
EXPLORING THE LIMITS OF CONCURRENCY IN ML TRAINING ON GOOGLE TPUS

Sameer Kumar¹ Yu Emma Wang¹ Cliff Young¹ James Bradbury¹ Anselm Levskaya¹ Blake Hechtman¹
Dehao Chen¹ HyoukJoong Lee¹ Mehmet Deveci¹ Naveen Kumar¹ Pankaj Kanwar¹ Shibo Wang¹
Skye Wanderman-Milne¹ Steve Lacy¹ Tao Wang¹ Tayo Oguntebi¹ Yazhou Zu¹ Yuanzhong Xu¹
Andy Swing¹

ABSTRACT

Recent results in language understanding using neural networks have required training hardware of unprecedented scale, with thousands of chips cooperating on a single training run. This paper presents techniques to scale ML models on the Google TPU Multipod, a mesh with 4096 TPU-v3 chips. We discuss model parallelism to overcome scaling limitations from the fixed batch size in data parallelism, communication/collective optimizations, distributed evaluation of training metrics, and host input processing scaling optimizations. These techniques are demonstrated in both the TensorFlow and JAX programming frameworks. We also present performance results from Google’s recent submission to the MLPerf-v0.7 benchmark contest, achieving record-breaking training times from 16 to 28 seconds in four MLPerf models on the Google TPU-v3 Multipod machine.

1 INTRODUCTION

The deep learning revolution is in the midst of a “space race” in the field of language understanding, with leading research labs training and publishing papers about a sequence of models of exponentially increasing size. One of the early breakthroughs was Google’s Neural Machine Translation System (Wu et al., 2016), which used LSTMs (Hochreiter & Schmidhuber, 1997) and Attention (Luong et al., 2014; Bahdanau et al., 2014) to achieve a significant quality improvement. GNMT was rapidly followed by Transformers (Vaswani et al., 2017) which parallelized over input sequences, allowing faster training than sequentially limited LSTMs. Transformers in turn are a fundamental component of BERT (Devlin et al., 2018) models, which are able to “pre-train” for general linguistic knowledge, then “fine-tune” to particular language tasks, including Translation. The latest GPT-3 model appears to be able to compose plausible essay-length arguments, albeit with some degree of human guidance or selection (Brown et al., 2020). The size of these models is growing exponentially; OpenAI observed that the training resources for state-of-the-art deep learning models appears to be doubling every 3.5 months (Amodei et al., 2018).

Training such models requires correspondingly large machines. In 2012, the breakthrough AlexNet paper (Krizhevsky et al., 2012) trained with model parallelism over two GPUs. That same year, Google harnessed their datacenter-scale CPU clusters to train asynchronously in the DistBelief system (Dean et al., 2012). The Deep Learning

revolution sparked huge investments in GPUs: NVIDIA revenues rose an average of 50% year-over-year every quarter from mid-2016 to mid-2018 (MacroTrends.net, 2020). By 2015, Google had built a specialized neural network accelerator, the Tensor Processing Unit (TPU), a single chip which offered over a 10x improvement in performance/watt, peak performance, and inference latency (Jouppi et al., 2017). Within two years, Google’s second-generation TPU used 256-chip pods to train a single model with near-perfect parallel scaling (Jouppi et al., 2020); the third-generation TPU increased pod size to 1024 (Jouppi et al., 2020; Kumar et al., 2019). NVIDIA and other GPU suppliers have fielded clusters of similar scale, with Microsoft and OpenAI constructing a 10,000-GPU cluster (Langston, 2020). The space-race uses increasingly accurate models to approach Artificial General Intelligence, but there is no doubt that the hardware being fielded is also astronomically ambitious.

Unlike the space race, where low-earth orbit and the moon make for obvious milestones, the best way to measure the accomplishments of these parallel machines is less concrete. Benchmarking competitions can serve this purpose: AlexNet surprised and transformed the vision community by winning the ImageNet Large-Scale Visual Recognition Competition (Russakovsky et al., 2015) in 2012. Computer architects and system builders recognized the need for a benchmark suite similar to SPEC and TPC in their field, and a broad coalition of universities and companies founded MLPerf in 2018 to serve this need (mlp). In particular, the MLPerf Training division (Mattson et al., 2019) attracts




Figure 1. TPU-v3 1-pod vs 4-pods in the Google datacenter.

HPC-scale entries, as submissions compete to reach state-of-the-art accuracy on parallel training problems on massively parallel clusters in minimum wall-clock time. The techniques used in MLPerf submissions generally benefit the deep learning community, as they are folded into systems, libraries, compilers, and best-practices application code. This paper focuses on Google’s MLPerf 0.7 Training submission, and explains the algorithmic, architectural, performance, and system-tuning techniques that demonstrated world-class training at scale.

MLPerf (Mattson et al., 2019) is a machine learning benchmark suite that is designed to benchmark different classes of ML accelerators and frameworks on state-of-the-art ML tasks. It has gained industry wide support and recognition. The recently concluded MLPerf-v0.7 Training submission round has submissions from NVIDIA, Google, AliBaba, Fujitsu, Shenzhen Institute and Intel. Along with CPUs and NVIDIA GPUs, benchmarked hardware included the Google TPU-v3 and TPU-v4 as well as an AI accelerator from Huawei. ML frameworks included PyTorch, TensorFlow, JAX, MXNet, MindSpore and Merlin HugeCTR.

Like systems benchmark suites which have come before it, the MLPerf benchmark suite is pushing performance forward and our MLPerf-v0.7 Training submission on Google TPU-v3 and TPU-v4 systems showcase the large scale we are able to achieve. The MLPerf-v0.7 rules add new models, namely: i. BERT, a large language model, ii. DLRM, a deep learning recommendation system, and iii. an enhanced larger version of MiniGo to achieve higher scalability. An MLPerf training benchmark involves training a model (e.g., BERT) on a specific dataset (a Wikipedia dump) to a pre-defined convergence test metric while following specific methodology for parameters, optimizations, and timing.

