Optimizing Data Collection in Deep Reinforcement Learning

James Gleeson
jgleeson@cs.toronto.edu
University of Toronto, Vector Institute
Toronto, Ontario, Canada

Daniel Snider
dans@cs.toronto.edu
University of Toronto, Vector Institute
Toronto, Ontario, Canada

Yvonne Yang
yvon.yang@mail.utoronto.ca
University of Toronto
Toronto, Ontario, Canada

Moshe Gabel
mgabel@cs.toronto.edu
University of Toronto
Toronto, Ontario, Canada

Eyal de Lara
delara@cs.toronto.edu
University of Toronto
Toronto, Ontario, Canada

Gennady Pekhimenko
pekhimenko@cs.toronto.edu
University of Toronto, Vector Institute
Toronto, Ontario, Canada

ABSTRACT

Reinforcement learning (RL) workloads take a notoriously long time to train due to the large number of samples collected at runtime from simulators. Unfortunately, cluster scale-up approaches remain expensive, and commonly used CPU implementations of simulators induce high overhead when switching back and forth between GPU computations. We explore two optimizations that increase RL data collection efficiency by increasing GPU utilization: (1) *GPU vectorization:* parallelizing simulation on the GPU for increased hardware parallelism, and (2) *simulator kernel fusion:* fusing multiple simulation steps to run in a single GPU kernel launch to reduce global memory bandwidth requirements. We find that GPU vectorization can achieve up to 1024× speedup over commonly used CPU simulators. We profile the performance of different implementations and show that for a simple simulator, ML compiler implementations (XLA) of GPU vectorization outperform a DNN framework (PyTorch) by 13.4× by reducing CPU overhead from repeated Python to DL backend API calls. We show that simulator kernel fusion speedups with a simple simulator are 11.3× and increase by up to 1024× as simulator complexity increases in terms of memory bandwidth requirements. We show that the speedups from simulator kernel fusion are orthogonal and combinable with GPU vectorization, leading to a multiplicative speedup.

CCS CONCEPTS

• **Computing methodologies → Reinforcement learning**.

KEYWORDS

reinforcement learning, GPU, simulation, kernel fusion

1 INTRODUCTION

Reinforcement learning (RL) workloads take a notoriously long time to train due to the large number of samples collected at runtime from simulators. Recent works address this problem by running multiple simulators in parallel across a cluster of accelerator-equipped machines. For example, AlphaZero [25] used 5000 TPUs to perform self-play in parallel reducing training time to 13 days. However, at an hourly cloud pricing of $6.50/hour for an on-demand v1 TPU [1] used at the time, that brings the cost of each trained model to $10, 140, 000. Hence, scaling up RL training is economically impractical outside of large-scale industrial research projects.

A recent survey of RL training workloads [10] demonstrated that the low GPU utilization of RL workloads is caused by data collection, which manifests in two ways. First, CPU simulation time takes up a large amount of training time, with at least 38.1% of training time spent in simulation across popular robotics simulators. Second, a large amount of CPU time originates from overheads induced by switching between CPU-based simulation and GPU-based inference. CUDA API calls alone account for 3.6× as much time on average as the GPU kernel execution in both PyTorch [21] and TensorFlow [2] implementations of an RL algorithm. Hence, to optimize data collection in RL frameworks, GPU utilization must be increased.

In this paper, we explore two potential optimizations for increasing GPU utilization during the time-consuming data collection phase of RL training workloads: (1) *GPU vectorization*, and (2) *simulator kernel fusion*. This work is preliminary, as it focuses on a simple simulator that eased implementation efforts across different ML frameworks thereby enabling a more in-depth analysis. Additional simulators will be explored in future work (Section 7).

GPU vectorization exploits the massive hardware parallelism of GPUs to run N simulator instances in parallel. This leads to performance benefits from a greater amount of simulators that can run in parallel and from reduced CPU overheads (e.g., data copying, API calls) when switching to and from the GPU. Using both DL framework (PyTorch) and ML compiler (XLA [11]) approaches, we are able to achieve up to 1024× speedup in data collection. Prior work [9] also observed that a GPU implementation can have large speedups, but limited their comparisons to only a single ML framework, and only compared against an inefficient single-core CPU implementation. In contrast, we compare and contrast the full gamut of commonly used and high-performance implementations of vectorization. We show that using multiple cores with C++ only accounts for a 4× speedup over the single-core OpenAI implementation. We compare different GPU implementations and

show that XLA can achieve 13.4× speedup over PyTorch. Profiling reveals that XLA amortizes CPU overheads by launching all data collection GPU kernels in a constant number of Python→XLA API calls regardless of the number of steps, whereas PyTorch requires a greater number of Python→PyTorch API calls as data collection steps increase.

Kernel fusion is a GPU optimization that benefits from (1) reduced kernel launch overhead from fewer kernel launches, and (2) increased cache efficiency by avoiding device memory transfers. Many common robotic physics simulators can be modeled as rigid body simulations that are parallelized over the joint states of the simulation [9]. Since these simulation GPU kernels are short, this makes them an attractive target for kernel fusion. To demonstrate this, we fused multiple simulation steps in the simple cartpole [20] simulator and achieved 11.3× speedup. For increasingly complex simulators, we show that speedups of kernel fusion are larger for more memory bandwidth bound simulators since fusion reduces global memory transfers. In particular, kernel fusion can have 8 − 1024× speedup depending on how memory bandwidth bound the simulator is. The speedup from kernel fusion is independent of the number of parallel simulators used, so kernel fusion can be combined with GPU vectorization to achieve massive simulation speedups.

In summary, our contributions are:
- We thoroughly compare the performance limitations of GPU vectorization implementations. Both ML compiler and deep neural network (DNN) frameworks are up to 1024× faster as parallel environments saturate. However, ML compiler approaches outperform DNN framework approaches by 13.4× for smaller parallel environments configurations by amortizing CPU overheads.
- We show that simulator kernel fusion achieves 11.3× speedup for a real simulator and that the maximum speedup increases up to 1024× for memory bandwidth bound simulators of increasing complexity.
- We demonstrate that simulator kernel fusion speedups and GPU vectorization speedups are independent, and both can be combined for massive multiplicative benefits.

