

Quickest Paths, Straight Skeletons, and the City Voronoi Diagram*

Oswin Aichholzer,¹ Franz Aurenhammer,¹ and Belén Palop²

¹Institute for Theoretical Computer Science, Graz University of Technology,
Inffeldgasse 16b, A-8010 Graz, Austria
{oaich, auren}@igi.tu-graz.ac.at

²Departamento de Ciencias de la Computación, Universidad Alcalá de Henares,
Madrid, Spain
b.palop@uah.es

Abstract. The city Voronoi diagram is induced by quickest paths in the L_1 plane, made faster by an isothetic transportation network. We investigate the rich geometric and algorithmic properties of city Voronoi diagrams, and report on their use in processing quickest-path queries. In doing so, we revisit the fact that not every Voronoi-type diagram has interpretations in both the distance model and the wavefront model. Especially, straight skeletons are a relevant example where an interpretation in the former model is lacking. We clarify the relationship between these models, and further draw a connection to the bisector-defined abstract Voronoi diagram model, with the particular goal of computing the city Voronoi diagram efficiently.

1. Introduction and Results

Imagine a large modern-style city, with streets arranged in north–south and east–west directions. The city is equipped with a public transportation network such as a subway or a bus system. Time is precious and people intend to follow the quickest path from their homes to their desired destinations, using the network whenever appropriate. For some people several facilities of the same kind are equally attractive (think of post offices or hospitals), and their wish is to find out which facility is reachable first. There may also be commercial interest (from real estate agents or from a tourist office) to make the area

* The research by O. Aichholzer was supported by APART (Austrian Programme for Advanced Research and Technology) of the Austrian Academy of Sciences. B. Palop's research was supported by Proyecto MEC-DGES-SEUID-PB98-0933. Part of the work was carried out when this author visited the Institute for Theoretical Computer Science, Graz University of Technology.

which can be reached in, say 1 hour, from a given location in the city (the apartment for sale or the recommended hotel) visible. Neuralgic places lying within this 1-hour zone, like the main square, train stations, shopping centers, or tourist attraction sites, should be displayed to the customer.

This paper offers a geometric and algorithmic study of questions of this kind. Their geometric complexity (and appeal) becomes apparent when noticing that quickest paths are inherently complex: once having accessed the transportation network, it may be too slow simply to follow it to an exit point close to the desired destination; taking intermediate shortcuts by foot may be advantageous at several places. Still, quickest paths induce a metric in the plane. We call the resulting Voronoi diagram for a given city transportation network C and a given set S of point sites in the plane, the *city Voronoi diagram*, $V_C(S)$, for C and S . For each individual site, the city Voronoi diagram makes explicit the region of all points reached first when starting from this site.

One of our results is that, under rather general assumptions, the size of the city Voronoi diagram is as small as can be hoped for: linear in the size of S and C . We further show how $V_C(S)$ can be augmented (without asymptotic increase of storage) to return the quickest paths from arbitrary query points to S in the optimal time. Finally, a construction algorithm is developed whose running time is optimal up to an additive term which is independent from the number of sites.

Computing a city Voronoi diagram is a highly non-standard task. Cyclic site bisectors may occur, and a single bisector may be as complex as the whole transportation network. To achieve a satisfactory performance we have to touch upon powerful concepts like *straight skeletons* and *abstract Voronoi diagrams*, the former not being Voronoi diagram interpretable (even in the abstract sense) unless certain conditions are fulfilled. Characterizing these conditions in a general setting is considered a result of separate interest to this paper.

1.1. Model Concretization

The kind of transportation network we consider is a planar straight line graph C with isothetic (horizontal or vertical) edges. No additional requirements are imposed on the network. In particular, C may be composed of many connected components and may contain cycles. We let c denote the number of *nodes* of C . The $\Theta(c)$ edges of C are termed *segments* to distinguish them from edges arising from other structures later.

By assumption, we are free to enter C at any point, be it at a node or on a segment. (This is not unrealistic for a bus system with densely arranged stops, and exactly meets the situation for shared taxis which regularly drive on predetermined routes and will stop for every customer.) In fact, postulating *fixed* entry points eases the problem considerably; see Section 2.2. Once having accessed C we travel at an arbitrary but fixed speed v in one of the (at most four) available directions. Movement off the network takes place with unit speed with respect to the L_1 metric. (Again, this is realistic when walking in a modern city.) Let $v > 1$ to ensure that C is of use at all. The problems in question simplify to standard proximity problems in the L_1 metric, otherwise.

In this form, the model was first introduced by Abellanas et al. [1] who derived basic properties of the quickest-path metric; we revisit some of them in Section 4. They also

made observations on networks C which are both horizontally and vertically monotone. Moreover, for the case where C is a single straight line, an $O(n \log n)$ time construction algorithm is given for the city Voronoi diagram $V_C(S)$ of a set S containing n sites. We are not aware of other work on this model. More general related models have been studied earlier, the most similar in [12]; see Section 2.1.

The present paper solves the problem of constructing $V_C(S)$ in the full generality of the model. By interpreting $V_C(S)$ as the straight skeleton of polygonal figures, we show the size of $V_C(S)$ to be $O(n + c)$. In fact, $V_C(S)$ is only a subgraph of this straight skeleton; see Fig. 10 for an illustration. Skeleton edges further refine $V_C(S)$ such that it allows the retrieval of quickest paths from query points to sites (multiple source quickest-path queries) in time $O(\log(n + c) + k)$, where k is the complexity of the path. By exploiting the special structure of such straight skeletons, we obtain an $O(n \log n + c^2 \log c)$ algorithm for their construction. The key observation is that the machinery of abstract Voronoi diagrams—which fails for straight skeletons in general—applies in this case, after careful adaption of the skeleton figures.

