SLAP: Making a Case for the Low-Powered Cluster by leveraging Mobile Processors

Dukyun Nam, Jik-Soo Kim, Hoon Ryu, Gibeom Gu, and Chan Yeol Park
Korea Institute of Science and Technology Information (KISTI), Daejeon, Republic of Korea
{dynam, jiksoo.kim, elec1020, gibeom.gu, chan}@kisti.re.kr

ABSTRACT
In this paper, we present our empirical study of building a low-powered cluster called SLAP by leveraging mobile processors. To investigate the reliability and usability of our system, we have conducted various performance analyses based on HPL benchmark, a real semiconductor engineering application and the application of distributed file system. Our experience can show potential benefits, possibilities and limits of applying energy efficient alternative solutions to supercomputing which can result in many interesting research issues.

Keywords
Low-powered cluster, Mobile processor, ARM, High-performance computing

1. Motivation and Experience of SLAP
In the operation of leadership-class supercomputing facilities, efficient management of power consumption is becoming a challenging issue that can affect not only the total cost but also the amount of space for cooling systems. Moreover, with the upcoming era of Exascale computers that are expected delivered in 2019-2020, we need to devise effective solutions that can fulfill 20MW of power consumption [1]. However, traditional multi-core processors will not be sufficient to satisfy the limited energy consumption requirement so that many-core coprocessors such as GPU or Xeon Phi should be appropriately utilized as we can see from the Top 500 list of supercomputers [2].

An alternative approach can be leveraging low-powered and lightweight mobile processors which have been mainly used in portable devices such as smartphones, tablet PCs. Although mobile processors are not originally designed to support High-Performance Computing, they are optimized to maintain sustainable performance under a certain degree of power consumption (i.e. batteries). Therefore, the European project called Mont-Blanc [3] has been conducted to design a new type of system built from energy efficient solutions used in embedded and mobile devices. Additionally, relatively low energy consuming ARM server processors are also becoming prevalent. Despite the potential computing powers of emerging mobile processors for HPC should be more investigated, effectively leveraging low-powered and lightweight processor technologies to build a supercomputing cluster can substantially minimize the overall system size along with the cooling device.

In this paper, we present our empirical study of building a small scale cluster called SLAP (Scalable, Low-powered, Autonomous & Robust, Pluggable) based on mobile processor technologies under different use cases and scenarios. Although mobile processor based cluster can have advantages in terms of physical space and energy efficiency, they have not mainly used for HPC servers so that we especially focus on the reliability and usability of our system.

We built a total 10 nodes of SLAP cluster (as we can see from Figure 1) where each node is an ODROID-XU board [4] which consists of Samsung Exynos 5410 SoC CPU, 2GB LPDDR3 RAM, 10/100 Ethernet Controller, USB 3.0 and 2.0 ports as seen from Figure 2.

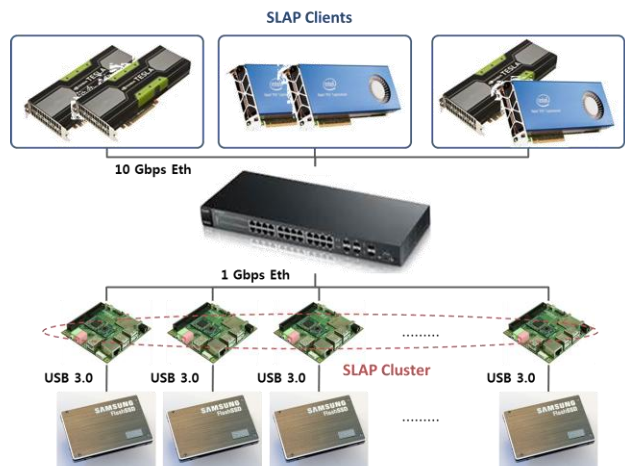

Figure 1: System Architecture of SLAP Cluster

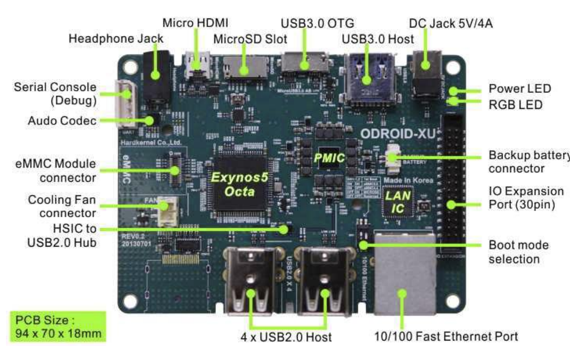

Figure 2: A Single System Board in the SLAP Cluster

Each SLAP cluster node utilizes eMMC 4.5 64GB and Samsung 500GB SSD (SATA 6Gb/s) connected through USB 3.0 as local storages. Figure 3 shows our final prototype version of SLAP cluster mounted on a half-size rack.

Figure 3: Final Prototype of SLAP Cluster in a Half Rack

We first measured the performance of a single node based on the HPL (High-Performance Linpack) benchmark [5] which is mainly used for ranking the Top 500 supercomputers [2]. In an idle state, a system board (Figure 2) consumes about 1.5W and when we performed HPL benchmark, it shows 2.739 GFlops/Watt which is similar to the 18th energy efficient supercomputer in the Green 500 list [6]. When we increase the number of nodes into total 10, benchmark result shows 47.44GFlops (83.25% compared to the linear scaling) while the energy efficiency drops into 1.35 GFlops/Watt mainly due to limited scalability of interconnect and some waste of energy other than CPU or Memory such as board, USB, NIC, cooling fans.

