Motion Planning for a Steering-Constrained Point Robot through Moderate Obstacles

PANKAJ K. AGARWAL* PRABHAKAR RAGHAVAN† HISAO TAMAKI‡

Abstract

Most mobile robots use steering mechanisms to guide their motion. Such mechanisms have stops that constrain the rate at which the robot can change its direction. We study a point robot in the plane subject to such constraints, in the presence of a class of obstacles we call *moderate obstacles*. We consider the case in which the robot has a reverse gear that allows it to back up, as well as the case when it does not. All our algorithms run in polynomial time and produce paths whose lengths are either optimal or within an *additive* constant of optimal.

1 Introduction

The *motion-planning* (or *path planning*) problem involves planning a collision-free path for a robot moving amid obstacles. This is one of the main problems in robotics, and it has been widely studied (see, e.g., the book by Latombe [22] and the survey paper by Schwartz and Sharir [36]). In the simplest form of the motion planning, given a moving robot B, a set of obstacles O, and a pair of placements I and F of B; we wish to find a continuous, collision-free path for B from I to F. This problem is known to be PSPACE-complete [5, 33], and efficient algorithms have been developed for several special cases [36]. Most of these algorithms, however, do not take into account the kinodynamic constraints (for instance, velocity/acceleration bounds, curvature bounds), the so called *nonholonomic constraints*, of a real robot imposed by its physical limitations. Although there has been considerable recent work in the robotics literature (see [1, 2, 3, 13, 17, 19, 20, 21, 24, 25, 26, 28, 39, 40] and references therein) on nonholonomic motion planning problems, relatively little theoretical work has been done on these important problems, because they are considerably harder than the holonomic motion-planning problems.

*Computer Science Department, Duke University, Box 90129, Durham, NC 27708-0129. A portion of this work was done while the author was visiting the IBM T.J. Watson Center. His work is supported by National Science Foundation Grant CCR-93–01259, an NYI award, and by matching funds from Xerox Corp.

†IBM T.J. Watson Research Center, Yorktown Heights, NY 10598. A portion of this work was done while the author was visiting the IBM Tokyo Research Laboratory.

‡IBM Tokyo Research Laboratory, 1623-14 Shimotsuruma, Yamato-shi, Kanagawa 242 Japan.

In holonomic motion planning, the placement of a robot with k degrees of freedom is determined by a tuple of k (typically real) parameters, each describing one degree of freedom. The set of all placements is called the *configuration space*, and the set of placements at which the robot does not intersect any obstacles is called the *free configuration space*. There exists a path between an initial placement and a final placement if and only if the two placements lie in the same (path-)connected component of the free configuration space. This is not necessarily true if the robot has to obey nonholonomic constraints. In nonholonomic motion planning, a placement is not enough to describe the robot. Instead, a robot is completely described by its state, which consists of the k parameters and their derivatives (see [21] for a more detailed discussion), which makes the problem significantly harder.

In this paper, we study the path planning problem for a point robot B whose path is constrained to have curvature at most 1. B may or may not be equipped with a reverse gear. If B is equipped with a reverse gear, then it can change its orientation instantaneously, so we require that path has an average curvature at most 1 in every interval except where it reverses its orientation. The curvature constraint corresponds naturally to constraints imposed by stops on the steering mechanism found in virtually every real-world robot; see e.g., [22]. We present efficient algorithms for computing optimal, or near-optimal, collision-free paths for B in presence of moderate obstacles. Roughly speaking an obstacle is moderate if it is convex and the curvature of its boundary is bounded by 1.

Dubins [11] considered this problem in the absence of obstacles, with the further restriction that the robot is not equipped with a *reverse gear*, so that it can never back up. His characterization shows that the shortest path from any start position to any final position consists of at most 3 segments, each of which is either a straight line or an arc of a unit-radius circle. Reeds and Shepp [32] extended this obstacle-free characterization to robots that are allowed to make reversals. (Boissonnat *et al.* [4] gave an alternative proof for both cases, using ideas from control theory.) For robots that must contend with polygonal obstacles, Fortune and Wilfong [12] gave a $2^{poly(n,m)}$-time algorithm, where n is the total number of vertices in the polygons defining the obstacles, and m the maximum number of bits required to specify any vertex of obstacles; their algorithm only decides whether a path is feasible, without necessarily finding one. Jacobs and Canny [16] gave an $O(n^3 \log n + (n+L)^2/\delta^2)$-time algorithm that finds an approximate shortest path provided the shortest path is δ-*robust*, where L is the total edge length of the obstacles. The running time was recently improved by Wang and Agarwal to $O((n^2/\delta^2)\log n)$ [38]. (Informally, a path is δ-robust if perturbations of any point of the path by δ — in distance or in angle — do not violate the feasibility of the path.) Wilfong [39] studies a restriction in which the robot must stay on one of m line-segments (thought of as “lanes”), except to turn between lanes. For a scene with n obstacle vertices, his algorithm preprocesses the scene in time $O(m^2(n^2 + \log m))$, following which queries are answered in time $O(m^2)$. Other, more general, dynamic constraints are considered in [6, 7, 10, 29, 35].

This paper is organized as follows. In Section 2 we formalize our notation (some of which are borrowed from the paper by Fortune and Wilfong [12]), define our model, and state our main results. Section 3 deals with the case when B is not allowed to make reversals, and Section 4 deals when B is allowed to make reversals. We conclude in Section 5 by suggesting some extensions and open problems.

2 Model and Main Results

An *orientation* is defined as a point on $\mathbb{S}^1$, the unit-radius circle centered at the origin. We represent an orientation θ as the angle that the ray emanating from the origin in the direction θ makes with the positive x-axis in the counterclockwise direction. A *position* X is a pair $(\text{LOC}(X), \psi(X))$, where $\text{LOC}(X)$ is a point in the plane representing the location of X and $\psi(X)$ is the orientation of X; we assume $0 \le \psi(X) < 2\pi$. The "reversal" of position X, denoted by $\text{REV}(X)$, is defined to be $(\text{LOC}(X), \psi(X) + \pi)$. We will use L_X to denote the (oriented) line passing through $\text{LOC}(X)$ in direction $\psi(X)$, and $C_L(X), C_R(X)$ to denote the two circles tangent to L_X at $\text{LOC}(X)$ — the center of the the former being to the left (and the center of the latter to the right) of L_X. We assume $C_L(X)$ is counterclockwise oriented and $C_R(X)$ is counterclockwise oriented.

The graph of a function $P : [0, l] \to \mathbb{R}^2$ is called a *smooth path* if $P(t) = (x_P(t), y_P(t))$, $x_P, y_P : [0, l] \to \mathbb{R}$ are differentiable functions, and the derivatives x'_P, y'_P are continuous and do not vanish simultaneously. Since any smooth path has finite length [**?**], we assume P is parametrized by its arc length. Let $\varphi_P(t)$ denote the orientation of the tangent to $P(t)$. P is a line segment if $\varphi_P(t)$ is the same for all $t \in [0, l]$, and a *closed curve* if $P(0) = P(l)$ and $\varphi_P(0) = \varphi_P(l)$. We say that P is a path from a position X to another position Y if $P(0) = \text{LOC}(X)$, $\varphi_P(0) = \psi(X)$, $P(l) = \text{LOC}(Y)$, and $\varphi_P(l) = \psi(Y)$. P is called *moderate* (or *curvature constrained*) if $|P'(t_1) - P'(t_2)| \le |t_1 - t_2|$ for all $t_1, t_2 \in [0, l]$; the curvature of a moderate path is at most 1 almost everywhere, except at finitely many points.

If the point robot B is not equipped with a reverse gear, then it can follow only a moderate path, but if it is equipped with a reverse gear, it can reverse its orientation instantaneously. Therefore, we need to define the notion of forward, backward, and composite paths between two positions. Let X and Y be two positions. A moderate path from X to Y is called a *forward path*, and a forward path from $\text{REV}(Y)$ to $\text{REV}(X)$ is a *backward path* from X to Y. A *composite path* from X to Y is a sequence of moderate paths $P_1, P_2, \ldots P_k$, alternating between forward and backward paths, such that if P_i is a forward path from a position X_{i-1} to another position X_i then P_{i+1} is a backward path from X_i to X_{i+1}, and vice-versa, Moreover, $X_0 = X$ and $X_k = Y$. See Figure 1 for examples of forward, backward, and composite paths. We call the positions X_i (or sometimes their locations), for $1 \le i < k$, the *cusps* of P. We call each P_i the *component path* of P. A composite path is moderate if each of its component path is moderate.

An *obstacle* is a convex region O on the plane. O is *moderate* if its boundary is a moderate (closed) curve. We assume that the boundary of each obstacle is piecewise algebraic—it is composed of a sequence of edges, each of which is an algebraic arc of constant degree.

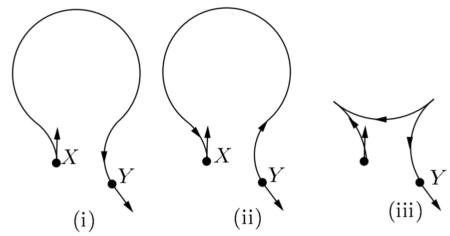

Figure 1: (i) a forward path from X to Y, (ii) a backward path from Y to X, and (iii) a composite path from X to Y with 3 component paths.

A *moderate scene* is a collection of disjoint moderate obstacles. See Figure 2 for an example. (This scene may resemble streets and traffic islands; indeed, algorithms for driving a steering-constrained robot through streets are the focus of a substantial European research initiative [13].) The *size* of a scene is the total number of edges in its obstacles.

Figure 2: A moderate scene.

A (composite or forward) path is *feasible* if it is moderate and does not intersect the interior of any obstacle in the scene. A composite path is *k-feasible* if it is feasible and contains at most k component paths. Given two positions S and F and a moderate scene, a forward (resp. composite) path from S to F is optimal if it is feasible, and its length is minimum among all forward (resp. composite) paths from S to F. (It can be argued that the minimum always exists.) A composite path P from S to F is *k-optimal* if it is k-feasible, and its length is minimum among all k-feasible paths from S_P to F_P. See Figure 1 for examples of forward, backward, and composite paths.

Finally, a few words about the model of computation. Our algorithms work under the arithmetic model of computation common in computational geometry [31], i.e., we assume infinite precision arithmetic operations. We also assume that the roots of a constant degree polynomial can be computed in $O(1)$ time, so that various primitive operations on the edges of obstacles — computing common tangents of a pair of edges, computing a unit-radius circle tangent to a pair of edges, computing the points of vertical tangencies on an edge — can be performed in $O(1)$ time.

In this paper we present efficient algorithms for computing near optimal paths in a moderate scene for a point robot B that may or may not be equipped with the reverse gear. In the first part of this paper (Section 3), we concentrate on the case when B does not have a reverse gear, so it can follow only a moderate forward path. We give a characterization of (forward) optimal paths in moderate scenes, and present efficient algorithm that computes a a near-optimal path for a given instance. We make no δ-robustness assumptions on the optimal paths.

Theorem 1 (i) *Given a moderate scene of size n, an initial position S, a final position F, and a real parameter ε; a feasible forward path from S to F, whose length is at most $\varepsilon > 0$ more than that of an optimal path, can be computed in time $O(n^2 \log n + (1/\varepsilon)\log(1/\varepsilon))$. Here n is the total complexity of the scene.*

(ii) *Given a moderate scene of complexity n, an initial position S, and a final position F such that the distance between $\textsc{loc}(S)$ and $\textsc{loc}(F)$ is at least 6; a forward optimal path from S to F can be computed in time $O(n^2 \log n)$.*

In the second part of the paper (Section 4), we consider the case when reversals are allowed. The main result of this section is the following theorem.

Theorem 2 *Given a moderate scene, an initial position, a final position, and a positive integer k, we can compute in time $O(n^2 \log n + k)$ a ck-feasible path P of length at most $\ell_k + c_1$ where ℓ_k is the length of the k-optimal path, and c_1, c_2 are constants independent of k.*

Note: In the above theorem, if we allow to find a path from S to either F or $\textsc{rev}(F)$, then the values of c_1, c_2 can be improved. See Section 4 for more details on this.

We conclude this section by defining a few more terms. Let $P : [0, l] \to \mathbb{R}^2$ be a path. We will use $\|P\|$ to denote the arc length of P, and P^{-1} to denote the reverse of P, i.e., $P^{-1}(t) = P(l - t)$. We say that a position X is on P if $\textsc{loc}(X) = P(t)$ and $\psi(X) = \varphi_P(t)$ for some $0 \le t \le l$. For two positions $X, Y \in P$, we use $P[X, Y]$ to denote the portion of P from X to Y. The orientation of an undirected curve can be defined by specifying a position X on C, with $\textsc{loc}(X) \in C$ and L_X being a tangent of C. For example, $C_L(X), C_R(X)$ are counterclockwise and clockwise directed, respected. Similarly, if a path touches the boundary of an obstacle at a position X, then X defines the oriented of ∂O. For two paths P_1, P_2 such that the final position of P_1 is the same as the initial position of P_2, $P_1 \circ P_2$ will denote the path resulting from their concatenation. We denote by $\theta(P) = \int_0^l \varphi_P(t)dt$ the cumulative angle of rotation (of the orientation vector) as we go along P from its initial position to its final position, with counterclockwise rotation considered positive. For example, if P is a circular arc of unit radius then $|\theta(P)|$ is equal to the length of P. If $P = P_1 \circ P_2$ then $\theta(P) = \theta(P_1) + \theta(P_2)$.

Let P_1 and P_2 be two paths. A *tangent segment* from P_1 to P_2 is a line segment from a position $X \in P_1$ to another position $Y \in P_2$ such that $\psi(X) = \psi(Y)$ (i.e, the line supporting the tangent segment is a tangent to both P_1 and P_2). In general, the tangent segment may not be defined, or there may be more than one tangent segment from P_1 to P_2, but if P_1 and P_2 are two disjoint, convex closed curves, then there is a unique tangent segment from P_1 to P_2.

3 Paths without Reversals

In this section, we study feasible paths without reversals, so only forward paths are allowed. Hence, by feasible (resp. optimal) paths, we will always mean foward feasible (resp. optimal) paths. We first characterize the optimal paths amid a moderate scene and then describe a simple, efficient algorithm for computing a path that is at most ε longer than the optimal path, for any given positive ε. Our approach is to reduce the problem to a shortest-path problem in a certain directed network.

Let P be a feasible path in a moderate scene. A non-empty subpath of P is called an *O-segment* of P if it lies on the boundary of an obstacle and is maximal with respect to this property; a *C-segment* if it is a maximal circular arc of unit radius that is not (a part of) an O-segment; an *L-segment* if it is a maximal line segment that does not contain or overlap an O-segment. Suppose, for example, that a path P consists of a C-segment, an L-segment, and a C-segment in this order. Following Dubins[11] convention, we say that P is of type CLC. This notation is generalized to an arbitrary sequence of O-/C-/L-segments.

3.1 Scenes without obstacles

We first consider optimal paths in an obstacle-free scene. Dubins [11] gave the following characterization of optimal paths.

Theorem 3 (Dubins [11]) *Any optimal path is of type CCC or CLC, or a substring thereof.*

We briefly sketch his proof and rederive some of his key lemmas because we need to generalize them for later use.

Sketch: The proof proceeds in two parts. He first establishes that if P is an optimal path whose length is bounded by a small constant (say, $\pi/8$), then P must be of type CLC or a substring thereof (Proposition 5.12 of [11]). Since a subpath of an optimal path itself is an optimal path, this implies that any optimal path is a finite sequence of C- and L-segments. He then shows that a path of type LCL, CCL, LCC, or CCCC cannot be an optimal path (this is established below in Lemma 4 for LCL/CCL/LCC, and in Lemma 7 for CCCC). This completes the proof, because the only types that remain possible among the types described by a string of at least 3 symbols are CCC and CLC. □

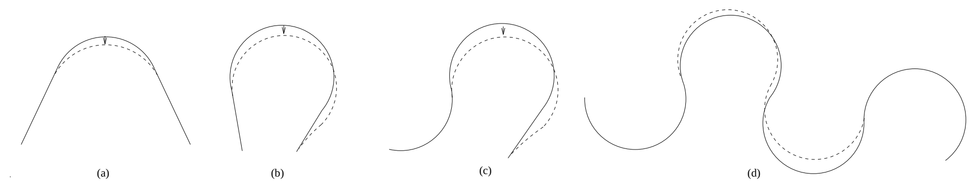

Figure 3: Length-reducing perturbations

Figure 3 illustrates the perturbations for eliminating subpaths of types LCL (Figure 3(a), (b)), CCL (Figure 3(c)), and CCCC (Figure 3(d)); they will be defined more formally in the lemmas to follow. In each case, the solid line path is perturbed into the dashed line path, which is strictly shorter than the original path. It is easy to see that perturbation (a) is indeed length-reducing as long as the length of the C-segment in the middle is at most π. Readers are referred to [4, 11] for perturbation (d) (we remark that in Lemma 6 below we give an elementary geometric proof of the claim that perturbation (d) is length-reducing). We formalize perturbations (b) and (c) in Lemma 7 below, in a generalized form to be used later. The proof of this lemma uses essentially the same argument as Dubins [11].