The situation changes if the model is even slightly altered. Non-isothetic networks, individual traveling speeds on network segments, or the Euclidean distance instead of L_1 as the underlying metric, are scenarios all leading to city Voronoi diagrams of size $\Theta(n \cdot c)$, as will be shown by simple examples. Some results for a single-line transportation network under the Euclidean metric are discussed in [13]. In view of the prohibitive (i.e., quadratic) diagram size, we refrain from further considering such models. In certain situations, our model may serve as a (rough) linear-size approximation of the seemingly most interesting alternative, the Euclidean city Voronoi diagram.

2. Related Concepts

2.1. *Weighted Region Problems*

The *weighted region problem* is based on a partition of the plane into polygonal regions, each associated with an individual weight that expresses the internal traveling speed (or the reciprocal, the cost per unit distance). Regions may be entered and left at any point. For two given points p and q , the quickest (Euclidean) path between p and q is sought. Mitchell and Papadimitriou [19] originally posed the problem and provided an interesting solution based on Snell's law of refraction. The query time is $O(c^7 \cdot \ell)$, where c is the complexity of the polygonal partition and ℓ denotes the precision of the problem instance. Allowing only weights 0 and 1 leads to the shortest obstacle-avoiding path problem (geodesic path problem) for which much more efficient solutions are known; see, e.g., [18].

For a given transportation network C , quickest-path queries can be viewed as a one-dimensional instance of the weighted region problem. Choose a weight v_e for each network segment e , and weight 1 on all connected parts of the complement of C in the plane. This variant (and others) have been studied by Gewali et al. [12], who construct the quickest path between two given points p and q in time $O(c^2)$. Our model constitutes the special case where C is isothetic and of constant speed v , and where the underlying metric is L_1 . Although we need $O(c^2 \log c)$ time for finding a quickest path between p

and q (by constructing $V_C(\{p\})$ and performing point location for q), we can retrieve multiple source quickest paths of complexity k in optimal time $O(k + \log(c + n))$, once the city Voronoi diagram $V_C(S)$ for a set S of n point sites (sources) has been constructed. We mention that Voronoi diagram (and wavefront expansion) techniques not unlike the methods in this paper have been used, e.g., in [17], for solving multiple source shortest path problems in polygonal domains.

2.2. Airlift Voronoi Diagrams

Another related concept is the *airlift distance* considered in [6]. Here an arbitrary graph A on c points (airports) in the plane is given, with positive edge weights (flight durations) that need not fulfill the triangle inequality. The graph may be entered and exited only at the airports. Given, in addition, a set S of n sites, we are interested in the Voronoi diagram that results from the distance function induced by quickest paths in the plane using A .

Though showing disconnected regions, compared with $V_C(S)$ this is a considerably simpler diagram. In fact, denying access to the transportation network C except at designated points yields an instance of the airlift Voronoi diagram, where the weight of a segment of C is its length divided by the network speed v . Note that such weights do obey the triangle inequality.

We argue below that the airlift Voronoi diagram (in its general form) can be constructed for any fixed L_p metric, $p \geq 1$, in time $O(n \log n + c^2)$ by a reduction to an L_p Voronoi diagram for $n + c$ circles in this metric. This has been also observed by Berman [7] who posed the question of whether a worst-case runtime sensitive to the number of *edges* of A can be achieved. L_p circles induce a so-called *nice metric*, so their Voronoi diagram is computable in time $O((n + c) \log(n + c))$; see [14]. To specify the circles, the sites in S get assigned a radius of zero. For each airport a its radius is

$$d(a) = -\min\{d(s, a) \mid s \in S\},$$

for $d(s, a)$ being the minimum time needed to reach a from site s , be it via A or not. (All radii can be made positive by adding a fixed constant; this does not change the resulting Voronoi diagram.) The airport radii can be calculated from the complete graph A' obtained from A by adding edges whose weights are their L_p lengths. Thereby, we initially determine which airports are reached first—without using A' —by each particular site, by locating all airports in the (standard) L_p Voronoi diagram of S . Starting with these first access times to A' , we conduct a simultaneous breadth-first search in A' to see how and when the sites conquer the airports.

Note that the reduction, rather than the diagram construction for the L_p circles, yields the dominant time factor.

3. Straight Skeletons

Straight skeletons are the key concept for a proper geometric and algorithmic understanding of the city Voronoi diagram. Originally intended as a linearization of the medial axis,

the concept has been introduced in [2] and has been generalized to planar straight-line graphs in [3] and [4]. Various applications exist, some being seemingly unrelated; see [4], [8], [9], [11], [20], and references therein. The present paper adds to this list.

Let us revisit the basic properties of straight skeletons in some detail for later use. Consider a simple polygon f in the plane. The polygon f is called a *figure* if each edge of f can be assigned an individual speed such that, when moving the boundary of f inwards at these speeds and in a self-parallel fashion, f contracts to a point without prior self-intersection. Viewing this process backwards, each figure f can be associated with a *place* x_0 and a *time* t_0 of birth, and with a parametrized set $W_t(f)$ of polygons, for $t \geq t_0$, called the *wavefront* sent out by f . We define $W_t(f) = \emptyset$ for $t < t_0$.

Let F be an arbitrary set of figures. The *common wavefront* $W_t(F)$ of F is defined as follows. For all t where the individual wavefronts $W_t(f)$, $f \in F$, are pairwise disjoint, we set $W_t(F) = \bigcup_{f \in F} W_t(f)$. For larger t , the movement of wavefronts ceases at all points in the plane where wavefronts come into contact or self-contact. That is, two polygons of $W_t(F)$ merge into one in the former case, and a polygon of $W_t(F)$ splits into two in the latter case (one polygon shrinking when t increases). Note that $W_t(F)$ consists of a single polygon for sufficiently large t .

