Towards Understanding HPC-Big Data Convergence Using Cloud Platforms

Shweta Salaria*, Kevin Brown, Hideyuki Jitsumoto and Satoshi Matsuoka

Tokyo Institute of Technology, Japan

*Email: salaria.s.aa@m.titech.ac.jp

Introduction

- Trend of recent decade is to extend the functionality of HPC to solve Big Data problems.
- Problems in HPC centers: Limited access; Low availability of resources.
- Cloud platforms: Public access; Constantly evolving to meet the needs of HPC users.
- Amazon EC2 announced the next generation of HPC-oriented instances called C4 which have not been tested.

Objective

- To study the capabilities of new generation of Amazon EC2 HPC-like C4 instances [1] for running HPC and Big Data applications.
- To compare the performance of these C4 instances with a leading supercomputer, TSUBAME-KFC.

Experimentation

- We evaluated the latest generation of compute-optimized instances c4.8xlarge by Amazon EC2 with three different kinds of applications:
  1. The Intel MPI Benchmark [2]
  2. NICAM-DC-MINI [3]
  3. Graph 500 Benchmark [4]

.

	Amazon EC2 c4.8xlarge instance	TSUBAME-KFC
Cores	36 Virtual cores	12 Physical cores 24 Virtual cores
Memory (GiB)	60	64
Processor	Intel Xeon E5-2666v3 (Haswell) 2.9 GHz	Intel Xeon E5-2620v2 (Ivy Bridge) 2.10 GHz
Network	10 Gbps Ethernet	54 Gbps (4xFDR Infiniband)
Storage	Elastic Block Storage	NFS over 10 Gigabit Ethernet (30 TB)

Specifications of Amazon EC2 c4.8xlarge instances and TSUBAME-KFC

Results: Intel MPI Benchmark

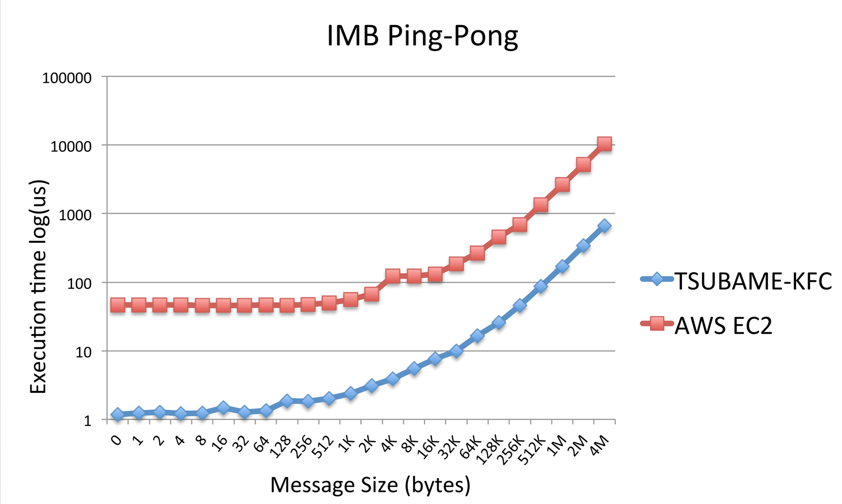

Ping-Pong: Significant performance gap on message latency between TSUBAME-KFC and AWS EC2.

