On Point Covers of c-Oriented Polygons*

Franck NIELSEN
SONY Computer Science Laboratory Inc.
E-mail: nielsen@csl.sony.co.jp

Abstract

Let $\mathcal{S}$ be any family of n *c-oriented* polygons of the 2-dimensional Euclidean plane $\mathbb{E}^2$, i.e., bounded intersection of halfplanes whose normal directions of edges belong to a fixed collection of c distinct directions. In this paper, we prove that $\tau(\mathcal{S}) \leq \mathfrak{G}(2,c)\phi(\mathcal{S})\log_2^{c-1}(\phi(\mathcal{S})+1)$, where $\phi(\mathcal{S})$ is the packing number, $\tau(\mathcal{S})$ the transversal number and $\mathfrak{G}(2,c)$ the Gallai number of pairwise intersecting c-oriented polygons. Our bound collapses to $\tau(\mathcal{S}) = O(\mathfrak{G}(2,c)\phi(\mathcal{S}))$ if objects are of more or less of the same size. We describe a $t(n,c) + O(nc\log\phi(\mathcal{S}))$-time algorithm with linear storage that computes such a 0-transversal, where $t(n,c)$ is the time required to pierce pairwise intersecting c-oriented polygons. We provide linear-time algorithms $t(n,c) = \Theta(nc)$ for α-fat c-oriented polytopes, translates or homothets of $\mathbb{E}^d$ proving that $\mathfrak{G}(2,c) = O(\alpha)^d$, $\mathfrak{G}(2,c) \leq d^d$ and $\mathfrak{G}(2,c) \leq (3d^{\frac{3}{2}})^d$ respectively.

Keywords: computational geometry, output-sensitive algorithms, precision-sensitive heuristics, transversal and packing numbers.

1 Introduction and preliminary notations.

Let $\mathcal{S} = \{P_1, ..., P_n\}$ be a set of n d-dimensional geometric objects of the Euclidean space $\mathbb{E}^d$ and $\mathcal{P}$ be a point set. We say that $\mathcal{P}$ is a *0-transversal* of $\mathcal{S}$ if and only if every object P_i is pierced by $\mathcal{P}$, i.e., $P_i \cap \mathcal{P} \neq \emptyset$. $\mathcal{P}$ is said a *covering* or *stabbing* point set. The transversal number of $\mathcal{S}$ is defined as the minimum size of any 0-transversal of $\mathcal{S}$. Finding the minimum k so that $\mathcal{S}$ can be *pierced* (i.e., stabbed) by k points has been shown to be NP-complete [FPT81] as soon as $d \geq 2$. Even in the one-dimensional case, this problem remains NP-complete for non-convex instances [HM87]. The geometric covering/piercing problem is also referenced in the literature as the *set covering problem* (SCP), or dually as the *hitting set problem* (HSP), where it is transformed into an optimization problem by means of matrix formulations. In this section, we will briefly sum up previously known complexity results (Section 1.1) and heuristics on abstract sets (Section 1.2). Then we will emphasize on its geometric counterpart (Section 1.3). Finally to conclude the first section, we will introduce Gallai numbers (Section 1.4). Section 2 defines c-oriented and b-stripped polytopes and give some of their fundamental properties. In Section 3, we give an output-sensitive precision-sensitive algorithm for the geometric SCP of c-oriented polygons and study its time complexity as well as its performance ratio. Section 4 describes how the algorithm can be slightly modified in order to compute independent sets (i.e., pairwise non-intersecting objects) of precision-sensitive size (although the counterpart on abstract graph is $\Omega(n^{\frac{1}{4}})$ hard to approximate in polynomial time). The *packing number* of $\mathcal{S}$ is defined as the maximal size of any independent set of $\mathcal{S}$. Finally, Section 5 summarizes up-to-date results on the geometric SCP/HSP.

1.1 Abstract set aspects.

Let $\mathcal{V} = \{S_i | i \in I\}$ be a collection of $v = |\mathcal{V}| = |I|$ subsets of $2^{\mathcal{S}}$ for a set $\mathcal{S}$ of n elements. We want to find a minimal covering collection, i.e., a subset $I' \subseteq I$ of indices such that $\mathcal{S} = \bigcup_{i \in I'} S_i$ with $|I'|$ as small as possible. In other words, we want to minimize $e_v^T \times x = |I'|$ subject to $Ax \geq e_n$ for x a $\{0,1\}^v$-vector,

*Part of this work has been done at École Polytechnique (France) and INRIA Sophia-Antipolis, Project PRISME (France).

e_k a k-dimensional vector of 1's and A a $(v \times n)$-binary matrix, each column of which is the incidence vector of one of the sets I_i, $1 \leq i \leq v$. The above formulation gives an *integer linear program*. Using these notations, the set cover system is said ϵ-*dense* (δ-*super-dense*) if $|S_i| \geq \epsilon|\mathcal{S}|$ (resp. $|S_i| \geq |\mathcal{S}| - o(m^\delta)$), for all $i \in I$. The SCP remains Max SNP-hard even on dense cases but can be solved polynomially in super-dense cases [AKK95, KZ96].

