Two Level Demand Response Planning for Retail Multi-Tenant Datacenters

Jonathan Muraña, Sergio Nesmachnow, Santiago Iturriaga,
Sebastián Montes de Oca, Gonzalo Belcredi, Pablo Monzón
Universidad de la República
Uruguay
{jmurana,sergion,siturria,smontes,gbelcredi,monzon}@fing.edu.uy

Andrei Tchernykh
CICESE Research Center, Ensenada, Mexico
South Ural State University, Chelyabinsk, Russia
Ivannikov Institute for System Programming, Moscow, Russia
chernykh@cicese.mx

***Abstract*—This article presents a two level planning approach for the participation of datacenters in the electricity market, by providing appropriate actions to demand response events, a relevant issue to contribute to the smart grid paradigm. A planning strategy is proposed for retail collocation datacenters, considering their heterogeneity. A negotiation is applied at the datacenter level and scheduling heuristics are used at the tenants level. Three heuristics are evaluated, accounting for different features of tasks submitted for execution. The proposed approach is evaluated considering realistic infrastructures (servers and air conditioning system) and workloads. The main results suggest that the proposed approach provides appropriate plannings for demand response events, significantly improving over Business-as-Usual operation and when there is not electricity load shedding. The Nash+PL heuristic computed the best results regarding cost while accounting for appropriate quality of service.**

Keywords*—*demand response*, *computational intelligence*, *planning*, *datacenters

I. Introduction

Within the modern smart grid paradigm [1], a large consumer with flexible power utilization can participate in the electricity market under different modalities. This participation is encouraged by grid operators as one of the main concepts for implementing strategies oriented to provide services in modern smart electric networks [2]. A large consumer can participate by adapting his electricity utilization to the needs of the grid, e.g., reducing consumption in peak periods to help flatten the demand curve of the electrical system or even save the grid from suffering a blackout. It can also provide ancillary services associated with specific events. Demand response events trigger a change in power consumption to better match the demand for power with the supply.

To account for proper active participation in the electricity market by executing effective actions to demand response events, a large consumer must plan his activities in a satisfactory manner. Thus, appropriate planning methods that take into consideration the dynamic situation of the market should be applied. Datacenters and high performance computing facilities are examples of plannable systems for which the proposed approach can be successfully applied. Datacenters have several methods for adjusting power consumption to help the electric network in demand response events, including consuming an available surplus of energy, deferring tasks execution, and regulating the utilization of thermal/cooling infrastructures.

The suggested planning approach is valuable for *multi-tenant* or *colocation* datacenters, usually operated by a company that provides services to multiple enterprise tenants. This article focuses on the retail multi-tenants model, which typically serves a relatively large number of tenants (in opposition to wholesale datacenters). Each tenant applies its own business model (regarding QoS offered to clients, reputation, etc.). Retail multi-tenants datacenters account for about 75% of the colocation market [3].

Tenants in a retail collocation datacenter are responsible for deploying and keeping full control of servers to provide services to clients. Clients are heterogeneous, and submitted workloads reflect that heterogeneity, including tasks that vary from being critical to best-effort. Non-critical tasks have flexibility for execution, and this fact can be exploited to adjust the power profile of colocation datacenters to attend demand response events. Furthermore, since multi-tenant datacenters usually consume significantly more energy than ad-hoc datacenters [4], developing and applying effective energy-aware planning techniques provides an improved operation model.

In this line of work, this article presents a two level approach for demand response planning to be used in retail multi-tenant datacenters. Two levels are considered: energy-aware scheduling at the tenants level and negotiation on the datacenter operator level. Three sources are considered to attain a specific energy consumption reduction requested by the electricity market: consolidating and/or deferring non-critical tasks (tenants), lowering the energy required for Heating, Ventilation and Air Conditioning (HVAC) according to the executed tasks, and on-site generation (operator). Two heuristics are proposed and compared with a Business as Usual (BaU) operation over realistic problem instances from real datacenters and using an empirical energy consumption model. Results demonstrate the effectiveness of the proposed approach, computing solutions that improve up to **60.7%** over the BaU operation.

The article is organized as follows. Section II presents the demand response planning problem for datacenters and a review of related works. Section III describes the proposed two-level planning approach to solve the problem. Section IV reports the experimental analysis of the developed model and heuristics. Finally, conclusions and future work are formulated in Section V.

II. Demand Response Planning Problem for Datacenters

This section describes the problem of energy-aware remand response planning for datacenters and reviews related works.

A. Problem model

The problem addressed in this article is to plan energy utilization of datacenters to attend demand response events from the electricity market. The planning procedure starts as a response to an energy reduction requested by the smart grid. Then, the datacenter operator initiates a negotiation with its tenants: monetary incentives are offered to each tenant to lower their energy consumption. Tenants apply a local energy-aware scheduling heuristic to determine the attainable energy consumption reduction for the proposed monetary incentive, considering QoS offered to clients. In turn, in view of the energy consumed by tenants, the datacenter operator determines the possible reduction by lowering or switching off the HVAC system, and finally determines if the on-site generation unit must be used to provide enough power to achieve the committed energy reduction. A conceptual diagram of the problem model is presented in Fig. 1, identifying the three agents in the model (market, datacenter operator, and tenants) and their interactions.

