Business Scenarios for Hierarchical Workload Management in Data Centers

Alfredo Cuzzocrea
University of Trieste
and ICAR-CNR
Italy
cuzzocrea@icar.cnr.it

Carmine De Napoli
eco4cloud srl
c/o Unical
Rende (CS), Italy
denapoli@eco4cloud.com

Agostino Forestiero
ICAR-CNR
via P. Bucci 7/11 C
Rende (CS), Italy
forestiero@icar.cnr.it

Demetrio Lagana'
eco4cloud srl
c/o Unical
Rende (CS), Italy
lagana@eco4cloud.com

Giovanni Lupi
eco4cloud srl
c/o Unical
Rende (CS), Italy
lupi@eco4cloud.com

Carlo Mastroianni
ICAR-CNR
via P. Bucci 7/11 C
Rende (CS), Italy
mastroianni@icar.cnr.it

Leonardo Spataro
eco4cloud srl
c/o Technest - Unical
Rende, Cosenza
spataro@eco4cloud.com

ABSTRACT

Increasing the efficiency of workload management in data centers is essential to achieve several business and technical goals, such as reduction of costs, energy consumption and carbon emissions. Many solutions are available for workload management in a single data center, but there is still much space for the development of frameworks that are able to manage the workload in a distributed scenario, with multiple sites and data centers. In this paper, we present the main benefits of hierarchical solutions, in which the problem is decomposed into two layers, a lower layer that focuses on the workload management within each single data center, and an upper layer that orchestrates the management of the workload on a multi-site environment. We focus on the main advantages of hierarchical approaches, i.e., autonomous management, scalability and modularity, and illustrate how and to which extent these advantages can be exploited in some emerging business scenarios, i.e., geographical data centers, software data centers and Hybrid Cloud.

CCS Concepts

•**Computing methodologies** → **Distributed computing methodologies;**

Keywords

Cloud Computing, Data Centers, Workload Management, Big Data

1. INTRODUCTION

The ever-increasing demand for computing resources has led companies and resource providers to build private warehouse-sized data centers, or to offload applications to the data centers owned by a Cloud company. Overall, data centers require a significant amount of power to be operated. The total electricity demand of data centers increased by about 56% from 2005 to 2010, and the electricity usage accounted for about 1.5% of the worldwide electricity usage in 2010, which is comparable to the aviation industry. The financial impact for the data center management is also huge, since a data center spends between 30% to 50% of its operational expense toward electricity[1]. The efficient utilization of resources in these data centers is therefore essential to reduce costs, energy consumption, carbon emissions and also to ensure a high quality of service to users.

Over the recent years, the ecological conscience of the European Union has led to a determination to restructure their Energy Policy. Energy represents one of the most significant operating costs in data centers. Moreover, based on current and foreseeable trends, the basic price of energy will continue to rise over time, and may become constrained as global demands rise, making energy use and efficiency a long term business priority. Carbon footprint and greenhouse gases are also becoming subject to governmental regulations and taxes. Companies are today strongly encouraged to reduce the amount of carbon emissions, not only to compel to laws and rules, but also to advertise their green effort and attract

[1]Updated information can be found on the Web portal of the US National Resources Defense Council, http://www.nrdc.org/energy/data-center-efficiency-assessment.asp

customers that are increasingly careful about sustainability issues.

In principle, Cloud computing has the environmentally important benefits of reduced energy consumption and carbon footprint in comparison with more traditional approaches, thanks to the following main features:

- **Hardware reduction**: one of the Cloud's green attributes take the form of "dematerializing" the economy which involves reducing the number of physical materials. Moreover, by reducing the need for hardware, companies can reduce costs and eliminate the need for maintenance and upgrades.

- **Virtualization**: through the allocation of multiple Virtual Machines (VMs) on the same physical server, the virtualization technology helps to increase the efficiency of DCs. A good level of efficiency must be guaranteed also in geographically distributed DCs, whose adoption is rapidly increasing.

- **Reduction of carbon emissions**: the Cloud reduces carbon emissions through minimized energy requirements and, in particular, offsite servers have the potential to prevent 85.7 million metric tons of annual carbon emissions by 2020 [1]. The environmental impact of these substantial reductions in energy are significant.

- **Reduction of energy costs** (depending on data centers location): the cost of electricity is generally different from site to site and also varies with time, even on a hour-to-hour basis, therefore the overall cost may be reduced by shifting portions of the workload to more convenient sites.

