An Online Algorithm for Smoothed Regression and LQR Control

Gautam Goel Adam Wierman

Abstract

We consider Online Convex Optimization (OCO) in the setting where the costs are m-strongly convex and the online learner pays a switching cost for changing decisions between rounds. We show that the recently proposed Online Balanced Descent (OBD) algorithm is constant competitive in this setting, with competitive ratio $3 + O(1/m)$, irrespective of the ambient dimension. We demonstrate the generality of our approach by showing that the OBD framework can be used to construct competitive algorithms for a variety of online problems across learning and control, including online variants of ridge regression, logistic regression, maximum likelihood estimation, and LQR control.

1. Introduction

In this paper we study the problem of *smoothed online convex optimization* (SOCO), a variant of OCO where the online learner incurs a *switching cost* for changing its actions between rounds. More concretely, the online learner plays a series of rounds $t = 1 \dots T$. In each round, the learner receives a convex loss function f_t, picks a point x_t from a convex action space $\chi \subset \mathbb{R}^d$, and pays a *hitting cost* $f_t(x_t)$ as well as a *switching cost* $c(x_t, x_{t-1})$ which penalizes the learner for changing its action between rounds.

This problem was first introduced in the context of the dynamic management of service capacity in data centers (LWAT11), where the switching costs represent the performance and wear-and-tear costs associated with changing server configurations. Since then, SOCO has attracted considerable interest, both theoretical and applied, due to its use in dozens of applications across learning, distributed systems, networking, and control, such as speech animation (KYTM15), video streaming (JdV12), management of electric vehicle charging (KG14), geographical load balancing (LLWA12), and multi-timescale control (GCW17). See (CGW18) for an extensive list of applications.

Unfortunately, despite a large and growing literature, all existing results identifying competitive algorithms for SOCO either (i) place strong restrictions on the action space, (ii), place strong restrictions on the class of loss functions, or (iii) require algorithms to make use of predictions of future cost functions. For example, a series of papers (LWAT11), (BGK$^+$15) developed competitive algorithms for one-dimensional action spaces. Until earlier this year there were no known algorithms that were competitive for SOCO beyond one dimension without requiring the use of predictions. Finally, (CGW18) presented the first algorithm that is constant-competitive beyond one dimension, but the algorithm was shown to be constant competitive only in the case of polyhedral cost functions, a restrictive class that does not include most loss functions used in machine learning. Beyond this result, the most general positive results all assume predictions of future cost functions are available, e.g. (LLWA12), (CAW$^+$15), (CCL$^+$16), (LQL18).

The existing work on SOCO highlights a crucial open question: *Does there exist a competitive algorithm for high-dimensional SOCO problems with cost functions that capture standard losses for online learning problems, e.g., logistic loss or least-squares loss?*

In this paper we answer this question by proving that the recently introduced Online Balanced Descent (OBD) algorithm is constant-competitive for SOCO with strongly convex costs. Additionally, highlighting the importance of the class of strongly convex costs, we show that the OBD framework can be used to construct the first competitive algorithms for problems as diverse as online ridge regression, online logistic regression, and LQR control, which was not possible with previous approaches.

Contributions of this paper. This paper makes two main contributions to the literature on SOCO.

First, in Section 3 we show that OBD is constant competitive for SOCO when the costs are strongly convex (Theorem 1). This establishes OBD as the first constant competitive algorithm for strongly convex costs beyond one-dimensional action spaces. The key to our proof is a novel potential function argument that exploits essential properties of ℓ_2 geometry. In particular, controlling the change in potential rests upon comparing the side lengths of certain triangles (see Figure 1), which can be done via the Law of Cosines.

Second, in Section 4 we show novel applications of OBD to problems arising in statistics, learning, and control, in-

cluding ridge regression, logistic regression, and maximum likelihood estimation. A highlight of this section is a reduction of LQR control to SOCO, giving the first competitive algorithm for LQR control (results for LQR control typically make strong distributional assumptions). We emphasize that none of these applications could be handled by previous work on SOCO, which highlights the importance of deriving a competitive bound for OBD in the strongly convex setting.

