LETTER

Disaggregated Accelerator Management System for Cloud Data Centers*

Ryousei TAKANO^{†a)} and Kuniyasu SUZAKI[†], *Members*

SUMMARY A conventional data center that consists of monolithic servers is confronted with limitations including lack of operational flexibility, low resource utilization, low maintainability, etc. Resource disaggregation is a promising solution to address the above issues. We propose a concept of disaggregated cloud data center architecture called Flow-in-Cloud (FiC) that enables an existing cluster computer system to expand an accelerator pool through a high-speed network. FlowOS-RM manages the entire pool resources, and deploys a user job on a dynamically constructed slice according to a user request. This slice consists of compute nodes and accelerators where each accelerator is attached to the corresponding compute node. This paper demonstrates the feasibility of FiC in a proof of concept experiment running a distributed deep learning application on the prototype system. The result successfully warrants the applicability of the proposed system.

key words: resource disaggregation, resource management, cloud computing

1. Introduction

The end of Moore's law is coming within a decade due to technical and economic limitations. No more drastic performance improvement for general purpose processors is expected and new computing paradigms and architectures are needed for the explosive growing computational workload such as big data analysis, deep learning training and inference, and so on. Specialization or, in other words, domain specific architecture (DSA) is a promising research direction in the Post-Moore era. Specifically, many task-specific accelerators including Google TPU, Fujitsu DLU, Microsoft BrainWave, and D-Wave Quantum Annealer were proposed recently. To take advantage of such accelerators, it is important to establish a resource management system to fully utilize a variety of hardware resources, including a generic processor, an accelerator, and storage, depending on the workloads. However, a conventional data center consists of monolithic servers and it cannot provide such flexible use of computing hardware resources. It also faces limitations including lack of operational flexibility, low resource utilization, low maintainability, etc [2], [3].

Manuscript received March 17, 2020.

Manuscript revised September 13, 2020.

Manuscript publicized December 7, 2020.

[†]The authors are with Information Technology Research Institute, National Institute of Advanced Industrial Science and Technology (AIST), Tsukuba-shi, 305-8560 Japan.

*This paper extends our preliminary work published at SC2018 research poster [1]. Specifically, we conducted additional experiments with more realistic deep learning workload. Moreover, we have clarified the novelty in related work section.

a) E-mail: takano-ryousei@aist.go.jp

DOI: 10.1587/transinf.2020EDL8040

To address the limitations of conventional data centers, there is an emerging interest in resource disaggregation [2]–[6] that decomposes monolithic servers into independent hardware components, including CPU, accelerator, memory, and storage, through a high-speed network. In a disaggregated data center, hardware components are separated in each resource pool and reconstructed to meet the user requirements. We have proposed a new concept of disaggregated data center architecture, Flow-in-Cloud (FiC), that enables an existing PC cluster to expand an accelerator pool through a high-speed network. FlowOS manages hardware resources and application jobs on FiC. To demonstrate the feasibility of FiC and FlowOS, currently we are developing FlowOS on the prototype system of FiC. This paper focuses on the resource management system of FlowOS called FlowOS-RM and reports on the effectiveness for distributed deep learning applications.

2. Related Work

Some papers [2], [3] have reported that the resource utilization, e.g., CPU or main memory, varies considerably for each server in commercial data centers. This is because it is quite difficult to assign various workloads in such a way that all resources are fully and equally consumed. As a result, the resource utilization remains low. To address this problem, many studies of resource disaggregation have emerged since the middle of the 2010s. This movement started with hardware-level resource disaggregation [4]–[7] and expanded to the OS-level resource disaggregation [2], [3] in recent years. The interconnection technology for enabling disaggregation, including Gen-Z, is being standardized and will be commercialized in the near future. The proposed system addresses device disaggregation with a special focus on accelerators, and it seamlessly extends an existing cloud data center by facilitating access an accelerator resource pool. This characteristic allows us to benefit from the software ecosystem of an existing cluster resource management system like Apache Mesos. To the best of our knowledge, there is no existing works that consider cooperation with cluster resource management systems. Although we use an electric interconnection network, an optical network is promising as several researchers have proposed in [4]–[7].

