

Jet: Fast quantum circuit simulations with parallel task-based tensor-network contraction

Trevor Vincent¹, Lee J. O’Riordan¹, Mikhail Andrenkov¹, Jack Brown¹, Nathan Killoran¹, Haoyu Qi¹, and Ish Dhand^{1,2}

¹Xanadu, 777 Bay Street, Toronto, Canada

²Institute of Theoretical Physics and IQST, Ulm University, Albert-Einstein-Allee 11, 89081 Ulm, Germany

We introduce a new open-source software library *Jet*, which uses task-based parallelism to obtain speed-ups in classical tensor-network simulations of quantum circuits. These speed-ups result from i) the increased parallelism introduced by mapping the tensor-network simulation to a task-based framework, ii) a novel method of reusing shared work between tensor-network contraction tasks, and iii) the concurrent contraction of tensor networks on all available hardware. We demonstrate the advantages of our method by benchmarking our code on several Sycamore-53 and Gaussian boson sampling (GBS) supremacy circuits against other simulators. We also provide and compare theoretical performance estimates for tensor-network simulations of Sycamore-53 and GBS supremacy circuits for the first time.

1 Introduction

Classical simulation of quantum systems are the workhorse of research in quantum many-body physics and quantum information processing. In particular, a large body of research exists for the classical simulations of random quantum circuits (RQCs) as they are good candidates for demonstrating quantum advantage [1–5]. In a landmark paper, Arute et al. [6] demonstrated a task that was beyond the capability of current classical computers using their Sycamore-53 device on a 53-qubit RQC with a depth (m) of 20 cycles. However, it was argued shortly thereafter that the classical simulation benchmarks used in [6] were overestimates and could be greatly improved upon by using secondary storage of the

Summit supercomputer [7]. Deshpande et al. [8] have proposed a quantum advantage experiment using Gaussian boson sampling (GBS) and a three-dimensional RQC. It appears these circuits could be much harder to simulate classically than Sycamore-53, with some simulations showing it would take $\approx 10^{14}$ seconds on the top supercomputer in the most idealized scenario where unlimited memory was available. While there are many different classical methods for simulating RQCs, currently the fastest known methods involve tensor networks, with Huang et al. [5] having the shortest estimates for simulating the Sycamore-53 device. Despite the level of complexity and innovation of the algorithms used in their classical simulations, the estimate is still orders of magnitude slower than the Sycamore-53 device. A major issue in the simulations of [5] was efficiency, as the simulator utilized only around 15 percent of the theoretical compute performance of the NVIDIA V100 GPU they performed the tests on. A second issue was the lack of CPU usage in their simulations, which would result in under-utilization of the Summit supercomputer, which has 9,216 IBM POWER9 22-core CPUs amongst the 27,648 NVIDIA Tesla V100 GPUs. These issues may be accentuated as supercomputers reach exascale performance in the next few years.

Current inefficiencies in tensor-network simulation on supercomputers are likely to worsen with the oncoming era of exascale computing and the challenges it will bring. Future exascale supercomputers are expected to increase in parallelism and heterogeneity, with billions of threads running concurrently [9, 10]. Specifically, the systems will not increase in the number of nodes, but rather in on-node concurrency with large multi-core CPUs and multiple GPUs per node [11–13]. In addition, memory hierarchies are expected to

become deeper; however, the memory size and memory bandwidth per core is expected to decrease. Due to the large number of components in an exascale machine, hardware failures are expected to increase and fault-tolerance will be a major issue. The programming models used in tensor-network simulators will likely have problems addressing these changes.

To address issues with current programming models, the high-performance computing community has started considering asynchronous task-parallelism (sometimes referred to as asynchronous many-task run-time models) as a solution [10, 14]. Asynchronous task-parallelism is founded on the idea of decomposing an algorithm into units of work, known as tasks, and executing them asynchronously. Asynchronicity allows the hiding of communication latencies, whilst dependencies between tasks enforce program correctness and allow for finer-grained synchronisation. Furthermore, to solve the problem of heterogeneity, tasking would be a sensible approach as it allows for implicitly handling load balancing by executing work when needed, as opposed to the current commonly-used methods of problem sub-division. In addition to solving problems that will exist in upcoming exascale supercomputers, task-based parallelism will allow for more irregular and complex scientific programs to be solved, which have previously been untenable due to a lack of support for a suitable programming model. Finally, asynchronous task-parallelism can also effectively handle fault-tolerance issues without global synchronization [15]. Several task-parallelism libraries are currently under active development, such as Taskflow [16], Charm++ [17], HPX [18] and Kokkos [19].

To date, there has been no prior work detailing how to execute tensor-network quantum circuit simulations using a task-based parallelism framework. In this work, we introduce the open-source software Jet, which builds the task-dependency graph for tensor-network-based quantum-circuit simulations to be executed efficiently on heterogeneous nodes using the Taskflow library [16]. The introduction of task-parallelism with Jet provides us with three key advantages over existing tensor-network simulators. Firstly, decomposing the tensor-network contraction problem into tasks allows for a greater degree of parallelism

because multiple tensor transposes and/or contractions can be executed concurrently. Secondly, by writing out the task graph we can take advantage of shared computation between multiple tensor-network contractions without any extra effort. Lastly, we can use Taskflow to target heterogeneous CPU+GPU nodes.

This paper is structured as follows. In Sec. 2 we overview the background needed to understand tensor-network simulators and task parallelism. In Sec. 3 we describe how to contract tensor networks using task-based parallelism and shared-work reuse on heterogeneous computing nodes. We follow this in Sec. 4 with numerical results showcasing our methods on Sycamore-53 and three-dimensional GBS circuits, including a first-ever comparison of theoretical estimates for Sycamore-53 and GBS quantum advantage circuits. We then conclude the paper in Sec. 5 with viewpoints on potential future improvements and extensions.

2 Background

2.1 Tensor Networks

A tensor network is a countable set of tensors bonded together by contracted indices [20]. We call the reduction of a tensor network through the summation of all contracted indices, a contraction of the tensor network. Contracting tensor networks with arbitrary structure remains at least $\#P$ -hard [21]. The time complexity of the contraction is heavily sensitive to the order of summations [22], with the determination of the optimal path being an NP-hard problem [23]. Despite these difficulties, several approximate and exact methods have been developed to quickly determine quasi-optimal contractions paths and contract the tensor network, with the most common methods being tree decomposition [24–26] and graph partitioning [3, 5, 27].

2.2 Quantum Circuits as Tensor Networks

A quantum circuit is a series of logical gates acting on qudits, a d -level generalization of qubit. In general a quantum circuit can be represented as a unitary operator \mathcal{U} acting on an input set of qudits each in initial state $|0\rangle$. A graphical representation of a general two-qudit circuit is shown in Fig. 1.

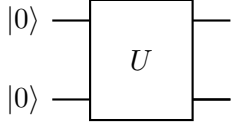
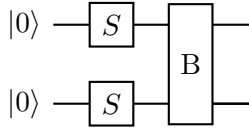


Figure 1: A general 2-qudit quantum circuit.

To demonstrate the mapping between a quantum circuit and a tensor network it is easiest to consider a concrete example of a quantum circuit, which we provide in the following example.

