Demand response and ancillary services for supercomputing and datacenters

Sergio Nesmachnow, Santiago Iturriaga, Jonathan Muraña, Sebastián Montes de Oca, Gonzalo Belcredi, Pablo Monzón, Pablo Belzarena, and Juan Bazerque

Facultad de Ingeniería, Universidad de la República, Uruguay
{sergion,siturria,jmurana,smontes,gbelcredi,monzon,belza,jbazerque}@fing.edu.uy

Abstract. This article describes a proposal for the participation of supercomputing platforms and datacenters in the electric market, by implementing demand response techniques and ancillary services. Supercomputing and datacenters are appropriate candidates to adjust their power consumption in order to help the electric network to fulfill specific goals, either by consuming available surplus of energy to execute complex tasks, or by deferring activities when energy is more expensive or generation is lower than normal. Their thermal/cooling infrastructures demand about half of the energy consumption and provide a large inertia that can be carefully used to interact with the power grid. These strategies allow implementing a smart management of the electric grid, achieving a rational utilization of renewable energy sources, and the correct utilization of information technologies to improve decision-making processes. A specific case study is presented: The National Supercomputing Center in Uruguay (Cluster-UY), for which strategies for optimal planning of the execution of tasks and energy utilization are proposed, taking into account the energy consumption, the Quality of Service provided to the users, and the thermal/cooling demands of the infrastructure. In addition, the business opportunities and business models for supercomputing and datacenters in the electric market are revisited. Results suggest the effectiveness of the proposed strategies to implement demand response techniques and provide ancillary services under the smart grid paradigm.

Keywords: Energy efficiency · demand response · datacenters.

1 Introduction

In modern electricity markets, a large consumer with flexible consumption of active and reactive power can participate in the market in different ways. This concept is key to implementing strategies oriented to smart networks, associating consumers with the roles of active clients and market agents [18]. As an active client, the consumer can adapt his demand to peak hours, reducing consumption in these periods and contributing to flattening the demand curve of the system. Multi-hour tariffs can also be implemented, handling time blocks where it is preferable to consume. 'Day-ahead agreements' (based on price announced in advance) can be set, or even a dynamic behavior can be stimulated, when the price of energy is available in real time. Acting as an agent, the consumer can participate in the electricity market and receive income by applying mechanisms that may be restricted or driven by regulations, e.g., by establishing bilateral agreements between a large consumer and a generation company (possible in the Uruguayan energy market) or by auctions, e.g., in a day-ahead market, offering a profile of hourly consumption and establishing maximum prices to pay [13] (not yet present in our country).

In this context, demand response planning strategies are needed to manage energy consumption and be able to participate in the market, on different roles. Specific techniques are needed to dimension the activities that consume energy, advance or defer their execution, analyze the impact on global energy efficiency, and the possible degradation of the Quality of Service (QoS) offered to users.

This article describes a proposal for developing and applying demand response strategies on large consumers allowing them to participate in the electric market and provide ancillary services. As a case study, the project proposes to address the planning of supercomputing and datacenters, conceived as an example of planned systems that have emerged in modern societies, linked to the smart grid paradigm (other relevant examples are fleets of electric cars, smart buildings, irrigation systems, etc.). Supercomputing and datacenters provide scenarios that allow the direct experimentation of demand response strategies in the academic and business environments. These platforms can adjust power consumption in order to help the electric network to fulfill specific goals, either by consuming available surplus of energy to execute complex tasks, or by deferring activities (i.e., tasks execution) when energy is more expensive or generation is lower than normal. Furthermore, their thermal/cooling infrastructures demand about half of the energy consumption and provide a large inertia, that can be used to interact with the power grid. The studied strategies allow implementing a smart management of the electric grid, achieving a rational utilization of renewable energy sources, and the correct utilization of information technologies to improve decision-making processes.

Strategies for optimal planning of the execution of tasks and energy utilization are proposed the National Supercomputing Center in Uruguay (Cluster-UY) [22], taking into account the energy consumption, the QoS provided to users, and the thermal/cooling demands of the infrastructure. In addition, the business opportunities and business models for supercomputing and datacenters in the electric market are revisited. Results suggest the effectiveness of the proposed strategies to implement demand response techniques and provide ancillary services under the smart grid paradigm.

The article is organized as follows. Next section describes the model applied to characterize the energy consumption on datacenters. Section 3 describes the opportunities for datacenters in the electric market. The proposed strategies for energy-aware planning of datacenters are summarized in Section 4, including some preliminary results for smart planning to follow a reference power profile. Finally, Section 5 formulates the main conclusions and current lines of work.

2 Modeling the energy consumption of datacenters

This section presents an analysis of the power consumption of the main components of a datacenter. Since servers are a key part of the datacenter energy usage, a power model for servers is introduced and a case study is evaluated. Finally, a specific power model for high-end multicores is introduced.

