Les Cahiers du GERAD ISSN: 0711–2440

Online dynamic submodular optimization

A. Lesage-Landry, J. Pallage

G–2023–17

June 2023

La publication de ces rapports de recherche est rendue possible grâce au soutien de HEC Montréal, Polytechnique Montréal, Université McGill, Université du Québec à Montréal, ainsi que du Fonds de recherche du Québec – Nature et technologies.

Dépôt légal – Bibliothèque et Archives nationales du Québec, 2023
– Bibliothèque et Archives Canada, 2023

The publication of these research reports is made possible thanks to the support of HEC Montréal, Polytechnique Montréal, McGill University, Université du Québec à Montréal, as well as the Fonds de recherche du Québec – Nature et technologies.

Legal deposit – Bibliothèque et Archives nationales du Québec, 2023
– Library and Archives Canada, 2023

GERAD HEC Montréal
3000, chemin de la Côte-Sainte-Catherine
Montréal (Québec) Canada H3T 2A7

Tél. : 514 340-6053
Téléc. : 514 340-5665
info@gerad.ca
www.gerad.ca

Online dynamic submodular optimization

Antoine Lesage-Landry[a, b, c]

Julien Pallage[a, b, c]

[a] *Department of Electrical Engineering, Polytechnique Montréal, Montréal (Qc), Canada, H3T 1J4*

[b] *GERAD, Montréal (Qc), Canada, H3T 1J4*

[c] *Mila, Montréal (Qc), Canada, H2S 3H1*

antoine.lesage-landry@polymtl.ca
julien.pallage@polymtl.ca

June 2023
Les Cahiers du GERAD
G–2023–17

Abstract : We propose new algorithms with provable performance for online binary optimization subject to general constraints and in dynamic settings. We consider the subset of problems in which the objective function is submodular. We propose the online submodular greedy algorithm (`OSGA`) which solves to optimality an approximation of the previous round's loss function to avoid the NP-hardness of the original problem. We extend `OSGA` to a generic approximation function. We show that `OSGA` has a dynamic regret bound similar to the tightest bounds in online convex optimization. For instances where no approximation exists or a computationally simpler implementation is desired, we design the online submodular projected gradient descent (`OSPGD`) by leveraging the Lovász extension. We obtain a regret bound that is akin to the conventional online gradient descent (`OGD`). Finally, we numerically test our algorithms in two power system applications: fast-timescale demand response and real-time distribution network reconfiguration.

Keywords: Eemand response, dynamic regret, network reconfiguration, online optimization, submodular minimization

---

Acknowledgements: This work was funded by the Institute for Data Valorization (IVADO) and by the National Science and Engineering Research Council of Canada (NSERC).

1 Introduction

Online dynamic decision-making aims at consecutively providing decisions to minimize each round's objective function while relying only on the outcome of previous rounds. The objective function is considered to be time-varying and decisions are made at each discretized time instance. Moreover, it is assumed to be unknown at the time decisions have to be made. The online convex optimization (OCO) [1–3] framework assumes a convex objective function and a convex and compact decision set. Provable performance guarantees can be established under some additional assumptions, e.g., boundedness of the objective function and its gradient [1], or the cumulative difference in round optima computed in hindsight [1, 4].

Online optimization is an appealing framework for real-time decision-making problems because it uses computationally efficient and scalable updates and provides performance guarantees. For example, it is used in the context of moving target tracking [5, 6], dynamic resource allocation in data centers [7, 8], dynamic pricing in power systems [9], or renewable generation intermittency mitigation [10, 11].

Online binary optimization [2, 12] considers a subset of problems in which the feasible set is the intersection of an application-specific constraint set and the binary set $\{0,1\}^n$, where $n \in \mathbb{N}$ is the decision variable's dimension. To tackle efficiently constrained, non-linear online binary problems, we further assume that the objective function is submodular. Specifically, in this work, we consider online dynamic submodular optimization for which the objective is to provide the round optimal binary decisions. We propose two types of algorithms: (i) greedy approaches that solve approximations of the previous round's objective function (Section 3) and (ii) a projected gradient-based approach using the continuous and convex Lovász extension of submodular functions (Section 4). For all algorithms, we provide a performance analysis based on the dynamic regret. The regret bounds are shown to be sublinear in the number of rounds under different conditions on the variation between round optima computed in hindsight. Under these assumptions, the time-averaged dynamic regret vanishes as the time horizon increases and are, therefore, Hannan-consistent [13].

Related work

We now review the relevant literature on online non-linear binary optimization. Linearity simplifies considerably the problem as argued by [12] and, for this reason, is not considered. Examples of online binary approaches for linear problems include [14, 15]. In [2], the authors first studied the online submodular optimization problem. They only considered the static setting in which the decisions are benchmarked with a single static decision computed in hindsight. This is referred to as static regret analysis [1]. They further restricted their analysis to unconstrained problems. Reference [12] then proposed approaches to integrate constraints within online submodular optimization. They also limit their analysis to the static setting. In both cases, greedy and projected gradient-based approaches are proposed. In this work, we provide dynamic regret analysis for all our approaches which then provides a performance guarantee with respect to the round optimum. This latter aspect is important in an engineering setting because one wants to achieve optimality at each round, e.g., to track a time-varying setpoint. In the dynamic setting, [11] used randomization and online convex optimization to solve problems with convex objective functions, i.e., convex with respect to the convex hull of the decision set. However, [11] do not admit other than binary constraints and the dynamic regret analysis does not hold asymptotically.

In the power system literature, several approaches based on time-varying optimization have been proposed to deal with binary decision variables. References [16, 17] apply the error diffusion algorithm to obtain binary decisions from continuous decisions computed via the relaxed problem. In [18], randomization is used to convert continuous decisions to binary ones. This body of literature does not compare the round minima with the algorithm's decisions like online optimization does using a dynamic regret analysis.

Specifically, in this work, we make the following contributions:

- We propose two algorithms for online dynamic submodular optimization. Under the submodularity assumption, we provide the online binary optimization algorithm with the tightest dynamic regret bound to this date.
- We formulate a greedy algorithm that solves a β-approximation of the previous round's objective function. When this approximation is not available, we show that a generic approximation can be used with limited impact on the performance bound.
- We provide a computationally efficient and scalable algorithm for very fast timescale online problems which only performs a single project gradient descent step on the Lovász extension of the objective function.
- We establish conditions under which our algorithms lead to a sublinear dynamic regret bound in the number of rounds, thus showing that our approaches are Hannan-consistent.
- We numerically evaluate the performance of our approaches in power system examples. First, we use the projected gradient-descent update to dispatch demand response resources for frequency regulation. Second, we apply the greedy update to real-time network configuration where line switches can be controlled (ON/OFF) to minimize the powerline congestion while spanning a radial network.

2 Preliminaries

In this section, we introduce the online optimization settings and provide the relevant background on submodular functions.

2.1 Online optimization

In online optimization, a round-dependent objective function must be minimized at each round $t \in \{1, 2, \ldots, T\}$, where $T \in \mathbb{N}$ is the time horizon. In this setting, the objective function is assumed to be observed only after the decision maker has implemented the round's decision, which must be provided on a fast-timescale.

We consider a subset of online binary optimization problems with the base set $V = \{1, 2, 3, \ldots, n\}$ in which the objective function is assumed to be submodular. Let the power set 2^V represent the set of all possible decisions. In each round $t \in \{1, 2, \ldots, T\}$, a decision $S \subseteq 2^V$ must be made. The problem takes the form:

$$\min_{S \in \mathcal{S}} f_t(S), \tag{1}$$

where $f_t : 2^V \mapsto \mathbb{R}$ is a submodular set function, $\mathcal{S} \subseteq 2^V$ is the feasible set, i.e., the set that expresses the problem's constraints, and $t \in \{1, 2, \ldots, T\}$.

As noted by [12], at time t, (1) is NP-hard if $\mathcal{S} \neq 2^V$. Because no offline optimization algorithm can solve (1) given f_t in polynomial time, we benchmark the decisions provided by our online optimization algorithm with an offline α-approximation algorithm [12]. Let $S_t^\star \in \arg\min_{S \subseteq \mathcal{S}} f_t(S)$. An α-approximation algorithm provides a solution S_t such that $f(S_t^\star) \leq f_t(S_t) \leq \alpha f(S_t^\star)$, where $\alpha \geq 1$. Building on [12], we define the *dynamic* α-regret to characterize the performance of our online optimization approaches.

Definition 1. The dynamic α-regret $\mathtt{R}_\alpha^{\mathrm{d}}(T)$ over a time horizon T is:

$$\mathtt{R}_\alpha^{\mathrm{d}}(T) = \sum_{t=1}^{T} \left(f_t(S_t) - \alpha f_t(S_t^\star) \right),$$

where S_t is the decision provided by the online optimization algorithm at round t.