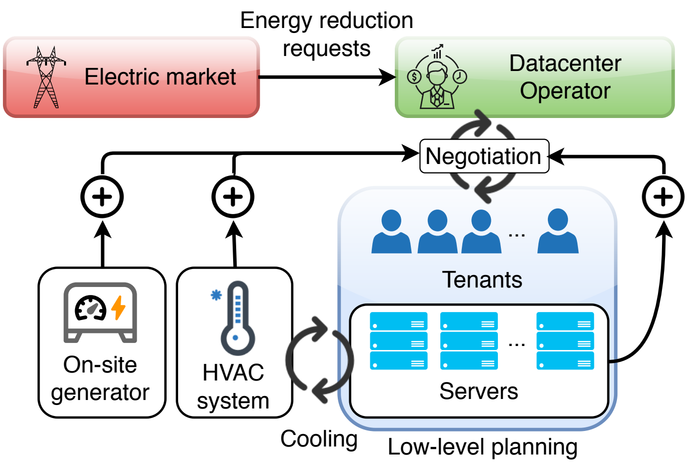

Fig. 1: Schema of the proposed model for demand response planning in retail multi-tenant datacenters

B. Problem formulation

The underlying optimization problem proposes minimizing the total operation cost for the datacenter to attend a demand response event, considering the following components: the total monetary incentive rewarded to tenants, the cost of using the HVAC system, and the cost of using the on-site generator in order meet the reduction target.

The problem formulation is as follows. Consider:

- A set of discrete timesteps t in the planning interval $t \in [0, T]$.
- A target reduction ρ, requested by the electric market, to be attained at every timestep of the planning interval (the reference energy consumption is the BaU operation).
- A set of tenants $C = \{c_1, ..., c_{|C|}\}$ and a monetary incentive RI_j offered to tenant c_j for each energy unit reduced.
- A workload of tasks W_j for each tenant c_j, $W_j = \{w_j^1, ..., w_j^{|W|}\}$. Each task w_j^i demands a certain number of operations l_j^i (the *length* of the task, expressed in millions of instructions). Each task w_j^i has an arrival date AD_j^i and also a due date or *deadline* DD_j^i.
- Function $\Phi_j(RI_j)$ determines a workload schedule s_j for each client c_j. FT_j^i is the finishing time of task w_j^i for schedule s_j. s_j^0=$\Phi_j(0)$ is the workload schedule for tenant c_j when no monetary incentives are offered (BaU operation). P_j^t is the servers' power requirement and H_j^t is the HVAC power requirement at each time t of each workload schedule s_j.
- The violated deadline variable is $VD_j^i = 0$ if $FT_j^i \leq DD_j^i$ for schedule s_j, otherwise $VD_j^i = 1$.
- Function Φ_j determines a new schedule $\hat{s}_j$ with a power requirement $\hat{P}_j^t$ and a new HVAC power requirement $\hat{H}_j^t$ for tenant c_j given incentive RI_j, $\Phi_j(RI_j) = \hat{s}_j$. Two functions are defined to measure the reduction between $\hat{s}_j$ with respect to schedule s_j^0: the energy requirement reduction is $\Delta P(\hat{s}_j) = \min_{\forall t \in T}(P_j^t - \hat{P}_j^t)$, and the HVAC requirement reduction is $\Delta H(\hat{s}_j) = \min_{\forall t \in T}(H_j^t - \hat{H}_j^t)$.
- Let GP^t be the power generated using the on-site generator at time t, and α be the monetary cost per unit of energy of using the on-site generator.
- Let φ be the monetary cost per unit of energy used by the HVAC.

The problem proposes the optimization of the objective function defined in Eq. 1.

$$\min \sum_{j=1...|C|} \Big(\Delta P(\Phi_j(RI_j)) \times RI - \Delta H(\Phi_j(RI_j)) \times \varphi\Big) + \sum_{t \in T} GP^t \times \alpha \tag{1a}$$

subject to:

$$\rho \leq \sum_{j=1...|C|} \Big(\Delta P(\Phi_j(RI_j)) + \Delta H(\Phi_j(RI_j))\Big) + \sum_{t \in T} GP^t \tag{1b}$$

The objective (1a) is to minimize cost for the datacenter operator: terms in the sum correspond to the money paid to tenants, the cost of using the HVAC, and the cost of using the on-site generator in order meet the reduction target ρ. Constraint (1b) states the total energy reduction must be at least ρ per timestep.

Function $\Phi_j(RI_j)$ is computed via simulations, considering realistic settings for each tenant, including how likely is he willing to admit deadline violations for tasks in its workloads and the monetary penalty each tenant must pay to his clients for violated tasks. A specific power consumption model based on empirical analysis of real applications is used [5], [6].

C. Related work

A common approach for the demand response planning problem is considering theoretical supply functions that assume an explicit mathematical relationship between the main features of the proposed model (e.g., energy reduction requested by the market and prices to offer to tenants [7], [8]). Supply functions result in a competition problem between tenants to meet the requested demand reduction [9],. Chen et al. [10] solved the problem using game-theory and a simple aggregation approach for cost of the datacenter operator and tenants, considering an on-site energy generation source and no thermal conditioning. After that, Tran et al. [11] proposed a method for promoting tenants to participate in demand response programs, based on economic rewards and bidding games and Wang et al. [12] studied a distributed cloud-based infrastructure with shared virtual machines between tenants. Several assumptions were applied in these proposals, e.g., requests are not explicitly modeled, but queuing theory is applied; very short tasks are considered; servers are turned-off; and synthetic scenarios were solved. Later, Bahramiet et al. [13] proposed an algorithm to respond to demands in deregulated electricity markets, where the datacenter administrator can choose the best energy supplier to fit its needs. Probabilistic models were used to describe the relation between workload, energy consumption, and QoS degradation (i.e. supply functions).

