CHAPTER 1

STRATEGIES FOR INCREASED ENERGY AWARENESS IN CLOUD FEDERATIONS

G. Kecskemeti[1], A. Kertesz[2], A. Cs. Marosi[2] and Zs. Nemeth[2]

[1]Distributed and Parallel Systems group, Institute of Informatics, University of Innsbruck, Innsbruck, Austria

[2]Laboratory of Parallel and Distributed Systems, MTA SZTAKI Computer and Automation Research Institute, Budapest, Hungary

1.1 INTRODUCTION

Cloud computing encompasses many aspects of sharing software and hardware solutions, including computing and storage resources, application runtimes or complex application functionalities. The cloud paradigm changed the way people look at computing infrastructures. First, one does not need to be expert in infrastructure administration, operation and maintenance even if large scale systems are utilized. Second, the elasticity of Infrastructure as a Service clouds allow these systems to better follow the users' actual demands. However, there is also an adversary effect: the virtualized nature of these systems detaches users from several operational issues like energy efficient usage, that has been addressed previously in the context of parallel and distributed systems, and largely remains unnoticed.

The Cloud computing technology made a qualitative breakthrough as it is present in many consumer appliances from mobile phones to television sets and thus, about a quantitative explosion, too. The illusion of infinite resources towards the consumers however, raises severe issues with energy consumption; the higher levels of quality and availability require irrational energy expenditures; according to some experts the consumed energy of resources spent for idling represent a considerable amount [1]. Current trends are claimed to be clearly unsustainable with respect to resource utilisation, CO_2 footprint and overall energy efficiency. It is anticipated that further growth is objected by energy consumption furthermore, competitiveness of companies are and will be strongly tied to these issues.

Several European data centers maintain infrastructures similar to IaaS cloud systems and provide some (quasi-)dynamic services to their users (e.g., virtual server hosting or virtual private server offerings). Even in this case, addressing – the currently neglected – energy related issues (e.g., improving energy and environmental performance of their data centers) could increase the competitiveness of these data centers even in a non-cloud scenario.

Energy awareness is a highlighted research topic and there are efforts and solutions for processor level, component level and datacenter level energy efficiency. For instance, new energy efficient approaches were proposed to automate the operation of data centers behind clouds [2], so that they help with rearranging the virtualized load from various users. Thus, smaller sized physical infrastructure is sufficient for the actual demand and momentarily unused capacities can be switched off. Nevertheless, these approaches are applicable to single data centers only.

Nowadays, cloud providers operate geographically distributed data centers as demands like disaster recovery and multisite backups became widespread. Recent solutions hide the diversity of multiple clouds and form a unified federation on top of them. Therefore, today's large systems are composed of multiple service providers per se that need new approaches to ensure their overall energy-aware operation, on one hand. On the other hand there is an unexplored potential for energy-aware operation in federated and interoperable clouds. Our work is targeted at examining what new aspects of energy awareness can be exploited in federative schemes.

This chapter first identifies three scenarios that current energy aware cloud solutions cannot handle as isolated IaaS, but their federative efforts offer opportunities to be explored. These scenarios are centered around: (i) multi-datacenter cloud operator, (ii) commercial cloud federations, (iii) academic cloud federations. Based on these scenarios, we identify energy-aware scheduling policies to be applied in the management solutions of cloud federations. Among others, these policies should consider the behavior of independent administrative domains, the frequently contradicting goals of the participating clouds and federation wide energy consumption.

Our earlier work introduced the Federated Cloud Management architecture to provide a unified interface over multiple cloud providers [3]. This architecture consists of three main components: meta-brokering (inter-cloud scheduling), cloud brokering (intra-cloud scheduling) and a generic virtual machine image repository. In this work, regarding the meta-brokering component, we propose a new scheduling policy that minimizes the overall energy consumption of a federation while maintaining its performance. Concerning the cloud brokering component, we offer new heuristics for virtual machine management like early destruction, which maintains the energy efficient serving of service calls within a global limit and aggressively shuts down idling or surplus virtual machines (VM).

This chapter is organized as follows. First relevant research is reviewed in Sect. 1.2. Afterwards, Sect. 1.3 discusses three scenarios to highlight the issues that impede the energy-aware operation of current systems. Next, Sect. 1.4 discusses the identification of the necessary cloud extensions that enable energy conscious federations, and extends the Federated Cloud Management architecture with multi-level and energy-aware scheduling strategies. Finally, Sect. 1.5 provides our conclusion and pinpoints some future research directions.

1.2 RELATED WORK

Cloud federation refers to a mesh of cloud providers that are interconnected based on open standards to provide a universal decentralized computing environment, where everything is driven by constraints and agreements in a ubiquitous, multi-provider infrastructure. Until now, the cloud ecosystem has been characterized by the steadily rising hundreds of independent and privately managed heterogeneous cloud providers offering various services to their clients.

Buyya et al. [4] suggest a federation-oriented, just in time, opportunistic and scalable application services provisioning environment called InterCloud. They envisioned utility-oriented federated IaaS systems that are able to predict application service behavior for intelligent infrastructures with down- and up-scaling abilities. They list research issues of flexible service-to-resource mapping, user and resource centric quality of service (QoS) optimization, integration with in-house systems of enterprises, scalable monitoring of system components. They present a market-oriented approach to offer InterClouds including cloud exchanges and brokers that bring together producers and consumers. Producers are offering domain specific enterprise Clouds that are connected and managed within the federation with their Cloud Coordinator component. Celesti et al. [5] proposed an approach for establishing federations considering generic cloud architectures according to a three-phase model, representing an architectural solution for federation by means of a Cross-Cloud Federation Manager (CCFM), a software component in charge of executing the three main functionalities required for a federation. In particular, the component explicitly manages: i) the discovery phase in which information about other clouds are received and sent, ii) the match-making phase performing the best choice of the provider according to some utility measure and iii) the authentication phase creating a secure channel between the federated clouds.