We note that the special case $\mathcal{S} = V$ can be solved to optimality in polynomial time. At this time, we let $\alpha = 1$ and we retrieve the standard dynamic regret definition from OCO [1]. This fact is later used to specialize our results.

The dynamic α-regret defers from the α-regret employed in [12] because it uses as comparators to the algorithm's decision (first term of the sum), the round optima instead of the best fixed decision in hindsight. The power system applications, later discussed in Section 5, motivate the use of a framework that targets round optimal decisions instead of an averaged, static decision. Dynamic regret bounds are given in terms of the cumulative round optimum variation V_T or a derivative of it [1, 19]. This term is used as a complexity measure in dynamic problem [4].

We conclude by adapting the definition of V_T to online set function optimization from which online submodular optimization is a subset of. Let $\boldsymbol{\chi}_A \in \{0,1\}^n$ where $\boldsymbol{\chi}_A$'s i^{th} component is one if and only if $i \in A$ and zero otherwise be the characteristic vector of the set $A \subseteq 2^V$. Consider the online binary optimization problem counterpart of (1):

$$\min_{\mathbf{x} \in \mathcal{X} \cap \{0,1\}^n} f_t^{\mathrm{b}}(\mathbf{x}),$$

where $\mathcal{X} \subseteq \mathbb{R}^n$ is the constraint set and $f^{\mathrm{b}} : \mathcal{X} \cap \{0,1\}^n \mapsto \mathbb{R}$ is the objective function. Let $\mathbf{x}_t^\star \in \arg\min_{\mathbf{x} \in \mathcal{X} \cap \{0,1\}^n} f^{\mathrm{b}}(\mathbf{x})$. The cumulative variation term V_T, as in standard online (convex) optimization, is [1, 11]:

$$\begin{aligned} V_T &= \sum_{t=2}^{T} \left\| \mathbf{x}_t^\star - \mathbf{x}_{t-1}^\star \right\|_2 \\ &= \sum_{t=2}^{T} \left\| \boldsymbol{\chi}_{S_t^\star} - \boldsymbol{\chi}_{S_{t-1}^\star} \right\|_2, \end{aligned}$$

where $S_t^\star \subseteq 2^V$ is the subset of $\mathbf{x}_t^\star$'s components with value one. Adapting V_T to set-valued objective functions, we, therefore, obtain:

$$V_T = \sum_{t=2}^{T} \sqrt{\operatorname{card}\left(S_t^\star \ominus S_{t-1}^\star\right)},$$

where $\ominus$ is the symmetric difference or disjunctive union of two sets. When the context requires it, we will introduce alternative V_T definitions, e.g., when the optima are defined from function approximations.

2.2 Submodularity

A function $f_t : 2^V \mapsto \mathbb{R}$ is submodular if it exhibits the diminishing marginal return property [12], i.e., if

$$f_t(A \cup \{i\}) - f_t(A) \geq f_t(B \cup \{i\}) - f_t(B),$$

for all $A \subseteq B \subseteq V$ and $i \in V$. The Lovász extension $\hat{f} : [0,1]^n \mapsto \mathbb{R}$ of a function f_t can be defined as:

$$\hat{f}_t(\mathbf{x}) = \sum_{i=1}^{n} x_i \left(f_t(\{1,2,\ldots,i\}) - f_t(\{1,2,\ldots,i-1\}) \right),$$

where x_i is the i^{th} largest component of $\mathbf{x}$, $\{0\} \equiv \emptyset$, and $f_t(\emptyset) \equiv 0$ [20]. Lastly, we will make use of two important properties of the Lovász extension: (i) $\hat{f}_t$ is convex if and only if f_t is submodular and (ii) $\hat{f}_t(\boldsymbol{\chi}_A) = f_t(A)$ for submodular functions.

A subgradient of $\hat{f}_t$ at a point $\mathbf{x} \in \operatorname{conv} \mathcal{S}$ can be computed using only evaluations of the original, submodular function f_t. Let $\pi : [0,1]^n \times V \mapsto V$ be a function where $\pi(\mathbf{x}, i) = j$ is such that the i^{th}

largest component of $\mathbf{x}$ is x_j. Let $\partial \hat{f}_t(\mathbf{x})$ be the subgradient set of $\hat{f}$ at $\mathbf{x}$. Then, we have the following definition for $\mathbf{g}_t \in \partial \hat{f}_t(\mathbf{x})$:

$$\mathbf{g}_t = \sum_{i=1}^{n} \left(f_t(\{1,2,\ldots,i\}) - f_t(\{1,2,\ldots,i-1\})\right) \boldsymbol{\chi}_{\{\pi(\mathbf{x},i)\}}. \tag{2}$$

Finally, a rounding algorithm can be employed to convert the Lovaśz extension's continuous input $\mathbf{x} \in [0,1]^n$ to the corresponding set of the original, set function f_t. For example, in Section 5 we will use $\text{rounding}_{2^V}(\mathbf{x}) : [0,1]^n \mapsto 2^V$, a standard rounding map for unconstrained problems defined as follows. Let $\mathbf{x} \in [0,1]^n$, then $\text{rounding}_{2^V}(\mathbf{x}) = S$ where $S = \{i \in V \,|\, x_i \geq p\}$ with $p \sim \text{Uniform}[0,1]$. Note that we get $\hat{f}_t(\mathbf{x}) = \mathbb{E}\left[f_t(S)\right]$ [2]. Alternatively, for some types of feasible sets $\mathcal{S}$ [21, 22], rounding algorithms can be characterized by their approximation guarantee [12]. For example, a rounding technique $\text{rounding}_{\mathcal{S}}$ with approximation guarantee α is such that $\alpha \hat{f}_t(\mathbf{x}_t) \geq f_t(S_t) = \hat{f}_t(\boldsymbol{\chi}_{S_t})$ for $S_t = \text{rounding}_{\mathcal{S}}(\mathbf{x}_t)$, where $S_t \in \mathcal{S}$ and $\mathbf{x}_t \in \text{conv}\,\mathcal{S}$.

3 Greedy approaches

We now propose greedy approaches for online binary optimization. These approaches are based on the previous round's objective function and an approximation that renders the submodular problem tractable. We first consider the following function approximation.

Definition 2 (β-approximation function [12]). The function $\tilde{f}_t : 2^V \mapsto \mathbb{R}$ is a β-approximation of f_t if it satisfies the following conditions:

1. $f_t(S) \leq \tilde{f}(S) \leq \beta f_t(S)$ for $\beta \geq 1$ and all $S \subseteq V$;
2. $\min_{S \subseteq \mathcal{S}} \tilde{f}(S)$ can be solved to optimality in polynomial time.

Based on Definition 2, (1) can be solved in an online fashion using the following update:

$$S_t \in \operatorname*{arg\,min}_{S \in \mathcal{S}} \tilde{f}_{t-1}(S), \tag{3}$$

where $\tilde{f}$ is a β-approximation function of f. We refer to an algorithm implementing (3) as online submodular greedy algorithm (`OSGA`).

For the next results, we make the following two assumptions.

Assumption 1. Let f_t be a bounded function over the set $\mathcal{S}$, i.e., there exists $M \in \mathbb{R}_{>0}$ such that $|f_t(S)| \leq M$ for all $t = 1, 2, \ldots, T$ and $S \in \mathcal{S}$.

Assumption 2. The set value function f_t is such that

$$|f_t(S_1) - f_t(S_2)| \leq L\,\text{card}\left(S_1 \ominus S_2\right),$$

for all $S_1, S_2 \subseteq \mathcal{V}$ and $0 < L < +\infty$.

In other words, we assume a Lipschitz continuity-like property for set functions. Assumption 2 holds for any submodular function if Assumption 1 does, e.g., the generic approximation defined below [23] or the β-approximation function for minimum spanning tree with submodular loss function $h(S)$, $\tilde{h}(S) = \sum_{i \in S} h(i)$ [24].

For the regret analysis, we let $\tilde{S}_t^\star \in \arg\min_{S \in \mathcal{S}} \tilde{f}_t(S)$ where $\tilde{f}_t$ is a β-approximation of f_t. We redefine the cumulative variation of the optima as $\tilde{V}_T = \sum_{t=2}^{T} \left\| \chi_{\tilde{S}_t^\star} - \chi_{\tilde{S}_{t-1}^\star} \right\|_2$. This definition is similar to the one used in standard dynamic online convex optimization [1, 19] and has the advantage of being a function of efficiently obtainable optima. The regret analysis of update (3) is provided in Theorem 1.

