High-Performance Custom Computing with FPGA Cluster as an Off-loading Engine

Takaaki Miyajima, Tomohiro Ueno, Atsushi Koshiba, Jens Huthmann, and Kentaro Sano
Center for Computational Science, RIKEN
Kobe, Hyogo, Japan
takaaki.miyajima@riken.jp

1 HIGH PERFORMANCE FPGA CLUSTER AND COMBINE IT WITH HPC SYSTEM

A heterogeneous system with Field Programmable Gate Arrays (FPGAs) is gathering attention in High-Performance Computing (HPC) area. We have been researching and developing a high-performance custom computing with FPGA cluster as an off-loading engine for an existing HPC system. We believe that one of the promising approaches is a combination of an existing HPC system and multiple FPGAs to cover wider applications by off-loading some of them. Although a few examples of large scale heterogeneous systems with FPGA are emerging, it is still challenging to efficiently utilize FPGAs combined with an existing HPC system. We have the following two main problems to achieve this goal. **How should we combine FPGA cluster with an existing HPC system? How should we achieve high-performance custom computing by off-loading tasks to FPGA clusters?** This poster overviews our motivation and current status of research topics.

1.1 Combine FPGA cluster with a HPC system

We are conducting a feasibility study of combining FPGA cluster with an existing HPC system such as Fugaku. Our FPGA cluster is a high-performance custom computing part and used as an off-load engine from the viewpoint of an existing HPC system. Figure 1 illustrates the architectural overview of our system. We assume that some kernels run on FPGAs but the others run on the host servers. An existing HPC system controls, manages and communicates with the FPGA cluster. FPGA cluster is loosely connected to an existing HPC system via InfiniBand HCA.

Our goal is to provide the following functionality. First, feature to enable host servers to access remote FPGAs to use remote FPGAs transparently and efficiently. Second, feature to make existing HPC node can handle arbitrary FPGAs. Third, feature to decouple a 1:1 ratio of host CPUs and FPGAs. Remote OPAE (rOPAE) is developed as an extension of Intel's OPAE library in order to provide the same functions from a remote node. rOPAE server library on FPGA-gateway server node works as a software bridge. It is flexible but causes overhead since the bridge is similar to a software switch. Software bridged data transfer (SBDT) function is developed to transfer data between two parts [3]. Preliminary evaluation on an older version FPGA shows that SBDT achieved 81.0% of theoretical bandwidth of PCIe Gen3 x8.

1.2 High performance custom computing with FPGA cluster

It is also a challenging problem how should we achieve high-performance custom computing with FPGA cluster. There are many problems to be solved for this goal. We have been researching and developing research topics such as a dedicated inter-FPGA network [2], programming model [1], and compiler for FPGAs [4]. Regarding the inter-FPGA network, we evaluate the is three kinds; (a) 1D ring network without a router, (b) 1D or 2D torus network with a router, and (c) Ethernet-based network with switch.

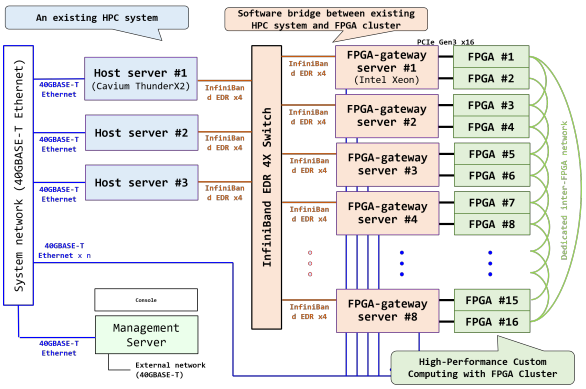

Figure 1: Our system consists of three main part; an existing HPC system, software bridge for FPGA cluster, and high-performance custom computing with FPGAs.

2 CONCLUSION

It is still a challenging how do we efficiently utilize multiple FPGAs and combine them with an existing HPC system. We are researching and developing multiple research topics to achieve this goal.

ACKNOWLEDGMENT

This work was partially supported by JSPS KAKENHI Grant Numbers 17H01706, and 19K20282.

REFERENCES

[1] Jinpil Lee, Tomohiro Ueno, Mitsuhisa Sato, and Kentaro Sano. 2018. High-productivity Programming and Optimization Framework for Stream Processing on FPGA. In *Proceedings of the 9th International Symposium on Highly-Efficient Accelerators and Reconfigurable Technologies, HEART 2018, Toronto, ON, Canada, June 20-22, 2018.* 5:1–5:6. https://doi.org/10.1145/3241793.3241798

[2] Antoniette Mondigo, Kentaro Sano, and Hiroyuki Takizawa. 2018. Performance Estimation of Deeply Pipelined Fluid Simulation on Multiple FPGAs with High-speed Communication Subsystem. In *29th IEEE International Conference on Application-specific Systems, Architectures and Processors, ASAP 2018, Milano, Italy, July 10-12, 2018.* 1–4. https://doi.org/10.1109/ASAP.2018.8445100

[3] Miyajima Takaaki, Hirao Tomoya, Miyamoto Naoya, Son Jeongdo, and Sano Kentaro. 2019. A software bridged data transfer on a FPGA cluster by using pipelining and InfiniBand verbs. In *Proceedings of the 9th International Symposium on Highly-Efficient Accelerators and Reconfigurable Technologies (HEART 2019).*

[4] Luzhou Wang, Kentaro Sano, and Satoru Yamamoto. 2012. Domain-Specific Language and Compiler for Stencil Computation on FPGA-Based Systolic Computational-Memory Array. In *Reconfigurable Computing: Architectures, Tools and Applications - 8th International Symposium, ARC 2012, Hong Kong, China, March 19-23, 2012. Proceedings.* 26–39.