SINR Diagrams: Towards Algorithmically Usable SINR Models of Wireless Networks

[Extended Abstract]

Chen Avin
Department of Communication Systems Engineering
Ben Gurion University
Beer-Sheva, Israel
avin@cse.bgu.ac.il

Yuval Emek*
School of Electrical Engineering
Tel Aviv University
Tel Aviv, Israel
yuvale@eng.tau.ac.il

Erez Kantor
Department of Computer Science and Applied Mathematics
Weizmann Institute of Science
Rehovot, Israel
erez.kantor@weizmann.ac.il

Zvi Lotker†
Department of Communication Systems Engineering
Ben Gurion University
Beer-Sheva, Israel
zvilo@cse.bgu.ac.il

David Peleg‡
Department of Computer Science and Applied Mathematics
Weizmann Institute of Science
Rehovot, Israel
david.peleg@weizmann.ac.il

Liam Roditty
Department of Computer Science
Bar Ilan University
Ramat-Gan, Israel
liamr@macs.biu.ac.il

ABSTRACT

The rules governing the availability and quality of connections in a wireless network are described by *physical* models such as the *signal-to-interference & noise ratio (SINR)* model. For a collection of simultaneously transmitting stations in the plane, it is possible to identify a *reception zone* for each station, consisting of the points where its transmission is received correctly. The resulting *SINR diagram* partitions the plane into a reception zone per station and the remaining plane where no station can be heard.

SINR diagrams appear to be fundamental to understanding the behavior of wireless networks, and may play a key role in the development of suitable algorithms for such networks, analogous perhaps to the role played by Voronoi diagrams in the study of proximity queries and related issues in computational geometry. So far, however, the properties of SINR diagrams have not been studied systematically, and most algorithmic studies in wireless networking rely on simplified *graph-based* models such as the *unit disk graph (UDG)* model, which conveniently abstract away interference-related complications, and make it easier to handle algorithmic issues, but consequently fail to capture accurately some important aspects of wireless networks.

The current paper focuses on obtaining some basic understanding of SINR diagrams, their properties and their usability in algorithmic applications. Specifically, based on some algebraic properties of the polynomials defining the reception zones we show that assuming uniform power transmissions, the reception zones are convex and relatively well-rounded. These results are then used to develop an efficient approximation algorithm for a fundamental point location problem in wireless networks.

Categories and Subject Descriptors

F.2.2 [**Analysis of Algorithms and Problem Complexity**]: Nonnumerical Algorithms and Problems—*geometrical problems and computations*

General Terms

Algorithms

Keywords

wireless networks, SINR, reception zone

*Partially supported by the Israel Science Foundation, grants 664/05 and 221/07.

†Partially supported by a gift from Cisco research center.

‡Supported in part by grants from the Minerva Foundation and the Israel Ministry of Science.

PODC'09, August 10–12, 2009, Calgary, Alberta, Canada.

1. INTRODUCTION

Background. Wireless networks are hard to represent faithfully, due to the fact that deciding whether a transmission by a station s is successfully received by another station s' is nontrivial, and depends on the positioning and activities of s, s', and nearby stations that might interfere with

the transmission and prevent its successful reception. Thus such a transmission from s may reach s' under certain circumstances but fail to reach it under other circumstances. Moreover, the question is not entirely "binary", in the sense that connections can be of varying quality and capacity.

The rules governing the availability and quality of wireless connections can be described by *physical* or *fading channel* models (cf. [14, 4, 15]). Among those, the most commonly studied is the *signal-to-interference & noise ratio (SINR)* model. In the SINR model, the energy of a signal fades with the distance to the power of the path-loss parameter α. If the signal strength received by a device divided by the interfering strength of other simultaneous transmissions (plus the fixed *background noise* N) is above some *reception threshold* β, then the receiver successfully receives the message, otherwise it does not. Formally, denote by $\mathrm{dist}(p, q)$ the Euclidean distance between p and q, and assume that each station s_i transmits with power ψ_i. (A *uniform power network* is one where all stations transmit with the same power.) At an arbitrary point p, the transmission of station s_i is correctly received if

$$\frac{\psi_i \cdot \mathrm{dist}(p, s_i)^{-\alpha}}{N + \sum_{j \neq i} \psi_j \cdot \mathrm{dist}(p, s_j)^{-\alpha}} \;\geq\; \beta \;.$$

Hence for a collection $S = \{s_0, \ldots, s_{n-1}\}$ of simultaneously transmitting stations in the plane, it is possible to identify with each station s_i a *reception zone* $\mathcal{H}_i$ consisting of the points where the transmission of s_i is received correctly. It is believed that the path-loss parameter $2 \leq \alpha \leq 4$, where $\alpha = 2$ is the common "textbook" choice, and that the reception threshold $\beta \approx 6$ (β is always assumed to be greater than 1).

To illustrate how reception depends on the locations and activities of other stations, consider (the numerically generated) Fig. 1. Fig. 1(A) depicts uniform stations s_1, s_2, s_3 and their reception zones. Point p (represented as a solid black square) falls inside $\mathcal{H}_2$. Fig. 1(B) depicts the same stations except station s_1 has moved, so that now p does not receive any transmission. Fig. 1(C) depicts the stations in the same positions as Fig. 1(B), but now s_3 is silent, and as a result, the other two stations have larger reception zones, and p receives the message of s_1.

Fig. 1 illustrates a concept central to this paper, namely, the *SINR diagram*. An SINR diagram is a "reception map" characterizing the reception zones of the stations, namely, partitioning the plane into n reception zones $\mathcal{H}_i$, $0 \leq i \leq n-1$, and a zone $\mathcal{H}_\emptyset$ where no station can be heard. In many scenarios the diagram changes dynamically with time, as the stations may choose to transmit or keep silent, adjust their transmission power level, or even change their location from time to time.

Notice that for the locations in which the stations themselves are positioned, the SINR diagram is meaningless since it follows from the definition that station s_i cannot receive the transmission of any station s_j, $j \neq i$ (unless the two stations coincide). However, SINR diagrams can be extremely useful for a listening device that does not belong to S and is located in an arbitrary point in the plane. Using the SINR diagram, it is possible to decide which of the station of S (if any) can be correctly received in the location of the listening device.

It is our belief that SINR diagrams are fundamental to understanding the dynamics of wireless networks, and will play a key role in the development of suitable (sequential or distributed) algorithms for such networks, analogous perhaps to the role played by Voronoi diagrams in the study of proximity queries and related issues in computational geometry. Yet, to the best of our knowledge, SINR diagrams have not been studied systematically so far, from either geometric, combinatorial, or algorithmic standpoints. In particular, in the SINR model it is not clear what shapes the reception zones may take, and it is not easy to construct an SINR diagram even in a static setting.

Taking a broader perspective, a closely related concern motivating this paper is that while a fair amount of research exists on the SINR model and other variants of the physical model, little has been done in such models in the *algorithmic* arena. (Some recent exceptions are [8, 9, 10, 11, 12, 13, 16].) The main reason for this is that SINR models are complex and hard to work with. In these models it is even hard to decide some of the most elementary questions on a given setting, and it is definitely more difficult to develop communication or design protocols, prove their correctness and analyze their efficiency.

Subsequently, most studies of higher-layer concepts in wireless multi-hop networking rely on simplified *graph-based* models rather than on the SINR model. In particular, the model of choice for many protocol designers is the *unit disk graph (UDG)* model [6]. In this model, also known as the *protocol* model [9], two stations are considered to be *neighbors* if their Euclidean distance is at most 1. This defines the *UDG graph*. A silent station s successfully receives the message of a transmitting station s' if s' is a neighbor of s and no other neighbor of s transmits concurrently.

Graph-based models are attractive for higher-layer protocol design, as they conveniently abstract away interference-related complications. Issues of topology control, scheduling and allocation are also handled more directly, since notions such as adjacency and overlap are easier to define and test, in turn making it simpler to employ also some useful derived concepts such as domination, independence, clusters, and so on. On the down side, it should be realized that graph-based models, and in particular the UDG model, ignore or do not accurately capture a number of important physical aspects of real wireless networks. In particular, such models oversimplify the physical laws of interference; in reality, several nodes slightly outside the reception range of a receiver station v (which consequently are not adjacent to v in the UDG graph) might still generate enough cumulative interference to prevent v from successfully receiving a message from a sender station adjacent to it in the UDG graph. Hence the UDG model might yield a "false positive" indication of reception. Conversely, a simultaneous transmission by two or more neighbors should not always end in collision and loss of the message; in reality this depends on other factors, such as the relative distances and the relative strength of the transmissions. Hence in this case the UDG model yields a "false negative" indication.