The approach using (parameterized) supply functions does not capture specific features of real datacenter operation, e.g., Service Level Agreement degradation, real energy consumption of nowadays server and processors, and cooling.

This article contributes by extending our previous works about holistic energy-aware planning of datacenters [14]–[18] applying computational intelligence [19]. Specifically, our previous research [20] is enhanced by including a market mechanism for the active participation of tenants [21], a realistic evaluation of power consumption of computing resources [6], and a simple thermal model for energy reductions of the HVAC operation.

III. The proposed approach for demand response planning

This section describes the proposed approach for demand response planning in retail multi-tenant datacenters.

A. Overall description

A two level planning approach is proposed [22]. On the lower level, each tenant schedules the workloads submitted by its clients applying energy-aware heuristic methods accounting for the main attributed of tasks: length, priority, and deadline. On the upper level a negotiation is performed between the operator and the tenants that use the datacenter, applying a game-theory algorithm. Considering the energy consumed by tenants according to the negotiated monetary incentive, the datacenter operator takes two relevant decisions to reduce energy consumption: lowering or switching off the HVAC system, and using the on-site generation unit. Finally, the cost of the three sources of energy reduction is evaluated, to achieve the committed energy reduction at the lowest possible cost.

The energy required by the HVAC system for cooling the datacenter at time t (H^t) is calculated using a simple formulation that takes into account the area of the servers room (R), the energy consumed by the servers of tenant j (P_j^t), and the coefficient of performance of the HVAC system (COP). The thermal balance must be maintained to keep the temperature of the server room. Hence, the HVAC system must deal with the thermal energy produced by the servers and the thermal energy required to cool the server room. Almost all the energy consumed by servers is dissipated as heat, so P_j^t is an accurate measure of the thermal energy introduced by servers of tenant j. According to the U.S. Department Of Energy "an air conditioner generally needs 20 British Thermal Units (BTU) for each square foot of living space" [23]. Finally, considering that COP is the ratio of thermal energy to energy consumed of a cooling or heating device, higher values of COP indicate lower energy consumption. Hence, each instant the thermal balance must be maintained, according to $H_j^t \times COP = P_j^t + R \times 20$, and $H^t = \sum_j H_j^t$

B. Proposed heuristics

The proposed planning methods for tenants (lower-level) and datacenter operator (higher-level) are described next.

a) Scheduling heuristics for tenants: Three simple heuristics are proposed and evaluated for scheduling on the tenants level. They are based on applying a greedy approach, which takes locally optimal decisions to determine a tasks-to-processor assignment, considering different criteria that accounts for relevant properties of submitted tasks. The criteria are used to sort the list of non-executed tasks in each timestep of the planning period:

- Penalty/Length (PL): the list of non-executed tasks is sorted by the ratio of monetary penalty (MP_j^i) and length (l_j^i). The rationale behind this method is to schedule first those short tasks that impact the most on the planning cost, increasing both the throughput and the revenue.
- Penalty/Deadline (PD): the list of non-executed tasks is sorted by the ratio of monetary penalty (MP_j^i) and deadline (DD_j^i). This method gives priority for execution to those tasks with larger monetary penalty and tighter deadline, increasing the revenue and the QoS provided to clients.
- Deadline/Length (DL): the list of non-executed tasks is sorted by the ratio of deadline (DD_j^i) and length (l_j^i). The main idea of this method is to schedule first those short tasks with lower due dates, increasing the throughput and the QoS provided to clients.

The three proposed low-level heuristics follow an online scheduling approach, i.e. scheduling decisions are taken during the system execution, as soon as new tasks arrive. This approach provides a more realistic view of the datacenter operation than applying a static offline approach. Furthermore, it does not require additional information or estimations about tasks features.

b) Negotiation mechanism for datacenter operator: On the upper level, the datacenter operator applies the negotiation based on Nash equilibrium of the related non-cooperative game, proposed in our previous work [20]. To achieve the requested energy reduction target ρ, a base supply function is used for the negotiation (defined in Eq. 2), where $\Delta P(\Phi_j)$ is the power reduction for tenant j, $\Delta H(\Phi_j(RI_j))$ is the power reduction for the HVAC according to the power reduction $\Delta P(\Phi_j)$, b_j is the tenant offer (*bidding*) for reducing the power consumption by $\Delta P(\Phi_j)$ and p is the market clearing price determined by the operator.

$$\Delta P(\Phi_j(RI_j)) + \Delta H(\Phi_j(RI_j)) = \rho - \frac{b_j}{p} \qquad (2)$$

The negotiation is an iterative process that aims at attaining the Nash equilibrium for the non-cooperative game between the operator and the tenants. Four steps are involved in the negotiation: i) the operator broadcasts the supply function to the tenants; ii) each tenant places its bid b_j for reducing its energy consumption $\Delta P(\Phi_j(RI_j))$ units, taking into account the monetary incentive and the QoS provided to its users; iii) the operator determines the market clearing price p, the HVAC power reduction $\Delta H(\Phi_j(RI_j))$, and the energy to generate on-site, in order to minimize the total cost, as proposed by the problem formulation (1a); and iv) the supply function is updated and the process repeats until convergence. When convergence is reached (considering an error threshold ϵ due to the iterative dynamic of the proposed negotiation), the more convenient of the latest two bids for the datacenter operator are accepted as agreed values for RI_j, tenants commit to the resulting energy reduction and accept to pay the penalty for deferred tasks to their clients.

The main difference with approaches previously proposed in the literature is that tenants obtain all the required information for the negotiation process via explicit simulations, considering a realistic energy consumption model [6] and the possible energy reductions by HVAC operation.

