Computer Networks 130 (2018) 65–80

Contents lists available at ScienceDirect

Computer Networks

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

A Joint Power Efficient Server and Network Consolidation approach for virtualized data centers

Antonio Marotta[a,b], Stefano Avallone[b,*], Andreas Kassler[a]

[a] Karlstad University, Universitetsgatan 2, Karlstad 65188, Sweden
[b] University of Naples Federico II, Via Claudio 21, Naples 80125, Italy

ARTICLE INFO

Article history:
Received 15 May 2017
Revised 6 October 2017
Accepted 14 November 2017
Available online 21 November 2017

Keywords:
Cloud
Virtualization
Power
Green computing
Simulated annealing

ABSTRACT

Cloud computing and virtualization are enabling technologies for designing energy-aware resource management mechanisms in virtualized data centers. Indeed, one of the main challenges of big data centers is to decrease the power consumption, both to cut costs and to reduce the environmental impact. To this extent, Virtual Machine (VM) consolidation is often used to smartly reallocate the VMs with the objective of reducing the power consumption, by exploiting the VM live migration. The consolidation problem consists in finding the set of migrations that allow to keep turned on the minimum number of servers needed to host all the VMs. However, most of the proposed consolidation approaches do not consider the network related consumption, which represents about 10–20% of the total energy consumed by IT equipment in real data centers. This paper proposes a novel joint server and network consolidation model that takes into account the power efficiency of both the switches forwarding the traffic and the servers hosting the VMs. It powers down switch ports and routes traffic along the most energy efficient path towards the least energy consuming server under QoS constraints. Since the model is complex, a fast Simulated Annealing based Resource Consolidation algorithm (SARC) is proposed. Our numerical results demonstrate that our approach is able to save on average 50% of the network related power consumption compared to a network unaware consolidation.

1. Introduction

One of the main objectives in cloud computing is to find the right compromise between the cost-efficiency of the underlying infrastructure and the Quality of Service (QoS) as perceived by the users running their virtualized applications [1]. In the cloud computing paradigm there is an emerging **green computing** [2] awareness, which is strictly related to the cost-efficiency in the utilization of the physical resources: the aim is to reduce the power consumption and achieve an appropriate level of energy-efficiency. Big data centers, along with the adoption of the virtualization technology, are increasingly experiencing the need to reduce the energy consumption, because of both the environmental pollution and the economic concern. Hence, the interest of the research community is moving towards the introduction of metrics that allow to evaluate how energy efficient the resource utilization is. In the literature, techniques that try to increase the resource utilization of data centers have been massively studied from the energy viewpoint. One of them is the *VM consolidation*, whose objective is to minimize the number of physical servers necessary to host a set of VMs. By leveraging the VM *live migration* [3], the allocation of VMs to physical nodes can be dynamically adjusted to achieve different goals such as load balancing, avoiding hot spots or hibernating under-utilized servers.

Many VM consolidation techniques do not take into account networking issues within a data center. Greenberg et al. [4] shows, that the network consumes around 20% of the total energy in a data center. Large data centers are providing cloud enabled services and applications, which typically consist of multiple cooperating VMs that need to exchange large data volumes over the data center network. Classical network infrastructures, based on the spanning tree topology, have been proved to suffer from severe performance limits, such as network throughput, equipment capital expenditure (CapEx), network diameter and so on. In [5] a new parameter, namely the *Network Power Effectiveness* (NPE), is introduced to evaluate the network efficiency as the ratio between the aggregate throughput and the total power consumption. Along with novel architectures, techniques for increasing energy efficiency in the network are also emerging. Commercial switches allow putting network interfaces in sleep-mode when no traffic is flowing, and

* Corresponding author.
E-mail addresses: antonio.marotta@kau.se (A. Marotta), stefano.avallone@unina.it (S. Avallone), andreas.kassler@kau.se (A. Kassler).

https://doi.org/10.1016/j.comnet.2017.11.003
1389-1286/

waking-up a link only when a packet arrives. This is useful since several papers argue that the traffic load is very day-time dependent in data centers. For example, [6] shows that the average link utilization in the connections to the aggregate switches is only 8% of the capacity for 95% of the time, while in the core layer the utilization is between 20% and 40%. Therefore, it is possible to dynamically switch-off under-utilized network devices and links in some time intervals in order to save energy. Network protocols can also benefit from energy efficiency. For example, the IEEE 802.3az [7] standard is aimed at achieving energy reduction in the Ethernet based communications by activating a link only when real data is being sent.