Theorem 1. Suppose $\tilde{f}_t$ is a β-approximation of f_t such that $\tilde{f}_t$ satisfies Assumptions 1 and 2. If $\alpha \geq \beta$, then the α-regret of `OSGA` is bounded by:

$$\mathtt{R}^{\mathrm{d}}_{\alpha}(T) \leq \frac{\alpha L}{\beta} \sum_{t=2}^{T} \sqrt{\mathrm{card}\left(\tilde{S}^{\star}_t \ominus \tilde{S}^{\star}_{t-1}\right)} = \frac{\alpha L}{\beta} \tilde{V}_T.$$

If $\tilde{V}_T$ is sublinear, then so is the α-regret.

Proof. We bound the α-regret using Definition 2 to obtain

$$\mathtt{R}^{\mathrm{d}}_{\alpha}(T) \leq \sum_{t=1}^{T} \tilde{f}_t(S_t) - \frac{\alpha}{\beta} \tilde{f}_t(S^{\star}_t). \tag{4}$$

We observe that $\tilde{f}_t(S_t) = \tilde{f}_t(\tilde{S}^{\star}_{t-1})$ because of (3) and $\tilde{S}^{\star}_t \in \arg\min_{S \in \mathcal{S}} \tilde{f}_t(S)$. Thus, we can rewrite (4) as

$$\begin{aligned} \mathtt{R}^{\mathrm{d}}_{\alpha}(T) &\leq \sum_{t=1}^{T} \tilde{f}_t(\tilde{S}^{\star}_{t-1}) - \frac{\alpha}{\beta} \tilde{f}_t(S^{\star}_t) \\ &\leq \sum_{t=1}^{T} \tilde{f}_t(\tilde{S}^{\star}_{t-1}) - \frac{\alpha}{\beta} \tilde{f}_t(\tilde{S}^{\star}_t), \end{aligned} \tag{5}$$

where the last inequality follows from the definition of $\tilde{S}^{\star}_t$. By assumption, $\alpha \geq \beta$ and $\tilde{f}$ satisfies Assumption 2. We can thus rewrite (5) as

$$\begin{aligned} \mathtt{R}^{\mathrm{d}}_{\alpha}(T) &\leq \frac{\alpha}{\beta} \sum_{t=1}^{T} \tilde{f}_t(\tilde{S}_t) - \frac{\alpha}{\beta} \tilde{f}_t(\tilde{S}^{\star}_t) \\ &\leq \frac{\alpha L}{\beta} \sum_{t=1}^{T} \sqrt{\mathrm{card}\left(\tilde{S}^{\star}_t \ominus \tilde{S}^{\star}_{t-1}\right)}, \end{aligned}$$

and we have completed the proof. □

We remark that contrarily to [12, Theorem 2], the approximation factor α does not need to be known to run the algorithm. Given a β-approximation of f_t, `OSGA` leads to a $O(\tilde{V}_T)$ regret bound similarly to the tightest dynamic bound in standard OCO [4]. Note that this latter work requires strong convexity. Theorem 1 bound's also improves on [11]'s expected bound because it is only a function of the cumulative variation and holds asymptotically.

Recall that unconstrained submodular minimization problem can be solved to optimality efficiently. Hence, for the special case where $\mathcal{S} = 2^V$, i.e., when (1) is an unconstrained binary problem, `OSGA` can be directly applied to the previous round loss function. This application leads to the following regret bound.

Corollary 1. If $\mathcal{S} = 2^V$ and $\alpha = \beta = 1$, then `OSGA`'s update reduces to

$$S_t \in \arg\min_{S \in 2^V} f_{t-1}(S), \tag{6}$$

and leads to:

$$\mathtt{R}^{\mathrm{d}}(T) \leq L \sum_{t=2}^{T} \sqrt{\mathrm{card}\left(S^{\star}_t \ominus S^{\star}_{t-1}\right)} < O\left(V_T\right).$$

Proof. The proof follows from Theorem 1 where (i) the regret is considered instead of the α-regret and (ii) f_t is used directly instead of $\tilde{f}_t$, the β-approximation because (6) can be solved efficiently. □

For some problem instances, finding an approximation that satisfies both Definition 2 and Assumption 2 is difficult. Alternatively, the generic approximation for submodular functions provided in Definition 3 is considered [12, 23].

Definition 3 (Generic approximation function [12, 23]). The function $\tilde{f}^{\mathrm{g}} : 2^V \mapsto \mathbb{R}$ is a generic approximation of f defined as: $\tilde{f}^{\mathrm{g}}_t(S) = \sqrt{\sum_{i \in S} c_i}$, for some $\mathbf{c} \in \mathbb{R}^n$, and satisfies

$$\left(\tilde{f}^{\mathrm{g}}_t(S)\right)^2 \leq (f_t(S))^2 \leq \gamma^2 \left(\tilde{f}^{\mathrm{g}}_t(S)\right)^2,$$

for all $S \subseteq 2^V$ and some $\gamma > 0$.

Interested readers are referred to [23] for details about the constant $\mathbf{c}$. We now consider the online submodular generic greedy algorithm (OSGGA), i.e., the generic approximation-based OSGA. OSGGA uses the following update:

$$S_t \in \arg\min_{S \in \mathcal{S}} \left(\tilde{f}^{\mathrm{g}}_{t-1}(S)\right)^2, \tag{7}$$

The update rule (7) is equivalent to solving a linear program and can, therefore, be solved efficiently [12].

For the next result, we utilize the variation term $\tilde{V}^{\mathrm{g}}_T = \sum_{t=2}^{T} \sqrt{\mathrm{card}\left(S^{\mathrm{g},\star}_t \ominus S^{\mathrm{g},\star}_{t-1}\right)}$. Similarly to OSGA, $\tilde{V}^{\mathrm{g}}_T$ is based only on the optima of tractable problems. Let $\nu \geq \min_{t, S \in \mathcal{S}} f_t(S) > 0$, be lower bound on all round minima. Lastly, we remark that the squared generic approximation function satisfies Assumption 2 with modulus $\tilde{L}^{\mathrm{g}}$ because it is linear and bounded by Assumption 1. The α-regret for the OSGGA is presented below.

Corollary 2. Suppose $\tilde{f}^{\mathrm{g}}_t$ is a generic approximation of f_t. Then the α-regret of OSGGA is bounded above by

$$\mathtt{R}^{\mathrm{d}}_\alpha(T) \leq \frac{4\alpha^2 \tilde{L}^{\mathrm{g}} L}{(1+\alpha)\nu} \sum_{t=2}^{T} \sqrt{\mathrm{card}\left(S^{\mathrm{g},\star}_t \ominus S^{\mathrm{g},\star}_{t-1}\right)} = \frac{4\alpha^2 \tilde{L}^{\mathrm{g}} L}{(1+\alpha)\nu} \tilde{V}^{\mathrm{g}}_T,$$

and is sublinear for $\tilde{V}^{\mathrm{g}}_T < O(T)$.

Proof. We based our proof on [12, Theorem 2 and Lemma 3]. The α-regret for update (7) is

$$\begin{aligned}
\mathtt{R}^{\mathrm{d}}_\alpha(T) &= \sum_{t=1}^{T} f_t(S_t) - \alpha f_t(S^\star_t) \\
&= \sum_{t=1}^{T} \frac{(f_t(S_t) - \alpha f_t(S^\star_t))(f_t(S_t) + \alpha f_t(S^\star_t))}{(f_t(S_t) + \alpha f_t(S^\star_t))} \\
&= \sum_{t=1}^{T} \frac{(f_t(S_t))^2 - \alpha^2 (f_t(S^\star_t))^2}{(f_t(S_t) + \alpha f_t(S^\star_t))} \\
&\leq \sum_{t=1}^{T} \frac{(f_t(S_t))^2 - \alpha^2 (f_t(S^\star_t))^2}{(1+\alpha)\nu},
\end{aligned}$$

where $\nu \geq \min_{t, S_t} f_t(S_t) > 0$. By Definition 3, we then have

$$\begin{aligned}
\mathtt{R}^{\mathrm{d}}_\alpha(T) &\leq \frac{1}{(1+\alpha)\nu} \sum_{t=1}^{T} \gamma^2 \left(\tilde{f}^{\mathrm{g}}_t(S_t)\right)^2 - \alpha^2 \left(\tilde{f}^{\mathrm{g}}_t(S^\star_t)\right)^2 \\
&\leq \frac{1}{(1+\alpha)\nu} \sum_{t=1}^{T} \gamma^2 \left(\tilde{f}^{\mathrm{g}}_t(S_t)\right)^2 - \alpha^2 \left(\tilde{f}^{\mathrm{g}}_t(\tilde{S}^{\mathrm{g},\star}_t)\right)^2,
\end{aligned}$$

where $\tilde{S}_t^{\mathrm{g},\star} \in \arg\min_{S\in\mathcal{S}} \tilde{f}_t^{\mathrm{g}}(S)$. Using the update rule (3), we obtain

$$\mathtt{R}_\alpha^d(T) \le \frac{\gamma^2}{(1+\alpha)\nu} \sum_{t=1}^{T} \left(\tilde{f}_t^{\mathrm{g}}(\tilde{S}_{t-1}^{\mathrm{g},\star})\right)^2 - \left(\tilde{f}_t^{g}(\tilde{S}_t^{\mathrm{g},\star})\right)^2. \tag{8}$$

