Titre: Title:	Online Second-order Methods for Time-Varying Equality-Constrained Optimization
Auteur: Author:	Jean-Luc Lupien
Date:	2023
Type:	Mémoire ou thèse / Dissertation or Thesis
Référence: Citation:	Lupien, J.-L. (2023). Online Second-order Methods for Time-Varying Equality-Constrained Optimization [Mémoire de maîtrise, Polytechnique Montréal]. PolyPublie. https://publications.polymtl.ca/54778/

URL de PolyPublie: PolyPublie URL:	https://publications.polymtl.ca/54778/
Directeurs de recherche: Advisors:	Antoine Lesage-Landry
Programme: Program:	Génie électrique

POLYTECHNIQUE MONTRÉAL
affiliée à l'Université de Montréal

Online Second-order Methods for Time-varying Equality-constrained Optimization

JEAN-LUC LUPIEN
Département de Génie électrique

Mémoire présenté en vue de l'obtention du diplôme de *Maîtrise ès sciences appliquées*
Génie électrique

Août 2023

POLYTECHNIQUE MONTRÉAL
affiliée à l'Université de Montréal

Ce mémoire intitulé :

Online Second-order Methods
for Time-varying Equality-constrained Optimization

présenté par **Jean-Luc LUPIEN**
en vue de l'obtention du diplôme de *Maîtrise ès sciences appliquées*
a été dûment accepté par le jury d'examen constitué de :

Roland MALHAMÉ, président
Antoine LESAGE-LANDRY, membre et directeur de recherche
Jérôme LE NY, membre

ACKNOWLEDGEMENTS

Research is made possible by readers like you. Thank you.

REMERCIEMENTS

J'aimerais remercier mon directeur de recherche Antoine Lesage-Landry pour sa patience et son attention au détail tout au long de cette maîtrise. Son enthousiasme et son énergie pour la recherche et les connaissances sont contagieux pour tous ceux qui l'entourent.

Merci au CRSNG (Conseil de recherches en sciences naturelles et en génie du Canada) et à IVADO pour leur support financier lors de cette maîtrise.

À Philippe Maisonneuve pour son intérêt et toutes nos conversations insolites; à Julien Pallage pour sa gentillesse et son écoute; à Loreley Sepho, Matthias Molénat, Abraham Kouamé N'Zi, Marie-Christine Paré, Olivier Daoust, Olivier Bélanger, Victor Darleguy, Samuel Mendoza, Christina Soldati, et Bouh Abdillahi: votre camaraderie et votre bonne humeur ont su ensoleiller nos journées dans un local, pourtant, sans fenêtre.

Merci à mes parents qui ont su inspirer en moi, dès un jeune âge, la curiosité et la persévérance. La réalisation de cette thèse aurait été impossible sans leur aide précieuse et constante.

À celle qui m'a doucement appuyé, soulevé et encouragé pendant cette épreuve, je dois ma plus sincère reconnaissance. Merci d'avoir été à mes côtés par les bons et mauvais moments. Patricia, je t'aime de tout mon cœur.

RÉSUMÉ

En optimisation convexe en temps réel (OCT), des décisions séquentielles sont prises par un agent, qui cherche à minimiser une fonction de coût déterminée par l'environnement. L'agent ne considère que de l'information passée du problème. La performance des algorithmes d'OCT est mesurée en comparant les coûts engendrés par la séquence de décisions choisie par l'algorithme et une séquence de référence. La métrique qui définit cette comparaison se nomme le regret. Si la séquence de référence employée est la séquence de décisions optimales pour chaque ronde et que le regret de l'algorithme est sous-linéaire, la moyenne temporelle du regret tend vers zéro lorsque l'horizon temporel tend vers l'infini. Ce faisant, la séquence de décisions de l'algorithme devient optimale en moyenne. Cette garantie de performance justifie l'utilisation d'algorithmes d'OCT dans des environnements variables, incertains ou antagonistes.

Si plusieurs procédés décisionnels peuvent être modélisés dans ce contexte, plusieurs applications réelles nécessitent que la séquence décisionnelle respecte des contraintes. Ces contraintes surviennent dans des problèmes tel que l'écoulement de puissance optimal dans les réseaux électriques; réseaux qui possèdent des limites physiques et de sécurité inhérentes. Violer ces contraintes de sécurité peut entraîner des conséquences désastreuses sur le réseau, dont des dommages sévères à l'infrastructure électrique. Pour de telles applications, des algorithmes d'OCT possédant à la fois un regret et une violation des contraintes (la distance minimale à l'ensemble réalisable d'une décision) sous-linéaires est l'objectif de conception.

Dans ce mémoire, plusieurs algorithmes d'OCT utilisant de l'information sur la dérivée seconde qui admettent des contraintes linéaires d'égalité variant dans le temps sont présentés. Tous ces algorithmes possèdent des bornes supérieures sous-linéaires sur leur regret et leur violation des contraintes. Les deux algorithmes principaux de ce mémoire, la méthode de Newton en temps-réel projetée et la méthode des points intérieurs en temps réel pour les contraintes d'égalité variant dans le temps, admettent une fonction objectif variant dans le temps et des inégalités généralisées, respectivement. Finalement, la performance supérieure de ces algorithmes par rapport à des algorithmes existants lorsqu'appliqués à des problèmes d'optimisation en temps réel pour l'allocation de ressources sur réseau et pour l'écoulement de puissance optimal, est illustrée en simulation.

Le but visé par ce mémoire est de faire un pas vers le développement d'algorithmes d'OCT pouvant simultanément traiter des fonctions objectif, des contraintes d'égalité, et des contraintes généralisées variant toutes dans le temps. Ce faisant, une telle approche permettrait

d'élargir énormément le champ d'application des algorithmes d'OCT.

ABSTRACT

In online convex optimization (OCO), an agent makes sequential decisions using only prior information with the goal of minimizing an environment-determined loss function. The performance of OCO algorithms is measured by comparing the loss suffered by the decision sequence made by the algorithm and a benchmark decision sequence which is captured by a performance indicator called regret. If the benchmark used is the round-optimal solution and the regret of the OCO algorithm is sublinear, the time-averaged difference between the algorithm's suffered loss and that of the round-optimal decisions goes to zero as the time horizon goes to infinity. This makes the algorithm's decision sequence optimal, on average. This performance guarantee motivates the use of OCO algorithms in changing, uncertain, or adversarial environments.

While many sequential decision-making processes can be modelled within this context, many real-world problems also require decisions to satisfy constraints. These constraints arise in problems such as the optimal power flow in electric power systems where the grid has inherent physical and safety limits. Violating these safety constraints can be catastrophic for the grid operator, potentially damaging electrical infrastructure. For such applications, OCO algorithms possessing simultaneous sublinear regret and constraint violation – a decision's distance from feasibility – are the design objective.

In this work, multiple OCO algorithms leveraging second-order information and admitting time-varying linear equality constraints are presented. These algorithms all possess sublinear dynamic regret and constraint violation bounds. The differences between these algorithms reside in the types of constraints they consider. The two main algorithms of this Master's thesis, the online projected Newton's method (`OPEN-M`) and the online interior-point method for time-varying equality constraints (`OIPM-TEC`) admit a time-varying loss function and generalized inequalities, respectively. These algorithms are shown to outperform previous OCO algorithms from the literature in online network resource allocation and online optimal power flow problems.

The hope for this Master's thesis is for it to become a building-block towards the development of online optimization algorithms capable of simultaneously tackling time-varying loss functions, time-varying equality constraints and time-varying generalized inequalities; in the process, immensely expanding the real-world applicability of OCO approaches.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF ALGORITHMS

LIST OF SYMBOLS AND ABBREVIATIONS

OCO	Online Convex Optimization
OCT	Optimization Convexe en Temps-réel
OPF	Optimal Power Flow
PSD	Positive Semi-definite
QCQP	Quadratically Constrained Quadratic Program
`OEN-M`	Online Equality-constrained Newton's Method
`OPEN-M`	Online Projected Equality-constrained Newton's Method
`MOSP`	Modified Online Saddle-Point Method
`MALM`	Modified Augmented Lagrangian Method
`OIPM-TEC`	Online Interior-point Method for Time-varying Equality Constraints
ϵ`OIPM-TEC`	Epsilon Online Interior-point Method for Time-varying Equality Constraints

CHAPTER 1 INTRODUCTION

Online convex optimization (OCO) is an emerging research field that lies at the intersection between optimization, game theory, and machine learning. In the OCO framework, an agent interacts with the environment by making sequential decisions [1–3]. At every round, after the agent has committed to a decision, an environment-determined loss (or objective) function is revealed. The agent then suffers the loss determined by the loss function evaluated at the set decision. The agent seeks to minimize his cumulative losses but only has access to past information. The loss function can be chosen randomly or even adversarially [2].

Although this problem may appear daunting or even impossible, OCO algorithms possess theoretical performance guarantees under mild assumptions [2–4]. Specifically, the sequence of decisions chosen by OCO algorithms are guaranteed to be indiscernible, on average, from the optimal sequence of decisions given a sufficiently long time-horizon. OCO algorithms are, therefore, highly adaptable and resilient even under uncertain, changing, or adversarial conditions. This can be captured by a sublinear cumulative difference between the sequence of losses suffered by the algorithm and a benchmark. This cumulative difference is referred to as regret and will be discussed extensively in this Master's thesis.

The assumption that the agent has access to only past information is one of the defining aspects of OCO. Instead of relying on a model, the nature and evolution of the environment are assumed to be unknown. The agent must, therefore, learn about the environment through experimentation. The ability to maintain provable performance in this context is the distinguishing feature of OCO approaches. Indeed, OCO presents a robust approach to interacting and learning from unpredictable or complex environments. Rather than purely seeking the optimal solution to a given loss function as in offline optimization, OCO uses optimization as a process [2]. Not only are OCO agents learning and improving their decision-making in unknown environments, they do so in such a way that their learning sequence never suffers an arbitrarily high loss. This advantage opens up the possibility of applying such methods in safety-critical dynamic environments.

1.1 OCO Applications

Many real-world scenarios can be modelled as OCO problems. A classic example is that of online portfolio selection [2, 5, 6]. In this context, the agent is an investor seeking to determine the best portfolio of assets to own at every moment with the objective of maximizing their profits – or, equivalently, minimizing their losses [5]. Given a finite set of assets, the decision vector is the proportion of the investor's wealth invested in each asset. The loss function represents the increase or decrease of wealth associated with holding the chosen assets over a certain period of time. The investor only has access to prior information on asset prices. Because asset prices are random or even directly affected by the investor's decisions, the assumption of a randomly or adversarially selected loss function is well-suited to the problem [7].

However, not all decision-making scenarios can be so easily modelled in this manner. The online optimal power flow (OPF) problem is one such example. In the online OPF, the operator of the electric power grid must determine the optimal generation profile that minimizes costs while serving all loads in the network, which will inevitably vary with time. Because of the nature of such power systems, however, certain physical and safety limits must be respected. These can be modelled as constraints and need to be integrated into the classical OCO framework. Moreover, much like the loss function, these constraints are revealed after decisions are made. In applications such as the online OPF, the performance of the algorithm cannot be solely based on obtaining limited cumulative losses. In this setting, algorithms must also yield feasible solutions – which can be measured by the cumulative distance from feasibility: constraint violation.

1.2 Objective of the Presented Work

Simultaneously achieving sublinear regret and constraint violation bounds is a difficult challenge in OCO [8–10]. In fact, in the general case, it is proven that it is impossible to have both sublinear regret and constraint violation bounds [11, 12]. However, if certain mild assumptions are met, this goal becomes possible [12, 13]. Tackling this important challenge is the objective of this work. In this Master's thesis, OCO algorithms capable of efficiently handling time-varying equality constraints are presented. Specifically, OCO algorithms achieving sublinear regret and constraint violation bounds are presented when subject to time-varying linear equality constraints. Such algorithms have many possible applications to constrained dynamic problems such as the online OPF and the online network resource allocation problem.

1.3 Structure of the Thesis

The remainder of this Master's thesis consists of six chapters. First, a review of the relevant literature is presented in Chapter 2. A general framework of the research forming the backbone of this thesis is presented in Chapter 3. Next, a published article pertaining to an online Newton's method with time-varying equality constraints is presented in Chapter 4.

An article submitted for publication on an online interior-point method with time-varying equality constraints is presented in Chapter 5. This method builds on the previously proposed second-order methods, adding the properties of being projection-free and admitting generalized inequalities. A general discussion on the impact of these two articles is then presented in Chapter 6. Finally, concluding remarks and future research avenues are explored in Chapter 7.

CHAPTER 2 LITERATURE REVIEW

In this chapter, a review of the OCO literature is presented. The types of problems typically considered in OCO are explained along with their corresponding solution concepts.

2.1 First Steps in OCO

The field of online convex optimization is relatively new in the literature. The foundational paper was published in 2003 by Zinkevich in which the author introduces the concept of iterative convex optimization and proposes the first online optimization algorithm [3]. Let $\mathbf{x}_t \in \mathbb{R}^N$, $N \in \mathbb{N}$ be the decision vector at every timestep $t = 1, 2, 3, ...T$. Let $f_t : \mathbb{R}^N \mapsto \mathbb{R}$ be the loss function and $\mathcal{X} \subseteq \mathbb{R}^N$ be a compact, convex set. The canonical OCO problem at every time t, can then be expressed as:

$$\min_{\mathbf{x}_t \in \mathcal{X}} f_t(\mathbf{x}_t).$$

Because [3] is interested mostly in time-invariant functions, the assumption that one does not have access to $f_t(\mathbf{x})$ before implementing the decision $\mathbf{x}_t$ is not explicitly stated. However, the modern understanding and formulation of OCO is such that f_t is revealed only after committing to a decision, i.e., any current-round information is only revealed at the subsequent round [1, 2]. This one-step delay in information is what enables modelling of stochastic, dynamic or adversarial environments and distinguishes OCO from time-varying optimization.

The static regret $R_{\mathrm{s}}(T)$ is most often used in this context, which can be defined as:

$$R_{\mathrm{s}}(T) = \sum_{t=1}^{T} f_t(\mathbf{x}_t) - \min_{\mathbf{x}^* \in \mathcal{X}} \sum_{t=1}^{T} f_t(\mathbf{x}^*).$$

The algorithm outlined by [3], based on a gradient descent approach, obtains sublinear static regret bounds. From this influential paper have stemmed multiple online optimization approaches to solving discrete-time iterative problems. More mature approaches leading to tighter regret bounds were subsequently developed in [14, 15]. Algorithms have also been proposed for distributed OCO extensions including in [16] and [17].

Another important extension of the OCO framework incorporates the additional challenge that the agent receives only limited feedback from the environment. So-called bandit OCO is also an active area of research in which fall the works of [18–20]. These approaches can be more representative of the real-world behaviour of electric grid sensors because reporting can be intermittent, noisy, or even erroneous in certain cases [21, 22].

2.2 From Static Regret to Dynamic Regret

An issue arising with the use of static regret is that it fails to capture dynamic changes in optima. An alternative performance metric to the static regret is the dynamic regret. Let $\mathbf{x}_t^* \in \underset{\mathbf{x}\in\mathcal{X}}{\operatorname{argmin}} f_t(\mathbf{x})$ be the round-optimal decision. The dynamic regret $R_{\mathrm{d}}(t)$ can then be expressed as:

$$R_{\mathrm{d}}(T) = \sum_{t=1}^{T} f_t(\mathbf{x}_t) - f_t(\mathbf{x}_t^*). \tag{2.1}$$

Comparing an algorithm's decision sequence to the round-optimal decisions rather than the best fixed decision in hindsight changes the fundamental design objective. Instead of seeking to converge towards a fixed solution, tracking a changing optima through time becomes the goal. Sublinear dynamic regret implies an algorithm will, on average, track the optimal solution exactly as the time horizon increases. Such algorithms have a much larger sphere of application as they are much better-suited to dynamic environments. Interestingly, the seminal paper [3] is also the first to introduce the dynamic regret as a measure of performance. More refined methods are presented in [4, 23, 24] that present improved sublinear dynamic regret bounds. The first online algorithm to achieve sublinear regret bounds predicated only on the accumulated variation of optima (also referred to as path-length) is presented in [25]. This regret bound is the tightest known to date in the OCO literature. Subsequently, an algorithm was presented in [26], the online Newton's method, that has matched these bounds and even provided an explicit bound on optima variation under which dynamic regret is $O(1)$ where O denotes the asymptotic order.

2.3 Integrating Time-varying Constraints

The integration of time-varying constraints into the OCO setting, which is referred to as constrained OCO in the sequel, is another extension of the OCO framework. Let $g_{i,t}(\mathbf{x})$ be convex functions for all $i = 1, 2, 3...I$, $I \in \mathbb{N}^+$ at every round t. The most common formulation of constrained OCO can be written as:

$$\begin{aligned} \min_{\mathbf{x}_t\in\mathcal{X}} \quad & f_t(\mathbf{x}_t) \\ \text{s.t.} \quad & g_{i,t}(\mathbf{x}_t) \le 0 \quad \forall i \end{aligned}$$

The constraint violation metric used for this problem, $\mathrm{Vio}(T)$, is defined as:

$$\mathrm{Vio}(T) = \sum_{t=1}^{T}\sum_{i=1}^{I}\left[g_{i,t}(\mathbf{x}_t)\right]^+ ,$$

where the operator $[x]^+$ denotes $\max\{x, 0\}$.

Similarly to the loss function, the constraint functions $g_{i,t}$ are only revealed once the agent has committed to the decision $\mathbf{x}_t$. In the most general case, the number of constraints I at every round can even be time-varying. This is what constitutes the primary challenge in constrained OCO. There is, therefore, no way to guarantee satisfaction of all constraints at any time step. However, just as for regret, under certain assumptions, constrained OCO algorithms can achieve sublinear regret; obtaining feasible decisions on average over long time horizons.

The first step taken in constrained OCO was to consider long-term or slowly varying constraints. This is done in [11, 12, 27], where sublinear constraint violation was obtained but only under static regret analysis. Another common method in constrained OCO is to employ virtual queues. This entails solving a Lagrangian system. Approaches that fall into this category include [11, 13, 18] that possess sublinear static regret and constraint violation bounds. A stronger performance guarantee is obtained when designing a constrained OCO algorithm posessing sublinear dynamic regret bounds. This has been done in [10] and [8] using Lagrangian approaches and [9] by using Polyak-lojasiewicz regions. However, the dynamic regret bounds presented are not as tight as those established for the algorithms presented in this work.

2.4 Second-order Optimization methods

One of the simplest optimization algorithms is the gradient descent algorithm. The idea behind this algorithm is seeking to decrease the loss by moving in the direction opposite that of the steepest slope [28]. Only first-order information (the gradient) is used to compute the update. The idea behind second-order methods is to use the information contained in the Hessian matrix ($\nabla^2 f_t(\mathbf{x})$) to obtain a better approximation of the function, and, therefore, to converge more rapidly towards the optimal solution [28, 29]. Such approaches have many inherent advantages such as not being sensitive to stepsize selection the same way first-order approaches are [29,30]. Another advantage of such methods is the superlinear convergence displayed by Newton's method, a second-order algorithm [30, 31]. Moreover, if the objective functions possess a bounded third derivative (see self-concordance [32, Section 2.3]), there exists an explicit formulation for the number of iterations required to obtain a solution given a certain tolerance [29,30]. Additionally, interior-point methods – second order methods for constrained optimization – have been developed to conserve this bound on the number of iterations for convex constrained optimization [30]. Interior-point methods have the advantage of only yielding feasible (interior) points at every iteration [29, 31, 32].

