Quickest Paths, Straight Skeletons, and the City Voronoi Diagram

Oswin Aichholzer[*]
Institute for Theoretical Computer Science
Graz University of Technology
Inffeldgasse 16b, A-8010
Graz, Austria
oaich@igi.tu-graz.ac.at

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

Belén Palop[†]
Departamento 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 speeded up 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 relation 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.

Categories and Subject Descriptors

F.2.2 [**Nonnumerical Algorithms**]: Geometric problems and computations—*computations on discrete structures*

General Terms

Algorithms, Theory

Keywords

Voronoi diagram, transportation network, shortest path, straight skeleton, wavefront model

[*]Research supported by APART [Austrian Programme for Advanced Research and Technology] of the Austrian Academy of Sciences

[†]Research supported by Proyecto MEC-DGES-SEUID-PB98-0933. Part of work was carried out when this author visited the Institute for Theoretical Computer Science, Graz University of Technology.

SoCG'02, June 5-7, 2002, Barcelona, Spain.

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 like 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 also may be commercial interest (from real estate agents, or from a tourist office) to make visible the area which can be reached in, say one hour, from a given location in the city (the apartment for sale, or the recommended hotel). 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.

The present 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 to simply follow it to an exit point close to the desired destination; taking intermediate shortcuts by foot walking may be advantageous at several places. Still, quickest paths induce a metric in the plane. We will 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 it 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 quickest paths from arbitrary query points to S in 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 nonstandard task. Cyclic site bisectors may occur, and a single bisector may be as complex as the whole transportation network. To achieve satisfactory performance we will 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 of 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 posed 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 on.

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 Subsection 2.2. Once having accessed C we travel at 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.

The model has been introduced first in Abellanas et al. [1] who derive basic properties of the quickest-path metric; we will revisit some of them in Section 3. They also make 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.

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 Figure 5 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 a careful adaption of the skeleton figures.

The situation changes if the model is even slightly altered. Non-isothetic networks, individual travelling 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 can be shown by simple examples. Some results for a single-line transportation network under the Euclidean metric are discussed in Hurtado et al. [12]. In view of the prohibitive (i.e., quadratic) diagram size, we refrain from further considering such models. In fact, the Euclidean city Voronoi diagram, which seems to be the most interesting alternative, is to some extent approximated in shape by our linear-size model. This gives additional confirmation of the chosen model.

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 speed of travelling inside the region. Regions may be entered and left at any point. For two given query points p and q, the quickest (Euclidean) path between them is sought. Mitchell and Papadimitriou [17] 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 (geodesic path) problem for which much more efficient solutions are known; see e.g. Mitchell [16].

For our transportation network C, quickest-path queries can be viewed as a one-dimensional instance of the weighted region problem. Choose constant weight v on all the network segments, and weight 1 on all connected parts of the complement of C in the plane. If the underlying metric is L_1, we are able to solve general quickest-path queries in $O(c^2 \log c)$ time, by constructing the (refined) city Voronoi diagram $V_C(\{p\})$ of the first query point p and performing point location for the second query point q.

2.2 Airlift Voronoi diagrams

Another related concept is the *airlift distance* considered in Aurenhammer and Klein [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 to $V_C(S)$ this is considerably simpler a 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 in Berman [7] who posed the question of whether a 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 Klein [13]. 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. 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.

2.3 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 Aichholzer et al. [2] and has been generalized to planar straight-line graphs in Aichholzer and Aurenhammer [3, 4]. Various applications exist, some being seemingly unrelated; see [4, 8, 9, 11, 18] and references therein. The present paper adds to this list.

Let F be an arbitrary set of simple polygons (called *figures*) in the plane. Each figure is associated with a birth time, and an individual speed for each of its edges to move outwards in a self-parallel fashion. In this way, each figure $f \in F$ sends out a wavefront $W(f)$. The *straight skeleton* $SK(F)$ of F is defined as the interference pattern of these wavefronts, under the requirement that their movement ceases at all points in the plane where wavefronts come into contact or self-contact. During the propagation of the wavefronts, their vertices trace out the *edges* of $SK(F)$, and their edges sweep out the *faces* of $SK(F)$.