Related work. There is a vast literature on OCO; for a recent survey see (H$^+$16). OCO with switching costs was first studied in the scalar setting in (LWAT11), which used SOCO to model dynamic right-sizing in data centers and gave a 3-competitive algorithm. In subsequent work, (BGK$^+$15) improved the competitive ratio to 2, also in the scalar setting. The first constant-competitive algorithm beyond one dimension was given in (CGW18), which introduced the OBD framework and showed that it was competitive for SOCO with polyhedral costs. The results in this paper highlight that OBD is also constant-competitive for strongly convex cost functions, a class that is particularly important for learning and control applications, and is wholly disjoint from the class of polyhedral cost functions when the minimizer of the cost function is zero.

A special case of SOCO is the Convex Body Chasing problem, first introduced in (FL93). The connection between Convex Body Chasing and SOCO was observed in (ABN$^+$16). A recent series of papers (BBE$^+$18), (ABC$^+$18) identified competitive algorithms in the setting where the bodies are nested.

Before this paper, the only class of SOCO problems for which positive results for strongly convex cost functions existed is when the learner had access to accurate predictions of future cost functions. The study of SOCO with predictions began with (LLWA12) and then continued with a stream of work in the following years, e.g., (CAW$^+$15; CCL$^+$16). The most relevant to this work is (LQL18), which shows a lower bound on the dynamic regret of SOCO with strongly convex cost functions.

In Section 4 we apply OBD to diverse problems like maximum likelihood estimation and LQR control. These problems have been widely studied; we refer the reader to (BV04) and (AM10) for a survey.

Finally, we note that SOCO can be viewed as a continuous version of the classic Metrical Task Systems (MTS) problem, one of the most widely studied problems in the online algorithms community, e.g (BLS92), (BBBT97), (BB00). A special case of the MTS is the celebrated k-server problem, first proposed in (MMS90), which has received significant attention in recent years, for example in (BCL$^+$18).

2. Smoothed Online Convex Optimization

An instance of SOCO consists of a convex action set $\chi \subset \mathbb{R}^d$, an initial point $x_0 \in \chi$, a sequence of non-negative convex costs $f_1 \dots f_T : \mathbb{R}^d \to \mathbb{R}^+$, and a non-negative function $c : \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}^+$. In each round t, the online learner observes the cost function f_t, picks a point x_t, and pays the sum of the *hitting cost* $f_t(x_t)$ and the *movement* or *switching cost* $c(x_t, x_{t-1})$. The switching cost acts as a regularizer, penalizing the online learner for changing its decisions between rounds. The goal of the online learner is to minimize its aggregate cost so as to approximate the offline optimal cost:

$$\min_{x_1 \dots x_T \in \chi} \sum_{t=1}^{T} f_t(x_t) + c(x_t, x_{t-1})$$

More generally, x_t could be matrix-valued and f_t, c could be functions on matrices. Note that we make no restrictions on the sequence of cost functions $f_1 \dots f_T$ other than strong convexity; they could be adversarial, or even adaptively chosen to hurt the online learner.

We emphasize that SOCO differs from OCO in two important ways. Firstly, unlike in OCO, the costs incurred in each round of SOCO depend on the previous choice, coupling the online learner's decisions across rounds. Secondly, the online learner can observe the cost function f_t before picking x_t. This is a standard assumption in the SOCO literature, e.g. in (BGK$^+$15), (LWAT11), (CGW18) and isolates the complexity of SOCO onto the coupling across timesteps due to the switching costs instead of the uncertainty in the costs.

In this paper, we measure the performance of OBD in terms of its *dynamic regret* and *competitive ratio*. The dynamic regret is defined as

$$\sum_{t=1}^{T} f_t(x_t) + c(x_t, x_{t-1}) - \left[\sum_{t=1}^{T} f_t(x_t^*) + c(x_t^*, x_{t-1}^*)\right]$$

Here $x_1 \dots x_T$ are the points picked by the online learner and $x_1^* \dots x_T^*$ are the offline optimal points. We note that this is a more natural performance metric for SOCO than static regret, since the main motivation for SOCO is to understand the effects of switching costs on online learning. In contrast, in the static regret setting the comparator never moves and hence incurs no switching cost, making it a less ideal performance metric for SOCO.

Instead of using an additive metric the competitive ratio uses a multiplicative metric:

$$\frac{\sum_{t=1}^{T} f_t(x_t) + c(x_t, x_{t-1})}{\sum_{t=1}^{T} f_t(x_t^*) + c(x_t^*, x_{t-1}^*)}$$

We note that (ABL$^+$13) showed that, in general, no online algorithm can have both sublinear static regret and constant competitive ratio.

