Optimal Motion Strategies to Track and Capture a Predictable Target

Alon Efrat   Héctor H. González-Baños   Stephen G. Kobourov   Lingeshwaran Palaniappan

Introduction In this paper, we present an $O(n \log^{1+\varepsilon} n)$-time algorithm for finding the optimal path for an observer tracking a target inside a polygonal region. The goal is to keep the target visible to a mobile observer for as long as possible. Variations of the tracking problem arise in different applications such as visual servoing [1], [2], computer assisted surgery [3] and visibility-based planning of sensor control strategies [4].

A planning problem for maintaining line-of-sight object visibility is considered in [5]. A general dynamic programming approach is used to generate motions for a mobile robot (as an observer) to keep track of a moving target. If the target moves predictably, an optimal tracking motion can be computed off-line. If the target's motion is unpredictable, on-line strategies are used instead [6], [7]. Although on-line techniques are usually more desirable in a robotics context, off-line strategies remain useful in applications involving graphical simulations, supervision of industrial robots (which often follow known paths), and surveillance of vehicles constrained to move in a road-map.

Most off-line algorithms reported in the literature are not computationally efficient, requiring several seconds of computation even for simple view. In [8], a problem similar to the off-line case is modeled in the configuration-time space relative to the target under practical considerations for on-line use. The improved planning time for finding a feasible tracking motion is a few seconds. Both of the algorithms in [5] and [8] discretize the space in order to approximate the optimal path. This process is computationally costly, although heuristic techniques can speed the computation.

In this paper, we consider the planning problem for maintaining line-of-sight target visibility inside a polygonal region P which may contain holes. We assume that the observer loses the the target as soon as the line of sight between them is broken. The observer is modeled as a holonomic robot with a maximal speed of $v_{\rm obs}$, equipped with an omnidirectional sensor of unlimited range. Target detection and robot localization are assumed to be perfect. The target is initially in view and moves along a known path until a time t_{stop}.

We present an $O(n \log^{1+\varepsilon} n)$ time algorithm for finding the optimal path π^* given an initial observer position M_0, the path π_T of the target, and the velocities v_{obs} and v_T of the observer and the target respectively. Here ε is an arbitrarily small positive constant, and n is the number of vertices of P. The path is optimal in the sense that it maximizes the time t_{esc} when the observer first loses the target. If the visibility requirements are always met, then π^* minimizes the distance between the observer and the target at every time $t < t_{stop}$. If this distance becomes zero, the target is said to be *captured.*

A. Efrat, S. Kobourov and L. Palaniappan are with the Department of Computer Science of the University of Arizona, Tucson, AZ. E-mail: {alon, kobourov, lingesh}@cs.arizona.edu

Héctor H. González-Banos is with Honda's Fundamental Research Labs, Mountain View, CA. E-mail: hhg@hra.com

Although not a necessary assumption for our algorithm, if the target follows the shortest path to its destination the presentation becomes simpler and it is easier to obtain a bound on the running time. However, the algorithm computes an optimal path even when this assumption does not hold.

The observer is assumed to be a point. The configuration space of the target is also assumed to be two-dimensional, but the target is forbidden from moving arbitrarily close to an obstacle. In order to simplify our discussion, we represent the target as a small axis-parallel square.

Visibility Model The observer is assumed to be fitted with an idealized omnidirectional sensor. Hence, a target is visible *iff* the line-of-sight between the target and the observer is un-obstructed.

Let $\mathcal{V}(q) \subseteq P$ be the set of all locations in P from which the target is visible to an observer located at q. The set $\mathcal{V}(q)$ is the *visibility region* at the observer position q. An important concept in the formulation of our algorithm is that of the *visibility sweeping line.* A visibility sweeping line $\ell(t)$ is defined as a chord (a segment fully contained in P and connects two points on ∂P) through a corner of the target and a reflex vertex of the polygon P. At any time t, the observer must be in one of the half-planes defined by $\ell(t)$ in order to maintain line of sight with the target; see figure in the next page. Though in general there could be $\Theta(n)$ visibility lines, we show that at each time at most two of them influence the observer's path.