Workload consolidation, i.e., using the minimum number of physical hosts to accommodate as more virtual machines as possible, is a powerful means to improve IT efficiency and reduce power consumption and carbon emissions within a data center [5] [15] [26] [3]. While there are a number of efficient solutions for workload management in a single data center, there is much space to the development of frameworks that are able to manage the workload in a distributed scenario, with multiple sites and data centers.

In this paper, we will describe the main benefits of hierarchical solutions, in which the problem is tackled through a lower layer, which focuses on the workload management within each single data center, and an upper layer, which dynamically orchestrates the management of the workload on a geographically distributed environment. Moreover, we will summarize the main features and advantages of EcoMultiCloud [14], a comprehensive solution recently developed and implemented by ICAR-CNR (Institute for High Performance Computing Networking of the Italian National Research Council) and by the company Eco4Cloud, a spinoff from CNR and the University of Calabria, Italy.

Subsequently, we will present the main business scenarios that are expected to take advantage from multi data center solutions for the workload management. While the most natural business scenario is the increasing adoption of geographically distributed data centers, other two business applications are expected to experiment an even larger growth in the next future. They are: (i) software defined data centers, a paradigm that aims to deliver full control of data centers via software, while the routing and communication procedures are no more hardwired on the routers but are defined on the edge hosts through the services offered by Software Defined Networking (SDN) and (ii) the Hybrid Cloud, which will allow companies to manage an ecosystem where the on-premises hardware resources are transparently interconnected to the facilities of external Cloud data centers, and it is possible to dynamically and rapidly move data and computational workload from internal to external resources and vice versa, depending on the technical and business requirements.

Along this line of research, the integration of Clouds with big data management is similarly relevant. Indeed, the workload-distribution solutions discussed in our paper can, from a side, be beneficial to this end, and, from another side, smoothly integrate Clouds and big data management. Besides, it is clearly enough that Cloud/big-data integration now represents a fundamental component of every modern big data application.

The rest of the paper is organized as follows: in Section 2 we summarize the architecture, the main characteristics and the benefits of hierarchical solutions for workload management, taking EcoMultiCloud as an example; in Section 3 we discuss the main business applications for such solutions and offer an overview of the market trends and expectations; Section 4 briefly explores some hot-topics related to the integration of Clouds with big data management; finally, Section 5 concludes the paper and suggests some avenues for future academic and industrial research on this field.

2. A HIERARCHICAL SOLUTION FOR MULTI-SITE DATA CENTERS

The problem of workload assignment and redistribution in geographically distributed data centers is a topical field today. Research efforts focus on two related but different aspects [21]: the routing of service requests to the most efficient data center, in the so called *assignment* phase, and the live *migration* of portions of the workload when conditions change and some data centers become preferable in terms of electricity costs, emission factors, or more renewable power generation (redistribution and allocation). Several studies explore the opportunity of energy cost-saving by routing jobs when/where the electricity prices are lower [22]. Rao et al. [25] tackle the problem taking into account the spatial and time diversity in dynamic electricity markets.

Workload management in a geographical scenario is typically solved as an optimization problem, often in a centralized way. This approach has three main implications: (i) poor scalability, due to the large number of parameters and servers; (ii) poor ability to adapt to changing conditions, as massive migrations of VMs may be needed to match a new decision on of the workload distribution; (iii) limitation to the autonomy of the sites, which are often required to share the same strategies and algorithms.

The workload assignment and migration decision processes are made particularly complex by the time-variability of electricity cost, and by the workload variability both within single sites and across interconnected sites. The dynamic migration of workload among data centers has become an opportunity to improve several aspects: better resiliency and failover management, improved load balancing, exploitation of the "follow the moon" paradigm (i.e., move the workload where the energy is cheaper/cleaner and/or cooling costs are lower). Inter-site migration is enabled by the availability of a much higher network capacity and guaranteed latency, thanks to both physical improvements (e.g., through techniques such as wavelength division multiplexing) and logical/functional enhancements (e.g., the adoption of Software Defined Networks). Reliable and low-latency connections can be used to shift significant amount of workload from one site to another through dedicated networks or even via regular Internet connections.