Marshall et al. proposed an IaaS cloud solution to elastically extend physical clusters with cloud resources [6]. They created a so-called elastic site manager on top of Nimbus, which interfaces directly with local cluster managers and three different policies were examined for elastic site addition. Rochwerger et al. [7] introduced the Reservoir project and its federated IaaS cloud management model, and proposed that commercial cloud providers could also temporarily lease excess capacities during high-demand periods. They investigated the following problems faced by federated cloud solutions: (i) dynamic service elasticity, (ii) admission control, (iii) policy-driven placement optimization, (iv) cross-cloud virtual networks (v) cross-cloud monitoring, and (vi) cross-cloud live migration. Bernstein et al. [8] defined two use case scenarios that exemplify the problems of multi-cloud systems like (i) VM Mobility where they identify the networking, the specific cloud VM management interfaces and the lack of mobility interfaces as the three major obstacles and (ii) storage interoperability and federation scenario in which storage provider replication policies are subject to change when a cloud provider initiates subcontracting.

J. L. Lucas-Simarro et al. [9] proposed different scheduling strategies for optimal deployment of services across multiple clouds based on various optimization criteria. The examined scheduling policies include budget, performance, load balancing and other dynamic conditions, but they neglected energy efficiency, which is the aim of our investigation.

Regarding energy efficiency in a single cloud, Cioara et al. in [2] introduced an energy-aware scheduling policy to consolidate power management by using reinforcement learning techniques to bring back the service center in an energy efficient state. Cardosa et al. [10] presented a novel suite of techniques for placement and power consolidation of VMs in data centres taking advantage of the min-max and shares features inherent in virtualization technologies, like VMware and Xen. These features let one specify the minimum and maximum amount of resources that can be allocated to a VM, and provide a shares based mechanism for the hypervisor to distribute spare resources among contending VMs. Lee et al. [11] discuss service request scheduling in Clouds based on achievable profits. They propose a pricing model using processor sharing for composite services in Clouds. Berral et al. [12] present a framework to address energy efficiency using an intelligent consolidation methodology, which applies various techniques such as machine learning on scheduling algorithms to improve server workload predictions, power aware consolidation algorithms, and turning off spare servers and thereby saving energy in a data center. However, their approach is limited to a private datacenter, and does not consider hybrid Clouds with engineering approaches like federation. Feller et. al. [13,14] proposed energy management algorithms and mechanisms of a novel holistic energy-aware VM management framework called Snooze for private clouds. Their solution uses power meters to monitor energy usage of cloud resources, and estimate the resource usage of VMs. Their mechanisms address VM placement, relocation and migration by keeping VMs on as few nodes as possible. This solution is able to dynamically consolidate the workload of a software and hardware heterogeneous large-scale cluster composed of resources using the virtualization. Also, IBM has proposed pMapper [15], which is a power-aware application placement controller in the context of an environment with heterogeneous virtualized server clusters. The placement component of the application management middleware takes into account the power and migration costs in addition to the performance benefit while placing the application containers on the physical servers. These approaches are focusing on consolidating power usage mostly within a single cloud. On the other hand, our goal is to consolidate power usage within a cloud federation by redistributing VM calls and utilizing federation-wide VM management policies.

Service level agreement (SLA) management is also an important issue in Clouds. An autonomic SLA violation detection solution is presented in [16] by V. Emeakaroha et al., that can be used to minimize user interaction. We also try to minimize user involvement in energy efficient service provisioning over multiple clouds by the use of SLAs.

In this chapter we propose different approaches for enabling an energy-aware cloud federation with various resource sharing strategies (e.g., energy-aware scheduling policies to affect carbon emissions). We also investigate, how service integrations among different providers can be performed with the application of these strategies, in order to enable enhanced reliability, scalability and utilization to broaden the market for smaller providers by bursting and enabling outsourcing towards such providers that could not sell their entire capacity previously.

1.3 SCENARIOS

1.3.1 INCREASED ENERGY AWARENESS ACROSS MULTIPLE DATA CENTERS WITHIN A SINGLE ADMINISTRATIVE DOMAIN

As small cloud providers and cloud startups are becoming more popular, they soon face user demands that cannot be satisfied with their current infrastructures. These user demands range from occasional needs for extreme amount of resources (compared to the provider's own infrastructure) to the need for multi-site virtual machine deployment options that allow disaster recovery. Providers thus, deemed to increase the size of their infrastructure by introducing multiple data centers on various locations and offering unprecedented amount of resources. Unfortunately, the increase in infrastructure size also increases resource heterogeneity. Hence, these providers have to deal with inhomogeneities by novel virtual machine placement and scheduling strategies. Current IaaS solutions offer these strategies allowing providers to focus their attention to non-technical issues like the increased operating cost of their data centers.

