Guarding a Terrain by Two Watchtowers*

Pankaj K. Agarwal† Sergey Bereg‡ Ovidiu Daescu‡ Haim Kaplan§
Simeon Ntafos‡ Binhai Zhu¶

ABSTRACT

Given a polyhedral terrain T with n vertices, the *two-watchtower* problem for T calls for finding two vertical segments, called *watchtowers*, of smallest common height, whose bottom endpoints (bases) lie on T, and whose top endpoints *guard* T, in the sense that each point on T is visible from at least one of them. In this paper we present the following results for the two-watchtower problem in $\mathbb{R}^2$ and $\mathbb{R}^3$: (1) We show that the *discrete* two-watchtowers problem in $\mathbb{R}^2$, where the bases are constrained to lie at vertices of T, can be solved in $O(n^2 \log^4 n)$ time, significantly improving previous solutions. The algorithm works, without increasing its asymptotic running time, even if one of the towers is allowed to be placed anywhere on T. (2) We show that the *continuous* two-watchtower problem in $\mathbb{R}^2$, where the bases can lie anywhere on T, can be solved in $O(n^3 \alpha(n) \log^3 n)$ time, again significantly improving previous results. (3) Still in $\mathbb{R}^2$, we show that the continuous version of the problem of guarding a finite set $P \subset T$ of m points by two watchtowers of smallest height can be solved in $O(mn \log^4 n)$ time. (4) The discrete version of the two-watchtower problem in $\mathbb{R}^3$ can be solved in $O(n^{11/3} \operatorname{polylog}(n))$ time; this is the first nontrivial result for this problem in $\mathbb{R}^3$.

†Department of Computer Science, Duke University, Box 90129, Durham NC 27708-0129, USA. E-mail: pankaj@cs.duke.edu

‡Department of Computer Science, University of Texas at Dallas, Box 830688, Richardson, TX 75083, USA. E-mail: {besp,daescu,simeon}@utdallas.edu

§School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel. E-mail: haimk@post.tau.ac.il

¶Department of Computer Science, Montana State University, Bozeman, MT 59717-3880, USA. E-mail: bhz@cs.montana.edu

*Research by Pankaj K. Agarwal is supported by NSF under grants CCR-00-86013, EIA-01-31905, CCR-02-04118, and DEB-04-25465, by ARO grants W911NF-04-1-0278 and DAAD19-03-1-0352, and by a grant from the U.S.–Israel Binational Science Foundation. Research by Haim Kaplan is supported in part by grant 548/00 from the Israel Science Foundation (ISF). Research by Binhai Zhu is supported by NSF under grant DMS-0138065. Research by Ovidiu Daescu is supported by NSF under grant CCF-0430366.

SCG'05, June 6–8, 2005, Pisa, Italy.

General Terms

Algorithms, Theory

Categories and Subject Descriptors

F.2.2 [**Nonnumerical Algorithms**]

Keywords

Computational geometry, visibility algorithms, terrain guarding, art gallery, parametric search.

1. INTRODUCTION

A polyhedral terrain in $\mathbb{R}^d$ is the graph of a continuous, piecewise-linear $(d-1)$-variate function—it is an x-monotone polygonal chain in $\mathbb{R}^2$ and an xy-monotone polyhedral surface in $\mathbb{R}^3$. In this paper we study the *two-watchtower* problem for polyhedral terrains in $\mathbb{R}^2$ and $\mathbb{R}^3$, defined as follows: Given a polyhedral terrain T with n edges, a *watchtower* is a vertical line segment whose bottom endpoint (*base*) lies on T. A point $w \in T$ is *seen* (or *guarded*) by a watchtower τ if the top endpoint of τ sees w, that is, the segment connecting it to w lies fully above T. The problem is to find the smallest height $h > 0$ for which there exist two points $u, v \in T$ such that the watchtowers of height h erected at u, v guard together the entire terrain T (i.e., every point of T is visible from the top of at least one of the watchtowers).

Two versions of the problem have been studied in the literature. In the *discrete* version, the bases u and v are restricted to be vertices of T (or, for that matter, could belong to any prespecified finite pointset). In the *continuous* version, u and v can be located anywhere on T. Another variant is when not all of T needs to be guarded. Instead, we specify a finite set P of m "critical" points on T, and the goal is to find two watchtowers of minimum common height that together guard P (i.e., every point of P is visible from at least one of the towers). Again, we may consider the discrete or the continuous versions of this problem. Of course, similar problems can be stated for any number of watchtowers, and one can also consider variants, such as the one where the height h of the watchtowers is specified, and one wishes to find the *smallest* number of watchtowers of that height that collectively guard T.

Previous results. Guarding a terrain by watchtowers is a special case of the general class of visibility problems in two and three dimensions, known as *art gallery problems*, which have been heavily studied for more than two decades. These problems have numerous applications in surveillance, navigation, computer vision, modeling and graphics, GIS, and many more. See [20] for a recent survey.

The problem of guarding a terrain in $\mathbb{R}^2$ by two watchtowers has been studied in several recent papers. Bespamyatnikh *et al.* [4]

show that for the discrete case, deciding whether there exist two watchtowers of given height h that guard T can be done in $O(n^3)$ time. Using parametric search [16], they obtain an $O(n^3 \log^2 n)$-time algorithm for the optimization problem. They also present an $O(n^4)$-time solution that avoids parametric search. Ben-Moshe *et al.* [2] also address the discrete version of the problem and give an $O(n^{2.688} \log^2 n)$-time algorithm, based on parametric search and on the computation of all dominances for a set of n points in $\mathbb{R}^n$. The continuous case is solved in [4], by an algorithm that takes $O(n^4 \log^2 n)$ time, using parametric search.

Much less is known about terrain guarding in $\mathbb{R}^3$. Early work on terrain guarding, due to Cole and Sharir [10], shows that the problem of finding the minimum number of guards is NP-complete, even if the guards are placed on the terrain (no elevation is allowed). However, the case of a single watchtower guarding the terrain has been shown by Sharir [18] to be solvable in $O(n \log^2 n)$ time. An $O(n \log n)$-time algorithm for this problem was later obtained by Zhu [21].

Attention has also been given to the problem of guarding a two-dimensional terrain with the minimum number of guards placed on the terrain. Recently, Ben-Moshe, Katz, and Mitchell [3], and, independently, Clarkson and Varadarajan [8], discovered constant-factor approximation algorithms for the problem. Eidenbenz et al. [12] show that the related problem of guarding a simple polygon with a minimum number of guards is APX-hard. So for the problem of guarding polygons there is an ϵ such that it is NP-hard to obtain a $(1+\epsilon)$-approximation for the minimum number of guards.

Our results. In this paper we study the problem of guarding a terrain by two watchtowers in $\mathbb{R}^2$ and $\mathbb{R}^3$. For the planar case we obtain the following results.

(i) We prove that the discrete two-watchtower problem can be solved in $O(n^2 \log^4 n)$ time, significantly improving the previous solutions cited earlier. The algorithm works, without affecting its asymptotic running time, for the *semi-continuous* version as well, in which one of the bases can be anywhere on T and the other has to be placed at a vertex of T.

(ii) We prove that the continuous two-watchtower problem can be solved in $O(n^3 \alpha(n) \log^3 n)$ time, again significantly improving previous results.

(iii) We prove that the continuous version of the problem of guarding a finite set $P \subset T$ of m points by two watchtowers can be solved in $O(mn \log^4 n)$ time.

We also study the problem of guarding a terrain by two watchtowers in $\mathbb{R}^3$, and present an $O(n^{11/3} \operatorname{polylog}(n))$ algorithm for the discrete two-watchtower problem. This is the first nontrivial algorithm for the problem: A trivial solution for the discrete problem takes about $O(n^4)$ time.

All the results derived in this paper are based on the parametric-searching technique [16]. For each result, we first design a *decision procedure* that, given T and a real value $h > 0$, determines whether T can be guarded by two watchtowers of height at most h. Next, we apply the parametric searching technique to the decision procedure, to obtain an algorithm that finds where to place two watchtowers of smallest possible height.

The parametric search step involves developing a parallel algorithm for the decision problem and simulating it generically at the unknown value of the smallest height. Since the main contributions of the paper lie in developing the decision procedures, and because of lack of space, we describe the decision procedures for all the problems that we consider, describe in detail the parametric search steps for the discrete and semi-continuous planar two-watchtower problems, and omit the details of the other parametric searches. We also omit some of the details of the algorithm for guarding a finite set of points on a terrain in $\mathbb{R}^2$.

Section 2 describes the algorithms for discrete, semicontinuous, and continuous two-watchtower problems in $\mathbb{R}^2$, and Section 3 describes the algorithm for the discrete two-watchtower problem in $\mathbb{R}^3$.

2. THE PLANAR CASE

Let T be an x-monotone polygonal chain with n edges, and let V be the set of its vertices. Given T and a real value $h > 0$, we describe an algorithm to determine whether there exist two points $u, v \in T$, so that the two points $u(h), v(h)$ that lie at vertical distance h above u, v, respectively (the tops of watchtowers of height h erected at u and v), guard the entire terrain T. Section 2.1 describes the procedure for the discrete case, i.e., when $u, v \in V$, and Section 2.2 describes the procedure when u and v are arbitrary points of T.

