

On Convex Body Chasing*

Joel Friedman¹ and Nathan Linial²

¹ Department of Computer Science, Princeton University,
Princeton, NJ 08544, USA

² Department of Computer Science, Hebrew University,
Givat Ram, Jerusalem 91904, Israel

Abstract. A player moving in the plane is given a sequence of instructions of the following type: at step i a planar convex set F_i is specified, and the player has to move to a point in F_i . The player is charged for the distance traveled. We provide a strategy for the player which is competitive, i.e., for any sequence F_i the cost to the player is within a constant (multiplicative) factor of the “off-line” cost (i.e., the least possible cost when all F_i are known in advance). We conjecture that similar strategies can be developed for this game in any Euclidean space and perhaps even in all metric spaces. The analogous statement where convex sets are replaced by more general families of sets in a metric space includes many on-line/off-line problems such as the k -server problem; we make some remarks on these more general problems.

1. Introduction

Consider a fixed metric space, (S, ρ) , and a family of subsets of S , \mathcal{F} . A *chasing problem instance* consists of a point $p_0 \in S$ and a sequence F_1, \dots, F_n of elements of \mathcal{F} . A *solution* to the instance is a sequence p_1, \dots, p_n of points of S such that $p_i \in F_i$; for such a solution we define its *cost* to be

$$\sum_{i=1}^n \rho(p_{i-1}, p_i).$$

* Joel Friedman wishes to acknowledge the National Science Foundation for supporting this research in part under a PYI Grant, CCR-8858788, and IBM for an IBM faculty development award. Nathan Linial's work was supported by a grant from the Binational Science Foundation Israel/US and from the Israel Academy of Sciences.

The problem at hand is to find a solution whose cost is as small as possible. As usual, this problem has an *off-line* version, where we know the F_i in advance, and an *on-line version*, where the F_i are given one at a time and p_i must be chosen before knowing F_{i+1} ; we seek to find a *competitive* on-line algorithm, i.e., one for which the cost is never more than a fixed constant times the cost of any (off-line) solution. A family \mathcal{F} is said to be *chaseable* if there exists an on-line algorithm competitive with the off-line algorithm.

We wish to study what families are chaseable, and what geometric properties guarantee that a family is chaseable or not. At this level of generality these questions are probably difficult, and contain many on-line/off-line questions (as in [1]–[11]).

For example, this problem contains the k -server problem of [9]. More generally, we can form a k -server version of the set-chasing problem for $k > 1$, but clearly this is again a set-chasing problem for a family of subsets in the k th cartesian product of the original metric space. In fact, one motivation for the set-chasing problem is to put the chaseability of families such as those arising from k -server problems into a simple geometric framework.

From the geometric point of view, it seems natural to first consider set chasing in \mathbf{R}^d . The main goal of this paper is to prove that the collection of convex sets in \mathbf{R}^2 is chaseable. We more generally pose:

Conjecture 1.1. *For any d , the family of closed convex sets, in the metric space \mathbf{R}^d , is chaseable.*

Question 1.2. *For which metric spaces is it true that the family of closed convex sets is chaseable? Same question for the family of unions of $\leq n$ closed convex sets, with n fixed.*

In the above, by a convex set in a metric space we mean a subset T which for any $x, y \in T$ contains all points z with $\rho(x, y) = \rho(x, z) + \rho(z, y)$. We remark that the condition of the second part of Question 1.2 also generalizes the k -server condition. We cannot really hope that the family of closed convex sets in every metric space is chaseable, for this would imply that there is a universal competitiveness ratio for all metric spaces (by gluing collections of metric spaces with bad ratios together); for example, the competitiveness ratio of convex set chasing in \mathbf{R}^n cannot be better than \sqrt{n} on the problem instance $p_0 = 0$, and $F_i = \{x_i = \pm 1\}$ (with \pm chosen according to which is further from the on-line player). Hence convex sets in \mathbf{R}^∞ are not chaseable. One can ask for geometric properties on a metric space which imply chaseability of convex sets, such as a Helly-type property, etc. For example, the analysis of the “move-to-front” rule for maintaining a linear list in [11] shows that the family of convex sets in the symmetric group, S_n (with metric given by the number of transpositions), is chaseable, in fact with competitiveness ratio 1 using the greedy algorithm. More generally geometric constraints on the family of subsets may be written down so that the greedy

algorithm is competitive,¹ but they do not give very general conditions. Question 1.2 is even interesting with “convex bodies” replaced by points; this is equivalent to the “layered graph traversal” problem of [10], for which Fiat *et al.* [6] have recently given upper and lower bounds exponential in n for the competitiveness ratio for on-line algorithms.

Our approach to chasing convex sets in \mathbf{R}^2 is first to notice that it suffices to be able to chase half-planes. Instead of chasing half-planes, we begin with the easier problem of chasing lines. The line-chasing problem already points out failures in two natural algorithms, and our half-plane-chasing algorithm, in some sense, builds on the ideas used for line chasing. It might seem that half-plane chasing is no harder than line chasing, at first; indeed, we can always assume that $p_i \notin F_{i+1}$ for the on-line player, and certainly there is no advantage in moving to the interior of F_{i+1} , so the on-line player is essentially always chasing lines, i.e., boundaries of half-planes. However, the off-line player need not move to this boundary when he lies in the requested half-plane, and so his cost could be substantially less than in the “corresponding” line-chasing problem.

Henceforth the “angle” of a line in the plane means the measure of the angle it makes with the x -axis.

So we turn to the problem of line chasing. We begin by describing two natural algorithms which fail. We are given an initial point $p \in \mathbf{R}^2$ and a sequence of lines, l_1, \dots, l_N . Consider the following algorithm, which we call the greedy algorithm: to satisfy the request l_i , we move to the point on l_i closest to the current position. To see that this is not competitive, fix a large number $B > 0$. Consider the point $p = (1, 0)$, and for a small $\varepsilon > 0$ take $N = \lfloor B/\varepsilon \rfloor$ and let l_i be the line which passes through the origin of angle εi . As $\varepsilon \rightarrow 0$ the limiting path taken by the greedy algorithm is to follow the circle of radius 1 about the origin, for a total cost of B . For a sufficiently small ε the greedy path’s cost comes arbitrarily close to B . On the other hand, moving to the origin and staying there has cost 1. Since B can be taken to be arbitrarily large, the greedy algorithm is not competitive.

We remark that the same thing happens if we take p as before, l_i for odd i to be the line through the origin of angle ε , and l_i for i even to be the x -axis. This situation is the previous one modified by taking mirror images each time.

Another natural algorithm, which would take care of the above situation, would be the “move-to-optimal” strategy: each request l_i is satisfied by computing where the off-line optimal for the requests $(p_0; l_1, \dots, l_i)$ would be at the end (i.e., after the l_i request), and to move to this location. To see that this strategy is not competitive, consider the following requests: $p_0 = (0, 1)$; l_1 is the x -axis; for $i \geq 2$, l_i is the line intersecting

$$Q_i = l_{i-1} \cap \{(x, y) | x = (-1)^i\}$$

¹ For example, say that a family, \mathcal{F} , of subsets is *linear* if for all $F \in \mathcal{F}$ and $p \notin F$ there is a closest point to p in F , say $f = \text{Greedy}(F, p)$ with the property that $\rho(f', p) \geq \rho(f, f') + \rho(f, p)$ for all $f' \in F$. Then any linear family is chaseable via the greedy algorithm with competitiveness ratio 1 (i.e., is optimal). The same holds if $\rho(f, p)$ is replaced by $\theta \rho(f, p)$ for any fixed $\theta > 0$, or if we allow $\text{Greedy}(F, p)$ to be any point within a fixed constant factor of p ’s distance to F (and the “greedy” algorithm moves to $\text{Greedy}(F, p)$) with different competitiveness ratios.

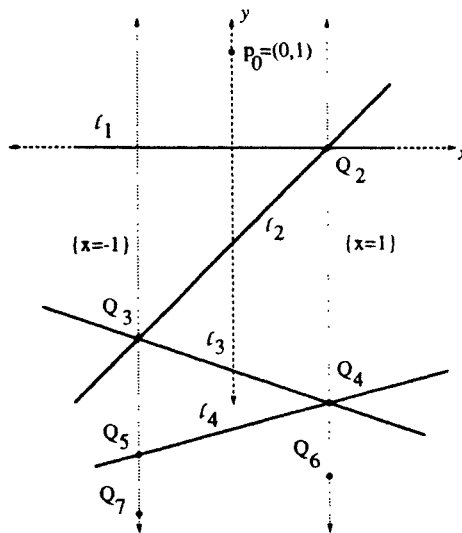


Fig. 1. "Move to optimal" is not competitive.

and which is perpendicular to the line from p_0 to Q_i (see Fig. 1). It is easy to see that if θ_i is the angle of l_i , then

$$\theta_2 = 45^\circ > \theta_4 > \theta_6 > \cdots > 0 > \cdots > \theta_5 > \theta_3 > -45^\circ$$

and that $|\theta_i| \rightarrow 0$ as $i \rightarrow \infty$. Furthermore, the move-to-optimal strategy is to satisfy l_i by moving to Q_i for $i \geq 2$, and so the cost of this strategy is more than $1/\beta$ times the optimal off-line strategy of moving down the y -axis, for any $\beta \geq \limsup |\theta_i|$. Since $\theta_i \rightarrow 0$, the move-to-optimal strategy is not competitive.

We conclude this section by describing in rough terms our algorithm. For certain periods of time we follow a locally greedy strategy. If the off-line player is near us, then his cost will be close to ours. Occasionally we decide that a far away off-line player could be gaining a lot in terms of cost, an amount that is proportional to his distance. We then declare the "round" to be over and make an updating move, i.e., a move in between requests, to move closer to the region where the off-line player could be gaining on us. The crucial point is to decide when to update, and then to guarantee that our updating move will ensure that the following locally greedy moves will not be too costly, no matter where the off-line player is.

