Computer Networks 57 (2013) 3479–3491

Contents lists available at ScienceDirect

Computer Networks

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

Energy-saving self-configuring networked data centers

Nicola Cordeschi *, Mohammad Shojafar, Enzo Baccarelli

Dpt. of Information Engineering, Electronic and Telecommunication, "Sapienza" University of Rome, via Eudossiana 18, 00184 Rome, Italy

ARTICLE INFO

Article history:
Received 24 May 2013
Accepted 5 August 2013
Available online 18 September 2013

Keywords:
Energy-saving
Networked data centers
Dynamic online communication–computing resource provisioning
Hard real-time applications

ABSTRACT

In this paper, we develop the optimal minimum-energy scheduler for the dynamic online joint allocation of the task sizes, computing rates, communication rates and communication powers in virtualized Networked Data Centers (NetDCs) that operates under hard per-job delay-constraints. The referred NetDC's infrastructure is composed by multiple frequency-scalable Virtual Machines (VMs), that are interconnected by a bandwidth and power-limited switched Local Area Network (LAN). Due to the nonlinear power-vs.-communication rate relationship, the resulting Computing-Communication Optimization Problem (CCOP) is inherently *nonconvex*. In order to analytically compute the exact solution of the CCOP, we develop a solving approach that relies on the following two main steps: (i) we prove that the CCOP retains a *loosely coupled* structure, that allows us to perform the lossless decomposition of the CCOP into the cascade of two simpler sub-problems; and, (ii) we prove that the coupling between the aforementioned sub-problems is provided by a (scalar) constraint, that is linear in the offered workload. The resulting optimal scheduler is amenable of scalable and distributed online implementation and its analytical characterization is in closed-form. After numerically testing its actual performance under randomly time-varying synthetically generated and real-world measured workload traces, we compare the obtained performance with the corresponding ones of some state-of-the-art static and sequential schedulers.

1. Introduction

Green computing over Networked Data Centers (NetDCs) is an emerging paradigm that aims at performing the dynamic energy-saving management of data center infrastructures. The goal is to provide QoS Internet services for large populations of clients, while minimizing the overall computing-plus-communication energy consumption [1–3, Chapter 6]. As recently pointed out in [3], the energy cost of communication gear for data centers may represent a large fraction of the overall system cost and it is induced primarily by switches, LAN infrastructures, routers and load balancers. In order to attain energy-saving, virtualization techniques are typically used for attaining resource isolation and online resource balancing [3]. However, current virtualized data centers are not designed for supporting communication–computing intensive real-time applications, such as, info-mobility applications, real-time video co-decoding, target recognition and tracking [28,33]. In fact, imposing hard-limits on the overall per-job delay does not allow the data center to perform data buffering over large time intervals, but forces the overall networked computing infrastructure to adapt quickly its resource utilization to the (possibly, unpredictable and abrupt) time fluctuations of the offered workload [4,9,10].

1.1. Goal of this paper

In order to attain the green paradigm, the joint balanced provisioning and scaling of the communication-plus-computing resources are demanded. This is, indeed, the

* Corresponding author. Tel.: +39 06 44585366; fax: +39 06 4873330.
E-mail addresses: nicola.cordeschi@diet.uniromal.it (N. Cordeschi), shojafar@diet.uniroma1.it (M. Shojafar), enzo.baccarelli@diet.uniroma1.it (E. Baccarelli).

1389-1286/$ - see front matter © 2013 Published by Elsevier B.V.
http://dx.doi.org/10.1016/j.comnet.2013.08.002

focus of this paper, whose main contributions can be summarized as follows:

i. the contrasting objectives of low consumption of both communication and computing energies in delay and bandwidth-constrained NetDCs affected by reconfiguration costs are cast in the form of a suitable constrained optimization problem, namely, the *Computing and Communication Optimization Problem* (CCOP);
ii. due to the nonlinear behavior of the rate-vs.-power-vs.-delay relationship, the CCOP is not a convex optimization problem and neither guaranteed-convergence iterative algorithms nor closed-form formulas are, to date, available for its solution. Hence. in order to solve the CCOP in exact and closed-form, we prove that it admits a loss-free (e.g., optimality preserving) decomposition into two simpler sub-problems, namely, the *CoMmunication Optimization Problem* (CMOP) and the *ComPuting Optimization Problem* (CPOP). Although the CMOP and CPOP are not guaranteed to be convex problems, their loosely coupled structure (in the sense of [25]) allows us to solve the (nonconvex) CCOP and develop analytical conditions for its feasibility;
iii. we develop a fully *autonomic* version of the proposed resource scheduler, that is capable to quickly adapt to the a priori unknown time-variations of the offered workload, without requiring workload forecasting;
iv. finally, we derive analytical conditions for the (possible) hibernation of the instantiated VMs, that highlight the tight inter-play between computing and communication resources.

A remarkable feature of the developed adaptive scheduler is its scalable and distributed structure, that makes the complexity of its online implementation independent from the size of the considered NetDC.

1.2. Related work

Updated surveys of the current technologies and open communication challenges related to the green cloud paradigm have been recently presented in [3,15]. Specifically, power management schemes that exploit Dynamic Voltage and Frequency Scaling (DVFS) techniques [15, Section 5] for performing resource provisioning are the focus of [26,30,31]. Although these contributions consider hard-deadline constraints, they do not consider, indeed, the performance penalty and the energy-vs.-delay tradeoff stemming from the finite capacity of the utilized LANs. Furthermore, the complexity of the online implementation of the therein proposed schedulers scales as $O(M\,log(M))$, where M is the minimum between the number of tasks to be executed in parallel and the number of available computing machines.

Energy-saving dynamic provisioning of the computing resources in virtualized green data centers is the topic of [10,16–20,27]. Specifically, [16] formulates the optimization problem as a feedback control problem that must converge to a priori known target performance level. While this approach is suitable for tracking problems, it cannot be employed for energy-minimization problems, where the target values are priori unknown. Roughly speaking, the common approach pursued by [10,17,18,20] is to formulate the afforded minimum-cost problems as sequential optimization problems and, then, solves them by using limited look-ahead control. Hence, the effectiveness of this approach relies on the ability to predict accurately future workload and degrades when the workload exhibits almost unpredictable time fluctuations [15].

Furthermore, the joint provisioning of communication and computing resources is not considered by these contributions, that mainly focus on the computing aspects. In order to avoid the prediction of future workload, [19] resorts to a Lyapunov-based technique, that dynamically optimizes the provisioning of the computing resources by exploiting the available queue information. Although the pursued approach is of interest, it relies on an inherent delay-vs.-utility tradeoff, that does not allow to account for hard-deadline constraints. The combined exploitation of DVFS and virtualization techniques is the focus of [27]. Although the parallel computing platform considered in [27] is, indeed, multi-core and managed by a Virtual Machine Manager (VMM), the framework developed in [27] does not consider load balancing and neglects both the energy and time overheads induced by the underlying LAN.

The joint analysis of the computing-plus-communication energy consumption in NetDCs is, indeed, the focus of [2,4,6], where delay-tolerant Internet-based applications are considered. Interestingly, the main lesson stemming from these contributions is that the energy consumption due to data communication may represent a large part of the overall energy demand, especially when the utilized network is power and/or bandwidth-limited. Overall, these works numerically analyze and test the energy performance of some state-of-the-art schedulers for NetDCs, but do not attempt to optimize it through the dynamic joint scaling of the available communication-plus-computing resources. As also recently pointed out in [3] and [15, Section 9], this is still an open research topic, especially when hard-delay requirements are also present.

The rest of this paper is organized as follows. After modeling in Section 2 the considered NetDC infrastructure, in Section 3 we formally state the afforded CCOP and, then, in Section 4, we solve it and provide the analytical conditions for its feasibility. In Section 5, we present the main structural properties of the resulting optimal scheduler, and analytically characterize the (possible) occurrence of hibernation phenomena of the instantiated VMs. In Section 6, we numerically test the average performance of the proposed scheduler at various values of the Peak-to-Mean Ratio (PMR) of the (randomly time-variant) offered workload, and, then, we compare the obtained performance against the corresponding ones of some state-of-the-art static and sequential schedulers. The conclusive Section 7 recaps the main results and outlines some hints for future research. The final Appendix reports the main analytical proofs.