SINR Diagram with Interference Cancellation

Chen Avin [*] Asaf Cohen [*] Yoram Haddad [*] [†] Erez Kantor [‡] Zvi Lotker [*¶]
Merav Parter [§¶] David Peleg [§¶]

Abstract

This paper studies the reception zones of a wireless network in the SINR model with receivers that employ interference cancellation (IC). IC is a recently developed technique that allows a receiver to decode interfering signals, and *cancel* them from the received signal in order to decode its intended message. We first derive the important topological properties of the reception zones and their relation to *high-order Voronoi diagrams* and other geometric objects. We then discuss the computational issues that arise when seeking an efficient description of the zones. Our main fundamental result states that although potentially there are exponentially many possible cancellation orderings, and as a result, reception zones, in fact there are much fewer nonempty such zones. We prove a linear bound (hence tight) on the number of zones and provide a polynomial time algorithm to describe the diagram. Moreover, we introduce a novel parameter, the *Compactness Parameter*, which influences the tightness of our bounds. We then utilize these properties to devise a logarithmic time algorithm to answer point-location queries for networks with IC.

1 Introduction

1.1 Background and Motivation Today wireless networks are embedded in our daily lives, with an ever-growing use of cellular, satellite and sensor networks. Thus, the capacity of wireless networks, i.e., the maximum *achievable rate* by which stations can communicate reliably, has received an increasing attention in recent years [11, 18, 14, 10, 4, 2, 12]. The great advantage of wireless communication, the broadcast nature of the medium, also creates its biggest obstacle, interference. When a station wants to decode a message (i.e., a signal) sent from a transmitter, it must cope with all other (legitimate) simultaneous neighboring transmissions.

Roughly speaking, two basic approaches to handle interference dominated the research community for many years [9]. One of the approaches was to use orthogonalization. By using time-division (TDMA) or frequency division (FDMA), the degrees of freedom in the channel can be divided between the participating transmitters. This generates an independent channel for each transmitter. The second approach was to treat the interference as noise. This way, together with the ambient (or background) noise, the interference disrupts the signal reception and decoding abilities. For the signal to be safely decoded, the *Signal to Interference & Noise Ratio (SINR)* must be large enough.

Due to the increasingly large number of users, the *achievable rate* or *utilization* of wireless networks has become the bottleneck of the communication. Therefore, one of the main challenges for wireless network designers is to increase this rate and try to reach the capacity of the network. In a sense, both of the aforementioned approaches treat interference in wireless communication as a foe, and try to either avoid it or overcome it. This paper focuses on a relatively recent and promising method for efficient decoding called *interference cancellation (IC)* [1].

The basic idea of interference cancellation, and in particular *successive interference cancellation (SIC)* is quite simple. First, the strongest interfering signal is detected and decoded. Once decoded, this signal can then be subtracted ("canceled") from the original signal. Subsequently, the next strongest interfering signal can be detected and decoded from the now "cleaner" signal, and so on. Optimally, this process continues until all interferences are cancelled and we are left with the desired transmitted signal, which can now be decoded. SIC is similar in spirit to several well known algorithms like the Gram-Schmidt process [21], solving triangular systems of linear equations, and fountain codes [5]. It should be noted that without using IC, every station can decode at most one transmitter (i.e., the strongest signal it receives). In contrast, with IC, every station can decode more transmitters, or expressed dually, every transmitter can reach more receivers. This clearly increases the utilization of the network.

[*]Ben Gurion University, Beer-Sheva, Israel. {avin,coasaf,zvilo}@cse.bgu.ac.il, yoram.haddad@gmail.com
[†]Jerusalem College of Technology,Jerusalem, Israel
[‡]The Technion, Haifa, Israel. erez.kantor@gmail.com
[§]The Weizmann Institute of Science, Rehovot, Israel. {merav.parter,david.peleg}@ weizmann.ac.il.
[¶]Supported by a grant of the Israel Science Foundation.

Interference cancellation is fairly well-studied from an information-theoretic point of view. In fact, it is the optimal strategy in several scenarios, such as strong interference [19, 6], corner points of a multiple access channel [7, Chapter 14], and spread spectrum communication (CDMA) [22], and it constitutes a key building block in the best known bounds for the capacity of the interference channel [9]. Nevertheless, to the best of our knowledge, reception zones under interference cancellation (areas in which transmitters can be decoded), as well as algorithmic issues for large networks, are a virgin land yet to be explored.

In this paper, we initiate the study of the topological properties of the reception zones in the context of IC setting, discuss the computational issues arising when trying to compute these reception zones or answer queries regarding specific points, and devise polynomial-time algorithms to address these problems. This is done by extending *SINR diagrams* [3] to the setting of stations that can apply successive interference cancellation. The SINR diagrams of a wireless network of n transmitters $s_1, s_2, \ldots s_n$ partitions the plane into reception zones $\mathcal{H}(s_1), \mathcal{H}(s_2), \ldots \mathcal{H}(s_n)$, one per station, and the complementary region of the plane where no station can be decoded, is denoted $\mathcal{H}(\emptyset)$. In [3], SINR diagrams have been studied for the specific case where all stations use the same transmission power, i.e., *uniform power*. It is shown there that the reception zones have some "nice" properties, like being convex (hence connected) and fat. In [13] it was established that for a *nonuniform power* setting, the reception zones are not necessarily connected, but are (perhaps surprisingly) hyperbolically convex in a space with dimension higher by one than the network's dimension.

When adding SIC to SINR diagrams, the resulting structures, denoted *SIC-SINR diagrams*, become much more complex to present. However they can reveal the benefits of the cancellation method. An example of this idea is illustrated in Figure 1. In all three parts of the figure we have a network with two transmitters s_1, s_2 and two receivers (or points in the plane) r_1, r_2, with the requirement that r_1 needs to decode the signal transmitted by s_1 and r_2 needs to receive the signal of s_2. All four nodes are ordered on a line with a special configuration similar to the known example given in [16]. This example, known as "nested links", shows that in order to achieve the requirements, a nonuniform power assignment must be used by the two transmitters, thus demonstrating that the capacity (achievable rate) of nonuniform power assignments is higher than that of uniform power assignments. Indeed, Figure 1(a) shows zones $\mathcal{H}(s_1)$ and $\mathcal{H}(s_2)$ for s_1 and s_2 respectively, which satisfy $r_1 \in \mathcal{H}(s_1)$ and $r_2 \in \mathcal{H}(s_2)$. As mentioned, it can be proved that these two demands cannot be satisfied when both s_1 and s_2 transmit with the same power. An SINR diagram with a uniform power assignment is shown in Figure 1(b). Note that here, $r_1 \notin \mathcal{H}(s_1)$, but $r_1 \in \mathcal{H}(s_2)$. In contrast, when SIC is possible at r_1, it can first decode s_2. Afterward it "cancels" s_2 from its interference and then decode s_1. Therefore, with SIC the two demands can be satisifed once again, even with uniform powers! The SIC-SINR diagram presented in Figure 1(c) illustrates this by showing an additional zone, $\mathcal{H}(s_2, s_1)$, the zone in which stations with SIC can decode s_1 after "canceling" s_2. Note that, as explained later, $\mathcal{H}(s_2, s_1)$ is the intersection of two convex shapes, $\mathcal{H}(s_2)$ and $\mathcal{H}^*(s_1)$, where the latter (shown as an empty circle) is the reception zone of s_1 if it had transmitted alone in the network. One clearly sees that the total reception area of s_1 with SIC is much larger than without SIC. In Subsection 1.3 we present an even more compelling motivating example, that shows the following.

Observation 1.1. *There exists a wireless network for which any power assignment requires n time slots to satisfy all the demands, while using SIC allows a satisfying schedule using a* single *time slot.*

Despite the importance of IC, not much is known about its complexity. The goal of this paper is to take a first step towards understanding it, by studying reception maps under the setting of SIC. The starting point of our work is the observation that under the SIC setting, reception zones are no longer guaranteed to be convex, fat or even connected. This holds true even for the most "simplified" settings where stations transmit at the same power level and are aligned on a straight line (one dimensional map). The zones are also not hyperbolically convex as was shown in the nonuniform power settings without IC [13]. Moreover, while for SINR diagrams without IC there is a single polynomial that represents each of their reception zones, with IC, the reception zone of each transmitter may depend on the cancellation order, which can lead to an exponential number of polynomials and cells. If this is the case, then even drawing the diagram might prove to be infeasible.

1.2 Our Contributions The study of SIC-SINR reception maps raises several immediate questions. The first is a simple "counting" question that has strong implications on our algorithmic question: What is the maximum number of reception cells that may occur in an SINR diagram of a wireless network with n stations where every point in the map is allowed to perform SIC? Is it indeed exponential? We address this question in two different ways. Initially we re-explore the

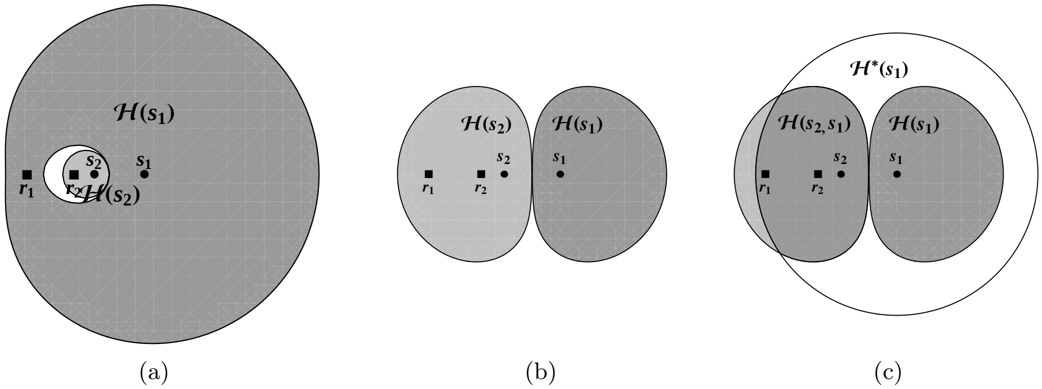

Figure 1: (a) **Nonuniform** power assignment where for $i = 1, 2$, receiver r_i is in the reception zone $\mathcal{H}(s_i)$ of transmitter s_i. (b) In every **uniform** power assignment, r_1 is not in the reception zone of s_1. (c) In a **uniform** power setting but with interference cancellation, station r_1 may be in the reception zone $\mathcal{H}(s_2, s_1)$ of transmitter s_1, i.e., it could decode transmitter s_1 after canceling transmitter s_2.

intimate connections between wireless communication and computational geometry methods like higher-order Voronoi diagrams [20, 17]. In particular we use a bound on the number of cells in *ordered order-k Voronoi diagram* [17] to upper bound the number of reception zones by $O(n^{2d})$, where n is the number of transmitters and d is the dimension. In the general case this bound is not tight, but interestingly we were able to tie the number of reception zones to a novel parameter of the network, named the *Compactness Parameter*, and achieve a much tighter bound when the compactness parameter is high enough. The compactness parameter is a function of the two most important parameters of the wireless network model, the *reception threshold* constant $\beta \geq 1$, and the *path-loss parameter* $\alpha > 0$, and its value is $\mathcal{CP} = \sqrt[\alpha]{\beta}$. We then prove that when $\mathcal{CP} \geq 5$, the number of reception zones is linear for any dimension! These bounds allow us to provide an efficient scheme for computing the cancellation order that gives the reception zones and therefore allows us to build and represent the diagram efficiently.

