

Transitions in Geometric Minimum Spanning Trees

Clyde Monma and Subhash Suri

Bell Communications Research, 445 South Street,
Morristown, NJ 07960, USA
clyde@bellcore.com
suri@bellcore.com

Abstract. We study some combinatorial and algorithmic problems associated with an arbitrary motion of input points in space. The motivation for such an investigation comes from two different sources: *computer modeling* and *sensitivity analysis*. In modeling, the dynamics enters the picture since geometric objects often model physical entities whose positions can change over time. In sensitivity analysis, the motion of the input points might represent uncertainties in the precise location of objects.

The main results of the paper deal with state transitions in the minimum spanning tree when one or more of the input points move arbitrarily in space. In particular, questions of the following form are addressed: (i) How many different minimum spanning trees can arise if one point moves while the others remain fixed? (ii) When does the minimum spanning tree change its topology if all points are allowed to move arbitrarily?

1. Introduction

A *minimum spanning tree* of a d -dimensional set of points S is a minimum cost tree connecting the set of points S , where the cost of an edge is the Euclidean distance between its endpoints. This paper studies state transitions in minimum spanning trees, where the state refers to the *topology* or the *graph structure* of the tree. Let S be a fixed and labeled set of points in \mathcal{E}^d for some fixed d . We make a standard nondegeneracy assumption that no $(d + 2)$ points are cospherical. This assumption is needed to ensure that not too many points are equidistant from a point. An appropriate tie-breaking rule easily extends our results to general points; however, we assume nondegeneracy to simplify the discussion. (A general tie-breaking scheme is described in [16].)

Let v be a variable point in d -space, and let $MST(S, v)$ denote the minimum spanning tree of the set $S \cup \{v\}$. We say that the trees $MST(S, p)$ and $MST(S, q)$

are *topologically equivalent* if their graphs are the same. Otherwise, the trees are topologically distinct. The main results of our paper are summarized below.

1. (Combinatorial Bounds.) We establish that at least $\Omega(n^d)$ and at most $O(n^{2d})$ topologically distinct minimum spanning trees $MST(S, x)$ can arise if x is allowed to vary over the d -space \mathcal{E}^d for $d \geq 3$. In the lower dimensions, $d \leq 2$, we prove an asymptotically tight bound of $\Theta(n^d)$ for this combinatorial quantity.
2. (Classification.) We compute a subdivision of \mathcal{E}^d into maximally connected cells such that, for all x in a particular cell C , minimum spanning trees $MST(S, x)$ are topologically equivalent. Our algorithm runs in time $O(n^d)$ for $d \leq 2$, and in time $O(n^{2d})$ for $d \geq 3$.
3. (Topology Zone.) We give a linear-time algorithm to check if a given topology T is a feasible. That is, given a topology T on $n + 1$ nodes, we want to check if a point $p \in \mathcal{E}^d$ can be found such that $MST(S, p)$ and T are topologically equivalent. If such a p exists, then we can also compute the *topology zone*, $Z(T)$, which is the maximal, connected subset of \mathcal{E}^d with the property that, $\forall x \in Z(T)$, $MST(S, x)$ has the same topology as T .
4. (Sensitivity Measure.) We consider the problem of determining the maximum distance by which the points of S can be (arbitrarily) moved without altering the topology of their minimum spanning tree. We call this distance the *sensitivity measure* $\mu(S)$, and give an $O(n \log n)$ -time algorithm for computing it in two dimensions.
5. (Embedding.) Finally, we consider the general problem of *embedding* an abstract topology as a minimum spanning tree: Given a (free) tree T on n nodes, does there exist a set of n points $S \subset \mathcal{E}^d$ such that $MST(S)$ is topologically equivalent to T ? Our result is the following theorem for embedding in two dimensions: any tree on n nodes, with maximum node degree five, can be realized as the minimum spanning tree of some set of n points in the plane. Together with the fact that any set of n points in the plane admits a minimum spanning tree with maximum node degree five, our theorem implies that the geometric embedding of a minimum spanning tree has only one (trivial) obstruction: maximum node degree.

There are several applications of the minimum spanning tree where the topology of the tree is of primary concern. In statistical clustering, for example, clusters are determined solely by the graph structure of the minimum spanning tree of the underlying points; see [1] for an application of minimum and maximum spanning trees in finding clusters that maximize (minimize) intercluster (intracluster) distances. In these applications, the structural form of the solution is unaffected by an arbitrary perturbation of the input points, as long as the topology of the minimum spanning tree does not change. Motivated by applications of this nature, we investigate several combinatorial and algorithmic problems related to transitions in geometric minimum spanning trees that arise when one or more of the input points are allowed to move in an arbitrary manner.

Although there is an extensive literature on computing minimum spanning trees, the main focus in the past has been on static input. A notable exception is the

paper by Atallah [2] in which a general framework for some dynamic problems in computational geometry is considered. In the framework of Atallah, points trace out paths that can be parametrized by some fixed-degree polynomials in time. The main algorithmic problems addressed in [2] are concerned with the *steady-state* behavior of the points.

In this paper we address dynamic problems of a different nature. Instead of restricting the motion to be polynomially describable, we allow completely arbitrary motion of the points. However, lest the problems should become totally hopeless, we either consider the motion of only one point (keeping the remaining points fixed) or allow all the points to move simultaneously, independently, and arbitrarily, but only seek the first instant at which the state transition occurs.

Parametric problems have also been considered extensively in the context of network optimization. Examples include work on maximum flow [8], minimum cut [14], minimum spanning tree [10], [11], and stable marriage [13] problems. In these problems the data (costs and capacities) are varied as a function of a parameter λ , and we want to compute an optimal solution as a function of λ . Parametric programming algorithms either produce a finite table of solutions for critical values of λ , or give a compact encoding of the output from which a solution can be determined fast once the actual value of λ is specified. Parametric programming has also found applications in many other problems, such as the rectilinear layout [18], arboricity [17], module distribution [12], and network reliability [5], [14]. Besides being interesting problems in their own right, they also play an important role in understanding/analyzing the sensitivity of optimal solutions to changes in problem data. Our work fits in this general framework, however, the network parameters in our case are induced by a mobile set of points. Not surprisingly, therefore, we are able to produce much sharper results by exploiting the geometry.

The organization of our paper is as follows. In Sections 2 and 3 we establish the combinatorial bounds on the number of topologically distinct minimum spanning trees. In Section 4 we give an algorithm to partition the space into maximally connected cells for which the minimum spanning tree topology is invariant. The results for checking the feasibility of a given topology and for computing the topology zone are presented in Section 5. Section 6 presents our result for computing the sensitivity measure, and Section 7 proves our embedding result (problem 5). Conclusions and open problems are discussed in Section 8.

2. A Lower Bound on the Number of Topologies

We show that there are at least $\Omega(n^d)$ topologically distinct minimum spanning trees $MST(S, x)$, where $x \in \mathcal{E}^d$. Our lower bound construction uses d groups of points, each of which has cardinality $\lfloor n/d \rfloor$. Let a_i denote the unit vector in the direction of the positive i th coordinate axis; the i th coordinate of a_i is 1 and the others are 0. In our construction, the i th group of points S_i lies near a_i . Let

$$b = \left(\frac{1}{d} + \alpha, \frac{1}{d} + \alpha, \dots, \frac{1}{d} + \alpha \right),$$

where $\alpha = 1/d^4$. (We have fixed α to be specific, although any suitably small constant, depending on the dimension, would work.) Observe that a_i 's form the standard basis of \mathcal{E}^d and b is equidistant from the a_i 's.

2.1. Construction of the Sets

Let B denote the d -simplex determined by $\{a_1, a_2, \dots, a_d, b\}$, and let H_i denote the hyperplane spanned by $\{a_1, \dots, a_{i-1}, b, a_{i+1}, \dots, a_d\}$. (Thus, H_i bounds the facet of B opposite a_i .) Observe that H_i satisfies the equation $\sum c_j x_j = 1$, where $c_j = 1$ for $j \neq i$ and $c_i = 1 - \alpha d^2/(1 + \alpha d)$. For any positive real number β , we define a β -family of hyperplanes \mathcal{H}_i^β , for $i = 1, 2, \dots, d$, such that

1. \mathcal{H}_i^β has $\lfloor n/d \rfloor + 1$ hyperplanes $H_i(0), H_i(1), \dots, H_i(\lfloor n/d \rfloor)$, all parallel to H_i .
2. The distance between the origin and $H_i(k)$ is $\Delta_i + k\beta$, where Δ_i is the distance between H_i and the origin.

Let \mathcal{A} denote the arrangement formed by the families of hyperplanes \mathcal{H}_i^β , $i = 1, 2, \dots, d$. Let $U(b, r)$ denote the d -dimensional ball of radius r centered on b . The following lemma shows that, given any $r > 0$, we can choose β such that the hyperplane families \mathcal{H}_i^β , $i = 1, 2, \dots, d$, partition $U(b, r)$ into $\Theta(n^d)$ cells.

Lemma 2.1. *Let $\beta = \frac{1}{3}rn^{-1}d^{-4}$. Then every d -tuple of hyperplanes $\{H^1, H^2, \dots, H^d\}$, where $H^i \in \mathcal{H}_i^\beta$, intersects in a point inside $U(b, r)$.*

Proof. Since $H_1(0) \cap H_2(0) \cap \dots \cap H_d(0) = b$, the corresponding system of linear equations

$$\sum c_{ij}x_j = 1, \quad i = 1, 2, \dots, d, \tag{1}$$

where $c_{ii} = 1 - \alpha d^2/(1 + \alpha d)$, and $c_{ij} = 1$ for $j \neq i$, has the unique solution $x_i = 1/d + \alpha$ for $i = 1, 2, \dots, d$. The k th hyperplane of the family \mathcal{H}_i^β is obtained by a parallel translation of $H_i(0)$ by the distance $k\beta$. Thus,

$$H_i(k): \sum c_j x_j = 1 + k\beta\sqrt{d - 1 + \gamma^2}, \tag{2}$$

where $\gamma = 1 - \alpha d^2/(1 + \alpha d)$, $c_i = \gamma$, and $c_j = 1$ for $j \neq i$.¹

¹ Observe that the Hessian normal form equation for H_i is

$$\sum n_j x_j = (d - 1 + \gamma^2)^{-1/2},$$

where (n_1, n_2, \dots, n_d) is the unit normal vector of H_i , with $n_i = \gamma(d - 1 + \gamma^2)^{-1/2}$ and $n_j = (d - 1 + \gamma^2)^{-1/2}$ for $j \neq i$. The distance between H_i and the origin is $(d - 1 + \gamma^2)^{-1/2}$. The equation for $H_i(k)$ is now easily obtained by replacing the right-hand side of the equation by $1/\sqrt{d - 1 + \gamma^2} + k\beta$.

