Batched Point Location in SINR Diagrams via Algebraic Tools

BORIS ARONOV, New York University
MATTHEW J. KATZ, Ben-Gurion University

41

The *SINR* (Signal to Interference plus Noise Ratio) model for the quality of wireless connections has been the subject of extensive recent study. It attempts to predict whether a particular transmitter is heard at a specific location, in a setting consisting of n simultaneous transmitters and background noise. The SINR model gives rise to a natural geometric object, the *SINR diagram*, which partitions the space into n regions where each of the transmitters can be heard and the remaining space where no transmitter can be heard.

Efficient *point location* in the SINR diagram, i.e., being able to build a data structure that facilitates determining, for a query point, whether any transmitter is heard there, and if so, which one, has been recently investigated in several articles. These planar data structures are constructed in time at least quadratic in n and support logarithmic-time approximate queries. Moreover, the performance of some of the proposed structures depends strongly not only on the number n of transmitters and on the approximation parameter ϵ , but also on some geometric parameters that cannot be bounded *a priori* as a function of n or ϵ .

In this article, we address the question of *batched* point location queries, i.e., answering many queries simultaneously. Specifically, in one dimension, we can answer n queries *exactly* in amortized polylogarithmic time per query, while in the plane we can do it approximately.

In another result, we show how to answer n^2 queries *exactly* in amortized polylogarithmic time per query, assuming the queries are located on a possibly non-uniform $n \times n$ grid.

All these results can handle *arbitrary* power assignments to the transmitters. Moreover, the amortized query time in these results depends only on n and ϵ .

We also show how to speed up the preprocessing in a previously proposed point-location structure in SINR diagram for uniform-power sites, by almost a full order of magnitude. For this, we obtain results on the sensitivity of the reception regions to slight changes in the reception threshold, which are of independent interest.

Finally, these results demonstrate the (so far underutilized) power of combining algebraic tools with those of computational geometry and other fields.

CCS Concepts: • **Theory of computation** → **Computational geometry**; • **Networks**; • **Computing methodologies** → **Symbolic and algebraic algorithms**;

Additional Key Words and Phrases: Wireless networks, SINR model, SINR diagram, batched point location, algebraic methods, fast polynomial multiplication, fast polynomial multipoint evaluation, range searching

An earlier version of this article (excluding Sections 4 and 5) was presented at ICALP'15 [3].

Work on this article by B.A. has been partially supported by the National Science Foundation Grants No. CCF-11-17336, No. CCF-12-18791, and No. CCF-15-40656, and by Grant No. 2014/170 from the US-Israel Binational Science Foundation.

Work on this article by M.K. has been supported by Grant No. 2014/170 from the US-Israel Binational Science Foundation, and by Grants No. 1045/10 and No. 1884/16 from the Israel Science Foundation.

Authors' addresses: B. Aronov, Department of Computer Science and Engineering, Tandon School of Engineering, New York University, 2 MetroTech Center, Brooklyn, NY 11201, USA; email: boris.aronov@nyu.edu; M. J. Katz, Department of Computer Science, Ben-Gurion University of the Negev, P.O.B. 653, Beer-Sheva 8410501, Israel; email: matya@cs.bgu.ac.il. Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2018 ACM 1549-6325/2018/08-ART41 \$15.00

<https://doi.org/10.1145/3209678>

ACM Reference format:

Boris Aronov and Matthew J. Katz. 2018. Batched Point Location in SINR Diagrams via Algebraic Tools. *ACM Trans. Algorithms* 14, 4, Article 41 (August 2018), 29 pages.
<https://doi.org/10.1145/3209678>

1 INTRODUCTION

The *SINR (Signal to Interference plus Noise Ratio) model* attempts to more realistically predict whether a wireless transmission is received successfully, in a setting consisting of multiple simultaneous transmitters in the presence of background noise. In particular, it takes into account the attenuation of electromagnetic signals. The SINR model has been explored extensively in the literature; see Reference [22] for a survey.

Let $\mathcal{S} = \{s_1, \dots, s_n\}$ be a set of n points in the plane representing the locations of n transmitters. Let $p_i > 0$ be the transmission power of transmitter s_i , $i = 1, \dots, n$. In the *SINR model*, a receiver located at point q is able to receive the signal transmitted by s_i if the following inequality holds:

$$\frac{\frac{p_i}{|q-s_i|^\alpha}}{\sum_{j \neq i} \frac{p_j}{|q-s_j|^\alpha} + N} \geq \beta,$$

where $|a - b|$ denotes the Euclidean distance between points a and b , and $\alpha > 0$, $\beta > 1$, and $N > 0$ are given constants; N represents the background noise. This inequality is also called the *SINR inequality*, and when it holds, we say that q *receives* (or *hears*) s_i ; we refer to the left-hand side of the inequality as *SIN ratio* (for receiver q with respect to transmitter s_i) and to the right-hand side as the *reception threshold*.

Notice that, since $\beta > 1$, a necessary condition for q to receive s_i is that $p_i/|q - s_i|^\alpha > p_j/|q - s_j|^\alpha$, for any $j \neq i$. In particular, in the *uniform power setting* where $p_1 = p_2 = \dots = p_n$, a necessary condition for q to receive s_i is that s_i is the closest to q among the transmitters in $\mathcal{S}$. This simple observation implies that, for any point q in the plane, either exactly one of the transmitters is received by q or none of them is. Thus, one can partition the plane into n not necessarily connected reception regions R_i , one per transmitter in $\mathcal{S}$, plus an additional region R_0 consisting of all points where none of the transmitters is received. This partition is called the *SINR diagram* of $\mathcal{S}$ [9]. Consider the *multiplicatively weighted Voronoi diagram* D of $\mathcal{S}$ in which the region V_i associated with s_i consists of all points q in the plane for which $\frac{1}{\sqrt[\alpha]{p_i}}|q - s_i| < \frac{1}{\sqrt[\alpha]{p_j}}|q - s_j|$, for any $j \neq i$ (see Section B.2). Then $R_i \subset V_i$.

In a seminal article, Avin et al. [9] studied properties of SINR diagrams, focusing on the uniform power setting. Their main result is that in this setting the reception regions R_i are convex and fat. (Here, R_i is *fat* if the ratio between the radii of the smallest disk centered at s_i containing R_i and the largest disk centered at s_i contained in R_i is bounded by some constant.) In the non-uniform power setting, on the other hand, the reception regions are not necessarily connected, and their connected components are not necessarily convex or fat. In fact, they may contain holes [20].

A natural question that one may ask is: “Given a point q in the plane, does q receive one of the transmitters in $\mathcal{S}$, and if yes which one?” Or equivalently: “Which region of the SINR diagram does q belong to?” The latter question is referred to as a *point-location query* in the SINR diagram of $\mathcal{S}$. We can answer it in linear time by first finding the sole candidate, s_i , as the transmitter for which the ratio $\frac{1}{\sqrt[\alpha]{p_i}}|q - s_i|$ is minimum, and then evaluating the SIN ratio and comparing it to β . To facilitate multiple queries, one may want to build a data structure that can guarantee faster response. We can expedite the first step by constructing the appropriate Voronoi diagram $D = D(\mathcal{S})$ together with a point-location structure, so that the sole candidate transmitter for a

point q can be found in $O(\log n)$ time; see, e.g., Reference [7]. However, the boundary of the region R_i is described by a degree- $\Theta(n)$ algebraic curve; it seems difficult (impossible, in general?) to build a data structure that can quickly determine the side of the curve a given point lies on. The answer is not even obvious in one dimension (where the transmitters and potential receivers all lie on a line), as there R_i is a collection of intervals delimited by roots of a polynomial of degree $\Theta(n)$.

The problem has been approached by constructing data structures for *approximate* point location in SINR diagrams. All approaches use essentially the same logic: first find the sole candidate s_i that the query point q may hear and then approximately locate q in R_i . This is done by constructing two sets R_i^+, R_i^- such that $R_i^+ \subset R_i \subset R_i^- \subset V_i$,¹ and preprocessing them for point location. In the region R_i^+ reception of s_i is guaranteed, so if $q \in R_i^+$, return “can hear s_i .” Outside of R_i^- one cannot hear s_i , so if $q \in V_i$ does not belong to R_i^- , return “cannot hear anything.” The set $R_i^- \setminus R_i^+$ is where the approximation occurs: s_i may or may not be heard there, so if $q \in R_i^-$ but $q \notin R_i^+$, return “may or may not hear s_i .”

Two different notions of approximation have appeared in the literature. In the first [9, 20], it is guaranteed that the uncertain answer is only given infrequently, namely that $\text{area}(R_i^- \setminus R_i^+) \leq \varepsilon \cdot \text{area}(R_i)$, for a prespecified parameter $\varepsilon > 0$. In the second [20], it is promised that the SIN ratio for every point in $R_i^- \setminus R_i^+$ lies within $[C_1\beta, C_2\beta]$ for suitable constants C_1, C_2 with $0 < C_1 < 1$, $C_2 > 1$. We show below (see end of Section 4) how the two notions are related, at least for the uniform power case.

We now briefly summarize previous work. Observing the difficulty of answering point-location queries exactly, Avin et al. [9] resorted to approximate query answers in the *uniform power* setting. Given an $\varepsilon > 0$, they build a data structure in total time² $O(n^2/\varepsilon)$ and space $O(n/\varepsilon)$ that can be wrong only in a region of area $\varepsilon \cdot \text{area}(R_i)$ for each s_i (i.e., approximation of the first type described above). It supports logarithmic-time queries.

In a subsequent paper, Kantor et al. [20] studied properties of SINR diagrams in the *non-uniform power* setting. After revealing several interesting and useful properties, such as that the reception regions in the $(d+1)$ -dimensional SINR diagram of a d -dimensional scene are connected, they present several solutions to the problem of efficiently answering point-location queries. One of them uses the second type of approximation, with $C_1 = (1 - \varepsilon)^{2\alpha}$ and $C_2 = (1 + \varepsilon)^{2\alpha}$, for a pre-specified $\varepsilon > 0$. Queries can be performed in time $O(\log(n \cdot \varphi/\varepsilon))$, where φ is an upper bound on the fatness parameters of the reception regions (which cannot be bounded as a function of n or ε). The size of this data structure is $O(n \cdot \varphi'/\varepsilon^2)$ and its construction time is $O(n^2 \cdot \varphi'/\varepsilon^2)$, where $\varphi' > \varphi^2$ is some function of the fatness parameters of the reception regions.

Although highly non-trivial, the known results for point location in the SINR model are unsatisfactory, in that they suffer from very large preprocessing times. Moreover, in the non-uniform setting, the bounds include geometric parameters such as φ and φ' above, which cannot be bounded as a function of n or ε . In this article, we focus on *batched* point location in the SINR model. That is, given a set Q of m query points, determine for each point $q \in Q$ whether it receives one of the transmitters in S , and if yes, which one. Often the set of query points is known in advance, for example, in the planning stage of a wireless network or when examining an existing network. In these cases, one would like to exploit the additional information to speed up query processing. We achieve this goal in the SINR model; that is, we devise efficient approximation and exact algorithms for batched point location in various settings. Our algorithms use a novel combination of geometric data structures and tools from computer algebra for multipoint evaluation, interpolation, and

¹Notice that we have not followed the original notation in the literature, for consistency with our notation below.

²See below for discussion of an apparent discrepancy between the claimed preprocessing time of their data structure and the analysis presented in their article; the actual preprocessing time seems to be $O(n^2(\log n + 1/\varepsilon))$.

fast multiplication of polynomials and rational functions. For example, consider a one-dimensional batched point location where $m = n$ and power is non-uniform. We can answer *exactly* a point-location query in amortized time $O(\log^2 n \log \log n)$. Considering the same problem in the plane, for any $\varepsilon > 0$, we can approximately answer a query in amortized time polylogarithmic in n and ε , as opposed to the result of Kantor et al. [20] mentioned above in which the bounds depend on additional geometric parameters which cannot be bounded as a function of n or ε . See Section 1.3 for a detailed list of our results.

1.1 Related Work

The articles most relevant to ours are those by Avin et al. [9] and Kantor et al. [20] discussed above. Avin et al. [8] also considered the problem of handling queries of the following form (in the uniform-power setting): Given a transmitter s_i and query point q , does q receive s_i by successively applying interference cancellation? (Interference cancellation is a technology that enables a point q to receive a transmitter s , even if s 's signal is not the strongest one received at q ; see Reference [8] for further details.)

Gupta and Kumar [15] initiated an extensive study of the *maximum capacity* and *scheduling* problems in the SINR model. Given a set L of sender-receiver pairs (i.e., directional links), the *maximum capacity* problem is to find a *feasible* subset of L of maximum cardinality, where $L' \subseteq L$ is *feasible* if, when only the senders of the links in L' are active, each of the links in L' is feasible according to the SINR inequality. The *scheduling* problem is to partition L into a minimum number of feasible subsets (i.e., rounds). We mention several articles and results dealing with the maximum capacity and scheduling problems. Goussevskaia et al. [14] showed that both problems are NP-complete, even in the uniform power setting. Goussevskaia et al. [13], Halldórsson and Wattenhofer [18], and Wan et al. [29] gave constant-factor approximation algorithms for the maximum-capacity problem yielding an $O(\log n)$ -approximation algorithm for the scheduling problem, assuming uniform power. In Reference [13], they note that their $O(1)$ -approximation algorithm also applies to the case where the ratio between the maximum and minimum power is bounded by a constant and for the case where the number of different power levels is constant. More recently, Halldórsson and Mitra [17] have considered the case of oblivious power. This is a special case of non-uniform power where the power of a link is a simple function of the link's length. They gave an $O(1)$ -approximation algorithm for the maximum capacity problem, yielding an $O(\log n)$ -approximation algorithm for scheduling. Finally, the version where one assigns powers to the senders (i.e., with power control) has also been studied, see, e.g., [2, 16, 17, 21, 24]. For additional work on the subject, see Reference [4].

1.2 Our Tools and Goals

Besides making progress on the actual problems being considered here, we view this work as another demonstration of what we hope to be a developing trend of combining tools from the computer algebra world with those of computational geometry and other fields. Several relatively recent representatives of such synergy show examples of seemingly impossible speedups in geometric algorithms by expressing a subproblem in algebraic terms [1, 23]. The algebraic tools themselves are mostly classical ones, such as Fast Fourier Transform (FFT), fast polynomial multiplication, multipoint evaluation, and interpolation [10, 28]; see Appendix A for details. We combine them with only slightly newer tools from computational geometry, such as Voronoi diagrams, point-location structures in the plane, fast exact and approximate nearest-neighbor query data structures, and range searching data structures [11]; refer to Appendix B. One very recent result we need is that of Har-Peled and Kumar [19] that, as a special case, allows one to build a compact data structure for approximately answering multiplicatively weighted nearest-neighbor queries

in the plane; the exact version appears to require building the classical multiplicatively weighted Voronoi diagram, which is a quadratic-size structure.

