HyperTune: Dynamic Hyperparameter Tuning for Efficient Distribution of DNN Training Over Heterogeneous Systems

Ali HeydariGorji
University of California Irvine
Irvine, California, USA
heydari@uci.edu

Siavash Rezaei
University of California Irvine
Irvine, California, USA
siavashr@uci.edu

Mahdi Torabzadehkashi
NGD Systems Inc.
Irvine, California, USA
mahdi.torabzadeh@ngdsystems.com

Hossein Bobarshad
NGD Systems Inc.
Irvine, California, USA
hossein.bobarshad@ngdsystems.com

Vladimir Alves
NGD Systems Inc.
Irvine, California, USA
vladimir.alves@ngdsystems.com

Pai H. Chou
University of California Irvine
Irvine, California, USA
phchou@uci.edu

ABSTRACT

Distributed training is a novel approach to accelerating training of Deep Neural Networks (DNN), but common training libraries fall short of addressing the distributed nature of heterogeneous processors or interruption by other workloads on the shared processing nodes. This paper describes distributed training of DNN on computational storage devices (CSD), which are NAND flash-based, high-capacity data storage with internal processing engines. A CSD-based distributed architecture incorporates the advantages of federated learning in terms of performance scalability, resiliency, and data privacy by eliminating the unnecessary data movement between the storage device and the host processor. The paper also describes Stannis, a DNN training framework that improves on the shortcomings of existing distributed training frameworks by dynamically tuning the training hyperparameters in heterogeneous systems to maintain the maximum overall processing speed in term of processed images per second and energy efficiency. Experimental results on image classification training benchmarks show up to 3.1x improvement in performance and 2.45x reduction in energy consumption when using Stannis plus CSD compare to the generic systems.

KEYWORDS

DNN training, distributed training, federated learning, dynamic tuning, computational storage devices, heterogeneous systems

ACM Reference Format:

Ali HeydariGorji, Siavash Rezaei, Mahdi Torabzadehkashi, Hossein Bobarshad, Vladimir Alves, and Pai H. Chou. 2020. HyperTune: Dynamic Hyperparameter Tuning for Efficient Distribution of DNN Training Over Heterogeneous Systems. In *IEEE/ACM International Conference on Computer-Aided Design (ICCAD '20)*, November 2–5, 2020, Virtual Event, USA. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3400302.3415699>

1 INTRODUCTION

Nowadays, huge volumes of data are being generated continually due to the exponential growth in the number of data-capturing devices, ranging from industrial to personal IoT devices. Data gain value after being processed, and machine learning (ML) has become a mainstream way of processing the captured data. To be effective, ML models need to be trained with a sufficiently large amount of data. The more data being used for the training phase, the more accurate the ML model can be in the inference phase. However, it often takes days or weeks to train a neural network to an acceptable accuracy level. To accelerate training, distributing the training on clusters of GPUs can be effective, although GPUs tend to be power hungry, making them unsuitable for edge devices or data centers.

In addition to performance and power, data privacy is a rising concern in a distributed system. Data encryption or approximation could help address some privacy issues [7], but these methods are prone to intrusion and hacking as well. Federated learning (FL) [1] is a new approach to sparse training. In FL, each device trains a local replica of the network and shares the parameter updates with a head node and other devices respectively. FL considered to be asynchronous in the way that the nodes work independently and do not wait for other nodes to finish the training. This method is not as efficient as synchronous training such as the GPU example.

Our approach to addressing the above concerns is in-storage processing (ISP), which brings the processing to data rather than moving data to the processing nodes. The benefit of using ISP is that the data does not have to leave the storage system and can be processed in place. As a result, it can cross off the privacy and security concern by keeping the data within the storage system. It also reduces the movements of data, since only the processed data or the updated parameters rather than the raw data are shared with the head node.

To support the ISP concept, we introduce the latest generation of Solid State Drives (SSD), called Computational Storage Devices (CSDs). CSDs are high-capacity, NAND flash-based data storage systems that are equipped with internal processing engines. Going back to the example of distributed processing, all the training procedures can be done on CSDs and only the final parameters that are significantly smaller in size than the raw input data are shared with other nodes.

This work is licensed under a Creative Commons Attribution International 4.0 License.

ICCAD '20, November 2–5, 2020, Virtual Event, USA

© 2020 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8026-3/20/11.

