

Online Convex Optimization and Integral Quadratic Constraints: An automated approach to regret analysis

Fabian Jakob and Andrea Iannelli

Abstract—We propose a novel approach for analyzing dynamic regret of first-order online convex optimization (OCO) algorithms for strongly convex and Lipschitz-smooth objectives. Crucially, we provide a general analysis that is applicable to a wide range of first-order algorithms that can be expressed as an interconnection of a linear dynamical system in feedback with a first-order oracle. By leveraging Integral Quadratic Constraints (IQCs), we derive a semi-definite program which, when feasible, provides a regret guarantee for the online algorithm. For this, the concept of variational IQCs is introduced as the generalization of IQCs to time-varying monotone operators. Our bounds capture the temporal rate of change of the problem in the form of the path length of the time-varying minimizer and the objective function variation. In contrast to standard results in OCO, our results do not require neither the assumption of gradient boundedness, nor that of a bounded feasible set. Numerical analyses showcase the ability of the approach to capture the dependence of the regret on the function class condition number.

I. INTRODUCTION

Online Convex Optimization (OCO) has emerged as a powerful framework for tackling real-time decision-making problems under uncertainty. Traditionally, the study of OCO has focused on proposing online algorithms whose performance is assessed in terms of their static or dynamic regret [1], [2]. In recent years, this framework has raised interest in the control community both for the design of OCO-inspired controllers [3] and for using the concept of regret as a performance metric to evaluate controllers dealing with uncertainty [4].

OCO algorithms aim to make a sequence of decisions in real-time minimizing a cumulative loss function that is revealed sequentially. Several algorithms have emerged, among which first-order algorithms represent a fundamental subclass. Each algorithm comes with its individual regret guarantees and proof techniques to verify them [5]. For the particular case of strongly convex and smooth objective functions, also accelerated methods [6] and multi-step methods with regret guarantees have been proposed [7]. However, a general methodology to approach their analysis does not exist.

In contrast, for static convex optimization, an automated approach to analyze first-order algorithms based on systems theory has been thoroughly investigated in the last years, see e.g. [8], [9], [10] and references therein. The idea is to

model the first-order algorithm as a Lur’e system, i.e. an interconnection of a linear dynamical system in feedback with a first-order oracle. Integral Quadratic Constraints (IQCs) can then be leveraged to formulate analysis conditions based on Semi-Definite Programs (SDPs). Recently, this framework has also been extended to time-varying optimization [11].

Time-varying optimization and OCO are in fact inherently related [12]. In this paper, we take up tools from [11] to propose an automated approach to dynamic regret analysis for first-order algorithms in OCO. By leveraging a system theoretic view on algorithms, we model an OCO algorithm as a dynamical system interconnected with a *time-varying* first-order oracle. We consider strongly-convex and smooth objective functions. To handle the time-varying nature of the oracle, we leverage variational Integral Quadratic Constraints (vIQCs), which in contrast to conventional IQCs depend on different measures of temporal variation of the problem. In line with the OCO literature we obtain a bound on the dynamic regret that accounts for the path length of the time-varying optimal solution, the objective function variation, and the gradient variation. In contrast to the common analysis approaches in OCO, our proofs are algorithm-agnostic and thus, applicable to a large number of first-order algorithms. The regret upper bound depends on decision variables of the SDP, leaving the possibility to tune the bound by trading off their magnitude. We show the implicit dependence of the regret on the function class condition ratio via a numeric case study.

The main contributions can be summarized as follows. We provide a general modeling framework for OCO algorithms and a new proof technique for bounding dynamic regret. Notably, our approach does neither require the typical bounded gradient assumption, nor does it require the bounded feasible set assumption with the use of vIQCs. A numerical study demonstrates the versatility of the general algorithm formulation and provides a comparative study of commonly used OCO-algorithms. We believe the strengths of this new approach consist of: weaker assumptions required for the analysis; generality; and insights provided by comparing the analysis results of different algorithms.

The paper is structured as follows. We state the preliminaries in section II, introducing the notion of regret and our algorithm formulation. In section III, we introduce the IQC formulation needed to establish our regret bounds and illustrate them with numerical examples. In section IV, we derive our regret bounds as main results. We conclude the paper in section V.

Notation. Let $\text{vec}(v_1, v_2)$ denote the vertical stack of the

Authors are with the University of Stuttgart, Institute for Systems Theory and Automatic Control, 70569 Stuttgart, Germany. {fabian.jakob, andrea.iannelli}@ist.uni-stuttgart.de. Fabian Jakob acknowledges the support of the International Max Planck Research School for Intelligent Systems (IMPRS-IS).

vectors v_1, v_2 . We denote $\mathbf{1}_p$ a column vector of ones of length p and \mathbb{I}_p the index set of integers from 1 to p . We indicate by I_d a $d \times d$ identity matrix. For positive definite P , we define the norm $\|v\|_P := \sqrt{v^\top P v}$. The Kronecker product between two matrices is written as $A \otimes B$. Let $\text{diag}(v)$ denote the diagonal matrix with the elements of a vector v and $\text{blkdiag}(P_1, P_2)$ a blockdiagonal matrix of P_1, P_2 . We write $f(T) = \mathcal{O}(g(T))$ if $\lim_{T \rightarrow \infty} \frac{f(T)}{g(T)} < \infty$. A linear dynamic mapping $x_{t+1} = Ax_t + Bu_t, y_t = Cx_t + Du_t$ is compactly expressed as $y_t = Gu_t, G = \begin{bmatrix} A & B \\ C & D \end{bmatrix}$.

