Distributed Scheduling Algorithms for the Smart Grid

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

George Kesidis, *Senior Member, IEEE*
CS&E and EE Depts
The Pennsylvania State University
University Park, PA, 16802
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 aggregate demand profile depending on how much information they share. When customers can share all their load profiles, we describe 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 has been developed according to a static, centralized structure: remote power plants transmit electrical power which is carried on long-distance high-voltage lines to substations (transmission network), which in turn adapt and deliver power 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. With the progressive integration of *smart meters* [11] and communicating appliances, and deployment of associated communication protocol and "demand response" architectures, this "blind" system is migrating to a more decentralized smart grid [24], which is foreseen as a way to save billions of dollars in energy consumption [10].

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 [12];
- 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 [20]; and
- Energy savings: studies [6] already suggest that giving customers access to real-time consumption information yields significant savings.

A shorter version of this paper was presented at IEEE SmartGridsComm, Gaithersburg, MD, Oct. 2010. The corresponding author is G. Kesidis.

Such load management is becoming even more crucial as *plug-in electric or hybrid electric vehicles* (PEVs or 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 [12].

Therefore, the design of appropriate incentives and efficient energy consumption scheduling algorithms is a main issue for the deployment of the upcoming smart grid. Such adaptive frameworks are described in [2], a survey with large scope and an emphasis on decision making with incomplete information (*i.e.*, a stochastic setting). For example, several authors have considered scheduling "smart" users of a spot-price system, *e.g.*, [14] used the Nash certainty equivalence principle for the LQG framework [9]. In [5], market models are described wherein aggregate demand of heterogeneous consumers is matched to power supply at equilibrium. In [17], the authors address scheduling 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 Demand Response architecture and the degree of knowledge appliances have on the state of the network. As in [17], we embed consumers in a local distribution network consisting of a single energy source (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. In a centralized system, there are dedicated circuits at consumers' premises which are under the control of the grid, *i.e.*, the grid can cycle power to these circuits to peak-shave aggregate demand to avoid overages and, given reliable information of consumer demand, can better schedule supply *a priori* to each consumer. In a distributed system, the smart grid would signal the consumers with periodically updated spot prices, or more detailed aggregate load information, allowing the consumers themselves to schedule the flexible portion of their demand to avoid overage costs. In the distributed setting, the consumers may need to configure scheduling software to manage their demand for power while they sleep, *e.g.*, that of PEVs or PHEVs.

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 even given complete knowledge and describe a simple greedy heuristic. In section IV, we briefly discuss "variable" supply that results from power sources such as wind or solar, or from bottlenecks of supply in the distribution system itself. A distributed, cooperative game between consumers with complete knowledge is studied in section V. We provide stochastic policies for a distributed setting with only partial knowledge (aggregate load) in section VI. In section VII, we derive the best distributed policy for synchronized users in a power grid with no knowledge. Both the idealized complete-knowledge cooperative game of section V and the fully "blind" setting of section VII are used as references in section VIII where we provide experimental results. We draw conclusions in section IX.

II. Problem set-up

We consider a finite, T-hour time period, *e.g.*, $T = 6$ hours from midnight to 6AM, during which N customers need to (perhaps automatically) schedule their electrical jobs.

A. Loads and Costs

The n^{th} customer is assumed to have a 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:

$$D_n(t) \;=\; d_n \mathbf{1}_{\{s_n \le t \le s_n + \tau_n\}}. \tag{1}$$

We assume $0 \le s_n \le T - \tau_n$ and that once the service start time s_n is selected, it cannot be interrupted by the user (see, *e.g.*, [19] for consideration of deferrable demand). The total instantaneous load on the network is then

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

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 (5)). C depends on the instantaneous load $\lambda(t)$ and is likely to depend on additional system parameters. For example, the two-step conservation rate model used by BC Hydro [4] (parametrized by load a power threshold L) is

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

Such parameters may be set, *e.g.*, by a Day-Ahead Demand Response Program (DADRP) of the smart grid Demand Response framework, *e.g.*, [16]. 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 $L > 0$:

$$C_L(\lambda(t)) \;=\; C_0 + C'(\lambda(t) - L)^+, \tag{2}$$

where the base cost C_0 and the overage rate C' are positive constants, and $x^+ := \max(0, x)$. Threshold L corresponds to the load upon which the utility experiences overages, and therefore raises the cost to dissuade customers from scheduling their jobs in a disributed control setting. Otherwise, $\lambda < L$ corresponds to the grid's nominal operational regime.

