Electric Power Systems Research 212 (2022) 108597

Contents lists available at ScienceDirect

Electric Power Systems Research

journal homepage: www.elsevier.com/locate/epsr

Dispatch-aware planning for feasible power system operation☆

Nicolas Christianson[1], Lucien Werner *,[1], Adam Wierman, Steven Low

Computing & Mathematical Sciences, California Institute of Technology, Pasadena, USA

A R T I C L E I N F O

Keywords:
Economic dispatch
Electricity markets
Online algorithms
Resource adequacy
Capacity planning

A B S T R A C T

Maintaining stable energy production with increasing penetration of variable renewable energy requires sufficient flexible generation resources and dispatch algorithms that accommodate renewables' uncertainty. In this work, we study the feasibility properties of real-time economic dispatch (RTED) algorithms and establish fundamental limits on their performance. We propose a joint methodology for resource procurement and online economic dispatch with guaranteed feasibility. Our algorithm, Feasible Fixed Horizon Control (FFHC) is a regularized form of Receding Horizon Control (RHC) that balances exploitation of good near-term demand predictions with feasibility requirements. Empirical evaluation of FFHC in comparison to the standard RHC on realistic load profiles highlights that FFHC achieves near-optimal performance while ensuring feasibility in high-ramp scenarios where RHC becomes infeasible.

1. Introduction

In power systems with high penetrations of variable renewable energy production, sufficient flexible and dispatchable generation resources are necessary to ensure a stable energy supply. However, conventional dispatchable thermal generators are ramp-constrained, limiting how quickly they can modulate their production to accommodate large fluctuations in net demand. This poses a challenge for system operators on two fronts: *resource procurement* and *real-time generation scheduling*.

Resource procurement refers to the system operator's task of planning for sufficient available capacity and ramp for the system to meet uncertain net demand. Resource procurement takes place on longer timescales (e.g., years to day-ahead) and includes several problems familiar to power system operators including security-constrained unit commitment (SCUC), resource adequacy, and capacity planning. On shorter timescales (e.g., 5 to 15 min), system operators must dispatch available generation resources efficiently to meet realized net demand. This is known as real-time economic dispatch (RTED).

Numerous methods have been devised in both of these domains to ensure robustness to uncertainty in net demand. For resource procurement problems, scenario-based optimization is common in practice, while other stochastic optimization techniques and robust optimization have been explored in the research community. For RTED, lookahead dispatch algorithms have been widely implemented by independent system operators (ISOs) in energy markets, and additional ancillary services such as flexible ramping products and load-following reserves have seen some adoption in markets with high demand variability.

The ultimate goal of both resource procurement and RTED is to deliver sufficient generation to meet realized demand while satisfying system constraints: that is, to guarantee feasibility of the dispatch in real time. However, a crucial challenge facing state-of-the-art methods today is that if resource procurement fails to account for the particular dispatch algorithm to be used, or if the RTED algorithm used does not appropriately consider procured resources (e.g., generation & ramp capacity) when making decisions, then feasibility is not assured.[2] Moreover, merging the problems of resource procurement and feasible RTED algorithm synthesis, i.e., optimizing over both system specifications and RTED algorithms, is intractable for the class of general dispatch algorithms.

This motivates the goal of this paper: *developing tractable methods for resource procurement and RTED that together yield provable guarantees of feasibility*.

☆ The authors are supported by the NSF Graduate Research Fellowship, USA (DGE-1745301), Resnick Sustainability Institute Graduate Fellowship, USA, Pimco Graduate Fellowship in Data Science, USA, NSF, USA grants CNS-2106403, NGSDI-2105648, ECCS 1931662, and CPS ECCS 1932611, and Amazon AWS, USA.

* Corresponding author.

E-mail addresses: nchristianson@caltech.edu (N. Christianson), lwerner@caltech.edu (L. Werner), adamw@caltech.edu (A. Wierman), slow@caltech.edu (S. Low).

[1] Equal contribution.

[2] In systems with no inter-temporal coupling constraints (e.g., ramp or state-of-charge constraints), this feasibility mismatch does not arise. However, in practice ramp limits matter, i.e., they constrain the set of operating points reachable by the system.

https://doi.org/10.1016/j.epsr.2022.108597
Received 4 October 2021; Received in revised form 16 April 2022; Accepted 2 July 2022
Available online 25 July 2022

1.1. Contributions

First, to motivate the need for a joint approach to resource procurement and RTED, we show in Section 3 that even on a single-bus system and with nearly full knowledge of future demand, offline feasibility over a set of demand trajectories is insufficient to guarantee the existence of *any* online dispatch algorithm that can feasibly meet those demand sequences.

Second, a practical joint algorithm for resource procurement and RTED is presented in Section 4. The first step is a robust optimization problem called *Dispatch-Aware Planning* (DAP), which determines adequate system capacity to ensure a feasible RTED algorithm exists. The second step is a dispatch algorithm called *Feasible Fixed Horizon Control (*FFHC*)* that minimally modifies the standard receding horizon control (RHC) algorithm to robustly use trusted predictions of demand.

Third, in Section 5 we give matching upper and lower bounds on the competitive ratio of *any* feasible online dispatch algorithm. These bounds imply identical bounds on FFHC.

Finally, we evaluate the proposed approach on a synthetic system derived from CAISO demand and generation data. We show that FFHC retains the excellent average-case performance of RHC but in cases where there are large demand fluctuations and the system is ramp constrained, RHC fails while FFHC remains feasible with minimal additional cost. The test cases presented in this work are designed to clearly demonstrate the feasibility properties of our approach. Although our algorithm is applicable to realistic problems like SCUC, large-scale simulations are not explored here and are the subject of future work.

1.2. Related work

This work bridges the online algorithms and power systems literatures. We briefly highlight some related work in each of these domains.

Online algorithms. RTED is an online decision-making problem characterized by a challenging combination of time-coupling and unknown time-varying constraints. None of the existing constrained online optimization literature, e.g., [1–4], directly handles our setting.

The authors in [1] explore a related ramp-constrained online optimization problem, yet feasibility does not pose an issue due to the lack of unknown time-varying constraints. Recent work in online optimal control considers time-invariant [2] as well as time-varying and coupling constraints [3] on state and action. However, the feasibility guarantees depend on advance knowledge of the constraints.

The work in [4] comes closest to our setting. The authors optimize over affine policies to design algorithms for online optimization with switching costs and ramp limits that are robust to polytopic uncertainty in certain constraints. However, their approach does not consider the problem of guaranteeing feasibility, and their proposed algorithm is unable to fully exploit good predictions of near-term uncertainty.

Power systems. Resource procurement for system reliability and multi-interval economic dispatch are two key problems in power systems operation addressed by this work.

Our formulation of the resource procurement problem has broad applicability to several problems in power system reliability: in particular security-constrained unit commitment (SCUC) for day-ahead markets [5–9], resource adequacy [10–12], and capacity planning [13–15]. Most of the approaches in this literature do not consider behavior of causal RTED algorithms or generally involve scenario-based optimization [16,17]. In practice, resource adequacy planning relies on regulatory standards and scenario-based studies and similarly ignore the behavior of the RTED algorithm.

Many ISOs have implemented multi-interval lookahead optimization for RTED, as it can better accommodate variability in forecasts for renewables and intertemporal constraints from conventional generation and storage [18–21] Ancillary services such as flexible ramping products [9,22,23] and load-following reserves [24–26] have been studied and implemented in some markets. To our knowledge, all of the aforementioned proposals for multi-interval dispatch do not provide provable guarantees for the feasibility of the lookahead optimization problem. There is prior work on utilizing affine policies to robustly dispatch reserves in the real-time market when ramp constraints are present [27,28], but this work similarly does not explicitly consider the question of feasibility, and the affine policies utilized may be more conservative than the lookahead dispatch algorithms used by operators.

