Online Hitting Sets In A Geometric Setting Via Vertex Ranking

Guy Even* Shakhar Smorodinsky†

Abstract

We consider the problem of hitting sets online. The hypergraph (i.e., range-space consisting of points and ranges) is known in advance. However, the ranges to be stabbed are input one-by-one in an online fashion. The online algorithm must stab each range upon arrival. An online algorithm may add points to the hitting set but may not remove already chosen points. The goal is to use the smallest number of points. The best known competitive ratio for online hitting sets by Alon et al. [1] is $O(\log n \cdot \log m)$ for general hypergraphs, where n and m denote the number of points and the number of ranges, respectively. We consider three special classes of hypergraphs.

The first setting consists of subsets of nodes of a given graph that induce connected subgraphs. We show how vertex ranking can be employed to design a simple online algorithm, the competitive ratio of which equals the number of colors used by the vertex ranking. When the underlying graph is a planar graph (e.g., a Delaunay triangulation) with n vertices, we obtain an optimal $O(\sqrt{n})$-competitive ratio. We remark that the analysis of the competitive ratio of the algorithm of [1] only proves an $O(n)$-competitive ratio for this case.

In the second setting, we consider subsets of a given set of n points in the Euclidean plane that are induced by half-planes. We apply the first setting to obtain an $O(\log n)$-competitive ratio. We also prove an $\Omega(\log n)$ lower bound for the competitive ratio in this setting.

In the third setting, we consider subsets of a given set of n points in the plane induced by unit discs. Since the number of subsets in this setting is $O(n^2)$, the competitive ratio obtained by Alon et al. is $O(\log^2 n)$. We introduce an algorithm with $O(\log n)$-competitive ratio. We also show that any online algorithm for this problem has a competitive ratio of $\Omega(\log n)$, and hence our algorithm is optimal.

1 Introduction

In the minimum hitting set problem, we are given a hypergraph (X, R), where X is the ground set of points and R is a set of hyperedges. The goal is to find a finite set $S \subseteq X$ such that every hyperedge is stabbed by S, namely, every hyperedge has a nonempty intersection with S.

The minimum hitting set problem is a classical NP-hard problem [15], and remains hard even for geometrically induced hypergraphs (see [12] for several references). A sharp logarithmic threshold for hardness of approximation was proved by Feige [11]. On the other hand, the greedy algorithm achieves a logarithmic approximation ratio [14, 7]. Better approximation ratios have been obtained for several geometrically induced hypergraphs using specific properties of the induced hypergraphs [12, 17, 2]. Other improved approximation ratios are obtained using the theory of VC-dimension and ε-nets [4, 10, 8]. Much less is known about online versions of the hitting set problem.

In this paper, we consider an online setting in which the set of points X is given in the beginning, and the ranges are introduced one by one. Upon arrival of a new range, the online algorithm may add points (from X) to the hitting set so that the hitting set also stabs the new range. However, the online algorithm may not remove points from the hitting set. We use the competitive ratio, a classical measure for the performance of online algorithms [19, 3], to analyze the performance of online algorithms.

Alon et al. [1] considered the online set-cover problem for arbitrary hypergraphs. In their setting, the ranges are known in advance, and the points are introduced one by one. Upon arrival of an uncovered point, the online algorithm must choose a range that covers the point. Hence, by replacing the roles of ranges and points, their setting is equivalent to our setting. The online set cover algorithm presented by Alon et al. [1] achieves a competitive ratio of $O(\log n \log m)$ where n and m are the number of points and the number of hyperedges respectively. Note that if $m \geq 2^{n/\log n}$, the analysis of the online algorithm only guarantees that the competitive ratio is $O(n)$; a trivial bound if one range is chosen for each point. On the other hand, in many geometric settings the underlying hypergraph has a bounded VC-dimension which implies that the number of hyperedges is polynomial in the number of points (see, e.g., [18]). If m is polynomial in n, their algorithm achieves a competitive ratio of $O(\log^2 n)$. Since there is no matching lower bound of $\Omega(\log^2 n)$ for the competitive ratio, it might be the case that there is an algorithm with an $O(\log n)$-competitive

*School of Electrical Engineering, Tel-Aviv Univ., Tel-Aviv 69978, Israel. guy@eng.tau.ac.il

†Mathematics Department, Ben-Gurion Univ., Be'er Sheva 84105, Israel. shakhar@math.bgu.ac.il

ratio for such hypergraphs.

We present online hitting set algorithms for special classes of hypergraphs, some with bounded VC-dimension and some with linear VC-dimension. These algorithms are based on a novel relation between vertex ranking [16] and hitting sets.

