PARTITIONING AND GEOMETRIC EMBEDDING OF RANGE SPACES OF FINITE VAPNIK-CHERVONENKIS DIMENSION

Noga Alon

Department of Mathematics,
Tel Aviv University, Ramat Aviv, TEL AVIV 69978, Israel.

David Haussler

Computer Science Department,
University of California at Santa Cruz, SANTA CRUZ, CA 95064, USA.

Emo Welzl

Institutes for Information Processing,
Technical University of Graz, Schiesstattgasse 4a, A-8010 GRAZ, Austria.

1. Introduction and statement of results.

Range search problems are widely studied in the area of computational geometry and they have applications for example in computer graphics and database management. In this section we will state our results and discuss the algorithmic motivation (from range search) for the results that are of purely combinatorial nature. (Definitions relevant for the results will be emphasized.) Sketches of proofs will be given in the next section.

Roughly speaking, we investigate the questions:

(1) Which range search problems can be attacked with the partition approach?

(2) Which range search problems allow a transformation into a halfspace-range search problem in some finite-dimensional Euclidean space?

(1) and (2) will be answered affirmatively for range search problems of Vapnik-Chervonenkis dimension 1, while there are (abstract!) range search problems of Vapnik-Chervonenkis dimension 2 or higher which allow no transformation as addressed in (2) and which do not allow partitions as required in the partition approach. (The Vapnik-Chervonenkis dimension of a range search problem will be defined below.)

In an abstract setting a <u>range search problem</u> can be specified by a four-tuple $P = (X,R,\sigma,G)$, where X is a set, R is a set of subsets (called <u>ranges</u>) of X, i.e. $R \subseteq 2^X$, G is a commutative semigroup (a set closed under a commutative and associative addition operation) and σ is a mapping σ: $X \rightarrow G$. The problem is now to find a method for designing a data structure for a given (arbitrary but fixed) finite set A of elements in X which allows us to compute $\sum\{\sigma(x) \mid x \in r \cap A\}$ efficiently for a query range $r \in R$. (We address here a static version of the problem, neglecting insertions and deletions in the set A, see [F].) A method for designing data structures for a range search problem (X,R,σ,G) <u>achieves query time</u> $t(n)$ if for every finite $A \subseteq X$ it produces a data structure for A such that, for all r in R, $\sum\{\sigma(x) \mid x \in r \cap A\}$ can be computed in time at most $t(|A|)$.

In our results we will be mainly concerned with the first two components of a range search problem.

Definition. A range space S is a pair (X,R), where X is a set and $R \subseteq 2^X$. Members of X are called elements or points of S and members of R are called ranges of S. S is finite (countable) if X is finite (countable). □

Example 1. Consider the range search problem $(\mathbb{R}^2, H_2^*, \chi, (\mathbb{N}_0, +))$, where $\mathbb{R}^2$ is the set of points in the real plane, H_2^* is the set of open halfplanes, $(\mathbb{N}_0,+)$ is the semigroup of nonnegative integers with normal addition and $\chi(x) = 1$ for every $x \in \mathbb{R}^2$. Hence, the problem is to design a data structure for a given finite set A of points in the plane such that the number of points of A that lie in a query halfplane $h^* \in H_2^*$ can be determined efficiently. This range search problem is commonly referred to as the halfplane-range counting problem.

The halfplane-range reporting problem is specified by $(\mathbb{R}^2, H_2^*, \varrho, (2^{\mathbb{R}^2}, \cup))$, where $\varrho(x) = \{x\}$ for every $x \in \mathbb{R}^2$. Obviously, here we are interested in the set of points of a given finiteset $A \subseteq \mathbb{R}^2$ that lie in a query halfplane. □

We make the following computational assumptions on a range search problem (X,R,σ,G):

(i) for any $x \in X$ and $r \in R$, $x \in r$? can be decided in constant time,
(ii) for any $A \subseteq X$ and $r \in R$, $A \subseteq r$? and $A \cap r = \emptyset$? can be decided in constant time (which implies (i)), and
(iii) addition in G can be performed in constant time.

(The reader might argue that these assumptions, except for (i), are not realistic in many cases. This is no problem for our lower bound arguments, but it has to be kept in mind if one wants to use the results for upper bounds.)

Taking into account the above assumptions on a range search problem (X,R,σ,G), simply storing a given finite subset A of X in a linear array gives a trivial data structure which needs linear time to answer a query. Unless we have some further information on the semigroup G, this data structure is already optimal, if, for all subsets A' of A, there is a range $r \in R$ with $A' = A \cap r$. Then the number of possible answers (e.g., for reporting) to a query might be as big as $2^{|A|}$ which gives $|A|$ as an information theoretic lower bound for the time needed to answer a query. This naturally leads to the following definitions (see [VC], where these definitions can be found, albeit, in a different context).

Definition. Let S = (X,R) be a range space and let A be a finite set of elements of X. Then $\Pi_R(A)$ denotes the set of all subsets of A that can be obtained by intersecting A with a range in R, i.e. $\Pi_R(A) = \{A \cap r \mid r \in R\}$. If $\Pi_R(A) = 2^A$, then we say that A is shattered by R. The Vapnik-Chervonenkis dimension of S (or simply the dimension of S) is the largest d such that there exists a subset A of X of cardinality d that is shattered by R (if R is empty, then the dimension is -1, and if no such maximal d exists, we say the dimension of S is infinite. □

The (Vapnik-Chervonenkis) dimension of a range search problem (X,R,σ,G) is the dimension of its underlying range space (X,R).