Thus, the α-regret can be re-expressed as

$$\mathtt{R}_\alpha^{\mathrm{d}}(T) \le \frac{\gamma^2}{(1+\alpha)\nu} \mathtt{R}^d \left(\left(\tilde{f}_t^{\mathrm{g}}(S)\right)^2, T \right),$$

where $\mathtt{R}^{\mathrm{d}} \left(\left(\tilde{f}_t^{g}(S)\right)^2, T \right)$ is the (1-)regret of update (7) when used on the problem $\min_{S\in\mathcal{S}} \left(\tilde{f}_t^{g}(S)\right)^2$.

Using Theorem 1 with $\alpha = \beta = 1$ in (8) yields

$$\mathtt{R}_\alpha^{\mathrm{d}}(T) \le \frac{\gamma^2 \tilde{L}^{\mathrm{g}} L}{\nu} \sum_{t=1}^{T} \sqrt{\operatorname{card}\left(\tilde{S}_{t-1}^{\mathrm{g},\star} \ominus \tilde{S}_t^{\mathrm{g},\star}\right)} = \frac{\gamma^2 \tilde{L}^{\mathrm{g}} L}{\nu} \tilde{V}_T^{\mathrm{g}},$$

which completes the proof. □

Hence, `OSGGA` leads to an α-regret bound that is similar to Theorem 1's. Albeit different variation terms are used, the bounds (i) only differ within a constant factor and (ii) both admit up to a linear variation term to be sublinear.

4 Projected gradient descent-based approach

In this section, we consider a convex optimization-based update to solve (1) [2, 12]. Our approach leverages the Lovász extension's convexity for submodular functions. We propose the online submodular projected gradient descent (`OSPGD`) based on the update defined as:

$$\mathbf{x}_{t+1} = \operatorname{proj}_{\operatorname{conv}(\mathcal{S})} \mathbf{x}_t - \eta \mathbf{g}_t \tag{9}$$

$$S_{t+1} = \operatorname{rounding}_{\mathcal{S}}(\mathbf{x}_{t+1}), \tag{10}$$

where $\mathbf{g}_t \in \partial \hat{f}_t(\mathbf{x}_t)$ is defined in (2) and $\operatorname{proj}_{\operatorname{conv}(\mathcal{S})}$ is the projection onto the convex hull of $\mathcal{S}$. `OSPGD` has the advantage over the greedy updates to be computationally very simple because only a single gradient descent step is performed. It however requires a rounding approach which might not be available for all constrained problems. For `OSPGD`, we obtain the following regret bound.

Theorem 2. Suppose that a rounding algorithm with approximation guarantee α is used. Then, `OSPGD` with $\eta = \frac{\delta}{\sqrt{T}}$ leads to an α-regret bounded from above by:

$$\begin{aligned}\mathtt{R}_\alpha^{\mathrm{d}}(T) &\le \alpha \left(\sqrt{n}\delta \sum_{t=2}^{T} \sqrt{\operatorname{card}\left(S_t^\star - S_{t-1}^\star\right)} + \frac{5n}{2\delta} + 4M\delta \right) \sqrt{T} \\ &= \alpha \left(\sqrt{n}\delta V_T + \frac{5n}{2\delta} + 4M\delta \right) \sqrt{T},\end{aligned}$$

and is sublinear if $V_T < O\left(\sqrt{T}\right)$.

Proof. By definition, we have

$$\mathtt{R}_\alpha^{\mathrm{d}}(T) = \sum_{t=1}^{T} f_t(S_t) - \alpha f_t(S_t^\star)$$

$$\leq \sum_{t=1}^{T} \alpha \hat{f}_t(x_t) - \alpha f_t(S_t^\star),$$

using the rounding algorithm approximation guarantee bound. Using the property of the Lovaśz extension, we obtain

$$\mathtt{R}_\alpha^{\mathrm{d}}(T) \leq \alpha \sum_{t=1}^{T} \hat{f}_t(x_t) - \hat{f}_t(\boldsymbol{\chi}_{S_t^\star}). \tag{11}$$

We then follow the standard proof techniques for the online gradient descent (OGD) from [1, 13]. We have

$$\begin{aligned}
\left\|\mathbf{x}_{t+1} - \boldsymbol{\chi}_{S_t^\star}\right\|_2^2 &= \left\|\left(\mathrm{proj}_{\mathrm{conv}(\mathcal{S})}\, \mathbf{x}_t - \eta \mathbf{g}_t\right) - \boldsymbol{\chi}_{S_t^\star}\right\|_2^2 \\
&\leq \left\|\mathbf{x}_t - \eta \mathbf{g}_t - \boldsymbol{\chi}_{S_t^\star}\right\|_2^2 \\
&= \left\|\mathbf{x}_t - \boldsymbol{\chi}_{S_t^\star}\right\|_2^2 - 2\eta_t \mathbf{g}_t^\top \left(\mathbf{x}_t - \boldsymbol{\chi}_{S_t^\star}\right) \\
&\quad + \eta^2 \left\|\mathbf{g}_t\right\|_2^2 \\
\Leftrightarrow \mathbf{g}_t^\top \left(\mathbf{x}_t - \boldsymbol{\chi}_{S_t^\star}\right) &\leq \frac{1}{2\eta}\left(\left\|\mathbf{x}_t - \boldsymbol{\chi}_{S_t^\star}\right\|_2^2 - \left\|\mathbf{x}_{t+1} - \boldsymbol{\chi}_{S_t^\star}\right\|_2^2\right) \\
&\quad + \frac{\eta^2}{2} \left\|\mathbf{g}_t\right\|_2^2 .
\end{aligned} \tag{12}$$

The convexity of $\hat{f}_t$ implies that for all $\mathbf{x}, \mathbf{y} \in [0,1]^n$, we have

$$\hat{f}_t(\mathbf{x}) \geq \hat{f}_t(\mathbf{y}) + \mathbf{g}_t^\top (\mathbf{x} - \mathbf{y}),$$

for $\mathbf{g}_t \in \partial \hat{f}_t(\mathbf{y})$. Using $\mathbf{x} = \boldsymbol{\chi}_{S_t^\star}$ and $\mathbf{y} = \mathbf{x}_t$, we obtain

$$\hat{f}_t(\mathbf{x}_t) - \hat{f}_t(\boldsymbol{\chi}_{S_t^\star}) \leq \mathbf{g}_t^\top (\mathbf{x} - \mathbf{y}). \tag{13}$$

Substituting (12) and (13) in (11) leads to

$$\begin{aligned}
\mathtt{R}_\alpha^{\mathrm{d}}(T) \leq& \alpha \sum_{t=1}^{T} \frac{1}{2\eta}\left(\left\|\mathbf{x}_t - \boldsymbol{\chi}_{S_t^\star}\right\|_2^2 - \left\|\mathbf{x}_{t+1} - \boldsymbol{\chi}_{S_t^\star}\right\|_2^2\right) \\
&+ \alpha \sum_{t=1}^{T} \frac{\eta^2}{2} \left\|\mathbf{g}_t\right\|_2^2 \\
=& \alpha \sum_{t=1}^{T} \frac{1}{2\eta}\left(\left\|\mathbf{x}_t\right\|_2^2 - \left\|\mathbf{x}_{t+1}\right\|_2^2\right) + \alpha \sum_{t=1}^{T} \frac{\eta}{2} \left\|\mathbf{g}_t\right\|_2^2 \\
&+ \alpha \sum_{t=2}^{T} \frac{1}{\eta} \mathbf{x}_t^\top \left(\boldsymbol{\chi}_{S_t^\star} - \boldsymbol{\chi}_{S_{t+1}^\star}\right) - \mathbf{x}_1^\top \boldsymbol{\chi}_{S_1^\star} + \mathbf{x}_{T+1}^\top \boldsymbol{\chi}_{S_T^\star} \\
=& \frac{\alpha}{2\eta}\left\|\mathbf{x}_1\right\|_2^2 - \frac{\alpha}{2\eta}\left\|\mathbf{x}_{T+1}\right\|_2^2 - \frac{\alpha}{\eta}\mathbf{x}_1^\top \boldsymbol{\chi}_{S_1^\star} + \frac{\alpha}{\eta}\mathbf{x}_{T+1}^\top \boldsymbol{\chi}_{S_T^\star} \\
&+ \frac{\alpha}{\eta} \sum_{t=2} \mathbf{x}_t^\top \left(\boldsymbol{\chi}_{S_t^\star} - \boldsymbol{\chi}_{S_{t-1}^\star}\right) + \alpha \sum_{t=1}^{T} \frac{\eta}{2} \left\|\mathbf{g}_t\right\|_2^2 ,
\end{aligned}$$