2.1 Breakdown of the power consumption of datacenters

Two main operational components account for most power consumption of datacenters: i) operation of the technological infrastructure (servers, network, storage, etc.) and ii) operation of the cooling system and other physical resources [23,27]. Both sources of power consumption are related because more power is required for the cooling system when servers operate a full capacity. Servers represent a significant percentage of datacenter power consumption and the variability of their power consumption in different load levels allows implementing specific techniques for energy savings. Moreover, variability can be used for demand response under external changes related to energy prices, temperature, etc.

Power models are used for predicting the servers power consumption and evaluating the efficacy of energy aware policies. Due to the high complexity and cost, the quality of energy aware policies is evaluated with simulation tools. Power consumption of high-end servers found in datacenters is broadly described by Eq. 1, where P_{idle} is the server power consumption without load and P_{peak} is the server power consumption at full (100%) utilization. The variable u is the current utilization percentage of the server and function f describes the relationship between utilization and power consumption [1,2].

$$P_{server} = P_{idle} + (P_{peak} - P_{idle})f(u) \tag{1}$$

Most of power consumption of servers corresponds to the CPU. However, power consumption of other computing resources (memory, disk, network) are not negligible. Through workload categorization by resource utilization highly precise power models can be built. Modeling power consumption considering resource utilization also allows taking advantage of task consolidation. Eq. 2 shows a server power model where u_{CPU} is the percentage of server capacity executing workload categorized as CPU-intensive, u_{mem} is the percentage of server capacity executing workload categorized as memory-intensive, and so on for each resource in the model.

$$P_{server} = P_{idle} + (P_{peak} - P_{idle})f(u_{CPU}, u_{mem}, u_{disk}, u_{inet}, \dots) \tag{2}$$

The empirical study of AMD and Intel multicores by Muraña et al. [20] showed that for CPU-intensive workloads, the server power consumption has a linear relationship with resource utilization. Furthermore, power consumption of memory-intensive workload decelerates as utilization increases. Power consumption of memory-intensive workload was greater than CPU-intensive workload.

Some works have proposed empirical energy models that consider types of computing resources, measuring power consumption using different benchmarks (intensive in one specific computing resource), such as Linpack [12,17], Abinit [9], and or Namd [17,25]. Power data can be collected through software tools that consult internal hardware counters–e.g., Running Average Power Limit (RAPL) interface on Intel servers–or by using an external power meter [8,20,28].

2.2 Empirical analysis of power consumption of servers in a datacenter

Power characterization measurements were performed over a HP ProLiant DL380 G9 server (2 Intel Xeon Gold 6138 CPUs, 20 cores each, 128 GB RAM) from Cluster-UY. The experiment consisted in executing a CPU-intensive benchmark and measuring its power consumption using *likwid* [29], a software tool that allows access to RAPL interface counters to estimate the power consumption.

Algorithm 1 presents the procedure applied for energy measurement. The power consumption reported by *likwid* is logged while executing an increasing number of benchmark instances to consider different levels of server utilization.

Algorithm 1 Procedure for power consumption measurement

```
process_per_level ← 5
utilization_levels ← 8
independent_executions ← 30
for j = 1 to independent_executions do
    likwid-power-meter -s 60s
end for
for i = 1 to utilization_levels do
    instances_current_level ← process_per_level × i
    for j = 1 to independent_executions do
        launch_benckmark_instances(instances_current_level,60)  ▷ launch in background the benchmark instances for 60 seconds
        likwid-power-meter -s 60s
    end for
end for
```

For the experiments, eight utilization levels (UL) were defined with five process per level. Utilization level zero corresponds to server without load. UL one corresponds to 12.5% of server utilization, UL two corresponds to 25% of server utilization, and so on. Power consumption of each level is measured 30 times to obtain statistically significant values. Measurements for each UL last 60 seconds.

The CPU-intensive benchmark utilized for experiment belongs to the Sysbench toolkit [16]. The benchmark is a procedure written in language c to calculate the prime number counting function using a backtracking technique.

Figure 1 reports the results of power consumption measurements of the CPU-intensive benchmark. The independent variable u corresponds to the percentage of server utilization and PC is the power consumption (in Watts) reported by *likwid* by consulting the RAPL interface. A significant difference in power consumption is measured between utilization zero and the following levels. This difference is explained by the internal power management of Intel chip (decreasing voltage of inactive resources).

u	*PC (W)*
0 %	30.06± 3.91
12.5 %	120.38 ± 0.51
25 %	142.36 ± 0.72
37.5 %	162.83 ± 0.90
50 %	176.17 ± 0.46
62.5%	180.03 ± 1.02
75 %	192.63 ± 1.01
87.5%	194.09 ± 0.50
100%	208.17± 0.44

