} |\mathbf{E}_k| &\leq \sum_{r=1}^{R^L} \left| w_{k_1 r_1} \cdots w_{k_L r_L} \mathbf{S}_r^{[L]} \mathbf{T}_r^{[L]} \right| \delta_k \epsilon + O(\epsilon^2) \\ &\leq \sum_{r=1}^{R^L} |w_{k_1 r_1} \cdots w_{k_L r_L}| \sum_{i=1}^{M_0^L K_0^L} |u_{i_1 r_1} \cdots u_{i_L r_L}| \|\mathbf{A}_i\| \sum_{j=1}^{K_0^L N_0^L} |v_{j_1 r_1} \cdots v_{j_L r_L}| \|\mathbf{B}_j\| \delta_k \epsilon \\ &\quad + O(\epsilon^2) \\ &\leq \sum_{r=1}^{R^L} |w_{k_1 r_1} \cdots w_{k_L r_L}| \sum_{i=1}^{M_0^L K_0^L} |u_{i_1 r_1} \cdots u_{i_L r_L}| \sum_{j=1}^{K_0^L N_0^L} |v_{j_1 r_1} \cdots v_{j_L r_L}| \\ &\quad \cdot (K/K_0^L) \|\mathbf{A}\| \|\mathbf{B}\| \delta_k \epsilon + O(\epsilon^2). \end{aligned}$$

Let $\xi_k = \sum_r |w_{k_1 r_1} \cdots w_{k_L r_L}| \sum_i |u_{i_1 r_1} \cdots u_{i_L r_L}| \sum_j |v_{j_1 r_1} \cdots v_{j_L r_L}|$. In order to determine the largest intermediate quantity, we compute the maximum over all output

blocks \mathbf{C}_k :

$$\begin{aligned} \max_k \xi_k &= \max_{k_1, \dots, k_L} \sum_{r_1, \dots, r_L} |w_{k_1 r_1} \cdots w_{k_L r_L}| \sum_{i_1, \dots, i_L} |u_{i_1 r_1} \cdots u_{i_L r_L}| \sum_{j_1, \dots, j_L} |v_{j_1 r_1} \cdots v_{j_L r_L}| \\ &= \left(\max_{k_1} \sum_{r_1} |w_{k_1 r_1}| \sum_{i_1} |u_{i_1 r_1}| \sum_{j_1} |v_{j_1 r_1}| \right) \\ &\quad \cdots \left(\max_{k_L} \sum_{r_L} |w_{k_L r_L}| \sum_{i_L} |u_{i_L r_L}| \sum_{j_L} |v_{j_L r_L}| \right) \\ &= \left(\max_{k_1} \sum_{r_1} |w_{k_1 r_1}| \sum_{i_1} |u_{i_1 r_1}| \sum_{j_1} |v_{j_1 r_1}| \right)^L = E^L, \end{aligned}$$

where E is given in Definition 2.

Computing $\max_k |\mathbf{E}_k|$ by maximizing over δ_k and ξ_k separately, we obtain our result. We note that the two quantities may not achieve their maxima for the same k , but we ignore the possible looseness, as the overall bound will typically be dominated by the value of E . \square

Note that if $L = \log_{K_0} K$ (full recursion), the bound in Theorem 3 becomes

$$\|\hat{\mathbf{C}} - \mathbf{C}\| \leq (1 + Q \cdot L) \cdot E^{\log_{K_0} K} \|\mathbf{A}\| \|\mathbf{B}\| \epsilon + O(\epsilon^2),$$

which is the bound provided by Demmel et al. [12], assuming $M_0 = K_0 = N_0$, $M = K = N$, all nonzeros of \mathbf{U} have the same value, all nonzeros of \mathbf{V} have the same value, and all nonzeros of \mathbf{W} have the same value. If $L = 0$ (no recursion), we get the familiar bound

$$\|\hat{\mathbf{C}} - \mathbf{C}\| \leq K^2 \|\mathbf{A}\| \|\mathbf{B}\| \epsilon + O(\epsilon^2).$$

Example 4. Because our definition of E (Definition 2) accounts for the magnitude of the entries \mathbf{U} , \mathbf{V} , and \mathbf{W} in situ, the bound from Theorem 3 can be tighter than previous analyses [4, 12] when \mathbf{U} , \mathbf{V} , or \mathbf{W} has entries outside of $\{-1, 0, 1\}$. As an example, we consider a $\langle 4, 4, 2 \rangle$ algorithm, where the \mathbf{U} and \mathbf{W} matrices have entries in $\{-0.5, 0.5\}$ [3] (see Appendix C). For this algorithm, E according to Definition 2 is 89, while E according to previous work is 125.

4.4. Forward error analysis of uniform, nonstationary algorithms. Recall that uniform, nonstationary algorithms use a single algorithm at each recursive level. We denote the prefactor vector, stability vector, and partition dimensions of algorithm $[\mathbf{U}^{[l]}, \mathbf{V}^{[l]}, \mathbf{W}^{[l]}]$ at level l by $\mathbf{q}^{[l]}$, $\mathbf{e}^{[l]}$ and $M_0^{[l]}$, $K_0^{[l]}$, and $N_0^{[l]}$. Using analysis similar to that in subsection 4.3, we get the following stability bound for this class of algorithms.

THEOREM 5. *Suppose that $\mathbf{C} = \mathbf{A} \cdot \mathbf{B}$ is computed by a uniform, nonstationary algorithm with L recursive steps of fast matrix multiplication, with the fast algorithm $[\mathbf{U}^{[l]}, \mathbf{V}^{[l]}, \mathbf{W}^{[l]}]$ used at level l and the classical algorithm used to multiply the matrices at the base case of the recursion. Then the computed matrix $\hat{\mathbf{C}}$ satisfies*

$$\|\hat{\mathbf{C}} - \mathbf{C}\| \leq \left(\frac{K}{\prod_{l=1}^L K_0^{[l]}} + \sum_{l=1}^L Q^{[l]} \right) \left(\frac{K}{\prod_{l=1}^L K_0^{[l]}} \right) \cdot \left(\prod_{l=1}^L E^{[l]} \right) \|\mathbf{A}\| \|\mathbf{B}\| \epsilon + O(\epsilon^2).$$

Proof. The proof is similar to the proof of Theorem 3. The largest accumulation error δ now satisfies

$$\begin{aligned} \max_k \delta_k &= \frac{K}{\prod_{l=1}^L K_0^{[l]}} + \max_{k_1} \left\{ \gamma_{k_1}^{[1]} + \max_{r_1} (\alpha_{r_1}^{[1]} + \beta_{r_1}^{[1]}) \mathbb{I}(w_{k_1 r_1}^{[1]} \neq 0) \right\} + \dots \\ &\quad + \max_{k_L} \left\{ \gamma_{k_L}^{[L]} + \max_{r_L} (\alpha_{r_L}^{[L]} + \beta_{r_L}^{[L]}) \mathbb{I}(w_{k_L r_L}^{[L]} \neq 0) \right\} = \frac{K}{\prod_{l=1}^L K_0^{[l]}} + \sum_{l=1}^L Q^{[l]}, \end{aligned}$$

and the largest intermediate growth quantity ξ satisfies

$$\begin{aligned} \max_k \xi_k &\left(\max_{k_1} \sum_{r_1=1}^{R^{[1]}} |w_{k_1 r_1}^{[1]}| \sum_{i_1=1}^{M_0^{[1]} K_0^{[1]}} |u_{i_1 r_1}^{[1]}| \sum_{j_1=1}^{K_0^{[1]} N_0^{[1]}} |v_{j_1 r_1}^{[1]}| \right) \\ &\dots \left(\max_{k_L} \sum_{r_L=1}^{R^{[L]}} |w_{k_L r_L}^{[L]}| \sum_{i_L=1}^{M_0^{[L]} K_0^{[L]}} |u_{i_L r_L}^{[L]}| \sum_{j_L=1}^{K_0^{[L]} N_0^{[L]}} |v_{j_L r_L}^{[L]}| \right) = \prod_{l=1}^L E^{[l]}. \quad \square \end{aligned}$$

4.5. Forward error analysis of nonuniform, nonstationary algorithms.

We now consider nonstationary algorithms, where the algorithm may be nonuniform at every given recursive level of fast matrix multiplication. That is, at any node in the recursion tree, we may choose a different fast algorithm. For simplicity, we assume that at level l in the recursion tree, all algorithms have the same partitioning scheme and rank (so that the $[[\mathbf{U}^{[l, r_1, \dots, r_{l-1}], \mathbf{V}^{[l, r_1, \dots, r_{l-1}], \mathbf{W}^{[l, r_1, \dots, r_{l-1}]]]$ representations have the same dimensions across all values r_1, \dots, r_{l-1}) and that after L levels of recursion, all leaf nodes use the classical algorithm.

In the case of stationary algorithms, one $[[\mathbf{U}, \mathbf{V}, \mathbf{W}]]$ defines the entire algorithm; in the case of uniform nonstationary algorithms, L choices of $[[\mathbf{U}^{[l]}, \mathbf{V}^{[l]}, \mathbf{W}^{[l]}]]$ define the entire algorithm; in this case, we have much more flexibility and can choose $1 + R^{[1]} + R^{[1]}R^{[2]} + \dots + \prod_{l=1}^{L-1} R^{[l]}$ different fast algorithms (the number of internal nodes of the recursion tree). Recall that we use the notation $[l, r_1, r_2, \dots, r_{l-1}]$ as a superscript to refer to the algorithm used at level l in the recursion tree, where r_1 defines subtree membership at level 1, r_2 defines subtree membership at level 2, and so on, and r_{l-1} defines the subtree node at the l th level.

Our analysis of these algorithms is fundamentally the same—we bound the accumulated error (δ) and then bound the number of terms (ξ). However, maximizing over all output blocks is not as straightforward and cannot be simplified as cleanly as in the previous cases. In particular, we define the largest accumulation error δ

recursively as $\max_k \delta_k^{[1]}$, where

$$\begin{aligned} \delta_k^{[1]} &= \frac{K}{\prod_{l=1}^L K_0^{[l]}} + \gamma_{k_1}^{[1]} + \max_{r_1} \delta_k^{[2,r_1]} \cdot \mathbb{I}(w_{k_1 r_1}^{[1]} \neq 0), \\ \delta_k^{[2,r_1]} &= \gamma_{k_2}^{[2,r_1]} + \max_{r_2} \delta_k^{[3,r_1,r_2]} \cdot \mathbb{I}(w_{k_2 r_2}^{[2,r_1]} \neq 0), \\ &\vdots \\ \delta_k^{[l,r_1,\dots,r_{l-1}]} &= \gamma_{k_l}^{[l,r_1,\dots,r_{l-1}]} + \max_{r_l} \delta_k^{[l+1,r_1,\dots,r_l]} \cdot \mathbb{I}(w_{k_l r_l}^{[l,r_1,\dots,r_{l-1}]} \neq 0), \\ &\vdots \\ \delta_k^{[L,r_1,\dots,r_{L-1}]} &= \gamma_{k_L}^{[L,r_1,\dots,r_{L-1}]} + \max_{r_L} \chi_r \cdot \mathbb{I}(w_{k_L r_L}^{[L,r_1,\dots,r_{L-1}]} \neq 0), \text{ and} \\ \chi_r &= \alpha_{r_1}^{[1]} + \beta_{r_1}^{[1]} + \alpha_{r_2}^{[2,r_1]} + \beta_{r_2}^{[2,r_1]} + \dots + \alpha_{r_L}^{[L,r_1,\dots,r_{L-1}]} + \beta_{r_L}^{[L,r_1,\dots,r_{L-1}]} \end{aligned}$$

This expression does not simplify as before. Note that for block k of the output matrix, node (r_1, \dots, r_{l-1}) at level l of the recursion tree accumulates error for the additions/subtractions required by matrix $\mathbf{W}^{[l,r_1,\dots,r_{l-1}]}$ at that node plus the maximum accumulated error from any of the combined terms. The expression for χ_r reflects the number of additions and subtractions required to produce the factor matrices $\mathbf{S}_r^{[L]}$ and $\mathbf{T}_r^{[L]}$ at the leaf nodes, and the error accumulated during the classical matrix multiplications is included in the definition of $\delta_k^{[1]}$.