The second question has a broader scope: Are there any "nice" properties of reception zones that can be established in the SIC settings? Specifically, we aim toward finding forms of convexity satisfied by reception cells in SIC reception maps. Apart from their theoretical interest, these questions also have considerable practical significance, since having reception zones with some form of convexity might ease the development of protocols for various design and communication tasks. We answer this question by using the key observation that zones are intersections of convex shapes giving us some "nice" geometric guarantees.

The third question is of algorithmic nature. We consider the point location task, where given a point $p \in \mathbb{R}^d$ and a station s_i, one wants to know whether p can receive s_i using SIC. Applying the trivial computation in $O(n \log n)$ time, one can compute the set of stations that p receives under the SIC setting. However, if the number of queries is large, an order of $O(n \log n)$ per query might be too costly. To approach this problem we use the guarantees of the first two questions and present a scheme for answering point location queries in logarithmic time.

We believe that the questions raised herein, as well as the results and techniques developed, can make a major contribution to the evolving topic of wireless topology and what we may refer to as computational wireless geometry.

The rest of the paper is organized as follows. In Section 3, we establish the basic properties of SIC-SINR diagrams and show its relation to the higher-order Voronoi diagrams. We then derive a tight bound on the number of connected components in the reception map of a given station under SIC. Section 4 describes how one can construct SIC-SINR reception maps in polynomial time. Finally, Section 5 considers the point-location task and provides an efficient construction of a data structure that answers point-location queries (with predefined approximation guarantees) in logarithmic time.

1.3 Motivating Example: Interference Cancellation vs. Power Control The following motivating example illustrates the power of interference cancellation even in a setting where all stations use the same power. Consider a set of n communication requests $L = \{(s_i, r_i) \mid i = 1, ..., n\}$ consisting of transmitter-receiver pairs embedded on a real line as follows. The n receivers are located at the origin and the set of transmitters are positioned on an exponential node-chain, e.g., s_i is positioned on $x_i = 2^{i/\alpha}$, see Fig 2. Since all receivers share the same position, without SIC, there exists no power assignment that can satisfy more than one request simultaneously, hence n time slots are necessary for satisfying all the requests. We claim that by using SIC, all requests can be satisfied in a single time slot even with a uniform power assignment. We focus on a given receiver r_j and show that it successfully decodes the signal from s_j after successive cancellations of the signal $\mathcal{S}_i$ for every $i < j$. Using the notation of Section 2.2, let $\mathcal{A}_i = \langle d, S_i = \{s_i, \ldots, s_n\}, \bar{1}, N \leq 1/2^n, \beta = 1, \alpha \rangle$ denote the network imposed on the last $n - i$ stations, whose positions are $2^{(n-i)/\alpha}$ to $2^{n/\alpha}$. Note that

$$\mathrm{SINR}_{\mathcal{A}_i}(s_i, r_j) \ = \ \frac{1/2^i}{\sum_{j=i+1}^{n} 1/2^j + N} \geq 1$$

for every $i \leq n$. We therefore establish that there exists an instance $L = \{(s_i, r_i) \mid i = 1, ..., n\}$ such that any power assignment for scheduling L requires n slots, whereas using SIC allows a satisfying schedule using a *single* time slot.

So far, the literature on capacity and scheduling addressed mostly nonuniform powers, showing that nonuniform power assignments can outperform a uniform assignment [16, 15] and increase the capacity of a network. In contrast, examples such as Fig. 2 illustrate the power of interference cancellation even with uniform power assignments, and motivate the study of this technique from an algorithmic point of view. Understanding SINR diagrams with SIC may play a role in the development of suitable algorithms (e.g. capacity, scheduling and power control), filling the current gap between the electrical engineering and algorithmic communities with respect to SIC research.

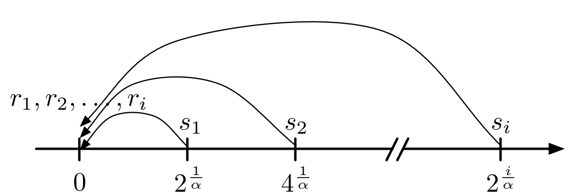

Figure 2: The power of interference cancellation. Any power assignment requires n time slots to schedule these n requests. Using SIC, all requests are satisfied in a single time slot (even when using *uniform* powers).

2 Preliminaries

2.1 Geometric notions We consider the d-dimensional Euclidean space $\mathbb{R}^d$ (for $d \in \mathbb{Z}_{\geq 1}$). The *distance* between points p and q is denoted by $\mathrm{dist}(p, q) = \|q - p\|$. A *ball* of radius r centered at point $p \in \mathbb{R}^d$ is the set of all points at distance at most r from p, denoted by $B^d(p, r) = \{q \in \mathbb{R}^d \mid \mathrm{dist}(p, q) \leq r\}$. Unless stated otherwise, we assume the 2-dimensional Euclidean plane, and omit d. The maximal distance between a point p and set of points Q is defined as $\max(\mathrm{dist}(p, Q)) = \max_{q \in Q}\{\mathrm{dist}(p, q)\}$. Analogously, the minimal distance between p and Q is defined as $\min(\mathrm{dist}(p, Q)) = \min_{q \in Q}\{\mathrm{dist}(p, q)\}$. The hyperplane $HP(q_i, q_j)$, for $q_i, q_j \in \mathbb{R}^d$, is defined by $HP(q_i, q_j) = \{p \in \mathbb{R}^d \mid \mathrm{dist}(p, q_i) = \mathrm{dist}(p, q_j)\}$. Given a set of n points $Q = \{q_i \in \mathbb{R}^d\}$, let the corresponding set of all $\binom{n}{2}$ hyperplanes be $HP(Q) = \{HP(q_i, q_j) \mid q_i, q_j \in Q\}$. A finite set Υ of hyperplanes defines a dissection of $\mathbb{R}^d$ into connected pieces of various dimensions, known as the *arrangement* $Ar(\Upsilon)$ of Υ. The basic notions of open, closed, bounded, compact and connected sets of points are defined in the standard manner (see [3]).

We use the term *zone* to describe a point set with some "nice" properties. Unless stated otherwise, a zone refers to the union of an open connected set and some subset of its boundary. It may also refer to a single point or to the finite union of zones. A polynomial $F : \mathbb{R}^d \to \mathbb{R}$ is the *characteristic polynomial* of a zone Z if $p \in Z \Leftrightarrow F(p) \leq 0$ for every $p \in \mathbb{R}^d$.

Denote the *area* of a bounded zone Z (assuming it is well-defined) by $\mathrm{area}(Z)$. A nonempty bounded zone $Z \neq \emptyset$ is *fat* if the ratio between the radii of the smallest circumscribed and largest inscribed circles with respect to Z is bounded by a constant.

2.2 Wireless Networks and SINR We consider a wireless network $\mathcal{A} = \langle d, S, \psi = \bar{1}, N, \beta > 1, \alpha \rangle$, where $d \in \mathbb{Z}_{\geq 1}$, $S = \{s_1, s_2, \ldots, s_n\}$ is a set of transmitting *radio stations* embedded in d-dimensional space, ψ is a mapping assigning a positive real *transmitting power* ψ_i to each station s_i, $N \geq 0$ is the *background noise*, $\beta > 1$ is a constant serving as the *reception threshold* (to be explained soon), and $\alpha > 0$ is the *path-loss parameter*. The *signal to interference & noise ratio (SINR)* of s_i at point p is defined as

$$\mathrm{SINR}_{\mathcal{A}}(s_i, p) = \frac{\psi_i \cdot \mathrm{dist}(s_i, p)^{-\alpha}}{\sum_{j \neq i} \psi_j \cdot \mathrm{dist}(s_j, p)^{-\alpha} + N}.$$

When the network $\mathcal{A}$ is clear from the context, we may omit it and write simply $\mathrm{SINR}(s_i, p)$. Let $\mathcal{A}_{d'}$ be a network identical to $\mathcal{A}$ except its dimension is $d' \neq d$. In our arguments, we sometimes refer to an ordered subset of stations, $\overrightarrow{S_i} = (s_{i_1}, \ldots, s_{i_k}) \subseteq S$. Denote the last element in $\overrightarrow{S_i}$ by $\mathsf{Last}(\overrightarrow{S_i})$. When the order is insignificant, we refer to this set as simply $S_i = \{s_{i_1}, \ldots, s_{i_k}\}$. The wireless network restricted to a subset of nodes S_i is given by $\mathcal{A}(S_i) = \langle d, S_i, \psi, N, \beta, \alpha \rangle$. The network is assumed to contain at least two stations, i.e., $n \geq 2$.

