

Joint Computing and Electric Systems Optimization for Green Datacenters

Ali Pahlevan, Maurizio Rossi, Pablo G. Del Valle, Davide Brunelli, and David Atienza

Abstract This chapter presents an optimization framework to manage green datacenters using multi-level energy reduction techniques in a joint approach. A green datacenter exploits renewable energy sources and active Uninterruptible Power Supply (UPS) units to reduce the energy intake from the grid while improving its Quality of Service (QoS). At server level, the state-of-the-art correlation-aware Virtual Machines (VMs) consolidation technique allows to maximize server's energy efficiency. At system level, heterogeneous Energy Storage Systems (ESS) replace standard UPSs while a dedicated optimization strategy aims at maximizing the lifetime of the battery banks and to reduce the energy bill, considering the load of the servers. Results demonstrate, under different number of VMs in the system, up to 11.6% energy savings, 10.4% improvement of QoS compared to existing correlation-aware VM allocation schemes for datacenters and up to 96% electricity bill savings.

35.1 Introduction

Ever increasing demands for computing and growing number of clusters and servers in datacenters have ramped up the power consumption costs as an undesirable effect [21]. On the other hand, traditional fossil fuel concerns, carbon emissions and global warming impose the introduction of more sustainable energy sources and behavioural change of people [41], since 10% of the global consumption of electrical energy has been estimated to be consumed by Information Technology (IT) infrastructures [14].

Ali Pahlevan, Pablo G. Del Valle, and David Atienza
Embedded Systems Laboratory (ESL), EPFL, Switzerland
e-mail: {ali.pahlevan, pablo.garciadelvalle, david.atienza}@epfl.ch

Maurizio Rossi and Davide Brunelli
Department of Industrial Engineering, University of Trento, Italy
e-mail: {maurizio.rossi, davide.brunelli}@unitn.it

To optimize the operation of a datacenter, it is crucial to minimize both IT and cooling energy consumptions. Server consolidation [26] is one of the widely used techniques to reduce the energy overheads, which minimizes the number of active servers by packing workloads or Virtual Machines (VMs) into the minimal number of active servers exploiting a virtualized environment. Large virtualized datacenters use renewable energy to reduce their dependence on costly and brown energy from the grid [33].

In the recent years, all the big energy consumers in the IT market (Amazon, Google, Rackspace, etc.) have already introduced renewable energy sources in their supply chain, locating their infrastructures in suitable geographical locations around the world. The penetration of renewable and green energy sources is almost none for company owned datacenters, IT infrastructures located in the same corporate building where the business is run, mostly in urban environments.

Solar energy is the most effective renewable source employed in green datacenters since Photovoltaic (PV) modules can be easily located close by the datacenter and the converted energy can be immediately used without distribution. Moreover it is the most suitable for small to medium datacenters (up to few hundreds kW of IT power) located in urban environments where wind turbines and water storage infrastructures may not be built, given the space required for such infrastructures.

Renewable energy sources are not constant over the time, their intensity depends on weather, geographical position of the plant and seasons, moreover a maximum in the energy intake rarely corresponds with a maximum in the demand. However, estimating their short-term trend (one day ahead) with small error (Mean Average Percentage Error (MAPE) close to 10%) is possible, as it has been demonstrated in [11]. Similar results can be expected when dealing with electricity demand prediction at building scale (few tens of kW) [32]. To tackle the imbalance between energy intake and demand, a widespread monitoring system of the produced and consumed power over time is necessary, as well as efficient forecasting algorithms of datacenters load consumption are required to optimize the usage of Energy Storage Systems (ESS) that collects the surplus of green energy for future needs.

Variability and fast-changing characteristics of applications, for instance scale-out applications [17] (e.g. web search, MapReduce, etc.), affect the energy consumption of servers due to the dependency on external factors, e.g., number of clients/queries in the system. To this end, the impact of servers' energy consumption on the usage of green energy becomes more substantial and management of consumed energy will play a major role in lifetime and operation of ESS. Consequently, without consideration of minimizing datacenter energy consumption, many existing approaches to management of green energy and batteries are sub optimal.

In this chapter, we introduce and propose a multi-level and multi-objective framework for the optimization of green virtualized datacenters, to jointly minimize the energy consumption and the carbon footprint, exploiting renewable energy sources, state-of-the-art VMs allocation schemes and Hybrid Electric Systems (HES). With HES, we refer to electrical ESS where different battery technologies are employed together, allowing to compensate for the inherent drawbacks of each technology. We incorporated dynamic VM allocation into servers' powers by novel HES, and

optimization methods to maximize the battery banks lifetime are used. The framework consists of two modules running concurrently, the Datacenter Energy Controller which minimizes the energy consumption of datacenter without any significant Quality-of-Service (QoS) degradation and shares the real energy consumption data with the Green Energy Controller; and the Green Energy Controller that manages renewable sources and HES, providing feedback to the Datacenter Energy Controller.

The Datacenter Energy Controller is based on a state-of-the-art correlation-aware VM allocation scheme [22] due to a high correlation within a cluster of applications in virtualized datacenters. Regarding load correlation, the authors demonstrate that, having detailed information about the applications characteristics, as opposed to using stationary load values for the VMs (e.g. peak or average values), gives the opportunity to further reduce the energy consumption of a datacenter. On the other side, the QoS degradation occurs when the aggregated utilization among colocated VMs is beyond the CPU capacity of a server. It means that there will be some workloads which cannot be executed at the right time. Therefore, datacenter providers take into account the Service-Level Agreements requirements to satisfy the customers. The Green Energy Controller, based on [31], is a two-phase controller that takes into account the cost policies of the grid energy and exploits forecasts of both the datacenter's load and of the incoming energy from renewables. The framework uses PV modules as green energy source and two battery technologies (lead-acid and lithium-ion) for the HES that are used with different priorities and roles.

In current datacenters, not enough efforts have been dedicated to implement adaptive energy reduction techniques and real-time resource scheduling to manage efficiently IT equipment and renewable energy sources. The novelty of our work consists in the introduction of a HES architecture to replace standard Uninterruptible Power Supply (UPS) systems, which allows an active management and the full exploitation of the energy buffers for the locally-generated renewable energy. We also designed a dedicated control loop which connects the VMs allocation scheme to the HES manager and optimizes the resources in real-time. At the same time, the modular structure allows to use both general purpose models and high-end ones for performance evaluation, model verification and feasibility analysis.