2.1 The Discrete Problem

Let T and h be as defined above. We wish to determine whether there are two vertices $u, v \in V$ so that the top endpoints $u(h), v(h)$ of the watchtowers at u, v of height h guard the entire T. We present a procedure that does much more than that: Given one tower $u(h)$, it finds the *shortest* second tower that can be placed *anywhere* on the terrain, so that both towers guard T. In particular, this also serves as a decision procedure for the semi-continuous problem mentioned in the Introduction.

Computing visibility pairs. Consider the following problem, which arises as a main step in the overall decision procedure: Let P be a set of points on T, and let V be the set of vertices of T. We want to find, for each vertex $v \in V$, a point $p \in P$ (if one exists), that satisfies the following conditions: (i) p lies to the left of v; (ii) p sees v; and (iii) p is the rightmost point in P that satisfies (i) and (ii). Alternatively, among the points of P that satisfy (i) and (ii), the segment pv has the largest slope (it is easily verified that these two formulations are equivalent).

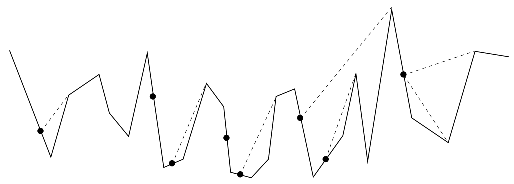

Denote by $\Pi(P, V)$ the set of these pairs. We assume that $|P| = O(n)$, which will be the case in our application. We solve this problem using the following divide-and-conquer technique. Take the set $P \cup V$ and draw a vertical line ℓ that splits it into two subsets of equal size. Let P_L, V_L (resp., P_R, V_R) denote the subsets of P and V that lie to the left (resp., right) of ℓ. We clearly have

$$\Pi(P, V) = \Pi(P_L, V_L) \cup \Pi(P_R, V_R) \cup \Pi(P_L, V_R'),$$

where V_R' is the subset of all points $v \in V_R$ for which there is no point $p \in P_R$ that forms with v a pair in $\Pi(P_R, V_R)$. Let n' denote the size of V_R'.

We compute $\Pi(P_L, V_L)$ and $\Pi(P_R, V_R)$ recursively, extract the subset V_R', and compute $\Pi(P_L, V_R')$ in the following direct way. Without loss of generality, assume that ℓ is the y-axis. We pass to the dual plane. For each $p \in P_L$, let γ_p denote the dual representation of the locus of all lines λ that pass through p, such that p sees the intercept $\lambda \cap \ell$. In the dual plane, γ_p is a rightward directed ray, contained in the line dual to p and emanating from the point w_p that represents the shallowest line through p with the above property. It is easy to show that the rays γ_p are pairwise openly disjoint. In complete analogy, for each point $q \in V'_R$, let δ_q denote the dual representation of the locus of all lines λ that pass through q, such that q sees the intercept $\lambda \cap \ell$ (here q sees the intercept to its left). In the dual plane, δ_q is a leftward directed ray, emanating from the point z_q that represents the shallowest line through q with the above property.

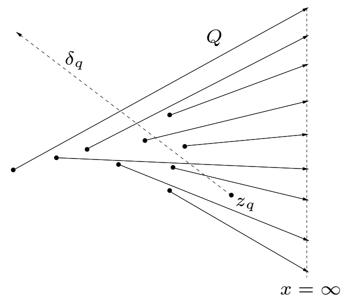

Let $q \in V'_R$, and let $p \in P_L$ be the point that satisfies $(p, q) \in \Pi(P_L, V'_R)$ (assuming that such a point exists). The point dual to the line that passes through p and q is a point that lies on both rays γ_p, δ_q, and is the *rightmost* point of intersection of δ_q with some ray γ_p. Let Q denote the unbounded (and degenerate) 'comb-like' simple polygon, whose boundary consists of all the rays γ_p, and of a vertical line ℓ_0 at $x = +\infty$; see the preceding figure. Then the point $p \in P_L$ that forms with q a pair in $\Pi(P_L, V'_R)$ is simply the output of a ray-shooting query inside Q along δ_q.

We first construct Q. For this, we need to compute the (apices of the) rays γ_p, for $p \in P_L$. Let T_L (resp., T_R) denote the region consisting of all points that lie above T and to the left (resp., right) of ℓ. Let b_ℓ denote the intersection point $T \cap \ell$. We construct the shortest-path map inside T_L with b_ℓ as a source, as in [13]. This allows us to find, in linear time, the terminal segments of all the shortest paths from b_ℓ to the points of P_L. For each $p \in P_L$, the apex of γ_p is then the dual of the line containing the respective terminal segment pa. The total cost of this step is linear.

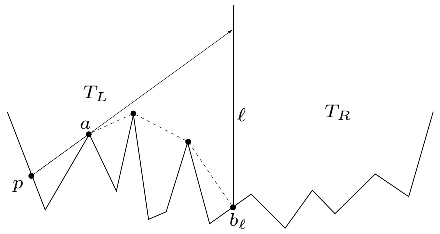

To complete the construction of Q, we need to sort the rays γ_p by their intercepts with ℓ_0. However, this order is equal to the order of the slopes of these rays, which in turn is the same as the order of the abscissas of the points of P_L, and this order can be computed prior to the beginning of the divide-and-conquer process. Hence, this substep also takes linear time.

Using the algorithm of Guibas *et al.* [13] (see also [5]), we next preprocess Q in linear time for ray-shooting queries that take logarithmic time each. We then construct the (apices of the) rays δ_q, for $q \in V'_R$. This is done, in a symmetric manner to the construction of the rays γ_p, using the polygon T_R instead of T_L. This step also runs in linear time. We then perform ray-shooting queries in Q with the rays δ_q, for $q \in V'_R$, thereby obtaining the set $\Pi(P_L, V'_R)$. Hence, the overall cost of the 'merge' step of the recursion is $O(n \log n)$ time, so the entire procedure takes $O(n \log^2 n)$ time.

In summary, we have an algorithm that computes the set $\Pi(P, V)$ in $O(n \log^2 n)$ time.

The overall decision procedure. We fix a vertex u of T, erect the first tower, of height h, over u, and compute the *visibility polygon* $W(u(h))$, which is the portion of the halfplane H lying above T that is visible from $u(h)$. The polygon $W(u(h))$ can be computed in linear time [13]. The portion of $\partial W(u(h))$ that lies on T consists of subsegments of the edges of T, at most one subsegment for each edge, which are delimited either by the vertices of T, or by intercepts of visibility rays that emanate from $u(h)$ and pass through some vertex of T. In either case, at least one endpoint of each visibility segment is a vertex of T.

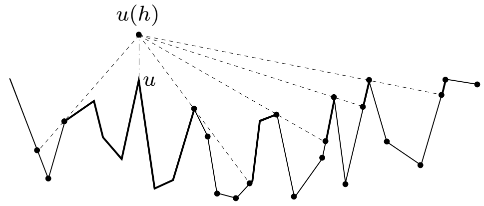

Let $P = P(u, h)$ denote the set of all the endpoints of the portions of the edges of T that are not visible from $u(h)$. Clearly, $u(h)$ and another tower top $v(h')$ (with a possibly different height h') guard T if and only if $v(h')$ sees all the points of P. This follows from the easy observation that if a point $v(h')$ above T sees both endpoints of a subsegment of some edge of T then it sees the entire subsegment.

The following lemma provides the crucial geometric property on which our algorithm relies.

LEMMA 2.1. *For a point v on T and a height h', the tower top $v(h')$ sees all the points of P that lie to the left of v if and only if $v(h')$ lies above all the lines pw such that $(p, w) \in \Pi(P, V)$ and w lies to the left of (or coincides with) v.*
Equivalently, $v(h')$ sees all the points of P that lie to the left of v if and only if, for each $p \in P$ to the left of v, $v(h')$ lies above the steepest line pw such that $(p, w) \in \Pi(P, V)$ and w lies to the left of (or coincides with) v.

Proof: If $v(h')$ sees all the points of P that lie to the left of v then it must lie above all the lines pw as in the first statement of the lemma, or else the corresponding point p would not be visible from $v(h')$. Conversely, suppose that $v(h')$ lies above all these lines, and let p be a point in P to the left of v. If $v(h')$ does not see p then the shortest path from p to $v(h')$ that lies above T must bend at some vertices of T. Let w be the leftmost such vertex. If $(p, w) \in \Pi(P, V)$, then (p, w) satisfies the conditions in the lemma, and $v(h')$ passes below the line pw, contrary to our assumption. If $(p, w) \notin \Pi(P, V)$ then there must exist another point $p' \in P$ that lies between p and w so that $(p', w) \in \Pi(P, V)$. The line $p'w$ is even steeper than pw, and $v(h')$ thus lies below it, again a contradiction. The second statement of the lemma is obvious. $\square$

Lemma 2.1 has a symmetric version, which handles visibility from a tower top $v(h')$ to the points of P to its right. For this version we need a symmetric version of $\Pi(P, V)$, involving pairs $(p, v) \in P \times V$ with p lying to the right of v, which is defined and constructed in a fully symmetric manner.

