Improved Bounds on Weak ε -Nets for Convex Sets*

B. Chazelle,¹ H. Edelsbrunner,² M. Grigni,¹ L. Guibas,³ M. Sharir,⁴ and
E. Welzl⁵

¹ Department of Computer Science, Princeton University,
Princeton, NJ 08544, USA
chazelle@princeton.edu

² Computer Science Department, University of Illinois,
Urbana, IL 61801, USA

³ DEC Systems Research Center, Laboratory for Computer Science,
Massachusetts Institute of Technology, Cambridge, MA 02139, USA
and
Computer Science Department, Stanford University,
Stanford, CA 94305, USA

⁴ School of Mathematical Sciences, Tel Aviv University,
Ramat Aviv 69 978, Israel
and
Courant Institute of Mathematical Sciences, New York University,
251 Mercer Street, New York, NY 10012, USA

⁵ Institut für Informatik, Freie Universität Berlin,
Arnimallee 2–6, D-14195 Berlin, Germany

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 $O((1/\varepsilon^d) \log^{\beta_d}(1/\varepsilon))$, 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.* Such a net can

* 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 NSERC Operating Grants and NSF Grant DMS-9206251. 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.

be computed effectively. We also consider two special cases: when S is a planar point set in convex position, we prove the existence of a net of size $O((1/\varepsilon) \log^{1.6}(1/\varepsilon))$. In 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 $O(1/\varepsilon)$, which improves a previous bound of Capowleas.

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 . This contrasts with the notion of ε -net for range spaces of finite VC-dimension, for which a complete theory has been developed [11], [13]. In our case the underlying range space has infinite VC-dimension, so none of that theory applies.

Alon *et al.* [1] (see also [2]) have recently shown that weak ε -nets for convex ranges always exist, which have size $O(1/\varepsilon^{(d+1)(1-1/s_d)})$, 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 $O(1/\varepsilon^2)$. In this paper we improve their bound for $d \geq 3$, showing the existence of weak ε -nets of size $O((1/\varepsilon^d) \log^{\beta_d}(1/\varepsilon))$, 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, such 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, a net point p which is an *approximate center point* of T exists. This means that any half-space that contains p also contains at least some fixed fraction of the points of T . While, for an exact center point, this fraction is at least $1/(d+1)$ [8], in our case it is only $\Omega(1/\log^{\beta_d}(1/\varepsilon))$.

Finally, we look at two special cases: If S consists of points in convex position in the plane, then a net of size $O((1/\varepsilon) \log^c(1/\varepsilon))$ can be found, where $c = \log_2 3 \approx 1.6$.

If the points of S lie uniformly on a circle (as do, say, the vertices of a regular polygon), then we exhibit an optimal weak ε -net of size $O(1/\varepsilon)$, which improves a previous bound of $O(1/\varepsilon) \cdot 2^{\log^*(1/\varepsilon)}$ due to Capowleas [5]. Interestingly, our weak ε -net consists of vertices from a tessellation of the hyperbolic plane. Thus, we show that the problem “lives” naturally in hyperbolic space.

The motivation for studying weak ε -nets is twofold: On a general level, the study addresses the central, largely unresolved, issue of sampling in spaces of infinite VC-dimension. This arises when dealing with problems whose dimensionality is a parameter that is not fixed but part of the input: examples are standard optimization problems, such as linear programming, or the problem of simulating quasi-uniform distribution for, say, learning boolean formulas. On a more specific level, weak ε -nets provide threshold tests for convex functions, which might be useful in Monte Carlo applications. For example, suppose that we are given a density measure in d -space and we wish to test quickly whether a convex polytope has volume above a certain threshold (in the given measure). We can do that in the following manner: First, sample very finely according to the given measure

and then compute a weak ε -net of that sample. It is easy to construct smooth but irregular measures where the standard Monte Carlo approach fails miserably: in those cases, building a weak ε -net appears to be the only solution at this point. To extend these threshold tests into full-fledged approximation schemes is open. We can show that weak ε -approximations (in the standard sense) do not exist, but this does not rule out other lines of attack.