IV. Experimental analysis

This section reports the experimental analysis of the proposed approach for demand response planning of datacenters.

A. Methodology

a) Problem instances: The problem instances used in the experimental evaluation extend the ones introduced in our previous work [24] to consider parameters regarding the cooling system: the datacenter area (*R*), the thermal ratio of HVAC (*COP*), and the operation cost of HVAC.

A total number of 36 problem instances are considered, classified by size and heterogeneity. The classification by size considers the number of tenants of the problem instance; small instances have 5 tenants, medium instances have 10 tenants, and large instances have 30 tenants. Two dimensions are considered for the classification by heterogeneity: tenants size and tenants tolerance. Tenants size is homogeneous if all tenants of the instance have the same size (i.e. equal servers number and workload size), and it is heterogeneous otherwise.

The four combinations of size heterogeneity and tenants heterogeneity generate four types of instances (size heterogeneity-tenants heterogeneity): HM-HM, HM-HT, HT-HM, and HT-HT. HM means homogeneous and HT means heterogeneous.

Workloads (publicly available at https://www.fing.edu.uy/inco/grupos/cecal/hpc/DRAS/) were built considering realistic data from HPC datacenters. Modern servers are considered too. The energy consumption model corresponds to HP ProliantDL380 G9 servers with 2 Intel Xeon E5-2643v3, 24 cores.

Regarding the HVAC system, it is assumed an operation in three levels: a maximum operation level accounting for 70% of the total capacity of the available computing resources, a medium level between 50-70%, and below 50%.

The cost of the fuel for on-site generation is assumed to be four monetary units (e.g., 4 USD) per Watt. The cost of the HVAC is assumed to be 0.4 monetary units per Watt.

b) Computational platform and software: Simulations were performed in a custom simulator built by extending the jMetal framework for optimization [24]. Specific modifications were included to account for the thermal balance calculation to compute ΔH, according to the description in the previous section. All experiments were performed in the high performance computing infrastructure of National Supercomputing Center, Uruguay [16].

c) Metrics and evaluation: The evaluation considers the datacenter operator cost as the main metric for the comparison of the results. Contribution of the main elements in the proposed model (monetary incentives to tenants, cost of the on-site generation, and cost of the HVAC system) are logged and reported. Additionally, two metrics related to the QoS offered to clients are reported and evaluated: number of postponed tasks and the total monetary penalty for violated tasks of all tenants.

A comparative analysis of the three studied scheduling heuristics for tenants is reported. In turn, results of the proposed approach are compared with a BaU operation. Other reference values are considered too, such as the cost when using the on-site generation without electricity load shedding (thus, achieving the best possible QoS levels).

B. Experimental results

The analysis is focused on analyzing the cost values for the datacenter operator, but QoS-realted metrics are also studied. Regarding operation costs, Fig. 2 presents a representative example of the three relevant components of the problem model: energy reduction by tenants (red), energy consumption of the HVAC system (blue), and on-site generation (orange). All these energy consumption sources have an associated cost for the datacenter operation to attend a demand response event. The interesting values for the reduction incentive are between 0.5 and 3.0 monetary units, where the financial equation for tenants allows increasing the energy consumption reduction and both HVAC and on-site generation reduce accordingly.

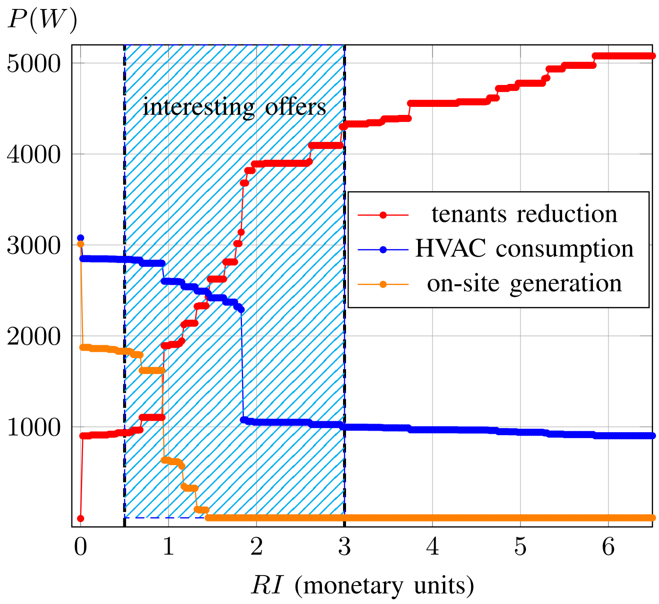

Fig. 2: Example of energy reduction by tenants (red), HVAC consumption (blue), and on-site generation (orange) for different RI values, for a representative problem instance

Table I reports the cost and QoS results of the compared heuristics for the studied problem instances. Values reported correspond to the average total cost and the relative standard deviation for the datacenter operator, as defined by Eq. 1a. In turn, several metrics related to the QoS are reported: the average number of tasks deferred for execution outside of the planning period (DT), the percentage they represent over the total number of tasks in the corresponding workloads ($\%D$), the average penalty cost that tenants must pay to users for deferred tasks (*MP*), and the sum of the average penalty costs weighted by the tenants reputation (CR).

Finally, to establish a baseline for the comparative analysis of the obtained results, two relative values are also reported in Table I: the average cost improvement over a BaU operation, where tenants do not consider any planning for tasks execution and all tasks are executed following a First-Come-First-Served approach (Δ_{BAU}), and the average cost improvement over a planning without load shedding (Δ_{NS}).

