# Communication-Optimal Parallel Algorithm for Strassen's Matrix Multiplication 

Grey Ballard ${ }^{*}$<br>EECS Department<br>UC Berkeley<br>Berkeley, CA 94720<br>ballard@eecs.berkeley.edu

James Demmel ${ }^{* \dagger}$<br>Mathematics Department<br>and CS Division UC Berkeley<br>Berkeley, CA 94720<br>demmel@cs.berkeley.edu

Olga Holtz ${ }^{\ddagger}$<br>Mathematics Departments<br>UC Berkeley<br>and TU Berlin<br>Berkeley, CA 94720<br>holtz@math.berkeley.edu

Benjamin Lipshitz*<br>EECS Department UC Berkeley<br>Berkeley, CA 94720<br>lipshitz@berkeley.edu

Oded Schwartz ${ }^{\S}$
EECS Department UC Berkeley
Berkeley, CA 94720
odedsc@eecs.berkeley.edu


#### Abstract

Parallel matrix multiplication is one of the most studied fundamental problems in distributed and high performance computing. We obtain a new parallel algorithm that is based on Strassen's fast matrix multiplication and minimizes communication. The algorithm outperforms all known parallel matrix multiplication algorithms, classical and Strassen-based, both asymptotically and in practice.

A critical bottleneck in parallelizing Strassen's algorithm is the communication between the processors. Ballard, Demmel, Holtz, and Schwartz (SPAA'11) prove lower bounds on these communication costs, using expansion properties of the underlying computation graph. Our algorithm matches these lower bounds, and so is communication-optimal. It exhibits perfect strong scaling within the maximum possible range.

Benchmarking our implementation on a Cray XT4, we obtain speedups over classical and Strassen-based algorithms


[^0]ranging from $24 \%$ to $184 \%$ for a fixed matrix dimension $n=94080$, where the number of nodes ranges from 49 to 7203.

Our parallelization approach generalizes to other fast matrix multiplication algorithms.
Categories and Subject Descriptors: F.2.1 [Analysis of Algorithms and Problem Complexity]: Numerical Algorithms and Problems: Computations on matrices
ACM General Terms: algorithms
Keywords: parallel algorithms, communication-avoiding algorithms, fast matrix multiplication

## 1. INTRODUCTION

Matrix multiplication is one of the most fundamental algorithmic problems in numerical linear algebra, distributed computing, scientific computing, and high-performance computing. Parallelization of matrix multiplication has been extensively studied (e.g., $9,6,11,1,25,20,34,10,23,30,4$, 3). It has been addressed using many theoretical approaches, algorithmic tools, and software engineering methods in order to optimize performance and obtain faster and more efficient parallel algorithms and implementations.

We obtain a new parallel algorithm based on Strassen's fast matrix multiplication ${ }^{1}$ It is more efficient than any other parallel matrix multiplication algorithm of which we are aware, including those that are based on classical $\left(\Theta\left(n^{3}\right)\right)$ multiplication, and those that are based on Strassen's and other Strassen-like matrix multiplications. We compare the efficiency of the new algorithm with previous algorithms, and provide both asymptotic analysis (Sections 3and 4) and benchmarking data (Section 5).

### 1.1 The communication bottleneck

To design efficient parallel algorithms, it is necessary not only to load balance the computation, but also to minimize the time spent communicating between processors. The inter-processor communication costs are in many cases significantly higher than the computational costs. Moreover, hardware trends predict that more problems will become communication-bound in the future [19, 18. Even matrix multiplication becomes communication-bound when run on sufficiently many processors. Given the importance of communication costs, it is preferable to match the performance of an algorithm to a communication lower bound, obtaining a communication-optimal algorithm.

### 1.2 Communication costs of matrix multiplication

We consider a distributed-memory parallel machine model as described in Section 2.1. The communication costs are measured as a function of the number of processors $P$, the local memory size $M$ in words, and the matrix dimension $n$. Irony, Toledo, and Tiskin [23] proved that in the distributed memory parallel model, the bandwidth cost of classical $n$-by- $n$ matrix multiplication is bounded by $\Omega\left(\frac{n^{3}}{P M^{1 / 2}}\right)$ words. Using their technique one can also deduce a memoryindependent bandwidth cost bound of $\Omega\left(\frac{n^{2}}{P^{2 / 3}}\right) 3$ and generalize it to other classes of algorithms 5. For a shared-memory model similar bounds were shown in [2]. Until recently, parallel classical matrix multiplication algorithms (e.g., "2D" 9 , 34, and "3D" 6, 1]) have minimized communication only for specific $M$ values. The first algorithm that minimizes the communication costs for the entire range of $M$ has recently been obtained by Solomonik and Demmel 30. See Section 4.1 for more details.

None of these lower bounding techniques and parallelization approaches generalize to fast matrix multiplication, such as $32,26,7,29,28,13,33,14,12,35$. A communication cost lower bound for fast matrix multiplication algorithms has only recently been obtained [4]: Strassen's algorithm run on a distributed-memory parallel machine has bandwidth

[^1]$\operatorname{cost} \Omega\left(\left(\frac{n}{M^{1 / 2}}\right)^{\omega_{0}} \cdot \frac{M}{P}\right)$ and latency $\operatorname{cost} \Omega\left(\left(\frac{n}{M^{1 / 2}}\right)^{\omega_{0}} \cdot \frac{1}{P}\right)$, where $\omega_{0}=\log _{2} 7$ (see Section 2.4). These bounds generalize to other, but not all, fast matrix multiplication algorithms, with $\omega_{0}$ being the exponent of the computational complexity.

In the sequential case ${ }^{2}$ the lower bounds are attained by the natural recursive implementation 32 which is thus optimal. However, a parallel communication-optimal Strassenbased algorithm was not previously known. Previous parallel algorithms that use Strassen (e.g., 20, 25, 17]), decrease the computational costs at the expense of higher communication costs. The factors by which these algorithms exceed the lower bounds are typically small powers of $P$ and $M$, as discussed in Section 4 However both $P$ and $M$ can be large (e.g. on a modern supercomputer, one may have $P \sim 10^{5}$ and $M \sim 10^{9}$ ).

### 1.3 Parallelizing Strassen's matrix multiplication in a communication efficient way

The main impetus for this work was the observation of the asymptotic gap between the communication costs of existing parallel Strassen-based algorithms and the communication lower bounds. Because of the attainability of the lower bounds in the sequential case, we hypothesized that the gap could be closed by finding a new algorithm rather than by tightening the lower bounds.

We made three observations from the lower bound results of 4 that lead to the new algorithm. First, the lower bounds for Strassen are lower than those for classical matrix multiplication. This implies that in order to obtain an optimal Strassen-based algorithm, the communication pattern for an optimal algorithm cannot be that of a classical algorithm but must reflect the properties of Strassen's algorithm. Second, the factor $M^{\omega_{0} / 2-1}$ that appears in the denominator of the communication cost lower bound implies that an optimal algorithm must use as much local memory as possible. That is, there is a tradeoff between memory usage and communication (the same is true in the classical case). Third, the proof of the lower bounds shows that in order to minimize communication costs relative to computation, it is necessary to perform each sub-matrix multiplication of size $\Theta(\sqrt{M}) \times \Theta(\sqrt{M})$ on a single processor.

With these observations and assisted by techniques from previous approaches to parallelizing Strassen, we developed a new parallel algorithm which achieves perfect load balance, minimizes communication costs, and in particular performs asymptotically less computation and communication than is possible using classical matrix multiplication.

### 1.4 Our contributions and paper organization

Our main contribution is a new algorithm we call Communication-Avoiding Parallel Strassen, or CAPS.

Theorem 1.1. CAPS asymptotically minimizes computational and bandwidth costs over all parallel Strassen-based algorithms. It also minimizes latency cost up to a logarithmic factor in the number of processors.

CAPS performs asymptotically better than any previous previous classical or Strassen-based parallel algorithm. It also runs faster in practice. The algorithm and its computational and communication cost analyses are presented in Section 3 There we show it matches the communication lower bounds.

[^2]We provide a review and analysis of previous algorithms in Section 4. We also consider two natural combinations of previously known algorithms (Sections 4.4 and 4.5). One of these new algorithms that we call "2.5D-Strassen" performs better than all previous algorithms, but is still not optimal, and performs worse than CAPS.

We discuss our implementations of the new algorithms and compare their performance with previous ones in Section 5 to show that our new CAPS algorithm outperforms previous algorithms not just asymptotically, but also in practice. Benchmarking our implementation on a Cray XT4, we obtain speedups over classical and Strassen-based algorithms ranging from $24 \%$ to $184 \%$ for a fixed matrix dimension $n=94080$, where the number of nodes ranges from 49 to 7203.

In Section 6 we show that our parallelization method applies to other fast matrix multiplication algorithms. It also applies to classical recursive matrix multiplication, thus obtaining a new optimal classical algorithm that matches the 2.5D algorithm of Solomonik and Demmel 30. In Section 6 we also discuss numerical stability, hardware scaling, and future work.

## 2. PRELIMINARIES

### 2.1 Communication model

We model communication of distributed-memory parallel architectures as follows. We assume the machine has $P$ processors, each with local memory of size $M$ words, which are connected via a network. Processors communicate via messages, and we assume that a message of $w$ words can be communicated in time $\alpha+\beta w$. The bandwidth cost of the algorithm is given by the word count and denoted by $B W(\cdot)$, and the latency cost is given by the message count and denoted by $L(\cdot)$. Similarly the computational cost is given by the number of floating point operations and denoted by $F(\cdot)$. We call the time per floating point operation $\gamma$.