The first class of hypergraphs that we consider, is induced by connected components of a given graph. The input of the algorithm is a graph $G = (V, E)$, and the adversary chooses subsets $V' \subseteq V$ such that the induced subgraph $G[V']$ is connected. An application for such a setting is the placement of servers in virtual private networks (VPNs). Each VPN is a subset of vertices that induce a connected subgraph, and requests for VPNs arrive online. The algorithm selects a location for each VPN, and the goal is to select as few servers as possible.

For the case of hypergraphs induced by connected components of graph, we show that one can use vertex ranking to design an efficient online algorithm where the competitive ratio equals the number of colors used by a vertex ranking. In particular, for forests on n vertices, our algorithm achieves an optimal $O(\log n)$-competitive ratio, and for planar graphs our algorithm achieves $O(\sqrt{n})$-competitive ratio. This class is of particular interest since the VC-dimension of such a hypergraph is not bounded. For example, if G is a star (i.e., a vertex v with $n-1$ neighbors), the number of subsets of vertices that induce a connected graph is 2^{n-1}. However, the star has a vertex ranking that uses just two colors, hence, the competitive ratio of our algorithm in this case is 2. This is easily seen to be the best competitive ratio that can be achieved. If G is a planar graph, then G admits a vertex ranking that uses $O(\sqrt{n})$ colors, and the competitive ratio of our algorithm is $O(\sqrt{n})$. This is an improvement over the analysis of the algorithm of Alon et al. [1] which only proves a competitive ratio of $O(n)$. Thus, our algorithm is useful even in hypergraphs whose VC-dimension is unbounded.

Two more classes of hypergraphs are obtained geometrically as follows. In both settings we are given a set X of n points in the plane. In one hypergraph, the hyperedges are intersections of X with half planes. In the other hypergraph, the hyperedges are intersections of X with unit discs. Our main result is an online algorithm for the hitting set problem for points in the plane and unit discs (or half-planes) with an optimal competitive ratio of $O(\log n)$. The competitive ratio of this algorithm improves the $O(\log^2 n)$-competitive ratio of Alon et al. by a logarithmic factor. An application for points and unit discs is the selection of access points or base stations in a wireless network. The points model base stations and the disc centers model clients. The reception range of each client is a disc, and the algorithm has to select a base station that serves a new uncovered client. The goal is to select as few base stations as possible.

2 Preliminaries

Let (X, R) denote a hypergraph, where R is a set of nonempty subsets of the ground set X. Members in X are referred to as *points*, and members in R are referred to as *ranges* (or *hyperedges*). A subset $S \subseteq X$ *stabs* a range r if $S \cap r \neq \emptyset$. A *hitting set* is a subset $S \subseteq X$ that stabs every range in R. In the minimum hitting set problem, the goal is to find a hitting set with the smallest cardinality.

In this paper, we consider the following online setting. The adversary introduces a sequence $\sigma \triangleq \{r_i\}_{i=1}^{s}$ of ranges. Let σ_i denote the prefix $\{r_1, \dots, r_i\}$. The online algorithm computes a chain of hitting sets $C_1 \subseteq C_2 \subseteq \cdots$ such that C_i is a hitting set with respect to the ranges in σ_i.

The competitive ratio of the algorithm is defined as follows. Let $\text{OPT}(\sigma) \subseteq X$ denote a minimum cardinality hitting set for the ranges in σ. Let $\text{ALG}(\sigma) \subseteq X$ denote the hitting set computed by an online algorithm ALG when the input sequence is σ. Note that the sequence of minimum hitting sets $\{\text{OPT}(\sigma_i)\}_i$ is not necessarily a chain of inclusions. The *competitive ratio* of an online hitting set algorithm ALG is defined as the supremum, over all sequences σ of ranges, of the ratio $|\text{ALG}(\sigma)|/|\text{OPT}(\sigma)|$.

3 Summary of Our Results

Connected Subgraphs. We consider the following setting of a hypergraph induced by connected subgraphs of a given graph. Formally, let $G = (V, E)$ be a graph. Let $H = (V, R)$ denote the hypergraph over the same set of vertices V. A subset $r \subseteq V$ is a hyperedge in R if and only if the subgraph $G[r]$ induced by r is connected.

We need the notion of a vertex ranking of a graph[16]. A *vertex ranking* is a function $c : V \to \mathbb{N}$ that satisfies the following property: For any pair of vertices $u, v \in V$, if $c(u) = c(v)$, then, for any simple path P in G connecting u and v, there is a vertex w in P such that $c(w) > c(u)$. Notice that, in particular, a vertex ranking of a graph G is also a proper coloring of G since adjacent vertices must get distinct colors. For a subset $X' \subset X$, let $c_{\max}(X') \triangleq \max\{c(v) \mid v \in X'\}$. It is easy to see that if c is a vertex ranking, then $|\{v \in r \mid c(v) = c_{\max}(r)\}| = 1$, for every $r \in R$. Thus, let $v_{\max}(r)$ denote the (unique) vertex v in r such that $c(v) = c_{\max}(r)$.
Our first result is an online hitting set algorithm for connected subgraphs.