1.3.1.1 Facing the increased energy consumption Energy consumption is a major component of operating costs. Despite its significance, current IaaS clouds barely provide energy-aware solutions. Providers are restricted to reduce their consumption at the hardware level – independently from the IaaS. These reductions range from the use of more energy efficient computer components to the upgrade of their heating, ventilation and air conditioning (HVAC) systems to increase their power usage efficiency (PUE). Although these improvements are crucial, the energy consumption could also be significantly reduced by software means in over-provisioned IaaS systems where more physical resources are available at the provider side than actually requested by users. Over-provisioning is a key behavior at smaller sized providers that offer services for users with occasional peaks in resource demands. We consider a provider small, if the number of its customers with such requirements does not reach approximate uniform distribution throughout the year. To reduce their energy costs, these providers should minimize their over-provisioning while they maintain a fluid experience towards their customers without violating the previously agreed service level. Energy consumption could be reduced with software techniques focusing on intra- and inter-datacenter issues.

First, let us consider intra-datacenter issues. State-of-the-art techniques that reduce over-provisioning range from basic ones, like switching on/off unused parts of the infrastructure, and to more elaborate ones, like the use of load migration between resources to reduce resource usage fragmentation. Although, some significant research efforts were devoted to the investigation of these techniques, they are not adopted by IaaS solutions (and as a consequence by cloud providers). Even nowadays, the adoption curve is still in its early stages because of several issues, e.g.: (i) computers that frequently switching on/off have smaller mean time between failures (MTBF), (ii) switching on/off introduces considerable amount of delays in infrastructure provisioning, (iii) weak migration support by underlying technologies, (iv) frequently large cost of migration, (v) migration might cause disruptions in service level, and (vi) providers usually do not apply software energy meters that could continuously monitor their infrastructure's consumption and thus support decisions towards over provisioning reduction.

Next, energy awareness raises further issues, if the provider has multiple datacenters. The intra-datacenter techniques are often not applicable in inter-datacenter situations (i.e., operations between datacenters on distant geographic locations). In such situations, the cost of migration significantly rises and frequently causes service level degradations for the affected customers. Also, providers often have a wide variety of energy sources to choose from for each of their data centers. This variety of choices is often temporal, and each energy source has different price and CO_2 emissions for a given amount of consumed energy. For example, recent wind activity could significantly change a datacenter's implicit CO_2 emission, if a wind turbine is amongst its energy sources. Despite the fact that even providers with a single datacenter could introduce additional usage policies to drive their users towards greener operations, these policies may lead to significant SLA changes (e.g., offering resources only for such time periods, when wind turbines are the main energy source for the datacenter). When multiple datacenters are at the disposal of the provider, these datacenters open new possibilities for the provider to maintain the service level while still increase energy efficiency of the user loads. These possibilities include cross datacenter energy-aware virtual machine placement strategies and scheduling prioritization of datacenters with green energy surplus.

1.3.1.2 Increasing green operations of inter-datacenter constructs The diverse locations of the available datacenters of a provider increase the likeliness of having one or more datacentres with available CO_2 emission free energy sources. In such cases customers with no specific requirements on resource location could be directed and hosted in these datacenters (e.g., the IaaS *scheduler could prioritize* the greener datacenters). As some green energy sources are quite spontaneous, the IaaS is advised to continuously check for better provisioning options even after a customer is directed to a particular datacenter. If the variety of potential energy sources tends to be less green, the datacenter's future resource schedule must be proactively rearranged. For example, if the agreed service level is not violated, the IaaS may migrate some workload.

A consumption profile defines how a particular resource behaves energy-wise under specific types of load and this information may offer differentiated solutions in energy-awareness. Even with strong central administration (i.e. hardware purchases are controlled for all the datacenters of the provider by a single authority), the multitude of datacenters increase the heterogeneity in the offered resources. For example, despite the administrative efforts, small differences in newly bought hardware are inevitable. A seemingly minuscule difference in the used resources (e.g. processor stepping) could lead to significant changes in the energy consumption profiles of the datacenters. IaaS solutions could reduce the overall energy consumption of the provider's infrastructure more efficiently by taking into consideration these consumption profiles across all the provider's datacenters. Therefore, to reduce the overall energy consumption of the entire cloud system, virtual machine placement policies have to be enhanced with awareness of available consumption profiles.

Albeit both the CO_2 emission and the diversification of energy sources is important for increased greenness of larger scale cloud providers, these details are not accessible and not even offered to be used by the IaaS solutions today. Fortunately, because of the strength of their central administration (i.e. all of their data centres are within the same administrative domain), these providers have a chance to enforce emission and consumption profile collection. Therefore, a natural next step would allow them to increase the greenness of their operations by offering the collected data to be used by IaaS solutions. The availability of these data allows future IaaS solutions to make energy conscious decisions even without the user's consent.

1.3.1.3 The view of the energy conscious user Nowadays, more and more users are getting energy-conscious and try to integrate green aspects into their requirements towards the cloud infrastructure providers. More and more providers advertise their green practices, but they offer minimal control over the greenness of resources granted for a particular consumer. Currently, users could assume that during the fulfillment of their requests the provider has the least amount of energy consumed, when its dynamic pricing scheme (e.g., spot prices at Amazon Elastic Compute Cloud[1]) indicates the smallest price/resource. This approach however, is not applicable with providers who do not have dynamic pricing, while it also leads to false concepts, because the pricing scheme might reflect other factors than energy consumption. For those users who plan to ensure some level of energy consumption reduction on an arbitrary cloud infrastructure, an alternative approach could be the optimization of their applications (even on the source code level if necessary) with energy awareness in mind (e.g. use of [17]). This approach is however, not practical in most of the user scenarios, and does not even solve the root of excessive energy consumption. Therefore, there is a need for some novel techniques that could provide measures or estimates for greenness and energy awareness. Through

[1] http://aws.amazon.com/ec2/

the use of these techniques, energy conscious users should be able to determine the circumstances, when their tasks will have the least effect on the environment or on the overall consumption of the provider.