2.3 SINR diagrams (without SIC) The fundamental rule of the SINR model is that the transmission of station s_i is received correctly at point $p \notin S$ if and only if its SINR at p reaches or exceeds the reception threshold of the network, i.e.,

$$\mathrm{SINR}(s_i, p) \geq \beta.$$

When this happens, we say that s_i is *heard* at p. We refer to the set of points that hear station s_i as the *reception zone* of s_i, defined as

$$\mathcal{H}_{\mathcal{A}}(s_i) \ = \ \{p \in \mathbb{R}^d - S \mid \mathrm{SINR}_{\mathcal{A}}(s_i, p) \geq \beta\} \cup \{s_i\} \ .$$

(Note that $\mathrm{SINR}(s_i, \cdot)$ is undefined at points in S and in particular at s_i itself.) Analogously, the set of points that hear no station $s_i \in S$ (due to the background noise and interference) is defined as

$$\mathcal{H}_{\mathcal{A}}(\emptyset) \ = \ \{p \in \mathbb{R}^d - S \mid \mathrm{SINR}(s_i, p) < \beta, \ \ \forall s_i \in S\}.$$

An SINR diagram

$$\mathcal{H}(\mathcal{A}) = \left(\bigcup_{s_i \in S} \mathcal{H}_{\mathcal{A}}(s_i)\right) \cup \mathcal{H}_{\mathcal{A}}(\emptyset)$$

is a "reception map" characterizing the reception zones of the stations. This map partitions the plane into $n+1$ zones; one for each station $\mathcal{H}_{\mathcal{A}}(s_i)$, $1 \leq i \leq n$, and the zone $\mathcal{H}_{\mathcal{A}}(\emptyset)$ where none of the stations is received. It is important to note that a reception zone $\mathcal{H}_{\mathcal{A}}(s_i)$ is not necessarily connected. A *maximal connected component* within a zone is referred to as a *cell*. Let $\mathcal{H}_{\mathcal{A}}(s_i, j)$ be the j^{th} cell in $\mathcal{H}_{\mathcal{A}}(s_i)$. Hereafter, the set of points where the transmissions of a given station are successfully received is referred to as its *reception zone*. Hence the reception zone is a set of cells, given by

$$\mathcal{H}_{\mathcal{A}}(s_i) = \{\mathcal{H}_{\mathcal{A}}(s_i, 1), \ldots \mathcal{H}_{\mathcal{A}}(s_i, \tau_i)\},$$

where $\tau_i = \tau_i(\mathcal{A})$ is the number of cells in $\mathcal{H}_{\mathcal{A}}(s_i)$. Analogously, $\mathcal{H}_{\mathcal{A}}(\emptyset)$ is composed of $\tau_{\emptyset}(\mathcal{A})$ connected cells $\mathcal{H}_{\mathcal{A}}(\emptyset, j)$. Overall, the topology of a wireless network $\mathcal{A}$ is arranged in three levels: The *reception map* is at the top of the hierarchy. It is composed of n reception zones, $\mathcal{H}_{\mathcal{A}}(s_i)$, $s_i \in S$ and $\mathcal{H}_{\mathcal{A}}(\emptyset)$. Each zone $\mathcal{H}_{\mathcal{A}}(s_i)$ is composed of $\tau_i(\mathcal{A})$ reception cells. The following is from [3].

Lemma 2.1. ([3]) *Let $\mathcal{A} = \langle d, S, \psi = \bar{1}, N, \beta > 1, \alpha \rangle$ be a uniform power network. Then $\mathcal{H}_{\mathcal{A}}(s_i)$ is convex and fat for every $s_i \in S$.*

We sometimes refer to the wireless network $\mathcal{A}$ induced on a subset of stations $S_j \subseteq S$. The reception zone of s_i in this induced network is denoted by $\mathcal{H}_{\mathcal{A}}(s_i \mid S_j)$. When $\mathcal{A}$ is clear from context, we may omit it and write $\mathcal{H}(s_i)$ and $\mathcal{H}(s_i \mid S_j)$.

2.4 Geometric diagrams in $\mathbb{R}^d$ Throughout the paper we make use of the following diagrams.

Hyperplane Arrangements. Given a set of Υ of n hyperplanes in $\mathbb{R}^d$, the arrangement $Ar(\Upsilon)$ of Υ dissects $\mathbb{R}^d$ into connected pieces of various dimensions. Let $\tau^A(\Upsilon)$ denote the number of connected components in $Ar(\Upsilon)$. The following facts about $Ar(\Upsilon)$ are taken from [8].

Lemma 2.2. ([8]) *(a) $\tau^A(\Upsilon) = \Theta(n^d)$.*
(b) $Ar(\Upsilon)$ can be constructed in $\Theta(n^d)$ time and maintained in $\Theta(n^d)$ space.

Given a set of n points $S \subset \mathbb{R}^d$, we define $Ar(S)$ to be the arrangement on $HP(S) = \{HP(s_i, s_j) \mid s_i, s_j \in S\}$, the set of all $\binom{n}{2}$ hyperplanes of pairs in S. $Ar(S)$ has an important role in constructing $SINR - SIC$ maps, as will be described later on.

Corollary 2.1. $\tau^A(S) = \Theta(n^{2d})$.

Voronoi diagrams. The ordinary Voronoi diagram on a given set of points S tessellates the space in such a way that every location in the space is assigned to the closest point in S, thus partitioning the space into *cells*, each consisting of the set of locations closest to one point in S (referred to as the cell's *generator*). This forms the ordinary Voronoi diagram, consisting of Voronoi edges and vertices. Let $\mathrm{Vor}(s_i)$ denote the Voronoi cell of s_i given a set of generators S. Let $\mathrm{Vor}(s_i \mid S_j)$, for $S_j \subseteq S$, denote the Voronoi cell of s_i in a system restricted to the points of S_j.

Avin et al. [3] discuss the relationships between the SINR diagram on a set of stations S with *uniform* powers and the corresponding *Voronoi diagram* on S. Specifically, it is shown that the n reception zones $\mathcal{H}_{\mathcal{A}}(s_i)$ are strictly contained in the corresponding Voronoi cells $\mathrm{Vor}(s_i)$. The following lemma summarizes these relations. Let $\mathcal{A} = \langle d, S, \bar{1}, N, \beta \geq 1, \alpha \rangle$.

Lemma 2.3. ([3]) *$\mathcal{H}_{\mathcal{A}}(s_i) \subseteq \mathrm{Vor}(s_i)$ for every $s_i \in S$.*

Higher order Voronoi diagrams. Higher order Voronoi diagrams are a natural extension of the ordinary Voronoi diagram, where cells are generated by more than one point. They provide tessellations where each region consists of the locations having the same k (ordered or unordered) closest points in S, for some given integer k.

Order-k Voronoi diagram. The order-k Voronoi diagram $\mathbb{V}^{(k)}(S)$ is a set of all non-empty order-k Voronoi regions $\mathbb{V}^{(k)}(S) = \{\textsc{Vor}(S_1^{(k)}), \ldots, \textsc{Vor}(S_m^{(k)})\}$, where the order-$k$ Voronoi zone $\textsc{Vor}(S_i^{(k)})$ for an unordered subset $S_i^{(k)} \subseteq S, |S_i^{(k)}| = k$ is defined as follows.

$$\textsc{Vor}(S_i^{(k)}) = \{p \in \mathbb{R}^d \mid \max(\text{dist}(p, S_i^{(k)})) \leq \min(\text{dist}(p, S \setminus S_i^{(k)}))\} \ .$$

This can alternatively be written as

$$\textsc{Vor}(S_i^{(k)}) = \bigcap_{s \in S_i^{(k)}} \textsc{Vor}(s \mid S \setminus S_i^{(k)} \cup s) \ . \tag{2.1}$$

Note that $\mathbb{V}^{(1)}(S)$ corresponds to the ordinary Voronoi diagram and that any $\mathbb{V}^{(k)}(S)$, for $k > 1$, is a refinement of $\mathbb{V}^{(1)}(S)$.

Ordered Order-k Voronoi diagram. Let $\overrightarrow{S_i} \subseteq S$ be an ordered set of k elements from S. When the k generators are ordered, the diagram becomes the *ordered order-k Voronoi diagram* $\mathcal{V}^{\langle k \rangle}(S)$ [17], defined as

$$\mathcal{V}^{\langle k \rangle}(S) = \{\textsc{Vor}(\overrightarrow{S_i})\},$$

where the ordered order-k Voronoi region $\textsc{Vor}(\overrightarrow{S_i})$, $|\overrightarrow{S_i}| = k$, is defined as

$$\textsc{Vor}(\overrightarrow{S_i}) = \{p \in \mathbb{R}^d \mid \text{dist}(p, s_{i_1}) \leq \text{dist}(p, s_{i_2}) \leq \ldots \leq \text{dist}(p, s_{i_k}) \leq \min(\text{dist}(p, S \setminus S_i))\}.$$

Alternatively, as in [17],

$$\textsc{Vor}(\overrightarrow{S_i}) = \bigcap_{j=1}^{k} \textsc{Vor}(s_{i_j} \mid S \setminus \{s_{i_1}, \ldots, s_{i_{j-1}}\}) \ . \tag{2.2}$$

Note that each $\textsc{Vor}(\overrightarrow{S_i})$ is an intersection of k convex shapes and hence it is convex as well.
The following claim is useful for our later arguments.