Example 2. Let Σ be an alphabet and let Σ^* be the set of all words over the alphabet Σ (including the empty word λ). For a word $w \in \Sigma^*$, pref(w) denotes the set of all prefixes of w and, for a language $L \subseteq \Sigma^*$, $PREF(L) = \{pref(w) \mid w \in L\}$. (Note that PREF(L) is not the set of all prefixes of words in L!) Obviously, the pair $(\Sigma^*, PREF(\Sigma^*))$ is a range space and we call it the Σ^*-prefix space.

We observe that the dimension of this range space is 1: consider two words w' and w" in Σ^*. If w' is a prefix of w", then {w',w"} cannot be shattered, because there is no word w with w" ∈ pref(w) and w' ∉ pref(w); analogously, if w" is a prefix of w'. If neither of the words w' and w" is a prefix of the other word, then there is no word w with {w',w"} ∩ pref(w) = {w',w"}. Hence no subset of Σ^* of cardinality two (or larger) can be shattered. Since a subset of cardinality one can be shattered, the dimension is 1. ▫

As we have observed above, sublinear query time cannot be achieved for all finite sets of a range search problem of infinite dimension. The natural follow-up question is, whether sublinear query time can be achieved if the dimension is finite. First we quote a result (due to [S], [VC]) which shows that the information theoretic lower bound is at most $\Omega(d \cdot \log n)$ for a range search problem of finite dimension d.

Definition. For d≥0, and n≥0 integers, $\Phi_d(n)$ is defined as follows: $\Phi_d(0) = 1$ for all d≥0, $\Phi_0(n) = 1$ for all n≥0, and $\Phi_d(n) = \Phi_d(n-1) + \Phi_{d-1}(n-1)$ for d,n≥1. ▫

> **PROPOSITION** 1.1. Let (X,R) be a range space of finite dimension d. For every finite subset A of X, $|\Pi_R(A)| \leq \Phi_d(|A|)$. ▫

It remains to observe that $\Phi_d(n) \in \Theta(n^d)$ for d fixed (and, hence, $\log |\Pi_R(A)|$ is $O(d \cdot \log |A|)$).

We will now investigate whether the partition approach always leads to sublinear query time if the dimension of a range search problem is finite. First we describe what we mean by a partition-tree and how a query is performed in such a tree.

Let P be a range search problem and let A be a finite set of elements of P. A partition-tree T_A for A is a tree with the following properties: if A contains exactly one element x, then T_A is a node that stores x and $\sigma(x)$. If $|A| \geq 2$, the root p of T_A stores a partition parts(p) = $(A_1,A_2,\ldots,A_k)$, k≥2, of A into nonempty sets, and for each i, 1≤i≤k, a pointer to a node child(p,i) that is a root of a partition-tree for A_i along with sum(p,i) = $\sum\{\sigma(x) \mid x \in A_i\}$.

In order to compute $\sum\{\sigma(x) \mid x \in r \cap A\}$ for a range r, we start with a global variable SUM, initially set to zero (the neutral element in G), in the root p of the partition tree T_A, add sum(p,i) to SUM for all parts A_i in parts(p) with $A_i \subseteq r$ and continue in the children child(p,i) for all parts A_i which are neither contained in nor disjoint from r. (Note that actually all the leaves in T_A are redundant for the queries!)

If we recall our computational assumptions, then a query in any partition-tree of A for a range r can be performed in linear time. The following definitions will lead to a condition on T_A which ensures sublinear query time.

Definition. Let r and A be sets. We say that r <u>avoids</u> A if either $A \subseteq r$ or $A \cap r = \emptyset$ holds. If r does not avoid A, then we say that r <u>cuts</u> A. ▫

Definition. Let S = (X,R) be a range space, let v and m be integers with 1≤m≤v, and let ε be a real number with $0<\varepsilon<1$.
(i) A partition $(A_1,A_2,\ldots,A_k)$ of a finite subset A of X is called (v,m,ε)-<u>partition</u> (<u>with respect to</u> S), if

- $k \leq v$,
- $|\{i \mid 1 \leq i \leq k,\ r \text{ cuts } A_i\}| \leq m$ for all r in R, and
- $|\bigcup\{A_i \mid 1 \leq i \leq k,\ r \text{ cuts } A_i\}| \leq \varepsilon \cdot |A|$ for all r in R.

(ii) A partition-tree T_A of A is called (v,m,ε)-partition tree of A (with respect to S), if parts(p) is a (v,m,ε)-partition for all nodes p in T_A. ▫

The following proposition can be obtained from calculations in [HW].

