2018 IEEE Conference on Decision and Control (CDC)
Miami Beach, FL, USA, Dec. 17-19, 2018

Automatic Power Exchange for Distributed Energy Resource Networks: Flexibility Scheduling and Pricing

Junjie Qin, Jonathan Mather, Jhi-Young Joo, Ram Rajagopal, Kameshwar Poolla and Pravin Varaiya

Abstract— **This paper proposes an Automatic Power Exchange (APEX) that enables monetization of underutilized distribution system energy resources. APEX features an open-gate forward market design to incorporate uncertainty from variable resources, and an explicit flexibility market that schedules flexible resources based on information submitted by users through a simple yet expressive order format. We study the non-convex non-preemptive scheduling problem in APEX, proposing polynomial time algorithms with finite and asymptotic performance guarantees. We then analyze the properties of marginal pricing, generalized to fit the APEX context with forward markets and distribution network constraints. We establish that it is revenue adequate but may lead to inadmissible prices for flexible orders. We then suggest a simple pricing mechanism that provably produces admissible prices for users and adequate revenue for APEX if implemented together with the proposed scheduling algorithms.**

I. INTRODUCTION

A. Background

Wholesale electricity markets have long enabled efficient trading of bulk energy and services at the transmission scale. But there are many significant resources and assets connected to the distribution network that have not been fully monetized [1]. Novel distribution system markets that match the local intermittent supply with flexible demand can potentially greatly increase the utilization of these assets.

However, designing such markets is challenging for a number of reasons. First, many distributed energy resources are *variable* resulting in intermittent and uncertain power generation that introduces both *quantity risk* (e.g., not enough supply to meet demand) and *price risk* (e.g., consumers may be charged highly volatile prices) into the market. Managing such risks require a sophisticated distribution system operator (DSO) solving stochastic dispatch programs, or *forward markets* in which the market participants can trade to hedge against uncertainty [2]–[4]. Second, although demand flexibility is ubiquitous, unlocking it usually requires upfront capital investments from users (e.g. for installing smart appliances and/or building energy management systems [5], [6]) that need to be justified by a clear expectation of (financial) benefits. Existing proposals around real time pricing could potentially provide such an expectation, but it may be blurred by difficulties in forecasting prices driven by both exogenous uncertainty (e.g. renewable generation) and endogenous uncertainty (e.g. other market participants' behaviors). Explicit *flexibility markets* that schedule flexible demand on behalf of the users could significantly reduce the burden of users and provide clear incentives for users to engage, reveal and trade their flexibility. Finally, as the market outcomes induce physical power flow on the distribution network, physical network constraints need to be managed to ensure the reliability of the distribution network.

B. Contributions

In this paper, we propose a scalable market platform, referred to as *Automatic Power Exchange* (APEX), that enables monetization of these underutilized distribution system assets. Our APEX platform allows distribution system participants to trade energy and services. It incorporates variable distributed energy resources by an *open-gate* forward market design. That is, for each delivery period, users can submit orders in anytime inside of a trading time window, which if possible will be cleared as submitted. Effectively, this introduces a continuum of forward markets, where users can hedge against uncertainty through adjustment orders based on most updated information. APEX also arranges an explicit flexibility market. Distribution system participants can submit the *availability* information of their flexible loads, and APEX will schedule these flexible loads on behalf of the participants. APEX will respect distribution network constraints on the flow of electricity either by directly managing the distribution network or by following a coordinated trading protocol [7], [8] operated by a third-party distribution system operator.

This paper contributes to the literature in the following ways. (i) We propose a novel design for a distribution system market that addresses uncertainty from Distributed Energy Resources (DERs) using an open-gate forward market design and solicits demand flexibility by efficient in-market flexible demand scheduling, while managing distribution network constraints. (ii) We study the non-convex problem for scheduling *non-preemptive* flexible loads and propose provably efficient algorithms to ensure the scalability of the APEX platform. (iii) We analyze the properties of a natural marginal pricing mechanism in the APEX context, establishing that it is revenue adequate but may lead to inadmissible prices for flexible orders. We then suggest a simple alternative that is guaranteed to produce admissible prices for users and adequate revenue for APEX when used together with the proposed scheduling algorithms.

J. Qin and J.Y. Joo are with the Grid Integration Group at Lawrence Berkeley National Laboratory {jqin, jjoo}@lbl.gov

J. Qin, J. Mather, K. Poolla and P. Varaiya are with the Department of Electrical Engineering and Computer Science, University of California at Berkeley {qinj, jonathan.mather, poolla, varaiya}@berkeley.edu

R. Rajagopal is with the Department of Civil and Environmental Engineering, Stanford University ramr@stanford.edu

978-1-5386-1395-5/18/$31.00 ©2018 IEEE

C. Literature

Forward markets have been implemented in many wholesale electricity markets. It is known that forward markets help manage uncertainty and incorporate generation technologies with different lead times [2]. When a sequence of forward markets are available, risk sensitive consumers (suppliers) may limit their risks by procuring from (offering to) multiple forward markets [9]. Although many of these studies have focused on the wholesale market, empirical studies demonstrate that for smaller consumers having the option of participating in forward markets helps them to hedge their bill volatility [3]. Open-gate forward markets, compared to fixed-time forward markets such as day-ahead and hour-ahead markets, are not common for electricity. Yet, they are widely implemented in financial industries [10].

The utilization of flexible energy resources in distribution networks has been studied in a number of papers. Existing studies usually exploit the flexibility in restrictive settings where only one attribute of the flexible resources is allowed to vary. Such treatments lead to interesting control and pricing problems for electricity services that are differentiated according to that particular attribute [11]–[13]. It is usually assumed in these papers that these services are organized and provided by aggregators instead of a flexibility market that matches flexible resources with other resources. Furthermore, in our flexibility market, the flexible orders are allowed to simultaneously have many different attributes, thus bridging a gap between prior studies and practical implementations.

As a whole package, APEX provides a novel design for distribution system electricity market with significant DER penetration. Existing alternative proposals for distribution system markets can be roughly categorized into *centralized* and *transactive*. Centralized designs seek to modify or augment existing utility companies' rate structures to align DERs' power consumption/production with wholesale price signals. Notable examples include real time pricing (RTP) and its variants (cf. [14]–[17] and references therein). The benefits of centralized design include tight management of distribution network through utility companies and the fact that they are relatively easy to implement given today's institutional structures of retail electricity markets. Such designs, nevertheless, are inflexible as it is difficult to incorporate differentiated electricity services. In contrast, transactive designs rely on bilateral or multilateral transactions (or contracts) among individual participants (cf. [18], [19] and references therein). As the terms and conditions in the contracts can be tailored according to individual needs, it is very easy to incorporate various flexible resources in transactive market designs. However, these designs represent a big structural departure from today's utility-managed distribution markets, require coordination to ensure reliability of the distribution network, [7] and may impose significant *search costs* on the participants. Compared to these two classes of distinct designs, we view APEX as a middle ground where flexible resources are incorporated by an explicit flexibility market and an expressive alphabet of standard commodities, and distribution network constraints are tightly managed (possibly by a coordinated trading protocol). Participants' search costs are also largely reduced in APEX.

