Competitively Chasing Convex Bodies

Sébastien Bubeck
Microsoft Research, USA
sebubeck@microsoft.com

Yuanzhi Li
Stanford University, USA
yuanzhil@stanford.edu

Yin Tat Lee
University of Washington and Microsoft Research, USA
yintat@uw.edu

Mark Sellke
Stanford University, USA
msellke@stanford.edu

ABSTRACT

Let $\mathcal{F}$ be a family of sets in some metric space. In the $\mathcal{F}$ -chasing problem, an online algorithm observes a request sequence of sets in $\mathcal{F}$ and responds (online) by giving a sequence of points in these sets. The movement cost is the distance between consecutive such points. The competitive ratio is the worst case ratio (over request sequences) between the total movement of the online algorithm and the smallest movement one could have achieved by knowing in advance the request sequence. The family $\mathcal{F}$ is said to be *chaseable* if there exists an online algorithm with finite competitive ratio. In 1991, Linial and Friedman conjectured that the family of convex sets in Euclidean space is chaseable. We prove this conjecture.

CCS CONCEPTS

• Theory of computation → Online algorithms.

KEYWORDS

chasing convex bodies, online algorithms

ACM Reference Format:

Sébastien Bubeck, Yin Tat Lee, Yuanzhi Li, and Mark Sellke. 2019. Competitively Chasing Convex Bodies. In *Proceedings of the 51st Annual ACM SIGACT Symposium on the Theory of Computing (STOC '19)*, June 23–26, 2019, Phoenix, AZ, USA. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3313276.3316314>

1 INTRODUCTION

Let $\mathcal{K}$ denote the family of convex sets in $\mathbb{R}^d$, and $\mathcal{K}^*$ the set of strings with alphabet $\mathcal{K}$. A map $S : \mathcal{K}^* \rightarrow \mathbb{R}^d$ is called an *online selector* if for any $t \in \mathbb{N}$ and $(K_1, \dots, K_t) \in \mathcal{K}^t$ one has $S(K_1, \dots, K_t) \in K_t$. In the convex body chasing problem, introduced in [9], the performance of an online selector on a sequence of convex sets $\mathbf{K} \in \mathcal{K}^{\mathbb{N}}$ and a starting state $x_0 \in \mathbb{R}^d$ is measured through its *movement cost* (with $S(\emptyset) := x_0$):

$$\text{cost}_S(\mathbf{K}) := \sum_{t \geq 1} \|S(K_1, \dots, K_{t-1}) - S(K_1, \dots, K_t)\|,$$

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
STOC '19, June 23–26, 2019, Phoenix, AZ, USA

© 2019 Association for Computing Machinery.
ACM ISBN 978-1-4503-6705-9/19/06...\$15.00
<https://doi.org/10.1145/3313276.3316314>

where $\|\cdot\|$ denotes the Euclidean norm. The cost of the *offline optimum*, denoted $\text{cost}^*(\mathbf{K})$, is defined as the infimum of the above quantity over all S , or equivalently (with $o_0 = x_0$):

$$\text{cost}^*(\mathbf{K}) = \inf_{(o_t) \in \mathbf{K}} \sum_{t \geq 1} \|o_{t-1} - o_t\|.$$

We say that an online selector S is ω -competitive if for any $\mathbf{K} \in \mathcal{K}^{\mathbb{N}}$ one has

$$\text{cost}_S(\mathbf{K}) \leq \omega \cdot \text{cost}^*(\mathbf{K}).$$

CONJECTURE 1.1 ([9]). *For any $d \in \mathbb{N}$, there exists an ω -competitive online selector for some $\omega > 1$.*

In [9] this conjecture was proved for $d = 2$, and it remained open until now for any $d \geq 3$. In this paper we resolve this conjecture affirmatively and obtain an exponential scaling on the dimension for the competitive ratio (it is known from [9] that the competitive ratio has to be $\Omega(\sqrt{d})$):

THEOREM 1.1. *For any $d \in \mathbb{N}$, there exists a 2^{30d} -competitive online selector.*

1.1 Motivation

The original motivation for convex body chasing was to better understand the influence of geometric structure in the general chasing problem. Indeed the most central problem in online algorithms, k -server ([14]), can be viewed as chasing sets with a certain combinatorial structure. However, while there has been lots of progress on the k -server problem since the late eighties (see [5, 12] for the current state of the art), the convex body chasing problem has been essentially unstirred (only some very special cases have been solved, see related works section below). Fortunately, more convincing applications for convex body chasing have been proposed in the last decade, in particular by considering the generalized problem of *chasing convex functions*.

In convex function chasing, a request corresponds to a convex cost function $f_t : \mathbb{R}^d \rightarrow \mathbb{R}_+ \cup \{+\infty\}$, instead of merely a convex set as in convex body chasing. In addition to the movement cost between two consecutive points x_{t-1} and x_t , there is now also a service cost of $f_t(x_t)$. It is easy to see that chasing convex functions in dimension d can be reduced to chasing convex bodies in dimension $d + 1$ (one can simply replace a function request f_t by requesting the epigraph of f_t followed by requesting $\mathbb{R}^d \times \{0\}$). This more general setting can now model various problems in resource management (e.g., powering data centers [13]). Interestingly it also applies naturally to online machine learning. In the latter case the movement cost is not explicitly part of the problem, however it

arises naturally from the uncertainty in the online learning model, namely the fact that the decision point x_t at time t is evaluated on the cost function f_{t+1} at time $t + 1$. In particular for Lipschitz cost functions the discrepancy between $f_{t+1}(x_t)$ and $f_{t+1}(x_{t+1})$ is bounded from above by the movement cost $\|x_t - x_{t+1}\|$. We note that, on the contrary to the classical regret analysis of online convex optimization [10], here the resulting online machine learning algorithm would be able to track a slowly shifting concept. The price to pay is a multiplicative guarantee instead of an additive guarantee.

