J. Parallel Distrib. Comput. 107 (2017) 87–100

Contents lists available at ScienceDirect

J. Parallel Distrib. Comput.

journal homepage: www.elsevier.com/locate/jpdc

Analysis of the efficiency characteristics of the first High-Temperature Direct Liquid Cooled Petascale supercomputer and its cooling infrastructure

Hayk Shoukourian [a,*], Torsten Wilde [a], Herbert Huber [a], Arndt Bode [a,b]

[a] *Leibniz Supercomputing Centre (LRZ) of the Bavarian Academy of Sciences and Humanities, Boltzmannstraße 1, 85748 Garching bei München, Germany*

[b] *Technische Universität München (TUM), Fakultät für Informatik I10, Boltzmannstraße 3, 85748 Garching bei München, Germany*

HIGHLIGHTS

- Impacts of water cooling temperatures on system performance and power consumption.
- Node power variability at large scale (Intel Sandy Bridge and Intel Haswell).
- Impacts of new Intel power and/or temperature capping on node performance.
- Impacts of AVX/AVX2 on application with different power limiting configurations.
- Data center's cooling efficiency at different water cooling temperatures.

ARTICLE INFO

Article history:
Received 18 January 2017
Received in revised form
11 April 2017
Accepted 17 April 2017
Available online 27 April 2017

Keywords:
HPC
Energy-efficiency
System characteristic
Cooling efficiency
High-Temperature Direct Liquid Cooling (HT-DLC)

ABSTRACT

SuperMUC, deployed at the Leibniz Supercomputing Centre, is the first High-Temperature (ASHRAE W4 chiller-less) DirectLiquid Cooled (HT-DLC) Petascale supercomputer installed worldwide. Chiller-less direct liquid cooling can save data centers a substantial amount of energy by reducing data center cooling overheads. An essential question remains unanswered — how to determine an optimal operational environment for balancing scientific discovery with the energy consumption of both the supercomputer and the cooling infrastructure? This paper shows, for the first time, how the new technologies (HT-DLC and chiller-less cooling) influence the performance and energy/power efficiency of large-scale HPC applications and how different inlet temperatures affect the overall system power consumption and the HT-DLC efficiency.

1. Introduction

Energy efficiency improvements of High Performance Computing (HPC) data centers continue to remain an active research area [24,25,14,34,33,4], mainly due to the continuous growth in energy consumption and associated costs [6]. Fig. 1 shows the trend of the energy costs at Leibniz Supercomputing Centre (LRZ) from 2000 till 2016 including an estimate for 2017. As can be seen, the energy costs keep rising. This is partly due to the German "Energie Wende" where additional taxes and fees impact the yearly energy price which is projected to increase further also due to the environmental improvements and protection. The other part is the increased power consumption over the generations of supercomputers (shown in Fig. 2). Even though each new generation of HPC system provides better power efficiency (i.e. better performance/watt ratio) the density and overall performance of the system increases substantially as well leading to an overall increased power/energy consumption.

Being energy efficient requires a data center to look at all data center areas affecting its power consumption. Fig. 3 shows the 4 Pillar Framework (see [40] for a full description) which was developed by LRZ to provide a common frame of reference identifying the major data center areas affecting the overall power/energy consumption. The 4 Pillars are: *Building Infrastructure*; *HPC System Hardware*; *HPC System Software*; and *Applications*. LRZ believes that in order to improve a data centers energy efficiency one needs to take a wholistic approach which includes all the 4 Pillars, as well as data center requirements and mission goals, combined with outside influences and constraints.

* Corresponding author. Fax: +49 89 35831 9700.
E-mail address: hayk.shoukourian@lrz.de (H. Shoukourian).

http://dx.doi.org/10.1016/j.jpdc.2017.04.005
0743-7315/

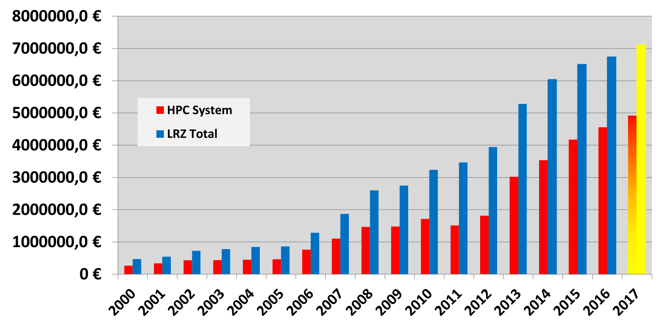

Fig. 1. Trend of energy costs for LRZ from 2000 till 2016, and prediction for 2017.

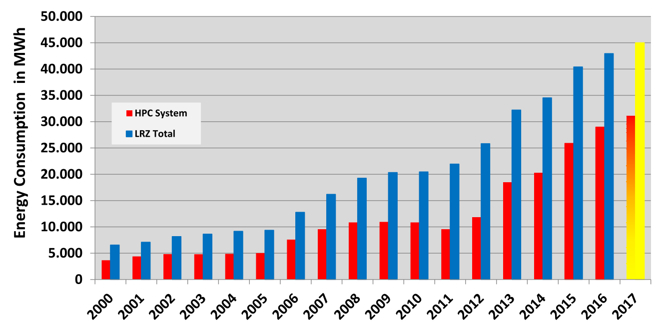

Fig. 2. Trend of energy consumption for LRZ from 2000 till 2016, and prediction for 2017.