We count the number of words, messages and floating point operations along the critical path as defined in 37. That is, two messages that are communicated between separate pairs of processors simultaneously are counted only once, as are two floating point operations performed in parallel on different processors. This metric is closely related to the total running time of the algorithm, which we model as

$$
\alpha L(\cdot)+\beta B W(\cdot)+\gamma F(\cdot)
$$

We assume that (1) the architecture is homogeneous (that is, $\gamma$ is the same on all processors and $\alpha$ and $\beta$ are the same between each pair of processors), (2) processors can send/receive only one message to/from one processor at a time and they cannot overlap computation with communication (this latter assumption can be dropped, affecting the running time by a factor of at most two), and (3) there is no communication resource contention among processors. That is, we assume that there is a link in the network between each pair of processors. Thus lower bounds derived in this model are valid for any network, but attainability of the lower bounds depends on the details of the network.

### 2.2 Strassen's algorithm

Strassen showed that $2 \times 2$ matrix multiplication can be performed using 7 multiplications and 18 additions, instead of the classical algorithm that does 8 multiplications and

4 additions [32]. By recursive application this yields an algorithm with multiplies two $n \times n$ matrices $O\left(n^{\omega_{0}}\right)$ flops, where $\omega_{0}=\log _{2} 7 \approx 2.81$. Winograd improved the algorithm to use 7 multiplications and 15 additions in the base case, thus decreasing the hidden constant in the $O$ notation 36 . We review the Strassen-Winograd algorithm in Appendix A

### 2.3 Previous work on parallel Strassen

In this section we breifly describe previous efforts to parallelize Strassen. More details, including communication analyses, are in Section 4 A summary appears in Table 1

Luo and Drake 25 explored Strassen-based parallel algorithms that use the communication patterns known for classical matrix multiplication. They considered using a classical 2D parallel algorithm and using Strassen locally, which corresponds to what we call the "2D-Strassen" approach (see Section 4.2). They also consider using Strassen at the highest level and performing a classical parallel algorithm for each subproblem generated, which corresponds to what we call the "Strassen-2D" approach. The size of the subproblems depends on the number of Strassen steps taken (see Section 4.3). Luo and Drake also analyzed the communication costs for these two approaches.

Soon after, Grayson, Shah, and van de Geijn 20 improved on the Strassen-2D approach of 25 by using a better classical parallel matrix multiplication algorithm and running on a more communication-efficient machine. They obtained better performance results compared to a purely classical algorithm for up to three levels of Strassen's recursion.

Kumar, Huang, Johnson, and Sadayappan 24 implemented Strassen's algorithm on a shared-memory machine. They identified the tradeoff between available parallelism and total memory footprint by differentiating between "partial" and "complete" evaluation of the algorithm, which corresponds to what we call depth-first and breadth-first traversal of the recursion tree (see Section 3.1). They show that by using $\ell$ DFS steps before using BFS steps, the memory footprint is reduced by a factor of $(7 / 4)^{\ell}$ compared to using all BFS steps. They did not consider communication costs in their work.

Other parallel approaches $17,22,31$ have used more complex parallel schemes and communication patterns. However, they restrict attention to only one or two steps of Strassen and obtain modest performance improvements over classical algorithms.

### 2.4 Strassen lower bounds

For Strassen-based algorithms, the bandwidth cost lower bound has been proved using expansion arguments on the computation graph, and the latency cost lower bound is an immediate corollary.

Theorem 2.1. (Memory-dependent lower bound) [4] Consider a Strassen-based algorithm running on $P$ processors each with local memory size $M$. Let $B W(n, P, M)$ be the bandwith cost and $L(n, P, M)$ be the latency cost of the algorithm. Assume that no intermediate values are computed twice. Then

$$
\begin{aligned}
B W(n, P, M) & =\Omega\left(\left(\frac{n}{\sqrt{M}}\right)^{\omega_{0}} \cdot \frac{M}{P}\right) \\
L(n, P, M) & =\Omega\left(\left(\frac{n}{\sqrt{M}}\right)^{\omega_{0}} \cdot \frac{1}{P}\right)
\end{aligned}
$$

A memory-independent lower bound has recently been proved using the same expansion approach:

Theorem 2.2. (Memory-independent lower bound) [3] Consider a Strassen-based algorithm running on $P$ processors. Let $B W(n, P)$ be the bandwith cost and $L(n, P)$ be the latency cost of the algorithm. Assume that no intermediate values are computed twice. Assume only one copy of the input data is stored at the start of the algorithm and the computation is load-balanced in an asymptotic sense. Then

$$
B W(n, P)=\Omega\left(\frac{n^{2}}{P^{2 / \omega_{0}}}\right),
$$

and the latency cost is $L(n, P)=\Omega(1)$.
Note that when $M=O\left(n^{2} / P^{2 / \omega_{0}}\right)$, the memorydependent lower bound is dominant, and when $M=$ $\Omega\left(n^{2} / P^{2 / \omega_{0}}\right)$, the memory-independent lower bound is dominant.

## 3. COMMUNICATION-AVOIDING PARALLEL STRASSEN

In this section we present the CAPS algorithm, and prove it is communication-optimal. See Algorithm 1 for a concise presentation and Algorithm 2 for a more detailed description.

Theorem 3.1. CAPS has computational cost $\Theta\left(\frac{n^{\omega_{0}}}{P}\right)$, bandwidth cost $\Theta\left(\max \left\{\frac{n^{\omega_{0}}}{P M^{\omega_{0} / 2-1}}, \frac{n^{2}}{P^{2 / \omega_{0}}}\right\}\right)$, and latency $\operatorname{cost} \Theta\left(\max \left\{\frac{n^{\omega_{0}}}{P M^{\omega_{0} / 2}} \log P, \log P\right\}\right)$.

By Theorems 2.1 and 2.2 we see that CAPS has optimal computational and bandwidth costs, and that its latency cost is at most $\log P$ away from optimal. Thus Theorem 1.1 follows. We prove Theorem 3.1 in Section 3.5

### 3.1 Overview of CAPS

Consider the recursion tree of Strassen's sequential algorithm. CAPS traverses it in parallel as follows. At each level of the tree, the algorithm proceeds in one of two ways. A "breadth-first-step" (BFS) divides the 7 subproblems among the processors, so that $\frac{1}{7}$ of the processors work on each subproblem independently and in parallel. A "depth-first-step" (DFS) uses all the processors on each subproblem, solving each one in sequence. See Figure 1.

In short, a BFS step requires more memory but reduces communication costs while a DFS step requires little extra memory but is less communication-efficient. In order to minimize communication costs, the algorithm must choose an ordering of BFS and DFS steps that uses as much memory as possible.

Let $k=\log _{7} P$ and $s \geq k$ be the number of distributed Strassen steps the algorithm will take. In this section, we assume that $n$ is a multiple of $2^{s} 7^{\lceil k / 2\rceil}$. If $k$ is even, the restriction simplifies to $n$ being a multiple of $2^{s} \sqrt{P}$. Since $P$ is a power of 7 , it is sometimes convenient to think of the processors as numbered in base 7. CAPS performs $s$ steps of Strassen's algorithm and finishes the calculation with local matrix multiplication. The algorithm can easily be generalized to other values of $n$ by padding or dynamic peeling.


Figure 1: Representation of BFS and DFS steps. In a BFS step, all seven subproblems are computed at once, each on $1 / 7$ of the processors. In a DFS step, the seven subproblems are computed in sequence, each using all the processors. The notation follows that of Appendix A.

We consider two simple schemes of traversing the recursion tree with BFS and DFS steps. The first scheme, which we call the Unlimited Memory ( $U M$ ) scheme, is to take $k$ BFS steps in a row. This approach is possible only if there is sufficient available memory. The second scheme, which we call the Limited Memory ( $L M$ ) scheme is to take $\ell$ DFS steps in a row followed by $k$ BFS steps in a row, where $\ell$ is minimized subject to the memory constraints.

It is possible to use a more complicated scheme that interleave BFS and DFS steps to reduce communication. We show that the LM scheme is optimal up to a constant factor, and hence no more than a constant factor improvement can be attained from interleaving.

```
Algorithm 1 CAPS, in brief. For more details, see Algo-
rithm 2.
Input: \(A, B, n\), where \(A\) and \(B\) are \(n \times n\) matrices
    \(P=\) number of processors
Output: \(C=A \cdot B\)
    procedure \(\mathrm{C}=\operatorname{CAPS}(A, B, n, P)\)
        if enough memory then \(\quad \triangleright\) Do a BFS step
                locally compute the \(S_{i}\) 's and \(T_{i}^{\prime} s\) from \(A\) and \(B\)
                parallel for \(i=1 \ldots 7\) do
                        redistribute \(S_{i}\) and \(T_{i}\)
                        \(Q_{i}=\operatorname{CAPS}\left(S_{i}, T_{i}, n / 2, P / 7\right)\)
                        redistribute \(Q_{i}\)
                locally compute \(C\) from all the \(Q_{i}\) 's
            else \(\quad \triangleright\) Do a DFS step
            for \(i=1 \ldots 7\) do
                        locally compute \(S_{i}\) and \(T_{i}\) from A and B
                        \(Q_{i}=\operatorname{CAPS}\left(S_{i}, T_{i}, n / 2, P\right)\)
                locally compute contribution of \(Q_{i}\) to \(C\)
    \(\triangleright\) The dependence of the \(S_{i}\) 's on \(A\), the \(T_{i}\) 's on \(B\) and \(C\)
    on the \(Q_{i}\) 's follows the Strassen or Strassen-Winograd
    algorithm. See Appendix A
```


### 3.2 Data layout

We require that the data layout of the matrices satisfies the following two properties:

1. At each of the $s$ Strassen recursion steps, the data layouts of the four sub-matrices of each of $A, B$, and $C$ must match so that the weighted additions of these sub-matrices can be performed locally. This technique follows 25 and allows communication-free DFS steps.
2. Each of these submatrices must be equally distributed among the $P$ processors for load balancing.