Now, consider an arbitrary set of d hyperplanes H^1, H^2, \dots, H^d , where $H^i \in \mathcal{H}_i^\beta$. This defines a system of linear equations:

$$\begin{aligned} \gamma x_1 + x_2 + \dots + x_d &= 1 + \delta_1, \\ x_1 + \gamma x_2 + \dots + x_d &= 1 + \delta_2, \\ &\vdots \\ x_1 + x_2 + \dots + \gamma x_d &= 1 + \delta_d, \end{aligned} \tag{3}$$

where $\max\{\delta_1, \delta_2, \dots, \delta_d\} < n\beta/\sqrt{d}$. The bound on δ_i follows from (2), together with the observations that $\gamma^2 < 1$ and $k \leq n/d$.

We now need to choose $\beta > 0$ so that the solution to the system of (3) lies within distance r of b . Elementary algebraic manipulations show that if $(x_1^*, x_2^*, \dots, x_d^*)$ satisfies the system (3), then

$$x_i^* \leq \left(\frac{1}{d} + \alpha\right) + \delta^* \left(1 \pm 2d \pm \frac{2}{\alpha}\right) \left(\frac{1}{d} + \alpha\right),$$

where $\delta^* = \max\{\delta_1, \delta_2, \dots, \delta_d\}$. The error term is strictly less than $3\delta^*d^3$, which we can bound to be at most r/\sqrt{d} by choosing

$$\beta \leq \frac{r}{3nd^3}. \tag{4}$$

Thus, each coordinate x_i^* is within distance r/\sqrt{d} from $x_i = 1/d + \alpha$, implying that the solution vector x^* lies within the ball $U(b, r)$. This completes the proof. \square

Lemma 2.1 says that, for an appropriate choice of β , we can partition the interior of $U(b, r)$ into $\Theta(n^d)$ cells by d β -families of parallel hyperplanes. We now describe how to choose the point-sets S_i 's.

We pick r to be a small constant, say, $r = 0.1d^{-4}$. By choosing the appropriate β , construct the hyperplane families \mathcal{H}_i^β , for $i = 1, 2, \dots, d$, so that Lemma 2.1 holds. We match the vertices of the d -simplex B to the hyperplanes so that a_i and H_j are paired only if $a_i \in H_j$. In particular, let π be a permutation on the set $\{1, 2, \dots, d\}$ such that $a_i \in H_{\pi(i)}$. The family $\mathcal{H}_{\pi(i)}^\beta$ is then assigned to a_i . Now, consider the line l_i through a_i that is directed normal to $H_{\pi(i)}$, and thus normal also to all the planes of the family $\mathcal{H}_{\pi(i)}^\beta$. Suppose that the planes of the family $\mathcal{H}_{\pi(i)}^\beta$ intersect l_i in points $u_0, u_1, \dots, u_{\lfloor n/d \rfloor}$. We choose S_i to be the set of $\lfloor n/d \rfloor$ points $\frac{1}{2}(u_i + u_{i+1})$ for $i = 0, 1, \dots, \lfloor n/d \rfloor - 1$. Figure 1 illustrates this construction for $d = 2$.

Next we show that a variable point x can induce $\Omega(n^d)$ topologically distinct minimum spanning trees on the set $S = \bigcup_{i=1}^d S_i$.

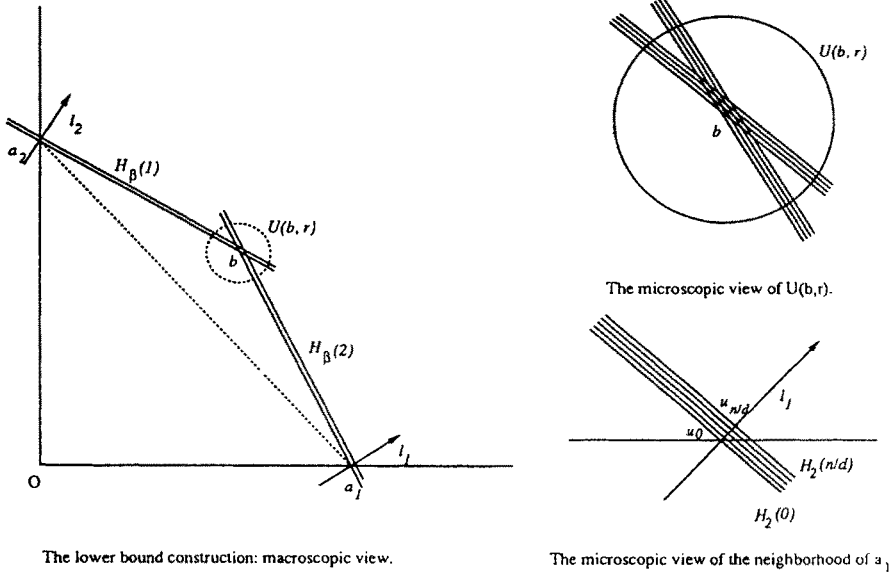


Fig. 1. The lower bound construction in the plane.

2.2. Proof of the Lower Bound

Let p be an arbitrary point in \mathcal{E}^d . Define the *neighbor-vector* of p to be the ordered d -tuple (q_1, q_2, \dots, q_d) where q_i is such that $dist(p, q_i) \leq dist(p, q)$ for any $q \in S_i$; q_i is a nearest neighbor of p in S_i .

Lemma 2.2. *As x varies over $U(b, r)$, it induces $\Omega(n^d)$ distinct neighbor-vectors.*

Proof. Consider $S_i = \{v_1, v_2, \dots, v_{\lfloor n/d \rfloor}\}$. The Voronoi diagram of S_i consists of $\lfloor n/d \rfloor$ parallel slabs, defined by the hyperplanes $H_{\pi(i)}(1), \dots, H_{\pi(i)}(\lfloor n/d \rfloor - 1)$. The nearest neighbor of x in S_i is v_j if and only if x and v_j lie in the same slab. Since the hyperplane arrangement formed by the planes $H_{\pi(i)}(1), \dots, H_{\pi(i)}(\lfloor n/d \rfloor - 1)$, for $i = 1, 2, \dots, d$, partitions the space into $\Theta(n^d)$ cells, there are $\Theta(n^d)$ distinct neighbor-vectors of x ; each cell of the arrangement gives rise to a distinct neighbor-vector. □

Finally, the number of distinct neighbor-vectors of x is also a lower bound on the number of distinct topologies. Our proof will show that each cell of the arrangement \mathcal{A} , formed by the hyperplanes $H_{\pi(i)}(1), \dots, H_{\pi(i)}(\lfloor n/d \rfloor - 1)$, for $i = 1, 2, \dots, d$, induces a distinct topology. Since there are $\Omega(n^d)$ cells in \mathcal{A} , the lower bound on the number of distinct topologies would follow.

To prove the claim, we consider an arbitrary cell C of \mathcal{A} and a point $x \in C$. Let (q_1, q_2, \dots, q_d) be the neighbor-vector of x . We show that the set of edges incident with x in the minimum spanning tree $MST(S, x)$ is precisely

$\{xq_1, xq_2, \dots, xq_d\}$, which suffices for the proof since the neighbor-vectors for each of the $\Theta(n^d)$ cells of \mathcal{A} are distinct. Let $\bar{S}_i = (S \cup \{x\}) - S_i$ denote the complement of the set S_i . Clearly, it suffices to show that $MST(S, x)$ has exactly one edge between S_i and \bar{S}_i , and that edge is incident to x . By construction, the maximum distance between any two points of S_i is

$$l_1 = \frac{\beta n}{d} = \left(\frac{0.1d^{-4}}{3nd^3}\right)\binom{n}{d} < \frac{1}{100d^8}. \tag{5}$$

The distance between a point $p_i \in S_i$ and any point $p_j \in S_j$, for $j \neq i$, is at least

$$l_2 \geq \text{dist}(a_i, a_j) - 2l_1 = \sqrt{2} - 2l_1. \tag{6}$$

On the other hand, the distance between S_i and any $x \in U(b, r)$ is at most

$$l_3 \leq \text{dist}(a_i, b) + l_1 + r < \sqrt{1 - \frac{1}{2d}} + \frac{1}{100d^8} + \frac{1}{10d^4}. \tag{7}$$

Inequalities (5) and (7) show that the diameter of S_i is much smaller than the distance between S_i and \bar{S}_i . Thus, exactly one edge connects S_i with \bar{S}_i , and, since $l_3 < l_2$, that edge is incident to x . This establishes the main result of this section.

Theorem 2.3. *As x varies over the d -dimensional space \mathcal{E}^d , it induces $\Omega(n^d)$ topologically distinct minimum spanning trees $MST(S, x)$ on the set of points S .*

Although our construction involves collinear points, it is easy to see that the points can be perturbed slightly so that they are in a general position, without affecting the lower bound.

Observe that if we restrict the motion of x to a k -dimensional subspace of \mathcal{E}^d , for $k \leq d$, then the number of distinct topologies is $\Omega(n^k)$. This follows from our construction since any k -dimensional subspace passing through b intersects $\Omega(n^k)$ cells of the arrangement \mathcal{A} , and each cell produces a distinct topology. We, therefore, have the following corollary of Theorem 2.3.

Corollary 2.4. *As x varies over a k -dimensional subspace of \mathcal{E}^d , for $k \leq d$, it induces $\Omega(n^k)$ topologically distinct minimum spanning trees $MST(S, x)$ on the set of points S .*

The following section addresses the question of an upper bound on the number of distinct topologies. We first give a general upper bound of $O(n^{2d})$ for any fixed dimension d , and then show how to obtain the tight bounds of $O(n^d)$ in the lower dimensions $d \leq 2$.