Likewise, the largest intermediate growth quantity ξ is $\max_k \xi_k$, where

$$\begin{aligned} \xi_k &= \sum_{r_1,\dots,r_L} \left| w_{k_1 r_1}^{[1]} w_{k_2 r_2}^{[2,r_1]} \dots w_{k_L r_L}^{[L,r_1,\dots,r_{L-1}]} \right| \\ &\quad \cdot \sum_{i_1,\dots,i_L} \left| u_{i_1 r_1}^{[1]} u_{i_2 r_2}^{[2,r_1]} \dots u_{i_L r_L}^{[L,r_1,\dots,r_{L-1}]} \right| \cdot \sum_{j_1,\dots,j_L} \left| v_{j_1 r_1}^{[1]} v_{j_2 r_2}^{[2,r_1]} \dots v_{j_L r_L}^{[L,r_1,\dots,r_{L-1}]} \right|, \end{aligned}$$

which we can simplify to

$$\begin{aligned} \xi_k &= \sum_{r_1} \left| w_{k_1 r_1}^{[1]} \right| a_{r_1}^{[1]} b_{r_1}^{[1]} \cdot \sum_{r_2} \left| w_{k_2 r_2}^{[2,r_1]} \right| a_{r_2}^{[2,r_1]} b_{r_2}^{[2,r_1]} \\ &\quad \dots \sum_{r_L} \left| w_{k_L r_L}^{[L,r_1,\dots,r_{L-1}]} \right| a_{r_L}^{[L,r_1,\dots,r_{L-1}]} b_{r_L}^{[L,r_1,\dots,r_{L-1}]}, \end{aligned}$$

where \mathbf{a} and \mathbf{b} vectors are defined as in (4). Note that we cannot simplify further as in the uniform case.

In subsection 5.2, we use nonuniform, nonstationary algorithms to improve the numerical stability of fast matrix multiplication algorithms.

5. Algorithm selection. Theorem 3 immediately provides several options for improving the numerical stability of fast matrix multiplication. First, we can look for algorithms with a smaller Q and E . Since prior work on finding fast algorithms focuses on performance, this provides a new dimension for algorithm design. In subsection 5.1, we compare several stationary algorithms for the same base case as a first step in this dimension of algorithm design. We then extend this analysis to nonuniform, nonstationary algorithms in subsection 5.2. Second, we can reduce the number of recursive levels before using standard matrix multiplication. However, fewer recursive levels means an asymptotically slower algorithm. We examine this trade-off in subsection 5.3. Finally, we can also reduce $\|A\|$ and $\|B\|$ by pre- and postprocessing the data, and we provide several such strategies in section 6.

5.1. Searching for better stationary algorithms. Typically, the only quantity of interest for finding fast matrix multiplication algorithms is the rank of the solution, which controls the asymptotic complexity. However, we can also search for algorithms to minimize the Q and E quantities while maintaining the same rank. This will improve the numerical stability of the algorithm without sacrificing (asymptotic) performance. We will also consider the number of nonzeros (nnz) in the solution, i.e., the sum of the number of nonzero entries in \mathbf{U} , \mathbf{V} , and \mathbf{W} , as this affects the constant in the asymptotic complexity and has noticeable impact on empirical performance [3]. Thus, the parameters of interest for these algorithms is a performance-stability 3-tuple (nnz, Q , E). In general, the number of nonzeros is positively correlated with Q and E , since these quantities directly depend on the nonzero patterns of \mathbf{U} , \mathbf{V} , and \mathbf{W} (see (5) and (6)).

We first examined the base case $\langle 4, 2, 3 \rangle$, which has outperformed Strassen's algorithm in practice [3]. We found 479 algorithms with rank $R = 20$ using numerical low-rank tensor decomposition search techniques [3]. Of these, there were 208 performance-stability tuples. The smallest nnz, Q , and E quantities over all algorithms were 130, 12, and 32, and the corresponding algorithms had performance-stability tuples (130, 14, 34), (138, 12, 34), and (134, 13, 32). No algorithm we found had parameters that achieved more than one of these minima, so we call these three algorithms *semi-optimal*. Consequently, there is a theoretical trade-off between performance and stability. We note that although this list of algorithms is not exhaustive, they are the only publicly available $\langle 4, 2, 3 \rangle$ algorithms.²

We tested the stability of these algorithms by computing the product of samples of random matrices $\mathbf{A} \in \mathbb{R}^{4096 \times 256}$ and $\mathbf{B} \in \mathbb{R}^{256 \times 2187}$. The distributions were $a_{ij}, b_{ij} \sim \text{Uniform}(0, 1)$ and $a_{ij}, b_{ij} \sim \text{Uniform}(-1, 1)$. In addition to the three semi-optimal algorithms described above, we also tested an algorithm with a much worse performance-stability tuple of (156, 26, 132), which we call a *suboptimal* algorithm. For each pair of matrices, we ran the four algorithms with number of recursive levels $L = 1, 2, \dots, 6$. Our goal here is to compare the errors of different algorithms with the same base case and varying number of recursive levels—we are not claiming that any of these algorithms are the best to use for these problem dimensions.

To estimate $\|\hat{\mathbf{C}} - \mathbf{C}\|$, we computed \mathbf{C} using the classical algorithm in quadruple precision arithmetic. All other computations used double precision arithmetic. Overall, we computed the errors for 100 random pairs \mathbf{A} and \mathbf{B} for each distribution. Figure 1 reports the maximum error over the 100 trials for each algorithm and each number of recursive levels as well as the upper bound on the error from Theorem 3. We see the following results:

1. The error bounds are still pessimistic, even with the improved analysis from Theorem 3. Furthermore, the error bounds for the three semi-optimal $\langle 4, 2, 3 \rangle$ algorithms are quite similar.
2. The true error increases with the number of recursive levels, as predicted by Theorem 3 and modeled by the error bound.
3. For both distributions, the suboptimal algorithm has larger errors than the semi-optimal algorithms, as modeled by the error bound.
4. The difference between the semi-optimal algorithms depends on the matrices. For the $\text{Uniform}(0, 1)$ distribution, there is a clear difference in error between the algorithms. Interestingly, the (134, 13, 32) semi-optimal algorithm has

²All of our algorithms, as well as the software for finding them, are publicly available at <https://github.com/arbenson/fast-matmul>.

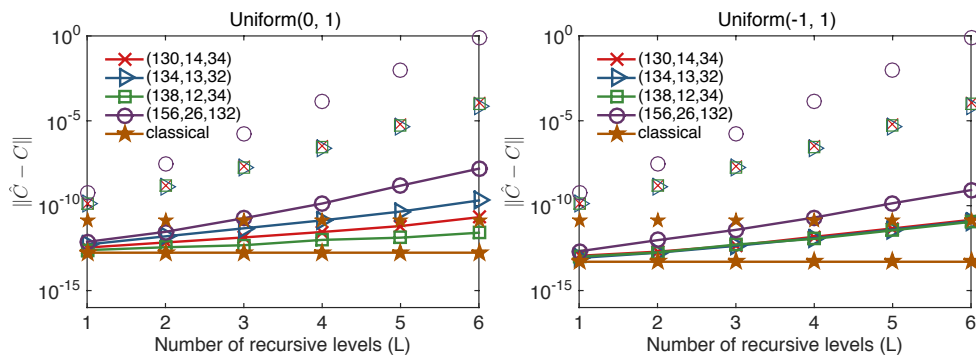


FIG. 1. Error for four $\langle 4, 2, 3 \rangle$ fast matrix multiplication algorithms with different stability parameters and the classical algorithm as a function of the number of recursive levels, L . Three algorithms are semi-optimal in the sense that they minimize one of the following quantities: number of nonzeros, Q , or E . The solid curves are the maximum experimental error over 100 pairs of random matrices, and the corresponding markers are the upper bounds from Theorem 3. The experimental error increases with L , as modeled by Theorem 3. The semi-optimal algorithms with minimal nnz , Q , and E all have similar performance, but the fast algorithm with a worse performance-stability tuple is noticeably less stable in theory and practice.

larger errors than the $(130, 14, 34)$, even though the former algorithm has strictly better Q and E parameters. For the $\text{Uniform}(-1, 1)$ distribution, the errors of the semi-optimal algorithms are practically indistinguishable.

We also considered the $\langle 2, 3, 2 \rangle$ base case, which has optimal rank $R = 11$ [5]. One known algorithm that achieves the optimal rank uses Strassen's algorithm on a 2×2 subblock and classical matrix multiplication on the remaining subblocks. The base case of the algorithm is small enough so that we could use a SAT solver [10] to find over 10,000 rank-11 $\langle 2, 3, 2 \rangle$ algorithms (ignoring symmetries). We found that the combination of Strassen's algorithm with the classical algorithm had a strictly smaller performance-stability triple than all of the other rank-11 solutions. We conclude that this algorithm is likely optimal in both a performance and a stability sense for the class of $\langle 2, 3, 2 \rangle$ algorithms where the scalar multiplications are ± 1 .

5.2. Searching for better nonuniform, nonstationary algorithms. Stationary algorithms benefit from their simplicity, but nonuniform, nonstationary algorithms provide a broader search space for algorithms with better numerical properties. We provide several examples below.

Example 6. D'Alberto [9] describes a nonuniform, nonstationary approach using Strassen's algorithm that obtains a smaller stability factor than that of the original stationary algorithm (for $L \geq 2$). Strassen's algorithm, with $[\mathbf{U}, \mathbf{V}, \mathbf{W}]$ as given in Appendix A, has stability vector $e = [12 \ 4 \ 4 \ 12]$; two levels of recursion with a stationary approach yield a two-level stability vector of $e \otimes e$ with maximum entry $12^2 = 144$. D'Alberto shows that, for $L = 2$, a stability factor of 96 can be obtained with a nonuniform approach using one variant of Strassen's algorithm. One way to achieve this stability factor is to use the alternative algorithm

$$[\tilde{\mathbf{U}}, \tilde{\mathbf{V}}, \tilde{\mathbf{W}}] = \left[\mathbf{U} \left(\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \otimes \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \right) \mathbf{V} \left(\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \otimes \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \right) \mathbf{W} \right]$$

for nodes $[2, 1]$, $[2, 3]$, and $[2, 4]$ of the recursion tree, while using the original algorithm at nodes $[1]$, $[2, 2]$, $[2, 5]$, $[2, 6]$, and $[2, 7]$. Similar improvements can be made based

on the Strassen–Winograd algorithm, which has a slightly larger stability factor.

A more generic nonuniform approach is described in a patent by Castrapel and Gustafson [8]. They consider eight variants of the Strassen–Winograd algorithm, defined by

$$\left[\left(\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}^x \otimes \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}^y \right) \mathbf{U} \left(\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}^z \otimes \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}^x \right) \mathbf{V} \left(\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}^y \otimes \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}^z \right) \mathbf{W} \right],$$

with $x, y, z \in \{1, 2\}$. The correctness of these variants can be derived from the work of Johnson and McLoughlin [17, equation (6)]. Castrapel and Gustafson suggest using random, round-robin, or matrix-dependent selections of algorithms to more evenly distribute the error, but they do not prove that any particular techniques will reduce the stability factor.

Example 7. We can improve the two-level stability factor for the $\langle 3, 2, 3 \rangle$ case in a similar manner. The smallest stability factor we have discovered for this case is $E = 20$, given by the $\llbracket \mathbf{U}, \mathbf{V}, \mathbf{W} \rrbracket$ in Appendix B, which has stability vector

$$e = [20 \quad 20 \quad 2 \quad 12 \quad 4 \quad 20 \quad 4 \quad 12 \quad 20].$$

Compared to a uniform two-level stability factor of $20^2 = 400$, we can achieve a stability factor of 352 using 3 variants of the algorithm. We use the original algorithm at nodes [1], [2, 2], [2, 6], [2, 8], [2, 14], and [2, 15], the variant

$$\left[\left(I_2 \otimes \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix} \right) \mathbf{U} \left(\begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix} \otimes I_2 \right) \mathbf{V} \left(\begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix} \otimes \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix} \right) \mathbf{W} \right]$$

at nodes [2, 1], [2, 3], [2, 10], and [2, 11], and the variant

$$\left[\left(I_2 \otimes \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{bmatrix} \right) \mathbf{U} \left(\begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{bmatrix} \otimes I_2 \right) \mathbf{V} \left(\begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{bmatrix} \otimes \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{bmatrix} \right) \mathbf{W} \right]$$

at nodes [2, 4], [2, 5], [2, 7], [2, 9], [2, 12], and [2, 13]. We suspect that better two-level stability factors are achievable.