We hope that the current work will lead to further productive collaborations between computational geometry and computer algebra.

After a preliminary version of our work has appeared [3], we became aware of the work of Ziegler [30] on approximate computation of forces and potentials in classical electrostatic and gravitational fields of multiple point sources. We found that the fundamental tools and techniques applied by Ziegler are quite similar to those applied by us in the first part of this article (i.e., in Sections 2 and 3); however, the context of SINR networks gives rise to several distinct challenges that distinguish between our work and his.

1.3 Our Results

We now summarize our main results. We use O^* notation to suppress logarithmic factors and O_ε to denote polynomial dependence on $1/\varepsilon$, where $\varepsilon > 0$ is the approximation parameter. We present algorithms for both the uniform-power and non-uniform-power settings, where the algorithms of the former type are usually somewhat simpler. Throughout the article, we assume that α is an integer. Our algorithms can handle both even and odd values of α , excluding the results in Sections 3.2 and 3.3, which require that α be even. Sometimes, the running time in the odd case is slightly slower than in the even case.

- In one dimension, we can perform n queries among n transmitters exactly in $O^*(n)$ total time; see Section 2.
- In two dimensions, we can perform n queries among n transmitters approximately in $O_\varepsilon^*(n)$ total time; see Section 3.4.
- If the underlying transmitters are using directional antennas, then we can perform n queries among n transmitters approximately in $O_\varepsilon^*(n^{4/3})$ total time, after modifying the SINR inequality to suit this case; see Section 3.5.
- We can also facilitate exact batch queries when queries or transmitters form a grid; see Sections 3.2 and 3.3 for the exact statements.
- We apply our methods to speed up the preprocessing stage of the uniform-power point-location algorithm of Avin et al. [9] mentioned above to $O_\varepsilon^*(n)$; see Section 5.
- The correctness of our faster construction is based on a couple of technical lemmas revealing interesting structural properties of SINR regions; see Section 4.

2 BATCHED POINT LOCATION ON THE LINE

In this section, $\mathcal{S}$ is a set of $n \geq 3$ point transmitters and $\mathcal{Q}$ is a set of m query points, both on the line. We first consider the *uniform-power version* of the problem, where each transmitter has transmission power 1 (i.e., $p_1 = \dots = p_n = 1$), and then extend the approach to the arbitrary power version.

2.1 Uniform Power

A query point q receives s_i if and only if

$$\frac{1}{|q-s_i|^\alpha} \geq \beta.$$

$$\frac{1}{\sum_{j \neq i} \frac{1}{|q-s_j|^\alpha} + N} \geq \beta.$$

Recall that, since $\beta > 1$, if q receives one of the transmitters, then it must be the transmitter that is closest to it; we call it the *candidate* transmitter for q and denote it by $s(q) = s(q, \mathcal{S})$.

Next, we define a univariate function f as

$$f(q) := \sum_{j=1}^n \frac{1}{|q - s_j|^\alpha}.$$

Then, q can hear its candidate transmitter $s(q)$ if and only if

$$E(q) := \frac{\frac{1}{|q-s(q)|^\alpha}}{f(q) - \frac{1}{|q-s(q)|^\alpha} + N} \geq \beta.$$

We distinguish between the case where α is even (Theorem 2.1) and the case where it is odd (Theorem 2.2).

THEOREM 2.1. *For any fixed positive even integer α , given a set $\mathcal{S}$ of transmitters (all of power 1) and a set $\mathcal{Q}$ of receivers, of sizes n and m , respectively, we can determine which, if any, transmitter is received by each receiver in total time $O((n+m) \log^2 n \log \log n)$.*

PROOF. As pointed out above, a receiver q can receive only the closest transmitter $s(q)$, if any, as the SINR inequality implies $\frac{1}{|q-s(q)|^\alpha} > \frac{1}{|q-s|^\alpha}$ for any $s \neq s(q)$, or equivalently, $|q - s(q)| < |q - s|$. So, as a first step, we identify the closest transmitter for each receiver, which can be done, for example, by sorting $\mathcal{S}$, and using binary search for each receiver, in total time $O((m+n) \log n)$. Moreover, we can compute the term $\frac{1}{|q-s(q)|^\alpha}$, for each $q \in \mathcal{Q}$, in the same amount of time.

Observe that since α is even, we can view f as a sum of n low-degree fractional functions of a single real variable q (i.e., the fractional functions $\frac{1}{(q-s_j)^\alpha}$, for $j = 1, \dots, n$), so, according to Corollary A.4, we can now evaluate f on all points of $\mathcal{Q}$ simultaneously in time $O((n+m) \log^2 n \log \log n)$.

In $O(m)$ additional operations, we can evaluate the expressions $E(q_1), \dots, E(q_m)$ and determine for which receivers the SINR inequality holds, so that the signal is actually received.

Computing and evaluating the fraction dominates the computation cost, so the total running time is $O((n+m) \log^2 n \log \log n)$. $\square$

THEOREM 2.2. *For any fixed positive odd integer α , given a set $\mathcal{S}$ of transmitters (all of power 1) and a set $\mathcal{Q}$ of receivers, of sizes n and m , respectively, we can determine which, if any, transmitter is received by each receiver in total time $O((n+m) \log^3 n \log \log n)$.*

PROOF. The proof is identical to that of Theorem 2.1 with one exception. Again, we need to evaluate f for each receiver q_i , but now, unlike the previous case, we cannot view f as a sum of proper rational functions, since, if we do so, the sign of some of the terms might be negative (rather than positive), depending on whether q_i lies to the left or to the right of s_j . To overcome this difficulty, we add another layer of processing.

We sort the transmitters s_i by their position along the line and build a balanced binary search tree with transmitters as leaves, each node corresponding to a set of transmitters. We associate with each node of the tree two sets of receivers, a *right set* and a *left set*, where each receiver in the right set (resp., left set) of a node v lies to the right (resp., to the left) of all transmitters associated with v . We do this by searching for each receiver q_i in the tree. The set of all transmitters lying left of q_i is naturally expressed as a disjoint union of $O(\log n)$ subsets corresponding to the nodes of the search tree, and we add q_i to the right set of each of these nodes. Similarly, the set of all transmitters lying right of q_i is naturally expressed as a disjoint union of $O(\log n)$ subsets corresponding to the nodes of the search tree, and we add q_i to the left set of each of these nodes.

Now, for each node of the tree, we apply Corollary A.4 twice; once for its right set of receivers and the partial sum of fractional functions of the form $\frac{1}{(q-s_j)^\alpha}$ associated with it, and once for its

left set of receivers and the partial sum of fractional functions of the form $\frac{1}{(s_j - q)^\alpha}$ associated with it.

We now proceed as in Theorem 2.1, where $f(q_i)$ is obtained as a sum of a logarithmic number of terms that have already been computed. The total cost is $O(\log n)$ times that of Theorem 2.1, as each transmitter and receiver appears in at most a logarithmic number of computations. $\square$

Notice that the above “trick” is a simpler, one-dimensional version of the range-tree machinery that we use in Section 3.4.

2.2 Arbitrary Power

We proceed in a similar manner, except for the construction of the multiplicatively weighted Voronoi diagram on a line, which is more subtle.

Observe that, if q hears any of the transmitters in $\mathcal{S}$, then it is the *candidate* transmitter $s(q)$, for which the expression $p(s)/|q - s|^\alpha$ is maximum, over all $s \in \mathcal{S}$; $p(s)$ is the power of transmitter s . Thus, we find for each $q \in \mathcal{Q}$ its corresponding transmitter, by computing the appropriate weighted Voronoi diagram and querying it. This can be done in $O((n + m) \log n)$ time (see Fact B.5). Next, we generalize the function f defined above to mean

$$f(q) := \sum_{j=1}^n \frac{p_j}{|q - s_j|^\alpha}.$$

Now, q hears $s(q)$ if and only if

$$E(q) := \frac{\frac{p(s(q))}{|q - s(q)|^\alpha}}{f(q) - \frac{p(s(q))}{|q - s(q)|^\alpha} + N} \geq \beta,$$

and we have

THEOREM 2.3. *For any fixed positive integer α , given a set $\mathcal{S}$ of transmitters (not necessarily all of the same power) and a set $\mathcal{Q}$ of receivers, of sizes n and m , respectively, we can determine which, if any, transmitter is received by each receiver in total time $O((n + m) \log^2 n \log \log n)$, if α is even, and $O((n + m) \log^3 n \log \log n)$, if α is odd.*

PROOF. Assume α is even. For odd α , the proof is identical, except for an additional layer of processing, which is described in the proof of Theorem 2.2 and which adds a logarithmic factor to the total running time. Instead of looking for the nearest transmitter to q , we need to look for the transmitter maximizing $p_i/|q - s_i|^\alpha$, which is the same as minimizing $|q - s_i|/p_i^{1/\alpha}$. This corresponds to the one-dimensional *multiplicatively weighted Voronoi diagram* of $\mathcal{S}$ with weights $p_i^{1/\alpha}$, which, as per Fact B.5 and surrounding discussion, has complexity $O(n)$ and can be constructed in $O(n \log n)$ time; see References [5] and [27, Theorem 6.1]. We once again use binary search to identify, for each receiver q_i , the only possible transmitter $s(q_i)$ that it may be able to hear, compute the term $p(s(q_i))/|q_i - s(q_i)|^\alpha$ needed to finish computing $E(q_i)$, and thereby check the reception.

The rest of the algorithm proceeds as in Theorem 2.1: the function f depends on the power of individual transmitters, but it is still a sum of n low-degree univariate fractions and the remainder of the analysis follows unchanged. $\square$

3 BATCHED POINT LOCATION IN THE PLANE

In this section, $\mathcal{S} = \{s_i\}$ is a set of n point transmitters in the plane. We consider three versions of batched point location, where in the first two the answers we obtain are exactly correct, while in the third one the answer to a query q may be either YES (meaning that q receives s), NO (meaning

that q does not receive any transmitter), or *MAYBE* (meaning that q may or may not be receiving some transmitter; the SIN ratio is too close to β and we are unable to decide quickly whether it is above or below β).

Specifically, we consider the following three versions of batched point location. In the first version, we assume that the *transmitters* form an $\sqrt{n} \times \sqrt{n}$ non-uniform grid and that each transmitter has power 1. We show how to solve a *single* point-location query in this setting in $O(\sqrt{n} \log^2 n \log \log n)$ (rather than linear) time; refer to Section 3.2. In the second version (Section 3.3), we assume that the *receivers* form an $n \times n$ non-uniform grid, but the n transmitters, on the other hand, are located anywhere in the plane. Moreover, we allow arbitrary transmission powers. We show how to answer the n^2 queries in near-quadratic (rather than cubic) time.

Finally, in the third version (Section 3.4), we do not make any assumptions on the location of the devices (either transmitters or receivers). As a result of this, we might not be able to give a definite answer in borderline instances. Specifically, given n transmitters and m receivers, we compute (in total time near-linear in $n + m$), for each receiver q , its unique candidate transmitter s and a value $\tilde{E}(q)$, such that (a) if $\tilde{E}(q)$ is sufficiently greater than β , then q surely receives s , (b) if $\tilde{E}(q)$ is sufficiently smaller than β , then q surely does not receive s , and (c) otherwise, $\tilde{E}(q)$ lies in the *uncertainty interval*, and q may or may not receive some transmitter. We first present a solution for which the uncertainty interval is $[2^{-\alpha/2}\beta, 2^{\alpha/2}\beta]$, i.e., a constant-factor approximation. We then generalize it so that the uncertainty region is $[(1 - \varepsilon)\beta, (1 + \varepsilon)\beta]$, for any $\varepsilon > 0$, i.e., a PTAS. We consider both the uniform- and arbitrary-power settings.

3.1 General Discussion

Once again, the SINR inequality determines which, if any, of the transmitters $s \in S$ can be heard by a receiver at point q and the only candidate transmitter $s(q)$ is the one that minimizes $|q - s|/p^{1/\alpha}$ among all transmitters s with corresponding power p . In the uniform-power case, this means the transmitter closest to q in Euclidean distance, and the matching space decomposition is the Euclidean Voronoi diagram which can be constructed in $O(n \log n)$ time (see Section B.2), where $n = |S|$. In the non-uniform-power case, this corresponds to the multiplicatively weighted Voronoi diagram in the plane, which is a structure of worst-case complexity $\Theta(n^2)$ that can be constructed in time $O(n^2)$; see Reference [6].

Once again we define the function $f(q)$, which represents the total signal strength at q from *all* transmitters, and express the decision of whether the transmitter $s(q)$ is received at q by computing $E(q)$ from $f(q)$ and $s(q)$ and comparing it with β . The difference from the one-dimensional case is that $f(q)$ is now a sum of low-degree *bivariate* fractions, with the two variables being the coordinates of q .

In all cases, the goal is to evaluate $f(q)$, for each receiver q , and to identify the suitable candidate transmitter $s(q)$, faster than by brute force. Given this information, the decision can be made in constant time per receiver.

We now present each two-dimensional problem in turn.

3.2 Transmitters on a Grid

In this version, we assume that the transmitters of S form an $\sqrt{n} \times \sqrt{n}$ non-uniform grid and have uniform power. We show how to answer a *single* arbitrary point-location query in roughly $\sqrt{n}$, rather than linear, time. We assume that $S = X \times Y$, where X and Y are two sets of $\sqrt{n}$ numbers (coordinates) each. We start by sorting X and Y .

As mentioned above, the problem reduces to computing $f(q)$ and identifying $s(q)$ for the query q . The latter is not difficult: since all powers are the same, it is sufficient to identify the point in $X \times Y$

closest to q . It is easy to check that the closest point is always one of the (at most) four corners of the grid cell containing q and can be found by one binary search on X and one on Y and comparing each of the at most four corresponding distances to q .

So we focus on the computation of $f(q)$. We rewrite it as

$$f(q; X; Y) = \sum_{x \in X} \sum_{y \in Y} \frac{1}{|q - (x, y)|^\alpha} = \sum_{x \in X} g_Y(q; x),$$

where we view

$$g_Y(q; x) := \sum_{y \in Y} \frac{1}{|q - (x, y)|^\alpha}$$

as a sum of fractional functions of x only, having substituted the actual values for the numbers of Y . Using Corollary A.4, we can sum these fractions symbolically and then evaluate the sum at the $\sqrt{n}$ distinct values $x_1, x_2, \dots, x_{\sqrt{n}}$ of X , obtaining $g_Y(q; x_1), g_Y(q; x_2), \dots, g_Y(q; x_{\sqrt{n}})$ and thereby $f(q)$, in $O(\sqrt{n} \log^2 n \log \log n)$ operations, which dominates the running time.

