Performance Evaluation 79 (2014) 104–119

Contents lists available at ScienceDirect

Performance Evaluation

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

Reducing the expenses of geo-distributed data centers with portable containerized modules☆

Marco Brocanelli *, Wenli Zheng, Xiaorui Wang

Power-Aware Computer Systems (PACS) Laboratory, Department of Electrical and Computer Engineering, The Ohio State University, Columbus, OH, 43210, USA

ARTICLE INFO

Article history:
Available online 10 July 2014

Keywords:
Containerized data centers
Traditional data centers
Container transportation
Workload dispatch
OpEx
CapEx

ABSTRACT

In order to maximize their profits, big IT companies need to reduce the operating expenses (OpEx) and capital expenses (CapEx) of their geo-distributed data centers. To reduce OpEx, recent studies have proposed algorithms that dynamically dispatch workload among different data center sites, to exploit the variability in local electricity price and green energy availability. However, the amount of OpEx savings achieved by workload dispatching is often limited by the local computing capacity. To reduce CapEx, data centers try to minimize the needed increases in size based on estimated workload growth. However, workload overestimation would result in a high CapEx, while underestimation would hurt business increases.

In this paper, we propose ExContainer, a novel strategy that leverages portable containerized modules to minimize both OpEx and CapEx for geo-distributed data centers. With ExContainer, containers are allocated monthly to different sites and workload is dispatched hourly, to minimize total OpEx. To cut CapEx, ExContainer finds the best trade-off between the degree of capacity expansion and the number of needed new containers to handle the workload peak. We evaluate ExContainer with real-world traces. Our experimental results show that compared to state-of-the-art solutions, ExContainer allows significant reduction in both OpEx and CapEx.

1. Introduction

In the era of cloud computing, big IT companies are rapidly expanding their data centers in order to meet the fast-increasing demands for computing. In order to achieve the maximized profits, those companies need to minimize the operating expenses (OpEx) and capital expenses (CapEx) of their geo-distributed data centers.

Dynamic workload dispatching has been recently proposed as an effective way to reduce data center OpEx. For example, it is well-known that different geographical locations may have different and time-varying electricity prices, so Internet requests can be dynamically distributed to different data center sites to take advantage of the lower prices in certain locations [1,2]. Likewise, some data centers may have on-site renewable energy generators (e.g., solar or wind farms) that can provide low-cost green energy. As a result, requests can be dynamically dispatched based on green energy availability as well for reduced OpEx [3,4]. Similarly, requests can also be dispatched based on the ambient temperatures of different data centers, because the ambient temperature can significantly impact data center costs if advanced cooling techniques such as

☆ This work was supported, in part, by the US NSF under grants CNS-1143607 (CAREER Award) and CCF-1331712.
* Corresponding author.
E-mail addresses: brocanelli.1@osu.edu (M. Brocanelli), zheng.691@osu.edu (W. Zheng), wang.3596@osu.edu (X. Wang).

http://dx.doi.org/10.1016/j.peva.2014.07.007
0166-5316/

free cooling [5] are used. Since cooling power accounts for 30%–40% of the total power consumption of a data center [6], the OpEx savings on cooling can also be significant.

However, the amount of OpEx savings that can be achieved by workload dispatching is often limited by the size of each data center site. This is because a data center cannot handle more workload if it already reaches the peak computing capacity (i.e., all its servers are 100% utilized), regardless of its low electricity price, abundant green energy, or cold weather. Simply increasing the servers hosted in every data center may lead to a significantly increased CapEx as we discuss later. In addition, different data centers are commonly configured to provide different services (e.g., search, email, video streaming). Hence, workload can be dispatched only to some data centers that provide the required services and have available capacities, which further limits the benefits of workload dispatching. Since different data centers rarely have their workload peaks at the same time, portable containerized modules [7–9] can be exploited to flexibly move servers among different sites monthly based on electricity price, green energy availability, and ambient temperatures, such that hourly workload dispatching can achieve a further reduced OpEx.