Several device disaggregation technologies have been proposed in [8], [9]. ExpEther [8] is a PCIe-over-Ethernet technology and it allows us to dynamically attach and detach remote PCIe devices through Ethernet. On the other

hand, rCUDA [9] is an OS-level disaggregation technology. Although it works with only NVIDIA GPUs, it can seamlessly access a remote GPU through Infiniband and Ethernet. Our preliminary experiment [10] shows the performance overhead of device disaggregation technologies. The host-to-device bandwidth of ExpEther is about 20% of that of a local PCIe device. While this performance degradation is a worse-case situation, some application performances are regulated by the amount of traffic among the host and devices. On the other hand, the impact on computation bound applications like GEMM and convolution is negligible. The practical problem with rCUDA is the lack of compatibility. For example, it does not support cuDNN, which is heavily used on deep learning applications.

3. Flow-in-Cloud

Flow-in-Cloud (FiC) is a proof of concept system for a disaggregated cloud data center, and provides the user with a slice of resources as an application execution environment. An application is then divided into several tasks and each task is optimized through the use of a suitable accelerator. In the case of deep learning, a convolution layer task is executed on GPU, and a full connected layer task is executed on FPGA. We call such a set of accelerators as a meta accelerator. A slice is dynamically configured by combining meta accelerators and attaching them to corresponding compute nodes according to a user request, as shown in Fig. 1. Compute nodes and accelerators are connected through a high-speed circuit-switched network called FiC network, and it comprises a set of FiC switch boards [11] that have a middle grade FPGA chip (Xilinx Kintex Ultrascale XCKU095), 32-10Gbps FireFly serial connections and DRAM. In addition, Raspberry Pi 3 is implemented on the board as a controller. Raspberry Pi 3 communicates with FPGA via GPIO for configuration of FPGA and data transmission. Note that we employ circuit switching instead of packet switching because friction-less transition from electric network to optical network is possible. A circuit-switching logic and a user-defined logic written in a high-level synthesis language are running on the FPGA, and the latter logic is partially reconfigurable in advance of application deployment.

FlowOS manages the entire pool of FiC resources, and supports the execution of a user job on provided slices. FlowOS employs a layered architecture including FlowOS-

Fig. 1 The overview of Flow-in-Cloud architecture

Job, FlowOS-RM, and FlowOS-drivers. FlowOS-Job is a heterogeneous programming framework that allows the users to describe a job composed of several tasks, where each task is optimized for a specific accelerator according to the workload. FlowOS-RM is a resource manager and the detail is described in Sect. 4. FlowOS-driver is a proxy component to access underlying hardware resources like accelerator pools and compute nodes.

Currently, we have implemented FlowOS on a small disaggregated data center environment where compute nodes and accelerators are connected through ExpEther [8] instead of FiC network. FlowOS provides two major features: disaggregated resource management and heterogeneous programming framework as mentioned above. In this paper, we focus on the latter and demonstrate disaggregate device management and cooperation with a cluster resource management system. Although ExpEther cannot support direct communication among accelerators as FiC originally addresses, it is a reasonable alternative technology to demonstrate the concept of FiC using commodity hardware.

4. FlowOS-RM

FlowOS-RM seamlessly works in cooperation with a cluster resource management system such as Apache Mesos [12], Kubernetes, SLURM, and so on. In other words, FlowOS-RM extends such systems to support accelerator disaggregation.

FlowOS-RM works in cooperation with the following components: (1) Disaggregate device management: ExpEther is a PCIe-over-Ethernet technology and it allows us to dynamically attach and detach remote PCIe devices through Ethernet. (2) OS deployment: Bare-Metal Container (BMC) [13] constructs an execution environment to run a Docker image with an application optimized OS kernel on a node. (3) Task scheduling and execution: FlowOS-RM is implemented on top of a Mesos framework, and it co-allocates nodes to meet a user requirement and launches a task on each node in the manner of Mesos.

