ISRAEL JOURNAL OF MATHEMATICS **101** (1997), 125–139

GUARDING GALLERIES WHERE EVERY POINT SEES A LARGE AREA

BY

GIL KALAI*

Institute of Mathematics, The Hebrew University of Jerusalem
Jerusalem 91904, Israel
and
Institute for Advanced Studies
Princeton, N.J. 08540, USA
e-mail: kalai@math.huji.ac.il

AND

JIŘÍ MATOUŠEK**

Department of Applied Mathematics, Charles University
Malostranské nám. 25, 118 00 Praha 1, Czech Republic
e-mail: matousek@kam.mff.cuni.cz

ABSTRACT

We prove a conjecture of Kavraki, Latombe, Motwani and Raghavan that if X is a compact simply connected set in the plane of Lebesgue measure 1, such that any point $x \in X$ sees a part of X of measure at least ε, then one can choose a set G of at most $const\frac{1}{\varepsilon}\log\frac{1}{\varepsilon}$ points in X such that any point of X is seen by some point of G. More generally, if for any k points in X there is a point seeing at least 3 of them, then all points of X can be seen from at most $O(k^3 \log k)$ points.

* Research supported by grants from the Sloan Foundation, the Israeli Academy of Sciences and Humanities, and by G.I.F.

** Research supported by Czech Republic Grant GAČR 201/94/2167 and Charles University grants No. 351 and 361. Part of the work was done while the author was visiting The Hebrew University of Jerusalem.

Received October 17, 1995

1. Introduction

Consider an art gallery of total area 1 in the plane such that every guard wherever located in the gallery can see an area of size at least ϵ. Kavraki, Latombe, Motwani and Raghavan (KLMR) [8] asked if the gallery can be guarded by $f(\varepsilon)$ guards, for some function f of ϵ.

Broder, Dyer, Frieze, Raghavan and Upfal [3] showed that for art galleries, even in dimension n the number of guards can be bounded by a function of ϵ, n and the *diameter* D of the gallery.

We will give an affirmative answer to KLMR's problem for $f(\epsilon) = C\frac{1}{\varepsilon}\log\frac{1}{\varepsilon}$, where C is an absolute constant (which is rather large).

A **gallery** is a compact set in the plane, X. A point $x \in X$ **sees** a point $y \in X$ if the segment xy is fully contained in X. A subset $G \subseteq X$ **guards** X if each point of X is seen by at least one point of G. For a point $x \in X$, denote by $V(x)$ the **visible region** of x in X, that is, the set of all points $y \in X$ seen by X.

THEOREM 1: *Let X be a simply connected gallery of Lebesgue measure $\lambda^2(X) = 1$, and let $\varepsilon > 0$ be a real number such that $\lambda^2(V(x)) \geq \varepsilon$ for all $x \in X$. Then X can be guarded by at most const $\frac{1}{\varepsilon}\log\frac{1}{\varepsilon}$ points. More generally, if X has h holes, and $\lambda^2(V(x)) \geq \varepsilon$ for all $x \in X$, then X can be guarded by $C(h)\frac{1}{\varepsilon}\log\frac{1}{\varepsilon}$ points. (In fact, a random sample of this many points, chosen from the uniform distribution on X given by the Lebesgue measure, guards X with high probability.)*

The proof uses concepts and results of Vapnik and Chervonenkis [17] and Haussler and Welzl [6] (VC-dimension and ε-nets). For a set with holes (i.e. for $h > 0$), we use a Ramsey-type argument which was suggested by Nešetřil in a slightly different context. The proof is also related to the so-called **visibility graphs**; see a remark in section 2.

By our method, we bound $C(h)$ in Theorem 1 by a quite fast-growing function. Kavraki et al. [8] conjectured the bound on the number of guards should be polynomial in h and $1/\varepsilon$. Continuing our work, Valtr [16] proved this conjecture in a very strong form, bounding the number of guards by $O\left(\log_2 h\,\frac{1}{\varepsilon}\log\frac{1}{\varepsilon}\right)$ for h sufficiently large. He also improved the numeric constant in the bound for simply connected galleries considerably (see section 2).

The following is a more general result for simply connected art galleries:

THEOREM 2: *Let X be a simply connected gallery and k an integer such that* *among any k points in X, there are some 3 which can be seen from a single point. Then X can be guarded by at most $O(k^3 \log k)$ points.*

The proof uses a technique developed by Alon and Kleitman [1], employing a fractional Helly theorem and a linear separation theorem.

