Computer Communications 107 (2017) 60–74

Contents lists available at ScienceDirect

Computer Communications

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

On the benefits of resource disaggregation for virtual data centre provisioning in optical data centres

Albert Pagès*, Rubén Serrano, Jordi Perelló, Salvatore Spadaro

Advanced Broadband Communications Center (CCABA), Universitat Politècnica de Catalunya (UPC), Jordi Girona 1-3, 08034 Barcelona Spain

ARTICLE INFO

Article history:
Received 4 July 2016
Revised 27 March 2017
Accepted 30 March 2017
Available online 4 April 2017

Keywords:
Data centres
Resource disaggregation
Virtualization
Optimization
Optical networks

ABSTRACT

Virtual Data Centre (VDC) allocation requires the provisioning of both computing and network resources. Their joint provisioning allows for an optimal utilization of the physical Data Centre (DC) infrastructure resources. However, traditional DCs can suffer from computing resource underutilization due to the rigid capacity configurations of the server units, resulting in high computing resource fragmentation across the DC servers. To overcome these limitations, the disaggregated DC paradigm has been recently introduced. Thanks to resource disaggregation, it is possible to allocate the exact amount of resources needed to provision a VDC instance. In this paper, we focus on the static planning of a shared optically interconnected disaggregated DC infrastructure to support a known set of VDC instances to be deployed on top. To this end, we provide optimal and sub-optimal techniques to determine the necessary capacity (both in terms of computing and network resources) required to support the expected set of VDC demands. Next, we quantitatively evaluate the benefits yielded by the disaggregated DC paradigm in front of traditional DC architectures, considering various VDC profiles and Data Centre Network (DCN) topologies.

1. Introduction

Data Centre (DC) infrastructures are a key element in nowadays' telecom and cloud infrastructures, allowing the access to enormous quantities of information anytime and anywhere. Thanks to the collaborative efforts of the thousands of servers hosted inside their premises, complex Internet and cloud services (e.g., search engines, cloud storage, etc.) can be realized. In traditional DCs, servers are arranged in racks, each one equipped with a Top of the Rack (ToR) switch that interconnects the several servers inside, and allows for the exchange of information between different racks across an intra-DC Network (DCN) fabric. Current DCN architectures are usually build upon commodity electrical switches (e.g., Ethernet), arranged in a multi-layer architecture, which provides several aggregation points and means of redundancy for enhanced utilization of the network resources and Quality of Service (QoS) guarantees [1]. Besides, racks on the DC are usually grouped in different regions, named clusters, to allow for a better scalability and management of the whole DC infrastructure.

However, the constant growth of the Internet traffic and cloud services fostered by bandwidth-hungry applications/paradigms such as Big Data, Internet of Things (IoT) and Video on Demand (VoD), pleads for bigger DC infrastructures in terms of both computing and network capacities, in order to accommodate all applications and workflows. For instance, it is forecast that the global IP traffic managed by DCs will almost double by the year 2019, rising from 5.6 ZB to 10.4 ZB per year, with around 75% of the traffic staying inside their premises [2]. This unprecedented traffic growth is pushing the capabilities of current electrical-based DCN fabrics beyond their limits. For this reason, special attention at improving the performance of intra-DCNs is being put in the development of future DC architectures. In this regard, optical technologies have gained considerable interest due to their superior scalability, bandwidth and latency, as well as reduced power consumption. Hence, lots of efforts are being devoted to integrate them in future DCNs [3], either based on hybrid electrical/optical (e.g., as in [4]) or all-optical (e.g., see [5,6]) network fabrics for the communication of servers inside the DC.

Despite such efforts on improving the performance of DCNs, current server-centric DCs still face some limitations toward efficient computing resource utilization. In general, services/tasks in DCs are executed on top of Virtual Machines (VMs) that are deployed at servers. Each VM is provisioned with a set of computing resources (i.e., CPU cores, storage and memory) tailored to the computational needs of the applications. These resources are then allocated and dedicated to VMs during their whole lifecycle. A co-existence of VMs inside the same server is possible if the total amount of resources requested by all of them does not exceed the server's total resource capacity. However, the heterogeneous VM

* Corresponding author.
E-mail address: albertpages@tsc.upc.edu (A. Pagès).

http://dx.doi.org/10.1016/j.comcom.2017.03.009
0140-3664/

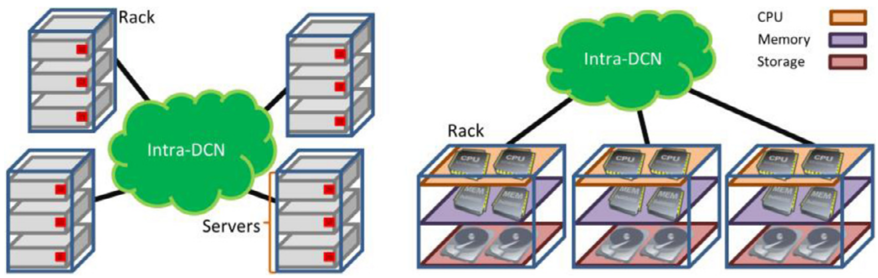

Fig. 1. Server-centric (left) vs. disaggregated (right) DC architecture.

computing resource demands can lead to server underutilization. For instance, it may happen that an application/service (i.e., a VM) running on a server employs almost the totality of one resource type (e.g., CPU cores), while imposing almost no requirements to the others (e.g., storage, memory). As a result, it may be impossible to allocate another application in the same server due to the scarcity of that resource type, letting the remainder underutilized. As an example of this phenomenon, Google has recently published data regarding the utilization of their DC infrastructures, disclosing high disparity of storage/memory to CPU usage for their tasks [7]. Furthermore, it becomes even more difficult to dynamically configure the DC resources under an unpredictable traffic profile.