The edges of $SK(F)$ are portions of (weighted) angle bisectors, and thus are straight line segments. The faces of $SK(F)$ are connected polygons, which define a partition of the plane with vertex degree at least three. 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. Due to this 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 $O(e^{\frac{8}{5}+\varepsilon})$, 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 3D.

Interestingly, Voronoi behavior is retained when F consists of so-called *non-piercing* figures. We prove in Section 6 that, for this case, $SK(F)$ is an *abstract Voronoi diagram* in the sense of Klein [13]. Beside these and related other general findings, the merit of Sections 4 through 6 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.

3. THE CITY VORONOI DIAGRAM

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

3.1 Quickest paths

Let a network C as described in Subsection 1.1 be given. For two points x and y in the plane, let $Q_C(x, y)$ denote a

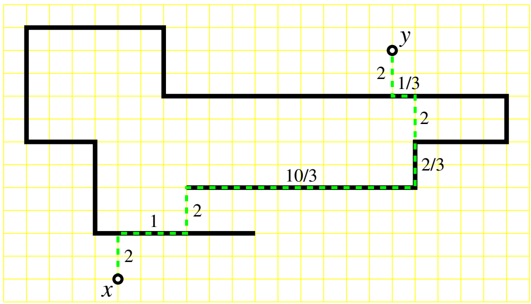

Figure 1: Quickest path for $v = 3$

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 $Q_C(x, y)$. Recall that d_C induces a metric (which we shall 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 for a simple-shaped network C which, though, is complex enough to enforce two *shortcuts* (intermediate pieces leaving C) as well as a non-monotonic behavior. 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.

Even worse, the number of potential shortcuts for C is $\Theta(c^2)$, 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 Figure 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 nontrivial limitation on the number of accesses to C results in a loss of the triangle inequality.

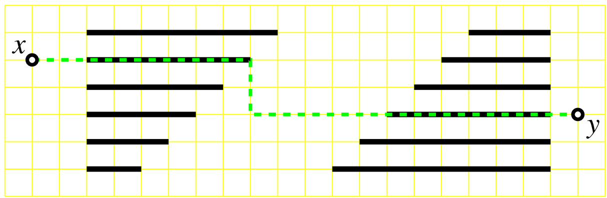

Figure 2: Many shortcuts between terminals

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

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

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. Still, the 'disk' interior to $K(x, r)$ is a connected set, as being the union of quickest paths; see Figure 4.

3.2 The diagram

Let now 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

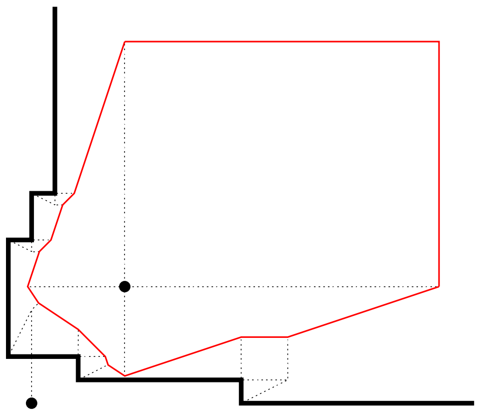

Figure 3: Cyclic bisector with $\Theta(c)$ edges

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 Figure 5 for an illustration.

OBSERVATION 1. *The regions of $V_C(S)$ are connected.*

PROOF. We show that $x \in reg(s)$ implies $Q_C(x,s) \subset reg(s)$, that is, the region $reg(s)$ is path-connected. 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 4. In fact, slight alterations of our model lead to diagrams with $\Theta(n \cdot c)$ vertices and edges, as can be shown by simple examples.