LRZ sets itself the goal of supporting advanced scientific discovery by running large scale parallel applications fast and energy efficiently. To improve energy efficiency, LRZ decided in 2009 to use High-Temperature (chiller-less) Direct Liquid Cooling (HT-DLC) in its data center and up to the time of this publication, SuperMUC (Section 2) is still one of the largest HT-DLC cooled HPC system installation in the world.

Another vital part when running a HPC data center is the characterization of the deployed HPC systems for the selection of an optimal operational environment. Looking at each pillar individually is a first step toward a global data center energy optimization strategy. Optimizing individual pillars can be contrarious, for example the use of chiller-less direct liquid cooling with its increased cooling temperature increases leakage currents. HT-DLC is more efficient with higher cooling temperatures, but the increase in leakage currents can negate the efficiency gain from the infrastructure (Pillar 1) when taking the IT power consumption (Pillar 2) into account. Since LRZ is operating two large scale HT-DLC systems (SuperMUC Phase1 and Phase2 described in Section 2) with different processor architectures, this paper presents a unique first look at:

- the impacts of water cooling temperatures on system node performance as well as power and energy consumption (***Pillar 1*** and ***Pillar 2***—Section 3.1);
- the detailed system compute node power variability using different hardware power limits and HPC applications (***Pillar 2***, ***Pillar 3***, and ***Pillar 4***—Sections 3.2 and 3.3);
- the efficiency of the heat removal from compute nodes at different inlet temperatures for two different compute node design approaches (***Pillar 1*** and ***Pillar 2***—Section 3.4);
- power consumption characteristics of a real world application using different Advanced Vector Extension (AVX) versions on platforms with different enforced power consumption limits (***Pillar 2***, ***Pillar 3***, and ***Pillar 4***—Section 3.5);
- the impacts of data center's cooling infrastructure efficiency at different inlet and return temperatures correspondingly flowing to and from the HPC systems (***Pillar 1*** and ***Pillar 2***—Section 4).

Most of the related works looked at characterizing a single CPU or a limited number of compute nodes. For example [31,12] describe energy/power efficiency characteristics of new generation Intel processor instances. This work will look at energy/power efficiency characteristics by outlining the efficiency fluctuations on two large-scale systems and the essential impacts brought by the supporting cooling infrastructure.

2. Description of SuperMUC system

SuperMUC, operated by the Leibniz Supercomputing Centre (LRZ) [26], consists of two installation phases, called Phase1 and Phase2, shown in Fig. 4. SuperMUC has a combined theoretical

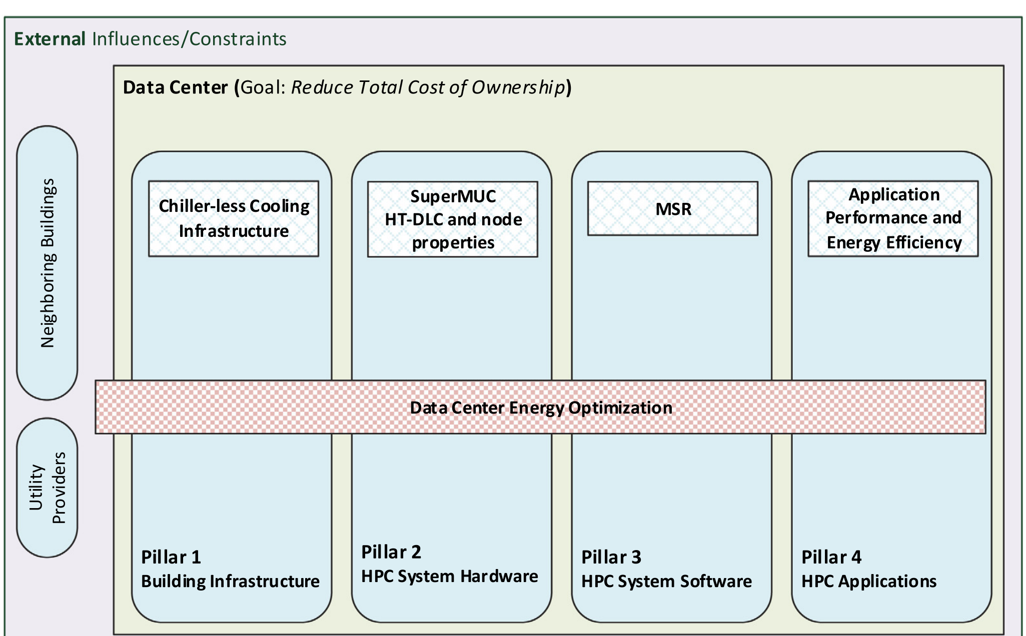

Fig. 3. 4 Pillar framework and the wholistic energy efficiency approach at LRZ.