Batched Point Location in SINR Diagrams via Algebraic Tools

Boris Aronov[1] and Matthew J. Katz[2(✉)]

[1] Department of Computer Science and Engineering,
Polytechnic School of Engineering, New York University, Brooklyn, NY 11201, USA
boris.aronov@nyu.edu

[2] Department of Computer Science, Ben-Gurion University, Beer-Sheva, Israel
matya@cs.bgu.ac.il

Abstract. The *SINR 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 papers. 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 paper, 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.

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.

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

1 Introduction

The *SINR (Signal to Interference plus Noise Ratio) model* attempts to more realistically predict whether a wireless transmission is received successfully, in

Work on this paper by B.A. has been partially supported by NSF Grants CCF-11-17336 and CCF-12-18791. Work on this paper by M.K. has been partially supported by grant 1045/10 from the Israel Science Foundation. A more complete version of this paper is available on arXiv [3].

M.M. Halldórsson et al. (Eds.): ICALP 2015, Part I, LNCS 9134, pp. 65–77, 2015.
DOI: 10.1007/978-3-662-47672-7_6

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 [19].

Let $\mathcal{S} = \{s_1, \ldots, s_n\}$ be a set of n points in the plane representing n transmitters. Let $p_i > 0$ be the transmission power of transmitter s_i, $i = 1, \ldots, n$. 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:

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

where $|a-b|$ denotes the Euclidean distance between points a and b, and $\alpha > 0$, $\beta > 1$,[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 w.r.t. transmitter s_i).

Notice that, since $\beta > 1$, a necessary condition for q to receive s_i is that $p_i/|q-s_i| > p_j/|q-s_j|$, for any $j \neq i$. In particular, in the *uniform power setting* where $p_1 = p_2 = \cdots = p_n$, 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 $R_\emptyset$ consisting of all points where none of the transmitters is received. This partition is called the *SINR diagram* of $\mathcal{S}$. 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 $j \neq i$ [4]. Then $R_i \subset V_i$.

In a seminal paper, Avin et al. [6] 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 [17].

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

[1] In this paper, we assume $\beta > 1$. A variant of our techniques applies also when $\beta < 1$: up to $1/\beta$ receivers can be heard simultaneously, multiple nearest neighbors need to be identified as the candidates, and the algorithms slow down correspondingly.

which the ratio $\frac{1}{\sqrt[\alpha]{p}}|q - s|$ 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. 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 $R_i^+ \subset R_i \subset R_i^- \subset V_i$,[2] 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 $q \notin R_i^-$, 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 [6,17], it is guaranteed that the uncertain answer is only given infrequently, namely that $\mathit{area}(R_i^- \setminus R_i^+) \leq \varepsilon \cdot \mathit{area}(R_i)$, for a suitable parameter $\varepsilon > 0$. In the second [17], 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 now briefly summarize previous work. Observing the difficulty of answering point-location queries exactly, Avin et al. [6] resorted to approximate query answers in the *uniform power* setting. Given an $\varepsilon > 0$ they build a data structure in total time $O(n^2/\varepsilon)$ and space $O(n/\varepsilon)$ that can be wrong only in a region of area $\varepsilon \cdot \mathit{area}(R_i)$ 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. [17] 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 prespecified $\varepsilon > 0$. 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.

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

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 paper we focus on *batched* point location in the SINR model. That is, given a set $\mathcal{Q}$ of m query points, determine for each point $q \in \mathcal{Q}$ whether it receives one of the transmitters in $\mathcal{S}$, 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 sophisticated geometric data structures and tools from computer algebra for multipoint evaluation, interpolation, and fast multiplication of polynomials and rational functions. For example, consider 1-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. [17] mentioned above in which the bounds depend on additional geometric parameters which cannot be bounded as a function of n or ε.

1.1 Related Work

The papers most relevant to ours are those by Avin et al. [6] and Kantor et al. [17] discussed above. Avin et al. [5] 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 [5] for further details.)

