PAPER Special Section on Parallel, Distributed, and Reconfigurable Computing, and Networking

An FPGA-Based Optimizer Design for Distributed Deep Learning with Multiple GPUs*

Tomoya ITSUBO[†a)], ***Nonmember***, **Michihiro KOIBUCHI**[††], ***Senior Member***, **Hideharu AMANO**[†], ***Fellow***, ***and*** **Hiroki MATSUTANI**[†], ***Member***

SUMMARY Since deep learning workloads perform a large number of matrix operations on training data, GPUs (Graphics Processing Units) are efficient especially for the training phase. A cluster of computers each of which equips multiple GPUs can significantly accelerate the deep learning workloads. More specifically, a back-propagation algorithm following a gradient descent approach is used for the training. Although the gradient computation is still a major bottleneck of the training, gradient aggregation and optimization impose both communication and computation overheads, which should also be reduced for further shortening the training time. To address this issue, in this paper, multiple GPUs are interconnected with a PCI Express (PCIe) over 10Gbit Ethernet (10GbE) technology. Since these remote GPUs are interconnected with network switches, gradient aggregation and optimizers (e.g., SGD, AdaGrad, Adam, and SMORMS3) are offloaded to FPGA-based 10GbE switches between remote GPUs; thus, the gradient aggregation and parameter optimization are completed in the network. The proposed FPGA-based 10GbE switches with the four optimizers are implemented on NetFPGA-SUME board. Their resource utilizations are increased by PEs for the optimizers, and they consume up to 56% of the resources. Evaluation results using four remote GPUs connected via the proposed FPGA-based switch demonstrate that these optimizers are accelerated by up to 3.0x and 1.25x compared to CPU and GPU implementations, respectively. Also, the gradient aggregation throughput by the FPGA-based switch achieves up to 98.3% of the 10GbE line rate.

key words: *deep learning, FPGA switch, remote GPU*

1. Introduction

A gradient descent optimization algorithm with a back-propagation algorithm is used for training deep neural networks. It iteratively computes gradients of a loss function and optimizes weight parameters of the neural networks. The training phase performs a large number of matrix operations, and thus the computation cost is extremely high. GPUs (Graphics Processing Units) are typically used for accelerating the gradient computation. Actually, a cluster of computers each of which equips multiple GPUs has been used for reducing the training time [1]–[3].

In the parallel and distributed deep learning [4], in addition to the gradient computation parallelized by multiple GPUs, their gradients are aggregated and the weight parameters are then optimized based on the gradients by using a parameter optimization algorithm. Although the gradient computation is still a major computation bottleneck of the training, the gradient aggregation and parameter optimization impose both communication and computation overheads, which should also be reduced for further shortening the training time.

In this paper, we focus on the gradient aggregation and parameter optimization and accelerate them by using FPGA-based 10Gbit Ethernet (10GbE) switches. More specifically, multiple GPUs are interconnected with a PCI Express (PCIe) over 10GbE technology [5]. Since these remote GPUs are interconnected via network switches, the gradient aggregation and optimization algorithms are offloaded to the FPGA-based 10GbE switches between these remote GPUs. In this case, the gradient aggregation and parameter optimization are completed in the network. We implement the gradient aggregation and four optimization algorithms (i.e., SGD, Adagrad, Adam, and SMORMS3) on NetFPGA-SUME card [6] that has four 10GbE interfaces. The proposed FPGA-based switch is evaluated in terms of the gradient aggregation performance, parameter optimization performance, and FPGA resource utilization [**].

The rest of this paper is organized as follows. Section 2 introduces the parameter optimization algorithms, parallel and distributed deep learning approaches, and remote GPU technologies over Ethernet. Section 3 proposes the network switch that connects remote GPUs and performs the gradient aggregation and parameter optimization. Section 4 describes the packet trace used in the experiment and Sect. 5 describes the implementation. Section 6 shows the evaluation results of performance and resource utilization. Section 7 concludes this paper.

2. Related Work

2.1 Parallel and Distributed Training Models

As a distributed deep learning approach, this paper employs a synchronous data parallel model [4] that combines data parallel training and synchronous parameter optimization. In the data parallel model, training data is divided and assigned to GPU workers, so that the training phase is per-

Manuscript received January 7, 2021.
Manuscript revised April 29, 2021.
Manuscript publicized July 1, 2021.

[†]The authors are with Graduate School of Science and Technology, Keio University, Yokohama-shi, 223–8522 Japan.
[††]The author is with National Institute of Informatics, Tokyo, 101–8430 Japan.

*This work was partially supported by JSPS KAKENHI Grant Number JP19H01106.

a) E-mail: itsubo@arc.ics.keio.ac.jp

DOI: 10.1587/transinf.2021PAP0008

**This paper is an extended version of our conference paper [7] by evaluating resource utilization of the entire 10GbE switch and revising the description.

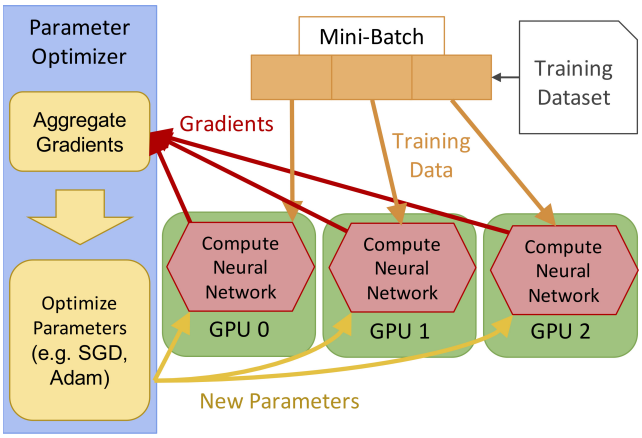

Fig. 1 Training with synchronous data parallel model

Algorithm 1 Algorithm of SGD

Require: lr: Learning rate
Require: θ_0: Initial parameter
Require: $f(\theta)$: Loss function with parameter θ
Require: $g(\theta)$: Gradient of loss function with parameter θ
$t \leftarrow 0$
while Exit condition is not satisfied **do**
  $t \leftarrow t + 1$
  $g(\theta_t) \leftarrow \nabla_{\theta_t} f(\theta_t)$
  $\theta_t \leftarrow \theta_{t-1} - lr \cdot g(\theta_t)$
end while

formed in parallel by sharing intermediate training results. Then, the gradients separately computed by GPU workers are combined so as to optimize weight parameters of a model. Since a synchronous model is assumed, all the GPU workers are synchronized when they finish the gradient computation of errors for their assigned mini-batch. Then, the gradients are aggregated and weight parameters are optimized.