35.2 Related Work

Renewable energy sources integration in the electricity grid and in particular green datacenters are currently a hot-topic. Different research ideas have been presented in the last few years that address the problem of exploiting local energy generation to mitigate grid energy demand of datacenters [18] and in general of any human activity [12]. At the same time, HES have been addressed in several works available in the literature. The fundamental idea behind HES management is to use batteries as energy buffers to store the amount of green energy that cannot be used directly by the connected loads. Different management approaches have been proposed to au-

tomatically control the energy flows from renewables to loads and storage units [16] and also hybrid solutions for battery banks have been demonstrated [40]. This is particularly of interest nowadays because of the large availability of second-life batteries from electric vehicles that can have up to 75% remaining capacity available for storage applications [25, 37]. Despite the market availability of hybrid storage systems is still far, the literature review demonstrates that these technologies worth the efforts for being implemented. In this work, we followed the approach proposed in [31] to shape the active-UPS (or HES) system presented in the following. The authors in [31] propose a two-phase control scheme that exploits intrinsic advantages of different battery technologies mitigating, at the same time, their drawbacks.

A number of research works present methods for server consolidation based on per-VM workload characteristics, i.e., the peak, off-peak, and average utilization of workload [26, 35]; which aim is to reduce heat dissipation of hot spot zones and improve overall power utilization in datacenters [9, 20]. In [34] authors propose abstract models to balance computing power in a datacenter by minimizing peak inlet temperatures. A holistic approach that manages IT, power and cooling equipment by dynamically migrating servers' workloads and adjusting cooling is presented in [13]. Experimental results for a virtual datacenter demonstrate a reduction by 35% in power consumption and 15% in cooling. Authors in [27] present a control-oriented model that considers cyber and physical dynamics in datacenters to study the potential impact of coordinating the IT and cooling controls. To achieve further power savings while maintaining the QoS level, joint relationships among VMs, like load correlations, have been exploited in recent works [36, 24, 19]. For instance, in [24], Meng et al. proposed a VM sizing technique that pairs two uncorrelated VMs into a super-VM by predicting the workloads. However, once the super-VMs are formed, this solution does not consider dynamic changes of the VMs' load, which limits further energy savings. Therefore, these approaches do not work well with non-stationary and fast-changing VM behaviors in particular for scale-out applications. In [22], a power-efficient solution is proposed based on the First-Fit-Decreasing heuristic to separate load correlated VMs especially targeting the characteristics of the scale-out applications. They also exploit server's Dynamic Voltage and Frequency Scaling (DVFS) techniques to achieve further energy savings. Note that, these schemes do not take into account the renewable energy sources and datacenter system model in modern green datacenters.

There is no evidence in the literature of the joint application of HES optimization and correlation-aware techniques to the optimization of datacenter energy consumption, and the potential savings (both from environmental and money perspectives) are clearly worth the effort for further investigation.

35.3 The System Modeling Framework

We introduce a novel green datacenter system model where datacenter equipment, PV modules, smart grid, and UPS are connected as shown in Figure 35.1. The IT

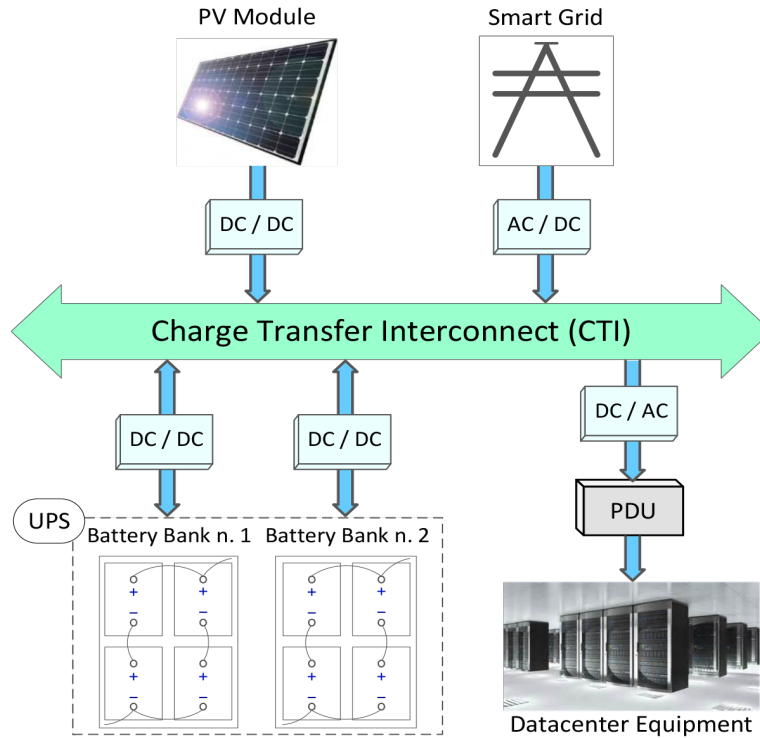


Figure 35.1: The complete system modeling framework.

equipment and cooling system inside this datacenter are the major contributors to power consumption than the other facilities. These components are combined using Power Distribution Units (PDUs) that eventually connect to the Charge Transfer Interconnect (CTI) bus that serve the whole facility [15]. In this framework the UPS is designed as a HES to provide both supply in case of grid outages and a buffer for green energy.

The system models two battery banks, a PV module and the bidirectional CTI bus, managed by a dedicated controller, not shown, as presented in [38]. Each unit is connected to the CTI by means of a bidirectional DC-DC converter for level shifting and charge routing, while the PV's one is unidirectional. Grid and PDUs are modeled in terms of power source and load, connected with the CTI by means of AC-DC and DC-AC converters, respectively.

We defined two constraints to the simulated system: (i) the exceeding renewable energy cannot be injected into the main grid (if it cannot be stored), and (ii) a peak/off-peak price scenario from a regulated electricity market for the energy taken from the grid (we considered the Zurich's tariff 7.5/14.9 CHFcent/kWh [1]).

Thus, renewable energy and batteries should completely sustain the load of the datacenter or, at least, provide supply during outages and periods with the highest price. These choices are justified by the fact that selling energy back to the grid, namely providing net-metering ancillary service to the Distribution System Operators (DSOs), follows rules that are country-specific and strongly depend on the interface between datacenter and energy network; moreover, datacenters are usually big energy consumers and it is unlikely to have enough excess green energy to justify the effort (economically and technologically) of improving the electric system to handle this task. The peak/off-peak price scenario in a regulated energy market instead, can be easily implemented also in a free energy market scenario where the energy price is continuously evolving; in this case, our assumption can be seen as a threshold on the freely variable price: while the free market price is below the threshold it is more convenient to buy from the grid, and the opposite when the price rises.

We developed a discrete-time framework (cf. see Chapter 6 for more details on different design space exploration options) that simulates the target green datacenter, with hourly time-steps. The Green Energy Controller manages the PV modules, the heterogeneous batteries and the CTI, the HES considered in this framework, and has been implemented using Matlab. The Datacenter Energy Controller, implemented in C++, manages the datacenter and VMs allocation scheme. Both components communicate using sockets for interprocess communication, while the time-step length of one hour guarantees that the time for VMs relocation (several GBs) does not overtake the actual execution time.

