Improved Bounds on Weak ε-Nets for Convex Sets *

Bernard Chazelle† Herbert Edelsbrunner‡ Michelangelo Grigni†
Leonidas Guibas§ Micha Sharir¶ Emo Welzl‖

June 27, 1994 (to appear in DCG)

Abstract

Let S be a set of n points in $\mathbb{R}^d$. A set W is a *weak ε-net* for (convex ranges of) S if for any $T \subseteq S$ containing εn points, the convex hull of T intersects W. We show the existence of weak ε-nets of size $O\left(\frac{1}{\varepsilon^d} \log^{\beta_d} \frac{1}{\varepsilon}\right)$, where $\beta_2 = 0$, $\beta_3 = 1$, and $\beta_d \approx 0.149 \cdot 2^{d-1}(d-1)!$, improving a previous bound of Alon *et al.* We present a deterministic algorithm for computing such a net in time $n(1/\varepsilon)^{O(1)}$. We also consider two special cases: when S is in convex position, we prove the existence of a net of size $O(\frac{1}{\varepsilon} \log^{1.6} \frac{1}{\varepsilon})$; for the case where S consists of the vertices of a regular polygon, we use an argument from hyperbolic geometry to exhibit an optimal net of size $O(1/\varepsilon)$.

*Work by Bernard Chazelle has been supported by NSF Grant CCR-90-02352 and the Geometry Center. Work by Herbert Edelsbrunner has been supported by NSF Grant CCR-89-21421. Work by Michelangelo Grigni has been supported by National Science and Engineering Research Council Operating Grants and BC Advanced Systems Institute Fellowship Awards. Work by Leonidas Guibas and Micha Sharir has been supported by a grant from the U.S.-Israeli Binational Science Foundation. Work by Emo Welzl and Micha Sharir has been supported by a grant from the G.I.F., the German-Israeli Foundation for Scientific Research and Development. Work by Micha Sharir has also been supported by NSF Grant CCR-91-22103, and by a grant from the Fund for Basic Research administered by the Israeli Academy of Sciences.

†Department of Computer Science, Princeton University

‡Computer Science Department, University of Illinois at Urbana-Champaign

§DEC Systems Research Center, Laboratory for Computer Science, M.I.T., and Computer Science Department, Stanford University

¶School of Mathematical Sciences, Tel Aviv University and Courant Institute of Mathematical Sciences, New York University

‖Institut für Informatik, Freie Universität Berlin

1 Introduction

Let S be a set of n points in $\mathbb{R}^d$. A set $W \subset \mathbb{R}^d$ is called a *weak ε-net* for (convex ranges of) S if, for any subset T of εn points of S, the convex hull of T intersects W. If the points of W could be chosen from among those of S, then W would have been called a *strong* ε-net (or just an ε-net) of S. Such nets, introduced by Haussler and Welzl [11], are defined for general *range spaces*, where a range space is a pair $(S, \mathcal{R})$, where S is a set and $\mathcal{R}$ is a set of subsets of S, called *ranges*. A set $N \subseteq S$ is an ε-net if every range of S that contains at least εn points intersects N. Haussler and Welzl have shown that if the range space has *finite VC dimension* then there always exists an ε-net of size $O(\frac{1}{\varepsilon}\log\frac{1}{\varepsilon})$; in particular, this size is independent of the size of S. Moreover, a random subset of S of this size will be an ε-net with high probability. See [11] for more details.

In the setup that we are concerned with, the ranges are all intersections of S with convex sets; we will refer to these ranges as *convex ranges*. Unfortunately, in this setup, the resulting range space has infinite VC dimension, which, in particular, is manifested by the fact that in general strong ε-nets must be very large, that is, contain more than $(1-\varepsilon)n$ points of S (this is the case if all points of S lie in convex position).

Alon *et al.* [1] (see also [2]) have recently shown that there always exist weak ε-nets for convex ranges, which have size $O(1/\varepsilon^{(d+1)(1-\frac{1}{s_d})})$, where $s_d = (4d+1)^{d+1}$. They remarked that their proof method cannot give a selection exponent s_d smaller than $(d+1)^{d+1}$. In the plane they obtained an improved bound of $O(1/\varepsilon^2)$ and they gave an $O(n^2)$ time algorithm for finding a weak net of that size. They also provided a faster algorithm, which finds a weak ε-net consisting of $O(\epsilon^{-4.88})$ points in time $O(n\log(1/\epsilon))$. In this paper we improve their bound for $d \geq 3$, showing the existence of weak ε-nets of size $O\left(\frac{1}{\varepsilon^d}\log^{\beta_d}\frac{1}{\varepsilon}\right)$, where $\beta_2 = 0$, $\beta_3 = 1$, and $\beta_d \approx 0.149 \cdot 2^{d-1}(d-1)!$. Our analysis actually produces a more powerful structure. Namely, we obtain a collection Q of $O(n^d)$ points (that depend on S but not on ε), so that, for any given $\varepsilon > 0$ we can obtain a weak ε-net by an appropriate sampling of points of Q. Moreover, our net has the following stronger property: For any subset T of S of at least εn points, there exists a net point which is an *approximate center point* of T. A point p is called an approximate center point of a point set T if every halfspace that contains p also contains at least some fixed fraction of the points of T. For a real center point, which always exists, this fraction is at least $1/(d+1)$; see [8]. In our case, the fraction is a function of ε, namely it is at least $\Omega(1/\log^{\beta_d/d}\frac{1}{\varepsilon})$.

We also consider the algorithmic problem of efficiently constructing a weak ε-net for a given set S. We present a deterministic algorithm that computes a net in time $n(1/\varepsilon)^{O(1)}$. Finally, we look at two special cases: If S consists of points in convex position, then a net of size $O(\frac{1}{\varepsilon}\log^{1.6}\frac{1}{\varepsilon})$ can be found. If the points of S lie uniformly on a circle (as do, say, the vertices of a regular polygon) then there exists an optimal ε-net of size $O(1/\varepsilon)$. Interestingly, the ε-net can be chosen as a subset of the vertices of a tesselation of the hyperbolic plane. The simplicity of the proof shows that the problem lives naturally in hyperbolic space.

The paper is organized as follows. In Section 2 we present our construction of weak ε-nets. Section 3 is devoted to the algorithmic issues of the construction. Section 4 and 5 treat the two special cases mentioned above.