<https://doi.org/10.1145/3400302.3415699>

To efficiently distribute the processing in a heterogeneous environment consisting of the CSDs and the host node, we examined the current distributing libraries such as Tensorflow [2] and Pytorch [17]. Since the processing capability of CSDs are vastly different from a high-end server processor such as Intel Xeon, such libraries fall short of efficiently distributing the training task. To address this problem, we introduce Stannis, a DNN training framework that can efficiently distribute the training process on homogeneous and heterogeneous systems. Stannis evaluates the processing capability of each node and assigns appropriate portion of the task to each individual node. It also monitors the training progress and reconfigures the hyperparameters to optimize the training's progress. Overall, we can summarize our contributions as follow:

- Introducing Laguna, an M.2 form factor CSD platform,
- Deploying the concept of ISP for distributed ML training,
- Proposing a hyperparameter-tuning algorithm to optimize training operations and a load-balancing algorithm that dynamically changes training hyperparameters to mitigate the effects of the interrupting tasks and workloads,
- Proposing a data assignment algorithm that considers data privacy (private data is processed on CSD),
- Evaluating the proposed framework on a cluster of 36 CSD prototypes.

The rest of this paper is organized as follows. Section 2 provides a background and discusses related work on distributed ML training and ISP. Section 3 introduces the proposed distributed training method. Section 4 presents Laguna, our prototype CSD platform. Section 5 presents our evaluation setup and the experimental results. Finally, Section 6 concludes this paper with directions for future work.

2 BACKGROUND AND RELATED WORK

2.1 Distributed Training

Long training times impose a severe limitation on the progress in deep neural network research. These limitations like using big dataset sizes and real-time DNN training have led researchers to emphasis on the problem of DNN training speedup [9, 10]. For inference acceleration, hardware-software co-design [19] as well as automated DNN design algorithms [11] have been studied in the literature. [19] proposes a framework that efficiently implements a neural network on FPGA in a way that requires far less memory than the conventional implementations. This is done by using reinforcement learning to determine an optimal encoding bit-width across network layers. Researchers in [15] show that increasing the batch size does not necessarily decrease the convergence rate, and thus large batch sizes can be used in paralleled stochastic gradient descent training. This helps reduce the communication within the training session. Authors in [5] propose an efficient, adaptive, practical engine named JointDNN for collaborative computation between a mobile device and cloud for DNNs in both inference and training phases. They provide optimization formulations at layer granularity for forward and backward propagation in DNNs. It can adapt to limits on mobile battery, constraints on cloud server load, and quality of service. [13] proposes an algorithm to speed up the training of the recurrent neural networks. This algorithm is based on sequence bucketing and multi-GPU data parallelization. Since

the processing complexity in training is determined by the length of the data sequence in a batch, equalizing the data length helps finish the training at the same time, thus improving the speed. This method has been shown effective with Long Short-Term Memory (LSTM) and Gated Recurrent Units (GRU) training in voice recognition, language modeling, and other perceptual ML tasks. FireCaffe [9] is an example of a framework that accelerates DNN training on a cluster of GPUs by reducing the communication overhead between the nodes. This is done by determining the best network communication hardware that fits the training session and using more efficient communication algorithms such as reduction trees instead of a generic parameter server. Another option is to increase the batch size and change other hyperparameters, which in turn decrease the number of updates and the communication cost. These methods enable FireCaffe to achieve a 36% scaling capability for specific neural networks.

A better alternative to FireCaffe is Horovod [21], which can run on most of the current DNN libraries such as Tensorflow [2] and Pytorch [17]. It uses the NVIDIA NCCL and MPI for the communication. Horovod improves the communication by using the ring allreduce algorithm [3] where each node communicates with only its two neighbor nodes in a circular fashion. It also changes the learning rate after a certain number of epochs have passed to preserve the accuracy with a better convergence rate. By doing so, Horovod can yield up to 88% scaling efficiency on CPUs. However, Horovod is effective on only homogeneous but not heterogeneous systems. We address Horovod's shortcomings by assigning a proper portion of processing to each node and continuously monitoring the system's performance.