There are many data layouts that satisfy these properties, perhaps the simplest being block-cyclic layout with a processor grid of size $7^{\lfloor k / 2\rfloor} \times 7^{\lceil k / 2\rceil}$ and block size $\frac{n}{2^{s} 7^{\lfloor k / 2\rfloor}} \times \frac{n}{2^{s} 7^{〔 k / 2\rceil}}$. (When $k=\log _{7} P$ is even these expressions simplify to a processor grid of size $\sqrt{P} \times \sqrt{P}$ and block size $\frac{n}{2^{s} \sqrt{P}}$.) See Figure 2

Any layout that we use is specified by three parameters, ( $n, P, s$ ), and intermediate stages of the computation use the same layout with smaller values of the parameters. A BFS step reduces a multiplication problem with layout parameters $(n, P, s)$ to seven subproblems with layout parameters $(n / 2, P / 7, s-1)$. A DFS step reduces a multiplication problem with layout parameters $(n, P, s)$ to seven subproblems with layout parameters $(n / 2, P, s-1)$.

Note that if the input data is initially load-balanced but distributed using a different layout, we can rearrange it to the above layout using a total of $O\left(\frac{n^{2}}{P}\right)$ words and $O\left(n^{2}\right)$ messages. This has no asymptotic effect on the bandwidth cost but significantly increases the latency cost in the worst case.


Figure 2: An example matrix layout for CAPS. Each of the 16 submatrices as shown on the left has exactly the same layout. The colored blocks are the ones owned by processor 00 . On the right is a zoomed-in view of one submatrix, showing which processor, numbered base 7, owns each block. This is block-cyclic layout with some blocksize $b$, and matches our layout requirements with parameters ( $n=4 \cdot 49 \cdot b, P=49, s=2$ ).

### 3.3 Unlimited Memory scheme

In the UM scheme, we take $k=\log _{7} P$ BFS steps in a row. Since a BFS step reduces the number of processors involved in each subproblem by a factor of 7 , after $k$ BFS steps each subproblem is assigned to a single processor, and so is computed locally with no further communication costs. We first describe a BFS step in more detail.

The matrices $A$ and $B$ are initially distributed as described in Section 3.2 In order to take a recursive step, the 14 matrices $S_{1}, \ldots S_{7}, T_{1}, \ldots, T_{7}$ must be computed. Each processor allocates space for all 14 matrices and performs local additions and subtractions to compute its portion of the matrices. Recall that the submatrices are distributed identically, so this step requires no communication. If the layouts of $A$ and $B$ have parameters $(n, P, s)$, the $S_{i}$ and the $T_{i}$ now have layout parameters $(n / 2, P, s-1)$.

```
Algorithm 2 CAPS, in detail
Input: \(A, B\), are \(n \times n\) matrices
        \(P=\) number of processors
        rank \(=\) processor number base- 7 as an array
        \(M=\) local memory size
Output: \(C=A \cdot B\)
    procedure \(\mathrm{C}=\operatorname{CAPS}(A, B, P\), rank, \(M)\)
        \(\ell=\left\lceil\log _{2} \frac{4 n}{P^{1 / \omega_{0} M^{1 / 2}}}\right\rceil\)
            \(\triangleright \ell\) is number of DFS steps to fit in memory
        \(k=\log _{7} P\)
        call \(\operatorname{DFS}(A, B, C, k, \ell\), rank \()\)
    procedure \(\operatorname{DFS}(A, B, C, k, \ell\), rank)
            \(\triangleright\) Do \(C=A \cdot B\) by \(\ell\) DFS, then \(k\) BFS steps
        if \(\ell \leq 0\) then call \(\operatorname{BFS}(A, B, C, k\), rank); return
        for \(i=1 \ldots 7\) do
            locally compute \(S_{i}\) and \(T_{i}\) from \(A\) and \(B\)
                \(\triangleright\) following Strassen's algorithm
            call \(\operatorname{DFS}\left(S_{i}, T_{i}, Q_{i}, k, \ell-1\right.\), rank )
            locally compute contribution of \(Q_{i}\) to \(C\)
                                    \(\triangleright\) following Strassen's algorithm
    procedure \(\operatorname{BFS}(A, B, C, k\), rank \()\)
        \(\triangleright\) Do \(C=A \cdot B\) by k BFS steps, then local Strassen
        if \(k==0\) then call localStrassen \((A, B, C)\); return
        for \(i=1 \ldots 7\) do
            locally compute \(S_{i}\) and \(T_{i}\) from \(A\) and \(B\)
                \(\triangleright\) following Strassen's algorithm
        for \(i=1 \ldots 7\) do
            target \(=\) rank
            \(\operatorname{target}[k]=\mathrm{i}\)
            send \(S_{i}\) to target
            receive into \(L\)
            \(\triangleright\) One part of \(L\) comes from each of 7 processors
            send \(T_{i}\) to target
            receive into \(R\)
            \(\triangleright\) One part of \(R\) comes from each of 7 processors
        call \(\operatorname{BFS}(L, R, P, k-1\), rank )
        for \(i=1 \ldots 7\) do
            target \(=\) rank
            \(\operatorname{target}[k]=i\)
            send \(i^{\text {th }}\) part of \(P\) to target
            receive from target into \(Q_{i}\)
        locally compute \(C\) from \(Q_{i}\)
            \(\triangleright\) following Strassen's algorithm
```

The next step is to redistribute these 14 matrices so that the 7 pairs of matrices $\left(S_{i}, T_{i}\right)$ exist on disjoint sets of $P / 7$ processors. This requires disjoint sets of 7 processors performing an all-to-all communication step (each processor must send and receieve a message from each of the other 6). To see this, consider the numbering of the processors base-7. On the $m^{\text {th }}$ BFS step, the communication is between the seven processors whose numbers agree on all digits except the $m^{\text {th }}$ (counting from the right). After the $m^{\text {th }}$ BFS step, the set of processors working on a given subproblem share the same $m$-digit suffix. After the above communication is performed, the layout of $S_{i}$ and $T_{i}$ has parameters ( $n / 2, P / 7, s-1$ ), and the sets of processors that own the $T_{i}$ and $S_{i}$ are disjoint
for different values of $i$. Note that since each all-to-all only involves seven processors no matter how large $P$ is, this algorithm does not have the scalability issues that typically come from an all-to-all communication pattern.

### 3.3.1 Memory requirements

The extra memory required to take one BFS step is the space to store all 7 triples $S_{j}, T_{j}, Q_{j}$. Since each of those matrices is $\frac{1}{4}$ the size of $A, B$, and $C$, the extra space required at a given step is $7 / 4$ the extra space required for the previous step. We assume that no extra memory is required for the local multiplications $3^{3}$ Thus, the total local memory requirement for taking $k$ BFS steps is given by

$$
\begin{aligned}
\operatorname{Mem}_{\mathrm{UM}}(n, P) & =\frac{3 n^{2}}{P} \sum_{i=0}^{k}\left(\frac{7}{4}\right)^{i}=\frac{7 n^{2}}{P^{2 / \omega_{0}}}-\frac{4 n^{2}}{P} \\
& =\Theta\left(\frac{n^{2}}{P^{2 / \omega_{0}}}\right)
\end{aligned}
$$

### 3.3.2 Computational costs

The computation required at a given BFS step is that of the local additions and subtractions associated with computing the $S_{i}$ and $T_{i}$ and updating the output matrix $C$ with the $Q_{i}$. Since Strassen performs 18 additions and subtractions, the computational cost recurrence is

$$
F_{\mathrm{UM}}(n, P)=18\left(\frac{n^{2}}{4 P}\right)+F_{\mathrm{UM}}\left(\frac{n}{2}, \frac{P}{7}\right)
$$

with base case $F_{\mathrm{UM}}(n, 1)=c_{s} n^{\omega_{0}}-6 n^{2}$, where $c_{s}$ is the constant of Strassen's algorithm. See Appendix A for more details. The solution to this recurrence is

$$
F_{\mathrm{UM}}(n, P)=\frac{c_{s} n^{\omega_{0}}-6 n^{2}}{P}=\Theta\left(\frac{n^{\omega_{0}}}{P}\right) .
$$

### 3.3.3 Communication costs

Consider the communication costs associated with the UM scheme. Given that the redistribution within a BFS step is performed by an all-to-all communication step among sets of 7 processors, each processor sends 6 messages and receives 6 messages to redistribute $S_{1}, \ldots, S_{7}$, and the same for $T_{1}, \ldots, T_{7}$. After the products $Q_{i}=S_{i} T_{i}$ are computed, each processor sends 6 messages and receive 6 messages to redistribute $Q_{1}, \ldots, Q_{7}$. The size of each message varies according to the recursion depth, and is the number of words a processor owns of any $S_{i}, T_{i}$, or $Q_{i}$, namely $\frac{n^{2}}{4 P}$ words.