2.2.1 Example Circuit

Consider the simple two-qudit circuit:



Here we apply two one-qudit gates we call S followed by a two-qudit gate we call B . The one-qudit S gate can be represented by a tensor of rank 2, S_{ba} , whereas B is a two-qudit gate represented by a tensor of rank 4, B_{cfbe} . Without loss of generality, we assume the initial state is the product of rank-one tensors: $|0\rangle_a |0\rangle_d$. All tensors in this example have indices with dimension D , the size of the qudit. We form a tensor network representation of this circuit by connecting these tensors, which represent the nodes of the network, together by the indices they share, which represent the edges of the network. For example, S_{ba} is connected to $|0\rangle_a$ and B_{cfbe} along the shared indices b and a .

For classical simulations of large random quantum circuits, it is typically impossible to store the final state in memory. Therefore we can only compute amplitudes or other low-memory quantities. If we want the amplitude of the quantum state $|a_1 a_2\rangle$, we attach the rank-1 $|a_1\rangle_c$ and $|a_2\rangle_f$ tensors to the B_{cfbe} tensor in the network.

The amplitude $\langle a_1 a_2 | \mathcal{U} | 00 \rangle$, which we will shorten to $\langle a | \mathcal{U} | 0 \rangle$ hereafter, can then be computed through the contraction of a tensor network. One way of computing such a contraction is by summing over all indices in the network:

$$\langle a | \mathcal{U} | 0 \rangle = \sum_{abcdef}^D \langle a_1 |_c \langle a_2 |_f B_{cfbe} S_{ba} S_{ed} |0\rangle_a |0\rangle_d, \quad (1)$$

which has a time complexity of $\mathcal{O}(D^6)$ and is not the most optimal way to contract the network, as we'll see in the next section.

2.3 Pairwise contraction

Computing the amplitude by calculating the sum in Eq. 1 has a time complexity of $\mathcal{O}(D^6)$ and in general has exponential cost in the number of shared indices. It has been shown that pairwise contraction of the tensors in a tensor network is orders of magnitude faster than the naive summation in Eq. 1 at computing a tensor network contraction [3, 5, 28]. For the tensor network defined by Eq. 1, one could follow the sequence of pairwise contractions to obtain the same resulting amplitude:

$$\begin{aligned} 1) \quad & \sum_c \langle a_1 |_c B_{cfbe} = T_{fbe}^1 \quad \mathcal{O}(D^4) \\ 2) \quad & \sum_d S_{ed} |0\rangle_d = T_e^2 \quad \mathcal{O}(D^2) \\ 3) \quad & \sum_a S_{ba} |0\rangle_a = T_b^3 \quad \mathcal{O}(D^2) \\ 4) \quad & \sum_f \langle a_2 |_f T_{fbe}^1 = T_{be}^4 \quad \mathcal{O}(D^3) \\ 5) \quad & \sum_e T_{be}^4 T_e^2 = T_b^5 \quad \mathcal{O}(D^2) \\ 6) \quad & \sum_b T_b^3 T_b^5 = \langle a | \mathcal{U} | 0 \rangle \quad \mathcal{O}(D). \end{aligned} \quad (2)$$

The pairwise contraction in Eq. 2 can naturally be described using a binary tree, which we show in Fig. 2. Changing the contraction path will also change the binary tree associated with it.

2.4 Slicing

As the circuits become larger in depth and in the number of qudits, the intermediary tensors in the contraction tree grow in memory. To avoid the need to hold the tensors in distributed memory, which would require inter-node communication to compute contractions, tensor-network simulators address the problem by using a method known as slicing or cutting [3, 28, 29].

Slicing is the simple method of subdividing a tensor network through fixing the value of a shared index. For example, if we consider our qudit circuit example, with the contraction defined by Eq. 1 and $D = 2$, then each tensor index can have a value of 0 or 1. If we slice index e in the

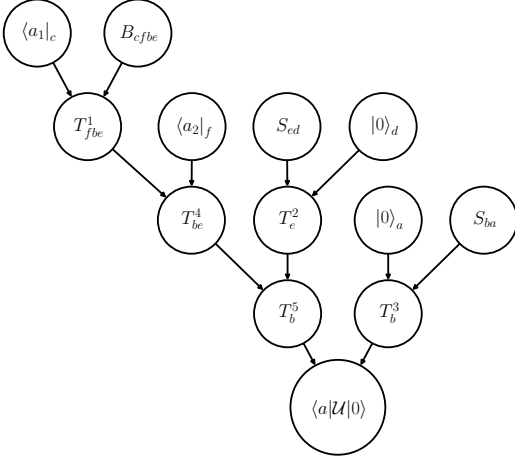


Figure 2: The binary contraction tree for the pairwise contraction in Eq. 2.

tensor network, then Eq. 1 reduces to two partial sums

$$1) \sum_{acdfb} \langle a_1 |_c \langle a_2 |_f B_{cfb0} S_{ba} S_{0d} |0\rangle_a |0\rangle_d = s_0 \quad (3)$$

$$2) \sum_{acdfb} \langle a_1 |_c \langle a_2 |_f B_{cfb1} S_{ba} S_{1d} |0\rangle_a |0\rangle_d = s_1 \quad (4)$$

$$3) \sum_i s_i = \langle a | \mathcal{U} | 0 \rangle. \quad (5)$$

Slicing the tensor network in this manner allows for embarrassingly (data-) parallel computation, as we can distribute the two partial sums to two processing units and collect the results through a single reduction operation after the two partial sums are computed. Equivalently, in terms of pairwise contraction, slicing the index e results in two different binary contraction trees which can be distributed to two different processing units. As an example of a sliced contraction tree, we show one of the two possible slices through the e index with a value of $e = 0$ in Fig. 3. The total amount of floating point operations (FLOP) for a sliced computation of an amplitude is

$$\text{FLOP}_{\text{amp}} = N_{\text{sl}} \text{FLOP}_{\text{sl}}, \quad (6)$$

where N_{sl} is the number of slices and FLOP_{sl} is the FLOP for the contraction of a single slice. The goal of previous tensor-network contraction algorithms was to find a set of slices and a contraction path that minimizes Eq. 6.

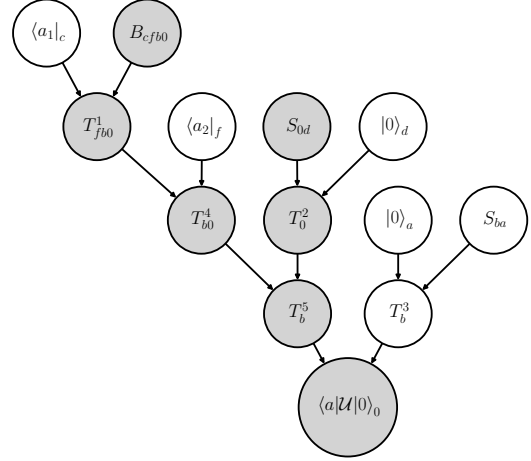


Figure 3: The binary contraction tree for a slice of Fig. 2 through the index e with value 0. We fill all nodes grey that are sliced or the result of a contraction involving a sliced tensor.

Importantly, for a fixed contraction order, slicing will typically increase the total FLOP as compared to the non-sliced contraction order, known as the slicing overhead [3]. However, by optimizing the contraction order after slicing the network, one can often minimize or eliminate the increase in total FLOP [5].