A question that these strategies lead to is can a "simple" strategy for these chasing problems be given, one that involves no updating move and that perhaps involves only local calculations of the optimal off-line costs at nearby points? If so this might lead to much simpler analysis for certain algorithms such as ours for convex body chasing in \mathbb{R}^2 .

As a matter of technical convenience, it is often simpler to deal with continuous versions of discrete problems. Here we describe our algorithms in terms of

continuous or piecewise-continuous versions of the problem (with a very restricted set of discontinuities, so as not to include the discrete problem trivially!). For another example where the continuous version is simpler (for somewhat different reasons) see [3].

In Section 2 we give a simple algorithm and analysis for line chasing in the plane, and give some variants of the algorithm which are also competitive. In Section 3 we solve the half-plane-chasing problem in the plane. In Section 4 we make some general remarks about set-chasing problems, and in particular explain that convex body chasing in the plane follows from Section 3.

2. Line Chasing

In this section we discuss the problem of line chasing. For this problem we give a simple algorithm and analysis, and the techniques used here are built upon for the half-plane-chasing algorithm.

2.1. Continuous Version

Consider the following continuous version of line chasing: we are given an initial point $p_0 \in \mathbf{R}^2$, and a family of lines in \mathbf{R}^2 , l_t , where $t \in [0, T]$ for some T . In addition, $p_0 \in l_0$, and the lines vary continuously and piecewise differentiable in t ; by the latter we mean that we can write the lines as

$$l_t = \{(x, y) | a(t)x + b(t)y + c(t) = 0\}$$

with a, b, c continuous and piecewise differentiable functions² of t with a, b never vanishing simultaneously. A solution to the problem is a Lipschitz continuous path³ $p: [0, T] \rightarrow \mathbf{R}^2$ with $p(t) \in l_t$ for all t . The cost of p is its length, i.e.,

$$\int_0^T |p'(t)| dt = \text{TV}(p),$$

where TV denotes the total variation. More generally, we allow p to be piecewise Lipschitz continuous, in which case the cost is interpreted as the total variation of p , i.e.,

$$\text{TV}(p) = \sum_{s \in D} |p(s+0) - p(s-0)| + \int_0^T |p'(t)| dt,$$

² For simplicity the reader can assume that the functions are piecewise linear; this suffices to deal with the reductions from the discrete problem.

³ For the piecewise-linear version, it suffices to define a solution as a piecewise-linear path, allowing a discrete set of discontinuities or jumps (for ease of discussion). The cost is the length of the path, i.e., sum of the lengths of the linear parts plus the lengths of the jumps.

where D is the set of discontinuities of p . We also refer to these discontinuities as “jumps.”

To avoid ambiguity, we refer to the original chasing problem as the *discrete* version of the problem.

Lemma 2.1. *Line chasing in the plane is reducible to continuous line chasing; i.e., given a line-chasing problem, the on-line player can form a continuous line-chasing instance given on-line whose optimal cost is, within a multiplicative factor of $\sqrt{2}$, that of the discrete problem.*

Proof. Consider an instance of discrete line chasing, p, l_1, l_2, \dots, l_N . For simplicity we assume that the problem instance contains an initial line l_0 which contains p ; l_0 may be chosen arbitrarily by the on-line player. We may assume successive lines l_i, l_{i+1} are distinct. For each $i = 0, 1, \dots, N - 1$ we define l_t for $t \in (i, i + 1)$ in the following natural way: if l_i and l_{i+1} are parallel, l_t sweeps through the parallel lines between l_i and l_{i+1} at a constant speed; if l_i and l_{i+1} intersect in a point, l_t pivots through the lines through I between l_i and l_{i+1} at a constant angular speed, between meaning through the smaller angle (and either angle if they are perpendicular).

Clearly a solution $p: [0, T] \rightarrow \mathbb{R}^2$ of the continuous problem gives a solution, by restriction to $\{1, 2, \dots, N\}$, to the discrete problem whose (discrete) cost is no more than the continuous cost. On the other hand, given a solution to the discrete problem, p_1, \dots, p_T , we form a solution to the continuous problem in the natural way: if the line segment from p_i to p_{i+1} lies in the region swept out by l_t with $t \in [i, i + 1]$, we take $p(t)$ in that range to follow that segment; if not, then we move along l_i until we intersect l_{i+1} (which defines $p(t)$ for $t < i + 1$) and then “jump” to p_{i+1} (yielding a discontinuity). The continuous cost of the jumping steps is easily seen to be no more than $\sqrt{2}$ times the cost of the discrete cost, and the continuous cost of the nonjumping steps is the same as the discrete cost. Hence a discrete solution yields a continuous solution whose total cost is no more than $\sqrt{2}$ times the discrete cost. \square

A useful simplification in the continuous line-chasing problem is to assume that the lines rotate counterclockwise with increasing t . Given a problem instance, l_t , with t either continuous or discrete, fix an $s \in (0, T)$, let R_s be the reflection of the plane through l_s , and consider the new problem instance given by lines \tilde{l}_t defined by

$$\tilde{l}_t = \begin{cases} l_t & \text{for } t \in [0, s], \\ R_s(l_t) & \text{for } t \in (s, T]. \end{cases}$$

We call the new problem instance the *reflection of l_t at s* ; it is merely the old problem with a reflection introduced at time s . A solution to the old problem gives a solution to the new one (or the same cost), and vice versa, by reflecting everything after time s (as in the above equation for \tilde{l}). This gives an equivalence of the two problems.

Given a continuous line-chasing problem l_t , let $\theta(t)$ denote the angle of l_t , chosen in a way to make it continuous in t . If $\theta(t)$ is piecewise monotone in t on $[0, T]$, i.e., if it “changes direction” a finite number of times on $[0, T]$, then applying a finite number of reflections as above we get an equivalent problem for which the new $\theta(t)$ is a monotone increasing function in t . In particular this can be applied to the continuous problems derived from discrete problems.

Although not needed for discrete line chasing, it is true that any continuous line-chasing problem is equivalent to a problem whose θ is monotone increasing, by taking a limit of problems which are reflections of piecewise differentiable approximations to the original problem. A limit of reflections of the original problem may also be taken directly, as follows. Let $S_1 \subset S_2 \subset \dots$ be an increasing sequence of finite subsets of $[0, T]$ with the property that for all i the number of points in S_{i+1} between any two consecutive points of S_i is even (consecutive with respect to $<$ on \mathbf{R}). Then it is easy to see that the limit of reflecting any problem by the S_i exists. We can clearly take S_i to contain enough points so that the total variation in $\theta(t)$ is no more than $1/i$ between consecutive points. The limit of this problem has θ monotone increasing.

For the algorithm described below, we assume that θ is monotone increasing and that its derivative exists and is nonzero (i.e., positive) everywhere. This allows us to reparametrize t so as to assume $\theta'(t) = 1$ for all t . Rotating coordinates we can further assume that $\theta(0) = 0$, so that $\theta(t) = t$ for all t , i.e., l_t makes an angle of t with the x -axis. We call a problem with such θ an instance of *angular line chasing*. The argument that solving such a problem is sufficiently general to solve all continuous line chasing can be handled by approximation or by modifying the algorithm, and is given in the Section 2.4.

Finally, notice that for the discrete line-chasing problem we can assume that l_i and l_{i+1} are never parallel; if T is the translation taking l_{i+1} to l_i , then the new problem, $l_1, \dots, l_i, T(l_{i+2}), T(l_{i+3}), \dots$, with an added cost to all players of the distance from l_i to l_{i+1} , is the same as the original problem (and charging an extra positive cost to everyone does not worsen the competitiveness ratio). A similar remark holds for the continuous problem over any interval where $\theta'(t) = 0$. Hence the reader interested only in the discrete problem and piecewise-linear version can automatically assume θ' is strictly positive and can skip the aforementioned approximation argument.

2.2. Greedy Coordinates

At this point we assume that the sequence of requests is a continuous family of lines, l_θ , such that at time θ the angle the line makes with the x -axis is θ . In addition, we assume that the instantaneous point of intersection

$$I(\theta) = \lim_{\theta_1, \theta_2 \rightarrow \theta} l_{\theta_1} \cap l_{\theta_2}$$

exists and is a finite point for each value of θ . Again, this assumption is justified later.

For each θ , the points in l_θ can be parameterized by a single coordinate, $x = x_\theta$, which is a function from l_θ to \mathbf{R} ; we write $x(\theta; t)$ for the point on l_θ with x_θ coordinate t . There are two natural systems of coordinates to choose, one in which the $I(\theta)$ is the origin, the other being the parallel transport coordinates. The latter give a simple description of a locally greedy strategy and are constructed as follows.

Definition 2.2. *Parallel transport coordinates with respect to l_θ are coordinates x_θ for which the trajectory $x(\theta; t)$ for fixed t , as a function of θ , is perpendicular to l_θ for all θ .*

Lemma 2.3. *For Lipschitz continuous l_θ , parallel transport coordinates exist and are uniquely determined (given an initial coordinate x_0 on l_0). Such coordinates are distance preserving.*

Proof. For any point $Q \in l_0$, there is a unique curve $q(\theta)$ such that $q(0) = Q$, $q(\theta) \in l_\theta$, and $q'(\theta)$ is perpendicular to l_θ for all θ ; indeed, clearly q is determined by the ordinary differential equation

$$q'(\theta) = f(\theta, q(\theta)), \quad q(0) = Q,$$

where f measures the speed and direction that the line l_θ moves away from $q(\theta)$ (f can be viewed as a function on $[0, T] \times \mathbf{R}^2$ by (pulling back f from the) orthogonal projection onto l_θ). We call $q(t)$ the parallel transport of $q(0)$ to l_t . If R is another point of l_0 and $r(\theta)$ is the corresponding curve as above, then

$$(|q(\theta) - r(\theta)|^2)' = 2(q - r, q' - r') = 0.$$

Hence parallel transport is an isometry. So choose a coordinate $x_0: l_0 \rightarrow \mathbf{R}$ which preserves distance. For $q \in l_\theta$, let $x_\theta(q) = x_0(\tilde{q})$ where \tilde{q} is the inverse parallel transport of q back to l_0 . These are the desired coordinates. \square

These coordinates are also called *greedy* coordinates for the reason that keeping $x(\theta; p(\theta))$ constant means following the locally greedy strategy, i.e., locally moving to the nearest point.

