High Performance Monte Carlo Simulation of Ising Model on TPU Clusters

Kun Yang*
kuny.work@gmail.com
Google Research

Yi-Fan Chen†
yifanchen@google.com
Google Research

Georgios Roumpos
roumposg@google.com
Google Research

Chris Colby
ccolby@google.com
Google Research

John Anderson
janders@google.com
Google Research

ABSTRACT

Large-scale deep learning benefits from an emerging class of AI accelerators. Some of these accelerators' designs are general enough for compute-intensive applications beyond AI and Cloud TPU is one such example. In this paper, we demonstrate a novel approach using TensorFlow on Cloud TPU to simulate the two-dimensional Ising Model. TensorFlow and Cloud TPU framework enable the simple and readable code to express the complicated distributed algorithm without compromising the performance. Our code implementation fits into a small Jupyter Notebook and fully utilizes Cloud TPU's efficient matrix operation and dedicated high speed inter-chip connection. The performance is highly competitive: it outperforms the best published benchmarks to our knowledge by 60% in single-core and 250% in multi-core with good linear scaling. When compared to Tesla V100 GPU, the single-core performance maintains a $\sim 10\%$ gain. We also demonstrate that using low precision arithmetic—bfloat16—does not compromise the correctness of the simulation results.

CCS CONCEPTS

• **Computing methodologies** $\rightarrow$ **Distributed simulation; Massively parallel and high-performance simulations; Applied computing** $\rightarrow$ *Physics; Mathematics and statistics;*

KEYWORDS

Ising Model, Cloud TPU, Markov Chain Monte Carlo

ACM Reference Format:

Kun Yang, Yi-Fan Chen, Georgios Roumpos, Chris Colby, and John Anderson. 2019. High Performance Monte Carlo Simulation of Ising Model on TPU Clusters. In *The International Conference for High Performance Computing, Networking, Storage, and Analysis (SC '19)*, November 17–22, 2019, Denver, CO, USA. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3295500.3356149>

*This author is no longer affiliated with Google.

†Corresponding author.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
SC '19, November 17–22, 2019, Denver, CO, USA
© 2019 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-6229-0/19/11.
<https://doi.org/10.1145/3295500.3356149>

1 INTRODUCTION

The Ising model [13], which considers short-range interactions between spin variables on the sites of a d -dimensional lattice, plays an important role in statistical physics as a prototyping system to study the universal behavior of critical phenomena. Many significant breakthroughs in statistical physics are attributed to the study of the model from either its computational or its theoretical perspective. It is well known that the Ising model has no phase transition in one dimension; however, it undergoes a second-order phase transition between an ordered and a disordered phase in two dimensions or more [5, 18]. The critical temperature T_c at which this phase transition occurs on a two-dimensional square lattice was analytically solved by Lars Onsager [18], but it is still an open problem in three or more dimensions. Computer simulation offers a powerful alternative to study such systems and determine critical temperatures, thanks to the development of finite scaling theory [3] and availability of increasing computational power. This approach ushered in a plethora of interdisciplinary applications outside of physics, including bioinformatics [2], economics [21] and operations research [11, 23].

Large-scale simulation of systems such as Ising model requires a large amount of high performance computing resources, which are usually available in multi-core computing architectures based on distributed shared memory, or distributed clusters (a.k.a data-centers) with homogeneous or heterogeneous nodes commonly seen in private or commercial clouds. Benefiting from the explosion of machine learning, especially deep learning, commercial clouds provide not only CPUs and GPUs, but also specialized chips such as FPGAs and other in-house processors. The Tensor Processing Unit (“Cloud TPU” or “TPU” for short)—an AI application-specific integrated circuit (ASIC) developed by Google for neural network machine learning—has received much attention in the machine learning community [15, 16]. Its latest release, Cloud TPU v3, offers 420×10^{12} floating-point operations per second (FLOPS) and 128GB of high bandwidth memory (HBM)¹. Multiple units are connected to form a “POD” (Cloud TPU v3 Pod) through a dedicated high speed 2-D toroidal mesh network, allowing up to 100+ peta-FLOPS and 32TB of HBM¹ to be accessed by the application with very low latency and in lockstep. TPU is programmable via software frontends such as TensorFlow [1] or PyTorch [20], and can be deployed both for training huge deep neural networks and for performing low-latency

¹cloud.google.com/tpu/

online prediction. [14] reports impressive acceleration of training and online prediction.

With the tremendous amount of computation resources that TPU offers, it is compelling to also consider the opportunities TPU brings for applications beyond machine learning. The programming frontends that are used for TPU, such as TensorFlow, also offer a rich set of functionalities that are highly relevant for scientific computations. The TensorFlow TPU programming stack also provides the additional benefits of allowing distributed algorithms to be expressed with simple and easy-to-understand code without sacrificing performance. In addition, the ability to program conventional scientific simulations in TensorFlow framework makes it easier to explore the hybrid approaches employing both conventional scientific computation methods and modern machine learning techniques on the same framework.

Motivated by these observations, we developed a Single Instruction, Multiple Data (SIMD) distributed Markov Chain Monte Carlo (MCMC) simulation of the two-dimensional Ising model that is programmed in TensorFlow to run on TPU. We demonstrate that such an approach for scientific simulations is very promising. Our code implementation is easy to understand, with entire source code ~ 600 lines, while also achieves competitive performance and scaling properties when compared with the state-of-the-art benchmarks in terms of both speed and scalability. Another interesting observation is that the lower precision arithmetic—bfloat16 (1 sign bit, 8 exponent bits and 7 mantissa bits²) instead of float32—does not compromise the accuracy of the result.

The implementation is open-sourced (URL available in AD Appendix) and can be run through Colaboratory³, a free cloud service based on Jupyter Notebooks for interactive data science. It is also worth pointing out that all results, including large-scale distributed computation, can be run through the Notebook frontend with minimal setup (mainly for allocating TPU backends and setting up the connection to the TPU backend).

In the following sections, we first review the high level architecture of TPU framework (Sec. 2), then discuss the adaption of a widely used algorithm for the Ising model simulation to better leverage TPU architecture (Sec. 3). We then follow up with benchmark results (Sec. 4) and in Section 5 highlight the implementation details and high-level performance analysis. Finally, we conclude the paper with Section 6, outlining our views on how the recent developments in both software and hardware for machine learning applications can impact the scientific computation applications.

2 TPU DEVICE ARCHITECTURE

TPU is a programmable linear algebra accelerator optimized for machine learning workloads. In the Cloud TPU v3 architecture, one TPU unit consists of four TPU chips on a board (we will use “TPU unit” or just “unit” to refer to such a board throughout this paper), and each TPU chip contains two TensorCores. Those TensorCores are treated as independent processors that communicate to each other through a dedicated high-bandwidth low-latency inter-chip network. In a larger system, more than 1000 TPU chips are packed

on a two-dimensional toroidal mesh inter-chip interconnect network to form a TPU cluster known as a “TPU Pod” [7]. In addition, each TPU unit is paired up with a TPU host server, with CPU, memory and disk resources attached. The communication links between TPU unit and its host server is through the regular data center connections. The host server can be used for data preparation, I/O task while TPU can be leveraged for compute-intensive tasks. The dedicated two-dimensional toroidal mesh allows all TensorCores in a “Pod” to work in lockstep efficiently, without going through the host servers. This large-scale high-bandwidth low-latency connectivity between TPU chips provides a significant advantage to achieve the strong linear-scaling performance reported in this paper.

The TensorCore, depicted in Figure 1, is optimized for dense linear algebra computations. It contains distinct classes of computing units, such as a scalar processor, a vector processor, accumulators and matrix units [9, 15, 16]. All those processors are backed by its 16GB High-Bandwidth Memory (HBM). Vectorized operations are handled either in the vector processor directly or forwarded to corresponding extended vector units. Each extended vector unit takes the input operands, performs the corresponding operations, and returns the results back to the vector processor. One of these extended vector units is the matrix unit (MXU), which is capable of performing 128×128 multiply-accumulate operations in each cycle [8]. The MXU is the main computing power of the TPU architecture, so it should be exploited as much as possible. While its inputs and outputs are 32-bit floating point values, the MXU rounds inputs down to bfloat16—a 16-bit floating point representation that provides better training and model accuracy than the IEEE half-precision representation—before multiplying.

Machine learning research shows that many machine learning models can tolerate lower precision arithmetic without degradation of converged accuracy. TPU supports storing values in bfloat16 format as a way to reduce the size of data and allow larger models to fit in memory [10]. For scientific computing, however, low precision is potentially dangerous because the increasing rounding error can introduce significant bias or divergent computation. In the case of the Ising model, while the binary spin values can be encoded in bfloat16 without loss of accuracy, the acceptance ratio and the random numbers used to determine acceptance in an MCMC simulation are more sensitive to reduced precision. However, our experiments show no noticeable differences in accuracy between bfloat16 and float32. By using bfloat16 instead of float32, we are able to simulate larger systems and leverage the MXU more effectively.

