

ECOGreen: Electricity Cost Optimization for Green Datacenters in Emerging Power Markets

Ali Pahlevan, *Student Member, IEEE*, Marina Zapater, *Member, IEEE*, Ayse K. Coskun, *Senior Member, IEEE*, and David Atienza, *Fellow, IEEE*

Abstract—Modern datacenters need to tackle efficiently the increasing demand for computing resources while minimizing energy usage and monetary costs. Power market operators have recently introduced emerging demand-response programs, in which electricity consumers regulate their power usage following provider requests to reduce monetary costs. Among different programs, regulation service (RS) reserves are particularly promising for datacenters due to the high credit gain possibilities and datacenters’ flexibility in regulating their power consumption. Therefore, it is essential to develop bidding strategies for datacenters to participate in emerging power markets together with power management policies that are aware of power market requirements at runtime. In this paper we propose ECOGreen, a holistic strategy to jointly optimize the datacenter RS problem and virtual machine (VM) allocation that satisfies the hour-ahead power market constraints in the presence of electrical energy storage (EES) and renewable energy. We first find the best power and reserve bidding values as well as the number of active servers in a fast analytical way that works well in practice. Then, we present an online adaptive policy that modulates datacenter power consumption by controlling VMs CPU resource limits and efficiently utilizing demand-side EES and renewable power, while guaranteeing quality-of-service (QoS) constraints. Our results demonstrate that ECOGreen can provide 76% of the datacenter power consumption on average as reserves to the market, due to largely operating on renewable sources and EES. This translates into ECOGreen saving up to 71% electricity costs when compared to other state-of-the-art datacenter electricity cost minimization techniques that participate in the power market.

Index Terms—Green datacenters, power market, regulation service reserves, VM allocation, EES and renewable power, power management, electricity cost, quality-of-service.

1 INTRODUCTION

EVER increasing demands for computing and the growing number of clusters and servers in datacenters have ramped up computing power consumption world-wide, which is estimated to be at 1.3% of the global usage, and growing at a yearly rate of 20% [1]. Consequently, with the increase in power usage, the electricity cost of datacenters doubles every five years [2]. In this context, the latest generation of datacenters tend to use on-site electrical energy storage (EES) systems and renewable energy sources to reduce costs, carbon emissions, and their dependency on energy from the power grid [3], [4], [5], becoming what we herein name *green datacenters*. However, due to the instability and high variability of renewable energy production (i.e., solar and wind), matching the demand-side renewable production and load in a green datacenter is a challenging task, which forces green datacenters to be connected to the power grid.

Renewable energy sources are also being integrated on the supply-side. In fact, the European Union (EU) aims to integrate over 20% share of renewables in gross energy production for carbon emission reduction by 2020 [6], and a growth of 52% is expected in the US by 2040 [7]. However, with growing integration of renewables into the grid, the volatility and intermittency of renewable generation provide higher uncertainty to independent system operators (ISOs), who need to match supply

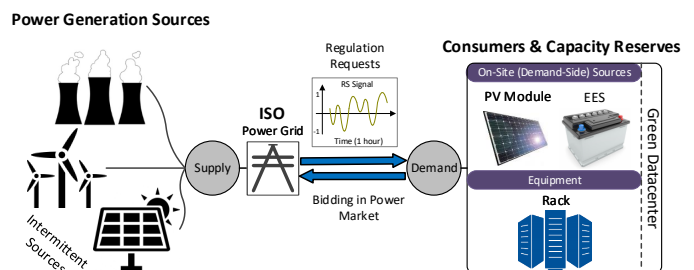


Fig. 1: Structure of the supply-side and demand-side (consumers) as capacity reserves.

and demand in the power grid in real-time. A solution that provides competitive prices and services for consumers is the use of demand-side capacity reserves [8], [9]. In this scenario, the ISO requests consumers to adapt their power consumption depending on its requirements (supply-demand matching). As datacenters are among the fastest growing electricity consumers, they are highly promising candidates to provide demand-side capacity reserves and reduce their electricity costs [10].

From the emerging power market perspective, large ISOs (such as PJM) allow the demand-side to provide reserves [8]. Among the various types of capacity reserves, regulation service (RS) reserves [9] are particularly interesting for green datacenters due to the relatively high value of such reserves and the capabilities of datacenters for providing high flexibility in their power consumption. In RS reserves provision, the demand-side (i.e., green datacenter) must dynamically modulate its power consumption to follow an RS signal broadcasted by the ISO every few seconds. RS signal is

- Ali Pahlevan, Marina Zapater, and David Atienza are with the Embedded Systems Laboratory (ESL), EPFL, Switzerland.
E-mail: {ali.pahlevan, marina.zapater, david.atienza}@epfl.ch
- Ayse K. Coskun is with the Department of Electrical and Computer Engineering, Boston University, Boston, MA 02215 USA.
E-mail: {acoskun}@bu.edu

the main factor used by ISO to balance the supply and demand in the power grid. In this scenario (depicted in Fig. 1), the demand-side acts as a capacity reserve that stabilizes the ISO power in the presence of the intermittency of renewable energies, and benefits from the power market rewards. However, the demand-side itself is also affected by the instability of on-site renewables. This poses an important challenge on the electricity cost and power minimization of green datacenters.

Recently, several studies have evaluated the capabilities and benefits of RS reserves provision in datacenters [11], [12], [13]. However, most of these studies disregard the use of demand-side renewables and EES when computing average power and reserve values in emerging power market bidding. In addition, an online policy that is aware of datacenter energy sources and workload constraints is needed to track the RS signal. The closest works to ours [14], [15] present an online tracking policy that exploits different server power modes to regulate server and datacenter power consumption. However, in contrast to our work, they do not consider the demand-side renewable and EES usage, neither do they jointly optimize the number of active servers and workload co-allocation.

A major challenge in integrating green datacenters in RS programs is that the computation of the best power consumption and reserve values (bidding), and the RS tracking problem are largely impacted by the availability of demand-side renewable and EES, incoming workload, efficient server selection, and virtual machine (VM) allocation policies. Therefore, to achieve the highest savings, a low-overhead method that incorporates all these aspects is required. This paper aims to advance the state-of-the-art in this direction. Our specific contributions are as follows:

- We introduce ECOGreen, a new Electricity Cost Optimization strategy for Green datacenters that computes the best average power and reserve bidding considering the renewable and EES energy for RS reserves provision in emerging power markets, along with determining the number of active servers.
- We develop a low-overhead online policy that enables a green datacenter to regulate its power and track the RS signal broadcasted every few seconds accurately, while also guaranteeing quality-of-service (QoS) constraints.
- We jointly manage VM allocation with the use of demand-side renewable and EES in RS reserves provision. To this end, we consider both the time-changing trends of renewable energy sources and the power loss in battery banks due to aging and charging sequences.
- Our results demonstrate that ECOGreen enables a green datacenter to provide 76% of its power consumption on average to the power market due to largely operating on renewable energy and EES. This allows the datacenter to save up to 71% in electricity cost compared to state-of-the-art power market participation strategies that do not use demand-side green energy sources. In addition, ECOGreen reduces monetary cost by 35% when compared to itself, but without participation in RS reserves (i.e., participation in traditional electric utilities), showing the combined benefit of using power market participation, renewable energy, and EES.

The remainder of this paper is organized as follows. Sec. 2 reviews related work. In Sec. 3, we provide an overview of the problem description and target optimization scenario. Sec. 4

describes the used green datacenter system modeling. In Sec. 5 we first solve the RS bidding problem, and then provide our runtime policy that dynamically regulates the green datacenter power to track the RS signal in real time. Sec. 6 and 7 present the experimental setup and results, respectively. Finally, a summary of the main conclusions of this work is provided in Sec. 8.

2 RELATED WORK

The research on power and cost management in green datacenters can be generally categorized in runtime policies and energy sources optimization in green datacenters, participation in emerging demand-response programs, and other energy-aware workload allocation methods.

2.1 Runtime Policies and Energy Sources Optimization in Green Datacenters

Pawlsh *et al.* [16] investigate the energy efficiency of datacenters considering different factors (e.g., utilization rates). The authors present a number of relevant performance metrics, which can be incorporated into the decision making process for datacenter's energy efficiency. Liu *et al.* [17] propose the GreenCloud architecture, which aims to minimize datacenter power consumption using live VM migration while meeting performance requirements. In addition, Wang *et al.* [18] discuss a taxonomy for performance, power, and thermal metrics to design a green datacenter considering economic and environmental effects (e.g., cost and carbon emission rate).

Differently from the previous works, to mitigate the harmful effects of carbon footprint, datacenters are equipped nowadays with renewable energy sources [3]. In this context, various research ideas have been presented in the last few years that address the problem of exploiting local energy generation and EES to reduce grid power demand of datacenters [19], [20]. To address this challenge, Goiri *et al.* [19] propose a parallel batch job scheduler to adjust the available solar energy to computational workload in a datacenter regardless of EES (battery) management. Also, Ghamkhari *et al.* [20] demonstrate how a convex-mathematical model can be used to maximize the total profit in datacenters with respect to wind energy sources and nature of workload.

Moreover, EES management has been addressed in several works available in the literature [21], [22]. The fundamental idea behind EES management is to use batteries as energy buffers to store the exceeding amount of green energy that cannot be used directly by the connected loads. In this context, various management approaches have been recently proposed to automatically control the energy flows from renewables to loads and storage units [22].

However, none of these works consider the impact of renewable and EES usage on electricity cost minimization in emerging power markets, which is our focus in this work.

2.2 Datacenter Cost Optimization On Emerging Power Markets

Power market operators have recently introduced smart grid demand-response programs, in which electricity consumers regulate their power usage following provider requirements [13]. Moreover, combining the optimization of the datacenter demand-response participation with the use of on-site renewable energy sources and EES is currently a popular research direction [23], [24].

A recent systematic comparison of multiple types of service markets [11] demonstrates that RS reserve provision is the most suitable and profitable program for datacenters. A few offline and online control policies for datacenters RS reserves provision are proposed in the literature [14], [15], [25], [26], [27]. Most of these studies [25], [26], [27] use highly simplified datacenter models for RS reserves provision. Chen *et al.* [14] propose an online policy that simply regulates the server power to track the instant value of the RS signal as accurately as possible. Chen *et al.* [15] also present a dynamic power control policy that modulates datacenter power consumption using server power capping techniques and different server power states. However, none of the previous works tackles this problem on a green datacenter equipped with on-site renewables and EES. Moreover, none of them has proposed a low overhead joint strategy that computes the market power and reserve bidding problem in a fast analytical way, along with determining the number of active servers needed for the allocation phase, while minimizing at the same time the electricity cost of the green datacenter.

2.3 Energy-Aware VM Allocation

Regarding energy-aware allocation methods, VM consolidation solutions that consider per-VM workload characteristics, like CPU-load correlation (i.e., the similarity of CPU utilization traces and the coincidence of their peaks), are widely used to achieve energy savings [28], [29], [30]. Kim *et al.* [28] present a CPU-load correlation-aware solution based on the First-Fit-Decreasing heuristic to separate CPU-load correlated VMs. Lin *et al.* [29] utilize the peak workload characteristics to measure the similarity of VMs' workload. The latest method achieves better results for VMs whose workload follows a Gaussian distribution. The main drawback of this approach is that it cannot be used for online management at large-scale datacenters due to its high computational overhead.

Dynamic allocation via migration is also used for minimizing datacenters cost and energy consumption. Ruan *et al.* [31] propose a dynamic migration-based VM allocation method to achieve the optimal balance between server utilization and energy consumption such that all servers operate at the highest performance-to-power levels. Wang *et al.* [32] also address a matching-based VM consolidation mechanism using migration such that active servers can operate close to a desirable utilization threshold. The main drawback of those approaches is their high VM migration overhead. Thus, as opposed to short-term decision, Chen *et al.* [33] propose a long-term VM consolidation mechanism such that the total demand of co-located VMs nearly reaches their host capacity during their lifetime period. This algorithm detects the utilization pattern of each VM based on four types of simple pulse functions. Nonetheless, this work disregards the original utilization pattern of the VMs, which is usually a combination of those simple types of functions, thus achieving lower energy savings.

None of the prior approaches considers jointly incorporating bidding, selection of the number of active servers, VM allocation, and the use of demand-side renewable and battery energy in an RS reserves provision context. To the best of our knowledge, ECOGreen is the first to address all these aspects together in a holistic strategy.