$W_t(F)$ uniquely describes an interference pattern of the individual wavefronts $W_t(f)$ —which is called the *straight skeleton* $SK(F)$ of F —in the following way. The vertices of $W_t(F)$ trace out the *edges* of $SK(F)$, and the edges of $W_t(F)$ sweep out the *faces* of $SK(F)$. The edges of $SK(F)$ are portions of (weighted) polygon angle bisectors, and thus are straight line segments. The faces of $SK(F)$ are connected polygonal regions, which define a partition of the plane with vertex degree at least 3. As each skeleton face stems from some figure edge, the size of $SK(F)$ is $O(e)$, where e counts the total number of figure edges.

Despite these “Voronoi-like” properties, $SK(F)$ does not admit any distance-from-site definition, in general. This justifies the (less convenient) procedural definition given above. Due to the same unpleasant fact, all construction algorithms for $SK(F)$ known to date are mere simulations of the wavefront expansion. The most (time) efficient implementation, by Eppstein and Erickson [11], runs in worst-case time $O(e^{8/5+\varepsilon})$, for any fixed $\varepsilon > 0$. For k -oriented figures, their algorithm takes $O(k^4 e \log^2 e)$ time (and superlinear space) by a reduction to $O(k^4)$ orthogonal range query problems in three dimensions.

Interestingly, Voronoi diagram behavior is retained when F consists of so-called *non-piercing* figures. We prove in Section 7 that, for this case, $SK(F)$ is an *abstract Voronoi diagram* in the sense of Klein [14]. Beside these and other related general findings, the merit of Sections 5–7 is a reduction of the city Voronoi diagram to the straight skeleton of a set of non-piercing (and convex) figures, whose birth times, shapes, and edge speeds have been carefully chosen.

4. The City Voronoi Diagram

We continue with basic observations on quickest paths and the city Voronoi diagram; they partially found an earlier mention in [1].




Fig. 1. Quickest path for $v = 3$.

4.1. *Quickest Paths*

Let a network C as described in Section 1.1 be given. For two points x and y in the plane, let $Q_C(x, y)$ denote a *quickest path* between them, that is, a path minimizing the travel duration between x and y . Let $d_C(x, y)$ be the time for traversing $Q_C(x, y)$. Observe that d_C induces a metric (which we call the *city metric*) in the plane: d_C is non-negative, symmetric, and obeys the triangle inequality because, for each point $z \in Q_C(x, y)$, the concatenation of $Q_C(x, z)$ and $Q_C(z, y)$ gives $Q_C(x, y)$.

Figure 1 displays a quickest path (shown dashed) whose linear pieces are labeled with their durations of traversal. The simple-shaped network C (shown in bold) enforces two *shortcuts*, i.e., intermediate pieces leaving C , as well as a non-monotonic behavior of the path. It is obvious that $Q_C(x, y)$ is an isothetic polygonal line which, apart from special cases, is not unique. Note further that $Q_C(x, y)$ may contain $\Theta(c)$ pieces on C plus $\Theta(c)$ shortcuts.

The number of potential shortcuts for C may be $\Theta(c^2)$, in the sense that removal of any such shortcut disables all possible paths $Q_C(x, y)$ between certain points x and y in the plane. For instance, in Fig. 2, every quickest path that connects a network segment to the left with a network segment to the right and below, has to run via an individual shortcut. On the other hand, we face the fact that any non-trivial limitation on the number of accesses to C results in a loss of triangle inequality. These observations do not imply an $\Omega(c^2)$ lower bound on computing quickest paths or the city Voronoi diagram, however.

The complexity of the city metric d_C is also apparent from the shape of the induced *circles*:

$$K(x, r) = \{y \mid d_C(x, y) = r\}.$$




Fig. 2. Many shortcuts between terminals.




Fig. 3. City Voronoi diagram for two sites.

Not only may $K(x, r)$ break into $\Theta(c)$ pieces, its shape also varies with its radius r , as well as with the relative position of its center x to C ; see Fig. 5. Still, the “disk” interior to $K(x, r)$ is a connected set, being the union of quickest paths.

4.2. The Diagram

Let S be a set of n point sites in the plane. The *city Voronoi diagram* $V_C(S)$ contains, for each site $s \in S$, the region of all points given by

$$reg(s) = \{x \mid d_C(x, s) < d_C(x, t), \forall t \in S \setminus \{s\}\}.$$

Because d_C is a distance function, $V_C(S)$ induces a partition of the plane. See Fig. 10 for a detailed illustration. Figure 3 displays the city Voronoi diagram induced by two sites. The regions are separated by a closed line; the region of the lower left site is not simply connected. However, regions are path-connected, by the observation below.

Observation 1. *reg(s) is the union of all quickest paths from $x \in reg(s)$ to the site s, and equivalently, to the set S.*

Proof. We have to show that $x \in reg(s)$ implies $Q_C(x, s) \subset reg(s)$. Assume there exists a point y with $y \in Q_C(x, s)$ but $y \in reg(t)$ for $t \neq s$. Then $d_C(x, s) = d_C(x, y) + d_C(y, s) > d_C(x, y) + d_C(y, t)$. The latter term is at least $d_C(x, t)$ by the triangle inequality, which implies $x \notin reg(s)$ and gives a contradiction. \square

Observation 1 does *not* imply that $V_C(S)$ is of size $O(n)$. Vertices of degree 2 may occur, whose number cannot be bounded by applying the Eulerian theorem to a planar graph with n faces. A trivial upper bound is $O(n \cdot c)$, by Observation 2 below. The true complexity of $V_C(S)$ is $O(n + c)$; we postpone the proof to Section 5.




Fig. 4. Non-isothetic networks or different speeds cause many diagram edges.