Online second-order methods have been developed which benefit from these properties. These include the aforementioned online Newton's step from [26]. In [33], the feasibility guarantee of offline interior-point methods is used to develop a projection-free constrained OCO algorithm with sublinear static regret bounds. A quasi-Newton approach was extended to the online setting and used to solve a simplified version of the online OPF in [34]. The algorithms presented in this Master's thesis build on these results and make full-use of the advantageous properties of second-order methods.

2.5 The Optimal Power Flow Problem

The online optimal power flow problem (OPF) is a promising application of OCO algorithms [35]. The OPF consists in determining the most efficient operational settings for generation, voltage, and power flows across transmission lines whilst ensuring all loads are adequately supplied. This problem is of paramount importance for operators of the electric grid. Indeed, finding the optimal solution to this problem enables the operator to minimize losses – resistive and monetary – while respecting the physical and safety limits imposed on the transmission network. The ability to quickly and efficiently solve this problem ensures reliability and stability in the electric grid, hence, reducing the likelihood of disruptions or outages.

The AC-OPF problem, so-called because no relaxations are employed in its definition, was first presented in [36]. The power flow equations, active and reactive generation limits, maximum apparent power limits in power lines, and voltage limits are all considered in this formulation. The objective of the AC-OPF is usually minimizing the generation costs, assumed to be convex quadratic in the quantity of active power generated. Consider the set of buses in the network $\mathcal{B} \subseteq \mathbb{N}^b$, $b \in \mathbb{N}$. Consider the set of generators $\mathcal{G} \subseteq \mathcal{B}$ connected to a set of loads $\mathcal{D} \subseteq \mathcal{B}$ through a network of power lines $\mathcal{L} \subseteq \mathcal{B} \times \mathcal{B}$. Let $\gamma_{ij} \in \mathbb{C}$ denote the admittance of the line $ij \in \mathcal{L}$. Let $p_i \in \mathbb{R}$ and $q_i \in \mathbb{R}$ indicate the active and reactive power balance at node $i \in \mathcal{B}$. Let $p_i^{\mathrm{g}} \in \mathbb{R}$ and $q_i^{\mathrm{g}} \in \mathbb{R}$ denote active and reactive power generation at bus i and $p_i^{\mathrm{d}} \in \mathbb{R}$ and $q_i^{\mathrm{d}} \in \mathbb{R}$ denote active and reactive power demand at bus i. Let $p_{ij} \in \mathbb{R}$ and $q_{ij} \in \mathbb{R}$ represent active and reactive power flow from bus i to bus j. Let $\underline{p}, \underline{q}, \overline{p}, \overline{q} \in \mathbb{R}$ denote the minimum and maximum active and reactive power generation, $\overline{k}_{ij} > 0$ denote the apparent power limits in the power line ij, and $\underline{v}, \overline{v} \geq 0$ denote the voltage magnitude limits.

Finally, let $\mathbf{v}_i \in \mathbb{C}$ denote the voltage phasor at bus i. The physical and safety limits of the generators and power lines can be written as:

$$\underline{p_i} \leq p_i^{\mathrm{g}} \leq \overline{p_i}, \forall i \in \mathcal{G}$$
$$\underline{q_i} \leq q_i^{\mathrm{g}} \leq \overline{q_i}, \forall i \in \mathcal{G}$$
$$p_{ij}^2 + q_{ij}^2 \leq \overline{k}_{ij}^2, \forall ij \in \mathcal{L}.$$

The safety limits on voltage can be expressed as:

$$\underline{v} \leq |\mathbf{v}_i| \leq \overline{v},$$

where $|x|$ denotes the magnitude of the complex number x.

The nodal power balance can be expressed as:

$$p_i = p_i^{\mathrm{g}} - p_i^{\mathrm{d}} = \sum_i p_{ij}, \quad \forall i \in \mathcal{B}$$
$$q_i = q_i^{\mathrm{g}} - q_i^{\mathrm{d}} = \sum_i q_{ij}, \quad \forall i \in \mathcal{B}.$$

The power flow equation can be expressed as:

$$p_{ij} + \jmath q_{ij} = \mathbf{v}_i(\mathbf{v}_i^* - \mathbf{v}_j^*)\gamma_{ij}^*, \quad \forall ij \in \mathcal{L}, \tag{2.2}$$

where $*$ denotes the complex conjugate.

The objective function (for cost minimization) is written as:

$$\sum_{i \in \mathcal{G}} a_i (p_i^{\mathrm{g}})^2 + b_i p_i^{\mathrm{g}} + c_i.$$

A useful reformulation of the constraint (2.2) employs the matrix $\mathbf{W} \in \mathbb{C}^{\mathrm{card}\mathcal{B}^2}$ defined so that $\mathbf{W}_{ij} = \mathbf{v}_i \mathbf{v}_j^*$. The constraint (2.2) can then be equivalently re-expressed with the following three constraints:

$$p_{ij} + \jmath q_{ij} = (\mathbf{W}_{ii} - \mathbf{W}_{ij})\gamma_{ij}^*, \quad \forall ij \in \mathcal{L}$$
$$\mathbf{W} \succeq 0$$
$$\text{rank } \mathbf{W} = 1.$$

The first constraint is simply a linear constraint which is convex. The second constraint is a positive semidefinite constraint. Although more difficult to address than a linear constraint, a semidefinite constraint is in fact a proper cone which can be treated using commercial solvers. The formal definition of a proper cone and its relevance to convex problems will be discussed in the next chapter.

The issue with this formulation is that it is a non-convex quadratically constrainted quadratic program (QCQP). As such, solving the AC-OPF problem is NP-hard [37, 38]. Because of the scale of power grids and the dynamic nature of loads, this makes the AC-OPF impractical in many cases. To alleviate this obstacle, convex relaxations of the AC-OPF problem have been proposed [39]. Being convex problems, they can be solved in polynomial time. Because they are relaxations, they can often yield exact solutions, i.e., solutions that are feasible and optimal with respect to the original AC-OPF.

The tightest convex relaxation of the OPF involves removing the rank 1 condition on the matrix $\mathbf{W}$. This transforms the problem into a semi-definite program. This relaxation was first proposed in [40]. Another convex relaxation of the OPF is a second-order cone relaxation. This was first done in [41] and remains one of the fastest relaxations to solve numerically due to the maturity of the associated commercial solvers. To obtain this relaxation, we remove the positive semidefinite constraint $\mathbf{W} \succeq 0$ and instead impose a non-negativity condition on all 2×2 principal minors of the matrix $\mathbf{W}$. The resulting constraint can be expressed as:

$$\mathbf{W}_{ii}\mathbf{W}_{jj} \geq |\mathbf{W}_{ij}|^2, \quad \forall ij \in \mathcal{L}.$$

This can also be expressed as a canonical second-order cone constraint which takes the following form:

$$\left\| \begin{matrix} 2\mathbf{W}_{ij} \\ \mathbf{W}_{ii} - \mathbf{W}_{jj} \end{matrix} \right\| \leq \mathbf{W}_{ii} + \mathbf{W}_{jj}, \quad \forall ij \in \mathcal{L}.$$

Convex relaxations of the OPF, especially the second-order cone relaxation, play a vital role in the numerical application of the OCO algorithms presented in Chapter 5. Other convex relaxations of the OPF have been proposed. A survey of such relaxations is provided in [42].

CHAPTER 3 DISCUSSION OF THE RESEARCH PROJECT AS A WHOLE AND GENERAL ORGANIZATION OF THE THESIS OR DISSERTATION INDICATING THE COHERENCE OF THE ARTICLES IN RELATION TO THE RESEARCH GOALS

In this thesis, OCO algorithms capable of efficiently tackling time-varying equality constraints are presented. Specifically, OCO algorithms achieving provable performance with respect to regret and constraint violation are developed. Two main categories of online optimization problems are considered in this work. Let $\mathbf{x}_t \in \mathbb{R}^N$, $N \in \mathbb{N}$ be the decision vector at every timestep $t = 1, 2, 3, ...T$. Let $\mathbf{A}_t \in \mathbb{R}^{P \times N}$, $P \in \mathbb{N}$ be a full row-rank matrix and $\mathbf{b}_t \in \mathbb{R}^P$ be a vector. In Chapter 4, a problem in which an agent seeks to play the best decision $\mathbf{x}_t$ with respect to a time-varying loss function $f_t(\mathbf{x}_t) : \mathbb{R}^N \mapsto \mathbb{R}$ while satisfying the time-varying linear equality constraints: $\mathbf{A}_t\mathbf{x}_t = \mathbf{b}_t$ is considered. Formally, at every round t, the following problem is being resolved:

$$\begin{aligned} \min_{\mathbf{x}_t} \quad & f_t(\mathbf{x}_t) \\ \text{s.t.} \quad & \mathbf{A}_t\mathbf{x}_t = \mathbf{b}_t. \end{aligned} \tag{3.1}$$

Notice in this formulation that f_t, $\mathbf{A}_t$ and $\mathbf{b}_t$ are time-varying. Because the problem is viewed in the constrained OCO framework, all three of these time-varying parameters are revealed after having committed to the decision $\mathbf{x}_t$.

The problem written in (3.1) is capable of representing many real-world engineering scenarios such as network flows, portfolio management under constraints, or resource allocation [7,13]. To illustrate the performance of the proposed algorithms, a network flow problem modelling electric grid distribution is presented. Specifically, the problem can be seen as a branch flow relaxation of the OPF problem [37]. A commodity, in this case electricity, is transmitted through a network composed of edges (power lines) and nodes (buses). Transporting a commodity through an edge incurs a cost, in this case resistive losses. Only real power is considered.

The problem presented in (3.1) does not integrate inequality constraints. In Chapter 5, generalized inequality constraints are considered. To understand these generalized inequalities, we must first define the concept of a proper cone. From [29, Section 2.4], we have that $K \subseteq \mathbb{R}^N$ is a proper cone if it is convex, closed, has a nonempty interior, and contains no line. Proper cones permit the definition of a partial ordering such that:

$$\mathbf{x} \preceq_K \mathbf{y} \Longleftrightarrow \mathbf{y} - \mathbf{x} \in K.$$

From [29, Section 3.6.2], a function $f : \mathbb{R}^N \mapsto \mathbb{R}^M$, $N, M \in \mathbb{N}$ is said to be convex with respect to the proper cone K if for all $\mathbf{x}, \mathbf{y} \in \mathbb{R}^N$ and $0 \leq \alpha \leq 1$:

$$f(\alpha\mathbf{x} + (1-\alpha)\mathbf{y}) \preceq_K \alpha f(\mathbf{x}) + (1+\alpha) f(\mathbf{y}).$$

If K is one-dimensional, we retrieve the standard definition of a convexity. These concepts can now be utilized to integrate generalized inequalities into our framework.

Let $K_i \in \mathbb{R}^{k_i}$, $k_i \in \mathbb{N}$, $i = 1, 2, 3, ..., I$ be proper cones. Let $g_i : \mathbb{R}^N \mapsto \mathbb{R}^{k_i}$, $i = 1, 2, 3, ..., I$ be convex functions with respect to the cone K_i. We add the generalized constraints $g_i(\mathbf{x}_t) \preceq_{K_i} \mathbf{0}$ to the problem considered in the previous chapter. However, because of the complexity of simultaneously considering equality and inequality constraints, the inequality constraints and the objective function are assumed to be time-invariant. This gives us the following problem to solve at every round t:

$$\begin{aligned} \min_{\mathbf{x}_t} \quad & f(\mathbf{x}_t) \\ \text{s.t.} \quad & \mathbf{A}\mathbf{x}_t = \mathbf{b}_t \\ & g_i(\mathbf{x}_t) \preceq_{K_i} \mathbf{0} \quad \forall i. \end{aligned} \tag{3.2}$$

This problem can be reformulated so that the objective function becomes linear by introducing a slack variable $s \in \mathbb{R}$. The problem is then expressed as:

$$\begin{aligned} \min_{s, \mathbf{x}_t} \quad & s \\ \text{s.t.} \quad & f(\mathbf{x}_t) \leq s \\ & \mathbf{A}\mathbf{x}_t = \mathbf{b}_t \\ & g_i(\mathbf{x}_t) \preceq_{K_i} \mathbf{0} \quad \forall i. \end{aligned}$$

Notice here the only time-varying parameter is $\mathbf{b}_t$. This will be the only parameter revealed after the decision $\mathbf{x}_t$ is implemented, all others being known before the start of the algorithm. Although this limits the scenarios that can be represented, this can be considered the first step towards more general approaches. Moreover, this formulation still remains useful notably because it is possible to express the online OPF in this way. Indeed, the only time-varying parameter are the loads because the network configuration is generally time-invariant as are the characteristics of the lines and generators. For the online OPF, the generalized inequalities take the form of either a PSD constraint or a second-order cone constraint, two different types of proper cones. Because of this, the algorithms in Chapter 5 are applied to an online OPF problem.

Both constrained OCO problems described in (3.1) and (3.2) require dedicated approaches to be solved efficiently. Such approaches are presented in the two chapters that follow.

CHAPTER 4 ARTICLE 1 : AN ONLINE NEWTON'S METHOD FOR TIME-VARYING EQUALITY CONSTRAINTS

Jean-Luc Lupien • Antoine Lesage-Landry

Published in: IEEE Control Systems Letters, 22 February 2023 [43].

Abstract

We consider online optimization problems with time-varying linear equality constraints. In this framework, an agent makes sequential decisions using only prior information. At every round, the agent suffers an environment-determined loss and must satisfy time-varying constraints. Both the loss functions and the constraints can be chosen adversarially. We propose the Online Projected Equality-constrained Newton Method (OPEN-M) to tackle this family of problems. We obtain sublinear dynamic regret and constraint violation bounds for OPEN-M under mild conditions. Namely, smoothness of the loss function and boundedness of the inverse Hessian at the optimum are required, but not convexity. Finally, we show OPEN-M outperforms state-of-the-art online constrained optimization algorithms in a numerical network flow application.

Keywords. Optimization algorithms, Time-varying systems, Machine learning.

4.1 Introduction

In online convex optimization (OCO), an agent aims to sequentially play the best decision with respect to a potentially adversarial loss function using only prior information [1, 3]. In other words, decisions must be made before observing the loss function and constraints. OCO algorithms have many applications including portfolio selection, artificial intelligence, and real-time control of power systems [2, 22, 44].

The preponderant performance metric for OCO algorithms is regret [1], the cumulative difference between the loss incurred by the agent and that of a comparator sequence. Two main types of regret exist: static and dynamic. For static regret, the comparator sequence is defined as the best fixed decision in hindsight [3]. For dynamic regret, the comparator sequence is the round-optimal decision in hindsight [2, 3]. In OCO, one intends to design an algorithm that possesses a sublinear regret bound. Sublinear regret implies that the time-averaged regret goes to zero as time increases. The OCO algorithm therefore plays the best decision with respect to the comparator sequence over sufficiently long time horizons [1, 2].

A persistent obstacle in OCO algorithmic design has been the integration of time-varying constraints. In this context, the decision sequence must also satisfy environment-determined constraints [10]. The performance of such algorithms is measured both in terms of its regret but also its constraint violation, the cumulative distance from feasibility of the agent's decisions. Similarly to the regret, a sublinear constraint violation is desired and implies that, on average, decisions will be feasible over a large time horizon [11]. When considering time-varying constraints, an analysis using the dynamic regret is preferable over one based on the static regret because the best fixed feasible decision can have an arbitrarily large loss or might not even exist [12].

Most OCO algorithms tackling time-varying constraints are analysed based only on static regret. In [11,12] and [33], sublinear regret bounds and violation bounds are achieved for long-term constraints. Sublinear static regret and constraint violation bounds are also achieved in [27] using virtual queues and in [18] using an online saddle-point algorithm. More recently, an augmented Lagrangian method [13] has been shown to outperform previous Lagragian-based methods [10, 18, 27] in numerical experiments.

Algorithms with dynamic regret bounds have also been developed such as the modified online saddle-point method (`MOSP`) [10]. This algorithm has simultaneous sublinear dynamic regret and constraint violation bounds. However, `MOSP`'s bounds are dependent on strict conditions on optimal primal and dual variable variations in addition to time-sensitive step sizes. An exact-penalty method for dealing with time-varying constraints possessing sublinear regret and constraint violation bounds was presented in [45] but shares `MOSP`'s step size limitation. Virtual queues are used in [8] with time-varying constraints achieving simultaneous sublinear regret and constraint violation bounds without requiring Slater's condition to hold. Sublinear regret and constraint violation are achieved in [9] including for some non-convex functions but with considerably looser bounds.

The application of an interior-point method to time-varying convex optimization is presented in [46]. However, this context differs from OCO because current-round information is available to the decision-maker.

In this work, our main contribution is the design of a novel online optimization algorithm that can efficiently deal with time-varying linear equality constraints in the OCO setting.

This method simultaneously possesses the tightest dynamic regret and constraint violation bounds presented thus far for constrained OCO problems. Additionally, the method does not require hyperparameters, time-dependent step sizes, or a predefined time horizon making it easily implementable.

4.2 Background

In recent work, a second-order method for online optimization yielded tighter dynamic regret bounds compared to first-order approaches [26]. Specifically, [26] proposes an online extension of Newton's method, `ONM`, applicable to non-convex problems that possesses a tight dynamic regret bound. However, this approach is only applicable to unconstrained problems. In an offline setting, the unconstrained and linear equality-constrained Newton's method's performance are the same [29, 32]. This result motivates the extension of `ONM` to a setting with time-varying linear equality constraints.

4.2.1 Problem definition

We consider online optimization problems of the following form. Let $\mathbf{x}_t \in \mathbb{R}^n$, $n \in \mathbb{N}$ be the decision vector at time t. Let f_t: $\mathbb{R}^n \mapsto \mathbb{R}$ be a twice-differentiable function. Let $\mathbf{A}_t \in \mathbb{R}^{p \times n}$ be a rank $p \in \mathbb{N}$ full row-rank matrix, and $\mathbf{b}_t \in \mathbb{R}^p$. The problem at round $t = 1, 2, ..., T$ can then be written as:

$$\begin{aligned} &\min_{\mathbf{x}_t} f_t(\mathbf{x}_t) \\ &\text{s.t.} \quad \mathbf{A}_t \mathbf{x}_t = \mathbf{b}_t. \end{aligned} \tag{4.1}$$

In this work, dynamic regret will be used as the performance metric because it is more stringent compared to static regret. Indeed, sublinear dynamic regret implies sublinear static regret [12]. Dynamic regret $R_\mathrm{d}(T)$ is defined as:

$$R_\mathrm{d}(T) = \sum_{t=1}^{T} \Big[f_t(\mathbf{x}_t) - f_t(\mathbf{x}_t^*) \Big], \tag{4.2}$$

where $\mathbf{x}_t^*$ is the round-optimal solution and $T \in \mathbb{N}$ is the time horizon. For (4.1), the round-optimum $\mathbf{x}_t^*$ is the solution to the following system of equations:

$$\begin{aligned} \nabla f_t(\mathbf{x}_t^*) + \mathbf{A}_t^\top \boldsymbol{\nu}_t^* &= 0 \\ \mathbf{A}_t \mathbf{x}_t^* - \mathbf{b}_t &= 0, \end{aligned}$$

where $\boldsymbol{\nu}_t^* \in \mathbb{R}^n$ is the dual variable associated with (4.1)'s equality constraints.