3 PROBLEM DESCRIPTION

In this section we provide a description of the overall scenario, the system we optimize, and the main assumptions taken. Figure 2

illustrates the proposed scenario and strategy for participating in the power market. Figure 2-System shows our green datacenter system, which comprises a green datacenter equipped with on-site (demand-side) renewable energy (PV modules), and an EES system, interconnected between them and to the power grid (ISO) via a charge transfer interconnect (CTI) bus.

From the green datacenter perspective, in this work the EES is used to provide both supply in case of grid outages and a buffer for green and grid power provided by the ISO. In addition, we assume that exceeding renewable energy cannot be injected back into the power grid (i.e., it can only be used or stored). Thus, renewable energy and EES are used to track the RS signal coming from the ISO or, at least, provide supply during outages.

From the emerging power market perspective, we focus on RS reserves in the hour-ahead power market as the price of reserves is high and the green datacenter can modulate its power consumption flexibly. In this program, the cost structure is different from the traditional electric utilities pricing in which datacenters' electricity cost is based on their peak power and energy consumption [34]. The RS program indeed provides incentives to participants; i.e., the electricity cost of a participant is determined by its average power consumption \bar{P} and the reserves R it can provision [35]. Green datacenter should declare both its \bar{P} and R , to the ISO an hour in advance. The datacenter is charged for the average power \bar{P} and credited for the provided reserves R . However, to be given a certain credit, each hour the datacenter is asked to modulate its power consumption dynamically to track the RS signal (z) every 4 seconds [36]. Part of the RS credit is reduced based on the magnitude of the tracking error. Moreover, if the tracking error exceeds a statistical tolerance constraint, the participant (i.e., datacenter) may lose its contract [37].

Furthermore, there can be other costs, such as transmission rights, if a datacenter wants to be completely in the wholesale market. Transmission rights are financial mediums related to congestion on power transmission lines. Congestion happens when the lowest-priced power cannot be conveyed to areas of the power grid that experience high demand. In this situation, the higher-priced electricity is dispatched to high-demanded areas by following a less congested path in the transmission system. As a result, the difference in the prices can be passed to consumers, leading to an increase in their electricity bills. However, power distribution companies provide a kind of insurance to protect their customers against these costs [38]. In our work, we focus on the operational cost of the datacenters with respect to any clearing price provided by the power market under a contract. In case of any price-changing scenario (i.e., price of average power consumption and reserve), our strategy is able to find a solution to optimize the datacenter's electricity cost.

In order to bid in the hour-ahead market (i.e., to find the values of \bar{P} and R), our ECOGreen strategy requires predicting, at the beginning of time slot T (i.e., every hour), sun irradiance forecasts, and per-VM CPU and memory utilization patterns (\tilde{U}_{cpu} and \tilde{U}_{mem}). Given the daily periodicity observed in the VMs of Google Cluster traces, this can be achieved by using the autoregressive integrated moving average (ARIMA) prediction model [39]. ARIMA considers the CPU and memory utilization from the previous week and forecasts the next-hour traces per VM. Moreover, among all the VMs, the worst-case prediction error is less than 10%.

Given the predicted VMs workloads, renewable energy forecasts, and battery status, we first compute the (\bar{P}, R) , number

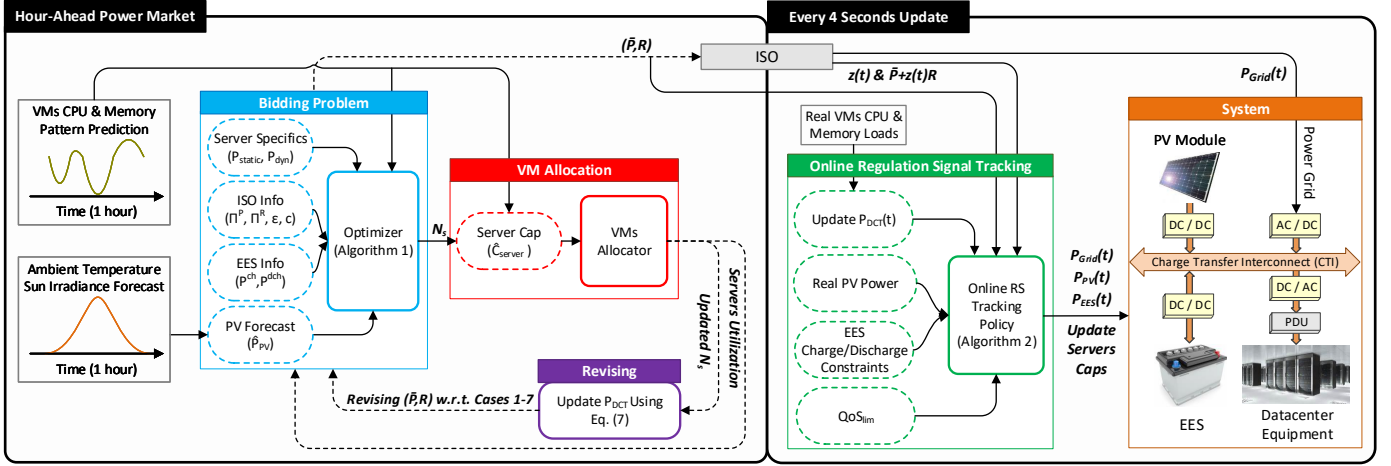


Fig. 2: Overall diagram of the proposed scenario and strategy, i.e., ECOGreen, including Bidding Problem, Allocation, Revising Bidding Values, and Online Policy phases.

of active servers (N_s), and the per-server capacity \hat{C}_{server} provided for the VM allocation phase. Then, the VMs are mapped to the servers based on \hat{C}_{server} every hour. Before sending bidding values to ISO, we need to recompute (\bar{P}, R) with respect to the estimated power consumption of the datacenter (\hat{P}_{DCT}), if all VMs would not have fit into N_s . This recalculation is only required once.

Once the ISO approves the bid, $z(t)$ is dynamically broadcasted to the green datacenter from the ISO every 4 seconds. Therefore, we need an online policy for tracking the power imposed by the ISO (i.e., $P_{Grid}(t) = \bar{P} + z(t)R$) with the smallest tracking error possible, given the real datacenter status (i.e., real VM's CPU utilization, memory footprint, and renewable power). We also aim to optimize the usage of renewable and battery power under the above-mentioned system constraints. To propose a solution to both the bidding in power market and the tracking of the RS signal, we need to accurately model the system.

4 OVERALL SYSTEM MODELING

In this section we describe all the models for the green datacenter system (as shown in Fig. 2-System). We use the green datacenter model proposed in our previous work [40] where datacenter equipment, PV modules, EES, and power grid are connected via a bidirectional CTI bus. In the following subsections we first present the used EES model, which is essential for correctly assessing the battery lifetime. Then, we define the model for the demand-side PV power generation. Next, we detail the considered green datacenter power model. Finally, we describe the power management model at the CTI bus that interconnects all components.

4.1 Electrical Energy Storage (EES) System

The EES comprises one homogeneous battery bank managed in a hierarchical fashion. The battery model is based on Peukert's law [41], and has been conceived as a plug-and-play component that can be easily replaced and adapted [40]. To correctly account for the benefits of the EES system we need to control the charging sequences of the battery bank, and to model battery ageing (i.e., power loss in EES).

Eq. (1) defines the state-of-the-health (SoH) of the battery as the ratio of currently available charge capacity (C_{ref}) to the nominal charge declared by the manufacturer (C_{nom}). Eq. (2) defines the charge capacity as a linear combination of the previous

charge and the charge that is drained. Z_b denotes the linear aging coefficient, which is dependant on the battery technology [42]. Finally, Eq. (3) and (4) determine the state-of-the-charge (SoC) and the equivalent battery current (I_{eq}), respectively. These are a function of the current flowing from batteries (I_b), and the nominal battery parameters: 1) the reference discharge current (I_{ref}) provided by the manufacturer and used to compute the reference charge, 2) the Peukert's coefficient (k_b), and 3) the charge actually used by the system and computed as current I_{eq} times time (t) in seconds. The SoH of the battery decreases only during discharge (therefore it is calculated only during discharge), whereas the SoC is updated during both charge and discharge cycles. Further details about the model and its usage can be found in literature [41], [42].

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

$$C_{ref}(t+1) = C_{ref}(t) - C_{nom} \cdot Z_b \cdot (SoC(t) - SoC(t+1)) \quad (2)$$

$$SoC(t+1) = \frac{C_{ref}(t) \cdot SoC(t) - I_{eq}(t) \cdot t}{C_{ref}(t)} \quad (3)$$

$$I_{eq}(t) = \left(\frac{|I_b(t)|}{I_{ref}} \right)^{(k_b-1)} \cdot I_b(t) \quad (4)$$

While lead-acid technology is cheaper, easier to recycle and has a wider working temperature range, it suffers from a limited number of sustainable cycles (i.e., lifetime). Therefore, we choose the lithium-ion technology instead, which offers at least one order of magnitude higher number of cycles (useful for tracking fast time-changing RS signal), but at a higher cost. To maximize the lifetime of the storage and make the battery bank work in the optimal range of SoC, we constraint the maximum depth of discharge (DoD) to 70% [40]. This indicates that only 30% of the total capacity of the battery can be used for bidding problem. The remaining capacity is however available in the event of an outage, thus providing standard UPS support. We tune the parameters of the general purpose model (maximum and reference charge/discharge currents) according to commercial devices (i.e., a StarkPower 'UltraEnergy' [43]).

4.2 Photovoltaic (PV) Module

The PV module provides energy proportional to the intensity of the solar irradiance impinging on it, which in turn depends

mostly on the weather. In our framework, we assume that the PV module has been modeled as a linearly-varying voltage source, with an integrated MPPT controller [44] and tuned accordingly to real device characteristics [45]. Real sun irradiance [46] and temperature profiles [47] are used for our experiments. Equation (5) presents the linear model of the PV array as follows:

$$P_{PV} = \left[P_{PV,STC} \cdot \left(\frac{G_T}{1000} \right) \cdot (1 - \gamma \cdot (T_j - 25)) \right] \cdot N_{PV,S} \cdot N_{PV,P} \quad (5)$$

$$T_j = T_{amb} + \left(\frac{G_T}{800} \right) \cdot NOCT - 20 \quad (6)$$

The output voltage of PV is affected if operating temperature, solar irradiance, and ambient temperature change. Therefore, in order to properly handle these changing parameters, i.e., in order to model the solar cells correctly, we consider the correction factor (γ). In this case, variation of solar irradiation and operating temperature are resulting in variable voltages. Therefore, to overcome this issue, a closed loop voltage control is assumed to be used to maintain the output voltage of PV at a particular desired value [48].

Based on these assumptions, the parameters are evaluated in nominal operating cell temperature (NOCT) and standard test conditions (STC), which yield the nominal output power ($P_{PV,STC}$) of 2.65W, irradiance level (G_T) of 1000W/m² @ 25°C, and a temperature coefficient (γ) of 0.0043%/°C. $N_{PV,S}$ and $N_{PV,P}$ are the number of series and parallel cells in the module. The cell temperature (T_j) is obtained using Eq. (6), where T_{amb} is the environmental temperature, $G_T = 800W/m^2$ @ 20°C and $NOCT = 45.5^\circ C$. We also tune the PV module size (the number of cells and panels) considering the peak power of the green datacenter with respect to its load, which is the most common approach to PV sizing [40].

