



ELSEVIER

Contents lists available at ScienceDirect

Ad Hoc Networks

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

Energy-efficient, thermal-aware modeling and simulation of data centers: The CoolEmAll approach and evaluation results



Leandro Cupertino^a, Georges Da Costa^a, Ariel Oleksiak^b, Wojciech Piątek^b,
Jean-Marc Pierson^a, Jaume Salom^c, Laura Sisó^c, Patricia Stolf^{a,*}, Hongyang Sun^a,
Thomas Zilio^a

^a Toulouse Institute of Computer Science Research (IRIT), University of Toulouse, France

^b Poznan Supercomputing and Networking Center (PSNC), Poznan, Poland

^c Catalonia Institute for Energy Research (IREC), Barcelona, Spain

ARTICLE INFO

Article history:

Received 31 March 2014

Received in revised form 26 September 2014

Accepted 3 November 2014

Available online 18 November 2014

Keywords:

Data centers

Energy efficiency

Metrics

Resource management policies

Scheduling

ABSTRACT

This paper describes the CoolEmAll project and its approach for modeling and simulating energy-efficient and thermal-aware data centers. The aim of the project was to address energy-thermal efficiency of data centers by combining the optimization of IT, cooling and workload management. This paper provides a complete data center model considering the workload profiles, the applications profiling, the power model and a cooling model. Different energy efficiency metrics are proposed and various resource management and scheduling policies are presented. The proposed strategies are validated through simulation at different levels of a data center.

© 2014 Elsevier B.V. All rights reserved.

1. Introduction

IT energy impact is now a major concern from the economical point of view but also from the sustainability one. IT was responsible for around 2% of the global energy consumption making it equal to the demand of aviation industry in 2008 [1]. Focusing on data centers, late 2012 numbers from the European Commission [2] shows that European data centers consumed 60 TW h during 2012. The same study expects this number to double before 2020.

While this aggregated consumption is high, still nearly a third of organizations (29%) owning data centers did not measure their efficiency in 2012 [3]. Out of this study, for the data centers that measure their Power Usage Effectiveness (PUE) [4], more than a third (34%) have a PUE over or equal to 2, meaning they consume more power on cooling, air movement and infrastructure than on computing itself. The average PUE over all data centers is between 1.8 and 1.89.

Large energy needs and significant CO₂ emissions cause the issues related to cooling, heat transfer, and IT infrastructure location more and more carefully studied during planning and operation of data centers. Even if we take ecological and footprint issues aside, the amount of consumed energy can impose strict limits on data centers. First of all, energy bills may reach millions of euros making computations expensive. Furthermore, available power supply is usually limited so it also may reduce data center development capabilities, especially looking at challenges

* Corresponding author.

E-mail addresses: fontoura@irit.fr (L. Cupertino), dacosta@irit.fr (G. Da Costa), ariel@man.poznan.pl (A. Oleksiak), piatek@man.poznan.pl (W. Piątek), pierson@irit.fr (J.-M. Pierson), jsalom@irec.cat (J. Salom), lsiso@irec.cat (L. Sisó), stolf@irit.fr (P. Stolf), sun@irit.fr (H. Sun), zilio@irit.fr (T. Zilio).

related to exascale computing breakthrough foreseen within this decade. For these reasons many efforts were undertaken to measure and improve energy efficiency of data centers. Some of those projects focused on data center monitoring and management [5–7] whereas others on prototypes of low power computing infrastructures [8,9]. Additionally, vendors offer a wide spectrum of energy efficient solutions for computing and cooling [10,11].

A variety of possibilities exist at the design level, which have to be simulated in order to be compared and to select the best one. During the lifetime of a data center, smart management can lead to better visibility of the platform behavior and to reduce energy consumption.

In order to optimize the design or configuration of a data center we need a thorough study using appropriate metrics and tools evaluating how much computation or data processing can be done within a given power and energy budget and how it affects temperatures, heat transfers, and airflows within the data center. Therefore, there is a need for simulation tools and models that approach the problem from a perspective of end users and take into account all the factors that are critical to understanding and improving the energy efficiency of data centers, in particular, hardware characteristics, applications, management policies, and cooling. To address these issues the CoolEmAll project [12] aimed at decreasing energy consumption of data centers by allowing data center designers, planners, and administrators to model and analyze energy efficiency of various configurations and solutions. To this end, the project provides models of data center building blocks and tools that apply these models to simulate, visualize and analyze data center energy efficiency.