II. PROBLEM SETUP

A. Online Convex Optimization

In OCO, an algorithm sequentially selects an action $x_t \in \mathcal{X}$ from a closed convex decision set \mathcal{X} at each time step t , based solely on the information available up to time $t-1$. Upon choosing x_t , a convex objective function f_t is revealed, and the algorithm incurs a loss $f_t(x_t)$. Throughout this work we assume the objective functions f_t are m -strongly convex and L -smooth, and that $\mathcal{X} \subseteq \mathbb{R}^d$. The goal is to minimize the cumulative loss over T rounds, and the performance of an OCO algorithm is typically assessed in terms of regret, which represents the cumulative suboptimality *w.r.t.* the best possible decisions in hindsight [13]. In this work, we define the best hindsight decision as the pointwise-in-time minimizer

$$x_t^* = \arg \min_{x \in \mathcal{X}} f_t(x), \quad t \in \mathbb{N}_+, \quad (1)$$

leading to the notion of dynamic regret [2]

$$\mathcal{R}_T = \sum_{t=1}^T f_t(x_t) - \sum_{t=1}^T f_t(x_t^*). \quad (2)$$

Classical analyses of OCO primarily focus on deriving upper bounds on the regret. Typically, dynamic regret bounds are characterized in terms of regularity measures, which quantify the temporal change of the problem [14], [15]. Notable regularity measures commonly employed in the literature include the path length and squared path length, respectively defined as

$$\mathcal{P}_T = \sum_{t=1}^T \|x_t^* - x_{t+1}^*\|, \quad \mathcal{S}_T = \sum_{t=1}^T \|x_t^* - x_{t+1}^*\|^2, \quad (3)$$

and the function and gradient variation

$$\mathcal{V}_T = \sum_{t=2}^T \sup_{x \in \mathcal{X}} |f_{t-1}(x) - f_t(x)|, \quad (4)$$

$$\mathcal{G}_T = \sum_{t=2}^T \sup_{x \in \mathcal{X}} \|\nabla f_{t-1}(x) - \nabla f_t(x)\|^2. \quad (5)$$

Table I provides a comparative overview of existing regret bounds for strongly convex and smooth objectives, alongside the bounds established in this work.

We will consider (1) in its equivalent composite form

$$\min_{x \in \mathbb{R}^d} f_t(x) + \mathcal{I}_{\mathcal{X}}(x), \quad (6)$$

TABLE I: Regret bounds for strongly convex and smooth objectives (OGD=online gradient descent, MD=mirror descent).

Ref.	Algorithm	Regret bound	Assumptions
[15]	OGD	$\mathcal{O}(\mathcal{P}_T)$	∇f_t bounded
[16]	Optimistic MD	$\mathcal{O}(\log \mathcal{G}_T)$	\mathcal{X} bounded
[7]	Multi-step OGD	$\mathcal{O}(\min\{\mathcal{P}_T, \mathcal{S}_T, \mathcal{V}_T\})$	∇f_t bounded
Thm. 4	(7)	$\mathcal{O}(\mathcal{P}_T)$	\mathcal{X} bounded
Thm. 5	(7)	$\mathcal{O}(\mathcal{S}_T + \mathcal{V}_T + \mathcal{G}_T)$	None

where $\mathcal{I}_{\mathcal{X}}$ is the indicator function

$$\mathcal{I}_{\mathcal{X}}(x) = \begin{cases} 0 & , \text{ if } x \in \mathcal{X}, \\ \infty & , \text{ if } x \notin \mathcal{X}. \end{cases}$$

As it is well known, the subdifferential of $\mathcal{I}_{\mathcal{X}}$ is the normal cone of the set \mathcal{X} , denoted as $\partial \mathcal{I}_{\mathcal{X}}(x) \triangleq \mathcal{N}_{\mathcal{X}}(x)$ [17].

B. General First-Order Algorithms

In line with [8]–[10], we consider general first-order algorithms that are expressed as a linear time-invariant system

$$\begin{aligned} \xi_{t+1} &= A\xi_t + Bu_t, \\ y_t &= C\xi_t + Du_t \end{aligned} \quad (7a)$$

with state $\xi \in \mathbb{R}^{n_\xi}$ in feedback with a first-order oracle $u_t = \varphi_t(y_t)$. For some integer $p \geq 1, q \geq 0$, we consider that the algorithm makes use of p gradient evaluations of f_t and q subgradient evaluations of $\mathcal{I}_{\mathcal{X}}$, and the input and output can be decomposed into

$$y_t = \begin{bmatrix} s_t \\ z_t \end{bmatrix}, \quad u_t = \begin{bmatrix} \delta_t \\ g_t \end{bmatrix}, \quad (7b)$$

where $s_t := \text{vec}(s_t^1, \dots, s_t^p)$ and $z_t := \text{vec}(z_t^1, \dots, z_t^q)$, with $s_t^i, z_t^j \in \mathbb{R}^d$ for all $i \in \mathbb{I}_p, j \in \mathbb{I}_q$, and

$$\delta_t = \begin{bmatrix} \nabla f_t(s_t^1) \\ \vdots \\ \nabla f_t(s_t^p) \end{bmatrix}, \quad g_t = \begin{bmatrix} g_t^1 \\ \vdots \\ g_t^q \end{bmatrix}, \quad g_t^j \in \mathcal{N}_{\mathcal{X}}(z_t^j). \quad (7c)$$

The case $q = 0$ is relevant for unconstrained optimization problems. The algorithm's iterate coincides with the first output, i.e., $x_t = s_t^1$. To enforce that x_t only depends on information up to $t-1$, we assume the readout is independent from u_t , that is, the first block-row in D is assumed to be zero. The first block-row of C is denoted as C_1 , i.e., $x_t = C_1 \xi_t$. We furthermore assume that the pair (C_1, A) is observable, so that if the set \mathcal{X} is bounded, then the set of internal states ξ_t is also bounded.