We sweep the plane with a vertical line ℓ from left to right, and construct a *height function* h_L, so that, for each $q \in T$, $h_L(q)$ is the height of the shortest tower erected over q that sees all the points of P to its left. Equivalently, in view of Lemma 2.1, $h_L(q)$ is the distance from q to the *upper envelope* E of the lines pw, for all the pairs $(p, w) \in \Pi(P, V)$ with w lying to the left of (or coinciding with) q.

As we sweep ℓ, $h_L(q)$ remains a linear function, equal to the vertical distance between an edge of E and an edge of T, until ℓ reaches a vertex of either E or T. If it reaches a vertex of E, we pass to another edge of E, and the corresponding update of h_L is easy and obvious. If ℓ reaches a vertex of T, in addition to the obvious local update of h_L, we need to dynamically update E, by adding the line pv, for the unique pair $(p, v) \in \Pi(P, V)$, to the set of lines that form E, or do nothing if such a pair does not exist. This has the effect of either leaving E unchanged, or creating two new vertices of E and deleting the old vertices between them, and can easily be done in overall time $O(\log n)$. After each update is performed, we also need to find the next vertex of E to the right of the current position of ℓ, which can be done in $O(1)$ time per update. Hence, the construction of h_L can be done in $O(n \log n)$ time (after the preprocessing stage of constructing $\Pi(P, V)$, which takes $O(n \log^2 n)$ time). Note that this algorithmic analysis also implies that the combinatorial complexity (number of vertices of the graph) of the function h_L is $O(n)$.

We then apply a symmetric process, in which we sweep the plane from right to left, and construct the symmetrically defined height function h_R. We then construct the upper envelope h^* of h_L and h_R, in time $O(n)$. The global minimum of h^* is the shortest height of a second tower, erected anywhere on T, that sees, together with $u(h)$, the entire terrain.

Since we need to repeat this step for each vertex u of T, the entire decision procedure runs in $O(n^2 \log^2 n)$ time. In summary, we thus have decision procedures for the discrete and semi-continuous problems, both running in $O(n^2 \log^2 n)$ time. Moreover, given one tower, a second tower of minimum height that guard T with the first tower, can be found in $O(n \log^2 n)$ time.

Applying parametric searching. We now show how to find the optimum height in the discrete and semi-continuous cases, by applying the parametric searching technique [16]. That is, we run a generic version of the decision procedure, in which h is left as an unknown parameter. Comparisons that depend on h are resolved by finding the few critical values of h at which the answer to the comparison may change, and by running (a concrete version of) the decision procedure at these critical heights, thereby finding the non-critical range that contains the optimum height h^* and is delimited by two consecutive critical heights. This determines the outcome of the comparison, and at the same time narrows down the range I that is known to contain h^*. We proceed in this manner through the execution of the generic decision procedure. If h^* is found during one of these comparison resolutions, we stop and report it. Otherwise, upon termination, we output the smaller endpoint of the final interval I.

To reduce the cost of the generic execution, we need to run a parallel version of it. More precisely, the only steps in the algorithm that need be parallelized are comparisons that depend on h. All such comparisons are resolved simultaneously, by running a binary search on all the resulting critical heights. If the algorithm runs in T_π parallel steps and uses p processors, then its overall cost is $O(T_\pi(p + D \log p))$, where $D = O(n^2 \log^2 n)$ is the cost of the decision procedure.

We describe two such generic parallel implementations of the decision procedure, one for the discrete problem and one for the semi-continuous problem. Both implementations require $O(\log n)$ parallel steps. The first algorithm requires $O(n^2 \log n)$ processors, and the second requires $O(n^2 \log^3 n)$ processors. Hence, both result in algorithms with $O(n^2 \log^4 n)$ running time.

We begin by preprocessing T as follows. For each vertex u of T, we compute the visibility map from the vertical halfline above u, using the algorithm of [13]. This takes overall time $O(n^2)$. For each u, the output contains all mutually visible pairs (v, v') of vertices, such that some point above u sees both v and v' along the line vv'. In addition, we also have the first point along that line, as we trace it from the point above u towards v and v', where it crosses T into the region below T (which may happen at the farthest vertex among v and v' or at another further point of T). As shown in [13], the number of such critical events is linear for each u. We thus obtain $O(n)$ critical heights over u, one for each of these visibility events. Repeating this step for each vertex u, we obtain a set H_0 of a total of $O(n^2)$ critical heights, and a corresponding set $\mathcal{I}$ of intercepts along the edges of T, where the above visibility rays vv' 'enter' the region below T.

We run a standard binary search through H_0, using the decision procedure described above. After $O(\log n)$ calls to that procedure, with overall cost $O(n^2 \log^3 n)$, we obtain an initial interval I_0 that contains the optimum height h^*. For each $u \in V$ and for each $h \in I_0$, the visibility polygon $W(u(h))$ has a fixed combinatorial structure. In particular, for each visible portion of an edge e of T, the nature of its endpoints is fixed, in the sense that each of them is either a fixed vertex of T, or the intercept along e of a visibility ray that emanates from $u(h)$ and passes through a fixed vertex of T before hitting e. Let $P(u, h)$ denote the set of those endpoints of the visibility segments of $W(u(h))$ that are not vertices of T, and set $P(h)$ to be the union of these sets, over all vertices u. The parameter h reminds us that these are parametric points that depend on h.

Our next step is to locate the points of $P(h)$ among the points of $\mathcal{I}$. We run the following process iteratively. At each stage we have a partition S of the edges of T into subsegments, each delimited by points of $\mathcal{I} \cup V$, an interval $I_S \subseteq I_0$, and a corresponding partition of $P(h)$ into subsets, each known to be contained in a fixed segment s of S, for all $h \in I_S$. Initially, S is the set of edges of T, $I_S := I_0$, and the sets $P_s(h) := P(h) \cap s$, for s an edge of T, form the desired partition of $P(h)$. (Recall that, for $h \in I_0$, the points of $P_s(h)$ move continuously along s, and no point enters or leaves this set.) Each iterative step refines S as follows.

For each segment $s \in S$, set $\mathcal{I}_s := \mathcal{I} \cap \mathrm{int}(s)$. Segments $s \in S$ with $\mathcal{I}_s = \emptyset$ are removed from S, and we report the corresponding sets $P_s(h)$ as part of the output of this stage. For any remaining segment s, let q_s denote the median point of $\mathcal{I}_s$. For each $p \in P_s(h)$, let h_p be the (unique) height at which p coincides with q_s. If there is no such height, or if h_p lies outside the range I_S, then we know which half of s contains p, and we ignore p in the remainder of the present step. Repeating this over all segments $s \in S$, we obtain a set H_S of $O(n^2)$ critical heights, and we run a binary search through them, as above, to obtain a subinterval $I' \subseteq I_S$, so that for each $h \in I'$, $s \in S$, and $p \in P_s(h)$, we know which half of s contains p. The next partition S' is then obtained by splitting each surviving segment $s \in S$ at the median point q. We set $I_{S'} := I'$, and get the new partition of $P(h)$ by assigning each point p to the new half-segment it belongs to. After at most $O(\log n)$ steps, all segments of S are removed, and each point $p \in P(h)$, for any h in the final range I_S, is located between two successive points of $\mathcal{I} \cup V$, as required. The overall cost of this step is $O(n^2 \log^4 n)$.

Let I_1 denote the final range I_S. For each $h \in I_1$ and for each $u \in V$, the set $\Pi(P(u, h), V)$ is now fixed, as is easily seen. This follows from the observation that whenever $\Pi(P(u, h), V)$ changes, some point in $P(h)$ must coincide with a point in $\mathcal{I} \cup V$. This allows us to run the left-to-right sweeping procedure, for each fixed $u \in V$, without having to fix the value of h, and obtain, for each swept vertex v, the set L_v of all the *active* lines pw, for $(p, w) \in \Pi(P(u, h), V)$, namely, the lines corresponding to those pairs (p, w) with w lying to the left of v. In the discrete problem, we need to find which of these lines attain the maximum height at v, and assert that this maximum height is $\leq h^*$; see below for more details. Handling the semi-continuous case is somewhat more involved, and will be discussed later.

To parallelize this step, we store the vertices of T at the leaves of a minimum-height binary tree τ. Each internal node ξ of τ represents the set $V(\xi)$ of all vertices stored at the leaves of the subtree $\tau(\xi)$ rooted at ξ. Let $L(\xi)$ denote the set of all the lines pw, for $(p, w) \in \Pi(P, V)$, and $w \in V(\xi)$, and let $E(\xi)$ denote the upper envelope of the lines in $L(\xi)$. Note that $E(\xi)$ is an envelope of $|L(\xi)|$ lines, so its complexity is $O(|L(\xi)|)$. We have $\sum_{\xi \in \tau} |L(\xi)| = O(n \log n)$. Hence, the overall complexity of all the envelopes $E(\xi)$ is $O(n \log n)$, for a fixed first tower base u. We construct each of the envelopes $E(\xi)$ as follows. We dualize the respective lines pw to points in the dual plane, and run the parallel algorithm of [1] for constructing the upper convex hull of these (parametric) points. For a set of m points, this algorithm uses $O(m)$ processors and runs in $O(\log m)$ time. Each generic comparison that the algorithm performs is either between the x-coordinates of two dual points (that is, between the slopes of two input lines), or a left-turn test involving three dual points. Clearly, all the envelopes $E(\xi)$, over all the nodes ξ of τ and over all the first tower bases u in V, can be constructed in parallel, using a total of $O(n^2 \log n)$ processors and $O(\log n)$ parallel steps. This yields an overall parametric searching stage that runs in time $O(n^2 \log^4 n)$, as follows from the general time bound for parametric searching given above. Note that the output envelopes are still parametric and vary with h. However, their combinatorial structure is fixed: For each envelope we know the sequence of lines (or, rather, pairs (p, w) defining those lines) that attain the envelope from left to right, and thus also know the nature of each breakpoint of the envelope.

