Characterizing Flexibility of an Aggregation of Deferrable Loads

He Hao and Wei Chen

***Abstract*— Flexibility from distributed deferrable loads presents an enormous potential to provide fast ramping resources that are necessary to vastly integrate renewable energy resources. In this paper, we study aggregation, characterization, and scheduling of a collection of deferrable loads to facilitate integrating renewable generation into the power system. A generation profile is called *exactly adequate* if there exists a scheduling policy that could allocate the power to satisfy the energy requirements of all deferrable loads without surplus or deficit. We provide sufficient and/or necessary characterizations on the adequacy of allocated generation profiles. Moreover, we propose a novel scheduling algorithm to service deferrable loads. Heuristic algorithms such as Earliest Deadline First (EDF) and Least Laxity First (LLF) policies are used in numerical experiments to compare with the proposed algorithm. Extensive simulations show that our scheduling policy generally can fulfill the energy requirements of all loads without surplus or deficit for exactly adequate generation profiles, while the EDF and LLF policies cannot meet this objective in most cases.**

I. Introduction

Renewable energy resources are uncertain, volatile, and intermittent. The vast integration of variable renewable energies into the grid creates daunting challenges for the system operator to maintain the power balance. Increasing reliance on renewable energy supplies will require more flexible ramping capacity to handle the fast ramps resulting from the variability of such uncontrollable resources. In a recent study [1], it is shown that with higher level of renewable penetration in California (e.g. 50%), the main operational challenge for power system stability is overgeneration, which is mainly due to the fast ramping of solar energy resources.

In this paper, we study coordinated aggregation of a large collection of deferrable loads such as residential pool pumps, electric vehicles, and deferrable appliances to provide fast ramping resources to the grid, without vastly increasing procurement of expensive and pollutive generation reserves. This paper is studied under the GRIP (Grid with Intelligent Periphery) architecture [2], which aggregates diverse geographically co-located distributed resources into clusters. Each resource cluster will be managed by a cluster manager who serves as an intermediary interfacing the cluster with the electricity grid.

The cluster manager is responsible for forecasting the aggregate load and renewable generation, purchasing energy in the forward energy market, and buying/selling reserves in the ancillary service markets, in order to procure a total generation profile to service a collection of loads. Moreover, the cluster manager implements local scheduling and control algorithms for distributed loads, and manage the power balance in the cluster while respecting resource constraints at run-time. A generation profile is called *exactly adequate* if there exists a scheduling policy that could allocate the generation to satisfy all loads' energy requirements without surplus or deficit. In this paper, we derive sufficient and/or necessary characterizations on the adequacy of generation profiles for servicing a collection of deferrable loads.

In a recent paper [3], it was shown that for an exactly adequate generation profile, in general there did not exist a feasible *causal* scheduling algorithm that can fulfill the energy requirements of a collection of deferrable loads without surplus or deficit. In this paper, we propose a *non-causal* scheduling algorithm, which we call the Earliest Deadline Priority (EDP) policy. Extensive simulations show that it is generally able to finish all tasks for exactly adequate generation profiles. Heuristic causal algorithms such as Earliest Deadline First and Least Laxity First policies are used in numerical experiments to compare with our proposed algorithm. We show that the EDF and LLF polices generally cannot achieve this goal. A caveat with our proposed method is that this policy is non-causal, which requires accurate prediction of the generation profiles. However, it will become useful when demand response with deferrable loads is bid and scheduled in the day-ahead energy market.

There are several papers that are closed related to this work. These include real-time scheduling of deferrable loads [3], coordination of a large population of electric vehicles [4], [5], aggregating flexibility of thermostatically control loads [6], distributed control of deferrable loads such as pool pumps for ancillary service provision [7], characterizing flexibility of fans in commercial bulding HVAC systems [8], and optimally supplying renewable generation to deferrable electricity loads in the presence of a spot market [9].

In particular, this paper is a follow-up work of [4], [3]. However, compared to [3], [4], there are several significant differences and important contributions in this paper. In [4], when each deferrable load has finite power limit, only a couple of necessary conditions are given for exactly adequate generation profiles. In this paper, we provide three additional necessary conditions. Additionally, if every load has the same rated power, arrival and departure times, we provide a sufficient characterizations on the set of adequate generation profiles. Moreover, our results hold when the energy requirements of deferrable loads are not fixed but flexible. Additionally, it is assumed in [4], [3] that the power

He Hao is with the Department of Electrical Engineering & Computer Sciences, University of California, Berkeley, 94720, USA. Corresponding author. Email: hehao@berkeley.edu.