2 BACKGROUND

We provide a high-level overview of the RL training procedure with a focus on data collection, and ML frameworks used to implement it. We describe a simple simulator that we use to explore GPU vectorization and simulator kernel fusion optimizations and demonstrate that it is representative of robotics physics simulators commonly used in RL.

2.1 RL Training

The RL training procedure uses backpropagation (shown in Figure 1a), and collects the training dataset at run-time by interacting with a simulator. The data collection process (Figure 1b) consists of a simulation/inference loop, whereby the state of the simulator is fed into the model learned so far to determine which action to take. The selected action determines the next state of the simulator and the resulting reward, and the reward is used to form labels for the collected data.

The data collection loop runs for a pre-determined number of steps K (i.e., a hyperparameter). Different ML frameworks implement the data collection loop differently, which we will show has performance implications. DNN frameworks (PyTorch) execute K separate DL backend API calls for each step, whereas ML compilers (XLA) can condense the entire loop into a single DL backend API call. Once enough data is collected from the simulator, the model is updated using the backpropagation algorithm, which completes a training epoch for the RL algorithm.

To accelerate data collection, RL training frameworks will run multiple simulator instances N. Simulators are typically CPU-based (e.g., Mujoco [26], PyBullet [7], OpenAI gym [4]). Some RL training frameworks use multiple CPU cores to parallelize simulators [14], whereas others opt for flexibility and use only a single CPU core[1]; in our evaluation we consider both. Multiple simulator states are combined into a single minibatch that is fed to inference running on the GPU. Training continues until a pre-determined number of training epochs has completed, after which the average reward per episode will converge to a maximum value.
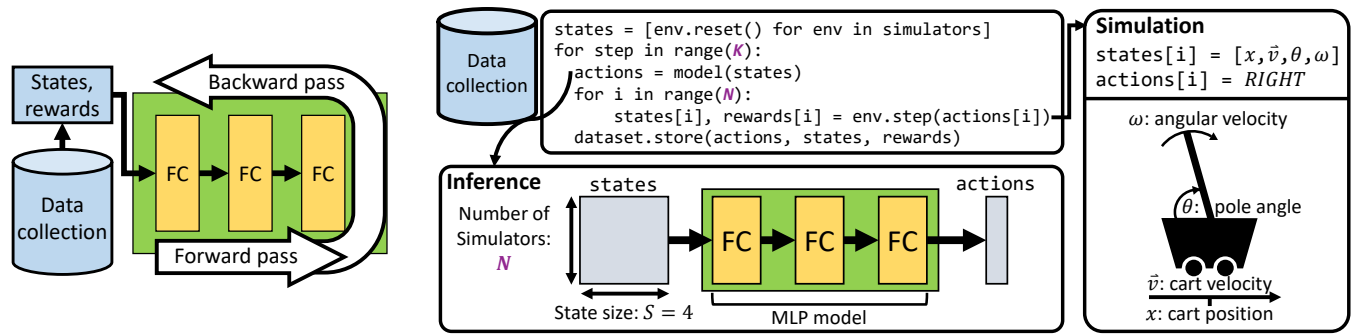
2.2 Cartpole Simulator

The cartpole [20] simulation models a cart moving left and right along a 2D plane to balance a pole, as illustrated in Figure 1b. The simulation state size is small consisting of $S = 4$ floating point numbers: the cart's velocity ($\vec{v}$), position (x), the pole's joint angle relative to the cart (θ), and pole's angular velocity (ω). The simulation dynamics equations consist of simple row-wise transformations that produce a new set of 4 floats.

In this study, we focus on the cartpole simulator since it is simple to implement in DNN frameworks and directly in CUDA [17], and its performance is very similar to other common robotics physics simulators. To test this, we directly compared the simulation throughput of XLA GPU implementations of cartpole against commonly used robotics physics simulators. From Figure 2, we observe that the simulation throughput of cartpole is similar to common robotics simulators used in RL. The main difference is that simulation throughput for cartpole scale up to 2^{19} environments, whereas Brax [9] environments only scale up to 2^{12} environments, which is attributable to the additional compute complexity of Brax simulations which also include collision detection. To understand the effect of simulator complexity on GPU vectorization and kernel fusion speedups, we also vary the compute and memory bandwidth requirements of the cartpole simulation in Section 4.2.

2.3 ML Frameworks

RL developers have a variety of ML frameworks to choose from when implementing RL training algorithms. While most RL developers choose ML frameworks solely based on convenience, in this paper we instead delineate ML frameworks by subtle design choices that affect their performance. As we will see, these performance differences become more pronounced when we investigate optimizing data collection.

---

[1]Python RL frameworks [12] opt to use only a single CPU core due to inefficiencies in Python's shared memory multi-threading implementation [23]

(a) RL training: RL training uses the backpropagation algorithm similar to supervised learning, except that the training data must be collected at runtime.

(b) Data collection: data collection is a simulation/inference loop. Common RL training frameworks use CPU-based simulator implementations (e.g., OpenAI gym), and assemble a single minibatch to perform inference on the GPU. Data collection hyperparameters (N simulator instances, K simulator steps) are shown.

