GCLab

Power Usage in Data-Intensive Applications using MapReduce over MPI

Paula Olaya, Josh Davis, Joe Teague, Tao Gao, Yanfei Guo, Boyu Zhang, Pietro Cicotti, Yutong Lu, Pavan Balaji, Michela Taufer

Project Overview

Data analytics and data-intensive workloads are gaining representation at peta- and exascale. MapReduce has gained the most traction in the HPC community. **Mimir**, a novel **MapReduce over MPI** framework tackles skewed data, imbalance in memory usage, and loss in data scalability with combiner optimizations, dynamic repartitions, and a split method to handle datasets with superkeys. All this data movement **is power intensive** but little work is available in providing quantitative evaluations of these costs. This project quantitatively measures the impact of power capping on performance metrics such as runtime and power usage over time for data-intensive applications on top of a MapReduce over MPI framework when executed on HPC systems.

Mimir: a MapReduce over MPI framework

Optimizations	Benchmarks
• Combiner Optimizations • Dynamic Partition • Superkeys and Splitting	• Word count (WC) • Octree clustering (OC) • Breadth-first search (BFS) • Join

Imbalanced datasets (Tianhe-2 128 24-core nodes)

Value-mapping imbalance

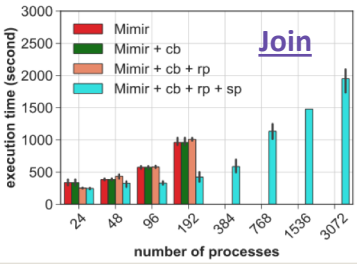

Key-mapping imbalance

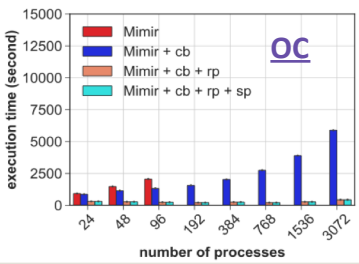

Power Capping on Mimir

Space of Interest:

We study the impact of HPC system, rate of unique words, and WC mini-apps on performance, measured by:

- Runtime (sec)
- Processor energy (joules)
- DRAM energy (joules)
- Power over time (watts)

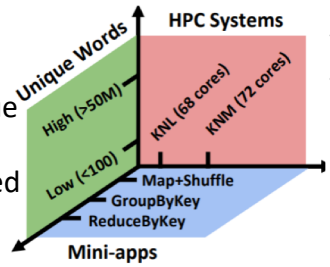

Without local combiner

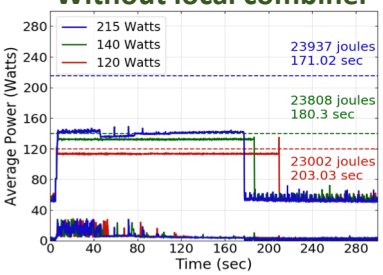

With local combiner

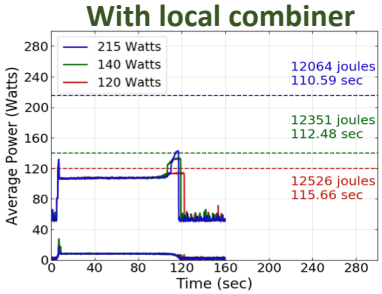

Case of Study

Mini-app: Reducebykey

HPC System: KNM

Local combiner optimizations save **up to 50%** energy

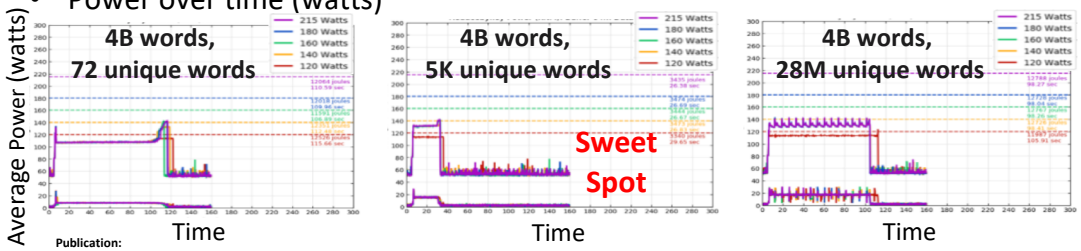

Lessons Learned

Combiner optimizations offer significant power performance benefits for our mini-apps.

Power usage gives insight into data management patterns indistinguishable by runtime performance.

Publication:

[1]T. Gao, Y. Guo, B. Zhang, P. Cicotti, Y. Lu, P. Balaji, and M. Taufer: **Mimir: Memory-Efficient and Scalable MapReduce for Large Supercomputing Systems.** IPDPS 2017: 1098-1108.

[2]J. Davis, T. Gao, S. Chandrasekaran, J. Heike, D. Anthony, B. Pavan, J. Dongarra, and M. Taufer. **Characterization of Power Usage and Performance in Data-Intensive Applications using MapReduce over MPI.** ParCo 2019.

THE UNIVERSITY OF TENNESSEE KNOXVILLE