Lemma 2.4. *For every $\overrightarrow{S_i}, \overrightarrow{S_j} \subseteq S$ such that $\overrightarrow{S_i} \not\subseteq \overrightarrow{S_j}$ there exist $s_{k_1}, s_{k_2} \in S$, such that the hyperplane $HP(s_{k_1}, s_{k_2})$ separates $\textsc{Vor}(\overrightarrow{S_i})$ and $\textsc{Vor}(\overrightarrow{S_j})$.*

Proof. We focus on the case where $\textsc{Vor}(\overrightarrow{S_i}) \neq \emptyset$ and $\textsc{Vor}(\overrightarrow{S_j}) \neq \emptyset$. Let m denote the first index such that $s_{i_m} \neq s_{j_m}$. First consider the case where $m = 1$. Then $\textsc{Vor}(\overrightarrow{S_i}) \subseteq \textsc{Vor}(s_{i_1})$ and $\textsc{Vor}(\overrightarrow{S_j}) \subseteq \textsc{Vor}(s_{j_1})$, so $HP(s_{i_1}, s_{j_1})$ separates the zones and the lemma holds. Otherwise, assume $m > 1$ and let $\overrightarrow{S^*} = \{s_1^*, \ldots, s_{m-1}^*\}$ denote the longest common prefix of $\overrightarrow{S_i}$ and $\overrightarrow{S_j}$. Let $X_0 = \textsc{Vor}(\overrightarrow{S^*})$, $X_1 = \textsc{Vor}(s_{i_m} \mid S \setminus S^*)$ and $X_2 = \textsc{Vor}(s_{j_m} \mid S \setminus S^*)$. First note that by Eq. (2.2), $\textsc{Vor}(\overrightarrow{S_i}) \subseteq X_0 \cap X_1$ and $\textsc{Vor}(\overrightarrow{S_j}) \subseteq X_0 \cap X_2$. In addition, X_0, X_1 and X_2 are convex. Next, observe that X_1, X_2 correspond to distinct Voronoi regions in the system of points $S \setminus S^*$ and therefore X_1 and X_2 are separated by $HP(s_{i_m}, s_{j_m})$. The lemma follows. ☐

3 SIC-SINR Diagrams in Uniform Power Networks

In this section, we first formally define the reception zones under interference cancellation, forming the *SIC-SINR Diagrams*. We then take a first step towards studying the properties of these diagrams. We elaborate on the relation between the SIC-SINR diagram and the *ordered order-k Voronoi diagram*, and use it to prove convexity properties of the diagram and to bound the number of connected components in the SIC-SINR Diagrams. We then define the *Compactness Parameter* of the diagrams and use it to achieve tighter bounds on the number of connected components.

3.1 SINR diagrams with SIC Let $\mathcal{A} = \langle d, S, \psi = \bar{1}, N, \beta > 1, \alpha \rangle$. We now focus on the reception zone of a single station, say s_1, under the setting of interference cancellation. In other words, we are interested in the area containing all points that can decode s_1 after successive cancellations. To warm up, we start with the case of a single point $p \in \mathbb{R}^d$ and ask the following question: does p successfully receive s_1 using SIC?

Let $\overrightarrow{S_p} = \{s_{p_1}, \ldots, s_{p_k}\}$ correspond to the set of stations S ordered in nonincreasing order of received signal strength at point p up to station s_1, i.e., $\mathrm{E}_{\mathcal{A}}(s_{p_1}, p) \geq \mathrm{E}_{\mathcal{A}}(s_{p_2}, p) \geq \cdots \geq \mathrm{E}_{\mathcal{A}}(s_{p_k}, p)$, where $s_{p_k} = s_1$. To receive s_1 correctly, p must cancel the signals $\mathcal{S}_{p_i}$ of station s_{p_i}, for $i < k$, in a successive manner. It therefore follows that p successfully receives s_1 following SIC iff

$$p \in \mathcal{H}(s_{p_i} \mid S \setminus \{s_{p_1}, \ldots, s_{p_{(i-1)}}\}), \tag{3.3}$$

for every $i \leq k$. The reception zone of s_1 in a wireless network $\mathcal{A}$ under the setting of SIC is denoted by $\mathcal{H}_{\mathcal{A}}^{SIC}(s_1)$, or simply $\mathcal{H}^{SIC}(s_1)$ when $\mathcal{A}$ is clear from the context. It contains s_1 and the set of points p obeying Equation (3.3), i.e.,

$$\mathcal{H}^{SIC}(s_1) = \{p \in \mathbb{R}^d - S \mid p \text{ satisfies Eq. (3.3)}\}.$$

We now provide a more constructive formulation for $\mathcal{H}^{SIC}(s_1)$, which becomes useful in our later arguments. Let $\overrightarrow{S_i} \subseteq S$ be an ordering of k stations. Let $\mathcal{H}(\overrightarrow{S_i})$ denote the reception area of all points that receive $\mathsf{Last}(\overrightarrow{S_i})$ correctly after successive cancellation of $s_{i_1}, \ldots, s_{i_{k-1}}$. Formally, the zone $\mathcal{H}(\overrightarrow{S_i})$ is defined in an inductive manner with respect to the length of the ordering $\overrightarrow{S_i}$, i.e., number of cancellations minus one. For $\overrightarrow{S_i} = \{s_j\}$, $\mathcal{H}(\overrightarrow{S_i}) = \mathcal{H}(s_j)$. Otherwise, for $k > 1$,

$$\mathcal{H}(\overrightarrow{S_i}) \ = \ \mathcal{H}(\overrightarrow{S_i} \setminus s_{i_k}) \cap \mathcal{H}(s_{i_k} \mid (S \setminus S_i) \cup s_{i_k}) \ ,$$

or,

$$\mathcal{H}(\overrightarrow{S_i}) = \bigcap_{j=1}^{k} \mathcal{H}(s_{i_j} \mid S \setminus \{s_{i_1}, \ldots, s_{i_{(j-1)}}\}) \ . \tag{3.4}$$

The following is a direct consequence of Eq. (3.4).

Corollary 3.1. *Let $\overrightarrow{S_i} \subseteq S$, $|\overrightarrow{S_i}| = k$. Then*

$$\begin{aligned} \mathcal{H}(\overrightarrow{S_i}) \ &\subseteq \ \mathcal{H}(s_{i_1}, \ldots, s_{i_{(k-1)}}) \subseteq \mathcal{H}(s_{i_1}, \ldots, s_{i_{(k-2)}}) \\ &\subseteq \ \ldots \subseteq \mathcal{H}(s_{i_1}) \ . \end{aligned}$$

Finally the reception zone of s_1 under SIC is given as follows. Let $\mathcal{CO}_j$ denote the collection of all cancellation orderings ending with s_j, namely,

$$\mathcal{CO}_j = \{\overrightarrow{S_i} \subseteq S \mid \mathsf{Last}(\overrightarrow{S_i}) = s_j\}.$$

Then

$$\mathcal{H}^{SIC}(s_1) \ = \ \bigcup_{\overrightarrow{S_i} \in \mathcal{CO}_1} \mathcal{H}(\overrightarrow{S_i}) \ . \tag{3.5}$$

The reception zone

$$\mathcal{H}^{SIC}(s_1) = \{\mathcal{H}^{SIC}(s_1, 1), \ldots \mathcal{H}^{SIC}(s_1, \tau_1^{SIC})\}$$

is a set of τ_1^{SIC} cells. Note that the region of unsuccessful reception to any of the points, namely, $\mathcal{H}(\emptyset)$, is unaffected by SIC. This follows by noting that SIC only affects the set of points $p \in \bigcup \mathcal{H}(s_i) \setminus S$. In other words, the successive signal cancellation allows points $p \in \mathbb{R}^d \setminus S$ to "migrate" from the reception zone of station s_i to that of station s_j. However, points that hear nobody can cancel none of the signals. Overall, the topology of a wireless network $\mathcal{A}$ under SIC is arranged in three levels: The *reception map* is at the top of the hierarchy. It is composed, at the next level, of n reception zones, $\mathcal{H}^{SIC}(s_i)$, $s_i \in S$ and $\mathcal{H}(\emptyset)$. Finally, at the lowest level, each zone $\mathcal{H}^{SIC}(s_i)$ is composed of τ_i^{SIC} reception cells.

Throughout the paper we consider a uniform power network of the form $\mathcal{A} = \langle d, S, \psi = \bar{1}, N, \beta > 1, \alpha \rangle$. Avin et al. [3] established that reception zones of uniform power maps are fat and convex. However, once signal cancellation enters the picture, the convexity (and connectivity) of zones is lost even for the simple case where stations are aligned on a line; see Figure 3 for an illustration of the SIC-SINR map of a 3-station system. In this section, we derive a bound on the number of connected cells in the zone $\mathcal{H}^{SIC}(s_1)$ and show that each of these cells is convex. In addition, we establish a relation between SIC-SINR diagrams and a generalized form of Voronoi diagrams.

3.2 Higher-order Voronoi diagrams and SIC-SINR maps To understand the structure and the topological properties of SIC-SINR reception maps, we begin our study by describing the relation between SIC-SINR reception maps and ordered order-k Voronoi diagram. Specifically, we prove that every SIC-SINR zone is composed of a collection of convex cells, each of which is related to a cell of the higher-order Voronoi diagram. To avoid complications, we assume our stations are embedded in *general positions*.

Toward the end of this section, we establish the following.

Lemma 3.1. *For every two reception cells $\mathcal{H}^{SIC}(s_1, i)$ and $\mathcal{H}^{SIC}(s_1, j)$, there are distinct orderings $\overrightarrow{S_i}, \overrightarrow{S_j} \in \mathcal{CO}_1$ such that $\mathcal{H}^{SIC}(s_1, i) \subseteq \textsc{Vor}(\overrightarrow{S_i})$ and $\mathcal{H}^{SIC}(s_1, j) \subseteq \textsc{Vor}(\overrightarrow{S_j})$.*

For illustration of these relations, see Figures 3 and 4.

We begin by describing the relation between a nonempty reception region $\mathcal{H}(\overrightarrow{S_i})$ and an nonempty ordered order-k polygon.

Lemma 3.2. $\mathcal{H}(\overrightarrow{S_i}) \subseteq \textsc{Vor}(\overrightarrow{S_i})$, *for $\beta \geq 1$.*

Proof. By Lemma 2.3, $\mathcal{H}(s_j \mid S') \subseteq \textsc{Vor}(s_j \mid S')$. Therefore by Eq. (3.4) it follows that

$$\begin{aligned} \mathcal{H}(\overrightarrow{S_i}) \ &\subseteq \ \bigcap_{i=1}^{k} \textsc{Vor}(s_{i_j} \mid (S \setminus \{s_{i_1}, \ldots, s_{i_{j-1}}\}) \\ &= \ \textsc{Vor}(\overrightarrow{S_i}) \ , \end{aligned}$$

where the last inequality follows by Eq. (2.2). □

We now show that reception regions in $\mathcal{H}^{SIC}(s_1)$ that result from different cancellation orderings correspond to distinct connected cells, which will establish Lemma 3.1.