1.2 Heuristics on abstract sets.

Chvátal [Chv79, Sla96] gave a quadratic-time greedy algorithm to find a cover set of size $\Delta(\mathcal{S}) \leq \tau(\mathcal{S})(\log k - \log\log k + 0.78)$, where k is the maximum column sum ($k \leq n$) of A. Notice that SCP can be solved in polynomial time whenever $k \leq 2$ (using a maximum matching algorithm in a bipartite graph) and that k-SCP have better heuristics using semi-local optimization [DF97]. An algorithm is *precision-sensitive* if its performance ratio does not depend on the input size but rather on the size of any optimal solution (greedy algorithm is not precision-sensitive). Hochbaum [Hoc82] proposed a cubic-time algorithm with a cover set of size at most $\tau(S)f$, where f is the maximum row sum of incidence matrix A using an analytical Russian method and a linear program relaxation. Interestingly, Feige [Fei96] showed that no polynomial-time algorithm can approximate the optimal solution within a factor of $(1-\epsilon)\log|\mathcal{S}|$, unless $NP \subseteq DTIME[n^{\log\log n}]$, where $1 > \epsilon > 0$. Considering A as an hypergraph, we have $\tau(\mathcal{S}) \leq 11\lambda(\mathcal{S})^2(\lambda(\mathcal{S}) + \phi(\mathcal{S}) + 3)\binom{\lambda(\mathcal{S})+\phi(\mathcal{S})}{\phi(\mathcal{S})}$, where $\lambda(\mathcal{S})$ is the maximum l so that the incidence matrix A has as a submatrix the incidence matrix of the complete graph $\mathcal{K}_l$ (clique of size l) [DSW94].

1.3 Heuristics on geometric instances.

One major drawback from the computational geometry point of view is that those methods do not consider geometrical objects nor their shapes and require matrix A. One way to proceed is to consider from the whole arrangement of the constant-size descriptive objects all the sets induced by k-faces, $0 \leq k \leq d$. We label each k-face with the set of objects fully containing it. A label is said *maximal* if it is not included in another one. We remove non-maximal labels and obtain a so-called *Sperner system*, still possibly of size $O(n^d)$ [KN96]. Hochbaum and Maass [HM84] considered the case of geometrical objects and gave a *polynomial-time approximation scheme*. Their algorithm allows us to consider sets of congruent star-shaped centrally symmetric objects T, or dually covering sets of points with star-shaped translates $T^* = T$. In that context, piercing families of c-oriented translates of a given polytope is of particular interest since it corresponds to covering a set of points by a minimum number of congruent copies of a given polytope. Brönninman and Goodrich [BG94] investigate these problems using the concept of the *Vapnik-Červonenkis dimension*. They obtain precision-sensitive set covers if the VC-dimension is bounded as it is usually the case when considering geometric scenes. Their evolutionary algorithm still relies on the fact that matrix A is computed beforehand. Recently, Efrat et al. [EKNS97] studied dynamic data-structures for fat objects and obtain efficient piercing algorithms under the fatness assumptions in dimension 2 and 3. As a byproduct, we have $\tau(\mathcal{S}) = O(\phi(\mathcal{S}))$ for fat objects in arbitrary fixed dimension. The case of isothetic boxes in arbitrary dimension has been studied in [Kár91, dFK93, KT95, Nie96]. It is worth noting that for d-dimensional boxes, we have the following inequality:

$$\tau(\mathcal{S}) \leq \phi(\mathcal{S}) \log_2^{d-1} \phi(\mathcal{S}) + d - \frac{1}{2}\phi(\mathcal{S}) \log_2^{d-2} \phi(\mathcal{S}).$$

Moreover, efficient $O(dn \log \phi(\mathcal{S}))$-time $O(dn)$-space output-sensitive algorithms have been designed to compute such a 0-transversal in that case. In this paper, we generalize the methodology we used for boxes to c-oriented polygons. Observe that for general objects, there is no relationship between $\tau(\mathcal{S})$ and $\phi(\mathcal{S})$. For example, let us take $\mathcal{S}$ as n pairwise intersecting segments in non-degenerate position then it is clear that $\tau(\mathcal{S}) = \lceil \frac{n}{2} \rceil$ and $\phi(\mathcal{S}) = 1$. This observation can be extended to polytopes by transforming each segment to a polytope obtained as the Minkowski sum of it plus a tiny d-dimensional cube so that the overall resulting set of polytopes has the same intersection graph, see Figure 1. Therefore, it is of particular interest to focus on the restricted case of c-oriented polygons for which we obtain the first (to our knowledge) non-trivial upper bounds.

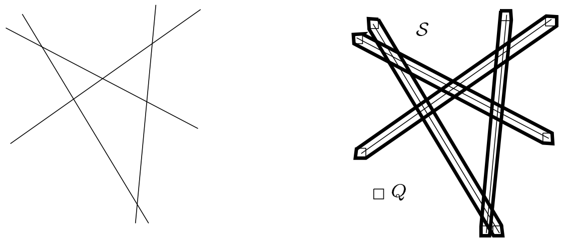

Figure 1: A family $\mathcal{S}$ of n pairwise intersecting polytopes ($\phi(\mathcal{S}) = 1$) that requires $\tau(\mathcal{S}) = \lceil \frac{n}{2} \rceil$ points to pierce them.

