Incentive-Based Energy Consumption Scheduling Algorithms for the Smart Grid

Stéphane Caron, George Kesidis

► To cite this version:

Stéphane Caron, George Kesidis. Incentive-Based Energy Consumption Scheduling Algorithms for the Smart Grid. 2010 1st IEEE International Conference on Smart Grid Communications (SmartGridComm), Oct 2010, Gaithersburg, United States. pp.391-396, 10.1109/SMARTGRID.2010.5622073 . hal-02108638

HAL Id: hal-02108638

<https://hal.science/hal-02108638>

Submitted on 24 Apr 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Incentive-based Energy Consumption Scheduling Algorithms for the Smart Grid

Stéphane Caron
Département d'Informatique
École Normale Supérieure
45 rue d'Ulm, Paris, France
Email: stephane.caron@ens.fr

George Kesidis
CS&E and EE Depts
The Pennsylvania State University
University Park, PA, 16802
Email: kesidis@engr.psu.edu

Abstract—In this paper, we study Demand Response (DR) problematics for different levels of information sharing in a smart grid. We propose a dynamic pricing scheme incentivizing consumers to achieve an aggregate load profile suitable for utilities, and study how close they can get to an ideal flat profile depending on how much information they share. When customers can share all their load profiles, we provide a distributed algorithm, set up as a cooperative game between consumers, which significantly reduces the total cost and *peak-to-average* ratio (PAR) of the system. In the absence of full information sharing (for reasons of privacy), when users have only access to the instantaneous total load on the grid, we provide distributed stochastic strategies that successfully exploit this information to improve the overall load profile. Simulation results confirm that these solutions efficiently benefit from information sharing within the grid and reduce both the total cost and PAR.

I. INTRODUCTION

The current U.S. electrical grid is built up according to a static, centralized structure: remote power plants transmit electrical power through long-distance high-voltage lines to substations (transmission network), which in turn adapt and deliver it to local end users (distribution network). In this model, the local network is often statically tuned to match a given average load profile from its consumers.

Yet this structure is about to undergo a major shift: the progressive integration of *smart meters* [1] and communicating appliances will upgrade this “blind” system to a decentralized “smart grid” [2], which is foreseen as a way to save billions dollars in energy consumption [3].

In a smart grid infrastructure, utilities can set up dynamic tariffs incentivizing customers to adjust their loads to the current state of the network. This key feature, known as *demand side management*, will yield several benefits, including:

- Integration of intermittent energy sources such as wind or solar power at the distribution level [4];
- Demand Response: customers will be encouraged to shift their heavy loads to *off-peak* hours;
- Resilience to attacks or power outages, with users spurred to turn off their “non-critical” devices in case of heavy load on the grid or upstream outages [5]; and
- Energy savings: studies [6] already suggest that giving customers access to real-time consumption information yields significant savings.

Such load management is becoming even more crucial as *plug-in hybrid electric vehicles* (PHEVs) are coming to the market. With battery capacities varying from 15 to 50 kWh, these vehicles are expected to double the average household load during charging time [4].

Therefore, the design of appropriate incentives and efficient *energy consumption scheduling* (ECS) algorithms is a main issue for the deployment of the upcoming smart grid.

In [7], the authors deal with ECS in the case of increasing strictly convex cost functions. They propose a distributed algorithm and show through a game-theoretic analysis that, for incentives satisfying certain properties, it yields optimal energy consumptions for end users. However, they implicitly assume that the daily load on the network is proportional to the daily cost for the utility, with a constant independent of load scheduling (*i.e.*, of game dynamics): this is a strong hypothesis which implies utilities’ costs are linearly bounded in any situation.

In this paper, we survey different scheduling problems depending on the DR architecture (centralized or distributed) and the degree of knowledge appliances have on the state of the network. As in [7], we embed consumers in a local distribution network consisting of a single energy source (*e.g.*, a step-down substation) supplying several load subscribers. Customers are incentivized to move their loads to off-peak hours through marginal costs which are linearly increasing with instantaneous value of network overload.