Lemma 3.3. *Every two regions $\mathcal{H}(\overrightarrow{S_1}), \mathcal{H}(\overrightarrow{S_2}) \subseteq \mathcal{H}^{SIC}(s_1)$ correspond to two distinct cells.*

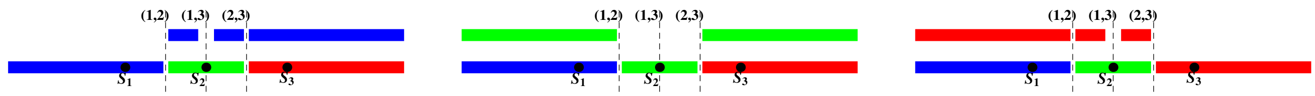

Figure 3: SIC-SINR reception map in $\mathbb{R}^1$ for a 3-station network aligned on a line. The first line of colored segments corresponds to reception zones with no cancellations. The second line of segments in each figure represents the added reception cells by signal cancellation. $\mathcal{H}^{SIC}_{\mathcal{A}_{d=1}}(s_1)$, $\mathcal{H}^{SIC}_{\mathcal{A}_{d=1}}(s_2)$ and $\mathcal{H}^{SIC}_{\mathcal{A}_{d=1}}(s_3)$ are in middle, light and dark grey respectively.

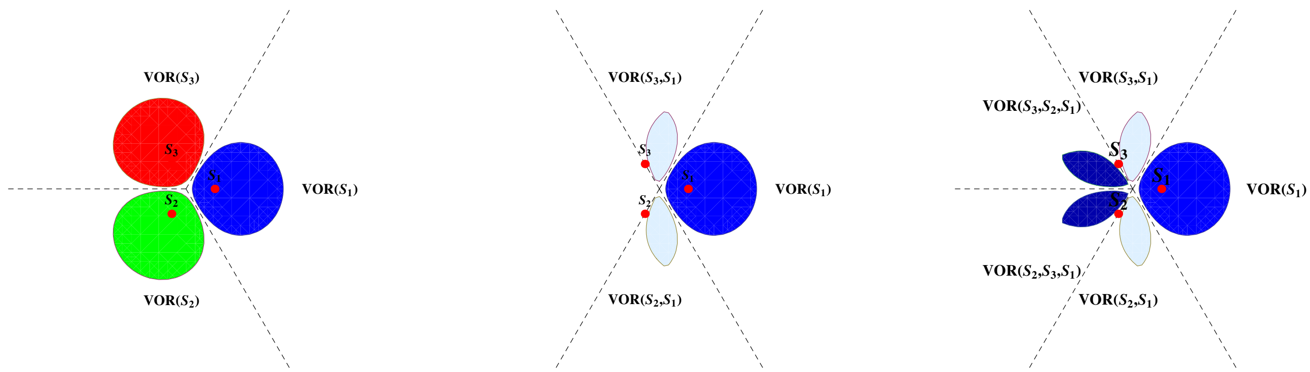

Figure 4: Reception map of $\mathcal{H}^{SIC}_{\mathcal{A}_{d=2}}(s_1)$ and ordered order-k Voronoi diagram (for $k \in [1,3]$). (a) SINR map with no cancellation. Shown are $\mathcal{H}(s_1)$, $\mathcal{H}(s_2)$ and $\mathcal{H}(s_3)$. (b) Intermediate map of $\mathcal{H}^{SIC}_{\mathcal{A}_{d=2}}(s_1)$: Reception cells of s_1 following at most one cancellation. (c) Final map of $\mathcal{H}^{SIC}_{\mathcal{A}_{d=2}}(s_1)$: Reception cells of s_1 following two cancellations.

Proof. By Eq. (3.5), $\mathcal{H}^{SIC}(s_1)$ is the union of $\mathcal{H}(\overrightarrow{S_i})$ regions for $\overrightarrow{S_i} \in \mathcal{CO}_1$, i.e., where $\mathsf{Last}(\overrightarrow{S_i}) = s_1$. By Lemma 3.2, $\mathcal{H}(\overrightarrow{S_1}) \subseteq \textsc{Vor}(\overrightarrow{S_1})$ and $\mathcal{H}(\overrightarrow{S_2}) \subseteq \textsc{Vor}(\overrightarrow{S_2})$. Due to Claim 2.4, $\textsc{Vor}(\overrightarrow{S_1}) \cap \textsc{Vor}(\overrightarrow{S_2}) = \emptyset$ and hence also $\mathcal{H}(\overrightarrow{S_1}) \cap \mathcal{H}(\overrightarrow{S_2}) = \emptyset$. The lemma follows. □

Next, this relation between the SIC-SINR reception maps and ordered order-k Voronoi diagram is used to establish the convexity of cells and to bound the number of connected components in the zone.

Convexity of SIC-SINR cells. We first show that the reception cells of $\mathcal{H}^{SIC}(s_1)$ are convex.

LEMMA 3.4. *Every reception cell* $\mathcal{H}^{SIC}(s_1, i), i \in [1, \tau_i^{SIC}]$ *of* $\mathcal{H}^{SIC}(s_1)$ *is convex.*

Proof. Due to Lemma 3.3 it is enough to show that every nonempty $\mathcal{H}(\overrightarrow{S_i})$ is convex. Let $k = |S_i|$ and let $X_j = \mathcal{H}(s_{i_j} \mid S \setminus \{s_{i_1}, \ldots, s_{i_{(j-1)}}\})$. By Lemma 2.1, X_j is convex for every $j \leq k$ and therefore by Eq. (3.4), $\mathcal{H}(\overrightarrow{S_i})$ is an intersection of k convex and bounded shapes, hence it is convex (and bounded) as well.

Number of connected components. In this subsection, we discuss the number of connected components in SIC-SINR diagrams. Toward the end of this section, we exploit the relation between SIC-SINR diagrams and high-order Voronoi diagrams to establish the following.

LEMMA 3.5. $\tau_i^{SIC} = O(n^{2d})$, *for every* $s_i \in S$.

Without loss of generality we focus on station s_1. By Lemma 3.3, every two distinct orderings correspond to distinct reception cells (though it might be empty). Since the number of distinct orderings of length $n-1$ is $(n-1)!$ (i.e., the size of $\mathcal{CO}_1$) and each of those orderings might correspond to a distinct cell, it follows that the number of connected cells in $\mathcal{H}^{SIC}(s_1)$ might be exponential. Fortunately, the situation is much better due to Lemma 3.1. An ordering $\overrightarrow{S_i}$ is defined as a *nonempty cancellation ordering* ($\mathcal{NCO}$) if and only if $\textsc{Vor}(\overrightarrow{S_i})$ is nonempty. Partition the collection of $\mathcal{NCO}$'s into sets $\mathcal{NCO}_1, \ldots, \mathcal{NCO}_n$ as follows. An ordering $\overrightarrow{S_i}$ is in the set $\mathcal{NCO}_j$ if and only if it is an $\mathcal{NCO}$ and in addition $\mathsf{Last}(\overrightarrow{S_i}) = s_j$. The following lemma shows that there are only polynomially many orderings in $\mathcal{NCO}_j$.

LEMMA 3.6. *(a)* $|\mathcal{NCO}_1| = O(n^{2d})$.
(b) $|\mathcal{NCO}_1| = \Omega(n^2)$ *for* $d = 1$.

Proof. To derive an upper bound on $\mathcal{NCO}_1$, proving claim (a), we use the diagram of $Ar(S)$. We claim that any given cell $f \in Ar(S)$ intersects with at most one high-order Voronoi region $\text{Vor}(\overrightarrow{S_i})$, where $\overrightarrow{S}_i \in \mathcal{NCO}_1$. (The proofs of the following two claims will appear in the journal version of the paper.)

Claim 3.1. *Let $f \in Ar(S)$. Then there exists at most one $\overrightarrow{S}_i \in \mathcal{NCO}_1$ such that $f \cap \text{Vor}(\overrightarrow{S_i}) \neq \emptyset$.*

This, together with Cor. 2.1, establishes part (a) of Lemma 3.6. To see part (b) we provide a construction of an n-station wireless network in $\mathbb{R}^1$ that has $\Omega(n^2)$ cells corresponding to $\mathcal{NCO}_1$. Consider a set of n points S where s_i is positioned on x_i. The following claim defines a sufficient condition on the points $x_1, \ldots, x_n$ so that $|\mathcal{NCO}_1| = \Omega(n^2)$.

Claim 3.2. *If $x_2 > x_1$ and $x_{i+1} > x_i + x_{i-1} - x_1$ (for every $i \geq 2$), then $|\mathcal{NCO}_1| = \Omega(n^2)$.*

This establishes claim (b) of Lemma 3.6. □

We are now ready to complete the proof of Lemma 3.5.

Proof. By Lemma 3.1 and Eq. (3.5), a cancellation ordering might correspond to a distinct cell in $\mathcal{H}^{SIC}(s_1)$ only if it is in $\mathcal{NCO}_1$. Therefore the lemma follows immediately by Lemma 3.6. □

3.3 A Tighter Bound on the Number of Connected Components

In this section we introduce a key parameter of a wireless network, termed the *Compactenss Parameter* of the network, and establish a tighter bound on the number of connected components in the diagram under a certain condition on this parameter. For a wireless network $\mathcal{A} = \langle d, S, \bar{1}, N, \beta, \alpha \rangle$, the *Compactness Parameter* of $\mathcal{A}$ is

$$\mathcal{CP}(\mathcal{A}) \ = \ \beta^{1/\alpha}.$$

In what follows, we show that in the SIC-SINR model, this parameter plays a key role affecting the complexity of the resulting diagram. In particular, it follows that when $\alpha \to \infty$, both $\mathcal{CP}(\mathcal{A}) \to 1$ and the number of components gets closer to the bounds dictated by the high order Voronoi diagram. However, for a certain threshold value of $\mathcal{CP}(\mathcal{A})$ the situation is guaranteed to be much better, as discussed in the coming sections. Our results motivate further study of the compactness parameter, towards better understanding of the dynamics of the SINR diagram as a function of its compactness.

Toward the end of this subsection, we establish the following.

Lemma 3.7. *If the compactness parameter of the network $\mathcal{A}$ satisfies $\mathcal{CP}(\mathcal{A}) \geq 5$, then $\tau_i^{SIC}(\mathcal{A}) = O(1)$.*

Corollary 3.2. *If the compactness parameter of the network $\mathcal{A}$ satisfies $\mathcal{CP}(\mathcal{A}) \geq 5$, then $\tau^{SIC}(\mathcal{A}) = O(n)$.*

Note this this lemma implies that despite the fact that there exist instances in $\mathbb{R}^1$ of station sets S admitting $\mathcal{NCO}_i$ sequences of length $\Omega(n^2)$ (see Lemma 3.6(b)), only a constant number of those orderings correspond to nonempty reception regions.