Finally, let us 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 due to Observation 1 and – apart from obvious L_1-specific degenerations – B_{st} is a piece-wise linear curve because the L_1 metric is affected by C only in an isothetic manner. The next property, exemplified in Figure 3, is undesirable in view of an algorithmic construction of $V_C(S)$.

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

Cyclic and non-constant sized bisectors are known to be a main obstacle for efficiently applying divide-and-conquer and randomized incremental insertion.

4. LINK TO STRAIGHT SKELETONS

4.1 The reduction

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

To aid the intuition of the reader, let us 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.

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 wavefront propagated from a certain set of polygonal figures.

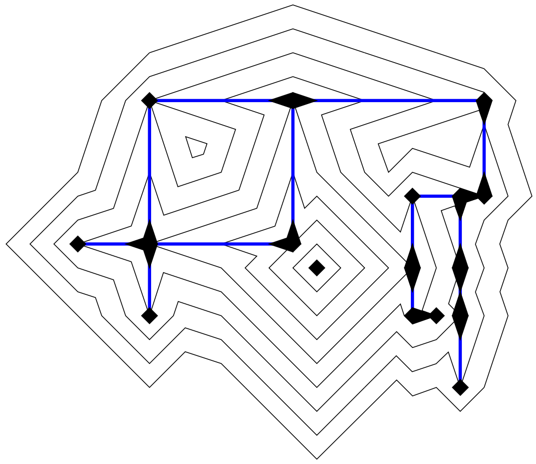

Figure 4: Concentric d_C circles and figure generation

See Figure 4. For sufficiently small t, $K(s,t)$ is a diamond with center s as before. Let F_s be a set of polygonal figures which initially contains this diamond. 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 $K(s,t)$ will change its shape. A diamond D with center x will appear at time $t + \varepsilon$, along with sharp-angled wedges tangent to D whose peaks move at speed v in all possible directions on the network that point to the exterior of $K(s,t)$. To construct the corresponding figure to be put into the set F_s, the union of D with its tangent wedges is taken. We term the point x the *center* (or place of birth) of this figure.

The new figures are all centered on C and may be of various shapes. We categorize them 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 four, and at least one of the four isothetic directions points to the interior of $K(s,t)$. Only for 0-, 1-, and 2-needles the diamond D may