1.2 Related Works

The lack of progress on the Friedman-Linial conjecture prompted the community to consider special cases of convex body chasing, where the request sequence is constrained in some ways. These types of results in fact go back to the original paper, where the question of chasing *lines* was resolved. See also [1] and references therein for more results in that vein. (However, beware that their Lemma 2 is flawed, see <https://people.cs.pitt.edu/~kirk/Errata.html>.) More recently the *nested* version has gotten some attention ([2, 3]): this is the scenario where the request sequence is nested. In a companion paper, [6], we solve this simpler problem and obtain a competitive ratio of $O(\sqrt{d} \log(d))$, thus almost matching the $\Omega(\sqrt{d})$ lower bound. Another natural question recently considered is whether dimension-free competitive ratio can be obtained for special type of convex sets. This question was addressed for convex function chasing in [8], where the dimension-free property was proved for “linearly growing” functions. See also [7] for logarithmic competitive ratio with linear costs and covering LP type sets.

1.3 Notation and Convex Geometry Reminders

We denote $B(x, R)$ for a Euclidean ball centered at x and of radius R , $\mathring{B}(x, R)$ for the corresponding open ball, and $B := B(0, 1)$. The minimum width or minwidth of a convex set K is denoted $\delta(K)$, and the center of gravity or centroid is denoted $\text{cg}(K)$. We also denote P_V for the projection on a subspace V , and $\text{dist}(x, \Omega)$ for the Euclidean distance between a point $x \in \mathbb{R}^d$ and a set $\Omega \subset \mathbb{R}^d$. In order to emphasize key aspects we will sometimes use the O_d and Ω_d notation to hide dimension-dependent constants (eventually all constants are made explicit).

LEMMA 1.2. *Let K and L be two convex bodies such that $\text{cg}(K) \notin L$. Then for any $\varepsilon > 0$ one has*

$$\text{vol}(K \cap (L + \varepsilon B)) \leq \left(1 - \frac{1}{e} + \frac{2d(d+1)}{\delta(K)}\varepsilon\right) \text{vol}(K).$$

PROOF. The approximate Grünbaum’s inequality states that if K is isotropic then the volume decreases by $1 - 1/e + \varepsilon$, [4, Theorem 3]. For non-isotropic K this implies that the volume decreases by $1 - 1/e + \frac{\varepsilon}{\sqrt{\lambda_{\min}}}$ where $\lambda_{\min}$ is the smallest eigenvalue of the covariance matrix of K . It only remains to observe that

$$\delta(K) \leq 2\sqrt{d(d+1)\lambda_{\min}}, \quad (1)$$

which is a consequence of the fact that an isotropic convex body is included in a ball of radius $\sqrt{d(d+1)}$ ([11, Theorem 4.1]). $\square$

LEMMA 1.3. *A convex body K contains a ball of radius $\frac{\delta(K)}{2d(d+1)}$.*

PROOF. Use that an isotropic body contains a ball ([11, Theorem 4.1]) together with (1). $\square$

2 PROOF SKELETON: A MULTISCALE CHASE

Our proposed online selector works in phases, where a phase corresponds to a block of requests. For each phase $b \in \mathbb{N}$, a center $z_b \in \mathbb{R}^d$ and a scale $r_b \geq 0$ will be chosen (with $r_1 = \text{dist}(x_0, K_1)$ and $z_1 = x_0$). As the phase b will always be clear from context, we drop the subscript and simply write $r \equiv r_b$ and $z \equiv z_b$. To describe the update procedure of those parameters we will use the notation $r^{\text{new}} \equiv r_{b+1}$ and $z^{\text{new}} \equiv z_{b+1}$. We adopt similar notation for other parameters. The t^{th} request in a phase is denoted K_t , and the online selector’s response is x_t . It will be useful for us to respond with a point x_t possibly outside of the request K_t , in which case the corresponding movement cost is at most $\|x_{t-1} - x_t\| + 2\text{dist}(x_t, K_t)$. Without loss of generality we can assume $x_{t-1} \notin K_t$.

2.1 Target Properties

The general idea is to ensure the following two properties in a phase:

- (i) The online selector’s total movement is $O_d(r)$. As a first step to ensure this, the online selector will remain in the ball $B(z, R)$ during the duration of the phase (for some $R = O_d(r)$ to be defined later).
- (ii) If the optimal selector’s movement is $\leq r$, then the “distance” between the online selector and the optimal selector gets reduced by $\Omega_d(r)$ during the phase.

With these two properties an estimate on the competitive ratio follows from a standard potential argument. Indeed, in each phase either the optimal selector pays $\Omega(\text{online cost})$, or some potential gets decreased by $\Omega(\text{online cost})$. See Lemma 2.3 for the details.

In Subsection 2.2 we explain the stopping condition for a phase, as well as how to update the parameters r and z . In Subsection 2.3 we give the notion of “distance” (i.e., the potential) that we use to satisfy the Property (ii). In Subsection 2.4 we explain how to respond to requests during a phase to satisfy Property (i). There, a critical and non-trivial “pancake issue” arises, and we devote Section 3 to resolving this issue.

2.2 Tracking a Small Cost Optimal Selector

The online selector keeps track of a convex set $\Omega_t \subset B(z, R)$ that contains all possible locations visited (during the current phase) by optimal selectors that:

- (i) stay in $B(z, R)$ during the phase, and
- (ii) have movement cost $\leq r$ during the phase.

An example for such a set would be to take the intersection of *padded requests*:

$$\tilde{\Omega}_t = B(z, R) \cap \bigcap_{s \leq t} (K_s + rB). \quad (2)$$

Indeed we know that a selector satisfying (i) and (ii) must necessarily lie in $\tilde{\Omega}_t$ (and furthermore $\tilde{\Omega}_t$ is convex). Due to the “pancake issue” (see Subsection 2.4) it will turn out that (2) is too coarse for our needs. However for now the reader is encouraged to think

Figure 1: The phase ends when the convex set Ω_t is trapped in a ball of radius $\alpha \cdot r$. The localization is not acting on the left figure (condition 1) and is acting on the right figure (condition 2).

$\Omega_t \simeq \tilde{\Omega}_t$. We will also require the following condition for Ω_t which is easy to ensure:

- (iii) if there is no path in Ω_t satisfying (i) and (ii) then $\Omega_t = \emptyset$.

