

Title	Constructing Optimal Highways
Author(s)	Ahn, Hee-Kap; Alt, Helmut; Asano, Tetsuo; Bae, Sang Won; Brass, Peter; Cheong, Otfried; Knauer, Christian; Na, Hyeon-Suk; Shin, Chan-Su; Wolff, Alexander
Citation	International Journal of Foundations of Computer Science, 20(1): 3-23
Issue Date	2009
Type	Journal Article
Text version	author
URL	http://hdl.handle.net/10119/9186
Rights	Electronic version of an article published as International Journal of Foundations of Computer Science, 20(1), 2009, 3-23. DOI:10.1142/S0129054109006425. Copyright World Scientific Publishing Company, http://dx.doi.org/10.1142/S0129054109006425
Description	



Constructing Optimal Highways*

Hee-Kap Ahn[†] Helmut Alt[‡] Tetsuo Asano[§] Sang Won Bae[¶] Peter Brass^{||}
Otfried Cheong[¶] Christian Knauer[‡] Hyeon-Suk Na^{**} Chan-Su Shin^{††}
Alexander Wolff^{‡‡}

Abstract

For two points p and q in the plane, a (unbounded) line h , called a highway, and a real $v > 1$, we define the *travel time* (also known as the *city distance*) from p and q to be the time needed to traverse a quickest path from p to q , where the distance is measured with speed v on h and with speed 1 in the underlying metric elsewhere.

Given a set S of n points in the plane and a highway speed v , we consider the problem of finding an axis-parallel line, the *highway*, that minimizes the maximum travel time over all pairs of points in S . We achieve a linear-time algorithm both for the L_1 - and the Euclidean metric as the underlying metric. We also consider the problem of computing an optimal pair of highways, one being horizontal, one vertical.

Keywords: geometric facility location, min-max-min problem, city metric, time metric, optimal highways

1 Introduction

Imagine that there are n cities that are represented by points on the plane with a metric, such as the L_1 or the Euclidean metric. To increase the cultural and the commercial interchange among those cities, they would have decided to build a straight and long

highway with fixed speed v and with fixed direction across the cities. Further, they might hope to minimize the maximum travel time in moving from one city to another. How could they locate such an optimal highway? This paper gives efficient and simple solutions to this question and to several variations.

Most previous results considering highways (also called *transportation networks* or *roads*) have focused on how to compute quickest paths among the cities and Voronoi diagrams for the cities under metrics induced by *given* highways.

Abellanas et al. (2001) considered the Voronoi diagram for point sets given an isothetic and monotone highway under the L_1 -metric. Abellanas et al. (2003) considered the problem also under the Euclidean metric and studied shortest paths. Aichholzer et al. (2004) introduced the city metric induced by the L_1 -metric and an isothetic highway network that consists of a number of axis-parallel line segments. They gave an efficient algorithm for constructing the Voronoi diagram and a quickest-path map for a set of points given the city metric.

More recently, Görke and Wolff (2005) and Bae et al. (2005) improved the results of Aichholzer et al. (2004) in terms of running time of the construction algorithm. Under the Euclidean metric, Bae and Chwa (2004) presented algorithms that compute the Voronoi diagram and shortest paths in more general highway networks whose segments can have arbitrary orientation and speed. Bae and Chwa (2005) recently also proved that their approach naturally extends to more general metrics including asymmetric convex distances.

Many well-known facility-location problems are so-called *min-max* problems where the task is to place a facility such that the maximum cost incurred by the customers is minimized. For example, if the customers are points in the plane and the cost is their Euclidean distance to the facility, the center of the smallest enclosing circle is the point that minimizes the maximum distance to the customers. Recently, Cardinal and Langerman (2006) introduced the subclass of so-called *min-max-min* facility location problems, where the cost caused by a customer is the minimum between the cost that arises from using the facility and the cost from *not* using the facility. Transport facility location problems are typical min-max-min problems: a facility like a railway line or a highway will usually not be used by a customer if the facility is too far

*This research was started during the 8th Korean Workshop on Computational Geometry, organized by T. Asano at JAIST, Kanazawa, Aug. 1–6, 2005

[†]Department of Computer Science and Engineering, Sejong University, Seoul, Korea. Email: heekap@gmail.com

[‡]Institute of Computer Science, Freie Universität Berlin, Germany. Email: {alt,christian.knauer}@inf.fu-berlin.de

[§]School of Information Science, Japan Advanced Institute of Science and Technology, Japan. Email: t-asano@jaist.ac.jp

[¶]Division of Computer Science, Korea Advanced Institute of Science and Technology, Korea. Email: {swbae,otfried}@tclab.kaist.ac.kr

^{||}Department of Computer Science, City College of New York, USA. Email: peter@cs.cuny.cuny.edu

^{**}School of Computing, Soongsil University, Seoul, South Korea. Email: hsnua@ssu.ac.kr

^{††}School of Electr. and Inform. Engineering, Hankuk University of Foreign Studies, Yongin, Korea. Email: cssin@hufs.ac.kr

^{‡‡}Fakultät für Informatik, Universität Karlsruhe, Germany. Email: awolff@ira.uka.de

Copyright ©2007, Australian Computer Society, Inc. This paper appeared at Computing: The Australasian Theory Symposium (CATS'07), Ballarat, Australia. Conferences in Research and Practice in Information Technology (CRPIT), Vol. 65. Barry Jay and Joachim Gudmundsson, Ed. Reproduction for academic, not-for profit purposes permitted provided this text is included.

away. Cardinal and Langerman (2006) consider three such problems, among them the following. Given a set P of pairs of points, they compute the highway that minimizes the maximum travel time over all pairs in P . The travel time is the minimum between the bee-line distance in the underlying metric (any L_p -metric in their case) and the travel-time via the highway. Assuming infinite speed along the highway and vertical access to the highway, they solve the problem in expected time linear in the number of pairs.

