Constant-Depth and Subcubic-Size Threshold Circuits for Matrix Multiplication

Ojas Parekh

Sandia National Laboratories
Albuquerque, NM, USA
odparek@sandia.gov

Conrad D. James

Sandia National Laboratories
Albuquerque, NM, USA
cdjame@sandia.gov

Cynthia A. Phillips

Sandia National Laboratories
Albuquerque, NM, USA
caphill@sandia.gov

James B. Aimone

Sandia National Laboratories
Albuquerque, NM, USA
jbaimon@sandia.gov

ABSTRACT

Boolean circuits of McCulloch-Pitts threshold gates are a classic model of neural computation studied heavily in the late 20th century as a model of general computation. Recent advances in large-scale neural computing hardware has made their practical implementation a near-term possibility. We describe a theoretical approach for multiplying two N by N matrices that integrates threshold gate logic with conventional fast matrix multiplication algorithms, that perform $O(N^\omega)$ arithmetic operations for a positive constant $\omega < 3$. Our approach converts such a fast matrix multiplication algorithm into a constant-depth threshold circuit with approximately $O(N^\omega)$ gates. Prior to our work, it was not known whether the $\Theta(N^3)$ -gate barrier for matrix multiplication was surmountable by constant-depth threshold circuits.

Dense matrix multiplication is a core operation in convolutional neural network training. Performing this work on a neural architecture instead of off-loading it to a GPU may be an appealing option.

KEYWORDS

threshold circuits; matrix multiplication; triangle counting; numerical algorithms; neural-inspired algorithms; neuromorphic computing; neural networks

ACM Reference Format:

Ojas Parekh, Cynthia A. Phillips, Conrad D. James, and James B. Aimone. 2018. Constant-Depth and Subcubic-Size Threshold Circuits for Matrix Multiplication. In *SPAA '18: 30th ACM Symposium on Parallelism in Algorithms and Architectures, July 16–18, 2018, Vienna, Austria*. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3210377.3210410>

1 INTRODUCTION

Neuromorphic computing devices (NCDs) are composed of massively parallel networks of artificial neurons that compute boolean functions, akin to conventional logic gates but with relatively larger fan-in. The simplicity of each neuron affords NCDs immense energy

efficiency and scalability, circumventing some of the data movement and energy bottlenecks associated with traditional HPC [17]. Within the last few years, large-scale commercial NCD platforms, such as SpiNNaker, IBM's True North, and Intel's Loihi, have emerged and currently offer up to hundreds of millions of neurons, with systems with billions of neurons a near-term likelihood [5, 17, 23]. Such NCDs were envisioned, for example, to accelerate deep learning; however, a specific application for which NCDs offer a clear and rigorously established advantage over conventional parallel computing remains elusive. Developing theoretical models for NCDs and identifying potential advantages and tradeoffs over other parallel models of computation remain largely open problems. We consider a well-known theoretical model that captures some of the features offered by NCDs, and we demonstrate resource-efficient constant-time algorithms for a fundamental problem, matrix multiplication.

We study the computational power of boolean circuits where the fundamental gates have unbounded fan-in and compute a linear threshold function¹. Such circuits are rooted in the classical McCulloch-Pitts neuronal model [12], with linear threshold functions serving as plausible models of spiking neurons. A boolean *threshold gate* with m binary inputs $y_1, y_2, \dots, y_m$ computes a *linear threshold function*, outputting 1 if and only if $\sum_{i=1}^m w_i y_i \geq t$, where the integer weights w_i and integer threshold t are constants associated with the gate. Rational w_i and t may be represented, for example, by multiplying such w_i and t with a sufficiently large number. There are several natural measures of complexity associated with boolean circuits, including *size*: the total number of gates, *depth*: the length of the longest directed path from an input node to an output node, *edges*: the total number of connections between gates, and *fan-in*: the maximum number of inputs to any gate.

Threshold gates are the most basic model for an artificial neuron, and as such, all currently available neuromorphic computing architectures provide a hardware implementation of threshold circuits. We consider threshold circuits with constant depth and polynomial size with respect to the total number of inputs; this class of circuits is called TC^0 . Such circuits represent a plausible model of constant-time parallel computing. This is a notion of perfect parallelizability, faster than the polylogarithmic time allowed in the complexity class

ACM acknowledges that this contribution was authored or co-authored by an employee, contractor, or affiliate of the United States government. As such, the United States government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for government purposes only.

SPAA'18, July 16–18, 2018, Vienna, Austria

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-5799-9/18/07...\$15.00

<https://doi.org/10.1145/3210377.3210410>

¹As we explain below, for convolutional neural networks, we can restrict the fan-in to a constant, or whatever practical fan-in the architecture supports.

NC . TC^0 circuits can compute a variety of functions including integer arithmetic, sorting, and matrix multiplication [19, 21]. There is also a TC^0 threshold-gate circuit of sublinear size to compute the parity of n bits [20]. In contrast, any constant-depth circuit with NOT gates and unbounded-fan-in AND and OR gates requires a superpolynomial number of gates [7, 26].

Understanding the power and limitations of TC^0 circuits has been a major research challenge over the last couple of decades. The 1990's saw a flurry of results showing what TC^0 circuits could do [19, 21], while more recent results have focused on lower bounds showing what TC^0 circuits cannot do [9]. TC^0 has been studied as a theoretical model. Its practicality is an open question. Currently, large-scale electronic circuits with high fan-in may be difficult to implement. However, neural-inspired architectures may offer hope.

The adult human brain contains about 100 billion neurons, with maximum fan-in 10,000 or larger some regions, operating at under 50 watts [1]. Though impressive, this figure represents a single class of instance size, so one might wonder how a synthetic system based on the physical limitations governing a brain might scale asymptotically. We are not aware of any generative brain models for which this has been analyzed. However, a fan-in that grows with the total system size seems plausible for a 3-dimensional circuit such as the brain. Although large fan-in is a critical resource in designing TC^0 algorithms, in practical neuromorphic devices resource requirements may grow as a function of fan-in. For example, the available numerical precision or dynamic range may decrease as fan-in increases, resulting in an overall increase in energy expenditure or execution time. Thus constant depth, in the TC^0 sense, may not practically equate to constant time. However, ideal theoretical algorithms may still guide the development of resource-efficient practical algorithms as neuromorphic architectures become more prevalent.

While neuromorphic computing has long focused on the use of analog computation to emulate neuronal dynamics [8], recent years have seen rapid development of novel digital CMOS neural hardware platforms which can scale to very large numbers of neurons [10, 13]. While initially designed for large biologically inspired circuits, these architectures are attracting attention as an alternative to conventional CMOS architectures for accelerating machine learning algorithms such as deep artificial neural networks [6]. Many of these neural architectures, such as TrueNorth and the SpiNNaker platform, achieve considerable benefits in energy and speed by using large numbers of simple digital spiking neurons instead of a relatively smaller number of powerful multi-purpose processors. These systems are almost configurable threshold gate circuits, except that they are capable of extended temporal dynamics. Scientific computing is an application domain for which neural architectures are often quickly dismissed. There is a perception that human cognition is better for data-centric functions, such as image recognition, and for abstract decision making than for precise numerical calculations, particularly at large scale. While biologically-inspired neural algorithms are often probabilistic or approximate, the neuronal-level computations in large scale neural architectures are sufficiently precise for numerical computation.

We consider a fundamental scientific-computing-inspired problem: can one produce constant-depth threshold circuits that compute the product of two $N \times N$ matrices using $O(N^{3-\epsilon})$ gates for

constant $\epsilon > 0$? For matrices with relatively large entries (say $\Omega(N)$ bits), this goal seems out of reach as we would need to output $\Omega(N^3)$ bits in the worst case. However, prior to our work, it was not known if this was possible even for *binary* matrices, those with entries that are all either 0 or 1.

We show how to multiply two $N \times N$ integer matrices with $O(\log N)$ -bit entries using $O(N^{3-\epsilon})$ gates in constant depth. The naïve algorithm based on the definition of matrix multiplication requires $\Theta(N^3)$ arithmetic operations. Our results are based on classical breakthroughs for fast matrix multiplication [22]: multiplying two $N \times N$ matrices using $O(N^\omega)$ arithmetic operations, for a positive constant $\omega < 3$, that depends on the particular fast matrix multiplication being employed. These techniques can be extended, in a relatively straightforward manner, to yield $O(\log N)$ -time conventional parallel algorithms (for architectures such as PRAMs) with $O(N^\omega)$ total work. In contrast, we give a constant-time algorithm, in the threshold circuit model, with approximately $O(N^\omega)$ total gates, which is a reasonable measure of total work.