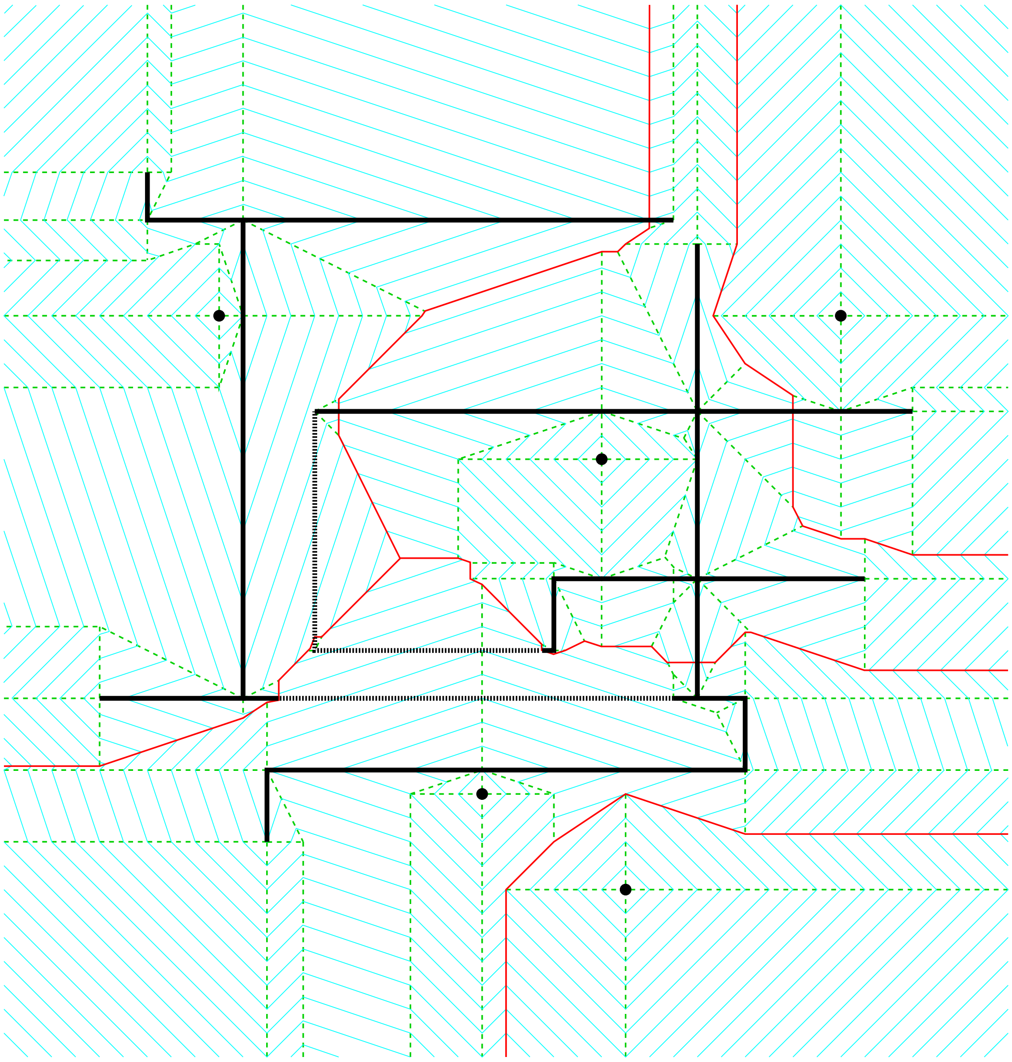

Figure 5: Illustration of a city Voronoi diagram

contribute edges to the figure. For such edges, the translational speed is $1/\sqrt{2}$ as before, whereas the translational speed of wedge 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.

The set F_s thus contains all figures (including their times and places of birth) defined by the site s and the network C. Let $W_t(F_s)$ denote the common wavefront propagated from these figures at time t. 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.

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.*

4.2 Path encoding

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 come back to this issue in Subsection 4.3.) Edges of $SK(F)$ which are no edges of $V_C(S)$ stem from non-redundant figures born in the *same* region of $V_C(S)$. We will alternatively call $SK(F)$ the *refined* city Voronoi diagram of S and C, or $\mathcal{V}_C(S)$ for short. The 2D 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 $SK(F)$.

Figure 5 displays 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). Quiet ranges (confer Subsection 4.3) occur in the network, shown as dotted subsegments. They could be closed down without delaying any quickest path to S. Also, one site is isolated from the network (the bottom-most site s), in the sense that for no point q in the region of s the quickest path $Q_C(q,s)$ takes advantage of C. Note that outside the smallest window covering $C \cup S$ all edges of $\mathcal{V}_C(S)$ are isothetic.

THEOREM 2. *$\mathcal{V}_C(S)$ encodes, for every point q in the plane, the quickest path from q to S.*

PROOF. 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 directly go from a site to the network), and others are of a third type. Let us orient all these edges in the direction they are traced out by the vertices of the wavefront $W_t(F)$ propagated from F. This yields a directed acyclic graph with the sites in S as sources because, intuitively speaking, time does not run backwards. Moreover, when selecting all the edges of the first two types, a quickest-path tree $T(s)$ for each site s within its region $reg(s)$ is obtained. Note that $T(s)$ is built from isothetic edges.