In summary, while the existing body of literature on models and algorithms for wireless networks represents a significant base containing a rich collection of tools and techniques, the state of affairs described above leaves us in the unfortunate situation where the more practical graph-based models (such as the UDG model) are not sufficiently accurate, and the more accurate SINR model is not well-understood and therefore difficult for protocol designers. Hence obtaining a better understanding of the SINR model, and consequently

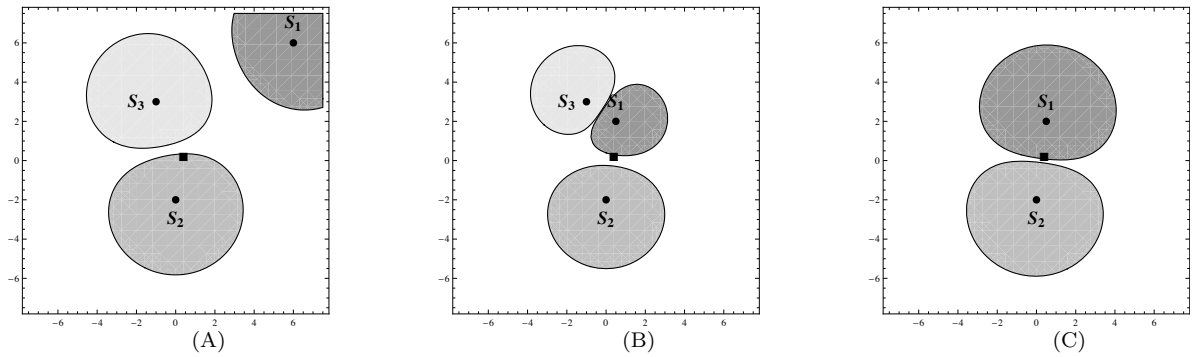

Figure 1: An example of SINR diagram with three transmitters s_1, s_2, s_3 and one receiver denoted by the solid black square. (A) The receiver can hear s_2. (B) Station s_1 moves and now the receiver cannot hear any transmission. (C) If, at the same locations as in (B), s_3 is silent, then the receiver can hear s_1.

bridging the gap between this physical model and the graph based models may have potentially significant (theoretical and practical) implications. This goal is the central motivation behind the current paper.

Related work. Some recent studies aim at achieving a better understanding of the SINR model. In their seminal work [9], Gupta and Kumar analyzed the capacity of wireless networks in the physical and protocol models. Moscibroda [11] analyzed the worst-case capacity of wireless networks, making no assumptions on the deployment of nodes in the plane, as opposed to almost all the previous work on this problem.

Thought provoking experimental results presented in [12] show that even basic wireless stations can achieve communication patterns that are impossible in graph-based models. Moreover, the paper presents certain situations in which it is possible to apply routing / transport schemes that may break the theoretical throughput limits of any protocol which obeys the laws of a graph-based model.

Another line of research, in which known results from the UDG model are analyzed under the SINR model, includes [13], which studies the problem of topology control in the SINR model, and [8], where impossibility results were proved in the SINR model for scheduling.

More elaborate graph-based models may employ two separate graphs, a connectivity graph $G_c = (S, E_c)$ and an interference graph $G_i = (S, E_i)$, such that a station s will successfully receive a message transmitted by a station s' if and only if s and s' are neighbors in the connectivity graph G_c and s does not have a concurrently transmitting neighbor in the interference graph G_i. Protocol designers often consider special cases of this more general model. For example, it is sometimes assumed that G_i is G_c augmented with all edges between 2-hop neighbors in G_c. Similarly, a variant of the UDG model handling transmissions and interference separately, named the *Quasi Unit Disk Graph (Q-UDG)* model, was introduced in [10]. In this model, two concentric circles are associated with each station, the smaller representing its reception zone and the larger representing its area of interference. An alternative interference model, also based on the UDG model, is proposed in [16].

Our results. A fundamental issue in wireless network modeling involves characterizing the reception zones of the stations and constructing the reception diagram. The current paper aims at gaining a better understanding of this issue in the SINR model, and as a consequence, deriving some algorithmic results. In particular, we consider the structure of reception zones in SINR diagrams corresponding to uniform power networks in d dimensional space ($d \geq 2$) with path-loss parameter $\alpha > 0$ and examine two specific properties of interest, namely, the *convexity* and *fatness*[1] of the reception zones. Apart from their theoretical interest, these properties are also of considerable practical significance, as obviously, having reception zones that are non-convex, or whose shape is arbitrarily skewed, twisted or skinny, might complicate the development of protocols for various design and communication tasks.

Our first result, which turns out to be surprisingly less trivial than we may have expected, is cast in Theorem 1. This theorem is proved in Section 3 for the central special case of $d = 2$ and $\alpha = 2$, by a complex analysis of the polynomials defining the reception zones, based on combining several observations with Sturm's condition for counting real roots. Proof of the general statement of the theorem is deferred to the full paper due to space considerations.

THEOREM 1. *The reception zones in an SINR diagram of a uniform power network in d dimensional space ($d \geq 2$) with path-loss parameter $\alpha > 0$ and reception threshold $\beta > 1$ are convex.*

Note that our convexity proof still holds when $\beta = 1$. In contrast, when $\beta < 1$, the reception zones of a uniform power network are not necessarily convex. This phenomenon is illustrated in (the numerically generated) Fig. 2. We then establish an additional attractive property of the reception zones (also proved in Section 3), which in a certain sense lends support to the model of *Quasi Unit Disk Graphs* suggested by Kuhn et al. in [10].

THEOREM 2. *The reception zones in an SINR diagram of a uniform power network with path-loss parameter $\alpha = 2$ and reception threshold $\beta > 1$ are fat.*

[1]The notion of fatness has received a number of non-equivalent technical definitions, all aiming at capturing the same intuition, namely, absence of long, skinny or twisted parts. In this paper we say that the reception zone of station s_i is *fat* if the ratio between the radii of the smallest ball centered at s_i that completely contains the zone and the largest ball centered at s_i that is completely contained by it is bounded by some constant. Refer to Section 2 for a formal definition.

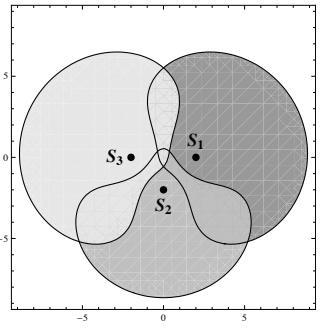

Figure 2: A uniform power network with path-loss parameter $\alpha = 2$, reception threshold $\beta = 0.3$, and background noise $N = 0.05$. The reception zones are clearly non-convex.

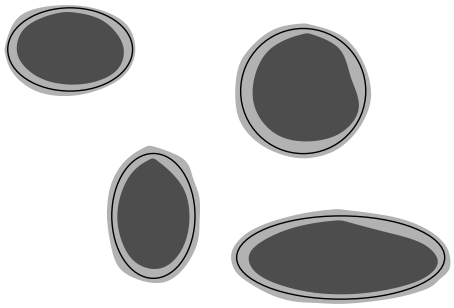

Figure 3: The reception zones $\mathcal{H}_i$ (enclosed by the bold lines) and the partition of the plane to disjoint zones $\mathcal{H}_i^+$ (dark gray), $\mathcal{H}_i^?$ (light gray), and $\mathcal{H}^-$ (the remaining white).

Armed with this characterization of the reception zones, we turn to a basic algorithmic task closely related to SINR diagrams, namely, answering *point location queries*. We address the following natural question: given a point in the plane, which reception zone contains this point (if any)? For UDG, this problem can be dealt with using known techniques (cf. [1, 2]). For arbitrary (non-unit) disk graphs, the problem is already harder, as the direct reduction to the technique of [2] no longer works. In the SINR model the problem becomes even harder. A naive solution will require computing the signal to interference & noise ratio for each station, yielding time $O(n^2)$. A more efficient ($O(n)$ time) querying algorithm can be based, for example, on the observation that there is a unique candidate $s_i \in S$ whose transmission may be received at p, namely, the one whose Voronoi cell contains p in the Voronoi diagram defined for S. However, it is not known if a sublinear query time can be obtained. This problem can in fact be thought of as part of a more general one, namely, point location over a general set of objects defined by polynomials and satisfying some "niceness" properties. Previous work on the problem dealt with Tarski cells, namely, objects whose boundaries are defined by a constant number of polynomials of constant degree [5, 1]. In contrast, the SINR diagram consists of objects (the reception zones) whose boundaries are defined by polynomials of degree proportional to n. We are unaware of a technique that answers point location queries for such objects in sublinear time.

Consider the SINR diagram of a uniform power network with path-loss parameter $\alpha = 2$ and reception threshold $\beta > 1$ and fix some performance parameter $0 < \epsilon < 1$. The following theorem is proved in Section 4 (refer to Fig. 3 for illustration).

THEOREM 3. *A data structure* DS *of size $O(n\epsilon^{-1})$ is constructed in $O(n^3\epsilon^{-1})$ preprocessing time. This data structure essentially partitions the Euclidean plane into disjoint zones $\mathbb{R}^2 = \bigcup_{i=0}^{n-1} \mathcal{H}_i^+ \cup \mathcal{H}^- \cup \bigcup_{i=0}^{n-1} \mathcal{H}_i^?$ such that for every $0 \le i \le n-1$:*
(1) $\mathcal{H}_i^+ \subseteq \mathcal{H}_i$;
(2) $\mathcal{H}^- \cap \mathcal{H}_i = \emptyset$; and
(3) $\mathcal{H}_i^?$ is bounded and its area is at most an ϵ-fraction of the area of $\mathcal{H}_i$.
On a query point $p \in \mathbb{R}^2$, DS *identifies the zone in $\{\mathcal{H}_i^+\}_{i=0}^{n-1} \cup \{\mathcal{H}^-\} \cup \{\mathcal{H}_i^?\}_{i=0}^{n-1}$ to which p belongs, in time $O(\log n)$.*

Open problems. Our results concern wireless networks with uniform power transmissions. General wireless networks are harder to deal with. For instance, the point location problem becomes considerably more difficult when different stations are allowed to use different transmission energy, since in this case, the appropriate graph-based model is no longer a unit-disk graph but a (*directed*) general disk graph, based on disks of arbitrary radii. An even more interesting case is the *variable power* setting, where the stations can adjust their transmission energy levels from time to time.

