Shortest Paths and Voronoi Diagrams with Transportation Networks Under General Distances
[Extended Abstract]

Sang Won Bae and Kyung-Yong Chwa

Division of Computer Science, Department of EECS,
Korea Advanced Institute of Science and Technology, Daejeon, Korea
{swbae, kychwa}@jupiter.kaist.ac.kr

Abstract. Transportation networks model facilities for fast movement on the plane. A transportation network, together with its underlying distance, induces a new distance. Previously, only the Euclidean and the L_1 distances have been considered as such underlying distances. However, this paper first considers distances induced by general distances and transportation networks, and present a unifying approach to compute Voronoi diagrams under such a general setting. With this approach, we show that an algorithm for convex distances can be easily obtained.

1 Introduction

Transportation networks model facilities for fast movement on the plane. They consist of roads and nodes; *roads* are assumed to be segments along which one can move at a certain fixed speed and *nodes* are endpoints of roads. We assume that there are no crossings among roads but roads can share their endpoints as nodes. We thus define a transportation network as a plane graph whose vertices are nodes and whose edges are roads with speeds assigned. Also, we assume that one can access or leave a road through any point on the road. This, as Aichholzer et al. [3] pointed out, makes the problem distinguishable from and more difficult than those in other similar settings such as the airlift distance.

In the presence of a transportation network, the distance between two points is defined to be the shortest elapsed time among all possible paths joining the two points using the roads of the network. We call such an induced distance a *transportation distance*. (In other literature [2, 3], it is called a *time distance* or a *city metric*.) More precisely, a transportation distance is induced on the plane by a transportation network and its underlying distance that measures the distance between two points without roads.

Since early considerations for roads [7, 17, 20], fundamental geometric problems, such as shortest paths and Voronoi diagrams, under transportation distances have been receiving much attention recently [1, 2, 3, 4, 8, 18]. However, underlying distances considered in the literature were only the Euclidean distance [2, 4] and the L_1 distance [1, 3, 8]. Since transportation distances have

X. Deng and D. Du (Eds.): ISAAC 2005, LNCS 3827, pp. 1007–1018, 2005.

quite different properties depending on their underlying distances, there has not been a common approach to extend such problems to more general underlying distances. This paper thus considers geometric problems, in particular, shortest paths and Voronoi diagrams under transportation distances induced from general distances.

The Results. This paper presents, to the best of our knowledge, the first result that studies general underlying distances and gives algorithms for computing Voronoi diagrams with transportation networks, in Section 3. More precisely, we classify *transportable* distance functions where transportation networks and transportation distances are well-defined. Transportable distances include asymmetric convex distances, nice metrics, and even transportation distances. In this general setting, a unifying approach to compute Voronoi diagrams is presented.

As a special case of transportable distances, we take the convex distances into account in Section 4. Based on the approach in Section 3 together with geometric and algorithmic observations on convex distances, we first obtain an efficient and practical algorithm that computes the Voronoi diagram with a transportation network under a convex distance.

For the L_1 metric, previous work considered only isothetic networks and a single or a constant number of speeds for roads [3, 8]. Our results first deal with more general transportation networks, which have no restriction except for straightness; the roads can have arbitrarily fixed speeds and directions. For the Euclidean metric, we obtain the same time and space bounds as those of the previously best results [4].

Note that a resulting diagram of our algorithm is in fact a refined diagram of the real Voronoi diagram so that it consists of shortest-paths information in each cell and it can also serve as a shortest path map structure.

2 Preliminaries

2.1 Transportation Networks Under General Distances

Here, we let $d : \mathbb{R}^2 \times \mathbb{R}^2 \to \mathbb{R}$ be a total distance function. For the Euclidean and the L_1 distances, a transportation network can be sufficiently represented as a planar straight-line graph. If, however, we consider more general distances, the meaning of "straight" should be reconsidered. Note that a straight segment is a shortest path or a geodesic on the Euclidean plane or on the L_1 plane. Geodesics, in general, naturally generalize straight segments, and a road can be defined to be a segment along a geodesic. Thus, in order to build a transportation network under d, d needs to admit a geodesic between any two points on the plane.

In this step, we define a *transportable distance* that satisfies several axioms. It is easy to observe that distances that admit geodesics and are possibly asymmetric are transportable. We will show that transportable distances admit a geodesic between any two points on the plane. We call a distance d over $\mathbb{R}^2$ *transportable* if the following properties hold:

1. d is non-negative and $d(p,q) = d(q,p) = 0$ iff $p = q$ and d satisfies the triangle inequality.
2. The backward topology induced by d on X induces the Euclidean topology.
3. The backward d-balls are bounded with respect to the Euclidean metric.
4. For any two points p and r, there exists a point $q \notin \{p,r\}$ such that $d(p,q) + d(q,r) = d(p,r)$.