> **PROPOSITION** 1.2. Let P be a range search problem with underlying range space S and let A be a set of n

elements of S. If T_A is a (v,m,ε)-partition tree of A with respect to S ($1 \leq m < v$, $0 < \varepsilon < 1$), then any query with a range of S in T_A visits at most

$$\frac{1+\varepsilon}{m-1}\cdot(mn^{\alpha}-1) \in O(n^{\alpha}),\ \alpha = \frac{\log_{1/\varepsilon} m}{1+\log_{1/\varepsilon} m}$$

nodes in T_A for $m \geq 2$, and at most $1 + \log_{1/\varepsilon} n$ nodes in T_A for $m = 1$. □

If we recall our computational assumptions, then Proposition 1.2 shows that a query in a (v,m,ε)-partition tree can be performed in $O(v \cdot n^{\alpha})$ time, α as defined above.

Example 3. Consider the range space $S_2{}^* = (\mathbb{R}^2, H_2{}^*)$ (recall Ex. 1). It is known (see [W]) that for any set A of n points in the plane in general position there are lines ℓ_1 and ℓ_2 (disjoint from A) such that the partition (A_1, A_2, A_3, A_4) induced on A by ℓ_1 and ℓ_2 has the property $\lfloor n/4 \rfloor \leq |A_i| \leq \lceil n/4 \rceil$ for $i = 1,2,3,4$. Note now that every halfplane h^* either contains or is disjoint from one of the quarters in the dissection of the plane induced by ℓ_1 and ℓ_2. Hence, for every h^* there is an i, $1 \leq i \leq 4$, such that h^* does not cut A_i. Now partition A_i in (A_1, A_2, A_3, A_4) with $|A_i| > \lfloor n/4 \rfloor$ into two parts $A_i - \{x\}$ and $\{x\}$ for an arbitrary x in A_i. Then we get a $(7,3,3/4)$-partition $(B_1, B_2, \ldots, B_k)$ of A with respect to $S_2{}^*$; actually such a partition exists for all finite point sets in the plane (not only for those in general position). (Note that the quarter in the dissection of ℓ_1 and ℓ_2 that is not cut by a query halfplane h^* can be determined in constant time!)

Consequently, for every set A of n points in the plane there is a $(7,3,3/4)$-partition tree; it allows to answer a halfplane-range counting query in $O(n^{\alpha})$ time, $\alpha = \log_4 3 \approx 0.79$. From a similar construction it follows that every finite set of points in $\mathbb{R}^d$ allows a $(2^{d+1}-1, 2^d-1, (2^d-1)/2^d)$-partition with respect to $(\mathbb{R}^d, H_d{}^*)$, $H_d{}^*$ the set of open halfspaces in $\mathbb{R}^d$ (see [YY]).

Willard [W] was the first to use the partition-tree approach for higher-dimensional nonorthogonal range search. He actually proved a query time for halfplane-range counting with $\alpha = \log_6 4 \approx 0.77$. Meanwhile this bound has been improved to $\alpha = 2/3+\gamma$ for each $\gamma > 0$, [HW]. □

We have now prepared terminology and context for the first result.

THEOREM A. (1) Let S be a range space of dimension 1. For every finite set A of elements of S there exists a $(15,7,7/8)$-partition of A. (2) There exists a countable range space S of dimension 2 such that for all v, m, and ε ($1 \leq m < v$, $0 < \varepsilon < 1$) there is a finite set A of elements of S which does not allow a (v,m,ε)-partition. □

We turn now to the second problem we want to consider. Yao and Yao [YY] demonstrated that a number of range search problems can be transformed into (embedded in) a halfspace range search problem of some finite (Euclidean) dimension d. We investigate, whether this is always possible for range search problems of finite Vapnik-Chervonenkis dimension. To be precise, we settle some terminology.

Definition. Let $S = (X,R)$ and $S' = (X',R')$ be range spaces. We say S is <u>embeddable in</u> S', in symbols $S \lhd S'$, if there is an injective mapping $\eta : X \to X'$ and a mapping $\varrho: R \to R'$, such that for all $r \in R$,

$$\eta\hat{}(r) = \eta\hat{}(X) \cap \varrho(r),$$

where $\eta\hat{}$ is the extension of η to subsets A of X by $\eta\hat{}(A) = \{\eta(x) \mid x \in A\}$. □

Definition. For $d \geq 1$, let $H_d{}^+$ be the set of positive halfspaces in $\mathbb{R}^d$. (A halfspace h^* is one of the two open regions in the dissection of $\mathbb{R}^d$ induced by a hyperplane h. h^* is a positive halfspace if either h is vertical or h^* is the open halfspace above h, i.e. h^* intersects the positive vertical axis in a halfline.)

By $S_d{}^+$ we denote the range space $(\mathbb{R}^d, H_d{}^+)$. □

Example 4. Consider the range space $S_1 = (\mathbb{R}, I)$ with I the set of open intervals on the real line. For $a \in \mathbb{R}$ let $\eta(a)$ be the point $(a, -a^2)$ in $\mathbb{R}^2$, and for the interval (a,b) in I, let $\varrho((a,b))$ be the positive halfplane bounded by the line through points $(a,-a^2)$ and $(b,-b^2)$. Then η and ϱ realize an embedding of S_1 in S_2^+; hence $S_1 \triangleleft S_2^+$. □

Example 5. Let $\Sigma = \{0,1\}$, and let $L = \{w_1, w_2, w_3, w_4, w_5, w_6\} \subseteq \Sigma^*$ where $w_1 = 0$, $w_2 = 00$, $w_3 = 01$, $w_4 = 10$, $w_5 = 001$, and $w_6 = 000$. Let $R = \Pi_{PREF(\Sigma^*)}(L)$, i.e. $R = \{r_i \mid i=1,2,\ldots,7\}$, where $r_1 = \emptyset$, $r_2 = \{0\}$, $r_3 = \{0,00\}$, $r_4 = \{0,00,001\}$, $r_5 = \{10\}$, $r_6 = \{0,01\}$, and $r_7 = \{0,00,000\}$.