The rest of this paper is organized as follows. We introduce our pricing scheme and notations in section II. In section III, we see that the general scheduling problem is NP-hard, and provide a distributed solution set-up as a cooperative game between consumers. We provide stochastic policies for a decentralized setting in section IV. In section V, we derive the best distributed policy for synchronized users in a power grid, which we use as a reference in section VI where we provide experimental results. We draw conclusions in section VII and suggest future work in section VIII, including considerations of interruptible and non-uniform demand profiles.

II. PROBLEM SET-UP

We consider a T -hour time period, *e.g.*, $T = 6$ hours from midnight to 6AM, during which N customers need to

automatically schedule their electrical jobs. Note that our time horizon is finite since users do not want their jobs to be delayed forever.

A. Loads and Costs

The n^{th} customer has demand profile D_n parametrized by flexible start time s_n and fixed (d_n, τ_n) parameters, where d_n denote the instantaneous power consumption of the job and τ_n its duration:

$$D_n(t) = d_n \mathbf{1}_{\{s_n \leq t \leq s_n + \tau_n\}}$$

where $0 \leq s_n \leq T - \tau_n$. We assume that once the service start time s_n is selected, it cannot be interrupted by the user. The total instantaneous load on the network is then

$$\lambda(t) := \sum_n D_n(t) = \sum_n d_n \mathbf{1}_{\{s_n \leq t \leq s_n + \tau_n\}}$$

We denote by $C(\lambda(t))$ the cost, in \$/kW, experienced by the utility at time t (for the costs charged to consumers, see equation (4)). It depends on the instantaneous load $\lambda(t)$, and the function C itself is likely to depend on additional system parameters. For example, the two-step conservation rate model used by BC Hydro [8] (parametrized by load threshold L) is

$$C_L(\lambda(t)) = C_0 \cdot \mathbf{1}_{\{\lambda(t) < L\}} + C_1 \cdot \mathbf{1}_{\{\lambda(t) \geq L\}}.$$

In a more general setting, C can be any smooth convex function of $\lambda(t)$. However, in this article, we will focus on a *ramp* cost function with load threshold $L > 0$:

$$C_L(\lambda(t)) = C_0 + C'(\lambda(t) - L)^+ \quad (1)$$

where the base cost C_0 and the overage rate C' are positive constants (x^+ denotes $\max(0, x)$). Threshold L corresponds to the load upon which the utility experiences overages, and therefore raises the cost to dissuade customers from scheduling their jobs (*i.e.*, the second term in sum (1) represents the marginal costs of producing electrical power by an overloaded plant). Otherwise, $\lambda < L$ corresponds to the plant's nominal operational regime.

Finally, we will call "Global Cost" the overall cost (in \$) for the utility:

$$\text{GC} := \int_{t=0}^T \sum_n D_n(t) C_L(\lambda(t)) dt = \int_{t=0}^T \lambda(t) C_L(\lambda(t)) dt$$

With a ramp pricing scheme, this global cost becomes:

$$\text{GC}_{\text{ramp}} = \text{GC}_0 + C' \int_0^T \lambda(t)(\lambda(t) - L)^+ dt. \quad (2)$$

where $\text{GC}_0 := C_0 \sum_n d_n \tau_n$ is a schedule-independent incompressible cost.

B. Non-triviality Criterion

We are interested in scenarios where there is too much demand for the system to avoid overages, and so it has to cope with such situations. A simple way to insure this is for jobs to meet the following criterion:

$$\sum_n d_n \tau_n > LT. \quad (3)$$

III. COMPLETE KNOWLEDGE SETTING

In this section, we survey scheduling when all the jobs' characteristics (d_n, τ_n) are known, either to all players or to a single entity who tries to find an optimal schedule for the whole system. This is for example the purpose of the global controller in [9]. Web portals like Google PowerMeter [10], OPOWER [11] or CustomerIQ [12] also centralize energy consumption data about their users, which they can use thereafter to derive an efficient schedule and advise consumers to conform to it¹.