The constraint violation term is defined as:

$$\mathrm{Vio}(T) = \sum_{t=1}^{T} \left\| \mathbf{A}_t \mathbf{x}_t - \mathbf{b}_t \right\|, \tag{4.3}$$

and quantifies the cumulative distance from feasibility, with respect to the Euclidean norm, of the decision sequence. All norms, $\|\cdot\|$, refer to the Euclidean norm in the sequel. Constraint violation is zero if decisions are feasible at all rounds. This definition is similar to that used in [9, 10] and is stricter than in [13] and [8] because the constraints must be satisfied at every timestep and not on average.

In the first part of this work, we investigate the case for which the feasible space is the same for all rounds, i.e., $\mathbf{A}_t = \mathbf{A}, \mathbf{b}_t = \mathbf{b}$ for all t. The second part builds on this result and extends it to time-varying equality constraints.

4.2.2 Preliminaries

We define the function $\mathbf{D}_t(\mathbf{x}) : \mathbb{R}^n \mapsto \mathbb{R}^{(n+p)^2}$ as:

$$\mathbf{D}_t(\mathbf{x}) = \begin{bmatrix} \nabla^2 f_t(\mathbf{x}) & \mathbf{A}_t^\top \\ \mathbf{A}_t & 0 \end{bmatrix},$$

where $\nabla^2 f_t(\mathbf{x})$ is the Hessian matrix of f_t. We assume that the Hessian is invertible for all t which implies that $\mathbf{D}_t(\mathbf{x})$ is also invertible [29, Section 10.1]. This guarantees that the Newton update is defined at every round.

Next, we present the online equality-constrained Newton (`OEN`) update. For any feasible point $\mathbf{x}_t$, the `OEN` update minimizes the second-order approximation of f_t around $\mathbf{x}_t$ subject to the equality constraints. An estimate of the optimal dual variable, $\boldsymbol{\nu}_t$, is also obtained from the update.

Definition 4.1 (`OEN` update)**.** *The `OEN` update is:*

$$\begin{aligned} \begin{bmatrix} \Delta \mathbf{x}_t \\ \boldsymbol{\nu}_t \end{bmatrix} &= -\mathbf{D}_t^{-1}(\mathbf{x}_t) \begin{bmatrix} \nabla f_t(\mathbf{x}_t) \\ 0 \end{bmatrix} \\ \mathbf{x}_{t+1} &= \mathbf{x}_t + \Delta \mathbf{x}_t. \end{aligned} \tag{4.4}$$

Let $\mathbf{v}_t \in \mathbb{R}^n$ be the difference between subsequent optima: $\mathbf{v}_t = \mathbf{x}_{t+1}^* - \mathbf{x}_t^*$. Throughout this work, we assume: $0 \leq \|\mathbf{v}_t\| \leq \overline{v}$ for all t. This limits the variation in optima between two subsequent rounds. It is a common assumption in dynamic OCO [9,10,26]. This assumption could be satisfied in real-world applications such as electric grids where the temporal continuity imposed by the underlying physics limits the variation in optima provided the timestep is sufficiently small. The total variation V_T is defined as:

$$V_T = \sum_{t=0}^{T-1} \left\| \mathbf{x}_{t+1}^* - \mathbf{x}_t^* \right\| = \sum_{t=0}^{T-1} \left\| \mathbf{v}_t \right\|,$$

and is bounded above by $V_T \leq \overline{v}T$.

An important tool for the analysis of the `OEN` update is the reduced function $\tilde{f}_t(\mathbf{z}) : \mathbb{R}^{n-p} \mapsto \mathbb{R}$ which is a representation of $f_t(\mathbf{x})$ over the feasible set.

Let $\mathbf{F}_t \in \mathbb{R}^{n\times(n-p)}$ be such that $\mathcal{R}(\mathbf{F}_t) = \mathcal{N}(\mathbf{A}_t)$ where $\mathcal{R}$ is the column space of a matrix and $\mathcal{N}$ is its null space. Let $\hat{\mathbf{x}} \in \mathbb{R}^n$ be such that $\mathbf{A}_t\hat{\mathbf{x}} - \mathbf{b}_t = \mathbf{0}$. Then, $\tilde{f}_t$ is defined as:

$$\tilde{f}_t(\mathbf{z}) = f_t(\mathbf{F}_t\mathbf{z} + \hat{\mathbf{x}}).$$

We remark that the reduced function shares minima with f_t, i.e., $\min_{\mathbf{z}} \tilde{f}_t(\mathbf{z}) = f_t(\mathbf{x}_t^*)$ [29]. This gives rise to an equivalent unconstrained, reduced problem: $\min_{\mathbf{z}} \tilde{f}_t(\mathbf{z})$, which can be solved using `ONM` [26].

Additionally, there exists a semi-unitary matrix $\overline{\mathbf{F}}_t \in \mathbb{R}^{n\times(n-p)}$ such that $\mathcal{R}\left(\overline{\mathbf{F}}_t\right) = \mathcal{N}(\mathbf{A}_t)$. Without loss of generality, we let $\mathbf{F}_t = \overline{\mathbf{F}}_t$ in our analysis. It follows that for any $\mathbf{z}_t$ satisfying $\mathbf{x}_t = \overline{\mathbf{F}}_t\mathbf{z}_t + \hat{\mathbf{x}}$, we have $\|\mathbf{x}_t - \mathbf{x}_t^*\| = \left\|\overline{\mathbf{F}}_t(\mathbf{z}_t - \mathbf{z}_t^*)\right\| = \|\mathbf{z}_t - \mathbf{z}_t^*\|$. In other words, the norm of the original problem's `OEN` update and the reduced problem's `ONM` update coincide when $\mathbf{F}_t = \overline{\mathbf{F}}_t$.

4.2.3 Assumptions

We now introduce three recurring assumptions that are used throughout this work. These assumptions are mild and notably do not require the objective function to be convex. These assumptions must hold for all $t = 1, 2, 3, ..., T$.

Assumption 4.1. *There exists a constant $h > 0$ such that:*

$$\left\|\nabla^2 f_t(\mathbf{x}_t^*)^{-1}\right\| \leq \frac{1}{h}.$$

Assumption 4.2. *There exists non-negative finite constants $\beta > 0$ and $0 < L < +\infty$ such that:*

$$\|\mathbf{x} - \mathbf{x}_t^*\| \leq \beta \Rightarrow \left\|\nabla^2 f_t(\mathbf{x}) - \nabla^2 f_t(\mathbf{x}_t^*)\right\| \leq L \|\mathbf{x} - \mathbf{x}_t^*\|.$$

Assumption 4.3. *There exists* $0 < l < +\infty$ *such that:*

$$|f_t(\mathbf{x}) - f_t(\mathbf{x}_t^*)| \leq l \left\| \mathbf{x} - \mathbf{x}_t^* \right\|.$$

Assumption 4.1 imposes an upper bound on the norm of the inverse Hessian at the optimum. This implies that the Hessian's eigenvalues can be positive or negative but must be bounded away from zero. For convex loss functions, this is equivalent to strong convexity which is a common assumption in OCO [1, 26, 27]. Assumptions 4.2 and 4.3 define a local Lipschitz continuity condition on the Hessian and a Lipschitz condition on the objective function around the optimum, respectively.

4.2.4 Reduced function identities

We now provide two lemmas which characterize the reduced function $\tilde{f}_t$.

Lemma 4.1. *[29, Section 10.2.3] Suppose*

1. $\exists \mathbf{F}_t$ *such that* $\mathcal{N}(\mathbf{A}_t) = \mathcal{R}(\mathbf{F}_t)$;
2. $\exists \hat{\mathbf{x}}$ *such that* $\mathbf{A}_t\hat{\mathbf{x}} = \mathbf{b}_t$;
3. $\mathbf{A}_t\mathbf{x}_t - \mathbf{b}_t = \mathbf{0}$.

Consider the Newton step applied to the reduced function $\tilde{f}_t(\mathbf{z}_t)$*:* $\Delta\mathbf{z}_t = -\nabla^2\tilde{f}_t(\mathbf{z}_t)^{-1}\nabla\tilde{f}_t(\mathbf{z}_t)$*. Then the following identity holds :*

$$\Delta\mathbf{x}_t = \mathbf{F}_t\Delta\mathbf{z}_t. \tag{4.5}$$

Lemma 4.1 implies that the Newton step applied to the constrained problem coincides with the Newton step applied to the reduced problem. By setting $\mathbf{F}_t = \overline{\mathbf{F}}_t$, we obtain $\|\Delta\mathbf{x}_t\| = \|\Delta\mathbf{z}_t\|$.

The second lemma characterizes the local strong convexity and Lipschitz continuity of the reduced function.

Lemma 4.2. *Suppose Assumptions 4.1 and 4.2 hold. Then we have:*

$$\left\|\nabla^2\tilde{f}_t(\mathbf{z}_t^*)^{-1}\right\| \leq \frac{1}{\sigma_{\min}(\mathbf{F}_t)^2 h} \tag{4.6}$$

$$\|\mathbf{z} - \mathbf{z}_t^*\| \leq \frac{\beta}{\|\mathbf{F}_t\|} \Rightarrow \left\|\nabla^2\tilde{f}_t(\mathbf{z}) - \nabla^2\tilde{f}_t(\mathbf{z}_t^*)\right\| \leq L\left\|\mathbf{F}_t\right\|^3\left\|\mathbf{z} - \mathbf{z}_t^*\right\|, \tag{4.7}$$

where $\sigma_{\min}(\mathbf{F}_t)$ is the minimum singular value of $\mathbf{F}_t$.

Proof. Differentiating $\tilde{f}_t$ twice, we obtain,

$$\nabla^2 \tilde{f}_t(\mathbf{z}) = \mathbf{F}_t \nabla^2 f_t(\mathbf{x}_t^*) \mathbf{F}_t^\top.$$

The inverse Hessian of $\tilde{f}_t$ is thus upper bounded by:

$$\begin{aligned}
\left\| \nabla^2 \tilde{f}_t(\mathbf{z}^*)^{-1} \right\| &= \left\| \mathbf{F}_t (\mathbf{F}_t^\top \mathbf{F}_t)^{-1} \nabla^2 f_t(\mathbf{x}_t^*)^{-1} (\mathbf{F}_t^\top \mathbf{F}_t)^{-1} \mathbf{F}_t^\top \right\| \\
&\leq \left\| (\mathbf{F}_t^\top \mathbf{F}_t)^{-1} \mathbf{F}_t^\top \right\|^2 \left\| \nabla^2 f_t(\mathbf{x}_t^*)^{-1} \right\| \\
&\leq \frac{1}{\sigma_{\min}(\mathbf{F}_t)^2 h},
\end{aligned}$$

which is (4.6).
For (4.7), we have,

$$\begin{aligned}
\left\| \nabla^2 \tilde{f}_t(\mathbf{z}) - \nabla^2 \tilde{f}_t(\mathbf{z}_t^*) \right\| &= \left\| \mathbf{F}_t \Big(\nabla^2 f_t(\mathbf{F}_t \mathbf{z} + \hat{\mathbf{x}}) - \nabla^2 f_t(\mathbf{F}_t \mathbf{z}_t^* + \hat{\mathbf{x}}) \Big) \mathbf{F}_t^\top \right\| \\
&\leq \|\mathbf{F}_t\|^2 \left\| \nabla^2 f_t(\mathbf{F}_t \mathbf{z} + \hat{\mathbf{x}}) - \nabla^2 f_t(\mathbf{F}_t \mathbf{z}_t^* + \hat{\mathbf{x}}) \right\| \\
&\leq \|\mathbf{F}_t\|^2 L \|\mathbf{F}_t(\mathbf{z} - \mathbf{z}_t^*)\| \\
&\leq L \|\mathbf{F}_t\|^3 \|\mathbf{z} - \mathbf{z}_t^*\|,
\end{aligned}$$

where the last inequalities follow from the Lipschitz continuity of $\nabla^2 f_t$ and the definition of $\mathbf{z}_t$, respectively. □

4.2.5 Feasible Newton update

Using Lemmas 4.1 and 4.2, we derive the following lemma for the original constrained problem:

Lemma 4.3 (Equality-constrained Newton identities)**.** *Suppose Assumptions 4.1, 4.2 hold and:*

1. $\|\mathbf{x}_t - \mathbf{x}_t^*\| < \min\left\{\beta, \frac{h}{2L}\right\}$;
2. $\mathbf{A}_t \mathbf{x}_t - \mathbf{b}_t = \mathbf{0}$.

Then we have the following two identities for `OEN`*:*

$$\|\mathbf{x}_{t+1} - \mathbf{x}_t^*\| < \|\mathbf{x}_t - \mathbf{x}_t^*\|\,; \tag{4.8}$$

$$\|\mathbf{x}_{t+1} - \mathbf{x}_t^*\| \leq \frac{2L}{h}\,\|\mathbf{x}_t - \mathbf{x}_t^*\|^2\,. \tag{4.9}$$

The first inequality guarantees the next iterate is strictly closer to the optimum compared to the current iterate. The second inequality provides an upper bound on this value.

Proof. By the definition of the `OEN` update and Lemma 4.1 we have:

$$\mathbf{x}_{t+1} - \mathbf{x}_t^* = \mathbf{x}_t - \mathbf{x}_t^* - \mathbf{F}_t \nabla^2 \tilde{f}_t(\mathbf{z}_t)^{-1} \nabla \tilde{f}_t(\mathbf{z}_t).$$

Rearranging and letting $\mathbf{F}_t = \overline{\mathbf{F}}_t$, we have

$$\begin{aligned}
\mathbf{x}_{t+1} - \mathbf{x}_t^* &= \mathbf{x}_t - \mathbf{x}_t^* - \overline{\mathbf{F}}_t \nabla^2 \tilde{f}_t(\mathbf{z}_t)^{-1}\Big(\nabla \tilde{f}_t(\mathbf{z}_t) - \nabla \tilde{f}_t(\mathbf{z}_t^*)\Big) \\
&= \mathbf{x}_t - \mathbf{x}_t^* - \overline{\mathbf{F}}_t \nabla^2 \tilde{f}_t(\mathbf{z}_t)^{-1} \cdot \int_0^1 \nabla^2 \tilde{f}_t\Big(\mathbf{z}_t + \tau(\mathbf{z}_t^* - \mathbf{z}_t)\Big)(\mathbf{z}_t^* - \mathbf{z}_t)\mathrm{d}\tau.
\end{aligned}$$

From the symmetry of the Hessian and its inverse we have,

$$\begin{aligned}
\mathbf{x}_{t+1} - \mathbf{x}_t^* &= \mathbf{x}_t - \mathbf{x}_t^* - \overline{\mathbf{F}}_t(\mathbf{z}_t^* - \mathbf{z}_t)\nabla^2 \tilde{f}_t(\mathbf{z}_t)^{-1} \cdot \int_0^1 \nabla^2 \tilde{f}_t\Big(\mathbf{z}_t + \tau(\mathbf{z}_t^* - \mathbf{z}_t)\Big)\mathrm{d}\tau \\
&= (\mathbf{x}_t^* - \mathbf{x}_t)\nabla^2 \tilde{f}_t(\mathbf{z})^{-1} \cdot \int_0^1 \Big[\nabla^2 \tilde{f}_t\Big(\mathbf{z}_t + \tau(\mathbf{z}_t^* - \mathbf{z}_t)\Big) - \nabla^2 \tilde{f}_t(\mathbf{z}_t)\Big]\mathrm{d}\tau.
\end{aligned}$$

Taking the norm on both sides and using Lemmas 4.1 and 4.2 yields

$$\begin{aligned}
\|\mathbf{x}_{t+1} - \mathbf{x}_t^*\| &\leq \|\mathbf{x}_t^* - \mathbf{x}_t\| \left\|\nabla^2 \tilde{f}_t(\mathbf{z}_t)^{-1}\right\| \cdot \int_0^1 \left\|\nabla^2 \tilde{f}_t\Big(\mathbf{z}_t + \tau(\mathbf{z}_t^* - \mathbf{z}_t)\Big) + \nabla^2 \tilde{f}_t(\mathbf{z}_t)\right\| \mathrm{d}\tau \\
&\leq \|\mathbf{x}_t^* - \mathbf{x}_t\| \left\|\nabla^2 \tilde{f}_t(\mathbf{z}_t)^{-1}\right\| \cdot \int_0^1 2L \left\|\overline{\mathbf{F}}_t\right\|^2 \tau \,\|\mathbf{x}_t^* - \mathbf{x}_t\| \,\mathrm{d}\tau \\
&\leq \left\|\nabla^2 \tilde{f}_t(\mathbf{z}_t)^{-1}\right\| L \,\|\mathbf{x}_t^* - \mathbf{x}_t\|^2\,.
\end{aligned} \tag{4.10}$$

Using [26, Lemma 2], we can bound the inverse Hessian as:

$$\begin{aligned}
\left\|\nabla^2 \tilde{f}(\mathbf{z}_t)^{-1}\right\| &\leq \frac{1}{\sigma_{\min}(\overline{\mathbf{F}}_t)^2 h - L\left\|\overline{\mathbf{F}}_t\right\|^2 \|\mathbf{x}_t^* - \mathbf{x}_t\|} \\
&\leq \frac{1}{h - L\,\|\mathbf{x}_t^* - \mathbf{x}_t\|},
\end{aligned} \tag{4.11}$$

because $\sigma_{\min}(\overline{\mathbf{F}}_t) = \left\|\overline{\mathbf{F}}_t\right\| = 1$. Substituting (4.10) into (4.11) leads to:

$$\|\mathbf{x}_{t+1} - \mathbf{x}_t^*\| \leq \frac{L}{h - L\,\|\mathbf{x}_t^* - \mathbf{x}_t\|}\,\|\mathbf{x}_t^* - \mathbf{x}_t\|^2 \tag{4.12}$$

Finally, (4.8) and (4.9) follows from (4.12) and [26, Lemma 2]'s proof. □

4.3 Online Equality-constrained Newton's Method

We now present our online optimization methods for problems with time-independent and time-varying equality constraints.

4.3.1 Online Equality-constrained Newton's Method

In this section, we propose the Online Equality-constrained Newton's Method (`OEN-M`) for online optimization subject to time-invariant linear equality constraints. This is the first online, second-order algorithm that admits constraints. `OEN-M` is presented in Algorithm 1.

Algorithm 1: Online Equality-constrained Newton's Method

Parameters: A, **b**
Initialization: Receive $\mathbf{x}_0 \in \mathbb{R}$ such that
$\|\mathbf{x}_0 - \mathbf{x}_0^*\| < \gamma = \min\left\{\beta, \frac{h}{2L}\right\}$ and $\mathbf{A}\mathbf{x}_0 - \mathbf{b} = \mathbf{0}$
for $t = 0, 1, 2...T$ **do**
  Play the decision $\mathbf{x}_t$.
  Observe the outcome $f_t(\mathbf{x}_t)$.
  Update decision:
  $\begin{bmatrix}\mathbf{x}_{t+1}\\ \boldsymbol{\nu}_t\end{bmatrix} = \begin{bmatrix}\mathbf{x}_t\\ \mathbf{0}\end{bmatrix} - \mathbf{D}_t(\mathbf{x}_t)^{-1}\begin{bmatrix}\nabla f_t(\mathbf{x}_t)\\ \mathbf{0}\end{bmatrix}.$
end for

We now show that `OEN-M` has a dynamic regret bounded above by $O(V_T + 1)$ and a null constraint violation.

Theorem 4.1. *If Assumptions 4.1–4.3 hold and the following conditions are respected:*

1. $\exists\mathbf{x}_0$ *such that* $\|\mathbf{x}_0 - \mathbf{x}_0^*\| < \gamma = \min\{\beta, \frac{h}{2L}\}$;
2. $\overline{v} < \gamma - \frac{2L}{h}\gamma^2$.