Now let $m(\theta)$ denote $x(\theta; I(\theta))$, i.e., the greedy coordinate of $I(\theta)$. A solution $p(\theta)$ is determined by its greedy coordinate $f(\theta) \equiv x(\theta; p(\theta))$, and the cost incurred by such an f is the total variation of p , which is just

$$\mathcal{C}_0^T(f) = \int_0^T \sqrt{(f(\theta) - m(\theta))^2 + (f'(\theta))^2} d\theta;$$

the integrand contains the parallel and perpendicular to l_θ components of p' . In the above equation we understand that f is piecewise Lipschitz continuous, and if f has a discontinuity, then the jump in f is accounted for by the f' in the above equation (integrated from $s - 0$ to $s + 0$ at discontinuities s); in other words, we add $|f(s + 0) - f(s - 0)|$ at all discontinuities.

We claim that the locally greedy strategy of keeping f constant is essentially optimal over fixed amounts of time (angular motion). More precisely, fix an $S \in [\Omega, T - 2\pi]$ and let

$$\mathcal{O}(\alpha, \beta) = \inf\{\mathcal{C}_S^{S+2\pi}(f) \mid f(S) = \alpha, f(S + 2\pi) = \beta\}, \quad (2.1)$$

where \mathcal{C}_a^b is the cost from time a to time b (i.e., the integral in (2.1) from a to b), and the infimum is taken over all piecewise Lipschitz f . Let

$$\mathcal{O}_* = \inf\{\mathcal{O}(\alpha, \beta) \mid \alpha, \beta \in \mathbf{R}\}.$$

Let

$$\tilde{\mathcal{O}}(\alpha) = \mathcal{C}_S^{S+2\pi}(\alpha),$$

which is the cost of the locally greedy algorithm which stays at α , and let α_* be a minimizer of $\tilde{\mathcal{O}}$. Notice that

$$\tilde{\mathcal{O}}(\alpha) = \int_S^{S+2\pi} |m(\theta) - \alpha| d\theta,$$

and so the set of all α_* 's is simply the set of all medians of m , i.e., α_* 's satisfying

$$\int_{\{\theta \mid m(\theta) \leq \alpha^*\}} d\theta = \int_{\{\theta \mid m(\theta) \geq \alpha^*\}} d\theta;$$

if $I(\theta)$ is, for example, continuous, then the median is unique.

Lemma 2.4. *For $\varepsilon = 1/(1 + 2\pi)$ the following hold:*

1. *For any α, β we have $\mathcal{O}(\alpha, \beta) \geq |\beta - \alpha|$.*
2. *For any α, β we have $\mathcal{O}(\alpha, \beta) \geq \varepsilon \tilde{\mathcal{O}}(\alpha)$.*
3. *$\mathcal{O}_* \geq \varepsilon \tilde{\mathcal{O}}(\alpha_*)$.*
4. *For any α, β we have $\mathcal{O}(\alpha, \beta) \geq \varepsilon \tilde{\mathcal{O}}(\beta)$.*
5. *For any α, α' we have $\tilde{\mathcal{O}}(\alpha) \leq \tilde{\mathcal{O}}(\alpha') + 2\pi|\alpha - \alpha'|$ and*

$$\tilde{\mathcal{O}}(\alpha) \geq 2\pi|\alpha - \alpha'| - \tilde{\mathcal{O}}(\alpha').$$

6. *For any α we have $\tilde{\mathcal{O}}(\alpha) \geq \pi|\alpha - \alpha_*|$.*
7. *For any α, β we have $\mathcal{O}(\alpha, \beta) \geq (\pi\varepsilon/2)(|\alpha - \alpha_*| + |\beta - \alpha_*|)$.*

Proof. For convenience we omit the sub- and superscripts S and $S + 2\pi$ from \mathcal{C} . Part (7) follows from (2), (4), and (6). Part (3) follows from (2). Part (1) follows from the observation that

$$\mathcal{C}(f) \geq \int |f'(\theta)| d\theta = \text{TV}(f),$$

where TV denotes the total variation. This observation and the observation that

$$\mathcal{C}(f) \geq \int |m(\theta) - f(\theta)| d\theta$$

imply (2) and (4)—for example, to prove (2) note that if $\mathcal{C}(f)$ is less than $\varepsilon \tilde{\mathcal{C}}(\alpha) = \varepsilon \mathcal{C}(\alpha)$, then

$$\int |(m - \alpha) - (f - \alpha)| = \int |m - f| \leq \varepsilon \int |m - \alpha|$$

and so

$$\int |f - \alpha| \geq (1 - \varepsilon) \int |m - \alpha|.$$

It follows that $|f - \alpha|$ attains the value $(1 - \varepsilon)M$ somewhere where M is the average value of $|m - \alpha|$, and hence

$$\mathcal{C}(f) \geq \text{TV}(f) \geq (1 - \varepsilon)M = \frac{1 - \varepsilon}{2\pi} \int |m - \alpha| = \frac{1 - \varepsilon}{2\pi} \mathcal{C}(\alpha) = \varepsilon \mathcal{C}(\alpha).$$

Part (4) follows similarly. Part (5) follows from integrating the fact that

$$|m - \alpha| + |m - \alpha'| \geq |\alpha - \alpha'| \geq |m - \alpha| - |m - \alpha'|.$$

Finally, (6) follows from adding the equations

$$\tilde{\mathcal{C}}(\alpha) \geq 2\pi|\alpha - \alpha_*| - \tilde{\mathcal{C}}(\alpha_*)$$

(obtained from (5)) and

$$\tilde{\mathcal{C}}(\alpha) \geq \tilde{\mathcal{C}}(\alpha_*).$$

□

2.3. Algorithm and Analysis

We now describe the on-line strategy. From time $\theta = 0$ until 2π we take the locally greedy solution, i.e., $f \equiv 0$, and then at $\theta = 2\pi$ we compute an α_* of the interval $[0, 2\pi]$ and jump to the point with that $x(2\pi)$ coordinate. We call this the first round. Each successive round lasts 2π in duration, and our strategy is to follow the locally greedy strategy and then jump to the α_* for that round.

We wish to explain and analyze our algorithm in a style which will help to understand our half-plane-chasing algorithm. If at some point the off-line player, F , stops moving, and the l_i start rotating around F , then the α_* 's for all later

rounds are precisely at F 's position; furthermore, any competitive player must eventually move closer and closer to F . In general, whenever there is a “cost-free” region, the on-line player, N , must move toward it, at least when this region persists. In particular, we introduce a potential function which incorporates N 's distance to F ; let $c_2 > c_1 > 1$ be two constants to be specified later, and set

$$\Phi(t; N, F) = c_2 C_F(t) - c_1 |F(t) - N(t)| - C_N(t),$$

where $C_N(t)$, $C_F(t)$ and $N(t)$, $F(t)$ denote the total cost and position of N , F up to time t . Our analysis shows that this potential function is *nondecreasing over rounds*, i.e., is at least as large at the end of any round than at its beginning. This then implies that $C_N \leq c_2 C_F$ at the end of any round. The point is that if F does no work during a round, then N 's jump at the end of the round decreases the distance to F and hence increases the potential function; furthermore, F cannot gratuitously lower the potential function by, say, jumping away from N after N 's jump, for the resulting increase in distance is compensated by the increase in C_F , provided that $c_2 \geq c_1$.

Theorem 2.5. *The above on-line strategy for line chasing is competitive.*

Proof. It suffices to show that for appropriate constants $c_1, c_2 > 0$ the above potential function never decreases from start to end of a round; indeed, if so then, since $\Phi(0) = 0$, at the end of any round $\Phi \geq 0$ and so $C_N \leq c_2 C_F$.

Let F_0, F_1 be F 's position at the start and end of the round, and similarly for N . Using Lemma 2.4 we see that the increase in C_F over the round is at least

$$\mathcal{O}(F_0, F_1) \geq \max\left(\frac{(|F_0 - \alpha_*| + |F_1 - \alpha_*|)\pi\epsilon}{2}, \mathcal{O}_*\right). \quad (2.2)$$

The change in C_N is the cost of the locally greedy path plus the jump, namely,

$$\tilde{\mathcal{O}}(N_0) + |N_0 - \alpha_*|$$

since $N_1 = \alpha_*$, which is

$$\leq \tilde{\mathcal{O}}(\alpha_*) + (2\pi + 1)|N_0 - \alpha_*| \leq K(\mathcal{O}_* + |N_0 - \alpha_*|) \quad (2.3)$$

for some constant K , by (5) and (3) of the lemma. Finally, the distance between N and F decreases by at least

$$|N_0 - F_0| - |\alpha_* - F_1| \geq |N_0 - \alpha_*| - |F_0 - \alpha_*| - |F_1 - \alpha_*|. \quad (2.4)$$

Using the right-hand sides of (2.2)–(2.4), we see that if $c_1 \geq K$ and

$$c_2 \geq 2 \max(K, c_1/\epsilon),$$

then the potential function's value at the end of the round is at least that at the beginning. \square

A more careful analysis (considering separately the cases where \varnothing_* is \leq or \geq the other argument of \max in (2.2)) shows that the constant yielded above for c_2 is $9 + 5\pi + (2/\pi)$. A similar analysis where the rounds last for L radians instead of 2π gives a constant of $9 + 5L + (4/L)$, for a minimum of $9 + 4\sqrt{5}$. For discrete problems reduced to continuous ones, we include the $\sqrt{2}$ factor from the reduction for a total competitiveness ratio of $9\sqrt{2} + 5\sqrt{10} = 28.53\dots$

2.4. Variants and the General Algorithm

We begin by describing a few modifications of the algorithm, more precisely of the updating step, that also give competitive results. More generally, we want to know if there is a "simple local strategy" with no updating step, e.g., some combination of local optimal and local greedy. Here we show that the previous algorithm is robust in that many variants of it are also competitive; perhaps these could lead to a simple strategy with no updates.