We consider the *bisector* of two sites $s, t \in S$ in the city metric, which is the set

$$B_{st} = \{x \mid d_C(x, s) = d_C(x, t)\}.$$

B_{st} is connected by Observation 1. However, due to well-known L_1 -specific degenerations, B_{st} may fail to be one-dimensional. A canonical one-dimensional form of B_{st} (which we use throughout this paper) results from the straight skeleton interpretation of $V_C(S)$ given in Section 5. In this form, B_{st} is always a piecewise linear curve. The following property of B_{st} , exemplified in Fig. 3, is undesirable in view of the algorithmic construction of $V_C(S)$.

Observation 2. *B_{st} is a (possibly) cyclic polygonal curve with $\Theta(c)$ edges in the worst case.*

Note that in Fig. 3 each node and each segment of C , respectively, causes a vertex and an edge of B_{st} . Cyclic and non-constant-sized bisectors are known to be a main obstacle for efficiently applying divide-and-conquer and randomized incremental insertion.

We finally point out that slight alterations to our model lead to diagrams with $\Omega(n \cdot c)$ vertices and edges. For example, in both of the situations shown in Fig. 4, $n - 1$ bisectors each consisting of $\Omega(c)$ linear components arise.

5. Link to Straight Skeletons

5.1. The Reduction

This subsection draws a connection between city Voronoi diagrams and straight skeletons. Its implications are manifold; they fill the contents of the rest of the paper.

To aid the intuition of the reader, we recall a well-known interpretation of the standard L_1 Voronoi diagram in the wavefront model. Imagine each site in the given set S sends out a wavefront, in the shape of an expanding L_1 circle (called a *diamond*), at time 0 and with translational speed $1/\sqrt{2}$ for all its edges. Wavefront movement stops at every point where two diamonds come into contact. In other words, the L_1 Voronoi diagram is a straight skeleton with these diamonds as its defining figures; see Section 3 for a definition of figures, wavefronts, and the straight skeleton.

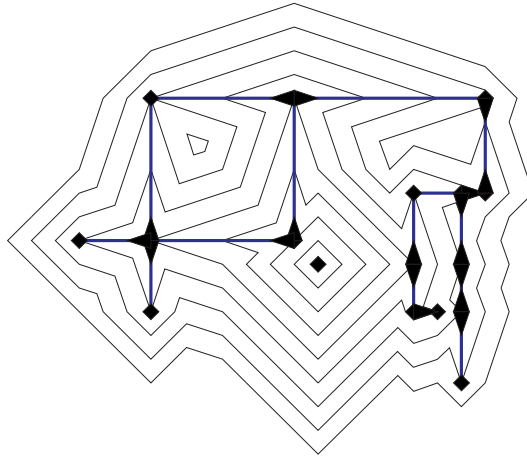


Fig. 5. Concentric d_C circles and figure generation. The small black polygons depict the figures in the set F_s .

How does the network C influence this wavefront model? The circles expanding around the sites change their shapes, in a manner determined by C which we want to explore next. Let $s \in S$ be a fixed site, and recall that $K(s, t)$ denotes the circle of radius t around s in the city metric. Our aim is to interpret $K(s, t)$ as the common wavefront propagated from a certain set F_s of figures.

See Fig. 5. For sufficiently small $t > 0$, $K(s, t)$ is a diamond with center s as before. Let F_s be a set of figures which initially contains this diamond, with place s and time 0 of birth. Now consider all points $x \in C$ where either a vertex of $K(s, t)$ runs into a network segment, or an edge of $K(s, t)$ slides into a network node. At times $t = d_C(s, x)$, and at no other point in time, the circle $K(s, t)$ changes its combinatorial shape. Accordingly, we add to the set F_s an appropriate figure with time t and place x of birth, constructed as below.

Observe that, at time $t + \varepsilon$, a diamond D with center x and diagonal length ε appears on $K(s, t)$, along with sharp-angled wedges tangent to D , whose peaks move at speed v in all possible directions on C that point to the exterior of $K(s, t)$. The figure we construct is simply the union of D and these wedges. Each wedge is an isosceles triangle with base d and height $\varepsilon \cdot v$, where d is a diagonal of D ; see Fig. 6.

We categorize the figures in F_s by their number j of wedges and call them j -needles. Note that $0 \leq j \leq 3$ holds, because the degree of a network node is at most 4, and at least one of the four isothetic directions points to the interior of $K(s, t)$. Only for a 0-, 1-, and 2-needle, may its diamond contribute edges to the figure. For such edges, the translational speed is $1/\sqrt{2}$ as before, whereas the translational speed of all other edges amounts to $v/\sqrt{v^2 + 1}$. Observe further that all figure vertices move in an isothetic manner during the expansion of their individual wavefronts.

Let $W_t(F_s)$ denote the common wavefront of all the figures in F_s . By construction, $W_t(F_s)$ enjoys the following property: $W_t(F_s)$ coincides with $K(s, t)$ at any time t . This property has an obvious but important consequence.




Fig. 6. Different types of needles.

Theorem 1. *Let $F = \bigcup_{s \in S} F_s$. The straight skeleton $SK(F)$ of F contains the city Voronoi diagram $V_C(S)$ as a subgraph.*

The set F in Theorem 1 contains *redundant* figures, i.e., figures which do not contribute to $SK(F)$. Their wavefronts solely and always consist of points swept over earlier (or at the same time) by wavefronts of other figures. It is clear that, for each site $s \in S$, all figures of F_s whose places of birth lie outside the region $reg(s)$ in $V_C(S)$ are redundant in F . (Even more figures are redundant; we will return to this issue in Section 5.3.) Edges of $SK(F)$ which are not edges of $V_C(S)$ stem from non-redundant figures born in the *same* region of $V_C(S)$. We alternatively call $SK(F)$ the *refined* city Voronoi diagram of S and C , or $\mathcal{V}_C(S)$ for short. The two-dimensional components of $\mathcal{V}_C(S)$ are its *regions*, which are the regions of $V_C(S)$ and are partitioned into *faces*, the faces of $SK(F)$. See Fig. 10 for an illustration.