2.5 Task-based parallelism and Taskflow

As high-performance computing (HPC) moves toward increasingly diverse architectures and larger processor counts in the exascale era, new programming models are being developed to mitigate the predicted challenges [10]. The current trend of the high-performance community is a move towards asynchronous task-parallelism [9–11, 30, 31], commonly referred to as asynchronous many-task parallelism. Asynchronous task-parallelism is based on the notion of a task, which represents an asynchronous operation. The size or granularity of a task can range from a single to many hardware instructions. Tasks are strung together based on dependencies to form a task graph. At run-time a scheduler continuously feeds the processing units with tasks from the task graph until the computation has completed. Task-parallelism has demonstrated great performance for a number of complex scientific applications; see [32–37], for a few selected examples.

While there are many libraries, APIs and language extensions that implement some form of asynchronous task-parallelism, we choose the C++ library Taskflow due to its ease of use, modern design, and improved performance over other competing libraries [16]. Unlike many other libraries, Taskflow supports general control flow, not just a simple directed acyclic task graph, as well as a work-stealing task scheduler that can efficiently utilize heterogeneous architectures such as a multi-core CPU connected to multiple GPUs, which is common amongst supercomputer nodes. On top of this, Taskflow can handle multiple GPU vendors through the SYCL portability layer [38]. While Taskflow currently only supports on-node parallelism, we can naturally use a MPI+Taskflow model on distributed memory clusters because our problem is embarrassingly parallel due to the technique called slicing we described earlier.

2.6 Overview of previous tensor-network simulators for quantum circuit simulation

Many high-performance tensor-network simulators have been developed for various types of quantum simulation and it would be impractical to overview them all here. Instead, we limit our focus to those simulators benchmarking amplitude or sampling calculations of massive quantum circuits, such as Sycamore-53 or GBS circuits. We also include simulators specifically developed for exascale computing. To date, there have been three major publications for tensor-network simulators benchmarking the Sycamore-53 circuit.

The first was the simulations published by Google using the tensor-network simulator qFlex for CPU simulations and TAL-SH for GPU simulations [6, 28]. These simulations used a contraction order and set of slices predetermined by human calculation. No algorithms were used to optimize these slices or paths. The qFlex simulator had one of the fastest available CPU tensor contraction engines (specialized for tensor networks) available at the time and likewise TAL-SH was one of the fastest available for GPU tensor contraction. In Arute et al. [6], they were not able to simulate the Sycamore-53 circuit to $m = 20$ cycles with qFlex/TAL-SH or even provide estimates because the intermediary tensors along the contraction path could not fit on a node. Instead

they computed 1 million samples to 0.5% fidelity in 4.644×10^3 seconds for the $m = 12$ circuit and estimated they could do 1 million samples to 0.5% fidelity in 5.875×10^6 seconds for the $m = 14$ circuit. This simulation was run on 4550 nodes out of 4608 of the Summit supercomputer (at the time the largest supercomputer) with 6 GPUs per node.

The second publication came the following year with Gray and Kourtis [3] showcasing the strength of hypergraph partitioning methods in finding optimal contraction paths. Together with a greedy slicing algorithm, Gray et al. dramatically reduced the runtimes for various Sycamore-53 circuits using their software library called CoTenGra [39]. On a single NVIDIA Quadro P2000 they were able to contract an amplitude for the Sycamore-53 $m = 12$ circuit in 5.74×10^2 seconds, the $m = 14$ circuit in an estimated 2.92×10^3 seconds and $m = 20$ in an estimated 7.17×10^9 seconds. This was the first estimate for the time complexity of the $m = 20$ circuit.

The third publication came in late 2020 with Huang et al. [5] fine-tuning the method of Gray et al. using new techniques, most notably, dynamic slicing and local optimization. Huang et al. started with a contraction order found using hypergraph partitioning and then alternated between finding a slice (using a greedy algorithm) and local optimization of the sliced contraction tree by doing exact solves of several sub-trees. Using this method in their AC-QDP simulator they ran one sample of the $m = 12$ and $m = 14$ circuits on a NVIDIA Tesla V100 SMX2 GPU with 16 GB of RAM. They used the results from these smaller simulations to estimate runtimes for a million samples with the appropriate fidelity on the Summit supercomputer. For $m = 12, 14$ they estimated it would take 18s and 88s respectively. For $m = 20$ they estimated they could compute the samples in 1.67×10^6 seconds. The alterations to the Gray and Kourtis methodology introduced by Huang et al. were later incorporated into the software library of Gray and Kourtis, CoTenGra [39]. In collaboration with NVIDIA, CoTenGra contracted a Sycamore-53 $m = 20$ sample in 558 seconds on the Selene GPU cluster [40]

In [8] a three-dimensional random GBS quantum circuit containing 216 modes (6 modes per dimension of the circuit architecture) and a sin-

gle cycle was proposed to showcase quantum advantage. In order to map the three-dimensional random GBS circuit to a tensor-network simulation, a truncated Fock space corresponding to a qudit size of 4 was assumed. Using the theoretical compute performance of the Fugaku supercomputer, and the CoTenGra contraction path finder [3, 39], Deshpande et al. [8] found that the 6x6x6 circuit could be computed in $\approx 10^{14}$ seconds assuming access to memory well beyond that of all the nodes of Fugaku together, let alone a single node of Fugaku. Slicing the circuit down to sizes that could fit in the 32GB RAM of a Fugaku node would come with a slicing overhead that was astronomical, making this circuit infeasible on current or future supercomputers. Ref. [8] also studied other three-dimensional GBS circuits and found that the circuit with 5 modes per dimension was also likely intractable, however an amplitude of the circuit with 4 modes per dimension could likely be computed in under an hour on Fugaku. Based on this evidence, the three-dimensional GBS circuit with 6 modes is well beyond the capabilities of current simulation methods.

Another notable tensor-network simulator is that of Lykov et al. [26], who computed 210-qubit Quantum Approximate Optimization Ansatz (QAOA) circuits with 1785 gates on 1024 nodes of the Cray XC 40 supercomputer Theta. Lykov et al. used a greedy path optimizer, which is known to perform slightly worse than the hypergraph partitioners that AQCDP and CoTenGra use [3, 5]. However, they coupled this greedy path optimizer with a novel step-dependent slicing method that determines indices to slice at run-time as the path is contracted. This type of method would be very suitable for dynamic task graphs and could be incorporated into the setup of Jet very naturally. Unfortunately, [26] did not benchmark on Sycamore-53 so we do not know the extent to which this novel slicing method would perform against interleaving slicing with subtree-reconfiguration as was first done in ACQDP [5] and now supported in the current version of CoTenGra [39]. Lastly, the tensor-network simulator ExaTN [41] has been used to perform large quantum circuit simulations on Summit. Nguyen et al. [41] benchmarked ExaTN on Sycamore-53 ($m = 14$), however it appears they did not perform pre-contraction simplifica-

tions or high-quality path optimization so their run-times are not competitive, despite the compute efficiency of ExaTN being quite high.

3 Overview of Jet

A major shortcoming of the simulators discussed in the previous section is that they do not take advantage of the duplicate work between slice calculations and amplitude calculations. Furthermore, they do not take advantage of the fact that some contractions in the contraction tree can be performed in parallel as they are independent. Lastly, there is no support for heterogeneous nodes and concurrent contraction on both the CPU and GPU, with the sole exception of TAL-SH and its successor ExaTN. In the following sections we introduce our simulator Jet and show how using a combination of task-based parallelism and a novel shared-work reuse we can mitigate the above difficulties.