We program TPU through TensorFlow. The flow to run programs on TPU is roughly depicted in Figure 2. More details are available in TensorFlow’s official XLA page⁴. In the first stage, TensorFlow constructs the computation graph and marks the graph for replication. Then the graph is rewritten to be TPU-compatible and compiled to a High Level Optimizer (HLO) program. Next, the Accelerated Linear Algebra (XLA) compiler takes over and converts HLO operations to Low Level Optimizer (LLO) code—effectively “TPU assembly code”, which can be readily executed on TPU. The graph construction and the compilation occurs on host server and incurs over-head cost. But once the compiled LLO code is deployed to TPU’s the

²en.wikipedia.org/wiki/Bfloat16_floating-point_format

³<https://colab.research.google.com>

⁴www.tensorflow.org/xla

Figure 1: One TPU chip consists of two TensorCores. A TensorCore in 3rd generation TPU (TPU v3) consists of a scalar processor, vector processor, and two matrix units. The arrows depict the datapaths across different processors/units and the high-bandwidth memory (HBM), the diagram is borrowed from [9], here “Core” is the “Tensorcore”.

Figure 2: TPU software architecture, consisting of the neural network model or other computational task, TensorFlow client, TensorFlow server and XLA compiler [9].

computation step can repeat as many times as required without the intervention from the host servers.

Importantly XLA also provides communication primitives such as `CollectivePermute`⁵ and `AllToAll`⁶. These primitives are implemented over the dedicated high-bandwidth low-latency inter-chip interconnect network so the communication between any TPU chips within a TPU Pod is extremely efficient.

Another important detail that can have a critical impact on performance and memory usage is the choice of the shape of the tensor variables used in the program (expressed as TensorFlow tensors). According to the performance guide [8], unlike most other architectures, arrays in TPU are tiled in two dimensions. This entails to padding one dimension to a multiple of 8, and the other dimension to a multiple of 128. XLA performs data layout transformations and data are arranged in memory such that the hardware can efficiently process them. Programs that operate on array sizes undividable by 8 will have sub-optimal performance.

3 THE ISING MODEL

Mathematically, the Ising spin Hamiltonian is given by

$$H(\sigma) = -J \sum_{\langle ij \rangle} \sigma_i \sigma_j - \mu \sum_{i=1}^N \sigma_i$$

where σ_i is a random variable assuming the values of ± 1 on sites $i = 1, \dots, N$ of a d -dimensional hypercubic lattice, and $\langle ij \rangle$ indicates that sites i and j are nearest neighbors. (Note that throughout this paper, we will use the bold σ to represent a tensor of spins, while the elements within it will be represented as regular σ). The first term, where the sum is over pairs of nearest-neighbor sites, represents the interaction energy that favors an ordered ferromagnetic state (if $J > 0$). The second term, involving the interaction between the applied field and the spin system, is of a paramagnetic character. The configuration probability is given by the Boltzmann distribution with inverse temperature $\beta = (k_B T)^{-1}$:

$$\pi(\sigma) = \frac{e^{-\beta H(\sigma)}}{Z_\beta}$$

⁵https://www.tensorflow.org/xla/operation_semantics#collectivepermute

⁶https://www.tensorflow.org/xla/operation_semantics#alltoall

where $Z_\beta = \sum_{\sigma} e^{-\beta H(\sigma)}$ is the partition function and k_B is the Boltzmann constant. For a function f of the spins (“observable”), we denote $\langle f \rangle = \sum_{\sigma} f(\sigma) \pi(\sigma)$ the expectation (mean value) of f . $\langle f \rangle$ can often be difficult to evaluate numerically if there are many states in the system. The Markov Chain Monte Carlo is the most commonly used monte carlo algorithm to calculate statistics on the Ising model.

Without loss of generality, in what follows, we assume no external magnetic field, i.e., $\mu = 0$ and $J = 1$, and the 2D lattice has circular boundary or in other words, is a torus. A given configuration of the lattice (spin values) is represented by matrix σ .

3.1 Checkerboard Algorithm

The Metropolis-Hastings algorithm is the standard algorithm to simulate the Ising model. At each step, it proposes a candidate spin and flips the candidate based on the energy difference and acceptance probability. Closely related to this vanilla version that flips one spin at each step, there is another similar algorithm by flipping non-interacting spins in parallel, i.e., the efficient checkerboard algorithm [22].

Like the checkerboard above (Figure 3-(1)), spins in the lattice are colored black and white. The energy difference by flipping a spin of one color is completely described by its 4 neighbors of the opposite color. Thus, by fixing all spins of one color, the spins of the opposite color have no interactions with each other, and can be updated independently using Metropolis-Hastings. This observation leads to a highly parallel algorithm by alternating the 2 sub-routines below:

- Fixing all black spins, flip each white spin based on Metropolis Hastings in parallel.
- Fixing all white spins, flip each black spin based on Metropolis Hastings in parallel.

Write $\pi(\sigma) = \pi(\sigma_w, \sigma_b)$, where σ_w are values of all white spins and σ_b are values of all black ones. The transition kernel is,

$$P\{(\sigma_w, \sigma_b) \rightarrow (\sigma_w^*, \sigma_b^*)\} = P\{(\sigma_w, \sigma_b) \rightarrow (\sigma_w^*, \sigma_b^*)\} P\{(\sigma_w^*, \sigma_b) \rightarrow (\sigma_w^*, \sigma_b^*)\} \quad (1)$$

and

$$P\{(\sigma_w, \sigma_b) \rightarrow (\sigma_w^*, \sigma_b^*)\} = \prod_{i \in w} \min\{1, e^{\beta(\sigma_i^* - \sigma_i)} \cdot \text{nn}(i)\}$$

$$P\{(\sigma_w^*, \sigma_b) \rightarrow (\sigma_w^*, \sigma_b^*)\} = \prod_{i \in b} \min\{1, e^{\beta(\sigma_i^* - \sigma_i)} \cdot \text{nn}(i)\}$$

where $\text{nn}(i)$ is the sum of neighbor spins of i . By conditional probability decomposition, it is easy to show that

$$\pi(\sigma_w^*, \sigma_b^*) = \sum_{\sigma_w, \sigma_b} P\{(\sigma_w, \sigma_b) \rightarrow (\sigma_w^*, \sigma_b^*)\} \pi(\sigma_w, \sigma_b)$$

thus, the transition kernel satisfies the stationary distribution. The proof is based on Metropolis-Within-Gibbs Sampler [17] and is included in the supplemental materials.

3.2 Computation

The compute-intensive part of the checkerboard algorithm is the computation of the sum of neighbor values for each spin. To leverage the MXU, for a given lattice with size $[128 \times m, 128 \times n]$ in a

TPU core, we divide it into a $[m, n]$ grid of 128×128 sub-lattices (Figure 3-(1)). Define the kernel matrix K as,

$$K = \begin{bmatrix} 0 & 1 & 0 & \dots & 0 & 0 & 0 \\ 1 & 0 & 1 & \dots & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & \dots & 1 & 0 & 1 \\ 0 & 0 & 0 & \dots & 0 & 1 & 0 \end{bmatrix}_{128 \times 128}$$

then for a given sub-lattice σ_{ij} , $\text{matmul}(\sigma_{ij}, K) + \text{matmul}(K, \sigma_{ij})$ calculates the sum of nearest neighbors for all its internal sites. However, for the boundary sites, half of their nearest neighbors are missing from the sums. Those neighbors are on the boundaries of neighboring sub-lattices and ought to be sliced out and added into the sums. The whole lattice is updated one color at a time. To fix the spins of one color on the checkerboard, we multiply flip probabilities with a binary mask M , where

$$M = \begin{bmatrix} 1 & 0 & 1 & \dots & 0 & 1 & 0 \\ 0 & 1 & 0 & \dots & 1 & 0 & 1 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 1 & 0 & 1 & \dots & 0 & 1 & 0 \\ 0 & 1 & 0 & \dots & 1 & 0 & 1 \end{bmatrix}_{128 \times 128}$$