Finally, we will call "Global Cost" the overall cost (in \$) for the utility (or that borne by the consumers):

$$\mathsf{GC} := \int_0^T \sum_n D_n(t) C_L(\lambda(t)) \mathrm{d}t = \int_0^T \lambda(t) C_L(\lambda(t)) \mathrm{d}t$$

With a ramp pricing scheme, this global cost becomes:

$$\mathsf{GC}_{\mathsf{ramp}} = \mathsf{GC}_0 + C' \int_0^T \lambda(t)(\lambda(t) - L)^+ \mathrm{d}t. \tag{3}$$

where $\mathsf{GC}_0 := C_0 \sum_n d_n \tau_n$ 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. That is, assume

$$\sum_n d_n \tau_n > LT. \tag{4}$$

III. Complete Knowledge Setting: Centralized Control

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 (*i.e.*, the grid itself) who tries to find an optimal schedule for the whole system. This is for example the purpose of the global controller in [18]. Web portals like Google PowerMeter [8], OPOWER [21] or CustomerIQ [22] also centralize energy consumption data about their users, which they can use thereafter to derive an efficient schedule and advise consumers to conform to it[1].

We will first remind that finding an optimal schedule is an NP-hard problem and discuss a simple greedy approach.

A. Greedy approach

When load profiles (d_n, τ_n) are different for different users, the problem of minimizing $\mathsf{GC}_{\mathsf{ramp}}$ over all start times $\{s_n\}_{n=1}^N$ is NP-hard [3]. 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).

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 (one can heuristically deal with local extrema by repeating greedy optimization at different

[1]Note that the data rates associated with the smart grid frameworks considered herein are small (at most a few kilobits/second per user), communication and strong authentication overhead (to *e.g.*, prevent or detect false data injection [13]) will be negligible.

randomly chosen initial job orderings, *i.e.*, random restart, not just a single intial ordering of jobs, say according to decreasing $d\tau$).

We consider inserting jobs in a given order $i_1, \ldots, i_N$ and denote by λ_k the load profile after jobs $i_1, \ldots, i_k$ have been scheduled (with $\lambda_0 \equiv 0$). Given λ_k, we want to schedule the job $k+1$ job so as to minimize the global cost incurred by λ_{k+1}.

Claim 1. *Given λ_k, minimal $\mathsf{GC}_{\mathsf{ramp}}$ overage for λ_{k+1} is achieved when s_{k+1} minimizes over s:*

$$\int_s^{s+\tau_{k+1}} [(\lambda_k + d_{k+1})(\lambda_k + d_{k+1} - L)^+ - \lambda_k(\lambda_k - L)^+]\mathrm{d}t.$$

Proof: By (3), the overage (divided by C') given λ_k is

$$\int_0^T \lambda_k(t)(\lambda_k(t) - L)^+\mathrm{d}t.$$

If job $k+1$ is scheduled at time s, the overage is

$$\int_0^T (\lambda_k(t) + D_{k+1}(t))(\lambda_k(t) + D_{k+1}(t) - L)^+\mathrm{d}t.$$

After substituting (1), the display in the claim is simply the difference of the previous two displays. ■

IV. Discussion: Variable supply

The addition of a variable power supply may be used to mitigate overages.

A. Wind or solar power

Consider the addition of a known "variable" power supply $S(t)$ to the grid. Such supply could be due to wind or solar power, the amount of which is only precisely known over a short time interval, say on an hourly basis [25]. Plug-in electric (or hybrid electric) vehicles are somewhat flexible energy storage systems which *are* capable of exploiting such variable energy resources [15], whether centralized or decentralized control is in effect. In, *e.g.*, [19], a centralized demand peak-shaving problem was considered assuming deferrable demand using a more realistic (non-constant) power-consumption profile for plug-in electric vehicles than (1).