These considerations naturally lead to a hierarchical infrastructure [24]. In this scenario, the single DCs need to manage the local workload autonomously and communicate with each other to route and migrate VMs among them. Recently, we have presented the EcoMultiCloud solution [14], a hierarchical framework for the efficient distribution of the workload on a multi-site platform, which allows for an integrated and homogeneous management of heterogeneous platforms but at the same time preserves the autonomy of single sites. Through the self-organizing and adaptive nature of the approach, the Virtual Machines migrations are performed asynchronously, both location-wise and time-wise, and with a tunable rate managed by data center administrators. The EcoMultiCloud framework was firstly presented in [13], where it was compared to [19], the reference of non-hierarchical approaches that have full visibility about all VMs and servers. The results of this comparison showed that the hierarchical approach leads to performance improvements with respect to single layer algorithms, and in addition it offers notable advantages in terms of efficiency, scalability, autonomy of sites, overall administration, information management.

The EcoMultiCloud hierarchical architecture is composed of two layers:

- The lower layer is used to allocate the workload within single DCs: each site adopts its own strategy to assign VMs internally, with local consolidation algorithms (possibly different from site to site). The lower layer collects information about the state of the local DC, and passes it to the upper layer;

- The upper layer is able to exchange information among homogeneous and interconnected sites and drive the redistribution of VMs among the DCs, through inter-data center VM migrations.

The reference scenario is depicted in Figure 1, which shows the upper and lower layer for two interconnected data centers, as well as the main involved components.

At each data center, a data center manager (DCM) runs the algorithms of the upper layer, while the local manager (LM) performs the functionalities of the lower layer. Three basic algorithms must be designed and implemented at each DCM: (i) the assignment algorithm that determines the appropriate target datacenter for each new VM; (ii) the migration algorithm that determines from which source site and to which target site the workload should be migrated; (iii) the redistribution algorithm that periodically evaluates whether the current load balance is appropriate and, if necessary, decides whether an amount of workload should be migrated to/from another site. The assignment algorithm is used to route a new VM to the best target datacenter. However, the workload distribution may become inefficient

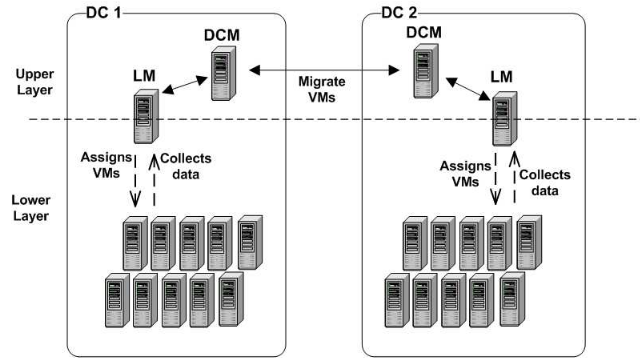

Figure 1: EcoMultiCloud **scenario: upper and lower layer of two interconnected data centers.**

when the conditions change, e.g., the overall load or the price of energy may vary in one or more data centers. In such cases inter-data center VM migrations are performed to redistribute the workload considering the new conditions.

2.1 Main Features of EcoMultiCloud

The EcoMultiCloud framework adds a higher level of inter-datacenter workload management to the tools that a previous solution, EcoCloud [23], has developed to consolidate virtual machines (VMs) onto the fewest number of servers on a single data center. The focus on workload management in a multi-site scenario is coherent with other efforts in this area, including a strong push from all the big players. EcoMultiCloud aims to offer the following advantages with respect to non-hierarchical solutions:

- Scalability: the EcoMultiCloud solution tackles the overall problem of workload management by decomposing it into two subproblems: intra-cloud consolidation within each data center and optimal distribution among multiple remote data centers. This approach allows to notably improve the scalability with respect to the number of sites and the overall number of applications and virtual machines. This feature is particularly advantageous in the geographical data center scenario, described in Section 3, especially when the number of interconnected data centers is large;

- Autonomous management: the EcoMultiCloud architecture leaves most of the intelligence to single DCs. At the lower layer, each DC is fully autonomous, and can manage the internal workload using either EcoCloud or any other consolidation algorithm. At the upper layer, coordinating decisions, for example about the necessity of migrating a fraction of the workload from one site to another, are taken combining the information related to single DCs;

- Modularity: the hierarchical architecture allows the protocols and algorithms of the two layers, upper and lower, to be separately modified and improved without affecting the behavior of the other layer. While the upper layer is expected to more stable, since it requires the consensus of all the data centers, this feature allows each data center to modify the internal management of the workload, as long as the interface between the upper and the lower layer is respected;