Wei Chen is with the Department of Electronic & Computer Engineering, The Hong Kong University of Science and Technology, Clear Water Bay, Kowloon, Hong Kong, China. Email: eeweichen@ust.hk.

consumption of each load can be continuously modulated, while in this paper we consider the case each load has on-off switching behaviors, which is more appropriate for residential loads such as pool pump, air-conditioner, and electric vehicle charging. We also propose a new scheduling algorithm which generally could finish all tasks without surplus or deficit for exactly adequate generation profiles. However, in [4], [3], only heuristic algorithms such as EDF, LLF policies and/or receding horizon control algorithms are considered, and these algorithms generally are not feasible to finish all tasks given an exactly adequate generation profile.

The remainder of this paper is organized as follows. Section II presents the modeling method for deferrable loads. Characterization of generation adequacy appears in Section III. Scheduling of deferrable loads and numerical experiments are reported in Section IV. The paper ends with conclusions and future work in Section V.

II. Problem Formulation

In this paper, time is designated as discrete with discretization step δ over a time horizon $[0,(T+1)\delta]$. We use $\{t : t = 0, 1, \cdots, T\}$ to index the time interval $[t\delta,(t+1)\delta)$, and denote u_t as the power supplied to a collection of loads over that time interval. The sequence $u = (u_0, \cdots, u_t, \cdots, u_T)$ is regarded as a power profile.

Consider a collection of N distributed loads, in which each load has on-off switching behavior. The power consumption of the i^{th} load over time interval $[t\delta,(t+1)\delta)$ is given by

$$u_t^i = \begin{cases} m^i, & \text{ON state}, \\ 0, & \text{OFF state}, \end{cases}$$

where m^i is its rated power. Each load has an energy requirement $E^i \in [\underline{E}^i, \overline{E}^i]$ over a desired time step window $[a^i, d^i) \subseteq \{0, 1, \cdots, T\}$, where $\overline{E}^i \leq (d^i - a^i)m^i\delta$. Moreover, we assume $\underline{E}^i$ and $\overline{E}^i$ are integer multiples of $m^i\delta$. A load is called *deferrable* if $E^i < (d^i - a^i)m^i\delta$. In this paper, we assume all loads are deferrable.

We model each deferrable load as a *task*. A task is parameterized by its total energy requirement E^i, arrival time a^i, departure time d^i, and rated power m^i. It is denoted by the quadruple $\mathbb{T}^i = (E^i, a^i, d^i, m^i)$. A task can be serviced at time $t \in [a^i, d^i)$, and it is required that

$$\sum_{k=a^i}^{d^i-1} u_k^i \delta = E^i, \quad u_k^i \in \{0, m^i\}.$$

The energy state of task $\mathbb{T}^i$ at time t is defined as

$$s_t^i := s_{t-1}^i + u_{t-1}^i \delta = \sum_{k=a^i}^{t-1} u_k^i \delta.$$

A population of deferrable loads can be described as a collection of tasks

$$\mathbb{T} := \{\mathbb{T}^i = (E^i, a^i, d^i, m^i) : i = 1, \cdots, N\}.$$

The set of *active tasks* at time t is denoted as

$$\mathbb{A}_t := \{\mathbb{T}^i : a^i \leq t < d^i\}.$$

In each resource cluster, the cluster manager is responsible for forecasting renewable generation, purchasing energy in the forward energy market, and buying/selling reserves in the ancillary service markets, in order to procure a total generation profile to service a collection of defferrable loads. The total generation profile (which results from renewable generation, bulk power, and reserve procurement) to service active tasks is denoted by $g = (g_0, \cdots, g_t, \cdots, g_T)$, and its cumulative energy is defined as

$$G_t := G_{t-1}^i + g_{t-1}^i \delta = \sum_{k=0}^{t-1} g_k \delta.$$

We next characterize the set of all generation profiles, g's, that are able to satisfy the energy requirements of all deferrable loads without surplus or deficit.

III. Characterizing Generation Adequacy

The cluster manager allocates the available generation g_t at time t to active tasks using some scheduling policy. A scheduling policy σ is an algorithm that allocates the available generation g_t to active tasks:

$$\sigma(g_t) = (u_t^1, u_t^2, \cdots, u_t^N)$$

such that $u_t^i = 0$ if $i \notin \mathbb{A}_t$, $u_t^i \in \{0, m^i\}$ if $i \in \mathbb{A}_t$, and $\sum_i u_t^i = g_t$. We call a scheduling policy *causal* if it only uses the past and current information. Otherwise, it is called non-causal.

A. Definitions of Generation Adequacy

We next define set-theoretic representations of adequate generation profiles.