2. Construction of Weak ε -Nets in $\mathbb{R}^d$

Let S be a given set of n points in $\mathbb{R}^d$. We 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 construct a multilevel 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 stopping each interval I at the unique node of the tree whose left subtree stores one endpoint of I and whose right subtree stores the other endpoint. In what follows we 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 $j \leq d - 1$. A node v of a $(j - 1)$ -level interval tree corresponds to a $(d - j + 1)$ -flat f_v of the form $x_1 = a_1, \dots, x_{j-1} = a_{j-1}$, and each segment stored at v corresponds to a $(j - 1)$ -simplex $p_1 p_2 \dots p_j$ that is spanned by S and intersects f_v at a point. Let N_v denote the set of all these simplices and let 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 $(p_1 \dots p_{j-1} q)$, $(p_1 \dots p_{j-1} r)$, 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 $[(p_1 p_2 \dots p_{j-1} p_j), (p_1 p_2 \dots p_{j-1} p_{j+1})]$, corresponds to the j -simplex $p_1 p_2 \dots p_{j+1}$, 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 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 $\ell \geq 0$ in its tree.

We define a sequence $\{\beta_j\}_{j \geq 2}$ of integers by $\beta_2 = 0$ and $\beta_{j+1} = 2j\beta_j + 1$, for $j \geq 2$. If we put $\beta_j = 2^{j-1}(j - 1)!$, ξ_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 $\varepsilon_v = \frac{1}{2}(\frac{3}{4})^\ell \varepsilon$, $M_v = (c/\varepsilon_v^d) \log^{\beta_d}(1/\varepsilon_v)$, for an appropriate constant c depending on d , 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 , all belonging to the subtree of u . 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 (this argument is omitted here). Moreover, none of these lists store simplices having a vertex outside S_u . 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

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

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 $2(\frac{2}{3})^d < 1$, 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 $O((1/\varepsilon^d) \log^{\beta_d}(1/\varepsilon))$ 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 $\text{conv}(T) \cap W \neq \emptyset$. We 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 $\frac{1}{4}(\frac{3}{4})^\ell \varepsilon n$ 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 $(\frac{3}{4})^\ell \varepsilon n$ points of T in its subtree. However, the number of points of S stored at that subtree is at most $n/2^\ell$, which is smaller than $(\frac{3}{4})^\ell \varepsilon n$ when ℓ is sufficiently large. Put $\varepsilon_0 = \frac{1}{2}(\frac{3}{4})^\ell \varepsilon$. 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 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 some node v_j exists whose associated list N_{v_j} contains at least $c \varepsilon_0^{j+1} n^{j+1} / \log^{\beta_{j+1}}(1/\varepsilon_0)$ 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 node v_{d-1} at the last level $d-1$, whose list $N_{v_{d-1}}$ contains at least $c_{d-1}\varepsilon_0^d n^d / \log^{\beta_d}(1/\varepsilon_0)$ $(d-1)$ -simplices spanned by the points of T_0 . All these simplices are contained in $\text{conv}(T)$, and $\text{conv}(T)$ 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 $\text{conv}(T)$, must contain a point of W .

To show the existence of the nodes v_j , we argue as follows. We 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 , some point on the line that is contained in at least $m^2/4n^2$ intervals of M exists.*

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 the case $j=2$ separately. 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 $> t/2\varepsilon_0 n$. 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 $t/2\varepsilon_0 n$; this process cannot remove more than $t/2$ 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 $|E'| \cdot (t/2\varepsilon_0 n)^2$. 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 that descends from a node u to its left (right) child if a_2 is small (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 $\geq 3 \log(1/\varepsilon_0) + \log(1/c') + 2$ 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

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

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 $c_2 \varepsilon_0^3 n^3 / \log(1/\varepsilon_0)$ 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, $t = |E| \geq c_{j-1} \varepsilon_0^j n^j / \log^{\beta_j}(1/\varepsilon_0)$. 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 $t/2 \binom{\lfloor \varepsilon_0 n \rfloor}{j-1}$ 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 $t/2 \binom{\lfloor \varepsilon_0 n \rfloor}{j-1}$; this process cannot remove more than $t/2$ edges of E , since there are at most $\binom{\lfloor \varepsilon_0 n \rfloor}{j-1}$ 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 $q_1 \neq p_1, \dots, p_j$ of T_0 such that $(p_2 \cdots p_j q_1)$ is also in E' , then choose any point $q_2 \neq p_1, \dots, p_j, q_1$ 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 $q_1, \dots, q_j$ are chosen (so the last edge is $(q_1 q_2 \cdots q_j) \in E'$), and add to $\mathcal{M}$ the segment connecting $(p_1 p_2 \cdots p_j)$ and $(q_1 q_2 \cdots q_j)$ in $f_{v_{j-1}}$.

