Extreme Big Data

Next Generation Big Data Infrastructure Technologies Towards Yottabyte / Year

TOKYO TECH Pursuing Excellence

Extreme Big Data (EBD) Overview

Our project[1], called EBD, aims to achieve the convergence of extreme supercomputing and big data in order to cope with explosion of data from multiple sources such as massive numbers of sensors whose resolution is increasing exponentially, high resolution simulations generating huge data results, as well as evolution of social infrastructures that allow for "opening up of data silos", i.e., data sources being confined within an institution, much as how scientific data are being handled in the modern era as common asset openly accessible within and across disciplines. Our primary target proxy applications include metagenomics, social simulation, and climate simulation with real-time data assimilation. Based on these EBD co-design applications, we define future EBD convergent SW/HW architecture and system. We have several on-going collaboration work with RIKEN AICS, ORNL, LLNL, ETH and JST Graph CREST / Univ. Kyushu.

100,000 Times EBD "Convergent" System Architecture

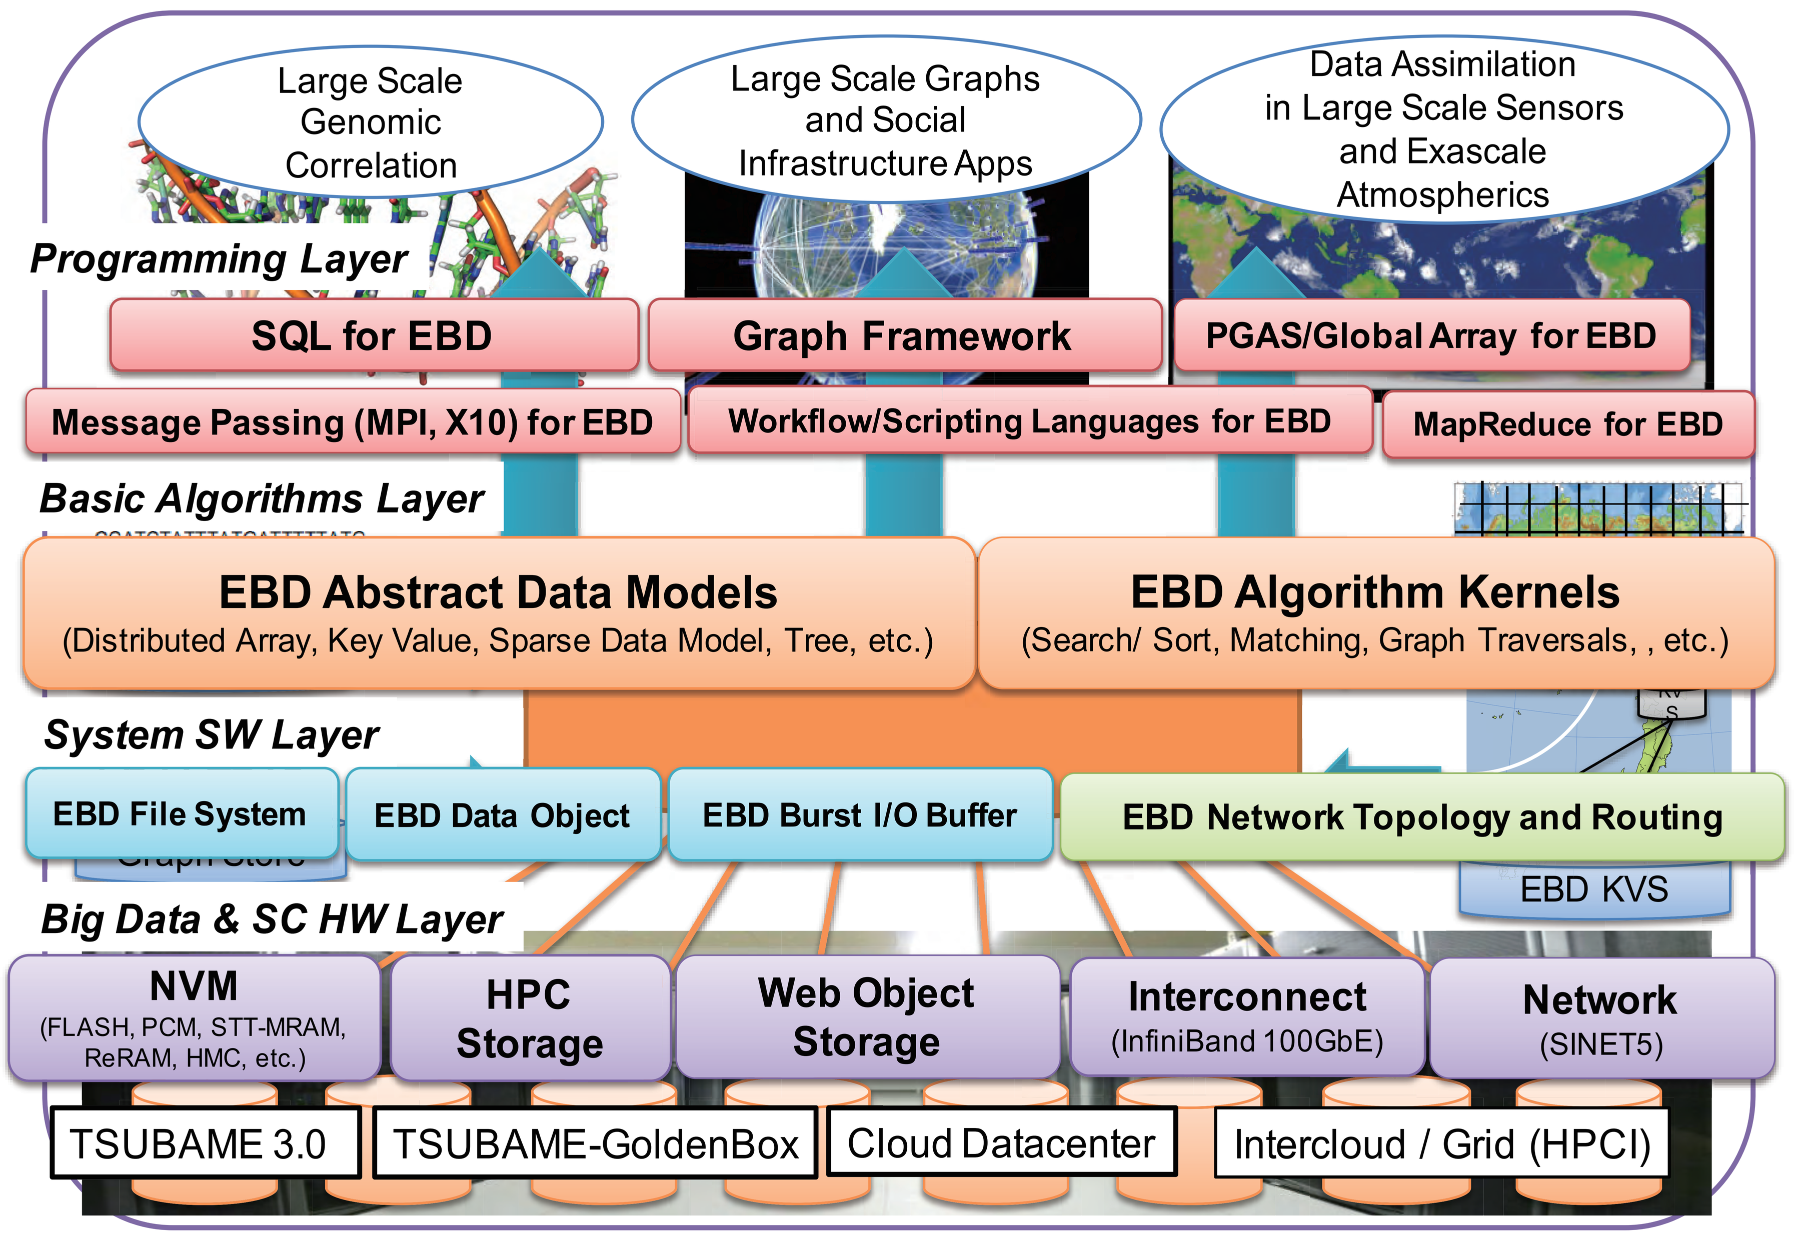

[1] S. Matsuoka, et al. "Extreme Big Data (EBD): Next generation big data infrastructure technologies towards yottabyte/year", Supercomputing frontiers and innovations, 2014.

Graph500

Collaboration work with JST Graph CREST / Univ. Kyushu and RIKEN AICS

We have developed extremely fast breadth first search (BFS) implementations for large-scale distributed environments and NVM-based hierarchical memory machines[2]. We have achieved several notable results on the Green500 and the Green Graph500, including becoming world #1 on the Graph500 (June 2014 and June 2015)[3] on K Computer and #1 on the Green Graph500 (November 2013)[4] on TSUBAME-KFC, based on our implementations.