We will first remind that finding an optimal schedule is an NP-hard problem and discuss a greedy approach for it. We will then consider a *distributed* algorithm set-up as a game between consumers, and derive an optimal strategy for it.

A. NP-hardness

When load profiles (d_n, τ_n) are different for different users, the problem of minimizing GC_{ramp} over all start times $\{s_n\}_{n=1}^N$ is NP-hard [13]. Even when all durations τ_n are equal (or similarly all d_n are equal), finding an optimal schedule is still an NP-hard problem (*i.e.*, the BIN PACKING problem).

B. Discussion : Greedy approach

Since the overall problem is NP-hard, one can consider approximating its optimal solutions, *e.g.*, using well-known metaheuristics such as *simulated annealing*. Though we won't investigate how these techniques would perform, we will give an incremental greedy solution that may get trapped in suboptimal local extrema.

We consider inserting jobs in a given order $i_1, \dots, i_N$ and denote by λ_k the load profile after jobs $i_1, \dots, i_k$ have been scheduled (with $\lambda_0 \equiv 0$). Given λ_k , we want to schedule the $(k+1)^{\text{th}}$ job so as to minimize the global cost incurred by λ_{k+1} . Let $I_{k+1}(s) := [s; s + \tau_{k+1}] \cap \{t | \lambda_k(t) > L - d_{k+1}\}$, $\lambda_k^+ = (\lambda_k - L)^+$ and $\lambda_k^- = (L - \lambda_k)^+$. Then, one can show that minimal GC_{ramp} for λ_{k+1} is achieved when s_{k+1} minimizes:

$$\int_{I_{k+1}(s_{k+1})} [(d_{k+1} - \lambda_k^-)^2 + 2d_{k+1}\lambda_k^+ + L(d_{k+1} - \lambda_k^-)]$$

C. Distributed Action

Now we assume consumers have complete knowledge of each others demands and play a game where they seek to minimize the global cost GC_{ramp} . We provide an effective strategy for players, derived in a pessimistic setting, which turns out to be efficient at *peak shaving* and yields very good results in practice (see section VI).

Here "complete" knowledge means players will either communicate their demand profiles or make inferences about others demands based on repeated observations (*e.g.*, night after night of $[0, T] = [12 \text{ AM} - 6 \text{ AM}]$ activity).

To incent customers to minimize GC_{ramp} (which correspond to the actual cost of supplying their demand, or an upper

¹Note that the data rates associated with this framework are so small (at most a few kilobits/second per user) that communication and security overhead is negligible.

bound of it), utilities may charge customer i with an amount b_i proportional to both the energy he consumed and the global cost, e.g.,

$$b_i := \frac{d_i \tau_i}{\sum_j d_j \tau_j} \times \text{GC}_{\text{ramp}} = C_0 d_i \tau_i \times \frac{\text{GC}_{\text{ramp}}}{\text{GC}_0}, \quad (4)$$

where $C_0 d_i \tau_i$ is the *minimal possible cost* for scheduling player i 's job.

In what follows, we denote by F_i the Cumulative Distribution Function (CDF) of the start time of player i , $F_i(t) := \mathbb{P}[s_i \leq t]$, and f_i the density of dF_i . We also define $\phi_i(t) := F_i(t) - F_i(t - \tau_i)$ which is the probability of job i being active at time t .

Considering equation (2), we can upper bound GC_{ramp} as follows:

$$\text{GC}_{\text{ramp}} \leq \text{GC}_0 + \int_0^T \lambda(t)^2 dt =: \text{GC}_{\text{bound}}.$$

Let us consider a game where users seek to minimize the expected value of GC_{bound} , which is the same as minimizing $\mathbb{E}[\int \lambda^2]$. Since $\int \lambda^2 = \int (\lambda - \mu)^2$ plus a constant, where μ denotes the temporal mean of λ , this goal is closely related to peak shaving.

