Computer Communications 50 (2014) 142–151

Contents lists available at ScienceDirect

Computer Communications

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

Elastic operations in federated datacenters for performance and cost optimization

L. Velasco [a,*], A. Asensio [a], J.Ll. Berral [a,b], E. Bonetto [c], F. Musumeci [d], V. López [e]

[a] Universitat Politècnica de Catalunya (UPC), Barcelona, Spain
[b] Barcelona Supercomputing Center (BSC), Barcelona, Spain
[c] Politecnico di Torino, Turin, Italy
[d] Politecnico di Milano, Milan, Italy
[e] Telefónica Investigación y Desarrollo (TID), Madrid, Spain

ARTICLE INFO

Article history:
Available online 12 March 2014

Keywords:
Federated datacenters
Energy costs minimization
Inter-datacenter networks

ABSTRACT

The huge energy consumption of datacenters providing cloud services over the Internet has motivated different studies regarding cost savings in datacenters. Since energy expenditure is a predominant part of the total operational expenditures for datacenter operators, energy aware policies for minimizing datacenters' energy consumption try to minimize energy costs while guaranteeing a certain quality of experience (QoE). Federated datacenters can take advantage of its geographically distributed infrastructure by managing appropriately the green energy resources available in each datacenter at a given time, in combination with workload consolidation and virtual machine migration policies. In this scenario, inter-datacenter networks play an important role and communications cost must be considered when minimizing operational expenditures. In this work we tackle the Elastic Operations in Federated Datacenter for Performance and Cost Optimization (ELFADO) problem for scheduling workload orchestrating federated datacenters. Two approaches, distributed and centralized, are studied and integer linear programming (ILP) formulations and heuristics are provided. Using those heuristics, we analyze cost savings with respect to a fixed workload placement. For the sake of a compelling analysis, exhaustive simulation experiments are carried out considering realistic scenarios. Results show that the centralized ELFADO approach can save up to 52% of energy cost and more than 44% when communication costs are also considered.

1. Introduction

Cloud computing has transformed the IT industry, shaping the way IT hardware is designed and purchased [1]. Datacenters contain hardware and software to provide services over the Internet. Because datacenters consume huge amount of energy [2], energy expenditure becomes a predominant part of total operational expenditures for their operators. Aiming at reducing energy expenditure, datacenter operators can use, or even generate themselves, *green* energy coming from solar or wind sources; green energy would replace either partially or totally energy coming from *brown*, polluting sources. The drawback is that green energy is not always available, depending on the hour of the day, weather and season, among others. In contrast, brown energy can be drawn from the grid at any time, although its cost might vary along the day.

Thanks to virtualization, workloads (e.g. web applications) can be easily consolidated and placed in the most proper server according to its performance goals. By encapsulating workloads in virtual machines (VM) a datacenter resource manager can migrate them from one server to another looking for optimizing some objective function, such as energy consumption, whilst ensuring the committed quality of experience (QoE) [3,4].

Large Internet companies, such as Google and Microsoft, have their own infrastructures consisting in a number of large datacenters. These datacenters, placed in geographically diverse locations, guarantee good QoE to users and are interconnected through a wide area network [5]. Using that scheme, those companies can move workloads among datacenters to take advantage of reduced energy cost during off-peak energy periods in some locations (in addition to load balancing) while using green energy when is available in some other locations and turning off servers when they are not used, thus minimizing their energy expenditure.

* Corresponding author.
E-mail address: lvelasco@ac.upc.edu (L. Velasco).

http://dx.doi.org/10.1016/j.comcom.2014.03.004
0140-3664/

Nonetheless, there is a large number of smaller independently operated infrastructures which cannot perform such elastic operations. Notwithstanding, those medium-size datacenters can cooperate by creating datacenter federations [6] to increase their revenue from using IT resources that would otherwise be underutilized, and to expand their geographic coverage without building new datacenters. Network providers can facilitate federated datacenters interconnection by allowing them to request connections' setup on demand with the desired bitrate, while tearing down those connections when they are not needed in a *pay as you go* model. To that end, network operators could use some automated interface to allow resource managers, in charge of each datacenter, to request such connections even in multi-domain network scenarios [7].

From the optical networking perspective, the advent of the flexgrid technology allows optical connections to be assigned an optical spectrum width according to their requested bitrate [8]. In addition, huge research and standardization work have been done defining control plane architectures and protocols to automate connection provisioning allowing to request them dynamically [9]. Led by the development of the software-defined network (SDN) concept, the IETF is also moving towards a centralized controller with the definition of the Application-Based Network Operations (ABNO) architecture [10]. In our previous works in [11,12], we studied the relation between datacenter management and flexgrid networks using the ABNO architecture.