Fig. 1.1 shows that $(L,R) \triangleleft S_2^+$, (with $\varrho(r_i) = h_i^+$, $1 \leqslant i \leqslant 7$).

More generally, it can be shown that every Σ^*-prefix space is embeddable in S_2^+. □

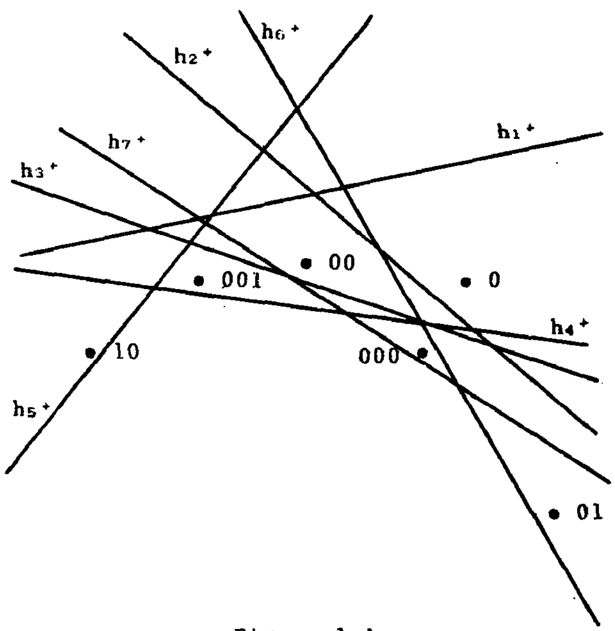

Figure 1.1

If $(X,R) \triangleleft (X',R')$ and every finite $A' \subseteq X'$ has a (v,m,ε)-partition with respect to (X',R'), then, of course, a (v,m,ε)-partition exists for every $A \subseteq X$ with respect to (X,R). More generally, we can transform data structures for a range search problem with underlying range space (X',R') into data structures for a range search problem with underlying range space (X,R); (the only crucial point is that the embedding mapping for the ranges is effective). Since partitioning results for the range spaces S_d^+ are known ([YY], [HW]), we are interested in embeddings of a range space in S_d^+. Our result is the following (all statements are best possible):

THEOREM B. (1) Let S be a countable range space of dimension 1. Then $S \triangleleft S_3^+$. (2) Let S be a countable range space of dimension 1 such that the empty set is a range of S. Then $S \triangleleft S_2^+$. (3) There is a countable range space S of dimension 2 such that $S \triangleleft S_d^+$ holds for no $d \geqslant 1$. □

Example 6. It is interesting to apply Theorem B(2) to Σ^*-prefix spaces: if the empty word is omitted in every range of a Σ^*-prefix space, then the assumptions of Theorem B(2) are satisfied and it follows that a $(7,3,3/4)$-partition exists for every finite language over Σ. □

2. Proofs of the theorems.

The reader might have already realized that the two results we presented are closely related. So Theorem B(1) immediately implies Theorem A(1), since we know that every finite point set in $\mathbb{R}^3$ allows a $(15,7,7/8)$-partition. It is also clear that Theorem A(2) entails Theorem B(3) (due to the partition results in [YY] and [HW]). So we will sketch the proofs for Theorems A(2) and B(2). The proof of B(1) is omitted.

<u>Proof of Theorem A(2)</u>.

In order to show Theorem A(2) we consider finite projective planes. Recall that the finite projective plane PG(2,q) is a set V of q^2+q+1 elements (called points) and a family L of q^2+q+1 subsets of V (called lines), such that each $\ell \in L$ has cardinality $q+1$, each x in V belongs to exactly $q+1$ lines in L, and for every pair x and x' of distinct points in V there is a unique line $\ell \in L$ with $x, x' \in \ell$. It is well known (see [H]) that PG(2,q) exists for every prime power q. Clearly, PG(2,q) = (V,L) is a range space in our sense. No set of three points of PG(2,q) is shattered, while two points can be shattered; hence, PG(2,q) has Vapnik-Chervonenkis dimension 2. The key property of finite projective planes (for our purposes) is formulated in the following proposition (see [A1], [A2], and also [L ,Problem 13.13]).

PROPOSITITON 2.1. Let (V,L) denote the range space PG(2,q) and let $n = q^2+q+1$, i.e. $|V| = |L| = n$. For every $A \subseteq V$, $A \neq \emptyset$, the inequality $|\{\ell \in L : \ell \cap A = \emptyset\}| < n^{3/2}/|A|$ holds. □

We use Proposition 2.1 to prove that, for PG(2,q) = (V,L), every (v,m,ε)-partition $(V_1, V_2, \ldots V_k)$ of V with respect to PG(2,q) satisfies $n = q^2+q+1 < (v^2/(1-\varepsilon))^2$.