4.3 Datacenter Power Model

We model total datacenter power consumption (P_{DCT}) as the sum of the power of servers: $P_{DCT} = \sum_{j=1}^{N_{DCT}} P_j$, where P_j and N_{DCT} specify the j^{th} server power and the total number of servers in the datacenter, respectively. Following the same methodology than in previous research [49], [50], the major contributors to power consumption in servers are the CPU, memory, fans and disks. Among these, CPU has the largest effect on power, and previous research shows that the power-frequency relation is linear for a given CPU-intensive workload [51]. Hence, server power can be calculated as [49]:

$$P_j = P_{j,static} + P_{j,dyn} \quad (7)$$

$$\begin{cases} P_{j,static} = P_{disk} + P_{fan} + P_{cpu}^{leak} + P_{cpu}^{idle} + P_{mem}^{idle} \\ P_{j,dyn} = P_{cpu}^{dyn} \cdot (U_{cpu_j}/100) + P_{mem}^{dyn} \cdot (U_{mem_j}/100) \end{cases}$$

where under $P_{j,static}$ we include all the contributions to power that are workload-independent. P_{disk} and P_{fan} are considered constants for our particular workload, and respectively account for the power consumption of disks and fans. P_{cpu}^{leak} refers to temperature-dependent leakage power. We consider a high fan speed and a low inlet temperature to reduce the effect of temperature-dependent leakage power, considering it as its worst-case constant. P_{cpu}^{idle} and P_{mem}^{idle} are constants that account for the idle power consumption of CPU and memory respectively. $P_{j,dyn}$ accounts for server dynamic power, and is proportional to the CPU and memory utilization (U_{cpu_j} and U_{mem_j} , respectively). P_{cpu}^{dyn} and P_{mem}^{dyn} are fitted constants

obtained under the same experimental conditions used in previous research [49] for the same set of CPU-intensive workloads.

4.4 Power Management Model on Charge Transfer Interconnect (CTI)

As we consider two types of power losses in the battery, i.e., charge capacity rate and SoH degradation, which are dependent on charge/discharge currents, a fine-grained system model is used to manage the energy sources and also simulate a realistic scenario.

In this framework, the IT equipment is connected via a power distribution unit (PDU) to the CTI bus that serves the whole facility [52]. The system comprises both AC and DC sources/loads, while the CTI is a DC path. Therefore, each unit is connected to the CTI through a converter interface. The EES is attached to the CTI by means of a bidirectional DC-DC converter for level shifting and charge routing, whereas the PV one is unidirectional. Power grid and datacenter are modeled as power source and load (i.e., P_{Grid} and P_{DCT}), connected to the CTI by means of AC-DC and DC-AC converters, respectively. All converters have an efficiency (η_X), defined as the ratio of power requested by the system w.r.t. the nominal power delivered by the converter. Equations (8) to (11) describe the AC-to-DC and DC-to-DC conversion functions used for each system component, as follows:

$$P_{Grid}^{CTI}(t) = P_{Grid}(t) \cdot \eta_{ACDC} \quad (8)$$

$$P_{PV}^{CTI}(t) = P_{PV}(t) \cdot \eta_{DCDC} \quad (9)$$

$$P_{EES}^{CTI}(t) = P_{EES}(t) \cdot \eta_{DCDC} \quad (10)$$

$$P_{DCT}^{CTI}(t) \cdot \eta_{DCAC} = P_{DCT}(t) \quad (11)$$

To control and optimize the amount of power usage of each source, the power management problem is solved at the CTI bus level. Equation (12) represents the power balance of the system and states that the sum of the input from the power grid, PV and EES must be equal to the datacenter requirements. α is a directional parameter that can be -/+1 depending on the charging/discharging status.

$$P_{DCT}^{CTI}(t) = P_{Grid}^{CTI}(t) + P_{PV}^{CTI}(t) + \alpha \cdot P_{EES}^{CTI}(t) \quad (12)$$

In order to reduce the computational complexity and generalize the system models, we consider a fixed CTI voltage level (V_{CTI}) and converters with $\eta_X = 90\%$ efficiency. Detailed efficiency curves for high-power equipment are not publicly provided by manufacturers [53] but still efficiency is claimed to stay within the 80-95% range for loads down to 20%.

5 ECOGREEN: PROPOSED ELECTRICITY COST OPTIMIZATION STRATEGY FOR GREEN DATACENTER

In order to optimally solve the cost minimization problem, we need to combine the bidding and allocation phases. Nonetheless, this optimization problem is non-linear and NP-hard. In our previous works [30], [54], we showed that the execution time of integer linear programming (ILP)-based VM allocation method for a small-scale scenario with 150 VMs is more than 2 hours; thus, runtime allocation would be unfeasible. In addition, combining allocation with bidding makes the execution time exponentially higher due to the increment of problem variables. Hence, in the following sections, we propose a strategy (i.e., ECOGreen) splitting these two phases, which solves this complex problem in a reasonable execution time.

5.1 Bidding Solution

5.1.1 General Bidding Problem - Initial Details

The first step in the RS hour-ahead power market is to compute the bidding for average power and reserves (\bar{P}, R) for every 1-hour time slot (T) (see Fig. 2). The best bid is the one that minimizes the monetary costs. This is achieved by trying to match \bar{P} and R while reducing the tracking error on the 4-second RS signal at each time t . As in our case the datacenter is equipped with on-site renewable power and EES, the power that needs to be provided by the grid at each time t (i.e., every 4 seconds), can be estimated based on the predicted datacenter power ($\hat{P}_{DCT}(t)$), the forecasted renewable power (\hat{P}_{PV}), and the current EES charge (P_{EES}), as follows:

$$\hat{P}_{Grid}(t) = \hat{P}_{DCT}(t) - \alpha \cdot P_{EES}(t) - \hat{P}_{PV}(t), \quad \alpha = \pm 1 \quad (13)$$

where $\hat{P}_{DCT}(t)$ is predicted based on the VMs loads (using the ARIMA model) and depends on the number of active servers (N_s). Due to the large time delay of booting a server, we assume that during the time slot neither the active servers are shut down nor new servers are turned on. That is, N_s is decided at the beginning of each time slot T and remains unchanged during it.

After finding the bidding values, during the time slot and as the real values for datacenter power and renewable energy are determined, we need to modulate the datacenter power and manage the green energy to track the RS signal (see Section 5.4). That is, to minimize error (and subsequently cost), we must ensure that $\hat{P}_{Grid}(t) \approx \bar{P} + z(t)R$ for every t .

In this scenario, and without loss of generalization, the electricity cost minimization problem for every t during the time slot T can be formulated as in recent work [10], as follows:

$$\min_{\bar{P}, R} Cost(t) = \pi^P \bar{P} - \pi^R R + \pi^R c \frac{|\hat{P}_{Grid}(t) - (\bar{P} + z(t)R)|}{R} \quad (14)$$

Subject to

$$1. \bar{P} + R \leq N_s \cdot P_s^{max} + P_{EES}^{ch}(t) = a \quad (15)$$

$$2. \bar{P} - R \geq \max(N_s \cdot P_{static} + \hat{P}_{DCT_{dyn}}(t) - \hat{P}_{PV}(t) - P_{EES}^{dch}(t), 0) = b \quad (16)$$

$$3. \frac{|\hat{P}_{Grid}(t) - (\bar{P} + z(t)R)|}{R} \leq \varepsilon \quad (17)$$

$$4. \bar{P} \geq 0 \quad (18)$$

$$5. R \geq 0 \quad (19)$$

The minimization problem (Eq. (14)) is subject to the constraints given in Eq. (15) to (19), where π^P is the hour-ahead price of power and π^R is the hour ahead price of reserves, both in \$/kWh, and c is the penalty coefficient on the second moment of the tracking error. As in literature, we assume $\pi^P \approx \pi^R$ [14].

Constraints 1 and 2 limit the upper and lower bound of the bidding. These limits are a function of the number of active servers (Eq. (15) and (16)). Upper limit (a) is computed based on the maximum datacenter power, given the number of turned-on servers (N_s), the per-server maximum power (P_s^{max}), and the amount of power that can be injected into the battery from the power grid. Lower bound (b) shows the excess power that cannot be provided by renewable and battery sources. The tracking error is measured during the hour by Constraint 3. Part of the credit (i.e., $\pi^R R$) is reduced proportionally to the tracking error. The reserve provider may lose its contract in further RS reserves provision if

the tracking error exceeds a limit (i.e., ε). Finally, Constraint 4 and 5 ensure that (\bar{P}, R) do not take negative values.

5.1.2 Specific Solution for ECOGreen

The previous formulation describes the general cost function and constraints for our problem. In what follows, we present a specific formulation for the worst-case scenario (i.e., the case when $z(t) = 1$ or $z(t) = -1$), to avoid losing the contract due to a too large error in following the RS signal.

Within a time slot, the worst-case error is achieved when the following conditions are met: (i) datacenter aggregated workload is maximum ($\hat{P}_{DCT_{dyn}}^{max}$, i.e., the utilization peaks of VMs coincide at the same time); and (ii) renewable energy generation is at a minimum. In such scenario, b can be rewritten as:

$$b = \max(N_s \cdot P_{static} + \hat{P}_{DCT_{dyn}}^{max} - \min_t(\hat{P}_{PV}(t)) - P_{EES}^{dch}(t), 0) \quad (20)$$

Also, as in this case $z(t) = 1$ or $z(t) = -1$, Constraint 3 can be expressed as:

$$\begin{cases} (\bar{P} + z(t)R) - \hat{P}_{Grid}(t) \leq \varepsilon R \rightarrow \bar{P} + (1 - \varepsilon)R \leq a \\ \hat{P}_{Grid}(t) - (\bar{P} + z(t)R) \leq \varepsilon R \rightarrow \bar{P} - (1 - \varepsilon)R \geq b \end{cases} \quad (21)$$

Hence, Constraints 1, 2 and 3 can be simply replaced by these two new constraints. Therefore, the final problem can be considered as follows:

$$\min_{\bar{P}, R} Cost(t) = \pi^P \bar{P} - \pi^R R + \pi^R c \frac{|\hat{P}_{Grid}(t) - (\bar{P} + z(t)R)|}{R} \quad (22)$$

Subject to

$$1. \bar{P} + (1 - \varepsilon)R \leq a \quad (23)$$

$$2. \bar{P} - (1 - \varepsilon)R \geq b \quad (24)$$

$$3. \bar{P} \geq 0 \quad (25)$$

$$4. R \geq 0 \quad (26)$$

5.1.3 Solving The Worst-Case Scenario

We solve the previous minimization by using the derivative method, in order to obtain the best solution for the whole time slot (integral of the objective function, i.e., Eq. (22), over 1 hour). In this problem, the RS signal ($z(t)$) is generated based on the ISO specified integral proportional filter of the area control error and frequency flow out of a tolerance, thus it is unpredictable [10]. However, while the RS signal is not fully predictable, its stochastic properties are known and it can follow any distribution (e.g., normal, log-normal distribution, etc.) [35].

Therefore, we propose a general method to optimize the bidding values without any further specific knowledge in order to provide robustness against different signal values and distributions. Hence, we solve the problem for the worst-case scenario, which leads to two cases: (i) $z(t) = -1$, and (ii) $z(t) = 1$.

Worst-case scenario 1: $z(t) = -1$

When $z(t) = -1$ the absolute value of Eq. (22) is $\hat{P}_{Grid}(t) \geq \bar{P} + z(t)R$, $\forall t \in T$, and we can write:

$$\begin{aligned} \int_t^{t+T} Cost(t) dt &= \int_t^{t+T} (\pi^P \bar{P} - \pi^R R) dt + \int_t^{t+T} \pi^R c \left(\frac{\hat{P}_{Grid}(t)}{R} \right) dt \\ &\quad - \int_t^{t+T} \pi^R c \left(\frac{\bar{P}}{R} \right) dt - \int_t^{t+T} \pi^R c z(t) dt \end{aligned} \quad (27)$$

As $z(t)$ is the real-time broadcaster power market signal, and it takes a real number in the interval $[-1, 1]$ with an average of zero over longer time intervals [14], we can compute the aggregated cost for the whole 1 hour time slot as:

$$Cost^{agr} = \pi^P \bar{P}T - \pi^R RT + \pi^R c \left(\frac{\hat{P}_{Grid}^{agr}}{R} \right) - \pi^R c \left(\frac{\bar{P}}{R} \right) T \quad (28)$$

To find the solution, we first compute the first and second derivatives from \bar{P} and R as follows:

$$\begin{cases} \frac{\partial Cost^{agr}}{\partial \bar{P}} = \pi^P T - \frac{\pi^R c T}{R} = 0 \rightarrow R = \frac{\pi^R c}{\pi^P} \\ \frac{\partial^2 Cost^{agr}}{\partial \bar{P}^2} = 0 \\ \frac{\partial Cost^{agr}}{\partial R} = -\pi^R T - \left(\frac{\pi^R c \hat{P}_{Grid}^{agr} - \bar{P} \pi^R c T}{R^2} \right) = 0 \rightarrow \\ R = \left[c \left(\bar{P} - \frac{\hat{P}_{Grid}^{agr}}{T} \right) \right]^{\frac{1}{2}} \rightarrow \bar{P} = c \left(\frac{\pi^R}{\pi^P} \right)^2 + \frac{\hat{P}_{Grid}^{agr}}{T} \\ \frac{\partial^2 Cost^{agr}}{\partial \bar{P}^2} = \frac{2\pi^R c \hat{P}_{Grid}^{agr} - \bar{P} \pi^R c T}{R^3} \end{cases} \quad (29)$$