Then the regret $R_d(T)$ and the constraint violation $Vio(T)$ are bounded above by :

$$R_d(T) \leq \frac{lh}{h - 2L\gamma}(V_T + \delta); \tag{4.13}$$

$$Vio(T) = 0, \tag{4.14}$$

where $\delta = \frac{2L}{h}\gamma(\|\mathbf{x}_0 - \mathbf{x}_0^*\| - \|\mathbf{x}_T - \mathbf{x}_T^*\|)$.

Remark 4.1. *The assumption that the decision-maker has access to $\mathbf{x}_0$ such that $\|\mathbf{x}_0 - \mathbf{x}_0^*\| \leq \gamma$ is standard in OCO [9, 18, 26]. Essentially, this means obtaining a good starting estimate of the initial optimal solution is required. It is assumed that a good estimate can be obtained before the start of the online process from, for example, offline calculations or a previously implemented decision.*

Proof. Using Assumption 4.3, the regret is bounded by:

$$R_{\mathrm{d}}(T) = \sum_{t=1}^{T} \left| f_t(\mathbf{x}_t) - f_t(\mathbf{x}_t^*) \right| \leq l \sum_{t=1}^{T} \|\mathbf{x}_t - \mathbf{x}_t^*\| \,. \tag{4.15}$$

Rearranging (4.15)'s sum we obtain:

$$\begin{aligned}
\sum_{t=1}^{T} \|\mathbf{x}_t - \mathbf{x}_t^*\| &= \sum_{t=1}^{T} \left\| \mathbf{x}_t - \mathbf{x}_{t-1}^* + \mathbf{x}_{t-1}^* - \mathbf{x}_t^* \right\| \\
&\leq \sum_{t=1}^{T} \left\| \mathbf{x}_t - \mathbf{x}_{t-1}^* \right\| + \sum_{t=1}^{T} \left\| \mathbf{x}_t^* - \mathbf{x}_{t-1}^* \right\| \\
&\leq \sum_{t=0}^{T-1} \frac{2L}{h} \|\mathbf{x}_t - \mathbf{x}_t^*\|^2 + V_T \\
&\leq \sum_{t=1}^{T} \frac{2L}{h}\gamma \|\mathbf{x}_t - \mathbf{x}_t^*\| + V_T + \delta,
\end{aligned}$$

where $\delta = \frac{2L}{h}\gamma(\|\mathbf{x}_0 - \mathbf{x}_0^*\| - \|\mathbf{x}_T - \mathbf{x}_T^*\|)$. Lemma 4.3 is used to obtain the second inequality. Solving for $\|\mathbf{x}_t - \mathbf{x}_t^*\|$, we have:

$$\sum_{t=1}^{T} \|\mathbf{x}_t - \mathbf{x}_t^*\| \leq \left(1 - \frac{2L}{h}\gamma\right)^{-1} (V_T + \delta). \tag{4.16}$$

This implies that the dynamic regret is bounded above by:

$$R_{\mathrm{d}}(T) \leq \frac{lh}{h - 2L\gamma}(V_T + \delta),$$

and hence, $R_{\mathrm{d}}(T) \leq O(V_T + 1)$.

As for the constraint violation, we have that $\mathbf{x}_0$ is feasible by assumption. Because every `OEN` update is such that $\mathbf{A}\Delta\mathbf{x} = 0$, every subsequent decision will also be feasible. We thus have:

$$\mathrm{Vio}(T) = \sum_{t=1}^{T} \|\mathbf{A}\mathbf{x}_t - \mathbf{b}\| = 0,$$

which completes the proof. $\square$

4.3.2 Online Projected Newton's Method

We now consider online optimization problems with time-varying equality constraints. `OEN` does not apply to this class of problems because the previously played decision might not be feasible under the new constraints. We propose the Online Projected Equality-constrained Newton's Method (`OPEN-M`) to address this limitation. `OPEN-M` consists of a projection of the previous decision onto the new feasible set followed by an `OEN` step from this point. `OPEN-M` is detailed in Algorithm 2. We now analyse the performance of `OPEN-M`.

Algorithm 2: Online Projected Eq.-const. Newton's Method

Initialization: Receive $\mathbf{x}_0 \in \mathbb{R}$ such that $\|\mathbf{x}_0 - \mathbf{x}_0^*\| < \gamma = \min\left\{\beta, \frac{h}{2L}\right\}$
for $t = 0, 1, 2...T$ **do**
Play the decision $\mathbf{x}_t$.
Observe the outcome $f_t(\mathbf{x}_t)$ and constraints $\mathbf{A}_t, \mathbf{b}_t$.
Project $\mathbf{x}_t$ onto the feasible set:
$\tilde{\mathbf{x}}_t = \mathbf{x}_t + \mathbf{A}_t^\top(\mathbf{A}_t\mathbf{A}_t^\top)^{-1}(\mathbf{b}_t - \mathbf{A}_t\mathbf{x}_t)$.
Update decision:
$$\begin{bmatrix} \mathbf{x}_{t+1} \\ \boldsymbol{\nu}_t \end{bmatrix} = \begin{bmatrix} \tilde{\mathbf{x}}_t \\ \mathbf{0} \end{bmatrix} - \mathbf{D}_t(\tilde{\mathbf{x}}_t)^{-1} \begin{bmatrix} \nabla f_t(\tilde{\mathbf{x}}_t) \\ \mathbf{0} \end{bmatrix}.$$
end for

Theorem 4.2. *Suppose Assumptions 4.1–4.3 hold and:*

1. $\exists \mathbf{x}_0$ *such that* $\|\mathbf{x}_0 - \mathbf{x}_0^*\| < \gamma = \min\{\beta, \frac{h}{2L}\}$;
2. $\overline{v} < \gamma - \frac{2L}{h}\gamma^2$;
3. $\exists a > 0$ *such that* $\|\mathbf{A}_t\| \leq a \quad \forall t$.

Then the regret $R_d(T)$ and the constraint violation $Vio(T)$ of `OPEN-M` *is bounded above by :*

$$R_d(T) \leq \frac{lh}{h - 2L\gamma}(V_T + \delta) \tag{4.17}$$

$$Vio(T) \leq \frac{ah}{h - 2L\gamma}(V_T + \delta), \tag{4.18}$$

where $\delta = \frac{2L}{h}\gamma(\|\mathbf{x}_0 - \mathbf{x}_0^*\| - \|\mathbf{x}_T - \mathbf{x}_T^*\|)$

Proof. We first show that the following inequality holds:

$$\|\tilde{\mathbf{x}}_t - \mathbf{x}_t^*\| \leq \|\mathbf{x}_t - \mathbf{x}_t^*\| . \tag{4.19}$$

Since $\tilde{\mathbf{x}}_t$ is the projection of $\mathbf{x}_t$ onto the feasible set at time t, $\mathbf{x}_t - \mathbf{x}_t^*$ and $\tilde{\mathbf{x}}_t - \mathbf{x}_t$ are orthogonal. It follows that: $\|\tilde{\mathbf{x}}_t - \mathbf{x}_t^*\|^2 = \|\mathbf{x}_t - \mathbf{x}_t^*\|^2 - \|\tilde{\mathbf{x}}_t - \mathbf{x}_t\|^2$. Because $\|\tilde{\mathbf{x}}_t - \mathbf{x}_t\| \geq 0$, we have $\|\tilde{\mathbf{x}}_t - \mathbf{x}_t^*\| \leq \|\mathbf{x}_t - \mathbf{x}_t^*\|$, which is (4.19).

This means that $\|\tilde{\mathbf{x}}_t - \mathbf{x}_t^*\| \leq \gamma$ which means the projected decision $\tilde{\mathbf{x}}_t$ satisfies all the requirements for `OEN-M`. The same analysis as for `OEN-M` therefore holds for `OPEN-M` and the same regret bound is obtained, thus leading to (4.17).

As for the constraint violation, we recall:

$$\text{Vio}(T) = \sum_{t=1}^{T} \|\mathbf{A}_t\mathbf{x}_t - \mathbf{b}_t\| .$$

Using the fact that $\mathbf{x}_t^*$ is feasible for every timestep,

$$\begin{aligned}\text{Vio}(T) &= \sum_{t=1}^{T} \|\mathbf{A}_t\mathbf{x}_t - \mathbf{b}_t - (\mathbf{A}_t\mathbf{x}_t^* - \mathbf{b}_t)\| \\ &= \sum_{t=1}^{T} \|\mathbf{A}_t(\mathbf{x}_t - \mathbf{x}_t^*)\| .\end{aligned}$$

By the Cauchy-Swartz inequality and the bound on $\mathbf{A}_t$:

$$\begin{aligned}\mathrm{Vio}(T) &\leq \sum_{t=1}^{T} \|\mathbf{A}_t\| \, \|\mathbf{x}_t - \mathbf{x}_t^*\| \\ &\leq a \sum_{t=1}^{T} \|\mathbf{x}_t - \mathbf{x}_t^*\| \\ &\leq \frac{ah}{h - 2L\gamma}(V_T + \delta),\end{aligned}$$

where the last inequality follows from (4.16). This yields (4.18) and completes the proof. □

Remark 4.2. *Because a closed-form projection step is possible for* `OPEN-M`*, the algorithmic time-complexity is the same as for* `OEN-M`*. This is because the time-complexity of* `OEN-M` *and* `OPEN-M` *are dominated by the matrix inversion step which is $O(n^5 \log(n))$ in the general case. Thus, there is no additional burden to* `OPEN-M`*'s projection step which is $O(n^3)$ in the worst case.*

Remark 4.3. `OPEN-M` *possesses the tightest dynamic regret bounds of any previously proposed online equality-constrained algorithm in the literature [8–10, 13]. Under the standard assumption in OCO that the variation of optima V_T is sublinear [13], constraint violation and regret will be sublinear. The method is also parameter-free which eliminates the need for time-dependent step sizes and hyperparameter tuning, e.g., the step size in gradient-based methods. The time horizon during which the algorithm is used is also arbitrary and does not need to be defined before execution. These advantages provide ample justification for the additional complexity of the inversion step.*

4.4 Numerical Experiment

We now illustrate the performance of `OPEN-M` and use it in an optimal network flow problem. This type of problem can model electric distribution grids when line losses are considered as negligible [47, 48]. In this context, a convex, quadratic cost is most commonly used. Note that, `OPEN-M` is also applicable to non-convex cost functions.

Consider the network flow problem over a directed graph $\mathcal{G} = (\mathcal{M}, \mathcal{L})$ with nodes $\mathcal{M}$ and directed edges $\mathcal{L}$. At every timestep t, load (sink) nodes, $i \in \mathcal{D}$, require a power supply b_t^i and generator (source) nodes, $j \in \mathcal{P}$, can produce a positive quantity of power. The power is distributed through the edges of the graph. The decision variable is $\mathbf{x}_t$ and models the power flowing through each edge. Assuming no active power losses, the power balance at each node leads to the constraint: $\mathbf{A}\mathbf{x}_t = \mathbf{b}_t$ where $b_t^{(i)}$ is the power demand at each load node and:

$$\mathbf{A}_{(i,l)} = \begin{cases} 1, & \text{if edge } l \text{ enters node } i \\ -1, & \text{if edge } l \text{ leaves node } i \\ 0, & \text{else.} \end{cases}$$

A numerical example is provided next using a fixed, radial network composed of 15 nodes connected via 30 arcs. A single power source is located at the root of the network. Every node's load is chosen independently as: $b_t^i = \frac{\zeta}{\sqrt{t}} + 10$ where t denotes the round and ζ is uniformly sampled in $[0, 5]$. The cost function f_i for each arc i is convex and adopts the following form: $f_i(x) = \alpha_i \mathrm{e}^{\beta_i |x|}$. The parameters α_i and β_i are chosen following: $\alpha_i = \frac{\eta}{\sqrt{t}} + 1$ and $\beta_i = \frac{\gamma}{\sqrt{t}} + 2$ where t denotes the round and η and γ are uniformly sampled in $[0, 10]$ at every round. This loss function is chosen because it is harder to solve than a quadratic function and yet approximately models electric grid costs. The temporal dependence of the parameters ensures that the total variation of optima (V_T) is bounded and sublinear. The time horizon is set to $T = 2500$. The fixed nature of the network and the diagonal Hessian matrix means that the inversion step only has to be done once. Note that `OPEN-M` also admits time-varying network topologies, i.e., using $\mathbf{A}_t$ instead of $\mathbf{A}$.

We use `MOSP` from [10] and the model-based augmented Lagrangian method (`MALM`) from [13] as benchmarks to establish `OPEN-M`'s performance. Because these algorithms can only admit inequality constraints, the equality constraint $\mathbf{A}\mathbf{x}_t = \mathbf{b}_t$ is relaxed to $\mathbf{A}\mathbf{x}_t - \mathbf{b}_t \leq 0$. This ensures that there is enough power at every node but lets the operator over-serve loads. This relaxation is mild because the constraint should be active at the optimum given that costs are minimized. Dynamic regret, defined in (4.2), and constraint violation, defined in (4.3), for this problem are presented in Figures 4.1 and 4.2, respectively.

We observe sublinear dynamic regret and constraint violation from all three algorithms illustrating they are well-adapted to this problem. We remark that `OPEN-M` exhibits a lower regret than both the `MOSP` and `MALM` algorithms. Indeed, the dynamic regret of `OPEN-M` is multiple order of magnitude smaller. `OPEN-M` also has significantly better constraint violation compared to the other two algorithms.

4.5 Conclusion

In this paper, a second-order approach for online constrained optimization is developed. Under linear time-varying equality constraints, the resulting algorithm, `OPEN-M`, achieves simultaneous $O(V_T + 1)$ dynamic regret and constraint violation bounds. These bounds are

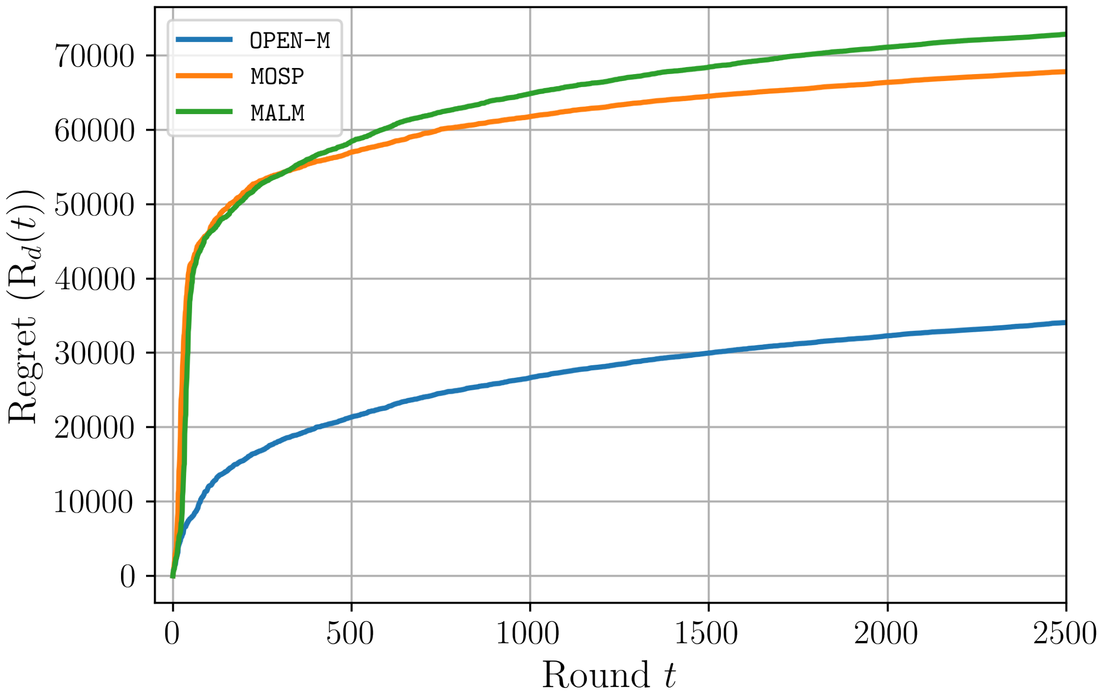

Figure 4.1 Dynamic regret comparison of `OPEN-M`, `MOSP` and `MALM`

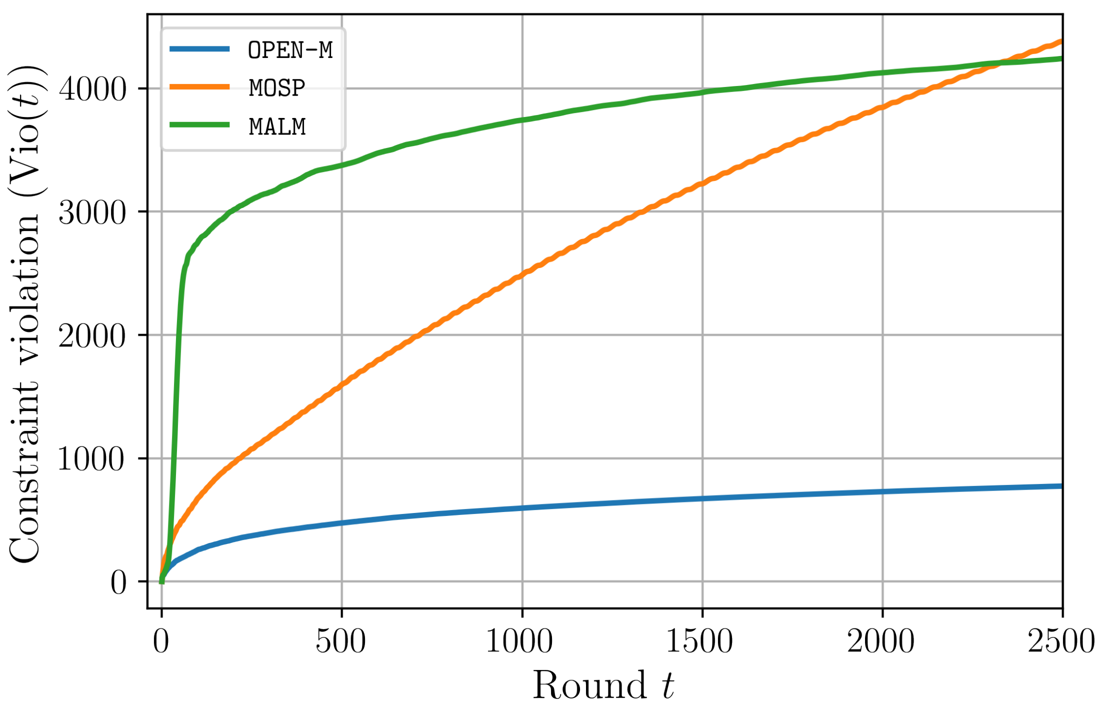

Figure 4.2 Constraint violation comparison of `OPEN-M`, `MOSP` and `MALM`

the tightest yet presented in the literature. A numerical network flow example is presented to showcase the performance of `OPEN-M` compared to other methods from the literature.

Considering the prevalence of interior-point methods in the offline optimization literature, an extension of the equality-constrained Newton's method which admits inequality constraints [32], a similar extension can be envisioned for `OPEN-M`. A second-order approach to an online optimization problem with time-varying inequalities has the potential to improve current dynamic regret and constraint violation bounds.

CHAPTER 5 ARTICLE 2 : ONLINE INTERIOR-POINT METHODS FOR TIME-VARYING EQUALITY-CONSTRAINED OPTIMIZATION

Jean-Luc Lupien • Iman Shames • Antoine Lesage-Landry

Submitted to: IEEE Transactions on Automatic Control, 18 July 2023 [49].