For any user i , let us define:

$$H_i(t) := \sum_{j \neq i} d_j (\phi_j(t) - \phi_j(t + \tau_i)),$$

which does not depend on F_i . We have:

$$\begin{aligned} \mathbb{E} \left[\int_0^T \lambda(t)^2 dt \right] &= \sum_{i,j} d_i d_j \int_0^T \phi_i(t) \phi_j(t) dt \\ &= \sum_i d_i \int_0^T F_i(t) H_i(t) dt =: \sum_i \gamma_i. \end{aligned}$$

The game between customers goes like this: users play asynchronously, and at his/her turn, player i updates F_i in order to minimize $\gamma_i \propto \int F_i H_i$.

Claim 1. *The optimal CDF F_i^* minimizing $\int F_i H_i$ for any given (right-continuous) function H_i is an indicator $F_i^*(t) = \mathbf{1}_{\{t \geq s_i\}}$ for some $s_i \in [0, T - \tau_i]$.*

To show this property, let us remark a few facts.

Lemma 1. *For any CDF F_i , there exist a staircase CDF $\hat{F}_i$ such that $\int \hat{F}_i H_i \leq \int F_i H_i$.*

Proof: Since H_i is right continuous, one can take a subdivision $0 = r_0 < r_1 < \dots < r_n = T$ of $[0, T]$ such that H_i is of constant sign on subintervals $I_k := [r_k, r_{k+1}[$, but changes sign between consecutive subintervals. Now define $\hat{F}_i$ on I_k as $\max_{I_k} F_i$ if H_i is negative on I_k , and $\min_{I_k} F_i$ otherwise (see Figure 1). This definition yields a new CDF such that $\hat{F}_i(t) H_i(t) \leq F_i(t) H_i(t)$ for all $t \in [0, T]$. ■

Lemma 2. *For any staircase CDF F_i , there exists a “one step” CDF F_i^* such that $\int F_i^* H_i \leq \int F_i H_i$.*

Proof: From the previous lemma, we can suppose without loss of generality that F_i is a staircase CDF, so that $F_i(t) =$

Fig. 1. Staircase CDF optimization.

$\sum_k p_k \mathbf{1}_{\{t \geq r_k\}}$ where $\sum_k p_k = 1$. Thus, $\int F_i H_i = \sum_k p_k A_k$ with $A_k := \int_{r_k}^T H_i(t) dt$. This is just a convex combination of real constants: if we denote by m the index of the minimum A_k , $s_i := r_m$ and $F_i^*(t) = \mathbf{1}_{\{t \geq s_i\}}$, then $\int F_i^* H_i = A_m \leq \sum_k p_k A_k \leq \int F_i H_i$. ■

Hence, given H_i , there is an optimum F_i^* which is an indicator $F_i^*(t) := \mathbf{1}_{\{t \geq s_i\}}$, where we know how to compute s_i from H_i . Furthermore,

$$\int F_i^* H_i = \int_{s_i}^{s_i + \tau_i} \sum_{j \neq i} d_j \phi_j(t) dt,$$

which means the best move for player i is to schedule his job deterministically at a time minimizing the (weighted) sum of the probabilities of other jobs being active during his span $[s_i, s_i + \tau_i]$.

This game seeks to minimize $\sum_i \gamma_i$ by optimizing each γ_i iteratively. It does not necessarily lead to the optimal solution since re-scheduling job i may increase any γ_j for $j \neq i$, yet we will see in section VI that it achieves its goal pretty well in practice.