In summary, the algorithm to update the lattice given the color c and configuration σ is given in Algorithm 1. By alternating colors with different c , the Ising model is properly simulated. However, in Algorithm 1, there are several redundant calculations that are subject to further optimization: 1) Line 1 generates the probability for all spins, while only the spins colored by c are eligible for flipping. 2) Lines 2-6 calculate the nearest neighbor sums of all spins, however, only spins colored by c are updated. 3) Lines 8-9 generate flips but multiplying mask to fix the opposite color is expensive. To eliminate all the redundancies above, we reorganize the lattice and represent it in a compact way, i.e., the lattice is instead split into a $[m', n']$ grid of $[256, 256]$ sub-lattices, and define 4 “compact” $[m', n', 128, 128]$ sub-lattices (Figure 3-(2)),

$$\hat{\sigma}_{00} = \sigma[:, :, 0 :: 2, 0 :: 2], \hat{\sigma}_{01} = \sigma[:, :, 0 :: 2, 1 :: 2]$$

$$\hat{\sigma}_{10} = \sigma[:, :, 1 :: 2, 0 :: 2], \hat{\sigma}_{11} = \sigma[:, :, 1 :: 2, 1 :: 2]$$

and kernel,

$$\hat{K} = \begin{bmatrix} 1 & 1 & 0 & \dots & 0 & 0 & 0 \\ 0 & 1 & 1 & \dots & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & \dots & 0 & 1 & 1 \\ 0 & 0 & 0 & \dots & 0 & 0 & 1 \end{bmatrix}_{128 \times 128}$$

thus, $\hat{\sigma}_{00}$ and $\hat{\sigma}_{11}$ are all “black” spins, and $\hat{\sigma}_{01}$ and $\hat{\sigma}_{10}$ are all “white” spins. Now, it is trivial to show that the nearest neighbor sums of internal spins of those 4 “compact” sub-lattices are

$$\text{nn}(\hat{\sigma}_{00}) = \text{matmul}(\hat{\sigma}_{01}, \hat{K}) + \text{matmul}(\hat{K}^T, \hat{\sigma}_{10})$$

$$\text{nn}(\hat{\sigma}_{11}) = \text{matmul}(\hat{K}, \hat{\sigma}_{01}) + \text{matmul}(\hat{\sigma}_{10}, \hat{K}^T)$$

$$\text{nn}(\hat{\sigma}_{01}) = \text{matmul}(\hat{\sigma}_{00}, \hat{K}^T) + \text{matmul}(\hat{K}^T, \hat{\sigma}_{11})$$

$$\text{nn}(\hat{\sigma}_{10}) = \text{matmul}(\hat{K}, \hat{\sigma}_{00}) + \text{matmul}(\hat{\sigma}_{11}, \hat{K})$$

and their boundary spins are corrected using a similar approach as in Algorithm 1. The complete algorithm is presented in Algorithm 2.

Figure 3: A 2-d checkerboard: (1) Original checkerboard: on the left, the 16×16 board is split into a 4×4 grid of 4×4 sub-lattices, i.e., it is represented by a $[4, 4, 4, 4]$ tensor, where $[l, k, :, :]$ is the sub-lattice at $[l, k]$ of the grid; on the right, the sub-lattice is zoomed in and the indices of its spin sites are shown; (2) Reorganized checkerboard: on the left, each 4×4 sub-lattice is reorganized by 4 “compact” 2×2 sub-lattices; on the right, 4 “compact” 2×2 sub-lattices are zoomed in and their original indices from the 4×4 sub-lattice are shown. In general, such alternate coloring of black and white can be extended to lattices with any dimensions.

According to our experiments, it is about 3x faster than Algorithm 1 and has less memory footprint as it uses less temporary HBM.

Algorithm 1: Subroutine **UpdateNaive**(c, σ)

```

input :  $c$ : color (black or white)
         $\sigma \in \{0, 1\}^{m \times n \times 128 \times 128}$ : a configuration
output: A new configuration
//  $\cdot$  is element-wise multiplication,  $:$  or  $:s$  is
// array slicing.
1 probs = random_uniform( $[m, n, 128, 128]$ )  $\in [0, 1]$ 
// nn( $\sigma$ ) is the sum of nearest neighbors of each
// site, it has the same size as the lattice.
// Here,  $K$  is applied to each sub-lattice.
2 nn( $\sigma$ ) = matmul( $\sigma, K$ ) + matmul( $K, \sigma$ )
// Compensate the northern boundaries of each
// sub-lattice.
3 nn( $\sigma$ )[ $0, :, 0, :$ ] += { $\sigma[-1, :, -1, :]$ ,  $\sigma[:, -1, :, -1, :]$ }
// Compensate the southern boundaries of each
// sub-lattice.
4 nn( $\sigma$ )[ $1, :, -1, :$ ] += { $\sigma[1, :, 0, :]$ ,  $\sigma[:, 1, :, 0, :]$ }
// Compensate the western boundaries of each
// sub-lattice.
5 nn( $\sigma$ )[ $0, :, :, 0$ ] += { $\sigma[:, -1, :, -1]$ ,  $\sigma[:, :, -1, -1]$ }
// Compensate the eastern boundaries of each
// sub-lattice.
6 nn( $\sigma$ )[ $0, :, :, -1$ ] += { $\sigma[:, 1, :, 0]$ ,  $\sigma[:, :, 1, 0]$ }
7 acceptance_ratio = exp( $-2 \cdot \text{nn}(\sigma) \cdot \sigma$ )
8 mask =  $M$  if  $c$  is black else  $1 - M$ 
9 flips = (probs < acceptance_ratio)  $\cdot$  mask
10 return  $\sigma - 2 \cdot \text{flips} \cdot \sigma$ 

```

4 SIMULATION RESULTS

4.1 Correctness

The average magnetization per spin for a given β is defined as,

$$m(T) = m(\beta) = \langle \sigma \rangle = \frac{1}{N} \sum_i \sigma_i$$

where $N = n^2$ for a square lattice with size n . It can be shown that below the critical temperature, i.e., $(k_B \beta)^{-1} = T < T_c = (k_B \beta_c)^{-1} = \frac{2}{k_B \ln(1+\sqrt{2})}$, there is spontaneous magnetization—the interaction among spins is sufficiently large to cause neighboring spins to spontaneously align. On the other hand, thermal fluctuations completely eliminate any alignment above the critical temperature. Moreover, at the critical temperature, there is a discontinuity in the first derivative of $\langle m(T) \rangle$ with respect to T_c . This discontinuity generates a downward drop in the average magnetization. The sudden loss of spontaneous magnetization above the critical temperature is a signature of phase transition. Besides average magnetization, a more sensitive test of correctness is the Binder parameter [4], which is given by

$$U_4(T) = 1 - \frac{\langle m(T)^4 \rangle}{3 \langle m(T)^2 \rangle^2}$$

namely, the *kurtosis* of $m(T)$. It is frequently used to accurately determine phase transition points in numerical simulations of various models.

We verify the correctness of our algorithm and implementation by computing both quantities at various temperatures on different sizes of square lattices. As shown in Figure 4, the curves of average magnetization, with subtle differences, overlap with each other, and those of Binder parameters cross the critical line almost perfectly. Additionally, we investigate the implications of lower precision, i.e., bfloat16, for the estimation accuracy. In MCMC simulation of Ising model, lower precision has impact on the calculation of acceptance ratio and the random number generation, the bias introduced might

Algorithm 2: Subroutine **UpdateOptim**(c, σ)

input : c : color (*black* or *white*)
 $\hat{\sigma}_{00}, \hat{\sigma}_{01}, \hat{\sigma}_{10}, \hat{\sigma}_{11} \in \{0, 1\}^{m' \times n' \times 128 \times 128}$: a configuration
output: A new configuration

```

1 probs0 = random_uniform([ $m', n', 128, 128$ ])  $\in [0, 1)$ 
2 probs1 = random_uniform([ $m', n', 128, 128$ ])  $\in [0, 1)$ 
3 if black then
4    $\hat{\sigma}_0 = \hat{\sigma}_{00}$ 
5    $\hat{\sigma}_1 = \hat{\sigma}_{11}$ 
6    $nn0 = \text{matmul}(\hat{\sigma}_{01}, \hat{K}) + \text{matmul}(\hat{K}^T, \hat{\sigma}_{10})$ 
7    $nn0[:, :, 0, :] += \{\hat{\sigma}_{10}[:, :, -1, :], \hat{\sigma}_{10}[:, -1, :, -1, :]\}$ 
8    $nn0[:, :, :, 0] += \{\hat{\sigma}_{01}[:, -1, :, -1], \hat{\sigma}_{01}[:, :, -1, :]\}$ 
9    $nn1 = \text{matmul}(\hat{K}, \hat{\sigma}_{01}) + \text{matmul}(\hat{\sigma}_{10}, \hat{K}^T)$ 
10   $nn1[:, :, -1, :] += \{\hat{\sigma}_{01}[:, 1, :, 0], \hat{\sigma}_{10}[:, 1, :, 0, :]\}$ 
11   $nn1[:, :, :, -1] += \{\hat{\sigma}_{01}[:, :, 1, :, 0], \hat{\sigma}_{01}[:, :, :, 1, :]\}$ 
12 else
13   $\hat{\sigma}_0 = \hat{\sigma}_{01}$ 
14   $\hat{\sigma}_1 = \hat{\sigma}_{10}$ 
15   $nn0 = \text{matmul}(\hat{\sigma}_{00}, \hat{K}^T) + \text{matmul}(\hat{K}, \hat{\sigma}_{11})$ 
16   $nn0[:, :, 0, :] = \{\hat{\sigma}_{11}[:, -1, :, -1, :], \hat{\sigma}_{11}[:, -1, :, -1, :]\}$ 
17   $nn0[:, :, :, -1] = \{\hat{\sigma}_{00}[:, 1, :, 0], \hat{\sigma}_{00}[:, :, 1, :, 0]\}$ 
18   $nn1 = \text{matmul}(\hat{K}, \hat{\sigma}_{00}) + \text{matmul}(\hat{\sigma}_{11}, \hat{K})$ 
19   $nn1[:, :, -1, :] = \{\hat{\sigma}_{00}[:, 1, :, 0], \hat{\sigma}_{00}[:, 1, :, 0, :]\}$ 
20   $nn1[:, :, :, 0] = \{\hat{\sigma}_{11}[:, -1, :, -1], \hat{\sigma}_{11}[:, -1, :, -1]\}$ 
21 end
22 acceptance_ratio0 =  $\exp(-2 \cdot nn0 \cdot \hat{\sigma}_0)$ 
23 acceptance_ratio1 =  $\exp(-2 \cdot nn1 \cdot \hat{\sigma}_1)$ 
24 flips0 = (probs0 < acceptance_ratio0)
25 flips1 = (probs1 < acceptance_ratio1)
26 return  $(\hat{\sigma}_0 - 2 \cdot \text{flips0} \cdot \hat{\sigma}_0), (\hat{\sigma}_1 - 2 \cdot \text{flips1} \cdot \hat{\sigma}_1)$ 