As a first variant notice that instead of moving to a median we can move to an approximate median, i.e., a point where the mass of m on either side is at least some fraction of the total mass. More precisely, we can fix a constant $\beta \in (0, \frac{1}{2})$ and move to any point $\hat{\alpha}$ such that

$$\beta \leq \frac{\int_{\{\theta | m(\theta) \leq \hat{\alpha}\}} d\theta}{2\pi} \leq 1 - \beta.$$

It is easy to see that

$$\tilde{\varnothing}_* \geq \varepsilon |\alpha_* - \hat{\alpha}|$$

for some ε depending on β ; it follows that Lemma 2.4 holds as well for α_* replaced by $\hat{\alpha}$ (only with different constants which depend on β). Hence this algorithm is also competitive.

The next modification we describe is that of lazily following the optimal. By this we mean that we pick a $\gamma \in (0, 1)$, and instead of moving to the median we move only a fraction γ of the distance to the median. Of course, Lemma 2.4 does not hold with α_* replaced by the lazy step, but the analysis of the potential function, Φ , goes through as before. The only real modification is that in the left-hand side of (2.4), i.e., the distance change estimate, the $|\alpha_* - F_1|$ term, is replaced by

$$|N_1 - F_1| \leq |N_1 - \alpha_*| + |\alpha_* - F_1| = (1 - \gamma)|N_0 - \alpha_*| + |\alpha_* - F_1|;$$

hence the change in distance is

$$\geq \gamma |N_0 - \alpha_*| - |F_0 - \alpha_*| - |F_1 - \alpha_*|.$$

In the lazy algorithm V_F 's cost is the same and C_N 's cost is no greater than in the original algorithm, and so the analysis goes through as before provided that $c_1 \geq C/\gamma$.

In addition to these two variants, we can form a hybrid algorithm which lazily moves to an approximate median. Clearly this also yields a competitive algorithm, combining Lemma 2.4 modified for the approximate median with the analysis of the potential function modified for the lazy step.

We finish by remarking that the algorithm given in the previous subsection, and the modifications of it given above, work for problem instances where the angle is monotone increasing but not necessarily strictly increasing (or with positive time derivative), and where it may not be the case that $I(\theta)$ exists for all θ . Indeed, for the more general instance we can still define the locally greedy coordinates and therefore algorithm, and define a median of m, α_* , in terms of these coordinates. Such a strategy gives rise to a solution, $N(t)$. We give an easy approximation argument to show that $N(t)$ is competitive with the same competitiveness ratio.

Namely, for any ε we can form a perturbed problem instance, l_t^ε , which satisfies the assumptions of the previous section and which differs from l_t by an ε -approximation; by this we mean that for all t , $l_t^\varepsilon = A^\varepsilon(t)l_t$ where A^ε is a family of rigid motions which, as 2×2 matrices, differ in each coordinate from the identity matrix followed by the zero translation by functions of less than ε norm in $H^{1,\infty}$. If the α_* 's are always unique, then $N(t)$ is the limit of the locally greedy strategies for l_t^ε , $N_\varepsilon(t)$ with $\varepsilon \rightarrow 0$; if not, then $N(t)$ is the limit of some N_β^ε solution given by a β approximate median solution with $\beta = \beta(\varepsilon) \rightarrow \frac{1}{2}$ as $\varepsilon \rightarrow 0$. In either case, it is easy to see that the cost of the limit is the limit of the costs; on the other hand, any off-line strategy $F(t)$ for l_t has a perturbed strategy, namely $A^\varepsilon(t)F(t)$, for l_t^ε , having the cost of the limit being the limit of the costs. Therefore the competitive ratio of the on-line algorithm is preserved in the more general situation.

3. Half-Plane Chasing

3.1. Preliminaries for Half-Plane Chasing

Our strategy for half-plane chasing, like that for line chasing, is locally greedy followed by an update at the end of each "round." What makes half-plane chasing harder than line chasing is the fact that, for line chasing, if during the entire round F is answering requests "for free," i.e., without even moving, then F 's position is limited to one point; for half-plane chasing the "free" region can be much larger (even unbounded). Consequently, the updating move in our half-plane-chasing algorithms is more complicated.

Before we describe the algorithm, we describe the "free" regions which arise, and for each we give an updating move designed to reduce our distance to any point in the region. We take time coordinates so that the round begins and ends at times $t = 0, 1$; we use $N(1), N(1+)$ respectively to distinguish between N 's position before and after the update.

In the first situation the cost-free region, R , is a half-plane intersected with a disk. Specifically, set up coordinates normalized by the condition that $N(0) = (0, 1)$ and H_1 is the half-plane of points with x -coordinate ≥ 1 . Our algorithm ensures that $N(1)$ is not far from $U = (1, 0)$, and our updating move is to U . We (i.e., the algorithm) discover that the cost-free region, R , is contained in a disk of radius C about the origin for some bounded C . So let

$$R_C = \{(x, y) \in \mathbf{R}^2 \mid x \geq 1, x^2 + y^2 \leq C\}$$

and let

$$\mathcal{U}_\delta = \{P \in \mathbf{R}^2 \mid |P - U| \leq |P - N(0)| - \delta\}. \quad (3.1)$$

This region is bounded by one half of a hyperbola, which as $\delta \rightarrow 0$ tends to the line $x = \frac{1}{2}$ (see Fig. 2).

Lemma 3.1. *For any $C > 0$ there is a $\delta > 0$ such that $R_C \subset \mathcal{U}_\delta$. There is also a $\mu > 0$ such that the same holds with R_C replaced by a translation and rotation of R_C of, respectively, distance and angle $\leq \mu$.*

The latter statement is to give us leeway when we make approximations in our statements.

Proof. A simple calculation. □

In the other situation, R will be a wedge of angle bounded away from 180° and whose vertex is “not far below the origin.” Specifically, take coordinates so that

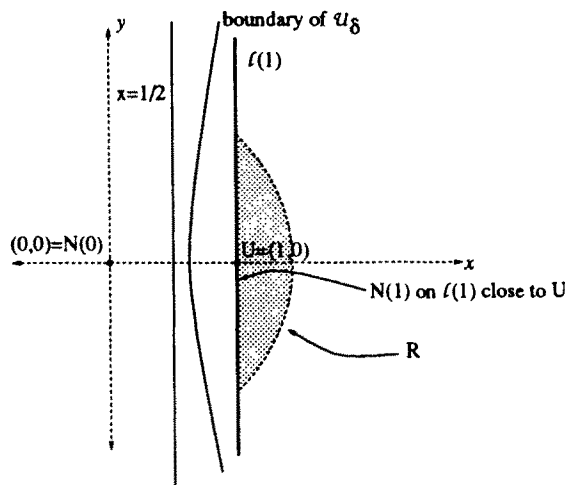
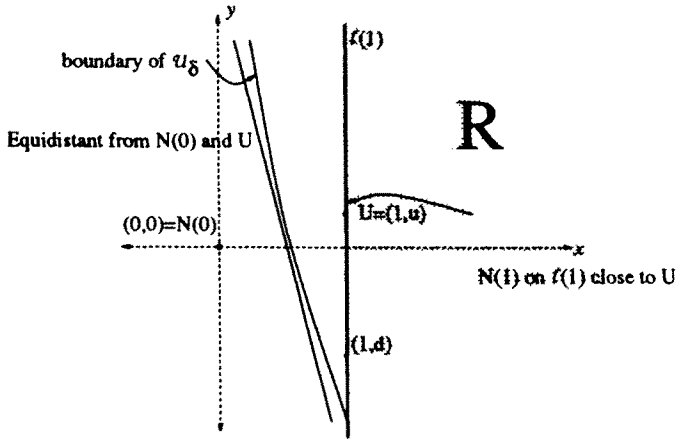


Fig. 2. \mathcal{U}_δ containing R_C .

Fig. 3. $\mathcal{U}_\delta(u)$ containing $R_{d,\varepsilon}$.

$N(0)$ and H_1 are as before. R will be a wedge of the form

$$R_{d,\varepsilon} = \{(x, y) \in \mathbf{R}^2 \mid x \geq 1, y - d \geq -\varepsilon(x - 1)\}$$

(see Fig. 3). Note that $R_{d,\varepsilon} \subset R_{d',\varepsilon'}$ if $\varepsilon' \leq \varepsilon$ and $d' \leq d$. By the vertex being “not far below the origin” we mean that we have a lower bound for d . In this situation our updating move will be to move to a point $U = (1, u)$ with u positive and small. Let $\mathcal{U}_\delta = \mathcal{U}_\delta(u)$ be as in (3.1).

Lemma 3.2. *For any $\varepsilon > 0$ and d there are $u, \delta > 0$ such that $R_{d,\varepsilon} \subset \mathcal{U}_\delta(u)$. There is also a $\mu > 0$ such that the same holds with $R_{d,\varepsilon}$ replaced by a translation and rotation of $R_{d,\varepsilon}$ of, respectively, distance and angle $\leq \mu$.*

Proof. Again, an easy calculation. □

In what follows, distance almost always means distance with respect to the original metric; when confusion can occur between this distance and distance in terms of the coordinates just described (e.g., $N(0) = (0, 0)$), we use the adjectives *original* and *coordinate* to distinguish the two. In fact, the original distance will always be w times the coordinate distance, where w is a quantity to be defined later which is always proportional to N 's work during the round.

