

On Interference Among Moving Sensors and Related Problems

Jean-Lou De Carufel¹, Matthew J. Katz², Matias Korman^{*3},
André van Renssen⁴, Marcel Roeloffzen⁵, and
Shakhar Smorodinsky^{†6}

- 1 University of Ottawa, Ottawa, Canada
jdecaruf@uottawa.ca
- 2 Ben-Gurion University of the Negev, Beer-Sheva, Israel
matya@cs.bgu.ac.il
- 3 Tohoku University, Sendai, Japan
mati@dais.is.tohoku.ac.jp
- 4 National Institute of Informatics, Tokyo, Japan, and
JST, ERATO, Kawarabayashi Large Graph Project
andre@nii.ac.jp
- 5 National Institute of Informatics, Tokyo, Japan, and
JST, ERATO, Kawarabayashi Large Graph Project
marcel@nii.ac.jp
- 6 Ben-Gurion University of the Negev, Beer-Sheva, Israel, and
École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland
shakhar@math.bgu.ac.il

Abstract

We show that for any set of n moving points in \mathbb{R}^d and any parameter $2 \leq k \leq n$, one can select a fixed non-empty subset of the points of size $O(k \log k)$, such that the Voronoi diagram of this subset is “balanced” at any given time (i.e., it contains $O(n/k)$ points per cell). We also show that the bound $O(k \log k)$ is near optimal even for the one dimensional case in which points move linearly in time. As an application, we show that one can assign communication radii to the sensors of a network of n moving sensors so that at any given time, their interference is $O(\sqrt{n \log n})$. This is optimal up to an $O(\sqrt{\log n})$ factor.

1998 ACM Subject Classification F.2.2 Nonnumerical Algorithms and Problems, G.2.2 Graph Theory

Keywords and phrases Range spaces, Voronoi diagrams, moving points, facility location, interference minimization

Digital Object Identifier 10.4230/LIPIcs.ESA.2016.34

1 Introduction

We consider the following kinetic facility location problem: given n clients (i.e., points) that are moving in \mathbb{R}^d along simple trajectories and a parameter $k \leq n$, we wish to select few of them to become facilities to serve the remaining clients. We follow the usual assumption

* M. K. was supported in part by the ELC project (MEXT KAKENHI Nos. 12H00855 and 15H02665).

† S. S. was partially supported by Grant 1136/12 from the Israel Science Foundation and by the Swiss National Science Foundation Grants 200020144531 and 200021-137574.

that at any instant of time a client is served by its nearest facility. Our aim is to select the facilities so that none serves too many customers. Specifically, we wish to maintain the invariant that at any given time the number of clients served by each of the chosen facilities is bounded by n/k .

The pigeon-hole principle directly implies that we cannot select fewer than k facilities. Our main result is that a subset of size $O(k \log k)$ will suffice. We also show that one cannot improve this bound to $O(k)$, even for $d = 1$. As an application, we show how to construct a communication graph among a set of n moving sensors such that at any given time, the interference of the communication graph is bounded by $O(\sqrt{n \log n})$ (and its hop-diameter is three). Intuitively speaking, the interference of a sensor is the in-degree (i.e., the number of sensors that can communicate to him directly, see more details in Section 5). This bound is near optimal as already, in the static case, there are examples where the interference is at least $\Omega(\sqrt{n})$ [6].

In order to obtain our results we use the machinery of geometric hypergraphs and the theory of VC-dimension and ε -nets. By a geometric hypergraph (also called a range-space) we mean the following: suppose we are given a finite set P of points in \mathbb{R}^d and a family of simple geometric regions, such as the family of all halfspaces in \mathbb{R}^d . Then we consider the combinatorial structure of the set system $(P, \{h \cap P\})$ where h is any halfspace. A key property of such hypergraphs is bounded VC-dimension (see Section 2 for exact definitions). In this paper we study a more general structure by allowing the underlying set of points to move along some “reasonable” trajectories (i.e., the coordinates of each point can be described with a polynomial function of bounded degree). Even though the static case is well-known, little research has been done for the case in which the points move. We show that those more complex hypergraphs, defined as the union of all hypergraphs obtained at all possible times, still have a bounded VC-dimension.

In addition to the above mentioned applications, we believe that the bounded VC-dimension of such hypergraphs is of independent interest and to the best of our knowledge has not been observed before. We hope that this paper will have many follow-up applications, since bounded VC-dimension has applications in many other areas of mathematics and computer science.

The paper is organized as follows: in Section 2 we introduce several key concepts as well as review known results that hold for static range spaces. In Section 3 we extend these results to the kinetic case. In Sections 4 and 5 we prove our main results concerning Voronoi diagrams for moving points and the interference problem mentioned above.