To ensure that this dynamical system is meaningful from an optimization standpoint we require that the fixed point of (7) satisfies the first-order optimality conditions of (II-A)

$$-\nabla f_t(x_t^*) \in \mathcal{N}_{\mathcal{X}}(x_t^*). \quad (8)$$

We will particularly constrain ourselves to algorithms whose fixed-points are of the form

$$\begin{aligned} \xi_t^* &= A\xi_t^*, & y_t^* &= C\xi_t^*, \\ y_t^* &= \begin{bmatrix} 1_p \otimes I_d \\ 1_q \otimes I_d \end{bmatrix} x_t^*, & u_t^* &= \begin{bmatrix} 1_p \otimes I_d \\ -1_q \otimes I_d \end{bmatrix} \nabla f_t(x_t^*), \end{aligned} \quad (9)$$

that is, $s_t^{i,*} = z_t^{j,*} = x_t^*$, and $\delta_t^{i,*} = -g_t^{j,*} = \nabla f_t(x_t^*)$, for all $i \in \mathbb{I}_p, j \in \mathbb{I}_q$. Moreover, we enforce $Bu_t^* = 0$ and $Du_t^* = 0$, which will be necessary for subsequent derivations. We make the following assumptions to ensure (7) exhibits (9) as fixed-point.

Assumption 1: There exists a matrix U such that

$$\begin{bmatrix} I - A \\ C \end{bmatrix} U = \begin{bmatrix} 0 \\ \begin{bmatrix} 1_p \\ 1_q \end{bmatrix} \otimes I_d \end{bmatrix}, \quad (10)$$

Assumption 2: If $q \geq 1$, then the kernels of the matrices B and D satisfy

$$\ker B = \begin{bmatrix} 1_p \otimes I_d \\ -1_q \otimes I_d \end{bmatrix}, \quad \ker D = \begin{bmatrix} 1_p \otimes I_d \\ -1_q \otimes I_d \end{bmatrix}. \quad (11)$$

Both assumptions essentially allow us, given an optimal point x_t^* , to reconstruct the optimal algorithmic state as $\xi_t^* = Ux_t^*$, which represents a special case of the conditions worked out in [18]. Assumption 2 specifically ensures $Bu_t^* = 0$, $Du_t^* = 0$ [11]. This is slightly restrictive, however, we still cover a large class of classical OCO algorithms. We illustrate the algorithm representation with the following example.

Example 1: Consider the example of the following two-step projected gradient descent

$$\begin{aligned} \hat{x}_t &= \Pi_{\mathcal{X}} [x_t - \alpha \nabla f_t(x_t)] \\ x_{t+1} &= \Pi_{\mathcal{X}} [\hat{x}_t - \alpha \nabla f_t(\hat{x}_t)], \end{aligned} \quad (12)$$

with $\Pi_{\mathcal{X}}(z) := \arg \min_{x \in \mathcal{X}} \|x - z\|^2$ and $\alpha > 0$. To bring (12) into the form (7) we leverage the first-order optimality condition arising from the projections, namely

$$\begin{aligned} x_t - \alpha \nabla f_t(x_t) - \hat{x}_t &\in \alpha \mathcal{N}_{\mathcal{X}}(\hat{x}_t) \\ \hat{x}_t - \alpha \nabla f_t(\hat{x}_t) - x_{t+1} &\in \alpha \mathcal{N}_{\mathcal{X}}(x_{t+1}). \end{aligned} \quad (13)$$

As a technicality, we scaled the cones by α so that (11) will be satisfied. Now define the outputs $s_t \triangleq \text{vec}(x_t, \hat{x}_t)$, $z_t \triangleq \text{vec}(\hat{x}_t, x_{t+1})$ and the inputs $\delta_t \triangleq \text{vec}(\nabla f_t(x_t), \nabla f_t(\hat{x}_t))$, $g_t^1 \in \mathcal{N}_{\mathcal{X}}(\hat{x}_t)$, $g_t^2 \in \mathcal{N}_{\mathcal{X}}(x_{t+1})$. Then, by letting $\xi_t := x_t$, we can write (12) as (7) with $p = 2, q = 2$ and

$$\begin{aligned} A &= 1 \otimes I_d, & B &= \begin{bmatrix} -\alpha & -\alpha & -\alpha & -\alpha \end{bmatrix} \otimes I_d, \\ C &= \begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \end{bmatrix} \otimes I_d, & D &= \begin{bmatrix} 0 & 0 & 0 & 0 \\ -\alpha & 0 & -\alpha & 0 \\ -\alpha & 0 & -\alpha & 0 \\ -\alpha & -\alpha & -\alpha & -\alpha \end{bmatrix} \otimes I_d. \end{aligned}$$

In the spirit of Example 1, many relevant algorithms can be formulated as (7), such as Online Gradient Descent (O-GD) [1], Online Accelerated Gradient Descent (O-AGD) [14], the Online Nesterov Method (O-NM) [13] or Online Mirror Descent [19].