Abstract

An important challenge in the online convex optimization (OCO) setting is to incorporate generalized inequalities and time-varying constraints. The inclusion of constraints in OCO widens the applicability of such algorithms to dynamic but safety-critical settings such as the online optimal power flow (OPF) problem. In this work, we propose the first projection-free OCO algorithm admitting time-varying linear constraints and convex generalized inequalities: the online interior-point method for time-varying equality constraints (`OIPM-TEC`). We derive simultaneous sublinear dynamic regret and constraint violation bounds for `OIPM-TEC` under standard assumptions. For applications where a given tolerance around the optima is accepted, we propose a new OCO performance metric – the epsilon-regret – and a more computationally efficient algorithm, the epsilon `OIPM-TEC`, that possesses sublinear bounds under this metric. Finally, we showcase the performance of these two algorithms on an online OPF problem and compare them to another OCO algorithm from the literature.

Keywords. Online convex optimization, Optimization algorithms, Time-varying systems, Machine learning.

5.1 Introduction

Online convex optimization (OCO) with time-varying constraints, henceforth referred to as constrained OCO, is an emerging field with many practical applications including network resource allocation [50], portfolio selection [2], and control of electric grids [35] or of multi-energy systems [45]. In OCO, an agent makes sequential decisions to minimize an environment-determined loss function [1–3]. Constrained OCO adds the requirement that the decision sequence is subject to time-varying constraints. Integrating online or time-varying constraints is essential to ensure effective implementation of OCO algorithms in safety-critical applications. The performance of a constrained OCO algorithm is measured both in terms of its regret and its constraint violation [1, 9]. These represent the cumulative distance from optimality and from feasibility, respectively. Regret can be calculated either by comparing the decision sequence to the best fixed decision in hindsight (static regret) or to the round-optimal decision (dynamic regret) [4]. The dynamic regret is a tighter metric than the static regret as a bounded dynamic regret guarantees a bounded static regret [10]. When designing an OCO algorithm, the objective is to obtain sublinear regret and constraint violation bounds which imply the time-averaged regret and constraint violation will go to zero as the time-horizon goes to infinity [1, 2, 10]. This means implemented decisions will be optimal and feasible, on average, over a long time horizon.

A promising application of constrained OCO is the online optimal power flow (OPF) problem. The OPF problem consists of finding the optimal power generation profile that minimizes costs for the operator of the electric grid while ensuring that all loads are adequately supplied and that physical and safety limits are respected [38]. This problem is non-convex and thus NP-hard to solve. For this reason, multiple convex relaxations of the OPF have been formulated to compute an approximated solution in polynomial time [37, 38, 51]. However, the tightest convex relaxations are conic relaxations that leverage generalized inequalities [37]. This makes solving these problems tedious and time-consuming in large-scale systems. In modern power grids, the integration of renewable energy sources introduces power fluctuations on a significantly faster timescale when compared to conventional generators [44]. This makes the traditional approach of solving the OPF every 15-30 minutes [35, 39] unable to conform to the renewable timescale. The efficiency, adaptability, and resiliency of constrained OCO algorithms can overcome this limitation.

However, solving a conic relaxation of the OPF problem in an OCO context is challenging for multiple reasons. First and foremost is the integration of time-varying equality constraints linked to changes in power demand. Addressing the generalized inequalities associated with the conic constraints is also an obstacle. Many OCO algorithms use a projection step [8,10,13] onto the feasible space defined by these constraints. This is problematic because when the feasible space is the cone of positive semi-definite (PSD) matrices or a second-order cone – as is the case in the relaxed OPF – the projection can be as costly as solving the offline problem. Additionally, the less stringent static regret is ill-adapted to an online OPF setting. A fixed solution satisfying the constraints at all timesteps may incur arbitrarily large losses or might not exist at all [10]. Therefore, a candidate OCO algorithm for the online OPF is *one admitting time-varying equality constraints, being projection-free, and possessing a sublinear* *dynamic regret and constraint violation.* We propose an online optimization approach based on a primal interior-point method that possesses these three qualities.

Related work Although many constrained OCO approaches can be found in the literature, none possesses all three required properties necessary to effectively solve the online OPF. The majority of OCO algorithms that admit time-varying constraints use static regret as a performance metric. This metric uses the best fixed decision in hindsight as the comparator sequence. A common approach in this setting is to use a virtual queue to ensure a bounded static regret while satisfying long-term or slowly varying constraints. This approach is equivalent to solving a modified Lagrangian system. Virtual queues are used in [11, 12, 18, 27] to achieve sublinear static regret and constraint violation bounds. The model-based augmented Lagrangian method (`MALM`) presented in [13], which possesses sublinear regret and constraint violation bounds, has been shown to experimentally outperform many of these previous methods. A projection-free OCO algorithm utilising barrier functions as regularizers is presented in [33] which achieves sublinear static regret. This approach does not admit time-varying equality constraints or a dynamic regret bound.

Algorithms with dynamic regret bounds have also been developed to handle time-varying constraints. An online saddle-point method (`MOSP`) proposed in [10] and an approach using virtual queues presented in [8] possess sublinear dynamic regret and constraint violation bounds. These bounds are, however, predicated on strict limits on objective and constraint function variations as well as time-varying stepsizes. Both these algorithms rely on a costly projection step for conic constraints. Sublinear dynamic regret and constraint violation is obtained in [9] even for non-convex functions. The bounds are looser than those of [8,10] and the algorithm requires a projection at every iteration. A primal-dual mirror descent approach was proposed in [17] for distributed OCO and analyzed using dynamic regret but the authors only obtained constraint violation bounds that grow linearly with time. A second-order method was presented in [26] for unconstrained optimization which was extended in [43] to admit time-varying equality constraints. Both these approaches have tight dynamic regret bounds but are incompatible with inequality constraints.

Approaches to iteratively solve the OPF have also been presented. An algorithm based on a quasi-Newton update was used to solve the OPF in real-time in [34]. Although this algorithm has good tracking performance under strict conditions, the authors did not derive a bound on constraint violation. A distributed feedback controller is designed in [52] but uses a linear approximation of the OPF rather than a conic relaxation. A model-free online feedback optimization algorithm is presented in [53]. This approach is shown to converge to optimal solutions of the OPF problem but does not have theoretical guarantees on tracking or constraint violation.

In this paper, we present the first constrained OCO algorithm capable of satisfying all three necessary conditions for effective resolution of the online OPF. This is achieved by extending the `OPEN-M` algorithm from [43] to integrate generalized inequalities using an interior-point method-like update. Our contributions can be summarized as follows:

- We present the first projection-free OCO algorithm that admits convex generalized inequalities and time-varying linear equality constraints while achieving tight sublinear dynamic regret and constraint violation bounds.
- We introduce the notion of epsilon-regret in which one accepts being in an epsilon-neighbourhood of the optimal solution and propose an efficient algorithm possessing tight epsilon-regret and constraint violation bounds.
- We illustrate the performance of our approaches by solving an online OPF based on the 33-bus Matpower case.

5.2 Preliminaries

In this section, we define useful terms, lay down important assumptions, and introduce the Newton's step which will be used in the analysis of our main results.

5.2.1 Definitions

In this work, our goal is to solve time-varying equality constrained OCO problems subject to generalized inequality constraints described in (5.1). The generalized inequality constraints can, for example, be defined with respect to the cone of positive semidefinite matrices or a second-order cone. Formally, let $\mathbf{c} \in \mathbb{R}^N$ be our objective vector with $N \in \mathbb{N}$ the dimension of the decision variable. Let $\mathbf{x}_t \in \mathbb{R}^N$ be the decision vector at any time $t = 0, 1, 2, ..., T$ where $T \in \mathbb{N}$ is the time horizon. Let K_i, $i = 1, 2, ..., I$, $I \in \mathbb{N}$ be a proper cone as defined in [29, Section 3.6.2]. All functions $g_i(\mathbf{x}) : \mathbb{R}^N \mapsto \mathbb{R}^{K_i}$, $i = 1, 2, ..., I$, are assumed to be convex with respect to K_i and self-concordant. Self-concordancy is defined in (5.4) and is a common assumption in the interior-point literature [29, 32]. Let the matrix $\mathbf{A} \in \mathbb{R}^{P \times N}$ be a full row-rank matrix with $P \in \mathbb{N}$, $P < N$. Let the vector $\mathbf{b}_t \in \mathbb{R}^P$ be time-varying. In the sequel, the norm $\|\cdot\|$ denotes the Euclidean norm. In summary, we tackle the following online optimization problem:

$$\begin{aligned}\min_{\mathbf{x}_t} \quad & \mathbf{c}^\top \mathbf{x}_t \\ \text{s.t.} \quad & g_i(\mathbf{x}_t) \preceq_{K_i} 0, \quad \text{for } i = 1, 2, 3, ..., I \\ & \mathbf{A}\mathbf{x}_t - \mathbf{b}_t = 0.\end{aligned} \tag{5.1}$$

We remark that the only time-dependent parameter is the vector $\mathbf{b}_t$ in this formulation. Although this limits the scope of OCO problems that can be represented, many applications are adequately modelled in this manner. This is the case for the optimal power flow problem, for example, where the only time-dependent parameter are the loads that need to be supplied [37, 38]. We define the optimal decision, $\mathbf{x}_t^*$, to (5.1) at every timestep as:

$$\begin{aligned}\mathbf{x}_t^* \in \operatorname*{argmin}_{\mathbf{x}} \quad & \mathbf{c}^\top \mathbf{x} \\ \text{s.t.} \quad & g_i(\mathbf{x}) \preceq_{K_i} 0, \quad \text{for } i = 1, 2, 3, ..., I \\ & \mathbf{A}\mathbf{x} - \mathbf{b}_t = 0.\end{aligned}$$

To measure the performance of constrained OCO algorithms, two metrics are employed: regret and constraint violation. These represent the cumulative distance from optimality and from feasibility, respectively. In this work, dynamic regret is used which uses the round-optimal solution $\mathbf{x}_t^*$ as a comparator. This can be formally defined as:

$$R_{\mathrm{d}}(T) = \sum_{t=1}^{T} \mathbf{c}^\top \mathbf{x}_t - \mathbf{c}^\top \mathbf{x}_t^*. \tag{5.2}$$

Let $d_{K_i}(\mathbf{x}) = \min_{\tilde{\mathbf{x}}} \|\mathbf{x} - \tilde{\mathbf{x}}\|$ s.t. $g_i(\tilde{\mathbf{x}}) \preceq_{K_i} 0$, the closest distance from the point $\mathbf{x}$ to a point in the interior of the proper cone K_i. We can then define constraint violation as:

$$\mathrm{Vio}(T) = \sum_{t=1}^{T} \left[\sum_i d_{K_i}(\mathbf{x}_t) + \|\mathbf{A}\mathbf{x}_t - \mathbf{b}_t\| \right]. \tag{5.3}$$

Constraint violation at t is null if decision $\mathbf{x}_t$ is feasible and positive otherwise. We define the cumulative variation between optimal decisions V_T as:

$$V_T = \sum_{t=1}^{T} \left\| \mathbf{x}_t^* - \mathbf{x}_{t-1}^* \right\|.$$

Similarly, we define the cumulative variation in the right-hand term of the time-varying equality constraint $\mathbf{b}$ as:

$$V_{\mathbf{b}} = \sum_{t=1}^{T} \|\mathbf{b}_t - \mathbf{b}_{t-1}\| .$$

Let $\|\cdot\|_{\mathbf{x}}$ denote the Hessian norm as in [32]. For any vector $\mathbf{z} \in \mathbb{R}^N$ in the domain of a function $f(\mathbf{x}) \in \mathbb{R}^N \mapsto \mathbb{R}$, the Hessian norm is:

$$\|\mathbf{z}\|_{\mathbf{x}}^2 = \mathbf{z}^\top \nabla^2 f(\mathbf{x})\mathbf{z}.$$

The Hessian norm will be used extensively in the analysis of our proposed constrained OCO algorithms.

An important property of functions that guarantees convergence of interior-point methods is self-concordance [29, 32]. Self-concordance can be defined as all functionals for which any $\mathbf{x}_1$, $\mathbf{x}_2 \in \mathbb{R}^N$ such that $\|\mathbf{x}_2 - \mathbf{x}_1\|_{\mathbf{x}_1} < 1$, satisfy the following condition:

$$\frac{\|\mathbf{v}\|_{\mathbf{x}_2}}{\|\mathbf{v}\|_{\mathbf{x}_1}} \leq \frac{1}{1 - \|\mathbf{x}_2 - \mathbf{x}_1\|_{\mathbf{x}_1}}, \tag{5.4}$$

for any non-zero vector $\mathbf{v} \in \mathbb{R}^N$. Next, we define the barrier functional $\phi(\mathbf{x}) \in \mathbb{R}^N \mapsto \mathbb{R}$ as follows:

$$\phi(\mathbf{x}) = \sum_{i=1}^{I} -\log\Big(-g_i(\mathbf{x})\Big).$$

Here the logarithm is either the natural logarithm if g_i is a scalar function or a generalized logarithm as defined in [29, Section 11.6]. This barrier functional is convex and self-concordant over its domain [29].

Interior-point methods do not attempt to solve (5.1) directly. Instead, a sequence of convex optimization problems related to the original problem are solved. This is achieved in the following way. Let $\eta \in \mathbb{R}^+$ be a positive barrier parameter. This barrier parameter is useful in defining the following equality-constrained problem:

$$\begin{aligned} \min_{\mathbf{x}_t} \quad & \eta \mathbf{c}^\top \mathbf{x}_t + \phi(\mathbf{x}_t) \\ \text{s.t.} \quad & \mathbf{A}\mathbf{x}_t - \mathbf{b}_t = 0. \end{aligned} \tag{5.5}$$

Problem (5.5) is of interest because when $\eta \to \infty$, the solutions of (5.1) and (5.5) coincide. The strategy is, therefore, to incrementally increase η, iteratively solving the original problem.

In the sequel, we use the following notation:

$$d(\mathbf{x}_t, \eta) \stackrel{\text{def}}{=} \eta \mathbf{c}^\top \mathbf{x}_t + \phi(\mathbf{x}_t). \tag{5.6}$$

This represents the objective function we seek to minimize at every timestep t given a barrier parameter η.

Finally, we denote the feasible set by $D_f \stackrel{\text{def}}{=} \left\{\mathbf{x} \in \mathbb{R}^N | g_i(\mathbf{x}) \preceq 0, \forall i, \mathbf{A}\mathbf{x} - \mathbf{b}_t = 0\right\}$.

5.2.2 Assumptions

Certain sufficient conditions need to be met to ensure the regret and constraint violation suffered by the algorithms are sublinear. These are standard in the interior-point literature [51] and will hold for any convex relaxation of the OPF problem [38].

1. The relative interior of the feasible set D_f forms a closed and convex set.

2. The feasible set D_f is non-empty. This implies there exists at least one feasible point and that Slater's condition holds.

3. The barrier functional $\phi(\mathbf{x})$ is a self-concordant barrier functional of complexity $v_f \in \mathbb{R}^+$. That is,

$$\sup_{\mathbf{x}} \quad \nabla\phi(\mathbf{x}) \nabla^2 \phi^{-1}(\mathbf{x}) \nabla\phi(\mathbf{x}) \leq v_f. \tag{5.7}$$

   The complexity of standard barrier functionals can be reasonably bounded by the dimension of the decision space, i.e., $v_f = O(N)$ [29]. For example, the barrier functionals used for the space of positive semi-definite (PSD) matrices, second-order cones, quadratic functions, and linear functions all satisfy (5.7) [32, Section 2.3].

4. For $\mathbf{x} \in D_f$ we have:

$$\left\| \begin{bmatrix} \nabla^2\phi(\mathbf{x}) & \mathbf{A}^\top \\ \mathbf{A} & 0 \end{bmatrix} \right\| \leq \frac{1}{m}.$$

   This is a similar assumption to [29, Chapter 10].

We assume that Assumptions 1-4 hold hereinafter.

5.2.3 Newton Step

To define the Newton's step, we use the following formalism:

$$\begin{aligned}
\mathbf{y}_t &= \begin{bmatrix} \mathbf{x}_t \\ \boldsymbol{\nu}_t \end{bmatrix} \\
r_t(\mathbf{y}_t, \eta) &= \begin{bmatrix} \nabla d(\mathbf{x}_t, \eta) + \mathbf{A}^\top \boldsymbol{\nu}_t \\ \mathbf{A}\mathbf{x}_t - \mathbf{b}_t \end{bmatrix} \\
\mathbf{D}(\mathbf{y}_t) &= \begin{bmatrix} \nabla^2 d(\mathbf{x}_t) & \mathbf{A}^\top \\ \mathbf{A} & 0 \end{bmatrix} = \begin{bmatrix} \nabla^2 \phi(\mathbf{x}_t) & \mathbf{A}^\top \\ \mathbf{A} & 0 \end{bmatrix}.
\end{aligned}$$

We notice that $\mathbf{D}$ is the gradient of r and is independent of both time t and barrier parameter η.

The infeasible Newton's step, $n_t(\mathbf{y}_t, \eta)$, – so-called because it does not require a feasible initial point – at round t associated to a barrier parameter η is defined as:

$$n_t(\mathbf{y}_t, \eta) = \Delta \mathbf{y}_t = -\mathbf{D}(\mathbf{y}_t)^{-1} r_t(\mathbf{y}_t, \eta),$$

where t denotes the current round.

We now present a sequence of lemmas that will later be used to to prove our main results in Section 5.3.

Lemma 5.1. *Consider the following functional:*

$$\Gamma(\mathbf{y}) = \Gamma(\mathbf{x}, \boldsymbol{\nu}) = d_\eta(\mathbf{x}) + \boldsymbol{\nu}^\top \mathbf{A}\mathbf{x} - \boldsymbol{\nu}^\top \mathbf{b}_t,$$

where $d_\eta(\mathbf{x}) = d(\mathbf{x}, \eta)$ *and* η *is set to a fixed constant. Then* Γ *is self-concordant and* $\nabla \Gamma(\mathbf{y}) = r(\mathbf{y})$ *and* $\nabla^2 \Gamma(\mathbf{y}) = \mathbf{D}(\mathbf{y})$.

Proof. The third derivative of $\Gamma(\mathbf{y})$ coincides with the third derivative of $d_\eta(\mathbf{x})$ having no dependency in $\boldsymbol{\nu}$. Because the third derivative of $d_\eta(\mathbf{x})$ is bounded by assumption, this provides a bound on the third derivative of $\Gamma(\mathbf{y})$. By [32, Section 2.5], we have that $\Gamma(\mathbf{y})$ is self-concordant. □

Lemma 5.2. *Let* $\Gamma(\mathbf{y}) = d(\mathbf{x}, \eta) + \boldsymbol{\nu}^\top \mathbf{A}\mathbf{x} - \boldsymbol{\nu}^\top \mathbf{b}_t$ *be a self-concordant function. The following holds:*

$$\left\| \mathbf{D}^{-1}(\mathbf{y}_1) \mathbf{D}(\mathbf{y}_2) \right\|_{\mathbf{y}_1} \leq \frac{1}{(1 - \|\mathbf{y}_1 - \mathbf{y}_2\|_{\mathbf{y}_1})^2}.$$

Proof. The inequalities follow directly from the definition of self-concordance [32]. We have:

$$\begin{aligned}\left\|\mathbf{D}^{-1}(\mathbf{y}_1)\mathbf{D}(\mathbf{y}_2)\right\|_{\mathbf{y}_1} &= \max_{\mathbf{v}} \frac{\mathbf{v}^\top \mathbf{D}(\mathbf{y}_1)\mathbf{D}^{-1}(\mathbf{y}_1)\mathbf{D}(\mathbf{y}_2)\mathbf{v}}{\|\mathbf{v}\|_{\mathbf{y}_1}^2} \\ &= \max_{\mathbf{v}} \frac{\|\mathbf{v}\|_{\mathbf{y}_2}^2}{\|\mathbf{v}\|_{\mathbf{y}_1}^2} \leq \left(\frac{1}{1-\|\mathbf{y}_2-\mathbf{y}_1\|_{\mathbf{y}_1}}\right)^2,\end{aligned}$$

and we have established the inequality. □

The dual-feasible nature of every point on the central path provides an estimation of the error for each iteration. Defining the next iterate as $\mathbf{y}^+ = \mathbf{y} + n_t(\mathbf{y}, \eta)$, Lemma 5.3 formalizes this concept.