2 Preliminaries and Previous Work

A hypergraph $H = (V, \mathcal{E})$ is a pair of sets such that $\mathcal{E} \subseteq 2^V$. A geometric hypergraph is one that can be realized in a geometric way. For example, consider the hypergraph $H = (V, \mathcal{E})$, where V is a finite subset of \mathbb{R}^d and \mathcal{E} consists of all subsets of V that can be cut-off from V by intersecting it with a shape belonging to some family of “nice” geometric shapes, such as the family of all halfspaces. The elements of V are called *vertices*, and the elements of \mathcal{E} are called *hyperedges*. For a subset $V' \subseteq V$, the hypergraph $H[V'] = (V', \{V' \cap S : S \in \mathcal{E}\})$ is the *sub-hypergraph* induced by V' .

We consider the following families of geometric hypergraphs: Let P be a set of points in \mathbb{R}^2 (or, in general, in \mathbb{R}^d) and let \mathcal{R} be a family of regions in the same space. We refer to the hypergraph $H = (P, \{P \cap r : r \in \mathcal{R}\})$ as the hypergraph induced by P with respect to \mathcal{R} . When \mathcal{R} is clear from the context, we sometimes refer to it as the hypergraph induced

by P . In the literature, hypergraphs that are induced by points with respect to geometric regions of some specific kind are also referred to as *range spaces*. We sometimes abuse the notation and write (P, \mathcal{R}) , instead of $H = (P, E)$, where $E = \{P \cap r : r \in \mathcal{R}\}$.

ε -nets and VC-dimension

A subset $T \subset V$ is called a *transversal* (or a *hitting set*) of a hypergraph $H = (V, \mathcal{E})$, if it intersects all sets of \mathcal{E} . The *transversal number* of H , denoted by $\tau(H)$, is the smallest possible cardinality of a transversal of H . The fundamental notion of a transversal of a hypergraph is central in many areas of combinatorics and its relatives. In computational geometry, there is a particular interest in transversals, since many geometric problems can be rephrased as questions on the transversal number of certain hypergraphs. An important special case arises when we are interested in finding a small size set $N \subset V$ that intersects all “relatively large” sets of \mathcal{E} . This is captured in the notion of an ε -net for a hypergraph:

► **Definition 1** (ε -net). Let $H = (V, \mathcal{E})$ be a hypergraph with V finite. Let $\varepsilon \in [0, 1]$ be a real number. A set $N \subseteq V$ (not necessarily in \mathcal{E}) is called an ε -net for H if for every hyperedge $S \in \mathcal{E}$ with $|S| \geq \varepsilon|V|$ we have $S \cap N \neq \emptyset$.¹

In other words, a set N is an ε -net for a hypergraph $H = (V, \mathcal{E})$ if it stabs all “large” hyperedges (i.e., those of cardinality at least $\varepsilon|V|$). The well-known result of Haussler and Welzl [7] provides a combinatorial condition on hypergraphs that guarantees the existence of small ε -nets (see below). This requires the following well-studied notion of the Vapnik-Chervonenkis dimension [16]:

► **Definition 2** (VC-dimension). Let $H = (V, \mathcal{E})$ be a hypergraph. A subset $X \subset V$ (not necessarily in \mathcal{E}) is said to be *shattered* by H if $\{X \cap S : S \in \mathcal{E}\} = 2^X$. The *Vapnik-Chervonenkis dimension*, also denoted the *VC-dimension* of H , is the maximum size of a subset of V shattered by H .

Relation between ε -nets and the VC-dimension

Haussler and Welzl [7] proved the following fundamental theorem regarding the existence of small ε -nets for hypergraphs with small VC-dimension.

► **Theorem 3** (ε -net theorem). Let $H = (V, \mathcal{E})$ be a hypergraph with VC-dimension d . For every $\varepsilon \in (0, 1)$, there exists an ε -net $N \subset V$ with cardinality at most $O\left(\frac{d}{\varepsilon} \log \frac{1}{\varepsilon}\right)$.

In fact, it can be shown that a random sample of vertices of size $O\left(\frac{d}{\varepsilon} \log \frac{1}{\varepsilon}\right)$ is an ε -net for H with a positive constant probability (see [10] for details on how to compute such nets).

Many hypergraphs studied in computational geometry and learning theory have a “small” VC-dimension, where by “small” we mean a constant independent of the number of vertices of the underlying hypergraph. It is known that whenever range spaces are defined through semi-algebraic sets of constant description complexity (i.e., sets defined as a Boolean combination of a constant number of polynomial equations and inequalities of constant maximum degree), the resulting hypergraph has finite VC-dimension. Halfspaces, balls, boxes, etc. are examples of such sets; see, e.g., [11, 13] for more details.