5.3. Performance and stability trade-offs with a small number of recursive levels. In addition to searching for better algorithms, we may also consider the effect of the number of recursive levels on the numerical stability. We now consider the performance and stability of fast matrix multiplication algorithms across several base cases and several values of L . Table 1 summarizes the best known (to us) stability factors (E) for several practical base case dimensions. The columns of the table represent the relevant performance and stability parameters for each algorithm, all of which can be computed from the $\llbracket \mathbf{U}, \mathbf{V}, \mathbf{W} \rrbracket$ representation.

The rank R and the number of nonzeros (nnz), along with the number of recursive levels used, determine the number of floating point operations performed by the stationary version of the algorithm. The rank can be compared to the product $M_0 K_0 N_0$, the rank of the classical algorithm for that base case. The quantities Q and E are computed using Definitions 1 and 2, respectively; for a given base case we report the algorithm with the best known E along with that algorithm’s Q . We do not report both $\langle M_0, K_0, N_0 \rangle$ and $\langle N_0, K_0, M_0 \rangle$ because the best algorithms for each




FIG. 2. Distribution of relative instability, $(E/K_0^2)^L$, and percentage of classical flops, $(R/(M_0K_0N_0))^L$, for the algorithms in Table 1 with $L = 1, 2, 3, 4$. A larger value on the y-axis means a less stable algorithm, and a smaller value on the x-axis means a faster algorithm for large problem sizes. There is a general log-linear trade-off between stability and number of floating point operations.

have identical nnz , E , and Q parameters, due to transformations corresponding to transposition of the matrix multiplication.

Although we stress that these algorithms will be used with only a few levels of recursion, we also report the asymptotic stability exponent (stab. exp.) in order to compare algorithms across different base case dimensions. If an algorithm for a square base case $\langle N_0, N_0, N_0 \rangle$ is used on square matrices of dimension N down to subproblems of constant dimension, the bound of Theorem 3 can be simplified to

$$(12) \quad \|\hat{\mathbf{C}} - \mathbf{C}\| \leq c \cdot N^{\log_{N_0} E} \log N \|\mathbf{A}\| \|\mathbf{B}\| \epsilon + O(\epsilon^2),$$

where c is a constant that depends in part on Q . In this case, the stability exponent is $\log_{N_0} E$. We note that the first two rows of Table 1 match the results of Bini and Lotti [4, Table 2]. The most stable rank-23 $\langle 3, 3, 3 \rangle$ algorithm of which we are aware is a cyclic rotation of the one given by Smirnov [24]. In the case of rectangular base cases $\langle M_0, K_0, N_0 \rangle$, we assume a uniform, nonstationary algorithm based on cyclic use of algorithms for $\langle M_0, K_0, N_0 \rangle$, $\langle N_0, M_0, K_0 \rangle$, and $\langle K_0, N_0, M_0 \rangle$, where the three recursive rules are transformations of each other, either by cyclic rotations or transposition (for more details, see Appendices B and C).

Figure 2 shows the distribution of relative instability and percentage of classical flops for the algorithms in Table 1 for $L = 1, 2, 3, 4$. We measure both terms asymptotically. Ignoring the quadratic cost of additions, the percentage of classical flops is given by $(R/(M_0K_0N_0))^L$. For large matrix dimension and L small, we can ignore Q by Theorem 3, and we define the *relative instability* to be $(E/K_0^2)^L$, which is the factor by which the error bound exceeds that of the classical algorithm. In general, most algorithms follow a narrow log-linear trade-off between these two parameters. However, there is still room to select algorithms for a fixed number of recursion levels. For example, with $L = 1$, the $\langle 3, 3, 3 \rangle$ algorithm has roughly the same stability and

does fewer floating point operations than Strassen’s algorithm.

6. Scaling. We now turn our attention to strategies for pre- and postprocessing matrices in order to improve numerical stability. The error bounds from section 4 can be summarized by the following elementwise absolute error bound:

$$(13) \quad |c_{ij} - \hat{c}_{ij}| \leq f_{\text{alg}}(K)\|A\|\|B\|\epsilon + O(\epsilon^2).$$

Recall that f_{alg} is the (at worst) polynomial function of the inner dimension that depends on the particular algorithm used. Unfortunately, these bounds can often be quite large when $|c_{ij}|$ is small relative to $\|A\|\|B\|$. The purpose of this section is to address the contribution of $\|A\|$ and $\|B\|$ to the error bound, ignoring the particular fast algorithm that is used. Thus, for the remainder of this section, we will ignore f_{alg} and consider it a fixed quantity, so that $f_{\text{alg}}(K)\epsilon = O(\epsilon)$, and we will focus on relative error.

The following example shows that the relative error from fast matrix multiplication computations can be large. We note that for the purposes of this example and subsequent examples throughout the paper, we assume that floating point operations incur an ϵ relative error even if the operands happen to be powers of two or if we are subtracting identical values. This assumption does not limit the generality of our examples; instead, we have chosen the matrix entries to make the examples as simple as possible. One could apply small independent relative perturbations on matrix entries for the examples to work in standard floating point arithmetic.

Example 8. Consider the matrices

$$(14) \quad \mathbf{A} = \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix}, \mathbf{B} = \begin{bmatrix} z & 1 \\ z & 1 \end{bmatrix}, \mathbf{C} = \mathbf{A} \cdot \mathbf{B} = \begin{bmatrix} 2z & 2 \\ 2z & 2 \end{bmatrix}$$

for small $z > 0$. By (13), we have the following relative error bound:

$$(15) \quad \frac{|c_{11} - \hat{c}_{11}|}{|c_{11}|} \leq O\left(\frac{\|\mathbf{A}\|\|\mathbf{B}\|\epsilon}{|c_{11}|}\right) = O(\epsilon/z),$$

which can be quite large for small z . Furthermore, this bound is actually achieved with Strassen’s algorithm (see Appendix A for the definition of Strassen’s algorithm). Specifically, Strassen’s algorithm computes

$$\begin{aligned} m_1 &= (a_{11} + a_{22})(b_{11} + b_{22}) = (1 + 1) \cdot (z + 1), \\ m_4 &= a_{22}(b_{21} - b_{11}) = 1 \cdot (z - z), \\ m_5 &= (a_{11} + a_{12})b_{22} = (1 + 1) \cdot 1, \\ m_7 &= (a_{12} - a_{22})(b_{21} + b_{22}) = (1 - 1) \cdot (z + 1), \\ c_{11} &= m_1 + m_4 - m_5 + m_7. \end{aligned}$$

There are terms of size $O(1)$ in computing $m_1, m_4, m_5,$ and m_7 , so the absolute error $|c_{11} - \hat{c}_{11}|$ is $O(\epsilon)$. Since $c_{11} = z$, the relative error is $O(\epsilon/z)$.

We now demonstrate several methods for improving numerical stability issues by preprocessing \mathbf{A} and \mathbf{B} and postprocessing \mathbf{C} . The idea underlying these methods is the following straightforward observation:

$$(16) \quad \mathbf{C} = \mathbf{D}_A \mathbf{D}_A^{-1} \mathbf{A} \mathbf{D}_B \mathbf{D}_B^{-1} \mathbf{B} \mathbf{D}_B^{-1} \mathbf{D}_B$$

for any nonsingular scaling matrices \mathbf{D}_A , \mathbf{D}_B , and \mathbf{D} . By taking advantage of the associativity of matrix multiplication (in exact arithmetic) and scaling matrices \mathbf{D}_A , \mathbf{D}_B , and \mathbf{D} that are easy to apply, we can improve the normwise bound in Theorem 3 without significantly affecting the performance of the algorithm.

For the algorithms and analysis in this section, we will consider diagonal scaling matrices with positive diagonal entries. In order to simplify the analysis, we will assume that there is no numerical error in applying or computing the scaling matrices. This could be achieved, for example, by rounding the scaling matrix entries to the nearest power of two. Regardless, the error introduced by the fast matrix multiplication algorithm has the larger impact on the stability, and the scaling matrices can curb numerical inaccuracies.

6.1. Outside scaling. In light of (16), Dumitrescu proposed the following outside scaling matrices [13]:

$$\mathbf{D}_A = \text{diag}\left(\max_j |a_{ij}|\right), \quad \mathbf{D}_B = \text{diag}\left(\max_i |b_{ij}|\right).$$

The resulting procedure is Algorithm 1.

Algorithm 1. Outside scaling for fast matrix multiplication.

Require: matrices \mathbf{A} and \mathbf{B}

Ensure: $\mathbf{C} = \mathbf{A} \cdot \mathbf{B}$

- 1: $\mathbf{D}_A \leftarrow \text{diag}(\max_j |a_{ij}|)$
 - 2: $\mathbf{A}' \leftarrow \mathbf{D}_A^{-1} \mathbf{A}$
 - 3: $\mathbf{D}_B \leftarrow \text{diag}(\max_i |b_{ij}|)$
 - 4: $\mathbf{B}' \leftarrow \mathbf{B} \mathbf{D}_B^{-1}$
 - 5: $\mathbf{C}' \leftarrow \mathbf{A}' \cdot \mathbf{B}'$ with fast matrix multiplication.
 - 6: $\mathbf{C} \leftarrow \mathbf{D}_A \mathbf{C}' \mathbf{D}_B$
-

Clearly, the algorithm correctly computes $\mathbf{C} = \mathbf{A} \cdot \mathbf{B}$ in exact arithmetic, provided there are no all-zero rows in \mathbf{A} or all-zero columns in \mathbf{B} . Importantly, the normwise bound in Theorem 3 applies to the scaled matrices \mathbf{A}' and \mathbf{B}' . In particular, we get the following improved bound [13].

PROPOSITION 9. *Using Algorithm 1,*

$$|c_{ij} - \hat{c}_{ij}| \leq O(\epsilon) \|a_{i,:}\| \|b_{:,j}\|.$$

Proof. Outside scaling ensures that $\|\mathbf{A}'\| = \|\mathbf{B}'\| = 1$, so by (13), $\|\mathbf{C}' - \hat{\mathbf{C}}'\| \leq O(\epsilon)$. Since $\mathbf{C}' - \hat{\mathbf{C}}' = \mathbf{D}_A(\hat{\mathbf{C}} - \mathbf{C})\mathbf{D}_B$, the result follows from the fact that the i th diagonal entry of \mathbf{D}_A is $\|a_{i,:}\|$ and j th diagonal entry of \mathbf{D}_B is $\|b_{:,j}\|$. \square

For the matrices in Example 8, the bound from Proposition 9 improves upon (15):

$$\frac{|c_{11} - \hat{c}_{11}|}{|c_{11}|} \leq O\left(\frac{\|a_{1,:}\| \|b_{:,1}\| \epsilon}{|c_{11}|}\right) = O(\epsilon).$$

This indeed improves the numerical stability of Strassen's algorithm. For the matrices in (14), the outside scaling is

$$\mathbf{A}' \leftarrow \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix}, \quad \mathbf{B}' \leftarrow \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix}, \quad \mathbf{C}' = \mathbf{A}' \cdot \mathbf{B}' = \begin{bmatrix} 2 & 2 \\ 2 & 2 \end{bmatrix},$$

and when computing \mathbf{C}' with Strassen's algorithm,

$$\begin{aligned} m_1 &= (a'_{11} + a'_{22})(b'_{11} + b'_{22}) = (1 + 1) \cdot (1 + 1), \\ m_4 &= a'_{22}(b'_{21} - b'_{11}) = 1 \cdot (1 - 1), \\ m_5 &= (a'_{11} + a'_{12})b'_{22} = (1 + 1) \cdot 1, \\ m_7 &= (a'_{12} - a'_{22})(b'_{21} + b'_{22}) = (1 - 1) \cdot (1 + 1), \\ c'_{11} &= m_1 + m_4 - m_5 + m_7. \end{aligned}$$

Now, all subterms are on the order of unity, so the relative error in computing c'_{11} is $O(\epsilon)$.

