2019 American Control Conference (ACC)
Philadelphia, PA, USA, July 10-12, 2019

Data-Driven Capacity Bidding for Frequency Regulation

Sen Li*, Junjie Qin†, Akhil Shetty†, Kameshwar Poolla†*, and Pravin Varaiya†

Abstract— **Frequency regulation is crucial for balancing the supply and demand of modern electricity grids. To provide regulation services, it is important to understand the capability of flexible resources to track regulation signals. This paper studies the problem of submitting capacity bids to a forward regulation market based on historical regulation signals. We consider an aggregator who manages a group of flexible resources with linear dynamic constraints. He seeks to find the optimal capacity bid, so that real-time regulation signals can be followed with an *arbitrary* guaranteed probability. We formulate this problem as a chance-constrained program with unknown regulation signal distributions. A sampling and discarding algorithm is proposed. It provably provides near-optimal solutions at a guaranteed probability of success without knowing the distribution of the regulation signals. This result holds for resources with *arbitrary* linear dynamics and allows *arbitrary* intra-hour data correlations. We validate the proposed algorithm with real data via numerical simulations. Two cases are studied: (1) CAISO market, where providers separately submit capacity estimates for regulation up and regulation down signals, (2) PJM market, where regulation up and down capacities are the same. Simulation results show that the proposed algorithm provides near-optimal capacity estimates for both cases.**

I. Introduction

Increasing penetration of renewable energy poses significant challenges for balancing supply and demand in modern power grids. Among various ancillary services, the most valuable balancing operation is frequency regulation [1], where service providers adjust their power consumption to counteract the second-to-second variation of power imbalances. Frequency regulation is procured in a forward market. Each service provider submits a capacity in a day-ahead or hour-ahead market, and receives regulation signals within this capacity range in real-time.

How does a regulation service provider decide the capacity size in the forward regulation market? The optimal capacity size depends on the cost of operating the flexible resources, their dynamic constraints, and the statistics of regulation signals. The challenges of capacity bidding arise from several aspects. First, the service provider faces a trade-off among operational costs, regulation revenue and market risks. The cost models and risk attitudes are distinct for different market participants, and market rules change very often [2]. Second, the resource constraints are diverse. Generators have ramping constraints, batteries have capacity constraints, and aggregated thermostatically controlled loads have time-varying constraints that depend on external temperature and occupancy level. Capturing all aspects of these dynamic constraints is hard. Third, the statistics of regulation signals are unknown. The regulation signals have very strong correlation in a short time window (a couple of seconds), and very weak correlation in a longer horizon (one hour). Since resources have dynamic constraints, intra-hour correlation of regulation signals is important. However, this is challenging to capture due to the large dimension of signals and the unknown stochasticity that drives the system-wide imbalances.

This research is supported by the National Science Foundation under grants EAGER-1549945 and CPS-1646612, and by the National Research Foundation of Singapore under a grant to the Berkeley Alliance for Research in Singapore.

*The Department of Mechanical Engineering, University of California, Berkeley, USA. Email: lisen1990@berkeley.edu.

†The Department of Electrical Engineering and Computer Science, University of California, Berkeley, USA.

A. Related Work

One way to address the capacity bidding problem is via robust optimization [3]. This approach pursues the optimal decision for the worst-case scenarios in the future. They typically assume the randomness is bounded in predetermined intervals, and look for the capacity bids that accommodate the worst case scenario within this bound [4], [5], [6], [7]. However, since the worst scenario rarely happens in practice, the robust approach often leads to over-conservative solutions. This is particularly undesirable for resources with high operational costs [8].

Another strand of work provides more aggressive capacity estimates by allowing the violation of regulation contract. In [9] and [10], the authors formulated a chance-constrained program for the PEV aggregator to submit capacity bids in a forward market. A scenario approach was adopted in [11], where finite samples are drawn to approximate the chance constraint. This work mainly focuses on hourly dynamic constraints, neglecting the intra-hour correlations of regulation signals. In [12], the authors formulated a Markov decision problem for electric vehicles to optimize the capacity bids. The integrated regulation signals are modeled as a Markov random process. A similar model was proposed in [13], where the regulation signals are represented by hourly aggregate parameters. However, these models critically depend on the fact that only instantaneous and cumulative regulation signals matter, and this is largely due to the special structure of battery dynamics. For more general cases (e.g., batteries with depreciation factor, generators, etc), the hourly aggregate parameters can not fully capture the statistics of regulation signals for capacity bidding. As an alternative, [8] proposed an optimal bidding policy that maximizes the

978-1-5386-7926-5/$31.00 ©2019 AACC

market profit of batteries without modeling the regulation signals. The result of this paper still depends on the structure of battery models, which is hard to generalize.

B. Our Contribution

This paper studies the problem of submitting capacity bids to a forward regulation market based on historical regulation data. We consider an aggregator who manages a group of flexible resources with dynamic constraints. He seeks to find the maximum regulation capacity, so that real-time regulation signals can be followed with an *arbitrary* guaranteed probability. The major contributions of this paper are summarized as follows:

- We formulate the problem as a chance-constrained program. An efficient sampling and discarding algorithm is proposed. It provably provides near-optimal solutions that address the trade-off between regulation payments and market risks at a guaranteed probability of success. The proposed method is completely model-free: we do not need any model for the regulation signal. The result holds for *arbitrary* resources with linearly constrained dynamics and allows *arbitrary* intra-hour correlations of regulation signals.
- We address some practical concerns for implementing the sampling and discarding algorithm. We show that enough data is available to run the proposed algorithm. We also propose an efficient algorithm that improves the optimality loss by accurately approximating optimal constraint removal.
- We validate the proposed algorithm with real data via numerical simulations. Two cases are studied: (1) CAISO market, where providers separately submit capacity estimates for regulation up and regulation down signals, (2) PJM market, where regulation up and down capacities are the same. Simulation results show that the proposed algorithm provides near-optimal capacity estimates for both cases.