Let now s be the site with $q \in reg(s)$. Then $Q_C(s,q)$ is the quickest path from q to S. $Q_C(s,q)$ completely consists of edges of the quickest-path tree $T(s)$, except possibly for the edge e ending at q, in case q does not lie on C or in S. In that case, e lies in the interior of a single face of $\mathcal{V}_C(S)$, and the second-last edge of $Q_C(s,q)$ (which still belongs to $T(s)$) lies on the boundary of that face. $\square$

COROLLARY 1. *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.*

Observe at this place that each face g of $\mathcal{V}_C(S)$ can be uniquely associated with one of the *defining features* of $\mathcal{V}_C(S)$: a site, a network node, or a (portion of a) network segment. This feature, call it $\beta(g)$, contains the place of birth of the very figure whose wavefront sweeps out the face g with one of its edges. In fact, g is just the union of all quickest paths $Q_C(x,q)$, for $x \in \beta(g)$ and $q \in g$. Quickest paths are isothetic, so every straight line normal to $\beta(g)$ intersects g in a connected set. This provides us with a property not shared by straight skeletons in general.

LEMMA 1. *The faces of $\mathcal{V}_C(S)$ are monotone polygons.*

Monotonicity implies simple connectedness; in contrast, general straight skeleton faces may contain holes.

4.3 The linear size

The diagram $\mathcal{V}_C(S) = SK(F)$ is a planar graph whose faces are simply connected by Lemma 1, and whose vertices have degree at least three, 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 center s. Moreover, by arguments in Subsection 4.2, each node of C gives rise to a single non-redundant j-needle centered at this node. This gives $n+c$ relevant figures so far. It remains to examine the 2-needles centered in the interior of network segments. They will be called *interior 2-needles*, in contrast to the 2-needles centered at network nodes. Their number is $\Theta(n \cdot c)$ in the worst case: each of the n sites may cause c such figures, centered on c segments that lie one behind the other; cf. Figure 4, right lower corner. However, the following fact comes to our rescue.

LEMMA 2. *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 Figure 6, with the common vertex x as the center of homothety. This vertex coincides, at time t, with the center of f'. $\square$

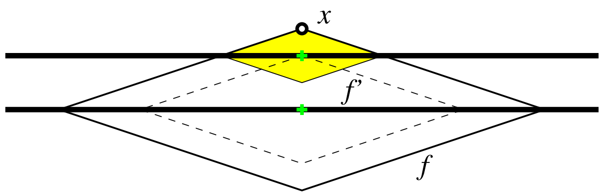

Figure 6: Redundant 2-needle

By Lemma 2, only the interior 2-needles that stem from contacts between network segments and *diamond vertices* (rather than needle vertices) are non-redundant in the set F. As diamond vertices disappear from the wavefront of F after 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 2, are the reason for *quiet ranges* on C, i.e., portions of C that never need to be used. A quiet range 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 Figure 5.

5. 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^*. This seems satisfactory in view of the existence of $\Theta(c^2)$ potential shortcuts in the network C; see Subsection 3.1.

The diamonds for the sites in S are calculated in a trivial manner, so let us concentrate on the figures caused by the network C; consult Subsection 4.1.

5.1 Figures at nodes

We treat figures centered 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 now is the union of grid segments. For a site $s \in S$, let $G(s)$ denote the grid box which contains s.

LEMMA 3. *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 an access point, p, of s to C. $\square$

Let us call the point p above a *primary access point*. The following algorithm is used to determine when (and from where) each grid node is conquered by the sites. This will give us the birth times and shapes of all non-redundand figures (j-needles) 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 $O(c^2 \log c)$ time.

5.2 Figures on segments

Now consider figures centered on network segments; we called them interior 2-needles in Subsection 4.3. Redundancy occurs for two reasons, one made precise in Lemma 2. But still an interior 2-needle which 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 be redundant: there might be a site s whose circle $K(s,t)$ has swept over p already. To single out 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 *all* access points (not necessarily primary) of the sites in the column, above and below L respectively. Whenever L reaches a grid edge g which belongs to C we do the following. For each *primary* access point on g, and for g's nodes, detect all time-covered neighbors (sites) in A and B. Remove them (and file the new neighbors for later use as sentinels, in Section 6). If a primary access point is covered itself by A or B then remove the corresponding site from both lists (and update sentinels accordingly). Note that access times of sites to L can be updated quickly during the sweep, by maintaining offset times t_A and t_B for A and B. 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 timed) 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 $O(n \log n + c^2 \log c)$.

