

Low-Depth Parallel Algorithms for the Binary-Forking Model without Atomics

Zafar Ahmad¹, Rezaul Chowdhury¹, Rathish Das¹, Pramod Ganapathi¹, Aaron Gregory²,
and Mohammad Mahdi Javanmard¹

¹Department of Computer Science, Stony Brook University

²Department of Applied Mathematics & Statistics, Stony Brook University

Abstract

The binary-forking model is a parallel computation model, formally defined by Blelloch et al. very recently, in which a thread can fork a concurrent child thread, recursively and asynchronously. The model incurs a cost of $\Theta(\log n)$ to spawn or synchronize n tasks or threads. The binary-forking model realistically captures the performance of parallel algorithms implemented using modern multithreaded programming languages on multicore shared-memory machines. In contrast, the widely studied theoretical PRAM model does not consider the cost of spawning and synchronizing threads, and as a result, algorithms achieving optimal performance bounds in the PRAM model may not be optimal in the binary-forking model. Often, algorithms need to be redesigned to achieve optimal performance bounds in the binary-forking model and the non-constant synchronization cost makes the task challenging.

Though the binary-forking model allows the use of atomic *test-and-set* (TS) instructions to reduce some synchronization overhead, assuming the availability of such instructions puts a stronger requirement on the hardware and may limit the portability of the algorithms using them. In this paper, we avoid the use of locks and atomic instructions in our algorithms except possibly inside the join operation which is implemented by the runtime system.

In this paper, we design efficient parallel algorithms in the binary-forking model without atomics for three fundamental problems: Strassen’s (and Strassen-like) matrix multiplication (MM), comparison-based sorting, and the Fast Fourier Transform (FFT). All our results improve over known results for the corresponding problem in the binary-forking model both with and without atomics.

We present techniques for designing efficient algorithms without using locks and atomic instructions. We use extra space to prevent the work blow-up in the highly parallel asynchronous computations performed by our MM and FFT algorithms. Extra space also allows us to avoid the use of atomic TS instructions in our sorting algorithm as well as to achieve a stronger bound (i.e., with high probability) on its work. Though space plays a major role in the design of all our algorithms, our MM and FFT algorithms do not use asymptotically more space than their PRAM counterparts. We present space-adaptive algorithms for MM and sorting that achieve provably good performance bounds for any given amount of space.

1 Introduction

We present efficient algorithms with optimal/near-optimal span¹ for several fundamental problems in the binary-forking model without locks and atomic instructions. The binary-forking model was introduced in Blelloch et al. [6] (see also [1, 3, 4, 5, 20]) to accurately capture the performance of algorithms designed for modern multi-core shared-memory machines. In this model, the computation starts with a single thread, and as the computation progresses, threads are created dynamically and asynchronously; the computation finishes when all threads end. A thread can spawn/fork a concurrent asynchronous child thread while it progresses simultaneously and such forking of threads can happen recursively; hence the model is called the binary-forking model. The model also includes a “join” operation to synchronize the threads. Though the model introduced in [6] allows the use of atomic instructions, we do not use them in this paper.

The binary-forking model is closely related to the well-studied PRAM model [29]. The PRAM model is strictly more powerful than the binary-forking model; however, it does not correlate well with modern architectures. In the PRAM model, computation progresses in synchronous steps. Modern architectures employ new techniques such as use of multiple caches, processor pipelining, branch prediction, hyper-threading, and many more, which give rise to many asynchronous events such as cache misses, varying clock speed, interrupts, etc., thus demanding the development of a parallel computation model where computation can proceed asynchronously. Asynchronous thread creation in the binary-forking model makes it an ideal candidate for modeling parallel computation in modern architectures. As pointed out in [6], this is the model underlying many widely used parallel programming languages/environments such as Cilk [25], the Java fork-join framework [22], Intel TBB [39], and the Microsoft Task Parallel Library [37].

One can trivially reduce any algorithm designed for the PRAM model to an algorithm for the binary-forking model, incurring an $O(\log n)$ -factor blow-up in the span while keeping the work² asymptotically the same as in the PRAM model. Spawning n threads takes $\Theta(1)$ time and $\Theta(n)$ work in the PRAM model—making the synchronization cost (span) constant. This synchronization can be simulated by using a binary tree of $\log n$ depth and $\Theta(n)$ nodes in the binary-forking model. Each internal node in the binary tree corresponds to a binary-forking operation, and the n leaves correspond to the n spawned threads.

A direct simulation of an optimal-span PRAM algorithm may not produce an algorithm with optimal span in the binary-forking model. For example, Cole’s parallel merge sort [16] achieves optimal $\Theta(\log n)$ span and $\Theta(n \log n)$ work in the PRAM model. The binary-tree reduction increases the span to $\Theta(\log^2 n)$ while keeping the work asymptotically the same. On the other hand, by increasing work to $\Theta(n^2)$, it becomes trivial to get a $\Theta(\log n)$ span sorting algorithm — each item independently computes its rank in the final sorted list in $\Theta(\log n)$ time and $\Theta(n)$ work by comparing itself with all n elements. However, neither algorithm is optimal in the binary-forking model — the former has non-optimal span while the latter performs non-optimal work. Cole and Ramachandran [17] presented a deterministic sorting algorithm with $\mathcal{O}(\log n \log \log n)$ span and optimal $\Theta(n \log n)$ work in the binary-forking model. Recently, Ramachandran and Shi [34] gave a data-oblivious sorting algorithm in the binary-forking model with optimal work and $\mathcal{O}(\log n \log \log n)$ span. Very recently, Blelloch et al. [6] used atomic test-and-set operations to design a randomized sorting algorithm with $\Theta(\log n)$ span w.h.p.³ in n and $\Theta(n \log n)$ work in expectation. Hence, finding an optimal (both in span and work) sorting algorithm without using atomic instructions

¹Span/depth is the running time of an algorithm with an unbounded number of processors.

²Work is the number of operations performed by a parallel algorithm on a serial computer.

³An event ξ occurs with high probability (w.h.p.) in n provided it occurs with probability at least $1 - \frac{c}{n^\alpha}$ for constants $\alpha \geq 1$ and $c > 0$.

remains an interesting and non-trivial open problem in the binary-forking model. We encounter the span blow-up problem when running other fundamental low-span PRAM algorithms, such as those for Strassen’s matrix multiplication and Fast Fourier Transform (FFT), in the binary-forking model. Both algorithms have $\Theta(\log n)$ span in the PRAM model, which blow up by a factor of $\log n$ and $\log \log n$, respectively, in the binary-forking model.

The binary-forking model introduced in [6] allows the use of atomic *test-and-set* (TS) operations. When performed on a shared memory location L , TS performs the following as a single undivided operation: it reads the value stored at L and if the value is zero, sets L to one and returns zero, otherwise, returns one without changing L . While a TS operation makes the binary-forking model arguably more powerful, it also puts a stronger requirement on the memory hardware. Hence, an algorithm may become more portable by avoiding the use of TS. Such an algorithm is also arguably more elegant [6]. While several parallel algorithms have been designed for the binary-forking model without locks and atomic instructions [24, 19, 13, 14, 4, 15, 7, 5, 40, 11, 12, 10, 30, 20, 42, 41, 26, 36], developing algorithms with optimal/near-optimal span and work for several fundamental problems in the binary-forking model without locks and atomic instructions remain open.

Algorithms for the binary-forking model without atomics face two major challenges: 1) how to avoid the blow-up in span (synchronization cost) without blowing up work? 2) how to avoid the use of atomic operations without asymptotically increasing span and work? Surprisingly, it turns out that using extra space, we can tackle both challenges. By extra space, we mean the space allocated from RAM (from heap memory), not from processors registers or stack memory.

Our Contributions. In this paper, we present results for three fundamental problems in the binary-forking model without atomics. Our major results include: **(1)** an optimal $O(\log n)$ span algorithm for Strassen’s Matrix Multiplication (MM) with only a $\Theta(\log \log n)$ -factor blow-up in work as well as a near-optimal $O(\log n \log \log n)$ span algorithm with no asymptotic blow-up in work; **(2)** a randomized comparison-based sorting algorithm with optimal $O(\log n)$ span and $O(n \log n)$ work, both with w.h.p. in n ; and **(3)** a near-optimal $O(\log n \log \log \log n)$ span algorithm for FFT with less than a $\log n$ -factor blow-up in work for all practical values of n (i.e., $n \leq 10^{10,000}$).

Though space played a major role in the design of all our algorithms in this paper, our algorithms for Strassen’s matrix multiplication and FFT do not use asymptotically more space than their standard PRAM counterparts. We present a space-adaptive algorithm for Strassen’s MM which always achieves a span within $\Theta(\log n)$ factor of optimal for any given amount of space.

We list our major results in Table 1.

Major Techniques. The extra $\log n$ factor in the span of the standard parallelization of Strassen’s MM algorithm in the binary-forking model arises from the fact that it spends $\Theta(\log \frac{n}{2^i})$ time (synchronization cost) computing intermediate results at recursion level i for each $i \in [0, \log_2 n]$ which requires only $\mathcal{O}(1)$ time in the PRAM model. We observe that none of those intermediate matrices need to be explicitly computed or stored to compute the final output. Indeed, each cell in the final output matrix can be computed directly in $\mathcal{O}(\log n)$ time from the two original input matrices of the algorithm. This *single-point computation method* can be used to compute all the cells in the output matrix simultaneously in $\Theta(\log n)$ span. However, this approach blows up the work performed by the algorithm by up to a $\Theta(n^2)$ factor because the approach does not reuse intermediate results. We avoid this work blow-up by computing and temporarily storing the intermediate results at $\Theta(\log \log n)$ carefully ‘chosen levels’ of recursion, which eliminates the need for implicitly recomputing the intermediate results over and over again. So, all single-point computations proceed in stages where a stage includes all levels of recursion between two consecutive ‘chosen levels,’

and synchronizations happen only at stage boundaries with all threads executing asynchronously within every stage. We show that this **stage-based approach** reduces the work blow-up from $\Theta(n^2)$ factor to only $\Theta(\log \log n)$ factor while achieving the optimal $\Theta(\log_2 n)$ span. A similar approach works for FFT.

The technique described above works for all Strassen-like algorithms, including Victor Pan's $\mathcal{O}(n^{2.795})$ work algorithm [33]. We remark that while we use additional techniques specific to the problems to achieve better work and span bounds, the main contribution is devising the general technique to enable limited work-sharing among the single-point computations using extra space.

Using extra space helps to avoid atomic TS operations, too. Blelloch et al. [6] achieve optimal span (with high probability) and optimal work (in expectation) for sorting, semisorting, and random permutation using TS. They use TS to distribute a set of n items into $m = o(n)$ buckets where each item knows its destination bucket, but how many items will fall in a bucket is unknown. They reduce the problem to a variant of a balls and bins problem. In particular, when n items try to find unoccupied cells randomly among $c \cdot n$ cells (c is a constant) in parallel, it is enough for each item to try $\Theta(\log n)$ times to find an unoccupied cell with high probability; the span of this process is thus $\Theta(\log n)$. An item tries to put itself in a random cell using a TS on a flag to reserve it. If the TS fails, it tries again since the cell is already taken. In the absence of TS, it would take $\Theta(\log n)$ synchronization time to figure out the items that fail to find a cell after each attempt, thus making the overall span $\Theta(\log^2 n)$ for the $\Theta(\log n)$ attempts. Our approach can avoid this $\log n$ synchronization steps by increasing the space by a factor of $\log n$ and by allowing each item to simultaneously attempt to place itself in $\log n$ random cells. We assume arbitrary concurrent writes meaning that if multiple concurrent threads try to write to any given shared memory location simultaneously only one arbitrary thread succeeds. This approach does not increase span and work.

While Blelloch et al.'s [6] sorting algorithm performs $\Theta(n \log n)$ work in expectation, we achieve the same bound w.h.p. in n . The success probability of the distribution step in a recursive call of Blelloch et al.'s algorithm is dependent on the size of the input to that recursive call. Since input size decreases doubly exponentially with the increase of recursion level, though the success probability is high in the corresponding input size, it reduces rapidly as execution moves deeper in the recursion tree and does not remain high w.r.t. the original input size n_0 .

To achieve a $\Theta(n \log n)$ work bound w.h.p. in n (along with a $\Theta(\log n)$ span also w.h.p. in n) our FULL-SORT algorithm proceeds in two phases — a recursive ALMOST-SORT phase and a non-recursive leftover integration phase. Extra space is used throughout FULL-SORT to simulate TS operations as described in the previous paragraph. Given an input of size n_0 , ALMOST-SORT sorts $n_0 - o(n_0)$ items of the input, then the remaining $o(n_0)$ items are merged with the sorted $n_0 - o(n_0)$ items in the integration step. ALMOST-SORT is a recursive bucketing algorithm with some similarity to Blelloch et al.'s algorithm. However, the bucket size in each of its recursive calls is an $r(n_0)$ factor larger than the ones used in Blelloch et al.'s paper, where $r(n_0) = \Theta\left(\frac{\log n_0 \log \log \log n_0}{\log \log n_0}\right)$, and recursion in ALMOST-SORT is terminated much sooner than in Blelloch et al.'s algorithm, so that throughout ALMOST-SORT partitioning succeeds w.h.p. in n_0 . Each item tries to put itself into its destination bucket twice and fails with probability $1/(r(n_0))^2$. We set aside the failed items to be incorporated into the final sorted sequence during the integration step later and move to the next level of recursion without them. By allowing some items to remain unsorted, we make the work done by each recursive step effectively independent of the probability that single items are successfully written to their chosen location. We show that with high probability in n_0 , at most $\Theta(n_0/(r(n_0))^2)$ items fail to move from any level of recursion to the next level. By switching to Cole-Ramachandran's deterministic sorting algorithm [17] after $\log \log \log n_0$ levels of recursion, we ensure that ALMOST-SORT performs $O(n_0 \log n_0)$ work w.h.p. in n_0 . The integration phase then

Algorithm	Work (T_1)	Space (S_∞)	Span (T_∞)	Result
Strassen's Matrix Multiplication				
Strassen's MM [35] [19]	$\mathcal{O}(n^w)$	$\Theta(n^w)$	$\mathcal{O}(\log^2 n)$	
♣ STRASSEN-S	$\mathcal{O}(n^w \log \log n)$	$\Theta(n^w)$	$\mathcal{O}(\log n)$	Th. 2.4
♣ STRASSEN-W	$\mathcal{O}(n^w)$	$\Theta(n^w / \log \log n)$	$\mathcal{O}(\log n \log \log \log n)$	Th. 2.7
♣ STRASSEN-S-ADAPTIVE	$\mathcal{O}(n^w \log \log n)$	$\Theta(s)$	$\mathcal{O}((n^w/s) \log n)$	Th. 2.5
♣ STRASSEN-W-ADAPTIVE	$\mathcal{O}(n^w)$	$\Theta(s)$	$\mathcal{O}((n^w/s) \log^2 n)$	Th. 2.6
Sorting				
COLE-RAMACHANDRAN [17]	$\mathcal{O}(n \log n)$	$\Theta(n)$	$\mathcal{O}(\log n \log \log n)$	
Blelloch et al. (atomic) [6]	$\mathcal{O}(n \log n)$ exp.	$\Theta(n)$	$\mathcal{O}(\log n)$ whp	
♣ FULL-SORT	$\mathcal{O}(n \log n)$ whp	$\Theta\left(\frac{n \log n \log \log \log n}{\log \log n}\right)$	$\mathcal{O}(\log n)$ whp	Th. 3.6
♣ SORT-ADAPTIVE	$\mathcal{O}(n \log n)$ whp	$\Theta(s)$	$\mathcal{O}((n/s) \log^2 n)$ whp	Th. 3.7
Fast Fourier Transform				
Cooley-Tukey \sqrt{n} -way [18]	$\mathcal{O}(n \log n)$	$\Theta(n)$	$\mathcal{O}(\log n \log \log n)$	
♣ FFT	$\mathcal{O}(n \log^{g(n)} n)$	$\Theta(n)$	$\mathcal{O}(\log n \log \log \log n)$	Cor. 4.3.1

Table 1: The complexity analyses of our new algorithms (marked with ♣) in the binary-forking model without atomics. Here, n = problem size, s = space \in [input size, work], $w = \log_2 7$, exp. = expected time, whp = with high probability in n , $g(n) < 2$ for $n < 10^{10,000}$.

combines the $\mathcal{O}(n_0 \log \log \log n_0 / (r(n_0))^2)$ (w.h.p. in n_0) leftover items with the already sorted sequence in $\mathcal{O}(n_0 \log n_0)$ work w.h.p. in n_0 .