Though variable supply may be predictable long-term with sufficient accuracy for demand scheduling, we will not explore this possibility herein. For nondeferrable demand (the context of this paper), note that in the previous claim, we do *not* assume that the scheduling time s_k of the k^{th} is monotonically nondecreasing in k. So, if new supply $S(t)$ is available at time t, the grid could *reschedule* demand from $[t, T]$ replacing the overage threshold L by $L + S(t)$ in the previous claim. A suitable distributed framework for the exploitation of variable supply is that of section VI below.

B. Bottlenecks of supply

In a local distribution system, substations provide power to distribution feeders along which step-down transformers are used to serve 3 or 4 consumer premises. Dealing with idealized "graphical" constraints of distribution is considered in, *e.g.*, [23].

A simple topological scenario involves two "connected" subgroups of consumers, A and B, and a power-transfer threshold $K_{BA} < L_B$. If the instantaneous aggregate demands $\lambda_A > L_A$ and $\lambda_B < L_B$, then the overage in A is reduced from $\lambda_A - L_A$ to

$$(\lambda_A - \min\{K_{BA},\ L_B - \lambda_B\} - L_A)^+,$$

i.e., only a maximal amount of power K_{BA} can be transferred from B to A (and vice versa if there is an overage at B and excess supply at A).

The additional supply to A $S_{BA}(t) := \min\{K_{BA},\ L_B - \lambda_B(t)\}$ may be known and can be easily factored when the jobs of A and B are *jointly* scheduled, as in the previous claim. So, this is a special case of "Demand Side Ancillary Service Program" (DSASP) of the smart grid Demand Response framework, *e.g.*, [16]. Otherwise, from the point of view of A, the supply S_{BA} could be viewed as variable, again *c.f.* the remark in section VI below.

V. Complete Knowledge Setting: Distributed Control

Though the following scenario is unlikely owing to privacy concerns, we develop it for use as a reference in the numerical study given below. Here assume consumers have complete knowledge of each others' demands and play a game where they seek to minimize the global cost $\mathsf{GC}_{\mathsf{ramp}}$. We provide an effective strategy for players, derived in a pessimistic setting, which turns out to be efficient at *peak shaving* [26] and yields very good results in practice (see section VIII).

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] = [$12 AM–6 AM$]$ activity).

To incent customers to minimize $\mathsf{GC}_{\mathsf{ramp}}$ (which correspond to the actual cost of supplying their demand, or an upper 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.*,

$$b_i \ := \ \frac{d_i\tau_i}{\sum_j d_j\tau_j} \times \mathsf{GC}_{\mathsf{ramp}} \ = \ C_0 d_i\tau_i \times \frac{\mathsf{GC}_{\mathsf{ramp}}}{\mathsf{GC}_0}, \qquad (5)$$

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, $F_i(t) := \mathsf{P}[s_i \le t]$, and f_i its density. 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 (3), we can upper bound $\mathsf{GC}_{\mathsf{ramp}}$ as follows:

$$\mathsf{GC}_{\mathsf{ramp}} \le \mathsf{GC}_0 + \int_0^T \lambda(t)^2\mathrm{d}t =: \mathsf{GC}_{\mathsf{bound}}.$$

Let us consider a game where users seek to minimize the expected value of $\mathsf{GC}_{\mathsf{bound}}$, which is the same as minimizing $\mathsf{E}[\int \lambda^2]$. 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 each user i, define:

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

which does not depend on F_i. We have:

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

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 2. *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 $\widehat{F}_i$ such that $\int \widehat{F}_i H_i \leq \int F_i H_i$.*

Proof: Since H_i is right continuous, one can take a subdivision $0 = r_0 < r_1 < \ldots < r_n = T$ 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 $\widehat{F}_i$ 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 $\widehat{F}_i(t) H_i(t) \leq F_i(t) H_i(t)$ for all $t \in [0, T]$. ∎

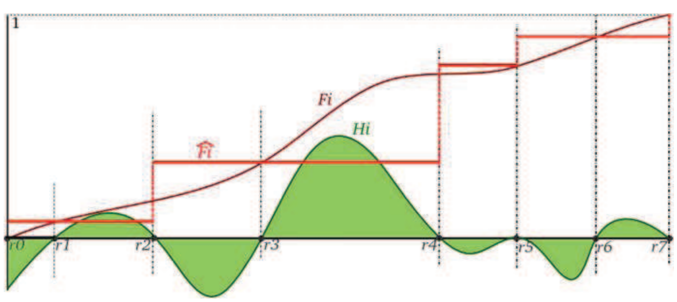

Fig. 1. Staircase CDF optimization.