6. ABSTRACT VORONOI SETTINGS

We are left with the problem of computing the straight skeleton of the set F^*. According to Subsection 4.3, 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 Subsection 2.3. We show below how $SK(F^*)$ 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.

6.1 Relation 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 relation between these models. We take the chance to clarify the situation, with the special goal to resolve the problem of computing $SK(F^*)$ efficiently.

The classical model is, of course, the *distance model*. Let $I_d = \{1, 2, \ldots, m\}$ 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 Edelsbrunner and Seidel [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 $t \geq 0$ measuring time, and the interference pattern of all these wavefronts constitutes the diagram; compare Subsection 2.3.

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, in Aurenhammer and Edelsbrunner [5]. Being more important in the present context, straight skeletons are known to admit no distance-from-site definition, in general; see Aichholzer and Aurenhammer [3].

We now give characterizing conditions for the equivalence of both models. Call two sites i and 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 $V(\{i,j\}) \neq \Pi(\{i,j\})$.

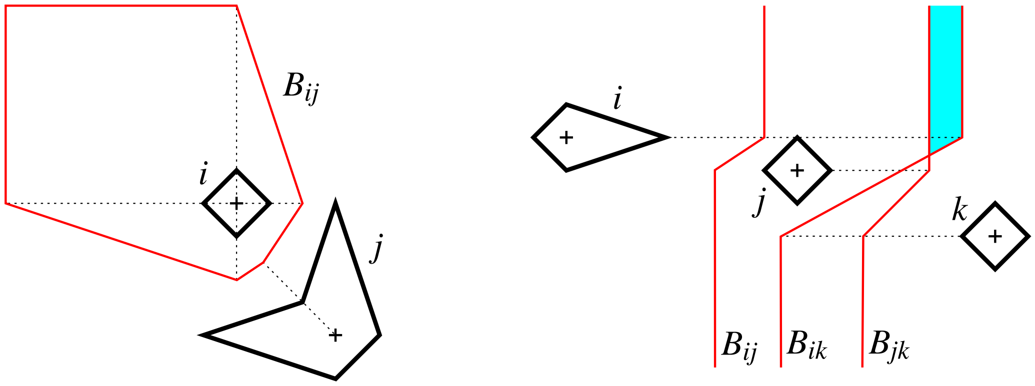

Figure 7: Diamonds and needles violating (1) and (2)

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. But this is equivalent to $V(J) = \Pi(J)$. $\square$

A third well-known model is the *abstract Voronoi diagram* introduced in Klein [13]. 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) $B_{ij} \cap B_{ik} \subset B_{jk}$

(3) Regions are path-connected

Actually, an additional technical condition is postulated,

(4) $B_{ij} \cap B_{ik}$ 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 7 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 lands (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 $B_{ij} = \{x \mid d(x, i) = d(x, i)\}$, 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 halfplane defined by B_{ij} is path-connected. This implies that B_{ij} is a connected curve. We show that B_{ij} is acyclic. Our assuption $V(I_d) = \Pi(I_d)$ 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. But 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 is a point t in time where $W_t(i)$ 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 2. *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.

6.2 Figure adaption

To reach the intended scenario in Corollary 2, we have to modify the figures in the set F^* (which has been constructed in Section 5) so as to become pairwise non-piercing, but without changing their straight skeleton.

