
Beyond Online Balanced Descent: An Optimal Algorithm for Smoothed Online Optimization

Gautam Goel^{*1} Yiheng Lin^{*2,1} Haoyuan Sun^{*1} Adam Wierman¹

¹California Institute of Technology

²Institute for Interdisciplinary Information Sciences, Tsinghua University

Abstract

We study online convex optimization in a setting where the learner seeks to minimize the sum of a per-round hitting cost and a movement cost which is incurred when changing decisions between rounds. We prove a new lower bound on the competitive ratio of any online algorithm in the setting where the costs are m -strongly convex and the movement costs are the squared ℓ_2 norm. This lower bound shows that no algorithm can achieve a competitive ratio that is $o(m^{-1/2})$ as m tends to zero. No existing algorithms have competitive ratios matching this bound, and we show that the state-of-the-art algorithm, Online Balanced Descent (OBD), has a competitive ratio that is $\Omega(m^{-2/3})$. We additionally propose two new algorithms, Greedy OBD (G-OBD) and Regularized OBD (R-OBD) and prove that both algorithms have an $O(m^{-1/2})$ competitive ratio. The result for G-OBD holds when the hitting costs are quasiconvex and the movement costs are the squared ℓ_2 norm, while the result for R-OBD holds when the hitting costs are m -strongly convex and the movement costs are Bregman Divergences. Further, we show that R-OBD simultaneously achieves constant, dimension-free competitive ratio and sublinear regret when hitting costs are strongly convex.

1 Introduction

We consider the problem of Smoothed Online Convex Optimization (SOCO), a variant of online convex optimization (OCO) where the online learner pays a movement cost for changing actions between rounds. More precisely, we consider a game where an online learner plays a series of rounds against an adaptive adversary. In each round, the adversary picks a convex cost function $f_t : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$ and shows it to the learner. After observing the cost function, the learner chooses an action x_t and pays a *hitting cost* $f_t(x_t)$, as well as a *movement cost* $c(x_t, x_{t-1})$, which penalizes the online learner for switching points between rounds.

SOCO was originally proposed in the context of dynamic power management in data centers [29]. Since then it has seen a wealth of applications, from speech animation to management of electric vehicle charging [24, 26, 27], and more recently applications in control [21, 22] and power systems [5, 28]. SOCO has been widely studied in the machine learning community with the special cases of online logistic regression and smoothed online maximum likelihood estimation receiving recent attention [22].

Additionally, SOCO has connections to a number of other important problems in online algorithms and learning. Convex Body Chasing (CBC), introduced in [20], is a special case of SOCO [14]. The problem of designing competitive algorithms for Convex Body Chasing has attracted much recent

^{*}Gautam Goel, Yiheng Lin, and Haoyuan Sun contributed equally to this work.

attention. e.g. [2, 6, 14]. SOCO can also be viewed as a continuous version of the Metrical Task System (MTS) problem (see [9, 11, 12]). A special case of MTS is the celebrated k -server problem, first proposed in [31], which has received significant attention in recent years (see [13, 15]).

Given these connections, the design and analysis of algorithms for SOCO and related problems has received considerable attention in the last decade. SOCO was first studied in the scalar setting in [30], which used SOCO to model dynamic “right-sizing” in data centers and gave a 3-competitive algorithm. A 2-competitive algorithm was shown in [8], also in the scalar setting, which matches the lower bound for online algorithms in this setting [1]. Another rich line of work studies how to design competitive algorithms for SOCO when the online algorithm has access to predictions of future cost functions (see [16, 17, 28, 29]).

Despite a large and growing literature on SOCO and related problems, for nearly a decade the only known constant-competitive algorithms that did not use predictions of future costs were for one-dimensional action spaces. In fact, the connections between SOCO and Convex Body Chasing highlight that, in general, one cannot expect dimension-free constant competitive algorithms due to a $\Omega(\sqrt{d})$ lower bound (see [18, 20]). However, recently there has been considerable progress moving beyond the one-dimensional setting for large, important classes of hitting and movement costs.

A breakthrough came in 2017 when [18] proposed a new algorithm, Online Balanced Descent (OBD), and showed that it is constant competitive in all dimensions in the setting where the hitting costs are locally polyhedral and movement costs are the ℓ_2 norm. The following year, [22] showed that OBD is also constant competitive, specifically $3 + O(1/m)$ -competitive, in the setting where the hitting costs are m -strongly convex and the movement costs are the squared ℓ_2 norm. Note that this setting is of particular interest because of its importance for online regression and LQR control (see [22]).

While OBD has proven to be a promising new algorithm, at this point it is not known whether OBD is *optimal* for the competitive ratio, or if there is more room for improvement. This is because there are no non-trivial lower bounds known for important classes of hitting costs, the most prominent of which is the class of strongly convex functions.

Contributions of this paper. In this paper we prove the first non-trivial lower bounds on SOCO with strongly convex hitting costs, both for general algorithms and for OBD specifically. These lower bounds show that OBD is not optimal and there is an order-of-magnitude gap between its performance and the general lower bound. Motivated by this gap and the construction of the lower bounds we present two new algorithms, both variations of OBD, which have competitive ratios that match the lower bound. More specifically, we make four main contributions in this paper.

First, we prove a new lower bound on the performance achievable by any online algorithm in the setting where the hitting costs are m -strongly convex and the movement costs are the squared ℓ_2 norm. In particular, in Theorem 1, we show that as m tends to zero, any online algorithm must have competitive ratio at least $\Omega(m^{-1/2})$.

Second, we show that the state-of-the-art algorithm, OBD, cannot match this lower bound. More precisely, in Theorem 2 we show that, as m tends to zero, the competitive ratio of OBD is $\Omega(m^{-2/3})$, an order-of-magnitude higher than the lower bound of $\Omega(m^{-1/2})$. This immediately begs the question: can any online algorithm close the gap and match the lower bound?

Our third contribution answers this question in the affirmative. In Section 4, we propose two novel algorithms, Greedy Online Balanced Descent (G-OBD) and Regularized Online Balanced Descent (R-OBD), which are able to close the gap left open by OBD and match the $\Omega(m^{-1/2})$ lower bound. Both algorithms can be viewed as “aggressive” variants of OBD, in the sense that they chase the minimizers of the hitting costs more aggressively than OBD. In Theorem 3 we show that G-OBD matches the lower bound up to constant factors for quasiconvex hitting costs (a more general class than m -strongly convex). In Theorem 4 we show that R-OBD has a competitive ratio that *precisely matches the lower bound, including the constant factors*, and hence can be viewed as an optimal algorithm for SOCO in the setting where the costs are m -strongly convex and the movement cost is the squared ℓ_2 norm. Further, our results for R-OBD hold not only for squared ℓ_2 movement costs; they also hold for movement costs that are Bregman Divergences, which commonly appear throughout information geometry, probability, and optimization.

Finally, in our last section we move beyond competitive ratio and additionally consider regret. We prove in Theorem 6 that R-OBDD can simultaneously achieve bounded, dimension-free competitive ratio and sublinear regret in the case of m -strongly convex hitting costs and squared ℓ_2 movement costs. This result helps close a crucial gap in the literature. Previous work has shown that it not possible for any algorithm to simultaneously achieve both a constant competitive ratio and sublinear regret in general SOCO problems [19]. However, this was shown through the use of linear hitting and movement costs. Thus, the question of whether it is possible to simultaneously achieve a dimension-free, constant competitive ratio and sublinear regret when hitting costs are strongly convex has remained open. The closest previous result is from [18], which showed that OBD can achieve either constant competitive ratio or sublinear regret with locally polyhedral cost functions depending on the “balance condition” used; however both cannot be achieved simultaneously. Our result (Theorem 6), shows that R-OBDD can simultaneously provide a constant competitive ratio and sublinear regret for strongly convex cost functions when the movement costs are the squared ℓ_2 norm.

2 Model & Preliminaries

An instance of Smoothed Online Convex Optimization (SOCO) consists of a convex action set $\mathcal{X} \subset \mathbb{R}^d$, an initial point $x_0 \in \mathcal{X}$, a sequence of non-negative convex cost functions $f_1 \dots f_T : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$, and a movement cost $c : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$. In every round, the environment picks a cost function f_t (potentially adversarially) for an online learner. After observing the cost function, the learner chooses an action $x_t \in \mathbb{R}^d$ and pays a cost that is the sum of the *hitting cost*, $f_t(x_t)$, and the *movement cost*, a.k.a., switching cost, $c(x_t, x_{t-1})$. The goal of the online learner is to minimize its total cost over T rounds: $\text{cost}(ALG) = \sum_{t=1}^T f_t(x_t) + c(x_t, x_{t-1})$.

We emphasize that it is the movement costs that make this problem interesting and challenging; if there were no movement costs, $c(x_t, x_{t-1}) = 0$, the problem would be trivial, since the learner could always pay the optimal cost simply by picking the action that minimizes the hitting cost in each round, i.e., by setting $x_t = \arg \min_x f_t(x)$. The movement cost couples the cost the learner pays across rounds, which means that the optimal action of the learner depends on unknown future costs.

There is a long literature on SOCO, both focusing on algorithmic questions, e.g., [8, 18, 22, 30], and applications, e.g., [24, 26, 27, 29]. The variety of applications studied means that a variety of assumptions about the movement costs have been considered. Motivated by applications to data center capacity management, movement costs have often been taken as the ℓ_1 norm, i.e., $c(x_1, x_2) = \|x_1 - x_2\|_1$, e.g. [8, 30]. However, recently, more general norms have been considered and the setting of squared ℓ_2 movement costs has gained attention due to its use in online regression problems and connections to LQR control, among other applications (see [3, 21, 22]).

In this paper, we focus on the setting of the squared ℓ_2 norm, i.e. $c(x_2, x_1) = \frac{1}{2} \|x_2 - x_1\|_2^2$; however, we also consider a generalization of the ℓ_2 norm in Section 4.2 where c is the Bregman divergence. Specifically, we consider $c(x_t, x_{t-1}) = D_h(x_t || x_{t-1}) = h(x_t) - h(x_{t-1}) - \langle \nabla h(x_{t-1}), x_t - x_{t-1} \rangle$, where both the potential h and its Fenchel Conjugate h^* are differentiable. Further, we assume that h is α -strongly convex and β -strongly smooth with respect to an underlying norm $\|\cdot\|$. Definitions of each of these properties can be found in the appendix.

Note that the squared ℓ_2 norm is itself a Bregman divergence, with $\alpha = \beta = 1$ and $\|\cdot\| = \|\cdot\|_2$, $D_h(x_t || x_{t-1}) = \frac{1}{2} \|x_t - x_{t-1}\|_2^2$. However, more generally, when $h(y) = \sum_i y_i \ln y_i$ with domain $\Delta_n = \{y \in [0, 1]^n \mid \sum_i y_i = 1\}$, $D_h(x_t || x_{t-1})$ is the Kullback-Liebler divergence (see [7]). Further, h is $\frac{1}{2 \ln 2}$ -strongly convex and $\frac{1}{\delta \ln 2}$ -strongly smooth in the domain $\mathcal{X} = P_\delta = \{y \in [0, 1]^n \mid \sum_i y_i = 1, y_i \geq \delta\}$ (see [18]). This extension is important given the role Bregman divergence plays across optimization and information theory, e.g., see [4, 32].

Like for movement costs, a variety of assumptions have been made about hitting costs. In particular, because of the emergence of pessimistic lower bounds when general convex hitting costs are considered, papers typically have considered restricted classes of functions, e.g., locally polyhedral [18] and strongly convex [22]. In this paper, we focus on hitting costs that are m -strongly convex; however our results in Section 4.1 generalize to the case of quasiconvex functions.

Competitive Ratio and Regret. The primary goal of the SOCO literature is to design online algorithms that (nearly) match the performance of the offline optimal algorithm. The performance

Algorithm 1 Online Balanced Descent (OBD)

- 1: **procedure** OBD(f_t, x_{t-1}) ▷ Procedure to select x_t
 - 2: $v_t \leftarrow \arg \min_x f_t(x)$
 - 3: Let $x(l) = \prod_{K_t^l}(x_{t-1})$. Initialize $l = f_t(v_t)$. Here $K_t^l = \{x | f_t(x) \leq l\}$.
 - 4: Increase l . Stop either when $x(l) = v_t$ or $c(x(l), x_{t-1}) = \gamma l$.
 - 5: $x_t \leftarrow x(l)$.
 - 6: **return** x_t
-

metric used to evaluate an algorithm is typically the *competitive ratio* because the goal is to learn in an environment that is changing dynamically and is potentially adversarial. The competitive ratio is the worst-case ratio of total cost incurred by the online learner and the offline optimal costs. The cost of the offline optimal is defined as the minimal cost an algorithm if it has full knowledge of the sequence of costs $\{f_t\}$, i.e. $cost(OPT) = \min_{x_1, \dots, x_T} \sum_{t=1}^T f_t(x_t) + c(x_t, x_{t-1})$. Using this, the *competitive ratio* is defined as $\sup_{f_1, \dots, f_T} cost(ALG) / cost(OPT)$.

Note that another important performance measure of interest is the *regret*. In this paper, we study a generalization of the classical regret called the L -constrained regret, which is defined as follows. The L -constrained dynamic regret of an online algorithm ALG is $\rho_L(T)$ if for all sequences of cost functions f_1, \dots, f_T , we have $cost(ALG) - cost(OPT(L)) \leq \rho_L(T)$ where $OPT(L)$ is the cost of an L -constrained offline optimal solution, i.e., one with movement cost upper bounded by L : $OPT(L) = \min_{x \in \mathcal{X}^T} \sum_{t=1}^T f_t(x_t) + c(x_t, x_{t-1})$ subject to $\sum_{t=1}^T c(x_t, x_{t-1}) \leq L$.

As the definitions above highlight, the regret and competitive ratio both compare with the cost of an offline optimal solution, however regret constrains the movement allowed by the offline optimal. The classical notion of regret focuses on the static optimal ($L = 0$), but relaxing that to allow limited movement bridges regret and the competitive ratio since, as L grows, the L -constrained offline optimal approaches the offline (dynamic) optimal. Intuitively, one can think of regret as being suited for evaluating learning algorithms in (nearly) static settings while the competitive ratio as being suited for evaluating learning algorithms in dynamic settings.

Online Balanced Descent. The state-of-the-art algorithm for SOCO is Online Balanced Descent (OBD). OBD is formally defined in Algorithm 1. Intuitively, it works as follows. In every round, OBD projects the previously chosen point x_{t-1} onto a carefully chosen level set of the current cost function f_t . The level set is chosen so that the hitting costs and movement costs are “balanced”: in every round, the movement cost is at most a constant γ times the hitting cost. The balance helps ensure that the online learner is matching the offline costs. Since neither cost is too high, OBD ensures that both are comparable to the offline optimal. The parameter γ can be tuned to give the optimal competitive ratio and the appropriate level set can be efficiently selected via binary search.

Implicitly, OBD can be viewed as a proximal algorithm with a dynamic step size [33], in the sense that, like proximal algorithms, OBD iteratively projects the previously chosen point onto a level set of the cost function. Unlike traditional proximal algorithms, OBD considers several different level sets, and carefully selects the level set in every round so as to balance the hitting and movement costs. We exploit this connection heavily when designing Regularized OBD (R-OBD), which is a proximal algorithm with a special regularization term added to the objective to help steer the online learner towards the hitting cost minimizer in each round.

OBD was proposed in [18], where the authors show that it has a constant, dimension-free competitive ratio in the setting where the movement costs are the ℓ_2 norm and the hitting costs are locally polyhedral, i.e. grow at least linearly away from the minimizer. This was the first time an algorithm had been shown to be constant competitive beyond one-dimensional action spaces. In the same paper, a variation of OBD that uses a different balance condition was proven to have $O(\sqrt{TL})$ L -constrained regret for locally polyhedral hitting costs. OBD has since been shown to also have a constant, dimension-free competitive ratio when movement costs are the squared ℓ_2 norm and hitting costs are strongly convex, which is the setting we consider in this paper. However, up until this paper, lower bounds for the strongly convex setting did not exist and it was not known whether the performance of OBD in this setting is optimal or if OBD can simultaneously achieve sublinear regret and a constant, dimension-free competitive ratio.