35.3.1 Energy Management Models

According to Figure 35.1, the power management problem is solved at the CTI bus level which is a DC path. Conversely, the system comprises both AC and DC sources/loads thus, for the former ones, it is required to consider the power factor component in the conversion. For example, considering the power intake from the grid, if we measure the total apparent power that enters the rectifier, for example on the grid side $P_{Grid}[VA] = V_{RMS} \cdot I_{RMS}$, this can be converted into active power (the useful power available on the DC side) according to the $P_{Grid}[W] = P_{Grid}[VA] \cdot \cos(\phi)$ where ϕ is the angle between Voltage and Current waveforms and $\cos(\phi)$ is called power factor. In addition, the converter's efficiency $\eta_X(\cdot)$ must be added to any transformation, since it depends on the actual power flowing with respect to the nominal one.

$$P_{Datacenter}^{CTI}(t) = P_{Grid}^{CTI}(t) + P_{PV}^{CTI}(t) + \sum_{n=1}^{nEES} \alpha \cdot P_{EES,n}^{CTI}(t) \quad (35.1)$$

$$P_{Grid}^{CTI}(t) = P_{Grid}(t) \cdot \cos(\phi) \cdot \eta_{ACDC}(\rho(t)) \quad (35.2)$$

$$P_{PV}^{CTI}(t) = P_{PV}(t) \cdot \eta_{DCDC}(\rho(t)) \quad (35.3)$$

$$P_{EES,n}^{CTI}(t) = P_{EES,n}(t) \cdot \eta_{DCDC}(\rho(t)) \quad (35.4)$$

$$P_{Datacenter}^{CTI}(t) \cdot \eta_{DCAC}(\rho(t)) = P_{Datacenter}(t) \quad (35.5)$$

$$\rho(t) = \frac{P_{out}}{P_{nom}} \cdot 100 \quad (35.6)$$

Equation (35.1) represents the power balance of the system, it states that the sum of the input from the grid, PV and battery arrays must be equal to the datacenter requirements, additionally the α is a directional parameter which can be $-/+1$ depending on the charging/discharging status (source or load of the system) and $nEES$ is the number of separated battery banks that compose the HES. Equation (35.2) to Equation (35.5) describe the AC-to-DC and DC-to-DC conversion functions used for each system component, where the conversion efficiency term $\eta_X(\cdot)$ depends on $\rho(t)$, the ratio of power requested by the system with respect to the nominal power delivered by the converter, which is expressed in percentage as defined by Equation (35.6).

In order to reduce the computational complexity and generalize the system's models we considered fixed power factor equal to one, fixed CTI voltage level and energy converters have been modeled considering a fixed 90% efficiency since detailed efficiency curves for high power equipment are not publicly provided by manufacturers [2] but still are claimed to work in the range of 80-95% (with loads down to $\rho(t) = 20\%$).

35.3.2 Electrical Energy Storage System

The HES can exploit two heterogeneous battery banks managed in hierarchical fashion: a lead-acid array (the battery bank n. 1) and a lithium-ion array (the battery bank n. 2). The battery model is based on the Peukert's law [29]. The goal is to model HES, that combine the advantages of the different battery technologies (lead-acid and lithium-ion). The module, as all the modules in the framework, has been conceived as a plug-and-play component; therefore, it can be easily replaced and adapted.

Equation (35.7) defines the State of Health (SoH) of the battery as a ratio between currently available charge capacity (C_{ref}) and the nominal one. Equation (35.8) defines the charge capacity as a linear combination of the previous charge and a term that depends on the charge which is drained, where C_{nom} is the nominal charge declared by manufacturer while Z_b , linear aging coefficient, is a parameter depending on the battery technology [30]. The following two equations (Equation (35.9) and

Equation (35.10)) allow to determine the State of Charge (SoC) and the equivalent battery current (I_{eq}), function of the current flowing from batteries (I), with respect to the nominal battery parameters: I_{ref} the reference discharge current (provided by the manufacturer and used to compute the reference charge), the Peukert's coefficient k_b and the charge actually used by the system, computed as current I_{eq} times time slot (t_{slot}) length in seconds. The SoH of the battery decreases only during discharge, so it is calculated only during discharge, whereas the SoC is updated during both charge and discharge cycles. More details about the model and its utilization can be found in [29, 30].

$$SoH(t+1) = \frac{C_{ref}(t+1)}{C_{nom}} \quad (35.7)$$

$$C_{ref}(t+1) = C_{ref}(t) - C_{nom} \cdot Z_b \cdot (SoC(t) - SoC(t+1)) \quad (35.8)$$

$$SoC(t+1) = \frac{C_{ref}(t) \cdot SoC(t) - (I_{eq}(t) \cdot t_{slot})}{C_{ref}(t)} \quad (35.9)$$

$$I_{eq}(t) = \left(\frac{|I(t)|}{I_{ref}} \right)^{(k_b-1)} \cdot I(t) \quad (35.10)$$

We tuned the parameters of the general purpose model (maximum and reference charge/discharge currents) according to commercial devices, a VARTA Professional Dual Power (230 Ah @ 12 V) [3] as lead-acid, and a StarkPower 'UltraEnergy' (100 Ah @ 12 V) [4] as the lithium-ion.

We preferred to double the size of the battery bank n. 1, with respect the lithium-ion one, because lead-acid technology is cheaper, easier to recycle and has a wider working temperature range. However, lead-acid batteries suffer from a limited number of sustainable cycles (i.e. lifetime). The lithium-ion technology instead offers at least one order of magnitude higher number of cycles, but it is also more expensive. To maximize the lifetime of the storage (in particular of the lead-acid bank), we put some constraints on the allowed Depth-of-Discharge (DoD) for both banks. To force both banks to work in the optimal range of SoC, we set the minimum SoC to 65% for the bank n. 1 and 70% for the bank n. 2. The remaining capacity is however available in the event of outage, thus providing standard UPS support.

Moreover, in the simulations we considered two configurations, the HES-1 where we have 48 kWh as lead-acid capacity (16.8 kWh available) and 24 kWh as lithium-ion capacity (7.2 kWh available); and the HES-2 with 96 kWh (33.6 kWh) and 48KWh capacity (14.4 kWh) respectively.