¹ An analogous definition applies when V is not necessarily finite and H is endowed with a probability measure.

Thus, by Theorem 3, these hypergraphs admit “small” size ε -nets. Kórnlos *et al.* [8] proved that the bound $O(\frac{d}{\varepsilon} \log \frac{1}{\varepsilon})$ on the size of an ε -net for hypergraphs with VC-dimension d is best possible. Namely, for a constant d , they construct a hypergraph H with VC-dimension d such that any ε -net for H must have size of at least $\Omega(\frac{1}{\varepsilon} \log \frac{1}{\varepsilon})$. Recently, several breakthrough results provided better lower and upper bounds on the size of ε -nets in several special cases [1, 2, 14].

3 Kinetic hypergraphs

We start by extending the concept of geometric hypergraphs to the kinetic model. Let $P = \{p_1, \dots, p_n\}$ denote a set of n moving points in \mathbb{R}^d , where each point is moving along some “simple” trajectory. That is, each p_i is a function $p_i : [0, \infty) \rightarrow \mathbb{R}^d$ of the form $p_i(t) = (x_1^i(t), \dots, x_d^i(t))$, where $x_j^i(t)$ is a univariate polynomial ($1 \leq j \leq d$). For a given real number $t \geq 0$ and a subset $P' \subset P$, we denote by $P'(t)$ the fixed set of points $\{p(t) : p \in P'\}$.

Let \mathcal{R} be a (not necessarily finite) family of ranges; for example, the family of all half-spaces in \mathbb{R}^d . We define the *kinetic hypergraph* induced by \mathcal{R} :

► **Definition 4** (kinetic hypergraph). Let P be a set of moving points in \mathbb{R}^d and let \mathcal{R} be a family of ranges. Let (P, \mathcal{E}) denote the hypergraph where \mathcal{E} consists of all subsets $P' \subseteq P$ for which there exists a time t and a range $r \in \mathcal{R}$ such that $P'(t) = P(t) \cap r$. We call (P, \mathcal{E}) the *kinetic hypergraph* induced by \mathcal{R} .

As in the static case we abuse the notation and denote the kinetic hypergraph by (P, \mathcal{R}) . In order to apply our techniques, we need the following “bounded description complexity” assumption concerning the movement of the points of P . We say that a point $p_i = p_i(t) = (x_1^i(t), \dots, x_d^i(t)) \in P$ moves with *description complexity* $s > 0$ if for each $1 \leq j \leq d$, the univariate polynomial $x_j^i(t)$ has degree at most s . In the remainder of this paper, we assume that $P(0)$ is in “general position”. That is, at time $t = 0$ no $d + 1$ points of $P(0)$ are on a common hyperplane. This assumption can be removed through usual symbolic perturbation techniques.

3.1 VC-Dimension of kinetic hypergraphs

In this section we prove that for many of the static range spaces that have small VC-dimension, their kinetic counterparts also have small VC-dimension. We start with the family \mathcal{H}_d of all halfspaces in \mathbb{R}^d .

► **Theorem 5.** *Let $P \subset \mathbb{R}^d$ be a set of moving points with bounded description complexity s . Then, the kinetic-range space (P, \mathcal{H}_d) has VC-dimension bounded by $O(d \log d)$.*

To prove Theorem 5, we need the following known definition and lemma (see, e.g., [11]). The *primal shatter function* of a hypergraph $H = (V, \mathcal{E})$ denoted by π_H is a function:

$$\pi_H : \{1, \dots, |V|\} \rightarrow \mathbb{N}$$

defined by $\pi_H(i) = \max_{V' \subseteq V, |V'|=i} |H[V']|$, where $|H[V']|$ denotes the number of hyperedges in the sub-hypergraph $H[V']$.

► **Lemma 6.** *Let $H = (V, \mathcal{E})$ be a hypergraph whose primal shatter function π_H satisfies $\pi_H(m) = O(m^c)$ for some constant $c \geq 2$. Then the VC-dimension of H is $O(c \log c)$.*

We provide a sketch of the proof of Lemma 6 for the sake of completeness.