2 Construction of Weak ε-Nets

Let S be a given set of n points in $\mathbb{R}^d$. We will assume, with no loss of generality, that the points of S lie in general position, meaning that their coordinates are nd real numbers that are *algebraically independent* over the rationals. If the points of S are not in general position, then we can choose a sequence of sets $\{S_i\}$ that converge to S pointwise and are in general position. Let W_i be a weak ε-net for S_i, all having the same size. By compactness, a subsequence of the W_i converges pointwise to a set W, having no more points than each of the W_i's, which is easily seen to be a weak ε-net for S.

We will construct a multi-level structure from the points of S, as follows. At the first level, we project the points of S on the x_1-axis, and denote by S_1 the resulting set. We consider the set of all intervals on the x_1-axis connecting pairs of points in S_1, and construct an *interval tree* over these intervals. In more detail, this is a binary tree whose root corresponds to a point a on the x_1-axis, such that at most half of the intervals lie fully to the left of a and at most half lie fully to its right. The intervals that contain a are stored at the root. The left (resp. right) subtree of the interval tree is obtained recursively by applying the same bisection step to the intervals lying fully to the left (resp. to the right) of a.

This definition of an interval tree is general, and applies to arbitrary collections of intervals on a line. In this first-level structure, though, the same tree can also be constructed by building a balanced binary tree on the projected points of S, and by storing each interval I at the unique node of the tree whose left subtree stores one endpoint of I and its right subtree stores the other endpoint. In what follows we will make use of this alternative representation.

Note that each interval is stored in exactly one node of the tree. Note also that, interpreting the structure back in $\mathbb{R}^d$, each node of the tree corresponds to a hyperplane of the form $x_1 = a$, and the intervals that are stored at the node are x_1-projections of those segments that connect pairs of points in S and cross the hyperplane.

Now consider the second level of our structure. Let v be a node of the first-level interval tree, corresponding to a hyperplane $f_v : \ x_1 = a_v$, and let N_v denote the set of segments stored at v. Map each segment pq in N_v to the point $pq \cap f_v$, and label the point by the unordered pair (pq). Let K_v denote the resulting collection of points. Consider now the collection of all segments in f_v that connect pairs of points of K_v of the form (pq), (pr) (i.e. pairs of points whose labels share a common point of S). We project these segments onto the x_2-axis within f_v, and construct an interval tree for these projected segments, in the same manner as described above. A node of this second-level tree corresponds to a $(d-2)$-flat f of the form $x_1 = a_v$, $x_2 = b$, and the intervals $[(pq), (pr)]$ stored at that node correspond to triangles pqr that are spanned by points of S and intersect f. It is easily verified that each triangle pqr spanned by S is stored in exactly one second-level node (over all first-level trees). To see this, suppose that pq is the edge whose x_1-projection is longest. Then the first-level node v where pq is stored must also store (exactly) one of pr, qr, say it stores pr. Then the segment $[(pq), (pr)]$ is one of the segments processed at v and thus is stored somewhere in the interval tree of v. Uniqueness also follows easily from this argument.

We continue in this manner, constructing one extra level of the structure for each dimension. Consider the j-th level of the structure, for $j \le d-1$. A node v of a $(j-1)$-level interval tree corresponds to a $(d-j+1)$-flat f_v of the form $x_1{=}a_1, \ldots, x_{j-1}{=}a_{j-1}$, and each segment stored at v corresponds to a $(j-1)$-simplex $p_1p_2\cdots p_j$ that is spanned by S and intersects f_v at a point. Let N_v denote the set of all these simplices and K_v denote the set of intersection points of these simplices with f_v, each labeled by the unordered set of vertices of the corresponding simplex. We form the collection of all segments within f_v that connect pairs of points of K_v of the form $(p_1\cdots p_{j-1}q)$, $(p_1\cdots p_{j-1}r)$, project these segments onto the x_j-axis (within f_v), and construct an interval tree on this set of projected segments. Again, a node u of the resulting tree corresponds to a $(d-j)$-flat f_u within f_v, and each segment that u stores, having the form $[(p_1p_2\cdots p_{j-1}p_j), (p_1p_2\cdots p_{j-1}p_{j+1})]$, corresponds to the j-simplex $p_1p_2\cdots p_{j+1}$, which is easily seen to intersect f_u.

Let v be a node of some interval tree of the last level $d-1$. The node v stores a list N_v of $(d-1)$-simplices that cross the line f_v associated with v, and a corresponding list K_v of the intersection points of these simplices with f_v. We sort the list K_v in increasing order of the x_d-coordinates of its points. Let Q denote the union of all lists K_v, over all nodes v of the last-level interval trees; note Q does not depend on ε. Given ε, we now describe a sampled subset W of Q which will be a weak ε-net.

Let v_1 denote the node of the first-level interval tree in whose substructure v lies, and suppose that v_1 lies at depth $\ell \ge 0$ in its tree.

We define a sequence $\{\beta_j\}_{j\ge 2}$ of integers by $\beta_2 = 0$ and $\beta_{j+1} = 2j\beta_j + 1$, for $j \ge 2$. If we put $\beta_j = 2^{j-1}(j-1)!\xi_j$, then we have $\xi_2 = 0$, and

$$\xi_{j+1} = \xi_j + \frac{1}{2^j j!} ,$$

so that

$$\xi_j = \sum_{k=2}^{j-1} \frac{1}{2^k k!} \approx \sqrt{e} - 1.5 < 0.149 .$$

We now put $\varepsilon_v = \frac{1}{2}\left(\frac{3}{4}\right)^{\ell}\varepsilon$, $M_v = \frac{c}{\varepsilon_v^d}\log^{\beta_d}\frac{1}{\varepsilon_v}$, for an appropriate constant c depending on d that will be determined later, and we sample every (n^d/M_v)-th point of

K_v. If K_v has fewer than M_v points, we do not sample any point of K_v. Let W be the union of all such samples from Q, over all K_v.