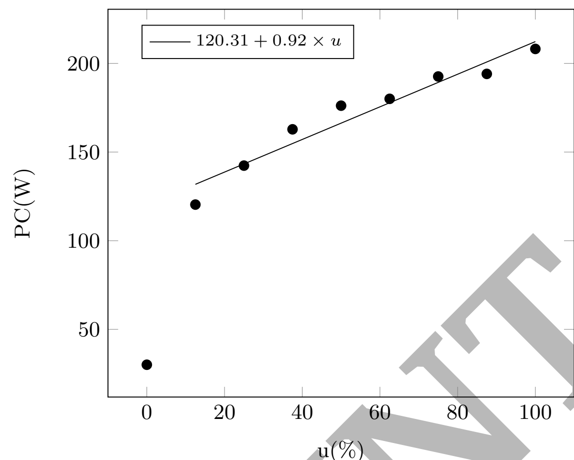

Fig. 1. Power consumption of CPU-intensive benchmark over Cluster-UY multicore

If utilization zero is not considered, the power consumption can be adjusted to a linear function, for example, using least squares. The derivative of the function (0.92) is coherent to the one reported in [20] (0.82), where the same benchmark was measured using a Power Distribution Unit over a similar high-end server. The same work also reports experimental result of memory-intensive benchmark in similar high-end servers. Eq. 3, introduced in [20], presents a linear combination of models of CPU-intensive and memory-intensive workload. Eq. 3, u_{CPU} is the server utilization corresponding to CPU-intensive workload and u_{mem} is the server utilization corresponding to memory-intensive workload. The variable u'_{CPU} is zero when u_{CPU} is less than 50% and $u_{CPU} - 50$ otherwise. An analog model can be built for the specific hardware of the case study.

$$\begin{aligned} P_{server} &= 0.802 \times u_{CPU} + 0.042 \times u'_{CPU} + 2.902 \times u_{mem} \\ &-0.02107 \times u_{mem}^2 + 7.644 \times 10^{-5} \times u_{mem}^3 + \tfrac{56.36+36.89}{2} + 57.0 \end{aligned} \tag{3}$$

Since the downside of energy savings is the degradation of system performance, the energy model must be complemented with a performance model. To empirically model the performance, similar experiments should be performed considering execution times instead of power consumption.

3 Opportunities for datacenters in the electric market

This section describes the different ways a datacenter may participate in demand response and ancillary services mechanisms and introduces the particular case of multi-tenants datacenters.

3.1 Participation in the electric market

A flexible consumer needs planning techniques to ensure a proper use of its energy resources and to response to the energy market signals. In a datacenter, the energy is used evenly distributed into two particular sectors: the operative hardware that provides the services required by the datacenter clients and the thermal/cooling infrastructure. These are the knobs that may be adjusted according some time-varying power consumption profile. In this way, the datacenter can participate as an active agent in the electric market.

A relative simple way is to implement a mechanism of demand response, using the thermal inertia of the building to increase or decrease the power consumption, letting move the building temperature between acceptable levels. In order to define the limits of an electric power band that the datacenter can offer to the system operator, a proper model of the building temperature dynamics must be used. The more accurate the model, the more wide the offered power band and the more profit can be obtained. Of course, the model that describes the temperature evolution should include the impact of the servers activity, and this fact leads to the inclusion of the tasks execution profile into the datacenter demand response strategy. As explained in Section 2, the execution of the tasks directly consumes electric power and also affects the building temperature. Maintaining that temperature within prescribed limits implies the utilization of the thermal/cooling units, that also consume electric power. In this way, an appropriate demand-response strategy should combine the flexibility of the thermal behavior and the tasks allocation.

3.2 Demand response in multi-tenants datacenters

Over the two main actions of a datacenter, new variants can be devised. This section focuses on a pricing mechanism for multi-tenant datacenter that allows the operator to obtain load shedding among tenants. Following the ideas of our previous work [19], a responsive scheme for the clients is proposed. Clients may choose to postpone or lose a task in exchange of some kind of economical reward provided by the datacenter, which is an active agent in the electric market.

We pay special attention on multi-tenants collocation datacenter, since the tenants deploy and keep full control of their own physical servers, while the datacenter operator provides facility support. Tenant's workloads in collocation datacenter are highly heterogeneous, and many tenants run non-critical workloads, with high scheduling flexibility, different delay sensitivities, different service level agreements with peak loads periods. This type of datacenters are often located in metropolitan areas, where demand response calls are most needed. They can participate actively in the energy market by modulating their power profile and helping maximize distribution grid resources. The main disadvantage is that each tenant manage its own servers independently and has very different incentives to cooperate with the operator during a demand response event.

In an electricity market with uncertainty in supply or price volatility, supply function as a strategic variable allows to adapt better to changing market conditions than a simple commitment to a fixed price or quantity does [15]. This is one reason why we propose to use supply function bidding, creating a market mechanism which fixed a uniform market clearing price. Other motivation is to respect practical informational constraints in the power network. A customer might not want to reveal its cost function because of incentive or security concerns, or the cost function may require a high description complexity, which means more communication. A properly chosen parameterized supply function *controls* information revelation while demands less communication.

Chen et al. [3] considered two abstract market models for designing demand response to match the supply and shape the demand, respectively. In the modeled situation, there is an inelastic supply deficit on electricity, and study a supply function bidding scheme for allocating load shedding among different users to match the supply. Each customer submits a linear parameterized supply function to the agent aggregator (i.e., the datacenter operator). In a competitive market where customers are price taking, the system achieves an efficient equilibrium that maximizes the social welfare. In an oligopolistic market where customers are price anticipating and strategic, the system achieves a unique Nash equilibrium that maximizes another additive, global objective function.