IV. PARTIAL KNOWLEDGE SETTING

In this section, we suppose players do not share information about each others demands (for privacy reasons), but can still make inferences through the instantaneous total load $\lambda(t)$ which is assumed actively communicated by the network.²

We consider an iterative decision process where, at time t , user i decides (stochastically) whether to schedule his job or not according to:

- his own parameters (d_i, τ_i) ,
- the past load profile $\{\lambda(t'), t' < t\}$.

Concerning the load profile, we will focus on protocols where the decision at time t only depends on the last known value of the load $\lambda(t^-)$.

In what follows, we suppose that all jobs' durations τ_n are integer multiples of a unit time slot duration τ_0 dividing T , so that we can without loss of generality schedule jobs at times multiples of τ_0 .

²We hence suppose that the utility is able to measure the effective state of the grid and compute its load, which is not a minor hypothesis since recent work [14] highlighted flaws in the *state estimation* techniques currently in use.

A. ALOHA Strategy

The first strategy we propose is inspired by the slotted ALOHA protocol [15]. At each time step, if his job has not been scheduled yet, player i applies the following decision procedure, which is parametrized by $0 < q_i < p_i < 1$:

Algorithm IV.1 ALOHA decision procedure for player i

```

if  $t = T - \tau_i$  (last possible scheduling slot) then
   $s_i \leftarrow t$ 
else if  $\lambda(t^-) + d_i \leq L$  then
   $s_i \leftarrow t$  with probability  $p_i$ 
else
   $s_i \leftarrow t$  with probability  $q_i$ 
end if

```

Parameters p_i should be low enough to avoid customers synchronization, but high enough to allow most of the jobs to induce no overage (keep $\lambda < L$).

When L is far below the mean load $\frac{1}{T} \int \lambda$ and all $q_i = 0$, the policy may keep too many jobs for the end, resulting in peak loads at times close to T . Suitable values $q_i > 0$ help deal with this unwanted behavior.

B. Decision Density

A way to generalize this approach is to set up a scheduling decision function for player i , $g_i(t) := g(\lambda(t), t, d_i, \tau_i) \in [0, 1]$, where we assume the form of g is the same for all players. Player i will therefore start at time t with probability $g_i(t)$, decisions being independently made by all players.

For example, under this formulation, the decision density for the ALOHA strategy is

$$g(\lambda, t, d, \tau) := \text{mux}(t = T - \tau, 1, \text{mux}(\lambda + d \leq L, p_i, q_i)),$$

where mux is the *multiplexer* function ($\text{mux}(c, a_1, a_2) := a_1$ if c is true, and a_2 otherwise). Reasonable assumptions about g include:

- g increases with t ;
- $g \rightarrow 1$ when $t \rightarrow T - \tau_i$;
- $g < 1$ when $\lambda \ll L$ and $t \ll T$; and
- g decreases with λ when $\lambda > L$.

With this in mind, we devised a new stochastic strategy improving the ALOHA one.

C. Time/Slackness Strategy

One of the issue of the ALOHA strategy lies in the way it discriminates jobs, since it focuses on the instantaneous load d_i and only takes τ_i into account as $t \rightarrow T - \tau_i$. To remediate this shortcoming, we instead use the *slackness* σ defined when $\lambda < L$ and $t < T - \tau$ as:

$$\sigma(\lambda, t, d, \tau) := \frac{d\tau}{(L - \lambda)(T - \tau - t)},$$

i.e., the ratio of the job's overall energy consumption $d\tau$ and the residual energy $(L - \lambda)(T - \tau - t)$ which corresponds to the energy available with no overage under the assumption that λ stays constant.

Now define $g_1(t) := \left(\frac{t}{T - \tau}\right)^\alpha$. We propose to use the simple density:

$$g(t, \sigma) = g_1(t) + (1 - g_1(t))(\beta + \gamma \cdot \mathbf{1}_{\{0 < \sigma < 1\}}), \quad (5)$$

so that just three parameters α , β and γ are in play. We call the associated policy *Time/Slackness*, since it consists of a BERNOULLI trial over $g_1(t)$ (ensuring the task is scheduled in time), followed by another trial based on slackness, giving a boost to the tasks for which there is enough residual energy.