Figure 1: Counterexample used to prove Theorem 2. In the figure, $\{x_t\}$ are the choices of OBD and $\{x_t^*\}$ are the choices of the offline optimal.

3 Lower Bounds

Our first set of results focuses on lower bounding the competitive ratio achievable by online algorithms for SOCO. While [18] proves a general lower bound for SOCO showing that the competitive ratio of any online algorithm is $\Omega(\sqrt{d})$, where d is the dimension of the action space, there are large classes of important problems where better performance is possible. In particular, when the hitting costs are m -strongly convex, [22] has shown that OBD provides a dimension-free competitive ratio of $3 + O(1/m)$. However, no non-trivial lower bounds are known for the strongly convex setting.

Our first result in this section shows a general lower bound on the competitive ratio of SOCO algorithms when the hitting costs are strongly convex and the movement costs are quadratic. Importantly, there is a gap between this bound and the competitive ratio for OBD proven in [22]. Our second result further explores this gap. We show a lower bound on the competitive ratio of OBD which highlights that OBD cannot achieve a competitive ratio that matches the general lower bound. This gap, and the construction used to show it, motivate us to propose new variations of OBD in the next section. We then prove that these new algorithms have competitive ratios that match the lower bound.

We begin by stating the first lower bound for strongly convex hitting costs in SOCO.

Theorem 1. *Consider hitting cost functions that are m -strongly convex with respect to ℓ_2 norm and movement costs given by $\frac{1}{2} \|x_t - x_{t-1}\|_2^2$. Any online algorithm must have a competitive ratio at least $\frac{1}{2} \left(1 + \sqrt{1 + \frac{4}{m}}\right)$.*

Theorem 1 is proven in the appendix using an argument that leverages the fact that, when the movement cost is quadratic, reaching a target point via one large step is more costly than reaching it by taking many small steps. More concretely, to prove the lower bound we consider a scenario on the real line where the online algorithm encounters a sequence of cost functions whose minimizers are at zero followed by a very steep cost function whose minimizer is at $x = 1$. Without knowledge of the future, the algorithm has no incentive to move away from zero until the last step, when it is forced to incur a large cost; however, the offline adversary, with full knowledge of the cost sequence, can divide the journey into multiple small steps.

Importantly, the lower bound in Theorem 1 highlights the dependence of the competitive ratio on m , the convexity parameter. It shows that the case where online algorithms do the worst is when m is small, and that algorithms that match the lower bound up to a constant are those for which the competitive ratio is $O(m^{-1/2})$ as $m \rightarrow 0^+$. Note that our results in Section 4 show that there exists online algorithms that precisely achieve the competitive ratio in Theorem 1. However, in contrast, the following shows that OBD cannot match the lower bound in Theorem 1.

Theorem 2. *Consider hitting cost functions that are m -strongly convex with respect to ℓ_2 norm and a movement costs given by $\frac{1}{2} \|x_t - x_{t-1}\|_2^2$. The competitive ratio of OBD is $\Omega(m^{-\frac{2}{3}})$ as $m \rightarrow 0^+$, for any fixed γ .*

As we have discussed, OBD is the state-of-the-art algorithm for SOCO, and has been shown to provide a competitive ratio of $3 + O(1/m)$ [22]. However, Theorem 2 highlights a gap between OBD and the general lower bound. If the lower bound is achievable (which we prove it is in the next section), this implies that OBD is a sub-optimal algorithm.

The proof of Theorem 2 gives important intuition about what goes wrong with OBD and how the algorithm can be improved. Specifically, our proof of Theorem 2 considers a scenario where the cost functions have minimizers very near each other, but OBD takes a series of steps without approaching the minimizing points. The optimal is able to pay little cost and stay near the minimizers, but OBD never moves enough to be close to the minimizers. Figure 1 illustrates the construction, showing OBD moving along the circumference of a circle, while the offline optimal stays near the origin.

4 Algorithms

The lower bounds in Theorem 1 and Theorem 2 suggest a gap between the competitive ratio of OBD and what is achievable via an online algorithm. Further, the construction used in the proof of Theorem 2 highlights the core issue that leads to inefficiency in OBD. In the construction, OBD takes a large step from x_{t-1} to x_t , but the offline optimal, x_t^* , only decreases by a very small amount. This means that OBD is continually chasing the offline optimal but never closing the gap. In this section, we take inspiration from this example and develop two new algorithms that build on OBD but ensure that the gap to the offline optimal x_t^* shrinks.

How to ensure that the gap to the offline optimal shrinks is not obvious since, without the knowledge about the future, it is impossible to determine how x_t^* will evolve. A natural idea is to determine an online estimate of x_t^* and then move towards that estimate. Motivated by the construction in the proof of Theorem 2, we use the minimizer of the hitting cost at round t , v_t , as a rough estimate of the offline optimal and ensure that we close the gap to v_t in each round.

There are a number of ways of implementing the goal of ensuring that OBD more aggressively moves toward the minimizer of the hitting cost each round. In this section, we consider two concrete approaches, each of which (nearly) matches the lower bound in Theorem 1.

The first approach, which we term Greedy OBD (Algorithm 2) is a two-stage algorithm, where the first stage applies OBD and then a second stage explicitly takes a step directly towards the minimizer (of carefully chosen size). We introduce the algorithm and analyze its performance in Section 4.1. Greedy OBD is order-optimal, i.e. matches the lower bound up to constant factors, in the setting of squared ℓ_2 norm movement costs and quasiconvex hitting costs.

The second approach for ensuring that OBD moves aggressively toward the minimizer uses a different view of OBD. In particular, Greedy OBD uses a *geometric view* of OBD, which is the way OBD has been presented previously in the literature. Our second view uses a “local view” of OBD that parallels the *local view* of gradient descent and mirror descent, e.g., see [7, 23]. In particular, the choice of an action in OBD can be viewed as the solution to a per-round *local* optimization. Given this view, we ensure that OBD more aggressively tracks the minimizer by adding a regularization term to this local optimization which penalizes points which are far from the minimizer. We term this approach Regularized OBD (Algorithm 3), and study it in Section 4.2. Note that Regularized OBD has a competitive ratio that precisely matches the lower bound, including the constant factors, when movement costs are Bregman divergences and hitting costs are m -strongly convex. Thus, it applies for more general movement costs than Greedy OBD but less general hitting costs.

4.1 Greedy OBD

The formal description of Greedy Online Balanced Descent (G-OBD) is given in Algorithm 2. G-OBD has two steps each round. First, the algorithm takes a standard OBD step from the previous point x_{t-1} to a new point x'_t , which is the projection of x_{t-1} onto a level set of the current hitting cost f_t , where the level set is chosen to balance hitting and movement costs. G-OBD then takes an additional step directly towards the minimizer of the hitting cost, v_t , with the size of the step chosen based on the convexity parameter m . G-OBD can be implemented efficiently using the same approach as described for OBD [18]. G-OBD has two parameters γ and μ . The first, γ , is the balance parameter in OBD and the second, μ , is a parameter controlling the size of the step towards the minimizer v_t . Note that the two-step approach of G-OBD is reminiscent of the two-stage algorithm used in [10]; however the resulting algorithms are quite distinct.

While the addition of a second step in G-OBD may seem like a small change, it improves performance by an order-of-magnitude. We prove that G-OBD asymptotically matches the lower bound proven in Theorem 2 not just for m -strongly convex hitting costs, but more broadly to quasiconvex costs.

Theorem 3. *Consider quasiconvex hitting costs such that $f_t(x) \geq f_t(v_t) + \frac{m}{2} \|x - v_t\|_2^2$ and movement costs $c(x_t, x_{t-1}) = \frac{1}{2} \|x_t - x_{t-1}\|_2^2$. G-OBD with $\gamma = 1, \mu = 1$ is an $O(m^{-1/2})$ -competitive algorithm as $m \rightarrow 0+$.*

Algorithm 2 Greedy Online Balanced Descent (G-OBDD)

1: **procedure** G-OBDD(f_t, x_{t-1}) ▷ Procedure to select x_t
2: Set $v_t = \arg \min_x f_t(x)$
3: $x'_t \leftarrow \text{OBD}(f_t, x_{t-1}, \gamma)$
4: **if** $\mu\sqrt{m} \geq 1$ **then**
5: $x_t \leftarrow v_t$
6: **else**
7: $x_t \leftarrow \mu\sqrt{m}v_t + (1 - \mu\sqrt{m})x'_t$
8: **return** x_t

Algorithm 3 Regularized OBD (R-OBDD)

1: **procedure** R-OBDD(f_t, x_{t-1}) ▷ Procedure to select x_t
2: $v_t \leftarrow \arg \min_x f_t(x)$
3: $x_t \leftarrow \arg \min_x f_t(x) + \lambda_1 c(x, x_{t-1}) + \lambda_2 c(x, v_t)$
4: **return** x_t

4.2 Regularized OBD

The G-OBDD framework is based on the geometric view of OBD used previously in literature. There are, however, two limitations to this approach. First, the competitive ratio obtained, while having optimal asymptotic dependence on m , does not match the constants in the lower bound of Theorem 1. Second, G-OBDD requires repeated projections, which makes efficient implementation of challenging when the functions f_t have complex geometry.

Here, we present a variation of OBD based on a *local view* that overcomes these limitations. Regularized OBD (R-OBDD) is computationally simpler and provides a competitive ratio that matches the constant factors in the lower bound in Theorem 1. However, unlike G-OBDD, our analysis of R-OBDD does not apply to quasiconvex hitting costs. R-OBDD is described formally in Algorithm 3. In each round, R-OBDD picks a point that minimizes a weighted sum of the hitting and movement costs, as well as a regularization term which encourages the algorithm to pick points close to the minimizer of the current hitting cost function, $v_t = \arg \min_x f_t(x)$. Thus, R-OBDD can be implemented efficiently using two invocations of a convex solver. Note that R-OBDD has two parameters λ_1 and λ_2 which adjust the weights of the movement cost and regularizer respectively.

While it may not be immediately clear how R-OBDD connects to OBD, it is straightforward to illustrate the connection in the squared ℓ_2 setting. In this case, computing $x_t = \arg \min_x f_t(x) + \frac{\lambda_1}{2} \|x - x_{t-1}\|_2^2$ is equivalent to doing a projection onto a level set of f_t , since the selection of the minimizer can be restated as the solution to $\nabla f_t(x_t) + \lambda_1(x_t - x_{t-1}) = 0$. Thus, without the regularizer, the optimization in R-OBDD gives a local view of OBD and then the regularizer provides more aggressive movement toward the minimizer of the hitting cost.

Not only does the local view lead to a computationally simpler algorithm, but we prove that R-OBDD matches the constant factors in Theorem 1 precisely, not just asymptotically. Further, it does this not just in the setting where movement costs are the squared ℓ_2 norm, but also in the case where movement costs are Bregman divergences.

Theorem 4. *Consider hitting costs that are m -strongly convex with respect to a norm $\|\cdot\|$ and movement costs defined as $c(x_t, x_{t-1}) = D_h(x_t \| x_{t-1})$, where h is α -strongly convex and β -strongly smooth with respect to the same norm. Additionally, assume $\{f_t\}$, h and its Fenchel Conjugate h^* are differentiable. Then, R-OBDD with parameters $1 \geq \lambda_1 > 0$ and $\lambda_2 \geq 0$ has a competitive ratio of $\max\left(\frac{m+\lambda_2\beta}{\lambda_1} \cdot \frac{1}{m}, 1 + \frac{\beta^2}{\alpha} \cdot \frac{\lambda_1}{\lambda_2\beta+m}\right)$. If λ_1 and λ_2 satisfy $m + \lambda_2\beta = \frac{\lambda_1 m}{2} \left(1 + \sqrt{1 + \frac{4\beta^2}{\alpha m}}\right)$ then the competitive ratio is $\frac{1}{2} \left(1 + \sqrt{1 + \frac{4\beta^2}{m\alpha}}\right)$.*

Theorem 4 focuses on movement costs that are Bregman divergences, which generalizes the case of squared ℓ_2 movement costs. To recover the squared ℓ_2 case, we use $\|\cdot\| = \|\cdot\|_2$ and $\alpha = \beta = 1$, which results in a competitive ratio of $\frac{1}{2}(1 + \sqrt{1 + 4/m})$. This competitive ratio matches exactly

with the lower bound claimed in Theorem 1. Further, in this case the assumption in Theorem 4 that the hitting cost functions are differentiable is not required (see Theorem 7 in the appendix).

It is also interesting to investigate the settings of λ_1 and λ_2 that yield the optimal competitive ratio. Surprisingly, setting $\lambda_2 = 0$ achieves the optimal competitive ratio as long as $\lambda_1 = 2 \left(1 + \sqrt{1 + \frac{4\beta^2}{\alpha m}}\right)^{-1}$. By restating the update rule in R-OBD as $\nabla f_t(x_t) = \lambda_1(\nabla h(x_{t-1}) - \nabla h(x_t))$, we see that R-OBD with $\lambda_2 = 0$ can be interpreted as “one step lookahead mirror descent”. Further R-OBD with $\lambda_2 = 0$ can be implemented even when we do not know the location of the minimizer v_t . For example, when $h(x) = \frac{1}{2} \|x\|_2^2$, we can run gradient descent starting at x_{t-1} to minimize the strongly convex function $f_t(x) + \frac{\lambda_1}{2} \|x - x_{t-1}\|_2^2$. Only local gradients will be queried in this process. However, the following lower bound highlights that this simple form comes at some cost in terms of generality when compared with our results for G-OBD.

Theorem 5. *Consider quasiconvex hitting costs such that $f_t(x) - f_t(v_t) \geq \frac{m}{2} \|x - v_t\|_2^2$ and movement costs given by $c(x_t, x_{t-1}) = \frac{1}{2} \|x_t - x_{t-1}\|_2^2$. Regularized OBD has a competitive ratio of $\Omega(1/m)$ when $\lambda_2 = 0$.*

5 Balancing Regret and Competitive Ratio

In the previous sections we have focused on the competitive ratio; however another important performance measure is *regret*. In this section, we consider the L -constrained dynamic regret. The motivation for our study is [19], which provides an impossibility result showing that no algorithm can simultaneously maintain a constant competitive ratio and a sub-linear regret in the general setting of SOCO. However, [19] utilizes linear hitting costs in its construction and thus it is an open question as to whether this impossibility result holds for strongly convex hitting costs. In this section, we show that the impossibility result does not hold for strongly convex hitting costs. To show this, we first characterize the parameters for which R-OBD gives sublinear regret.

Theorem 6. *Consider hitting costs that are m -strongly convex with respect to a norm $\|\cdot\|$ and movement costs defined as $c(x_t, x_{t-1}) = D_h(x_t | x_{t-1})$, where h is α -strongly convex and β -strongly smooth with respect to the same norm. Additionally, assume $\{f_t\}$, h and its Fenchel Conjugate h^* are differentiable. Further, suppose that $\|\nabla h(x)\|_*$ is bounded above by $G < \infty$, the diameter of the feasible set \mathcal{X} is bounded above by D , and $\nabla h(0) = 0$. Then, for λ_1, λ_2 such that $\lambda_1 \geq 1 - \frac{m}{4\beta}$ and $\lambda_2 = \eta(T, L, D, G)$, where $\eta(T, L, D, G)$ is such that $\lim_{T \rightarrow \infty} \eta(T, L, D, G) \cdot \frac{D^2}{G} \sqrt{\frac{L}{T}} < \infty$, the L -constrained regret of R-OBD is $O(G\sqrt{TL})$.*

Theorem 6 highlights that $O(G\sqrt{TL})$ regret can be achieved when $\lambda_1 \geq 1 - \frac{m}{4\beta}$ and $\lambda_2 \leq \frac{KG}{D^2} \cdot \sqrt{\frac{L}{T}}$ for some constant K . This suggests that the tendency to aggressively move towards the minimizer should shrink over time in order to achieve a small regret. It is not possible to use Theorem 6 to simultaneously achieve the optimal competitive ratio and $O(G\sqrt{TL})$ regret for all strongly convex hitting costs ($m > 0$). However, the corollary below shows that it is possible to simultaneously achieve a dimension-free, constant competitive ratio and an $O(G\sqrt{TL})$ regret for all $m > 0$. An interesting open question that remains is whether it is possible to develop an algorithm that has sublinear regret and matches the optimal order for competitive ratio.