Figure 1: RL training loop: A high-level breakdown of how RL training is performed.
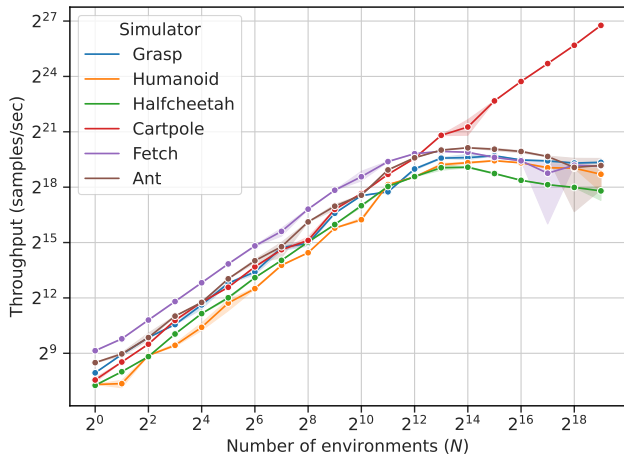

Figure 2: Comparing popular RL simulators: simulation throughput of XLA GPU implementations of popular robotics physics simulators are shown. The throughput of the simple cartpole simulator is similar to robotics simulators, which allows us to focus our analysis on cartpole to reduce implementation complexity.

Figure 3: Optimization 1. GPU vectorization: simulation is vectorized to operate on multiple parallel environments N on the GPU. Multiple implementations can achieve this optimization, either through DNN operators (XLA, PyTorch), or by manually writing kernels to parallelize simulators across GPU threads (CUDA).

DNN frameworks provide a pre-compiled library of GPU kernels for common deep learning operators (e.g., matrix multiplication, convolution, etc.), allowing developers to compose and build model architectures. GPU kernels are either provided by the framework itself (e.g., element-wise operations), or accelerator-specific vendor libraries that have tuned specific operations for high-performance (e.g., matrix multiplication operators from the cuBLAS [16] library when using NVIDIA GPUs). The DNN framework we study in this paper is PyTorch, which is popular for its developer-friendly eager execution model. Eager execution executes operations as GPU kernels as they are issued from Python
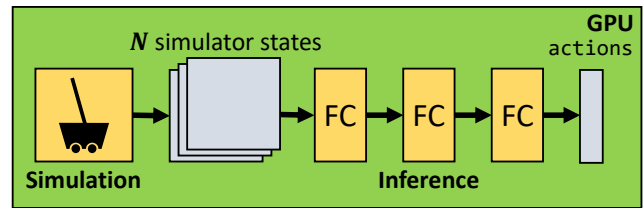
[22], which conveniently allows developers to inspect intermediate DNN computations. While convenient, a greater number of GPU kernel launches can suffer from kernel launch overhead and increased memory bandwidth requirements from reading/writing device memory between kernel launches.

ML compilers parse the entire DNN computational graph and perform global optimizations over that graph to increase execution efficiency. In this paper, we focus on the XLA compiler [11], which targets an intermediate representation consisting of elementary DNN operations which are amenable to execution on both TPU and GPU accelerators. In particular, we use the JAX [3] Python frontend library to XLA. XLA provides an API analogous to high-level linear algebra libraries (i.e., Python's NumPy [15]), but additionally supports just-in-time compilation to accelerator kernels using a symbolic tracing of tensor shapes. As a result, certain simple operations like element-wise operations can be fused into a single GPU kernel launch to increase performance compared to DNN frameworks.

3 GPU VECTORIZATION

We study the performance limitations of GPU vectorization (illustrated in Figure 3) through implementations spanning multiple ML frameworks: an ML compiler (XLA), and a DNN framework (PyTorch). We demonstrate that both XLA and PyTorch have a 1024× speedup when saturating parallel environments. However, at many smaller configurations of parallel environments, XLA outperforms PyTorch by 13.4× by amortizing CPU overheads when combining multiple data collection steps in a single Python→XLA API call.

3.1 Hardware and Software Configuration

All analysis in this paper was performed on hardware consisting of an AMD EPYC 7371 CPU running at 3.1 GHz with 128 GB of RAM and an NVIDIA 2080Ti GPU. For software configuration, we used Ubuntu 20.04, CUDA 11.2.0, GCC 9.3.0, PyTorch v1.8.1, JAX frontend v0.2.13 with `jaxlib` backend (XLA) v0.1.67, and Python 3.8.10.

3.2 Data Collection Implementations

To measure the benefit of GPU vectorization, we created data collection implementations spanning multiple hardware types and ML frameworks. For the CPU implementations of simulation, we consider both the status-quo single-core CPU approach used by popular RL training frameworks (OpenAI gym), and a custom C++ implementation that utilizes multiple CPU cores. For the GPU implementations of simulation, we consider commonly used DNN frameworks (PyTorch), and more performance-oriented ML compilers (XLA).

OpenAI gym: OpenAI gym provides a scalar (single instance) implementation of the cartpole simulator that stores the simulator state within a numpy array of $S = 4$ floating point state values. Each step of the simulator is performed using an object-oriented API with a $step()$ state transition function, with multiple simulator instances implemented using multiple object instances. Due to inefficient shared memory multi-threading in the Python high-level language, RL training frameworks opt to run multiple simulator instances on a single CPU core to maximize their performance. MLP inference uses a single PyTorch model and combines multiple simulator outputs into a single inference minibatch that is copied to a GPU tensor input. This common approach to data collection is illustrated in Figure 1b.

C++: To explore the limits of multi-core CPU-based simulator implementations without being limited by the high-level language overheads of OpenAI gym, we implement the cartpole simulator using shared memory multi-threading. While the simulation runs on the CPU, the inference component still uses PyTorch, with inference/simulation output being shared efficiently through a shared device memory allocation.

PyTorch: We converted the OpenAI gym cartpole simulator into a GPU implementation using PyTorch operators. This is achieved by vectorizing the numpy implementation by storing multiple (N) simulator instances in an $N \times S$ state matrix, and replacing scalar operations with vectorized ones. PyTorch is an eagerly executed DL framework; that is, Python→DL backend API calls are executed on the GPU as the user invokes them from Python.

