

# Toward Sustainable Data Centers: a Comprehensive Energy Management Strategy

Jordi Guitart

Received: date / Accepted: date

**Abstract** Data centers are major contributors to the emission of carbon dioxide to the atmosphere, and this contribution is expected to increase in the following years. This has encouraged the development of techniques to reduce the energy consumption and the environmental footprint of data centers. Whereas some of these techniques have succeeded to reduce the energy consumption of the hardware equipment of data centers (including IT, cooling, and power supply systems), we claim that sustainable data centers will be only possible if the problem is faced by means of a holistic approach that includes not only the aforementioned techniques but also intelligent and unifying solutions that enable a synergistic and energy-aware management of data centers. In this paper, we propose a comprehensive strategy to reduce the carbon footprint of data centers that uses the energy as a driver of their management procedures. In addition, we present a holistic management architecture for sustainable data centers that implements the aforementioned strategy, and we propose design guidelines to accomplish each step of the proposed strategy, referring to related achievements and enumerating the main challenges that must be still solved.

**Keywords** Green computing · Energy efficiency · Energy management · Energy measurement · Sustainability · Resource management · Data centers

**Mathematics Subject Classification (2000)** 68U01 · 68M01 · 68-02

---

This work is supported by the Ministry of Science and Technology of Spain under contract TIN2012-34557, by the Generalitat de Catalunya under contract 2014-SGR-1051, and by the European Commission under FP7-SMARTCITIES-2013 contract 608679 (RenewIT) and FP7-ICT-2013-10 contracts 610874 (ASCETiC) and 610456 (EuroServer).

---

Barcelona Supercomputing Center and Universitat Politecnica de Catalunya  
Barcelona, Spain, E-08034  
Tel.: +34-93-4054047  
Fax: +34-93-4017055  
E-mail: jordi.guitart@bsc.es

## 1 Introduction

Energy used by data centers worldwide increased by about 56% from 2005 to 2010, accounting for 1.5% of total energy use in 2010 according to Koomey's study [47]. Greenpeace [5] estimates that data centers energy use can grow up to 1012 billion kWh by 2020, which is a 3x increment regarding their energy consumption in 2007. The cost of this enormous amount of energy has turned into the primary cost driver for data centers. In particular, Belady [17] estimates that the annual amortized energy costs in a data center for a single server exceeded the cost of the server itself in 2008, and that the combined cost of the cooling infrastructure and energy would be 75% of the cost in 2014, while the IT equipment would be only 25%. That is a significant shift from the 20% – 80% ratio of the early 90s. In addition to the energy cost burden, the energy consumption of the data centers contributes also to the climate change by increasing the  $CO_2$  emissions to the atmosphere. In particular, The Climate Group [3] states that worldwide data centers emitted 116 million metric tons of  $CO_2$  ( $MtCO_2$ ) in 2007, slightly more than the entire country of Nigeria, and claims that this figure could increase to 259  $MtCO_2$  by 2020, even considering the recent advances in virtualization, cooling, and power supply that are being introduced in data centers.

For these reasons, there has been recently a great interest in the development of techniques to reduce the energy consumption and the environmental footprint of data centers. They range from proposals to enhance specific aspects of data centers design and operation (for instance, usage of low-power components, energy-aware resource management, and free-cooling solutions) to overall strategies to achieve sustainable data centers, as proposed by Google [43] and HP [14]. This position paper follows the line of those latter approaches and proposes a comprehensive strategy to reduce the carbon footprint of data centers, by using the energy as a driver of the management procedures of the data center. In addition, we present a holistic management architecture for sustainable data centers that implements the aforementioned strategy.

Given the scenario described in the previous paragraphs, we claim that sustainable data centers will be only possible if the problem is faced by means of a holistic approach that includes not only solutions to reduce the energy consumption of the hardware equipment of data centers (including IT, cooling, and power supply systems) but also intelligent and unifying solutions that enable a synergistic and energy-aware management of the hardware equipment and the software infrastructure. In particular, we envision the following steps to come up with a strategy that can enable sustainable data centers:

1. Enable awareness of energy impact and carbon footprint.
2. Increase the energy efficiency of the IT equipment.
3. Increase the energy efficiency of the cooling and power supply subsystems.
4. Increase the use of renewable energy.
5. Exploit opportunities in the energy markets.
6. Customize this energy strategy to the data center reality.

In the following sections, we present the management architecture to implement such strategy and we describe in detail the several steps of our strategy and their interaction. For each step, we propose design guidelines to accomplish that step, we present related achievements and cite relevant works, and we enumerate the main challenges that must be still solved and the areas that need further research.

## 2 Architecture

Figure 1 presents an architecture for managing sustainable data centers that implements the proposed strategy. The management subsystem is led by the *energy-aware manager* and comprises several configurable controllers to manage the various subsystems in the data center. As described in Section 8, the controlling capabilities of the data center determine what management controllers are required. In addition, the *energy-aware manager* is in charge of customizing the management strategy to fulfill the objectives and constraints during operation of the whole data center and configuring and coordinating the management controllers accordingly.

The *monitoring & metrics* and *modeling & prediction* components are responsible for the energy and carbon awareness of the data center. As described in Section 3, the former monitors the operation of the data center and calculates relevant metrics that provide an indication of its sustainability. The latter uses this monitored information to build models that capture the system behavior and uses them to forecast that behavior in the future.

The *IT equipment controller* manages the IT equipment as described in Section 4, deciding about the energetic status and the operating frequency of the physical hosts in the data center, as well as the configuration parameters of the networking elements (i.e. routers, switches, ...).

The *workload shaping & task scheduling* component is responsible for the execution of the IT workload in the data center. As described in Sections 4 and 6, it decides about what tasks are accepted for execution, when such tasks will be executed, and their placement in the physical hosts of the data center.

The *remote controller* selects the most appropriate data center where to deploy a workload when remote execution in the ecosystem is expected to provide a better outcome regarding energy efficiency or carbon footprint, as described in Sections 4.3 and 6.

The *cooling controller* configures the cooling subsystem over different operating points to maintain a target temperature while reducing the power consumption, as described in Section 5.

The *power supply controller* manages the power supply subsystem as described in Section 6, deciding for each time interval the amount of energy to be generated on-site, the amount of energy to be taken from the power grid, and the amount of energy to be stored for later use, and also how to distribute the available energy among the rest of subsystems of the data center.