Many of our technical results extend the growing literature on scheduling non-preemptive deferrable loads [20]–[22]. Most of these prior studies focus on an aggregator setting while we consider scheduling in a two-sided market. Furthermore, the fluid relaxation based scheduling algorithm that we propose is novel and well-suited for the region where there is a large number of flexible loads. As the scheduling problem is non-convex, the problem of pricing these flexible loads with distribution network constraints is challenging and understudied in the literature. Our results on understanding the properties of marginal pricing with sub-optimal scheduling algorithms may pave the way for future development on this topic.

D. Organization

The rest of the paper is organized as follows. Section II describes the APEX market platform and states the order matching problem in APEX. Section III proposes efficient algorithms for the combinatorial optimization of scheduling non-preemptive flexible orders and establishes their performance guarantees. The associated pricing problem for APEX order matching is then considered in Section IV. Section V concludes the paper.

II. APEX Platform

In a nutshell, APEX receives orders (Section II-C) from users (Section II-B), forms and maintains an orderbook (Section II-D), and solves an order matching problem (Section II-E) that fulfills standing orders in the orderbook by matching supply with demand respecting distribution network constraints (Section II-A). The schematic of the trading process in APEX is demonstrated in Figure 1.

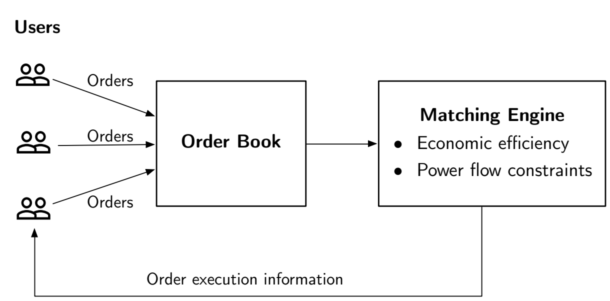

Fig. 1. Trading process in APEX.

Trading in APEX happens within the following temporal structure. Time is slotted into time intervals of Δt length (e.g., Δt can be 5 minutes). Power delivery in each of these time intervals is traded. We thus work with a discrete time model, using $t \in \mathbb{Z}$ to denote each time period. At any time instance t, users can submit orders regarding power delivery in any future time intervals belonging to a *trading time window* $\mathcal{T}_t$ that includes T time intervals. The trading time window may start with the next time interval $t+1$ and ends 24 hours after the next time interval, i.e., in this case T is 24 hours/Δt. As a result, orders regarding the power delivery in any time interval t' can be submitted at any t such that $t' \in \mathcal{T}_t$. This implements an open-gate forward market. Figure 2 gives an example of the open-gate forward markets and trading windows for two delivery intervals.

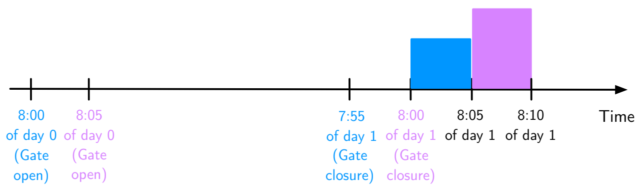

Fig. 2. Example of APEX trading windows

We proceed to introduce individual components of the APEX platform.

A. Distribution network model

Consider a distribution network specified by a graph with buses $\mathcal{N} = \{1, \ldots, N\}$ and lines $\mathcal{E}$. As distribution systems are usually radial, the graph has a tree structure so that $|\mathcal{E}| = N-1$. For the bulk of the paper, we focus on real power flow and adopt the *Simplified DistFlow* model [23] that allows us to write the nodal voltage $\mathbf{v} \in \mathbb{R}^N$ and the branch flow $\boldsymbol{\ell} \in \mathbb{R}^{N-1}$ as linear functions of the nodal power injection $\mathbf{p} \in \mathbb{R}^N$:

$$\mathbf{v} = \widehat{\mathbf{v}} + R\mathbf{p}, \tag{1a}$$

$$\boldsymbol{\ell} = H\mathbf{p}, \tag{1b}$$

where $\widehat{\mathbf{v}} \in \mathbb{R}^{N+1}$ is the vector of reference voltages that include the substation voltage and the voltage contribution from reactive power injections, $R \in \mathbb{R}^{N\times N}$ is a matrix that depends on the topology of the distribution network and the line resistances (see [24] for more details), $H \in \mathbb{R}^{(N-1)\times N}$ is the shift-factor matrix for the network.

Operational constraints for the distribution network usually include bounds on the nodal voltages and line flows (to avoid transformer overloading). Under (1), these constraints specify a polyhedral real power injection region for a distribution network

$$\mathcal{P}_{\mathrm{D}} := \{\mathbf{p} \in \mathbb{R}^N : \mathbf{1}^\top \mathbf{p} = 0,\ \underline{\mathbf{v}} \le \widehat{\mathbf{v}} + R\mathbf{p} \le \overline{\mathbf{v}},\ \underline{\boldsymbol{\ell}} \le H\mathbf{p} \le \overline{\boldsymbol{\ell}}\},$$

where $\mathbf{1} \in \mathbb{R}^N$ is the all-one vector, $\underline{\mathbf{v}}$ and $\overline{\mathbf{v}}$ are the voltage bounds, and $\underline{\boldsymbol{\ell}}$ and $\overline{\boldsymbol{\ell}}$ are the line flow bounds. As is common the case in practice, we assume $\underline{\mathbf{v}} \le \widehat{\mathbf{v}} \le \overline{\mathbf{v}}$ and $\underline{\boldsymbol{\ell}} \le \mathbf{0} \le \overline{\boldsymbol{\ell}}$, and therefore $\mathbf{0} \in \mathcal{P}_{\mathrm{D}}$.

B. User model

We denote the set of users by $\mathcal{I} := \{1, \ldots, I\}$. Let the bus that user $i \in \mathcal{I}$ resides be denoted by n_i and the set of users located at bus n be denoted by $\mathcal{I}_n$. Each user may model an individual home, a commercial building, or an aggregation of many buildings coordinated by an aggregator or as a micro-grid. In this paper, APEX is agnostic to the level of aggregation inside of each user.

C. Order formats

At any instance in time, referred to as $t = 0$, a participant located at bus n (i.e. $i \in \mathcal{I}_n$) can submit buy or sell orders for a trading window of time periods $t \in \mathcal{T} := \mathcal{T}_t = \{1, \ldots, T\}$. Considering typical supply-side and demand-side characteristics, we allow buy and sell orders in the formats specified as follows.