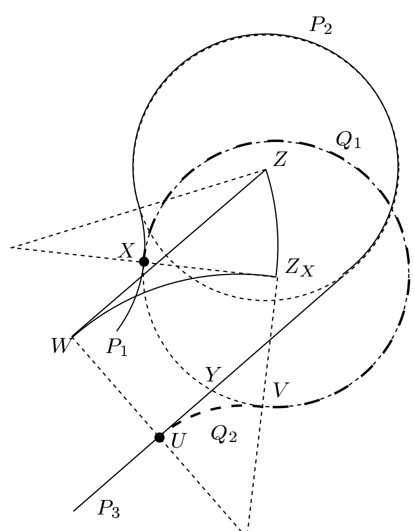

Figure 4: Perturbations (b) and (c)

Lemma 4 (Dubins [11]**)** *A path of type LCL, CCL, or LCC cannot be an optimal path.*

Proof: The claim is easy to prove if the length of the middle C-segment is at most π — there is an obvious perturbation of the path (shown in Figure 3(a)) that reduces the length while remaining moderate and preserving the initial and final positions. Suppose the length of the C-segment in the middle is greater than π. Dubins proves the lemma for LCL-type paths and leaves CCL for the reader (the case of LCC is symmetric to CCL). We prove the lemma for CCL-type paths, because we need to generalize it later. (In fact, as we will see below, both LCL and CCL are special cases of this generalization.) Our proof is similar to Dubin's proof. In particular, we will show that the perturbation in Figure 3(c) reduces the length of an CCL-type path. We first define this perturbation formally, and then argue that it reduces the length.

Let $P = P_1 \circ P_2 \circ P_3$ be a CCL-type path, where P_1 and P_2 are C-segments, P_3 is an L-segment, and $\pi < \|P_2\| < 2\pi$. Without loss of generality, we assume that P_1 is counterclockwise oriented, so P_2 is clockwise oriented. Let X be a position on P_1 and $C_X = C_R(X)$; by our convention C_X is clockwise oriented. We choose X close enough to the final point of P_1 so that C_X intersects with both P_2 and P_3. (Since $\|P_2\| > \pi$, a perturbation argument shows that such an X always exists.) Let P_1' be the part of P_1 starting from X and ending at the final point of P_1. Let Y be the intersection of C_X and P_3. (Since C_X intersects P_2 also, it can intersect P_3 at only one point.) We draw a unit-radius circle C tangent to both P_3 and the arc $C_X[X, Y]$. Let U, V be the positions on C at which C touches P_3 and C_X, respectively. Let P_3' be the initial segment of P_3 ending at U, $Q_1 = C_X[X, V]$, and $Q_2 = C[V, U]$; see Figure 4. Consider the perturbation of the original path obtained by replacing $P' = P_1' \circ P_2 \circ P_3'$ with $Q = Q_1 \circ Q_2$.

In order to prove that the length of the perturbed path is less than that of P, it suffices to show that $\|P'\| - \|Q\| > 0$. First observe that

$$\|P_2\| = |\theta(P')| + |\theta(P_1')|, \quad (1)$$

$$\|Q_1\| = |\theta(Q)| + |\theta(Q_2)|. \quad (2)$$

Therefore,

$$\begin{aligned} \|P'\| - \|Q\| &= \|P_1'\| + \|P_2\| + \|P_3\| - \|Q_1\| - \|Q_2\| \\ &= \|P_1'\| + |\theta(P')| + |\theta(P_1')| + \|P_3\| - |\theta(Q)| - |\theta(Q_2)| - \|Q_2\| \\ &\qquad \text{(Using (1) and (2))} \\ &= \|P_1'\| + |\theta(P_1')| + \|P_3'\| - \|Q_2\| - |\theta(Q_2)|, \end{aligned} \quad (3)$$

where the last equality follows from the fact that $\theta(P') = \theta(Q)$.

Since P_1' is unit-radius circular arc, $\|P_1'\| + |\theta(P_1')| = 2\|P_1'\|$, which is the same as the length of the circular arc $\widehat{ZZ_X}$ (of radius 2) in Figure 4, where Z is the center of the circle containing the arc P_2, Z_X is the center of C_X, and the center of the arc $\widehat{ZZ_X}$ is the same as that of P_1. Similarly, we represent the quantity $\|Q_2\| + |\theta(Q_2)|$ by the length of the circular arc $\widehat{Z_X W}$, which has the same center and and the same angle as Q_2. Substituting these quantities in (**??**), we obtain

$$\|P'\| - \|Q\| = \|\widehat{ZZ_X}\| + \|P_3'\| - \widehat{Z_X W}. \tag{4}$$

Note that the line segment WZ has the same length as P_3'. Since the path $WZ \circ \widehat{ZZ_X}$ lies above the circular arc $\widehat{Z_X W}$ (of length less than π), it can be shown that $\|WZ\| + \|\widehat{ZZ_X}\| > \|\widehat{Z_X W}\|$, and therefore, by (4), $\|P'\| - \|Q\| > 0$, as required. □

Next, we will use the above length-reducing perturbation in a more general form, but we first need a technical lemma whose proof is omitted from here.

How about adding a short proof?

Lemma 5 *Let R be a moderate path. Place a unit-radius circle so that it is tangent to R at its initial point and it is on the right hand side of R with respect to the orientation of R. Let R' be the trajectory of the center of this circle as we roll this circle on R. Then, $\|R'\| = \|R\| + \theta(R)$.*

We say that a path P is *right-convex* (*left-convex*, respectively) if its curvature is non-negative (non-positive, respectively) everywhere along P. For example, a counterclockwise oriented circular arc is right-convex, as well as a line segment.

the term "curvature" may be confusing because we have been using it for the curvature in absolute value

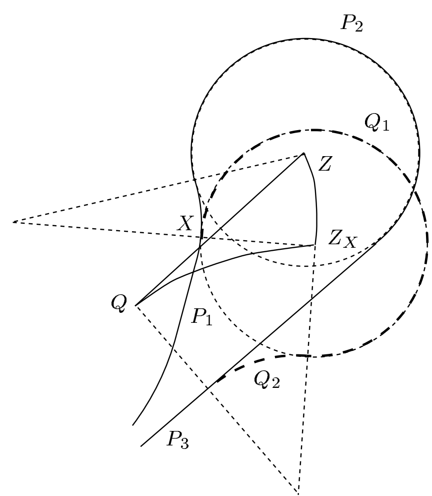

Figure 5: Illustration to Lemma 6

Lemma 6 *Let P and Q be two distinct moderate paths from X to Y such that*

(i) $P = P_1 \circ P_2 \circ P_3$, *where P_1 is a right-convex moderate path, P_2 is a clockwise-oriented circular arc of unit radius, and P_3 is a line segment;*

(ii) $Q = Q_1 \circ Q_2$, *where* Q_1 *is a clockwise-oriented circular arc of unit radius and* Q_2 *is a right-convex moderate path;*

(iii) P *and* Q *intersect exactly once and the intersection point lies on* Q_1.

Then $\|P\| > \|Q\|$. *See Fig 5.*

Note: Lemma 4 (i) is a special case of this lemma with P_1 being a line segment, and Lemma 4 (ii) is another special case with P_1 being a unit-radius disc.

Proof: The proof proceeds along the same lines as of Lemma 4 (ii). P (Q, respectively) in this lemma corresponds to the path P' (Q, respectively) of Lemma 4. The only difference is that P_1 and Q_2 here are not necessarily circular arcs. We can, however, still obtain the equality (3), and then apply Lemma 5. □

We now eliminate the possibility of an optimal path of type CCCC. Our proof is purely geometric and more intuitive than the analytic proof of Dubins.

Lemma 7 (Dubins [11]) *A path of type CCCC cannot be an optimal path.*

Proof: Let P be a moderate path of type CCCC. We define the perturbation shown in Figure 3(d) formally, and show that its length is less than that of P. Suppose $P = P_0 \circ P_1 \circ P_2 \circ P_3$, where each P_i is a C-segment of P. We can assume that the lengths of both P_1 and P_2 are greater than π, because otherwise we can apply perturbation (a) of Figure 3. For later use, we generalize the situation as before: P_0 is a left-convex moderate path with non-zero curvature at its final point (not necessarily a unit-radius arc) and P_3 is a right-convex moderate path. Let C_1, C_2 be the circles containing P_1, P_2, respectively, and let X_1 and X_2 be their centers. P_1 and P_2 respectively. Let Z_{i-1}, Z_i denote the initial and final positions of P_i. See Figure 6. We assume that line segment $\textsc{loc}(Z_0)X_1$ is on the x-axis, X_2 is above the x-axis, and the angle τ of the ray $X_2\textsc{loc}(Z_2)$ with the x-axis is non-positive, as in Figure 6. (Otherwise, we can rotate, reverse the direction of the $+y$-axis, and reverse the orientation of P, respectively.) Now, roll the circle containing P_1 by a small amount on the path P_0. Let C_1' denote the resulting (counter-clockwise oriented) circle, Y_1 the center of C_1', and Z_0' the position at which P_0 is tangent to C_1'. Let θ be the angle (in absolute value) between the line segments X_1Z_0 and Y_1Z'. Let Y_1Y_2 be the edge opposite to X_1X_2 in a parallelogram. Let Q_2' be the translation of P_2 with center at Y_2, let Z_1', Z_2' denote the initial and final positions of Q_2'; by construction Q_2' is tangent to C_1' at Z_1'. Let Q_3 be the tangent segment from Q_2 to P_3; if Z_0' is chosen sufficiently close to Z_0, Q_3 is defined. Let U_1, U_2 be the initial and final positions of Q_3. Set $Q_1 = C_1'[Z_0', Z_1']$, $Q_2 = Q_2[Z_1', U_1]$, $P_0' = P_0[Z_0', Z_0]$, and $P_3' = P_3[Z_2, U_2]$. We shortcut P by replacing $P[Z_1', U_2]$ with $Q_1 \circ Q_2 \circ Q_3$.

We need to prove that the new path is shorter. Let Q_3' be the segment connecting $\textsc{loc}(Z_2')$ to $\textsc{loc}(Z_2)$. Then

$$\|Q_1 \circ Q_2 \circ Q_3\| < \|Q_1\| + \|Q_2'\| + \|Q_3'\| + \|P_3'\|.$$

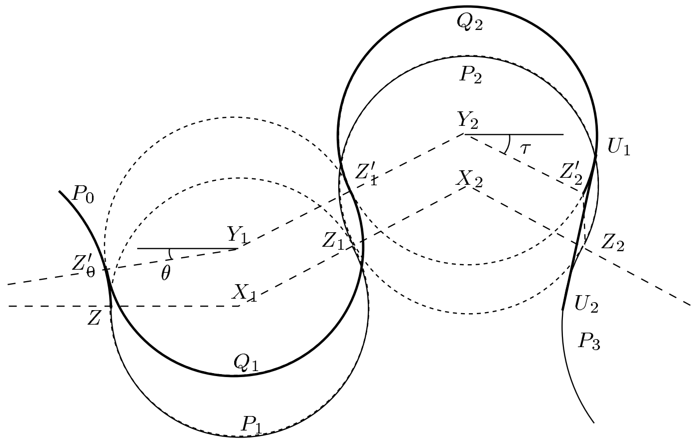

Figure 6: Eliminating type CCCC

Hence it suffices to show that $\|Q_1\| + \|Q_2'\| + \|Q_3'\| < \|P_0'\| + \|P_1\| + \|P_2\|$. Note that

$$\|Q_1\| \quad = \quad \|P_1\| - \theta, \quad \|Q_2'\| = \|P_2\|.$$

On the other hand, by Lemma 5,

$$\|Q_3'\| < \|P_0'\| + \theta.$$

This completes the proof of the lemma. □

3.2 Scenes with moderate obstacles

In the presence of (arbitrary) obstacles, Fortune and Wilfong [12] and Jacobs and Canny [16] observe that any subpath of an optimal path is a Dubins path if the subpath does not touch any obstacle except at the endpoints. Thus, any optimal path is a finite sequence of C-, L-, or O-segments. This observation alone, unfortunately, does not narrow our search to a finite set of Dubins paths. Exploiting our moderate-scene assumption, we develop below a characterization of optimal paths that leads to an efficient algorithm. Our approach, extending that of Dubins, is to exclude certain sequences of C-/L-/O- segments as valid subpaths of an optimal path.

We first give a few definitions. We refer to the circles $C_L(S)$ and $C_R(S)$ as *initial circles*, and to $C_L(F)$ and $C_R(F)$ as *final circles*. A unit-radius circle is called *anchored*, if it is either one of the initial/final circles or is tangent to at least two objects, each of which is either an initial/final circle or an obstacle. A C-segment of a path is called *terminal* if it is the first or the last segment of the path; otherwise, it is *nonterminal*. A C-segment

$\mathcal{S}$ of a path is called *anchored* if either $\mathcal{S}$ is terminal, or it touches obstacles or terminal C-segments at two distinct points (notice that $\mathcal{S}$ can touch a terminal C-segment only at one of its endpoints, i.e., $\mathcal{S}$ is adjacent to a terminal C-segment). Otherwise $\mathcal{S}$ is called *floating*. Each anchored C-segment must be an arc of an anchored circle.

As proved below (see Lemma 21), there are only $O(n)$ anchored circles, but there are infinite number of floating circles. One might hope, by further investigation, that an optimal path does not contain any floating C-segment. Unfortunately this is not the case. For example, in Figure 7, it has been numerically verified that the path from S to F represented by a thick line, which has floating C-segments, is shorter than any feasible path without floating C-segments. We thus seek as restrictive a characterization of floating segments as possible. The following two theorems give such a characterization.

Can we say something more convincing here?

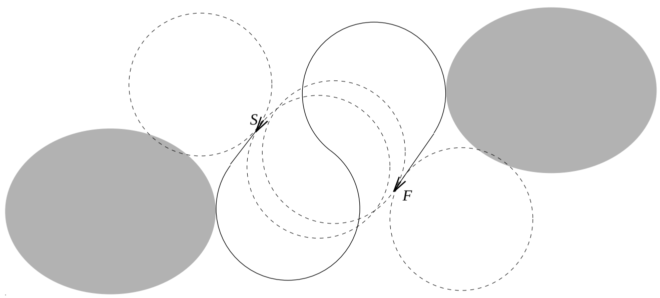

Figure 7: Floating C-segments: the path follows the arcs of two circles each of which is tangent to only a single obstacle.

Theorem 8 *Any optimal path contains at most two floating C-segments, each touching an obtscale or an initial/final circle. Moreover, if it contains two floating C-segments, then they lie on a single subpath of type CC.*

Theorem 9 *Let P_1 and P_2 be consecutive C-segments of an optimal path from S to F. Then the center of P_1 is within distance 3 of $\textsc{loc}(F)$, and the center of P_2 is within distance 3 of $\textsc{loc}(S)$.*

We prove these theorems by a sequence of lemmas.

Lemma 10 *The length of every nonterminal C-segment in any optimal path is more than π.*

Proof: Suppose a moderate path has a nonterminal C-segment of length at most π. Since each obstacle is moderate, no obstacle touches this C-segment from its concave side. Thus, we can still apply perturbation (a) of Figure 3. $\square$

Lemma 11 *If an optimal path contains a subpath of type CCC, then either the first or the last C-segment of this subpath is terminal.*

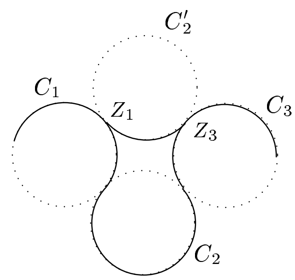

Figure 8: Shortcutting a CCC-type subpath

Proof: Let P be an optimal path, and let $P_1 \circ P_2 \circ P_3$ be a subpath of type CCC, such that neither P_1 nor P_3 is a terminal C-segment. Then, by Lemma 10, $\|P_i\| > \pi$ for each $1 \leq i \leq 3$. Assume that P_2 is counterclockwise oriented. Let C_i be the circle containing P_i, and let $C_2' \neq C_2$ be the second circle tangent to both C_1 and C_3. Set $Z_1 = C_1 \cap C_2'$ and $Z_3 = C_3 \cap C_2'$. Since $\|P_i\| > \pi$, $Z_1 \in P_1$ and $Z_3 \in P_3$. Our moderate-scene assumption implies that the circular arc $C_2'[Z_1, Z_3]$ is feasible. We can therefore shortcut P by replacing its subpath from Z_1 to Z_3 with the circular arc $C_2'[Z_1, Z_3]$. □

Lemma 12 *Each nonterminal C-segment of an optimal path is either tangent to at least one obstacle, or is adjacent (on the path) to a terminal C-segment.*

Proof: Fix a nonterminal C-segment of an optimal path. Let P be the subpath of this optimal path consisting of this C-segment together with the two segments immediately adjacent. If P contains a terminal C-segment or an O-segment, we are done. Suppose otherwise. By Lemma 11, P cannot be of type CCC. Thus, the type of P must be LCL, CCL, or LCC. Unless the middle C-segment touches any obstacle, we can apply a length-reducing perturbation (a), (b), or (c) in Figure 3, which contradicts the fact that P is optimal. □

Lemma 13 *If an optimal path contains a subpath of type LCL, OCL, LCO, or OCO, then the C-segment in this subpath is anchored.*

Proof: Let P be a subpath of an optimal path consisting of 3 consecutive segments, with a C-segment in the middle. If P is of type *OCO*, this C-segment is anchored, by definition. If P is of type OCL or LCO and the middle C-segment were floating, then the length-reducing perturbation of Lemma 6 would apply. Finally, suppose P is of type LCL.