Figure 1 illustrates an outline of a training phase of the synchronous data parallel model. Below is the procedure using multiple GPU workers.

1. A certain amount of training data is picked up as a mini-batch.
2. The mini-batch is assigned to each GPU worker so that it processes the assigned training data with a neural network.
3. Each GPU worker computes error gradients, and they are aggregated in a host machine.
4. Weight parameters are optimized with the aggregated gradients, and then the optimized parameters are distributed to all the GPU workers.

2.2 Parameter Optimization Algorithms

A gradient method is used for the parameter optimization of neural networks so that an error calculated by a loss function is minimized. There are various algorithms to compute new weight parameters based on gradients of the loss function and current parameters. A well-known algorithm is SGD (Stochastic Gradient Descent) listed in Algorithm 1. In SGD, gradients multiplied by a learning rate are subtracted from current weight parameters to compute new parameters.

Algorithm 2 Algorithm of Adagrad

Require: lr: Initial learning rate
Require: θ_0: Initial parameter
Require: $f(\theta)$: Loss function with parameter θ
Require: $g(\theta)$: Gradient of loss function with parameter θ
Require: ε: Small constant
$t \leftarrow 0$
$h \leftarrow 0$
while Exit condition is not satisfied **do**
  $t \leftarrow t + 1$
  $g(\theta_t) \leftarrow \nabla_{\theta_t} f(\theta_t)$
  $h_t \leftarrow h_{t-1} + g^2(\theta_t)$
  $lr_t \leftarrow \frac{lr}{\sqrt{h_t}+\varepsilon}$
  $\theta_t \leftarrow \theta_{t-1} - lr_t \cdot g(\theta_t)$
end while

Algorithm 3 Algorithm of Adam

Require: lr: Learning rate
Require: θ_0: Initial parameter
Require: $f(\theta)$: Loss function with parameter θ
Require: $g(\theta)$: Gradient of loss function with parameter θ
Require: ε: Small constant
Require: $\beta_1, \beta_2 \in [0, 1)$:
$t \leftarrow 0$
$m_0 \leftarrow 0$
$v_0 \leftarrow 0$
while Exit condition is not satisfied **do**
  $t \leftarrow t + 1$
  $g(\theta_t) \leftarrow \nabla_{\theta_t} f(\theta_t)$
  $m_t \leftarrow \beta_1 \cdot m_{t-1} + (1 - \beta_1) \cdot g(\theta_t)$
  $v_t \leftarrow \beta_2 \cdot v_{t-1} + (1 - \beta_2) \cdot g^2(\theta_t)$
  $\widehat{m_t} = \frac{m_t}{1-\beta_1^t}$
  $\widehat{v_t} = \frac{v_t}{1-\beta_2^t}$
  $\theta_t \leftarrow \theta_{t-1} - lr \cdot \frac{\widehat{m_t}}{\sqrt{\widehat{v_t}}+\varepsilon}$
end while

Algorithm 4 Algorithm of SMORMS3

Require: lr: Learning rate
Require: θ_0: Initial parameter
Require: $f(\theta)$: Loss function with parameter θ
Require: $g(\theta)$: Gradient of loss function with parameter θ
Require: ε: Small constant
$t \leftarrow 0$
$m_0 \leftarrow 0$
$v_0 \leftarrow 0$
$s_0 \leftarrow 1$
while Exit condition is not satisfied **do**
  $t \leftarrow t + 1$
  $g(\theta_t) \leftarrow \nabla_{\theta_t} f(\theta_t)$
  $s_t \leftarrow 1 + (1 - x_{t-1}) \cdot s_{t-1}$
  $\rho_t \leftarrow \frac{1}{s_t+1}$
  $m_t \leftarrow (1 - \rho_t) \cdot m_{t-1} + \rho_t \cdot g(\theta_t)$
  $v_t \leftarrow (1 - \rho_t) \cdot v_{t-1} + \rho_t \cdot g^2(\theta_t)$
  $x_t \leftarrow \frac{m_t^2}{v_t+\varepsilon}$
  $\theta_t \leftarrow \theta_{t-1} - \frac{\min\{lr, x_t\}}{\sqrt{v_t}+\varepsilon} \cdot g(\theta_t)$
end while

Although SGD is simple, it may converge to a local solution with low accuracy or require a number of iterations to convergence, depending on a selected learning rate. A careful tuning is thus required for selecting the hyperparameters.

To improve SGD in terms of accuracy and the number of iterations to convergence, variants of SGD have been invented. Typical examples of such algorithms are Adagrad [8], Adam [9], and SMORMS3 [10]. They are listed in Algorithms 2, 3, and 4. They are widely used to optimize a model with higher accuracy and smaller number of iterations to convergence. Since computation costs for these algorithms are higher than that of SGD, although they can reduce the number of iterations, they incur a longer computation time for each iteration.

2.3 Distributed Deep Learning Using GPUs

A large-scale distributed deep learning has been efficiently executed on GPU clusters. In a GPU cluster, nodes consist of a host machine with several GPUs and they are interconnected with a high-speed network, such as 10GbE. In [1], a GPU cluster consisting of 128 nodes with 1,024 GPUs completed a training phase of ImageNet in 15 minutes. In [2] and [3], GPU clusters using 4,352 GPUs and 2,048 GPUs completed the training phase of ImageNet in 122 seconds and 74.7 seconds, respectively. They use MPI AllReduce for communication between the nodes. Although the communication is efficiently done by using Allreduce, the communication overhead increases as the number of nodes increases, and the execution time per iteration increases. In [1], approximately 15% of execution time is spent for communication at 128 nodes.

In [11], it is reported that a significant portion of training time is spent for the communication. To reduce the communication overhead, nodes in a cluster circulate their gradients to aggregate them within a cluster without a specific aggregation node. Also, since the gradients are more tolerant of accuracy loss than weight parameters, a compression technique is applied to the gradients to reduce the communication overhead.

2.4 Network-Attached GPUs

Since the number of PCIe (PCI Express) slots in a single machine is limited, the number of required host machines is increased as the number of GPUs increases, resulting in a higher cost and power consumption. To mitigate this limitation, GPUs can be connected to network switches and accessed via a high-speed network remotely by using PCIe over Ethernet technology.

As a PCIe over 10GbE (10Gbit Ethernet) technology, in this paper we use ExpEther 10G [5] that can extend PCIe over 10GbE. PCIe devices in a 10GbE network can be assigned to the host machines as shown in Fig. 2. Such network-attached GPUs have been used to accelerate a large-scale graph processing [12]. In this paper, we employ network-attached GPUs for accelerating the gradient computation, while the gradient aggregation and parameter optimization are accelerated by using network-attached FPGA.