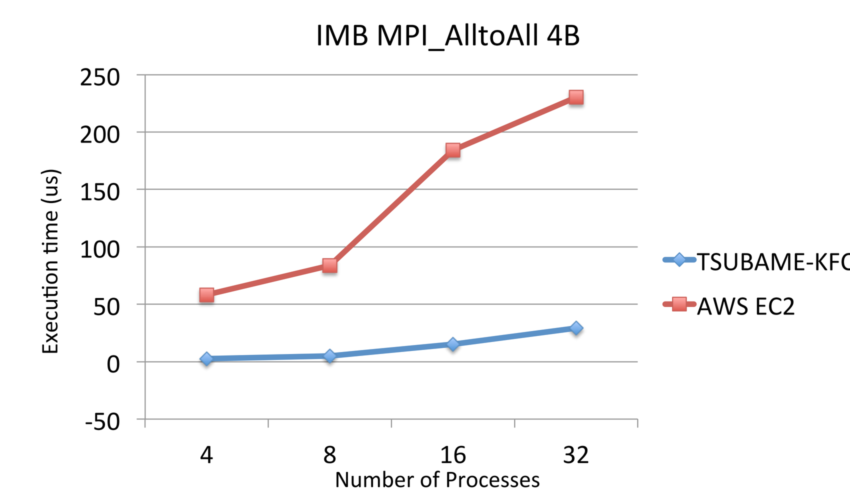

MPI_AlltoAll: When the message size is 4 bytes, AWS EC2 performance lags behind TSUBAME-KFC significantly and and the performance gap widens as the number of processes increases.

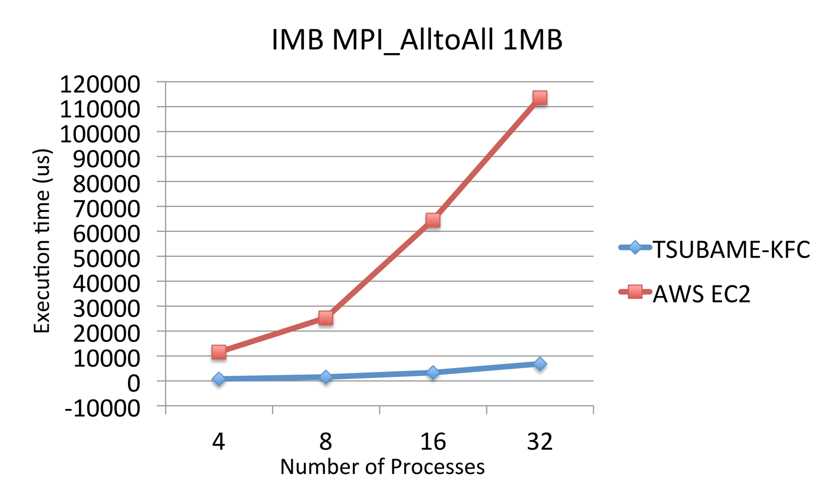

MPI_AlltoAll: When the message size is 1MB, TSUBAME-KFC continues to outperform AWS EC2. It exhibits the similar pattern with the performance gap rapidly growing with the increasing number of processes.

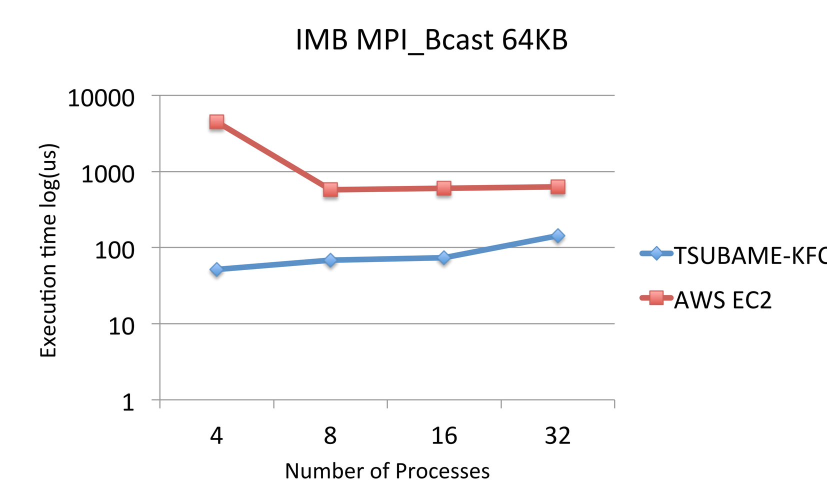

MPI_Bcast: With 64KB message size, still there is notable performance difference between TSUBAME-KFC and AWS EC2.