As each of the $Q_{i}$ is computed simultaneously on disjoint sets of $P / 7$ processors, we obtain a cost recurrence for the entire UM scheme:

$$
\begin{aligned}
B W_{\mathrm{UM}}(n, P) & =36 \frac{n^{2}}{4 P}+B W_{\mathrm{UM}}\left(\frac{n}{2}, \frac{P}{7}\right) \\
L_{\mathrm{UM}}(n, P) & =36+L_{\mathrm{UM}}\left(\frac{n}{2}, \frac{P}{7}\right)
\end{aligned}
$$

with base case $L_{\mathrm{UM}}(n, 1)=B W_{\mathrm{UM}}(n, 1)=0$. Thus

$$
\begin{align*}
B W_{\mathrm{UM}}(n, P) & =\frac{12 n^{2}}{P^{2 / \omega_{0}}}-\frac{12 n^{2}}{P}=\Theta\left(\frac{n^{2}}{P^{2 / \omega_{0}}}\right) \\
L_{\mathrm{UM}}(n, P) & =36 \log _{7} P=\Theta(\log P) \tag{1}
\end{align*}
$$

[^3]
### 3.4 Limited Memory scheme

In this section we discuss a scheme for traversing Strassen's recursion tree in the context of limited memory. In the LM scheme, we take $\ell$ DFS steps in a row followed by $k$ BFS steps in a row, where $\ell$ is minimized subject to the memory constraints. That is, we use a sequence of DFS steps to reduce the problem size so that we can use the UM scheme on each subproblem without exceeding the available memory.

Consider taking a single DFS step. Rather than allocating space for and computing all 14 matrices $S_{1}, T_{1}, \ldots, S_{7}, T_{7}$ at once, the DFS step requires allocation of only one subproblem, and each of the $Q_{i}$ will be computed in sequence.

Consider the $i^{\text {th }}$ subproblem: as before, both $S_{i}$ and $T_{i}$ can be computed locally. After $Q_{i}$ is computed, it is used to update the corresponding quadrants of $C$ and then discarded so that its space in memory (as well as the space for $S_{i}$ and $T_{i}$ ) can be re-used for the next subproblem. In a DFS step, no redistribution occurs. After $S_{i}$ and $T_{i}$ are computed, all processors participate in the computation of $Q_{i}$.

We assume that some extra memory is available. To be precise, assume the matrices $A, B$, and $C$ require only $\frac{1}{3}$ of the available memory:

$$
\begin{equation*}
\frac{3 n^{2}}{P} \leq \frac{1}{3} M \tag{2}
\end{equation*}
$$

In the LM scheme, we set

$$
\begin{equation*}
\ell=\max \left\{0,\left\lceil\log _{2} \frac{4 n}{P^{1 / \omega_{0} M^{1 / 2}}}\right\rceil\right\} . \tag{3}
\end{equation*}
$$

The following subsection shows that this choice of $\ell$ is sufficient not to exceed the memory capacity.

### 3.4.1 Memory requirements

The extra memory requirement for a DFS step is the space to store one subproblem. Thus, the extra space required at this step is $1 / 4$ the space required to store $A, B$, and $C$. The local memory requirements for the LM scheme is given by

$$
\begin{aligned}
\operatorname{Mem}_{\mathrm{LM}}(n, P) & =\frac{3 n^{2}}{P} \sum_{i=0}^{\ell-1}\left(\frac{1}{4}\right)^{i}+\operatorname{Mem}_{\mathrm{UM}}\left(\frac{n}{2^{\ell}}, P\right) \\
& \leq \frac{M}{3} \sum_{i=0}^{\ell-1}\left(\frac{1}{4}\right)^{i}+\frac{7\left(\frac{n}{2^{\ell}}\right)^{2}}{P^{2 / \omega_{0}}} \\
& \leq \frac{127}{144} M<M
\end{aligned}
$$

where the last line follows from (3) and (2). Thus, the limited memory scheme does not exceed the available memory.

### 3.4.2 Computational costs

As in the UM case, the computation required at a given DFS step is that of the local additions and subtractions associated with computing the $S_{i}$ and $T_{i}$ and updating the output matrix $C$ with the $Q_{i}$. However, since all processors participate in each subproblem and the subproblems are computed in sequence, the recurrence is given by

$$
F_{\mathrm{LM}}(n, P)=18\left(\frac{n^{2}}{4 P}\right)+7 \cdot F_{\mathrm{LM}}\left(\frac{n}{2}, P\right)
$$

After $\ell$ steps of DFS, the size of a subproblems is $\frac{n}{2^{\ell}} \times \frac{n}{2^{\ell}}$, and there are $P$ processors involved. We take $k$ BFS steps to compute each of these $7^{\ell}$ subproblems. Thus

$$
F_{\mathrm{LM}}\left(\frac{n}{2^{\ell}}, P\right)=F_{\mathrm{UM}}\left(\frac{n}{2^{\ell}}, P\right),
$$

and

$$
\begin{aligned}
F_{\mathrm{LM}}(n, P) & =\frac{18 n^{2}}{4 P} \sum_{i=0}^{\ell-1}\left(\frac{7}{4}\right)^{i}+7^{\ell} \cdot F_{\mathrm{UM}}\left(\frac{n}{2^{\ell}}, P\right) \\
& =\frac{c_{s} n^{\omega_{0}}-6 n^{2}}{P}=\Theta\left(\frac{n_{0}^{\omega}}{P}\right)
\end{aligned}
$$

### 3.4.3 Communication costs

Since there are no communication costs associated with a DFS step, the recurrence is simply

$$
\begin{aligned}
B W_{\mathrm{LM}}(n, P) & =7 \cdot B W_{\mathrm{LM}}\left(\frac{n}{2}, P\right) \\
L_{\mathrm{LM}}(n, P) & =7 \cdot L_{\mathrm{LM}}\left(\frac{n}{2}, P\right)
\end{aligned}
$$

with base cases

$$
\begin{aligned}
B W_{\mathrm{LM}}\left(\frac{n}{2^{\ell}}, P\right) & =B W_{\mathrm{UM}}\left(\frac{n}{2^{\ell}}, P\right) \\
L_{\mathrm{LM}}\left(\frac{n}{2^{\ell}}, P\right) & =L_{\mathrm{UM}}\left(\frac{n}{2^{\ell}}, P\right) .
\end{aligned}
$$

Thus the total communication costs are given by

$$
\begin{align*}
B W_{\mathrm{LM}}(n, P) & =7^{\ell} \cdot B W_{\mathrm{UM}}\left(\frac{n}{2^{\ell}}, P\right) \\
& \leq \frac{12 \cdot 4^{\omega_{0}-2} n^{\omega_{0}}}{P M^{\omega_{0} / 2-1}} \\
& =\Theta\left(\frac{n^{\omega_{0}}}{P M^{\omega_{0} / 2-1}}\right) . \\
L_{\mathrm{LM}}(n, P) & =7^{\ell} \cdot L_{\mathrm{UM}}\left(\frac{n}{2^{\ell}}, P\right) \\
& \leq \frac{(4 n)^{\omega_{0}}}{P M^{\omega_{0} / 2}} 36 \log _{7} P \\
& =\Theta\left(\frac{n^{\omega_{0}}}{P M^{\omega_{0} / 2}} \log P\right) . \tag{4}
\end{align*}
$$

### 3.5 Communication optimality

Proof. (of Theorem 3.1). In the case that $M \geq$ $\operatorname{Mem}_{\mathrm{UM}}(n, P)=\Omega\left(\frac{n^{2}}{P^{2 / \omega_{0}}}\right)$ the UM scheme is possible. Then the communication costs are given by (1) which matches the lower bound of Theorem 2.2. Thus the UM scheme is communication-optimal (up to a logarithmic factor in the latency cost and assuming that the data is initially distributed as described in Section 3.2. For smaller values of $M$, the LM scheme must be used. Then the communication costs are given by $(4)$ and match the lower bound of Theorem 2.1 so the LM scheme is also communication-optimal.

We note that for the LM scheme, since both the computational and communication costs are proportional to $\frac{1}{P}$, we can expect perfect strong scaling: given a fixed problem size, increasing the number of processors by some factor will decrease each cost by the same factor. However, this strong scaling property has a limited range. As $P$ increases, holding everything else constant, the global memory size $P M$ increases as well. The limit of perfect strong scaling is exactly when there is enough memory for the UM scheme. See 3] for details.

## 4. ANALYSIS OF OTHER ALGORITHMS

In the section we detail the asymptotic communication costs of other matrix multiplication algorithms, both classical
and Strassen-based. These communication costs and the corresponding lower bounds are summarized in Table 1

Many of the algorithms described in this section are hybrids of two different algorithms. We use the convention that the names of the hybrid algorithms are composed of the names of the two component algorithms, hyphenated. The first name describes the algorithm used at the top level, on the largest problems, and the second describes the algorithm used at the base level on smaller problems.

### 4.1 Classical Algorithms

Classical algorithms must communicate asymptotically more than an optimal Strassen-based algorithm. To compare the lower bounds, it is necessary to consider three cases for the memory size: when the memory-dependent bounds dominate for both classical and Strassen, when the memorydependent bound dominates for classical, but the memoryindependent bound dominates for Strassen, and when the memory-independent bounds dominate for both classical and Strassen. This analysis is detailed in Appendix B Briefly, the factor by which the classical bandwidth cost exceeds the Strassen bandwidth cost is $P^{a}$ where $a$ ranges from $\frac{2}{\omega_{0}}-\frac{2}{3} \approx 0.046$ to $\frac{3-\omega_{0}}{2} \approx 0.10$ depending on the relative problem size. The same sort of analysis is used throughout Section 4 to compare each algorithm with the Strassen-based lower bounds.