XLA: Similar to PyTorch, we created a vectorized GPU implementation using XLA operators by storing multiple (N) simulator instances in an $N \times S$ state matrix. In contrast to PyTorch which eagerly executes Python→DL backend calls on the GPU, JAX's numpy-like interface builds an entire symbolic computational graph composed of XLA operators. This graph-based approach allows us to perform multiple data collection *steps* in a single Python→DL backend call which, as we will see, has performance implications.

3.3 Data Collection Throughput

In Figure 4, we run the data collection loop and measure the throughput in samples/second of data collected. Each implementation has a maximum data collection throughput plateau: C++ at 2^{18}, OpenAI gym at 2^{16}, PyTorch at 2^{26}, XLA at 2^{26}. Moving simulation to the GPU using either PyTorch or XLA can provide up to a 1024× speedup over the status quo approach of OpenAI gym's CPU simulation. Using multiple CPU cores with C++ only accounts for at most a 4× speedup over the single-core OpenAI implementation.

> [F.1] **Both DL framework (PyTorch) and ML compiler (XLA) GPU vectorization approaches can provide up to a 1024× speedup in data collection over OpenAI gym.**

Since XLA can express loop control-flow structures, we can define the entire data collection loop as a computational graph, which is invoked with a single Python→XLA API call. Since PyTorch is eagerly executed, control-flow constructs are expressed in Python resulting in multiple calls from Python into PyTorch for each step of data collection. As a result, the data collection throughput for PyTorch does not change as we increase the steps. On the other hand, XLA benefits from increased throughput.

Increasing the number of consecutive data collection steps executed in a single Python→XLA API call from $10^0 \rightarrow 10^3$ results in a 53.8× speedup on average for most environment sizes ($\leq 2^{13}$), which is a 13.4× speedup over PyTorch. An important question this raised was whether XLA could be performing kernel fusion across multiple steps, given that it has complete computational graph knowledge. To investigate this, we performed a profiling deep dive of the XLA implementation.

3.4 Time Breakdown of XLA GPU Data Collection

To understand the underlying bottlenecks that exist in the XLA GPU implementation of data collection, it is helpful to get a time breakdown of where total execution time is spent across the stack. To obtain this time breakdown, we used RL-Scope [10]. RL-Scope provides a full-stack breakdown of time across the CPU, GPU, and within different parts of RL computation, such as inference and simulation.

> [F.2] **In contrast to DNN frameworks (e.g., PyTorch), ML compilers that can express control-flow constructs (e.g., XLA) can reduce CPU-based overhead by reducing Python→Backend API calls. However, more complete computational graph information is not being used to fuse simulation and inference GPU kernels.**
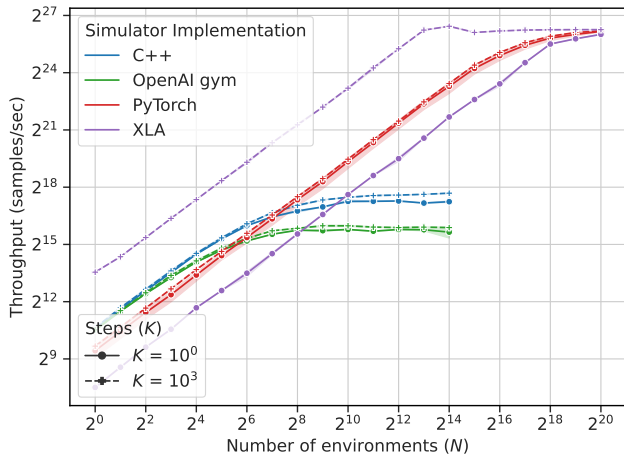
Figure 4: GPU data collection: DL frameworks can express simulation as GPU computation, exploiting GPU hardware parallelism and reducing CPU overheads from switching back and forth to GPU computation. GPU implementations achieve much higher throughput than both CPU implementations (OpenAI gym, C++) of simulation.
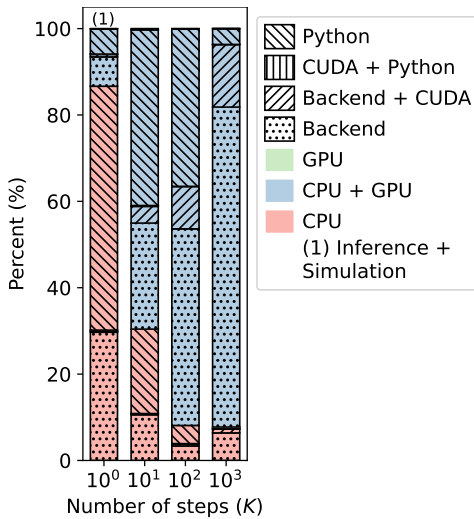

Figure 5: XLA GPU time breakdown: as we increase the number of data collection steps completed in a single Python→XLA API call, CPU overheads from Python and Backend C++ code are amortized, and GPU time saturates the computation. Number of environments $N = 2^{16}$.

In Figure 5, we observe that at small step sizes like 10^0, CPU-bound operations make up 87% of data collection time split across backend C++ calls (30%) and Python time (57%). As we increase the number of steps executed in a single XLA API call, CPU overheads are amortized, which contributes to the overall speedup achieved by the XLA implementation at higher step sizes.
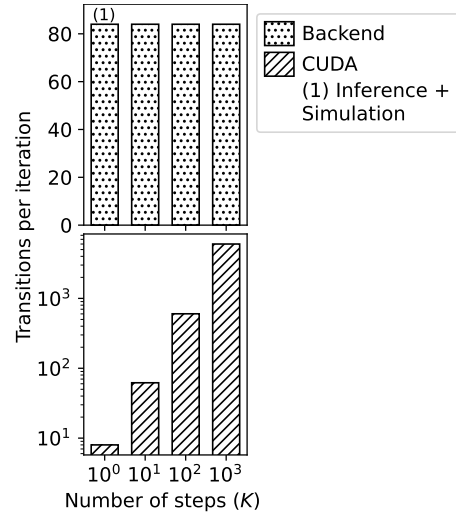

Figure 6: XLA API transition counts: *CUDA* shows the number of GPU kernel launches, and *Backend* shows the number of Python→DL backend (i.e., XLA) API calls. Python→XLA API calls (*Backend*) remain constant even as we increase data collection steps, explaining amortized CPU overheads. However, GPU kernel launches (*CUDA*) increase linearly, indicating kernel fusion does not occur across data collection steps.

