ПРОГРАММИРОВАНИЕ, 2020, № ?, с. 1–21

УДК 004.421.6

Negotiation approach for the participation of datacenters and supercomputing facilities in smart electricity markets

© 2020 г. **Jonathan Muraña*, Sergio Nesmachnow*‡, Santiago Iturriaga*, Sebastián Montes de Oca*, Gonzalo Belcredi*, Pablo Monzón*, Vladimir Shepelev‡, Andrei Tchernykh†‡**

* *Universidad de la República, Uruguay,* ‡ *South Ural State University, Russia,* † *CICESE, México*

E-mail:{jmurana,sergion,siturria,smontes,gbelcredi,monzon}@fing.edu.uy
v.shepelevvd@susu.ru, chernykh@cicese.mx

Поступила в редакцию 30.01.2020

This article presents an approach for the participation of datacenters and supercomputing facilities in smart electricity markets. This is a relevant problem in modern smart grid systems to implement demand response strategies for a better use of resources to guarantee energy efficiency. The proposed approach includes a datacenter model based on empirical information to determine the power consumption of CPU-intensive and memory-intensive tasks. A negotiation approach between the datacenter and its clients and a heuristic planning method for energy reduction optimization are proposed. The experimental evaluation is performed over realistic problem instances modeling the operation of the National Supercomputing Center in Uruguay. The obtained results indicate that the proposed approach is effective to provide appropriate demand response actions according to monetary incentives. Accurate results are reported for realistic problem instances and different type of clients.

1 Introduction

Smart electricity network, or *smart grid*, refers to a electrical grid that includes operation and management features to improve the controlling of production and distribution of energy [21]. Smart grids are the current state-of-the-art technology for electricity networks, the last step in their evolution from unidirectional systems of electric power transmission and distribution to holistic approaches that provide different services for demand-driven control. The main goal of smart grids is to maintain a reliable and secure infrastructure to properly satisfy the demand growth, the integration of distributed energy resources generation and storage, and several other features related to smart devices and real-time information provided to clients [11]. Information and Communication Technologies have provided a key foundation for communicating and processing information that is very useful at different levels to implement the aforementioned services [10].

Within the smart grid paradigm, a large consumer with flexible power utilization can participate in the electricity market. This is one of the main ideas behind the implementation of strategies oriented to modern smart electric networks, where consumers are associated to the roles of both active clients and market agents [21]. On the one hand, as an active client, a consumer can adapt his electricity demand to peak hours, e.g., by reducing power consumption in peak periods and contributing to flattening the demand curve of the whole electrical system. Paradigms and strategies applying multi-hour tariffs can also be implemented, handling time periods where it is preferable that consumers use energy. On the other hand, as a market agent, a consumer can participate in the electricity market and receive an income by providing different services. Several mechanisms can be applied in this scenario, including establishing bilateral agreements with an electricity generation company or by participating in periodic auctions for smart grid management.

Wholesale electricity markets were designed to meet short-term and future requirements of operating the electric power system reliably and at the lowest cost. Policy makers saw competition among suppliers as a mean to control pricing by attracting new sources and technologies from the private sector in an open, competitive, and transparent market. The pricing system facilitates using the lowest-cost generation options, i.e., by encouraging suppliers and consumers to buy and sell electricity at the lowest prices while still ensuring reliability. The wholesale market is structured in several sub-markets, with different horizon times in advance of the electricity purchase. Markets range from a few seconds or minutes for ancillary services, correcting mismatches between generation and demand, to a few year in advance for the capacity market (assuring the capacity of the future demand), including real-time and day-ahead markets for the energy purchase [35].

In any case, within the smart grid paradigm, 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 plan those activities that consume energy, i.e. by advancing or deferring their execution. In addition, the impact on global energy efficiency, and the possible degradation of QoS offered to users must be analyzed. These planning techniques are essential to ensure the correct use of energy resources and to guarantee the energy efficiency of large flexible consumers.

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 relevant case study, the planning of datacenters and supercomputing infrastructures is addressed. These platforms are conceived as examples of planned systems that have emerged in modern societies, linked to the smart grid paradigm. Datacenter and supercomputing facilities can adjust power consumption to help the electric network to fulfill specific goals: they are able to consume available surplus of energy by executing complex tasks that demand large execution times, or they can defer activities (i.e., tasks execution) in periods where energy is more expensive and/or power generation is lower than normal. Furthermore, they thermal/cooling infrastructures demand significant energy consumption and provide a large inertia. Thus, they can be used to interact with a smart electric grid.

While existing studies show promising benefits and progress of the datacenters participation on energy markets, most of them focus on owner-operated datacenter, whose operator have full control over scheduling task, servers and facilities [31]. However, this type of datacenter represent a minor percentage of datacenters power consumption (e.g., in USA, it is less than 10%), while *colocation datacenters* represent a significantly larger contribution (almost 40% in USA) and the rest correspond to enterprise in-house datacenters [5]. This article focuses on multi-tenants colocation datacenters. In this type of datacenters, multiple tenants deploy and keep full control of their own physical servers in a shared space, while the datacenter operator provides facility support (e.g. high-availability power and cooling). In addition to consuming a significant amount of energy, workloads from tenants in colocation datacenter are highly heterogeneous, and many tenants run non-critical workloads which has high scheduling flexibility, different delay sensitivities, different service level agreements with peak loads periods. Thus, colocation datacenters can participate actively in the energy market by modulating their power profile and helping maximize distribution grid resources. The potential of shedding the load of tenants is ideal for demand response participation. The main disadvantage is that each tenant manages its own servers independently and have very different incentives to participate or cooperate with the operator during a demand response event.

In this line of work, the research reported in this article is based on a negotiation using a pricing mechanism that guarantees that datacenter operators can extract load shedding from tenants. This market mechanism can effectively provide incentives for tenants to reduce energy consumption during demand response events, complementing, and even substituting, the high-cost of diesel generation. Furthermore, this price mechanism is modeled as a negotiation [27, 16]. The evolution of the client reduction functions and the trade-off of datacenter costs and QoS is analyzed.

The proposed strategy for demand response allow implementing a smart management of the electric grid, achieving a rational utilization of energy sources, and the correct use of information technologies to improve decision-making processes within modern smart grids.

The main contributions of the research reported in this article are summarized next. Strategies for smart planning of tasks execution and management of energy utilization are proposed for the National Supercomputing Center in Uruguay (Cluster-UY) [25], taking into account the energy consumption and the QoS provided to users. A specific negotiation approach for the participation of colocation datacenters in the ancillary services market, based on a utility function for the tenants and a heuristic for energy consumption optimization. A specific feature of the proposed model and algorithms is that, unlike other works in the related literature, real tasks data and real energy consumption evaluation are considered for the planning instead of using theoretical models. The experimental evaluation is performed through simulations that consider realistic workloads, high-end servers, and a power consumption model built from real data. Results suggest the effectiveness of the proposed strategies to implement demand response techniques and provide ancillary services under the smart grid paradigm. Participants achieve a significant reduction of their operative power consumption, keeping their profit in positive values. Furthermore, the proposed non-cooperative negotiation shows consistent results in all studied scenarios, which allow efficient energy management and business viability at different problem scales.

The article is structured as follows. The description of the proposed model to characterize the energy consumption for datacenters and supercomputing facilities is presented in Section 2. Section 3 describes the problem addressed in this article and a review of related works. The proposed algorithms are introduced in Section 4, including the non-cooperative negotiation algorithm and a greedy heuristic for optimizing the energy utilization. The experimental evaluation over realistic problem instances is reported in Section 5. Finally, Section 6 presents the conclusions and the main lines for future work.

2 Market behavior and energy consumption modelling

This section presents the main concepts about the applied models for simulating the market behavior and the energy consumption in datacenters and supercomputing facilities.

2.1 Market modeling

In open deregulated electricity markets, demand response is considered a supply-side resource in capacity and ancillary services. In particular, in Pennsylvania-New Jersey-Maryland Interconnection (PJM) electricity market, demand response services are concentrated in the synchronize reserves and the capacity markets through the emergency and economic load program. In other electricity markets (e.g., UK or Germany) there are similar demand response programs for supply-side resource, mainly focusing on frequency restoration or for demand turn up services [34, 2]. These approaches are effective for customers to manage their electricity costs and for the wholesale market to incorporate load-reduction actions. From the perspective of the grid operator, the flexible power demand of datacenters serves as a valuable energy buffer, helping balance the grid power supply and demand at runtime.

This article focuses on the PJM electricity market, particularly on the synchronize reserve market and the emergency load program. In both cases, if the grid operator anticipates an emergency (e.g., wrong forecast of demand, extreme weather, or a generation unit out of work), participants are notified, usually at least 10 minutes in advance, and obliged to fulfill their contracted amounts of energy reduction during the event, which may span a few minutes to a few hours.

The main differences between the two programs are the time horizon of the offers and the mechanism to participate. For the emergency and economic load program, participants typically sign contracts with a load serving entity in advance (e.g., three years ahead in PJM) [4, 6] and receive financial rebates for their committed energy reduction even if no emergency demand response signals are triggered during the participation year, whereas non-compliance (i.e., failure to cut load as required during an emergency demand response) incurs a heavy penalty. In the synchronized reserve market, the demand resource participates in a real time market through auctions for each time gap of the market, defining in the offer the price and the capacity to be reduced if necessary. If the offer is selected, the participant is paid for availability and in case of an event, for the energy retired from the grid. In this market, the price of primary reserves results from the demand curve for primary reserves and the supply of primary reserves.