Various parallel classical matrix multiplication algorithms minimize communication relative to the classical lower bounds for certain amounts of local memory $M$. For example, Cannon's algorithm 9 minimizes communication for $M=O\left(n^{2} / P\right)$. Several more practical algorithms exist (such as SUMMA 34) which use the same amount of local memory and have the same asymptotic communication costs. We call this class of algorithms " 2 D " because the communication patterns follow a two-dimensional processor grid.

Another class of algorithms, known as "3D" 6, 1] because the communication pattern maps to a three-dimensional processor grid, uses more local memory and reduces communication relative to 2 D algorithms. This class of algorithms minimizes communication relative to the classical lower bounds for $M=\Omega\left(n^{2} / P^{2 / 3}\right)$. As shown in 3 , it is not possible to use more memory than $M=\Theta\left(n^{2} / \frac{3}{P^{2 / 3}}\right)$ to reduce communication.

Recently, a more general algorithm has been developed which minimizes communication in all cases. Because it reduces to a 2D and 3D for the extreme values of $M$ but interpolates for the values between, it is known as the " 2.5 D " algorithm 30.

### 4.2 2D-Strassen

One idea to parallelize Strassen-based algorithms is to use a 2D classical algorithm for the inter-processor communication, and use the fast matrix multiplication algorithm locally 25. We call such an algorithm "2D-Strassen". It is straightforward to implement, but cannot attain all the computational speedup from Strassen since it uses a classical algorithm for part of the computation. In particular, it does not use Strassen for the largest matrices, when Strassen provides the greatest reduction in computation. As a result, the computational cost exceeds $\Theta\left(n^{\omega_{0}} / P\right)$ by a factor of $P^{\left(3-\omega_{0}\right) / 2} \approx P^{0.10}$. The 2D-Strassen algorithm has the same communication cost as 2D algorithms, and hence does not match the communication costs of CAPS. In comparing

|  |  | Flops | Bandwidth | Latency |
| :---: | :---: | :---: | :---: | :---: |
|  | Lower Bound 23 | $\frac{n^{3}}{P}$ | $\max \left\{\frac{n^{3}}{P M^{1 / 2}}, \frac{n^{2}}{P^{2 / 3}}\right\}$ | $\max \left\{\frac{n^{3}}{P M^{3 / 2}}, 1\right\}$ |
|  | 2D 9,34 | $\frac{n^{3}}{P}$ | $\frac{n^{2}}{P^{1 / 2}}$ | $P^{1 / 2}$ |
|  | $3 \mathrm{D}, 1,6$ | $\frac{n^{3}}{P}$ | $\frac{n^{2}}{P^{2 / 3}}$ | $\log P$ |
|  | 2.5 D (optimal) 30 | $\frac{n^{3}}{P}$ | $\max \left\{\frac{n^{3}}{P M^{1 / 2}}, \frac{n^{2}}{P^{2 / 3}}\right\}$ | $\frac{n^{3}}{P M^{3 / 2}}+\log P$ |
|  | Lower Bound 3, 4 | $\frac{n^{\omega_{0}}}{P}$ | $\max \left\{\frac{n^{\omega_{0}}}{P M^{\omega_{0} / 2-1}}, \frac{n^{2}}{P^{2 / \omega_{0}}}\right\}$ | $\max \left\{\frac{n^{\omega_{0}}}{P M^{\omega_{0} / 2}}, 1\right\}$ |
|  | 2D-Strassen 25 | $\frac{n^{\omega_{0}}}{P^{\left(\omega_{0}-1\right) / 2}}$ | $\frac{n^{2}}{P^{1 / 2}}$ | $P^{1 / 2}$ |
|  | Strassen-2D 20, 25 | $\left(\frac{7}{8}\right)^{\ell} \frac{n^{3}}{P}$ | $\left(\frac{7}{4}\right)^{\ell} \frac{n^{2}}{P^{1 / 2}}$ | $7^{\ell} P^{1 / 2}$ |
|  | 2.5D-Strassen | $\max \left\{\frac{n^{3}}{P M^{3 / 2-\omega_{0} / 2}}, \frac{n^{\omega_{0}}}{P \omega_{0} / 3}\right\}$ | $\max \left\{\frac{n^{3}}{P M^{1 / 2}}, \frac{n^{2}}{P^{2 / 3}}\right\}$ | $\frac{n^{3}}{P M^{3 / 2}}+\log P$ |
|  | Strassen-2.5D | $\left(\frac{7}{8}\right)^{\ell} \frac{n^{3}}{P}$ | $\max \left\{\left(\frac{7}{8}\right)^{\ell} \frac{n^{3}}{P M^{1 / 2}},\left(\frac{7}{4}\right)^{\ell} \frac{n^{2}}{P^{2 / 3}}\right\}$ | $\left(\frac{7}{8}\right)^{\ell} \frac{n^{3}}{P M^{3 / 2}}+7^{\ell} \log P$ |
|  | CAPS (optimal) | $\frac{n^{\omega_{0}}}{P}$ | $\max \left\{\frac{n^{\omega_{0}}}{P M^{\omega_{0} / 2-1}}, \frac{n^{2}}{P^{2 / \omega_{0}}}\right\}$ | $\max \left\{\frac{n^{\omega_{0}}}{P M^{\omega_{0} / 2}} \log P, \log P\right\}$ |

Table 1: Asymptotic matrix multiplication computational and communication costs of algorithms and corresponding lower bounds. Here $\omega_{0}=\log _{2} 7 \approx 2.81$ is the exponent of Strassen; $\ell$ is the number of Strassen steps taken. None of the Strassen-based algorithms except for CAPS attain the lower bounds of Section 2.4 , see Section 4 for a discussion of each.
the 2D-Strassen bandwidth cost, $\Theta\left(n^{2} / P^{1 / 2}\right)$, to the CAPS bandwidth cost in Section 3 note that for the problem to fit in memory we always have $M=\Omega\left(n^{2} / P\right)$. The bandwidth cost exceeds that of CAPS by a factor of $P^{a}$, where $a$ ranges from $\left(3-\omega_{0}\right) / 2 \approx .10$ to $2 / \omega_{0}-1 / 2 \approx .21$, depending on the relative problem size. Similarly, the latency cost, $\Theta\left(P^{1 / 2}\right)$, exceeds that of CAPS by a factor of $P^{a}$ where $a$ ranges from $\left(3-\omega_{0}\right) / 2 \approx .10$ to $1 / 2=.5$.

### 4.3 Strassen-2D

The "Strassen-2D" algorithm applies $\ell$ DFS steps of Strassen's algorithm at the top level, and performs the $7^{\ell}$ smaller matrix multiplications using a 2D algorithm. By choosing certain data layouts as in Section 3.2, it is possible to do the additions and subtractions for Strassen's algorithm without any communication 25. However, Strassen-2D is also unable to match the communication costs of CAPS. Moreover, the speedup of Strassen-2D in computation comes at the expense of extra communication. For large numbers of Strassen steps $\ell$, Strassen-2D can approach the computational lower bound of Strassen, but each step increases the bandwidth cost by a factor of $\frac{7}{4}$ and the latency cost by a factor of 7 . Thus the bandwidth cost of Strassen-2D is a factor of $\left(\frac{7}{4}\right)^{l}$ higher than 2D-Strassen, which is already higher than that of CAPS. The latency cost is even worse: Strassen-2D is a factor of $7^{\ell}$ higher than 2D-Strassen.

One can reduce the latency cost of Strassen-2D at the expense of a larger memory footprint. Since Strassen-2D runs a 2 D algorithm $7^{\ell}$ times on the same set of processors, it is possible to pack together messages from independent matrix multiplications. In the best case, the latency cost is reduced to the cost of 2D-Strassen, which is still above that of CAPS, at the expense of using a factor of $\left(\frac{7}{4}\right)^{\ell}$ more memory.

### 4.4 2.5D-Strassen

A natural idea is to replace a 2D classical algorithm in 2D-Strassen with the superior 2.5D classical algorithm to obtain an algorithm we call 2.5D-Strassen. This algorithm uses the 2.5 D algorithm for the inter-processor communication, and then uses Strassen for the local computation. When $M=\Theta\left(n^{2} / P\right), 2.5$ D-Strassen is exactly the same as 2D-Strassen, but when there is extra memory it both decreases the communication cost and decreases the computational cost since the local matrix multiplications are performed (using Strassen) on larger matrices. To be precise, the computational cost exceeds the lower bound by a factor of $P^{a}$ where $a$ ranges from $1-\frac{\omega_{0}}{3} \approx 0.064$ to $\frac{3-\omega_{0}}{2} \approx 0.10$ depending on the relative problem size. The bandwidth cost exceeds the bandwidth cost of CAPS by a factor of $P^{a}$ where $a$ ranges from $\frac{2}{\omega_{0}}-\frac{2}{3} \approx 0.046$ to $\frac{3-\omega_{0}}{2} \approx 0.10$. In terms of latency, the cost of $\frac{n^{3}}{P M^{3 / 2}}+\log P$ exceeds the latency cost of CAPS by a factor ranging from $\log P$ to $P^{\left(3-\omega_{0}\right) / 2} \approx P^{0.10}$, depending on the relative problem size.