Lemma 5.3. *For any pair $\{\mathbf{y}, \eta\}$ such that $\|n_t(\mathbf{y}, \eta)\|_{\mathbf{y}} \leq 1$, we have:*

$$\left\|n_t(\mathbf{y}^+, \eta)\right\|_{\mathbf{y}^+} \leq \frac{\|n_t(\mathbf{y}, \eta)\|_{\mathbf{y}}^2}{(1-\|n_t(\mathbf{y}, \eta)\|_{\mathbf{y}})^2}. \tag{5.8}$$

Proof. From the definition of the Newton's step, we have:

$$\|n_t(\mathbf{y}, \eta)\|_{\mathbf{y}}^2 = r_t(\mathbf{y}, \eta)^\top \mathbf{D}(\mathbf{y})^{-1} r_t(\mathbf{y}, \eta).$$

We can now express $\|n_t(\mathbf{y}^+, \eta)\|_{\mathbf{y}^+}$ as:

$$\begin{aligned}\left\|n_t(\mathbf{y}^+, \eta)\right\|_{\mathbf{y}^+}^2 &= r_t(\mathbf{y}^+, \eta)^\top \mathbf{D}(\mathbf{y}^+)^{-1} r_t(\mathbf{y}^+, \eta) \\ &= r_t(\mathbf{y}^+, \eta)^\top \mathbf{D}(\mathbf{y})^{-1}\mathbf{D}(\mathbf{y})\mathbf{D}(\mathbf{y}^+)^{-1} r_t(\mathbf{y}^+, \eta) \\ &\leq \left\|\mathbf{D}(\mathbf{y})\mathbf{D}(\mathbf{y}^+)^{-1}\right\|_{\mathbf{y}} r_t(\mathbf{y}^+, \eta)^\top \mathbf{D}(\mathbf{y})^{-1} r_t(\mathbf{y}^+, \eta) \\ &= \left\|\mathbf{D}(\mathbf{y})\mathbf{D}(\mathbf{y}^+)^{-1}\right\|_{\mathbf{y}} \left\|\mathbf{D}(\mathbf{y})^{-1} r_t(\mathbf{y}^+, \eta)\right\|_{\mathbf{y}}^2,\end{aligned}$$

where $\mathbf{D}(\mathbf{y})^\top = \mathbf{D}(\mathbf{y})$ because $\mathbf{D}$ is symmetric. Using Assumption 4, we can bound the first term by:

$$\begin{aligned}\left\|\mathbf{D}(\mathbf{y})\mathbf{D}(\mathbf{y}^+)^{-1}\right\|_{\mathbf{y}} &\leq \frac{1}{(1-\|\mathbf{y}^+ - \mathbf{y}\|_{\mathbf{y}})^2} \\ &= \frac{1}{(1-\|n_t(\mathbf{y}, \eta)\|_{\mathbf{y}})^2.}\end{aligned}$$

Therefore, we have that:

$$\left\| n_t(\mathbf{y}^+, \eta) \right\|_{\mathbf{y}^+} \leq \frac{\|\mathbf{D}(\mathbf{y})^{-1} r_t(\mathbf{y}^+, \eta)\|_{\mathbf{y}}}{1 - \|n_t(\mathbf{y}, \eta)\|_{\mathbf{y}}}. \tag{5.9}$$

We now bound the numerator of (5.9). We have:

$$\begin{aligned} \left\| \mathbf{D}(\mathbf{y})^{-1} r_t(\mathbf{y}^+, \eta) \right\|_{\mathbf{y}} &= \left\| \mathbf{D}(\mathbf{y})^{-1} r_t(\mathbf{y}^+, \eta) - \mathbf{D}(\mathbf{y})^{-1} r_t(\mathbf{y}, \eta) - n_t(\mathbf{y}, \eta) \right\|_{\mathbf{y}} \\ &\leq \int_0^1 \left\| \mathbf{D}(\mathbf{y})^{-1} \mathbf{D}\Big(\mathbf{y} - \tau n_t(\mathbf{y}, \eta)\Big) \right\|_{\mathbf{y}} \|n_t(\mathbf{y}, \eta)\|_{\mathbf{y}} \, \mathrm{d}\tau + \int_0^1 \|n_t(\mathbf{y}, \eta)\|_{\mathbf{y}} \, \mathrm{d}\tau. \end{aligned}$$

Using Lemma 5.2 yields:

$$\begin{aligned} \left\| \mathbf{D}(\mathbf{y})^{-1} r_t(\mathbf{y}^+, \eta) \right\|_{\mathbf{y}} &\leq \int_0^1 \frac{\|n_t(\mathbf{y}, \eta)\|_{\mathbf{y}}}{(1 - \tau \|n_t(\mathbf{y}, \eta)\|_{\mathbf{y}})^2} \mathrm{d}\tau + \|n_t(\mathbf{y}, \eta)\|_{\mathbf{y}} \\ &= \left[\frac{-1}{1 - \tau \|n_t(\mathbf{y}, \eta)\|_{\mathbf{y}}} \right]_0^1 + \|n_t(\mathbf{y}, \eta)\|_{\mathbf{y}} \\ &= \frac{\|n_t(\mathbf{y}, \eta)\|_{\mathbf{y}}^2}{(1 - \|n_t(\mathbf{y}, \eta)\|_{\mathbf{y}})}, \end{aligned} \tag{5.10}$$

Substituting (5.10) into (5.9) completes the proof. □

An important consequence of Lemma 5.3 is if we have an initial pair $\{\mathbf{y}, \eta\}$ such that $\|n_t(\mathbf{y}, \eta)\|_{\mathbf{y}} \leq \frac{1}{4}$, the next iterate satisfies the following inequality:

$$\left\| n_t(\mathbf{y}^+, \eta) \right\|_{\mathbf{y}^+} \leq \frac{1}{9}.$$

It is possible to determine a point's proximity to the optimum using the norm of the Newton's step computed from this point. This result is presented in Lemma 5.4.

Lemma 5.4. *For any pair $\{\mathbf{y}, \eta\}$ such that $\|n_t(\mathbf{y}, \eta)\|_{\mathbf{y}} \leq 1$ we have:*

$$\|\mathbf{y}^\eta - \mathbf{y}\|_{\mathbf{y}} \leq \|n_t(\mathbf{y}, \eta)\|_{\mathbf{y}} + \frac{3 \|n_t(\mathbf{y}, \eta)\|_{\mathbf{y}}^2}{(1 - \|n_t(\mathbf{y}, \eta)\|_{\mathbf{y}})^3},$$

where $\mathbf{y}^\eta = \begin{bmatrix} \mathbf{x}^\eta \\ \boldsymbol{\nu}^\eta \end{bmatrix}$ *denotes the optimal primal and dual variables such that:*

$$\begin{aligned} \mathbf{x}^\eta \in \operatorname*{argmin}_{\mathbf{x}} & \, \eta \mathbf{c}^\top \mathbf{x} - \phi(\mathbf{x}) \\ \text{s.t.} \quad & \eta \mathbf{c} - \nabla \phi(\mathbf{x}^\eta) + \mathbf{A}^\top \boldsymbol{\nu}^\eta = 0. \end{aligned}$$

Proof. The proof follows from Lemma 5.3 and the analysis from [32, Section 2.2.5]. □

In the specific case where $\|n_t(\mathbf{y}^+, \eta)\|_{\mathbf{y}^+} \leq \frac{1}{9}$, this implies:

$$\left\| \mathbf{y}^\eta - \mathbf{y}^+ \right\|_{\mathbf{y}^+} \leq \frac{1}{6}.$$

Lemma 5.5. *If* $\|\mathbf{y}^\eta - \mathbf{y}\|_{\mathbf{y}} \leq \frac{1}{6}$*, where* $\mathbf{y}$ *is a feasible point, then:*

$$\|\mathbf{y}^\eta - \mathbf{y}\|_{\mathbf{y}^\eta} \leq \frac{1}{5}.$$

Proof. From the definition of self-concordance we have:

$$\|\mathbf{y}^\eta - \mathbf{y}\|_{\mathbf{y}^\eta} \leq \frac{\|\mathbf{y} - \mathbf{y}^\eta\|_{\mathbf{y}}}{1 - \|\mathbf{y} - \mathbf{y}^\eta\|_{\mathbf{y}}} = \frac{\frac{1}{6}}{\frac{5}{6}} = \frac{1}{5}.$$

□

Lemma 5.6. *Consider problem* (5.5)*. Let* $\Delta \mathbf{b}_t \stackrel{def}{=} \mathbf{b}_{t+1} - \mathbf{b}_t$ *be the change in the online parameter* $\mathbf{b}_t$ *between rounds. If* $\|n_t(\mathbf{y}, \eta)\| \leq \frac{1}{9}$ *and* $\|\Delta \mathbf{b}_t\| \leq \sqrt{\frac{3m}{160}}$ *for all* t*, then:*

$$\|n_{t+1}(\mathbf{y}, \eta)\| \leq \frac{1}{4}.$$

Proof. First, we remark that the only change between rounds is the parameter $\mathbf{b}_t$. Therefore,

$$r_{t+1}(\mathbf{y}, \eta) = r_t(\mathbf{y}, \eta) + \begin{bmatrix} \mathbf{0} \\ \Delta \mathbf{b}_t \end{bmatrix}.$$

Let $\Delta\mathbf{r}_t = \begin{bmatrix}\mathbf{0}\\ \Delta\mathbf{b}\end{bmatrix}$. Thus, $\|\Delta\mathbf{r}_t\| = \|\Delta\mathbf{b}_t\|$. The Hessian norm of the Newton's step becomes:

$$
\begin{aligned}
\|n_{t+1}(\mathbf{y},\eta)\|_{\mathbf{y}}^2 &= (r_t(\mathbf{y},\eta)+\Delta\mathbf{r})^\top \mathbf{D}(\mathbf{y})^{-1}(r_t(\mathbf{y},\eta)+\Delta\mathbf{r}_t)\\
&\leq 2r_t(\mathbf{y},\eta)^\top \mathbf{D}(\mathbf{y})^{-1} r_t(\mathbf{y},\eta) + 2\Delta\mathbf{r}_t^\top \mathbf{D}(\mathbf{y})^{-1}\Delta\mathbf{r}_t\\
&= \frac{2}{81} + 2\left\|\Delta\mathbf{b}_t\right\|^2 \left\|\mathbf{D}(\mathbf{y})^{-1}\right\|\\
&= \frac{2}{81} + \frac{2\left\|\Delta\mathbf{b}_t\right\|^2}{m} \leq \frac{1}{16},
\end{aligned}
$$

where we used Assumption 4 and the bound on $\|\Delta\mathbf{b}\|$ to obtain the final inequality. □

From [32, Section 2.3.1], we have that the restriction of any barrier functional to a subspace or translation of a subspace results in a barrier functional whose complexity barrier is smaller than that of the original functional. This implies that if $\|\phi(\mathbf{x})\|_{\mathbf{x}}^2 \leq \upsilon_f$ any feasible pair $\mathbf{x},\boldsymbol{\nu}$ satisfies:

$$
\begin{bmatrix}\nabla\phi(\mathbf{x}) + \mathbf{A}^\top\boldsymbol{\nu}\\ \mathbf{0}\end{bmatrix}^\top \mathbf{D}(\mathbf{x})^{-1} \begin{bmatrix}\nabla\phi(\mathbf{x}) + \mathbf{A}^\top\boldsymbol{\nu}\\ \mathbf{0}\end{bmatrix} \leq \upsilon_f.
$$

Defining

$$
h(\mathbf{y}) \stackrel{\text{def}}{=} \mathbf{D}(\mathbf{x})^{-1}\begin{bmatrix}\nabla\phi(\mathbf{x}) + \mathbf{A}^\top\boldsymbol{\nu}\\ \mathbf{0}\end{bmatrix},
$$

we notice that $\|h(\mathbf{y})\|_{\mathbf{y}} \leq \sqrt{\upsilon_f}$. We will use this result in Lemma 5.7.

Lemma 5.7. *For a feasible pair $\{\mathbf{y},\eta\}$ where $\|n_t(\mathbf{y},\eta)\|_{\mathbf{y}} \leq \frac{1}{9}$, there exists a parameter $1 < \beta \leq 1 + \frac{1}{8\sqrt{\upsilon_f}}$ such that for $\eta^+ = \beta\eta$, we have:*

$$
\left\|n_t(\mathbf{y},\eta^+)\right\|_{\mathbf{y}} \leq \frac{1}{4}.
$$

Proof. We can decompose the Newton's step taken at any feasible point in the following way:

$$
\begin{aligned}
n_t(\mathbf{y},\eta) &= -\mathbf{D}(\mathbf{x})^{-1}\begin{bmatrix}\eta\mathbf{c} + \nabla\phi(\mathbf{x}) + \mathbf{A}^\top\boldsymbol{\nu}\\ \mathbf{0}\end{bmatrix}\\
&= -\mathbf{D}(\mathbf{x})^{-1}\left(\begin{bmatrix}\eta\mathbf{c}\\ \mathbf{0}\end{bmatrix} + \begin{bmatrix}\nabla\phi(\mathbf{x}) + \mathbf{A}^\top\boldsymbol{\nu}\\ \mathbf{0}\end{bmatrix}\right).
\end{aligned}
$$

This allows us to express $n_t(\mathbf{y},\eta^+)$ as a function of $n_t(\mathbf{y},\eta)$:

$$
n(\mathbf{y},\eta^+) = \frac{\eta^+}{\eta} n_t(\mathbf{y},\eta) + \left(\frac{\eta^+}{\eta} - 1\right) h(\mathbf{y}).
$$

Taking the Hessian norm on both sides leads to:

$$\begin{aligned}\left\|n_t(\mathbf{y},\eta^+)\right\|_{\mathbf{y}} &\leq \frac{\eta^+}{\eta}\left\|n_t(\mathbf{y},\eta)\right\|_{\mathbf{y}} + \left|1-\frac{\eta^+}{\eta}\right| \left\|h(\mathbf{y})\right\|_{\mathbf{y}} \\ &\leq \frac{\eta^+}{\eta}\left\|n_t(\mathbf{y},\eta)\right\|_{\mathbf{y}} + \left|1-\frac{\eta^+}{\eta}\right| \sqrt{v_f}.\end{aligned}$$

By assumption, $\|n_t(\mathbf{y},\eta)\|_{\mathbf{y}} \leq \frac{1}{9}$. Letting $\frac{\eta^+}{\eta} = \beta \leq 1 + \frac{1}{8\sqrt{v_f}}$ we obtain:

$$\left\|n_t(\mathbf{y},\eta^+)\right\|_{\mathbf{y}} \leq \frac{1}{9}\left(1+\frac{1}{8\sqrt{v_f}}\right) + \frac{\sqrt{v_f}}{8\sqrt{v_f}} \leq \frac{1}{4},$$

where we use the fact that $\sqrt{v_f} \geq 1$ [32, Section 2.2] to complete the proof. □

Lemma 5.8. *Let* $\mathbf{x}_t^\eta \in \operatorname{argmin}_{\mathbf{x}}\{\eta\mathbf{c}^\top\mathbf{x} + \phi(\mathbf{x}) \text{ s.t. } \mathbf{A}\mathbf{x} - \mathbf{b}_t = 0\}$*. If* $\bar{\mathbf{x}}$ *is a feasible point then,*

$$\mathbf{c}^\top(\bar{\mathbf{x}} - \mathbf{x}_t^\eta) \leq \frac{v_f}{\eta_t}(1 + \|\bar{\mathbf{x}} - \mathbf{x}_t^\eta\|_{\mathbf{x}_t^\eta}). \tag{5.11}$$

Proof. The above identity follows from [32, Section 2.4.1] for a linear objective. □

Finally, we relate $\|\mathbf{y}_t - \mathbf{y}_t^\eta\|_{\mathbf{y}}$ to $\|\mathbf{x}_t - \mathbf{x}_t^\eta\|_{\mathbf{x}}$ in Lemma 5.9.

Lemma 5.9. *For any primal-dual feasible point* $\mathbf{y}$*, we have:*

$$\|\mathbf{y}_t - \mathbf{y}_t^\eta\|_{\mathbf{y}} \geq \|\mathbf{x}_t - \mathbf{x}_t^\eta\|_{\mathbf{x}}.$$

Proof. Recall from the definition of the Hessian norm:

$$\begin{aligned}\|\mathbf{y}_t - \mathbf{y}_t^\eta\|_{\mathbf{y}}^2 &= (\mathbf{y}_t - \mathbf{y}_t^\eta)^\top \mathbf{D}(\mathbf{y})(\mathbf{y}_t - \mathbf{y}_t^\eta) \\ &\geq \begin{bmatrix}\mathbf{x}_t - \mathbf{x}_t^\eta \\ \mathbf{0}\end{bmatrix}^\top \mathbf{D}(\mathbf{y}) \begin{bmatrix}\mathbf{x}_t - \mathbf{x}_t^\eta \\ \mathbf{0}\end{bmatrix} \\ &= (\mathbf{x}_t - \mathbf{x}_t^\eta)^\top \nabla^2\phi(\mathbf{x}_t)(\mathbf{x}_t - \mathbf{x}_t^\eta) \\ &= \|\mathbf{x}_t - \mathbf{x}_t^\eta\|_{\mathbf{x}_t}^2.\end{aligned}$$

Taking the square root on both sides, we obtain the desired result completing the proof. □

5.3 Algorithms

Suppose we have an initial primal-dual point $\mathbf{x}_0$, $\boldsymbol{\nu}_0$ and an initial barrier parameter η_0 such that $\|n_0(\mathbf{y}_0,\eta_0)\|_{\mathbf{y}_0} \leq \frac{1}{9}$ and $\mathbf{A}\mathbf{x}_0 = \mathbf{b}_0$. We propose the online interior-point method for time-varying equality constraints (`OIPM-TEC`) presented in Algorithm 3.

Algorithm 3: `OIPM-TEC`

Parameters: $\mathbf{A}$, $\mathbf{c}$, β

Initialization: Receive $\mathbf{x}_0 \in \mathbb{R}^n$, $\boldsymbol{\nu}_0 \in \mathbb{R}^p$, $\eta_0 \in \mathbb{R}^+$ such that $\|n_0(\mathbf{y}_0, \eta_0)\|_{\mathbf{y}_0} \leq \frac{1}{9}$.

for $t = 0, 1, 2...T$ **do**

Implement the decision $\mathbf{x}_t$.

Observe the outcome $\mathbf{c}^\top \mathbf{x}_t$ and the new constraint $\mathbf{b}_t$.