In case of an emergency event is triggered, the datacenter operator can reduce load in response to a demand response event either through extracting IT energy reductions from the tenants (re-scheduling workflow) or by turning on an on-site generator. Since the mandatory demand response reduction target is fixed, the operator must balance between paying tenants for reduction and using on-site generation in order to minimize the datacenter operator cost.

2.2 Energy consumption model for datacenters and supercomputing

The datacenter follows a colocation model with a set of tenants $C = \{c_1, ..., c_{|C|}\}$ where each tenant owns a subset of the total computing resources of the datacenter. Computing resources of any single tenant are considered to be homogeneous in the proposed model, however resources are heterogeneous when considering multiple tenants. The energy consumed by the datacenter is determined by a set of workload schedules $WS = \{ws_1, ..., ws_{|C|}\}$, one for each tenant. Fig. 1 presents a schema of the proposed model.

Each workload schedule ws_j is a tenant-level scheduling of the computing workload of tenant c_j. The algorithm for computing each schedule is tenant-specific and solves an underlying optimization to minimize energy budget, violation of due dates and execution time, among other goals. This article considers the computing workload of each tenant to be comprised of a set of independent tasks with due dates, which indicate the expected or desired completion time. Two types of due dates are modeled: deferrable and non-deferrable. A task with a deferrable due date is a flexible task that may violate its due date and run later, subject to specific tenant-related business conditions. A task with a non-deferrable due date is a time-critical task that must finish its execution within its due date or not execute at all. Each task is owned by a tenant and must be executed in the computing resources owned by that tenant.

Regarding the power consumption model, a computing resource of tenant c_j consumes SE_j^{max} at maximum utilization, SE_j^{idle} when idle, and consumes no power when off. A computing resource of tenant c_j executing a set Ω of tasks at time t consumes an amount of energy given by the function $\sigma_j^t(\Omega)$ with $SE_j^{idle} \leq \sigma_j^t(\Omega) \leq SE_j^{max}$. In this article, function σ_j^t is defined empirically by applying a polynomial regression technique and considering real-world computing resources with two types of tasks: CPU-intensive tasks and memory-intensive tasks. The energy consumed by the computing resources of tenant c_j at time t for schedule ws_j is TP_j^t. Hence, the total energy consumption of the computing resources at time t for the whole datacenter is defined by Eq. 2.1.

$$DP^t = \sum_{j=1}^{|C|} TP_j^t \tag{2.1}$$

When the electric market operator requests the datacenter to reduce its energy consumption by a target Φ, the datacenter must initiate a negotiation phase with its tenants. During this negotiation, the datacenter offers a monetary incentive per each unit of energy reduced (RI). Considering this incentive, each tenant c_j may choose to modify its planned scheduling ws_j and compute a new schedule $\overline{WS}_j$ by postponing or even cancelling the execution of some of its tasks to reduce its energy consumption.

In addition, an on-site energy (e.g., fossil fuel) generator is considered, with a monetary cost β per unit of energy generated. This generator can be used to reduce the energy the datacenter consumes from the grid, in case the energy reduction from its tenants is not enough to meet the reduction target from the electric market. The on-site generator is controlled by the datacenter owner and it generates GP^t energy at time t. Hence, the energy consumed from the grid by the datacenter at time t is the difference between the energy consumed by the computing resources and the energy generated by the on-site generator: $P^t = DP^t - GP^t$.

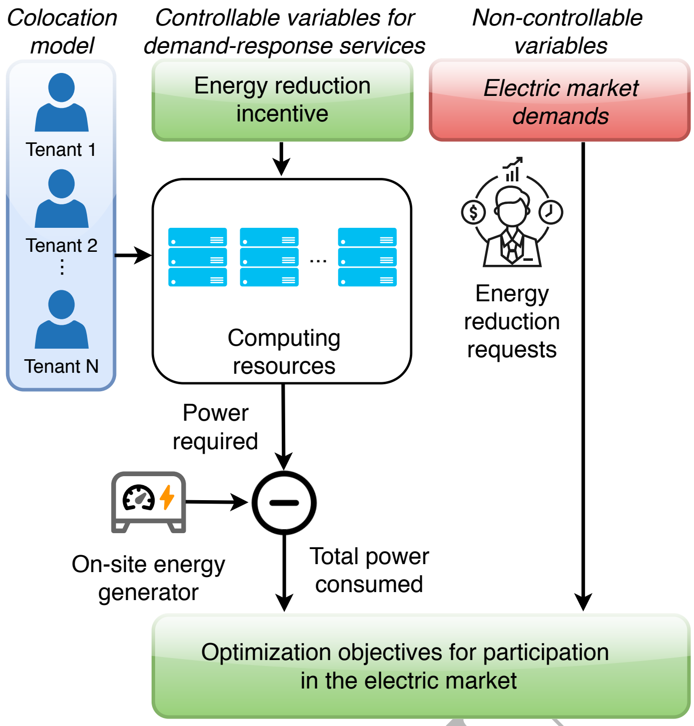

Fig. 1.: Schema of the proposed model for energy consumption in datacenters and supercomputing facilities

3 Scheduling problem for datacenters and related work

This section introduces the formulation of the scheduling problem formulation for the datacenter and a review of related work about strategies for the participation of datacenters in the energy market.

3.1 Scheduling problem formulation

Given a reduction request from the electric market, the optimization problem consists in minimizing the total monetary incentive rewarded to tenants and the cost of using the on-site generator in order meet the reduction target. The energy reduction must maintained for a time horizon T. be The formulation is as follows.

- A set of discrete timesteps t comprising the time horizon T.
- A target reduction of β requested by the electric market.
- A set of tenants (or clients), $C = \{c_1, ..., c_{|C|}\}$.
- A workload of tasks for each tenant c_j, $W_j = \{w_j^1, ..., w_j^{|W|}\}$.
- Let $DF_j^i = 1$ if task w_j^i is deferrable and $DF_i^j = 0$ if its non-deferrable.
- Let DD_j^i be the due date of task w_j^i.
- Let MP_j^i be the monetary penalty received by tenant c_j if the due date of task w_j^i is not met.
- Let RI be the monetary incentive for each tenant for each energy unit reduced.
- A workload schedule for each tenant c_j with no monetary incentives (i.e. $RI = 0$), $WS_j = \{ws_j, ..., ws_{|C|}\}$.
- Let DP_j^t be the power requirement at each time t of each workload schedule ws_j.
- Let FT_j^i be the finishing time of task w_j^i for schedule ws_j.
- Let $VD_j^i = 0$ if the $FT_j^i \leq DD_j^i$ for schedule ws_j, otherwise $VD_j^i = 1$.
- Let define the total monetary penalty of a schedule ws_j for a tenant c_j as $\Upsilon_j = \sum_{i=1...|W|} VD_j^i \times P_j^i$.

- Let the function γ_j determine the new schedule $\overline{ws}_j$ with a power requirement $\overline{DP}_j^t$ for tenant c_j given incentive RI, $\gamma_j(RI) = \overline{ws}_j$.
- Let define the energy reduction function between $\overline{ws}_j$ with respect to the schedule ws_j as $\delta(\overline{ws}_j) = \min_{\forall t \in T}(\overline{DP}_j^t - DP_j^t)$.
- Let GP^t be the energy generated using the on-site generator at each time t.
- Let α be the monetary cost per unit of energy of using the on-site generator.

Formally, the problem is defined by Eq. 3.1.

$$\min z = \sum_{j=1\ldots|C|} \delta(\gamma_j(RI)) \times RI + \sum_{t \in T} GP^t \times \alpha \quad (3.1a)$$

subject to:

$$\beta \leq \sum_{j=1\ldots|C|} \delta(\gamma_j(RI)) + \sum_{t \in T} GP^t \quad (3.1b)$$

$$z \leq \sum_{t \in T} \overline{DP}^t \times \alpha \quad (3.1c)$$

The objective (3.1a) is to minimize cost for the datacenter operator (i.e., the money paid to tenants plus the cost of using the on-site generator) in order meet the reduction target. Constraint (3.1b) states the total energy reduction must be at least β per timestep. Finally, constraint (3.1c) indicates that the total monetary incentive to be paid to tenants must be less than the cost of powering the whole datacenter using the on-site generator.

Formally, the multicriteria planning problem for demand response in datacenters proposes the simultaneous optimization of the objective functions defined by Equations **??**–**??**.

3.2 Related work

Regarding market mechanisms for demand response, they usually are variants of the classical economic problem of allocating joint costs, which arise in many situations and can be tackled through the use of supply functions [17, 37]. Particularly, Johari and Tsitsiklis [16] considered a special type of supply function for a problem where a finite number of producers compete to meet an infinitely divisible but inelastic demand reduction. The authors showed that the Nash equilibrium can be obtained as the solution of an allocation problem.