At a high level, Jet takes in a special file describing the tensor network, which stores the contraction path and the raw tensor network data, builds a task dependency graph from this data which is then mapped to the CPU and GPU or other devices using Taskflow’s task scheduler. Our code does not perform the search for optimal slices or paths, to allow ourselves flexibility in which methods we use to compute these. Currently, to compute high-quality paths and slices for tensor networks we use CoTenGra [39], which uses hypergraph partitioning methods along with greedy slicing and subtree reconfiguration. We will now overview how we build the task-dependency graph from tensor-network data step by step.

3.1 Building the task graph for a contraction

The basic building block of the task-dependency graph is the pairwise contraction of two tensors. This task can be decomposed into two independent (partial) tensor transposes and a single matrix multiply [42]. While tensor contractions can be computed with other approaches, the transpose-transpose-matrix-multiply method is useful in a task-based setting as it allows for more parallelism due to the fact that transposes from different contractions in the contraction tree can be executed in parallel. For contractions on

the CPU we use a slightly more optimized version of the transpose used by qFlex [43], which is the fastest transpose on the CPU for quantum simulation we are aware of [28], and we use the Intel Math Kernel Library (MKL) or OpenBLAS for the matrix multiplication. On the GPU, we use v1.3.1 of cuTENSOR [44]. cuTENSOR allows us to leverage the tensor cores of the GPU when available.

To create the task graph we loop through the contraction path in the tensor network file, adding the tensor transposes and matrix multiplies for each contraction in the path. Each transpose is only dependent on the tensor it is transposing being available and therefore has one dependency. The matrix multiply requires the two transposes to finish and therefore has two dependencies. The task graph for the calculation of a single amplitude corresponding to the contraction tree in Fig. 2 is shown in Fig. 4. The nodes with superscript T correspond to transposes and the nodes without superscript T correspond to pure matrix multiplies whose result gives the tensor named in the node. In this simple contrived example many of the tensors do not need to be transposed before the multiplication because their indices are already in the correct order, so there are far fewer transposes than would occur in a more realistic example. Regardless, this example clearly shows that nodes highlighted in grey correspond to tasks which can be computed in parallel. In previous simulators, which did not use task-based parallelism, taking advantage of this parallelism would not be possible.

3.2 Building the task graph for multi-contractions and reusing shared work

In quantum circuit simulation almost all computations typically require multi-contractions, that is, multiple tensor networks will need to be contracted in order to get some final quantity of interest. Often these contractions will share work. The simplest example of this is for computing slices. Slicing breaks a single contraction problem into a multi-contraction problem. For large circuits, there will always be shared work between these sliced contractions. To build the task graph for this problem we loop through the contraction path for each of the slices, ignoring redundant calculations by not adding tasks with the same name. We name the tasks by concatenating

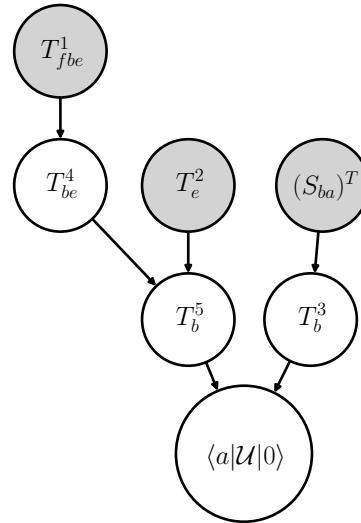


Figure 4: The task dependency graph for a single tensor-network contraction corresponding to the contraction tree in Fig. 2. All nodes of this task graph are named by the computational result of the task. All nodes correspond to the result of matrix-multiply calls except those with superscript T , which correspond to tensor transposes. In this simple contrived example, we need to transpose the tensor S_{ba} because it is on the right side of the contraction between $|0\rangle_a$ and S_{ba} . For this simple example, it would be possible to avoid the transpose by swapping the positions of $|0\rangle_a$ and S_{ba} in the contraction tree, but in general such an operation is not possible without affecting the tensors in the rest of the tree.

their contraction path order number, with their indices and the slice value associated to that index. This naming scheme avoids collisions when adding multiple slices to the same task graph. We showcase the task graph in Fig. 5. Two things are immediately of note. Firstly, the task graph now contains more filled nodes, meaning there are more independent computations that can be contracted concurrently. Secondly, the nodes filled with a light blue color are shared between the two slice contractions. This means that the blue work only needs to be computed once and then can be reused for the other slice. If we were to compute more slices, we could keep re-using this work.

Previous simulators did not take advantage of this duplicate work, but with task-based parallelism it is not only possible, it is straightforward. We can actually take this a step further and start maximizing the amount of shared work between multi-contractions. Thus, the goal of path and slice finding within a task-based parallelism framework is then to minimize the amount

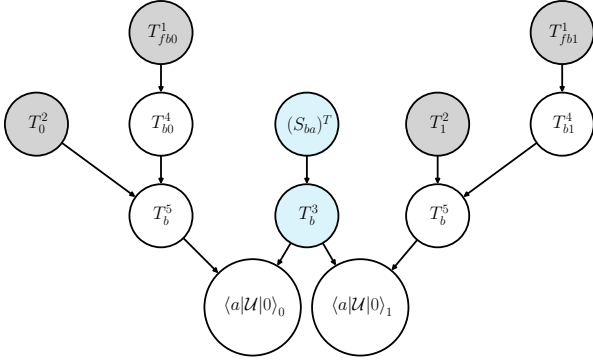


Figure 5: The task dependency graph for a multi-contraction of two slices corresponding to the fixed values of the index e : $e = 0$ and $e = 1$. All the filled nodes that are independent can be computed concurrently. Furthermore, the nodes colored light blue represent shared work between the two contractions. In previous simulators the light blue work would be computed twice, once for each slice. However, by using tasking we only need to compute it once.

of FLOP per contraction while maximizing the amount of duplicate tasks (shared work) between contractions.

The equation that should be minimized is not Eq. 6 but the following:

$$\text{FLOP}_{\text{amp}} = f_{\text{sl}}\text{FLOP}_{\text{sl}} + N_{\text{sl}}(1 - f_{\text{sl}})\text{FLOP}_{\text{sl}}, \quad (7)$$

where f_{sl} is the fraction of FLOP_{sl} that contains duplicate tasks between slices. This generalizes to other forms of multi-contraction, such as computing batches of amplitudes and various methods of sampling quantum circuits which always require multi-contractions. The shared work can also be stored onto disk for future computation.

4 Numerical Results

In this section we benchmark Jet’s task-based approach, and offer a comparison with existing simulators. All benchmarks can be reproduced with scripts provided in the open-source Jet repository [45]. While there are many interesting circuits and computations that require tensor-network contractions, we restrict ourselves to contractions of random amplitudes or slices of Sycamore-53 and GBS circuits, which have both been used or proposed for quantum advantage tasks. For the Sycamore circuits, we use the circuit files provided in the attached data to Ref. [6]; these circuits are also provided in the CoTenGra reposi-

tory [39]. To generate the high-dimensional random GBS circuits we use the algorithm provided in Alg. 1.

Algorithm 1 Random n -dimensional GBS circuit.