To determine whether these performance benefits could be due to kernel fusion by XLA, we measured the number of GPU kernel launches during an iteration of data collection. In Figure 6, *CUDA* shows the number of GPU kernel launches, and *Backend* shows the number of Python→DL backend (i.e., XLA) API calls. In Figure 6, as expected, as we increase the number of steps handled in a single XLA API call, the number of Python→Backend calls remains constant, explaining why CPU overhead was amortized. If we look at the number of GPU kernel launches (*CUDA*), they scale linearly with the step size. This tells us that XLA is not performing kernel fusion at higher step sizes. Further inspection of kernel launches reveals additional calls to simulator kernels, and also cuBLAS matrix multiplication kernels from inference which are inherently infusible since they are closed source.

Our profiling demonstrates that kernel fusion remains an unexplored optimization for the RL training loop. As an initial step in this direction, we next investigate fusing multiple simulation steps to understand the potential performance benefit of fusion.

4 SIMULATOR KERNEL FUSION

GPU kernel fusion can benefit from (1) reduced kernel launch overhead from fewer kernel launches, and (2) increased cache efficiency by avoiding device memory transfers. Based on our profiling results (Section 3.4), more performant ML compiler implementations cannot fuse GPU kernel launches across multiple data collection steps. To study the effects of kernel fusion, we implement a purely CUDA implementation of simulation where we can precisely control the number of GPU kernel launches during simulation (illustrated in
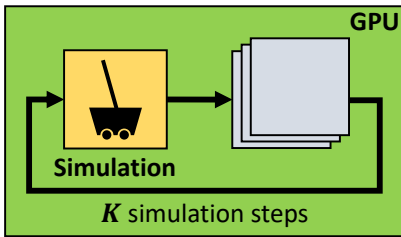
Figure 7: Optimization 2. Simulator kernel fusion: simulation is a short duration kernel that performs row-wise operations, so fusing it into one kernel launch in CUDA will benefit from (1) reduced launch overhead from fewer kernel launches, and (2) increased cache efficiency of simulator state. K simulator steps are executed in a single fused GPU kernel launch.

Figure 7). We demonstrate that kernel fusion is an orthogonal optimization that can be combined with GPU vectorization leading to multiplicative benefits.

4.1 Simulator Kernel Fusion Throughput

Since the simulation kernel for cartpole is short and consists of simple row-wise operations, fusing simulation steps will benefit both from reduced launch overhead and increased efficiency of keeping intermediate simulator state in cache. To explore kernel fusion, we added a CUDA implementation, in addition to the GPU-based PyTorch and XLA implementations of simulation (Section 3.2):

CUDA: To explore the limits of GPU-based kernel fusion and control kernel fusion decisions, we manually implemented a GPU kernel for simulating multiple cartpole instances in parallel. Multiple simulation steps are executed within a single kernel launch, minimizing kernel launch overhead and maximizing data reuse.

In Figure 8, we show the simulation throughput of different implementations of multi-step simulation. Since PyTorch has an eager API, the number of GPU kernel launches scales linearly with the number of steps and cannot benefit from kernel fusion, causing the $K = 10^0$ and $K = 10^3$ lines of PyTorch to overlap. Since XLA is provided with the full computational graph for iteratively performing simulation, multiple simulation steps can be condensed into a single Python→XLA call, amortizing CPU overheads, leading to a 83.3× speedup for XLA going from $K = 10^0 \rightarrow 10^3$.

Since CUDA is handwritten to execute multiple simulator steps in a single kernel launch, it can benefit both from reduced kernel launch overhead and better caching of simulator state in registers. The 112.8× speedup for CUDA→XLA for $K = 10^3$ benefits both from kernel fusion and reduced framework API overhead. To isolate the benefit from kernel fusion alone, we can consider the speedup for CUDA going from $K = 10^0 \rightarrow 10^3$ steps; we see that kernel fusion accounts for a 11.3× speedup.

> [F.3] **For repeated simulation steps, kernel fusion can provide** 11.3× **speedup over an unfused implementation of the cartpole simulator.**
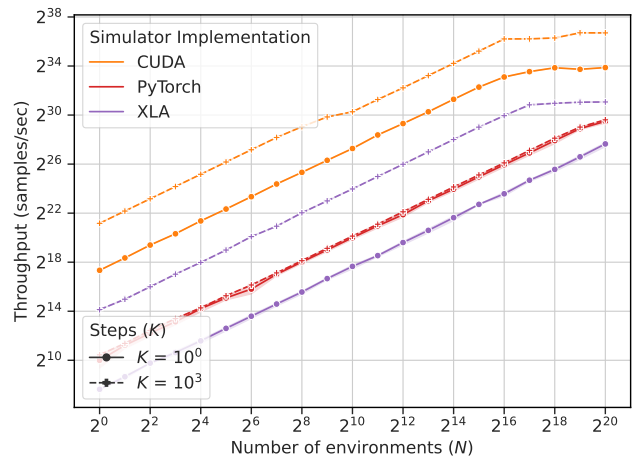

Figure 8: Simulator kernel fusion: simulation throughput benefits from increased fusion at a higher number of fused simulation steps K. XLA does not fuse GPU kernels but still benefits from amortizing CPU overheads. CUDA benefits from fusing kernel launches (i.e., reduced launch overheads, cached simulator state).