Distances with Condition 1 are called *quasi-metrics* and they induce two associated topologies by two families of open balls, $B_d^+(x,\epsilon) = \{y | d(x,y) < \epsilon\}$ and $B_d^-(x,\epsilon) = \{y | d(y,x) < \epsilon\}$, called *forward* and *backward*, respectively [9, 16]. Here, we consider only the backward case since we shall take only the *inward* Voronoi diagrams into account; Voronoi sites are supposed to be static and fixed. In fact, Conditions 2-4 of the definition of transportable distances mimic those of *nice metrics* in the sense of Klein and Wood [13]. By these conditions, $(\mathbb{R}^2, d)$ is known to be backward-complete [16]. Also, we can find a geodesic, whose length is the same as the distance, between two points in $\mathbb{R}^2$.

Lemma 1. *Let d be a transportable distance. Then, for any two points p and r, there exists a path π from p to r such that for each point q on π the equality $d(p,r) = d(p,q) + d(q,r)$ holds.*

Such paths are called *d-straight* and they generalize straight line segments with respect to the Euclidean metric, indeed. This lemma can be shown by Menger's *Verbindbarkeitssatz* that implies the existence of d-straight paths in a complete metric space [19].

Now, we are able to define a transportation network under a transportable distance d. Since d can be asymmetric, roads in a transportation network may have an orientation. Thus, throughout this paper, a transportation network under d is defined to be a directed plane graph $G = (V,E)$ such that any edge e in E is a segment of a d-straight line and has its own weight $v(e) > 1$, called *speed*. We note that an edge has an orientation so that it can be regarded as a one-way road. And we call edges in E *roads* and vertices in V *nodes*, and roads and nodes may denote d-straight paths and points by themselves referenced. Two incident nodes of a road e is identified by $p_1(e)$ and $p_2(e)$, where e has the orientation toward $p_2(e)$ from $p_1(e)$. Note that any crossings among roads can be removed by introducing additional nodes. An anomaly occurs when we think of a two-way road. Thus, we allow coincidence only for two roads having the same incident nodes.

Now, we consider a distance d_G induced by a transportable distance d and a transportation network $G = (V,E)$ under d, which can be defined as follows:

$$d_G(p,q) = \min_{P=(p_1,\cdots,p_\ell)\in\mathcal{P}(p,q)} \sum_{i=1}^{\ell-1} \frac{1}{v_i} d(p_i, p_{i+1}),$$

where $\mathcal{P}(p,q)$ is the set of all piecewise d-straight paths from p to q and $v_i = v(e)$ if there exists an oriented road $e \in E$ such that the path from p_i to p_{i+1} passes along e, otherwise, $v_i = 1$. We call d_G the *transportation distance* induced by d

and G. Note that transportable distances include nice metrics, convex distances, and even transportation distances.

2.2 Needles

Bae and Chwa [4, 4] defined a needle as a generalized Voronoi site, which is very useful for a transportation network. In fact, the concept of needles was first proposed by Aichholzer, Aurenhammer, and Palop [3] but needles were not thought of as Voronoi sites in their results.

Here, we consider a needle under a transportable distance d. A needle $\mathbf{p}$ under d can be represented by a 4-tuple $(p_1(\mathbf{p}), p_2(\mathbf{p}), t_1(\mathbf{p}), t_2(\mathbf{p}))$ with $t_2(\mathbf{p}) \geq t_1(\mathbf{p}) \geq 0$, where $p_1(\mathbf{p}), p_2(\mathbf{p})$ are two endpoints and $t_1(\mathbf{p}), t_2(\mathbf{p})$ are additive weights of the two endpoints, respectively. In addition, a needle under d is d-straight in a sense that it can be viewed as a set of weighted points on the d-straight path from $p_2(\mathbf{p})$ to $p_1(\mathbf{p})$. Other terms associated with a needle under d are determined in a similar way with the Euclidean case. Thus, let $s(\mathbf{p})$ be the set of points on the d-straight path from $p_2(\mathbf{p})$ to $p_1(\mathbf{p})$, and $v(\mathbf{p})$ be the *speed* of $\mathbf{p}$, defined by $d(p_2(\mathbf{p}), p_1(\mathbf{p}))/(t_2(\mathbf{p}) - t_1(\mathbf{p}))$.

The distance from any point x to a needle $\mathbf{p}$ is measured as $d(x, \mathbf{p}) = \min_{y \in s(\mathbf{p})}\{d(x, y) + w_{\mathbf{p}}(y)\}$, where $w_{\mathbf{p}}(y)$ is the weight assigned to y on $\mathbf{p}$, given as $w_{\mathbf{p}}(y) = t_1(\mathbf{p}) + d(y, p_1(\mathbf{p}))/v(\mathbf{p})$, for all $y \in s(\mathbf{p})$.

For the Euclidean case, the Voronoi diagram for pairwise non-piercing needles has been shown to be an abstract Voronoi diagram. Two needles are called *non-piercing* if, and only if, the bisector between them contains at most one connected component. For more details, we refer to [4, 4].