where we have evaluated the telescoping sums to obtain the last line [13]. Using [12, Lemma 1], we have $\left\|\mathbf{g}_t\right\|_2 \leq 4M$. The regret becomes

$$\mathtt{R}_\alpha^{\mathrm{d}}(T) \leq \frac{5n\alpha}{2\eta} + 4\alpha M \eta T + \frac{\alpha\sqrt{n}}{\eta} \sum_{t=2}^{T} \left\|\boldsymbol{\chi}_{S_t^\star} - \boldsymbol{\chi}_{S_{t-1}^\star}\right\|_2 ,$$

where we also used the fact that $\|\mathbf{x}_t\|_2 \leq \sqrt{n}$. We remark that for a submodular function and its Lovaśz extension pair, $\left\|\boldsymbol{\chi}_{S_t^\star} - \boldsymbol{\chi}_{S_{t-1}^\star}\right\|_2$ is equivalent to $\sqrt{\operatorname{card} S_t^\star \ominus S_{t-1}^\star}$. We now have

$$\mathtt{R}_\alpha^{\mathrm{d}}(T) \leq \frac{5n\alpha}{2\eta} + 4\alpha M\eta T + \frac{\alpha\sqrt{n}}{\eta}\sum_{t=2}^{T}\sqrt{\operatorname{card}\left(S_t^\star \ominus S_{t-1}^\star\right)}.$$

Setting $\eta = \frac{\delta}{\sqrt{T}}$, $\delta \in \mathbb{R}_{>0}$ completes the proof. □

In sum, we obtain an α-regret bound that is of the same order as the standard online gradient descent for OCO problems [1], i.e., $O\left(\sqrt{T}\left(1+V_T\right)\right)$. Hence, in comparison to Section 3's approaches, we have traded higher algorithmic simplicity for a stricter regret bound.

Lastly, if a randomized rounding technique is used to convert a continuous decision vector to a binary one, expected and high-probability regret bounds, i.e., where $\alpha = 1$ as opposed to previous results, can be derived.

Corollary 3. Consider a random rounding technique $\operatorname{rounding}_{\mathcal{S}} : \operatorname{conv}(\mathcal{S}) \mapsto \mathcal{S}$ such that for $S = \operatorname{rounding}_{\mathcal{S}}(\mathbf{x})$ we have $\mathbb{E}\left[f_t(S)\right] = \hat{f}_t(\mathbf{x})$. The expected and high-probability dynamic regret for `OSPGD` with $\frac{\delta}{\sqrt{T}}$ are bounded from above:

$$\mathbb{E}\left[\mathtt{R}^{\mathrm{d}}(T)\right] \leq \sqrt{nT}\delta V_T + \left(\frac{5n}{2\delta} + 4M\delta\right)\sqrt{T},$$

and

$$\mathtt{R}^{\mathrm{d}}(T) \leq \sqrt{nT}\delta V_T + \left(\frac{5n}{2\delta} + 4M\delta + 2M\delta\log\frac{1}{\epsilon}\right)\sqrt{T},$$

with probability of at least $1-\epsilon$.

Proof. We adapt Theorem 2's and [2, Theorem 1]'s proofs. First, for the expected bound, we have

$$\begin{aligned}\mathbb{E}\left[\mathtt{R}^{\mathrm{d}}(T)\right] &= \sum_{t=1}^{T}\mathbb{E}\left[f_t(S_t)\right] - \mathbb{E}\left[f_t(S_t^\star)\right] \\ &= \sum_{t=1}^{T}\hat{f}_t(\mathbf{x}_t) - f_t(S_t^\star) \\ &= \sum_{t=1}^{T}\hat{f}_t(\mathbf{x}_t) - \hat{f}_t(\boldsymbol{\chi}_{S_t^\star}).\end{aligned} \tag{14}$$

The bound then follows from Theorem 2. Second, for the high probability bound, we use Hœffding inequality [2, Theorem 13]. With a probability of a least $1-\epsilon$, we have

$$\sum_{t=1}^{T} f_t(S_t) \leq \sum_{t=1}^{T}\mathbb{E}[f_t(S_t)] + M\sqrt{2T\log\frac{1}{\epsilon}}. \tag{15}$$

Substituting (15) in the regret definition, we obtain

$$\mathtt{R}^{\mathrm{d}} \leq \sum_{t=1}^{T}\left(\mathbb{E}[f_t(S_t)] - f_t(S_t^\star)\right) + M\sqrt{2T\log\frac{1}{\epsilon}}. \tag{16}$$

We observe that the first term of (16)'s right-hand side and (14)'s are identical. Using Theorem 2 with $\eta = \frac{\delta}{\sqrt{T}}$ in (16) yields the high probability regret bound. □

For unconstrained problems, we have $\mathcal{S} = 2^V$ and $\operatorname{conv}\mathcal{S} = [0,1]^n$. Then, Corollary 3 holds when the randomized rounding procedure described in Section 2.2 is implemented in `OSPGD`.

5 Applications to electric power systems

We now apply `OSPGD` and `OSGA` to power system problems.

5.1 Demand response for frequency regulation

In demand response, a load aggregator is contracted by the system operator [25–27]. The aggregator's mandate is to modulate the load power consumption to help out the grid, e.g., to mitigate renewable intermittency or reduce peak demand. Specifically, we consider frequency regulation services [10, 28–30], i.e., load balancing on a fast timescale, e.g., 4 seconds. Advantages of demand response over other frequency regulation approaches, like battery energy storage and fast-ramping fuel-burning generation, include low deployment costs and sustainability [30].

5.1.1 Setting

Consider N thermostatically controlled loads (TCLs), e.g., residential loads equipped with electric water heaters, heaters, or air conditioners, enrolled in the demand response program. Consider a program of duration T in which decision rounds are indexed by t. Let $p_{n,t} \geq 0$ and $\tilde{p}_{n,t} \geq 0$ be the power consumption of TCL $n \in \{1, 2, \dots, N\}$ when the load is flexible and inflexible, respectively. This formulation is similar to [11]'s. Each load must stay in an acceptable temperature range, e.g., $\pm 0.5°$C of the desired user temperature, to be flexible, i.e., to be controlled according to the aggregator's need. If the load temperature is too high or too low, the backup controller forces the load to be active or inactive accordingly, and its power consumption must be accounted for.

At time t, the aggregator's objective is to track a regulation setpoint r_t provided by the system operator by adjusting the TCL power consumption. In this work, we consider a setting in which the aggregator wants to deploy the minimum number of flexible loads such that the regulation signal is met. This problem can be formulated as an online dynamic submodular optimization problem using the objective function $f_t^{\mathrm{DR}} : 2^V \mapsto \mathbb{R}$,

$$f_t^{\mathrm{DR}}(S_t) = \sum_{A \subseteq 2^V} \left[\left(\sum_{n \in A} u_{n,t} \right)^2 - \left(\sum_{n \in V} u_{n,t} \right)^2 \right] \cdot \max\{0, |S_t \cap A| - |S_t \cup A| + 1\} \mathbb{I}_{A \subseteq \mathcal{R}_t}, \tag{17}$$

where $u_{n,t} = p_{n,t} + \tilde{p}_{n,t}$, $\mathbb{I}$ is the indicator function which returns 1 if the subscript is true and 0 otherwise, and

$$\mathcal{R}_t = \left\{ S \subseteq 2^V \,\middle|\, \sum_{n \in S} u_{n,t} \geq r_t \right\}.$$

In (17), the term between brackets promotes partitions $A \subseteq 2^V$ with lower aggregated power. Then, the maximum term identifies to which partition A the set S belongs to, because it is equal to one if and only if $S = A$. Lastly, the indicator function ensures that the set of dispatched loads is at least equal to the regulation signal.

We apply `OSPGD` to this problem. In terms of standard online optimization, this corresponds to a quadratic program with time-dependent binary constraints, which, to this day, has not been investigated. To the author's best knowledge, no other approach has been shown to have provable performance in this context.