2.2 In-Storage Processing

In-storage processing is a very promising solution to overcome the limited processing performance of the available multi-core processors by not only eliminating the data movement, but also adding extra processing power to the servers. We can categorize the available in-storage processing platforms into two main groups: 1) those that time multiplex the available processing units dedicated to the storage controller for the purpose of in-storage processing [14, 16], and 2) those that implant dedicated processing engines for in-storage processing [6, 12, 22]. While the first approach is power- and cost-efficient, it provides limited processing power, and also it can negatively impact the performance of the SSD controller. In the previous works focusing on the second approach, researchers have investigated different types of dedicated processing engines such as FPGAs [14] and embedded processors [23, 24]. Although FPGAs are power efficient, they have several limitations for being used in ISPs, including the challenging step of implementing an RTL design of the tasks [18], the time-consuming step of reconfiguration of FPGAs for running different tasks, and the lack of supporting file systems to allow tasks to deal with the concept of file when accessing the storage. Recently, the new concept of deploying a dedicated processor in the storage unit has been investigated and it has shown promising features. More detailed discussion about this approach is presented in Section 4.

Figure 1: Processing speed vs batch size for MobileNetV2 on Intel® Xeon Silver 4108 CPU

3 METHOD

Most of current DNN training applications are optimized for homogeneous architectures such as a cluster of GPUs where all GPUs have the same processing performance and memory sizes. Applications such as Horovod go further and enable us to use heterogeneous systems such as a mixture of CPUs and GPUs for training. However, a major obstacle associated with these applications is that they use static hyperparameters for the training. It means that the training hyperparameters such as the batch size or the input dataset are determined at the beginning of the process and will remain the same throughout the entire training. Our solution, Stannis, or System for Training of Neural Networks In Storage, is an extension to the current distribution engines [8]. It targets heterogeneous systems such as servers with a master CPU node and clusters of CSDs or multiple servers with different CPUs. Stannis assigns a proper portion of processing to each node, monitors the training performance through a feedback loop, and tunes the hyperparameters to maximize the processing speed while maintaining the accuracy. For now we use Stannis on Horovod and Keras, but it can be paired with other applications as well.

3.1 Batch Size and Dataset Size Tuning

Stannis starts with benchmarking the network to be trained on all distinct processing engines by running a small training session with different batch sizes. This procedure helps us determine the best batch size to achieve the highest processing speed on one node. It also enables us to generate a function to convert the batch size to the processing speed by curve fitting on the pairs of [batch size, speed]. Fig. 1 shows the example of such pairs for MobileNetV2 on an Intel® Xeon® Silver 4108 CPU. As the trend shows in Fig. 1, the processing speed remains the same beyond a certain batch size, which means the operation is getting more computation bound rather than communication bound. Having the speed function for all nodes, we select the most influencing device and maximize its processing speed by choosing the proper batch size. The most influencing device is determined by multiplying a single device's processing speed by the number of such type of devices. Having the batch size and the speed, we can calculate the required time

Figure 2: Stannis's architecture

per step. Since in the synchronized training, the synchronization happens at the end of each step, the elapsed time per step should be the same for all nodes. This helps us determine the best batch size on all other nodes such that all nodes finish at the same time, thus eliminating the idle time caused by rank stall. The portion of the dataset assigned to each node and the number of steps per epoch is calculated as follows:

$$\begin{aligned} Dataset_i &= \frac{BS_i}{\sum_{i=1}^N BS_i} \times Dataset \\ N_{steps} &= \frac{Dataset}{\sum_{i=1}^N BS_i} \end{aligned} \quad (1)$$

where $Dataset_i$ is the portion of input data assigned to node i and BS_i is the batch size on node i and $Dataset$ is the input dataset size. Assuming there are two types of data, private and public, Stannis considers the permission of the data and assigns the private data to the local CSD and the public data to the local CSD and other processing nodes. This guarantees the security of the data by eliminating the need to move sensitive data out of the storage system. Since the input data on one node is shuffled before training, the parameters that are updated as a result of the training on the private data is mixed with parameters of the public data and thus ensures that no private data is transmitted outside of the storage system. This is a fundamental characteristic of Federated Learning [1]. Fig. 2 shows the overall architecture of the system.

3.2 Dynamic Batch and Dataset Allocation

So far, Stannis provides the optimized hyperparameters for maximizing the overall speed and ensures the privacy of the data is preserved. Considering the nature of the CPU, it is possible for the processing nodes to get busy with other tasks that can affect the available processing power to the training session. Since distributed training is a synchronous operation, an interrupt to any node means the entire operation halts until the busy node finishes its portion of the operation. Hence, it is crucial to ensure that all nodes can finish each step at the same time. Stannis tries to compensate for the workload interrupt in nodes by using our implemented function called HyperTune. HyperTune reschedules the portion of operation that is assigned to each node based on the availability of the processing power on that node. This is done by measuring the local

processing speed and available processing power and consequently, updating the batch size list, by either decreasing the batch size on the busy node or increasing it on the other nodes.