The pruning procedure ensures that the number of segments in $\mathcal{M}$ is at least $|E'| \cdot \left(t / 2 \binom{\lfloor \varepsilon_0 n \rfloor}{j-1} - j \right)^j$. 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

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

projected segments of $\mathcal{M}$, for an appropriate positive constant c' , depending on j and on c_{j-1} . Clearly, if $[(p_1 p_2 \cdots p_j), (q_1 q_2 \cdots q_j)]$ 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 $[(p_1 p_2 \cdots p_j), (p_2 \cdots p_j q_1)]$, $[(p_2 \cdots p_j q_1), (p_3 \cdots p_j q_1 q_2)]$, $\dots$, $[(p_j q_1 q_2 \cdots q_{j-1}), (q_1 q_2 \cdots q_j)]$ must have an x_j -projection that also contains a_j . Hence, a_j is contained in at least $c' \varepsilon_0^{2j} n^{2j} / \log^{2j\beta_j}(1/\varepsilon_0)$ projected segments of the latter kind, which, by construction, are all sorted 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 $c'' \varepsilon_0^{j+1} n^{j+1} / \log^{2j\beta_j}(1/\varepsilon_0)$ 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 $(j+1) \log(1/\varepsilon_0)$, to conclude that a node v_j exists in that interval tree whose associated $(d-j)$ -flat f_{v_j} crosses at least $c_j \varepsilon_0^{j+1} n^{j+1} / \log^{2j\beta_j+1}(1/\varepsilon_0)$ distinct j -simplices spanned by the points of T_0 , for an appropriate 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. $\square$

Theorem 2.4. *Any finite point set in $\mathbb{R}^d$ admits a weak ε -net for convex sets, of size $O((1/\varepsilon^d) \log^{\beta_d}(1/\varepsilon))$.*

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 $\text{conv}(T)$ contains a point z of W so that there are at least $c \varepsilon^d n^d / \log^{\beta_d}(1/\varepsilon)$ $(d-1)$ -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 half-space bounded by a hyperplane passing through z must contain at least αen points of T , where $\alpha = \Omega(1/\log^{d/d}(1/\varepsilon))$. It is easily checked that this also holds for any subset $T \subseteq S$ that contains at least en points. This property is weaker than being a real center point of T , which is a point having the property that each half-space 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 en points.

3. Weak ε -Nets for Planar Point Sets on Convex Position

For $0 < \varepsilon < 1$, let $\delta(\varepsilon)$ be the smallest integer for which any finite planar point set S in convex position admits a weak ε -net N for convex sets, of size $\delta(\varepsilon)$. We show that

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

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

$$\delta(\varepsilon) \leq \binom{\ell}{2} + \ell \delta\left(\frac{\ell \varepsilon}{3}\right), \quad (1)$$

with the trivial “boundary condition” $\delta(1) = 1$. The bound on $\delta(\varepsilon)$ follows by choosing $\ell = 3/\sqrt{\varepsilon}$, which gives

$$\delta(\varepsilon) \leq O\left(\frac{1}{\varepsilon}\right) + \frac{3}{\sqrt{\varepsilon}} \delta(\sqrt{\varepsilon}),$$

and, consequently,

$$\begin{aligned} \delta(\varepsilon) &\leq (1 + 3 + 3^2 + \cdots + 3^{\log_2 \log_2(1/\varepsilon)}) \cdot O\left(\frac{1}{\varepsilon}\right) \\ &= O\left(\frac{1}{\varepsilon} \log^{\log_2 3} \frac{1}{\varepsilon}\right). \end{aligned}$$