The remainder of this paper proceeds as follows. The capacity problem is formulated in Section II, followed by the main results presented in Section III and Section IV. Section V presents the numerical studies to validate the proposed approach. Last, concluding remarks and future directions are given in Section VI.

II. Capacity Bidding Problem

Consider an aggregator who manages a group of flexible resources to provide frequency regulation service. The resources are diverse and have dynamic constraints, such as energy storage, generators, aggregation of responsive loads [14], etc. At time t_0, the aggregator submits a capacity bid to the regulation market. During the delivery window $[0, T]$, the system operator sends back regulation signals within the capacity range. After receiving the regulation signal, the aggregator dispatches its resources to collectively track the regulation signal. Regulation markets are typically operated on an hourly basis: each delivery window is one hour, and the capacity is constant throughout this hour. As capacity bidding for consecutive hours is not strongly coupled, we focus on a single delivery window (hour) throughout the paper. The timeline of the problem is illustrated in Figure 1.

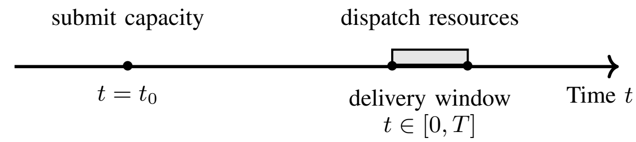

Fig. 1. Timeline of the capacity bididng problem.

A. Modeling Signals and Resources

Divide the delivery window $[0, T]$ into T contiguous periods. At each period t, the system operator sends out a regulation signal $e^t \in [-1, 1]$ to all market participants. Upon receiving the signal, each provider multiplies e^t by the respective capacity[1], and is obligated to track this power trajectory. Let $e = (e^1, \dots e^T)$ denote the signal trajectory for the hour. Assume that $e \in \Delta$, with $\Delta \subset \mathbb{R}^T$, and let $\mathbb{P}$ be a probability measure on Δ. We assume measure $\mathbb{P}$ is *unknown*.

Consider a group of N flexible resources with linear constraints. These constraints define a set of feasible energy outputs throughout the horizon. Let s_i^t be the net energy output of resource i with respect to a reference output at time t. Note hat s_i^t can be either positive or negative. For flexible resources that only consume energy (e.g., demand response) or produce energy (e.g., generator), s_i^t is the difference between the energy output and the nominal reference point. Let $s_i = (s_i^1, \dots, s_i^T)$ be the energy output trajectory in the horizon T. We assume that s_i takes values in the set S_i, which satisfies the following assumption:

Assumption 1: S_i is a bounded and convex polytope in $\mathbb{R}^T$, and $0 \in S_i$ for all $i = 1, \dots, N$.

The condition $0 \in S_i$ trivially holds if the resources simply follow the reference power trajectory. This polytope model captures a rich variety of flexible resources. It includes any resource with linear dynamic constraints. Meanwhile, it can also capture resources with time-varying dynamic constraints. For instance, the flexibility of commercial buildings is different from hour to hour due to change in occupancy level and outside temperature. This is significantly different from a vast number of works, which focus on batteries [5], [8], [15].

[1]In CAISO market, the regulation signal is multiplied by regulation up capacity if $e^t > 0$, and regulation down capacity if $e^t < 0$

B. Problem Statement

Regulation markets do not require all participants to exactly follow the regulation signals. For instance, PJM market maintains a performance score for all market participants [16]. Each participant is qualified as long as its service performance is scored above 75%. At the end of the day, market participants are compensated by a pay-for-performance scheme: they receive higher payments for better performance. This provides the opportunity to bid capacities aggressively. In this paper, we seek the optimal capacity bidding strategy that maximizes overall revenue of the service provider. The solution shall address the trade-offs between regulation payments and the risk of being disqualified.

To formulate this problem, we define $\alpha \in \mathbb{R}^d$ as the capacity bid of the service provider (aggregator). Note that α is not necessarily a scalar. For instance, in CAISO market, each participant submits separate capacity bids for regulation up signals and down signals. This leads to a vector capacity bid, i.e., $d = 2$. In contrast, the PJM market requires resources to be symmetric: the regulation up capacity equals the regulation down capacity. In this case, the capacity bid is a scalar, i.e., $d = 1$. For notation convenience, we define $L : \mathbb{R}^d \times \mathbb{R}^T \to \mathbb{R}^T$ as a mapping $L(\alpha, e)$ from α and e to the power trajectory to be followed by the aggregator. In the PJM case, $L(\alpha, e) = \alpha e$, whereas in the CAISO case, the positive signals are multiplied by $\alpha(1)$ and the negative signals are multiplied by $\alpha(2)$, i.e., $L(\alpha, e) = \alpha(1)e^+ + \alpha(2)e^-$.