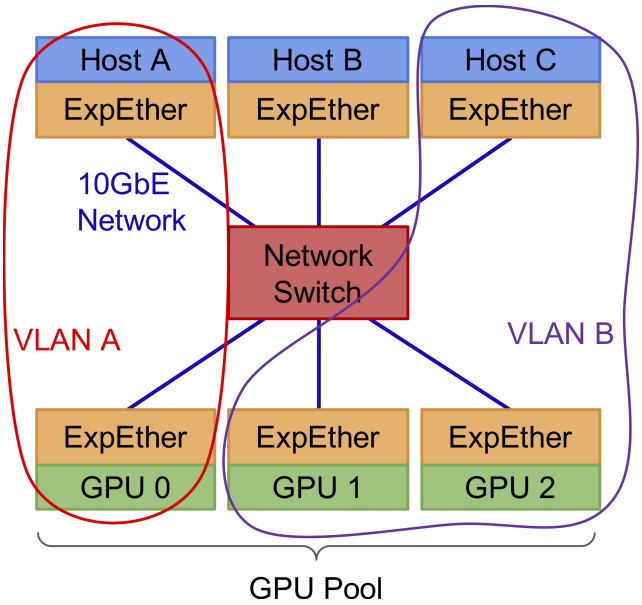

Fig. 2 Connection between hosts and GPUs using ExpEther [5]

3. FPGA-Based Switch Design

3.1 Preliminary Evaluations

This section proposes an FPGA-based acceleration of the gradient aggregation and parameter optimization in distributed deep learning. First, a preliminary evaluation is conducted to show execution times for the aggregation and parameter optimization in an overall training phase. To measure the execution time, Chainer [13] is used as a deep learning framework. Table 1 shows the preliminary evaluation environment. GoogleNet is used as the DNN model, and Adam is used as a parameter optimization algorithm. A synchronous data parallel model is used in this evaluation.

Figure 3 shows a breakdown of the execution time. The execution time for each iteration is divided into six parts: data preparation, forward propagation, back propagation, gradient aggregation, parameter optimization, and parameter distribution. Gradient computation (i.e., forward propagation and back propagation) are executed by four GPUs, while parameter optimization using Adam is done by a single GPU. The other parts, such as gradient aggregation, are done by a host CPU.

As shown in Fig. 3, the execution times for the compute-intensive parts, such as forward propagation and back propagation, are large. Those for the gradient aggregation and parameter optimization account for approximately 8.1% even with four GPUs. Although they are not a major bottleneck in the case of four GPUs, their execution times would be increased as the number of GPUs increases. Thus, their execution times should be reduced for further shortening the training time.

Table 1 Preliminary evaluation environment

CPU	Intel Core i7-6850K @3.6GHz
Memory	32GB
GPU	NVIDIA Geforce GTX 1080 (8GB RAM) x4
CUDA	version 10.0
Chainer	version 6.2.0

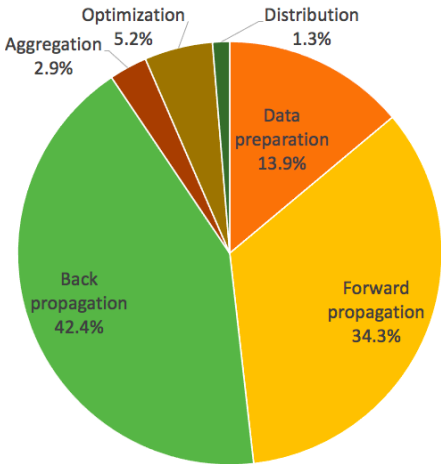

Fig. 3 Breakdown of execution times in training phase

3.2 System Overview

To further reduce the execution times of distributed deep learning, in this paper we propose to use network-attached GPUs for the gradient computation, and offload the gradient aggregation and parameter optimization to network-attached FPGA in a network.

When a host machine accesses a remote GPU, input data and output data to/from the GPU go through one or more network switches. A high speed data processing can be performed during a communication between the host machine and the GPU by incorporating in the network switch. In this paper, we propose to offload the gradient aggregation and parameter optimization to an FPGA-based 10GbE switch. Conventionally, this gradient aggregation and parameter optimization are processed by CPU. Figure 4 illustrates the proposed system, where remote GPUs using PCIe over 10GbE are interconnected via an FPGA-based network switch. The gradient aggregation and parameter optimization are thus completed in the middle of communication with low overheads.

The PCIe over 10GbE used in this paper is transparent to software layer, which means that remote GPUs connected via the PCIe over 10GbE can be accessed as local GPUs directly attached to PCIe slots. The proposed FPGA-based switch performs aggregation or parameter optimization depending on incoming packet information, which means that the software layer does not explicitly control the proposed FPGA-based switch. Since the aggregation and parameter optimization are done by the proposed FPGA-based switch, the software layer should be modified not to execute these

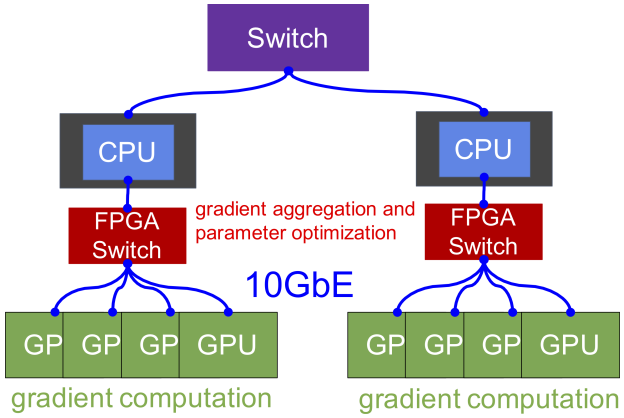

Fig. 4 Proposed system (gradient computation is done by "GPUs" and gradient aggregation and parameter optimization are done by "Switches")

operations though our approach has not been implemented in well-known DNN frameworks.

Since the 10GbE bandwidth is narrower than PCIe Gen3 x16, data transfer time may increase in the proposed system. This issue can be mitigated by using recent ExpEther 40G product. Actually, offloading the gradient aggregation and parameter optimization to a network switch can compensate for the lower bandwidth and efficiently improve their execution times, as shown in Sect. 6.

In our implementation, packets are encapsulated as 10GbE frames, so 10GbE-based layer-2 switches can be inserted between remote GPUs and host machine. However, layer-3 packet routing cannot be supported and thus layer-3 routing devices cannot be inserted between them.

3.3 Gradient Aggregation Function

Here, data flow of the gradient aggregation in the proposed system is described below. As shown in Fig. 4, GPU workers compute gradients and then send them to a host machine (denoted as "CPU") via the FPGA-based network switch (denoted as "Switch"). The FPGA-based switch extracts the gradients from ExpEther packets sent from GPUs to a host. Then the gradients are aggregated and sent to the host. The host machine thus receives already-aggregated gradients from the switch.

