Azerbaijan Journal of High Performance Computing, Vol 1, Issue 1, 2018, pp. 126-128
https://doi.org/10.32010/26166127.2018.1.1.126.128

The "CPU+GPU+MIC" Hybrid Supercomputer π

James Lin[1], Minhua Wen[1]

[1]*Center for HPC, Shanghai Jiao Tong University, Shanghai, China, james@sjtu.edu.cn, wenminghua@sjtu.edu.cn*

*Correspondence: James Lin, Center for HPC, Shanghai Jiao Tong University, Shanghai, China, james@sjtu.edu.cn

Abstract

The supercomputer π at Center for HPC (CHPC) of Shanghai Jiao Tong University (SJTU) was the fastest supercomputer among all the universities in China during 2013 to 2015. With a highly hybrid architecture, it is the first "CPU+GPU+MIC" supercomputer in China. It was also the largest Kepler GPU cluster in China when deployed in 2013. In addition, it firstly utilized the "FDR+DDN" combination in China. In this paper, we will introduce the design of supercomputer π.

1. Introduction

SJTU is one of few top Universities in China all strong in science, engineering and medicine. To meet higher and higher demanding of compute resource from scientist and researchers, CHPC built the π supercomputer in Year 2012 using best part of heterogeneous computing architecture creatively to address forefront scientific, engineering and medicine problems [1, 2]. This high performance computing system was designed and built as the basic platform to support the key disciplines and key laboratories in SJTU [3]. The application software should cover various fields, such as computational chemistry, molecular dynamics, CAD/CAE, etc., and the characteristics of different software are also very different. To this end, how to design a system to satisfy versatile computational tasks is a major challenge for CHPC. Based on this understanding, we designed the "CPU+GPU+MIC" high performance system [4].

2. Design specification

π, with a peak performance of 367 TFLOPS, ranked 158th of TOP500 in June 2013. π utilized a "CPU+GPU+MIC+FAT" hybrid architecture, with over 400 compute nodes. In order to increase the computing density, the use of high-density servers is already the trend of high-performance computing. For the CPU computing system, up to 4 compute nodes can be deposited in a 2U space, including a total of 8 CPUs and 64 computing cores. The ultra-high computational density also brings convenience to computational power and future scalability. The use of GPU+CPU heterogeneous systems to perform general scientific and engineering computation has become mainstream [5]. Utilizing GPUs for parallel computing can greatly improve the computational efficiency, although there are great challenges. For the computational tasks with large memory requirements, the general CPU nodes

cannot meet the requirements. The design of fat nodes with 256GB memory is used to handle such tasks.

π is also supported by a high-speed Infiniband 56G FDR network, which enables the carriage of data-intensive applications. Thanks to the space and electric power reserved during π 1.0 design, we can keep updating the computational resource, such as K40, K80 and P100. By 2016, near 40% performance is increased compared with the π 1.0 system [6, 7]. In 2016, we also designed the 4PB high performance secondary storage system to supplement the initial 1PB DDN storage, which greatly meets the users' demands of large data storage. Another highlight of the secondary storage is that it reserved the hardware interface to serve the upcoming π 2.0 system. The detailed specification of π is as Table 1.

By now, π has been supporting research projects various from science, engineering to medicine, with over 200 users from all the 16 scientific and engineering schools of SJTU.

TABLE 1 π specification

Performance	367 TFLOPS
CPU	830 Intel SandyBridge E5-2670 30 Intel Sandy Bridge E5-2680 V3
GPU	100 Nvidia Kepler K20, 10 Nvidia Kepler K40, 30 Nvidia Kepler K80, 4 Nvidia Kepler P10
Interconnection	Mellanox IB FDR 56Gbps
Parallel File Storage	DDN SFA12K 1PB with Lustre DELL MD3420 4PB with Lustre
SSD	80 Intel SSD 400G
Cooling System	Rittal Liquid

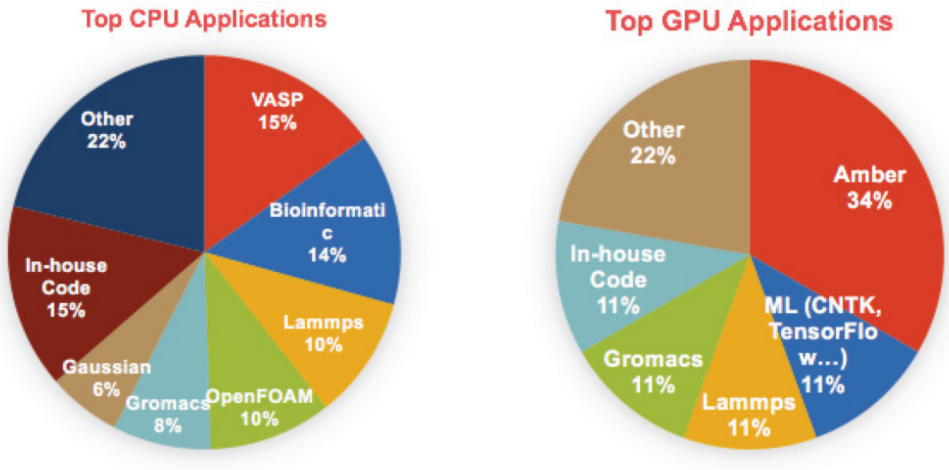

Figure 1 Top applications on Supercomputer π

With 5 years' high quality HPC service and lots of high end research, π has been one of the most influential supercomputer in China. In the near future, we will design and set up a new hybrid supercomputer with over 2 PFLOPS computational resource.

3. Conclusion

Nowadays requirements of scientific problems are so high that CPU is not able for their solution. For this purpose it is necessary to apply new architectures. π supercomputer used "CPU+GPU+MIC+FAT" hybrid architecture for eliminating of this problem. Therefore, this architecture allows us to solve scientific and engineering problems in less time, simultaneously its facilities let to run GPU-based applications on π supercomputer. Also hybrid architecture of π supercomputer gives opportunity for solution of new scientific problems.

Reference

[1]. Schulte, M. J. (2015) Achieving exascale capabilities through heterogeneous computing. IEEE Micro 35(4), 26-36.

[2]. Mittal, S., Jeffrey S. V. (2015) A survey of CPU-GPU heterogeneous computing techniques. ACM Computing Surveys (CSUR) 47(4), 69.

[3]. Jin, G., Lin, J., Endo, T. (2014) Efficient utilization of memory hierarchy to enable the computation on bigger domains for stencil computation in CPU-GPU based systems. Proceedings of the International Conference on High Performance Computing and Applications.

[4]. Sîrbu, A., Babaoglu, O. (2016) Power consumption modeling and prediction in a hybrid CPU-GPU-MIC supercomputer. European Conference on Parallel Processing.

[5]. Bientinesi, P. (2015) Parallel computing on graphics processing units and heterogeneous platforms. Concurrency and Computation: Practice and Experience 27(6), 1525-1527.

[6]. Sawadsitang, S. (2015) Understanding Performance Portability of OpenACC for Supercomputers. IEEE International Parallel and Distributed Processing Symposium Workshop.

[7]. Fu, H. (2017) Solving global shallow water equations on heterogeneous supercomputers. PloS one 12(3).

Submitted 26.02.2018
Accepted 30.05.2018