Let S be a planar set of n points in convex position, and let $\ell \leq n$. We select ℓ points in S enumerated by $p_0, p_1, \dots, p_{\ell-1}, p_\ell = p_0$ 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:

- (a) The points $P = \{p_0, p_1, \dots, p_{\ell-1}\}$.
- (b) A weak $(\ell\varepsilon/3)$ -net obtained recursively for each of the sets S_i , $1 \leq i \leq \ell$.
- (c) The intersection of segment $p_j p_i$ with segment $p_0 p_{j-1}$, for each pair i, j with $1 \leq i < j - 1 \leq \ell - 2$.

Clearly, this will yield the recursion (1). It remains to verify that this collection of points forms a weak ε -net. Consider $T \subseteq S$, $|T| \geq \varepsilon n$. If $T \cap P \neq \emptyset$, then the point set in (a) will hit the convex hull of T . If T is contained in at most three of the S_i 's, then, for one i , $|T \cap S_i| \geq \varepsilon n/3$, and $\text{conv}(T)$ is thus hit by a point selected in (b). So, it remains to consider the case when T has points in at least four of the S_i 's, say for $i = a, b, c, d$, $1 \leq a < b < c < d \leq \ell$. Then all four quadrants formed by the lines through $p_c p_a$ and $p_0 p_{c-1}$, respectively, contain points from T , and thus the intersection of the respective segments lies in the convex hull of T . This intersection has been chosen in (c) (note that, indeed, $1 \leq a < c - 1 \leq \ell - 2$). We thus have established our claim.

Theorem 3.1. *Given a planar set S of n points in convex position, a weak ε -net for convex sets of S exists, of size $O((1/\varepsilon) \log^{\log_2 3}(1/\varepsilon))$.*

4. Weak ε -Nets for Points Uniformly Distributed on a Circle

We next show that if S consists of points with a *quasi-uniform* distribution on the unit circle $\mathcal{U}$, then a weak ε -net for S of size $O(1/\varepsilon)$ exists. 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 S , for some constant $c > 0$. This result, which improves a previous bound of $O(1/\varepsilon 2^{\log^*(1/\varepsilon)})$ by Capowleas [5], is a simple corollary of the following theorem.

Theorem 4.1. *Given any $\varepsilon > 0$, a set of P of size $O(1/\varepsilon)$ such that any triangle whose vertices lie in $\mathcal{U}$ and whose side lengths all exceed ε must intersect P .*

A Brief Sketch of the Proof. A first attempt might be to put within the disk $\mathcal{D}$ bounded by $\mathcal{U}$ a sufficiently fine square grid that has $O(1/\varepsilon)$ vertices, and take the set of these vertices as our net. This will not work, unfortunately, because the area of a triangle as in the theorem might be much too small to be hit by a grid vertex. As an extreme case, consider a triangle uvw where uv, vw are two chords in $\mathcal{D}$ of length ε each; the area of uvw is easily seen to be $O(\varepsilon^3)$, which is indeed two orders of magnitude smaller than areas for which the grid is guaranteed to hit the triangle. The trick is to distribute the net points within $\mathcal{D}$ so that the density of the distribution is larger as we approach the boundary $\mathcal{U}$. We thus take a nonuniform 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 area of the hyperbolic plane in its projective (Klein) model. Although our proof of Theorem 4.1 can be interpreted in Euclidean terms [14], it is nonconstructive and more complicated. Strikingly, the proof is completely trivial in hyperbolic geometry, which thus appears to be its natural “habitat.” We assume that the reader is familiar with the basic properties of the hyperbolic plane. See [4], [7], [9], [10], [12], [15], and [16] for background material.

Lemma 4.2. *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. In the Poincaré model the metric ds^2 is of the form $4(1 - r^2)^{-2}(dx^2 + dy^2)$. Because δ is much smaller than δ' , the hyperbolic distance between p and q can be estimated accurately by integrating ds along the geodesic from p to q and pretending that the value of r is 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 $1/(1 - r^2)$ is on the order of $1/\delta'$, and therefore the hyperbolic distance between p and q is on the order of $\delta/\delta' = 1/100$. $\square$

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 , then 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 the hyperbolic plane H^2 has compared with 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 of 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). Assuming that $n > 3$, this means that, at some point during the shrinking, α becomes equal to $2\pi/n$ (Fig. 1). 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 the entire hyperbolic plane.