Chen et al. [27] applied the previous work to obtain a market mechanism for datacenters where the operator can reduce power consumption from clients, without using on-site generation. The goal of the mechanism is to induce the tenants to reduce their energy consumption during emergency demand response events. When an emergency demand response arrives, tenants bid an energy consumption reduction using a parameterized supply function. This is a simple scheme for tenants to participate in the market, since they only need to bid for a reward on energy reduction but keeping its cost function private. Following these ideas, Tran et al. [38] addressed the problem of promoting tenants in colocation datacenters to participate in demand response programs. The authors proposed two mechanisms based on economic rewards and bidding games, considering elastic and fixed energy reduction targets. The workload of each tenant is considered to be comprised of short-lived Internet requests, hence individual requests are not explicitly modelled and queuing theory is applied to tackle the problem. Servers are turned off when energy reduction is necessary and SLA-related economic penalties are considered for each tenant according to the average response time of their workloads. Scenarios were synthetically generated using two well-known real-world traces and simulation results validate the effectiveness of the proposed algorithms. Later, Tran et al. [39] extended their previous work by considering multiple colocation datacenters and modeling the negotiation between the datacenters and the energy provider.

Wang et al. [40] addressed a problem similar to the fixed-reduction problem presented by Tran et al. [38]. Both works considered servers to be homogeneous and do not model individual requests explicitly. However, Wang et al. considered a cloud-oriented scenario where servers may be shared to execute virtual machines from different tenants. Based on this scenario, the authors proposed an incentive mechanism that encourages coordination between the tenants themselves in order to further reduce energy consumption. Simulations based on real-world traces showed that the mechanism proposed by Wang et al. is able to meet the desired energy reduction target at lower monetary costs when comparing to state of the art mechanisms.

Our group has developed research on the holistic energy-aware planning of datacenters, also including the use of renewable energy sources [7, 15, 25, 26]. The general approach consists in applying computational intelligence methods [24] to solve the underlying optimization problem that proposes determining the tasks schedule and the energy consumption of both infrastructure and cooling systems, subject to QoS and operation (e.g., temperature) constraints. Another line of research was the development of strategies for the participation of consumers in energy markets [29, 32], showing how an active consumer may obtain profits offering demand response services. The particular role of an intermediate agent that concentrates consumers and interacts with the system operator was analyzed in our previous article [28].

In this line of work, this article extends our previous research [15, 26] by including a market mechanism for the active participation of tenants, using a realistic evaluation of power consumption of nowadays high performance computing servers [22].

4 The proposed negotiation approach for participation in a smart electricity market

This section introduces the proposed approach for the participation of datacenters and supercomputing facilities in smart electricity markets.

4.1 Overall description

The proposed approach considers a negotiation between the datacenter administrator and the clients, in order to fulfill some specific goals regarding power consumption of the computational infrastructure.

Clients are assumed to be focused on executing scientific applications, which are the ones that demand significant energy consumption [26]. Applications are modeled as *computing tasks*. Two types of tasks are considered: CPU-intensive and memory-intensive, which accounts for the most common types of scientific applications, according to the related literature [8, 12, 22].

The approach considers a two-level procedure for negotiation and power consumption management. In the higher level of abstraction, a negotiation is modeled between the datacenter administrator and each of its clients. In a lower level, each client applies a specific heuristic to analyze if it is able to execute its tasks considering power consumption reductions, according to specific monetary offers by the datacenter administrator. Decisions are taken by considering specific Service Level Agreements (SLA) between the datacenter clients and their users, which is modeled by the deadline of the submitted tasks for execution. The main components of the proposed approach are described in the next subsection.

4.2 Datacenter administration procedure

According to the electricity market model explained in Section 2, the grid operator offers a monetary incentive to the datacenter administrator for reducing a certain amount of energy from its consumption during a certain time. To achieve the necessary reduction, the next market mechanism was implemented.

The proposed market mechanism for colocation datacenters follows the main idea from Chen et al. [27], characterizing the Nash optimum of the non-cooperative game as an optimization problem known as allocation problem [37]. In this approach, the operator can induce a reduction on client's power consumption diminishing the need of brown energy, using a parameterized supply function represented in Eq. 4.1, where r_i is the power reduction for client i, D is data center's power reduction target, b_i is the client offer for reducing the power consumption by r_i and p is the market clearing price determined by the operator.

$$r_i(b_i, p) = D - \frac{b_i}{p} \tag{4.1}$$

The market mechanism for reducing D amount of energy is exercised in four steps in an iterative approach:

(i) The datacenter broadcasts the supply function to the clients, $r_i(b_i, p) = D - \frac{b_i}{p}$.

(ii) Each client i bids a reward b_i for reducing r_i amount of power, in order to maximize its utility and can be interpreted as the IT revenue that client is willing to give up.

(iii) The datacenter determines the market clearing price p and the amount of energy to produce via on-site generation y (with generation cost α) by minimizing the total cost, represented by the cost of generation and the rewards payed to the tenants.

$$p(b_i, y) = \frac{\sum_i b_i}{(N-1)D + y} \tag{4.2}$$

$$y = \underset{0 \leq y \leq D}{\arg\min} (D - y)p + \alpha y \tag{4.3}$$

The first-order optimality condition for Eq. 4.3 gives the value for y:

$$y = \sqrt{\frac{(\sum_{i=1}^{N} b_i) N D}{\alpha}} - (N-1)D \tag{4.4}$$

(iv) If p and y converges, latest bids are accepted and energy reduction is scheduled by the clients, else the operator broadcast the new supply function with the updated value for p.

Algorithm 1 describes the strategy used by the datacenter, by implementing a solution of the allocation problem based on a proximal method [33]. A distributed solution is generated for each agent. These solutions are coupled by the power balance equation $D = \sum_{i=1}^{N} reduction[i] + y_k$. This equality constraint is relaxed in the proximal method. In the algorithm, D is the power reduction target, *price* is the market clearing price per Watt, N is the number of tenants and j is the tenant id. The function $client\_evaluation(price, j)$ corresponds to the offer evaluation of the tenant j, considering its SLAs. This function returns the energy reduction committed by the tenant ($reduction[j]$), according to the price, $bid[j]$ is the offer of tenant j for reducing the power consumption, y_k is the iteration variable, which at the end of the negotiation corresponds to the power generated by the on-site generator. The cost of generate one Watt using the generator is denoted α. The parameter ϵ is a measure of the compliance of the coupling restriction.

Algorithm 1 Datacenter market mechanism

INPUT: D (power reduction target), $price_0$
OUTPUT: *price*, *on-site-generation*

$k \leftarrow 0$ ▷ iteration step
$price_k \leftarrow price_0$
while $\epsilon \geq \epsilon_{min}$ **do**
  for j=1 to N **do**
    $reduction[j] \leftarrow client\_evaluation(price, j)$
    $bid[j] \leftarrow (D - reduction[j]) \times price_k$
  end for
  $y_k \leftarrow \max(\sqrt{(\sum bid) N.D/\alpha} - (N-1)D, 0)$
  $price_k \leftarrow \sum_j bid/((N-1)D + y_k)$
  $\epsilon \leftarrow \|(y_k + \sum_j reduction - D)/D\|$
  $k \leftarrow k + 1$
end while
$on\text{-}site\_generation \leftarrow y_k$

4.3 Client scheduling and optimization

This subsection describes the procedure performed by each client to implement the corresponding power consumption reductions according to the offers from the datacenter administration.

Client offer evaluation. To evaluate the monetary offer of the datacenter administrator and determine the amount of power to be reduced, clients simulate the execution of their workload, applying an energy optimization strategy. The monetary offer of the datacenter administrator is accepted if the net income obtained from the energy reduction minus the loss the client must pay in case of not complying with the SLA with his users, is greater than zero. In any case, different trade-offs are obtained for different monetary offers from the negotiation. These trade-offs can be considered in case the datacenter cannot meet the desired power consumption reduction, to account for different compromises between the problem objectives (energy reduction and cost).

The energy optimization strategy proposed in this article aims at maximizing the profit of clients by reducing the active cores of the servers, thus lowering the energy consumption according to the offer received from the datacenter administration. A heuristic procedure is applied: Active Cores Reduction (ACR) whose main idea is to select the best scheduling according to its profit, considering all combinations of active cores. The details of the ACR heuristic are presented in Algorithm 2.

Algorithm 2 Energy optimization strategy

INPUT: *price* **OUTPUT:** *reduction*

```
profit ← 0
reduction ← 0
cores ← server_cores × servers_number
for cores_number in cores do
    sol ← schedule(cores_number)
    reduction_ax, profit_ax ← evaluate(price, sol)
    if profit_ax > profit then
        reduction ← reduction_ax
        profit ← profit_ax
    end if
end for
```

In Algorithm 2, *price* is the offer per reduced Watt and *reduction* is the amount of watts that the client is willing to reduce according the offer. Function *schedule* simulates the execution of the workload considering *cores_number* active cores, out of the *server_cores* total cores available in the client. In turn, function *eval* evaluates the profit and the reduction of the solution scheduling *sol*.

Client scheduling simulation. The considered clients are providers of high performance computing services to single users. In this type of services, batch tasks arrive to the system and they are queued until a server has the capacity to execute it. A server has capacity to execute a task if it meets the task requirements, such as available cores, memory, and estimated execution time. Other task requirements of high performance computing systems are related to the distribution of assigned cores on the server, for example, all cores required by a specific task must be on the same server. Tasks included in the workloads considered in this article are assumed to require only one core to be completed and they have no specific memory requirements (i.e., they can execute with the usual memory available per core on each server, at least 1–3 GB). Servers are considered to be high performance multicore architectures, such as the ones available in nowadays datacenters and supercomputing facilities [25].