In section 4 we give an example showing that there exist galleries (with a large number of holes) in which each point sees a constant fraction of the area of the gallery (1/10, say) but arbitrarily many points are needed to guard the whole gallery. In fact, calculation shows that the number of guards in our example grows as $\Omega(\log h)$, where h is the number of holes, hence Valtr's upper bound mentioned above is nearly tight, at least for a fixed ε. We were informed that an example with similar properties was constructed also by Broder, Dyer, Frieze, Raghavan and Upfal [3].

2. VC-dimension of galleries

First we recall definitions and results from [17], [6]. Let $\mathcal{S}$ be a set system on a set X. We say that a subset $A \subseteq X$ is **shattered** (by $\mathcal{S}$) if every possible subset of A can be obtained as the intersection of some $S \in \mathcal{S}$ with A. The **VC-dimension** of $\mathcal{S}$ is the supremum of the sizes of all finite shattered subsets of X.

Let μ be a probability measure on X such that all sets of $\mathcal{S}$ are measurable. A set $N \subseteq X$ is called an ε-net for $\mathcal{S}$ (with respect to μ) if it intersects each $S \in \mathcal{S}$ with $\mu(S) \geq \varepsilon$ ($\varepsilon > 0$ is a real number). Haussler and Welzl [6], extending ideas of Vapnik and Chervonenkis [17], proved the following:

THEOREM 3: *Let X be a set, μ a probability measure on X, and $\mathcal{S}$ a system of measurable sets on X of VC-dimension at most d. Then for any $\varepsilon \in (0,1)$ there exists an ε-net for $\mathcal{S}$ (with respect to μ) of size at most $C(d)\frac{1}{\varepsilon}\log\frac{1}{\varepsilon}$, where the number $C(d)$ depends on d only. A random sample of this size is an ε-net with probability exponentially small in $\frac{1}{\varepsilon}$.*

Let us remark that [6] proves this result for the special case when X is finite and μ is the uniform distribution on X. However, the same proof goes through almost literally for an arbitrary probabilistic measure (see also [17] for a proof of a related result for general probability measures). Another fact we need is as follows:

LEMMA 4 ([13], [14], [17]): *If $\mathcal{S}$ is a set system of VC-dimension d on an n-point* *set, then*

$$|S| \leq \binom{n}{0} + \binom{n}{1} + \cdots + \binom{n}{d}.$$

In particular, if d is fixed, $|S|$ is bounded by a fixed polynomial in n.

Our proof of Theorem 1 is based on the following:

PROPOSITION 5: *Let $X \subseteq \mathbb{R}^2$ be compact and simply connected. Then the VC-dimension of the set system $\mathcal{V}(X) = \{V(x);\, x \in X\}$ is bounded by a constant. More generally, if X has at most h holes, then the VC-dimension of $\mathcal{V}(X)$ is bounded by a function of h.*

Proof of Theorem 1: Consider a gallery X as in Theorem 1. Since each $V(x)$ has Lebesgue measure at least ε, an ε-net for the set system $\mathcal{V}(X)$ (with respect to Lebesgue measure) intersects each $V(x)$, and thus guards X. By Theorem 3 and Proposition 5, an ε-net of size $O_h(\frac{1}{\varepsilon}\log\frac{1}{\varepsilon})$ exists in this situation. ∎

Let us introduce some terminology concerning visibility graphs. If A is a subset of X, we define the **visibility graph of A in X**, denoted by $VG_X(A)$, as the graph with vertex set A and with two distinct points u, v forming an edge iff they see each other (within X). For two sets $A, B \subseteq X$, we define the **bipartite visibility graph of A, B in X**, denoted by $BVG_X(A, B)$, as the bipartite graph (A, B, E) (A, B are the color classes and $E \subseteq A \times B$ is the edge set), where $(a, b) \in E$ iff a and b see each other.

Proof of Proposition 5: Let d be a sufficiently large number, and suppose that there exists a d-point set $A \subseteq X$ shattered by $\mathcal{V}(X)$. This means that for each subset $S \subseteq A$ there exists a point $y_S \in X$ which sees all points of S but no point of $A \setminus S$; put $B = \{y_S;\, S \subseteq A\}$. In such a situation, we say that A is **shattered by B**.

Consider the bipartite visibility $BVG_X(A, B)$. For later use we note that if $G = (U, V, E)$ is any fixed bipartite graph, A is sufficiently large, and is shattered by B, then $BVG_X(A, B)$ contains an isomorphic copy of G as an induced subgraph (with vertices of U mapped into A and vertices of V mapped into B).