Montes de Oca et al. [19] proposed a distributed algorithm to optimize social welfare over a distribution network considering AC physical constraints over the grid but with several users aggregators. However, these forms of parameterized supply function do not admit treatable analysis. Johari and Tsitsiklis [14] considered an alternative supply function model (Eq. 4) where a finite number of producers compete to meet an infinitely divisible but inelastic demand reduce δ. Each user (or tenant) is characterized by a production cost, convex in the output produced, and the customers act as profit maximizers. The mechanism yields bounded efficiency loss at a Nash equilibrium and also characterize the problem of finding the Nash equilibrium as the solution of a collocation problem.

$$S_n(b_n, p) = \delta - \frac{b_n}{p} \tag{4}$$

Chen et al. [4] extended the previous work by proposing a uniform pricing mechanism for collocation datacenters where the operator can extract load shedding from tenants, without using the backup generator. The goal is to effectively provide incentives for tenants to reduce energy consumption during emergency demand response events. When an emergency demand response arrives, tenants bid using a parameterized supply function (Eq. 4), and then the datacenter operator announces a market clearing price which when plugged into the bids, specifies how much energy tenants will reduce and how much they will be paid. The main advantage of this mechanism is that for the tenants is very easy to participate in the market since they are only asked to bid a parameter but keeping the integrity of the private information. The authors propose a market mechanism and prove existence and uniqueness of the best strategy for each tenant. In addition, they characterize the Nash optima of the non-cooperative game as an optimization problem, which can be solved in a distributed manner between participants, preserving private information. A mathematical model for this approach is presented in Section 4.2.

4 Smart strategies for effective planning of datacenters

This section describes strategies for datacenter planning and operation and a proposal for a demand response scheme in a multi-tenant datacenter.

4.1 Datacenter planning and operation

One of the key issues related with energy-aware datacenter planning refers to the problem of following a reference power profile for energy consumption. The main goal is to appropriately plan the execution of tasks and the operation of the cooling system to minimize the deviation with respect to the reference power profile. This way, the datacenter can adapt its operation and participate in the energy market as an agent with the capabilities of fulfilling specific goals.

Our group has developed research on the holistic energy-aware planning of datacenters, and also including the use of renewable energy sources [7,10,11,22,23,24]. The general approach consists in applying computational intelligence methods [21] to solve the underlying optimization problem that proposes determine the tasks scheduling and the energy consumption of both infrastructure and cooling systems, subject to QoS and operation (e.g., temperature) constraints.

Fig. 2 presents an overview of the proposed system model, including their two key components: the computing infrastructure and the cooling system.

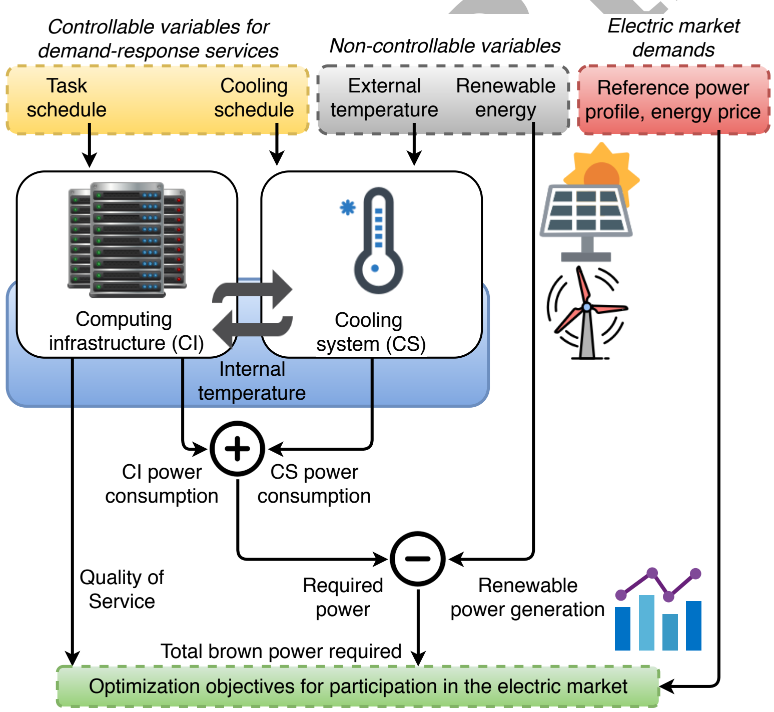

Fig. 2. Schema of the proposed model for energy-aware planning in datacenters

A realistic energy consumption model is considered for the computing infrastructure in which each computing resource may be executing, idle, or asleep. On the one hand, when a computing resource is executing a task, it is considered to be at its peak performance. On the other hand, when a computing resource is idle it is considered to be consuming the minimum amount of energy required of its operation. Finally, when asleep, a computing resource considered to be consuming a marginal amount of energy.