We have two ways of applying the length-reducing perturbation of Lemma 6, taking the first or the last line segment as a special case of a convex path P_1 in the lemma. A single obstacle cannot obstruct both of these perturbations simultaneously. $\square$

Corollary 14 *Any floating C-segment in an optimal path must be adjacent to another C-segment.*

We can further narrow down the occurrences of floating C-segments.

Lemma 15 *Let $P = P_1 \circ P_2 \circ P_3 \circ P_4$ be an optimal path, where P_2 and P_3 are C-segments. Let C_i denote the circle containing the C-segment P_i, $i = 2, 3$. Suppose P_2 is oriented counterclockwise and P_3 clockwise. Then, for every position X on P_1, $C_L(X)$ intersects C_3 and, for every position Y on P_4, $C_R(Y)$ intersects C_2.*

Proof: Let Z_{i-1}, Z_i denote the endpoints of P_i. For the sake of simplicity, let us assume that the Z_1 is the lowest point of C_2 (i.e., $\psi(Z_1) = 0$). Let C_2 and C_3 be counterclockwise and clockwise oriented, respectively. Suppose there is some position X on P_1 such that $C_L(X)$ does not intersect the interior of C_3. Let X_0 be the last such position on P_1, i.e., $C_L(X)$ intersects the interior of C_3 for every position $X \in P_1$ after X_0. Then, $C_0 = C_L(X_0)$ is tangent to C_3 at a position Y and the length of the arc $C_0[X_0, Y]$ is less than π (the second claim uses the fact that the curvature of P_1 is at most 1). Let $P_1' = P[X_0, Z_1]$ and $Q_1 = C_0[X_0, Y]$.

We will construct a new path P' from P, which is shorter than P. There are three cases to consider: (i) $Y \in P_3$, (ii) $Y \notin P_3$ and C_3 does not intersect the interior of any obstacle, and (iii) $Y \notin P_3$ and C_3 intersects the interior of some obstacle; see Figure 9.

Case 1: $Y \in P_3$. We obtain a new path P' by replacing $P[X_0, Y]$ with the circular arc Q_1 (see Figure 10). We claim that Q_1 does not intersect the interior of any obstacle, which implies that P' is feasible. Suppose, for the sake of contradiction, that Q_1 intersects an obstacle O. It can be shown that the relative interior of Q_1 does not intersect $P[X_0, Y]$, because otherwise we can obtain a shorter feasible path from X_0 to Y. Since X_0 and Y do not lie in the interior of O, and since the curvature of ∂O is at most 1; ∂O intersects C_0 at two points, and both of them lie in the relative interior of Q_1. Let R be the connected component of the region bounded by the path $P[X_0, Y]$ and the arc $C_0[Y, X_0]$ ($C_0[Y, X_0]$ is the closure of $C_0 \setminus Q_1$; R is the shaded region in Figure 10); note that Q_1 lies inside R. Since P does not intersect the interior of O and ∂O does not intersect $C_0[Y, X_0]$, $\partial O \subset R$. This is, however, impossible, because then we can find a position $X_0 \neq X \in P_1'$ for which $C_L(X)$ does not intersect the interior of C_3. Hence, Q_1 does not intersect the interior of any obstacle. Finally, P' is shorter than P because $\|P_2\| > \pi$.

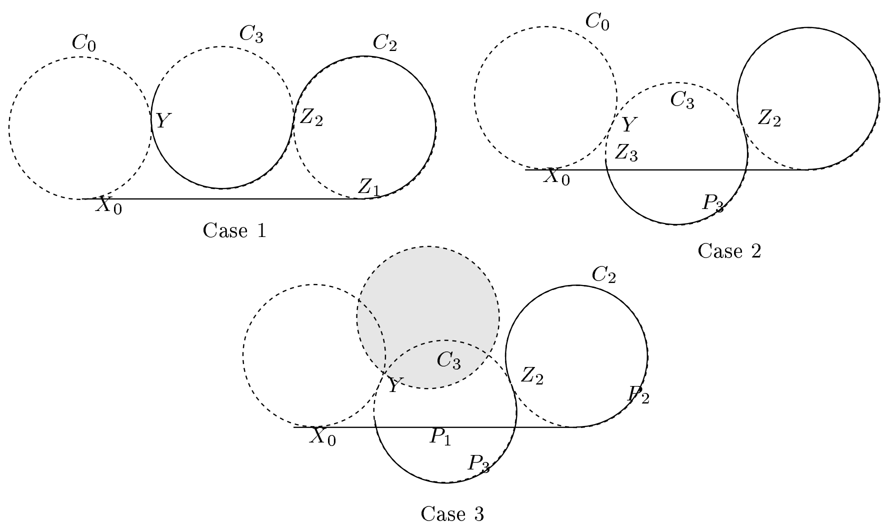

Figure 9: Three different cases

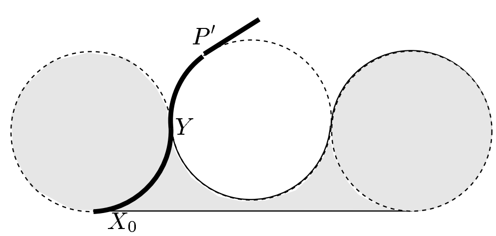

Figure 10: Case 1: $Y \in P_2$.

Case 2: $Y \notin P_3$ and C_3 does not intersect the interior of any obstacle. Assume that C_3 is clockwise oriented. Set $Q_2 = C_3[Y, Z_2]$. We obtain a new path P' from P by replacing $P_1' \circ P_2$ with the path $Q_1 \circ Q_2$; see Figure 11. Since C_3 does not intersect the interior of any obstacle, both $C_3[Y, Z_2]$ and Q_1 are feasible. Hence, P' is also feasible.

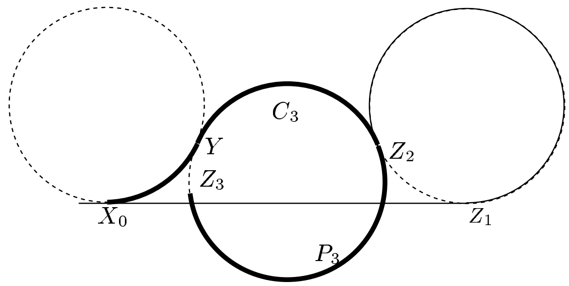

Figure 11: Case 2: $Y \notin P_2$ and C_3 does not intersect any obstacle.

Next, we prove that $\|P'\| < \|P\|$. Set $Q_3 = C_2[Z_2, Z_1]$ (i.e., Q_3 is the closure of $C_1 \setminus P_1$) and $Q = Q_1 \circ Q_2 \circ Q_3$. In order to prove $\|P'\| < \|P\|$, we need to prove that $\|Q_1\| + \|Q_2\| < \|P_1'\| + \|P_1\|$. But $\|P_1\| = 2\pi - \|Q_3\|$, so it suffices to prove that

$$\|Q_1\| + \|Q_2\| + \|Q_3\| < \|P_1'\| + 2\pi. \tag{5}$$

Note that if $\|Q_1\| + \|Q_3\| < \|Q_2\|$, (5) is obvious because $\|Q_2\| < \pi$. So assume that $\|Q_1\| + \|Q_3\| \geq \|Q_2\|$.

Q and P_1' have the same initial and final positions, therefore $\theta(P_1') = \theta(Q) + 2c\pi$ for some integer c. But $\theta(Q) = \|Q_1\| + \|Q_3\| - \|Q_2\|$. Moreover, $\|P_1'\| \geq |\theta(P_1')|$, for P_1' is a moderate path. Hence,

$$\|P_1'\| \geq \Big| \|Q_1\| + \|Q_3\| - \|Q_2\| + 2c\pi \Big|.$$

Using the fact the each of $0 < \|Q_i\| < \pi$, for $1 \leq i \leq 3$, the right hand side of the above inequality is minimized when $c = 0$ or $c = -1$. First consider the case when $c = 0$. Since $\|Q_1\| + \|Q_3\| \geq \|Q_2\|$,

$$\begin{aligned} \|P_1'\| + 2\pi &\geq \|Q_1\| + \|Q_3\| - \|Q_2\| + 2\pi \\ &= \|Q_1\| + \|Q_2\| + \|Q_3\| - 2\|Q_2\| + 2\pi \\ &> \|Q_1\| + \|Q_2\| + \|Q_3\|, \end{aligned}$$

because $\|Q_2\| < \pi$. Similarly when $c = -1$,

$$\begin{aligned} \|P_1'\| + 2\pi &\geq 4\pi - (\|Q_1\| + \|Q_3\|) + \|Q_2\|) \\ &= \|Q_1\| + \|Q_2\| + \|Q_3\| - 2(\|Q_1\| + \|Q_3\|)) + 4\pi \\ &> \|Q_1\| + \|Q_2\| + \|Q_3\|, \end{aligned}$$

because $\|Q_1\|, \|Q_3\| < \pi$. This completes the proof for this case.

Case 3: $Y \notin P_2$ and C_3 intersects an obstacle. Since P_3 touches an obstacle and $Y \notin P_3$, P_3 does not lie completely above P_1' (because this would imply that the interior of the region bounded by $P[X_0, Y]$ and the arc $C_0[Y, X_0]$ contains an obstacle, which is impossible, as argued earlier). Let σ be the last intersection point of P_1' and P_3 (along P_1'), and let θ_1 (resp. θ_3) be the orientation of P_1 (resp. P_3) at σ. Set $\Sigma_1 = (\sigma, \theta_1)$ and $\Sigma_3 = (\sigma, \theta_3)$. If P_1' touches P_3 at σ (i.e., $\Sigma_1 = \text{REV}(\Sigma_3)$), then let $U = \Sigma_1$ and $V = \Sigma_3$. Otherwise, let C_4 be the (counterclockwise oriented) circle that is tangent to both $P_1' \circ P_2$ and P_3, and that lies below both of them; if there is more than one such circle, we choose the last one along P_1'. Let U (resp. V) be the position on $P_1' \circ P_2$ (resp. P_3) tangent to C_4, and let $Q_1 = C_4[U, V]$. Obviously $\|Q_1\| \leq \|P_3[\Sigma_3, V]\| + \|P[\Sigma_1, U]\|$. Next, let Q' be the path (not necessarily smooth) $P_1'[X_0, \Sigma_1] \circ P_3^{-1}[\text{REV}(\Sigma_3), \text{REV}(U)]$, and let Q_2 be the optimal, feasible path from X_0 to $\text{REV}(U)$ that is homotopic to Q' — Q' initially follows C_0 in the counterclockwise direction, and then, after a sequence of LO subpaths, it follows P_2 and P_1' in the opposite direction; see Figure 12. Obviously $\|Q_2\| < \|Q'\|$. We construct a new path P' from P by replacing $P[X_0, V]$ with $Q_2 \circ Q_1$. It is easily seen that P' is feasible. Finally,

$$\begin{aligned}
\|Q_2 \circ Q_1\| &= \|Q_2\| + \|Q_1\| \\
&< \|P_1'[X_0, \Sigma_1]\| + \|P_3[Z_2, \Sigma_3]\| + \|P[U, Z_2]\| + \|P[\Sigma_1, U]\| + \|P_3[\Sigma_3, V]\| \\
&= \|P[X_0, V]\|.
\end{aligned}$$

Hence P' is shorter than P.

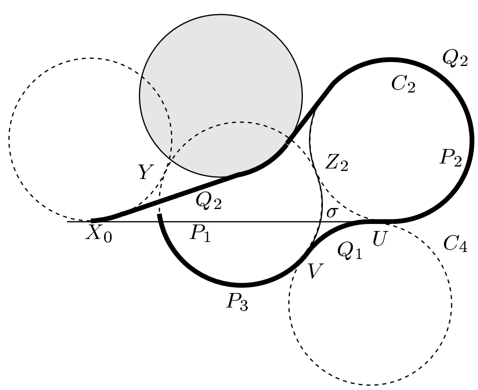

Figure 12: Case 3: $Y \notin P_2$ and C_3 intersects an obstacle

This completes the proof of the lemma. □

Corollary 16 *If an optimal path P contains a nonterminal subpath of type CC then it does not contain any other nonterminal C-segment.*

Proof: Suppose P contains a CC-type path $P_1 \circ P_2$. Suppose P also contain a nonterminal C-segment P_3 after P_2. The length of each P_i is more than π. Let C_i be the circle containing P_i. Suppose C_1 is counterclockwise oriented. Then for every point $X \in P_3$, $C_L(X)$ intersects C_1. Since $\|P_1\| > \pi$, there are two positions X, Y on C_3, with $\psi(X) = \psi(Y) + \pi$ (i.e., $\textsc{loc}(X), \textsc{loc}(Y)$ are antipodal points), such that both $C_L(X)$ and $C_L(Y)$ intersect C_1. But this is impossible, for there is no unit-radius circle whose interior intersects both $C_L(X)$ and $C_L(Y)$ (in fact, the only unit-radius circle that intersects both of them is C_3). □

Theorem 9 is now an easy consequence of Lemma 15. We now complete the proof of Theorem 8, the main characterization result of this section.

Proof of Theorem 8: Let P be an optimal path. First note that a floating circle cannot occur adjacent to the initial or the final C-segment. Suppose for the sake of contradiction, a floating C-segment P_1 follows the initial C-segment. By definition, the circle containing P_1 cannot touch any obstacle, so P_1 is followed by a L-segment. Let X be the final position of this L-segment. Then we have a CCL-type path from the initial position of P to X that does not touch any obstacle, so by Lemma 4, P is not an optimal path. Hence, we can assume that the C-segment adjacent to each floating C-segment is a nonterminal C-segment. If there are two nonadjacent floating C-segments or there are more than two floating C-segments, then P contains a CC-type subpath and another C-segment, which contradicts Corollary 16. □

Corollary 17 *Let $P = P_1 \circ P_2 \circ P_3 \circ P_4$ be an optimal path, where P_2 and P_3 are C-segments. Suppose P_2 is oriented counterclockwise and P_3 clockwise. Then P_1 and P_4 satisfy the following properties:*

(i) *Every point of P lies within distance 6 from S or F.*

(ii) *Every O-segment in P_1 (P_4, respectively) is clockwise (counterclockwise, respectively) oriented.*

Proof: Part (i) is an obvious consequence of Theorem 9. As for (ii), suppose P_4 has a clockwise-oriented O-segment. Let X be the first position of P_4 on this O-segment, and let O be the obstacle whose boundary contains this O-segment. Since the obstacles are moderate, $C_R(X) \subseteq O$, and therefore does not intersect P_2 properly, contradicting Theorem 8. □

Next, we prove a simple theorem, which will be useful in proving the correctness of our algorithm.

Theorem 18 *Suppose there is a subpath of type OCCO in an optimal path. Then, it is impossible for both of the two C-segments in the subpath to be floating.*

Proof: Suppose that the two C-segments in an optimal path of type OCCO are both floating. Then, the perturbation used in the proof of Lemma 7 provides a contradiction. Recall the generalization we had in the proof, which makes this application possible. □

3.3 Algorithm

We now describe a simple algorithm that, given a moderate scene, an initial position S, a final position F, and a real parameter $\varepsilon > 0$; produces a feasible (forward) path from S to F whose length is at most ε more than that of an optimal (forward) path from S to F. As mentioned in the beginning of this section, we reduce the problem to computing a shortest path in a certain network.

Let δ be a real parameter. We define $\mathcal{C} = \mathcal{C}(\delta)$ to be the set of circles that satisfy the following properties.

(P1) C is an anchored circle; Figure 13(i).

(P2) C is tangent to an obstacle O, and the line supporting the common tangent to C and O is tangent to another obstacle or an initial/final circle; Figure 13(ii).

(P3) C is tangent to the line L_S (L_F, respectively) and an obstacle O; Figure 13(iii).

(P4) There exists a pair of obstacles O_1, O_2 and a circle C' such that (C, C') bridges (O_1, O_2) and the points $O_1 \cap C$ and $C \cap C'$ are antipodal; Figure 13(iv).

(P5) C is tangent to an obstacle O at a point p such that $\|O[p_O, p]\| = i\delta$, where p_O is a fixed representative point on ∂O; Figure 13(v).

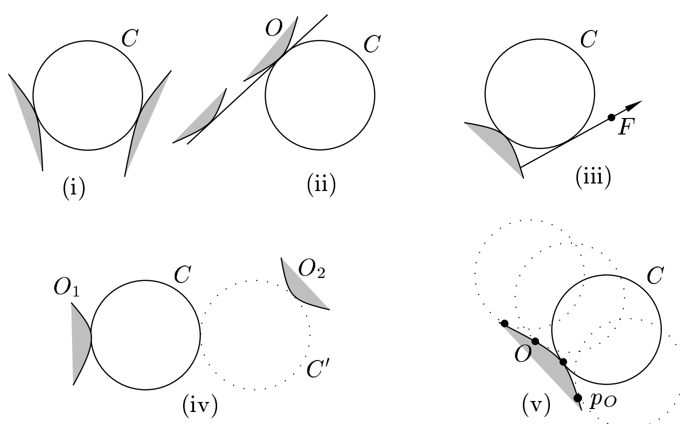

Figure 13: Examples of (P1)–(P5) properties

Let O_i, $i = 1, 2$, be either an obstacle or an initial/final circle. We say that a pair of unit-radius circles C_1 and C_2 *bridges* O_1 and O_2 if C_1 and C_2 are tangent to each other, C_1 is tangent to O_1, and C_2 is tangent to O_2. We call such a pair *relevant* if the center of C_1 is within distance 3 of the final point $\text{LOC}(F)$ and the center of C_2 is within distance 3 of the initial point $\textsc{loc}(S)$. By Theorem 9, a floating C-segment of any optimal path must be a part of a relevant pair of circles. For a circle C tangent to some obstacle O, let R_C denote the set of relevant pairs of the form (C, C'), or of the form (C', C), where C' is a circle tangent to C. Define $\mathcal{R}(\delta) = \bigcup_{C \in \mathcal{C}} R_C$.