To this end let $c = (1-\varepsilon)\cdot n/v$. If $c \leq q+1$, then $n/(q+1) \leq v/(1-\varepsilon)$ which implies by a short calculation that $n < (v^2/(1-\varepsilon))^2$ (use $v \geq 2$), and we are done. So let us assume that $c > q+1$. First we observe that every ℓ in L avoids a V_i, $1 \leq i \leq k$, with $|V_i| \geq c$; otherwise the union of all parts avoided by ℓ contains less than $k \cdot c \leq v \cdot c = (1-\varepsilon)\cdot n$ elements and so the union of all parts cut by ℓ contains more than $\varepsilon \cdot n$ elements; a contradiction to the fact that we are dealing with a (v,m,ε)-partition.

Let L' be the union of all $\{\ell \in L : \ell \cap V_i = \emptyset\}$ over all i with $|V_i| \geq c$. Then $|L'| < v \cdot n^{3/2}/c$, by Proposition 2.1, i.e. $|L'| < v^2 \cdot n^{1/2}/(1-\varepsilon)$. Since we assumed that $c > q+1$, no ℓ in L can avoid a V_i with $|V_i| \geq c$ by containing it. Consequently, by our first observation, L' = L. Thus $n < v^2 \cdot n^{1/2}/(1-\varepsilon)$ which gives the claimed bound.

If we let now $\{q_i : i=1,2,\ldots\}$ be an infinite set of prime powers, then the range space S that is the disjoint union of all $PG(2,q_i)$, i=1,2,..., demonstrates Theorem A(2).

<u>Proof of Theorem B(2)</u>.

Before we show how to embed a countable range space S = (X,R) of dimension 1 with $\emptyset \in R$, we "force a rigid structure" on S by adding ranges to R without increasing the dimension.

LEMMA 2.2. Let S = (X,R) be a range space of dimension 1 with $\emptyset \in R$, let $\emptyset \neq T \subseteq R$, let $r_T = \bigcap\{r : r \in T\}$ and let $R' = R \cup \{r_T\}$. Then $\Pi_R(A) = \Pi_{R'}(A)$ for every $A \subseteq X$ of cardinality two.

<u>Proof</u>. Consider two elements x and y of S. We distinguish four cases.
<u>Case</u> 1. $\{x,y\} \cap r_T = \{x,y\}$: For $r \in T$, this implies $\{x,y\} \cap r = \{x,y\}$ and so $\Pi_R(\{x,y\}) = \Pi_{R'}(\{x,y\})$.
<u>Case</u> 2. $\{x,y\} \cap r_T = \emptyset$: As $\{x,y\} \cap \emptyset = \emptyset$ the assertion holds.
<u>Case</u> 3. $\{x,y\} \cap r_T = \{x\}$: Then $x \in r$ for all r in T and there is a r_1 in T with $y \notin r_1$. For this range r_1 we have $\{x,y\} \cap r_1 = \{x\}$ and so $\Pi_R(\{x,y\}) = \Pi_{R'}(\{x,y\})$.
<u>Case</u> 4. $\{x,y\} \cap r_T = \{y\}$: Analogous to Case 3. □

Definition. Let S = (X,R) be a range space.
(i) S is <u>intersection-closed</u> if $\bigcap\{r : r \in T\}$ is in R for all $\emptyset \neq T \subseteq R$.
(ii) S is a <u>standard range space</u>, if
- S is of dimension 1,
- $\emptyset \in R$,
- S is intersection-closed,
- for all $x \in X$ there is an $r \in R$ with $x \in r$, and
- for all $x, y \in X$, $x \neq y$, there is an $r \in R$ with $|r \cap \{x,y\}| = 1$. □

LEMMA 2.3. Let S = (X,R) be a range space of dimension 1 with $\emptyset \in R$. Then there is a standard range space S' = (X,R') with $R \subseteq R'$.

<u>Proof</u>. Let $R'' = R \cup \{r_T : \emptyset \neq T \subseteq R\} \cup \{\{x\} : x \in X$ and there is no $r \in R$ with $x \in r\}$, where $r_T = \bigcap\{r : r \in T\}$. Then (X,R") is intersection closed and, by Lemma 2.2, (X,R") is still of dimension 1. It remains to add ranges such that for all x,y in X, $x \neq y$, there is a range r such that $|r \cap$

$\{x,y\}| = 1$ (without violating any of the other properties of a standard range space we have obtained already).