The general idea is now to “follow” Ω_t (see Subsection 2.4) until it becomes small enough. Precisely we stop the phase at the first time T when Ω_T is included in a ball of radius $\alpha \cdot r$ for some $\alpha \geq 1$ to be defined later. The radius of the localization ball $B(z, R)$ will turn out to be

$$R = 7\alpha r. \quad (3)$$

When the phase stops there are two possibilities: either $\Omega_T \subset \dot{B}(z, R-r)$ (“localization is not acting”) or $\Omega_T \cap (B(z, R) \setminus \dot{B}(z, R-r)) \neq \emptyset$ (“localization acting”). For reasons to be explained in Subsection 2.5 we also distinguish the case where $\Omega_T = \emptyset$.

- (1) If the localization is not acting and $\Omega_T \neq \emptyset$, then we have trapped the optimal selector (provided that it moved $\leq r$, see Lemma 2.1 below for the details) and thus we can start a new phase at a lower scale, say $r^{\text{new}} = r/2$. We also choose z^{new} to be the center of the enclosing ball for Ω_T .
- (2) If the localization is acting then the optimal selector might be outside of the localized ball $B(z, R)$, and thus we start a new phase at a larger scale, say $r^{\text{new}} = 2r$ and $z^{\text{new}} = z$. We also perform this update when $\Omega_T = \emptyset$.

LEMMA 2.1. *If the phase ends with condition 1 above, then any selector with total movement in the phase $\leq r$ must end the phase in Ω_T .*

PROOF. Assume on the contrary that there exists a path $(o_t)_{t \in [T]}$ with movement $\leq r$ and that visits $\mathbb{R}^d \setminus B(z, R)$ at some point during the phase (by definition of Ω_T we know that a path that stays in $B(z, R)$ and with movement $\leq r$ must be in Ω_T). Since $\Omega_T \neq \emptyset$ we also know by property (iii) above that there exists a path $(o'_t)_{t \in [T]}$ with movement $\leq r$ that stays in $B(z, R)$, and ends the phase in $B(z, R-r)$ (since $\Omega_T \subset B(z, R-r)$). Let $p \in [0, 1]$ be such that $\|po'_T + (1-p)o_T - z\| = R-r$. Then clearly $po' + (1-p)o$ has movement $\leq r$, stays in $B(z, R)$, yet ends in $\partial B(z, R-r)$ which contradicts the fact that $\Omega_T \subset \dot{B}(z, R-r)$. $\square$

2.3 The Potential Argument

Let us denote o (resp. o^{new}) for the location of the optimal selector at the start (resp. end) of the phase. We consider the following potential: $\Phi := \max(\|z - o\| - \alpha r, \alpha r)$. Denote $\Phi^{\text{new}} := \max(\|z^{\text{new}} - o^{\text{new}}\| - \alpha r^{\text{new}}, \alpha r^{\text{new}})$.

LEMMA 2.2. *Let $R = 7\alpha \cdot r$. In each phase (as described in Subsection 2.2) one has*

$$(1 + 8.5\alpha) \cdot \text{cost}^*|_{\text{phase}} + \Phi - \Phi^{\text{new}} \geq \frac{\alpha}{2} r, \quad (4)$$

where $\text{cost}^*|_{\text{phase}}$ denotes the cost of the optimal selector during the current phase.

PROOF. We distinguish three cases:

- If $\text{cost}^*|_{\text{phase}} \geq r$ then we use that (recall that $\|z^{\text{new}} - z\| \leq R = 7\alpha r$):

$$\begin{aligned} \Phi^{\text{new}} &\leq \max(2\alpha r, \|z^{\text{new}} - o^{\text{new}}\|) \\ &\leq \max(2\alpha r, \|z - o\| + \|z^{\text{new}} - z\| + \|o - o^{\text{new}}\|) \\ &\leq \Phi + (1 + 8\alpha) \cdot \text{cost}^*|_{\text{phase}}, \end{aligned}$$

which proves (4).

We now assume that $\text{cost}^*|_{\text{phase}} \leq r$.

- If we end the phase with condition 1 from Subsection 2.2 then by Lemma 2.1 we have that $\|z^{\text{new}} - o^{\text{new}}\| \leq \alpha r$. Thus $\Phi^{\text{new}} = \alpha r/2$ (indeed $r^{\text{new}} = r/2$), while on the other hand $\Phi \geq \alpha r$. This proves (4).
- Finally if we end the phase with condition 2 from Subsection 2.2 then we must have $\|z - o^{\text{new}}\| \geq R - r - 2\alpha r \geq 4\alpha r$, which implies

$$\max(\|z - o^{\text{new}}\| - \alpha r, \alpha r) = \alpha r + \max(\|z - o^{\text{new}}\| - 2\alpha r, 2\alpha r) \geq \alpha r + \Phi^{\text{new}}.$$

where the last inequality uses that $z^{\text{new}} = z$. Thus we have:

$$\begin{aligned} \Phi^{\text{new}} &\leq \max(\|z - o^{\text{new}}\| - \alpha r, \alpha r) - \alpha r \\ &\leq \|o - o^{\text{new}}\| + \Phi - \alpha r \\ &\leq \text{cost}^*|_{\text{phase}} + \Phi - \alpha r, \end{aligned}$$

which concludes the proof of (4). $\square$

We note that Lemma 2.2 together with Property (i) in Subsection 2.1 imply a finite competitive ratio:

LEMMA 2.3. *Consider a phase-based online selector as described in Subsection 2.1 and 2.2. Assume that it has movement at most βr per phase (Property (i) in Subsection 2.1), for some $\beta \geq 1$. Then it is 21β -competitive.*

PROOF. Equation (4) implies that (since $\alpha \geq 1$)

$$r \leq 19 \cdot \text{cost}^*|_{\text{phase}} + \frac{2}{\alpha}(\Phi - \Phi^{\text{new}}).$$

Together with the assumption of βr -movement per phase, as well as the fact that the potential is non-negative and its initial value is smaller than $\alpha \cdot \text{cost}^*$, one obtains:

$$\text{cost}_S \leq 21 \cdot \beta \cdot \text{cost}^*.$$

□

2.4 In-phase Movement with Center of Gravity

Upon receiving a new request K_t (such that $x_{t-1} \notin K_t$) we shall update Ω_{t-1} to Ω_t which will be a subset of $\Omega_{t-1} \cap (K_t + rB)$. Recall also that a phase ends when Ω_t is included in a ball of radius $\alpha \cdot r$. Thus the hope is that x_{t-1} is “deep” inside Ω_{t-1} so that Ω_t gets significantly smaller than Ω_{t-1} . Importantly we note that, since we “cut” with $K_t + rB$ but we only have the guarantee that $x_{t-1} \notin K_t$, we will also want that Ω_{t-1} is “not too small” (the “pancake issue”), so that cutting x_{t-1} or cutting r -close to x_{t-1} has a similar effect.