- Risk Mitigation: EcoMultiCloud performs preemptive VM migrations to prevent overload of hardware resources. The maximum utilization thresholds (of CPU, RAM etc.) can be tuned to optimize efficiency and maximize the Quality of Service (QoS) at the same time. Risk mitigation is essential in all the application scenarios described in Section 3, specifically in Hybrid Cloud, since on-premises resources are interconnected with resources held by third parties;

- Real-Time Adaptation: capacity of quickly adapting the assignment of VMs (e.g., through live migrations) to workload changes. EcoMultiCloud has the ability of quickly adapting the assignment of VMs to workload changes, e.g., through asynchronous live migrations and/or proper allocation of new applications. Real-time adaptation is a very important requirement in all the application scenarios described in Section 3, especially in the cases that client applications need to be moved from one site to another with minimum downtime;

- Combination of multiple business goals: the EcoMultiCloud solution can be specialized and tuned to match diverse business goals, depending on the administrators' requirements. Some of these goals are load balancing among data centers, reduction of monetary costs, consumed energy and carbon emissions, etc. This feature is particularly advantageous for the geographical data center scenario and the Hybrid Cloud scenario, as described in Section 3.

3. BUSINESS SCENARIOS FOR ECOMULTICLOUD

In the following, we describe the main business scenarios for which hierarchical solutions, and in particular EcoMultiCloud, are suitable, i.e.: geographical data centers, Hybrid Cloud and software defined data centers. We will highlight which specific benefits of hierarchical solutions can be exploited in each of the three business scenarios. The three main benefits considered here are: (i) *high scalability*, because the problem is decomposed into two layers and its size is thus notably reduced; (ii) *modularity*, since the algorithms of the lower layer can be designed independently from those of the upper layer, and vice versa; (ii) *autonomy* of data centers, since any data center can choose its own solution for the internal management of the workload.

3.1 Geographical Data Centers

Globalization, security and disaster recovery considerations are driving business to diversify locations across multiple regions. In addition, organizations are looking to distribute workloads between computers, share network resources effectively and increase the availability of applications. With the ultimate goal of eliminating downtime and sharing data across regions, companies and institutions are deploying geographically dispersed data centers to minimize planned or unplanned downtime. The size of these data centers vary depending on the organization's business needs. Some enterprizes are building multiple data centers, each having up to 500,000 square feet of floor space that house thousands of servers, storage and networking equipment including switches, routers, tape backup libraries and disk arrays. These data centers run standard and mission critical applications and continuously process and store data.

An infrastructure with mobile, active virtual machines can respond to new requirements much more quickly and very cost-effectively. Cloud computing users can gain even greater advantages from mobile virtual machines when they can be moved not only within a Cloud data center, but over greater distances to connect multiple Cloud data centers. Virtual machine movement between Cloud data centers enables applications such as disaster recovery and data replication. Intra- and inter-cloud data center migration of virtual machines dictate very specific network design requirements: networks must be 'flat'- which means that they have to be designed to connect potentially thousands of physical servers hosting tens of thousands of virtual servers within a single layer 2 network domain.

Geographical data centers are deployed by major Cloud service providers, such as Amazon, Google, and Microsoft, to match the increasing demand for resilient and low-latency Cloud services, or to interconnect heterogeneous data centers owned by different companies in so called "Inter-Cloud" scenarios. Solutions for the management of geographical environments are being introduced by major vendors (e.g. Cisco's Intercloud Fabric or VMware's vCloud Air) to build highly secure distributed Clouds and extend private data centers to public Clouds as needed. In such environments, data centers offer different and time-varying energy prices, and workload variability is experienced both within single sites and across the whole infrastructure. Applications and services are forwarded on demand to the best available location(s) and with consistent network and security policies.

Figure 2: Global Data Center Equipment Market Revenue from 2012 to 2020.

Figure 2 shows that the global data center equipment market is forecast to reach USD 72.07 billion by 2020 and is forecast to grow at a CAGR of 12.9% from 2014 to 2020[2]. While it is difficult to predict the fraction of data centers that will have more than one sites, the figure confirms that the worldwide market addressable by software solutions for the management of these environments is huge.