For each x in X, let r(x) denote the range $\bigcap\{r \in R'' \mid x \in r\}$ in R" (recall that R" is closed under intersections!) and let [x] denote $\{y \in X \mid r(x)=r(y)\}$. For every $E = [x]$, $x \in X$, let $r_E = r(x)$, let $<_E$ be a total order on E, and we set $R_E = \{r_E - \{y \in E \mid z <_E y\} \mid z \in E\}$. Then it can be shown that $R' = R'' \cup \bigcup\{R_E \mid E=[x], x \in X\}$ satisfies the statement of the lemma. (We omit here further datails, since they lead to a case analysis like in the proof of Lemma 2.2. □

Lemmma 2.3 shows that we can restrict ourselves to countable *standard* range spaces, if we are interested in embeddings of countable range spaces of dimension 1 with the empty set a range. Obviously, if S = (X,R) is a standard range space, then $(A,\Pi_R(A))$ is a standard range space for every subset A of X. We will see now that finite standard range spaces have a simple correspondence to directed out-trees.

Definition. The 1-inclusion graph I_S of a range space S = (X,R) is the directed edge labeled graph with the set R as node set and $\{(r,r') \mid r,r' \in R,\ r \neq r',\text{ and } r'=r \cup \{x\}$ for some $x \in X\}$ as edge set. The label of the edge (r,r') is the unique element x in r' - r. □

> **LEMMA** 2.4. Let S = (X,R) be a finite standard range space. Then I_S is a directed out-tree (all but exactly one node, the root, have exactly one ingoing edge) where every element in X occurs exactly once as label.

Proof. For x in X, let $r(x) = \bigcap\{r \in R \mid x \in r\}$ which is a range in R since S is intersection-closed. It also follows that if $x \neq y$, then $r(x) \neq r(y)$ since there is a range r in R that contains exactly one of the elements x and y. We conclude that $|R| \geq |X|+1$ (recall that also $\emptyset$ is a range in R!). By Proposition 1.1, this entails actually that $|R| = |X|+1$ (since $\Phi_1(n) = n+1$). A simple case analysis shows that if we add r(x) - {x} to R (for any x in X) then this does not increase the dimension of the resulting range space. Since R has already the maximal cardinality for a range space of dimension 1, this implies that r(x) - {x} is already in R and (r(x)-{x},r(x)) is an edge in I_S with label x. Assume now that there is an y in X, $y \neq x$, such that r(x) - {y} is in R; then {x,y} is shattered by $\{\emptyset, r(x)-\{x\}, r(x)-\{y\}, r(x)\} \subseteq R$; a contradiction. Consequently, I_S has exactly |X| edges, i.e. exactly one edge less than there are nodes. It is easily seen that there is a path from the node in I_S that corresponds to $\emptyset$ to every other node in I_S; the lemma follows. □

Before we proceed now to the embedding of countable standard range spaces, we still need some geometrical notation.

Definition. (i) If y and y' are two distinct points in $\mathbb{R}^2$, then h(y,y') denotes the line through y and y'.
(ii) If h is a non-vertical line in $\mathbb{R}^2$, then h^+ denotes the open halfplane bounded by h that lies above h and h^- denotes the open halfplane bounded by h that lies below h.
(iii) If s is a region in $\mathbb{R}^2$, i.e. $s \subseteq \mathbb{R}^2$, then int(s) denotes the interior of s and cl(s) denotes the closure of s. □

Definition. Let y, y', and y" be three distinct points in $\mathbb{R}^2$, no two on a common vertical line. If $y \in h(y',y'')^+$, $y' \in h(y,y'')^-$ and $y'' \in h(y,y')^-$, then the convex hull t of {y,y',y"} is termed a <u>normal triangle</u>. y is called the <u>top corner of</u> t. The <u>corridor of</u> t, corr(t), is the closure of

$$(h(y,y')^- \cup h(y,y'')^-) - (h(y,y')^- \cap h(y,y'')^- \cap h(y',y'')^-).\ □$$

Fig 2.1 depicts a normal triangle t with top corner y and its corridor c.

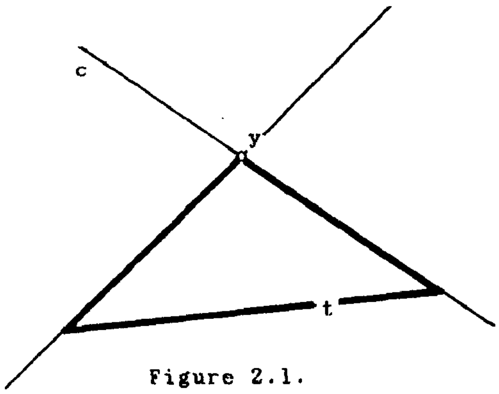

Figure 2.1.

Let now $S = (X,R)$ be an arbitrary but fixed countable range space and let $\{x_1, x_2, \dots\}$ be an enumeration of the elements in X. For x_i in X, let $r_i = \bigcap\{r \in R \mid x_i \in r\}$. Since S is a standard range space, r_i is in R for every positive integer i.

To obtain our embedding of S in S_2^+, we will associate with every x_j two normal triangles t_j and $\hat{t}_j$ with their corridors $c_j = corr(t_j)$ and $\hat{c}_j = corr(\hat{t}_j)$ such that the following holds:

(1) $\hat{t}_j \subseteq int(t_j)$ and $\hat{c}_j \subseteq int(c_j)$ for all positive integers j.

(2) if $x_j \in r_i$, then $t_i \subseteq int(\hat{t}_j)$ and $c_i \subseteq int(\hat{c}_j)$ for all positive integers i and j, $i \neq j$.