Experimental results (see section VI) confirm this new strategy yields better results than the ALOHA one, suggesting energy is a better discrimination criterion than power.

V. BLIND SETTING

In this section, we survey a power-grid setting where there is no communication layer between users. We also assume all customers have the same demand profile (d, τ) and decide to schedule their jobs at times multiples of τ (where $T = K\tau$, $K \in \mathbb{N}$) in a discrete time setting. We show that, in this simplified setting, the best strategy for customers is to choose their time slot uniformly at random.

Note that, here, broadcasting $\lambda(k\tau^-)$ to the users would be useless since this value is independent from $\lambda(k\tau)$.

Claim 2. *The expected overall cost $\mathbb{E}[\text{GC}]$ is minimized when (independent) start times are chosen uniformly distributed on $\{k\tau, k \in \llbracket 0, K - 1 \rrbracket\}$.*

Proof: Let p_k be the Probability Mass Function (PMF) of start-time s , common to all customers by symmetry. Limited information implies independent scheduling decisions. Therefore, the number of customers that select a given service epoch is binomially distributed, *i.e.*,

$$\mathbb{P}[\lambda(k\tau) = nd] = \binom{N}{n} p_k^n (1 - p_k)^{N-n}.$$

So, the overall expected cost

$$\sum_{n=1}^N \sum_{k=0}^{K-1} \tau n d C_L(nd) \binom{N}{n} p_k^n (1 - p_k)^{N-n} =: \sum_{k=1}^K G(p_k)$$

is to be minimized subject to the PMF $\underline{p}$ in the K -dimensional simplex $\sum_{k=1}^{K-1} p_k = 1$. Note that G , defined by swapping order of summation, does not depend on the time-index k . The Lagrangian for this problem is

$$\sum_{k=0}^{K-1} G(p_k) + c \left(1 - \sum_{k=0}^{K-1} p_k \right),$$

with Lagrange multiplier c , leading to the first-order necessary conditions whose solution is

$$\forall k \in \llbracket 0, K - 1 \rrbracket, \quad p_k = (G')^{-1}(c),$$

i.e., p_k is constant in k . (One can check that G' is indeed bijective.) Condition $\sum_{k=1}^K p_k = 1$ therefore yields $p_k = 1/K$, so $\underline{p}$ is the PMF of a uniform distribution. ■

Recall that the conditions in which this uniform policy is optimal are different from the other settings we studied: we

are here in a power-grid with no communication layer. We use this policy only as a reference in our numerical experiments.

VI. NUMERICAL EXPERIMENTS

Experiments on DR scenarios can involve embedding users in one of the IEEE test systems used in [14] (which can be found in MATPOWER, a MATLAB package) with a shared-resources game between them, taking into account on the characteristics of the buses. For our experiments we chose the simpler model, used in [7], of a local distribution network with one energy source and several load subscribers (see Figure 2).

Fig. 2. Setup of the test system.

We implemented a simulator in PYTHON working on a six hours time frame divided into a customizable number of time slots. It implements the different policies we encountered:

- *Uniform*: the best solution in the blind setting;
- *ALOHA I*: the ALOHA strategy where all users share the same probabilities $p_i = p$ and $q_i = 0$;
- *ALOHA II*: same strategy with $\forall i, p_i = p > q_i = q > 0$;
- *Time/Slackness*: the policy from section IV with decision density (5) parametrized by α, β and γ ;
- *Game*: the game from section III.

Our simulator is open-source and available online at [16].

We set-up different test settings and ensured criterion (3) was met in each of them. For the Game policy, optimal behavior was reached for an average of 3 moves per player, which suggests this strategy converges quickly.