1.3.2 ENERGY CONSIDERATIONS IN COMMERCIAL CLOUD FEDERATIONS

In the previous subsection, we have discussed how a single cloud provider could increase its energy efficient behavior, while we have also introduced the way energy conscious users should approach these providers. However, multiple cloud providers add further challenges to the picture, regarding the improvement of energy consumption and greenness.

1.3.2.1 Challenges for cloud infrastructure providers within a federation Federations can be formed in various ways but commercial cloud providers are obviously driven by clear financial benefits. These providers prefer to be the primary contact to their users thus, commercial federations will mostly be formed by those cloud providers with a large enough user base. These larger providers will make contracts with some smaller ones to serve as an outsourcing target, in case the users of the large provider request some special kind or amount of resources. Although these contracts will not bring the smaller providers within the same administrative domain, they simplify the interfacing between two providers and define the service level that is necessary for the larger provider to fulfill its SLA requirements.

With the help of these contracts the large providers can introduce new policies which determine outsourcing to their contractual partners. The initial decision of the large provider is based purely on its partners prices. However, with the introduction of green policies this initial simple policy might change:

- As legislation proceeds towards cloud providers, large providers will soon face CO_2 quotas. With these quotas, the decisions made by large providers are not simply based on local and partner prices but they also include the price of their carbon credits (or their actual price in emission trading systems). This step does not inherently reduce the emissions, while performing user requests as the smaller provider might have different operating costs or could just value its available carbon credits much more cheaper. Therefore, in case the users ask for specific guarantees in terms of CO_2 emissions, smaller providers must share their CO_2 emissions for their offered resources. We assume, the number of such users will rise over time as CO_2 quotas are introduced more widely (i.e., more and more infrastructure users have their own CO_2 emission caps which should include the emitted CO_2 even at the cloud provider side).

- Even before the legislation reaches providers, they can start experimenting and offering green options to their users. As a result, large providers can see the demand for green resources without immediate investments. In such environments, the federated partners could even compete on the level of greenness offered to the large provider.

However, in both cases the providers have to trust each other regarding the reported CO_2 emissions. Thus, there is an increasing need for third parties (e.g., auditors as envisioned by the European Commission in [18]) to independently assess the energy and CO_2 efficiency of a particular infrastructure. This third party infrastructure evaluation increases the providers credibility, and enables the construction of new decision making tools that use the public evaluation records (if they are offered to the general public).

1.3.2.2 Pursuing energy awareness by users of cloud federations Although federations created by large providers could enable a simple way for users to reduce their ecological footprint, the decisions made by a large provider are not necessarily the most beneficial for users. Thus, users might try make decisions themselves. This is especially important when there are providers worth considering but not present in federations. Therefore a user may have the incentive to create a different kind of federation based on providers a particular user has access to. However, to construct such a federation, the users have to face several issues that were previously hidden by large providers. These issues range from the differences between the applied IaaS solutions by the accessible providers to the inability of interoperation between the various providers participating in the user's federation. To overcome these issues, users frequently turn towards third party federative solutions that are capable of hiding the differences of the providers underneath, but allow users to optimally reach them. Currently, these federative solutions are barely aimed at energy awareness.

The previously envisioned interfaces to publish CO_2 emissions and third party estimates for greenness allow the creation of basic energy consumption and CO_2 emission profiles for providers, datacenters or even individual cloud resources (that are provisioned to the user). These profiles pave the way for new and more energy conscious federative solutions that allow their users to construct federations of infrastructure clouds on a way that not only optimizes for the performance or the reliability of the user's tasks, but also their energy consumption and CO_2 emissions.

1.3.3 REDUCED ENERGY FOOTPRINT OF ACADEMIC CLOUD FEDERATIONS

Academic cloud infrastructures are offered with a non-profit effort to fellow academics. These infrastructures usually do not consider energy efficiency and green operations as priorities. Also, compared to commercial providers, academic cloud infrastructures are relatively small-sized. Academics tend to use infrastructures in bursts (causing likely overloaded infrastructures in certain periods of the year). Thus, they frequently find infrastructure limitations (e.g., temporal underprovisioning), even though these infrastuctures are often underutilized. To be more energy efficient, when the resources are barely utilized, academic IaaS should switch to an energy-saving virtual and physical machine management strategy. To reduce the effects of underprovisioning during computing bursts, academics could use federations over the currently existing small islands of academic clouds. Both the previously mentioned federation approaches are used within academic clouds. A few academic providers offer bursting capabilities which outsource some of the resource requests to other academic clouds on the users request. On the other hand, user-oriented federations are also supported with similar or the same federative solutions mentioned in the previous subsection. Unfortunately, because of their non-profit nature, academic providers do not have the incentive to reduce their carbon footprint or energy consumption. This is also the case for their users.

Since users will soon start using the federative solutions widely, energy efficiency could be enforced by the software accomplishing federations (federative/outsourcing solutions). If other circumstances are identical or indifferent, this software should prefer the more energy efficient cloud operators and base its decision in favor of a provider using third party estimates for greenness. With this approach, users will end up with more energy efficient/green resources without knowing it. At the provider side this could result in significantly less load on less energy aware academic clouds. The loss of demand for these providers would indicate that they should either increase their greenness or retire their infrastructure.