(3) if $x_j \notin r_i$ and $x_i \notin r_j$ then t_j lies below c_i and t_i lies below c_j.

Figs. 2.2 and 2.3 illustrate the situations (2) and (3), respectively.

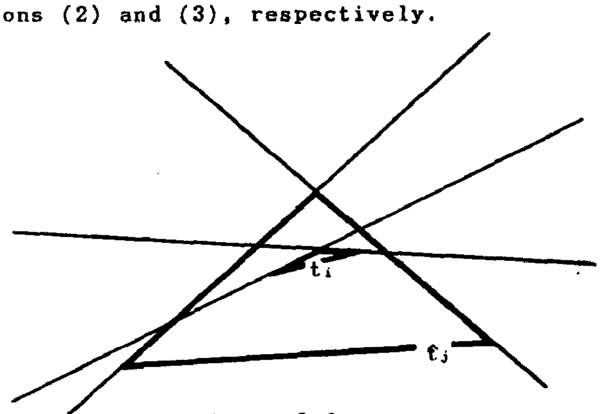

Figure 2.2.

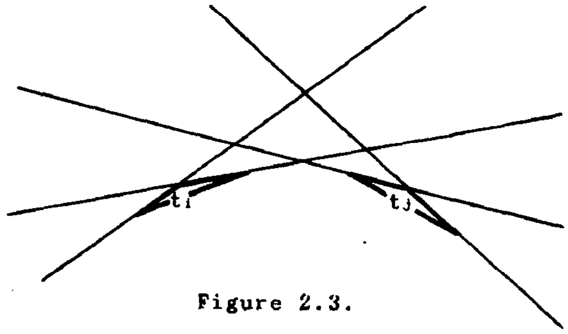

Figure 2.3.

Before we demonstrate how we can find the triangles t_j and $\hat{t}_j$, we show that they imply already the required embedding. To this end let y_j be the top corner of t_j and let h_j be a line through the top corner of $\hat{t}_j$ such that $h_j^+ \cap \hat{t}_j = \emptyset$ and such that $h_j \subseteq \hat{c}_j$. Then it follows from the conditions (1) through (3) above that $y_j \in h_i^+$ if and only if $x_j \in r_i$. Hence, the y_j's and h_j^+'s constitute already an embeddding of X for all ranges r_i in R. In order to show that we find positive half-planes h_r^+ for every r in R such that $\{j \mid y_j \in h_r^+\} = \{j \mid x_j \in r\}$, consider an abitrary r in R and let $t = \bigcap\{\hat{t}_i \mid x_i \in r\}$ and $c = \bigcap\{\hat{c}_i \mid x_i \in r\}$. It is not too difficult to show that there is a line $h_r \subseteq c$ such that $t \cap h_r^+ = \emptyset$. Let x_i be an element in r (and so $r_i \subseteq r$). Then y_i lies above $\hat{c}_i$; since $c \subseteq \hat{c}_i$, this shows that $y_i \in h_r^+$. If x_i is not an element in r, then we distinguish two cases:

<u>Case</u> 1. $x_j \in r_i$ for all x_j in r: Then $t_i \subseteq int(\hat{t}_j)$ for all x_j in r and so y_i lies in t; consequently, $y_i \notin h_r^+$.

<u>Case</u> 2. $x_j \notin r_i$ for a x_j in r: Since $x_i \notin r$, we have also $x_i \notin r_j$. Thus, t_i lies below c_j and so it lies below h_r; consequently, again $y_i \notin h_r^+$.

We conclude, that if we find triangles satisfying the conditions (1) through (3) then we have established an embedding of S in S_2^+. For this last step in the proof we use the correspondence to directed out-trees as stated in Lemma 2.4; however, it will be necessary to deal with ordered out-trees and we need the following definition.

Definition. Let T be an ordered directed out-tree, let $e = (p,q)$ be an edge in T, let $(p_1, p_2, \ldots, p_k)$ and $(q_1, q_2, \ldots, q_m)$ be the sequences of children of p and q, respectively, and let $q = p_i$. Then the tree T-e is obtained from T by contracting the edge e, i.e. we identify q and p and assign to p the sequence $(p_1, p_2, \ldots, p_{i-1}, q_1, q_2, \ldots, q_m, p_{i+1}, \ldots, p_k)$ of children. □

For every non-negative integer j, let $A_j = \{x_i \in X : i \leq j\}$ (thus $A_0 = \emptyset$) and let T_j be an ordered directed labeled out-tree such that (i) for all non-negative integers j, the underlying unordered tree of T_j is isomorphic to the 1-inclusion graph of $(A_j, \Pi_R(A_j))$, and (ii) for all positive integers j, T_{j-1} is isomorphic to $T_j - e_j$, where e_j is the edge in T_j that is labeled by x_j. By Lemma 2.4, (i) can be obtained. It is not hard to achieve (ii), so we omit here the details.