For the ALOHA and Time/Slackness policies, we manually chose good values of the parameters for each setting. In fact, all settings turned out to share approximately the same efficient values of the parameters, *i.e.*,

- *ALOHA I*: $p \approx 0.2$
- *ALOHA II*: $p \approx 0.145$ and $q \approx 0.0175$
- *Time/Slackness*: $\alpha \approx 45$, $\beta \approx 0.006$ and $\gamma \approx 0.12$. This value of α implies time considerations are neglected while $t < 90\% T$. In the last decile however, $g_1(t)$ yields more balanced schedules than a simple time-over check.

A. Residential Setting

The first scenario we considered is the case where all jobs have the same duration τ and instantaneous cost d , *i.e.*, a residential area where houses have the same first-order load profile. For 1,000 users with a demand profile of 20 kW for 1 hour, the system's nominal load was set to $L = 3,000$ kW, while we chose $C_0 = 2.8 \times 10^{-6}$ \$/kW/s (which is the first step in the model used at BC Hydro [8]) and $C_1 = 2.8 \times 10^{-8}$ \$/kW²/s.

Fig. 3. Average global costs in the residential setting.

We focused on the global cost GC_{ramp} experienced by the utility for each policy. Results averaged over several runs are shown in Figure 3 with confidence bars.

The best load profile for this configuration is a flat one. The Game strategy achieves a nearly optimal result, which comes from the fact that it is the only policy with enough information to *actively* seek a flat profile. Stochastic heuristics just try to approximate it (again with only limited information) while the uniform one tends to underload the borders (times close to 0 and T). We also see that, in this setting where all customers are identical, the uniform policy yields better schedules than the heuristics from section IV.

B. Heterogeneous Setting

We also investigated the case where a lot of different profiles coexist on the network, including:

- a few “big” users demanding 100–400 kW for 2-5 hours,
- about 100 users demanding 10–50 kW for 1-3 hours,
- about 100 users demanding 10 kW for about 1 hour,
- a few “peak” users demanding > 800 kW for < 30 min.

System-wide parameters were set to $L = 1000$ kW, $C_0 = 2.8 \times 10^{-6}$ \$/kW/s and $C_1 = 2.8 \times 10^{-7}$ \$/kW²/s. Results are show in Figure 4.

Again, the Game policy achieves the best behavior, but this time our heuristics perform better than the uniform strategy. Sample load profiles (which we won't produce here but are available online at [16]) indicate that:

- the Uniform strategy tends to make expensive mistakes, scheduling “big” players when the grid is already stressed and unloading the borders;
- ALOHA I achieves a rather flat load, but is likely to keep big players for the end, yielding a final peak;
- ALOHA II partially avoids this behaviour when q is high enough, but does not discriminate players in case of overage;

Fig. 4. Average global costs in the heterogeneous setting.

- Time/Slackness is the best of the three heuristics and achieves a good compromise in scheduling both small and big users at each time step.

C. Peak-to-average ratio

We mentioned that the Game strategy is effective at peak shaving. We may illustrate this with figures, comparing its *peak-to-average* ratio (PAR) with the one of the Uniform strategy, averaged over several runs.

	Game	Uniform	Improvement
Domestic	1.02	1.29	20.9%
Heterogeneous	1.38	2.28	39.5%

TABLE I
PAR BENEFITS OF THE GAME POLICY.

VII. CONCLUSIONS

In this paper, we studied Demand Response problematics on multiple architectures for the dynamic pricing scheme (1). We saw that the general problem of finding an optimal schedule under this cost is NP-hard. We then surveyed different strategies depending on the degree of information sharing in the network.

When all demand information is shared, we proposed a game played by customers yielding good results in practice. When only the instantaneous load is known, we provided distributed strategies using the instantaneous load to reduce their costs. To experimentally evaluate all these policies, we developed our own open-source simulator which we released at [16]. Simulation results confirm that all these strategies perform better than when consumers do not communicate, especially the distributed game which significantly reduces the global cost and PAR, given the required information is available.