Assume that $\pi \in \mathbb{R}^d$ is the per-unit capacity price for regulation services. We cast the capacity bidding problem as follows:

$$\max_{\alpha} \quad \pi \cdot \alpha \tag{1}$$

$$\begin{cases} \mathbb{P}\{L(\alpha, e) \in S_1 \oplus S_2 \oplus \cdots \oplus S_N\} \geq 1 - \epsilon, & \text{(2a)} \\ 0 \leq \alpha \leq \bar{\alpha}, & \text{(2b)} \end{cases}$$

where $\pi \cdot \alpha$ denotes the inner product of π and α, $\bar{\alpha}$ is sufficiently large upper bound on α to guarantee existence of a finite solution, and $\oplus$ represents the Minkowski sum of sets, i.e., $A \oplus B = \{a + b | a \in A, b \in B\}$. This is a chance-constrained program. The chance constraint (2a) dictates that under capacity α, the aggregator can cover the regulation signals with probability of greater than $1 - \epsilon$.

Remark 1: The capacity bidding problem (1) is parametrized by ϵ. When ϵ is bigger, service reliability is worse, and regulation capacity bid is more aggressive. Therefore, the reliability parameter ϵ essentially captures the trade off between regulation payment and market risks. In this paper, we only consider the case where the aggregator has already come up with a target reliability ϵ. Our end product will be an algorithmic procedure that determines the optimal capacity for any ϵ in the practical regime of interest.

III. Randomized Solution

For notation convenience, let us define $\mathcal{X}_e$ as a set of capacity bids that cover the signal e:

$$\mathcal{X}_e = \{\alpha \in \mathbb{R}^d | L(\alpha, e) \in S_1 \oplus \cdots \oplus S_N\}.$$

$\mathcal{X}_e$ is a set parameterized by e. Since S_i is convex and $L(\alpha, e)$ is linear in α, we can easily verify that $\mathcal{X}_e$ is closed and convex for any $e \in \mathcal{R}^T$. Under this notation, the capacity bidding problem (1) can be reformulated as:

$$\begin{aligned} J_\epsilon^* = & \max_{\alpha \leq \bar{\alpha}} \quad \pi \cdot \alpha \\ & \text{s.t. } \mathbb{P}\{\alpha \in \mathcal{X}_e\} \geq 1 - \epsilon, \end{aligned} \tag{3}$$

where J_ϵ^* is the optimal value of (3).

In general, finding a solution carrying an arbitrary probability of constraint satisfaction is a non-trivial task, especially when the probability measure $\mathbb{P}$ is unknown. In this paper, we consider a randomized solution that uses historical data to provide an accurate approximation of the optimal solution to (3). We divide historical regulation signals into vectors corresponding to contiguous hours, i.e., $e_m \in \mathbb{R}^T$, where $m = 1, 2, \ldots$ corresponds to each hour. Assume that e_m is independently drawn from $\mathbb{P}$:

Assumption 2: Assume that the data $e_1, e_2, \ldots$ are independently drawn from the same distribution $\mathbb{P}$.

Since each data e_m is a vector in $\mathbb{R}^T$, i.e., $e_m = (e_m^1, \ldots, e_m^T)$, we shall emphasize that Assumption 2 only imposes i.i.d. assumptions between random vectors e_i and e_j, but does not impose any conditions on intra-vector correlations. In fact, we allow arbitrary intra-hour correlations between regulation signals $e_m^{t_1}$ and $e_m^{t_2}$. This is distinct from many model-based approaches in the literature.

Remark 2: We impose the i.i.d. assumption to establish theoretical guarantees for the performance of our proposed approach. However, actual regulation signals are not necessarily i.i.d. We will show that our approach performs well on real data, thus justifying the use of our approach in (possibly) non-i.i.d. cases.

From the regulation data, we randomly select M samples, $e_1, \ldots, e_M$, and solve the following program that simultaneously enforces the M sampled constraints:

$$\begin{aligned} & \max_{\alpha \leq \bar{\alpha}} \quad \pi \cdot \alpha \\ & \text{s.t. } \alpha \in \mathcal{X}_{e_m}, \quad \forall m \in \mathcal{M}. \end{aligned} \tag{4}$$

where $\mathcal{M} = \{1, \ldots, M\}$. A distinct feature of (4) is that it has finite number of linear constraints. When the number of M is not too large, it can be efficiently solved as a linear program [17]. On the other hand, a natural question is whether the solution to (4), denoted as $\alpha_{\mathcal{M}}^*$, respects the feasibility constraint of (3). In other words, whether the proportion of constraints that is violated by $\alpha_{\mathcal{M}}^*$ is greater than $1-\epsilon$ under the probability measure $\mathbb{P}$. We define the notion of violation probability as follows:

Definition 1: The violation probability of a given $\alpha \in \mathbb{R}^d$ is defined as $V(\alpha) = \mathbb{P}\{e \in \Delta : \alpha \notin \mathcal{X}_e\}$.

The feasibility constraint of (3) requires that $V(\alpha^*_{\mathcal{M}}) \leq \epsilon$. To this end, we can solve the scenario program (4) that satisfies the M sampled constraints, hoping that a vast majority of unseen constraints automatically take care of themselves. Intuitively, the more samples we draw, the smaller violation probability we have. This indicates that a desired violation probability can be obtained by controlling the size of samples.