Binary-Forking Model. Binary-forking model captures the current multi-core shared-memory systems. Many parallel algorithms are based on binary-forking model [1, 3, 4, 5, 20]. Computations in the binary-forking model can be viewed as a series-parallel DAG where each node represents a thread's instruction. The root of the tree is the first instruction of the starting thread. Each node has at most two children. If node u denotes the i -th instruction of thread t and u has only one child v , then v denotes the $(i+1)$ -th instruction of thread t . If node u has two children v and w , then v represents the $(i+1)$ -th instruction of thread t and w represents the first instruction of the new forked thread t' . The binary-forking model includes “join” instructions to join the forking threads. They are modeled as a node with two incoming edges. The work of the computation is the number of nodes in the series-parallel DAG and the span of the computation is the length of the longest path in the DAG assuming unbounded resources such as processors and space.

Performance Metrics of a Parallel Program. We use the *work-span* model [19] to analyze the performance of parallel programs executed on shared-memory multicore machines. The *work* of a multithreaded program, denoted by $T_1(n)$, where n is the input parameter, is defined as the total number of CPU operations it performs when executed on a single processor. The *span* $T_\infty(n)$ of a program which is also known as its *critical-path length* or *depth*, is the maximum number of operations performed on any single processor when the program is run on an unbounded number of processors. The *parallel running time* $T_p(n)$ of a program when run on p processors under a greedy scheduler is given by $T_p(n) = \mathcal{O}(T_1(n)/p + T_\infty(n))$. The *parallelism*, computed by the ratio of $T_1(n)$ and $T_\infty(n)$, is the average amount of work performed by the program in each step of its critical path.

2 Strassen's Matrix Multiplication

Suppose $w = \log_2 7$. Strassen's matrix multiplication (MM) algorithm [35] performs $\mathcal{O}(n^w)$ work (i.e., multiplications and additions), unlike the classic MM algorithm that performs $\mathcal{O}(n^3)$ work. A straightforward parallelization of Strassen's MM leads to $\Theta(\log^2 n)$ span. Our goal is to design a parallel Strassen's MM in the binary-forking model without using locks and atomic instructions to achieve an optimal span of $\mathcal{O}(\log n)$ without affecting the work bound of $\Theta(n^w)$.

In this paper, we present parallel Strassen MM algorithms (i) having optimal $\mathcal{O}(\log n)$ span and $\mathcal{O}(n^w \log \log n)$ work, i.e., work very close to that of the standard Strassen's MM; and (ii) having $\mathcal{O}(n^w)$ work and $\mathcal{O}(\log n \log \log \log n)$ span, i.e., very close to optimal span.

The core ideas and techniques used in our algorithms are as follows. We first perform **single-point computation**, i.e., computation of a single cell of the output matrix independently from that of other cells/entries. This implies that all cells of the output matrix are computed independently in $\mathcal{O}(\log n)$ span. However, as there is no work-sharing across multiple threads, the total work blows up to $\mathcal{O}(n^{w+2})$. We enable partial work-sharing across threads by saving intermediate computations at carefully selected levels of recursion. By splitting the recursion tree into **stages** and allowing work-sharing across stages, we are able to reduce the work to very close to $\mathcal{O}(n^w)$. Hence, by using single-point computations in stages, we are able to obtain good work and span bounds. We use this algorithm to design other parallel Strassen's MM algorithms with different advantages.

k -way Strassen's MM [35, 19]. The k -way Strassen's MM, for $k \in [1, 7]$, executes the child nodes in exactly $\lceil 7/k \rceil$ parallel steps without executing more than k child nodes at a time.

Lemma 2.1 ([35, 19]). *The k -way Strassen's MM has a complexity of $\mathcal{O}(n^w)$ work, $\mathcal{O}(\log^2 n)$ span if $k = 7$, $\mathcal{O}(n^{\log_2 \lceil 7/k \rceil})$ span if $k \neq 7$, $\mathcal{O}(n^2 \log n)$ space if $k = 4$, and $\mathcal{O}(n^{\max(2, \log_2 k)})$ space if $k \neq 4$.*

Proof. The work, span, and extra space recurrences for the k -way Strassen's MM are as follows. If $n = 1$, then $T_1(n) = \mathcal{O}(1)$ and $T_\infty(n) = \mathcal{O}(1)$. If $n > 1$, then

$$T_1(n) = 7T_1(n/2) + \mathcal{O}(n^2), \quad T_\infty(n) = \lceil 7/k \rceil T_\infty(n/2) + \mathcal{O}(\log n), \quad S_\infty(n) = kS_\infty(n/2) + \mathcal{O}(n^2).$$

Solving these recurrences, we have the lemma. \square

The work of the k -way Strassen's MM for any value of k is $\mathcal{O}(n^{\log_2 7})$. The k -way algorithm gives a trade-off between span and space. When $k = 1$, we get the standard Strassen's algorithm [35]. When $k = 7$, we get the standard parallel Strassen's MM [19] that spawns all the child nodes in parallel achieving $\mathcal{O}(\log^2 n)$ span and occupying $\mathcal{O}(n^{\log_2 7})$ space.

Strassen-S MM. In this section, we present a parallel Strassen's MM algorithm, as shown in Figure 1, that achieves the optimal span of $\mathcal{O}(\log n)$ with only a $\mathcal{O}(\log \log n)$ factor increase in the work compared with the classical sequential Strassen's MM algorithm. In this algorithm, we multiply two matrices U and V and store the matrix product in X . We first construct the required data structures as shown in Figure 2. We then compute the input matrices (U and V) at all nodes in the recursion tree in parallel in $\mathcal{O}(\log n)$ span. Finally, we compute the output matrix (X) at all nodes in the recursion tree in $\mathcal{O}(\log n)$ span.

[Step 1. Compute the Input Matrices.] Consider the standard 7-way parallel Strassen's MM. The height of the recursion tree is $\mathcal{O}(\log n)$ and in each level, the total cost of forking

STRASSEN-S(X, U, V)	$\triangleright X \leftarrow U \times V$
1. $(\overline{U}, \overline{V}, \overline{X}, U_{\text{quad}}, V_{\text{quad}}, X_{\text{branch}}) \leftarrow \text{CONSTRUCT-DATA-STRUCTURES}(U, V)$ 2. $\text{COMPUTE-INPUT-MATRICES}(\overline{U}, 0, 0, U_{\text{quad}}); \text{COMPUTE-INPUT-MATRICES}(\overline{V}, 0, 0, V_{\text{quad}})$ 3. $\text{COMPUTE-OUTPUT-MATRICES}(\overline{X}, 0, 0, X_{\text{branch}})$ 4. $X \leftarrow \overline{X}[0][0]$	
COMPUTE-INPUT-MATRICES($Z, \text{stage_id}, \text{root_id}, \text{quad}$)	
1. $\text{height} \leftarrow \# \text{levels in the stage}; \# \text{leaves} \leftarrow 7^{\text{height}}$ 2. parallel for $\text{node} \leftarrow 0$ to $\# \text{leaves} - 1$ do 3. $\text{leaf_id} \leftarrow (\text{root_id} - 1) \times \# \text{leaves} + \text{node}$ 4. parallel for $i \leftarrow 0$ to $n - 1$ do 5. parallel for $j \leftarrow 0$ to $n - 1$ do 6. $Z[\text{stage_id}][\text{leaf_id}][i, j] \leftarrow \text{COMPUTE-INPUT-CELL}(Z, \text{stage_id}, \text{leaf_id}, i, j, n, \text{height}, \text{quad})$ 7. if not last stage then $\text{COMPUTE-INPUT-MATRICES}(Z, \text{stage_id} + 1, \text{leaf_id}, \text{quad})$	
COMPUTE-INPUT-CELL($Z, \text{stage_id}, \text{node_id}, i, j, n, \text{height}, \text{quad}$)	
1. if $\text{height} = 0$ then return $Z[\text{stage_id}][\text{node_id}][i, j]$ 2. $\text{parent_id} \leftarrow (\text{node_id} / 7), \text{branch_id} \leftarrow \text{node_id} \bmod 7$ 3. parallel for $k \leftarrow 0$ to 1 do 4. $t[k] \leftarrow 0; \text{coeff} \leftarrow \text{quad}[\text{branch_id}][k].\text{coeff}$ 5. if $\text{coeff} \neq 0$ then 6. $\text{new_i} \leftarrow n \times \text{quad}[\text{branch_id}][k].\text{shift_i} + i; \text{new_j} \leftarrow n \times \text{quad}[\text{branch_id}][k].\text{shift_j} + j$ 7. $t[k] \leftarrow \text{coeff} \times \text{COMPUTE-INPUT-CELL}(Z, \text{stage_id}, \text{parent_id}, \text{new_i}, \text{new_j}, 2n, \text{height} - 1, \text{quad})$ 8. return $t[0] + t[1]$	
COMPUTE-OUTPUT-MATRICES($Z, \text{stage_id}, \text{root_id}, \text{branch}$)	
1. $\text{height} \leftarrow \# \text{levels in the stage}; \# \text{leaves} \leftarrow 7^{\text{height}}$ 2. parallel for $\text{node} \leftarrow 0$ to $\# \text{leaves} - 1$ do 3. $\text{leaf_id} \leftarrow (\text{root_id} - 1) \times \# \text{leaves} + \text{node}$ 4. if last stage then $Z[\text{stage_id}][\text{leaf_id}][0, 0] \leftarrow U[\text{stage_id}][\text{leaf_id}][0, 0] \times V[\text{stage_id}][\text{leaf_id}][0, 0]$ 5. else $\text{COMPUTE-OUTPUT-MATRICES}(Z, \text{stage_id} + 1, \text{leaf_id}, \text{branch})$ 6. parallel for $i \leftarrow 0$ to $n - 1$ do 7. parallel for $j \leftarrow 0$ to $n - 1$ do 8. $\text{COMPUTE-OUTPUT-CELL}(Z, \text{stage_id}, \text{root_id}, i, j, n, \text{height}, \text{branch})$	
COMPUTE-OUTPUT-CELL($Z, \text{stage_id}, \text{node_id}, i, j, n, \text{height}, \text{branch}$)	
1. if $\text{height} = 0$ then return $Z[\text{stage_id} + 1][\text{node_id}][i, j]$ 2. $\text{shift_i} \leftarrow [i > n/2]; \text{shift_j} \leftarrow [j > n/2]$ $\triangleright []$ is the Iverson bracket 3. $\text{quad_id} \leftarrow 2 \text{ shift_i} + \text{shift_j}; \text{new_i} \leftarrow i - (n/2) \text{ shift_i}; \text{new_j} \leftarrow j - (n/2) \text{ shift_j}$ 4. parallel for $k \leftarrow 0$ to 6 do 5. $t[k] \leftarrow 0; \text{coeff} \leftarrow \text{branch}[\text{quad_id}][k]$ 6. if $\text{coeff} \neq 0$ then 7. $\text{child_id} \leftarrow (\text{node_id} - 1) \times 7 + k$ 8. $t[k] \leftarrow \text{coeff} \times \text{COMPUTE-OUTPUT-CELL}(Z, \text{stage_id}, \text{child_id}, \text{new_i}, \text{new_j}, n/2, \text{height} - 1, \text{branch})$ 9. return $\sum_{k=0}^6 t[k]$	

Figure 1: The STRASSEN-S MM algorithm.

branch_id	$k = 0$		$k = 1$	
	$\langle \text{shift}_i, \text{shift}_j, \text{coeff} \rangle$	$\langle \text{shift}_i, \text{shift}_j, \text{coeff} \rangle$	$\langle \text{shift}_i, \text{shift}_j, \text{coeff} \rangle$	$\langle \text{shift}_i, \text{shift}_j, \text{coeff} \rangle$
0	$\langle 0, 0, 1 \rangle$	$\langle 1, 1, 1 \rangle$	0	$\langle 0, 0, 1 \rangle$
1	$\langle 1, 0, 1 \rangle$	$\langle 1, 1, 1 \rangle$	1	$\langle 0, 0, 1 \rangle$
2	$\langle 0, 0, 1 \rangle$	$\langle -, -, 0 \rangle$	2	$\langle 0, 1, 1 \rangle$
3	$\langle 1, 1, 1 \rangle$	$\langle -, -, 0 \rangle$	3	$\langle 1, 0, 1 \rangle$
4	$\langle 0, 0, 1 \rangle$	$\langle 0, 1, 1 \rangle$	4	$\langle 1, 1, 1 \rangle$
5	$\langle 1, 0, 1 \rangle$	$\langle 0, 0, -1 \rangle$	5	$\langle 0, 0, 1 \rangle$
6	$\langle 0, 1, 1 \rangle$	$\langle 1, 1, -1 \rangle$	6	$\langle 1, 0, 1 \rangle$

quad_id	k						
	0	1	2	3	4	5	6
0	1	0	0	1	-1	0	1
1	0	0	1	0	1	0	0
2	0	1	0	1	0	0	0
3	1	-1	1	0	0	1	0

Figure 2: Data structures required for the STRASSEN-S MM algorithm. Left: $Uquad$ and $Vquad$. Right: $Xbranch$.

and synchronizing threads to compute the input matrices is $\mathcal{O}(\log n)$. Hence, the total span for computing input matrices at all nodes in the recursion tree is $\mathcal{O}(\log^2 n)$. We can reduce the span to $\mathcal{O}(\log n)$ using single-point computation.

A cell of an input matrix (U or V) at a node of the recursion depends on at most two cells of the corresponding input matrix at its parent node. This implies that each cell in an input matrix at a leaf node depends on at most $2^{\log n} = n$ cells in the corresponding input matrix at the root node. If we were to compute all input cells of all input matrices at all nodes, the total work would explode to $\mathcal{O}(n^w \times n) = \mathcal{O}(n^{w+1})$. To keep the work very close to $\mathcal{O}(n^w)$, we split the entire recursion tree into stages. We then use single-point computation of input cells in stages.

For this algorithm, we have $\mathcal{O}(\log \log n)$ stages so that the work performed in each stage is $\mathcal{O}(n^w)$. Using single-point computation in each stage, we are able to achieve the desired optimal span of $\mathcal{O}(\log n)$ limiting the total work to $\mathcal{O}(n^w \log \log n)$.

In this step, we compute the input matrices of all nodes in the recursion tree. The step consists of $h + 1$ sequential stages: $0, 1, \dots, h$, as shown in Figure 3, such that the height of stage i is fixed at $c_i \log n$, where h and c_i are given below:

$$c_i = \begin{cases} 0 & \text{if } i = -1, \\ 1 - \alpha^{i+1} & \text{if } i \in [0, h-1], \text{ such that } w = \log_2 7, \alpha = \frac{1}{w-1}, \text{ and } h = \log_{w-1} \frac{\log n}{\log \log \log n}. \\ 1 & \text{if } i = h. \end{cases} \quad (1)$$

Please refer to Figure 1 (the COMPUTE-INPUT-MATRICES algorithm) for computing the input matrices (U and V) for all the leaf nodes in all stages. We start from stage 0. For any given stage, denoted by stage id, we can easily compute the topmost level, called the root level and the bottommost level, called the leaf level. It is also easy to list out all indices of the leaf nodes in a given stage. So, for all leaf nodes, for all cells in the input matrix in a particular leaf node, we invoke the function COMPUTE-INPUT-CELL. This function computes the value of a specific cell in the input matrix of a leaf node.

The working of the COMPUTE-INPUT-CELL algorithm is as shown in Figure 4 (left). The figure shows the way in which a highlighted cell in the U matrix at a leaf node with id 05 (in base-7 system, for simplicity) is computed. As the last digit of the index is 5, it means that the leaf node is the 5th child of its parent. From the logic of the Strassen's MM algorithm, we know that the U matrix in the 5th child node of a parent node is computed by subtracting the first quadrant (-11) from the third quadrant ($+21$) of the U matrix of the parent node. Hence, we can exactly know the two cells in the U matrix of the parent node on which the highlighted cell in the U matrix of




Figure 3: Stages in the STRASSEN-S MM algorithm for computing the input matrices.




Figure 4: Left: Single-point computation of a cell in the input matrix U at a leaf node in a stage. Right: Single-point computation of a cell in the output matrix X at a root node in a stage. (If there is an arrow from cell a to cell b , it means that cell a depends on cell b .)

the leaf node depends. Also, we can compute the highlighted cell in $\mathcal{O}(1)$ time using the two cells of the parent node. The first digit of the index of the leaf node is 0. This means that the parent node of the leaf node is the 0th child of its parent (i.e., the leaf node's grandparent). From the logic of the Strassen's MM algorithm, we know that the U matrix in the 0th child node of a parent node is computed by adding the first quadrant (+11) to the fourth quadrant (+22) of the U matrix of the parent node. Using this approach, we can trace the path from the leaf node to its ancestor at the root level. So, each cell in the leaf node depends on 2 cells in its parent node which in turn depends on 4 cells in its parent node and so on until we reach a node at the root level. In this way, we can spawn multiple threads that recursively compute each cell at the leaf node using cells from its ancestor at the root level of the stage. The span for computing each cell is simply the height of the stage i.e., the number of levels in that stage.