We next apply another parallel stage, in which we locate, for each envelope $E(\xi)$ and each of its breakpoints q, the edge of T that lies vertically above or below q. We process these breakpoints in parallel, and for each breakpoint q we run a binary search with its (parametric) x-coordinate among the x-coordinates of the vertices of T. This stage uses $O(n^2 \log n)$ processors and runs in $O(\log n)$ parallel steps, so its overall cost is also $O(n^2 \log^4 n)$.

This initial part of the algorithm applies to both the discrete and the semi-continuous cases.

Proceeding with the discrete case is now easy: Keeping the first tower base u fixed, we process all the vertices v of T in parallel. For each vertex v, we obtain the set of vertices that lie to the left of v as the (disjoint) union of $O(\log n)$ subtrees of τ. For each of these trees $\tau(\xi)$, we locate v among the vertices of $E(\xi)$, using binary search. Since we have already located the vertices of each $E(\xi)$ among the vertices of T in their left-to-right order, these binary searches can be performed explicitly, without having to use generic parametric searching. Hence this step can be performed in overall $O(n^2 \log^2 n)$ time. We now have, for each pair u, v, a set of $O(\log n)$ lines, each attaining a respective subenvelope $E(\xi)$ at v, and we compute their maximum height at v. This is still a parametric step, which can be easily performed in parallel, using $O(n^2 \log n)$ processors and $O(\log n)$ parallel depth, so that its overall cost is, as above, $O(n^2 \log^4 n)$. To end the algorithm, we output a pair $u, v \in V$ for which the height computed at v, as described above, is $\leq h$. This step is also parametric, and can be performed within the time bound for the preceding step. It narrows down the range of h to its final value, and we terminate by returning the minimum of this interval (unless the optimum height h^* has already been detected in one of the comparison resolution steps). Hence we obtain:

THEOREM 2.2. *The discrete two-watchtower problem for a terrain in $\mathbb{R}^2$ with n edges can be solved in $O(n^2 \log^4 n)$ time.*

We next handle the semi-continuous case. We proceed through the preceding stages, up to the point where all the envelopes $E(\xi)$ have been constructed, and their vertices have been located over the edges of T. We then proceed as follows. Consider a fixed first tower base u. For each edge $e = v'v$ of T, take its right endpoint v and obtain, as above, the set of vertices that lie to the left of v as the (disjoint) union of $O(\log n)$ subtrees of τ. We need to find the shortest vertical distance from e to the upper envelope of the $O(\log n)$ envelopes $E(\xi)$, over the corresponding subtrees $\tau(\xi)$.

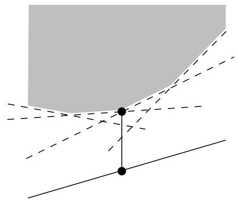

Note that the shortest vertical distance is attained (i) at an endpoint of e, (ii) at a vertex q of one of the envelopes $E(\xi)$, or (iii) at a point q of intersection between two envelopes $E(\xi)$, $E(\xi')$. In case (ii), the edge of $E(\xi)$ incident to q and lying to its left (resp., right) must have slope smaller than (resp., larger than) that of e. In case (iii), q is incident to an edge of $E(\xi)$ and to an edge of $E(\xi')$, so that the slope of e lies in between the slopes of these two edges.

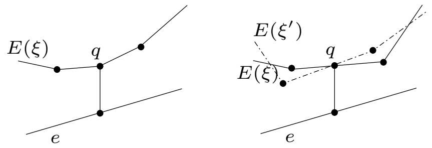

Consider such a shortest vertical distance of type (iii). Both subenvelopes $E(\xi)$, $E(\xi')$ correspond to subtrees $\tau(\xi)$, $\tau(\xi')$ that are left children of nodes along the path in τ from the leaf storing v to the root. Without loss of generality, suppose that ξ is deeper in the tree than ξ'. Then if $\tau(\xi)$ is part of the output for a vertex v of T, $\tau(\xi')$ must also be part of that output. We will refer to this case by saying that ξ' is a *left great uncle* of ξ. In other words, we have argued that, independent of v, a subenvelope $E(\xi)$ can form intersection points of type (iii) with only $O(\log n)$ other subenvelopes $E(\xi')$ for which ξ' is higher in τ, and all these nodes τ' are left great uncles of ξ. We therefore proceed as follows. We only fix the first tower base u, and obtain the corresponding set $P = P(u, h)$ and tree τ. We fix a pair (ξ, ξ') of nodes of τ such that ξ' is a left great uncle of ξ. We construct an implicit representation of the upper envelope of $E(\xi) \cup E(\xi')$, by merging the vertices and edges of $E(\xi)$ into $E(\xi')$, as follows. For each vertex q of $E(\xi)$, locate it among the vertices of $E(\xi')$, with respect to the x-order, using binary search (which is parametric). Next, for each edge e of $E(\xi)$, find its intersections, if any, with $E(\xi')$; since $E(\xi')$ is a convex polygonal chain, e can meet it in at most two points, and they can be found using (parametric) binary search. We assign a processor to each vertex and edge of $E(\xi)$ for each such pair (ξ, ξ'). Since the overall size of all the envelopes $E(\xi)$ is $O(n \log n)$, and each has only $O(\log n)$ left great uncles, we need a total of $O(n \log^2 n)$ processors, and $O(\log n)$ parallel steps. We run this procedure for all first tower bases u in parallel, and thus use $O(n^2 \log^2 n)$ processors and $O(\log n)$ parallel steps, so the overall cost of this parametric searching step is, as above, $O(n^2 \log^4 n)$.

We next construct the set B of all breakpoints of the individual envelopes $E(\xi)$, of the breakpoints of envelopes of pairs of envelopes $E(\xi)$, $E(\xi')$, and of the vertices of T, over all the first tower bases u. The preceding analysis implies that the size of B is $O(n^2 \log^2 n)$. Using the same parametric binary search as in a previous stage, we locate the new breakpoints (of envelopes of pairs of envelopes) among the vertices of T. This step too takes $O(n^2 \log^4 n)$ time.

For each fixed first tower base u, we next process each edge e of T. We query in T with the right vertex v of e and obtain the set $L(v)$ of lines pw, as defined above, as the union of $O(\log n)$ disjoint subtrees of τ. Let $B(v)$ denote the subset of B consisting of the two endpoints of e, and of all the breakpoints of subenvelopes and of pairs of subenvelopes that correspond to the above $O(\log n)$ subtrees, with the additional requirement that they lie above or below e. Except for vertices of T, the sets $B(v)$ are pairwise disjoint, and their overall size is at most twice the size of B. With careful preprocessing, all the sets $B(v)$, over all tower bases u, can be retrieved in $O(|B|) = O(n^2 \log^2 n)$ time (note that this step is no longer parametric).

With u and v fixed, we now assign a processor to each pair of a point $q \in B(v)$ and a subenvelope $E(\xi'')$ in the output of v. The overall number of processors is $O(n^2 \log^3 n)$. The processor assigned to q and $E(\xi'')$ has to determine whether q lies below $E(\xi'')$, which it can do using (parametric) binary search over the vertices of $E(\xi'')$. Hence, with $O(n^2 \log^3 n)$ processors and $O(\log n)$ parallel depth, we can collect all points $q \in B$ that are not hidden from above by another subenvelope. It is easily verified that, for each fixed u and v, among the points that pass these tests, exactly one point q has the property that the slope of e lies between the slopes of the two envelope edges incident to q. This surviving point yields the desired shortest vertical distance from e to the envelope. We find this point by testing the slope condition at each of the surviving points.

To end the algorithm, we need to test (parametrically) that among the surviving points there exists one for which the corresponding vertical distance is $\leq h$. This step narrows down the range of h to its final value, and we terminate by returning the minimum of this interval. We thus have:

THEOREM 2.3. *The semi-continuous two-watchtower problem for a terrain in $\mathbb{R}^2$ with n edges can be solved in $O(n^2 \log^4 n)$ time.*

2.2 The Continuous Problem

We next consider the continuous version of the problem. Given T and a real value h, the decision procedure determines whether there exist two towers of height h that together see the entire terrain.

Let e_1, e_2 be two fixed distinct edges of T, and consider the subproblem of determining whether there exist points $p \in e_1$, $q \in e_2$ such that the two tower tops $p(h), q(h)$ see together the entire T. Parametrize the locations of p and q by their respective x-coordinates s and t.