6.2. Inside scaling. There are pairs of matrices in which outside scaling is not sufficient for numerical stability.

Example 10. Consider the matrices

$$(17) \quad \mathbf{A} = \begin{bmatrix} 1 & z \\ 1 & z \end{bmatrix}, \mathbf{B} = \begin{bmatrix} z & z \\ 1 & 1 \end{bmatrix}, \mathbf{C} = \mathbf{A} \cdot \mathbf{B} = \begin{bmatrix} 2z & 2z \\ 2z & 2z \end{bmatrix}$$

for small $z > 0$. Using outside scaling on these matrices does nothing since $\mathbf{D}_A = \mathbf{D}_B = \mathbf{I}$. However, using a fast algorithm can still have severe numerical stability issues. Computing c_{12} with Strassen's algorithm uses the following computations:

$$\begin{aligned} m_3 &= a_{11}(b_{12} - b_{22}) = 1 \cdot (z - 1), \\ m_5 &= (a_{11} + a_{12})b_{22} = (1 + z) \cdot 1, \\ c_{12} &= m_3 + m_5. \end{aligned}$$

The computation of m_3 and m_5 has terms of unit size, so $|c_{12} - \hat{c}_{12}|$ is $O(\epsilon)$ and the relative error is $O(\epsilon/z)$. This is reflected in the bound from (13),

$$\|\mathbf{A}\|\|\mathbf{B}\|/|c_{12}| = 1/(2z).$$

We now propose a technique called *inside scaling* based on the following matrix:

$$(18) \quad \mathbf{D} = \text{diag} \left(\sqrt{\frac{\max_j |b_{kj}|}{\max_i |a_{ik}|}} \right).$$

The resulting procedure is in Algorithm 2. The idea is to scale the columns of \mathbf{A} and the corresponding rows of \mathbf{B} to have the same norm. In general, we get an improved error bound, as detailed in Proposition 11. We note that there exist several references to inside scaling [1, 11, 16, 22], though to the best of our knowledge this is the first explicit statement of the diagonal values in (18). A crude inside scaling method was proposed earlier by Brent [7], where the inside scaling matrix is $\mathbf{D} = \sqrt{\|\mathbf{B}\|/\|\mathbf{A}\|\mathbf{I}}$.

PROPOSITION 11. *Using Algorithm 2,*

$$\|\hat{\mathbf{C}} - \mathbf{C}\| \leq O(\epsilon) \max_{i,k,j} |a_{ik}| |b_{kj}|.$$

Proof. By (13),

$$\|\hat{\mathbf{C}} - \mathbf{C}\| \leq O(\epsilon) \|\mathbf{A}\mathbf{D}\|\|\mathbf{D}^{-1}\mathbf{B}\| = O(\epsilon) \left(\max_k \|a_{:,k}\| d_{kk} \right) \left(\max_l d_{ll}^{-1} \|b_{l,:}\| \right).$$

Algorithm 2. Inside scaling for fast matrix multiplication.

Require: matrices \mathbf{A} and \mathbf{B}

Ensure: $\mathbf{C} = \mathbf{A} \cdot \mathbf{B}$

1: $\mathbf{D} \leftarrow \text{diag}\left(\sqrt{\frac{\max_j |b_{kj}|}{\max_i |a_{ik}|}}\right)$

2: $\mathbf{A}' \leftarrow \mathbf{A}\mathbf{D}$

3: $\mathbf{B}' \leftarrow \mathbf{D}^{-1}\mathbf{B}$

4: $\mathbf{C} \leftarrow \mathbf{A}' \cdot \mathbf{B}'$ with fast matrix multiplication.

By the definition of \mathbf{D} ,

$$\|a_{:,k}\|d_{kk} = d_{kk}^{-1}\|b_{k,:}\| = \sqrt{\|a_{:,k}\|\|b_{k,:}\|},$$

so the two maxima are attained at the same index. The result then follows from the fact that $\|a_{:,k}\|\|b_{k,:}\| = \max_{i,k,j} |a_{ik}||b_{kj}|$. \square

For \mathbf{A} and \mathbf{B} in Example 10, $\max_{i,k,j} |a_{ik}||b_{kj}| = z$, and we get an $O(\epsilon)$ relative error bound for computing each entry in \mathbf{C} . The inside scaling updates to the matrices in (17) are

$$\mathbf{D} = \begin{bmatrix} \sqrt{z} & 0 \\ 0 & 1/\sqrt{z} \end{bmatrix}, \mathbf{A}' \leftarrow \begin{bmatrix} \sqrt{z} & \sqrt{z} \\ \sqrt{z} & \sqrt{z} \end{bmatrix}, \mathbf{B}' \leftarrow \begin{bmatrix} \sqrt{z} & \sqrt{z} \\ \sqrt{z} & \sqrt{z} \end{bmatrix}.$$

Strassen's algorithm now computes

$$m_3 = a'_{11}(b'_{12} - b'_{22}) = \sqrt{z} \cdot (\sqrt{z} - \sqrt{z}),$$

$$m_5 = (a'_{11} + a'_{12})b'_{22} = (\sqrt{z} + \sqrt{z}) \cdot \sqrt{z},$$

$$c'_{12} = m_3 + m_5.$$

This time, the computation of m_3 and m_5 involves terms on the order of z instead of on the order of unity, and we get an $O(\epsilon)$ relative error in the computation.

6.3. Repeated outside-inside scaling. Next, we consider repeatedly applying outside and inside scaling in alternating order, as shown in Algorithm 3. This process can only improve the error bounds, and it eventually converges. Outside and inside scaling can simply be applied several times, or the user can specify a cheaply computed stopping criterion that will guarantee a relative distance from the limit point.

We start with our accuracy analysis. In our analysis, we use $\mathbf{A}^{(t)}$ and $\mathbf{B}^{(t)}$ to denote the values of \mathbf{A}' and \mathbf{B}' , respectively, after t steps of Algorithm 3. We also use $r_i^{(t)}$ and $s_j^{(t)}$ to denote the diagonal elements of \mathbf{D}_A and \mathbf{D}_B , respectively, after t steps. The initial values of these variables correspond to $t = 0$ in our notation.

PROPOSITION 12. *Let t be the number of steps of Algorithm 3 that we complete. The computed product satisfies*

$$|c_{ij} - \hat{c}_{ij}| \leq O(\epsilon) r_i^{(t)} s_j^{(t)} \|\mathbf{A}^{(t)}\| \|\mathbf{B}^{(t)}\|.$$

Proof. If the last step is an O step, then following the proof of Proposition 9, $\|\mathbf{A}^{(t)}\| = \|\mathbf{B}^{(t)}\| = 1$ and $|c'_{ij} - \hat{c}'_{ij}| \leq O(\epsilon)$. If the last step is an I step, then by Proposition 11,

$$|c'_{ij} - \hat{c}'_{ij}| \leq O(\epsilon) \max_{i,k,j} |a_{ik}^{(t)}| |b_{kj}^{(t)}| \leq O(\epsilon) \|\mathbf{A}^{(t)}\| \|\mathbf{B}^{(t)}\|.$$

The result then follows from the fact that $\mathbf{C} - \hat{\mathbf{C}} = \mathbf{D}_A(\mathbf{C}' - \hat{\mathbf{C}}')\mathbf{D}_B$. \square

Algorithm 3. Repeated outside-inside scaling for fast matrix multiplication.

```

1:  $\mathbf{A}' \leftarrow \mathbf{A}, \mathbf{B}' \leftarrow \mathbf{B}, \mathbf{D}_A \leftarrow \mathbf{I}, \mathbf{D}_B \leftarrow \mathbf{I}$ 
2: alternate
3:   O step
4:      $\mathbf{D}'_A \leftarrow \text{diag}(\max_k |a'_{ik}|)$ 
5:      $\mathbf{D}_A \leftarrow \mathbf{D}_A \mathbf{D}'_A$ 
6:      $\mathbf{A}' \leftarrow (\mathbf{D}'_A)^{-1} \mathbf{A}'$ 
7:      $\mathbf{D}'_B \leftarrow \text{diag}(\max_k |b'_{kj}|)$ 
8:      $\mathbf{D}_B \leftarrow \mathbf{D}'_B \mathbf{D}_B$ 
9:      $\mathbf{B}' \leftarrow \mathbf{B}' (\mathbf{D}'_B)^{-1}$ 
10:  end
11:  I step
12:     $\mathbf{D} \leftarrow \text{diag}\left(\sqrt{\frac{\max_j |b'_{kj}|}{\max_i |a'_{ik}|}}\right)$ 
13:     $\mathbf{A}' \leftarrow \mathbf{A}' \mathbf{D}$ 
14:     $\mathbf{B}' \leftarrow \mathbf{D}^{-1} \mathbf{B}'$ 
15:  end
16: until converged
17:  $\mathbf{C}' \leftarrow \mathbf{A}' \cdot \mathbf{B}'$  with fast matrix multiplication.
18:  $\mathbf{C} \leftarrow \mathbf{D}_A \mathbf{C}' \mathbf{D}_B$ 

```

We now state the main result of this section.

THEOREM 13. *The sequence*

$$r_i^{(t)} s_j^{(t)} \|\mathbf{A}^{(t)}\| \|\mathbf{B}^{(t)}\| \quad \text{for } t = 0, 1, \dots$$

is monotonically nonincreasing and converges linearly.

Proof. See Appendix D. □

This result implies that we can safely iteratively apply inside and outside scaling to improve the error bounds, but this process provides diminishing returns.

Next, we introduce our stopping criterion, which requires the following additional notation. Whenever step t is an O step, we use $r_i^{(t)}$ and $s_j^{(t)}$ to denote the diagonal elements of the matrices \mathbf{D}'_A and \mathbf{D}'_B , respectively, that we compute in step t . Similarly, if step t is an I step, $p_k^{(t)}$ denotes the diagonal elements of \mathbf{D} .

The stopping criterion works as follows. We test the intermediate scaling factors $p_k^{(t)}$, $r_i^{(t)}$, and $s_j^{(t)}$ in each iteration starting with the one that immediately follows the first O step. In the I steps, we test whether all of the $p_k^{(t)}$ fall within the interval $[(1 + \tau)^{-\frac{1}{4}}, (1 + \tau)^{\frac{1}{4}}]$, and in the O steps, we test whether all of the $r_i^{(t)}$ and $s_j^{(t)}$ are greater than the threshold $(1 + \tau)^{-1/2}$. Whenever one of these conditions is true, Theorem 14 below states that we are within a relative distance τ from the limit, and so we stop iterating. In practice, we may just specify t steps of scaling a priori, so as to have a better handle on the overhead of the preprocessing. We explore the performance overhead of the preprocessing in subsection 6.5.

We now state the theorem that justifies the stopping criterion. As we show in Appendix D, the sequences $r_i^{(t)}$, $s_j^{(t)}$, $\|\mathbf{A}^{(t)}\|$, and $\|\mathbf{B}^{(t)}\|$ converge. We use a superscript \star to denote their limits, so that

$$r_i^{(t)} \rightarrow r_i^{(\star)}, \quad s_j^{(t)} \rightarrow s_j^{(\star)}, \quad \|\mathbf{A}^{(t)}\| \rightarrow \|\mathbf{A}^{(\star)}\|, \quad \|\mathbf{B}^{(t)}\| \rightarrow \|\mathbf{B}^{(\star)}\|,$$

and we let

$$\mu_{ij}^{(t)} = \frac{|r_i^{(t)} s_j^{(t)} \| \mathbf{A}^{(t)} \| \| \mathbf{B}^{(t)} \| - r_i^{(*)} s_j^{(*)} \| \mathbf{A}^{(*)} \| \| \mathbf{B}^{(*)} \|}{|r_i^{(*)} s_j^{(*)} \| \mathbf{A}^{(*)} \| \| \mathbf{B}^{(*)} \|}$$

be the relative distance of $r_i^{(t)} s_j^{(t)} \| \mathbf{A}^{(t)} \| \| \mathbf{B}^{(t)} \|$ from the limit. We use t_0 to denote the index of the first O step of the iteration.