Definition 1: The generation profile g is called *exactly adequate* if there exists a (possibly non-causal) scheduling policy that completes all the tasks without surplus or deficit. The set of all exactly adequate generation profiles is described by

$$\mathbf{G} = \left\{ g \;\middle|\; \begin{array}{c} \exists\ \sigma,\ \sigma(g_t) = (u_t^1, u_t^2, \cdots, u_t^N),\ \forall\ t, \\ \sum_{t=0}^{T} g_t \delta = \sum_{i=1}^{N} E^i, \\ s_{a^i}^i = 0,\ s_{d^i}^i = E^i, \quad \forall\ i, \end{array} \right\}. \qquad \square$$

We propose a generalized battery model to succinctly model the generation adequacy. It serves as a convenient and portable model that facilitates the cluster manager to decide how to purchase and/or sell in the energy and ancillary service markets.

Definition 2: A *Generalized Battery Model* $\mathbb{B}$ is a set of power signals n that satisfy

$$\left\{ n \;\middle|\; \begin{array}{c} n_t^- \leq n_t \leq n_t^+, \\ x_{t+1} = x_t + n_t\delta,\ x_0 = 0\ \Rightarrow\ x_t^- \leq x_t \leq x_t^+ \end{array} \right\},$$

where $t = \{0, 1, \cdots, T\}$. The model is specified by non-negative parameters $\phi_t = (n_t^-, n_t^+, x_t^-, x_t^+)$, and it is written compactly as $\mathbb{B}(\phi)$. $\square$

We can regard n_t as the power draw of the battery, x_t as its state-of-charge, n_t^-, n_t^+ as its discharge/charge power limits, and x_t^-, x_t^+ as its energy capacity limits.

B. Characterization under General Scenario

We first consider a collection of deferrable loads with heterogeneous energy requirements, arrival and departure times, and rated powers.

Theorem 1: Consider a collection of deferrable loads with heterogenous energy requirement E, arrival time a, departure time d, and rated power m. Then the set of exactly adequate generations satisfies

$$\mathbf{G} \subseteq \mathbb{B}(\phi),$$

where the parameters $\phi_t = (n_t^-, n_t^+, x_t^-, x_t^+)$ of the battery model are given by

$$n_t^- = 0, \qquad n_t^+ = \sum_{i \in \mathbb{A}_t} m^i, \tag{1}$$

$$x_t^- = \sum_{i:d^i \le t} \underline{E}^i + \sum_{i \in \mathbb{A}_t} \max\{\underline{E}^i - (d^i - t)m^i\delta, 0\}, \tag{2}$$

$$x_t^+ = \sum_{i:d^i \le t} \overline{E}^i + \sum_{i \in \mathbb{A}_t} \min\{\overline{E}^i, (t - a^i)m^i\delta\}. \tag{3}$$

Moreover, the power draw of the battery, n_t, satisfies

$$\sum_{a^i \le k < d^i} \mathbb{I}(n_k \ge m^i) \ge \underline{E}^i/(m^i\delta), \quad \forall\, i, \tag{4}$$

$$\sum_{a^i \le k < d^i} \mathbb{I}(n_k = \sum_{j \in \mathbb{A}_t} m^j) \le \overline{E}^i/(m^i\delta), \quad \forall\, i, \tag{5}$$

where $\mathbb{I}(\cdot)$ is an indicator function. □

Proof: The proof borrows idea from [4]. See Appendix A for details. ■

C. Characterization for Tasks with Homogeneous Rated Power, Arrival and Departure times

In general, it is very challenging to give an analytic sufficient characterization on the generation adequacy for a collection of heterogenous tasks. However, if the deferrable loads have the same rated power, arrival and departure times, a sufficient and necessary characterization is possible. In particular, we will find a battery model $\mathbb{B}(\phi)$ such that for all $g \in \mathbb{G}$, we have $g \in \mathbb{B}(\phi)$, and vice versa.

Before we proceed further, it is necessary to give some preliminaries on majorization theory, which is the mathematical foundation for our characterization. For any vector $y \in \mathbb{R}^N$, let $y_{[1]} \ge \cdots \ge y_{[N]}$ denote the components of y in decreasing order. For any vectors $y, z \in \mathbb{R}^N$, if $\sum_{i=1}^k y_{[i]} \le \sum_{i=1}^k z_{[i]}$ for all $k \in \{1, \cdots, N\}$, we say y is weakly majorized by z, and we write $y \prec_w z$. Additionally, if $\sum_{i=1}^N y_{[i]} = \sum_{i=1}^N z_{[i]}$, we say y is majorized by z, and we write $y \prec z$. Given nonnegative integers $T+1 \ge a_1 \ge a_2 \ge \cdots \ge a_N$, and let A be the $N \times (T+1)$ matrix with i-th row δ_{a_i}, where δ_{a_i} is a $(T+1)$-dimensional vector with first a_i components equal to 1, and the rest components equal to 0. The conjugate vector of $(a_1, a_2, \cdots, a_N)$ is just the vector of the column sums of A, which has $(T+1)$ components. It is denoted by $(a_1, a_2, \cdots, a_N)^*$. See [10] for more details on majorization theory.