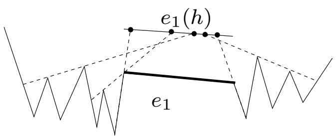

Let $e_1(h)$ denote the segment obtained by translating e_1 upwards by distance h. Compute the *visibility structure* of T from $e_1(h)$, which can be done in $O(n \log n)$ time [6], by adapting the algorithm of [13]. The output is a partition of $e_1(h)$ into $O(n)$ intervals, delimited by points that see two vertices of T along a common ray. We extend each such ray to the point where it first crosses T and enters the region below it. The sequence of all these crossing points that lie on a specific edge g of T is denoted by $\Sigma(e_1, g)$. If the line containing e_1 intersects T on g, we add this intersection point to $\Sigma(e_1, g)$. (This adds at most two points to the union of $\Sigma(e_1, g)$ over all $g \in T$.)

As the point $p = p(s)$ moves along e_1 from left to right, the corresponding tower top $p(s, h)$ traces the segment $e_1(h)$. Let g be another edge of T. The point $p(s, h)$ sees a portion of g which, if not empty, is delimited by the endpoint of g farthest from e_1, and by a point $z(s) = z_g(s)$ that moves continuously with p. As long as $p(s, h)$ does not cross a critical point of the visibility structure, $z(s)$ is either the endpoint of g nearest to e_1, or is the intercept of a visibility ray that emanates from $p(s, h)$ and passes through a fixed vertex v of T. When $p(s, h)$ crosses a critical point, and $z_g(s)$ crosses the matching point in $\Sigma(e_1, g)$, the 'pivot' vertex v may change, but the motion of $z(s)$ remains continuous, and *monotone*: The direction of motion of $z(s)$ depends on whether g lies to the left or to the right of e_1, and on whether the pivot vertex through which the segments $p(s, h)z_g(s)$ pass lies above or below the line containing $e_1(h)$. It is easily seen that when the pivot vertex changes, the new vertex continues to lie on the same side of the line through $e_1(h)$, and thus the motion of $z_g(s)$ continues in the same direction.

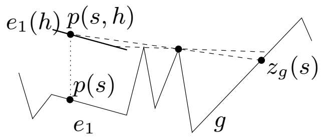

Abusing the notation, let us denote by $z_g(s)$ the x-coordinate of this point. Then the collection of functions $z_g(s)$, over all edges g of T, is a collection of continuous piecewise linear-rational functions of s. We omit the easy proof in this version. The overall number of linear rational portions of these functions, over all edges g, is $O(n)$. Indeed, such a portion ends either at an endpoint of $e_1(h)$, or at a critical visibility point on $e_1(h)$ that sees two vertices of T along a common ray, and the number of such points on $e_1(h)$ is only $O(n)$.

We apply an analogous construction to the edge e_2, denote the point that moves along e_2 by $q = q(t)$, the corresponding tower top by $q(t, h)$, and the collection of functions that trace the (x-coordinates of the) endpoints of the visibility subsegments of edges g of T, as seen from $q(t, h)$, by $\{w_g(t)\}$. Let C denote the rectangle $e_1^* \times e_2^*$ in the st-plane, where e_1^*, e_2^* denote respectively the x-projections of e_1, e_2. Partition C, only for the purpose of analysis, into $O(n^2)$ subrectangles, by the vertical and horizontal lines through the critical points of the visibility structures on $e_1(h)$ and $e_2(h)$, respectively.

Fix an edge g of T, different from e_1, e_2. Let β_g denote the curve $z_g(s) = w_g(t)$, drawn in C. Within each subrectangle of C that β_g crosses, it is a hyperbolic arc. Moreover, β_g is a connected curve which is both s- and t-monotone, and its endpoints lie on ∂C. Finally, the number of hyperbolic arcs that constitute the curves β_g, over all edges g of T, is only $O(n)$. We omit the easy proofs of these facts. Each curve β_g thus partitions C into two portions, one of which, denoted MV_g, consists of all points (s, t) that represent placements of two towers on e_1 and e_2 that guard g, whereas the complement of MV_g consists of points representing placements of towers where not all of g is visible. Figure 1 exhibits four types of such portions of C. The classification depends (i) on the left-to-right order of the edges e_1, e_2, g, (ii) on whether the pivot vertex through which the segments $p(s, h)z_g(s)$ pass lies above or below the line containing $e_1(h)$, and (iii) on whether the pivot vertex through which the segments $q(t, h)w_g(t)$ pass lies above or below

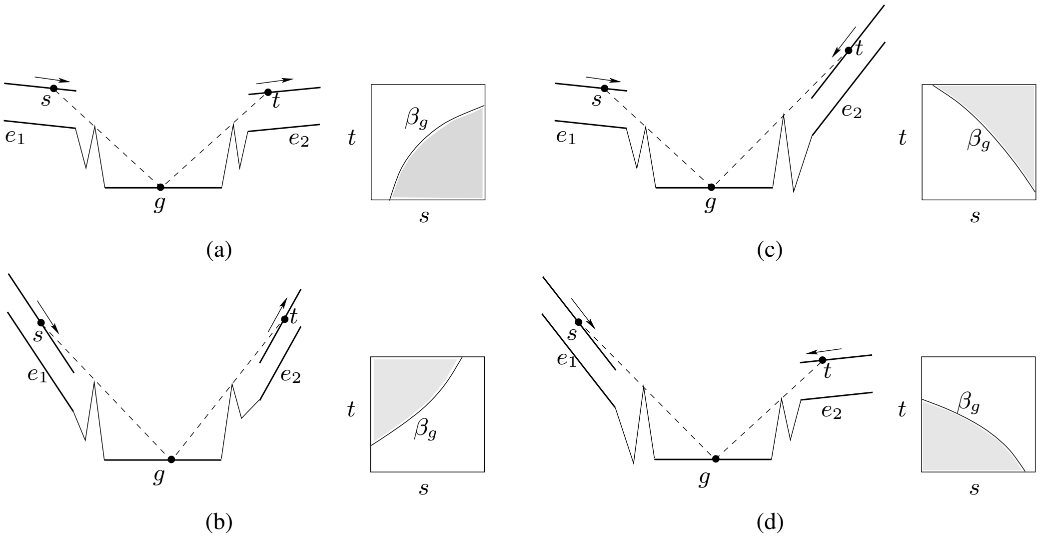

Figure 1. The four types of mutual visibility of an edge g from a pair of moving towers, and the corresponding st-regions where the entire g is guarded.

the line containing $e_2(h)$. As already remarked, when either of these pivot vertices changes, the new vertex continues to lie on the same side of the line through $e_1(h)$ or $e_2(h)$.

The decision procedure then continues as follows. Keeping the pair of edges e_1, e_2 fixed, we first construct, in linear time, the regions MV_g, or rather the curves β_g, over all the edges g of T. This is done by merging the sequences $\Sigma(e_1, g)$ and $\Sigma(e_2, g)$ into a common sequence $\Sigma(e_1, e_2, g)$. We then process the subintervals of g delimited by the points of $\Sigma(e_1, e_2, g)$. Each such interval I defines a piece of β_g, in which the common visible point $z_g(s) = w_g(t)$ moves in I. Clearly, only intervals I that lie in the range of both functions contribute nonempty portions to β_g. The computation of all the regions MV_g takes $O(n \log n)$ time for fixed e_1 and e_2.

Next, we compute the intersection $\bigcap_g MV_g$ of these regions. If this intersection is nonempty, any point (s, t) in it represents a placement of two towers on e_1 and e_2 that guard T. Conversely, if the intersection is empty, no such placement exists. Since each region is bounded by a curve that is s-monotone, $\bigcap_g MV_g$ is a *sandwich region* between the upper envelope of the curves β_g, for which MV_g lies above β_g (in the t-direction), as depicted in Figure 1(b,c), and the lower envelope of the curves β_g, for which MV_g lies below β_g, as depicted in Figure 1(a,d). Any pair of our hyperbolic arcs intersect at most twice. Hence the complexity of either envelope, and thus also of the sandwich region, is $O(\lambda_4(n))$, and it can be computed in time $O(\lambda_3(n) \log n)$, using the algorithm of Hershberger [14, 19]. This dominates the time for computing the regions MV_g. We repeat this procedure for every pair e_1, e_2 of edges of T. We stop as soon as we find a pair for which the intersection $\bigcap_g MV_g$ is nonempty, and then report a corresponding placement of the two towers. If all these intersections are found to be empty, we report that no pair of towers of height h can see the entire T. Hence, the overall cost of the decision procedure is $O(n^3 \alpha(n) \log n)$.

Again, by applying the parametric searching technique to this decision procedure, whose details are omitted from this version, we obtain the following result.

THEOREM 2.4. *The continuous two-watchtower problem for an x-monotone polygonal chain in $\mathbb{R}^2$ with n edges can be solved in $O(n^3 \alpha(n) \log^3 n)$ time.*

2.3 Guarding a Fixed Set of Points

Let T be a terrain with n vertices, and let $P = \{p_1, \ldots, p_m\}$ be a set of m points, all lying on T, and sorted in this left-to-right order along T. We assume that m is polynomial in n, so that $\log m = O(\log n)$. The goal is to place two towers $u(h), v(h)$ of the smallest possible height *anywhere* on T, so that they guard the entire set P. We show that this problem can be solved in time close to $O(mn)$. We first develop a decision procedure that for a fixed h finds placements for two towers of height h over the terrain that together cover the entire set P, or determines that no such pair of towers exists, and then apply the parametric search technique.