Proof. Let d denote the VC-dimension of H , and let $V' \subseteq V$ be a shattered subset of cardinality d . On one hand it means that the number of possible subsets of V' that can be realized as the intersection of V' and a hyperedge in \mathcal{E} is 2^d . On the other hand, by our assumption on π_H , for a subset of size d , there can be at most Ad^c hyperedges in the sub-hypergraph induced by it, for some appropriate constant A . In other words we have $2^d \leq \pi_H(d) \leq Ad^c$. This is easily seen to imply that $d = O(c \log c)$. Indeed, by choosing, say, $d = 10Ac \log c$, the above inequality does not hold, a contradiction. This completes the proof of the lemma. ◀

Proof of Theorem 5. By Lemma 6 it suffices to bound the primal shatter function $\pi_{\mathcal{H}_d}(m)$ by a polynomial of constant degree. It is a well known fact that the number of combinatorially distinct half-spaces determined by n (static) points in \mathbb{R}^d is $O(n^d)$. This can be easily seen by charging hyperplanes to d -tuples of points (using rotations and translations) and observing that each tuple can be charged at most a constant (depending on the dimension d) number of times. Thus, at any given time, the number of hyperedges is bounded by $O(n^d)$. Next, note that as t varies, a combinatorial change in the hypergraph $(P(t), \mathcal{R})$ can occur only when $d + 1$ points $p_1(t), \dots, p_{d+1}(t)$ become affinely dependent. Indeed, a hyperedge is defined by a hyperplane that contains d points of $P(t)$, and that hyperedge changes when an additional point of $P(t)$ crosses the hyperplane (and thus $d + 1$ points become affinely dependent). This happens if and only if the following determinant condition holds:

$$\begin{vmatrix} x_1^1(t) & x_2^1(t) & \cdots & x_d^1(t) & 1 \\ x_1^2(t) & x_2^2(t) & \cdots & x_d^2(t) & 1 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ x_1^{d+1}(t) & x_2^{d+1}(t) & \cdots & x_d^{d+1}(t) & 1 \end{vmatrix} = 0 \tag{1}$$

where $x_i^j(t)$ denotes the i 'th coordinate of $p_j(t)$. The left side of the equation is a univariate polynomial of degree at most ds . By our general position assumption this polynomial is not identically zero and thus can have at most ds solutions.

That is, a tuple of $d + 1$ points of $P(t)$ generates at most ds events. Hence, the total number of such events is bounded by $O(\binom{n}{d+1}) = O(n^{d+1})$. Between any two events we have a fixed set of at most $O(n^d)$ distinct hyperedges, thus we can have $O(n^{2d+1})$ distinct hyperedges along all instants of time.

Since each hyperedge is defined by the points on its boundary, this property is hereditary. That is, for any subset $P' \subseteq P$ the hypergraph $H[P']$ has at most $O(|P'|^{2d+1})$ hyperedges. Thus, the shatter function satisfies $\pi_H(m) = O(m^{2d+1})$. Then by Lemma 6, (P, \mathcal{H}_d) has bounded constant VC-dimension, where the constant depends only on d and s . ◀

Theorem 5 can be further generalized to arbitrary ranges with so-called bounded description complexity as defined below:

► **Theorem 7.** *Let \mathcal{R} be a collection of semi-algebraic subsets of \mathbb{R}^d , each of which can be expressed as a Boolean combination of a constant number of polynomial equations and inequalities of maximum degree c (for some constant c). Let P be a set of moving points in \mathbb{R}^d with bounded description complexity. Then the kinetic range-space (P, \mathcal{R}) has bounded VC-dimension.*

Proof. The proof combines Lemma 6 with Theorem 5 and the so-called Veronese lifting map from Algebraic Geometry. We omit the details as it is very similar to the proof for the static case. See, e.g., [11]. ◀

4 Balanced Voronoi cells for moving points

In this section we tackle the facility location problem for a set of moving clients, where the goal is to ensure a balanced division of the load among the facilities at any instance of time. Given a set P of moving points or *clients* in \mathbb{R}^d , locate a small number of the points to serve as *facilities* so that at every instance of time no facility is serving more than n/k clients. We make the usual assumption that each client goes to its nearest facility. In the following we show how to obtain an almost optimal balancing (up to a $\log k$ factor), even under the restriction that facilities may be located only at points of P .

► **Theorem 8.** *Let $P = \{p_1, \dots, p_n\}$ be any set of n moving points in \mathbb{R}^d with bounded description complexity. For any integer $2 \leq k \leq n$, there exists a subset $N \subset P$ of cardinality $O(k \log k)$, such that for any time $t \geq 0$, each cell of the Voronoi diagram $\text{Vor}(N(t))$ contains at most $O(n/k)$ points of $P(t)$.*

Before proceeding with the proof of Theorem 8 we need the following result. An *infinite cone* with apex $a \in \mathbb{R}^d$ and angle $\theta \in \mathbb{R}$ is defined as the union of all halflines emanating from a whose orientations belong to some fixed cap of the sphere \mathcal{S}^{d-1} . Equivalently, it can be defined as the set:

$$\{x \in \mathbb{R}^d : (x - a) \cdot (b - a) \geq \|x - a\| \cos \theta\},$$

where “ $\|\cdot\|$ ” denotes the Euclidean norm, “ \cdot ” denotes the scalar product and b is a vector such that $\|b - a\| = 1$. A *bounded cone* is the intersection of an infinite cone with a ball centered at its apex.