Since changing the batch sizes also requires a recalculation for the dataset assignment, Stannis reassigns the dataset based on the new batch size to prevent rank stall on the training session. New batch sizes, data set indexes and lengths are passed to each node using MPIscatter function.

3.3 Monitoring and Decision Making

The initial approach was to monitor the processing speed at the end of each epoch and update the hyperparameters accordingly. This method was sub-optimal since the interruption could happen anywhere within an epoch and could finish before the end of the epoch. Hence, we decided to implement the monitoring session after each step within the epoch. The speed measurements from all nodes are gathered on the one arbitrary node using MPIgather and passed to a decision making function.

To make a better decision, the speed change along with the progress percentage of the current epoch are passed to the function and converted to a decline index based on the following weighted sum:

$$index_i = 0.7 \times \frac{SP - SP_i}{SP} + 0.3 \times \left(\frac{N_{step} - step_i}{N_{step}} \right) \quad (2)$$

SP is the normal speed that is obtained from the *batchsize_to_speed()* function. SP_i and $step_i$ are the current speed and step. and N is the number of steps per epoch. To avoid chattering, a hysteresis-like algorithm is implemented to ignore glitches and mis-measurements in speed. When the decline index exceeds 25%, the step will be flagged as under-utilized and this report will be saved in a separate array. Having five consecutive under-utilization flaggings terminates the current epoch and triggers the *batchsize_controller()* function to determine a new batch size. Our initial approach to determine the new batch size was to use the inverse of the *batchsize_to_speed()* function we got from the tuning section and determine the new batch size based on the current processing speed of the interrupted node. Although this method sounds promising, our evaluations showed non-negligible errors that could worsen the performance. As a second solution, we decided to use a weighted averaging method from the nearest values to the current speed. We use the following equation to calculate the new batch size:

$$BS_i = BS_n \times \frac{SP_i - SP_n}{SP_{n+1} - SP_n} + BS_{n+1} \times \frac{SP_{n+1} - SP_i}{SP_{n+1} - SP_n} \quad (3)$$

where BS_i is the new optimal batch size and SP_i is the current speed that is between SP_n and SP_{n+1} , and BS_n and BS_{n+1} are the batch sizes corresponding to SP_n and SP_{n+1} . A third method that yields almost the same result with no further processing cost is to monitor the CPU usage of the training session on each node. For this approach, we have implemented a sliding window to keep track of the CPU usage for the last 10 steps. The new batch size is proportional to the average of the last five steps (second half of the sliding window) with the declined CPU usage and the normal CPU usage (stored at the first half of the sliding window). The window size should be large enough to ignore possible glitches, but also not too large so it can detect the utilization decline pattern.

Figure 3: Laguna's high-level architecture

Figure 4: Laguna's software architecture

Note that parameters such as the size of the sliding window or the margin for speed decline detection are experimental and can be changed based on the required precision. A benefit of using CPU utilization as a gauge is that the system can also increase the batch size in case there is extra processing power available. When the system frees the occupied processing power, the training session can claim it back by increasing the batch size; whereas the case of speed monitoring cannot detect the availability of extra processing power.

By rapidly monitoring the performance and updating the hyper-parameters, Stannis assigns a larger portion of processing to the free nodes to mitigate the overall performance decline from the interruption. One concern is that by early termination of epoch, the network will lose part of the dataset in the training process and multiple occurrences of such epoch terminations might lead us to completely miss that portion of the dataset. The solution to this concern is to use the shuffle function when creating the input batch of the data. This results in the randomness of the input data list and statistically ensures that all input data will go through the training after a sufficient number of epochs. Another concern is that the training might not converge as a result of constant changing of batch size. To address this concern, we change the batch size in a limited range such that it will not affect the convergence. Careful design of the model's architecture and dynamic selection of hyper parameters such as batch size and learning rate results in better convergence rate and higher achieved accuracy[4]. Dynamic selection of learning rate will be added to the decision-making function in the future work.