3.2. Restricted Half-Plane Chasing

In this section we analyze a restricted version of continuous half-plane chasing, which illustrates many points of the general algorithm. The requests are a family

of half-planes, H_t for $t \in \mathbb{R}$, $t \geq 0$, given by

$$H_t = \{(x, y) | a(t)x + b(t)y + c(t) \geq 0\},$$

with a , b , c Lipschitz continuous functions of t with a , b never vanishing simultaneously. We also assume that the half-planes always move away from the on-line player.

It is interesting to note that, unlike line chasing, insisting that the half-planes are Lipschitz continuous already makes the problem too restrictive to subsume discrete plane chasing clearly, unlike the situation in Lemma 2.1—for example, it is not clear how to replace the discrete requests $H_i = \{x \geq 1\}$ and $H_{i+1} = \{x \leq 1\}$ by a sequence of continuous ones “interpolating” the two (e.g., if we rotate clockwise, that will destroy the cost ratio of paths which move in the counter-clockwise direction above the pivot point).

Let l_t denote the boundary of H_t , and let $\omega(t)$ be the angle of l_t . For a round beginning at time t_0 , we look at the first time, t_1 , that ω differs from $\omega(t_0)$ by $\alpha = 30^\circ$,⁴ and declare the round over then. For the duration of the round we follow the boundary, l_t , by the locally greedy algorithm described in Section 2. We then make an updating move described below. Essentially, we see if $\omega(t)$ remains roughly the same for most of the round, where time is normalized according to N 's cost. If so we call the round *concentrated*; in this case the “for free” region can be limited to a wedge of an angle bounded away from 180° . In the other case we call the round *spread out*, and the “for free” region can be limited to a bounded portion of a half-plane. Our updating move is taken according to whether the round is concentrated or spread out.

Before describing the updating move, notice that if $W = N(t_1) - N(t_0)$ is the displacement of N following the locally greedy strategy, and if $w = |W|$, then N 's cost for the locally greedy motion is no more than $w/\cos \alpha$; indeed, $N'(t)$'s component in the direction perpendicular to l_{t_0} is at least $|N'(t)| \cos \alpha$, and so W 's component in this direction is at least $\cos \alpha$ times $\int |N'(t)| = N$'s cost for the motion. N 's updating step is also bounded by $O(w)$, and hence N 's total cost for the round as well.

From now on we assume, for simplicity, that $t_0 = 0$ and $t_1 = 1$. We normalize time with respect to N 's cost, so that $C'_N(t) = |N'(t)|$ is constant for $t \in (0, 1)$ (this constant is between $w \cos \alpha$ and w). Aside from this we keep the same notation as in the previous section.

We say the round (or work) is *concentrated* if there is an ω_0 for which $\omega(t) \in [\omega_0, \omega_0 + 4^\circ]$ for at least half of the time. If not we say the work is *spread out*, and the updating move consists of moving to the image of the orthogonal projection of $N(0)$ onto l_1 .

If the work is concentrated, set up coordinates as in the last subsection, i.e., $N(0) = (0, 0)$ and l_1 is the line $x = 1$. For some $\tau \in [0, 1]$ we must have that

⁴ We sometimes use α in what follows, even though it remains 30° in this subsection, since it will be modified in the next section and we wish to point out which equations are modified accordingly.

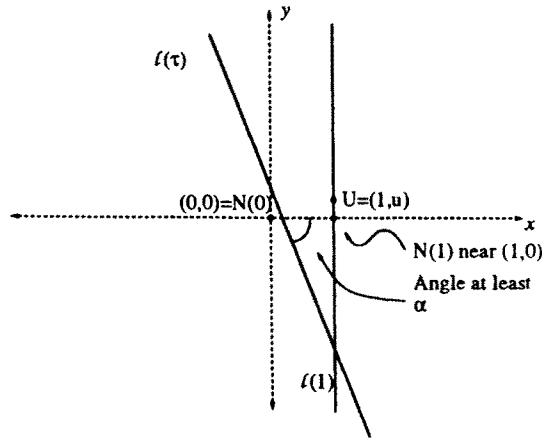


Fig. 4. Determination of U .

$|\omega(\tau)| \geq \alpha$; fix one such τ , for which we can assume that $\omega(\tau)$ is negative. Our updating move is to move to the point $U = (1, u)$, with u a small positive constant to be determined later. See Fig. 4.

Theorem 3.3. *The above “locally greedy with update” algorithm is competitive, for an appropriate choice of u .*

Proof. As in line chasing, we introduce the same potential function, Φ , and it suffices to show that, for appropriate $c_1, c_2 > 0$, Φ ’s value at the end of each round is at least that at the beginning.

Henceforth we let C_F, C_N denote the cost to F and N of the entire round. As mentioned before, N ’s updating move never costs more than $O(w)$ —when the work is spread out, then the distance of $N(1)$ to the orthogonal projection of $N(0)$ onto any line through $N(1)$ is no greater than w . When the work is concentrated, $N(1) = (1, \eta)$ in the above coordinates, with $|\eta| \leq \sin \alpha = \frac{1}{2}$, and so the coordinate metric is proportional to w times the original metric; the updating step costs $|\eta - u|$ in these coordinates, which is $O(w)$ assuming that, say $|u| \leq 1$. Under this assumption, then, $C_N \leq O(w)$.

For the analysis we can make certain simplifying assumptions. First, if C_F during the round is $\geq \mu w$ for a fixed constant $\mu > 0$, then we are done, provided that

$$c_2 \geq \frac{c_1 + 1}{\mu} + c_1. \quad (3.2)$$

Indeed, setting $d(t) = |N(t) - F(t)|$ and $\Delta d = d(1) - d(0)$, then since

$$|\Delta d| \leq C_N + C_F$$

we have

$$c_2 C_F - c_1 \Delta d - C_N \geq C_N(\mu(c_2 - c_1) - (c_1 + 1)).$$

Hence we may assume that $C_F \leq \mu w$; in particular the poistion of F throughout the round remains in a ball of radius μw bout, say, $F(0)$.

Secondly, consider the angle, $\theta(t)$, which the line $F(t)N(t)$ makes with l_t , measured between 0 and 90° . We say that N's work has *small* θ if $\theta(t) \in [0, 1^\circ]$ for at least two-thirds of the time. We claim that we are done if this is not so, provided that

$$c_2 \geq c_1 \geq \frac{K}{\cos \alpha} \quad (3.3)$$

for some constant K . Indeed, when $\theta(t) > 0$, it is necessarily the case that $F(t)$ lies in the interior of H_t and is therefore not moving, and also $d'(t) = -|N'(t)|\sin \theta(t)$ (see Fig. 5). So N is gaining in his distance to F, "paying" a factor of $1/\sin \theta$ for this gain. When $\theta = 0$ and both N and F are moving, the distance can increase, but only at a cost to F; more precisely $d'(t) \leq C_F(t)$, since greedy coordinates are an isometry and so $d'(t)$ is precisely \pm the parallel to l_t component of $F'(t)$. So if N does not have small θ , then

$$\Delta d \leq -\frac{1}{3}|N'|\sin 1^\circ + C_F,$$

$|N'|$ being constant and proportional to w (by a factor between 1 and $\cos \alpha$). Hence the potential function does not decrease if (3.3) holds.

We now finish the analysis under the assumptions that $C_F \leq \mu w$ and N has small θ .

First consider the case of spread-out work. We claim that $d(0)$ (and $d(1)$ for that matter) are bounded by wC where C is an absolute constant assuming that that $\mu < 1$. Indeed, there exist $s_1, s_2 \in [0, 1]$ such that $\omega(s_1)$ and $\omega(s_2)$ differ by at

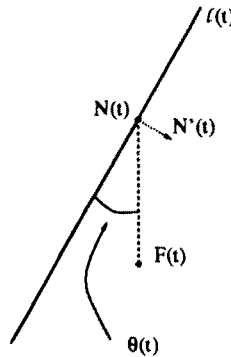


Fig. 5. N moves closer for $\theta > 0$.

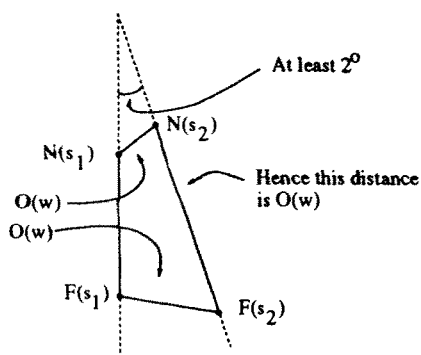


Fig. 6. Spread-out work implies F close to N .

least 4° while both $\theta(s_i)$'s are $\leq 1^\circ$ since θ is small more than two-thirds of the time and yet ω stays in any interval of 4° for no more than half of the time. This gives the picture in Fig. 6, of the quadrilateral with vertices $F(s_i)$, $N(s_i)$, with $|F(s_1) - F(s_2)|$ and $|N(s_1) - N(s_2)|$ of distance bounded by $O(w)$, and with the difference in angle of the opposite sides $F(s_i)N(s_i)$, $i = 1, 2$, being at least 2° ; this clearly implies $|F(s_i) - N(s_i)|$ is bounded by $O(w)$.