3 Voronoi Diagrams Under Transportation Distances

3.1 d_G-Straight Paths and Needles

As noted in the previous section, a transportable distance d and a transportation network G induce a new distance d_G and d_G-straight paths. The structure of any d_G-straight path can be represented by a string of $\{\mathsf{S}, \mathsf{T}\}$, where S denotes a d-straight path without using any road and T denotes that along a road.

Let us consider a single road e as a simpler case. Given a transportation network G with only one road e, a d_G-straight path is of the form STS or its substring except for SS. This is quite immediate; paths represented by longer strings than STS can be reduced since a road is d-straight and d satisfies the triangle inequality. Thus, any d_G-straight path P from p to q using a road e can be represented as $P = (p, p', q', q)$, where p' is the entering point to e and q' is the exiting point to q. We then call q' a *footpoint* of q on e. A point may have several or infinitely many footpoints on a road. Let $FP_e(q)$ be the set of footpoints of q on e for all d_G-straight paths from any point to q using e. Let us consider a total order $\prec_e$ on the points on a road e, where $x \prec_e y$ for $x, y \in e$ if the orientation of the d-straight path from x to y is equivalent to that of e. Then, the following property of footpoints can be shown.

Lemma 2. *Let $F \subseteq e$ be connected. Then, the following are equivalent.*

1. *F is a connected component of $FP_e(q)$.*
2. *F has the least point q_0 with respect to $\prec_e$ such that $q_0 \in FP_e(q)$, and $d(q_1, q) = d(q_1, q_2)/v(e) + d(q_2, q)$ for any $q_1, q_2 \in F$ with $q_1 \prec_e q_2$.*

For such a shortest path $P = (p, p', q', q)$, we can find a needle $\mathbf{q}$ such that $d(p, \mathbf{q}) = d_G(p, q)$. Such a needle $\mathbf{q}$ is said to be *produced on a road e from a point q for a footpoint q'*, and can be defined by setting parameters as follows; $p_1(\mathbf{q}) = q'$, $p_2(\mathbf{q}) = p_1(e)$, $t_1(\mathbf{q}) = d(q', q)$, and $t_2(\mathbf{q}) = d(p'q')/v(e) + d(q', q)$. We let $\sigma_e(q)$ be the set of needles produced on e from q for the least footpoint of every connected component in $FP_e(q)$. Also, both of $FP_e(\cdot)$ and $\sigma_e(\cdot)$ can be naturally extended for needles since shortest paths to a needle under d_G are also d_G-straight paths.

Lemma 3. *If a transportation network G under d contains only one road e, for a point x and a needle $\mathbf{p}$, $d_G(x, \mathbf{p}) = d(x, \sigma_e(\mathbf{p}) \cup \{\mathbf{p}\})$.*

Now, we consider multiple roads. Let $\sigma_G(\mathbf{p}) = \bigcup_{e \in E} \sigma_e(\mathbf{p})$ and $\sigma_G(A) = \bigcup_{\mathbf{p} \in A} \sigma_G(\mathbf{p})$ for a set A of needles. Since d_G-straight paths may pass through several roads, we apply $\sigma_G(\cdot)$ repeatedly. We thus let $\sigma_G^k(\mathbf{p}) = \sigma_G(\sigma_G^{k-1}(\mathbf{p}))$ and $\sigma_G^0(\mathbf{p}) = \{\mathbf{p}\}$. Also, we let $\mathcal{S}_\mathbf{p}^k$ denote $\bigcup_{i=0}^k \sigma_G^i(\mathbf{p})$ and $\mathcal{S}_\mathbf{p}$ denote $\mathcal{S}_\mathbf{p}^\infty$.

Theorem 4. *Given a transportable distance d and a transportation network G under d, for a point x and a needle $\mathbf{p}$,*

$$d_G(x, \mathbf{p}) = d(x, \mathcal{S}_\mathbf{p}).$$

Proof. We first define $d_G^k(p, q)$ be the length of a shortest path from p to q where the path passes through at most k roads in G. Surely, $d_G(p, q) = d_G^\infty(p, q)$.

We claim that $d_G^k(x, \mathbf{p}) = d(x, \mathcal{S}_\mathbf{p}^k)$, which directly implies the theorem. We prove this by induction. Lemma 3 gives us an inductive basis. We have $d(x, \mathcal{S}_\mathbf{p}^{\ell+1}) = \min\{d(x, \mathcal{S}_\mathbf{p}^\ell), d(x, \sigma_G^{\ell+1}(\mathbf{p}))\} = \min\{d(x, \mathcal{S}_\mathbf{p}^\ell), d(x, \sigma_G(\sigma_G^\ell(\mathbf{p})))\}$. By inductive hypothesis and Lemma 3, the equation is evaluated as $d(x, \mathcal{S}_\mathbf{p}^{\ell+1}) = \min\{d_G^\ell(x, \mathbf{p}), d_G^1(x, \sigma_G^\ell(\mathbf{p}))\}$.