One of our motivations for neural-circuit-based matrix multiplication is convolutional neural networks for deep learning. See Section 5 for more details. Deep learning is a major driver for neural-inspired architectures. A current vision for using these architectures for deep learning requires the matrix multiplication to be moved off-system to a GPU. If circuit-based matrix multiplication can be made practical, perhaps this computation can be left on-chip, avoiding energy-intensive and slow I/O.

We also consider the somewhat simpler problem of determining whether the trace of A^3 is at least τ , for an $N \times N$ integer matrix A with entries of size $O(\log N)$ bits. This case illustrates the fundamental ideas of our approach and has applications in social network analysis, particularly to triangle counting. The problem we solve allows one to answer: "Does a graph G have at least τ triangles?" The user may select a relevant value of τ . See Section 5 for more details on triangles, social network analysis, and picking τ .

There is a simple depth-2 threshold circuit to solve this problem for a graph $G = (V, E)$. The circuit has an input variable, x_{ij} for $i, j \in V$ with $i < j$; the variable x_{ij} is 1 if $ij \in E$ and 0 otherwise. The first layer of the circuit consists of a gate, g_{ijk} for each triple $i, j, k \in V$ with $i < j < k$. The gate g_{ijk} computes the value of the linear threshold function $x_{ij} + x_{ik} + x_{jk} \geq 3$ as an output y_{ijk} . That is, the gate fires ($y_{ijk} = 1$) if and only if all edges in the triangle on i, j , and k are in the graph. The second layer consists of a single output gate that computes the linear threshold function $\sum_{i,j,k \in V: i < j < k} y_{ijk} \geq \tau$; this gate fires if and only if the number of triangles in G is at least τ . The circuit has $\binom{N}{3} + 1 = \Theta(N^3)$ gates.

We ask (and answer) whether it is possible to beat the size of this threshold circuit in constant depth. This is akin to asking if it is possible to beat the naïve matrix multiplication algorithm with an algorithm that performs $O(N^\omega)$ operations for $\omega < 3$. In fact the above threshold circuit is a specialization of the naïve matrix multiplication algorithm.

The analysis of our new threshold circuits is more involved than analyzing conventional fast matrix multiplication methods. We must explicitly consider sparsity (see Definition 2.1), a measure of how many times a matrix element or intermediate result is part of a computation during the fast multiplication. Thus, while we

use existing fast matrix multiplication techniques to achieve our results, we use them in a new context. Our performance results exploit different features of fast matrix multiplication techniques than those traditionally used.

Results and contributions

Consider a fast recursive or divide-and-conquer matrix multiplication algorithm like Strassen's, with run-time complexity $O(N^\omega)$. We consistently use ω as the exponent in the runtime complexity of a base non-circuit fast matrix multiplication algorithm. Our results leverage such a fast matrix multiplication to construct a constant-depth threshold circuit with $\tilde{O}(N^{\omega+\varepsilon})$ -gates, where ε depends on the depth of the circuit.

Our main result is an $O(d)$ -depth, $\tilde{O}(N^{\omega+O(\gamma^d)})$ -gate threshold circuit for multiplying two $N \times N$ integer matrices with $O(\log N)$ -bit entries, for a positive integer d and constant $\gamma < 1$. Specifically, for a given integer d , the depth is $4d + 1$. The constant d is a multiplicative factor hidden in the $\tilde{O}$ for the number of gates. Section 4.3 gives a more detailed discussion of the value of γ . For Strassen's algorithm it is about 0.491. The constant multiplier of γ^d is about 1.581 for Strassen's algorithm. Thus for $d > 3$, this circuit will have $O(N^{3-\varepsilon})$ gates for positive constant $\varepsilon > 0$. We also give a $O(\log \log N)$ -depth, $\tilde{O}(N^\omega)$ -gate circuit for this task.

We present a simplified circuit of the same gate complexity and slightly lower depth ($2d + 2$) for computing the trace of A^3 , for an $N \times N$ integer matrix A . This gives triangle counts for a graph G with adjacency matrix A (see Section 2.3). Our circuits implement limited-depth versions of fast divide-and-conquer matrix multiplication, and our techniques should extend to other types of algebraic divide-and-conquer algorithms.

Our contributions are:

- This work revives and redirects research on designing algorithms for a classical theoretical model of parallel computation to a data science problem on an emerging class of neural-inspired parallel architectures. We show that threshold circuits, comprised of threshold gates that model neurons, might be applicable to linear-algebra computations for deep learning on new neuromorphic hardware.
- We give $O(\log \log n)$ -depth and $\tilde{O}(N^\omega)$ -gate threshold circuits for computing the product of two $N \times N$ matrices with $O(\log N)$ -bit entries, where N^ω is the complexity of a fast conventional matrix multiplication algorithm like Strassen's.
- We give constant-depth threshold circuits with $\tilde{O}(N^{\omega+\varepsilon})$ gates, where ε is exponentially small in the depth of the circuit.
- Our circuits are elementary and are composed entirely of copies of a relatively simple depth-2 threshold circuit that performs addition. We hope this will facilitate implementation of our approach in neural-inspired hardware.

2 PRELIMINARIES AND PROBLEM STATEMENT

2.1 Fast matrix multiplication algorithms

Strassen developed the first matrix multiplication algorithm requiring $O(N^{3-\varepsilon})$ multiplications [22]. Strassen observed that one can

$$\begin{aligned}
 M_1 &= A_{11}(B_{12} - B_{22}) \\
 M_2 &= (A_{21} + A_{22})B_{11} \\
 M_3 &= (A_{11} + A_{22})(B_{11} + B_{22}) \\
 M_4 &= A_{22}(B_{21} - B_{11}) \\
 M_5 &= (A_{11} + A_{12})B_{22} \\
 M_6 &= (A_{21} - A_{11})(B_{11} + B_{12}) \\
 M_7 &= (A_{12} - A_{22})(B_{21} + B_{22}). \\
 C_{11} &= M_3 + M_4 - M_5 + M_7 \\
 C_{12} &= M_1 + M_5 \\
 C_{21} &= M_2 + M_4 \\
 C_{22} &= M_1 - M_2 + M_3 + M_6.
 \end{aligned}$$

Figure 1: Strassen's algorithm for multiplying two 2×2 matrices A and B . The 7 multiplications computed in Strassen's algorithm are represented by $M_1, \dots, M_7$. Each M_i is the product of weighted sums of entries of matrices A and B . The entries of product matrix C are then computed from the M_i using only addition and subtraction. One can verify by substitution and expansion that the entries of C are set to the proper expressions involving entries of A and B .

compute the matrix product, $C = AB$ for 2×2 matrices A and B using 7 multiplications rather than the 8 multiplications required by the naïve algorithm. The reduction in multiplications comes at the expense of additional additions and subtractions.

Figure 1 gives Strassen's algorithm for 2×2 matrices. The algorithm is generalized to $N \times N$ matrices A and B , where $N = 2^l$ for some positive integer l , as follows. We partition A and B into 4 blocks, each of size $N/2 \times N/2$, and let A_{ij} and B_{ij} refer to these blocks, for $i, j \in \{1, 2\}$. The above equations remain correct. However, each M_i now represents a multiplication of two $N/2 \times N/2$ matrices. We can recursively apply the above procedure to perform each of these multiplications until the blocks are individual matrix elements or, for more practical applications, sufficiently small. For each of the $l = \log_2 N$ levels of recursion, we invoke 7 recursive matrix multiplications, resulting in a total of $7^{\log_2 N} = N^{\log_2 7} \approx N^{2.81}$ scalar multiplications. The recurrence relation for the total number of arithmetic operations is $T(N) = 7 \cdot T(N/2) + 18 \cdot (N/2)^2$ and $T(1) = O(1)$. The $18 \cdot (N/2)^2$ term arises from the 18 additions or subtractions on $N/2 \times N/2$ blocks in the expressions above. This recurrence shows the total number of scalar additions or subtractions is also $O(N^{\log_2 7})$.

Although Strassen's seminal approach was based on a fast matrix multiplication algorithm for 2×2 matrices, subsequent work has yielded improved algorithms by employing a fast matrix multiplication algorithm involving larger square matrices as well as more sophisticated techniques. The currently best known algorithm requires $O(N^{2.373})$ operations [11]. For the sake of exposition, we view fast matrix multiplication algorithms as recursive divide-and-conquer approaches, yet our techniques extend to the more general tensor perspective of fast matrix multiplication. See the survey by Bläser [4] for a detailed introduction to and history of this area, including the connection between the (border) rank of the matrix multiplication tensor and fast matrix multiplication algorithms.