Starting with A, B as above, we find a smaller shattered set in a special position. Draw a line thru each pair of points of A. The arrangement of these at most $\binom{d}{2}$ lines has $O(d^4)$ cells (vertices, edges, and open convex polygons), so there is one such cell containing a subset $B' \subseteq B$ of at least $2^d/O(d^4)$ points of B. These points correspond to subsets of A, so they define a set system $\mathcal{S}_1$ on

A. If d_1, the VC-dimension of $\mathcal{S}_1$, were bounded by a constant independent of d, then the number of sets in $\mathcal{S}_1$ would grow at most polynomially with d (by Lemma 4), but we know it grows exonentially, hence d_1 grows to infinity with $d \to \infty$. Thus, we may assume that some subset $A_1 \subseteq A$ is shattered by a subset $B_1 \subseteq B'$, with $d_1 = |A_1|$ large.

By the observation made in the beginning of the proof, we know that the bipartite visibility graph of A_1 and B_1 contains any prescribed bipartite induced subgraph (up to some size). In particular, we can select subsets $B_2 \subseteq B_1$ and $A_2 \subseteq A_1$ such that $d_2 = |B_2|$ is large, $|A_2| = 2^{d_2}$ and B_2 is shattered by A_2 (so we reverse the sides; d_2 can be chosen $\lfloor \log_2 d_1 \rfloor$ in this situation — this is an observation due to Assouad [2]).

Next, we repeat the procedure from the first step of the proof, this time selecting a set $B_3 \subseteq B_2$ of size d_3 (still sufficiently large), and $A_3 \subseteq A_2$, such that B_3 is shattered by A_3 and A_3 lies in a single cell of the arrangement of all lines defined by pairs of points of B_3. This cell must be 2-dimensional (if it were an edge, we would get that all the points of A_3 and B_3 are collinear, which is impossible), so no line determined by two points of A_3 intersects $\operatorname{conv}(B_3)$, and vice versa (in particular, $\operatorname{conv}(A_3) \cap \operatorname{conv}(B_3) = \emptyset$). Hence each point of B_3 sees all points of A_3 within an angle smaller than π, and in the same clockwise angular order; let $\leq_A$ be this linear order of the points of A_3. Similarly we have a common counterclockwise angular order, $\leq_B$, of points of B_3 around any point of A_3.

Let us consider the case of X simply connected. Here it suffices to have $d_3 = 5$. We put $B' = B_3$, and for each $b \in B'$ we consider the point $a = a(b) \in A_3$ which sees all points of B' but b. Let these 5 points form a set $A' \subset A_3$.

Since we have 5 points on each side, we may choose a $b \in B'$ such that b is neither the first nor the last point of B' in the $\leq_B$ ordering, and at the same time $a(b) \in A'$ is not the first or last point in the $\leq_A$ ordering of A'. We get a situation as in Figure 1, namely that b sees both the predecessor a' and the successor a'' of $a(b)$, and $a = a(b)$ sees both the predecessor b' and the successor b'' of b. It is easy to check that the segments ab' and $a'b$ intersect as shown (because of the restrictions on the relative position of A' and B'), and similarly for the segments $a''b$ and ab''. These four segments are contained in X, and since X is simply connected, also the shaded region bounded by the segments must be a part of X, hence a and b see each other — a contradiction.

Next, let X have h holes. We use a Ramsey-type result of Nešetřil and Rödl [10]. An **ordered bipartite graph** is a bipartite graph (U, V, E) with some linear orderings on U and on V. (Nešetřil [private communication] earlier suggested this kind of approach for exhibiting a forbidden induced bipartite subgraph of visibility graphs of simple polygons; see a remark at the end of this section.)

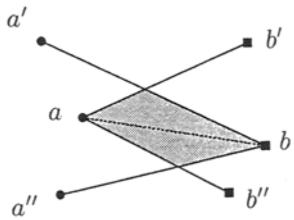

Figure 1. A contradiction to the invisibility of a and b.

LEMMA 6 ([10]): *Let (U, V, E) be a fixed ordered bipartite graph. There exists a bipartite graph (R, S, F) such that for any linear orders $\leq_R$ on R and $\leq_S$ on S, the corresponding ordered bipartite graph contains an ordered induced copy of (U, V, E) (i.e. the induced embedding sends U into R and V into S in an order-preserving manner)**.