Papers [18] and [19] pioneered a feasibility theory that provides a sample complexity upper bound to guarantee any violation probability ϵ. In [19], the authors show that the violation probability of α^*_M satisfies the following condition:

$$\mathbb{P}^M\{V(\alpha^*_{\mathcal{M}}) \geq \epsilon\} \leq \binom{M}{d}(1-\epsilon)^{M-d} \tag{5}$$

where $\mathbb{P}^M$ is the product measure on the space $\Delta \times \cdots \times \Delta$. It also shows that this upper bound is tight for all fully-supported problems (see definition in [19]). However, when the problem is not fully-supported, this bound may be over-conservative: the actual violation probability of the randomized solution is lower than desired, leading to small capacity estimates. Unfortunately, this is the case for the capacity bidding problem: using real data from PJM market, we find that the optimality loss of this approach can be greater than 50% compared to the optimal solution to (3).

Rooted in this optimality loss is the fact that the sample complexity bound (5) guarantees feasibility (chance constraint) but ignores optimality (objective value). This reveals a fundamental trade-off between feasibility and optimality: more samples provide a better feasibility guarantee, but lead to a higher optimality loss. One way to address this trade-off is *sampling and discarding*: we first generate M samples in Δ, then select k out of M samples to be discarded, and solve the scenario program (4) with the remaining $M-k$ samples. The optimal value of (4) will be improved due to constraint removal, and feasibility can be maintained if $M-k$ is large. Clearly, this method works best if the constraints are optimally removed: discarded constraints lead to the best improvement in the cost objective among all possible eliminations of k samples out of M. For notation convenience, denote $\mathcal{A}$ as the set of all subsets of $\mathcal{M}$ with cardinality k, i.e., $\mathcal{A} = \{A | A \subseteq \mathcal{M}, |A| = k\}$, then optimal removal can be defined as follows:

Definition 2: Given M samples $e_1, \ldots, e_M$, $A_{opt} \in \mathcal{A}$ is an optimal removal if it is the optimal solution to the following subset selection problem:

$$\begin{aligned} \max_{A,\alpha} \quad & \pi \cdot \alpha \\ \text{s.t. } & \alpha \in \mathcal{X}_{e_i}, \quad \forall i \in \mathcal{M} \backslash A, A \in \mathcal{A} \end{aligned} \tag{6}$$

Remark 3: In general, optimal removal (6) is a non-trival combinatorial problem. The computational burden of (6) can be prohibitive with large numbers of M and k. One way to address this is a greedy algorithm that removes constraints sequentially, which however, is still impractical for medium-sized problems. In this paper, we will develop an alternative algorithm that provides efficient and accurate approximation for the capacity bidding problem in Section IV.

The proposed sampling and discarding algorithm is summarized in Algorithm 1. When constraint removal is optimal, we refer to it as the *optimal sampling and discarding algorithm*.

Algorithm 1 The optimal sampling and discarding algorihtm

Initialization: Number of samples M and number of discards k.

1: Randomly draw M samples independently from the distribution $\mathbb{P}$.
2: Remove k samples out of M by solving the optimal removal problem (6) and obtain A_{opt}.
3: Solve the following scenario program with the remaining $M-k$ constraints:

$$\begin{aligned} \max_{\alpha \leq \bar{\alpha}} \quad & \pi \cdot \alpha \\ \text{s.t. } & \alpha \in \mathcal{X}_{e_i}, \quad \forall i \in \mathcal{M} \backslash A_{opt}. \end{aligned} \tag{7}$$

Output: The optimal solution α^*.

Intuitively, Algorithm 1 can guarantee feasibility by tuning M (and potentially k), and improves optimality loss by tuning k. When constraint removal is handled optimally, the algorithm can provide near optimal solutions to the chance-constrained problem (3):

Theorem 1: Let $\beta \in (0,1)$ be a small confidence parameter and $v \in (0, \epsilon)$ be a performance degradation parameter. If M and k are such that:

$$\binom{k+d-1}{k} \sum_{i=0}^{k+d-1} \phi(\epsilon, i) + \sum_{i=k+1}^{M} \phi(\epsilon - v, i) \leq \beta. \tag{8}$$

With probability at least $1-\beta$, the solution to Algorithm 1 simultaneously satisfies:
(i): $V(\alpha^*) \leq \epsilon$
(ii): $\pi \cdot \alpha^* \geq J^*_{\epsilon - v}$,
where $\phi(\epsilon, i) = \binom{M}{i}\epsilon^i(1-\epsilon)^{M-i}$, and $J^*_{\epsilon-v}$ is the optimal value to the perturbed problem:

$$\begin{aligned} J^*_{\epsilon-v} = \max_{\alpha} \quad & \pi \cdot \alpha \\ \text{s.t. } & \mathbb{P}\{\alpha \in \mathcal{X}_e\} \geq 1 - \epsilon + v. \end{aligned} \tag{9}$$

Furthermore, result (i) holds even if constraint removal is not optimal.

The proof of Theorem 1 can be found in [20]. Result (i) is a feasibility result, and it is important to note that it holds regardless of whether constraint removal is optimal or not. Result (ii) states that the performance of the optimal sampling and discarding is no worse than that of the perturbed problem (9), where v is the degradation margin. Clearly, v is not a direct measure for the optimality loss, However, when the optimal value of (9) is Lipschitz continuous with respect to v, the degradation margin can be used to control optimality loss. The result of Theorem 1 is very general: it holds for any d and any convex sets $\mathcal{X}_{e_i}$. It is further proved that for any β and $v < \epsilon$, M and k always exist to satisfy (8), and can be found in a principled manner [20].

IV. Implementing Randomized Solution in Capacity Bidding

This section addresses some practical concerns for implementing the sampling and discarding algorithm in the capacity bidding problem. Since historical data on regulation signals is limited [16], and the optimal removal is generally intractable, we will look into the following questions:
Is there enough data for Algorithm 1 to provide good guarantee for reliability?
How to approximately implement optimal removal?