The problems discussed above become harder in a dynamic setting, and in particular, if we assume the stations are mobile, and extending our approach to the dynamic and mobile settings are the natural next steps.

2. PRELIMINARIES

Geometric notions. In the Euclidean plane $\mathbb{R}^2$, the *distance* between points p, q is denoted by $\mathrm{dist}(p, q) = \mathrm{dist}(q, p) = \|q - p\|$. A *ball* of radius r centered at point p is the set of all points at distance at most r from p, denoted by $B(p, r) = \{q \in \mathbb{R}^2 \mid \mathrm{dist}(p, q) \le r\}$. Point $p \in \mathbb{R}^2$ is *internal* to the point set P if there exists some $\epsilon > 0$ such that $B(p, \epsilon) \subseteq P$.

Consider some point set P. P is said to be an *open set* if all points $p \in P$ are internal points. P is said to be a *closed set* if the complement of P is an open set. If there exists some real r such that $\mathrm{dist}(p, q) \le r$ for every two points $p, q \in P$, then P is said to be *bounded*. A *compact* set is a set which is both closed and bounded. The *closure* of P is the smallest closed set containing P. The *boundary* of P, denoted by ∂P, is the intersection of the closure of P and the closure of its complement. A *connected set* is a point set that cannot be partitioned to two non-empty subsets such that each of the subsets has no point in common with the closure of the other. We refer to the closure of an open bounded connected set as a *thick* set. By definition, every thick set is compact.

A point set P is said to be *convex* if the segment $\overline{p\,q}$ is contained in P for every two points $p, q \in P$. The point set P is said to be *star-shaped* [7] with respect to point $p \in P$ if the segment $\overline{p\,q}$ is contained in P for every point $q \in P$. Note that a convex point set P is star-shaped with respect to any point $p \in P$; the converse is not necessarily true.

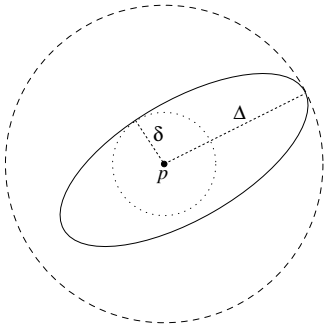

Figure 4: The zone Z (enclosed by the solid line) with the ball defining $\delta(p, Z)$ (dotted line) and the ball defining $\Delta(p, Z)$ (dashed line).

LEMMA 1. *A thick set P is convex if and only if every line intersects ∂P at most twice.*

We frequently use the term *zone* to describe a point set with some "niceness" properties. Unless stated otherwise, a zone refers to the union of an open connected set and some subset of its boundary. (A thick set is a special case of a zone.) It may also refer to a single point or to the finite union of zones. Given some bounded zone Z, we denote the *area* and *perimeter* of Z (assuming that they are well defined) by $\text{area}(Z)$ and $\text{per}(Z)$, respectively. Let Z be a non-empty bounded zone and let p be some internal point of Z. Denote

$$\delta(p, Z) = \sup\{r > 0 \mid Z \supseteq B(p, r)\}$$
$$\Delta(p, Z) = \inf\{r > 0 \mid Z \subseteq B(p, r)\}$$

and define the *fatness parameter* of Z with respect to p to be the ratio of $\Delta(p, Z)$ and $\delta(p, Z)$, denoted by $\varphi(p, Z) = \Delta(p, Z)/\delta(p, Z)$. (See Fig. 4.) The zone Z is *fat* w.r.t. p if $\varphi(p, Z) \leq c$ for some constant $c > 0$.

The *separation line* of two points p_1 and p_2 in the plane is the set of points $\{q \mid \text{dist}(p_1, q) = \text{dist}(p_2, q)\}$.

Wireless networks. We consider a wireless network $\mathcal{A} = \langle S, \psi, N, \beta\rangle$, where $S = \{s_0, s_1, \ldots, s_{n-1}\}$ is a set of transmitting *radio stations* embedded in the Euclidean plane, ψ is an assignment of a positive real *transmitting power* ψ_i to each station s_i, $N \geq 0$ is the *background noise*, and $\beta \geq 1$ is a constant that serves as the *reception threshold* (will be explained soon). For the sake of notational simplicity, s_i also refers to the point (a_i, b_i) in the plane where the station s_i resides. The network is assumed to contain at least two stations, i.e., $n \geq 2$. We say that $\mathcal{A}$ is a *uniform power network (UPN)* if $\psi = \bar{1}$, namely, if $\psi_i = 1$ for every i.

The *energy* of station s_i at point $p \neq s_i$ is defined to be $\text{E}_{\mathcal{A}}(s_i, p) = \psi_i \cdot \text{dist}(s_i, p)^{-2}$. The *energy* of a station set $T \subseteq S$ at point p, where $p \neq s_i$ for every $i \in T$, is defined to be $\text{E}_{\mathcal{A}}(T, p) = \sum_{i \in T} \text{E}_{\mathcal{A}}(s_i, p)$. Fix some station s_i and consider some point $p \notin S$. We define the *interference* to s_i at point p to be the energies at p of all stations other than s_i, denoted $\text{I}_{\mathcal{A}}(s_i, p) = \text{E}_{\mathcal{A}}(S - \{s_i\}, p)$. The *signal to interference & noise ratio (SINR)* of s_i at point p is defined as

$$\begin{aligned}\text{SINR}_{\mathcal{A}}(s_i, p) &= \frac{\text{E}_{\mathcal{A}}(s_i, p)}{\text{I}_{\mathcal{A}}(s_i, p) + N} \\ &= \frac{\psi_i \cdot \text{dist}(s_i, p)^{-2}}{\sum_{j \neq i} \psi_j \cdot \text{dist}(s_j, p)^{-2} + N} \ .\end{aligned} \tag{1}$$

Observe that $\text{SINR}_{\mathcal{A}}(s_i, p)$ is always positive since the transmitting powers and the distances of the stations from p are always positive and the background noise is non-negative. When the network $\mathcal{A}$ is clear from the context, we may omit it and write simply $\text{E}(s_i, p)$, $\text{I}(s_i, p)$, and $\text{SINR}(s_i, p)$.

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 is not smaller than the reception threshold of the network, i.e., $\text{SINR}(s_i, p) \geq \beta$. If this is the case, then 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 (somewhat tediously) as

$$\mathcal{H}_i = \{p \in \mathbb{R}^2 - S \mid \text{SINR}(s_i, p) \geq \beta\} \cup \{s_i\} \ .$$

(This definition is necessary as $\text{SINR}(s_i, \cdot)$ is not defined at any point in S and in particular, at s_i itself.)

Consider station s_0 and an arbitrary point $p = (x, y) \in \mathbb{R}^2$. By rearranging the expression in (1), we correlate the fundamental rule of the SINR model to the 2-variate polynomial $H(x, y)$ s.t. s_0 is heard at p iff

$$\begin{aligned}H(x, y) = \ \beta\Big[&\sum_{i>0} \psi_i \cdot \prod_{j \neq i} \left((a_j - x)^2 + (b_j - y)^2\right) \\ &+ N \cdot \prod_{i} \left((a_i - x)^2 + (b_i - y)^2\right) \Big] \\ &- \psi_0 \cdot \prod_{i>0} \left((a_i - x)^2 + (b_i - y)^2\right) \ \leq \ 0 \ .\end{aligned}$$

Note that this condition holds even for points $p \in S$. Consequently, we can rewrite $\mathcal{H}_0 = \{(x, y) \in \mathbb{R}^2 \mid H(x, y) \leq 0\}$, where the boundary points of $\mathcal{H}_0$ are exactly the roots of $H(x, y)$. In general, the polynomial $H(x, y)$ has degree $2n$; the degree is $2n - 2$ if the background noise $N = 0$. This polynomial plays a key role in the analysis carried out in Section 3.2.

A UPN $\mathcal{A} = \langle S, \bar{1}, N, \beta\rangle$ is said to be *trivial* if $|S| = 2$, $N = 0$, and $\beta = 1$. Note that for $i = 0, 1$, the reception zone $\mathcal{H}_i$ of station s_i in a trivial UPN is the half-plane consisting of all points whose distance to s_i is not greater than their distance to s_{1-i}. In particular, $\mathcal{H}_i$ is unbounded. For non-trivial networks, we have the following observation that relies on the fact that $\text{SINR}(s_i, \cdot)$ is a continuous function in $\mathbb{R}^2 - S$.