In our situation, we let (U, V, E) be the ordered bipartite graph with $U = (u_0, u_1, \ldots, u_{3h+2})$, $V = (v_0, \ldots, v_{3h+2})$, where the subgraph on each two triples $(u_{3i}, u_{3i+1}, u_{3i+2})$ and $(v_{3i}, v_{3i+1}, v_{3i+2})$ $(i = 0, 1, \ldots, h)$ is as the one in Figure 1, i.e. u_{3i+1} is connected to v_{3i}, v_{3i+2} but not to v_{3i+1}, v_{3i+1} is connected to u_{3i} and u_{3i+2}, and the remaining edges do not matter. We choose d_3 so large that the bipartite visibility graph $BVG_X(A_3, B_3)$ contains an induced copy (non-ordered) of the graph (R, S, F) constructed for (U, V, E) as in Theorem 6. In the drawing of $BVG_X(A_3, B_3)$, we then obtain $h+1$ situations as in Figure 1, with the $h+1$ shaded regions being pairwise disjoint. At most h of these regions may contain a hole of X, and the remaining one gives a contradiction to the supposed invisibility as for the simply connected case. ∎

For simply connected galleries, our proof yields a bound of roughly 10^{12} for the VC-dimension. It seems reasonable to conjecture that the VC-dimension is

* As was noted by Noga Alon [private communication], it is easy to check that a sufficiently large random bipartite graph has the required property with a positive probability (but the Nešetřil–Rödl construction can be applied also in more general situations, where a probabilistic proof seems difficult).

in fact a small number, perhaps 6. A floor plan of an art gallery with a 5-point shattered subset is shown in Figure 2. The shattered set is indicated by dots numbered 1 thru 5. Crosses mark points which see only certain subsets; e.g., a cross labeled by 134 sees points 1, 3, and 4 only. All types of subsets up to symmetry are shown, with an exception of the empty set (for which we can always make a tiny niche somewhere in the wall, so that a point there sees no-one).

Recently Valtr [16] found a more complicated example with a 6-point shattered subset and succeeded in improving the upper bound on the VC-dimension to 23.

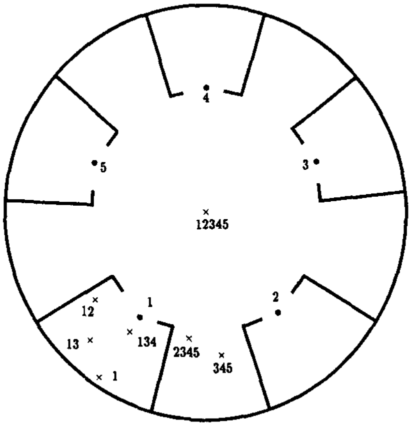

Figure 2. A 5-point shattered set.

Remark: In her thesis [5], Everett asked whether there is a bipartite graph which cannot occur as an induced subgraph of the visibility graph of a simple polygon (i.e. a graph of the form $VG_X(A)$, where X is a simple polygon and A is the set of its vertices). Shermer [15] exhibited such a forbidden bipartite subgraph; in fact he constructed a bipartite graph on 30 vertices which cannot occur as $VG_X(A)$ for any simply connected X and any $A \subseteq X$. For the VC-dimension result we need more — namely a bipartite graph which cannot occur as $BVG_X(A, B)$, i.e. it is not a visibility graph even if we add any subset of edges on A and on B.

3. Fractional Helly and duality

The extra step we need for the proof of Theorem 2 is the following

PROPOSITION 7: *If X and k are as in Theorem 2, then there exists a Borel measure μ on X such that $\mu(V(x)) \geq \varepsilon = k^{-3}$ for all $x \in X$.*

We will first show:

LEMMA 8: *Let X and k be as in Theorem 2. Let $\mathcal{F} = \{V(x) : x \in X\}$, and $\varepsilon = k^{-3}$. Then for any finite collection $F_1, F_2, \ldots, F_n$ of sets in $\mathcal{F}$ and any choice of nonnegative reals $t_1, \ldots, t_n$ with $t_1 + t_2 + \cdots + t_n = 1$ there exists a point $x \in X$ such that*

$$\sum_{i;\, x \in F_i} t_i \geq \varepsilon.$$

Alon and Kleitman proved the assertion of Lemma 8 for a family $\mathcal{F}$ of convex compact sets in the plane so that from every k members of the family 3 have a nonempty intersection. Alon and Kleitman derived their result from a "fractional Helly theorem" of Katchalski and Liu asserting (in a sharp form proved by Kalai and by Eckhoff) that for a family of n convex sets in the plane, if the number of triples of sets in the family having non-empty intersection is $\alpha\binom{n}{3}$ then there is an intersecting subfamily of size δn, with $\delta \geq 1 - (1-\alpha)^{1/3}$. The fractional Helly theorem gives the assertion of the Lemma for $t_i = 1/n$ for $i = 1, 2, \ldots, n$. Applying the fractional Helly theorem for families obtained from $\mathcal{F}$ by taking several copies of each set gives, with some calculations, the case where $t_1, t_2, \ldots, t_n$ are rational numbers, and the general case follows by a limiting argument.