A. Data Adequacy

This section investigates the data adequacy problem. In order to implement Algorithm 1, we need at least M data points from the past. More data is needed to evaluate the performance of Algorithm 1. This is because we further need to separate the data set into training data and testing data.

PJM regulation market provides historical regulation signals for the year of 2017 [16]. The data is recorded every 2 seconds, and normalized between -1 and $+1$. It indicates the imbalance of the overall system between the supply and demand. Since each year has 8760 hours, we have up to 8760 samples. To implement Algorithm 1 with these data, we shall make sure that for β, v and ϵ in the practical regime of interest, the minimum number of samples needed for condition (8) does not exceed 8760.

To verify data adequacy, we let $\beta = 0.01$, and $v = 0.05$. For $d = 1$ and $d = 2$, we draw the minimum number of samples and discards under different target reliability $1 - \epsilon$. The result is shown in Figure 2. It is clear that 8760 data is enough for our problem.

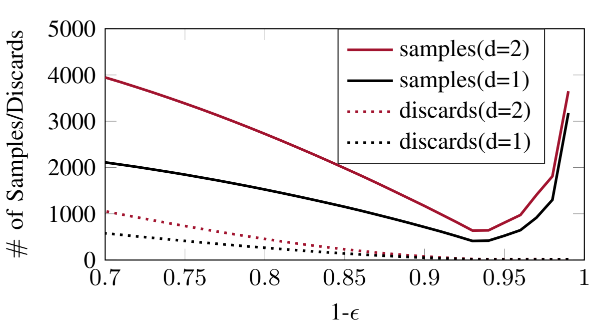

Fig. 2. The number of samples and discards.

Remark 4: In Figure 2, the number of samples first decreases and then increases. The increase is because v has to be smaller than ϵ, so we decrease v as ϵ is smaller than 0.05. However, the decrease of sample complexity in Figure 2 is counter-intuitive: within certain range, to achieve a higher reliability, the algorithm needs less samples. This is because reliability is not the only criteria. When $1 - \epsilon$ is small, the algorithm has to discard more samples to bound the *performance degradation*. In order to remove an extra sample without affecting feasibility, the algorithm may have to draw more than one extra sample points.

B. Optimal Constraint Removal

Optimal constraint removal involves a subset selection problem (6). This is an NP hard combinatorial problem. In this subsection, we discuss (approximate) solutions to (6) for PJM market and CAISO market separately.

1) PJM Market: In the context of PJM market, the *exact* solution to the optimal removal problem (6) can be easily solved. This is due to the following proposition:

Proposition 1: When $d = 1$, $\alpha \in \mathcal{X}_{e_i}$ if and only if $\alpha \leq \alpha_i^*$, where α_i^* satisfies:

$$\begin{aligned} \alpha_i^* = &\max_{\alpha \leq \bar{\alpha}} \quad \alpha \\ &\text{s.t. } \alpha \in \mathcal{X}_{e_i}. \end{aligned} \tag{10}$$

Proposition 1 simply follows from the fact that the Minkowski sum of S_i is convex, and 0 is contained in all the set S_i. It indicates that the scenario program (4) can be transformed to the following form:

$$\begin{aligned} &\max_{0 \leq \alpha \leq \bar{\alpha}} \quad \alpha \\ &\text{s.t. } \alpha \leq \alpha_i^*, \quad \forall i \in \mathcal{M}. \end{aligned}$$

This can be solved by computing α_i^* for each e_i individually, and choose the smallest α_i^* among all $i \in \mathcal{M}$. It further implies that the subset selection problem (6) can be solved by computing α_i^* for all $i \in \mathcal{M}$ and discarding k samples with the smallest α_i, which only involves computing α_i^* and sorting the vector $(\alpha_1^*, \ldots, \alpha_M^*)$.

2) CAISO Market: In CAISO market, each provider submits separate capacity estimates for regulation up and down signals, i.e., $d = 2$. Due to the combinatorial nature of the optimal removal problem, we can not derive the optimal solution to (6) within reasonable computation time. To approximate the solution to (6), one way is to recursively and optimally eliminate groups of p constants ($p \ll k$) at a time. When $p = 1$, this reduces to a greedy algorithm. In greedy algorithms, the scenario program is solved MK times. This is rather challenging given the size of M and k. Another choice is to progressively update the solution by eliminating all the active constraints at the currently reached solution. This approach requires light computation, but we find that its optimality loss can be more than 15%.

In this paper, we propose Algorithm 2 to approximate the solution to (6). It is motivated by the active-set method in convex optimization [21], where we remove samples one by one, and at each time, we choose the sample that corresponds the "most" binding constraints. In our algorithm, the "most" binding sample e_i is measured by the infinity norm of Lagrange multiplier λ_i associated with the *inequality* constraints of this sample. Formally, λ_i is defined as the Lagrange multiplier for (12b) (note that based on [17], (11)-(12) is equivalent to (6)). To implement Algorithm 2, we roughly need to run the scenario program (11) for k times. This is much faster than greedy algorithms.

Algorithm 2 Heuristics for optimal constraint removal

Initialization: M samples and number of discards k.