History of Graph500 Scores on TSUBAME2

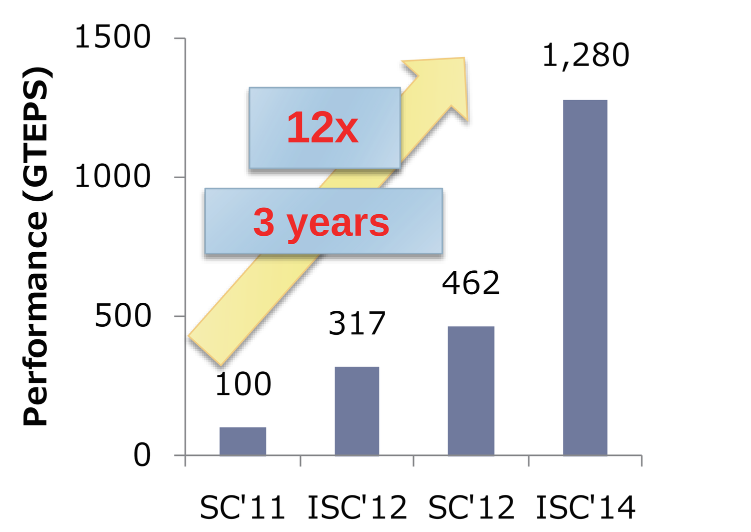

Graph500 Benchmark Performance

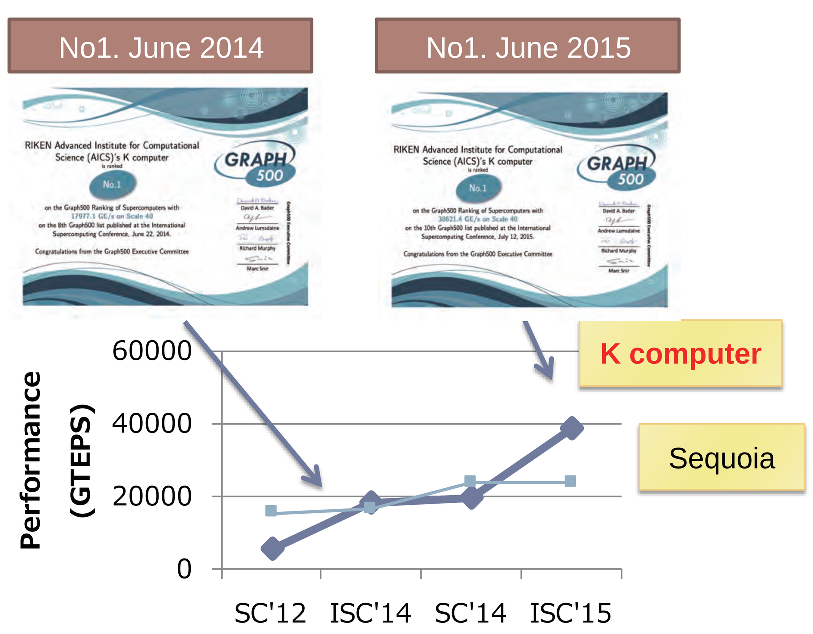

Optimization Techniques

Optimizations	SC11	ISC12	SC12	ISC14	ISC15
2D decomposition	✓	✓	✓	✓	✓
vertex sorting	✓				✓
direction optimization				✓	✓
data compression	✓	✓	✓		
bitmap graph representation				✓	✓
adaptive data representation				✓	✓
overlapped communication	✓	✓	✓	✓	✓
shared memory				✓	✓
GPGPU		✓	✓		

[2] K. Iwabuchi, H. Sato, Y. Yasui, K. Fujisawa, S. Matsuoka,"NVM-based Hybrid BFS with Memory Efficient Data Structure", IEEE BigData 2014.
[3] Graph500, http://www.graph500.org.
[4] Green Graph500, http://green.graph500.org.

GPU Sorting

We have developed several GPU-based sorting implementations, including the support for large-scale distributed environments over 1024 nodes[5], multi-level memory hierarchies with NVM, CPU and GPU[6], and variable length keys[7].