Figure 5: Laguna CSD prototype

4 COMPUTATIONAL STORAGE PLATFORM

In this section, we introduce Laguna, a full-featured ASIC-based CSD that enables us to run the ML training experiments in a realistic environment. The proposed and prototyped CSD interacts with the host via the NVMe over PCIe protocol and has 8 TB of flash storage. We first describe the hardware and software architecture of Laguna CSD, including different components and how they communicate with each other. We also demonstrate a fully functional Laguna prototype and provide implementation details and features of the CSD.

4.1 Laguna Hardware Architecture

Fig. 3 shows the high-level architecture of Laguna CSD. The proposed CSD is composed of three main components: front-end, back-end, and in-storage processing (ISP) subsystems. The front-end (FE) and back-end (BE) subsystems are similar to the subsystems in the off-the-shelf flash-based storage devices, i.e., SSDs. In Laguna, a conventional subsystem runs the flash translation layer (FTL) processes that manage flash memory and transfer data between the host and NAND flash modules. Besides this conventional subsystem, Laguna also has a complete ISP engine inside, which runs user applications in-place without sending data to the host. While ISP technology is simple in definition and concept, designing and implementing a practical CSD platform is complicated due to the coexistence and interaction of the conventional components and the innovative ISP engine.

Category	Feature	Description
shared	device form factor	M.2 22110
	memory	4 GB
conventional subsystems (FE and BE)	host interface	PCIe Gen 3.0 x4
	processing cores	3x ARM Cortex-M7 (BE) 1x ARM Cortex-M7 (FE)
ISP engine	host interface	TCP/IP tunnel over NVMe/PCIe
	Flash data link	direct intra-chip link to BE
	processing cores	Quad-core ARM A53 @1.0GHz

Table 1: Laguna prototype features

Figure 6: The 1U server equipped with 36 CSDs

4.1.1 Back-End. For managing the NAND flash modules, BE has three ARM Cortex-M7 processing cores together with a fast-release buffer (FRB), an error-correction unit (ECC), and a memory controller interface unit (MIC). The ARM Cortex M-7 cores run the flash translation layer (FTL) processes, including logical and physical address translation, garbage collection, and wear-leveling. The ECC unit is responsible for handling the errors that happen in the flash memory modules. The fast-release buffer (FRB) transfers data between the memory interface controller (MIC) and other components, while MIC issues low-level I/O commands to NAND flash modules. These modules are organized in 16 channels, so 16 I/O transfers can be performed simultaneously.

4.1.2 Front-End. The FE is responsible for communicating with the host via the NVMe protocol. It receives the I/O commands from the hosts, interprets them, checks the integrity of the commands, and populates internal registers to notify the BE that a new I/O command is received. The FE also de/packetizes the data that is transferred between the host and the CSD. This subsystem consists of a single-core ARM Cortex-M7 processor and an ASIC-based NVMe/PCIe interface.

4.1.3 ISP Engine. Besides the FE and BE subsystems, there is an ISP engine inside Laguna, which is composed of a quad-core ARM Cortex-A53 processor together with a software stack that provides a seamless environment for running user applications. The ISP engine has access to the shared 4 GB memory. A full-fledged Linux OS has been ported to the ISP engine, so the ISP engine supports a vast spectrum of programming languages and models. The ISP engine has a low-latency, power-efficient direct link to the BE subsystem to read and write the flash memory. In other words, the data transferred to the ISP engine bypasses the whole FE subsystem and the power-hungry NVMe/PCIe interface. In Section 5, we show how this data link together with the efficient ISP processing subsystem and the Laguna software stack benefits the ML applications.

4.2 Laguna Software Architecture

Fig. 4 depicts the Laguna software architecture that enables running user applications in-place. From the ISP user’s point of view, it is crucial to have a standard communication link between the host and the ISP engine. The user on the host side can use this communication link to initiate the execution of the application inside the ISP engine and monitor the outcome of the executions. To provide such a standard link, a TCP/IP tunnel through the NVMe/PCIe is supplied inside the Laguna ISP engine. In other words, host-side applications can communicate to the applications running inside the CSD via a

Figure 7: Evaluation of HyperTune

TCP/IP link. This link is also essential for distributed processing applications where processes running on multiple nodes require a TCP/IP link to communicate. To implement this link, we have developed a software layer on both Laguna and the host to de/packetize the TCP/IP data inside the NVMe packets. On the Laguna side, the TCP/IP tunnel software layer runs on the FE subsystem.