FlowOS-RM provides users with the REST API to configure a slice and execute a job on it. A single-node job as well as an MPI type multi-node job is supported. Figure 2

Fig. 2 Job execution flow in FlowOS-RM

Table 1 Major operations in FlowOS-RM

attach-device	Attach devices to a node
launch-machine	Boot a node with a specific OS kernel and container, and it joins active nodes under Mesos
prepare-task	Do housekeeping for launching a task, including submitting a task to the corresponding node through Mesos
launch-task	Launch a task in a node (running state)
detach-device	Detach devices from a node
destroy-machine	Shutdown a node and it leaves from active nodes

Table 2 Experimental setting

Compute Node Configuration	
CPU	10-core Intel Xeon E5-2630v4/2.2GHz
M/B	Supermicro X10SRG-F
Memory	128 GB DDR4-2133
Network	ExpEther 40G HBA
NIC	Intel I350 (Gigabit Ethernet)
Disaggregated Resource (PCIe device)	
GPU	NVIDIA Tesla P100 x4, P40 x1
NVMe	Intel SSD 750 x4
Software Configuration	
OS	CentOS 7.4
Kernel	Linux 3.10.0-514.26.2.el7.x86_64 Mesos 1.4.1, ChainerMN, OpenMPI 3.1.0 NVIDIA CUDA 8.0.61

presents a job execution flow in FlowOS-RM, where a job is a set of tasks and each task runs on a node belonging to a slice. Table 1 summarizes each operation of FlowOS-RM. First, a slice is constructed in two steps: *attach-device* and *launch-machine*. Second, a job is launched in the following two steps: *prepare-task* and *launch-task*. After job execution, the slice is destructed in two steps: *detach-device* and *destroy-machine*.

5. Experiment

5.1 Experimental Setting

In order to demonstrate the feasibility of FlowOS-RM, we have conducted distributed deep learning training experiments on a four-node cluster environment as shown in Fig. 3(a). Each compute node has two ExpEther HBAs to connect PCIe devices, e.g., NVIDIA Tesla P100 GPU and Intel NVMe SSD, on I/O Boxes through a 40 GbE switch. Linux is running on each compute node, and FlowOS-RM and Mesos are installed on this environment. We used two applications, a handwriting character recognition (MNIST) and a large-scale image classification (ImageNet), as benchmark programs, and they are implemented with a distributed deep learning framework ChainerMN [14].

5.2 Experimental Results

5.2.1 Slice Construction and Destruction Overheads

We have demonstrated a flexible resource management and

Fig. 3 Experimental configuration

the performance overhead of FlowOS-RM. Firstly, we ran an MNIST application on three different slice configurations as shown in Fig. 3(b), and the breakdown in the execution time for each slice is shown in Fig. 4(a). An MNIST training runs faster as the number of GPUs per node increases. The *run-task* elapsed times of 4node-1gpu, 2node-2gpu, and 1node-4gpu are 366.36, 237.31, and 104.57 seconds, respectively. It is a relatively lightweight workload and the slice construction and destruction operations account for 32% to 45% of the total execution time. Specifically, a *launch-machine* operation takes longer as the number of nodes increases, because downloading a container image that is about 3GB in size through GbE becomes the bottleneck. Some operations including *attach/detach-device* and *launch-task* take longer as the number of GPUs per node increases, because these operations are not parallelized. We plan to reduce the above overhead by using a faster network and parallelizing operations.

Secondly, we ran ImageNet, a more practical application on the same slice configurations. We used the ResNet-50 model and ILSVRC2012 dataset. In this experiment, a slice has not only GPUs but also two NVMe SSDs to store ILSVRC2012 dataset. Unlike an MNIST experiment, the slice construction and destruction operations account for 0.15% to 0.17% of the total execution time as shown in Figs. 4(c) and 4(b). Generally speaking, deep learning training execution time tends to significantly increase and the overhead of FlowOS-RM can be negligible.

5.2.2 Resource Sharing

We confirmed disaggregated resources are shared among several slices according to a user requirement. In this experiment, a user submitted four MNIST application jobs and FlowOS-RM allocated resources into each slice in the FIFO manner. The slice configurations of each job are as follows: Slice1 and 2 consist of 2node-2gpu (P100), Slice3 consists of 1node-1gpu (P40), and Slice4 consists of 4node-1gpu (P100). Figure 5(a) shows that resource sharing among

Fig. 4 Slice execution life cycle

Fig. 5 Resource sharing

slices works as expected.