Corollary 1. *Consider the same conditions as in Theorem 6 and fix $m > 0$. R-OBD with pa-*

rameters $\lambda_1 = \max \left(2 \left(1 + \sqrt{1 + \frac{4\beta^2}{\alpha m}} \right)^{-1}, 1 - \frac{m}{4\beta} \right)$, $\lambda_2 = 0$ has an $O(G\sqrt{TL})$ regret and is $\max \left(\frac{1}{2} \left(1 + \sqrt{1 + \frac{4\beta^2}{\alpha m}} \right), 1 - \frac{\beta}{4\alpha} + \frac{\beta^2}{\alpha m} \right)$ -competitive.

References

- [1] A. Antoniadis and K. Schewior. A tight lower bound for online convex optimization with switching costs. In *Proceedings of the International Workshop on Approximation and Online Algorithms*, pages 164–175. Springer, 2017.
- [2] C. Argue, S. Bubeck, M. B. Cohen, A. Gupta, and Y. T. Lee. A nearly-linear bound for chasing nested convex bodies. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 117–122, 2019.
- [3] K. J. Aström and R. M. Murray. *Feedback systems: an introduction for scientists and engineers*. Princeton university press, 2010.
- [4] N. Azizan and B. Hassibi. Stochastic gradient/mirror descent: Minimax optimality and implicit regularization. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2019.
- [5] M. Badiei, N. Li, and A. Wierman. Online convex optimization with ramp constraints. In *IEEE Conference on Decision and Control (CDC)*, pages 6730–6736, 2015.
- [6] N. Bansal, M. Böhm, M. Eliáš, G. Koumoutsos, and S. W. Umboh. Nested convex bodies are chaseable. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1253–1260, 2018.
- [7] N. Bansal and A. Gupta. Potential-function proofs for first-order methods. *arXiv preprint arXiv:1712.04581*, 2017.
- [8] N. Bansal, A. Gupta, R. Krishnaswamy, K. Pruhs, K. Schewior, and C. Stein. A 2-competitive algorithm for online convex optimization with switching costs. In *Proceedings of the Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2015.
- [9] Y. Bartal, A. Blum, C. Burch, and A. Tomkins. A polylog(n)-competitive algorithm for metrical task systems. In *Proceedings of the ACM Symposium on Theory of Computing (STOC)*, pages 711–719, 1997.
- [10] M. Bienkowski, J. Byrka, M. Chrobak, C. Coester, L. Jez, and E. Koutsoupias. Better bounds for online line chasing. *arXiv preprint arXiv:1811.09233*, 2018.
- [11] A. Blum and C. Burch. On-line learning and the metrical task system problem. *Machine Learning*, 39(1):35–58, 2000.
- [12] A. Borodin, N. Linial, and M. E. Saks. An optimal on-line algorithm for metrical task system. *Journal of the ACM*, 39(4):745–763, 1992.
- [13] S. Bubeck, M. B. Cohen, Y. T. Lee, J. R. Lee, and A. Mądry. k-server via multiscale entropic regularization. In *Proceedings of the ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 3–16, 2018.
- [14] S. Bubeck, Y. T. Lee, Y. Li, and M. Sellke. Competitively chasing convex bodies. In *Proceedings of the ACM SIGACT Symposium on Theory of Computing (STOC)*, 2019.
- [15] N. Buchbinder, A. Gupta, M. Molinaro, and J. Naor. k-servers with a smile: online algorithms via projections. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 98–116, 2019.
- [16] N. Chen, A. Agarwal, A. Wierman, S. Barman, and L. L. Andrew. Online convex optimization using predictions. *ACM SIGMETRICS Performance Evaluation Review*, 43(1):191–204, 2015.
- [17] N. Chen, J. Comden, Z. Liu, A. Gandhi, and A. Wierman. Using predictions in online optimization: Looking forward with an eye on the past. *ACM SIGMETRICS Performance Evaluation Review*, 44(1):193–206, 2016.
- [18] N. Chen, G. Goel, and A. Wierman. Smoothed online convex optimization in high dimensions via online balanced descent. In *Proceedings of Conference On Learning Theory (COLT)*, pages 1574–1594, 2018.
- [19] A. Daniely and Y. Mansour. Competitive ratio vs regret minimization: achieving the best of both worlds. In *Proceedings of Algorithmic Learning Theory*, pages 333–368, 2019.
- [20] J. Friedman and N. Linial. On convex body chasing. *Discrete & Computational Geometry*, 9(3):293–321, 1993.

- [21] G. Goel, N. Chen, and A. Wierman. Thinking fast and slow: Optimization decomposition across timescales. In *Proceedings of the IEEE Conference on Decision and Control (CDC)*, pages 1291–1298, 2017.
- [22] G. Goel and A. Wierman. An online algorithm for smoothed regression and LQR control. In *Proceedings of the Machine Learning Research*, volume 89, pages 2504–2513, 2019.
- [23] E. Hazan et al. Introduction to online convex optimization. *Foundations and Trends in Optimization*, 2(3-4):157–325, 2016.
- [24] V. Joseph and G. de Veciana. Jointly optimizing multi-user rate adaptation for video transport over wireless systems: Mean-fairness-variability tradeoffs. In *Proceedings of the IEEE INFOCOM*, pages 567–575, 2012.
- [25] S. Kakade, S. Shalev-Shwartz, and A. Tewari. On the duality of strong convexity and strong smoothness: Learning applications and matrix regularization. *Unpublished Manuscript*, <http://ttic.uchicago.edu/%7eshai/papers/KakadeShalevTewari09.pdf>, 2009.
- [26] S. Kim and G. B. Giannakis. An online convex optimization approach to real-time energy pricing for demand response. *IEEE Transactions on Smart Grid*, 8(6):2784–2793, 2017.
- [27] T. Kim, Y. Yue, S. Taylor, and I. Matthews. A decision tree framework for spatiotemporal sequence prediction. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 577–586, 2015.
- [28] Y. Li, G. Qu, and N. Li. Using predictions in online optimization with switching costs: A fast algorithm and a fundamental limit. In *Proceedings of the American Control Conference (ACC)*, pages 3008–3013. IEEE, 2018.
- [29] M. Lin, Z. Liu, A. Wierman, and L. L. Andrew. Online algorithms for geographical load balancing. In *Proceedings of the International Green Computing Conference (IGCC)*, pages 1–10, 2012.
- [30] M. Lin, A. Wierman, L. L. Andrew, and E. Thereska. Dynamic right-sizing for power-proportional data centers. *IEEE/ACM Transactions on Networking (TON)*, 21(5):1378–1391, 2013.
- [31] M. S. Manasse, L. A. McGeoch, and D. D. Sleator. Competitive algorithms for server problems. *Journal of Algorithms*, 11(2):208–230, 1990.
- [32] N. Murata, T. Takenouchi, T. Kanamori, and S. Eguchi. Information geometry of u-boost and Bregman divergence. *Neural Computation*, 16(7):1437–1481, 2004.
- [33] N. Parikh and S. Boyd. Proximal algorithms. *Foundations and Trends in Optimization*, 1(3):127–239, 2014.
- [34] S. Shalev-Shwartz and Y. Singer. On the equivalence of weak learnability and linear separability: New relaxations and efficient boosting algorithms. *Machine learning*, 80(2-3):141–163, 2010.

Appendices

The appendices that follow provide the proofs of the results in the body of the paper. Throughout the proofs in the appendix we use the following notation to denote the hitting and movement costs of the online learner: $H_t := f_t(x_t)$ and $M_t := c(x_t, x_{t-1})$, where x_t is the point chosen by the online algorithm at time t . Similarly, we denote the hitting and movement costs of the offline optimal (adversary) as $H_t^* := f_t(x_t^*)$ and $M_t^* := c(x_t^*, x_{t-1}^*)$, where x_t^* is the point chosen by the offline optimal at time t .

Before moving to the proofs, we summarize a few standard definitions that are used throughout the paper.

Definition 1. A function $f : \mathcal{X} \rightarrow \mathbb{R}$ is α -strongly convex with respect to a norm $\|\cdot\|$ if for all x, y in the relative interior of the domain of f and $\lambda \in (0, 1)$, we have

$$f(\lambda x + (1 - \lambda)y) \leq \lambda f(x) + (1 - \lambda)f(y) - \frac{\alpha}{2}\lambda(1 - \lambda)\|x - y\|^2.$$

Definition 2. A function $f : \mathcal{X} \rightarrow \mathbb{R}$ is β -strongly smooth with respect to a norm $\|\cdot\|$ if f is everywhere differentiable and if for all x, y we have

$$f(y) \leq f(x) + \langle \nabla f(x), y - x \rangle + \frac{\beta}{2}\|y - x\|^2.$$

Definition 3. A function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is quasiconvex if its domain \mathcal{X} and all its sublevel sets

$$S_\alpha = \{x \in \mathcal{X} \mid f(x) \leq \alpha\},$$

for $\alpha \in \mathbb{R}$, is convex.

Definition 4. For a norm $\|\cdot\|$ in \mathcal{X} , its dual norm (on \mathcal{X}) $\|\cdot\|_*$ is defined to be

$$\|y\|_* = \sup\{\langle x, y \rangle \mid \|x\| \leq 1\}.$$

Definition 5. For a convex function $f : \mathcal{X} \rightarrow \mathbb{R}$, its Fenchel Conjugate f^* is defined to be

$$f^*(y) = \sup\{\langle x, y \rangle - f(x) \mid x \in \mathcal{X}\}.$$

Next, we introduce a few technical lemmas that are important throughout our analysis.

The first technical lemma is a characterization of strongly convex functions.

Lemma 1. Suppose f is α -strongly convex for some $\alpha > 0$ with respect to some norm $\|\cdot\|$ and both f and f^* are differentiable, then the first condition implies the second condition and the third condition:

1. $\forall x, y, f(y) \leq f(x) + \langle \nabla f(x), y - x \rangle + \frac{\beta}{2}\|x - y\|^2$;
2. $\forall x, y, f(y) \geq f(x) + \langle \nabla f(x), y - x \rangle + \frac{1}{2\beta}\|\nabla f(x) - \nabla f(y)\|_*^2$;
3. $\forall x, y, \|\nabla f(x) - \nabla f(y)\|_* \leq \beta\|x - y\|$.

To prove Lemma 1, we use Lemma 2, Lemma 3, and Lemma 4 below.

The following lemma is Theorem 6 in [25].

Lemma 2. If f is convex and closed, the following two conditions are equivalent:

1. $\forall x, y, f(y) \geq f(x) + \langle \nabla f(x), y - x \rangle + \frac{\beta}{2}\|x - y\|^2$;
2. $\forall x, y, f^*(y) \leq f^*(x) + \langle \nabla f^*(x), y - x \rangle + \frac{1}{2\beta}\|x - y\|_*^2$

i.e. f is β -strongly convex w.r.t some norm $\|\cdot\|$ if and only if f^* is $\frac{1}{\beta}$ -strongly smooth w.r.t the dual norm $\|\cdot\|_*$.

The next lemma is a special case of Lemma 17 in [34].

Lemma 3. Let f be a closed, convex, and differentiable function. Then we have

$$f^*(\nabla f(x)) + f(x) = \langle \nabla f(x), x \rangle.$$

Now we prove a technical result that describes a property of the gradient of the Fenchel Conjugate.

Lemma 4. Suppose f is α -strongly convex for some $\alpha > 0$ with respect to some norm $\|\cdot\|$ and both f and f^* are differentiable. Then we have

$$x = \nabla f^*(\nabla f(x)), \forall x.$$

Proof. For convenience, we define $y = \nabla f(x)$ and $x' = \nabla f^*(y)$. It suffices to prove that $x' = x$.

By Lemma 3, we obtain

$$f^*(y) + f(x) = \langle y, x \rangle = \langle x, y \rangle. \quad (1)$$

Again by Lemma 3, we have

$$f(x') + f^*(y) = f^{**}(x') + f^*(y) = \langle x', y \rangle, \quad (2)$$

where we use the fact that $f^{**} = f$.

Combining inequalities (1) and (2), we obtain

$$0 = f(x) - f(x') - \langle x - x', y \rangle = f(x) + \langle x' - x, \nabla f(x) \rangle - f(x') \leq -\frac{\alpha}{2} \|x - x'\|^2,$$

where in the last inequality we use the definition of α -strongly convex. Therefore we have proved that $x = x'$. \square

Using the three lemmas above, we now prove Lemma 1.

Proof of Lemma 1. By the first condition and Lemma 2, we know f^* is $\frac{1}{\beta}$ -strongly convex with respect to $\|\cdot\|_*$. Therefore we see

$$f^*(\nabla f(y)) \geq f^*(\nabla f(x)) + \langle \nabla f^*(\nabla f(x)), \nabla f(y) - \nabla f(x) \rangle + \frac{1}{2\beta} \|\nabla f(x) - \nabla f(y)\|_*^2.$$

Using Lemma 3 and Lemma 4, we obtain

$$\langle y, \nabla f(y) \rangle - f(y) \geq (\langle x, \nabla f(x) \rangle - f(x)) + \langle x, \nabla f(y) - \nabla f(x) \rangle + \frac{1}{2\beta} \|\nabla f(x) - \nabla f(y)\|_*^2.$$

Rearranging the terms, we get

$$f(x) \geq f(y) + \langle x - y, \nabla f(y) \rangle + \frac{1}{2\beta} \|\nabla f(x) - \nabla f(y)\|_*^2,$$

which is the second condition.

The third condition follows from subtracting the second condition from the first condition. \square

Finally, before moving to the proofs of our main results, we prove two properties of the Bregman Divergence that play an important role in the analysis.

Lemma 5. $\forall a, b, c \in \mathbb{R}^d$ and potential h , we have

$$\langle \nabla h(b) - \nabla h(c), c - a \rangle = D_h(a||b) - D_h(a||c) - D_h(c||b).$$

Proof. By the definition of Bregman Divergence, we obtain

$$\begin{aligned} & D_h(a||b) - D_h(a||c) - D_h(c||b) \\ &= (h(a) - h(b) - \langle \nabla h(b), a - b \rangle) - (h(a) - h(c) - \langle \nabla h(c), a - c \rangle) \\ &\quad - (h(c) - h(b) - \langle \nabla h(b), c - b \rangle) \\ &= -\langle \nabla h(b), a - b \rangle + \langle \nabla h(c), a - c \rangle + \langle \nabla h(b), c - b \rangle \\ &= (-\langle \nabla h(b), a - b \rangle + \langle \nabla h(b), c - b \rangle) + \langle \nabla h(c), a - c \rangle \\ &= \langle \nabla h(b), c - a \rangle + \langle \nabla h(c), a - c \rangle \\ &= \langle \nabla h(b) - \nabla h(c), c - a \rangle. \end{aligned}$$

\square

Lemma 6. For all $a, b, c \in \mathbb{R}^d$, we have

$$D_h(c||a) - D_h(c||b) = D_h(0||a) - D_h(0||b) + \langle \nabla h(b) - \nabla h(a), c \rangle.$$

Proof. Using the definition of Bregman divergence, we obtain

$$\begin{aligned} D_h(c||a) - D_h(c||b) &= h(c) - h(a) - \langle \nabla h(a), c - a \rangle - h(c) + h(b) + \langle \nabla h(b), c - b \rangle \\ &= (h(b) - \langle \nabla h(b), b \rangle) - (h(a) - \langle \nabla h(a), a \rangle) + \langle \nabla h(b) - \nabla h(a), c \rangle \\ &= D_h(0||a) - D_h(0||b) + \langle \nabla h(b) - \nabla h(a), c \rangle. \end{aligned}$$

□

A Proof of Theorem 1

We consider a sequence of hitting cost functions on the real line such that the algorithm stays at the starting point through time steps $t = 1, 2, \dots, n$ and is forced to incur a huge movement cost at time step $t = n + 1$, whereas the offline adversary can pay relatively little cost by dividing the long trek between x_0 and v_{n+1} into multiple small steps through time steps $t = 1, 2, \dots, n + 1$.

