Evaluation of Data-intensive Applications on Intel Knights Landing Cluster

Student: Tao Gao[1,2]
Advisors: Pietro Cicotti[3], Yanfei Guo[4], Pavan Balaji[4], Yutong Lu[2], Michela Taufer[1]
[1]University of Delaware, [2]National University of Defense Technology,
[3]San Diego Supercomputer Center, [4]Argonne National Laboratory

***Abstract*—Analyzing and understanding large datasets on high performance computing platforms is becoming more and more important in various scientific domains. MapReduce is the dominant programming model for processing these datasets. Platforms for data processing are empowered by many-core nodes with cutting-edge processing units. Intel Knights Landing (KNL) is the new arrival in the field. However, this new architecture has not been fully evaluated for data-intensive applications. In this poster, we present the assess of KNL on the performance of three key data-intensive applications based on a high-performance MapReduce programming framework on the latest KNL-cluster, Stampede2. We focus on the impact of different KNL memory models, we compare Stampede2 with other clusters such as Tianhe-2 and Mira, and we measure the scalability of large datasets. We observe how KNL-based clusters are a promising architecture for data-intensive applications. We also identify key aspects to enable more efficient usage of KNL-based clusters.**

I. Motivation and Contributions

Analyzing large volumes of data is becoming increasingly important in various scientific computing domains (e.g., genome analysis). Among the programming models for data analytics, MapReduce [1] has gained the most attractions due to its simplicity and scalability. Promoted by some successful implementations (e.g., Spark and Hadoop), MapReduce programming model has been successfully used in various domains and has become the `de facto` standard for big data processing.

At the same time, HPC platforms increasingly use many-core chips (e.g., GPU and Intel Xeon Phi) to gain high performance and low power consumption. Accelerator-based architectures (i.e., in which many-core chips are used as accelerators) have been successfully used in HPC domains. However, this architectures often are not suitable for data-intensive applications because the data transfer overhead between accelerators and host CPUs is often a bottleneck. Emerging trends move towards using the many-core processors as the host node directly. Clusters based on Intel Knights Landing (KNL) [3] represents the latest trend update.

As scientists start moving their data-intensive applications to platforms such as Stampede2, it is important to understand if this architecture is indeed suitable for these applications. This is still an open research question. To the best of our knowledge, no previous work has performed a fully evaluation of data-intensive applications on KNL clusters. This poster presents our evaluation for data-intensive applications on a large-scale KNL cluster.

The contributions of this poster are as follows.

1) We perform the full evaluation of three different data-intensive applications on a KNL cluster such as Stampede2.
2) We compare and contrast the KNL performance of the selected data-intensive applications with other cutting-edge clusters (i.e., Tianhe-2 and Mira).
3) We update the state of the practice on KNL for data-intensive applications based on our evaluation.

II. Evaluation Methods

For our evaluation, we used three proxy benchmarks: *WordCount* (WC), K*-mer Counting* (KC), and *Breadth-first Search* (BFS). WC counts the number of occurrences of each unique word in given input files. KC is a fundamental operation in genome analytics. BFS is a graph traversal algorithm that generates a tree rooted at a source vertex. These benchmarks are chosen to represent different kinds of data-intensive applications.

We implemented these three benchmarks with the high-end MPI-based MapReduce framework, Mimir [2]. Mimir is optimized for memory-efficiency and has been proved to scale to large-scale systems. The workflow of Mimir is shown in Figure 1.

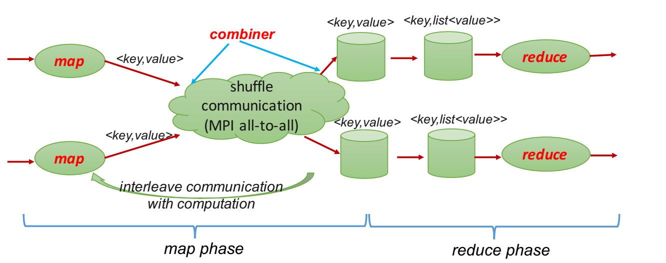

Fig. 1: Workflow of the high-end MPI-based MapReduce framework Mimir.

Our tests were performed on the XSEDE cluster Stampede2. Each KNL has 68 cores running at 1.4 GHz and each core supports 4 hardware threads. Each node has 96 GB of DDR4 memory and 16 GB of Multi-Channel Dynamic Random Access Memory (MCDRAM). MCDRAM can be configured as the level 3 cache (cache model) or as the direct access memory (flat model). Other clusters considered were Tianhe-2 and Mira.