6. Conclusion and Future Work

We have demonstrated flexible and effective resource sharing on the proposed disaggregated resource management system (FlowOS-RM) for AI and Big Data applications in next generation cloud data centers. We found some performance issues, but the impact is limited for long, hours-running applications like distributed deep learning training. Our future work is replacing ExpEther with the FiC network, and then opening up a new perspective of heterogeneous accelerator computing by leveraging resource disaggregation. Furthermore, in this experiment, we cannot take advantage of the potential of bare metal containers. Thus, we plan to evaluate various applications with applying performance optimization techniques such as a profile-guided optimization on this system.

Acknowledgements

The authors would like to thank Hidetaka Koie, SURIGIKEN for support on the engineering effort, and Jason Haga for his valuable comments. This paper is partially based on results obtained from a project commissioned by the New Energy and Industrial Technology Development Organization (NEDO).

References

- [1] R. Takano, K. Suzuki, and H. Koie, "FlowOS-RM: Disaggregated resource management system," ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis (Poster), 2018.
- [2] J. Gu, Y. Lee, Y. Zhang, M. Chowdhury, and K.G. Shin, "Efficient memory disaggregation with infiniswap," 14th USENIX Symposium on Networked Systems Design and Implementation (NSDI), pp.649–667, March 2017.
- [3] Y. Shan, Y. Huang, Y. Chen, and Y. Zhang, "LegoOS: A disseminated, distributed OS for hardware resource disaggregation," 13th USENIX Symposium on Operating Systems Design and Implementation (OSDI), pp.69–87, Oct. 2018.
- [4] K. Asanovic, "Firebox: A hardware building block for 2020 warehouse-scale computers," 2014.
- [5] K. Katrinis, D. Syrivelis, D. Pnevmatikatos, G. Zervas, D. Theodoropoulos, I. Koutsopoulos, K. Hasharoni, D. Raho, C. Pinto, F. Espina, S. Lopez-Buedo, Q. Chen, M. Nemirovsky, D. Roca, H. Klos, and T. Berends, "Rack-scale disaggregated cloud data centers: The dReDBox project vision," Design, Automation Test in Europe Conference Exhibition (DATE), pp.690–695, March 2016.
- [6] X. Guo, F. Yan, X. Xue, G. Exarchakos, and N. Calabretta, "Performance assessment of a novel rack-scale disaggregated data center with fast optical switch," Optical Fiber Communications Conference and Exhibition, pp.1–3, 2019.
- [7] HP Labs, "Memory-driven computing," <https://www.labs.hpe.com/memory-driven-computing>, 2018.
- [8] J. Suzuki, Y. Hidaka, J. Higuchi, T. Yoshikawa, and A. Iwata, "ExpressEther - Ethernet-Based Virtualization Technology for Reconfigurable Hardware Platform," 14th IEEE Symposium on High-Performance Interconnects (HOTI), pp.45–51, 2006.
- [9] J. Duato, A.J. Pena, F. Silla, R. Mayo, and E.S. Quintana-Ortí, "rCUDA: Reducing the number of GPU-based accelerators in high performance clusters," IEEE Int. Conf. High Performance Computing and Simulation (HPCS), pp.224–231, 2010.
- [10] R. Takano, T. Ikegami, K. Suzuki, A. Tanaka, and T. Hirofuchi, "FlowOS: A conception of system software for accelerator clouds," IPSJ SIG Technical Report 2017-HPC-163, pp.1–7, 2018. (in Japanese).
- [11] K. Hironaka, A.B. Ahmed, and H. Amano, "Multi-FPGA Management on Flow-in-Cloud Prototype System," 20th IEEE/ACIS Int. Conf. Software Engineering, Artificial Intelligence, Networking and Parallel/Distributed Computing (SNPD), pp.443–448, July 2019.
- [12] B. Hindman, A. Konwinski, M. Zaharia, A. Ghodsi, A.D. Joseph, R. Katz, S. Shenker, and I. Stoica, "Mesos: A platform for fine-grained resource sharing in the data center," NSDI, pp.1–14, 2011.
- [13] K. Suzuki, H. Koie, and R. Takano, "Bare-metal container — direct execution of a container image on a remote machine with an optimized kernel —," 18th IEEE Int. Conf. High Performance Computing and Communications (HPCC), pp.25–36, 2016.
- [14] T. Akiba, K. Fukuda, and S. Suzuki, "ChainerMN: Scalable distributed deep learning framework," Workshop on ML Systems, 2017.