Geographical data centers is the business scenario for which the hierarchical EcoMultiCloud solution has primarily been designed, thanks to the significant benefits that the solution can offer. Indeed, EcoMultiCloud allows for migrating the workload from remote data centers depending

[2]Source: LockerDome, Company Annual Report, Industrial Journals, https://lockerdome.com

on the environmental conditions and the business priorities. The need for *autonomous management* is self-explanatory in geographical data centers, since it allows different algorithms to be adopted in the lower layers of different sites. The *modularity* feature is also very important, because it offers the concrete opportunity of designing the appropriate strategy at the upper layer, independently from the adopted algorithms at the lower layer, and vice versa. The *scalability* of hierarchical solutions is also essential, due to the presence of a huge number of servers, up to tens of thousands, in multi-site data centers, and an even larger number of running virtual machines.

3.2 Hybrid Cloud

In the Hybrid Cloud scenario, private and public data centers are interconnected and integrated. An IDC (International Data Corporation) study predicted that in 2015, the majority of chief information managers will move to Hybrid Cloud and, as part of this migration, existing deficiencies in service management will become evident, forcing investment in automation and consumption of externally managed services as alternatives to on-premises deployment of Cloud [9]. Key trends that will emerge in the Hybrid Cloud market include demand for integrated software development methods that stress communication, collaboration, integration, automation and measurement of cooperation between software developers and other IT professionals. On March 24th, 2014, Cisco announced plans to build the world's largest global inter-cloud together with a set of partners to create a network of Clouds called "Intercloud Fabric" [8]. Cisco's Intercloud strategy will essentially facilitate the move toward Hybrid Cloud, where private and public Clouds are integrated. VMware, with its Cross-Cloud Architecture, aims to supports ITŠs adoption of public or private Clouds without creating Cloud silos [31]. This solution will provide a control plane for common management, policies, networking, and security across private and public Clouds. IT organizations will be able to extend the capabilities of private Cloud technologies, discover what services exist across different Clouds, and enforce security and governance while efficiently managing costs.

The adoption of hierarchical solutions like EcoMultiCloud can be highly beneficial in Hybrid Cloud infrastructures, for example in the case that one or several sites are the former asset of an acquired company, or are hosted by co-located multi-tenant facilities. The high *scalability* of EcoMultiCloud can be a significant advantage for big companies that need to cope with a highly dynamic workload. The migration of applications and VMs from the private to the public components and vice versa, as envisioned by EcoMultiCloud, can help to tackle this issue efficiently. The *autonomy* characteristic is also important, as it is possible to use different algorithms of the lower layer in the private and in the public components. Moreover, *modularity* can also become a key feature in this scenario, because we can envision the proposal of software solutions that will try to build a uniform upper layer that helps to facilitate the exchange of workload among private and public facilities.

3.3 Software Defined Data Centers

Software data centers or software-defined data centers (SDC or SDDC) are data centers where all the infrastructure is virtualized and delivered as a service. The control of the data center is fully automated by software, meaning that the hardware configuration is maintained through intelligent software systems. This is in contrast to traditional data centers where the infrastructure is typically defined by hardware and devices. Software-defined data centers are considered by many to be the next step in the evolution of virtualization and Cloud computing as it provides a solution to support both legacy enterprize applications and new Cloud computing services. There are three core components of the software-defined data center: network virtualization based on Software Defined Networking (SDN), server virtualization and storage virtualization. A business logic layer is also required to translate application requirements, SLAs, policies and cost considerations [32].

A primary goal of the software-defined data center is to help IT organizations be more agile and deliver rapid, Cloud-like services to users. The concept of software defined data center involves data centers that can combine private, public, and Hybrid Clouds. Software data centers will manage applications and all the required resources – including compute, storage, networking and security features – to create a "logical" infrastructure. Commonly cited benefits of software-defined data centers include improved efficiencies from extending virtualization throughout the data center; increased agility from provisioning applications quickly; improved control over application availability and security through policy-based governance; and the flexibility to run new and existing applications in multiple platforms and Clouds. Furthermore, software data centers improve security by giving organizations more control over their hosted data and security levels, compared to security provided by hosted Cloud providers. An idea on the market potential of software data centers can be provided by the market expectations of SDN, which are shown in Figure 3 for the five most important regional areas[3].