Specifically, suppose the starting point of the algorithm and the offline adversary is $x_0 = x_0^* = 0$, and the hitting cost functions are

$$f_t(x) = \begin{cases} \frac{m}{2}x^2 & t \in \{1, 2, \dots, n\} \\ \frac{m'}{2}(x-1)^2 & t = n+1 \end{cases}$$

for some large parameter m' that we choose later.

Suppose the algorithm first moves at time step t_0 . If $t_0 < n + 1$, we stop the game at time step t_0 and compare the algorithm with an offline adversary which always stays at $x = 0$. The total cost of offline adversary is 0, but the total cost of the algorithm is non-zero. So, the competitive ratio is unbounded.

Next we consider the case where $t_0 \geq n + 1$. This implies that $x_1, \dots, x_n = 0$ and x_{n+1} is some non-zero point, say x . We see that the cost incurred by the online algorithm is

$$\text{cost}(ALG) \geq \min_{x_{n+1}} (M_{n+1} + H_{n+1}) = \min_x \left(\frac{1}{2}x^2 + \frac{m'}{2}(x-1)^2 \right).$$

Notice that the right hand side tends to $\frac{1}{2}$ as m' tends to infinity; specifically, we have

$$\text{cost}(ALG) \geq \min_x \left(\frac{1}{2}x^2 + \frac{m'}{2}(x-1)^2 \right) = \frac{1}{2\left(1 + \frac{1}{m'}\right)}. \quad (3)$$

Now let us consider the offline optimal. Notice that, in the limit as m' tends to infinity, the offline optimal must satisfy $x_0^* = 0$ and $x_{n+1}^* = 1$; otherwise it would incur unbounded cost. Our lower bound is derived by considering the case when $m' \rightarrow \infty$ and so we constrain the adversary to satisfy the above, knowing that the adversary is not optimal for finite m' , i.e., $\text{cost}(ADV) \geq \text{cost}(OPT)$ with $\text{cost}(ADV) \rightarrow \text{cost}(OPT)$ as $m' \rightarrow \infty$.

Let the sequence of points the adversary chooses as $x^* = (x_0^*, x_1^*, \dots, x_{n+1}^*) \in \mathbb{R}^{n+2}$. We compute the cost incurred by the adversary as follows where, to simplify presentation, we define $\mathcal{K}(n, y)$ to be the set $\{x \in \mathbb{R}^{n+2} \mid x_i \leq x_{i+1}, x_0 = 0, x_{n+1} = y\}$.

$$\begin{aligned} a_n &= 2 \min_{x^* \in \mathcal{K}(n, 1)} \sum_{i=1}^{n+1} (H_i^* + M_i^*) \\ &= 2 \min_{x^* \in \mathcal{K}(n, 1)} \left(\sum_{i=1}^n \frac{m}{2}(x_i^*)^2 + \sum_{i=1}^{n+1} \frac{1}{2}(x_i^* - x_{i-1}^*)^2 \right). \end{aligned}$$

In words, a_n is twice the minimal offline cost subject to the constraints $x_0^* = 0, x_{n+1}^* = 1$. We derive the limiting behavior of the offline costs as $n \rightarrow \infty$ in the following lemma.

Lemma 7. For $m > 0$, define

$$a_n = 2 \min_{x^* \in \mathcal{K}(n,1)} \left(\sum_{i=1}^n \frac{m}{2} (x_i^*)^2 + \sum_{i=1}^{n+1} \frac{1}{2} (x_i^* - x_{i-1}^*)^2 \right).$$

Then we have $\lim_{n \rightarrow \infty} a_n = \frac{-m + \sqrt{m^2 + 4m}}{2}$.

Given the lemma, the total cost of the offline adversary will be $\frac{a_n}{2}$. Finally, applying (3), we know $\forall n$ and $\forall m' > 0$,

$$\frac{\text{cost}(ALG)}{\text{cost}(ADV)} \geq \frac{\frac{1}{2(1+\frac{1}{m'})}}{\frac{a_n}{2}} = \frac{1}{(1+\frac{1}{m'})a_n}.$$

By taking the limit $n \rightarrow \infty$ and $m' \rightarrow \infty$ and using Lemma 7, we obtain

$$\frac{\text{cost}(ALG)}{\text{cost}(OPT)} = \lim_{n, m' \rightarrow \infty} \frac{\text{cost}(ALG)}{\text{cost}(ADV)} \geq \left(\frac{-m + \sqrt{m^2 + 4m}}{2} \right)^{-1} = \frac{1 + \sqrt{1 + \frac{4}{m}}}{2}.$$

All that remains is to prove Lemma 7, which describes the cost of the offline adversary in the limit as n tends to infinity.

Proof of Lemma 7. Using the fact that the costs are all homogeneous of degree 2, we see that for all $y \in [0, 1]$, we have

$$\begin{aligned} & \min_{x^* \in \mathcal{K}(n,y)} \left(\sum_{i=1}^n \frac{m}{2} (x_i^*)^2 + \sum_{i=1}^{n+1} \frac{1}{2} (x_i^* - x_{i-1}^*)^2 \right) \\ &= y^2 \min_{x^* \in \mathcal{K}(n,1)} \left(\sum_{i=1}^n \frac{m}{2} (x_i^*)^2 + \sum_{i=1}^{n+1} \frac{1}{2} (x_i^* - x_{i-1}^*)^2 \right). \end{aligned} \quad (4)$$

The sequence $\{a_n\}$, $n \geq 0$ has a recursive relationship as follows:

$$\begin{aligned} a_{n+1} &= 2 \min_{x^* \in \mathcal{K}(n+1,1)} \left(\sum_{i=1}^{n+1} \frac{m}{2} (x_i^*)^2 + \sum_{i=1}^{n+2} \frac{1}{2} (x_i^* - x_{i-1}^*)^2 \right) \\ &= 2 \min_{0 \leq x \leq 1} \left(\min_{x^* \in \mathcal{K}(n,x)} \left(\sum_{i=1}^n \frac{m}{2} (x_i^*)^2 + \sum_{i=1}^{n+1} \frac{1}{2} (x_i^* - x_{i-1}^*)^2 \right) \right. \\ &\quad \left. + \frac{m}{2} x^2 + \frac{1}{2} (1-x)^2 \right) \\ &= 2 \min_{0 \leq x \leq 1} \left(x^2 \min_{x^* \in \mathcal{K}(n,1)} \left(\sum_{i=1}^n \frac{m}{2} (x_i^*)^2 + \sum_{i=1}^{n+1} \frac{1}{2} (x_i^* - x_{i-1}^*)^2 \right) \right. \\ &\quad \left. + \frac{m}{2} x^2 + \frac{1}{2} (1-x)^2 \right) \\ &= 2 \min_{0 \leq x \leq 1} \left(\frac{a_n}{2} x^2 + \frac{m}{2} x^2 + \frac{1}{2} (1-x)^2 \right) \\ &= \frac{a_n + m}{a_n + m + 1}. \end{aligned} \quad (5)$$

Solving the equation $x = \frac{x+m}{x+m+1}$, we find the two fixed points of the recursive relationship $a_{n+1} = \frac{a_n+m}{a_n+m+1}$ are

$$x_1 = \frac{-m + \sqrt{m^2 + 4m}}{2},$$

and

$$x_2 = \frac{-m - \sqrt{m^2 + 4m}}{2}.$$

Notice that for $i = 1, 2$, we have

$$m - (m + 1)x_i = -(1 - x_i)x_i.$$

Using this property, we obtain

$$a_{n+1} - x_1 = \frac{a_n + m}{a_n + m + 1} - x_1 = \frac{(1 - x_1)a_n + m - (m + 1)x_1}{a_n + m + 1} = \frac{(1 - x_1)(a_n - x_1)}{a_n + m + 1}, \quad (6)$$

and

$$a_{n+1} - x_2 = \frac{a_n + m}{a_n + m + 1} - x_2 = \frac{(1 - x_2)a_n + m - (m + 1)x_2}{a_n + m + 1} = \frac{(1 - x_2)(a_n - x_2)}{a_n + m + 1}. \quad (7)$$

Notice that $a_{n+1} - x_2 > 0$. By dividing equations (6) and (7), we obtain

$$\left(\frac{a_{n+1} - x_1}{a_{n+1} - x_2} \right) = \frac{1 - x_1}{1 - x_2} \cdot \left(\frac{a_n - x_1}{a_n - x_2} \right), \forall n \geq 0.$$

Remember that $a_0 = 1$. Therefore we have

$$\left(\frac{a_n - x_1}{a_n - x_2} \right) = \left(\frac{1 - x_1}{1 - x_2} \right)^n \left(\frac{a_0 - x_1}{a_0 - x_2} \right) = \left(\frac{1 - x_1}{1 - x_2} \right)^{n+1}.$$

Rearranging this equation, we get

$$a_n = \left(1 - \left(\frac{1 - x_1}{1 - x_2} \right)^{n+1} \right)^{-1} \left(x_1 - x_2 \cdot \left(\frac{1 - x_1}{1 - x_2} \right)^{n+1} \right).$$

Since $0 < \left(\frac{1 - x_1}{1 - x_2} \right) < 1$, we have

$$\lim_{n \rightarrow \infty} a_n = x_1 = \frac{-m + \sqrt{m^2 + 4m}}{2}. \quad (8)$$

□

B Proof of Theorem 2

Our proof of Theorem 2 relies on a set of technical lemmas, which follow. Lemma 8 and Lemma 10 work together to establish a lower bound on the competitive ratio as m tends to zero when the balance parameter γ is set to be $o(1/m)$, while Lemma 11 lower bound on the competitive ratio as m tends to zero when the balance parameter γ is set to be $\Omega(1/m)$.

Lemma 8. *If $\gamma = o(1/m)$, the competitive ratio of OBD is $\Omega(1/(\gamma m))$ when $m \rightarrow 0^+$.*

Proof. Our approach is to construct a scenario where OBD is forced to move along the circumference of a large circle, but the offline adversary moves along the circumference of a much smaller circle (see Figure 1). The adversary is hence able to pay much smaller movements costs, forcing the competitive ratio to be large.

We propose a series of costs which force OBD to move in a circle. The idea is to construct a cost function so that, at the end of every round, the relative positions of the OBD algorithm, the offline adversary, and the minimizer are fixed. Since OBD is memoryless, we can simply input this function arbitrarily many times and the positions of OBD and the offline adversary will trace out a pair of concentric circles (see Figure 1).

Suppose that, at the start of a round, OBD is at the point A . Let ℓ be the distance between OBD and the adversary. Consider a right triangle ABC such that $|AB| = h = \sqrt{\gamma m \ell}$, the offline adversary is at some point D on the hypotenuse AC and $|AD| = |BC| = \ell$ (see Figure 2). Let us introduce a coordinate system such that the origin lies at C , the x -axis contains BC and the y -axis is parallel

Figure 2: In the right triangle $\triangle ABC$, $\angle ABC = 90^\circ$, $|BC| = \ell$, $|AB| = h = \sqrt{\gamma m} \ell$. Point D is on the line segment AC such that $|AD| = \ell$. OBD starts at point A and selects point E . The offline adversary starts at point D and selects point F . G is the projection point of E on line segment AB .

to AB , such that the positive part of the axis lies on the same side of BC as the segment AC . Our goal is to construct a cost function which forces OBD towards B . This will preserve the relative positions of OBD and the adversary, since we assumed that they were a distance ℓ away at the start of the round. Consider the costs $g(u) = \frac{m}{2} \|u - C\|^2$, $h(u) = \alpha \cdot d(u, BC)$ where $d(u, BC)$ is the distance from the point u to the line passing through B and C and $\alpha > 0$ is a parameter we will pick later. Define $f_t(u) = h(u) + g(u)$. Notice that f_t is m -strongly convex because it is the sum of an m -strongly convex function and a convex function. Intuitively, when α is large, the function f_t is infinity outside of the line BC but is equal to $g(u) = \frac{m}{2} \|u - C\|^2$ when restricted to points u on the line. After observing the cost f_t , OBD will pick some new point E .

The following lemma highlights that E can be driven arbitrarily close to B by taking α to be sufficiently large.

Lemma 9. *Let $\varepsilon > 0$, and suppose α is picked to that $\alpha > \frac{hm\ell^2}{\varepsilon^2}$. Then the point E picked by OBD satisfies $|EB| < \varepsilon$.*

We instruct the adversary to pick the point F on the line BC (the x -axis) such that $EF = \ell$ (see Figure 2). Notice that $|CF| = |BF| - |BC| \leq |BE| + |EF| - |BC| = |EB| + \ell - \ell < \varepsilon$, where we used the triangle inequality. Let $z = |DC|$. We see that the total cost incurred by the offline adversary is

$$M_t^* + H_t^* = \frac{1}{2} |DF|^2 + \frac{m}{2} |CF|^2 \leq \frac{1}{2} (|DC| + |CF|)^2 + \frac{m}{2} |CF|^2 \leq \frac{1}{2} (z + \varepsilon)^2 + \frac{m\varepsilon^2}{2},$$

where we applied the triangle inequality.

Notice that $h = |AB| = \sqrt{|AC|^2 - |BC|^2}$ by the Pythagorean theorem (recall that ABC is a right triangle). Since $|AC| = \ell + z$ and $|BC| = \ell$, we see that $h = \sqrt{2z\ell + z^2}$. Hence the movement cost incurred by the OBD is

$$M_t \geq \frac{1}{2} (h - \varepsilon)^2 = \frac{1}{2} (\sqrt{2z\ell + z^2} - \varepsilon)^2.$$

Hence the ratio of the costs is

$$\frac{M_t + H_t}{M_t^* + H_t^*} \geq \frac{M_t}{M_t^* + H_t^*} \geq \frac{\frac{1}{2} (\sqrt{2z\ell + z^2} - \varepsilon)^2}{\frac{1}{2} (z + \varepsilon)^2 + \frac{m\varepsilon^2}{2}}.$$

Since the limit of this expression as $\varepsilon \rightarrow 0$ is $\frac{2z\ell + z^2}{z^2}$, for sufficiently small ε this will be at least $\frac{1}{2} \frac{2z\ell + z^2}{z^2} \geq \frac{\ell}{z}$. Since $z = \sqrt{h^2 + \ell^2} - \ell$ and $h = \sqrt{\gamma m} \ell$, the ratio of costs is at least

$$\frac{\ell}{\sqrt{\gamma m \ell^2 + \ell^2} - \ell} = \frac{1}{\sqrt{\gamma m + 1} - 1} = \frac{\sqrt{\gamma m + 1} + 1}{\gamma m} \geq \frac{2}{\gamma m}.$$

Now, we describe the whole process. When $t = 1$, the hitting cost function is $f_1(x) = \frac{m}{2} \|x\|_2^2$. While OBD stays at $x = 0$, the adversary moves to the point $(\ell, 0)$; it incurs a one-time cost of $M_1^* + H_1^* = \frac{1}{2} \ell^2 + \frac{m}{2} \ell^2$. On all subsequent steps $t = 2 \dots T$, we repeatedly apply the construction, which forces OBD to move in a circle. The one-time cost incurred by the adversary to setup the game is negligible in the limit as T is large, and the per-round ratio of costs is $\Omega(\frac{1}{\gamma m})$, so the competitive ratio is also $\Omega(\frac{1}{\gamma m})$ as claimed. \square

The key technical lemma used in the proof is Lemma 9, and we now provide a proof of that result.