III. IQCS FOR VARYING OPERATORS

Integral Quadratic Constraints provide a framework to characterize unknown or nonlinear input-output mappings in terms of inequalities [20]. Informally, an operator φ satisfies an IQC defined by a dynamic filter Ψ and symmetric matrix M , if for all square summable sequences y and $u = \varphi(y)$ it holds

$$\sum_{t=1}^T \psi_t^\top M \psi_t \geq 0, \quad \psi_t = \Psi \begin{bmatrix} y_t \\ u_t \end{bmatrix} \quad (14)$$

for all $T \geq 1$. Such descriptions have proven particularly useful in the context of first-order algorithms for static optimization [8], [10]. We will introduce an adapted formulation to characterize the input-output relation of time-varying gradients, which will be leveraged to establish our regret bounds.

A. Pointwise IQCs

Many IQCs have been derived for gradients of convex and strongly-convex-smooth functions [8]. For time-varying operators, we can recover their pointwise-in-time formulation.

Lemma 1: Let f_t be m -strongly convex and L -smooth. Take $x_t \in \mathcal{X}$, x_t^* as in (1) and define

$$\psi_t = \begin{bmatrix} LI_d & -I_d \\ -mI_d & I_d \end{bmatrix} \begin{bmatrix} x_t - x_t^* \\ \nabla f_t(x_t) \end{bmatrix} \quad (15a)$$

Then for all t , it holds

$$\psi_t^\top M_1 \psi_t \geq 0 \quad \text{with} \quad M_1 = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \otimes I_d. \quad (15b)$$

Lemma 1 is a simple consequence of [8, Prop. 5]. Notably, (15) satisfies the inequality *pointwise*, i.e. the filter Ψ is static and every single summand is nonnegative.

We will also need a similar characterization of $\mathcal{I}_{\mathcal{X}}$.

Lemma 2: Take $x_t \in \mathcal{X}$ and x_t^* as in (1). Define $\hat{\psi}_t = \text{vec}(x_t - x_t^*, \beta_t)$ for some $\beta_t \in \mathcal{N}_{\mathcal{X}}(x_t)$. Then for M_1 as in (15b) and all t , it holds $\hat{\psi}_t^\top M_1 \hat{\psi}_t \geq 0$.

Note that Lemma 2 is simply the statement that the normal cone of a convex set is a monotone operator.

B. Variational IQCs

It is well known that introducing a dynamic mapping from $(x_t - x_t^*, \nabla f_t(x_t)) \mapsto \psi_t$ can lead to a less conservative input-output characterization of the gradient operator [8], [10]. Unfortunately, most dynamic IQC results are not applicable to time-varying operators. However, recently the notion of variational IQCs has been introduced [11].

Proposition 3: Let f_t be m -strongly convex and L -smooth. Define the variational measures $\Delta x_t^* = x_t^* - x_{t+1}^*$ and $\Delta \delta_t(\cdot) = \nabla f_t(\cdot) - \nabla f_{t+1}(\cdot)$. Consider the mapping

$$\psi_t = \Psi_f \begin{bmatrix} x_t - x_t^* \\ \nabla f_t(x_t) \\ \Delta x_t^* \\ \Delta \delta_t(x_t) \end{bmatrix} \quad (16a)$$

with the linear filter

$$\Psi_f = \left[\begin{array}{cccc|cccc} 0 & 0 & 0 & 0 & I_d & 0 & I_d & 0 \\ 0 & 0 & 0 & 0 & 0 & I_d & 0 & -I_d \\ 0 & 0 & 0 & 0 & -mI_d & I_d & 0 & 0 \\ 0 & 0 & 0 & 0 & aI_d & 0 & 0 & 0 \\ \hline -LI_d & I_d & 0 & 0 & LI_d & -I_d & 0 & 0 \\ 0 & 0 & 0 & 0 & -mI_d & I_d & 0 & 0 \\ 0 & 0 & I_d & 0 & 0 & 0 & 0 & 0 \\ aI_d & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & I_d & 0 & 0 & 0 & 0 \\ -mI_d & I_d & 0 & 0 & 0 & 0 & 0 & 0 \end{array} \right], \quad (16b)$$

where $a := \sqrt{\frac{m(L-m)}{2}}$. Then the quadratic inequality

$$\sum_{t=1}^T \psi_t^\top M_2 \psi_t \geq -4(L-m) \mathcal{V}_T \quad (16c)$$

holds with \mathcal{V}_T as defined in (4) and the blockdiagonal matrix $M_2 = \text{blkdiag}(\frac{1}{2}M_1, [\frac{1}{2}I_d \ -I_d], \frac{1}{2}[\frac{1}{2}I_d \ -I_d])$.

Proof: The proof is the result of applying [11, Prop. 4.2] with $\rho = 1$ and leveraging the definition of \mathcal{V}_T (4). ■

Proposition 3 resembles the classical notion of IQCs, with the difference that (i), the integral quadratic term is additionally dependent on the time-variations of the minimizer and gradient, and (ii), the right hand side of the constraint depends on the function value variation. Crucially, Proposition 3 captures multiple sources of time-variation that inherently change the input-output behaviour of the operator. We denote (16) as an instance of a *variational IQC* (vIQC).

IV. MAIN RESULTS

A. Regret with pointwise IQCs

It is instructive to start by showing regret bounds using pointwise IQCs. We assume the following assumption holds.

Assumption 3: The feasible set \mathcal{X} is bounded.