THEOREM 3.1. *For any fixed positive even integer α , given a set S of n transmitters (all of power 1) forming a $\sqrt{n} \times \sqrt{n}$ non-uniform grid and a receiver q , we can determine which, if any, transmitter, is received by q in time $O(\sqrt{n} \log^2 n \log \log n)$.*

Notes. (a) For a $k \times l$ grid, with $l \leq k$, this would take $O((k + l) \log^2 l \log \log l) + O((k + l) \log(k + l)) = O(k \log^2 l \log \log l + k \log k)$ time.

(b) Odd integer α cannot be handled using these methods, as $|q - (x, y)|^\alpha$ is not a polynomial (informally, it involves square roots).

3.3 Receivers on a Grid

In this version, we assume that the receivers of Q form an $n \times n$ non-uniform grid and that α is an even integer. The n transmitters, on the other hand, are located anywhere in the plane. Moreover, we allow arbitrary transmission powers. We show how to answer the n^2 queries in near-quadratic (rather than cubic) time.

In this case, we need to evaluate $f(q)$, which is a sum of n bivariate low-degree fractions at all points q_{ij} of a two-dimensional, possibly non-uniform grid $X \times Y$.³ This is taken care of in $O(n^2 \log^2 n \log \log n)$ time by Corollary A.8.

The only missing ingredient is identifying the candidate transmitter $s(q)$ for each $q \in Q$. This can be done, for example, by computing the weighted Voronoi diagram, preprocessing it for point location, and executing n^2 such queries, one for each $q \in Q$.

It turns out that the following alternative may be simpler to implement: Observe that on each line $x = x_j$, the functions $|q - s_i|/p_i^{1/\alpha}$ behave very similarly to the univariate case in Section 2.2 above, so the problem can be solved in $O(n \log n)$ time per line, for a total of $O(n^2 \log n)$ time; refer to Fact B.6.

Combining the computation of $f(q)$ over all points of $X \times Y$ and identification of $s(q)$, we obtain the claimed result in $O(n^2 \log^2 n \log \log n)$ time.

THEOREM 3.2. *For any fixed positive even integer α , given a set S of n transmitters (not necessarily all of the same power), and a set Q of n^2 receivers forming an $n \times n$ non-uniform grid, we can determine which, if any, transmitter, is received by each receiver in total time $O(n^2 \log^2 n \log \log n)$.*

Notes. (a) There are two obstacles to speeding this up for a smaller number of receivers: (i) An explicit representation of $f(q)$ has $\Theta(n^2)$ coefficients and thus cannot be processed in subquadratic

³Slanted and sheared grids can be handled by a simple extension of this observation; we omit the easy details.

time. (ii) It seems harder to identify the candidate transmitter in the non-uniform power case; the first solution builds a quadratic-size space decomposition and the second requires linear time for each set of collinear points; both result in near-quadratic performance. A new idea is needed.

(b) Odd integer α cannot be handled using these methods, as $|q - (x, y)|^\alpha$ is not a polynomial (informally, it involves square roots of polynomials).

(c) If we do not require the n^2 receivers to lie on an $n \times n$ grid, we can still obtain a subcubic solution by applying a different algebraic technique. Indeed, according to Fact A.9, we can determine which, if any, transmitter, is received by each receiver in total time $O(n^{1+\omega_2/2+\varepsilon})$, for any $\varepsilon > 0$, where $\omega_2 < 3.334$ is a constant related to the efficiency of matrix multiplication.

3.4 Approximating the General Case

We now abandon the ambition to get exact answers and aim for an approximation algorithm, in the sense we will make precise below. Again, $\mathcal{S} = \{s_i\}$ is the set of n transmitters, with each s_i a point in the plane with power p_i ; similarly $\mathcal{Q} = \{q_j\}$ is the set of m receivers (queries), where a generic receiver is $q = (q_x, q_y)$.

For a query point q and a transmitter $s = (s_x, s_y)$ of power p , set $l(q, s) = \max\{|q_x - s_x|, |q_y - s_y|\}$; in other words, $l(q, s)$ is the L^∞ distance between points q and s . In complete analogy to our previous approach, put

$$\tilde{f}(q) := \sum_{i=1}^n \frac{p_i}{l(q, s_i)^\alpha} \quad \text{and} \quad \tilde{E}(q) := \frac{\frac{p}{l(q, s)^\alpha}}{\tilde{f}(q) - \frac{p}{l(q, s)^\alpha} + N}.$$

What is the significance of the quantity $\tilde{E}(q)$? Since for any two points s, q , $l(q, s) \leq |q - s| \leq \sqrt{2}l(q, s)$,

$$2^{-\alpha/2} \frac{p_j}{l(q, s_j)^\alpha} \leq \frac{p_j}{|q - s_j|^\alpha} \leq \frac{p_j}{l(q, s_j)^\alpha},$$

so $2^{-\alpha/2}\tilde{f}(q) \leq f(q) \leq \tilde{f}(q)$, and therefore $2^{-\alpha/2}\tilde{E}(q) \leq E(q) \leq 2^{\alpha/2}\tilde{E}(q)$. Informally, $\tilde{E}(q)$ is “pretty close” to $E(q)$.

This suggests an approximation strategy that begins by computing $\tilde{E}(q)$ instead of $E(q)$. If $\tilde{E}(q) \geq 2^{\alpha/2}\beta$, we know that $E(q) \geq \beta$ and the signal from the unique candidate transmitter $s(q)$ is received (YES). If $\tilde{E}(q) < 2^{-\alpha/2}\beta$, then $E(q) < \beta$ and the signal from $s(q)$ is *not* received and therefore no signal is received by q (NO). For intermediate values of $\tilde{E}(q)$, we cannot definitely determine whether $s(q)$ ’s signal is received at q (MAYBE).

Now we turn to the actual batch computation of $\tilde{E}(q)$ for all receivers in $\mathcal{Q}$ and point out a few additional caveats.

Computationally, $\tilde{E}(q)$ can be evaluated in constant time, given $\tilde{f}(q)$ and point $s(q) = s(q, \mathcal{S})$. So we focus on these two subproblems. For the uniform-power case, we can construct the Voronoi diagram of $\mathcal{S}$, preprocess it for point location, and query it for $s(q)$ with each receiver q , for a total cost of $O((n + m) \log n)$; see Fact B.3 and Reference [11]. In the case of non-uniform power, if we are content with near-quadratic running time, we can determine $s(q)$ by computing the multiplicatively weighted Voronoi diagram of $\mathcal{S}$ as outlined above, and then querying it with each receiver in total time $O(n^2 + m \log n)$ (see Fact B.4 and [6, 11]), which is too expensive when $m \approx n$. We provide an alternative below.

We show how to compute the values $\tilde{f}(q_1), \dots, \tilde{f}(q_m)$ in near-linear time, using a two-dimensional orthogonal range search tree. Indeed, observe that $l(s, q) = |q_x - s_x|$ provided $|q_x - s_x| \geq |q_y - s_y|$. For a fixed q , the region W_q containing the transmitters of $\mathcal{S}$ satisfying this inequality is a 90° double wedge. Using (a 45° tilted version of) the orthogonal range search tree

[11] (see Section B.1), we can construct a pair decomposition $\{(S_i, Q_i)\}$ of small size, so that each pair (s, q) with $s \in W_q$ appears in exactly one product $S_i \times Q_i$.

We now denote by $\tilde{f}(q, Z)$ the sum analogous to $\tilde{f}(q)$, where the summation goes over the elements of the supplied set Z rather than those of $\mathcal{S}$. Clearly,

$$\tilde{f}(q, \mathcal{S} \cap W_q) = \sum_{i: q \in Q_i} \tilde{f}(q, S_i), \quad (1)$$

by the definition of the pair decomposition. The number of terms in the last sum is $O(\log^2 n)$. Notice that $\tilde{f}(q, S_i)$, for a fixed i , is a sum of small fractional *univariate* functions, with $|S_i|$ terms in it, since the expression for transmitters in W_q depends only on q_x and not on q_y . Now for each pair (Q_i, S_i) , we use Corollary A.4 to batch evaluate $\tilde{f}(q, S_i)$ on each $q \in Q_i$ in total time $O((|Q_i| + |S_i|) \log^2 |S_i| \log \log |S_i|) = O((|Q_i| + |S_i|) \log^2 n \log \log n)$. This gives us all the summands of Equation (1) and therefore allows us to evaluate $\tilde{f}(q, \mathcal{S} \cap W_q)$ for all $q \in \mathcal{Q}$, in total time at most proportional to $\sum_i (|Q_i| + |S_i|) \log^2 n \log \log n = (\sum_i (|Q_i| + |S_i|)) \log^2 n \log \log n = O((m + n) \log^4 n \log \log n)$.

Of course, we have only treated those s that lie in W_q . But the calculation is repeated in the complementary double wedge, where now only the y -coordinates matter and $\tilde{f}(q)$ is the sum of the two values thus obtained.

THEOREM 3.3. *For any fixed positive even integer α , given a set $\mathcal{S}$ of n transmitters (all of power 1) and a set $\mathcal{Q}$ of m receivers, we can do the following in total time $O((m + n) \log^4 n \log \log n)$ and $O((m + n) \log^2 n)$ space. For each $q \in \mathcal{Q}$, we find its unique candidate transmitter $s(q)$ and compute a value $\tilde{E}(q)$, such that (i) if $\tilde{E}(q) \geq 2^{\alpha/2} \beta$, then q can definitely hear $s(q)$, (ii) if $\tilde{E}(q) < 2^{-\alpha/2} \beta$, then q definitely cannot hear $s(q)$, and (iii) if $2^{-\alpha/2} \beta \leq \tilde{E}(q) < 2^{\alpha/2} \beta$, then q may or may not hear $s(q)$.*

The algorithm for the non-uniform power case is hampered by the fact that the obvious way to identify the candidate transmitter each receiver might hear seems to involve constructing the multiplicatively weighted Voronoi diagram of quadratic complexity. However, we do not need the exact multiplicatively closest neighbor, but rather a reasonably-close approximation of the value $|q - s|/p(s)^{1/\alpha}$, over all $s \in \mathcal{S}$ (being off by a multiplicative factor of at most $2^{1/2}$ is sufficient; see the discussion below). Such an approximation is provided by the first algorithm in Fact B.7 (see Har-Peled and Kumar [19]), for a constant value of the approximation parameter ε (namely, $\varepsilon = 2^{1/2} - 1$), yielding the following:

THEOREM 3.4. *For any fixed positive even integer α and any $\beta > 2^{\alpha/2}$, given a set $\mathcal{S}$ of n transmitters of arbitrary powers and a set $\mathcal{Q}$ of m receivers, we can do the following in total time $O(n \log^7 n + m \log^4 n \log \log n)$ and $O(n \log^4 n + m \log^2 n)$ space: For each $q \in \mathcal{Q}$, we find a transmitter s_q and compute a value $\tilde{E}(q)$, such that (i) if $\tilde{E}(q) \geq 2^{\alpha/2} \beta$, then q can definitely hear s_q (implying that $s_q = s(q)$), (ii) if $\tilde{E}(q) < 2^{-\alpha/2} \beta$, then q definitely cannot hear any transmitter, and (iii) if $2^{-\alpha/2} \beta \leq \tilde{E}(q) < 2^{\alpha/2} \beta$, then q may or may not hear one of the transmitters.*

Note. The transmitter s_q in the theorem above is not necessarily the unique candidate transmitter $s(q)$. We would like to show that if $\tilde{E}(q) \geq 2^{\alpha/2} \beta$ (and therefore $E(q) \geq \beta$), then s_q is necessarily $s(q)$. Assume that they are different (i.e., that $s_q \neq s(q)$), and let e_q (resp., $e(q)$) be the strength of s_q 's signal (resp., $s(q)$'s signal) at q . Then, we know that $e_q \leq e(q) \leq 2^{\alpha/2} e_q$. Notice that $E(q) \leq e(q)/e_q$, since $E(q)$ is maximized when there is no third transmitter and no noise, so $e(q)/e_q \geq \beta$ since $E(q) \geq \beta$. Recall that we are assuming that $\beta > 2^{\alpha/2}$, so we get that $e(q)/e_q > 2^{\alpha/2}$, which is a contradiction.

Therefore, in fact, it must be the case that $s_q = s(q)$, as claimed.

We now turn the algorithm described above into a PTAS, in the sense that we will confine $\tilde{E}(q)$ to the range $((1 - \varepsilon)E(q), (1 + \varepsilon)E(q))$, for a prespecified $\varepsilon > 0$. We outline the approach below. Consider the regular k -gon K_k circumscribed around the Euclidean unit disk, for a large enough even $k \geq 4$ specified below. We modify the above algorithm, replacing the L^∞ -norm whose “unit disk” is a 2×2 square, with the norm $|\cdot|_k$ with K_k as the unit disk. Then $|v|_k \leq |v| \leq (1 + \Theta(k^{-2}))|v|_k$, for any vector v in the plane. In the range-searching data structure, wedges with opening angle $\pi/2 = 2\pi/4$ are replaced by wedges with opening angle $2\pi/k$, and we need $k/2$ copies of the structure.

In terms of the quality of approximation, the factor $2^{\alpha/2} = (\sqrt{2})^\alpha$ is replaced by $(1 + \Theta(k^{-2}))^\alpha \approx 1 + \alpha\Theta(k^{-2})$. Hence to obtain an approximation factor of $1 + \varepsilon$, we set $\varepsilon = \alpha\Theta(k^{-2})$, or $k = c(\alpha/\varepsilon)^{1/2}$, for a suitable absolute constant c . In other words, it is sufficient to create $O(\varepsilon^{-1/2})$ copies of the data structure. To summarize, we have:

THEOREM 3.5. *For a positive ε , any fixed positive even integer α , given a set S of n transmitters (all of power 1) and a set Q of m receivers, we can do the following in total time $O((m + n)\varepsilon^{-1/2} \log^4 n \log \log n)$ and $O((m + n)\varepsilon^{-1/2} \log^2 n)$ space. For each $q \in Q$, we find its unique candidate transmitter $s(q)$ and compute a value $\tilde{E}(q)$, such that (i) if $\tilde{E}(q) \geq (1 + \varepsilon)\beta$, then q can definitely hear $s(q)$, (ii) if $\tilde{E}(q) < (1 - \varepsilon)\beta$, then q definitely cannot hear $s(q)$, and (iii) if $(1 - \varepsilon)\beta \leq \tilde{E}(q) < (1 + \varepsilon)\beta$, then q may or may not hear $s(q)$.*

THEOREM 3.6. *For a positive $\varepsilon < 1 - \beta$, any fixed positive even integer α ,⁴ given a set S of n transmitters of arbitrary powers and a set Q of m receivers, we can do the following in total time $O(n\varepsilon^{-6} \log^7 n + m\varepsilon^{-1/2} \log^4 n \log \log n)$ and $O(n\varepsilon^{-6} \log^4 n + m\varepsilon^{-1/2} \log^2 n)$ space: For each $q \in Q$, we find a transmitter s_q and compute a value $\tilde{E}(q)$, such that (i) if $\tilde{E}(q) \geq (1 + \varepsilon)\beta$, then q can definitely hear s_q (implying that $s_q = s(q)$), (ii) if $\tilde{E}(q) < (1 - \varepsilon)\beta$, then q definitely cannot hear any transmitter, and (iii) if $(1 - \varepsilon)\beta \leq \tilde{E}(q) < (1 + \varepsilon)\beta$, then q may or may not hear one of the transmitters.*

Notes. (a) Hereafter, we refer to cases (i), (ii), and (iii) as our approximate SINR diagram point location oracle returning YES, NO, and MAYBE, respectively.