Theorem 2: Consider a collection of deferrable loads with homogeneous rated power m, arrival time a, and departure time d. Then the set of exactly adequate generations satisfies

$$\mathbf{G} = \mathbb{B}(\phi),$$

where the parameters ϕ_t of the battery model satisfies conditions (1)- (3) in Theorem 1. Moreover, the power draw of the battery $n = (n_0, \cdots, n_T)$ satisfies

$$(\underline{E}^1/(m\delta), \cdots, \underline{E}^N/(m\delta)) \prec_w (n_0/m, \cdots, n_T/m)^*,$$
$$(n_0/m, \cdots, n_T/m) \prec_w (\overline{E}^1/(m\delta), \cdots, \overline{E}^N/(m\delta))^*,$$

where y^* denotes the conjugate vector of y. Moreover, if $\underline{E}^i = E^i = \overline{E}^i$ for all i, the above two weak majorization conditions are reduced to

$$(n_0/m, \cdots, n_T/m) \prec (E^1/(m\delta), \cdots, E^N/(m\delta))^*. \quad \square$$

Proof: See Appendix B. ■

Remark 1: Note that in the case of homogeneous rated power, arrival and departure times, the adequacy of generation profiles is in fact translated into the existence of a $(0, 1)$-matrix with given or approximately given row and column sums. The existence of such matrices is widely studied in the literature, for instance, [11], [12]. In particular, the work in [13] deploys such matrices to study a duration-differentiated energy service.

Also note that majorization is an old mathematical tool which has been studied in mathematics by giants like Hardy, Littlewood, and Pólya in their masterpiece [14] and has been widely used in statistics in the past 100 years. Recently, its engineering applications are attracting considerable attention in many different fields, such as wireless communication [15], information theory [16], control theory [17], operations research [18], and smart grid [13], etc. The study in this paper provides another powerful application of majorization in smart grid. □

IV. Scheduling of Deferrable Loads

In the above section, we have characterized the set of exactly adequate generation profiles. However, given an exactly adequate generation profile, the question of how to allocate it to deferrable loads is not answered. In this section, we present three priority-based scheduling algorithms to service a collection of deferrable loads.

A. Load Scheduling Policies

One of the well-known *causal* scheduling policies is Earliest Deadline First (EDF) [19], [20]. It is one of the simplest policies which allocates available generation to the tasks based on their deadlines d^i's. The task with earlier departure time receives higher priority. At each time step t, available generation g_t is first assigned to the task T_j with the highest priority, where $j = \arg\min_{i \in \mathbb{A}_t} d^i$. Available generation in excess of the rated power for task T_j is allocated to the active task with the next higher priority.

Another well known *causal* scheduling policy is Least Laxity First (LLF) [21], [3], which allocates available generation to active tasks based a priority that is measured by *slack* time, which is defined as $\delta_t^i = d^i - (E^i - s_t^i)/m^i$. The task with smaller slack time is granted with higher priority. The LLF policy first assigns available generation to the active task T_j, where $j = \arg\min_{i \in \mathbb{A}_t} \delta_t^i$. Available generation g_t in excess of the rated power for task T_j is allocated to the active task with the next least slack time.

It was shown in [3] that given any exactly adequate generation profile, there generally do not exist a feasible *causal* scheduling algorithm that completes all tasks without surplus or deficit. Therefore, we propose in this paper a *non-causal* heuristic scheduling policy, which we call the *Earliest Deadline Priority* (EDP) policy. It first starts with the task with the earliest deadline, say T^j. For each time step t during its life span $[a^j, d^j)$, a priority that is measured by the difference between the current available power g_t and the maximum acceptable power for the collection $\overline{g}_t = \sum_{i \in \mathbb{A}_t} m^i$ is calculated. The smaller the difference is, the higher priority the time slot receives. We then assign power to task T^j at the first E^j time slots that have the highest priorities. A new generation profile is then obtained by subtracting the assigned powers from the previous generation profile, and the finished task is removed from the set of active tasks. This process is repeated with remaining tasks with later deadlines using the new generation profile. The heuristics behind the EDP algorithm is that if g_t is equal to the maximum acceptable power $\overline{g}_t$, then it must be allocated to all tasks at time t. Otherwise, there is some flexibility in assigning the power. Moreover, the smaller the difference is, the smaller flexibility it has, thus it receives a higher priority. Additionally, we aim to finish the task with the earliest deadline first. This motivates the EDP algorithm.