In order to explore the limits of concurrency in the MLPerf models we assembled a TPU-v3 Multipod with 4096 chips, with 105 TFLOPS per chip at peak. It is four times larger than the TPU-v3 pod used for the MLPerf-v0.6 training benchmark submission. A 4-pod Multipod configuration




Figure 2. TPU-v3 4-pod configuration where cross-pod links connect neighboring TPU-v3 pods in the Google datacenter.

with 4096 TPU-v3 chips is shown in Figure 1. Here the two pods are connected along the X-dimension of the mesh by the cross-pod optical links (Figure 2). These links are longer than standard TPU-v3 within-pod links. The MLPerf benchmarking was done on a 4-pod Multipod with 4096 chips in a 128x32 2-D mesh topology (with within-pod torus links at the Y edges). As the TPU-v3 chip had only 1024 entries in the routing table, we used a sparse routing scheme where only neighbors along rows and columns were visible to each chip. This was sufficient for achieving peak throughput in the all-reduce communication operations.

We chose a subset of MLPerf models to benchmark at the Multipod scale. These included i) BERT, ii) ResNet-50, iii) Transformer and iv) Single Shot Detector (SSD). In BERT and ResNet-50 we used batch parallelism to scale to the Multipod, while in Transformer and SSD we used a combination of model parallelism and batch parallelism techniques. The Mask-RCNN and DLRM models are also discussed in this paper. For the Mask-RCNN model, the available batch parallelism is extremely limited and we therefore present results on a slice with 512 TPU-v3 chips. For the DLRM model, scalability is capped by limited global batch size and communication overheads quickly outweigh scale-out benefits. We present results on a slice with 256 TPU-v3 chips.

This paper made the following major contributions.

- A world-record scale ML architecture with 4096 nodes. It was the biggest machine for MLPerf-v0.7 and it put extra pressure on the dedicated interconnect. This machine extends the X-dimension with cross-pod optical links that have higher latency and lower bandwidth than the links within pods. To mitigate the link speed difference, we designed a novel all-reduce algorithm that pushes most of the all-reduce payload along the Y dimension that results in high throughput as communication along X-dimension is reduced by a factor that is the same as the Y-dimension size (32).
- Optimized global summation for model parallelism. The current state-of-the-art MeshTF (Shazeer et al.,

2018) maps language models along batch and model dimensions that are then mapped to the physical 2-D mesh of TPU-v3. We found this approach had significantly high communication overheads as the gradient all-reduce step is executed on a 1-D ring. We present a novel strided communication optimization scheme that enables high throughput in both the forward and the gradient reduction steps, that results in the MLPerf Transformer model training in 16 seconds.

- We scale weight update sharding (distributed optimizer) in a complex hybrid of data and model parallelism scenario via model parallelism and spatial partitioning.
- Analysis of the JAX programming model and comparison with TensorFlow. This is the first paper that studies JAX at scale, uses JAX on TPU (multi) pods, and uses model parallelism techniques (SPMD partitioning and weight update sharding) in JAX. The JAX results demonstrate the generality of TPUs and the enhancements added to XLA, and provide a useful comparison for multi- vs. single-controller design points in distributed ML systems.
- Multipod Performance Results. Four models finish training in under 30 seconds. BERT and DLRM, the models recently added to MLPerf-v0.7, are optimized at a TPU Multipod scale for the first time.

2 MULTIPLE FRAMEWORKS

While the primary frontend for TPUs has historically been TensorFlow (Abadi et al., 2016), the hardware and XLA compiler are general enough to support other programming environments. Therefore in this paper, we chose to benchmark both TensorFlow and JAX (Frostig et al., 2018), a new, research-oriented numerical computing system based on XLA (TensorFlow.org, 2020). Both systems required additional software engineering to scale effectively to the Multipod, but they ultimately achieved similar benchmark results.

As shown in Figure 3, two architectural differences between TensorFlow and JAX differentiate their performance at scale. First, they have different staging approaches. TensorFlow embeds an expressive and dynamic intermediate language (TensorFlow graphs that can span both accelerators and CPU hosts) in Python, and then JIT-compiles subsets of these graphs with XLA. Meanwhile, JAX has one fewer stage: it is a staged programming environment that embeds JIT-compiled XLA programs (for static compiled performance on accelerators and parallelism on the accelerator network) in the Python host language (used for dynamic and on-accelerated computations). As a consequence, TensorFlow has additional compilation steps, which we acceler-

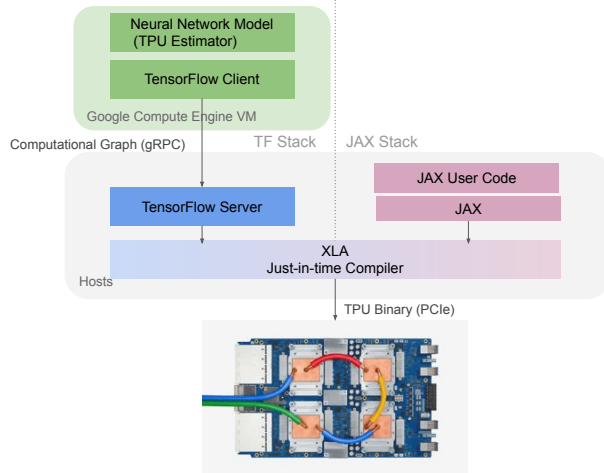


Figure 3. Stack view of the TF and JAX frameworks on the TPU-v3 machines.

ated using multithreading, while JAX requires more careful management of Python bottlenecks (for instance, moving blocking tasks like data infeed off of the main thread).