Let σ be a simplex spanned by the points of S, and let u be the node of smallest depth ℓ of the first-level interval tree which stores some edge of σ. It is easily verified that all vertices of σ are stored at the subtree rooted at u. Let S_u denote the set of points stored at this subtree. The size of S_u is $n/2^\ell$. Next we claim that each $(d-1)$-simplex spanned by the points of S_u is stored in at most a constant number of last-level lists K_v. This is best proved by showing, using induction on j, that each j-simplex spanned by S_u is stored in at most m_j j-level nodes, for an appropriate constant m_j. It follows that the total number of sampled points at lists K_v in the substructure of the same first-level node u is at most

$$O\left(\frac{\left(\frac{n}{2^\ell}\right)^d}{\frac{n^d}{M_v}}\right) = O\left(\frac{1}{2^{d\ell}} \cdot \frac{1}{\varepsilon_v^d} \log^{\beta_d} \frac{1}{\varepsilon_v}\right) =$$

$$O\left(\left(\frac{2}{3}\right)^{d\ell} \cdot \frac{1}{\varepsilon^d}\left[\log\frac{1}{\varepsilon} + \ell \log\frac{4}{3}\right]^{\beta_d}\right) .$$

We now sum these bounds over all 2^ℓ first-level nodes u at the same depth ℓ, and then sum over ℓ, to obtain an overall bound of

$$\sum_{\ell \geq 0} O\left(\left[2\left(\frac{2}{3}\right)^d\right]^\ell \frac{1}{\varepsilon^d}\left[\log\frac{1}{\varepsilon} + \ell\log\frac{4}{3}\right]^{\beta_d}\right) .$$

Since $d \geq 2$, we have $2\left(\frac{2}{3}\right)^d < 1$, so the sum is easily seen to be dominated by its leading term $\ell = 0$, which implies:

Lemma 2.1 *The set W consists of at most $O\left(\frac{1}{\varepsilon^d}\log^{\beta_d}\frac{1}{\varepsilon}\right)$ points.*

Lemma 2.2 *W is a weak ε-net for S.*

Proof: Let T be a subset of S consisting of εn points; we need to show that $conv(T) \cap W \neq \emptyset$. We will proceed through the structure level by level, but, for technical reasons, we give the first level separate treatment. Let v_1 be a node of the first-level interval tree of smallest depth $\ell \geq 0$, such that at least $\frac{1}{4}\left(\frac{3}{4}\right)^\ell \varepsilon n$ points of T are stored at each of the two subtrees of v_1. Such a node must exist, for otherwise we would obtain a single path π in the tree, so that the node of π at depth ℓ stores at least $\left(\frac{3}{4}\right)^\ell \varepsilon n$ points of T in its subtree. However, the number of points of S stored at that subtree is at most $\frac{n}{2^\ell}$, which is smaller than $\left(\frac{3}{4}\right)^\ell \varepsilon n$ when ℓ is sufficiently large. Put $\varepsilon_0 = \frac{1}{2}\left(\frac{3}{4}\right)^\ell \varepsilon$. Let T_0 denote the subset of T consisting of those points of T stored at the subtree of v_1. By removing some points from T_0, if necessary, we will assume that the size of T_0 is exactly $\varepsilon_0 n$, and that exactly half the points of T_0 are stored at the subtree of each child of v_1.

We claim that at each level j there exists some node v_j whose associated list N_{v_j} contains at least $c_j \varepsilon_0^{j+1} n^{j+1} / \log^{\beta_{j+1}} \frac{1}{\varepsilon_0}$ j-simplices spanned by the points of T_0, for some positive constant c_j (with v_1 being the node just defined). This will imply the existence of a node v_{d-1} at the last level $d-1$, whose list $N_{v_{d-1}}$ contains at least $c_{d-1}\varepsilon_0^d n^d / \log^{\beta_d} \frac{1}{\varepsilon_0}$ $(d-1)$-simplices spanned by the points of T_0. All these simplices are contained in $conv(T)$, and $conv(T)$ intersects the line $f_{v_{d-1}}$ in an interval I which therefore contains all the points of $K_{v_{d-1}}$ corresponding to these simplices. Since v_{d-1} lies at a substructure of the node v_1, it follows by construction (if we choose the constant c in the definition of M_v to be c_{d-1}), that I, hence $conv(T)$, must contain a point of W.

To show the existence of the nodes v_j, we argue as follows. We will make use of the following elementary Selection Lemma:

Lemma 2.3 (Selection Lemma [3, 6]) *Given a set N of n points on the line, and a set M of m intervals delimited by the points of N, there exists some point on the line that is contained in at least $\frac{m^2}{4n^2}$ intervals of M.*

We now proceed by induction on the level j. For $j = 1$, since $\beta_2 = 0$, we need to show that v_1 stores at least $c_1 \varepsilon_0^2 n^2$ segments connecting pairs of points of T. This follows, with $c_1 = \frac{1}{4}$, from the fact that both the left and right subtrees of v_1 store $\frac{1}{2}\varepsilon_0 n$ points of T.

For the sake of exposition, we treat separately the case $j = 2$. Let E be the set of intervals spanned by the x_1-projections of the points of T_0 and stored at the node v_1 obtained above. Let $t = |E| \geq c_1 \varepsilon_0^2 n^2$. Regard E as the edge set of an undirected graph, whose nodes are the points of T_0.

We claim that by deleting no more than half the elements of E (and some points of T_0) we can guarantee that every remaining point of T_0 has degree $> \frac{t}{2\varepsilon_0 n}$. This is proved by a simple pruning process, that iteratively removes a point and all its incident edges if the point has degree smaller than or equal to $\frac{t}{2\varepsilon_0 n}$; this process cannot remove more than $\frac{t}{2}$ edges of E.

Now consider the resulting pruned set E' as a set of points in the hyperplane $f_{v_1} : x_1 = a_1$ corresponding to v_1 (thus E' is a subset of K_{v_1}). We construct a set $\mathcal{M}$ of segments in f_{v_1}, as follows. Take each point $(pq) \in E'$, choose any point $r \neq p, q$ of T_0 such that (qr) is also in E', then choose any point $s \neq p, q, r$ of T_0 such that (rs) is also in E', and add to $\mathcal{M}$ the segment connecting (pq) to (rs) in f_{v_1}. The pruning procedure ensures that the number of segments in $\mathcal{M}$ is at least $|E'| \cdot \left(\frac{t}{2\varepsilon_0 n}\right)^2$. Now apply the Selection Lemma to E' and $\mathcal{M}$, projected onto the x_2-axis within f_{v_1}, to obtain a point $x_2 = a_2$ contained in at least

$$\frac{1}{4}\left(\frac{t}{2\varepsilon_0 n}\right)^4 \geq \frac{1}{64} c_1^4 \varepsilon_0^4 n^4 = c' \varepsilon_0^4 n^4$$

projected segments of $\mathcal{M}$. Clearly, if $[(pq), (rs)]$ is a segment of $\mathcal{M}$ whose x_2-projection contains a_2, which is formed through the intermediate point (qr), then a_2 is also contained in the x_2-projection of either $[(pq),(qr)]$ or $[(qr),(rs)]$ (or both). Hence a_2 is contained in at least $c'\varepsilon_0^4 n^4$ projected segments of the latter kind, which, by construction, are all stored in the second-level interval tree of f_{v_1}. However, these segments are not necessarily distinct, and each may be counted with multiplicity at most $2\varepsilon_0 n$ (a segment $[(pq),(qr)]$ may be counted once for each point s of T_0 that induces a segment $[(qr),(rs)]$, and once for each point s that induces a segment $[(ps),(pq)]$). Hence, a_2 is contained in at least $\frac{1}{2}c'\varepsilon_0^3 n^3$ distinct projected segments of this kind.