Results in Table I indicate that heuristics including the penalty values in the ratio considered as sorting criteria allow computing the better results. Nash+PL and Nash+PD achieved significant improvements over the BaU operation: Nash+PL computed the best cost values and Nash+PD was the second best option. In turn, the heuristic considering deadline and length was ineffective to compute accurate cost values, but it can be useful in specific scenarios when QoS metrics have more importance. Nash+PL scaled properly with the size of the instances, improvements over Nash+PD increased from 2.09% in small instances to 8.65% in large instances. Large values of standard deviation indicate that a wide range of problem instances were considered in each dimension.

The comparison with the BaU operation indicates that significant cost reductions are obtained when properly planning in the lower level; improvements ranged from 54.3% in small instances to 57.4% in large instances. The cost improvement over a datacenter planning without load shedding (Δ_{NS}) are significant too. Improvements of up to 62.0% were computed by Nash+PL, demonstrating the usefulness of the proposed approach as business model for a datacenter participating in the electricity market. Finally, QoS values were reasonable for the proposed approach, especially in small and medium instances; large instances, associated to a more critical demand response event, require deferring a larger number of tasks.

TABLE I: Cost and QoS results for the studied problem instances, averaged by size

Heuristic	*cost*		QoS				Δ_{BAU}	Δ_{NS}
	avg.	*std.*	DT	$\%D$	*MP*	*CR*		
small instances								
Nash+PL	**4625.7**	0.4	4744.9	2.8%	10280.7	6745.7	**54.3%**	**60.4%**
Nash+PD	4724.2	0.4	4803.4	2.8%	10338.4	6793.9	53.3%	59.5%
Nash+DL	10103.4	0.4	761.8	0.4%	6296.8	4426.7	0.1%	13.4%
BaU	10118.3	0.4	772.9	0.4%	6431.6	4503.3	-	-
medium instances								
Nash+PL	**14176.8**	0.3	15608.0	9.0%	33859.7	23266.1	**55.2%**	**59.5%**
Nash+PD	15018.2	0.3	15992.1	9.3%	34212.3	23608.4	52.5%	57.1%
Nash+DL	31803.1	0.2	2418.1	1.4%	20978.7	15237.4	-0.6%	9.1%
BaU	31616.7	0.2	2501.2	1.4%	21636.9	15690.0	-	-
large instances								
Nash+PL	**39426.0**	0.3	46011.6	26.7%	115149.8	80272.6	**57.4%**	**62.0%**
Nash+PD	43159.4	0.2	50143.8	29.1%	119314.3	83034.4	53.3%	58.4%
Nash+DL	92852.4	0.2	9268.1	5.4%	79344.5	58301.8	-0.4%	10.4%
BaU	92491.7	0.2	9453.2	5.5%	80880.6	59303.7	-	-