### 4.5 Strassen-2.5D

Similarly, by replacing a 2D algorithm with 2.5D in Strassen-2D, one obtains the new algorithm we call Strassen2.5D. First one takes $\ell$ DFS steps of Strassen, which can be done without communication, and then one applies the 2.5D algorithm to each of the $7^{\ell}$ subproblems. The computational cost is exactly the same as Strassen-2D, but the communication cost will typically be lower. Each of the $7^{\ell}$ subproblems is multiplication of $n / 2^{\ell} \times n / 2^{\ell}$ matrices. Each subproblem uses only $1 / 4^{\ell}$ as much memory as the original problem. Thus there may be a large amount of extra memory available for each subproblem, and the lower communication costs of the 2.5 D algorithm help. The choice of $\ell$ that minimizes the
bandwidth cost is

$$
\ell_{\mathrm{opt}}=\max \left\{0,\left\lceil\log _{2} \frac{n}{M^{1 / 2} P^{1 / 3}}\right\rceil\right\} .
$$

The same choice minimizes the latency cost. Note that when $M \geq \frac{n^{2}}{P^{2 / 3}}$, taking zero Strassen steps minimizes the communication within the constraints of the Strassen-2.5D algorithm. With $\ell=\ell_{\text {opt }}$, the bandwidth cost is a factor of $P^{1-\omega_{0} / 3} \approx P^{0.064}$ above that of CAPS. Additionally, the computational cost is not optimal, and using $\ell=\ell_{\mathrm{opt}}$, the computational cost exceeds the optimal by a factor of $P^{1-\omega_{0} / 3} M^{3 / 2-\omega_{0} / 2} \approx P^{0.064} M^{0.096}$.

It is also possible to take $\ell>\ell_{\mathrm{opt}}$ steps of Strassen to decrease the comptutational cost further. However the decreased computational cost comes at the expense of higher communication cost, as in the case of Strassen-2D. In particular, each extra step over $\ell_{\text {opt }}$ increases the bandwidth cost by a factor of $\frac{7}{4}$ and the latency cost by a factor of 7 . As with Strassen-2D, it is possible to use extra memory to pack together messages from several subproblems and decrease the latency cost, but not the bandwidth cost.

## 5. PERFORMANCE RESULTS

We have implemented CAPS using MPI on a Cray XT4, and compared it to various previous classical and Strassenbased algorithms. The benchmarking data is shown in Figure 3

### 5.1 Experimental setup

The nodes of the Cray XT4 have 8GB of memory and a quad-core AMD "Bupdapest" processor running at 2.3 GHz . We treat the entire node as a single processor, and when we use the classical algorithm we call the optimized threaded BLAS in Cray's LibSci to provide parallelism between the four cores in a node. The peak flop rate is 9.2 GFLOPS per core, or 36.8 GFLOPS per node. The machine consists of 9,572 nodes. All the data in Figure 3 is for multiplying two square matrices with $n=94080$.

### 5.2 Performance

Note that the vertical scale of Figure 3 is "effective GFLOPS", which is a useful measure for comparing classical and fast matrix multiplication algorithms. It is calculated as

$$
\begin{equation*}
\text { Effective GFLOPS }=\frac{2 n^{3}}{(\text { Execution time in seconds }) 10^{9}} \tag{5}
\end{equation*}
$$

For classical algorithms, which perform $2 n^{3}$ floating point operations, this gives the actual GFLOPS. For fast matrix multiplication algorithms it gives the performance relative to classical algorithms, but does not accurately represent the number of floating point operations performed.

Our algorithm outperforms all previous algorithms, and attains performance as high as 49.1 effective GFLOPS/node, which is $33 \%$ above the theoretical maximum for all classical algorithms. Compared with the best classical implementation, our speedup ranges from $51 \%$ for small values of $P$ up to $94 \%$ when using most of the machine. Compared with the best previous parallel Strassen algorithms, our speedup ranges from $24 \%$ up to $184 \%$. Unlike previous Strassen algorithms, we are able to attain substantial speedups over the entire range of processors.

### 5.3 Strong scaling

Figure 3 is a strong scaling plot: the problem size is fixed and each algorithm is run with $P$ ranging from the minimum that provides enough memory up to the largest allowed value of $P$ smaller than the size of the machine. Perfect strong scaling corresponds to a horizontal line in the plot. As the communication analysis predicts, CAPS exhibits better strong scaling than any of the other algorithms (with the exception of ScaLAPACK, which obtains very good strong scaling by having poor performance for small values of $P$ ).

### 5.4 Details of the implementations

### 5.4.1 CAPS

This implementation is the CAPS algorithm, with a few modifications from the presentation in Section 3 First, when computing locally it switches to classical matrix multiplication below some size $n_{0}$. Second, it is generalized to run on $P=c 7^{k}$ processors for $c \in\{1,2,3,6\}$ rather than just $7^{k}$ processors. As a result, the base-case classical matrix multiplication is done on $c$ processors rather than 1. Finally, implementation uses the Winograd variant of Strassen; see Appendix A for more details. Every point in the plot is tuned to use the best interleaving pattern of BFS and DFS steps, and the best total number of Strassen steps. For points in the figure, the optimal total number of Strassen steps is always 5 or 6 .

### 5.4.2 ScaLAPACK

We use ScaLAPACK 8 as optimized by Cray in LibSci. This is an implementation of the SUMMA algorithm, and can run on an arbitrary number of processors. It should give the best performance if $P$ is a perfect square so the processors can be placed on a square 2D grid. All the runs shown in Figure 3 are with $P$ a perfect square.

### 5.4.3 2.5D classical

This is the code of 30 . It places the $P$ processors in a grid of size $\sqrt{P / c} \times \sqrt{P / c} \times c$, and requires that $\sqrt{P / c}$ and $c$ are integers with $1 \leq c \leq P^{1 / 3}$, and $c$ divides $\sqrt{P / c}$. Additionally, it gets the best performance if $c$ is as large as possible, given the constraint that $c$ copies of the input and output matrices fit in memory. In the case that $c=1$ this code is an optimized implementation of SUMMA. The values of $P$ and $c$ for the runs in Figure 3 are chosen to get the best performance.

### 5.4.4 Strassen-2D

Following the algorithm of 20,25 , this implementation uses the DFS code from the implementation of CAPS at the top level, and then uses the optimized SUMMA code from the 2.5 D implementation with $c=1$. Since the classical code requires that $P$ is a perfect square, this requirement applies here as well. The number of Strassen steps taken is tuned to give the best performance for each $P$ value, and the optimal number varies from 0 to 2 .

### 5.4.5 2D-Strassen

Following the algorithm of 25], the 2D-Strassen implementation is analagous to the Strassen-2D implementation, but with the classical algorithm run before taking local Strassen steps. Similarly the same code is used for local Strassen steps here and in our implementation of CAPS. This code also


Figure 3: Strong scaling performance of various matrix multiplication algorithms on Cray XT4 for fixed problem size $n=94080$. The top line is CAPS as described in Section 3, and substantially outperforms all the other classical and Strassen-based algorithms. The horizontal axis is the number of nodes in log-scale. The vertical axis is effective GFLOPS, which are a performance measure rather than a flop rate, as discussed in Section 5.2. See Section 5.4 for a description of each implementation.
requires that $P$ is a perfect square. The number of Strassen steps is tuned for each $P$ value, and the optimal number varies from 0 to 3 .

### 5.4.6 2.5D-Strassen

This implementation uses the 2.5D implementation to reduce the problem to one processor, then takes several Strassen steps. The processor requirements are the same as for the 2.5D implementation. The number of Strassen steps is tuned for each number of processors, and the optimal number varies from 0 to 3 . We also tested the Strassen2.5 D algorithm, but its performance was always lower than $2.5 \mathrm{D}-$ Strassen in our experiments.

## 6. CONCLUSIONS/FUTURE WORK

### 6.1 Stability of fast matrix multiplication

CAPS has the same stability properties as sequential versions of Strassen. For a complete discussion of the stability of fast matrix multiplication algorithms, see 21,16 . We highlight a few main points here. The tightest error bounds for classical matrix multiplication are componentwise: $|C-\hat{C}| \leq n \epsilon|A| \cdot|B|$, where $\hat{C}$ is the computed result and
$\epsilon$ is the machine precision. Strassen and other fast algorithms do not satisfy component-wise bounds but do satisfy the slightly weaker norm-wise bounds: $\|C-\hat{C}\| \leq f(n) \epsilon\|A\|\|B\|$, where $\|A\|=\max _{i, j} A_{i j}$ and $f$ is polynomial in $n$ 21. Accuracy can be improved with the use of diagonal scaling matrices: $D_{1} C D_{3}=D_{1} A D_{2} \cdot D_{2}^{-1} B D_{3}$. It is possible to choose $D_{1}, D_{2}, D_{3}$ so that the error bounds satisfy either $\left|C_{i j}-\hat{C}_{i j}\right| \leq f(n) \epsilon\|A(i,:)\|\|B(:, j)\|$ or $\|C-\hat{C}\| \leq$ $f(n) \epsilon\||A| \cdot|B|\|$. By scaling, the error bounds on Strassen become comparable to those of many other dense linear algebra algorithms, such as LU and QR decomposition 15 . Thus using Strassen for the matrix multiplications in a larger computation will often not harm the stability at all.

### 6.2 Hardware scaling

Although Strassen performs asymptotically less computation and communication than classical matrix multiplication, it is more demanding on the hardware. That is, if one wants to run matrix multiplication near the peak CPU speed, Strassen is more demanding of the memory size and communication bandwidth. This is because the ratio of computational cost to bandwidth cost is lower for Strassen than for classical. From the lower bounds in Section 2.4 the asymptotic
ratio of computational cost to bandwidth cost is $M^{\omega_{0} / 2-1}$ for Strassen-based algorithms, versus $M^{1 / 2}$ for classical algorithms. This means that it is harder to run Strassen near peak than it is to run classical matrix multiplication near peak. In terms of the machine parameters $\beta$ and $\gamma$ introduced in Section 2.1, the condition to be able to be computationbound is $\gamma M^{1 / 2} \geq c \beta$ for classical matrix multiplication and $\gamma M^{\omega_{0} / 2-1} \geq c^{\prime} \beta$ for Strassen. Here $c$ and $c^{\prime}$ are constants that depend on the constants in the communication and computational costs of classical and Strassen-based matrix multiplication.