At this point we can apply Lemma 3.1. In terms of the coordinates set there we know that $N(1) = (1, \eta)$ with $|\eta| \leq \sin \alpha = \frac{1}{2}$ and the unit length of the coordinates is proportional to w . The lemma implies that

$$\Delta d \leq |N(1+) - F(1)| - |N(0) - F(1)| + |F(1) - F(0)| \leq -\delta_1 w + \mu w$$

for some constant δ_1 depending only on μ . Hence the potential function is nondecreasing if

$$\delta_1 > \mu \quad \text{and} \quad c_1 \geq \frac{1}{\delta_1 - \mu}. \quad (3.4)$$

In the case of concentrated work, we apply Lemma 3.2. Again, in terms of the coordinates there, $N(1) = (1, \eta)$ with $|\eta| \leq \sin \alpha = \frac{1}{2}$ and the coordinate unit is proportional to w . The line $l(\tau)$, τ as before, passes through $N(\tau) = (x_\tau, y_\tau)$ with $|x_\tau| \leq 1$ and $|y_\tau| \leq \frac{1}{2}$. Therefore l_τ intersects the line l_1 at the point $(1, \beta)$ with β bounded by $O(1/\sin \alpha)$. F 's position at all times lies within $O(\mu w)$ of the wedge $R = H_\tau \cap H_1$, and so Lemma 3.2 implies the existence of an appropriate u and δ_2 for which we know the potential function is nondecreasing if

$$\delta_2 > \mu \quad \text{and} \quad c_1 \geq \frac{1}{\delta_2 - \mu}. \quad (3.5)$$

Notice that u and δ_2 depend only on μ .

We now conclude by choosing the constants. Begin by choosing a small positive value for μ and calculate the two δ_i 's of (3.4) and (3.5); the latter δ and u depend on μ . Now choose another μ smaller than both δ_i 's; the same δ_i 's and u work for the smaller μ . Then choose c_1 and then c_2 so that (3.5), (3.4), (3.2), and (3.3) hold. For any such c_1, c_2 the resulting potential function is nondecreasing, and so our algorithm is competitive with competitiveness ratio c_2 . \square

3.3. Piecewise-Continuous Half-Plane Chasing

In this section we describe how to modify the previous algorithm in the case of "piecewise-continuous" half-plane chasing. A half-plane chasing problem is *piecewise continuous* if it is a continuous half-plane-chasing problem except that we allow a discrete set of discontinuities, $S \subset T$. Obviously this problem includes the discrete problem, namely when H_t is piecewise constant in t .

We remark that, to find a competitive algorithm for this problem, it suffices to do so under the restrictions that

- (i) the half-planes always move away from N ,
- (ii) for all $t \in [0, T]$, $N(t) \in l_{t \pm 0}$, and
- (iii) at the discontinuities s , H_{s-0} and H_{s+0} intersect in a wedge of angle $\leq \nu$ for some small, fixed, positive ν , i.e., that $H_{s \pm 0}$ are almost opposite.

Indeed, consider a general instance where (i)–(iii) do not necessarily hold. At any time when (i) ceases to hold we ignore all further requests until N must move. Next, assuming that (i) holds, at any point, s , of discontinuity where (ii) does not hold, we translate H_{s+0} so that l_{s+0} contains $N(s)$, and then insert an interval of time between s and $s + 0$ where H_t moves continuously from H_s to H_{s+0} . Finally, assuming that (i) and (ii) hold, at all remaining points of discontinuity, s , if H_{s-0} and H_{s+0} are not almost opposite in the above sense, we can insert an interval of time where H_t rotates from H_{s-0} to H_{s+0} through an angle of measure $\leq 180^\circ - \nu$; this change may introduce additional violations of (i), which we correct as before. The changes we have made so that (i) and (ii) hold introduce no extra cost; the changes to ensure (iii) introduce a multiplicative factor in the cost of at most $2 \sin((\pi - \nu)/2) / \sin \nu$ —indeed, if H_{s-0} and H_{s+0} are not almost opposite and a player (namely F) on $l_{s-0} - H_{s+0}$ wants to move to a point on $l_{s+0} - H_{s-0}$, he is forced to move along two sides of a triangle making an angle ν instead of traversing the third side; the sum of the two sides divided by the third is easily seen to be at most $2 \sin((\pi - \nu)/2) / \sin \nu$, i.e., in the case of an isosceles triangle. As at the end of Section 2.1, we have described how to handle a finite number of violations of the desired conditions; the general case follows either by an approximation argument or by a limiting argument similar to the ones given before.

We may now assume that the problem instance satisfies (i)–(iii), with $\nu > 0$ to be fixed later. A round beginning at time t_0 ends at time t_1 which is the smallest $t > t_0$ when either of the following happen: (1) $|\omega(t) - \omega(t_0)|$ becomes $\geq \alpha = 10^\circ$, or (2) the amount of "backward motion" becomes at least ψ times the "forward motion," where $\psi < 1$ is a constant to be specified later. More precisely, a

half-plane H_t with $t \in [t_0, t_1]$ is called *forward (backward) moving* if its intersection with H_{t_0+0} is a wedge of measure $>90^\circ$ ($<90^\circ$); accordingly, at points where H_t is continuous, N or $N(t)$ is said to be *moving forward (backward)*. Furthermore, measure the amount of forward and backward motion in terms of time, normalized according to N 's cost (so that $|N'(t)|$ is constant). Of course, when a round ends because of case (1), there are two qualitatively different possibilities, namely H_{t_1} can be forward or backward.

We now describe the updating steps. As before, normalize time so that the round begins and ends at $t = 0, 1$ (and, of course, $|N'(t)|$ constant). We always set coordinates so that $N(0) = (0, 0)$ and l_1 is the line $x = 1$; here, of course, we have the added possibility that $H_1 = \{x \leq 1\}$, i.e., that H_1 could be backward. As before, the coordinate unit will always be proportional to $w = |N(0) - N(1)|$.

First take the case that the round ends with H_1 forward, and therefore $\omega(1)$ differs from $\omega(0)$ by something between α and $\alpha + v$. Since the amount of backward motion is small, the angle of $W = N(1) - N(0)$ is less than α' in absolute value, with $\alpha' = \alpha + v + O(\psi)$ for ψ small. We now fix α', v, ψ by choosing any positive values of v, ψ so that $\alpha' = 15^\circ$ and $\psi \leq \frac{1}{2}$. We define "concentrated round," "spread-out round," and "small θ " with respect to time, as before, and the case analysis goes through the same. Again, we first argue that we can assume $C_F(1) \leq \mu w$ for a small constant μ and that θ is small for two-thirds of the time. When the work is spread out we use Lemma 3.1; when the work is concentrated, we use Lemma 3.2. The only essential difference here is that α is replaced by α' .

The second case is when the round ends in case (1) but with H_1 backward. Set coordinates with $N(0) = (0, 0)$ and l_1 being the line $x = 1$. Let $t \in [0, 1]$ be any value such that $N(t)$ and $N(0)$ lie on the opposite sides of l_1 and such that N is moving forward at time t ; we can take, for example, t to be the last time that l_t is moving forward (see Fig. 7⁵). Without loss of generality we can assume that the

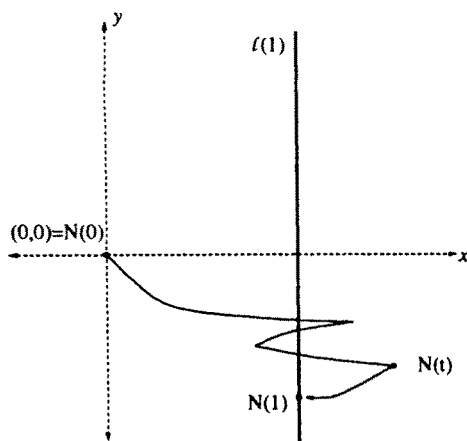


Fig. 7. t is last forward moving time.

⁵ The figures in this subsection will be drawn as if α and v were a lot bigger than they actually are (for ease of illustration).

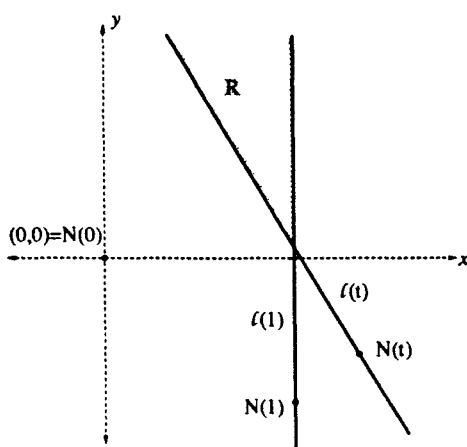


Fig. 8. F is restricted to the wedge R .

angle of l_0 is $< 90^\circ$ (i.e., $\in [90 - \alpha - \nu, 90 - \alpha]$) or equivalently that the angle of l_t is $> 90^\circ$. Hence throughout the round F must remain within $O(w\mu)$ in the region $R = H_t \cap H_1$ (see Fig. 8), which is a wedge of angle $\leq 2\alpha + \nu$ and whose vertex is above a bounded distance “below perpendicular,” in the sense of Lemma 3.2, where l_t here plays the role of l_1 in the lemma. Hence we can apply Lemma 3.2, with l_1 replaced by l_t , and make the updating step indicated there, i.e., move to the point on l_t which is slightly higher than the image of the projection of $N(0)$.

In the last case, that of backward work being proportional to forward work, we limit the “free” region as follows. Set coordinates as before. Let s_0 be a point where

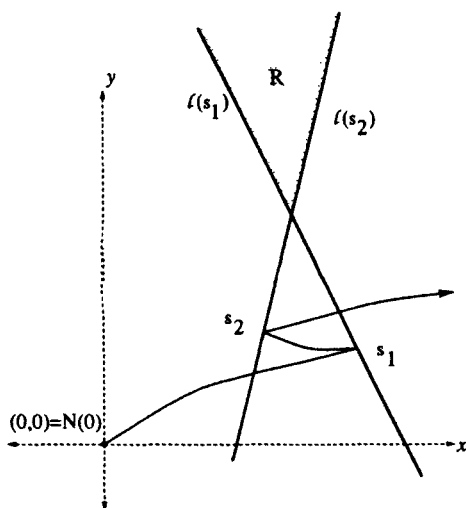


Fig. 9. A wedge which points up.