1.4 Introducing Gallai numbers.

Let $\prod^q$ denotes the q-pierceability property, i.e., being pierceable by q points. We define Gallai numbers as follows:

Definition 1 *If every subset of $\mathcal{S}$ ($\mathcal{S} \in \mathcal{C}$) of size p is $\prod^q$ (let $\prod_p^q$ denote that property) then $\mathcal{S}$ is pierceable with $\mathfrak{G}(p,q,\mathcal{C})$ points, where $\mathfrak{G}(p,q,\mathcal{C})$ is the Gallai number of class $\mathcal{C}$ for the $\prod^q$ property of its p-subsets.*

For example $\mathfrak{G}(2,1,\mathcal{D}^2) = 4$, where $\mathcal{D}^2$ is the class of planar disks.[1] In the sequel, $\mathfrak{G}(2,c)$ will denote the Gallai number of pairwise intersecting c-oriented polytopes.

2 Properties of c-oriented polytopes of $\mathbb{E}^d$.

An object $P \in \mathcal{S}$ is said to be c-oriented if it can be expressed as the intersection of translated halfspaces of a given family $\mathcal{H} = \{H_1, ..., H_c\}$ of c halfspaces. We denote by $\partial\mathcal{H}$ the corresponding family of c bounding hyperplanes: $\partial\mathcal{H} = \{\partial H_1, ..., \partial H_c\} = \{h_1, ..., h_c\}$. A geometric object P is a *c-oriented polytope* if it is both c-oriented and bounded. Next, we introduce *b-stripped* objects as polytopes defined by the intersection of at most b strips, where a *strip* is the intersection of two halfspaces whose corresponding bounding hyperplanes are translates of each other. Thus, d-boxes are $2d$-oriented but d-stripped. A strip S is described by means of a d-dimensional vector a and two real values $\alpha \le \beta$ such that $S = \{x \in \mathbb{E}^d | \alpha \le a \cdot x \le \beta\}$ (see Figure 2). Each c-oriented polytope $P_i \in \mathcal{S}$ is the product of $b \le c$ strips: $P_i = \prod_{j \in \{1,...,b\}} strip(a_j, \alpha_{i,j}, \beta_{i,j})$. The

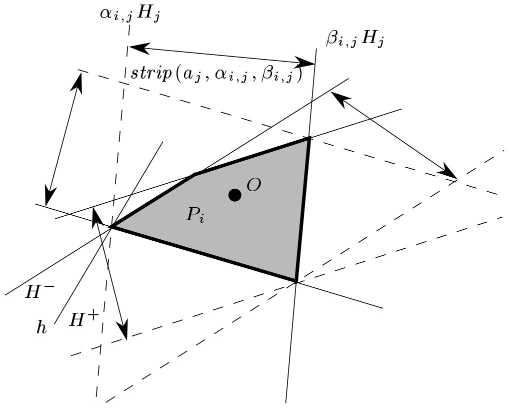

Figure 2: A 4-oriented polytope P_i, the supporting hyperplanes of its facets and the four induced strips.

[1] For congruent pairwise intersecting disks, three points suffice. Surprisingly, it has been shown by Grünbaum that $\mathfrak{G}(2,1,\mathcal{D}^d) > 1.003^d$, where $\mathcal{D}^d$ is the family of d-dimensional balls.

intersection of two b-stripped polytopes $P \cap P'$ is also a b stripped polytopes. Note that $P \cap P'$ may be empty even if the cartesian product of its strips is not empty, i.e.

$$\prod_{j \in \{1,\ldots,b\}} [\max\{\alpha_j, \alpha'_j\}, \min\{\beta_j, \beta'_j\}] \neq \emptyset.$$

Clearly, the converse is true only for boxes (up to a projective transformation, see [PS85], page 329). We denote by $strip_l(P)$ the real interval $[\alpha_l, \beta_l]$. A **non-empty polytope** P' lies within P if and only if $strip_l(P') \subseteq strip_l(P)$ for all $l \in \{1, ..., b\}$.

Theorem 1 *Two non-empty c-oriented polygons P and P' have an empty intersection if and only if there exists $l \in \{1, ..., c\}$ such that $strip_l(P') \cap strip_l(P) = \emptyset$.*

The key point is to prove that disjoint non-empty c-oriented polygons have at least a pair of disjoint strips.
Proof. Assume that $P \cap P' = \emptyset$ and that the origin O is within P. Let λP be the smallest homothet of P with scaling factor λ such that $\lambda P \cap P' \neq \emptyset$. Clearly for $1 \leq \lambda' < \lambda$, we have $P \subseteq \lambda' P \subset \lambda P$. Since $P \cap P' = \emptyset$, we have $\lambda > 1$. λP and P' have a either an edge-vertex (Figure 3, left) or vertex-vertex contact (Figure 3, right). In the first case, let h be the line passing through that edge then h separates P from P' and thus the corresponding strips of P and P' of same orientation are disjoint. In the second case, there is always one line passing through an edge incident at v that separate P from P'. **Q.E.D.**

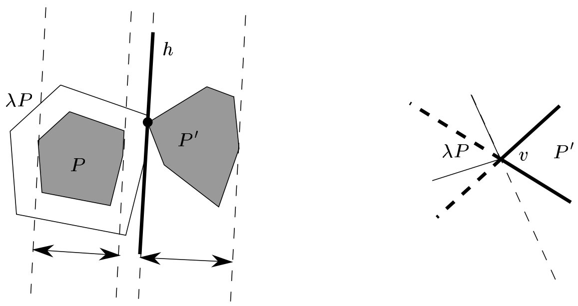

Figure 3: $P \cap P' = \emptyset$ implies two disjoint strips. **Left**: vertex-edge contact. **Right**: vertex-vertex contact.