One needs to argue that even in approximate shortest path distance 3 (not $3+\varepsilon$) is enough for relevant pairs

Let $\mathcal{A}$ be the set of anchored circles and $\mathcal{L}_1$ the set of lines tangent to two objects in $\mathcal{A} \cup \mathcal{O}$. Set $\delta = \varepsilon / c$, where $c > 1$ is a constant independent of ε, and let $\mathcal{R}$ be the set of circles in $\mathcal{R}(\delta)$. Let $\mathcal{O}_1$ be the set of obstacles that are within distance 6 from S or F and $\mathcal{L}_2$ the set of lines tangent to a pair (C, Γ), where $C \in \mathcal{R}$ and $\Gamma \in \mathcal{O}_1 \cup \{C_L(S), C_R(S), C_L(F), C_R(F)\}$. Set $\Sigma = \mathcal{A} \cup \mathcal{O} \cup \mathcal{R} \cup \mathcal{L}_1 \cup \mathcal{L}_2$. We construct a weighted directed graph G as follows. Let $\mathcal{X}$ be the set of positions at which two objects (i.e. circles, lines, or obstacles) in Σ are tangent. The vertex set V of G is $\{X, \textsc{rev}(X) \mid X \in \mathcal{X}\} \cup \{S, F\}$. There is an edge from X to Y, $X, Y \in V$, if position Y is reachable from position X via some part of an object of Σ without going through any other position in V. The weight of this edge is the length of such a path from position X to position Y.

Lemma 19 *Let P be an optimal path from S to F. If every C-segment of P is contained in some circle of $\mathcal{A} \cup \mathcal{R}$, then P is an S–F path in G.*

Proof: If the circles containing all C-segments of P are anchored, then all L-segments are supported by the lines of $\mathcal{L}_1$, and therefore endpositions of each segment in P belongs to $\mathcal{X}$. It is now easily seen that P is an S–F path in G. If P contains a C-segment of $\mathcal{R}$, then, by Theorem 9, every point on P lies within distance 3 from S or F, and, by Corollary 17, the length of each L-segment is at most 6. Therefore, the line supporting each L-segment belongs to $\mathcal{L}_1 \cup \mathcal{L}_2$, and the lemma follows. □

Theorem 20 *The length of a shortest S–F path in G is at most ε longer than that of an optimal path from the initial position S to the final position F.*

Proof: Let P be an optimal path from S to F. We show the existence of a path P' that corresponds to an S–F path of graph G whose length is at most $\|P\| + \varepsilon$. In view of Lemma 19, we only consider the case when P has a floating C-segment. This C-segment must occur in a subpath of P that has type LCCL, XCCL, or LCCX, where X is either an O-segment or a terminal C-segment; by Theorem 18, a floating C-segment cannot occur in an OCCO-type subpath.

We consider the OCCL-type subpath; other cases can be handled in an analogous manner. Let P_0, P_1, P_2, P_3 be the 4 segments constituting this subpath, and let C_1 and C_2 be the unit-radius circles containing the C-segments P_1 and P_2. If $(C_1, C_2) \in \mathcal{R}(\delta)$, then again, by Lemma 19, P is an S–F path in G, so assume that $(C_1, C_2) \notin \mathcal{R}(\delta)$. This assumption implies that P_3 is not a terminal segment, because otherwise C_2 satisfies property (P3) and therefore $(C_1, C_2) \in \mathcal{R}(\delta)$. Let P_{-1}, P_4 be the subpaths of P preceding P_0, and following P_3, respectively.

Without loss of generality, we can assume that P_3 is a horizontal line segment oriented in the $+x$-direction, and that P_2 is clockwise oriented (then P_0 and P_1 are clockwise and counterclockwise oriented, respectively); see Figure 14(i). For $0 \le i \le 3$, let X_i denote the common position of P_{i-1} and P_i. By Lemma 12, for $i = 1, 2$, P_i is tangent to an obstacle, say, O_i (P_0 is contained in ∂O_1); by Corollary 16, P_4 does not contain any nonterminal C-segment; and by Corollary 17, all O-segments in P_4 are counterclockwise oriented. For the sake of simplicity, let us assume that if the final segment of P_4 is a C-segment, then it is also counterclockwise oriented. This implies that P_4 is right-convex. (A slight modification of the proof works even when the terminal C-segment is clockwise oriented.)

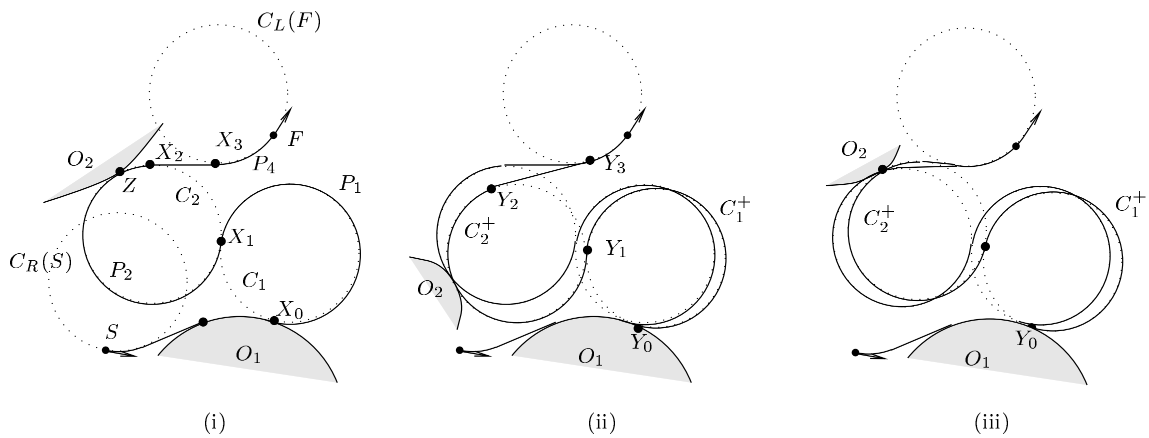

Figure 14: Approximating an optimal path in G

We choose a relevant pair $(C_1^+, C_2^+) \in \mathcal{R}(\delta)$ bridging (O_1, O_2), as follows. Roughly speaking, we roll C_1, C_2 along O_1 and O_2, keeping C_1 and C_2 in contact and moving C_1 in the clockwise direction along ∂O_1, until we find a pair $(C_1^+, C_2^+) \in \mathcal{R}(\delta)$. More precisely, for $0 \le t \le \delta$, let X^t denote the position on ∂O_1 so that $\|\partial O_1[X_0, X^t]\| = t$. Let $C_1(t) = C_L(X^t)$. There are at most two unit-radius circles tangent to both $C_1(t)$ and ∂O_2. If there is no such circle, $C_2(t)$ is undefined, and if there is one such circle then we set $C_2(t)$ to that circle. Finally, if there are two tangent circles, we define $C_2(t)$ to be the one such that $C_1(t), C_2(t)$ can be rolled back to C_1 and C_2, keeping them in contact with each other and with O_1 and O_2, respectively. Let t_0 to be the minimum value of t for which $(C_1(t_0), C_2(t_0)) \in \mathcal{R}(\delta)$; set $C_1^+ = C_1(t_0)$ and $C_2^+ = C_2(t_0)$. Property (P4) ensures that t_0 exists, and that for every $t < t_0$ and for $i = 1, 2$, $C_i(t) \notin \mathcal{C}$. Let $Z = C_2 \cap O_2$, $Z^+ = C_2^+ \cap O_2$, $Y_0 = X_1^{t_0}$, $Y_1 = C_1^+ \cap C_2^+$, $P_0' = \partial O_1[X_{-1}, Y_0]$, $P_1' = C_1^+[Y_0, Y_1]$. We set $P' = P_{-1} \circ P_0' \circ P_1' \circ P_2' \circ P_3' \circ P_4'$, where P_2', P_3', and P_4' are defined as follows. There are two cases to consider.

1. If Z^+ lies before Z along ∂O_2, then let P_3' be the tangent segment from ∂O_2 to P_4 (see Figure 14(ii)). P_3' is defined, because otherwise, by property (P3), the circle tangent

to L_F and O_2 (and lying to the right of L_F) would touch the arc $\partial O_2[Z^+, Z]$, i.e., there exists a $t < t_0$ with $C_2(t) \in \mathcal{C}$. Let Y_2, Y_3 be the initial and final positions of P_3'. Set $P_2' = C_2^+[Y_1, Y_2]$ and $P_4' = P_4[Y_3, F]$.

2. If Z^+ lies after Z along ∂O_2, then we define P_3' to be the feasible, optimal path from Z^+ to X_3 that is homotopic to the composite path $\Gamma \circ P[Z, X_3]$, where Γ is the backward path from Z^+ to Z along ∂O_2. Define $P_2' = C_2^+[Y_1, Z^+]$ and $P_4' = P_4$.

Using property (P1) (we will be need (P2) if the original path was LCCL-type) and the facts that P is feasible and obstacles are moderate, it can be shown that the new path P' is feasible. Moreover, P' is obviously an S–F path in G. Hence, it remains to prove that $\|P'\| - \|P\| \le \varepsilon$. We consider the two cases separately.

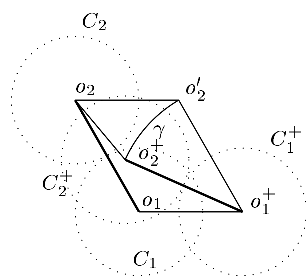

Figure 15: Bounding $|\psi(Y_1) - \psi(X_1)|$.

Case 1: Z^+ lies before Z along ∂O_2. Set $Q_1 = P[X_0, Z]$, $Q_2 = P[Z, F]$, $Q_1' = P'[Y_0, Z^+]$, and $Q_2' = P'[Z^+, F]$.

$$\|P'\| - \|P\| \le (\|P_0'\| - \|P_0\|) + (\|Q_1'\| - \|Q_1\|) + (\|Q_2'\| - \|Q_2\|). \tag{6}$$

The first term is at most δ, by construction, and the perturbation argument of Lemma 6 shows that

$$\|Q_2'\| \le \|Q_2\| + \|\partial O_2[Z^+, Z]\| \le \|Q_2\| + \delta.$$

We now bound the quantity $\|Q_1'\| - \|Q_1\|$. An easy calculations shows that

$$\|Q_1'\| - \|Q_1\| \le 2|\psi(Y_1) - \psi(X_1)| + \delta. \tag{7}$$

Let o_i, o_i^+ (for $i = 1, 2$) denote the centers of C_i, C_i^+, respectively; see Figure 15. Then $\|o_1 o_2\| = \|o_1^+ o_2^+\| = 2$. By Lemma 5,

$$\|o_1 o_1^+\| \le \|\partial O_1[Y_1, X_1]\| + \theta(\partial O_1[Y_1, X_1]) \le 2\delta. \tag{8}$$

Similarly, we have $\|o_2 o_2^+\| \le 2\delta$. Let $o_1^+ o_2'$ be the translate of the segment $o_1 o_2$, and let γ be the circular arc of radius 2, spanning from o_2^+ to o_2' in the clockwise direction, and centered at o_1^+. Then $|\psi(Y_1) - \psi(X_1)| = \angle o_2' o_1^+ o_2^+$. Hence,

$$|\psi(Y_1) - \psi(X_1)| \quad = \quad \angle o_2' o_1^+ o_2^+ \quad = \quad \frac{\|\gamma\|}{2}$$

Note that γ is a convex arc and lies inside the triangle $o_2' o_2 o_2^+$. Therefore,

$$\|\gamma\| < \|o_2' o_2^+\| + \|o_2 o_2^+\| = \|o_1 o_1^+\| + \|o_2 o_2^+\| \le 2\delta.$$

Putting everything together, we obtain

$$\|P'\| - \|P\| \le 7\delta \le \varepsilon,$$

provided that $\delta \le \varepsilon/7$.

Case 2: Z^+ lies after Z along ∂O_2. Let Y_2 be the position on C_2^+ such that $\psi(Y_2) = \psi(X_2) = 0$. Set $Q_1 = P[X_0, Z]$, $Q_2 = P[Z, X_2]$, $Q_1' = P'[Y_0, Z^+]$, $Q_2' = P'[Z^+, Y_2]$, and $Q_2' = P'[Y_2, X_3]$. Then

$$\|P'\| - \|P\| = (\|Q_1'\| - \|Q_1\|) + (\|Q_2'\| - \|Q_2\|) + (\|Q_3'\| - \|Q_3\|).$$

As in Case 1, we can argue that

$$\|Q_1'\| - \|Q_1\| \le 5\delta \quad \text{and} \quad \|Q_2'\| - \|Q_2\| \le \delta.$$

Hence, it suffices to bound $\|Q_3'\| - \|P_3\|$.

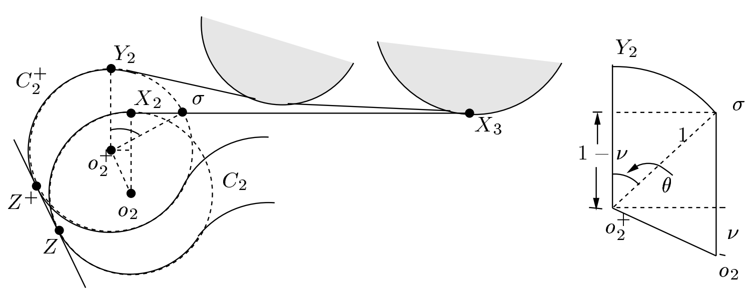

Figure 16: Bounding $\|Q_3'\| - \|P_3\|$.

Property (P2) ensures that X_3 lies outside C_2^+, therefore C_2^+ intersects P_3; see Figure 16. Let σ be the intersection point of C_2^+ and P_3, and let Σ, Σ^+ be the positions on C_2^+ and P_3 with $\textsc{loc}(\Sigma) = \textsc{loc}(\Sigma^+) = \sigma$. Then $\|Q_3'\| < \|C_2^+[Y_2, \Sigma^+]\| + \|P_3[\Sigma, X_3]\|$. Therefore,

$$\|Q_3'\| - \|P_3\| < \|C_2^+[Y_2, \Sigma^+]\| - \|P_3[X_2, \Sigma]\|. \tag{9}$$

Let o_2, o_2^+ denote the centers of C_2 and C_2^+, respectively. By (8), $\|o_2 o_2^+\| \leq 2\delta$. Let $\nu \leq 2\delta$ denote the vertical distance between o_2 and o_2^+, and let $\|C_2^+[Y_2, \Sigma^+]\| = \theta$. Then $\sin\theta = \sqrt{2\nu - \nu^2}$ and $\|P_3[X_2, \Sigma]\| \geq \sin\theta - 2\delta$; see Figure 16. Hence,

$$\|Q_3'\| - \|P_3\| \leq \sin^{-1}\sqrt{2\nu - \nu^2} - \sqrt{2\nu - \nu^2} + 2\delta$$

Using Taylor's expansion, one can prove that $\sin^{-1} x - x < x^2$ for $x \leq 1/2$, which implies $\|Q_3'\| - \|P_3\| < 2\nu - \nu^2 + 2\delta \leq 6\delta$. Putting everything together, we obtain $\|P'\| - \|P\| \leq \varepsilon$, provided $\delta \leq \varepsilon/12$. This completes the proof of the theorem. □

Next, we prove that G has $O(n^2 + 1/\varepsilon)$ vertices and edges, and that it can be constructed in time $O(n^2 \log n + (1/\varepsilon)\log(1/\varepsilon))$.

Lemma 21 *There are $O(n)$ anchored circles, and they can be computed in $O(n \log n)$ time.*

Proof: For each obstacle O_i, we define the *expanded obstacle*, O_i^*, to be the Minkwoski sum of O_i and the unit-radius disc centered at origin. The number of anchored circles is the same as the number of intersection points between the boundaries of expanded obstacles. Suppose there is a point p in the plane that lies in k expanded obstacles. Then the unit-radius disc B_p centered at p intersects k obstacles, say, $O_1, \ldots, O_k$. If $B_p \subseteq O_i$, for any $1 \leq i \leq k$, then no other obstacle can intersect B_p, because the obstacles are pairwise disjoint. Otherwise, for each $1 \leq i \leq k$, let D_i be the unit-radius disc tangent at one of the points in $\partial B_p \cap \partial O_i$ and contained in O_i. (Since each O_i is moderate, such a D_i exists.) By construction, the interiors of D_i's are pairwise disjoint, and all of them intersect B_p. This implies that $k \leq 6$, as at most 6 unit-radius discs with pairwise disjoint interiors can intersect a unit-radius disc.