2.2 Technical challenges

The divide-and-conquer Strassen's algorithm has a natural $O(\log N)$ -time parallel (PRAM) implementation with a total work of $O(N^{\log_2 7})$

arithmetic operations. The main question we consider is whether Strassen's approach and subsequent improvements of it can yield a *constant-time* algorithm implemented using threshold circuits with $O(N^{3-\epsilon})$ gates, where the latter is a measure of total work. The recursive ($O(\log N)$ -depth) implementation of Strassen's algorithm only performs scalar multiplications during the base case. However, it performs matrix additions and subtractions at each level before and after the recursion, reusing computed results.

If we attempt to implement Strassen's approach without recursion and the consequent reuse of computed results, we must first compute $N^{\log_2 7}$ scalars that are linear combinations of entries of A and another $N^{\log_2 7}$ scalars representing linear combinations of entries of B . The main technical hurdle is that such linear combinations involve up to N entries of A or B , and we seek to compute $O(N^{\log_2 7})$ of these sums with constant-depth threshold circuits. A naïve implementation would require $\Omega(N^{1+\log_2 7})$ gates.

We overcome this hurdle by selecting a constant or $O(\log \log N)$ number of levels of recursion out of the $O(\log N)$ levels suggested by the standard implementation of Strassen's approach. We make the notion of selecting a level of recursion precise in Section 4. This allows us enough reuse of computed results within the confines of a constant-depth or $O(\log \log N)$ -depth circuit to achieve our results. We note that we must carefully select the levels of recursion employed; for instance, simply selecting every k th level does not achieve our best results.

2.3 Problem statement

We develop threshold circuits to compute the matrix product $C = AB$ of two $N \times N$ integer matrices A and B . Our results assume the entries of A and B require at most $O(\log N)$ bits. We also consider a related problem: given an integer matrix A as above and an integer τ , determine whether the matrix trace of A^3 is at least τ . This problem is solved by a simpler threshold circuit than for computing matrix product and serves to illustrate our main ideas. It also has applications to triangle counting in graphs as we describe below.

Let A be the $N \times N$ symmetric adjacency matrix of a simple graph $G = (V, E)$ with N nodes: for $i, j \in V$, $A_{ij} = A_{ji} = 1$ if $ij \in E$, and $A_{ij} = A_{ji} = 0$ otherwise. Since there are no self-loops in the graph, we have $A_{ii} = 0$ for all nodes i . Consider the square of the adjacency matrix, $C = A^2$. For $i, j \in V$ with $i \neq j$, $C_{ij} = \sum_{k \in V} A_{ik} A_{kj} = |\{k \in V \mid ik \in E \text{ and } kj \in E \text{ and } k \neq i, j\}|$, which is the number of paths of length 2 between i and j . If there is an edge between the nodes i and j , then each path of length 2 between them, along with the edge ij , forms a triangle in G . Moreover, every triangle containing i and j arises in this way. Suppose G has Δ triangles. Then,

$$3\Delta = \sum_{i, j \in V: i < j} A_{ij} C_{ij}, \quad (1)$$

since the sum counts each triangle once for each of its edges. Thus one can count the triangles in G by summing some of the entries of A^2 . An equivalent computation is the trace of A^3 , $\text{trace}(A^3)$, which, from (1), is equal to 6Δ . This counts the loop from each vertex in each direction.

We employ a threshold circuit implementation of fast matrix multiplication algorithms to compute this sum in constant depth using $O(N^{3-\epsilon})$ gates. In fact the exponent of our gate count can be

made arbitrarily close to the exponent of the arithmetic operation count for the best possible fast matrix multiplication algorithm.

We explain our notion of a fast matrix multiplication algorithm. We assume we are given an algorithm for multiplying two $T \times T$ matrices using a total of r multiplications (for Strassen's algorithm, $T = 2$ and $r = 7$). We assume $N = T^l$ for some positive integer l . As outlined in Section 2.1, this yields a recursive algorithm for computing the product of two $N \times N$ matrices, $C = AB$, using a total of $r^l = r^{\log_T N} = N^{\log_T r}$ scalar multiplications.

As with Strassen's algorithm, we assume we are given a list of r expressions for each of the multiplications, $M_1, \dots, M_r$; we view each M_i as an expression involving the T^2 different $N/T \times N/T$ blocks of A and B . In particular each M_i is a product of a $\{-1, 1\}$ -weighted sum of blocks of A with a $\{-1, 1\}$ -weighted sum of blocks of B . We also assume the fast matrix multiplication algorithm provides a list of T^2 expressions, each representing a $N/T \times N/T$ block of C as a $\{-1, 1\}$ -weighted sum of the M_i . More general fast matrix multiplication algorithms may allow the M_i to be products of linear combinations with rational weights beyond $\{-1, 1\}$ (likewise for the entries of C). Although we do not present details here, our techniques can be extended for such fast matrix multiplication algorithms (these weights correspond to the w_i in Lemma 3.2 in Section 3).

For $1 \leq i \leq r$, let a_i be the number of distinct blocks of A that appear in the expression M_i , and let b_i be defined analogously with respect to B . We let c_i be the number of expressions for blocks of C in which M_i appears.

Definition 2.1. We let

$$s_A = \sum_{1 \leq i \leq r} a_i, \quad s_B = \sum_{1 \leq i \leq r} b_i, \quad \text{and} \quad s_C = \sum_{1 \leq i \leq r} c_i.$$

We define the *sparsity* of a fast matrix multiplication algorithm as $s = \max\{s_A, s_B, s_C\}$.

Sparsity will be an essential ingredient of our analysis, and we better motivate it in Section 4.3. Others [2, 3] consider sparsity in analyzing and improving the numerical stability of fast matrix multiplication algorithms, though they do not refer to it by the same name.

We use the following notation in proofs of circuit quality. We define $\text{bits}(m)$ as the minimum number of bits required to express the nonnegative integer m in binary, i.e. the least integer l such that $m < 2^l$.

3 BASIC TC^0 ARITHMETIC CIRCUITS

We first develop the fundamental TC^0 arithmetic circuits on which our results rely. Our circuits are designed with neuromorphic implementation in mind, and we try to favor simple constructions over those that offer the least depth or gate count. The bulk of the computation performed by our circuits is computing the bits of integer-weighted sums of nonnegative integers, $\sum_i w_i x_i$, where the nonnegative x_i depend upon the inputs to the circuit but the weights w_i are constants associated with the circuit.

Our first circuit follows from a classical technique to compute symmetric functions in TC^0 by Muroga from 1959 [14, 15]; it is also a special case of a more general result by Siu et al. [20]. We include a proof to demonstrate the simplicity of the construction.

LEMMA 3.1. *Let $s = \sum_i w_i x_i$ be an integer-weighted sum of bits, $x_i \in \{0, 1\}$. We assume $s \geq 0$ and fix an integer l such that $s \in [0, 2^l)$. For $1 \leq k \leq l$, the k th most significant bit of s can be computed by a depth-2 threshold circuit using $2^k + 1$ gates.*

PROOF. We define $\text{bool}(P)$, for a predicate P , to be 1 if predicate P is true and 0 otherwise.

The k th most significant bit of s is 1 precisely when s lies in one of the intervals $[i2^{l-k}, (i+1)2^{l-k})$ for some odd integer $1 \leq i < 2^k$. The interval enumerates over all combinations of bits less significant than the k th and the odd multipliers i enumerate over all combinations of bits more significant than k . The first layer of our circuit computes the function $y_i := \text{bool}(s \geq i2^{l-k})$, for $1 \leq i \leq 2^k$. The output of the circuit is $\text{bool}(\sum_{i \text{ odd}} (y_i - y_{i+1}) \geq 1)$, since $y_i - y_{i+1}$ is 1 if $s \in [i2^{l-k}, (i+1)2^{l-k})$ and is 0 otherwise. $\square$

The circuit construction for the above lemma requires an integer l such that the sum s is guaranteed to be in $[0, 2^l)$. Note that if $s \notin [0, 2^l)$, the circuit for any k outputs 0. We build upon the above to obtain our primary addition circuit. The next lemma is a generalized and strengthened version of Siu et al.'s Lemma 3.1 in [20] for the depth-2 case.

LEMMA 3.2. *Let $s = \sum_i w_i z_i$ be an integer-weighted sum of n nonnegative numbers z_i , each with at most b bits. We assume $s \geq 0$ and let $w = \max_i |w_i|$. The sum s can be computed by a depth-2 threshold circuit with $O(wbn)$ gates.*

PROOF. The sum s requires at most $\text{bits}(nw(2^b - 1)) \leq \text{bits}(n) + \text{bits}(w) + b$ bits. (Recall the definition of $\text{bits}()$ in Section 2.3).