The problems we consider also fall in the class of min-max-min facility location problems. We are interested in the case where n points are given, and we wish to place a highway such that the maximum travel time over *all* pairs of points is minimized. For this case Cardinal and Langerman’s algorithm takes expected quadratic time. We show how to make use of the coherence between the pairs of points to get deterministic near-linear-time algorithms.

In particular, we achieve linear-time algorithms for finding the optimal *vertical* highway under the L_1 - and the L_2 -metric, see Sections 4 and 5, respectively. If we allow arbitrary orientation, we can determine the optimal highway in $O(n \log n)$ time under the L_1 -metric, see Section 4.

We also consider placing a *highway cross*, i.e., a pair of highways that intersect perpendicularly. Under the L_1 -metric we can determine the optimal axis-aligned highway cross with infinite speed in $O(n \log n)$ time, see Section 3. For constant speed the problem becomes considerably harder—even under the L_1 -metric, see Section 6. We give a generic exact $O(n^{4+\epsilon})$ -time algorithm based on computing minima of upper envelopes. We also consider approximative solutions, see Section 7. All our results are summarized in Table 1.

Throughout the paper we assume that the input point set S contains at least three points and that not all points have the same y -coordinate. If $|S| < 3$, it is trivial to get an optimal highway. Also, if all points of S have the same y -coordinate, no vertical highway reduces the maximum travel time.

2 The optimal highway for infinite speed

As a warm-up exercise, let us consider the problem of finding the optimal placement of a vertical highway, assuming the highway speed is infinite.

Theorem 1 *Given n points in the plane, the middle line of the smallest enclosing vertical strip is an optimal vertical highway of infinite speed. It can be computed in linear time.*

The easy proof is left to the reader. What is interesting is that the optimal highway corresponds to a smallest enclosing figure—we will see this theme repeatedly in the following, see Figure 1.

Note that the result holds in any L_p -metric, as all travel to and from the highway is parallel to the x -axis.

The theorem generalizes to highways of arbitrary orientation in the Euclidean metric (that is, travel to and from the highway is orthogonal to it):

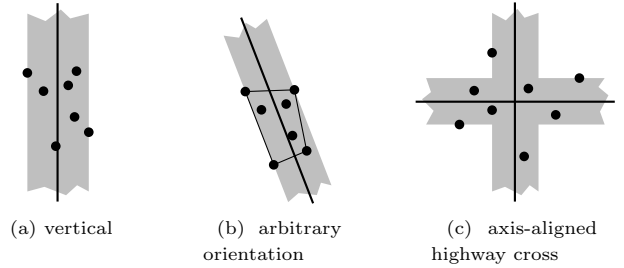


Figure 1: Optimal infinite-speed highways (solid lines) and corresponding enclosing figures (shaded).

Theorem 2 *Given n points in the plane, the middle line of the smallest enclosing strip is an optimal highway of infinite speed. It can be computed in $O(n \log n)$ time.*

The algorithm used here is the rotating calipers algorithm by Toussaint (1983). After computing the convex hull of the point set, it runs in linear time.

3 The optimal highway cross for infinite speed

Now we consider the problem of placing more than one highway. Observe that multiple parallel highways with the same speed do not reduce the maximum travel time because the quickest path using several highways can be simulated with only one highway. Instead we investigate highway crosses, i.e., pairs of highways that intersect perpendicularly. We give algorithms for computing the optimal *axis-aligned* highway cross, see Figure 1c.

Definition 1 *An enclosing cross for a point set S is the union of a horizontal and a vertical strip of equal width containing S .*

Lemma 1 *The travel-time diameter of the optimal axis-aligned highway cross with infinite speed equals the width of the smallest enclosing cross.*

Proof. Let δ be the travel-time diameter, and let δ' be the width of a smallest enclosing cross C .

We first show that $\delta' \leq \delta$: Let h_1, h_2 be a pair of optimal highways. We assign each point in S to its closest highway so that S is partitioned into two subsets: one consisting of points closer to the horizontal highway and the other consisting of points closer to the vertical highway. We put around each highway the narrowest strip containing all the points assigned to the highway. Then both strips have width at most δ , otherwise there are two points in the wider strip whose travel-time distance is larger than δ . Therefore we can obtain an enclosing cross of width δ by widening each strip until its width becomes δ . Since δ' was minimal, we have $\delta' \leq \delta$.

It remains to show $\delta \leq \delta'$: We place highways in the middle of each strip of C . This results in a pair of highways with travel-time diameter at most δ' . Since δ is optimal, we have $\delta \leq \delta'$. \square

v	facility	vertical highway	highway w/arbitrary orientation		axis-aligned highway cross
	underlying metric	L_1 & L_2	ortho	L_2	L_1
∞	exact	$O(n)$	$O(n \log n)$		$O(n \log n)$
const.	exact	$O(n)$	$O(n \log n)$	open	$O(n^{4+\epsilon})$ for any $\epsilon > 0$
	$(1 + \sqrt{2})$ -approx.				$O(n \log n)$
	$(2 + \epsilon)$ -approx.				$O(\log(1/\epsilon)\alpha(n)n \log n)^*$
	$(1 + \epsilon)$ -approx.				$O(\log(1/\epsilon)\alpha(n)n^2 \log n)^*$

Table 1: Overview over our results. “Ortho” means orthogonal travel to and from the highway. *) Yields an approximation of the maximum travel time, no highway.

Note that once again the optimal facility corresponds to a minimal enclosing shape. This shape can be computed efficiently.

Theorem 3 *Given n points in the plane, the optimal axis-aligned highway cross for infinite speed corresponds to the smallest enclosing strip cross. It can be computed in $O(n \log n)$ time.*

Proof. The characterization follows from Lemma 1. The smallest enclosing cross of a set S of n points can be found as follows.