5.2. Path Encoding

We show next that the refined city Voronoi diagram $\mathcal{V}_C(S)$ encodes, for every point q in the plane, the quickest path from q to S .

By construction of F , certain edges of $\mathcal{V}_C(S)$ are (portions of) network segments, some edges are shortcuts in the network C (or go directly from a site to the network), and others are of a third type. We orient all these edges in the direction they are traced out by the vertices of the common wavefront $W_t(F)$ of F . This yields a directed acyclic graph with the sites in S as sources. When selecting all edges of the first two types in this graph, a quickest-path tree $T(s)$ for each site s within its region $reg(s)$ is obtained. Note that $T(s)$ consists of isothetic edges.

Let s be a site with $q \in reg(s)$. Then $Q_C(q, s)$ is the quickest path from q to S . If q lies on C or in S , then $Q_C(q, s)$ consists entirely of edges of the quickest-path tree $T(s)$. Otherwise, q lies in the interior of a face g of $\mathcal{V}_C(S)$. The face g can be uniquely associated with its *defining feature*, $\beta(g)$. This is the site, network node, or (portion of) a network segment that contains the place of birth of the unique figure whose wavefront sweeps out g with one of its edges. In fact, g is just the union of all the quickest paths between points $q \in g$ and $x \in \beta(g)$. This implies that the first edge of the path $Q_C(q, s)$

lies in the interior of g , whereas the second edge of $Q_C(q, s)$ (if it exists) lies on the boundary of g and belongs to $T(s)$. We conclude:

Theorem 2. *Given an arbitrary query point q , the quickest path from q to S can be reported by point location in $\mathcal{V}_C(S)$, in optimal time $O(\log(n + c) + k)$ where k is the path complexity.*

5.3. The Linear Size

The diagram $\mathcal{V}_C(S) = SK(F)$ is a planar graph whose faces are connected and whose vertices have degree at least 3, as in every straight skeleton. The number of edges and vertices of $\mathcal{V}_C(S)$ therefore is linear in the number of its faces. This number, in turn, is bounded by the total number of edges of all figures which are non-redundant in the set F , because each such edge (possibly) sweeps out a single face of $\mathcal{V}_C(S)$.

How many figures are non-redundant? Clearly, each site $s \in S$ defines a non-redundant diamond with birth place s . Moreover, by arguments in Section 5.1, each node of C gives rise to a single non-redundant j -needle born at this node. This gives $n + c$ relevant figures so far. It remains to examine the 2-needles born in the interior of network segments. They are called *interior 2-needles*, in contrast to the 2-needles born at network nodes. Their number is $\Theta(n \cdot c)$ in the worst case: each of the n sites may cause c such figures, born on c parallel segments; see Fig. 5, right lower corner. However, the following fact comes to our rescue.

Lemma 1. *Let f be an interior 2-needle on network segment σ . Then the interior 2-needle f' born when the wavefront of f hits some segment parallel to σ is redundant in the set $\{f, f'\}$.*

Proof. Let t be the birth time of f' . For any time $u > t$, the wavefronts $W_u(f)$ and $W_u(f')$ are homothetic rhombs, see Fig. 7, with the common vertex x as the center of homothety. This vertex coincides, at time t , with the center of f' . □

An interior 2-needle never stems from a contact between a network segment and some *peak* (sharp-angled vertex) of a j -needle. So Lemma 1 implies that only the interior 2-needles that stem from contacts with vertices incident to some diamond edge are non-redundant in the set F . As such vertices disappear from the common wavefront of F after




Fig. 7. Redundant 2-needle.

their first contact with C , we have at most $4(n + c)$ figures of this kind. In summary, F contains at most $5(n + c)$ non-redundant figures. Each figure is of constant complexity, which gives a total of $O(n + c)$ figure edges, and with it, skeleton faces.

Theorem 3. *The refined city Voronoi diagram $\mathcal{V}_C(S)$ consists of $O(n + c)$ faces, edges, and vertices.*

We remark that redundant 2-needles as in Lemma 1 are the reason for portions of C that never need to be used. An unused portion ϱ is characterized by containment in a face g whose defining feature $\beta(g)$ is a (portion of a) network segment different from ϱ . See Fig. 10 where unused portions are shown as dotted bold lines.

6. Conquering the Network

We now turn to algorithmic issues. Our first step in the construction of $\mathcal{V}_C(S)$ is finding a set F^* of non-redundant figures whose straight skeleton gives $\mathcal{V}_C(S)$. When tolerating redundancy, this is a quite simple task, leading to a worst-case runtime of $\Omega(n \cdot c)$, however. We show below that $O(n \log n + c^2 \log c)$ time suffices for computing F^* . In view of the existence of $\Theta(c^2)$ potential shortcuts in the network C (see Section 4.1), it might be difficult to achieve a subquadratic complexity in c .

The diamonds for the sites in S are calculated in a trivial manner, so we concentrate on the figures caused by the network C ; see Section 5.1.

6.1. Figures at Nodes

We treat figures born at network nodes first. Consider the grid G obtained by putting a horizontal and a vertical line through each node of C . Each network segment is now the union of grid segments. For a site $s \in S$, let $G(s)$ denote the grid box that contains s . An *access point* of s to C is a point where C is crossed by one of the two isothetic lines through s . A *primary access point* is a first access point as seen from s .