As pointed out in the proof of Lemma 3, $d_G^1(x, \sigma_G^\ell(\mathbf{p}))$ implies a shortest path to a needle in $\sigma_G^\ell(\mathbf{p})$ using exactly one road in G, and further a shortest path to $\mathbf{p}$ using exactly $\ell + 1$ roads. Therefore, we conclude $d(x, \mathcal{S}_\mathbf{p}^{\ell+1}) = d_G^{\ell+1}(x, \mathbf{p})$, implying the theorem. □

Theorem 4 says a nice relation between needles and roads. Furthermore, it directly implies that the Voronoi diagram $\mathcal{V}_d(\mathcal{S})$ under d for $\mathcal{S}$ induces the Voronoi diagram $\mathcal{V}_{d_G}(S)$ under d_G for S, where $\mathcal{S}$ denotes $\bigcup_{p \in S} \mathcal{S}_p$. In other words, any Voronoi region in $\mathcal{V}_d(\mathcal{S})$ is completely contained in a Voronoi region in $\mathcal{V}_{d_G}(S)$, i.e., $\mathcal{V}_{d_G}(S)$ is a sub-diagram of $\mathcal{V}_d(\mathcal{S})$.

Corollary 5. *$\mathcal{V}_{d_G}(S)$ can be extracted from $\mathcal{V}_d(\mathcal{S})$ in time linear in the size of $\mathcal{V}_d(\mathcal{S})$.*

3.2 Computing Effective Needles

In this subsection, we present an algorithm for computing the Voronoi diagram $\mathcal{V}_{d_G}(S)$ for a given set S of sites under d_G. The algorithm consists of three phases; it first computes the set $\mathcal{S}$ of needles from G and S, secondly, the Voronoi diagram $\mathcal{V}_d(\mathcal{S})$ for $\mathcal{S}$ under d is constructed, and the Voronoi diagram $\mathcal{V}_{d_G}(S)$ for S under d_G is finally obtained from $\mathcal{V}_d(\mathcal{S})$.

The second phase, computing $\mathcal{V}_d(\mathcal{S})$, would be solved by several technics and general approaches to compute Voronoi diagrams, such as the abstract Voronoi diagram [10]. Also, the third phase can be done by Corollary 5. We therefore focus on the first phase, computing $\mathcal{S}$—in fact, its finite subset—from G and S.

Recall that $\mathcal{S}$ is defined as all the needles recursively produced from given sites and contains infinitely many useless needles, that is, needles that do not constitute the Voronoi diagram $\mathcal{V}_d(\mathcal{S})$. We call a needle $\mathbf{p} \in \mathcal{S}$ *effective with respect to* $\mathcal{S}$ if the Voronoi region of $\mathbf{p}$ in $\mathcal{V}_d(\mathcal{S})$ is not empty. Let $\mathcal{S}^* \subseteq \mathcal{S}$ be the largest set of effective needles with respect to $\mathcal{S}$, i.e., $\mathcal{V}_d(\mathcal{S}^*) = \mathcal{V}_d(\mathcal{S})$. Our algorithm computes $\mathcal{S}^*$ from G and S.

The algorithm works with handling events, which are defined by a certain situation at a time. Here, at each time t, we implicitly maintain the (backward) d_G-balls of the given sites, where the backward d_G-ball of a site p is defined as the set $B^-_{d_G}(p,t) = \{x | d_G(x,p) < t\}$, and d_G-balls expand as time t increases. Here, we have only one kind of events, called *birth events* which occur when a d_G-ball touches any footpoint on a road during their expansions; at that time a new needle will be produced in the algorithm. We can determine a birth event associated with a footpoint of a needle on a road. In order to handle events, we need two data structures: Let $\mathcal{Q}$ be an event queue implemented as a priority queue such that the priority of an event e is its occurring time and $\mathcal{Q}$ supports inserting, deleting, and extracting-minimum in logarithmic time with linear space. And, let $\mathcal{T}_1, \mathcal{T}_2, \cdots, \mathcal{T}_m$ be balanced binary search trees, each associated with e_i, where the road set E is given as $\{e_1, e_2, \cdots, e_m\}$. Each $\mathcal{T}_i$ stores needles on e_i in order and the precedence for a needle $\mathbf{p}$ follows from that of $p_1(\mathbf{p})$ with respect to $\prec_{e_i}$. $\mathcal{T}_i$ supports inserting and deleting of a needle in logarithmic time, and also a linear scan for needles currently in $\mathcal{T}_i$ in linear time and space.

Now, we are ready to describe the algorithm ComputeEffectiveNeedles. First, the algorithm computes $\sigma_G(S)$ and the associating birth events, and insert events into $\mathcal{Q}$. Then, while the event queue $\mathcal{Q}$ is not empty, repeat the following procedure: (1)Extract the next upcoming event b, say that b is a birth event on a road e_i associated with a needle $\mathbf{p}$. (2)Test the effectiveness of $\mathbf{p}$ and, (3)if the test has passed, compute birth events associated with $\sigma_G(\mathbf{p})$, and insert the events into $\mathcal{Q}$.