We randomly generate loads' parameters similarly to [31]. We use the same constraints and logical rules as [11] and omit the lockout constraint. The thermodynamic model is based on [29]. We consider TCLs equipped with air conditioners.

5.1.2 Numerical results

We deploy 15 TCLs to track a vanishing sinusoidal regulation signal subject to Perlin noise [32]. We compare `OSPGD` to the closest work to ours, `bOGD` [11]. We note that, in this setting, `bOGD`'s regret analysis does not hold. Lastly, we provide the round optimum, which we denote $\texttt{OSPGD}_t^\star$.

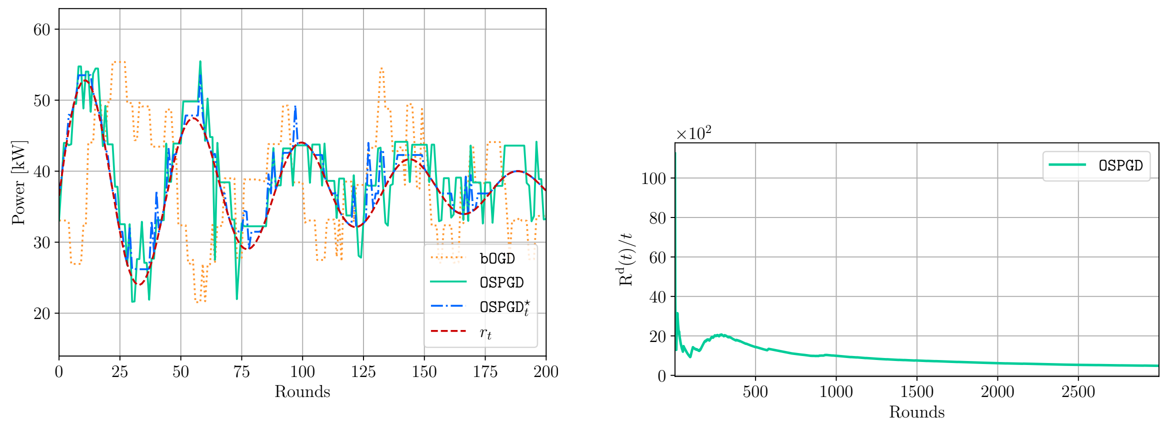

(a) Setpoint tracking **(b) Time-averaged experimental dynamic regret for `OSPGD`**

Figure 1: Demand response with 15 Loads

Figure (1a) presents `OSPGD`'s time-averaged dynamic α-regret. The vanishing time-averaged regret implies a sublinear regret. The tracking root-mean-square error (RMSE) over 3000 rounds is 1.406 kW for $\texttt{OSPGD}_t^\star$, 3.227 kW for `OSPGD`, and 6.928 kW for `bOGD`. As shown in Figure (1b), `OSPGD` outperforms `bOGD` and offers good setpoint tracking.

5.2 Real-Time Network Reconfiguration

Electric distribution networks generally possess a radial topology [33]. Their topology is controlled via switches located throughout the network. By opening and closing different switches, the topology can be modified, for example, to minimize active power losses or line congestion, and, thus, to increase the grid efficiency [34–36]. The set of switch statuses must always induce a radial network topology while assuring that all loads are supplied.

Distribution grids with high penetration of grid-edge/behind-the-meter technologies [37, 38], e.g., electric vehicles, residential solar panels, or demand response, can experience large, fast-ramping variations in power demand at the different buses. These rapid changes in loading are out of the distribution system operator's control and can lead to network perturbations, e.g., over/under-voltage, line congestion, etc. [39–41]. To mitigate incidents, the system operator can preemptively configure the distribution network by altering its topology. Remotely-activated switches allow fast network reconfiguration (NR) and can be used to adapt to the load demand in real-time, viz., to prevent line congestion or higher power losses.

5.2.1 Setting

We consider a distribution network consisting of a set of static powerlines $\overline{\mathcal{L}}$, a set of loads $\mathcal{N}$, and a set of lines equipped with switches $V = \{1, 2, \ldots, \overline{S}\}$, $\overline{S} \in \mathbb{N}$. Let $\mathcal{L}(S_t) = \overline{\mathcal{L}} \cup S_t$ where $S_t \subseteq 2^V$ be the set of all powerlines active at time t, i.e., the static line set augmented by the lines with closed switches S_t. The set $\mathcal{L}(S_t)$ is subject to two constraints; it must be such that (i) the network topology is radial and (ii) all loads are connected.

Let $p_{i,t} \geq 0$ and $q_{i,t} \geq 0$, be the active and reactive power demand, respectively, at bus $i \in \mathcal{N}$ and time t. Let $\mathcal{N}_r$ be the set of feeder nodes. Let $P_{ij,t} \in \mathbb{R}$ and $Q_{ij,t} \in \mathbb{R}$ be, respectively, the active and reactive power flowing from node i to j if $ij \in \mathcal{L}(S_t)$. Let $P_{ij,t} = Q_{ij,t} = 0$ if $ij \notin \mathcal{L}(S_t)$. Line ij's apparent power is denoted by $A_{ij,t} = P_{ij,t} + jQ_{ij,t}$. Let $v_i \in \mathbb{C}$ be the voltage at node i, $I_{ij} \in \mathbb{C}$ be the current flowing in line ij, and $y_{ij} \in \mathbb{C}$ be the admittance of line ij. Let notation $\overline{x}$ and $\underline{x}$ represent upper and lower bounds on any given parameter x.

To minimize active power losses in distribution grids, the NR problem can be cast using the objective function $f_t^{\mathrm{NR}} : 2^V \mapsto \mathbb{R}$ presented in (18) where power losses on line ij at time t are defined as $|v_{i,t} - v_{j,t}|^2 y_{ij}^*$. In (18), the spanning tree constraint ensures that the network topology is radial and connects all loads to the source node. The other constraint ensure that the power flow (PF), which models the electric network's physics, respects all operational constraints while meeting power demand.

$$
\begin{aligned}
\min_{S_t \subseteq 2^V} \quad & f_t^{\mathrm{NR}}(S) = \sum_{ij \in \mathcal{L}(S)} |v_{i,t} - v_{j,t}|^2 y_{ij}^* \\
\text{s.t.} \quad & \mathcal{L}(S_t) \subseteq \mathrm{SpanningTree}(\mathcal{N}) \\
& \left\{ \begin{matrix} \{v_{i,t}\}_{i \in \mathcal{N}} \\ \{P_{ij,t}, Q_{ij,t}\}_{ij \in \mathcal{L}(S_t)} \end{matrix} \right\} \in \mathrm{PF}(\{p_{i,t}, q_{i,t}\}_{i \in \mathcal{N}}, \mathcal{L}(S_t))
\end{aligned} \tag{18}
$$

5.2.2 Weakly-Meshed Approximation

Finding the optimal configuration of a radial network is NP-hard. We re-express (18) as an online dynamic submodular optimization problem, which can then be solved in real-time.

When the radiality constraint is relaxed, the network, in which the set of active powerlines is $\overline{\mathcal{L}} \cup V$, referred to as the weakly-meshed network (WMN), is a good solution, if not optimal, for loss minimization [42]. Using the WMN as a starting point, our goal is to find the radial network that best imitates its power flow. This can be done by first computing the WMN power flow. Then, a minimum spanning tree (MST) algorithm (e.g., Prim's algorithm [43]) with edge weights set as the negative line currents $-I_{ij}$ obtained from the power flow, is used. The MST is fast and guarantees radiality. By removing the edges with lower currents, the MST returns a radial network with a power flow pattern similar to the WMN as demonstrated by [42]. We note that in all evaluations, the resulting topology admitted a feasible power flow with respect to the original AC power flow constraints. If infeasible, the resulting topology could be projected onto the set induced by these constraints. Finally, we can approximate (18) by the following online dynamic submodular problem:

$$
\begin{aligned}
\min_{S_t \subseteq 2^V} \quad & f_t^{\mathrm{WM}}(S_t) = \sum_{ij \in \mathcal{L}(S_t) \cup V} -I_{ij,t} \\
\text{s.t.} \quad & \mathcal{L}(S_t) \subseteq \mathrm{SpanningTree}(\mathcal{N}),
\end{aligned} \tag{19}
$$

where $I_{ij,t}$ is an online parameter extracted from power flow computations, e.g. [44], of the WMN.

Because (19) is submodular, and can be solved to optimality in polynomial time using a MST algorithm like Prim's [43] over the WMN, we apply our `OSGA` for online reconfiguration. The process is summarized in Algorithm 1.