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 $F_i(t) = \sum_k p_k \mathbf{1}_{\{t \geq r_k\}}$ where $\sum_k p_k = 1$. Thus, $\int F_i H_i = \sum_k p_k A_k$ with $A_k := \int_{r_k}^T H_i(t)\mathrm{d}t$. 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,

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

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 the numerical study of section VIII that it achieves its goal pretty well in practice.

VI. Partial Knowledge Setting: Distributed Control

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 know 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 discrete times that are multiples of τ_0.

A. ALOHA Strategy

The first strategy we propose is that of slotted ALOHA [1]. 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 VI.1 ALOHA decision procedure for player i

```
if t = T − τ_i (last possible scheduling slot) then
    s_i ← t
else if λ(t⁻) + d_i ≤ L then
    s_i ← t with probability p_i
else
    s_i ← t with probability q_i
end if
```

Parameters p_i should be low enough to avoid customers synchronization, but high enough for minimal overage (again, when $\lambda > L$).

[2] We 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 [13] highlighted flaws in the *state estimation* techniques currently in use.

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 of the interval, 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

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

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

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

With this in mind, the following strategy improves upon that based on ALOHA.

C. Time/Slackness Strategy

The ALOHA focuses on the instantaneous load d_i and only takes τ_i into account as $t \to T - \tau_i$. 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:

$$g(t, \sigma) = g_1(t) + (1 - g_1(t))(\beta + \gamma \cdot \mathbf{1}_{\{0<\sigma<1\}}), \qquad (6)$$

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.

The experimental results of section VIII confirm this new strategy yields better results than the ALOHA one, suggesting energy is a better discrimination criterion than power.

D. Remark: Variable Supply

If the extent to which the total demand $\lambda(s)$, $s \in [s, s + \tau]$ is known, so as to make an informed scheduling decision at time t, a variable power supply S is also known, then one can simply replace L by $L + S(t)$ in the above decision density functions.

VII. Blind Setting: Distributed Control

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 $T = K\tau$ for $K \in \mathbb{N}$) 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 3. *The expected overall cost* $\mathsf{E}[\mathsf{GC}]$ *is minimized when (independent) start times are chosen uniformly distributed on* $\{k\tau,\, k \in [\![0, K-1]\!]\}$.

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.*,

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

So, the overall expected cost

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

is to be minimized subject to the PMF $\underline{p}$ in the K-dimensional simplex $\sum_{k=0}^{K-1} p_k = 1$. 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 multipier c, leading to the first-order necessary conditions whose solution is

$$\forall k \in [\![0, K-1]\!], \quad p_k \;=\; (G')^{-1}(c),$$

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: here the grid has no communication layer. We use this policy only as a reference in our numerical experiments.

VIII. Numerical Experiments on Distributed Control

Experiments on Demand Response scenarios can involve embedding users in one of the IEEE test systems used in [13] (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 [17], of a

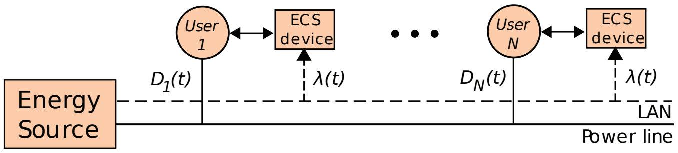

Fig. 2. Setup of the test system.

local distribution network with one energy source and several load subscribers (see Figure 2).

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 distributed 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 VI with decision density (6) parametrized by α, β and γ;
- *Game:* the game from section V.

Our simulator is open-source and available online at [7].

We set-up different test settings and ensured criterion (4) 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 by BC Hydro [4]) and $C_1 = 2.8 \times 10^{-8}$ \$/kW^2/s.

We focused on the global cost $\mathsf{GC}_{\mathsf{ramp}}$ 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 VI.