III. Results Analysis

In our state of the practice analysis we answer four key questions: (1) How does one KNL node perform as the dataset size increases? (2) Can data-intensive applications benefit from MCDRAM? (3) How does one KNL-based node perform for data-intensive applications compared with other platforms? (i.e., Intel CPU node in Tianhe-2 and IBM CPU node in Mira)? and (4) Can data-intensive applications scale to large number of nodes on the KNL-based cluster?

Because of space constraint, we only show the results for WC or BFS and present our key observations here. The full set of results can be found in the poster.

A. Single-node Scalability

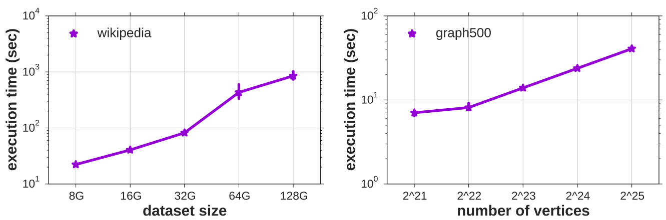

(a) WC (b) BFS

Fig. 2: Single-node performance of different dataset sizes.

As shown in Figure 2, one KNL node gains linear or close to linear performance as the dataset size increases.

B. Impact of Memory Model

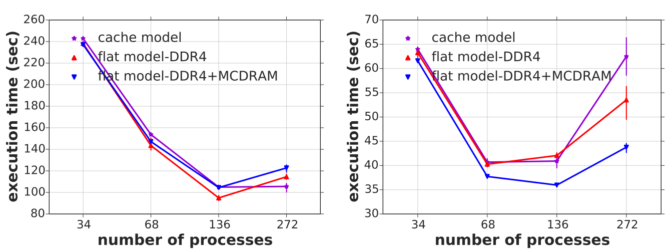

(a) WC (b) BFS

Fig. 3: Performance of different memory model and process count.

As shown in Figure 3, the KNL's cache model perform poorly due to high memory access latency while the flat model is suitable for data-intensive applications (i.e., it provides more memory to use, its DDR has lower memory access latency in flat model). Data-intensive applications benefits from MCDRAM by putting latency-bound structure on DDR and bandwidth-bound structure on MCDRAM (e.g., BFS).

C. Comparison with other platforms

Figure 4 shows a single KNL node' performance; KNL better performs than a single Intel Xeon node on Tianhe-2 (two sockets) and much better than a single IBM node on Mira (16 cores). Thus, KNL is a promising platform for data-intensive applications.

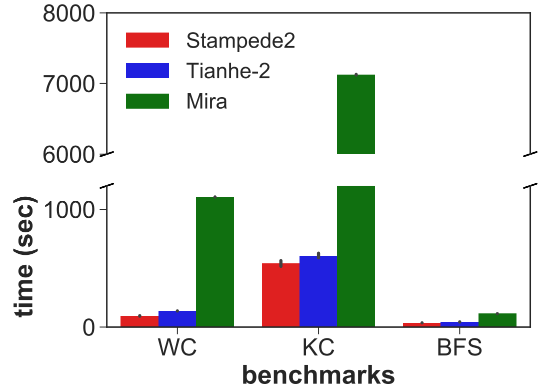

Fig. 4: Single-node performance compared with other platforms.

D. Large-scale Scalability

Figure 5 shows the scalability of our benchmarks up to at least 128 nodes (i.e., 8704 cores). However, the high global synchronization and communication (due to I/O variability and imbalance load) still limits the scalability starting at 32 nodes.

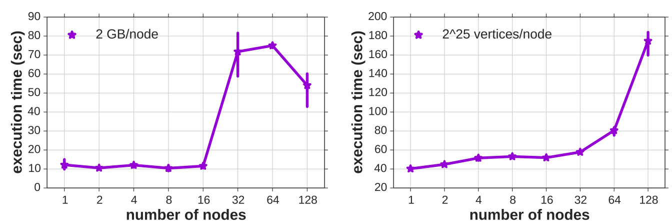

(a) WC (b) BFS

Fig. 5: Weak scalability.

IV. Conclusions

Our evaluation shows that KNL is a promising platform for data-intensive applications. To maximize the KNL benefits, applications should integrate these important aspects: (1) make use of MCDRAM; (2) provide I/O optimizations to adapt the performance variability; (3) minimize the load imbalance in the applications; and (4) optimize global communications.

References

[1] J. Dean and S. Ghemawat. MapReduce: Simplified Data Processing on Large Clusters. *Communications of the ACM*, 51(1):107–113, 2008.

[2] 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. In *Proceedings of the 31th IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, 2017.

[3] Avinash Sodani. Knights landing (knl): 2nd generation intel® xeon phi processor. In *Hot Chips 27 Symposium (HCS), 2015 IEEE*, pages 1–24. IEEE, 2015.