Much attention has been focused on the setting where the switching cost is a norm: $c(x_t, x_{t-1}) = \|x_t - x_{t-1}\|$, e.g. (LWAT11), (BGK$^+$15). Note that in the one-dimensional setting, all ℓ_p norms are identical, making the choice of norm somewhat vacuous. The first algorithm to work beyond the one dimensional setting was proposed in (CGW18), which considered a setting where the switching cost is given by the Euclidean distance and the loss functions are *polyhedral*, meaning that they at grow at least linearly as one moves away from the minimizer.

We instead focus on the setting where the cost functions $f_1 \dots f_T$ are m-strongly convex with respect to the Euclidean norm and the switching cost is quadratic:

$$c(x_t, x_{t-1}) = \frac{1}{2}\|x_t - x_{t-1}\|_2^2$$

In Section 4, we show that OBD can be used with many important loss functions, such as the least-squares loss and the ℓ_2 regularized logistic loss, none of which could be handled by previous work.

We assume that the domain χ is all of $\mathcal{R}^d$. Note that this presents no real restriction, since we can always define $f_t(x) = \infty$ for all $x \notin \chi$. The objective becomes

$$\min_{x_1 \dots x_T \in \mathbb{R}^d} \sum_{t=1}^{T} f_t(x_t) + \frac{1}{2}\|x_t - x_{t-1}\|_2^2 \tag{1}$$

Notation. We use $\|\cdot\|$ to denote the ℓ_2 norm. We often use H_t and M_t to denote the hitting cost $f_t(x_t)$ and the movement cost $\frac{1}{2}\|x_t - x_{t-1}\|^2$, respectively. The offline costs H_t^* and M_t^* are defined analogously. We let ALG denote the total cost incurred by OBD across all rounds and define OPT to be the analogous offline cost. We let v_t denote the minimizer of the cost function f_t.

Algorithm 1 Online Balanced Descent (OBD)

1: **for** $t = 1, \dots, T$ **do**
2: Receive f_t. Let v_t be the minimizer of f_t.
3: Let $x(l) = \Pi_{K_t^l}(x_{t-1})$. Initialize $l = f_t(v_t)$. Here K_t^l is the l-sublevel set of f_t, i.e., $K_t^l = \{x \mid f_t(x) \le l\}$.
4: Increase l. Stop either when $x(\ell) = v_t$ or $\frac{1}{2}\|x(l) - x_{t-1}\|_2^2 = \beta l$.
5: Set $x_t = x(l)$.
6: **end for**

3. A Competitive Algorithm

Our main technical result shows that a recently proposed algorithm, Online Balanced Descent (OBD), is constant competitive for SOCO problems with strongly convex cost functions.

OBD was introduced in (CGW18), where it was analyzed for the class of polyhedral costs. The detailed workings of OBD are summarized in Algorithm 1. The key insight of OBD is to exploit the full geometry of the level sets of the current cost function f_t when choosing the point x_t in such a way as to take switching costs into account.

OBD works by iteratively projecting the previously chosen point onto a level set of the current cost function. The level set K_t picked by OBD is the level set such that the switching cost incurred while traveling from x_{t-1} to K_t is equal to $\beta f_t(x_t)$, where x_t is the projection onto K_t and β is the *balance parameter* which can be tuned to get different performance guarantees. We note that OBD can be efficiently implemented via a binary search over the level sets (CGW18).

We can now state our main result, a bound on the competitive ratio of OBD for strongly convex costs.

Theorem 1. *OBD is competitive for the problem (1) for all $\beta > \frac{4}{m}$. Furthermore, if β is set to be $2 + \frac{10}{m}$, the competitive ratio of OBD is at most $3 + O(1/m)$, irrespective of the ambient dimension.*

We note that (CGW18) proved a bound on the competitive ratio of OBD of the form $3+O(1/\alpha)$ where α measures the "steepness" of the costs. While this superficially resembles the bound in Theorem 1, we emphasize that the settings are quite different; their work applied to the class of polyhedral cost functions while we focus on strongly convex cost functions. In the case where the cost functions have minimum value zero these classes are wholly disjoint. We are led to consider strongly convex costs due to the fact many common learning and control problems have loss functions that are strongly convex (e.g., see Section 5). Until this paper, there existed no competitive algorithms for SOCO problems with strongly convex costs.