Once all the cells in a leaf node with id 05 are computed, the algorithm recursively and asynchronously invokes COMPUTE-INPUT-MATRICES for the next stage with this leaf node as the new root. The base case of the COMPUTE-INPUT-MATRICES algorithm is when the algorithm reaches the last stage at which we compute the cells at the leaf nodes using the exact same idea.

Lemma 2.2. COMPUTE-INPUT-MATRICES has a complexity of $\mathcal{O}(n^w \log \log n)$ work, $\mathcal{O}(\log n)$ span, and $\mathcal{O}(n^w)$ space.

Proof. [Work.] We compute the input matrices in $h + 1$ stages S_0, S_1, \dots, S_h , where $h = \log_{w-1}(\log n / \log \log \log n)$. Suppose W_i define the work done at stage S_i . We first come up

with a generic formula for W_i . We use a direct proof to show that $W_i = \mathcal{O}(n^w)$, which implies that the total work is $\sum_{i=0}^h W_i = \mathcal{O}(n^w \log \log n)$.

We compute W_i for $i \in [0, h-1]$. #Nodes at the leaf level of stage S_i is $7^{c_i \log n} = n^{wc_i}$. The #cells in a matrix at the leaf level is $(n/2^{c_i \log n})^2 = n^{2(1-c_i)}$. Each cell in a matrix at the leaf level depends on $\mathcal{O}(2^{(c_i-c_{i-1}) \log n}) = \mathcal{O}(n^{c_i-c_{i-1}})$ cells in a matrix at the root level of the stage. Hence, $W_i = \mathcal{O}(n^{wc_i} n^{2(1-c_i)} n^{c_i-c_{i-1}}) = \mathcal{O}(n^{c_i(w-1)-c_{i-1}+2})$. To show that $W_i = \mathcal{O}(n^w)$ for all $i \in [0, h-1]$, it is enough to prove that $c_i(w-1) - c_{i-1} + 2 = w$. We substitute the values of c_i and c_{i-1} from equation [1] to get: $c_i(w-1) - c_{i-1} + 2 = (1 - (1/(w-1))^{i+1})(w-1) - (1 - (1/(w-1))^i) + 2 = w$.

We now compute W_h . The height of the first h stages is $c_{h-1} \log n$. So, the height of the last stage S_h is $\log n - c_{h-1} \log n$. Substituting the value of c_{h-1} from equation [1] and simplifying, we get the height of stage S_h as $\log \log \log n$. There are n^w nodes at S_h . The size of a matrix at a leaf node is 1×1 . Each cell depends on $2^{\log \log \log n} = \log \log \log n$ cells in a matrix at the root level of stage S_h . Hence, work done at the last stage is $W_h = \mathcal{O}(n^w \log \log n)$.

Combining the work of the first h stages and the last stage, we get $T_1(n) = \sum_{i=0}^{h-1} W_i + W_h = \mathcal{O}(n^w \log \log n)$.

[Span.] Let $T_\infty(m, i)$ denote the span of the COMPUTE-INPUT-MATRICES algorithm starting from stage i where a matrix at the root level is of size $m \times m$. We give a recursive formula to compute $T_\infty(m, i)$. Then, the total span for the algorithm is $T_\infty(n, 0)$.

Consider the COMPUTE-INPUT-MATRICES algorithm. Let $\Delta c_i = c_i - c_{i-1}$. #Nodes at the leaf level of stage S_i is $7^{\Delta c_i \log n}$. Launching these nodes in parallel (line [2]) incurs a span of $\mathcal{O}(\Delta c_i w \log n)$. A matrix at the leaf level will be of size $m/(2^{\Delta c_i \log n}) \times m/(2^{\Delta c_i \log n})$. Spawning COMPUTE-INPUT-CELL function for all cells (lines [4], [5]) incur a span of $\mathcal{O}(2 \log m - 2 \Delta c_i \log n)$. Executing the COMPUTE-INPUT-CELL algorithm incurs $\mathcal{O}(\Delta c_i \log n)$ span. Adding all these spans give us $\mathcal{O}(\Delta c_i \log n + \log m)$.

The span of stage S_i recursively depends upon the span of stage S_{i+1} as the matrices at the leaf level of stage S_{i+1} are constructed from the leaf level matrices of stage S_i . Hence, $T_\infty(m, i)$ can be recursively defined using the previous analysis as: $T_\infty(m, i) = \mathcal{O}(\log \log \log n)$ if $i = h$ and $T_\infty(m, i) = \mathcal{O}(\Delta c_i \log n + \log m) + T_\infty(m^{1-\Delta c_i}, i+1)$ if $i < h$. Substituting the values of c_i from equation [1], we get $\Delta c_i = 1 - \alpha^{i+1} - (1 - \alpha^i) = \alpha^i(1 - \alpha) = \mathcal{O}(\alpha^i)$. We know that m starts with n and decreases by a factor of $n^{\Delta c_i}$ for every stage. Hence, $m = n^{1-c_{i-1}} = n^{\alpha^i}$, which implies that $\log m = \alpha^i \log n$.

By unrolling the recursion and using the fact that α^i is a geometric series and $\alpha < 1$, we compute the total span as $T_\infty(n, 0) = \sum_{i=0}^{h-1} \alpha^i \log n + T_\infty(n, h) = \mathcal{O}(\log n)$.

[Space.] The total space is dominated by the space used by the data structures. There are n^w matrices at the leaf level for each of the input matrices U and V . Each such matrix is of size 1×1 . Hence, space usage is $\mathcal{O}(n^w)$. \square

[Step 2. Compute the Output Matrices.] The logic used to compute the output matrices is very similar to that used to compute the input matrices. A cell of the output matrix (X) at a node of the recursion depends on at most four cells of the corresponding output matrices at its child nodes. This implies that each cell in the output matrix at the root node depends on at most $4^{\log n} = n^2$ cells in the corresponding output matrices at the leaf nodes. If we were to compute all output cells of all output matrices at all nodes, the total work would explode to $\mathcal{O}(n^w \times n^2) = \mathcal{O}(n^{w+2})$. To keep the work very close to $\mathcal{O}(n^w)$, we split the entire recursion tree into stages. We then use single-point computation of output cells in stages.

In this step, we compute the output matrix of all nodes in the recursion tree. The phase consists of $h+1$ sequential stages: $0, 1, \dots, h$, as shown in Figure [3] (replace c_i 's with d_i 's), such that the

height of stage i is fixed at $d_i \log n$, where h and d_i are given below:

$$d_i = \begin{cases} 0 & \text{if } i = -1, \\ 1 - \beta^{i+1} & \text{if } i \in [0, h-1], \text{ such that } w = \log_2 7, \beta = \frac{4-w}{2}, \text{ and } h = \log_{\frac{2}{4-w}} \frac{2 \log n}{\log \log \log n} - 1. \\ 1 & \text{if } i = h. \end{cases} \quad (2)$$

In step 1, we computed the input matrices in the top-down fashion. In contrast, in this step, we construct the output matrices at different recursion levels in a bottom-up fashion. In other words, we compute the last stage S_h first, then stage S_{h-1} , and so on until stage S_0 . At stage S_0 , the final output matrix X will be of size $n \times n$.

Please refer to Figure (the COMPUTE-OUTPUT-MATRICES algorithm) for computing the output matrix for all leaf nodes in all stages. We first descend the tree until we reach the last stage S_h . We know that all cells in the leaf nodes of this stage (or the recursion tree) already store the input U and V matrices using which we can compute the output matrices at that level. Using these output matrices at the leaf level of S_h , we compute the output matrices at the root level of S_h (or the leaf level of S_{h-1}). Using these matrices at the leaf level of S_{h-1} , we compute the output matrices at the root level of S_{h-1} . This process continues until we reach the root level of S_0 (or the root node of the entire recursion tree), which is the desired matrix product.

The way an output matrix at the root level is computed from the output matrices at the leaf level of stage S_i is as follows. For all cells in the output matrix at the root level, we invoke the function COMPUTE-OUTPUT-CELL. This function computes the final value at that cell.

The working of the COMPUTE-OUTPUT-CELL algorithm is shown in Figure 4 (right). The figure shows the way in which a highlighted cell in the output matrix at the root level is computed. The highlighted cell belongs to the first quadrant (11) of the output matrix. From the logic of the Strassen's MM algorithm, we know that the first quadrant of the output matrix of a node is computed by adding the output matrices of the 0th, 3rd, and 6th child nodes and subtracting that of the 4th child node. We can compute the highlighted cell from four cells in the next level in $\mathcal{O}(1)$ time. Now, consider the output cell in the 3rd child node of the root node. This cell belongs to the third quadrant (22) of that matrix. Again, from the logic of the Strassen's MM algorithm, the fourth quadrant of the matrix is computed by adding the output matrices of the 0th, 2nd, and 5th child nodes and subtracting that of the 1st child node. We continue the process until we reach the leaf level of that stage.

Once cells in the output matrices at the root level of a stage S_i are computed, the algorithm will proceed to computing the cells in the output matrices at the root level of stage S_{i-1} recursively until we reach the root node of the entire recursion tree.

Lemma 2.3. *COMPUTE-OUTPUT-MATRICES has a complexity of $\mathcal{O}(n^w \log \log n)$ work, $\mathcal{O}(\log n)$ span, and $\mathcal{O}(n^w)$ space.*

Proof. [Work.] We compute the output matrix in $h+1$ stages S_0, S_1, \dots, S_h , where $h = \log_{w-1}(\log n / \log \log \log n)$. Suppose W_i defines the work done at stage S_i . We first come up with a generic formula for W_i . We use a direct proof to show that $W_i = \mathcal{O}(n^w)$, which implies that the total work is $\sum_{i=0}^h W_i = \mathcal{O}(n^w \log \log n)$.

We compute W_i for $i \in [0, h-1]$. Each output matrix at the root level of stage S_i is constructed from the output matrices at the leaf level of the stage. All cells in all output matrices at the root level of stage S_i are computed in parallel. #Nodes at the root level of stage S_i is $7^{d_{i-1} \log n} = n^{w d_{i-1}}$. Each such matrix has $(n^{1-d_{i-1}})^2 = n^{2(1-d_{i-1})}$ cells. Each cell at a recursion level ℓ depends on at most 4 output cells in recursion level $\ell+1$. Hence, each cell in a matrix at the root level of stage

S_i depends on $\mathcal{O}(4^{(d_i-d_{i-1})\log n}) = \mathcal{O}(n^{2(d_i-d_{i-1})})$ cells in a matrix at the leaf level of the stage. Hence, $W_i = \mathcal{O}(n^{wd_{i-1}} n^{2(1-d_{i-1})} n^{2(d_i-d_{i-1})}) = \mathcal{O}(n^{2d_i+(w-4)d_{i-1}+2})$.

To show that $W_i = \mathcal{O}(n^w)$ for all $i \in [0, h-1]$, it is enough to prove that $2d_i+(w-4)d_{i-1}+2 = w$. We substitute the values of d_i and d_{i-1} from equation [2] and simplify to get: $2d_i+(w-4)d_{i-1}+2 = w$.

We now compute W_h . We see that $W_h = \mathcal{O}(n^{2d_h+(w-4)d_{h-1}+2})$ using the equation aforementioned. We substitute the values of d_h , d_{h-1} , h , and β from equation [2] and simplify to obtain $W_h = \mathcal{O}(n^w \cdot n^{(4-w)(1-d_{h-1})}) = \mathcal{O}(n^w \cdot n^{2\beta^{h+1}}) = \mathcal{O}(n^w \log \log n)$.

Combining the work of the first h stages and the last stage, we get $T_1(n) = \sum_{i=0}^{h-1} W_i + W_h = \mathcal{O}(n^w \log \log n)$.

[Span.] Let $T_\infty(m, i)$ denote the span of the COMPUTE-OUTPUT-MATRICES algorithm starting from stage i where a matrix at the leaf level is of size $m \times m$. We give a recursive formula to compute $T_\infty(m, i)$. Then, the total span for the algorithm is $T_\infty(n, 0)$.

Consider the COMPUTE-OUTPUT-MATRICES algorithm. Let $\Delta d_i = d_{i+1} - d_i$. #Nodes at the leaf level of stage S_i is $7^{\Delta d_i \log n}$. Launching these nodes in parallel (line [2]) incurs a span of $\mathcal{O}(\Delta d_i w \log n)$. A matrix at the leaf level will be of size $m/(2^{\Delta d_i \log n}) \times m/(2^{\Delta d_i \log n})$. Spawning COMPUTE-OUTPUT-CELL function for all cells (lines [6], [7]) incur a span of $\mathcal{O}(2 \log m - 2\Delta d_i \log n)$. Executing the COMPUTE-OUTPUT-CELL algorithm incurs $\mathcal{O}(\Delta d_i \log n)$ span. Adding all these spans give us $\mathcal{O}(\Delta d_i \log n + \log m)$.

The span of stage S_i recursively depends upon the span of stage S_{i+1} as the output matrices at the leaf level of stage S_{i+1} are constructed from the leaf level matrices of stage S_i . Hence, $T_\infty(m, i)$ can be recursively defined using the previous analysis as: $T_\infty(m, i) = \mathcal{O}(\log \log \log n)$ if $i = h$ and $T_\infty(m, i) = \mathcal{O}(\Delta d_i \log n + \log m) + T_\infty(m^{1-\Delta d_i}, i+1)$ if $i < h$. Substituting the values of d_i from equation [2], we get $\Delta d_i = 1 - \beta^{i+2} - (1 - \beta^{i+1}) = \beta^{i+1}(1 - \beta) = \mathcal{O}(\beta^i)$. We know that m starts with n and decreases by a factor of $n^{\Delta d_i}$ for every stage. Hence, $m = n^{1-d_{i-1}} = n^{\beta^i}$, which implies that $\log m = \beta^i \log n$.

By unrolling the recursion and using the fact that β^i is a geometric series and $\beta < 1$, we compute the total span as $T_\infty(n, 0) = \sum_{i=0}^{h-1} \beta^i \log n + T_\infty(n, h) = \mathcal{O}(\log n)$.

[Space.] Using a similar analysis as given in Lemma [2.3], space usage is $\mathcal{O}(n^w)$. □

Theorem 2.4. *The STRASSEN-S MM algorithm has a complexity of $\mathcal{O}(n^w \log \log n)$ work, $\mathcal{O}(\log n)$ span, and $\mathcal{O}(n^w)$ space.*

Proof. The theorem follows from lemmas [2.2] and [2.3]. □

Strassen-S-Adaptive MM. We design a parallel Strassen's MM algorithm STRASSEN-S-ADAPTIVE with space-span trade-off, which for any given s amount of space in the range $[n^2, n^w]$, achieves the optimal span for that space and performing work very close to $\mathcal{O}(n^w)$. Suppose we are given the input matrices U and V . We need to compute the output matrix X using space $s \in [n^2, n^w]$. Then, the algorithm works as follows. Observe that there are $\log n$ levels in the recursion tree of the Strassen's MM algorithm. We split the entire recursion tree, at level t , into two parts: the top part and the bottom part. The threshold level t depends on the value s . We execute the classical sequential Strassen's MM algorithm in the top portion of the recursion tree and the STRASSEN-S algorithm in the bottom portion of the tree.

Theorem 2.5. *The STRASSEN-S-ADAPTIVE MM algorithm has a complexity of $\mathcal{O}(n^w \log \log n)$ work and $\mathcal{O}((n^w/s) \log n)$ span, given $\Theta(s)$ amount of space.*

Proof. Let $T_1(n, s)$, $T_\infty(n, s)$, and $S_\infty(n, s)$ denote work, span, and space of STRASSEN-S-TUNABLE. Let $T_1(n)$, $T_\infty(n)$, and $S_\infty(n)$ denote work, span, and space of STRASSEN-S. Note that STRASSEN-S does not have a space parameter.