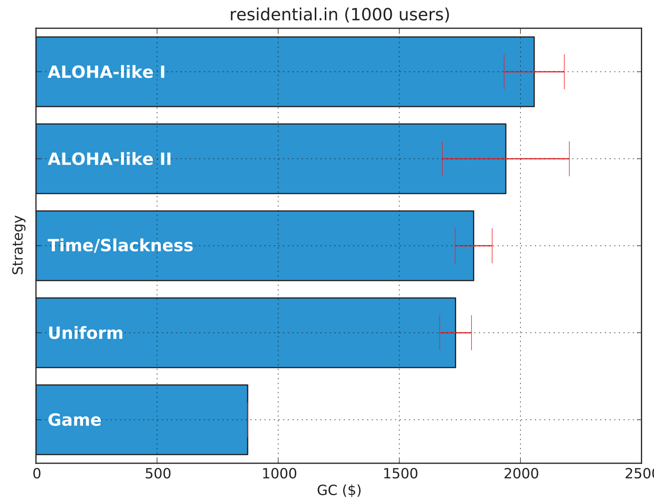

Fig. 3. Average global costs in the residential setting.

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 $C_1 = 2.8 \times 10^{-7}$ \$/kW^2/s. Results are show in Figure 4.

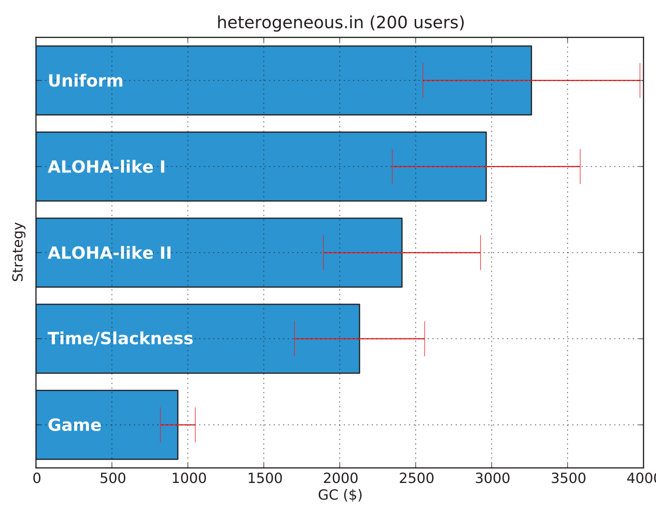

Fig. 4. Average global costs in the heterogeneous setting.

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 [7]) 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;
- 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 the following table, 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.

IX. CONCLUSIONS

In this paper, we studied Demand Response problematics on multiple architectures for the dynamic pricing scheme (2). We saw that the general centralized 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. We also discussed how variable power supply could be exploited in these frameworks. The above methods directly generalize to the scheduling (without interruption) of non-constant power consumptions profiles ranging from 0.01 kW (light bulb) to 1 kW (dishwasher, cloths dryer), including those of plug-in electric or hybrid-electric vehicles.

When all demand information is shared, we proposed a game played by customers yielding good results in our numerical study. When only the instantaneous load (or equivalent spot price information) 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 [7]. Simulation results confirm that all these strategies perform better than when consumers do not communicate, especially the distributed, complete information game which significantly reduces the global cost and PAR, given the required information is available.