OBSERVATION 1. *Let $\mathcal{A} = \langle S, \bar{1}, N, \beta\rangle$ be a non-trivial UPN. Then the reception zone $\mathcal{H}_i$ is compact for every $s_i \in S$. Moreover, every point in $\mathcal{H}_i$ is closer to s_i than it is to any other station in S (i.e., $\mathcal{H}_i$ is strictly contained in the Voronoi cell of s_i in the Voronoi diagram of S).*

LEMMA 2. *Let $f : \mathbb{R}^2 \to \mathbb{R}^2$ be a mapping consisting of rotation, translation, and scaling by a factor of $\sigma > 0$. Consider some network $\mathcal{A} = \langle S, \psi, N, \beta\rangle$ and let $f(\mathcal{A}) = \langle f(S), \psi, N/\sigma^2, \beta\rangle$, where $f(S) = \{f(s_i) \mid s_i \in S\}$. Then for every station s_i and for all points $p \notin S$, we have $\text{SINR}_{\mathcal{A}}(s_i, p) = \text{SINR}_{f(\mathcal{A})}(f(s_i), f(p))$.*

3. CONVEXITY AND FATNESS OF THE RECEPTION ZONES

In this section we consider the SINR diagram of a uniform power network $\mathcal{A} = \langle S, \bar{1}, N, \beta \rangle$ and establish Theorems 1 and 2 for the special case of $d = 2$ and $\alpha = 2$. As all stations admit the same transmitting power, it is sufficient to focus on s_0 and to prove that the reception zone $\mathcal{H}_0$ is convex and fat. The fatness property is established in Section 3.4. For the convexity proof, we consider some arbitrary two points $p_1, p_2 \in \mathbb{R}^2$ and argue that if s_0 is heard at p_i for $i = 1, 2$, then s_0 is heard at all points in the segment $\overline{p_1\, p_2}$. This argument is established in three steps.

First, as a warmup, we prove that $\mathcal{H}_0$ is star-shaped with respect to s_0. This proof, presented in Section 3.1, establishes our argument for the case that p_1 and p_2 are colinear with s_0. Next, we prove that in the absence of a background noise (i.e., $N = 0$), if $p_i \in \mathcal{H}_0$ for $i = 1, 2$, then $\overline{p_1\, p_2} \subseteq \mathcal{H}_0$. This proof, presented in Section 3.3, relies on the analysis of a special case of a network consisting of only three stations which is analyzed in Section 3.2 and in a sense, constitutes the main technical challenge of this paper. Finally, we reduce the convexity proof of a UPN with n stations and arbitrary background noise, to that of a UPN with $n + 1$ stations and no background noise. This reduction is omitted from this extended abstract. While the analysis in Section 3.3 is consistent with some "physical intuition", the proof presented in Section 3.2 is based purely on algebraic arguments.

3.1 Star-shape

In this section we consider a UPN $\mathcal{A} = \langle S, \bar{1}, N, \beta \rangle$ and show that the reception zone $\mathcal{H}_0$ is star-shaped w.r.t. the station s_0. In fact, we establish the following slightly stronger lemma, whose proof is omitted from this extended abstract.

LEMMA 3. *Consider some point $p \in \mathbb{R}^2$. If $\mathrm{SINR}(s_0, p) \geq 1$, then $\mathrm{SINR}(s_0, q) > \mathrm{SINR}(s_0, p)$ for all internal points q in the segment $\overline{s_0\, p}$.*

Consider a non-trivial UPN $\mathcal{A} = \langle S, \bar{1}, N, \beta \rangle$ and suppose that $s_0 \neq s_j$ for every $j > 0$, i.e., the location of s_0 is not shared by other stations. Lemma 3 implies that the point set $\mathcal{H}'_0 = \{p \in \mathbb{R}^2 - S \mid \mathrm{SINR}_{\mathcal{A}}(s_0, p) > \beta\} \cup \{s_0\}$ is star-shaped w.r.t. s_0, and in particular, connected. Moreover, since SINR is a continuous function in $\mathbb{R}^2 - S$, it follows that $\mathcal{H}'_0$ is an open set. As $\mathcal{H}_0$ is the closure of $\mathcal{H}'_0$, we have the following.

COROLLARY 1. *In a nontrivial network, if s_0's location is not shared by other stations, then $\mathcal{H}_0$ is a thick set.*

3.2 Three stations with no background noise

In this section we analyze the special case of the UPN $\mathcal{A}_3 = \langle S, \bar{1}, N, \beta \rangle$, where $S = \{s_0, s_1, s_2\}$, $N = 0$, and $\beta = 1$. Our goal is to establish the following lemma, which constitutes the main technical challenge in the course of proving Theorem 1.

LEMMA 4. *The reception zone $\mathcal{H}_0$ of station s_0 in $\mathcal{A}_3$ is convex.*

Lemma 4 clearly holds if $s_j = s_0$ for some $j \in \{1, 2\}$, as this implies that $\mathcal{H}_0 = \{s_0\}$. So, in what follows we assume that no other station shares the location of s_0. By Corollary 1 we know that $\mathcal{H}_0$ is a thick set. Lemma 1 can now be employed to establish Lemma 4. To do that, it is required to show that every line intersects $\partial\mathcal{H}_0$ at most twice.

Consider an arbitrary line L in $\mathbb{R}^2$. We claim that L and $\partial\mathcal{H}_0$ have no more than two intersection points. If $s_0 \in L$, then the claim holds due to Lemma 3. Hence in the remainder of this section, assume that $s_0 \notin L$. Recall (see Section 2) that point $(x, y) \in \mathbb{R}^2$ is on the boundary of $\mathcal{H}_0$ iff it is a root of the polynomial $H(x, y) = ((a_0 - x)^2 + (b_0 - y)^2)((a_1 - x)^2 + (b_1 - y)^2 + (a_2 - x)^2 + (b_2 - y)^2) - ((a_1 - x)^2 + (b_1 - y)^2)((a_2 - x)^2 + (b_2 - y)^2)$ so it is sufficient to prove that the projection of $H(x, y)$ on the line L has at most two distinct real roots.

Employing Lemma 2, we may assume that s_0 is located at the origin and that L is the line $y = 1$. By substituting $y = 1$ into the expression of $H(x, y)$ and rearranging the resulting expression, we get

$$
\begin{aligned}
&H(x) = \\
&\left(x^2 + 1\right)\left((a_1 - x)^2 + (b_1 - 1)^2 + (a_2 - x)^2 + (b_2 - 1)^2\right) \\
&- \left((a_1 - x)^2 + (b_1 - 1)^2\right)\left((a_2 - x)^2 + (b_2 - 1)^2\right) = \\
&x^4 + (2 - 4a_1a_2)x^2 \\
&+ (2a_2a_1^2 + 2a_2^2a_1 + 2(1 - b_2)^2a_1 \\
&- 2a_1 + 2a_2(1 - b_1)^2 - 2a_2)x \\
&+ a_1^2 - a_1^2a_2^2 + a_2^2 - a_2^2(1 - b_1)^2 + (1 - b_1)^2 \\
&- a_1^2(1 - b_2)^2 - (1 - b_1)^2(1 - b_2)^2 + (1 - b_2)^2 \ ,
\end{aligned}
$$

so that $(x, 1)$ is on the boundary of $\mathcal{H}_0$ if and only if x is a root of $H(x)$.

Our goal in the remainder of this section is to show that $H(x)$ has at most two distinct real roots, and towards this goal we will first invest some effort in simplifying this polynomial. As a first step we show that we can restrict our attention to the case where both s_1 and s_2 are in the first quarter above the line $y = 1$, that is, $a_j > 0$ and $b_j \geq 1$ for $j = 1, 2$. The latter restriction is trivial due to the symmetry of interference along the line $y = 1$, which implies that if $b_j < 1$ for some $j \in \{1, 2\}$, then relocating s_j in $(a_j, 1 + |1 - b_j|)$ does not affect the interference at q for all points q on the line $y = 1$, and in particular, does not affect the number of simple real roots of $H(x)$. For the former restriction we establish the following proposition whose proof is omitted from this extended abstract.

