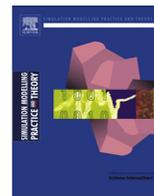


Contents lists available at [ScienceDirect](http://www.sciencedirect.com)

Simulation Modelling Practice and Theory

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

DCworms – A tool for simulation of energy efficiency in distributed computing infrastructures

K. Kurowski^a, A. Oleksiak^{a,*}, W. Piątek^a, T. Piontek^a, A. Przybyszewski^a, J. Węglarz^{a,b}^a Poznan Supercomputing and Networking Center, Noskowskiego 10, Poznan, Poland^b Institute of Computing Science, Poznan University of Technology, Piotrowo 2, Poznan, Poland

ARTICLE INFO

Article history:

Available online 26 September 2013

Keywords:

Simulations
Workload and resource modeling
Energy-efficiency
Data centers

ABSTRACT

In the recent years, energy-efficiency of computing infrastructures has gained a great attention. For this reason, proper estimation and evaluation of energy that is required to execute data center workloads became an important research problem. In this paper we present a Data Center Workload and Resource Management Simulator (DCworms) which enables modeling and simulation of computing infrastructures to estimate their performance, energy consumption, and energy-efficiency metrics for diverse workloads and management policies. We discuss methods of power usage modeling available in the simulator. To this end, we compare results of simulations to measurements of real servers. To demonstrate DCworms capabilities we evaluate impact of several resource management policies on overall energy-efficiency of specific workloads executed on heterogeneous resources.

© 2013 Elsevier B.V. All rights reserved.

1. Introduction

Rising popularity of large-scale computing infrastructures caused quick development of data centers. Nowadays, data centers are responsible for around 2% of the global energy consumption making it equal to the demand of aviation industry [14]. Moreover, in many current data centers the actual IT equipment uses only half of the total energy whereas most of the remaining part is required for cooling and air movement resulting in poor Power Usage Effectiveness (PUE) [31] values. Large energy needs of data centers led to increased interest in cooling, heat transfer, and IT infrastructure location during planning and operation of data centers.

For these reasons many efforts were undertaken to measure and study energy efficiency of data centers. There are projects focused on data center monitoring and management [10,1] whereas others on energy efficiency of networks [8] or distributed computing infrastructures, like grids [16]. Additionally, vendors offer a wide spectrum of energy efficient solutions for computing and cooling [32,27,29]. However, a variety of solutions and configuration options can be applied planning new or upgrading existing data centers. In order to optimize a design or configuration of data center we need a thorough study using appropriate metrics and tools evaluating how much computation or data processing can be done within given power and energy budget and how it affects temperatures, heat transfers, and airflows within 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. These tools should support data center designers and operators by answering questions how specific application types, levels of load, hardware specifications, physical arrangements,

* Corresponding author. Tel.: +48 618582187; fax: +48 618582151.

E-mail address: ariel@man.poznan.pl (A. Oleksiak).

cooling technology, etc. impact overall data center energy efficiency. There are various tools that allow simulation of computing infrastructures, like SimGrid [6]. Some of them include advanced packages for modeling heat transfer and energy consumption in data centers [30] or tools concentrating on their financial analysis [7]. On the other hand, there are simulators focusing on computations such as CloudSim [4]. The CoolEmAll project aims to integrate these approaches and enable advanced analysis of data center efficiency taking into account all these aspects [3,28].

One of the results of the CoolEmAll project is the Data Center Workload and Resource Management Simulator (DCworms) which enables modeling and simulation of computing infrastructures to estimate their performance, energy consumption, and energy-efficiency metrics for diverse workloads and management policies. We discuss methods of power usage modeling available in the simulator. To this end, we compare results of simulations to measurements of real servers. To demonstrate DCworms capabilities we evaluate impact of several resource management policies on overall energy-efficiency of specific workloads executed on heterogeneous resources.

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 modeling and simulation of computing systems in terms of energy efficiency. Section 3 discusses the main features of DCworms. In particular, it introduces our approach to workload and resource management, presents the concept of energy efficiency modeling and explains how to incorporate a specific application performance model into simulations. Section 4 discusses energy models adopted within the DCworms. In Section 5 we assess the energy models by comparison of simulation results with real measurements. We also present experiments that were performed using DCworms to show various types of resource and scheduling technics allowing decreasing the total energy consumption of the execution of a set of tasks. Final conclusions and directions for future work are given in Section 6.

2. Related work

The growing importance of energy-efficiency in information technologies led to significant interest in energy saving methods for computing systems. Nevertheless, studies of impact of resource management policies on energy-efficiency of IT infrastructures require a large effort and are difficult to perform in real distributed environments. To overcome these issues, extensive research has been conducted in the area of modeling and simulation and variety of tools that address the green computing have emerged. The most popular ones are: GreenCloud [11], CloudSim [4] and DCSG Simulator [5].

GreenCloud is a C++ based simulation environment for studying the energy-efficiency of cloud computing data centers. CloudSim is a simulation tool that allows modeling of cloud computing environments and evaluation of resource provisioning algorithms. Finally, the DCSG Simulator is a data center cost and energy simulator calculating the power and cooling schema of the data center equipment.

The scope of the aforementioned toolkits concerns the data center environments. However, all of them, except DCworms presented in this paper, restricts the simulated architecture in terms of types of modeled resources. In this way, they impose the use of predefined sets of resources and relations between them. GreenCloud defines switches, links and servers that are responsible for task execution and may contain different scheduling strategies. Contrary to what the GreenCloud name may suggest, it does not allow testing the impact of virtualization-based approaches. CloudSim allows creating a simple resources hierarchy consisting of machines and processors. To simulate a real cloud computing data center, it provides an extra virtualization layer responsible for the virtual machines (VM) provisioning process and managing the VM life cycle. In DCSG Simulator user is able to take into account a variety of mechanical and electrical devices as well as the IT equipment and define for each of them numerous factors, including device capacity and efficiency as well as the data center conditions.

The general idea behind all of the analyzed tools is to enable studies concerning energy efficiency in distributed infrastructures. GreenCloud approach enables simulation of energy usage associated with computing servers and network components. For example, the server power consumption model implemented in GreenCloud depends on the server state as well as its utilization. The CloudSim framework provides basic models to evaluate energy-conscious provisioning policies. Each computing node can be extended with a power model that estimates the current power consumption. Within the DCSG Simulator, performance of each of the data center equipment (facility and IT) is determined by a combination of factors, including workload, local conditions, the manufacturer's specifications and the way in which it is utilized. In DCworms, the plugin idea has been introduced that offers emulating the behavior of computing resources in terms of power consumption. Additionally, it delivers detailed information concerning resource and application characteristics needed to define more sophisticated power draw models.

In order to emulate the behavior of real computing systems, green computing simulator should address also the energy-aware resource management. In this term, GreenCloud offers capturing the effects of both of the Dynamic Voltage and Frequency Scaling (DVFS) and Dynamic Power Management schemes. At the links and switches level, it supports downgrading the transmission rate and putting network equipment into a sleep mode. CloudSim comes with a set of predefined and extensible policies that manage the process of VM migrations in order to optimize the power consumption. However, the proposed approach is not sufficient for modeling more sophisticated policies like frequency scaling techniques and managing resource power states. DCSG Simulator is told to implement a set of basic energy-efficient rules that have been developed on the basis of detailed understanding of the data center as a system. The output of this simulation is a set of energy metrics, like PUE, and cost data representing the IT devices. DCworms introduces a dedicated interface that provides methods to obtain the detailed information about each resource and its components energy consumption and allows changing its current

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