Consider (7c) and define for $i \in \mathbb{I}_p, j \in \mathbb{I}_q$ the filtered vectors

$$\psi_t^i = \begin{bmatrix} LI_d & -I_d \\ -mI_d & I_d \end{bmatrix} \begin{bmatrix} s_t^i - x_t^* \\ \delta_t^i \end{bmatrix}, \quad \hat{\psi}_t^j = \begin{bmatrix} z_t^j - x_t^* \\ g_t^j \end{bmatrix}. \quad (17)$$

Lemma 1 and 2 imply that each $\psi_t^i, \hat{\psi}_t^j$ satisfy the quadratic inequalities $(\psi_t^i)^\top M_1 \psi_t^i \geq 0$ and $(\hat{\psi}_t^j)^\top M_1 \hat{\psi}_t^j \geq 0$, respectively. By stacking all ψ_t^i and $\hat{\psi}_t^j$ into a vector ψ_t , and defining suitable matrices D_Ψ^y and D_Ψ^u , we can write down the more compact relation

$$\psi_t = \begin{bmatrix} D_\Psi^y & D_\Psi^u \end{bmatrix} \begin{bmatrix} y_t - y_t^* \\ u_t \end{bmatrix},$$

where the block rows of D_Ψ^y and D_Ψ^u select the respective components of $y_t - y_t^*$ and u_t to realize $\psi_t^i, \hat{\psi}_t^j$. Moreover, recall that $y_t^* = C\xi_t^*$, such that

$$\psi_t = \underbrace{\begin{bmatrix} D_\Psi^y C & D_\Psi^y D + D_\Psi^u \end{bmatrix}}_{=:\begin{bmatrix} \hat{C} & \hat{D} \end{bmatrix}} \begin{bmatrix} \xi_t - \xi_t^* \\ u_t \end{bmatrix}.$$

We state the following theorem.

Theorem 4: Let Assumptions 1, 2 and 3 hold. If there exists a symmetric matrix $P \in \mathbb{S}^{n_\xi}$ and non-negative vectors $\lambda_p \in \mathbb{R}_{\geq 0}^p, \lambda_q \in \mathbb{R}_{\geq 0}^q$, such that $P \succ 0$ and the inequality

$$\begin{bmatrix} A^\top P A - P & A^\top P B \\ B^\top P A & B^\top P B \end{bmatrix} + \frac{1}{2} \begin{bmatrix} 0 & [C_1^\top \ 0] \\ [C_1] & 0 \end{bmatrix} + \begin{bmatrix} \hat{C} & \hat{D} \end{bmatrix}^\top M_\lambda \begin{bmatrix} \hat{C} & \hat{D} \end{bmatrix} \preceq 0 \quad (18)$$

holds with $M_\lambda = M_1 \otimes \text{diag}(\lambda_p, \lambda_q)$, then we have

$$\mathcal{R}_T = \mathcal{O}(\lambda_{\max}(P) \mathcal{P}_T). \quad (19)$$

Proof: We start by bounding the optimality difference in the P -normed state space, that is

$$\begin{aligned} \|\xi_{t+1} - \xi_{t+1}^*\|_P^2 &= \|\xi_{t+1} - \xi_t^* + \xi_t^* - \xi_{t+1}^*\|_P^2 \\ &= \|\xi_{t+1} - \xi_t^*\|_P^2 + \|\xi_t^* - \xi_{t+1}^*\|_P^2 \\ &\quad + 2(\xi_{t+1} - \xi_t^*)^\top P(\xi_t^* - \xi_{t+1}^*) \\ &\leq \|\xi_{t+1} - \xi_t^*\|_P^2 + 3\lambda_{\max}(P)R\|\xi_t^* - \xi_{t+1}^*\|, \end{aligned} \quad (20)$$

where R is the diameter of the state space, which is finite by Assumption 3 and the observability of (C_1, A) . Next, we can bound the distance $\|\xi_{t+1} - \xi_t^*\|_P^2$ by

$$\begin{aligned} \|\xi_{t+1} - \xi_t^*\|_P^2 &= \|A(\xi_t - \xi_t^*) + Bu_t\|_P^2 \\ &= \begin{bmatrix} \xi_t - \xi_t^* \\ u_t \end{bmatrix}^\top \begin{bmatrix} A^\top P A & A^\top P B \\ B^\top P A & B^\top P B \end{bmatrix} \begin{bmatrix} \xi_t - \xi_t^* \\ u_t \end{bmatrix} \\ &\stackrel{(18)}{\leq} \|\xi_t - \xi_t^*\|_P^2 - (x_t - x_t^*)^\top \nabla f_t(x_t) \\ &\quad - \sum_{i=1}^p \lambda_p^i (\psi_t^i)^\top M_1 \psi_t^i - \sum_{j=1}^q \lambda_q^j (\hat{\psi}_t^j)^\top M_1 \hat{\psi}_t^j. \end{aligned}$$

The inequality follows by left and right multiplying (18) with $\text{vec}(\xi_t - \xi_t^*, u_t)$. We use Lemma 1 and 2, and the fact that by convexity $(x_t - x_t^*)^\top \nabla f_t(x_t) \geq f(x_t) - f(x_t^*)$, to get

$$\|\xi_{t+1} - \xi_t^*\|_P^2 \leq \|\xi_t - \xi_t^*\|_P^2 - (f(x_t) - f(x_t^*)). \quad (21)$$

Combining (20) and (21) yields

$$\begin{aligned} f(x_t) - f(x_t^*) &\leq \|\xi_t - \xi_t^*\|_P^2 - \|\xi_{t+1} - \xi_{t+1}^*\|_P^2 \\ &\quad + 3\lambda_{\max}(P)R\|\xi_t^* - \xi_{t+1}^*\|. \end{aligned} \quad (22)$$