Definition 1 (Simple sell order): A simple sell order from participant $i \in \mathcal{I}_n$ is a tuple

$$s = (n,\ t,\ \overline{q},\ \underline{\pi}),$$

where n is the bus index, $t \in \mathcal{T}$ is time of electricity delivery, $\overline{q} \in \mathbb{R}_+$ is the maximum amount of electricity to be sold and $\underline{\pi} \in \mathbb{R}_+$ is the minimum acceptable price of electricity for the sell order.

Symmetrically, we have simple buy order defined.

Definition 2 (Simple buy order): A simple buy order from participant $i \in \mathcal{I}_n$ is a tuple

$$b = (n,\ t,\ \overline{q},\ \overline{\pi}),$$

where n is the bus index, $t \in \mathcal{T}$ is time of electricity delivery, $\overline{q} \in \mathbb{R}_+$ is the maximum amount of electricity to be bought and $\overline{\pi} \in \mathbb{R}_+$ is the maximum acceptable price of electricity for the buy order.

Simple orders may not be expressive enough to incorporate certain flexible loads such as *non-preemptive* shiftable loads. Many such loads consume pre-defined load shapes but are indifferent to the time at which the loads are served as long as they are served in a certain time window. This motivates us to incorporate flexible buy order as follows.

Definition 3 (Flexible buy order): A flexible buy order from participant $i \in \mathcal{I}_n$ is a tuple

$$f = (n,\ t^{\mathrm{ES}},\ t^{\mathrm{LC}},\ \tau^{\mathrm{D}},\ \widehat{\mathbf{q}},\ \overline{\pi}),$$

where n is the bus index, $t^{\mathrm{ES}} \in \mathcal{T}$, $t^{\mathrm{LC}} \in \mathcal{T}$ and $\tau^{\mathrm{D}} \in \mathcal{T}$ denote the earliest starting time, latest completion time and duration of the flexible load, respectively, $\widehat{\mathbf{q}} \in \mathbb{R}_+^{\tau^{\mathrm{D}}}$ is the load shape to be consumed, and $\overline{\pi} \in \mathbb{R}_+$ is the maximum acceptable price of electricity[1] for the buy order.

Figure 3 depicts the parameters used to define a flexible buy order.

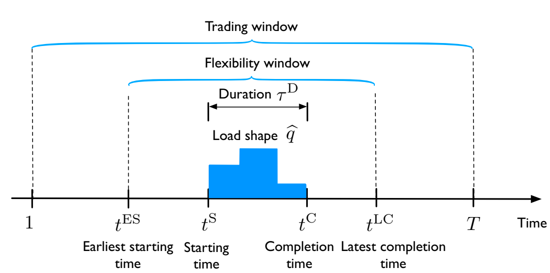

Fig. 3. Parameters of a flexible buy order

[1]The maximum amount of payment associated with the order is $\overline{\pi}\mathbf{1}^\top\widehat{\mathbf{q}}$.

D. Orderbook

In practice, buy and sell orders arrive continuously in time. Whenever a new order arrives, APEX runs an efficient matching algorithm with all the standing orders and the newly arrived order, which determines the fulfillment of a subset of these orders. All unfulfilled orders, which cannot be matched due to (i) lack of supply for a buy order or lack of demand for a sell order, (ii) lack of a mutually acceptable price, and (iii) network constraints, remain standing and are recorded into an *orderbook*.

Definition 4 (Orderbook): The orderbook at time t is defined to be the triple $(\mathcal{B}, \mathcal{S}, \mathcal{F})$, where $\mathcal{B}$ is the collection of standing simple buy orders, $\mathcal{S}$ is the collection of standing simple sell orders, and $\mathcal{F}$ is the collection of standing flexible buy orders.

In the order matching process, some simple orders might be partially fulfilled. These orders remain in orderbook with the desirable quantities ($\overline{q}_b$ or $\overline{q}_s$) updated by subtracting the fulfilled amounts.

E. Order matching problem

To fulfill the orders in the orderbook $(\mathcal{B}, \mathcal{S}, \mathcal{F})$, the order matching process aim to determine an *admissible schedule* and an *admissible price* for each (partially) fulfilled order in the orderbook.

Definition 5 (Admissible schedule): A power production or consumption schedule $\mathbf{q} \in \mathbb{R}^T$ over the trading window $\mathcal{T}$ is deemed admissible, if the following conditions hold.

- For simple sell order $s = (n,\ t,\ \overline{q},\ \underline{\pi})$:
$$\mathbf{q} = \mathbf{q}_s \in \mathcal{Q}_s := \{q\mathbf{1}_t \in \mathbb{R}^T : 0 \le q \le \overline{q}\},$$
where $\mathbf{1}_t \in \mathbb{R}^T$ is the elementary vector with a 1 at t-th element and 0's elsewhere.
- For simple buy order $b = (n,\ t,\ \overline{q},\ \overline{\pi})$:
$$\mathbf{q} = \mathbf{q}_b \in \mathcal{Q}_b := \{q\mathbf{1}_t \in \mathbb{R}^T : 0 \le q \le \overline{q}\}.$$
- For flexible buy order $f = (n,\ t^{\mathrm{ES}},\ t^{\mathrm{LC}},\ \tau^{\mathrm{D}},\ \widehat{\mathbf{q}},\ \overline{\pi})$:
$$\mathbf{q} = \mathbf{q}_f \in \mathcal{Q}_f,$$
where $\mathcal{Q}_f$ is defined as the set of power profiles $\mathbf{q} \in \mathbb{R}^T$ such that there exists a starting time t^{S} so $(\mathbf{q}, t^{\mathrm{S}})$ satisfies
$$t^{\mathrm{S}} \in \{t^{\mathrm{ES}}, \dots, t^{\mathrm{LC}} - \tau^{\mathrm{D}} + 1\}, \tag{2a}$$
$$q(t) = \begin{cases} \widehat{q}(t - t^{\mathrm{S}} + 1), & \text{if } t^{\mathrm{S}} \le t < t^{\mathrm{S}} + \tau^{\mathrm{D}}, \\ 0, & \text{otherwise.} \end{cases} \tag{2b}$$
If the order is not to be scheduled, we denote $t^{\mathrm{S}} = 0$ and so $\mathbf{0} \in \mathcal{Q}_f$ by definition.

Definition 6 (Admissible price): For a (partially) scheduled order (i.e., $\mathbf{q} \neq 0$), a clearing price $\pi \in \mathbb{R}$ is deemed admissible, if the following conditions hold.

- For simple sell order $s = (n,\ t,\ \overline{q},\ \underline{\pi})$: $\pi \ge \underline{\pi}$.
- For simple buy order $b = (n,\ t,\ \overline{q},\ \overline{\pi})$: $\pi \le \overline{\pi}$.
- For flexible buy order $f = (n,\ t^{\mathrm{ES}},\ t^{\mathrm{LC}},\ \tau^{\mathrm{D}},\ \widehat{\mathbf{q}},\ \overline{\pi})$: $\pi \le \overline{\pi}$.