In Algorithm 1, M is a large constant. We remark that in the case of multiple feeders, we temporarily add virtual lines between the different sources (generators) in the MST algorithm to ensure radiality, see steps 5−6. These lines are then removed from S_t, see step 8.

5.2.3 Numerical Results

For this section, we consider the IEEE standardized 33-bus/1-feeder (33b/1f) [45] and the 135-bus/8-feeder (135b/8f) [46] distribution networks with the added modification, on both networks, that every line is equipped with a switch to fully benefit from the flexibility of online optimization. At each round, we add randomly generated Perlin noise [32] on $p_{i,t}$, $q_{i,t}, \forall i \in \mathcal{N}$ to model uncertainty. Figure 2 illustrates `OSGA`'s sublinear dynamic α-regret. This is depicted by the vanishing time-averaged regret.

Algorithm 1 `OSGA` **for Real-time Network Reconfiguration**

for $t = 1, 2, ..., T$ **do**
  Reconfigure the network according to S_t.
  Observe new online parameters: $p_{i,t}$, $q_{i,t}, \forall i \in \mathcal{N}$.
  Compute the power flow of the WMN with the Newton-Raphson algorithm and extract $I_{ij,t}$ $\forall ij \in \overline{\mathcal{L}} \cup V$.
  if card$(\mathcal{N}_r) > 1$ **then**
    Set $I_{ij,t} = M$, $\forall ij$ where $\{i, j \in \mathcal{N}_r \cap i \neq j\}$
  end if
  Update S_{t+1} via a greedy MST algorithm:

$$\tilde{S}_{t+1} = \arg\min \; f_t^{\text{WM}}(S) \text{ s.t. } \mathcal{L}(S) \subseteq \text{SpanningTree}(\mathcal{N})$$
$$S_{t+1} \in \tilde{S}_{t+1} \setminus \{\forall \; ij \text{ where } i, j \in \mathcal{N}_r\}$$

end for

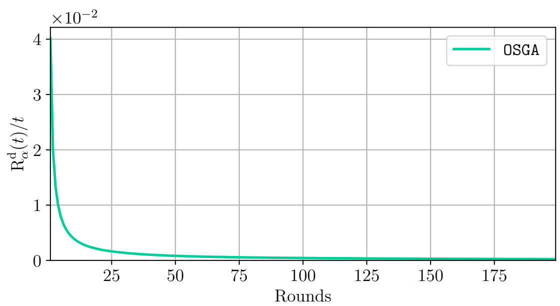

Figure 2: Time-averaged experimental dynamic regret for `OSGA`

We now compare `OSGA` to its offline counterparts solved in hindsight both dynamically ($\texttt{OSGA}_t^\star$) and statically ($\texttt{OSGA}^\star$) over the time horizon. We note that hindsight solutions only serve analysis purposes and have no practical application. We benchmark our approach, in the simpler network (33b/1f), to the closest work in OCO (`bOGD`) [11] to which we must add a projection on the feasible power flow set to handle operational constraints of the grid. We also compare `OSGA` to a round-optimal offline configuration, with a limited flexibility of 9 switches, based on the second-order cone relaxation power flow (`SOCR-9SW`) [33], which require much more computational power. Lastly, we present the case where a random feasible reconfiguration is implemented each round.

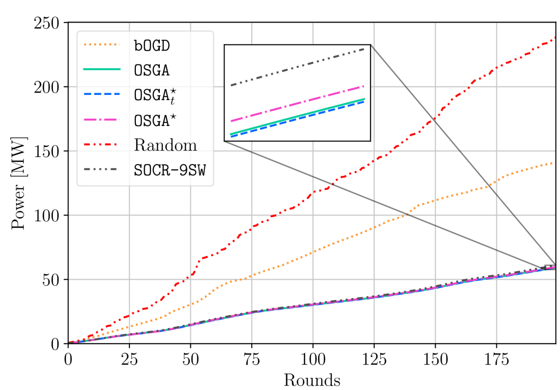

Figure 3: Cumulative power losses (33b/1f)

Figure 4 presents snapshots of the 135/8f NR at different rounds according to the apparent power demand at each node. The demand is represented by a light-dark scale: the darker the node the higher the demand is. Closed and open switches are pictured in green and red, respectively. Squares are generators. Radiality is always preserved.

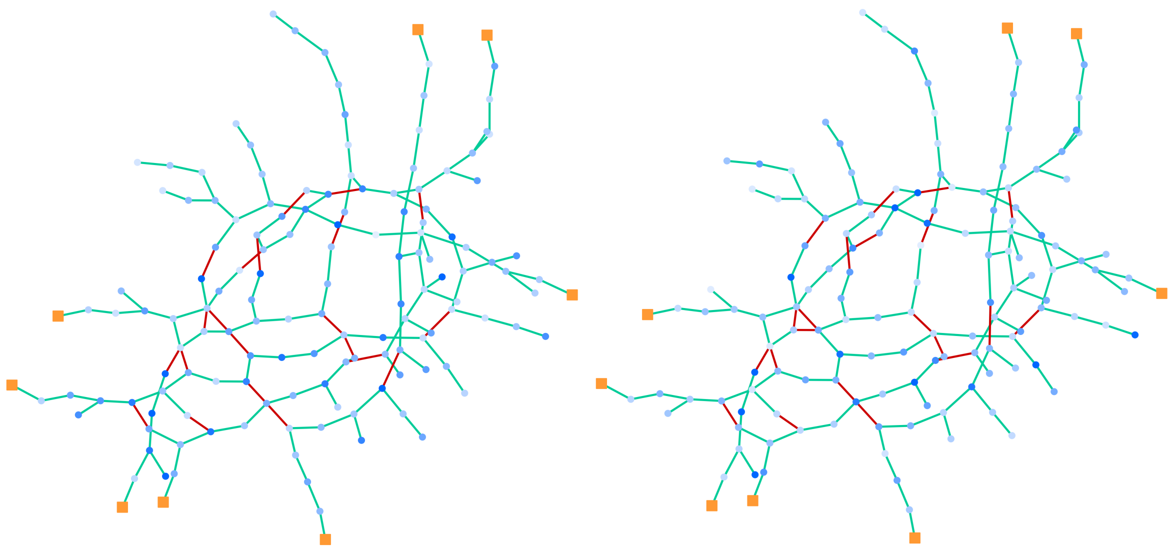

(a) Round 21

(b) Round 23

Figure 4: Network reconfiguration (135b/8f)

In sum, `OSGA` performs systematically better than `bOGD` and is considerably faster because it stands on a fast MST heuristic. It also scales easily to bigger networks and guarantees radiality without the need for a projection step. `OSGA` also outperforms $\texttt{OSGA}^\star$ while maintaining a small performance gap with $\texttt{OSGA}_t^\star$, the round optimal solution. For example, we observed a total power loss increase of 0.038% for `OSGA` and of 1.545% for $\texttt{OSGA}^\star$, after 400 rounds, when compared to $\texttt{OSGA}_t^\star$ on 135b/8f in the simulation leading to Figure 4.

6 Conclusion

In this work, we investigate online binary optimization in dynamic settings. We consider submodular objective functions and general binary constraints. We first assume an approximation of the objective function which can be minimized in polynomial time exists. We propose `OSGA` that solves the previous round approximation as a proxy and in doing so, circumvents the NP-hardness of the original problem. We adapt our approach to a generic but weaker approximation that can be used to recast general submodular problems in a simpler form. Second, aiming at algorithmic simplicity, we formulate `OSPGD` which leverages the Lovász extension and convex optimization. For all our algorithms, we provide a dynamic regret analysis. We show that `OSGA` and `OSPGD` possess, respectively, a dynamic regret bound that is similar to the tightest bound in the literature and to the `OGD` used in online convex optimization.

Finally, we present two applications of our approaches in electric power systems. First, `OSPGD` is employed to dispatch demand response resources, viz., thermostatic loads, to mitigate fast-timescale power imbalances. Second, `OSGA` is used to minimize active power losses in distribution networks via real-time reconfiguration, i.e., closing and opening switches in the network to better shape its topology. Our numerical study illustrates the performance of our approaches and their ability to outperform prior work [11] in both applications by harnessing submodularity properties. On another note, our experimental NR results show that the WMN approximation paired with Prim's algorithm is great for online decision-making, and for leveraging grid flexibility, thanks to its light computational needs, near-optimal performances, and scalability.

Next, time-varying binary constraints, i.e., constraints that, similarly to the objective function, are observed only at the end of a round while needing to be satisfied in average, will be investigated.

References

[1] M. Zinkevich, Online convex programming and generalized infinitesimal gradient ascent, in Proceedings of the 20th international conference on machine learning (icml-03), 2003, pp. 928–936.