1. We resort the points by their x - and by their y -coordinates.
2. For a given width $\omega > 0$, we can decide in linear time whether an enclosing cross of width ω exists. If it is the case, the enclosing cross can be found in the same time. Our decision algorithm is as follows. We slide a vertical strip V of width ω across the point set from left to right. We maintain a horizontal strip H of smallest width containing all the points not in V . For each point entering V from the right or leaving V from the left, we update H accordingly. If the width of H ever becomes ω or less, we answer “yes” and report an enclosing cross. Otherwise, we answer “no”.
3. The width of the smallest enclosing cross is in the list of numbers $L = L_x \cup L_y$, where $L_x = \{x_j - x_i \mid 1 \leq i < j \leq n\}$ and $x_1 \leq x_2 \leq \dots \leq x_n$ is the sorted sequence of x -coordinates. The list L_y is defined analogously based on the sorted sequence of y -coordinates.
4. Consider the matrix A with $A[i, j] = x_j - x_{n-i+1}$. The rows and columns of A are sorted in ascending order. Using the technique of Frederickson and Johnson (1984), we can determine the k -th element of such a *sorted matrix* in $O(n)$ time without constructing A explicitly. This gives us a way to do binary search on L_x . This search consists of $O(\log n)$ steps, each of which first invokes the algorithm of Frederickson and Johnson to find the median of the remaining elements in L and then calls the decision algorithm of stage 2. Thus, the total runtime is $O(n \log n)$. Likewise we can search for the smallest value in L_y for which an enclosing cross exists. Finally we return the minimum of the two values. Notice that

the algorithm in stage 2 computes not only the width of the smallest enclosing cross but also the cross itself. □

4 The optimal highway in the L_1 -metric

We say that a v -rhombus is a rhombus of aspect ratio v .

Theorem 4 *Given n points in the plane, the vertical axis of symmetry of the smallest enclosing axis-aligned v -rhombus is an optimal vertical highway for the L_1 -metric and highway speed v . It can be computed in $O(n)$ time.*

Proof. For a pair of points p, q , let $d(p, q) := |y_p - y_q|/v + |x_p - x_q|$. Clearly, $d(p, q)$ is a lower bound for the diameter of the point set for any vertical highway, and therefore $\delta := \max_{p, q \in P} d(p, q)$ is also a lower bound. We show that in fact this bound can be obtained, resulting in an optimal highway.

We observe that the point set can be enclosed in a rhombus with horizontal diagonal δ and vertical diagonal δv . For an example with $v = 2$, see Figure 2a. If we place a vertical highway along the vertical diagonal of this rhombus, then any point in the rhombus has travel time at most $\delta/2$ to the center of the rhombus. This implies that the travel-time diameter is at most δ .

The computation boils down to computing minimum and maximum y -axis intercepts among all lines of slope v and among all lines of slope $-v$ that go through input points. □

Consider now the case where we allow arbitrary orientation of the highway, but travel to and from the highway is still orthogonal to it. The characterization above still applies, and so we can find the optimal highway by finding the smallest v -rhombus that encloses the point set, see Figure 2b.

Theorem 5 *Given n points in the plane, the optimal highway with speed v and orthogonal travel to and from the highway can be found in $O(n \log n)$ time.*

Proof. First we compute the convex hull C of the point set. Then we use the rotating-calipers algorithm to compute the function $w : [0, \pi) \rightarrow \mathbb{R}^+$

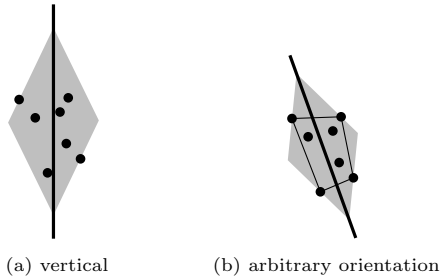


Figure 2: Optimal constant-speed highways (solid lines) and corresponding enclosing figures (shaded).

that maps an angle ϕ to the minimum width of a strip that contains C and forms an angle of size ϕ with the positive x -axis. This function consists of at most n pieces each of which is a trigonometric function that can be computed explicitly in constant time. Let $\alpha := 2 \arctan v$ be the size of the larger two inner angles formed by a v -rhombus. Then the function $w'(\phi) = \max\{w(\phi), w(\phi + \alpha \bmod \pi)\}$ maps ϕ to the width of a smallest v -rhombus that contains C and forms an angle of ϕ with the positive x -axis. The minimum over $w'(\phi)$ with $\phi \in [0, \pi)$ corresponds to the width of a smallest v -rhombus that contains C . By our above characterization the highway that goes through the longer diagonal of such a rhombus is optimal. \square

5 The optimal vertical highway in the Euclidean Metric

In this section we consider the Euclidean metric in the plane, and a vertical highway of speed $v > 1$.

Theorem 6 *Given n points in the plane, an optimal vertical speed- v highway under the Euclidean metric can be computed in $O(n)$ time.*

Proof. As Abellanas et al. (2003) showed, the quickest path (i.e., the one with shortest travel-time) between two points p and q has one of two forms. The quickest p - q path is either the segment pq or a path consisting of three segments pp^+ , p^+q^- , q^-q , where p^+ and q^- are points on the highway, and the lines pp^+ and qq^- form an angle of $\alpha = \arccos 1/v$ with the highway, see Figure 3a.

Now let us define a norm $\eta(x, y)$ on \mathbb{R}^2 as

$$\eta(x, y) = a|x| + (|y| - b|x|)/v = (a - b/v)|x| + (1/v)|y|,$$

where $a = 1/\sin \alpha$ and $b = 1/\tan \alpha$. Since $a > b > 0$ and $v > 1$, $a - b/v > 0$, and so $\eta(x, y) > 0$ unless $(x, y) = (0, 0)$, and η is indeed a norm.

Given two points p and q such that the highway is inbetween p and q and such that the shortest path between p and q makes use of the highway. Then the travel time from p to q is $\eta(q - p)$. When the highway cannot be used because the line pq forms an angle larger than α with the highway, then the travel time is simply the Euclidean distance $d(p, q)$, and $\eta(q - p)$ is an underestimate.