Figure 3: Balance condition at time step t in Lemma 10. Starting from x_{t-1} , OBD picks x_t after observing the hitting cost function $f_t(x) = \frac{m}{2}(x-t)^2$.

Proof of Lemma 9. Suppose $\alpha > \frac{hm\ell^2}{\varepsilon^2}$. We first show that OBD selects the point E strictly contained by the $\frac{m}{2}\ell^2$ -level set, which is the one B lies on. First observe that the point B satisfies the balance condition: $\frac{1}{2}|AB|^2 = \gamma\frac{m}{2}|BC|^2$, because we constructed ABC so that $|AB| = h = \sqrt{\gamma m}\ell$ and $|BC| = \ell$. However, the point B is not necessarily a projection of A onto any level set of f_t . If OBD projected onto the level set which B lies on, it would incur less cost than if it moved to B ; however then the balance condition would be violated. To restore the balance condition, we must increase the movement cost while decreasing the hitting cost – which means we must move to a strictly smaller level set, say the $\frac{m}{2}l_1^2$ -level set, where $l_1 < l$.

Let E_y denote the y -coordinate of E , using the coordinate system we define in the proof of Lemma 8. Notice that $E_y = \frac{g(E)}{\alpha}$, since $g(E)$ was defined to be the vertical distance to the x -axis times α . Since $g(E) \leq f_t(E)$, we see that $E_y \leq \frac{f_t(E)}{\alpha} = \frac{ml_1^2}{2\alpha} \leq \frac{ml^2}{2\alpha}$, where we used the fact that E lies on the $\frac{m}{2}l_1^2$ level set and $l_1 \leq l$. By the balance condition, $\frac{1}{2}|AE|^2 = \frac{\gamma m}{2}l_1^2 \leq \frac{\gamma m}{2}l^2 = \frac{1}{2}h^2$. Let G be the point with coordinates (B_x, E_y) . Applying the Pythagorean theorem successively to the right triangle BEG and the right triangle AGE , we see that

$$\begin{aligned} |EB|^2 &= |E_x - B_x|^2 + E_y^2 \leq (|AE|^2 - (|AB| - E_y)^2) + E_y^2 \\ &\leq (|AB|^2 - (|AB| - E_y)^2) + E_y^2 \leq 2h \cdot E_y \leq h \frac{ml^2}{\alpha}, \end{aligned} \quad (9)$$

where we used the fact that $|AB| \geq |AE|$ and $|AB| = h$. Since we picked $\alpha > \frac{hm\ell^2}{\varepsilon^2}$, we see that $|EB| < \varepsilon$. □

Now we move on to the next technical lemma in the proof of Theorem 2.

Lemma 10. *When $\gamma = o(\frac{1}{m})$, the competitive ratio of OBD is $\Omega(\sqrt{\frac{\gamma}{m}})$.*

Proof. We consider a sequence of cost functions on the real line such that the OBD algorithm moves far away from the starting point, incurring significant movement costs, whereas the offline adversary could pay relatively little cost by staying at the starting point. More specifically, we consider the sequence of hitting cost functions $f_t(x) = \frac{m}{2}(x-t)^2$, $t = 1, 2, \dots, n$. The value of n will be picked later. We assume the starting point is at zero.

Notice that by the balance condition we always have $M_t = \gamma H_t$, so $\frac{1}{2}\|x_t - x_{t-1}\|^2 = \gamma\frac{m}{2}\|x_t - t\|^2$. We can rearrange this expression to obtain $\frac{x_t - x_{t-1}}{t - x_{t-1}} = \sqrt{\gamma m}$. Define

$$\lambda = \frac{x_t - x_{t-1}}{t - x_{t-1}} = \frac{\sqrt{\gamma m}}{1 + \sqrt{\gamma m}}.$$

We obtain the recursive equation $x_t = x_{t-1} + (t - x_{t-1})\lambda$ with initial condition $x_0 = 0$. Solving this equation, we obtain $x_t = t - \frac{1-\lambda}{\lambda}(1 - (1-\lambda)^t)$.

Suppose we picked n to be $\lceil \frac{1}{\lambda} \rceil$. By assumption, $\gamma = o(\frac{1}{m})$; hence in the limit as m tends to zero, λ also tends to zero. Notice that $x_n = n - \frac{1-\lambda}{\lambda}(1 - (1-\lambda)^n) \geq \frac{1}{\lambda} \frac{1}{2e} - (1 - \frac{1}{e}) \geq \frac{1}{6\lambda}$ for sufficiently small λ . Here we used the fact that $(1-\lambda)^{\frac{1}{\lambda}} \rightarrow e^{-1}$.

Suppose the next cost function is $f_{n+1}(x) = m'x^2$. Notice that if the offline adversary simply stays at zero throughout the game, the total cost it incurs would be

$$\text{cost}(ADV) = \frac{m}{2}(1^2 + 2^2 + \dots + n^2) \leq \frac{mn^3}{2} = \Theta\left(\frac{m}{\lambda^3}\right) = \Theta\left(\frac{1}{\sqrt{\gamma^3 m}}\right).$$

In the last step, we used the fact that λ tends to $\sqrt{\gamma m}$ when $\gamma = o(\frac{1}{m})$ and m tends to zero.

If we pick m' large enough that OBD is forced to incur movement cost at least $\frac{1}{2}(\frac{x_n}{2})^2$, the total cost incurred by OBD is

$$\text{cost}(OBD) \geq \frac{1}{2} \left(\frac{x_n}{2} \right)^2 = \Theta \left(\frac{1}{\lambda^2} \right) = \Theta \left(\frac{1}{\gamma m} \right).$$

Putting these facts together, we see that the competitive ratio is at least $\Theta(\sqrt{\frac{\gamma}{m}})$. \square

The last technical lemma used to proof Theorem 2 is the following.

Lemma 11. *When $\gamma = \Omega(\frac{1}{m})$, the competitive ratio of OBD is $\Omega(\frac{1}{m})$.*

Proof. Since $\gamma = \Omega(\frac{1}{m})$, we can assume there exists $C > 0$ such that $\gamma \geq C/m$. We again consider a situation such that the OBD algorithm moves far away from the starting point, incurring significant movement cost, whereas the offline adversary could pay relatively little cost by staying at the starting point. More specifically, suppose the starting point is zero and the first cost function is $f_1(x) = \frac{m}{2}(1-x)^2$. Suppose the adversary stays at zero. The cost incurred by the adversary will be

$$\text{cost}(ADV) = \frac{m}{2}.$$

Notice that by the balance condition ($M_t = \gamma H_t$), the point x_1 picked by OBD satisfies $\frac{x_1^2}{2} = \gamma \frac{m}{2}(1-x_1)^2$. So the cost incurred by OBD is lower bounded by

$$\text{cost}(OBD) \geq M_1 = \frac{1}{2} \left(\frac{\sqrt{\gamma m}}{1 + \sqrt{\gamma m}} \right)^2 \geq \frac{1}{2} \left(\frac{\sqrt{C}}{1 + \sqrt{C}} \right)^2.$$

Since C is a positive constant, the competitive ratio of OBD is lower bounded by $\frac{OBD}{ADV} = \Theta(\frac{1}{m})$. \square

Now we return to the proof of Theorem 2. This proof is a straightforward combination of the above lemmas. When $\gamma = o(\frac{1}{m})$, by combining Lemma 8 and Lemma 10, we know the competitive ratio is at least $\max\left(\frac{C_1}{\gamma m}, C_2 \sqrt{\frac{\gamma}{m}}\right)$ for some positive constants C_1, C_2 . Notice that function $\frac{C_1}{\gamma m}$ is monotonically decreasing in γ and $C_2 \sqrt{\frac{\gamma}{m}}$ is monotonically increasing in γ . Solving the equation $\frac{C_1}{\gamma m} = C_2 \sqrt{\frac{\gamma}{m}}$, we get $\gamma = \left(\frac{C_1}{C_2}\right)^{\frac{2}{3}} m^{-\frac{1}{3}}$. Therefore we see that

$$\max\left\{\frac{C_1}{\gamma m}, C_2 \sqrt{\frac{\gamma}{m}}\right\} \geq C_1^{\frac{1}{3}} C_2^{\frac{2}{3}} m^{-\frac{2}{3}} = \Theta(m^{-\frac{2}{3}}).$$

On the other hand, when $\gamma = \Omega(\frac{1}{m})$, by Lemma 11, we know the competitive ratio of OBD is lower bounded by $\Theta(\frac{1}{m})$.

Together, the above implies that the competitive ratio of OBD is at least $\Theta(m^{-\frac{2}{3}})$ when $m \rightarrow 0^+$.

C Proof of Theorem 3

To begin, note that it is sufficient to prove result for all positive $m \leq \frac{9}{64}$. Similarly, it also suffices to show Theorem 3 when the minimum of every hitting cost function is zero, since otherwise the competitive ratio can only improve if this is not the case.

Our argument makes use of the following potential function: $\phi(x_t, x_t^*) = \eta \|x_t - x_t^*\|^2$. We define $\Delta\phi = \phi(x_t, x_t^*) - \phi(x_{t-1}, x_{t-1}^*)$ and $\Delta\phi' = \phi(x_t', x_t^*) - \phi(x_{t-1}, x_{t-1}^*)$. It suffices to show that $H_t + M_t + \Delta\phi \leq C(H_t^* + M_t^*)$, for some positive constant C . From this inequality, we can sum over all timesteps t to yield that the competitive ratio is upper bounded by C :

$$\sum_{t=0}^T H_t + M_t \leq \sum_{t=0}^T H_t + M_t + \Delta\phi \leq C \sum_{t=0}^T (H_t^* + M_t^*).$$

Throughout the proof, we fix $\eta = 4$ and use $\|\cdot\|$ to denote ℓ_2 norm. When we refer to generalized mean inequality, we mean

$$(a + b)^2 \leq 2a^2 + 2b^2, \forall a, b \in \mathbb{R}.$$

We define $H'_t := f_t(x'_t)$ and $M'_t := c(x'_t, x_{t-1}) = \frac{1}{2} \|x'_t - x_{t-1}\|_2^2$, where x'_t is the point chosen by the first OBD phase (line 3) of Algorithm 2.

Before we move to the main casework in the proof, we begin with a technical lemma that we use to bound the change in the potential function.

Lemma 12. *Suppose the potential function $\phi : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$ is defined as $\phi(a, b) = \eta \|a - b\|^2$, where $\eta > 0$. Then $\forall \lambda > 0$, the change in potential satisfies*

$$\phi(a, c) - \phi(a, b) \leq (1 + \lambda^2)\phi(b, c) + \frac{1}{\lambda^2}\phi(a, b),$$

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

Proof. Using the triangle inequality, we obtain

$$\|a - c\|^2 \leq (\|a - b\| + \|b - c\|)^2 = \|a - b\|^2 + \|b - c\|^2 + 2\|a - b\|\|b - c\|.$$

Rearranging the terms, we obtain

$$\begin{aligned} \|a - c\|^2 - \|a - b\|^2 &\leq \|b - c\|^2 + 2\|a - b\|\|b - c\| \\ &= \|b - c\|^2 + 2\left(\frac{1}{\lambda}\|a - b\|\right)(\lambda\|b - c\|) \\ &\leq (1 + \lambda^2)\|b - c\|^2 + \frac{1}{\lambda^2}\|a - b\|^2, \end{aligned}$$

where in the last line we use the AM-GM inequality. \square

We are now ready to precede with the proof, which is divided up into two cases based on the relationship between the hitting cost of the algorithm and that of the adversary.

Case 1: $H'_t \leq H_t^*$

Since the hitting cost function satisfies $f_t(x) \geq \frac{m}{2} \|x - v_t\|^2$, by the triangle inequality, we have

$$\|x'_t - x_t^*\| \leq \|x'_t - v_t\| + \|x_t^* - v_t\| \leq \left(\sqrt{\frac{2H'_t}{m}} + \sqrt{\frac{2H_t^*}{m}} \right). \quad (10)$$

Thus the change in potential satisfies

$$\begin{aligned} \frac{1}{\eta} \Delta \phi' &= \|x'_t - x_t^*\|^2 - \|x_{t-1} - x_{t-1}^*\|^2 \\ &= (\|x'_t - x_t^*\| - \|x_{t-1} - x_{t-1}^*\|)(\|x'_t - x_t^*\| + \|x_{t-1} - x_{t-1}^*\|) \\ &\leq (\|x'_t - x_{t-1}\| + \|x_t^* - x_{t-1}^*\|)(\|x'_t - x_{t-1}\| + \|x_t^* - x_{t-1}^*\| + 2\|x'_t - x_t^*\|) \quad (11a) \end{aligned}$$

$$\begin{aligned} &= (\|x'_t - x_{t-1}\| + \|x_t^* - x_{t-1}^*\|)^2 + 2(\|x'_t - x_{t-1}\| + \|x_t^* - x_{t-1}^*\|)\|x'_t - x_t^*\| \\ &\leq 2\|x'_t - x_{t-1}\|^2 + 2\|x_t^* - x_{t-1}^*\|^2 + 2(\|x'_t - x_{t-1}\| + \|x_t^* - x_{t-1}^*\|)\|x'_t - x_t^*\| \quad (11b) \end{aligned}$$

$$\leq 4M'_t + 4M_t^* + 2(\sqrt{2M'_t} + \sqrt{2M_t^*}) \left(\sqrt{\frac{2H'_t}{m}} + \sqrt{\frac{2H_t^*}{m}} \right) \quad (11c)$$

$$\leq 4M'_t + 4M_t^* + \sqrt{\frac{1}{m}} ((\sqrt{2M'_t} + \sqrt{2M_t^*})^2 + (\sqrt{2H'_t} + \sqrt{2H_t^*})^2) \quad (11d)$$

$$\leq 4M'_t + 4M_t^* + \sqrt{\frac{1}{m}} ((4M'_t + 4M_t^*) + (4H'_t + 4H_t^*)) \quad (11e)$$

$$= \left(4 + 4\sqrt{\frac{1}{m}} \right) M'_t + \left(4 + 4\sqrt{\frac{1}{m}} \right) M_t^* + 4\sqrt{\frac{1}{m}} H'_t + 4\sqrt{\frac{1}{m}} H_t^*,$$

where we use the triangle inequality in line (11a); the generalized mean inequality in lines (11b), (11d) and (11e) and inequality (10) in line (11c).