► **Lemma 9.** *Let P be a set of moving points in \mathbb{R}^d with bounded description complexity s , and let \mathcal{R} be the family of all bounded cones. The kinetic hypergraph (P, \mathcal{R}) has bounded VC-dimension.*

Proof. As shown above, the boundary surface of an infinite cone is a quadric (i.e., a polynomial of degree 2). In particular, the ranges of \mathcal{R} can be expressed as semi-algebraic sets of constant description complexity. Thus, by Theorem 7 the hypergraph (P, \mathcal{R}) has constant VC-dimension as claimed. ◀

Proof of Theorem 8. Let \mathcal{W} be the family of all bounded cones in \mathbb{R}^d . Let $H = (P, \mathcal{W})$ be the corresponding kinetic hypergraph. By Lemma 9, H has constant VC-dimension.

We fix $\varepsilon = \frac{1}{k}$ and let $N \subset P$ be an ε -net for H of size $O(k \log k)$ (refer to Theorem 3). We show that N satisfies the desired property. That is, for any time $t \geq 0$ and point $q \in N$, the Voronoi cell of $q(t)$ in the Voronoi diagram $\text{Vor}(N(t))$ contains at most $O(n/k)$ points of $P(t)$. Let C_d be the minimum number of sixty-degree caps that are needed to cover the unit sphere \mathcal{S}^{d-1} . Using packing arguments it is easily seen that C_d is a constant that depends only on d ; see, e.g., [3].

Assume to the contrary that the Voronoi cell of $q(t)$ contains a subset $P'(t) \subset P(t)$ of more than $C_d n/k$ points. By definition, each of the points in $P'(t)$ is closer to $q(t)$ than to any other point in $N(t)$. By the pigeonhole principle, there is an infinite sixty-degree cone W which has $q(t)$ as its apex and that contains at least $n/k + 1$ of the points of $P'(t)$. Sort the points of $P'(t) \cap W$ in increasing distance from $q(t)$; let $p_1(t), \dots, p_j(t)$ be the obtained order (note that by assumption, we have $j \geq n/k + 1$). Slightly perturb the cone W and bound it to obtain a bounded cone W' that contains the points $p_1(t), \dots, p_j(t)$ but does not contain $q(t)$ (or any other point of $P(t)$). This can always be done by choosing a sufficiently large

radius, doing an infinitesimally small translation of the apex and (if necessary) changing the angle of the cone. Since N is an ε -net with respect to bounded cones, W' must contain a point $q'(t) \in N(t)$ (other than $q(t)$).

Since any point in $P(t) \cap W'$ also belongs to W , which is a cone of sixty degrees, any point $p(t) \in P(t) \cap W'$ for which $d(p(t), q(t)) \geq d(q'(t), q(t))$ must be closer to $q'(t)$ than to $q(t)$. In particular, $p_j(t)$ satisfies this inequality and thus belongs to the Voronoi cell of $q'(t)$ (and not of $q(t)$), which is a contradiction. \blacktriangleleft

In fact, a careful look at the proof above shows that the following stronger result holds:

► **Corollary 10.** *Let $N \subset P$, $|N| = O(k \log k)$, as in Theorem 8. Then, for any finite point set $S \subset \mathbb{R}^d$, and for any $t \geq 0$, the cell of any $q \in S$ in the Voronoi diagram $\text{Vor}(N(t) \cup S)$ contains at most $O(n/k)$ points of $P(t)$.*

Remark

We note that the bound of $O(k \log k)$ in Theorem 8 is near optimal. Clearly, if there are only $o(k)$ points in N then by the pigeonhole principle one of the Voronoi cells must contain $\omega(n/k)$ points of P . We also note that reducing the size of the set N seems to be out of reach and maybe impossible, even for the one dimensional case where the points move with constant speed. This follows from a recent lower-bound construction of Alon [1] for ε -nets for static hypergraphs consisting of points with respect to strips in the plane.