On the other hand, the ISP applications should be able to read data from the flash memory in a power-efficient and low-latency manner. The ISP block-device driver (BDD) module is developed as an interface between the BE subsystem and the applications running in the ISP engine. This BDD module makes it possible to mount the flash media on the ISP operating system, so the ISP applications can read/write data from/to flash media similar to when they run on a conventional machine. In other words, the ISP hardware and software architecture is entirely invisible to the applications, so the ISP users do not need to modify the applications to offload them to Laguna. Since both the OS running on the ISP engine and the host's OS can access the flash memory simultaneously, we have used the Oracle cluster filesystem second version (OCFS2) for synchronization between these two operating systems. This approach takes advantage of the TCP/IP tunnel to send the synchronization control data and enforce data integrity.

4.3 Prototype and Features

To demonstrate the feasibility of the Laguna design, we developed a fully functional prototype, as shown in Fig. 5. As previously described, the prototype is an NVMe over PCIe CSD. The categorized features of the CSD comes in Table 1.

To the best of our knowledge, Laguna is the first fully functional M.2 form factor, ASIC-based CSD prototype that can run a vast spectrum of user applications in-place. This CSD provides a complete software stack that makes the ISP implementation details invisible to the user, so the applications can be offloaded to run in Laguna CSD without modification. The user can initiate and monitor the execution of the ISP applications using the TCP/IP tunnel. Additionally, multiple Laguna CSDs can be orchestrated in a distributed processing platform. They can join the host machine in such a distributed environment and seamlessly augment their efficient processing horsepower to the host. Later in this paper, we

use this feature to run ML training on the host and multiple Laguna CSDs in a distributed fashion.

5 EXPERIMENTAL RESULTS

In this section, we present two experiments to evaluate different features of Stannis and Laguna. In the first experiment, we evaluate the performance of the Stannis on three similar nodes in term of processing capabilities. In the second experiment, we put both Stannis and CSDs in use and run multiple training sessions with different configurations to investigate the performance improvement in different situations.

5.1 HyperTune Evaluation

To evaluate Stannis and its HyperTune function, we run a training session on three similar nodes in the term of processing power. The reason for choosing similar processing power is that the significance of HyperTune can be shown better when the nodes are equally important in the term of processing power. In this test, each node is an AIC 2U-FB201-LX server and is equipped with an 8-core, 16-thread Intel® Xeon® Silver 4108 CPU with 64 GB of DRAM. We choose the MobileNetV2 neural network to run the test on, with 300,000 images as the input for the system. To simulate external workloads, we use the Gzip compression application, which enables us to occupy the desired number of cores on the processors. We decided to interrupt one node at a time to make the evaluation less complicated. In theory, interrupting multiple nodes is not different from interrupting a single node since the master node collects the interrupt reports from all nodes and can easily detect if the slowdown is due to the local workload or an external node.

Stannis starts with finding the best batch size for each node, which is 180. Then, in two separate steps on separate nodes, we set the Gzip to occupy 4 and 6 cores out of 8 cores and monitor the speed change. We define the performance as the average number of processed images per second for the training. Fig. 7 shows the overall processing speed of the three nodes during the test. As can be seen, the total processing speed is 93.4 images/second in normal operation mode. As the workload increases, the speed drops to 75.6 img/sec for 50% external workload and 53.3 img/sec for 75% external workloads. These speeds remain nearly constant as long as

the workload sustains thereafter. In contrast, if HyperTune detects interrupts by new workload and determines non-transient decline in the speed, it recomputes the new batch size for the busy node and updates the hyperparameters, which are 140 and 100 for 4-core and 6-core workloads, respectively. By changing the batch size, the other two nodes regain the initial processing speed, which shows that Equation (3) works fine in determining the new batch size. The processing speed declines respectively to 83.7 and 85.8 img/sec for 75% and 50% external workloads. That is, HyperTune is 14% and 57% faster than the baseline at no cost of power or accuracy and with negligible computation.