Using the OBD's balance condition $M'_t = \gamma H'_t$ and the assumption $H'_t \leq H_t^*$ based on inequality (11), we have

$$\begin{aligned} \frac{1}{\eta} \Delta \phi' &\leq \left(4 + 4\sqrt{\frac{1}{m}}\right) \gamma H'_t + \left(4 + 4\sqrt{\frac{1}{m}}\right) M'_t + 4\sqrt{\frac{1}{m}} H'_t + 4\sqrt{\frac{1}{m}} H_t^* \\ &\leq \left(4 + 4\sqrt{\frac{1}{m}}\right) \gamma H_t^* + \left(4 + 4\sqrt{\frac{1}{m}}\right) M_t^* + 8\sqrt{\frac{1}{m}} H_t^*. \end{aligned}$$

Notice that by the triangle inequality and the generalized mean inequality, we have that

$$M_t = \frac{1}{2} \|x_t - x_{t-1}\|^2 \leq \frac{1}{2} (\|x'_t - x_{t-1}\| + \|x_t - x'_t\|)^2 \leq \frac{1}{2} (2\|x'_t - x_{t-1}\|^2 + 2\|x_t - x'_t\|^2).$$

Remember that since $\mu = 1$, we have $\|x_t - x'_t\|^2 = m \|x'_t - v_t\|^2$. Using this fact, we derive the following bound on $H_t + M_t + \Delta \phi$:

$$\begin{aligned} H_t + M_t + \Delta \phi &\leq H'_t + \frac{1}{2} \left(2\|x'_t - x_{t-1}\|^2 + 2\|x_t - x'_t\|^2\right) \\ &\quad + \eta(\|x_t - x_t^*\|^2 - \|x'_t - x_t^*\|^2) + \Delta \phi' \\ &\leq H'_t + (2M'_t + m\|x'_t - v_t\|^2) \\ &\quad + \left(\eta \left(1 + \frac{1}{\sqrt{m}}\right) \|x_t - x'_t\|^2 + \eta\sqrt{m} \|x'_t - x_t^*\|^2\right) + \Delta \phi' \end{aligned} \quad (12a)$$

$$\begin{aligned} &\leq H'_t + (2M'_t + m\|x'_t - v_t\|^2) \\ &\quad + \left(\eta \left(1 + \frac{1}{\sqrt{m}}\right) m\|x'_t - v_t\|^2 + \eta\sqrt{m} \left(2\|x'_t - v_t\|^2 + 2\|x_t^* - v_t\|^2\right)\right) \\ &\quad + \Delta \phi' \end{aligned} \quad (12b)$$

$$\leq H'_t + (2M'_t + 2H'_t) + \left(\eta \left(1 + \frac{1}{\sqrt{m}}\right) 2H'_t + \eta\sqrt{m} \left(\frac{4H'_t}{m} + \frac{4H_t^*}{m}\right)\right) + \Delta \phi' \quad (12c)$$

$$\begin{aligned} &= (3 + 2\eta + \frac{6\eta}{\sqrt{m}})H'_t + 2M'_t + 4\eta\frac{H_t^*}{\sqrt{m}} + \Delta \phi' \\ &= \left(3 + 2\eta + \frac{6\eta}{\sqrt{m}} + 2\gamma\right)H'_t + 4\eta\frac{H_t^*}{\sqrt{m}} + \Delta \phi' \\ &\leq \left(3 + 2\eta + \frac{6\eta}{\sqrt{m}} + 2\gamma\right)H_t^* + 4\eta\frac{H_t^*}{\sqrt{m}} + \Delta \phi' \\ &= \left(3 + 2\eta + \frac{10\eta}{\sqrt{m}} + 2\gamma\right)H_t^* + \Delta \phi', \end{aligned} \quad (12d)$$

where we use Lemma 12 in line (12a); the triangle inequality in line (12b); m -strongly convexity of f_t in line (12c); and the assumption $H'_t \leq H_t^*$ in line (12d).

Combining inequalities (11) and (12), we obtain

$$H_t + M_t + \Delta \phi \leq (3 + 2\eta + 2\gamma + 4\eta\gamma + \frac{\eta}{\sqrt{m}}(18 + 4\gamma))H_t^* + \eta(4 + 4\sqrt{\frac{1}{m}})M_t^*. \quad (13)$$

Case 2: $H'_t \geq H_t^*$

In this case, we prove that for any $x_t^*, x_{t-1}^* \in \mathbb{R}^d$, we have

$$H_t + M_t + \Delta \phi \leq \frac{C}{\sqrt{m}}(H_t^* + M_t^*), \quad (14)$$

Figure 4: Starting at x_{t-1} , G -OBD first does projection on to the H'_t level set (red dashed line) in the first phase. The projection point is x'_t . Then G -OBD moves toward the minimizer to obtain point x_t in the second phase. Let the minimizer v_t be the origin. Notice that the three points x_{t-1}, x'_t, v_t defines a plane S . Without loss of generality, we can let axis D_2 be parallel to line $x'_t x_{t-1}$; and let axis D_1 be parallel to the projection hyperplane.

for some positive constant C .

In the proof, we use D_1, D_2, \dots, D_d to represent the d axes in the coordinate system.

As shown in Figure 4, without loss of generality, let $v_t = (0, 0, \dots, 0)$, $x'_t = (h_1, h_2, 0, \dots, 0)$ and $D_2 = h_2$ be the projection hyper plane, where $h_1 \geq 0, h_2 \geq 0$. And let $l = \|x_{t-1} - x'_t\| > 0$. Note that our analysis still holds in one-dimension because we can restrict ourselves to the D_2 axis.

Then we know $x_{t-1} = (h_1, h_2 + l, 0, \dots, 0)$, $x_t = (h_1(1 - \sqrt{m}), h_2(1 - \sqrt{m}), 0, \dots, 0)$. Since we know x_t^* must lie below the projection hyper plane, we can let $x_t^* = (x, h_2 - y, a_3, a_4, \dots, a_d)$, where $y > 0$.

Now we show that it suffices to prove the statement when x_{t-1}^* is on the line segment $x_t^* x_{t-1}$. Suppose x_{t-1}^* is not on the line segment $x_t^* x_{t-1}$. If $\|x_{t-1}^* - x_{t-1}\| > \|x_t^* - x_{t-1}\|$, by moving x_{t-1}^* to x_t^* , $\Delta\phi$ increases and M_t^* decreases. Otherwise, we can choose a point K on line segment $x_t^* x_{t-1}$ such that $\|K - x_{t-1}\| = \|x_{t-1}^* - x_{t-1}\|$. By moving x_{t-1}^* to K , $\Delta\phi$ remains unchanged and M_t^* decreases. Therefore if inequality (14) holds for x_{t-1}^* on the segment $x_t^* x_{t-1}$, then it must also hold for any other $x_{t-1}^* \in \mathbb{R}^d$.

Now we suppose x_{t-1}^* is on the line segment $x_t^* x_{t-1}$, and $\|x_t^* - x_{t-1}^*\| = \lambda \|x_t^* - x_{t-1}\|$.

Recall that we set $\gamma = 1$, so $M'_t = \gamma H'_t = H'_t$. It follows that

$$M_t \leq l^2 + \|x_t - x'_t\|^2 = l^2 + m(h_1^2 + h_2^2) \leq l^2 + 2H'_t = l^2 + 2M'_t \leq 2l^2,$$

and

$$H_t \leq H'_t = M'_t = \frac{l^2}{2}.$$

We can separate $\Delta\phi$ into two parts:

$$\frac{\Delta\phi}{\eta} = \left(\|x_t^* - x_t\|^2 - \|x_t^* - x_{t-1}\|^2 \right) + \left(\|x_t^* - x_{t-1}\|^2 - \|x_{t-1}^* - x_{t-1}\|^2 \right).$$

For convenience, we define

$$\Delta\phi_1 := \left(\|x_t^* - x_t\|^2 - \|x_t^* - x_{t-1}\|^2 \right),$$

and

$$\Delta\phi_2 := \left(\|x_t^* - x_{t-1}\|^2 - \|x_{t-1}^* - x_{t-1}\|^2 \right).$$

We further notice that from the triangle inequality,

$$\Delta\phi_2 \leq (1 - (1 - \lambda)^2) \|x_t^* - x_{t-1}\|^2 = \lambda(2 - \lambda) \left((x - h_1)^2 + (y + l)^2 + \sum_{i=3}^d a_i^2 \right). \quad (15)$$

Now we express M_t^* and H_t^* in terms of the variables we define, which are

$$M_t^* = \frac{1}{2} (\lambda \|x_t^* - x_{t-1}\|)^2 = \frac{\lambda^2}{2} \left((x - h_1)^2 + (y + l)^2 + \sum_{i=3}^d a_i^2 \right), \quad (16)$$

and

$$H_t^* \geq \frac{m}{2} \|x_t^* - v_t\|^2 = \frac{m}{2} \left(x^2 + (h_2 - y)^2 + \sum_{i=3}^d a_i^2 \right). \quad (17)$$

We also expand $\Delta\phi_1$:

$$\begin{aligned}
\Delta\phi_1 &= \|x_t^* - x_t\|^2 - \|x_t^* - x_{t-1}\|^2 \\
&= (x - h_1 + h_1\sqrt{m})^2 + (y - h_2\sqrt{m})^2 + \sum_{i=3}^d a_i^2 - (x - h_1)^2 - (y + l)^2 - \sum_{i=3}^d a_i^2 \\
&= ((x - h_1 + h_1\sqrt{m})^2 - (x - h_1)^2) + ((y - h_2\sqrt{m})^2 - (y + l)^2) \\
&= h_1\sqrt{m}(2x - 2h_1 + h_1\sqrt{m}) - (h_2\sqrt{m} + l)(2y + l - h_2\sqrt{m}) \\
&= 2xh_1\sqrt{m} - 2h_1^2\sqrt{m} + h_1^2m - 2y(h_2\sqrt{m} + l) - l^2 + h_2^2m.
\end{aligned} \tag{18}$$

Using the condition that $m \leq \frac{9}{64} < 1$, we derive the following bound:

$$\begin{aligned}
\Delta\phi_1 &\leq 2xh_1\sqrt{m} - 2h_1^2\sqrt{m} + h_1^2m - 2y(h_2\sqrt{m} + l) - l^2 + h_2^2\sqrt{m} \\
&= 2xh_1\sqrt{m} - 2h_1^2\sqrt{m} + h_1^2m + \sqrt{m}(h_2 - y)^2 - \sqrt{m}y^2 - 2yl - l^2.
\end{aligned} \tag{19}$$

Substituting equations (16) and (17) into inequality (14), we know that it suffices to show that for some constant C ,

$$M_t + H_t + \eta\Delta\phi_1 + \eta\Delta\phi_2 \leq \frac{C}{\sqrt{m}} \left(\frac{m}{2} (x^2 + (h_2 - y)^2 + \sum_{i=3}^d a_i^2) + \frac{\lambda^2}{2} ((x - h_1)^2 + (y + l)^2 + \sum_{i=3}^d a_i^2) \right). \tag{20}$$

Subcase 2.1: $\lambda \leq \frac{\sqrt{m}}{2}$

We can bound equation (15) as follows:

$$\begin{aligned}
\Delta\phi_2 &= \lambda(2 - \lambda) \left((x - h_1)^2 + (y + l)^2 + \sum_{i=3}^d a_i^2 \right) \\
&\leq \sqrt{m}(x - h_1)^2 + \sqrt{m}(y + l)^2 + \sqrt{m} \sum_{i=3}^d a_i^2 \\
&= \sqrt{m}x^2 - 2\sqrt{m}xh_1 + \sqrt{m}h_1^2 + \sqrt{m}y^2 + 2\sqrt{m}yl + \sqrt{m}l^2 + \sqrt{m} \sum_{i=3}^d a_i^2.
\end{aligned} \tag{21}$$

Summing inequalities (19) and (21), we get

$$\begin{aligned}
\Delta\phi_1 + \Delta\phi_2 &\leq \sqrt{m}x^2 + (-h_1^2\sqrt{m} + h_1^2m) + \sqrt{m}(h_2 - y)^2 \\
&\quad + (2\sqrt{m}yl - 2yl) + (\sqrt{m}l^2 - l^2) + \sqrt{m} \sum_{i=3}^d a_i^2 \\
&\leq \sqrt{m}x^2 + 0 + \sqrt{m}(h_2 - y)^2 + 0 - \frac{5}{8}l^2 + \sqrt{m} \sum_{i=3}^d a_i^2 \\
&\leq \sqrt{m}x^2 + \sqrt{m}(h_2 - y)^2 - \frac{5}{8}l^2 + \sqrt{m} \sum_{i=3}^d a_i^2,
\end{aligned} \tag{22a}$$

where we use the condition that $m \leq \frac{9}{64}$ in line (22a). We further obtain

$$\begin{aligned}
M_t + H_t + \eta(\Delta\phi_1 + \Delta\phi_2) &\leq 2l^2 + \frac{l^2}{2} + \eta \left(\sqrt{m}x^2 + \sqrt{m}(h_2 - y)^2 - \frac{5}{8}l^2 + \sqrt{m} \sum_{i=3}^d a_i^2 \right) \\
&= \frac{5l^2}{2} + 4 \left(\sqrt{m}x^2 + \sqrt{m}(h_2 - y)^2 - \frac{5}{8}l^2 + \sqrt{m} \sum_{i=3}^d a_i^2 \right) \\
&= 4 \left(\sqrt{m}x^2 + \sqrt{m}(h_2 - y)^2 + \sqrt{m} \sum_{i=3}^d a_i^2 \right).
\end{aligned}$$

Therefore, for $C \geq 8$, we have

$$M_t + H_t + \eta \Delta \phi_1 + \eta \Delta \phi_2 \leq \frac{C}{\sqrt{m}} \left(\frac{m}{2} \left(x^2 + (h_2 - y)^2 + \sum_{i=3}^d a_i^2 \right) + \frac{\lambda^2}{2} \left((x - h_1)^2 + (y + l)^2 + \sum_{i=3}^d a_i^2 \right) \right),$$

which establishes inequality (20).

Subcase 2.2: $\lambda \geq \frac{\sqrt{m}}{2}$

Notice that when $C \geq 32$, we have

$$\frac{C}{2\sqrt{m}} \lambda^2 \geq \frac{16}{\sqrt{m}} \lambda^2 \geq \frac{16}{\sqrt{m}} \cdot \frac{\sqrt{m}}{2} \lambda = 8\lambda \geq 4\lambda(2 - \lambda) = \eta\lambda(2 - \lambda).$$

Substituting this inequality into equation (15), we know that for $C \geq 32$,

$$\eta \Delta \phi_2 \leq \frac{C}{\sqrt{m}} \cdot \frac{\lambda^2}{2} \left((x - h_1)^2 + (y + l)^2 + \sum_{i=3}^d a_i^2 \right). \quad (23)$$

We can further bound inequality (19):

$$\begin{aligned} \Delta \phi_1 &\leq 2xh_1\sqrt{m} - 2h_1^2\sqrt{m} + h_1^2m + \sqrt{m}(h_2 - y)^2 - \sqrt{m}y^2 - 2yl - l^2 \\ &\leq \sqrt{m}x^2 + \sqrt{m}h_1^2 - 2h_1^2\sqrt{m} + h_1^2m + \sqrt{m}(h_2 - y)^2 - l^2 \\ &\leq \sqrt{m}x^2 + \sqrt{m}(h_2 - y)^2 - l^2, \end{aligned}$$

where we apply the AM-GM inequality in step 2 and use the condition $m < 1$ in step 3.

Therefore we have

$$\begin{aligned} H_t + M_t + \eta \Delta \phi_1 &\leq \frac{5l^2}{2} + 4(\sqrt{m}x^2 + \sqrt{m}(h_2 - y)^2 - l^2) \\ &\leq 4(\sqrt{m}x^2 + \sqrt{m}(h_2 - y)^2). \end{aligned} \quad (24)$$

Summing inequalities (24) and (23), we yield that for $C \geq 32$,

$$M_t + H_t + \eta \Delta \phi_1 + \eta \Delta \phi_2 \leq \frac{C}{\sqrt{m}} \left(\frac{m}{2} \left(x^2 + (h_2 - y)^2 + \sum_{i=3}^d a_i^2 \right) + \frac{\lambda^2}{2} \left((x - h_1)^2 + (y + l)^2 + \sum_{i=3}^d a_i^2 \right) \right),$$

which establishes inequality (20).

Combining all cases above, we conclude that G-OBDA is an $O(\frac{1}{\sqrt{m}})$ -competitive algorithm.

D Proof of Theorem 4

To prove Theorem 4 we make use of Lemma 1 and 5.

Our approach is to make use of strong convexity and properties of Bregman Divergences to derive an inequality in the form of $H_t + M_t + \Delta \phi \leq C(H_t^* + M_t^*)$ for some positive constant C , where $\Delta \phi$ is the change in potential, which we will define later. The constant C is then an upper bound for the competitive ratio.