For each point $u \in T_h$, let $r(u)$ (resp., $t(u)$) denote the leftmost (resp., rightmost) point $p \in P$ that lies to the right of u and is not visible from u, or $+\infty$ if there is no such point. For each point $u \in T_h$, let $P_u \subset P$ be the set of points that u does not see. We have the following lemma.

LEMMA 2.5. *Let $u, v \in T$ with u to the left of v. Then $u(h)$ and $v(h)$ guard the set $P^+ \subseteq P$ of all points of P to the right of v if and only if $r(v) > t(u)$.*

For each point $p_i \in P$, let $W_i = \partial W(p_i)$ denote the boundary of the visibility polygon of p_i with respect to the terrain T. This boundary, W_i, is a sequence of connected portions of T, interleaved with segments qr, where q is a vertex of T, $r \in T$, and the ray emanating from p_i towards q sees both q and r and crosses T at r.

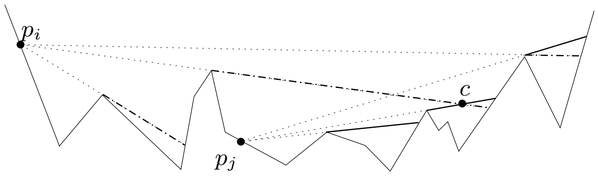

LEMMA 2.6. *Let s be a segment that lies fully above T. Then s intersects W_i at most once to the left of p_i and at most once to its right.*

Let E_i be the set of points where T_h intersects W_i. By Lemma 2.6, each edge of T_h intersects W_i at most once, with the exception of the edge that lies directly above p_i, which does not intersect W_i. It thus follows that $|E_i| \leq n-1$. Let $E = \bigcup_{i=1}^m E_i$. Let u and v be two consecutive points of E along T_h. Clearly, from any point z on T_h that lies between u and v, we see the same subset of P. Therefore if there are two tower-tops on T_h that guard P, then there are two points in E that do the same. Our algorithm will therefore determine whether there exist two points $u, v \in E$ that guard P.

The algorithm first computes W_i, for $i = 1, \ldots, m$ using the algorithm of [13]. This takes a total of $O(mn)$ time. Then we compute the sets E_i, for $1 \leq i \leq m$, by traversing in parallel W_i and T_h from left to right, locating intersections between edges of T_h and W_i whose x-projections overlap. Since both T_h and W_i are connected x-monotone polygonal chains with $O(n)$ edges each, this step takes $O(n)$ time for each W_i, for a total of $O(mn)$ time. We organize E in a list sorted by x-coordinate, represented as a balanced search tree.

We next compute the pointers $r(u)$ and $t(u)$ for every $u \in E$ by traversing the points in E from left to right along T_h, while maintaining the subset $P_u \subseteq P$ of points that are not visible from the current point u. We start out from $-\infty$, with $P_{-\infty} = \emptyset$. When we advance from a point $u \in E$ to the next point v on T_h, we add and/or delete a point to/from P_u to obtain P_v as follows. Let p_j be the point such that u is an intersection of W_j with T_h, and let p_k be the point such that v is an intersection of W_k with T_h. If to the right of u we cannot see p_j, then we add p_j to P_u. Similarly if to the right of v we see p_k then we delete p_k from P_u. Note that if $j = k$ and between u and v we cannot see p_j then $P_v = P_u$ and we do not have to add and delete p_j. After these updates we have the set P_v, and we set $r(v)$ (resp., $t(v)$) to be the leftmost (resp., rightmost) point in P_v which is to the right of v, or to $+\infty$ if P_v is empty. We perform this computation in $O(mn \log n)$ time, by maintaining P_u in a search tree, sorted by increasing x-coordinate.

Let W_i^r (resp., W_i^ℓ) be the portion of W_i that lies to the right (resp., left) of p_i. The core of our algorithm is based on the following lemma and its corollary.

LEMMA 2.7. *Let $i < j$, and let u be the leftmost intersection point of W_i^r and W_j^r. Then to the left of u, W_j^r lies below W_i^r, and to the right of u, W_j^r lies above or overlaps with W_i^r.*

COROLLARY 2.8. *Let $p_{i_1}, \ldots, p_{i_k}$ be a subset of the points of P, ordered from left to right. Then the upper envelope of $W_{i_1}^r, \ldots, W_{i_k}^r$ is a concatenation of connected portions of $W_{i_1}^r, \ldots, W_{i_k}^r$, such that, for each $i = 1, \ldots, k-1$, the portion of $W_{i_j}^r$ precedes the portion of $W_{i_{j+1}}^r$. Some of the portions may be empty.*

For each point $u \in E$, recall that $P_u \subset P$ is the set of points that u does not see. As in the computation of $r(u)$ and $t(u)$, our algorithm traverses the points of E from left to right along T_h, while maintaining the sets P_u incrementally. For each point u that we traverse, we decide whether there is another point $v \in E$ *to the right* of u, such that all points in P_u are visible from v. To that end we also maintain the upper envelope $\mathcal{E}(u)$ of $\{W_i^r \mid p_i \in P_u\}$. By definition, any point of E on or above $\mathcal{E}(u)$ sees all the points of P_u *to its left*. Let $v \in E$ be the point on or above $\mathcal{E}(u)$ that lies to the right of u and has maximum $r(v)$. By Lemma 2.5, if $r(v) > t(u)$, then u and v cover the entire set P, and otherwise there is no point in E that can sees the entire P_u.

We maintain P_u as a binary search tree $\mathcal{T}$. A leaf in $\mathcal{T}$ corresponds to W_i^r, for some $p_i \in P_u$. Each internal node x of $\mathcal{T}$ represents the upper envelope $\mathcal{E}_x$ of the chains W_i^r, for all points p_i in its subtree. It follows that if r is the root of $\mathcal{T}$ then $\mathcal{E}_r = \mathcal{E}(u)$. For $x \in \mathcal{T}$, the envelope $\mathcal{E}_x$ is represented as a search tree $\mathcal{T}_x$. Each node $\alpha \in \mathcal{T}_x$ represents an edge e_α of $\mathcal{E}_x$. $\mathcal{T}_x$ is organized such that if we traverse it in symmetric order we obtain the edges on the envelope from left to right. We associate with each node $\alpha \in \mathcal{T}_x$ another search tree $\mathcal{T}_{x,\alpha}$, that represents the points of E that lie on or above the corresponding edge, ordered from left to right.

The representation of $\mathcal{T}_{x,\alpha}$ is as follows. Recall that the points of E are stored in a static search tree, sorted from left to right along T_h. T_h intersects e_α at points that belong to E. These points delimit contiguous subsequences of E, which alternate between lying above e_α and lying below it. We collect the subsequences that lie above e_α, and represent each of them by a leaf of $\mathcal{T}_{x,\alpha}$, where these leaves are sorted from left to right in $\mathcal{T}_{x,\alpha}$. Each such leaf simply stores pointers to the two nodes of the E-tree that delimit the corresponding subsequence. In addition, each leaf of $\mathcal{T}_{x,\alpha}$ stores the maximum value $r(v)$ of the points v in the subsequence of E that it represents, and the point v where this maximum is attained. Each internal node $\xi \in \mathcal{T}_{x,\alpha}$ stores the maximum value $r(v)$ of the points $v \in E$ which are above e_α in all the subsequences represented by the leaves of the subtree rooted at ξ, as well as the point v where this maximum is attained.

We propagate points with maximum $r(v)$ value further up in $\mathcal{T}_x$ as well. Each node $\alpha \in \mathcal{T}_x$ has, in the root of $\mathcal{T}_{x,\alpha}$, the maximum value $r(v)$ of the points $v \in E$ above e_α. We also store in each node $\alpha \in \mathcal{T}_x$ the maximum $r(v)$ value of a point $v \in E$ which is above any of the edges of $\mathcal{E}_x$ that reside in the subtree rooted by α.

Consider $\mathcal{T}$ when we are at a point u as we move along T_h. The root r of $\mathcal{T}$ stores in $\mathcal{T}_r$ the upper envelope $\mathcal{E}(u)$ of W_i^r for all points $p_i \in P_u$. We use $\mathcal{T}_r$ to find the point $v \in E$ on or above $\mathcal{E}(u)$ that lies to the right of u and has maximum $r(v)$. We retrieve v from $\mathcal{T}_r$ essentially by searching $\mathcal{T}_r$ with the x-coordinate of v, and inspecting roots of subtrees hanging to the right of this search paths. (details are provided in the full version.) By what has been argued above, we test whether $r(v) > t(u)$. If so, v and u cover the entire set P, and we stop and report them. Otherwise there is no point in E that, together with v, covers P, and then we move to the next point of E.

As we move from a point $u \in E$ to the next point $v \in E$ to the right along T_h, we update $\mathcal{T}$ so that it stores $\mathcal{E}(v)$ — the upper envelope of the functions W_i^r, for $p_i \in P_v$. We do this by at most one deletion and one insertion of a polygon W_j^r to the tree $\mathcal{T}$, in $O(\log^3 n)$ time; we defer the details of the update to the full version of the paper.