In order to reduce CapEx, data centers try to minimize the needed increases in computing, cooling, and power delivery capacities based on estimated workload growth over the lifetime of the data center infrastructure (about 15 years [6,10]). However, it is well-known that workload growth can be highly unpredictable with significant fluctuations. As a result, workload overestimation would result in over-provisioning of data center capacities and thus a high CapEx [10], while underestimation would hurt business increases due to insufficient capacities. Therefore, if the size of a data center can be increased incrementally with container modules to handle short-term workload peaks, the data center can wait for the demand to grow to a certain level that best justifies an upgrade of the data center power and cooling capacities. Hence, expensive upgrades can be deferred, resulting in reduced CapEx.

Recently, big IT companies such as Microsoft [11], eBay [12], and Google [13] have started to adopt portable containerized modules, which are designed mainly to provide additional computing capacity for disaster recovery, sport events, or military applications. In sharp contrast to traditional (room-based) data centers, containerized modules can be quickly built (typically within 1–6 months), shipped with trucks, and deployed with power and network in a few hours [9]. IT equipment and cooling system can be placed within one module (all-in-one, e.g., HP POD 240a, Huawei IDS1000) or within two different modules (e.g., IBM PMDC). Moreover, the cooling system is usually built with an economizer mode (i.e., free cooling), which uses outdoor air (when the temperature is lower than a certain threshold) to cool down servers. As a result, containers can generally achieve a lower Power Usage Effectiveness (*PUE*, as low as 1.05 [8]).

In this paper, we propose ExContainer, a novel strategy that leverages portable containerized modules to minimize both OpEx[1] and CapEx for geo-distributed data centers. Although containerized modules have already been used in industry, to our best knowledge, they have not yet been used to fully exploit the diversity and variability in electricity price, green energy availability, and ambient temperature across different geographical locations for minimized OpEx. Neither have they been used to optimize data center CapEx by handling short-term workload spikes and deferring the expensive upgrades of data center infrastructure. ExContainer is composed of an OpEx optimizer and a CapEx optimizer. The OpEx optimizer features a two-level approach, which allocates containers to different sites on a coarse time scale (e.g., one month) and conducts dynamic workload dispatching and server turning on/off on a fine time scale (e.g., every hour), for the minimization of electricity and network delay (from dispatcher to data center) cost. The CapEx optimizer uses containers to handle short-term workload peaks that would exceed the capacity of a data center. When the demand grows to a level that justifies an upgrade of the data center capacities, the CapEx optimizer finds the best trade-off between the degree of capacity expansion and the number of needed new containers, to minimize the total CapEx. Specifically, this paper makes the following major contributions:

- While existing work focuses mainly on workload dispatching for reduced OpEx, we identify its two limitations: limited capacity of each site and service heterogeneity. Accordingly, we propose to leverage portable containerized modules for further reduced OpEx.
- While current practice commonly uses containers in a stationery way as modules and only moves containers for special events, we propose to move them periodically for minimized OpEx, based on the electricity prices, green energy availability, and ambient temperatures.
- We propose a new CapEx minimization approach to defer the expensive upgrades of data center infrastructure by finding the best trade-off between the degree of capacity expansion and the number of needed new containers.
- We formulate three constrained optimization problems to minimize OpEx by (1) allocating containers monthly to different locations, (2) conducting hourly dynamic workload dispatching and server turning on/off, and to minimize CapEx. Our results show that compared to state-of-the-art solutions [1,2,14,6,10], ExContainer allows significant reduction in both OpEx and CapEx.

The rest of this paper is organized as follows. Section 2 discusses the related work. Section 3 presents the overview of ExContainer. Sections 4 and 5 describe the detailed designs of the OpEx and CapEx optimizers, respectively. Section 6 evaluates ExContainer with real-world traces. Section 7 discusses the implications of ExContanainer and Section 8 concludes the paper.

[1] We focus on electricity bill in this paper, though OpEx can include maintenance, software license, salaries, etc.