1: **while** Cardinality of A is less than k **do**

2: Solve the scenario program to obtain λ_i:

$$\max_{\alpha, \{V_n^i\}_{i,n}} \quad \pi \cdot \alpha \tag{11}$$

$$\begin{cases} L(\alpha, e_i) = \sum_{n=1}^{N} V_n^i, \quad \forall i \in \mathcal{M} \backslash A, & \text{(12a)} \\ V_n^i \in S_i, \quad \forall i \in \mathcal{M} \backslash A, \forall n \in \mathcal{N}. & \text{(12b)} \end{cases}$$

3: Find sample i^* that contributes the most binding constraint, i.e., $||\lambda_{i^*}||_\infty \geq ||\lambda_i||_\infty$ for $\forall i \in \mathcal{M} \setminus A$.

4: Update the set of discarded samples, $A = \{A, i^*\}$.

5: **end while**

Output: The set of removed samples A.

V. Case Studies

This section validates the proposed sampling and discarding algorithm for the capacity bidding problem with real data. We consider both PJM and CAISO regulation markets and show that the proposed algorithm attains satisfying performance in both cases.

A. Simulation Setup

Consider an aggregator with N batteries. Each battery is fully described by the following parameters: capacity constraint C_i, maximum charging rate constraint c_i, maximum discharging rate constraint d_i, and initial state of change s_i^0. Given this, resource model S_i is a polytope, and each $s_i \in S_i$ satisfies the following constraints:

$$\begin{cases} -d_i \leq s_i^t \leq c_i, \quad \forall i, t. \\ -C_i s_i^0 \leq \sum_{k=1}^{t} s_i^t \leq C_i - C_i s_i^0, \quad \forall i, t. \end{cases} \tag{13}$$

To setup realistic battery parameters, we take the Nissan Leaf EV as an example. The battery of Nissan Leaf has a capacity of 40kWh. A Level 3 charger can provide 80% capacity within 30 minutes [22], so the maximum charging rate is roughly 1kWh/Min. Typically the maximum discharging rate for battery is greater than the charging rate, so we set it to be around 1.5kWh/Min. After these base parameters are set, we randomly generate the parameters of N batteries according to a uniform distribution in $\pm 25\%$ neighborhood of the base value. The initial state of charge is uniformly generated between $20\% - 80\%$.

To evaluate the performance of the proposed algorithm, we set $N = 5$ and $T = 12$. This corresponds to a regulation signal every 5 minutes. We emphasize that the proposed algorithm only involves linear programs, and scales polynomially with respect to N and T (even including $T = 1800$: one signal every 2 seconds). We choose these parameters to be small, as for *evaluating* the proposed approach, we have to run Algorithm 1 many times to validate the chance-constraint feasibility and obtain the average optimality loss. In addition, we also need to compute the empirical optimal solution as the benchmark, which is rather computationally intensive for large values of N and T. When Algorithm 1 is applied in practice, it only runs once for each delivery hour, and there is no need to compute the empirical optimal solution. Therefore, the proposed algorithm is still tractable when N and T scale up to much a larger size.

We use regulation signal data from PJM for year 2017 [16]. The signal is normalized between -1 and 1, and recorded at a 2 second resolution. Since $T = 12$, we pick the signal at the beginning of each 5 minute interval, and obtain 8760 signal trajectories. Throughout this section, we set $\beta = 0.01$, $v = 0.05$, and ϵ varies from 0.01 to 0.3. Since $v < \epsilon$, when $\epsilon \leq 0.05$, we scale down v accordingly.

B. PJM Regulation Market

To test the sampling and discarding approach, we run Algorithm 1 in the context of PJM market with the previously specified data. In this case, $d = 1$ and optimal removal can be efficiently realized via the method described in Section IV-B.1. The numerical simulation focuses on feasibility and optimality: we will validate the violation probability of the solution to Algorithm 1, and compute its optimlality loss compared to the empirical optimal solution.

We compute the violation probability of the solution to Algorithm 1 in the following manner: at each hour τ of the year 2017, we run Algorithm 1 by randomly sampling data points from the *past*, and discarding samples optimally. For any given ϵ, the number of samples and discards are selected according to Figure 2. We denote the outcome of this algorithm as α_τ^*, and test whether the signal for time $\tau + 1$ is contained in this capacity range, i.e., whether $\alpha_\tau^* e_{\tau+1} \in S_1 \oplus \cdots \oplus S_N$ or not. Next, roll the horizon one step forward. Repeat this procedure for $\tau + 1$, and then do the same for $\tau + 2, \tau + 3, \ldots$. We use the proportion of time instance at which capacity estimates covers the next signal as an indicator of service reliability $1 - \epsilon$. Figure 3 shows the empirical reliability as a function of the anticipated reliability. The curve of empirical reliability is slightly above

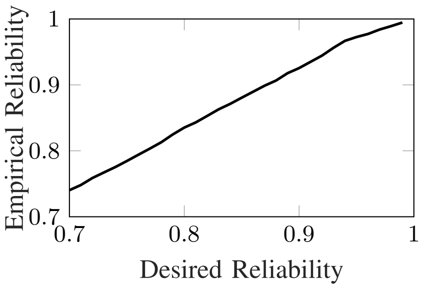

Fig. 3. The empirical service reliability of the capacity estimates for PJM market.

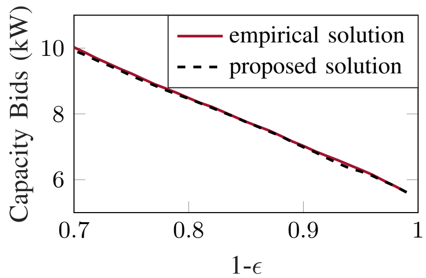

Fig. 4. The capacity bids by the proposed algorithm and brute-force empirical approach for PJM market.