[2] E. Hazan and S. Kale, Online submodular minimization. Journal of Machine Learning Research, 13(10), 2012.

[3] S. Shalev-Shwartz et al., Online learning and online convex optimization, Foundations and Trends® in Machine Learning, 4(2):107–194, 2012.

[4] A. Mokhtari, S. Shahrampour, A. Jadbabaie, and A. Ribeiro, Online optimization in dynamic environments: Improved regret rates for strongly convex problems, in 2016 IEEE 55th Conference on Decision and Control (CDC). IEEE, 2016, pp. 7195–7201.

[5] A. Lesage-Landry, J. A. Taylor, and I. Shames, Second-order online nonconvex optimization, IEEE Transactions on Automatic Control, 66(10):4866–4872, 2020.

[6] Y.-M. Pun and A. M.-C. So, Dynamic regret bound for moving target tracking based on online time-of-arrival measurements, in 2020 59th IEEE Conference on Decision and Control (CDC). IEEE, 2020, pp. 5968–5973.

[7] T. Chen, Q. Ling, and G. B. Giannakis, An online convex optimization approach to proactive network resource allocation, IEEE Transactions on Signal Processing, 65(24):6350–6364, 2017.

[8] X. Cao, J. Zhang, and H. V. Poor, A virtual-queue-based algorithm for constrained online convex optimization with applications to data center resource allocation, IEEE Journal of Selected Topics in Signal Processing,12(4):703–716, 2018.

[9] S.-J. Kim and G. B. Giannakis, An online convex optimization approach to real-time energy pricing for demand response, IEEE Transactions on Smart Grid, 8(6):2784–2793, 2016.

[10] A. Lesage-Landry and J. A. Taylor, Setpoint tracking with partially observed loads, IEEE Transactions on Power Systems, 33(5):5615–5627, 2018.

[11] A. Lesage-Landry, J. A. Taylor, and D. S. Callaway, Online convex optimization with binary constraints, IEEE Transactions on Automatic Control, 66(12):6164–6170, 2021.

[12] S. Jegelka and J. A. Bilmes, Online submodular minimization for combinatorial structures. in ICML. Citeseer, 2011, pp. 345–352.

[13] E. Hazan et al., Introduction to online convex optimization, Foundations and Trends® in Optimization, 2(3–4):57–325, 2016.

[14] A. Kalai and S. Vempala, Efficient algorithms for online decision problems, Journal of Computer and System Sciences, 71(3):291–307, 2005.

[15] W. M. Koolen, M. K. Warmuth, J. Kivinen et al., Hedging structured concepts. in COLT. Citeseer, 2010, pp. 93–105.

[16] A. Bernstein, N. J. Bouman, and J.-Y. L. Boudec, Real-time minimization of average error in the presence of uncertainty and convexification of feasible sets, arXiv preprint arXiv:1612.07287, 2016.

[17] A. Bernstein and E. Dall'Anese, Real-time feedback-based optimization of distribution grids: A unified approach, IEEE Transactions on Control of Network Systems, 6(3):1197–1209, 2019.

[18] X. Zhou, E. Dall'Anese, and L. Chen, Online stochastic optimization of networked distributed energy resources, IEEE Transactions on Automatic Control, 65(6):2387–2401, 2019.

[19] E. C. Hall and R. M. Willett, Online convex optimization in dynamic environments, IEEE Journal of Selected Topics in Signal Processing, 9(4):647–662, 2015.

[20] F. Bach et al., Learning with submodular functions: A convex optimization perspective, Foundations and Trends® in Machine Learning, 6(2–3)1:45–373, 2013.

[21] S. Iwata and K. Nagano, Submodular function minimization under covering constraints, in 2009 50th Annual IEEE Symposium on Foundations of Computer Science. IEEE, 2009, pp. 671–680.

[22] S. Jegelka and J. A. Bilmes, Approximation bounds for inference using cooperative cuts, in Proceedings of the 28th International Conference on Machine Learning (ICML-11). Citeseer, 2011, pp. 577–584.

[23] M. X. Goemans, N. J. Harvey, S. Iwata, and V. Mirrokni, Approximating submodular functions everywhere, in Proceedings of the twentieth annual ACM-SIAM symposium on Discrete algorithms. SIAM, 2009, pp. 535–544.

[24] G. Goel, C. Karande, P. Tripathi, and L. Wang, Approximability of combinatorial problems with multi-agent submodular cost functions, in 2009 50th Annual IEEE Symposium on Foundations of Computer Science. IEEE, 2009, pp. 755–764.

[25] D. S. Callaway and I. A. Hiskens, Achieving controllability of electric loads, Proceedings of the IEEE,99(1):184–199, 2010.

[26] P. Palensky and D. Dietrich, Demand side management: Demand response, intelligent energy systems, and smart loads, IEEE transactions on industrial informatics, 7(3):381–388, 2011.

[27] P. Siano, Demand response and smart grids—a survey, Renewable and sustainable energy reviews, 30:461–478, 2014.

[28] J. L. Mathieu, S. Koch, and D. S. Callaway, State estimation and control of electric loads to manage real-time energy imbalance, IEEE Transactions on power systems, 28(1):430–440, 2012.

[29] D. S. Callaway, Tapping the energy storage potential in electric loads to deliver load following and regulation, with application to wind energy, Energy Conversion and Management, 50(5):1389–1400, 2009.

[30] J. A. Taylor, S. V. Dhople, and D. S. Callaway, Power systems without fuel, Renewable and Sustainable Energy Reviews, 57:1322–1336, 2016.

[31] J. L. Mathieu, M. Kamgarpour, J. Lygeros, G. Andersson, and D. S. Callaway, Arbitraging intraday wholesale energy market prices with aggregations of thermostatic loads, IEEE Transactions on Power Systems, 30(2):763–772, 2015.

[32] K. Perlin, Improving noise, ACM Trans. Graph., 21(3):681–682, jul 2002. [Online]. Available: https://doi.org/10.1145/566654.566636

[33] J. A. Taylor, Convex optimization of power systems. Cambridge University Press, 2015.

[34] R. S. Rao, K. Ravindra, K. Satish, and S. Narasimham, Power loss minimization in distribution system using network reconfiguration in the presence of distributed generation, IEEE transactions on power systems, 28(1):317–325, 2012.

[35] S. Mishra, D. Das, and S. Paul, A comprehensive review on power distribution network reconfiguration, Energy Systems, 8(2):227–284, 2017.

[36] A. Khodabakhsh, G. Yang, S. Basu, E. Nikolova, M. C. Caramanis, T. Lianeas, and E. Pountourakis, A submodular approach for electricity distribution network reconfiguration, arXiv preprint arXiv:1711.03517, 2017.

[37] P. Paudyal, F. Ding, S. Ghosh, M. Baggu, M. Symko-Davies, C. Bilby, and B. Hannegan, The impact of behind-the-meter heterogeneous distributed energy resources on distribution grids, in 2020 47th IEEE Photovoltaic Specialists Conference (PVSC). IEEE, 2020, pp. 0857–0862.

[38] S. Lin and H. Zhu, Enhancing the spatio-temporal observability of grid-edge resources in distribution grids, IEEE Transactions on Smart Grid, 12(6):5434–5443, 2021.

[39] J. A. P. Lopes, F. J. Soares, and P. M. R. Almeida, Integration of electric vehicles in the electric power system, Proceedings of the IEEE, 99(1):168–183, 2010.

[40] R.-C. Leou, C.-L. Su, and C.-N. Lu, Stochastic analyses of electric vehicle charging impacts on distribution network, IEEE Transactions on Power Systems, 29(3):1055–1063, 2013.

[41] A. Navarro-Espinosa and L. F. Ochoa, Probabilistic impact assessment of low carbon technologies in lv distribution systems, IEEE Transactions on Power Systems, 31(3):2192–2203, 2015.

[42] H. Ahmadi and J. R. Martí, Minimum-loss network reconfiguration: A minimum spanning tree problem, Sustainable Energy, Grids and Networks, 1:1–9, 2015.

[43] R. C. Prim, Shortest connection networks and some generalizations, The Bell System Technical Journal, 36(6):1389–1401, 1957.

[44] L. Thurner, A. Scheidler, F. Schafer, J. H. Menke, J. Dollichon, F. Meier, S. Meinecke, and M. Braun, pandapower - an open source python tool for convenient modeling, analysis and optimization of electric power systems, IEEE Transactions on Power Systems, 2018.

[45] M. Baran and F. Wu, Network reconfiguration in distribution systems for loss reduction and load balancing, IEEE Transactions on Power Delivery, 4(2):1401–1407, 1989.

[46] R. Kavasseri and C. Ababei, Reds: Repository of distribution systems.