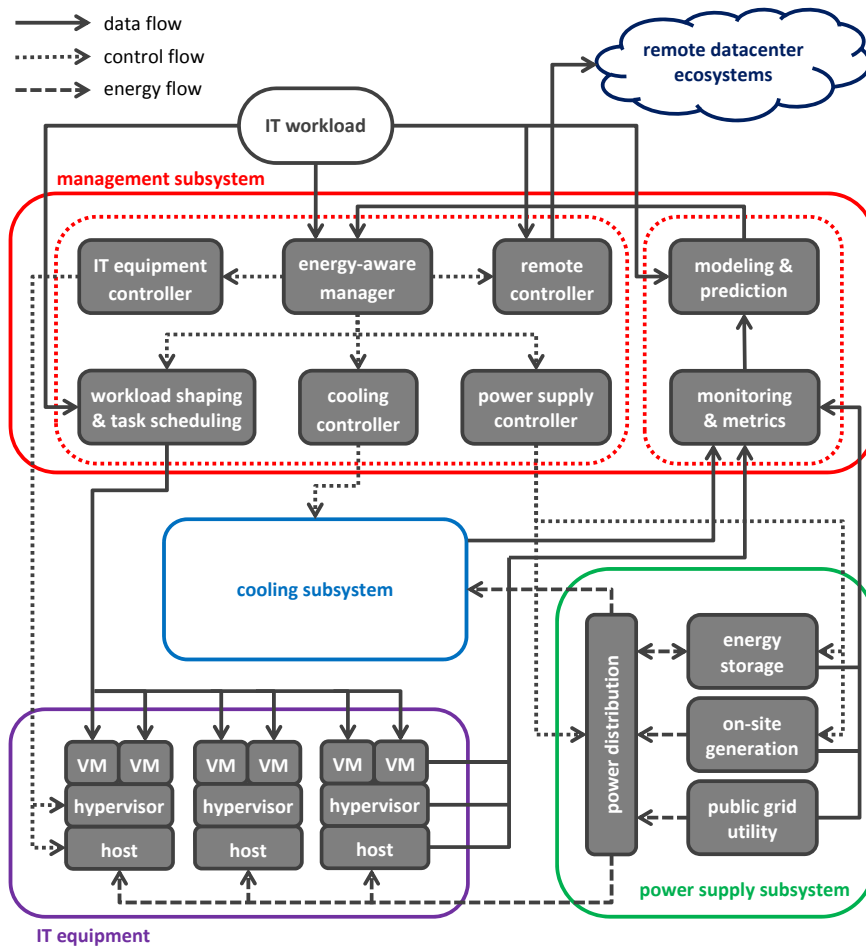


Fig. 1 Architecture for the management of sustainable data centers

### 3 Energy and Carbon Awareness

Energy and carbon awareness refers to the ability of data centers to measure their energy consumption and carbon footprint. It is a fundamental capability because “consumers and businesses can’t manage what they can’t measure”, as stated by The Climate Group [3].

#### 3.1 Metrics

First, we must define what to measure. Usually, data centers account for metrics that provide an indication of their sustainability degree, and to this end, they measure parameters such as greenhouse gas emissions, power consump-

tion, temperatures, humidity, etc. [72]. For instance, a common metric is the Power Usage Effectiveness (PUE) [6], which is the ratio of the total amount of energy used by a data center (including the power supply and cooling subsystems in addition to the IT equipment) to the energy delivered to the IT equipment. A similar rationale is used to calculate the Carbon Usage Effectiveness (CUE) [4], but using the total  $CO_2$  emissions caused by the total amount of energy used by a data center in the numerator.

PUE is useful as an indicator on how efficiently the data center uses the energy in the long-term (in fact, it is generally reported from annual information aggregates), but it does not capture well the dynamic nature of the energy usage in data centers. PUE Scalability [6] has been suggested to show how well a data center total energy consumption scales with dynamic changes in its IT equipment loads. Furthermore, PUE is not intended to assess how energy is being used to generate efficiently the useful work of the data center. Data Center energy Productivity (DCeP) [2] has been suggested to quantify the useful work output that a data center provides in relation to the amount of energy expended to produce it. This metric allows the continuous monitoring of the work product as a function of energy consumed during data center operation, but it can be sometimes difficult to implement, since the useful work in a data center can take many forms. This hardens also the comparison of the productivity of different data centers by using this metric. In addition, some authors claim that metrics such as DCeP can favor algorithms that achieve high sustained performance, which are also power hungry algorithms, and propose an alternative metric (so-called FTTSE), which prioritizes the total energy reduction and the minimization of time to solution [16].

There is an ongoing effort trying to harmonize the metrics to be used to measure data centers energy and ecological efficiency. Whereas some consensus has been achieved regarding the metrics [8], their adoption as drivers of the energy management strategies of data centers is still in the initial stages. An important challenge for the standardization of this kind of metrics comes up because of the different actors involved in the data center operation (i.e. owners, administrators, users), who have different requirements regarding the information they want to obtain and about the granularity of such information (per function, per service, per host, per rack, or the whole data center) [72].

### 3.2 Measurements vs. models

Second, we must define how to measure. The total amount of energy used by a data center can be monitored with Metered Power Distributions Units (PDU), which provide the ability to measure power usage of a system at the socket. Modern IT equipment also allows to measure power usage at the server by means of the Intelligent Platform Management Interface (IPMI) [39]. Some servers can even report the power consumed by some individual components (CPU, RAM, etc.) if the server comes equipped with the appropriate sensors. Alternatively, Intel Sandy Bridge processors include an onboard energy and

power metering capability for processor packages and DRAM so-called Running Average Power Limit (RAPL), which estimates energy usage by using hardware performance counters and a software power model [40].

The amount of carbon emissions must be determined for the actual mix of energy delivered to the data center, including on-site and grid-based energy sources, by using the carbon emission factor ( $kgCO_2/kWh$ ) associated to each energy source [7]. Whereas it is easy to identify energy sources when energy is generated on-site, when the energy consumed comes from the grid the data center must rely on the information provided by the corresponding energy supplier. Generally, they provide aggregate values for a given region.

The information about energy consumption and carbon footprint should be integrated in the existing monitoring platforms of data centers (e.g., Ganglia [1]), allowing in that way a seamless access to all the information that is relevant to understand the data center behavior.

Power metering mechanisms described above provide real measures of the energy consumption of individual hosts. However, they cannot account for the energy consumption of individual virtual machines (VMs) or services running on each host, which is required for instance to implement effective power-aware scheduling policies. This can be accomplished by using power models, which correlate the resource usage and the power consumption of individual VMs based on monitored resource usage information [42, 75]. Models also allow to estimate the current power consumption of hosts and VMs when direct measurement is not feasible (due to lack of scalability in large clusters, performance overhead of measurements, and cost of devices [59]) and to forecast the power consumption in the future (see Section 3.3).

Models are generated offline by collecting power measurements and resource usage indicators periodically while the system runs a special training workload that includes micro-benchmarks that selectively stress each resource at varying levels of utilization. Given the heterogeneity of applications running in current data centers [60], the training workload must be generic, since including only a given kind of applications [41] will result in inaccurate estimations for the rest, and able to capture the essentials of the power behavior of the modeled host. To this end, power models must consider the impact on energy consumption of several resources, including not only the processor (as models for traditional HPC applications did in the past), but also the memory, the disk, and the network. Those subsystems apart from the processor have been reported to make up 40%-60% of the total power consumption depending on the workload [21]. Similarly, power models must also support heterogeneous virtualization hypervisors, as they have been reported to impact differently the performance and energy consumption of the applications [68].

Only appropriate indicators to account for the usage of each resource must be selected. The indicator must clearly represent the usage intensity of the resource (i.e. if the usage increases, the value of the indicator should increase correspondingly) and it must be unequivocally correlated with the energy consumption incurred by that resource. For instance, the utilization is not the best indicator of the processor usage regarding its correlation with energy consump-

tion, because applications with the same utilization can have different processor energy consumption depending on what instructions they are executing [42]. According to this, the utilization cannot be the unique determinant of the processor power consumption, though it could be used to further refine a model that considers, for instance, the number of operations executed.