Indeed, assume that $d = 1$ and each point $p \in P$ is described with a linear equation of the form $p(t) = at + b$ (i.e., a line). Assume that there exists a subset $N \subset P$ such that for any $t > 0$ and $q \in N$, the Voronoi cell of $q(t)$ contains at most n/k points of $P(t)$. In particular, this implies that there are at most $2n/k$ points of $P(t)$ between any pair of consecutive points of $N(t)$. If we view the moving points in \mathbb{R} as lines in \mathbb{R}^2 , this is equivalent to choosing a subset of the lines with the property that any vertical segment (i.e., a range of the form $t_0 \times [c, d]$ for constants $t_0 > 0$, $c, d \in \mathbb{R}$) that intersects more than $2n/k$ of the above lines will also intersect one of the chosen lines. By standard point-line duality in two dimensions, this is equivalent to the problem of finding an $\varepsilon = \frac{2}{k}$ -net for points with respect to strips in the plane, which still remains an open problem. Recently, Alon [1] gave a construction showing that such hypergraphs cannot have linear (in $\frac{1}{\varepsilon}$) size ε -nets. Since their problem can be reduced to ours, the same lower bound holds for our problem.

5 Low interference for moving transmitters

Here we show how to tackle the problem of minimizing interference among a set of wireless moving transmitters while keeping the number of topological changes of the underlying network subquadratic. In the following we define the concept of (receiver-based) *interference* of a set of ad-hoc sensors [17] (see Figure 1).

► **Definition 11.** Let $P = \{p_1, \dots, p_n\}$ be a set of n points in \mathbb{R}^d and let r_1, \dots, r_n be n non-negative reals representing power levels (or transmission radii) assigned to the points p_1, \dots, p_n , respectively. Let $G = (P, E)$ be the graph associated with this power assignment, where $E = \{\{p, q\} : d(p, q) \leq \min\{r_p, r_q\}\}$. That is, points p, q are neighbors in G if and only if p is contained in the ball centered at q with radius r_q and *vice versa*. Let $D = \{d_1, \dots, d_n\}$ denote the set of balls where d_i is the ball centered at p_i and having radius r_i .

Let $I(D)$ denote the maximum depth of the arrangement of the balls in D . That is $I(D) = \max_{q \in \mathbb{R}^d} |\{d \in D : q \in d\}|$. We call $I(D)$ the *interference* of D , which is also the



■ **Figure 1** Given a set of fixed points in \mathbb{R}^2 and their power assignments represented by disks, the interference is the deepest point in the arrangement of the disks (the highlighted region in the figure). The underlying communication graph is shown with solid edges.

interference of the network. Note that both G and $I(D)$ are determined by P and r_1, \dots, r_n . Given a set P of points in \mathbb{R}^d , the *interference* of P (denoted $I(P)$) is the smallest possible interference $I(D)$, where D corresponds to a power assignment whose associated graph is connected. The *interference minimization problem* asks for the power assignment for which $I(P) = I(D)$.

Empirically, (in dimension two) it has been observed that networks with high interference have high rates of message collision. This requires messages to be repeated often, which slows down the network and reduces battery life of the sensors [17]. Thus, a significant amount of research has focused in the creation of connected networks with low interference (see, e.g., [6, 9]). It is known that computing $I(P)$ (or even approximating it by a constant factor) is an NP-complete problem [4], but some worst-case bounds are known.

► **Theorem 12** ([6]). *Let P be a set of n points in the plane. Then $I(P) = O(\sqrt{n})$. Furthermore, this bound is asymptotically tight, in the sense that for any n there exists a set P of n points such that $I(P) = \Omega(\sqrt{n})$.*

Here, we turn our attention to the kinetic version of the interference problem in arbitrary but fixed dimension. We wish to maintain a connected graph of a set of moving points (representing moving sensors) that always has low interference. Unless the distances between sensors remain constant, no static radii assignment can work for a long period of time (since points will eventually be far from each other). Instead, we describe the network in a combinatorial way. That is, we look for a function $f : P \times [0, \infty) \rightarrow P$ that determines, for each sensor of P and instance of time, its furthest away sensor that must be reached. Then, at time t the communication radius of a sensor $p \in P$ is simply set equal to the distance between p and $f(p, t)$. Ideally, we would like to construct a network that not only has small interference at any instance of time, but also the underlying graph has a small amount of combinatorial changes along time.

Our algorithm to maintain a connected graph is based on the ideas used in [6] for the static case. We first pick a subset $N \subset P$ of “hubs”. Those hubs will never change along time and will always have transmission radius big enough to cover all other points. Each other point in $P \setminus N$ will be assigned at every instance of time to its nearest hub. In the following we show that a careful choice of hubs will ensure a small interference, and overall small number of combinatorial changes in the radii assignment protocol. To bound the number of combinatorial changes, we need to use the machinery of Davenport-Schinzel sequences: A finite sequence $\sigma = (e_1, \dots, e_m)$ over an alphabet of n symbols is called a *Davenport-Schinzel sequence* of order t when no two consecutive elements of σ are equal, and for any two distinct symbols x, y , there does not exist a subsequence where x and y alternate $t + 2$

times. Several bounds are known on the maximum length of Davenport-Schinzel sequences of a given order. In particular, we are interested in upper bounds. See [15] for more details on Davenport-Schinzel sequences.