4. Applications

In this section we show several applications of OBD to diverse problems across learning and control. We emphasize that none of these applications would be possible without a competitive algorithm for strongly convex costs, which has not been attainable with previous approaches.

4.1. Smoothed Online Regression

We consider the problem of a learner who wishes to fit a series of regularized regressors or classifiers to a changing dataset, without changing the estimators too much between rounds. This is naturally modeled by the objective

$$\min_{\theta_1 \dots \theta_T \in \mathbb{R}^d} \sum_{t=1}^{T} f_t(\theta_t) + \frac{\lambda_1}{2}\|\theta_t\|_2^2 + \frac{\lambda_2}{2}\|\theta_t - \theta_{t-1}\|_2^2 \quad (2)$$

Here f_t represents an estimation or classification task at timestep t, θ_t is the regressor at time t, and λ_1, λ_2 are parameters that control the strength of the ℓ_2 and smoothing regularizations, respectively. We impose no constraint on f_t other than convexity; in particular, f_t need not be strongly convex (though if f_t happens to be strongly convex, we can optionally drop the regularization term $\frac{\lambda_1}{2}\|\theta_t\|^2$). OBD gives a constant-competitive algorithm in this setting:

Corollary 1. *The competitive ratio of OBD with balance parameter $\beta = 2 + \frac{10}{m}$ on problem 2 is $3 + O(\lambda_2/\lambda_1)$.*

Before we turn to the proof, we emphasize that the bound on competitive ratio does not vary with respect to dimension; hence OBD can be applied to estimation problems with thousands or millions of parameters.

Proof. We first divide the objective by λ_2. Notice that the function $\tilde{f}_t(\theta) = \frac{1}{\lambda_2} f_t(\theta) + \frac{\lambda_1}{2\lambda_2}\|\theta_t\|_2^2$ is (λ_1/λ_2)-strongly convex in θ whenever f_t is convex, hence Theorem 1 implies that OBD achieves competitive ratio $3 + O(\lambda_2/\lambda_1)$. □

Our approach applies to many common learning problems:

- **Ridge Regression**. We take $f_t(\theta) = \|X_t\theta - y_t\|^2$, where $X_t \in \mathbb{R}^{n\times d}$ is a data matrix and $y_t \in \mathbb{R}^n$ is the response variable.
- **Logistic Regression**. We take $f_t(\theta) = -\frac{1}{n_t}\sum_{i=1}^{n_t} \log\left(1 + e^{-y_{i,t}\theta^\top x_{i,t}}\right)$ where $x_{i,t} \in \mathbb{R}^d$ is a vector of features, $y_{i,t} \in \{0, 1\}$ is a binary outcome, and n_t is the number of samples in round t. OBD hence fits a series of binary classifiers which don't vary too much between rounds. Our approach easily extends to the multiclass setting as well.
- **Maximum Likelihood Estimation.** More generally, we can perform smoothed online maximum likelihood estimation using OBD. Here θ_t are parameters of a model and $f_t(\theta_t)$ is the likelihood function of some dataset at time t. If the likelihood function is convex then OBD can be applied. For example, the problem of estimating a series of covariance matrices Σ_t of a series of Gaussian distributions $\mathcal{N}(0, \Sigma_t)$ given independent samples arranged as a matrix Y_t can be posed as the problem

  $$\min_{\Sigma}(\Sigma Y) - \log\det\Sigma$$

  which is a convex optimization problem (see (BV04), p. 357). We can apply OBD over the set of positive definite matrices to find a series of covariance matrices Σ_t that fit the data well but don't vary too much between rounds.

4.2. Linear Quadratic Regulator (LQR) control

Our second application comes from the controls community. Consider the classical problem of LQR control:

$$\min_{u_1 \dots u_T} \sum_{t=1}^{T} \frac{1}{2} x_t' Q_t x_t + \frac{1}{2} u_t' R u_t$$

with dynamics given by

$$x_{t+1} = Ax_t + Bu_t + w_t$$

Here u is a control action, x a state variable, and Q_t, R are assumed to be positive definite. Usually, the noise increments w_t are assumed to be i.i.d. Gaussian, and the goal is to design a control policy to minimize the expected cost. Instead of an in-expectation result, we can use OBD to design a controller with a strong pathwise guarantee, with no distributional or boundedness assumptions on the noise. We focus on the setting where $A = I$, i.e. the system is stationary in the absence of noise or control actions.