THEOREM 14. *Let $\tau > 0$ be a user-specified tolerance parameter. We have that for $t = t_0, t_0 + 2, \dots$,*

$$\max_{i,j} \mu_{ij}^{(t+1)} \leq \tau \quad \text{if} \quad \min_k p_k^{(t+1)} \geq (1 + \tau)^{-\frac{1}{4}} \quad \text{and} \quad \max_k p_k^{(t+1)} \leq (1 + \tau)^{\frac{1}{4}},$$

and

$$\max_{i,j} \mu_{ij}^{(t+2)} \leq \tau \quad \text{if} \quad \min_i r_i^{(t+2)} \geq (1 + \tau)^{-\frac{1}{2}} \quad \text{and} \quad \min_j s_j^{(t+2)} \geq (1 + \tau)^{-\frac{1}{2}}.$$

Proof. See Appendix E. □

Finally, we note that Algorithm 3 does not specify which form of scaling to apply first. While the analysis in this section applies to either choice, we note that the limits of the two sequences are not identical—they depend on which step is applied first. We conclude this subsection with examples demonstrating that these two choices can produce significantly different results (in the case of one iteration of Algorithm 3) and that neither choice is always preferable. Example 15 shows a case where performing outside followed by inside scaling is more accurate than performing inside followed by outside scaling; the opposite is true for the case of Example 16.

Example 15. Consider the matrices

$$\mathbf{A} = \begin{bmatrix} 1 & z^{-1} \\ 1 & 1 \end{bmatrix}, \quad \mathbf{B} = \begin{bmatrix} z & 1 \\ z & 1 \end{bmatrix}, \quad \mathbf{C} = \mathbf{A} \cdot \mathbf{B} = \begin{bmatrix} 1 + z & 1 + z^{-1} \\ 2z & 2 \end{bmatrix}$$

for small $z > 0$. We consider one step of alternating scaling. Performing outside and then inside scaling computes

$$(19) \quad \mathbf{A}' \leftarrow \begin{bmatrix} z & 1 \\ 1 & 1 \end{bmatrix}, \quad \mathbf{B}' \leftarrow \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix}, \quad \mathbf{C}' \leftarrow \begin{bmatrix} 1 + z & 1 + z \\ 2 & 2 \end{bmatrix},$$

and inside followed by outside scaling computes

$$(20) \quad \mathbf{A}' \leftarrow \begin{bmatrix} z^{1/2} & 1 \\ 1 & z^{1/2} \end{bmatrix}, \quad \mathbf{B}' \leftarrow \begin{bmatrix} z^{1/2} & z^{1/2} \\ 1 & 1 \end{bmatrix}, \quad \mathbf{C}' \leftarrow \begin{bmatrix} 1 + z^{1/4} & 1 + z^{1/4} \\ 2 \cdot z^{1/2} & 2 \cdot z^{1/2} \end{bmatrix}.$$

Consider the computation of entry c_{21} with Strassen’s algorithm:

$$m_2 = (a'_{21} + a'_{22})b'_{11}, \quad m_4 = a'_{22}(b'_{21} - b'_{11}), \quad c'_{21} = m_2 + m_4.$$

With \mathbf{A}' and \mathbf{B}' in (19), all subterms are $O(1)$ and c'_{21} is $O(1)$, whereas for \mathbf{A}' and \mathbf{B}' in (20), there are $O(z^{1/4})$ subterms and c'_{21} is $O(z^{1/2})$.

From Proposition 12, the absolute error bound for entry $c_{21} = O(z)$ with no scaling is $O(1/z)$, with outside and then inside is $O(z)$, and with inside and then outside is $O(1/z^{1/2})$. Thus, using only one step of Algorithm 3, the accuracy of starting with an O step can be much better than that of starting with an I step.

Example 16. Consider the matrices

$$\mathbf{A} = \begin{bmatrix} 1 & z \\ z & z \end{bmatrix}, \quad \mathbf{B} = \begin{bmatrix} z & 1 \\ 1 & z^{-1} \end{bmatrix}, \quad \mathbf{C} = \mathbf{A} \cdot \mathbf{B} = \begin{bmatrix} 2z & 2 \\ z + z^2 & 1 + z \end{bmatrix}$$

for small $z > 0$. We consider one step of alternating scaling. Performing outside and then inside scaling computes

$$(21) \quad \mathbf{A}' \leftarrow \begin{bmatrix} 1 & z \\ 1 & 1 \end{bmatrix}, \quad \mathbf{B}' \leftarrow \begin{bmatrix} z & 1 \\ z & 1 \end{bmatrix}, \quad \mathbf{C}' \leftarrow \begin{bmatrix} z + z^2 & 1 + z \\ 2z & 2 \end{bmatrix},$$

and inside followed by outside scaling computes

$$(22) \quad \mathbf{A}' \leftarrow \begin{bmatrix} 1 & 1 \\ z & 1 \end{bmatrix}, \quad \mathbf{B}' \leftarrow \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix}, \quad \mathbf{C}' \leftarrow \begin{bmatrix} 2 & 2 \\ 1 + z & 1 + z \end{bmatrix}.$$

Consider the computation of entry c_{21} with Strassen’s algorithm:

$$m_2 = (a'_{21} + a'_{22})b'_{11}, \quad m_4 = a'_{22}(b'_{21} - b'_{11}), \quad c'_{21} = m_2 + m_4.$$

With \mathbf{A}' and \mathbf{B}' in (21), c'_{21} is $O(z)$ but there are $O(1)$ subterms, whereas for \mathbf{A}' and \mathbf{B}' in (22), c'_{21} and all subterms are $O(1)$. This case illustrates that when using only one step of Algorithm 3, the accuracy of starting with an I step can be much better than that of starting with an O step.

6.4. Scaling is not always enough. We now provide a simple example that shows how Strassen’s algorithm computes a result with large relative error, using any of the scaling algorithms presented in this section.

Example 17. Consider the matrices

$$(23) \quad \mathbf{A} = \begin{bmatrix} 1 & z \\ z & 1 \end{bmatrix}, \quad \mathbf{B} = \begin{bmatrix} 1 & z \\ z & 1 \end{bmatrix}, \quad \mathbf{C} = \mathbf{A} \cdot \mathbf{B} = \begin{bmatrix} 1 + z^2 & 2z \\ 2z & 1 + z^2 \end{bmatrix}$$

for small $z > 0$. In this case, both outside and inside scaling leave the matrix unchanged. When computing c_{12} ,

$$\begin{aligned} m_3 &= a_{11}(b_{12} - b_{22}) = 1(z - 1), \\ m_5 &= (a_{11} + a_{12})b_{22} = (1 + z)1, \\ c_{12} &= m_3 + m_5. \end{aligned}$$

There are subterms on the order of unity, so the relative error is $O(1/z)$.

6.5. Numerical experiments. We tested the scaling algorithms on samples of random matrices whose entries were not as contrived as those in the prior sections. We used a sample of $\mathbf{A} \in \mathbb{R}^{N \times N}$ and $\mathbf{B}^{N \times N}$ from the following distributions:

1. $a_{ij}, b_{ij} \sim \text{Uniform}(0, 1)$
2. $a_{ij} \sim \text{Uniform}(0, 1/N^2)$ if $j > N/2$; otherwise, $a_{ij} \sim \text{Uniform}(0, 1)$; $b_{ij} \sim \text{Uniform}(0, 1/N^2)$ if $i < N/2$; otherwise, $b_{ij} \sim \text{Uniform}(0, 1)$
3. $a_{ij} \sim \text{Uniform}(0, N^2)$ if $i < N/2$ and $j > N/2$; otherwise, $a_{ij} \sim \text{Uniform}(0, 1)$; $b_{ij} \sim \text{Uniform}(0, 1/N^2)$ if $j < N/2$; otherwise, $b_{ij} \sim \text{Uniform}(0, 1)$.

Samples from the first distribution are well-behaved for fast matrix multiplication algorithms. On the other hand, samples from the second and third distributions are adversarial and model the matrices in Examples 10 and 15, respectively.




FIG. 3. Relative error of Strassen's algorithm as a function of the number of recursive steps, L , for several scaling techniques. The results in each plot are for matrices \mathbf{A} and \mathbf{B} sampled from different probability distributions. Left: Stability is well-behaved, and no scaling is necessary for small relative errors. Middle: The matrices are adversarial, and inside-outside and 2-times repeated outside-inside scaling have the smallest relative errors. Right: The matrices are adversarial, and outside, outside-inside, and 2-times repeated outside-inside scaling have the smallest relative errors.

We sampled 100 pairs of matrices ($N = 2000$) from each distribution and computed the error of Strassen's algorithm with L recursive levels, $L = 1, 2, \dots, 6$. Specifically, the error was the maximum value of $\max_{ij} |\hat{c}_{ij} - c_{ij}|/|c_{ij}|$ over the 100 samples, where \mathbf{C} was computed with quadruple precision. Figure 3 plots these errors. For the first probability distribution, the relative errors are all roughly the same. With the second distribution, only inside-outside scaling and 2-times repeated outside-inside scaling compute relatively accurate solutions. In this case, inside and outside-inside scaling are moderately more accurate than no scaling or outside scaling, but they still produce relative errors several orders of magnitude larger than the best case. Finally, for the third distribution, inside scaling and no scaling result in much larger relative errors, and inside-outside scaling is slightly worse than outside, outside-inside, or 10-times repeated outside-inside scaling. These experiments demonstrate that with no prior knowledge of the distribution, repeated outside-inside scaling is the safe choice for fast matrix multiplication.

Each iteration of outside or inside scaling is $O(MK + KN + MN)$ flops, so scaling does not affect the asymptotic performance. However, quadratic costs do affect the practical implementation of fast matrix multiplication [3]. Subsequently, we tested the performance impact of scaling. We use *effective gflops* [3, 22] to measure the performance of multiplying an $M \times K$ matrix by a $K \times N$ matrix:

$$(24) \quad \frac{2 \cdot MKN - MN}{\text{time in seconds}} \cdot 1e-9.$$

This lets us compare fast matrix multiplication algorithms to the classical algorithm on a familiar inverse-time scale. All experiments were conducted on a single compute node on NERSC's Edison machine. Each node has two 12-core Intel 2.4 GHz Ivy Bridge processors and 64 GB of memory. Our experiments were single-threaded. We report the median of five trials for each timing result.

Figure 4 summarizes the performance results for Strassen's algorithm ($L = 1$), with and without two steps of Algorithm 3, for multiplying square matrices of dimension N . There is a noticeable impact on performance. Strassen's algorithm without scaling outperforms the classical algorithm for $N \geq 2500$, while scaling pushes this threshold to $N \geq 3500$. As N grows, the performance impact of scaling gets smaller. This follows from the asymptotic analysis—as N grows, the impact from quadratic terms shrinks.

7. Discussion. One of the central components of our algorithmic error analysis is that two data-independent quantities drive the error bounds for fast matrix




FIG. 4. Performance of Strassen’s algorithm ($L = 1$), with and without two steps of inside-outside scaling, and the classical algorithm.

multiplication. First, Q captures the accumulation error from adding matrices. Second, E bounds the growth in the magnitude of intermediate terms. Our results in section 5 show that having a small E is important, but this does not fully characterize stability in practice. The same result has been observed when comparing Strassen’s algorithm and the Winograd variant [16]. An encouraging result from our experiments is that the number of nonzeros in the \mathbf{U} , \mathbf{V} , and \mathbf{W} matrices, which determines the constant in the computational complexity, is positively correlated with E . In other words, for a given base case and algorithm rank, by improving performance, we generally also improve stability. Another lesson from our analysis is that we should not think of using fast algorithms asymptotically but rather as having a fixed number of recursive levels. This leads to better performance in practice [3] and also to the improved error bounds and numerical stability presented in sections 4 and 5. Finally, because the principal quantities for understanding algorithmic error (E and Q) are independent of the asymptotic complexity, we have new metrics over which to optimize when searching for fast matrix multiplication algorithms.