In large-scale distributed deep learning, a single host machine will be a bottleneck, and thus multiple host machines should be used as shown in Fig. 4. In this case, these nodes perform AllReduce to exchange their already-aggregated gradients. Even with such multiple host cases, gradient aggregation overheads are greatly reduced by the proposed network switch.

3.4 Parameter Optimization Function

The parameter optimization is performed in the proposed switch after all the gradients are aggregated in the network. As parameter optimization algorithms, SGD, Adagrad, Adam, and SMORMS3 are implemented on the

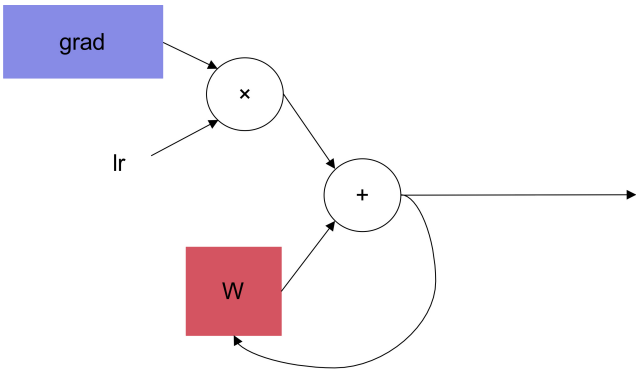

Fig. 5 Computation graph of SGD

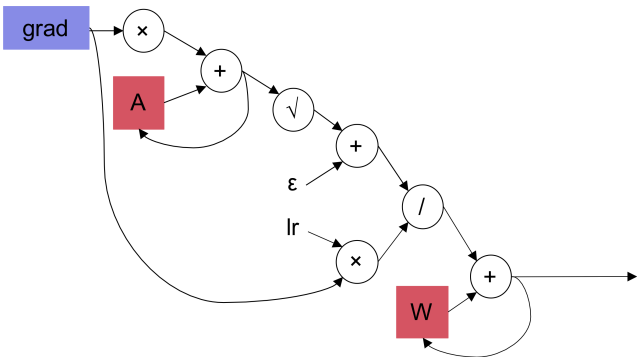

Fig. 6 Computation graph of Adagrad

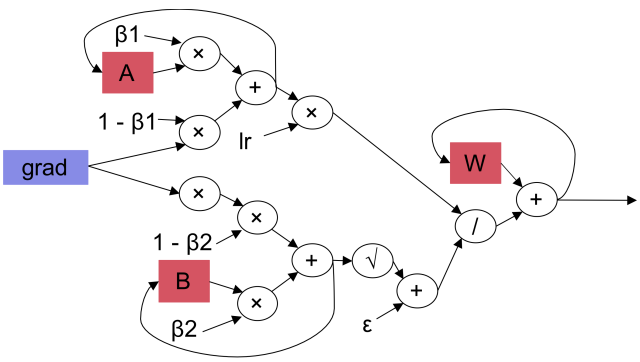

Fig. 7 Computation graph of Adam

FPGA-based switch. These algorithms are selected so as to cover both the simple algorithm (i.e., SGD) and sophisticated algorithm that requires relatively higher computation cost (i.e., SMORMS3).

Figures 5, 6, 7, and 8 show their computation graphs, respectively. In the computation graphs, circle symbols represent computational operations, such as addition, subtraction, multiplication, division, square root, and minimum. Square symbols represent storage elements, each of which is corresponding to a floating point number. *grad* represents input gradients and W represents weight parameters under optimization. The other symbols, such as A, B, and C, are algorithm-specific parameters.

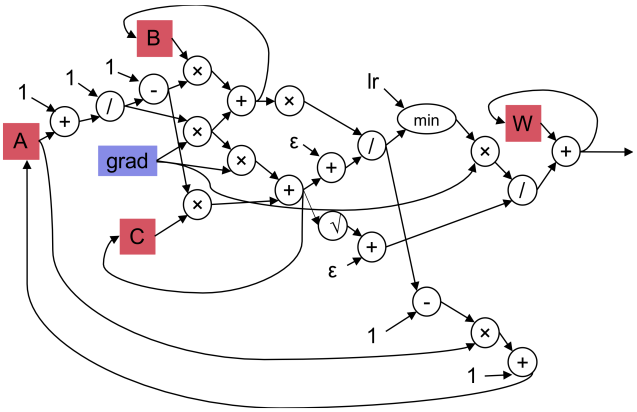

Fig. 8 Computation graph of SMORMS3

4. Packet Trace

4.1 GPU Packet Trace

Prior to the implementation, this section describes packet traces of remote GPUs connected via a PCIe over 10GbE technology. The packet traces were analyzed so that we can extract the gradients from ExpEther packets and write back optimized weight parameters to the packets. The packets were also used for the evaluations of the proposed FPGA-based network switch.

Since NVIDIA's GPUs are used for the gradient computation, CUDA (Compute Unified Device Architecture) [14] is used as an integrated development environment for the remote GPUs. In a CUDA program, a cudaMemcpy function copies data from host main memory to GPU device memory. The instruction and data are transferred as PCIe over 10GbE packets and go through the proposed FPGA-based network switch. We collected the PCIe over 10GbE packet traces between a host machine and remote GPU devices.

Figure 9 shows the packet capture environment. As shown, there are three machines: a host machine, a switch machine, and a capture machine. The host machine is equipped with an ExpEther host adapter. It executes a CUDA program for a remote GPU connected via the proposed network switch. The switch machine is equipped with a NetFPGA-SUME card that has a Xilinx Virtex-7 XC7VX690T FPGA and four SPF+ connectors for 10GbE interfaces. This FPGA card is used as a 4-port 10GbE switch. The packet capture machine is equipped with a 10GbE network interface card connected to the network switch. Wireshark is used as a packet capture software at the capture machine.

In this environment, when a CUDA program is executed on the host machine for the remote GPU, PCIe over 10GbE packets are transferred between the host machine and the remote GPU. These packets go through the network switch and are captured by the capture machine.