Our work is most closely related to research on adaptive robust unit commitment with causal affine real-time policies in [29,30]. Here, robust policy-aware economic dispatch is combined with robust unit commitment, and an algorithmic framework for efficient computation of large-scale problems is proposed. Like [4], robustness comes at the expense of fully utilizing predictions. In contrast, we bring an online algorithms perspective to the problem of feasible RTED, focusing primarily on (a) designing feasible RTED algorithms that can fruitfully exploit predictions, and (b) characterizing the performance of feasible RTED algorithms in general.

1.3. Notation

$N \in \mathbb{Z}_+$ is the number of dispatchable generators and $T \in \mathbb{Z}_+$ is the length of the time horizon. We denote the ordered set of time intervals between indices a and b by $[a, b] := \{a, \dots, b\} \subset \mathbb{Z}_+$. The inequalities in (1) and subsequent optimization problems are element-wise.

2. Problem formulation

The problem of optimal power system planning and operation can be cast as a sequential optimization problem robust to uncertainty revealed prior to each stage.[3]

$$
\begin{aligned}
&\min_{\mathbf{y}} \max_{d_1} \min_{\mathbf{x}_1} \cdots \max_{d_T} \min_{\mathbf{x}_T} \quad \bar{\mathbf{c}}^\top \mathbf{y} + \sum_{t=1}^{T} \mathbf{c}_t^\top \mathbf{x}_t && \text{(1a)}\\
&\quad \text{s.t.} \quad \mathbf{1}^\top \mathbf{x}_t = d_t \qquad \forall t \in [1, T] && \text{(1b)}\\
&\qquad\quad g_t(\mathbf{x}_t, \mathbf{y}) \le \mathbf{0} \qquad \forall t \in [1, T] && \text{(1c)}\\
&\qquad\quad h_t(\mathbf{x}_{t-1}, \mathbf{x}_t) \le \mathbf{0} \qquad \forall t \in [1, T] && \text{(1d)}\\
&\qquad\quad (d_1, \dots, d_T) \in \mathcal{D} && \text{(1e)}
\end{aligned}
$$

For concreteness, we limit our presentation to a single planning stage with decision variables $\mathbf{y} \in \mathbb{R}^K$ (e.g., generator capacities, ramp/line limits, unit commitments) followed by T generation dispatch stages, each with decision variables $\mathbf{x}_t \in \mathbb{R}^N, t = 1, \dots, T$, where the initial operating point $\mathbf{x}_0$ is fixed. We assume the cost functions for planning variables $\bar{\mathbf{c}}$ and dispatches $\mathbf{c}_t$ are linear and known by the system operator. Constraint (1b) is the supply–demand balance constraint where d_t is the demand at time t. Constraints (1c) and (1d) are affine and represent capacity/planning constraints and intertemporal (ramp, state-of-charge) constraints respectively. We focus on a single-bus network with all dispatchable generators satisfying a net load trajectory $\mathbf{d} = (d_1, \dots, d_T)$, which is contained in a bounded, known uncertainty set $\mathcal{D}$.[4] We assume that $\mathcal{D}$ is polytopic, i.e. it takes the form $\mathcal{D} = \{\mathbf{d} \in \mathbb{R}^T : \mathbf{E}\mathbf{d} \le \mathbf{f}\}$ with parameters $\mathbf{E} \in \mathbb{R}^{L \times T}$ and $\mathbf{f} \in \mathbb{R}^L$ known to the system operator prior to solving the planning and dispatch problem.

Each dispatch stage depends on the planning decision as well as the previous dispatch. An example of a problem falling under this framework is SCUC followed by multi-interval real-time dispatch. However, this framework can be extended to include several planning stages in advance of dispatch, such as capacity planning and intraday unit commitment.

[3] For simplicity, in this work we assume that the only uncertainty is the demand, although uncertainty in generation (e.g., solar, wind) can also be accommodated.

[4] Everything that follows can be extended to multi-bus setting with network constraints, as in [29,30].

2.1. Planning problem

The goal of the planning problem is to determine a choice $\mathbf{y}^*$ of the planning decisions. Given that problem (1) is intractable due to the sequential $\min-\max-\min$ operators, an approach taken by power system operators in practice is to choose $\mathbf{y}^*$ by solving an "offline" form of the problem where the demand sequence $\mathbf{d} = (d_1, \dots, d_T) \in \mathcal{D}$ (or a small number of scenarios) is assumed known in advance.

$$\min_{\substack{\mathbf{y}\in\mathbb{R}^K\\ \mathbf{x}_t\in\mathbb{R}^N}} \quad \bar{\mathbf{c}}^\top \mathbf{y} + \sum_{t=1}^{T} \mathbf{c}_t^\top \mathbf{x}_t \tag{2a}$$

$$\text{s.t.} \quad \mathbf{1}^\top \mathbf{x}_t = d_t \qquad \forall t \in [1,T] \tag{2b}$$

$$g_t(\mathbf{x}_t, \mathbf{y}) \le \mathbf{0} \qquad \forall t \in [1,T] \tag{2c}$$

$$h_t(\mathbf{x}_{t-1}, \mathbf{x}_t) \le \mathbf{0} \qquad \forall t \in [1,T] \tag{2d}$$

As written, (2) is a linear program; when $\mathbf{y}$ represents unit commitments, (2) becomes a MILP with the addition of integrality constraints on $\mathbf{y}$ (not shown).

While the resulting $\mathbf{y}^*$ from this offline optimization would be ex-post optimal, were the assumed demand sequence the true demand, this will not generally be the case, as the planning problem is typically solved far in advance, when there is still uncertainty in future demand. To provide stronger guarantees in the face of demand uncertainty, system operators may wish for $\mathbf{y}^*$ to satisfy (2b)–(2d) *for any* $\mathbf{d} \in \mathcal{D}$. This motivates the following definition of offline feasibility.

Definition 2.1. A planning decision $\mathbf{y}^*$ is *offline feasible* if and only if for all $\mathbf{d} \in \mathcal{D}$ there exists a dispatch sequence $\mathbf{x}_1, \dots, \mathbf{x}_T$ satisfying the dispatch feasibility constraints (1b)–(1d).

2.2. Online dispatch problem

After the planning variables $\mathbf{y}^*$ are chosen, the task of the system operator is to determine real-time dispatches $\mathbf{x}_t$. They do so via an *online dispatch algorithm*: a sequence of functions $X_1, \dots, X_T$, each of which maps a demand sequence to a dispatch for time t: $X_t : \mathcal{D} \to \mathbb{R}^N$. Crucially, the collection of functions $\{X_t\}_{t=1}^T$ must be *causal*, so the decision $X_t(\mathbf{d})$ at time t can only depend on information known to the system operator at time t. We will assume that the system operator knows the exact demand d_t at time t, and also has access to perfect predictions of demand $d_{t+1}, \dots, d_{t+h}$ within a short lookahead window of length h. Thus $X_t(\mathbf{d})$ may only depend on demands through time $\min\{t+h, T\}$.

A desirable objective for an online dispatch algorithm is the satisfaction of dispatch feasibility constraints. This motivates the following definition of *online feasibility* of a dispatch algorithm as well as of a planning decision $\mathbf{y}^*$.

Definition 2.2.

(1) Given a fixed planning decision $\mathbf{y}^*$, a *feasible* online dispatch algorithm is a sequence of causal policies $\{X_t\}_{t=1}^T$ with the property that for any demand sequence $\mathbf{d} \in \mathcal{D}$, the produced decisions $X_1(\mathbf{d}), \dots, X_T(\mathbf{d})$ satisfy the constraints (1b)–(1d).

(2) If $\mathbf{y}^*$ admits a feasible online dispatch algorithm, then $\mathbf{y}^*$ is said to be an *online feasible* planning decision.