In this case, (\bar{P}, R) is a saddle point, since $\frac{\partial^2 Cost^{agr}}{\partial \bar{P}^2} \frac{\partial^2 Cost^{agr}}{\partial R^2} - \left(\frac{\partial^2 Cost^{agr}}{\partial \bar{P} \partial R} \right)^2 < 0$. Therefore, we also need to check the boundaries to find the solution under the worst-case scenario and assumptions. According to the defined constraints, there are four different explicit solutions considering these four boundaries: 1) $\bar{P} = b + (1 - \varepsilon)R$, 2) $\bar{P} = a - (1 - \varepsilon)R$, 3) $R = 0^+$, and 4) the intersection point of boundaries 1 and 2, as follows:

Case 1: $\bar{P} = b + (1 - \varepsilon)R$

$$\begin{cases} R = \left[\frac{\pi^R c \left(\frac{\hat{P}_{Grid}^{agr}}{T} - b \right)}{(1 - \varepsilon)\pi^P - \pi^R} \right]^{\frac{1}{2}} \\ \bar{P} = b + (1 - \varepsilon) \left[\frac{\pi^R c \left(\frac{\hat{P}_{Grid}^{agr}}{T} - b \right)}{(1 - \varepsilon)\pi^P - \pi^R} \right]^{\frac{1}{2}} \end{cases} \quad (30)$$

Case 2: $\bar{P} = a - (1 - \varepsilon)R$

$$\begin{cases} R = \left[\frac{\pi^R c \left(a - \frac{\hat{P}_{Grid}^{agr}}{T} \right)}{(1 - \varepsilon)\pi^P + \pi^R} \right]^{\frac{1}{2}} \\ \bar{P} = a - (1 - \varepsilon) \left[\frac{\pi^R c \left(a - \frac{\hat{P}_{Grid}^{agr}}{T} \right)}{(1 - \varepsilon)\pi^P + \pi^R} \right]^{\frac{1}{2}} \end{cases} \quad (31)$$

Case 3: $R = 0^+$

$$\begin{cases} \bar{P} = a \\ R = 0 \end{cases} \quad (32)$$

Case 4: Finally for the intersection point of Cases 1 and 2:

$$\begin{cases} R = \frac{a-b}{2(1-\varepsilon)} \\ \bar{P} = \frac{a+b}{2} \end{cases} \quad (33)$$

Worst-case scenario 2: $z(t) = 1$

With the same approach, we also solve the problem when $z(t) = 1$ and $\hat{P}_{Grid}(t) < \bar{P} + z(t)R$, $\forall t \in T$. In this case, three additional solutions are obtained under the aforementioned boundaries (the intersection point is the same for both situations). As in general $\pi^P \approx \pi^R$, Case 4 (Eq. (33)) always provides the best solution, for both conditions of the absolute value. Nevertheless, this problem and solution are still valid for the case where the $\pi^P \not\approx \pi^R$. However, under these circumstances, we should compute all 7 solutions for all the possible number of active servers, as Case 4 is not guaranteed to be the best.

Algorithm 1 Bidding - Find (\bar{P}, R) and N_s

Input: $P_{EES}^{ch}(t)$, $P_{EES}^{dch}(t)$, $\min_t(\hat{P}_{PV}(t))$, π^R , and π^R
Output: (\bar{P}, R) and N_s

```

1:  $\varepsilon \leftarrow 0$  (Find best solution without error)
2:  $Cost_{min}^{Tot} \leftarrow$  Maximum real value
3: for  $i = 1 : N_s^{DCT}$  do
4:    $\hat{P}_{Grid}(t) \leftarrow$  Compute  $\max_t(\hat{P}_{DCT}(t))$  and  $\min_t(\hat{P}_{DCT}(t))$  for two
   situations separately
5:   Compute upper ( $a$ ) and lower ( $b$ ) bounds
6:    $Cost_{min} \leftarrow$  Minimum cost among the solutions of two situations
   (i.e., Case 4, Eq. (33) when  $\pi^P \approx \pi^R$ )
7:   if  $Cost_{min} < Cost_{min}^{Tot}$  then
8:      $Cost_{min}^{Tot} \leftarrow Cost_{min}$ 
9:     Update  $(\bar{P}, R)$ 
10:     $N_s \leftarrow i$ 
11:   end if
12: end for

```

5.1.4 Jointly Selecting \bar{P} , R and N_s

\bar{P} and R are functions of N_s (i.e., the number of active servers determined in the VM allocation phase).

To select N_s , we first compute $P_{EES}^{ch}(t)$, $P_{EES}^{dch}(t)$, and the minimum renewable energy ($\min_t(\hat{P}_{PV}(t))$) during the next time slot. For $z(t) = -1$ (Eq. (28)), we consider $\frac{\hat{P}_{Grid}^{agr}}{T}$ as the maximum predicted datacenter power (i.e., $\max_t(\hat{P}_{DCT}(t))$) during the time slot. Similarly, when $z(t) = 1$ we consider $\min_t(\hat{P}_{DCT}(t))$ as the minimum predicted power, which we dynamically compute based on the different number of active servers and dynamic power. This assumption provides two benefits:

- Following the RS signal with better reliability and lower error due to the capability of masking the prediction error on the VMs workloads and renewable energy.
- Obtaining less QoS degradation using battery and renewable energy sources for the following regulation signal.

As the datacenter provider may lose its contract if the tracking error exceeds a limit, we should compute the seven solutions in the worst cases, assuming ε zero, to avoid maximizing R by increasing the error (see Eq. (21)). By doing so, we can control the tracking error even when the prediction error on the VMs workloads and renewable energy is high (i.e., during abrupt changes). As stated before, since we assume $\pi^P \approx \pi^R$, Case 4 is always the best solution, and we only need to iterate on the number of active servers, from 1 to N_s^{DCT} . Therefore, for each number of servers we compute the estimated power taken from the grid ($\hat{P}_{Grid}(t)$) for different a and b values, together with the associated cost. Finally, we choose the (\bar{P}, R) and N_s values that minimize cost while satisfying the constraints, as shown in Algorithm 1. We conclude that the proposed method has a time complexity of $O(N_s^{DCT})$.

5.2 Workload Allocation

After finding the number of turned-on servers (i.e., N_s) we define the same server cap (\hat{C}_{server}) for all active servers, as follows:

$$\hat{C}_{server} = \max_t \left(\sum_{k=1}^{N_{VM}} VMcpu_{k,t}^T / N_s \right) \quad (34)$$

where N_{VM} is the total number of VMs in the datacenter. Matrix $VMcpu_{k,t}^T$ contains the predicted k^{th} VM's CPU utilization at time t during the T^{th} time slot. Setting the same cap for all active servers allows a better control on the QoS. This is because server overutilization can only occur due to under-predictions on the VM usage, when the VMs require more CPU resources than predicted.

To allocate VMs to servers, we use a state-of-the-art correlation-aware VM allocation method [28]. Correlation refers to the similarity of VMs CPU utilization traces and the coincidence of their peaks. At the beginning of each time slot, the correlation between any two VMs is updated based on the history. Therefore, the time complexity is $O(N_{VM}^2)$. In this algorithm, the VMs are allocated to servers such that the correlation among the allocated VMs in the corresponding server is minimized, while the server does not exceed its defined cap (\hat{C}_{server}) with a time complexity of $O(N_s \cdot N_{VM}^2)$. This favors consolidation and leads to allocating VMs to the determined number of active servers (N_s). In this regard, in the worst case, the computational complexity is $O(N_{VM}^2 + N_s \cdot N_{VM}^2)$. In conclusion, the time complexity of the algorithm is $O(N_s \cdot N_{VM}^2)$. This correlation-aware VM allocation method is periodically invoked at every time slot T .

5.3 Revising Average Power (\bar{P}) and Reserves (R)

The previous step tries to map the VMs (both newly arrived and already running VMs on the system) to servers considering the servers cap. However, due to the allocation error (i.e., potentially allocating VMs to a higher number of servers than N_s), we need to revise (\bar{P}, R). Therefore, after computing the initial allocation and determining N_s but before sending the reserve value to ISO, we estimate the power consumption of datacenter using Eq. (7) based on the predicted VMs CPU and memory loads. Then, we calculate $\hat{P}_{Grid}(t)$ using Eq. (13) with respect to current battery status and predicted renewable power available in time slot T . Finally, we update (\bar{P}, R) once using the solution with the minimum cost that satisfies the constraints.

5.4 Online Regulation Signal (RS) Tracking

In this section, we describe our methodology for dynamically modulating the datacenter, battery and renewable power consumption to track the ISO RS signal (i.e., $P_{Grid}(t) = \bar{P} + z(t)R$) for a given (\bar{P}, R) with the smallest error possible. Our proposed method receives three inputs: 1) (\bar{P}, R) and $z(t)$ signals, 2) the real VM's CPU utilization and memory footprint, and 3) the current available renewable power. The output of our method is the CPU resource limit per VM, charge/discharge of the battery, and the renewable power usage, required to meet the QoS requirements.

Before formulating the online policy, we consider the following additional assumptions on the EES system:

- We set a limit on the battery discharge current, forcing it to be less than the maximum value (i.e., I_{dch}^{max}). We also assume that battery can be used to power the datacenter.
- The EES can be charged (P_{EES}^{ch}) using the power grid when the datacenter power consumption is less than the power provided by the grid (i.e., P_{Grid}). In this case, we define SoC^{lim} as the maximum battery level that can be reached by charging it using renewable power. From SoC^{lim} to full capacity, the battery is charged using the grid power.
- To charge the battery using the grid, we define a charge current limit (I_{ch}^{lim}) as follows:

$$\begin{cases} I_{ch}^{lim} = I_{ch}^{max} & P_{EES}^{ch} \geq I_{ch}^{max} \cdot V_{EES} \\ I_{ch}^{lim} = \frac{P_{EES}^{ch}}{V_{EES}} & otherwise \end{cases} \quad (35)$$

where V_{EES} indicates the battery voltage level. This guarantees that P_{EES}^{ch} is used to react to the ISO 4-second RS signal fluctuations.

In our method, as shown in Algorithm 2, we first compute the real datacenter power consumption ($P_{DCT}(t)$) using the real VMs CPU utilization and memory footprint. If the datacenter power consumption is higher than the provided grid power (lines 4–6), i.e., datacenter power dominance, we rely on Algorithm 3. Otherwise, if the datacenter power consumption is lower than grid power (lines 6–8), we use Algorithm 4.