```

accumulate and decrease the overall accuracy. However, in our experiments, all curves generated in bfloat16 and float32, especially those of Binder parameters, are almost the same and sharp turns around critical line are clearly observable. Based on this evidence, we argue that using bfloat16 has negligible impact on Ising model simulation, and in turn it offers two benefits: 1). We are able to simulate a larger lattice on a single TPU core because of smaller memory footprint, and 2). bfloat16 matrix multiplication with 32-bit accumulation is very efficient in MXU, while float32 matrix multiplication is more expensive as several bfloat16 passes are required.

4.2 Benchmarks

Preis et al. [22] showed impressive acceleration in their single GPU implementation of checkerboard algorithm using the compute unified device architecture (CUDA). By exploiting GPU's large pool of threads, its single-instruction multiple thread (SIMT) unit and memory hierarchy, their algorithm achieved 60x speedup on single GPU compared to its CPU counterpart. In their follow-up paper [6], their original algorithm was modified to overcome the memory limitations of single GPU. The improved algorithm was able to

simulate significantly larger systems and reached a performance of ~ 7.9774 flips/ns throughput in its best performing variant. By combining CUDA with MPI on the CPU level, their distributed algorithm achieved 206 flips/ns on a $800,000^2$ lattice. Besides the work on GPU, another encouraging line of research is the use of field-programmable gate array (FPGA). A recent implementation achieves ~ 614.4 flips/ns throughput, see [19] and references therein.

To quantify the performance of our implementation, we run our algorithm on TPU v3 using single core and multiple cores on TPU v3 clusters. As in [6, 19, 22], we measure the time spent on one sweep update, i.e., one update on all “black” spins plus one update on all “white” ones, and compute the average number of flips per nanosecond by dividing with n^2 .

4.2.1 Single TPU Core. First, we simulate the system on a single TPU v3 core (half TPU v3 chip). Our choice of lattice size is compatible with MXU's registers and achieves 100% memory capacity utilization according to our profilings. We can simulate lattice with size up to $(656 \times 128)^2 = 83,968^2$, which consumes 96% of the memory. The benchmarks in various sizes from $(20 \times 128)^2$ to $(640 \times 128)^2$ are summarized in Table. 1, the lattice size and flips in nanoseconds increase in tandem, as more computation is spent on matrix multiplication. For comparison, we also implemented the algorithms based on [6, 22] under CUDA 10.1, and using its cuRand package for random number generation and its Thrust package for reductions. To avoid the excessive temporary memory allocation on GPU, we wrote a custom memory allocator to reuse temporary memory. The benchmark we obtained on NVIDIA's Tesla V100, which powered by NVIDIA's latest Volta architecture, is 11.3704 flips/ns. Other older benchmarks published in [6, 19, 22] are also listed for reference.

It is also interesting to estimate the energy consumption in the computation. Assuming the average power consumption during the operation to be $\mathcal{P}W$ and the throughput achieved to be $\mathcal{F}$ flips/ns, the corresponding energy used is $(\mathcal{P}/\mathcal{F})nJ/\text{flip}$. The actual average power consumption depends on many factors and usually sophisticated modeling and measurements are needed. However, for a rough estimate of the upper bound, we use 250 W for GPU (based on NVIDIA's Tesla V100 Spec for PCIe version max power consumption⁷). While Google doesn't publish the number for TPU v3, it has been estimated to be 200W⁸ for a TPU v3 chip, or equivalently 100W for a TPU v3 core.

4.2.2 Linear Scaling on Multiple TPU cores. The idea to simulate Ising model on TPU v3 clusters is to split the whole lattice into sub-lattices, and exchange their boundary values to calculate the nearest-neighbor sums and update the sub-lattices in each core independently. The challenges are to handle the synchronization among the cores: 1). Block the update of sub-lattices until all the nearest-neighbor sums are calculated. 2). Block next iteration before the update of all sub-lattices is completed. Fortunately, such synchronization is already implemented in TensorFlow op `collective_permute`, which is used to exchange data by specifying source and target cores.

⁷images.nvidia.com/content/technologies/volta/pdf/437317-Volta-V100-DS-NV-US-WEB.pdf

⁸www.nextplatform.com/2018/05/10/tearing-apart-googles-tpu-3-0-ai-coprocessor/

Figure 4: Binder parameter $U_4(T)$ and average magnetization $m(T)$ as a function of T/T_c for various sizes of the two dimensional square lattice Ising model. Each data point is calculated by a Markov Chain of 1,000,000 samples using checkerboard update, where the first 100,000 samples are discarded for burn-in, and the rest 900,000 samples are used for the calculation. The plots of $U_4(T)$ for various lattice sizes cross almost perfectly at the critical temperature, which is shown additionally as a black dashed line, and their float32 and bfloat16 versions almost completely match. The plots of $m(T)$ show vanishing magnetization above critical temperature, but there are subtle differences between float32 and bfloat16 as $m(T)$ is a less sensitive test.

In its current release, TPUs in a TPU Pod are organized into a grid and each TPU core has an associated coordinate. All cores communicate through a specialized high-speed interconnect. In our experiments, we use smaller sections of a pod called *slices* [7], and show that the TPU Pod interconnect makes the overhead of exchanging boundary values between cores negligible. As a result, we observe strict linear scaling of flips/ns to the number of TPU cores, as shown in Table. 2.

In [6], the communications between GPUs are handled by MPI through the hosts, which is potentially a bottleneck in the simulations. Another interesting comparison is to benchmark their multi-GPU algorithm using Nvidia NVLink Fabric, which enables the interconnect of 8 GPUs, or Nvidia NVSWITCH that support the interconnect of 16 GPUs [12]. However, as the code for multi-GPU simulation is not available and the limitations on the number of interconnect GPUs, we leave it for our future work.

5 IMPLEMENTATION HIGHLIGHTS AND PERFORMANCE ANALYSIS

In the previous section, we report the competitive performance numbers achieved in our experiments. It is worthwhile going into some depth to highlight the implementation design choices and analyze the performance, in order to provide more insights on this new approach towards scientific simulations. The entire source code is also available through AD appendix.

5.1 Implementation Highlights: TensorFlow, SIMD, Highspeed Mesh Network

While most commonly seen applications of TensorFlow are in the context of machine learning, the functionalities available are highly relevant to various scientific applications. At a high level, TensorFlow provides various transformation operations that can be

lattice size n^2	(flips/ns)	(nJ/flip)
$(20 \times 128)^2$	8.1920	12.2070
$(40 \times 128)^2$	9.3623	10.6811
$(80 \times 128)^2$	12.3362	8.1062
$(160 \times 128)^2$	12.8266	7.7963
$(320 \times 128)^2$	12.9056	7.7486
$(640 \times 128)^2$	12.8783	7.7650
GPU in [6, 22]	7.9774	–
Nvidia Tesla V100	11.3704	21.9869
FPGA in [19]	614.4	–

Table 1: The computation throughput (flips/ns) and the estimated energy consumption upper bound (nJ/flip) with different sizes of the square lattice on a single TPU v3 core (half TPU v3 chip). Not comparing to FPGA, a single TensorCore sustains more flips/ns at all but the two smallest lattice sizes and consistently shows better energy efficiency.