Lemma 2. *Let s and x be a site and a network node, respectively. Then $Q_C(s, x)$ can be chosen to follow the grid G , except for the first edge e .*

Proof. Edge e is an isothetic edge which connects site s to some grid segment g , and $Q_C(s, x)$ has to pass through one of the nodes of g . Either $g \not\subseteq C$ or $g \subseteq C$. In the former case, g is chosen to be a segment of $G(s)$, and in the latter case g contains a primary access point of s to C . \square

The following algorithm determines when (and from where) each grid node is conquered by the sites. This will give us the birth times and shapes of all non-redundant figures (j -needles) born at the nodes of C . First, each grid segment is given a weight 1

or v depending on its containment in C . Then, for each site $s \in S$, the nodes of $G(s)$ are *initialized*, as are the nodes of the at most four grid segments on C that contain the primary access points of s to C in their interior. That is, for each such node x , the minimum time $d_C(x, S)$ (and the corresponding direction) is stored. This takes time $O(n \log c)$. To process the remaining nodes, we simply run the continuous Dijkstra method on the initialized grid G , in $O(c^2 \log c)$ additional time.

6.2. Figures on Segments

Now consider figures born on network segments; we called them interior 2-needles in Section 5.3. Redundancy occurs for two reasons, one made precise in Lemma 1. However, an interior 2-needle that stems from a contact p with a site diamond (that is, from a primary access point p whose time t has been implicitly calculated earlier, namely when initializing G) may still be redundant: there might be a site s whose circle $K(s, t)$ has swept over p already. To identify all non-redundant interior 2-needles, we proceed as below.

Consider each grid column (and, later, each grid row) separately. Sweep the column bottom-up with a horizontal line L , starting below S and C . Keep horizontally sorted lists A and B of the sites in the column, above and below L , respectively. Whenever L reaches a grid edge $g \subseteq C$ we do the following. For each *primary* access point p on g , and for g 's endpoints, detect all access points on g for sites in A and in B , respectively, which are *time covered* by p , that is, which can be reached earlier via p . Remove these sites from their lists (and file the new access points neighboring p on g , for later use as sentinels in Section 7). If a primary access point on g is time-covered itself by access points for sites in A or B , then remove the corresponding site from its list (and update sentinels accordingly). Note that the access times of all the sites to L can be updated quickly during the sweep, as these are just the vertical distances of the sites to L . Correctness is guaranteed because a site s outside the column cannot influence the situation, not even in a horizontal manner, as quickest paths from s enter the column via some (correctly initialized) grid node.

The total runtime for processing all columns and rows is $O((n+c) \log n)$. This finally gives a minimal set F^* with $SK(F^*) = \mathcal{V}_C(S)$, in time $O(n \log n + c^2 \log c)$.

7. Abstract Voronoi Settings

We are left with the problem of computing the straight skeleton $SK(F^*)$ of the set F^* . According to Section 5.3, the figures in F^* have a total of $O(n+c)$ edges. Moreover, these figures are of simple shape: they are k -oriented (for $k = 12$, as can be easily checked). Thus we could apply the algorithm of Eppstein and Erickson [11] which achieves a runtime of $O((n+c) \log^2(n+c))$, though with large constants and superlinear space; see Section 3. We show below that $SK(F^*)$ can be computed optimally, in time $O((n+c) \log(n+c))$ and space $O(n+c)$, by a simpler method which comes as a byproduct of the following more general results.

7.1. Relationship between Models

During our investigations so far, we have encountered several “Voronoi-like” diagrams, which are interpretable in one or the other model. In fact, the variety of approaches taken in the literature to model such diagrams is confusing at first glance, and certain peculiarities are hidden in the relationship between these models. We take the chance to clarify the situation, with the special goal of resolving the problem of computing $SK(F^*)$ efficiently.

The classical model is, of course, the *distance model*. Let $I_d = \{1, 2, \dots, m\}$ be a set of sites, each being augmented by some distance function $d(x, i)$. The resulting Voronoi diagram $V(I_d)$ divides the plane among the sites, according to the nearest neighbor rule. Distances may be defined in an abstract way, namely by bivariate functions $f_i(x)$ on the plane. Then $V(I_d)$ corresponds to the lower envelope of these functions; see [10].

Another frequently used concept is the *wavefront model* (or growth model, as it is sometimes called in the applied natural sciences). Each site i sends out some wavefront $W_t(i)$, for t measuring time, and the interference pattern of all these wavefronts constitutes the diagram; compare Section 3.

Wavefronts induce distances, and vice versa, in an obvious way. Simply let $d(x, i)$ be the time t when $W_t(i)$ reaches x . Conversely, put $W_t(i) = \{x \mid d(x, i) = t\}$, which is the circle of radius t around i . We denote with $\Pi(I_d)$ the interference pattern that arises—in this way—from the augmented set I_d of sites.

Still, the two models above are not equivalent. This becomes apparent when noticing that Voronoi diagrams with disconnected regions have no interpretation in the wavefront model. An example is the Voronoi diagram for point sites with multiplicative weights [5]. Being more important in the present context, straight skeletons are known to admit no distance-from-site definition, in general; see [3].

We now give characterizing conditions for the equivalence of both models. For fixed t , define the *interior* of a wavefront W_t as $\bigcup_{u < t} W_u$. Call two sites $i, j \in I_d$ *non-piercing* if, for all times t , the interior of their common wavefront $W_t(\{i, j\})$ coincides with the interior of the union of their individual wavefronts $W_t(i)$ and $W_t(j)$. For instance, L_p circles for $p \geq 1$ witness this property, even if radii are different—a scenario arising in the construction of airlift Voronoi diagrams, in Section 2.2. The property is lost for $p < 1$, and, indeed, disconnected Voronoi regions do arise.

The following general assertion can be made.