Perform t-step:

$\tilde{\mathbf{y}}_{t+1} = \mathbf{y}_t - \begin{bmatrix} \nabla^2 \phi(\mathbf{x}_t) & \mathbf{A}^\top \\ \mathbf{A} & \mathbf{0} \end{bmatrix}^{-1} \begin{bmatrix} \nabla d(\mathbf{x}_t, \eta_t) + \mathbf{A}^\top \boldsymbol{\nu}_t \\ \mathbf{A}\mathbf{x} - \mathbf{b}_t \end{bmatrix}.$

Perform η-step:

$\eta_{t+1} = \eta_t \beta,$

$\mathbf{y}_{t+1} = \tilde{\mathbf{y}}_{t+1} - \begin{bmatrix} \nabla^2 \phi(\tilde{\mathbf{x}}_t) & \mathbf{A}^\top \\ \mathbf{A} & \mathbf{0} \end{bmatrix}^{-1} \begin{bmatrix} \nabla d(\tilde{\mathbf{x}}_t, \eta_{t+1}) + \mathbf{A}^\top \boldsymbol{\nu}_t \\ \mathbf{0} \end{bmatrix}.$

end for

Theorem 5.1. *If Assumptions 1-4 hold,* $\max_t \|\mathbf{b}_t - \mathbf{b}_{t-1}\| \leq \sqrt{\frac{3m}{160}}$*, and supposing an initial point* $\mathbf{x}_0$*,* $\boldsymbol{\nu}_0$ *and an initial barrier parameter* η_0 *such that* $\|n_0(\mathbf{y}_0, \eta_0)\|_{\mathbf{y}_0} \leq \frac{1}{9}$*, then the regret and constraint violation of Algorithm 3 are bounded by:*

$$R_{\mathrm{d}}(T) \leq \frac{11 v_f \beta}{5\eta_0(\beta - 1)} + cV_T \tag{5.12}$$

$$\mathrm{Vio}(T) \leq V_{\mathbf{b}}. \tag{5.13}$$

The dynamic regret and constraint violation are, therefore, $O(V_T + 1)$ *and* $O(V_{\mathbf{b}})$*, respectively.*

Proof. We first derive some properties of the algorithm. Suppose at some round we have $\mathbf{y}_{t-1}$ such that $\|n_{t-1}(\mathbf{y}_{t-1}, \eta_{t-1})\|_{\mathbf{y}_{t-1}} \leq \frac{1}{9}$. Then, when observing the new vector $\mathbf{b}_t$, Lemma 5.6 yields $\|n_t(\mathbf{y}_{t-1}, \eta_{t-1})\|_{\mathbf{y}_{t-1}} \leq \frac{1}{4}$. By Lemma 5.3, this implies that the point $\tilde{\mathbf{y}}$ obtained after the t-step will respect:

$$\|n_t(\tilde{\mathbf{y}}_t, \eta_{t-1})\|_{\tilde{\mathbf{y}}_t} \leq \frac{1}{9}.$$

After updating η, Lemma 5.7 yields:

$$\|n_t(\tilde{\mathbf{y}}_t, \eta_t)\|_{\tilde{\mathbf{y}}_t} \leq \frac{1}{4}.$$

After the η-step from Lemma 5.3 we have:

$$\|n_t(\mathbf{y}_t, \eta_t)\|_{\mathbf{y}_t} \leq \frac{1}{9}.$$

Lemma 5.4 leads to:

$$\|\mathbf{y}_t - \mathbf{y}_{t-1}^{\eta_t}\|_{\mathbf{y}_t} \leq \frac{1}{6}.$$

Now, applying Lemma 5.5, we get:

$$\|\mathbf{y}_t - \mathbf{y}_{t-1}^{\eta_t}\|_{\mathbf{y}_{t-1}^{\eta_t}} \leq \frac{1}{5}.$$

Finally, Lemma 5.9 yields,

$$\|\mathbf{x}_t - \mathbf{x}_{t-1}^{\eta_t}\|_{\mathbf{x}_{t-1}^{\eta_t}} \leq \frac{1}{5}. \tag{5.14}$$

It follows that because we initially have $\|n_0(\mathbf{y}_0, \eta_0)\|_{\mathbf{y}_0} \leq \frac{1}{9}$ by assumption, then $\|n_{t-1}(\mathbf{y}_{t-1}, \eta_{t-1})\|_{\mathbf{y}_{t-1}} \leq \frac{1}{9}$ for all t.

Defining $c = \|\mathbf{c}\|$, we now re-express the regret as:

$$\begin{aligned}
R_{\mathrm{d}}(T) &= \sum_{t=1}^{T} f(\mathbf{x}_t) - f(\mathbf{x}_t^*) \\
&\leq \sum_{t=1}^{T} \mathbf{c}^\top \left(\mathbf{x}_t - \mathbf{x}_{t-1}^{\eta_t} + \mathbf{x}_{t-1}^{\eta_t} - \mathbf{x}_{t-1}^* + \mathbf{x}_{t-1}^* - \mathbf{x}_t^*\right) \\
&\leq \sum_{t=1}^{T} \mathbf{c}^\top \left(\mathbf{x}_t - \mathbf{x}_{t-1}^{\eta_t} + \mathbf{x}_{t-1}^{\eta_t} - \mathbf{x}_{t-1}^*\right) + c \left\|\mathbf{x}_{t-1}^* - \mathbf{x}_t^*\right\| \\
&= \sum_{t=1}^{T} \mathbf{c}^\top \left(\mathbf{x}_t - \mathbf{x}_{t-1}^{\eta_t}\right) + \sum_{t=1}^{T} \mathbf{c}^\top \left(\mathbf{x}_{t-1}^{\eta_t} - \mathbf{x}_{t-1}^*\right) + cV_T.
\end{aligned} \tag{5.15}$$

Next, we bound the first terms of (5.15). The first sum represents the estimation error between the current decision and the optimum at a given η. Using Lemma 5.8,

$$\begin{aligned}
\sum_{t=1}^{T} \mathbf{c}^\top \left(\mathbf{x}_t - \mathbf{x}_{t-1}^{\eta_t}\right) &\leq \sum_{t=1}^{T} \frac{v_f}{\eta_t} \left(1 + \|\mathbf{x}_t - \mathbf{x}_{t-1}^{\eta_t}\|_{\mathbf{x}_{t-1}^{\eta_t}}\right) \\
&\leq v_f \sum_{t=1}^{T} \frac{1}{\eta_t} \left(1 + \frac{1}{5}\right) \\
&= \frac{6v_f}{5\eta_0} \sum_{t=1}^{T} \left(\frac{1}{\beta}\right)^t \leq \frac{6v_f\beta}{5\eta_0(\beta - 1)},
\end{aligned}$$

where we used (5.14) to obtain the second inequality.

From [29, Chapter 9], we have that any optimizer $\mathbf{x}_t^\eta$ of the functional $d_t(\mathbf{x}_t, \eta)$ for a particular

η has an optimality gap of at most $\frac{v_f}{\eta}$, i.e., $\mathbf{c}^\top(\mathbf{x}_t^\eta - \mathbf{x}_t^*) \le \frac{v_f}{\eta}$. Therefore,

$$\begin{aligned}\sum_{t=1}^{T} \mathbf{c}^\top(\mathbf{x}_{t-1}^{\eta_t} - \mathbf{x}_{t-1}^*) &\le \sum_{t=1}^{T} \frac{v_f}{\eta_t} \\ &= \frac{v_f}{\eta_0} \sum_{t=1}^{T} \left(\frac{1}{\beta}\right)^t \le \frac{v_f \beta}{\eta_0(\beta - 1)}.\end{aligned}$$

Thus, the regret bound is :

$$R_{\mathrm{d}}(T) \le cV_T + \frac{11 v_f \beta}{5\eta_0(\beta - 1)}.$$

As for constraint violation, we have:

$$\mathrm{Vio}(T) = \sum_{t=1}^{T} \left[\sum_i d_{K_i} + \|\mathbf{A}\mathbf{x}_t - \mathbf{b}_t\|\right].$$

We note that iterates are always in the feasible space defined by the inequality constraints $g_i(\mathbf{x})$. This is because, at each round, a unit-step is taken in the Newton direction which lies in the feasible space of the self-concordant functional [32]. It follows that, $g_i(\mathbf{x}_t) \preceq 0 \forall t, i$ and $d_{K_i} = 0$. We also observe that iterates are always feasible with respect to the previous round's equality constraints. This is again because of the unit-step taken during the t-step [29]. Using this fact, we can write:

$$\begin{aligned}\mathrm{Vio}(T) &= \sum_{t=1}^{T} \Big[0 + \|\mathbf{A}\mathbf{x}_t - \mathbf{b}_t\|\,\Big] \\ &= \sum_{t=1}^{T} \|\mathbf{b}_{t-1} - \mathbf{b}_t\| = V_{\mathbf{b}},\end{aligned}$$

which completes the proof. □

When operating power grids, the exact optimal solution is not always required. Indeed there is a non-zero tolerance-region around the optimum that the electric grid operator can accept. For example, if the objective is cost-minimization, a solution to the OPF within a few cost units would fall within the tolerance of the grid operator. To translate this to a more traditional OCO framework, we can define an epsilon-regret $R_\epsilon(T)$ as:

$$R_\epsilon(T) = \sum_{t=1}^{T} \left[f_t(\mathbf{x}_t) - f_t(\mathbf{x}_t^*) - \epsilon\right]^+, \tag{5.16}$$

where $0 \le \epsilon$ is the desired tolerance. We remark that the epsilon-regret is null if all decisions are within our tolerance and positive otherwise. Note that epsilon-regret is a weaker metric than dynamic regret. A bounded dynamic regret implies a bounded epsilon-regret.

If instead of requiring sublinear dynamic regret, we are only interested in sublinear epsilon-regret, an algorithm requiring only a t-step can be used: the epsilon online interior-point method for time-varying equality constraints (ϵ`OIPM-TEC`). This algorithm is presented in Algorithm 4.

Algorithm 4: ϵ`OIPM-TEC`

Parameters: $\mathbf{A}$, $\mathbf{c}$

Initialization: Receive $\mathbf{x}_0 \in \mathbb{R}^n$, $\boldsymbol{\nu}_0 \in \mathbb{R}^p$, $\eta \in \mathbb{R}^+$ such that $\|n_0(\mathbf{y}_0, \eta)\|_{\mathbf{y}_0} \leq \frac{1}{9}$.

for $t = 0, 1, 2...T$ **do**

Implement the decision $\mathbf{x}_t$.

Observe the outcome $\mathbf{c}^\top \mathbf{x}_t$ and the new constraint $\mathbf{b}_t$.

$$\mathbf{y}_{t+1} = \mathbf{y}_t - \begin{bmatrix} \nabla^2 \phi(\mathbf{x}_t) & \mathbf{A}^\top \\ \mathbf{A} & \mathbf{0} \end{bmatrix}^{-1} \begin{bmatrix} \nabla d(\mathbf{x}_t, \eta_t) + \mathbf{A}^\top \boldsymbol{\nu}_t \\ \mathbf{A}\mathbf{x} - \mathbf{b}_t \end{bmatrix}.$$

end for

Theorem 5.2. *Let $\epsilon \geq 0$ be the tolerance. If Assumptions 1-4 hold,* $\max_t \|\mathbf{b}_t - \mathbf{b}_{t-1}\| \leq \sqrt{\frac{3m}{160}}$*, and we have the initial point $\mathbf{x}_0$, $\boldsymbol{\nu}_0$, and η such that $\eta \geq \frac{11 v_f}{5\epsilon}$, $\|n_0(\mathbf{y}_0, \eta)\|_{\mathbf{y}_0} < \frac{1}{9}$, then the epsilon-regret and constraint violation of Algorithm 4 are bounded by:*

$$R_\epsilon(T) \leq c V_T \tag{5.17}$$

$$\text{Vio}(T) \leq V_{\mathbf{b}}. \tag{5.18}$$

Proof. The constraint violation is bounded by $V_{\mathbf{b}}$ following the same reasoning as for Algorithm 3.

Similarly to Algorithm 3, the point $\mathbf{y}_{t+1}$ obtained after the Newton's step will respect: $\|n_t(\mathbf{y}_t, \eta_{t-1})\|_{\tilde{\mathbf{y}}_t} \leq \frac{1}{9}$. Applying Lemmas 5.4, 5.5, and 5.9 successively, we have: $\|\mathbf{x}_{t+1} - \mathbf{x}_t^\eta\|_{\mathbf{x}_t^\eta} \leq \frac{1}{5}$. We can therefore re-express the epsilon-regret as:

$$\begin{aligned} R_\epsilon(T) &= \sum_{t=1}^{T} [f_t(\mathbf{x}_t) - f_t(\mathbf{x}_t^*) - \epsilon]^+ \\ &\leq \sum_{t=1}^{T} \left[\frac{v_f}{\eta} \left(1 + \|\mathbf{x}_{t+1} - \mathbf{x}_t^\eta\|_{\mathbf{x}_t^\eta} \right) + \frac{v_f}{\eta} - \epsilon \right]^+ + c V_T \\ &\leq \sum_{t=1}^{T} \left[\frac{11 v_f}{5\eta} - \epsilon \right]^+ + c V_T. \end{aligned} \tag{5.19}$$

Because $\eta \geq \frac{11 v_f}{5\epsilon}$ by assumption, $\epsilon \geq \frac{11 v_f}{5\eta}$. Therefore the leftmost sum of (5.19) is zero. This implies that $R_\epsilon(T) \leq c V_T$, thus completing the proof. $\square$

5.4 Numerical Experiment

We now illustrate the performance of our algorithms in an online OPF problem. In the online OPF, the grid operator seeks to serve time-varying loads within a fixed network while minimizing the generation costs. Consider a set of buses $\mathcal{B}$ on which lies a set of generators $\mathcal{G}$ connected to a set of loads $\mathcal{D}$ by power lines $\mathcal{L}$. The admittance of every line $ij \in \mathcal{L}$ is given by $\gamma_{ij} \in \mathbb{C}$.

At time t, this problem's decision variables are the real and reactive power injections at every bus given by $p_{i,t}$ and $q_{i,t}$, respectively. The dependent variables are the voltages, $v_{i,t}$ at each bus. The parameters are the minimum and maximum power limits on active and reactive generation $(\underline{p}, \underline{q}, \overline{p}, \overline{q})$, the apparent power limits in the power lines $(\overline{k}_{ij})$ and the voltage limits $(\underline{v}, \overline{v})$. The active and reactive demand at every bus at time t is given by $p_{d,t}$ and $q_{d,t}$, respectively. Finally, in order to obtain a second-order cone relaxation of the problem, the matrix $\mathbf{W} \in \mathbb{C}^{\mathrm{card}\mathcal{B}^2}$ is formed where every element $\mathbf{W}_{ij}$ for $ij \in \mathcal{L}$ represents the product $v_i v_j^*$, with * denoting the complex conjugate [37]. At time t, the online OPF problem takes the following form:

$$
\begin{aligned}
\min \quad & s_t \\
\text{s.t.} \quad & \sum_{i \in \mathcal{G}} a_i p_{i,t}^2 + b_i p_{i,t} - s_t \leq 0 \\
& \underline{p_i} \leq p_{i,t} \leq \overline{p_i} \quad \forall i \in \mathcal{G} \\
& \underline{q_i} \leq q_{i,t} \leq \overline{q_i} \quad \forall i \in \mathcal{G} \\
& \underline{v}^2 \leq \mathbf{W}_{ii,t} \leq \overline{v}^2 \quad \forall i \in \mathcal{B} \\
& \left\| (\mathbf{W}_{ii,t} - \mathbf{W}_{ij,t}) \gamma_{ij}^* \right\| \leq \overline{k}_{ij} \quad \forall ij \in \mathcal{L} \\
& \left\| \begin{matrix} 2\mathbf{W}_{ij,t} \\ \mathbf{W}_{ii,t} - \mathbf{W}_{jj,t} \end{matrix} \right\| \leq \mathbf{W}_{ii,t} + \mathbf{W}_{jj,t} \quad \forall ij \in \mathcal{L} \\
& p_{i,t} + p_{d,t} + \jmath q_{i,t} + \jmath q_{d,t} = \sum_{j \in \mathcal{N}} (\mathbf{W}_{ii,t} - \mathbf{W}_{ij,t}) \gamma_{ij}^* \quad \forall ij \in \mathcal{L}
\end{aligned}
\tag{5.20}
$$

where $\jmath$ is the imaginary number. We consider a time-varying extension of the second-order cone relaxation of the Matpower 33-bus case network presented in [54]. To represent demand fluctuation, a random variation $\Delta p_{d,t}$ is added to every load at every timestep. For this experiment, every variation is independent and taken as: $\Delta p_{d,t} = \frac{0.01\zeta}{\sqrt{t}}$ where ζ is uniformly sampled in $[0, 1]$ at every round. The load variation diminishes with the square root of t ensuring that the variation in optima (V_T) and in the online parameter $\mathbf{b}_t$ $(V_{\mathbf{b}})$ are sublinear. If the offline problem, at a given timestep, is unsolvable, new random variables ζ are selected.

The time horizon is set at $T = 2000$. For `OIPM-TEC`, the initial barrier parameter is set to $\eta_0 = 1$ and the parameter β is set to 1.02. For ϵ`OIPM-TEC`, the barrier parameter is set to $\eta = 10^4$. The initial primal-dual point $\mathbf{x}_0, \boldsymbol{\nu}_0$ for both algorithms is determined from the offline optimal solution of the Matpower 33-bus network. The acceptable tolerance $\epsilon = 0.015$ is selected for the epsilon-regret which represents less than 0.001% of the offline optimal cost. This value is equivalent to 0.015$/hr in this case. As a comparison, the `MOSP` algorithm [10] is also applied to problem (5.20) using the same initial primal point as `OIPM-TEC`. In the `MOSP` implementation, the time-varying equality constraints are relaxed to inequality constraints such that $\mathbf{A}\mathbf{x}_t - \mathbf{b}_t \leq \mathbf{0}$. This enables the operator to over-serve loads and should be strictly respected at the optimum. The other linear inequality constraints are implemented directly while a projection is used to satisfy the second-order cone constraints. The parameters of the algorithm are set to : $\alpha = \mu = t^{-\frac{1}{3}}$, where t denotes the current round.

The resulting regret (dynamic regret and epsilon regret are considered) and constraint violation suffered by the three algorithms are presented in Figures 5.1 and 5.2. We remark that all algorithms achieve sublinear epsilon-regret and constraint violation. However, the epsilon-regret suffered by `OIPM-TEC`, while higher than that of ϵ`OIPM-TEC`, is much lower than for `MOSP`. Both `OIPM-TEC` and ϵ`OIPM-TEC` exhibit sublinear constraint violation terms that outperform `MOSP`.

5.5 Conclusion

In this work, an online optimization algorithm using an interior-point method-like update capable of solving the online OPF is presented. This algorithm, the `OIPM-TEC`, admits time-varying equality constraints and generalized inequalities while possessing a $O(V_T+1)$ dynamic regret bound and a $O(V_{\mathbf{b}})$ constraint violation bound. Additionally, if there exists an acceptable tolerance when solving an online optimization problem, the ϵ`OIPM-TEC` algorithm is capable of ensuring the tightest possible bound with respect to T on regret and constraint violation. These two algorithms are applied to an online OPF example showcasing their performance compared to other methods from the literature.

The possibility of using OCO approaches to solve the OPF bridges the gap between traditional OCO algorithms with performance guarantees and real-time OPF approaches like [34] and [52] that enjoy good real-world performance. A possible future direction is the development of online primal-dual interior-point methods which could increase numerical stability.