The collected resource usage indicators are fitted to the power measurements by using machine learning techniques [13], which include mainly traditional auto-regressive methods (e.g. linear regression) [59] but also more novel methods such as artificial neural networks [27]. Those are very powerful techniques to capture data correlations, but they might require driving the mining process to achieve better accuracy and reduce the computational complexity. According to this, the modeling methodology should systematically filter the relevant indicators and include derivatives that capture better the existing relations. Models have assumed often a linear relationship between the power consumption and the resource usage. Linear models provided a reasonable accuracy with low computational complexity in traditional platforms, but they are weak in modern multicore platforms [58], where non-linear models provide better accuracy, although they incur also higher computational complexity.

Power models must support heterogeneous platforms (e.g. low-power such as Intel Atom vs. high-performance such as Intel Xeon), which are the norm in current data centers [60], by incorporating a specific model for each type of platform, as they can present very different energy consumption patterns [76]. One could think that this would lead to different models for each application-platform pair, which would entail a noticeable effort to generate those models. However, we can reduce the complexity incurred due to platform heterogeneity by deriving approximate models of new platforms using the model of existing platforms with the same architecture, thus avoiding the benchmarking of the new platforms. Similarly, we can reduce the complexity incurred due to application heterogeneity by using the platform model generated with the training workload and, if needed, refining it online for each application (or group of applications sharing common resource usage patterns) using captured data.

Finally, power models must consider the impact of co-locating several VMs in the same resources. Sharing resources generates an interference between the VMs, which induces some overhead [48]. This overhead has an obvious impact on the performance of the VMs, but also on the energy consumption.

### 3.3 Forecasts

As power models correlate the resource usage and the power consumption, if one can predict the resource usage of a host or a VM in the future, those models could be then used to forecast the energy consumption of that host or VM [52]. Resource usage predictions can be obtained using time series analysis techniques [22]. Note that this approach supports predictions of the energy consumption at different periods ahead, from short-term to long-term, provided that the model is supplied with the corresponding resource usage

predictions. Contrariwise, many other model-based predictors only allow for very short-term forecasts, as they feed the model with current measurements of resource usage instead of predictions.

It should be also helpful if power models can be used to forecast the energy consumption for potential actions. For instance, predicting the energy consumption of the data center upon a VM deployment, VM migration, or VM cancellation before the actual action takes place [64]. The system scheduler could use these forecasts to guide its resource management decisions. Predicting the future resource usage is harder with this kind of forecasts because the effect of the potential action must be incorporated to the predicted resource usage in the data center. The complexity increases considerably when forecasting the energy impact of a combination of potential actions.

The amount of carbon emissions can be predicted from the expected energy consumption and the carbon emission factor associated to each energy source, using the same rationale described in Section 3.2. This requires predicting the expected workload, but also the expected energy mix to be delivered to the data center, as the energy mix will vary over time when using intermittent energy sources such as solar and wind energy.

## 4 Energy Efficiency of the IT Equipment

Energy efficiency describes the ratio between the produced outcomes and the energy used. Something is more energy efficient if it delivers more services using the same amount of energy, or the same services using less energy. IT equipment can be more energy efficient by means of energy-driven management strategies that reduce its energy consumption and carbon footprint.

### 4.1 Energy-related actuators

Energy-driven management strategies require actuators that can modify the behavior of the IT equipment regarding their energy footprint. For instance, data centers can take profit of virtualization technology to consolidate workloads from different customers into a smaller number of physical hosts. This allows turning off (or suspending) unused hosts (see ACPI C-states [66]) and therefore reducing energy consumption. Data centers can apply also elasticity of VMs regarding the use of resources to control the number of VMs in each physical host. Moreover, they can use also Dynamic Voltage and Frequency Scaling (DVFS) to adjust the voltage and/or frequency of the processors on the fly (see ACPI P-states [66]) aiming to reduce power consumption (dynamic power consumption changes linearly with frequency and in a quadratic manner with voltage) or to reduce the amount of heat dissipated. Note that running an application at lower frequency increases its execution time, which could lead to higher energy consumption depending on the relationship between the reduction of power consumption and the increase in the execution time for



a given frequency reduction [30]. Smart management solutions that trade off these variables to reduce the overall energy consumption are needed here.

## 4.2 Energy-driven management

Energy-awareness and energy-related actuators allow putting energy-driven management strategies for the IT equipment into operation, as they can be the basis for management algorithms that optimize their operation according to energy-based objectives (e.g. maximize the energy productivity of the data center). On one side, management algorithms can decide the placement of VMs in the physical hosts of the data center, as well as the energetic status of those hosts (on(P-state)/suspended(C-state)/off) [56]. Each VM must be allocated with enough resources of each type to fulfill its demand. This complicates the allocation problem as several dimensions have to be considered. The impact of co-locating several VMs in the same physical host must be also considered as this can degrade their performance due to interferences accessing the resources [45]. As the energy consumption of the data center is correlated with the number of active physical hosts, a common strategy to improve energy efficiency is to consolidate all the VMs running in the data center in the minimum number of hosts (the necessary to fulfill the desired performance) while suspending the rest of hosts [73]. These idle hosts could be turned on again if they were needed when a peak load occurs. However, management algorithms can achieve better results if, instead of the number of active hosts, they consider the dynamic power consumption of the physical hosts (which depends on their load) that would result from each configuration [19]. For that, they will use the power models described in the previous section to forecast the energy consumption of VMs and physical hosts according to their predicted resource usage and the status of the infrastructure. In addition, DVFS and elasticity of VMs can be used to adapt the computing capacity of the hosts and the VMs to the intensity of the workloads they have to execute [65].

On the other side, management algorithms can decide the configuration parameters of the networking elements (i.e. routers, switches, ...), as well as their energetic status, to optimize the energy needed to route messages within the data center. This can have a noticeable impact, especially in big data centers, which organize in complex hierarchical switching topologies. According to this, common strategies to improve energy efficiency consist of adapting dynamically the operating data transmission rate of communication links [20], finding routes that minimize the total number of switches and turning off unused ones [77], and exploiting low energy modes in physical hosts and networking devices together in a single management algorithm [11].

The outcome of those optimization algorithms can be significant when applied in heterogeneous environments [76]. On one side, data centers can have heterogeneous platforms (e.g. low-power vs. high-performance), which have very different energy consumption profiles. Low-power platforms offer reduced power consumption in the idle state and mild increments with resource usage,

but they provide also reduced performance. For this reason, heterogeneous platforms present interesting trade-offs for the management algorithms, i.e. run briefly in a high-performance high-power platform and suspend this platform for longer vs. run for a longer time in a low-performance low-power platform and only suspend the platform briefly.

On the other side, this platform heterogeneity fits very well with the heterogeneity of the applications running in today's data centers, which include interactive services, high performance computing jobs and big data applications. As these applications present very different resource usage patterns (and consequently, different energy consumption profiles), they allow for several placement options in the data center according to the required performance by each application and the energy consumption incurred to provide that performance [26]. Again, interesting trade-offs appear here, i.e. place all the applications of a given type in the platform where their energy productivity is better but foster their interference because they are using the same kind of resources vs. co-locate applications of different types in a given platform to minimize interference even if their energy productivity is not optimal.