Theorem 4. *Let I_d be a set of pairwise non-piercing sites. Then the two diagrams $V(J)$ and $\Pi(J)$ coincide, for all $J \subseteq I_d$. The converse is true, too.*

Proof. Let $d(x, J)$ be the time t when the common wavefront $W_t(J)$ reaches the point x . If i and j are non-piercing sites, and only in this case, we have $d(x, \{i, j\}) = \min\{d(x, i), d(x, j)\}$, for all points x in the plane. So, assuming i and j to be piercing sites implies $V(\{i, j\}) \neq \Pi(\{i, j\})$. Conversely, if the sites in the set I_d are pairwise non-piercing, then each subset $J \subseteq I_d$ has this property, such that

$$d(x, J) = \min_{i \in J} d(x, i)$$

holds for all points x in the plane. However, this is equivalent to $V(J) = \Pi(J)$. \square




Fig. 8. Diamonds and needles violating requirements (1) and (2).

A third well-known model is the *abstract Voronoi diagram* introduced in [14]. In this setting, a bisector system $(B_{ij})_{1 \leq i < j \leq m}$ is given, which is called *admissible* if it fulfills the following requirements:

- (1) B_{ij} is homeomorphic to a line.
- (2) $B_{ik} \cap B_{jk} \subset B_{ij}$.
- (3) Regions are path-connected.

Sometimes, an additional technical condition is postulated,

- (4) $B_{ik} \cap B_{jk}$ has finitely many components,

which is mostly satisfied—in particular, by all the diagrams mentioned in the present paper. We therefore exclude (4) from further considerations.

Condition (1) is violated by the (unrefined) city Voronoi diagram $V_C(S)$. Straight skeletons violate both (1) and (2). Figure 8 reveals that even the simple-shaped figures in the set F^* yield this behavior. Note that the bisector of two figures i and j is defined as $B_{ij} = SK(\{i, j\})$ in this case. Violation of (2) is most serious, because “no-mans land” (shaded) may occur which belong to no site. This leads to anomalies like the *expansion* of a site’s region caused by the insertion of another site.

The following result tells us where to place the third model.

Theorem 5. *The bisectors $B_{ij} = \{x \mid d(x, i) = d(x, j)\}$, for an augmented set I_d of sites, constitute an admissible system if and only if $V(I_d) = \Pi(I_d)$.*

Proof. Suppose $V(I_d) = \Pi(I_d)$ first. To prove that condition (2) holds in any distance model, observe that $x \in B_{ij} \cap B_{ik}$ implies $d(x, i) = d(x, j) = d(x, k)$, hence $x \in B_{jk}$. Similarly, condition (3) holds in the wavefront model, because no part of a wavefront lost in the interference process can reappear later.

Condition (1) remains to be proven. From (3) we know that either half-plane defined by B_{ij} is path-connected. This implies that B_{ij} is a connected curve. We show that B_{ij} is acyclic. Our assumption $V(I_d) = \Pi(I_d)$ implies that i and j are non-piercing sites, by Theorem 4. Hence, in their common wavefront $W_t(\{i, j\})$, the two parts $W_t(\{i, j\}) \cap W_t(i)$ and $W_t(\{i, j\}) \cap W_t(j)$ are connected and non-empty, for all times t . However, these parts touch at B_{ij} , for all t , which prevents B_{ij} from being cyclic.

To prove the converse, assume $V(I_d) \neq \Pi(I_d)$ now. Then, by Theorem 4, there exists a pair $i, j \in I_d$ that violates the non-piercing property. That is, there exists some t such that $W_t(i)$ either disconnects or completely encloses $W_t(j)$ (or vice versa). In the former case, B_{ij} is disconnected, too, whereas B_{ij} is a closed curve in the latter. Thus condition (1) is violated. \square

By Theorems 4 and 5, the well-developed machinery for computing abstract Voronoi diagrams applies to certain straight skeletons.

Corollary 1. *Let F be a set of pairwise non-piercing figures. Then the straight skeleton of F is an abstract Voronoi diagram, with the figures in F as its sites, and with $(SK(\{f, g\}))_{\{f, g\} \subset F}$ as its bisector system.*

Interestingly, the city Voronoi diagram can be defined as a straight skeleton of non-piercing figures, as is shown in the next subsection.

7.2. Figure Adaption

To reach the intended scenario in Corollary 1, we have to modify the figures in the set F^* (which has been constructed in Section 6) to become pairwise non-piercing, but without changing their straight skeleton.

In a first step, each j -needle f is replaced by the j individual 1-needles whose union constitutes f ; compare Fig. 6. This modification clearly has no effect on the straight skeleton. Secondly, each 1-needle g is *domesticated* in the following way. We limit the range of g 's peak p (the unique vertex that moves with network speed v) and let p take on a diamond shape beyond that point. This defines a new figure g' whose wavefront $W_t(g')$ has the shape of a 1-needle in the beginning and contains six edges when fully developed; see Fig. 9. Limitation of p is at its first time-covering network node, or at a sentinel precomputed in Section 6.2. In both cases, p would not have passed further without domestication, which implies that the skeleton is left unchanged.

To see that peak p does no harm *within* its range, observe that p might pierce into, but not through, a wavefront having swept over p 's limitation. When fully developed, all domesticated figures, together with the original diamonds in F^* , are pairwise non-piercing: consider any such pair f, g and assume that $W_t(f)$ and $W_t(g)$ have their first contact in, say, the horizontal direction. Then piercing can take place only horizontally. However, this cannot happen either, because the speed of all vertices of $W_t(f)$ and $W_t(g)$, and particularly of those moving horizontally, is 1.

An abstract Voronoi diagram with m sites and constant bisector complexity can be computed in $O(m \log m)$ time. Particularly attractive is a randomized incremental construction; see [16] and [15]. To adapt this algorithm to our situation, it is sufficient to have a subroutine that accepts five figures as input and returns their straight skeleton. Any trivial skeleton algorithm may be implemented in this subroutine, because all input figures are convex, 12-oriented, and have at most six edges. We conclude a main result of this paper.




Fig. 9. Domesticated 1-needle.