To begin, recall that h is assumed to be α -strongly convex and β -strongly smooth with respect to norm $\|\cdot\|$. Thus we can give a trivial bound on Bregman Divergence, namely

$$\forall x, y, \frac{\alpha}{2} \|x - y\|^2 \leq D_h(x||y) \leq \frac{\beta}{2} \|x - y\|^2. \quad (25)$$

Recall that the update rule in Algorithm 3 can be stated as:

$$x_t = \arg \min_x f_t(x) + \lambda_1 D_h(x||x_{t-1}) + \lambda_2 D_h(x||v_t).$$

Since the function $f_t(x) + \lambda_1 D_h(x||x_{t-1}) + \lambda_2 D_h(x||v_t)$ is strongly convex, the minimizer x_t exists and is unique. Furthermore, it must satisfy the first-order condition

$$\nabla f_t(x_t) + \lambda_1(\nabla h(x_t) - \nabla h(x_{t-1})) + \lambda_2(\nabla h(x_t) - \nabla h(v_t)) = 0.$$

Further, since $f_t(x)$ is m -strongly convex, we have

$$\begin{aligned} f_t(x_t^*) &\geq f_t(x_t) + \langle \nabla f_t(x_t), x_t^* - x_t \rangle + \frac{m}{2} \|x_t^* - x_t\|^2 \\ &= f_t(x_t) - \lambda_1 \langle \nabla h(x_{t-1}) - \nabla h(x_t), x_t - x_t^* \rangle \\ &\quad - \lambda_2 \langle \nabla h(v_t) - \nabla h(x_t), x_t - x_t^* \rangle + \frac{m}{2} \|x_t^* - x_t\|^2. \end{aligned} \quad (26)$$

Using Lemma 5, we obtain

$$\langle \nabla h(x_{t-1}) - \nabla h(x_t), x_t - x_t^* \rangle = D_h(x_t^*||x_{t-1}) - D_h(x_t^*||x_t) - D_h(x_t||x_{t-1}),$$

and

$$\langle \nabla h(v_t) - \nabla h(x_t), x_t - x_t^* \rangle = D_h(x_t^*||v_t) - D_h(x_t^*||x_t) - D_h(x_t||v_t).$$

Substituting the two above identities into inequality (26), we get

$$\begin{aligned} &f_t(x_t) + \lambda_1 D_h(x_t||x_{t-1}) + \lambda_2 D_h(x_t||v_t) + (\lambda_1 + \lambda_2) D_h(x_t^*||x_t) + \frac{m}{2} \|x_t^* - x_t\|^2 \\ &\leq f_t(x_t^*) + \lambda_1 D_h(x_t^*||x_{t-1}) + \lambda_2 D_h(x_t^*||v_t). \end{aligned}$$

It follows that

$$\begin{aligned} &f_t(x_t) + \lambda_1 D_h(x_t||x_{t-1}) + (\lambda_1 + \lambda_2) D_h(x_t^*||x_t) + \frac{m}{2} \|x_t^* - x_t\|^2 \\ &\leq f_t(x_t^*) + \lambda_1 D_h(x_t^*||x_{t-1}) + \lambda_2 D_h(x_t^*||v_t). \end{aligned} \quad (27)$$

We define the potential function as $\phi(x_t, x_t^*) = (\lambda_1 + \lambda_2) D_h(x_t^*||x_t) + \frac{m}{2} \|x_t^* - x_t\|^2$, and let $\Delta\phi = \phi(x_t, x_t^*) - \phi(x_{t-1}, x_{t-1}^*)$. Applying this notation to inequality (27) and rearranging terms, we obtain

$$\begin{aligned} &H_t + \lambda_1 M_t + \Delta\phi \\ &\leq (H_t^* + \lambda_2 D_h(x_t^*||v_t)) + \lambda_1 D_h(x_t^*||x_{t-1}) - (\lambda_1 + \lambda_2) D_h(x_{t-1}^*||x_{t-1}) - \frac{m}{2} \|x_{t-1}^* - x_{t-1}\|^2. \end{aligned} \quad (28)$$

Using Lemma 1, we get

$$\frac{1}{2\beta} \|\nabla h(x_{t-1}) - \nabla h(x_{t-1}^*)\|_*^2 \leq D_h(x_{t-1}^*||x_{t-1}), \quad (29)$$

and

$$\|\nabla h(x_{t-1}) - \nabla h(x_{t-1}^*)\|_* \leq \beta \|x_{t-1} - x_{t-1}^*\|. \quad (30)$$

Using Lemma 5 and the two above inequalities, we get

$$\begin{aligned} & \lambda_1 D_h(x_t^* \| x_{t-1}) - (\lambda_1 + \lambda_2) D_h(x_{t-1}^* \| x_{t-1}) - \frac{m}{2} \|x_{t-1}^* - x_{t-1}\|^2 \\ = & \lambda_1 (D_h(x_t^* \| x_{t-1}) - D_h(x_{t-1}^* \| x_{t-1})) - \lambda_2 D_h(x_{t-1}^* \| x_{t-1}) - \frac{m}{2} \|x_{t-1}^* - x_{t-1}\|^2 \end{aligned} \quad (31a)$$

$$\begin{aligned} = & \lambda_1 D_h(x_t^* \| x_{t-1}^*) + \lambda_1 \langle \nabla h(x_{t-1}) - \nabla h(x_{t-1}^*), x_{t-1}^* - x_t^* \rangle \\ & - \lambda_2 D_h(x_{t-1}^* \| x_{t-1}) - \frac{m}{2} \|x_{t-1}^* - x_{t-1}\|^2 \end{aligned} \quad (31b)$$

$$\begin{aligned} \leq & \lambda_1 D_h(x_t^* \| x_{t-1}^*) + \lambda_1 \|\nabla h(x_{t-1}) - \nabla h(x_{t-1}^*)\|_* \|x_{t-1}^* - x_t^*\| \\ & - \lambda_2 D_h(x_{t-1}^* \| x_{t-1}) - \frac{m}{2} \|x_{t-1}^* - x_{t-1}\|^2 \end{aligned} \quad (31c)$$

$$\begin{aligned} \leq & \lambda_1 D_h(x_t^* \| x_{t-1}^*) + \frac{\lambda_2 \beta + m}{2\beta^2} \|\nabla h(x_{t-1}) - \nabla h(x_{t-1}^*)\|_*^2 + \frac{\lambda_1^2 \beta^2}{2(\lambda_2 \beta + m)} \|x_{t-1}^* - x_t^*\|^2 \\ & - \lambda_2 D_h(x_{t-1}^* \| x_{t-1}) - \frac{m}{2} \|x_{t-1}^* - x_{t-1}\|^2 \end{aligned}$$

$$\begin{aligned} = & \lambda_1 D_h(x_t^* \| x_{t-1}^*) + \frac{\lambda_1^2 \beta^2}{2(\lambda_2 \beta + m)} \|x_{t-1}^* - x_t^*\|^2 \\ & + \left(\frac{\lambda_2}{2\beta} \|\nabla h(x_{t-1}) - \nabla h(x_{t-1}^*)\|_*^2 - \lambda_2 D_h(x_{t-1}^* \| x_{t-1}) \right) \\ & + \left(\frac{m}{2\beta^2} \|\nabla h(x_{t-1}) - \nabla h(x_{t-1}^*)\|_*^2 - \frac{m}{2} \|x_{t-1}^* - x_{t-1}\|^2 \right) \end{aligned} \quad (31d)$$

$$\begin{aligned} \leq & \lambda_1 D_h(x_t^* \| x_{t-1}^*) + \frac{\lambda_1^2 \beta^2}{2(\lambda_2 \beta + m)} \|x_{t-1}^* - x_t^*\|^2 \\ \leq & \lambda_1 \left(1 + \frac{\lambda_1 \beta^2}{\alpha(\lambda_2 \beta + m)} \right) D_h(x_t^* \| x_{t-1}^*), \end{aligned} \quad (31e)$$

where we use Lemma 5 in line (31a); Cauchy-Schwartz inequality in line (31b); the AM-GM inequality in the line (31c); inequalities (29) and (30) in line (31d); and inequality (25) in line (31e).

Substituting inequality (31) into inequality (28), we obtain

$$H_t + \lambda_1 M_t + \Delta\phi \leq (H_t^* + \lambda_2 D_h(x_t^* \| v_t)) + \lambda_1 \left(1 + \frac{\lambda_1 \beta^2}{\alpha(\lambda_2 \beta + m)} \right) M_t^*.$$

Using inequality (25) and the fact that f_t is m -strongly convex, we obtain

$$\lambda_2 D_h(x_t^* \| v_t) \leq \frac{\lambda_2 \beta}{2} \|x_t^* - v_t\|^2 \leq \frac{\lambda_2 \beta}{m} H_t^*.$$

Therefore we have

$$H_t + \lambda_1 M_t + \Delta\phi \leq \left(1 + \frac{\lambda_2 \beta}{m} \right) H_t^* + \lambda_1 \left(1 + \frac{\lambda_1 \beta^2}{\alpha(\lambda_2 \beta + m)} \right) M_t^*.$$

Since $0 < \lambda_1 \leq 1$, we have

$$H_t + M_t + \frac{1}{\lambda_1} \Delta\phi \leq \frac{H_t + \lambda_1 M_t + \Delta\phi}{\lambda_1} \leq \frac{m + \lambda_2 \beta}{m \lambda_1} H_t^* + \left(1 + \frac{\beta^2}{\alpha} \cdot \frac{\lambda_1}{\lambda_2 \beta + m} \right) M_t^*.$$

Theorem 4 follows from summing the above inequality over all timesteps t .

E R-OBD with Squared ℓ_2 Norm

When $h(x) = \frac{1}{2} \|x\|_2^2$, the Bregman Divergence $D_h(x\|y)$ is equal to the squared ℓ_2 norm $\frac{1}{2} \|x - y\|_2^2$. Hence, setting $h(x) = \frac{1}{2} \|x\|_2^2$ in Algorithm 3 gives us R-OBD in the squared ℓ_2 setting. In this section, we present a separate proof of Regularized OBD with squared ℓ_2 norm, in order to remove the assumption that the hitting costs $\{f_t\}$ are differentiable.

Theorem 7. Consider hitting cost functions that are m -strongly convex with respect to ℓ_2 norm and movement costs given by $\frac{1}{2} \|x_t - x_{t-1}\|_2^2$. There exists a choice λ_1, λ_2 such that the competitive ratio of Regularized OBD matches the lower bound proved in Theorem 1, i.e. the competitive ratio is at most $\frac{1}{2} \left(1 + \sqrt{1 + \frac{4}{m}}\right)$.

This result follows from the more general bound in Theorem 8 below, which describes the competitive ratio of Algorithm 3 as a function of λ_1, λ_2 .

Theorem 8. Consider hitting cost functions that are m -strongly convex with respect to ℓ_2 norm and movement costs given by $\frac{1}{2} \|x_t - x_{t-1}\|_2^2$. Regularized-OBD (Algorithm 3 with $h(x) = \frac{1}{2} \|x\|_2^2$) with parameters $1 \geq \lambda_1 > 0, \lambda_2 \geq 0$ has competitive ratio at most

$$\max \left(\frac{m + \lambda_2}{\lambda_1} \cdot \frac{1}{m}, 1 + \frac{\lambda_1}{\lambda_2 + m} \right).$$

Notice that Theorem 7 follows immediately by setting $\frac{m + \lambda_2}{\lambda_1} = \frac{m}{2} \left(1 + \sqrt{1 + \frac{4}{m}}\right)$ in Theorem 8.

Before proving Theorem 8, we first prove a technical lemma which gives a lower bound of the value of hitting cost as a function of the distance to the minimizer.

Lemma 13. If $f : \mathcal{X} \rightarrow \mathbb{R}$ is a m -strongly convex function with respect to some norm $\|\cdot\|$, and v is the minimizer of f (i.e. $v = \arg \min_{x \in \mathcal{X}} f(x)$), then we have $\forall x \in \mathcal{X}$,

$$f(x) \geq f(v) + \frac{m}{2} \|x - v\|^2.$$

Proof. By the definition of m -strongly convex, we obtain that $\forall \alpha \in (0, 1)$,

$$f(\alpha x + (1 - \alpha)v) \leq \alpha f(x) + (1 - \alpha)f(v) - \frac{m}{2} \alpha(1 - \alpha) \|x - v\|^2. \quad (32)$$

Notice that $f(v) \leq f(\alpha x + (1 - \alpha)v)$. Combining this with inequality (32), we obtain that $\forall \alpha \in (0, 1)$,

$$f(v) \leq \alpha f(x) + (1 - \alpha)f(v) - \frac{m}{2} \alpha(1 - \alpha) \|x - v\|^2.$$

Rearranging the terms, we observe that $\forall \alpha \in (0, 1)$,

$$f(x) \geq f(v) + \frac{m}{2} (1 - \alpha) \|x - v\|^2.$$

Therefore

$$f(x) \geq \lim_{\alpha \rightarrow 0^+} \left(f(v) + \frac{m}{2} (1 - \alpha) \|x - v\|^2 \right) = f(v) + \frac{m}{2} \|x - v\|^2.$$

□

Now we return to the proof of Theorem 8.

Proof of Theorem 8. In the proof, we use the property of strongly convex to derive an inequality in the form of $H_t + M_t + \Delta\phi \leq C(H_t^* + M_t^*)$, where $\Delta\phi$ is the change in potential and C is an upper bound for the competitive ratio.

Throughout the proof, we use $\|\cdot\|$ to denote ℓ_2 norm.

Notice that when $h(x) = \frac{1}{2} \|x\|^2$, the update rule in Algorithm 3 is:

$$x_t = \arg \min_x f_t(x) + \frac{\lambda_1}{2} \|x - x_{t-1}\|^2 + \frac{\lambda_2}{2} \|x - v_t\|^2.$$

For convenience, we define

$$F_t(x) = f_t(x) + \frac{\lambda_1}{2} \|x - x_{t-1}\|^2 + \frac{\lambda_2}{2} \|x - v_t\|^2.$$

Since $f_t(x)$ is m -strongly convex, $\frac{\lambda_1}{2} \|x - x_{t-1}\|^2$ is λ_1 -strongly convex, and $\frac{\lambda_2}{2} \|x - v_t\|^2$ is λ_2 -strongly convex, $F_t(x)$ is $(m + \lambda_1 + \lambda_2)$ -strongly convex. Since $x_t = \arg \min_x F_t(x)$, by Lemma 13, we obtain

$$F_t(x_t^*) \geq F_t(x_t) + \frac{m + \lambda_1 + \lambda_2}{2} \|x_t^* - x_t\|^2,$$

which implies

$$\begin{aligned} & H_t + \lambda_1 M_t + \frac{m + \lambda_1 + \lambda_2}{2} \|x_t^* - x_t\|^2 \\ & \leq H_t + \lambda_1 M_t + \frac{\lambda_2}{2} \|x - v_t\|^2 + \frac{m + \lambda_1 + \lambda_2}{2} \|x_t^* - x_t\|^2 \\ & \leq H_t^* + \frac{\lambda_1}{2} \|x_t^* - x_{t-1}\|^2 + \frac{\lambda_2}{2} \|x_t^* - v_t\|^2. \end{aligned} \quad (33)$$

We define the potential function as $\phi(x_t, x_t^*) = \frac{m + \lambda_1 + \lambda_2}{2} \|x_t^* - x_t\|^2$ and $\Delta\phi = \phi(x_t, x_t^*) - \phi(x_{t-1}, x_{t-1}^*)$. We then can rewrite inequality (33) as

$$H_t + \lambda_1 M_t + \Delta\phi \leq \left(H_t^* + \frac{\lambda_2}{2} \|x_t^* - v_t\|^2 \right) + \frac{\lambda_1}{2} \|x_t^* - x_{t-1}\|^2 - \frac{m + \lambda_1 + \lambda_2}{2} \|x_{t-1}^* - x_{t-1}\|^2. \quad (34)$$

Additionally

$$\begin{aligned} & \frac{\lambda_1}{2} \|x_t^* - x_{t-1}\|^2 - \frac{m + \lambda_1 + \lambda_2}{2} \|x_{t-1}^* - x_{t-1}\|^2 \\ & \leq \frac{\lambda_1}{2} (\|x_t^* - x_{t-1}^*\| + \|x_{t-1}^* - x_{t-1}\|)^2 - \frac{m + \lambda_1 + \lambda_2}{2} \|x_{t-1}^* - x_{t-1}\|^2 \end{aligned} \quad (35a)$$

$$\begin{aligned} & = \frac{\lambda_1}{2} \|x_t^* - x_{t-1}^*\|^2 + \lambda_1 \|x_t^* - x_{t-1}^*\| \cdot \|x_{t-1}^* - x_{t-1}\| - \frac{m + \lambda_2}{2} \|x_{t-1}^* - x_{t-1}\|^2 \\ & \leq \frac{\lambda_1}{2} \|x_t^* - x_{t-1}^*\|^2 + \frac{\lambda_1^2}{2(m + \lambda_2)} \|x_t^* - x_{t-1}^*\|^2 + \frac{m + \lambda_2}{2} \|x_{t-1}^* - x_{t-1}\|^2 \\ & \quad - \frac{m + \lambda_2}{2} \|x_{t-1}^* - x_{t-1}\|^2 \end{aligned} \quad (35b)$$

$$\begin{aligned} & = \frac{\lambda_1(\lambda_1 + \lambda_2 + m)}{2(\lambda_2 + m)} \|x_{t-1}^* - x_{t-1}\|^2 \\ & = \lambda_1 \left(1 + \frac{\lambda_1}{\lambda_2 + m} \right) M_t^*, \end{aligned}$$

where we apply the triangle inequality in line (35a) and AM-GM in line (35b).