A particular online dispatch algorithm that is widely used in practice is Receding Horizon Control (RHC), where at time t, dispatches are optimized over the h-step perfect lookahead horizon $[t, t+h]$. Only the first dispatch $\mathbf{x}_t^*$ is committed at each step of RHC; the remaining dispatches over the lookahead horizon are merely "advisory". This process is then repeated for the subsequent interval $t+1$, and so on. However, RHC has a significant downside in that it is not necessarily feasible, even if the planning decision $\mathbf{y}^*$ is online feasible. We demonstrate situations when RHC lacks feasibility in Section 6.

3. Offline feasibility does not imply online feasibility

Although we presented the planning and dispatch problems separately in the previous section, we show now why the common practice of solving for the planning variables $\mathbf{y}^*$ offline in a dispatch-unaware fashion can ultimately cause online dispatch infeasibility. We answer the following question: *For a particular demand uncertainty set, does offline feasibility of planning decisions* $\mathbf{y}^*$ *necessarily imply their online feasibility?*

We answer this question in the negative in the following theorem. This establishes that anything short of full knowledge of the demand sequence is insufficient for offline feasibility to imply online feasibility.

Theorem 3.1. *There exist choices of affine system constraints* $\{g_t\}$, $\{h_t\}$, *a polytopic demand uncertainty set* $\mathcal{D}$, *and fixed planning decisions* $\mathbf{y}^*$ *that are offline feasible, yet which are not online feasible if* $h < T-1$.

For the sake of brevity, we do not present the detailed proof of Theorem 3.1 here. The gist of the argument is as follows: we construct a ramp-limited system and two demand sequences $\mathbf{d}$ and $\hat{\mathbf{d}}$ differing only in their final value d_T. For each of these sequences $\mathbf{d}, \hat{\mathbf{d}}$, there is only a single offline feasible dispatch trajectory $\{\mathbf{x}_t\}, \{\hat{\mathbf{x}}_t\}$ (respectively). Moreover, these demand sequences require different initial dispatch decisions: $\mathbf{x}_1 \neq \hat{\mathbf{x}}_1$. Any online dispatch algorithm with lookahead $h < T-1$ cannot know the final demand d_T when it makes a decision for time $t = 1$. Thus any such online dispatch algorithm has no way of knowing whether the true demand sequence is $\mathbf{d}$ or $\hat{\mathbf{d}}$, and since the *necessary* feasible decisions $\mathbf{x}_1, \hat{\mathbf{x}}_1$ that must be chosen in either of these scenarios are different, the online dispatch algorithm cannot choose correctly.

4. Joint algorithm for system planning and online dispatch

The counterexample in Theorem 3.1 establishes that the offline feasibility of a planning decision $\mathbf{y}^*$ for a particular demand uncertainty set does not imply its online feasibility. This motivates the joint approach in Algorithm 1, where we use affine policies to guarantee the existence of an online feasible $\mathbf{y}^*$ as well as an online feasible dispatch algorithm. Affine policies approximate online decision making during the planning stage (called *Dispatch-aware Planning (*DAP*)*), before passing the optimal planning variables to the RTED algorithm (called *Feasible Fixed Horizon Control (*FFHC*)*). Rather than using the (conservative) affine policies for determining actual dispatch schedules in real time, we subtly modify the standard RHC dispatch algorithm to include an affine-policy-based regularization term on the last decision of each subhorizon. This allows for online scheduling to exploit accurate short-term predictions without taking decisions that are too myopic. Other variants of fixed horizon control, like AFHC [31] or CHC [32], can be substituted for RHC in our algorithm at the expense of more burdensome notation.

Algorithm 1 Joint algorithm for planning and dispatch

```
1: input: Cost functions c̄, c_t and constraints g_t, h_t
2: Solve DAP problem (4)
3: Fix optimal planning variables y*
4: for t = 1, ..., T do
5:     d_{t+h} revealed
6:     Solve FFHC problem (5) with d_1, ..., d_{t+h} and x*_{t-1} as parameters
7:     return x*_t
8: end for
```

4.1. Dispatch-aware planning (DAP)

The dispatch-aware planning problem is defined in (4) below. The demand sequence $\mathbf{d} = (d_1, \dots, d_T)$ resides in a known polytopic demand uncertainty set $\mathcal{D}$, and the linear planning/dispatch cost functions are known as well. Piecewise-linear cost functions can also be accommodated with additional notation.

The real-time scheduling policies $\{X_t(\cdot)\}_{t=1}^{T}$ are defined to be affine in the demand trajectory:

$$X_t(\mathbf{d}) := \mathbf{A}_t\mathbf{d} + \mathbf{b}_t \quad \forall t \in [1, T] \tag{3}$$

To optimize over the X_t is to optimize over the matrices $\mathbf{A}_t \in \mathbb{R}^{N\times T}$ and vectors $\mathbf{b}_t \in \mathbb{R}^N$. The X_t are causal, meaning $\mathbf{A}_t$ have 0's for all columns with index greater than t. This requirement can be enforced with entrywise constraints on the matrices. It is assumed that (4) has a feasible solution.

$$\min_{\substack{\mathbf{y} \\ X_1,\dots,X_T}} \max_{\mathbf{d}\in\mathcal{D}} \quad \bar{\mathbf{c}}^\top \mathbf{y} + \sum_{t=1}^{T} \mathbf{c}_t^\top X_t(\mathbf{d}) \tag{4a}$$

$$\text{s.t.} \quad \mathbf{1}^\top X_t(\mathbf{d}) = d_t \qquad \forall t \in [1, T] \tag{4b}$$

$$g_t(X_t(\mathbf{d}), \mathbf{y}) \leq \mathbf{0} \qquad \forall t \in [1, T] \tag{4c}$$

$$h_t(X_{t-1}(\mathbf{d}), X_t(\mathbf{d})) \leq \mathbf{0} \quad \forall t \in [1, T] \tag{4d}$$

As convention we assume $X_0(\mathbf{d}) = \mathbf{x}_0$. Problem (4) is a linear program with semi-infinite constraints resulting from the "$\forall$" qualification on $\mathbf{d}$. Using strong duality of linear programs, (4) can be equivalently posed as a linear program with a finite number of additional variables and constraints [33]. The result is a tractable linear program that can be solved with off-the-shelf optimization solvers. However, depending on the length of the time horizon and the complexity of $\mathcal{D}$, (4) can be challenging to scale to large problem sizes. Although this scaling is *not* the focus of this paper, strategies for improving scaling are discussed in Section 6.3.

After solving (4), the optimal planning variables $\mathbf{y}^*$ are fixed. It is then possible to schedule in real time using the optimal affine policies $X_1^*, \dots, X_T^*$ applied to the real-time demand sequence. The resulting generation schedules are guaranteed to be feasible – that is, they satisfy (1b)–(1d) for any $\mathbf{d} \in \mathcal{D}$ – but because of their robustness and inability to incorporate more refined demand predictions, the cost of the dispatch is likely to be quite conservative. In contrast, the algorithm we propose next uses the policies to constrain online dispatches to an always-feasible region while still allowing accurate short-term predictions to be exploited.

4.2. Feasible fixed horizon control (FFHC)

Economic dispatch in real-time (e.g., 5-min, 15-min) electricity markets is often a multi-interval optimization problem over an $(h+1)$-step horizon from which only the first dispatch decision is binding and the remaining are advisory. In the control literature this algorithm is referred to as receding horizon control (RHC) or model predictive control (MPC).