For performance reasons, the best choice of fast algorithm depends on the shape of the matrices being multiplied [3]. In general, a choice of algorithm can be made at each recursive level. Subsequently, we believe that uniform nonstationary algorithms are the right choice in practice for achieving the best performance. Theorem 5 provides the appropriate error bounds for this case.

The analysis in subsection 4.5 formalizes the error analysis for existing techniques to improve stability of Strassen’s algorithm and the Winograd variant [8, 9] and also generalizes the approach for all fast matrix multiplication algorithms. The analysis provides the formula over which to optimize when considering nonuniform, nonstationary algorithms. However, finding the best algorithm is a combinatorial optimization problem that grows exponentially in the number of recursive levels. Algorithm design in this space is an interesting avenue for future research.

Using these algorithmic techniques improves the normwise accuracy of the computed product. However, because the errors are normwise, small elements of the product can be computed less accurately than warranted by their condition numbers. By pre- and postprocessing the data, we can improve componentwise accuracy as well. Specifically, we analyzed a hierarchy of diagonal scaling techniques that reduce the number of cases where fast matrix multiplication yields relatively inaccurate small entries in the product. Nevertheless, there are cases that cannot be solved by our

diagonal scaling algorithms (e.g., Example 17). When scaling helps, a few iterations are sufficient, and this is backed up by Theorem 13.

The asymptotic operation cost of diagonal scaling is proportional to the size of the matrices, so it is dominated by the cost of current matrix multiplication algorithms. In our experiments, we found that scaling does incur a noticeable performance penalty for reasonably sized matrices, but fast matrix multiplication with diagonal scaling still can outperform the classical algorithm. We note that our diagonal scaling implementation is not fully optimized; for example, it is possible to overlap inside and outside scaling and delay updating actual matrix entries, both of which can reduce the memory traffic overhead.

Appendix A. Strassen's algorithm. Strassen's algorithm [25] is a $\langle 2, 2, 2 \rangle$ algorithm specified by the following \mathbf{U} , \mathbf{V} , and \mathbf{W} matrices:

$$\mathbf{U} = \begin{bmatrix} 1 & 0 & 1 & 0 & 1 & -1 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 1 \\ 0 & 1 & 0 & 0 & 0 & 1 & 0 \\ 1 & 1 & 0 & 1 & 0 & 0 & -1 \end{bmatrix},$$

$$\mathbf{V} = \begin{bmatrix} 1 & 1 & 0 & -1 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 1 \\ 1 & 0 & -1 & 0 & 1 & 0 & 1 \end{bmatrix},$$

$$\mathbf{W} = \begin{bmatrix} 1 & 0 & 0 & 1 & -1 & 0 & 1 \\ 0 & 1 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 1 & 0 & 0 \\ 1 & -1 & 1 & 0 & 0 & 1 & 0 \end{bmatrix}.$$

Note that the rows of \mathbf{U} and \mathbf{V} correspond to a column-major ordering of the entries of the input matrices, and the rows of \mathbf{W} correspond to a row-major ordering of the output matrix, following the convention of previous work [6, 17]. We point out that this algorithm is cyclic-invariant, so that $\llbracket \mathbf{U}, \mathbf{V}, \mathbf{W} \rrbracket = \llbracket \mathbf{W}, \mathbf{U}, \mathbf{V} \rrbracket = \llbracket \mathbf{V}, \mathbf{W}, \mathbf{U} \rrbracket$ (up to permutations on the columns of the matrices), which implies that all three rotations have the same Q and E values.

Note that the rows of \mathbf{U} and \mathbf{V} correspond to a column-major ordering of the entries of the input matrices, and the rows of \mathbf{W} correspond to a row-major ordering of the output matrix, which implies that $\llbracket \mathbf{W}, \mathbf{U}, \mathbf{V} \rrbracket$ is an algorithm for $\langle 4, 2, 4 \rangle$ and $\llbracket \mathbf{V}, \mathbf{W}, \mathbf{U} \rrbracket$ is an algorithm for $\langle 2, 4, 4 \rangle$. However, $\llbracket \mathbf{V}, \mathbf{W}, \mathbf{U} \rrbracket$ yields an $E = 102$, which is greater than $\llbracket \mathbf{U}, \mathbf{V}, \mathbf{W} \rrbracket$'s 89. The E value can be maintained for base case $\langle 2, 4, 4 \rangle$ by using the following different transformation that corresponds to transposing the matrix multiplication: $\llbracket \mathcal{P}_{4,2} \mathbf{V}, \mathcal{P}_{4,4} \mathbf{U}, \mathcal{P}_{2,4} \mathbf{W} \rrbracket$, where $\mathcal{P}_{m,n}$ is the so-called vec permutation matrix [15], exchanging column-ordering for row-ordering in a vectorized $m \times n$ matrix.

Appendix D. Convergence analysis of alternating scaling. In this appendix we prove Theorem 13. We start with its first part.

LEMMA 18. *The sequence*

$$r_i^{(t)} s_j^{(t)} \|\mathbf{A}^{(t)}\| \|\mathbf{B}^{(t)}\| \quad \text{for } t = 0, 1, \dots$$

is monotonically nonincreasing.

Proof. If step t is an O step, then

$$\|\mathbf{A}^{(t)}\| = \|\mathbf{B}^{(t)}\| = 1, \quad r_i^{(t)} = r_i^{(t-1)} \|a_{i,:}^{(t-1)}\|, \quad s_j^{(t)} = s_j^{(t-1)} \|b_{:,j}^{(t-1)}\|,$$

and therefore

$$\begin{aligned} r_i^{(t)} s_j^{(t)} \|\mathbf{A}^{(t)}\| \|\mathbf{B}^{(t)}\| &= r_i^{(t-1)} s_j^{(t-1)} \|a_{i,:}^{(t-1)}\| \|b_{:,j}^{(t-1)}\| \\ &\leq r_i^{(t-1)} s_j^{(t-1)} \|\mathbf{A}^{(t-1)}\| \|\mathbf{B}^{(t-1)}\|. \end{aligned}$$

Next, assume that step t is an I step. Column k of \mathbf{A}' is transformed so that

$$a_{ik}^{(t)} = \left(\frac{\|b_{k,:}^{(t-1)}\|}{\|a_{:,k}^{(t-1)}\|} \right)^{\frac{1}{2}} a_{ik}^{(t-1)},$$

and therefore,

$$(25) \quad \|a_{:,k}^{(t)}\| = \left(\frac{\|b_{k,:}^{(t-1)}\|}{\|a_{:,k}^{(t-1)}\|} \right)^{\frac{1}{2}} \|a_{:,k}^{(t-1)}\| = \sqrt{\|a_{:,k}^{(t-1)}\| \|b_{k,:}^{(t-1)}\|},$$

and similarly

$$(26) \quad \|b_{k,:}^{(t)}\| = \sqrt{\|a_{:,k}^{(t-1)}\| \|b_{k,:}^{(t-1)}\|}.$$

Hence,

$$\|\mathbf{A}^{(t)}\| = \|\mathbf{B}^{(t)}\| = \max_k \sqrt{\|a_{:,k}^{(t-1)}\| \|b_{k,:}^{(t-1)}\|}, \quad r_i^{(t)} = r_i^{(t-1)}, \quad s_j^{(t)} = s_j^{(t-1)},$$

and therefore,

$$\begin{aligned} r_i^{(t)} s_j^{(t)} \|\mathbf{A}^{(t)}\| \|\mathbf{B}^{(t)}\| &= r_i^{(t-1)} s_j^{(t-1)} \left(\max_k \sqrt{\|a_{:,k}^{(t-1)}\| \|b_{k,:}^{(t-1)}\|} \right)^2 \\ &= r_i^{(t-1)} s_j^{(t-1)} \max_k (\|a_{:,k}^{(t-1)}\| \|b_{k,:}^{(t-1)}\|) \\ &\leq r_i^{(t-1)} s_j^{(t-1)} \|\mathbf{A}^{(t-1)}\| \|\mathbf{B}^{(t-1)}\|. \quad \square \end{aligned}$$

Next, we prove that the factors in the sequence of Theorem 13 converge individually. This is required in the subsequent analysis.

LEMMA 19. *The sequences $r_i^{(t)}$, $s_j^{(t)}$, $\|\mathbf{A}^{(t)}\|$, and $\|\mathbf{B}^{(t)}\|$ for $t = 0, 1, \dots$ converge.*

Proof. As we showed in the proof of Lemma 18,

$$\begin{aligned} \|\mathbf{A}^{(t)}\| &= \|\mathbf{B}^{(t)}\| = 1, \\ \|\mathbf{A}^{(t+1)}\| &= \|\mathbf{B}^{(t+1)}\| = \max_k \sqrt{\|a_{:,k}^{(t)}\| \|b_{k,:}^{(t)}\|} \quad \text{for } t = t_0, t_0 + 2, \dots \end{aligned}$$

Therefore,

$$\|\mathbf{A}^{(t+1)}\| = \max_k \sqrt{\|a_{:,k}^{(t)}\| \|b_{k,:}^{(t)}\|} \leq \sqrt{\|\mathbf{A}^{(t)}\| \|\mathbf{B}^{(t)}\|} = 1,$$

and hence

$$(27) \quad r_i'^{(t+2)} = \|a_{i,:}^{(t+1)}\| \leq \|\mathbf{A}^{(t+1)}\| \leq 1.$$

The sequence $r_i^{(t)}$ satisfies

$$\begin{aligned} r_i^{(t_0)} &= r_i^{(t_0+1)} = r_i'^{(t_0)}, \\ r_i^{(t_0+2)} &= r_i^{(t_0+3)} = r_i'^{(t_0)} r_i'^{(t_0+2)} \\ &\vdots \quad \quad \quad \vdots \end{aligned}$$

It is nonnegative because $r_i'^{(t)} \geq 0$, it is monotonically nonincreasing by (27), and hence it must converge. The same is true for $s_j^{(t)}$.

Consider the effect of the first t steps of the iteration on \mathbf{A}' and \mathbf{B}' . The cumulative effect of the O steps is to divide the rows of \mathbf{A}' by $r_i^{(t)}$ and the columns of \mathbf{B}' by $s_j^{(t)}$, and that of the I steps is to make sure that every column of \mathbf{A}' is equal in norm to the corresponding row of \mathbf{B}' . Therefore,

$$a_{ik}^{(t)} = a_{ik} \frac{1}{r_i^{(t)}} \left(\frac{\max_j |b_{kj}/s_j^{(t-1)}|}{\max_i |a_{ik}/r_i^{(t-1)}|} \right)^{\frac{1}{2}} \quad \text{for } t = t_0 + 1, t_0 + 2, \dots,$$

which shows that the convergence of $r_i^{(t)}$ and $s_j^{(t)}$ guarantees the convergence of $a_{ik}^{(t)}$, and hence also the convergence of $\|\mathbf{A}^{(t)}\|$. The same is true for $\|\mathbf{B}^{(t)}\|$. \square

The following lemma shows that the intermediate scaling factors that we compute in each step rapidly converge to 1. We use the notation

$$\begin{aligned} w^{(t)} &= \max \left(\max_i |\log r_i'^{(t)}|, \max_j |\log s_j'^{(t)}| \right), \\ w^{(t+1)} &= \max_k |\log p_k'^{(t+1)}| \quad \text{for } t = t_0, t_0 + 2, \dots \end{aligned}$$

LEMMA 20. *The following bounds hold:*

$$\begin{aligned} w^{(t)} &\leq w^{(t-1)}, \\ w^{(t+1)} &\leq 0.5w^{(t)} \quad \text{for } t = t_0 + 2, t_0 + 4, \dots \end{aligned}$$

Proof. Assume that $t = t_0 + 2, t_0 + 4, \dots$. Because step $t - 2$ is an O step, there is a column g so that $|a_{ig}^{(t-2)}| = 1$, and therefore,

$$r_i^{(t)} = \max_k |a_{ik}^{(t-1)}| = \max_k |a_{ik}^{(t-2)} p_k^{(t-1)}| \geq |a_{ig}^{(t-2)} p_g^{(t-1)}| = p_g^{(t-1)}.$$

Taking logarithms yields

$$\log r_i^{(t)} \geq \log p_g^{(t-1)}.$$

Both sides of this inequality are nonpositive because $r_i^{(t)} \leq 1$ by (27), and so

$$|\log r_i^{(t)}| = -\log r_i^{(t)} \leq -\log p_g^{(t-1)} = |\log p_g^{(t-1)}|.$$

A similar analysis shows that

$$|\log s_j^{(t)}| \leq |\log p_f^{(t-1)}|$$

for a suitably defined row f , and these two inequalities imply the first bound in the statement of the lemma.