PROPOSITION 1. *If $\mathrm{sign}(a_1) \cdot \mathrm{sign}(a_2) \neq 1$, then $H(x)$ has at most two distinct real roots.*

By Proposition 1, we may hereafter assume that $\mathrm{sign}(a_1) = \mathrm{sign}(a_2) \neq 0$. The case where $a_j < 0$ for $j = 1, 2$ is redundant, since relocating station s_j in $(-a_j, b_j)$ turns $H(x)$ into $H(-x)$, and in particular, does not affect the number of distinct real roots of the polynomial. Therefore, in what follows we assume that $a_j > 0$ and $b_j \geq 1$ for $j = 1, 2$. Our next step is to rewrite $H(x)$ as

$$H(x) = (x^2 + 1)^2 - J(x) \ , \tag{2}$$

where $J(x) = 4a_2a_1x^2 - (2a_2a_1^2 + 2a_2^2a_1 + 2b_2^2a_1 - 4b_2a_1 + 2a_2b_1^2 - 4a_2b_1)x + a_2^2a_1^2 + b_2^2a_1^2 - 2b_2a_1^2 + a_2^2b_1^2 + b_1^2b_2^2 - 2b_1b_2^2 - 2a_2^2b_1 - 2b_1^2b_2 + 4b_1b_2$ is a polynomial of degree 2. Under

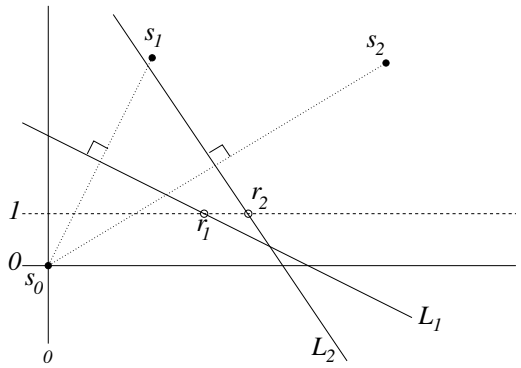

Figure 5: The point $(r_j, 1)$ is on the separation line L_j of s_0 and s_j.

the assumption that a_1 and a_2 are positive, $J(x)$ has the following (not necessarily distinct) real roots:

$$r_1 = \frac{a_1^2 + (b_1 - 2)b_1}{2a_1} \qquad \text{and} \qquad r_2 = \frac{a_2^2 + (b_2 - 2)b_2}{2a_2} \ .$$

Perhaps surprisingly, the root r_j corresponds to the x-coordinate of the intersection point of the line $y = 1$ and the separation line L_j of s_0 and s_j for $j = 1, 2$ (see Fig. 5). To validate this observation, note that the point (x, y) is on L_j if and only if $(x - a_j)^2 + (y - b_j)^2 = x^2 + y^2$, or equivalently, if and only if $a_j^2 + b_j^2 = 2(a_j x + b_j y)$. Fixing $y = 1$, we get that $x = \frac{a_j^2 + (b_j - 2)b_j}{2a_j} = r_j$.

Moreover, since x has a negative coefficient when L_j is expressed as $y = -\frac{a_j}{b_j}x + a_j^2 + b_j^2$, we conclude that the point $(x', 1)$ is as close to s_j at least as it is to s_0 for all $x' \geq r_j$. The next corollary follows since the real roots of $H(x)$ correspond to points on the boundary of $\mathcal{H}_0$, and since s_0 is heard at all such points.

Corollary 2. *The real x' is not a root of the polynomial $H(x)$ for any $x' \geq \min\{r_1, r_2\}$.*

Incidentally, let us comment without proof that the point $\min\{r_1, r_2\}$ is the intersection point of the line $y = 1$ and the boundary of the Voronoi cell of s_0 in the Voronoi diagram of S.

Fix $\bar{r} = (r_1 + r_2)/2$ and define the shifted variable $z = x - \bar{r}$. Since $\bar{r}$ is the center of the parabola $J(x)$, it follows that when expressing $J(x)$ in terms of the shifted variable z, we get the form $J(z) = \gamma z^2 + \delta$, where $\gamma > 0$ since the leading coefficient of $J(x)$ is positive, and $\delta \leq 0$ since $J(x)$ has at least one real root. We can now express $H(x)$, as represented in (2), in terms of the shifted variable z, introducing the polynomial

$$\widehat{H}(z) \ = \ \left((z + \bar{r})^2 + 1\right)^2 - \gamma z^2 - \delta \ ,$$

which is obviously much simpler. Clearly, $H(x)$ and $\widehat{H}(z)$ have the same number of distinct real roots.

In what follows, we employ *Sturm's Theorem* in order to bound the number of distinct real roots of $\widehat{H}(z)$. Consider some degree n polynomial $P(x) = \alpha_n x^n + \cdots + \alpha_1 x + \alpha_0$ over the reals. The *Sturm sequence* of $P(x)$ is a sequence of polynomials denoted by $P_0(x), P_1(x), \ldots, P_m(x)$, where $P_0(x) = P(x)$, $P_1(x) = P'(x)$, and $P_i(x) = -\mathrm{rem}(P_{i-2}(x)/P_{i-1}(x))$ for $i > 1$. This recursive definition terminates at step m such that $-\mathrm{rem}(P_{m-1}(x)/P_m(x)) = 0$. Since the degree of $P_i(x)$ is at most $n - i$, we conclude that $m \leq n$. Define $\mathrm{SC}_P(t)$ to be the number of sign changes in the sequence $P_0(t), P_1(t), \ldots, P_m(t)$. We are now ready to state the following theorem attributed to Jacques Sturm, 1829 (cf. [3]).

Theorem 4 (Sturm's condition). *Consider two reals a, b, where $a < b$ and neither of them is a root of $P(x)$. Then the number of distinct real roots of $P(x)$ in the interval (a, b) is $\mathrm{SC}_P(a) - \mathrm{SC}_P(b)$.*

We will soon show that[2] $\mathrm{SC}_{\widehat{H}}(-\infty) \leq 3$ and $\mathrm{SC}_{\widehat{H}}(\infty) \geq 1$, hence $\mathrm{SC}_{\widehat{H}}(-\infty) - \mathrm{SC}_{\widehat{H}}(\infty) \leq 2$. Therefore, by Theorem 4, we conclude that $\widehat{H}(z)$ has at most two distinct real roots. It is sufficient for our purposes to consider the first three polynomials in the Sturm sequence of $\widehat{H}(z)$:

$$\begin{aligned}
\widehat{H}_0(z) =& z^4 + 4\bar{r}z^3 + (6\bar{r}^2 - \gamma + 2)z^2 + (4\bar{r}^3 + 4\bar{r})z + \bar{r}^4 \\
&+ 2\bar{r}^2 - \delta + 1 \\
\widehat{H}_1(z) =& 4z^3 + 12\bar{r}z^2 + 2\left(6\bar{r}^2 - \gamma + 2\right)z + 4\bar{r}^3 + 4\bar{r} \\
\widehat{H}_2(z) =& (\gamma/2 - 1)z^2 - \bar{r}(2 + \gamma/2)z - \bar{r}^2 - 1 + \delta \ .
\end{aligned}$$

Proposition 2. *The polynomial $\widehat{H}(z)$ satisfies $\mathrm{SC}_{\widehat{H}}(\infty) \geq 1$.*

Proof. We first argue that $\widehat{H}(z)$ does not have any root in the interval $[0, \infty)$. This argument holds due to Corollary 2 since by the definition of $z = x - \bar{r}$, $z \geq 0$ implies $x \geq \bar{r} \geq \min\{r_1, r_2\}$. Therefore Theorem 4 guarantees that $\mathrm{SC}_{\widehat{H}}(\infty) = \mathrm{SC}_{\widehat{H}}(0)$. Now, the sign of $\widehat{H}_0(0)$ is positive while the sign of $\widehat{H}_2(0)$ is negative, so there must be at least one sign change when the Sturm sequence of $\widehat{H}(z)$ is evaluated on 0, hence $\mathrm{SC}_{\widehat{H}}(\infty) = \mathrm{SC}_{\widehat{H}}(0) \geq 1$. □

Proposition 3. *The polynomial $\widehat{H}(z)$ satisfies $\mathrm{SC}_{\widehat{H}}(-\infty) \leq 3$.*

Proof. First note that there are at most five polynomials in the Sturm sequence of $\widehat{H}(z)$, hence $\mathrm{SC}_{\widehat{H}}(-\infty)$ cannot be greater than 4. Suppose, towards deriving contradiction, that $\mathrm{SC}_{\widehat{H}}(-\infty) = 4$. This implies that there are exactly 5 polynomials in the Sturm sequence of $\widehat{H}(z)$ and the degree of $\widehat{H}_i(z)$ is $4 - i$ for every $0 \leq i \leq 4$. Clearly, both $\mathrm{SC}_{\widehat{H}}(-\infty)$ and $\mathrm{SC}_{\widehat{H}}(\infty)$ depend solely on the signs of the leading coefficients of the polynomials in the Sturm sequence of $\widehat{H}(z)$. Since $\mathrm{sign}(\widehat{H}_0(-\infty)) = 1$, we must have $\mathrm{sign}(\widehat{H}_i(-\infty)) = -1$ for $i = 1, 3$ and $\mathrm{sign}(\widehat{H}_i(-\infty)) = 1$ for $i = 2, 4$. As $\mathrm{sign}(\widehat{H}_i(\infty)) = \mathrm{sign}(\widehat{H}_i(-\infty))$ for $i = 0, 2, 4$ and $\mathrm{sign}(\widehat{H}_i(\infty)) = -\mathrm{sign}(\widehat{H}_i(-\infty))$ for $i = 1, 3$, we conclude that $\mathrm{sign}(\widehat{H}_i(\infty)) = 1$ for every $0 \leq i \leq 4$. Therefore $\mathrm{SC}_{\widehat{H}}(\infty) = 0$, in contradiction to Proposition 2. □

To conclude, combining Propositions 2 and 3 with Theorem 4, we get that $\widehat{H}(z)$ has at most two distinct real roots, and thus $H(x)$ has at most two distinct real roots. It follows that every line has at most two intersection points with $\partial\mathcal{H}_0$, which completes the proof of Lemma 4.

[2]We write $\mathrm{SC}_{\widehat{H}}(\infty)$ and $\mathrm{SC}_{\widehat{H}}(-\infty)$ as a convenient shorthand for $\lim_{z\to\infty} \mathrm{SC}_{\widehat{H}}(z)$ and $\lim_{z\to-\infty} \mathrm{SC}_{\widehat{H}}(z)$, respectively.