The version of fixed horizon control that we propose here, called *Feasible Fixed Horizon Control* (FFHC), is RHC with the addition of a robust affine constraint composed from the optimal policies from (4). FFHC is parameterized by the optimal solutions $\mathbf{y}^*, X_1^*, \dots, X_T^*$ from (4).[5] FFHC has access to d_t and h (perfect predictions of) future demand $d_{t+1}, \dots, d_{t+h}$, as well as the previously committed dispatch $\mathbf{x}_{t-1}^*$. The first decision $\mathbf{x}_t^*$ in the subhorizon of the following optimization problem determines the decision of FFHC in time t, denoted henceforth by FFHC(t).[6]

$$\operatorname*{arg\,min}_{\mathbf{x}_t,\dots,\mathbf{x}_{t+h}} \quad \sum_{s=t}^{t+h} c_s^\top \mathbf{x}_s \tag{5a}$$

$$\text{s.t.} \quad \mathbf{1}^\top \mathbf{x}_s = d_s \qquad \forall s \in [t, t+h] \tag{5b}$$

$$g_s(\mathbf{x}_s, \mathbf{y}^*) \leq \mathbf{0} \qquad \forall s \in [t, t+h] \tag{5c}$$

$$h_s(\mathbf{x}_s, \mathbf{x}_{s-1}, \mathbf{y}^*) \leq \mathbf{0} \qquad \forall s \in [t+1, t+h] \tag{5d}$$

$$h_t(\mathbf{x}_t, \mathbf{x}_{t-1}^*, \mathbf{y}^*) \leq \mathbf{0} \tag{5e}$$

$$h_{t+h+1}(X_{t+h+1}^*(\mathbf{d}), \mathbf{x}_{t+h}, \mathbf{y}^*) \leq \mathbf{0} \tag{5f}$$

$$\forall \mathbf{d} \in \mathcal{D}_{0:t+h}$$

$\mathcal{D}_{0:t+h}$ is the restricted set of demand sequences in $\mathcal{D}$ that are possible given the already-revealed demand values from time 0 to $t+h$. In general for a pair of indices $r, s \in [0, T]$ and $r \leq s$ and $\tilde{\mathbf{d}}_{r:s} = (\tilde{d}_r, \dots, \tilde{d}_s)$ a subsequence of realized values, we define

$$\mathcal{D}_{r:s} := \{\mathbf{d} \in \mathcal{D} \mid d_t = \tilde{d}_t \ \forall t = r, \dots, s\} \tag{6}$$

[7] Clearly, $\mathcal{D}_{r:s} \subseteq \mathcal{D} \subseteq \mathbb{R}^{T+1}$.

As was explored in Theorem 3.1, just enforcing constraints (5b)–(5e) does not always yield a feasible solution. Because the ramping constraint ties the previously committed decision $\mathbf{x}_{t-1}^*$ to all subsequent dispatches, a short-sighted dispatch early on could lead to infeasibility for a subsequent round. The addition of robust constraints (5f) on the last decision $\mathbf{x}_{t+h}$ ensures that earlier decisions are robust to future uncertainty. As in (4), the robust constraint in (5) can be transformed into auxiliary linear constraints on $\mathbf{x}_{t+h}$. Taking the optimal solutions of the first variables $\mathbf{x}_t^*$ from each subhorizon for $t = 1, \dots, T$ gives the dispatch sequence from FFHC, which, as presented in the following theorem, is feasible.

Theorem 4.1. FFHC *is a feasible online dispatch algorithm. That is, for any* $\mathbf{d} \in \mathcal{D}$, $T-h$ *successive rounds of* FFHC(t) *produce a dispatch sequence* $\mathbf{x}_1^*, \dots, \mathbf{x}_T^*$ *that satisfies* (1b)–(1d).

Proof. The result is shown by inductively constructing feasible solutions for each round of FFHC, starting with the dispatch provided by the affine policies in the first round and matching constraints between problems (4) and (5) in the subsequent rounds. The full proof is omitted here. □

Remark. The terminal constraint (5f) is only applied through FFHC($T-h-1$). Subsequent rounds of FFHC have no demand uncertainty. FFHC's guaranteed feasibility distinguishes it from [4]. Moreover, the placement of constraint (5f) on the terminal decision enable FFHC to fully exploit all perfect predictions of demand, in contrast to work in [4,29,30].

5. Upper and lower bounds on feasible online dispatch algorithms

We now turn to bounding the worst-case performance of the class of feasible dispatch algorithms, which contain our proposed FFHC as an instance. The exactly matching upper and lower bounds we obtain establish fundamental limits on the performance of algorithms for RTED. They also establish that, in general, feasibility of an online algorithm implies optimality. In other words, *feasibility is the best*

[5] X_t^* refers to $(\mathbf{A}_t^*, \mathbf{b}_t^*)$.

[6] At $t = T-h$, the optimal solution of (5) determines FFHC's remaining dispatch decisions $\mathbf{x}_{T-h}^*, \dots, \mathbf{x}_T^*$ because by that time, the entire demand sequence is known.

[7] The explicit dependence of $\mathcal{D}_{r:s}$ on $\tilde{\mathbf{d}}_{r:s}$ is suppressed in the notation for simplicity.

you can do. Nonetheless, different algorithms can be distinguished in their average-case performance, as we examine in the experiments in Section 6.

We evaluate performance via the metric of competitive ratio, which has recently seen increasing use in the control and power systems communities [34–36]:

$$\mathrm{CR}_{\mathrm{ALG}} = \sup_{\mathbf{d}\in\mathcal{D}} \frac{\mathrm{Cost}_{\mathrm{ALG}}(\mathbf{d})}{\mathrm{Cost}_{\mathrm{OPT}}(\mathbf{d})} \tag{7}$$

A competitive ratio of 1 signifies optimal performance of the online algorithm, whereas a competitive ratio larger than 1 indicates suboptimal performance. We choose to focus on the competitive ratio because it is unitless and time-independent, thus facilitating fair comparison of algorithm performance across different problem instances and system parameters.[8]

In the following theorems we assume for clarity of exposition that $\mathcal{D} \subseteq \mathbb{R}^T_{\geq 0}$, dispatch variables $\mathbf{x}_t$ are always nonnegative, and costs are linear, positive, and potentially time-varying. However, performance bounds can be obtained in more general settings.

Theorem 5.1. *Suppose* ALG $:= \{X_t\}_{t=1}^T$ *is a feasible online dispatch algorithm for demand uncertainty set $\mathcal{D}$ on some arbitrary system, where costs are linear and time-varying, $\mathbf{c}_1, \dots, \mathbf{c}_T \in \mathbb{R}^N_{>0}$. Then, the competitive ratio of* ALG *is bounded above as:*

$$\mathrm{CR}_{\mathrm{ALG}} \leq \max_{s,t\in[1,T]} \frac{c_{\max,s}}{c_{\min,t}}. \tag{8}$$

where $c_{\max,t} = \max_{i\in[1,N]} c_{i,t}$ *and* $c_{\min,t} = \min_{i\in[1,N]} c_{i,t}$

Proof. The upper bound follows immediately from upper (lower) bounding the online (offline) unit cost in each timestep by the unit cost of the most (least) expensive generator at any time. The full proof is omitted here due to space limitations. □

In the following, we construct a system with linear, time-invariant costs and generators which have capacity and ramp constraints. We assume planning decisions $\mathbf{y}^*$ result in a certain generator having arbitrarily slow ramp limit ϵ. Thereby, in the regime where $\epsilon \to 0$, there is a demand trajectory on which any feasible dispatch algorithm must have competitive ratio arbitrarily close to the upper bound (8). This constitutes an exactly tight lower bound on the competitive ratio of any feasible dispatch algorithm.