Let π denote the path in the interval tree of f_{v_1} leading to a_2, that is, the path descends from a node u to its left (right) child if a_2 is smaller (larger) than the x_2-value associated with u. It is easily verified that each projected segment containing a_2 must be stored at some node along π. Moreover, the total number of projected segments stored at nodes of π at depth $\geq 3\log\frac{1}{\varepsilon_0} + \log\frac{1}{c'} + 2$ is at most $\frac{1}{4}c'\varepsilon_0^3 n^3$, so at least half of the projected segments containing a_2 are stored at higher nodes along π, which implies that some node v_2 of π stores at least

$$\frac{\frac{1}{4}c'\varepsilon_0^3 n^3}{3\log\frac{1}{\varepsilon_0} + \log\frac{1}{c'} + 2} \geq \frac{c_2\varepsilon_0^3 n^3}{\log\frac{1}{\varepsilon_0}}$$

such projected segments, for an appropriate positive constant c_2. In other words, the $(d-2)$-flat f_{v_2} associated with v_2 crosses at least $c_2\varepsilon_0^3 n^3/\log\frac{1}{\varepsilon_0}$ distinct triangles spanned by the points of T_0, all of which are passed to the third-level substructure of v_2.

The general inductive step at a level j is argued in much the same way as in the case $j=2$. That is, we consider the set E of $(j-1)$-simplices spanned by the points of T_0 and stored at the node v_{j-1} produced at the preceding induction step. By induction hypothesis, $t = |E| \geq c_{j-1}\varepsilon_0^j n^j/\log^{\beta_j}\frac{1}{\varepsilon_0}$. We regard E as the edge set of an unordered j-hypergraph, whose nodes are the points of T_0. We say that a $(j-1)$-set $\{p_1 p_2 \cdots p_{j-1}\}$ is *present* in the hypergraph if the hypergraph contains at least one edge that contains that set.

We claim that, by deleting no more than half the elements of E, we can guarantee that for every $(j-1)$-set that is still present in the pruned hypergraph, there are at least $\frac{t}{2\binom{\varepsilon_0 n}{j-1}}$ remaining edges containing that set. The proof is similar to that used in the case $j=2$: iteratively remove a $(j-1)$-set and all its containing edges if the number of such edges is smaller than $\frac{t}{2\binom{\varepsilon_0 n}{j-1}}$; this process cannot remove more than $\frac{t}{2}$ edges of E, since there are at most $\binom{\varepsilon_0 n}{j-1}$ distinct $(j-1)$-sets, and each can be removed (with its containing edges) at most once.

Now consider the resulting pruned hypergraph E' as a set of points in the flat $f_{v_{j-1}}$ corresponding to v_{j-1} (thus E' is a subset of $K_{v_{j-1}}$). We construct a set $\mathcal{M}$ of segments in $f_{v_{j-1}}$, as follows. Take each point $(p_1 p_2 \cdots p_j) \in E'$, choose any point $q_1 \neq p_1,\ldots,p_j$ of T_0 such that $(p_2 \cdots p_j q_1)$ is also in E', then choose any point $q_2 \neq p_1,\ldots,p_j,q_1$ of T_0 such that $(p_3 \cdots p_j q_1 q_2)$ is also in E', continue in this manner until j new points $q_1,\ldots,q_j$ are chosen (so the last edge is $(q_1 q_2 \cdots q_j) \in E'$), and add to $\mathcal{M}$ the segment connecting $(p_1p_2\cdots p_j)$ and $(q_1q_2\cdots q_j)$ in $f_{v_{j-1}}$. The pruning procedure ensures that the number of segments in $\mathcal{M}$ is at least $|E'|\cdot\left(\frac{t}{2\binom{\varepsilon_0 n}{j-1}}-j\right)^j$. Now apply the Selection Lemma to E' and $\mathcal{M}$, projected onto the x_j-axis within $f_{v_{j-1}}$, to obtain a point $x_j = a_j$ contained in at least

$$\frac{1}{4}\left(\frac{t}{2\binom{\varepsilon_0 n}{j-1}}-j\right)^{2j} \geq \frac{c'\varepsilon_0^{2j}n^{2j}}{\log^{2j\beta_j}\frac{1}{\varepsilon_0}}$$

projected segments of $\mathcal{M}$, for an appropriate positive constant c', depending on j and on c_{j-1}. Clearly, if $[(p_1p_2\cdots p_j),\ (q_1q_2\cdots q_j)]$ is a segment of $\mathcal{M}$ whose x_j-projection contains a_j, which is formed, say, through the chain of 'point replacements' used above, then at least one of the segments

$$[(p_1p_2\cdots p_j),(p_2\cdots p_jq_1)],\ [(p_2\cdots p_jq_1),(p_3\cdots p_jq_1q_2)],\ \ldots,\ [(p_jq_1q_2\cdots q_{j-1}),(q_1q_2\cdots q_j)]$$

must have an x_j-projection that also contains a_j. Hence a_j is contained in at least $c'\varepsilon_0^{2j}n^{2j}/\log^{2j\beta_j}\frac{1}{\varepsilon_0}$ projected segments of the latter kind, which, by construction, are all stored in the j-level interval tree constructed within $f_{v_{j-1}}$. As above, these segments are not necessarily distinct, and each may be counted with multiplicity at most $O(\varepsilon_0^{j-1}n^{j-1})$ (which is the number of times such a segment can be extended, via the point-replacement mechanism described above, to a segment connecting two points of E' whose labels share no point of T_0). Hence, a_j is contained in at least $c''\varepsilon_0^{j+1}n^{j+1}/\log^{2j\beta_j}\frac{1}{\varepsilon_0}$ distinct segments of this kind, for another constant c''. We now repeat the argument involving the path in the interval tree leading to a_j, truncated at depth roughly $(j+1)\log\frac{1}{\varepsilon_0}$, to conclude that there exists a node v_j in that interval tree whose associated $(d-j)$-flat f_{v_j} crosses at least $c_j\varepsilon_0^{j+1}n^{j+1}/\log^{2j\beta_j+1}\frac{1}{\varepsilon_0}$ distinct j-simplices spanned by the points of T_0, for an approriate positive constant c_j, and all these simplices are stored at N_{v_j}.