Gupta and Kumar [11] 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 papers and results dealing with the maximum capacity and scheduling problems. Goussevskaia et al. [10] showed that both problems are NP-complete, even in the uniform power setting. Goussevskaia et al. [9], Halldórsson and Wattenhofer [14], and Wan et al. [24] 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 [9] 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 [13] 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., [2,12,13,18,22].

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 speed-ups in geometric algorithms by expressing a subproblem in algebraic terms [1,20,21]. The algebraic tools themselves are mostly classical ones, such as Fast Fourier Transform, fast polynomial multiplication, multipoint evaluation, and interpolation [7,23]; see [3, 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 [8]; refer to [3, Appendix B]. One very recent result we need is that of Har-Peled and Kumar [15] that, as a special case, allows one to build a compact data structure for approximating 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 object.

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

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. In general, we present algorithms for both the uniform-power and non-uniform-power settings, where the algorithms of the former type are usually somewhat simpler.

- 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 3.2.
- We can also facilitate exact batch queries when queries or transmitters form a grid; we omit the details in this version; see [3].

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., $p_1 = \cdots = p_n = 1$), 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

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

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

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

Theorem 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 $q \in Q$, in the same amount of time.

Observe that f is a sum of n low-degree fractional functions of a single real variable q, so according to [3, Corollary 1], 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 $E(q_1), \ldots, E(q_m)$ 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 $O((n+m)\log^2 n \log\log n)$. □

2.2 Arbitrary Power

We proceed in a similar manner, except the construction of the multiplicatively weighted Voronoi diagram on a line, which is more subtle; see [3].

Theorem 2. *For any fixed positive even 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 $O((n+m)\log^2 n \log\log n)$.*

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 "s" (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. In the second version, 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. The details of these two versions are omitted due to space limitations; see [3].

Finally, in the third version (Section 3.2), 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, if $\tilde{E}(q)$ is sufficiently greater than β, then q surely receives s, if $\tilde{E}(q)$ is sufficiently smaller than β, then q surely does not receive s, and otherwise, q may or may not receive some transmitter (i.e., $\tilde{E}(q)$ lies in the *uncertainty interval*). 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 [8]), where $n = |\mathcal{S}|$. 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 [4].

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 $E(q)$ 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.

Due to space constraints, we omit the discussion of transmitters on a grid and of receivers on a grid; see [3]. Therefore in the remainder of the section we focus on the last version of the problem.

3.2 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, 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

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

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

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

so $2^{-\alpha/2} \tilde{f}(q) \le f(q) \le \tilde{f}(q)$, and therefore $2^{-\alpha/2} \tilde{E}(q) \le E(q) \le 2^{\alpha/2} \tilde{E}(q)$. 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 $\tilde{E}(q) \ge 2^{\alpha/2}\beta$, we know that $E(q) \ge \beta$ and the signal from the unique candidate transmitter $s(q)$ *is* received. 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. For intermediate values of $\tilde{E}(q)$, we cannot definitely determine whether $s(q)$'s signal is received at q.

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 with each receiver, for a total cost of $O((n+m)\log n)$ [8]. 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 [4,8], which is too much for $m \approx n$. We provide an alternative below.

We show how to compute the values $\tilde{f}(q_1), \ldots, \tilde{f}(q_m)$ 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 tilted version of) the orthogonal range search tree [8] (see, [Section B.1, Fact 14]), we can construct a pair decomposition $\{(\mathcal{S}_i, \mathcal{Q}_i)\}$ of small size, so that each pair (s,q) with $s \in W_q$ appears in exactly one product $\mathcal{S}_i \times \mathcal{Q}_i$.