Theorem 5.2. *Fix* $\epsilon \in (0, 1)$. *There exists a choice of system parameters $\mathbf{y}^*(\epsilon)$ with linear, time-invariant costs $\mathbf{c} \in \mathbb{R}^N_{\geq 0}$ and a polytopic demand uncertainty set $\mathcal{D}$, as well as a distinguished demand sequence $\hat{\mathbf{d}} \in \mathcal{D}$, such that for any feasible online dispatch algorithm* ALG $:= \{X_t\}_{t=1}^T$,

$$\frac{\mathrm{Cost}_{\mathrm{ALG}}(\hat{\mathbf{d}})}{\mathrm{Cost}_{\mathrm{OPT}}(\hat{\mathbf{d}})} \geq \epsilon + (1-\epsilon)\frac{c_{\max}}{c_{\min}} \tag{9}$$

where $c_{\max} := \max_{i\in[1,N]} c_i$ *and* $c_{\min} := \min_{i\in[1,N]} c_i$.

Proof of Theorem 5.2. Fix h to be some positive integer, independent of T. We construct a 2-generator system with costs $\mathbf{c} = (c_{\max}, c_{\min})$; capacity lower and upper bounds $\underline{\mathbf{x}} = (0, 0)$, $\overline{\mathbf{x}} = (2h, 2h)$; ramp lower and upper bounds $\overline{\Delta} = (\epsilon, 2-\epsilon)$ and $\underline{\Delta} = (-\epsilon, -2+\epsilon)$; and initial operating point $\mathbf{x}_0 = ((2-\epsilon)h, \epsilon h)$. We define the demand uncertainty set $\mathcal{D} \subset \mathbb{R}^T$ as follows:

$$\mathcal{D} = \left\{ \mathbf{d} : d_0 = 2h, d_t \leq d_{t+1} \leq d_t + 2, d_t \leq 4h \,\forall t \in [0, T] \right\},$$

where we define $d_0 = \mathbf{1}^\top \mathbf{x}_0 = 2h$. Observe that $\mathcal{D}$ admits an online feasible algorithm: specifically, the online algorithm that chooses its operating point at time t in the set $\left\{ \mathbf{x} : \mathbf{1}^\top \mathbf{x} = d_t, \mathbf{x} = \lambda \mathbf{x}_0 + (1-\lambda)\overline{\mathbf{x}}, \lambda \in [0, 1] \right\}$ is feasible for $\mathcal{D}$.

Now consider the specific demand trajectory $\hat{\mathbf{d}}$ with $\hat{d}_t = 2h$ for all $t \in [1, T]$. We claim that, for all times $t \in [1, T-2h]$, $X_t(\hat{\mathbf{d}})_1 \geq (2-2\epsilon)h$. We prove this by contradiction: suppose alternatively that $X_t(\hat{\mathbf{d}})_1 < (2-2\epsilon)h$ for some $t \in [1, T-2h]$. But consider another demand trajectory $\tilde{\mathbf{d}} \in \mathcal{D}$ defined by $\tilde{d}_s = 2h$ for $s \in [1, t+h]$ and $\tilde{d}_s = \min\{2h+2(s-t-h), 4h\}$ for $s \in [t+h+1, T]$. As $\hat{\mathbf{d}}$ and $\tilde{\mathbf{d}}$ coincide in their entries through time $t+h$, causality dictates that $X_t(\hat{\mathbf{d}}) = X_t(\tilde{\mathbf{d}})$, so $X_t(\tilde{\mathbf{d}})_1 < (2-2\epsilon)h$ as well. But then the online algorithm cannot remain feasible for $\tilde{\mathbf{d}}$ for the rest of time: this is because feasibly meeting $\tilde{\mathbf{d}}$ for the rest of time, and in particular remaining feasible for the sequence of demand increases beginning at time $t+h+1$, requires $X_{t+h}(\tilde{\mathbf{d}}) = \mathbf{x}_0$ due to the ramp and capacity constraints. However, since $X_t(\tilde{\mathbf{d}})_1 < (2-2\epsilon)h$ and up-ramp on the first generator is bounded by ϵ, it is impossible for the online algorithm to reach $\mathbf{x}_0$ at time $t+h$. Thus the online algorithm cannot be feasible for $\mathcal{D}$, yielding a contradiction.

By the last paragraph's result, we know that for $t \in [1, T-2h]$, $X_t(\hat{\mathbf{d}})_1 \geq (2-2\epsilon)h$; as $X_t(\hat{\mathbf{d}})_1 = d_t - X_t(\hat{\mathbf{d}})_2$, it follows that the cost of the online algorithm on each of the first $T-2h$ timesteps is lower bounded by $(2-2\epsilon)hc_{\max} + 2h\epsilon c_{\min}$. On each of the last $2h$ timesteps, we trivially lower bound the online algorithm cost by $2hc_{\min}$. Thus we obtain

$$\mathrm{Cost}_{\mathrm{ALG}}(\hat{\mathbf{d}}) \geq (T-2h)\left((2-2\epsilon)hc_{\max} + 2h\epsilon c_{\min}\right) + (2h)^2 c_{\min} \tag{10}$$

Now we turn to providing an upper bound on $\mathrm{Cost}_{\mathrm{OPT}}(\hat{\mathbf{d}})$. Since the offline optimal knows all demands in advance, it will ramp maximally to transfer all generation onto the second, cheaper generator and will remain at this operating point for the rest of time. It will take $\frac{(2-\epsilon)h}{\epsilon}$ timesteps to ramp to the operating point $(0, 2h)$, since generator 1 has ramp limit ϵ and demand is constant through time. For each of these first $\frac{(2-\epsilon)h}{\epsilon}$ timesteps, we upper bound the offline optimal cost trivially by $2hc_{\max}$ in each step. Once the offline optimal reaches $(0, 2h)$, its cost in each step for the rest of time is exactly $2hc_{\min}$. We get:

$$\mathrm{Cost}_{\mathrm{OPT}}(\hat{\mathbf{d}}) \leq \frac{(2-\epsilon)h}{\epsilon}(2hc_{\max}) + \left(T - \frac{(2-\epsilon)h}{\epsilon}\right)(2hc_{\min}) \tag{11}$$

Forming the ratio of (10) with (11) and taking the limit as $T \to \infty$ yields the lower bound (9). □

6. Experiments

In this section we explore through simulations on simple systems how the proposed algorithm handles infeasibilities that otherwise arise when resource procurement is done in a dispatch-agnostic fashion. We also include a discussion about the scalability of the method to larger, more realistic power systems (see Fig. 1).

6.1. A two-generator case

We use a two-generator case (same as the one presented in the proof of Theorem 3.1) to show how FFHC is able to compute a feasible dispatch when the standard RHC algorithm cannot.

Setting $h = 2$ and $T = 4$, we define a demand uncertainty set $\mathcal{D} = \{(2, 2, 2, d) : d \in [1, 4]\}$. We run DAP on this system with costs $\mathbf{c} = (1, 3/4)$, $\overline{\mathbf{c}} = (10, 11)$, nominal max capacity $\overline{\mathbf{x}} = (2, 2)$, nominal ramp rates $\underline{\Delta}^{\mathrm{nom}} = \overline{\Delta}^{\mathrm{nom}} = (2, 1/4)$, and starting point $\mathbf{x}_0 = (1, 1)$. For offline planning and dispatch, the nominal max capacity of $(2, 2)$ is sufficient to satisfy all $\mathbf{d} \in \mathcal{D}$. DAP procures an additional 37.5% capacity and a proportional amount of ramp capacity on (lower cost) Generator 1.

Figs. 1(a) and 1(b) show the performance of the algorithms RHC, FFHC, and RAP (affine policies synthesized in DAP), as well as the offline optimal, on the two demand sequences $\mathbf{d}^{(A)}$ and $\mathbf{d}^{(B)}$ distinguished in Theorem 3.1. $\mathbf{d}^{(A)}$ is an "easy" demand sequence: all algorithms are feasible, and FFHC takes more conservative (i.e., costlier) decisions

[8] Competitive difference upper and lower bounds can be obtained for FFHC, and more generally for arbitrary feasible online dispatch algorithms, that essentially match those in [1], with slight modifications due to the inclusion of supply–demand balance constraints in our setting. Further, [1] obtains upper and lower bounds on competitive difference matching up to a factor of 4, whereas our competitive ratio upper and lower bounds match exactly.