In addition to basic HPL benchmark tests, we also performed evaluations based on a real application from semiconductor engineering and the application of the GlusterFS [7] on top of SLAP nodes (as a compute cluster for HPC applications and a storage cluster based on distributed file system). The performance has been benchmarked with a parallel 3D Schrödinger-equation solver, where we have computed five eigenvalues of a complex, symmetric 393040x393040 matrix with the LANCZOS algorithm to find the optical gap of InAs/GaAs quantum dots [8]. While the computing time increases by 3~6 times, the energy consumption reduces by 5~107 times representing the potential strength of the SLAP cluster for scientific computing. Increasing MTU (Maximum Transmission Unit) improves the performance, causing reduction in both energy consumption and computing time due to less communication loads. The aggregated throughput of 10 storage nodes in SLAP cluster (based on GlusterFS) shows 5Gbyte/s which is very competitive compared to the results with traditional CPU architectures [9].

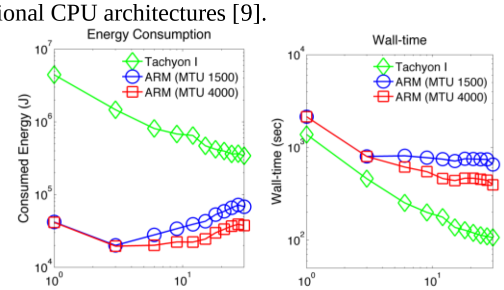

Figure 4: LANCZOS Algorithm Evaluation Results

Through our empirical study, we could find potential benefits, possibility along with limits of applying embedded technology for supercomputing area. However, as the mobile processor and ARM server technologies evolve, energy efficient processors can become a competitive alternative to the road of Exascale computing.

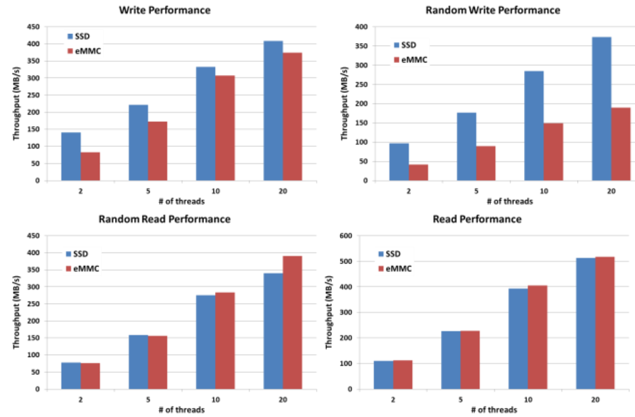

Figure 5: GlusterFS Evaluation Result

In the near future, we plan to build a SLAP-v2 cluster with 64-bit ARM processors as well. Lately we complete a preliminary design and thermal analysis of the next-version cluster which is a server appliance type in Figure 6.

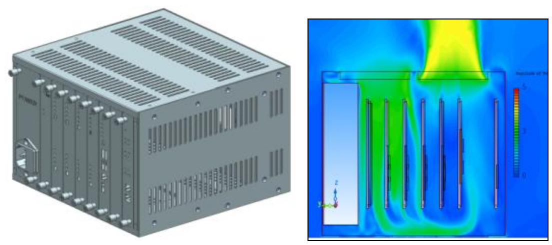

Figure 6: Preliminary Design and Thermal Analysis Result of a Server Appliance-type Cluster using 64-bit ARM Processors

2. REFERENCES

[1] Exascale Computing Study: Technology Challenges in Achieving Exascale Systems, in: *DARPA Technical Report 2008-13* (July 2008).

[2] TOP500 SUPERCOMPUTING SITES: http://www.top500.org/

[3] Rajovica, N., Ricoa, A., Puzovica, N., Adeniyi-Jonesc, C., and Ramirez, A. 2014. Tibidabo: Making the case for an ARM-based HPC system, *Future Generation Computer Systems* 36 (July 2014), 322–334.

[4] ODROID-XU board site: http://www.hardkernel.com/

[5] HPL - A Portable Implementation of the High-Performance Linpack Benchmark for Distributed-Memory Computers: http://www.netlib.org/benchmark/hpl/

[6] The Green500 List: http://www.green500.org/

[7] GlusterFS, a scalable network file system: http://www.gluster.org/

[8] Ryu, H., Nam, D., Ahn, B., Lee, J. R., Cho, K., Lee., S., Klimeck, G., and Shin. M. 2013. Optical TCAD on the Net: A tight-binding study of inter-band light transitions in self-assembled InAs/GaAs quantum dot photodetectors. *Mathematical and Computer Modeling* 58, 1-2 (July 2013), 288-299.

[9] Donvito, G., Marzulli, G. and Diacono, D. 2014. Testing of several distributed file-systems (HDFS, Ceph and GlusterFS) for supporting the HEP experiments analysis. *Journal of Physics: Conference Series 513* (July 2014).