We next show by numerical experiments that for exactly adequate generation profiles satisfying conditions (1)-(5), the EDP policy generally can complete all tasks (that have homogeneous rated power) without surplus or deficit, while the EDF and LLF policies generally cannot meet this objective. An analytical study on feasibility of EDP policy for heterogenous tasks is a future work.

B. Numerical Experiments

In this section, we focus on a collection of residential pool pumps. According to the International Aquatic Foundation, there are about 8 million residential pools in the United States [22]. In particular, California and Florida account for 1.3 million and 1.1 million respectively [23], [24], [25]. On average, each pool pump consumes 1 kW power, and needs to run 8 hours to filter the water in the pool once in a day. If these resources are intelligently controlled, there is an enormous potential to provide fast ramping resources to absorb the variability of renewable energies. In fact, the control and communication infrastructure have been in place. For example, the SmartAC™ program in PG&E (Pacific Gas and Electric Company) and the OnCall® program in FPL (Florida Power and Light Company) gathered over one million customers, who cede their control of residential pool pumps, water heaters and air-conditioners to the aggregator to mange energy emergency situations[26], [27]. These programs can be exploited to provide various grid services with minimal investment.

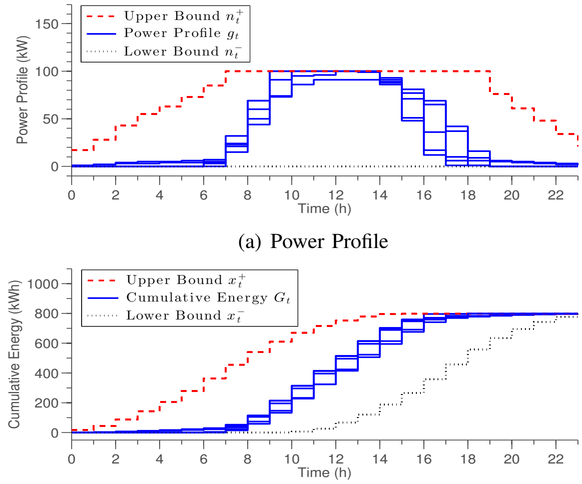

(a) Power Profile

(b) Cumulative Energy

Fig. 1. Generation profiles created by using CAISO solar generation data.

We assume the cluster manager manages resources with a timeline similar to the day-ahead energy market. Time is discrete with 1-hour increment over a 24-hour window, i.e. $t \in \{0, 1, \cdots, 23\}$. We consider a population of residential pool pumps, in which each pool pump needs to filter the water in the pool once in a day. When a pool pump is on, it has certain flow rate (gallon per minute) and consumes each amount of power. Given the volume capacity (gallons) of each pool, the pool pump needs to run at least h hours to circulate all the water once through the pool's filter. Therefore, each pool pump can be modeled as a task. We consider a collection of 100 pool pumps with homogeneous rated power, i.e. $m^i = 1$ kW for all $i \in \{1, 2, \cdots, N\}$. Each pool pump's energy requirement E^i, desired start time a^i, and desired end time d^i are assumed to have discrete uniform distributions as $E^i \sim U(7, 9), a^i \sim U(0, 7), d^i \sim U(19, 24)$.

We create the generation profiles based on the hourly solar PV generations in California Independent System Operator (CAISO). We first scale the magnitude of the hourly renewable generation profile in a day so that its total energy is equal to the total energy requirements of 100 pool pumps. We next adjust its power profiles by slightly moving the excessive peak powers (those larger than $n_t^+ = \sum_{i \in \mathbb{A}_t} m^i$) to its neighboring time slots so that it is within the upper power limits n_t^+. We randomly select a day's hourly solar PV generation profiles in CAISO in the year of 2013. Numerical experiments reveal that for exactly adequate generation profiles generated in this way, such as those shown in Fig. 1 (a), the EDP policy can finish all the tasks without surplus or deficit. However, the EDF and LLF policies could not finish all task. In particular, with the EDF policy, about 11% of the tasks are not finished, with an average of 24% of their energy requirements are not satisfied. Similarly, with the LLF policy, 7% of the tasks are not finished, with an average of 13% of their energy requirements are not satisfied.