1.4 ENERGY-AWARE CLOUD FEDERATIONS

As we learned in [18], reducing the carbon footprint of European countries is a must, and expected by the European Commission, as well as to increase the number and size of European Cloud providers. By federating these providers, more competitive initiatives can be founded, that can be sophistically managed to meet these expectations. The general goal of the management layer in a Cloud federation is to distribute load among the participating cloud providers, to enhance user satisfaction by filtering out underperforming providers, and to schedule and execute service calls with minimized energy consumption within the selected IaaS system. To achieve this, we proposed an architecture called Federated Cloud Management (FCM – as introduced in [3]). In this holistic approach a two-level brokering solution is used: a meta-brokering component is used to direct service calls to providers, and then a cloud brokering component to map these calls onto an optimized number of virtual machines.

In order to address green aspects, i.e. energy consumption and CO_2 emissions, enhanced call scheduling algorithms should be developed. These approaches may focus on different aspects. At the meta-brokering layer, relying on an enhanced monitoring system within the federation, service executions can be directed to data centers of providers consuming less energy, having higher CO_2 emission quotas, or have produced less amount of CO_2 that expected within some timeframe. In this way, the issues raised in the second and third scenarios can be managed in practice.

At the cloud brokering layer, if the energy consumption parameters of a cloud *suddenly change*, there should be strategies to limit or move around calls and even (if necessary) VMs federation-wise. The changes here may mean the introduction of new hardware, or just switching on/off some parts of the datacenters, or changing the number of VMs. Realigning calls may not have immediate effects, however migration of VMs across the federation is also an energy consuming operation, that needs to be measured and considered when decisions are made (thus this operation should not happen only in case of really drastic changes). The *system should prefer* data centers where the difference between the highest load and the average load is small because a VM has the smallest impact on those resources [19]. In Sect. 1.4.3, we introduce strategies to be followed by a Cloud-Broker (CB) acting in this layer, which can solve energy utilization problems mentioned in all three scenarios.

1.4.1 AVAILABILITY OF ENERGY CONSUMPTION RELATED INFORMATION

In order to provide a solution that is able to handle the previously introduced scenarios, we extend our FCM concept [3] towards energy-awareness by taking into account energy consumption metrics for decision making at both levels. We introduce the following model into FCM.

Energy use of a computer system is usually expressed as the energy necessary for a certain unit of work to be performed, defined and measured differently at various levels. There are component (processor, node, network) and facility (e.g. a whole data center) related metrics that differ in granularity, detail and precision of observation. Commonly they can be characterized as an energy devoted to carry out certain activity (workload), where workload may be an instruction, a certain type and number of instructions, transactions, queries, storing or transferring a certain amount of data and so on, expressed as an $\frac{Energy}{Workload}$ ratio. Commonly both $Energy$ and $Workload$ are normalized to unit time hence,

$$\frac{Energy}{Workload} = \frac{\frac{Energy}{time}}{\frac{Workload}{time}} = \frac{Power}{Performance},$$

where performance may be expressed as MIPS, FLOPS, MFLOPS, BPS and other well known quantities. Energy efficiency, on the other hand, is a quantity that should be higher if the same amount of work is done using less energy or more work is done using the same energy. For this purpose, energy efficiency of a computer system can be characterized as $\frac{Workload}{Energy}$ or $\frac{Performance}{Power}$ [19] although, this definition is not entirely precise as efficiency should be the proportion of two energy quantities and hence, dimensionless.

In our case we have to select the right metrics carefully. Our work is aimed at a federated cloud system hence, it is pointless to take into consideration the energy used for each instruction and also, cummulative metrics of a whole datacenter would not give the necessary details for decision making. Since the cloud infrastructure is service-oriented, we consider the $\frac{Service}{Energy}$ or $\frac{Service\,throughput}{Power}$ fraction as our definition for energy efficiency.

In our FCM architecture we try to improve energy awareness by optimizing (i) the number of VMs per provider and (ii) direct service calls to these VMs. Each service is associated with a virtual appliance stored in a repository. Appliances are automatically transfered and deployed at a provider's IaaS when needed. In order to maintain an "energy saving state" of the whole federation, we keep the values of the above introduced *Call Throughput Per Power* (*ctpp*) $\frac{Service\,throughput}{Power}$ [$\frac{1}{W}$] of all participating providers within an interval. Later we will see how this metric is dependent on the number of VMs and the number of service calls.

1.4.2 SERVICE CALL SCHEDULING AT THE META-BROKERING LEVEL OF FCM

Regarding the energy efficient management of Cloud federations at the meta-brokering level, the user calls have to be balanced over the cloud providers (and their datacenters) participating in the federation based on their energy consumption parameters. Therefore the meta-broker component [20] of FCM forwards the service calls to providers having the lowest $ctpp$ value. The exact number of physical and virtual machines serving the actual load of requests for the user services should be intelligently managed locally by cloud brokers of the corresponding cloud infrastructure providers (discussed in detail in the next subsection).

Let $m(ctpp(i,t))$ denote the average service call throughput per unit of power for a service type i. Therefore in the simplest case the 'energy-balancing' algorithm chooses a provider j, for a given new service call for service i, from N participating providers based on the $ctpp$ measure with the following formula: $E_{tot}[j] * m(ctpp(i,t))[j]$. Once the call arrives to a Cloud-Broker that manages the selected provider, it tries to execute the call in a way that the overall energy consumption of the provider stays optimal.