3. An Upper Bound on the Number of Topologies

We use a result of Yao [20] to reduce the minimum spanning tree problem to the problem of computing certain *geographic neighbors*. We begin with some definitions.

Given a basis $B = \{b_1, b_2, \dots, b_d\}$ of \mathcal{E}^d , the *convex cone* of B is defined as

$$\text{cone}(B) = \left\{ \sum_{i=1}^d \lambda_i b_i \mid \lambda_i \geq 0 \text{ for all } i \right\}.$$

The *geographic set* of a point $v \in \mathcal{E}^d$, relative to B , is defined as

$$G(v; B) = \{s \in S \mid s - v \in \text{cone}(B)\}.$$

The *geographic nearest neighbor* of v for the basis B , denoted as $GN(v; B)$, is defined to be a point $w \in G(v; B)$ such that $\text{dist}(v, w) \leq \text{dist}(v, y)$ for all $y \in G(v; B)$. Following Yao [20], we call a (finite) family of bases \mathcal{B} a *frame* if $\bigcup_{B \in \mathcal{B}} \text{cone}(B) = \mathcal{E}^d$. The *set of geographic nearest neighbors* of a point v for a frame \mathcal{B} is defined as

$$GN(v; \mathcal{B}) = \{GN(v; B) \mid B \in \mathcal{B}\}.$$

Figure 2 illustrates these definitions in the plane.

Let $E(S; \mathcal{B})$ denote the set of edges connecting the points of S with their geographic nearest neighbors for the frame \mathcal{B} :

$$E(S; \mathcal{B}) = \bigcup_{v \in S} \{(v, u) \mid u \in GN(v; \mathcal{B})\}.$$




Fig. 2. Basis, frame, and geographic sets. Bases $B_1 = \{b_1, b_2\}$, $B_2 = \{b_2, b_4\}$, $B_3 = \{b_4, b_5\}$, and $B_4 = \{b_4, b_1\}$. Frame $\mathcal{B} = \{B_1, B_2, B_4\}$. Members of $G(v; B_1)$ are drawn as circles, while the remaining points are drawn as crosses.

We need one more definition. The *angular diameter* of a basis B is given by $Ang(B) = \max\{\theta(b_i, b_j) \mid b_i, b_j \in B\}$, where $\theta(b_i, b_j) = \arccos(b_i \cdot b_j) / (\|b_i\| \|b_j\|)$. The angular diameter of the frame \mathcal{B} is given as $Ang(\mathcal{B}) = \max\{Ang(B) \mid B \in \mathcal{B}\}$. The following result of Yao [20] says that for sufficiently narrow frames the set $E(S; \mathcal{B})$ contains $MST(S)$.

Lemma 3.1 [21]. *If \mathcal{B} is a frame with $Ang(\mathcal{B}) < \arcsin 1/(2d)$, then $E(S; \mathcal{B})$ contains $MST(S)$. We can construct a frame \mathcal{B} with $Ang(\mathcal{B}) < \arcsin 1/(2d)$ in $O(1)$ time for any fixed d .*

Remark. The maximum degree of a node in $MST(S)$ is of the order $2^{O(d)}$, which is $O(1)$ for fixed d . This follows from a packing argument, but can also be observed from Lemma 3.1.

We now proceed to prove our upper bound on the number of minimum spanning tree topologies.

3.1. Proof of the Upper Bound

Let $\mathcal{B} = \{B_1, B_2, \dots, B_k\}$, where $k = O(1)$, be a frame satisfying the condition of Lemma 3.1, and let $B_i = \{b_i(1), b_i(2), \dots, b_i(d)\}$ be a basis in \mathcal{B} . Denote by $F_i(j)$ the linear space spanned by $B_i - \{b_i(j)\}$, for $j = 1, 2, \dots, d$. Call the hyperplanes $F_i(j)$'s the *basis planes* of B_i . Let \mathcal{F} denote the family of basis planes for all the bases in \mathcal{B} :

$$\mathcal{F} = \{F_i(j) \mid 1 \leq i \leq k, 1 \leq j \leq d\}. \tag{8}$$

Since \mathcal{B} has size $O(1)$, it follows that $|\mathcal{F}| = O(1)$. We now construct two families of hyperplanes \mathcal{H}_1 and \mathcal{H}_2 , where the former consists of the bisectors of pairs of points in S and the latter consists of n copies of the family \mathcal{F} , one copy for each point $s \in S$. More precisely, these families are defined as follows:

1. $\mathcal{H}_1 = \{H(s_i, s_j) \mid s_i, s_j \in S\}$, where

$$H(s_i, s_j) = \{x \in \mathcal{E}^d \mid \text{dist}(x, s_i) = \text{dist}(x, s_j)\}.$$

2. $\mathcal{H}_2 = \{F + s \mid s \in S \text{ and } F \in \mathcal{F}\}$.

Observe that $|\mathcal{H}_1| = O(n^2)$ and $|\mathcal{H}_2| = O(n)$. Let $\mathcal{A}(\mathcal{H}_1 \cup \mathcal{H}_2)$ denote the arrangement formed by the set of hyperplanes $\mathcal{H}_1 \cup \mathcal{H}_2$. This arrangement partitions the space \mathcal{E}^d into $\Theta(n^{2d})$ convex, polyhedral cells with disjoint interiors. We prove our upper bound by showing that each cell in the arrangement $\mathcal{A}(\mathcal{H}_1 \cup \mathcal{H}_2)$ gives rise to only a constant number of distinct topologies.

In particular, let C be a cell in the arrangement $\mathcal{A}(\mathcal{H}_1 \cup \mathcal{H}_2)$. We show that the set of geographic nearest neighbors $GN(x; \mathcal{B})$, and hence the set of edges $E(S \cup \{x\}; \mathcal{B})$, is invariant for all $x \in C$. Each $GN(x; \mathcal{B})$ can give rise to $O(2^{|\mathcal{B}|}) = O(1)$ distinct topologies, one per subset of $GN(x; \mathcal{B})$, see Lemma 3.1. Since there are $O(n^{2d})$ cells in $\mathcal{A}(\mathcal{H}_1 \cup \mathcal{H}_2)$, the number of distinct minimum

spanning tree topologies is also bounded by $O(n^{2d})$. We now prove the claim that the set $GN(x; \mathcal{B})$ is invariant for the cell C .

Let x and y be two arbitrary points in C , and let $B \in \mathcal{B}$ be a basis. We first observe that $G(x; B) = G(y; B)$, where $G(x; B)$ is the geographic set of x with respect to B . This follows since for $G(x; B) \neq G(y; B)$ to hold there must be a point $s \in S$ that lies in $\text{cone}(B) + x$ but not in $\text{cone}(B) + y$, which in turn would imply that there is a basis plane through s separating x and y . That is impossible since \mathcal{H}_2 includes all the basis planes, and yet x and y lie in the same cell of the arrangement $\mathcal{A}(\mathcal{H}_1 \cup \mathcal{H}_2)$. Therefore, we must have $G(x; B) = G(y; B)$.

Next, we observe that the ordering of the distances between x and the points of $G(x; B)$ is also invariant for all $x \in C$. This follows since \mathcal{H}_1 includes the bisector planes for all pairs of points. Since the basis B was picked arbitrarily, the claim holds for all $B \in \mathcal{B}$, and we have shown that the geographic nearest neighbor set $GN(x; \mathcal{B})$ is invariant for all $x \in C$.

The following upper bound theorem for the number of distinct minimum spanning tree topologies has now been established.

Theorem 3.2. *Let S be a fixed set of n points in \mathcal{E}^d . If a point x moves arbitrarily in \mathcal{E}^d , then the number of topologically distinct minimum spanning trees $MST(S, x)$ induced by it is $O(n^{2d})$.*

The upper bound in Theorem 3.2 is not tight. In the lower dimensions, for $d \leq 2$, we can prove an upper bound of $O(n^d)$. This is the subject of the next section.

3.2. A Tight Upper Bound for Dimensions $d \leq 2$

In one dimension the number of distinct topologies is $n + 1$: $MST(S, x)$ is a path that joins the points of $S \cup \{x\}$ in the left-to-right order, and there are $n + 1$ topologically distinct positions of x among the points of S .

The two-dimensional case is more interesting. Using Voronoi diagrams based on convex distance functions, we can show that there are $\Theta(n^2)$ distinct topologies of the minimum spanning tree $MST(S, x)$. We exploit the fact that this Voronoi diagram has $O(n)$ complexity in the plane, a fact that does not hold in higher dimensions.

Let $B = \{b_1, b_2\}$ be a basis in \mathcal{E}^2 , and let p be an arbitrary point. Define the distance from p to another point q as follows:

$$\sigma(p, q; B) = \begin{cases} \text{dist}(p, q) & \text{if } q - p \in \text{cone}(B), \\ +\infty & \text{otherwise.} \end{cases}$$

In other words, $\sigma(p, q; B) = \text{dist}(p, q)$ if q is a geographic neighbor of p with basis B ; otherwise the σ -distance is infinity. We easily check that σ is a convex distance function, whose unit circle is a pie-shaped figure, and that it obeys the triangle inequality—observe, however, that σ is not symmetric.

Let $VOR(S; \sigma, B)$ be the Voronoi diagram of S using the distance function σ and basis B . We need the following result of Chew and Drysdale [4] concerning the Voronoi diagrams under a convex distance function.

Lemma 3.3 [4]. *Let S be a set of n points in the plane and let d be any convex distance function for which the triangle inequality holds. The Voronoi diagram of S for distance function d is planar, has size $O(n)$, and can be computed in $O(n \log n)$ time.*

We are now ready to prove the upper bound on the number of distinct topologies.