This establishes the induction step for j, since $\beta_{j+1} = 2j\beta_j + 1$, by definition, and thus completes the inductive proof of the lemma. □

There are several interesting consequences of our construction. First, let Q denote the collection of all points of intersection between the lines f_v, for nodes v of the last-level interval trees, and the $(d-1)$-simplices spanned by the points of S and stored at v. The analysis given above implies that the size of Q is $O(n^d)$. Note that the set Q depends only on S and is independent of ε. We have thus shown the existence of a fixed set Q of $O(n^d)$ points, depending only on S, so that, for any $\varepsilon > 0$, a weak ε-net for S can be obtained by an appropriate sampling of the points of Q.

Second, if we construct the weak ε-net W by sampling more points of Q, say three times more densely, then W has the following additional property. If T is a subset of S containing εn points, then $conv(T)$ contains a point z of W so that there are at least $c\varepsilon^d n^d/\log^{\beta_d}\frac{1}{\varepsilon}$ $(d-1)$-simplices spanned by points of T and lying above z, and at least that many such simplices lying below z. This in turn implies that z is an *approximate center point* of T, meaning that any halfspace bounded by a hyperplane passing through z must contain at least $\alpha\varepsilon n$ points of T, where $\alpha = \Omega(1/\log^{\beta_d/d}\frac{1}{\varepsilon})$. It is easily checked that this also holds for any subset $T \subseteq S$ that contains at least εn points. This property is weaker than being a real center point of T, which is a point having the property that each halfspace bounded by a hyperplane passing through p contains at least $1/(d+1)$ of the points of T (it is well known that such a point always exists; see [8]). Still it is interesting that the fixed, and reasonably small, set Q contains an approximate center point for every subset T of S that contains at least εn points.

3 Algorithms for Computing Weak ε-Nets

In this section we develop efficient algorithms for computing a weak ε-net of a given point set S in $\mathbb{R}^d$. We first present a deterministic algorithm that runs in time $O(n^d \log(1/\varepsilon))$, and then present a randomized algorithm.

4 Weak ε-nets for Planar point sets in Convex Position

For $0 < \varepsilon < 1$, let $\delta(\varepsilon)$ be the smallest integer that guarantees the existence of a weak ε-net N of size $\delta(\varepsilon)$ of every finite planar point set S in convex position (for convex sets); (that is, if $A \subseteq S$, $|A| > \varepsilon|S|$, then the convex hull of A contains a point in N). We show that

$$\delta(\varepsilon) = O(\frac{1}{\varepsilon}\log^{\log_2 3}\frac{1}{\varepsilon}) .$$

To do this, we prove that for any real number ε, $0 < \varepsilon < 1$, and any positive integer ℓ, the function δ obeys the inequality

$$\delta(\varepsilon) \le \binom{\ell}{2} + \ell\delta(\frac{\ell\varepsilon}{3}) . \qquad (1)$$

The bound on $\delta(\varepsilon)$ follows by choosing $\ell = 3/\sqrt{\varepsilon}$, which gives

$$\delta(\varepsilon) \le O(\frac{1}{\varepsilon}) + \frac{3}{\sqrt{\varepsilon}}\delta(\sqrt{\varepsilon}),$$

and consequently,

$$\delta(\varepsilon) \le (1 + 3 + 3^2 + \cdots + 3^{\log_2\log_2(1/\varepsilon)})O(1/\varepsilon).$$

Let S be a set of n points and let $\ell \le n$. We select ℓ points in S enumerated by $p_0, p_1, \ldots, p_{\ell-1}, p_\ell = p_0$ in counterclockwise direction; the choice has to be made so that between any two consecutive points p_{i-1} and p_i there are at most n/ℓ points of S. Let S_i denote the points from S between p_{i-1} and p_i (without the points p_{i-1} and p_i themselves!). A weak ε-net of S can now be constructed by choosing

1. the points $P = \{p_0, p_1, \ldots, p_{\ell-1}\}$;
2. a weak $(\ell\varepsilon/3)$-net for each of the sets S_i, $1 \le i \le \ell$;
3. the intersection of segment $p_j p_i$ with segment $p_0 p_{j-1}$, for each pair i, j with $1 \le i < j - 1 \le \ell - 2$.

Clearly, this will yield the recursion (1). It remains to verify, that this collection of points forms a weak ε-net. Consider $A \subseteq S$, $|A| > \varepsilon n$. If $A \cap P \neq \emptyset$, then the point set in 1. will hit the convex hull of A. If A is contained in at most three of the S_i's, then for one i, $|A \cap S_i| > \varepsilon n/3$, and the case is handled by the points selected in 2. So it remains to consider the case when A has points in at least four of the S_i's, say for $i = a, b, c, d$, $1 \le a < b < c < d \le \ell$. Then all four quadrants formed by the lines through $p_c p_a$ and $p_0 p_{c-1}$, respectively, contain points from A, and thus the intersection of the respective segments lies in the convex hull of A. This intersection has been chosen in 3; (note that, indeed, $1 \le a < c - 1 \le \ell - 2$). We thus have established our claim.

Theorem 4.1 *Given n points in convex position, there exists a weak ε-net for convex sets of size $O(\frac{1}{\varepsilon} \log^{\log_2 3} \frac{1}{\varepsilon})$.*

5 Weak ε-Nets for Points Uniformly Distributed on a Circle

We show that if P consists of points with a *quasi-uniform* distribution on the unit circle $\mathcal{U}$ then there exists a weak ε-net for P of size $O(1/\varepsilon)$. By a quasi-uniform distribution we mean that any arc of $\mathcal{U}$ of length λ should contain at most $\lceil c\lambda n \rceil$ points of P, for some constant $c > 0$. This result is a simple corollary of the following theorem (Fig.1).