We run the sequential Strassen's MM algorithm for the first t levels of the recursion tree. At level t , the size of the matrices is $n/2^t \times n/2^t$ and the number of matrices is 7^t . We have

$$\begin{aligned} T_1(n, s) &= 7T_1(n/2, s) + \mathcal{O}(n^2) = \dots = 7^t T_1(n/2^t) + \mathcal{O}(7^t n^2) \\ T_\infty(n, s) &= 7T_\infty(n/2, s) + \mathcal{O}(\log n) = \dots = 7^t T_\infty(n/2^t) + \mathcal{O}(7^t \log n) \\ S_\infty(n, s) &= S_\infty(n/2, s) + \mathcal{O}(n^2) = \dots = S_\infty(n/2^t) + \mathcal{O}(n^2) = \mathcal{O}((n/2^t)^w + n^2) \end{aligned}$$

Equating the total space usage with s , we get $s = \Theta((n/2^t)^w + n^2)$. We simplify this expression to get the two expressions $n/2^t = \Theta((s - n^2)^{1/w})$ and $7^t = \Theta(n^w/(s - n^2))$. Substituting the two expressions in the span and work equations, we have

$$\begin{aligned} T_\infty(n/2^t) &= \mathcal{O}(\log^2(n/2^t)) = \mathcal{O}(\log^2 s) = \mathcal{O}(\log n) \\ T_\infty(n, s) &= 7^t T_\infty(n/2^t) + \mathcal{O}(7^t \log n) = \mathcal{O}((n^w/s) \log n) \\ T_1(n/2^t) &= (n/2^t)^w \log \log(n/2^t) \\ T_1(n, s) &= 7^t \cdot T_1(n/2^t) + \mathcal{O}(7^t n^2) = \mathcal{O}(n^w \log \log n) \end{aligned}$$

□

Strassen-W-Adaptive MM. We can show that by using the sequential 1-way Strassen's MM until recursion level t (depends on s units of space) and then switching to the 7-way Strassen's MM instead of STRASSEN-S, we can achieve work bound the same as that of the classical Strassen's MM, but the span increases by an extra $\mathcal{O}(\log n)$ factor.

Theorem 2.6. *The STRASSEN-W-ADAPTIVE MM algorithm has a complexity of $\mathcal{O}(n^w)$ work and $\mathcal{O}((n^w/s) \log^2 n)$ span, given $\Theta(s)$ amount of space.*

Proof. The proof is similar to that of Theorem 2.5 except that $T_\infty(n/2^t) = \mathcal{O}(\log^2 n)$ and $T_1(n/2^t) = (n/2^t)^w$. □

Corollary 2.6.1. *With $s = \Theta(n^2)$ units of space, (i) STRASSEN-S-ADAPTIVE has a complexity of $\mathcal{O}(n^w \log \log n)$ work and $\mathcal{O}(n^{w-2} \log n)$ span and (ii) STRASSEN-W-ADAPTIVE has a complexity of $\mathcal{O}(n^w)$ work and $\mathcal{O}(n^{w-2} \log^2 n)$ span.*

Strassen-W MM. We ask the following question. If the work is bounded by $\mathcal{O}(n^w)$, what is the best span achievable by a parallel Strassen's MM algorithm? It turns out that with $\mathcal{O}(n^w)$ work bound one can achieve $\mathcal{O}(\log n \log \log n)$ span.

We split the entire recursion tree, at level $\log(n/(\log \log n)^{1/(w-2)})$, into two parts: the top and the bottom parts. We execute the STRASSEN-S in the top portion and the quadratic space STRASSEN-W-ADAPTIVE algorithm in the bottom portion.

Theorem 2.7. *The STRASSEN-W MM algorithm has a complexity of $\mathcal{O}(n^w)$ work, $\mathcal{O}(\log n \cdot \log \log \log n)$ span, and $\mathcal{O}(n^w / \log \log n)$ space.*

Proof. There are m matrices of size $(n/m) \times (n/m)$ at the switching level t , where $m = n/(\log \log n)^{1/(w-2)}$. At each node at the threshold level, we add two matrices with two for loops. Adding two matrices has $\mathcal{O}(\log(n/m))$ span and $\mathcal{O}((n/m)^2)$ work. From Theorem 2.4, the span and work for STRASSEN-S are $\mathcal{O}(\log m \cdot \log(n/m))$ and $\mathcal{O}((m^w \log \log m) \cdot (n/m)^2)$ respectively. The span and work for

STRASSEN-W-TUNABLE in m^w leaves are $\mathcal{O}((n/m)^{(w-2)} \log^2(n/m))$ and $\mathcal{O}(m^w(n/m)^w)$ respectively. Combining the span from STRASSEN-S and STRASSEN-W-TUNABLE, we get the expression for span for the hybrid algorithm as follows.

When $m = n/(\log \log n)^{1/(w-2)}$, then $n/m = (\log \log n)^{1/(w-2)}$. We compute span as

$$\begin{aligned} T_\infty(n) &= \mathcal{O}(\log m \cdot \log(n/m) + (n/m)^{w-2} \log^2(n/m)) \\ &= \mathcal{O}(((\log n - 1/(w-2)) \log \log \log n)(1/(w-2)) \log \log \log n + ((1/(w-2)) \log \log \log n)^2 \log \log n) \\ &= \mathcal{O}(\log n \cdot \log \log \log n). \end{aligned}$$

Combining work from both STRASSEN-S and STRASSEN-W-TUNABLE, we get the total work as:

$$\begin{aligned} T_1(n) &= \mathcal{O}((m^w \log \log m)(n/m)^2 + m^w(n/m)^w) \\ &= \mathcal{O}\left((n^w/(\log \log n)^{\frac{w}{w-2}})(\log \log n)^{\frac{2}{w-2}} + n^w\right) = \mathcal{O}(n^w). \end{aligned}$$

We use $s = \Theta(n^w / \log \log n)$ space for the whole algorithm.

$$S_\infty(n) = \mathcal{O}(m^w(n/m)^2) = \mathcal{O}\left((n^w/(\log \log n)^{\frac{w}{w-2}})(\log \log n)^{\frac{2}{w-2}}\right) = \mathcal{O}(n^w / \log \log n).$$

□

Strassen-like MM Algorithms. Let recursive algorithm ALG multiply two input matrices U and V of size $n \times n$ and produce output matrix X , that is $X = U \cdot V$. ALG divides U into $m \times m$ blocks each of size $(n/m) \times (n/m)$. Similarly, ALG divides the other input matrix V and output matrix X . Suppose that algorithm ALG has R recursive calls in each level of recursion. Then ALG creates R temporary matrices each for both input matrices and the output matrix. In particular, the computation of ALG in each level of recursion is as follows. For each $r = 0, 1, \dots, R-1$, Here $A^{(r)}$, $B^{(r)}$ and $C^{(r)}$ represent temporary matrices.

$$\begin{aligned} A^{(r)} &\leftarrow 0; B_{(r)} \leftarrow 0 \\ A^{(r)} &\leftarrow A^{(r)} + \alpha_{i,k}^{(r)} U_{i,k} \text{ for } i, k = 0, 1, \dots, m-1 \\ B^{(r)} &\leftarrow B^{(r)} + \beta_{k,j}^{(r)} V_{k,j} \text{ for } k, j = 0, 1, \dots, m-1 \\ C^{(r)} &\leftarrow A^{(r)} \cdot B^{(r)} \\ X_{i,j} &\leftarrow X_{i,j} + \gamma_{i,j}^{(r)} C^{(r)} \text{ for } i, j = 0, 1, \dots, m-1. \end{aligned}$$

We call ALG a Strassen-like MM algorithm [32]. Sequential algorithm ALG does $\mathcal{O}(n^w)$ work where $w = \log_m R$. In Strassen, $m = 2$ and $R = 7$.

It is important to observe that the approaches used by STRASSEN-S, STRASSEN-W, STRASSEN-S-ADAPTIVE, and STRASSEN-W-ADAPTIVE MM algorithms apply to all Strassen-like MM algorithms.

Lower Bounds. We give the following lower bound of any parallel version of Strassen's MM algorithm using s units of space in the binary-forking model.

Theorem 2.8. *Let A be a parallel version of Strassen's MM algorithm which uses s units of extra space. Then A 's span is $\Omega(\max(n^w/s, \log n))$ in the binary-forking model.*

Proof. We consider the binary-forking model without atomics [6]. In this model, every binary operation (addition, subtraction, multiplication and division) is associated with a memory location. Specifically, the output of such binary operation needs to be written to a memory location. Concurrent reads from a memory location are allowed, while concurrent writes are not.

Let A be a matrix multiplication serial algorithm with $\Theta(n^w)$ work where $w \geq 2$. Let $T_\infty(n, w, s)$ denote the span of any parallel algorithm B that parallelize the serial algorithm A using s units of space. We do not make any restriction on the number of processors used. We remark that only heap space is counted in s (all our previous algorithms also allocate space from heap memory).

We get the first lower bound as follows. When we use s units of memory locations, then from the pigeonhole principle, there must exist a memory location that is subjected to $\Theta(n^w/s)$ write operations. As concurrent writes to a memory location are not allowed, the following lower bound holds: $T_\infty(n, w, s) = \Omega(n^w/s)$.

We get the second lower bound from a more general PRAM CREW model. Thus, it holds for a more restricted binary-forking model. When we have an unbounded #memory locations and unbounded #processors, the following lower bound [29] holds for the computation of $x_1 + x_2 + \dots + x_n$ (array-sum) where every $s_i \in \{0, 1\}$: $T_\infty(n, w, \infty) = \Omega(\log n)$.

As matrix multiplication is at least as hard as array-sum, it must have $\Omega(\log n)$ span. Combining both the lower bounds, we get the following: $T_\infty(n, w, s) = \Omega(\max(n^w/s, \log n))$. \square

3 Sorting

Sorting a set of n numbers is one of the most fundamental problems in computer science. Several efficient sorting algorithms exist [31, 29, 2, 21, 8, 28]. In this section, we consider work-optimal (i.e., performing $\mathcal{O}(n \log n)$ work) comparison sort algorithms in the binary-forking model. Our goal is to design a work-optimal comparison sort parallel algorithm with optimal span of $\mathcal{O}(\log n)$ in the binary-forking model without atomics.

Cole's pipelined merge sort [16] has optimal span of $\mathcal{O}(\log n)$ in the PRAM model. However, it achieves a span of $\mathcal{O}(\log^2 n)$ when analyzed in the binary-forking model. Blelloch et al.'s randomized sorting algorithm [7] achieves $\mathcal{O}(n \log n)$ work and $\mathcal{O}(\log^{1.5} n)$ span, both bounds w.h.p. in n . Cole and Ramachandran's sample-partition-merge-sort algorithm [17] is based on multi-way merge and sample sort [23] and has a low span of $\mathcal{O}(\log n \log \log n)$. Blelloch et al.'s randomized sorting algorithm [6] is based on sample sort and has $\mathcal{O}(n \log n)$ expected work and $\mathcal{O}(\log n)$ span w.h.p. in n , but it makes use of the test-and-set instruction and as a result cannot be straightforwardly adapted to the binary-forking model without atomics.

Avoiding TS and Retaining Expected Work. Our primary strategy for avoiding atomic TS operations is to use extra space and randomize write locations, rendering collisions unlikely. In Blelloch et al.'s [6] randomized sorting algorithm, n elements are partitioned by repeatedly attempting to place them into random locations in $\Theta(n)$. Since there is a $\Theta(1)$ likelihood of success on every write attempt, it is enough for each item to try $\Theta(\log n)$ times to find an unoccupied cell with high probability; the span of this process is $\Theta(\log n)$. This method cannot be applied in $\Theta(\log n)$ span without TS.

Instead of using test-and-set, we use $2n \log n$ extra space. Each element randomly picks $\log n$ cells and tries to put itself into all of them in parallel. If multiple elements try to put themselves into the same cell, they collide and only one succeeds. For each element e set a memory location f_e that holds the cell where e should go. After the attempt is completed, if a cell at memory location i holds element e , then it tries to write its index i into f_e . If element e succeeds in putting itself into

at least one of the $\log n$ randomly chosen cells, then f_e will w.h.p. in n hold an index where element e is stored. After all memory locations f_e are populated, compact the elements into contiguous locations using the content of the memory locations.

In the algorithm by Blelloch et al., every element tries to put itself into a memory location $\log n$ times serially where the success probability of each attempt is $1/2$. In our algorithm, each element tries in parallel $\log n$ times where the success probability of each attempt is $1/2$. Hence, the work remains the same, that is $\Theta(n \log n)$. The span is also the same, that is $\Theta(\log n)$. We increase the success probability for an item to land in an unoccupied place by using $\Theta(\log n)$ times extra space instead of trying $\log n$ times.

Improving Work Bound from Expectation to High Probability. Here we present a sorting algorithm which takes optimal work and span, both with high probability in n . We achieve this by first sorting all but $o(n)$ elements of an array in optimal work and span, w.h.p. in n and later integrating the leftover $o(n)$ elements with the sorted $n - o(n)$ elements. We also give a space-adaptive sorting algorithm that achieves near-optimal work and span w.h.p. in n for any given amount of space s .

First we show that all but $o(n)$ elements can be sorted in optimal work and span w.h.p. in n by using ALMOST-SORT, a very shallow divide and conquer algorithm of similar design as Blelloch et al.'s, in that it recursively partitions elements into buckets. Although Blelloch et al.'s algorithm gives high probability at all levels of recursion, it is only high probability in the size of arrays passed to the recursive steps, which becomes subpolynomial in n (the original input size) at the lowest levels of recursion. This causes the related work bound to hold only in expectation. In ALMOST-SORT we only use recursive partitioning to depth $\log \log \log n$, after which we apply the Cole-Ramachandran deterministic sorting algorithm. This keeps the array size large enough to achieve a high probability bound on work at all levels of recursion.

Partitioning an array of size n in ALMOST-SORT is done in two steps. First we sample a set of $\sqrt{n} \log^3 n$ potential pivots from the array, sort this subset using n^ϵ -WAY-SORT (see Figure 11), and select \sqrt{n} uniformly spaced elements from it to serve as pivots. Second we attempt to place every element in the array into two randomly chosen locations in their respective buckets. If two elements attempt to write to the same location, then both attempts fail. If both of a single element's attempts to be written into its partition fail, then that element is dropped and will not appear in the next recursive step. The pseudocode listing for ALMOST-SORT is given in figure 5.

We then give a method for recombining the $o(n)$ unsorted elements (those that were dropped by ALMOST-SORT) into the sorted array in $\mathcal{O}(n \log n)$ work and $\mathcal{O}(\log n)$ span, w.h.p. in n . This method also depends on using extra space to convert the problem of merging a small unsorted array with a large sorted array into a series of bins-and-balls type scenarios. The union of ALMOST-SORT with this integration method gives us the FULL-SORT algorithm, which w.h.p. in n sorts an array of size n in optimal $\mathcal{O}(n \log n)$ work and $\Theta(\log n)$ span, using a $\Theta\left(\frac{\log n \log \log \log n}{\log \log n}\right)$ factor extra space.

Our main results are given in theorem 3.3 and theorem 3.6.

Lemma 3.1. *Partitioning an array of size n into \sqrt{n} blocks with oversampling factor $\log^3 n$ will produce no blocks with size falling outside of an $\epsilon = 1/\log^2 \log n$ factor of \sqrt{n} , with probability at least $1 - n^{-\log n}$.*

Proof. Let $\langle a_1, \dots, a_n \rangle$ be the input array A as a sorted sequence, from which $q = \sqrt{n} - 1$ pivots are chosen (with repetition) with oversampling factor $s = \log^3 n$. If one of the resulting blocks contains at least (resp. less than) k elements, then there must be a subsequence $\langle a_j, \dots, a_{j+k-1} \rangle$ from which less than (resp. at least) s elements were sampled.

Let $\{x_1, \dots, x_{(q+1)s}\}$ be the elements sampled from A , and define the Bernoulli variables $X_i = 1$ if $x_i \in \langle a_j, \dots, a_{j+k-1} \rangle$, $X_i = 0$ otherwise, and $X = \sum_{i=1}^{(q+1)s} X_i$ for some arbitrary fixed j . Note that $E(X) = (q+1)sE(X_i) = \frac{sk}{\sqrt{n}}$. Let b_i be the sizes of the blocks of the partition of A with pivots $x_s, x_{2s}, \dots, x_{qs}$: we then have $P(\max_i b_i \geq k) \leq P(\exists j \ X < s)$ (resp. $P(\min_i b_i < k) \leq P(\exists j \ X \geq s)$).