We note that η is simply a rescaled version of the L_1 -norm, and so we can find the smallest unit circle of this norm enclosing a given set of n points in linear time. This means we have the smallest factor $\delta > 0$ such that the entire point set fits in the rhombus R with corners $(0, \delta v)$, $(0, -\delta v)$, $(\delta/(a - b/v), 0)$, and $(-\delta/(a - b/v), 0)$ (after translating the point set).

We claim that the y -axis is now an optimal highway. We already know that there is a pair of points whose η -distance is 2δ , so this is a lower bound on the diameter. We now show that for any pair of points, either their travel time (with respect to the highway at $x = 0$) is at most 2δ , or they cannot use any vertical highway.

For any two points p, q in the rhombus R , we have $\eta(q - p) \leq 2\delta$. This means that if the highway lies inbetween, then we are already done. So assume that both p and q lie to the left of the highway, and such that they *can* use a vertical highway. Let q' be the reflection of q around the highway. Since R is symmetric with respect to $x = 0$, q' is also in R , and if p, q' can use a vertical highway, then their travel-time distance is at most 2δ , implying that the travel time from p to q is also at most 2δ .

It remains to consider the case that p, q' cannot use a vertical highway. This means that the line pq' forms an angle larger than α with the y -axis, see Figure 3b. Note that $\eta(q' - p)$ still has a geometric meaning: There is a path from p to the highway, then *backwards* along the highway, then straight to q , see Figure 3c.

The η -distance measures the whole travel time, but counting the time on the highway *negative*. Reflecting the last segment of this path back around the highway, we obtain a path from p to q with travel time $\eta(q' - p)$, still counting time spent on the highway negative. But now observe that this path self-intersects in a point x , see Figure 3d. Let x' be the reflection of x . Then $\eta(q' - p) = d(p, x) + \eta(x' - x) + d(x', q') = d(p, x) + \eta(x' - x) + d(x, q) \geq d(p, x) + d(x, q) \geq d(p, q)$ (using $\eta \geq 0$). It follows $d(p, q) \leq \eta(q' - p) \leq 2\delta$. \square

Again we found an optimal highway by computing a minimal enclosing shape. Interestingly, the shape to be minimized is not the unit circle under the travel-time metric. If the highway is the y -axis then the unit circle in the travel-time metric is the convex hull of the points $(0, v)$ and $(0, -v)$ and the Euclidean unit circle centered at the origin.

It is remarkable that we can compute a highway that realizes the optimal diameter in linear time, without actually computing the travel-time diameter. In fact, computing the diameter cannot be done in linear time in the algebraic decision-tree model.

Lemma 2 *In the algebraic decision-tree model the computation of the travel-time diameter takes $\Omega(n \log n)$ time.*

Proof. The following problem has a lower bound of $\Omega(n \log n)$ in the algebraic decision-tree model: Given two sets A and B of n real numbers, is $A \cap B = \emptyset$? We show how to transform this problem in linear time into a decision instance of the diameter problem.

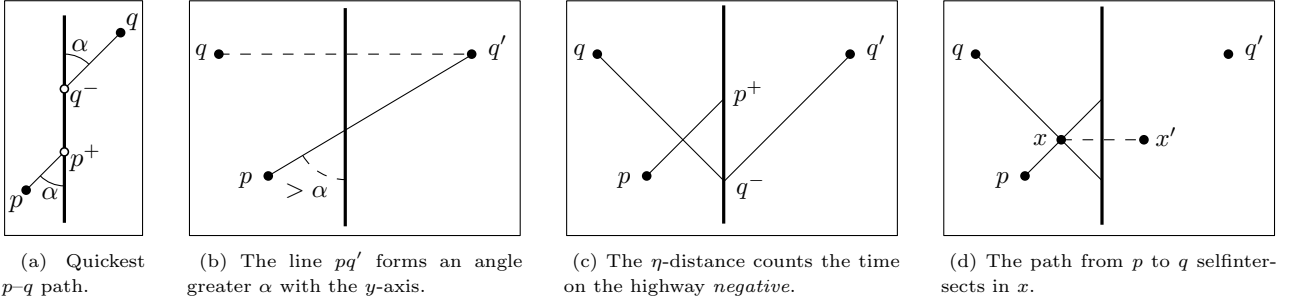


Figure 3: Optimal highway under the L_2 -metric.

Our instance consists of a set A' of n points and a set B' of n points, computed from A and B . All points lie on the unit circle. We first scale all numbers in A and B so that they are close to zero (depending on v). For each a in A , we create the point (x, a) , where $x = \sqrt{1 - a^2}$. For each b in B , we create the point $(x, -b)$, where $x = -\sqrt{1 - b^2}$. Note that since the points are close to the x -axis, no vertical highway can be used to speed up the connection between A' and B' , and so the diameter of the set is simply the Euclidean diameter. It follows that the diameter of $A' \cup B'$ is 2 if and only if A and B contain a common number. \square

To summarize we observe that for the Euclidean metric the smallest enclosing figure characterizes the optimal highway, but other than in the case of orthogonal travel to and from the highway, the size of the smallest enclosing figure does not give us the travel-time diameter. Thus we cannot use the rotating-calipers algorithm as in the proof of Theorem 5 to find the the optimal orientation of the highway for the Euclidean metric.

6 The optimal axis-aligned highway cross in the L_1 -metric

In this problem we have been unable to characterize the optimal solution by a smallest enclosing shape. There is no natural “center” for the problem: sometimes there are critical paths where a point connects to the highway that is *further* away.