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,

$$\tilde{f}(q, \mathcal{S} \cap W_q) = \sum_{i: q \in \mathcal{Q}_i} \tilde{f}(q, \mathcal{S}_i), \tag{1}$$

by the definition of the pair decomposition. The number of terms in the last sum is $O(\log^2 n)$. Notice that $\tilde{f}(q, \mathcal{S}_i)$, for a fixed i, is a sum of small fractional *univariate* functions, with $|\mathcal{S}_i|$ 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 $(\mathcal{Q}_i, \mathcal{S}_i)$, we use [3, Corollary 1] to evaluate $\tilde{f}(q, \mathcal{S}_i)$ on each $q \in \mathcal{Q}_i$ in total time $O((|\mathcal{Q}_i| + |\mathcal{S}_i|)\log^2 |\mathcal{S}_i| \log\log |\mathcal{S}_i|) = O((|\mathcal{Q}_i| + |\mathcal{S}_i|)\log^2 n \log\log n)$. This gives us all the summands of (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 $\sum_i (|\mathcal{Q}_i| + |\mathcal{S}_i|)\log^2 n \log\log n = (\sum_i (|\mathcal{Q}_i| + |\mathcal{S}_i|))\log^2 n \log\log n = O((m+n)\log^4 n \log\log n)$.

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. *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 $O((m+n)\log^4 n \log\log n)$. 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 an algorithm of Har-Peled and Kumar [15,16], by setting $\varepsilon = 2^{1/2} - 1$ (see [3]), yielding the following:

Theorem 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.

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 given $\varepsilon > 0$. 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 square, with the norm $|\ldots|_k$ 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 $1 + \varepsilon = 1 + \alpha\Theta(k^{-2})$, 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 5. *For a positive ε, 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 $O((m+n)\varepsilon^{-1/2} \log^4 n \log\log n)$. 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 $\tilde{E}(q) \geq (1+\varepsilon)\beta$, 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 6. *For a positive ε, any fixed positive even integer α, and any $\beta > 1+\varepsilon$,*[3] *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 $O(n\varepsilon^{-6}\log^7 n + m\varepsilon^{-1/2}\log^4 n \log\log n)$ and $O(n\varepsilon^{-6}\log^4 n + m\varepsilon^{-1/2}\log^2 n)$ space: For each $q \in \mathcal{Q}$, 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.*

4 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.

We believe that Theorems 5 and 6 can be applied to speed up the preprocessing stage of existing point-location results. Consider, e.g., the data structure presented by Avin et al. [6] for a set of n uniform-power transmitters, whose construction time is $O(n^2/\delta)$. This data structure is actually a collection of n data structures, one per transmitter, where 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 $\mathit{area}(R_i^- \setminus R_i^+) \leq \delta \cdot \mathit{area}(R_i)$, see the definitions in the introduction. The construction of DS_i is based on the convexity and fatness of region R_i and consists of two stages. In the first, explicit estimates for the radii of the largest disk centered at s_i and contained in R_i and the smallest such disk containing R_i are obtained, by applying a binary-search-like procedure (beginning with the distance between s_i to its nearest (other) transmitter in $\mathcal{S}$), where each comparison is resolved by explicitly evaluating the SIN ratio at some point q and comparing it to β, i.e., by an *in/out* test. In the second stage, a $1/\delta \times 1/\delta$ grid 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 $\Theta(\log n + 1/\delta)$ in/out tests per transmitter, at a cost of $\Theta(n)$ operations each; the high cost of each test is the bottleneck.

We believe that it is possible to speed up the algorithm by constructing the n individual data structures in parallel. During the construction, we will form $O(\log n + 1/\delta)$ batches of n queries each, and use Theorem 5 to deal with each of them in near-linear time. The only problem is that our query answers are not exact, but approximate; for some queries, instead of "in" or "out," we answer "maybe. We think that there is a way to overcome this problem, but we leave it for a full version.

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

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

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 further applications where algebraic and geometric tools can be combined to achieve significant improvements.

Acknowledgments. B.A. would like to acknowledge the help of Sariel Har-Peled in matters of approximation and of Guillaume Moroz in matters of algebra. He would also like to thank Pankaj K. Agarwal for general encouragement.