Second, they enable different distributed programming models. JAX adopts a multi-client approach to distributed programming, running a separate copy of the same JAX code (including the Python interpreter) on each host in the pod. The programs communicate with each other in only two ways: at startup time, to coordinate TPU mesh setup, and in XLA-compiled collectives such as all-reduce that operate over the dedicated TPU network during model training. On the other hand, TensorFlow programs TPUs with a single-client approach, giving one Python process (running either on one of the hosts in the pod or elsewhere) global visibility and control over the entire distributed system. The rest of the TPU hosts run a TensorFlow server that executes partitioned subsets of TensorFlow graphs sent via RPCs from the client over the datacenter network.

These two approaches differ in usability and performance characteristics. While TensorFlow’s single-client distributed system enables user code that directly reflects the overall workload, JAX’s multi-client approach enables more direct control of the code that runs on each worker. JAX invokes the XLA compiler independently on each host—relying on deterministic compilation to avoid incompatibilities between the resulting programs—while TensorFlow compiles once and distributes the binaries to the workers. The TensorFlow representation of multi-device graphs can also cause Amdahl’s law bottlenecks, as the client process incurs graph construction and optimization time proportional to the number of workers, while JAX setup times (other than TPU topological mesh initialization) do not change significantly with an increase in the number of workers.

3 SCALABILITY TECHNIQUES

In this section, we describe the optimization techniques required to scale MLPerf-v0.7 models implemented in both frameworks to the 4096-chip TPU-v3 Multipod machine. Optimization of MLPerf-v0.6 models to a single TPU-v3 pod is presented in (Kumar et al., 2019). To achieve higher scale on the Multipod, we next present novel all-reduce optimizations, aggressive model parallelism and input pipeline optimizations.

3.1 Model Parallelism

In models where data parallelism is limited, we use model parallelism to achieve higher concurrency on TPU-v3 Multipod. We leverage XLA’s Single Program Multiple Data (SPMD) partitioner (Lepikhin et al., 2020) to automatically partition model graphs based on light-weight annotations. In the segmentation models, SSD and MaskRCNN, we implement spatial partitioning by annotating input images. The SPMD partitioner can automatically parallelize computation along the spatial dimensions. These models have relatively large spatial dimensions (8000x1333 for MaskRCNN and 300x300 for SSD). The SPMD partitioner inserts halo exchange communication operations to compute the activations for the next step from spatially partitioned computations. Both of these models enable spatial partitioning along 8 cores to achieve the highest level of concurrency. Communication optimization and elimination of Amdahl bottlenecks via the XLA compiler SPMD approach (Lepikhin et al., 2020) enabled higher concurrency in spatial partitioning. For example, in MaskRCNN the largest batch size is 256, but we were able to parallelize the training on up to 1024 accelerator cores.

In the language models such as the MLPerf transformer benchmark, where the spatial dimensions are small, we explore partitioning the feature dimension as described in (Shazeer et al., 2018), but implemented as annotations for the SPMD partitioner. In this approach, the model weights and activations are split on a tile of the TPU mesh. In the forward pass, partial matrix multiplication operations are computed on each core of the tiled sub-mesh. The activation contributions from each core are reduced via an all-reduce operation on the tiled submesh to execute the next layer of the model. The backward pass has a similar partial matrix multiplication followed by all-reduce producing both activations and gradients. As the weights are also partitioned, the gradients are summed between a partitioned core and its corresponding peer on every other tiled sub-mesh of the TPU machine. Techniques to optimize gradient summation on Multipod are presented in Section 3.3.

3.2 Weight Update Sharding

In traditional data parallelism, model weights are replicated and updated by the optimizer at the end of each training step. However, this computation can become significant when the mini batch size per core is small. For example, we measured in the MLPerf BERT model, the LAMB optimizer weight-update time is about 18% of the step time on 512 TPU-v3 chips. The weight-update-sharding technique (Xu et al., 2020) distributes this computation by first executing a global reduce-scatter after which each accelerator has a shard of summed gradients. This is used to compute a shard of updated weights. In the next step, the shard of updated weights is globally broadcast to update all replicas. To achieve higher speedups we enable weight update sharding in both data and model parallelism. In the segmentation models, where the weights are replicated, the weight-update-sharing scheme is similar to data parallelism. However, when the weights are distributed, we execute multiple concurrent weight-update-sharding steps in each model parallel core across all the replicas.

3.3 Optimized Global Summation

The gradient summation step is critical to achieve strong scaling with MLPerf benchmarks (Mattson et al., 2019). In order to optimize gradient summation on the large TPU-v3 Multipod, we take advantage of the torus wrap links along the Y-dimension. A bidirectional ring is used to execute a reduce-scatter operation along the Y-dimension with the output being a shard of the summed gradients along the Y-ring. Next, a reduce-scatter is executed along the X-dimension. This is followed by a weight update computation with the gradient shard as the input. The updated weights are broadcast first along X and then Y in two steps. Note, in data parallelism, the payload transferred along the X-dimension is 32 times less than the data transferred along the Y-dimension.

In the MLPerf transformer benchmark, we execute distributed matrix multiplication operations by sharding model weights on up to 4 neighboring TPU cores. These cores are placed along a line on the X-dimension. In the forward pass of ML training, all-reduce calls are executed along short rings of X-neighbors. The gradient summation on the Y-dimension stays unchanged as with data parallelism. However, the gradient summation along the X-dimension hops over peers that are model parallelism neighbors. The different ring reductions in the Transformer benchmark are illustrated in Figure 4. In the BERT and transformer models, we also used the brain-float 16-bit floating point precision (bfloat16) to further reduce gradient summation overheads.