#cores	lattice size	time of whole lattice update (ms)	throughput (flips/ns)	energy consumption (nJ/flip)
$1 \times 1 \times 2$	$(896 \times 128)^2$	574.7	22.8873	8.7385
$2 \times 2 \times 2$	$(1792 \times 128)^2$	574.9	91.5174	8.7415
$4 \times 4 \times 2$	$(3584 \times 128)^2$	575.0	366.0059	8.7430
$8 \times 8 \times 2$	$(7168 \times 128)^2$	575.2	1463.5146	8.7461
$16 \times 16 \times 2$	$(14336 \times 128)^2$	575.3	5853.0408	8.7476
64 GPUs [6]	$800,000^2$	~ 3000	206	–

Table 2: Each core contains a $[896 \times 128, 448 \times 128]$ sub-lattice. Hence, for a $n \times n \times 2$ -core cluster, the lattice size is $(512 \times 128 \times n)^2$. Dividing flips/ns by number of cores, the flips per nanosecond per TPU v3 core is roughly 11.4337, compared to 3.2188 per GPU—a 250% speedup. Note that the energy consumption estimate is an upper bound.

applied on tensors. These tensors can be either stateful variables or temporary values.

It is easy to see that the two-dimensional spin sites can be easily represented as a stateful tensor variable. In this work, we have chosen to use arrays of rank-4 tensors to represent the supergrid structure used in checkerboard algorithm 2. More specifically, the partial spin lattice on each core is represented as a four-dimensional array: `super_grids $[N_x, N_y, 2, 2]$` where each of the array elements is a rank-4 tensor variable with shape of $[m, n, 128 \times i, 128 \times j]$.

The tensor shape is chosen so the last two dimensions are always integer multiples of 128, to better match TPU HBM tiling and the MXU structure[8]. We also choose to use `super_grids[:, :, 0, 0]`, `super_grids[:, :, 1, 1]` to represent the black compact sub-lattices, and use `super_grids[:, :, 0, 1]` and `super_grids[:, :, 1, 0]` to represent the white compact sub-lattices, as depicted in Figure 3-(2). In the results reported in Table 2 we use $(N_x, N_y, m, n, i, j) = (2, 2, 224, 112, 1, 1)$, which gives us the per-core lattice size of $128 \times [896, 448]$.

To determine the acceptance for the flipping of each spin site, a random tensor generation operation available in TensorFlow is used: `tf.random_uniform`⁹. It generates random tensor for a given shape with uniform probability between $[0, 1]$ for each element. While this process is not the most compute-intensive, it does take

up about $\sim 10\%$ of the step time (more discussions on this in next sub-section).

The most compute-intensive part of the simulation is the computation of acceptance ratio, which involves summing on the nearest neighbor spin values. As pointed out previously, we leverage MXU’s matrix multiplication to achieve this. Since each TensorCore provides a raw computational power of ~ 50 TFLOPS, it greatly helps the efficiency of our simulation. This part takes up about $\sim 60\%$ of the step time.

In TensorFlow framework, the expression of the computation is converted into a graph, and for TPU, it is further compiled through XLA into the executable programs (LLO) and deployed to TPU during run time. This Just-In-Time (JIT) compilation can incur overhead but usually it is fairly small (under a few secs) for smaller problems, and while it can sometimes take longer time (up to minutes) for larger problems, usually it is well-amortized for these larger problems as typically millions of steps are executed.

For the distributed multi-core case, TensorFlow also provides the primitive for SIMD programming on TPU. Using `tf.tpu.replicate`¹⁰, one can easily replicate the computation across multiple TPU TensorCores with simple syntax (typically a few lines of code) to map and assign core ids using an object returned during the TPU system

⁹www.tensorflow.org/api_docs/python/tf/random/uniform

¹⁰www.tensorflow.org/api_docs/python/tf/tpu/replicate

initialization call, which encapsulates the mesh topology information.

Figure 5: Illustration of tensor values permute across 3 TensorCores and the procedure of acquiring the neighboring boundary grid values for each TensorCore. Each TensorCore contains a local sub-grid. The right boundary and left boundary are represented by tensors R and L . And the extended boundaries that would need to be filled with the values from their neighbors are represented by tensors R' and L' . (1) Showing the tensor values permutation with $L' = \text{tpu_ops.collective_permute}(R, [[0, 1], [1, 2], [2, 0]])$ and $R' = \text{tpu_ops.collective_permute}(L, [[0, 2], [2, 1], [1, 0]])$. (2) After the permutation, each core gets the boundary values from the neighboring core and the extended boundaries are filled in with the correct values on each core.

Another critical component for the distributed case is that in the acceptance ratio computation, each core needs to exchange its border spin values with the neighboring cores. Again, TensorFlow provides a simple primitive that is syntactically simple and leverages TPU Pod’s high speed dedicated inter-chip mesh network, `tpu_ops.collective_permute`¹¹ which allows a tensor’s value to be permuted through different cores according to the source-destination pair specification. Note that in this case, each core has the same instruction and in the call to `tpu_ops.collective_permute`, the source-destination mapping contains globally identical specifications. Each core that is involved in the operation, when executing this part of program, will block until it sends and receives the corresponding values according to the source-destination specification.

Figure 5 shows an example of a collection of 3 cores that are exchanging the ‘boundaries’ with periodic boundary condition. Each core has tensors R and L representing the right and left boundaries of the internal grid. Figure 5 shows how each core can acquire the extended boundaries from its neighboring cores. The highly efficient inter-core communications in TPU Pod allows all TensorCores to work in lockstep with minimal latency even when a large

number of cores participate in the communications. In fact, the time spent on this step in our experiments is well below 0.15% in all cases and this explains the linear-scaling performance we observed. Note that in all our experiments, no attempt is made to match the logical layout of the lattice with the physical TPU cores: two logically neighboring sub-lattices could be distributed to two cores physically far apart, requiring multiple hops for the communication between them and, yet this doesn’t cause any noticeable performance loss.

5.2 Performance Analysis

Figure 6: A screen grab of the TPU profiling tool’s trace viewer. Shown here is for the case with $16 \times 16 \times 2$ cores and only showing the traces from a few cores. It can also be seen that these cores progress in a lockstep fashion.

We look into several key components of the computations and their respective performance in terms of the time spent on each. We use the tool developed for TPU profiling (available also in Google Cloud¹²). The profiling tool is able to provide fine-grained analysis of the utilization of the hardware, the efficiency of the operations at program level and more. Figure 6 gives an example output from the tool’s trace viewer.

Using the tool, we took measurements of the breakdown of the key operations at the HLO level: the time spent on the computation of the nearest neighbor sum (mostly using MXU), the time spent on generating the random uniform tensors (mostly using VPU), the time spent of data formatting and reshaping for the computation, and the time spent on data exchange between cores. Detailed breakdown is in Table 3. Note that from Table 2 we know that for all cases, the time step is essentially identical ($\sim 580\text{ms}$), so percentage numbers in this table can be used directly for comparison across different cases.

The first observation is that the time breakdown is fairly stable across different scales. We also see that the most expensive part is the computation of nearest neighbor sum, leveraging the MXU, which accounts for $\sim 60\%$ of the time. The time spent on generating random tensors is also significant: it is $\sim 12\%$ of the time. The data formatting takes up $\sim 28\%$ of the time. It is also worth noting that this part can be significantly worse if the shape of the tensor variables do not conform to the tiling in TPU HBM. The most interesting part is the time spent on inter-core communication, and it can be seen that it is taking very insignificant amount of time, even for the very large case when 512 cores are involved. It is a key property of TPU Pod that allows the linear scaling. Note that since the cores participating in the “collective_permute” operations

¹¹ www.github.com/tensorflow/tensorflow/blob/master/tensorflow/python/tpu/ops/tpu_ops.py

¹² cloud.google.com/tpu/docs/cloud-tpu-tools

#cores	lattice size	MXU time (%)	VPU time (%)	data formatting time (%)	collective_permute time (%)
$1 \times 1 \times 2$	$(896 \times 128)^2$	59.6	12	28.2	0.024
$2 \times 2 \times 2$	$(1792 \times 128)^2$	59.6	12	28.1	0.038
$4 \times 4 \times 2$	$(3584 \times 128)^2$	59.5	11.9	28.2	0.063
$8 \times 8 \times 2$	$(7168 \times 128)^2$	59.5	12	28.1	0.08
$16 \times 16 \times 2$	$(14336 \times 128)^2$	59.4	12	28.1	0.11

Table 3: Percentage time breakdown of the computation. Note that since the step time for all cases are basically all ~ 580 ms, the numbers in the table are proportional to absolute time. MXU time is most from the matrix multiplication operations employed for nearest-neighbor computation and it is the biggest portion. VPU time is the time spent in vector unit. In this case, it is mostly for generating random uniform tensors. Data formatting is the time spent on moving data, reshaping tensors, slicing etc. within a core. The inter-core communication time is very negligible for all cases and this is consistent with the strong linear-scaling we observed. In all cases, the amount of the data (the edges) that are moved between cores are $896 \times 128 \times 2 = 229,376$ bytes per edge in one direction and $448 \times 128 \times 2 = 114,688$ bytes per edge in another direction, for each core. These are very small data and the observed time are primarily dominated by other factors such as synchronization and latency.