Next, let us prove the second bound. We have that

$$(p_k^{(t+1)})^2 = \frac{\max_j |b_{kj}^{(t)}|}{\max_i |a_{ik}^{(t)}|} = \frac{\max_j |b_{kj}^{(t-1)} / s_j^{(t)}|}{\max_i |a_{ik}^{(t-1)} / r_i^{(t)}|} \leq \frac{\max_j |b_{kj}^{(t-1)}| \max_j (1/s_j^{(t)})}{\max_i |a_{ik}^{(t-1)} / r_i^{(t)}|}.$$

Inequality (27) states that $r_i^{(t)} \leq 1$, and therefore

$$\max_i |a_{ik}^{(t-1)} / r_i^{(t)}| \geq \max_i |a_{ik}^{(t-1)}|,$$

which we substitute into the previous inequality, obtaining

$$(p_k^{(t+1)})^2 \leq \frac{\max_j |b_{kj}^{(t-1)}| \max_j (1/s_j^{(t)})}{\max_i |a_{ik}^{(t-1)}|}.$$

By (25) and (26),

$$\max_i |a_{ik}^{(t-1)}| = \|a_{:,k}^{(t-1)}\| = \|b_{k,:}^{(t-1)}\| = \max_j |b_{kj}^{(t-1)}|,$$

which implies

$$(p_k^{(t+1)})^2 \leq \max_j (1/s_j^{(t)}).$$

A similar analysis shows that

$$(p_k^{(t+1)})^2 \geq \frac{1}{\max_i (1/r_i^{(t)})}.$$

Taking the logarithm of these two bounds and interchanging the positions of the logarithms with those of the max operators yields

$$-\max_i \left(\log(1/r_i^{(t)}) \right) \leq 2 \log p_k^{(t+1)} \leq \max_j \left(\log(1/s_j^{(t)}) \right).$$

Because $r_i^{(t)} \leq 1$, we have that $\log(1/r_i^{(t)}) = |\log r_i^{(t)}|$, and similarly for $s_j^{(t)}$. Applying this to the previous inequality yields

$$-\max_i |\log r_i^{(t)}| \leq 2 \log p_k^{(t+1)} \leq \max_j |\log s_j^{(t)}|,$$

and therefore

$$2|\log p_k^{(t+1)}| \leq \max\left(\max_i |\log r_i^{(t)}|, \max_j |\log s_j^{(t)}|\right),$$

which implies the second bound in the statement of the lemma. \square

The following lemma proves linear convergence, and therefore completes the proof of Theorem 13.

LEMMA 21. *There is a sequence $\nu^{(t)}$ so that*

$$\mu_{ij}^{(t)} \leq \nu^{(t)}, \quad \mu_{ij}^{(t+1)} \leq \nu^{(t+1)},$$

and

$$\nu^{(t+1)} = \nu^{(t)}, \quad \nu^{(t+2)} \leq 0.5 \nu^{(t)} \quad \text{for } t = t_0, t_0 + 2, \dots$$

Proof. Assume that $t = t_0, t_0 + 2, \dots$. Rearranging the definition of $\mu_{ij}^{(t)}$, we may write

$$\mu_{ij}^{(t)} = (r_i^{(t)} s_j^{(t)} \|\mathbf{A}^{(t)}\| \|\mathbf{B}^{(t)}\|) (r_i^{(*)} s_j^{(*)} \|\mathbf{A}^{(*)}\| \|\mathbf{B}^{(*)}\|)^{-1} - 1.$$

We have that

$$r_i^{(t)} = r_i^{(t_0)} r_i^{(t_0+2)} \dots r_i^{(t)}, \quad s_j^{(t)} = s_j^{(t_0)} s_j^{(t_0+2)} \dots s_j^{(t)}, \quad \|\mathbf{A}^{(t)}\| = \|\mathbf{B}^{(t)}\| = 1,$$

and therefore

$$r_i^{(*)} = r_i^{(t_0)} r_i^{(t_0+2)} \dots, \quad s_j^{(*)} = s_j^{(t_0)} s_j^{(t_0+2)} \dots, \quad \|\mathbf{A}^{(*)}\| = \|\mathbf{B}^{(*)}\| = 1.$$

Substituting this into the above yields

$$\begin{aligned} \mu_{ij}^{(t)} &= (r_i^{(t_0)} s_j^{(t_0)} r_i^{(t_0+2)} s_j^{(t_0+2)} \dots r_i^{(t)} s_j^{(t)}) (r_i^{(t_0)} s_j^{(t_0)} r_i^{(t_0+2)} s_j^{(t_0+2)} \dots)^{-1} - 1 \\ &= (r_i^{(t+2)} s_j^{(t+2)} r_i^{(t+4)} s_j^{(t+4)} \dots)^{-1} - 1. \end{aligned}$$

Applying the definition of $w^{(t)}$ to this, we obtain

$$\begin{aligned} \mu_{ij}^{(t)} &= (r_i^{(t+2)} s_j^{(t+2)} r_i^{(t+4)} s_j^{(t+4)} \dots)^{-1} - 1 \\ &= \exp(-\log r_i^{(t+2)} - \log s_j^{(t+2)} - \log r_i^{(t+4)} - \log s_j^{(t+4)} - \dots) - 1 \\ &= \exp(|\log r_i^{(t+2)}| + |\log s_j^{(t+2)}| + |\log r_i^{(t+4)}| + |\log s_j^{(t+4)}| + \dots) - 1 \\ &\leq \exp(2w^{(t+2)} + 2w^{(t+4)} + \dots) - 1. \end{aligned}$$

We define

$$\nu^{(t)} = \exp(2w^{(t+2)} + 2w^{(t+4)} + \dots) - 1, \quad \nu^{(t+1)} = \nu^{(t)},$$

thereby guaranteeing that $\mu_{ij}^{(t)} \leq \nu^{(t)}$, as the lemma states. Since $r_i^{(t)} s_j^{(t)} \|\mathbf{A}^{(t)}\| \|\mathbf{B}^{(t)}\|$ is monotonically nonincreasing, so is its relative distance to its limit, meaning that

$$\mu_{ij}^{(t+1)} \leq \mu_{ij}^{(t)},$$

and hence

$$\mu_{ij}^{(t+1)} \leq \mu_{ij}^{(t)} \leq \nu^{(t)} = \nu^{(t+1)}.$$

This proves another condition in the statement of the lemma, leaving us with only the condition $\nu^{(t+2)} \leq 0.5 \nu^{(t)}$ to prove.

By Lemma 20,

$$\begin{array}{ll} w^{(t+3)} \leq 0.5w^{(t+2)}, & w^{(t+4)} \leq w^{(t+3)} \leq 0.5w^{(t+2)}, \\ w^{(t+5)} \leq 0.5w^{(t+4)} \leq 0.25w^{(t+2)}, & w^{(t+6)} \leq w^{(t+5)} \leq 0.25w^{(t+2)} \\ \vdots & \vdots \end{array}$$

Therefore,

$$\begin{aligned} w^{(t+2)} &= (0.5 + 0.25 + \dots)w^{(t+2)} \\ &= 0.5w^{(t+2)} + 0.25w^{(t+2)} + \dots \\ &\geq w^{(t+4)} + w^{(t+6)} + \dots, \end{aligned}$$

and thus

$$\exp(2w^{(t+2)}) \geq \exp(2w^{(t+4)} + 2w^{(t+6)} + \dots).$$

Applying this bound to the definition of $\nu^{(t)}$, we obtain

$$\begin{aligned} \nu^{(t)} &= \exp(2w^{(t+2)} + 2w^{(t+4)} + \dots) - 1 \\ &= \exp(2w^{(t+2)}) \exp(2w^{(t+4)} + 2w^{(t+6)} + \dots) - 1 \\ &\geq \left(\exp(2w^{(t+4)} + 2w^{(t+6)} + \dots) \right)^2 - 1. \end{aligned}$$

We define $x = \exp(2w^{(t+4)} + 2w^{(t+6)} + \dots)$, so that the above expression has the form $x^2 - 1$ and $\nu^{(t+2)} = x - 1$. Then we have

$$\nu^{(t)} \geq x^2 - 1 = (2 + x - 1)(x - 1) = (2 + \nu^{(t+2)}) \nu^{(t+2)} \geq 2 \nu^{(t+2)},$$

which implies that $\nu^{(t+2)} \leq 0.5 \nu^{(t)}$ as the lemma states. \square

Finally, we show that the analysis in Lemma 21 is asymptotically sharp.

LEMMA 22. *There are matrices \mathbf{A} and \mathbf{B} and indices i and j so that*

$$\mu_{ij}^{(t+1)} = \mu_{ij}^{(t)}, \quad \mu_{ij}^{(t+2)} / \mu_{ij}^{(t)} \rightarrow 0.5 \quad \text{for } t = t_0, t_0 + 2, \dots$$

Proof. Let

$$\mathbf{A} = \begin{bmatrix} 1 & 0 \\ 1 & 2^{-2^v} \end{bmatrix}, \quad \mathbf{B} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

for some integer v . Let us start the iteration with an O step, and assume that $t = t_0, t_0 + 2, \dots$. A straightforward calculation, which we omit for brevity, shows

that

$$\begin{aligned} \mathbf{A}^{(t)} &= \begin{bmatrix} 1 & 0 \\ 1 & 2^{-2^v-(t-t_0)/2} \end{bmatrix}, & \begin{bmatrix} r_1^{(t)} \\ r_2^{(t)} \end{bmatrix} &= \begin{bmatrix} 1 \\ 1 \end{bmatrix}, \\ \mathbf{B}^{(t)} &= \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, & \begin{bmatrix} s_1^{(t)} \\ s_2^{(t)} \end{bmatrix} &= \begin{bmatrix} 1 \\ 2^{-2^v(1-2^{-(t-t_0)/2})} \end{bmatrix}, \\ \mathbf{A}^{(t+1)} &= \begin{bmatrix} 1 & 0 \\ 1 & 2^{-2^v-(t-t_0)/2-1} \end{bmatrix}, & \begin{bmatrix} r_1^{(t+1)} \\ r_2^{(t+1)} \end{bmatrix} &= \begin{bmatrix} r_1^{(t)} \\ r_2^{(t)} \end{bmatrix}, \\ \mathbf{B}^{(t+1)} &= \begin{bmatrix} 1 & 0 \\ 0 & 2^{-2^v-(t-t_0)/2-1} \end{bmatrix}, & \begin{bmatrix} s_1^{(t+1)} \\ s_2^{(t+1)} \end{bmatrix} &= \begin{bmatrix} s_1^{(t)} \\ s_2^{(t)} \end{bmatrix}, \end{aligned}$$

and therefore

$$\mathbf{A}^{(*)} = \begin{bmatrix} 1 & 0 \\ 1 & 1 \end{bmatrix}, \quad \mathbf{B}^{(*)} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad \begin{bmatrix} r_1^{(*)} \\ r_2^{(*)} \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \end{bmatrix}, \quad \begin{bmatrix} s_1^{(*)} \\ s_2^{(*)} \end{bmatrix} = \begin{bmatrix} 1 \\ 2^{-2^v} \end{bmatrix},$$

and

$$\|\mathbf{A}^{(t)}\| = \|\mathbf{B}^{(t)}\| = \|\mathbf{A}^{(t+1)}\| = \|\mathbf{B}^{(t+1)}\| = \|\mathbf{A}^{(*)}\| = \|\mathbf{B}^{(*)}\| = 1.$$