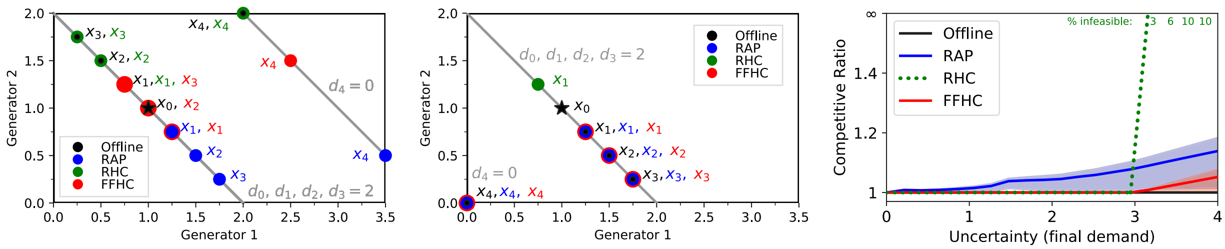

(a) “Easy” case: $\mathbf{d}^{(A)}$ (b) “Hard” case: $\mathbf{d}^{(B)}$ (c) Competitive ratio vs. uncertainty

Fig. 1. Counterexample from Theorem 3.1 revisited. Panel (a) shows a scenario where all algorithms are feasible. Panel (b) shows a scenario where RHC is unable to remain feasible, whereas FFHC remains feasible. Panel (c) illustrates the performance of the RHC, FFHC, and RAP against the offline optimal as the size of the uncertainty set grows. The dotted green line going to ∞ indicates that RHC becomes infeasible on some trajectories beginning at an uncertainty value of 3.

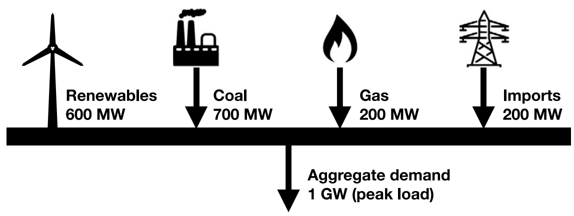

Fig. 2. Power system for 4-generator cast study. Capacities shown are peak values.

Table 1
Parameters for generation sources. Costs are from [39, Table 1]. Max capacity values indicate the maximum available generation for each type. CAISO generation mix is used to derive a import cost [40]. Ramp rates are taken from reasonable ranges given in [41,42].

Generation type	Max Cap. (MW)	Ramp rate (% cap./min)	Variable cost (\$/MWh)	Cap. cost (\$/MW)
Imports	200	±5	1.93	0
Gas	200	±2	2.56	1.08×10^6
Coal	700	±0.5	4.52	3.67×10^6
Renewables	600	Instantaneous	0	NA

than RHC, which immediately moves up to the top left of the capacity region to exploit the lower cost of Generator 2. On the other hand, $\mathbf{d}^{(B)}$ details the “hard” demand sequence, for which RHC is unable to remain feasible, since it mistakenly chooses to exploit the lower cost of Generator 2 production at $t = 1$, leaving it unable to meet d_4. FFHC is able to remain feasible in contrast.

Fig. 1(c) compares the performance of the algorithms via competitive ratio as the demand uncertainty set is scaled. We parameterize the uncertainty set by u:

$$\mathcal{D}(u) := \left\{(2, 2, 2, d) : d \in \left[2 - u/2, 2 + u/2\right]\right\}$$

For $u \in [0, 4]$, we run DAP on $\mathcal{D}(u)$ to determine system parameters. We then sample trajectories from $\mathcal{D}(u)$ set using a hit-and-run sampler for polytopes [37], and compute the dispatch of each algorithm. The mean empirical competitive ratio of the trajectories along with upper/lower bounds (shaded) for each algorithm are shown in Fig. 1(c).

While the performance of RAP suffers in comparison to that of the offline optimal, both FFHC and RHC exactly match the offline performance for $u < 3$. For $u > 3$, RHC begins encountering infeasibility on some of the demand trajectories, and by $u = 4$ is infeasible for 10% of the sampled trajectories. Meanwhile, FFHC, just like RAP, always remains feasible, though its performance degrades slightly from that of the offline since its dispatches are influenced by the robust constraint.

6.2. Scenario based on CAISO load profile

The purpose of this example is to show the necessity of dispatch-aware planning to maintain feasibility under realistic net load profiles. For simplicity, we do not incorporate integer unit commitment variables and associated cost functions and therefore the setting is not intended to represent the particular variety of unit commitment problems solved by system operators.

We consider the small power system shown in Fig. 2, which has 1 GW of peak load and four generation sources: variable renewables (wind & solar), a fast-ramping gas turbine, a slow-ramping coal plant, and a transmission interconnection. This setup, while stylized, represents the scenario of a transmission-constrained zone within a larger grid where local infeasibilities could arise under high fluctuations of net demand and ramp shortages.

Details on capacity, ramp rates, and costs for the generation sources are given in Table 1. The 24 hr nominal generation profile sampled at 15 min intervals is taken from CAISO’s aggregate demand on Sept. 9, 2021 [38]. We subtract the variable renewable generation profile (from [38]) to get a nominal *net* demand curve (solid black line in Fig. 3(a)), around which a demand uncertainty set is constructed. Trajectories are sampled uniformly from this set using a hit-and-run sampler [37]. Capacity costs are used in the DAP problem to determine an optimal robust generation mix (neither renewables nor imports are included in this step as they are considered already fixed).

Due to the fast, sustained afternoon ramp event in the net load profile, the standard RHC dispatch runs into infeasibility for 29% of sampled demand trajectories. In contrast, FFHC is always feasible at little to no extra cost beyond that incurred by the offline optimal. Fig. 3(b) shows the optimal solutions for a particular net demand trajectory. Prior to 20.5 h, both algorithms return identical solutions. After that, RHC becomes infeasible whereas FFHC does not. At the start of the ramp event, imports are already at their maximum and RHC chooses to myopically exploit the lower cost of the gas generator. In contrast, FFHC “pre-ramps” the slow coal generator and saves ramping capacity on the gas generator to accommodate later fluctuations. It is also notable that even though the planning variables (generator capacities and ramps) in this example are online feasible, RHC is still unable to remain feasible. Thus, the existence guarantee of an online feasible dispatch algorithm does not imply that even a good (on average) policy like RHC can produce a feasible sequence of dispatches.

Fig. 3(c) shows that the feasibility guarantee of FFHC comes at a very minimal efficiency loss. When RHC is can stay feasible, both algorithms attain near-optimal cost with average empirical competitive ratios $\mathrm{CR}_{\mathrm{RHC}} = 1.0000$ and $\mathrm{CR}_{\mathrm{FFHC}} = 1.0002$. In comparison, the RAP algorithm, while robustly feasible, has a significantly higher average competitive ratio of 1.0934, indicating the value of using predictions (robustly!) in our approach. The robust resource procurement step (DAP) in this simulation procures 1066.6 MW of total generation capacity, which is 14% above peak demand of 938.8 MW. In ramp constrained power systems, additional capacity may be required to accommodate long high ramp events. DAP provides a way to directly optimize for this margin.