Fig. 10. A complex transportation network C (bold segments) and the city Voronoi diagram $V_C(S)$ (full edges) for a set S of five sites scattered among C . The diagram is refined to $\mathcal{V}_C(S)$ by straight skeleton edges (dashed). Unused portions occur in the network, shown as dotted subsegments. They could be eliminated without delaying any quickest path to S . One site is isolated from the network (the bottommost site s), in the sense that for no point q in the region of s does the quickest path $Q_C(q, s)$ take advantage of C . Outside the smallest rectangle covering $C \cup S$, all edges of $\mathcal{V}_C(S)$ are isothetic.

Theorem 6. *Let C be an isothetic transportation network with constant speed, and let S be a set of sites in the plane. The refined city Voronoi diagram $\mathcal{V}_C(S)$ (Fig. 10) can be computed in $O(n \log n + c^2 \log c)$ time and optimal space $O(n + c)$, where c and n denote the size of C and S , respectively.*

The time complexity clearly is optimal in the number n of sites. We raise the question of whether a subquadratic dependency on the network size c can be achieved, preferably retaining the computational simplicity as in our algorithm. Also, in applications where c is much larger than n , an output-sensitive algorithm for constructing the (unrefined) city Voronoi diagram $V_C(S)$ is of interest.

Acknowledgements

We thank the referees for their valuable comments. The third author gratefully acknowledges discussions on the presented topic with the authors of [1].

References

1. M. Abellanas, F. Hurtado, C. Icking, R. Klein, E. Langetepe, L. Ma, B. Palop, V. Sacristan, Proximity problems for time metrics induced by the L_1 metric and isothetic networks, in: *IX Encuentros en Geometría Computacional*, N. Coll et al. (eds.), Universitat de Girona, 2001, pp. 175–181.
2. O. Aichholzer, D. Alberts, F. Aurenhammer, B. Gärtner, A novel type of skeleton for polygons, *Journal of Universal Computer Science* **1** (1995), 752–761.
3. O. Aichholzer, F. Aurenhammer, Straight skeletons for general polygonal figures, in: *Proceedings of the 2nd Annual International Computing and Combinatorics Conference (COCOON '96)*, Lecture Notes in Computer Science 1090, Springer-Verlag, Berlin, 1996, pp. 117–126.
4. O. Aichholzer, F. Aurenhammer, Straight skeletons for general polygonal figures in the plane, in: *Voronoi's Impact on Modern Sciences II*, A.M. Samoilenko (ed.), Proc. Institute of Mathematics of the National Academy of Sciences of Ukraine 21, Kiev, 1998, pp. 7–21.
5. F. Aurenhammer, H. Edelsbrunner, An optimal algorithm for constructing the weighted Voronoi diagram in the plane, *Pattern Recognition* **17** (1984), 251–257.
6. F. Aurenhammer, R. Klein, Voronoi diagrams, in: *Handbook of Computational Geometry*, J. Sack, G. Urrutia (eds.), Elsevier Science, Amsterdam, 2000, Chapter 5, pp. 201–290.
7. Y. Berman, Transportation Voronoi diagrams, Manuscript, School of Computer Science and Engineering, The Hebrew University of Jerusalem, 2002.
8. S.-W. Cheng, A. Vigneron, Motorcycle graphs and straight skeletons, *Proceedings of the 13th Annual ACM-SIAM Symposium on Discrete Algorithms*, 2002, pp. 156–165.
9. E.D. Demaine, M.L. Demaine, J.S.B. Mitchell, Folding flat silhouettes and wrapping polyhedral packages: new results in computational origami, *Computational Geometry: Theory and Applications* **16** (2000), 3–21.
10. H. Edelsbrunner, R. Seidel, Voronoi diagrams and arrangements, *Discrete & Computational Geometry* **1** (1986), 25–44.
11. D. Eppstein, J. Erickson, Raising roofs, crashing cycles, and playing pool: applications of a data structure for finding pairwise interactions, *Discrete & Computational Geometry* **22** (1999), 569–592.
12. L. Gewali, A. Meng, J.S.B. Mitchell, S. Ntafos, Path planning in $0/1/\infty$ weighted regions with applications, *ORSA Journal on Computing* **2** (1990), 253–272.
13. F. Hurtado, B. Palop, V. Sacristan, Diagramas de Voronoi con funciones temporales, in: *VIII Encuentros en Geometría Computacional*, J. Badenas et al. (eds.), Universitat Jaume I, 1999, pp. 279–287.
14. R. Klein, *Concrete and Abstract Voronoi Diagrams*, Lecture Notes in Computer Science 400, Springer-Verlag, Berlin, 1989.

15. R. Klein, K. Mehlhorn, S. Meiser, Randomized incremental construction of abstract Voronoi diagrams, *Computational Geometry: Theory and Applications* **3** (1993), 157–184.
16. K. Mehlhorn, S. Meiser, C. O’Dunlaing, On the construction of abstract Voronoi diagrams, *Discrete & Computational Geometry* **6** (1991), 211–224.
17. J.S.B. Mitchell, On maximum flows in polyhedral domains, *Proceedings of the 4th Annual ACM Symposium on Computational Geometry*, 1988, pp. 341–351.
18. J.S.B. Mitchell, Shortest paths among obstacles in the plane, *Proceedings of the 9th Annual ACM Symposium on Computational Geometry*, 1993, pp. 308–317.
19. J.S.B. Mitchell, C.H. Papadimitriou, The weighted region problem: finding shortest paths through a weighted planar subdivision, *Journal of the ACM* **38** (1991), 18–73.
20. E. Papadopoulou, D.T. Lee, The L_∞ Voronoi diagram of segments and VLSI applications, *International Journal of Computational Geometry & Applications* **11** (2001), 503–528.

Received June 25, 2002, and in revised form April 17, 2003. Online publication November 10, 2003.