A simulation-driven approach is applied to determine the cost of implementing a certain energy optimization strategy. An ad-hoc simulator is utilized due to the limited capabilities of existing datacenter/cloud simulators to provide an accurate environment for implementing the main features of the proposed approach. Two well-known simulators were evaluated: CloudSim [3] and DCworms [19]. CloudSim lacks the flexibility to model the real operation of datacenters and HPC infrastructures, since it is oriented to model Cloud systems. DCworms is in fact oriented to model HPC infrastructures. However, DCworms does not allow to easily include a temperature model, necessary for our future work. Furthermore, both simulators offer poorly-documented application interfaces, thus they require a significant effort to implement custom-made heuristics and models. This is a major drawback in the context of the research reported in this article, where a set of novel realistic power consumption models are proposed and evaluated. For the aforementioned reasons, a new ad-hoc simulator was implemented. The main features of the new simulator include a simple design, strongly oriented to the operation of HPC systems, and its flexibility for implementing complex scheduling strategies, allowing to consider variables such as energy consumption, temperature and QoS. In addition, the ad-hoc simulator is implemented in Java and is designed for future integration with meta-heuristic frameworks such as jMetal [9]. The source code of the ad-hoc simulator is public and is available at `https://www.fing.edu.uy/inco/grupos/cecal/hpc/DRAS/`.

The simulation period is divided into intervals of equal duration int_d. At each interval, the scheduler assigns the arrived tasks to the servers, taking into account the current capacity of each server and the scheduling strategy criteria. The number of intervals that a task is running on a server is calculated as $ct/int\_d + 1$, where ct is the completion time of a task (in seconds) and int_d is the duration of each interval (in seconds). The completion time of a task is defined as its size in millions of instructions per second (MIPS) divided by the assigned core speed (also, in MIPS).

Fig. 2 presents a sample simulation of a client where 13 tasks (gray rectangles) are assigned on two servers with three cores each. In the three-dimensional graph, servers cores are on the vertical axis (cores). The time intervals of the simulation are shown on the horizontal axis (steps). The oblique axis (power) represents the whole power consumption of the client in each time interval (which depends on the workload of its servers).

As shown in the sample simulation in Fig. 2, the reduction in power consumption achieved after applying the scheduling strategy is calculated as the reference power consumption minus the maximum period consumption, considering all intervals. The reference power consumption considered in the proposed scheduling model is the maximum energy consumption of the period, considering all the intervals, obtained when executing the simulation without applying any energy optimization. This reference is used to calculate the reduction achieved by the energy optimization strategy.

To estimate the power consumption of servers in

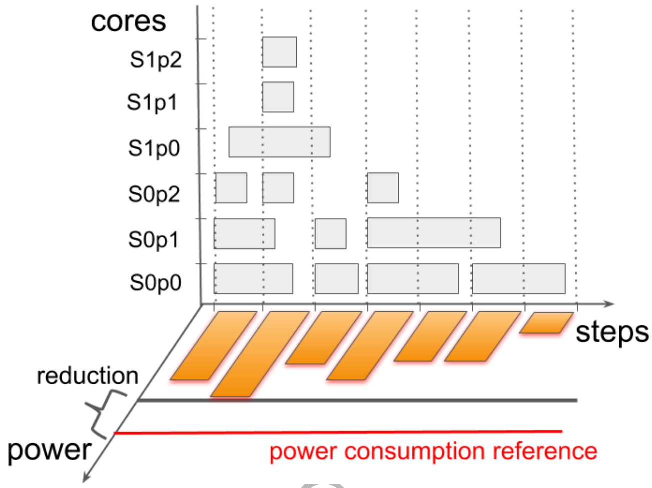

Fig. 2.: Sample energy scheduling simulation for a client of the datacenter

the simulation, specific power consumption models are needed. The next paragraph describes how the strategies evaluate the power consumption of a given planning.

Power consumption model. In order to estimate the power consumption of a scheduling, novel power consumption models are applied. The models were generated by applying polynomial regression to the experimental data published in our previous work [22].

The experimental data corresponds to the power consumption measurements of a high-end HP Proliant DL380 G9 server (2 Intel Xeon E5-2643v3 CPUs, 12 cores each, 128 GB RAM) from the National Supercomputing Center (Cluster-UY), Uruguay [25], registered through a non-invasive power monitoring setup [22, 23].

The power consumption models consider the combined utilization of two computing resources (CPU and memory) by tasks specialized in either of the available resources. This is a novel approach for modeling real datacenter operation, which allows improving the accuracy of previously proposed power consumption models based on overall server utilization or CPU utilization only. In addition, the proposed models extract characteristics of power consumption from the simultaneous execution of tasks that makes use of different computing resources (CPU or memory). These features cannot be extracted through models generated considering computing resource independently, such as the ones commonly applied in the related literature [13, 14, 26]. In addition, it allows implementing scheduling algorithms to take advantage of tasks consolidation.

The empirical model obtained using a linear regression approach for the Intel server is presented in Eq. 4.5.

$$PC = \begin{cases} 0.8904x + 1.5167y \\ \qquad +115.5292 & \text{if } x,y \neq 0 \\ PC_{IDLE} & \text{if } x,y = 0 \end{cases} \tag{4.5}$$

In turn, the empirical model obtained using a polynomial regression approach for the Intel server is presented in Eq. 4.6.

$$PC = \begin{cases} 116.09724 + 0.73349x + \\ 1.72632y + 0.00786\,xy + & \text{if } x,y \neq 0 \\ 0.00043x^2 - 0.00388y^2 \\ PC_{IDLE} & \text{if } x,y = 0 \end{cases} \tag{4.6}$$

In Equations 4.5 and 4.6, variable x represents the percentage of CPU-intensive tasks in execution and variable y represents the percentage of memory-intensive tasks in execution. PC is the power consumption of the server, in Watts, and PC_{IDLE} is the power consumption of the server when idle, i.e., when it is not executing any task.

Figure **??** presents the 3D graph of the linear regression model on Intel server, for different utilization levels (UL, in percentage values) of CPU and memory resources. In turn, Figure 3 presents the 3D graph of the quadratic regression model on the Intel server for different ULs.

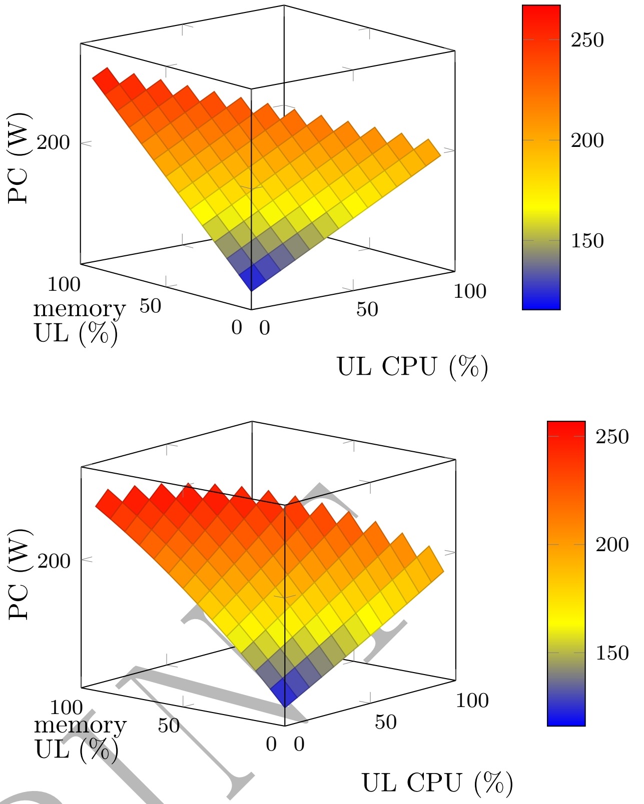

Fig. 3.: Quadratic power consumption model for combined CPU–memory workloads on Intel server

Two relevant metrics to assess the quality of statistical models were applied to analyze the models: *coefficient of determination* (R-squared or R^2) and *adjusted R-squared* ($\overline{R}^2$). Both metrics evaluate the forecasting capabilities of the studied model (i.e., prediction of future values in a given temporal series). $\overline{R}^2$ is an extension of R^2 proposed to avoid spurious increasing when using a larger number of independent variables [36].

Table 1 reports the statistics of the proposed power consumption models for combined CPU and memory intensive tasks. Results indicate that the quadratic model achieves a better fit to the data.

Table 1.: Statistics of linear and quadratic power consumption models for combined CPU–memory workloads on Intel server

model	R^2	$\overline{R}^2$
linear	0.94252	0.93207
quadratic	$\boxed{0.99423}$	$\boxed{0.99062}$

The proposed models are based on experimental data of power consumption in real hardware. This approach, unlike theoretical models, allows considering the holistic behavior of power consumption in servers, which depends on multiple factors (chip design, voltage, cooling system, etc.).

5 Experimental analysis

This section reports the experimental evaluation of the proposed negotiation approach for the participation of datacenters and supercomputing facilities in the electricity market.

5.1 Problem instances

The instances created for the evaluation of the proposed approach are described in this subsection.

Instances description. Problem instances were created considering real data from both workloads executed and computing resources available on nowadays datacenters and supercomputing facilities.