Theorem 1 *Let $c : V \to \mathbb{N}$ denote a vertex ranking of a graph $G = (V, E)$. Then there exists an online* *hitting set algorithm for the connected subgraphs of G with a competitive ratio of $c_{\max}(V)$.*

By [16], planar graphs admit vertex rankings with $c_{\max}(V) = O(\sqrt{|V|})$. Therefore, Theorem 1 implies that the competitive ratio of our algorithm for connected subgraphs of planar graphs is $O(\sqrt{n})$. We also prove that this competitive is optimal.

Theorem 2 *The competitive ratio of every online hitting set algorithm for connected subgraphs of planar graphs is $\Omega(\sqrt{n})$.*

Points and Half-Planes. We prove the following results for hypergraphs in which the ground set X is a finite set of n points in $\mathbb{R}^2$ and the ranges are all subsets of X that can be cut off by a half-plane. Namely, each range r is induced by a line L_r such that r is the set of points of X in the half-plane below (respectively, above) the line L_r.

Theorem 3 *The competitive ratio of every online hitting set algorithm for points and half-planes is $\Omega(\log n)$.*

Theorem 4 *There exists an online hitting set algorithm for points and half-planes that achieves a competitive ratio of $O(\log n)$.*

Points and Congruent Discs. We prove the following results for hypergraphs in which the ground set X is a finite subset of n points in $\mathbb{R}^2$ and the ranges are intersections of X with unit discs. Namely, a unit disc d induces a range $r = r(d)$ defined by $r = d \cap X$.

Theorem 5 *The competitive ratio of every online hitting set algorithm for points and unit discs is $\Omega(\log n)$.*

Theorem 6 *There exists an online hitting set algorithm for points and unit discs that achieves a competitive ratio of $O(\log n)$.*

4 Vertices and Connected Subgraphs

4.1 Algorithm Description

A listing of Algorithm HS appears as Algorithm 1. The algorithm is input a graph $G = (V, E)$ and a vertex ranking $c : V \to \mathbb{N}$. The sequence $\sigma = \{r_i\}_i$ of subsets of vertices that induce connected subgraphs is input online.

Algorithm 1 HS(G, c) - an online hitting set for connected subgraphs, given a vertex ranking c.

Require: $G = (V, E)$ is a graph and $c : V \to \mathbb{N}$ is a vertex ranking.

```
C_0 ← ∅
for i = 1 to ∞ do {arrival of a range r_i}
    if r_i is not stabbed by C_{i−1} then
        C_i ← C_{i−1} ∪ {v_max(r_i)} {add the vertex with the max color in r_i}
    else
        C_i ← C_{i−1}
    end if
end for
```

4.2 Analysis of The Competitive Ratio

Definition 1 *For a color a, let $\sigma(a)$ denote the subsequence of σ that consists of ranges that satisfy: (i) r_i is not stabbed by C_{i-1}, and (ii) $c_{\max}(r_i) = a$.*

The following lemma implies a lower bound on the (offline) minimum hitting set of the ranges in $\sigma(a)$.

Lemma 7 *If $r_i, r_j \in \sigma(a)$, then the subgraph $G[r_i \cup r_j]$ induced by $r_i \cup r_j$ is not connected. Hence, the ranges in $\sigma(a)$ are pairwise disjoint.*

Proof. Clearly, $c_{\max}(r_i \cup r_j) = \max\{c_{\max}(r_i), c_{\max}(r_j)\} = a$. Assume that $r_i \cup r_j$ induces a connected subgraph. Since c is a vertex ranking, we conclude that $r_i \cup r_j$ contains exactly one vertex colored a. This implies that $v_{\max}(r_i) = v_{\max}(r_j)$. If $j > i$, then the range r_j is stabbed by C_{j-1} since it is stabbed by C_i, a contradiction. ☐

Proof. [Proof of Theorem 1] Algorithm HS satisfies $|\text{HS}(\sigma)| = \sum_{a \in \mathbb{N}} |\sigma(a)|$. But $\sum_{a \in \mathbb{N}} |\sigma(a)| \le c_{\max}(V) \cdot \max_{a \in \mathbb{N}} |\sigma(a)|$. By Lemma 7, each range in $\sigma(a)$ must be stabbed by a distinct vertex, thus $|\text{OPT}(\sigma)| \ge \max_{a \in \mathbb{N}} |\sigma(a)|$, and the theorem follows. ☐

Corollary 8 *Let $G = (V, E)$ be a planar graph and let H be the hypergraph consisting of V together with all subsets of vertices inducing connected subgraphs. Then there is an online hitting set for H with competitive ratio of $O(\sqrt{n})$.*