We begin with two general claims that hold for every distance metric. Consider two stations $s_i, s_j \in S$ and a point $p \in \mathbb{R}^d$.

Claim 3.3. *If $p \in \mathcal{H}^{SIC}(s_i)$ and $\text{dist}(s_i, p) < \text{dist}(s_j, p)$, then $\text{dist}(s_j, p) \geq \mathcal{CP}(\mathcal{A}) \cdot \text{dist}(s_i, p)$.*

Proof. If $p \in \mathcal{H}^{SIC}(s_i)$ and $\text{dist}(s_i, p) < \text{dist}(s_j, p)$, then there exists $\overrightarrow{S}_k$ such that $p \in \mathcal{H}(\overrightarrow{S_k})$, $\textsf{Last}(S_k) = s_i$ and $s_j \notin S_k$. This implies that

$$\left(\frac{\text{dist}(s_j, p)}{\text{dist}(s_i, p)}\right)^{\alpha} \geq \frac{\text{dist}(s_i, p)^{-\alpha}}{\sum_{s \in S \setminus (S_k \setminus \{s_i\})} \text{dist}(s, p) + N} \geq \beta,$$

which yields the claim. □

We proceed with the second claim. Let $\overrightarrow{S}_i = (s_{i_1}, \ldots, s_{i_{k_1}})$ and $\overrightarrow{S}_j = (s_{j_1}, \ldots, s_{j_{k_2}})$, such that $k_1 \leq k_2$ and $S_i \nsubseteq S_j$.

Claim 3.4. *If $\mathcal{H}(\overrightarrow{S_i}), \mathcal{H}(\overrightarrow{S_j}) \neq \emptyset$, then $s_{i_1} \neq s_{j_1}$ or $\textsf{Last}(S_i) \neq \textsf{Last}(S_j)$, assuming $\mathcal{CP}(\mathcal{A}) \geq 5$.*

Proof. Assume, toward contradiction, that there exist i and j such that $\mathcal{H}(\overrightarrow{S_i}), \mathcal{H}(\overrightarrow{S_j}) \neq \emptyset$ and yet $s_{i_1} = s_{j_1}$ and $\textsf{Last}(S_i) = \textsf{Last}(S_j)$. Let m be the first index such that $s_{i_m} \neq s_{j_m}$. Let $S^* = \{s_{i_1}, \ldots s_{i_{m-2}}\}$ and set $s_1^* = s_{i_{m-1}}$, $s_2^* = s_{i_m}$ and $s_3^* = s_{j_m}$. Consider the reception zones $X1 = \mathcal{H}((s_1^*, s_2^*) \mid S \setminus S^*)$ and $X2 = \mathcal{H}((s_1^*, s_3^*) \mid S \setminus S^*)$. Since $X1 \subseteq \mathcal{H}(\overrightarrow{S_i})$ and $X2 \subseteq \mathcal{H}(\overrightarrow{S_j})$, it follows that $X1, X2 \neq \emptyset$. Consider points $p \in X1$ and $q \in X2$. For ease of notation, let $\text{dist}(s_1^*, p) = 1$, $\text{dist}(s_2^*, p) = c_2$ and $\text{dist}(s_3^*, p) = c_2 \cdot c_3$. Since $p \in X1$, it holds that $c_2, c_3 \geq \mathcal{CP}(\mathcal{A})$. The following inequalities are direct consequences of the fact that $p \in X1$.

$$\text{dist}(s_1^*, s_2^*) \ \leq \ c_2 + 1 \tag{3.6}$$

$$\text{dist}(s_2^*, s_3^*) \ \geq \ c_2 \cdot c_3 - c_2 \tag{3.7}$$

In addition, note that point q satisfies $\text{dist}(s_3^*, q) < \text{dist}(s_2^*, q)$ (since q cancels s_3^* before s_2^*). Combining this with Ineq. (3.7), we get that

$$\text{dist}(s_2^*, q) \ \geq \ \frac{c_2 \cdot c_3 - c_2}{2} \tag{3.8}$$

By Ineq. (3.6,3.8) we have that

$$\text{dist}(s_1^*, q) \geq \frac{c_2 \cdot c_3 - 3c_2 - 2}{2} \tag{3.9}$$

Next, note that $\text{dist}(s_2^*, q) \leq \text{dist}(s_1^*, q) + c_2 + 1$ by Ineq. (3.6). Combining this with Ineq. (3.8), we get that

$$\frac{\text{dist}(s_2^*, q)}{\text{dist}(s_1^*, q)} \leq 1 + \frac{2(c_2+1)}{c_3 \cdot c_2 - 3c_2 - 2} < \mathcal{CP}(\mathcal{A}) ,$$

for $\mathcal{CP}(\mathcal{A}) \geq 5$. It therefore follows that $q \notin \mathcal{H}(s_1^* \mid S \setminus S^* \cup \{s_3^*\})$ and therefore also $q \notin X2$, contradiction.

□

Hereafter, we focus on station s_i, and show that $\tau_i^{SIC} = O(1)$. The station s_j is called a *contributor* for i if $\mathcal{H}^{SIC}(s_i) \cap \text{VOR}(s_j) \neq \emptyset$. By Claim 3.4, each contributor can contribute at most one reception cell to $\mathcal{H}^{SIC}(s_i)$. We proceed by arguing that the number of contributors for i is bounded by a constant, for any dimension $d > 0$. (The proof of this claim will appear in the journal version of the paper.)

CLAIM 3.5. *The number of contributors for i is $O(1)$.*

We are now ready to conclude the proof of Lemma 3.7, by combining the above two claims.

Proof. By Claim 3.4, a given Voronoi cell can contribute at most one cell to a given $\mathcal{H}^{SIC}(s_i)$. By Claim 3.5, at most constant number of Voronoi cells can contribute to $\mathcal{H}^{SIC}(s_i)$. Overall, each $\mathcal{H}^{SIC}(s_i)$ is composed of constant number of cells. □

4 Construction of SIC-SINR Maps

The goal of this section is to provide an efficient scheme for constructing $\mathcal{H}^{SIC}(s_1)$, the reception zone under SIC for station s_1. Recall that $\mathcal{H}^{SIC}(s_1)$ is a collection of cells, each corresponding to a unique cancellation ordering $\overrightarrow{S_i}$. For a given network $\mathcal{A}$, the reception map without SIC can be drawn by using the characteristic polynomial of each zone $\mathcal{H}(s_i)$. In the SIC setting, however, the characteristic polynomial of a given cell depends on the cancellation ordering that generated it. This is due to the fact that $\mathcal{H}(\overrightarrow{S_i})$ is characterized by $|S_i| \leq n$ intersections of convex regions (see Eq. (3.4)) and the characteristic polynomial of each such region is known. The main task in drawing $\mathcal{H}^{SIC}(s_1)$ is therefore determining the (at most) $O(n^{2d})$ orderings of $\mathcal{NCO}_1$ among the collection of a-priori $(n-1)!$ orderings (i.e., $\mathcal{CO}_1$). We address this challenge by constructing the arrangement $Ar(S)$ and modifying it into a data structure that contains the information of all $\mathcal{NCO}_i$. Towards the end of the section we establish the following.

THEOREM 4.1. *A data structure HDS of size $O(n^{2d+1})$ can be constructed in time $O(n^{2d+1})$. Using HDS, $\mathcal{NCO}_i$ can be computed in time $O(n^{2d+1})$.*

4.1 Algorithm Description We begin by providing some notation. Associate with every point $p \in \mathbb{R}^d$ a *label* $\mathcal{L}(p)$, given by $\overrightarrow{S}^p = (s_1^p, \ldots, s_n^p)$, a sorted array of S stations such that $i < j$ if $\text{dist}(s_i^p, p) \leq \text{dist}(s_j^p, p)$. In Observation 4.1 we prove that for a cell f in the arrangement $Ar(S)$, all points in f have the same label. Hence we associate with each cell $f \in Ar(S)$ a unique label by setting $\mathcal{L}(f) = \mathcal{L}(p)$ for some point $p \in f$. (Observation 4.1 proves also that all cell labels are distinct.) Let $\overrightarrow{S}_i^p = (s_1^p, \ldots, s_k^p)$ be the prefix of $\overrightarrow{S}^p$ such that $\mathsf{Last}(\overrightarrow{S}_i^p) = s_k^p = s_i$. We then define $\mathcal{L}_i(p) = \overrightarrow{S}_i^p$ (and $\mathcal{L}_i(f)$ is defined accordingly).

We proceed by describing Algorithm BuildHDS. The algorithm is composed of two steps:
(1) building $Ar(S)$ (see Chapter 7 of [8]), and
(2) computing the labels $\mathcal{L}(f)$ for the cells $f \in Ar(S)$.
The resulting data structure HDS is a labeled arrangement denoted by $\mathcal{L}(Ar(S))$. We now describe the labeling process, given by Algorithm LabelArrangment. The algorithm starts from an arbitrary cell $f \in Ar(S)$ and computes $\mathcal{L}(f)$ by ordering the distances of stations in S with respect to some arbitrary point $p \in f$. Starting from f, $Ar(S)$ is now traversed in a DFS fashion, where the label of a newly encountered cell g, $\mathcal{L}(g)$, is computed using the label of its parent in the DFS tree, $\mathcal{L}(parent(g))$. Given that g and $parent(g)$ are separated by the hyperplane $HP(s_i, s_j)$, Algorithm LabelCell sets $\mathcal{L}(g) = \mathcal{L}(parent(g))$ and swaps the relevant positions s_i and s_j in $\mathcal{L}(parent(g))$. Finally, Algorithm ExtractNCO describes how $\mathcal{NCO}_i$ can be extracted from HDS. The algorithm constructs a hash-table HDS_i to maintain $\mathcal{NCO}_i$. To do that, the algorithm traverses the labeled arrangement $\mathcal{L}(Ar(S))$ and appends the truncated labels $\mathcal{L}_i(f)$ to HDS_i.