A problem instance is divided in three components, simulation parameters, datacenter details, and workloads:

- *Simulation parameters* include the simulation period and the number of scheduling steps, and other relevant parameters of the negotiation: the demanded power reduction, the price per Watt of power generation using gasoline, and the error threshold used as stopping criterion of the negotiation.
- *Datacenter details* include the number of clients and their number of servers, and the total number of servers. In addition, the *tolerance*, a normalized value in [0,1] that defines the flexibility of each client for deferring submitted tasks, is indicated.
- *Workloads description* corresponds to the set of tasks to be executed by the clients. The format uses for workloads extends the Standard Workload Format (swf) [1]. The extension proposed in this article to model the addressed problem adds two new fields to the format: i) *penalty*, which represents the monetary cost to be paid by the datacenter client, when operating as IT system provider, to the user in case of not meeting the SLA (specified for the deadline of each task), a larger penalty implies a large monetary cost, and ii) *tolerance*, which defines the flexibility of the client for deferring submitted tasks, according to a normalized value where 1.0 represents the lower flexibility, i.e., a totally strict client.

The relevant fields of the extended swf format used in this article to describe workloads include:

- *id:* identifier of the job.
- *owner:* identifier of the user.
- *size:* a value that is proportional to the number of micro-instructions needed to execute to complete the task (in MIPS).
- *arrive:* arrival time of the job (in seconds).
- *deadline:* deadline time of the job (in seconds).
- *deferrable:* indicates if the job can be postponed or not.
- *penalty:* monetary cost in case of not meeting the service level agreement.
- *type:* indicates the type of computing resource in which the task is intensive. Type 0 corresponds to CPU intensive tasks and type 1 corresponds to memory intensive tasks.

Listing 1 presents an example of the workload format.

Specific instances. Specific instances were generated to evaluate and validate the proposed model for datacenter participation in the electricity market. Instances are organized by size, according to the number of clients and servers, which defines the complexity of the underlying planning problem. A total number of nine instances were generated by combining different sizes (small, medium, and large) with defined tolerance values (1.0, 0.8, 0.4). Small instances contain 5 clients, with 5 servers and 1500 tasks each. In turn, medium instances contain 10 clients, with 10 servers and 6000 tasks each.

Finally, large instances contain 30 clients, with 30 servers and 9000 tasks each. The nomenclature user for instances is DPEM-[*size*]-[*tolerance*]. DPEM is an acronym of 'Datacenter Participation in Electric Market' problem.

Regarding the computational infrastructure, the following parameters considered for all instances, according to the described justifications and related literature:

- servers with 24 cores are considered, similar to the characteristics of the high-end servers included in Cluster-UY [25]. For simplicity, identical servers are included in each simulation (HP Proliant DL380 G9 server with 2 Intel Xeon E5-2643v3 CPUs).
- each core processes 3000 MIPS. The chosen value for core speed is realistic according to nowadays processors and follows the related literature about datacenter simulations [3, 26].
- the power consumption model of the servers corresponds to the quadratic regression model in Eq. 4.6. This model corresponds to high-end servers currently existing in Cluster-UY and presents the best fit to the experimental data.

Regarding the workloads of each client, they were generated applying the following methodology:

- size randomly generated according to a uniform distribution in [180000,540000] MIPS. These size values implies task completion times between 1 and 3 minutes, according to the processor speed. Similar task sizes are used in the related literature [1, 3]. The considered values allow defining a reasonable number of tasks to schedule, accounting for about 80% of system utilization, considering a time horizon of one hour and the number of resources in the simulated computational platform.
- arrival times randomly generated according to a uniform distribution in [0,3600]. The upper bound (3600) was chosen so that all tasks arrive within the scheduling period.
- deadlines randomly generated according to a uniform distribution in $[ct{\times}20, ct{\times}30]$, where ct is the task completion time, computed as size/core speed.
- deferrable is 1 (i.e., all task are assumed to be deferrable), as previous works considering the scheduling of both, deferrable and no-deferrable tasks, confirmed that the impact of non-deferrable tasks is negligible for the overall quality of the obtained schedules [15, 26].
- penalty values randomly selected in {1,10,50}, according to a non-uniform distribution (probabilities 0.6, 0.3, and 0.1, respectively). Values were chosen to generate different QoS levels for tasks.
- type of task randomly generated according to a uniform distribution in {0,1}. The number of CPU-intensive tasks and memory-intensive tasks are assumed to be similar, following the analysis by Armenta et al. [1].

Regarding simulation parameters, the following values were set: the simulation time horizon for task planning is T=60 minutes, the demanded power reduction is 2000 W, 10000 W, and 30000 W for small, medium and large instances, respectively; a price per Watt for the on-site power generator of 2 monetary units; an error threshold of 0.005; and an initial offer price of 0.01 monetary units.

The files describing the instances and the source code of the instances generator implemented are publicly available at `https://www.fing.edu.uy/inco/grupos/cecal/hpc/DRAS/`.

id	owner	size	arrive	deadline	deferrable	penalty	type
1	0	44979	1	1124	1	10	0
2	0	141241	4	3531	1	2	1
3	0	15143	10	378	0	6	0
4	0	143270	50	3581	0	10	1
5	0	139761	51	3494	1	6	1
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮

Listing 1: Workload format example

5.2 Development and execution platform

The simulator was implemented in Java SE 1.8, using JMetal [9] as an auxiliary library. Instance files and data analysis reports were generated on the Python language [30], using the Pandas library [20] and the Jupyter Notebook environment [18] following the reproducible research paradigm.

Experiments were performed on an HP Proliant DL380 G9 server (2 Intel Xeon E5-2643v3 CPUs, 12 cores each, 128 GB RAM) from the National Supercomputing Center (Cluster-UY), Uruguay [25]

5.3 Numerical results

This subsection presents and discusses the numerical results of the experiments, which corresponds to simulations of the datacenter operation described in the previous sections.

Results are presented in two subsections. Subsection 5.3.1 corresponds to the evaluation of the client optimization strategy, where the reduction function (which depends on the offered price) of a client is presented. Subsection 5.3.2 corresponds to the results of the negotiation of the datacenter administrator with its clients.

5.3.1 Client offer evaluation function

The offer evaluation function of a client computes its power consumption reduction according to the the monetary offer received. Fig. 4a presents the simulation results of the offer evaluation function for three sample clients for small instances, one client for each tolerance value considered. . Results show that the power optimization strategy manages to reduce the power consumption for all offer. This means that clients have positive profit values in all offers. The comparison between clients with different tolerances and the same instance size allows observing that less flexible customers demand higher prices to reduce the same amount of power consumption.

Fig. 4b presents the offer evaluation function of sample clients of medium instances. In this instance size, clients are able to reduce their power consumption above of 2000 W. As well as observed for the small instance samples, clients with low

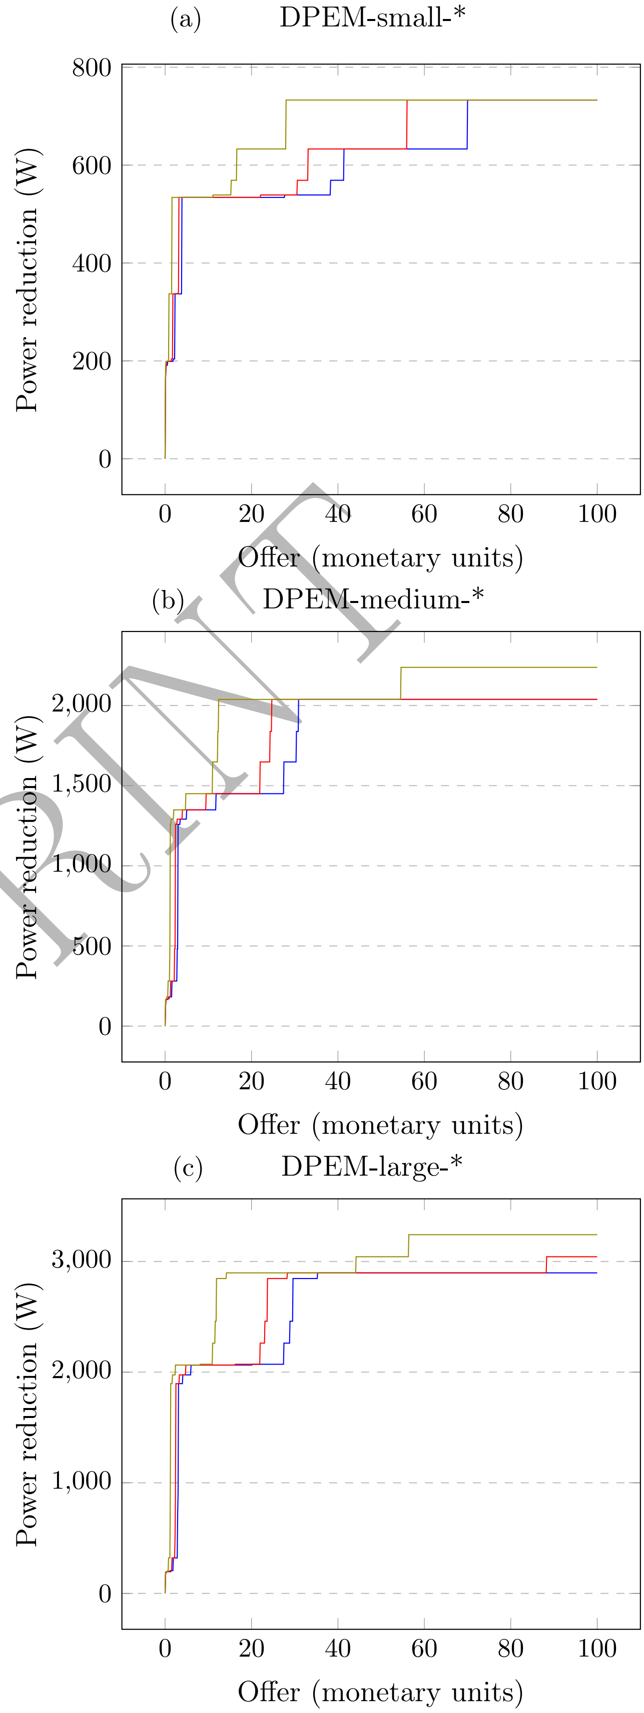

Fig. 4.: Offer evaluation function of nine sample clients: *-tol-0.4, *-tol-0.8, *-tol-1.0

tolerance demand high price per reduced Watt. In turn, Fig. 4c shows the offer evaluation function of sample clients of large instances. In these instances, clients are able to reduce their power consumption above of 3000 W. All reported client evaluation functions have a similar behavior, characterized by high values of the first derivative near zero. This similarity is due to the fact that all client workloads are generated using the same distribution function and all clients use the same scheduling strategies.