Finally, summing from $t = 1$ to $t = T$ we get

$$\begin{aligned} \sum_{t=1}^T f(x_t) - f(x_t^*) &\leq \|\xi_1 - \xi_1^*\|_P^2 - \|\xi_{T+1} - \xi_{T+1}^*\|_P^2 \\ &\quad + 3\lambda_{\max}(P)R \sum_{t=1}^T \|\xi_t^* - \xi_{t+1}^*\|. \end{aligned} \quad (23)$$

The left-hand side corresponds to \mathcal{R}_T and, since $\xi_t^* = Ux_t^*$, the last sum is $\mathcal{O}(\mathcal{P}_T)$. ■

Theorem 4 shows that, given any algorithm satisfying the structural assumptions, proving a regret upper bound boils down to a feasibility problem in the form of a Linear Matrix Inequality (LMI). Therefore, Theorem 4 offers an *automated* way to establish regret bounds, without the need for ad-hoc individual proofs.

Qualitatively, (23) shows that the regret grows sublinearly in T if the path length \mathcal{P}_T does. Meanwhile, the sensitivity of the regret bound *w.r.t.* the path length is determined by the maximum eigenvalue of the variable P , which can be optimized for by framing (18) as the feasible set of an SDP whose objective it is to minimize the spectral radius of P .

We investigate this path length sensitivity for different algorithms with a numerical study¹. In particular, we compare O-GD, Multi-step O-GD, O-NM, and O-AGD [14]. We

¹The source code for all numerical experiments can be accessed at: <https://github.com/col-tasas/2025-oco-with-iqcs>.




Fig. 1: Regret sensitivity to \mathcal{P}_T over function class condition ratio, according to Theorem 4.

study the value of $\lambda_{\max}(P)$ for different convexity/Lipschitz moduli m and L , and visualize the result over the function class condition ratio $\frac{L}{m}$. The results are shown in Fig. 1. We observe that most algorithms attain a finite regret bound, with only O-NM growing unboundedly around $\frac{L}{m} \approx 8.5$. We observe how the factor $\lambda_{\max}(P)$ is proportional to $\frac{L}{m}$, indicating a fundamental dependence of (19) on the function class condition ratio. Note that detrimental effects of large $\frac{L}{m}$ are well known in static optimization, but has not been particularly emphasized in the OCO field. Interestingly, we observe that, in terms of the upper bound (19), accelerated O-GD performs best among all compared algorithms, and also that 10-Step O-GD performs better than its counterparts with fewer steps.

B. Regret with variational IQCs

The use of pointwise IQC for the characterization of slope-restricted nonlinearities is a known source of conservatism [8]. We therefore now develop a regret bound that leverages the variational IQCs.

We introduce the fixed-point variation $\Delta\xi_t^* := \xi_t^* - \xi_{t+1}^*$. Using this, we can write the evolution of (7) in error coordinates $\tilde{\xi}_t := \xi_t - \xi_t^*$ as

$$\begin{aligned}\tilde{\xi}_{t+1} &= A\tilde{\xi}_t + Bu_t + \Delta\xi_t^* \\ \tilde{y}_t &= C\tilde{\xi}_t + Du_t.\end{aligned}\quad (24)$$

We now apply Proposition 3 to each subcomponent of (s_t, δ_t) . We moreover apply Proposition 3 to (z_t, g_t) by letting $f_t \leftarrow \mathcal{I}_{\mathcal{X}}$, $\nabla f_t \leftarrow \mathcal{N}_{\mathcal{X}}$ and $m \rightarrow 0$ and $L \rightarrow \infty$. Since $\mathcal{I}_{\mathcal{X}}$ is not time-varying, the gradient and function variation is not needed, and the vIQC simplifies drastically. That is, we can define

$$\psi_t^i = \Psi_f \begin{bmatrix} s_t^i - x_t^* \\ \delta_t^i \\ \Delta x_t^* \\ \Delta\delta_t(s_t^i) \end{bmatrix}, \quad \psi_t^j = \Psi_{\mathcal{I}} \begin{bmatrix} z_t^j - x_t^* \\ g_t^j \\ \Delta x_t^* \end{bmatrix} \quad (25)$$

with Ψ_f and $\Delta\delta_t$ as defined in Proposition 3 and $\Psi_{\mathcal{I}}$ adapted accordingly, and it holds $\sum_{t=1}^T (\psi_t^i)^\top M_2 \psi_t^i \geq -4(L-m)\mathcal{V}_T$ and $\sum_{t=1}^T (\hat{\psi}_t^j)^\top M_1 \hat{\psi}_t^j \geq 0$ for all $i \in \mathbb{I}_p$, $j \in \mathbb{I}_q$. We refer to the extended version of this work [21, Appendix A] for the discussion on $\Psi_{\mathcal{I}}$.