35.3.3 Photovoltaic Module

The PV module provides green energy accordingly to the intensity of the solar irradiance impinging on it, which in turn depends on the weather mostly. In this framework, we implemented it as a linearly varying voltage source, with integrated MPPT controller [31] and tuned accordingly to real device's characteristics [5]. Sun irradiance [6] and temperature profiles [7] for the year 2005 in Zurich have been used for tests.

$$P_{PV} = \left[P_{PV,STC} \cdot \left(\frac{G_T}{1000} \right) \cdot (1 - \gamma \cdot (T_j - 25)) \right] \cdot N_{PV,S} \cdot N_{PV,P} \quad (35.11)$$

$$T_j = T_{amb} + \left(\frac{G_T}{800} \right) \cdot NOCT - 20 \quad (35.12)$$

Equation (35.11) presents the linear model of the PV array, the parameters were evaluated in Nominal Operating Cell Temperature (NOCT) and Standard Test Conditions (STC) which are the nominal output power ($P_{PV,STC} = 2.65W$) in this case, the cell temperature (T_j), irradiance level ($G_T = 1000W/m^2 @ 25^\circ C$) and the temperature coefficient ($\gamma = 0.0043\%/^\circ C$), while $N_{PV,S}$ and $N_{PV,P}$ are the number of series and parallel cells in the module. The cell temperature is then obtained using Equation (35.12), where T_{amb} is the environmental temperature, $G_T = 800W/m^2 @ 20^\circ C$ and $NOCT = 45.5^\circ C$.

We tuned the PV module size considering two different cases of peak power production (hence the number of cells and panels) that are 10 kWp for the HES-1 simulating scenario and 30 kWp for the HES-2.

35.4 Simulation Framework Description

The overall diagram of our simulating framework, that jointly manages the Green Energy and Datacenter Energy Controllers, is shown in Figure 35.2. At the beginning of the simulation time horizon (off-line phase), the Green Energy Controller computes the expected energy budget for the datacenter, processing historical datacenter power profiles as well as the sun irradiance forecasts. This task is executed only once and provides a preliminary energy budget for the whole simulation horizon.

The on-line phase starts when the off-line phase of the Green Energy Controller sends the available energy budget to the Datacenter Energy Controller for the first time slot. Next, it waits until the VMs allocation to be completed according to the prediction of upcoming loads of VMs, then, receives back the real energy demand of the datacenter computed based on the real workload. Therefore, the Green Energy Controller compensates the differences between: (i) expected and available green

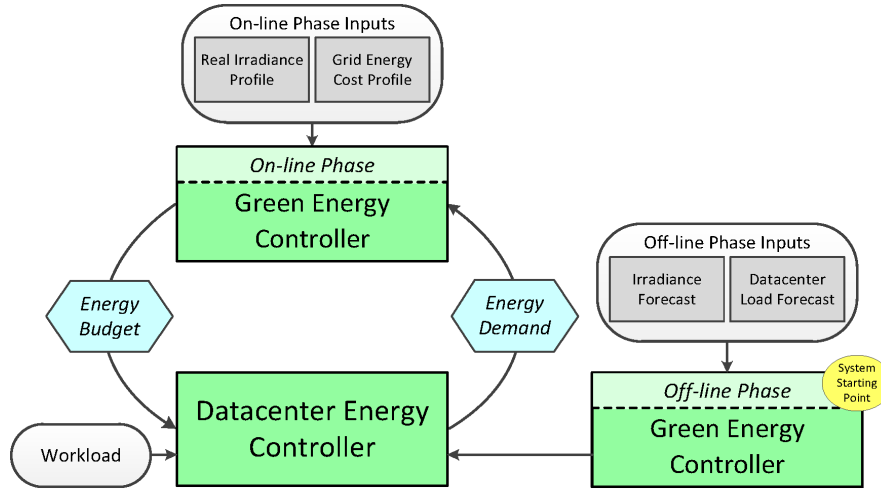


Figure 35.2: The simulating framework that jointly manages the Green Energy and Datacenter Energy Controllers. Off-line phase, as a starting point of simulation, is executed once at the beginning of the simulation time to compute expected energy budget for datacenter. In On-line phase, at each time slot, Datacenter Energy Controller first receives forecasted workload and energy budget from Green Energy Controller to allocate VMs to servers, then, sends back the real energy demand to Green Energy Controller.

energy, and (ii) real energy consumption and energy budget for the datacenter, using the lithium-ion battery as additional energy reserve or the grid if both banks in the HES have been drained. To this end, if the actual energy consumed by datacenter is higher than the expected, the Green Energy Controller compensates the datacenter energy requirements. At the end of each time slot the Green Controller provides an updated budget to the Datacenter Energy Controller for the VMs allocation of the next time slot.

On the other side, the Datacenter Energy Controller tries to find the best allocation for VMs on the servers at each time slot using the VMs specification from the previous time slot as incoming workload and the energy budget provided by Green Energy Controller. The goal is to allocate VMs to the minimal number of servers that yields in optimized total energy consumption of datacenter, as it will be explained in the following. After the allocation was completed, the Datacenter Energy Controller communicates the actual energy demand for the current time slot to the Green Energy Controller. Both of the controllers are invoked periodically, at every time slot, i.e., t_{slot} . The overall process of the framework and two controllers' communication has been shown in Figure 35.3. In the following sections, we describe these two controllers in detail.

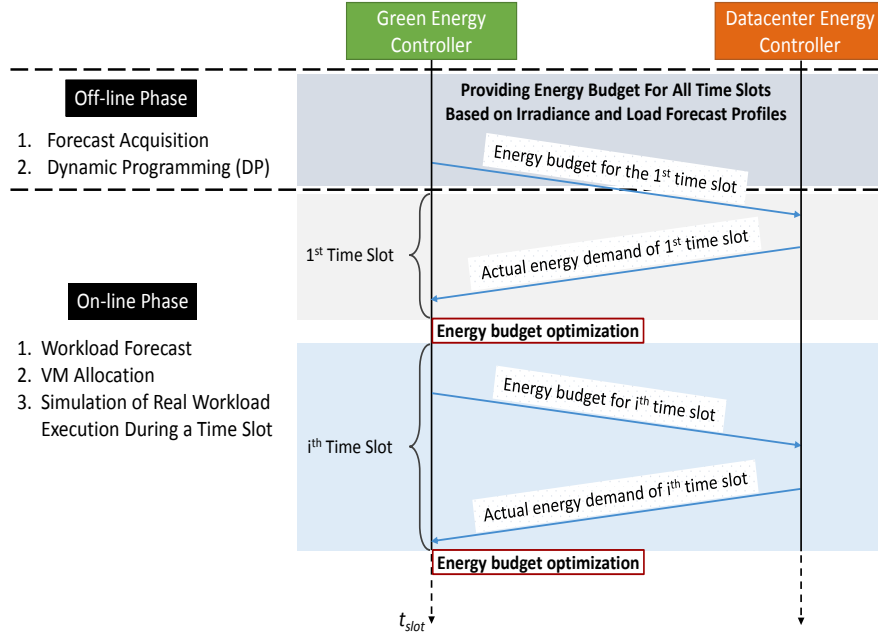


Figure 35.3: Overall process of the proposed framework - joint Datacenter and Green Energy Controllers.