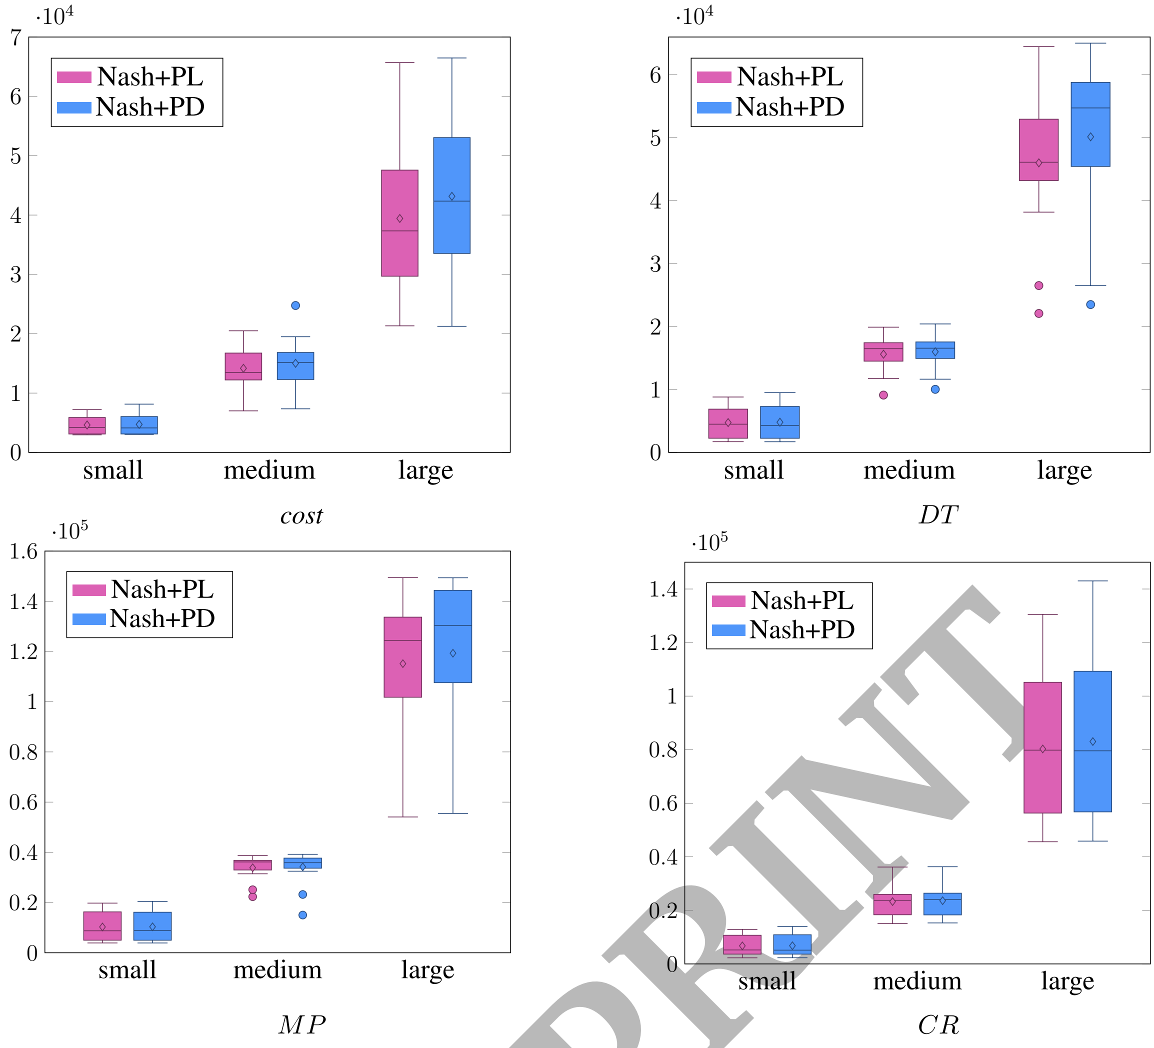

Fig. 3: Statistical analysis of cost and QoS results for Nash+PL and Nash+PD, averaged by scenario size

Boxplots in Fig. 3 compare cost and QoS results for Nash+PL and Nash+PD, the best two heuristics among the studied. Results demonstrate that both heuristics are equally accurate when addressing small- and medium-sized instances. However, Nash+PL is more accurate than Nash+PD for cost and DT when addressing large-sized instances. Nash+PL computes 8.7% lower costs and 8.2% lower DT values than Nash+PD, on average.

A relevant issue to analyze is how heterogeneity is accounted in the problem model, since hosting heterogeneous tenants is one of the main features of the business model in retail multi-tenant datacenters. Table II reports the cost and QoS results of the compared heuristics for different heterogeneity levels for both size of the computing resources and tolerance of tenants.

Results in Table II indicate that high heterogeneity problem instances are harder to solve, mainly because they pose different conditions for low-level scheduling. Nash+PL computed the best results regarding the heterogeneity classification, but differences with Nash+PD reduced in high heterogeneity instances. Nash+PL outperformed Nash+PD in 6.8% regarding the Δ_{BAU} metric in HM-HM instances, but the difference reduced to just 2.1% in HT-HT instances. Regarding QoS, the monetary penalty paid by tenants is not significantly affected by the heterogeneity levels, but when considering the reputation, values of the CR metric significantly reduced for heterogeneous scenarios, indicating that both Nash+PL and Nash+PD are able to properly plan the tasks execution considering heterogeneity, accounting for accurate plannings from the point of view of tenants.

Fig. 4 graphically compares the results of Nash+PL, the best planning heuristic studied, for the four heterogeneity classes in the workloads. Results show the impact of heterogeneity in the results, with high tenants heterogeneity reducing the cost improvements over BaU from 60.7% to 54.8%, and size heterogeneity significantly impacting in QoS related metrics (e.g., improving in up to 22% in MP and 27% in CR metrics when comparing HM-HM and HM-HT instances).

TABLE II: Cost and QoS analysis according to heterogeneity classes

Heuristic	*cost*	QoS			Δ_{BAU}	Δ_{NS}	Heuristic	*cost*	QoS			Δ_{BAU}	Δ_{NS}
		%*D*	*MP*	*CR*					%*D*	*MP*	*CR*		
HM-HM instances							HT-HM instances						
Nash+PL	19824.8	13.5%	58772.6	42473.7	**60.7%**	**64.6%**	Nash+PL	19637.8	11.4%	45902.1	30777.6	**57.7%**	**70.3%**
Nash+PD	23258.7	15.1%	61622.2	45013.4	53.9%	58.5%	Nash+PD	17628.2	11.6%	46342.6	31207.0	55.1%	68.5%
Nash+DL	50612.0	2.8%	40949.8	29925.8	-0.4%	9.6%	Nash+DL	39456.8	2.0%	30242.0	20235.6	-0.6%	29.5%
BaU	50402.0	2.9%	41843.3	30510.4	-	10.0%	BaU	39228.0	2.1%	30762.2	20515.6	-	30.0%
HM-HT instances							HT-HT instances						
Nash+PL	23475.4	14.7%	60899.8	42790.4	**53.2%**	**58.1%**	Nash+PL	17725.4	11.7%	46812.4	31004.1	**54.8%**	**68.3%**
Nash+PD	24465.9	16.1%	63222.0	43905.7	51.2%	56.3%	Nash+PD	18516.4	12.0%	47299.8	31122.7	52.7%	66.9%
Nash+DL	50225.9	2.80%	40678.6	31441.7	-0.1%	10.3%	Nash+DL	39383.9	2.0%	30289.8	22351.4	-0.5%	29.7%
BaU	50151.6	2.8%	41610.0	32117.5	-	10.4%	BaU	39187.3	2.1%	31049.9	22852.5	-	30.0%