Theorem 7 *The optimal axis-aligned speed- v highway cross under the L_1 -metric can be computed in $O(n^{4+\epsilon})$ time, for any $\epsilon > 0$.*

Proof. The optimal solution corresponds to the lowest point on the upper envelope of the pairwise distance functions, see Figure 4. Since these functions are of constant description complexity, their upper envelope can be computed in $O(n^{4+\epsilon})$ time for any $\epsilon > 0$, as Sharir and Agarwal (1995) show. \square

Suppose we could characterize the travel-time diameter given the optimal highway cross to get a compact list L of candidate values as in the case of the infinite-speed highway cross. Then we could do binary search on L using the decision algorithm de-

scribed in the following subsection. Given the travel-time graph Γ_{pq} for each pair of points p and q (see Figure 4), consider the upper envelope over all these graphs. It seems plausible that the minimum of the upper envelope is determined by three graphs, i.e., by at most six points. However, we do not know whether the minimum of the upper envelope of these three graphs is the same as the minimum of the upper envelope over all graphs—it could be less. If we could show that the two minima are equal we could even apply Chan’s technique (1999), just as Cardinal and Langerman (2006), in order to get a randomized algorithm for the optimization problem whose expected running time is asymptotically the same as the running time of the decision algorithm.

6.1 The decision problem

We now present an algorithm that decides for a given $\delta > 0$ whether there is an axis-aligned speed- v highway cross such that the resulting travel-time diameter is at most δ . This is used as a subroutine for one of the approximation algorithms in Section 7.

Theorem 8 *Given a set S of n points in the plane, a speed $v > 1$, and a parameter $\delta > 0$, we can decide in $O(n^2 \alpha(n) \log n)$ whether there is an axis-aligned highway cross of speed v such that the travel-time diameter of S is at most δ .*

Proof. For points $\sigma, p, q \in \mathbb{R}^2$, let $d_\sigma(p, q)$ denote the travel-time distance between p and q , assuming an axis-aligned highway cross with speed v has been placed (with center) at σ . We define the region

$$R(p, q) := \{\sigma \in \mathbb{R}^2 \mid d_\sigma(p, q) \leq \delta\}.$$

We observe that the answer to the decision problem is positive if and only if $\bigcap_{p, q \in S} R(p, q)$ is not empty.

The shape of the region $R(p, q)$ depends on δ . Let w, h be the horizontal and vertical distance of points p and q . If $\delta < (w + h)/v$, then $R(p, q)$ is empty. If $(w + h)/v \leq \delta < \min\{w + h/v, h + w/v\}$, then $R(p, q)$ consists of two convex quadrilaterals. If $\min\{w + h/v, h + w/v\} \leq \delta < \max\{w + h/v, h + w/v\}$, then $R(p, q)$ is infinite in one (axis-parallel) direction. If $\max\{w + h/v, h + w/v\} \leq \delta < w + h$, then $R(p, q)$ is infinite in both axis-parallel directions. Finally, if $w + h \leq \delta$, then $R(p, q) = \mathbb{R}^2$. See Figure 5.

Let us call a planar region F (a, b) -monotone if for every point $(x, y) \in F$ and any $\lambda \geq 0$ the

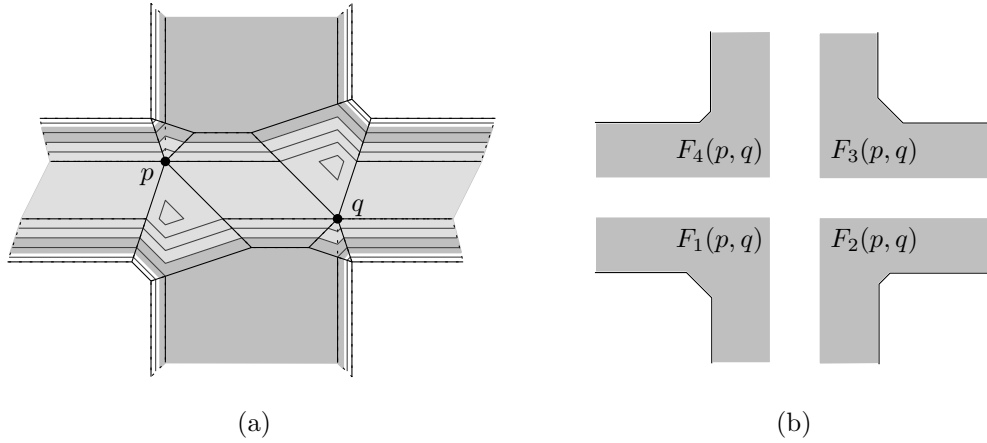


Figure 5: (a) the regions $R(p, q)$ when $\min\{w + h/v, h + w/v\} \leq \delta < \max\{w + h/v, h + w/v\}$ (light gray region) and when $\max\{w + h/v, h + w/v\} \leq \delta < w + h$ (dark and light gray regions). (b) the dark and light gray regions can be expressed as the intersection of the four types of regions $F_1, F_2, F_3,$ and F_4 .

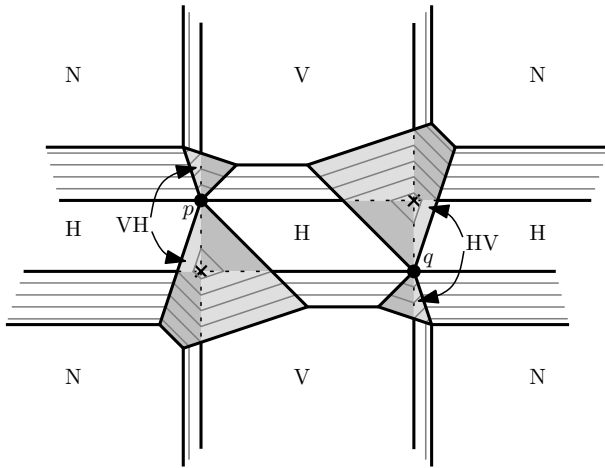


Figure 4: Travel-time distance between two points p and q for an axis-aligned highway cross centered at (x, y) . The graph of this function has 31 faces of 15 different orientations. The two points marked \times are the lowest, i.e., those where the corresponding highway crosses minimize the travel time from p to q . The thin dark gray lines are contour lines. If a highway cross is centered in a V- or H-region, the vertical and horizontal highway is used by a quickest p - q path. In the VH- and HV-region both highways are used in the corresponding order. In the N-regions no highway is used.