Since the obstacles are pairwise disjoint, a result of Kedem et al. [18] implies that the boundaries of any pair of O_i^*'s intersect in at most two points. Plugging this observation and the above claim to a result of Sharir [37], one can show that the number of intersection points between the boundaries of expanded obstacles is $O(n)$, which proves the lemma. Finally, these intersection points, and thus the anchored circles, can be computed by a sweep-line algorithm in $O(n \log n)$ time. □

Lemma 22 *The set $\mathcal{R}(\delta)$ has $O(1/\delta)$ relevant pairs, and it can be computed in time $O(n + 1/\delta)$.*

Proof: Let D be the disc of radius 5 centered at $\textsc{loc}(S)$. In view of Theorem 9, any relevant pair of circles lies entirely inside D. Let $\mathcal{O}'$ be the set of obstacles that intersect D. Since the area of each obstacle is at least 1 and they are pairwise disjoint, only $O(1)$ obstacles belong to $|\mathcal{O}'|$. This implies that $|R_C| = O(1)$ for any circle C, because the circle C' in any relevant pair of the form (C, C') has to be tangent to an obstacle in $\mathcal{O}'$. There are $O(1)$ circles that lie in D and satisfy properties (P1)–(P4), and there are $O(1/\delta)$ circles that lie in D and satisfy (P5), $\mathcal{R}(\delta)$ has $O(1/\delta)$ circles. We can compute $\mathcal{O}'$ in $O(n)$ time, and then $\mathcal{R}(\delta)$ in additional $O(1)$ time by a brute-force approach. $\square$ *Should we say more??*

The previous two lemmas imply that $|\mathcal{A}\cup\mathcal{O}\cup\mathcal{R}| = O(n+1/\varepsilon)$. Obviously $|\mathcal{L}_1| = O(n^2)$. Since $|\mathcal{R}| = O(1/\varepsilon)$, we have $|\mathcal{L}_2| = O(1/\varepsilon)$. While computing the circles in $\mathcal{A}\cup\mathcal{R}$ and lines in $\mathcal{L}_1 \cup \mathcal{L}_2$, we can also compute the two objects touched by these circles or lines. Hence, $\mathcal{X}$ can be constructed in time $O(n^2 + 1/\varepsilon)$. Next, we need to compute the edges of G. We divide the edges into three subsets: E_1 is the set of edges that lie along the boundary of obstacles, or along the initial and final circles; E_2 is the set of edges that lie on the lines in $\mathcal{L}_1 \cup \mathcal{L}_2$; and E_3 is the edges that lie on the circles in $\mathcal{A}\cup\mathcal{R}$. E_1 can easily be computed in $O(n^2 \log n)$ time, and E_2 can be computed in $O(n^2 \log n + 1/\varepsilon)$ time by adapting any algorithm for computing the visibility graph of a set of circles; see e.g., [**?**]. Each edge $\gamma \in E_2$ is contained in a circle of $\mathcal{A}\cup\mathcal{R}$. For each $C \in \mathcal{A}$ (resp. $C \in \mathcal{R}$), we compute all intersection points between C and edges of $\mathcal{O}$ (resp. $\mathcal{O}_1$). By sorting these intersection points along C, we can compute all maximal connected portions of C that do not intersect any obstacle. An edge $\gamma \subset C$ of C is feasible if and only it lies within one such connected component. Hence, we can compute in $O(n^2 \log n + (1/\varepsilon)\log(1/\varepsilon))$ time all feasible edges of E_2.

Finally, we compute an S–F shortest path in G, in time $O((n^2 + 1/\varepsilon)\log(n + 1/\varepsilon))$, using the algorithm by Fredman and Tarjan [14]. The running time can be improved to $O(n^2 \log n + (1/\varepsilon)\log(1/\varepsilon))$ using Theorem 9 and Lemma 22, as follows. We construct two graphs: We construct G_1, in time $O(n^2 \log n)$, as above by setting $\mathcal{C} = \mathcal{O}\cup\mathcal{A}\cup\mathcal{L}_1$. We construct another graph G_2, in time $O((1/\varepsilon)\log(1/\varepsilon))$, by setting $\mathcal{C} = \mathcal{O}_1 \cup \mathcal{R} \cup \mathcal{L}_2$. We compute S–F shortest paths in both graphs, and choose the shorter of the two. The total time spent is $O(n^2 \log n + (1/\varepsilon)\log(1/\varepsilon))$. If an optimal path from S to F does not contain a floating circle, then an S–F shortest path in G_1 gives such a path. Otherwise, by Theorem 9, the length of an S–F shortest path is at most ε more than that of an optimal path from S to F. This completes the proof of Theorem 1(i).

As a consequence of Theorem 9, if the initial and the final locations are sufficiently far apart then an optimal path between them does not contain a floating C-segment, so we can construct it by computing a S-F path in G_1. This proves Theorem 1 (ii).

4 Paths with Reversals

In this section we study the problem of obtaining an approximately optimal path when reversals are allowed. Let P be a feasible path and X a position. We employ the following approach. First note that if the orientation vectors of the initial and final positions are not specified then the optimization problem can be solved easily by the standard method.

Lemma 23 *The shortest obstacle-avoiding path from* LOC(S) *to* LOC(F) *can be found in time* $O(n^2 \log n)$*. Moreover, the path found is moderate.*

Sketch: We solve the shortest path problem in the network of tangents and obstacle boundaries in the standard way. Since the solution path consists of line segments and arcs along obstacle boundaries, with each line segment tangent to the arcs adjacent to it on the path, it is moderate. □

Our approach is to use the path from LOC(S) to LOC(F) as above, modifying it in the neighborhood of S and F so as to make it feasible. In the next subsection, we study some special (but typical, in practical sense) cases where we can achieve this modification with a small number of reversals and an additive constant term in the path length. In the subsequent subsections, we will turn to more subtle cases where indefinitely many number of reversals are required in this approach.

4.1 Simple cases

We start with some lemmas providing the way to adjust the orientation vector of a position by a path that ends with the same location as it starts with.

Lemma 24 *Let* X *and* Y *be positions such that* LOC$(X) =$ LOC(Y) *and no obstacle is within distance* 1 *of* LOC(X)*. Then, there is a* 3*-feasible path from* X *to* Y *of length at most* π*.*

Proof: A path consisting of three C-segments as in Figure 17 satisfies the requirement. □

When the free space available around LOC(X) is smaller, we may repreat the path of the above type as many times as necessary, to obtain the following.

Lemma 25 *Let* X *and* Y *be positions such that* LOC$(X) =$ LOC(Y) *and no obstacle is within distance* $\epsilon \leq 1$ *of* LOC(X)*. Then, there is a* k*-feasible path from* X *to* Y *of length at most* π *such that* $k \leq 3\lceil 1/\epsilon \rceil$*. The path can be computed in time* $O(k)$*.*

Theorem 26 *Suppose that the distance from* LOC(S) *to any obstacle is at least* $1/k_1$ *and the distance from* LOC(F) *to any obstacle is at least* $1/k_2$*. Then, in time* $O(n^2 \log n + k_1 + k_2)$ *we can construct a* $3(k_1 + k_2)$*-feasible path from* S *to* F *with length at most* $l + 2\pi$*, where* l *is the length of the shortest obstacle-avoiding path from* LOC(S) *to* LOC(F)*.*

If we want our solution to have only a constant number of reversals, then the above theorem works only if both LOC(S) and LOC(F) are some constant-distance away from obstacles. In the following, we study some cases where an obstacles is arbitrarily close

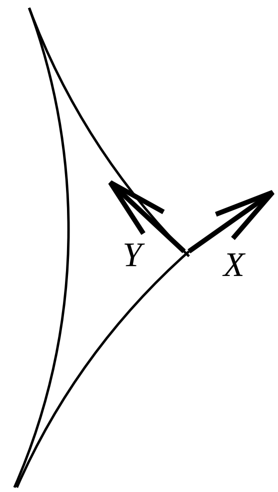

Figure 17: A 3-feasible path for rotation

to the initial or final position and yet we can efficiently find an near-optimal path with a constant number of reversals.

We say that a position X is *blocked* by an obstacle O *from the front*, if every unbounded moderate simple path that goes forward from X strictly intersects O. Equivalently, X is blocked by O from the front if there is a bounded region whose boundary consists of an arc of $C_R(X)$ starting from X, an arc of $C_L(X)$ starting from X, and a portion of the boundary of O. Similarly, X is blocked by O *from the back*, if every unbounded moderate simple path that goes backward from X strictly intersects O. If a position X is blocked by O_1 from the front and O_2, then we say that X is *doubly blocked* by the *blocking pair* (O_1, O_2). The goal of the rest of this subsection is to prove the following theorem.

Theorem 27 *Suppose neither the initial nor the final position is doubly-blocked. Then, in time* $O(n^2 \log n)$ *we can find a 7-feasible path of length at most* $l + 7\pi$ *from* S *to* F*, where* l *is the length of the shortest obstacle-avoiding path from* $\textsc{loc}(S)$ *to* $\textsc{loc}(F)$*.*

We need a few lemmas for the proof. Let us say that a position X is tangent to an obstacle O if $\textsc{loc}(X)$ is on the boundary of O and the line L_X is tangent to O. We next consider the case where the initial and final positions are tangent to obstacles.

Lemma 28 *Suppose both the initial position* S *and the final position* F *are tangent to obstacles. Let* l *be the length. of the shortest obstacle-avoiding path from* $\textsc{loc}(S)$ *to* $\textsc{loc}(F)$*. Then, a 3-feasible path from* S *to either* F *or* $\textsc{rev}(F)$ *of length at most* $l + 3\pi$ *can be constructed in time* $O(n^2 \log n)$*.*

Proof: Let O_S and O_F be the obstacles S and F are on respectively. Let P be a shortest obstacle-avoiding path from $\textsc{loc}(S)$ to $\textsc{loc}(F)$ obtained as in Lemma 23. Then

P is moderate and consists of line segments and arcs along obstacle boundaries. Suppose first that the length of P is at least 2. Let X be the first position on P such that either $C_L(X)$ or $C_R(X)$ is tangent to O_S. We claim that such an X always exists and moreover the length of the arc along P from $\textsc{loc}(S)$ to $\textsc{loc}(X)$ is at most 1. This is certainly the case if P is tangent to O_S at $\textsc{loc}(S)$, because then $\textsc{loc}(X) = \textsc{loc}(S)$. Otherwise, P must start with a line segment; let us call this line segment L. If the length of L is at least 1, then we can draw a unit-radius circle tangent to L and O_S, confirming the claim. So suppose that the length of L is smaller than 1 and let p be the endpoint of L opposite $\textsc{loc}(S)$. At p, P switches from L to an arc around the boundary of some obstacle O, which contains a unit-radius circle C' tangent to P at p. Because C' does not intersect with the interior of O_S, $\textsc{loc}(X)$ must lie between $\textsc{loc}(S)$ and p on P, confirming the claim again.

Similarly, let Y be the last position on P such that either $C_L(Y)$ or $C_R(Y)$ is tangent to O_F. Then the length of the arc from $\textsc{loc}(Y)$ to $\textsc{loc}(F)$ along P is at most 1. Our path P' starts with an arc around O_S from $\textsc{loc}(S)$ to the point at which $C_R(X)$ or $C_L(X)$ is tangent to O_S, then reverses to $\textsc{loc}(X)$ along $C_R(X)$ or $C_L(X)$, continues to $\textsc{loc}(Y)$ along P, to the boundary of O_F along $C_R(Y)$ or $C_L(Y)$, and then finally reverses along the boundary of O_F to F. The moderate scene assumption ensures that this path is obstacle-avoiding and hence 3-feasible. Being generous, it is easy to confirm that the length of P' is at most 2π greater than that of P.

Now suppose that the length of P is smaller than 2. In this case, we provide the following alternative path. Draw a unit-radius circle C tangent to O_S at $\textsc{loc}(S)$. Of the two intersection points between C and the boundary of O_F, let q denote the one closer to $\textsc{loc}(S)$. Let A denote the shorter arc of C from $\textsc{loc}(S)$ to q. Suppose first that A intersects with some obstacle other than O_S or O_F. Then, from the moderateness of this obstacle, we may draw a unit radius circle C' tangent to both O_S and O_F that intersects with A. Let p_1 and p_2 be the points at which C' is tangent to O_S and O_F respectively. Our path P' in this case starts with the arc from $\textsc{loc}(S)$ to p_1 along the boundary of O_S, reverses along C' to p_2, and finishes with an arc along the boundary of O_F to $\textsc{loc}(F)$. From the moderate scene assumption, P' is obstacle avoiding and hence 3-feasible. It is easy to confirm that the length of P' is at most $4 + \pi$ larger than that of P, due to the condition that C' intersects A. Suppose next that A does not intersect with any obstacle except at the endpoints. Draw a unit-radius circle C'' tangent to A and O_F. Note that A is long enough to allow such a common tangent: otherwise O_S would intersect the interior of O_F. Let p_1 and p_2 be the points of tangency on A and O_S respectively. Our path P' in this case consists of the part of arc A from $\textsc{loc}(S)$ to p_1, the shorter arc from p_1 to p_2 along C'', and then to $\textsc{loc}(F)$ along the boundary of O_F. Then, P' is obstacle-avoiding due to the moderate scene assumption. Since the length of the arc from $\textsc{loc}(S)$ to p_1 is at most $2\pi/3$ and that of the arc from p_1 to p_2 is at most $\pi/2$, the increase of the length when replacing P by P' is at most $2(2\pi/3 + \pi/2) = 7\pi/3$. □

Lemma 29 *Let X be a position that is not doubly-blocked. Suppose furthermore that there is an obstacle within distance 1 of $\textsc{loc}(X)$. Then, in time $O(\log n)$ we can find a*

2-feasible path of length at most π from X to some position Y tangent to some obstacle.

Proof: Let O be the obstacle whose distance from $\textsc{loc}(X)$ is at most 1 and moreover the smallest among all such obstacles. First suppose that neither $C_L(X)$ nor $C_R(X)$ intersects the interior of O. Then, we can draw a moderate simple path of type LC from X (either forward or backward) to a position tangent to O. This path cannot be obstructed, due to the assumption that O is the obstacle closest to $\textsc{loc}(X)$. The length of the path is clearly smaller than π. Next suppose exactly one of $C_R(X)$ and $C_L(X)$ intersects the interior of O; without loss of generality we assume that $C_R(X)$ does but $C_L(X)$ does not. Then, there is a position Z on $C_R(X)$ on the shorter arc between $\textsc{loc}(X)$ and O such that $C_L(Z)$ is tangent to O, providing a path of type CC that satisfies the requirement of the lemma. Finally suppose that both $C_R(X)$ and $C_L(X)$ intersect O. Since the distance of O from $\textsc{loc}(X)$ is at most 1, this implies that X is blocked by O either from the front or from the back; assume it is from the front without loss of generality. Let A and B be backward paths of length $\pi/3$ from position X along $C_L(X)$ and $C_R(X)$ respectively. By the assumption that X is not doubly blocked, no obstacle intersects both A and B in its interior. If A intersects with the interior of some obstacle O', then B does not with O' and hence we may apply the above construction for the second case to O' instead of O. Finally, suppose A does not intersect with the interior of any obstacle. Then, there must be a position Z' on A such that $C_R(Z')$ is tangent to O, providing a path with one reversal that satisfies the requirement. □

Combining Lemmas 24, 29, and the proof of 28, we can now prove Theorem 27.

Proof of Theorem 27: If there is no obstacle within distance 1 from S, then by Lemma 24, we may adjust the orientation of $\psi(S)$ aribitrarily with at most 2 reversals. Otherwise, by Lemma 29, we may get tangent to some obstacle from S with at most 2 reversals. Applying a similar observation to F, we reduce the problem to that of finding a 1^+-feasible path from S' to F' for some positions S' and F'; either S' is tangent to an obstacle or its orientation is unspecified, and similarly for F'. Similarly to the proof of Lemma 28 we can construct a 3-feasible path for this problem, which is at most 3π longer than the shortest obstacle-avoiding path from S' to F'. Since the lengths of our paths from S to S' and from F' to F is at most π each, our combined path has length at most 7π longer than the optimal. □

4.2 Unblocking

We now focus on the case where either S or F is doubly-blocked. In this situation, one important component of our solution is the method of *unblocking.* Suppose a position X is doubly blocked, by obstacle O_1 from the front and by O_2 from the back. We call a feasible path from X *unblocking* if it contains a position not doubly blocked. We sometimes say that such a path *unblocks.* An unblocking path is *minimal* if every position on the path except for the final is doubly blocked. A *greedily rotating path* is a feasible path such that all its cusps are on obstacle boundaries and its component paths are unit-radius arcs, with each pair of consecutive arcs tangent to each other. See Figure 18. Thus, the rotation of the orientation vector along a greedy rotating path is at the maximal allowed rate and in a consistent direction. Note that, given an initial point, there are 4 ways to start a greedy rotating path: either forward or backward and either clockwise or counterclockwise. We call a minimal unblocking path *greedy* if it is greedily rotating. The goal of this subsection is to prove the following theorems.