The objective are controlled by two input variables, the task schedule and the cooling schedule. The task schedule determines the execution of tasks on the computing infrastructure for the scheduling horizon. Likewise, the cooling schedule determines the on/off of the cooling system for the scheduling horizon. Three non-controllable input variables are considered: external temperature (the air temperature outside the datacenter), renewable energy generation (amount of available energy generated by renewable energy source such as solar panels, wind turbines, etc.), and power reference profile of the electric market, used by the datacenter to provide ancillary services and to consider demand response to match the supply and shape the demand. These variables are not controllable by the system and have uncertainty.

The optimization objectives and constraints are defined as follows. The maximization of the QoS is related to the number of tasks with unmet due dates. The total brown energy required by the datacenter and the reference power profile requested by the electric market for maximizing the profit. The internal temperature of the datacenter is constrained to a maximum operating value.

We proposed a number of exact methods, stochastic and deterministic heuristics, and single- and multi-objective metaheuristics for addressing several variants of this optimization problem with promising results [10,11,21,23,24]. As an example, our previous work [24] proposed the following mathematical model.

Controllable variables: *cooling schedule* (c_k), controls the operation of the cooling system; and the *power schedule* (s_k), controls the computing infrastructure power consumption. It controls the number of servers running, load constraints, and specific user requirements.

Non-controllable variables. External temperature (α_k), air temperature outside the datacenter; *target reference power profile* (R_k), the desired total power consumption for each time step; and the *target reference temperature profile* (T_{ref}), the desired internal temperature of the datacenter for each time step.

Other variables. The *internal temperature* (T_k) in the datacenter; the power consumption of the cooling system (C_k); the *power consumption of the computing infrastructure* (I_k); and the *total power consumption of the datacenter* (P_k).

The total computing infrastructure power is defined by $I_k = S_k^{max} + S_k^{idle} + S_k^{sleep}$. Where S_k^{max}, S_k^{idle} and S_k^{sleep} are the total power of all servers that are executing, idle, and sleep at time k, respectively.

The datacenter must execute a set of n tasks in a simulation period of K time steps. Each task i must finish before a deadline $D(i)$. The actual finishing time of a task $FT(i)$ and its deadline $D(i)$ define whether a deadline is satisfied or violated and contributes to the QoS of the schedule.

The main goal is to schedule the operation of a datacenter in order to follow as closely as possible a predefined power and temperature reference while simultaneously minimizing its impact on the QoS of the system. Formally, this means to minimize the deviation from the reference power profile (Eq. (5)) and the deviation from the reference temperature profile (Eq. (6)), while simultaneously minimizing the total exceeding time of deadline violations (Eq. (7)).

$$\sum_{k=1}^{K} \frac{|P_k - R_k|}{\max(R_k)} \tag{5}$$

$$\sum_{k=1}^{K} |T_{ref} - T_k| \tag{6}$$

$$\sum_{i=1}^{K} \max(0, FT(i) - D(i)) \tag{7}$$

Our previous work [24] proposed a multiobjective evolutionary approach for solving the proposed problem. The experimental results show the proposed approach computes accurate schedules for all objectives as well as competitive trade-off schedules. Fig. 3 show the computed solution for the reference power profile objective. It shows the power consumption closely follows the reference power, enabling the datacenter to potentially reduce electricity costs, maximize renewable energy use, or participate in the electricity market.

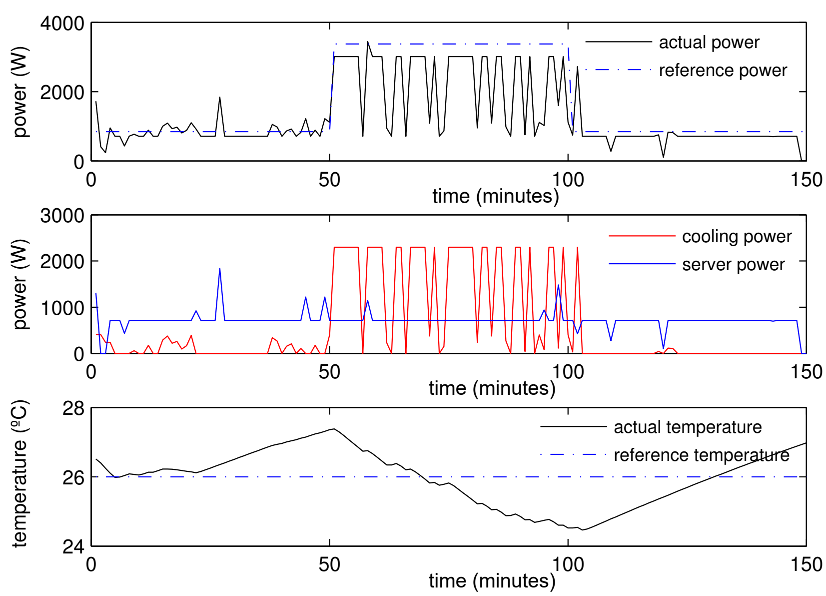

Fig. 3. Best computed solution for the reference power profile objective

4.2 Proposal for a demand response scheme in a multi-tenant datacenter

This section proposes an optimization model for the demand response scheme described in Section 3.2.