4.2 Analysis We next sketch the correctness proof of the algorithm. (Complete proofs will be given in the journal version of the paper.) We begin by showing that the labels of all points in a given cell $f \in Ar(S)$ are the same and the face labels are distinct.

OBSERVATION 4.1. *The points $p_1, p_2 \in \mathbb{R}^d$ belong to the same face in $Ar(S)$ iff $\mathcal{L}(p_1) = \mathcal{L}(p_2)$.*

To show that Algorithm LabelCell is correct, we establish the following claims.

LEMMA 4.1. *Let f_1, f_2 be two neighboring cells in $Ar(S)$. Then given $\mathcal{L}(f_1)$, Algorithm LabelCell computes $\mathcal{L}(f_2)$ in time $O(1)$.*

Finally we show that HDS_1 contains all $\mathcal{NCO}_1$, proving the correctness of Algorithm ExtractNCO.

LEMMA 4.2. *$\overrightarrow{S}_i \in \mathcal{NCO}_1$ iff $\overrightarrow{S}_i \in HDS_1$.*

It is left to consider the construction time and memory size of HDS. By Lemma 2.2 , $Ar(S)$ is constructed in $O(n^{2d})$ and maintained in $O(n^{2d})$ space. Clearly, the labeled arrangement $\mathcal{L}(Ar(S))$ is of size $O(n^{2d+1})$ as each label $\mathcal{L}(f)$ is of size $O(n)$. We now consider the time it takes to label $Ar(S)$ (i.e., the running time of Algorithm LabelArrangment). Note that the labeling of the first cell $f \in Ar(S)$ takes $O(n \log n)$ time as it involves sorting. Subsequent labels, however, are cheaper as they are computed directly using the label of their neighbor and the hyperplane that separates them. Overall, the labeling requires time linear in the size of the labeled arrangement and bounded by $O(n^{2d+1})$. Finally we consider Algorithm ExtractNCO. The extraction of $\mathcal{NCO}_i$ requires one pass over the labels of $Ar(S)$ cells and therefore takes $O(n^{2d+1})$ time. This completes the proof of Theorem 4.1.

5 Approximate Point Location

In this section, we utilize the topological properties derived thus far in order to address the problem of efficiently answering *point location queries* under interference cancellation. We first briefly review the topological and computational properties of the reception zones. In Eq. (3.5), we described the reception zone $\mathcal{H}^{SIC}(s_i)$ as a union of cells. By Lemmas 3.3 and 3.4, all cells in this union are distinct and convex. Moreover, Eq. (3.4) described each cell as the intersection of at most n SINR reception zones with no ordered cancellation. In Corollary 3.5, the number of possible cells in $\mathcal{H}^{SIC}(s_i)$, τ_i^{SIC} was shown to be $O(n^{2d})$. In fact, by Section 3.3, this bound can be tightened for threshold value of $\mathcal{CP}(\mathcal{A})$. Theorem 4.1 then established the existence of a polynomial-time algorithm to compute the non-empty cancellation ordering responsible for each cell. We will see here that when all these properties of the reception zone are put to use, a point location algorithm with logarithmic running time can be devised. For ease of illustration, we focus here only on the 2-dimensional case.

At this point, a few remarks are in order. First, when no cancellation is used, a station s_i can be heard at point p only if the signal from s_i is the strongest among all transmitting stations. In the uniform power scenario, this means s_i is heard at p only if p is in the Voronoi cell of s_i. As a result, one could devise a point location algorithm that for a given point p returns the nearest station s_i and whether $p \in \mathcal{H}(s_i)$ or not (with some slack). However, when cancellation is possible, several stations can be heard at p simultaneously (even for $\beta > 1$). Consequently, we consider here only *joint station-location* queries, that is, we wish to answer the following question: given a point p in the plane and a station s_i, is s_i heard at p under *some ordering* of cancellations?

Second, note that without offline preprocessing, $\Omega(n \log n)$ time is required to answer a single point location query. When processing a large number of queries, this might be too costly, hence the need for a tailored data structure that will facilitate $O(\log n)$ time for each query.

Toward our goal, we use a number of results and data structures derived for the SINR model with no cancellation [3]. For completeness, we include the basic concepts herein. For further details, the readers are referred to [3]. In [3], the authors use the following procedure: for a given reception zone $\mathcal{H}(s_i) \subset \mathbb{R}^2$, a square grid is drawn (see Figure 5(a)). Then, the boundary of $\mathcal{H}(s_i)$ is traversed, marking the grid squares that intersect with the boundary (with possibly $O(1)$ additional squares in each step). These marked grid squares form the region $\mathcal{H}^?(s_i)$, for which no conclusive answer can be returned. The interior grid squares form the region $\mathcal{H}^+(s_i)$, for which an affirmative answer is returned. The rest of squares form $\mathcal{H}^-(s_i)$, for which a negative answer is returned. It is proved that since the region $\mathcal{H}(s_i)$ is convex and fat (Lemma 2.1), for any given ϵ, one can choose the grid granularity such that $\text{area}(\mathcal{H}^?(s_i)) \leq \epsilon \cdot \text{area}(\mathcal{H}(s_i))$.

We now turn to our original problem. For each station s_i, we construct a data structure $\texttt{DS}(s_i)$ representing the reception zone $\mathcal{H}^{SIC}(s_i)$, together with two binary search trees on its (now possibly slightly overlapping) cells. Using these data structures, we are able to design an algorithm answering joint station-location queries in logarithmic time. Our main result in this section is the following.

THEOREM 5.1. *Let $\mathcal{A} = \langle d = 2, S, \psi = \bar{1}, N > 0, \beta > 1, \alpha = 2 \rangle$. Fix a station s_i. A data structure $\texttt{DS}(s_i)$ of size $O(n^9 \epsilon^{-1})$ is constructed in $O(n^{11} \epsilon^{-1})$ processing time. This data structure partitions the Euclidean plane into disjoint zones $\mathbb{R}^2 = \mathcal{H}^{SIC,+}(s_i) \cup \mathcal{H}^{SIC,-}(s_i) \cup \mathcal{H}^{SIC,?}(s_i)$ such that*

1. $\mathcal{H}^{SIC,+}(s_i) \subseteq \mathcal{H}^{SIC}(s_i)$
2. $\mathcal{H}^{SIC,-}(s_i) \cap \mathcal{H}^{SIC}(s_i) = \emptyset$
3. $\text{area}(\mathcal{H}^{SIC,?}(s_i)) \leq \epsilon \cdot \text{area}(\mathcal{H}^{SIC}(s_i))$.

$\texttt{DS}(s_i)$ identifies the zone to which a query point $p \in \mathbb{R}^2$ belongs in time $O(\log n)$.

Let us start by describing the construction of the data structure $\texttt{DS}(s_i)$. Fix a station in the network.

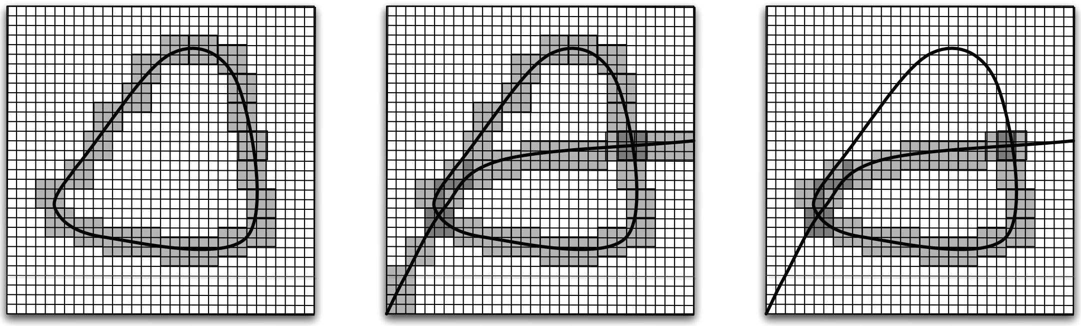

(a) One reception zone (b) Zones intersection (c) After one iteration

Figure 5: (a) The grid structure for the representation of $\mathcal{H}(s_i)$. The region boundary is in bold line. The undetermined squares, forming $\mathcal{H}^?(s_i)$, are marked in gray. The inner squares form $\mathcal{H}^+(s_i)$, while the outer form $\mathcal{H}^-(s_i)$. (b) Adding the boundary of a second region. (c) The gray areas marking the undetermined region of the intersection. Each grayed square was grayed out in at least one of the original shapes.

Without loss of generality, assume it is s_1. Represent the reception zone of s_1 as a union of cells $\mathcal{H}(\overrightarrow{S_i})$ indexed by the orderings $\overrightarrow{S_i} \in \mathcal{CO}_1$, as in Eq. (3.5). Recalling that $\mathcal{NCO}_1 = \{\overrightarrow{S_i} \subseteq S \mid \mathsf{Last}(\overrightarrow{S_i}) = s_1 \text{ and } \textsc{Vor}(\overrightarrow{S_i}) \neq \emptyset\}$, in the rest of this proof, sums and unions over $\overrightarrow{S_i} \in \mathcal{CO}_1$ in the representation of $\mathcal{H}^{SIC}(s_1)$ refer only to the distinct ordered subsets that define $\mathcal{H}^{SIC}(s_1)$, given in $\mathcal{NCO}_1$, rather than to all possible subsets.

By Lemma 3.3, all cells in the union in the right hand side of (3.5) are distinct, and each has the form

$$(5.10) \quad \mathcal{H}(\overrightarrow{S_i}) = \bigcap_{j=1}^{k} \mathcal{H}(s_{i_j} \mid S \setminus \{s_{i_1}, \ldots, s_{i_{(j-1)}}\}) ,$$

where $k = |\overrightarrow{S_i}|$. $\mathcal{H}(\overrightarrow{S_i})$ is thus the intersection of at most n reception zones of the form $\mathcal{H}(s_{i_j} \mid S \setminus \{s_{i_1}, \ldots, s_{i_{(j-1)}}\})$, that is, SINR reception zones without interference cancellation (though with, possibly, some stations turned off). Since there are $O(n^4)$ such distinct cells, whose cancellation orders are stored in the data structure HDS, our data structure $\mathtt{DS}(s_i)$ is constructed in three main steps: (1) retrieve the orderings $\overrightarrow{S_i} \subseteq S$ that form $\mathcal{H}^{SIC}(s_1)$ from the data structure HDS. Now all cancellation orders in $\mathcal{NCO}_1$ are known. (2) Construct the data structures that represent $\mathcal{H}(\overrightarrow{S_i})$, $\overrightarrow{S_i} \in \mathcal{NCO}_1$, by intersecting the required maps according to (5.10). (3) Assemble all data structures built in step (2) in a search tree facilitating logarithmic queries.