35.4.1 Datacenter Energy Controller

In this section, we have considered the state-of-the-art correlation-aware VM allocation scheme as a datacenter power management solution [22]. Correlation refers to the VMs' utilizations when the peaks of two VMs occur at the same time during a certain time interval. Therefore, for using the servers resources efficiently during a time slot, highly correlated VMs should be placed apart, in different servers. Thereby, based on the VMs utilization patterns, the aggregated utilization of colocated VMs nearly reaches their server's capacity during a time slot. This favors consolidation and leads to power savings by lowering the number of active servers. In this context, due to the distributed operations of multiple VMs in a cluster, a high correlation within a cluster of VMs is observed, called intra-cluster correlation, rather than the correlation among different clusters targeted in other correlation-aware schemes [36, 19]. The correlation-aware VM allocation method has been proposed, in [22], while sharing cores among colocated VMs based on defining a cost function depending on QoS requirement to efficiently quantify the correlation between the VMs across a certain time horizon. Finally, a way to scale the Voltage/Frequency (V/f) level is provided to achieve more power savings without any QoS degradation. In this algorithm, the VMs are allocated such that the correlation among the allocated VMs in the server is minimized, while the server does not

exceed its total CPU capability, as well as the number of the active servers is minimized while satisfying performance requirements. Once all the VMs are allocated into servers, an optimal V/f level for each server is determined. This correlation-aware VM allocation algorithm is periodically invoked at every t_{slot} .

35.4.2 Green Energy Controller

The Green Energy Controller is a two-phase scheduler - off-line and on-line phases - that manages the CTI bus and provides guidelines to the Datacenter Energy Controller, by recursively solving the set of equations presented in Section 35.3. Moreover, see Chapter 10 for more details about combined design- and run-time-exploration and -adaptation approaches for computing systems.

The off-line phase's goal is to find the best resource allocation strategy to minimize the energy intake from the grid (i) and to maximize the lifetime of the lead-acid battery bank (ii) by minimizing the number of charge-discharge cycles and using as much as possible uninterrupted cycles. This is based on Dynamic Programming (DP) that is a strategy to solve complex problems by splitting them into lower complexity ones, solving and storing each solution, thus when a previously solved problem occurs the system looks up the previous solution saving computational time. It takes as input the expected workload of the datacenter, the price profile of the energy from the grid, and the irradiance forecasts for the whole time horizon [39, 31]; in this phase, the scheduler manages the battery bank n. 1 only. The algorithm ranks all the possible system states (charge to discharge, charge to charge, discharge to charge and discharge to discharge) for each time slot in the simulation horizon, that fulfills the above constraints. For each state transition it assigns a weight based on the battery usage, the higher the weight the lower the ranking. At the end, it provides an optimal energy budget for each time slot and the best utilization strategy for the lead-acid bank for the whole time horizon. Only the budget for the first time slot is then sent to the Datacenter Energy Controller and this message triggers the on-line phase. All the other energy budgets computed are kept in memory for the on-line phase to use them when the off-line concludes.

The on-line phase, for each time slot, optimizes the initial energy budget, computed by the off-line phase, trying to compensate the difference between expected workload and irradiance forecast with respect to the real data measured by the system. In the on-line phase the scheduler manages also the battery bank n. 2 mainly to compensate error in the forecasts and to maximize the lifetime of the lead-acid bank. This is a constrained multivariate optimization problem that has been solved numerically using the Matlab's `fmincon` [8] solver. For each time slot, the Green Scheduler must find the optimal currents balance in the CTI to minimize the energy taken from the grid (optimization goal), to fulfill the off-line lead-acid battery scheduling and to supply the load. For each component of the system (grid, PV, batteries and load) we set constrained boundaries for the currents and the input power from the grid, linear constraints for the CTI based on the Kirchhoff currents law and non-linear

constraints to compute the effect of energy converters and batteries' SoC. Problem's constraints (current flow direction for batteries and use of the grid) change in accordance with the system state, in this way it is possible, for example, to force the lithium-ion battery to be discharged when the lead-acid battery is recharged and the green energy is unavailable or lower than the load. At the end of the time slot, the actual energy balance is updated to the datacenter and this triggers also a new cycle of the simulator with the following t_{slot} .

35.5 Experimental Results

We validated the effectiveness and applicability of the proposed framework to larger scale problems using two weeks simulation horizon, workload traces obtained from a real datacenter setup and real irradiance and temperature profiles. We arranged the simulations in two separate sets, firstly we evaluated the best VMs allocation algorithm in terms of energy and QoS; secondly we placed this best scheme into the Datacenter Energy Controller and we executed the joint optimization framework.

35.5.1 Setup

We modeled a green urban datacenter consisting of medium sized facilities with two components: computing power consumption (IT equipment) (i) and Computer Room Air Conditioning (CRAC) power consumption as the cooling unit (ii). We evaluated the effectiveness of the proposed solution with a virtual testbed consisting of 250 servers where the servers are homogeneous. We targeted an Intel Xeon E5410 server configuration which consists of 8 cores and two frequency levels (2.0GHz and 2.3GHz), and used the power model proposed in [28].

To simulate the datacenter workload and energy demand we sampled the CPU utilization of a real datacenter setup every 5 min. for one day, then we duplicated the samples up to 14 days. Such assumption has been proved by real-trace studies, since the real datacenter's workload shows significant variability and a daily pattern during one week [23]. Finally, to generate different samples for each day, we synthesized fine-grained samples per 5 sec. with a lognormal random number generator [10], whose mean is the same as the collected value for the corresponding 5-minute sample rate.

We computed the irradiance forecasts implementing the algorithm presented in [11], an example of the two resulting sequences is depicted in Figure 35.4. At the same time, we used hourly averaged energy consumption profile from the real datacenter as forecast, which results in a smoothed profile compared to the original one.

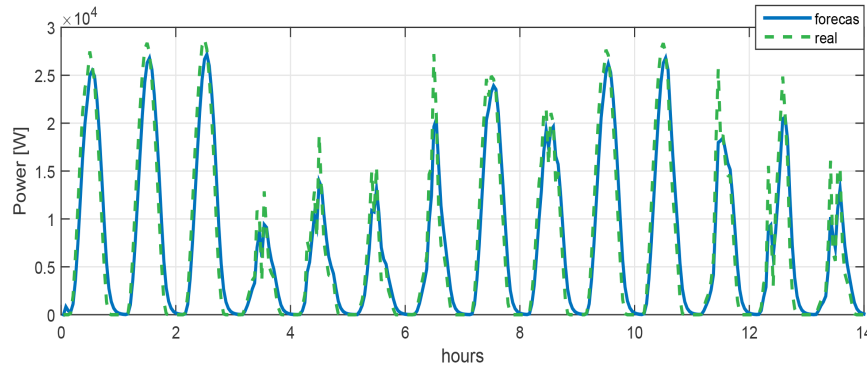


Figure 35.4: Solar power profile, forecasted vs. real.