(step time, collective_permute time) with various per-core lattice size (ms)				
#cores	$[896 \times 128, 448 \times 128]$	$[448 \times 128, 224 \times 128]$	$[224 \times 128, 112 \times 128]$	
$4 \times 4 \times 2$	(575.0, 0.37)	(255, 0.36)	(64.61, 0.18)	
$8 \times 8 \times 2$	(575.2, 0.47)	(255.11, 0.41)	(64.69, 0.25)	
$16 \times 16 \times 2$	(575.3, 0.65)	(255.03, 0.64)	(64.92, 0.58)	

Table 4: The measured step time and collective_permute time (in ms) at various per-core lattice size. The data amount exchanged are small as only edges of the sub-lattice are exchanged between cores so the measured time is not bandwidth limited (the largest edge is only 229,376 bytes, and would take just ~ 0.023 msec to transport on a relatively moderate 10GB network). We also see that for a given total lattice size, as the number of cores increase, the step time decreases.

will perform both sending and receiving and the time measured includes both the synchronization overhead between the cores as well as the time that the data travels between cores. However, since the data amount (the edges of the sub-lattice) is small, the time is not dominated by the data propagation and not bandwidth bound (the largest edge is only 229,376 bytes, and would take just ~ 0.023 msec to transport even on a relatively moderate 10GB network). Table 4 shows the step time and the “collective_permute” time with various per-core lattice size at different number of cores. The “collective_permute” time in all cases are insignificant when compared with the step time. We also note that this time is more directly affected by the number of cores than the size of the sub-lattice, indicating that the communication is not in the bandwidth limited regime.

Another interesting aspect of the data in Table 4 is that, for a constant full lattice size (i.e., the entries along the diagonal direction in Table 4), while the step time consistently decreases as the number of cores increases, we can see that there are two different regimes of the rate of step time decrease. When the per-core lattice size decreases from $[896 \times 128, 448 \times 128]$ to $[448 \times 128, 224 \times 128]$, a $4\times$ decrease, the step time decreases from ~ 575 ms to ~ 255 ms, or $\sim$

#cores	% of roofline optimal	% of HW peak
$1 \times 1 \times 2$	76.68	9.31
$2 \times 2 \times 2$	76.65	9.3
$4 \times 4 \times 2$	76.51	9.28
$8 \times 8 \times 2$	76.52	9.27
$16 \times 16 \times 2$	76.43	9.26

Table 5: Top: Achieved program FLOPS compared against the roofline model optimal performance and hardware peak FLOPS. All measured with per-core lattice size of $128 \times [896, 448]$. All measurements shown here are memory bound. Bottom: the roofline model plot for $16 \times 16 \times 2$ cores with per-core lattice size of $128 \times [896, 448]$.

44%, instead of 25%. But when the per-core lattice size decrease from $[896 \times 128, 448 \times 128]$ to $[224 \times 128, 112 \times 128]$, the step time changes from ~ 255 ms to ~ 65 ms, or $\sim 25.5\%$. This is due to higher MXU utilization when the per-core lattice size is $[896 \times 128, 448 \times 128]$, and the utilization pattern changes (lower MXU utilization) when the per-core lattice size changes to $[448 \times 128, 224 \times 128]$, and stays about the same when the per-core lattice size is $[224 \times 128, 112 \times 128]$.

Finally, we use the profiling tool to measure the FLOPS performance of our program. The results are summarized in Table 5. In

all cases, the achieved FLOPS is about 76% of the roofline (memory bound), or roughly 5.89 TFLOPS. A rough estimate can be done using the number operations in matrix multiplication per core used for nearest neighbor sum: there are total $4 \times 896 \times 448$ of matrix multiplications of size 128 ($2 \times$ from the inner grid nearest neighbor computations, and $2 \times$ from the boundary nearest neighbor). Total number of operations is $896 \times 448 \times 128^3$. Using step time ~ 580 ms, we got the estimate of 5.8 TFLOPS, very close to the measured program FLOPS. It is also interesting to note that, from the slope of the roofline plot in Table 5, we can estimate the HBM bandwidth to be at least ~ 300 GB/sec.

We believe that the efficiency of the computation can be improved by further optimizing the data formatting operations: identifying bottlenecks and rearranging the layout of the tensors. It is also worth noting that the matrix multiplication involves sparse diagonal band kernel with shape of 128×128 and we could potentially explore smaller size of kernel to improve efficiency. A possible direction is trying to utilize the convolution operation in TensorFlow to improve the efficiency further. We also expect that as XLA being actively developed and improved over time, it will deliver higher and higher performance.

6 CONCLUSION

We demonstrate a novel approach to simulate the two-dimensional ferromagnetic Ising model using TensorFlow on Cloud TPU. We adapted the standard checkerboard algorithm to exploit the TPU architecture, in particularly to leverage its efficient matrix operation and the dedicated high-bandwidth low-latency inter-chip interconnect network of the TPU Pod. We calculate the average magnetization and Binder parameter at various temperatures with different lattice sizes. Our numeric estimates on those size-independent quantities match the theoretical results using both float32 and bfloat16 precision. Our benchmarks also demonstrate competitive and linear-scaling performance. The algorithm used in this work can be generalized for three-dimensional Ising model. An interesting direction to follow up would be applying the approach on some of the recent works that push the three-dimensional Ising model simulations to limits, for example[24].

This work demonstrates how the new Cloud TPU computation resources could be efficiently employed for conventional scientific simulation problems. However, even more significantly, by implementing the entire simulation using TensorFlow framework, we point a direction where the direct integration of machine learning approaches with conventional simulations is possible and could be done easily. For example, the automatic differentiation in TensorFlow¹³ is readily applicable for optimization of parameters in design problems that use simulations. In the context of this current work, an interesting followup would be finding the optimal $J_{i,j}$ given material properties for the case where J is not uniform across all spin sites. In our view, the research in this direction will bring many interesting advancements and will continue to shape the future of computing.

¹³www.tensorflow.org/tutorials/eager/automatic_differentiation

7 APPENDICES

7.1 Proof of Stationarity

$$\begin{aligned}
& \sum_{\sigma_w, \sigma_b} P\{(\sigma_w, \sigma_b) \rightarrow (\sigma_w^*, \sigma_b^*)\} \pi(\sigma_w, \sigma_b) \\
&= \sum_{\sigma_w, \sigma_b} P\{(\sigma_w, \sigma_b) \rightarrow (\sigma_w^*, \sigma_b^*)\} P\{(\sigma_w^*, \sigma_b^*) \rightarrow (\sigma_w, \sigma_b)\} \pi(\sigma_w, \sigma_b) \\
&= \sum_{\sigma_w, \sigma_b} \prod_{i \in w} P(\sigma_i \rightarrow \sigma_i^* | \sigma_b) \prod_{i \in b} P(\sigma_i \rightarrow \sigma_i^* | \sigma_w^*) \prod_{i \in w} \pi(\sigma_i | \sigma_b) \cdot \pi(\sigma_b) \\
&= \sum_{\sigma_w, \sigma_b} \prod_{i \in w} P(\sigma_i \rightarrow \sigma_i^* | \sigma_b) \pi(\sigma_i | \sigma_b) \prod_{i \in b} P(\sigma_i \rightarrow \sigma_i^* | \sigma_w^*) \cdot \pi(\sigma_b) \\
&\quad \text{By detailed balance of single spin flip:} \\
&P(\sigma_i \rightarrow \sigma_i^* | \sigma_b) \pi(\sigma_i | \sigma_b) = P(\sigma_i^* \rightarrow \sigma_i | \sigma_b) \pi(\sigma_i^* | \sigma_b) \\
&= \sum_{\sigma_w, \sigma_b} \prod_{i \in w} P(\sigma_i^* \rightarrow \sigma_i | \sigma_b) \pi(\sigma_i^* | \sigma_b) \prod_{i \in b} P(\sigma_i \rightarrow \sigma_i^* | \sigma_w^*) \cdot \pi(\sigma_b) \\
&= \sum_{\sigma_w, \sigma_b} \prod_{i \in w} P(\sigma_i^* \rightarrow \sigma_i | \sigma_b) \prod_{i \in b} P(\sigma_i \rightarrow \sigma_i^* | \sigma_w^*) \prod_{i \in w} \pi(\sigma_i^* | \sigma_b) \cdot \pi(\sigma_b) \\
&= \sum_{\sigma_w, \sigma_b} \prod_{i \in w} P(\sigma_i^* \rightarrow \sigma_i | \sigma_b) \prod_{i \in b} P(\sigma_i \rightarrow \sigma_i^* | \sigma_w^*) \pi(\sigma_i | \sigma_w^*) \cdot \pi(\sigma_w^*) \\
&\quad \text{By detailed balance of single spin flip:} \\
&P(\sigma_i \rightarrow \sigma_i^* | \sigma_w^*) \pi(\sigma_i | \sigma_w^*) = P(\sigma_i^* \rightarrow \sigma_i | \sigma_w^*) \pi(\sigma_i^* | \sigma_w^*) \\
&= \sum_{\sigma_w, \sigma_b} \prod_{i \in w} P(\sigma_i^* \rightarrow \sigma_i | \sigma_b) \prod_{i \in b} P(\sigma_i^* \rightarrow \sigma_i | \sigma_w^*) \pi(\sigma_i^* | \sigma_w^*) \cdot \pi(\sigma_w^*) \\
&= \pi(\sigma_w^*, \sigma_b^*) \sum_{\sigma_w, \sigma_b} \prod_{i \in w} P(\sigma_i^* \rightarrow \sigma_i | \sigma_b) \prod_{i \in b} P(\sigma_i^* \rightarrow \sigma_i | \sigma_w^*) \\
&= \pi(\sigma_w^*, \sigma_b^*)
\end{aligned}$$