First we compute the j th (least significant) bit of s , for $1 \leq j \leq b$. Let $s_j = \sum_i w_i z_{ij}$, where z_{ij} is obtained from z_i by ignoring all but the least significant j bits. Note that s_j requires at most $\text{bits}(n) + \text{bits}(w) + j$ bits and that s and s_j have the same least significant j bits. Furthermore, since we are given the bits of each z_i , we may treat s_j as an integer-weighted sum of bits, where each weight is a product of some w_i with a power of 2. Thus we may compute the j th bit of s_j by appealing to Lemma 3.1 on s_j with $k = \text{bits}(n) + \text{bits}(w) + 1$. To see this recall that k represents the k th most significant bit. The j th least significant bit of s_j is the $(\text{bits}(s_j) - j + 1) = n + w + j - j + 1$ most significant bit. This circuit requires $2^k + 1 = 2 \cdot 2^{\text{bits}(n)} \cdot 2^{\text{bits}(w)} + 1 = O(wn)$ gates, hence $O(bwn)$ gates suffice to compute the b least significant bits of s .

Appealing to Lemma 3.1 to compute each of the remaining $a = \text{bits}(n) + \text{bits}(w)$ most significant bits of s requires $O(\sum_{k=1}^a 2^k) = O(2^a) = O(wn)$ gates. This is improved in practice by observing that the functions y_i computed for $k = \text{bits}(n) + \text{bits}(w)$ in the proof of Lemma 3.1 include those required for all the most significant $\text{bits}(n) + \text{bits}(w)$ bits of s . $\square$

We need to compute products of numbers as well. However, the products we compute are only used as inputs to other threshold gates, and we do not need an explicit base-2 representation of the bits of these products. A more generic representation suffices: a representation of an integer x is an integer-weighted sum of binary variables, $x = \sum_{1 \leq i \leq d} w_i x_i$ with $x_i \in \{0, 1\}$ and d polynomial in $\text{bits}(x)$.

LEMMA 3.3. *A representation of the product of three m -bit nonnegative integers can be computed by a depth-1 threshold circuit with m^3 gates.*

PROOF. We compute a representation of the product of $x = \sum_{1 \leq i \leq m} 2^{i-1} x_i$, $y = \sum_{1 \leq j \leq m} 2^{j-1} y_j$, and $z = \sum_{1 \leq k \leq m} 2^{k-1} z_k$, with $x_i, y_j, z_k \in \{0, 1\}$. Thus $xyz = \sum_{1 \leq i, j, k \leq m} 2^{i+j+k-3} x_i y_j z_k$, which is a representation for xyz . This representation differs from the standard binary representation in that $2^{i+j+k-3}$ can represent the same power of 2 for different values of i, j , and k . We use m^3 gates in a single layer with predicates $x_i + y_j + z_k \geq 3$ to compute $x_i y_j z_k \in \{0, 1\}$ for $1 \leq i, j, k \leq m$. $\square$

Our results require only the above relatively simple arithmetic circuits. This facilitates practical implementation.

Negative numbers. The above lemmas give circuits to compute products and weighted sums of nonnegative integers. However, they can be extended to handle negative integers. We represent each integer x as $x = x^+ - x^-$, where x^+ and x^- are each non-negative. Other more efficient representations are possible, but we select this one as it makes for a relatively simple presentation and implementation at the cost of a constant-factor overhead in gate and wire count.

The workhorse subroutine of our circuits, captured by Lemma 3.2, is computing the bits of integer-weighted sums of nonnegative integers, $s = \sum_i w_i x_i$, where the $x_i = x_i^+ - x_i^-$ depend upon the inputs to the circuit but the w_i are constant with respect to the inputs. Let W^+ be the set of indices with $w_i > 0$, and let W^- be those indices with $w_i < 0$. We define $s^+ = \sum_{i \in W^+} w_i x_i^+ + \sum_{i \in W^-} (-w_i) x_i^-$ to be the positive terms in sum s , and we define $s^- = \sum_{i \in W^+} w_i x_i^- + \sum_{i \in W^-} (-w_i) x_i^+$ to be the negation of the negative terms in sum s . We have $s = s^+ - s^-$ and $s^+, s^- \geq 0$. Moreover, each of s^+ and s^- is an integer-weighted sum of nonnegative integers, hence the bits of each of s^+ and s^- may be computed using two separate instances of the circuit of Lemma 3.2. Each of these circuit instances only depends on the x_i^+ and x_i^- hence we may apply them in parallel without increasing the depth of the resulting overall circuit.

Computing products also incurs extra constant overhead. The representation of the product of three numbers, xyz described in the proof of Lemma 3.3 becomes $xyz = \sum_{1 \leq i, j, k \leq m} 2^{i+j+k-3} (x_i^+ - x_i^-)(y_j^+ - y_j^-)(z_k^+ - z_k^-)$. Thus $\sum_{1 \leq i, j, k \leq m} [2^{i+j+k-3} (x_i^+ y_j^+ z_k^+ + x_i^+ y_j^+ z_k^- + x_i^+ y_j^- z_k^+ + x_i^+ y_j^- z_k^- + x_i^- y_j^+ z_k^+ + x_i^- y_j^+ z_k^- + x_i^- y_j^- z_k^+ + x_i^- y_j^- z_k^-)]$ is also a representation of xyz that requires eight times as many gates to compute, which is still $O(m^3)$.

For ease of exposition, we proceed as if we were only computing positive quantities. From this point on, we assume that a number x requires at most b bits, by which we mean that each of x^+ and x^- requires at most b bits.

4 SUBCUBIC TC^0 CIRCUITS FOR TRACE AND MATRIX MULTIPLICATION

4.1 Overview

Our circuits for matrix trace and matrix multiplication implement a given conventional fast matrix multiplication algorithm in both a depth-efficient and gate-efficient manner. We assume we are given

Figure 2: The r -ary tree $\mathcal{T}_A$ for Strassen's Algorithm ($r = 7, T = 2$). For $K \times K$ matrices U and V , the notation U_{ij} or $(U)_{ij}$ refers to the (i, j) th $\frac{K}{T} \times \frac{K}{T}$ block of U . Observe that $(U + V)_{ij} = U_{ij} + V_{ij}$. Each node has children corresponding to the r multiplication expressions M_i (see Figure 1). An edge associated with M_i is labeled with the number of terms of A that appear in M_i . Each node u on level h , starting with the root as level 0, corresponds to a matrix that is a weighted sum of $\frac{N}{T^h} \times \frac{N}{T^h}$ blocks of A . The number of blocks of A appearing in such a sum is the product of the edge labels on the path from u to the root of the tree. For example, $(A_{12} - A_{22})_{12} - (A_{12} - A_{22})_{22} = (A_{12})_{12} - (A_{22})_{12} - (A_{12})_{22} + (A_{22})_{22}$ is a weighted sum of $4 \frac{N}{T^2} \times \frac{N}{T^2}$ blocks of A . The $N^{\log_T r}$ leaves of $\mathcal{T}_A$ correspond to scalars that are weighted sums of entries of A .

$N \times N$ integer matrices A and B with entries of size $O(\log N)$. We consider two problems: (1) determining whether $\text{trace}(A^3) \geq \tau$, for an integer τ , and (2) computing the bits of the matrix product $C = AB$. We consider the first problem in this section, and the second problem in the next section.

We define trees $\mathcal{T}_A$ and $\mathcal{T}_B$ for the input matrices A and B , respectively, based on the recursive or divide-and-conquer structure of the fast matrix multiplication algorithm. The nodes in $\mathcal{T}_A$ represent weighted sums of blocks of A and likewise for $\mathcal{T}_B$. The root of $\mathcal{T}_A$ represents the matrix A , while the leaves represent weighted sums of its entries. See Figure 2 for a detailed explanation.

In a conventional PRAM implementation of a fast matrix multiplication algorithm, all the matrices at each level of $\mathcal{T}_A$ and $\mathcal{T}_B$ are computed, and the results are reused. Since there are $O(\log N)$ levels, we cannot hope to compute all the matrices at each level in a constant-depth circuit. We give constant-depth threshold circuits that compute all nodes on only a constant number of levels of $\mathcal{T}_A$ and $\mathcal{T}_B$ while using a number of gates arbitrarily close to the total work performed by the fast matrix multiplication algorithm.

Our circuit computes the same $O(N^w)$ scalar products as the underlying fast matrix multiplication algorithm. These scalars correspond to the leaves of $\mathcal{T}_A$ and those of $\mathcal{T}_B$ respectively. Our algorithm processes these trees in a top-down fashion to compute the scalars, corresponding to the leaves. We then appeal to Lemma 3.3 to compute the product of each scalar corresponding to a leaf of $\mathcal{T}_A$ with the corresponding leaf of $\mathcal{T}_B$. For both the problems we solve, we will need to consider another tree with similar structure to $\mathcal{T}_A$ and $\mathcal{T}_B$; however, these trees will be used in different ways.