35.5.2 Results

As previously introduced, we split the performance evaluation in two separate sets of experiments. To select the best VMs allocation scheme for power management to use with the Datacenter Energy Controller we compared the following three approaches:

- Best-Fit-Decreasing (BFD): a conventional best-fit-decreasing heuristic approach. In detail, after sorting VMs in decreasing order of their utilization, the algorithm allocates each VM to a server that provides the closest resource requirements with respect to this VM utilization (i.e., the server with the smallest remaining capacity is sufficient to contain the VM).
- Peak Clustering-based Placement (PCP) [36]: a correlation-aware VM allocation which clusters VMs using its Envelope-based correlation classification. The authors presented a static clustering-based VM allocation method by defining VMs' utilization in a time series as a binary sequence where the value becomes '1' when utilization is higher than a threshold value, otherwise, '0'. This algorithm first clusters VMs such that the envelopes of VMs' utilization included in different clusters do not overlap. Then, it allocates VMs to servers in order to colocate VMs in different clusters.
- Correlation-aware VM Placement (CVMP) [22] the correlation-aware VM allocation considered as the state-of-the-art approach and explained in Section 35.4.1.

Figure 35.5 compares the total energy consumption of the three approaches under different number of VMs (obtained by duplicating the trace for 250 VMs) in the system for a horizon of 14 days when we set the V/f level at the time of VM placement t_{slot} . The CVMP algorithm provides up to 11.6% and 7.3% energy savings compared to BFD and PCP respectively due to using the lower frequency levels more frequently. It is noteworthy that PCP provides almost similar results with BFD because, due to high and fast-changing correlations among VMs in our utilization

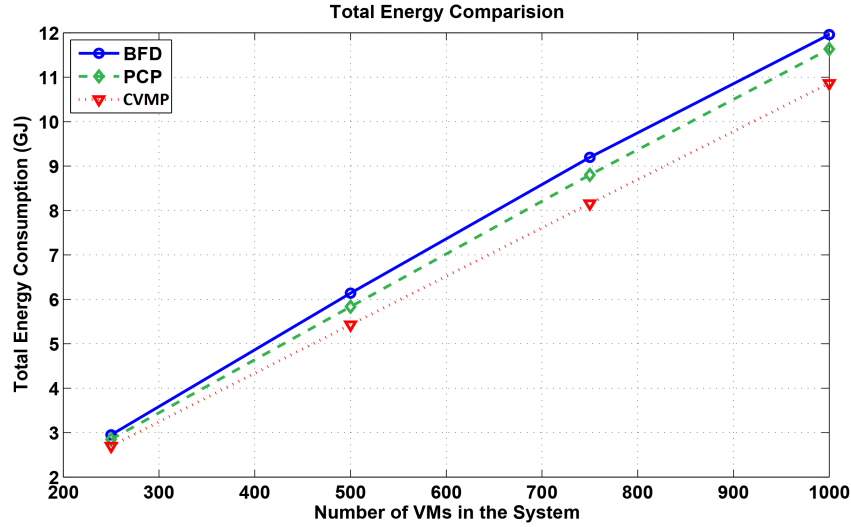


Figure 35.5: Total energy consumption of datacenter under different number of VMs for a horizon of 14 days.

traces, PCP classifies VMs into only 1 cluster during most of the time periods. When the number of clusters is 1, PCP behaves exactly the same as BFD. Note that the semi-linear trend of the energy consumption depends on the analogous behavior of the workload among different days, in a typical datacenter.

Table 35.1 shows the maximum violation defined as maximum per-period ratio of the number of over-utilized time instances (i.e., when the aggregated utilization among colocated VMs is beyond the CPU capacity of a corresponding server) to t_{slot} , during the two weeks under different number of VMs in the system. A graphical representation of these data is provided in Figure 35.6. As a result, the CVMP scheme provides a drastic reduction of the violations, up to 10.4% and 9.6% compared to BFD and PCP respectively. In CVMP method, VMs are allocated based

Table 35.1: Maximum violations (%) of ratio of over-utilized time instances to t_{slot} , during the entire periods, i.e., 336 hours (14 days) under different number of VMs scenario.

Approach	Number of VMs			
	250	500	750	1000
BFD	2.1	4.9	9.6	18.4
PCP	1.1	2.8	3.4	17.6
CVMP	0.85	2	3.1	8

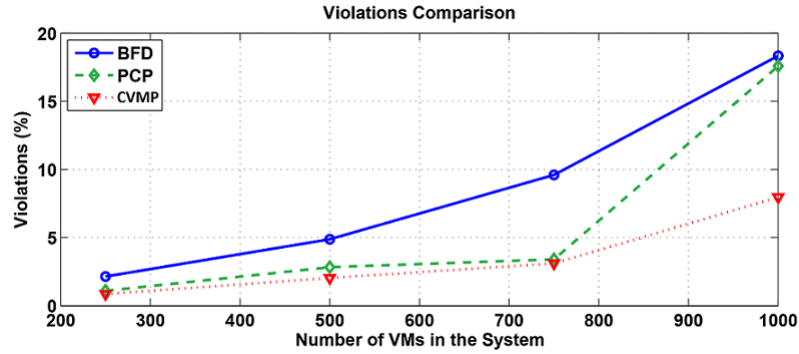


Figure 35.6: Trend of maximum violations (%) under different number of VMs for a horizon of 14 days.

on their peak utilizations, which were predicted from their history. Despite the provision based on the peak utilization, we observed quality degradation over the three approaches due to the mis-predictions of the peak utilization, especially during abrupt workload changes under increasing the number of VMs in the system. However, the CVMP method can statistically reduce the probability of the violation by colocating uncorrelated VMs. Thus, the probability of joint under-predictions among the colocated VMs is drastically decreased. Using the CVMP algorithm, we performed the complete framework simulation (VM allocation, green energy scheduling and communication between the two controllers) with $t_{slot} = 1$ hour, with predictions of upcoming workloads of datacenter using a last-value predictor.