► **Theorem 13** (Upper bound on Davenport-Schinzel sequences [12]). *A Davenport-Schinzel sequence of order t on n symbols has length at most $O(n2^{O(\alpha(n)^{\lfloor (t-2)/2 \rfloor})})$, where $\alpha(n)$ is the inverse of the Ackermann function.*

The Ackermann function is a function that grows very rapidly, hence its inverse is usually regarded as a small constant (indeed, it is known that $\alpha(n) \leq 5$ for any input that can be stored explicitly in current computers). Davenport-Schinzel sequences are often used to bound the complexity of upper (or lower) envelopes of polynomial functions. Whenever we have a family of n functions such that no two graphs of those functions cross more than t times (for some bounded constant t), we can use Theorem 13 to bound the complexity of their upper and lower envelope.

► **Theorem 14.** *Let P be a set of n moving points in \mathbb{R}^d with bounded description complexity s . Then, there is a power assignment with updates, such that at any given time t the interference of the network is at most $O(\sqrt{n \log n})$. Moreover, the total number of combinatorial changes in the network is at most $O^*(n^{1.5} \sqrt{\log n})$, where the O^* notation hides a term involving the inverse Ackermann function that depends on d and s .*

Proof. We use Theorem 8 for some value of k that will be determined later. We obtain a set N of size $O(k \log k)$ with the properties guaranteed by Theorem 8. The elements of N are called *hubs*, and we assign to each of them the largest possible radius. That is, at any instance of time $t \geq 0$, a point $p \in N$ is assigned the distance to its furthest point in P . In other words, $f(p, t)$ is equal to the point $q \in P$ that maximizes the distance $d(p(t), q(t))$. Other points of P are assigned the distance to their nearest hub. More formally, $f(p, t)$, for a point $p \in P \setminus N$, is equal to the point $q \in N$ that minimizes the distance $d(p(t), q(t))$. Equivalently, if we consider the Voronoi diagram with sites $N(t)$, the function $f(p, t)$ will match $p(t)$ with the site associated to the Voronoi cell that contains $p(t)$ at time t .

First observe that the network is connected: indeed, all hubs are connected to each other forming a clique. Moreover, each point of $P \setminus N$ has radius large enough to reach one point of N . In particular, any two points of P can reach each other after hopping through at most two intermediate sensors of N (thus, the constructed network has diameter 3).

We now pick the correct value of k so that the interference of this protocol is minimized. Since N has $O(k \log k)$ points, the overall interference contribution by hubs is bounded by the same amount. By Corollary 10, we also know that no point $q \in \mathbb{R}^d$ can be reached by more than $O(n/k)$ points of $P \setminus N$ at any instance of time. That is, the total interference of any point $q \in \mathbb{R}^d$ is at most $O(k \log k)$ from hubs, and at most $O(n/k)$ from non-hubs. Thus, by setting $k = \sqrt{n/\log n}$ we obtain the claimed bound.

We now bound the total number of combinatorial changes that will happen to the network along time. Let $p \in P$, we will show that the number of combinatorial changes of p is bounded. Recall that, if p is a hub it will connect to its furthest away point of P . Otherwise, p will connect to its nearest hub. In either case, it suffices to bound the number of combinatorial changes of the nearest/furthest point within a group of moving points with respect to the moving point p . Equivalently, we are looking at the number of combinatorial changes of the upper envelope of the family of functions $\mathcal{F}_1 = \{d(p(t), p'(t)) : p' \in P\}$ for points $p \in N$, or the lower envelope of the family of functions $\mathcal{F}_2 = \{d(p(t), p'(t)) : p' \in N\}$ for points $p \notin N$. By the bounded description complexity assumption, functions of \mathcal{F}_1 and

\mathcal{F}_2 are such that the graphs of any pair of them cross $O(s)$ times. Thus, by the Davenport-Schinzel Theorem we can bound the number of combinatorial changes of the upper envelope of \mathcal{F}_1 by $O(\lambda_{O(s)}(n))$, where $\lambda_t(m)$ denotes the maximum length of a Davenport-Schinzel sequence of order t on m symbols. Similarly, the number of changes of the lower envelope of \mathcal{F}_2 is bounded by $O(\lambda_{O(s)}(\sqrt{n \log n}))$.