In a first step, we replace each j-needle f by the j individual 1-needles whose union constitutes f. This modification clearly has no affect on the straight skeleton. Secondly, we *domesticate* each 1-needle g, by limiting the range of its peak p (that moves with network speed v) and letting p take on diamond shape beyond that point. Note that this gives a new figure g' whose wavefront $W_t(g')$ is well defined. In particular, $W_t(g')$ contains six edges after having fully developed; see Figure 8. Limitation of p is at its first time-covering network node, or at a sentinel precomputed in Section 5. In both cases, p would not have passed further

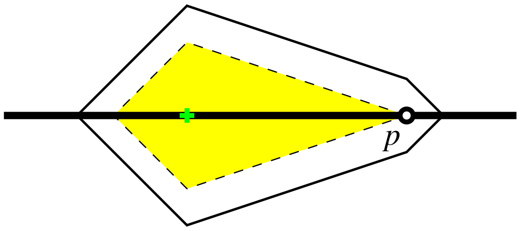

Figure 8: Domesticated 1-needle

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. Outside their ranges all domesticated figures, together with the original diamonds in F^*, are pairwise non-piercing: consider any such pair f and g and assume that $W_t(f)$ and $W_t(g)$ have their first contact in, say, horizontal direction. Then piercing can take place only horizontally. But 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 Mehlhorn et al. [15] and Klein et al. [14]. To adapt this algorithm to our situation, it it 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 6 edges. We conclude a main result of this paper.

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 $\mathcal{V}_C(S)$ 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.*

Note that the time complexity is optimal to within an additive term which does not depend on the sites. We raise the question of whether a subquadratic dependency on the network size can be achieved, preferably with retaining computational simplicity as in our algorithm.

Acknowledgements

The third author gratefully acknowledges discussions on the presented topic with the authors of the paper [1].

7. 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.* IX Encuentros en Geometri a Computacional, N. Coll et al. (eds.), Universitat de Girona, 2001, 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: Proc. 2nd Ann. Int'l Computing and Combinatorics Conf. COCOON'96, Springer Lecture Notes in Computer Science 1090 (1996), 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, Ukraine, 1998, 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.* Handbook of Computational Geometry, Chapter 5, J.Sack, G.Urrutia (eds.), Elsevier Science Publishing, 2000, 201-290.

[7] Y.Berman, Personal communication. School of Computer Science and Engineering, The Hebrew University of Jerusalem, Israel, 2001.

[8] S.-W.Cheng, A.Vigneron, *Motorcycle graphs and straight skeletons.* Proc. 13th ACM-SIAM Symposium on Discrete Algorithms 2002.

[9] E.D.Demaine, M.L.Demaine, J.S.B.Mitchell, *Folding flat silhouettes and wrapping polyhedral packages: new results in computational origami.* Proc. 15th Ann. ACM Symp. Computational Geometry, 1999, 105-114.

[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.* Proc. 14th Ann. ACM Symp. Computational Geometry, 1998, 58-67.

[12] F.Hurtado, B.Palop, V.Sacristan, *Diagramas de Voronoi con funciones temporales.* VIII Encuentros en Geometri a Computational, J.Badenas et al.(eds.), Universitat Jaume I, 1999, 279-287.

[13] R.Klein, *Concrete and Abstract Voronoi Diagrams.* Springer Lecture Notes in Computer Science 400 (1989).

[14] R.Klein, K.Mehlhorn, S.Meiser, *Randomized incremental construction of abstract Voronoi diagrams.* Computational Geometry: Theory and Applications 3 (1993), 157-184.

[15] K.Mehlhorn, S.Meiser, C.O'Dunlaing, *On the construction of abstract Voronoi diagrams.* Discrete & Computational Geometry 6 (1991), 211-224.

[16] J.S.B.Mitchell, *Shortest paths among obstacles in the plane.* Proc. 9th Ann. ACM Symp. Computational Geometry, 1993, 308-317.

[17] 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.

[18] E.Papadopoulou, D.T.Lee, *The L_∞ Voronoi diagram of segments and VLSI applications.* Int'l Journal of Computational Geometry & Applications 11 (2001), 503-528.