References

1. Ajwani, D., Ray, S., Seidel, R., Tiwary, H.R.: On computing the centroid of the vertices of an arrangement and related problems. In: Dehne, F., Sack, J.-R., Zeh, N. (eds.) WADS 2007. LNCS, vol. 4619, pp. 519–528. Springer, Heidelberg (2007)
2. Andrews, M., Dinitz, M.: Maximizing capacity in arbitrary wireless networks in the SINR model: Complexity and game theory. In: INFOCOM, pp. 1332–1340 (2009)
3. Aronov, B., Katz, M.J.: Batched point location in SINR diagrams via algebraic tools (2014). arXiv:1412.0962 [cs.CG]
4. Aurenhammer, F., Edelsbrunner, H.: An optimal algorithm for constructing the weighted Voronoi diagram in the plane. Pattern Recognition 251–257 (1984)
5. Avin, C., Cohen, A., Haddad, Y., Kantor, E., Lotker, Z., Parter, M., Peleg, D.: SINR diagram with interference cancellation. In: SODA, pp. 502–515 (2012)
6. Avin, C., Emek, Y., Kantor, E., Lotker, Z., Peleg, D., Roditty, L.: SINR diagrams: Convexity and its applications in wireless networks. J. ACM **59**(4), 18:1–318:4 (2012)
7. Bini, D., Pan, V.Y.: Polynomial and Matrix Computations: Fundamental Algorithms, vol. 1. Birkhauser Verlag, Basel (1994)
8. de Berg, M., Cheong, O., van Kreveld, M., Overmars, M.H.: Computational Geometry: Algorithms and Applications, 3rd edn. Springer-Verlag, Berlin (2008)
9. Goussevskaia, O., Halldórsson, M.M., Wattenhofer, R., Welzl, E.: Capacity of arbitrary wireless networks. In: INFOCOM, pp. 1872–1880 (2009)
10. Goussevskaia, O., Oswald, Y.A., Wattenhofer, R.: Complexity in geometric SINR. In: MobiHoc, pp. 100–109 (2007)
11. Gupta, P., Kumar, P.R.: The capacity of wireless networks. IEEE Trans. Information Theory **46**(2), 388–404 (2000)
12. Halldórsson, M.M.: Wireless scheduling with power control. ACM Transactions on Algorithms **9**(1) (2012)
13. Halldórsson, M.M., Mitra, P.: Wireless capacity with oblivious power in general metrics. In: SODA, pp. 1538–1548 (2011)
14. Halldórsson, M.M., Wattenhofer, R.: Wireless communication is in APX. In: Albers, S., Marchetti-Spaccamela, A., Matias, Y., Nikoletseas, S., Thomas, W. (eds.) ICALP 2009, Part I. LNCS, vol. 5555, pp. 525–536. Springer, Heidelberg (2009)
15. Har-Peled, S., Kumar, N.: Approximating minimization diagrams and generalized proximity search. SIAM J. Comput. Accepted for publication. http://sarielhp.org/p/12/wann/wann.pdf
16. Har-Peled, S., Kumar, N.: Approximating minimization diagrams and generalized proximity search. In: FOCS, pp. 717–726 (2013)

17. Kantor, E., Lotker, Z., Parter, M., Peleg, D.: The topology of wireless communication. In: STOC, pp. 383–392 (2011)
18. Kesselheim, T.: A constant-factor approximation for wireless capacity maximization with power control in the SINR model. In: SODA, pp. 1549–1559 (2011)
19. Lotker, Z., Peleg, D.: Structure and algorithms in the SINR wireless model. SIGACT News **41**(2), 74–84 (2010)
20. Moroz, G., Aronov, B.: Computing the distance between piecewise-linear bivariate functions. In: SODA, pp. 288–293 (2012)
21. Moroz, G., Aronov, B.: Computing the distance between piecewise-linear bivariate functions. ACM Transactions on Algorithms (2013). Accepted for publication arXiv:1107.2312 [cs.CG]
22. Moscibroda, T., Wattenhofer, R.: The complexity of connectivity in wireless networks. In: INFOCOM, pp. 23–29 (2006)
23. von zur Gathen, J.: Modern Computer Algebra. Cambridge University Press, Cambridge (1999)
24. Wan, P.-J., Jia, X., Yao, F.: Maximum independent set of links under physical interference model. In: Liu, B., Bestavros, A., Du, D.-Z., Wang, J. (eds.) WASA 2009. LNCS, vol. 5682, pp. 169–178. Springer, Heidelberg (2009)