We handle the upper tail by setting $k = (1 + \varepsilon)\sqrt{n}$ gives us $E(X) = (1 + \varepsilon)s$, from which we can upper bound the probability of an unexpectedly large partition block by using a Chernoff bound for Poisson Binomial distributions:

$$\begin{aligned}
P\left(\max_i b_i \geq (1 + \varepsilon)\sqrt{n}\right) &\leq P(\exists j \ X < s) \\
&\leq nP(X < (1 - \varepsilon/2)E(X)) && \left(1 - \varepsilon/2 > \frac{1}{1 + \varepsilon}\right) \\
&\leq n \exp(-\varepsilon^2 s/8) \\
&= n \exp\left(\frac{-\log^3 n}{8 \log^4 \log n}\right) \\
&< (1/2)n^{-\log n} && (\text{for } n \geq 4).
\end{aligned}$$

Similar to the treatment of upper tail, for the lower tail we set $k = (1 - \varepsilon)\sqrt{n}$, which gives $E(X) = (1 - \varepsilon)s$. Applying the Chernoff bound for lower tails then shows that

$$\begin{aligned}
P\left(\min_i b_i < (1 - \varepsilon)\sqrt{n}\right) &\leq P(\exists j \ X \geq s) \\
&\leq nP(X > (1 + \varepsilon)E(X)) && \left(1 + \varepsilon < \frac{1}{1 - \varepsilon}\right) \\
&\leq n \exp(-\varepsilon^2(1 - \varepsilon)s/4) \\
&\leq n \exp\left(\frac{-\log^3 n}{8 \log^4 \log n}\right) \\
&< (1/2)n^{-\log n} && (\text{for } n \geq 4).
\end{aligned}$$

Combining the upper and lower tails immediately gives a bound on the probability of a sufficiently regular partition, $P(\forall i \ (1 - \varepsilon)\sqrt{n} \leq b_i \leq (1 + \varepsilon)\sqrt{n}) > 1 - n^{-\log n}$ for $n \geq 4$. \square

Corollary 3.1.1. *Lemma [3.1](#) continues to hold with high probability in n for arrays of size $m < n$ being partitioned into \sqrt{m} blocks with fixed $\varepsilon = 1/\log^2 \log n$ and oversampling factor $\log^3 n$.*

Proof. Decreasing the number of elements to be partitioned while holding the oversampling factor constant can only increase the likelihood of a suitably uniform partition. Therefore, if Lemma [3.1](#) guarantees high probability of a uniform partition for an array of size n , then that guarantee will continue to if we lower the number of elements in the array to $m < n$. \square

Lemma 3.2. *Attempting to place n elements in $nm = \Omega\left(\frac{n \log n}{\log \log n}\right)$ space will with high probability in n take $\mathcal{O}(\log n)$ span, $\Theta(n)$ work, and result in $\Theta\left(\frac{n}{m}\right)$ collisions.*

Proof. [Collisions.] The likelihood of any given element experiencing a collision is less than $1/m$, and therefore the fact that $\Theta\left(\frac{n}{m}\right)$ collisions will occur with high probability follows directly from a pair of Chernoff bounds, $P(X > 2E(X)) < e^{-E(X)/3}$ and $P(X < E(X)/2) < e^{-E(X)/8}$, where we take X to be the number of collisions.

[Work and Span.] We assume that k collisions at the same location will produce a span of $\mathcal{O}(k)$. Consider the set of $\Theta(n/m)$ elements which were not placed due to a collision: they are randomly distributed throughout the $\Theta(n)$ locations that were selected by at least one element, so with probability greater than $1 - \exp(-\Theta(n/m^2))$ there are $\Theta(n/m^2)$ of them that collided with one another. By induction, there are $\Theta(n/m^d)$ locations where d elements collided, w.h.p. in n for all d such that $n/m^d = \Omega(\log n)$. The remaining $\mathcal{O}(\log n)$ locations where $\Omega(\log_m n)$ elements collided contribute, at most, $\mathcal{O}(\log^2 n)$ work and $\mathcal{O}(\log n)$ span.

The span is therefore dominated by the time taken to spawn n processes, $T_\infty(n) = \Theta(\log n)$, and work is given by $T_1(n) = \Theta\left(n + n \sum_{d=1}^{\log_m n} \frac{d^2}{m^d}\right) = \Theta(n)$, where we have used the identity $\sum_{j=1}^{\infty} j^2 x^j = \frac{x(1+x)}{(1-x)^3} \sim x$ for x small. □

Theorem 3.3. *The ALMOST-SORT algorithm (Figure 5) takes $\mathcal{O}(n \log n)$ work w.h.p. in n , $\Theta(\log n)$ span w.h.p. in n , and $\Theta\left(\frac{n \log n \log \log \log n}{\log \log n}\right)$ space to sort all but $\Theta\left(\frac{n \log^2 \log n}{\log^2 n \log \log \log n}\right)$ elements w.h.p. in n of an array of size n .*

Proof. [Work.] Let n be the size of the array passed to ALMOST-SORT in the initial call, and n_c be the size of the array passed at some point in the recursion tree. We terminate recursion at depth $\log \log \log n$, at which point the COLE-RAMACHANDRAN sorting algorithm is applied, taking $\mathcal{O}(n_c \log n_c)$ work. All higher levels of recursion sample $\sqrt{n_c}$ pivots with oversampling factor $\log^3 n$, which are sorted using n^ε -WAY-SORT with $\varepsilon = 1/2$, after which ALMOST-SORT is called on the $\sqrt{n_c}$ partition blocks, each of which will be no larger than $(1 + 1/\log^2 \log n)\sqrt{n_c}$. For every call to ALMOST-SORT there is $\mathcal{O}\left(n_c \frac{\log n \log \log \log n}{\log \log n}\right)$ work done on prefix sums and $\mathcal{O}(n_c \log n_c)$ work done on the binary searches by which elements find which partition block to be placed in. Thus the work is $T_1(n_c, n) = \mathcal{O}(n_c \log n_c)$ when depth is greater than $\log \log \log n$, and $T_1(n_c, n) = \mathcal{O}\left((\sqrt{n_c} \log^3 n)^{3/2} + n_c \frac{\log n \log \log \log n}{\log \log n} + n_c \log n_c\right) + \sqrt{n_c} T_1((1 + 1/\log^2 \log n)\sqrt{n_c}, n)$ otherwise.

Note that ALMOST-SORT is called $\mathcal{O}(n \log \log \log n)$ times, and the likelihood of a suitably uniform partition during any particular call of ALMOST-SORT is bounded below by $1 - n^{-\log n}$ (via corollary 3.1.1), so all partitions will be suitably uniform with high probability in n .

We can upper bound w.h.p. the size of arrays n_c passed to depth d with $(1 + 1/\log^2 \log n)^d n^{2^{-d}} \leq (1 + 1/\log^2 \log n)^{\log_{3/2} \log n} n^{2^{-d}} \lesssim \exp(1/\log \log n) n^{2^{-d}} \sim (1 + 1/\log \log n) n^{2^{-d}}$. Therefore in order to be able to apply Lemma 3.2 to the array partitioning which occurs at depth d , we will need $\Omega\left(n^{2^{-d}} \frac{\log n}{\log \log n}\right)$ space per function instance.

There are $n^{1-2^{-d}}$ instances of n^ε -WAY-SORT being called at depth d , which cumulatively take $\mathcal{O}\left(\frac{n}{n_c} (\sqrt{n_c} \log^3 n)^{3/2}\right) = \mathcal{O}(n)$ work for⁴ $n_c = \Omega(\log^{18} n)$. The size of arrays being processed at the lowest level of recursion is $\mathcal{O}(n^{1/\log \log n})$, so the total work done by the n/n_c calls to COLE-RAMACHANDRAN is $\mathcal{O}\left(\frac{n \log n}{\log \log n}\right)$. Summing across all levels, we have a total of $\mathcal{O}(n \log n)$ work (primarily from the binary searches that occur during partitioning).

[Span.] Following a similar line of reasoning as was used for the work bound, we note that n^ε -WAY-SORT (still with $\varepsilon = 1/2$) has span $\mathcal{O}(\log n_c)$, and COLE-RAMACHANDRAN has span

⁴Note that when $d < \log \log \log n$, we have $n_c^{-1/4} < (n^{1/\log \log n})^{-1/4} < \log^{-c} n$ for all positive constants c .

ALMOST-SORT($n_c, n, l_0, d, m, B, C, D$)

n_c : size of array to sort (only $B[l_0, \dots, l_0 + n_c - 1]$ is occupied)

n : size of the array at the highest level of recursion

l_0 : location where the array to sort begins

d : depth of current call in recursion tree

m : multiple of extra memory to use

$B[l_0, \dots, l_0 + n_c m - 1]$: contains array to be sorted

$C[l_0, \dots, l_0 + n_c m - 1]$: ancillary space

$D[l_0, \dots, l_0 + n_c m - 1]$: where prefix sums will be stored (for indexing)

1. **if** $d \geq \log \log \log n$ **do** COLE-RAMACHANDRAN($B[l_0, l_0 + n_c - 1]$) and **return**
 { Partitioning B ; choose $\sqrt{n_c} - 1$ pivots with oversampling factor s }
2. $s \leftarrow \log^3 n_c$
3. $P \leftarrow$ sample with repetition $(\sqrt{n_c} + 1)s$ elements from $B[l_0, \dots, l_0 + n_c - 1]$
4. $P \leftarrow n^\epsilon$ -WAY-SORT($P, 1/2$)
5. $P[0] \leftarrow -\infty$; $P[(\sqrt{n_c} + 1)s] \leftarrow +\infty$
 { Nondeterministically partition B into C }
6. **parallel for** $a \in B[l_0, \dots, l_0 + n_c - 1]$ **do**
7. Find some i s.t. $P[i \cdot s] \leq a \leq P[(i + 1)s]$
8. Choose a random number $j \in [0, \dots, m\sqrt{n_c} - 1]$
9. Attempt to assign $C[im\sqrt{n_c} + j] \leftarrow a$; in case of collision do nothing
10. **parallel for** $i \leftarrow l_0$ **to** $l_0 + n_c - 1$ **do** $B[i] \leftarrow \text{null}$
 { Compact the partitions of C into B }
11. **parallel for** $i \leftarrow 0$ **to** $\sqrt{n_c} - 1$ **do**
12. $\text{lo} \leftarrow l_0 + im\sqrt{n_c}$; $\text{hi} \leftarrow l_0 + (i + 1)m\sqrt{n_c} - 1$
13. $D[\text{lo}, \dots, \text{hi}] \leftarrow \text{INDICATOR-PREFIX-SUM}(C[\text{lo}, \dots, \text{hi}])$
14. **parallel for** $j \leftarrow \text{lo}$ **to** hi **do**
15. $D[j] \leftarrow D[j] + im\sqrt{n_c}$
16. **parallel for** $i \in [l_0, \dots, l_0 + n_c m - 1]$ **do**
17. **if** $C[i]$ is not null **do**
18. $B[D[i]] \leftarrow C[i]$
19. $C[i] \leftarrow \text{null}$
 { Divide and Conquer }
20. **parallel for** $i \leftarrow 0$ **to** $\sqrt{n_c} - 1$ **do**
21. ALMOST-SORT($\sqrt{n_c}, n, l_0 + im\sqrt{n_c}, d + 1, m, B, C, D$)
 { Compact the final result }
22. $C \leftarrow B$; $D \leftarrow \text{INDICATOR-PREFIX-SUM}(C)$
23. **parallel for** $i \leftarrow 0$ **to** $n_c m - 1$ **do**
24. **if** $C[i]$ is not null **then**
25. $B[D[i]] \leftarrow C[i]$

Figure 5: Listing for the randomized sorting algorithm ALMOST-SORT. BOOLEAN-PREFIX-SUM(A) is a standard prefix sum taken over the indicator function of A , i.e. the function which is 0 where $A[i]$ is null, and 1 elsewhere. Note that on line 7 we have an implicit binary search across $\sqrt{n_c}$ elements.

$\mathcal{O}(\log n_c \log \log n_c)$, so the recurrence relation for span is

$$T_\infty(n_c, n) \leq \begin{cases} \mathcal{O}(\log n_c \log \log n_c) & \text{if depth} \leq \log \log \log n, \\ \mathcal{O}(\log(n_c \log^3 n)) + T_\infty((1 + 1/\log^2 \log n)\sqrt{n_c}, n) & \text{else,} \end{cases}$$

which, on substituting $n_c = n^{1/\log \log n}$ at lowest depth, is solved by $T_\infty(n, n) = \mathcal{O}(\log n)$.

[Space.] ALMOST-SORT uses $\Theta\left(\frac{n \log n \log \log \log n}{\log \log n}\right)$ space, by construction. All partitioning for a node is done within a chunk of memory inherited from the parent which instantiated it, and the root node begins with $\Theta\left(\frac{n \log n \log \log \log n}{\log \log n}\right)$ memory.

[Unsorted Elements.] At recursive depth d there will be $n^{1-2^{-d}}$ arrays of size $\Theta(n^{2^{-d}})$ from which we are selecting pivots then partitioning. By Lemma 3.2 this will produce $\Theta\left(\frac{n \log \log n}{\log n \log \log \log n}\right)$ collisions, and attempting to place elements twice each will lower this number to $\Theta\left(\frac{n \log^2 \log n}{\log^2 n \log^2 \log \log n}\right)$ elements which failed to find a place in their respective partitions. Summing across all levels of recursion then gives a total of $\Theta\left(\frac{n \log^2 \log n}{\log^2 n \log \log \log n}\right)$ elements left unsorted due to collisions. \square

Lemma 3.4. *With high probability in n , there are no sequences of length $\Theta\left(\frac{\log n}{\log \log n}\right)$ in the array A that have no elements which appear in $B = \text{ALMOST-SORT}(A, n)$.*

Proof. We know with high probability that the bins at all stages of ALMOST-SORT's divide and conquer process are filled no more than within a $(1 + 1/\log \log n)$ factor above or below expectation (see theorem 3.3), and this gives us an upper bound on the likelihood that any given element $a_j \in A$ fails to be included in B :

$$P(a_j \notin B) \leq 1 - \left(1 - \frac{1 + 1/\log \log n}{\frac{\log^2 n \log^2 \log \log n}{\log^2 \log n}}\right)^{\log \log \log n} \sim \frac{\log^2 \log n}{\log^2 n \log \log \log n} \equiv \delta.$$

This result also follows from our high probability bound on the number of collisions which occur during ALMOST-SORT.

Now we can repeat the argument made for Lemma 3.1, interpreting B as having been uniformly randomly sampled⁵ from A , and thereby bounding the likelihood that a subsequence of length ℓ exists in A from which no elements are included in B .

We define the indicator variables Y_i to be 1 if $b_i \in \langle a_j, \dots, a_{j+\ell-1} \rangle$ and 0 otherwise, and $Y \equiv \sum_i Y_i$, from which we have $P(Y = 0) \sim \prod_i P(a_{j+i} \notin B) \leq \delta^\ell$. The probability of a subsequence of length ℓ being contained in A but completely absent from B is then bounded with $P(\exists j Y = 0) \leq nP(Y = 0) \leq n\delta^\ell$, so we find that there is a polynomially small chance of there being a subsequence of length $\ell = \frac{\log n}{\log 1/\delta} = \Theta\left(\frac{\log n}{\log \log n + \log \log \log \log n - \log \log \log n}\right)$ which is entirely absent from B . \square

Recombining Unsorted Elements. We now present a method for merging the unsorted elements $A \setminus B$ with the sorted subarray $B = \text{ALMOST-SORT}(A)$, using a sequence of balls-and-bins type arguments to achieve high probability that no element remains unsorted. The rough outline is as follows: (1) Partition $A \setminus B$ by the elements in B , and sample elements into their respective partition blocks (found via a binary search through B). Use this to index the nonempty partition

⁵Using $P(a_j \notin B) \leq \delta$ allows us to unconditionally upper bound the chance that none of a sequence of elements will end up in B .

block. (2) Partition $A \setminus B$ again by the elements in B , this time distributing space only to the blocks that are known to be nonempty. Use the result to estimate the exact size of each partition block. (3) Partition A by the elements in B , now distributing memory so that every block has $\Omega\left(\frac{\log n}{\log \log n}\right)$ times more space than needed to compactly store the elements that are bound for it. Sample elements of A into their respective partitions, making $\log n$ parallel placement attempts for each element. (4) Remove all but one copy of each element from their associated partition blocks, sort each partition block, and use a prefix sum to compact the entire array. With high probability the result is a sorted copy of the entire array A .

Lemma 3.5. *The unsorted set of elements $A \setminus B$ can with high probability be merged with the sorted subarray $B = \text{ALMOST-SORT}(A)$ in $\mathcal{O}(n \log n)$ work, $\Theta(\log n)$ span, and $\Theta\left(\frac{n \log n \log \log \log n}{\log \log n}\right)$ space.*

Proof. [Step 1.] We start by associating every element of B with a space of size $\Theta\left(\frac{\log n \log \log \log n}{\log \log n}\right)$, and then place elements from $A \setminus B$ into a single random location in their respective buckets, where collisions of k elements in the same location result in $\mathcal{O}(k)$ span and no element being successfully placed. This will take $\mathcal{O}(n)$ work and $\Theta(\log n)$ span, even if the maximum number of collisions occur (which are bounded by lemma 3.4). We then run $\Theta(n)$ parallel prefix sums to find the number of elements in each of the $\Theta(n)$ partition blocks, taking $\Theta\left(\frac{n \log n \log \log \log n}{\log \log n}\right)$ work, after which we index the nonempty blocks by way of a prefix sum taking $\Theta(n)$ work. Partition blocks with no associated elements do not need any space devoted to them, so the number of buckets we have to consider has been reduced from $\Theta(n)$ to $\mathcal{O}\left(\frac{n \log^2 \log n}{\log^2 n \log \log n}\right)$.