Let us consider what happens with the simple proposal of moving to the centroid, that is $x_t = \text{cg}(\Omega_t)$ (provided that Ω_t is non-empty, see Subsection 2.5 for the empty case). Note that x_t is not necessarily in K_t , and thus in addition to moving from x_{t-1} to x_t one might incur an extra r movement to actually satisfy the request. Thus after T requests the online selector has moved at most $T \cdot (R + r)$. In fact thanks to Lemma 2.4 below this bound will hold true even if the phase ends with $\Omega_T = \emptyset$. Now the approximate Grünbaum’s inequality (Lemma 1.2) implies that

$$\delta(\Omega_{t-1}) \geq 8d(d+1)r \Rightarrow \text{vol}(\Omega_t) \leq \frac{9}{10} \text{vol}(\Omega_{t-1}). \quad (5)$$

Furthermore, if we can guarantee that the minwidth remains $8d(d+1)r$, then we have that Ω_t always contain a ball of radius $4r$ (see Lemma 1.3), which in turn means that the total volume decrease is at most $(R/(4r))^d \leq (2\alpha)^d$ (by (3)), at which point the phase will stop (we will take $\alpha \geq 4$). Thus we see that the length of a phase is $T \leq d \log_{10/9}(2\alpha)$, which in turn guarantees that the total movement in a phase is $(R + r) \cdot d \log_{10/9}(2\alpha) = O_d(r)$.

We see that the *only* remaining difficulty is to ensure that the minwidth of the set Ω_t remains $\Omega_d(r)$. In our former nested convex body chasing work, [2], we dealt with the small width directions via a recursive argument. Here the situation is much more delicate, and we dedicate the next section to this “pancake issue”. We also note that, at this point of the argument, one has a poly(d)-competitive ratio. The exponential dependency in Theorem 1.1 comes from the recursive argument (which could potentially be improved).

2.5 The Empty Scenario

When $\Omega_t = \emptyset$ we know that a new phase will begin after the current request K_t . However it might be that to satisfy the request K_t one has to pay a movement much larger than r . The next lemma shows a simple reduction which allows us to assume that this never happens.

LEMMA 2.4. *Any request sequence can be modified online so that:*

Figure 2: We use V_t to denote the subspace spanned by small width directions in Ω_t

- Each original request K is replaced by a finite nested sequence ending with K (in particular the value of the offline optimum is the same on the original and modified sequences).
- On the modified sequence, in each phase one has that all requests satisfy $\text{dist}(z, K) \leq R + r$.

PROOF. Let us consider a request K which violates the second point. Then we know that the current phase will end with condition 2 from Subsection 2.2 since $(K + rB) \cap B(z, R) = \emptyset$. In the modified sequence we replace K by $K + hB$ where $h > 0$ is such that $\text{dist}(z, K + hB) = R + r$, so that the current phase still ends with condition 2 but now the “virtual” request $K + hB$ satisfy the second point in Lemma 2.4.

In the next phase we start by giving tentatively K and repeat the procedure above, that is we potentially replace K by $K + h^{\text{new}}B$ if $\text{dist}(z^{\text{new}}, K) > R + r^{\text{new}}$. The key point is that eventually (in a finite number of steps), the request K will be valid, since $r^{\text{new}} = 2r$ will keep doubling while $z^{\text{new}} = z$ remains constant. □

3 INDUCTION ON DIMENSION ARGUMENT

In this section we discuss how to deal with the small width directions in Ω_t . By induction we assume that we have access to an ω_k -competitive algorithm for convex body chasing in $\mathbb{R}^k$, $k < d$ (note that the case $k = 1$ is trivial).

Let us denote V_t for a subspace spanned by a maximal set of orthogonal “small” width directions in Ω_t . More precisely, given V_t and Ω_{t+1} , we define V_{t+1} by the following iterative procedure:

- $V \leftarrow V_t$.
- While there exists $v \in V^\perp$ such that the width of Ω_{t+1} in the direction v is $\leq \frac{\alpha}{2^{\dim V^\perp}}$
 - Set $V \leftarrow \text{span}(V, v)$
- $V_{t+1} \leftarrow V$

Note that if $\dim(V_t) = d$ then we know that Ω_t is included in a hyperrectangle with lengths $\frac{\alpha}{2d}r, \frac{\alpha}{2(d-1)}r, \dots$ and thus it is also included in a Euclidean ball of radius

$$\frac{\alpha r}{2} \sqrt{\frac{1}{d^2} + \frac{1}{(d-1)^2} + \dots + 1} \leq \alpha \cdot r.$$

In particular we obtain that $\dim(V_t) = d$ ensures that the global phase ends at time t . Thus let us assume $\dim(V_t) < d$.

3.1 The Nested Intuition

To set the stage, let us briefly recall the *nested* convex body chasing strategy we proposed in our prior work [2]. In this subsection we view the sequence $\Omega_t, \Omega_{t+1}, \dots$ as a nested convex body chasing request sequence. The entire algorithm and analysis can be summarized in the following two points:

- (1) Let x_t be any point such $P_{V_t^\perp} x_t = \text{cg}(P_{V_t^\perp} \Omega_t)$. One can now play the $\dim(V_t)$ -algorithm on $x_t + V_t$, until one gets an empty set on that affine subspace, that is a time T such that $\Omega_T \cap (x_t + V_t) = \emptyset$. For notational simplicity let us say that $T = t + 1$. Crucially at this point one obtains $P_{V_t^\perp} x_t \notin P_{V_t^\perp} \Omega_{t+1}$, so that by Grünbaum's inequality:

$$\text{vol}(P_{V_t^\perp} \Omega_{t+1}) \leq (1 - 1/e) \cdot \text{vol}(P_{V_t^\perp} \Omega_t). \quad (6)$$

- (2) The natural potential to track progress is $\Psi_t := \text{vol}(P_{V_t^\perp} \Omega_t)$.