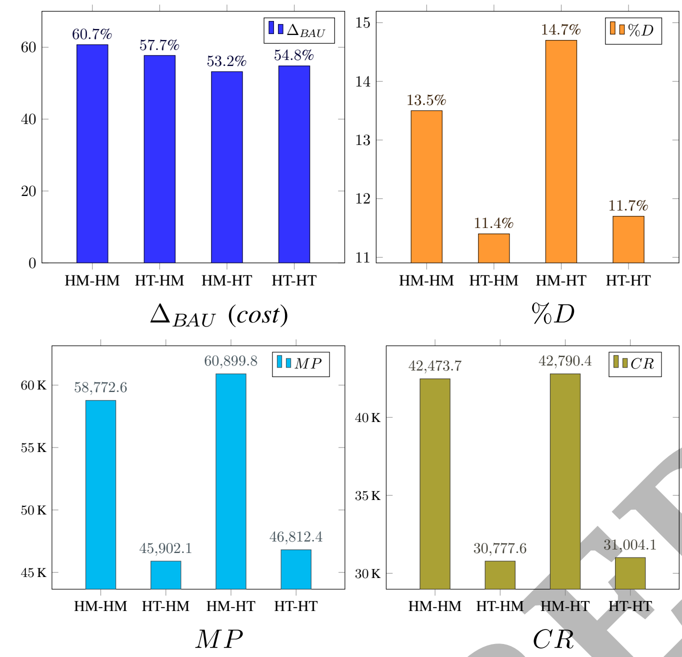

Fig. 4: Analysis of the impact of heterogeneity for the Nash+PL heuristic results

From a global point of view, the demand response planning problem is an inherent multiobjective optimization problem, regarding cost and QoS offered to users [24]. In this sense, trade-off analysis between cost and QoS is relevant to understand the dynamics of the underlying optimization problem.

The multiobjective nature of the datacenter planning problem is also a relevant issue to consider when the market operates in the case of *voluntary* demand response events, which are not usually associated to emergency situations but to profitable load reduction according to given compensations associated to a grid request. Then, both the datacenter operator and the tenants have more elasticity to provide energy reduction and can take advantage of several other situations (e.g., priority events, renewable energy availability, etc.).

Fig. 5 presents a sample trade-off analysis between cost and QoS for a representative problem instance with high tenants heterogeneity (the instance includes five tenants, three small tenants, one medium, and one large). Points C1,..., C5 in the graph represent the cost and QoS for each tenant in the considered instance. QoS is evaluated by the inverse of the cost-reputation metric (CR) in the y axis and the x axis represents the datacenter operator cost. Cost values are identical for all tenants of a given solution, since the problem model assumes price-taking tenants and the same offer is broadcasted to all tenants.

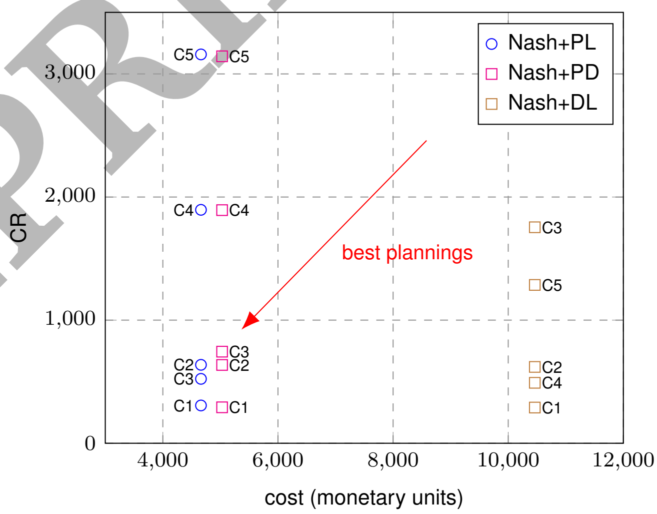

Fig. 5: Sample trade-off analysis for a representative scenario with high tenants heterogeneity

Results in Fig. 5 indicate that the studied heuristics provide different trade-off values between datacenter operation cost and the QoS evaluated by the inverse of the CR metric. Nash+PL heuristics allowed computing the best plannings, The heterogeneous nature of tenants allows them to provide different contributions to energy reduction, as accounted for the CR values.

V. Conclusions and Future Work

This article addressed the problem of demand response planning for retail multi-tenant datacenters, a relevant issue to contribute with the development of improved electric network within the modern smart grid paradigm.

A two level planning approach was proposed for the participation of datacenters in the electricity market, by providing appropriate actions to demand response events. The planning strategy considers the heterogeneity of retail collocation datacenters to account for an improved execution plan to reduce energy consumption as requested by the electricity market.

The two level approach applies a game theory negotiation at the (upper) datacenter level and scheduling heuristics at the (lower) tenants level. In the lower level, three heuristics are evaluated, accounting for different features of tasks submitted for execution (penalty, deadlines, and task length).

The experimental evaluation of the proposed approach is performed over a set of realistic scenarios and problem instances including the main features of nowadays datacenters (computing infrastructure and air conditioning system) and workloads build using real data. The main results indicate that the proposed planning methods are able to compute accurate solution to the problem. The heuristic applying the ratio between priority and length to sort the list of non-executed tasks (Nash+PL) was able to compute the best results regarding cost, while accounting for appropriate quality of service. Different trade-off levels are computed by the proposed planing methods, which are a useful input for decision making at the datacenter operator level.

Results suggest that the proposed approach provide appropriate plannings for demand response events, significantly improving over Business-as-Usual operation and the situation where there is not load shedding. These results demonstrate the viability of the proposed approach for demand response planning in retail multi-tenant datacenters.

The main lines for future work are related to extend the experimental evaluation of the proposed approach by considering different demand response events and even on-line actions at micro temporal levels. A more realistic approach can be applied to model the thermal control of a realistic datacenter facility, possibly by applying computational intelligence techniques to learn an empirical distribution of the temperature values for different servers utilization levels. Finally, different approaches using local and global information can be applied at the datacenter operator level to propose more powerful planning methods, e.g., using metaheuristic algorithms.