We assume, as in Section 2.3, that we have a fast matrix multiplication algorithm that multiplies two $T \times T$ matrices using r multiplications. We describe an improved TC^0 circuit for computing the values at the leaves of $\mathcal{T}_A$. Our results extend naturally to computing the leaves of $\mathcal{T}_B$. Level h of $\mathcal{T}_A$ contains r^h nodes, each corresponding to an $N/T^h \times N/T^h$ matrix. Moreover, each entry of each matrix at level h is the $\{-1, 1\}$ -weighted sum of at most T^{2h} entries of the root matrix, A . Hence if each entry of the integer matrix A requires at most b bits, the number of bits required for each entry of a matrix at level h is at most

$$\text{bits}((2^b - 1)T^{2h}) \leq b + \text{bits}(T^{2h}) = b + O(h \log T). \quad (2)$$

For our results we assume $b = O(\log N)$ bits. Moreover, T is a constant associated with the fast matrix multiplication selected, and $h \leq \log_T N$, hence the scalar values at the leaves of $\mathcal{T}_A$ require $O(\log N)$ bits.

We give a subcubic TC^0 circuit for computing the matrix product $C = AB$ in the next section, but first, we illustrate our main ideas by showing how to check $\text{trace}(A^3) \geq \tau$. As mentioned in Section 2.3, this allows us, for example, to count triangles in a graph. The bulk of our work lies in showing how to compute the $N^{\log_T r}$ scalar products prescribed by the fast matrix multiplication algorithm. Each such scalar product is between a weighted sum of entries of A and a weighted sum of entries of B . We next show how to compute these weighted sums for A with a circuit that computes a constant number of levels of $\mathcal{T}_A$. An analogous construction works for B .

4.2 Approach

Our main goal in the following sections is to give $O(\log \log N)$ -depth, $\tilde{O}(N^\omega)$ -gate and $O(d)$ -depth, $\tilde{O}(N^{\omega + O(\gamma^d)})$ -gate threshold circuits for multiplying two $N \times N$ integer matrices with $O(\log N)$ -bit entries, for a positive integer d and constant $\gamma < 1$.

We first motivate our approach by attempting to construct a constant-depth and $O(N^{3-\epsilon})$ -gate threshold circuit for matrix multiplication using Strassen's algorithm as a guide. As described in the previous section and Figure 2, our immediate goal is to compute the scalars associated with the leaves of $\mathcal{T}_A$ (and $\mathcal{T}_B$) for Strassen's algorithm by selecting a constant number of levels of $\mathcal{T}_A$ to compute. The most natural approach is perhaps to directly compute the leaves, at level $\log_2 N$, without computing any other level. The leaves of $\mathcal{T}_A$ correspond to $\{-1, 1\}$ -weighted sums of at most N entries of A . By recalling (2) and invoking Lemma 3.2, we can compute each such sum in depth 2 using $O(N \log N)$ gates. However, we must compute $O(N^{\log_2 7})$ such sums, yielding a total of $\tilde{O}(N^{1 + \log_2 7}) \approx \tilde{O}(N^{3.81})$ gates. This can be improved to $\approx \tilde{O}(N^{3.58})$ by observing that not all sums have the same number of summands; however, this still fails to achieve our goal.

We can improve the approach by employing addition circuits of depth greater than 2 due to Siu et al. [20] (Corollary 2). This allows computation of the desired sums in depth $O(d)$ using $O(dN^{1/d})$ gates, yielding the following result.

THEOREM 4.1. *Suppose we are given $N \times N$ integer matrices A and B with entries of size $O(\log N)$ bits. There is a threshold circuit of depth $O(d)$ that computes the matrix product AB using $\tilde{O}(dN^{\omega + 1/d})$ gates.*

We do not include a full proof of this theorem as our main results give superior results, both in terms of gate count and the simplicity of the resulting circuits. The results of the following sections show how to more carefully select a constant number of levels of $\mathcal{T}_A$ and $\mathcal{T}_B$ in order to improve the exponent in the gate count from $\omega + 1/d$ to $\omega + O(\gamma^d)$ for a constant $\gamma < 1$.

4.3 Matrix trace

We select t levels, $0 = h_0 < h_1 < h_2 < \dots < h_t$ of the tree $\mathcal{T}_A$. Our TC^0 circuit computes all of the matrices at these t levels of $\mathcal{T}_A$. Our goal is to compute the scalars corresponding to the leaves of $\mathcal{T}_A$, hence $h_t = \log_T N$. The benefit of computing level h_i is that each entry of each matrix at level h_{i+1} is then a $\{-1, 1\}$ -weighted sum of at most $T^{2(h_{i+1}-h_i)}$ matrices at level h_i .

Our results rely on parameters associated with our fast matrix multiplication algorithm. Recall s_A from Definition 2.1. We define $\alpha = r/s_A$ and $\beta = s_A/T^2$, and we have that $0 < \alpha \leq 1$ and $\beta \geq 1$ (for Strassen's algorithm, $\alpha = 7/12$ and $\beta = 3$).

LEMMA 4.2. *For $1 \leq i \leq t$, if the matrices at level h_{i-1} of $\mathcal{T}_A$ have been computed, then the matrices at level h_i can be computed in depth 2 using $O((b + h_{i-1})\alpha^{h_{i-1}}\beta^{h_i}N^2)$ gates.*

PROOF. The r^{h_i} nodes at level h_i of $\mathcal{T}_A$ each correspond to an $N/T^{h_i} \times N/T^{h_i}$ matrix. We set $\delta_i = h_i - h_{i-1}$ for convenience. We can associate each node u at level h_i with the unique subtree rooted at level h_{i-1} that contains it. The $N/T^{h_i} \times N/T^{h_i}$ matrix corresponding to u is a sum of at most $T^{2\delta_i}$ blocks of the $N/T^{h_{i-1}} \times N/T^{h_{i-1}}$ matrix associated with the root of the subtree containing u .

We seek a better bound on the number of such blocks we must sum to obtain the matrix associated with u . Let $\text{size}(u)$ represent this quantity, and let $\text{root}(u)$ be the node at level h_{i-1} on the path from u to the root of $\mathcal{T}_A$. Recall that each edge of $\mathcal{T}_A$ corresponds to one of the fast matrix multiplication expressions M_i and that a_i is the number of distinct blocks of A that appear in M_i (defined in Section 2.3). The quantity $\text{size}(u)$ is the product of the a_i associated with the edges on the path from u to $\text{root}(u)$ (see Figure 2). Thus for each node v at level h_{i-1} , we have:

$$\sum_{\{u | \text{root}(u)=v\}} \text{size}(u) = \sum_{m_1 + \dots + m_r = \delta_i} \binom{\delta_i}{m_1, \dots, m_r} \prod_{1 \leq j \leq r} a_j^{m_j} = s_A^{\delta_i}, \quad (3)$$

where the last equality follows from the multinomial theorem. We now bound the number of gates required to compute the matrices at level h_i . Since we assume the matrices at level h_{i-1} have been computed, by Lemma 3.2, each entry of the matrix associated with node u at level h_i can be computed using $O((b + h_{i-1})\text{size}(u))$ gates in depth 2. We charge the gate count for u to $\text{root}(u)$, and by (3) and (2), we have that the number of gates charged to each node at level h_{i-1} is

$$O((b + h_{i-1})s_A^{h_{i-1}}N^2/T^{2h_i}),$$

hence the total number of gates required for level h_i is

$$\begin{aligned} O((b + h_{i-1})r^{h_{i-1}}s_A^{h_{i-1}}N^2/T^{2h_i}) &= \\ O((b + h_{i-1})(r/s_A)^{h_{i-1}}(s_A/T^2)^{h_i}N^2), \end{aligned}$$

as desired. $\square$

Next we show how to set the h_i so that the number of gates required at each level is approximately balanced. This yields a total gate count that is, at worst, within a factor of t of the gate count for an optimal setting of the h_i . We must assume the number of multiplications our fast $T \times T$ matrix multiplication algorithm requires is greater than T^2 . The results, as stated and proven below, do not hold if we have an optimal fast matrix multiplication algorithm where the number of multiplications, $r = T^2$. We set $\gamma = \log_\beta(1/\alpha)$. Note that $0 < \gamma < 1$ since $r > T^2$ is equivalent to $\alpha\beta > 1$ (for Strassen's algorithm, $\gamma \approx 0.491$).

LEMMA 4.3. *Let $h_i = \lceil (1 - \gamma^i)\rho \rceil$, for some $\rho > 0$. Then all the matrices at levels $h_1, \dots, h_t$ of $\mathcal{T}_A$ can be computed in depth $2t$ using $O(t(\alpha\beta)^\rho(b + \log N)N^2)$ gates.*

PROOF. We have $h_i \leq \log_T N$ for all $0 \leq i \leq t$ since the latter is the height of $\mathcal{T}_A$. By Lemma 4.2, level h_i can be computed in depth 2 using $O((b + \log N)\alpha^{h_{i-1}}\beta^{h_i}N^2)$ gates.