VIII. FUTURE WORK

In our jobs model, we chose a first-order approximation of the aggregate profile. Further study could take into account several devices per user with interruptible, non-constant load profiles and power consumptions ranging from 0.01 kW (light bulb) to 1 kW (dishwasher, clothes dryer). Furthermore, multiple hierarchy levels can be considered: devices, users, substations (trying to optimize their own costs, providing incentives to the customers).

Also, in cost (1) we took a threshold L constant within the timeframe of study, which does not encompass the integration of renewable energy sources that may deliver additional power during short time intervals. It would thus be interesting to survey how heuristics perform with a time-varying threshold $L(t)$.

Finally, though the smart grid marketplace information flow is small in volume (and can therefore be “strongly” authenticated with low cost), security problems such as *false data injection* attacks [14] arise, which warrant additional consideration.

REFERENCES

- [1] R. Krishnan, “Meters of Tomorrow”, *IEEE Power and Energy Magazine*, pp. 92-94, Mar. 2008.
- [2] U.S. Department of Energy, *The Smart Grid: An Introduction*, 2009.
- [3] L. D. Kannberg, M. C. Kintner-Meyer, D. P. Chassin, R. G. Pratt, J. G. DeSteele, L. A. Schienbein, S. G. Hauser, W. M. Warwick, *GridWise: The Benefits of a Transformed Energy System*, p. 25, Nov. 2003. Pacific Northwest National Laboratory under contract with the United States Department of Energy. <http://arxiv.org/pdf/nlin/0409035>
- [4] A. Ipakchi and F. Albuyeh, “Grid of the future”, *IEEE Power and Energy Magazine*, pp. 52-62, Mar. 2009.
- [5] M. Pipattanasomporn, H. Feroze, S. Rahman, “Multi-Agent Systems in a Distributed Smart Grid: Design and Implementation”, *Proc. IEEE PES 2009 Power Systems Conference and Exposition*, Mar. 2009.
- [6] S. Darby, *The Effectiveness of Feedback on Energy Consumption*, Environmental Change Institute, University of Oxford.
- [7] A.-H. Mohsenian-Rad, V. W.S. Wong, J. Jatskevich, R. Schober, “Optimal and Autonomous Incentive-based Energy Consumption Scheduling Algorithm for Smart Grid”, *Proc. of IEEE PES Conference on Innovative Smart Grid Technologies*, Jan. 2010.
- [8] BC Hydro, *Electricity Rates*, 2009.
- [9] A. Molderink, V. Bakker, M. G.C. Bosman, J. L. Hurink, G. J.M. Smit, “A Three-Step Methodology to Improve Domestic Energy Efficiency”, *Proceedings of the 2010 IEEE Innovative Smart Grid Technologies Conference*, Jan. 2010.
- [10] Google PowerMeter, <http://www.google.com/powermeter>
- [11] OPOWER, <http://www.opower.com>
- [12] Silver Sprint CustomerIQ web portal, <http://www.silverspringnet.com/products/customeriq.html>
- [13] G. Ausiello, P. Crescenzi, G. Gambosi, V. Kann, A. Marchetti-Spaccamela, M. Protasi, *Complexity and Approximation: Combinatorial Optimization Problems and Their Approximability Properties*, Springer-Verlag, Berlin, 1999.
- [14] Y. Liu and P. Ning and M. K. Reiter, “False Data Injection Attacks against State Estimation in Electric Power Grids”, *Proc. of the 16th ACM conference on Computer and communications security*, Nov. 2009.
- [15] N. Abramson, “The ALOHA System: another alternative for computer communications”, *AFIPS '70 (Fall): Proceedings of the November 17-19, 1970, fall joint computer conference (ACM)*, pp. 281-285.
- [16] *DR StratComp*, a Demand-Response Strategies Comparator. <https://www.ohloh.net/p/dr-stratcomp>