Corollary 2. *Suppose that $A = I$ and B is invertible, and the matrices Q_t each have their lowest eigenvalue bounded below by $\lambda > 0$. The LQR problem can be rewritten as a SOCO problem, and the competitive ratio of OBD is*

$$3 + O\left(\frac{\lambda_{max}(R)}{\lambda_{min}(B)^2\lambda}\right).$$

Note that $\lambda_{min}(B)$ can be interpreted as a lower bound on the gain of the control action u; intuitively, systems with high control gain are easier to regulate, since each control action gets amplified. Similarly, it is intuitive that as $\lambda_{max}(R)$ decreases the competitive ratio improves, since R controls the cost incurred by using the controller.

Appendices

Proof of Theorem 1

To prove Theorem 1, we use the potential function $\phi(x, x^*) = \eta\|x - x^*\|^2$. Clearly $\phi(x, x^*) \geq 0$ and $\phi(x_0, x_0^*) = 0$. Before we turn to the proof of Theorem 1, we prove a series of crucial lemmas relating to the potential function. Lemmas 1 and 2 show how the potential changes depending on the relative positions of its arguments, and highlight the role of the geometry associated with the ℓ_2 norm. Lemma 3 relates the potential to the hitting costs at every timestep.

Lemma 1. *The change in potential satisfies*

$$\phi(a, c) - \phi(a, b) \leq -\phi(b, c)$$

for all $a, b, c \in \mathbb{R}^d$ such that the angle θ between the vectors $a - c$ and $b - c$ lies in $[\pi/2, 3\pi/2]$.

Proof. Consider the triangle with vertices a, b, c. According to the Law of Cosines we have:

$$\|a - b\|^2 = \|a - c\|^2 + \|b - c\|^2 - 2\|a - c\|\|b - c\| \cos\theta.$$

Rearranging gives

$$\|a - c\|^2 - \|a - b\|^2 = -\|b - c\|^2 + 2\|a - c\|\|b - c\| \cos\theta.$$

Since θ lies in $[\pi/2, 3\pi/2]$, the cosine term must be nonpositive, immediately yielding the claim. □

Lemma 2. *The change in potential satisfies*

$$\phi(a, c) - \phi(a, b) \leq 2\phi(b, c) + \phi(a, b)$$

for all $a, b, c \in \mathbb{R}^d$.

Proof. We apply the Law of Cosines again:

$$\|a - c\|^2 - \|a - b\|^2 = \|b - c\|^2 - 2\|a - b\|\|b - c\| \cos\theta$$

where θ is the angle between the vectors $a - b$ and $b - c$. The second term on the right is at most $2\|a - b\|\|b - c\|$; applying the AM-GM inequality to this expression gives the claim. □

Lemma 3. *At all timesteps t, the potential satisfies*

$$\phi(x_t, x_t^*) \leq \frac{4\eta}{m} H_t + \frac{4\eta}{m} H_t^*$$

Proof. We have

$$\begin{aligned} \phi(x_t, x_t^*) &= \eta\|x_t - x_t^*\|^2 \\ &\leq \eta(\|x_t - v_t\| + \|x_t^* - v_t\|)^2 \\ &\leq 2\eta\|x_t - v_t\|^2 + 2\eta\|x_t^* - v_t\|^2 \\ &\leq \frac{4\eta}{m} H_t + \frac{4\eta}{m} H_t^* \end{aligned}$$

The first inequality is just the triangle inequality; the second follows from the AM-GM inequality; and in the last step we used the fact that $f_t(x) \geq \frac{m}{2}\|x - v_t\|^2$. □

Now we return to the proof of Theorem 1. Note that it suffices to show that OBD is constant competitive in the case where minimum value of each cost function is zero, since otherwise the competitive ratio can only improve. In this case, we always have $M_t = \beta H_t$ since H_t shrinks to zero as we move towards the minimizer while M_t increases.

Proof. To bound the cost charged to OBD in each step, we first consider two cases.