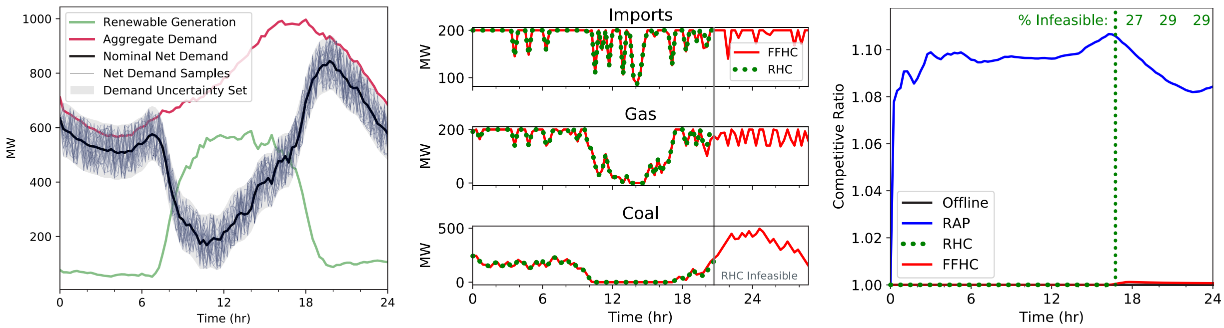

(a) Aggregate demand, renewable generation, and net nominal demand, along with additional net demand trajectories from demand uncertainty set. U/L bounds are $\pm 20\%$ of nominal.

(b) Dispatch solutions for RHC vs. FFHC on a particular demand trajectory that is infeasible for RHC. Lookahead horizon is set at 1 hr for both.

(c) Average competitive ratio of each algorithm plotted vs. time.

Fig. 3. Results from four-generator system using CAISO load and renewable generation profiles. Approximately 29% of 300 sampled demand trajectories are infeasible for RHC. The vertical dotted green line in panel (c) indicates at which time RHC first becomes infeasible and the values at the top of the frame display the percentage of infeasible trajectories (out of the total sampled) at regular time intervals.

6.3. Discussion of algorithm scalability

DAP requires solving a robust linear program, a problem known to suffer from scalability issues. While scalability is not the focus of this work, in this section we discuss effective strategies for reducing the problem dimension and highlight relevant existing literature on this subject.[9]

Optimization problems for large N-generator power systems (e.g., SCUC, capacity planning) already present operators with a demanding computational task, with $O(NT)$ variables and $O(M)$ constraints where M can be as large as $O(N^2T)$ for mesh network topologies. We are concerned with the *additional* complexity our robust linear formulation adds to this baseline, which arises from (1) the robust description of the uncertainty set $\mathcal{D}$ and (2) expressiveness of the causal affine policy class.

For (1), we take the reasonable assumption that correlations between elements of the demand vector are limited to neighbors. This means that the number of constraints in $\mathcal{D}$ is $O(T)$, rather than $O(T^2)$ which would arise if full correlations were allowed. For (2) we observe that synthesized policies often only make use of a few previous demand steps, which we call memory m with $m = O(1)$. This allows us to limit the size of the affine policies to Nm variables, as opposed to NT for full-history policies. Using limited memory policies necessitates a careful reformulation of $\mathcal{D}$ and the problem constraints, but the downstream benefits for the size of the robust LP are significant, as the total number of constraints ultimately scales with the number of policy variables. Restricting policy memory also eliminates the $O(NT^2)$ causality constraints required for the full policies.

Table 2 summarizes the number of constraints and variables (in order sense) for each problem setting. Limited-memory policies allow for the multiplicative factor of T in both variables and constraints for the full-memory robust formulation to be reduced to a (tunable and small) constant factor m.

After reducing the problem size in the proposed manner, the resulting problem may still be a large LP. We point the interested reader to the excellent discussion of this issue in [29,30] where a constraint generation approach along with various other algorithmic tweaks allow for efficient solutions to large LP/MILP power system problems. All of the proposed methods therein are applicable to our setting.

[9] We focus on DAP for scalability; the FFHC stage of our joint algorithm only includes a small robust constraint that does not appreciably affect computation.

Table 2
Comparison of number of variables and constraints for offline and two robust formulations.

	Offline	Memory-T policies	Memory-m policies
Variables	$O(NT)$	$O(NT^2 + MT)$	$O(mNT + mM)$
Constraints	$O(M)$	$O(NT^2 + MT)$	$O(mM)$

7. Conclusion

In this work, we analyze properties of feasible online dispatch algorithms in general, and specifically propose a joint algorithm for resource procurement and RTED that exploits lookahead predictions for good performance while also guaranteeing feasibility. Our framework is applicable to several types of resource procurement problem including SCUC and resource adequacy, and is compatible with arbitrary fixed-horizon lookahead optimization problems. We further present exactly matching upper and lower bounds on the competitive ratio for the problem class of RTED. Finally, our computational results demonstrate that FFHC nearly matches the performance of the offline optimal while always remains feasible, which contrasts with the frequent infeasibility of RHC. Thus the proposed approach provides feasible RTED with nearly no loss of efficiency compared to the standard algorithm.

Future work includes applying this algorithmic framework to problems with energy storage and time-varying state-of-charge requirements, as well as designing incentive compatible prices for dispatches computed by feasible RTED algorithms.

CRediT authorship contribution statement

Nicolas Christianson: Conceptualization, Methodology, Software, Formal analysis, Writing. **Lucien Werner:** Conceptualization, Methodology, Software, Formal analysis, Writing. **Adam Wierman:** Writing. **Steven Low:** Writing.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests:

Misha Chertkov
Chenye Wu
Pascal van Hentenryck.

References

[1] M. Badiei, N. Li, A. Wierman, Online convex optimization with ramp constraints, in: 2015 54th IEEE Conference on Decision and Control (CDC), IEEE, Osaka, 2015, pp. 6730–6736, [Online]. Available: http://ieeexplore.ieee.org/document/7403279/.

[2] Y. Li, S. Das, N. Li, Online optimal control with affine constraints, 2020, arXiv:2010.04891 [cs, eess, math]. [Online]. Available: http://arxiv.org/abs/2010.04891.

[3] T. Li, Y. Chen, B. Sun, A. Wierman, S.H. Low, Information aggregation for constrained online control, Proc. ACM Meas. Anal. Comput. Syst. 5 (2) (2021) 18:1–18:35, [Online]. Available: https://doi.org/10.1145/3460085.

[4] M. Shi, X. Lin, S. Fahmy, Competitive online convex optimization with switching costs and ramp constraints, IEEE/ACM Trans. Netw. 29 (2) (2021) 876–889, Conference Name: IEEE/ACM Transactions on Networking.

[5] R. Billinton, M. Fotuhi-Firuzabad, A reliability framework for generating unit commitment, Electr. Power Syst. Res. 56 (1) (2000) 81–88, [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0378779600001048.

[6] U. Ozturk, M. Mazumdar, B. Norman, A solution to the stochastic unit commitment problem using chance constrained programming, IEEE Trans. Power Syst. 19 (3) (2004) 1589–1598.

[7] L. Wu, M. Shahidehpour, T. Li, Stochastic security-constrained unit commitment, IEEE Trans. Power Syst. 22 (2) (2007) 800–811.

[8] D. Bertsimas, E. Litvinov, X.A. Sun, J. Zhao, T. Zheng, Adaptive robust optimization for the security constrained unit commitment problem, IEEE Trans. Power Syst. 28 (1) (2013) 52–63, Conference Name: IEEE Transactions on Power Systems.

[9] Q. Wang, B.-M. Hodge, Enhancing power system operational flexibility with flexible ramping products: A review, IEEE Trans. Ind. Inf. 13 (4) (2017) 1652–1664.

[10] J. Bushnell, Electricity resource adequacy: Matching policies and goals, Electr. J. 18 (8) (2005) 11–21, [Online]. Available: https://www.sciencedirect.com/science/article/pii/S104061900500103X.