[Step 2.] We give each nonempty partition block a space of size $\Theta\left(\frac{\log^3 n \log^2 \log \log n}{\log^3 \log n}\right)$ and again place every element in $A \setminus B$ into a single location in its associated bucket. By Lemma 3.2 the number of elements in every bucket will be within a $\Theta\left(\frac{\log^2 \log n}{\log^2 n \log^2 \log \log n}\right)$ factor of its maximum value. Running a prefix sum through each bucket gives the total number of elements associated with it to within a $\Theta\left(\frac{\log^2 \log n}{\log^2 n \log^2 \log \log n}\right)$ factor; this cumulatively takes $\Theta\left(\frac{n \log n \log \log \log n}{\log \log n}\right)$ work. A prefix sum over the $\Theta(n)$ individual bucket sizes then gives an approximation of the bounds of every bucket in the final sorted array.

[Step 3.] Using the bounds found in step 2, we give each bucket $\Theta\left(\frac{\log^3 \log n}{\log^3 n \log^2 \log \log n}\right)$ times more space than the number of elements that will be going into it. Now we can sample every unsorted item into its associated partition block $\log n$ times, as we have $m = \Theta\left(\frac{\log^3 \log n}{\log^2 n \log^2 \log \log n}\right)$ times extra space for Lemma 3.2 which is sufficient for a high probability bound. No more than 1 in $\Theta\left(\frac{\log^3 \log n}{\log^2 n \log^2 \log \log n}\right)$ write attempts will produce a collision, so we have with probability greater than $1 - \frac{n}{(\log n)^{\log n}}$ that every unsorted element will be placed into its bucket at least once.

[Step 4.] All duplicate items in each partition are found by using $\Theta(n)$ prefix-sum-like methods to search over the $\log n$ locations where each element $a \in A$ was placed. We run $\frac{n \log \log \log n}{\log \log n}$ of these methods in parallel at a time (so there are $\Theta\left(\frac{\log \log n}{\log \log \log n}\right)$ chunks of them run serially) in order to stay within our space bound, taking $\mathcal{O}(n \log n)$ work and $\Theta\left(\frac{\log^2 \log n}{\log \log \log n}\right)$ span. After removing duplicate elements we index all nonempty memory locations by running a prefix sum across the entirety of the $\Theta\left(\frac{n \log n \log \log \log n}{\log \log n}\right)$ space being used, and then use those indices to compact down all elements into our final sorted array. \square

```

FULL-SORT( $A, n$ )
{ Part 1: Sorting the majority of  $A$  } .....
1.  $m \leftarrow \log n \log \log \log n / \log \log n$ ; Allocate arrays  $B, C, D$  of size  $nm$ ;  $B[0, \dots, n-1] \leftarrow A$ 
2. ALMOST-SORT( $n, n, 0, 0, m, B, C, D$ )
3. num_sorted  $\leftarrow$  smallest  $i$  s.t.  $B[i] = \text{null}$ 
{ Part 2: Combining unsorted elements into  $B$  } .....
4. Set all elements of  $C$  and  $D$  to null
5. Allocate arrays  $E, F$  of size  $n$ 
{ Pass 1: Finding and indexing all nonempty buckets } .....
6. parallel for  $a \in A[0, \dots, n-1]$  do
7.   Find smallest  $i$  s.t.  $B[i] \leq a < B[i+1]$ 
8.   if  $a \neq B[i]$  then
9.     Choose a random number  $j \in [0, \dots, m-1]$ 
10.    Attempt to assign  $C[i \cdot m + j] \leftarrow a$ ; in case of collision do nothing
11.  $D \leftarrow \text{INDICATOR-PREFIX-SUM}(C)$ 
12. parallel for  $i \leftarrow 0$  to  $n$  do
13.   if  $D[i \cdot \text{block\_size}] \neq D[(i+1) \cdot \text{block\_size} - 1]$  then  $E[i] = 1$ 
14.  $F \leftarrow \text{PREFIX-SUM}(E)$ 
15. Set all elements of  $C$  and  $D$  to null
{ Pass 2: Approximating the size of each bucket } .....
16.  $\text{block\_size} \leftarrow \log^3 n \log^2 \log \log n / \log^3 \log n$ 
17. parallel for  $a \in A[0, \dots, n-1]$  do
18.   Find smallest  $i$  s.t.  $B[i] \leq a < B[i+1]$ 
19.   if  $a \neq B[i]$  and  $E[i] = 1$  do
20.     Choose a random number  $j \in [0, \dots, \text{block\_size} - 1]$ 
21.     Attempt to assign  $C[(F[i] - 1) \cdot \text{block\_size} + j] \leftarrow a$ ; in case of collision do nothing
22.  $D \leftarrow \text{INDICATOR-PREFIX-SUM}(C)$ 
23. parallel for  $i \leftarrow 0$  to num_sorted - 1 do
24.   if  $E[i] = 0$  then  $E[i] = 1$ 
25.   else  $E[i] = D[F[i] \cdot \text{block\_size} - 1] - D[(F[i] - 1) \cdot \text{block\_size}]$ 
26.  $F \leftarrow \text{PREFIX-SUM}(E)$ 
27. clear  $C$  and  $D$ 
{ Pass 3: Sampling elements into their buckets with  $\log n$  repetitions } .....
28. parallel for  $i \leftarrow 0$  to  $n-1$  do
29.   parallel for  $j \leftarrow 0$  to  $\log n - 1$  do
30.     Choose a random number  $k = H(j, (F[i] - E[i])m, F[i]m - 1)$ 
31.     Attempt to assign  $C[k] \leftarrow a$ ; in case of collision do nothing
{ Removing duplicates and compacting } .....
32.  $\text{chunk\_size} \leftarrow n \log \log \log n / \log \log n$ 
33. for  $i \leftarrow 0$  to  $n/\text{chunk\_size} - 1$ 
34.   parallel for  $a \in A[i \cdot \text{chunk\_size}, \dots, (i+1) \cdot \text{chunk\_size} - 1]$  do
35.     Find smallest  $i$  s.t.  $B[i] \leq a < B[i+1]$ 
36.     KEEP-SINGLE( $C, a, n, H, (F[i] - E[i])m, F[i]m - 1$ )
37. parallel for  $i \leftarrow 0$  to num_sorted - 1 do
38.    $D[(F[i] - E[i])m, \dots, F[i]m - 1] \leftarrow \text{INDICATOR-PREFIX-SUM}(C[(F[i] - E[i])m, \dots, F[i]m - 1])$ 
39.   parallel for  $j \leftarrow (F[i] - E[i])m$  to  $F[i]m - 1$  do
40.      $D[j] \leftarrow D[j] + (F[i] - E[i])m$ 
41. parallel for  $i \leftarrow 0$  to  $nm - 1$  do
42.   if  $C[i] \neq \text{null}$  then {  $A[D[i]] \leftarrow C[i]; C[i] \leftarrow \text{null}$  }
43. parallel for  $i \leftarrow 0$  to num_sorted - 1 do
44.    $\text{lo} \leftarrow (F[i] - E[i])m$ ;  $\text{hi} \leftarrow \text{lo} + E[i] - 1$ 
45.    $A[\text{lo}, \dots, \text{hi}] \leftarrow \text{COLE-RAMACHANDRAN}(A[\text{lo}, \dots, \text{hi}])$ 
46.  $D \leftarrow \text{INDICATOR-PREFIX-SUM}(A)$ 
47. parallel for  $i \leftarrow 0$  to  $nm - 1$  do
48.   if  $A[i] \neq \text{null}$  then {  $C[D[i]] \leftarrow A[i]; A[i] \leftarrow \text{null}$  }
49.  $A \leftarrow C$ 

```

Figure 6: Listing for FULL-SORT. $H(i, \text{lo}, \text{hi})$ on line [30](#) is a hash function which maps the triple $(i, \text{lo}, \text{hi})$ to a single value in the range $[\text{lo}, \dots, \text{hi}]$. $\text{KEEP-SINGLE}(A, a, n, H, \text{lo}, \text{hi})$ on line [36](#) is a function which performs a PREFIX-SUM-like operation to index the instances of a at the (assumed distinct) locations $H(i, \text{lo}, \text{hi})$ ($i = 0, \dots, \log n - 1$), and then sets all but one of the instances to null.

Theorem 3.6. *The FULL-SORT algorithm (Figure 6) sorts an array of size n in $\mathcal{O}(n \log n)$ work w.h.p., $\Theta(\log n)$ span w.h.p., and $\Theta\left(\frac{n \log n \log \log \log n}{\log \log n}\right)$ space.*

Proof. Directly follows from theorem 3.3 and lemma 3.5. \square

Space-Adaptive Sorting Algorithm. In the previous paragraphs, we present a $\Theta(\log n)$ span parallel sorting algorithm in the binary-forking model without atomic TS using $n \cdot x$ extra space, where $x = \Theta(\log n \log \log \log n / \log \log n)$. We now present a space-adaptive sorting algorithm which, given $s \in [n, n \cdot x]$ space, achieves near-optimal span.

1. We divide the input into m equal-sized segments A_i such that $(n/m)x = s$.
2. For each segment A_i of size n/m , use FULL-SORT with s amount of extra space.
3. Merge the m segments pairwise recursively.

Theorem 3.7. *The SORT-ADAPTIVE algorithm sorts an array of size n in $\mathcal{O}(n \log n)$ work w.h.p. in n , $\mathcal{O}((n/s) \log^2 n)$ span w.h.p. in n , using $\Theta(s)$ amount of space.*

Proof. Sorting each segment takes $\Theta(\log(n/m))$ span and $\Theta((n/m) \log n)$ work w.h.p. in n . As sorting each segment consumes the entire extra space s , these m calls to FULL-SORT are made serially. Hence, these m sorting steps take $\Theta(m \cdot \log(n/m)) = \Theta((nx/s) \log n)$ span w.h.p. in n . Merging m sorted segments takes the following span.

$$\sum_{i=1}^{\log m} \Theta\left(\log \frac{m}{2^i} + \log 2^i n m\right) = \sum_{i=1}^{\log m} \Theta\left(\log m + \log \frac{n}{m}\right) = \Theta(\log n \log m).$$

The span of SORT-ADAPTIVE is $\mathcal{O}(n \log^2 n/s)$ w.h.p. in n . The work needed to merge m sorted segments is $\Theta(n \log \log n)$ w.h.p. in n . Hence, the overall work of SORT-ADAPTIVE is $\Theta(n \log n)$ w.h.p. in n . \square

4 Fast Fourier Transform

The Discrete Fourier Transform (DFT) of an array a of n complex numbers is the array y computed as $y[i] = \sum_{j=0}^{n-1} a[j] w_n^{-ij}$ for all $i \in [0, n-1]$, where $w_n = e^{2\pi\sqrt{-1}/n}$ is a primitive n th root of unity. A Fast Fourier Transform (FFT) is an algorithm that computes the DFT of an array rapidly. FFT is often considered as one of the most important algorithms of the 20th century. It is extensively used in digital signal processing. Several FFT algorithms have been designed that perform $\mathcal{O}(n \log n)$ work. For example, prime-factor algorithm (or Good-Thomas' FFT) [27, 38], Bruun's FFT [9], and Winograd's FFT [43]. Designing an FFT algorithm with $o(n \log n)$ work is an open problem.

Consider the recursive divide-and-conquer Cooley-Tukey FFT algorithm [18] and its special case, the radix-2 FFT algorithm [19]. A straightforward parallelization of the generic Cooley-Tukey algorithm [18] has a complexity of $\mathcal{O}(n \log n)$ work (and space) and $\mathcal{O}(\log n \log \log n)$ span. A simple parallelization of the radix-2 algorithm [19] has $\mathcal{O}(\log^2 n)$ span. In this section, we aim to design a parallel FFT algorithm with close to optimal span for the binary-forking model without atomics keeping work as closely as possible to $\mathcal{O}(n \log n)$. To this end, we first design a simple single-point FFT algorithm that can be used to compute a single entry of the DFT. We then carefully combine an efficient variant of this algorithm with the radix-2 FFT and mixed-radix FFT.




Figure 7: The n^ϕ -way FFT algorithm. Here, $r = n^\phi$.

n^ϕ -way FFT (Cooley-Tukey Algorithm, [18]). The n^ϕ -way FFT algorithm, for $\phi \in [1/\log n, 1/2]$, is defined as follows. We view the n -sized array as a $n^{1-\phi} \times n^\phi$ matrix, as shown in Figure 7. In the first phase, we compute FFT of each of the columns recursively in parallel. We then multiply all entries of the matrix by appropriate twiddle factors. In the second phase, we compute FFT of each of the rows recursively in parallel. Finally, the resultant matrix read in the column-major order is the required DFT.

The work and the span recurrences for the n^ϕ -way FFT algorithm are as follows. Suppose $c \geq 1$ is a fixed constant. If $n \leq c$, then $T_1(n) = \mathcal{O}(1)$ and $T_\infty(n) = \mathcal{O}(1)$. If $n > c$, then

$$\begin{aligned} T_1(n) &= n^\phi T_1(n^{1-\phi}) + n^{1-\phi} T_1(n^\phi) + \mathcal{O}(n), \\ T_\infty(n) &= T_\infty(n^{1-\phi}) + T_\infty(n^\phi) + \mathcal{O}(\log n). \end{aligned}$$

The work performed by this algorithm is $\mathcal{O}(n \log n)$ for all values of ϕ . If we set $\phi = 1/\log n$, we obtain the 2-way FFT algorithm [18, 19], with span $\mathcal{O}(\log^2 n)$. On the other hand, if we set $\phi = 1/2$, we get the \sqrt{n} -way FFT algorithm, a special case of the algorithm given in [18], with span $\mathcal{O}(\log n \log \log n)$. To the best of our knowledge, the \sqrt{n} -way FFT algorithm achieves the best span.

Single-point 2-way FFT. A single-point FFT evaluation means that we can compute a single entry of the DFT independently without computing the entire DFT. In other words, the computation of $y[i]$ does not share work with the computation of $y[j]$ for all $i, j \in [0, n-1]$ and $i \neq j$. The single-point 2-way algorithm FFT-SP is shown in Figure 8 (left). Figure 8 (right) shows a visual depiction of how an entry corresponding to the i th entry of the DFT is computed at recursion level ℓ using entries corresponding to the i th entries of the DFT's of the two child nodes at recursion level $\ell + 1$, for $\ell \in [0, \log n]$.

Consider the binary recursion tree produced by the algorithm. It is important to note that an entry in the DFT of a node can be computed with two entries in the DFT's of the child nodes. We do not need to store the DFT entries at every level. They can be computed recursively and on-the-fly. A DFT entry at the root node of the recursion tree can be computed with $\mathcal{O}(2^{\log n}) = \mathcal{O}(n)$ work and $\mathcal{O}(\log n)$ span. This implies that we can compute all entries of the DFT at the root node performing $\mathcal{O}(n \cdot 2^{\log n}) = \mathcal{O}(n^2)$ work in $\mathcal{O}(\log n)$ span.

Lemma 4.1. *The single-point 2-way FFT algorithm has a complexity of $\mathcal{O}(n^2)$ work, $\mathcal{O}(n)$ space, and $\mathcal{O}(\log n)$ span.*

2-way FFT with Stages. The standard 2-way FFT algorithm has a complexity of $\mathcal{O}(n \log n)$ work due to high work-sharing across threads and $\mathcal{O}(\log^2 n)$ span due to expensive local synchronization points. In contrast, the single-point 2-way FFT algorithm has a complexity of $\mathcal{O}(n^2)$ work due to no work-sharing across threads and $\mathcal{O}(\log n)$ span due to inexpensive local synchronization


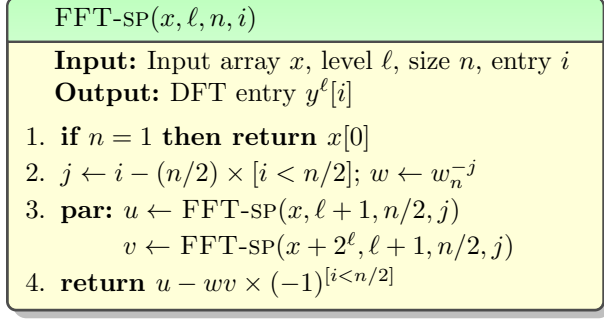


Figure 8: Left: Single-point FFT algorithm to compute the i th entry of the DFT of n -sized array x . Symbol $[]$ represents the Iversion bracket. Initial invocation is $\text{FFT-SP}(a, 1, n, i)$. Right: Single-point FFT evaluation of DFT entry $y^\ell[i]$ at level ℓ using the DFT entries $y^{\ell+1}[i]$ of the child nodes at level $\ell + 1$. Note that $y^\ell[n/2 + i]$ also depends on the $y^{\ell+1}[i]$ entries.