Figure 4. Figure shows a 16(mesh) \times 8 (torus) with model parallelism along 4 chips. Three different ring reductions are shown here. i) a black ring reduction for the model parallel forward pass. ii) red rings do the bulk of the gradient reduce scatter along the Y dimension iii) the dotted blue line shows gradient reduction among model peers (only peer id = 0 is shown).

3.4 Distributed Computation of Evaluation Metrics

The train and evaluation computations are executed in a tight loop on the TPU accelerators. The result of the train loop updates model weights in the HBM storage on each TPU accelerator. The updated weights are then used to evaluate the output metrics for the number of epochs specified in the MLPerf rules. In benchmarks where the evaluation batch size is larger than the number of examples in the evaluation dataset, the evaluation dataset is padded with dummy examples. In the TensorFlow implementation, the eval output tensors are used to compute the evaluation metric (for example top-1 accuracy in the Resnet-50 benchmark) on the TPU master host. However, in the JAX implementation the computation of the evaluation quality metric is fully distributed via global summation calls.

3.5 Input pipeline optimizations

One of the challenges of scaling the ResNet-50 model is load-imbalance in the host input pipeline. With the massive scale of a Multipod, some host input pipelines have high overheads of decompressing large JPEG images. Our solution is to store uncompressed images in the host memory so that the host input pipelines execute only i) random crop, ii) random flip and iii) image normalization with a constant mean and variance as specified in the MLPerf reference. This significantly increases the throughput of the host input pipeline allowing it to create a large prefetch buffer. So, when the host pipeline preprocesses a large input image it can feed TPUs with images in the prefetch buffer, thus, eliminating the input pipeline load-imbalance on the Multipod system. This optimization increases the training throughput of ResNet-50 by 35% on a Multipod. With uncompressed images, although the need for memory capacity increases,

the available memory capacity in the system is sufficient. The Multipod has about a thousand CPU host servers and the input is sharded across all these. Using uncompressed images does not incur extra memory throughput overhead, since decompressing images in host memory results in more memory transfers.

For BERT, one of the key techniques to improve convergence is to guarantee randomness and coverage in data shuffling. We find two things very helpful for BERT, using the `tf.data.shuffle` function before the `tf.data.repeat` function at file level, and increasing the shuffle buffer size at sequence level. At file level, proper data shuffling is especially important as the system scale increases, where every host has fewer data files to work with. For example, the 500 files in the BERT reference model will result in a medium-scale system with 128 hosts having only about 4 files per host. Executing a `tf.data.repeat` before `tf.data.shuffle` gives better randomness and coverage of the whole dataset, where the former guarantees the stochasticity, and the latter guarantees the model catches all information available in the dataset. At sequence level, shuffling with small buffer size incurs large run-to-run convergence difference, which originates from the difference of biased training batch at each training iteration, leading to very different convergence trajectories of different runs. With larger buffer sizes, every training batch of different runs can be more uniformly sampled from the whole dataset, which therefore reduces run-to-run difference.

DLRM, like many other recommendation engines, can quickly become input bound as the model accommodates a large per-core batch size while having a small step latency. One key input pipeline optimization for such models is to use host parallel processing to parse data at batch granularity, instead of per-sample. In the case of the dataset used for this model, each training sample is composed of about 40 input features. An additional optimization is to transmit input features over the PCIe bus in a stacked form, reducing the overhead of transmitting many features separately. Finally, batching overhead can be mitigated by shuffling and pre-serializing data in batch form.

4 MODEL OPTIMIZATIONS

In addition to the optimizations mentioned previously, in this section we present the optimizations applied to each MLPerf model. With the exception of MaskRCNN and DLRM, all other models are implemented in both TF and JAX. Note, the JAX implementations use the same scalability and convergence techniques as TF models, resulting in very similar step times as well as number of convergence steps. There are subtle differences w.r.t. TF implementations due to JAX's multi-client design. For example, JAX allows initializing datasets and input pipelines concurrently

in each TPU worker. Also, the global evaluation accuracy is computed via a global all reduce operation on TPUs, compared to TF, where the coordinator CPU process computes the sum after gathering all local accuracy metrics via host RPC calls.

4.1 BERT

BERT (Devlin et al., 2018) with the wikipedia dataset is newly added in MLPerf-v0.7. It is a pre-training task for language understanding with bi-directional transformer architecture. Thanks to the LAMB optimizer (You et al., 2019), BERT can scale very well to large batch sizes, and we are able to use data parallelism at a 4096-chip scale. Scaling BERT in large systems involves optimizations of two aspects, step time and steps to converge.

Other than the optimizations in Section 3, at model level, we optimize the step time of BERT by reducing the stress on architectural bottlenecks, including memory bandwidth, vector units, and registers allowing computation to be executed on the TPU-v3 matrix units with minimal pipeline bottlenecks. To reduce memory bandwidth, we utilize the bfloat16 data type (Wang & Kanwar, 2019) for model activations and gradients aggregation, which largely improves the step time and does not have negative effects on model convergence. To reduce the stress on vector units, we move the scalar multiplications and divisions to the smaller side of matrix multiplication by leveraging the commutativity of scalar multiplication and matrix multiplication. To reduce register spilling, we combine small variables, such as layernorm variables, into one large TensorFlow tensor. This largely reduces the number of variable addresses to store in the registers and therefore speeds up the training step time.

To reduce the steps to converge, we optimize hyperparameters and data shuffling in the input pipeline. First, we use Google Vizier (Golovin et al., 2017) to fine tune the hyperparameters for large batch training, enabled by the scalability of LAMB optimizer. This allows us to leverage maximum data parallelism, which gives better time to accuracy compared to model parallelism. Second, as detailed in Section 3.5, we optimize the way data is shuffled in the data input pipeline in order to ensure the convergence in large systems. This is critical for guaranteeing the stochasticity of the optimizer because large systems with thousands of hosts typically assign less data to each host.