Table 35.2 summarizes the results in terms of cost savings depending on the number of VMs, the HES-size and the season. The cost savings are computed as the difference between electricity cost to sustain the datacenter workload with or without the renewable energy sources. As expected with larger battery capacities (HES-2 configuration) we get higher savings. We compared also with the cost saving of using the PV panels without any storage (between brackets) to demonstrate the advantage of the proposed approach. Although in winter scenario the low irradiance and the cold weather strongly impact the renewable energy generation, causing the batteries to rarely reach the full charge, they still provide advantages in terms of savings. During summer instead the batteries are fully exploited resulting in higher savings with respect to the previous scenario. According to the model, during summer, when the HES system's usage is more intensive, we experienced a maximum SoH decreasing of 0.07% (ratio between nominal and remaining capacity), which means a lifetime longer than 15 years to reach the 70% of nominal capacity (lead-acid battery near the end of life). Finally, Figure 35.7 shows a two-days view (48 time slots) of the framework evolution with 500 VMs, summer irradiance and HES-2 configuration. We can observe the role of the energy buffer that allows to use green energy when there is no input from the PV panels (Figure 35.7-top) and the resulting money saving (Figure 35.7-bottom). In the specific time horizon depicted

Table 35.2: Overall framework results in terms of economic benefit of renewable-enabled datacenter with respect to a grid connected one. Two HES configurations are evaluated, HES-1 with 48 kWh as lead-acid and 24 kWh as lithium-ion capacity; HES-2 with 96 kWh and 48 kWh capacity respectively.

Configuration	Winter Savings (PV only)	Summer Savings (PV only)
250 VMs		
HES-1	29.30% (25.54%)	76.46% (57.86%)
HES-2	62.22% (38.72%)	96.13% (66.45%)
500 VMs		
HES-1	14.30% (13.16%)	55.92% (48.00%)
HES-2	38.43% (31.30%)	85.28% (61.59%)
750 VMs		
HES-1	9.53% (8.76%)	43.49% (40.16%)
HES-2	27.69% (24.86%)	73.39% (57.35%)
1000 VMs		
HES-1	7.05% (6.57%)	33.34% (32.51%)
HES-2	20.64% (19.16%)	65.28% (53.96%)

(Figure 35.7-middle), we experienced a low level of irradiance compared to other days in the overall horizon (cfr. Figure 35.4), it results in a lower amount of energy available to recharge the batteries, in particular the battery bank n. 1 which has a bigger capacity and a smaller recharge current with respect to the lithium-ion one. Similar considerations can be made for the other 3 cases that are not reported for the sake of summary.

35.6 Conclusion

In this chapter, we have presented a novel dynamic and multi-objective framework to manage the energy consumption of datacenter, battery banks lifetime and energy bill cost. The Datacenter Energy Controller minimizes the total energy consumption using the state-of-the-art correlation-aware VM allocation scheme for the given VMs specifications and energy budget provided by the Green Energy Controller while improving QoS requirements. In the Green Energy Controller, we use a real-time optimization technique to maximize the lifetime of battery banks and to reduce the energy bill by managing the PV source, in price-varying scenarios, and considering the energy consumed by the datacenter. Finally, we validated the effectiveness and applicability of our proposed system with the utilization traces obtained from a real datacenter setups. Our experimental results show that the proposed framework

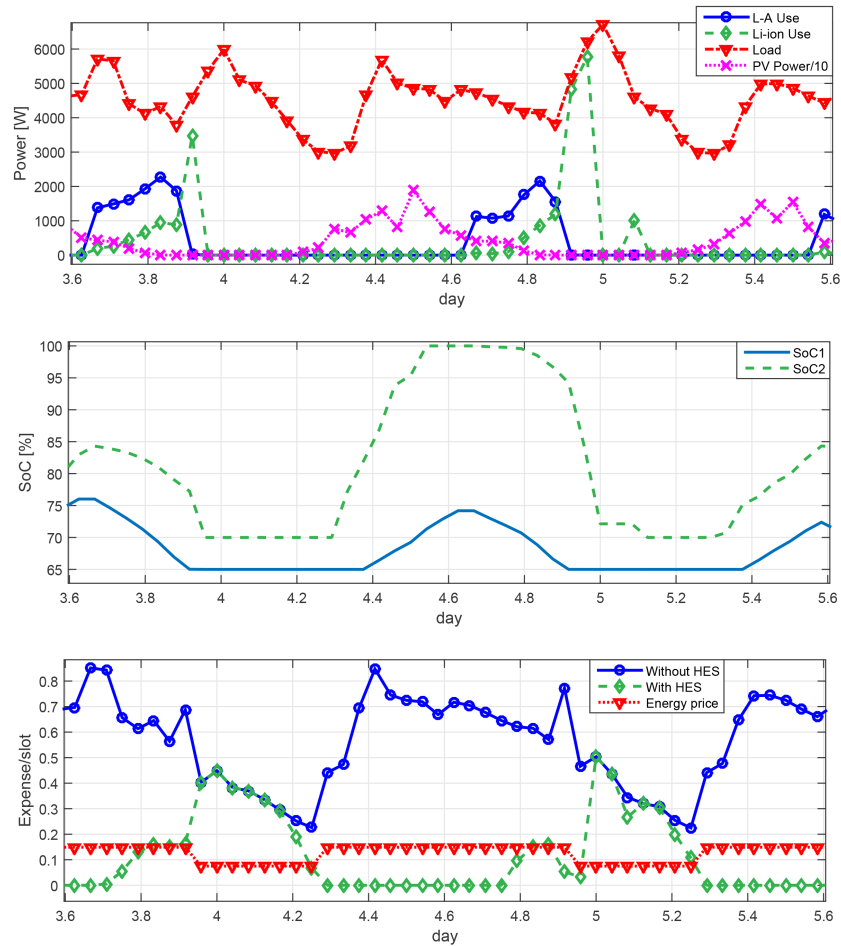


Figure 35.7: Two days framework evolution with 500 VMs, HES-2 (96 kWh lead-acid and 48 kWh lithium-ion capacity) configuration and summer irradiance (48 time slots). Power profile of the datacenter components (top); percentage SoC of the battery bank n. 1 (SoC1) and n. 2 (SoC2) (middle); cost per time slot (bottom).

provides up to 11.6% energy savings and up to 10.4% improvement of QoS level compared to existing conventional solutions under different number of VMs in the system, and up to 96% money saving in the electricity bill.

Acknowledgements This work has been partially supported by the EC FP7 GreenDataNet STREP project (Agreement No. 609000), and the YINS RTD project (no. 20NA21_150939), funded by Nano-Tera.ch with Swiss Confederation Financing and scientifically evaluated by SNSF.