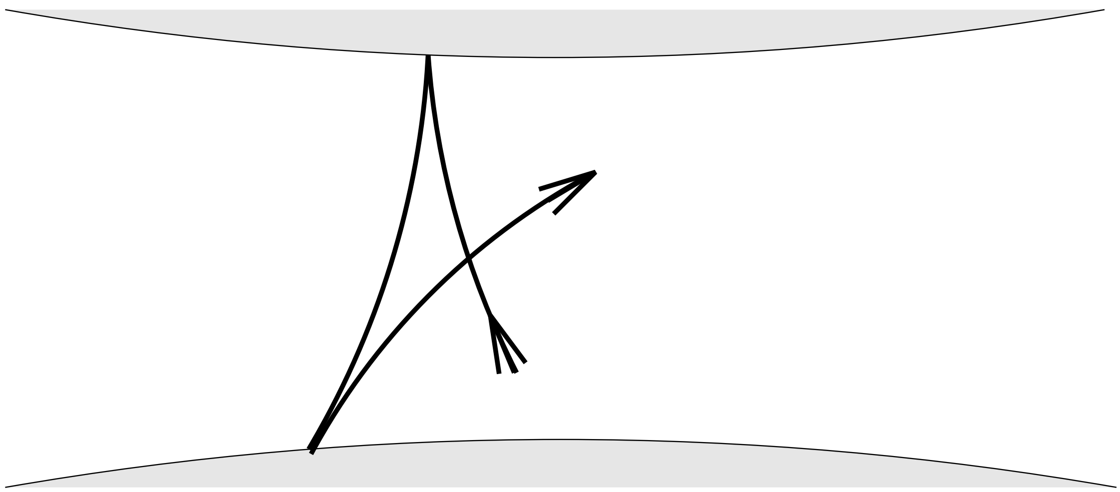

Figure 18: A greedily rotating path

Theorem 30 *Let X be a doubly-blocked position such that a k-feasible unblocking path exists from X. Then there exists a ck-feasible greedy unblocking path from X, where c is a universal constant.*

Theorem 31 *Let X and Y be positions doubly-blocked by the same pair of obstacles such that there exists a k-feasible path from X to Y. Then, there exists a ck-feasible greedily rotating path from X, where c is the constant in the previous theorem, whose final position is some Y' with $\psi(Y') = \psi(Y)$.*

The proof of Theorem 30 is accomplished through a series of lemmas.

Lemma 32 (Butterfly Lemma) *Let X and Y be two positions such that both $\textsc{loc}(X)$ and $\textsc{loc}(Y)$ are on the x-axis, with $\textsc{loc}(X)$ to the left of $\textsc{loc}(Y)$, and both $\psi(X)$ and $\psi(Y)$ have positive y–component. Suppose that $C_L(X)$ and $C_L(Y)$ have an intersection p above the x-axis and that $C_R(X)$ and $C_R(Y)$ have an intersection q above the x-axis. Define a new coordinate system (x', y') as follows. Let the line through the centers of $C_L(X)$ and of $C_R(Y)$ be the x'-axis and let the perpendicular bisector of the line segment between these* *two centers be the y'-axis, with its positive half-axis extending to the direction in which y–coordinate increases. Let h_1 and h_2 be the maximum y'–coordinates of a point on the arc from* LOC(Y) *to p along $C_L(Y)$ and the arc from* LOC(X) *to q along $C_R(X)$, respectively. Then $h_1 > h_2$ if the slope of the x'-axis against the x-axis is negative.*

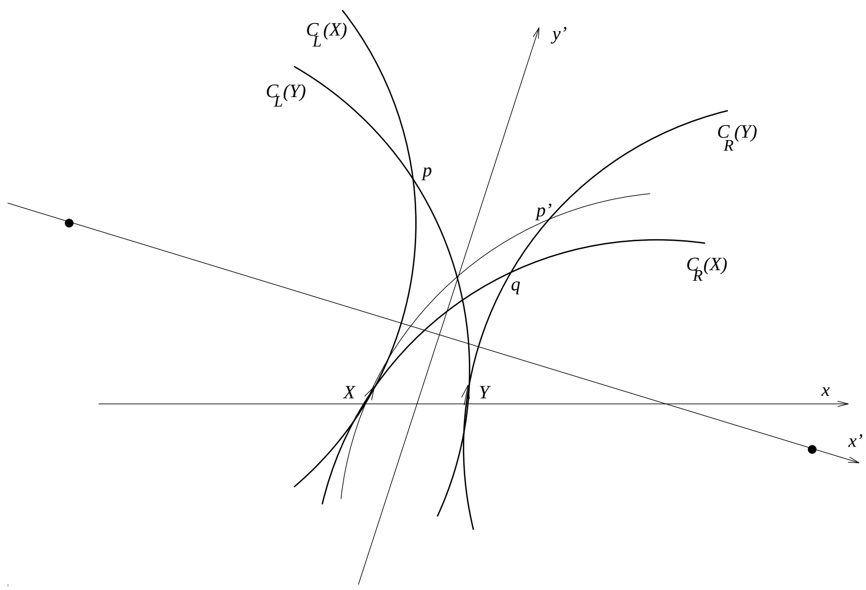

Figure 19: Butterfly Lemma

Proof: In Figure 19, let p' be the mirror image of p with respect to the y'-axis. Since the mirror image of $C_L(Y)$ is obtained by rolling $C_R(X)$ on $C_L(X)$ counterclockwise, the mirror image of the arc from LOC(Y) to p is clearly above the arc from LOC(X) to q. □

Let O_1 and O_2 be two obstacles. We define the *standard coordinate system determined by O_1 and O_2* as follows. Let point p_i be on the boundary of O_i, $i = 1, 2$, so that the distance between p_1 and p_2 is minimized. When such a choice is not unique, then p_i must be on a line segment L_i that is a part of the boundary of O_i, $i = 1, 2$, such that L_1 and L_2 are opposing edges of a rectangle: we take each p_i so as to bisect L_i. We take the line through p_1 and p_2 as the y-axis and the perpendicular bisector of the line segment between p_1 and p_2 as the x-axis. The orientations of the axes are chosen so that O_1 is in the half plane with positive y-coordinates.

Let X be a position, O an obstacle, and p a point on the boundary of O. We say that position X *injects* into O at p, if the line L_X goes through p and the line segment between LOC(X) and p does not intersect any obstacle except at p. Suppose a position X injects into an obstacle O at p and let L denote the line perpendicular to the boundary of O at p. The *injection angle* of X against O is defined to be the amount θ of counterclockwise rotation of L, $-\pi/2 < \theta < \pi/2$, that makes L coincide with L_X.

I changed the def of injection angle. HT

Let X be a position doubly blocked by a blocking pair (O_1, O_2). We say that X is *rightward-inclined* if $\psi(X) \leq \pi/2$, where angle $\psi(X)$ is with respect to the standard coordinate system determined by O_1 and O_2; X is *leftward-inclined* if $\psi(X) \geq \pi/2$. We call a path P *rightward-inclined* (*leftward-inclined*), respectively) if every position on P, except possibly for the final position, is doubly blocked and rightward-inclined (leftward-inclined, respectively). A rightward-inclined (leftward-inclined) path is called *pure* if the injections angles of each position on the path into the blocking obstacles are always both negative (positive, respectively). Note that, in this situation, two of the four greedy unblocing paths from X are pure. The following lemma is the heart of the proof of Theorem 30.

Lemma 33 *Suppose that a position X is doubly blocked and that the injection angles of X against the two blocking obstacles are of the same sign. Then, among all pure minimal unblocking paths from X, at least one of the two greedy unblocking paths has the minimum number of reversals.*

Proof: Let O_1 and O_2 be the obstacles blocking X from the front and the back respectively. Assume without loss of generality that the injection angle of X agains each of O_1 and O_2 is negative. Thus, pure unblocking paths from X rotate the orientation vector clockwise. Let P be the minimal unblocking path from X that is clockwise greedily rotating and starts forward, and suppose P has k cusps. We show that any pure minimal unblocking path from X that starts forward has k or more cusps. Since a similar claim holds for the greedy unblocking path that starts backward, the statement in the lemma will follow.

Let X_i, $1 \leq i \leq k$, be the position on P at its ith cusp, and P_i the ith component path of P. Let Q be an arbitrary pure minimal unblocking path from X that starts forward, let Y_i, $1 \leq i \leq k'$, be the position on Q at its ith cusp, and Q_i the ith component path of Q, where k' is the number of cusps of Q. We may assume that all cusps of Q are on the boundaries of O_1 and O_2, because if a cusp is not on an obstacle boundary, we may add a straight-line round-trip excursion from the cusp to one of the obstacle boundaries. Thus, both $\textsc{loc}(X_i)$ and $\textsc{loc}(Y_i)$ are on O_1 if i is odd and O_2 if i is even. Note also that P_i is along $C_R(X_i)$ if i is odd and along $C_L(X_i)$ if even. For odd i, $1 \leq i \leq k$, let L_i denote the line tangent to O_1 at $\textsc{loc}(X_i)$ and H_i the open half plane bounded by L_i and disjoint from O_1. Similarly, for even i in the range, let L_i denote the line tangent to O_2 at $\textsc{loc}(X_i)$ and H_i the open half plane bounded by L_i and disjoint from O_2. For convenience, define half plane H_0 as follows. Let L'_0 be the line tangent to O_2 at the intersection of L_X and O_2, let L_0 be the line through $\textsc{loc}(X)$ parallel to L'_0, and then let H_0 be the half plane bounded by L_0 and disjoint from O_2.

We show by induction on i that the following four conditions hold for $1 \leq i \leq k$. The result immediately will follow from the first condition for $i = k$.

1. Cusp Y_i exists, i.e., $i \leq k'$.

2. The shorter arc from $\text{LOC}(X_i)$ to $\text{LOC}(Y_i)$, along O_1 if i is odd and O_2 if i is even, is oriented clockwise, unless these two points are identical.

3. Either P_i and Q_i are identical or P_i intersects with $C_R(Y_i)$ if i is odd and with $C_L(Y_i)$ if i is even.

4. If i is odd then the center of $C_R(X_i)$ is in the complement of H_{i-1} and the center of $C_L(Y_i)$ is in H_{i-1}; if i is even then the center of $C_L(X_i)$ is in the complement of H_{i-1} and the center of $C_R(Y_i)$ is in H_{i-1}.

BASIS : We show that the induction hypothesis holds for $i = 1$. The first condition of the induction hypothesis is trivial since X is blocked from the front. If $C_R(Y_1) = C_R(X)$, which means that the first component path of Q is identical to that of P and hence $Y_1 = X_1$, the second and the third conditions are also trivially satisfied. The fourth condition is also satisfied, due to our assumption that the injection angle of X against O_2 is negative. See Figure 20 (a). To deal with the case $C_R(Y_1) \neq C_R(X)$, consider rolling the circle $C_R(Z)$ starting with $Z = X$ and ending with $Z = Y_1$. When this circle "rolls away" from $C_R(X)$, it does so in the counterclockwise direction, creating an intersection with $C_R(X)$ in the arc segment between $\text{LOC}(X)$ and $\text{LOC}(X_1)$. See Figure 20 (b). This intersection monotonically advances towards $\text{LOC}(X_1)$ and stays between $\text{LOC}(X)$ and $\text{LOC}(X_1)$ (otherwise Q would unblock without a reversal at all), establishing the third condition. The intersection of the moving circle with the boundary of O_1, that is initially at $\text{LOC}(X_1)$ moves around O_1 monotonically in the clockwise orientation, establishing the second condition. Finally, for the fourth condition, observe first that the center of $C_R(X_1)$ is the same as that of $C_R(X)$ and hence is in the complement of H_0. For the center of $C_L(Y_1)$, consider again rolling the circle $C_L(Z)$ as position Z moves from X to Y_1 along Q: since the component path of Q is moderate, the center of the circle never becomes closer to L_0 and hence stays in H_0, establishing the fourth condition.

INDUCTION : Suppose that the induction hypothesis is satisfied for i, $1 \leq i < k$. We prove that the hypothesis is satisfied for $i + 1$. We assume that i is even. The case where i is odd is similar. We also assume that P_i and Q_i are not identical; when P_i and Q_i are identical, the induction step is proved similarly to the base case. See Figure 21, where M_i denotes the straight line through the center of $C_L(X_i)$ and the center of $C_R(Y_i)$ and N_i denotes the straight line through $\text{LOC}(X_i)$ and $\text{LOC}(Y_i)$. Since Q is pure, the injection angle of position Y_i against O_2 is negative. Therefore, the slope of M_i against N_i is negative as shown in the figure. By the induction hypothesis, $C_L(Y_i)$ has an intersection point Z_i with P_i. Thus, we may apply the Butterfly Lemma to see that $C_R(Y_i)$ intersects with $C_R(X_i)$ at a point Z_i' that is closer to M_i than Z_i. We claim that Z_i' is on P_{i+1}. To see this, first observe that the slope of L_{i-1} against M_i is positive, due to condition 3 of the induction hypothesis. Therefore, since Z_i is in H_{i-1} also by the induction hypothesis, we have Z_i' in H_{i-1} and hence on P_{i+1}. Thus, $C_R(Y_i)$ intersects with O_1 and contains $\text{LOC}(X_{i+1})$ in its interior. This proves the first and the second conditions of the induction hypothesis for $i+1$, since Q_{i+1} stays outside of $C_R(Y_i)$ and must therefore intersect with the boundary of

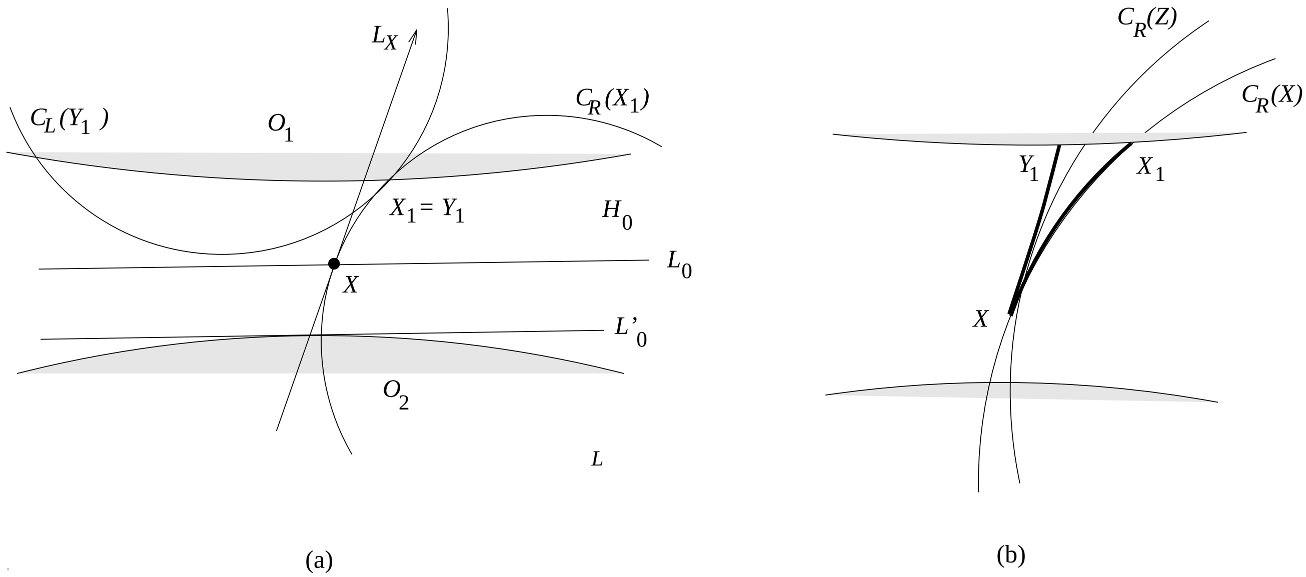

Figure 20: Base case of Lemma 33

O_1 at a point to the left of X_{i+1}. This in turn implies that P_{i+1} intersects with Q_{i+1} and hence with $C_R(Y_{i+1})$, establishing the third condition. That the fourth condition holds for $i+1$ is shown similarly to the base case. □

The following lemma can be proved by a similar induction argument. The *progress* of a path from X to Y that is leftward-inclined (rightward-inclined) is $\psi(X) - \psi(Y)$ ($\psi(Y) - \psi(X)$, respectively). This definition is intended to measure the progress toward unblocking from position X.

Lemma 34 *Let P be an arbitrary pure k-feasible path from position X with progress θ. Then, there is a k-feasible greedily rotating path that is also pure and have progress θ or larger.*

Suppose X is doubly blocked by a blocking pair (O_1, O_2). We denote by $\text{SEG}(X)$ the maximal line segment of the line L_X that contains X and does not intersect the interior of O_1 or O_2. We denote by $\text{GAP}(X)$ the length of the maximal line segment through $\text{LOC}(X)$ and parallel to the y-axis of the standard coordinate system determined by O_1 and O_2, that does not intersect the interior of O_1 or O_2.