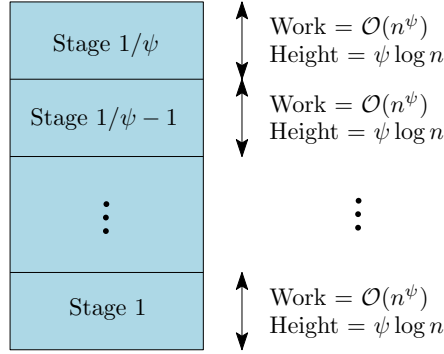


Figure 9: Stages in the 2-way FFT algorithm with single-point computations.

points. We now carefully combine the two algorithms to get advantages of both the worlds: low work and low span.

Consider Figure 9, in which the $\log n$ levels are split into $1/\psi$ stages, each stage containing $\psi \log n$ levels, for $\psi \in [1/\log n, 1]$. The core idea of the algorithm is as follows. There will be $1/\psi$ global synchronization points, one per stage. We compute the DFT at a root node of stage i using the DFT's at the root nodes of stage $i + 1$ using the single-point 2-way FFT algorithm. We compute and store the DFT's in the root nodes of all stages.

Lemma 4.2. *The 2-way FFT algorithm with stages has a complexity of $\mathcal{O}((1/\psi)n^{1+\psi})$ work, $\mathcal{O}(n)$ space, and $\mathcal{O}((1/\psi)\log n)$ span, where $\psi \in [1/\log n, 1]$.*

Proof. [Work.] There are $1/\psi$ stages. The height of each stage is $\psi \log n$. Total number of cells at the root nodes of each stage is $\mathcal{O}(n)$. The computation of each cell requires $\mathcal{O}(2^{\psi \log n}) = \mathcal{O}(n^\psi)$ work. So, the total work is $\mathcal{O}((1/\psi)n^{1+\psi})$. [Span.] The execution of stages is sequential. As there are $\mathcal{O}(n)$ cells at the root nodes of each stage, the span for launching these cells is $\mathcal{O}(\log n)$. The span for computing a DFT entry (i.e., FFT-SP) is $\mathcal{O}((1/\psi)\log n)$. [Space.] Storing the DFT array at each of the $1/\psi$ stages requires $\mathcal{O}((1/\psi)n)$ space. However, we can reuse two arrays to perform all computations. Hence, we just need $\mathcal{O}(n)$ space. \square

\sqrt{n} -way FFT and 2-way FFT with Stages. The core idea of the algorithm is as follows. We execute the \sqrt{n} -way FFT algorithm for the first $\log(1/\varepsilon)$ levels of the recursion tree, where

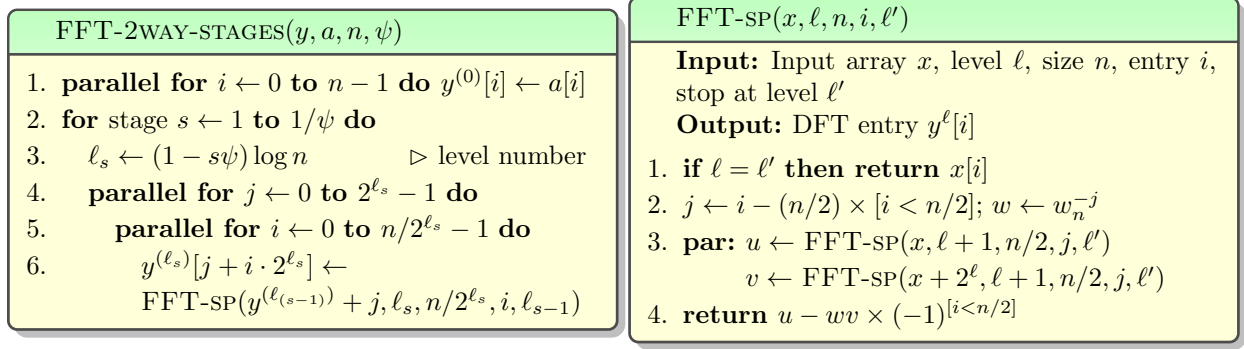


Figure 10: The 2-way FFT algorithm with stages.

$1/\varepsilon \in [2, \log n]$. We then switch to the 2-way FFT algorithm with stages.

Theorem 4.3. *The \sqrt{n} -way combined with 2-way with stages FFT algorithm has a complexity of $\mathcal{O}(n \log^{f(n, \varepsilon)} n)$ work, $\mathcal{O}(n)$ space, and $\mathcal{O}(\log(1/\varepsilon) \log n)$ span, where $1/\varepsilon \in [2, \log n]$ and $f(n, \varepsilon) = (1/\log \log n)(\log n / ((1/\varepsilon) \log(1/\varepsilon)) + \log((1/\varepsilon) \log(1/\varepsilon)))$.*

Proof. From Lemma 4.2, we have the following bounds for the 2-way FFT: $T'_1(m, \psi) = \mathcal{O}((1/\psi)m^{1+\psi})$ and $T'_\infty(m, \psi) = \mathcal{O}((1/\psi) \log m)$, where, $\psi \in [1/\log m, 1]$. Then, the work and span recurrences for the \sqrt{n} -way combined with 2-way with stages algorithm are as follows:

$$T_1(m) \leq \begin{cases} T'_1(m, \psi) & \text{if } m \leq n^\varepsilon, \\ 2m^{1/2}T_1(m^{1/2}) + \mathcal{O}(m) & \text{if } m > n^\varepsilon. \end{cases} \quad T_\infty(m) \leq \begin{cases} T'_\infty(m, \psi) & \text{if } m \leq n^\varepsilon, \\ 2T_\infty(m^{1/2}) + \mathcal{O}(\log m) & \text{if } m > n^\varepsilon. \end{cases}$$

Expanding the recurrences, we get

$$\begin{aligned} T_1(n) &\leq (1/\varepsilon)n^{1-\varepsilon}T'_1(n^\varepsilon, \psi) + c \cdot n/\varepsilon \\ T_\infty(n) &\leq (1/\varepsilon)T'_\infty(n^\varepsilon, \psi) + c' \cdot \log(1/\varepsilon) \log n \end{aligned}$$

for positive constants c and c' . Substituting the values of T'_1 and T'_∞ , we obtain $T_1(n) = \mathcal{O}((1/(\varepsilon\psi))n^{1+\varepsilon\psi})$ and $T_\infty(n) = \mathcal{O}((1/\psi + \log(1/\varepsilon)) \log n)$. We set $\psi = 1/\log(1/\varepsilon)$. Writing our work bound in the form $T_1(n) = \mathcal{O}(n \log^{f(n, \varepsilon)} n)$, we can easily find a corresponding function $f(n, \varepsilon)$. Direct computation shows that $f(n, \varepsilon) = (1/\log \log n)(\log n / ((1/\varepsilon) \log(1/\varepsilon)) + \log((1/\varepsilon) \log(1/\varepsilon)))$.

$$\begin{aligned} n^{1+\varepsilon\psi} / (\varepsilon\psi) &= n \log^{f(n, \varepsilon)} n \implies 1 + \varepsilon\psi \log n + \log(1/(\varepsilon\psi)) = 1 + f(n, \varepsilon) \log \log n \text{ (taking log)} \\ f(n, \varepsilon) &= (\varepsilon\psi \log n + \log(1/(\varepsilon\psi))) / \log \log n = (1/\log \log n)(\log n / ((1/\varepsilon) \log(1/\varepsilon)) + \log((1/\varepsilon) \log(1/\varepsilon))) \end{aligned}$$

When $\psi = 1/\log(1/\varepsilon)$, the span is $\mathcal{O}(\log(1/\varepsilon) \log n)$. Hence, the theorem follows. \square

Corollary 4.3.1. *The \sqrt{n} -way combined with 2-way with stages FFT algorithm has a complexity of $\mathcal{O}(n \log^{g(n)} n)$ work, $\mathcal{O}(n)$ space, and $\mathcal{O}(\log n \log \log \log n)$ span, where $g(n) < 2$ for $n < 10^{10,000}$.*

Proof. Set $1/\varepsilon = (\log \log n)^2$ and $g(n) = f(n, \varepsilon)$ in Theorem 4.3. \square

5 Conclusion

In this paper, we presented several fundamental low-span algorithms in the binary-forking model without using locks and atomic instructions. Our parallel algorithms perform work (almost) the same as that of the serial algorithms from which they are derived. All our results improve known results in the binary-forking model with and without atomics.

We introduced the technique of *single-point computation in stages* through Strassen's MM and FFT to carefully set a balance between high work-sharing and high span of the given algorithm and low work-sharing and low-span of the single-point computation variant to obtain parallel algorithms with optimal/near-optimal span without work blow-up. This technique can be used to design efficient parallel algorithms for other problems too.

We also presented a randomized sorting algorithm with optimal span and optimal work, both bounds are w.h.p. in the number of elements being sorted.

A few interesting problems (in the binary-forking model without using locks and atomic instructions) that one could aim to solve in the future are as follows: (1) Our parallel Strassen's MM algorithm achieves $\mathcal{O}(n^w \log \log n)$ work and optimal $\mathcal{O}(\log n)$ span (or $\mathcal{O}(n^w)$ work and $\mathcal{O}(\log n \log \log \log n)$ span). Design a Strassen's MM algorithm with $\mathcal{O}(n^w)$ work and optimal span. (2) Our randomized comparison sorting algorithm uses concurrent writes and achieves optimal $\mathcal{O}(n \log n)$ work and optimal $\mathcal{O}(\log n)$ span simultaneously, both bounds are with high probability in n , and uses $\omega(n)$ space. Design a randomized sorting algorithm that uses exclusive writes and achieves optimal work and optimal span bounds, both bounds w.h.p., and uses linear space. (3) Design a comparison sorting algorithm that uses exclusive writes and achieves optimal work and optimal span. (4) Our FFT algorithm achieves $\mathcal{O}(n \log^{g(n)} n)$ work and $\mathcal{O}(\log n \log \log \log n)$ span, where $g(n) < 2$ for $n < 10^{10,000}$. Design an FFT algorithm with $\mathcal{O}(n \log n)$ work and $\mathcal{O}(\log n)$ span.

Acknowledgments

This research was supported by NSF grants CNS-1553510, CCF-1439084, CNS-1938709, CCF-1617618, CCF-1716252, CCF-1725543, and CCF-1725428.

References

- [1] Umut A Acar, Guy E Blelloch, and Robert D Blumofe. The data locality of work stealing. In *ACM symposium on Parallel algorithms and architectures*, pages 1–12, 2000.
- [2] Selim G Akl. *Parallel sorting algorithms*, volume 12. Academic press, 2014.
- [3] Naama Ben-David, Guy E Blelloch, Jeremy T Fineman, Phillip B Gibbons, Yan Gu, Charles McGuffey, and Julian Shun. Parallel algorithms for asymmetric read-write costs. In *ACM Symposium on Parallelism in Algorithms and Architectures*, pages 145–156, 2016.
- [4] Guy E Blelloch, Rezaul Alam Chowdhury, Phillip B Gibbons, Vijaya Ramachandran, Shimin Chen, and Michael Kozuch. Provably good multicore cache performance for divide-and-conquer algorithms. In *SODA*, volume 8, pages 501–510. Citeseer, 2008.
- [5] Guy E Blelloch, Jeremy T Fineman, Phillip B Gibbons, and Harsha Vardhan Simhadri. Scheduling irregular parallel computations on hierarchical caches. In *ACM symposium on Parallelism in algorithms and architectures*, pages 355–366, 2011.
- [6] Guy E Blelloch, Jeremy T Fineman, Yan Gu, and Yihan Sun. Optimal parallel algorithms in the binary-forking model. In *ACM Symposium on Parallelism in Algorithms and Architectures*, pages 89–102, 2020.
- [7] Guy E Blelloch, Phillip B Gibbons, and Harsha Vardhan Simhadri. Low depth cache-oblivious algorithms. In *ACM Symposium on Parallelism in Algorithms and Architectures*, pages 189–199, 2010.
- [8] Guy E. Blelloch, Charles E. Leiserson, Bruce M Maggs, C Greg Plaxton, Stephen J Smith, and Marco Zagha. An experimental analysis of parallel sorting algorithms. *Theory of Computing Systems*, 31(2):135–167, 1998.
- [9] Georg Bruun. z-transform dft filters and fft’s. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 26(1):56–63, 1978.
- [10] Rezaul Chowdhury, Pramod Ganapathi, Yuan Tang, and Jesmin Jahan Tithi. Provably efficient scheduling of cache-oblivious wavefront algorithms. In *ACM Symposium on Parallelism in Algorithms and Architectures*, pages 339–350, 2017.
- [11] Rezaul Chowdhury, Pramod Ganapathi, Jesmin Jahan Tithi, Charles Bachmeier, Bradley C Kuszmaul, Charles E Leiserson, Armando Solar-Lezama, and Yuan Tang. Autogen: Automatic discovery of cache-oblivious parallel recursive algorithms for solving dynamic programs. In *ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, page 10. ACM, 2016.
- [12] Rezaul Chowdhury, Pramod Ganapathi, Stephen Tschudi, Jesmin Jahan Tithi, Charles Bachmeier, Charles E Leiserson, Armando Solar-Lezama, Bradley C Kuszmaul, and Yuan Tang. Autogen: Automatic discovery of efficient recursive divide-8-conquer algorithms for solving dynamic programming problems. *ACM Transactions on Parallel Computing*, 4(1):4, 2017.
- [13] Rezaul Alam Chowdhury and Vijaya Ramachandran. Cache-oblivious dynamic programming. In *SODA*, pages 591–600, 2006.

- [14] Rezaul Alam Chowdhury and Vijaya Ramachandran. Cache-efficient dynamic programming algorithms for multicores. In *SPAA*, pages 207–216, 2008.
- [15] Rezaul Alam Chowdhury and Vijaya Ramachandran. The cache-oblivious gaussian elimination paradigm: theoretical framework, parallelization and experimental evaluation. *Theory of Computing Systems*, pages 47(4):878–919, 2010.
- [16] Richard Cole. Parallel merge sort. *SIAM Journal on Computing*, 17(4):770–785, 1988.
- [17] Richard Cole and Vijaya Ramachandran. Resource oblivious sorting on multicores. *ACM Transactions on Parallel Computing*, 3(4):1–31, 2017.
- [18] James W Cooley and John W Tukey. An algorithm for the machine calculation of complex Fourier series. *Mathematics of Computation*, 19(90):297–301, 1965.
- [19] Thomas H Cormen, Charles E Leiserson, Ronald L Rivest, and Clifford Stein. *Introduction to algorithms*. MIT press, 2009.
- [20] Rathish Das, Shih-Yu Tsai, Sharmila Duppala, Jayson Lynch, Esther M Arkin, Rezaul Chowdhury, Joseph SB Mitchell, and Steven Skiena. Data races and the discrete resource-time tradeoff problem with resource reuse over paths. In *ACM Symposium on Parallelism in Algorithms and Architectures*, pages 359–368, 2019.
- [21] Vladimir Estivill-Castro and Derick Wood. A survey of adaptive sorting algorithms. *ACM Computing Surveys*, 24(4):441–476, 1992.
- [22] <http://docs.oracle.com/javase/tutorial/essential/concurrency/forkjoin.html>.
- [23] W Donald Frazer and Archie C McKellar. Samplesort: A sampling approach to minimal storage tree sorting. *Journal of the ACM*, 17(3):496–507, 1970.
- [24] Matteo Frigo, Charles E Leiserson, Harald Prokop, and Sridhar Ramachandran. Cache-oblivious algorithms. *ACM Transactions on Algorithms (TALG)*, 8(1):4, 2012.
- [25] Matteo Frigo, Charles E Leiserson, and Keith H Randall. The implementation of the cilk-5 multithreaded language. In *ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 212–223, 1998.
- [26] Pramod Ganapathi. *Automatic Discovery of Efficient Divide-&-Conquer Algorithms for Dynamic Programming Problems*. PhD thesis, Department of Computer Science, Stony Brook University, 2016.
- [27] Irving John Good. The interaction algorithm and practical fourier analysis. *Journal of the Royal Statistical Society: Series B (Methodological)*, 20(2):361–372, 1958.
- [28] Daniel S. Hirschberg. Fast parallel sorting algorithms. *Communications of the ACM*, 21(8):657–661, 1978.
- [29] J. JaJa. *An Introduction to Parallel Algorithms*. Addison Wesley, 1997. URL: <https://books.google.com/books?id=9BpYtwAACAAJ>.
- [30] Mohammad Mahdi Javanmard, Pramod Ganapathi, Rathish Das, Zafar Ahmad, Stephen Tschudi, and Rezaul Chowdhury. Toward efficient architecture-independent algorithms for dynamic programs. In *International Conference on High Performance Computing*, pages 143–164, 2019.