Theorem 5.1 *Given any $\varepsilon > 0$, there exists a set P of size $O(1/\varepsilon)$ such that any triangle whose vertices lie in $\mathcal{U}$ and whose side lengths exceed ε must intersect P.*

Here is a brief sketch of the proof. Consider a regular $\lceil c/\varepsilon \rceil$-gon R inscribed in the unit disk $\mathcal{D}$ (for some large constant c) and tile the plane with a square grid fine enough to have have about $1/\varepsilon$ vertices inside R: make these grid points the set P. Obviously this does not work, because the distribution of points should be denser near the boundary of the disk. So, instead, take a non-uniform grid where the density of points at a distance r from the center is roughly $r(1 - r^2)^{-3/2}$. Why such an odd-looking density? It is the intrinsic measure of the so-called *Lorentz space* of special relativity, which in turn is derived from the metric of the hyperbolic plane in its projective (Klein) model. Although the proof of Theorem 5.1 can be given solely in Euclidean terms, it is technical and ugly. Strikingly, the proof is completely trivial in hyperbolic geometry, which thus appears to be its natural "habitat." To provide a better understanding of what is happening and make the presentation self-contained, we begin with a brief overview of the geometry of the hyperbolic plane. See [4, 7, 9, 10, 12, 14, 15] for background material.

5.1 Hyperbolic Geometry

Euclid's fifth postulate states that through a given point exactly one line can be parallel to another given line. Hyperbolic geometry postulates that an infinite number of such lines exist. For our purposes the most interesting consequence of that fact is that the area of any triangle is bounded above by a constant. In particular, the area of a convex polygon with its vertices on the unit circle depends only on its number of sides and not on the position of the vertices. This allows us to treat the "test" convex sets for weak ε-nets in a uniform manner, and thus substitute simple area-based geometric arguments for otherwise complicated combinatorial proofs. The hyperbolic plane is typically modeled in one of three ways: the Klein model, the Poincaré model, or the halfplane model. We briefly discuss each of them.

The hyperbolic plane H^2 can be defined as the hyperboloid $\mathcal{H} : 1 - z^2 + x^2 + y^2 = 0$ with its associated Lorentz metric:

$$ds^2 = -dz^2 + dx^2 + dy^2.$$

To make H^2 into a part of the real projective plane, we identify antipodal points, or more simply, we restrict ourselves to the upper component $\mathcal{H}^+ = \mathcal{H} \cap \{z > 0\}$. In this way, a point of H^2 is represented by a line in E^3 passing through the origin and piercing the hyperboloid (Fig.2). To be able to view the hyperbolic plane as a plane and not as a set of lines, three planar models have been devised.

The Klein Model. It is obtained by centrally projecting $\mathcal{H}^+$ onto the plane $z = 1$. Namely, the point $(x, y, z) \in \mathcal{H}^+$ maps to $(x/z, y/z)$ on the plane $z = 1$ (Fig.2). The unit circle $\mathcal{U}$ bounding $\mathcal{D}$ corresponds to points at infinity. Viewed in E^3, a line of H^2 is the plane spanned by two lines passing through O. The intersection of that plane with $z = 1$ is a line, therefore lines still look straight in the Klein model. The underlying metric, however, must be transferred from the hyperboloid to the unit disk $\mathcal{D}$, and so it is not the familiar Euclidean metric. In particular, angles and distances cannot be read off by simply looking at them in the Klein disk.

It is actually quite easy to express the Lorentz metric in the Klein disk supplied with polar coordinates. Let (r, θ) be a point in the Klein disk with Euclidean polar coordinates (ρ, θ). The point is the image of a point $(t \cos\theta, t \sin\theta, z) \in \mathcal{H}^+$, therefore $t^2 - z^2 = -1$. Since $\rho = t/z$, we find that $t = \rho(1 - \rho^2)^{-1/2}$ and hence

$$dt^2 = (1 - \rho^2)^{-3}\, d\rho^2.$$

From $z = t/\rho$ we derive that $z = (1 - \rho^2)^{-1/2}$, and hence,

$$dz^2 = \rho^2 (1 - \rho^2)^{-3}\, d\rho^2.$$

Using the Lorentz metric, we find that

$$dr = \sqrt{dt^2 - dz^2} = \frac{d\rho}{1-\rho^2}.$$

A deviation of $d\theta$ produces an arc of length $td\theta$ (and not $rd\theta$!), therefore the area measure is $tdrd\theta$, which is

$$\rho(1-\rho^2)^{-3/2} d\rho d\theta,$$

as claimed earlier. We can now derive that any triangle with angles α, β, γ has area $\pi - (\alpha + \beta + \gamma)$. This fact can be derived much more simply in the halfplane model (see Appendix).

The metric of the hyperbolic plane is invariant under projective transformations. Since cross-ratios are known to be invariant under such transformations, it is natural that the hyperbolic distance between p and q should be a function of the cross-ratio of (a, b, p, q) (Fig.3). Indeed, it is equal to

$$\frac{1}{2} \log \frac{|q-a||b-p|}{|p-a||b-q|}.$$

The logarithm is needed to make the distance function additive. We should also point out that the factor $1/2$ is inserted to make the curvature -1. The notion of curvature is central to hyperbolic geometry: Indeed, in H^2 the circumference of a circle of radius r is greater than $2\pi r$, which is similar to what happens (locally) in the Riemannian metric of a negatively curved surface, e.g., a hyperbolic paraboloid. Note that, on the contrary, the circumference of a small circle taken on a positively curved surface, such as a sphere, is less than $2\pi r$. It is possible to construct negatively surfaces in $\mathbf{R}^3$ whose intrinsic metric is hyperbolic, but this can be only a local property because by a theorem of Hilbert no surface in $\mathbf{R}^3$ can be isometric to all of H^2. Intuitively, there is not enough room in Euclidean space to accommodate the hyperbolic metric.

The Klein model is useful to deal with point and lines but it does not help much in dealing with angles. For this we must turn to *conformal* models, which preserve angles.