Let $\tilde{h}_i = (1 - \gamma^i)\rho$. Observe that

$$\sum_{1 \leq i \leq t} \alpha^{h_{i-1}}\beta^{h_i} < \beta \sum_{1 \leq i \leq t} \alpha^{\tilde{h}_{i-1}}\beta^{\tilde{h}_i},$$

hence it suffices to bound $\sum_{1 \leq i \leq t} \alpha^{\tilde{h}_{i-1}}\beta^{\tilde{h}_i}$. The terms in this sum are all equal:

$$\alpha^{\tilde{h}_{i-1}}\beta^{\tilde{h}_i} = \left(\frac{\alpha\beta}{\alpha^{\gamma^{i-1}}\beta^{\gamma^i}} \right)^\rho = \left(\frac{\alpha\beta}{\alpha^{\gamma^{i-1}}(\beta\gamma)^{\gamma^{i-1}}} \right)^\rho = (\alpha\beta)^\rho,$$

so that $\sum_{1 \leq i \leq t} \alpha^{h_{i-1}}\beta^{h_i} = O(t(\alpha\beta)^\rho)$, from which the claim follows. $\square$

The above lemma establishes a tradeoff in the following sense. The value ρ impacts the total number of gates. However, we require that $h_t = \log_T N$, which imposes constraints on t and, consequently, the depth of the circuit. The larger ρ , the smaller t needs to be in order for $h_t = \log_T N$.

The natural strategy of taking $h_i = \lceil i \log_T N / t \rceil$ yields a weaker result, comparable to Theorem 4.1. We now establish our main theorems by better quantifying the tradeoff between ρ and t . For these theorems we assume we are given a fast matrix multiplication algorithm and take $\omega = \log_T r$.

THEOREM 4.4. *Suppose we are given an integer τ and an $N \times N$ integer matrix A with entries of size $O(\log N)$ bits. There is a threshold circuit of depth $O(\log \log N)$ that determines whether $\text{trace}(A^3) \geq \tau$ using $\tilde{O}(N^\omega)$ gates.*

PROOF. We appeal to Lemma 4.3, setting $\rho = \log_T N$. The gate bound follows from $(\alpha\beta)^\rho = (r/T^2)^{\log_T N} = N^{\omega-2}$. To bound the depth, we set $t = \lfloor \log_{1/\gamma} \log_T N \rfloor + 1 > \log_{1/\gamma} \log_T N$. This implies:

$$\log_T N - (1 - \gamma^t) \log_T N < \log_T N - (1 - 1/\log_T N) \log_T N = 1.$$

Thus $h_t = \lceil (1 - \gamma^t) \log_T N \rceil = \log_T N$ as desired.

This shows that we can compute the values corresponding to the leaves of $\mathcal{T}_A$ and $\mathcal{T}_B$ in the stated gate and depth bounds. One may see that each entry C_{ij} is a weighted sum of products, $\sum_{k \in I_{ij}} w_{ijk} p_k$, where each p_k corresponds to a product between a leaf of $\mathcal{T}_A$ and

the corresponding leaf of $\mathcal{T}_B$ with each weight $w_{ijk} \in \{-1, 1\}$. We seek to compute

$$\begin{aligned} \frac{\text{trace}(A^3)}{2} &= \sum_{i < j} A_{ij} C_{ij} = \sum_{i < j} A_{ij} \left(\sum_{k \in I_{ij}} w_{ijk} p_k \right) \\ &= \sum_k p_k \left(\sum_{i < j: k \in I_{ij}} w_{ijk} A_{ij} \right). \end{aligned} \quad (4)$$

Thus for each product, p_k , we want to multiply it with a $\{-1, 1\}$ -weighted sum over entries of A . We may compute these weighted sums independently and in parallel with those for A and B using the same techniques. Thus we seek to compute N^ω products of 3 $O(\log N)$ -bit numbers, and we appeal to Lemma 3.3 to accomplish this in depth 1 using a total of $\tilde{O}(N^\omega)$ gates. A final output gate sums the representations of the products computed by Lemma 3.3 and compares with the threshold τ . $\square$

We now prove our main theorem by exhibiting a more refined tradeoff between ρ and t .

THEOREM 4.5. *Suppose we are given an integer τ , an $N \times N$ integer matrix A with entries of size $O(\log N)$ bits, and a positive integer d . There is a threshold circuit of depth at most $2d + 5$ that determines whether $\text{trace}(A^3) \geq \tau$ using $\tilde{O}(dN^{\omega+c\gamma^d})$ gates, where $c > 0$ and $\gamma < 1$ are constants with respect to N and d that depend on the parameters of the fast matrix multiplication algorithm employed.*

PROOF. As for the previous theorem, we appeal to Lemma 4.3, this time setting $\rho = \log_T N + \varepsilon \log_{\alpha\beta} N$, for a constant $\varepsilon > 0$ whose value is given below. We have $(\alpha\beta)^\rho = (r/T^2)^{\log_T N} N^\varepsilon = N^{\omega-2+\varepsilon}$.

We set $\varepsilon = \gamma^d \log_T(\alpha\beta)/(1-\gamma) > \gamma^d \log_T(\alpha\beta)/(1-\gamma^d)$. This implies:

$$\begin{aligned} &\log_T N - (1-\gamma^d)(\log_T N + \varepsilon \log_{\alpha\beta} N) \\ &< \log_T N - (1-\gamma^d)(\log_T N + (\gamma^d/(1-\gamma^d)) \log_T(\alpha\beta) \log_{\alpha\beta} N) \\ &= \log_T N - (1-\gamma^d) \log_T N - \gamma^d \log_T N \\ &= 0, \end{aligned}$$

hence we may take $t < d$ in Lemma 4.3 in order to have $h_t = \log_T N$. The theorem follows from the argument used in the proof of Theorem 4.4 and taking $c = \log_T(\alpha\beta)/(1-\gamma)$ (for Strassen's algorithm, $c \approx 1.585$). $\square$

4.4 Matrix product

Now we describe how to compute the entries of the matrix product $C = AB$, where we assume the entries of the $N \times N$ matrices A and B require $O(\log N)$ bits. We define a tree $\mathcal{T}_{AB}$ with the same structure as $\mathcal{T}_A$ and $\mathcal{T}_B$. Each node of $\mathcal{T}_{AB}$ represents the product of the matrices of the corresponding nodes of $\mathcal{T}_A$ and $\mathcal{T}_B$. Hence the root of $\mathcal{T}_{AB}$ represents the matrix $C = AB$, and the leaves represent the $N^{\log_T r}$ scalar products computed by our fast matrix multiplication algorithm. We compute the root of $\mathcal{T}_{AB}$ in a bottom-up manner assuming that we are only computing the nodes at levels $\log_T N = h_t > h_{t-1} > \dots > h_1 > h_0 = 0$.

We let $\alpha_C = r/s_C$ and $\beta_C = s_C/T^2$ be parameters that are a function of the fast matrix multiplication algorithm employed. Recall

that from (2) we have that the scalars at the leaves of $\mathcal{T}_A$ and $\mathcal{T}_B$ each require $O(\log N)$ bits. Therefore the products of these scalars represented by the leaves of $\mathcal{T}_{AB}$ also require $O(\log N)$ bits.

We show that $\mathcal{T}_{AB}$ can be computed in a bottom-up manner with depth and gate bounds comparable to those we obtained for computing the leaves of $\mathcal{T}_A$ and $\mathcal{T}_B$ in the previous section. We will need a lemma analogous to Lemma 4.2.

LEMMA 4.6. *For $1 \leq i \leq t$, if the matrices at level h_i of $\mathcal{T}_{AB}$ have been computed, then the matrices at level h_{i-1} can be computed in depth 2 using $\tilde{O}(\alpha_C^{h_{i-1}} \beta_C^{h_i} N^2)$ gates.*

Proof: See the appendix.

Using the above lemma, the proof of Lemma 4.3 yields the following.

LEMMA 4.7. *Let $h_i = \lceil (1-\gamma^i)\rho \rceil$, for some $\rho > 0$. Then all the matrices at levels $h_1, \dots, h_t$ of $\mathcal{T}_{AB}$ can be computed in depth $2t$ using $\tilde{O}(t(\alpha\beta)^\rho N^2)$ gates.*

Armed with the above lemmas, we obtain our main results in similar fashion to those for the trace of A^3 . The structure of our circuit is that we compute the scalars corresponding to the leaves of $\mathcal{T}_A$ and $\mathcal{T}_B$ as described in the previous section. We then use Lemma 3.3 to compute the scalar products between corresponding leaves of $\mathcal{T}_A$ and $\mathcal{T}_B$ in depth 1. We finally apply the procedure outline above to compute the root of $\mathcal{T}_{AB}$, representing the matrix product AB , in a bottom-up manner. This essentially doubles the depth of the circuits we obtain compared to the corresponding circuits for the trace of A^3 .

