Scalable Distributed Training of Large Neural Networks with LBANN

Naoya Maruyama maruyama3@llnl.gov

7th Annual MVAPICH User Group Meeting August 20, 2019







Acknowledgement

- Nikoli Dryden (UIUC / LLNL)
- Yosuke Oyama (Tokyo Tech / LLNL)
- Brian Van Essen (LLNL)
- Tom Benson (LLNL)
- Tim Moon (LLNL)
- LBANN Team





Training Deep Convolutional Neural Networks

- Iterative computations until error rate becomes sufficiently small
- Each iteration is a statically fixed DAG of matrix computations
 - Convolutions
 - Matrix multiplications
 - Element-wise filtering
 - Reductions
- Training sweeps a large collection of labeled samples by picking up a subset of samples ("mini-batch")





while unconverged				
pick a mini-batch				
traverse DAG from input to output				
traverse DAG from output to input				
adjusts network parameters				
endwhile				





LBANN: Livermore Big Artificial Neural Network Toolkit

- <u>https://github.com/LLNL/lbann</u>
- Deep Neural Network training / classification
 - Optimized distributed memory algorithms
 - Including spatially decomposed convolutions
 - Optimized asynchronous communication library
 - Compose parallelism at multiple levels
 - Optimize for strong & weak scaling
- Unique HPC resources at scale
 - InfiniBand or Omnipath interconnect
 - Tightly-coupled GPU accelerators
 - Node-local NVRAM
 - High bandwidth Parallel File System
 - State-of-the art distributed linear algebra library









Parallel Training is Critical to Meet Growing Compute Demand

- Training is an extremely computeintensive task
- And growing exponentially
 - Doubling every 3.5 months
- Distributed training is essential and is proven to be successful

Table 1: Large-scale ResNet-50 training results.							
	Hardware	Chips	Batch	Optimizer	BN	Accuracy	Time
Goyal et al. [6] Smith et al. [16] Akiba et al. [2] Jia et al. [10]	P100 TPU v2 P100 P40	256 128 1024 1024	$\begin{array}{c} 8192 \\ 8192 \rightarrow 16384 \\ 32768 \\ 65536 \end{array}$	Momentum Momentum RMS + Mom. LARS	Local Local Local Local	76.3% 76.1% 74.9% 76.2%	1 hour 30 mins. 15 mins. 8.7 mins.
Baseline Ours Ours Ours Ours	TPU v2 TPU v2 TPU v2 TPU v3 TPU v3	4 256 256 512 1024	1024 16384 32768 32768 32768 32768	Momentum Momentum LARS LARS LARS	Local Local Local Local Distributed	76.3% 75.1% 76.3% 76.4% 76.3%	8.0 hours10 mins.8.5 mins.3.3 mins.2.2 mins.

Ying et al., "Image Classification at Supercomputer Scale," Systems for ML Workshop @ NIPS 2018



Source: OpenAI, https://blog.openai.com/ai-and-compute/









Generalized Parallel Convolution in LBANN



Scaling up Deep Learning for Scientific Data

- The DOE has large scientific data sets that are unlike any commercial data set
 - May not be natural images
 - Large samples
 - May be generated by computational simulations
- Even a mini-batch with just one sample may require more than O(10) GB of memory → Does not fit in a single GPU memory

Mesh Tangling Detection **Cosmological Analysis** 100 100 80 60 O(100)³ volumetric data O(1000)² mesh [Mathuriya18]





10x Better Prediction Accuracy with Large Samples







Scaling Performance beyond Data Parallel Training







Parallel Training

Not parallelizable with standard SGD

Parallelism lies inside the training loop

- Parallel convolution
 - *N*: number of samples, *C*: number of channels, F: number of filters, H: height, W: width
 - Input: images of NxCxHxW
 - Filters: FxCxKxK
 - Output: *NxFxHxW*
 - Partitioning the samples is the most common approach







with some

comm.