ACKNOWLEDGMENTS

We would like to thank Blake Hechtman, Brian Patton, Cliff Young, David Patterson, Naveen Kumar, Norm Jouppi, Rif A. Saurous, Yunxing Dai, Zak Stone, and many more for valuable discussions and helpful comments, which have greatly improved the paper.

REFERENCES

- [1] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, and Michael Isard. Tensorflow: a system for large-scale machine learning. In *OSDI*, volume 16, pages 265–283, 2016.
- [2] Didier Barradas-Bautista, Matias Alvarado-Mentado, Mark Agostino, and Germinal Cocho. Cancer growth and metastasis as a metaphor of go gaming: An ising model approach. *PloS one*, 13(5):e0195654, 2018.
- [3] Kurt Binder. Finite size scaling analysis of ising model block distribution functions. *Zeitschrift für Physik B Condensed Matter*, 43(2):119–140, 1981.
- [4] Kurt Binder, Dieter Heermann, Lyle Roelofs, A. John Mallinckrodt, and Susan McKay. Monte carlo simulation in statistical physics. *Computers in Physics*, 7(2):156–157, 1993.
- [5] Kurt Binder and Erik Luijten. Monte carlo tests of renormalization-group predictions for critical phenomena in ising models. *Physics Reports*, 344(4-6):179–253, 2001.
- [6] Benjamin Block, Peter Virnau, and Tobias Preis. Multi-gpu accelerated multi-spin monte carlo simulations of the 2d ising model. *Computer Physics Communications*, 181(9):1549–1556, 2010.
- [7] Google Cloud. Choosing between a single cloud tpu device and a cloud tpu pod (alpha). <https://cloud.google.com/tpu/docs/deciding-pod-versus-tpu>, 2019.
- [8] Google Cloud. Performance guide. <https://cloud.google.com/tpu/docs/performance-guide>, 2019.
- [9] Google Cloud. System architecture. <https://cloud.google.com/tpu/docs/system-architecture>, 2019.
- [10] Google Cloud. Using bfloat16 with tensorflow models. <https://cloud.google.com/tpu/docs/bfloat16>, 2019.
- [11] Chase Cook, Hengyang Zhao, Takashi Sato, Masayuki Hiromoto, and Sheldon X.-D. Tan. Gpu based parallel ising computing for combinatorial optimization problems in vlsi physical design. *arXiv preprint arXiv:1807.10750*, 2018.

- [12] Nvidia NVLink Fabric. Nvlink fabric: Advancing multi-gpu processing. <https://www.nvidia.com/en-us/data-center/nvlink/>, 2019.
- [13] Ernst Ising. Beitrag zur theorie des ferromagnetismus. *Zeitschrift für Physik*, 31(1):253–258, 1925.
- [14] Norm Jouppi. Quantifying the performance of the tpu, our first machine learning chip. *Google Cloud*: <https://cloud.google.com/blog/products/gcp/quantifying-the-performance-of-the-tpu-our-first-machine-learning-chip>, 2017.
- [15] Norman P. Jouppi, Cliff Young, Nishant Patil, David Patterson, Gaurav Agrawal, Raminder Bajwa, Sarah Bates, Suresh Bhatia, Nan Boden, and Al Borchers. In-datacenter performance analysis of a tensor processing unit. In *Computer Architecture (ISCA), 2017 ACM/IEEE 44th Annual International Symposium on*, pages 1–12. IEEE, 2017.
- [16] Norman P. Jouppi, Doe Hyun Yoon, George Kurian, Sheng Li, Nishant Patil, James Landon, Cliff Young, and David Patterson. A Domain-Specific Supercomputer for Training Deep Neural Networks. *Submitted to Communications of the ACM*.
- [17] Peter Müller. *A generic approach to posterior integration and Gibbs sampling*. Purdue University, Department of Statistics, 1991.
- [18] Lars Onsager. Crystal statistics. i. a two-dimensional model with an order-disorder transition. *Physical Review*, 65(3-4):117, 1944.
- [19] Francisco Ortega-Zamorano, Marcelo A. Montemurro, Sergio Alejandro Cannas, José M. Jerez, and Leonardo Franco. Fpga hardware acceleration of monte carlo simulations for the ising model. *IEEE Trans. Parallel Distrib. Syst.*, 27(9):2618–2627, 2016.
- [20] Adam Paszke, Sam Gross, Soumith Chintala, and Gregory Chanan. Pytorch: Tensors and dynamic neural networks in python with strong gpu acceleration. *PyTorch: Tensors and dynamic neural networks in Python with strong GPU acceleration*, 2017.
- [21] Tobias Preis, Wolfgang Paul, and Johannes J. Schneider. Fluctuation patterns in high-frequency financial asset returns. *EPL (Europhysics Letters)*, 82(6):68005, 2008.
- [22] Tobias Preis, Peter Virnau, Wolfgang Paul, and Johannes J. Schneider. Gpu accelerated monte carlo simulation of the 2d and 3d ising model. *Journal of Computational Physics*, 228(12):4468–4477, 2009.
- [23] Francisco Prieto-Castrillo, Amin Shokri Gzafroudi, Javier Prieto, and Juan Manuel Corchado. An ising spin-based model to explore efficient flexibility in distributed power systems. *Complexity* 2018, 2018.
- [24] Alan M. Ferrenberg, Jiahao Xu, and David P. Landau Pushing the Limits of Monte Carlo Simulations for the 3d Ising Model. *arXiv preprint arXiv:1806.03558*, 2018.

Appendix: Artifact Description/Artifact Evaluation

SUMMARY OF THE EXPERIMENTS REPORTED

We used Google Cloud TPU with TensorFlow r1.13 to perform the experiments reported in the paper. The experiments reported can be repeated using the Jupyter Python Notebook on Google Cloud Platform following the detailed instructions available here: https://github.com/google-research/google-research/tree/master/simulation_research/ising_model

ARTIFACT AVAILABILITY

Software Artifact Availability: All author-created software artifacts are maintained in a public repository under an OSI-approved license.

Hardware Artifact Availability: There are no author-created hardware artifacts.

Data Artifact Availability: There are no author-created data artifacts.

Proprietary Artifacts: None of the associated artifacts, author-created or otherwise, are proprietary.

List of URLs and/or DOIs where artifacts are available:

https://github.com/google-research/google-research/tree/master/simulation_research/ising_model

BASELINE EXPERIMENTAL SETUP, AND MODIFICATIONS MADE FOR THE PAPER

Relevant hardware details: TPUv3, Nvidia V100, 8 vCPUs, 30 GB memory

Operating systems and versions: Ubuntu 16.04.5 LTS (GNU/Linux 4.15.0-1028-gcp x86_64)

Compilers and versions: NVCC 10.1

Applications and versions: NA

Libraries and versions: cuda 10.1, tensorflow 1.13.2

Key algorithms: checkerboard algorithm

Input datasets and versions: NA

Paper Modifications: In order to squeeze every possible bit of power from TPU/TensorFlow, we compiled the tensorflow from the source with optimization enabled and with assert() calls disabled (-O2 -DNDEBUG): "bazel -c opt"

Output from scripts that gathers execution environment information.

Note that the Jupyter notebook can be executed under
→ the Colaboratory environment provided at
→ <http://colab.research.google.com> (see here: <http://s://research.google.com/colaboratory/faq.html> as
→ well).