[11] P.A. Ruiz, G. Gross, Short-term resource adequacy in electricity market design, IEEE Trans. Power Syst. 23 (3) (2008) 916–926.

[12] J.P. Pfeifenberger, K. Spees, Resource Adequacy Requirements: Reliability and Economic Implications, Tech. Rep., The Brattle Group, 2013, [Online]. Available: http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.641.5126&rep=rep1&type=pdf.

[13] M. Xu, X. Zhuan, Optimal planning for wind power capacity in an electric power system, Renew. Energy 53 (2013) 280–286, [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0960148112007185.

[14] P. Cramton, A. Ockenfels, S. Stoft, Capacity market fundamentals, Econ. Energy Environ. Policy 2 (2) (2013) 27–46, [Online]. Available: https://www.jstor.org/stable/26189455.

[15] S. Dehghan, N. Amjady, A.J. Conejo, Reliability-constrained robust power system expansion planning, IEEE Trans. Power Syst. 31 (3) (2016) 2383–2392.

[16] K. Dvijotham, M. Chertkov, S. Backhaus, Storage sizing and placement through operational and uncertainty-aware simulations, in: 2014 47th Hawaii International Conference on System Sciences, 2014, pp. 2408–2416, ISSN: 1530-1605.

[17] K. Baker, G. Hug, X. Li, Energy storage sizing taking into account forecast uncertainties and receding horizon operation, IEEE Trans. Sustain. Energy 8 (1) (2017) 331–340.

[18] X. Han, H. Gooi, D. Kirschen, Dynamic economic dispatch: feasible and optimal solutions, IEEE Trans. Power Syst. 16 (1) (2001) 22–28, Conference Name: IEEE Transactions on Power Systems.

[19] L. Xie, M.D. Ilic, Model predictive economic/environmental dispatch of power systems with intermittent resources, in: 2009 IEEE Power Energy Society General Meeting, 2009, pp. 1–6, ISSN: 1932-5517.

[20] J.E. Price, M. Rothleder, Recognition of extended dispatch horizons in California's energy markets, in: 2011 IEEE Power and Energy Society General Meeting, 2011, pp. 1–5, ISSN: 1944-9925.

[21] E. Ela, M. O'Malley, Scheduling and pricing for expected ramp capability in real-time power markets, IEEE Trans. Power Syst. 31 (3) (2016) 1681–1691, Conference Name: IEEE Transactions on Power Systems.

[22] B. Wang, B.F. Hobbs, A flexible ramping product: Can it help real-time dispatch markets approach the stochastic dispatch ideal? Electr. Power Syst. Res. 109 (2014) 128–140, [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0378779613003386.

[23] M. Ghaljehei, M. Khorsand, Day-ahead operational scheduling with enhanced flexible ramping product: Design and analysis, IEEE Trans. Power Syst. (2021) 1, Conference Name: IEEE Transactions on Power Systems.

[24] E. Ela, M. Milligan, B. Kirby, Operating Reserves and Variable Generation, TP-5500-51978, NREL, 2011, [Online]. Available: https://www.nrel.gov/docs/fy11osti/51978.pdf.

[25] N.G. Paterakis, O. Erdinc, A.G. Bakirtzis, J.P.S. Catalão, Load-following reserves procurement considering flexible demand-side resources under high wind power penetration, IEEE Trans. Power Syst. 30 (3) (2015) 1337–1350.

[26] A. Muzhikyan, A.M. Farid, K. Youcef-Toumi, Relative merits of load following reserves & energy storage market integration towards power system imbalances, Int. J. Electr. Power Energy Syst. 74 (2016) 222–229, [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0142061515002999.

[27] J. Warrington, P.J. Goulart, S. Mariéthoz, M. Morari, Robust reserve operation in power systems using affine policies, in: 2012 IEEE 51st IEEE Conference on Decision and Control (CDC), 2012, pp. 1111–1117, ISSN: 0743-1546.

[28] J. Warrington, C. Hohl, P.J. Goulart, M. Morari, Rolling unit commitment and dispatch with multi-stage recourse policies for heterogeneous devices, IEEE Trans. Power Syst. 31 (1) (2016) 187–197.

[29] A. Lorca, X.A. Sun, E. Litvinov, T. Zheng, Multistage adaptive robust optimization for the unit commitment problem, Oper. Res. 64 (1) (2016) 32–51, [Online]. Available: https://search.ebscohost.com/login.aspx?direct=true&db=bsx&AN=114900408&site=eds-live&scope=site&authtype=ip,guest&custid=s8984125&groupid=main.

[30] A. Lorca, X.A. Sun, Multistage robust unit commitment with dynamic uncertainty sets and energy storage, IEEE Trans. Power Syst. 32 (3) (2017) 1678–1688.

[31] M. Lin, Z. Liu, A. Wierman, L.L.H. Andrew, Online algorithms for geographical load balancing, in: 2012 International Green Computing Conference (IGCC), 2012, pp. 1–10.

[32] N. Chen, J. Comden, Z. Liu, A. Gandhi, A. Wierman, Using predictions in online optimization: Looking forward with an eye on the past, in: Proceedings of the 2016 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Science, ACM, Antibes Juan-les-Pins France, 2016, pp. 193–206, [Online]. Available: https://dl.acm.org/doi/10.1145/2896377.2901464.

[33] A. Ben-Tal, A. Nemirovski, Robust solutions of uncertain linear programs, Oper. Res. Lett. 25 (1) (1999) 1–13, [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0167637799000164.

[34] S. Chen, Y. Ji, L. Tong, Large scale charging of electric vehicles, in: 2012 IEEE Power and Energy Society General Meeting, 2012, pp. 1–9, ISSN: 1944-9925.

[35] V. Chau, S. Feng, N.K. Thang, Competitive algorithms for demand response management in smart grid, in: M.A. Bender, M. Farach-Colton, M.A. Mosteiro (Eds.), LATIN 2018: Theoretical Informatics, in: Lecture Notes in Computer Science, vol. 10807, Springer International Publishing, Cham, 2018, pp. 303–316, [Online]. Available: http://link.springer.com/10.1007/978-3-319-77404-6_23.

[36] B. Sun, T. Li, S.H. Low, D.H.K. Tsang, ORC: An online competitive algorithm for recommendation and charging schedule in electric vehicle charging network, in: Proceedings of the Eleventh ACM International Conference on Future Energy Systems, ACM, Virtual Event Australia, 2020, pp. 144–155, [Online]. Available: https://dl.acm.org/doi/10.1145/3396851.3397727.

[37] F. Font-Clos, Hit and run, 2021, https://github.com/fontclos/hitandrun.

[38] California ISO - Today's Outlook. [Online]. Available: http://www.caiso.com/TodaysOutlook/Pages/default.aspx.

[39] Cost and Performance Characteristics of New Generating Technologies, Annual Energy Outlook 2021, Tech. Rep., Energy Information Administration (EIA), 2021, [Online]. Available: https://www.eia.gov/outlooks/aeo/assumptions/pdf/table_8.2.pdf.

[40] 2021 Summer Loads and Resource Assessment, Tech. Rep., CAISO, 2021, [Online]. Available: http://www.caiso.com/Documents/2021-Summer-Loads-and-Resources-Assessment.pdf.

[41] B. Kirby, M. Milligan, Method and Case Study for Estimating the Ramping Capability of a Control Area or Balancing Authority and Implications for Moderate or High Wind Penetration: Preprint, Tech. Rep. NREL/CP-500-38153, National Renewable Energy Lab., Golden, CO (US), 2005, [Online]. Available: https://www.osti.gov/biblio/15016292.

[42] A. Feldmüller, Flexibility of Coal and Gas Fired Power Plants, Paris, 2017, [Online]. Available: https://www.cleanenergyministerial.org/sites/default/files/2018-01/Andreas%20Feldmueller%20Siemens.pdf.