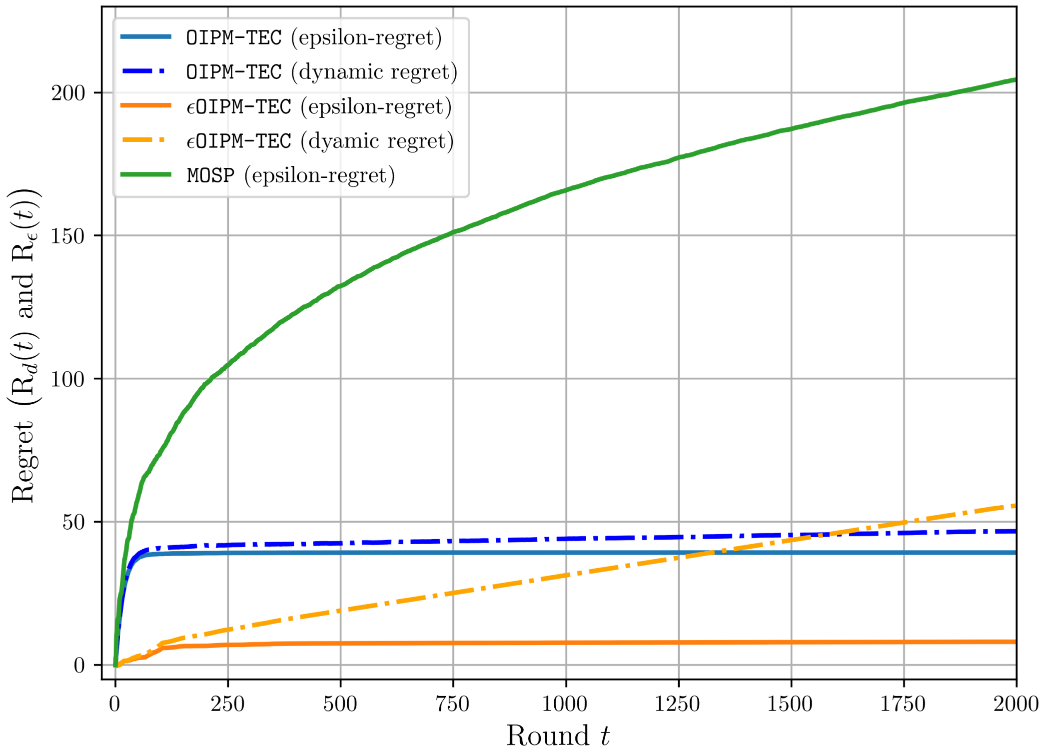

Figure 5.1 Regret comparison of `OIPM-TEC`, ϵ`OIPM-TEC`, and `MOSP`

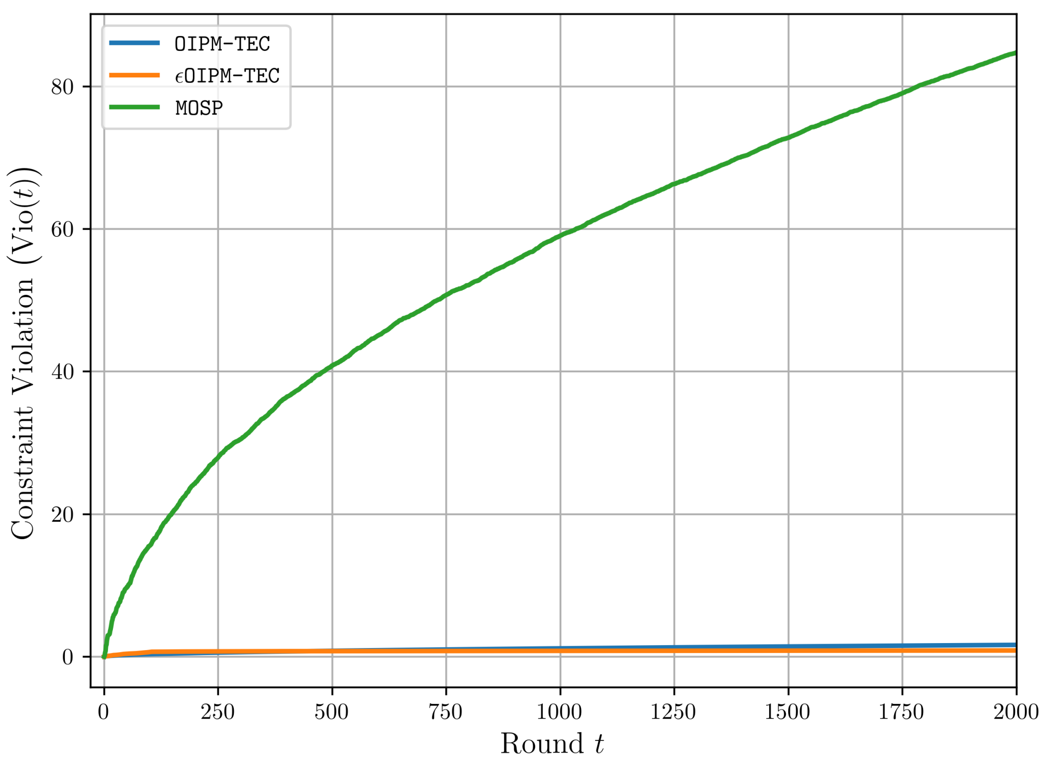

Figure 5.2 Constraint violation comparison of `OIPM-TEC`, ϵ`OIPM-TEC`, and `MOSP`

CHAPTER 6 GENERAL DISCUSSION

The importance of the algorithms presented in this Master's thesis in regards to the OCO literature can be decomposed into three main components.

The first aspect is the tight regret bounds with respect to the time horizon T obtained by these second-order algorithms. All regret bounds obtained for the second-order methods presented in this thesis are of the order of $O(V_T + 1)$ or smaller. This is the same order as that of the algorithms from [25, 26] which are the tightest proposed for dynamic regret. Simply put, when the required conditions hold, no other OCO algorithm in the literature can achieve a better asymptotic tracking behaviour than those presented in this work.

Second, the `OPEN-M` and ϵ`OIPM-TEC` algorithms presented are free of any time-dependent parameters. This facilitates initialization but is especially useful for restarting the algorithm without penalty. These algorithms are also free from hyper-parameters meaning there is no need for hyper-parameter tuning. As long as the necessary assumptions are met, these algorithms will perform well. This allows easy implementation of the algorithms.

Finally, the `OIPM-TEC` algorithm is the first OCO algorithm capable of tracking moving optima while ensuring implemented decisions are feasible with respect to generalized inequalities at all rounds without the need for a projection step. Operators can, therefore, employ this algorithm in safety-critical applications like power systems knowing that the implemented decisions are guaranteed to respect safety limits. The importance of this achievement cannot be understated. Safety concerns represent one of the principal barriers to implementation in real-world applications faced by OCO algorithms. Addressing this issue, thus, opens up many novel use-cases that would otherwise have been deemed too risky for OCO approaches.

The cost for these three benefits is requiring twice-differentiable loss functions and the computational power to perform a matrix inversion. Although for some extremely short-timescale problems or for applications with very limited computational resources these requirements might be prohibitive, online second-order methods present unique characteristics which cannot be replicated by any other approach in the literature.

CHAPTER 7 CONCLUSION AND RECOMMENDATIONS

Constrained OCO is an extension of the traditional OCO framework in which decisions must satisfy time-varying constraints. Although a difficult problem, offline Newton's methods and interior-point algorithms present very powerful tools to address the issue.

In this work, four algorithms pertaining to constrained OCO using second-order information are presented. These algorithms tackle time-varying linear equality constraints in different OCO problem formulations. Tight regret and constraint violation bounds are obtained for these algorithms under mild conditions which are common in the OCO literature [26,35]. As illustrated in numerical applications in online network resource allocation and OPF, these algorithms are competitive with respect to previously proposed methods in the literature.

Future Work

The next challenge suggested by these results is to design OCO algorithms capable of simultaneously handling time-varying loss functions, equality constraints and generalized inequalities. However, the precarious numerical stability of the primal interior-point methods – as observed by [29, 32] and during the numerical experiments presented in this work – would indicate a slightly different approach is necessary to achieve this goal. Fortunately, a potential solution already exists in the offline optimization literature: primal-dual interior-point methods. Unlike primal methods, these do not require feasibility at every timestep and are more numerically stable [32, Section 3.7]. In particular, an online potential-reduction method would appear the most promising.

Another avenue of interest would be the use of online second-order methods for bandit feedback OCO or for use in distributed optimization settings. These types of approaches are of particular interest for applications in power systems where feedback from the sensors distributed in the grid can be intermittent [22, 55]. Distributed optimization is also very useful in large-scale systems where computational power is decentralized and privacy is a major consideration [56].

REFERENCES

[1] S. Shalev-Shwartz, "Online learning and online convex optimization," *Foundations and Trends® in Machine Learning*, vol. 4, no. 2, pp. 107–194, 2012.

[2] E. Hazan, "Introduction to online convex optimization," *Foundations and Trends® in Machine Learning*, vol. 2, no. 3-4, pp. 157–325, 2015.

[3] 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.

[4] E. C. Hall and R. M. Willett, "Online convex optimization in dynamic environments," *IEEE Journal of Selected Topics in Signal Processing*, vol. 9, no. 4, pp. 647–662, 2015.

[5] A. Agarwal, E. Hazan, S. Kale, and R. E. Schapire, "Algorithms for portfolio management based on the newton method," in *Proceedings of the 23rd International Conference on Machine Learning*, ser. ICML '06. New York, NY, USA: Association for Computing Machinery, 2006, p. 9–16.

[6] P. H. Algoet and T. M. Cover, "Asymptotic optimality and asymptotic equipartition properties of log-optimum investment," *The Annals of Probability*, pp. 876–898, 1988.

[7] A. Agarwal and E. Hazan, "New algorithms for repeated play and universal portfolio management," Princeton University Technical Report TR-740-05, Tech. Rep., 2005.

[8] Q. Liu, W. Wu, L. Huang, and Z. Fang, "Simultaneously achieving sublinear regret and constraint violations for online convex optimization with time-varying constraints," *Performance Evaluation*, vol. 152, 2021.

[9] J. Mulvaney-Kemp, S. Park, M. Jin, and J. Lavaei, "Dynamic regret bounds for constrained online nonconvex optimization based on Polyak-Lojasiewicz regions," *IEEE Transactions on Control of Network Systems*, pp. 1 – 12, 2022.

[10] T.Chen, Q. Ling, and G. B. Giannakis, "An online convex optimization approach to proactive network resource allocation," *IEEE Transactions on Signal Processing*, vol. 65, pp. 6350–6364, 2017.

[11] H. Yu and M. J. Neely, “A low complexity algorithm with O($\sqrt{T}$) regret and O(1) constraint violations for online convex optimization with long term constraints,” *Journal of Machine Learning Research*, vol. 21, no. 1, pp. 1–24, 2020.

[12] D. J. Leith and G. Iosifidis, “Penalized ftrl with time-varying constraints,” in *Machine Learning and Knowledge Discovery in Databases*, M.-R. Amini, S. Canu, A. Fischer, T. Guns, P. Kralj Novak, and G. Tsoumakas, Eds. Cham: Springer Nature Switzerland, 2023, pp. 311–326.

[13] H. Liu, X. Xiao, and L. Zhang, “Augmented Langragian methods for time-varying constrained online convex optimization,” *arXiv preprint arXiv:2205.09571*, 2022.

[14] A. Rakhlin, O. Shamir, and K. Sridharan, “Making gradient descent optimal for strongly convex stochastic optimization,” *arXiv preprint arXiv:1109.5647*, 2012.

[15] E. Hazan, A. Argawal, and S. Kale, “Logarithmic regret algorithms for online convex optimization,” *Machine Learning*, vol. 69, pp. 169–192, 2007.

[16] X. Yi, X. Li, L. Xie, and K. Johansson, “Distributed online convex optimization with time-varying coupled inequality constraints,” *IEEE Transaction on Signal Processing*, vol. 68, pp. 731–746, 2020.

[17] P. Sharma, P. Khanduri, L. Shen, D. J. Bucci, and P. K. Varshney, “On distributed online convex optimization with sublinear dynamic regret and fit,” in *2021 55th Asilomar Conference on Signals, Systems, and Computers*, 2021, pp. 1013–1017.

[18] X. Cao and K. J. R. Liu, “Online convex optimization with time-varying constraints and bandit feedback,” *IEEE Transactions on Automatic Control*, vol. 64, pp. 2665–2680, 2018.

[19] A. D. Flaxman, A. T. Kalai, and H. B. McMahan, “Online convex optimization in the bandit setting: Gradient descent without a gradient,” in *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, ser. SODA ’05. USA: Society for Industrial and Applied Mathematics, 2005, p. 385–394.

[20] S. Bubeck, O. Dekel, T. Koren, and Y. Peres, “Bandit convex optimization: $\sqrt{T}$ regret in one dimension,” *Conference on Learning Theory*, pp. 266–278, 2015.

[21] 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*, vol. 28, no. 1, pp. 317–325, 2012.

[22] F. Badal, S. Sarker, and S. Das, "A survey on control issues in renewable energy integration and microgrid," *Protection and Control of Modern Power Systems*, vol. 4, no. 1, 2019.

[23] O. Besbes, Y. Gur, and A. Zeevi, "Non-stationary stochastic optimization," *Operations Research*, vol. 63, no. 5, pp. 1227–1244, oct 2015.

[24] L. Zhang, T. Yang, J. Yi, R. Jin, and Z.-H. Zhou, "Improved dynamic regret for non-degenerate functions," *Advances in Neural Information Processing Systems*, pp. 732–741, 2017.

[25] 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)*, 2016, pp. 7195–7201.

[26] A. Lesage-Landry, J. A. Taylor, and I. Shames, "Second-order online nonconvex optimization," *IEEE Transactions on Automatic Control*, vol. 66, no. 10, pp. 4866–4872, 2021.

[27] M. J. Neely and H. Yu, "Online convex optimization with time-varying constraints," *arXiv preprint arXiv:1702.04783*, 2017.

[28] J. Nocedal and S. J. Wright, *Numerical Optimization.* Springer Science & Business Media, 2006.

[29] S. Boyd and L. Vandenberghe, *Convex Optimization.* Cambridge University Press, 2004.

[30] Y. Nesterov and A. Nemirovski, "Interior-point polynomial algorithms in convex programming," in *SIAM studies in applied mathematics*, 1994.

[31] A. Ben-Tal and A. Nemirovski, *Lectures on Modern Convex Optimization.* Society for Industrial and Applied Mathematics, 2001.

[32] J. Renegar, *A Mathematical View of Interior-Point Methods in Convex Optimization.* Society for Industrial and Applied Mathematics, 2001.

[33] J. D. Abernethy, E. Hazan, and A. Rakhlin, "Interior-point methods for full-information and bandit online learning," *IEEE Transactions on Information Theory*, vol. 58, no. 7, pp. 4164–4175, 2012.

[34] Y. Tang, K. Dvijotham, and S. Low, "Real-time optimal power flow," *IEEE Transactions on Smart Grid*, vol. 8, no. 6, 2017.

[35] Z. Wang, W. Wei, J. Z. F. Pang, F. Liu, B. Yang, X. Guan, and S. Mei, "Online optimization in power systems with high penetration of renewable generation: Advances and prospects," *IEEE/CAA Journal of Automatica Sinica*, vol. 10, no. 4, pp. 839–858, 2023.

[36] J. Carpentier, "Contribution to the economic dispatch problem," *Bulletin de la Société française des électriciens*, vol. 3, no. 8, pp. 431, 447, 1962.

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

[38] S. H. Low, "Convex relaxation of optimal power flow—part I: Formulations and equivalence," *IEEE Transactions on Control of Network Systems*, vol. 1, no. 1, pp. 15–27, 2014.

[39] F. Zohrizadeh, C. Josz, M. Jin, R. Madani, J. Lavaei, and S. Sojoudi, "A survey on conic relaxations of optimal power flow problem," *European Journal of Operational Research*, vol. 287, no. 2, pp. 391–409, 2020.

[40] X. Bai, H. Wei, K. Fujisawa, and Y. Wang, "Semidefinite programming for optimal power flow problems," *Internation Journal of Electrical power & Energy Systems*, vol. 30, no. 6, pp. 383–392, 2008.

[41] R. Jabr, "Radial distribution load flow using conic programming," *IEEE Transactions on Power Systems*, vol. 21, no. 3, pp. 1458,1459, 2006.

[42] D. K. Molzahn and I. A. Hiskens, "A survey of relaxations and approximations of the power flow equations," *Foundations and Trends® in Electric Energy Systems*, vol. 4, no. 1-2, pp. 1–221, 2019.

[43] J.-L. Lupien and A. Lesage-Landry, "An online newton's method for time-varying linear equality constraints," *IEEE Control Systems Letters*, vol. 7, pp. 1423–1428, 2023.

[44] J. A. Taylor, S. V. Dhople, and D. S. Callaway, "Power systems without fuel," *Renewable and Sustainable Energy Reviews*, vol. 57, pp. 1322–1336, 2016.

[45] A. Lesage-Landry, H. Wang, I. Shames, P. Mancarella, and J. A. Taylor, "Online convex optimization of multi-energy building-to-grid ancillary services," *IEEE Transactions on Control Systems Technology*, vol. 28, no. 6, pp. 2416–2431, 2020.

[46] M. Fazlyab, S. Paternain, V. M. Preciado, and A. Ribeiro, "Prediction-correction interior-point method for time-varying convex optimization," *IEEE Transactions on Automatic Control*, vol. 63, no. 7, pp. 1973–1986, 2018.

[47] R. Ahuja, T. Magnanti, and J. Orlin, *Network Flows Theory, Algorithms and Applications.* Prentice-Hall, 1993.

[48] P. Nardelli, N. Rubido, C. Wang, M. Baptista, C. Pomalaza-Raez, P. Cardieri, and M. Latva-aho., "Models for the modern power grid," *The European Physical Journal Special Topics*, vol. 223, 2014.

[49] J.-L. Lupien, I. Shames, and A. Lesage-Landry, "Online interior-point methods for time-varying equality-constrained optimization," in *IEEE Transactions on Automatic Control*, 2023, pp. 1–8, submitted.

[50] 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*, vol. 12, no. 4, pp. 703–716, 2018.

[51] J. Kardoš, D. Kourounis, O. Schenk, and R. Zimmerman, "Beltistos: A robust interior point method for large-scale optimal power flow problems," *Electric Power Systems Research*, vol. 212, p. 108613, 2022.

[52] E. Dall'Anese and A. Simonetto, "Optimal power flow pursuit," *IEEE transactions on Smart Grid*, vol. 9, no. 2, pp. 942–952, 2018.

[53] M. Picallo, L. Ortmann, S. Bolognani, and F. Dörfler, "Adaptive real-time grid operation via online feedback optimization with sensitivity estimation," *Electric Power Systems Research*, vol. 212, p. 108405, 2022.

[54] R. D. Zimmerman, C. E. Murillo-Sánchez, and R. J. Thomas, "Matpower: Steady-state operations, planning, and analysis tools for power systems research and education," *IEEE Transactions on Power Systems*, vol. 26, no. 1, pp. 12–19, 2011.

[55] 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*, vol. 50, no. 5, pp. 1389–1400, 2009.

[56] A. Lesage-Landry and D. S. Callaway, "Dynamic and distributed online convex optimization for demand response of commercial buildings," *IEEE Control Systems Letters*, vol. 4, no. 3, pp. 632–637, 2020.