The set of parameters dim , width , modes , cycles and r are user-defined. We use $q(i)$ to represent the i -th qumode. S and BS are the squeezer and beam splitter gates respectively. The vertical bar $|$ represents the operation of a gate on a qudit or set of qudits.

```

1: procedure RQC
2:   for  $i \leftarrow 1$  to  $\text{dim}$  do
3:      $\text{lens}[i] = (\text{width})^{i-1}$ 
4:   end for
5:   for  $k \leftarrow 1$  to  $\text{modes}$  do
6:      $S(r) | q(k)$ 
7:   end for
8:   for  $c \leftarrow 1$  to  $\text{cycles}$  do
9:     for  $l \leftarrow 1$  to  $\text{dim}$  do
10:      for  $i \leftarrow 1$  to  $\text{modes} - \text{lens}[l]$  do
11:         $\theta = \text{random}(0, 2\pi)$ 
12:         $\phi = \text{random}(0, 2\pi)$ 
13:         $BS(\theta, \phi) | q(i), q(i + \text{lens}[l])$ 
14:      end for
15:    end for
16:  end for
17: end procedure

```

In Tab. 1 we showcase run-time estimates for a random amplitude contraction of a variety of circuits on the fastest supercomputer, Fugaku [13]. To find the run-time estimates, we used the contraction path finder CoTenGra [3, 39] and ran 200 searches each for an hour, using the KaHyPar hypergraph partitioner [46] as the search method with default CoTenGra settings for slicing and sub-tree reconfiguration for a max tensor size of 2^{27} . More details of these settings can be found in the reference [3, 39]. We re-ran this set of searches several times to test for any fluctuations in our fastest estimates, and we found that the fluctuations were within a factor of 2 or less. We show the fastest runs in the table. The code name Sycamore-53- m_x stands for Sycamore-53 to x cycles, GBS- zzz - m_1 stands for a three-dimensional $z \times z \times z$ GBS circuit to 1 cycle. For all the GBS circuits we assume a qudit size of 4, which is likely unrealistic for accurate simulation [8] and there-

Circuit	Slices	Max Size	Time (s)
Syc-53-m20	2^{25}	2.7×10^8	3.5×10^1
Syc-53-m20	0	9.0×10^{15}	1.6×10^1
GBS-444-m1	4^{14}	6.7×10^7	3.6×10^2
GBS-444-m1	0	4.4×10^{12}	1.7×10^{-1}
GBS-666-m1	0	7.9×10^{28}	2.1×10^{14}

Table 1: Estimated runtimes on the Fugaku supercomputer to contract one amplitude for various Sycamore-53 circuits and the three-dimensional GBS circuits GBS-444-m1 and GBS-666-m1, which have qudits arranged in $4 \times 4 \times 4$ and $6 \times 6 \times 6$ lattices and we apply gates to depth 1 ($m=1$). For the GBS circuits we use a qudit size of 4. The time estimate assumes we can compute contractions at the LINPACK benchmark speed for Fugaku, which is 442 PetaFLOP/s [12] and that the contraction path FLOP is a good estimate of the computational work. We showcase estimates that slice no indices (see the rows with slices = 0); these estimates typically would be lower bounds on the possible run-time, but they are unrealistic because they would need more memory than is available on any supercomputer. We also showcase estimates that searched for slices which constrain the rank of intermediary tensors to be below or equal to 27, as contraction trees with this bound on tensor rank usually fit in the memory of a GPU or CPU on a supercomputer node. All estimates assume single-precision complex numbers.

fore will only serve as a lower bound on the actual computation time for an accurate amplitude calculation. One obvious outcome of the results in this table is the fact that GBS-444-m1 is only around an order of magnitude more difficult than Sycamore-53-m20 for random amplitude computation.

In order to benchmark Jet’s task-based parallelism, we run it on the SciNet [47] clusters Niagara [48, 49] and Rouge [50] for our CPU benchmarks, and an AWS EC2 P4d virtual instance for our GPU benchmarks. Niagara is a large cluster of 2,024 Lenovo SD530 servers with 40 Intel Skylake cores at 2.4 GHz each (with 20 cores per socket). The Rouge cluster is a single-socket AMD Epyc 48 core processor at 2.3 GHz each. The AWS instance has two Cascade Lake 24-core processors at 3 GHz, with 8 NVIDIA A100 GPUs each having 40 GB of device memory. For all GPU benchmarks we restrict ourselves to a single device.

In Fig. 7 we showcase benchmarks for the computation of one amplitude of Sycamore-53-m10 across a variety of different systems, with the parameters defining the optimal minimum contrac-

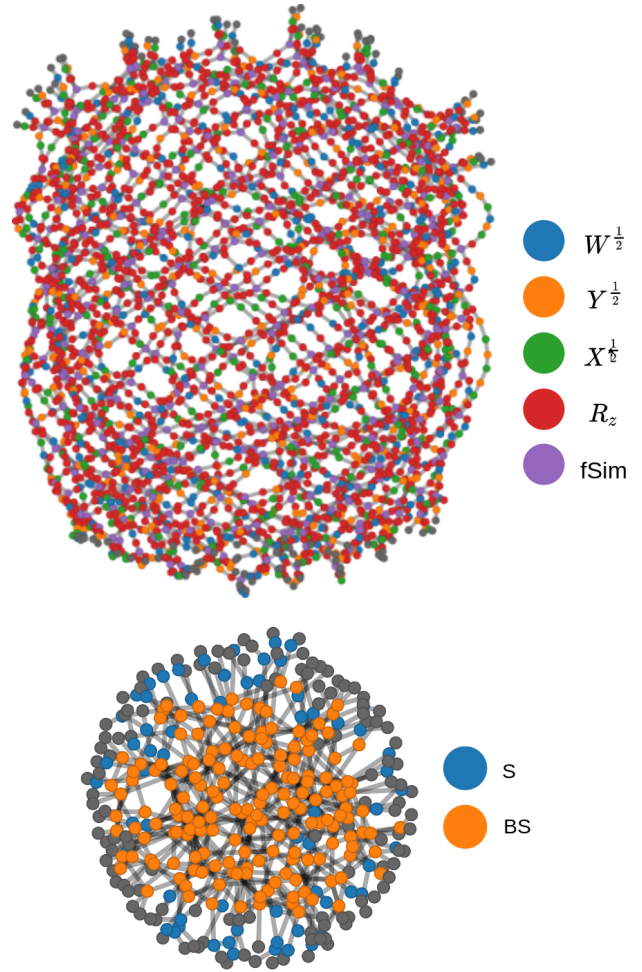


Figure 6: **top**: A plot of the tensor-network for the Sycamore-53 $m = 20$ circuit. The only two qubit gate is fSim. **bottom**: A $4 \times 4 \times 4$ GBS circuit with beam splitters (BS) which act as two-qudit gates and squeezers (S), which act as one-qudit gates.

tion time for the given problem and hardware, averaged over 10 runs. We can see that Jet is comparable to CoTenGra for CPU-based tensor contractions, where we show the work on a full network contraction using both Intel and AMD CPUs. In addition, we demonstrate the effect of slicing for performance, as well as using GPUs. From here we can see that the performance of the contractions are drastically improved by slicing the network, due to the inherent parallelism offered by the method. For the AMD Epyc system, slicing 6 edges (64 slices), and using 1 OpenMP thread and 16 concurrent Taskflow threads reduced the runtime by over a factor of 2. For the Skylake node, the improvement was almost a factor of 4, using 2 OpenMP threads and 16 concurrent tasks. Additionally, for the A100 GPU example, we see