3.3 Convexity with n stations and no background noise

We now return to a UPN $\mathcal{A} = \langle S, \bar{1}, N, \beta \rangle$ with an arbitrary number of stations and with an arbitrary reception threshold $\beta \geq 1$, but still, with no background noise (i.e., $N = 0$), and establish the convexity of $\mathcal{H}_0$.

LEMMA 5. *The reception zone $\mathcal{H}_0$ of station s_0 in $\mathcal{A}$ is convex.*

Lemma 5 is proved by induction on the number of stations in the network, $n = |S|$. For the base of the induction, we consider the case where $n = 2$. The theorem clearly holds if s_0 and s_1 share the same location, as this implies that $\mathcal{H}_0 = \{s_0\}$. Furthermore, if $\beta = 1$, which means that $\mathcal{A}$ is trivial, then $\mathcal{H}_0$ is a half-plane and in particular, convex. So, in what follows we assume that $s_0 \neq s_1$ and that $\beta > 1$.

Corollary 1 implies that $\mathcal{H}_0$ is a thick set, thus, by Lemma 1, it is sufficient to argue that every line L has at most two intersection points with $\partial\mathcal{H}_0$. If $s_0 \in L$, then the argument holds due to Lemma 3. If $s_0 \notin L$, then $\frac{\mathrm{E}(s_0,q)}{\mathrm{I}(s_0,q)} = \beta$ is essentially a quadratic equation, thus it has at most two real solutions which correspond to at most two intersection points of L and $\partial\mathcal{H}_0$.

The inductive step of the proof of Lemma 5 is more involved. We consider some arbitrary two points $p_1, p_2 \in \mathcal{H}_0$ and prove that $\overline{p_1\,p_2} \subseteq \mathcal{H}_0$. Informally, we will show that if there exist at least two stations other than s_0, then we can discard one station and relocate the rest so that the interference at p_i remains unchanged for $i = 1, 2$ and the interference at q does not decrease for all points $q \in \overline{p_1\,p_2}$. By the inductive hypothesis, the segment $\overline{p_1\,p_2}$ is contained in $\mathcal{H}_0$ in the new setting, hence it is also contained in $\mathcal{H}_0$ in the original setting. This idea relies on the following lemma, whose proof is omitted from this extended abstract.

LEMMA 6. *Consider the stations s_0, s_1, s_2 and some distinct two points $p_1, p_2 \in \mathbb{R}^2$. If $\mathrm{E}(s_0, p_i) \geq \mathrm{E}(\{s_1, s_2\}, p_i)$ for $i = 1, 2$, then there exists a location $s^* \in \mathbb{R}^2$ such that*
(1) $\mathrm{E}(s^*, p_i) = \mathrm{E}(\{s_1, s_2\}, p_i)$ for $i = 1, 2$; and
(2) $\mathrm{E}(s^*, q) \geq \mathrm{E}(\{s_1, s_2\}, q)$, for all points q in the segment $\overline{p_1\,p_2}$.

We now turn to describe the inductive step in the proof of Lemma 5. Assume by induction that the assertion of the theorem holds for $n \geq 2$ stations, i.e., that in a UPN with $n \geq 2$ stations and no background noise we have $\overline{p_1\,p_2} \subseteq \mathcal{H}_0$ for every $p_1, p_2 \in \mathcal{H}_0$. Now consider a UPN $\mathcal{A}$ with $n + 1$ stations $s_0, \ldots, s_n$ and no background noise. Let $p_1, p_2 \in \mathcal{H}_0$. Suppose that s_1 is closest to, say, p_1 among all stations $s_1, \ldots, s_n$. Since $p_1, p_2 \in \mathcal{H}_0$, we know that $\mathrm{E}(s_0, p_i) > \mathrm{E}(\{s_1, s_2\}, p_i)$ for $i = 1, 2$. Lemma 6 guarantees that there exists a station location $s^* \in \mathbb{R}^2$ such that (1) $\mathrm{E}(s^*, p_i) = \mathrm{E}(\{s_1, s_2\}, p_i)$ for $i = 1, 2$; and (2) $\mathrm{E}(s^*, q) \geq \mathrm{E}(\{s_1, s_2\}, q)$, for all points q in the segment $\overline{p_1\,p_2}$.

Note that the station location s^* must differ from s_0. This is because $\mathrm{E}(s^*, p_i) = \mathrm{E}(\{s_1, s_2\}, p_i)$ while $\mathrm{E}(s_0, p_i) > \mathrm{E}(\{s_1, s_2\}, p_i)$ for $i = 1, 2$, thus $\mathrm{dist}(s^*, p_i) > \mathrm{dist}(s_0, p_i)$.

Consider the n-station UPN $\mathcal{A}^*$ obtained from $\mathcal{A}$ by replacing s_1 and s_2 with a single station located at s^*. Note that $\mathrm{I}_{\mathcal{A}^*}(s_0, p_i) = \mathrm{I}_{\mathcal{A}}(s_0, p_i)$ for $i = 1, 2$ and $\mathrm{I}_{\mathcal{A}^*}(s_0, q) \geq \mathrm{I}_{\mathcal{A}}(s_0, q)$ for all points $q \in \overline{p_1\,p_2}$, hence $\mathrm{SINR}_{\mathcal{A}^*}(s_0, p_i) = \mathrm{SINR}_{\mathcal{A}}(s_0, p_i)$ for $i = 1, 2$ and $\mathrm{SINR}_{\mathcal{A}^*}(s_0, q) \leq \mathrm{SINR}_{\mathcal{A}}(s_0, q)$. By the inductive hypothesis, $\mathrm{SINR}_{\mathcal{A}^*}(s_0, q) \geq \beta$ for all points $q \in \overline{p_1\,p_2}$, therefore $\mathrm{SINR}_{\mathcal{A}}(s_0, q) \geq \beta$ and s_0 is heard at q in $\mathcal{A}$. It follows that every $q \in \overline{p_1\,p_2}$ belongs to $\mathcal{H}_0$ in $\mathcal{A}$, which establishes the assertion and completes the proof of Lemma 5.

3.4 The fatness of the reception zones

In this section we develop a deeper understanding of the "shape" of the reception zones by analyzing their fatness. Consider a UPN $\mathcal{A} = \langle S, \bar{1}, N, \beta \rangle$, where $S = \{s_0, \ldots, s_{n-1}\}$ and[3] $\beta > 1$ is a constant. We focus on s_0 and assume that its location is not shared by any other station (otherwise, the reception zone $\mathcal{H}_0 = \{s_0\}$). Theorem 5, whose proof is omitted from this extended abstract, exhibits explicit bounds on $\Delta(s_0, \mathcal{H}_0)$ and $\delta(s_0, \mathcal{H}_0)$.

THEOREM 5. *In a uniform energy network $\mathcal{A} = \langle S, \bar{1}, N, \beta \rangle$, where $S = \{s_0, \ldots, s_{n-1}\}$ and $\beta > 1$ is a constant, if the minimum distance from s_0 to any other station is $\kappa > 0$, then $\delta(s_0, \mathcal{H}_0) \geq \frac{\kappa}{\sqrt{\beta(n-1+N\cdot\kappa^2)}+1}$ and $\Delta(s_0, \mathcal{H}_0) \leq \frac{\kappa}{\sqrt{\beta(1+N\cdot\kappa^2)}-1}$.*

These bounds imply that $\varphi(s_0, \mathcal{H}_0) = O(\sqrt{n})$. In fact, a slightly more careful analysis (omitted from this extended abstract as well) shows that $\varphi(s_0, \mathcal{H}_0) = O(1)$, thus establishing Theorem 2.

4. HANDLING APPROXIMATE POINT LOCATION QUERIES

Our goal in this section is to prove Theorem 3. In fact, our technique for approximate point location queries is suitable for a more general framework of zones (and diagrams). Let $Q(x, y)$ be a 2-variate polynomial of degree m and suppose that the zone

$$\mathcal{Q} = \big\{(x, y) \in \mathbb{R}^2 \mid Q(x, y) \leq 0\big\}$$

is a thick set and that $Q(x, y)$ is strictly negative for all internal points (x, y) of $\mathcal{Q}$. Moreover, suppose that we are given an internal point s of $\mathcal{Q}$, a lower bound $\tilde{\delta}$ on $\delta(s, \mathcal{Q})$, and an upper bound $\tilde{\Delta}$ on $\Delta(s, \mathcal{Q})$. Let $0 < \epsilon < 1$ be a predetermined performance parameter. We construct in $O(m^2(\tilde{\Delta}/\tilde{\delta})^2\epsilon^{-1})$ preprocessing time a data structure `QDS` of size $O((\tilde{\Delta}/\tilde{\delta})^2\epsilon^{-1})$. `QDS` essentially partitions the Euclidean plane into three disjoint zones $\mathbb{R}^2 = \mathcal{Q}^+ \cup \mathcal{Q}^- \cup \mathcal{Q}^?$ such that
(1) $\mathcal{Q}^+ \subseteq \mathcal{Q}$;
(2) $\mathcal{Q}^- \cap \mathcal{Q} = \emptyset$; and
(3) $\mathcal{Q}^?$ is bounded and its area is at most an ϵ-fraction of the area of $\mathcal{Q}$.
Given a query point $p \in \mathbb{R}^2$, `QDS` answers in constant time whether p is in $\mathcal{Q}^+$, $\mathcal{Q}^-$, or $\mathcal{Q}^?$.