Case 1: $H_t \leq H_t^*$

This case is easy; the cost charged to OBD is

$$\begin{aligned} H_t + M_t + \Delta\phi &\leq H_t + M_t + \frac{4\eta}{m} H_t + \frac{4\eta}{m} H_t^* \\ &\leq \left(1 + \beta + \frac{8\eta}{m}\right) H_t^* \end{aligned}$$

Here in the first inequality we threw away the negative potential term and used Lemma 3. In the second inequality we used the fact that $M_t = \beta H_t$ and the inequality defining the case.

Case 2: $H_t > H_t^*$

This is the hard case. Unlike in the previous case, we cannot directly bound the cost charged to OBD in terms of the offline cost, since H_t^* is less than H_t. Our strategy will be to show that the change in potential was negative, offsetting the hitting and movement costs incurred by OBD.

Since $H_t^* < H_t$, the offline point x_t^* must lie strictly in the interior of K_t, where K_t is the H_t-level set of f_t. Notice that the angle θ made between the line segments $\overline{x_{t-1}x_t}$ and $\overline{x_t x_t^*}$ must be obtuse, since x_t was the projection onto the level set, and x_t^* lies strictly on the opposite side of the supporting hyperplane tangent to K_t at x_t (see Figure 1). We have

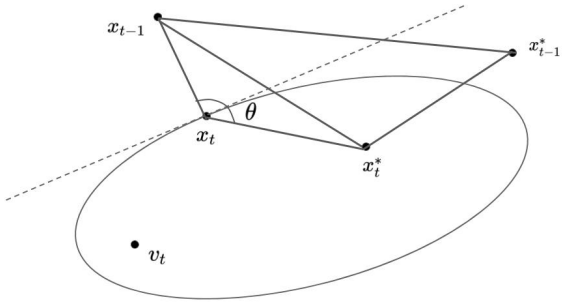

Figure 1. An illustration of the situation arising in Case 2. The ellipse represents the level set K_t of f_t which OBD has selected; x_t is the projection of x_{t-1} onto this level set. The offline point x^*_t must lie strictly inside the level set, since $H^*_t < H_t$. The angle θ must be obtuse, since x^*_t lies on the other side of the supporting hyperplane at x_t, shown by the dotted line.

$$\begin{aligned} H_t + M_t + \Delta\phi &= H_t + M_t + (\phi(x_t, x^*_t) - \phi(x_{t-1}, x^*_t)) \\ &\quad + \big(\phi(x_{t-1}, x^*_t) - \phi(x_{t-1}, x^*_{t-1})\big) \\ &\leq H_t + M_t - \phi(x_t, x_{t-1}) \\ &\quad + 2\phi(x^*_t, x^*_{t-1}) + \phi(x_{t-1}, x^*_{t-1}) \\ &= H_t + M_t - \eta M_t + 2\eta M^*_t \\ &\quad + \phi(x_{t-1}, x^*_{t-1}) \\ &\leq \left(1 + \frac{1}{\beta} - \eta\right) M_t + 2\eta M^*_t \\ &\quad + \frac{4\eta}{m} H_{t-1} + \frac{4\eta}{m} H^*_{t-1} \end{aligned}$$

In the first inequality we use Lemma 1 to bound the first change in potential and Lemma 2 to bound the second. In the second inequality we apply the fact that $M_t = \beta H_t$ and Lemma 3.

Bounding the competitive ratio

We have now bounded the cost charged to OBD in each of the two cases. Putting both cases together, we see that we always have

$$\begin{aligned} H_t + M_t + \Delta\phi &\leq \left(1 + \beta + \frac{8\eta}{m}\right) H^*_t + 2\eta M^*_t \\ &\quad + \frac{4\eta}{m} H_{t-1} + \frac{4\eta}{m} H^*_{t-1} \end{aligned}$$

where we assume that β, η were picked so that $\eta \geq 1 + \frac{1}{\beta}$. Adding up across all timesteps and collecting terms, we have:

$$\begin{aligned} \sum_{t=1}^{T} H_t + M_t &\leq \sum_{t=1}^{T} \left(1 + \beta + \frac{12\eta}{m}\right) H^*_t \\ &\quad + \sum_{t=1}^{T} 2\eta M^*_t + \sum_{t=1}^{T} \frac{4\eta}{m} H_{t-1} \end{aligned}$$