Analogously to the last section, we stack all ψ_t^i and $\hat{\psi}_t^j$ into a vector ψ_t , to get the compact vIQC

$$\psi_t = \left[\begin{array}{c|c|c|c|c} A_\Psi & B_\Psi^y & B_\Psi^u & B_\Psi^{\Delta\xi} & B_\Psi^{\Delta\delta} \\ \hline C_\Psi & D_\Psi^y & D_\Psi^u & D_\Psi^{\Delta\xi} & D_\Psi^{\Delta\delta} \end{array} \right] \begin{bmatrix} \tilde{y}_t \\ u_t \\ \Delta\xi_t^* \\ \Delta\delta_t \end{bmatrix}$$

where we defined $\Delta\delta_t := \text{vec}(\Delta\delta_t(s_t^1), \dots, \Delta\delta_t(s_t^p))$. The respective matrices result from a straightforward rearrangement and stacking of the filter matrices. In particular, we leveraged that $(1_{p+q} \otimes I_d)\Delta x_t^* = C\Delta\xi_t^*$. Together with (24), we can build the augmented plant

$$\left[\begin{array}{c|c|c|c|c} \hat{A} & \hat{B}_u & \hat{B}_{\Delta\xi} & \hat{B}_{\Delta\delta} \\ \hline \hat{C} & \hat{D}_u & \hat{D}_{\Delta\xi} & \hat{D}_{\Delta\delta} \end{array} \right] \triangleq \left[\begin{array}{c|c|c|c|c} A & 0 & B & I & 0 \\ \hline B_\Psi^y C & A_\Psi & B_\Psi^u & B_\Psi^{\Delta\xi} & B_\Psi^{\Delta\delta} \\ \hline D_\Psi^y C & C_\Psi & D_\Psi^u & D_\Psi^{\Delta\xi} & D_\Psi^{\Delta\delta} \end{array} \right], \quad (26)$$

defining the mapping $(u_t, \Delta\xi_t^*, \Delta\delta_t) \mapsto \psi_t$.

Theorem 5: Let Assumptions 1 and 2 hold. Consider the augmented plant (26). If there exists a symmetric matrix $P \in \mathbb{S}^{n_\xi + n_\zeta}$, non-negative vectors $\lambda_p \in \mathbb{R}_{\geq 0}^p$, $\lambda_q \in \mathbb{R}_{\geq 0}^q$ and scalars $\gamma_{\Delta\xi}, \gamma_{\Delta\delta} > 0$, such that $P \succ 0$ and the LMI

$$\begin{aligned}[\star]^\top & \begin{bmatrix} -P & & & & \\ & P & & & \\ & & M_\lambda & & \\ & & & -\gamma_{\Delta\xi} I & \\ & & & & -\gamma_{\Delta\delta} I \end{bmatrix} \begin{bmatrix} I & 0 & 0 & 0 \\ \hat{A} & \hat{B}_u & \hat{B}_{\Delta\xi} & \hat{B}_{\Delta\delta} \\ \hat{C} & \hat{D}_u & \hat{D}_{\Delta\xi} & \hat{D}_{\Delta\delta} \\ 0 & 0 & I & 0 \\ 0 & 0 & 0 & I \end{bmatrix} \\ & + \frac{1}{2} \begin{bmatrix} 0 & [C_1^\top \ 0] \\ [C_1 \ 0] & 0 \end{bmatrix} \leq 0 \quad (27)\end{aligned}$$

holds with $M_\lambda = \text{blkdiag}(M_2 \otimes \text{diag}(\lambda_p), M_1 \otimes \text{diag}(\lambda_q))$, then we have

$$\mathcal{R}_T = \mathcal{O}(\gamma_1 \mathcal{S}_T + \gamma_2 \mathcal{G}_T + \gamma_3 \mathcal{V}_T), \quad (28)$$

with $\gamma_1 = \gamma_{\Delta\xi}$, $\gamma_2 = p\gamma_{\Delta\delta}$, $\gamma_3 = (L-m) \sum_{i=1}^p \lambda_p^i$.

Proof: Define η_t as the state of (26). Left and right multiply (27) by $\text{vec}(\eta_t, u_t, \Delta\xi_t^*, \Delta\delta_t)$, to obtain the inequality

$$\begin{aligned}-\|\eta_t\|_P^2 + \|\eta_{t+1}\|_P^2 - \gamma_{\Delta\xi} \|\Delta\xi_t^*\|^2 - \gamma_{\Delta\delta} \sum_{i=1}^p \|\Delta\delta_t(s_t^i)\|^2 \\ + \sum_{i=1}^p \lambda_p^i (\psi_t^i)^\top M_2 \psi_t^i + \sum_{j=1}^q \lambda_q^j (\hat{\psi}_t^j)^\top M_1 \hat{\psi}_t^j \\ + (x_t - x_t^*)^\top \nabla f_t(x_t) \leq 0\end{aligned}$$

Note that $\sum_{t=1}^T \sum_{i=1}^p \|\Delta\delta_t(s_t^i)\|^2 \leq p\mathcal{G}_T$. Leveraging again convexity and Proposition 3, and summing from $t = 1$ to T gives, after telescoping and rearranging terms, that

$$\mathcal{R}_T \leq \|\eta_1\|_P^2 - \|\eta_{T+1}\|_P^2 + \gamma_1 \sum_{t=1}^T \|\Delta\xi_t^*\|^2 + \gamma_2 \mathcal{G}_T + \gamma_3 \mathcal{V}_T. \quad (29)$$

By $\xi_t^* = Ux_t^*$, the sum is $\mathcal{O}(\mathcal{S}_T)$, which gives the result. ■

In contrast to (19), the bound in (28) depends on the squared path length, gradient variation, and function variation, with sensitivities γ_1 , γ_2 , and γ_3 that are decision




Fig. 2: Regret sensitivity to \mathcal{S}_T (γ_1), \mathcal{G}_T (γ_2) and \mathcal{V}_T (γ_3) over function class condition ratio, according to Theorem 5.

variables in the LMI (27). These degrees of freedom can be optimized to tighten the upper bound (29), for example by minimizing a weighted sum $k_1\gamma_1 + k_2\gamma_2 + k_3\gamma_3$ for some $k_1, k_2, k_3 > 0$. However, the trade-off between those values depends on problem-specific considerations, i.e. on the foresight on \mathcal{S}_T , \mathcal{G}_T , and \mathcal{V}_T . I.e., setting $k_1 = \mathcal{S}_T$, $k_2 = \mathcal{G}_T$, $k_3 = \mathcal{V}_T$ may be beneficial.