4.2 Simulator Complexity

An important consideration is how the speedup from fusion varies as we increase the complexity of the simulator. In particular, the simulator could be more compute-bound, more memory bandwidth bound, or both. To answer this question, we took the cartpole simulator and added configurable parameters to increase the compute and memory requirements of the simulation. The compute factor C increases the compute requirements of cartpole by $C\times$ by increasing the number of simulation iterations. The state factor S increases the cartpole state size by a factor of $S\times$, resulting in $S\times$ as many global memory loads and stores. To measure the speedup due to fusion, we normalize the simulation throughput with respect to a fusionless run (i.e., steps $K = 10^0$). Figure 9 shows the speedup from fusion along all configurable workload dimensions (parallel environments N, fused steps K, state factor S, compute factor C).

Number of parallel environments (N): As we increase N (left to right across subplots), the speedup from fusion remains the same. For example, the last row of plots (steps $K = 10^3$) all plateau to a fusion speedup of 2^{10}. Hence, the benefit from kernel fusion is independent of how many parallel environments are used; kernel fusion is an orthogonal and combinable optimization with GPU vectorization.

State size factor ($S\times$): As we increase S (left to right along the x-axis of each subplot), the simulation becomes increasingly memory bandwidth bound due to increased global memory loads/stores. Increasingly memory bandwidth bound kernels benefit more from fusion since fusion reduces global memory transfers. For example, for the last row of plots (steps $K = 10^3$), the speedup varies from 2^4 to 2^{10} as S increases.

Number of fused steps (K): As we increase K (top to bottom across subplots) the speedup increases more significantly in memory bandwidth bound workloads. For example, less fusion
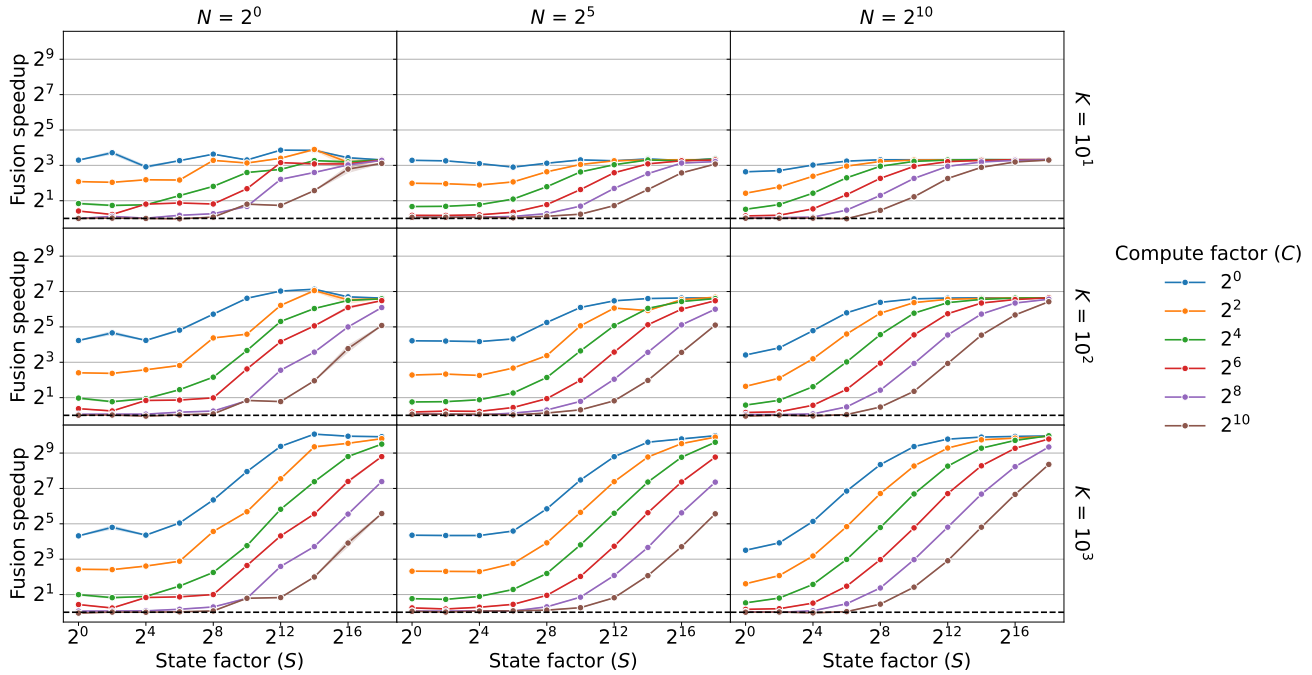
Figure 9: Simulator complexity: the compute factor C increases cartpole compute requirements by $C\times$ and state factor S increases memory bandwidth requirements by $S\times$. The fusion speedup is measured by normalizing with respect to a fusionless run (i.e., steps $K = 10^0$). Parallel environments N increase from left to right subplots, and number of fused steps K from top to bottom subplots.

(steps $K = 10^1$) achieves at most 2^3 speedup, whereas more fusion (steps $K = 10^3$) achieves at most 2^{10} speedup at large state factors ($S \geq 2^{11}$).

Compute factor (C): As we increase C (the lines of each subplot), the simulation becomes increasingly compute-bound due to additional simulation iterations. The more compute-bound the workload is, the less speedup that comes from fusion since the workload is not benefitting from greater data reuse by reducing global memory transfers. In the worst case where the simulation is entirely compute bound ($C = 2^{10}$) and not memory bandwidth bound ($S = 2^0$), fusion performs just as poorly as without fusion.