Results: NICAM-DC-MINI

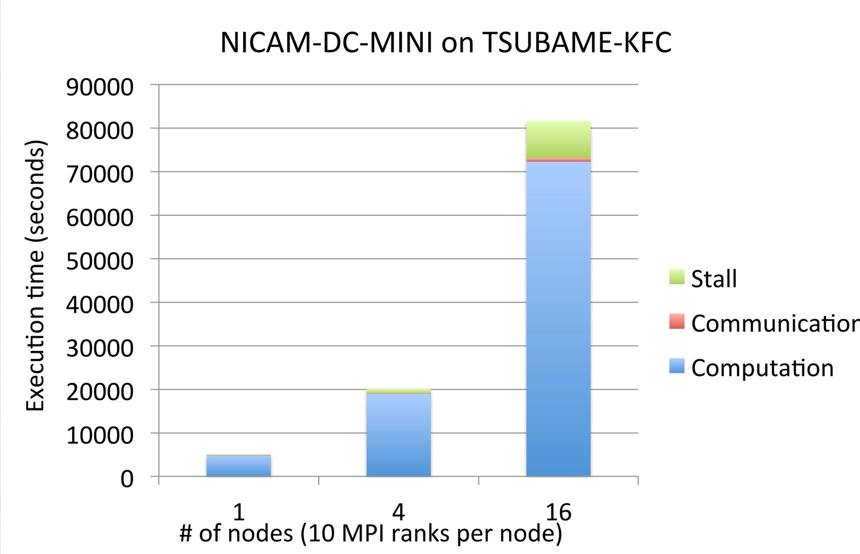

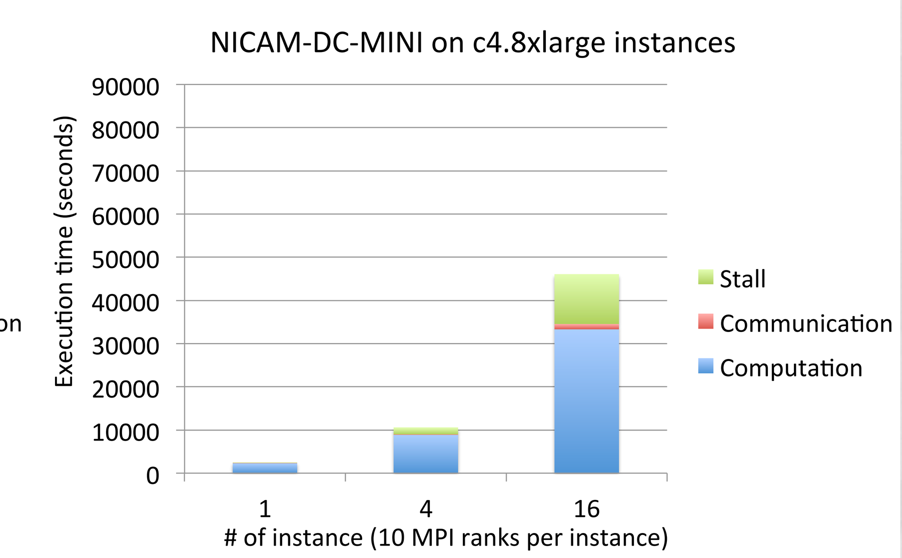

Since NICAM-DC-MINI is a compute-intensive miniapp and AWS EC2 c4.8xlarge instances have more powerful processors than TSUBAME-KFC, it performs 2x faster on AWS EC2 than TSUBAME-KFC. However, if we look at the communication performance, TSUBAME-KFC outperforms AWS EC2. The total execution time is dominated by the time spent on computation. Stall time corresponds to time spent in MPI_Waitall.