Theorem 3.4. *Let S be a fixed set of n points in the plane. Over all choices of a variable point $x \in \mathcal{E}^2$, there are $O(n^2)$ topologically distinct minimum spanning trees $MST(S, x)$.*

Proof. Let $\mathcal{B} = \{B_1, B_2, \dots, B_k\}$, where $k = O(1)$, be a frame satisfying the condition of Lemma 3.1. Let

$$D = \bigcap_{B \in \mathcal{B}} \text{VOR}(S; \sigma, B).$$

D is the subdivision of the plane induced by the intersection of k Voronoi diagrams. It is easy to see that, in each cell C of this subdivision, the set of geographic nearest neighbors $GN(x; \mathcal{B})$ is invariant for all $x \in C$. Since $|\mathcal{B}| = O(1)$, the number of distinct topologies corresponding to each geographic nearest neighbor set is $O(1)$. The total number of distinct minimum spanning tree topologies is therefore of the same order as the number of cells in D , which is $O(n^2)$ since D is induced by the intersection of $|\mathcal{B}| = O(1)$ planar maps, each with $O(n)$ straight line edges. This completes the proof. \square

It turns out that the bound of Theorem 3.4 was also proved by Georgakopoulos and Papadimitriou [9], who considered this problem under the guise of a 1-Steiner tree problem. Their approach is similar to ours.

Our method yields in fact a slightly stronger upper bound theorem: if x varies over a k -dimensional subspace of \mathcal{E}^d , for $k \leq d$, then the number of distinct topologies of $MST(S, x)$ is $O(n^{2k})$. This follows from the observation that when restricted to a k -dimensional subspace the arrangement $\mathcal{A}(\mathcal{H}_1 \cup \mathcal{H}_2)$ has cell complexity $O(n^{2k})$. In the special case where x moves along a straight line in the two-dimensional space, the number of topologies is $O(n)$; this follows since the cell complex D (see Section 3.2) intersects a line in $O(n)$ points. We have the following corollary of Theorems 3.2 and 3.4.

Corollary 3.5. *Let S be a fixed set of n points in \mathcal{E}^d . If a point x moves arbitrarily in a k -dimensional subspace of \mathcal{E}^d , for $k \leq d$, then the number of topologically distinct minimum spanning trees $MST(S, x)$ induced by it is $O(n^{2k})$. If $d = 2$ and $k = 1$, the number of topologies is $O(n)$.*

4. Classification of Topologies

The arrangement $\mathcal{A}(\mathcal{H}_1 \cup \mathcal{H}_2)$ partitions \mathcal{E}^d into cells that have the same geographic neighbor sets. A cell may still correspond to $\Omega(2^{|\mathcal{B}|})$ distinct topologies, one for each subset of $GN(x; \mathcal{B})$. The purpose of this section is to obtain a finer subdivision of the space in which each region corresponds to the same topology.

Our starting point is the arrangement $\mathcal{A}(\mathcal{H}_1 \cup \mathcal{H}_2)$. Let C be a cell in this arrangement, and let $G(C) = GN(x; \mathcal{B})$, where x is an arbitrary point in C —observe that the geographic neighbor set is invariant over all $x \in C$. Our partitioning scheme considers each of the $2^{|G(C)|}$ subsets of $G(C)$ in turn. For a subset $V \subseteq G(C)$, let $C(V)$ denote the locus of all those points $x \in C$ for which the following holds: in the minimum spanning tree $MST(S, x)$, x is joined by an edge to each point of V , and to no other point of S . The connected components of $C(V)$ form the elementary cells of our subdivision. By repeating this procedure for all the cells of $\mathcal{A}(\mathcal{H}_1 \cup \mathcal{H}_2)$ and, within each cell, for all subsets of the geographic neighbor sets, we obtain the desired subdivision of \mathcal{E}^d . In the following, we outline how to construct $C(V)$.

Let $v_i \in V$ be a point, and let $C(v_i, V)$ denote the locus of all points $x \in C$ for which the following three conditions holds:

- (1) $(x, v) \in MST(S, x)$ for all $v \in V$.
- (2) $(x, s) \notin MST(S, x)$ for any $s \in S - V$.
- (3) $dist(x, v_i) \leq dist(x, v_j)$ for $v_j \in V$.

Clearly, $C(V) = \bigcup_{v \in V} C(v, V)$. The following procedure computes $C(v, V)$ for an arbitrary $v \in V$.

Procedure Compute-Region (C, v, V)

Input: A cell $C \in \mathcal{A}(\mathcal{H}_1 \cup \mathcal{H}_2)$, its geographic neighbor set $G(C)$, a subset $V \subseteq G(C)$, and an element $v \in V$.

Output: $C(v, V)$.

1. Compute C' , the set of points of C that are closer to v than to any other point in V . Observe that $C' = C \cap VOR(v; V)$, where $VOR(v; V)$ is the Voronoi cell of v with respect to V .
2. Let $T = MST(S) \cup \{(x, v)\}$. That is, T is the tree obtained by adding the edge (x, v) to the minimum spanning tree of S .
3. An edge (x, v_i) , where $v_i \in G(C)$ and $v_i \neq v$, forms a unique cycle with T . Let r_i denote the length of the *longest* edge in this cycle whose endpoints are both in S .
4. Set $C(v, V) = C' \cap_{v_i \in V - \{v\}} U(v_i, r_i) \cap_{v_i \in G(C) - V} \bar{U}(v_i, r_i)$, where $U(v, r)$ is the closed ball of radius r centered on v and \bar{U} is the complement of U .

end Procedure

The correctness of the procedure Compute-Region (C, v, V) follows from the observation that step 1 takes care of condition (3) and steps 2–4 ensure that conditions (1) and (2) are satisfied. The following lemma establishes its running time.

Lemma 4.1. *After $O(n)$ -time preprocessing, each call to the procedure Compute-Region can be implemented in $O(1)$ time.*

Proof. Since $|G(C)| = O(1)$, we can compute the Voronoi diagram in Step 1 and intersect a constant number of balls in Step 4 in $O(1)$ time. This leaves the operations involving the tree T , where we need to find the longest edge in a cycle created by a nontree edge. We build an auxiliary binary tree T' from T and use the fast lowest-common ancestor finding algorithm of Harel and Tarjan [15] as follows. The tree T' is built recursively from T by taking the longest edge ℓ of T , making ℓ the root of T' , and then recursively building the subtrees for the two subtrees of T that result from the deletion of ℓ . Each edge of T is associated with a unique internal node of T' , and the leaves of T' are the vertices of T . Given a nontree edge (a, b) , the longest edge in the path from a to b in T is stored at the lowest common ancestor of a and b in T' . Thus, we can preprocess T' in linear time to support lowest-common ancestor queries in constant time [15]. Since only a constant number of these operations are required in the computation of $C(v, V)$, the lemma follows. □

We repeat the procedure Compute-Region for each cell $C \in \mathcal{A}(\mathcal{H}_1 \cup \mathcal{H}_2)$ and, within each cell, for each subset of the geographic neighbor set $G(C)$. The algorithm may return a finer subdivision than what is needed. In particular, many regions may correspond to the same topology. To obtain a maximally connected cell decomposition, we may coalesce adjacent regions having the same topology in a postprocessing step. This does not affect the asymptotic time bound. The main steps of our algorithm, therefore, can be outlined as follows:

Algorithm Topology-Subdivision

```

Compute the arrangement  $\mathcal{A}(\mathcal{H}_1 \cup \mathcal{H}_2)$ .
for each cell  $C \in \mathcal{A}(\mathcal{H}_1 \cup \mathcal{H}_2)$  do
    let  $G(C)$  be the geographic nearest neighbor set for  $C$ .
    for each subset  $V \subseteq G(C)$  do
        for each element  $v_i \in V$  do
            Call Compute-Region ( $C, v_i, V$ ).

```

end Algorithm.

In dimensions $d \leq 2$, a more efficient algorithm is possible. If $d = 1$, the desired subdivision is just the Voronoi diagram of S , which has complexity $O(n)$. For $d = 2$ we use the algorithm Topology-Subdivide, but instead of starting with the arrangement $\mathcal{A}(\mathcal{H}_1 \cup \mathcal{H}_2)$, we start with an $O(n^2)$ size subdivision of the plane, used to establish the upper bound of Section 3.2.

The following theorem summarizes the main result of this section.

Theorem 4.2. *There is an $O(n^{2d})$ -time algorithm that computes a subdivision of \mathcal{E}^d into $O(n^{2d})$ regions such that, for each region R , the topology of the minimum spanning tree $MST(S, x)$ is invariant for all $x \in R$. The time complexity improves to $O(n^d)$ for $d \leq 2$.*

5. Zone of a Single Topology

Suppose we are given a spanning tree T on the set of $n + 1$ (labeled) nodes $S \cup \{x\}$, where S is a fixed set of points in \mathbb{E}^d and x is a variable point. We want to determine if there exists a point $p \in \mathbb{E}^d$ such that $MST(S, p)$ is topologically equivalent to T . The point p is a geometric realization of the node x . If T is a feasible topology, we are also interested in computing its *topology zone* $Z(T) \subseteq \mathbb{E}^d$ such that $\forall x \in Z(T)$, $MST(S, x)$ has the same topology as T .

Let $N(x) = \{v_1, v_2, \dots, v_k\}$ denote the set of neighbors of x in T ; recall that $k = O(1)$. Let F denote the forest obtained by deleting x and all its incident edges from T ; let T_i denote the component of F containing the node v_i ; and let $S_i \subseteq S$ denote the nodes of T_i . The following observation is straightforward.

Observation 5.1. $\bigcup_{1 \leq i \leq k} T_i \subset MST(S)$.

Let X denote the set of edges that connect the subtree components T_i in $MST(S)$:

$$X = MST(S) - \bigcup_{i=1}^k T_i.$$

For each component T_i , where $1 \leq i \leq k$, let r_i denote the length of the shortest edge in X that is incident with T_i . (Figure 3 illustrates this in two dimensions.) As before, let $U(v, r)$ denote the ball of radius r centered on a point v , and let $\bar{U}(v, r)$ denote the complement of this ball.

For each $u \in S - N(x)$, the edge (x, u) forms a unique cycle with T . Let $s(u)$ denote the length of the longest edge in this cycle both of whose endpoints are in S . Finally, let $VOR(v_i; S_i)$ denote the Voronoi cell of v_i in the Voronoi diagram of point set S_i .