The structure of the paper is as follows. Section 2 presents relevant related works. Section 3 contains a brief description of the CoolEmAll project. In Section 4 we present the models that are used in the design and management tools. In Section 5 the metrics used to assess the quality of design and management are presented. Section 6 describes smart data center management techniques. In Section 7 we show the results of the simulation experiments and the impact of the proposed models and tools. Section 8 concludes the paper.

2. Related work

Issues related to cooling, heat transfer, IT infrastructure configuration, IT-management, arrangement of IT-infrastructure as well as workload management are gaining more and more interest and importance, as reflected in many ongoing works both in industry and research. There are already software tools available on the market capable to simulate and analyze thermal processes in data centers. Examples of such software include simulation codes along with more than 600 models of servers from Future Facilities [13] with its DC6sigma products, CA tools [14], or the TileFlow [15] application. In most cases these simulation tools are complex and expensive solutions that allow modeling and simulation of heat transfer processes in data centers. To simplify the analysis process Romonet [16] introduced a simulator, which concentrates only on costs

analysis using simplified computational and cost models, disclaiming analysis of heat transfer processes using Computational Fluid Dynamics (CFD) simulations. Common problem in case of commercial data center modeling tools is that they use closed limited databases of data center hardware. Although some of providers as Future Facilities have impressive databases, extensions of these databases and use of models across various tools is limited. To cope with this issue Schneider have introduced the GENOME Project that aims at collecting “genes” which are used to build data centers. They contain details of data center components and are publicly available on the Schneider website [17]. Nevertheless, the components are described by static parameters such as “nameplate” power values rather than details that enable simulating and assessing their energy efficiency in various conditions. Another initiative aiming at collection of designs of data centers is the Open Compute Project [18]. Started by Facebook which published its data center design details, it consists of multiple members describing data centers’ designs. However, Open Compute Project blueprints are designed for description of good practices rather than to be applied to simulations.

In addition to industrial solutions significant research effort was performed in the area of energy efficiency modeling and optimization. For example, models of servers’ power usage were presented in [19] whereas application of these models to energy-aware scheduling in [20]. Additionally, authors in [21,22] proposed methodologies of modeling and estimation of power by specific application classes. There were also attempts to use thermodynamic information in scheduling [23]. Nevertheless, the above works are focused on research aspects and optimization rather than providing models to simulate real data centers. In [24], the authors propose a power management solution that coordinates different individual approaches. The solution is validated using simulations based on 180 server traces from nine different real-world enterprises. Second, using a unified architecture as the base, they perform a quantitative sensitivity analysis on the impact of different architectures, implementations, workloads, and system design choices. Shah and Krishnan [25] explores the possibility of globally staggering compute workloads to take advantage of local climatic conditions as a means to reduce cooling energy costs, by performing an in-depth analysis of the environmental and economic burden of managing the thermal infrastructure of a globally connected data center network. SimWare [26] is a data warehouse simulator which compute its energy efficiency by: (a) decoupling the fan power from the computer power by using a fan power model; (b) taking into account the air travel time from the CRAC to the nodes; and (c) considering the relationship between nodes by the use of a heat distribution matrix.

3. The CoolEmAll project

CoolEmAll was an European Commission funded project which addresses the complex problem of how to make data centers more energy and resource efficient. CoolEmAll developed a range of tools to enable data center designers,

متن کامل مقاله

دریافت فوری ←

ISIArticles

مرجع مقالات تخصصی ایران

- ✓ امکان دانلود نسخه تمام متن مقالات انگلیسی
- ✓ امکان دانلود نسخه ترجمه شده مقالات
- ✓ پذیرش سفارش ترجمه تخصصی
- ✓ امکان جستجو در آرشیو جامعی از صدها موضوع و هزاران مقاله
- ✓ امکان دانلود رایگان ۲ صفحه اول هر مقاله
- ✓ امکان پرداخت اینترنتی با کلیه کارت های عضو شتاب
- ✓ دانلود فوری مقاله پس از پرداخت آنلاین
- ✓ پشتیبانی کامل خرید با بهره مندی از سیستم هوشمند رهگیری سفارشات