THEOREM 4.8. *Suppose we are given $N \times N$ integer matrices A and B with entries of size $O(\log N)$ bits. There is a threshold circuit of depth $O(\log \log N)$ that computes the matrix product AB using $\tilde{O}(N^\omega)$ gates.*

THEOREM 4.9. *Suppose we are given $N \times N$ integer matrices A and B with entries of size $O(\log N)$ bits, and a positive integer d . There is a threshold circuit of depth at most $4d + 1$ that computes the matrix product AB using $\tilde{O}(dN^{\omega+c\gamma^d})$ gates, where $c > 0$ and $\gamma < 1$ are constants with respect to N and d that depend on the parameters of the fast matrix multiplication algorithm employed.*

5 MATRIX-MULTIPLICATION APPLICATION BACKGROUND

In this section, we provide more background on the relevance of dense matrix multiplication to deep learning. We also discuss the relevance of matrix multiplication to triangle counting in graphs and triangle counting's relevance to social network analysis.

Deep Learning. As mentioned in Section 1, our primary motivation for neural-circuit-based matrix multiplication is convolutional neural networks for deep learning. See Warden's clear explanation of the role of matrix multiplication in convolution steps for neural networks [25], which we summarize here. For more details see the Stanford course notes at <http://cs231n.github.io>. These networks assume the input is a two-dimensional image, with an $n \times n$ grid of pixels, each with ℓ channels. The neural networks usually refer to the number of channels as depth, but in this paper, "depth" refers

to the number of layers in our circuit. Typically the number of channels ℓ is a constant, but not necessarily just the three classic color channels (red, green, blue). In a convolutional step, we apply a set of K kernels to the image. Each kernel looks for a particular subpattern such as a horizontal edge or a splash of red. The kernel considers a small constant $q \times q$ submatrix of pixels (with ℓ channels) at a time and is applied across the whole image based on a stride. This recognizes the pattern no matter where it is in the image. For example, if the stride is four, then the kernel is applied to every fourth column and every fourth row. A place where the kernel is applied is called a *patch*. For each patch, for each kernel, a dot product scores the extent to which the patch matches the kernel. Computing all the kernels simultaneously is a matrix multiplication. The first matrix is $P \times Q$, where $P = O(n^2)$ is the number of patches and $Q = q \times q \times \ell$ is the number of elements in a kernel. The second matrix is $Q \times K$. This gives a $P \times K$ output matrix, giving the score for each patch for each kernel.

Let N be the largest matrix dimension and suppose we use a fast matrix multiplication algorithm that can multiply two $N \times N$ matrices in time $O(N^\omega)$. Our circuit requires fan-in as large as $O(N^\omega)$. These are gates at the end that compute the final output matrix entries. Two of the relevant matrix dimensions for convolutional neural networks, K and Q , are generally constants. The third dimension P is not. However, if the particular architecture can only support fan in x , we can break the matrix multiplication into independent pieces, each with at most $\sqrt{x}$ rows in the first matrix. These can run in parallel, so they have the same depth, given a large enough architecture. Thus the unbounded fan-in in our algorithm is not necessarily a practical limitation for our motivating application.

Social Network Analysis. Social networks of current interest are too large for our circuit methods to be practical for neuromorphic architectures in the near future. Also social network adjacency matrices are sparse, unlike the dense small matrices for convolutional neural networks we described above. Nevertheless, we briefly review the motivation for matrix multiplication in this setting. One application is computing the clustering coefficient of an N -node graph (or subgraph). The global clustering coefficient is the ratio of the number of triangles in the graph to the number of *wedges* (length-2 paths) in the graph. A degree- δ node is at the center of $\binom{\delta}{2}$ wedges. The global clustering coefficient is the fraction of total wedges in the graph that close into triangles. These triangles are common in social networks, where the central node of a wedge may introduce two neighbors. Social-network-analysis researchers believe a high global clustering coefficient (also called transitivity) means the graph has community structure. For example, Seshadri, Kolda and Pinar [18] assumed constant global clustering coefficients when proving a structural property of social networks they used for their BTER (Block Two-Level Erdős-Renyi) generative model. Orman, Labatut and Cherifi [16] empirically studied the relationship between community structure and clustering coefficient. They found that high clustering coefficients did imply community structure, although low clustering coefficients did not preclude it.

This paper considered the question: “Does a graph G have at least τ triangles?” The user can pick a value of τ that represents reasonable community structure for their particular kind of graph.

Usually they compute the total number of wedges D in $O(N)$ time and set τ to some function of D (perhaps just scaling by a constant).

6 OPEN PROBLEMS

The main open problem is whether we can do matrix multiplication with $O(N^\omega)$ gates in constant depth. Theorem 4.4 shows this can be done in $O(\log \log N)$ depth. Another open question is lower bounds: What is the minimum depth of a threshold circuit for computing matrix products using $O(N^{3-\epsilon})$ gates? Can one show that a constant-depth threshold circuit using $\tilde{O}(N^\omega)$ gates yields an $O(\log N)$ PRAM algorithm with $O(N^\omega)$ work?

One may show that our circuits are L -uniform. Can a stronger uniformity condition be imposed?

One advantage of neural networks is their low energy relative to CMOS-based electronics. One possible energy model for threshold gates is to charge a gate only if it fires [24]. That is, charge a gate one unit of energy for sending a signal if and only if the weighted sum of the inputs exceeds the threshold. What is the energy complexity of the kinds of matrix-multiplication circuits we consider?

ACKNOWLEDGEMENTS

This research was supported by the Laboratory Directed Research and Development program at Sandia National Laboratories, a multi-mission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy’s National Nuclear Security Administration under contract DE-NA0003525.

REFERENCES

- [1] Frederico A. C. Azevedo, Ludmila R. B. Carvalho, Lea T. Grinberg, José Marcelo Farfel, Renata E. L. Ferretti, Renata E. P. Leite, Wilson Jacob Filho, Roberto Lent, and Suzana Herculano-Houzel. 2009. Equal numbers of neuronal and nonneuronal cells make the human brain an isometrically scaled-up primate brain. *The Journal of Comparative Neurology* 513, 5 (April 2009), 532–541. <https://doi.org/10.1002/cne.21974>
- [2] Grey Ballard, Austin R. Benson, Alex Druinsky, Benjamin Lipshitz, and Oded Schwartz. 2016. Improving the Numerical Stability of Fast Matrix Multiplication. *SIAM J. Matrix Anal. Appl.* 37, 4 (Jan. 2016), 1382–1418. <https://doi.org/10.1137/15M1032168>
- [3] Dario Bini and Grazia Lotti. 1980. Stability of fast algorithms for matrix multiplication. *Numer. Math.* 36 (1980), 63–72.
- [4] Markus Bläser. 2013. *Fast Matrix Multiplication*. Number 5 in Graduate Surveys. Theory of Computing Library. <http://theoryofcomputing.org/articles/gs005/>
- [5] M. Davies, N. Srinivasa, T. H. Lin, G. Chinya, Y. Cao, S. H. Choday, G. Dimou, P. Joshi, N. Imam, S. Jain, Y. Liao, C. K. Lin, A. Lines, R. Liu, D. Mathaikutty, S. McCoy, A. Paul, J. Tse, G. Venkataramanan, Y. H. Weng, A. Wild, Y. Yang, and H. Wang. 2018. Loihi: A Neuromorphic Manycore Processor with On-Chip Learning. *IEEE Micro* 38, 1 (Jan. 2018), 82–99. <https://doi.org/10.1109/MM.2018.112130359>
- [6] Steve K Esser, Rathinakumar Appuswamy, Paul Merolla, John V Arthur, and Dharmendra S Modha. 2015. Backpropagation for energy-efficient neuromorphic computing. In *Advances in Neural Information Processing Systems*. 1117–1125.
- [7] Merrick Furst, James B. Saxe, and Michael Sipser. 1984. Parity, circuits, and the polynomial-time hierarchy. *Mathematical systems theory* 17, 1 (Dec. 1984), 13–27. <https://doi.org/10.1007/BF01744431>
- [8] Giacomo Indiveri, Bernabe Linares-Barranco, Tara Julia Hamilton, André van Schaik, Ralph Etienne-Cummings, Tobi Delbruck, Shih-Chii Liu, Piotr Dudek, Philipp Häfliger, Sylvie Renaud, Johannes Schemmel, Gert Cauwenberghs, John Arthur, Kai Hynna, Fopefolu Folowosele, Sylvain Saighi, Teresa Serrano-Gotarredona, Jayawan Wijekoon, Yingxue Wang, and Kwabena Boahen. 2011. Neuromorphic silicon neuron circuits. *Frontiers in Neuroscience* 5, 73 (2011). <https://doi.org/10.3389/fnins.2011.00073>
- [9] Daniel M. Kane and Ryan Williams. 2015. Super-Linear Gate and Super-Quadratic Wire Lower Bounds for Depth-Two and Depth-Three Threshold Circuits. *arXiv:1511.07860 [cs]* (Nov. 2015). <http://arxiv.org/abs/1511.07860> arXiv: 1511.07860.