The above inequalities may guide hardware design as long as classical and Strassen matrix multiplication are considered important computations. They apply both to the distributed case, where $M$ is the local memory size and $\beta$ is the inverse network bandwidth, and to the sequential/shared-memory case where $M$ is the cache size and $\beta$ is the inverse memorycache bandwidth.

### 6.3 Optimizing on-node performance

Note that although our implementation performs above the classical peak performance, it performs well below the corresponding Strassen-Winograd peak, defined by the time it takes to perform $c_{s} n^{\omega_{0}} / P$ flops at the peak speed of each processor. To some extent, this is because Strassen is more demanding on the hardware, as noted above. However we have not yet analyzed whether the amount our performance is below Strassen peak can be entirely accounted for based on machine parameters. It is also possible that a high performance shared-memory Strassen implementation might provide substantial speedups for our implementation, as well as for $2 \mathrm{D}-$ Strassen and $2.5 \mathrm{D}-$ Strassen.

### 6.4 Testing on various architectures

We have implemented and benchmarked CAPS on only one architecture, a Cray XT4. It remains to check that it outperforms other matrix multiplication algorithms on a variety of architectures. On some architectures it may be more important to consider the topology of the network and redesign the algorithm to minimize contention, which we have not done.

### 6.5 Improvements to the algorithm

To be practically useful, it is important to generalize the number of processors on which CAPS can run. To attain the communication lower bounds, CAPS as presented in Section 3 must run on $P$ a power of seven processors. Of course, CAPS can then be run on any number of processors by simply ignoring no more than $\frac{6}{7}$ of them and incurring a constant factor overhead. Thus we can run on arbitrary $P$ and attain the communication and computation lower bounds up to a constant factor. However the computation time is still dominant in most cases, and it is preferable to attain the computation lower bound exactly. It is an open question whether any algorithm can run on arbitrary $P$, attain the computation lower bound exactly, and attain the communication lower bound up to a constant factor.

Moreover, although the communication costs of this algorithm match the lower bound up to a constant factor in bandwidth, and up to a $\log P$ factor in latency, it is an open question to determine the optimal constant in the lower bound and perhaps provide a new algorithm that matches it
exactly. Note that in the analysis of CAPS in Section 3, the constants could be slightly improved.

### 6.6 Parallelizing other algorithms

### 6.6.1 Another optimal classical algorithm

We can apply our parallelization approach to recursive classical matrix multiplication to obtain a communicationoptimal algorithm. This algorithm has the same asymptotic communication costs as the 2.5D algorithm 30. We observed comparable performance to the 2.5D algorithm on our experimental platform. As with CAPS, this algorithm has not been optimized for contention, whereas the 2.5D algorithm is very well optimized for contention on torus networks.

### 6.6.2 Other fast matrix multiplication algorithms

Our approach of executing a recursive algorithm in parallel by traversing the recursion tree in DFS (sequential) or BFS (parallel) manners is not limited to Strassen's algorithm. All fast square matrix multiplication algorithms are built out of ways to multiply $n_{0} \times n_{0}$ matrices using $q<n_{0}^{3}$ multiplications. Like with Strassen and Strassen-Winograd, they compute $q$ linear combinations of entries of each of $A$ and $B$, multiply these pairwise, then compute the entries of $C$ as linear combinations of these ${ }^{4}$ CAPS can be easily generalized to any such multiplication, with the following modifications:

- The number of processors $P$ is a power of $q$.
- The data layout must be such that all $n_{0}^{2}$ blocks of $A, B$, and $C$ are distributed equally among the $P$ processors with the same layout.
- The BFS and DFS determine whether the $q$ multiplications are performed in parallel or sequentially.

The communication costs are then exactly as above, but with $\omega_{0}=\log _{n_{0}} q$.

It is unclear whether any of the faster matrix multiplication algorithms are useful in practice. One reason is that the fastest algorithms are not explicit. Rather, there are non-constructive proofs that the algorithms exist. To implement these algorithms, they would have to be found, which appears to be a difficult problem. The generalization of CAPS described in this section does apply to all of them, so we have proved the existence of a communication-avoiding non-explicit parallel algorithm corresponding to every fast matrix multiplication algorithm. We conjecture that the algorithms are all communication-optimal, but that is not yet proved since the lower bound proofs in 4, 3, may not apply to all fast matrix multiplication algorithms. In cases where the lower bounds do apply, they match the performance of the generalization of CAPS, and so they are communicationoptimal.

[^4]
## 7. REFERENCES

[1] R. C. Agarwal, S. M. Balle, F. G. Gustavson, M. Joshi, and P. Palkar. A three-dimensional approach to parallel matrix multiplication. IBM Journal of Research and Development, 39:39-5, 1995.
[2] A. Aggarwal, A. K. Chandra, and M. Snir. Communication complexity of PRAMs. Theor. Comput. Sci., 71:3-28, March 1990.
[3] G. Ballard, J. Demmel, O. Holtz, B. Lipshitz, and O. Schwartz. Strong scaling of matrix multiplication algorithms and memory-independent communication lower bounds, 2012. Manuscript, submitted to SPAA 2012.
[4] G. Ballard, J. Demmel, O. Holtz, and O. Schwartz. Graph expansion and communication costs of fast matrix multiplication. In SPAA '11: Proceedings of the 23rd annual symposium on parallelism in algorithms and architectures, pages $1-12$, New York, NY, USA, 2011. ACM.
[5] G. Ballard, J. Demmel, O. Holtz, and O. Schwartz. Minimizing communication in numerical linear algebra. SIAM J. Matrix Analysis Applications, 32(3):866-901, 2011.
[6] J. Berntsen. Communication efficient matrix multiplication on hypercubes. Parallel Computing, 12(3):335-342, 1989.
[7] D. Bini. Relations between exact and approximate bilinear algorithms. applications. Calcolo, 17:87-97, 1980. 10.1007/BF02575865.
[8] L. S. Blackford, J. Choi, A. Cleary, E. DŠAzevedo, J. Demmel, I. Dhillon, J. Dongarra, S. Hammarling, G. Henry, A. Petitet, K. Stanley, D. Walker, and R. C. Whaley. ScaLAPACK Users' Guide. SIAM, Philadelphia, PA, USA, May 1997. Also available from http://www.netlib.org/scalapack/.
[9] L. Cannon. A cellular computer to implement the Kalman filter algorithm. PhD thesis, Montana State University, Bozeman, MN, 1969.
[10] J. Choi. A new parallel matrix multiplication algorithm on distributed-memory concurrent computers.
Concurrency: Practice and Experience, 10(8):655-670, 1998.
[11] J. Choi, D. W. Walker, and J. J. Dongarra. PUMMA: Parallel universal matrix multiplication algorithms on distributed memory concurrent computers. Concurrency: Practice and Experience, 6(7):543-570, 1994.
[12] H. Cohn, R. D. Kleinberg, B. Szegedy, and C. Umans. Group-theoretic algorithms for matrix multiplication. In $F O C S$, pages 379-388, 2005.
[13] D. Coppersmith and S. Winograd. On the asymptotic complexity of matrix multiplication. SIAM Journal on Computing, 11(3):472-492, 1982.
[14] D. Coppersmith and S. Winograd. Matrix multiplication via arithmetic progressions. In Proceedings of the Nineteenth Annual ACM Symposium on Theory of Computing, STOC '87, pages 1-6, New York, NY, USA, 1987. ACM.
[15] J. Demmel, I. Dumitriu, and O. Holtz. Fast linear algebra is stable. Numerische Mathematik, 108(1):59-91, 2007.
[16] J. Demmel, I. Dumitriu, O. Holtz, and R. Kleinberg.