Lemma 35 *Let X and Y be two positions such that both $\text{LOC}(X)$ and $\text{LOC}(Y)$ are on the x-axis and both $\psi(X)$ and $\psi(Y)$ have positive y-component. Let P be an arbitrary moderate path from X to Y with exactly one cusp p, which lies in the $y+$ half-plane. Let d denote the y-coordinate of p, v_1 the vector leading from $\text{LOC}(X)$ to p, and v_2 the vector*

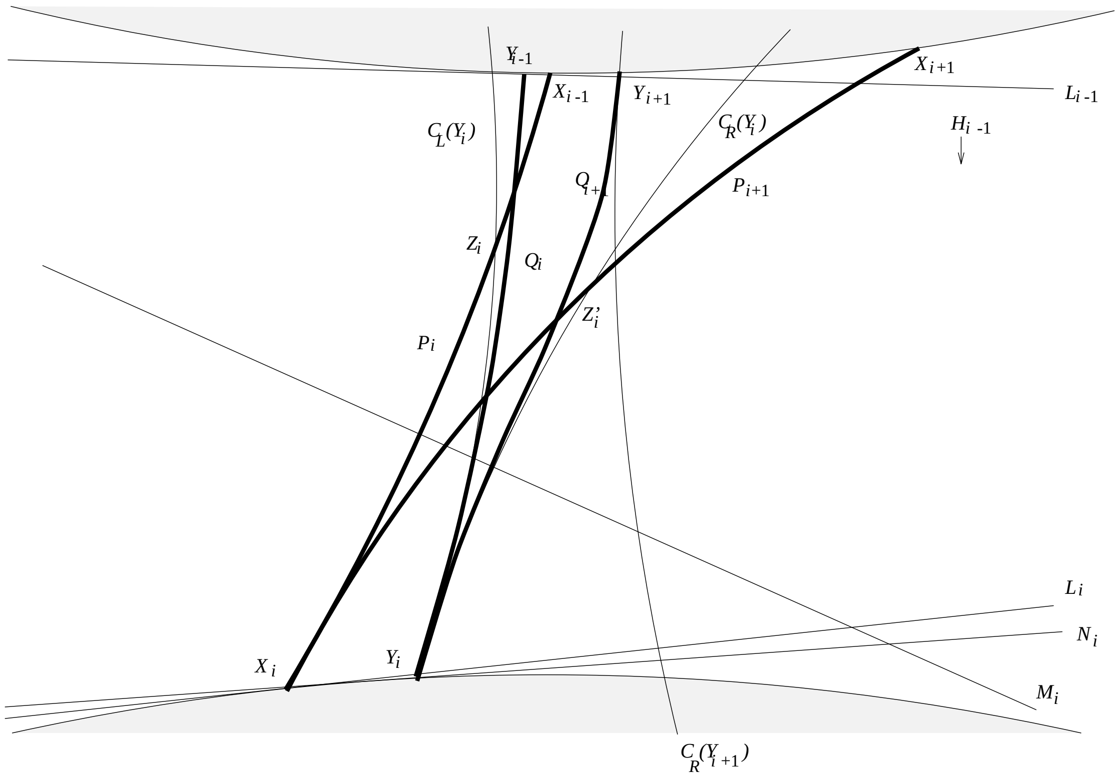

Figure 21: Induction step of Lemma 33

leading from LOC(Y) *to* p. *Let* $\psi_0 = \min\{\psi(v_1), \pi - \psi(v_1), \psi(v_2), \pi - \psi(v_2)\}$ *and suppose the arc-length of* P *is at most* $\psi_0/20$. *See Figure 22. Then, the distance between* LOC(X) *and* LOC(Y) *is at most* $5d^2 \sin^{-3}\psi_0$.

Proof: If we let (x, y) denote the coordinates of LOC(W) for a position W on P and let ψ denote $\psi(W)$, we have $\frac{dx}{dy} = \cot\psi$. To estimate the maximum difference of $\frac{dx}{dy}$ along the path, we estimate $\frac{d^2x}{dy^2}$. Let s denote the arc-length of the part of P from X to W. When W moves along each component path of P, we have:

$$\begin{aligned}
|\frac{d^2x}{dy^2}| &= |\frac{d}{dy}\cot\psi| \\
&= |\frac{d\psi}{dy}\frac{d\cot\psi}{d\psi}| \\
&= |\frac{ds}{dy}\frac{d\psi}{ds}\sin^{-2}\psi| \\
&= |\frac{d\psi}{ds}\sin^{-3}\psi| \\
&\leq \sin^{-3}(19\psi_0/20),
\end{aligned}$$

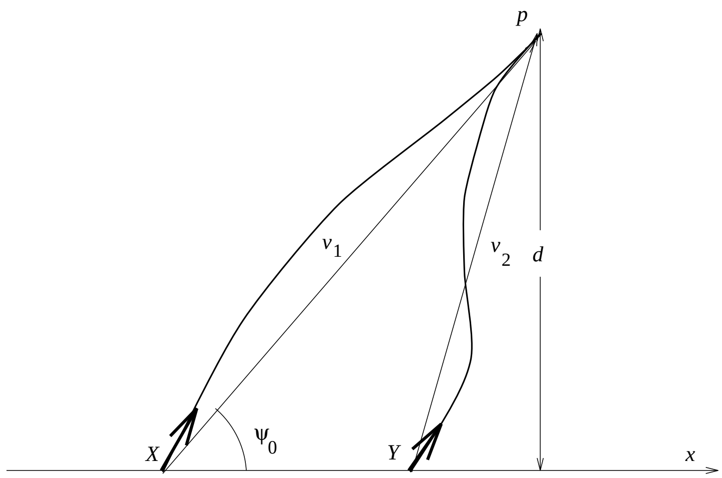

Figure 22: Lemma 35

where we used, in the last inequality, the assumption on the length of P and the fact that each component path of P is moderate. It follows that the difference between the minimum and the maximum values of $\frac{dx}{dy}$ along the path P is at most $2d\sin^{-3}(19\psi_0/20)$ and hence the distance between $\textsc{loc}(X)$ and $\textsc{loc}(Y)$ is at most $4d^2\sin^{-3}(19\psi_0/20) \leq 5d^2\sin^{-3}\psi_0$. □

Proof of Theorem 30: Suppose X is doubly blocked by O_1 and O_2. We use the standard coordinate system determined by O_1 and O_2. Assume without loss of generality that $\psi(X) \leq \pi/2$. Let P be a minimal unblocking path from X with the minimum number k of component paths. If P is pure, then we are done because, by Lemma 33, there is a greedily rotating path with k component paths that unblocks. So, suppose P is not pure. In the following, we give a constant c that satisfies the statement of the theorem for $k \geq 200$. This suffices, since it is clear that there is a constant that satisfies the statement for $k < 200$.

Let P' denote the maximal initial subpath of P that is pure. Note that P' is empty if position X is not pure. Let X' be the final position of P'. Let us first consider the simple case where $\textsc{gap}(X') > 1/(10k)$. A crude estimate shows that each component path (except for the first and the last ones) of the greedy clockwise rotating path from X' (that starts either forward or backward) has length at least $\textsc{gap}(X')/2 \geq 1/(20k)$. Therefore, $\pi/\textsc{gap}(X) + 2 + k \leq 70k$ component paths suffice to unblock from X by a pure unblocking path: P' followed by the greedy unblocking path from X'. Thus, taking $c > 70$ satisfies the theorem in this case. In the following, we assume that $\textsc{gap}(X') \leq 1/(10k)$.

Assume without loss of generality that the x-coordinate of $\textsc{loc}(X')$ is positive. Let p be the endpoint of $\textsc{seg}(X')$ on O_2 and let θ be the signed angle formed by the line tangent to O_2 at p against the x-axis, $-\pi/2 \leq \theta \leq \pi/2$. Then, since X' is not pure, we must have $\theta < 0$ and moreover $\psi(X') - \theta > \pi/2$. But, since $\textsc{gap}(X') \leq 1/(10k) \leq 1/2000$, we have

$\theta < 1/40$ and hence $\psi(X') \geq \pi/3$ as a very crude bound.

For each position Z on P, let $x(Z)$ denote the x-coordinate of the intersection of L_Z with the x-axis. Let Y be the first position on P after X' such that one of the following conditions holds.

1. $|\psi(Y) - \psi(X')| \geq \pi/6$;

2. $|x(Y) - x(X')| \geq 1/k$;

3. $\textsc{gap}(Y) \geq 1/(5k)$.

We show that in fact none of these conditions can happen. Let P_Y denote the subpath of P from X' to Y. First suppose that condition (1) happens. Then, since $\psi(Z) \geq \pi/6$ and $\textsc{gap}(Z) \leq 1/(5k)$ for every position Z on P_Y, each component path of P_Y has length at most $2/(5k)$ and therefore P_Y must have at least $(\pi/6)/(2/(5k)) > 1.3k$ component paths, a contradiction. Now suppose that condition (2) happens. Consider the path P'_Y obtained from P_Y by adding a segment of $L_{X'}$ at the beginning and a segment of L_Y at the end so that P'_Y is a k'-feasible path from a point on the x-axis to a point on the x-axis, where $k' \leq k+2$. Consider the intersection points of P with the x-axis. By Lemma 35, the distance between two such points that are consecutive along P is at most $(5k)^{-2} \sin^{-3} \pi/6 \leq 1/3k^2$. Therefore, P'_Y must have at least $3k$ component paths, a contradiction. Finally, suppose condition (3) happens. Because $\textsc{gap}(X') \leq 1/(10k) \leq 1/2000$, the slope of the boundary of each of O_1 and O_2 is at most $1/40$ in absolute value, when measured at the point with x-coordinate being equal to that of $\textsc{loc}(X)$. Let Z be an arbitrary position on P_Y. Since $|x(Z) - x(X')| \leq 1/k$, $\pi/6 \leq \psi(Z) \leq 5\pi/6$, and $\textsc{gap}(Z) \leq 1/(5k)$, the difference between the x-coordinates of $\textsc{loc}(Z)$ and $\textsc{loc}(X')$ is at most $1/k + 4/(5k) \leq 1/100$. Therefore, the slope of the boundary of each of O_1 and O_2 is at most $1/25$ in absolute value, when measured at any point with x-coordinate being equal to that of some point of P_Y. But this implies that $\textsc{gap}(Y) \leq \textsc{gap}(X') + (2/25)(1/k) < 1/(5k)$, a contradiction. Therefore, there cannot be a position Y on P after X' satisfying any of the conditions (1), (2), (3), and hence P cannot be unblocking, a contradiction. □

The proof of Theorem 31 is similar and is omitted.

4.3 Doubly-blocked terminal positions: weak-sense approximation

We now deal with the case where S or F is doubly blocked. We first consider the optimization problem in weak sense, i.e., we are required to find a path from S to either F or $\textsc{rev}(F)$. In this subsection, "k-optimal" is to be read in weak sense.

We first describe the construction of *basic paths* for the given problem instance. If S is doubly-blocked then let P_S be a minimal unblocking path from S that is greedily rotating and let S' be the final position of P_S. If S is not doubly-blocked then let P_S be an empty path and let $S' = S$. Similarly, if F is doubly blocked then let P_F be a minimal unblocking path from F that is greedily rotating and let F' be the final position of P_F; P_F is emtpy and $F' = F$ if F is not doubly blocked. Let P be the 7-feasible path from S' to F' obtained by Theorem 27. A basic path is constructed by concatenating P_S, P, and the reversal of P_F. Note that there are up to 16 basic paths for each given instance, depending on the four choices of the unblocking paths for each of S and F.

Lemma 36 *Suppose that a k-optimal path (in the weak sense) has arc-length l and that it unblocks. Then there is a $(ck+6)$-feasible basic path with length at most $l + 9\pi$, where c is the constant in Theorem 30.*

Proof: Since there are unblocking paths from S and from F with k reversals in total, the number of reversals in P_S and P_F together is at most ck, by Theorem 30, for an appropriate choices of the greedy rotating paths P_S and P_F. At most 6 more reversals are necessary in applying Theorem 27. For the length bound, the overhead of 7π comes from Theorem 27 and 2π comes from Theorem 30, π for the unblocking path at each end. □

When the k-optimal path does not unblock, we need to provide an alternative path. Consider the following example. We have a long narrow corridor between two parallel walls, with $\psi(S) = \psi(F)$ almost perpendicular to the wall. Let w be the width of the corridor and d the distance between $\text{LOC}(S)$ and $\text{LOC}(F)$. We can fix k and make w and d approach 0 in such a way that there is always a k-feasible path from S to F. It is clear that the basic path for this example is not the approximation we want, since the number of reversals required for unblocking from S tends to ∞.

Technically, we say that a problem instance is *non-basic*, if

1. S and F are doubly blocked by the same blocking pair, and
2. no $c(k + 200)$-feasible greedily rotating path from S or F unblocks, where c is the constant in Theorems 30 and 31.

Otherwise, we say the instance is *basic*. Note that if an instance is basic then one of the basic paths provides the approximation claimed in Theorem 2, in the weak sense. This is because if there is a $c(k + 200)$-feasible unblocking path from, say, S, then there is a $((c + 1)k + 200c + 1)$-feasible unblocking path from F and therefore we can unblock from both ends by greedily rotating paths with at most $c((2c+1)k+400c+1)$ reversals in total.

The following lemma provides a tool for constructing a path for a non-basic instance.

Lemma 37 *Let X be a position doubly blocked by a pair (O_1, O_2) such that $\text{LOC}(X)$ is on the x-axis of the standard coordinate system determined by this pair. Let l be the length of $\text{SEG}(X)$. Suppose $l \leq 1/10$. Then, for any position Y such that $\psi(Y) = \psi(X)$, $\text{LOC}(Y)$ is on the x-axis, and the distance between $\text{LOC}(X)$ and $\text{LOC}(Y)$ is at most $l^2/(200\psi(X))$, we can find a 2-feasible path from X to Y with length at most $3l$ in constant time.*

Proof: We may assume without loss of generality that $\psi(X) < \pi/2$. We furthermore assume that the x-coordinate of $\textsc{loc}(Y)$ is greater than that of $\textsc{loc}(X)$; the other case is similar. Let w denote the distance between $\textsc{loc}(X)$ and $\textsc{loc}(Y)$. Let p_1 and p_2 be the endpoints of the line segment $S = \textsc{seg}(X)$, with p_1 on O_1 and p_2 on O_2. We assume that p_1 is more distant from $\textsc{loc}(X)$ than p_2; the other case is similar. Let q be the intersection of $C_L(Y)$ with S in the same side of $\textsc{loc}(X)$ as p_1. Since the distance of $\textsc{loc}(Y)$ from L_X is $w \sin \psi(X) \leq l^2/200$ and furthermore $l \leq 1/10$, a crude calculation shows that the length of the chord formed by L_X on $C_L(Y)$ is at most $l/4$. Therefore, q lies on S and the distance between p_1 and q is at least $l/4$. Let Y' be a position on $C_L(Y)$ between $\textsc{loc}(Y)$ and q such that $C_R(Y')$ is tangent with L_X: let q' be the point at which S is tangent with $C_R(Y')$. Since the length of the arc of $C_R(Y')$ from $\textsc{loc}(Y')$ to q cannot be greater than that of the arc of $C_L(Y)$ from $\textsc{loc}(Y)$ to $\textsc{loc}(Y')$, which is crudely at most $l/4$, q' must lie on S. Our 2-feasible path first follows L_X up to q', reverses along $C_R(Y')$ to $\textsc{loc}(Y')$, and continues along $C_L(Y)$ backward up to $\textsc{loc}(Y)$. See Figure 23. □

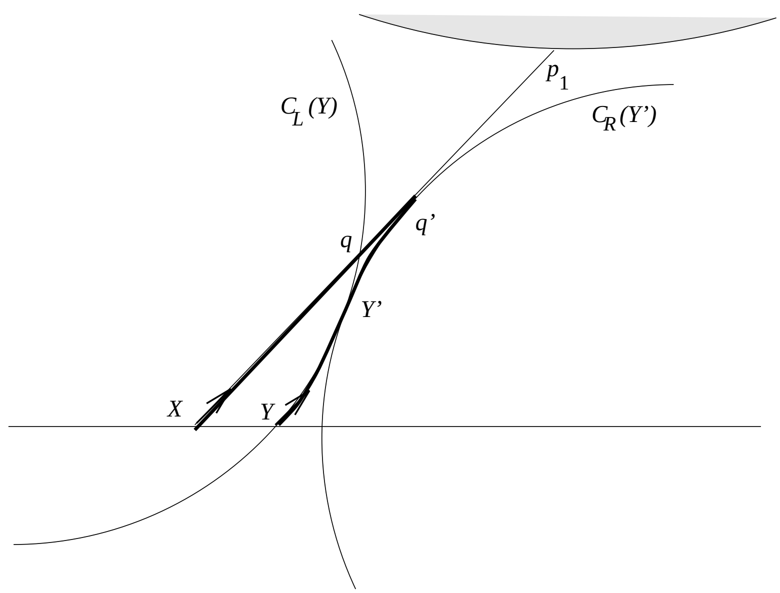

Figure 23: The path in Lemma 37

Let X and Y be positions doubly blocked by the same blocking pair, such that $\psi(X) = \psi(Y)$ and both $\textsc{loc}(X)$ and $\textsc{loc}(Y)$ lie on the x-axis of the standard coordinate system determined by the blocking pair. The *zigzag path* from X to Y is defined as follows. We define positions X_i, $i = 0, 1, \ldots$ on the x-axis and 2-feasible paths P_i, $i = 0, 1, \ldots$ by inductin on i. It will be maintained that $\psi(X_i) = \psi(X) = \psi(Y)$. We set $X_0 = X$ for the base case. Suppose X_i is defined and $X_i \neq Y$. Let l_i be the length of $\textsc{seg}(X_i)$ If the distance from $\textsc{loc}(X)$ to $\textsc{loc}(Y)$ is at most $l_i^2/(200\psi(X_i))$ then set X_{i+1} to Y; otherwise let X_{i+1} be the position such that $\psi(X_{i+1}) = \psi(X)$ and $\textsc{loc}(X_{i+1})$ is on the x-axis and at distance $l_i^2/(200\psi(X_i))$ from $\textsc{loc}(X_{i+1})$ toward $\textsc{loc}(Y)$. Define path P_i to be the 2-feasible path from X_i to X_{i+1} provided by Lemma 37. Let k be such that $X_k = Y$. Then, the zigzag path from X to Y is defined to be the concatenation of the k paths $P_0, \ldots, P_{k-1}$, which is $2k$-feasible. Note the zigzag path can be computed in time $O(n + k)$.