point $(x + \lambda a, y + \lambda b)$ is also in F . We observe that $R(p, q)$ can be expressed as the intersection of four regions $F_i(p, q)$, $i = 1, 2, 3, 4$, where $F_1(p, q)$ is $(1, 1)$ -monotone, $F_2(p, q)$ is $(-1, 1)$ -monotone, $F_3(p, q)$ is $(-1, -1)$ -monotone, and $F_4(p, q)$ is $(1, -1)$ -monotone. Figure 5 shows an example of $R(p, q)$, which can be expressed as the intersection of $F_1, F_2, F_3,$ and F_4 . Each region is bounded by a polygonal curve of constant complexity, and so we can compute $F_i := \bigcap_{p, q \in S} F_i(p, q)$ by a simple plane sweep. It is the lower envelope of a set of $O(n^2)$ line segments, and hence has complexity $O(n^2 \alpha(n))$, as observed by Sharir and Agarwal (1995). The intersections $F_1 \cap F_3$ and $F_2 \cap F_4$ can again be computed by a plane sweep in this time. We are left with two regions of complexity $O(n^2 \alpha(n))$, and we need to determine whether their intersection is empty. While we do not know how to bound the complexity of this region, we can test emptiness in $O(n^2 \alpha(n) \log n)$ time, by a simple plane sweep that stops as soon as a point in the intersection is found. Since any intersection between edges of the two regions implies that the intersection is not empty, this runs in the claimed time bound. \square

6.2 Further observations

The optimum axis-aligned highway cross for finite speed need not be contained in the strip cross for infinite speed, see Figure 6: take the points $(-2, 1), (-1, 2), (1, 2), (2, 1), (2, -1), (1, -2), (-1, -2), (-2, -1)$ —i.e., an octagon, contained in the strip cross $([-1, 1] \times \mathbb{R}) \cup (\mathbb{R} \times [-1, 1])$ —plus the two points $(-a, 0)$ and $(a, 1)$ for $a > 3v$. Then the coordinate axes are the optimum highway cross for infinite speed. It has travel-time diameter 2, so the above two strips are the optimum cover. But for any finite speed $v > 1$, the highway cross $x = a$ and $y = 0$ yields a diameter of $(2a + 1)/v$, which is better than the diameter $(2a/v) + 1$ caused by the coordinate axes being highways.

The following argument also rules out any simple

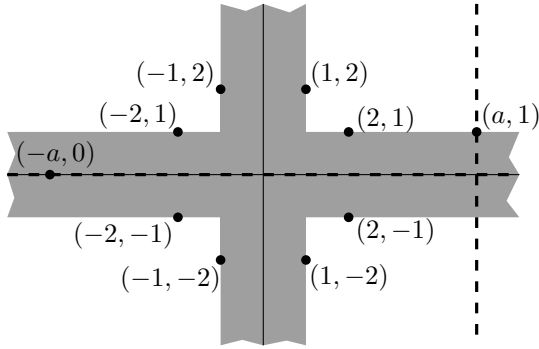


Figure 6: Here the optimum speed- v highway cross (dashed) is not contained in the optimum speed- ∞ strip cross (shaded).

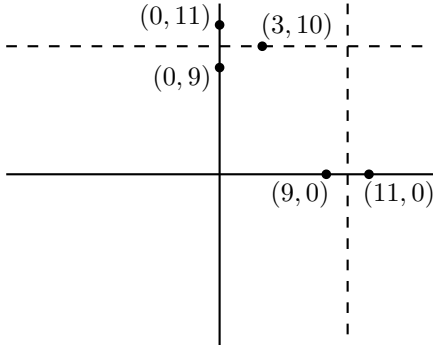


Figure 7: After adding the point $(3, 10)$ the optimal speed- ∞ highway cross changes (from solid to dashed), but the new point does not occur in any diametral pair.

incremental algorithm. Note that the minimum must occur at a vertex of the upper envelope. Thus there are always at least three diametral pairs which simultaneously realize that optimum distance. However, even for infinite speed there are point sets such that the addition of one point changes the diameter, and the new point does not occur in any diametral pair. An example (see Figure 7) for infinite speed is given by the points $(0, 9)$, $(0, 11)$, $(9, 0)$, $(11, 0)$ since now the coordinate axes are an optimal highway cross—with a diameter of 0. If we add the point $(3, 10)$, the optimal highway cross is centered at $(10, 10)$ and has a diameter of 2.

7 Fast constant-factor approximations for the optimal axis-aligned highway cross

Given a set S of n points. Let C be the smallest enclosing cross for S , and let h_1, h_2 be the middle line of each strip of C . We call h_1, h_2 the *median highways* for S .

Lemma 3 *The travel-time diameter δ_{med} of the median highways (with speed v) is at most $2 + 1/v$ times the travel-time diameter δ_{opt} of an optimal axis-aligned speed- v highway cross for S . There are point sets S where for $v \geq \sqrt{3}$ the travel-time diameter of the median highways is at least $2 - 1/(v + 2)$ times the optimum.*

Proof. We can scale S such that its L_1 -diameter is 2—this does not change the travel-time ratio. Let w be the width of C after scaling. Observe that $\delta_{\text{opt}} \geq 2/v$, as there are points at L_1 -distance 2. Furthermore, we have $\delta_{\text{opt}} \geq w$, since using the optimal highways at infinite speed cannot achieve diameter less than w .

On the other hand, $\delta_{\text{med}} \leq w + (2+w)/v$, since any point can reach a point on the highways at distance at most $w/2$, and the maximum distance of such points on the highways is at most $2+w$. This implies $\delta_{\text{med}} \leq (1+1/v)w + 2/v \leq (1+1/v)\delta_{\text{opt}} + \delta_{\text{opt}} = (2+1/v)\delta_{\text{opt}}$.