In [8], we proposed a new MILP model and a heuristic to consolidate a virtualized data center by reallocating VMs on the smallest subset servers in order to minimize the total energy consumption due to compute resources. The key contribution of this work is a new mathematical model extending [8], that also considers the network characteristics in order to jointly optimize the VM placement and energy efficient network routing. By knowing the current VM placement and the network routing for sending VM-to-VM traffic, we consider the energy efficiency profiles of servers and networking devices in order to find the VM migrations and network paths that jointly minimize the server and network power consumption, powering down unused switches and link ports which are not necessary for routing the traffic demands. As the proposed model is complex and very hard to solve, we develop a fast Simulated Annealing based Resource Consolidation (SARC) heuristic. The parameter, called *migration desirability*, is used in the perturbation phase of the heuristic: it allows us to balance between the resource utilization of the physical servers, the power efficiency and the impact the migrations may have on the power-aware established network paths. We show that our heuristic is able to save on average 50% of the total network power consumption when compared to the heuristic that simply consolidates the active servers.

The paper is structured as follows: in Section 2, the related work is discussed. In Section 3, the problem is formulated and the objective is explained, while in Section 4, a Mixed Integer Programming Model is described. In Section 5, the heuristic is presented and Section 6 shows the experimental evaluation. Finally, the paper concludes in chapter VII.

2. Related work

Several papers address the consolidation problem in virtualized environment by either optimizing the physical resources utilization or the network efficiency. In [8], we presented a novel model for power efficient VM consolidation: the problem is to find the set of migrations that minimize both the overall power consumption of the active servers after the consolidation and the number of migrations. The proposed heuristic was shown to be very effective in approximating the optimal solution. In [9–11], VM live migration is exploited to increase energy efficiency. Setzer and Wolke [12] addresses the reallocation of VMs over time through an ILP model that takes into account the migration overhead. In [13] an interesting approach for VM consolidation is discussed, which is based on the idea that virtual appliances exchanging a lot of traffic should be allocated in the same rack. This is because the traffic is not sent to physical switches, thus minimizing the communication overhead. In [14], authors present an architectural framework for energy-efficient cloud computing. The aim of the proposed heuristics is to consolidate the physical resources in such a way that the energy efficiency is maximized, while meeting the users' expectations in terms of Quality of Service (QoS). Therefore, they describe different energy-efficient resource allocation policies and algorithms which take into account the QoS needs and the power consumption of the resources, and evaluate them in the CloudSim toolkit [15]. The VM consolidation problem is formulated as a stochastic bin packing problem in [16] and an online packing algorithm with proven guarantees on the number of active servers is presented. An LP formulation and heuristics to avoid unnecessary VM migration by prioritizing VMs with steady capacity are proposed in [17].

Live migration of multiple VMs with resource reservation technology is investigated in [18]. A self-organizing and adaptive approach for the consolidation of VMs based on probabilistic decisions is presented in [19]. In [20], the Ant Colony System meta-heuristic is used to solve the VM consolidation problem. A profiling-based server consolidation framework is proposed in [21] to minimize the number of physical machines (PMs) used in data centers. This approach is based on profiling the performance losses of various workloads under two situations: running in co-location and experiencing migrations. In order to get improved disk I/O performance, in [22] authors address the VM placement problem with the additional constraint that, for every VM assigned to a physical node, its virtual disks should be spread out across the physical disks of the node. Nevertheless, none of the previous cited papers jointly solves the VM consolidation and the flow routing problems. A good overview of the approaches that only address the VM migration and server consolidation problem can be found in [23].

The problem of determining when to reallocate VMs from an overloaded host is addressed in [24], where a solution based on a Markov chain model is proposed to maximize the mean inter-migration time under the specified QoS goal. We consider this problem as complementary to the one we address in this paper. Only a few works take the network power consumption of data centers into account. For instance, Shang et al. [5] presents an interesting study of the network power effectiveness of the most common data center architectures. Authors evaluate the network power consumption of the classical and power-aware routing, both applied to well-known topologies, by assessing the impacts of the topology size and the traffic load. VMFlow [25] formulates the *Network Aware VM Placement (NAPV)*. It is solved through a greedy heuristic, which considers one VM traffic demand at a time, for which a VM placement and a flow-path assignment are computed, with the objective of minimizing the total network power consumption. The algorithm selects a single demand from a list, sorted in decreasing order of amount of requested bandwidth. Given the source and destination of a flow, a placement is determined among the servers that have enough resources to host the VMs. After the allocation, the algorithm computes the flow routing path which causes the minimum increase in the energy consumption, trying to guarantee the highest number of VM allocations and, as a consequence, of satisfied traffic demands. However, they do not consider multiple objectives in their formulation, as well as the migrations.

In [26], authors derive formulas for the minimum number of active switches needed in a fat-tree data center network. In [27], instead, authors develop a power-aware routing algorithm to maximize the network power conservation. However, the VM consolidation problem is not addressed in these papers. In [28], authors propose a network power manager that shuts down unused switches and/or links and forwards traffic flows through the active networking elements. They focus on the network power efficiency without taking into account the possibility of re-allocating VMs. Authors use a heuristic to locally determine the power state of each switch and link. The work in [29] presents a formulation for the flow assignment problem and an evaluation of the path consolidation algorithms, by considering different characteristics of the network topology. Given the total power consumption over the network elements in a time frame, the problem is formulated as a classical Multi-Commodity Flow Problem with the aim of minimiz-