Sample-Parallel Convolution

- Split the samples of a mini-batch between processes (or GPUs)
- Each process computes independently with reductions of gradients
- Implemented in many of DL frameworks such as LBANN, TensorFlow, PyTorch, Chainer, etc.



Sample Parallelism







Scalability Limitations of Sample-Parallel Training

Limited memory capacity

- Neural networks are becoming deeper
- Sample size is becoming bigger when dealing with scientific data rather than cats and dogs (^._.^)
- Does not fit GPU memory

Limited parallelism

- The degree of parallelism depends on the number of samples in a mini-batch (i.e., N)
- Can't make N arbitrary large as learning accuracy significantly drops
- Usually, N is O(100)-O(1000)





Goyal et al., Accurate, Large Minibatch SGD: Training ImageNet in 1 Hour, 2017.





Parallelism is not Limited to the Sample Dimension







Boundary Data Exchange

- Convolution/pooling needs adjacent data for each point \rightarrow Halo exchange
- Halo size depends on filter size and number of channels
- Also on stride and dilation
 - Does not depend on grouped convolution
- Pooling involves a "reverse halo exchange"
 - Push/accumulate values to remote











Hybrid Parallel Convolution

- Sample and spatial parallelisms are orthogonal and can be used simultaneously
- Can use $P_N * P_S$ GPUs
- If N > #GPUs && sizeof(activations+weights) < GPU memory size</p>
 - Use sample parallel only
- Otherwise:
 - Spatial parallel to fit GPU memory
 - Sample parallel up to N
- Example
 - Spatial parallel with intra-node 4 GPUs of Sierra nodes
 - Sample parallel with as many nodes as the mini-batch size



An example case with nested partitioning along sample and spatial domains





Spatial Parallelism

Implementation

- Extend the LBANN toolkit (<u>https://github.com/LLNL/lbann</u>)
 - Annotate layers with parallel execution strategies
- Distconv: Custom distributed tensor library for convolution
 - Similar approach as stencil libraries
- Aluminum for communication (<u>https://github.com/LLNL/Aluminum</u>)
 - CUDA peer-to-peer intra-node
 - Aluminum's custom MPI+CUDA backend for inter-node
 - Collective routines with MPI and NCCL
- NVIDIA cuDNN for CNN GPU kernels







Optimizing Halo Exchange

- Standard MPI + CUDA communication model is inefficient for CNNs
 - GPU stream synchronization before MPI_Send is mandatory
 - GPU stream synchronization can be very costly
 - Can't keep issuing asynchronous GPU tasks, causing GPU kernel launch overhead (~10 us)
- Not an issue when doing weak scaling
 - Common in scientific computing
- Only strong scaling is possible with CNNs

Sending data from GPU

<pre>do_something<<<,,,st>>>()</pre>
(cudaMemcpyDeviceToHost)
<pre>cudaStreamSynchronize(st)</pre>
<pre>MPI_Send()</pre>
<pre>do_something<<<,,,,st>>>()</pre>





Asynchronous Inter-Node Data Transfer

- CUDA stream memory operations allow GPU data transfer without GPU synchronization
- No GPU synchronization
- Implemented in the Aluminum communication library

Sendrecv with GPU buffers

Main thread	<pre>do_something<<<,,,st>>>()</pre>					
	(cudaMemcpyDeviceToHost)					
	<pre>cudaEventRecord(ev, st)</pre>					
	cuStreamWaitValue32(st, …)					
	<pre>do_something<<<,,,st>>>()</pre>					
1						
Comm thread	<pre>while (cudaEventQuery(ev)) {}</pre>					
	<pre>MPI_Sendrecv()</pre>					
	(cudaMemcpyHostToDevice)					
	cuStreamWriteValue32(st,)					