(b) Most likely the data structure from Reference [19] that we use for approximate multiplicatively weighted nearest neighbors is not the best possible for our purposes, as it is designed to handle a more general situation.

(c) Also, we could have replaced the exact Euclidean Voronoi diagram in Theorems 3.3 and 3.5 by an approximate nearest-neighbor structure. While unnecessary in the plane, it would allow for polylogarithmic amortized batched queries in 3D. In fact, everything we described in this section generalizes to 3D and, moreover, to $\mathbb{R}^d$, for any constant $d > 1$, to yield amortized polylogarithmic-time batched queries. One needs to approximate the Euclidean ball by a suitable polyhedron, use d -dimensional range trees and the d -dimensional data structure from Reference [19] for multiplicatively-weighted approximate neighbors (in the uniform case, one can use any one of the simpler approximate nearest-neighbor structures, such as the Approximate Voronoi Diagram; see further discussion in Reference [19]).

(d) The algorithms in Theorems 3.3, 3.4, 3.5, and 3.6 extend to odd values of α with very minor modifications. The difficulty with such values is that an expression such as $|q_x - s_{j,x}|^\alpha$ is not a polynomial, as it is equal to $\pm(q_x - s_{j,x})^\alpha$ depending on whether s_j is located to the left or to the right of q , in the double wedge W_q : we need to know if s_j lies in the left wedge of W_q or in the right one. In other words, the algorithm will proceed as before (with an appropriate sign correction), as

⁴This requirement is analogous to that in Theorem 3.4 to guarantee that the approximately highest-strength transmitter returned by the data structure is in fact the right one.

long as we replace double wedges by single wedges in the range searching data structure. (In fact, the range search already uses single wedges internally, so very little actual change is required.)

3.5 Handling Directional Antennas

Heretofore we have assumed that the transmitters in $\mathcal{S}$ use omnidirectional antennas. If directional antennas are employed instead, the SINR inequality must be modified accordingly. The coverage area of the directional antenna of transmitter s is a wedge $W = W(s)$ of angle $\theta = \theta(s)$ and apex at s . Only points within W can receive s . Assuming that $q \in W$, we need to evaluate the left side of the SINR inequality, taking into account only the transmitters of $\mathcal{S}$ whose wedges cover q . That is, q receives s of power p if and only if $q \in W(s)$ and

$$\frac{\frac{p}{|q-s|^\alpha}}{\sum_{s_j \neq s: q \in W_j} \frac{p_j}{|q-s_j|^\alpha} + N} \geq \beta.$$

The results of Section 3.4 are still relevant (as we show below), but they cannot be applied immediately. We first need to construct a data structure for range searching with wedges among the set of n points Q : Given a query wedge $W = W(s)$, return $Q \cap W$ as a union of a small number of canonical subsets of Q . Then, we perform a query with each of the wedges $W(s)$, $s \in \mathcal{S}$, to obtain a collection of pairs of sets $\{(S_i, Q_i)\}$ with the property that, for any $(s, q) \in \mathcal{S} \times Q$, if $q \in W(s)$, then there exists a single pair (S_i, Q_i) such that $(s, q) \in S_i \times Q_i$, and, if $q \notin W(s)$, then $(s, q) \notin S_i \times Q_i$, for any i . We now can apply the results of Section 3.4 to each of the pairs (S_i, Q_i) to obtain theorems analogous to Theorems 3.3–3.6. The time bound in each of these theorems will be $O_\varepsilon^*(n^{4/3})$, assuming $|\mathcal{S}| = |Q| = n$. We omit the remaining details.

Note that the additional overhead of dealing with directional antennas essentially disappears if there is only a fixed set of directions bounding the wedges $W(s)$.

4 SOME STRUCTURAL PROPERTIES OF SINR REGIONS

In this section, we obtain results on the sensitivity of the reception regions to slight changes in the reception threshold. Specifically, we prove two technical lemmas (assuming uniform power), which roughly state that scaling a reception region R_i centered at s_i by a factor of $(1 + \varepsilon)$ (respectively, $(1 - \varepsilon)$) is equivalent to decreasing β (respectively, increasing β) by a factor of $(1 - \Theta(\varepsilon))$ (respectively, $(1 + \Theta(\varepsilon))$). We use these lemmas in the next section to speed up the construction time of a previous data structure for approximate point location in the SINR diagram. These lemmas also imply that the two notions of approximating R_i mentioned in the introduction are actually equivalent, at least for the uniform power case; see the discussion below.

We introduce the following auxiliary notation: First, for a positive number δ , we let δR_i be a copy of the set R_i scaled by a factor δ ; the center of the scaling transformation is s_i , the location of the i th site. For example, $2R_i$ is a copy of R_i scaled so that its linear dimensions are twice as large, with the center at s_i . Notice that $\text{area}(\delta R_i) = \delta^2 \text{area}(R_i)$.

Consider modifying an instance of the SINR problem, in the following manner: Keep the locations and the power of the transmitter s_i the same, but replace the ratio threshold β by some value $\gamma > 1$. We use $R_i(\gamma)$ to denote the region of s_i in the resulting SINR diagram. By definition, $R_i = R_i(\beta)$.

We are now ready to state and prove the two technical lemmas; their proofs are similar in spirit, but differ in detail.

LEMMA 4.1. *Consider a set of n uniform-power transmitters at locations s_i , $i = 1, \dots, n$. Assume $\beta > 1$. There exist values $c_2 > c_1 > 0$, that depend only on α and β , such that, for any $\varepsilon \in (0, 1/c_2)$*

and all i ,⁵

$$R_i((1 - c_1\varepsilon)\beta) \subset (1 + \varepsilon)R_i \subset R_i((1 - c_2\varepsilon)\beta).$$

PROOF. Fix a transmitter s_i . Fix a point q on the boundary of R_i . Its SIN ratio $E(q)$ relative to s_i , defined as

$$E(q) := \frac{\frac{1}{|q-s_i|^\alpha}}{\sum_{j \neq i} \frac{1}{|q-s_j|^\alpha} + N},$$

is precisely β , since R_i is convex [9]. Without loss of generality, s_i lies at the origin and q on the positive x -axis at distance d from it. Put $d' = (1 + \varepsilon)d$. Consider the point $q' = (d', 0)$ on the boundary of $(1 + \varepsilon)R_i$. We measure how different the SIN ratio $E(q')$ with respect to s_i is from $E(q)$ by bounding $E(q')/E(q)$. Let $N(q)$ and $D(q)$ denote the numerator and denominator of the above expression for $E(q)$, respectively. Then, $E(q')/E(q) = (N(q')/N(q)) \cdot (D(q)/D(q'))$, and $N(q')/N(q) = (1/d')^\alpha / (1/d)^\alpha = (d/d')^\alpha = (1 + \varepsilon)^{-\alpha}$. We now consider the ratio of the denominators (i.e., $D(q)/D(q')$).

We first examine the effect of the noise N on the ratio of the denominators. It is easy to check that the ratio is a monotone function of N and approaches 1 (either from below or from above) as N grows arbitrarily large. Hence the extreme values of the ratio of the denominators can be computed by setting $N = 0$; one of the extremes corresponds to the denominator decreasing as we move from q to q' and the other to it increasing.

The maximum possible *decrease* in the denominator corresponds to the maximum *increase* in the distances to all sites other than s_i , which in turn corresponds to every such site s_j lying on the negative x -axis (we will see below that s_j cannot lie on the segment s_iq). Thus,

$$\begin{aligned} \frac{D(q)}{D(q')} &= \frac{\sum_{j \neq i} \frac{1}{|q-s_j|^\alpha}}{\sum_{j \neq i} \frac{1}{|q'-s_j|^\alpha}} \leq \max_{j \neq i} \frac{\frac{1}{|q-s_j|^\alpha}}{\frac{1}{|q'-s_j|^\alpha}} = \max_{j \neq i} \left(\frac{|s_j - q'|}{|s_j - q|} \right)^\alpha \\ &= \max_{j \neq i} \left(\frac{|s_j - q| + |q' - q|}{|s_j - q|} \right)^\alpha. \end{aligned}$$

We now observe that, since $E(q) = \beta > 1$, not only $|s_i - q| \leq |s_j - q|$ for all $j \neq i$, but also $|s_i - q|\beta^{1/\alpha} \leq |s_j - q|$ (this follows from the fact that, since the SIN ratio of q is β , the signal strength of s_i at q is at least β times that of s_j at q , i.e., $1/|s_i - q|^\alpha \geq \beta/|s_j - q|^\alpha$).⁶ Therefore, for all $j \neq i$,

$$\frac{|s_j - q| + |q' - q|}{|s_j - q|} \leq \frac{\beta^{1/\alpha}|s_i - q| + \varepsilon|s_i - q|}{\beta^{1/\alpha}|s_i - q|} = 1 + \varepsilon\beta^{-1/\alpha},$$

and

$$\begin{aligned} \frac{E(q')}{E(q)} &= (1 + \varepsilon)^{-\alpha} \cdot \frac{D(q)}{D(q')} \leq (1 + \varepsilon)^{-\alpha} \cdot (1 + \varepsilon\beta^{-1/\alpha})^\alpha \\ &= \left(\frac{1 + \varepsilon\beta^{-1/\alpha}}{1 + \varepsilon} \right)^\alpha = \left(\frac{1 + \varepsilon - (\varepsilon - \beta^{-1/\alpha}\varepsilon)}{1 + \varepsilon} \right)^\alpha \\ &= \left(1 - \frac{(1 - \beta^{-1/\alpha})\varepsilon}{1 + \varepsilon} \right)^\alpha \leq \left(1 - \frac{1 - \beta^{-1/\alpha}}{2}\varepsilon \right)^\alpha \leq 1 - \frac{(1 - \beta^{-1/\alpha})}{2}\varepsilon, \end{aligned}$$

provided $\varepsilon \leq 1$ and $\alpha \geq 1$.

⁵The first inclusion holds for any $\varepsilon \in (0, 1/c_1)$.

⁶Notice that this implies $s_j \notin s_iq$, as promised.

At the other extreme, the distance to s_j is maximally decreased if it lies on the positive x -axis beyond q' . So we assume that this is the case for all sites s_j , $j \neq i$. Thus,

$$\frac{D(q)}{D(q')} \geq \min_{j \neq i} \frac{\frac{1}{|q-s_j|^\alpha}}{\frac{1}{|q'-s_j|^\alpha}} = \min_{j \neq i} \left(\frac{|s_j - q'|}{|s_j - q|} \right)^\alpha = \min_{j \neq i} \left(\frac{|s_j - q| - |q' - q|}{|s_j - q|} \right)^\alpha.$$

Once again using $|s_i - q| \beta^{1/\alpha} \leq |s_j - q|$, we obtain, for all $j \neq i$,

$$\frac{|s_j - q| - |q' - q|}{|s_j - q|} \geq \frac{\beta^{1/\alpha} |s_i - q| - \varepsilon |s_i - q|}{\beta^{1/\alpha} |s_i - q|} = 1 - \varepsilon \beta^{-1/\alpha},$$

yielding

$$\begin{aligned} \frac{E(q')}{E(q)} &= (1 + \varepsilon)^{-\alpha} \cdot \frac{D(q)}{D(q')} \geq (1 + \varepsilon)^{-\alpha} \cdot (1 - \varepsilon \beta^{-1/\alpha})^\alpha = \left(\frac{1 - \varepsilon \beta^{-1/\alpha}}{1 + \varepsilon} \right)^\alpha \\ &= \left(\frac{1 + \varepsilon - (\varepsilon + \varepsilon \beta^{-1/\alpha})}{1 + \varepsilon} \right)^\alpha = \left(\frac{1 + \varepsilon - (1 + \beta^{-1/\alpha})\varepsilon}{1 + \varepsilon} \right)^\alpha \\ &= \left(1 - \frac{1 + \beta^{-1/\alpha}}{1 + \varepsilon} \varepsilon \right)^\alpha \geq (1 - (1 + \beta^{-1/\alpha})\varepsilon)^\alpha \geq 1 - \alpha(1 + \beta^{-1/\alpha})\varepsilon, \end{aligned}$$

where we used $\varepsilon > 0$ in the second to last step and $(1 - x)^c \geq 1 - xc$, for $0 < x < 1$ and $c \geq 1$, in the last step. The last estimate is only helpful if $\alpha(1 + \beta^{-1/\alpha})\varepsilon < 1$, or $\varepsilon < 1/(\alpha(1 + \beta^{-1/\alpha}))$.

Now recall R_i and $(1 + \varepsilon)R_i$ are star-shaped with center s_i , and q (respectively, q') is a boundary point of R_i (respectively, $(1 + \varepsilon)R_i$). We have therefore proven that

$$R_i((1 - c_1\varepsilon)\beta) \subset (1 + \varepsilon)R_i \subset R_i((1 - c_2\varepsilon)\beta),$$

with $c_1 = \frac{1 - \beta^{-1/\alpha}}{2}$ and $c_2 = \alpha(1 + \beta^{-1/\alpha})$, as promised. $\square$

We will also need an analogous statement for scaling R_i down rather than up.

LEMMA 4.2. *Consider a set of n uniform-power transmitters at locations s_i , $i = 1, \dots, n$. Assume $\beta > 1$. There exist values $c_4 > c_3 > 0$ that depend only on α and β so that, for any $\varepsilon \in (0, 1/2)$,*

$$R_i((1 + c_4\varepsilon)\beta) \subset (1 - \varepsilon)R_i \subset R_i((1 + c_3\varepsilon)\beta).$$

PROOF. Fix a transmitter s_i . Fix a point q on the boundary of R_i . Its SIN ratio $E(q)$ relative to s_i is precisely β , since R_i is convex [9]. Without loss of generality, s_i lies at the origin and q on the positive x -axis at distance d from it. Put $d' = (1 - \varepsilon)d$. Consider the point $q' = (d', 0)$ on the boundary of $(1 - \varepsilon)R_i$. We measure how different the SIN ratio $E(q')$ with respect to s_i is from $E(q)$ by bounding $E(q')/E(q)$. As before, let $N(q)$ and $D(q)$ denote the numerator and denominator of $E(q)$, respectively. Then, $E(q')/E(q) = (N(q')/N(q)) \cdot (D(q)/D(q'))$, and $N(q')/N(q) = (1/d')^\alpha / (1/d)^\alpha = (d/d')^\alpha = (1 - \varepsilon)^{-\alpha}$. We now consider the ratio of the denominators (i.e., $D(q)/D(q')$).