Fig. 2 presents a numerical study investigating the algorithm and function class dependence of these sensitivities, using the same list of algorithms as in Fig. 1. We observe analogously to Fig. 1 a deteriorating effect of large condition ratios $\frac{L}{m}$. Note especially that the relative performance of algorithms differs w.r.t. which variation metric is considered. As an example, 2-step O-GD tends to be favourable when \mathcal{G}_T is expected to dominate (because it achieves consistently the lowest value of γ_2), but may perform worse when \mathcal{S}_T is the primary source of variation. Thus, this worst-case bound may provide guidance on the algorithm choice when prior information is available. Moreover, Theorem 5 extends the regret bound for O-NM to higher condition ratios, indicating a finite regret bound over the whole function class. We observe that the bounds from both theorems offer distinct information, highlighting the complementary nature of both.

V. CONCLUSION

In this paper, we have presented a novel framework for an automated regret analysis of first-order algorithms in OCO. By recasting the algorithm as a feedback interconnection of

a linear system and a time-varying oracle we are able to provide an alternative proof strategy and a computational tool to quantify the regret of generalized first-order optimization algorithms. This new analysis framework represents a new viewpoint on OCO and can contribute to obtain a more systematic way to show regret. Future work may include the use of further robust control tools, such as algorithm synthesis or robustness analyses for gradient errors.

REFERENCES

- [1] M. Zinkevich, "Online convex programming and generalized infinitesimal gradient ascent," in *Proceedings of the 20th International Conference on Machine Learning*, 2003.
- [2] A. Jadbabaie, A. Rakhlin, S. Shahrampour, and K. Sridharan, "Online Optimization: Competing with Dynamic Comparators," in *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Statistics*, 2015.
- [3] M. Nonhoff and M. A. Müller, "An online convex optimization algorithm for controlling linear systems with state and input constraints," in *2021 American Control Conference (ACC)*, 2021.
- [4] A. Karapetyan, A. Iannelli, and J. Lygeros, "On the regret of H_∞ control," in *2022 IEEE 61st Conference on Decision and Control (CDC)*, 2022.
- [5] E. Hazan, A. Agarwal, and S. Kale, "Logarithmic regret algorithms for online convex optimization," *Machine Learning*, 2007.
- [6] C. Hu, W. Pan, and J. Kwok, "Accelerated gradient methods for stochastic optimization and online learning," in *Advances in Neural Information Processing Systems*, 2009.
- [7] P. Zhao and L. Zhang, "Improved analysis for dynamic regret of strongly convex and smooth functions," in *Proceedings of the 3rd Conference on Learning for Dynamics and Control*, 2021.
- [8] L. Lessard, B. Recht, and A. Packard, "Analysis and design of optimization algorithms via integral quadratic constraints," *SIAM Journal on Optimization*, 2016.
- [9] S. Michalowsky, C. Scherer, and C. Ebenbauer, "Robust and structure exploiting optimisation algorithms: an integral quadratic constraint approach," *International Journal of Control*, 2021.
- [10] C. Scherer and C. Ebenbauer, "Convex synthesis of accelerated gradient algorithms," *SIAM Journal on Control and Optimization*, 2021.
- [11] F. Jakob and A. Iannelli, "A linear parameter-varying framework for the analysis of time-varying optimization algorithms." *arXiv:2501.07461*, v2, 2025.
- [12] E. Dall'Anese, A. Simonetto, S. Becker, and L. Madden, "Optimization and learning with information streams: Time-varying algorithms and applications," *IEEE Signal Processing Magazine*, 2020.
- [13] E. Hazan, *Introduction to Online Convex Optimization*. Cambridge, MA: The MIT Press, 2022.
- [14] T. Yang, "Online accelerated gradient descent and variational regret bounds," in <https://homepage.cs.uiowa.edu/~tyng/papers/online-accelerated-gradient-method.pdf>, 2016. Accessed: 2025-03-25.
- [15] 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.
- [16] Y.-F. Xie, P. Zhao, and Z.-H. Zhou, "Gradient-variation online learning under generalized smoothness," in *Advances in Neural Information Processing Systems*, 2024.
- [17] A. Beck, *First-order methods in optimization*, vol. 25 of *MOS-SIAM series on optimization*. Philadelphia: Society for Industrial and Applied Mathematics (SIAM), 2017.
- [18] M. Upadhyaya, S. Banert, A. Taylor, and P. Giselsson, "Automated tight Lyapunov analysis for first-order methods," *Math. Program.*, 2025.
- [19] E. C. Hall and R. M. Willett, "Online convex optimization in dynamic environments," *IEEE Journal of Selected Topics in Signal Processing*, 2015.
- [20] A. Megretski and A. Rantzer, "System analysis via integral quadratic constraints," *IEEE Transactions on Automatic Control*, 1997.
- [21] F. Jakob and A. Iannelli, "Online convex optimization and integral quadratic constraints: An automated approach to regret analysis." *arXiv:2503.23600*, 2025.