Here, we would like to point out the main differences between boxes (whenever $b = d$) and c-oriented polytopes. In the case of boxes, a set $\mathcal{S}$ of boxes has a non-empty intersection ($\cap \mathcal{S} \neq \emptyset$) iff $\forall i \in \{1, ..., d\}$ $strip_i(\mathcal{S}) \neq \emptyset$. In terms of Gallai numbers, we have $\mathfrak{G}(2, 1, \mathcal{B}) = \mathfrak{G}(2, d) = 1$ for the class $\mathcal{B}$ of d-dimensional boxes. However, we have $\mathfrak{G}(2, 1, \mathcal{H}) = O(1)$ for the class $\mathcal{H}$ of translates or homothets. Theorem 1, which can be seen as a refinement of Farkas' lemma on the plane, is not anymore true for c-oriented polytopes (because of the possibly non-degenerate edge-edge contact, see Figure 4).

3 An output-sensitive algorithm.

In [Nie96], an optimal $\Theta(n \log \phi(\mathcal{S}))$-time algorithm with linear storage is given for computing an optimal covering point set of size $\Delta(\mathcal{S}) = \tau(\mathcal{S}) = \phi(\mathcal{S})$ for a family of n intervals. The algorithm for piercing c-oriented polygons proceeds through two stages as described below:

Partitioning. Create a partition $\mathcal{Q}$ of $\mathcal{S}$ such that any set $Q \in \mathcal{Q}$ has the following weak *pairwise intersection property*:

$$\forall P,\ P',\ P'' \in Q,\ (P \cap P' \neq \emptyset) \text{ and } (P' \cap P'' \neq \emptyset) \Rightarrow (P \cap P'' \neq \emptyset).$$

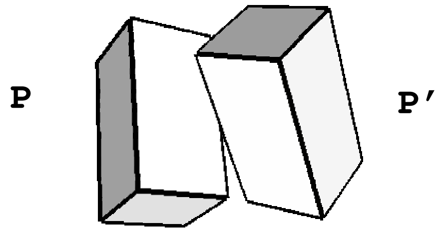

Figure 4: Two non-empty 6-stripped polytopes of $\mathbb{E}^3$ which do not admit a pair of non-empty strips.

Viewing Q as its intersection graph, it means that Q has a trivial disjoint clique covering (*cluster covering*) such that for any clique $K \in Q$ (cluster), we have: $\forall i \in \{1, ..., c\}, strip_i(K) \neq \emptyset$.

Piercing. Pierce c-oriented polygons whose corresponding induced families of strips intersect pairwise. That is, pierce families of pairwise intersecting c-oriented polygons (clusters, see Theorem 1).

3.1 Getting a partition into pairwise intersecting polytopes.

We choose an arbitrary direction, say h_1, and consider the induced family $\mathcal{S}_1$ of n intervals (perpendicular projection):

$$\mathcal{S}_1 = \{strip_1(S) | S \in \mathcal{S}\} = \Big\{[\alpha_{i,1}, \beta_{i,1}] | i \in \{1, ..., n\}\Big\}.$$

Then, we compute an optimal covering point set of $\mathcal{S}_1$ using the algorithm in [Nie96]. It is clear that $\Delta(\mathcal{S}_1) = \phi(\mathcal{S}_1) \leq \phi(\mathcal{S})$. Let $x_1 \leq ... \leq x_{l_1}$ be the l_1 ordered real points piercing $\mathcal{S}_1$ with $l_1 \leq \phi(\mathcal{S})$. We consider the median value x_m of $X = \{x_1, ..., x_{l_1}\}$, i.e., the $\lfloor \frac{l_1}{2} \rfloor$-th smallest value of X. We get an induced partition of $\mathcal{S}$ into three subsets $\mathcal{S}_l = \{\ P_i | \beta_{i,1} < x_m\}$, $\mathcal{S}' = \{\ P_i | \alpha_{i,1} \leq x_m \leq \beta_{i,1}\}$ and $\mathcal{S}_r = \{\ P_i | \alpha_{i,1} > x_m\}$. Observe also that for any pair of polygons $P_l \in \mathcal{S}_l$ and $P_r \in \mathcal{S}_r$ then $P_l \cap P_r = \emptyset$ since $strip_1(P_l) \cap strip_1(P_r) = \emptyset$. We first recurse on $\mathcal{S}'$ until at some stage $\phi(\mathcal{S}'_1) \leq 1$, i.e., $\mathcal{S}_l = \mathcal{S}_r = \emptyset$. Then, we perform a recursion on the $c-1$ remaining distinct orientations.