Proof. The proof uses the fact that there exists a vertex ranking for G with a total of $O(\sqrt{n})$ colors [16]. Combining such a vertex ranking with Algorithm HS and using Theorem 1 completes the proof. ☐

Remark 1 *The above corollary applies to arbitrary graphs with small balanced separators. Let $G = (V, E)$ be a graph such that every subgraph with m vertices has a balanced separator with $O(m^\alpha)$ vertices, for some fixed $0 < \alpha \le 1$. Then there is an* *online hitting set algorithm for the connected subgraphs with competitive ratio of $O(n^\alpha)$. The proof uses all of the above mentioned ingredients, replacing $\sqrt{n}$ with n^α.*

For the special case of trees, since there is always a balanced separator with 1 vertex, it is easily seen that trees admit vertex ranking with $O(\log n)$ colors and hence, Algorithm HS *achieves competitive ratio of $O(\log n)$. This bound is optimal as it holds as a lower bound even when the tree is a simple path.*

5 Open Problems

The main challenge left in this paper is to extend the result for unit discs to arbitrary discs. We conjecture that there exists an online hitting set algorithm for hitting arbitrary discs with a subset of a given set of n points in the plane, which has competitive ratio of $O(\log n)$. A more difficult problem is to design an online hitting set algorithm with a logarithmic competitive ratio for any hypergraph with bounded VC-dimension. To the best of our knowledge, there is no better lower bound except for $\Omega(\log n)$ for this problem.

References

[1] N. Alon, B. Awerbuch, Y. Azar, N. Buchbinder, and J. Naor. The Online Set Cover Problem. *SIAM Journal on Computing*, 39:361, 2009.

[2] B. Ben-Moshe, M. Katz, and J. Mitchell. A constant-factor approximation algorithm for optimal terrain guarding. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 515–524. Society for Industrial and Applied Mathematics, 2005.

[3] A. Borodin and R. El-Yaniv. *Online computation and competitive analysis*, volume 2. Cambridge University Press Cambridge, 1998.

[4] H. Brönnimann and M. Goodrich. Almost optimal set covers in finite VC-dimension. *Discrete and Computational Geometry*, 14(1):463–479, 1995.

[5] P. Cheilaris and S. Smorodinsky. Choosability in geometric hypergraphs. *Arxiv preprint arXiv:1005.5520*, 2010.

[6] K. Chen, H. Kaplan, and M. Sharir. Online conflict-free coloring for halfplanes, congruent disks, and axis-parallel rectangles. *ACM Transactions on Algorithms (TALG)*, 5(2):1–24, 2009.

[7] V. Chvatal. A greedy heuristic for the set-covering problem. *Mathematics of operations research*, 4(3):233–235, 1979.

[8] K. Clarkson and K. Varadarajan. Improved approximation algorithms for geometric set cover. *Discrete and Computational Geometry*, 37(1):43–58, 2007.

[9] G. Even, Z. Lotker, D. Ron, and S. Smorodinsky. Conflict-Free Colorings of Simple Geometric Regions with Applications to Frequency Assignment in Cellular Networks. *SIAM Journal on Computing*, 33:94, 2003.

[10] G. Even, D. Rawitz, and S. Shahar. Hitting sets when the VC-dimension is small. *Information Processing Letters*, 95(2):358–362, 2005.

[11] U. Feige. A threshold of ln n for approximating set cover. *Journal of the ACM (JACM)*, 45(4):634–652, 1998.

[12] D. Hochbaum and W. Maass. Approximation schemes for covering and packing problems in image processing and VLSI. *Journal of the ACM (JACM)*, 32(1):136, 1985.

[13] A. Iyer, H. Ratliff, and G. Vijayan. Optimal node ranking of trees. *Information Processing Letters*, 28(5):225–229, 1988.

[14] D. Johnson. Approximation algorithms for combinatorial problems*. *Journal of Computer and System Sciences*, 9(3):256–278, 1974.

[15] R. Karp. Reducibility among combinatorial problems, 85–103. In *Proc. Sympos., IBM Thomas J. Watson Res. Center, Yorktown Heights, NY*, 1972.

[16] M. Katchalski, W. McCuaig, and S. Seager. Ordered colourings. *Discrete Mathematics*, 142(1-3):141–154, 1995.

[17] V. Kumar and H. Ramesh. Covering rectilinear polygons with axis-parallel rectangles. In *Proceedings of the thirty-first annual ACM symposium on Theory of computing*, pages 445–454. ACM, 1999.

[18] J. Matoušek. *Lectures on Discrete Geometry*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2002.

[19] D. Sleator and R. Tarjan. Amortized efficiency of list update and paging rules. *Communications of the ACM*, 28(2):202–208, 1985.