By the balance condition, $H_t = \frac{1}{1+\beta}(H_t + M_t)$. We immediately obtain

$$\begin{aligned} ALG &\leq \max\left(\left(1 + \beta + \frac{12\eta}{m}\right), 2\eta\right) OPT \\ &\quad + \frac{4\eta}{m(1+\beta)} ALG \end{aligned}$$

Let us assume that that β, η are picked so that $\frac{4\eta}{m(1+\beta)} < 1$; rearranging gives

$$\frac{ALG}{OPT} \leq \frac{\max\left(\left(1 + \beta + \frac{12\eta}{m}\right), 2\eta\right)}{1 - \frac{4\eta}{m(1+\beta)}}$$

Now we seek to minimize the competitive ratio by appropriately picking β, η, subject to the constraints. Notice that the competitive ratio is always increasing in η. Since we know that we must have $\eta \geq 1 + \frac{1}{\beta}$, this must be the optimal value of η. We can hence immediately rewrite our optimization problem purely in terms of β:

$$\min_{\beta > \frac{4}{m}} \frac{\max\left(\left(1 + \beta + \frac{12}{m}(1 + \frac{1}{\beta})\right), 2(1 + \frac{1}{\beta})\right)}{1 - \frac{4}{m\beta}}$$

Note that this proves that OBD is competitive for all $\beta > \frac{4}{m}$. Instead of trying to find the exact optimal solution, we instead select a simple choice of β which gives a small competitive ratio. Setting $\beta = 2 + \frac{10}{m}$ immediately gives an upper bound on the competitive ratio of $3 + O(1/m)$ as claimed. □

Proof of Corollary 2

Proof. Define

$$y_t = \sum_{i=1}^{t} B u_i, \qquad v_t = -\sum_{i=1}^{t} w_i.$$

Notice that

$$x_t = y_t - v_t, \qquad u_t = B^{-1}(y_t - y_{t-1}),$$

so the LQR problem can be rewritten as

$$\min_{y_1 \ldots y_T} \sum_{t=1}^{T} \frac{1}{2}(y_t - v_t)' Q_t (y_t - v_t) + \frac{1}{2}(y_t - y_{t-1})' Z (y_t - y_{t-1})$$

where

$$Z = (B^{-1})' R B^{-1}$$

Now define

$$z_t = R^{\frac{1}{2}} B^{-1} y_t, \qquad s_t = R^{\frac{1}{2}} B^{-1} v_t.$$

The optimization problem becomes

$$\min_{z_1 \dots z_T} \sum_{t=1}^{T} \frac{1}{2}(z_t - s_t)' P_t (z_t - s_t) + \frac{1}{2}\|z_t - z_{t-1}\|_2^2$$

where

$$P_t = (B R^{-\frac{1}{2}})' Q_t B R^{-\frac{1}{2}}$$

This is just a special case of the SOCO problem. Notice that the costs are strongly convex with parameter $\lambda_{min}(P_t)$, which is bounded below by

$$\frac{\lambda_{min} B^2 \lambda}{\lambda_{max}(R)}$$

which in light of Theorem 1 proves the claim.

□

References

[ABC+18] CJ Argue, Sébastien Bubeck, Michael B Cohen, Anupam Gupta, and Yin Tat Lee. A nearly-linear bound for chasing nested convex bodies. *arXiv preprint arXiv:1806.08865*, 2018.

[ABL+13] Lachlan Andrew, Siddharth Barman, Katrina Ligett, Minghong Lin, Adam Meyerson, Alan Roytman, and Adam Wierman. A tale of two metrics: Simultaneous bounds on competitiveness and regret. In *Conference on Learning Theory*, pages 741–763, 2013.

[ABN+16] Antonios Antoniadis, Neal Barcelo, Michael Nugent, Kirk Pruhs, Kevin Schewior, and Michele Scquizzato. Chasing convex bodies and functions. In *Latin American Symposium on Theoretical Informatics*, pages 68–81. Springer, 2016.

[AM10] Karl Johan Aström and Richard M Murray. *Feedback systems: an introduction for scientists and engineers*. Princeton university press, 2010.

[BB00] Avrim Blum and Carl Burch. On-line learning and the metrical task system problem. *Machine Learning*, 39(1):35–58, 2000.