Results: Graph 500 Benchmark

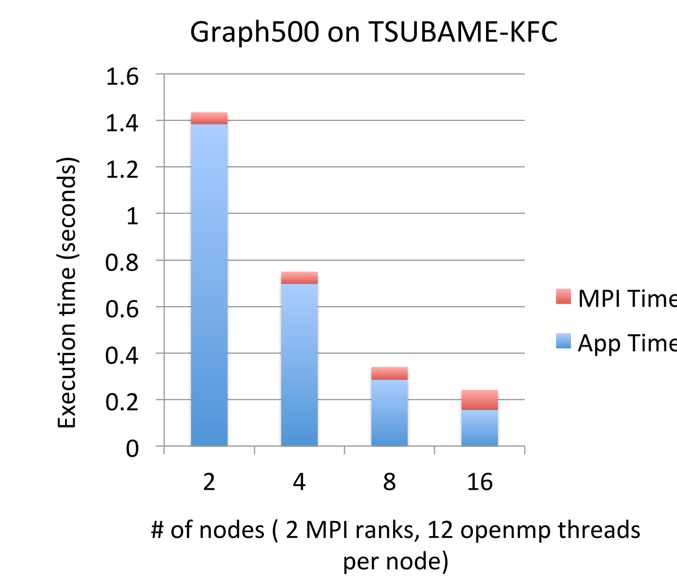

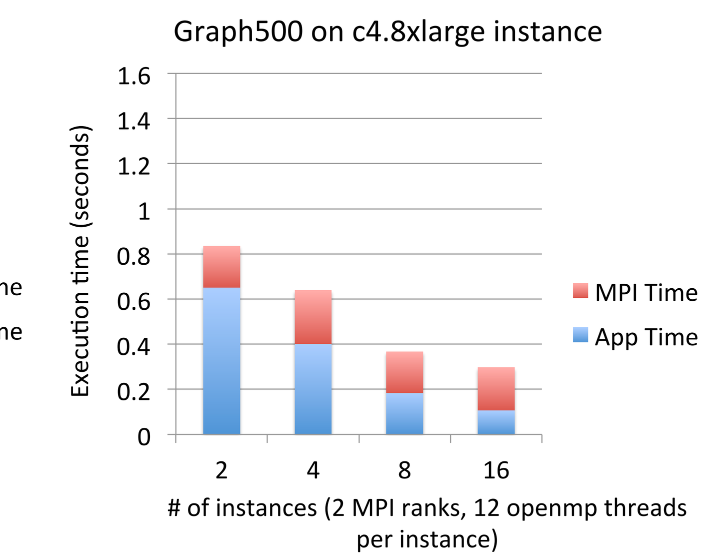

The above results corresponds to replicated-csr implementation of Graph 500 benchmark using scale 26. MPI time and App Time correspond to time spent on MPI calls and Non-MPI calls respectively during BFS phase of the benchmark. As expected, Non-MPI performance of AWS is better than that of TSUBAME-KFC. As the number of nodes increase, MPI time becomes more significant which drops the performance in AWS instances.

Conclusion and Future Work

AWS EC2 with new c4.8xlarge instances which are intended for high performance computing prove to be effective for compute-intensive applications with large message size exchanges among nodes. The 10-Gigabit Ethernet improves the bandwidth but exhibits high latency.

We plan to conduct extensive experimentation of the evaluated & new HPC Dwarfs and Big Ogres at larger scales with varying workloads. Also, we will establish the performance-cost tradeoff between C4 instances and TSUBAME-KFC.

References

1. Amazon EC2 C4 instances, https://aws.amazon.com/blogs/aws/new-c4-instances/
2. Intel MPI Benchmarks, https://software.intel.com/en-us/articles/intel-mpi-benchmarks
3. NICAM-DC-MINI, Fiber Miniapp Suite, https://github.com/fiber-miniapp/nicam-dc-mini
4. Graph 500 Benchmark, http://www.graph500.org/specifications

This research was partially supported by JST, CREST (Research Area: Advanced Core Technologies for Big Data Integration) and AWS Cloud Credits for Research Program.