5.2 Stannis and CSD Evaluation

To evaluate Stannis on CSD, we run a training session on a custom-built FlacheSAN1N36M-UN server equipped with the same Intel Xeon Silver 4108 CPU and 64 GB of DRAM. We put 36 Laguna CSDs, each with 8 TB capacity on the server to make a 1U class server with 288 TB storage capacity, as shown in Fig. 6. To measure the power and energy consumption, we connect an HPM-100A power meter between the server and the power source and measure the total power consumed by the server, including the cooling system and the storage devices. The power meter logs power at 1 Hz. Since the performance and the computation load depend on the network type and size, we run our test for two different networks of different sizes, MobilenetV2 [20] and ShuffleNet [25]. The former is a neural network for image classification, feature extraction, and segmentation with 3.4 M parameters and 300 M multiply-accumulate (MAC) operations. The latter is a high-accuracy neural network with 5.4 M parameters and 524 M MACs, where the relatively low number of FLOPS makes it suitable for edge devices.

Just like the previous experiment, we run multiple training sessions with different numbers of CSDs and interrupt the master node by taking 6 out of 8 cores from the training session. We only interrupt the host processor since the CSDs provide an augmented processing power and are not counted as the main processor. The initial benchmarking results give the batch size of 180 and 15 for the host and Laguna CSDs for MobilenetV2 and 300 and 25 for the ShuffleNet. Fig. 8a and Fig. 8b show the processing speed based on the number of cards in default operation mode, with external workload when HyperTune is not engaged, and when HyperTune is enabled. The results show that for MobileNet training, the processing speed improves from 33.4 img/sec to 99.83 img/sec when the training is distributed over the host and 36 CSDs, which is 3.1x improvement. When the external workloads occupy 6 out of 8 cores, the processing speed declines from 99.83 img/sec to 49.26 img/sec. However, by engaging HyperTune, the processing speed declines from 99.83 img/sec to 74.89 img/sec. This shows that the HyperTune improved the processing speed by about 1.5x while having no significant effect on the training accuracy and loss numbers or the power consumption. Same analysis for ShuffleNet gives us 2.82x and 1.45x performance improvement.

We also compare the energy consumption with and without the CSD. We estimated the overall energy consumption in terms of joules per image (J/img) by integrating the power consumption over time for the entire epoch and divide it by the number of processed images per epoch. For the MobileNetV2 the energy consumption

Figure 8: Benchmark results

is 1.32 J/img for the host only vs. 0.54 J/img with 36 CSDs added to the process. It means that the entire system consumes 2.45x less energy by training on the CSDs compared to a standalone CPU.

6 CONCLUSIONS

In this paper, we introduced a hardware and software combination to tackle the problems associated with the distributed training of neural networks over heterogeneous systems. Our hardware named Laguna, a NAND flash based computational storage device, can accelerate DNN training while reducing energy consumption with the privacy and security benefits of federated learning. However, much of the potential gains in energy and performance are diminished in straightforward schemes especially on heterogeneous nodes unless optimized dynamically. To address this challenge, our software framework named Stannis optimizes the training hyperparameters dynamically by monitoring workload variations. Stannis is complementary to current DNN training engines by optimizing training at the distribution level. Our experimental results show that Stannis along with Laguna can increase the processing speed by up to 3.1x while reduces the energy consumption by 2.45x at no cost or loss of accuracy. For future work, we plan to extend hyperparameter tuning to include learning rate to increase the accuracy while providing more support for federated learning.

7 ACKNOWLEDGEMENTS

This work is supported by the US National Science Foundation (NSF) grant 1660071.