We note that while $k = \dim(V_t^\perp)$ stays fixed, the potential drops by $(1 - 1/e)$ at each time step, starting from a value of at most $R^k = (7\alpha r)^k$ and ending with a value of at least $(r\alpha/(4k^2(k+1)))^k$ (by Lemma 1.3 and the minimum width $\geq \alpha/(2k)$). In particular the number of steps where the dimension of $V_t^\perp$ remains k is at most $k \cdot \log_{e/(e-1)}(28k^2(k+1))$.

Combining both points we would obtain a $O(d^2 \log d)$ -competitive algorithm for nested convex body chasing. We note that a slightly more careful analysis of Ψ_t easily yields a $O(d \log d)$ -competitive analysis, but it is irrelevant for the purpose of the present paper.

The essence of the difficulty in the non-nested case (in fact the *only* remaining obstacle at this point) is that one cannot afford to wait until a time where $\Omega_T \cap (x_t + V_t)$ gets empty. Indeed the lower-dimensional algorithm in the affine subspace is only competitive with respect to an optimal selector that stays within that affine subspace. What if such an optimal selector has in fact much worse performance than an unconstrained selector? This is indeed a critical situation for convex body chasing: all requests could have an intersection point far away from the current affine subspace, so that the lower-dimensional algorithm pays a lot while the optimal selector pays nothing. Perhaps the most important conceptual contribution of this work is to realize that in this case one can extract from the request sequence a cutting plane that is *parallel* to the current affine subspace, and thus actually obtain a large volume decrease in the orthogonal subspace just as in (6). This is detailed in Subsection 3.2.

3.2 Leveraging Lower Dimensional Selectors

Let us assume that we start a lower-dimensional phase at some time t_0 in the current global phase, by which we mean that $\dim(V_{t_0}) \in \{1, \dots, d-1\}$. Denote $k = \dim(V_{t_0}^\perp)$. Just like in the nested case we first move to any point x_{t_0} such that $P_{V_{t_0}^\perp} x_{t_0} = \text{cg}(P_{V_{t_0}^\perp} \Omega_{t_0})$, and we now run the $(d-k)$ -dimensional online selector in $x_{t_0} + V_{t_0}$. We now stop the lower-dimensional phase when the movement cost of the $(d-k)$ -dimensional online selector reaches $\zeta_k \cdot \omega_{d-k} \cdot r$, for some ζ_k to be defined later.

For notational simplicity we will assume that $x_{t_0} = 0$, and write $V \equiv V_{t_0}$, $\Omega \equiv \Omega_{t_0}$. Again with an abuse of notation we will now only use a “local time” for the current lower-dimensional phase. In

particular $K_1, K_2, \dots$ now denotes the sequence of requests in this lower-dimensional phase.

At the beginning of the lower-dimensional phase we assume by induction that Ω satisfies the following case analysis (recall property (i) and (ii) in Subsection 2.2). In the global phase, any selector must either:

- have a total cost $\geq r$;
- have visited a point outside of $B(z, R)$; or finally
- remain in Ω .

Our hope is that by the end of the phase (i.e., when the lower-dimensional algorithm has moved at least $\zeta_k \cdot r$) we will be able to refine this set Ω of possible locations for a localized small cost selector into a set Ω^{new} such that

$$\text{vol}(P_{V^\perp} \Omega^{\text{new}}) \leq \frac{9}{10} \cdot \text{vol}(P_{V^\perp} \Omega). \quad (7)$$

LEMMA 3.1. *Assume that (7) holds true for any $k \in [d]$. Then, in a global phase, the total number of calls to the $(d-k)$ -dimensional online selector¹ is at most $50k \log(k+1)$.*

PROOF. The argument is the same as in the second point of Subsection 3.1. Namely we consider the potential $\Psi := \text{vol}(P_{V^\perp} \Omega)$. While $k = \dim(V^\perp)$ stays fixed, the potential drops by $9/10$ with each call to the $(d-k)$ -dimensional online selector. Furthermore the potential starts from a value of at most $R^k = (7\alpha r)^k$, and ends with a value of at least $(r\alpha/(4k^2(k+1)))^k$ (by Lemma 1.3 and the minimum width $\geq \alpha/(2k)$). Thus the number of calls is at most $k \cdot \log_{10/9}(28k^2(k+1)) \leq 50k \log(k+1)$. $\square$

This in turns give the following (recall that after each call to a lower-dimensional selector one has to move to the center of gravity of $P_{(V^{\text{new}})^\perp} \Omega^{\text{new}}$, which adds a movement of $R+r$):

LEMMA 3.2. *Assume that (7) holds true for any $k \in [d]$. Then the total movement in a global phase is at most $\beta \cdot r$, with*

$$\beta = \sum_{k=1}^d 50k \log(k+1) \cdot (\zeta_k \cdot \omega_{d-k} + 7\alpha + 1).$$

Thus we see that it only remains to estimate the value of ζ_k (and of α) so that one can guarantee (7). Note for example that (7) is satisfied for $k = d$ (i.e., $V = \{0\}$, $\delta(\Omega) \geq \frac{\alpha}{2d}$, and there is no call to a lower-dimensional online selector) with $\Omega^{\text{new}} = \Omega \cap \{K + rB\}$ provided that (recall (5))

$$\frac{\alpha}{2d} \geq 8d(d+1). \quad (8)$$

3.3 A Refined Localization of Small Cost Optimal Selectors

Consider the set

$$\Omega'_t = \{y \in \Omega \text{ s.t. } \exists (y_s)_{s \in \{1, \dots, t\}} \text{ with } y_1 \in \Omega, y_t = y, y_s \in K_s \cap \Omega, \text{ and } \sum_{s=1}^{t-1} \|y_s - y_{s+1}\| \leq r\}.$$

Crucially we note that Ω'_t contains the set of possible locations at time t for a selector with movement $\leq r$ and that starts the

¹A call to a 0-dimensional selector simply means receiving a request.