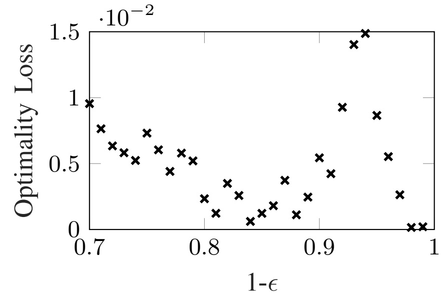

Fig. 5. The percentage optimality loss of the proposed algorithm for PJM market.

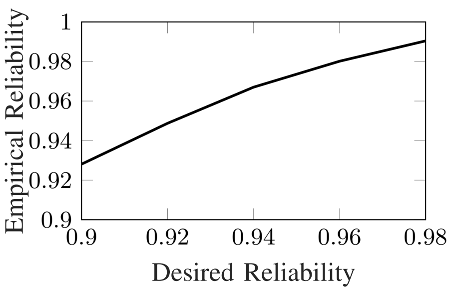

Fig. 6. The empirical service reliability of the capacity estimates for CAISO market.

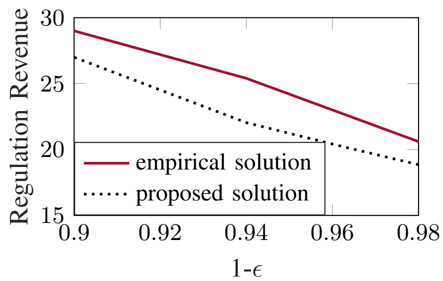

Fig. 7. The capacity bids by the proposed algorithm and brute-force empirical approach for CAISO market.

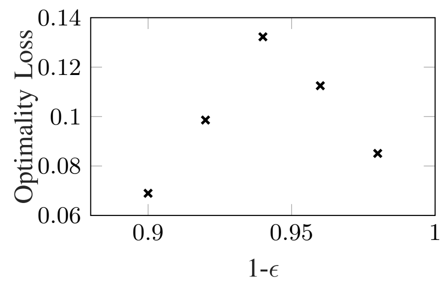

Fig. 8. The percentage optimality loss of the proposed algorithm for CAISO market.

the $45°$ line, indicating that the capacity estimate meets the feasibility requirement.

To obtain the optimality loss, we compare the optimal value of Algorithm 1 with that of (3). Since solution to Algorithm 1 is random, we use the mean of $\pi\alpha_\tau^*$ to represent its optimal value. As the probability measure $\mathbb{P}$ is unknown, we compute the empirical solution to (3) as the benchmark. In particular, we discretize the space of α, and for each α in the space, we empirically compute the violation probability of α by testing whether $\alpha e_i \in S_1 \oplus \cdots \oplus S_N$ for all signals, i.e., $i = 1, \ldots, 8760$. This leads to a monotonically decreasing curve for service reliability as a function of α. Based on this curve, one can find the capacity estimates for any service reliability. Figure 4 show the capacity bids for the proposed algorithm and the empirical optimal solution as a function of $1-\epsilon$. Clearly, the proposed method closely approximates the optimal solution. According to Figure 5, the sampling and discarding algorithm has at most 1.5% error as compared to the empirical optimal solution.

C. CAISO Regulation Market

In CAISO regulation market, each participant submits separate capacity bids for regulation up and down signals. In this paper, we use PJM regulation data to evaluate the algorithm performance in the CAISO market. This can be done by applying a simple transformation of data: if the regulation signal is positive, each agent multiplies the signal by the regulation up capacity. Otherwise, the signal is multiplied by the regulation down capacity.

When $d = 2$, the optimal constraint removal is not tractable. Therefore, we test the feasibility and optimality of Algorithm 1 by a heuristic approximation of optimal constraint removal (Algorithm 2). We let the regulation up price and down price to be the same and unitless, i.e., $\pi(1) = \pi(2) = 1$.

To reduce computation time, we compute the violation probability and optimality loss in the following manner: first, we randomly draw some samples from the *entire* historical data. The number of samples and discards is selected according to Figure 2. Second, we discard samples by running the heuristic approximation in Algorithm 2. The outcome of this procedure provides an optimal solution α^*. Third, repeat this entire process multiple times. With multiple draws of α^*, we compute their average and denote the mean value as $\bar{\alpha}^*$. Fourth, evaluate the empirical violation probability by testing whether $L(\bar{\alpha}^*, e_i) \in S_1 \oplus \cdots \oplus S_N$ is true for all signals, i.e., $i = 1, \ldots, 8760$. Last, compare the optimal value $\pi \cdot \bar{\alpha}^*$ with the empirical optimal value, which is derived by discretizing the space of α and empirically evaluating the violation probability over all data. Figure 6 shows the empirical service reliability as compared to the anticipated reliability. It indicates that the capacity estimate of the proposed algorithm meets the feasibility requirement. Figure 7 show the capacity bids for the proposed algorithm and the empirical optimal solution. The optimality loss is between $6\% - 13\%$ (see Figure 8).

Remark 5: For $d = 1$, we evaluate the feasibility of Algorithm 1 by repeatedly evaluating the solution for thousands of times. This is tractable mainly because we only need to solve (10) once, and then obtain each α_τ^* by simply sorting α_i^*. However, this is not possible as the result in Section IV-B.1 is no longer true for $d = 2$. Therefore, we evaluate the performance of CAISO market in a slightly different way, which requires less computation.