In this work we assume a set of federated datacenters strategically placed around the globe so as to provide worldwide, high QoE services, interconnected by a flexgrid-based network. Each datacenter has access to some amount of energy coming from green sources which can cover some percentage of total energy consumption (*green coverage*), being the rest drawn from the grid. We study two approaches to orchestrate such datacenter federation to provide committed QoE while minimizing operational expenditures: distributed and centralized. In the distributed approach, datacenters schedule VM placement so as to minimize an estimation of the energy cost plus communication costs while ensuring QoE. In the centralized approach, a centralized orchestrator computes the global optima from placing VM to take full advantage from green energy availability in the federated datacenters.

The internal datacenters architecture has become crucial to deploy energy-efficient infrastructures. A certain number of switches is necessary to provide connectivity between servers in the datacenter and to interface the datacenter with the Internet. Consequently, according to the datacenter architecture being adopted, a corresponding power is consumed, basically dependent on the number and type of switches used. Among the various intra-datacenter architectures studied in literature (see [13] for a detailed survey), the so-called *flattened butterfly* architecture has been identified as the most power-efficient datacenter architecture, thanks to its power-proportional behavior, i.e. its power consumption is proportional to the number of currently used servers. However, the most widely-deployed architecture for datacenter is the so-called *fat-tree* topology [14], which is based on a hierarchical structure where large higher-order switches represent the interface of the datacenter towards the network infrastructure, and are connected to the servers via a series of lower-order switches, providing the intra-datacenter connectivity.

Since minimizing energy expenditures is really important for datacenter operators, many papers can be found in the literature partially addressing that problem [15–18]. In [15], the authors propose scheduling workload in a datacenter coinciding with the availability of green energy, consolidating all the jobs on time slots with solar energy available, increasing green energy consumption up to 31%. Authors in [16] present a datacenter architecture to reduce power consumption, while guarantee QoE. They consider online-monitoring and VM placement optimization achieving energy savings up to 27%. Other works, e.g. [17], refer to the problem of load balance datacenter workloads geographically, following green energy availability, to reduce the amount of brown energy consumed focusing mainly on wind energy and the capability of store energy. Other works focus on the importance of counting as "energy expenditure" every element in the datacenter, not only computing machinery. The author in [18] remarks the idea that all IT equipment counts when consuming energy, also the fluctuation of green energy production and energy transportation are important factors.

As elastic operations for VM migration require huge bitrate to be available among datacenters for some time periods, the inter-datacenter network can be based on the optical technology and must provide automated interfaces to set-up and tear down optical connections with the required bitrate. Some works consider optical networks to interconnect datacenters. For instance, the authors in [19] present routing algorithms considering both routing and scheduling and compare energy savings with respect to a scenario where routing and scheduling problems are solved separately.

However, to the best of our knowledge, no work compares the way to compute scheduling considering both energy and communications costs in a single framework. In addition, we focus on solar energy, which is more predictable, and take more advantage of our network capabilities to migrate workload. Regarding our power model, we rely on using the Power Usage Effectiveness (PUE) ratio [20], where the consumed power becomes the computational power plus all the extra IT power directly derived from the first one. All the above is considered in the Elastic Operations in Federated Datacenter for Performance and Cost Optimization (ELFADO) problem. Solving ELFADO we reach energy consumption reductions up to 52%, outperforming previous works.

The rest of this article is organized as follows. Section 2 describes a power model for the fat-tree intra-datacenter architecture and presents the motivation of this work: to tackle the ELFADO problem. Two approaches for solving the ELFADO problem are presented: distributed and centralized. In Section 3, the ELFADO problem is formally stated and mathematical models and heuristics algorithms to solve it for both, the distributed and the centralized approach are presented. Illustrative results are provided in Section 4 for a realistic scenario with five datacenters strategically placed around the globe. Finally, Section 5 concludes the article.

2. Orchestrating federated datacenters

In this section we first present the considered power model to evaluate the energy consumption of each individual datacenter when they are based on the fat-tree architecture. Next, we present the main objective of elastic operations, i.e. minimizing operational costs by taking advantage from available green energy and cheap brown energy.

2.1. Datacenter power model

Two main contributions to the power consumption of a datacenter can be distinguished: (i) the power consumed by IT devices, P_{IT}, which comprises both the servers located in the datacenter as well as the switches employed to interconnect those servers; (ii) the power consumption of the non-IT equipment, $P_{non\text{-}IT}$, such as cooling, power supplies and power distribution systems. Thus, total power consumption of a datacenter can be computed as $P_{DC} = P_{IT} + P_{non\text{-}IT}$. P_{IT} can be easily estimated by counting the number of servers and switches of a datacenter. However, it is difficult