ComputeEffectiveNeedles returns exactly $\mathcal{S}^*$ by the effectiveness test in (2). This test can be done by checking if the associating footpoint of the current event has been already dominated by d_G-balls of other sites. Thus, if the test is passed, we decide that the new needle should be effective and insert it into $\mathcal{T}_i$. The following lemma shows that the effectiveness test is necessary and sufficient to compute $\mathcal{S}^*$.

Lemma 6. *For every birth event and its associating needle* $\mathbf{p}$*,* $\mathbf{p}$ *is effective with respect to* $\mathcal{S}$ *if and only if it passes the effectiveness test of Algorithm* COMPUTE-EFFECTIVENEEDLES.

The algorithm always ends, which can be shown by the finiteness of $\mathcal{S}^*$. The effectiveness test always fails if all the roads are covered with the d_G-balls after some large time T and there exists such T since distances between any two points in the space are always defined. Hence, $\mathcal{S}^*$ contains a finite number of needles. Moreover, the number of events handled while running the algorithm is also bounded by the following lemma.

Lemma 7. $\mathcal{S}^*$ *is finite and the number of handled events is* $O(s \cdot |\mathcal{S}^*|)$*, where* s *is the maximum cardinality of* $\sigma_G(\mathbf{p})$ *for any needle* $\mathbf{p}$*.*

We end this section with the following conclusion.

Theorem 8. *Given a transportable distance* d*, a transportation network* G *under* d*, and a set* S *of sites,* $\mathcal{S}^*$ *can be computed in* $O(\epsilon(\log \epsilon + T_{ef}))$ *time, where* ϵ *is the number of events handled while running the algorithm, the same as* $O(s \cdot |\mathcal{S}^*|)$*, and* T_{ef} *denotes time taken to test the effectiveness.*

In most natural cases, such as the Euclidean metric and convex distances, sufficiently $T_{ef} = O(\log \epsilon)$ so that the total running time becomes $O(\epsilon \log \epsilon)$.

4 Transportation Networks Under Convex Distances

In this section, we deal with convex distances as a special case of transportable distances. We thus investigate geometric and algorithmic properties of the induced distance by a convex distance and a transportation network, and construct algorithms that compute the Voronoi diagram for given sites under the induced distance.

In order to devise such an algorithm, we apply the abstract scheme described in the previous section; a bundle of properties have to be shown: how to compute needles produced from a needle, how to check the effectiveness of needles, how many needles and events to handle, how to compute the Voronoi diagram for needles, and some technical lemmas to reduce the complexity.

A *convex distance* is defined by a compact and convex body C containing the origin, or the *center*, and is measured as the factor that C centered at the source should be expanded or contracted for its boundary to touch the destination. Note that a convex distance is symmetric, i.e. being a metric, if and only if C is symmetric at its center.

We consider the convex C as a black box which supports some kinds of elementary operations. These are finding the Euclidean distance from the center to the boundary in a given direction, finding two lines which meet at a given point and are tangent to C, finding the footpoint for a needle and a road, and computing the bisecting curve between two sites under the convex distance based on C. Here, we assume that these operations consume reasonable time bounds.

Throughout this section, for a convex body C, we denote by $C+p$ a translation of C by a vector p and by λC an expansion or a contraction of C by a factor λ. And, we may denote by C' the reflected body of C at its center and by ∂C the boundary of C.

4.1 Roads and Needles Under a Convex Distance

For the Euclidean distance, to take a shortest path with a road e, one should enter or exit e with angle $\pi/2 \pm \alpha$, where $\sin \alpha = 1/v(e)$ [2]. For a convex distance d based on C, there also exist such entering or exiting angles for a road that they lead to a shortest path. We can obtain these angles by simple operations on C. First, we pick a point x on e not $p_2(e)$ and consider $(d(x, p_2(e))/v(e))C' + p_2(e)$, say D. We then compute two lines which meet at x and are tangent to D. Each of the two lines is above or below e. We let α_e^+ and α_e^- be two inward directions perpendicular to the two lines. And, let us consider two directions from $p_2(e)$ to two meeting points between the two lines and D. We then let β_e^+ and β_e^- be reflections of the two directions at the center of D. The symbols + and − mean "above" and "below" with respect to e, respectively. See Figure 1(a).