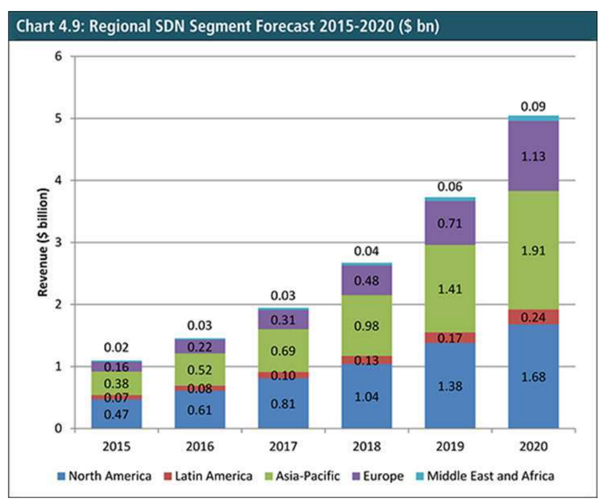

Figure 3: Regional SDN Segment Forecast from 2015 to 2020.

Hierarchical solutions like EcoMultiCloud will benefit from the software-driven management of the workload provided by software defined data centers and will smoothly in-

[3]Source: https://www.visiongain.com/Report/1457/Software-Defined-Data-Centre-(SDDC)-Market-Report-2015-2020.

tegrate with Software Defined Networking solutions. Specifically, a fraction of multi-site data centers will adopt the SDN technology, especially the large data centers owned by big companies. For this reason, the *scalability* characteristic is a key factor for the software data center business scenario, while the *modularity* and *autonomy* characteristics naturally match the objectives of the software data center paradigm.

4. INTEGRATION WITH BIG DATA MANAGEMENT

Data centers and, more generally, Cloud Computing have a tight relation with big data management techniques and algorithms (e.g., [2, 12]). This manly because modern big data applications are now mostly delivered via Clouds (e.g., [7, 30]). According to this evidence, in this section we briefly explore some hot-topics related to the integration of Clouds with big data management, for which the workload-distribution solutions discussed in our paper can, from a side, be beneficial and, from another side, smoothly support their full-integration. Besides, it is clearly enough that Cloud/big-data integration now represents a fundamental component of every modern big data application.

From this, a plethora of research problems that heavily found on this integration derive. First, storage alternatives for effectively and efficiently representing big data over Clouds are studied (e.g., [6, 27]). Similarly, the problem of developing applications that manage big data over Clouds by advocating the well-known *Hadoop* framework [33] for performance improvement has received considerable attention to date (e.g., [16, 17]). Performance (e.g., [35, 29]) and uncertain data management (e.g., [11, 18]) are correlated to the design and developing of these applications, and still capture the attention of larger and larger communities of researchers.

When processing large-scale amounts of big data over Clouds, novel paradigms arise. Among these, some interesting are: collaborative approaches (e.g., [34]), multi-store solutions (e.g., [4]), and, also, human-centric computing frameworks (e.g., [20]). These clearly denote a florid family of initiatives that stimulate further research efforts in the field.

Finally, privacy-preserving methods for managing big data over Clouds (e.g., [10, 36]) play a first-class role in the investigated scenario. Indeed, a possible interesting extension of workload-distribution over Hybrid Clouds would be that of considering the privacy of data as a basic criterion to guide the distribution task itself.

5. CONCLUSIONS AND FUTURE WORK

This paper has examined the main characteristics of the software platforms that will need allocate and distribute the workload on multi-site data centers. It emerged that the most efficient solutions should be built upon a hierarchical architecture, with a lower layer that manages the workload at the local data center, and an upper layer that orchestrates the operations of multiple data centers, and is able to migrate portions of the workload from one data center to another. Furthermore, we have presented and analyzed the main business scenarios for which hierarchical solutions, and in particular EcoMultiCloud, are particularly suitable, i.e.: geographical data centers, Hybrid Cloud and software defined data centers. For each of these scenarios, we highlighted the specific benefits of hierarchical solutions in terms of (i) *autonomy* (any data center can choose its own solution for the internal management of the workload); (ii) *modularity*, (the algorithms of upper and lower layers are designed and can modified independently from each other); (iii) *scalability*, since the problem is decomposed into two layers and the overall size of the problem is reduced. In addition to this, we have explored some hot-topics related to the integration of Clouds with big data management, for which the workload-distribution solutions discussed in our paper play a critical role. Our current efforts are devoted to develop algorithms and protocols for the efficient connection of hierarchical frameworks to the energy grids. This issue has rapidly emerging (e.g., [28], and its main objective is to improve the efficiency of energy distribution and increase the fraction of green energy, for example by dynamically moving portions of the workload to data centers that are close to sites where renewable energy is being produced.