Another important issue is the reduction of CO_2 emissions. In order to minimize this measure among the providers of a federation, one should find a way to measure the CO_2 emissions and modify the algorithm. A more aggressive, rebalancing strategy could also be used within the federation by migrating VMs from overloaded providers to less loaded ones. If we consider migrations to adapt to changing conditions during deployment and execution phases at the datacenters, we also have to consider the estimated costs of migrations from one provider to another. In this work we refrain from discussing algorithms considering migrations, instead we focus on energy-aware VM management strategies detailed in the next subsection.

1.4.3 SERVICE CALL SCHEDULING AND VM MANAGEMENT AT THE CLOUD-BROKERING LEVEL OF FCM

A Cloud-Broker (CB) manages VMs and dispatches received service calls in a single cloud of a provider. If the green aspects of the architecture are prioritized, the CB should apply VM management strategies more aggressively towards energy awareness. Incoming service calls are queued first, but served only if (a) they pass a threshold for waiting time (T_q^{up}) or (b) the energy usage for serving the same type of service calls exceeds the energy required for creating and terminating the VM by a given factor. VM's are terminated immediately when no service calls are queued for them.

We introduce the service efficiency factor $\alpha(i,t)$ as a measure of efficiency for serving queued service calls per service type in the CB. It is formulated as follows:

$$\begin{aligned}\alpha(i,t) &= \frac{\textit{Energy for creating or terminating VMs}}{\textit{Energy for running services}} = \\ &= \frac{vmc(i,t) * (E_{creat} + E_{term})}{scc(i,t) * \dfrac{1}{m(ctpp(i,t))}} = \\ &= \frac{vmc(i,t) * m(ctpp(i,t)) * (E_{creat} + E_{term})}{scc(i,t)} \qquad \left[\frac{1 * \dfrac{1}{[J]} * ([J] + [J])}{1}\right],\end{aligned} \tag{1.1}$$

where $vmc(i,t)$ denotes the number of VMs; $scc(i,t)$ represents the total number of service calls in the CB for a given service i at time t; E_{creat} and E_{term} denotes the energy required for starting and terminating VMs in the cloud. Note, that the value of α is calculated for each service type and should be recalculated as $scc(i,t)$ or $vmc(i,t)$ changes over time. Hence, $\alpha(i,t)$ is defined as:

$$\alpha(i,t) = \begin{cases} 0 & \text{if } scc(i,t) = 0 \\ \dfrac{vmc(i,t) * m(ctpp(i,t)) * (E_{creat} + E_{term})}{scc(i,t)} & \text{if } scc(i,t) > 0 \end{cases} \tag{1.2}$$

In certain cases these functions may be replaced by constants for the sake of simplicity, e.g., we may assume $scc = scc(i_0, \tau_0)$ and $vmc = vmc(i_0, \tau_0)$ for the decision making at $t = \tau_0$ or simply replace them by constants independent of i and t thus, α is defined as:

$$\alpha(i,t,vmc,scc) = \begin{cases} 0 & \text{if } scc = 0 \\ \dfrac{vmc}{scc} * m(ctpp(i,t)) * (E_{creat} + E_{term}) & \text{if } scc > 0 \end{cases} \tag{1.3}$$

Let α_i denote $\alpha(i,t)$, α_i^{up} denote an upper threshold and α_i^{low} denote a lower threshold for α_i at τ_0. These thresholds represent an upper and a lower efficiency barrier for service types at a CB instance and should be determined by the administrators.

Let vmc_i denote $vmc(i,t)$ and scc_i denote $scc(i,t)$ at τ_0:

$$\alpha_i = \alpha(i, \tau_0) \tag{1.4}$$
$$\alpha_i(vmc, scc) = \alpha(i, t_0, vmc, scc) \tag{1.5}$$
$$\alpha^{up} \geq max_{1 \leq i \leq z}(\alpha_i) \tag{1.6}$$
$$\alpha^{low} \geq min_{1 \leq i \leq z}(\alpha_i) \tag{1.7}$$
$$vmc_i = vmc(i, \tau_0) \tag{1.8}$$
$$scc_i = scc(i, \tau_0) \tag{1.9}$$

Let $\Theta = \{\theta_1, \theta_2, \ldots, \theta_z\}$ be a permutation of $\{1, 2, \ldots, z\}$ so that:

$$\begin{cases} \alpha_{\theta_i} \leq \alpha_{\theta_j} & \text{if } i < j \text{ and} \\ T_{q,\theta_i} \geq T_{q,\theta_j} & \text{if } i < j \text{ and } \alpha_{\theta_i} = \alpha_{\theta_j} \end{cases} \tag{1.10}$$

that is we define an order of services so that α values are in increasing order and waiting times in decreasing order should the α values be equal. We say that Θ is a CB instance represented as an ordered set of z type of services. Let $w_{ctpp}(i,t)$ denote the reciprocal of $m(ctpp(i,t))$ and let $w_{ctpp}(i)$ denote $w_{ctpp}(i,t)$ at τ_0:

$$w_{ctpp}(i,t) = \begin{cases} \dfrac{1}{m(ctpp(i,t))} & \text{if } m(ctpp(i,t)) \neq 0 \\ 0 & \text{else} \end{cases} \qquad [J] \tag{1.11}$$
$$w_{ctpp}(i) = w_{ctpp}(i, \tau_0) \qquad [J] \tag{1.12}$$