Lemma 1 *Partition $\mathcal{Q}$ has at most $\phi(\mathcal{S}) \log_2^{c-1}(\phi(\mathcal{S}) + 1)$ cluster coverings.*

Proof. Let $f(\phi, c)$ denotes the minimum number of covering points required for covering a set $\mathcal{S}$ of c-oriented polygons so that $\phi(\mathcal{S}) = \phi$. More precisely,

$$f(\phi, c) = \max_{\mathcal{S} \in \mathcal{P}_c} \{\tau(\mathcal{S}) \mid \phi(\mathcal{S}) = \phi \text{ and } \mathcal{S} \text{ is } c\text{-oriented}\},$$

where $\mathcal{P}_c$ is the class of c-oriented polygons of $\mathbb{E}^2$ $(c \geq 3)$.

From our algorithm, we get $f(\phi, c) \leq \min_{0 \leq k \leq \phi-2}\{f(k, c) + f(\phi-k-1, c) + f(\phi, c-1)\}$, for $\phi \geq 2$ and $c \geq 2$. If $c = 1$, we have $f(\phi, 1) = \phi$ [Nie96]. Note that $f(1, c) = \mathfrak{G}(2, c)$ for $c \geq 1$. Let $h(\phi, c) = \phi \log_2^{c-1}(\phi + 1)$. Then, we claim that $f(\phi, c) \leq \mathfrak{G}(2, c) h(\phi, c)$. We do an induction on the lexicographically ordered vector (c, ϕ). For $c = 1$, we have $f(\phi, 1) = \phi = \mathfrak{G}(2, 1)\phi \log^0(\phi + 1)$. If $\phi = 1$ then this means that we cannot find two disjoint c-oriented polytopes and therefore $f(1,c)=\mathfrak{G}(2,c)\leq \mathfrak{G}(2,c)(\phi\log_2^{c-1}2)$. Otherwise ($c>1$ and $\phi>1$), we have:

$$f(\phi,c)\leq f(\frac{\phi-1}{2},c)+f(\frac{\phi-1}{2},c)+f(\phi,c-1).$$

Therefore, we get:

$$h(\phi,c)\leq 2\frac{\phi}{2}\log_2^{c-1}\frac{\phi+1}{2}+\phi\log_2^{c-2}(\phi+1)\leq \phi\log_2^{c-2}(\phi+1)(\log_2\frac{\phi+1}{2}+1),$$

$$h(\phi,c)\leq \phi\log_2^{c-1}(\phi+1).$$

Q.E.D.

A polytope P is said *α-fat* if the ratio of the widths of a smallest hypersquare $H^+(P)$ containing P and a greatest hypersquares $H^-(P)$ included in P is bounded by α. A set $\mathcal{S}$ is said to have the *bounded aspect ratio property* (or alternatively *β-sized*) if there exists a constant β so that $\frac{H^+(\mathcal{S})}{H^-(\mathcal{S})}\leq\beta$, where $H^+(\mathcal{S})=\max\{H^+(P)\mid P\in\mathcal{S}\}$ and $H^-(\mathcal{S})=\min\{H^-(P)\mid P\in\mathcal{S}\}$.

Corollary 1 *For a set $\mathcal{S}$ of β-sized c-oriented polygons, partition $\mathcal{Q}$ has at most $\phi(\mathcal{S})O(\beta)^{c-1}$ cluster coverings.*

Proof. The key idea is to analyze the left to right sequence of generated subsets $Q_1,...,Q_{l_1}$ induced when separating the objects using some direction c_1. $\mathcal{S}_1$, the set of intervals obtained by considering for each polygon P_i its projection $strip_1(P_i)$, has the bounded aspect ratio property since objects of $\mathcal{S}$ are more or less the same size (i.e., $O(1)$-sized). Denote by $r^+=\max_{O\in\mathcal{S}}\{H^+(O)\}$ and $r^-=\min_{O\in\mathcal{S}}\{H^-(O)\}$. Our goal is to prove that we can create at most $(2\lceil\beta\rceil+1)\phi(\mathcal{S})$ sub-partitions after separating $\mathcal{S}$ with direction c_1. Note that if $\mathcal{S}$ is β-sized then so is $\mathcal{S}_1$. It suffices to notice that $Q_i\cap Q_j=\emptyset$ if $j\geq i+2\lceil\beta\rceil+1$, where $Q_i\in\mathcal{Q}_i$ and $Q_j\in\mathcal{Q}_j$. This comes from the fact that $x_{i+2}-x_i>\min_{O\in\mathcal{S}}\{H^-(O)\}$. Thus, $x_{i+2\lceil\beta\rceil+1}-x_i>\beta r^-\geq r^+$. This enables us to prove that $l\leq\phi(\mathcal{S})O(\beta)^{c-1}$ whenever $\mathcal{S}$ is a family of objects which have more or less the same size. **Q.E.D.**

Lemma 2 *The algorithm runs in $O(nc\log\phi(\mathcal{S}))$-time using linear storage $O(nc)$.*