We now describe the approximate solution for a non-basic instance. Suppose a non-basic instance is given. We assume the standard coordinate system determined by the blocking pair that blocks S and F simultaneously. Furthermore, we assume that $\textsc{loc}(S)$ and $\textsc{loc}(F)$ are on the x-axis. This is justified as follows. If $\textsc{loc}(S)$ is not on the x-axis, then let S' be the position such that $\psi(S') = \psi(S)$ and $\textsc{loc}(S')$ is the intersection of L_S with the x-axis. Define F' similarly; we have reduced the problem to the instance where S and F are replaced by S' and F'. It is easily seen that if the original instance admits a k-optimal path with length l then the modified instance admits a $(k+2)$-optimal path with length at most $l + O(1)$. Therefore, it suffices to consider the case where $\textsc{loc}(S)$ and $\textsc{loc}(F)$ are on the x-axis.

Consider all the greedily rotating paths with ck reversals (where c is the constant in Theorems 30 and 31), starting from S and from F. Since the given instance is non-basic, and therefore none of these paths is unblocking (recall the definition of a non-basic instance requires that no $(ck + 200)$-feasible greedily rotating path from S or F is unblocking). Let $\hat{\psi}$ be the smallest positive real number such that, for either $\hat{\psi}' = \hat{\psi}$ or $\hat{\psi}' = \pi - \hat{\psi}$, at least one of the above greedily rotating paths from each of S and F has a position X with $\psi(X) = \hat{\psi}'$. Without loss of generality, assume that this condition is satisfied with $\hat{\psi}' = \hat{\psi}$ and hence $0 < \hat{\psi} \leq \pi/2$. Let P_S be one of the greedily rotating paths from S as above with final position S_0 such that $\psi(S_0) = \hat{\psi}$. Let P'_S be P_S followed by a line segment of L_{S_0} so that the location of the final position S_1 of P'_S is on the x-axis. Define P'_F and F_1 similarly, replacing S by F in the above definition. Let P' be the zigzag path from S_1 to F_1. Then, our solution path from S to F is defined as the concatenation of P'_S, P', and the reversal of P'_F.

Lemma 38 *There exists a universal constant c_3 such that, for each non-basic instance that admits a k-feasible solution, the path constructed above is c_3k-feasible.*

Proof: We use the notation above and keep the assumption that $\hat{\psi} \leq \pi/2$. Let Q be the k-optimal path from S to F. The non-basicness assumption implies that Q is not unblocking, due to Theorem 30. Let Q' be the path from S_1 to F_1 obtained by concatenating the reversal of P'_S, Q, and P'_F in this order. Note that Q' is $((2c+1)k+4)$-feasible. Define positions $X_0, \ldots, X_{k'}$ inductively as follows. Set $X_0 = S_1$. Suppose X_i is defined and $X_i \neq F_1$. Let X'_{i+1} denote the last position on Q' such that $\textsc{loc}(X'_{i+1})$ is on the x-axis and on the side of $\textsc{loc}(X_i)$ opposite to $\textsc{loc}(F_1)$. Then, X_{i+1} is the first position after X'_{i+1} on Q' such that $\textsc{loc}(X)$ is on the x-axis. Note that $k' \leq ((2c+1)k+4)$. For $1 \leq i \leq k'$, let w_i denote the distance between $\textsc{loc}(X'_i)$ and $\textsc{loc}(X_i)$, p_i the cusp of Q' between X'_i and X_i, ψ_i the angle formed by the line through $\textsc{loc}(X_i)$ and p_i against the x-axis, ψ'_i the angle formed by the line through $\textsc{loc}(X'_i)$ and p_i against the x-axis, and $\hat{\psi}_i$ the smallest of ψ_i, $\pi - \psi_i$, ψ'_i, and $\pi - \psi'_i$.

We claim that the length of the subpath of Q' from X_i' to X_i is at most $\hat{\psi}_i/100$. This is because otherwise there would be a 200-feasible greedily rotating path from either X_i' or X_i that is unblocking, giving a $(k+200)$-feasible unblocking path from S or F: a contradiction to the assumption that the given instance is non-basic since this would implies a $c(k+200)$-feasible greedy unblocking path from S or F by Theorem 30. Thus, Lemma 35 applies to the 2-feasible path from X_i' to X_i yielding $w_i \leq 5d_i^2 \sin^{-3} \hat{\psi}_i$.

For $1 \leq i \leq k'$, let m_i denote the number of intersection points of our zigzag path P' with the x-axis that lie in the interval between $\textsc{loc}(X_i')$ and $\textsc{loc}(X_i)$, including $\textsc{loc}(X_i')$ but not $\textsc{loc}(X_i)$. We claim that $m_i \leq 8000$ for every i. To show this, fix i and let Y_1, ..., Y_{m_i} be those intersection points in the interval between between $\textsc{loc}(X_i')$ and $\textsc{loc}(X_i)$. For convenience, let Y_{m_i+1} denote the intersection of P' and the x-axis that immediately follows Y_{m_i} (and on or beyond $\textsc{loc}(X_i)$). Let l_j denote the length of $\textsc{seg}(Y_j)$, $1 \leq j \leq m_i$. Then, since $\hat{\psi}_i \geq \hat{\psi}$ it is easily verified that $l_j \geq d_i/\sin\hat{\psi}_i$ for every j, $1 \leq j \leq m_i$. Using this inequality and the definition of the zigzag path, the distance between Y_j and Y_{j+1}, for $1 \leq j \leq m_i$, is at least

$$(l_j/2)^2/(200\hat{\psi}) \geq d_i^2 \hat{\psi}_i^{-1} \sin^{-2}\hat{\psi}_i/800.$$

Therefore, combining this with the upper bound on w_i above, we have

$$\begin{aligned} m_i &\leq w_i/(d_i^2 \hat{\psi}_i^{-1} \sin^{-2}\hat{\psi}_i/800) \\ &\leq 5d_i^2 \sin^{-3}\hat{\psi}_i/(d_i^2 \hat{\psi}_i^{-1} \sin^{-2}\hat{\psi}_i/800) \\ &\leq 8000. \end{aligned}$$

Therefore, the number of reversals in the zigzag path P' is at most 16000 times that of Q'. Thus, our solution path for the given non-basic instance has less than $2ck+16000(2ck+k) = c_3k$ reversals, where $c_3 = 32002c + 16000$.

□

We are now ready to prove Theorem 2 in the weak sense.

Proof of the weak-sense version of Theorem 2: Given a position-position instance and k, we first construct the basic path. Let $c_2 = c(2c_3+4)$, where c is the constant in Theorem 30 and c_3 is the constant in Lemma 38. If the number of reversals in the basic path is at most c_2k, then we adopt this path. This path is at most 9π longer than the k-optimal path, due to Lemma 36. Otherwise, by the observation following the definition of a non-basic instance, the given instance is non-basic: we construct the path for this non-basic path as described above. The number of component paths of this path is at most c_3k, by Lemma 38. We claim that the length of this non-unblocking path is at most c_3. Suppose otherwise. Then, some component path of the non-unblocking path must have length greater than $1/k$. This means that we can unblock from S with roughly $c_3k + \pi k/2$ reversals, by following the non-unblocking path up to the long component path and then switching to greedy rotation. This is a contradiction, because then the basic path must have at most c_2k reversals. □

4.4 The strong-sense approximation

To approximate a k-optimal path in the strong sense, we first construct the path for weak-sense optimization as described in Section 4.3. If the path happens to end with position F, then we are done. Suppose the path ends with $\text{REV}(F)$. First note that, in this case, the given instance is basic, because if it is non-basic then S and $\text{REV}(F)$ must be doubly-blocked by an identical blocking pair and therefore any k-feasible path from S to F must unblock, implying that the instance is basic, a contradiction.

Call a point *exposed*, if the circle of radius $1/c_4k$ centered at the point intersects with at most one obstacle, where c_4 is some constant to be determined later. If there is an exposed point on the above path, then we modify it by inserting a path with at most $4c_4k$ reversals at this exposed point so that the orientation of the path is reversed there. Thus, the only problematic case is when there is no exposed point on the constructed path. We claim that, in this case, the k-optimal path in the strong sense must have an exposed point. Suppose otherwise. Then, our path together with the k-optimal path forms a path from F to $\text{REV}(F)$ with $(c_2+1)k$ reversals on which no point is exposed. This is a contradiction if we set, say, $c_4 = 10c_2$. Now the case we are considering (the path we have constructed lacks an exposed point) may happen only when S and F are prevented from being exposed by the same pair of obstacles, say O_1 and O_2. Using the standard coordinate system determined (O_1, O_2), let x_S and x_F be the x-coordinates of $\text{LOC}(S)$ and $\text{LOC}(F)$ respectively, and assume $x_S \leq x_F$ without loss of generality. Let p_1 be an exposed point with the largest x-coordinate smaller than x_S and p_2 be an exposed point with the smallest x-coordinate larger than x_F. Such points exist because the obstacles are of finite size. We construct a path that goes from S to p_1 and a path from F to p_1 (using the method of Section 4.3), and concatenate these paths via some suitable rotation of orientation at p_1. Construct a similar path using p_2 instead of p_1. Of these two paths, we choose the one with at most $10c_4k$ reversals and, when both qualify, the shorter one. Since the k-optimal path must go through (a neighborhood of) either p_1 or p_2 by the above claim, at least one of the two paths must have at most $10c_4k$ reversals. The strong-sense version of Theorem 2 follows (with bigger constants than in the weak-sense version).

5 Further Work

We now mention some extensions and directions for further work; as we point out, all of these apply only if the robot cannot reverse. Consider our problem in the setting of *online navigation* [30]: the robot learns about each obstacle only when it can see the obstacle. For the simple case where every obstacle is a unit-radius disc, a straightforward but tedious geometric calculation shows that the following algorithm achieves a constant competitiveness: imagine that there are no obstacles and travel along the Dubins path from S to F. When an obstacle that obstructs the path becomes visible, attempt to circumvent it by going around it in the direction that minimizes the travel on its boundary. On becoming tangent to the obstacle at position X, form a new Dubins path from X to F and continue. It may be possible to pose this problem with a restriction on the total number of reversals; allowing arbitrarily many reversals reduces the problem to the unconstrained case.

In the robotics literature the following approach to steering-constrained motion planning is frequently described [24, 25]: first use a conventional (unconstrained) path planner to generate a tentative path and then use heuristics to "round" this path into a topologically equivalent path that meets the curvature constraint. Given a scene and an unconstrained path, what is the complexity of deciding whether a topologically equivalent reversal-free path exists? For scenes with only moderate obstacles we have a fast decision procedure.

Consider the following twist on another online problem: we have a 1-server problem in the plane, and each request is a point. The robot services a request by moving to the point; there is no constraint on the orientation with which it arrives at the requested point. Clearly the orientation with which it arrives at any point affects its future cost. Even in the absence of obstacles, we know of no competitive algorithm for this problem.

How well can one approximate the traveling salesman problem in the plane subject to curvature constraints, without reversals (even in the absence of obstacles)? Of course the biggest question left open by this work is: to what extent can the requirement of moderate obstacles be relaxed?

References

[1] J. Barraquand and J-C. Latombe. Nonholonomic mobile robots and optimal maneuvering. *Revue d'Intelligence Artificielle*, 3:77–103, 1989.

[2] J. Barraquand and J-C. Latombe. Nonholonomic multibody mobile robots: controllability and motion planning in the presence of obstacles. *Algorithmica*, 10:121–155, 1993.

[3] X. Bui, J. Boissonnat, P. Souères, and J. Leblond. Shortest path synthesis for Dubins non-holonomic robot. *Proceedings of the IEEE International Conference on Robotics and Automation*, 1994.

[4] J. Boissonnat, A. Cérézo, and J. Leblond. Shortest paths of bounded curvature in the plane. *Proceedings of the IEEE International Conference on Robotics and Automation*, 1992.

[5] J. Canny. Some algebraic and geometric configurations in PSPACE. In *Proc. 20th Annu. ACM Sympos. Theory Comput.*, pages 460–467, 1988.

[6] J. Canny, B. Donald, J. Reif, and P. Xavier. On the complexity of kinodynamic planning. In *Proc. 29th Symp. on Foundations of Computer Science*, 1988.

[7] J. Canny, A. Rege, and J. Reif. An exact algorithm for kinodynamic planning in the plane. *Discrete Comput. Geom.*, 6:461–484, 1991.

[8] J. Canny and J. Reif. New lower bound techniques for robot motion planning. In *Proc. 28th Symp. on Theory of Computing*, pages 49–60, 1987.

[9] E. Cockayne and G. Hall. Plane motion of a particle subject to curvature constraints. *SIAM J. Control*, 43:197–220, 1975.

[10] B. Donald and P. Xavier. Near-optimal kinodynamic planning for robots with coupled dynamics bounds. In *IEEE Int. Symp. on Intelligent Controls*, 1989.

[11] L.E. Dubins. On curves of minimal length with a constraint on average curvature and with prescribed initial and terminal positions and tangents. *American Journal of Mathematics*, 79:497–516, 1957.

[12] S. Fortune and G. Wilfong. Planning constrained motion. In *Annals of Math. and Art. Intell.*, pages 21–82, 1991.

[13] T. Fraichard. Smooth trajectory planning for a car in a structured world. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 318–323, 1991.

[14] M Fredman and R. Tarjan. Fibonacci heaps and their uses in improved network optimization. 34:596–615, 1987.
J. ACM

[15] J. Hopcroft, J. Schwartz, and M. Sharir, editors. *Planning, Geometry, and Complexity of Robot Motion.* Ablex Publishing Corp., Norwood, New Jersy, 1984.

[16] P. Jacobs and J. Canny. Planning smooth paths for mobile robots. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 2–7, 1989.

[17] P. Jacobs, J-P. Laumond, and M. Taix. Efficient motion planners for nonholonomic mobile robots. In *Proceedings of the IEEE/RSJ International Workshop on Intelligent Robots and Systems*, pages 1229–1235, 1991.

[18] K. Kedem, R. Livne, J. Pach, and M. Sharir. On the union of Jordan regions and collision-free translational motion amidst polygonal obstacles, *Discrete Comput. Geom.*, 1:59–71, 1986.

[19] V. Kostov and E. Degtiariova-Kostova. Suboptimal paths in the problem of a planar motion with bounded curvature. Tech. Rept. 2051, INRIA, Sophia-Antipolis, 1993.

[20] V. Kostov and E. Degtiariova-Kostova. The planar motion with bounded derivative of the curvature and its suboptimal paths. Tech. Rept. 2189, INRIA, Sophia-Antipolis, 1994.

[21] J-C. Latombe. A fast path-planner for a car-like indoor mobile robot. In *Proceedings of the 9th National Conference on Artificial Intelligence*, pages 659–665, 1991.

[22] J-C. Latombe. *Robot Motion Planning.* Kluwer-Academic, Boston, 1991.

[23] J.P. Laumond. Finding collision free smooth trajectories for a nonholonomic mobile robot. In *Proc. 10th Int. Joint Conf. on Artificial Intelligence*, pages 1120–1123, 1987.

[24] J-P. Laumond and T. Siméon. Motion planning for a two degrees of freedom mobile robot with towing. Technical Report 89148, LAAS/CNRS, Toulouse, 1989.

[25] J-P. Laumond, M. Taix, and P. Jacobs. A motion planner for car-like robots based on a global/local approach. In *Proceedings of the IEEE/RSJ International Workshop on Intelligent Robots and Systems*, pages 765–773, 1990.

[26] J-P. Laumond. Finding collision-free smooth trajectories for a non-holonomic mobile robot. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 1120–1123, 1987.

[27] B. Mirtich and J. Canny. Using skeletons for nonholonomic path planning among obstacles. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 2533–2540, 1992.

[28] Y. Nakamura and R. Mukherjee. Nonholonomic path planning and automation. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 1050–1055, 1989.

[29] C. Ó'Dúnlaing. Motion planning with inertial constraints. *Algorithmica*, 2:431–475, 1987.

[30] C.H. Papadimitriou and M. Yannakakis. Shortest paths without a map. *Theoretical Computer Science*, 84:127–150, 1991.

[31] F. Preparata and M. Shamos, *Computational Geometry: An Introduction*, Springer Verlag, Heidelberg, 1985

[32] J.A. Reeds and L.A. Shepp. Optimal paths for a car that goes both forwards and backwards. *Pacific Journal of Mathematics*, 145:367–393, 1990.

[33] J. H. Reif. Complexity of the generalized movers problem. In J. Hopcroft, J. Schwartz, and M. Sharir, editors, *Planning, Geometry and Complexity of Robot Motion*, pages 267–281. Ablex Pub. Corp., Norwood, NJ, 1987.

[34] J. H. Reif and M. Sharir. Motion planning in the presence of moving obstacles. In *Proc. 26th Annu. IEEE Sympos. Found. Comput. Sci.*, pages 144–154, 1985.

[35] J.H. Reif and S.R. Tate. Approximate kinodynamic planning using L_2-norm dynamic bounds. *Computers Math. Applic.*, 27(5):29–44, 1994.

[36] J.T. Schwartz and M. Sharir. Algorithmic motion planning in robotics. In J. Van Leeuwen, editor, *Algorithms and Complexity*, volume A of *Handbook of Theoretical Computer Science*, pages 391–430. Elsevier, Amsterdam, 1990.

[37] M. Sharir. On k-sets in arrangements of curves and surfaces, *Discrete Comput. Geom.*, 6 (1991), 593–613.

[38] H. Wang and P. Agarwal. Approximation algorithms for shortest paths with bounded curvature, manuscript.

[39] G. Wilfong. Motion planning for an autonomous vehicle. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 529–533, 1988.

[40] G. Wilfong. Shortest paths for autonomous vehicles. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 15–20, 1989.