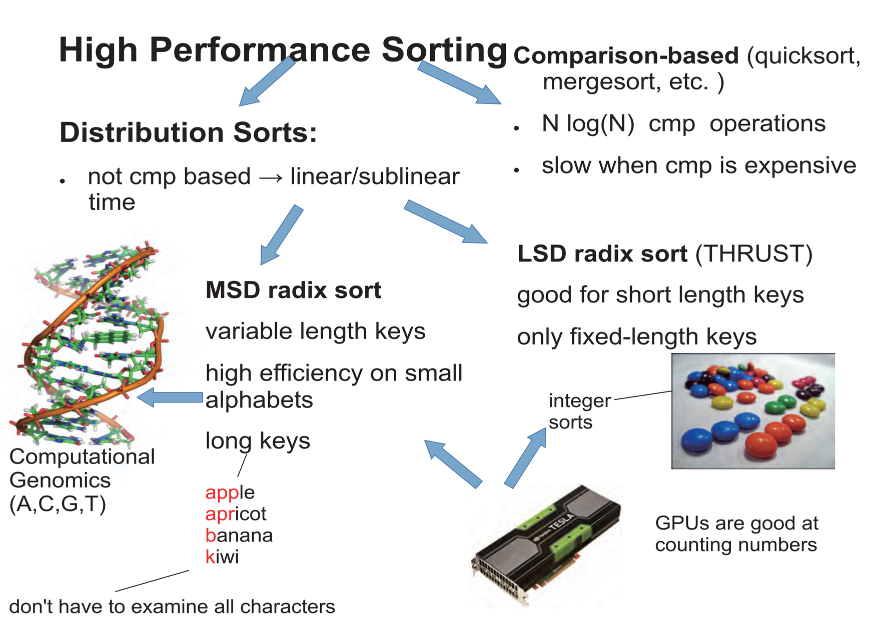

Single node/device sorting performance for variable length keys

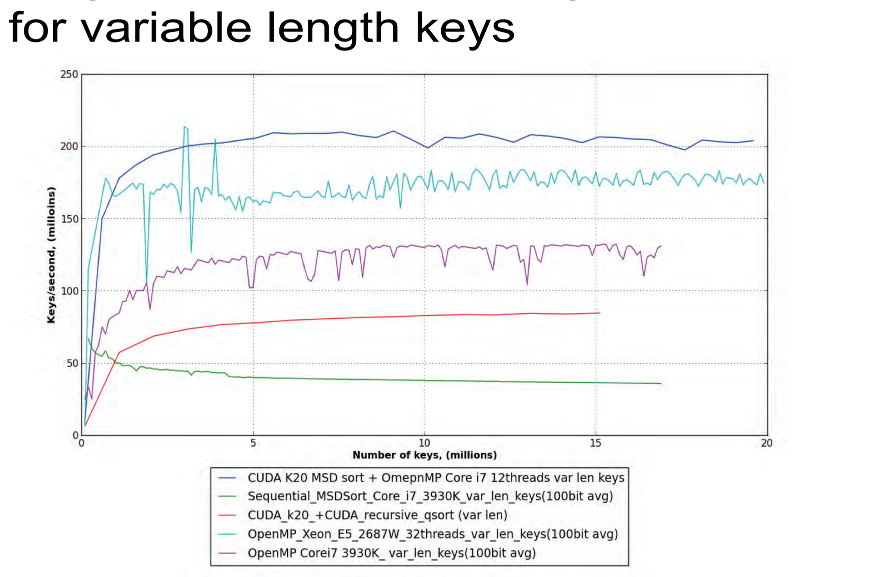

Large-scale GPU-based distributed sorting on TSUBAME2

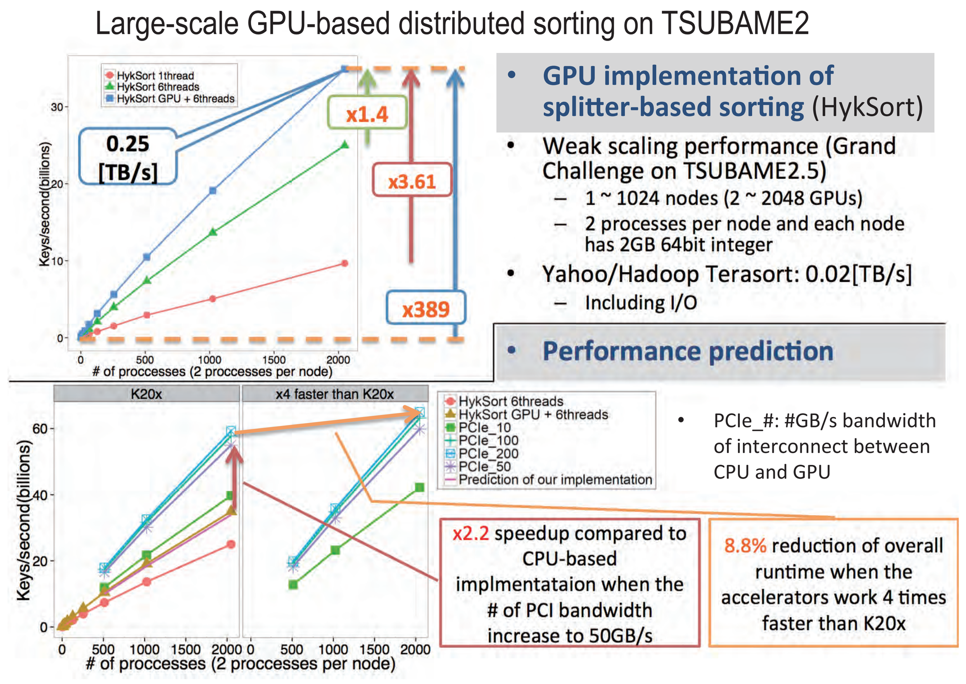

[5] H. Shamoto, K. Shirahata, A. Drozd, H. Sato, S. Matsuoka,"Large-scale Distributed Sorting for GPU-based Heterogeneous Supercomputers", IEEE BigData 2014.
[6] H. Sato, R. Mizote, S.Matsuoka, "Out-of-core Sorting Acceleration using GPU and Flash NVM.
[7] A. Drozd, M. Pericas, S. Matsuoka, "Efficient String Sorting on Multi- and Many-Core Architectures", BigData Congress 2014.

http://www.gsic.titech.ac.jp/sc15