First let t_0 and $\hat{t}_0$ be two normal triangles with $\hat{t}_0 \subseteq int(t_0)$ and $corr(\hat{t}_0) \subseteq int(corr(t_0))$. While constructing now t_j and $\hat{t}_j$ for $j = 1, 2, \ldots$ we observe the following additional condition. Let p be a node in T_j, let i be the index of the label of its ingoing edge (or $i = 0$ if p is the root of T_j) and let $i_1, i_2, \ldots, i_k$ be the indices of the labels of the edges outgoing from p (in this order). Then the triangles $t_{i_1}, t_{i_2}, \ldots, t_{i_k}$ form a "convex sequence" inside $\hat{t}_i$ in this order. Note that the only additional condition we impose here is that also the ordering has to be observed. (See Fig. 2.4 for an illustration.)

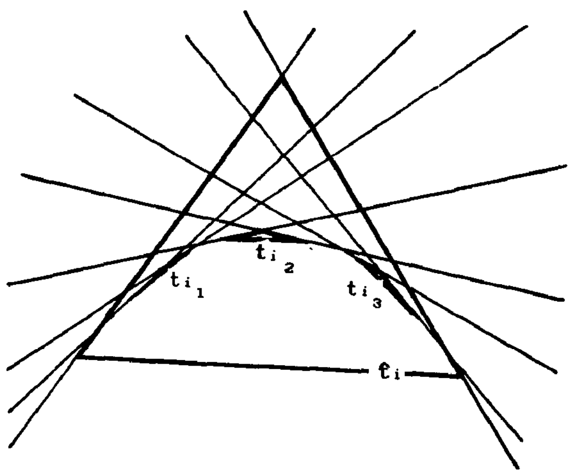

Figure 2.4.

Suppose now we have already defined all triangles t_i for $i \leq j-1$. Let (p,q) be the edge in T_j with label x_j, and let i be the index of the incoming edge of p (or $i = 0$ if p has no incoming edge). We distinguish two cases:

<u>Case</u> 1. q has at least one child in T_j: t_j has to be chosen inside $\hat{t}_i$ such that it contains all t_k, for which x_k is a label of an outgoing edge of q in T_j (and the same relations must hold for the corresponding corridors). Since space is limited, we cannot go into further details here; we only mention that the essential point is that all the triangles that must lie inside t_j form a (closed) convex subsequence inside $\hat{t}_i$ due to our order condition.

<u>Case</u> 2. q has no child in T_j: t_j has to be chosen inside $\hat{t}_i$ in such a way that the order condition is satisfied and, moreover, all triangles inside $\hat{t}_i$ lie below the corridor c_j of t_j and t_j lies below the corridors of all the triangles inside $\hat{t}_i$.

Finally, the triangle $\hat{t}_j$ can be obtained by a small (appropriate) perturbation of t_j.

This eventually concludes the proof of Theorem B(2).

Concerning the proof of Theorem B(1), we mention here only that it is based on B(2) by the following transformation: For sets A and B, let $A \Delta B$ denote the symmetric difference $(A-B) \cup (B-A)$ of A and B. For a range space $S = (X,R)$ and $A \subseteq X$, let $S \Delta A = (X, R\Delta A)$, where $R \Delta A = \{r\Delta A : r \in R\}$. It can be shown that the Vapnik-Chervonenkis dimension of $S \Delta A$ equals the Vapnik-Cervonenkis dimension of S, and that if $S \Delta A$ is embeddable in S_d^+ then S is embeddable in S_{d+1}^+. Now it remains to observe that if A is a range in S, then $\emptyset$ is a range in $S \Delta A$.

3. Open problems.

We conclude the paper with two problems which arise in connection with the results we presented.

(1) Is it possible to extend the positive embedding results beyond "countable"?

(2) Is it possible, that -- despite of our negative partition result for finite projective planes -- there are partition-trees for finite projective planes (or for range spaces of dimension 2 in general) where every range query visits only a sublinear number of nodes (and the number of children in the trees is bounded)? Note that this does *not* follow from our result!

References.

[A1] Alon, N., Expanders, sorting in rounds and superconcentrators of limited depth, *in* "Proc. 17th ACM Symp. on Theory of Comput." (1985) 98-102.

[A2] Alon, N., Eigenvalues, geometric expanders, sorting in rounds and Ramsey theory, *Combinatorica* 6 (1986) 207-219.

[F] Fredman, M.L., Lower bounds on the complexity of some optimal datastructures, *SIAM J. Comput.* 10 (1981) 1-10.

[H] Hall, M.Jr., *Combinatorial Theory*, Wiley and Sons, New York, London (1967).

[HW] Haussler, D. and Welzl, E., Epsilon-nets and simplex range search, *Discrete and Comput. Geom.*, to appear.

[L] Lovasz, L., *Combinatorial Problems and Exercises*, North-Holland, Amsterdam, New York, Oxford (1979).

[S] Sauer, N., On the density of families of sets, *J. Combin. Theory Ser. A* 13 (1972) 145-147.

[VC] Vapnik, V.N. and Chervonenkis, A.YA., On the uniform convergence of relative frequencies of events to their probabilities, *Theory Probab. Appl.* 16 (1971) 164-280.

[W] Willard, D., Polygon retrieval, *SIAM J. Comput.* 11 (1982) 149-165.

[YY] Yao, A. and Yao, F., A general approach to d-dimensional geometric queries, *in* "Proc. 17th ACM Symp. on Theory of Comput." (1985) 163-169.