In the order matching problem, we try to identify admissible fulfillment of all orders $(\{\mathbf{q}_s\}, \{\mathbf{q}_b\}, \{\mathbf{q}_f\}) := (\{\mathbf{q}_s\}_{s\in\mathcal{S}}, \{\mathbf{q}_b\}_{b\in\mathcal{B}}, \{\mathbf{q}_f\}_{f\in\mathcal{F}})$ in a way that maximizes certain criteria designed by the operator of APEX, denoted by $U(\{\mathbf{q}_s\}, \{\mathbf{q}_b\}, \{\mathbf{q}_f\})$, while respecting the distribution network constraints. This can be written as the following optimization problem

$$\max_{\{\mathbf{q}_s\},\{\mathbf{q}_b\},\{\mathbf{q}_f\}} \quad U\big(\{\mathbf{q}_s\}, \{\mathbf{q}_b\}, \{\mathbf{q}_f\}\big) \tag{3a}$$
$$\text{s.t.} \quad \mathbf{q}_s \in \mathcal{Q}_s, \quad s \in \mathcal{S}, \tag{3b}$$
$$\mathbf{q}_b \in \mathcal{Q}_b, \quad b \in \mathcal{B}, \tag{3c}$$
$$\mathbf{q}_f \in \mathcal{Q}_f, \quad f \in \mathcal{F}, \tag{3d}$$
$$\mathbf{p}_n = \sum_{s\in\mathcal{S}_n} \mathbf{q}_s - \sum_{b\in\mathcal{B}_n} \mathbf{q}_b - \sum_{f\in\mathcal{F}_n} \mathbf{q}_f, n \in \mathcal{N}, \tag{3e}$$
$$\mathbf{p}(t) \in \mathcal{P}_{\mathrm{D}}, \quad t \in \mathcal{T}. \tag{3f}$$

where (3e) is the local power balance equation at each node n, with $\mathcal{S}_n$, $\mathcal{B}_n$ and $\mathcal{F}_n$ denoting the set of simple sell orders, simple buy order and flexible buy order submitted by users at bus n, respectively.

APEX may optimize different criteria in the order matching problem depending on its real-world implementation (e.g., whether it is implemented by a for-profit platform company or by a regulated utility company). Here we list two possible objective functions to optimize.

- Total surplus:
$$\begin{aligned} &U(\{\mathbf{q}_s\}, \{\mathbf{q}_b\}, \{\mathbf{q}_f\}) \\ &= \sum_{b\in\mathcal{B}} \overline{\pi}_b \mathbf{1}^\top \mathbf{q}_b + \sum_{f\in\mathcal{F}} \overline{\pi}_f \mathbf{1}^\top \mathbf{q}_f - \sum_{s\in\mathcal{S}} \underline{\pi}_s \mathbf{1}^\top \mathbf{q}_s. \end{aligned} \tag{4}$$
- Total volume:
$$U(\{\mathbf{q}_s\}, \{\mathbf{q}_b\}, \{\mathbf{q}_f\}) = \sum_{b\in\mathcal{B}} \mathbf{1}^\top \mathbf{q}_b + \sum_{f\in\mathcal{F}} \mathbf{1}^\top \mathbf{q}_f. \tag{5}$$

We note that with criteria (4) or (5), the order matching problem has a linear objective function. Meanwhile, constraints (3b), (3c), (3e) and (3f) are all linear inequality or equality constraints. However, (3) is challenging to solve due to non-convex constraint (3d). In fact, the combinatorial nature is evident if we return to the characterization of $\mathcal{Q}_f$ using the starting times (2). Next section introduces two algorithms solve (3) approximately.

While solving (3) gives an admissible schedule for each order in the orderbook, it does not provide admissible prices. Thus the second part of the order matching problem is to identify an admissible price for each fulfilled order. Given the non-convex nature of (3) and the fact that we can only obtain approximate solutions of (3) in practice, the pricing problem for APEX is challenging. In particular, the natural application of the *marginal pricing* idea to this context requires a re-examination because its nice properties established for convex settings may no longer hold (see Section IV).

III. Scheduling Algorithms

Fixing the schedule of flexible buy orders, the order matching problem is a linear program thus efficiently solvable. We therefore denote

$$
\begin{aligned}
&J(\{\mathbf{q}_f\}) := \\
&\max_{\{\mathbf{q}_s\},\{\mathbf{q}_b\}} \quad U(\{\mathbf{q}_s\},\{\mathbf{q}_b\},\{\mathbf{q}_f\}) && \text{(6a)}\\
&\quad\text{s.t.} \quad \mathbf{q}_s \in \mathcal{Q}_s, \quad s \in \mathcal{S}, && \text{(6b)}\\
&\qquad\quad\ \ \mathbf{q}_b \in \mathcal{Q}_b, \quad b \in \mathcal{B}, && \text{(6c)}\\
&\qquad\quad\ \ \mathbf{p}_n = \sum_{s\in\mathcal{S}_n} \mathbf{q}_s - \sum_{b\in\mathcal{B}_n} \mathbf{q}_b - \sum_{f\in\mathcal{F}_n} \mathbf{q}_f, n \in \mathcal{N}, && \text{(6d)}\\
&\qquad\quad\ \ \mathbf{p}(t) \in \mathcal{P}_{\mathrm{D}}, \quad t \in \mathcal{T}. && \text{(6e)}
\end{aligned}
$$

and focus on the optimization for scheduling flexible orders:

$$
\begin{aligned}
&\max_{\{\mathbf{q}_f\}} \quad J(\{\mathbf{q}_f\}) && \text{(7a)}\\
&\quad\text{s.t.} \quad \mathbf{q}_f \in \mathcal{Q}_f, \quad f \in \mathcal{F}. && \text{(7b)}
\end{aligned}
$$

We proceed to describe a way to solve this problem based on a greedy heuristic.

A. Greedy scheduling

We start by reformulating (7) to a set function maximization. Let $\mathbf{t}^{\mathrm{S}} \in \mathcal{T}^{|\mathcal{F}|}$ be the vector of starting times of all flexible buy orders which uniquely determines the schedule of all flexible buy orders $\{\mathbf{q}_f\}$. Denote the value of (7) with some fixed $\mathbf{t}^{\mathrm{S}}$ by $V(\mathbf{t}^{\mathrm{S}})$, i.e.,

$$
V(\mathbf{t}^{\mathrm{S}}) = \begin{cases} J(\{\mathbf{q}_f(t_f^{\mathrm{S}})\}), & \text{if } \mathbf{q}_f(t_f^{\mathrm{S}}) \in \mathcal{Q}_f, f \in \mathcal{F}, \\ -\infty, & \text{otherwise,} \end{cases}
$$

where $\mathbf{q}_f(t_f^{\mathrm{S}})$ is the power consumption profile induced by starting time t_f^{S}. Consider the pairs of flexible buy orders and their starting times in the set

$$
\Omega = \{(f, t_f^{\mathrm{S}}) : f \in \mathcal{F}, t_f^{\mathrm{S}} \in \mathcal{T} \cup \{0\}\}. \tag{8}
$$