Substituting the above into the definition of $\mu_{ij}^{(t)}$ yields

$$\begin{aligned} \mu_{22}^{(t)} = \mu_{22}^{(t+1)} &= \frac{|r_2^{(t)} s_2^{(t)} \|\mathbf{A}^{(t)}\| \|\mathbf{B}^{(t)}\| - r_2^{(*)} s_2^{(*)} \|\mathbf{A}^{(*)}\| \|\mathbf{B}^{(*)}\|}{|r_2^{(*)} s_2^{(*)} \|\mathbf{A}^{(*)}\| \|\mathbf{B}^{(*)}\|} \\ (28) \quad &= \frac{|2^{-2^v(1-2^{-(t-t_0)/2})} - 2^{-2^v}|}{|2^{-2^v}|} \\ &= 2^{2^v-(t-t_0)/2} - 1. \end{aligned}$$

Let $x = 2^{2^v-(t-t_0)/2-1}$, so that the above expression has the form $x^2 - 1$. Then $\mu_{22}^{(t+2)} = x - 1$, and we have that

$$\mu_{22}^{(t)} = x^2 - 1 = (2 + x - 1)(x - 1) = (2 + \mu_{22}^{(t+2)}) \mu_{22}^{(t+2)},$$

and therefore

$$\mu_{22}^{(t+2)} / \mu_{22}^{(t)} = 1 / (2 + \mu_{22}^{(t+2)}).$$

From (28) we conclude that $\mu_{22}^{(t+2)} \rightarrow 0$, and therefore $\mu_{22}^{(t+2)} / \mu_{22}^{(t)} \rightarrow 0.5$. □

Appendix E. Proof of Theorem 14.

Proof. In the proof of Lemma 21 (see Appendix D), we show that for all i, j and $t = t_0, t_0 + 2, \dots$,

$$\begin{aligned} \mu_{ij}^{(t+1)} &\leq \mu_{ij}^{(t)}, \\ \mu_{ij}^{(t)} &\leq \exp(2w^{(t+2)} + 2w^{(t+4)} + \dots) - 1, \\ \exp(2w^{(t+4)} + 2w^{(t+6)} + \dots) &\leq \exp(2w^{(t+2)}). \end{aligned}$$

Putting these three statements together yields

$$\begin{aligned}\mu_{ij}^{(t+1)} &\leq \mu_{ij}^{(t)} \leq \exp(2w^{(t+2)} + 2w^{(t+4)} + \dots) - 1 \\ &= \exp(2w^{(t+2)}) \exp(2w^{(t+4)} + 2w^{(t+6)} + \dots) - 1 \\ &\leq \exp(2w^{(t+2)}) \exp(2w^{(t+2)}) - 1 \\ &= \exp(4w^{(t+2)}) - 1.\end{aligned}$$

Lemma 20 guarantees that $w^{(t+2)} \leq w^{(t+1)}$, and substituting this into the above yields

$$(29) \quad \mu_{ij}^{(t+1)} \leq \exp(4w^{(t+1)}) - 1.$$

Similarly,

$$(30) \quad \mu_{ij}^{(t+2)} \leq \exp(2w^{(t+4)} + 2w^{(t+6)} + \dots) - 1 \leq \exp(2w^{(t+2)}) - 1.$$

Next, let us prove the first statement of the theorem. Assume that

$$(1 + \tau)^{-\frac{1}{4}} \leq p_k^{(t+1)} \leq (1 + \tau)^{\frac{1}{4}}$$

for all k . Taking logarithms yields

$$-0.25 \log(1 + \tau) \leq \log p_k^{(t+1)} \leq 0.25 \log(1 + \tau)$$

or, equivalently,

$$|\log p_k^{(t+1)}| \leq 0.25 \log(1 + \tau),$$

and therefore

$$w^{(t+1)} = \max_k |\log p_k^{(t+1)}| \leq 0.25 \log(1 + \tau).$$

Substituting this into (29), we find that

$$\mu_{ij}^{(t+1)} \leq \exp(4w^{(t+1)}) - 1 \leq \exp(4 \cdot 0.25 \log(1 + \tau)) - 1 = \tau$$

for all i, j , which proves the first statement of the theorem.

Let us prove the second statement of the theorem. Assume that

$$(1 + \tau)^{-\frac{1}{2}} \leq r_i^{(t+2)}, \quad (1 + \tau)^{-\frac{1}{2}} \leq s_j^{(t+2)}$$

for all i and j . Taking logarithms yields

$$-0.5 \log(1 + \tau) \leq \log r_i^{(t+2)}, \quad -0.5 \log(1 + \tau) \leq \log s_j^{(t+2)}.$$

We show in the proof of Lemma 19 that $r_i^{(t+2)} \leq 1$, and therefore $\log r_i^{(t+2)} \leq 0$ and similarly for $s_j^{(t+2)}$. Hence

$$|\log r_i^{(t+2)}| \leq 0.5 \log(1 + \tau), \quad |\log s_j^{(t+2)}| \leq 0.5 \log(1 + \tau),$$

and hence

$$w^{(t+2)} = \max \left(\max_i |\log r_i^{(t+2)}|, \max_j |\log s_j^{(t+2)}| \right) \leq 0.5 \log(1 + \tau).$$

Substituting this into (30) yields

$$\mu_{ij}^{(t+2)} \leq \exp(2w^{(t+2)}) - 1 \leq \exp(2 \cdot 0.5 \log(1 + \tau)) - 1 = \tau$$

for all i, j , which proves the second statement of the theorem. \square

Acknowledgments. We thank Madan Musuvathi for providing the $\langle 2, 3, 2 \rangle$ algorithms from the SAT solver. We also thank the editor and anonymous reviewers for their suggestions, which improved the paper's quality and readability.

REFERENCES

- [1] G. BALLARD, J. DEMMEL, O. HOLTZ, B. LIPSHITZ, AND O. SCHWARTZ, *Communication-optimal parallel algorithm for Strassen's matrix multiplication*, in Proceedings of the 24th ACM Symposium on Parallelism in Algorithms and Architectures, ACM, New York, 2012, pp. 193–204, doi:10.1145/2312005.2312044.
- [2] G. BALLARD, J. DEMMEL, O. HOLTZ, AND O. SCHWARTZ, *Graph expansion and communication costs of fast matrix multiplication: Regular submission*, in Proceedings of the 23rd ACM Symposium on Parallelism in Algorithms and Architectures, ACM, New York, 2011, pp. 1–12, doi:10.1145/1989493.1989495.
- [3] A. R. BENSON AND G. BALLARD, *A framework for practical parallel fast matrix multiplication*, in Proceedings of the 20th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, ACM, New York, 2015, pp. 42–53, doi:10.1145/2688500.2688513.
- [4] D. BINI AND G. LOTTI, *Stability of fast algorithms for matrix multiplication*, Numer. Math., 36 (1980), pp. 63–72, doi:10.1007/BF01395989.
- [5] M. BLÄSER, *On the complexity of the multiplication of matrices of small formats*, J. Complexity, 19 (2003), pp. 43–60, doi:10.1016/S0885-064X(02)00007-9.
- [6] R. P. BRENT, *Algorithms for Matrix Multiplication*, Tech. report, Stanford University, Stanford, CA, 1970.
- [7] R. P. BRENT, *Error analysis of algorithms for matrix multiplication and triangular decomposition using Winograd's identity*, Numer. Math., 16 (1970), pp. 145–156, doi:10.1007/BF02308867.
- [8] R. R. CASTRAPEL AND J. L. GUSTAFSON, *Precision Improvement Method for the Strassen/Winograd Matrix Multiplication Method*, U.S. Patent 7,209,939, April 24, 2007, <http://www.google.com/patents/US7209939>.
- [9] P. D'ALBERTO, *The better accuracy of Strassen-Winograd algorithms (FastMMW)*, Adv. Linear Algebra Matrix Theory, 4 (2014), pp. 9–39, doi:10.4236/alamt.2014.41002.
- [10] L. DE MOURA AND N. BJÖRNER, *Z3: An efficient SMT solver*, in Tools and Algorithms for the Construction and Analysis of Systems, Lecture Notes in Comput. Sci. 4963, Springer, Berlin, 2008, pp. 337–340, http://link.springer.com/chapter/10.1007/2F978-3-540-78800-3_24.
- [11] J. DEMMEL, I. DUMITRIU, AND O. HOLTZ, *Fast linear algebra is stable*, Numer. Math., 108 (2007), pp. 59–91, doi:10.1007/s00211-007-0114-x.
- [12] J. DEMMEL, I. DUMITRIU, O. HOLTZ, AND R. KLEINBERG, *Fast matrix multiplication is stable*, Numer. Math., 106 (2007), pp. 199–224, doi:10.1007/s00211-007-0061-6.
- [13] B. DUMITRESCU, *Improving and estimating the accuracy of Strassen's algorithm*, Numer. Math., 79 (1998), pp. 485–499, doi:10.1007/s002110050348.
- [14] F. LE GALL, *Powers of tensors and fast matrix multiplication*, in Proceedings of the 39th International Symposium on Symbolic and Algebraic Computation, ACM, New York, 2014, pp. 296–303, doi:10.1145/2608628.2608664.
- [15] H. V. HENDERSON AND S. R. SEARLE, *The vec-permutation matrix, the vec operator and Kronecker products: A review*, Linear Multilinear Algebra, 9 (1981), pp. 271–288, doi:10.1080/03081088108817379.
- [16] N. J. HIGHAM, *Accuracy and Stability of Numerical Algorithms*, SIAM, Philadelphia, 2002, doi:10.1137/1.9780898718027.
- [17] R. W. JOHNSON AND A. M. MCLOUGHLIN, *Noncommutative bilinear algorithms for 3×3 matrix multiplication*, SIAM J. Comput., 15 (1986), pp. 595–603, doi:10.1137/0215043.
- [18] I. KAPORIN, *A practical algorithm for faster matrix multiplication*, Numer. Linear Algebra Appl., 6 (1999), pp. 687–700, [http://onlinelibrary.wiley.com/doi/10.1002/\(SICI\)1099-1506\(199912\)6:8%3C687::AID-NLA177%3E3.0.CO;2-I/abstract](http://onlinelibrary.wiley.com/doi/10.1002/(SICI)1099-1506(199912)6:8%3C687::AID-NLA177%3E3.0.CO;2-I/abstract).
- [19] I. KAPORIN, *The aggregation and cancellation techniques as a practical tool for faster matrix multiplication*, Theoret. Comput. Sci., 315 (2004), pp. 469–510, doi:10.1016/j.tcs.2004.01.004.
- [20] J. LADERMAN, V. PAN, AND X.-H. SHA, *On practical algorithms for accelerated matrix multiplication*, Linear Algebra Appl., 162–164 (1992), pp. 557–588, doi:10.1016/0024-3795(92)90393-O.

- [21] J. W. LANDSBERG, *New lower bounds for the rank of matrix multiplication*, SIAM J. Comput., 43 (2014), pp. 144–149, doi:10.1137/120880276.
- [22] B. LIPSHITZ, G. BALLARD, J. DEMMEL, AND O. SCHWARTZ, *Communication-avoiding parallel Strassen: Implementation and performance*, in Proceedings of the International Conference for High Performance Computing, Networking, Storage, and Analysis, IEEE Press, Piscataway, NJ, 2012, pp. 1–11, doi:10.1109/SC.2012.33.
- [23] W. MILLER, *Computational complexity and numerical stability*, SIAM J. Comput., 4 (1975), pp. 97–107, doi:10.1137/0204009.
- [24] A. V. SMIRNOV, *The bilinear complexity and practical algorithms for matrix multiplication*, Comput. Math. Math. Phys., 53 (2013), pp. 1781–1795, doi:10.1134/S0965542513120129.
- [25] V. STRASSEN, *Gaussian elimination is not optimal*, Numer. Math., 13 (1969), pp. 354–356, doi:10.1007/BF02165411.