A simple model of the cooling infrastructure and the thermal inertia of the building to increase or decrease the power consumption is proposed, letting move the temperature into the room between acceptable levels. For that, a multistage setting is considered. The proposal is based on a simple mechanism under which each consumer submits a single bid that reflects the willingness to adjust the consumer's demand over the entire T stages. Such mechanisms are easy to implement with a parameterized supply function, and would require the minimum effort from the tenants. In this line of word, results must be established on equilibrium characterization and bounded efficiency loss, analogous to those derived in related works [4,14].

Overview of market mechanism. A market mechanism was conceived, where tenants bid for the next T stages using parameterized supply functions (Eq. 4) and then, given the bids, the operator decides how much load to shed via tenants and how much to shed via on-site generation and cooling system.

The operation of the market is summarized below:

- The datacenter operator receives an emergency demand response event for a reduction target $\delta := \{\delta^1, .., \delta^T\}$ and broadcasts the supply function $\mathbf{S}(\cdot, \mathbf{p})$, specified by Eq. 4, to tenants;
- Participating tenants respond by placing their bids $\mathbf{b}_n := \{b_n^1, .., b_n^T\}$;
- The operator decides the amount of on-site generation an the temperature scheduling and calculate market clearing price $\mathbf{p}$ to minimize its cost for T stages, using Eq. 8 to set the market clearing price p and Eq. 9 to set $\mathbf{y}$ and $\Delta\mathbf{P}_c$, minimizing the cost of the operation during the demand response event;
- Demand response event is exercised. Tenant n sheds $\mathbf{S}_n(\mathbf{b}_n, \mathbf{p})$, and receives $\mathbf{S}_n(\mathbf{b}_n, \mathbf{p}) \cdot \mathbf{p}^t$ as a reward.

The clearing market price is given by Eq. 8. This mechanism is illustrated in Figure 4.

$$p^t(b_n^t, y^t, \Delta\mathbf{P}_c^t) = \frac{\sum_n b_n^t}{(N-1)\delta^t + y^t + \Delta\mathbf{P}_c^t} \quad (8)$$

To determine the vector of local generation amount $\mathbf{y}$ and power cooling reduction $\Delta\mathbf{P}_c$, the operator minimizes the cost of the three load-reduction options, given by Eq. 9.

$$(\mathbf{y}, \Delta\mathbf{P}_c) = \arg\min(\delta + \mathbf{y} + \Delta\mathbf{P}_c) \cdot \mathbf{p}(\mathbf{b}_n, \mathbf{y}, \Delta\mathbf{P}_c)^T + \alpha \cdot \mathbf{y}^T + \mathcal{C}(\Delta\mathbf{P}_c; t) \quad (9)$$

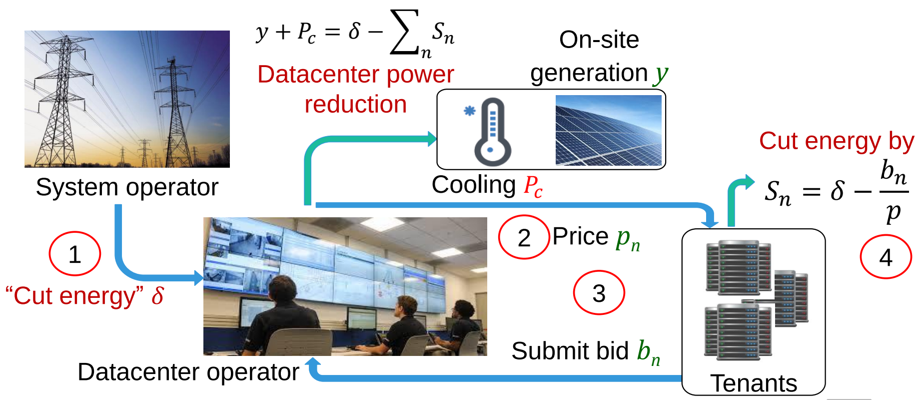

Fig. 4. Market mechanism for the proposed demand response scheme

Modeling cooling power. A simple temperature model can be considered as a function of the power for cooling P_c and the outdoor temperature T_{out} (Eq. 10).

$$\dot{T}(t) = a_1 \left[T(t) - T_{out}(t)\right] + a_2 P_c(t) \tag{10}$$

The model needs to penalize the misalignment between the actual temperature T_{in} and a set-point temperature T_{set}. The cost function in Eq. 11 is considered, where $\Delta\mathbf{P}_c(t)$ is the power difference between the power consumption at time t and the reference power for cooling before demand response takes place.

$$\mathcal{C}(\Delta\mathbf{P}_c; t) = \kappa \| T(\Delta\mathbf{P}_c; t) - \mathbf{T}_{set}(t) \| \tag{11}$$