Notice that any feasible scheduling can be represented by a subset of Ω; conversely, subsets of Ω that select no more than one starting time for each flexible buy order f can represent all feasible scheduling decisions. Define set-to-matrix mapping $\mathbb{I} : 2^{\Omega} \mapsto \mathbb{R}^{|\mathcal{F}|\times T}$

$$
[\mathbb{I}(X)]_{f,t} = \begin{cases} 1, & \text{if } (f,t) \in X, \\ 0, & \text{otherwise,} \end{cases} \tag{9}
$$

and normalized objective function $g(X) = V(\mathbb{I}(X)\boldsymbol{\delta}) - V(\mathbb{I}(\emptyset)\boldsymbol{\delta})$, where $\boldsymbol{\delta} \in \mathbb{R}^{T\times 1}$ is such that $\delta_t = t$, and the matrix vector product $\mathbb{I}(X)\boldsymbol{\delta}$ converts a subset $X \subset \Omega$ into the corresponding starting time vector t^{S}. With these definitions, problem (7) is equivalent to the following *subset selection problem*:

$$
\begin{aligned}
&\max_{X\subset\Omega} \quad g(X), && \text{(10a)}\\
&\quad\text{s.t.} \quad \sum_{t=1}^{T} [\mathbb{I}(X)]_{f,t} \le 1, \quad f \in \mathcal{F}, && \text{(10b)}
\end{aligned}
$$

where the constraint ensures that X selects at most one starting time for each flexible buy order. The problem in general is NP hard as the number of subsets is $2^{|\Omega|}$.

The greedy approach for solving (11) amounts to scheduling flexible orders one-by-one according to the incremental benefit that scheduling a new order brings as measured by function $g(X)$. Algorithm 1 lists the steps for greedy scheduling. After this algorithm terminates, admissible schedules for flexible buy orders are obtained. We can then obtain admissible schedules for simple orders by solving (6) with the resulting $\{\mathbf{q}_f\}$ from the greedy algorithm.

Algorithm 1: Greedy scheduling

$\mathcal{F}^{\mathrm{S}} = \emptyset$, $\mathcal{F}^{\mathrm{TBS}} = \mathcal{F}$;
$X \leftarrow \emptyset$;
while $\mathcal{F}^{\mathrm{TBS}} \neq \emptyset$ **do**
  $\mathcal{C} \leftarrow \{(f, t^{\mathrm{S}}) : f \in \mathcal{F}^{\mathrm{TBS}}, t^{\mathrm{S}} \in \{t_f^{\mathrm{ES}}, \dots, t_f^{\mathrm{LC}}+1-\tau_f^{\mathrm{D}}\}\}$;
  $(\widehat{f}, \widehat{t^{\mathrm{S}}}) \leftarrow \mathrm{argmax}_{(f,t^{\mathrm{S}})\in\mathcal{C}}\, g(X \cup \{(f, t^{\mathrm{S}})\}) - g(X)$;
  if $g(X \cup \{(\widehat{f}, \widehat{t^{\mathrm{S}}})\}) > g(X)$ **then**
    $X \leftarrow X \cup \{(\widehat{f}, \widehat{t^{\mathrm{S}}})\}$;
    $\mathcal{F}^{\mathrm{S}} \leftarrow \mathcal{F}^{\mathrm{S}} \cup \{\widehat{f}\}$, $\mathcal{F}^{\mathrm{TBS}} \leftarrow \mathcal{F}^{\mathrm{TBS}} \backslash \{\widehat{f}\}$
  else
    break;
  end
end

B. Fluid relaxation

Although the complexity of greedy scheduling is polynomial in the number of flexible orders, it becomes relatively slow when there are a large number of orders because it needs to loop over the remaining orders and their feasible starting times in each step. In this section, we consider an alternative scheduling algorithm based on relaxing the non-convex constraint $\mathbf{q}_f \in \mathcal{Q}_f$. In particular, it first relaxes the requirement that each load shape needs to follow the exact load shape $\widehat{\mathbf{q}}_f$ submitted by the user, solves a convex optimization to determine the schedule $\widetilde{\mathbf{q}}_f$, and then "projects" the schedule $\widetilde{\mathbf{q}}_f$ to a feasible schedule $\Pi_{\mathcal{Q}_f}(\widetilde{\mathbf{q}}_f) \in \mathcal{Q}_f$. As the key step in this algorithm is to remove the load shape requirement by allowing any profile to be scheduled in the time window $\{t_f^{\mathrm{ES}}, \dots, t_f^{\mathrm{LC}}\}$, we refer to this algorithm as *fluid relaxation*. Similar ideas have been used to develop approximation algorithms for job shop scheduling problems (cf. [25]).

In fluid relaxation, we replace the constraint $\mathbf{q}_f \in \mathcal{Q}_f$ by $\mathbf{q}_f \in \widetilde{\mathcal{Q}}_f$, with $\widetilde{\mathcal{Q}}_f$ defined as the set of power consumption profiles $\mathbf{q}_f \in \mathbb{R}^T$ satisfying the following constraints

$$
\begin{aligned}
&\mathbf{1}^\top \mathbf{q}_f \le \mathbf{1}^\top \widehat{\mathbf{q}}_f, && \text{(11a)}\\
&\mathrm{TV}(\mathbf{q}_f) \le \mathrm{TV}(\widehat{\mathbf{q}}_f), && \text{(11b)}\\
&q_f(t) \ge 0, \quad t \in \{t_f^{\mathrm{ES}}, \dots, t_f^{\mathrm{LC}}\}, && \text{(11c)}\\
&q_f(t) = 0, \quad t \notin \{t_f^{\mathrm{ES}}, \dots, t_f^{\mathrm{LC}}\}, && \text{(11d)}
\end{aligned}
$$

where the total variation of a vector $\mathbf{x} \in \mathbb{R}^d$ is defined as

$$
\mathrm{TV}(\mathbf{x}) = \sum_{t=0}^{d+1} |x(d+1) - x(t)|, \quad x(0) = x(d+1) := 0.
$$

In the definition of $\widetilde{\mathcal{Q}}_f$, constraint (12a) requires that the total energy of the scheduled consumption profile is no larger than the total energy of the submitted consumption profile; constraints (12c) and (12d) impose the non-negativity requirement for the consumption profile and restrict the profile can only have positive consumption in time periods specified by the time window submitted by the user; constraint (12b) controls the flexibility of the scheduled profile, by limiting the total variation of the scheduled profile with that of the submitted profile. It is easy to see that $\mathcal{Q}_f \subset \widetilde{\mathcal{Q}}_f$.

As $\widetilde{\mathcal{Q}}_f$ is a convex polytope for each f, the resulting convex relaxation for scheduling flexible orders is

$$\max_{\{\mathbf{q}_f\}} \quad J(\{\mathbf{q}_f\}) \tag{12a}$$

$$\text{s.t.} \quad \mathbf{q}_f \in \widetilde{\mathcal{Q}}_f, \quad f \in \mathcal{F}, \tag{12b}$$

whose solution is denoted by $\{\widetilde{\mathbf{q}}_f\}$.