5.4.1 Datacenter Power Consumption Dominancy

If $P_{DCT}(t) > P_{Grid}(t)$, as shown in Algorithm 3, we optimize the current renewable and battery power usage to compensate the

Algorithm 2 Online RS Tracking Policy (every 4 seconds)

Input: $z(t)$, SoC^{lim} , and I_{ch}^{lim}
Output: Following $P_{Grid}(t)$ with minimum error and computing final power taken from the grid

- 1: Update VMs CPU utilization and memory footprint
- 2: Update servers utilization (U_{cpu} & U_{mem})
- 3: Update datacenter power consumption ($P_{DCT}(t)$)
- 4: **if** $P_{DCT}(t) > P_{Grid}(t)$ **then**
- 5: **Call Algorithm 3**
 /* when the datacenter power consumption is higher than the provided grid power: the rest of power of datacenter should be supplied by managing PV and EES */
- 6: **else if** $P_{DCT}(t) < P_{Grid}(t)$ **then**
- 7: **Call Algorithm 4**
 /* when the datacenter power consumption is lower: storing grid power in the EES and increasing datacenter power consumption*/
- 8: **end if**

Algorithm 3 Tracking Policy - Datacenter Power Dominancy

- 1: [$P_{PV}(t)$, $P_{EES}(t)$, P_{rem}] \leftarrow GreenController(**S1**, $P_{DCT}(t)$, $P_{Grid}(t)$, SoC^{lim} , I_{ch}^{lim})
- 2: $\alpha \leftarrow +1$
- 3: **if** $P_{rem} > 0$ **then**
- 4: $QoS_{deg} \leftarrow P_{DCT\ dyn} / (P_{DCT\ dyn} - P_{rem})$
- 5: **if** $QoS_{deg} < QoS_{lim}$ **then**
- 6: $C_j^{update} \leftarrow (1/QoS_{deg}) \cdot U_{cpu_j} \quad \forall j \in 1 \dots N_s$
- 7: **else**
- 8: $C_j^{update} \leftarrow (1/QoS_{lim}) \cdot U_{cpu_j} \quad \forall j \in 1 \dots N_s$
- 9: **end if**
- 10: Update P_{rem}
- 11: $P_{Grid}(t) \leftarrow P_{Grid}(t) + P_{rem}$
- 12: **end if**

Algorithm 4 Tracking Policy - Grid Power Dominancy

- 1: [$P_{PV}(t)$, $P_{EES}(t)$, P_{rem}] \leftarrow GreenController(**S2**, $P_{DCT}(t)$, $P_{Grid}(t)$, SoC^{lim} , I_{ch}^{lim})
- 2: $\alpha \leftarrow -1$
- 3: **if** $P_{rem} > 0$ **then**
- 4: **for** $j = 1 : N_s$ **do**
- 5: $P_{dyn_j}^{rem} \leftarrow (100 - U_{cpu_j}) \cdot P_{cpu}^{dyn}$
- 6: **if** $P_{rem} < P_{dyn_j}^{rem}$ **then**
- 7: $C_j^{update} \leftarrow C_j + (P_{rem} \cdot 100) / P_{cpu}^{dyn}$
- 8: Update P_{rem}
- 9: **break**
- 10: **else**
- 11: $C_j^{update} \leftarrow 100$
- 12: $P_{rem} \leftarrow P_{rem} - P_{dyn_j}^{rem}$
- 13: **end if**
- 14: **end for**
- 15: $P_{Grid}(t) \leftarrow P_{Grid}(t) - P_{rem}$
- 16: **end if**

excess power consumed by the datacenter (i.e., $P_{DCT}(t) - P_{Grid}(t)$) using the GreenController function (line 1). After using green energy, if the datacenter power is not completely provisioned (i.e., $P_{rem} > 0$), we reduce the VMs CPU resources equally on the servers to meet the available power with minimum tracking error, allowing QoS degradation (QoS_{deg}) until reaching its limit (lines 3–12). For our virtualized applications, the QoS constraints are defined in terms of the maximum allowable degradation (i.e., increase in their execution time), which in our case is defined as $2x(QoS_{lim})$ [55], with respect to their baseline execution time. When the QoS constraint is tight and the VMs resources cannot be further reduced, we can additionally discharge the battery (even going below the DoD) to provision the extra power needed by the datacenter. Finally, we update $P_{Grid}(t)$ based on the remaining energy (line 11).

5.4.2 Grid Power Dominancy

If $P_{DCT}(t) < P_{Grid}(t)$, as shown in Algorithm 4, we attempt to charge the battery with the excess power (i.e., $P_{Grid}(t) - P_{DCT}(t)$) provided by the power grid using the GreenController function (line 1). Our priority is to fill the battery when the VMs meet the QoS limit. After doing so, if excess power still exists ($P_{rem} > 0$), we increase the resource limit of the VMs equally on the servers, one by one, until meeting the ISO power constraint with the minimum tracking error (lines 3–16), or until no more power can be used or stored.

5.4.3 Proposed Green Controller

To optimize the usage of renewable sources and battery power (charge/discharge), we specify different linear and nonlinear constraints and objective functions for different situations (i.e., **S1** and **S2** as shown in Algorithm 2). For all situations, we optimize vector x , which consists of SoC , I_b , I_b^{CTI} , I_{Rem}^{CTI} , I_{PV}^{CTI} , I_{DCT}^{CTI} , I_{Grid}^{CTI} , and I_{PV}^{waste} , respectively, for the defined system model in Sections 4 and 3. lb and ub indicate a set of lower and upper bounds on the design variables in x (i.e., $lb \leq x \leq up$).

S1-Discharge: this state occurs when $P_{PV}^{CTI}(t) \cdot \eta_{DCAC}(\rho(t)) \leq P_{DCT}(t) - P_{Grid}(t)$ and battery needs to be discharged ($\alpha \leftarrow +1$). We define lb and ub on x as follows:

$$lb = [DoD, 0, 0, 0, 0, 0, 0, 0]$$

$$ub = [1, I_{dch}^{max}, I_{CTI}^{max}, I_{CTI}^{max}, I_{CTI}^{max}, I_{CTI}^{max}, I_{CTI}^{max}, 0]$$

where I_{CTI}^{max} and I_{dch}^{max} denotes the maximum allowable current of the CTI and maximum discharge current of the battery, respectively. We indicate the linear equalities on x as $A_{eq} \cdot x = b_{eq}$, to solve the power management model at the CTI level. We simply name converters efficiency as η since this value is the same for all converters.

$$\begin{bmatrix} 0 & 0 & 1 & 1 & 1 & -1 & 1 & 0 \\ 0 & -V_{EES} \cdot \eta & V_{CTI} & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & V_{CTI} & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & V_{CTI} & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & V_{CTI} & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} x = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ P_{PV}^{CTI} \\ P_{DCT}^{CTI} \\ P_{Grid}^{CTI} \\ 0 \end{bmatrix}$$

Our strategy aims at minimizing the following function under the aforementioned constraints, in order to maximize the EES and renewable usage:

$$\min F = x(4)^2 + (ub(2) - x(2))^2 \quad (36)$$

S1-Charge: takes place when $P_{PV}^{CTI}(t) \cdot \eta_{DCAC}(\rho(t)) > P_{DCT}(t) - P_{Grid}(t)$ and battery can be charged ($\alpha \leftarrow -1$) using renewable energy. Therefore, lb and ub are as follows:

$$lb = [SoC^{cur}, -I_{ch}^{max}, -I_{CTI}^{max}, 0, 0, 0, 0, 0]$$

$$ub = [SoC^{lim}, 0, 0, 0, I_{CTI}^{max}, I_{CTI}^{max}, I_{CTI}^{max}, I_{CTI}^{max}]$$

where SoC^{cur} indicates the current SoC of the EES. We represent the linear equalities on x as (changes with respect to S1-Discharge are bolded):

$$\begin{bmatrix} 0 & 0 & 1 & \mathbf{0} & 1 & -1 & 1 & -1 \\ 0 & -V_{EES} & V_{CTI} \cdot \eta & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & \mathbf{0} & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & V_{CTI} & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & V_{CTI} & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} x = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ P_{DCT}^{CTI} \\ P_{Grid}^{CTI} \\ 0 \end{bmatrix}$$

Furthermore, the following linear inequality stands when renewable power cannot be fully used:

$$V_{CTI} \cdot I_{PV}^{CTI} \leq P_{PV}^{CTI} \quad (37)$$

We also minimize the following function under the aforementioned constraints:

$$\min F = -x(5)^2 + (lb(2) - x(2))^2 \quad (38)$$

S2-Charge: takes place when $P_{Grid}(t) - P_{DCT}(t) > 0$ and the generated grid power is higher than datacenter power consumption. Therefore, battery can be charged using the power grid ($\alpha \leftarrow -1$) and lb and ub are defined as follows:

$$lb = [SoC^{cur}, -I_{ch}^{lim}, -I_{CTI}^{max}, 0, 0, 0, 0, 0]$$

$$ub = [1, 0, 0, I_{CTI}^{max}, I_{CTI}^{max}, I_{CTI}^{max}, I_{CTI}^{max}, I_{CTI}^{max}]$$

We represent the linear equalities on x as (changes with respect to S1-Discharge are bolded):

$$\begin{bmatrix} 0 & 0 & 1 & -1 & \mathbf{0} & -1 & 1 & 0 \\ 0 & -V_{EES} & V_{CTI} \cdot \eta & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & \mathbf{0} & 0 & 0 & V_{CTI} \\ 0 & 0 & 0 & 0 & 0 & V_{CTI} & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \mathbf{0} & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & -1 \end{bmatrix} x = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ P_{PV}^{CTI} \\ P_{DCT}^{CTI} \\ 0 \end{bmatrix}$$

We also show the linear inequality on x when grid power is not completely used to supply the datacenter and charge the battery:

$$V_{CTI} \cdot I_{Grid}^{CTI} \leq P_{Grid} \cdot \eta \quad (39)$$

The minimization function is as follows:

$$\min F = -x(7)^2 + (lb(2) - x(2))^2 \quad (40)$$

Finally, we define a nonlinear equality ($C_{eq}(x) = 0$) according to Eq. (3) and (4) used for all the situations to compute the SoC considering the available charge capacity impacted by ageing as:

$$C_{eq} = [x(1) - \frac{C_{ref} \cdot SoC - I_{eq} \cdot dt}{C_{ref}}] \quad (41)$$

$$\text{where } I_{eq} = \left(\frac{|x(2)|}{I_{ref}} \right)^{(k_b-1)} \cdot x(2)$$

To solve the optimization problem, we utilize the `fmincon` function, which is part of a nonlinear programming solver, to find the minimum of the constrained nonlinear multivariable problem. Moreover, in our problem formulation, this method is a function of




Fig. 3: Forecasted vs. real PV power profile.

the number of power sources, i.e., PV, EES, Grid, and datacenter as a consumer and it is not a function of the datacenter size. Therefore, our proposed algorithm needs to be sufficiently fast to make decisions on renewable and battery (charge/discharge) usage to track the RS signal. As a result, the worst-case execution time of our method is 0.0635 seconds, carried out on a separate server equipped with a 24-core Intel CPU@1.60GHz and 50GB of memory.

6 EXPERIMENTAL SETUP AND SCENARIOS

In this section we present the experimental setup and introduce two scenarios to compare the proposed strategy, i.e., ECOGreen.

6.1 Experimental Setup

6.1.1 Green Datacenter Configuration

We model 30 racks, each rack with 10 Intel S2600GZ servers equipped with a 6-core CPU (Intel E5-2620) and 32GB of memory (RAM). The server power consumption is modeled as in Section 4.3. Each server consumes constant 16W and 27.2W for disk (P_{disk}) and cooling (P_{fan}) power, respectively. We consider a high fan speed (8000rpm) and a low inlet temperature (22°C) to reduce the effect of temperature-dependent leakage power. Under this condition, leakage power (P_{cpu}^{leak}) is almost constant and 3.1W in the worst case. Idle power for CPU (P_{cpu}^{idle}) and memory (P_{mem}^{idle}) are 50W and 4W, respectively. The CPU and memory dynamic power (P_{cpu}^{dyn} and P_{mem}^{dyn}) range from 0 to 42.5W and 0 to 56W, respectively (with their maximum values occurring at 100% utilization) [49].

6.1.2 Simulation Framework

In order to consider realistic CPU and memory usage traces, we use one week of traces of Google Cluster [56], which provides the CPU and memory utilization for VMs every 5 minutes (memory utilization is varying in the range of 2% to 32%). Arrival and total time (life time) of each VM, given in time slots, are generated by poisson and exponential distributions, respectively. VM allocation and power market bidding are invoked every hour, and the online RS tracking policy is invoked every 4 seconds. We solve the optimization problem using the following values: $\pi^P = \pi^R = 0.1$ $\$/kWh$, and $c = 1$ based on typical values of today's markets [37]. We use typical trajectories of $z(t)$ from PJM historical data [36], for a time horizon of one week.

We tune the PV module size considering the peak power production (i.e., the number of cells and panels) defining it to 35kWp. We compute the irradiance forecasts implementing the algorithm presented in previous work [57]. Figure 3 shows the real

versus forecasted PV power traces for one week, indicating less than 5% error in the worst case for a short-time ahead prediction. We have also analyzed the error of PV power prediction for a whole year. The worst-case error is still less than 10% over all the days of the year even in winter with less irradiance. As a result, predicting a periodical traces for renewable energy is sufficiently accurate for a short-time ahead, which in our case is one hour. Moreover, in the simulations, we consider a lithium-ion EES with 96kWh capacity and 70% of DoD, while keeping the remaining capacity in case of outage.