REFERENCES

- [1] [n.d.]. Google AI Blog: Federated Learning. <https://ai.googleblog.com/2017/04/federated-learning-collaborative.html>
- [2] Abadi et al. 2016. Tensorflow: A system for large-scale machine learning. In *12th USENIX Symposium on Operating Systems Design and Implementation OSDI*. 265–283.
- [3] Alekh Agarwal et al. 2014. A Reliable Effective Terascale Linear Learning System. *Journal of Machine Learning Research* 15 (2014), 1111–1133.
- [4] Reza Asadi and Amelia C Regan. 2020. A spatio-temporal decomposition based deep neural network for time series forecasting. *Applied Soft Computing* 87 (2020), 105963.
- [5] Amir Erfan Eshratifar, Mohammad Saeed Abrishami, and Massoud Pedram. 2019. JointDNN: an efficient training and inference engine for intelligent mobile cloud computing services. *IEEE Transactions on Mobile Computing* (2019).
- [6] Boncheol Gu et al. 2016. Biscuit: A framework for near-data processing of big data workloads. In *ACM SIGARCH Computer Architecture News*, Vol. 44. IEEE Press, 153–165.
- [7] Ehsan Hesamifard, Hassan Takabi, and Mehdi Ghasemi. 2019. Deep Neural Networks Classification over Encrypted Data. In *Proceedings of the Ninth ACM Conference on Data and Application Security and Privacy* (Richardson, Texas, USA) (CODASPY '19). Association for Computing Machinery, New York, NY, USA, 97–108. <https://doi.org/10.1145/3292006.3300044>
- [8] Ali HeydariGorji, Mahdi Torabzadehkashi, Siavash Rezaei, Hossein Bobarshad, Vladimir Alves, and Pai H. Chou. 2020. STANNIS: Low-Power Acceleration of Deep Neural Network Training Using Computational Storage. *arXiv:cs.DC/2002.07215*
- [9] Forrest N. Iandola et al. 2016. Firecaffe: near-linear acceleration of deep neural network training on compute clusters. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2592–2600.
- [10] Sergey Ioffe and Christian Szegedy. 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167* (2015).
- [11] Mojan Javaheripi, Mohammad Samragh, Tara Javidi, and Farinaz Koushanfar. 2020. GeneCAI: Genetic Evolution for Acquiring Compact AI. *arXiv preprint arXiv:2004.04249* (2020).
- [12] Sang-Woo Jun et al. 2015. Bluedbm: An appliance for big data analytics. In *2015 ACM/IEEE 42nd Annual International Symposium on Computer Architecture (ISCA)*. IEEE, 1–13.
- [13] Viacheslav Khomenko et al. 2016. Accelerating recurrent neural network training using sequence bucketing and multi-GPU data parallelization. In *2016 IEEE First International Conference on Data Stream Mining & Processing (DSMP)*. IEEE, 100–103.
- [14] Sungchan Kim et al. 2016. In-storage processing of database scans and joins. *Information Sciences* 327 (2016), 183–200.
- [15] Mu Li et al. 2014. Efficient mini-batch training for stochastic optimization. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. 661–670.
- [16] Dongchul Park et al. 2016. In-storage computing for Hadoop MapReduce framework: Challenges and possibilities. *IEEE Trans. Comput.* (2016).
- [17] Adam Paszke et al. 2017. Automatic Differentiation in PyTorch. In *NIPS Autodiff Workshop*.
- [18] S. Rezaei et al. 2019. UltraShare: FPGA-based Dynamic Accelerator Sharing and Allocation. In *International Conference on ReConfigurable Computing and FPGAs (ReConFig)*. 1–5.
- [19] Mohammad Samragh, Mojan Javaheripi, and Farinaz Koushanfar. 2019. CodeX: Bit-flexible encoding for streaming-based FPGA acceleration of DNNs. *arXiv preprint arXiv:1901.05582* (2019).
- [20] Mark Sandler et al. 2018. Inverted Residuals and Linear Bottlenecks: Mobile Networks for Classification, Detection and Segmentation. *CoRR* abs/1801.04381 (2018). *arXiv:1801.04381* <http://arxiv.org/abs/1801.04381>
- [21] Alexander Sergeev and Mike Del Balso. 2018. Horovod: fast and easy distributed deep learning in TensorFlow. *arXiv preprint arXiv:1802.05799* (2018).
- [22] Mahdi Torabzadehkashi et al. 2019. Catalina: in-storage processing acceleration for scalable big data analytics. In *27th Euromicro International Conference on Parallel, Distributed and Network-Based Processing (PDP)*. IEEE, 430–437.
- [23] Mahdi Torabzadehkashi et al. 2019. Computational storage: an efficient and scalable platform for big data and HPC applications. *Journal of Big Data* 6, 1 (2019), 100.
- [24] M. Torabzadehkashi, A. Heydarigorji, S. Rezaei, H. Bobarshad, V. Alves, and N. Bagherzadeh. 2019. Accelerating HPC Applications Using Computational Storage Devices. In *2019 IEEE 21st International Conference on High Performance Computing and Communications; IEEE 17th International Conference on Smart City; IEEE 5th International Conference on Data Science and Systems (HPCC/SmartCity/DSS)*. 1878–1885.
- [25] Xiangyu Zhang et al. 2017. ShuffleNet: An Extremely Efficient Convolutional Neural Network for Mobile Devices. *CoRR* abs/1707.01083 (2017). *arXiv:1707.01083*