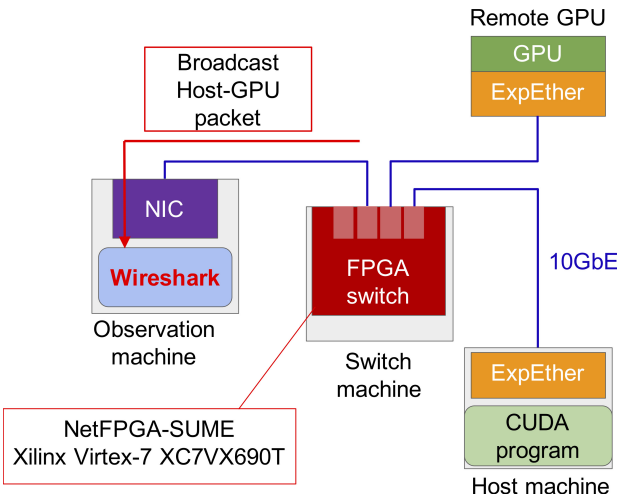

Fig. 9 Packet capture environment

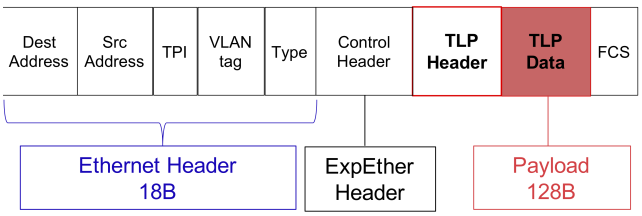

Fig. 10 The packet format of ExpEther [15]

4.2 ExpEther Packets

Figure 10 shows the format of the captured ExpEther packet. ExpEther adds the Ethernet header to the packet communicated by PCI Express, and communicates it as an Ethernet frame. The first part of the packet is the Ethernet header, and the VLAN-tag is inserted. Next is the control header of ExpEther [15]. Finally, the main part contains TLP (Transaction Layer Packet), which is packet used for sending and receiving data in PCIe.

When cudaMemcpy is executed to copy values on the GPU to the host machine, packets containing the value in TLP data pass through the switch. In the implementation of this paper, the data communicated between the host machine and the remote GPU is obtained from TLP data in packets that pass through the switch.

5. Implementation

5.1 Baseline Network Switch

Reference Switch Lite design provided by NetFPGA team [6] is used as a baseline 10GbE switch implemented on NetFPGA-SUME card. The gradient aggregation and parameter optimization modules are inserted to this baseline switch. Figure 11 shows a block diagram of the FPGA-based network switch. When packets are injected to the network switch, they are passed from one of four 10GbE in-

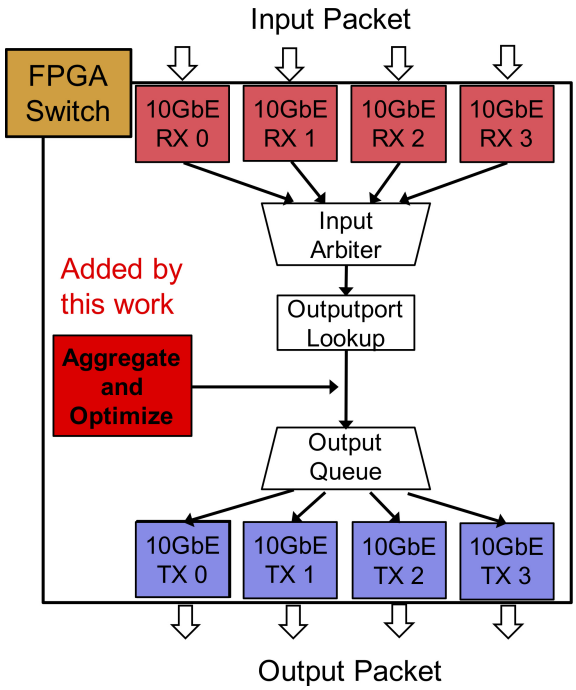

Fig. 11 Block diagram of FPGA-based network switch

terfaces (i.e., RX0 to RX3). Input Arbiter module receives packets one by one in a round robin manner from these interfaces. Next, Outputport Lookup module performs a routing function based on a packet header to determine a destination port of the switch. The gradient aggregation and parameter optimization modules judge if incoming packets are subject to the gradient aggregation, parameter optimization, or none of them. The gradient aggregation and parameter optimization modules are explained in the next subsections. Finally, Output Queue module distributes packets to one of the four 10GbE ports according to the routing result by Outputport Lookup module.

5.2 Gradient Aggregation and Parameter Optimization Modules

Figure 12 illustrates the gradient aggregation and parameter optimizations modules proposed in this paper. In these modules, first, based on the ExpEther packet format captured in Sect. 4.1, incoming packets are analyzed to see if they are subject to the gradient aggregation or parameter optimization modules. If the packet is not related to the aggregation nor parameter optimization, the packet simply skips these modules.

If the packet is subject to the gradient aggregation, the sequence number of the gradient field in the packet is read, and the corresponding gradient is retrieved from a BRAM that stores the aggregated gradients. Then, the gradients of the packet and those retrieved from the BRAM are added. After the addition, the result is written back to both the packet and BRAM, and then the packet is sent to the next module.

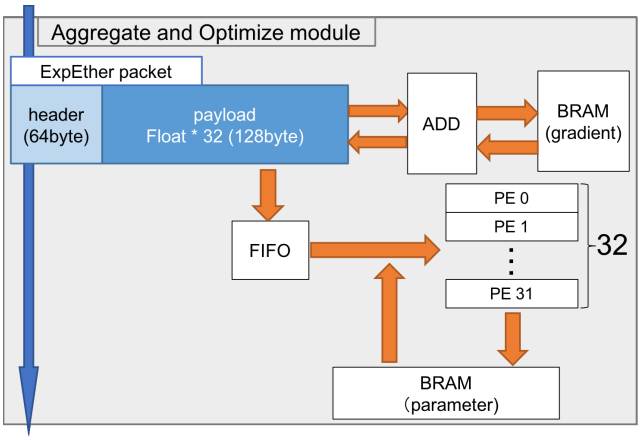

Fig. 12 Gradient aggregation and parameter optimization modules

If the packet is subject to the parameter optimization, in the same way, the sequence number of the gradient field in the packet is read, and the corresponding parameter is retrieved from a BRAM that stores the parameter. The gradients extracted from the packet and the parameters extracted from the BRAM are fed to PEs, and the parameter optimization is performed. After the parameter update is completed, new parameters are stored in the BRAM.

5.3 Parameter Optimization Algorithms

The parameter optimizer in the 10GbE switch is implemented on NetFPGA-SUME card. The target FPGA device is Xilinx Virtex-7 XC7VX690T. Xilinx Vivado v2016.4 is used for logic synthesis and implementation. The target operating frequency is 200MHz. As arithmetic IP cores, Floating-Point Operator v7.0 provided by Xilinx is used for these algorithms. These IP cores include addition, subtraction, multiplication, division, size comparison, and square root operations. They use 32-bit single-precision floating-point numbers. The four optimization algorithms (i.e., SGD, Adagrad, Adam, and SMORMS3) are implemented on NetFPGA-SUME card. These algorithms are implemented by combining and/or cascading the above-mentioned arithmetic IP cores as shown in Figs. 5 to 8.