Cost function $c_n(s)$. Chen et. al. [4] proposed a cost function $c_n(\cdot)$ that captures the effect of switching off m computers in a M/G/1/Processor-sharing queue, let's first consider an auxiliary function $\bar{c_n}(\cdot)$ defined as: $\bar{c_n}(m) = \frac{\beta T}{\frac{1}{vM} - \frac{1}{M-m}}$, where λ is the workload arrival rate, $v = \frac{\lambda}{\mu M}$ is the normalized workload arrival, μ is the service rate, β is a cost parameter ($/time unit/job), T is the duration of the power reduction event, M is the total of available servers and m the number of switched off servers for tenant n. The power reduction model is considered linear in m, so that $S_n = \theta m$. Then the cost function for a tenant's energy reduction is written as: $c_n(S_n) = \bar{c_n}(S_n/\theta) - \bar{c_n}(0)$ and 0 otherwise.

Efficiency analysis. The next step is to characterize the efficiency of the mechanism. There are two potential causes of inefficiency: the cost minimizing behavior of the operator and the strategic behavior of the tenants. In particular, since the forms of the tenant's cost functions are likely more complex than the supply function bids, tenants cannot bid their true cost function even if they wanted to. This means that evaluating the equilibrium outcome is crucial to understanding the efficiency of the mechanism. The equilibrium outcome depends highly on the behavior of the tenants whether they are price-taking or price-anticipating. The key to our analysis is the observation that the equilibrium can be characterized by an optimization problem. Once we have this optimization, we can use it to characterize the efficiency of the equilibrium outcome. This approach parallels the one proposed by Chen et al. [4] and Johari and Tsitsiklis [14].

Adding uncertainty. The task arrivals could introduce uncertainty that would be better captured by probabilistic models. We are interested in deriving these models for the uncertainty in the costs and prices from the queuing theory modelling arrivals. Previous work relating workloads with prices and power resource allocation can be found in [5]. Another line of research is the negotiation of the power reduction levels (δ) between the grid and the datacenter operators. Under uncertainty of random effects, and constraints in the power level provided by the diesel generators, this δ may not be accommodated and should be negotiated taking into account its conditional value at risk [6,26].

5 Conclusions and future work

This article introduced a proposal for supercomputing platforms and datacenters to participate in the electric market, by implementing demand response techniques and ancillary services.

A methodology was introduced for supercomputing and datacenters to adjust their power consumption in order to help the electric network to fulfill specific goals, either by consuming available surplus of energy to execute complex tasks, or by deferring activities when energy is more expensive or generation is lower than normal.

Smart strategies for effective energy-aware planning of datacenters were described, including a methodology applying computational intelligence for the problem of following a reference power profile, subject to QoS and temperature constraints, considering the power consumption of computing infrastructure and thermal/cooling system. A specific model is introduced for demand response in a multi-tenant datacenter applying a multistage procedure.

Preliminary results demonstrate that the proposed strategies allow implementing a smart management of the electric grid, achieving a rational utilization of renewable energy sources, and the correct utilization of information technologies to improve decision-making processes.

The main lines for current and future work are related to develop the proposed model and apply it to a relevant case study: The National Supercomputing Center in Uruguay (Cluster-UY), for which preliminary studies on evaluation and characterization of the power consumption of the computing infrastructure were also presented. The proposed models should be further improved to capture the reality of the case studies. Furthermore, more complex strategies are being studied to implement demand response techniques and provide ancillary services under the smart grid paradigm, including the application of single-objective and multi-objective computational intelligence methods.