Ignoring the terms that depend on the inverse of the Ackermann function, we have that for any fixed constant s , $\lambda_t(m) = O^*(m)$. Combining this with the fact that we have $O(\sqrt{n \log n})$ hubs and at most n non-hub points, the overall number of combinatorial changes is bounded by $O^*(n \times \sqrt{n \log n} + \sqrt{n \log n} \times n) = O^*(n^{1.5} \sqrt{\log n})$ as claimed. ◀

6 Conclusion

Using the the machinery of VC-dimension we have shown that the difference between static and kinetic environments for our facility location problem is small. We believe that a similar approach can be used for other problems. Some directly follow from Theorem 7 (such as kinetic range counting or discrepancy, see details in the extended version of this paper [5]). We hope that future research will show other interesting applications.

In Section 4 we argued that it is unlikely that the “balanced” property can be significantly improved. Similarly, it seems unlikely that the “reasonable” constraint can be removed, even in one dimension. Indeed, if points are allowed to move arbitrarily, they can create all $n!$ orderings along time. In particular, for any set $N \subset P$ we can always find a time and range that contains all points of $P \setminus N$. Thus, no subset $N \subset P$ can act as an ε -net for all instances of time. Further note that, since the alternation in orderings can be repeated arbitrarily many times, the number of times that we need to change the set N can also be unbounded. This behaviour can be created with trigonometric functions of low description complexity.

Acknowledgements. This work was initiated during the Second Sendai Winter Workshop on Discrete and Computational Geometry. The authors would like to thank the other participants for interesting discussions during the workshop, as well as Alexandre Rok for helpful discussions.

References

- 1 N. Alon. A non-linear lower bound for planar epsilon-nets. *Discrete & Computational Geometry*, 47(2):235–244, 2012. doi:10.1007/s00454-010-9323-7.
- 2 B. Aronov, E. Ezra, and M. Sharir. Small-size epsilon-nets for axis-parallel rectangles and boxes. *SIAM Journal on Computing*, 39(7):3248–3282, 2010.
- 3 K. Böröczky. Packing of spheres in spaces of constant curvature. *Acta Mathematica Academiae Scientiarum Hungarica*, 32(3-4):243–261, 1978. doi:10.1007/BF01902361.
- 4 Y. Brise, K. Buchin, D. Eversmann, M. Hoffmann, and W. Mulzer. Interference minimization in asymmetric sensor networks. In *ALGOSENSORS 2014*, pages 136–151, 2014. doi:10.1007/978-3-662-46018-4_9.
- 5 J.-L. De Carufel, M. Katz, M. Korman, A. van Renssen, M. Roeloffzen, and S. Smorodinsky. On kinetic range spaces and their applications. *CoRR*, abs/1507.02130, 2015. URL: <http://arxiv.org/abs/1507.02130>.
- 6 M.M. Halldórsson and T. Tokuyama. Minimizing interference of a wireless ad-hoc network in a plane. *Theoretical Computer Science*, 402(1):29–42, 2008. doi:10.1016/j.tcs.2008.03.003.

- 7 D. Haussler and E. Welzl. Epsilon-nets and simplex range queries. *Discrete & Computational Geometry*, 2:127–151, 1987.
- 8 J. Komlós, J. Pach, and G.J. Woeginger. Almost tight bounds for epsilon-nets. *Discrete & Computational Geometry*, 7:163–173, 1992.
- 9 M. Korman. Minimizing interference in ad-hoc networks with bounded communication radius. *Information Processing Letters*, 112(19):748–752, 2012. doi:10.1016/j.ipl.2012.06.021.
- 10 J. Matoušek. Approximations and optimal geometric divide-and-conquer. In *Symposium on the Theory of Computing*, pages 505–511, 1991.
- 11 J. Matoušek. *Lectures on Discrete Geometry*. Springer-Verlag New York, Inc., 2002.
- 12 G. Nivasch. Improved bounds and new techniques for Davenport–Schinzel sequences and their generalizations. *Journal of the ACM*, 57(3), 2010.
- 13 J. Pach and P. K. Agarwal. *Combinatorial Geometry*. Wiley Interscience, 1995.
- 14 J. Pach and G. Tardos. Tight lower bounds for the size of epsilon-nets. In *Symposium on Computational Geometry*, pages 458–463, 2011.
- 15 M. Sharir and P. K. Agarwal. *Davenport-Schinzel Sequences and Their Geometric Applications*. Cambridge University Press, New York, 1995.
- 16 V. N. Vapnik and A. Ya. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and its Applications*, 16(2):264–280, 1971.
- 17 P. von Rickenbach, R. Wattenhofer, and A. Zollinger. Algorithmic models of interference in wireless ad hoc and sensor networks. *IEEE/ACM transactions on sensor networks*, 17(1):172–185, 2009. doi:10.1109/TNET.2008.926506.