We also create the generation profiles based on the hourly

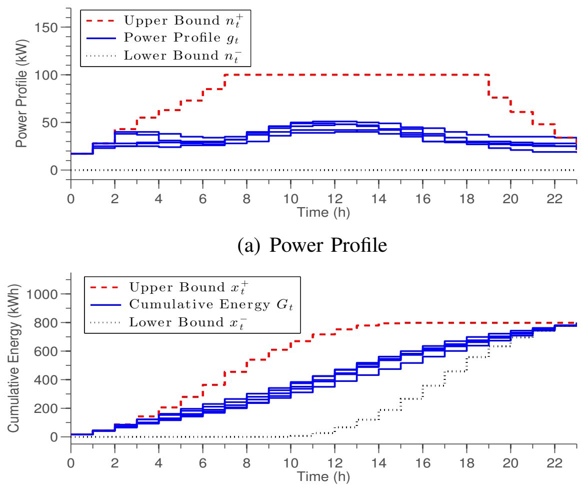

(a) Power Profile

(b) Cumulative Energy

Fig. 2. Generation profiles created by using CAISO total renewable generation data.

total renewable generations in CAISO. We create the generation profile the same way as that using the hourly solar generation data. We show that for exactly adequate generation profiles g generated in this way, such as those shown in Fig. 2, the EDP policy can finish all the tasks without surplus or deficit. However, the EDF and LLF policies could not finish all task without surplus or deficit. In particular, with the EDF policy, about 17% of the tasks are not finished, with an average of 24% of their energy requirements are not satisfied. Similarly, with the LLF policy, 20% of the tasks are not finished, with an average of 13% of their energy requirements are not satisfied.

In summary, we show by numerical experiments that for exactly adequate generation profiles constructed from CAISO solar or renewable generations, the EDP policy generally could finish all tasks without surplus or deficit, while the other two heuristic scheduling policies (EDF and LLF) fail. However, we should notice that the EDP policy is non-causal, which requires accurate prediction of generation profiles, while the EDF and LLF policies are causal, which only use the past and current information. However, the EDP policy will become useful when deferrable loads are scheduled in the day-ahead energy market, in which the generation profile is procured one day ahead of delivery time.

V. Conclusions and Future Work

We studied aggregation, characterization, and scheduling of a collection of deferrable loads for renewable integration. We proposed a generalized battery model to describe the generation adequacy, and provided necessary and/or sufficient characterizations on the set of exactly adequate generation profiles. We also proposed a non-causal scheduling algorithm. Numerical experiments showed that the EDP algorithm was generally able to allocate exactly adequate generation profiles to satisfy the energy requirements of all deferrable loads with homogenous rated power, whereas the EDF and LLF policies could not achieve this goal.

In the future, we plan to quantify the value of deferrable loads, in terms of the cost savings associated with the amount of generation reserves that are saved by deferrable loads. Additionally, studying the effect of prediction uncertainty of generation profiles on the scheduling performance of our proposed algorithm is a direction of future research. An analytic study on the feasibility of EDP policy for heterogenous tasks is also a future work.

Appendix

A. Proof of Theorem 1

If the available generation g is exactly adequate, then there exists a scheduling policy σ that completes all the tasks without surplus or deficit. First of all, the instantaneous power g_t cannot exceed the sum of the rated powers of all active tasks, i.e.

$$0 \le g_t \le \sum_{i \in \mathbb{A}_t} m^i.$$

Moreover, for each task i, the number of time steps in which the available generation g is greater than its rated power m^i during its life span $[a^i, d^i)$ must be no smaller than the minimum number of time steps to complete the task,

$$\sum_{a^i \le k < d^i} \mathbb{I}(g_k \ge m^i) \ge \underline{E}^i / m^i.$$

Similarly, for each task i, the number of time steps for the available generation g to be equal to its maximum possible power $\sum_{i \in \mathbb{A}_t} m^i$ during its life span $[a^i, d^i)$ must be no larger than the required maximum number of time steps to complete that task,

$$\sum_{a^i \le k \le d^i} \mathbb{I}(g_k = \sum_{i \in \mathbb{A}_t} m^i) \le \overline{E}^i / m^i.$$

Additionally, on the one hand, the cumulative energy G_t at time t must exceed the total energy requirements of all departed tasks. Moreover, any active task at time t with departure time $d^i > t$ can receive at most $(d^i - t)m^i$ units of energy on the time interval $[t, d^i)$. As a result, the cumulative generation at time t must supply at least $\max\{0, E^i - (d^i - t)m^i\}$ to any task active at time t. Hence

$$\begin{aligned} G_t &\ge \sum_{i: d^i \le t+1} E^i + \max\{0, E^i - (d^i - t)m^i\} \\ &\ge \sum_{i: d^i \le t+1} \overline{E}^i + \max\{0, \overline{E}^i - (d^i - t)m^i\}. \end{aligned}$$