Lemma 5.2. $Z(T) = \bigcap_{v_i \in N(x)} U(v_i, r_i) \cap \bigcap_{v_i \in N(x)} \text{VOR}(v_i; S_i) \cap \bigcap_{u \in S - N(x)} \bar{U}(u, s(u))$.




Fig. 3. The subtrees T_i and the set X . The figure on the left shows the tree T and the one on the right shows $MST(S)$. $X = \{e_1, e_2, e_3, e_4\}$; $r_1 = \min\{e_1, e_2, e_3\}$, $r_2 = e_3$, $r_3 = e_2$, $r_4 = e_4$, and $r_5 = \min\{e_1, e_4\}$.

Proof. $Z(T)$ is the locus of all points $p \in \mathcal{E}^d$ that satisfy the following two constraints:

- (i) $(p, v) \in MST(S, p)$ for all $v \in N(x)$.
- (ii) $(p, s) \notin MST(S, p)$ for any $s \in S - N(x)$.

First, consider constraint (i). An edge (p, v_i) , for $v_i \in N(x)$, appears in $MST(S, p)$ iff

$$dist(p, v_i) = \min\{dist(u, v) | u \in (S \cup \{p\}) - S_i \text{ and } v \in S_i\}.$$

Recalling that r_i is the minimum distance between S_i and $S - S_i$, this constraint is equivalent to the following:

$$p \in \bigcap_{v_i \in N(x)} U(v_i, r_i). \tag{9}$$

(Clearly, if $dist(p, v_i) > r_i$, then it is better to join T_i to the rest of the spanning tree using the edge corresponding to r_i than through p .)

Now, for the constraint (ii), consider an arbitrary point $u \in S_i, u \neq v_i$. We need to ensure that the edge (p, u) does not appear in $MST(S, p)$. Let $e(p, u)$ denote the longest edge in the (unique) cycle formed by the addition of (p, u) to the tree T . Our constraint requires that $dist(p, u)$ should be longer than the length of $e(p, u)$. There is, however, one catch: since p is a variable point, the lengths of both (p, u) and $e(p, u)$ may be variable. We, therefore, need to distinguish between the following two cases: (Case 1) both endpoints of $e(p, u)$ are in S_i , and (Case 2) $e(p, u)$ is incident to p , namely, $e(p, u) = (p, v_i)$. In Case 1 the length of $e(p, u)$ is given by $s(u)$, and so the required condition is that $s(u) < dist(p, u)$. If this condition is to be satisfied by all $u \in S - N(x)$, we get the constraint

$$p \notin \bigcup_{u \in S - N(x)} U(u, s(u)). \tag{10}$$

In Case 2 the required condition is $dist(p, u) \leq dist(p, v_i)$. Since it must be satisfied for all $u \in S_i$, our final constraint is

$$p \in \bigcap_{v_i \in N(x)} VOR(v_i; S_i). \tag{11}$$

The lemma now follows by combining the constraints (9), (10), and (11). □

Therefore, to check if T is a feasible minimum spanning tree topology, we need to determine whether a family of hyperplanes, balls, and complements of balls has a nonempty intersection. We now consider the time complexity of deciding if $Z(T) \neq \emptyset$.

The Voronoi cell $VOR(v_i; S_i)$ is the common intersection of $|S_i| - 1$ half-spaces, each bounded by the bisecting hyperplane for a pair $\{v_i, v_j\}, v_i, v_j \in S_i$. Since the S_i 's are disjoint sets, $\bigcap_{v_i \in N(x)} VOR(v_i; S_i)$ involves at most n hyperplanes. Next, we need to find the radii of the various balls. The radii r_i are easily computed in $O(1)$ time, since they only involve the edges in X and $|X| = O(1)$. Computation of the $s(u)$'s

is only slightly more involved. Suppose that $u \in S_i$. Then $s(u)$ is the length of the longest edge on the path between u and v_i in T —recall that v_i is the neighbor of x in S_i . We can easily compute $s(u)$, for all u , by a simple depth-first traversal of T , which takes linear time. We omit the easy details of this procedure.

Once we have determined the r_i 's and the $s(u)$'s, our remaining task is to check if the $O(n)$ half-spaces, balls, and complements of balls in Lemma 5.2 have a common point. A ball B with radius ρ and center (a_1, a_2, \dots, a_d) is given by the quadratic form $\sum_{i=1}^d (x_i - a_i)^2 \leq \rho^2$. The complement of the ball is obtained by replacing the “ \leq ” sign with the “ \geq ” sign. We can “linearize” our system of equations by introducing an auxiliary variable x_{d+1} , and a new constraint

$$x_{d+1} = x_1^2 + x_2^2 + \dots + x_d^2. \tag{12}$$

The (closed) interior of B is then given by

$$x_{d+1} \leq \sum_{i=1}^d 2a_i x_i - \sum_{i=1}^d a_i^2 + \rho^2. \tag{13}$$

The coefficient of x_{d+1} is zero in the linear constraints corresponding to the hyperplanes in \mathcal{E}^d ; thus a hyperplane in \mathcal{E}^d transforms to a *vertical* hyperplane in \mathcal{E}^{d+1} . By introducing x_{d+1} , we have transformed our system of constraints to a system of linear constraints, plus one convex quadratic constraint (12). The feasibility of such a system can be tested in $O(n)$ time, using an algorithm due to Dyer [6]. We, therefore, have established the following theorem.

Theorem 5.3. *Given a set of n points $S \in \mathcal{E}^d$, its minimum spanning tree $MST(S)$, and a topology tree T on $S \cup \{x\}$, where x is a variable node, we can determine in $O(n)$ time if there exists a point $p \in \mathcal{E}^d$ for which $MST(S, p)$ and T are topologically equivalent.*

Constructing $Z(T)$ requires intersecting $O(n)$ balls and hyperplanes. By the method outlined above, this is equivalent to intersecting $O(n)$ hyperplanes and a paraboloid, which can be done in $O(n \log n)$ time for the plane, and in $O(n^{\lceil (d+1)/2 \rceil})$ for $d \geq 3$ (see [7]).

6. Sensitivity of a Topology

We consider the sensitivity measure of a planar set of points S . Call a set S' an ε -perturbation of S if $|S'| = |S|$ and each point of S' is within distance ε of its corresponding point in S . The sensitivity measure $\mu(S)$ is the largest ε for which all ε -perturbations of S have the same minimum spanning tree topology as S . (We assume that S admits a unique minimum spanning tree topology; otherwise $\mu(S)$ is trivially zero.)

Let T denote the topology tree corresponding to $MST(S)$. For an edge $e \in T$, let $\mu(e)$ denote the sensitivity of e ; that is, for $\varepsilon < \mu(e)$, the minimum spanning tree of any ε -perturbation of S includes the edge e , and there is a $\mu(e)$ -perturbation of

S whose minimum spanning tree does not include the edge e . Since a topology change occurs only when a nontree edge swaps a tree edge, the following relation is straightforward:

$$\mu(S) = \min_{e \in T} \mu(e).$$

Our problem, therefore, reduces to computing $\mu(e)$ for all $e \in T$. We break the analysis of edge-sensitivity into two cases, depending on whether the edge replacing e shares an endpoint with e or not. In particular, let $\mu_1(e)$ be the sensitivity of e subject to the condition that the edge replacing e does not share an endpoint with e , and let $\mu_2(e)$ be the sensitivity of e subject to the condition that the edge replacing e shares an endpoint with e . Obviously, $\mu(e) = \min\{\mu_1(e), \mu_2(e)\}$.

Let $e = (a, b)$ be an edge of T , and let T_a and T_b denote the two components of $T - \{e\}$, where $a \in T_a$ and $b \in T_b$. Observe that

$$\text{dist}(a, b) = \min\{\text{dist}(c, d) \mid c \in T_a, d \in T_b\}.$$

The following two sections describe our algorithms for computing $\mu_1(e)$ and $\mu_2(e)$.

6.1. Computing the Sensitivity Measure μ_1

Our first lemma reduces the task of finding replacement edges to that of finding certain closest pairs. The next two lemmas produce in $O(n \log n)$ time an $O(n)$ -size set of edges that contains all the replacement edges (for the sensitivity measure μ_1). Our final lemma shows how to find the replacement edge for each $e \in T$.

Lemma 6.1. *If $\mu(S) = \mu_1(a, b)$, then*

$$\mu_1(a, b) = \min_{\substack{c \in T_a - \{a\} \\ d \in T_b - \{b\}}} \frac{1}{4}(\text{dist}(c, d) - \text{dist}(a, b)).$$

Proof. Consider a pair c, d , where $c \in T_a - \{a\}$ and $d \in T_b - \{b\}$. An ε -perturbation can decrease the distance between c and d by at most 2ε , and increase the distance between a and b by the same amount. Thus, (c, d) can replace (a, b) only if

$$\text{dist}(c, d) - 2\varepsilon < \text{dist}(a, b) + 2\varepsilon.$$

The sensitivity $\mu_1(e)$ equals the minimum value of ε satisfying the above inequality, over all choices of c, d . This completes the proof. \square

Thus, $\mu_1(e)$ can be calculated by finding the *closest pair* of points c, d , with $c \in T_a - \{a\}$ and $d \in T_b - \{b\}$. We call the corresponding edge (c, d) the *replacement edge* for (a, b) (for the sensitivity measure μ_1). The *diametral circle* of (c, d) is the circle whose diameter is $\text{dist}(c, d)$ and whose center is on the midpoint of (c, d) .

The following lemma demonstrates a sparse set of edge containing all the replacement edges.