The solution of the convex relaxation may be infeasible with respect to the original constraint $\mathbf{q}_f \in \mathcal{Q}_f$. One possibility is that for some f, constraint (12a) may not hold with equality at the solution $\widetilde{\mathbf{q}}_f$. In this case, we simply round down such that $\mathbf{q}_f(t) = 0$ for all t. If constraint (12a) holds with equality at the solution and so the energy requirement of the flexible load is satisfied, we identify a feasible starting time t_f^{S} for each $f \in \mathcal{F}$ by finding the time window with τ_f^{D} periods that contains the maximum total power consumption according to $\widetilde{\mathbf{q}}_f$, i.e.,

$$t_f^{\mathrm{S}} = \underset{t_f^{\mathrm{S}} \in \{t_f^{\mathrm{ES}}, \ldots, t_f^{\mathrm{LC}} - \tau_f^{\mathrm{D}} + 1\}}{\operatorname{argmax}} \sum_{t = t_f^{\mathrm{S}}}^{t_f^{\mathrm{S}} + \tau_f^{\mathrm{D}} - 1} \widetilde{q}_f(t). \tag{13}$$

The "projected" power consumption schedule is thus the power consumption profile induced by this starting time, denoted by $\mathbf{q}_f(t_f^{\mathrm{S}})$:

$$\Pi_{\mathcal{Q}_f}(\widetilde{\mathbf{q}}_f) := \begin{cases} \mathbf{q}_f(t_f^{\mathrm{S}}), & \text{if } \mathbf{1}^\top \widetilde{\mathbf{q}}_f = \mathbf{1}^\top \widehat{\mathbf{q}}_f, \\ \mathbf{0}, & \text{otherwise.} \end{cases} \tag{14}$$

with t_f^{S} defined in (14).

C. Performance guarantees

In this section, we provide a stylized analysis of the proposed scheduling algorithms for solving the non-convex (combinatorial) optimization (7) with the following two assumptions.

A1 The network can be represented as a single-bus network.
A2 The maximization criteria is the total surplus (4).

Assumption **A**1 allows us to focus on the combinatorial nature of the problem rather than the network constraints. We stress, however, both of the proposed algorithms result in admissible schedules that respect distribution network constraints for general networks. Analyzing the performance of approximation algorithms with network constraints represents a major challenge and is left for future work. Assumption **A**2 is introduced without loss of generality. In fact, the total volume objective function (5) can be thought of as a special case of the total surplus objective, with prices suitably modified.

We start with the performance guarantee for the greedy algorithm[2]. As is often the case, the greedy algorithm is $(1-1/e)$-optimal if the underlying problem is *submodular*. For our problem of scheduling flexible orders, by extending the analysis in [21], we can show submodularity indeed holds and therefore we have the following performance guarantee for the greedy algorithm. Proofs are omitted due to the page limit.

Lemma 1 (Greedy performance): Under Assumptions **A**1 and **A**2, the greedy scheduling algorithm is $(1-1/e)$-optimal:

$$\frac{J(\{\mathbf{q}_f^{\mathrm{g}}\})}{J(\{\mathbf{q}_f^{\star}\})} \geq 1 - \frac{1}{e}, \tag{15}$$

where $\{\mathbf{q}_f^{\mathrm{g}}\}$ is the greedy schedule and $\{\mathbf{q}_f^{\star}\}$ is the optimal schedule.

One important feature of the performance bound (16) is that it is independent of the problem instance. It is thus referred to as an *a priori* bound as it can be stated before seeing the actual problem data.

For the fluid heuristic, we can derive general *a posteriori* bounds based on the following observation:

Proposition 1: Let $\{\widetilde{\mathbf{q}}_f\}$ and $\{\Pi_{\mathcal{Q}_f}(\widetilde{\mathbf{q}}_f)\}$ be the solutions of fluid relaxation (13) and its projected solution, respectively. Then

$$J(\{\Pi_{\mathcal{Q}_f}(\widetilde{\mathbf{q}}_f)\}) \leq J(\{\mathbf{q}_f^{\star}\}) \leq J(\{\widetilde{\mathbf{q}}_f\}). \tag{16}$$

This result bounds the unknown quantity $J(\{\mathbf{q}_f^{\star}\})$ by quantities that are computable in the fluid relaxation steps. This bound is general in that it holds without assuming **A**1 and **A**2. As such, given any problem instance and having computed $J(\{\Pi_{\mathcal{Q}_f}(\widetilde{\mathbf{q}}_f)\})$ and $J(\{\widetilde{\mathbf{q}}_f\})$, we can gauge the (sub-)optimality of the solution by comparing these two quantities. If they are close, we can assert *a posteriori* that the fluid algorithm has produced a $J(\{\Pi_{\mathcal{Q}_f}(\widetilde{\mathbf{q}}_f)\})/J(\{\widetilde{\mathbf{q}}_f\})$ optimal solution.

We can in fact show that the fluid heuristic is asymptotically optimal, with the following additional assumption:

A3 For all $f \in \mathcal{F}$, $\widehat{q}_f(t) \equiv \bar{q}_f$ does not change with t, $t = 1, \ldots, \tau_f^{\mathrm{D}}$. Furthermore, $\bar{q}^{\min} \leq \bar{q}_f \leq \bar{q}^{\max}$ for all $f \in \mathcal{F}$, where $0 < \bar{q}^{\min} < \bar{q}^{\max} < \infty$.
A4 The supply is sufficient as for each period $t \in \mathcal{T}$ there is a supply order with unbounded $\bar{q}$ and a large acceptable price $\underline{\pi}_{\mathrm{u}} \geq \max_{s \in \mathcal{S}} \underline{\pi}_s$. Furthermore, among $|\mathcal{F}|$ flexible orders, there are at least $\alpha|\mathcal{F}|$ flexible orders such that $\overline{\pi}_f \geq \underline{\pi}_{\mathrm{u}}$, where $\alpha \in (0,1)$.

Assumption **A**3 replaces possibly time-varying load shapes by rectangles. This certainly limits the practicality of our next result. Removing it is possible as any time-varying load shapes can be approximated, with any desired

[2] A slight modification (i.e., using a specific initialization of $\mathcal{F}^{\mathrm{S}}$ instead of $\mathcal{F}^{\mathrm{S}} = \emptyset$) of the greedy algorithm presented in (1) may be needed due to the knapsack constraint (11b). The modified greedy algorithm is still polynomial time but is much slower due to the time-consuming initialization step. We stated the simple greedy algorithm in Algorithm 1 because we have been observing similar performances with and without the modification. See [26] and [27] for more discussions.

level of accuracy, by rectangles. Assumption **A**4 focuses our analysis on the case where there is enough supply for each future delivery interval. This can be the case when a utility company (or load serving entity) participates in the platform and offers to sell sufficient amount of energy at a relatively high price $\underline{\pi}_{\mathrm{u}}$. Under these assumptions, we can establish the (weak) asymptotic optimality of the proposed fluid heuristic, in the following sense.