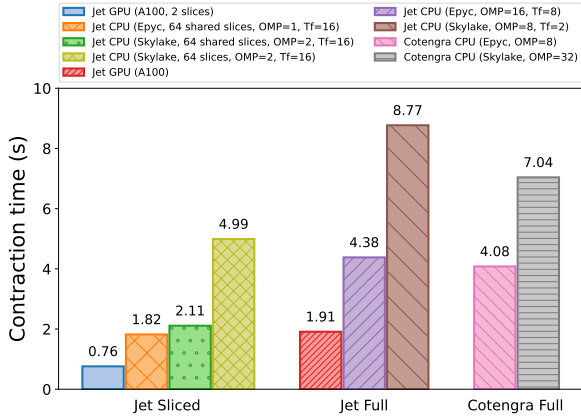


Figure 7: Jet $m = 10$ contraction times for a single amplitude for the full network, and for the sliced network. For comparison, we also show CoTenGra contraction times for the full network contraction on the available hardware. The OMP values correspond to the optimal OpenMP threads for BLAS operations and parallelised code-paths, and Tf to the Taskflow concurrent threads for the task-graph. We can see significant benefit introduced by slicing the network and using shared work, which in the case of the CPU gives approx. $2.4\times$ and $4\times$ speedup on each respective platform, and approx. $2.5\times$ for the GPU.

approximately a 2.5-fold difference. This improvement in performance can be attributed to several factors:

1. By slicing the network, we may fill the task graph with work for individual slices, to be operated on in parallel, and reduce the results once all contractions are completed. This embarrassingly parallel approach allows for the exploitation of multiple-cores/devices.
2. The task-based framework allows for easy control of the number of concurrent threads, enabling the above tasks to be consumed as the cores become available.
3. Our contraction operations are largely a combination of BLAS-3 calls (compute bound) and data permutations (memory-bandwidth bound). Slicing reduces overall tensor sizes, and hence allows more data through the memory bus during permutations, and hence reduced BLAS-3 overhead, enabling better performance all around. The runtime difference between Skylake and Epyc nodes is likely attributed to the memory bus differences between these hardware.

4. The tasks resulting from slicing that have shared work with other slices, and enable a reduction in the overall operations required to contract the network.

For the $m = 10$ dataset, we greedily selected our slices along a fixed contraction path to maximize shared work as much as possible. Our chosen contraction path and set of slices achieved a saving of 45.2 GFLOP (out of a total of 97.2 GFLOP) through use of shared work. We also save computation time because the shared transpose tasks are also no longer needed. In total, our use of shared work provides a $\approx 47\%$ saving in FLOP which is reflected in the more than two-fold speedup between the sliced Skylake run that used shared work (the green bar with 2.11s in Fig. 7) and the sliced Skylake run that did not use the shared work (the yellow bar with 4.99s in Fig. 7). Whether we can get the same savings in FLOP for larger circuits or tensor-networks with different geometry would require having a more accurate method of finding slices and contraction paths that maximize shared work, which is outside the scope of this work.

In Fig. 8 we showcase the peak memory consumption of the $m=10$ runs on a Skylake node. For each of the three different runs, we plot two bars, one bar for the contraction without deletion of intermediary tensors and a second bar for the contraction with deletion. From Fig. 8, we see that deleting intermediary tensors results in drastic savings in memory. In the case of the sliced run that does not use shared work, we see a decrease in the memory consumption by a factor of ≈ 120 . Secondly, we notice that the sliced run that uses shared work (the .68 GB orange bar in Fig. 8), requires more memory consumption than when the shared work is not used (the .14 GB red bar in Fig. 8). This is due to the extra shared intermediary tensors that are stored in memory, which is a minor drawback of using shared-work when slicing. For tensor-networks bigger than the one associated with the $m = 10$ circuit, the amount of slices will grow so that only a subset can be held in memory at a given time. For such cases we can determine a priori the size of this subset. Thus, we do not run into difficulties with this task-based approach as the tensor network grows larger.

We also examine the use of Jet for single slices of $m = 12$, comparing our CPU and GPU back-

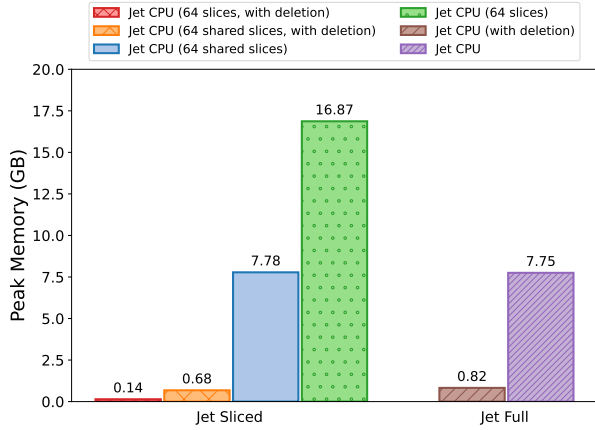


Figure 8: Jet $m = 10$ peak memory consumption in gigabytes (GB) for a single amplitude of the full network and the sliced network. Each simulation is run on a single Skylake node using one thread. The bars labelled "with deletion" correspond to runs where intermediary tensors are deleted after they are no longer required.

ends. Fig. 9 showcases a significant performance difference between running Jet on an NVIDIA A100 GPU and a node of the Niagara supercomputer (Skylake CPU). For the optimal runtime parameters, the GPU is approximately 2 orders of magnitude faster than the CPU, and showcases the gains one can make using GPUs for these contraction problems.

Finally, to date there are no known tensor-network simulators that can run GBS supremacy circuits. We therefore benchmark against TheWalrus code, which computes amplitudes of GBS circuits by calculating a hafnian [51]. Since the 3-dimensional random GBS circuits are computationally expensive, we benchmark on the 2-dimensional random GBS circuit GBS-88-m1 with a squeezing parameter $r = 0.5$. For our tensor-network simulations we use a Fock basis cutoff of 4 and 8 for our tensors. We compute the same random amplitude using both TheWalrus and Jet with differing total photon number, l , and show the results in Fig. 10. The Fock states used for the amplitudes are generated by computing a sequence of random integers in the range $[0, 4)$ with the total sum of these integers representing the total photon number, l .

Figure 10 shows that TheWalrus computation time grows exponentially whereas Jet's computation time remains constant across total photon number at a fixed Fock cutoff. The constant nature of Jet's computation is due to the fact that at a fixed Fock cutoff, changing the to-

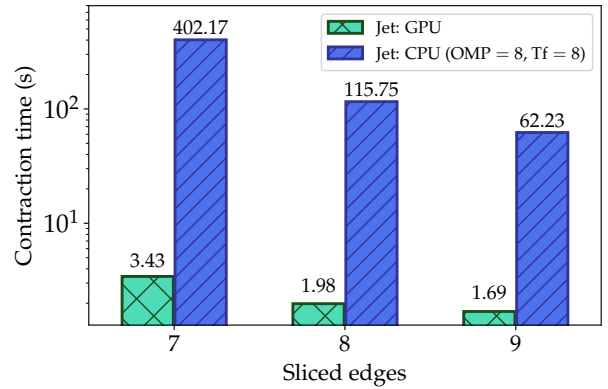


Figure 9: Runtime performance of Jet using CPU and GPU backends to contract a single slice of $m = 12$. The number of sliced edges determines the sliced width, with lower numbers requiring a larger amount of data to be contracted. The CPU code ran using 8 OpenMP threads and 8 concurrent Taskflow threads, giving the optimal performance on a Niagara Skylake node. The GPU code ran on an NVIDIA A100 node with 40GB of on-board memory. Both codes used 32-bit floating point values for the complex components (`complex<float>` and `cuComplex`).

tal photon number only changes the values of some of the tensors in the network, but does not change the sizes of these tensors, or the contraction path. However, due to the Fock basis truncation, there is approximation error inherent in the Jet amplitude calculation. This approximation error becomes larger as the total photon number increases and can make Jet's computation differ from TheWalrus by orders of magnitude. Therefore it would be unfair to make concluding statements about the computation speed of TheWalrus versus Jet at fixed cutoff. We leave a more thorough analysis of the approximation error and computation speed between Jet and TheWalrus to future work. For more details on the approximation error associated with tensor-network computation of optical quantum circuits, see Refs. [52, 53].