Figure 3: The left figure illustrates that one can construct a selector Y_t in V by taking average of selectors $y_t(\theta)$ starting at $\theta \in V^\perp$ with $\|\theta\| = \gamma r$. The right figure gives a case where there is a good selector (red) outside of V but all selectors (blue) in V are bad. In this case, we can discard a large portion of $V^\perp$.

lower-dimensional phase in Ω . Furthermore Ω'_t is convex (by convexity of Ω , K_s and the function $y \mapsto \|y\|$). Our tentative choice is $\Omega^{\text{new}} := \Omega \cap (\Omega'_T + rB)$ at the time T when the lower-dimensional phase ends. Note that Ω^{new} is indeed a valid refinement of Ω (in the sense that any selector with movement $\leq r$ and that remains in $B(z, R)$ must in fact remain in Ω^{new} for the whole global phase). We now want to show that (7) holds true.

The key non-trivial observation is the following: during the lower-dimensional phase, if the best selector that remains in V has total movement ηr (for some $\eta = \Theta_k(1)$), then it must be that there exists $x \in V^\perp$ with $\|x\| \leq \gamma r$ (for some $\gamma = \Theta_d(1)$) and $x \notin \Omega'_t$ (the precise formulation is given in Lemma 3.3 below). This means that for $\zeta_k \geq \eta$ we know that $x \notin \Omega^{\text{new}}$, and thus since x is close (for γ small enough compared to $\delta(P_{V^\perp}\Omega) \geq \frac{\alpha}{2k}$) to the center of gravity of $P_{V^\perp}\Omega$ (which is assumed to be 0), we expect that (7) will be satisfied, see Subsection 3.4 for the details.

In words this key lemma says that: if when restricted to a subspace V one has to move a lot to satisfy the requests, while at the same time there was a way to satisfy all those requests without moving much, then it must be that we can discard a large portion of $V^\perp$ as possible good locations.

LEMMA 3.3. *Let $V \subset \mathbb{R}^d$ be a $n - k$ dimensional linear subspace. Consider the cylinder $C = \{y \in \mathbb{R}^d : \|P_V y\| \leq \gamma' r\}$ and let*

$$\tilde{\Omega} = \{y \in C \text{ s.t. } \exists (y_t)_{t \in \{1, \dots, T\}} \text{ with } y_T = y, y_t \in K_t \cap C, \text{ and } \sum_{t=1}^{T-1} \|y_t - y_{t+1}\| \leq r\}.$$

Assume that $\mathbb{S} := \{\theta \in V^\perp : \|\theta\| = \gamma r\} \subset P_{V^\perp} \tilde{\Omega}$, where $\gamma \geq 12k$. Then there exists $(Y_t)_{t \in \{1, \dots, T\}}$ with $Y_t \in K_t \cap V$ and $\sum_{t=1}^{T-1} \|Y_t - Y_{t+1}\| \leq \left(1 + \frac{2+4\gamma'}{\gamma} k\right) r$.

PROOF. For any $\theta \in \mathbb{S}$, let us consider the path $(y_t(\theta))$ witnessing that $\theta \in \tilde{\Omega}$. Denote $u_t(\theta) = P_{V^\perp} y_t(\theta)$ and $v_t(\theta) = P_V y_t(\theta)$. Our tentative path (Y_t) will be taken of the form:

$$Y_t = \mathbb{E}[\mu_t(\theta) y_t(\theta)],$$

where the expectation is with respect to a random θ in $\mathbb{S}$, and $\mu_t : \mathbb{S} \rightarrow \mathbb{R}_+$ is a non-negative weight function satisfying $\mathbb{E}[\mu_t(\theta)] = 1$ (so that $Y_t \in K_t$) and $\mathbb{E}[\mu_t(\theta) u_t(\theta)] = 0$ (so that $Y_t \in V$). Let us define $m_t = \mathbb{E}[u_t(\theta)]$ and $\Sigma_t = \mathbb{E}[(u_t(\theta) - m_t)(u_t(\theta) - m_t)^\top]$ (which we view as a linear operator on $V^\perp$). We propose to take the following weight function (which can be derived as the minimum energy function satisfying the constraints):

$$\mu_t(\theta) = 1 + (m_t - u_t(\theta))^\top \Sigma_t^{-1} m_t.$$

Assuming that μ_t is well-defined (i.e., Σ_t is invertible as an operator on $V^\perp$) one easily checks that $\mathbb{E}[\mu_t(\theta)] = 1$ and $\mathbb{E}[\mu_t(\theta) u_t(\theta)] = 0$. Let us now prove that μ_t is indeed well-defined, and also that it is non-negative.

First note that by definition, since $y_T(\theta) = \theta$, one has $m_T = 0$ and $\Sigma_T = (\gamma r / \sqrt{k})^2 \text{Id}$. Let $\xi_t(\theta) = u_t(\theta) - m_t - \theta$, and note that $\|\xi_t\| \leq 2r$ (since the total movement of the $y_t(\theta)$ path is at most r). With this notation we have:

$$\Sigma_t = (\gamma r / \sqrt{k})^2 \text{Id} + \mathbb{E}[\xi_t(\theta) \xi_t(\theta)^\top] + \theta \xi_t(\theta)^\top + \xi_t(\theta) \theta^\top.$$

Note that the spectral norm of the second term is bounded by $4r^2 + 4\gamma r^2$ and thus we see that:

$$\Sigma_t \geq (\gamma^2/k - 4(\gamma + 1))r^2 \text{Id} \geq \frac{\gamma^2 r^2}{2k} \text{Id}, \quad (9)$$

where the second inequality uses that $\gamma \geq 12k$. In particular we see that $\Sigma_t > 0$ and thus μ_t is well-defined. Next using that

$$\|m_t\| \leq r \text{ and } \|u_t(\theta)\| \leq (\gamma + 1)r, \quad (10)$$

we obtain:

$$|u_t(\theta)^\top \Sigma_t^{-1} m_t| \leq 2k(\gamma + 1)/\gamma^2.$$

Thus we have that $\mu_t(\theta) \geq 0$ provided that $\gamma^2 \geq 2k(\gamma + 1)$ (which is true since $\gamma \geq 12k$). We note that (9) and (10) also give:

$$\mu_t(\theta) \leq 1 + 2k(\gamma + 2)/\gamma^2. \quad (11)$$

Finally it remains to estimate the movement of Y_t . For this it will be slightly more convenient to assume that the path is defined in continuous time, so that we want to estimate $\int_0^T \left\| \frac{d}{dt} Y_t \right\| dt$. Note

that since $P_{V^\perp} Y_t = 0$, in fact we only need to estimate the movement of $P_V Y_t = \mathbb{E}[\mu_t(\theta)v_t(\theta)]$. Recall that by assumption $v_t(\theta)$ has movement $\leq r$, and thus with (11) one obtains:

$$\int_0^T \|\mathbb{E}[\mu_t(\theta) \frac{d}{dt} v_t(\theta)]\| dt \leq \left(1 + 2k(\gamma + 2)/\gamma^2\right) r. \quad (12)$$