4.2 ResNet-50

ResNet-50 (He et al., 2016) is one of the most widely-used models for ML benchmarking. MLPerf uses the ResNet-50 model with the ImageNet-1K (Russakovsky et al., 2015) dataset as the image classification benchmark. Specifically, MLPerf uses the variant termed “version 1.5” (Goyal et al., 2017) to indicate a slight modification to the original model

architecture which is commonly found in practice. In order to scale the ResNet-50 benchmark to the TPU-v3 Multipod system, we apply optimizations including distributed evaluation, distributed batch normalization, weight update sharding, and gradient summation. The MLPerf-v0.7 reference model uses the LARS optimizer (You et al., 2017) that adaptively scales learning rates, which enables training ResNet-50 with data parallelism on large batch sizes. After the momentum hyperparameters are tuned, we are able to finish training in 88 epochs with batch 65536 on the Multipod.

4.3 Transformer

Transformer represents the state-of-the-art language translation in the MLPerf suite and is one of the two translation models. Trained on the WMT English to German dataset, Transformer uses an attention-based model which differentiates it from the other language model in MLPerf, GNMT. It has been observed that it is hard to scale Transformer with a fixed epoch budget beyond a global batch size threshold given the current dataset (Shallue et al., 2018). Therefore both data and model parallelism are applied to scale the Transformer model to a TPU-v3 Multipod system. With model parallelism, the model is able to run with fewer than batch one per core, using a fixed global batch size of 2048 where the hyperparameters have been well tuned.

SPMD sharding is employed to enable model parallelism. Unlike spatial partitioning (sharding the images) nor Gshard (Lepikhin et al., 2020) (which has sparse components and all-to-all communications), dense sharding is applied to the Transformer model. Shared embedding layers, multi-heads attention projection layers and feed-forward layers are sharded, along with vocab, num_heads, and hidden dimensions, respectively. To speed up gradient all-reduce, 2D cross-replica all-reduce is enabled for SPMD sharding with X-dimension hops over model parallelism neighbor replica. The all-reduce communication is performed in bfloat16 floating point precision to further improve the performance.

4.4 SSD

Single Shot Detection (SSD) is one of two image segmentation models in MLPerf; SSD is intended to reflect a simpler and lower latency model for interactive use cases such as in end-point and non-server situations. Notably, SSD uses a pre-trained ResNet-34 backbone as part of the architecture. The MLPerf SSD benchmark is trained on the COCO dataset (Lin et al., 2014). In the MLPerf-v0.6 submission we had used a global batch size of 2048 and 4-way model parallelism. In this round of MLPerf submissions, we are able to train with a batch size of 4096 using new hyperparameters. Note, this is still much smaller than the batch size of 65536

available in the ResNet-50 model. We used XLA’s SPMD partitioner to enable scaling up to eight TPU cores via model parallelism, replacing XLA’s MPMD spatial partitioner used in MLPerf-v0.6. SPMD has better scalability in compilation time and enabled us to increase the largest scale for SSD training from 2048 TPU-v3 cores in MLPerf-v0.6 to 8192 cores in MLPerf-v0.7. A unique benefit of the SPMD partitioner is that it enables the weight-update-sharding optimization even with model parallelism, that results in a 10% speedup. It is challenging to get high speedups in the SSD model from model parallelism as there are communication overheads from halo exchange and load imbalance as different workers may get uneven tiles of work. In addition, the input image to the SSD model is relatively small (300×300) in the first layer and is further reduced to 1×1 in the last layer. In spite of the above, we are able to get speedups on up to 8 TPU cores used for spatial partitioning.

The SSD implementation with JAX also follows the similar scalability and convergence techniques as its TF counterpart. In addition to the listed differences, the COCO eval execution is slightly different. In TF SSD, the results of the predictions are all brought to the TF coordinator process via host calls, and COCO eval is executed by the TF coordinator process’s CPUs. Since JAX does not have a separate coordinator process, COCO eval is executed on the worker processes in a round robin fashion to improve the load-imbalance, e.g., first worker executes the first COCO eval, the second worker executes the second one, and so on.

4.5 MaskRCNN

Mask-RCNN is the heavy weight object detection benchmark in MLPerf. Besides object detection, Mask-RCNN also performs instance segmentation, which assigns a semantic label as well as an instance index to each pixel in the image. Unlike SSD, which is a one stage detector, Mask-RCNN has two stages: one for proposing instance candidates and the other for fine-tuning the proposals. Also, Mask-RCNN uses a larger image size (800×1333) over SSD (300×300) even though they both train the COCO dataset. Furthermore, Mask-RCNN uses a Resnet-50 backbone plus Feature Pyramid Network contrasted with SSD’s use of Resnet-34.

Scaling MaskRCNN training is challenging as the largest batch size that achieves the MLPerf model quality is quite small. In MLPerf-v0.6 it was trained with a batch size of 128 while in MLPerf-v0.7 we are able to increase the batch size to 256 with new hyperparameters. Data parallelism is used uptill 128 TPU cores and then model parallelism is applied to scale further 1024 TPU cores. The following optimizations are enabled in XLA’s SPMD partitioner to achieve high throughput on TPUs:

- Model parallelism optimized Gather: ROIAlign operation in Mask-RCNN’s is dominated by non contiguous gather operations. These are optimized by one-hot-matmul calls that execute on the TPU matrix unit achieving linear speedups when increasing the number of model parallelism partitions.
- Resharding: the convolutions are split on spatial dimensions. However, when executing the einsum computation, the inputs are re-partitioned on the non-contracting dimension with minimal communication overheads
- Partitioning more ops: there was no XLA partitioner support for some mask-RCNN operations such as top-k, gather and special case convolutions, which can become an Amdahl bottleneck. We added support in the XLA compiler to aggressively split the computation between the TPU cores and increase speedup.
- Communication optimization: there are significant communication overheads between the model parallelism peer cores from resharding, gradient reductions and halo exchange in convolutions. We optimized these in the XLA compiler to reduce communication overheads from 30% to about 10%. Optimizations include minimizing the number of resharding steps, executing a single gradient all-reduce across model cores and replicas and barrier optimizations for the halo exchange.