5.3.2 Negotiation between datacenter administration and clients

Fig. 5 shows the result of the power reduction negotiation between the datacenter administrator and the clients, for small instances. The offer (per Watt) to the clients is on the independent axis. Blue squares are the evolution of the clients reduction (CR) through the negotiation, red circles are the values of y_k, and green triangles are the power generated on-site to cover the power reduction target ($OG = D - CR$). For the small instances, the negotiation quickly reaches a power reduction of 600 W, with a low monetary cost. However, larger power reductions require more iterations and a greater monetary offer. This slowdown in reduction is due to the fact that, at this point, clients cannot reduce their power consumption without a significant impact on user performance.

Table 2 reports a summary of the negotiation for small instances. Three negotiation steps are considered: the first offer, an intermediate offer, and the last offer (that is, when the negotiation algorithm ends according to the stopping criteria). Column k is the negotiation step, *price* is the offer in step k, CR is the reduction obtained from clients, OG is the on-site generation, ϵ evaluates the compliance of the coupling restriction, and *cost* is the monetary value that the datacenter administrator must invest to achieve the target reduction, defined in Eq. 5.1 .

$$cost = CR \times price + OG \times \alpha \qquad (5.1)$$

The comparison between instances with different tolerance values shows that when clients are less flexible, the negotiation determines in the last step low values in CR column and high values in OG column. This behavior corresponds to the intuitive

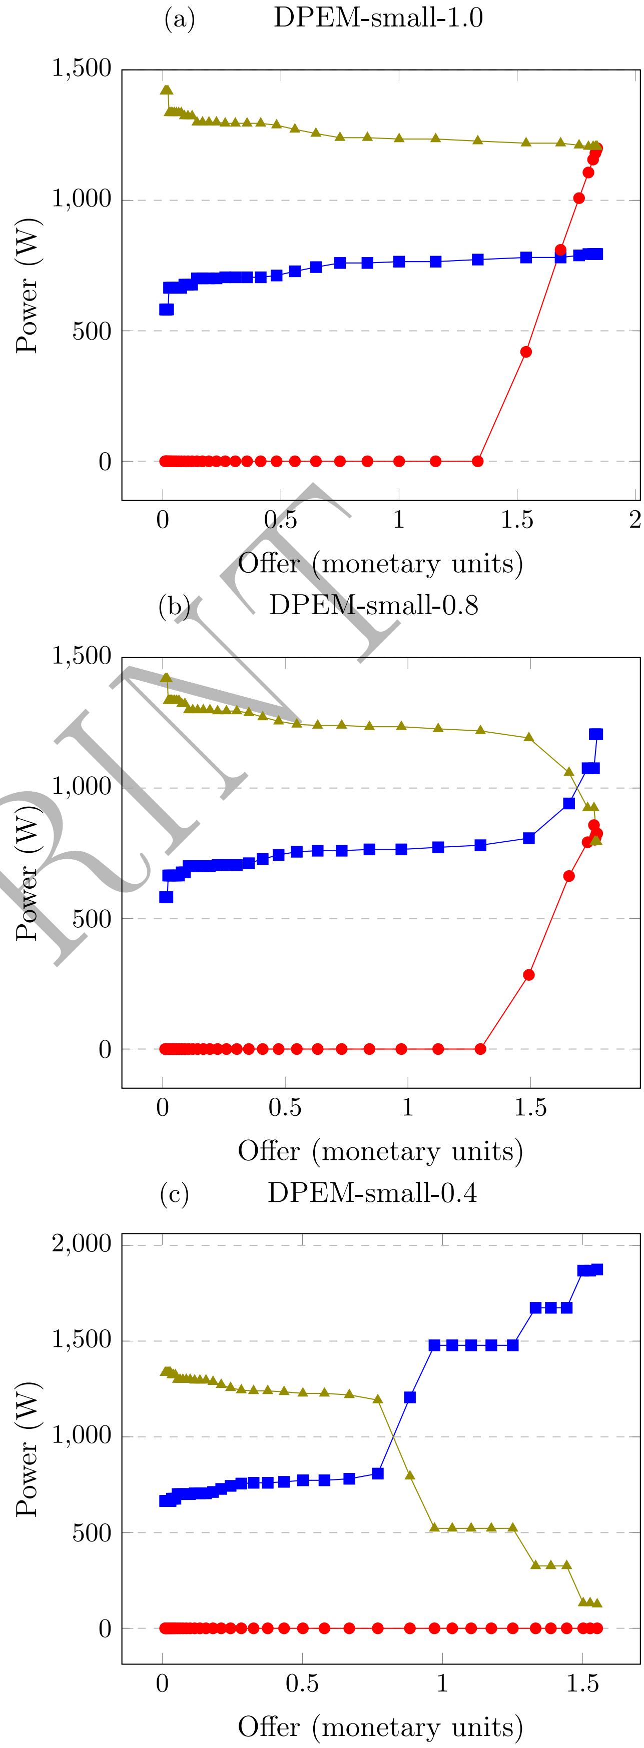

Fig. 5.: Negotiation for small instances: ●–y_k, ■–CR, ▲–on-site generation

DPEM-small-1.0							
k	*price*	*CR*	y_k	ϵ	*OG*	*cost*	*swc*
1	0.012	582	0	70.90%	1418	2842	5360
20	0.227	701	0	64.95%	1299	2757	5128
41	1.840	794	1200	0.03%	1206	3872	5024
DPEM-small-0.8							
k	*price*	*CR*	y_k	ϵ	*OG*	*cost*	*swc*
1	0.012	582	0	70.90%	1418	2842	4855
21	0.260	705	0	64.75%	1295	2773	4615
40	1.760	1206	802	0.04%	794	3710	4376
DPEM-small-0.4							
k	*price*	*CR*	y_k	ϵ	*OG*	*cost*	*swc*
1	0.012	665	0	66.75%	1335	2677	3680
21	0.244	728	0	63.60%	1272	2721	3561
42	1.576	1874	0	0.06%	126	3205	2442

Table 2.: Negotiation summary for small instances

idea that in datacenters where clients have less flexible SLAs, the on-site generation is the main option to achive the target reduction established by the electric provider. Moreover, less flexible instance (i.e. small-1.0) implies large offers.

Another interesting result arose regarding the cost for the datacenter to fulfill the target reduction: the best cost values are not computed in the last steps of the negotiation. This result indicates that the traditional approach for negotiations based on theoretical client reduction functions is not applicable when modeling realistic workloads for datacenters, which is a direct contribution of the research reported in this article. The problem of datacenter participation in the electricity market is inherently multiobjective. At least two objectives must be considered, regarding operation costs and QoS offered to users, and Pareto analysis are needed to determine different trade-offs between total cost and negotiation offers. Studying the multiobjective version of the problem is proposed as one of the main lines for future work.

Results of the negotiation for medium instances are presented in Fig. 6.

In medium instances, starting from the same initial price, negotiation requires more interactions with clients than in small instances to achieve the desired reduction, because in medium instances the target reduction is 8000 W, four times the target