6.2 Scenarios

6.2.1 Scenario I - Bidding and RS Signal Tracking Analysis

In this scenario, we evaluate the impact of using demand-side EES and renewable energy with and without datacenter participation in the RS program. We also investigate the effectiveness of our strategy compared to the state-of-the-art. These state-of-the-art approaches mostly focus on optimizing the operational cost of the datacenters. All approaches consist of an allocation policy, and a bidding plus online tracking policy, as shown in each method name, as follows:

- PSA+DynPow [15]: this work uses a server Power-State-Aware method (PSA) to estimate the \bar{P} and R , by determining the number of active servers for serving the workloads. Also, a Dynamic Power control policy (DynPow) is proposed to track the RS signal using power capping and power level adjustment using different server power modes. We adapt this method to our work by considering three power modes: turned-off, idle, and active.
- PSA+DynPow w/o Bid: PSA+DynPow without bidding and RS tracking. In this case, we assume $R = 0$ and \bar{P} as the average datacenter power consumption (i.e., $\bar{P} = \frac{\int_t^{t+T} P_{DCR}(t)dt}{T}$) taken from the grid during a time slot, as determined by the allocation policy.
- COAT+DynPow: for a further comparison with the aforementioned work [15], we consider the bidding and tracking solution methods proposed in this work [15] jointly with the COnsolidation-Aware allocaTION method (COAT) [28], as it is one of the best energy-aware VM consolidation strategies in the state-of-the-art.
- COAT+DynPow w/o Bid: COAT+DynPow without bidding and RS tracking, to evaluate the impact of allocation.
- ECOGreen: our proposed Electricity Cost Optimization for Green datacenters.
- ECOGreen w/o Bid: ECOGreen without bidding and RS tracking, optimizing green power usage to minimize the cost and reduce grid power (i.e., $\bar{P} = \frac{\int_t^{t+T} P_{Grid}(t)dt}{T}$).
- Oracle ECOGreen: ECOGreen with oracle prediction for the workloads and renewable energy. The Oracle ECOGreen assumes that, at the beginning of each time slot, all the VM characteristics and PV generation for the time slot are known.

6.2.2 Scenario II - Impact of Workload Allocation Methods

To evaluate and isolate the impact of VM allocation from the effect of bidding on online tracking policies, we also compare ECOGreen against different VM allocation policies, namely:

- COAT: COnsolidation-Aware allocaTION [28].




Fig. 4: Average power consumption (\bar{P}) for a time horizon of one week.

- LB: Load Balancing strategy that aims to spread VMs across servers, reaching an average server utilization close to 50%.
- ECOGreen: the proposed optimization strategy.
- ECOGreen w/o Green: ECOGreen without renewable and battery sources (green energy).

All the above-mentioned workload allocation methods are used in conjunction with our proposed bidding, online RS signal tracking, and green (renewable and EES) controller.

Finally, in order to evaluate the efficiency of our strategy under potential trade-offs between different objectives, we compute the euclidean distance for each method i in a normalized multi-dimensional space from the optimal values per each dimension (objective), i.e., vector \odot . This efficiency metric is given by Eq. (42), and the higher E is, the higher the efficiency of the method:

$$E_i = 1 - \|\mathbb{M}_i - \odot\|_2 \quad (42)$$

7 EXPERIMENTAL RESULTS

7.1 Scenario I - Bidding and RS Signal Tracking Analysis

In this section we compare our strategy, i.e., ECOGreen, in terms of bidding (\bar{P} and R), monetary cost, and QoS against different state-of-the-art approaches, introduced in Section 6.2.1.

7.1.1 Bidding (\bar{P}, R) Analysis

The (\bar{P}, R) values of different approaches for a time horizon of one week are shown in Fig. 4 and 5, respectively. Approaches w/o Bid are not shown in Fig. 5, as $R = 0$ and they only optimize the power consumption of the datacenter. Due to the nature of consolidation, which packs VMs into the minimum number of servers, COAT+DynPow provides lower \bar{P} , but also lower R , as it has less slack to dynamically change the server resources and meet VMs requirements. On the contrary, PSA+DynPow provisions higher R due to the larger power range achievable. This is achieved at the expense of higher \bar{P} , as VMs are distributed among a more servers to avoid further QoS degradation. We observe that ECOGreen provides RS reserves (R) of 35% and 76% of \bar{P} in the worst case and average, respectively, over one week. In the best case (i.e., when available renewable energy is high and the load of the datacenter is low), ECOGreen provides 100% of \bar{P} as R , drastically reducing cost when compared to other approaches. With the same reason, Oracle ECOGreen and ECOGreen reach almost similar results due to using the same strategy. In this sense,




Fig. 5: The amount of reserves (R) for a time horizon of one week.

the best solution to minimize electricity cost is not necessarily trying to achieve the lowest \bar{P} or the highest R , as the best bidding varies depending on the availability of demand-side renewable energy and the EES system status.

Also, these figures illustrate that approaches without bidding try to minimize total datacenter power consumption in order to reduce the electricity cost. However, the capacity reserves offered by the power grid bring the opportunity to reduce monetary costs even when increasing the average power consumption, as shown in the next section, due to the high credit obtained for the reserves.

7.1.2 Monetary Cost

The monetary cost in the RS provision case is calculated using the objective function defined in Eq. (22). In this calculation, installation costs of demand-side EES and renewables equipment are not taken into account for all approaches. This is because even if this installation has non-negligible costs, in this paper we focus on datacenter operational expenditure (OPEX) reduction, not on capital expenses. Without RS provision (i.e., w/o Bid), the monetary costs are calculated based solely on the power consumed by the datacenter during the time slot, which is $\bar{P} \cdot \pi^P \cdot T$, where $\bar{P} = \frac{\int_t^{t+T} P_{DCR}(t) dt}{T}$. Prior work shows that the peak power cost encompasses around 30% of the total electricity bill [34]. Therefore, we only focus on the energy consumption cost for comparison. However, with participation in the power market, the datacenter providers do not pay for the maximum power of the month and the cost function is different (see Eq. (22)).

To better evaluate the savings, we normalize the total monetary cost of each method to the largest value among all the methods, as shown in Fig. 6. All in all, ECOGreen uses the PV and EES power to optimize the bidding values and monetary costs, while ECOGreen w/o Bid only optimizes the green power usage for the datacenter. In contrast to COAT+DynPow and PSA+DynPow, ECOGreen tries to match \bar{P} and R , enabling the green datacenter to provide reserves close to its average power consumption. Given that the price for average power is the same than the credit obtained for the reserves (i.e., $\pi^P \approx \pi^R$), having similar values for \bar{P} and R yields the lowest cost, reduces the tracking error and efficiently utilizes renewable and EES power.

In addition, the results of ECOGreen w/o Bid show how renewables and battery power sources can save monetary costs compared to COAT+DynPow and PSA+DynPow, even without bidding. Also, COAT+DynPow and PSA+DynPow provide better results than COAT+DynPow w/o Bid and PSA+DynPow w/o Bid but with higher QoS degradation in order to regulate the datacenter power consumption. As a consequence, ECOGreen reduces the power costs by 35%, 61%, and 71% in comparison to ECOGreen

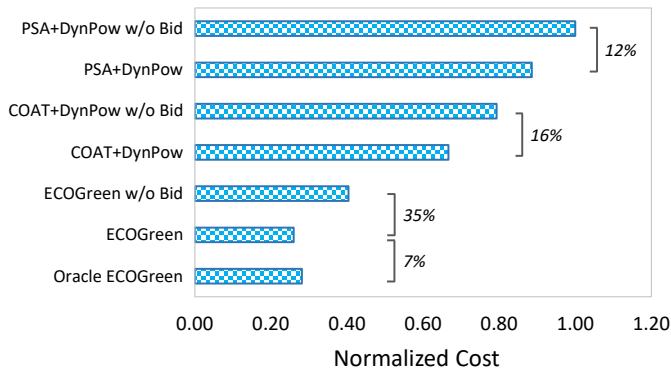


Fig. 6: Normalized monetary cost over a time horizon of one week.

w/o Bid, COAT+DynPow, and PSA+DynPow, respectively. Furthermore, ECOGreen achieves 7% cost improvement compared to Oracle ECOGreen because of better utilizing the battery energy, but at the expense of QoS degradation in case of miss-predictions.

7.1.3 Quality of Service (QoS)

For virtualized applications, the QoS constraint is defined in terms of the maximum allowable degradation (i.e., increase in execution time). In our case this limit is defined as $2x (QoS_{lim})$ [55], with respect to the VMs' baseline requirements. Fig. 7 shows the average degradation among all degraded VMs for a time horizon of one week. Degradation occurs due to miss-predictions on VMs workloads (especially during abrupt workload changes) and renewable energy production. Degradation increases when all power generation sources are not able to provide the required power consumption of the datacenter.

Overall, our experimental results indicate that all the approaches meet the QoS limit in the worst case. In fact, COAT+DynPow has less control on servers overutilization during peak loads, and the behaviour worsens due to the need for tracking the RS signal. To reduce QoS degradation, PSA+DynPow uses idle servers as QoS guarantee slack for immediately serving coming workloads. Due to the nature of Oracle ECOGreen, this method does not present any QoS degradation. Finally, ECOGreen reduces the average degradation compensating the tracking error using renewable and EES. ECOGreen w/o Bid and COAT+DynPow w/o Bid reach similar result due to using the same policy for allocating VMs to the servers. In conclusion, our strategy is able to meet the QoS degradation constraint for virtualized applications in RS reserves provision market, and obtains 19% and 10% less

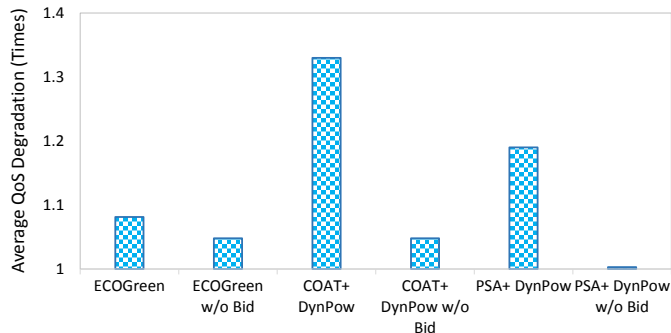


Fig. 7: Average QoS degradation.




Fig. 8: Normalized monetary cost over a time horizon of one week.

degradation on average than COAT+DynPow and PSA+DynPow, respectively.

7.2 Scenario II - Impact of Workload Allocation

In this section we compare our holistic strategy in terms of monetary cost, total green datacenter power consumption breakdown, EES efficiency, and potential trade-offs against the different state-of-the-art VM allocation policies introduced in Section 6.2.2.

7.2.1 Monetary Cost

Figure 8 shows the normalized monetary costs of the green datacenter for a time horizon of one week. Due to the nature of consolidation, COAT reduces the average power (\bar{P}) by co-allocating VMs into servers until reaching the maximum server capacity. This leads to reducing the RS reserves due to the decreased flexibility in increasing the CPU resources of each VM, as well as efficiently using the renewable energy. On the contrary, by distributing the VMs to a larger number of servers, LB provides higher reserves at the expense of an increase in the average power. Finally, ECOGreen jointly optimizes the bidding power market values and the number of active servers while considering the current state of EES and the predicted renewable power. Therefore, results show that the lowest average power or highest reserves do not provide, by themselves only, lower monetary cost in today's RS reserves provision. On the other hand, LB can still provide better savings compared to our proposed strategy when not using demand-side EES and renewable sources (i.e., LB vs. ECOGreen w/o Green) due to the optimization of green power usage.

In summary, ECOGreen obtains 7%, 53%, and 61% monetary savings compared to COAT, LB, and ECOGreen w/o Green, respectively.