Combining inequalities (34) and (35), we obtain

$$H_t + \lambda_1 M_t + \Delta\phi \leq \left(H_t^* + \frac{\lambda_2}{2} \|x_t^* - v_t\|^2 \right) + \lambda_1 \left(1 + \frac{\lambda_1}{\lambda_2 + m} \right) M_t^*. \quad (36)$$

And since $f_t(x)$ is m -strongly convex, we have

$$\frac{\lambda_2}{2} \|x_t^* - v_t\|^2 \leq \frac{\lambda_2}{m} H_t^*.$$

Substituting the above identity into inequality (36) yields

$$H_t + \lambda_1 M_t + \Delta\phi \leq \frac{m + \lambda_2}{m} H_t^* + \lambda_1 \left(1 + \frac{\lambda_1}{m + \lambda_2} \right) M_t^*. \quad (37)$$

Using inequality (37), we obtain

$$H_t + M_t + \frac{1}{\lambda_1} \Delta\phi \leq \frac{H_t + \lambda_1 M_t + \Delta\phi}{\lambda_1} \leq \frac{m + \lambda_2}{\lambda_1 m} H_t^* + \left(1 + \frac{\lambda_1}{m + \lambda_2} \right) M_t^*.$$

Theorem 8 follows from summing the above inequality over all timesteps t . \square

F Proof of Theorem 5

In this proof, we construct counterexamples for two separate cases, based on whether λ_1 is larger or smaller than m . Recall that $\lambda_2 = 0$ throughout the proof.

Case 1: $\lambda_1 > m$

In this case, we show the competitive ratio can be unbounded by proposing a series of identical hitting cost functions on the real number line. We construct a hitting cost function f with minimizer v so that there exists a fixed point $K \neq v$ (i.e. when $x_{t-1} = K$ and $f_t = f$, the algorithm selects $x_t = x_{t-1}$). Since R-OBD is independent of timestep t , we can propose $f_t = f$ for $t = 1, 2, \dots, T$ and let $x_0 = K$. In this scenario, the total cost of R-OBD grows linearly in T . However, by choosing $x_1 = x_2 = \dots = x_T = v$, the total cost incurred by the offline adversary is a constant. Therefore the competitive ratio of R-OBD will be unbounded.

Specifically, consider the hitting cost function

$$f(x) = \begin{cases} \frac{m}{2}(1 - (x+1)^2) & -1 \leq x \leq 0 \\ \frac{m}{2}x^2 & \text{otherwise} \end{cases}.$$

Suppose $x_{t-1} = -1$, then R-OBD will choose x_t such that

$$x_t = \arg \min_x f(x) + \frac{\lambda_1}{2}(x+1)^2.$$

Notice that

$$f(x) + \frac{\lambda_1}{2}(x+1)^2 = \begin{cases} \frac{m}{2} + \frac{\lambda_1 - m}{2}(x+1)^2 & -1 \leq x \leq 0 \\ \frac{m}{2}x^2 + \frac{\lambda_1}{2}(x+1)^2 & \text{otherwise} \end{cases}.$$

Since $\lambda_1 > m$, we see that the quantity above is $\geq \frac{m}{2}$ for all real x , where equality only holds when $x = -1$. It follows that $x_t = x_{t-1} = -1 \neq 0 = v$. Thus $K = -1$ is a fixed point satisfying the requirements described as above.

Case 2: $\lambda_1 \leq m$

We consider a situation such that the R-OBD algorithm moves far away from the starting point, incurring significant movement cost, whereas the offline adversary could pay relatively little cost by staying at the starting point. More specifically, suppose the starting point $x_0 = 0$ and the first hitting cost function is $f_1(x) = \frac{m}{2}(1-x)^2$. Consider an adversary which chooses $x_0 = x_1 = \dots = x_T$. The cost incurred by the adversary is

$$\text{cost}(ADV) = \frac{m}{2}.$$

Using the update rule, the R-OBD algorithm chooses

$$x_1 = \arg \min_x \frac{m}{2}(1-x)^2 + \frac{\lambda_1}{2}x^2 = \frac{m}{m + \lambda_1} \geq \frac{1}{2}.$$

The movement cost incurred by R-OBD is at least

$$\text{cost}(ALG) \geq M_1 = \frac{1}{2}x_1^2 \geq \frac{1}{8}.$$

Thus the competitive ratio is at least

$$\frac{\text{cost}(ALG)}{\text{cost}(ADV)} \geq \frac{1}{4m}.$$

Theorem 5 follows from combining these two cases.

G Proof of Theorem 6

Let $\{x_t^L\}$ be the sequence of points achieving the L -constrained offline optimal. We first prove an upper bound on the difference of hitting costs $f_t(x_t) - f_t(x_t^L)$, and then use this bound to prove a $O\left(G\sqrt{TL}\right)$ upper bound on the regret $\sum_{t=1}^T (f_t(x_t) - f_t(x_t^L) + c(x_t, x_{t-1})) - \sum_{t=1}^T c(x_t^L, x_{t-1}^L)$.

Since the function $f_t(x) + \lambda_1 D_h(x||x_{t-1}) + \lambda_2 D_h(x||v_t)$ is strongly convex, it has a unique minimizer, at which point the gradient vanishes. This is the point x_t which Algorithm 3 picks in round t . We can rearrange the vanishing gradient condition to obtain

$$\nabla f_t(x_t) = \lambda_1 (\nabla h(x_{t-1}) - \nabla h(x_t)) + \lambda_2 (\nabla h(v_t) - \nabla h(x_t)).$$

Therefore by Lemma 5, we have

$$\begin{aligned} \langle \nabla f_t(x_t), x_t - x_t^L \rangle &= \lambda_1 \langle \nabla h(x_{t-1}) - \nabla h(x_t), x_t - x_t^L \rangle + \lambda_2 \langle \nabla h(v_t) - \nabla h(x_t), x_t - x_t^L \rangle \\ &= \lambda_1 (D_h(x_t^L||x_{t-1}) - D_h(x_t^L||x_t) - D_h(x_t||x_{t-1})) \\ &\quad + \lambda_2 (D_h(x_t^L||v_t) - D_h(x_t^L||x_t) - D_h(x_t||v_t)). \end{aligned} \quad (38)$$

Recall that h is α -strongly convex and β -strongly smooth with respect to the norm $\|\cdot\|$, hence

$$\forall x, y, \frac{\alpha}{2} \|x - y\|^2 \leq D_h(x||y) \leq \frac{\beta}{2} \|x - y\|^2. \quad (39)$$

Therefore

$$D_h(x_t^L||v_t) - D_h(x_t^L||x_t) - D_h(x_t||v_t) \leq D_h(x_t^L||v_t) \leq \frac{\beta}{2} \|x_t^L - v_t\|^2 \leq \frac{\beta D^2}{2}.$$

In light of equation (38), we obtain

$$\langle \nabla f_t(x_t), x_t - x_t^L \rangle \leq \lambda_1 (D_h(x_t^L||x_{t-1}) - D_h(x_t^L||x_t) - D_h(x_t||x_{t-1})) + \frac{\beta D^2}{2} \cdot \lambda_2. \quad (40)$$

Let $q > 0$ be a parameter which we will pick later. For all $q > 0$, it holds that

$$\begin{aligned} &f_t(x_t) - f_t(x_t^L) \\ &\leq \langle \nabla f_t(x_t), x_t - x_t^L \rangle - \frac{m}{2} \|x_t - x_t^L\|^2 \end{aligned} \quad (41a)$$

$$\begin{aligned} &\leq \lambda_1 (D_h(x_t^L||x_{t-1}) - D_h(x_t^L||x_t) - D_h(x_t||x_{t-1})) - \frac{m}{2} \|x_t - x_t^L\|^2 + \frac{\beta D^2}{2} \cdot \lambda_2 \quad (41b) \\ &= (\lambda_1 + q) (D_h(x_t^L||x_{t-1}) - D_h(x_t^L||x_t)) - \lambda_1 D_h(x_t||x_{t-1}) \\ &\quad - \left(q D_h(x_t^L||x_{t-1}) - q D_h(x_t^L||x_t) + \frac{m}{2} \|x_t - x_t^L\|^2 \right) \\ &\quad + \frac{\beta D^2}{2} \cdot \lambda_2. \end{aligned}$$

where we apply strong convexity in line (41a), and equation (40) in line (41b). Using Lemma 5, we obtain

$$\begin{aligned}
& qD_h(x_t^L||x_{t-1}) - qD_h(x_t^L||x_t) + \frac{m}{2} \|x_t - x_t^L\|^2 \\
&= qD_h(x_t||x_{t-1}) + q\langle \nabla h(x_{t-1}) - \nabla h(x_t), x_t - x_t^L \rangle + \frac{m}{2} \|x_t - x_t^L\|^2 \\
&\geq qD_h(x_t||x_{t-1}) - q\|\nabla h(x_{t-1}) - \nabla h(x_t)\|_* \|x_t - x_t^L\| + \frac{m}{2} \|x_t - x_t^L\|^2 \tag{42a}
\end{aligned}$$

$$\geq qD_h(x_t||x_{t-1}) - \left(\frac{q^2}{2m} \|\nabla h(x_{t-1}) - \nabla h(x_t)\|_*^2 + \frac{m}{2} \|x_t - x_t^L\|^2 \right) + \frac{m}{2} \|x_t - x_t^L\|^2 \tag{42b}$$

$$\begin{aligned}
&= qD_h(x_t||x_{t-1}) - \frac{q^2}{2m} \|\nabla h(x_{t-1}) - \nabla h(x_t)\|_*^2 \\
&\geq qD_h(x_t||x_{t-1}) - \frac{\beta q^2}{m} D_h(x_t||x_{t-1}) \tag{42c} \\
&= \left(q - \frac{\beta q^2}{m} \right) D_h(x_t||x_{t-1}),
\end{aligned}$$

where we apply the Cauchy-Schwartz inequality in line (42a), the AM-GM inequality in line (42b), and Lemma 1 in line (42c).

In order to maximize the coefficient $\left(q - \frac{\beta q^2}{m} \right)$, we set $q = \frac{m}{2\beta}$. By substituting inequality (42) into inequality (41), we obtain

$$\begin{aligned}
& f_t(x_t) - f_t(x_t^L) \\
&\leq \left(\lambda_1 + \frac{m}{2\beta} \right) \left(D_h(x_t^L||x_{t-1}) - D_h(x_t^L||x_t) \right) - \left(\lambda_1 + \frac{m}{4\beta} \right) D_h(x_t||x_{t-1}) + \frac{\beta D^2}{2} \cdot \lambda_2. \tag{43}
\end{aligned}$$

Using the condition $\lambda_1 + \frac{m}{4\beta} \geq 1$, we observe that

$$f_t(x_t) - f_t(x_t^L) + D_h(x_t||x_{t-1}) \left(\lambda_1 + \frac{m}{2\beta} \right) \left(D_h(x_t^L||x_{t-1}) - D_h(x_t^L||x_t) \right) + \frac{\beta D^2}{2} \cdot \lambda_2. \tag{44}$$

Notice that

$$\sum_{t=1}^T \|x_t^L - x_{t+1}^L\| \leq \sqrt{T \left(\sum_{t=1}^T \|x_t^L - x_{t+1}^L\|^2 \right)} \leq \sqrt{T \left(\sum_{t=1}^T \frac{2D_h(x_{t+1}^L||x_t^L)}{\alpha} \right)} \leq \sqrt{\frac{2TL}{\alpha}}. \tag{45}$$

where we use the generalized mean inequality in the first step and α -strong convexity of h in the second step (cf. equation (39)). By Lemma 6, we can give the following upper bound:

$$\begin{aligned}
& \sum_{t=1}^T D_h(x_t^L || x_{t-1}) - D_h(x_t^L || x_t) \\
&= \sum_{t=1}^T (D_h(0 || x_{t-1}) - D_h(0 || x_t) + \langle \nabla h(x_t) - \nabla h(x_{t-1}), x_t^L \rangle) \\
&= D_h(0 || x_0) - D_h(0 || x_T) + \sum_{t=1}^{T-1} \langle \nabla h(x_t), x_t^L - x_{t+1}^L \rangle - \langle \nabla h(x_0), x_1^L \rangle + \langle \nabla h(x_T), x_T^L \rangle \\
&\leq \sum_{t=1}^T \langle \nabla h(x_t), x_t^L - x_{t+1}^L \rangle \tag{46a}
\end{aligned}$$

$$\leq \sum_{t=1}^T \|\nabla h(x_t)\|_* \|x_t^L - x_{t+1}^L\| \tag{46b}$$

$$\begin{aligned}
&\leq G \sum_{t=1}^T \|x_t^L - x_{t+1}^L\| \\
&\leq G \sqrt{\frac{2TL}{\alpha}}, \tag{46c}
\end{aligned}$$

where we use the facts $x_0 = x_0^L = x_{T+1}^L = 0, \nabla h(0) = 0$ in line (46a), the Cauchy-Schwartz inequality in line (46b), and inequality (45) in line (46c).

Therefore we obtain

$$\begin{aligned}
& \text{cost}(OBD) - \text{cost}(OPT(L)) \\
&= \sum_{t=1}^T (f_t(x_t) + D_h(x_t || x_{t-1})) - (f_t(x_t^L) + D_h(x_t^L || x_{t-1}^L)) \tag{47a}
\end{aligned}$$

$$\begin{aligned}
&\leq \left(\sum_{t=1}^T f_t(x_t) - f_t(x_t^L) + D_h(x_t || x_{t-1}) \right) - L \\
&\leq \left(\lambda_1 + \frac{m}{2\beta} \right) G \sqrt{\frac{2TL}{\alpha}} + T \cdot \frac{\beta D^2}{2} \cdot \lambda_2 - L, \tag{47b}
\end{aligned}$$

where we use the definition of $OPT(L)$ in line (47a); inequalities (44) and (46) in line (47b).

Since by assumption we have $G < \infty, \lambda_2 = \eta(T, L, D, G) \leq \frac{KG}{D^2} \cdot \sqrt{\frac{L}{T}}$ for some constant K , by inequality (47), we obtain

$$\text{cost}(OBD) - \text{cost}(OPT(L)) = O(G\sqrt{TL}),$$

which completes the proof.