Lemma 6.2. *If (c, d) is a replacement edge for (a, b) , then the diametral circle of (c, d) contains no point from the set $S - \{a, b\}$.*

Proof. Suppose there exists a point $z \in S, z \notin \{a, b\}$, that lies in the circle. If $z \in T_a$, then $\text{dist}(z, d) < \text{dist}(c, d)$, and if $z \in T_b$, then $\text{dist}(c, z) < \text{dist}(c, d)$. In either case, $\text{dist}(c, d) > \min\{\text{dist}(u, v) \mid u \in T_a - \{a\}, v \in T_b - \{b\}\}$ which contradicts the assumption that (c, d) minimizes the distance between $T_a - \{a\}$ and $T_b - \{b\}$. This completes the proof. \square

Lemma 6.2, therefore, ensures that all replacement edges for the sensitivity measure μ_1 correspond to pairs of points that have a nonempty region in the fourth-order Voronoi diagram of S . The fourth-order Voronoi diagram of S can be computed in $O(n \log n)$ time and $O(n)$ space. Each cell C of this diagram corresponds to a quadruple of points $\{s_1, s_2, s_3, s_4\}$ such that, for any point $x \in C$, the four nearest neighbors of x are s_1, s_2, s_3 , and s_4 . (An interested reader may consult the book by Edelsbrunner [7] for more details on Voronoi diagrams.) The linear space complexity of the order-4 Voronoi diagram implies that there are $O(n)$ such quadruples.

Let $R(S)$ denote the set of edges (p, q) corresponding to the pairs p, q that belong to a quadruple having a nonempty cell in the fourth-order Voronoi diagram of S . The preceding discussion has established the following lemma.

Lemma 6.3. *In $O(n \log n)$ time and $O(n)$ space, we can find a set of edges $R(S)$ such that the replacement edges for all $e \in T$ are in $R(S)$.*

With these preliminaries in place, we can now describe our algorithm for finding all the replacements edges. We represent the minimum spanning tree T using the “linking and cutting tree” data structure of Sleater and Tarjan [19]. Each edge of the tree has a cost associated with it, which is initially set to be its length; that is, $\text{cost}(e) = \text{dist}(a, b)$, where $e = (a, b)$. At some point during the course of the algorithm, we reset the cost of an edge to be zero, but we do it only once per edge.

Our algorithm processes the edges of $R(S)$ in the nondecreasing order of weight (length). While processing the edge $f = (c, d)$, the following computational steps are carried out:

- Step 1. We find a maximum-cost *internal* edge e on the path between c and d ; an internal edge is one that does not have c or d as an endpoint.
- Step 2. If $\text{cost}(e) > 0$, we record f as the replacement edge of e , and set $\text{cost}(e) = 0$. Go back to Step 1.
- Step 3. If $\text{cost}(e) = 0$, the processing of f is complete, and we move on to the next edge of $R(S)$.

We first establish the correctness of the preceding algorithm. Consider an edge $e \in T$, where $e = (a, b)$. By Lemma 6.3, the replacement edge for e is the minimum-

cost edge in $R(S)$ with one endpoint in each of $T_a - \{a\}$ and $T_b - \{b\}$. An edge $f = (c, d)$ meets this condition only if e is an internal edge in the path from c to d . Since we process the edges of $R(S)$ in the order of nondecreasing weight, the first edge to meet the condition is the replacement edge for e . Having found the replacement edge for e , we reset its cost to zero, taking it out of the game; this obviously does not prevent other edges from correctly finding their replacement edges. This proves that the algorithm correctly identifies the replacement edges for all $e \in T$.

It remains to show that the algorithm can be implemented in $O(n \log n)$ time using $O(n)$ space. The linking and cutting data structure of Sleator and Tarjan [19] allows the two operations used in Steps 1 and 2 to be performed at the cost of $O(\log n)$ each. That is, given a pair u and v , the data structure can find a maximum-cost edge on the tree path from u to v , and given an edge e and a real number x , the data structure can execute $cost(e) := x$, each in the worst-case time $O(\log n)$. Each execution of Steps 1 and 2 either sets the cost of an edge e to zero (and finds the replacement edge for e) or finishes the processing of an edge of $R(S)$. Since there are altogether $O(n)$ edges, the total running time of the algorithm is $O(n \log n)$. The linking and cutting tree data structure uses $O(n)$ space. We have established the following lemma.

Lemma 6.4. *There is an $O(n \log n)$ -time and $O(n)$ -space algorithm for computing $\mu_1(e)$, for all $e \in T$.*

6.2. Computing the Sensitivity Measure μ_2

Consider the case where the edge (a, b) is replaced by (a, c) for $c \in T_b - \{b\}$. Let $\mu_2(a, b, c)$ denote the largest ε such that $dist(a', b') \leq dist(a', c')$ holds for any ε -perturbation $\{a', b', c'\}$ of the set $\{a, b, c\}$. Similarly, $\mu_2(b, a, c)$ is the largest ε such that $dist(b', a') \leq dist(b', c')$. Clearly,

$$\mu_2(e) = \min_{c \notin \{a, b\}} \{ \mu_2(a, b, c), \mu_2(b, a, c) \}.$$

The following lemma gives a geometric interpretation of $\mu_2(a, b, c)$.

Consider the locus of all points x in the plane satisfying the equation

$$dist(c, x) - dist(b, x) = 2\varepsilon. \tag{14}$$

This locus is a hyperbola, denoted $h(b, c, \varepsilon)$, whose foci are b and c . Actually, since we are using the signed distance, the locus is the one arm of the hyperbola whose vertex is closer to b .

Lemma 6.5. *$\mu_2(a, b, c)$ equals the largest ε for which $U(a, \varepsilon)$, the ε -disk centered on a , lies completely on the same side of $h(b, c, \varepsilon)$ as b . The value of $\mu_2(a, b, c)$ can be computed in $O(1)$ time.*

Proof. Consider an ε -perturbation $\{a', b', c'\}$ of $\{a, b, c\}$ for some $\varepsilon \geq 0$. Then, $\text{dist}(a', b') \leq \text{dist}(a, b) + \varepsilon$ and $\text{dist}(a', c') \geq \text{dist}(a, c) - \varepsilon$. We want to determine the minimum $\varepsilon > 0$ subject to the following constraints:

1. $\text{dist}(a', c) - \varepsilon < \text{dist}(a', b) + \varepsilon$.
2. $\text{dist}(a, a') \leq \varepsilon$.

Since the locus of all points x satisfying the equation $\text{dist}(c, x) - \text{dist}(b, x) = 2\varepsilon$ is the hyperbola $h(b, c, \varepsilon)$, our problem is to determine the largest ε such that the ε -disk centered on a lies completely on the same side of $h(b, c, \varepsilon)$ as b . We assume that this one-variable, one-constraint optimization problem can be solved in some constant time. □

Lemma 6.5 describes how to calculate the sensitivity of $e = (a, b)$ with respect to a fixed point c in constant time. In order to compute $\mu_2(e)$, however, we may have to check against $\Omega(n)$ choices of c , which leads to an overall quadratic algorithm. In the remainder of this section we improve this time complexity to $O(n \log n)$ by showing that we only need to check $O(n)$ triples (a, b, c) .

Let $\delta(S)$ denote the minimum distance between any two points of S . Then the following simple observation gives an absolute upper bound on $\mu(S)$.

Observation 6.6. $\mu(S) \leq \frac{1}{2}\delta(S)$.

Call (a, b, c) a *candidate triple* for (a, b) if

- (1) $(a, b) \in T$,
- (2) $c \in T_b$, and
- (3) $\text{dist}(a, c) \leq \text{dist}(a, b) + 2\delta(S)$.

Lemma 6.5 and Observation 6.6 together imply that if we know all the candidate triples for (a, b) , then $\mu_2(a, b)$ can be computed in time proportional to the number of triples. Thus, the sensitivity measure $\mu_2(S)$ can be computed in time proportional to the total number of candidate triples (a, b, c) , where (a, b) varies over all the edges of the minimum spanning tree T . In the following we show that there are only a linear number of such triples and they can be found in $O(n \log n)$ time using $O(n)$ space.

Consider again the computation of $\mu_2(a, b, c)$. We claim that if (a, b, c) is a candidate triple, then a must be among a constant number of *geographic* neighbors of c . We consider a *wedge* w defined by two half-lines originating at c and making an angle of $\pi/6$. We can cover the entire plane with 12 such wedges w_1, \dots, w_{12} placed at c . Let $S(w_i, c)$ denote the set of points of S lying in the wedge w_i . Without loss of generality, assume that $a \in S(w_1, c)$. We claim that a is among a fixed number of nearest neighbors of c in the wedge w_1 .

To prove this claim, first observe that since $d(a, c) \leq d(a, b) + 2\delta(S)$, c lies within the circle of C_1 of radius $d(a, b) + 2\delta(S)$ centered at a . Furthermore, since (a, b) is an edge of the minimum spanning tree, no point of T_b lies inside the circle C_2 of




Fig. 4. Computing μ_2 .

radius $d(a, b)$ centered at a . Thus, c and any other point of T_b lies in the common intersection of the wedge w_1 and the annulus between circles C_1 and C_2 . This intersection consists of two pieces, I_1 and I_2 , as shown in Fig. 4. We are only interested in I_1 , since the points in I_2 are even farther from c than $d(c, a)$. Since no two points are closer than $\delta(S)$, only a constant number of points of T_b can be in I_1 .

However, what about points of T_a ? We show that no point $d \in T_a$ can satisfy $d(c, d) \leq d(c, a)$. Indeed, if $d(c, d)$ were smaller than $d(c, a)$, then $d(c, d) \leq d(c, a) \leq d(a, b) + 2\delta(S)$, which is manifestly false since Lemma 6.1 guarantees that $d(c, d) \geq d(a, b) + 4\mu_1$. Thus, all points of T_a in $S(w_1, c)$ are farther from c than $d(c, a)$.

We have, thus, shown that a is among the $O(1)$ nearest neighbors of c in the wedge w_1 . Since we have 12 wedges for c , and in each wedge we need to remember $O(1)$ nearest neighbors, c contributes $O(1)$ point-pairs (c, a) . Since a has at most six minimum spanning tree edges (a, b) incident to it, the total number of triples (a, b, c) is $O(n)$. Finally, we can compute all the test triples in $O(n \log n)$ time, using Voronoi diagrams [4].

The following theorem summarizes the main result of this section.

Theorem 6.7. *Given a set S of n points in the plane, we can compute its sensitivity measure $\mu(S)$ in $O(n \log n)$ time using $O(n)$ space.*

7. Embedding a Topology

Let T be a free tree on n nodes. Does there exist a set of n points $S \subset \mathcal{E}^2$ such that $MST(S)$ has the same topology as T ? We begin with some simple observations on minimum spanning trees of points in the plane.