Lemma 9. *Given a road e under a convex distance d, any shortest path from p to q passing through e has the following properties:*

- *If p is above e, either the access direction is α_e^+ or the access point is $p_1(e)$.*
- *If p is below e, either the access direction is α_e^- or the access point is $p_1(e)$.*
- *If q is above e, either the leaving direction is β_e^+ or the exiting point is $p_2(e)$.*
- *If q is below e, either the leaving direction is β_e^- or the exiting point is $p_2(e)$.*

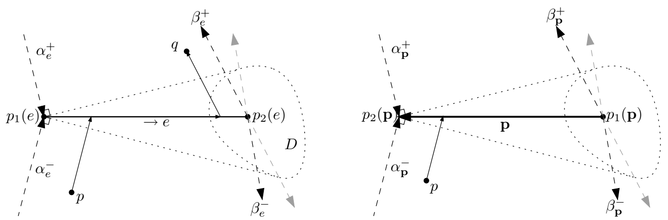

(a) A road e and a shortest path from p to q

(b) A needle $\mathbf{p}$ and a shortest path from p to $\mathbf{p}$

Fig. 1. Directions defined for a road and for a needle

By the observation of Lemma 9, we can easily find a shortest path with one road. For instance, Figure 1(a) shows a shortest path from p to q using road e.

We can define similar terms for a needle as we did for a road. The (backward) d-ball $B_d^-(\mathbf{p}, t)$ of $\mathbf{p}$ can be computed as follows: If $0 < t \leq t_1(\mathbf{p})$, $B_d^-(\mathbf{p}, t)$ is empty. If $t_1(\mathbf{p}) < t \leq t_2(\mathbf{p})$, $B_d^-(\mathbf{p}, t)$ is the convex hull of $(t - t_1(\mathbf{p}))C' + p_1(\mathbf{p})$ and the point x on $\mathbf{p}$ such that $d(x, p_1(\mathbf{p}))/v(\mathbf{t}) = w_{\mathbf{p}}(x) - t_1(\mathbf{p})$. And, if $t > t_2(\mathbf{p})$, $B_d^-(\mathbf{p}, t)$ is the convex hull of $(t - t_1(\mathbf{p}))C' + p_1(\mathbf{p})$ and $(t - t_2(\mathbf{p}))C' + p_2(\mathbf{p})$. Thus, when $t > t_1(\mathbf{p})$, $\partial B_d^-(\mathbf{p}, t)$ contains two line segments tangent to two scaled C' and the slopes of these line segments do not change even if t changes. Hence, the meeting points between the line segments and the convex bodies scaled from C' make 4 rays, and they have two directions $-\beta_{\mathbf{p}}^+$ and $-\beta_{\mathbf{p}}^-$. Then, we can define $\alpha_{\mathbf{p}}^+$, $\alpha_{\mathbf{p}}^-$, $\beta_{\mathbf{p}}^+$, and $\beta_{\mathbf{p}}^-$, equivalently as for a road, see Figure 1(b). Note that if $\mathbf{p} \in \sigma_e(\mathbf{q})$ for any needle $\mathbf{q}$, then $\alpha_{\mathbf{p}}^+ = \alpha_e^+$, $\alpha_{\mathbf{p}}^- = \alpha_e^-$, $\beta_{\mathbf{p}}^+ = \beta_e^+$, and $\beta_{\mathbf{p}}^- = \beta_e^-$, by definition of $\sigma_e(\mathbf{q})$.

4.2 Computing $\mathcal{S}^*$

In this subsection, we discuss how to compute $\mathcal{S}^*$ for given sites S from the algorithm ComputeEffectiveNeedles described in the previous section.

The footpoints. Under a convex distance, any needle has at most one connected component of footpoints on a road. Actually, Lemma 9 tells us how to compute the least footpoint of a point p on a road e; it can be obtained as either the intersecting point of the road e and the ray with direction $-\beta_e^+$ or $-\beta_e^-$ from p, or just $p_2(e)$.

Lemma 10. *Let d be a convex distance based on a convex C and G be a transportation network under d. For a road e in G and a needle $\mathbf{p}$, the number of connected components of footpoints of $\mathbf{p}$ on e is at most one and the least footpoint is either the least footpoint of $p_1(\mathbf{p})$ or $p_2(\mathbf{p})$, or just $p_2(e)$.*

The effectiveness test. At every time a birth event occurs, the effectiveness test is done by testing if the point where the event occurs is dominated by other already produced effective needles at that time. Under a convex distance, d-balls of any needle are convex so that we can test the effectiveness in logarithmic time by maintaining the tree structures $\mathcal{T}_i$ and doing a couple of operations on them.

The number of needles and events. By convexity of convex distances, we can show a couple of lemmas that prove the number of needles and events we handle.

Lemma 11. *Let $\mathbf{p}$ be a needle produced on a road e from a needle $\mathbf{q}$. For another road $e' \in E$, if $\mathbf{p}$ does not dominate any node of e or e', no needles in $\sigma_{e'}(\mathbf{p})$ are effective with respect to $\mathcal{S}$.*

Lemma 12. *$\mathcal{S}^*$ contains at most $O(m(n+m))$ needles, where n is the number of sites in S and m is the number of roads in G. Further, the number of handled events while the algorithm* ComputeEffectiveNeedles *running is $O(m^2(n+m))$.*

Remarks. Consequently, we have an algorithm to compute $\mathcal{S}^*$ in $O(m^2(n+m)(\log(n+m) + T_{op}(C)))$ time with $O(m^2(n+m))$ space by Theorem 8, where $T_{op}(C)$ is time taken during a simple operation on C. This complexity, however, can be reduced by small modifications on the algorithm. For the Euclidean case, Bae and Chwa [4] additionally maintains *node events*, which occur when $B_{d_G}^-(p, t)$

first touches a node, to reduce the number of events to $O(m(n+m))$. Here, we also can apply the approach by Lemma 11.