For the lower bound example, let speed $v > 1$ be given, and set parameter $\omega = 1/(v + 2)$. We will construct a point set S such that the smallest enclosing cross has width 2ω , the median highway has travel-time diameter $4/v - 2/(v(v + 2))$, and the optimal highway cross has travel-time diameter at most $2/v$, implying the lower bound.

Let $\varepsilon > 0$ be very small. Our point set S consists of the points $(1, 0)$, $(-1, 0)$, $(0, 1)$, $(0, -1)$, $(-2\omega, \omega + \varepsilon)$, $(-2\omega, -\omega - \varepsilon)$, $(\varepsilon, -2\omega)$, as in Figure 8. We claim that S has a unique smallest enclosing strip of width 2ω , centered around the lines $x = -\omega$ and $y = -\omega$. Indeed, the line $x = 0$ must be in the vertical strip (otherwise the horizontal strip would have width at least 2), while the line $y = 0$ must be in the horizontal strip. Similarly, the line $x = -2\omega$ must be in the vertical strip as well, and this now fixes the vertical strip of width 2ω around the line $x = -\omega$. It follows that the remaining point $(\varepsilon, -2\omega)$ is in the horizontal strip, fixing that strip around $y = -\omega$.

The median highway cross has travel-time diameter $2\omega + (2+2\omega)/v = 4/v - 2/(v(v+2))$ if the diameter is less than or equal to the L_1 distance of these points (note that the diameter is determined by $(1, 0)$ and $(0, 1)$.) That is, for $v \geq \sqrt{3}$ the median highway cross has travel-time diameter $4/v - 2/(v(v + 2))$. Consider now a highway cross with center at the origin. The four outer points and the point $(\varepsilon, -2\omega)$ can be reached from the origin within travel time $1/v$. The remaining two points can be reached with travel time $\omega(1 + 2/v)$ (ignoring all ε -terms). The travel-time distance between these two points is 2ω , and so the travel-time diameter is bounded by

$$\max\{2/v, 1/v + \omega(1 + 2/v), 2\omega\} = 2/v.$$

□

We can improve the result in Lemma 3 by a simple observation.

Theorem 9 *Given a set S of n points we can compute in $O(n \log n)$ time an axis-aligned highway cross whose travel-time diameter is at most $1 + \sqrt{2}$ times the travel-time diameter of an optimal axis-aligned speed- v highway cross for S .*

Proof. According to Theorem 3 the median highways can be computed in $O(n \log n)$ time. According to Lemma 3 they yield a factor- $(2 + 1/v)$ approximation for the optimal travel-time diameter.

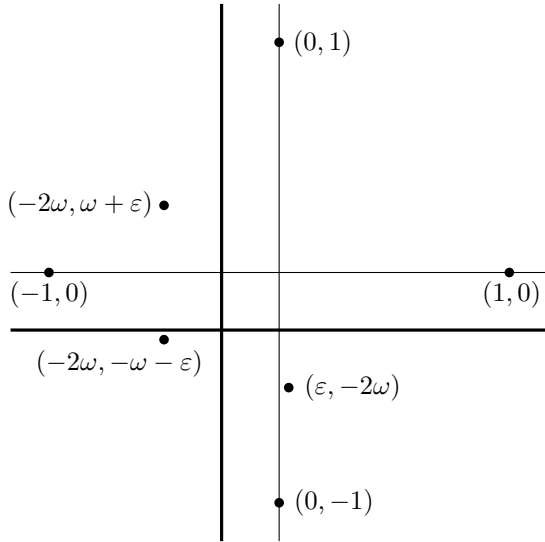


Figure 8: The median highway for $v = 2$.

Note that the approximation factor tends to 3 when the speed goes to 1. Clearly *not* building a highway cross is a factor- v approximation. Balancing out the two terms yields $\min\{2 + 1/v, v\} \leq 1 + \sqrt{2}$. \square

We can do better if we are only interested in the optimal travel-time diameter.

Lemma 4 *Let s be any point in S . Let H_s be the highway cross that minimizes the maximum travel time to s . The travel-time diameter of H_s is at most twice the travel-time diameter δ_{opt} of an optimal axis-aligned speed- v highway cross for S .*

Proof. Let $\{p, q\}$ be a pair of points in S and let $s' \in S$ be a point of maximum travel-time distance from s given H_s . We denote by d_s the metric induced by H_s and by d_{opt} the metric induced by the optimal highway cross. Then $d_s(s, p) \leq d_s(s, s') \leq d_{\text{opt}}(s, s') \leq \delta_{\text{opt}}$ and, by symmetry, $d_s(s, q) \leq \delta_{\text{opt}}$. This yields $d_s(p, q) \leq d_s(p, s) + d_s(s, q) \leq 2\delta_{\text{opt}}$. \square

Note that H_s is usually *not* centered at s (e.g. consider the set $S = \{(0, 1), (1, 0)\}$ whose optimal highway cross is centered at the origin).

Based on the constant-factor approximation from Theorem 9 we can use a modification of the decision procedure described in Section 6.1 combined with binary search to get the following.

Theorem 10 *Given a set S of n points in the plane, we can compute in $O(\log(1/\varepsilon)\alpha(n)n \log n)$ time a $(2 + \varepsilon)$ -approximation for the travel-time diameter of S under the optimal axis-aligned speed- v highway cross.*

Proof. Let s be any point in S and H_s be the highway cross that minimizes the maximum travel time to s .

According to Theorem 4 the travel-time diameter δ_s given H_s is at most twice the travel-time diameter δ_{opt} given an optimal axis-aligned speed- v highway cross for S , i.e., $\delta_s \leq 2\delta_{\text{opt}}$, so a $(1 + \varepsilon/2)$ -

approximation δ_ε to δ_s is a $(2 + \varepsilon)$ -approximation for δ_{opt} .