Observation 7.1. *Any minimum spanning tree of a finite set of points $S \subset \mathcal{E}^2$ has maximum node degree six.*

Lemma 7.2. *Every finite set of points S in the plane admits a minimum spanning tree whose maximum node degree is five.*

Proof. Suppose $MST(S)$ has at least one node u of degree six. Let v_0, \dots, v_5 be the neighbors of u , in counterclockwise order. We prove that the spanning tree obtained by swapping the edge (u, v_0) with (v_0, v_1) is also a minimum spanning tree, but it has one fewer degree-six node. By repeating this procedure at each degree-six node, we obtain a minimum spanning tree with maximum node degree five. To establish the validity of our edge-swap, observe that the six edges incident to u are all equal in length and that the angles $\angle v_i u v_{i+1}$ are all equal to $\pi/3$, for $i = 0, 1, \dots, 5$. Therefore, swapping (u, v_0) with (v_0, v_1) does not change the cost of the spanning tree and, furthermore, the resulting tree is a valid spanning tree since (v_0, v_1) is not an edge of the original spanning tree $MST(S)$; otherwise u, v_0 , and v_1 would form a cycle. We have reduced the degree of u to five, but it remains to show that the degree of v_1 does not increase to six. Suppose, for the sake of a contradiction, that the addition of (v_0, v_1) increases degree of v_1 to six. Again, by the preceding argument, the six neighbors of v_1 must form the vertices of a regular hexagon, with center on v_1 . This implies that v_2 must also be a neighbor of v_1 , which is impossible since $(v_1, v_2) \notin MST(S)$. This completes the proof that our edge-swap reduces the number of degree-six nodes by one. \square

We prove in the following that any tree with maximum node-degree five can be embedded as a minimum spanning tree of some set of points in the plane. Specifically, we give a scheme for embedding a complete degree-five tree T , rooted at node r . (This clearly suffices for realizing any tree with maximum node degree five.) We embed the nodes of T in the *breadth-first* order; nodes at distance i from r are called *level- i* nodes. We place the root node r at the origin, and embed its five neighbors evenly spaced on a unit circle, centered on r . To describe the embedding of higher-level nodes, we introduce the following notation.

Let θ_i denote the minimum angle between two consecutive neighbors of an i th level node minus $\pi/3$, and let ℓ_i denote the length of edges at level i . Thus, for $i = 0$, we have

$$\theta_0 = \frac{\pi}{15}, \quad \ell_0 = 1.$$




Fig. 5. Embedding of v 's neighbors; not drawn to scale. $\alpha = \pi/3 + \theta_i$ and $\beta = \pi/2 - \theta_{i-1}$.

The values of θ_i and ℓ_i , for $i \geq 1$, are determined recursively as follows:

$$\theta_{i+1} = \frac{\theta}{15}, \quad \ell_{i+1} = \frac{\ell_i \theta_i}{3}. \tag{15}$$

The closed form expressions for ℓ_i and θ_i are the following.

$$\theta_i = \frac{\pi}{15^{i+1}}, \quad \ell_i = \frac{(\pi/3)^i}{15^{i(i+1)/2}}. \tag{16}$$

Given a vertex v at level i , our embedding scheme determines the position of v 's neighbors as follows. See Fig. 5 for an illustration of the construction.

Embedding Rule. Let $\{v_0, v_1, \dots, v_4\}$ be the neighbors of v in counterclockwise order around v , and assume that v_0 is the parent of v . By induction, $dist(v, v_0) = \ell_{i-1}$. We place v_1, v_2, v_3, v_4 on a circle with center v and radius ℓ_i such that:

1. $\angle v_1 v v_0 = \angle v_0 v v_4 = \pi/2 - \theta_{i-1}/10$.
2. $\angle v_j v v_{j+1} = \pi/3 + \theta_i$ for $j = 1, 2, 3$.

Let $S(T)$ denote the set of points obtained using this embedding rule. Our next lemma shows that, indeed, $S(T)$ is an *embedding* of T , that is, the minimum spanning tree of $S(T)$ has the same topology as T . In our analysis we make extensive use of the following bounds, which follow from the Taylor series expansions of $\sin \theta$ and $\cos \theta$, for $0 < \theta < 1$:

$$\theta - \frac{\theta^3}{3!} < \sin(\theta) < \theta, \tag{17}$$

$$1 - \frac{\theta^2}{2!} < \cos(\theta) < 1, \tag{18}$$

$$1 + \frac{\theta}{2} - \frac{\theta^2}{8} < (1 + \theta)^{1/2} < 1 + \frac{\theta}{2}. \tag{19}$$




Fig. 6. Case 1 of Lemma 7.3, $\alpha = \pi/2 - \theta_i$. The subtree rooted at w_2 is embedded within the disk of radius $2\ell_{i+2}$ centered on $g(w_2)$.

Lemma 7.3. *The minimum spanning tree of $S(T)$ is topologically equivalent to T .*

Proof. Given a vertex $x \in T$, let $g(v) \in S(T)$ denote the embedding or the geometric realization of v . Clearly, to establish the lemma we only need to show that, for any nontree edge $(u, v) \notin T$, $dist(g(u), g(v))$ is greater than or equal to the length of the longest edge in the cycle induced by (u, v) . Let $w = lca(u, v)$ be the lowest common ancestor of u and v in T , and let i denote the level of w in the breadth-first ordering used by our embedding scheme. It is easy to see that in the cycle induced by (u, v) the longest edge is incident to $g(w)$ and its length is ℓ_i . Thus, we only need to show that $dist(g(u), g(v)) \geq \ell_i$. We break up our analysis into two cases:

- (1) $w = u$ or $w = v$.
- (2) $w \neq u$ and $w \neq v$.

Case 1: $w = u$ or $w = v$. See Fig. 6. We prove that if z is any descendant of w , then $dist(g(w), g(z)) \geq \ell_i$. Let $w = w_0, w_1, \dots, w_k = z$ be the path from w to z , where we may assume that $k \geq 2$. (Otherwise, $dist(w, z) = \ell_i$.) First, consider the case where, in the cyclic order around w_1 , the vertices w_0 and w_2 are nonadjacent. In this case, we easily verify that

$$dist(w_0, w_2) > \ell_i + 2\ell_{i+2}.$$

This suffices for our purpose since $dist(g(w_2), g(z)) < 2\ell_{i+2}$ for any descendant z of w_2 . (The inequality follows from the descending geometric series formed by the edge-lengths, see (15).) We, therefore, may assume that w_0 and w_2 are cyclic neighbors around w_1 . To simplify the calculations, assume that w_0 is placed at the origin $(0, 0)$, w_1 is placed on the x -axis at $(\ell_i, 0)$, and that w_2 lies above the x -axis. Recalling that $\angle w_0w_1w_2 = \pi/2 - \theta_i/10$, the x - and y -coordinates of $g(w_2)$ can be written as follows:

$$x(w_2) = \ell_i - \ell_{i+1} \cos\left(\frac{\pi}{2} - \frac{\theta_i}{10}\right), \quad y(w_2) = \ell_{i+1} \sin\left(\frac{\pi}{2} - \frac{\theta_i}{10}\right). \tag{20}$$

The Taylor series bounds for $\sin \theta$ and $\cos \theta$, (17) and (18), give the following

lower bounds:

$$x(w_2) > \ell_i - \frac{\ell_{i+1}\theta_i}{10}, \quad y(w_2) > \ell_{i+1}\left(1 - \frac{\theta_i^2}{200}\right). \tag{21}$$

Using (21), we can bound the distance between $g(w_0)$ and $g(w_2)$ from below:

$$\text{dist}(g(w_0), g(w_2)) > \sqrt{\left(\ell_i - \frac{\ell_{i+1}\theta_i}{10}\right)^2 + \left(\ell_{i+1}\left(1 - \frac{\theta_i^2}{200}\right)\right)^2}. \tag{22}$$

By substituting $\ell_{i+1} = \ell_i\theta_i/3$, inequality (22) simplifies to

$$\text{dist}(g(w_0), g(w_2)) > \ell_i\left(1 + \frac{2}{45}\theta_i^2\right)^{1/2}. \tag{23}$$

Using the Taylor series expansion of $(1 + x)^{1/2}$ in (23), we obtain

$$\text{dist}(g(w_0), g(w_2)) > \ell_i\left(1 + \frac{1}{45}\theta_i^2 - \frac{1}{8}\left(\frac{2}{45}\right)^2\theta_i^4\right). \tag{24}$$

In particular,

$$\text{dist}(g(w_0), g(w_2)) > \ell_i\left(1 + \frac{\theta_i^2}{50}\right). \tag{25}$$

In order to ensure that all descendants of w_2 are at least ℓ_i distance away from $g(w_0)$, we need to show that no such vertex is within distance $\frac{1}{50}\ell_i\theta_i^2$ of $g(w_2)$. The geometrically decreasing edge-lengths guarantee that the entire subtree rooted at w_2 is embedded within a disk of radius $2\ell_{i+2}$ centered on $g(w_2)$. Thus, the required inequality is

$$\frac{\ell_i\theta_i^2}{50} > 2\ell_{i+2}, \tag{26}$$

which is easily verified by plugging in the value of ℓ_{i+2} from (15). We have shown that $\text{dist}(g(w), g(z)) > \ell_i$, for any descendant z of w , and this completes the proof of Case 1.

Case 2: $w \neq u$ and $w \neq v$. See Fig. 7. In this case we estimate the distance between $g(w_1)$ and $g(w_2)$, where w_1 and w_2 are children of w and u (resp. v) is a descendant of w_1 (resp. w_2). Assume that w is placed at the origin and w_1 at $(\ell_i, 0)$. By construction, $\angle w_1ww_2 \geq \pi/3 + \theta_i$. The x - and y -coordinates of $g(w_2)$ are given by

$$x(w_2) < \ell_i \cos\left(\frac{\pi}{3} + \theta_i\right), \quad y(w_2) > \ell_i \sin\left(\frac{\pi}{3} + \theta_i\right). \tag{27}$$




Fig. 7. Case 2 of Lemma 7.3, $\alpha \geq \pi/2 - \theta_i$. The subtree rooted at w_j , for $j = 1, 2$, is embedded within the disk of radius $\ell_{i+1} + 2\ell_{i+2}$ centered on $g(w_j)$.