### 4.3 Data center ecosystems

The aforementioned management algorithms can contribute to the energy efficiency of data centers when considered in isolation. However, current data centers normally work within ecosystems where they interact with other data centers. For example, Cloud data centers can be organized in federations where a given data center can lease capacity from other data centers to meet peaks in demand or to outsource less critical workloads. Service owners can also use multi-cloud scenarios to deploy their service across multiple data centers [31].

In such scenarios, energy-driven management strategies can also have a noticeable impact. In this case, management algorithms will aim to optimize the operation of data centers according to energy-based objectives by selecting the most appropriate data center where to deploy every workload (a.k.a. Geographical Load Balancing (GLB)). For instance, GLB allows exploiting data centers located in different time zones and energy price variability [50] by enabling a follow-the-moon strategy [38], which deploys workloads in data centers by following the night-time, when data centers may be in an off-peak demand interval, the energy is cheaper, and lower outside temperature opens the possibility to use free cooling. GLB also allows exploiting data centers powered by renewable sources to reduce carbon emissions, as we will discuss in Section 6. In addition, GLB can be integrated with energy buffering management in order to shave peak power draw from the grid [10].

Whereas data center ecosystems offer new opportunities for energy-driven management, they also encompass new challenges that must be considered [44, 56, 10], such as the distant geographic locations of data centers, which have an impact in the migration of VMs (increasing the cost significantly (and the consumed energy) and frequently causing service level degradations for the

affected customers) and in the data exchange among VMs located in different data centers (increasing the communication latency among them); the independent administrative domains involved in the ecosystem, which have frequently conflicting goals and do not generally disclose information about their energy consumption and energy mix, thus increasing the need for third parties to independently assess energy data of data centers and share this information within the ecosystem; and the importance of the prediction accuracy of the input data (e.g. workload, energy price, renewable energy), which depends on the predictability of each data source and the prediction window length, and can be a downgrading factor on the efficiency of the management algorithms.

#### 4.4 Multi-objective optimization

Energy-related objectives are important for data centers administrators, but they are also interested in other aspects, such as availability, reliability, profit, and performance. For this reason, the management algorithms should optimize the operation of data centers according to multiple objectives. This complicates the optimization problem, especially when some objectives are mutually exclusive and both cannot be maximized at the same time. For instance, consolidating VMs onto a small number of physical hosts and turning off idle hosts is an effective way to reduce energy consumption, but this can cause heat imbalances and create hot spots, which may impact cooling costs and degrade host life and performance [74].

Multi-objective optimization problems can be solved via scalarization, that is, converting the original problem with multiple objectives into a single-objective optimization problem, for instance, by considering their impact on the data center profit [34]. The problem can be also solved as a single-objective problem when all but one of the objectives have a target value, as a constraint can be placed on those objectives. When a target value can be identified for all the objectives, the problem can be solved by means of goal-oriented adaptation, which is able to learn from monitored data the impact of repair actions on the value of the goals and apply the most convenient action when any of them deviates from its target value [70]. Another possibility for solving true multi-objective problems consists of computing all or a representative set of Pareto optimal solutions (those in which it is impossible to make any objective better off without making at least another one worse off), which are usually derived by means of evolutionary algorithms [74,32].

A similar issue arises when we consider not only the objectives of the data center, but also the objectives of its clients. Hopefully, a data center should fulfill the goals of its clients in addition to its own goals, but there could be incompatibilities between the goals of the data center and its clients or among the goals of the clients themselves. For instance, this occurs when data centers offer distinct quality of service categories (gold, silver, etc.) to their customers [71]. Again, trade-off solutions are required.

#### 4.5 Optimization problem solving

All the described optimization algorithms, whether they are single- or multi-objective, must solve a packing problem [56]. A way to obtain optimal solutions consists of formulating the problem in terms of some mathematical program, for instance by using Integer Linear Programming, and using an existing solver to calculate the solution [37]. However, packing problems are known to be NP-hard. This means that brute-force methods to find the optimal solution become infeasible when the scenario starts to grow. According to this, construction heuristics and local search methods are required to allow good solutions (although maybe not optimal) to be decided in real time.

Given the similarity with the well-known bin-packing problem, simple bin-packing heuristics (such as First Fit, Best Fit, First Fit Decreasing ...) have been adapted to the VM allocation problem [63]. Nevertheless, the VM allocation problem is more complicated than bin packing [57]. This obliges to analyze to what extent the bin-packing heuristics can be applied and what adaptations are required. The main differences of the VM allocation problem refer to its multi-dimensionality regarding the type of resources that determine the host capacities and VM sizes, the costs incurred by VM migrations, the performance degradation when a host is overloaded (or close to), the heterogeneity of the hosts, the impact of the load of the host in its energy consumption, the finite number of hosts in the data center which constraints the number of VMs that can be allocated, and the volatility of the system regarding the number of VMs to allocate and the number of available hosts [57].

### 5 Energy Efficiency of the Cooling and Power Supply Subsystems

The Climate Group states that in 2002 only about half of the energy used by data centers powered the IT equipment, while the rest was needed to run backup, uninterruptible power supplies and cooling systems [3]. Despite the recent advances in cooling and power supply systems, The Climate Group estimates that in 2020 the energy footprint of these systems will still account for a 39% of the total footprint of data centers (about 101  $MtCO_2$ ).

According to this, sustainable data centers must also aim to increase the energy efficiency of the power and cooling systems, or in other words, aim to reduce the PUE. This goal requires the usage of novel techniques to reduce the energy consumption of these subsystems (such as free cooling, which consists of using ambient air or water to cool data center space and equipment), but also including these systems in a comprehensive energy-aware management strategy that optimizes the IT equipment, as described in Section 4, while reducing the energy consumption of the cooling and energy supply subsystems.

On one side, this can be accomplished by means of smarter algorithms to manage the cooling subsystem, for instance, by adapting proactively the amount of cooling to the expected workload [51]. If we envision a low utilization in the data center in a near future, we can already throttle down the cooling

subsystem, because less hosts will be required and we will be able to suspend some of them (or use DVFS), thus reducing the amount of dissipated heat and the cooling needs. This differs from traditional reactive management approaches, which are triggered when the temperature surpasses predefined thresholds, and can result in delayed response to temperature changes. In addition, the optimal threshold range to avoid component damage while at the same time avoiding energy waste in unnecessary cooling is very hard to determine [51].

On the other side, the algorithms managing the IT workload can profile the tasks according to their impact on the temperature and consider the resulting temperature of the hosts as a criterion to decide the allocations. Again, a lower resulting temperature allows the cooling subsystem to be throttled down, further decreasing power consumption. For instance, the algorithms can apply this idea to decide when tasks must be executed, and run *cold* tasks (i.e. those with low impact on the processor temperature) after *hot* tasks to lower the resulting temperature (the order of execution of the tasks has an impact on the temperature) [25]. Similarly, they can also decide where tasks must be executed, and place *hot* tasks in the hosts of the data center that are easier to cool [24]. Naturally, both decisions on when and in which host to execute a task can be taken together to minimize the resulting temperature [15].