On the other hand, the cumulative generation cannot exceed the total energy needs of all active and past tasks. Moreover, any task active at time t with arrival time $a^i \le t$ could have received at most $\min\{E^i, (t - a^i)m^i\}$ units of energy. Thus

$$\begin{aligned} G_t &\le \sum_{i: d^i \le t+1} E^i + \min\{E^i, (t - a^i)m^i\} \\ &\le \sum_{i: d^i \le t+1} \overline{E}^i + \min\{\overline{E}^i, (t - a^i)m^i\} \end{aligned}$$

Roughly speaking, the above lower bound on G_t is corresponding to the power profile scheduled just before departure, and the upper bound is corresponding to the power profile scheduled immediately at arrival.

If $g \in \mathbf{G}$, the $g \in \mathbb{B}(\phi)$, where the power draw and the state of charge of the battery are respectively given by $n_t = g_t$ and $x_t = G_t$. This completes the proof. ■

B. Proof of Theorem 2

First of all, we know from Theorem 1 that $\mathbf{G} \subseteq \mathbb{B}(\phi)$, where the battery model parameters ϕ_t satisfy conditions (1)- (3). We next show that if the power draw of the battery satisfies conditions in Theorem 2, it can also guarantee sufficiency, i.e. $\mathbb{B}(\phi) \subseteq \mathbf{G}$.

Without loss of generality, we assume $a^i = 0, d^i = T + 1, m^i = 1, \delta = 1$, and $\underline{E}^i, E^i, \overline{E}^i$ are positive integers for all i. In this case, the two weak majorization conditions in Theorem 2 can be rewritten as

$$(\underline{E}^1, \cdots, \underline{E}^N) \prec_w (n_0, \cdots, n_T)^*,$$
$$(n_0, \cdots, n_T) \prec_w (\overline{E}^1, \cdots, \overline{E}^N)^*.$$

We now convert the original problem as a $(0, 1)$-matrix filling problem. If a generation profile g is exactly adequate, then there exists a scheduling policy σ that could allocate it to all tasks without surplus or deficit. It is equivalent to that there exists a $(0, 1)$ matrix with N rows and $T + 1$ columns, whose i^{th} row sum is E^i (where $i = 1, 2, \cdots, N$), and t^{th} column sum is g_t (where $t = 0, 1, \cdots, T$). Note that the distribution of entries with element 1 at the t^{th} column corresponds to a scheduling policy at time t. If E^i is flexible, i.e. $E^i \in [\underline{E}^i, \overline{E}^i]$, then the problem is equivalent to the existence of a $(0, 1)$ matrix with bounded row sum $[\underline{E}^i, \overline{E}^i]$ and fixed column sum g_t.

In [12] and [10, Page 252], it is shown the existence of such matrix is equivalent to the following weak majorization conditions.

$$(\underline{E}^1, \cdots, \underline{E}^N) \prec_w (g_0, \cdots, g_T)^*,$$
$$(g_0, \cdots, g_T) \prec_w (\overline{E}^1, \cdots, \overline{E}^N)^*.$$

Moreover, if $\underline{E}^i = E^i = \overline{E}^i$ for all $i = 1, \cdots, N$, the above two weak conditions are equivalent to the following majorization condition

$$(g_0, \cdots, g_T) \prec (E^1, \cdots, E^N)^*.$$

Since we assume $\underline{E}^i, E^i$, and $\overline{E}^i$ are integer multiples of m^i for all i, therefore the weak majorization and majorization conditions given in Theorem 2 also hold for general a, d, m and δ. ■

REFERENCES

[1] Energy and Environmental Economics, Inc., "Investigating a higher renewables portfolio standard in california," 2014. [Online]. Available: https://ethree.com/documents/E3_Final_RPS_Report_2014_01_06_with_appendices.pdf

[2] D. Bakken, A. Bose, K. Chandy, P. Khargonekar, A. Kuh, S. Low, A. von Meier, K. Poolla, P. Varaiya, and F. Wu, "GRIP - Grids with intelligent periphery: Control architectures for Grid2050," in *IEEE International Conference on Smart Grid Communications*, 2011, pp. 7–12.

[3] A. Subramanian, M. Garcia, A. Dominguez-Garcia, D. Callaway, K. Poolla, and P. Varaiya, "Real-time scheduling of deferrable electric loads," in *American Control Conference*, 2012, pp. 3643–3650.

[4] A. Nayyar, J. Taylor, A. Subramanian, K. Poolla, and P. Varaiya, "Aggregate felxibility of a collection of loads," in *IEEE conference on Decision and Control*, 2013, pp. 5600–5607.