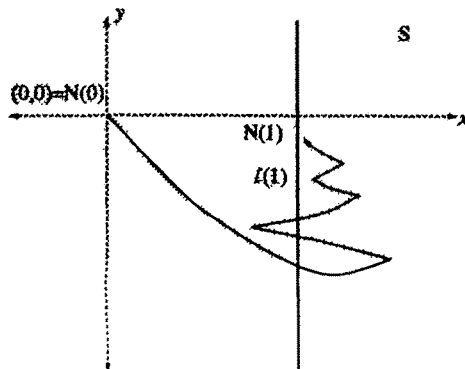


Fig. 10. F not in S incurs too much cost.

H changes direction from forward to backward, and let s_1 be the next point where H changes from backward to forward. The intersection of H_{s_0-0} and H_{s_1-0} is a wedge of angle $\leq \nu + \alpha$ that either “points up” or “points down” in the sense that it contains points either of arbitrarily large positive or negative y -coordinate (see Fig. 9). If, say, the wedge points up, and if $F(s_0)$ has the y -coordinate smaller than both those of $N(s_0)$ and $N(s_1)$, then it is easy to see that F must incur a cost at least proportional to $C_N(s_1) - C_N(s_0)$, and similarly for wedges that point down. Without loss of generality we can assume that at least half of the time (during all the backward motion) the wedges point up (see Fig. 10). Then $C_F \leq \mu C_N$, with μ small implies that F must lie in the region, S , of points with the y -coordinate bounded below (by a function of μ and α'). Intersecting this with l_1 gives a wedge of angle $= 90^\circ$ whose vertex is bounded below with respect to l_1 . Hence we can apply Lemma 3.2 and take the appropriate updating step.

The final step of selecting μ , the constants in the updating step, and those for the potential function is as before.

4. General Remarks and Consequences

4.1. Intersections and Convex Bodies

As in the introduction, let \mathcal{F} be a family of subsets of S in a metric space (S, ρ) . Let $\mathcal{I}(\mathcal{F})$ denote the family of finite intersections of \mathcal{F} and let $\bar{\mathcal{F}}$ denote the set of limit points of decreasing sequences of $\mathcal{I}(\mathcal{F})$ in the Hausdorff metric; by this we mean that ρ is extended to subsets of S via

$$\rho(A, B) = \sup_{a \in A} \inf_{b \in B} \rho(a, b) + \sup_{b \in B} \inf_{a \in A} \rho(a, b),$$

and by a limit point of decreasing sequences of a family \mathcal{G} we mean any closed

$G \subset S$ such that $\rho(G, G_i) \rightarrow 0$ as $i \rightarrow \infty$ for some $G_i \in \mathcal{G}$ with $G_1 \supseteq G_2 \supseteq \dots$. Under certain conditions, if \mathcal{F} is chaseable, then automatically $\mathcal{J}(\mathcal{F})$ will be as well.

Definition 4.1. \mathcal{F} is said to have *nice intersections* if for any positive integer n , any $F_1, F_2, \dots, F_n \in \mathcal{F}$, and for any $\delta > 0$ there is an $\varepsilon > 0$ such that

$$\rho(x, F_1 \cap \dots \cap F_n) \geq \delta \Rightarrow \max(\rho(x, F_1), \dots, \rho(x, F_n)) \geq \varepsilon.$$

Proposition 4.2. *If \mathcal{F} is chaseable with multiplicative constant K , then, for every $\mu > 0$, \mathcal{F} is chaseable with multiplicative constant $K + \mu$, if either (1) the underlying metric space is complete, or (2) \mathcal{F} has nice intersections.*

Proof. Consider a problem instance, $p_0 \in S$, and requests \bar{F}_1, \dots from \mathcal{F} . We can assume that $m = \rho(p_0, \bar{F}_1)$ is positive; this gives us a lower bound for the total cost of the problem. By definition we can find $\tilde{F}_1 = F_1^1 \cap \dots \cap F_1^{n_1}$ for some n_1 with $F_1^i \in \mathcal{F}$, $\tilde{F}_1 \supseteq \bar{F}_1$, and $\rho(\bar{F}_1, \tilde{F}_1) \leq m\mu/8$. We start by running the \mathcal{F} -chasing algorithm on the initial point p_0 and the request sequence $F_1^1, \dots, F_1^{n_1}, F_1^1, \dots, F_1^{n_1}, F_1^1, \dots$, stopping the first time we reach a point p_1 which is within $m\mu/8$ of \tilde{F}_1 ; we claim this must occur in a finite amount of time, either in case (1) or (2). Indeed, the competitiveness of the \mathcal{F} -chasing algorithm implies that the solution it generates, $q_0 = p_0, q_1, q_2, \dots$, has finite cost

$$\text{Cost}(\{q_i\}) = \sum_{i=1}^{\infty} |q_i - q_{i-1}|,$$

since all requests can be satisfied by moving to \tilde{F}_1 . In case (1), for any positive integer k , if q_{kn_1} is of distance $\geq \delta$ from \tilde{F}_1 , the cost of the next round of $F_1^1, \dots, F_1^{n_1}$ costs at least ε , for any δ and ε as in Definition 4.1. Hence, for sufficiently large k all the q_{kn_1} must lie within δ of \tilde{F}_1 for any $\delta > 0$, in particular for $\delta = m\mu/8$. In case (2), since the q_i are a Cauchy sequence, their limit q exists; clearly q lies in \tilde{F}_1 , and so p_1 occurs no later than when one q_i lies within $m\mu/8$ of \tilde{F}_1 .

At this point, from p_1 we are free to jump to the nearest point in \bar{F}_1 and back to p_1 , for a cost of $m\mu/2$. We continue the algorithm by approximating \bar{F}_2 to within $m\mu/16$, and continuing the \mathcal{F} -chasing algorithm from p_1 on the F_2^j repeatedly. Continuing in this way, we reach points p_2, p_3, \dots , such that the total costs of jumping from p_i to \bar{F}_i and back is $\leq m\mu$. Let M be the cost of the algorithm which does not jump back and forth. The optimal (off-line) solution to the \bar{F}_i problem, with cost C_* , trivially yields a solution to the \mathcal{F} -chasing problem constructed above, with the same cost. Hence $M \leq KC_*$. On the other hand, the algorithm which jumps from p_i to \bar{F}_i and back yields a solution to the \bar{F}_i problem of cost

$$\leq M + m\mu \leq (K + \mu)C_*,$$

since m is a lower bound for C_* .

Corollary 4.3. *If the family of affine half-spaces in \mathbf{R}^n is chaseable, then so is the family of convex bodies in \mathbf{R}^n .*

Corollary 4.4. *The family of convex bodies in the plane is chaseable.*

4.2. Plane Chasing in \mathbf{R}^3 , Lazy Line Chasing in \mathbf{R}^2 , and Function Chasing

At the time of writing we do not know whether or not convex bodies in \mathbf{R}^3 are chaseable. However, we define a “lazy set-chasing” problem and show that chasing planes in \mathbf{R}^3 is equivalent to the problem of lazy line chasing in \mathbf{R}^2 .

The problem of lazy set chasing differs from the set-chasing problem in that a positive $\varepsilon \leq 1$ is given as part of the input, and it is not required at time i to move to F_i . Instead, the cost of a solution p_1, \dots, p_n (here a solution is any collection of points in S) is

$$\sum_{i=1}^n \rho(p_{i-1}, p_i) + \varepsilon \rho(p_i, F_i),$$

where the second ρ is the minimal distance of a point in F_i to p_i .

For a fixed ε , or more generally if ε is restricted to a range bounded away from zero, the lazy set-chasing problem is equivalent, to within a factor of $1 + (2/\varepsilon)$, to the set-chasing problem. Indeed, if $p_i \notin F_i$ the cost of jumping to F_i and back to p_i is no more than $2/\varepsilon$ of the cost of that move. So if we require p_i to be in F_i , we can do so by adding two jumps per request which gives a total cost of at most $1 + (2/\varepsilon)$ of the original cost.

On the other hand, if there is no *a priori* bound on how near ε can be to zero, then the problem appears to be harder. To give precise examples, it becomes more convenient to allow ε to depend on i . Henceforth we take this as our definition of a lazy set-chasing problem.

Now we show that plane chasing in \mathbf{R}^3 is equivalent to lazy line chasing in \mathbf{R}^2 . For a continuous instance of plane chasing in \mathbf{R}^3 , locally greedy (= parallel transport) coordinates can be set up for the planes, and the cost function becomes

$$\mathcal{C}(f) = \int_0^T \sqrt{\rho(f(\theta), l(\theta))^2 + |f'(\theta)|^2} d\theta;$$

where θ is the arc length of the unit normal to the plane, ρ is the extension of $||$ to sets, $l(\theta)$ is the line of instantaneous intersection of the planes, and f takes values in \mathbf{R}^2 . To within a factor of $\sqrt{2}$ this is equivalent to

$$\tilde{\mathcal{C}}(f) = \max \left(\int_0^T \rho(f(\theta), l(\theta)) d\theta, \text{TV}(f) \right). \quad (4.1)$$

Given an instance of continuous plane chasing, which is just to give $l(\theta)$, we can

approximate l by a step function and therefore, by (4.1), get an instance of lazy line chasing. Conversely, given an instance of lazy line chasing we can form the corresponding step function l , whose value for θ between $\varepsilon_1 + \dots + \varepsilon_i$ and $\varepsilon_1 + \dots + \varepsilon_{i+1}$ is l_i , it is easy to see that we can form a plane-chasing problem whose instantaneous intersection is precisely l , and therefore reduce this problem to a plane-chasing problem.