Alon and Kleitman's proof would apply without any change if we can show that finite subfamilies of the family of sets $\{V(x) : x \in X\}$ satisfy the "fractional Helly theorem" (although the set $V(x)$ need not be convex in general). It follows from a theorem of Eckhoff [4] that if $\mathcal{F}$ is a finite family of contractible sets in the plane such that every nonempty intersection of sets in the family is contractible, then the f-vector of the *nerve* of K is equal to the f-vector of the nerve of a family of convex sets. (The f-vector $f = (f_0, f_1, \ldots)$ of the nerve of $\mathcal{F}$ is defined as follows: f_i is the number of subfamilies of $\mathcal{F}$ of size $i+1$ with nonempty intersection.) It follows that the "fractional Helly theorem" applies to families of contractible planar sets such that all nonempty intersections are contractible. We thus need

LEMMA 9: *For any points $x_1, \ldots, x_n \in X$, the set $V(x_1) \cap \cdots \cap V(x_n)$ is either empty or contractible.*

Proof: By a theorem of Molnár [9], it suffices to show that each $V(x)$ is simply connected (which is clear, as $V(x)$ is star-shaped), and that each pairwise intersection of the form $V(x) \cap V(y)$ is connected. This follows easily from the simple connectedness of X. Let $u, v \in X$ be two points in $V(x) \cap V(y)$; by examining the few possible relative positions of u, v, x, and y we find that u and v can always be connected by a path within $V(x) \cap V(y)$; see Figure 3. ∎

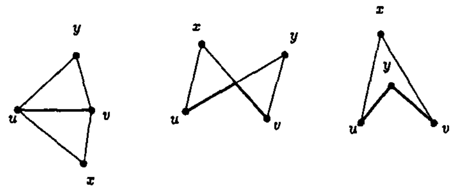

Figure 3. Illustration of the proof of connectedness of $V(x) \cap V(y)$.

The last step in the proof of Proposition 7 also follows the Alon–Kleitman proof. They use the duality theorem of linear programming (which is essentially the separation of disjoint convex sets in a finite-dimensional space by a hyperplane). We need to deal with an infinite family of sets, so we apply an infinite-dimensional separation theorem (alternatively, one could use linear programming duality and a limit argument), in a way suggested to us by E. Matoušková.

LEMMA 10: *Let $\mathcal{F}$ be a family of closed sets in a compact Hausdorff space X, and let $\varepsilon \in (0,1]$ be a real number. Suppose that for any finite collection $F_1, F_2, \ldots, F_n$ of sets in $\mathcal{F}$ and any choice of nonnegative reals $t_1, \ldots, t_n$ with $t_1 + t_2 + \cdots + t_n = 1$ there exists a point $x \in X$ such that*

$$\sum_{i;\, x \in F_i} t_i \geq \varepsilon .$$

Then there exists a Borel measure, μ, on X, such that $\mu(F) \geq \varepsilon$ for all $F \in \mathcal{F}$.

Proof: Let $C(X)$ be the Banach space of all continuous real functions $X \to \mathbb{R}$ with the supremum norm (i.e. $\|f\| = \max_{x \in X} |f(x)|$). For each set $F \in \mathcal{F}$, define a set $D_F \subseteq C(X)$ as the set of all continuous functions $X \to [0,1]$ which are 1 at all points of F. Let $D \subseteq C(X)$ be the convex hull of $\bigcup_{F \in \mathcal{F}} D_F$.

We claim that D contains no function f with $\|f\| < \varepsilon$. Suppose the contrary; this means that there exists a finite convex combination $f = \sum_{i=1}^{n} t_i f_i$, where $t_i \geq 0$, $\sum t_i = 1$, and $f_i \in D_{F_i}$ for some $F_1, \ldots, F_n \in \mathcal{F}$, such that $\|f\| < \varepsilon$. Since each D_{F_i} is convex, we may assume that the F_i's are all distinct. Apply the assumption of the lemma to the collection $F_1, \ldots, F_n$ and the numbers $t_1, \ldots, t_n$ corresponding to this convex combination. This yields a point x with $\sum_{i; x \in F_i} t_i \geq$

ε. If $x \in F_i$, then $f_i(x) = 1$ by the definition of D_{F_i}, so we get

$$\varepsilon > \| f\| \geq f(x) = \sum_{i=1}^{n} t_i f_i(x) \geq \sum_{i;x\in F_i} t_i \geq \varepsilon\,,$$

which is a contradiction.