Theorem 1 (Asymptotic optimality of fluid heuristic): Under Assumptions **A**1-**A**4, consider a set of flexible orders with increasing size $|\mathcal{F}| \to \infty$. For each $\mathcal{F}$, there exists an optimal schedule $\{\widetilde{\mathbf{q}}_f\}$ for the fluid relaxation (13) such that

$$\frac{J(\{\Pi_{\mathcal{Q}_f}(\widetilde{\mathbf{q}}_f)\})}{J(\{\mathbf{q}_f^\star\})} \to 1, \quad \text{as} \quad |\mathcal{F}| \to \infty, \tag{17}$$

where $\{\mathbf{q}_f^\star\}$ is an optimal schedule for the original problem (7).

The key intuition behind this theorem is that when the number of flexible orders becomes large, individual load shapes no longer play a significant role in the quality of the solution because (with a single-bus network) it is the aggregate load shape of all the flexible loads that matters in the optimization (7).

IV. PRICING

Algorithms presented in Section III only provides admissible schedules for orders in the orderbook, with clearing prices to be determined. Denote the collection of prices for orders by $\{\pi_s\}, \{\pi_b\}, \{\pi_f\}$. Basic requirements for these prices include (i) admissible, as defined in Definition 6, and (ii) *revenue adequate* so that the *merchandising surplus* of the platform, denoted by MS, is non-negative:

$$\mathrm{MS} = \sum_{b\in\mathcal{B}} \pi_b \mathbf{1}^\top \mathbf{q}_b + \sum_{f\in\mathcal{F}} \pi_f \mathbf{1}^\top \mathbf{q}_f - \sum_{s\in\mathcal{S}} \pi_s \mathbf{1}^\top \mathbf{q}_s \geq 0. \tag{18}$$

Without flexible orders, it can be shown that both requirements above can be met with a generalization of the simple idea of *marginal pricing*. In the rest of this section, we first state this generalization and then exam its properties when used with the scheduling algorithms proposed in Section III.

A. Marginal pricing

Given any schedule of the flexible orders $\{\mathbf{q}_f\}$, we consider the optimization for scheduling the remaining simple orders (6). Denote the optimal dual variable associated with constraint (6d) of the linear program by $\boldsymbol{\lambda}_n \in \mathbb{R}^T$, $n \in \mathcal{N}$. This is a collection of $N \times T$ prices, one for each *(bus, future deliver time)* pair. Thus these prices may be referred to as *temporal and locational marginal prices* [28], which are functions of the schedules of flexible orders in our setting.

In particular, for orders with non-zero cleared quantities, we define *marginal pricing rule* as

$$\pi_s = \lambda_n(t), \quad s = (n,\ t,\ \overline{q},\ \underline{\pi}), \tag{19a}$$

$$\pi_b = \lambda_n(t), \quad b = (n,\ t,\ \overline{q},\ \overline{\pi}), \tag{19b}$$

$$\pi_f = \frac{\boldsymbol{\lambda}_n^\top \mathbf{q}_f}{\mathbf{1}^\top \mathbf{q}_f}, \quad f = (n,\ t^{\mathrm{ES}},\ t^{\mathrm{LC}},\ \tau^{\mathrm{D}},\ \widehat{\mathbf{q}},\ \overline{\pi}). \tag{19c}$$

Thus under the marginal pricing rule, simple orders are paid or charged the locational marginal price for the future delivery time interval. As flexible orders usually span multiple delivery time intervals, they are charged the average locational marginal prices weighted by the amount of power they consume in different time intervals.

B. Properties of marginal pricing

We analyze properties of the marginal pricing rule based on Assumption **A**2, otherwise the dual variables of optimization (6) may not have a clear economic meaning.

Using the KKT condition of the linear program (6) and strong duality, we can establish the following property of the marginal pricing rule, given any schedules for the flexible orders $\{\mathbf{q}_f\}$:

Lemma 2 (Properties of marginal pricing): Under Assumption **A**2, the marginal pricing rule is revenue adequate and leads to admissible prices for simple buy orders and simple sell orders.

As the schedules of flexible orders $\{\mathbf{q}_f\}$ are not decision variables of the optimization (6), little can be said about whether the marginal prices will be admissible for flexible orders without considering the actual algorithms used to determine these schedules. Considering the algorithms proposed in Section III, we have the following negative results established using counter examples:

Lemma 3 (Inadmissibility for flexible orders): Under Assumption **A**2, the marginal pricing rule with greedy or fluid schedule is not guaranteed to produce admissible prices for flexible buy orders.

The observation above stems from the fundamental difficulties in non-convex pricing problems and the fact that the proposed algorithms use more than the marginal information to determine the schedules of flexible orders. To be precise, in each step of the greedy algorithm, it determines whether to schedule a flexible order based on the *total benefit* about which it brings measured by the change of function value $J(\{\mathbf{q}_f\})$ with and without the newly scheduled order. Charging it with the marginal cost, which corresponds to the cost of producing the last $\epsilon > 0$ amount of power, may result in a price higher than $\overline{\pi}_f$. For the fluid algorithm, the projection step does not use price information and may result in inadmissible prices.

While marginal pricing has many desirable properties as studied in the transmission market literature (cf. [29] and references therein), modifying it to extend some of these properties to the non-convex setting here require much additional work. Instead, in the next subsection, we provide a simple mechanism that will ensure price admissibility and budget adequacy.

C. Pay-as-bid mechanism

In the *pay-as-bid* mechanism, we simply pay or charge users based on the prices they submit with their order

$$\pi_s = \underline{\pi}_s, \quad \pi_b = \overline{\pi}_b, \quad \pi_f = \overline{\pi}_f, \tag{20}$$

for all (partially) fulfilled orders. By definition, this pricing mechanism produces admissible prices for all orders. We can also show that revenue adequacy is achieved if we use greedy scheduling algorithm with the total surplus objective:

Lemma 4: Under Assumption **A**2, the pay-as-bid mechanism is revenue adequate with greedy scheduling.

It is an easy consequence of Theorem 1 that if we use the fluid scheduling algorithm we are guaranteed revenue adequacy in an asymptotic sense.

Corollary 1: Under the same assumptions of Theorem 1 and with the pay-as-bid mechanism, consider a set of flexible orders with increasing size $|\mathcal{F}| \to \infty$. For each $\mathcal{F}$, there exists an optimal schedule $\{\widetilde{\mathbf{q}}_f\}$ for the fluid relaxation (13) such that the merchandizing surplus induced by the resulting schedule $\{\Pi_{\mathcal{Q}}(\widetilde{\mathbf{q}}_f)\}$ satisfies $\lim_{|\mathcal{F}|\to\infty} \mathrm{MS} > 0$.