Let $\widehat{\gamma}$ denote the current total and γ^{up} the maximum energy as a function of call throughput, of all the VM's managed by the CB at τ_0:

$$\widehat{\gamma} = \sum_{i=1}^{z} vmc_i * w_{ctpp}(i) \qquad [J] \tag{1.13}$$
$$\gamma^{up} \geq max_{1 \leq i \leq z}\left(w_{ctpp}(i)\right) \qquad [J] \tag{1.14}$$

The value of γ^{up} is determined by the administrators and it is a characteristic of the cloud the CB is accessing. It represents the maximum energy consumption limit allowed to be consumed by the cloud: the definition (c.f., 1.14) only states that if it is set (> 0) then the cloud should be able to execute at least one service call of each available service type.

The mapping of queued service calls to available VMs is also the responsibility of the CB, however in this work we focus on VM management. We assume that (i) a VM is handling a single service call at a time; (ii) a VM is only able to serve a single type of service and (iii) the time required for VM startup and termination is insignificant compared to service call execution times. We propose an example algorithm (called *ERG-gamma* - c.f., Algorithm 1.1) based on the metrics defined in this section to demonstrate their usability. This algorithm is evaluated periodically, for the following we assume it is evaluated at τ_0 and Θ is re-calculated only before the execution of the algorithm, not continuously.

Algorithm 1.1

```
ERG-gamma {
        stopCond(k) {
            return (α_k > α^up ∨ γ̂ > γ^up ∨ scc_k = 0) ∧ vmc_k > 0
        }
        startCond(k) {
            return t_q(k) > 0 ∧ (γ̂ + w_ctpp(k) ≤ γ^up) ∧ α_k < α^low
        }
        for i = z → 1
        {
            while stopCond(θ_i)
            {
                stopVM(θ_i)
            }
        }
        for i = 1 → z
        {
            while startCond(θ_i)
            {
                startVM(θ_i)
            }
        }
}
```

Algorithm 1.1 (*ERG-gamma*) shows an example how to satisfy γ^{up}, α^{up} and α^{low} constraints by first stopping service instances with the highest α when required. The stop condition function ($stopCond(i)$) states that an instance must be terminated if the upper limits (α^{up} and γ^{up}) are reached, or if there are no more service calls for the service (scc_i). This ensures that the limits are honored and instances with no service requests are terminated. New instances for a service type i are started, if there are any service calls waiting ($t_q(i)$), the global energy limit (γ^{up}) allows their execution and the service efficiency factor is below the threshold.

This approach enables for cloud administrators defining an interval ($[\alpha^{low}, \alpha^{up}]$) for setting the service efficiency of their infrastructure and an overall limit for all services running combined. This allows maintaining the energy efficient serving of service calls within a global limit and also aggressively shutting down idling or surplus VM's. Administrators may choose the interval based on their knowledge of their system to conform to internal or third party requirements regarding energy awareness.

1.5 CONCLUSIONS

Current Cloud infrastructure solutions are very rigid and do not exhibit the flexibility and configurability potential to address energy consumption reductions. To be competitive, infrastructures are required to set up their available hardware to be offered publicly (similarly to already available commercial public cloud infrastructures) in a flexible and configurable way. Thus, there is a need for sophisticated infrastructure management techniques instantly applicable for providers in federations to reduce their overall carbon footprint. On the other hand there is an unexplored potential for energy-aware operation in federated and interoperable clouds. Our work is targeted at examining what new aspects of energy awareness can be exploited in federative schemes.

In this chapter we identified three scenarios that current energy-aware cloud solutions cannot handle as isolated IaaS, but their federative efforts offer opportunities to be explored. These scenarios include multi-datacenter cloud operator, commercial cloud federations and academic cloud federations. Based on these scenarios, we identified energy-aware scheduling policies to be applied in the management solutions of cloud federations. We applied these approaches in a Federated Cloud Management architecture that provides a unified interface over multiple cloud providers. In its high-level meta-brokering component, we proposed a new scheduling policy that minimizes the overall energy consumption of a federation while maintaining its performance, and in the cloud brokering component, we discussed new heuristics for energy-aware virtual machine management like early destruction, which maintains the energy efficient serving of service calls within a global limit and aggressively shuts down idling or surplus VMs.

For future research directions we plan to investigate methods, e.g., fuzzy and pliant control methods for better determining and handling of the upper and lower thresholds (α^{up} and α^{low}), involve consumption profiles in the calculations and establish prediction methods for power management. Also we plan to refine our model using these approaches to include non-linearity instead of the current simplistic linear approach for determining the service efficiency factor itself.

ACKNOWLEDGMENTS

The research leading to these results has received partial funding from the European Community's Seventh Framework Programme FP7/2007-2013 under grant agreement N^o RI-283481 (SCI-BUS), from the Institute for Computer Science and Control of the Hungarian Academy of Sciences (MTA SZTAKI) under grant agreement N^o 121183 (SZTAKI Cloud) and from the EU COST Action IC0805 "Open European Network for High Performance Computing on Complex Environments".

REFERENCES

1. L. Lefèvre and A.-C. Orgerie, “Towards energy aware reservation infrastructure for large-scale experimental distributed systems,” *Parallel Processing Letters*, vol. 19, no. 03, pp. 419–433, 2009.