By applying the parametric searching techniques to this decision procedure we obtain the following theorem.

THEOREM 2.9. *Given a terrain T in $\mathbb{R}^2$ with n edges, and a set $P \subset T$ of m points, we can find, in $O(mn \log^4 n)$ time, two towers of smallest height that can be placed anywhere on T and together cover the entire set P.*

3. THE DISCRETE TWO-WATCHTOWER PROBLEM IN 3-SPACE

Let T be a polyhedral terrain in $\mathbb{R}^3$ with n edges, and let $h > 0$ be a real parameter. Without loss of generality we can assume that each face of T is a triangle. We wish to determine whether there exist two vertices u, v of T, so that the watchtowers $u(h), v(h)$ of height h erected at u, v guard the entire terrain, as defined in the introduction. We call any such pair (u, v) a *guarding pair* of T. Let V denote the set of vertices of T. For each $v \in V$ and a facet f of T, let $H_v(f)$ denote the portion of f that is invisible from $v(h)$. The regions $H_v(f)$, called the *invisibility region* of v in f, can be constructed as follows.

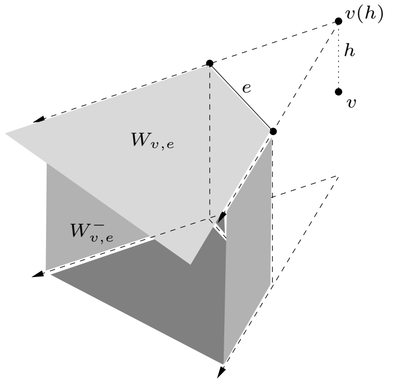

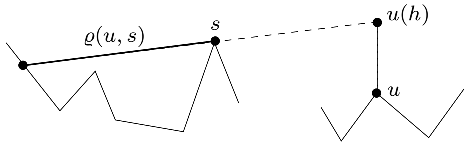

T

For each vertex $v \in V$ and for each edge e of T not adjacent to v, let $W_{v,e}$ denote the truncated planar wedge that is the union of all rays emanating from $v(h)$ and passing through e, minus the triangle spanned by $v(h)$ and e. Let $\pi_{v,e}$ denote the plane containing $W_{v,e}$, and let $W_{v,e}^-$ denote the unbounded prismatic region consisting of all points that lie vertically below $W_{v,e}$. The prism $W_{v,e}^-$ is bounded by $W_{v,e}$, and by three vertical walls (semi-unbounded vertical strips) bounded from above by the three edges of $W_{v,e}$. Informally, $W_{v,e}^-$ is the portion of $\mathbb{R}^3$ that is invisible from $v(h)$ if T degenerates to the single vertical wall bounded from above by e.

For each vertex $v \in V$, let $\mathcal{W}_v$ denote the set of all wedges $W_{v,e}$, for edges e of T that are not adjacent to v. Let Π_v and $\mathcal{W}_v^-$ denote the set of all corresponding planes $\pi_{v,e}$ and prisms $W_{v,e}^-$, respectively. Put $\mathcal{W} := \bigcup_{v\in V} \mathcal{W}_v$, $\Pi := \bigcup_{v\in V} \Pi_v$, and $\mathcal{W}^- := \bigcup_{v\in V} \mathcal{W}_v^-$.

For each face f of T and each vertex $v \in V$ not adjacent to f, we can construct $H_v(f)$ by intersecting each prism $W_{v,e}^- \in \mathcal{W}_v^-$ with f, and by taking the union of all the resulting regions. The complement $f \setminus H_v(f)$ is the *visibility region* of v (in f). The set of visibility regions, over all faces f of T, is the so-called *visibility map* from $v(h)$; its complexity is $\Theta(n^2)$ in the worst-case, and it can be computed in $O(n^2)$ time [19].

It is easy to establish the following properties of $H_v(f)$. The intersection of $H_v(f)$ with any vertical halfplane bounded by the vertical line λ_v through v, is a (possibly empty) line segment contained in f, and having one endpoint (the one nearer to v) on ∂f. As the halfplane rotates about λ_v, the other endpoint of the invisibility segment traces a polygonal path $\gamma_v(f) \subset f$, which is monotone with respect to the horizontal polar orientation θ of the halfplane about λ_v. The edges of $H_v(f)$ that lie in $\mathrm{int}(f)$ are portions of intersections of wedges $W_{v,e}$ with f. Moreover, $\gamma_v(f)$ can be interpreted as the upper envelope of these intersection edges, in an appropriate polar coordinate system within f. Hence, the combinatorial complexity of $H_v(f)$ is $O(n\alpha(n))$ [19]. The overall complexity of all these regions, for a fixed vertex v, is $O(n^2)$, as each of its vertices corresponds to a vertex of the visibility map of T from $v(h)$; as noted above, all the invisibility regions $H_v(f)$, over all faces f of T, can be computed in $O(n^2)$ time.

By definition of $H_v(f)$, two vertices u, v of T form a guarding pair of T if and only if $H_u(f) \cap H_v(f) = \emptyset$ for every facet f of T, or equivalently, $H_v(f)$ is fully visible from $u(h)$ for every f. Moreover, as is easily checked, $H_v(f)$ is entirely visible from $u(h)$ if and only if $u(h)$ sees every point on its boundary $\partial H_v(f)$.

For a pair $u, s \in V$, if s is visible from $u(h)$, let $\varrho(u,s)$ be the segment whose endpoints are s and the first intersection point with T of the ray from s in direction $\vec{u(h)s}$. If s is not visible from $u(h)$, $\varrho(u,s)$ is not defined.

LEMMA 3.1. *Let $u, v \in V$, and let f be a face of T. $\partial H_v(f)$ is entirely visible from $u(h)$ if and only if the following two conditions hold:*

(V1) $u(h)$ sees every vertex of $H_v(f)$.

(V2) There does not exist a vertex $s \in V$ such that the segment $\varrho(u,s)$ passes vertically above an edge of $H_v(f)$.

PROOF. Indeed, if $u(h)$ sees the entire boundary $\partial H_v(f)$ then (V1) and (V2) hold. Conversely, suppose for the sake of contradiction that (V1) and (V2) hold but an edge g of $H_v(f)$ is not fully visible from $u(h)$. Since g is not fully visible from $u(h)$ it must intersect some prism $W_{u,e}^-$. Let $p(W_{u,e}^-)$ be the point in the intersection of $W_{u,e}^-$ with the vertical plane through g of largest vertical distance from g. Among the prisms $W_{u,e}^-$ intersecting g let $W_{u,e'}^-$ be the one such that the distance of $p(W_{u,e}^-)$ from g is the largest. It is easy to verify that the ray emanating from $u(h)$ and passing through one of the endpoints s of e' (and through $p(W_{u,e}^-)$) defines a segment $\varrho(u,s)$ that passes vertically above g, and thus violates (V2). $\square$

Conditions (V1) and (V2) are equivalent to the following respective conditions.

(V1') There does not exist a wedge $W_{u,e}$, for an edge e of T not adjacent to u, such that $W_{u,e}^-$ contains a vertex of $H_v(f)$.

(V2') There does not exist a vertex $s \in V$ such that the xy-projections of the segment $\varrho(u,s)$ and of some edge of $H_v(f)$ intersect.

The equivalence of (V2) and (V2') follows from the fact that the segments $\varrho(u,s)$ lie above T, whereas all the edges of $H_v(f)$ lie on T. Recall that conditions (V1') and (V2') are formulated for a fixed triple u, v, f of two vertices and a face of T. The algorithm has to determine whether there exists a pair u, v for which (V1') and (V2') hold for every face f of T.

The algorithm proceeds in two stages. The first (resp., second) stage reports the set $\mathcal{N}_1$ of all pairs $(u,v) \in V \times V$ for which condition (V1') (resp., (V2')) is violated for some face f. Any pair of vertices $(u,v) \notin \mathcal{N}_1 \cup \mathcal{N}_2$ is a guarding pair of T, and thus constitutes a solution to the decision procedure. If all pairs of vertices are disqualified, the decision procedure has a negative answer.

Computing $\mathcal{N}_2$. Let $\mathcal{R}$ be the set of segments $\varrho(u,s)$ over all pairs $u, s \in V$ where s is visible from u. We compute $\mathcal{R}$ as follows. We fix a vertex $u \in V$ and preprocess T in $O(n \log n)$ time into a data structure, so that the first intersection point of T with a ray emanating from $u(h)$ can be computed in $O(\log n)$ time [19]. For each vertex $s \in V \setminus \{u\}$, we determine the first point ξ hit by the ray $u(h)s$. If ξ lies between s and $u(h)$ on the ray, $\varrho(u,s)$ is not defined; otherwise we set $\varrho(u,s) = s\xi$. We repeat this procedure for all vertices $u \in T$. The total time spent in this step is $O(n^2 \log n)$.