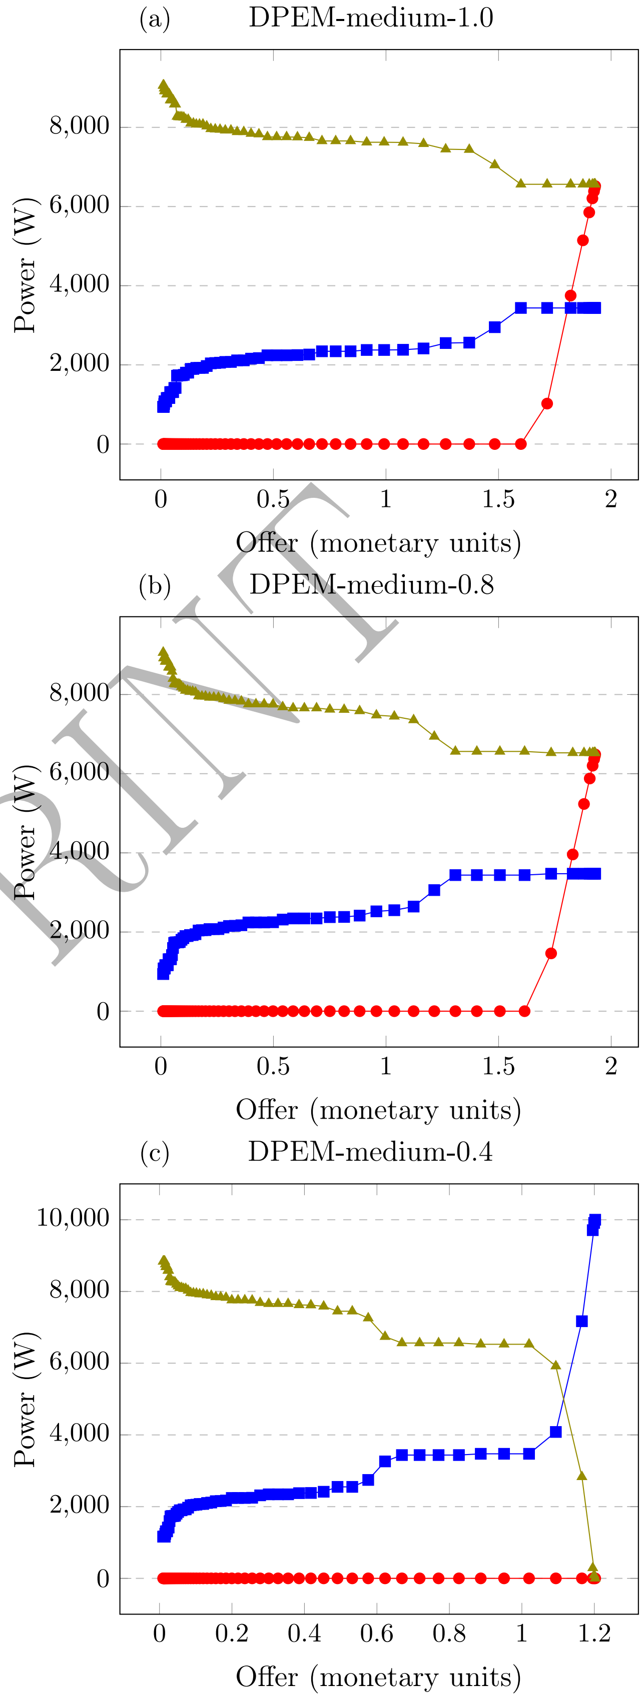

Fig. 6.: Negotiation for medium instances: ●–y_k, ■–CR, ▲–on-site generation

specified for small instances. Client reductions quickly achieve a stable value for the instance with the lower flexibility. However, more flexible instances allow finding negotiation values that significantly reduces the power used by clients, thus reducing the on-sit generation.

Table 3 reports the summary of the negotiation for a medium instances. As in small instances, the negotiation with less flexible clients requires more on-site generation and higher prices. Similar arguments holds for those instances. However, for instance DPEM-medium-0.4, results indicate that a significant reduction on the datacenter cost (32%) is obtained for the values computed at the end of the negotiation. This result confirm that when clients have the largest flexibility (i.e., they penalize deadline violations with the lowest values), the negotiation provides significant profit for the datacenter administrator. On the other instances, trade-off analysis are needed to determine the best price to offer to clients.

DPEM-medium-1.0						
k	*price*	*CR*	y_k	ϵ	*OG*	*cost*
1	0.011	940	0	90.60%	9060	18130
33	0.204	1924	0	80.76%	8076	16544
67	1.930	3438	6517	0.04%	6562	19759
DPEM-medium-0.8						
k	*price*	*CR*	y_k	ϵ	*OG*	*cost*
1	0.011	940	0	90.60%	9060	18130
34	0.216	2047	0	79.53%	7953	16348
68	1.930	3473	6486	0.04%	6527	19756
DPEM-medium-0.4						
k	*price*	*CR*	y_k	ϵ	*OG*	*cost*
1	0.011	1162	0	88.38%	8838	17688
30	0.143	2102	0	78.98%	7898	16096
61	1.201	9910	0	0.09%	90	12081

Table 3.: Negotiation summary for medium instances

Results of the negotiation for large instances are presented in Fig. 7.

Experimental results for large-sized instances confirm previous behaviour. Because of the increased number of clients and the increased target reduction, the negotiation requires about three times more iterations than the medium-sized

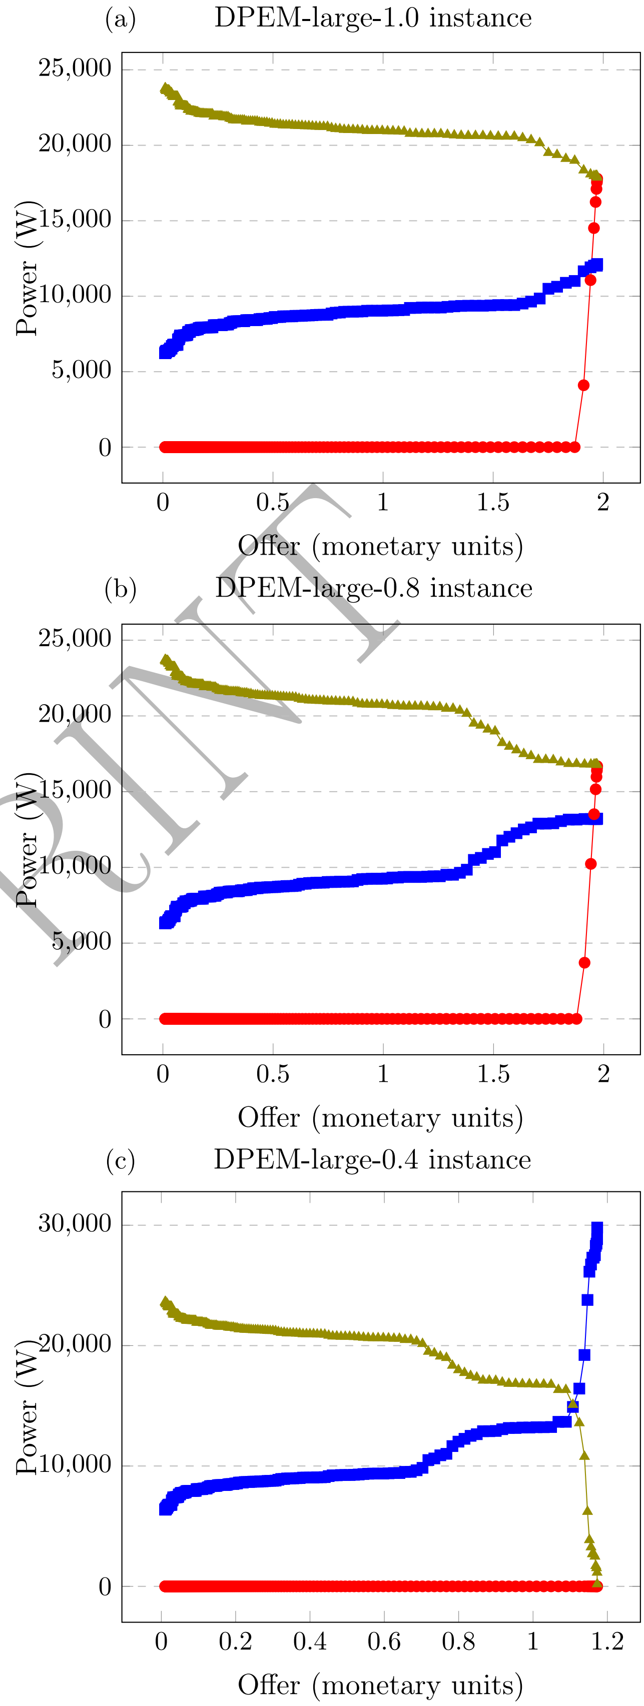

Fig. 7.: Negotiation for large instances: ●–y_k, ■–CR, ▲–on-site generation

instances and about five times more iterations than the small-sized instances. Clients quickly reach a reduction of about 9000 W in problem instances DPEM-large-1.0 and DPEM-medium-0.8. After that, clients require a significant increase in the offer from the datacenter administrator to surpass the 10000 W reduction mark. In instance DPEM-medium-0.4 reductions are over 13.000 W with a maximum value of 30.000 W.

Table 4 reports the negotiation summary for large instances. The negotiation algorithm is equally effective for the less flexible instances, producing similar results for DPEM-large-1.0 and DPEM-medium-0.8 instances. However, a notable improvement is obtained when clients have more flexibility for tasks execution. For the DPEM-medium-0.4 instance, the negotiation algorithm is able to reduce 41% when compared to DPEM-large-1.0 and DPEM-medium-0.8 instances. These results confirms that the proposed negotiation approach is able to properly take advantage of deferring execution tasks to fulfill the requested power consumption reduction.

DPEM-large-1.0						
k	*price*	*CR*	y_k	ϵ	*OG*	*cost*
1	0.010	6224	0	79.25%	23776	47614
108	0.172	7887	0	73.71%	22113	45582
216	1.973	12128	17789	0.03%	17872	59673
DPEM-large-0.8						
k	*price*	*CR*	y_k	ϵ	*OG*	*cost*
1	0.010	6310	0	78.96%	23690	47443
110	0.177	7938	0	73.54%	22062	45529
219	1.970	13221	16682	0.03%	16779	59603
DPEM-large-0.4						
k	*price*	*CR*	y_k	ϵ	*OG*	*cost*
1	0.010	6376	0	78.75%	23624	47311
106	0.152	8399	0	72.00%	21601	44478
205	1.173	29809	0	0.06%	191	35336

Table 4.: Negotiation summary for large instances

6 Conclusions and future work

This article studied a negotiation approach for the participation of datacenters and supercomputing facilities in smart electricity markets, an important problem in modern smart grid systems.