The described approaches to optimize the cooling subsystem must be tightly integrated with the optimization of the IT equipment. Isolated strategies to reduce the energy consumption of the IT equipment tend to consolidate the workload on fewer hosts and turn off the rest, which can cause heat imbalances and create hot spots, thus increasing the cooling needs. Isolated strategies to reduce the energy consumption of the cooling subsystem tend to spread the workload over all the hosts, thus increasing the idle power consumption of the IT equipment, as more hosts are turned on. A joint optimization that balances the number of active hosts to trade off the power needed by the cooling subsystem and the idle power needed by the IT equipment can reduce the total power consumption of the data center [12]. As the cooling subsystem could require some time to adapt changing conditions, joint approaches to optimize the cooling subsystem and the IT equipment can configure the cooling subsystem over different operating points to deal with long-term fluctuations and use thermal-aware workload placement to deal with short-term ones [69].

There is also room for the optimization of the power supply subsystem, especially when the data center can generate energy on-site (see Section 6).

## 6 Ecological Efficiency of the IT Equipment

Enhancements in energy efficiency, as described in previous sections, will contribute to the reduction of the energy consumed by a data center to accomplish a given amount of work. Although this is very important to reduce the carbon footprint of data centers, it is not enough. As the workload demand over data centers is expected to increase, their energy consumption will increase also, even if energy efficiency measures are applied. In fact, some argue that energy

efficiency could encourage further increments in the workload demand due to lower energy costs (which is known as Jevons' Paradox [18]). According to this, sustainable data centers require not only to reduce their energy consumption but also to use as much as possible energy sources that do not contribute to carbon emissions, that is, renewable energy sources.

The use of renewable energy sources introduces new challenges into the management of data centers. One of them appears when using intermittent energy sources, that is, sources that are not always available, such as solar and wind energy. In these scenarios, the management algorithms must adapt the IT workload to the energy availability (i.e. workload shaping). On one side, workload shaping can be accomplished by using techniques such as DVFS or server power state tuning to adjust the power demand to the time-varying renewable power budget [33]. On the other side, workload shaping can be also based on task scheduling, for instance by postponing tasks when not enough renewable energy is available to execute them, or by bringing forward tasks if there is more renewable energy than needed [35].

Data center ecosystems offer more opportunities to increase the use of renewable energy (and reduce the carbon footprint), as they comprise data centers located in different geographical locations, and even in different time zones, which results in different local weather conditions at some particular time. This allows enabling a follow-the-renewables strategy [54], where workloads are deployed in data centers according to the availability of renewable energy in their geographical locations [78], or to the expected carbon emissions, which are calculated from the data center energy mix [23].

Co-location (data centers draw renewable energy directly from an existing nearby plant) and self-generation (data centers generate their own renewable energy) have been reported as the preferred methods (versus grid-centric approaches) for data centers to exploit renewable energy [36], as they allow decreasing the transmission and conversion losses (because the power is generated close to where it is consumed and DC power can be directly supplied to the IT equipment instead of AC power) and avoid the grid-transmission charges imposed by power utilities. Self-generation allows also the utilization of the power plant waste heat to generate cooling for the data center [29]. These features make self-generation an appealing option to obtain cheaper power than the grid once the important initial capital cost of installing the needed infrastructure has been amortized.

Self-generation offers also additional management opportunities because the management algorithms can control the power supply system. In particular, they can use a load following approach that adapts the amount of energy generated to the expected workload [46]. As energy generation systems could require some time to adapt the amount of energy they generate, load following is not adequate to fit the energy demands of the IT workload when this fluctuates in short time intervals. In these situations, load following can be combined with workload shaping. The former defines the amount of energy to be generated during the next (coarse) time period according to the estimated workload demand, which is calculated from past information, and the latter

works within that time period to adapt the IT workload to fit its deviations to the available energy [53].

## 7 Exploit Opportunities in the Energy Markets

Even after applying the previous steps, data centers can still have some carbon footprint. This can be further reduced by exploiting the opportunities that are available in the current energy markets, provided that the required expenses to reduce its carbon footprint by using these opportunities fit in the data center business strategy. For instance, Renewable Energy Certificates (RECs) (a.k.a. Renewable Energy Credits) [67] represent proof that one MWh of electricity was generated from a renewable energy source. RECs are tradable commodities in energy markets and they allow their owner to claim that the corresponding portion of its overall energy consumption comes from renewable sources.

A data center can reduce its carbon footprint by purchasing RECs [28, 61]. Similarly, a data center that generates renewable energy on-site can offer its spare green energy in the market by issuing RECs, allowing other data centers to offset their carbon impact [61]. Note that, in addition to the benefit of a global reduction of the carbon impact, these strategies can offer also business opportunities to the data centers to increase their profit.

## 8 Customization of the Energy Strategy

Energy management strategies should be flexible in the sense that they must be easily customizable to each data center reality. On one side, they should adapt to the energy management mechanisms that the data center has. For instance, whether it can control the energy supply (because some energy is generated on-site), the operation point of the cooling subsystem, or the energetic status of the IT equipment. This requires a management architecture with independent controllers for each subsystem that can be plugged in/out as necessary and can work together to manage the data center.

On the other side, they should allow the definition of the objectives and constraints during operation of the whole data center (or any of the applications running there) and they should be able to fulfill them. For instance, ensuring a minimum performance (e.g. a minimum number of requests per second for a Web server), a maximum amount of energy that can be consumed (e.g. the amount of available green energy is limited), a maximum amount of carbon that can be emitted (e.g. renewable energy credits are limited), a maximum amount of power that can be provided (e.g. there is a power budget), or a maximum value for the operational costs (e.g. there is an economic budget).

The data center administrators define the high-level objectives and constraints according to the business interests of the provider. Those high-level objectives are commonly referred as Business Level Objectives (BLO). The Green Grid Data Center Maturity Model [9], which suggests best practices for

energy efficient data centers, can be also used as reference to define constraints [70]. The energy-aware manager configures the management controllers in order to fulfill the BLOs. Some controllers support their configuration directly by means of BLOs and can adapt their behavior to fulfill them under changing conditions [55]. Other controllers are configured by means of low-level objectives related to its own domain or adjustable low-level policies [62]. In this case, the energy-aware manager must be able to translate the BLOs to low-level terms that the management controllers can understand [49].

Management controllers operate independently and adhere to their specified local objectives by means of a self-adaptation loop. However, the impact of their decisions goes beyond their own domain, and can affect the rest of controllers. In fact, these decisions could be conflicting and their combined effect could fail to fulfill the data center BLOs. According to this, the energy-aware manager must implement a self-adaptive loop that monitors the system status, evaluates the BLOs fulfillment under those conditions, and aligns the configuration of the management controllers to ensure the fulfillment of the BLOs [62]. The effect of different configurations over the BLOs can be learned from available data using machine learning techniques [70].

## 9 Conclusions

In this paper, we have presented a comprehensive energy management strategy for data centers that aim for sustainability and a holistic architecture that implements such strategy. We claim that those data centers must first be aware of their energy impact and carbon footprint by measuring appropriate metrics and being able to forecast their value by using models. They must also increase the energy efficiency of their IT equipment by optimizing the placement of tasks in the data center physical hosts, as well as the energetic status of those hosts, while exploiting the opportunities given by heterogeneous platforms and applications and by data center ecosystems. They must increase the energy efficiency of the cooling and power supply subsystems by integrating them in the energy strategy applied with the IT equipment. As carbon footprint is directly impacted by the used energy sources, data centers must increase also the use of renewable energy by generating it on-site and optimizing its management depending on the available renewable energy, or by purchasing carbon credits to avoid accounting their emissions as real carbon impact. Finally, all the previous steps must fit in the frame of each data center and have to be customized accordingly.