Fix a vertex $v \in V$. We compute in $O(n^2)$ time the visibility map of T from $v(h)$ and thus the set $\mathcal{E}$ of edges of all the regions $H_v(f)$; $|\mathcal{E}| = O(n^2)$. For a geometric object γ in $\mathbb{R}^3$, let γ^* denote its xy-projection. Set $\mathcal{R}^* = \{\varrho^* \mid \varrho \in \mathcal{R}\}$ and $\mathcal{E}^* = \{\gamma \mid \gamma \in \mathcal{E}\}$. Each of $\mathcal{R}^*$ and $\mathcal{E}^*$ is a set of $O(n^2)$ segments in $\mathbb{R}^2$. Using the algorithm described in [15], we compute in $O(n^{8/3}\,\mathrm{polylog}(n))$ time, the set of all intersecting pairs $(\varrho^*, e^*) \in \mathcal{R}^* \times \mathcal{E}^*$ as the disjoint union of complete bipartite graphs, so that the overall size of their vertex sets is $O(n^{8/3}\log n)$.

For each complete bipartite subgraph $\mathcal{R}_i^* \times \mathcal{E}_i^* \subseteq \mathcal{R}^* \times \mathcal{E}^*$ in the output, we output all pairs (u,v), such that $\mathcal{R}_i^*$ contains (the projection of) a segment $\varrho(u,s)$, for some $s \in V$. The total cost of this step is $O(n^{8/3}\,\mathrm{polylog}(n))$, and we repeat it for each $v \in V$, to obtain an output collection $\mathcal{N}_2$ of all pairs $(u,v) \in V \times V$, with the property that the projection of some segment of the form $\varrho(u,s)$ intersects the projection of an edge of some $H_v(f)$. In view of the preceding discussion, no pair (u,v) in $\mathcal{N}_2$ is a guarding pair of T. The total cost of this step is $O(n^{11/3}\,\mathrm{polylog}(n))$.

Computing $\mathcal{N}_1$**.** We wish to report all pairs $(u,v) \in V \times V$ for which there exist an edge $e \in T$ and a face $f \in T$ such that $W_{u,e}^-$ contains a vertex of $H_v(f)$. A vertex $\chi \in H_v(f)$ lies in $W_{u,e}^-$ if and only if (i) the xy-projection χ^* of χ lies inside $W_{u,e}^*$, and (ii) χ lies below the plane $\pi_{u,e}$ containing $W_{u,e}$.

We fix a vertex $v \in T$ and report all pairs $(u,v) \in \mathcal{N}_1$ in two stages. Let Ξ be the set of vertices in $H_v(f)$, and let $\mathcal{W}$ be the set of all wedges $W_{u,e}$, as defined in the beginning of the section. The first stage reports all pairs in $\Xi \times \mathcal{W}$ that satisfy (i). The second stage reports a pair $(u,v) \in V \times V$ if there is a pair $(\chi, W_{u,e})$ reported in the first stage for which χ lies below $\pi_{u,e}$. In more detail, let $\Xi^* = \{\chi^* \mid \chi \in \Xi\}$, and let $\mathcal{W}^* = \{W^* \mid W \in \mathcal{W}\}$; $|\Xi^*|, |\mathcal{W}^*| = O(n^2)$. Each W^* is an unbounded triangle in the xy-plane. Using a triangle range-searching data structure [15], we report in $O(n^{8/3}\,\mathrm{polylog}(n))$ all pairs $(\chi, W) \in \Xi \times \mathcal{W}$ such that $\chi^* \in \mathcal{W}^*$, as the disjoint union of complete subgraphs. That is, we report a family $\mathcal{F} = \{(\Xi_1, \mathcal{W}_1), \ldots, (\Xi_r, \mathcal{W}_r)\}$ where (i) $\Xi_i \subseteq \Xi$ and $\mathcal{W}_i \subseteq \mathcal{W}$, (ii) for any $(\chi, W) \in \Xi_i \times \mathcal{W}_i$, $\chi^* \in W^*$, and for every pair $(\chi, W) \in \Xi \times \mathcal{W}$ such that $\chi^* \in W^*$, there is an index i with $\chi \in \Xi_i$ and $W \in \mathcal{W}_i$. Moreover, $\sum_i(|\Xi_i| + |\mathcal{W}_i|) = O(n^{8/3}\,\mathrm{polylog}\, n)$.

Fix a complete subgraph $(\Xi_i, \mathcal{W}_i) \in \mathcal{F}$. We preprocess Ξ_i, in $O(|\Xi_i|\log|\Xi_i|)$ time, into a data structure so that a halfspace-emptiness query (i.e., determine whether a query halfspace contains any point of Ξ_i) can be answered in $O(\log|\Xi_i|)$ time. This can be accomplished by constructing the Dobkin-Kirkaptrick hierarchy [11] on the convex hull of Ξ_i. For each $W_{u,e} \in \mathcal{W}_i$, we query the data structure with the halfspace $\pi_{u,e}^-$ lying below $\pi_{u,e}$. If $\pi_{u,e}^- \cap \Xi_i \neq \emptyset$, we add the pair (u,v) to $\mathcal{N}_1$. The total time spent over all complete bipartite graphs of $\mathcal{F}$ is $O(n^{8/3}\,\mathrm{polylog}(n))$. Repeating this procedure for all vertices $v \in T$, we construct the set $\mathcal{N}_1$ in $O(n^{11/3}\,\mathrm{polylog}(n))$.

Putting everything together, we can find in $O(n^{11/3}\,\mathrm{polylog}(n))$ time whether T can be guarded by two watchtowers of height at most h placed at two vertices of T. Finally, plugging this procedure to parametric searching, we conclude the following.

THEOREM 3.2. *For a polyhedral terrain in $\mathbb{R}^3$ with n edges the discrete version of the two-watchtower problem can be solved in $O(n^{11/3}\,\mathrm{polylog}(n))$ time.*

4. REFERENCES

[1] A. Aggarwal, B. Chazelle, L. Guibas, C. Ó'Dúnlaing and C. Yap, Parallel computational geometry, *Algorithmica* 3 (1988), 293–328.

[2] B. Ben-Moshe, P. Carmi, and M. J. Katz, Computing all large sums-of-pairs in $\mathbb{R}^n$ and the discrete planar two-watchtower problem, *Inform. Process. Lett.* 89 (2004), 137–139,

[3] B. Ben-Moshe, M. J. Katz, and J.S.B. Mitchell, A constant-factor approximation algorithm for optimal terrain guarding, In *Proc. 16th ACM-SIAM Symp. Discrete Algorithms*, 2005, pp. 515–524.

[4] S. Bespamyatnikh, Z. Chen, K. Wang, and B. Zhu, On the planar two-watchtower problem, In *Proc. 7th Ann. Internat. Conf. Computing and Combinatorics*, pp. 121–130, 2001.

[5] B. Chazelle and L. Guibas, Visibility and intersection problems in plane geometry, *Discrete Comput. Geom.* 4 (1989), 551–581.

[6] D.Z. Chen and O. Daescu, Maintaining visibility of a polygon with a moving point of view, *Inform. Process. Lett.*, 65 (1998), 269–275.

[7] D.Z. Chen, V. Estivvill-Castro and J. Urrutia, Optimal guarding of polygons and monotone chains, *Proc. 7th Canadian Conf. on Comput. Geom.*, 133–138, 1995.

[8] K. Clarkson and K. Varadarajan, Improved approximation algorithms for geometric set cover, these proceedings.

[9] R. Cole, Slowing down sorting networks to obtain faster sorting algorithms, *J. ACM* 34 (1987), 200–208.

[10] R. Cole and M. Sharir, Visibility problems for polyhedral terrains, *J. Symbolic Comput.*, 7 (1989), 11–30.

[11] D. P. Dobkin and D. G. Kirkpatrick, Fast detection of polyhedral intersection, *Theoretical Computer Science* 27 (1983), 241–253.

[12] S.J. Eidenbenz, C. Stamm, and P. Widmayer, Inapproximability results for guarding polygons and terrains. *Algorithmica* 31 (2001), 79–113

[13] L. Guibas, J. Hershberger, D. Leven, M. Sharir and R. Tarjan, Linear-time algorithms for visibility and shortest path problems inside triangulated simple polygons, *Algorithmica* 2 (1987), 209–233.

[14] J. Hershberger, Finding the upper envelope of n line segments in $O(n\log n)$ time, *Inform. Process. Lett.* 33 (1989), 169–174.

[15] J. Matoušek, Range searching with efficient hierarchical cuttings, *Discrete Comput. Geom.*, 10 (1993), 157–182.

[16] N. Megiddo, Applying parallel computation algorithms in the design of serial algorithms, *J. ACM* 30 (1983), 852–865.

[17] M.H. Overmars and J. van Leeuwen, Maintenance of configurations in the plane, *J. Comput. System Sci.* 23 (1981), 166–204.

[18] M. Sharir, The shortest watchtower and related problems for polyhedral terrains, *Inform. Process. Lett.*, 29(5) (1988), 265–270.

[19] M. Sharir and P.K. Agarwal, *Davenport Schinzel Sequences and Their Geometric Applications*, Cambridge University Press, New York, 1995.

[20] J. Urrutia, Art gallery and illumination problems, in *Handbook on Computational Geometry,* Elsevier Science Publishers, J.R. Sack and J. Urrutia, eds., pp. 973-1026, 2000.

[21] B. Zhu, Computing the shortest watchtower of a polyhedral terrain in $O(n\log n)$ time, *Comput. Geom. Theory Appls.*, 8 (1997), 181-193.