Acknowledgment

The research is part of project "Demand response and ancillary services for supercomputing and datacenters", supported by ANII, Uruguay. The work of S. Nesmachnow and S. Iturriaga is partly supported by ANII and PEDECIBA.

References

[1] J. Momoh, *Smart Grid: Fundamentals of Design and Analysis*. Wiley-IEEE Press, 2012.

[2] P. Andersen, E. Hauksson, A. Pedersen, D. Gantenbein, B. Jansen, C. Andersen, and J. Dall, *Smart Grid Applications, Communications, and Security*. Germany: Wiley-VCH, 2012.

[3] Y. Sverdlik, "How cloud is changing the colocation data center market," January 19, 2017, www.datacenterknowledge.com/ (June 2020).

[4] P. Delforge and J. Whitney, "Scaling up energy efficiency across the data center industry: Evaluating key drivers and barriers," Natural Resources Defense Council and Anthesis, Tech. Rep., 2014.

[5] J. Muraña, S. Nesmachnow, S. Iturriaga, and A. Tchernykh, "Power consumption characterization of synthetic benchmarks in multicores," in *High Performance Computing*, 2017, pp. 21–37.

[6] J. Muraña, S. Nesmachnow, F. Armenta, and A. Tchernykh, "Characterization, modeling and scheduling of power consumption of scientific computing applications in multicores," *Cluster Computing*, vol. 22, no. 3, pp. 839–859, Sep 2019.

[7] P. Klemperer and M. Meyer, "Supply function equilibria in oligopoly under uncertainty," *Econometrica*, vol. 57, no. 6, pp. 1243–1277, 1989.

[8] S. Tijs and T. Driessen, "Game theory and cost allocation problems," *Management science*, vol. 32, no. 8, pp. 1015–1028, 1986.

[9] R. Johari and J. Tsitsiklis, "Parameterized supply function bidding: Equilibrium and efficiency," *Operations Research*, vol. 59, no. 5, pp. 1079–1089, 2011.

[10] N. Chen, X. Ren, S. Ren, and A. Wierman, "Greening multi-tenant data center demand response," *Performance Evaluation*, vol. 91, pp. 229–254, 2015.

[11] N. Tran, C. Do, S. Ren, Z. Han, and C. Hong, "Incentive mechanisms for economic and emergency demand responses of colocation datacenters," *IEEE Journal on Selected Areas in Communications*, vol. 33, no. 12, pp. 2892–2905, 2015.

[12] Y. Wang, F. Zhang, C. Chi, S. Ren, F. Liu, R. Wang, and Z. Liu, "A market-oriented incentive mechanism for emergency demand response in colocation data centers," *Sustainable Computing: Informatics and Systems*, vol. 22, pp. 13–25, 2019.

[13] S. Bahrami, V. Wong, and J. Huang, "Data center demand response in deregulated electricity markets," *IEEE Transactions on Smart Grid*, vol. 10, no. 3, pp. 2820–2832, 2019.

[14] B. Dorronsoro, S. Nesmachnow, J. Taheri, A. Zomaya, E.-G. Talbi, and P. Bouvry, "A hierarchical approach for energy-efficient scheduling of large workloads in multicore distributed systems," *Sustainable Computing: Informatics and Systems*, vol. 4, no. 4, pp. 252–261, 2014.

[15] S. Iturriaga and S. Nesmachnow, "Scheduling energy efficient data centers using renewable energy," *Electronics*, vol. 5, no. 4, p. 71, 2016.

[16] S. Nesmachnow and S. Iturriaga, "Cluster-UY: scientific HPC in Uruguay," in *International Supercomputing in México*, 2019, pp. 1–15.

[17] S. Nesmachnow, C. Perfumo, and Í. Goiri, "Holistic multiobjective planning of datacenters powered by renewable energy," *Cluster Computing*, vol. 18, no. 4, pp. 1379–1397, 2015.

[18] F. Armenta, A. Tchernykh, J. Cortes, R. Yahyapour, A. Drozdov, P. Bouvry, D. Kliazovich, A. Avetisyan, and S. Nesmachnow, "Min_c: Heterogeneous concentration policy for energy-aware scheduling of jobs with resource contention," *Programming and Computer Software*, vol. 43, no. 3, pp. 204–215, 2017.

[19] S. Nesmachnow, "An overview of metaheuristics: accurate and efficient methods for optimisation," *International Journal of Metaheuristics*, vol. 3, no. 4, pp. 320–347, 2014.

[20] J. Muraña, S. Nesmachnow, S. Iturriaga, S. M. de Oca, G. Belcredi, P. Monzón, V. Shepelev, and A. Tchernykh, "Negotiation approach for the participation of datacenters and supercomputing facilities in smart electricity markets," *Programming and Computer Software*, 2020.

[21] S. Montes de Oca, J. Murana, P. Monzón, S. Nesmachnow, S. Iturriaga, and G. Belcredi, "Demand response program for supercomputing and datacenters providing ancillary services in the electricity market," in *IEEE PES Transmission and Distribution Conference*, (submitted).

[22] A. Tchernykh, J. Ramírez, A. Avetisyan, N. Kuzjurin, D. Grushin, and S. Zhuk, "Two level job-scheduling strategies for a computational grid," in *Parallel Processing and Applied Mathematics*, 2006, pp. 774–781.

[23] U.S. Department of Energy, Office of Energy Efficiency and Renewable Energy, "Room air conditioners," https://www.energy.gov/energysaver/room-air-conditioners (June 2020).

[24] J. Muraña and S. Nesmachnow, "Simulation and evaluation of multicriteria planning heuristics for demand response in datacenters," *Simulation*, 2020, (submitted).