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

Fig. 1. An ideal octagon shrinking toward O .

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 2 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.

Fig. 2. The $T^*(2, 3, 7)$ tiling.

Fig. 3. Barycentric subdivision.

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 $(1 - \frac{1}{2} - \frac{1}{3} - \frac{1}{7})\pi \approx 0.07479$. 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. 2).

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 4.2) 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. 3), 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 4.3. *A triangulation of the Poincaré disk using triangles of hyperbolic diameter below any desired positive constant exists such that, given any $0 < r < 1$, the number of triangles overlapping the disk $\mathcal{D}_r$ of Euclidean radius r is $O(1/(1-r))$.*

We are now ready to prove Theorem 4.1. Set $r = 1 - \varepsilon/10$ in Lemma 4.3, and choose a triangulation $\mathcal{T}$ of the Poincaré disk with triangles of hyperbolic diameter less than some suitable constant $d > 0$. We claim that, for d small enough, the set P consisting of the vertices of $\mathcal{T}$ within $\mathcal{D}_r$, when mapped back to the Klein disk satisfies the conditions of Theorem 4.1. To begin with, observe that the set contains $O(1/(1-r)) = O(1/\varepsilon)$ points.

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. 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. 4). We can assume that uv and vw are congruent, for, otherwise, if w is further from v than u is, then sliding w toward v shrinks the triangle uvh , within which we will seek a net point. Let $\lambda > \varepsilon$ be the Euclidean distance from u to v . Elementary calculations show that the triangle uvh contains

Fig. 4. The proof of Theorem 4.1.

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 4.2. This implies that their hyperbolic distance exceeds a fixed positive constant. Thus, choosing d small enough leads to a contradiction, and the proof of Theorem 4.1 is now complete. $\square$

Acknowledgments

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

References

1. Alon, N., Bárány, I., Füredi, Z., and Kleitman, D. Point selections and weak ε -nets for convex hulls, *Combin. Probab. Comput.* **3** (1992), 189–200.
2. Alon, N., and Kleitman, D. Piercing convex sets and the Hadwiger Debrunner (p, q) -problem, *Adv. in Math.* **96** (1992), 103–112.
3. Aronov, B., Chazelle, B., Edelsbrunner, H., Guibas, L. J., Sharir, M., and 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, New York, 1987.
5. Capoteas, V. An almost linear upper bound for weak ε -nets of points in convex position, Manuscript, 1992.
6. Chazelle, B., Edelsbrunner, H., Guibas, L. J., Hershberger, J., Seidel, R., and Sharir, M. Slimming down by adding; selecting heavily covered points, *Proc. 6th ACM Symp. on Computational Geometry*, 1990, pp. 116–127.
7. Coxeter, H. M. S. *Non-Euclidean Geometry*, University of 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 Mathematical Society Lecture Notes Series, Vol. 111, Cambridge University Press, Cambridge, 1984.
10. Fenchel, W. *Elementary Geometry in Hyperbolic Space*, de Gruyter Studies in Mathematics, Vol. 11, de Gruyter, Berlin, 1989.
11. Haussler, D., and Welzl, E. Epsilon nets and simplex range queries, *Discrete Comput. Geom.* **2** (1987), 127–151.
12. Magnus, W. *Noneuclidean Tessellations and Their Groups*, Academic Press, New York, 1974.
13. Matoušek, J. Approximations and optimal geometric divide-and-conquer, *Proc. 23rd Annual ACM Symp. on Theory of Computing*, 1991, pp. 506–511.
14. Matoušek, J. Private communication, 1992.
15. Milnor, J. Hyperbolic geometry: the first 150 years, *Bull. Amer. Math. Soc.* **6** (1982), 9–24.
16. Thurston, W. P. Three-dimensional geometry and topology, Preprint, 1993.

Received August 24, 1993, and in revised form December 20, 1993.