The convex sets D and $\{f \in C(X);\, \| f\| < \varepsilon\}$ (the open ε-ball) are thus disjoint and the latter one has nonempty interior, hence there exists a hyperplane separating them, that is, a bounded linear functional $h : C(X) \to \mathbb{R}$ such that $h(f) \geq \varepsilon$ for $f \in D$, while $h(f) \leq \varepsilon$ for $\| f\| \leq \varepsilon$ (essentially by the Hahn–Banach theorem; see e.g. [12] for an appropriate version of the separation result and references for other results referred to in the rest of this proof). The latter condition gives $\| h\| \leq 1$, where $\| h\| = \sup\{h(f);\, f \in C(X), \| f\| = 1\}$.

By the Riesz Representation theorem, there exists a unique regular Borel signed measure ν on X such that $h(f) = \int_X f \mathrm{d}\nu$ for each $f \in C(X)$. Let μ be the variation of ν, i.e. the measure defined by $\mu(E) = \sup\{\sum_{i=1}^{k} \nu(E_i)\}$, where $E_1, \ldots, E_k$ are disjoint measurable subsets of a set $E \subseteq X$. We have $\mu(X) = \| h\| \leq 1$.

We claim that $\mu(F) \geq \varepsilon$ for all $F \in \mathcal{F}$. Indeed, if $\mu(F) < \varepsilon$, choose a $G \supseteq F$ open with $\mu(G) < \varepsilon$. Then Tietze's theorem provides a continuous function $f : X \to [0,1]$ which is 0 on $X \setminus G$ and 1 on F, so $f \in D_F$. On the one hand, we should have $h(f) \geq \varepsilon$ by the choice of h, but on the other hand, we have $h(f) = \int_X f \mathrm{d}\nu \leq \int_X f \mathrm{d}\mu \leq \mu(G) < \varepsilon$. This contradiction concludes the proof. (Alternatively, we could add the set $\{f \in C(X);\, f(x) \geq \varepsilon\, \forall x \in X\}$ to D in the beginning; then the functional h provided by the separation is nonnegative and we get a measure right away.) ∎

This finishes the proof of Proposition 7 and thus also the proof of Theorem 2.

4. An example with many holes

Example 11: There exists a constant $\varepsilon_0 > 0$ such that for any integer k there is a gallery X of measure 1, such that each point of X sees an area at least ε_0, and more than k points are needed to guard X.

Proof: We use a probabilistic construction. Let k be given; we choose two sufficiently large integers $n = n(k)$ and $Q = Q(k)$. The construction starts by choosing a trapezoid $ABCD$ and appending n small triangular niches to the top side AB (the construction is illustrated in Figure 4 for $n = 4$). Next, we choose n integers $q_1, \ldots, q_n$ uniformly and independently at random in range $1, \ldots, Q$. We place $q_i + 1$ small triangular holes at the base of the ith niche (see the detail of the first niche in Figure 4), as follows: Assuming that the base is identified with the interval $[0, 1]$, the bases of the holes occupy intervals

$$\left[0, \frac{1}{4q_i}\right], \left[\frac{3}{4q_i}, \frac{5}{4q_i}\right], \ldots, \left[\frac{j}{q_i} - \frac{1}{4q_i}, \frac{j}{q_i} + \frac{1}{4q_i}\right], \ldots, \left[1 - \frac{1}{4q_i}, 1\right].$$

The angles at the top vertices of the niches are $\pi/2$, and the angles at the top vertices of the holes are $\pi/3$, say. This finishes the construction of the gallery X.