[BBBT97] Yair Bartal, Avrim Blum, Carl Burch, and Andrew Tomkins. A polylog (n)-competitive algorithm for metrical task systems. In *Proceedings of the twenty-ninth annual ACM symposium on Theory of computing*, pages 711–719. ACM, 1997.

[BBE+18] Nikhil Bansa, Martin Böhm, Marek Eliáš, Grigorios Koumoutsos, and Seeun William Umboh. Nested convex bodies are chaseable. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1253–1260. SIAM, 2018.

[BCL+18] Sebastien Bubeck, Michael B Cohen, Yin Tat Lee, James R Lee, and Aleksander Madry. k-server via multiscale entropic regularization. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 3–16. ACM, 2018.

[BGK+15] Nikhil Bansal, Anupam Gupta, Ravishankar Krishnaswamy, Kirk Pruhs, Kevin Schewior, and Cliff Stein. A 2-competitive algorithm for online convex optimization with switching costs. In *LIPIcs-Leibniz International Proceedings in Informatics*, volume 40, 2015.

[BLS92] Allan Borodin, Nathan Linial, and Michael E Saks. An optimal on-line algorithm for metrical task system. *Journal of the ACM (JACM)*, 39(4):745–763, 1992.

[BV04] Stephen Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004.

[CAW+15] Niangjun Chen, Anish Agarwal, Adam Wierman, Siddharth Barman, and Lachlan LH Andrew. Online convex optimization using predictions. In *ACM SIGMETRICS Performance Evaluation Review*, volume 43, pages 191–204. ACM, 2015.

[CCL+16] Niangjun Chen, Joshua Comden, Zhenhua Liu, Anshul Gandhi, and Adam Wierman. Using predictions in online optimization: Looking forward with an eye on the past. *ACM SIGMETRICS Performance Evaluation Review*, 44(1):193–206, 2016.

[CGW18] Niangjun Chen, Gautam Goel, and Adam Wierman. Smoothed online convex optimization in high dimensions via online balanced descent. In *Proceedings of the 31st Conference On Learning Theory*, volume 75 of *Proceedings of Machine Learning Research*, pages 1574–1594. PMLR, 2018.

[FL93] Joel Friedman and Nathan Linial. On convex body chasing. *Discrete & Computational Geometry*, 9(3):293–321, 1993.

[GCW17] Gautam Goel, Niangjun Chen, and Adam Wierman. Thinking fast and slow: Optimization decomposition across timescales. In *Decision and Control (CDC), 2017 IEEE 56th Annual Conference on*, pages 1291–1298. IEEE, 2017.

[H+16] Elad Hazan et al. Introduction to online convex optimization. *Foundations and Trends® in Optimization*, 2(3-4):157–325, 2016.

[JdV12] V. Joseph and G. de Veciana. Jointly optimizing multi-user rate adaptation for video transport over wireless systems: Mean-fairness-variability tradeoffs. In *2012 Proceedings IEEE INFOCOM*, pages 567–575, March 2012.

[KG14] Seung-Jun Kim and Geogios B Giannakis. Real-time electricity pricing for demand response using online convex optimization. In *Innovative Smart Grid Technologies Conference (ISGT), 2014 IEEE PES*, pages 1–5. IEEE, 2014.

[KYTM15] Taehwan Kim, Yisong Yue, Sarah Taylor, and Iain Matthews. A decision tree framework for spatiotemporal sequence prediction. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 577–586. ACM, 2015.

[LLWA12] M. Lin, Z. Liu, A. Wierman, and L. L. H. Andrew. Online algorithms for geographical load balancing. In *2012 International Green Computing Conference (IGCC)*, pages 1–10, June 2012.

[LQL18] Yingying Li, Guannan Qu, and Na Li. Using predictions in online optimization with switching costs: A fast algorithm and a fundamental limit. In *2018 Annual American Control Conference (ACC)*, pages 3008–3013. IEEE, 2018.

[LWAT11] M. Lin, A. Wierman, L. L. H. Andrew, and E. Thereska. Dynamic right-sizing for power-proportional data centers. In *2011 Proceedings IEEE INFOCOM*, pages 1098–1106, April 2011.

[MMS90] Mark S Manasse, Lyle A McGeoch, and Daniel D Sleator. Competitive algorithms for server problems. *Journal of Algorithms*, 11(2):208–230, 1990.