In Section 4.1 we describe the construction of `QDS`. In Section 4.2 we explain how the reception zones and the SINR diagram fall into the above framework and establish Theorem 3.

[3]Unlike the convexity proof which holds for any $\beta \geq 1$, the analysis presented in the current section is only suitable for β being a constant strictly greater than 1. In fact, when $\beta = 1$, the fatness parameter is not necessarily defined (think of a trivial network).

4.1 The construction of QDS

In this section we describe the construction of QDS. Let γ be a positive real to be determined later on. The data structure QDS is based upon imposing a γ-spaced *grid*, denoted by G_γ, on the Euclidean plane. The grid is aligned so that the point s is a grid vertex. The Euclidean plane is partitioned to grid *cells* with respect to G_γ in the natural manner, where ties are broken such that each cell contains all points on its south edge except its south east corner and all points on its west edge except its north west corner (the cell does contain its south west corner). Given some cell C, we define its *9-cell*, denoted by $\sharp C$, as the collection of 3×3 cells containing C and the eight cells surrounding it.

The grid cells will be classified to three types corresponding to the zones $\mathcal{Q}^+$, $\mathcal{Q}^-$, and $\mathcal{Q}^?$: cells of type T^+ are fully contained in $\mathcal{Q}$; cells of type T^- do not intersect $\mathcal{Q}$; and cells of type $\mathrm{T}^?$ are suspect of partially overlapping $\mathcal{Q}$, i.e., having some points in $\mathcal{Q}$ and some points not in $\mathcal{Q}$. A query on point $p \in \mathbb{R}^2$ is handled merely by computing the cell to which p belongs and returning its type. Our analysis relies on bounding the number of $\mathrm{T}^?$ cells.

By definition, the zone $\mathcal{Q}$ contains a ball of radius $\tilde{\delta}$ and it is contained in a ball of radius $\tilde{\Delta}$, both centered at s. Clearly, the area of $\mathcal{Q}$ is lower bounded by the area of any ball it contains. Since $\mathcal{Q}$ is convex, it follows that its perimeter is upper bounded by the perimeter of any ball that contains it. Therefore the zone $\mathcal{Q}$ satisfies

$$\mathrm{area}(\mathcal{Q}) \geq \pi\tilde{\delta}^2 \quad \text{and} \quad \mathrm{per}(\mathcal{Q}) \leq 2\pi\tilde{\Delta} \ . \tag{3}$$

We will soon present an iterative process, referred to as the *Boundary Reconstruction Process (BRP)*, which identifies the $\mathrm{T}^?$ cells. The union of the $\mathrm{T}^?$ cells form the zone $\mathcal{Q}^?$ that contains $\mathcal{Q}$'s boundary $\partial\mathcal{Q} = \{(x, y) \in \mathbb{R}^2 \mid Q(x, y) = 0\}$. In fact, the zone $\mathcal{Q}^?$ is isomorphic to a ring and in particular, it partitions $\mathbb{R}^2 - \mathcal{Q}^?$ to a zone *enclosed* by $\mathcal{Q}^?$ and a zone *outside* $\mathcal{Q}^?$. The cells in the former zone (respectively, latter zone) are subsequently classified as T^+ cells (resp., T^- cells). We shall conclude by bounding the area of $\mathcal{Q}^?$, showing that it is at most an ϵ-fraction of the area of $\mathcal{Q}$.

The main ingredient of BRP is a procedure referred to as the *segment test*. On input segment σ (which may be open or closed in each endpoint), the segment test returns the number of distinct intersection points of $\partial\mathcal{Q}$ and σ. (Since $\mathcal{Q}$ is convex, this number is either 0, 1, or 2.) The segment test is implemented to run in time $O(m^2)$ by employing Sturm's condition of the projection of the polynomial $Q(x, y)$ on σ and by direct calculations of the SINR function in the endpoints of σ. Typically, the segment test will be invoked on segments consisting of edges of the grid G_γ.

We now turn to describe BRP. Informally, the process traverses the boundary of $\mathcal{Q}$ in the clockwise direction and identifies the grid cells that intersect it (with some slack). Let C_s be the grid cell that contains the point s. (We will choose the parameter γ to ensure that $\gamma < \tilde{\delta}/\sqrt{2}$ so that C_s and the three other cells that share the vertex s are fully contained in $\mathcal{Q}$.) BRP begins by identifying the unique cell C_1 north to C_s (C_1 and C_s are in the same grid column) which contains a point of $\partial\mathcal{Q}$ along its west edge. Note that all grid vertices between s and the south west corner of C_1 are in $\mathcal{Q}$, while the north west corner of C_1 and all the grid vertices to its north are not in $\mathcal{Q}$. The computation of C_1 is performed by direct calculations of the SINR function at grid vertices north of s in a binary search fashion, starting with a grid vertex at distance at most $\tilde{\Delta}$ from s, and ending with a grid vertex at distance at least $\tilde{\delta}$ from s, so that the total number of SINR calculations is $O(\log(\tilde{\Delta}/\tilde{\delta}))$.

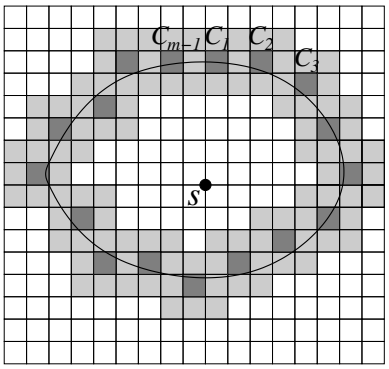

Figure 6: The cell collection $C_1, \ldots, C_{m-1}$ (dark gray) on top of the boundary of $\mathcal{Q}$ (bold curve). The $\mathrm{T}^?$ cells are the union of $C_1, \ldots, C_{m-1}$ and the 8 cells surrounding each one of them (in light gray).

Let q denote the (unique) intersection point of $\partial\mathcal{Q}$ and the west edge of C_1. Consider some continuous (injective) curve function $\phi : [0, 2\pi) \to \partial\mathcal{Q}$ that traverses $\partial\mathcal{Q}$ in the clockwise direction, aligned so that $\phi(0) = q$. For the sake of formality, we extend the domain of ϕ to $[0, \infty)$ by setting $\phi(z) = \phi(z - \lfloor z/(2\pi) \rfloor \cdot 2\pi)$ for every $z > 2\pi$. Let $z_1 = 0$. Given the cell C_{j-1} and the real $z_{j-1} \in [0, 2\pi)$, we define $z_j = \inf\{z > z_{j-1} \mid \phi(z) \notin \sharp C_{j-1}\}$. (Informally, $\phi(z_j)$ is the first point out of $\sharp C_{j-1}$ encountered along a clockwise traversal of $\partial\mathcal{Q}$ that begins at $\phi(z_{j-1})$.)

If $z_j \geq 2\pi$ (which means that the process have completed a full encirclement of $\partial\mathcal{Q}$), then we fix $m = j$ and BRP is over. Assume that $z_j < 2\pi$. If $\phi(z_j) \notin \sharp C_{j-1}$, then the cell C_j is defined to be the cell containing $\phi(z_j)$. Otherwise, the cell C_j is defined to be the cell containing the point $\phi(z_j + \delta)$ for sufficiently small $\delta > 0$. BRP then continues, gradually constructing the collection of $\mathrm{T}^?$ cells, consisting of all cells in the 9-cell of C_j for every $1 \leq j < m$. The choice of cells $C_1, \ldots, C_{m-1}$ is illustrated in Fig. 6.

It should be clarified that from an algorithmic perspective, we do not explicitly compute the real sequence $z_1, \ldots, z_m$, but rather the cell sequence $C_1, \ldots, C_{m-1}$. This is done by $O(1)$ applications of the segment test for every 9-cell involved in the process. Since $\partial\mathcal{Q}$ is a closed curve, and since $\mathcal{Q}$ is convex, these applications are sufficient to identify the grid edges (and vertices) crossed by (or tangent to) $\partial\mathcal{Q}$, and hence to compute the desired cell sequence $C_1, \ldots, C_{m-1}$.