Acknowledgements

This project has received funding from the European Union's Horizon 2020 research and innovation program under grant agreement No. 712363

6. ADDITIONAL AUTHORS

7. REFERENCES

[1] S. Agarwal and A. Nath. Desktop virtualization and green computing solutions. In *Second International Conference on Soft Computing for Problem Solving*, pages 1439–1449. Springer, 2014.

[2] D. Agrawal, S. Das, and A. El Abbadi. Big data and cloud computing: New wine or just new bottles? *PVLDB*, 3(2):1647–1648, 2010.

[3] A. Beloglazov, J. Abawajy, and R. Buyya. Energy-aware resource allocation heuristics for efficient management of data centers for cloud computing. *Future Generation Computer Systems*, 28(5):755–768, 2012.

[4] C. Bondiombouy, B. Kolev, O. Levchenko, and P. Valduriez. Multistore big data integration with cloudmdsql. *T. Large-Scale Data- and Knowledge-Centered Systems*, 28:48–74, 2016.

[5] M. Cardosa, M. R. Korupolu, and A. Singh. Shares and utilities based power consolidation in virtualized server environments. In *Proceedings of the 11th IFIP/IEEE Integrated Network Management (IM 2009)*, Long Island, NY, USA, June 2009.

[6] V. Chang and G. Wills. A model to compare cloud and non-cloud storage of big data. *Future Generation Comp. Syst.*, 57:56–76, 2016.

[7] Z. Chen, G. Xu, V. Mahalingam, L. Ge, J. Nguyen, W. Yu, and C. Lu. A cloud computing based network monitoring and threat detection system for critical infrastructures. *Big Data Research*, 3:10–23, 2016.

[8] Cisco. Cisco and partners to build world's largest global intercloud, 2014. http://newsroom.cisco.com/press-release-content?articleId=1373639.

[9] I. D. Corporation. Idc reveals cloud predictions for 2015, 2014. https://www.idc.com/.

[10] A. Cuzzocrea. Privacy and security of big data: Current challenges and future research perspectives. In *Proceedings of the First International Workshop on Privacy and Secuirty of Big Data, PSBD@CIKM 2014, Shanghai, China, November 7, 2014*, pages 45–47, 2014.

[11] A. Cuzzocrea, C. K. Leung, and R. K. MacKinnon. Mining constrained frequent itemsets from distributed uncertain data. *Future Generation Comp. Syst.*, 37:117–126, 2014.

[12] A. Cuzzocrea, D. Saccà, and J. D. Ullman. Big data: a research agenda. In *17th International Database Engineering & Applications Symposium, IDEAS '13, Barcelona, Spain - October 09 - 11, 2013*, pages 198–203, 2013.

[13] A. Forestiero, C. Mastroianni, M. Meo, G. Papuzzo, and M. Sheikhalishahi. Hierarchical approach for green workload management in distributed data centers. In *20th International European Conference on Parallel and Distributed Computing, Euro-Par 2014*, volume 8805 of *LNCS*, pages 323–334, Porto, Portugal, August 2014. Springer.

[14] A. Forestiero, C. Mastroianni, M. Meo, G. Papuzzo, and M. Sheikhalishahi. Hierarchical approach for efficient workload management in geo-distributed data centers. *IEEE Transactions on Green Communications and Networking*, August 2016. Early Access, https://doi.org/10.1109/TGCN.2016.2603586.

[15] P. Graubner, M. Schmidt, and B. Freisleben. Energy-efficient virtual machine consolidation. *IT Professional*, 15(2):28–34, 2013.

[16] R. Gu, X. Yang, J. Yan, Y. Sun, B. Wang, C. Yuan, and Y. Huang. Shadoop: Improving mapreduce performance by optimizing job execution mechanism in hadoop clusters. *J. Parallel Distrib. Comput.*, 74(3):2166–2179, 2014.

[17] D. H. Hagos. Software-defined networking for scalable cloud-based services to improve system performance of hadoop-based big data applications. *IJGHPC*, 8(2):1–22, 2016.