2. T. Cioara, I. Anghel, I. Salomie, G. Copil, D. Moldovan, and A. Kipp, “Energy aware dynamic resource consolidation algorithm for virtualized service centers based on reinforcement learning,” in *Parallel and Distributed Computing (ISPDC), 2011 10th International Symposium on*, pp. 163–169, IEEE, 2011.

3. G. Kecskemeti, A. Kertesz, A. Marosi, and P. Kacsuk, “Interoperable resource management for establishing federated clouds,” in *Achieving Federated and Self-Manageable Cloud Infrastructures: Theory and Practice* (M. Villari, I. Brandic, and F. Tusa, eds.), ch. 2, pp. 18–35, IGI Global, 2012.

4. R. Buyya, R. Ranjan, and R. N. Calheiros, “Intercloud: Utility-oriented federation of cloud computing environments for scaling of application services,” in *Algorithms and architectures for parallel processing*, pp. 13–31, Springer, 2010.

5. A. Celesti, F. Tusa, M. Villari, and A. Puliafito, “How to enhance cloud architectures to enable cross-federation,” in *Cloud Computing (CLOUD), 2010 IEEE 3rd International Conference on*, pp. 337–345, IEEE, 2010.

6. P. Marshall, K. Keahey, and T. Freeman, “Elastic site: Using clouds to elastically extend site resources,” in *Proceedings of the 2010 10th IEEE/ACM International Conference on Cluster, Cloud and Grid Computing*, pp. 43–52, IEEE Computer Society, 2010.

7. B. Rochwerger, D. Breitgand, A. Epstein, D. Hadas, I. Loy, K. Nagin, J. Tordsson, C. Ragusa, M. Villari, S. Clayman, *et al.*, "Reservoir-when one cloud is not enough," *Computer*, vol. 44, no. 3, pp. 44–51, 2011.
8. D. Bernstein, E. Ludvigson, K. Sankar, S. Diamond, and M. Morrow, "Blueprint for the intercloud-protocols and formats for cloud computing interoperability," in *Internet and Web Applications and Services, 2009. ICIW'09. Fourth International Conference on*, pp. 328–336, IEEE, 2009.
9. J. L. Lucas-Simarro, R. Moreno-Vozmediano, R. S. Montero, and I. M. Llorente, "Scheduling strategies for optimal service deployment across multiple clouds," *Future Generation Computer Systems*, vol. 29, no. 6, pp. 1431–1441, 2013.
10. M. Cardosa, M. R. Korupolu, and A. Singh, "Shares and utilities based power consolidation in virtualized server environments," in *Integrated Network Management, 2009. IM'09. IFIP/IEEE International Symposium on*, pp. 327–334, IEEE, 2009.
11. Y. C. Lee, C. Wang, A. Y. Zomaya, and B. B. Zhou, "Profit-driven service request scheduling in clouds," in *Proceedings of the 2010 10th IEEE/ACM International Conference on Cluster, Cloud and Grid Computing*, pp. 15–24, IEEE Computer Society, 2010.
12. J. L. Berral, Í. Goiri, R. Nou, F. Julià, J. Guitart, R. Gavaldà, and J. Torres, "Towards energy-aware scheduling in data centers using machine learning," in *Proceedings of the 1st International Conference on energy-Efficient Computing and Networking*, pp. 215–224, ACM, 2010.
13. E. Feller, C. Rohr, D. Margery, and C. Morin, "Energy Management in IaaS Clouds: A Holistic Approach," in *Cloud Computing (CLOUD), 2012 IEEE 5th International Conference on*, pp. 204–212, IEEE, 2012.
14. E. Feller, L. Rilling, C. Morin, R. Lottiaux, and D. Leprince, "Snooze: a scalable, fault-tolerant and distributed consolidation manager for large-scale clusters," in *Proceedings of the 2010 IEEE/ACM Int'l Conference on Green Computing and Communications & Int'l Conference on Cyber, Physical and Social Computing*, pp. 125–132, IEEE Computer Society, 2010.
15. A. Verma, P. Ahuja, and A. Neogi, "pMapper: power and migration cost aware application placement in virtualized systems," in *Middleware 2008*, pp. 243–264, Springer, 2008.
16. V. C. Emeakaroha, M. A. Netto, R. N. Calheiros, I. Brandic, R. Buyya, and C. A. De Rose, "Towards autonomic detection of sla violations in cloud infrastructures," *Future Generation Computer Systems*, vol. 28, no. 7, pp. 1017–1029, 2012.
17. Z. Herczeg, D. Schmidt, Á. Kiss, N. Wehn, and T. Gyimóthy, "Energy simulation of embedded XScale systems with XEEMU," *Journal of Embedded Computing*, vol. 3, no. 3, pp. 209–219, 2009.
18. L. Schubert, K. G. Jeffery, and B. Neidecker-Lutz, *The Future of Cloud Computing: Opportunities for European Cloud Computing Beyond 2010:–expert Group Report*. European Commission, Information Society and Media, 2010.
19. D. Tsirogiannis, S. Harizopoulos, and M. A. Shah, "Analyzing the energy efficiency of a database server," in *Proceedings of the 2010 ACM SIGMOD International Conference on Management of data*, pp. 231–242, ACM, 2010.
20. A. Kertész and P. Kacsuk, "GMBS: A new middleware service for making grids interoperable," *Future Generation Computer Systems*, vol. 26, no. 4, pp. 542–553, 2010.