The Poincaré Model. Let O be the center of the Klein disk $\mathcal{D}$ and let $\mathcal{S}$ be the unit sphere centered around O. A point p in the Klein disk is mapped to the point p' in the Poincaré disk (which is also $\mathcal{D}$) as follows: project p vertically down to the southern hemisphere of $\mathcal{S}$ and then centrally onto $\mathcal{D}$ towards the north pole of $\mathcal{S}$ (Fig.4). Straight lines in the Poincaré disk model appear as circular arcs orthogonal to $\mathcal{U}$. A line segment joining two points p and q is shown in Fig.5. In this model, isometries are obtained by composing inversions in circles. Given a circle C, an inversion about it leaves C invariant (pointwise) as well as any circle orthogonal to it (globally), e.g., L and L' (Fig.6). Let q (resp. q') be the image of p (resp. p') under the inversion. Since we claim that inversions are isometries, the (hyperbolic) length of pp' must be equal to the length of qq'. More generally, the lengths of all the transversals shown in Fig.6 must be equal, and therefore can depend only on the angle between L and L' and not on their particular position. The metric ds^2 is easily expressed as a function of the Euclidean metric $dx^2 + dy^2$:

$$ds^2 = \frac{4}{(1 - r^2)^2} (dx^2 + dy^2),$$

where r is the distance from the point (x, y) to O.

Lemma 5.2 *Let p and q be two points in the Poincaré disk whose Euclidean distance δ is equal to $\delta'/100$, where δ' is the Euclidean distance between $\{p, q\}$ and $\mathcal{U}$. Then the hyperbolic distance between p and q exceeds a positive constant (independent of p and q).*

Proof: Because δ is much smaller than δ', the hyperbolic distance between p and q can be estimated accurately by integrating $4(1 - r^2)^{-2} (dx^2 + dy^2)$ while keeping the value of r fixed. For the same reason, if we carry the integration along the Euclidean segment pq (instead of the geodesic), we lose at most another constant factor in the estimation. Clearly, we can assume that δ' is very small, which implies that r is close to 1. Then, a Taylor expansion around $r = 1$ shows that $1/(1 - r^2)^2$ is on the order of $1/\delta'$, and therefore the hyperbolic distance between p and q is on the order of $\delta/\delta' = \frac{1}{100}$. $\square$

A triangle in the Poincaré model is the region of $\mathcal{D}$ bounded by three circular arcs orthogonal to $\mathcal{U}$. The triangle is called *ideal* if its three vertices lie on $\mathcal{U}$. If its angles are α, β, γ, we already mentioned that its area is $\pi - (\alpha + \beta + \gamma)$. Note that ideal triangles have zero angles, so their area is exactly π (even though their sides have infinite length). Unlike its Euclidean counterpart, a triangle is completely characterized (up to congruency) by its three angles. The Poincaré model is conformal, so we can reason directly about angles. For example, it is easy to show that any regular n-gon can be used to tile the whole hyperbolic plane (which shows how much more room H^2 has compared to E^2). Indeed, consider a regular n-gon centered at O. By triangulating it we immediately derive that its area is equal to $(n - 2)\pi - n\alpha$, where α is its vertex angle. If the polygon is ideal, i.e., if all its vertices lie on the unit circle $\mathcal{U}$, then $\alpha = 0$. If we continuously "shrink" the polygon towards O, however, its area goes to 0, and therefore, α tends to $(1 - 2/n)\pi$. (Note that near the origin the hyperbolic plane behaves like the Euclidean plane.) This means that, at some point during the shrinking, α becomes equal to $2\pi/n$ (Fig.7). We can now draw the polygon at that position and reflect it about its edges (since angles sum up to 2π around the vertices). Iterating these reflections (which from a Euclidean standpoint are circle inversions) tiles all H^2.

A particularly interesting class of tilings is obtained by reflecting triangles around their edges. It is a standard theorem [12] that given any positive integers l, m, n such that

$$\frac{1}{l} + \frac{1}{m} + \frac{1}{n} < 1,$$

the triangle with angles π/l, π/m, and π/n (which is unique up to congruency) can tile H^2. Figure 8 shows a tiling with $l = 2$, $m = 3$, $n = 7$. The tiling is infinite so it is only shown within a finite disk. If X denotes reflection around edge x, the tiling is generated by the group, denoted $T^*(l, m, n)$, with generators L, M, N and relations,

$$(MN)^l = (NL)^m = (LM)^n = 1 \quad \text{and} \quad L^2 = M^2 = N^2 = 1.$$

The first group of relations express the fact that reflected images incident upon a fixed vertex cycle back after a while. The second group says that reflections are involutory.

The characterization of hyperbolic triangle groups given above immediately implies that (unfortunately) triangular tilings must be made of triangles of diameter higher than some fixed constant. In other words, triangles involved in a tiling cannot be too small. In fact, $T^*(2, 3, 7)$ is the tiling whose fundamental region has the smallest possible triangle: from what we said earlier, its area is $(1-\frac{1}{2}-\frac{1}{3}-\frac{1}{7})\pi \approx .0238$. For concreteness, place the center O of $\mathcal{U}$ at a vertex of degree 14 in the tiling generated by $T^*(2, 3, 7)$ (Fig.8).

Fix $0 < r < 1$, and let $\mathcal{D}_r$ be the disk centered at O of Euclidean radius $r < 1$. It is immediate to verify (for example, by using Lemma 5.2) that the number of triangles intersecting $\mathcal{D}_r$ is $O(1/(1 - r))$. Suppose that we wish to have smaller triangles. By decomposing each triangle barycentrically (Fig.9), and iterating in this fashion a constant number of times, we can bring down the hyperbolic diameter of every triangle below any desired positive constant. Note that the total number of triangles (intersecting $\mathcal{D}_r$) remains $O(1/(1 - r))$. Although it is not a problem here, it is worth noting that this operation is no longer a tiling since the new triangles are not congruent (and interestingly can *never* be made congruent). We summarize our results:

Lemma 5.3 *There exists a triangulation of the Poincaré disk using triangles of hyperbolic diameter below any desired positive constant such that, given any $0 < r < 1$, the number of triangles overlapping the disk $\mathcal{D}_r$ of Euclidean radius r is $O(1/(1-r))$.*

The Proof of the Theorem. Set $r = 1 - \varepsilon/10$ in Lemma 5.3, and choose a triangulation $\mathcal{T}$ of the Poincaré disk with triangles of hyperbolic diameter less than some suitable constant $c > 0$. We claim that, for c small enough, the set P consisting of the vertices of $\mathcal{T}$ within $\mathcal{D}_r$ satisfies the conditions of Theorem 5.1. To begin with, observe that the set contains $O(1/(1 - r))$ points, which is the desired size.