REFERENCES

[1] N. Abramson. The ALOHA System: another alternative for computer communications. In *Proc. ACM AFIPS*, pp. 281-285, Nov. 1970.
[2] R.N. Anderson, A. Boulanger, W.B. Powell and W. Scott. Adaptive stochastic control for the smart grid. *Proceedings of the IEEE* **99**(6), June 2011.
[3] G. Ausiello, P. Crescenzi, G. Gambosi, V. Kann. A. Marcheti-Spaccamela, M. Protasi, *Complexity and Approximation: Combinatorial Optimization Problems and Their Approximability Properties*. Springer-Verlag, Berlin, 1999.
[4] BC Hydro. *Electricity Rates*. Available at https://www.bchydro.com/youraccount/content/residential_rates.jsp
[5] L. Chen, N. Li, S.H. Low and J.C. Doyle. Two Market Models for Demand Response in Power Networks. In *Proc. IEEE SmartGridComm*, 2010.
[6] S. Darby. *The Effectiveness of Feedback on Energy Consumption*. Environmental Change Institute, University of Oxford.
[7] SECS: Smart Energy Consumption Strategies. Available at https://gitorious.org/dr-stratcomp
[8] Google PowerMeter. Available at http://www.google.com/powermeter
[9] M. Huang. Large-population LQG games involving a major player: The Nash certainty equivalence principle. *SIAM J. Control Opt.*, **48**(5), 2010.
[10] L.D. Kannberg, M.C. Kintner-Meyer, D.P. Chassin, R.G. Pratt, J.G. DeSteese, L.A. Schienbein, S.G. Hauser, W. M. Warwick, *GridWise: The Benefits of a Transformed Energy System*, p. 25, Nov. 2003. Available at http://arxiv.org/pdf/nlin/0409035
[11] R. Krishnan. Meters of Tomorrow. *IEEE Power and Energy Magazine*, pp. 92-94, Mar. 2008.
[12] A. Ipakchi and F. Albuyeh. Grid of the future. *IEEE Power and Energy Magazine*, pp. 52-62, Mar. 2009.
[13] Y. Liu, P. Ning and M. K. Reiter. False Data Injection Attacks against State Estimation in Electric Power Grids. In *Proc. ACM Conference on Computer and Communications Security (CCS)*, Nov. 2009.
[14] Z. Ma, D. Callaway and I. Hiskens. Decentralized Charging Control for Large Populations of Plug-in Electric Vehicle s. In *Proc. IEEE CDC*, Atlanta, 2010.
[15] S. Meyn, M. Negrete-Pincetic, G. Wang, A. Kowli, and E. Shafieepoorfard. The value of volatile resources in electricity markets. In *IEEE CDC*, Atlanta, Dec. 2010.
[16] F. Rahimi and A. Ipakchi. Overview of Demand Response under the smart grid and market paradigms. In *Proc. IEEE PES Conf. on Innovative Smart Grid Technologies*, Gaithersburg, MD, Jan. 2010.
[17] A.-H. Mohsenian-Rad, V.W.S. Wong, J. Jatskevich, R. Schober. Optimal and Autonomous Incentive-based Energy Consumption Scheduling Algorithm for Smart Grid. In *Proc. IEEE PES Conference on Innovative Smart Grid Technologies*, Jan. 2010.
[18] A. Molderink, V. Bakker, M.G.C. Bosman, J.L. Hurink, G.J.M. Smit. A Three-Step Methodology to Improve Domestic Energy Efficiency. In *Proc. IEEE Innovative Smart Grid Technologies Conference*, Jan. 2010.
[19] G. Pang, G. Kesidis and T. Konstantopoulos. Avoiding Overages by Deferred Aggregate Demand for PEV Charging on the Smart Grid. CSE Tech. Rpt. 11-009, Aug. 4, 2011; Available at http://www.cse.psu.edu/research/publications/tech-reports/2011/CSE-11-009.pdf
[20] M. Pipattanasomporn, H. Feroze, and S. Rahman. Multi-Agent Systems in a Distributed Smart Grid: Design and Implementation. In *Proc. IEEE PES Power Systems Conference and Exposition*, Mar. 2009.
[21] OPOWER. Available at http://www.opower.com
[22] Silver Sprint CustomerIQ web portal. Available at http://www.silverspringnet.com/products/customeriq.html
[23] J.A. Taylor and F.S. Hover. Lift-and-project relaxations of AC microgrid distribution system planning. In *Grand Challenges in Modeling and Simulation*, June 2011.
[24] U.S. Department of Energy. *The Smart Grid: An Introduction*. 2009.
[25] Wind Energy: The Facts. Part 2, Chapter 2 - Understanding Variable Output Characteristics of Wind Power. Available at http://www.wind-energy-the-facts.org/en/part-2-grid-integration/
[26] R.N. Wu, T.H. Lee and E.F. Hill. An investigation of the accuracy and the charac teristics of the peak-shaving method applied to production cost calculations. *IEEE Trans Power Syst*, **4**(3): pp. 1043-1049, 1989.