[18] Q. He, H. Wang, F. Zhuang, T. Shang, and Z. Shi. Parallel sampling from big data with uncertainty distribution. *Fuzzy Sets and Systems*, 258:117–133, 2015.

[19] A. Khosravi, S. Garg, and R. Buyya. Energy and carbon-efficient placement of virtual machines in distributed cloud data centers. In F. Wolf, B. Mohr, and D. Mey, editors, *Euro-Par 2013 Parallel Processing*, volume 8097 of *Lecture Notes in Computer Science*, pages 317–328. Springer Berlin Heidelberg, 2013.

[20] H. Kim, J. H. Park, and Y. Jeong. Human-centric storage resource mechanism for big data on cloud service architecture. *The Journal of Supercomputing*, 72(7):2437–2452, 2016.

[21] F. Kong and X. Liu. A survey on green-energy-aware power management for datacenters. *ACM Comput. Surv.*, 47(2):30:1–30:38, Nov. 2014.

[22] Z. Liu, M. Lin, A. Wierman, S. Low, and L. L. H. Andrew. Greening geographical load balancing. *IEEE/ACM Transactions on Networking*, 23(2):657–671, April 2015.

[23] C. Mastroianni, M. Meo, and G. Papuzzo. Probabilistic consolidation of virtual machines in self-organizing cloud data centers. *IEEE Transactions on Cloud Computing*, 1(2):215–228, July 2013.

[24] C. Mastroianni, D. Talia, and O. Verta. Evaluating resource discovery protocols for hierarchical and super-peer grid information systems. In *Proc. of the 15th Euromicro Int. Conf. on Parallel, Distributed and Network-Based Processing, PDP '07*, pages 147–154, Naples, Italy, February 2007.

[25] L. Rao, X. Liu, and W. Liu. Minimizing electricity cost: Optimization of distributed internet data centers in a multi-electricity-market environment. In *In Proc. of INFOCOM*, 2010.

[26] K. Schröder and W. Nebel. Behavioral model for cloud aware load and power management. In *Proc. of HotTopiCS '13, 2013 international workshop on Hot topics in cloud services*, pages 19–26. ACM, May 2013.

[27] K. Sekaran and P. V. Krishna. Big cloud: a hybrid cloud model for secure data storage through cloud space. *IJAIP*, 8(2):229–241, 2016.

[28] A. Singhee, M. A. Lavin, U. Finkler, F. Heng, J. M. Qu, and S. Hirsch. HAMS: A memory-efficient representation of power grids using hierarchical and multi-scenario graphs. In *Proceedings of the 2015 ACM Sixth International Conference on Future Energy Systems, e-Energy 2015, Bangalore, India, July 14-17, 2015*, pages 171–178, 2015.

[29] C. Smowton, A. Balla, D. Antoniades, C. J. Miller, G. Pallis, M. D. Dikaiakos, and W. Xing. A cost-effective approach to improving performance of big genomic data analyses in clouds. *Future Generation Comp. Syst.*, 67:368–381, 2017.

[30] R. Tudoran, A. Costan, and G. Antoniu. Overflow: Multi-site aware big data management for scientific workflows on clouds. *IEEE Trans. Cloud Computing*, 4(1):76–89, 2016.

[31] VMware. Vmware's cloud strategy: Hybrid cloud, 2016. https://www.vmware.com/solutions/cloud-computing.html#hybridcloud.

[32] Webopedia. SDDC - software-defined data center, 2016. http://www.webopedia.com/TERM/S/software_defined_data_center_SDDC.html.

[33] T. White. *Hadoop: The Definitive Guide*. O'Reilly Media, Inc., 1st edition, 2009.

[34] Q. Xia, Z. Xu, W. Liang, and A. Y. Zomaya. Collaboration- and fairness-aware big data management in distributed clouds. *IEEE Trans. Parallel Distrib. Syst.*, 27(7):1941–1953, 2016.

[35] B. Yu, A. Cuzzocrea, D. H. Jeong, and S. Maydebura. On managing very large sensor-network data using bigtable. In *12th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing, CCGrid 2012, Ottawa, Canada, May 13-16, 2012*, pages 918–922, 2012.

[36] X. Zhang, W. Dou, J. Pei, S. Nepal, C. Yang, C. Liu, and J. Chen. Proximity-aware local-recoding anonymization with mapreduce for scalable big data privacy preservation in cloud. *IEEE Trans. Computers*, 64(8):2293–2307, 2015.