The latencies (the number of clock cycles) to complete these algorithms are listed in Table 2. Here, the latency is a duration between when a single-precision floating-point input data is injected and when the corresponding computation result is generated.

Please note that the above-mentioned parameter optimization cores are in charge of single input data only. To accelerate the parameter optimization algorithms, multiple instances or PEs (processing elements) of these optimizer cores are implemented, as shown in Fig. 13. Floating-point numbers of input data are distributed to these PEs in a round-robin manner and processed in parallel. Throughput of the parameter optimization increases as the number of the optimizer PEs is increased, as long as the parallelism of input data can be exploited; thus there is a trade-off between the

Table 2 Latency for processing a single parameter by a single PE for each optimization algorithm

Optimizer	# of cycles
SGD	21
Adagrad	99
Adam	137
SMORMS3	164

Table 3 FPGA resource utilization of a single PE for each optimization algorithm

Optimizer	LUTs	FFs	DSPs
SGD	535 (0.12%)	953 (0.11%)	3 (0.08%)
Adagrad	2,695 (0.62%)	4,880 (0.56%)	9 (0.25%)
Adam	5,409 (1.25%)	10,278 (1.19%)	36 (1.00%)
SMORMS3	6,100 (1.41%)	11,510 (1.33%)	40 (1.11%)

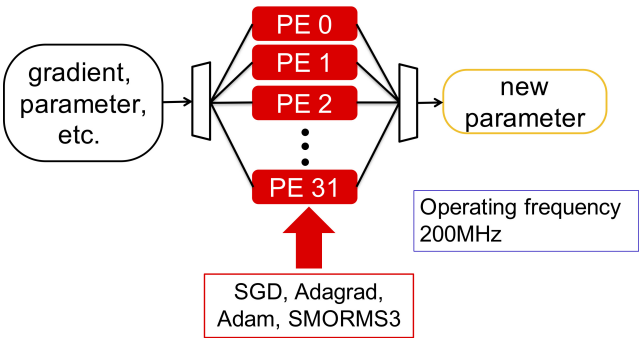

Fig. 13 Multi-PE implementation of optimization algorithms

throughput and the resource utilizations. Because of underlying technology of the PCIe over 10GbE, since 32 gradients can be stored in a single packet, we implemented 32 PEs on the proposed FPGA-based switch to process such packets at a time. It may be possible to implement 64 PEs to process two packets at a time.

The four optimization algorithms are evaluated in terms of the FPGA resource utilizations of LUTs (Look Up Tables), FFs (Flip Flops), and DSP (Digital Signal Processing) slices. Table 3 shows the result. In this implementation, 32 gradients are included in a single ExpEther packet, and thus 32 PEs should be implemented to fully exploit the parallelism of input data. As shown in Table 3, the PE sizes are quite small, and thus we implemented 32 PEs for each algorithm on the FPGA-based network switch.

6. Evaluations

6.1 Gradient Aggregation Throughput

First, the proposed FPGA-based network switch is evaluated in terms of the gradient aggregation throughput. Test packets including gradients in 10GbE line rate are generated by using Open Source Network Tester [16] and sent to the proposed network switch. The gradients are represented as an array of 32-bit single-precision floating-point numbers. In this evaluation, each packet contains 32 gradients and thus the packet length is 192 bytes including a packet header

and the payload. Open Source Network Tester is directly connected to the proposed network switch with a 10GbE SFP+ cable, and the gradient aggregation is executed on the switch. The aggregation throughput is measured by counting the number of packets processed by the proposed switch.

The measurements are performed ten times and the average throughput is 8.92Gbps. Assuming the packet length is 192 bytes, the 10GbE line rate in our environment is 9.07Gbps when considering the Ethernet preamble and interframe gap inserted for each packet. In this case, the measured throughput of the gradient aggregation is corresponding to 98.3% of the 10GbE line rate in our environment, and thus almost the line rate is achieved.

In the experiments, a single FPGA-based switch is inserted between a single remote GPU and a single host machine (i.e., single source and single destination case). It is possible to connect up to three remote GPUs to the FPGA-based switch since it has four 10GbE interfaces (i.e., multiple sources and single destination case). In this case, however, communication bandwidth (e.g., 10Gbps) between host machine and the FPGA-based switch would be a bottleneck since multiple remote GPUs share the same 10GbE link connected to the host machine. Please note that our proposed aggregation and parameter optimization functions achieve almost 10GbE line rate. In other words, the 10GbE host-link would be the bottleneck rather than the aggregation and parameter optimization overheads when the number of remote GPUs is increased.

6.2 Resource Utilization

The proposed network switch including the gradient aggregation and parameter optimization modules is evaluated in terms of FPGA resource utilizations of LUTs, BRAM, and DSP slices. Figure 14, 15, and 16 shows the resource utilizations of an entire switch in the cases of the four parameter optimization algorithms. The blue bar represents the resource utilization of the baseline network switch introduced in Sect. 5.1. The orange bar represents the resources increased by adding gradient aggregation and parameter optimization module using 32 PEs to the baseline switch. The addition of the gradient aggregation and parameter optimization module consumed a small amount of BRAM. The resource utilization of SGD version is the lowest. It consumes approximately 15% of LUTs. Adam and SMORMS3 versions consume more resources. Especially, SMORMS3 version consumes approximately 56% of LUTs, but even with Reference Switch modules, their resource utilizations still have room to add more PEs.

6.3 Parameter Optimization Latency

In general, the parameter optimization is done by a host CPU or GPU after the workers compute gradients. By introducing the proposed network switch, the parameter optimization is completed in the middle of communication path between host CPU and remote GPUs. In this section, the proposed network switch is evaluated in terms of the execution time of the parameter optimization.