Proof. Let $T(n,c)$ denote the running time of the algorithm. We have the recursive time-complexity system:

$$T(n,c)=\begin{cases}O(n\log\phi) & \text{if } c=1,\\ O(n\log\phi)+\sum_i T(n_i,c-1) & \text{otherwise.}\end{cases}$$

with $\sum_i n_i=n$ (that is $\phi(\mathcal{S}_1)$ terms in the sum, with $\phi(\mathcal{S}_1)\leq\phi(\mathcal{S})$). Clearly, a simple induction on the lexicographic ordered vector (c,n) proves that $T(n,c)=O(nc\log\phi(\mathcal{S}))$ [Nie96]. Trivially, the algorithm uses linear storage $O(N)=O(nc)$. Note that this algorithm is robust to numerical tests since it performs only comparisons of algebraic degree 1. **Q.E.D.**

3.2 Piercing pairwise intersecting c-oriented polytopes.

We consider below the problem in arbitrary dimension d.

3.2.1 The case of fat (c-coriented) polytopes.

Let $\mathcal{S} = \{P_1, ..., P_n\}$ be a set of n pairwise intersecting (c-oriented) polytopes. Since $\mathcal{S}$ is α-fat, we are able to pierce $\mathcal{S}$ with $O(\alpha)^d$ points as follows: choose a polytope, say P_1, which has a smallest enclosing hypersquare H_1^+. Consider the 'parallel' Minkowski hypersquare μH_1^+, where μ is a constant depending on α. We draw regularly $O(\alpha)^d$ points inside μH_1^+ using a Fredman's sampling process. The key property is that all other fat objects P_j's intersecting P_1 contains at least one of these points (see Figure 5). Hence, we get the following lemma:

Lemma 3 *Let $\mathcal{S}$ be a set of pairwise intersecting c-oriented α-fat polytopes then $\tau(\mathcal{S}) = O(\alpha)^d \phi(\mathcal{S}) = \mathfrak{G}(2, c)$. Moreover, there exists a simple algorithm with running time $t(n, c) = O_d(nc)$ for piercing $\mathcal{S}$ with $O(\alpha)^d$ points.*

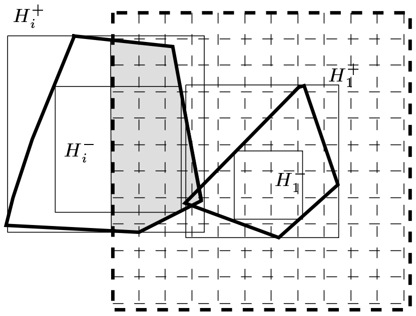

Figure 5: Sampling $O(1)$ covering points.

Lemma 4 *Let $\mathcal{S}$ be a collection of n α-fat c-oriented polygons of $\mathbb{E}^2$. Then, we get:*

$$\tau(\mathcal{S}) \leq O(\alpha)^2 \phi(\mathcal{S}) \log^{c-1}(\phi(\mathcal{S}) + 1).$$

In [EKNS97], we showed that $\tau(\mathcal{S}) = \phi(\mathcal{S}) O(\alpha)^d$ and give various output-sensitive precision-sensitive algorithms for arbitrary fat, not necessarily c-oriented, polytopes and balls of $\mathbb{E}^3$.

3.2.2 The case of homothets.

Let $\mathcal{S} = \{P_1, ..., P_n\}$ be a set of n pairwise intersecting homothets of a c-oriented polytope. Grünbaum showed that $\mathfrak{G}(2, c) \leq d^d = O_d(1)$ [Grü59].

Theorem 2 *Let $\mathcal{S}$ be a collection of c-oriented homothets in the Euclidean plane $\mathbb{E}^2$. Then,*

$$\tau(\mathcal{S}) \leq 4\phi(\mathcal{S}) \log^{c-1}(\phi(\mathcal{S}) + 1).$$

In the case of translates, it has been shown that $\mathfrak{G}(2, c) \leq (3d^{\frac{3}{2}})^d$ [Grü59]. For centrally symmetric translates, the bound can be lowered to $(\sqrt{2}d)^d$. Interestingly, there is a linear-time algorithm (that is, $t(n, c) = \Theta_d(nc)$) for piercing pairwise intersecting c-oriented homothets of a polytope with at most d^d points.

3.2.3 The general case.

When $\mathcal{S}$ is neither α-fat nor a collection of homothets of a given c-oriented polytope then the problem becomes far more difficult to tackle. The main open question to answer is whether $\mathfrak{G}(2, c)$ is bounded or not (assuming c as a constant parameter) for arbitrary infinitely many c-oriented polytopes. We noticed that $\mathfrak{G}(2, c) \geq \lceil \frac{c-d+1}{2} \rceil$ (see Figure 1).

4 Computing an independent set.

Given a graph $G=(V,E)$, an independent set I of G is a set of nodes $I \subseteq V$ such that there is no edge of E between two nodes of I. The intersection graph of a geometric set of objects is defined as follows: to each object we associate a corresponding node and we set an edge between two nodes if and only if the corresponding objects intersect. Therefore, a maximal independent set of the intersection graph corresponds to a maximal set of pairwise non-intersecting objects. Although finding a maximal independent set on general abstract graph G has been shown to be $\Omega(n^{\frac{1}{4}})$-hard to approximate [BS94] in polynomial time, we describe below a $\frac{1}{\log_2^{c-1}(\phi+1)}$ approximation heuristic based on the same partitionning scheme.