5 Conclusion

The use of task-based parallelism for tensor-network simulation of quantum circuits provides multiple benefits which have not been utilized in previous simulators. Firstly, mapping the problem to a task-based framework allows for more parallelism without any extra work. Secondly, we can make use of shared work and potentially max-

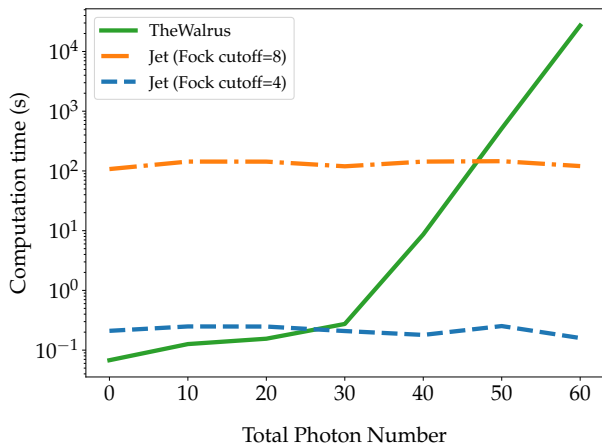


Figure 10: Jet and TheWalrus computation times for random amplitudes with different total photon numbers on a single node of the Niagara supercomputer. The amplitudes are generated by computing a sequence of random integers in the range $[0,4)$ with the total sum of these integers representing the total photon number. We compare times versus TheWalrus. We compute results with Jet using single-precision complex numbers, whereas we use double-precision complex numbers in TheWalrus as this is natively supported. This would likely only slow down TheWalrus by a constant factor of 2, which would not explain the orders of magnitude difference in speed or scaling behaviour observed

imize it during the search for low-cost contraction orders and slices. Furthermore, this shared work can be pre-computed and stored on disk to be used for any slice or amplitude calculation for a given circuit. Thirdly, a task-based method allows for better memory management as the tensors associated with tasks that no longer have any dependencies can be deleted on the fly, which in turn will allow for less slicing as the total memory needed is reduced. Lastly, task-based libraries allow for device agnostic computation, which will be a necessity as supercomputer nodes grow more heterogeneous. In this paper we also, for the first time, benchmark and compare Sycamore-53 supremacy circuits with GBS circuits.

For future work, there are many possible areas of further inquiry. As the problems become more complex (e.g., sampling), the task graphs can become dynamic and the shared work can grow dramatically. One could also extend our tasking to tensor networks with structure such as matrix product states (MPS), projected entangled pair states (PEPS), tree tensor networks and multi-scale entanglement renormalization ansatz tensor networks. We also did not thoroughly

investigate methods for maximizing the shared work between multi-contractions through Eq. 7. While we did use a basic greedy approach for maximizing the shared work between slices of the $m = 10$ runs, this is very rudimentary and tends to give worse results as the tensor networks become larger. Maximizing shared work between multi-contraction is in general a very hard problem to develop heuristics for as it requires optimizing along several different areas over many different contraction trees. We also did not investigate using the SYCL portability layer Taskflow offers for other GPU vendors or more complicated CPU-GPU workflows that Taskflow supports. Investigation of these additional directions is left for future work. Lastly, it would be interesting to develop a framework for tensor-network simulations using distributed asynchronous tasking. Taskflow aims to support this in a future release and this may provide advantages over our current scheme which requires slicing to prevent memory overflow.

6 Acknowledgements

The authors thank Johnnie Gray, Yohai Meiron, Fei Mao, Erik Spence and Ramses Van Zon for helpful discussions. The authors thank SOSCIP for their computational resources and financial support. We acknowledge the computational resources and support from SciNet. SciNet is funded by: the Canada Foundation for Innovation; the Government of Ontario; Ontario Research Fund - Research Excellence; and the University of Toronto. SOSCIP is funded by the Federal Economic Development Agency of Southern Ontario, IBM Canada Ltd. and Ontario academic member institutions.

References

- [1] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, R. Babbush, N. Ding, Z. Jiang, M. J. Bremner, J. M. Martinis, and H. Neven, *Nature Physics* **14**, 595 (2018).
- [2] B. Villalonga, D. Lyakh, S. Boixo, H. Neven, T. S. Humble, R. Biswas, E. G. Rieffel, A. Ho, and S. Mandrà, *Quantum Science and Technology* **5**, 034003 (2020).