To retrieve the orderings of cancellations that piece together $\mathcal{H}^{SIC}(s_1)$, we traverse the data structure HDS. According to Theorem 4.1, in $O(n^4)$ time, $\mathcal{NCO}_1$ is computed and the orderings of cancellations are retrieved.

We now consider the data structure $\mathtt{DS}(\overrightarrow{S_i})$ required to represent $\mathcal{H}(\overrightarrow{S_i})$. Let $\tilde{\epsilon}$ be a small positive parameter, to be defined later. By [3], for each of the reception zones in the right hand side of Eq. (5.10), an $\tilde{\epsilon}$ approximation is achieved using a data structure of size $O(\tilde{\epsilon}^{-1})$. The time required to construct such a data structure is $O(n\tilde{\epsilon}^{-1})$. Let $\mathtt{DS}^m(\overrightarrow{S_i})$, $1 \leq m \leq k$, be the data structure for the mth region in Eq. (5.10). $\mathtt{DS}^m(\overrightarrow{S_i})$ partitions the space into three regions, $\mathbb{R}^2 = \mathcal{H}_m^+ \cup \mathcal{H}_m^- \cup \mathcal{H}_m^?$ (see Figure 5(a)). In particular, it is represented as a vector with $O(\tilde{\epsilon}^{-1})$ entries (indexed by the x-axis value of the grid columns). Each entry stores the locations of the upper (high y-axis values) and lower marked squares, that is, the squares forming the boundary of the reception zone. In this way, given a point p, one can compute the grid square in which p resides, access the data structure at the entry corresponding to the column, and based on the y-axis values of the upper and lower marked squares decide in $O(1)$ whether s_i is heard at p, unheard or a conclusive answer cannot be returned. Note, however, that in order to keep all structures $\{\mathtt{DS}^m(\overrightarrow{S_i})\}_{m=1}^k$ on the same grid, all should be constructed according to the finest grid resolution. On the other hand, since we are interested only in the intersection, all grid columns corresponding to x locations outside the region of $\mathtt{DS}^1(\overrightarrow{S_i})$ can be discarded.

To construct $\mathtt{DS}(\overrightarrow{S_i})$, we proceed as follows. We start with $\mathtt{DS}^1(\overrightarrow{S_i})$, and iterate over its $\tilde{\epsilon}^{-1}$ entries. Let $\mathtt{DS}[l, \mathit{upper}]$ and $\mathtt{DS}[l, \mathit{lower}]$ denote the coordinates of the undetermined upper and lower (respectively) squares in column l of $\mathtt{DS}$ (in [3], each such region was represented by at least three squares). We say that

$\mathtt{DS}^i[l, upper] \geq \mathtt{DS}^j[l, upper]$ if the y index of the lowermost square in $\mathtt{DS}^i[l, upper]$ is greater than or equal to the y index of the lowermost square in $\mathtt{DS}^j[l, upper]$. For each entry, we iterate over m, the number of cells to intersect, updating $\mathtt{DS}^1(\overrightarrow{S_i})[l, upper]$ and $\mathtt{DS}^1(\overrightarrow{S_i})[l, lower]$ in each iteration to represent the intersection of all regions at the specific entry. This is done by Algorithm $\mathsf{Intersect}$, which compares the upper and lower values of the current region (describing the intersection thus far) with the ones of the new region we intersect with, and updates the region according to one of the 6 possible intersection patters. Note that the procedure follows a simple "hierarchy" among the three types of squares: A square that was tagged as a $'-'$ is any one of the data structures $\{\mathtt{DS}^m(\overrightarrow{S_i})\}_{m=1}^k$ will be tagged as such in $\mathtt{DS}(\overrightarrow{S_i})$. A square will be tagged as a $'+'$ iff it is tagged as such in all intersecting data structures. Finally, a square tagged as a $'?'$ in $\mathtt{DS}(\overrightarrow{S_i})$ must have been tagged as such in at least one of the intersecting structures. An example is given in Figures 5(b) and 5(c). In this case, a new region intersects the current one in such a way that the upper boundary of the intersection is that of the new region, and the lower boundary of the intersection is that of the current region. The comparison done in each grid column, for each new region added takes $O(1)$ processing time. Thus, the whole process of intersecting at most n given data structures requires $O(n\tilde{\epsilon}^{-1})$ processing time. The resulting data structure $\mathtt{DS}(\overrightarrow{S_i})$ representing the intersection is also of size $O(\tilde{\epsilon}^{-1})$, as it is not required to be larger than the largest among the intersecting structures. Since this method requires having all data structures beforehand, these are constructed in $O(n^2\tilde{\epsilon}^{-1})$ processing time. Clearly, a data structure $\mathtt{DS}(s_1)$ to represent $\mathcal{H}^{SIC}(s_1)$ is built from the $O(n^4)$ data structures, representing all cells of $\mathcal{H}^{SIC}(s_1)$. $\mathtt{DS}(s_1)$ requires $O(n^4\tilde{\epsilon}^{-1})$ memory and $O(n^6\tilde{\epsilon}^{-1})$ processing time to construct. Items 1 and 2 of the theorem thus follow by construction.

Consider now the area of the undetermined regions in $\mathtt{DS}(s_1)$. To establish item 3 of the theorem, we wish to show that this region can be made arbitrarily small compared to $\mathcal{H}^{SIC}(s_1)$ with a proper choice of $\tilde{\epsilon}$. The required choices, and the ensuing analysis establishing the theorem, will appear in the journal version of the paper.

To achieve a query time that is logarithmic in n, simply arrange the $O(n^4)$ data structures representing $\mathcal{H}^{SIC}(s_1)$ in two binary search trees, one according to right-most grid point each structure represents, and one according to the lowest grid point each structure represents. Since this procedure is merely technical, we skip the details. Given a point $p \in \mathbb{R}^2$, one can identify the data structure to which p may belong in $O(\log n)$, and query the relevant data structure in $O(1)$.

References

[1] J.G. Andrews. Interference cancellation for cellular systems: a contemporary overview. *Wireless Communications, IEEE*, 12(2):19 – 29, April 2005.

[2] M. Andrews and M. Dinitz. Maximizing capacity in arbitrary wireless networks in the SINR model: Complexity and game theory. In *Proc. 28th Conf. IEEE Computer and Communications Soc. (INFOCOM)*, 2009.

[3] C. Avin, Y. Emek, E. Kantor, Z. Lotker, D. Peleg, and L. Roditty. SINR diagrams: Towards algorithmically usable SINR models of wireless networks. In *Proc. 28th Symp. on Principles of Distrib. Computing*, 2009. Full version available at `http://www.bgu.ac.il/~avin/papers/sinr_full.pdf`.

[4] C. Avin, Z. Lotker, and Y.-A. Pignolet. On the power of uniform power: Capacity of wireless networks with bounded resources. In *ESA*, pages 373–384, 2009.

[5] J.W. Byers, M. Luby, M. Mitzenmacher, and A. Rege. A digital fountain approach to reliable distribution of bulk data. In *Proc. ACM SIGCOMM'98*, pages 56–67, 1998.

[6] M. Costa and A. El-Gamal. The capacity region of the discrete memoryless interference channel with strong interference. *IEEE Trans. Inf. Th.*, 33:710–711, 1987.

[7] T. Cover and J. Thomas. *Elements of Information Theory.* Wiley, NY, 1991.

[8] H. Edelsbrunner. *Algorithms in Combinatorial Geometry.* Springer-Verlag, 1987.

[9] R.H. Etkin, D.N.C. Tse, and H. Wang. Gaussian interference channel capacity to within one bit. *IEEE Trans. Inf. Th.*, 54(12):5534–5562, 2008.

[10] O. Goussevskaia, R. Wattenhofer, M.M. Halldórsson, and E. Welzl. Capacity of arbitrary wireless networks. In *INFOCOM*, pages 1872–1880, 2009.

[11] P. Gupta and P.R. Kumar. The capacity of wireless networks. *IEEE Trans. Inf. Th.*, 46(2):388–404, 2000.

[12] M.M. Halldórsson and R. Wattenhofer. Wireless communication is in APX. In *ICALP*, pages 525–536, 2009.

[13] E. Kantor, Z. Lotker, M. Parter, and D. Peleg. The topology of wireless communication. In *STOC*, 2011.

[14] T. Moscibroda. The worst-case capacity of wireless sensor networks. In *Proc. 6th Int. Conf. on Info. Process. in Sensor Networks (IPSN)*, pages 1–10, 2007.

[15] T. Moscibroda and R. Wattenhofer. The complexity of connectivity in wireless networks. In *Proc. 25th INFOCOM*, 2006.

[16] T. Moscibroda, R. Wattenhofer, and Y. Weber. Protocol design beyond graph-based models. In *Proc. 5th Workshop on Hot Topics in Networks (Hotnets)*, 2006.

[17] A. Okabe, B. Boots, K. Sugihara, and S.-N. Chiu. *Spatial Tesselations.* Princeton University Press, 1992.

[18] A. Ozgur, O. Leveque, and D. Tse. Hierarchical cooperation achieves optimal capacity scaling in ad hoc networks. *IEEE Trans. Inf. Th.*, 53:3549–3572, 2007.

[19] H. Sato. The capacity of the Gaussian interference channel under strong interference (corresp.). *IEEE Trans. Inf. Th.*, 27(6):786–788, 1981.

[20] M. Ian Shamos and D. Hoey. Closest-point problems. In *FOCS*, 1975.

[21] L.N. Trefethen and D. Bau. *Numerical linear algebra*. SIAM, 1997.

[22] A.J. Viterbi. Very low rate convolution codes for maximum theoretical performance of spread-spectrum multiple-access channels. *Selected Areas in Communications, IEEE J.*, 8(4):641–649, 1990.