As before, the extreme values of the ratio of the denominators can be computed by setting $N = 0$; one of the extremes corresponds to the denominator shrinking as we move from q to q' and the other to it increasing.

The maximum possible *increase* in the denominator corresponds to the maximum *decrease* in the distances to all sites other than s_i , which in turn corresponds to every such site s_j lying on the

negative x -axis (we will see below that s_j cannot lie on the segment $s_i q$). Thus,

$$\begin{aligned} \frac{D(q)}{D(q')} &= \frac{\sum_{j \neq i} \frac{1}{|q-s_j|^\alpha}}{\sum_{j \neq i} \frac{1}{|q'-s_j|^\alpha}} \geq \min_{j \neq i} \frac{\frac{1}{|q-s_j|^\alpha}}{\frac{1}{|q'-s_j|^\alpha}} = \min_{j \neq i} \left(\frac{|s_j - q'|}{|s_j - q|} \right)^\alpha \\ &= \min_{j \neq i} \left(\frac{|s_j - q| - |q' - q|}{|s_j - q|} \right)^\alpha. \end{aligned}$$

As in the proof of Lemma 4.1, we note that $E(q) = \beta > 1$ implies $|s_i - q|\beta^{1/\alpha} \leq |s_j - q|$.⁷ Therefore, for all $j \neq i$,

$$\frac{|s_j - q| - |q' - q|}{|s_j - q|} \geq \frac{\beta^{1/\alpha}|s_i - q| - \varepsilon|s_i - q|}{\beta^{1/\alpha}|s_i - q|} = 1 - \varepsilon\beta^{-1/\alpha},$$

and

$$\begin{aligned} \frac{E(q')}{E(q)} &\geq (1 - \varepsilon)^{-\alpha} (1 - \varepsilon\beta^{-1/\alpha})^\alpha = \left(\frac{1 - \varepsilon\beta^{-1/\alpha}}{1 - \varepsilon} \right)^\alpha \\ &= \left(\frac{1 - \varepsilon + (\varepsilon - \beta^{-1/\alpha}\varepsilon)}{1 - \varepsilon} \right)^\alpha = \left(1 + \frac{(1 - \beta^{-1/\alpha})\varepsilon}{1 - \varepsilon} \right)^\alpha \\ &\geq (1 + (1 - \beta^{-1/\alpha})\varepsilon)^\alpha \geq 1 + (1 - \beta^{-1/\alpha})\varepsilon, \end{aligned}$$

provided $\varepsilon < 1$ and $\alpha \geq 1$.

At the other extreme, the distance to s_j is maximally increased if it lies on the positive x -axis beyond q . So we assume that this is the case for every site $s_j, j \neq i$. Thus,

$$\frac{D(q)}{D(q')} \leq \max_{j \neq i} \left(\frac{|s_j - q'|}{|s_j - q|} \right)^\alpha = \max_{j \neq i} \left(\frac{|s_j - q| + |q' - q|}{|s_j - q|} \right)^\alpha.$$

Once again using $|s_i - q|\beta^{1/\alpha} \leq |s_j - q|$, we obtain, for all $j \neq i$,

$$\frac{|s_j - q| + |q' - q|}{|s_j - q|} \leq \frac{\beta^{1/\alpha}|s_i - q| + \varepsilon|s_i - q|}{\beta^{1/\alpha}|s_i - q|} = 1 + \varepsilon\beta^{-1/\alpha},$$

yielding

$$\begin{aligned} \frac{E(q')}{E(q)} &\leq \frac{(1 + \varepsilon\beta^{-1/\alpha})^\alpha}{(1 - \varepsilon)^\alpha} = \frac{(1 - \varepsilon + (\varepsilon + \varepsilon\beta^{-1/\alpha}))^\alpha}{(1 - \varepsilon)^\alpha} \\ &= \left(\frac{1 - \varepsilon + (1 + \beta^{-1/\alpha})\varepsilon}{1 - \varepsilon} \right)^\alpha = \left(1 + \frac{1 + \beta^{-1/\alpha}}{1 - \varepsilon} \varepsilon \right)^\alpha \\ &\leq (1 + 2(1 + \beta^{-1/\alpha})\varepsilon)^\alpha \leq 1 + 2((2 + \beta^{-1/\alpha})^\alpha - 1)\varepsilon, \end{aligned}$$

where we used $0 < \varepsilon < 1/2$ and convexity of x^α .

Now recall R_i and $(1 - \varepsilon)R_i$ are star-shaped with center s_i , and q (respectively, q') is a boundary point of R_i (respectively, $(1 - \varepsilon)R_i$). We have therefore proven that

$$R_i((1 + c_4\varepsilon)\beta) \subset (1 - \varepsilon)R_i \subset R_i((1 + c_3\varepsilon)\beta),$$

with $c_3 = 1 - \beta^{-1/\alpha}$ and $c_4 = 2((2 + \beta^{-1/\alpha})^\alpha - 1)$, as promised. $\square$

Notice that Lemmas 4.1 and 4.2 imply that (up to a factor that depends only on $\beta > 1$ and α), the two notions of approximating R_i defined in the introduction are actually equivalent, at least for the uniform power case. Roughly, being off by a factor of $1 \pm \delta$ in the scale of R_i (which is the same

⁷Notice that this implies $s_j \notin qs_i$, as promised.

as being off by a factor of $(1 \pm \delta)^2 = 1 \pm \Theta(\delta)$, when $\delta \leq 1$, in its area), corresponds to being off by a factor $1 \mp C\delta$ in the SINR ratio $E(q)$ for the query point q , where C depends on the constants c_1, c_2, c_3 , and c_4 from the above lemmas, which in turn depend only on α and β .

5 SPEEDING UP PREPROCESSING FOR A PREVIOUS POINT-LOCATION RESULT

In this section, we explain how Theorem 3.5 together with the observations from Section 4 can be applied to speed up the preprocessing stage of an existing point-location algorithm. (Similar techniques can be used in other situations, as long as they involve a large number of independent SIN ratio queries and can be modified to tolerate an approximate answer.)

Specifically, consider the data structure presented by Avin et al. [9] for a set of n uniform-power transmitters, with construction time⁸ $O(n^2/\delta)$ and query time $O(\log n)$, where $\delta > 0$ is a given approximation parameter; the query returns YES/NO/MAYBE, where YES/NO are guaranteed to be correct and MAYBE occurs rarely in the sense that (the regions corresponding to each of the YES/NO/MAYBE answers have well-defined area, see below, and) the area of the locus of points where the data structure returns MAYBE is at most $\delta \cdot \text{area}(R_i)$.

The structure of Reference [9] is actually a collection of n data structures DS_i , one per transmitter. The data structure DS_i for transmitter s_i consists of an inner (R_i^+) and outer (R_i^-) approximation for reception region R_i , so that $\text{area}(R_i^- \setminus R_i^+) \leq \delta \cdot \text{area}(R_i)$, see the definitions in the introduction. The construction of DS_i uses the fact that the region R_i is convex and fat; the latter in this case means that the ratio of the radius of the smallest disk enclosing R_i and centered at s_i and the radius of the largest disk enclosed in R_i and centered at s_i is bounded by an expression that depends only on α and β and not on n or the geometry of the input [9]. After building the Voronoi diagram of the transmitter sites in $O(n \log n)$ time [11], the construction of DS_i consists of two stages. In the first, they compute an explicit estimate on the size of R_i , by applying an exponential-search-like procedure along the segment connecting s_i to its nearest other transmitter s_j in $\mathcal{S}$, where each comparison is resolved by explicitly evaluating the SIN ratio at some point $q \in s_i s_j$ and comparing it to β , i.e., by an (*exact*) IN/OUT test. In the second stage, a grid of size roughly $1/\delta \times 1/\delta$, scaled to exactly cover the outer disk is laid, and, by performing $O(1/\delta)$ additional IN/OUT tests, the sets R_i^+ and R_i^- are obtained, as collections of grid cells. This algorithm thus performs $\Theta(\log n + 1/\delta)$ exact IN/OUT tests per transmitter, at a cost of $\Theta(n)$ arithmetic operations each; the high cost of each test is the bottleneck of the preprocessing step.

Conceptually, we speed up the above algorithm by constructing the n individual data structures DS_i in parallel, forming $O(\log n + 1/\delta^2)$ batches of n queries each, and use Theorem 3.5 to evaluate each batch in near-linear time. However, our batch query algorithm is not exact, but approximate; for some queries, instead of IN or OUT, it answers MAYBE. Below we explain how to handle this complication. We first give a more detailed description of the procedure used by Avin et al., and highlight the necessary modifications. Along the way, we will need to remind the reader how queries are performed in Reference [9], so as to make sure that the same query mechanism still applies after our changes.

(i) The Voronoi diagram of the transmitter locations is built in $O(n \log n)$ time and adds $O(\log n)$ time to the query, as already mentioned above. It is used at the query stage to identify the unique candidate site $s_i = s(q)$ for the query q . Both versions focus on deciding whether q can hear s_i , approximately, by constructing the data structure DS_i for handling those queries whose closest site is s_i , for every i . Our description will hereafter be confined to building and querying of DS_i .

(ii) They first estimate the extent of R_i by, conceptually, approximating the radii of the smallest enclosing and the largest enclosed circles of R_i , centered at s_i . Since R_i is fat, however, it is enough

⁸Again, the actual preprocessing time seems to be $O(n^2(\log n + 1/\delta))$, see below.

to restrict the calculations to a single ray emanating from s_i , namely, the one connecting it to the nearest transmitter location s_j . Without loss of generality, put $s_i = (0, 0)$ and assume that this ray is the positive ray of the x -axis, emanating from s_i .

They first estimate the extent of R_i by finding two values d_i, D_i , $0 < d_i < D_i < |s_i - s_j|$, so that $(d_i, 0) \in R_i$ and $(D_i, 0) \notin R_i$. Theorem 4.1 in Reference [9] gives explicit expressions for the initial values of d_i and D_i and proves that D_i/d_i is bounded by a polynomial in n . This calculation takes constant time (per transmitter), once the Voronoi diagram of the sites has been computed and nearest-neighbor information for each site extracted. It does not involve computing the SIN ratio for any point.

(iii) The next step refines the estimates on d_i and D_i to ensure $D_i/d_i = O(1)$. In Reference [9], this is done by performing an exponential search, each time doubling the distance from s_i to a point known to be in R_i until a pair of consecutive points, one in R_i and one outside of it, are obtained. Since the initial estimate ratio is known to be polynomial in n , this requires $O(\log n)$ exact oracle calls (i.e., IN/OUT tests), per transmitter.⁹

We now describe how to perform the same procedure using an approximate oracle with precision parameter Δ , to be fixed below. We call the oracle on a sequence of points $(d_i \lambda^k, 0)$, $k = 1, 2, \dots$, on $s_i s_j$ whose distance from s_i increases by a factor of λ at each step, where $\lambda > 1$ is a small constant. We artificially add $(d_i, 0)$ (from step (ii)) to the beginning of the sequence as a YES (even though the approximate oracle may not return a YES there) and $(D_i, 0)$ (from step (ii)) to the end of the sequence as a NO (even though it may not return a NO), to guarantee termination.

We return a pair of possibly non-consecutive points in this sequence so that the first is a YES, the last is a NO, and all intervening points, if any, are MAYBES. By construction, the first point in our sequence is a YES and the last a NO, hence the claimed pair exists. We evaluate the entire sequence and return such a pair. Notice that there cannot be two such pairs, as the sequence starts with a YES and a NO cannot appear before a YES, for that would imply the existence of a point outside of R_i between two points in it, contradicting the convexity of R_i .

We again refer to the resulting pair of values as d_i and D_i , respectively, although, of course, they may be different from the original d_i, D_i pair found in step (ii). Our goal now is to show that $D_i/d_i = O(1)$. For this, we need the following lemma whose proof is similar to the first part of the proof of Lemma 4.1.

LEMMA 5.1. *Assume $\beta > 1$, and let q be a point such that $E(q) \geq \beta/\rho$ with respect to s_i , where $1 \leq \rho < \beta$, and let q' be the point on the ray emanating from s_i and passing through q at distance $\tau|q - s_i|$ from s_i , $\tau > 1$. Then*

$$\frac{E(q')}{E(q)} \leq (1 - (1 - (\rho/\beta)^{1/\alpha})(1 - 1/\tau))^\alpha.$$

PROOF. Without loss of generality, s_i lies at the origin and q and q' on the positive x -axis at distances d and $d' = \tau d$, respectively, from it. For a point p , let $N(p)$ and $D(p)$ denote the numerator and denominator of the expression for $E(p)$ with respect to s_i , respectively. Then, $E(q')/E(q) = (N(q')/N(q)) \cdot (D(q)/D(q'))$, and $N(q')/N(q) = (1/d')^\alpha / (1/d)^\alpha = (d/d')^\alpha = \tau^{-\alpha}$. We now consider the ratio $D(q)/D(q')$ of the denominators.

As noted in the proof of Lemma 4.1, the maximum value of the ratio of the denominators is achieved when $N = 0$ and the denominator $D()$ decreases as much as possible as we move from q to q' , i.e., when the distances to all sites other than s_i increase as much as possible as we move from

⁹The stated total preprocessing time in Reference [9] seems to neglect the fact that this computation needs to be done for each transmitter, at the cost of $\Theta(n)$ per evaluation of the SIN ratio, for a total of $O(n^2 \log n)$; this term appears to be missing from the total running time in their analysis.

q to q' . This in turn corresponds to the case where each such site s_j lies on the negative x -axis (we will see below that s_j cannot lie on the segment $s_i q$). Thus,

$$\begin{aligned} \frac{D(q)}{D(q')} &= \frac{\sum_{j \neq i} \frac{1}{|q-s_j|^\alpha}}{\sum_{j \neq i} \frac{1}{|q'-s_j|^\alpha}} \leq \max_{j \neq i} \frac{\frac{1}{|q-s_j|^\alpha}}{\frac{1}{|q'-s_j|^\alpha}} = \max_{j \neq i} \left(\frac{|s_j - q'|}{|s_j - q|} \right)^\alpha \\ &= \max_{j \neq i} \left(\frac{|s_j - q| + |q' - q|}{|s_j - q|} \right)^\alpha. \end{aligned}$$

We now observe that $|s_i - q|(\beta/\rho)^{1/\alpha} \leq |s_j - q|$ for all $j \neq i$ (this follows from the fact that, since the SIN ratio of q is at least β/ρ , the signal strength of s_i at q is at least β/ρ times that of s_j at q , i.e., $1/|s_i - q|^\alpha \geq (\beta/\rho)/|s_j - q|^\alpha$).¹⁰ Therefore, for all $j \neq i$,

$$\begin{aligned} \frac{|s_j - q| + |q' - q|}{|s_j - q|} &\leq \frac{(\beta/\rho)^{1/\alpha} |s_i - q| + (\tau - 1) |s_i - q|}{(\beta/\rho)^{1/\alpha} |s_i - q|} \\ &= 1 + (\tau - 1)(\rho/\beta)^{1/\alpha}, \end{aligned}$$

and

$$\begin{aligned} \frac{E(q')}{E(q)} &= \tau^{-\alpha} \cdot \frac{D(q)}{D(q')} \leq \tau^{-\alpha} \cdot (1 + (\tau - 1)(\rho/\beta)^{1/\alpha})^\alpha \\ &= \left(\frac{1 + (\tau - 1)(\rho/\beta)^{1/\alpha}}{\tau} \right)^\alpha \\ &= \left(\frac{\tau - ((\tau - 1) - (\rho/\beta)^{1/\alpha}(\tau - 1))}{\tau} \right)^\alpha \\ &= (1 - (1 - (\rho/\beta)^{1/\alpha})(1 - 1/\tau))^\alpha. \end{aligned} \quad \square$$

We are now ready to show that $D_i/d_i = O(1)$. First, so that we do not have to distinguish between two cases, if D_i was not determined from the oracle calls on the points $(d_i \lambda^k, 0)$, $k = 1, 2, \dots$, i.e., if D_i is still the value obtained in step (ii), then we replace it by the first value above it of the form $d_i \lambda^k$. Next, we may assume that $D_i/d_i > \lambda^2$ (since otherwise we are done), and therefore the points $q = (d_i \lambda, 0)$ and $q' = (D_i/\lambda, 0)$ are distinct.