In summary, both proposed scheduling algorithms result in admissible and revenue-adequate prices if the pay-as-bid mechanism is used.

V. Conclusion

In this paper, we propose APEX, a market platform that enables monetization of underutilized distribution system assets. It features an open-gate forward market design and an explicit flexibility market. The forward markets help to incorporate variable distributed energy resources and reduce risks of market participants. In the flexibility market, resources submit their flexibility information with a simple yet expressive order format and APEX schedule these resources on behalf of the users efficiently. All functionalities of APEX are executed while ensuring the reliability of the distribution network, either by directly managing the distribution system constraints, or by communicating with a minDSO through a coordinated trading protocol.

For the proposed market platform, we study the non-convex problem of scheduling non-preemptive flexible resources and propose polynomial time algorithms that have finite or asymptotic performance guarantees. We then analyze the properties of marginal pricing together with the proposed algorithms and suggest a simple alternative that leads to admissible prices for all users and guarantees adequate revenue for the platform.

References

[1] North American Electric Reliability Corporation, "Distributed energy resources: Connection modeling and reliability considerations," Tech. Rep., 2017.

[2] P. P. Varaiya, F. F. Wu, and J. W. Bialek, "Smart operation of smart grid: Risk-limiting dispatch," *Proceedings of the IEEE*, vol. 99, no. 1, pp. 40–57, Jan 2011.

[3] S. Borenstein, "Customer risk from real-time retail electricity pricing: Bill volatility and hedgability," National Bureau of Economic Research, Tech. Rep., 2006.

[4] K. Alshehri, S. Bose, and T. Başar, "Cash-settled options for wholesale electricity markets," *IFAC-PapersOnLine*, vol. 50, no. 1, pp. 13 605–13 611, 2017.

[5] A. Radovanović *et al.*, "Powernet for distributed energy resource networks," in *Power and Energy Society General Meeting (PESGM), 2016*. IEEE, 2016, pp. 1–5.

[6] A. I. Dounis and C. Caraiscos, "Advanced control systems engineering for energy and comfort management in a building environmenta review," *Renewable and Sustainable Energy Reviews*, vol. 13, no. 6-7, pp. 1246–1261, 2009.

[7] F. Wu and P. Varaiya, "Coordinated multilateral trades for electric power networks: theory and implementation," *International Journal of Electrical Power & Energy Systems*, vol. 21, no. 2, pp. 75 – 102, 1999.

[8] J. Qin, R. Rajagopal, and P. P. Varaiya, "Flexible market for smart grid: Coordinated trading of contingent contracts," *IEEE Transactions on Control of Network Systems*, vol. PP, no. 99, pp. 1–1, 2017.

[9] J. Qin and R. Rajagopal, "Price of uncertainty in multistage stochastic power dispatch," in *53rd IEEE Conference on Decision and Control*, Dec 2014, pp. 4065–4070.

[10] Grinblatt and Titman, *Financial markets and corporate strategy*. McGraw Hill, 2004.

[11] C.-W. Tan and P. Varaiya, "Interruptible electric power service contracts," *Journal of Economic Dynamics and Control*, vol. 17, no. 3, pp. 495–517, 1993.

[12] E. Bitar and S. Low, "Deadline differentiated pricing of deferrable electric power service," in *51st IEEE Conference on Decision and Control (CDC)*, Dec 2012, pp. 4991–4997.

[13] A. Nayyar, M. Negrete-Pincetic, K. Poolla, and P. Varaiya, "Duration-differentiated energy services with a continuum of loads," *IEEE Transactions on Control of Network Systems*, vol. 3, no. 2, pp. 182–191, 2016.

[14] S. Borenstein, "Wealth transfers among large customers from implementing real-time retail electricity pricing," *The Energy Journal*, pp. 131–149, 2007.

[15] ——, "The long-run efficiency of real-time electricity pricing," *The Energy Journal*, pp. 93–116, 2005.

[16] A. W. Berger and F. C. Schweppe, "Real time pricing to assist in load frequency control," *IEEE Transactions on Power Systems*, vol. 4, no. 3, pp. 920–926, 1989.

[17] P. L. Joskow and C. D. Wolfram, "Dynamic pricing of electricity," *American Economic Review*, vol. 102, no. 3, pp. 381–85, May 2012.

[18] D. J. Hammerstrom, R. Ambrosio, T. A. Carlon *et al.*, "Pacific Northwest GridWise Testbed Demonstration Projects; Part I. Olympic Peninsula Project," PNNL, Tech. Rep., 2008.

[19] E. G. Cazalet, "Automated transactive energy (TEMIX)," in *Grid-Interop Forum*, 2011.

[20] A. Subramanian, M. J. Garcia, D. S. Callaway, K. Poolla, and P. Varaiya, "Real-time scheduling of distributed resources," *IEEE Transactions on Smart Grid*, vol. 4, no. 4, pp. 2122–2130, 2013.

[21] G. O'Brien and R. Rajagopal, "Scheduling non-preemptive deferrable loads," *IEEE Transactions on Power Systems*, vol. 31, no. 2, pp. 835–845, March 2016.

[22] A. Gupta, R. Jain, and R. Rajagopal, "Scheduling, pricing, and efficiency of non-preemptive flexible loads under direct load control," in *2015 53rd Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, Sept 2015, pp. 1008–1015.

[23] M. E. Baran and F. F. Wu, "Network reconfiguration in distribution systems for loss reduction and load balancing," *IEEE Transactions on Power Delivery*, vol. 4, no. 2, pp. 1401–1407, Apr 1989.

[24] M. Farivar, L. Chen, and S. Low, "Equilibrium and dynamics of local voltage control in distribution systems," in *Decision and Control (CDC), 2013 IEEE 52nd Annual Conference on*. IEEE, 2013, pp. 4329–4334.

[25] D. Bertsimas, D. Gamarnik, and J. Sethuraman, "From fluid relaxations to practical algorithms for high-multiplicity job-shop scheduling: The holding cost objective," *Operations Research*, vol. 51, no. 5, pp. 798–813, 2003.

[26] J. Qin, I. Yang, and R. Rajagopal, "Submodularity of energy storage placement in power networks," in *2016 IEEE 55th Conference on Decision and Control (CDC)*, Dec 2016, pp. 686–693.

[27] M. Sviridenko, "A note on maximizing a submodular set function subject to a knapsack constraint," *Operations Research Letters*, vol. 32, no. 1, pp. 41 – 43, 2004.

[28] W. Tang, J. Qin, R. Jain, and R. Rajagopal, "Pricing sequential forward power contracts," in *Smart Grid Communications (SmartGridComm), 2015 IEEE International Conference on*. IEEE, 2015, pp. 563–568.

[29] H.-P. Chao and S. Peck, "A market mechanism for electric power transmission," *Journal of Regulatory Economics*, vol. 10, no. 1, pp. 25–59, 1996. [Online]. Available: http://dx.doi.org/10.1007/BF00133357