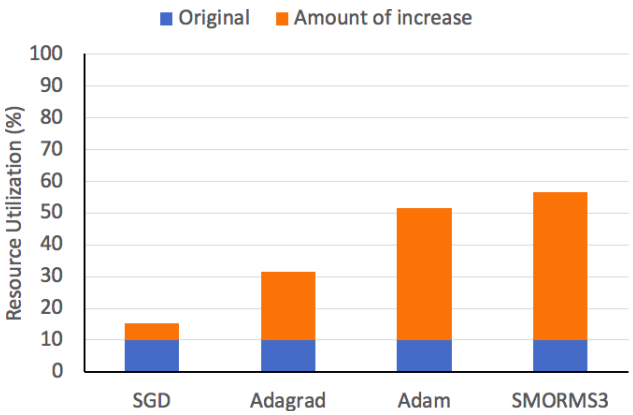

Fig. 14 Resource utilization (LUTs)

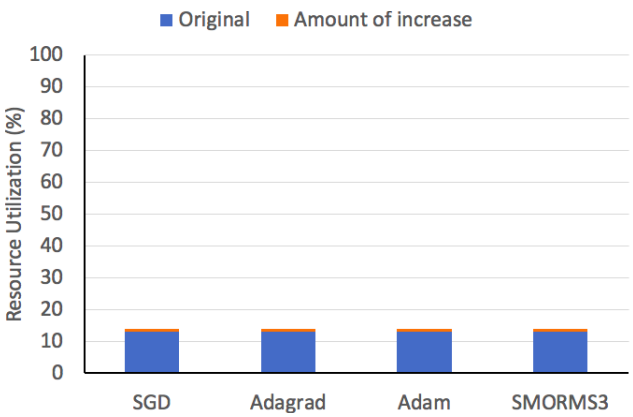

Fig. 15 Resource utilization (BRAM)

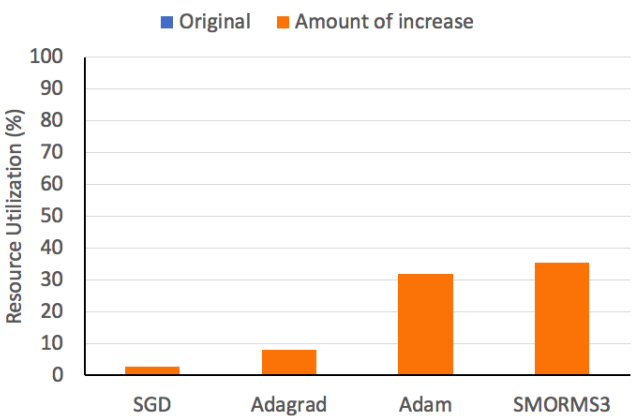

Fig. 16 Resource utilization (DSP)

The proposed network switch is compared with the following CPU and GPU-based approaches in terms of the parameter optimization.

1. CPU-based approach: Assuming a host CPU has already-aggregated gradients, the execution time for the parameter optimization by the CPU is measured.
2. GPU-based approach: After a GPU device receives

Table 4 CPU- and GPU-based execution environment

OS	Ubuntu 16.04
CPU	Intel Core i7-6800K @3.4GHz
Memory	32GB
GPU	NVIDIA Geforce GTX 1080Ti (11GB RAM) x4
CUDA	version 9.0
Chainer	version 5.2.0

already-aggregated gradients from a host machine, the execution time for the parameter optimization by the GPU is measured.

3. Proposed FPGA-based network switch: Assuming the proposed network switch is placed in a communication path between a host CPU and a remote GPU device, the number of cycles for the parameter optimizations by the network switch is measured.

The CPU- and GPU-based approaches are implemented as a software program in Chainer [13]. Their evaluation environment is listed in Table 4 [†]. In this evaluation, 800,000 parameters are assumed, so that the latest DNN models, such as DenseNet-BN [17], can be supported. SGD, Adagrad, Adam, and SMORMS3 algorithms are executed for 800,000 gradients using these three approaches.

Figure 17 shows the execution times of the parameter optimization of the four algorithms with the three approaches: CPU-based, GPU-based, and the proposed FPGA-based network switch approaches. In this graph, X-axis represents the algorithms and Y-axis represents their execution times. In the case of SGD, differences between the three approaches are small. However, the differences become large as the algorithm becomes complicated; in the case of SMORMS3, the differences are the largest. As a result, the proposed network switch approach outperforms the CPU- and GPU-based approaches by 1.2-3.0x and 1.05-1.25x, respectively.

In the proposed network switch approach, the differences between the four optimization algorithms are also quite small. This is because the serialization of input gradients is a major bottleneck in the proposed network switch regardless of the optimization algorithm selected. Among the four optimization algorithms, it turns out that the use of sophisticated algorithms (e.g., SMORMS3) is beneficial in the proposed network switch approach compared to the CPU- and GPU-based approaches.

6.4 Discussions

Although a major bottleneck of the training phase is still the gradient computation by GPU workers, here we estimate how much an entire training phase can be accelerated by introducing the proposed network switch in the case of the preliminary evaluation in Sect. 3.1. Since the aggregation and parameter optimization are done by the proposed FPGA-based switch during packet transfer between the remote GPUs and host machine, "Aggregation and Optimization" of Fig. 3 are accelerated in this paper. These parts account for 8.1% (2.9% for Aggregation and 5.2% for Optimization) of total execution time. The proposed network switch accelerates the parameter optimization of Adam algorithm by approximately 1.2x compared to the GPU-based approach. Also, the parameter aggregation is done in a stream processing manner in 98.3% of the 10GbE line rate, which can eliminate most computation time for the aggregation. Based on these results, an entire training phase is shortened by approximately 5% in the case of the preliminary evaluation in Sect. 3.1. In addition to this speedup, since the gradient aggregation and parameter optimization are offloaded to the network switch, the saved CPU and GPU resources can be used for the other tasks.

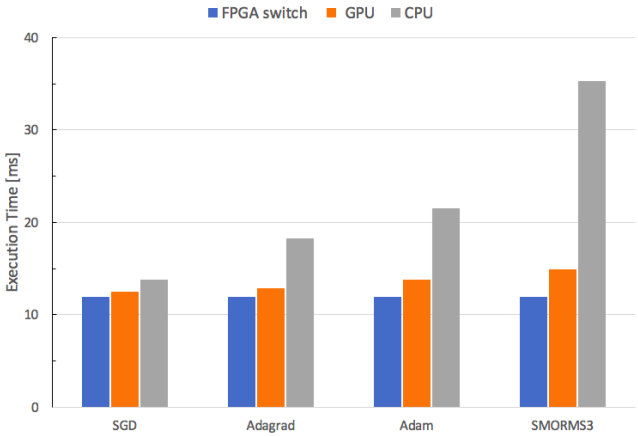

Fig. 17 Execution time of parameter optimization

In this paper, only four GPUs are used for the evaluations, but more GPUs are typically used in distributed deep learning. As the number of GPUs increases, execution times for the forward propagation and back propagation decrease, but that for the gradient aggregation increases. Furthermore, since the execution time for the parameter optimization cannot be accelerated as the number of GPUs is increased, the proposed network switch can accelerate especially such large-scale distributed deep learning with many GPUs.

Our target FPGA board (i.e., NetFPGA-SUME) has only four 10GbE interfaces, so the proposed FPGA-based switch currently implements four ports: one is connected to host machine and the others are connected to up to three remote GPUs. Our current implementation thus supports up to three GPUs. To support more remote GPUs, more FPGA-based switches will be necessary. However, the number of the proposed switches is limited by the number of PCIe slots of the host machine. In addition, communication bandwidth (e.g., 10Gbps) between host machine and the proposed FPGA-based switch would be a bottleneck when multiple remote GPUs are connected to the switch.