Next, we bound the number of $\mathrm{T}^?$ cells. In every iteration of BRP, we introduce at most 9 new $\mathrm{T}^?$ cells, hence the total number of $\mathrm{T}^?$ cells is at most $9(m-1)$. Recall our choice of reals $z_1, \ldots, z_m$. As $\phi(z_{j-1})$ lies on the boundary of C_{j-1} and $\phi(z_j)$ lies on the boundary of $\sharp C_{j-1}$, we conclude that $\mathrm{dist}(\phi(z_j), \phi(z_{j-1})) \geq \gamma$ for every $1 < j \leq m$. Therefore in each iteration, at least γ units of length are "consumed" from $\mathrm{per}(\mathcal{Q})$. Inequality (3) implies that $m \leq \lceil \mathrm{per}(\mathcal{Q})/\gamma \rceil \leq \lceil 2\pi\tilde{\Delta}/\gamma \rceil$, thus the number of $\mathrm{T}^?$ cells is at most $9(m-1) < 18\pi\tilde{\Delta}/\gamma$. Since the area of each grid cell is γ^2, it follows that the total area of $\mathcal{Q}^?$ (which is the union of the $\mathrm{T}^?$ cells) is smaller than $18\pi\tilde{\Delta}\gamma$. In order to guarantee that $\mathrm{area}(\mathcal{Q}^?) \leq$

$\epsilon \cdot \text{area}(\mathcal{Q})$, we employ inequality (3) once more and demand that $18\pi\tilde{\Delta}\gamma \leq \epsilon\pi\tilde{\delta}^2$. Therefore it is sufficient to fix $\gamma = \frac{\epsilon\tilde{\delta}^2}{18\tilde{\Delta}}$, which means that the number of T$^?$ cells is $O((\tilde{\Delta}/\tilde{\delta})^2\epsilon^{-1})$.

Let Q be the collection of grid columns with at least one T$^?$ cell. Clearly, $|Q| = O((\tilde{\Delta}/\tilde{\delta})^2\epsilon^{-1})$. Each column χ in Q contains at most 6 T$^?$ cells. Consider some cell C in χ which is not a T$^?$ cell. If there is a T$^?$ cell to the north of C and a T$^?$ cell to its south, then C is a T$^+$ cell; otherwise, C is a T$^-$ cell. It follows that the data structure `QDS` can be represented as a vector with one entry per each grid column in Q ($O((\tilde{\Delta}/\tilde{\delta})^2\epsilon^{-1})$ entries altogether), where the entry corresponding to the grid column $\chi \in Q$ stores the T$^?$ cells of χ (at most 6 of them). On input point $p \in \mathbb{R}^2$, we merely have to compute the grid cell to which p belongs and (possibly) look up at the relevant entry of `QDS`.

4.2 Approximate point location queries in the SINR diagram

In this section we explain the relevance of the construction presented in Section 4.1 to ϵ-approximate point location queries in the SINR diagram and establish Theorem 3. Consider some UPN $\langle S, \bar{1}, N, \beta\rangle$, where $S = \{s_0, \ldots, s_{n-1}\}$ and $\beta > 1$ is a constant. Recall that the reception zone $\mathcal{H}_i = \{(x,y) \in \mathbb{R}^2 \mid H_i(x,y) \leq 0\}$ for every $0 \leq i \leq n-1$, where $H_i(x,y)$ is a 2-variate polynomial of degree at most $2n$ that is strictly negative for all internal points (x,y) of $\mathcal{H}_i$ (see Section 2). Assuming that the location of s_i is not shared by any other station (if it is, then $\mathcal{H}_i = \{s_i\}$ and point location queries are answered trivially), we know that s_i is an internal point of $\mathcal{H}_i$. Furthermore, Theorem 1 guarantees that the reception zone $\mathcal{H}_i$ is a bounded convex zone and Theorem 5 provides us with a lower bound $\tilde{\delta}$ on $\delta(s_i, \mathcal{H}_i)$, and an upper bound $\tilde{\Delta}$ on $\Delta(s_i, \mathcal{H}_i)$ such that $\tilde{\Delta}/\tilde{\delta} = O(\sqrt{n})$.

In fact, we can obtain much tighter bounds on $\delta(s_i, \mathcal{H}_i)$ and $\Delta(s_i, \mathcal{H}_i)$. Let r be some positive real and assume that we are promised that $\delta(s_i, \mathcal{H}_i) = O(r)$ and that $\Delta(s_i, \mathcal{H}_i) = \Omega(r)$. Recall that $\Delta(s_i, \mathcal{H}_i)/\delta(s_i, \mathcal{H}_i) = O(1)$, hence both $\delta(s_i, \mathcal{H}_i)$ and $\Delta(s_i, \mathcal{H}_i)$ are $\Theta(r)$. Such a positive real r is found via an iterative binary-search-like process that directly computes the values of the SINR function of s_i at points to the, say, north of s_i, starting with a point at distance $\tilde{\Delta}$ form s_i, and ending with a point at distance at least $\tilde{\delta}$ from s_i. Since $\tilde{\Delta}/\tilde{\delta} = O(\sqrt{n})$, it follows that this process is bound to end within $O(\log n)$ iterations. Each iteration takes $O(n)$ time, thus the improved bounds for $\delta(s_i, \mathcal{H}_i)$ and $\Delta(s_i, \mathcal{H}_i)$ are computed in time $O(n \log n)$.

Given a performance parameter $0 < \epsilon < 1$, we apply the technique of Section 4.1 to $\mathcal{H}_i$ and its corresponding polynomial H_i with the improved bounds on $\delta(s_i, \mathcal{H}_i)$ and $\Delta(s_i, \mathcal{H}_i)$ and construct in time $O(n^2\epsilon^{-1})$ a data structure $\texttt{QDS}_i$ of size $O(\epsilon^{-1})$ that partitions the Euclidean plane to disjoint zones $\mathbb{R}^2 = \mathcal{H}_i^+ \cup \mathcal{H}_i^- \cup \mathcal{H}_i^?$ such that (1) $\mathcal{H}_i^+ \subseteq \mathcal{H}_i$; (2) $\mathcal{H}_i^- \cap \mathcal{H}_i = \emptyset$; and (3) $\mathcal{H}_i^?$ is bounded and its area is at most an ϵ-fraction of $\mathcal{H}_i$. Given a query point $p \in \mathbb{R}^2$, $\texttt{QDS}_i$ answers in constant time whether p is in $\mathcal{H}_i^+$, $\mathcal{H}_i^-$, or $\mathcal{H}_i^?$. (We construct a separate data structure $\texttt{QDS}_i$ for every $0 \leq i \leq n-1$.)

Recall that by Observation 1, point p cannot be in $\mathcal{H}_i$ unless it is closer to s_i than it is to any other station in S. Thus for such a point p there is no need to query the data structure $\texttt{QDS}_j$ for any $j \neq i$. A Voronoi diagram of linear size for the n stations is constructed in $O(n \log n)$ preprocessing time, so that given a query point $p \in \mathbb{R}^2$, we can identify the closest station s_i in time $O(\log n)$ and invoke the appropriate data structure $\texttt{QDS}_i$.

Combining the Voronoi diagram with the data structures $\texttt{QDS}_i$ for all $0 \leq i \leq n-1$, we obtain a data structure `DS` of size $O(n\epsilon^{-1})$, constructed in $O(n^3\epsilon^{-1})$ preprocessing time, that decides in time $O(\log n)$ whether the query point p is in $\mathcal{H}_i^+$ for some i, in $\mathcal{H}_i^?$ for some i, or neither, which means that $p \in \mathcal{H}^- = \bigcap_{i=0}^{n-1} \mathcal{H}_i^-$. Theorem 3 follows.

5. REFERENCES

[1] P.K. Agarwal and J. Erickson. Geometric range searching and its relatives. In *Advances in Discrete and Computational Geometry*, pages 1–56. American Mathematical Society, 1999.

[2] A. Aggarwal, M. Hansen, and F.T. Leighton. Solving query-retrieval problems by compacting voronoi diagrams. In *STOC*, pages 331–340, 1990.

[3] S. Basu, R. Pollack, and M.-F. Roy. *Algorithms in Real Algebraic Geometry*. Springer-Verlag, 2003.

[4] U. Black. *Mobile and Wireless Networks*. Prentice Hall, 1996.

[5] B. Chazelle, H. Edelsbrunner, L.J. Guibas, and M. Sharir. A singly exponential stratification scheme for real semi-algebraic varieties and its applications. *Theor. Comput. Sci.*, 84:77–105, 1991.

[6] B.N. Clark, C.J. Colbourn, and D.S. Johnson. Unit disk graphs. *Discrete Math.*, 86:165–177, 1990.

[7] M. de Berg, O. Cheong, M. van Kreveld, and M. Overmars. *Computational Geometry: Algorithms and Applications*. Springer-Verlag, 2008.

[8] O. Goussevskaia, Y.A. Oswald, and R. Wattenhofer. Complexity in geometric SINR. In *Proc. 8th ACM Int. Symp. on Mobile Ad Hoc Networking and Computing (MobiHoc)*, pages 100–109, 2007.

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

[10] F. Kuhn, R. Wattenhofer, and A. Zollinger. Ad-Hoc Networks Beyond Unit Disk Graphs. In *1st ACM Workshop on Foundations of Mobile Computing (DIALM-POMC)*, 2003.

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

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

[13] T. Moscibroda, R. Wattenhofer, and A. Zollinger. Topology control meets SINR: the scheduling complexity of arbitrary topologies. In *Proc. 7th ACM Int. Symp. on Mobile Ad Hoc Networking and Computing (MobiHoc)*, pages 310–321, 2006.

[14] K. Pahlavan and A. Levesque. *Wireless information networks*. Wiley, 1995.

[15] T. S. Rappaport. *Wireless Communications-Principles and Practice*. Prentice Hall, 1996.

[16] P. von Rickenbach, S. Schmid, R. Wattenhofer, and A. Zollinger. A robust interference model for wireless ad-hoc networks. In *Proc. 19th Int. Parallel and Distributed Processing Symp.*, 2005.