Remark 6: Computing the empirical optimal solution for (3) is rather challenging when $d \geq 2$. This computational burden can be slightly mitigated if we have a good estimate of the optimal solution and only search the space within a small neighborhood around the estimated value. When computational resource permits, the best heuristics for solving (3) is to combine Algorithm 1 and the brute-force search: we run Algorithm 1 as a warm-up, which provides an accountable estimation of the optimal solution, then we search within 10% of its neighborhood by computing the empirical violation probability. We emphasize that this approach cannot be realized without the proposed algorithm providing an initial guess, and the brute-force search does not scale for d greater than 2. In contrast, the proposed sampling and discarding algorithm has great potential in higher dimensions.

VI. CONCLUSION

This paper studies capacity bidding problem for providing frequency regulation services. The service provider estimates the capacity of his resources in advance, so that the real-time regulation signals can be covered with a guaranteed probability. As the regulation signal has an unknown distribution, we proposed a constraint sampling and discarding algorithm, which draws random samples from historical data, and optimally removes samples to maximize the regulation revenue. We show that the algorithm can obtain reliable solutions with bounded optimality loss. We also address the practical concerns for implementation of the algorithm in the capacity bidding problem, including data adequacy and optimal constraint removal. The algorithm is validated with real data. Numerical simulations show that it has promising performance. In future work, we will extend the proposed framework to the energy reserve procurement [17], where the dimension of the problem can be much larger.

REFERENCES

[1] B. J. Kirby. *Frequency regulation basics and trends*. United State Department of Energy, 2005.

[2] PJM. Energy & ancillary services market operations, 2015.

[3] A. Ben-Tal, L. El Ghaoui, and A. Nemirovski. *Robust optimization*. Princeton University Press, 2009.

[4] E. Yao, V. W. Wong, and R. Schober. Robust frequency regulation capacity scheduling algorithm for electric vehicles. *IEEE Transactions on Smart Grid*, 8(2):984–997, 2017.

[5] M. Kazemi, H. Zareipour, N. Amjady, W. D. Rosehart, and M. Ehsan. Operation scheduling of battery storage systems in joint energy and ancillary services markets. *IEEE Transactions on Sustainable Energy*, 8(4):1726–1735, 2017.

[6] E. Vrettos, F. Oldewurtel, and G. Andersson. Robust energy-constrained frequency reserves from aggregations of commercial buildings. *IEEE Transactions on Power Systems*, 31(6):4272–4285, 2016.

[7] L. Fabietti, T. T. Gorecki, F. A. Qureshi, A. Bitlislio**u**glu, I. Lymperopoulos, and C. N. Jones. Experimental implementation of frequency regulation services using commercial buildings. *IEEE Transactions on Smart Grid*, 9(3):1657–1666, 2018.

[8] B. Xu, Y. Shi, D. S. Kirschen, and B. Zhang. Optimal battery participation in frequency regulation markets. *IEEE Transactions on Power Systems*, 2018.

[9] E. Yao, V. W. Wong, and R. Schober. Optimization of aggregate capacity of pevs for frequency regulation service in day-ahead market. *IEEE Transactions on Smart Grid*, 9(4):3519–3529, 2018.

[10] E. Yao, V. W. Wong, and R. Schober. Risk-averse forward contract for electric vehicle frequency regulation service. In *IEEE International Conference on Smart Grid Communications*, pages 750–755. IEEE, 2015.

[11] A. Nemirovski and A. Shapiro. Scenario approximations of chance constraints. In *Probabilistic and randomized methods for design under uncertainty*, pages 3–47. Springer, 2006.

[12] J. Donadee and M. D. Ilić. Stochastic optimization of grid to vehicle frequency regulation capacity bids. *IEEE Transactions on Smart Grid*, 5(2):1061–1069, 2014.

[13] H. Zhang, Z. Hu, E. Munsing, S. J. Moura, and Y. Song. Data-driven chance-constrained regulation capacity offering for distributed energy resources. *IEEE Transactions on Smart Grid*, 2018.

[14] S. Li, W. Zhang, J. Lian, and K. Kalsi. On reverse stackelberg game and optimal mean field control for a large population of thermostatically controlled loads. In *American Control Conference*, pages 3545–3550. IEEE, 2016.

[15] S. I. Vagropoulos and A. G. Bakirtzis. Optimal bidding strategy for electric vehicle aggregators in electricity markets. *IEEE Transactions on power systems*, 28(4):4031–4041, 2013.

[16] http://www.pjm.com/markets-and-operations/ancillary-services.aspx.

[17] A. Shetty, S. Li, K. Poolla, and P. Varaiya. Optimal energy reserve procurement. In *European Control Conference*, 2018.

[18] G. Calafiore and M. C. Campi. Uncertain convex programs: randomized solutions and confidence levels. *Mathematical Programming*, 102(1):25–46, 2005.

[19] M. C. Campi and S. Garatti. The exact feasibility of randomized solutions of uncertain convex programs. *SIAM Journal on Optimization*, 19(3):1211–1230, 2008.

[20] M. C. Campi and S. Garatti. A sampling-and-discarding approach to chance-constrained optimization: feasibility and optimality. *Journal of Optimization Theory and Applications*, 148(2):257–280, 2011.

[21] S. Wright and J. Nocedal. Numerical optimization. *Springer Science*, 35(67-68):7, 1999.

[22] http://www.evtown.org/about-ev-town/ev-charging/charging-levels.html.