4.6 DLRM

The Deep Learning Recommendation Model (DLRM) (Nau-mov et al., 2019) is a neural-net based recommendation engine that is newly added in MLPerf-v0.7. The task of the recommender is pCTR click prediction using a combination of categorical and integer features and a single binary label. The MLPerf dataset is the open-source Criteo Terabyte click logs dataset (Ferns, 2015), a corpus of over 4B examples. Embedding table lookups and updates play a significant role in performance due to large unique vocabulary sizes for some of the categorical features. The remainder of the model is dominated by fully connected layers.

We use a global batch of 65536, the maximum with converging hyperparameters, in order to enable scalability. Despite the larger batch size, scalability is still limited for this problem, as step latency is small and communication overheads become a significant portion of runtime. As a result, we do not use a full multi-pod for this model, but rather a fraction of a pod. The input pipeline is also challenging due to the large batch size and small step latency. Section 3.5 discussed relevant input optimizations. Other optimizations used for DLRM include:

- Partition large embedding tables: This is actually necessary to run the model, due to the large memory footprint of the tables. The optimization involves choosing to replicate small tables and partition large ones.

Benchmark	TPU-v3 Chips	TF Runtime (mins.)	Speedup over MLPerf-v0.6	JAX Runtime (mins.)
Resnet-50	4096	0.48	2.67	0.47
BERT	4096	0.39	N/A	0.4
SSD	4096	0.46	2.63	N/A
SSD	2048	0.623	1.94	0.55
Transformer	4096	0.32	2.65	0.26
MaskRCNN	512	8.1	4.4	N/A
DLRM	256	2.4	N/A	N/A

Table 1. End-to-end time achieved with TF 1.x and JAX MLPerf-v0.7 benchmarks on the TPU Multipod machine. The table also presents the speedups achieved over the MLPerf-v0.6 TF submissions.

Benchmark	TPU-v3 TensorFlow (TPU Chips)	TPU-v3 JAX (TPU Chips)
Resnet-50	8.30 (4096)	2.23 (4096)
BERT	17.33 (4096)	3.17 (4096)
SSD	12.87 (4096)	2.03 (2048)
Transformer	14.47 (4096)	4.90 (4096)

Table 2. Initialization time (minutes) comparison between TensorFlow and JAX. Because of the multi-client distributed system of JAX, it shows lower initialization time than TensorFlow, which compiles a multi-device graph on the master host.

- Optimize gather overheads: The model includes a feature self-interaction function that uses a gather to eliminate redundant features. We mask the redundant features with zeros and modify the downstream fully connected layers to ignore the null features during initialization.
- Evaluate multiple steps without host communication: The inference step latency of the model is small enough that the PCIe communication with the host and network gather poses an unacceptable overhead. Instead, we perform multiple inference steps on device and accumulate them.
- Custom evaluation metric library: The evaluation metric is AUC (ROC) on a dataset composed of 90M samples. Popular python libraries scale poorly to this size, requiring 60 seconds per metric computation on a modern workstation. We write a custom C++ CLIF-wrapped implementation that relies on multithreaded sorting and loop fusion to compute the metric in 2 seconds per call.

5 PERFORMANCE ANALYSIS

In this section, we show the performance we are able to achieve with all the framework, infrastructure and model optimizations discussed. Table 1 presents the end-to-end times in six MLPerf-v0.7 benchmarks on Google TPU-v3




Figure 5. Speedup of ResNet-50 with the number of TPU chips.

Multipod. End-to-end time is defined as the time from data touch to computation of the evaluation metric that reaches target quality. Note that the presented optimizations and the larger number of accelerators in the Multipod result in four benchmarks training in under half a minute. The table also presents speedups over the Google TPU submission to the MLPerf-v0.6 benchmarks. Note that for DLRM, the best result of 1.21 minutes was achieved on a TPU-v4 machine. As TPU-v4 is not the focus of this paper and it is not discussed here, the TPU-v3 result (2.4 minutes) is presented in its place.

Table 2 compares the initialization time of TensorFlow and JAX on large-scale systems with 4096 and 2048 TPU chips. TensorFlow’s initialization time ranges from 498 seconds to 1040 seconds, while that of JAX’ is much lower, ranging from 122 seconds to 294 seconds. This is because of their different distributed programming models, which are tradeoffs between performance and usability, as mentioned in Section 2. JAX uses a multi-client distributed system, where each client has one copy of the code and compiles its own graph. This generates constant initialization time for varied sizes of systems. On the other hand, TensorFlow gives one Python process global visibility and control of the whole system, which makes it easier for user code to be




Figure 6. TF ResNet-50’s computation and communication (all-reduce) time on TPUs for executing a single min-batch. Note, the mini-batch size is decreased from 256 to 16 per TPU chip as the scale is increased.




Figure 8. TF BERT’s computation and communication (all-reduce) time on TPUs. Note, mini-batch per TPU chip is 2 at the 4096 chip scale and varies between 4 and 48 at other chip counts.




Figure 7. Speedup of TF BERT with the number of TPU chips.

reflected on the whole workload. But this Python process needs to compile a large multi-device graph for the whole system, which takes longer for larger systems.