Aside from poor resource utilization, server-centric architectures also suffer from a limited modularity that impacts on the system-wide performance. Traditional servers are usually built by tightly integrating their components (CPU, memory modules, disk, network interface card, etc...) into a single motherboard. This has been the basis of computer manufacturing for many years. However, this tight integration is responsible for the limited improvement possibilities of the overall system performance. This mainly happens because the rate at which the several components scale (in size, speed, etc.) is substantially different. For instance, the rate per year at which CPU performance has increased has been about 60%, while the rate of improvement in DRAM memory performance has merely been around 7% per year. This fact leads to a performance gap between CPU and memory of about 50% per year [8]. Such a disparity on the evolution of the different kinds of server components prevents utilizing the most advanced technology in some cases, since compromise decisions have to be taken in favour of a good system performance.

To overcome these challenges, new DC architectures have to be designed. An interesting approach to this end is the resource disaggregation concept [9], which proposes to disaggregate the computing resource components by physically decoupling and mounting them in separated blades, instead of tightly coupling them in a single integrated system. By physically decoupling the components, it is possible to adopt state-of-the-art technologies for each one of them, thus allowing for system optimization and customization. Such a concept has resulted in the disaggregated DC paradigm [10–13], where computing resources are no longer hosted in server units, but spread over standalone hardware blades. Resource blades can be grouped in racks hosting all types of computing resources (see Fig. 1, right), or in mono-hardware racks where only a single type of resource is held. Then, resource blades are interconnected through the intra-DCN fabric. To meet the strict latency and bandwidth requirements for communicating the different hardware modules, intra-DCN optical technologies are envisioned [10–12].

Through disaggregation, computing resources can be tightly assigned to VMs according to their needs, requiring fewer resources to satisfy a demand set while reducing the associated CAPEX. Moreover, disaggregation brings modularity to systems, enabling easier hardware upgrades when desired. For these reasons, optically interconnected disaggregated DCs are seen as a solution for future DCs.

Despite the benefits that the resource disaggregation paradigm promises to bring, there is little work in the literature on analysing the enhanced computing resource utilization of disaggregated DCs in front of nowadays' server-centric ones. In view of this, in this work we quantify the required computing resources to be equipped at DC infrastructures following the resource disaggregation paradigm when allocating service requests, and compare it to legacy server-centric DC ones. To this end, we focus on a static capacity planning of an optically interconnected disaggregated DC, aiming to give insight into the reduction of computing resource requirements that such paradigm can yield.

Given that Virtual Data Centre (VDC) has been identified as a key service that modern DCs have to offer to be able to efficiently implement multi-tenancy in a cloud environment, we will focus our efforts on analysing the planning of both disaggregated and server-centric DC infrastructures when supporting VDC services. To this goal, the remainder of the paper is structured as follows: Section 2 introduces the concept of VDC service and the issues involved. Next, Section 3 reviews the related work in the literature regarding disaggregated DCs and VDC provisioning, highlighting the contributions of this work. Section 4 elaborates on the considered DC planning scenario, presenting the optimization problem under consideration, while Section 5 details the different solutions proposed to tackle it, both for server-centric and disaggregated DC architectures. Next, Section 6 numerically evaluates and compares their computing resource requirements. Finally, Section 7 draws up the main conclusions of the work.

2. Virtual data centre provisioning

Contemporary DC infrastructures must allocate a plethora of customers, ranging from business/companies or public institutions to individual users, having all of them heterogeneous needs in terms of resource necessities, QoS, degree of control over the employed resources, and so on. In this regard, multi-tenancy becomes a pillar requirement that modern DC infrastructures must provide. However, traditional telecom and cloud infrastructure architectures present some drawbacks compromising the efficient implementation of multi-tenancy, even more in cloud environments, where a high degree of customization and dynamicity is present. Besides, scalability and resource provisioning specific challenges must be solved [14]. Indeed, the service structure is very rigid, with infrastructure owners focusing on the services offered on top of their infrastructures, with limited adaptability to the service to be deployed. This has led to architectures incapable to adapt to dynamic traffic patterns, with high heterogeneity on the characteristics of the services/applications to be deployed.

To overcome these limitations, the concept of Infrastructure as a Service (IaaS) has been introduced [15]. The emergence of IaaS arises from the need to provide telecom and cloud infrastruc-