For using Google Cloud Platform, we also collect the
→ information from the VM used and below is the
→ output (please note that the configuration of
→ this VM is less important than the actual TPU Node
→ used).

```
=====
LS_COLORS=rs=0:di=01;34:ln=01;36:mh=00:pi=40;33:so=0
1;35:do=01;35:bd=40;33;01:cd=40;33;01:or=40;31;0
1:mi=00:su=37;41:sg=30;43:ca=30;41:tw=30;42:ow=3
4;42:st=37;44:ex=01;32:*.tar=01;31:*.tgz=01;31:*.
arc=01;31:*.arj=01;31:*.taz=01;31:*.lha=01;31:*.
lz4=01;31:*.lzh=01;31:*.lzma=01;31:*.tlz=01;31:
*.txz=01;31:*.tzo=01;31:*.t7z=01;31:*.zip=01;31:
*.z=01;31:*.Z=01;31:*.dz=01;31:*.gz=01;31:*.lrz=
01;31:*.lz=01;31:*.lzo=01;31:*.xz=01;31:*.zst=01
;31:*.tzst=01;31:*.bz2=01;31:*.bz=01;31:*.tbz=01
;31:*.tbz2=01;31:*.tz=01;31:*.deb=01;31:*.rpm=01
;31:*.jar=01;31:*.war=01;31:*.ear=01;31:*.sar=01
;31:*.rar=01;31:*.alz=01;31:*.ace=01;31:*.zoo=01
;31:*.cpio=01;31:*.7z=01;31:*.rz=01;31:*.cab=01;
31:*.jpg=01;35:*.jpeg=01;35:*.mjpg=01;35:*.mjpeg
=01;35:*.gif=01;35:*.bmp=01;35:*.pbm=01;35:*.pgm
=01;35:*.ppm=01;35:*.tga=01;35:*.xbm=01;35:*.xpm
=01;35:*.tif=01;35:*.tiff=01;35:*.png=01;35:*.sv
g=01;35:*.svgz=01;35:*.mng=01;35:*.pcx=01;35:*.m
ov=01;35:*.mpg=01;35:*.mpeg=01;35:*.m2v=01;35:*.
mkv=01;35:*.webm=01;35:*.ogm=01;35:*.mp4=01;35:*.
m4v=01;35:*.mp4v=01;35:*.vob=01;35:*.qt=01;35:*.
nuv=01;35:*.wmv=01;35:*.asf=01;35:*.rm=01;35:*.
rmvb=01;35:*.flc=01;35:*.avi=01;35:*.fli=01;35:*.
flv=01;35:*.gl=01;35:*.dl=01;35:*.xcf=01;35:*.x
wd=01;35:*.yuv=01;35:*.cgm=01;35:*.emf=01;35:*.o
gv=01;35:*.ogx=01;35:*.aac=00;36:*.au=00;36:*.fl
ac=00;36:*.m4a=00;36:*.mid=00;36:*.midi=00;36:*.
mka=00;36:*.mp3=00;36:*.mpc=00;36:*.ogg=00;36:*.
ra=00;36:*.wav=00;36:*.oga=00;36:*.opus=00;36:*.
spx=00;36:*.xspf=00;36:
LANG=en_US.UTF-8
SUDO_GID=1001
USERNAME=USER
SUDO_COMMAND=/bin/bash collect.sh
USER=USER
PWD=/home/yifanchen
HOME=/USER
SUDO_USER=yifanchen
SUDO_UID=1000
MAIL=/var/mail/USER
SHELL=/bin/bash
TERM=xterm-256color
SHLVL=1
LOGNAME=USER
```

```

PATH=/usr/local/sbin:/usr/local/bin:/usr/sbin:/usr/b_
↳ in:/sbin:/bin
_=/usr/bin/env
Distributor ID:      Debian
Description:         Debian GNU/Linux 9.9 (stretch)
Release:             9.9
Codename:            stretch
Linux ising-large 4.9.0-9-amd64 #1 SMP Debian
↳ 4.9.168-1+deb9u5 (2019-08-11) x86_64 GNU/Linux
Architecture:       x86_64
CPU op-mode(s):     32-bit, 64-bit
Byte Order:         Little Endian
CPU(s):             8
On-line CPU(s) list: 0-7
Thread(s) per core: 2
Core(s) per socket: 4
Socket(s):          1
NUMA node(s):       1
Vendor ID:          GenuineIntel
CPU family:         6
Model:              85
Model name:         Intel(R) Xeon(R) CPU @ 2.00GHz
Stepping:           3
CPU MHz:            2000.148
BogoMIPS:           4000.29
Hypervisor vendor:  KVM
Virtualization type: full
L1d cache:          32K
L1i cache:          32K
L2 cache:           256K
L3 cache:           56320K
NUMA node0 CPU(s): 0-7
Flags:              fpu vme de pse tsc msr pae mce
↳ cx8 apic sep mtrr pge mca cmov pat pse36 clflush
↳ mmx fxsr sse sse2 ss ht syscall nx pdpe1gb rdtscp
↳ lm constant_tsc rep_good nopl xtopology
↳ nonstop_tsc pni pclmulqdq ssse3 fma cx16 pcid
↳ sse4_1 sse4_2 x2apic movbe popcnt aes xsave avx
↳ f16c rdrand hypervisor lahf_lm abm 3dnowprefetch
↳ invpcid_single ssbd ibrs ibpb stibp kaiser
↳ fsgsbase tsc_adjust bmi1 hle avx2 smep bmi2 erms
↳ invpcid rtm mpx avx512f avx512dq rdseed adx smap
↳ clflushopt clwb avx512cd avx512bw avx512vl
↳ xsaveopt xsavec xgetbv1 xsaves arat md_clear
↳ arch_capabilities
MemTotal:           30887032 kB
MemFree:            29741164 kB
MemAvailable:       29990136 kB
Buffers:            61956 kB
Cached:             522988 kB
SwapCached:         0 kB
Active:             653668 kB
Inactive:           276844 kB
Active(anon):       345752 kB
Inactive(anon):     10924 kB
Active(file):       307916 kB

```

```

Inactive(file):     265920 kB
Unevictable:        0 kB
Mlocked:            0 kB
SwapTotal:          0 kB
SwapFree:           0 kB
Dirty:              8 kB
Writeback:          0 kB
AnonPages:          345620 kB
Mapped:             188724 kB
Shmem:              11112 kB
Slab:               81624 kB
SReclaimable:       54620 kB
SUnreclaim:         27004 kB
KernelStack:        3328 kB
PageTables:          5592 kB
NFS_Unstable:       0 kB
Bounce:             0 kB
WritebackTmp:       0 kB
CommitLimit:        15443516 kB
Committed_AS:       1442836 kB
VmallocTotal:       34359738367 kB
VmallocUsed:         0 kB
VmallocChunk:        0 kB
HardwareCorrupted:  0 kB
AnonHugePages:      0 kB
ShmemHugePages:     0 kB
ShmemPmdMapped:     0 kB
HugePages_Total:    0
HugePages_Free:     0
HugePages_Rsvd:     0
HugePages_Surp:     0
Hugepagesize:       2048 kB
DirectMap4k:        120820 kB
DirectMap2M:        5122048 kB
DirectMap1G:        28311552 kB
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
sda 8:0 0 120G 0 disk
└─sda1 8:1 0 120G 0 part /
[0:0:1:0] disk Google PersistentDisk 1
↳ /dev/sda 128GB

```

H/W path	Device	Class	Description
=====			
		system	Google Compute Engine
/0		bus	Google Compute Engine
/0/0		memory	96KiB BIOS
/0/1001		processor	Intel(R) Xeon(R)
↳ CPU @ 2.00GHz			
/0/1002		processor	CPU
/0/1003		processor	CPU
/0/1004		processor	CPU
/0/1005		processor	CPU
/0/1006		processor	CPU
/0/1007		processor	CPU
/0/1008		processor	CPU
/0/200		memory	30GiB System Memory
/0/200/0		memory	16GiB DIMM RAM

High Performance Monte Carlo Simulation of Ising Model on TPU Clusters

/0/200/1		memory	14GiB DIMM RAM
/0/100		bridge	440FX - 82441FX PMC
↪ [Natoma]			
/0/100/1		bridge	82371AB/EB/MB
↪ PIIX4 ISA			
/0/100/1.3		bridge	82371AB/EB/MB
↪ PIIX4 ACPI			
/0/100/3		generic	Virtio SCSI
/0/100/3/0		generic	Virtual I/O device
/0/100/4		network	Virtio network
↪ device			
/0/100/4/0	eth0	network	Ethernet interface
/0/1	scsi0	storage	
/0/1/0.1.0	/dev/sda	disk	128GB
↪ PersistentDisk			
/0/1/0.1.0/1	/dev/sda1	volume	119GiB EXT4 volume
/1	docker0	network	Ethernet interface