> $\boxed{F.4}$ **Kernel fusion benefits memory bandwidth bound simulation kernels with low to moderate compute complexity by reducing global memory transfers. The speedup from kernel fusion is independent of how many parallel simulators are used and should be considered an orthogonal and hence combinable optimization.**

5 RELATED WORK

Related work falls into two categories, which we summarize here and discuss in detail below. Prior work has demonstrated training time speedups from moving simulation to the GPU by performing *GPU vectorization* either by expressing simulation as DNN operators [9] or with manually written CUDA code [8, 24]. In contrast to those works, we perform an apples-to-apples comparison of multiple vectorization implementations to understand inherent

performance limitations associated with hardware and ML framework choice. We demonstrate that ML compiler frameworks can outperform DNN framework implementations by amortizing CPU overheads, but that current ML compiler frameworks cannot perform kernel fusion that would benefit RL training workloads. *Kernel fusion* is a common optimization pass for ML compilers but is typically limited to element-wise operations [5], with some compilers opting to offload high-performance operations to infusible cuBLAS libraries [11]. In this paper, we do not limit ourselves to the current capabilities of ML compilers, and instead measure the potential benefits achievable if kernel fusion were implemented today in these ML compilers by exploring manual fusion of simulation kernels implemented in CUDA.

GPU vectorization: Brax [9] implements a general physics engine in JAX that can support many environments consisting of connected rigid bodies, allowing them to support common robotics physics simulators. The physics engine implementation parallelizes the computation across both environments and joint angles. Brax leverages accelerator (i.e., GPU, TPU) parallelism during simulation in RL training, allowing them to achieve significant training time speedups (266× on a V100 GPU) over simple single CPU core simulators [26]. In contrast to Brax, our evaluation of the performance benefit of GPU vectorization is more thorough for four reasons. First, we consider a multi-core C++ implementation for the CPU instead of just a naive single-core CPU implementation. Second, we demonstrate a 53.8× speedup in XLA by combining multiple data collection steps in a single Python→XLA API call to amortize

CPU overheads. Third, besides an ML compiler approach (XLA), we also compare against common DL framework implementations (PyTorch) and demonstrate how the superior 13.4× performance of XLA over PyTorch on the GPU is achieved through multiple data collection steps. Finally, through profiling, we show that XLA cannot perform kernel fusion in data collection, which motivates us to investigate kernel fusion as a novel and orthogonal optimization in RL training.

CuLE [8] explores porting the C++ Atari emulator to CUDA to run on GPUs, with each GPU thread executing an Atari hardware emulator instance. Emulating the Atari CPU architecture on the GPU leads to degradation in performance due to branch divergence. As a result, the observed speedup over CPU is at most 2.54× when training Atari Assault on a Titan V GPU. In contrast to this paper, CuLE implements the Atari simulation only in CUDA; it is likely not possible to express non-data-parallel operations like hardware instruction emulation using DNN framework or ML compiler supported operators. We instead focus on physics simulators that are more amenable to the GPU architecture, allowing us to investigate the performance limitations of both DNN frameworks and ML compiler implementations, and explore kernel fusion optimizations that are not explored in CuLE.

Kernel fusion: TVM [5] is an ML compiler that enables developers to write DNN operators in a high-level domain-specific language (DSL), with the final output being lowered to various different hardware backends (e.g., GPU, TPU, FPGA). TVM implements several optimization passes over the computational graph, including operator fusion. TVM operator fusion is limited to fusing with element-wise operations (called "injective" operators) as they are the simplest to fuse. In particular, injective operators fuse with: other injective operators (e.g., $BatchNorm \circ relu$), the input of a reduction (e.g., $scale \circ sum$), or the output of a complex-out-fusible (e.g., $conv2d \circ relu$). We chose not to use TVM to implement fusion since our simulator is row-wise not element-wise, and we wanted to tightly control how the simulator state is managed during fusion and explore performance implications.

XLA [11] is a compiler for linear algebra programs (e.g., DNNs) with operators that can be lowered to backend targets (i.e., TPU, GPU, CPU). Frontend APIs like JAX [3] builds on top of XLA by offering JIT compilation of operator graphs into optimized backend code. The operator fusion compiler pass heuristically tries to reduce memory bandwidth [13]. Producer-consumer fusion allow fusion of element-wise operations (e.g. fused $add \circ relu$). Sibling fusion fuses operators that share a common input to reduce global memory reads, interleaving the operator outputs as a tuple (e.g., fusing $sum(x)$ and $sum(x^2)$ for computing $BatchNorm$). In this paper, we found that kernel fusion across the inference and simulation boundary is not possible since XLA still relies on closed-source cuBLAS matrix multiplication kernel implementations. This motivated us to explore the potential benefits of performing kernel fusion in the simulation phase of training by manually implementing simulation CUDA kernels, and how these speedups vary with respect to simulator complexity. Our exploration of the circumstances under which simulators benefit from kernel fusion are important for incorporating compiler-based kernel fusion optimizations that target RL workloads.

6 CONCLUSIONS

We demonstrated that the speedups from GPU vectorization are substantial, with up to 1024× speedup for a simple simulator. We showed that ML compilers (XLA) outperform DNN frameworks (PyTorch) by 13.4× by amortizing CPU overheads across multiple backend API calls. For simulator kernel fusion, we showed that the speedups from kernel fusion for a simple simulator are 11.3× and increase by up to 1024× as simulator complexity increases in terms of memory bandwidth requirements. The speedups from kernel fusion are orthogonal and combinable with GPU vectorization, leading to a multiplicative speedup. We hope our study spurs greater interest in specialized optimizations targeting emerging RL workloads.

7 FUTURE WORK

Our initial analysis demonstrated that large speedups are possible with both GPU vectorization and simulator kernel fusion, and that these speedups are combinable. In future work, we will study additional simulators such as robotics physics simulations, and perform fusion across simulation and inference GPU kernels to speedup the data collection process in today's RL workloads.

Simulation/inference fusion: We limited our implementation of kernel fusion in CUDA to multiple simulation steps, but did not include inference. The benefit of this approach is that it reduced engineering complexity and allowed us to compare the performance and inherent limitations of several approaches to optimizing multi-step simulation (XLA, PyTorch, CUDA). However, to accelerate RL training, in the future we must apply fusion to the full data collection simulation/inference loop. Given that we cannot manually fuse closed-source GPU kernels, we will need to obtain high-performance open-source GPU kernels with performance comparable to cuBLAS. Based on our analysis of large fusion benefits in memory bandwidth bound simulators, we suspect that kernel fusion will still benefit DNN inference computations that have memory bandwidth bound behaviour. Memory bandwidth bound kernels are known to occur in matrix multiplication for irregularly shaped tall/skinny matrices [6] that are typical for RL inference due to the simulator state matrix.