Our performance of BERT, ResNet-50, Transformer, MarkRCNN, DLRM and SSD outperform the competitors. Figure 5 shows the end-to-end and throughput speedup of ResNet-50 with the number of TPU chips. It is not surprising that the throughput speedup is closer to ideal scaling than the end-to-end speedup. Note, at batch 64K the MLPerf ResNet-50 model takes 88 epochs to train, while at batch 4K it only needs 44 epochs to train to the target top-1 accuracy of 75.9%.

To further examine the Amdahl’s Law bottleneck on large systems, Figure 6 breakdowns the computation and all-reduce (communication) overhead in the step time, where the blue area shows the computation time, red area shows the communication time, and the sum of the two is the step time on the device. Note that both x- and y-axes are in log scale. With scaling to larger systems, the computation time keeps decreasing while the communication time, i.e., the

all-reduce time, stays almost constant. Using 4096 chips, the all-reduce operations take 22% of the total device step time.

Figure 7 shows the BERT speedup varying the number of TPU chips. BERT shows the highest scalability on systems with 16 to 4096 chips. Figure 8 shows the computation and communication time breakdown for BERT. Compared to ResNet-50, the Amdahl’s Law bottleneck for BERT takes larger percentages, in all scales ranging from 16 to 4096 chips. With the 4096-chip configuration, the all-reduce communication takes 27.3% of the total device step time.

Figure 9 shows speedups via model parallelism in the SSD, MaskRCNN and Transformer benchmarks. Both SSD and MaskRCNN improved the scalability over MLPerf-v0.6 with additional optimizations on the model (e.g., changing gather/scatter to einsum) as well as on the system (e.g., mixing SPMD partitioning with weight update sharding). The scaling is limited by communication overhead introduced for partitioning and inefficiencies from smaller dimensions after partitioning, such as spatial dimensions of later layers. The transformer model also achieves comparable speedup of 2.3× on four TPU-v3 cores. Speedup is limited by significant communication overheads in the all-reduce calls.

Figure 10 compares MLPerf benchmark end-to-end time using the TPU-v3 Multipod and NVIDIA Volta V100 and Ampere A100 GPU systems, reported by Google and NVIDIA to MLPerf-v0.7, respectively. The TPU results are in the “Research/Development/Internal” (RDI) category, while the GPU results are in the “Available On-prem” category. To compare the scalability of TPU and GPU systems, Figure 11 shows the topline end-to-end time speedups in MLPerf-v0.7 benchmarks over 16 accelerator chips of their own types (TPU-v3 chips or GPUs). The techniques presented in this paper enable TPUs to achieve lower end-to-end times and




Figure 9. Speedup via model parallelism in MLPerf-v0.7 TF.




Figure 11. End-to-end time speedups of MLPerf-v0.7 benchmarks over 16 accelerator chips of their own types.




Figure 10. MLPerf-v0.7 end-to-end times in minutes.

higher speedups.

6 CONCLUSION AND DISCUSSION

In this paper, we scaled ML models to the 4k-chip Google TPU-v3 Multipod machine. We invented and refined a number of techniques to achieve the best scale in six MLPerf submissions. The best parallelization varied: data parallelism in the BERT and Resnet-50 models; and model parallelism in the MLPerf SSD, MaskRCNN and Transformer models enabled training at the largest scale. A combination of aggressive compiler optimization of model parallelism, fast gradient summation at scale on the Multipod mesh, distributed evaluation, input pipeline optimizations, and model-specific optimizations contributed to reaching the highest scale. We demonstrated performance in both the TensorFlow and JAX programming frameworks. JAX’s multi-client approach further reduced startup and compilation overheads. We view the current competition in language understanding as a modern-day Space Race, with competing organizations assembling both giant machines and giant models in the quest for an Artificial General Intelligence breakthrough. The techniques in this paper are general and can be applied

to other models, frameworks, and ML accelerator architectures.

We distill the lessons learnt through this cross-stack optimization experience for large-scale ML machines. Firstly, since this machine connects 4096 nodes, a fast all-reduce communication primitive is needed. This can be challenging at a large scale, especially with model parallelism, when enabling the distributed optimizer. Secondly, parallelism mechanisms need to be chosen based on optimizers and models. For example, we showcased data parallelism for ResNet-50 and BERT with large batch optimizers (LARS and LAMB) and a combination of efficient data and model parallelism for SSD, MaskRCNN and Transformer. It needs to be coupled with extensive hyperparameter tuning to optimize the time to convergence. Thirdly, programming frameworks perform the best for different models. We explore and compare both JAX and TensorFlow and each framework excels at different models for higher efficiency and scale. For example, JAX has the best MLPerf-v0.7 results in two benchmarks. Finally, each model has different scalability issues that need to be addressed. For example, the data shuffling optimization for BERT and model parallelism challenges in SSD in Section 4.

7 ACKNOWLEDGEMENT

The authors would like to thank Amit Sabne, Benjamin Lee, Berkin Ilbeyi, Bramandia Ramadhana, Ce Zheng, Chi Chen, Chiachen Chou, David Majnemer, David Chen, Dimitris Vardoulakis, Haoyu Zhang, George Kurian, Jay Shi, Jeff Dean, Jinliang Wei, Jose Baiocchi Paredes, Manasi Joshi, Marcello Maggioni, Peter Gavin, Peter Hawkins, Peter Mattson, Qing Yi, Xiangyu Dong, Yuechao Pan, and Yunxing Dai for their support and feedback.