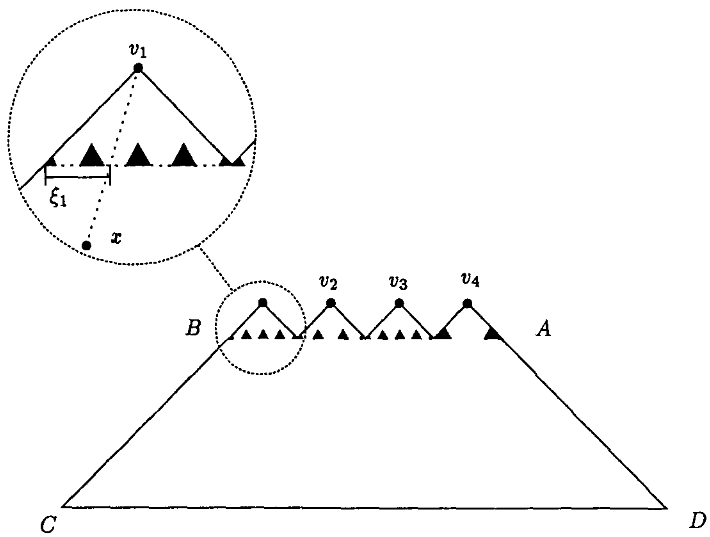

Figure 4. An example requiring many guards.

One can check that every point of the gallery sees at least some constant proportion ε_0 of the area; we omit the details of this. It remains to show that, with a positive probability, X cannot be guarded by k guards. We show that even the vertices $v_1, \ldots, v_n$ cannot be guarded by k guards.

Let x be a fixed point of the trapezoid $ABCD$. Let p_i be the intersection of the line AB with the line $v_i x$, and let ξ_i be the x-coordinate of p_i, where the coordinate system is chosen so that the base of the ith niche occupies the interval $[0, 1]$. If we have $|m_i \xi_i - q_i| \leq \frac{1}{4}$, for some integer m_i, then x cannot see v_i. If we consider k points $x_1, \ldots, x_k \in ABCD$, then the probability (over a random choice of the q_i) that none of x_j sees v_i is at least the probability that there exist integers $m_{i1}, \ldots, m_{ik}$ with $|m_{ij}\xi_{ij} - q_i| \leq \frac{1}{4}$, $j = 1, 2, \ldots, k$ (here ξ_{ij} corresponds to x_j similarly as ξ_i corresponds to x).

If $\xi_{i1}, \dots, \xi_{ik}$ are arbitrary real numbers, a theorem on simultaneous approximation by rationals (see, e.g., [7]) guarantees that for any given natural number Q there exist $q_i \in \{1, 2, \dots, Q\}$ and integers $m_{i1}, \dots, m_{ik}$ with $|m_{ij}\xi_{ij} - q_i| \le Q^{-1/k}$. Hence if we let $Q = 4^k$, we get that for any fixed k-tuple $x_1, \dots, x_k$, the probability that one particular v_i is guarded by at least one x_j is no more than $1 - 1/Q$. Since the choices of the q_i are independent, the probability that all the v_i are guarded by any fixed k-tuple $x_1, \dots, x_k \in X$ is at most $(1 - 1/Q)^{n-k}$ (since at most k points x_j can be placed inside the niches, and such x_j only see one v_i each).

We now want to bound the probability that there exists any placement of $x_1, \dots, x_k$ at all guarding all the v_i. For every i, we can divide the trapezoid $ABCD$ into at most $(2Q)^2$ angular sectors in such a way that points placed in one sector either all see v_i or none does, for any choice of q_i. Hence the number of possibly nonequivalent placements of a single point within $ABCD$ is no larger than the number of cells in an arrangement of $n.4Q^2$ lines, which is bounded by $5n^2Q^4$ (say). The number of nonequivalent positions for a k-tuple of points is then at most $(5n^2Q^4)^k$. If n is chosen so large that $(5n^2Q^4)^k(1 - 4^{-k})^{n-k} < 1$, then the probability that X can be guarded by k points is smaller than 1 as claimed. ∎

5. A remark on a greedy algorithm

One might suspect that under the conditions of Theorem 1, a guarding set of a size bounded in terms of ε could be obtained by a greedy algorithm: Select a guard which sees the maximum possible area, then select the second guard as one seeing the largest part of the area not seen by the first guard, etc. We present an example that this procedure might fail, i.e. select arbitrarily many guards for some simply connected galleries. An example of such a gallery is depicted in Figure 5.

The boundary of the gallery is drawn by a full line, the dotted lines are only auxiliary. The little spikes ("fins") $F_1, F_1', F_2, F_2', \dots$ are chosen so that the area of F_i and F_i' is much larger than the area of F_{i+1} and F_{i+1}'. The guards placed at A and B suffice to guard all the gallery. However, the first greedily placed guard comes to the point G_1, where it sees both F_1 and F_1' and the largest possible piece of the other fins (all points of the gallery, except for the fins, see everything but possibly parts of the fins, and since F_1, F_1' dominate, we look for a position where both can be seen). Now only the shaded parts of the other fins remain unguarded, with the dominating portion of the area being in F_2, F_2', so the next guard is placed in G_2, etc.