The authors also introduced primitive paths; a path is called *primitive* if the path contains no nodes in its interior and passes through at most one road. We can show the lemma of primitive paths in our setting, following from Lemma 11.

Lemma 13. *Given a transportation network G under a convex distance d, for two points p and q, there exists a d_G-straight path P from p to q such that P is a sequence of shortest primitive paths whose endpoints are p, q, or nodes in V.*

By Lemma 13 together with node events, we can improve our algorithm to be more efficient. Indeed, we can compute a shortest path for two given points by constructing an edge-weighted complete graph such that vertices are nodes and two given points, and edges are shortest primitive paths among vertices. Furthermore, we can use the graph to avoid useless computations during the algorithm. For more details, we refer to [4].

Lemma 14. *One can compute $\mathcal{S}^*$ in $O(m(n+m)(m+\log n+T_{op}(C)))$ time with $O(m(n+m))$ space.*

4.3 Voronoi Diagrams for Needles

In general, bisectors between two needles under a convex distance can be parted into two connected components. However, it will be shown that $\mathcal{S}^*$ can be replaced by such nice needles, so called *non-piercing*, that the Voronoi diagram for them is an abstract Voronoi diagram which can be computed in the optimal time and space.

Computing the Voronoi diagrams for non-piercing needles. The *abstract Voronoi diagram* is a unifying approach to define and compute general Voronoi diagrams, introduced by Klein [10]. In this model, we deal with not a distance but bisecting curves $J(p,q)$ defined in an abstract fashion between two sites p and q. A system $(S, \{J(p,q) | p,q \in S, p \neq q\})$ of bisecting curves for S is called *admissible* if the following conditions are fulfilled: (1) $J(p,q)$ is homeomorphic to a line or empty, (2) $R(p,q) \cap R(q,r) \subset R(p,r)$, (3) for any subset $S' \subseteq S$ and $p \in S'$, $R(p,S')$ is path-connected if it is nonempty, and (4) the intersection of any two bisectors consists of finitely many components, where $R(p,q) = \{x \in \mathbb{R}^2 | d(x,p) < d(x,q)\}$, $R(p,S) = \bigcap_{q \in S, p \neq q} R(p,q)$, and $J(p,q)$ is the bisector between p and q.

In fact, the first three conditions are enough to handle abstract Voronoi diagrams theoretically but the fourth one is necessary in a technical sense [11]. Though all convex distances satisfy the first three ones, there exist convex distances violating the fourth one. We guarantee the fourth condition by postulating that ∂C is semialgebraic [5]. We also note that two-dimensional bisectors can be avoided by a total order on given sites [13].

Lemma 15. *Let S be a set of pairwise non-piercing needles under a convex distance d based on C whose boundary is semialgebraic. Then, the system $(S, \{J(\mathbf{p},\mathbf{q}) | \mathbf{p},\mathbf{q} \in S, \mathbf{p} \neq \mathbf{q}\})$ of bisecting curves for S is admissible.*

There are several optimal algorithms computing abstract Voronoi diagrams [10, 15, 12, 6]. These algorithms assume that the bisector between two sites can be computed in constant time, and construct the Voronoi diagram in $O(n \log n)$ time with $O(n)$ space when n sites are given. In our setting, we assume that a representation of the bisector between two needles has the complexity $S_b(C)$, and can be computed in $T_b(C)$ time.

Corollary 16. *Let S be a set of n pairwise non-piercing needles under a convex distance d based on a convex C. Then, the Voronoi diagram $\mathcal{V}_d(S)$ for S can be computed in $O(T_b(C) \cdot n \log n)$ time and $O(S_b(C) \cdot n)$ space.*

Making $\mathcal{S}^*$ pairwise non-piercing. We can make $\mathcal{S}^*$ pairwise non-piercing by the procedure introduced in the proof of Lemma 2 in [4]. The only difference arises when we consider non-piercing needles as input sites; the produced needles on the roads may pierce the original needles. This problem can be solved by cutting a pierced original needle into two non-pierced needles. Since these piercing cases can occur only between original needles and produced needles dominating a node, we can check all the cases in $O(T_{op}(C) \cdot mn)$ time and the asymptotic number of needles does not increase. We denote by $\mathcal{S}^*_{np}$ the resulting set of pairwise non-piercing needles for a given set S of pairwise non-piercing needles. Note that $\mathcal{V}_d(\mathcal{S}^*_{np})$ is a refined diagram of $\mathcal{V}_d(\mathcal{S}^*)$.