## References

1. Ganglia Monitoring System. <http://ganglia.sourceforge.net/>
2. A Framework for Data Center Energy Productivity. White Paper n.13, The Green Grid (2008)
3. SMART 2020: Enabling the Low Carbon Economy in the Information Age. Report, The Climate Group (2008)



4. Carbon Usage Effectiveness (CUE): A Green Grid Data Center Sustainability Metric. White Paper n.32, The Green Grid (2010)
5. Make IT Green: Cloud Computing and its Contribution to Climate Change. Report, Greenpeace International (2010)
6. PUE: A Comprehensive Examination of the Metric. White Paper n.49, The Green Grid (2012)
7. Energy and Carbon Conversion Factors. Tech. Rep. CTL 153, The Carbon Trust (2013)
8. Harmonizing Global Metrics for Data Center Energy Efficiency. Statement, Global Taskforce (2014)
9. Data Center Maturity Model. White Paper n.56, The Green Grid (2015)
10. Abbasi, Z., Pore, M., Gupta, S.: Impact of Workload and Renewable Prediction on the Value of Geographical Workload Management. In: Energy-Efficient Data Centers: 2nd Intl. Workshop, E2DC 2013. Revised Selected Papers, *Lecture Notes in Computer Science*, vol. 8343, pp. 1–15. Springer (2014). DOI 10.1007/978-3-642-55149-9\_1
11. Addis, B., Ardagna, D., Capone, A., Carello, G.: Energy-aware Joint Management of Networks and Cloud Infrastructures. *Computer Networks* **70**, 75–95 (2014). DOI 10.1016/j.comnet.2014.04.011
12. Ahmad, F., Vijaykumar, T.N.: Joint Optimization of Idle and Cooling Power in Data Centers While Maintaining Response Time. In: Proc. 15th Intl. Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS'10), Pittsburgh, USA, pp. 243–256. ACM (2010). DOI 10.1145/1736020.1736048
13. Alpaydin, E.: Introduction to Machine Learning, Third Edition. The MIT Press (2014)
14. Arlitt, M., Bash, C., Blagodurov, S., Chen, Y., Christian, T., Gmach, D., Hyser, C., Kumari, N., Liu, Z., Marwah, M., McCreynolds, A., Patel, R., Shah, A., Wang, Z., Zhou, R.: Towards the Design and Operation of Net-Zero Energy Data Centers. In: Proc. 13th IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITHERM'12), San Diego, USA, pp. 552–561. IEEE (2012). DOI 10.1109/ITHERM.2012.6231479
15. Banerjee, A., Mukherjee, T., Varsamopoulos, G., Gupta, S.K.: Integrating Cooling Awareness with Thermal Aware Workload Placement for HPC Data Centers. *Sustainable Computing: Informatics and Systems* **1**(2), 134–150 (2011). DOI 10.1016/j.suscom.2011.02.003
16. Bekas, C., Curioni, A.: A New Energy Aware Performance Metric. *Computer Science - Research and Development* **25**(3-4), 187–195 (2010). DOI 10.1007/s00450-010-0119-z
17. Belady, C.L.: In the Data Center, Power and Cooling Costs More Than the IT Equipment It Supports. *Electronics Cooling* (2007). <http://www.electronics-cooling.com/2007/02/in-the-data-center-power-and-cooling-costs-more-than-the-it-equipment-it-supports>
18. Belady, C.L.: Does Efficiency in the Data Center Give Us What We Need? *Mission Critical Magazine* (2008). <http://www.missioncriticalmagazine.com/articles/does-efficiency-in-the-data-center-give-us-what-we-need>
19. Beloglazov, A., Abawajy, J., Buyya, R.: Energy-aware Resource Allocation Heuristics for Efficient Management of Data Centers for Cloud Computing. *Future Generation Computer Systems* **28**(5), 755–768 (2012). DOI 10.1016/j.future.2011.04.017
20. Bilal, K., Khan, S., Madani, S., Hayat, K., Khan, M., Min-Allah, N., Kolodziej, J., Wang, L., Zeadally, S., Chen, D.: A Survey on Green Communications using Adaptive Link Rate. *Cluster Computing* **16**(3), 575–589 (2013). DOI 10.1007/s10586-012-0225-8
21. Bircher, W.L., John, L.K.: Complete System Power Estimation Using Processor Performance Events. *IEEE Transactions on Computers* **61**(4), 563–577 (2012). DOI 10.1109/TC.2011.47
22. Brockwell, P., Davis, R.: Introduction to Time Series and Forecasting, 2nd Edition. Springer-Verlag New York (2002). DOI 10.1007/b97391
23. Cappiello, C., Melia, P., Pernici, B., Plebani, P., Vitali, M.: Sustainable Choices for Cloud Applications: a Focus on CO2 Emissions. In: Proc. 2nd Intl. Conference on ICT for Sustainability (ICT4S'14), Stockholm, Sweden, pp. 352–358. Atlantis Press (2014). DOI 10.2991/ict4s-14.2014.43
24. Chen, Y., Gmach, D., Hyser, C., Wang, Z., Bash, C., Hoover, C., Singhal, S.: Integrated Management of Application Performance, Power and Cooling in Data Centers. In:

- Proc. 2010 IEEE Network Operations and Management Symposium (NOMS'10), Osaka, Japan, pp. 615–622. IEEE (2010). DOI 10.1109/NOMS.2010.5488433
25. Chrobak, M., Durr, C., Hurand, M., Robert, J.: Algorithms for Temperature-aware Task Scheduling in Microprocessor Systems. *Sustainable Computing: Informatics and Systems* **1**(3), 241–247 (2011). DOI 10.1016/j.suscom.2011.05.009
  26. Chun, B.G., Iannaccone, G., Iannaccone, G., Katz, R., Lee, G., Niccolini, L.: An Energy Case for Hybrid Datacenters. *SIGOPS Operating Systems Review* **44**(1), 76–80 (2010). DOI 10.1145/1740390.1740408
  27. Cupertino, L., Da Costa, G., Pierson, J.M.: Towards a Generic Power Estimator. *Computer Science – Research and Development* **30**(2), 145–153 (2015). DOI 10.1007/s00450-014-0264-x
  28. Deng, N., Stewart, C., Gmach, D., Arlitt, M., Kelley, J.: Adaptive Green Hosting. In: Proc. 9th Intl. Conference on Autonomic Computing (ICAC'12), San Jose, USA, pp. 135–144. ACM (2012). DOI 10.1145/2371536.2371561
  29. Erden, H.S., Khalifa, H.E.: Energy and Environmental Assessment of On-Site Power and Cooling for Data Centers. *HVAC & R Research* **18**(1–2), 51–66 (2012). DOI 10.1080/10789669.2011.585422
  30. Etinski, M., Corbalan, J., Labarta, J., Valero, M.: Understanding the Future of Energy-Performance Trade-off via DVFS in HPC Environments. *Journal of Parallel and Distributed Computing* **72**(4), 579–590 (2012). DOI 10.1016/j.jpdc.2012.01.006
  31. Ferrer, A.J., Hernandez, F., Tordsson, J., Elmroth, E., Ali-Eldin, A., Zsigri, C., Sirvent, R., Guitart, J., Badia, R.M., Djemame, K., Ziegler, W., Dimitrakos, T., Nair, S.K., Kousiouris, G., Konstanteli, K., Varvarigou, T., Hudzia, B., Kipp, A., Wesner, S., Corrales, M., Forg, N., Sharif, T., Sheridan, C.: OPTIMIS: A Holistic Approach to Cloud Service Provisioning. *Future Generation Computer Systems* **28**(1), 66–77 (2012). DOI 10.1016/j.future.2011.05.022
  32. Gao, Y., Guan, H., Qi, Z., Hou, Y., Liu, L.: A Multi-objective Ant Colony System Algorithm for Virtual Machine Placement in Cloud Computing. *Journal of Computer and System Sciences* **79**(8), 1230–1242 (2013). DOI 10.1016/j.jcss.2013.02.004
  33. Gmach, D., Rolia, J., Bash, C., Chen, Y., Christian, T., Shah, A., Sharma, R., Wang, Z.: Capacity Planning and Power Management to Exploit Sustainable Energy. In: Proc. 2010 Intl. Conference on Network and Service Management (CNSM'10), Niagara Falls, Canada, pp. 96–103. IEEE (2010). DOI 10.1109/CNSM.2010.5691329
  34. Goiri, I., Berral, J.L., Fitó, J.O., Julià, F., Nou, R., Guitart, J., Gavaldà, R., Torres, J.: Energy-efficient and Multifaceted Resource Management for Profit-driven Virtualized Data Centers. *Future Generation Computer Systems* **28**(5), 718–731 (2012). DOI 10.1016/j.future.2011.12.002
  35. Goiri, I., Haque, M.E., Le, K., Beauchea, R., Nguyen, T.D., Guitart, J., Torres, J., Bianchini, R.: Matching Renewable Energy Supply and Demand in Green Datacenters. *Ad Hoc Networks* **25**, Part B, 520–534 (2015). DOI 10.1016/j.adhoc.2014.11.012
  36. Goiri, I., Katsak, W., Le, K., Nguyen, T.D., Bianchini, R.: Parasol and GreenSwitch: Managing Datacenters Powered by Renewable Energy. In: Proc. 18th Intl. Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS'13), Houston, USA, pp. 51–64. ACM (2013). DOI 10.1145/2451116.2451123
  37. Guenter, B., Jain, N., Williams, C.: Managing Cost, Performance, and Reliability Trade-offs for Energy-aware Server Provisioning. In: Proc. IEEE INFOCOM 2011, Shanghai, China, pp. 1332–1340. IEEE (2011). DOI 10.1109/INFCOM.2011.5934917
  38. Hatcher, J.: Follow the Moon. *Data Centre Management Magazine* (2013). <http://datacentremangement.com/news/view/follow-the-moon>
  39. Intel Corporation: Intelligent Platform Management Interface (IPMI). <http://www.intel.com/content/www/us/en/servers/ipmi/ipmi-home.html>
  40. Intel Corporation: Intel 64 and IA-32 Architectures Software Developer's Manual, Volume 3B: System Programming Guide, Part 2 (2015)
  41. Jarus, M., Oleksiak, A., Piontek, T., Weglarz, J.: Runtime Power Usage Estimation of HPC Servers for Various Classes of Real-life Applications. *Future Generation Computer Systems* **36**, 299–310 (2014). DOI 10.1016/j.future.2013.07.012
  42. Kansal, A., Zhao, F., Liu, J., Kothari, N., Bhattacharya, A.A.: Virtual Machine Power Metering and Provisioning. In: Proc. 1st ACM Symposium on Cloud Computing (SoCC'10), Indianapolis, USA, pp. 39–50. ACM (2010). DOI 10.1145/1807128.1807136

43. Kava, J.: Google's Three steps for a Zero Carbon Green Data Center. *DatacenterDynamics* London (2011). <http://www.greenm3.com/gdcblog/2011/12/6/googles-three-steps-for-a-zero-carbon-green-data-center.html>
44. Kecskemeti, G., Kertesz, A., Cs. Marosi, A., Nemeth, Z.: Strategies for Increased Energy Awareness in Cloud Federations. In: *High-Performance Computing on Complex Environments*, chap. 19, pp. 365–382. John Wiley & Sons, Inc. (2014). DOI 10.1002/9781118711897.ch19
45. Kim, S.g., Eom, H., Yeom, H.Y.: Virtual Machine Consolidation Based on Interference Modeling. *Journal of Supercomputing* **66**(3), 1489–1506 (2013). DOI 10.1007/s11227-013-0939-2
46. Kirby, B., Hirst, E.: Customer-Specific Metrics for the Regulation and Load-Following Ancillary Services. Tech. Rep. CON-474, Oak Ridge National Laboratory (2000)
47. Koomey, J.: Growth in Data Center Electricity Use 2005 to 2010. Report, Analytics Press, Oakland, USA (2011)
48. Kousiouris, G., Cucinotta, T., Varvarigou, T.: The Effects of Scheduling, Workload Type and Consolidation Scenarios on Virtual Machine Performance and Their Prediction Through Optimized Artificial Neural Networks. *Journal on Systems and Software* **84**(8), 1270–1291 (2011). DOI 10.1016/j.jss.2011.04.013
49. Kousiouris, G., Menychtas, A., Kyriazis, D., Gogouvtis, S., Varvarigou, T.: Dynamic, Behavioral-based Estimation of Resource Provisioning Based on High-level Application Terms in Cloud Platforms. *Future Generation Computer Systems* **32**, 27–40 (2014). DOI 10.1016/j.future.2012.05.009
50. Le, K., Bianchini, R., Martonosi, M., Nguyen, T.D.: Cost- and Energy-Aware Load Distribution Across Data Centers. In: *Proc. 2009 SOSP Workshop on Power Aware Computing and Systems (HotPower'09)*, Big Sky, USA. USENIX Association (2009)
51. Lee, E.K., Kulkarni, I., Pompili, D., Parashar, M.: Proactive Thermal Management in Green Datacenters. *Journal of Supercomputing* **60**(2), 165–195 (2012). DOI 10.1007/s11227-010-0453-8
52. Lewis, A.W., Tzeng, N.F., Ghosh, S.: Runtime Energy Consumption Estimation for Server Workloads Based on Chaotic Time-series Approximation. *ACM Trans. on Architecture and Code Optimization* **9**(3), 15:1–15:26 (2012). DOI 10.1145/2355585.2355588
53. Li, C., Zhou, R., Li, T.: Enabling Distributed Generation Powered Sustainable High-performance Data Center. In: *Proc. 19th Intl. Symposium on High Performance Computer Architecture (HPCA'13)*, Shenzhen, China, pp. 35–46. IEEE (2013). DOI 10.1109/HPCA.2013.6522305
54. Liu, Z., Lin, M., Wierman, A., Low, S.H., Andrew, L.L.: Geographical Load Balancing with Renewables. *SIGMETRICS Performance Evaluation Review* **39**(3), 62–66 (2011). DOI 10.1145/2160803.2160862
55. Macias, M., Guitart, J.: SLA Negotiation and Enforcement Policies for Revenue Maximization and Client Classification in Cloud Providers. *Future Generation Computer Systems* **41**, 19–31 (2014). DOI 10.1016/j.future.2014.03.004
56. Mann, Z.A.: Allocation of Virtual Machines in Cloud Data Centers – A Survey of Problem Models and Optimization Algorithms. *ACM Computing Surveys* **48**(1), 11:1–11:34 (2015). DOI 10.1145/2797211
57. Mann, Z.A.: Approximability of Virtual Machine Allocation: Much Harder Than Bin Packing. In: *Proc. 9th Hungarian-Japanese Symposium on Discrete Mathematics and Its Applications*, Fukuoka, Japan, pp. 21–30 (2015)
58. McCullough, J.C., Agarwal, Y., Chandrashekar, J., Kuppuswamy, S., Snoeren, A.C., Gupta, R.K.: Evaluating the Effectiveness of Model-based Power Characterization. In: *Proc. 2011 USENIX Annual Technical Conference (ATC'11)*, Portland, USA, pp. 159–172. USENIX Association (2011)
59. Mobius, C., Dargie, W., Schill, A.: Power Consumption Estimation Models for Processors, Virtual Machines, and Servers. *IEEE Transactions on Parallel & Distributed Systems* **25**(6), 1600–1614 (2014). DOI 10.1109/TPDS.2013.183
60. Reiss, C., Tumanov, A., Ganger, G.R., Katz, R.H., Kozuch, M.A.: Heterogeneity and Dynamicity of Clouds at Scale: Google Trace Analysis. In: *Proc. 3rd ACM Symposium on Cloud Computing (SoCC'12)*, San Jose, USA, pp. 7:1–7:13. ACM (2012). DOI 10.1145/2391229.2391236