7.2.2 Power Consumption Analysis

Figure 9 shows the total power consumption breakdown of the green datacenter for a time horizon of one week. As a consequence of maximizing renewable and battery power utilization, all approaches show an overall green power usage improvement higher than ECOGreen w/o Green. ECOGreen uses more grid power than COAT to provide better bidding values (i.e. higher RS reserves). ECOGreen also maximizes green energy usage, achieving 48% and 22% renewable and battery utilization improvement, respectively, compared to COAT. This situation is due to the larger reserves achieved by ECOGreen, which is a consequence of exploiting the green energy sources to track the RS signal. Furthermore, ECOGreen outperforms LB in terms of power consumption as in average it allocates VMs to a fewer number of

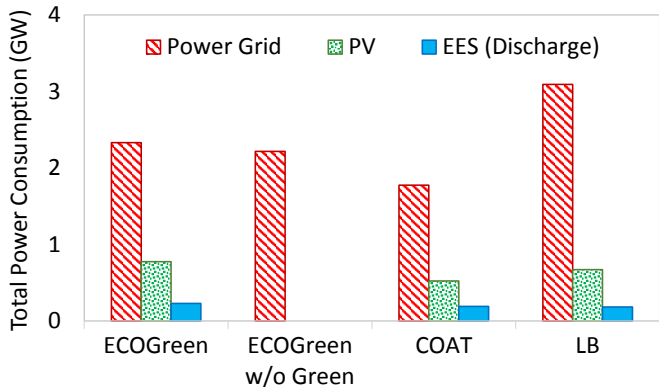


Fig. 9: Total power consumption breakdown of the green datacenter for different power supply sources for a time horizon of one week.

servers. In addition, because of its lower tracking error, the battery is not significantly discharged. We observe up to 16% and 28% PV and battery power utilization improvements for ECOGreen compared to LB, respectively.

7.2.3 EES Performance Analysis

Figure 10 shows a one-week view of the system evolution with battery charge/discharge cycles. We add some constraints on the allowed DoD for the battery bank. In particular, to force the battery bank to work in the optimal range of SoC, we set the minimum SoC to 70%. In this figure we observe how the energy buffer in RS reserves allows us to use battery power to optimize the bidding values and follow the RS signal in different situations. ECOGreen better utilizes the battery power than the other approaches, especially when compensating the excess power needed by the datacenter. However, this leads to decreasing the SoH of the battery, in particular during discharge cycles. As ECOGreen w/o Green does not have an EES system, this method is not shown in Fig. 10.

7.2.4 Performance Metrics Trade-offs Analysis

Our experimental results confirm that, by having a holistic strategy we can obtain better overall results by exploiting renewable and battery sources. Figures 11, 12 and 13 summarize the benefits of ECOGreen in comparison with other state-of-the-art techniques. First, Fig. 11 depicts the cost vs. green power sources trade-off, showing the best performance for ECOGreen in terms of both monetary cost and green power usage. Table 1 summarizes the




Fig. 10: State-of-the-charge (SoC) of battery bank for a time horizon of one week.

TABLE 1: Cost, grid, PV and EES (battery) usage improvements for ECOGreen compared to other approaches

	COAT	LB	ECOGreen w/o Green
Cost	7%	53%	61%
Grid Power	-29%	25%	-5%
PV Power	48%	16%	w/o PV
Battery Power	22%	28%	w/o EES

efficiency of ECOGreen compared to other methods in terms of power usage of different sources in RS reserves provision.

Figure 12 shows the normalized cost vs. QoS degradation trade-off, with ECOGreen providing 7%, 53%, and 61% improvements for cost compared to COAT, LB, and ECOGreen w/o Green, respectively. Our results imply that all the methods can meet the QoS limit, with ECOGreen w/o Green exhibiting the highest degradation (1.3x on average). This is because ECOGreen w/o Green cannot tolerate power reduction without QoS degradation to track the RS signal; while ECOGreen can use the renewable and EES to supply the additional power consumed by the datacenter. Compared to COAT and LB, it achieves less than 2% QoS degradation.

Finally, Fig. 13 depicts the cost vs. SoH trade-off: our algorithm results in better monetary cost at the expense of a slightly higher (0.03%) battery SoH decrease (ratio between nominal and remaining capacity, as computed in Eq. (1)). If we consider that batteries reach their end-of-life when working at 70% of their nominal capacity, ECOGreen enables a battery lifetime above 15 years.

7.2.5 Discussion - Efficiency of The Proposed Strategy

As in this work we deal with a multi-objective problem, we define an efficiency metric (Eq. (42)) to evaluate ECOGreen in a holistic way. The objectives considered (i.e., vector \mathbb{M} in the metric) are monetary cost, renewable usage, EES utilization, QoS and battery lifetime. Table 2 shows the normalized efficiency with respect to the worst-case value among all methods. As a result, in the considered green datacenter scenario, ECOGreen achieves the best overall performance (highest efficiency value) compared to other approaches in today's emerging power markets.

Real-world market services vary between countries due to the regulatory rules concerns [58]. Hence, there are several power markets and demand-response programs with different timescales, as well as several frequencies of request for power regulation. In particular, RS is one of the most widely extended power market in the US, which at the same time is a pioneering country in power regulation. Therefore, in this work, we focused on realistic pricing scenario and trajectories of RS signal from PJM (a large ISO in the US). In addition, RS reserve provision (hour-ahead markets) is the most suitable and profitable program for datacenters due to the high credit gain possibilities and datacenters' flexibility in regulating their power consumption. Thus, as datacenters are among the fastest growing electricity consumers, they are highly

TABLE 2: Overall efficiency for different methods

	ECOGreen	COAT	LB	ECOGreen w/o Green
Efficiency (E)	0.53	0.45	0.36	0

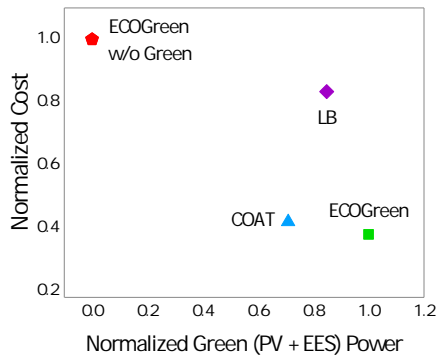


Fig. 11: Cost vs. green power trade-off.

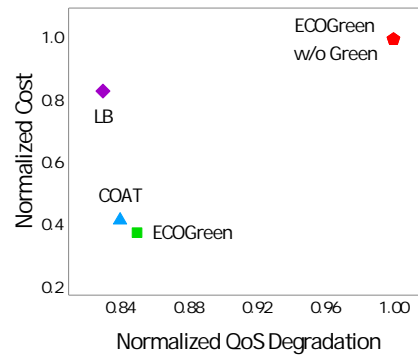


Fig. 12: Cost vs. QoS degradation trade-off.

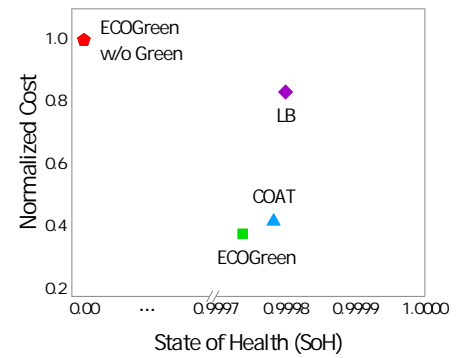


Fig. 13: Cost vs. SoH trade-off.

promising candidates to provide demand-side capacity reserves and reduce their electricity costs. From the point of view of datacenter providers, there is no competition as they only act as demand (i.e., consumers) to gain credits by making a contract for one of the demand response programs.

To optimize the green datacenters cost when participating in power markets RS reserves, we proposed in this work a strategy including Bidding Problem, Allocation, and Online Policy phases. Moreover, in order to simulate a realistic datacenter scenario, we utilized the real load traces for VMs, real sun irradiance and temperature profiles for PV generation, and general purpose model according to commercial devices for EES.

8 CONCLUSION

In this paper we have proposed ECOGreen, a novel strategy to tackle the challenge of green datacenter participation in RS reserves provision. We have first presented a fast analytical solution to jointly find the best average power and reserve bidding values, and decide on the required number of active servers in the datacenter, while optimizing the EES and renewable power usage. Moreover, we have proposed a runtime approach that dynamically regulates the datacenter power consumption following the RS signal, while also guaranteeing the QoS limit. Our runtime policy dynamically changes the server resources allocated to VMs, and utilizes EES and renewable power in decision making to minimize the signal tracking error. We have compared ECOGreen in terms of monetary cost, total power consumption breakdown, QoS degradation, and EES efficiency analysis against different state-of-the-art approaches. Our experimental results have shown that ECOGreen obtains up to 71%, 48%, and 28% monetary cost, renewable and EES utilization improvements, respectively, at the expense of a 0.03% battery SoH decrease when compared to other approaches. Nonetheless, ECOGreen enables battery lifetimes above 15 years. In summary, the holistic approach of ECOGreen ensures achieving the best trade-off between different objectives when compared to all other methods.

Our work can be applied for any datacenter running a wide range of applications found in private and public clouds. We cover time-varying workload setups ranging from virtualized banking applications to traditional scale-out applications, such as, web search or media streaming, in the context of dynamic datacenter requirements. On the other hand, a possible consideration for future work would be the use of stochastic methods to optimize the policy by predicting the RS signal.

ACKNOWLEDGMENTS

This work has been partially supported by the EC H2020 RECIPE FET-HPC project (Agreement No. 801137), the ERC Consolidator Grant COMPUSAPIEN (Agreement No. 725657), the EC H2020 AENEAS project (Agreement No. 731016), and the BU College of Engineering Dean's Catalyst Award. We would like to thank Yijia Zhang at Boston University for his comments that helped to improve this paper.

REFERENCES

- [1] J. Koomey, "Growth in data center electricity use 2005 to 2010," Analytics Press, Oakland, CA, Tech. Rep., 2011.
- [2] M. Dayarathna, Y. Wen, and R. Fan, "Data center energy consumption modeling: A survey," *IEEE Communications Surveys Tutorials*, vol. 18, no. 1, pp. 732–794, 2016.
- [3] Y. Zhang, Y. Wang, and X. Wang, "Greenware: Greening cloud-scale data centers to maximize the use of renewable energy," in *Springer Middleware*, 2011, pp. 143–164.
- [4] C. Stewart and K. Shen, "Some joules are more precious than others: Managing renewable energy in the datacenter," in *Workshop on Power Aware Computing and Systems*, 2009.
- [5] A. Pahlevan, P. G. D. Valle, and D. Aienza, "Exploiting cpu-load and data correlations in multi-objective vm placement for geo-distributed data centers," in *Design, Automation & Test in Europe Conference & Exhibition (DATE)*, March 2016, pp. 1333–1338.
- [6] C. Bhlinger, A. Lschel, U. Moslener, and T. F. Rutherford, "Eu climate policy up to 2020: An economic impact assessment," *Energy Economics*, vol. 31, pp. 295 – 305, 2009.
- [7] EIA, "Annual energy outlook," <http://www.eia.gov/forecasts/aeo>, 2014.
- [8] A. L. Ott, "Experience with pjm market operation, system design, and implementation," *IEEE Transactions on Power Systems*, vol. 18, no. 2, pp. 528–534, 2003.
- [9] NYISO, "Manual 2: Ancillary services manual, v3.26," <http://www.eia.gov/forecasts/aeo>, 2013.
- [10] H. Chen, Y. Zhang, M. Caramanis, and A. K. Coskun, "EnergyQARE: Qos-aware data center participation in smart grid regulation service reserve provision," *ACM Transactions on Modeling and Performance Evaluation of Computing Systems (TOMPECS)*, 2018.
- [11] D. Aikema, R. Simmonds, and H. Zareipour, "Data centres in the ancillary services market," in *IEEE International Green Computing Conference (IGCC)*, 2012, pp. 1–10.
- [12] H. Chen, A. K. Coskun, and M. C. Caramanis, "Real-time power control of data centers for providing regulation service," in *IEEE Conference on Decision and Control (CDC)*, 2013, pp. 4314 – 4321.
- [13] B. Kirpes and S. Klingert, "Evaluation process of demand response compensation models for data centers," in *ACM Proceedings of the International Workshop on Energy Efficient Data Centres*, 2016, pp. 1–6.
- [14] H. Chen, C. Hankendi, M. C. Caramanis, and A. K. Coskun, "Dynamic server power capping for enabling data center participation in power markets," in *IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, 2013, pp. 122–129.