Theorem 4.5. *Lazy line chasing in the \mathbf{R}^2 is equivalent to plane chasing in \mathbf{R}^3 (to within fixed proportionality constants).*

The argument given in the preceding paragraph is a rough description of the proof. To make the above rigorous, as well as to understand better lazy set-chasing problems, we introduce a type of “generalized set” or “function” chasing problem. To motivate these problems, consider an instance of a lazy set-chasing problem with the $\varepsilon_i = \varepsilon$ all equal and $B = 1/\varepsilon$ an integer. It is easy to see that any solution p_1, p_2, \dots is, within a constant factor, no better than the solution p'_1, p'_2, \dots , which only changes every B moves, i.e.,

$$p'_0 = p'_1 = \dots = p'_{B-1} = p_0, \quad p'_B = \dots = p'_{2B-1} = p_B, \dots$$

Denoting p_{B_i} by q_i , the total cost of such a solution is just

$$\sum_{i=1}^n \rho(q_{i-1}, q_i) + g_i(q_{i-1}),$$

where g_i is the function

$$g_i(q) = \frac{1}{B} \sum_{j=B(i-1)+1}^{B_i} \rho(q, F_j).$$

So given a metric space, (S, ρ) , we consider a family \mathcal{G} of functions from S to $\mathbf{R}_{\geq 0}$. We define a *function-chasing* problem to be an initial point, p_0 , and a sequence of functions, g_1, \dots , with $g_i \in \mathcal{G}$, and with the cost of a solution $\{p_i\}$ being

$$\sum_{i=1}^n \rho(p_{i-1}, p_i) + g_i(p_i);$$

accordingly we can speak of \mathcal{G} being chaseable or not.

Definition 4.6. Given a family of functions, \mathcal{G} , from S to \mathbf{R} , we define the hull of \mathcal{G} , $\text{Hull}(\mathcal{G})$, to be the family of all finite convex linear combinations of elements of \mathcal{G} . For any $K > 1$ we define the K approximations to \mathcal{G} , \mathcal{G}_K , to be the family of functions g' for which there is a $g \in \mathcal{G}$ such that $g(s)/K \leq g'(s) \leq Kg(s)$ for all $s \in S$. Given a family of subsets, \mathcal{F} , of S , we define the family of functions associated with \mathcal{F} , \mathcal{F} , to be the family of functions

$$\{g(s) = \rho(s, F) \mid F \in \mathcal{F}\}.$$

Theorem 4.7. *For any $K \geq 1$, a family \mathcal{G} is chaseable iff \mathcal{G}_K is chaseable. The lazy \mathcal{F} -chasing problem is equivalent to the $\text{Hull}(\mathcal{F})$ function-chasing problem.*

Proof. The first statement is obvious. The second statement follows essentially from the previous discussion. Namely, given a set-chasing problem with a sequence ε_i , we can easily reduce to the case that $\varepsilon_i \leq 1$. Letting I_1 be the smallest i with $\varepsilon_1 + \cdots + \varepsilon_i \geq 1$, we replace the first I_1 requests with the function

$$g_1(s) = \sum_{i=1}^{I_1} \varepsilon_i \rho(s, F_i),$$

which lies in $\text{Hull}(\mathcal{F})_2$. If no such I_1 exists, then it is easy to see that it is competitive to stay at p_0 for all time. We continue by defining I_2 as the smallest i with $\varepsilon_{I_1+1} + \cdots + \varepsilon_i \geq 1$, and forming g_2 . Continuing in this way we get a $\text{Hull}(\mathcal{F})_2$ -chasing problem equivalent to the lazy problem. Reducing a $\text{Hull}(\mathcal{F})$ function-chasing problem to a lazy \mathcal{F} problem is similar. \square

At this point it is clear how to prove rigorously that lazy line chasing in \mathbf{R}^2 and plane chasing in \mathbf{R}^3 are both equivalent to chasing the hull of the family of functions associated with lines in the plane.

We finish this subsection with a question on chasing functions. It might be hoped that competitive function-chasing algorithms might be simpler than set-chasing algorithms, in that some sort of “simple local strategy” might work for them. We define a generalized convex set to be a function, $g: S \rightarrow \mathbf{R}_{\geq 0}$, which is Lipschitz continuous with constant 1, i.e.,

$$|g(s) - g(t)| \leq \rho(s, t)$$

for all s, t , and which is convex, i.e.,

$$g(s_2) \leq \frac{\rho(s_2, s_3)}{\rho(s_1, s_3)} g(s_1) + \frac{\rho(s_2, s_1)}{\rho(s_1, s_3)} g(s_3)$$

for all s_1, s_2, s_3 with $\rho(s_1, s_2) + \rho(s_2, s_3) = \rho(s_1, s_3)$. Note that, for example, in \mathbf{R}^n , the family of generalized convex sets is larger than the hull of the family of functions associated with convex sets; the latter have certain growth conditions at ∞ with many symmetries.

Question 4.8. Is the family of all generalized convex sets in \mathbf{R}^n chaseable? In what metric spaces is the family of all generalized convex sets chaseable? If not, what are enough conditions to give chaseability for a lazy problem or family of functions?

The Lipschitz condition for generalized convex sets is certainly necessary; for example, consider $\mathbf{R}_{\geq 0}$ with initial point $p_0 = 0$ and function $f_1(x)$ which is 1 at $x = 0$ and x elsewhere.

The problem of chasing the family of all nonnegative functions is precisely the metrical task systems of [3]; there upper and lower bounds proportional to n are given for the competitiveness ratio of on-line algorithms, where $n = |S|$. We suggest that if restrictions such as convexity and Lipschitz continuity are placed on the functions, then competitive algorithms can exist even when $|S|$ is infinite.

4.3. Line Chasing in \mathbf{R}^n

Although we do not know whether or not convex bodies in \mathbf{R}^3 can be chased, we can say that lines in \mathbf{R}^3 can be competitively chased. Indeed, we can easily reduce a line-chasing problem in three dimensions to one in two dimensions. Namely, first we reduce chasing a sequence of lines, $\{l_i\}$, to one in which l_i and l_{i+1} are coplanar for all i —for any skew lines l_1 and l_2 we let P be the plane parallel to l_2 containing l_1 , and let l' be the projection of l_2 onto P . Inserting l' as a request between l_1 and l_2 does not change the cost by more than a factor of two, since for any $p_i \in l_i$, $i = 1, 2$, if p' is the projection of p_2 onto P , then both $|p_1 - p'|$ and $|p' - p_2|$ are clearly $\leq |p_1 - p_2|$. Now given a problem instance in which l_i, l_{i+1} are always coplanar, we can clearly define lines $\{m_i\}$ all lying in \mathbf{R}^2 , isometries π_i from l_i to m_i , inductively on i , which preserves the distance from points in l_i to points in l_{i+1} . This gives the desired reduction.

More generally the same technique shows that lines can be competitively chased in \mathbf{R}^n for any n . In general a slightly weaker claim can be made. Namely, let \mathcal{F} be any family of subsets of \mathbf{R}^d such that we can competitively chase the subsets, \mathcal{F}_{2d+1} , of all rigid motions of \mathcal{F} in \mathbf{R}^{2d+1} . Then it easily follows that we can chase \mathcal{F}_n in \mathbf{R}^n for all $n \geq 2d + 1$. (For line chasing we can replace $2d + 1$ by $2d$ (here $d = 1$), which is why the more general claim is weaker.)

It follows from the observations here and in the first subsection that if we are returning tennis balls at the net projected by a machine, assuming negligible gravity, then we can efficiently move to return them. Here the requests are half-lines in \mathbf{R}^3 , being that we are not allowed to cross the net and, in particular, hold the racket at the barrel of the machine (not a practice recommended by the authors). This application, under the guise of the racket game “matkot,” was pointed out to the authors by A. Wigderson.

Acknowledgment

Nati Linial wishes to thank Mike Saks for conversations which were critical for asking the correct questions.

References

1. S. Ben-David, A. Borodin, R. M. Karp, G. Tardos, and A. Wigderson. On the Power of Randomization in On-line Algorithms. In *Proc. of the 22nd Ann. ACM Symp. on Theory of Computing*, pages 379–386, May 1990.

2. P. Berman, H. J. Karloff, and G. Tardos. A Competitive Three-Server Algorithm. In *Proc. of the 1st Ann. ACM-SIAM Symp. on Discrete Algorithms*, pages 280–290, January 1990.
3. A. Borodin, N. Linial, and M. Saks. An Optimal On-Line Algorithm for Metrical Task Systems. In *Proc. of the 19th Ann. ACM Symp. on Theory of Computing*, pages 373–382, May 1987. Final version to appear in *J. Assoc. Comput. Mach.*
4. M. Chrobak, H. J. Karloff, T. Payne, and S. Vishwanathan. New Results on Server Problems. In *Proc. of the 1st Ann. ACM-SIAM Symp. on Discrete Algorithms*, pages 291–300, January 1990.
5. D. Coppersmith, P. Doyle, P. Raghavan, and M. Snir. Random Walks on Weighted Graphs and Applications to On-line Algorithms. In *Proc. of the 22nd Ann. ACM Symp. on Theory of Computing*, pages 369–378, May 1990.
6. A. Fiat, D. P. Foster, H. Karloff, Y. Rabani, Y. Ravid, and S. Vishwanathan. Competitive Algorithms for Layered Graph Traversal. Manuscript, 1991.
7. A. Fiat, Y. Rabani, and Y. Ravid. Competitive k -Server Algorithms. In *Proc. of the 31st Ann. IEEE Symp. on Foundations of Computer Science*, pages 454–463, October 1990.
8. E. Grove. The Harmonic Online k -Server Algorithm Is Competitive. Manuscript, 1990.
9. M. S. Manasse, L. A. McGeoch, and D. D. Sleator. Competitive Algorithms for On-Line Problems. In *Proc. of the 20th Ann. ACM Symp. on Theory of Computing*, pages 322–333, May 1988.
10. C. H. Papadimitriou and M. Yannakakis. Shortest Paths Without a Map. In *Proc. 16th ICALP*, pages 610–620, July 1989.
11. D. D. Sleator and R. E. Tarjan. Amortized Efficiency of List Update and Paging Rules. *Comm. ACM*, 28:202–208, 1985.

Received June 10, 1991, and in revised form August 4, 1992.