- [3] J. Gray and S. Kourtis, *Quantum* **5**, 410 (2021).
- [4] E. Pednault, J. A. Gunnels, G. Nannicini, L. Horesh, T. Magerlein, E. Solomonik, E. W. Draeger, E. T. Holland, and R. Wisnieff, *arXiv preprint* (2017), 10.48550/ARXIV.1710.05867.
- [5] C. Huang, F. Zhang, M. Newman, J. Cai, X. Gao, Z. Tian, J. Wu, H. Xu, H. Yu, B. Yuan, M. Szegedy, Y. Shi, and J. Chen, *arXiv preprint* (2020), 10.48550/ARXIV.2005.06787.
- [6] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. Brandao, D. A. Buell, *et al.*, *Nature* **574**, 505 (2019).
- [7] E. Pednault, J. A. Gunnels, G. Nannicini, L. Horesh, and R. Wisnieff, *arXiv preprint* (2019), 10.48550/ARXIV.1910.09534.
- [8] A. Deshpande, A. Mehta, T. Vincent, N. Quesada, M. Hinsche, M. Ioannou, L. Madsen, J. Lavoie, H. Qi, J. Eisert, D. Hangleiter, B. Fefferman, and I. Dhand, *Science Advances* **8**, eabi7894 (2022).
- [9] K. Bergman, S. Borkar, D. Campbell, W. Carlson, W. Dally, M. Denneau, P. Franzon, W. Harrod, K. Hill, J. Hiller, *et al.*, Defense Advanced Research Projects Agency Information Processing Techniques Office (DARPA IPTO), Tech. Rep **15** (2008).
- [10] S. Heldens, P. Hijma, B. V. Werkhoven, J. Maassen, A. S. Belloum, and R. V. Van Nieuwpoort, *ACM Computing Surveys (CSUR)* **53**, 1 (2020).
- [11] J. Dongarra, J. Hittinger, J. Bell, L. Chacon, R. Falgout, M. Heroux, P. Hovland, E. Ng, C. Webster, and S. Wild, *Applied mathematics research for exascale computing*, Tech. Rep. (Lawrence Livermore National Lab.(LLNL), Livermore, CA (United States), 2014).
- [12] “Top500 Benchmark,” <https://www.top500.org> (2021).
- [13] J. Dongarra, University of Tennessee-Knoxville Innovative Computing Laboratory, Tech. Rep. ICLUT-20-06 (2020).
- [14] P. Thoman, K. Dichev, T. Heller, R. Iakymchuk, X. Aguilar, K. Hasanov, P. Gschwandtner, P. Lemarinier, S. Markidis, H. Jordan, *et al.*, *The Journal of Supercomputing* **74**, 1422 (2018).
- [15] S. R. Paul, A. Hayashi, N. Slattengren, H. Kolla, M. Whitlock, S. Bak, K. Teranishi, J. Mayo, and V. Sarkar, in *European Conference on Parallel Processing* (Springer, 2019) pp. 346–360.
- [16] T.-W. Huang, Y. Lin, C.-X. Lin, G. Guo, and M. D. Wong, *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* (2020), 10.1109/TCAD.2020.3025075.
- [17] L. V. Kale and S. Krishnan, in *Proceedings of the eighth annual conference on Object-oriented programming systems, languages, and applications* (1993) pp. 91–108.
- [18] H. Kaiser, P. Diehl, A. S. Lemoine, B. A. Lelbach, P. Amini, A. Berge, J. Biddiscombe, S. R. Brandt, N. Gupta, T. Heller, *et al.*, *Journal of Open Source Software* **5**, 2352 (2020).
- [19] H. C. Edwards, C. R. Trott, and D. Sunderland, *Journal of Parallel and Distributed Computing* **74**, 3202 (2014).
- [20] J. C. Bridgeman and C. T. Chubb, *Journal of Physics A: Mathematical and Theoretical* **50**, 223001 (2017).
- [21] C. Damm, M. Holzer, and P. McKenzie, *Computational Complexity* **11**, 54 (2002).
- [22] F. Schindler and A. Jermyn, *Machine Learning: Science and Technology* (2020), 10.1088/2632-2153/ab94c5.
- [23] L. Chi-Chung, P. Sadayappan, and R. Wenger, *Parallel Processing Letters* **7**, 157 (1997).
- [24] I. L. Markov and Y. Shi, *SIAM Journal on Computing* **38**, 963 (2008).
- [25] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, and H. Neven, *arXiv preprint* (2017), 10.48550/ARXIV.1712.05384.
- [26] D. Lykov, R. Schutski, A. Galda, V. Vinokur, and Y. Alexeev, *arXiv preprint* (2020), 10.48550/ARXIV.2012.02430.
- [27] S. Kourtis, C. Chamon, E. R. Mucciolo, and A. E. Ruckenstein, *SciPost Phys.* **7**, 60 (2019).
- [28] B. Villalonga, S. Boixo, B. Nelson, C. Henze, E. Rieffel, R. Biswas, and S. Mandrà, *npj Quantum Information* **5**, 1 (2019).
- [29] J. Chen, F. Zhang, C. Huang, M. Newman, and Y. Shi, *arXiv preprint* (2018), 10.48550/ARXIV.1805.01450.

- [30] T. G. Mattson, R. Cledat, V. Cavé, V. Sarkar, Z. Budimlić, S. Chatterjee, J. Fryman, I. Ganey, R. Knauerhase, M. Lee, *et al.*, in *2016 IEEE High Performance Extreme Computing Conference (HPEC)* (IEEE, 2016) pp. 1–7.
- [31] J. Dongarra, L. Grigori, and N. J. Higham, *Philosophical Transactions of the Royal Society A* **378**, 20190066 (2020).
- [32] T. Heller, B. A. Lelbach, K. A. Huck, J. Biddiscombe, P. Grubel, A. E. Koniges, M. Kretz, D. Marcello, D. Pfander, A. Serio, *et al.*, *The International Journal of High Performance Computing Applications* **33**, 699 (2019).
- [33] L. E. Kidder, S. E. Field, F. Foucart, E. Schnetter, S. A. Teukolsky, A. Bohn, N. Deppe, P. Diener, F. Hébert, J. Lippuner, *et al.*, *Journal of Computational Physics* **335**, 84 (2017).
- [34] T.-W. Huang, G. Guo, C.-X. Lin, and M. D. F. Wong, *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* **40**, 776 (2021).
- [35] J. C. Phillips, D. J. Hardy, J. D. Maia, J. E. Stone, J. V. Ribeiro, R. C. Bernardi, R. Buch, G. Fiorin, J. Hémin, W. Jiang, *et al.*, *The Journal of chemical physics* **153**, 044130 (2020).
- [36] P. Jetley, F. Gioachin, C. Mendes, L. V. Kale, and T. Quinn, in *2008 IEEE International Symposium on Parallel and Distributed Processing* (2008) pp. 1–12.
- [37] D. C. Marcello, S. Shiber, O. De Marco, J. Frank, G. C. Clayton, P. M. Motl, P. Diehl, and H. Kaiser, *Monthly Notices of the Royal Astronomical Society* **504**, 5345 (2021).
- [38] A. Alpay and V. Heuveline, in *Proceedings of the International Workshop on OpenCL* (2020) pp. 1–1.
- [39] <https://github.com/jcmgray/cotengra> (2021).
- [40] <https://blogs.nvidia.com/blog/2021/04/12/what-is-quantum-computing/> (2021).
- [41] T. Nguyen, D. Lyakh, E. Dumitrescu, D. Clark, J. Larkin, and A. McCaskey, *arXiv preprint* (2021), 10.48550/ARXIV.2104.10523.
- [42] D. I. Lyakh, *Computer Physics Communications* **189**, 84 (2015).
- [43] <https://github.com/ngnr/saa/qflex> (2021).
- [44] <https://developer.nvidia.com/cutensor> (2021).
- [45] <https://github.com/XanaduAI/jet> (2021).
- [46] S. Schlag, V. Henne, T. Heuer, H. Meyerhenke, P. Sanders, and C. Schulz, in *2016 Proceedings of the Eighteenth Workshop on Algorithm Engineering and Experiments (ALENEX)* (SIAM, 2016) pp. 53–67.
- [47] C. Loken, D. Gruner, L. Groer, R. Peltier, N. Bunn, M. Craig, T. Henriques, J. Dempsey, C.-H. Yu, J. Chen, *et al.*, in *Journal of Physics: Conference Series*, Vol. 256 (IOP Publishing, 2010) p. 012026.
- [48] https://docs.scinet.utoronto.ca/index.php/Niagara_Quickstart (2019).
- [49] M. Ponce, R. van Zon, S. Northrup, D. Gruner, J. Chen, F. Ertinaz, A. Fedoseev, L. Groer, F. Mao, B. C. Mundim, *et al.*, *Proceedings of the Practice and Experience in Advanced Research Computing on Rise of the Machines (learning)*, 1 (2019).
- [50] <https://docs.scinet.utoronto.ca/index.php/Rouge> (2021).
- [51] B. Gupt, J. Izaac, and N. Quesada, *Journal of Open Source Software* **4**, 1705 (2019).
- [52] M. Lubasch, A. A. Valido, J. J. Renema, W. S. Kolthammer, D. Jaksch, M. S. Kim, I. Walmsley, and R. García-Patrón, *Phys. Rev. A* **97**, 062304 (2018).
- [53] R. García-Patrón, J. J. Renema, and V. Shchesnovich, *Quantum* **3**, 169 (2019).