- [10] Muhammad Mukaram Khan, David R Lester, Luis A Plana, A Rast, Xin Jin, Eustace Painkras, and Stephen B Furber. 2008. SpiNNaker: mapping neural networks onto a massively-parallel chip multiprocessor. In *Neural Networks, 2008. IJCNN 2008. (IEEE World Congress on Computational Intelligence). IEEE International Joint Conference on. IEEE*, 2849–2856.
- [11] François Le Gall. 2014. Powers of Tensors and Fast Matrix Multiplication. In *Proceedings of the 39th International Symposium on Symbolic and Algebraic Computation (ISSAC '14)*. ACM, New York, NY, USA, 296–303. <https://doi.org/10.1145/2608628.2608664>
- [12] Warren S. McCulloch and Walter Pitts. 1943. A logical calculus of the ideas immanent in nervous activity. *The bulletin of mathematical biophysics* 5, 4 (Dec. 1943), 115–133. <https://doi.org/10.1007/BF02478259>
- [13] Paul A Merolla, John V Arthur, Rodrigo Alvarez-Icaza, Andrew S Cassidy, Jun Sawada, Filipp Akopyan, Bryan L Jackson, Nabil Imam, Chen Guo, Yutaka Nakamura, et al. 2014. A million spiking-neuron integrated circuit with a scalable communication network and interface. *Science* 345, 6197 (2014), 668–673.
- [14] Robert C. Minnick. 1961. Linear-Input Logic. *IRE Trans. Electronic Computers* 10, 1 (1961), 6–16. <https://doi.org/10.1109/TEC.1961.5219146>
- [15] Saburo Muroga. 1959. The principle of majority decision logical elements and the complexity of their circuits. In *IFIP Congress*, 400–406.
- [16] Günce Keziban Orman, Vincent Labatut, and Hocine Cherifi. 2013. An empirical study of the relation between community structure and transitivity. *Studies in Computational Intelligence* 424 (2013), 99–110.
- [17] Catherine D. Schuman, Thomas E. Potok, Robert M. Patton, J. Douglas Birdwell, Mark E. Dean, Garrett S. Rose, and James S. Plank. 2017. A Survey of Neuromorphic Computing and Neural Networks in Hardware. (May 2017). <https://arxiv.org/abs/1705.06963>
- [18] C Seshadhri, Tamara G. Kolda, and Ali Pinar. 2012. Community structure and scale-free collections of Erdős-Rényi graphs. *Physical Review E* 85, 056109 (2012).
- [19] Jiri Šima and Pekka Orponen. 2003. General-Purpose Computation with Neural Networks: A Survey of Complexity Theoretic Results. *Neural Computation* 15, 12 (Dec. 2003), 2727–2778. <https://doi.org/10.1162/089976603322518731>
- [20] Kai-Yeung Siu, Vwani Roychowdhury, and Thomas Kailath. 1991. Depth-size tradeoffs for neural computation. *IEEE Trans. Comput.* 40, 12 (Dec. 1991), 1402–1412. <https://doi.org/10.1109/12.106225>
- [21] Kai-Yeung Siu, Vwani Roychowdhury, and Thomas Kailath. 1995. *Discrete Neural Computation: A Theoretical Foundation*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA.
- [22] Volker Strassen. 1969. Gaussian elimination is not optimal. *Numer. Math.* 13, 4 (Aug 1969), 354–356.
- [23] I. Sugiarto, G. Liu, S. Davidson, L. A. Plana, and S. B. Furber. 2016. High performance computing on SpiNNaker neuromorphic platform: A case study for energy efficient image processing. In *2016 IEEE 35th International Performance Computing and Communications Conference (IPCCC)*, 1–8. <https://doi.org/10.1109/IPCCC.2016.7820645>
- [24] Kei Uchizawa, Rodney Douglas, and Wolfgang Maass. 2006. Energy Complexity and Entropy of Threshold Circuits. In *Automata, Languages and Programming (Lecture Notes in Computer Science)*. Springer, Berlin, Heidelberg, 631–642. https://doi.org/10.1007/11786986_55
- [25] Pete Warden. 2015. Why GEMM is at the heart of deep learning. <https://petewarden.com/2015/04/20/why-gemm-is-at-the-heart-of-deep-learning/>. online, accessed February 9, 2017.
- [26] Andrew C. C. Yao. 1985. Separating the polynomial-time hierarchy by oracles. In *26th Annual Symposium on Foundations of Computer Science, 1985*. 1–10. <https://doi.org/10.1109/SFCS.1985.49>

APPENDIX

Proof of Lemma 4.6:

Our proof uses a similar analysis to that of Lemma 4.2. We need new parameters derived from our fast matrix multiplication algorithm. For $1 \leq j \leq T^2$, we use j to index the T^2 expressions for entries of C . We define c'_j as the number of M_i terms that appear in the j th expression for an entry of C . For Strassen's algorithm (Figure 1), we have $c'_1 = 4$, $c'_2 = 2$, $c'_3 = 2$, and $c'_4 = 4$. Recall the sparsity parameter s_C from Definition 2.1, and observe that $s_C = \sum_{1 \leq j \leq T^2} c'_j$.

We assume the matrix products at level h_i of $\mathcal{T}_{AB}$ have been computed and compute a node u at level h_{i-1} . We again define $\delta_i = h_i - h_{i-1}$ for convenience. As an example, suppose we were using the scalars corresponding to the leaves of $\mathcal{T}_{AB}$, at level h_t , to construct a $T \times T$ matrix at level $h_t - 1$. Each of the T^2 entries

of this matrix consists of a $\{-1, 1\}$ -weighted sum of the scalars at level h_t .

More generally, the matrix corresponding to node u at level h_{i-1} is composed of $T^{2\delta_i}$ blocks that are each $\{-1, 1\}$ -weighted sums of matrices of size $N/T^{h_i} \times N/T^{h_i}$ from level h_i . We seek to bound the number of terms in each such sum. Let u_l for $1 \leq l \leq T^{2\delta_i}$ correspond to the blocks of the matrix at node u , and let $\text{size}(u_l)$ be the number of terms in the weighted sum of matrices from level h_i that is equal to the block u_l . Using an approach similar to that from the proof of Lemma 4.2, we have:

$$\begin{aligned} \sum_{1 \leq l \leq T^{2\delta_i}} \text{size}(u_l) &= \sum_{m_1 + \dots + m_{T^2} = \delta_i} \binom{\delta_i}{m_1, \dots, m_{T^2}} \prod_{1 \leq j \leq T^2} (c'_j)^{m_j} \\ &= \left(\sum_{1 \leq j \leq T^2} c'_j \right)^{\delta_i} = s_C^{\delta_i}, \end{aligned} \quad (5)$$

where, as in the proof of Lemma 4.2, the penultimate equality follows from the multinomial theorem.

Each block u_l is of size $N/T^{h_i} \times N/T^{h_i}$, and by the discussion preceding this lemma and (2), we have that each entry of a block u_l requires $O(\log N)$ bits. Thus, by Lemma 3.2, we may compute all of the blocks u_l of u in depth 2 with a gate count of:

$$\begin{aligned} &\sum_{1 \leq l \leq T^{2\delta_i}} O(N^2/T^{2h_i} \log N \text{size}(u_l)) \\ &= O(N^2/T^{2h_i} \log N \sum_{1 \leq l \leq T^{2\delta_i}} \text{size}(u_l)) \\ &= O(N^2/T^{2h_i} \log N s_C^{\delta_i}), \end{aligned} \quad (6)$$

where the last equality follows from the above equation. Since there are $r^{h_{i-1}}$ nodes in total on level h_{i-1} , we may compute all the matrices on level h_{i-1} with a total gate count of,

$$\begin{aligned} O(r^{h_{i-1}} s_C^{h_i - h_{i-1}} N^2/T^{2h_i} \log N) &= O((r/s_C)^{h_{i-1}} (s_C/T^2)^{h_i} N^2 \log N) \\ &= \tilde{O}(\alpha_C^{h_{i-1}} \beta_C^{h_i} N^2). \end{aligned}$$

□