References

1. [online] <http://www.strompreis.elcom.admin.ch/PriceDetail.aspx?placeNumber=261&OpID=565&-Period=2014&CatID=12>
2. [online] <http://www.schaeferpower.de/cms/en/produkte.html>
3. [online] <http://www.varta-automotive.com/en-gb/products/industrial/industrial-professional-dual-purpose>
4. [online] <http://www.starkpower.com/spnews/energystoragebatt>
5. [online] <http://www.enfsolar.com/pv/cell-datasheet/429>
6. [online] http://www.soda-is.com/eng/services/services_radiation_free_eng.php
7. [online] <http://www.tutienpo.net/en/Climate>
8. [online] <http://www.mathworks.com/help/optim/ug/fmincon.html>
9. Bash, C., Forman, G.: Cool job allocation: Measuring the power savings of placing jobs at cooling-efficient locations in the data center. In: 2007 USENIX Annual Technical Conference on Proceedings of the USENIX Annual Technical Conference, ATC'07, pp. 29:1–29:6. USENIX Association, Berkeley, CA, USA (2007)
10. Benson, T., et al.: Understanding data center traffic characteristics. *ACM SIGCOMM Computer Communication Review* **40**(1), 92–99 (2010)
11. Bergonzini, C., et al.: Comparison of energy intake prediction algorithms for systems powered by photovoltaic harvesters. *Microelectronics Journal* **41**(11), 766–777 (2010)
12. Carpinelli, G., Celli, G., Mocci, S., Mottola, F., Pilo, F., Proto, D.: Optimal integration of distributed energy storage devices in smart grids. *IEEE Transactions on Smart Grid* **4**(2), 985–995 (2013)
13. 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: Network Operations and Management Symposium (NOMS), 2010 IEEE, pp. 615–622 (2010)
14. Clark, J.: It now 10 percent of world's electricity consumption, report finds (2013). [online] http://www.theregister.co.uk/2013/08/16/it_electricity_use_worse_than_you_thought/
15. Deng, N., et al.: Concentrating renewable energy in grid-tied datacenters. In: Sustainable Systems and Technology (ISSST), 2011 IEEE International Symposium on, pp. 1–6. IEEE (2011)
16. Farhangi, H.: The path of the smart grid. *Power and Energy Magazine, IEEE* **8**(1), 18–28 (2010)
17. Ferdman, M., et al.: Clearing the clouds: a study of emerging scale-out workloads on modern hardware. *ACM SIGARCH Computer Architecture News* **40**(1), 37–48 (2012)
18. Goiri, I., Katsak, W., Le, K., Nguyen, T.D., Bianchini, R.: Designing and managing datacenters powered by renewable energy. *IEEE Micro* (3), 8–16 (2014)
19. Halder, K., et al.: Risk aware provisioning and resource aggregation based consolidation of virtual machines. In: Cloud Computing (CLOUD), 2012 IEEE 5th International Conference on, pp. 598–605
20. J. Leverich M. Monchiero, V.T.P.R., Kozyrakis, C.: Power management of datacenter workloads using per-core power gating. *Computer Architecture Letter* **8**(2), 48–51 (2009)
21. Katz, R.H.: Tech titans building boom (2009)
22. Kim, J., et al.: Correlation-aware virtual machine allocation for energy-efficient datacenters. In: Design, Automation & Test in Europe (DATE) Conference, pp. 1345–1350 (2013)
23. Liu, Z., Chen, Y., Bash, C., Wierman, A., Gmach, D., Wang, Z., Marwah, M., Hyser, C.: Renewable and cooling aware workload management for sustainable data centers. In: Proceedings of the 12th ACM SIGMETRICS/PERFORMANCE Joint International Conference on Measurement and Modeling of Computer Systems, SIGMETRICS '12, pp. 175–186. ACM (2012)
24. Meng, X., et al.: Efficient resource provisioning in compute clouds via vm multiplexing. In: Proceedings of the 7th international conference on Autonomic computing, pp. 11–20. ACM (2010)

25. Mukherjee, N., Strickland, D.: Control of cascaded dc-dc converter-based hybrid battery energy storage systems - part i: Stability issue. *IEEE Transactions on Industrial Electronics* **63**(4), 2340–2349 (2016)
26. Pakbaznia, E., et al.: Minimizing data center cooling and server power costs. In: Proceedings of the 14th ACM/IEEE international symposium on Low power electronics and design, pp. 145–150. ACM (2009)
27. Parolini, L., Sinopoli, B., Krogh, B., Wang, Z.: A cyber-physical systems approach to data center modeling and control for energy efficiency. *Proceedings of the IEEE* **100**(1), 254–268 (2012)
28. Pedram, M., et al.: Power and performance modeling in a virtualized server system. In: Parallel Processing Workshops (ICPPW), 2010 39th International Conference on, pp. 520–526
29. Riffonneau, Y., et al.: System modelling and energy management for grid connected pv systems with storage. In: 23rd European Photovoltaic Solar Energy Conference and Exhibition, pp. 3447–3451 (2008)
30. Riffonneau, Y., Bacha, S., Barruel, F., Ploix, S.: Optimal power flow management for grid connected pv systems with batteries. *Sustainable Energy, IEEE Transactions on* **2**(3), 309–320 (2011)
31. Rossi, M., et al.: Real-time optimization of the battery banks lifetime in hybrid residential electrical systems. In: Design, Automation & Test in Europe (DATE) Conference, pp. 139–145 (2014)
32. Rossi, M., Brunelli, D.: Electricity demand forecasting of single residential units. In: Environmental Energy and Structural Monitoring Systems (EESMS), 2013 IEEE Workshop on, pp. 1–6. IEEE (2013)
33. Stewart, C., et al.: Some joules are more precious than others: Managing renewable energy in the datacenter. In: Proceedings of the Workshop on Power Aware Computing and Systems (2009)
34. Tang, Q., Gupta, S., Varsamopoulos, G.: Energy-efficient thermal-aware task scheduling for homogeneous high-performance computing data centers: A cyber-physical approach. *Parallel and Distributed Systems, IEEE Transactions on* **19**(11), 1458–1472 (2008)
35. Verma, A., et al.: pmapper: power and migration cost aware application placement in virtualized systems. In: *Middleware 2008*, pp. 243–264. Springer
36. Verma, A., et al.: Server workload analysis for power minimization using consolidation. In: Proceedings of the 2009 conference on USENIX Annual technical conference, pp. 28–28. USENIX Association
37. Wang, L.Y., Wang, C., Yin, G., Lin, F., Polis, M.P., Zhang, C., Jiang, J.: Balanced control strategies for interconnected heterogeneous battery systems. *IEEE Transactions on Sustainable Energy* **7**(1), 189–199 (2016)
38. Wang, Y., et al.: Charge migration efficiency optimization in hybrid electrical energy storage (hees) systems. In: *Low Power Electronics and Design (ISLPED) 2011 International Symposium on*, pp. 103–108 (2011)
39. Wang, Y., et al.: Optimal control of a grid-connected hybrid electrical energy storage system for homes. In: Design, Automation & Test in Europe Conference & Exhibition (DATE), 2013, pp. 881–886. IEEE (2013)
40. Wang, Y., Lin, X., Kim, Y., Xie, Q., Pedram, M., Chang, N.: Single-source, single-destination charge migration in hybrid electrical energy storage systems. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems* **22**(12), 2752–2765 (2014)
41. Zhang, Y., et al.: Greenware: Greening cloud-scale data centers to maximize the use of renewable energy. In: *Middleware 2011*, pp. 143–164. Springer