Aluminum: GPU-Centric Communication Library

- Generic C++ interface to MPI, NCCL, custom algorithms

 Asynchrony via dedicated communication engine
- Communication is "just another kernel" (GPU-centric, like NCCL)
 - Runtime handles all details (like non-blocking MPI- "progress is magic")
- Associate a "stream of computation" with a communicator
 - CUDA stream; implicit on CPU, but could be a (lightweight) thread, ...
 - Aluminum ensures communication does not begin until data on the stream is ready
- Blocking operations (like MPI_Allreduce):
 - Ensure no subsequent operations start until communication complete
 - Block only the associated stream
- Non-blocking operations (like MPI_Iallreduce):
 - Block no stream
 - Explicit completion operation (like MPI_Wait)
- https://github.com/llnl/aluminum





Aluminum

```
// Synchronous MPI:
for (int step = 0; step < num_steps; ++step) {
    load_mini_batch();
    for (auto&& layer : layers) layer.forward();
    for (auto&& layer : layers) {
      layer.backprop_data();
      if (layer.has_weights()) {
         layer.backprop_filter();
         cudaStreamSynchronize(stream);
      MPI_Allreduce(MPI_IN_PLACE, layer.weights,
         layer.size, MPI_FLOAT, MPI_SUM, comm);
         layer.sgd_step();
      }
    }
}
```











Evaluation

- Lassen 700 nodes
 - 2x POWER9 + 4X V100 + NVLink2 + 2x
 InfiniBand EDR
 - Experiments use up to 2048 GPUs
- cuDNN v7.5.1
- MVAPICH2-GDR 2.3.2
- Spectrum MPI/2019.01.30
- Mesh tangling data:
 - 1K: 1024 x 1024 x 18
 - 2K: 2048 x 2048 x 18
- Same spatial distribution for all layers









Performance of Spatial-Parallel Convolution

Convolution of a single 1024² image with 16 channels and 16 filters of 3x3 kernels







End-to-End Hybrid Parallel Training of Mesh Model

- 2048^2 mesh tangling model
 - Needs 2 GPUs per sample at least
- Weak scaling with sample parallelism
 - Excellent scaling by overlapped gradient allreduces
 - Small increase at 2048 GPUs due to decreased work to hide allreduces
- Strong scaling with spatial parallelism
 - Close to 4x with 16 GPUs
 - Speedup is smaller at later layers due to decreasing spatial dimensions







Conclusion

- Generalized parallel convolution
 - Sample/spatial/channel/filter parallelism
 - Address memory capacity constraint \rightarrow Enables more accurate models
 - Allow strong scaling ightarrow Faster training time
- LBANN: Scalable deep learning software stack for large-scale science and engineering problems
 - https://github.com/llnl/lbann
 - Aluminum: GPU-centric communication library (https://github.com/llnl/aluminum)
 - Implements the generalized parallel convolution algorithms





Ongoing Work

- Channel/filter parallelism
 - Early results to appear at SC19
- Performance modeling to identify optimal parallelization strategies
- Optimizing performance of reading training samples
- References
 - N. Dryden, N. Maruyama, T. Moon, T. Benson, A. Yoo, M. Snir, and B Van Essen, "Aluminum: An Asynchronous, GPU-Aware Communication Library Optimized for Large-Scale Training of Deep Neural Networks on HPC Systems," *Workshop on Machine Learning in HPC Environments (MLHPC'18)*, 2018.
 - N. Dryden, N. Maruyama, T. Benson, T. Moon, M. Snir, and B Van Essen, "Improving Strong-Scaling of CNN Training by Exploiting Finer-Grained Parallelism," *International Parallel and Distributed Processing Symposium (IPDPS'19)*, 2019
 - N. Dryden, N. Maruyama, T. Moon, T. Benson, M. Snir, and B Van Essen, "Channel and Filter Parallelism for Large-Scale CNN Training," *International Conference for High Performance Computing, Networking, Storage, and Analysis (SC'19)*, 2019 (to appear)







Lawrence Livermore National Laboratory

Center for Applied

Scientific Computing

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.