- [31] Donald E Knuth. *The art of computer programming: Volume 3: Sorting and Searching*. Addison-Wesley, 1998.
- [32] William F McColl and Alexandre Tiskin. Memory-efficient matrix multiplication in the bsp model. *Algorithmica*, 24(3-4):287–297, 1999.
- [33] V. Y. Pan. Strassen’s algorithm is not optimal trilinear technique of aggregating, uniting and canceling for constructing fast algorithms for matrix operations. In *Symposium on Foundations of Computer Science*, pages 166–176, 1978.
- [34] Vijaya Ramachandran and Elaine Shi. Data oblivious algorithms for multicores. *arXiv preprint arXiv:2008.00332*, 2020.
- [35] Volker Strassen. Gaussian elimination is not optimal. *Numerische mathematik*, 13(4):354–356, 1969.
- [36] Yuan Tang, Ronghui You, Haibin Kan, Jesmin Jahan Tithi, Pramod Ganapathi, and Rezaul A Chowdhury. Improving parallelism of recursive stencil computations without sacrificing cache performance. In *Workshop on Optimizing Stencil Computations*, pages 1–7, 2014.
- [37] <https://msdn.microsoft.com/en-us/library/dd460717%28v=vs.110%29.aspx>.
- [38] Llewellyn H Thomas. Using a computer to solve problems in physics. *Applications of digital computers*, pages 44–45, 1963.
- [39] <https://www.threadingbuildingblocks.org>.
- [40] J. J. Tithi, D. Matani, G. Menghani, and R. A. Chowdhury. Avoiding locks and atomic instructions in shared-memory parallel bfs using optimistic parallelization. In *International Symposium on Parallel Distributed Processing, Workshops and PhD Forum*, pages 1628–1637, 2013.
- [41] Jesmin Jahan Tithi. *Engineering High-performance Parallel Algorithms with Applications to Bioinformatics*. PhD thesis, Department of Computer Science, Stony Brook University, 2015.
- [42] Jesmin Jahan Tithi, Pramod Ganapathi, Aakrati Talati, Sonal Aggarwal, and Rezaul Chowdhury. High-performance energy-efficient recursive dynamic programming with matrix-multiplication-like flexible kernels. In *IPDPS*, pages 303–312, 2015.
- [43] Shmuel Winograd. On computing the discrete fourier transform. *Mathematics of computation*, 32(141):175–199, 1978.

6 Appendix

6.1 n^ε -way Merge Sort

In this section, we present a simple merge-based parallel sorting algorithm parameterized on a fixed constant $\varepsilon \in (0, 1]$, which achieves the span of $\mathcal{O}((\varepsilon + 1/\varepsilon) \log n)$ but performs suboptimal work of $\mathcal{O}((1/\varepsilon)n^{1+\varepsilon})$. We call this algorithm n^ε -WAY-SORT.

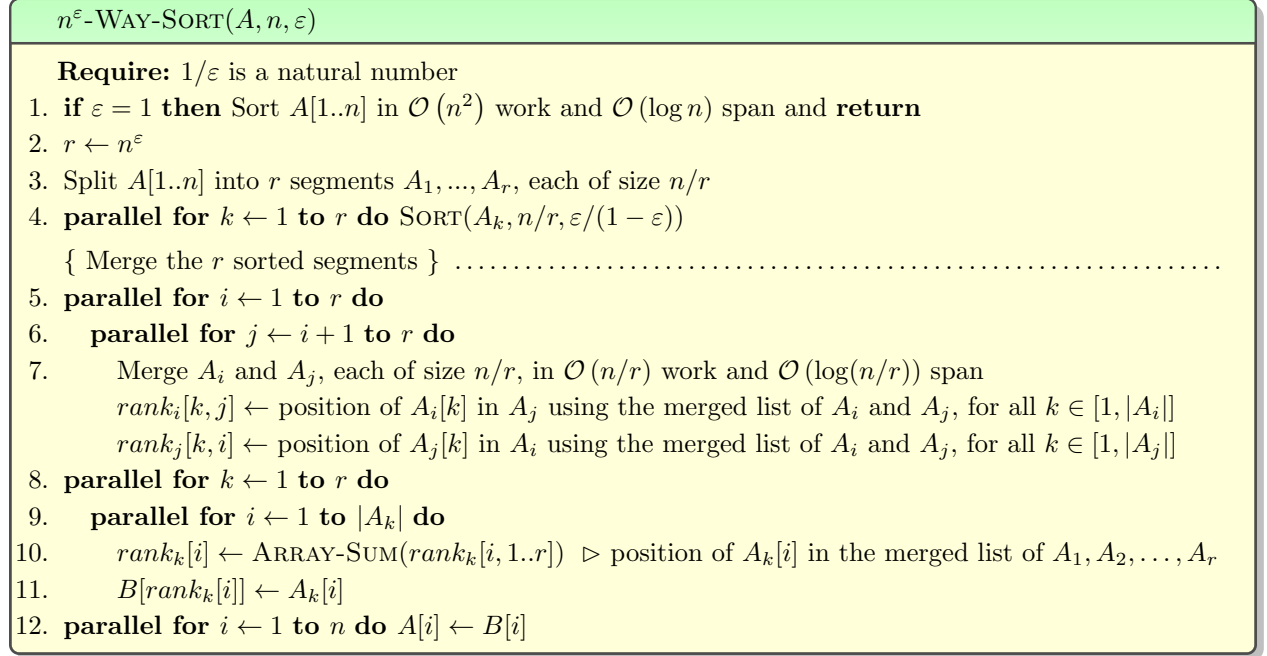


Figure 11: The n^ε -way sorting algorithm.

Figure 11 gives a pseudocode of the sorting algorithm. Choose a fixed constant $\varepsilon \in (0, 1]$ such that $1/\varepsilon$ is a natural number and suppose that $r = n^\varepsilon$. The input array $A[1..n]$ is split into r subarrays, each having n/r elements. All r subarrays are sorted recursively. These r sorted subarrays are then merged. The merging process consists of two stages.

In the first stage, we compute $rank_i[k, j]$ in parallel, for $i \in [1, r]$, $j \in [i + 1, r]$, and $k \in [1, |A_i|]$, where $rank_i[k, j]$ represents the rank or position of $A_i[k]$ (the k th element of subarray A_i) in A_j during the process of merging. Similarly, we compute $rank_j[k, i]$ in parallel. We do not need to store the merged list. We simply need to find the ranks of different elements of the array segments in $\mathcal{O}(n)$ time and $\mathcal{O}(\log n)$ span, which is easy to achieve.

In the second stage, we compute $rank_k[i]$ in parallel, for $k \in [1, r]$ and $i \in [1, |A_k|]$, where $rank_k[i]$ represents the position of $A_k[i]$ (the k th element of subarray A_i) in the merged list of the r subarrays A_1, A_2, \dots, A_r . We then use these values to sort A by corresponding assignments.

Lemma 6.1. n^ε -WAY-SORT has a complexity of $\mathcal{O}((1/\varepsilon)n^{1+\varepsilon})$ work and $\mathcal{O}((\varepsilon + 1/\varepsilon) \log n)$ span, for $\varepsilon \in (0, 1]$.

Proof. The work and span recurrences for the algorithm are:

$$T_1(n, \varepsilon) \leq \begin{cases} c_1 n^2 & \text{if } \varepsilon = 1, \\ n^\varepsilon T_1(n^{1-\varepsilon}, \frac{\varepsilon}{1-\varepsilon}) + c_2 n^{1+\varepsilon} & \text{if } \varepsilon \neq 1. \end{cases} \quad T_\infty(n, \varepsilon) \leq \begin{cases} c'_1 \log n & \text{if } \varepsilon = 1, \\ T_\infty(n^{1-\varepsilon}, \frac{\varepsilon}{1-\varepsilon}) + c'_2 \log n & \text{if } \varepsilon \neq 1. \end{cases}$$

Expanding the recurrences and setting $k = 1/\varepsilon - 1$, we get

$$\begin{aligned}
T_1(n, \varepsilon) &\leq n^\varepsilon \left[n^\varepsilon T_1(n^{1-2\varepsilon}, \frac{\varepsilon}{1-2\varepsilon}) + c_2 n \right] + c_2 n^{1+\varepsilon} = n^{2\varepsilon} T_1(n^{1-2\varepsilon}, \frac{\varepsilon}{1-2\varepsilon}) + 2c_2 n^{1+\varepsilon} \\
&\leq n^{2\varepsilon} \left[n^\varepsilon T_1(n^{1-3\varepsilon}, \frac{\varepsilon}{1-3\varepsilon}) + c_2 n^{1-\varepsilon} \right] + 2c_2 n^{1+\varepsilon} = n^{3\varepsilon} T_1(n^{1-3\varepsilon}, \frac{\varepsilon}{1-3\varepsilon}) + 3c_2 n^{1+\varepsilon} \\
&\leq n^{k\varepsilon} T_1(n^{1-k\varepsilon}, \frac{\varepsilon}{1-k\varepsilon}) + c_2 k n^{1+\varepsilon} = n^{1-\varepsilon} T_1(n^\varepsilon, 1) + c_2 k n^{1+\varepsilon} \\
&\leq n^{1-\varepsilon} (n^\varepsilon)^2 + c_2 k n^{1+\varepsilon} = n^{1+\varepsilon} + c_2 k n^{1+\varepsilon} = \mathcal{O}\left(\frac{1}{\varepsilon} n^{1+\varepsilon}\right) \\
T_\infty(n, \varepsilon) &\leq T_\infty(n^{1-k\varepsilon}, \frac{\varepsilon}{1-k\varepsilon}) + c'_2 \log n [1 + (1-\varepsilon) + (1-2\varepsilon) + \dots + (1-(k-1)\varepsilon)] \\
&= T_\infty(n^{1-k\varepsilon}, \frac{\varepsilon}{1-k\varepsilon}) + c'_2 \log n \times \frac{k}{2} [2 - (k-1)\varepsilon] \\
&= T_\infty(n^\varepsilon, 1) + c'_2 \left(\frac{1}{\varepsilon} + 1 - 2\varepsilon\right) \log n \leq c'_1 \log n^\varepsilon + c'_2 \left(\frac{1}{\varepsilon} + 1\right) \log n = \mathcal{O}\left((\varepsilon + \frac{1}{\varepsilon}) \log n\right)
\end{aligned}$$

□

6.2 d -D FFT

In this section, we develop a straightforward generalization of the \sqrt{n} -way FFT used in theorem 4.3. Instead of splitting up an FFT $y[i] = \sum_{j=0}^{n-1} a[j] w_n^{-ij}$ into two nested sums $y[i] = \sum_{j_1=0}^{\sqrt{n}-1} \sum_{j_2=0}^{\sqrt{n}-1} a[j_1 + \sqrt{n}j_2] w_n^{-i(j_1 + \sqrt{n}j_2)}$, we will break it into d nested sums

$$y[i] = \sum_{j_1=0}^{\sqrt[d]{n}-1} \dots \sum_{j_d=0}^{\sqrt[d]{n}-1} a \left[\sum_{k=1}^d j_k n^{(k-1)/d} \right] w_n^{-i \sum_{k=1}^d j_k n^{(k-1)/d}}.$$

Lemma 6.2. *The d -D-FFT algorithm takes $\mathcal{O}(dn \log n)$ work and $\mathcal{O}(d \log n \log_d \log n)$ span to run on an array of size n .*

Proof. The code given in Figure 12 shows the structure of the algorithm being used. We reindex the array as a column-major d -D hypercube with side length $\sqrt[d]{n}$; for every dimension we recursively apply d -D-FFT across all $n^{1-1/d}$ subarrays found by holding all but one of the d indices constant, followed by a single pointwise multiplication by twiddle factors. The recurrences from this for work and span are

$$\begin{aligned}
T_1(n) &= \begin{cases} \mathcal{O}(1) & \text{if } n \leq 1 \\ d n^{1-1/d} T_1(n^{1/d}) + \mathcal{O}(dn) & \text{if } n > 1 \end{cases}, \\
T_\infty(n) &= \begin{cases} \mathcal{O}(1) & \text{if } n \leq 1 \\ d T_\infty(n^{1/d}) + \mathcal{O}(d \log n) & \text{if } n > 1 \end{cases},
\end{aligned}$$

the solutions to these are given by $T_\infty(n) = \mathcal{O}(d \log n \log_d \log n)$ and $T_1(n) = \mathcal{O}(dn \log n)$. □

d -D-FFT(A, n)

1. **if** $n = 1$ **return** A
2. **for** $k \leftarrow 1$ to $\sqrt[d]{n}$ **do**
3. **parallel for** $j_1 \leftarrow 1$ to $\sqrt[d]{n} - 1$ **do**
- \vdots
4. **parallel for** $\hat{j}_k \leftarrow 1$ to $\sqrt[d]{n} - 1$ **do** {this line omitted}
- \vdots
5. **parallel for** $j_d \leftarrow 1$ to $\sqrt[d]{n} - 1$ **do**
6. d -D-FFT($A[\sum_{k'=1}^d j_{k'} n^{(k-1)/d}], n^{(k-1)/d}$) {where the array is indexed over j_k }
7. Multiply $A[\dots]$ by twiddle factors
8. **return** A

Figure 12: The d -D FFT algorithm.**FFT(X, n)**

(Input is a vector of length $n = 2k$ for some integer $k \geq 0$. Output is the in-place FFT of X .)

1. **Base Case:** If n is a small constant then compute FFT using the direct formula and return.
2. **Divide-and-Conquer:**
 - (a) **Divide:** Let $n_1 = 2^{\lceil \frac{k}{2} \rceil}$ and $n_2 = 2^{\lfloor \frac{k}{2} \rfloor}$. Observe that $n_2 \in \{n_1, 2n_1\}$.
 - (b) **Transpose:** Treat X as a row-major $n_1 \times n_2$ matrix. Transpose X in-place.
 - (c) **Conquer:** **for** $i \leftarrow 0$ to $n_2 - 1$ **do** FFT($X[i \times n_1, i \times n_1 + n_1 - 1], n_1$)
 - (d) **Multiply:** Multiply each entry of X by the appropriate twiddle factor
 - (e) **Transpose:** Treat X as a row-major $n_2 \times n_1$ matrix. Transpose X in-place.
 - (f) **Conquer:** **for** $i \leftarrow 0$ to $n_1 - 1$ **do** FFT($X[i \times n_2, i \times n_2 + n_2 - 1], n_2$)
 - (g) **Transpose:** Treat X as a row-major $n_1 \times n_2$ matrix. Transpose X in-place.
 - (h) **return** X

Figure 13: The generic FFT algorithm.

STRASSEN(U, V, n)

1. **if** $n = 1$ **then return** $U[0] \times V[0]$
 { Divide }
2. $U_{r1} \leftarrow U_{11} + U_{12}; U_{r2} \leftarrow U_{21} + U_{22}; U_{c1} \leftarrow U_{21} - U_{11}; U_{c2} \leftarrow U_{12} - U_{22}; U_{d1} \leftarrow U_{11} + U_{22}$
3. $V_{r1} \leftarrow V_{11} + V_{12}; V_{r2} \leftarrow V_{21} + V_{22}; V_{c1} \leftarrow V_{21} - V_{11}; V_{c2} \leftarrow V_{12} - V_{22}; V_{d1} \leftarrow V_{11} + V_{22}$
 { Conquer }
4. $P_1 \leftarrow \text{STRASSEN}(U_{d1}, V_{d1}, n/2); P_2 \leftarrow \text{STRASSEN}(U_{r2}, V_{11}, n/2); P_3 \leftarrow \text{STRASSEN}(U_{11}, V_{c2}, n/2)$
5. $P_4 \leftarrow \text{STRASSEN}(U_{22}, V_{c1}, n/2); P_5 \leftarrow \text{STRASSEN}(U_{r1}, V_{22}, n/2); P_6 \leftarrow \text{STRASSEN}(U_{c1}, V_{r1}, n/2)$
6. $P_7 \leftarrow \text{STRASSEN}(U_{c2}, V_{r2}, n/2)$
 { Merge }
7. Allocate X
8. $X_{11} \leftarrow P_1 + P_4 - P_5 + P_7; X_{12} \leftarrow P_3 + P_5; X_{21} \leftarrow P_2 + P_4; X_{22} \leftarrow P_1 - P_2 + P_3 + P_6$
9. **return** X

Figure 14: Strassen's MM algorithm.