4.4 Putting It All Together

From the previous discussions, we finally conclude the following theorem.

Theorem 17. *Let d be a convex distance based on a convex C whose boundary is semialgebraic, G be a transportation network with m roads under d, and S be a set of n sites. Then, the Voronoi diagram $\mathcal{V}_{d_G}(S)$ under d_G can be computed in $O(m(n+m)(m+T_b(C)\log(n+m)+T_{op}(C)))$ time with $O(S_b(C)m(n+m))$ space, where $T_{op}(C)$, $T_b(C)$, and $S_b(C)$ are defined as before.*

If C is a k-gon, we can see that $T_{op}(C) = O(\log k)$, $T_b(C) = S_b(C) = O(k)$ [14].

Corollary 18. *Let d be a convex distance based on a convex k-gon C, G be a transportation network with m roads under d, and S be a set of n sites. Then, the Voronoi diagram $\mathcal{V}_{d_G}(S)$ under d_G can be computed in $O(m(n+m)(m+k\log(n+m)))$ time with $O(km(n+m))$ space.*

References

1. M. Abellanas, F. Hurtado, C. Icking, R. Klein, E. Langetepe, L. Ma, B. Palop, and V. Sacristán. Proximity problems for time metrics induced by the l_1 metric and isothetic networks. *IX Encuetros en Geometria Computacional*, 2001.
2. M. Abellanas, F. Hurtado, C. Icking, R. Klein, E. Langetepe, L. Ma, B. Palop, and V. Sacristán. Voronoi diagram for services neighboring a highway. *Information Processing Letters*, 86:283–288, 2003.

3. O. Aichholzer, F. Aurenhammer, and B. Palop. Quickest paths, straight skeletons, and the city voronoi diagram. In *Proceedings of the 8th SoCG*, pages 151–159, 2002.
4. S. W. Bae and K.-Y. Chwa. Voronoi diagrams with a transportation network on the euclidean plane. Technical report, Korea Advanced Institute of Science and Technology, 2005. A preliminary version appeared in proceedings of ISAAC 2004.
5. A. G. Corbalan, M. Mazon, and T. Recio. Geometry of bisectors for strictly convex distances. *International Journal of Computertational Gemoetry and Applications*, 6(1):45–58, 1996.
6. F. Dehne and R. Klein. "the big sweep": On the power of the wavefront approach to Voronoi diagrams. *Algorithmica*, 17:19–32, 1997.
7. L. Gewali, A. Meng, Joseph S. B. Mitchell, and S. Ntafos. Path planning in $0/1/\infty$ weighted regions with applications. *ORSA J. Comput.*, 2(3):253–272, 1990.
8. R. Görke and A. Wolff. Computing the city voronoi diagram faster. In *Proc. 21st Euro. Workshop on Comput. Geom.*, pages 155–158, 2005.
9. J. C. Kelly. Bitopological spaces. *Proc. London Math. Soc.*, 13(3):71–89, 1963.
10. R. Klein. *Concrete and Abstract Voronoi Diagrams*, volume 400 of *LNCS*. Springer-Verlag, Berlin, Germany, 1989.
11. R. Klein. Private communication, 2005.
12. R. Klein, K. Mehlhorn, and S. Meiser. Randomized incremental construction of abstract voronoi diagrams. *Computational Geometry: Theory and Applications*, 3:157–184, 1993.
13. R. Klein and D. Wood. Voronoi diagrams based on general metrics in the plane. In *Proc. 5th STACS*, volume 294 of *LNCS*, pages 281–291. Springer-Verlag, 1988.
14. L. Ma. *Bisectors and Voronoi Diagams for Convex Distance Functions*. PhD thesis, Fern Unversität Hagen, 2000.
15. K. Mehlhorn, S. Meiser, and C. O'Dunlaing. On the construction of abstract Voronoi diagrams. *Discrete Comput. Geom.*, 6:211–224, 1991.
16. A. C. G. Mennucci. On asymmetric distances. *Preprint*, 2004. available at `http://cvgmt.sns.it/cgi/get.cgi/papers/and04/`.
17. J. S. B. Mitchell. Shortest paths among obstacles, zero-cost regions, and "roads". Technical Report 764, School Oper. Res. Indust. Engrg., Cornell Univ., Ithaca, NY, 1987.
18. B. Palop. *Algorithmic problems on proximity and location under metric constraints*. PhD thesis, U. Politécnica de Catalunya, 2003.
19. W. Rinow. *Die innere Geometrie der Metrischen Räume*, volume 105 of *Grundlehren der Mathematischen Wissenschaften in Einzeldarstellungen*. Springer-Verlag, Berlin, 1961.
20. N. C. Rowe. Roads, rivers, and obstacles: optimal two-dimensional path planning around linear features for a mobile agent. *Internat. J. Robot. Res.*, 9:67–73, 1990.