Next, let uvw be an ideal triangle whose Euclidean side lengths exceed ε. Since a constant number of random points hit every triangle with big enough sides, we can certainly assume that the sides of uvw are fairly short. Also, we can assume that uv and vw are congruent. (Because if w is further from v than u is, then sliding w towards v shrinks the triangle.) Now, let uvh be the triangle obtained as the intersection of the triangles uvw and $v'uv$, where v' is the reflection of v around u (Fig.10). Let $\lambda > \varepsilon$ be the Euclidean distance from u to v. Elementary calculations show that the triangle uvh contains a disk D^* whose Euclidean radius and Euclidean distance to $\mathcal{U}$ are both greater than, say, $\lambda/10$. By our choice of r, the disk is entirely contained in $\mathcal{D}_r$. Suppose for the sake of contradiction that D^* does not contain any point of P. Then, the triangle of $\mathcal{T}$ that contains the (Euclidean) center p of D^* must also contain a point at Euclidean distance at least $\lambda/10$ from p. Since the Euclidean distance from p to $\mathcal{U}$ is less than 2λ, the triangle must contain a point q such that the pair p, q satisfies the conditions of Lemma 5.2. This implies that their hyperbolic distance exceeds a fixed positive constant. Thus, choosing c small enough causes a contradiction, and the proof of Theorem 5.1 is now complete. □

6 Appendix

The *halfplane model* is obtained by applying to the Poincaré disk an inversion whose pole lies on the unit circle (Fig.11). Because inversions are conformal mappings, angles are still accurately represented in the halfplane model. The model is particulary convenient to work with because it can be thought of as the upper part of the complex plane. Pursuing this analogy, it can be shown that the group of isometries is precisely the group of fractional linear transformations (also know as Möbius transformations)

$$z \in \mathbf{C} \mapsto \frac{az+b}{cz+d},$$

with real coefficients and determinant $ad - bc \neq 0$.

It is easy to verify that the invariant area measure is $dxdy/y^2$, which is much simpler than in the Klein model. To illustrate this point, let us evaluate the area of a triangle in the halfplane model. First, observe that all ideal triangles are congruent. To verify that their common area is π, consider the triangle with vertices $(-1, 0)$, $(1, 0)$ and $(0, \infty)$ (Fig.12). Its area is

$$\int_{-1}^{1} \int_{\sqrt{1-x^2}}^{\infty} \frac{1}{y^2}\, dy\, dx = \int_{-1}^{1} \frac{1}{\sqrt{1-x^2}}\, dx = \int_{-\pi/2}^{\pi/2} \frac{\cos\theta}{\sqrt{1-\sin^2\theta}}\, d\theta = \pi.$$

We can treat the general case by a very simple argument due to Gauss [15]: Given $0 \leq \theta \leq \pi$, let $f(\theta)$ be the area of a triangle with two ideal vertices and one vertex with angle θ (we can verify that, for fixed θ, all these triangles are congruent). From Fig.13, we immediately derive $f(\alpha)+f(\pi-\alpha) = \pi$ and $f(\alpha)+f(\beta)+f(2\pi-\alpha-\beta) = \pi$, therefore

$$f((\pi-\alpha)+(\pi-\beta)) = f(\pi-\alpha)+(\pi-\beta)-\pi.$$

It follows that f is linear and, more precisely, that $f(\alpha) = \pi - \alpha$. To deal with the general case, consider an arbitrary triangle with angles α, β, γ. Extend each side in one direction towards $\mathcal{U}$ and form the ideal triangle with the three new vertices (Fig.14). In view of the previous result, the area A of the triangle satisfies $A + f(\pi - \alpha) + f(\pi-\beta) + f(\pi-\gamma) = \pi$, hence

$$A = \pi - (\alpha+\beta+\gamma),$$

as claimed earlier.

Acknowledgments

The authors wish to express their gratitude for the generous support and hospitality of the DEC Palo Alto Systems Research Center.

References

[1] Alon, N., Bárány, I., Füredi, Z., Kleitman, D. *Point selections and weak ε-nets for convex hulls*, Combin. Prob. and Computing 1, submitted.

[2] Alon, N., Kleitman, D. *Piercing convex sets and the Hadwiger Debrunner (p, q)-problem*, Adv. Math. (to appear).

[3] Aronov, B., Chazelle, B., Edelsbrunner, H., Guibas, L.J., Sharir, M. Wenger, R. *Points and triangles in the plane and halving planes in space*, Discrete Comput. Geom. 6 (1991), 435–442.

[4] Berger, M. *Geometry II*, Springer-Verlag, 1987.

[5] Chazelle, B. *Cutting hyperplanes for divide-and-conquer*, Disc. Comput. Geom. (1991), in press. Also, CS-TR-335-91, Princeton University, 1991. Prelim. version in Proc. 32nd FOCS, 1991.

[6] Chazelle, B., Edelsbrunner, H., Guibas, L.J., Hershberger, J., Seidel, R., Sharir, M. *Slimming down by adding; selecting heavily covered points*, Proc. 6th ACM Symp. on Computational Geometry (1990), 116–127.

[7] Coxeter, H.M.S. *Non-Euclidean Geometry*, Univ. Toronto Press, Toronto, 1942.

[8] Edelsbrunner, H. *Algorithms in Combinatorial Geometry*, Springer-Verlag, Heidelberg 1987.

[9] Epstein, D.B.A. *Analytical and Geometric Aspects of Hyperbolic Space*, London Math. Society Lecture Notes Series 111, 1984.

[10] Fenchel, W. *Elementary Geometry in Hyperbolic Space*, de Gruyter Studies in Mathematics 11, 1989.

[11] Haussler, D., Welzl, E. *Epsilon nets and simplex range queries*, Discrete Comput. Geom. 2 (1987), 127–151.

[12] Magnus, W. *Noneuclidean tesselations and their groups*, Academic Press, New York and London, 1974.

[13] Matoušek, J. *Approximations and optimal geometric divide-and-conquer*, Proc. 23rd Annu. ACM Symp. Theory Comput. (1991), 506–511.

[14] Milnor, J. *Hyperbolic geometry: the first 150 years*, Bull. Amer. Math. Soc. 6 (1982), 9–24.

[15] Thurston, W.P. *The Geometry and Topology of 3-manifolds*, Princeton Univ. Math. Dept. 1979.