By using the Taylor series bounds for $\sin \theta$ and $\cos \theta$, (17) and (18), we obtain the following:

$$\begin{aligned}
 x(w_2) &< \ell_i \left(\frac{1}{2} - \frac{\sqrt{3}\theta_i}{2} + \frac{\sqrt{3}\theta_i^3}{12} \right), \\
 y(w_2) &> \ell_i \left(\frac{\sqrt{3}}{2} + \frac{\theta_i}{2} - \frac{\sqrt{3}\theta_i^2}{4} - \frac{\theta_i^3}{12} \right).
 \end{aligned}$$

We, therefore, have the following estimate of the distance $dist(w_1, w_2)$:

$$\begin{aligned}
 dist(g(w_1), g(w_2)) &> \ell_i \sqrt{\left(\frac{1}{2} + \frac{\sqrt{3}\theta_i}{2} - \frac{\sqrt{3}\theta_i^3}{12} \right)^2 + \left(\frac{\sqrt{3}}{2} + \frac{\theta_i}{2} - \frac{\sqrt{3}\theta_i^2}{4} - \frac{\theta_i^3}{12} \right)^2} \\
 &> \ell_i \sqrt{1 + \sqrt{3}\theta_i + \frac{\theta_i^2}{4} - \frac{5\sqrt{3}\theta_i^3}{12} - O(\theta_i^4)} \\
 &> \ell_i \sqrt{1 + \sqrt{3}\theta_i + \frac{\theta_i^2}{12}} \\
 &> \ell_i \left(1 + \frac{\sqrt{3}\theta_i}{2} + \frac{\theta_i^2}{24} - \frac{1}{8} \left(3\theta_i^2 + \frac{\theta_i^3}{2\sqrt{3}} + \frac{\theta_i^4}{144} \right) \right) \quad (\text{using (19)}) \\
 &> \ell_i \left(1 + \left(\frac{\sqrt{3}}{2} - \frac{1}{8} \right) \theta_i \right), \tag{28}
 \end{aligned}$$

where we have used $\frac{1}{8}\theta_i$ to absorb the higher-order negative terms. Finally, observe that all descendants of w_j , for $j = 1, 2$, are embedded within the disk of radius $\ell_{i+1} + 2\ell_{i+2}$ centered on $g(w_j)$. Thus, to ensure that no descendant of w_1 is closer than ℓ_i to any descendant of w_2 , it suffices to show that

$$\left(\frac{\sqrt{3}}{2} - \frac{1}{8}\right)\ell_i\theta_i > 2\ell_{i+1} + 4\ell_{i+2}. \tag{29}$$

This is easily verified by plugging in the values of ℓ_{i+1} and ℓ_{i+2} from (15). This completes the proof of Case 2, and the lemma. \square

We summarize the main result of this section in the following theorem.

Theorem 7.4. *Let T be a tree on n nodes, with maximum node degree five. There exists a set of n points $S \subset \mathcal{E}^2$ such that $MST(S)$ is topologically equivalent to T .*

Remark. The maximum node degree five is important, since there do exist trees with degree-six nodes that cannot be realized as minimum spanning tree topologies. For instance, a tree containing a degree-six node adjacent to a degree-five node cannot be realized.

Exclusion of degree-six nodes does not seem like a serious drawback, since Lemma 7.2 guarantees that every finite set of points in the plane admits a minimum spanning tree with maximum node degree five. Thus, Theorem 7.4 implies that the only obstruction that prevents a tree from being realized as a minimum spanning tree topology in the plane is the trivial one: maximum node degree.

Thus, in the worst-case, our embedding scheme requires a $\Theta(2^{n^2}) \times \Theta(2^{n^2})$ size grid to embed an n -node tree.

8. Conclusions and Open Problems

We have studied a new class of problems in this paper, namely, the occurrence of state transitions in a geometric structure due to an arbitrary motion of one or more points. Since geometric objects often model physical entities that can move in space, the problems considered in our paper are fairly natural ones. We established combinatorial bounds on the number of topologically distinct minimum spanning trees induced on a fixed set of points by the motion of a single point. In particular, we showed that if a point x moves arbitrarily in a k -dimensional subspace of \mathcal{E}^d , $k \leq d$, then it can induce at least $\Omega(n^k)$ and at most $O(n^{2k})$ distinct minimum spanning tree topologies on a fixed set of n points. In the lower dimensions, $d \leq 2$, we proved an asymptotically tight bound of $\Theta(n^k)$ for $k \leq d$. Closing the gap between the lower and upper bounds for dimensions $d \geq 3$ remains an open problem. We conjecture that our lower bound is the right answer.

A second open problem concerns the structure of topology zones. Our Lemma 5.2 shows that a topology zone $Z(T)$ is the common intersection of $O(n)$ balls,

hyperplanes, and complement of balls, but we do not know much about its finer geometry. The questions whether $Z(T)$ is connected or hole-free remain open at this point.

Our algorithm for computing the sensitivity measure in the plane runs in time $O(n \log n)$. Our method for computing the sensitivity measure extends quite naturally to higher dimensions, although the resource bounds degrade. A straightforward implementation runs in $O(n^2 \log n)$ time, however, slightly subquadratic bounds are possible by using more sophisticated algorithms for bichromatic nearest neighbors and spherical range searching.

The embedding scheme described in Section 7 requires a grid size depending on the height of the tree T . If T has height k , then our algorithm embeds T on a grid of size $\Theta(2^{k^2}) \times \Theta(2^{k^2})$. Our conjecture is that this bound cannot be improved below $c^n \times c^n$, for some constant $c > 0$. We also have not addressed the question of embedding topologies in higher-dimensional spaces.

Finally, many of the questions addressed in this paper can be asked also about other geometric structures, such as geometric matchings or Euclidean traveling salesman tours, among others. We can show that, unlike the minimum spanning tree, the number of topologically distinct minimum-weight matchings or optimal TSP tours is independent of the dimension of the underlying space. Specifically, the number of distinct minimum-weight matchings that can arise if one of the $2n$ points moves arbitrarily through the d -dimension Euclidean space is $O(n)$, and the number of distinct optimal TSP tours is $O(n^2)$. In fact, for the matching problem in two dimensions, we can partition the plane in $O(n)$ *star-shaped* regions so that, for all positions of the moving point x in a region, the optimal matching is the same.

Acknowledgment

We wish to thank two anonymous referees for their careful reading of the manuscript and many helpful comments.

References

1. T. Asano, B. Bhattacharya, M. Keil, and F. F. Yao. Clustering algorithms based on minimum and maximum spanning trees. *Proceedings of the 4th Annual Symposium on Computational Geometry*, 1988, pp. 252–257.
2. M. J. Atallah. Some dynamic computational geometry problems. *Journal of Computational and Applied Mathematics*, Vol. 11, 1985, pp. 1117–1181.
3. B. Chazelle, R. Cole, F. P. Preparata, and C. Yap. New upper bounds for neighbor searching. *Information and Control*, Vol. 68, 1986, pp. 105–124.
4. L. P. Chew and R. L. Drysdale. Voronoi diagrams based on convex distance functions. *Proceedings of the 1st ACM Symposium on Computational Geometry*, 1985, pp. 235–244.
5. W. Cunningham. Optimal attack and reinforcement of a network. *Journal of the Association for Computing Machinery*, Vol. 32, No. 3, 1985, pp. 549–561.
6. M. E. Dyer. On a multidimensional search technique and its application to the Euclidean one-center problem. *SIAM Journal on Computing*, Vol. 15, 1986, pp. 725–738.
7. H. Edelsbrunner. *Algorithms in Combinatorial Geometry*. Springer-Verlag, Heidelberg, 1987.

8. G. Gallo, M. D. Grigoriadis, and R. E. Tarjan. A fast parametric maximum flow algorithm. *SIAM Journal on Computing*, Vol. 18, 1989, pp. 30–55.
9. G. Georgakopoulos and C. H. Papadimitriou. The 1-Steiner tree problem. *Journal of Algorithms*, Vol. 8, 1987, pp. 122–130.
10. D. Gusfield. Bounds for the parametric spanning tree problem. *Proceedings of the Humboldt Conference on Graph Theory, Combinatorics and Computing*, 1979, pp. 173–183.
11. D. Gusfield. Sensitivity analysis for combinatorial optimization. Ph.D. thesis, University of California, Berkeley, 1980.
12. D. Gusfield. Parametric combinatorial optimization and a problem of program module distribution. *Journal of the Association for Computing Machinery*, Vol. 30, No. 3, 1983, pp. 551–563.
13. D. Gusfield and R. Irving. Parametric stable marriage and minimum cuts. *Information Processing Letters*, Vol. 30, 1989, pp. 255–259.
14. D. Gusfield and C. Martel. A fast algorithm for the generalized parametric minimum cut problem and applications. Technical Report CSE-89-21, University of California, Davis, 1989.
15. D. Harel and R. E. Tarjan. Fast algorithms for finding nearest common ancestors. *SIAM Journal on Computing*, Vol. 13, No. 2, 1984, pp. 338–355.
16. C. Monma, M. Paterson, S. Suri, and F. Yao. Computing Euclidean maximum spanning trees. *Algorithmica*, Vol. 5, 1990, pp. 407–419.
17. J. C. Picard and M. Queyranne. Selected applications of minimum cuts in a network. *INFOR—Canadian Journal of Operations Research and Information Processing*, Vol. 20, 1982, pp. 394–422.
18. J. C. Picard and H. Ratliff. A cut approach to the rectilinear distance facility location problem. *Operations Research*, Vol. 26, 1978, pp. 422–433.
19. D. D. Sleater and R. E. Tarjan. A data structure for dynamic trees. *Journal of Computer and Systems Sciences*, Vol. 26, 1983, pp. 362–391.
20. A. C. C. Yao. On constructing minimum spanning trees in k -dimensional space and related problems. *SIAM Journal on Computing*, Vol. 11, No. 4, 1982, pp. 721–736.

Received April 10, 1991.