We also have (since $y_t(\theta) \in C$):

$$\int_0^T \|\mathbb{E}[v_t(\theta) \frac{d}{dt} \mu_t(\theta)]\| dt \leq \gamma' r \mathbb{E} \int_0^T \left| \frac{d}{dt} \mu_t(\theta) \right| dt. \quad (13)$$

Let us calculate:

$$\begin{aligned} \frac{d}{dt} \mu_t(\theta) &= \left(\frac{d}{dt} (m_t - u_t(\theta)) \right)^\top \Sigma_t^{-1} m_t \\ &\quad + (m_t - u_t(\theta))^\top \left(\frac{d}{dt} \Sigma_t^{-1} \right) m_t \\ &\quad + (m_t - u_t(\theta))^\top \Sigma_t^{-1} \frac{d}{dt} m_t. \end{aligned} \quad (14)$$

The first and last term in (14) are easy to control using that m_t and $u_t(\theta)$ have total movement $\leq r$, together with (9) and (10). Precisely we get that the contribution of those two terms to the integral in the right hand side of (13) is at most:

$$\frac{2k}{\gamma^2 r^2} \cdot (2r^2 + (\gamma + 2)r^2) = 2k(\gamma + 4)/\gamma^2. \quad (15)$$

For the middle term in (14) one has

$$\begin{aligned} \frac{d}{dt} \Sigma_t^{-1} &= -\Sigma_t^{-1} \left(\frac{d}{dt} \Sigma_t \right) \Sigma_t^{-1} \\ &= 2\Sigma_t^{-1} \mathbb{E} \left[(u_t(\theta) - m_t) \left(\frac{d}{dt} (u_t(\theta) - m_t) \right)^\top \right] \Sigma_t^{-1}. \end{aligned}$$

In particular we have that the spectral norm of $\frac{d}{dt} \Sigma_t^{-1}$ is bounded by

$$2 \cdot \left(\frac{2k}{\gamma^2 r^2} \right)^2 \cdot (\gamma + 2)r \cdot \mathbb{E} \left[\left\| \frac{d}{dt} (u_t(\theta) - m_t) \right\| \right]$$

so that the contribution of the middle term in (14) to the integral in the right hand side of (13) is at most:

$$(\gamma + 2)r^2 \cdot 2 \cdot (2k/(\gamma^2 r^2))^2 \cdot (\gamma + 2)r \cdot r = 2 \left(2k(\gamma + 2)/\gamma^2 \right)^2. \quad (16)$$

Putting together (15) and (16) with (13), and adding (12), one finally obtains:

$$\frac{1}{r} \int_0^T \left\| \frac{d}{dt} Y_t \right\| dt \leq 1 + \frac{2k(\gamma + 2)}{\gamma^2} + \gamma' \left(\frac{2k(\gamma + 4)}{\gamma^2} + 2 \left(\frac{2k(\gamma + 2)}{\gamma^2} \right)^2 \right),$$

which concludes the proof (up to a straightforward numerical verification using $\gamma \geq 12k$). $\square$

3.4 Proof of Theorem 1.1

Let us apply Lemma 3.3 with $\gamma' = \alpha$ and $\gamma = \frac{\alpha}{16k^2(k+1)}$, and consider ending the phase with $\zeta_k = 65(k+1)^4 \geq \left(1 + k \frac{2+4\gamma'}{\gamma}\right)$. Note that to satisfy the condition $\gamma \geq 12k$ we can take

$$\alpha = 192(d+1)^4, \quad (17)$$

which also satisfies (8). We obtain that when the lower dimensional phase ends, we have a point $\theta \in V^\perp$ such that $\|\theta\| = \gamma r \leq$

$\frac{\delta(P_{V^\perp}(\Omega))}{8k(k+1)} r$ and $\theta \notin P_{V^\perp} \Omega'_T$. Thus by Lemma 1.2 we obtain that (recall that $\Omega^{\text{new}} = \Omega \cap (\Omega'_T + rB)$)

$$\text{vol}(P_{V^\perp} \Omega^{\text{new}}) \leq \frac{9}{10} \text{vol}(P_{V^\perp} \Omega).$$

We can now apply Lemma 3.2 with

$$\begin{aligned} \beta &= \sum_{k=1}^d 50k \log(k+1) \cdot (\zeta_k \cdot \omega_{d-k} + 7\alpha + 1) \\ &\leq (7d)^7 + \sum_{k=1}^d (6k)^6 \omega_{d-k} \leq \sum_{k=1}^d (14k)^6 \omega_{d-k}. \end{aligned}$$

and plug it into Lemma 2.3 to obtain:

$$\omega_d \leq \sum_{k=1}^d (24k)^6 \omega_{d-k}.$$

The proof of Theorem 1.1 is now concluded with a straightforward numerical verification.