[5] J. Liu, S. Li, W. Zhang, J. L. Mathieu, and G. Rizzoni, "Planning and control of electric vehicles using dynamic energy capacity models," in *IEEE Conference on Decision and Control*, 2013, pp. 379–384.

[6] H. Hao, B. M. Sanandaji, K. Poolla, and T. L. Vincent, "Aggregate flexibility of thermostatically controlled loads," *IEEE Transactions on Power Systems*, vol. PP, no. 99, pp. 1–10, 2014. [Online]. Available: http://plaza.ufl.edu/hehao/publication.html

[7] S. Meyn, P. Barooah, A. Busic, and J. Ehren, "Ancillary service to the grid from deferrable loads: the case for intelligent pool pumps in florida," in *IEEE conference on Decision and Control*, 2013, pp. 6946–6953.

[8] H. Hao, Y. Lin, A. Kowli, P. Barooah, and S. Meyn, "Ancillary service to the grid through control of fans in commercial building HVAC systems," *IEEE Transactions on Smart Grid*, vol. 5, no. 4, pp. 2066–2074, 2014.

[9] A. Papavasiliou and S. S. Oren, "Supplying renewable energy to deferrable loads: Algorithms and economic analysis," in *IEEE Power and Energy Society General Meeting*, 2010, pp. 1–8.

[10] A. W. Marshall, I. Olkin, and B. C. Arnold, *Inequalities: Theory of Majorization and Its Applications*. Springer, 2010.

[11] H. J. Ryser, "Combinatorial properties of matrices of zeros and ones," *Canadian Journal of Mathematics*, vol. 9, pp. 371–377, 1957.

[12] D. Fulkerson, "A network flow feasibility theorem and combinatorial applications," *Can. J. Math*, vol. 11, pp. 440–451, 1959.

[13] A. Nayyar, M. N. Pincetic, K. Poolla, and P. Varaiya, "Duration-differentiated energy services with a continuum of loads," *accepted by 53rd IEEE Conference on Decision and Control*, 2014.

[14] G. H. Hardy, J. E. Littlewood, and G. Pólya, *Inequalities*. Cambridge University Press, 1952.

[15] D. P. Palomar and Y. Jiang, *MIMO Transceiver Design via Majorization Theory*. Now Publishers Inc., 2007.

[16] F. Cicalese and U. Vaccaro, "Supermodularity and subadditivity properties of the entropy on the majorization lattice," *IEEE Transactions on Information Theory*, vol. 48, no. 4, pp. 933–938, 2002.

[17] W. Chen, S. Wang, and L. Qiu, "When MIMO control meets MIMO communication: a majorization condition for networked stabilizability," *arXiv preprint arXiv:1408.3500*, 2014.

[18] C.-S. Chang and D. D. Yao, "Rearrangement, majorization and stochastic scheduling," *Mathematics of Operations Research*, vol. 18, no. 3, pp. 658–684, 1993.

[19] S. Baruah and J. Goossens, "Scheduling real-time tasks: Algorithms and complexity," *Handbook of scheduling: Algorithms, models, and performance analysis*, vol. 3, 2004.

[20] C. L. Liu and J. W. Layland, "Scheduling algorithms for multiprogramming in a hard-real-time environment," *Journal of the ACM (JACM)*, vol. 20, no. 1, pp. 46–61, 1973.

[21] M. L. Dertouzos and A. K. Mok, "Multiprocessor online scheduling of hard-real-time tasks," *IEEE Transactions on Software Engineering*, vol. 15, no. 12, pp. 1497–1506, 1989.

[22] "Media: Statistics: How many pools are there in the U.S.?" [Online]. Available: http://www.aquaticnet.com/media-statistics3.htm

[23] C. Palmgren, N. Stevens, M. Goldberg, R. Barnes, and K. Rothkin, "2009 california residential appliance saturation study, cec-200-2010-004-es," Tech. Rep., Octobor 2010. [Online]. Available: www.energy.ca.gov

[24] United States Census Bureau, U.S. Department of Commerce, "California quickfacts." [Online]. Available: http://quickfacts.census.gov/qfd/states/06000.html

[25] K. Hall and M. Lo, "Survey of residential swimming pools assessed by florida county property appraisers," October 2006. [Online]. Available: http://www.doh.state.fl.us/injury/PDF/ResidentialSwimmingPoolSurvey.pdf

[26] Pacific Gas and Electric Company, "Smart AC Program." [Online]. Available: http://www.pge.com/en/myhome/saveenergymoney/energysavingprograms/smartac/index.page

[27] Florida Power and Light Company, "On call savings program." [Online]. Available: http://www.fpl.com/residential/energy_saving/programs/oncall.shtml