[†]The machine used slightly differs from that used in the preliminary evaluation in Sect. 3.1 due to availability of the machine.

7. Conclusions

In this paper, the gradient aggregation and parameter optimization were accelerated by an FPGA-based network switch for distributed deep learning using remote GPUs via PCIe over 10GbE. In distributed deep learning, the number of GPUs is typically increased in order to increase the degree of parallelism and shorten the training time. However, communication overheads including gradient aggregation cannot be ignored in large-scale distributed deep learning. In this paper, we thus introduced remote GPUs via PCIe over 10GbE and reduced the overhead by offloading the gradient aggregation and parameter optimization to the FPGA-based network switch placed in the middle of communication. There are several parameter optimization algorithms with different characteristics, such as computation cost and the number of iterations to convergence. Four parameter optimization algorithms including SGD, Adagrad, Adam, and SMORMS3 were implemented in the proposed network switch.

Evaluation results of the proposed FPGA-based network switch demonstrated that the gradient aggregation achieved 98.3% of the 10GbE line rate. The resource utilization was less than 56% at the maximum when 32 parallel PEs were implemented. As for the parameter optimization, the proposed network switch outperformed CPU- and GPU-based approaches by approximately 1.2-3.0x and 1.05-1.25x, respectively. Also, we estimated that the overall training phase would be accelerated by approximately 5% by introducing the proposed network switch. As a future work, we are planning to demonstrate the performance improvement of overall training phase in a real environment that consists of a host machine, four remote GPUs via PCIe over 10GbE, and the proposed FPGA-based network switch.

References

[1] T. Akiba, S. Suzuki, and K. Fukuda, "Extremely large minibatch SGD: Training ResNet-50 on ImageNet in 15 minutes," CoRR, vol.abs/1711.04325, Nov. 2017.
[2] H. Mikami, H. Suganuma, P. U-chupala, Y. Tanaka, and Y. Kageyama, "Massively distributed SGD: ImageNet/ResNet-50 training in a flash," CoRR, vol.abs/1811.05233, Nov. 2018.
[3] M. Yamazaki, A. Kasagi, A. Tabuchi, T. Honda, M. Miwa, N. Fukumoto, T. Tabaru, A. Ike, and K. Nakashima, "Yet another accelerated SGD: ResNet-50 training on ImageNet in 74.7 seconds," CoRR, vol.abs/1903.12650, March 2019.
[4] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, M. Kudlur, J. Levenberg, R. Monga, S. Moore, D. G. Murray, B. Steiner, P. Tucker, V. Vasudevan, P. Warden, M. Wicke, Y. Yu, and X. Zheng, "TensorFlow: A system for large-scale machine learning," Proc. USENIX Symposium on Operating Systems Design and Implementation (OSDI'16), pp.265–283, Nov. 2016.
[5] J. Suzuki, Y. Hidaka, J. Higuchi, T. Yoshikawa, and A. Iwata, "ExpressEther - Ethernet-based virtualization technology for reconfigurable hardware platform," Proc. IEEE Symposium on High-Performance Interconnects (HOTI'06), pp.45–51, Aug. 2006.
[6] "The NetFPGA Project." https://netfpga.org.
[7] T. Itsubo, M. Koibuchi, H. Amano, and H. Matsutani, "Accelerating deep learning using multiple GPUs and FPGA-based 10GbE switch," Proc. Int. Conf. Parallel, Distributed and Network-Based Processing (PDP'2020), pp.102–109, March 2020.
[8] J. Duchi, E. Hazan, and Y. Singer, "Adaptive subgradient methods for online learning and stochastic optimization," J. Mach. Learn. Res., vol.12, pp.2121–2159, July 2011.
[9] D.P. Kingma and J. Ba, "Adam: A method for stochastic optimization," CoRR, vol.abs/1412.6980, Dec. 2014.
[10] "RMSprop loses to SMORMS3." https://sifter.org/~simon/journal/20150420.html.
[11] Y. Li, J. Park, M. Alian, Y. Yuan, Z. Qu, P. Pan, R. Wang, A. Schwing, H. Esmaeilzadeh, and N.S. Kim, "A network-centric hardware/algorithm co-design to accelerate distributed training of deep neural networks," Proc. Int. Symp. Microarchitecture (MICRO'18), pp.175–188, Oct. 2018.
[12] S. Morishima and H. Matsutani, "High-performance with an In-GPU graph database cache," IEEE IT Professional, vol.19, no.6, pp.58–64, Nov. 2017.
[13] "Chainer: A powerful, flexible, and intuitive framework for neural networks." https://chainer.org.
[14] "NVIDIA CUDA." https://developer.nvidia.com/cuda-zone.
[15] J. Suzuki, Y. Hayashi, M. Kan, S. Miyakawa, and T. Yoshikawa, "End-to-end adaptive packet aggregation for high-throughput I/O bus network using ethernet," Proc. IEEE Symposium on High-Performance Interconnects (HOTI'14), pp.17–24, Aug. 2014.
[16] "OSNT 10G Home." https://github.com/NetFPGA/OSNT-Public/wiki/OSNT-10G-Home.
[17] G. Huang, Z. Liu, L. van der Matten, and K.Q. Weinberger, "Densely connected convolutional networks," Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR'17), pp.2261–2269, July 2017.

Tomoya Itsubo received the B.E. degree from Keio University in 2019. He is currently a master course student in Keio University.

Michihiro Koibuchi received the B.E., M.E., and Ph.D. degrees from Keio University, Yokohama, Japan, in 2000, 2002 and 2003, respectively. Currently, he is an associate professor in the Information Systems Architecture Research Division, National Institute of Informatics and the Graduate University of Advanced Studies, Tokyo, Japan. His research interests include the area of high-performance computing and interconnection networks.

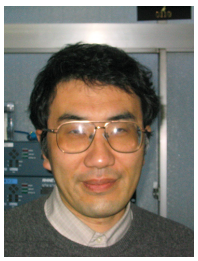

Hideharu Amano received the Ph.D. degree from the Department of Electronic Engineering, Keio University, Japan in 1986. He is currently a professor in the Department of Information and Computer Science, Keio University. His research interests include the area of parallel architectures and reconfigurable systems.

Hiroki Matsutani received the B.A., M.E., and Ph.D. degrees from Keio University in 2004, 2006, and 2008, respectively. He is currently an associate professor in the Department of Information and Computer Science, Keio University. His research interests include the areas of computer architecture and interconnection networks.