61. Ren, C., Wang, D., Urgaonkar, B., Sivasubramaniam, A.: Carbon-Aware Energy Capacity Planning for Datacenters. In: Proc. 20th Intl. Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS'12), Arlington, USA, pp. 391–400. IEEE (2012). DOI 10.1109/MASCOTS.2012.51
62. Sedaghat, M., Hernandez, F., Elmroth, E.: Unifying Cloud Management: Towards Overall Governance of Business Level Objectives. In: Proc. 11th IEEE/ACM Intl. Symposium on Cluster, Cloud and Grid Computing (CCGrid'11), Newport Beach, USA, pp. 591–597. IEEE (2011). DOI 10.1109/CCGrid.2011.65
63. Shi, L., Furlong, J., Wang, R.: Empirical Evaluation of Vector Bin Packing Algorithms for Energy Efficient Data Centers. In: Proc. 18th IEEE Symposium on Computers and Communications (ISCC'13), Split, Croatia, pp. 9–15. IEEE (2013). DOI 10.1109/ISCC.2013.6754915
64. Subirats, J., Guitart, J.: Assessing and Forecasting Energy Efficiency on Cloud Computing Platforms. *Future Generation Computer Systems* **45**, 70–94 (2015). DOI 10.1016/j.future.2014.11.008
65. Tesfatsion, S., Wadbro, E., Tordsson, J.: A Combined Frequency Scaling and Application Elasticity Approach for Energy-efficient Cloud Computing. *Sustainable Computing: Informatics and Systems* **4**(4), 205–214 (2014). DOI 10.1016/j.suscom.2014.08.007
66. UEFI Forum: Advanced Configuration and Power Interface Specification Revision 5.0a (2013). <http://www.acpi.info/spec50a.htm>
67. United States Environmental Protection Agency (EPA): Renewable Energy Certificates (RECs). <http://www.epa.gov/greenpower/gpmarket/rec.htm>
68. Varrette, S., Guzek, M., Plugaru, V., Besson, X., Bouvry, P.: HPC Performance and Energy-Efficiency of Xen, KVM and VMware Hypervisors. In: Proc. 25th Intl. Symposium on Computer Architecture and High Performance Computing (SBAC-PAD'13), Porto de Galinhas, Brazil, pp. 89–96. IEEE (2013). DOI 10.1109/SBAC-PAD.2013.18
69. Vasic, N., Scherer, T., Schott, W.: Thermal-aware Workload Scheduling for Energy Efficient Data Centers. In: Proc. 7th Intl. Conference on Autonomic Computing (ICAC'10), Washington, USA, pp. 169–174. ACM (2010). DOI 10.1145/1809049.1809076
70. Vitali, M., Pernici, B., O'Reilly, U.M.: Learning a Goal-Oriented Model for Energy Efficient Adaptive Applications in Data Centers. *Information Sciences* **319**, 152–170 (2015). DOI 10.1016/j.ins.2015.01.023
71. Wada, H., Suzuki, J., Yamano, Y., Oba, K.: E<sup>3</sup>: A Multiobjective Optimization Framework for SLA-Aware Service Composition. *IEEE Transactions on Services Computing* **5**(3), 358–372 (2012). DOI 10.1109/TSC.2011.6
72. Wang, L., Khan, S.U.: Review of Performance Metrics for Green Data Centers: A Taxonomy Study. *Journal of Supercomputing* **63**(3), 639–656 (2013). DOI 10.1007/s11227-011-0704-3
73. Xiao, Z., Song, W., Chen, Q.: Dynamic Resource Allocation Using Virtual Machines for Cloud Computing Environment. *IEEE Transactions on Parallel & Distributed Systems* **24**(6), 1107–1117 (2013). DOI 10.1109/TPDS.2012.283
74. Xu, J., Fortes, J.A.B.: Multi-Objective Virtual Machine Placement in Virtualized Data Center Environments. In: Proc. 2010 Intl. Conference on Green Computing and Communications & Intl. Conference on Cyber, Physical and Social Computing, Hangzhou, China, pp. 179–188. IEEE (2010). DOI 10.1109/GreenCom-CPSCoM.2010.137
75. Yang, H., Zhao, Q., Luan, Z., Qian, D.: iMeter: An Integrated VM Power Model Based on Performance Profiling. *Future Generation Computer Systems* **36**, 267–286 (2014). DOI 10.1016/j.future.2013.07.008
76. Zhang, Q., Zhani, M.F., Boutaba, R., Hellerstein, J.L.: Dynamic Heterogeneity-Aware Resource Provisioning in the Cloud. *IEEE Transactions on Cloud Computing* **2**(1), 14–28 (2014). DOI 10.1109/TCC.2014.2306427
77. Zhang, Y., Ansari, N.: HERO: Hierarchical Energy Optimization for Data Center Networks. *IEEE Systems Journal* **9**(2), 406–415 (2015). DOI 10.1109/JSYST.2013.2285606
78. Zhang, Y., Wang, Y., Wang, X.: GreenWare: Greening Cloud-Scale Data Centers to Maximize the Use of Renewable Energy. In: Proc. 12th ACM/IFIP/USENIX Intl. Middleware Conference, Lisbon, Portugal, *Lecture Notes in Computer Science*, vol. 7049, pp. 143–164. Springer (2011). DOI 10.1007/978-3-642-25821-3\_8