Target Predictability and Optimal Paths The target is supposed to be *predictable.* That is, the configuration of the target $q^t(t) \in \mathcal{C}^t$ is known for all time $t < t_{stop}$, where t_{stop} is the target *stopping time.* Therefore, our algorithm computes an *off-line* strategy. It is important to note that although the target is predictable, it may be unavoidable for the observer to lose the target for a particular choice of initial conditions and target velocities. In this scenario, a useful notion is that of *escape time* (t_{esc}) —the time when the observer first loses the target.

Alternatively, it is sometimes possible for the observer to capture the target —i.e., $q(t) = q^t(t)$ for some $t < t_{stop}$. The *capture time* (t_{cap}) is the time when the target is first captured.

In this paper, a path is optimal if it satisfies the following definition:

Definition .1: An observer path π^* is optimal if it is the shortest path among all possible paths that maximize the *escape time* (t_{esc}). For the cases when t_{esc} does not exist (i.e., the observer never loses track of the target), then π^* is optimal if it minimizes the *capture time* (t_{cap}). If capture before t_{stop} is not possible, then π^* minimizes the final separation between the target and the observer.

Overview of the Algorithm In the full version of this paper, we show that the optimal path π^* is composed of *straight-line segments* and *leaning* curves, as described below:

Straight-line segments: The observer follows a straight-line segment once it knows some (carefully chosen) point M, to which it can go directly without ever losing sight of the target.

Leaning curves: This is a curve defined by the a connected part of the shortest path that the observer can take to reach a vertex C of P while moving at maximal velocity, and at any time t being on a visibility sweeping line $\ell(t)$, rotating around C, as the target moves along a straight segment; see Figure on the right.

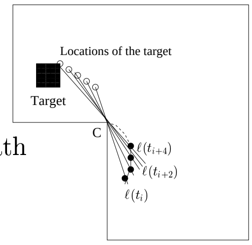

Without any loss of generality, assume that the observer is at the origin at time $t = 0$. We now proceed to define the concept of a π-path.

Definition .2 (π-path) The path $\pi(q,t) \subset P$ is a *π-path* if it connects the origin to q and satisfies the following properties:

a. The target is always visible to the observer.
b. π is the shortest path that satisfies the above.
c. The observer traverses π at the top speed v_{obs}.
d. The observer arrives at q in time t.

Note that a path $\pi(q,t)$ does not necessarily exists for every combination of q and t, but it can shown that if it exists, it is unique.

Hence in the rest of this paper we drop t in the notation of $\pi(q,t)$, and denote this path as $\pi(q)$. We define $\Gamma(t)$ as the subset of all points q such that the length of $\pi(q)$ is $v_{obs}t$. $\Gamma(t)$ contains the arrival point for the optimal path π^* at time t.

The algorithm computes the evolution of the set $\Gamma(t)$ as the target moves at maximum speed v_T along π_T. At time t, $\Gamma(t)$ contains the candidate points through which π^* might pass. It can be shown that $\Gamma(t)$ is a convex path in the plane composed of pieces we refer to as *edges*. Each edge has a constant description complexity.

Our algorithm maintains a set $\emptyset$ of obstacles which implicitly define $\Gamma(t)$ in the following sense: Every π-path $\pi(q)$ can be described as the shortest path $\pi'(q)$ connecting M_0 to q, avoids all obstacles in $\emptyset$, and homotopic to π_T (i.e., can be continuously deformed to $\pi(q)$ while staying in P, without disconnecting it or crossing any obstacle. The algorithm finds all the critical events that influence the shape of $\Gamma(t)$ as the target moves along π_T, and stores this information in a queue Q of future events. After each combinatorial change, we get the next event from Q, update $\Gamma(t)$ and $\emptyset$, compute future events, and insert them into Q. It is shown in the full version of the paper that if the target moves along a shortest path, then there are only $O(n)$ events, and each can be handled in time $O(\log^{\varepsilon} n)$.