References

1. Beloglazov, A., Buyya, R., Choon Lee, Y., Zomaya, A.: A taxonomy and survey of energy-efficient data centers and cloud computing systems. Advances in Computers **82**, 47–111 (2010)
2. Chen, F., Grundy, J., Yang, Y., Schneider, J.G., He, Q.: Experimental analysis of task-based energy consumption in cloud computing systems. In: 4^{th} ACM/SPEC International Conference on Performance Engineering. pp. 295–306 (2013)
3. Chen, L., Li, N., Low, S., Doyle, J.: Two market models for demand response in power networks. In: First IEEE International Conference on Smart Grid Communications (2010)
4. Chen, N., Ren, X., Ren, S., Wierman, A.: Greening multi-tenant data center demand response. Performance Evaluation **91**, 229–254 (2015)
5. Chen, T., Marques, A., Giannakis, G.: Dglb: Distributed stochastic geographical load balancing over cloud networks. IEEE Transactions on Parallel Distributed Systems **7**, 1866–1880 (2017)
6. Dall'Anese, E., Baker, K., Summers, T.: Chance-constrained ac op-timal power flow for distribution systems with renewables. IEEE Transactions on Power Systems **32**(5), 3427 – 3438 (2017)
7. Dorronsoro, B., Nesmachnow, S., Taheri, J., Zomaya, A.Y., Talbi, E.G., Bouvry, P.: A hierarchical approach for energy-efficient scheduling of large workloads in multicore distributed systems. Sustainable Computing: Informatics and Systems **4**(4), 252–261 (2014)
8. Du Bois, K., Schaeps, T., Polfliet, S., Ryckbosch, F., Eeckhout, L.: Sweep: Evaluating computer system energy efficiency using synthetic workloads. In: 6^{th} International Conference on High Performance and Embedded Architectures and Compilers. pp. 159–166 (2011)
9. Gonze, X., Amadon, B., Anglade, P., Beuken, J., Bottin, F., Boulanger, P., Bruneval, F., Caliste, D., Caracas, R., Côté, M., Deutsch, T., Genovese, L., Ghosez, P., Giantomassi, M., Goedecker, S., Hamann, D., Hermet, P., Jollet, F., Jomard, G., Leroux, S., Mancini, M., Mazevet, S., Oliveira, M., Onida, G., Pouillon, Y., Rangel, T., Rignanese, G., Sangalli, D., Shaltaf, R., Torrent, M., Verstraete, M., Zerah, G., Zwanziger, J.: Abinit: First-principles approach to material and nanosystem properties. Computer Physics Communications **180**(12), 2582 – 2615 (2009)
10. Iturriaga, S., Dorronsoro, B., Nesmachnow, S.: Multiobjective evolutionary algorithms for energy and service level scheduling in a federation of distributed datacenters. International Transactions on Operational Research **24**(1-2), 199–228 (2017)
11. Iturriaga, S., Nesmachnow, S.: Scheduling energy efficient data centers using renewable energy. Electronics **5**(4), 71 (2016)
12. Iturriaga, S., García, S., Nesmachnow, S.: An empirical study of the robustness of energy-aware schedulers for high performance computing systems under uncertainty. In: High Performance Computing. pp. 143–157 (2014)
13. Jie, B., Tsuji, T.: An analysis of market mechanism and bidding strategy for power balancing market in micro-grid. In: China International Conference on Electricity Distribution (2016)
14. Johari, R., Tsitsiklis, J.N.: Parameterized supply function bidding: Equilibrium and efficiency. Operations Research **59**(5), 1079–1089 (2011)
15. Klemper, P.D., Meyer, M.A.: Supply function equilibria in oligopoly under uncertainty. Econometrica **57**(6), 1243–1277 (1989)

16. Kopytov, A.: Sysbench repository. Online: https://github.com/akopytov/sysbench, [01/2019]
17. Kurowski, K., Oleksiak, A., Piatek, W., Piontek, T., Przybyszewski, A., Weglarz, J.: Dcworms–a tool for simulation of energy efficiency in distributed computing infrastructures. Simulation Modelling Practice and Theory **39**, 135–151 (2013)
18. Momoh, J.: Smart Grid: Fundamentals of Design and Analysis. Wiley-IEEE Press (2012)
19. Montes de Oca, S., Belzarena, P., Monzon, P.: Optimal demand response in distribution networks with several energy retail companies. In: IEEE Multi-Conference on Systems and Control. pp. 1092–1097 (2016)
20. Muraña, J., Nesmachnow, S., Armenta, F., Tchernykh, A.: Characterization, modeling and scheduling of power consumption of scientific computing applications in multicores. Cluster Computing (2019), online first, https://doi.org/10.1007/s10586-018-2882-8 [01/2019]
21. Nesmachnow, S.: An overview of metaheuristics: accurate and efficient methods for optimisation. International Journal of Metaheuristics **3**(4), 320–347 (2014)
22. Nesmachnow, S., Iturriaga, S.: Cluster-UY: scientific HPC in Uruguay. In: International Supercomputing in México. pp. 1–15 (2019)
23. Nesmachnow, S., Perfumo, C., Goiri, Í.: Multiobjective energy-aware datacenter planning accounting for power consumption profiles. In: High Performance Computing. CARLA 2014. Communications in Computer and Information Science, vol. 485, pp. 128–142. Springer (2014)
24. Nesmachnow, S., Perfumo, C., Goiri, Í.: Holistic multiobjective planning of datacenters powered by renewable energy. Cluster Computing **18**(4), 1379–1397 (2015)
25. Phillips, J., Braun, R., Wang, W., Gumbart, J., Tajkhorshid, E., Villa, E., Chipot, C., Skeel, R.D., Kale, L., Schulten, K.: Scalable molecular dynamics with namd. Journal of Computational Chemistry **26**(16), 1781–1802 (2005)
26. Rockafellar, R., Uryasev, S.: Optimization of conditional value-at-risk. Journal of Risk **2**, 21–42 (2000)
27. Rong, H., Zhang, H., Xiao, S., Li, C., Hu, C.: Optimizing energy consumption for data centers. Renewable and Sustainable Energy Reviews pp. 674–691 (2016)
28. Srikantaiah, S., Kansal, A., Zhao, F.: Energy aware consolidation for cloud computing. In: Conference on Power Aware Computing and Systems. pp. 1–5 (2008)
29. Treibig, J., Hager, G., Wellein, G.: Likwid: A lightweight performance-oriented tool suite for x86 multicore environments. In: 39^{th} International Conference on Parallel Processing Workshops. pp. 207–216 (2010)