Exploring additional simulators: Our analysis focused on the simple cartpole simulator, which allowed us to explore the influence of ML framework and hardware on kernel fusion across different implementations. Further, our study of simulator complexity demonstrated that increasingly memory bandwidth bound simulators benefit the most from kernel fusion. Hence, our immediate next step will be to explore additional simulators of varying memory/compute complexity to investigate how many existing simulators can benefit from kernel fusion. Our first step in this direction will be to re-implement popular robotics simulators from Brax [9] in CUDA and see how manual kernel fusion compares to the kernels produced by XLA. We would also like to explore photorealistic simulators built on industrial video game engines used in autonomous driving [18] and factory robotics scenarios [19] since these simulators are large contributors to total training time.

REFERENCES

[1] 2018. Cloud TPU machine learning accelerators now available in beta. https://cloud.google.com/blog/products/gcp/cloud-tpu-machine-learning-accelerators-now-available-in-beta.

[2] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. 2016. TensorFlow: A system for large-scale machine learning. In *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*. 265–283.

[3] James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. 2018. *JAX: composable transformations of Python+NumPy programs.* http://github.com/google/jax

[4] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. 2016. OpenAI gym. *arXiv preprint arXiv:1606.01540* (2016).

[5] Tianqi Chen, Thierry Moreau, Ziheng Jiang, Haichen Shen, Eddie Yan, Leyuan Wang, Yuwei Hu, Luis Ceze, Carlos Guestrin, and Arvind Krishnamurthy. 2018. TVM: end-to-end optimization stack for deep learning. *13th USENIX Symposium on Operating Systems Design and Implementation (OSDI 18)* (2018).

[6] Benjamin Y Cho, Jeageun Jung, and Mattan Erez. 2021. Accelerating bandwidth-bound deep learning inference with main-memory accelerators. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis.* 1–14.

[7] Erwin Coumans and Yunfei Bai. 2019. Bullet Real-Time Physics Simulation | Home of Bullet and PyBullet: physics simulation for games, visual effects, robotics and reinforcement learning. https://pybullet.org.

[8] Steven Dalton and Iuri Frosio. 2020. Accelerating reinforcement learning through GPU atari emulation. *Advances in Neural Information Processing Systems* 33 (2020), 19773–19782.

[9] C. Daniel Freeman, Erik Frey, Anton Raichuk, Sertan Girgin, Igor Mordatch, and Olivier Bachem. 2021. *Brax - A Differentiable Physics Engine for Large Scale Rigid Body Simulation.* http://github.com/google/brax

[10] James Gleeson, Srivatsan Krishnan, Moshe Gabel, Vijay Janapa Reddi, Eyal de Lara, and Gennady Pekhimenko. 2021. RL-Scope: Cross-Stack Profiling for Deep Reinforcement Learning Workloads. In *Proceedings of Machine Learning and Systems.*

[11] Google. 2022. XLA: Optimizing Compiler for Machine Learning. https://www.tensorflow.org/xla.

[12] Ashley Hill, Antonin Raffin, Maximilian Ernestus, Adam Gleave, Anssi Kanervisto, Rene Traore, Prafulla Dhariwal, Christopher Hesse, Oleg Klimov, Alex Nichol, Matthias Plappert, Alec Radford, John Schulman, Szymon Sidor, and Yuhuai Wu. 2018. Stable Baselines. https://github.com/hill-a/stable-baselines.

[13] Thomas Joerg. 2019. Automated GPU Kernel Fusion with XLA. https://www.youtube.com/watch?v=u3cWOd99xX0. EuroLLVM Developers' Meeting.

[14] Philipp Moritz, Robert Nishihara, Stephanie Wang, Alexey Tumanov, Richard Liaw, Eric Liang, Melih Elibol, Zongheng Yang, William Paul, Michael I Jordan, et al. 2018. Ray: A distributed framework for emerging AI applications. In *13th USENIX Symposium on Operating Systems Design and Implementation (OSDI 18)*. 561–577.

[15] NumPy. 2022. *NumPy.* https://numpy.org/

[16] NVIDIA. 2022. *cuBLAS :: CUDA Toolkit Documentation.* https://docs.nvidia.com/cuda/cublas/index.html

[17] NVIDIA. 2022. CUDA Toolkit. https://developer.nvidia.com/cuda-toolkit.

[18] NVIDIA. 2022. *NVIDIA DRIVE Sim.* https://developer.nvidia.com/drive/drive-sim

[19] NVIDIA. 2022. *NVIDIA Isaac Sim.* https://developer.nvidia.com/isaac-sim

[20] OpenAI. 2022. Gym - CartPole-v1. https://gym.openai.com/envs/CartPole-v1.

[21] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems 32*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (Eds.). Curran Associates, Inc., 8024–8035. http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf

[22] Python. 2022. Welcome to Python.org. https://www.python.org/.

[23] Antonin Raffin. 2018. RL Baselines Zoo - Single CPU Core. https://github.com/araffin/rl-baselines-zoo/blob/ff84f398a1fae65e18819490bb4e41a201322759/train.py#L261.

[24] Brennan Shacklett, Erik Wijmans, Aleksei Petrenko, Manolis Savva, Dhruv Batra, Vladlen Koltun, and Kayvon Fatahalian. 2021. Large batch simulation for deep reinforcement learning. *International Conference on Learning Representations* (2021).

[25] David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dharshan Kumaran, Thore Graepel, et al. 2018. A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. *Science* 362, 6419 (2018), 1140–1144.

[26] Emanuel Todorov, Tom Erez, and Yuval Tassa. 2012. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems.* IEEE, 5026–5033.