We now describe how to compute a $(1 + \varepsilon)$ -approximation for δ_s by binary search. Recall that the median highways yield a travel-time diameter of $\delta_{\text{med}} \leq (1 + \sqrt{2})\delta_{\text{opt}} \leq 3\delta_{\text{opt}} \leq 3\delta_s$ of δ_s , see Theorem 9. The median highways can be computed in $O(n \log n)$ time according to Theorem 3.

Now we conceptually subdivide the interval $I = [0, 2\delta_{\text{med}}]$ into at most $N = 6/\varepsilon$ pieces of length $\delta_{\text{med}} \cdot \varepsilon/3$, and denote the increasing sequence of interval endpoints by $\Delta = (\delta_1, \dots, \delta_N)$. Since $\delta_s \leq 2\delta_{\text{opt}} \leq 2\delta_{\text{med}}$, we know that δ_s lies in I . Hence there is an index $i \in \{1, \dots, N\}$ such that $\delta_i < \delta_s \leq \delta_{i+1}$.

Setting $\delta_\varepsilon = \delta_{i+1}$ we find that $\delta_s \geq \delta_i = \delta_\varepsilon - \delta_{\text{med}} \cdot \varepsilon/3$. This yields $\delta_\varepsilon \leq \delta_s + \delta_{\text{med}} \cdot \varepsilon/3 \leq \delta_s + 3\delta_s \cdot \varepsilon/3 \leq (1 + \varepsilon)\delta_s$. Thus δ_ε is indeed a $(1 + \varepsilon)$ -approximation of δ_s .

For a given $\delta > 0$ we can run a modification of the decision algorithm described in the proof of Theorem 8. Our modification considers for each point $q \in S \setminus \{s\}$ the $(n - 1)$ travel-time distances $d_\delta(s, q)$ between s and q in order to decide whether there is a highway cross such that the maximum travel-time distance to s is at most δ . Each such test takes $O(\alpha(n)n \log n)$ time.

Using $O(\log(1/\varepsilon))$ calls to this decision procedure, we can determine δ_ε by binary search on Δ . \square

If we are willing to invest more time, we can even get a $(1 + \varepsilon)$ -approximation of the optimal travel-time diameter δ_{opt} .

Theorem 11 *Given a set S of n points in the plane, we can compute in $O(\log(1/\varepsilon)\alpha(n)n^2 \log n)$ time a $(1 + \varepsilon)$ -approximation for the travel-time diameter of S under the optimal axis-aligned speed- v highway cross.*

Proof. We again first compute the median highways to get an upper bound δ_{med} for the optimal travel-time diameter δ_{opt} and then do binary search. We can now use the interval $[0, \delta_{\text{med}}]$, which contains δ_{opt} . We stop when the interval size is sufficiently small, i.e., at most $\delta_{\text{med}} \cdot \varepsilon/3$. This time we use the decision algorithm described in the proof of Theorem 8 without any modifications. \square

8 Concluding Remarks

There are many ways how this problem can be extended. First, can we compute an optimal highway with arbitrary orientation under the Euclidean metric in $o(n^2)$ (worst-case) time? Second, consider highways with different speeds, different slopes, or bounded lengths. Third, suppose an existing network of (axis-parallel) highways and a real $\ell > 0$ is given. Where to place a new (axis-parallel) highway segment of length ℓ in order to minimize the travel-time diameter of the resulting network?

References

- Abellanas, M., Hurtado, F., Icking, C., Klein, R., Langetepe, E., Ma, L., Palop del Río, B. & Sácristan, V. (2001), Proximity problems for time metrics induced by the L_1 metric and isothetic networks, in 'Actas de los IX Encuentros de Geometría Computacional', Girona, pp. 175–182.
- Abellanas, M., Hurtado, F., Icking, C., Klein, R., Langetepe, E., Ma, L., Palop del Río, B. & Sácristan, V. (2003), 'Voronoi diagram for services neighboring a highway', *Information Processing Letters* **86**, 283–288.
- Aichholzer, O., Aurenhammer, F. & Palop del Río, B. (2004), 'Quickest paths, straight skeletons, and the city Voronoi diagram', *Discrete Comput. Geom.* **31**(1), 17–35.
- Bae, S. W. & Chwa, K.-Y. (2004), Voronoi diagrams with a transportation network on the euclidean plane, in 'Proc. 15th Annu. Internat. Sympos. Algorithms Comput. (ISAAC'04)', Vol. 3341 of *Lecture Notes in Computer Science*, Springer-Verlag, pp. 101–112.
- Bae, S. W. & Chwa, K.-Y. (2005), Shortest paths and Voronoi diagrams with transportation networks under general distances, in 'Proc. 16th Annu. Internat. Sympos. Algorithms Comput. (ISAAC'05)', Vol. 3827 of *Lecture Notes in Computer Science*, Springer-Verlag, pp. 1007–1018.
- Bae, S. W., Kim, J.-H. & Chwa, K.-Y. (2005), L_1 shortest paths with isothetic roads, Technical report, Korea Advanced Institute of Science and Technology.
- Cardinal, J. & Langerman, S. (2006), Min-max-min geometric facility location problems, in 'Proc. of the European Workshop on Computational Geometry (EWCG'06)', Delphi, pp. 149–152.
- Chan, T. M. (1999), 'Geometric applications of a randomized optimization technique', *Discrete Comput. Geom.* **22**(4), 547–567.
- Frederickson, G. N. & Johnson, D. B. (1984), 'Generalized selection and ranking: sorted matrices', *SIAM J. Comput.* **13**, 14–30.
- Görke, R. & Wolff, A. (2005), Constructing the city Voronoi diagram faster, in 'Proc. 2nd Int. Symp. on Voronoi Diagrams in Science and Engineering (VD'05)', pp. 162–172.
- Sharir, M. & Agarwal, P. K. (1995), *Davenport-Schinzel Sequences and Their Geometric Applications*, Cambridge University Press, Cambridge.
- Toussaint, G. T. (1983), Solving geometric problems with the rotating calipers, in 'IEEE MELECON', pp. 1–4.