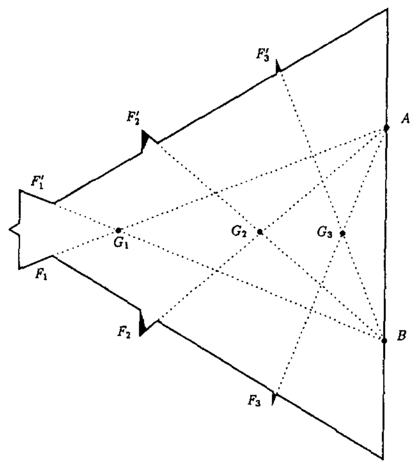

Figure 5. The greedy algorithm fails.

The gallery in the figure requires 3 greedily placed guards. It is problematic to actually draw examples of this type forcing the greedy algorithm to place more guards, but the construction method is extended easily. Namely, we start with fins F_1, F_1' much smaller and much closer to the tip of the large triangle, and then we adjoin progressively smaller fins along the sides of the triangle, the next pair always coming to the right of the intersection of the lines connecting the previous pair to A and B (as in Figure 5). Points in each fin F_i see the portion of the triangle above the horizontal level of A and to the right of the vertical level of the last fin, F_k. If F_k is placed sufficiently far from the vertical side of the triangle (i.e., if we start close enough to the tip with the first pair of fins), this represents a constant fraction of the area (we can get any fraction below $\frac{1}{2}$ by adjusting the proportions appropriately).

ACKNOWLEDGEMENT: We thank Eva Matoušková for help with Lemma 10. Thanks go also to Nati Linial and Jaroslav Nešetřil for useful discussions, and to Hazel Everett for pointing out the reference [15].

References

[1] N. Alon and D. Kleitman, *Piercing convex sets and the Hadwiger Debrunner* (p,q) *problem*, Advances in Mathematics **96** (1992), 103–112.

[2] P. Assouad, *Densite et dimension*, Annales de l'Institut Fourier (Grenoble) **33** (1983), 233–282.

[3] A. Broder, M. Dyer, A. Frieze, P. Raghavan and E. Upfal, unpublished.

[4] J. Eckhoff, *Über kombinatorisch-geometrische Eigenschaften von Komplexen und Familien konvexer Mengen*, Journal für die reine und angewandte Mathematik **313** (1980), 171–188.

[5] H. Everett, *Visibility graph recognition*, PhD thesis, Technical Report 231/90, University of Toronto, 1990.

[6] D. Haussler and E. Welzl, *Epsilon-nets and simplex range queries*, Discrete and Computational Geometry **2** (1987), 127–151.

[7] E. Hlawka, J. Schoißengeier and R. Taschner, *Geometric and Analytic Number Theory*, Springer-Verlag, Berlin, 1991.

[8] L. Kavraki, J-C. Latombe, R. Motwani and P. Raghavan, *Randomized processing in robot path planning*, Proceedings of the 27th ACM Symposium on Theoretical Computation, 1995, pp. 353–362.

[9] J. Molnár, *Über den zweidimensionalen topologischen Satz von Helly*, Matematikai Lapok **8** (1957), 108–114.

[10] J. Nešetřil and V. Rödl, *Simple proof of the existence of restricted Ramsey graphs by means of a partite construction*, Combinatorica **1** (1981), 199–202.

[11] J. O'Rourke, *Computational geometry column 18*, International Journal of Computational Geometry and Applications **3** (1993), 107–113; also in SIGACT News **24**:1 (1993), 20–25.

[12] W. Rudin, *Functional Analysis*, 2nd edn., McGraw-Hill, New York, 1991.

[13] N. Sauer, *On the density of families of sets*, Journal of Combinatorial Theory. Series A **13** (1972), 145–147.

[14] S. Shelah, *A combinatorial problem, stability and order for models and theories in infinitary languages*, Pacific Journal of Mathematics **41** (1972), 247–261.

[15] T. Shermer, *Several short results in the combinatorics of visibility*, Technical Report 91-02, School of Computing Science, Simon Fraser University, British Columbia.

[16] P. Valtr, *Guarding galleries where no point sees a small area*, KAM Series (technical report) 96–305, Department of Applied Mathematics, Charles University, Prague, 1996.

[17] V. N. Vapnik and A. Ya. Chervonenkis, *On the uniform convergence of relative frequencies of events to their probabilities*, Theory of Probability and its Applications **16** (1971), 264–280.