- [15] H. Chen, M. C. Caramanis, and A. K. Coskun, "The data center as a grid load stabilizer," in *Asia and South Pacific Design Automation Conference (ASP-DAC)*, 2014, pp. 105–112.
- [16] M. Pawlish, A. S. Varde, and S. A. Robila, "Analyzing utilization rates in data centers for optimizing energy management," in *International Green Computing Conference (IGCC)*, 2012, pp. 1–6.
- [17] L. Liu, H. Wang, X. Liu, X. Jin, W. B. He, Q. B. Wang, and Y. Chen, "Greencloud: A new architecture for green data center," in *ACM International Conference Industry Session on Autonomic Computing and Communications Industry Session (ICAC-INDST)*, 2009, pp. 29–38.
- [18] L. Wang and S. U. Khan, "Review of performance metrics for green data centers: a taxonomy study," *The Journal of Supercomputing*, vol. 63, no. 3, pp. 639–656, 2013.
- [19] I. Goiri and et al., "Greenslot: Scheduling energy consumption in green datacenters," in *ACM SC*, 2011.
- [20] M. Ghamkhari and H. Mohsenian-Rad, "Energy and performance management of green data centers: A profit maximization approach," *IEEE Trans. on Smart Grid*, 2013.
- [21] X. Deng and et al., "Eco-aware online power management and load scheduling for green cloud datacenters," *IEEE Systems Journal*, 2014.
- [22] H. Farhangi, "The path of the smart grid," *IEEE Power and Energy Magazine*, vol. 8, no. 1, pp. 18–28, 2010.
- [23] H. Dou, Y. Qi, W. Wei, and H. Song, "Carbon-aware electricity cost minimization for sustainable data centers," *IEEE Transactions on Sustainable Computing (T-SUSC)*, vol. 2, no. 2, pp. 211–223, 2017.
- [24] H. Chen, Z. Liu, A. K. Coskun, and A. Wierman, "Optimizing energy storage participation in emerging power markets," in *International Green and Sustainable Computing Conference (IGSC)*, 2015, pp. 1–6.
- [25] M. Ghasemi-Gol, Y. Wang, and M. Pedram, "An optimization framework for data centers to minimize electric bill under day-ahead dynamic energy prices while providing regulation services," in *International Green Computing Conference (IGCC)*, 2014, pp. 1–9.
- [26] B. Aksanli and T. Rosing, "Providing regulation services and managing data center peak power budgets," in *Design Automation Test in Europe Conference Exhibition (DATE)*, 2014, pp. 1–4.
- [27] Z. Liu, I. Liu, S. Low, and A. Wierman, "Pricing data center demand response," *ACM SIGMETRICS Perform. Eval. Rev.*, vol. 42, no. 1, pp. 111–123, 2014.
- [28] J. Kim, M. Ruggiero, D. Atienza, and M. Lederberger, "Correlation-aware virtual machine allocation for energy-efficient datacenters," in *Design, Automation & Test in Europe Conference & Exhibition (DATE)*, 2013, pp. 1345–1350.
- [29] W. Lin, S. Xu, J. Li, L. Xu, and Z. Peng, "Design and theoretical analysis of virtual machine placement algorithm based on peak workload characteristics," *Springer Soft Computing*, vol. 21, pp. 1301–1314, 2017.
- [30] A. Pahlevan, X. Qu, M. Zapater, and D. Atienza, "Integrating heuristic and machine-learning methods for efficient virtual machine allocation in data centers," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 37, no. 8, pp. 1667–1680, 2018.
- [31] X. Ruan and H. Chen, "Performance-to-power ratio aware virtual machine (VM) allocation in energy-efficient clouds," in *IEEE International Conference on Cluster Computing (CLUSTER)*, 2015, pp. 264–273.
- [32] J. V. Wang, K. Y. Fok, C. T. Cheng, and C. K. Tse, "A stable matching-based virtual machine allocation mechanism for cloud data centers," in *IEEE World Congress on Services (SERVICES)*, 2016, pp. 103–106.
- [33] L. Chen and H. Shen, "Consolidating complementary VMs with spatial/temporal-awareness in cloud datacenters," in *IEEE Conference on Computer Communications*, 2014, pp. 1033–1041.
- [34] M. Zapater, A. Turk, J. M. Moya, J. L. Ayala, and A. K. Coskun, "Dynamic workload and cooling management in high-efficiency data centers," in *International Green and Sustainable Computing Conference (IGSC)*, 2015, pp. 1–8.
- [35] Y. Zhang, I. C. Paschalidis, and A. K. Coskun, "Data center participation in demand response programs with quality-of-service guarantees," in *Proceedings of the ACM International Conference on Future Energy Systems*, 2019, pp. 285–302.
- [36] PJM, "Description of regulation signals," www.pjm.com, 2014.
- [37] —, "Pjm manual 12: Balancing operations," www.pjm.com, 2018.
- [38] —, "Ftrs: protection against congestion charges," <https://learn.pjm.com/three-priorities/buying-and-selling-energy/ftf-faqs.aspx>, 2020.
- [39] G. E. P. Box, G. M. Jenkins, and G. C. Reinsel, *Time Series Analysis, Forecasting and Control*. Fourth Edition, Hoboken, NJ, USA: Wiley, 2013.
- [40] A. Pahlevan, M. Rossi, P. G. D. Valle, D. Brunelli, and D. Atienza, *Joint Computing and Electric Systems Optimization for Green Datacenters*. Springer Handbook of Hardware/Software Codesign, 2017.
- [41] Y. Riffonneau, S. Bacha, F. Barruel, and A. Delaille, "Energy flow management in grid connected pv systems with storage - a deterministic approach," in *IEEE International Conference on Industrial Technology*, Feb 2009, pp. 1–6.
- [42] Y. Riffonneau, S. Bacha, F. Barruel, and S. Ploix, "Optimal power flow management for grid connected pv systems with batteries," *IEEE Transactions on Sustainable Energy*, vol. 2, no. 3, pp. 309–320, 2011.
- [43] StarkPower UltraEnergy as lithium-ion battery, <http://www.starkpower.com/spnews/energystoragebatt>.
- [44] M. Rossi, A. Toppano, and D. Brunelli, "Real-time optimization of the battery banks lifetime in hybrid residential electrical systems," in *Design, Automation & Test in Europe Conference & Exhibition (DATE)*, March 2014, pp. 1–6.
- [45] PV device characteristics, <http://www.enfsolar.com/pv/cell-datasheet/429>.
- [46] Real sun irradiance, http://www.soda-is.com/eng/services/services_radiation_free_eng.php.
- [47] Temperature profiles, <http://www.tutiempo.net/en/Climate>.
- [48] N. Upadhyay, V. K. Singh, and S. Urooj, "Closed loop voltage control design for photovoltaic inverter," in *IEEE Conference Nanotechnology for Instrumentation and Measurement*, 2017.
- [49] J. C. Salinas-Hilburg, M. Zapater, J. L. Risco-Martín, J. M. Moya, and J. L. Ayala, "Unsupervised power modeling of co-allocated workloads for energy efficiency in data centers," in *Design, Automation & Test in Europe Conference & Exhibition (DATE)*, 2016, pp. 1345–1350.
- [50] M. Zapater, O. Tuncer, J. L. Ayala, J. M. Moya, K. Vaidyanathan, K. Gross, and A. K. Coskun, "Leakage-aware cooling management for improving server energy efficiency," *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, vol. 26, no. 10, pp. 2764–2777, 2015.
- [51] D. Kusic, J. O. Kephart, J. E. Hanson, N. Kandasamy, and G. Jiang, "Power and performance management of virtualized computing environments via lookahead control," in *International Conference on Autonomic Computing*, 2008, pp. 3–12.
- [52] N. Deng, C. Stewart, and J. Li, "Concentrating renewable energy in grid-tied datacenters," in *Proceedings of the IEEE International Symposium on Sustainable Systems and Technology (ISSST)*, 2011, pp. 1–6.
- [53] Converter efficiency, <http://www.schaeferpower.de/cms/en/produkte.html>.
- [54] A. Pahlevan, M. Momtazpour, and M. Goudarzi, "Power reduction in hpc data centers: a joint server placement and chassis consolidation approach," *Journal of Supercomputing*, vol. 70, pp. 845–879, 2014.
- [55] C. Delimitro and C. Kozyrakis, "Optimizing resource provisioning in shared cloud systems," Stanford University, Tech. Rep., 2014.
- [56] J. Wilkes, "More google cluster data," *Google research blog*, November, 2011.
- [57] C. Bergonzini, D. Brunelli, and L. Benini, "Comparison of energy intake prediction algorithms for systems powered by photovoltaic harvesters," *Elsevier Microelectronics Journal*, vol. 41, no. 11, pp. 766–777, 2010.
- [58] M. H. Shoreh, P. Siano, M. Shafie-khah, V. Loia, and J. P. S. Catalão, "A survey of industrial applications of demand response," *Electric Power Systems Research*, vol. 141, pp. 31 – 49, 2016.



Ali Pahlevan is currently a post-doctoral researcher in the Embedded Systems Laboratory (ESL) at Swiss Federal Institute of Technology Lausanne (EPFL). He received his Ph.D. degree in Electrical Engineering (EE) from EPFL in 2019, M.Sc. degree in Computer Engineering from Sharif University of Technology (SUT) in 2012, and B.Sc. degree in Computer Engineering from Ferdowsi University of Mashhad (FUM) in 2010. His research interests focus on system-level energy optimization techniques and reliability in the area of computing systems, datacenters, cloud and green computing. He has published over 15 research papers in top international journals and conferences such as the IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD), IEEE Transactions on Services Computing (TSC), Springer Handbook of Hardware/Software Codesign, and Design, Automation & Test in Europe Conference & Exhibition (DATE).

ability in the area of computing systems, datacenters, cloud and green computing. He has published over 15 research papers in top international journals and conferences such as the IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD), IEEE Transactions on Services Computing (TSC), Springer Handbook of Hardware/Software Codesign, and Design, Automation & Test in Europe Conference & Exhibition (DATE).



Marina Zapater is associate professor in the School of Engineering and Management of Vaud (HEIG-VD) at the University of Applied Sciences Western Switzerland (HES-SO) since 2020, and research associate in the Embedded System Laboratory (ESL) at the Swiss Federal Institute of Technology Lausanne (EPFL), Switzerland, since 2016. She received her Ph.D. degree in electronic engineering from Universidad Politécnica de Madrid, Spain, in 2015. Her research interests include thermal, power and performance design and optimization of complex heterogeneous architectures, from embedded edge devices to high-performance computing processors; and energy efficiency in servers and datacenters. In these fields, she has co-authored more than 50 papers in top-notch conferences and journals. She is an IEEE and CEDA member, and has served as CEDA YP representative (2019-2020).



Ayse K. Coskun (M'06) is an associate professor in the electrical and computer engineering department at Boston University (BU). She received her MS and Ph.D. degrees in Computer Science and Engineering from University of California, San Diego. Coskun's research interests broadly span energy and temperature awareness in computing systems, including novel architectures such as 3D-stacked systems, cloud and HPC datacenters, and mobile/embedded systems. Prof. Coskun worked at Sun Microsystems (now Oracle), San Diego prior to her current position at BU. Coskun is a recipient of the NSF CAREER award and the IEEE CEDA Early Career Award, and currently serves as an associate editor for IEEE Transactions on CAD.



David Atienza (M'05-SM'13-F'16) is associate professor of electrical and computer engineering, and director of the Embedded Systems Laboratory (ESL) at the Swiss Federal Institute of Technology Lausanne (EPFL), Switzerland. He received his PhD in computer science and engineering from UCM, Spain, and IMEC, Belgium, in 2005. His research interests include system-level design methodologies for high-performance multi-processor system-on-chip (MPSoC) and low power Internet-of-

Things (IoT) systems, including new 2-D/3-D thermal-aware design for MPSoCs and many-core servers, and edge AI architectures for wireless body sensor nodes and smart consumer devices. He is a co-author of more than 300 papers in peer-reviewed international journals and conferences, several book chapters, and nine patents. Dr. Atienza received the DAC Under-40 Innovators Award in 2018, IEEE TCCPS Mid-Career Award in 2018, an ERC Consolidator Grant in 2016, the IEEE CEDA Early Career Award in 2013, the ACM SIGDA Outstanding New Faculty Award in 2012, and a Faculty Award from Sun Labs at Oracle in 2011. He served as DATE 2015 Program Chair and DATE 2017 General Chair. He is an IEEE Fellow, an ACM Distinguished Member, and has served as IEEE CEDA President (period 2019-2020).