Accepted Manuscript

Marcher: A Heterogeneous System Supporting Energy-Aware High Performance Computing and Big Data Analytics

Ziliang Zong, Rong Ge, Qijun Gu

PII: S2214-5796(16)30048-X
DOI: http://dx.doi.org/10.1016/j.bdr.2017.01.003
Reference: BDR 55

To appear in: *Big Data Research*

Received date: 29 April 2016
Revised date: 18 November 2016
Accepted date: 2 January 2017

Please cite this article in press as: Z. Zong et al., Marcher: A Heterogeneous System Supporting Energy-Aware High Performance Computing and Big Data Analytics, *Big Data Res.* (2017), http://dx.doi.org/10.1016/j.bdr.2017.01.003

Marcher: A Heterogeneous System Supporting Energy-Aware High Performance Computing and Big Data Analytics

Ziliang Zong[1], Rong Ge[2], Qijun Gu[1]

[1]Computer Science Department, Texas State University

[2]School of Computing, Clemson University

Abstract - Excessive energy consumption is a major constraint in designing and deploying the next generation of supercomputers. Minimizing energy consumption of high performance computing and big data applications requires novel energy-conscious technologies (both hardware and software) at multiple layers from architecture, system support, and applications. In the past decade, we have witnessed the significant progress toward developing more energy-efficient hardware and facility infrastructure. However, the energy efficiency of software has not been improved much. One obstacle that hinders the exploration of green software technologies is the lack of tools and systems that can provide accurate, fine-grained, and real-time power and energy measurement for technology evaluation and verification. Marcher, a heterogeneous high performance computing infrastructure, is built to fill the gap by providing support to research in energy-aware high performance computing and big data analytics. The Marcher system is equipped with Intel Xeon CPUs, Intel Many Integrated Cores (Xeon Phi), Nvidia GPUs, power-aware memory systems and hybrid storage with Hard Disk Drives (HDDs) and Solid State Disks (SSDs). It provides easy-to-use tools and interfaces for researchers to obtain decomposed and fine-grained power consumption data of these primary computing components. This paper presents the design of the Marcher system and demonstrates the usage of Marcher power measurement tools to obtain detailed power consumption data in various research projects.

Keywords –energy efficient high performance computing; energy-aware big data analytics; power-measurable systems; power profiling

1. INTRODUCTION

Excessive energy consumption is now a first-class constraint in designing and deploying the next generation of supercomputers. For example, the servers and data centers used about 2% of U.S. energy in 2011, which was almost equal to the annual energy usage of 7.7 million U.S. households [21]. In 2015, the average power consumption of the top 10 supercomputers was 8.2 MW, which cost about $8 million per year per system [23] [24]. Tianhe-2, the most powerful supercomputer in 2015, consumes approximately 17.8 MW of power [24] at peak performance, which corresponds to 95,000 metric tons of annual CO_2 emissions [22]. As a result, energy efficiency has become a top concern for HPC systems.

In recent years, we have witnessed the great progress in producing more energy-efficient hardware, e.g., low power CPUs, Intel Many Integrated Cores (Xeon Phi), GPGPUs, and Solid State Disks (SSDs). However, the energy efficiency of software running on high performance computing (HPC) platforms was not improved much due to the lack of systems and tools that can provide accurate power consumption data of all major components including CPUs, DRAMs, disks, accelerators and coprocessors. Direct fine-grained power measurements are necessary to identify energy inefficiencies and bottlenecks in software, and evaluate the impact of different software designs and optimizations on performance and energy. Although many exiting HPC systems can measure power at the system or rack level, very few enable power measurement at the granularity level that can provide valuable information