Fast matrix multiplication is stable. Numerische Mathematik, 106(2):199-224, 2007.
[17] F. Desprez and F. Suter. Impact of mixed-parallelism on parallel implementations of the Strassen and Winograd matrix multiplication algorithms: Research articles. Concurr. Comput. : Pract. Exper., 16:771-797, July 2004.
[18] S. H. Fuller and L. I. Millett, editors. The Future of Computing Performance: Game Over or Next Level? The National Academies Press, Washington, D.C., 2011. 200 pages, http://www.nap.edu.
[19] S. L. Graham, M. Snir, and C. A. Patterson, editors. Getting up to Speed: The Future of Supercomputing. Report of National Research Council of the National Academies Sciences. The National Academies Press, Washington, D.C., 2004. 289 pages, http://www.nap.edu.
[20] B. Grayson, A. Shah, and R. van de Geijn. A high performance parallel Strassen implementation. In Parallel Processing Letters, volume 6, pages 3-12, 1995.
[21] N. J. Higham. Accuracy and Stability of Numerical Algorithms. SIAM, Philadelphia, PA, 2nd edition, 2002.
[22] S. Hunold, T. Rauber, and G. Rünger. Combining building blocks for parallel multi-level matrix multiplication. Parallel Computing, 34:411-426, July 2008.
[23] D. Irony, S. Toledo, and A. Tiskin. Communication lower bounds for distributed-memory matrix multiplication. J. Parallel Distrib. Comput., 64(9):1017-1026, 2004.
[24] B. Kumar, C.-H. Huang, R. Johnson, and P. Sadayappan. A tensor product formulation of Strassen's matrix multiplication algorithm with memory reduction. In Proceedings of Seventh International Parallel Processing Symposium, pages 582 -588 , apr 1993.
[25] Q. Luo and J. Drake. A scalable parallel Strassen's matrix multiplication algorithm for distributed-memory computers. In Proceedings of the 1995 ACM symposium on Applied computing, SAC '95, pages 221-226, New York, NY, USA, 1995. ACM.
[26] V. Y. Pan. New fast algorithms for matrix operations. SIAM Journal on Computing, 9(2):321-342, 1980.
[27] R. Raz. On the complexity of matrix product. SIAM J. Comput., 32(5):1356-1369 (electronic), 2003.
[28] F. Romani. Some properties of disjoint sums of tensors related to matrix multiplication. SIAM Journal on Computing, 11(2):263-267, 1982.
[29] A. Schönhage. Partial and total matrix multiplication. SIAM Journal on Computing, 10(3):434-455, 1981.
[30] E. Solomonik and J. Demmel. Communication-optimal parallel 2.5 D matrix multiplication and LU factorization algorithms. In Euro-Par'11: Proceedings of the 17th International European Conference on Parallel and Distributed Computing. Springer, 2011.
[31] F. Song, J. Dongarra, and S. Moore. Experiments with Strassen's algorithm: From sequential to parallel. In Proceedings of Parallel and Distributed Computing and Systems (PDCS). ACTA, Nov. 2006.
[32] V. Strassen. Gaussian elimination is not optimal. Numer. Math., 13:354-356, 1969.
[33] V. Strassen. Relative bilinear complexity and matrix multiplication. Journal für die reine und angewandte Mathematik (Crelles Journal), 1987(375-376):406-443, 1987.
[34] R. A. van de Geijn and J. Watts. SUMMA: scalable universal matrix multiplication algorithm. Concurrency - Practice and Experience, 9(4):255-274, 1997.
[35] V. V. Williams. Breaking the Coppersmith-Winograd barrier, 2012. Manuscript.
[36] S. Winograd. On the multiplication of $2 \times 2$ matrices. Linear Algebra Appl., 4(4):381-388., October 1971.
[37] C.-Q. Yang and B. Miller. Critical path analysis for the execution of parallel and distributed programs. In Proceedings of the 8th International Conference on Distributed Computing Systems, pages 366-373, Jun. 1988.

## APPENDIX

## A. STRASSEN-WINOGRAD ALGORITHM

The Strassen-Winograd algorithm is usually preferred to Strassen's algorithm in practice since it requires fewer additions. We use it for our implementation of CAPS. Divide the input matrices $A, B$ and output matrix $C$ into 4 submatrices:

$$
A=\left[\begin{array}{ll}
A_{11} & A_{12} \\
A_{21} & A_{22}
\end{array}\right] B=\left[\begin{array}{ll}
B_{11} & B_{12} \\
B_{21} & B_{22}
\end{array}\right] C=\left[\begin{array}{ll}
C_{11} & C_{12} \\
C_{21} & C_{22}
\end{array}\right]
$$

Then form 7 linear combinations of the submatrices of each of $A$ and $B$, call these $T_{i}$ and $S_{i}$, respectively; multiply them pairwise; then form the submatrices of $C$ as linear combinations of these products:

$$
\begin{array}{llll}
T_{0}=A_{11} & S_{0}=B_{11} & Q_{0}=T_{0} \cdot S_{0} & U_{1}=Q_{0}+Q_{3} \\
T_{1}=A_{12} & S_{1}=B_{21} & Q_{1}=T_{1} \cdot S_{1} & U_{2}=U_{1}+Q_{4} \\
T_{2}=A_{21}+A_{22} & S_{2}=B_{12}+B_{11} & Q_{2}=T_{2} \cdot S_{2} & U_{3}=U_{1}+Q_{2} \\
T_{3}=T_{2}-A_{12} & S_{3}=B_{22}-S_{2} & Q_{3}=T_{3} \cdot S_{3} & C_{11}=Q_{0}+Q_{1} \\
T_{4}=A_{11}-A_{21} & S_{4}=B_{22}-B_{12} & Q_{4}=T_{4} \cdot S_{4} & C_{12}=U_{3}+Q_{5} \\
T_{5}=A_{12}+T_{3} & S_{5}=B_{22} & Q_{5}=T_{5} \cdot S_{5} & C_{21}=U_{2}-Q_{6} \\
T_{6}=A_{22} & S_{6}=S_{3}-B_{21} & Q_{6}=T_{6} \cdot S_{6} & C_{22}=U_{2}+Q_{2}
\end{array}
$$

This is one step of Strassen-Winograd. The algorithm is recursive since it can be used for each of the 7 smaller matrix multiplications. In practice, one often uses only a few steps of Strassen-Winograd, although to attain $O\left(n^{\omega_{0}}\right)$ computational cost, it is necessary to recursively apply it all the way down to matrices of size $O(1) \times O(1)$. The precise computational cost of Strassen-Winograd is

$$
\mathrm{F}(n)=c_{s} n^{\omega_{0}}-5 n^{2}
$$

Here $c_{s}$ is a constant depending on the cutoff point at which one switches to the classical algorithm. For a cutoff size of $n_{0}$, the constant is $c_{s}=\left(2 n_{0}+4\right) / n_{0}^{\omega_{0}-2}$ which is minimized at $n_{0}=8$ yielding a computational cost of approximately $3.73 n^{\omega_{0}}-5 n^{2}$. If using Strassen-Winograd with cutoff $n_{0}$ this should be substituted into the computational cost expressions of Section 3

## B. COMMUNICATION-COST RATIOS

In this section we derive the ratio $R$ of classical to Strassenbased bandwidth cost lower bounds that appear in the beginning of Section 4 . Note that both classical and Strassen-based lower bounds are attained by optimal algorithms. Similar derivations apply to the other ratios quoted in that section. Because the bandwidth cost lower bounds are different in the memory-dependent and the memory-independent cases, and the threshold between these is different for the classical
and Strassen-based bounds, it is necessary to consider three cases.

Case 1. $M=\Omega\left(n^{2} / P\right)$ and $M=O\left(n^{2} / P^{2 / \omega_{0}}\right)$. The first condition is necessary for there to be enough memory to hold the input and output matrices; the second condition puts both classical and Strassen-based algorithms in the memory-dependent case. Then the ratio of the bandwidth costs is:

$$
R=\Theta\left(\frac{n^{3}}{P M^{1 / 2}} / \frac{n^{\omega_{0}}}{P M^{\omega_{0} / 2-1}}\right)=\Theta\left(\left(\frac{n^{2}}{M}\right)^{\left(3-\omega_{0}\right) / 2}\right)
$$

Using the two bounds that define this case, we obtain $R=$ $O\left(P^{\left(3-\omega_{0}\right) / 2}\right)$ and $R=\Omega\left(P^{3 / \omega_{0}-1}\right)$.

Case 2. $M=\Omega\left(n^{2} / P^{2 / \omega_{0}}\right)$ and $M=O\left(n^{2} / P^{2 / 3}\right)$. This means that in the classical case the memory-dependent lower bound dominates, but in the Strassen-based case the memoryindependent lower bound dominates. Then the ratio is:

$$
R=\Theta\left(\frac{n^{3}}{P M^{1 / 2}} / \frac{n^{2}}{P^{2 / \omega_{0}}}\right)=\Theta\left(\left(\frac{n^{2}}{M}\right)^{1 / 2} P^{2 / \omega_{0}-1}\right)
$$

Using the two bounds that define this case, we obtain $R=$ $O\left(P^{3 / \omega_{0}-1}\right)$ and $R=\Omega\left(P^{2 / \omega_{0}-2 / 3}\right)$.

Case 3. $M=O\left(P^{2 / 3}\right)$. This means that both the classical and Strassen-based lower bounds are dominated by the memory-independent cases. Then the ratio is:

$$
R=\Theta\left(\frac{n^{2}}{P^{2 / 3}} / \frac{n^{2}}{P^{2 / \omega_{0}}}\right)=\Theta\left(P^{2 / \omega_{0}-2 / 3}\right)
$$

Overall, depending on the ratio of the problem size to the available memory, the factor by which the classical bandwidth costs exceed the Strassen-based bandwidth costs is between $\Theta\left(P^{2 / \omega_{0}-2 / 3}\right)$ and $\Theta\left(P^{\left(3-\omega_{0}\right) / 2}\right)$.


[^0]:    *Research supported by Microsoft (Award \#024263) and Intel (Award \#024894) funding and by matching funding by U.C. Discovery (Award \#DIG07-10227). Additional support comes from Par Lab affiliates National Instruments, Nokia, NVIDIA, Oracle, and Samsung.
    $\dagger$ Research is also supported by DOE grants DE-SC0003959, DE-SC0004938, and DE-AC02-05CH11231.
    ${ }^{\ddagger}$ Research supported by the Sofja Kovalevskaja programme of Alexander von Humboldt Foundation and by the National Science Foundation under agreement DMS-0635607, while visiting the Institute for Advanced Study.
    ${ }^{\S}$ Research supported by U.S. Department of Energy grants under Grant Numbers DE-SC0003959.

[^1]:    ${ }^{1}$ Our actual implementation uses the Winograd variant 36; see Appendix A for details.

[^2]:    ${ }^{2}$ See [4] for a discussion of the sequential memory model.

[^3]:    ${ }^{3}$ If one does not overwrite the input, it is impossible to run Strassen in place; however using a few temporary matrices affects the analysis here by a constant factor only.

[^4]:    ${ }^{4}$ By 27, all fast matrix multiplication algorithms can be expressed in this bilinear form.