We upper bound D_i/d_i by bounding the ratio $E(q')/E(q)$ both from below and from above. We know that our approximate oracle returned the answer `MAYBE` for both these points. Therefore, since it did not return `YES` for q , we may conclude that $\tilde{E}(q) < (1 + \Delta)\beta$, and since it did not return `NO` for q' , we may conclude that $\tilde{E}(q') > (1 - \Delta)\beta$. This allows us to lower bound the ratio $E(q')/E(q)$:

$$\frac{E(q')}{E(q)} \geq \frac{\tilde{E}(q')/(1 + \Delta)}{\tilde{E}(q)/(1 - \Delta)} \geq \left(\frac{1 - \Delta}{1 + \Delta} \right)^2.$$

We now apply Lemma 5.1 to upper bound the ratio $E(q')/E(q)$. Since our approximate oracle did not return the answer `NO` for q , we may conclude that $\tilde{E}(q) > (1 - \Delta)\beta$ and set $\rho = 1 + \Delta$ in Lemma 5.1. Moreover, since $|q' - s_i| = (D_i/(d_i \lambda^2))|q - s_i|$, we set $\tau = D_i/(d_i \lambda^2)$ in Lemma 5.1, to obtain

$$\frac{E(q')}{E(q)} \leq \left(1 - \left(1 - \left(\frac{1 + \Delta}{\beta} \right)^{1/\alpha} \right) \left(1 - \frac{d_i \lambda^2}{D_i} \right) \right)^\alpha.$$

¹⁰Notice that this implies $s_j \notin s_i q$.

Combining the two estimates, we obtain

$$\left(\frac{1-\Delta}{1+\Delta}\right)^2 \leq \left(1 - \left(1 - \left(\frac{1+\Delta}{\beta}\right)^{1/\alpha}\right)\left(1 - \frac{d_i \lambda^2}{D_i}\right)\right)^\alpha,$$

which after some manipulation yields

$$\frac{D_i}{d_i} \leq \lambda^2 \cdot \frac{1 - \left(\frac{1+\Delta}{\beta}\right)^{1/\alpha}}{\left(\frac{1-\Delta}{1+\Delta}\right)^{2/\alpha} - \left(\frac{1+\Delta}{\beta}\right)^{1/\alpha}},$$

provided both the denominator and the numerator of the last fraction are positive, which introduces the restriction $\beta \geq \frac{(1+\Delta)^3}{(1-\Delta)^2}$, which in turn is less strict than the requirement $\beta \geq 1 + 6\Delta$ (assuming $\Delta \leq 1/14$).

(iv) At this point, the authors of Reference [9] have obtained a constant-factor approximation of D_i . Then they construct a square just large enough to cover a disc of radius D_i times the maximum fatness allowed by Corollary 5.4 from Reference [9], centered at s_i , and therefore fully containing R_i ; since R_i is fat, it occupies a constant fraction of the square. This square is subdivided into a $k \times k$ grid, where $k = \Theta(1/\delta)$ and δ is the desired approximation parameter. Avin et al. [9] now classify the grid squares as *internal*, *external*, or *mixed*, depending on whether all/none/some (but not all or none) of its corners are contained in R_i , respectively. Conceptually, in their construction, the inner approximation R_i^+ of R_i consists of the internal squares only, mixed squares intersect the boundary of R_i , and external squares may or may not intersect R_i , so an additional test is needed. They show [9, Lemma 5.3 and Corollary 5.4] that any external cell at distance¹¹ more than $C = C(\alpha, \beta)$ grid steps from any mixed cell must lie entirely outside of R_i . Therefore, an outer approximation R_i^- can be obtained by collecting all internal and mixed cells, and all the external cells within this distance of any mixed cell.

Identification of internal/mixed/external cells can be performed by brute force, evaluating the SIN ratio at every one of the $k^2 = \Theta(1/\delta^2)$ grid nodes. The algorithm in Reference [9] is cleverer in that it only examines $\Theta(k)$ cells by taking advantage of the fact that R_i is convex, speeding up preprocessing to only $\Theta(k) = \Theta(1/\delta)$ oracle calls per transmitter, for a total of $O(n^2/\delta)$ work.

Since both R_i^+ and R_i^- are represented as (unions of) collections of grid squares, such that the squares in any one row are contiguous, they can be stored in an array of size $O(1/\delta)$, with two entries for each grid row, identifying the cells where each set starts and ends. Thus (this portion of) the query q can be performed in constant time, provided *floor* operation is available, by identifying the grid row containing q , followed by a look-up in the above array, to verify whether or not q belongs to the range of cells covered by R_i^+ and/or R_i^- .

This completes the description of the preprocessing and query algorithms as presented in Reference [9]. In our adaptation of this step of their algorithm, we encounter several difficulties. The first one is that our oracle only provides approximate answers. The second is that, even though R_i itself is convex, we cannot guarantee that the region where the oracle of Theorem 3.5 provides, say, YES answers, is convex. Therefore we will not use the cleverer grid-traversal idea of Avin et al. [9] and will settle for just calling the approximate oracle with approximation parameter Δ , to be specified below, on every single grid node. This requires k^2 oracle calls for each transmitter, but

¹¹The *distance* between two compact sets is the minimum Euclidean distance between any pair of points, one from each set.

Fig. 1. Left: The region R_i associated with transmitter s_i , the enclosing circle D_i , and the $k \times k$ grid. The region R_i^{++} (shaded) is contained in R_i , where $\mathbf{I}$ denotes an internal grid cell and $\mathbf{M}$ denotes a mixed grid cell. The region R_i^{--} (depicted by its boundary) contains R_i . Right: The top two rows of R_i^{++} . $\mathbf{Y}$ stands for YES and $\mathbf{N}$ stands for NO. Notice that the mixed grid cell in the second row is included in R_i^{++} , since it is “sandwiched” between internal cells.

will be done in a batched manner, each batch consisting of one call for each transmitter. The total cost of all these queries will be $O(n\delta^{-2}\Delta^{-1/2}\log^4 n \log \log n)$, as per Theorem 3.5.

We now modify the internal/mixed/external classification of Avin et al. for the situation where we use an approximate oracle, as follows: A grid cell where the *approximate* oracle reports YES at all four corners is *internal*. Intuitively, the new inner approximation R_i^{++} of R_i is again defined as the set of all internal grid squares. For technical reasons explained below, we will extend the definition slightly, as follows: In every row of the grid, add to R_i^{++} all squares from the leftmost internal square to the rightmost one, inclusively. In other words, we also include possibly non-internal squares “sandwiched” between internal ones; see Figure 1 (right). Observe that by convexity of R_i , both an internal square and a “sandwiched” non-internal square are fully contained in R_i . Therefore $R_i^{++} \subset R_i$.¹²

We claim that $\text{area}(R_i \setminus R_i^{++}) \leq O(\delta) \cdot \text{area}(R_i)$. Consider a grid point $q \in R_i((1 + 3\Delta)\beta)$, so that $E(q) \geq (1 + 3\Delta)\beta$. By the approximation guarantee of the oracle, $\tilde{E}(q) \geq (1 - \Delta)(1 + 3\Delta)\beta \geq (1 + \Delta)\beta$, provided $0 < \Delta \leq 1/3$, and the approximate oracle returns YES on q . Hence every grid square all of whose corners lie in $R_i((1 + 3\Delta)\beta)$ is contained in R_i^{++} . We cannot conclude that $R_i((1 + 3\Delta)\beta) \subseteq R_i^{++}$, but we may conclude that $R_i((1 + 3\Delta)\beta) \setminus R_i^{++}$ is fully contained in the union of grid squares intersected by the boundary of $R_i((1 + 3\Delta)\beta)$. The latter set is convex, so its boundary intersects the grid lines at most $4k$ times and therefore meets at most $4k$ grid squares. In particular, $\text{area}(R_i((1 + 3\Delta)\beta) \setminus R_i^{++}) \leq O(\delta) \cdot \text{area}(R_i)$, and, therefore, $\text{area}(R_i \setminus R_i^{++}) \leq \text{area}(R_i \setminus R_i((1 + 3\Delta)\beta)) + O(\delta) \cdot \text{area}(R_i)$.

Now observe that, by Lemma 4.2, $(1 - 3\Delta/c_3)R_i \subset R_i((1 + 3\Delta)\beta)$, where $c_3 = c_3(\alpha, \beta)$ is the constant from the lemma, as long as $0 < 3\Delta/c_3 < 1/2$, or $0 < \Delta < c_3/6$. So, $\text{area}(R_i \setminus R_i^{++}) \leq \text{area}(R_i \setminus (1 - 3\Delta/c_3)R_i) + O(\delta) \cdot \text{area}(R_i)$.

¹²It is not difficult to see that, in fact, $R_i^{++} \subset R_i^+$, though we will not use this fact below.

Note that

$$\begin{aligned}
 \text{area}(R_i \setminus (1 - 3\Delta/c_3)R_i) &= \text{area}(R_i) - \text{area}((1 - 3\Delta/c_3)R_i) \\
 &= \text{area}(R_i) - (1 - 3\Delta/c_3)^2 \text{area}(R_i) \\
 &= (6\Delta/c_3 - 9\Delta^2/c_3^2) \text{area}(R_i) \\
 &= \Theta(\Delta) \text{area}(R_i),
 \end{aligned}$$

for the range of Δ under consideration.

Therefore we may conclude that

$$\text{area}(R_i \setminus R_i^{++}) \leq \Theta(\Delta) \text{area}(R_i) + O(\delta) \text{area}(R_i) = O(\delta) \text{area}(R_i),$$

provided we pick $\Delta = O(\delta)$.

We now construct the upper approximation R_i^- to R_i ; see Figure 1 (left). An *external* grid square is one where the approximate oracle returns NO on all corners. Grid squares that are not internal or external are *mixed*. Let X be the set of internal and mixed grid squares, i.e., the squares where the approximate oracle did not return NO on all corners. We define R_i^- as the union of the cells in X plus all external grid cells within $C + \sqrt{2}$ grid cells of a mixed cell in X , i.e., the external cells whose Euclidean distance to a mixed cell is at most $C + \sqrt{2}$ times the edge length of a grid cell; refer to (iv) for the definition of C . Once again, to speed up the query (see below), in every grid row, we also add to R_i^- all the cells “sandwiched” between the leftmost and the rightmost identified by the above rules, even if they don’t satisfy the rules themselves.

By the approximation guarantee of the oracle, an external grid cell is also external relative to the exact oracle. Notice that the set of mixed cells according to the approximate oracle includes all the mixed cells according to the exact one. Therefore, the set of cells our algorithm includes in R_i^- is a superset of those included in R_i^- , so R_i^- contains R_i . (Note that including the additional “sandwiched” cells only enlarges R_i^- and therefore does not affect this statement.) We now estimate $\text{area}(R_i^- \setminus R_i)$.

Observe that $R_i((1 - 2\Delta)\beta) \subset (1 + 2\Delta/c_1)R_i$, by Lemma 4.1, provided $0 < 2\Delta/c_1 < 1/c_1$ or $0 < \Delta < 1/2$.

Consider a corner q of a grid cell in X for which the approximate oracle did not return a NO. By definition, $\tilde{E}(q) \geq (1 - \Delta)\beta$ and therefore $E(q) \geq (1 - \Delta)\tilde{E}(q) \geq (1 - \Delta)^2\beta > (1 - 2\Delta)\beta$. Hence, every such corner lies in $R_i((1 - 2\Delta)\beta)$ and therefore in $(1 + 2\Delta/c_1)R_i$. Therefore, by convexity of $(1 + 2\Delta/c_1)R_i$, any square in X without NO corners is contained in it. A square of X that is not fully contained in $(1 + 2\Delta/c_1)R_i$ must necessarily have at least one corner outside of $(1 + 2\Delta/c_1)R_i$ and at least one inside it. Therefore, such a square must meet the boundary of $(1 + 2\Delta/c_1)R_i$. Hence any cell of X lies either in $(1 + 2\Delta/c_1)R_i$ or in $(1 + 2\Delta/c_1)R_i$ Minkowski expanded by a disk of radius $\sqrt{2}$ grid cells. In particular, all cells of R_i^- are contained in $(1 + 2\Delta/c_1)R_i$ expanded by a disk of radius $C + 2\sqrt{2}$ grid cells. (Notice that this also applies to the “sandwiched” cells, as the latter set is convex.)

Recall that $\text{area}((1 + 2\Delta/c_1)R_i) = (1 + 2\Delta/c_1)^2 \text{area}(R_i)$ and therefore $\text{area}((1 + 2\Delta/c_1)R_i \setminus R_i) = (4\Delta/c_1 + 4\Delta^2/c_1^2) \text{area}(R_i) = \Theta(\Delta) \text{area}(R_i)$. We now estimate the area gained by expanding $(1 + 2\Delta/c_1)R_i$ by the disk of radius $C + 2\sqrt{2}$ grid cells. (The additional $\sqrt{2}$ appears because a grid cell lying at a distance x from a set lies fully in an expansion of the set by a disk of radius $x + \sqrt{2}$.)