A specific case of demand response strategy was studied for colocation datacenters to commit power reductions, according to offers proposed to clients. A decentralized approach was applied for negotiation, where clients do not need to provide strategic information to the datacenter administrator. Instead, each client negotiates a price considering a planning heuristic and the features of the tasks submitted for execution. Furthermore, a model based on empirical information was presented to determine the power consumption of CPU-intensive and memory-intensive tasks, using data from real datacenters.

The negotiation algorithm and a heuristic planning method for energy reduction optimization were experimentally validated over nine realistic problem instances that model different problem dimensions and flexibility of the datacenter clients.

The obtained results indicate that the proposed approach is effective to provide appropriate demand response actions according to monetary incentives. Not only the system achieved economic benefits to the datacenter operator, but also for the tenants (by providing rewards for reductions) and for the environment, due to the reductions of diesel usage.

Summarizing, clients quickly reached appropriate power reductions, thus limiting the need of using on-site generation by the datacenter. Results confirmed that the problem is inherently multiobjective. Both operation costs and QoS offered to users must be considered in the formulation, and trade-offs between total cost and negotiation offers must be studied. The proposed approach is realistic and efficient, to be implemented in nowadays datacenters and supercomputing facilities.

Two main lines are proposed for future work. On the one hand, studying the multiobjective version of the problem of datacenter participation in the electricity market, to account and characterize the trade-off between the total cost and the negotiation offers, including Pareto-based analysis of the results.

On the other hand, the proposed model can be improved. The presented analysis for reduction on energy consumption of the computational infrastructure should be extended to include the air conditioning system, thermal and electric energy storage, in line with our previous works for solar-based low-consumption datacenters [15, 26], by applying a thermal model to compute realistic values for the temperature in the datacenter. Furthermore, different time horizons must be studied for the optimization problem and specific features of the Uruguayan electricity market must be considered for the application of the proposed method in the National Supercomputing Center.

References

[1] Fermín Armenta-Cano et al. "Min_c: Heterogeneous concentration policy for energy-aware scheduling of jobs with resource contention". In: *Programming and Computer Software* 43.3 (2017), pp. 204–215.

[2] Holger Berndt et al. *Network and system rules of German transmission system operator*. 1.02 ed. Association of network operators. Germany, 2019.

[3] Rodrigo Calheiros et al. "CloudSim: a toolkit for modeling and simulation of cloud computing environments and evaluation of resource provisioning algorithms". In: *Software: Practice and experience* 41.1 (2011), pp. 23–50.

[4] PJM Day-Ahead and Real-Time Market Operations Division. *Manual 11: Energy & Ancillary Services Market Operations*. 108th ed. PJM Interconnection LLC. Mar. 2018.

[5] Pierre Delforge and Josh Whitney. *Scaling up energy efficiency across the data center industry: Evaluating key drivers and barriers*. Tech. rep. Natural Resources Defense Council and Anthesis, 2014.

[6] PJM Dispatch Operations Division. *Manual 12: Balancing Operations*. 39th ed. PJM Interconnection LLC. Feb. 2019.

[7] Bernabé Dorronsoro et al. "A hierarchical approach for energy-efficient scheduling of large workloads in multicore distributed systems". In: *Sustainable Computing: Informatics and Systems* 4.4 (2014), pp. 252–261.

[8] K. Du Bois et al. "SWEEP: Evaluating Computer System Energy Efficiency Using Synthetic Workloads". In: *6th International Conference on High Performance and Embedded Architectures and Compilers*. 2011, pp. 159–166.

[9] Juan J. Durillo and Antonio J. Nebro. "jMetal: A Java framework for multi-objective optimization". In: *Advances in Engineering Software* 42 (2011), pp. 760–771.

[10] European Automotive Research Partner Association. *Smart Grids European Technology Platform*. https://www.etip-snet.eu/. January 2020.

[11] Federal Energy Regulatory Commission. *Assessment of demand response & advanced metering*. Tech. rep. AD-06-2-00. 2006.

[12] Xixhou Feng, Rong Ge, and K.W. Cameron. "Power and Energy Profiling of Scientific Applications on Distributed Systems". In: *19th IEEE International Parallel and Distributed Processing Symposium*. 2005, pp. 34–44.

[13] Chung-Hsing Hsu and Stephen W. Poole. "Power signature analysis of the SPECpower_ssj2008 benchmark". In: *IEEE International Symposium on Performance Analysis of Systems and Software*. Apr. 2011, pp. 227–236.

[14] Santiago Iturriaga, Sebastián García, and Sergio Nesmachnow. "An Empirical Study of the Robustness of Energy-Aware Schedulers for High Performance Computing Systems under Uncertainty". In: *High Performance Computing*. 2014, pp. 143–157.

[15] Santiago Iturriaga and Sergio Nesmachnow. "Scheduling Energy Efficient Data Centers Using Renewable Energy". In: *Electronics* 5.4 (2016), p. 71.

[16] Ramesh Johari and John Tsitsiklis. "Parameterized Supply Function Bidding: Equilibrium and Efficiency". In: *Operations Research* 59.5 (2011), pp. 1079–1089.

[17] Paul Klemperer and Margaret Meyer. "Supply Function Equilibria in Oligopoly under Uncertainty". In: *Econometrica* 57.6 (Nov. 1989), pp. 1243–1277.

[18] Thomas Kluyver et al. "Jupyter Notebooks - a publishing format for reproducible computational workflows". In: *Positioning and Power in Academic Publishing: Players, Agents and Agendas, 20th International Conference on Electronic Publishing*. Ed. by Fernando Loizides and Birgit Schmidt. 2016, pp. 87–90.

[19] K. Kurowski et al. "DCworms–A tool for simulation of energy efficiency in distributed computing infrastructures". In: *Simulation Modelling Practice and Theory* 39 (2013), pp. 135–151.

[20] Wes McKinney. "pandas: a foundational Python library for data analysis and statistics". In: *Python for High Performance and Scientific Computing* (2011), pp. 1–9.

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

[22] Jonathan Muraña et al. "Characterization, modeling and scheduling of power consumption of scientific computing applications in multicores". In: *Cluster Computing* 22.3 (Sept. 2019), pp. 839–859.

[23] Jonathan Muraña et al. "Power Consumption Characterization of Synthetic Benchmarks in Multicores". In: *High Performance Computing. Communications in Computer and Information Science*. 2017, pp. 21–37.

[24] Sergio Nesmachnow. "An overview of metaheuristics: accurate and efficient methods for optimisation". In: *International Journal of Metaheuristics* 3.4 (2014), pp. 320–347.

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

[26] Sergio Nesmachnow, Cristian Perfumo, and Íñigo Goiri. "Holistic multiobjective planning of datacenters powered by renewable energy". In: *Cluster Computing* 18.4 (2015), pp. 1379–1397.

[27] Niangjun Chen and Xiaoqi Ren and Shaolei Ren and Adam Wierman. "Greening multitenant data center demand response". In: *Performance Evaluation* 91 (2015), pp. 229–254.

[28] Sebastián Montes de Oca, Pablo Belzarena, and Pablo Monzón. "Benefits of optimal demand response in distribution networks in a competitive retail market". In: *IEEE URUCON*. 2017, pp. 1–4.

[29] Sebastián Montes de Oca, Pablo Belzarena, and Pablo Monzón. "Optimal demand response in distribution networks with several energy retail companies". In: *IEEE Conference on Control Applications*. 2016, pp. 1092–1097.

[30] Travis E. Oliphant. "Python for Scientific Computing". In: *Computing in Science Engineering* 9.3 (2007), pp. 10–20.

[31] Thant Zin Oo et al. *A Survey on Coordinated Power Management in Multi-Tenant Data Centers*. 1st ed. Springer, 2018.

[32] Fernando Paganini, Pablo Belzarena, and Pablo Monzón. "Decision making in forward power markets with supply and demand uncertainty". In: *International Conference on Information Sciences and Systems*. 2014, pp. 1–6.

[33] Neal Parikh and Stephen Boyd. "Proximal Algorithms". In: *Foundations and Trends in Optimization* 1.3 (2014), pp. 127–239.

[34] National Grid PLC. *Product Roadmap For Frequency Response and Reserve*. London, England, UK, Dec. 2017.

[35] Steven Stoft. *Power System Economics: Designing Markets for Electricity*. 1st ed. Wiley-IEEE Press, 2002.

[36] Henri Theil. *Economic forecasts and policy*. 2nd ed. Amsterdam North-Holland Publishing Company, 1961.

[37] Stef Tijs and Theo Driessen. "Game theory and cost allocation problems". In: *Management science* 32.8 (1986), pp. 1015–1028.

[38] Nguyen H. Tran et al. “Incentive Mechanisms for Economic and Emergency Demand Responses of Colocation Datacenters”. In: *IEEE Journal on Selected Areas in Communications* 33.12 (2015), pp. 2892–2905.

[39] Nguyen Tran et al. “Reward-to-Reduce: An Incentive Mechanism for Economic Demand Response of Colocation Datacenters”. In: *IEEE Journal on Selected Areas in Communications* 34.12 (2016), pp. 3941–3953.

[40] Youshi Wang et al. “A market-oriented incentive mechanism for emergency demand response in colocation data centers”. In: *Sustainable Computing: Informatics and Systems* 22 (2019), pp. 13–25.