Simulation Modelling Practice and Theory 58 (2015) 40–54

Contents lists available at ScienceDirect

Simulation Modelling Practice and Theory

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

Energy and thermal models for simulation of workload and resource management in computing systems

Wojciech Piątek [a,*], Ariel Oleksiak [a,b], Georges Da Costa [c]

[a] Poznan Supercomputing and Networking Center, Noskowskiego 10, Poznan, Poland
[b] Institute of Computing Science, Poznan University of Technology, Poland
[c] Institute of Computer Science Research, University of Toulouse, France

ARTICLE INFO

Article history:
Available online 16 May 2015

Keywords:
Simulations
Thermal models
Energy-efficiency
Data centers

ABSTRACT

In the recent years, we have faced the evolution of high-performance computing (HPC) systems towards higher scale, density and heterogeneity. In particular, hardware vendors along with software providers, HPC centers, and scientists are struggling with the exascale computing challenge. As the density of both computing power and heat is growing, proper energy and thermal management becomes crucial in terms of overall system efficiency. Moreover, an accurate and relatively fast method to evaluate such large scale computing systems is needed. In this paper we present a way to model energy and thermal behavior of computing system. The proposed model can be used to effectively estimate system performance, energy consumption, and energy-efficiency metrics. We evaluate their accuracy by comparing the values calculated based on these models against the measurements obtained on real hardware. Finally, we show how the proposed models can be applied to workload scheduling and resource management in large scale computing systems by integrating them in the DCworms simulation framework.

1. Introduction

One of the main obstacles in building and cost-effective use of large scale HPC systems is their significant energy consumption. Thus, appropriate energy consumption modeling, monitoring, and optimization is essential to enable a fast growth of HPC systems computing power and to decrease their operational costs, which should pave the way for the exascale HPC systems.

Importantly, optimizing energy consumption of processors and servers is only part of the equation. Data center infrastructure (with the major contribution of cooling) may consume even the same amount of energy as IT equipment. To include cooling into optimization of the computing system, thermal aspects must be taken into consideration. These aspects influence both local energy consumption of servers and global consumption of the whole data center. Impact on local energy consumption is caused by differences in operation of fans and changes in power usage of CPUs dependent on temperature. Since the energy usage has a strong correlation with the thermal efficiency of the computing system [2], it is valuable, for its effectiveness, to take both power and temperature distribution into consideration while managing the data center equipment. On a global level the amount and distribution of heat dissipated by servers affect the overall energy consumption by the cooling system. Moreover thermal-aware management is crucial not only from the perspective of energy costs, but also for the

---

[*] Corresponding author. Tel.: +48 618582169; fax: +48 618582151.
E-mail address: piatek@man.poznan.pl (W. Piątek).

http://dx.doi.org/10.1016/j.simpat.2015.04.008
1569-190X/

system maintenance, as excess heat can cause downtimes and shorten equipment durability. These issues lead also to additional costs of system operation and upgrades. Thus, energy consumption and heat dissipation transforms obviously into hard limits of HPC system development due to maximal power constraints as well as maintenance costs.

Often experiments on management of large scale computing systems are difficult or not possible to conduct in real environments. Especially for new planned centers, new computing architectures, and future exascale systems simulation studies are needed in advance to propose optimization of workload and resource management. Hence, simulation models and tools are needed to study architectures of new computing systems, their management, cooling and application deployment. To enable realistic simulation of such systems many challenges must be met including on one hand dealing with large scales and sizes of data, and on the other hand, taking into consideration specific aspects and technologies of new architectures such as special focus on power usage and heat dissipation, various heterogeneous and low power architectures, diversity of cooling techniques, and higher probability of failures.

In this paper we propose models and simulation environment to address the problem of simulation of large scale computing systems including energy-efficiency and thermal aspects. The main contributions of this paper include: (i) power usage and thermal models of air-cooled servers along with validation on real hardware, (ii) approach to enabling transient simulations of dynamic systems (allowing studying on-line reactions on temperature and dynamic cooling control), and (iii) demonstration of the use of models in management policies: showing effect of fan management, managing power and cooling capacity, and taking into account characteristics of temperature changes in time. All these contributions are applied to the DCworms simulator [18].

The work presented in his paper is focused on the airflow cooling approach. Proposed models should easily cover both air and heatsink-based solutions as well as any combination of them. However, nowadays there exist alternatives including variety of gaining popularity liquid-based approaches. Although the methods below cannot be directly applied to address other cooling techniques, they can provide a good basis for such extension.

The remaining part of this paper is organized as follows. In Section 2 we give a brief overview of the current state of the art concerning power and thermal models. Section 3 presents the models we propose to use in the simulation of energy consumption and temperature of computing systems. In Section 4 we assess the proposed models by comparison to results of experiments performed on a real hardware. In Section 5 we demonstrate the application of these models to a few types of resource and workload management policies in the DCworms simulator. Final conclusions and directions for future work are given in Section 6.

2. Related work

Understanding thermal phenomena and then describing it in the analytical way has been the subject of intense studies for decades. Modeling such processes has gained significant interest in the area of physics and recently also in the scope of computer science. Providing accurate models is important not only in terms of the design of IT and non-IT devices but also in terms of evaluation of the system efficiency and management policies. Moreover, temperature depends on several factors, like time-varying IT heat loads, physical room layout and performance of cooling facilities. Rapid growth of computing capabilities starting from single servers, to typical data centers and to the ultrascale systems, has make it even more challenging with respect to the precision and time complexity of the simulation tools benefiting from these models.

In order to characterize the thermal distribution, there are several approaches that are differentiated by their accuracy and required model size. Fig. 1 gives the general overview [9].

For now, Computation Fluid Dynamics (CFD) simulations are considered as the most accurate approach. However, they require lots of effort, while preparing model and even more time to obtain rewarding results, which makes it expensive to use in terms of big systems simulations. As an alternative, Potential Flow Model (PFM) has been proposed [11], that benefits from the reduction of the model. Another approach follows proper orthogonal decomposition methodology [10] and corresponds to Reduced Order Models entity in Fig. 1.

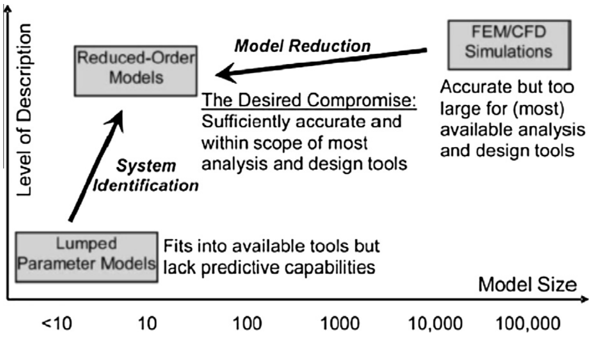

Fig. 1. Existing approaches complexity [9].