Invoking Steiner formula,¹³ the area of the expanded set, in grid squares, is the area of $(1 + 2\Delta/c_1)R_i$ plus the area of the disk $(\pi(C + 2\sqrt{2})^2 = O(1)$ grid cells) plus the radius of the latter

¹³Steiner formula states that $\text{area}(A \oplus tB) = \text{area}(A) + t \cdot \text{perimeter}(B) + t^2 \cdot \text{area}(B)$, if A is convex, B is a unit circle, and tB is a circle of radius t , in the plane; $\oplus$ denotes Minkowski addition of sets.

times the perimeter of the former, which are $(C + 2\sqrt{2}) = O(1)$ and $\Theta(k)$, respectively. Since the expanded set contains all the grid cells of R_i^{--} , we conclude that it gains at most $\Theta(k)$ additional grid squares over its unexpanded version, i.e., $\text{area } \Theta(\delta) \text{area}(R_i)$, as claimed.

Thus

$$\begin{aligned} \text{area}(R_i^{--} \setminus R_i) &\leq \text{area}((1 + 2\Delta/c_1)R_i \setminus R_i) + \Theta(\delta)\text{area}(R_i) \\ &\leq (\Theta(\Delta) + \Theta(\delta))\text{area}(R_i) \\ &= \Theta(\delta)\text{area}(R_i), \end{aligned}$$

provided $\Delta = O(\delta)$.

Therefore, $\text{area}(R_i^{--} \setminus R_i^{++}) = \text{area}(R_i^{--} \setminus R_i) + \text{area}(R_i \setminus R_i^{++}) = O(\delta)\text{area}(R_i)$, as claimed, for a sufficiently small $\Delta = O(\delta)$, concluding the proof of the approximation correctness for our algorithm for preprocessing R_i for point location.

Moreover, since both R_i^{++} and R_i^{--} contain only sets of contiguous squares (if any at all) in every grid row, a point can be located in each set in $O(1)$ time provided *floor* operation is available, by identifying the row containing the point, using its y -coordinate and then comparing its x -coordinate to that of the leftmost and rightmost grid cell in the row.

The following theorem summarizes the main result of this section:

THEOREM 5.1. *Fix positive integer α and $\beta > 1$. Given a set $\mathcal{S}$ of n transmitters (all of power 1) and a parameter $\delta > 0$, we can do the following in total time $O(n\delta^{-2.5} \log^4 n \log \log n)$.¹⁴ Compute, for each transmitter s_i , two sets R_i^{++} and R_i^{--} (each represented as a union of a collection of grid cells), such that (i) $R_i^{++} \subseteq R_i \subseteq R_i^{--}$ and (ii) $\text{area}(R_i^{--} \setminus R_i^{++}) = O(\delta)\text{area}(R_i)$. Moreover, given a query point q and R_i^{--} and R_i^{++} , we can determine in $O(1)$ time whether q is in R_i^{++} , $R_i^{--} \setminus R_i^{++}$, or $\mathbb{R} \setminus R_i^{--}$. Given q , we can thus approximately determine if it can receive any transmitter, and if so, which one, in $O(\log n)$ time.*

6 CONCLUDING REMARKS

We described several algorithms that combine computational geometry techniques and methods of computer algebra to obtain very fast batched SINR diagram point-location queries.

Besides speeding up the construction time of known structures, we would like to identify further applications of batched point location to other problems studied in the SINR model.

We note that our results are general, in the sense that analogous results can be obtained for diagrams that are induced by other inequalities similar to the SINR inequality.

Finally, on a larger scale, we are interested in other instances where algebraic and geometric tools can be combined to achieve significant results that are seemingly impossible without it.

APPENDICES

A TOOLS AND DEFINITIONS FROM THE WORLD OF COMPUTER ALGEBRA

In this section, we state several well-known results from computer algebra and refer to Reference [28] for details; see Reference [10] for an alternative presentation.

A.1 Definitions

A *univariate polynomial* $A(x)$ of degree at most n is defined by an expression of the form $A(x) = a_0 + a_1x + a_2x^2 + \dots + a_nx^n$. The tuple $C(A) := \langle a_0, a_1, \dots, a_n \rangle$ is the *coefficient*

¹⁴Our previous discussion implies that one can set the approximation parameter Δ for the approximate oracle to $\Delta = c\delta$, with $c = c(\alpha, \beta)$, a sufficiently small constant. In fact, setting $\Delta = \min\{\delta, (\beta - 1)/6, (1 - \beta^{-1/\alpha})/6, 1/14\}$ satisfies all the requirements.

representation of A . Given a set $X = \{x_0, \dots, x_n\}$ of $n + 1$ numbers, we let $V(A) = V(A, X)$ denote $\{(x_0, A(x_0)), \dots, (x_n, A(x_n))\}$, a *value representation* of A .

A *fractional function* $F(x)$ of degree at most n is a function that can be written in the form $F(x) = A(x)/B(x)$ for two polynomials A and B of degree at most n , with $B \neq 0$. *Coefficient representation* of F is simply $(C(A), C(B))$, the pair of coefficient representations of its numerator and denominator.

Analogously, a *bivariate polynomial* $B(x, y)$ of degree at most n in each variable is defined by an expression of the form $B(x, y) = \sum_{0 \leq i, j \leq n} b_{ij} x^i y^j$. The set $\{\langle i, j, b_{ij} \rangle\}$ is the *coefficient representation* of B . The set $\{(x_i, y_i, B(x_i, y_i))\}$ is a *value representation* of B , for a suitably large set of points $\{(x_i, y_i)\} \in \mathbb{R}^2$.

One can analogously define bivariate fractional functions.

A.2 Univariate Tools and Facts

The following two results can be found in Corollaries 10.8 and 10.12 in Reference [28], respectively.

FACT A.1 (UNIVARIATE MULTIPOINT EVALUATION AND INTERPOLATION). *If A is a univariate polynomial of degree at most n and X a tuple of $n + 1$ distinct numbers, then the coefficient representation clearly determines A , but so does its value representation $V(A, X)$.*

Multipoint Evaluation *Given the coefficient representation $C(A)$ of A , A can be evaluated at the $n + 1$ points of X , yielding $V(A, X)$ using $O(n \log^2 n \log \log n)$ arithmetic operations. If X contains m points, the computation can be done in $O((n + m) \log^2 n \log \log n)$ arithmetic operations.¹⁵*

Interpolation *Given a value representation $V(A, X)$ of a polynomial A of degree at most n on a set X of $n + 1$ points, the coefficient representation $C(A)$ of A , can be constructed using $O(n \log^2 n \log \log n)$ arithmetic operations.*

Given two univariate polynomials A and B of degree at most n each, let their *product* AB be the polynomial D defined by $D(x) = A(x)B(x)$ for all x . Multiplying two polynomials in value representation is easy: for each $x_i \in X$, $D(x_i) = A(x_i) \cdot B(x_i)$, by definition. Somewhat surprisingly, one can also quickly multiply univariate polynomials in coefficient representation (see Theorem 8.23 in Reference [28]):

FACT A.2 (FAST MULTIPLICATION OF UNIVARIATE POLYNOMIALS). *Given two univariate polynomials A, B of degree at most n each, in coefficient representation, one can construct their product AB , using $O(n \log n \log \log n)$ arithmetic operations.*

Note. The FFT is essentially an evaluation of a polynomial on a special set of complex numbers, the n th roots of unity (more precisely, the $2^{\lceil \log_2 n \rceil}$ th roots of unity). It can be performed in $O(n \log n \log \log n)$ operations, due to the special structure of the set of roots of unity; technically, the $\log \log n$ term appearing in many of our bounds is due to the assumption that the appropriate primitive root of unity is not available explicitly; otherwise the bounds improve by a factor of $\log \log n$. The Inverse FFT reverses the process (by using a variant of the FFT code), reconstructing a polynomial from its values at the roots of unity, again in $O(n \log n \log \log n)$ operations. The fast polynomial multiplication algorithm is an FFT, followed by point-wise multiplication, followed by the Inverse FFT.

The following observation has been made by previous authors; we provide a proof for completeness; we follow the construction from Reference [23].

¹⁵The second statement follows immediately from the first by extending X to length n with dummy data if $m < n$, or cutting it into chunks of length n if $m > n$.

LEMMA A.3 (SUM OF FRACTIONS [23]). *Given a set of n fractional functions $P_i(x)/Q_i(x)$ of constant degree each, in coefficient representation, the coefficient representation of their sum can be constructed using $O(n \log^2 n \log \log n)$ arithmetic operations.*

PROOF. Given two fractions $\frac{A}{B}$ and $\frac{C}{D}$, their sum can be written as $\frac{AD+BC}{BD}$. Therefore, given two fractions of degree at most d in coefficient representation, we can obtain their sum, which is a fraction of degree at most $2d$, in coefficient representation, using three polynomial multiplications and one addition, for a total of $O(d \log d \log \log d)$ arithmetic operations.

We start with the initial fractions P_i/Q_i , $i = 1, \dots, n$ and add them in pairs, then add the resulting sums in pairs, and so forth. One round costs $O(n \log n \log \log n)$ operations, since the sum of the degrees of intermediate polynomials at each level is $O(n)$. In each round, the number of fractions reduces by a factor of two, so the procedure stops after $O(\log n)$ rounds. Hence the final result will be obtained after $O(n \log^2 n \log \log n)$ operations. $\square$

COROLLARY A.4 (EVALUATION OF SUM OF FRACTIONS). *Given a set of n fractional functions $P_i(x)/Q_i(x)$ of constant degree each, in coefficient representation, and a set of m values x_j , one can compute the m values $\sum_i P_i(x_j)/Q_i(x_j)$ in time $O((n+m) \log^2 n \log \log n)$.*

PROOF. Symbolically sum the fractions to construct one degree- $O(n)$ fraction, using Lemma A.3, evaluate the numerator and denominator at each x_j using Fact A.1, and divide. $\square$

A.3 Bivariate Tools and Facts

It would be helpful to have analogous tools for the bivariate case. The difficulty is that a general bivariate polynomial of degree at most n in each variable is described by $(n+1)^2 = \Theta(n^2)$ coefficients, so its explicit coefficient representation necessarily has size $\Theta(n^2)$. Thus, in general one cannot expect that direct analogs of Facts A.2 and A.2, and Lemma A.3 exist, at least not in the same form.

We define the *product* $D = AB$ of two bivariate polynomials A and B similarly to that of univariate ones. Once again, multiplying two polynomials in value form is easy by multiplying the corresponding values. Somewhat surprisingly, one can also quickly multiply bivariate polynomials in coefficient representation:

FACT A.5 (FAST MULTIPLICATION OF BIVARIATE POLYNOMIALS; [10, 26] OR [28, COROLLARY 8.28]). *Given two bivariate polynomials $A(x, y)$, $B(x, y)$ of degree at most n in each variable, in coefficient representation (which has size $\Theta(n^2)$ in general), one can construct their product AB , using $O(n^2 \log n \log \log n)$ arithmetic operations.*

Notice that, once we have the multiplication result, one can use the same reasoning as in Lemma A.3 to obtain the following:

FACT A.6 (SUM OF BIVARIATE FRACTIONS). *Given a set of n bivariate fractional functions $P_i(x, y)/Q_i(x, y)$ of constant degree each, in coefficient representation, the coefficient representation of their sum $P(x, y)/Q(x, y)$, can be constructed using $O(n^2 \log n \log \log n)$ arithmetic operations.*

PROOF. For completeness, we outline the analysis. We repeat the reasoning from the proof of Lemma A.3, using Fact A.5 for bivariate multiplication. Round i involves polynomials of degree $O(2^i)$ with degrees summing to $O(n)$, so the number of arithmetic operations used in round i is $O(2^i n \log n \log \log n)$. This sums up to $O(n^2 \log n \log \log n)$ for the $\log n$ rounds, as claimed. $\square$

FACT A.7 (BIVARIATE EVALUATION AND INTERPOLATION OVER A GRID [10, 26]). *Consider a bivariate polynomial $A(x, y)$ of degree at most n in each variable, and two sets $X = \{x_i\}$ and $Y = \{y_j\}$ of n numbers each.*

Grid Evaluation Given the $O(n^2)$ coefficients of A , A can be evaluated at the points of $X \times Y$ (i.e., values $A(x_i, y_j)$ computed, for all combinations of i and j) in time $O(n^2 \log^2 n \log \log n)$.

Interpolation from Values on the Grid Given $(n+1)^2$ values $v_{ij} = A(x_i, y_j)$, one can reconstruct the coefficient representation of A in time $O(n^2 \log^2 n \log \log n)$.

Note. We outline one way to achieve the above bounds: Write $A(x, y) = \sum_{j=0}^n B_j(x)y^j$, where B_j is a univariate polynomial of degree at most n and use the univariate algorithm from Fact A.1 to multi-evaluate B_j on all values x_i ; the remaining arithmetic can be done in $O(n^2)$ operations; this results in an algorithm with $O(n^2 \log^2 n \log \log n)$ arithmetic operations.

Interpolation can be accomplished by invoking univariate tools as well. Treating $A(x_i, y)$ as a univariate polynomial in y , for every x_i separately, and applying univariate interpolation (Fact A.1) allows us to write it as $\sum B_j(x_i)y^j$, where $B_j(x)$ is as above and we have obtained the values $B_j(x_i)$ for all i, j . We now reconstruct each B_j by using univariate interpolation again. Once we are done, we have all the coefficients of all the B_j 's, which are the coefficients of A .

This requires $2n$ invocations of the univariate interpolation routine for a total of $O(n^2 \log^2 n \log \log n)$ arithmetic operations, as claimed.

COROLLARY A.8 (EVALUATION OF SUM OF FRACTIONS). Given a set of n bivariate fractional functions $P_i(x, y)/Q_i(x, y)$ of constant degree each, in coefficient representation and two sets $X = \{x_j\}$, $Y = \{y_k\}$ of n values each, one can compute the $n \times n$ values $\sum_i P_i(x_j, y_k)/Q_i(x_j, y_k)$ in time $O(n^2 \log^2 n \log \log n)$.

PROOF. Combine the fractions into a single fraction using Fact A.6 and then evaluate it on the grid $X \times Y$ using Fact A.7. $\square$

There does not seem to be an easy analogue of fast univariate evaluation on an arbitrary set of values, as in Fact A.1, though a simple observation (see, for example, Reference [25]) shows that one can evaluate a bivariate polynomial on n^2 points in roughly cubic time by just extending each set of n points into a grid and using the algorithm from Fact A.7. However, a stronger result is known.

FACT A.9 (GENERAL BIVARIATE MULTIPOINT EVALUATION [25]). If $A(x, y)$ is a bivariate polynomial of degree at most n in x and degree at most m in y , in coefficient representation (nm coefficients), A can be evaluated simultaneously at N different points of $\mathbb{R}^2$ in $O((N + nm)m^{\omega_2/2-1+\varepsilon})$ arithmetic operations, for any $\varepsilon > 0$, where ω_2 is such that the product of any $n \times n$ and $n \times n^2$ matrices can be performed in $O(n^{\omega_2})$ time; it is known that $\omega_2 < 3.334$.