$\Gamma(t)$ is made up of two kinds of *edges*: circular arcs and ρ-edges. The observe lie on a point q of ρ-edge at time t if it $\pi(q)$ has a non-empty intersection with a leaning curve γ, and $\pi(q)$ and the portion of $\pi(q)$ after leaving γ consists of a straight segment connecting q to a tangent point on γ.

At any instant, one or two visibility sweeping lines define the endpoints of one or two edges. Let $\gamma(t)$ be one of these edges, and let $p(t)$ be its endpoint on the sweeping line $\ell(t)$. We say that $p(t)$ *stretches* $\gamma(t)$ if there is a point $q \in \Gamma(t+\varepsilon)$ such that $\pi(q)$ goes through $p(t)$, for some sufficiently small $\varepsilon > 0$; Otherwise, we say that $p(t)$ *shrinks* $\gamma(t)$; The events at which an edge switches form stretching to shrinking or vice versa forms the most interesting events handled by the algorithm, and the corresponding points form vertices of the observer optimum path π^*.

We assume that given the parameters of a ρ-edge or a circular edge, the M-point for this edge can be computed in time $O(1)$.

In the full version of the paper we show how one can trace all types of events that $\Gamma(t)$ goes through. Eventually, either $\Gamma(t)$ becomes empty at a point $p \in P$, which implies that the target is doomed to loose the target at p. Otherwise, the algorithm ends when either the target reaches is destination, or when the location of the target coincides with the location of the observer. In either case, the optimal path of the observer can be reconstruct by tracing the obstacles that this path lies on.

References

[1] S. Hutchinson, G. D. Gager, and P. I. Corke, "A tutorial on visual servo control," *IEEE Transactions on Robotics and Automation*, vol. 12, no. 5, pp. 651–670, 1996.

[2] N. P. Papanikolopoulos, P. K. Khosla, and T. Kanade, "Visual tracking of a moving target by a camera mounted on a robot: A combination of control and vision," *IEEE Transactions on Robotics and Automation*, 1993.

[3] S. M. Lavallée, J. Troccaz, L. Gaborit, A. L. Benabid P. Cinquin, and D. Hoffmann, "Image guided operating robot: A clinical application in stereotactic neurosurgery," in *Computer Integrated Surgery: Technology and Clinical Applications*, R.H.Taylor, S. Lavallée, G. Burdea, and R. Mösges, Eds. 1995, pp. 342–351, MIT Press.

[4] Amy J. Briggs and Bruce Randall Donald, "Visibility-based planning of sensor control strategies," *Algorithmica*, vol. 26, no. 3-4, pp. 364–388, 2000.

[5] S. M. LaValle, H. H. González-Baños, C. Becker, and J.-C. Latombe, "Motion strategies for maintaining visibility of a moving target," in *ICRA*, 1997.

[6] H. Gonzalez-Banos, C.-Y. Lee, and J.-C. Latombe, "Real-time combinatorial tracking of a target moving unpredictably among obstacles," in *Proceedings of the 2002 International Conference on Robotics and Automation*, Washington, DC, USA, May 2002, pp. 1683–1690.

[7] P. Fabiani and J.C. Latombe, "Dealing with geometric constraints in game-theoretic planning," in *Proc. Int. Joint Conf. on Artif. Intell.*, 1999, pp. 942–947.

[8] T.-Y. Li and T.-H. Yu, "Planning object tracking motions," in *ICRA*, 1999.

[9] Kiriakos N. Kutulakos, Charles R. Dyer, and Vladimir J. Lumelsky, "Provable strategies for vision-guided exploration in three dimensions," in *ICRA* 1994, pp. 1365–1372.

[10] R. Sharma and S. Hutchinson, "On the observability of robot motion under active camera control," in *ICRA* 1994 pp. 162–167.