Algorithm 1: GlobalPhase(z, r)

```

1 Input: scale  $r > 0$  and center  $z \in \mathbb{R}^d$ .
2 Initialization:  $\alpha \leftarrow 192(d+1)^4$ ,  $R \leftarrow 7\alpha r$ ,  $\Omega \leftarrow B(z, R)$ ,
    $V \leftarrow \{0\}$ ,  $x \leftarrow z$ . //  $x$  represents
   the online selector's location
3 while  $V \neq \mathbb{R}^d$  do
4   Receive request  $K \subset \mathbb{R}^d$  such that  $x \notin K$ .
5   if  $\text{dist}(z, K) > R + r$  then // Lemma 2.4
6     End call and start GlobalPhase( $z, 2r$ ) with  $K$  as first
       request.
7   end
8   Move  $x$  to the closest point in  $K$ . // Satisfy the
       request
9   Let  $\text{cost}^*$  be the minimum cost to service the requests
        $K \cap B(z, R)$  so far.
10  if  $\text{cost}^* \geq r$  then // [Subsection 2.2, (iii)]
11    End call and start GlobalPhase( $z, 2r$ ).
12  end
13  Start LowerDimPhase( $\Omega, V$ ) with  $K$  as first request, and
       obtain  $\Omega'$ .
14   $\Omega \leftarrow \Omega \cap \{\Omega' + rB\}$ 
15  while there is  $v \in V^\perp$  such that  $\delta_v(\Omega) < \frac{\alpha}{2 \dim V^\perp}$  do
16     $V \leftarrow \text{span}(v, V)$ .
17  end
18  Move  $x$  to the closest point such that  $P_{V^\perp} x = \text{cg}(P_{V^\perp} \Omega)$ .
19 end
20 if  $\Omega \subset \overset{\circ}{B}(z, R - r)$  then // Condition 1 in Subsection
    2.2
21   End call and start GlobalPhase( $z^{\text{new}}, r/2$ ) where  $z^{\text{new}}$ 
       is a point s.t.  $\Omega \subset B(z^{\text{new}}, \alpha r)$ .
22 else // Condition 2 in Subsection 2.2
23   End call and start GlobalPhase( $z, 2r$ ).
24 end

```

Algorithm 2: Chasing(x_0)

-
- 1 **Input:** starting point $x_0 \in \mathbb{R}^d$.
 - 2 Receive a first request $K \subset \mathbb{R}^d$ and set $r = \text{dist}(x_0, K)$.
 - 3 Run GlobalPhase(x_0, r).
-

Algorithm 3: LowerDimPhase(Ω, V, r)

-
- 1 **Input:** current set $\Omega \subset \mathbb{R}^d$, subspace $V \subset \mathbb{R}^d$, scale r , and request sequence K_t .
 - 2 **Initialization:** $\zeta \leftarrow 65(\dim(V^\perp) + 1)^2$
 - 3 **if** $\dim(V) = 0$ **then**
 - 4 Receive request K_1 and **Return** $\Omega' = K_1$
 - 5 **else**
 - 6 Let $c \leftarrow \text{cg}(P_{V^\perp} \Omega)$ and run Chasing(c) on the request sequence $K'_t = K_t \cap (c + V)$ until a total movement of $\zeta \cdot \omega_{\dim(V)} \cdot r$.
 - 7 **Return** Ω' where // Used Lemma 3.3 with the fact $\Omega \subset C$
 - 8
$$\Omega' = \{y \in \Omega \text{ s.t. } \exists (y_t)_{t \in \{1, \dots, T\}} \text{ with}$$

$$y_T = y, y_t \in K'_t \cap \Omega, \text{ and } \sum_{t=1}^{T-1} \|y_t - y_{t+1}\| \leq r\}.$$
 - 9 **end**
-

REFERENCES

- [1] Antonios Antoniadis, Neal Barcelo, Michael Nugent, Kirk Pruhs, Kevin Schewior, and Michele Scquizzato. Chasing convex bodies and functions. In *LATIN 2016*:

- theoretical informatics*, volume 9644 of *Lecture Notes in Comput. Sci.*, pages 68–81. Springer, Berlin, 2016.
- [2] CJ Argue, Sébastien Bubeck, Michael B Cohen, Anupam Gupta, and Yin Tat Lee. A nearly-linear bound for chasing nested convex bodies. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, 2019.
- [3] Nikhil Bansal, Martin Böhm, Marek Eliáš, Grigorios Koumoutsos, and Seeun William Umboh. Nested convex bodies are chaseable. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1253–1260. SIAM, 2018.
- [4] Dimitris Bertsimas and Santosh Vempala. Solving convex programs by random walks. *Journal of the ACM (JACM)*, 51(4):540–556, 2004.
- [5] Sébastien Bubeck, Michael B. Cohen, Yin Tat Lee, James R. Lee, and Aleksander Madry. K-server via multiscale entropic regularization. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2018, pages 3–16. ACM, 2018.
- [6] Sébastien Bubeck, Yin-Tat Lee, Yuanzhi Li, and Mark Sellke. Chasing nested convex bodies nearly optimally. 2018.
- [7] Niv Buchbinder, Shahar Chen, and Joseph (Seffi) Naor. Competitive analysis via regularization. In *Proceedings of the Twenty-fifth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '14, pages 436–444, 2014.
- [8] Nianjun Chen, Gautam Goel, and Adam Wierman. Smoothed online convex optimization in high dimensions via online balanced descent. In *Proceedings of the 31st Conference On Learning Theory*, volume 75 of *Proceedings of Machine Learning Research*, pages 1574–1594. PMLR, 2018.
- [9] Joel Friedman and Nathan Linial. On convex body chasing. *Discrete & Computational Geometry*, 9(3):293–321, 1993.
- [10] Elad Hazan. Introduction to online convex optimization. *Foundations and Trends® in Optimization*, 2(3-4):157–325, 2016.
- [11] Ravi Kannan, László Lovász, and Miklós Simonovits. Isoperimetric problems for convex bodies and a localization lemma. *Discrete & Computational Geometry*, 13(3-4):541–559, 1995.
- [12] James R. Lee. Fusible hsts and the randomized k-server conjecture. In *Proceedings of the 59th Annual IEEE Symposium on Foundations of Computer Science*, FOCS 2018, 2018.
- [13] Minghong Lin, Adam Wierman, Lachlan LH Andrew, and Eno Thereska. Dynamic right-sizing for power-proportional data centers. *IEEE/ACM Transactions on Networking (TON)*, 21(5):1378–1391, 2013.
- [14] Mark S. Manasse, Lyle A. McGeoch, and Daniel D. Sleator. Competitive algorithms for server problems. *Journal of Algorithms*, 11(2):208–230, 1990.