B DEFINITIONS AND TOOLS FROM GEOMETRY AND DATA STRUCTURES

B.1 Orthogonal Range Search Trees

We say that point $p \in \mathbb{R}^2$ is *dominated* by point $q \in \mathbb{R}^2$ if $x(p) \leq x(q)$ and $y(p) \leq y(q)$; we write $p \leq q$.

FACT B.1. Given a set $P = \{p_1, \dots, p_n\}$ of n points in $\mathbb{R}^2$, a (two-dimensional) orthogonal range tree is a data structure of size $O(n \log n)$ that supports queries of the following type:¹⁶ Given a query point q , report all the points $p \in P$ dominated by q , in time $O(k + \log^2 n)$, where k is the size of the answer. More specifically, internally the range tree is a collection of $O(n \log n)$ canonical subsets of P , so that each point of P lies in $O(\log^2 n)$ canonical subsets, and the answer to the query q is represented as a disjoint union of $O(\log^2 n)$ such subsets.

¹⁶A range tree is actually more powerful, but we will not need all of its power here.

FACT B.2 (DOMINATING PAIR DECOMPOSITION). *Given two sets $P, Q \in \mathbb{R}^2$ each of size n , it is possible to construct a collection of pairs of subsets $\{(P_i, Q_i)\}$ with the following properties:*

- (1) $P_i \subseteq P, Q_i \subseteq Q$.
- (2) *For each pair of points $p \in P, q \in Q$ with $p \leq q$, there is a unique i so that $p \in P_i$ and $q \in Q_i$.*
- (3) *Each pair in $\bigcup_i (P_i \times Q_i)$ is a dominating pair.*
- (4) *There are at most $O(n \log n)$ pairs of sets in the collection.*
- (5) $\sum_i (|P_i| + |Q_i|) = O(n \log^2 n)$.
- (6) *The collection of pairs can be constructed in time $O(n \log^2 n)$.*

The above construction is not difficult if one starts with a range tree for P . The P_i are the canonical sets of the range tree. Initialize each Q_i to the empty set. Then execute a query for each $q \in Q$ and add q to the set Q_i associated with each canonical set P_i that participates in the query.

B.2 Voronoi Diagrams and Friends

Given a set P of n points in the plane, one can partition the plane into convex polygonal regions, so that each region $V(p)$ consists entirely of points q closer to a given point $p \in P$ than to any other point of P , where “closer” is measured with respect to the usual Euclidean distance. Boundaries of the regions contain points where the nearest neighbor is non-unique. The partition is called the *Voronoi diagram* of P and the sets $V(p)$ are the *Voronoi cells*. The Voronoi diagram consists of n cells and a linear number of edges and vertices separating the cells.

FACT B.3 (TWO-DIMENSIONAL VORONOI DIAGRAM). *Given an n -point set $P \subset \mathbb{R}^2$, the Voronoi diagram of P has linear complexity and, in time $O(n \log n)$, can be constructed and preprocessed for logarithmic-time point location queries: Given a point q , determine the point $p \in P$ whose region $V(p)$ contains q .*

If the Euclidean distance $|q - p|$ in the above construction is replaced by $|q - p|/w_p$, where w_p is a positive multiplicative weight of p , we obtain the *multiplicatively weighted Voronoi diagram*, which is similar to the Euclidean version, except that the complexity can be quadratic in n , in the worst case, the edges are circular arcs, and in general the region of a given Voronoi site can be disconnected and may contain holes.

FACT B.4 (TWO-DIMENSIONAL MULTIPLICATIVELY WEIGHTED VORONOI DIAGRAM [6]). *Given an n -point set $P \subset \mathbb{R}^2$ with positive weights w_p , the multiplicatively weighted Voronoi diagram of P has complexity $O(n^2)$ and, in time $O(n^2)$, can be constructed and preprocessed for logarithmic-time point-location queries: Given a point q , determine the point $p \in P$ whose region $V(p)$ contains q .*

The one-dimensional analogs of the Euclidean and multiplicatively weighted Voronoi diagrams for an n -point set have linear complexity and can be constructed and preprocessed for point location in time $O(n \log n)$; the weighted version is specifically addressed in Reference [5]:

FACT B.5 (ONE-DIMENSIONAL MULTIPLICATIVELY WEIGHTED VORONOI DIAGRAM). *Given an n -point set $P \subset \mathbb{R}$ with positive weights w_p , the multiplicatively weighted Voronoi diagram of P has linear complexity and, in time $O(n \log n)$, can be constructed and preprocessed for logarithmic-time point-location queries: Given a point $q \in \mathbb{R}$, determine the point $p \in P$ whose region $V(p)$ contains q .*

We will also need the following slightly esoteric variant of the weighted diagram (which, of course, is a generalization of Fact B.5):

FACT B.6 (ONE-DIMENSIONAL SLICE OF A TWO-DIMENSIONAL MULTIPLICATIVELY WEIGHTED VORONOI DIAGRAM). *Given an n -point set $P \subset \mathbb{R}^2$ with positive weights w_p , the multiplicatively*

weighted Voronoi diagram of P restricted to a given line ℓ has complexity $O(n)$ and, in time $O(n \log n)$, can be constructed and preprocessed for logarithmic-time point-location queries: Given a point q , determine the point $p \in P$ whose region $V(p)$ contains q .

PROOF. Without loss of generality, after a suitable rigid transformation, we can assume that the line coincides with the x -axis. The squared weighted-distance function $f_p(x)$ from point $(x, 0)$ to site p is given by

$$f_p(x) = |p - (x, 0)|^2 / w_p^2 = ((x(p) - x)^2 + y(p)^2) / w_p^2,$$

which is a quadratic function of x . As observed, for example, by Edelsbrunner and Seidel [12], the minimization diagram of these functions coincides with the desired restricted weighted Voronoi diagram. Being quadratic, the graphs of two such functions $f_p(x)$ and $f_q(x)$ intersect at most twice and therefore their lower envelope and minimization diagram has linear complexity and can be constructed by a straightforward $O(n \log n)$ time algorithm [27, Theorem 6.1]. $\square$

We will also need an approximate version of the Voronoi diagram. We will only state the weighted version of the problem, as that is what we need for our purposes [19, Theorem 2.16]:

FACT B.7 (APPROXIMATE TWO-DIMENSIONAL MULTIPLICATIVELY WEIGHTED NEAREST NEIGHBOR [19]). *Given an n -point set $P \subset \mathbb{R}^2$ with positive weights w_p and a positive number ϵ , one can preprocess it into a data structure of space $O(n\epsilon^{-6} \log^4 n)$ in time $O(n\epsilon^{-6} \log^7 n)$ to support $O(\log n + 1/\epsilon^3)$ time queries of the form: Given a point q , return $p' \in P$ so that $|p' - q|/w_{p'} \leq (1 + \epsilon)|p^* - q|/w_{p^*}$, where p^* is the point in P minimizing $|p - q|/w_p$.*

Alternatively, there is a data structure of space $O(n\epsilon^{-6} \log^4 n)$ built in time $O(n\epsilon^{-6} \log^7 n)$ with $O(\log(n/\epsilon))$ query time.

ACKNOWLEDGMENTS

The authors would like to acknowledge the significant help of Sarel Har-Peled in matters of approximation and the tremendous assistance of Guillaume Moroz in matters of computer algebra. B.A. would also like to thank Pankaj K. Agarwal for general encouragement and moral support.

REFERENCES

- [1] Deepak Ajwani, Saurabh Ray, Raimund Seidel, and Hans Raj Tiwary. 2007. On computing the centroid of the vertices of an arrangement and related problems. In *Proceedings of the 10th International Workshop on Algorithms and Data Structures (WADS'07)*. 519–528. DOI: http://dx.doi.org/10.1007/978-3-540-73951-7_45
- [2] M. Andrews and M. Dinitz. 2009. Maximizing capacity in arbitrary wireless networks in the SINR model: Complexity and game theory. In *Proceedings of the 28th IEEE International Conference on Computer Communications (INFOCOM'09)*. 1332–1340.
- [3] Boris Aronov and Matthew J. Katz. 2015. Batched point location in SINR diagrams via algebraic tools. In *Proceedings of Automata, Languages, and Programming—42nd International Colloquium (ICALP'15), Part I*. 65–77. DOI: http://dx.doi.org/10.1007/978-3-662-47672-7_6 See also arXiv:1412.0962 [cs.CG], 2014.
- [4] Rom Aschner, Gui Citovsky, and Matthew J. Katz. 2014. Exploiting geometry in the SINRk model. In *Proceedings of Algorithms for Sensor Systems—10th International Symposium on Algorithms and Experiments for Sensor Systems, Wireless Networks and Distributed Robotics (ALGOSENSORS'14)*. 125–135. DOI: http://dx.doi.org/10.1007/978-3-662-46018-4_8
- [5] Franz Aurenhammer. 1986. The one-dimensional weighted Voronoi diagram. *Inform. Process. Lett.* 22, 3 (1986), 119–123. [http://dx.doi.org/10.1016/0020-0190\(86\)90055-4](http://dx.doi.org/10.1016/0020-0190(86)90055-4).
- [6] Franz Aurenhammer and Herbert Edelsbrunner. 1984. An optimal algorithm for constructing the weighted Voronoi diagram in the plane. *Patt. Recog.* 17, 2 (1984), 251–257. DOI: [http://dx.doi.org/10.1016/0031-3203\(84\)90064-5](http://dx.doi.org/10.1016/0031-3203(84)90064-5)
- [7] Franz Aurenhammer, Rolf Klein, and D.-T. Lee. 2013. *Voronoi Diagrams and Delaunay Triangulations*. World Scientific. <http://www.worldscientific.com/worldscibooks/10.1142/8685>.
- [8] Chen Avin, Asaf Cohen, Yoram Haddad, Erez Kantor, Zvi Lotker, Merav Parter, and David Peleg. 2017. SINR diagram with interference cancellation. *Ad Hoc Netw.* 54 (2017), 1–16. DOI: <http://dx.doi.org/10.1016/j.adhoc.2016.08.003>

- [9] Chen Avin, Yuval Emek, Erez Kantor, Zvi Lotker, David Peleg, and Liam Roditty. 2012. SINR diagrams: Convexity and its applications in wireless networks. *J. ACM* 59, 4, Article 18 (2012), 34 pages. DOI : <http://dx.doi.org/10.1145/2339123.2339125>
- [10] Dario Bini and Victor Y. Pan. 1994. *Polynomial and Matrix Computations*, Vol. 1: Fundamental Algorithms. Birkhäuser.
- [11] M. de Berg, O. Cheong, M. van Kreveld, and M. H. Overmars. 2008. *Computational Geometry: Algorithms and Applications* (3rd ed.). Springer-Verlag, Berlin. <http://www.cs.ruu.nl/geobook>.
- [12] H. Edelsbrunner and R. Seidel. 1986. Voronoi diagrams and arrangements. *Discrete Comput. Geom.* 1 (1986), 25–44.
- [13] O. Goussevskaia, M. M. Halldórsson, R. Wattenhofer, and E. Welzl. 2009. Capacity of arbitrary wireless networks. In *Proceedings of the 28th IEEE International Conference on Computer Communications (INFOCOM'09)*. 1872–1880.
- [14] O. Goussevskaia, Y. A. Oswald, and R. Wattenhofer. 2007. Complexity in geometric SINR. In *Proceedings of the 8th ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc'07)*. 100–109.
- [15] P. Gupta and P. R. Kumar. 2000. The capacity of wireless networks. *IEEE Trans. Inform. Theor.* 46, 2 (2000), 388–404.
- [16] M. M. Halldórsson. 2012. Wireless scheduling with power control. *ACM Trans. Algorithms* 9, 1 (2012), 7:1–7:20.
- [17] M. M. Halldórsson and P. Mitra. 2011. Wireless capacity with oblivious power in general metrics. In *Proceedings of the 22nd ACM-SIAM Symposium on Discrete Algorithms (SODA'11)*. 1538–1548.
- [18] M. M. Halldórsson and R. Wattenhofer. 2009. Wireless communication is in APX. In *Proceedings Automata, Languages and Programming—36th International Colloquium (ICALP'09), Part I*. 525–536.
- [19] Sarel Har-Peled and Nirman Kumar. 2015. Approximating minimization diagrams and generalized proximity search. *SIAM J. Comput.* 44, 4 (2015), 944–974. DOI : <http://dx.doi.org/10.1137/140959067>
- [20] Erez Kantor, Zvi Lotker, Merav Parter, and David Peleg. 2011. The topology of wireless communication. In *Proceedings of the 43rd ACM Symposium on Theory of Computing (STOC'11)*. 383–392. DOI : <http://dx.doi.org/10.1145/1993636.1993688>
- [21] T. Kesselheim. 2011. A constant-factor approximation for wireless capacity maximization with power control in the SINR model. In *Proceedings of the 22nd ACM-SIAM Symposium on Discrete Algorithms (SODA'11)*. 1549–1559.
- [22] Zvi Lotker and David Peleg. 2010. Structure and algorithms in the SINR wireless model. *SIGACT News* 41, 2 (2010), 74–84. DOI : <http://dx.doi.org/10.1145/1814370.1814391>
- [23] Guillaume Moroz and Boris Aronov. 2016. Computing the distance between piecewise-linear bivariate functions. *ACM Trans. Algorithms* 12, 1 (2016), 3:1–3:13.
- [24] T. Moscibroda and R. Wattenhofer. 2006. The complexity of connectivity in wireless networks. In *Proceedings of the 25th IEEE International Conference on Computer Communications (INFOCOM'06)*. 25–37. DOI : <https://doi.org/10.1109/infocom.2006.23>
- [25] Michael Nüsken and Martin Ziegler. 2004. Fast multipoint evaluation of bivariate polynomials. In *Proceedings 12th European Symposium on Algorithms (ESA'04)*. 544–555. DOI : http://dx.doi.org/10.1007/978-3-540-30140-0_49
- [26] Victor Y. Pan. 1994. Simple multivariate polynomial multiplication. *J. Symb. Comput.* 18, 3 (1994), 183–186. DOI : <http://dx.doi.org/10.1006/jscs.1994.1042>
- [27] M. Sharir and P. K. Agarwal. 1995. *Davenport-Schinzel Sequences and Their Geometric Applications*. Cambridge University Press, New York. <http://us.cambridge.org/titles/catalogue.asp?isbn=0521470250>.
- [28] Joachim von zur Gathen. 1999. *Modern Computer Algebra*. Cambridge University Press, Cambridge.
- [29] P.-J. Wan, X. Jia, and F. F. Yao. 2009. Maximum independent set of links under physical interference model. In *Proceedings of the 4th International Conference on Wireless Algorithms, Systems, and Applications (WASA'09)*. 169–178.
- [30] Martin Ziegler. 2003. Fast relative approximation of potential fields. In *Proceedings of the 8th International Workshop on Algorithms and Data Structures (WADS'03)*. 140–149. DOI : http://dx.doi.org/10.1007/978-3-540-45078-8_13

Received January 2017; revised April 2018; accepted April 2018