REFERENCES

- MLPerf results used: 0.7-1,17-56,64-70. MLPerf is a trademark of mlcommons.org. <https://mlcommons.org/>.
- Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., et al. TensorFlow: A system for large-scale machine learning. In *12th {USENIX} symposium on operating systems design and implementation ({OSDI} 16)*, pp. 265–283, 2016.
- Amodei, D., Hernandez, D., SastryJack, G., Brockman, C., and Sutskever, I. AI and compute. *OpenAI Blog*, 2018.
- Bahdanau, D., Cho, K., and Bengio, Y. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*, 2014.
- Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*, 2020.
- Dean, J., Corrado, G., Monga, R., Chen, K., Devin, M., Mao, M., Ranzato, M., Senior, A., Tucker, P., Yang, K., et al. Large scale distributed deep networks. In *Advances in neural information processing systems*, pp. 1223–1231, 2012.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- Ferns, E. Criteo releases industry’s largest-ever dataset for machine learning to academic community. *criteo.com*, 2015.
- Frostig, R., Johnson, M. J., and Leary, C. Compiling machine learning programs via high-level tracing. *Systems for Machine Learning*, 2018.
- Golovin, D., Solnik, B., Moitra, S., Kochanski, G., Karro, J., and Sculley, D. Google Vizier: A service for black-box optimization. In *Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining*, pp. 1487–1495, 2017.
- Goyal, P., Dollár, P., Girshick, R., Noordhuis, P., Wesolowski, L., Kyrola, A., Tulloch, A., Jia, Y., and He, K. Accurate, large minibatch SGD: Training imagenet in 1 hour. *arXiv preprint arXiv:1706.02677*, 2017.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Hochreiter, S. and Schmidhuber, J. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- Jouppi, N. P., Young, C., Patil, N., Patterson, D., Agrawal, G., Bajwa, R., Bates, S., Bhatia, S., Boden, N., Borchers, A., et al. In-datacenter performance analysis of a tensor processing unit. In *Proceedings of the 44th Annual International Symposium on Computer Architecture*, pp. 1–12, 2017.
- Jouppi, N. P., Yoon, D. H., Kurian, G., Li, S., Patil, N., Laudon, J., Young, C., and Patterson, D. A domain-specific supercomputer for training deep neural networks. *Communications of the ACM*, 63(7):67–78, 2020.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. ImageNet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pp. 1097–1105, 2012.
- Kumar, S., Bitorff, V., Chen, D., Chou, C., Hechtman, B., Lee, H., Kumar, N., Mattson, P., Wang, S., Wang, T., et al. Scale MLPerf-0.6 models on Google TPU-v3 pods. *arXiv preprint arXiv:1909.09756*, 2019.
- Langston, J. Microsoft announces new supercomputer, lays out vision for future AI work. *Microsoft Blog*, 2020.
- Lepikhin, D., Lee, H., Xu, Y., Chen, D., Firat, O., Huang, Y., Krikun, M., Shazeer, N., and Chen, Z. Gshard: Scaling giant models with conditional computation and automatic sharding. *arXiv preprint arXiv:2006.16668*, 2020.
- Lin, T.-Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., and Zitnick, C. L. Microsoft COCO: Common objects in context. In *European conference on computer vision*, pp. 740–755. Springer, 2014.
- Luong, M.-T., Sutskever, I., Le, Q. V., Vinyals, O., and Zaremba, W. Addressing the rare word problem in neural machine translation. *arXiv preprint arXiv:1410.8206*, 2014.
- MacroTrends.net. NVIDIA Revenue 2006-2020. 2020.
- Mattson, P., Cheng, C., Coleman, C., Diamos, G., Micikevicius, P., Patterson, D., Tang, H., Wei, G.-Y., Bailis, P., Bittorf, V., et al. MLPerf training benchmark. *arXiv preprint arXiv:1910.01500*, 2019.
- Naumov, M., Mudigere, D., Shi, H.-J. M., Huang, J., Sundaraman, N., Park, J., Wang, X., Gupta, U., Wu, C.-J., Azzolini, A. G., et al. Deep learning recommendation model for personalization and recommendation systems. *arXiv preprint arXiv:1906.00091*, 2019.
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein,

- M., et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3): 211–252, 2015.
- Shallue, C. J., Lee, J., Antognini, J., Sohl-Dickstein, J., Frostig, R., and Dahl, G. E. Measuring the effects of data parallelism on neural network training. *arXiv preprint arXiv:1811.03600*, 2018.
- Shazeer, N., Cheng, Y., Parmar, N., Tran, D., Vaswani, A., Koanantakool, P., Hawkins, P., Lee, H., Hong, M., Young, C., et al. Mesh-TensorFlow: Deep learning for supercomputers. In *Advances in Neural Information Processing Systems*, pp. 10414–10423, 2018.
- TensorFlow.org. XLA: Optimizing compiler for machine learning. 2020. <https://www.tensorflow.org/xla>.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. Attention is all you need. In *Advances in neural information processing systems*, pp. 5998–6008, 2017.
- Wang, S. and Kanwar, P. BFloat16: The secret to high performance on Cloud TPUs. *Google Cloud Blog*, 2019.
- Wu, Y., Schuster, M., Chen, Z., Le, Q. V., Norouzi, M., Macherey, W., Krikun, M., Cao, Y., Gao, Q., Macherey, K., et al. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*, 2016.
- Xu, Y., Lee, H., Chen, D., Choi, H., Hechtman, B., and Wang, S. Automatic cross-replica sharding of weight update in data-parallel training. *arXiv preprint arXiv:2004.13336*, 2020.
- You, Y., Gitman, I., and Ginsburg, B. Large batch training of convolutional networks. *arXiv preprint arXiv:1708.03888*, 2017.
- You, Y., Li, J., Reddi, S., Hseu, J., Kumar, S., Bhojanapalli, S., Song, X. d., Demmel, J., Keutzer, K., and Hsieh, C.-J. Large batch optimization for deep learning: Training BERT in 76 minutes. *arXiv preprint arXiv:1904.00962*, 2019.