Let $\Delta'(\mathcal{S})$ be the size of an independent set found by some heuristic. Then, we have:

$$\Delta'(\mathcal{S}) \leq \phi(\mathcal{S}) \leq \tau(\mathcal{S}) \leq \Delta(\mathcal{S}).$$

On the other hand, we previously showed that $\Delta(\mathcal{S}) \leq \mathfrak{G}(2,c)\phi(\mathcal{S})\log_2^{c-1}(\phi(\mathcal{S})+1)$. Taking a closer look at the complexity analysis, we are able to derive *far better* bounds;

Indeed, we may refine the analysis of section 3.1 as:

$$\Delta(\mathcal{S}) \leq \max_i\{b_i\}\mathfrak{G}(2,c)\log_2^{c-1}(\phi(\mathcal{S})+1),$$

where the b_i's are the maximal number of pairwise disjoint intervals of subsets of $\mathcal{S}$ induced by the partitionnning scheme. We bounded the piercing number of each cluster (a clique) by a Gallai number $\mathfrak{G}(2,c)$. Considering the independent set problem, we can pick arbitrary any object inside a given cluster. Therefore, a simple heuristic consists in giving an independent set of objects of $\mathcal{S}$ from the subsets of $\mathcal{S}$ whose independent set of projected intervals is of maximal size. we get

$$\Delta(\mathcal{S}) = \max\{b_i\} \geq \frac{\phi(\mathcal{S})}{\log_2^{c-1}(\phi(\mathcal{S})+1)}.$$

A better heuristic, in practice, can be modeled as a binary tree where the root is set $\mathcal{S}$ and children correspond to recursive calls on subsets generated by partitionning the objects of $\mathcal{S}$ according to some orientation. Each internal node $\mathcal{S}$ has at most two children labeled with sets $\mathcal{S}_l$ and $\mathcal{S}_r$. We compute recursively in a bottom-to-top fashion an independent set of $\mathcal{S}$ by choosing either, at some node $\mathcal{S}$, the maximal independent set of $\mathcal{S}'$ or the union of the independent sets computed so far of sets $\mathcal{S}_l$ and $\mathcal{S}_r$. Let $\Delta'(\mathcal{S})$ denote the size of such an independent set found by this heuristic. Then, we can refine the analysis in section 3.1 as follows:

$$\phi(\mathcal{S}) \leq \tau(\mathcal{S}) \leq \mathfrak{G}(2,c)\Delta'(\mathcal{S})\log_2^{c-1}(\phi(\mathcal{S})+1).$$

In term of clusters, this means that $\phi(\mathcal{S}) \leq \Delta'(\mathcal{S})\log_2^{c-1}(\phi(\mathcal{S})+1)$, from which we get

$$\Delta'(\mathcal{S}) \geq \frac{\phi(\mathcal{S})}{\log_2^{c-1}(\phi(\mathcal{S})+1)}.$$

A program written in Java[2] demonstrates such a heuristic on planar boxes.

Finding maximal independet sets naturally arise in practice as for example in map labeling where one wants to maximize the number of labels so that they are pairwise non-intersecting (See [DMM$^+$97, PZC97, WW97, AvKS97]). Agarwal et al. [AvKS97] gave a geometric polynomial-time approximation scheme for computing a maximal independent set of boxes on the plane in $O(n\log n+n^{2k-1})$ time having a performance ratio of $1+\frac{1}{k}$ for any integer $k \geq 1$ using the shifting lemma combined with dynamic programming.

[2]See the applet located at`http://www.csl.sony.co.jp/person/nielsen/PCD/pierce.html`

5 Concluding remarks.

Our precision-sensitive heuristic can easily be parallelized on the PRAM model of computation following the work of Akl and Lyons [AL93]. We get an $O(N^{1-\epsilon} \log \Delta(\mathcal{S}) + T(N, N^\epsilon))$-time algorithm using $O(N^\epsilon)$ processors, where $T(N,P)$ is the time required to pierce n c-oriented pairwise intersecting polygons on P processors, for any constant $\epsilon > 0$. Below, we give a table stating the main results obtained on the geometric SCP.

Class	$\tau \leq f(\phi)$?	Performance ratio $\frac{\Delta(\mathcal{S})}{\tau(\mathcal{S})}$	Reference
Boxes	$\tau \leq \phi(o(1) + \log \phi)^{d-1}$	$(o(1) + \log \phi)^{d-1}$	[Kár91, Nie96]
Fat objects	$\tau = O_{d,\alpha}(\phi)$	$O_{d,\alpha}(1)$	[EKNS97]
c-Oriented Polygons	$\tau \leq \mathfrak{G}(2,c)\phi \log^{c-1}(\phi+1)$	$\mathfrak{G}(2,c) \log^{c-1}(\phi+1)$	This paper
Bounded VC-dim.	None	$O(\log \tau), O(1)$	[BG94]
Abstract sets	None	$\log n - \log\log n + 0.78$	[Chv79, Sla96]

Finally, we list some known Gallai numbers:

Type	$\mathfrak{G}(2,c), c \geq d+1$
Homothets	d^d
Translates	$3d^{\frac{3}{2}}$
Centrally symmetric translates	$(\sqrt{2}d)^d$
Fat objects	$O(\alpha)^d$

References

[AKK95] S. Arora, D. Karger, and M. Karpinski. Polynomial-time approximation schemes for dense instances of NP-hard problems. In *Symp. on Theory of Comp.*, pages 284–293, 1995.

[AL93] S. G. Akl and K. A. Lyons. *Parallel Computational Geometry.* Prentice Hall, 1993.

[AvKS97] Pankaj K. Agarwal, Mark van Kreveld, and Subhash Suri. Label placement by maximum independent set in rectangles. In *Proc. 9th Canad. Conf. Comput. Geom.*, pages 233–238, 1997.

[BG94] H. Brönnimann and M. T. Goodrich. Almost optimal set covers in finite VC-dimension. In *Proc. 10th Annu. ACM Sympos. Comput. Geom.*, pages 293–302, 1994.

[BS94] M. Bellare and M. Sudan. Improved non-approximability results. In *Proc. 26th Annu. ACM Sympos. Theory Comput.*, pages 184–193, 1994.

[Chv79] V. Chvátal. A greedy heuristic for the set-covering problem. *Math. Oper. Res.*, 4:233–235, 1979.

[DF97] Rong-chii Duh and Martin Fürer. Approximation of k-set cover by semi-local optimization. In *Proc. 29th Annu. ACM Sympos. Theory Comput.*, pages 256–264, 1997.

[dFK93] D. G. Fon der Flass and A. V. Kostochka. Covering boxes by points. *Discr. Math.*, 120:269–275, 1993.

[DMM+97] S. Doddi, M. Marathe, A. Mirzaian, B. Moret, and B. Zhu. Map labeling and its generalization. In *Proc. 8th ACM-SIAM Sympos. Discrete Algorithms*, pages 148–157, 1997.

[DSW94] G. L. Ding, P. Seymour, and P. Winkler. Bounding the vertex cover number of a hypergraph. *Combinatorica*, 14:23–34, 1994.

[EKNS97] Alon Efrat, Matthew J. Katz, Franck Nielsen, and Micha Sharir. Dynamic data structures for fat objects and their applications. In *Proc. Work. on Alg. and Data Struct. 97, LNCS 1272*, pages 297 – 306. Halifax, Nova Scotia, Canada, 1997.

[Fei96] U. Feige. A Threshold of $\log n$ for Approximating Set Cover. In *Proc. of the 28-th ACM Symposium on Theory of Computing*, 1996.

[FPT81] R. J. Fowler, M. S. Paterson, and S. L. Tanimoto. Optimal packing and covering in the plane are NP-complete. *Inform. Process. Lett.*, 12(3):133–137, 1981.

[Grü59] B. Grünbaum. On intersections of similar sets. *Portugal Math.*, 18:155–164, 1959.

[HM84] D. S. Hochbaum and W. Maass. Approximations schemes for covering and packing problems in image processing and VLSI. *J. ACM*, 31:130–136, 1984.

[HM87] D. S. Hochbaum and W. Maass. Fast approximation algorithms for a nonconvex covering problem. *J. Algorithms*, 8:305–323, 1987.

[Hoc82] D. S. Hochbaum. Approximation Algorithms of the Set Covering and Vertex Cover Problems. *SIAM J. COMPUT.*, 11(3):555 – 556, August 1982.

[Kár91] GY. Károlyi. On point covers of parallel rectangles. *Periodica Mathematica Hungarica*, 23(2):105–107, 1991.

[KN96] M. J. Katz and F. Nielsen. On piercing sets of objects. In *Proc. 12th Annu. ACM Sympos. Comput. Geom.*, pages 113–121, 1996.

[KT95] G. Károlyi and G. Tardos. On point covers of multiples intervals and axis-parallel rectangles. Technical report, DIMACS TR 95-45, 1995.

[KZ96] M. Karpinski and A. Zelikovsky. Approximating dense cases of covering problems. Technical report, DIMACS TR 96-59, 1996.

[Nie96] F. Nielsen. Fast stabbing of boxes in high dimensions. In *Proc. 8th Canad. Conf. Comput. Geom.*, pages 87–92, 1996.

[PS85] F. P. Preparata and M. I. Shamos. *Computational Geometry: An Introduction*. Springer-Verlag, New York, NY, 1985.

[PZC97] C. Poon, B. Zhu, and F. Chin. A polynomial time solution for labeling a rectilinear map. In *Proc. 13th Annu. ACM Sympos. Comput. Geom.*, pages 451–453, 1997.

[Sla96] P. Slavík. A Tight Analysis of the Greedy Algorithm for Set Cover. In *Proc. of the* 28-*th ACM Symposium on Theory of Computing*, 1996.

[WW97] F. Wagner and A. Wolff. A practical map labeling algorithm. *Comput. Geom. Theory Appl.*, 7:387–404, 1997.