

Q Cells in Wireless Networks

Martin Haenggi, *Fellow, IEEE*

Abstract—For a given set of transmitters such as cellular base stations or WiFi access points, is it possible to analytically characterize the set of locations that are “covered” in the sense that users at these locations experience a certain minimum quality of service? In this paper, we affirmatively answer this question, by providing explicit simple outer bounds and estimates for the coverage manifold. The key geometric elements of our analytical method are the *Q cells*, defined as the intersections of a small number of disks. The Q cell of a transmitter is an outer bound to the service region of the transmitter, and, in turn, the union of Q cells is an outer bound to the coverage manifold. In infinite networks, connections to the meta distribution of the signal-to-interference ratio allow for a scaling of the Q cells to obtain accurate estimates of the coverage manifold.

Index Terms—Wireless networks, geometry, interference, coverage, stochastic geometry, meta distribution.

I. INTRODUCTION

A. Problem Formulation

Coverage is a key performance metric of wireless networks. It has been studied extensively in terms of the fraction of the area of a region that is “covered” by wireless service, see, e.g., [1], [2] and references therein. This paper addresses the more challenging problem of analytically characterizing the *coverage manifold* $\mathcal{C} \subset \mathbb{R}^2$, which is the subset of the plane that is covered in the sense that the probability that the signal-to-interference ratio (SIR) exceeds a threshold θ is at least u . Specifically, for a set of transmitter locations, the problem is to find tight bounds and estimates for $\mathcal{C}$ that are described using only elementary geometric shapes—polygons and disk segments—and can be calculated very efficiently.

B. Approach and Related Work

Our approach is based on the so-called *Q cells*, which consists of the locations in which a serving transmitter provides enough power such that no single interferer can cause the location to be uncovered. Conversely, for any location outside the Q cell of a transmitter, there exists an interferer that causes the location to be uncovered.

Q cells generalize the standard Voronoi cells and provide an outer bound on the region that is covered by a transmitter. They are related to the SIR cells introduced in [1, Ch. 5], defined as the set of locations where the SIR exceeds a certain threshold. The SIR cells are random sets that depend on the instantaneous values of the fading coefficients in the propagation model, which results in a much more complex and time-dependent structure. In contrast, the Q cells do not depend on the fading coefficients but merely the fading statistics. As a result, they are deterministic (for a given set of transmitters).

Simplified interference modeling using only the nearest interferer(s) dates back to at least [3], whose model became known as the “Wyner model”. More recently, [4] used the concept of *dominant nodes* to bound the so-called transmission capacity of ad hoc networks, where an interferer is said to be dominant if its interference alone can cause an outage. In this terminology, our Q cells are defined as the locations for which there are no dominant interferers. In [5], the total interference is approximated using the nearest interferer only, and [6] uses a dominant-interferer approach for THz networks, with the additional assumption that non-dominant interferers by themselves cannot cause outages. Considering only the dominant interferer results can yield asymptotically exact or scaling results in terms of the tail of the interference or in the low-outage regime, as established in [7], [8].

Another line of work models the additional interference from non-dominant interferers by adding a deterministic term [9] or a Gaussian random variable [10]. This is essential for path loss exponents near 2, where the interference from distant transmitters is significant [11]. These articles focus on a form of coverage probability or coverage area fraction, not estimates of the coverage manifold.

Related in their objective but rather different in their approach are the works on using machine learning (ML) techniques for coverage estimation, using convolutional neural networks [12] and generative adversarial networks [13]. They require large data sets and extensive training, cannot be used to obtain bounds, and do not lend themselves for simple interpretations, in contrast to our direct geometric method.

The SIR meta distribution (MD) [14], [15] is also relevant for our work since it equals the fraction of the network area that is covered in infinite networks where the transmitters are modeled as a point process. We will take advantage of this connection to turn the coverage manifold bound that is based on Q cells into an accurate estimate of $\mathcal{C}$. Conversely, if the MD is not available, it can be upper bounded by the area fraction covered by the Q cells. Of particular interest is the notion of *separability* [16] for Poisson networks, which asserts that the MD separates into the product of a function of θ and a function of u for a certain set of values of (θ, u) . The separable regime corresponds to the *stringent QoS regime* in this paper, which is the regime of practical interest for the quality-of-service (QoS) parameters θ and u .

C. Contributions

- We show that the QoS parameters θ and u can be lumped together into a single *stringency* parameter σ . Based on σ , we define lax and stringent QoS regimes.
- We introduce *Q cells*, which depend on σ and generalize Voronoi cells. They are defined as the intersection of

disks given by of Möbius transformed half-planes (instead of the intersection of half-planes forming the Voronoi cells), and their union outer bounds the coverage manifold (Thm. 1).

- Focusing on locations with two equidistant interferers, we show that cutting off corner pieces of the Q cell results in *refined Q cells* that are tighter outer bounds on the coverage regions (Thm. 2).
- For infinite networks, we give a simple universal lower bound on the fraction of uncovered users (Thm. 3) and explore and exploit the connection to the SIR MD, which leads to a universal upper bound on the MD (Cor. 1) and permits the scaling of Q cells to obtain accurate estimates of the coverage manifold (Recipe 1).

II. PRELIMINARIES

A. Notation

Since the map $x \mapsto x/(1+x)$, introduced in [17] as the MH (Möbius homeomorphic) transform, and its inverse appear frequently throughout the paper, we use the compact notation

$$x^\oplus \triangleq x/(1+x), \quad x \in \mathbb{R} \setminus \{-1\}$$

$$x^\ominus \triangleq x/(1-x), \quad x \in \mathbb{R} \setminus \{1\}.$$

Important identities include $(x^\oplus)^\ominus \equiv x$ and $(1/x)^\oplus \equiv 1-x^\oplus$.

The open disk centered at x with radius r is denoted by $b(x, r)$, the cumulative distribution function (cdf) of a random variable X is F_X , its complement (ccdf) is $\bar{F}_X$, and its inverse (quantile function) is F_X^{-1} . For $n \in \mathbb{N}$, $[n] \triangleq \{1, 2, \dots, n\}$. For an uncountable set $B \subset \mathbb{R}^d$, ∂B is its boundary, and $|B|$ is its d -dimensional Lebesgue measure. For a countable set $\mathcal{X}$, $|\mathcal{X}|$ is its cardinality.

For the standard mapping from $\mathbb{R}^2 \mapsto \mathbb{C}$ and its inverse, we use $\sharp$ and $\flat$, respectively, i.e., for $(x, y) \in \mathbb{R}^2$, $(x, y)^\sharp \triangleq x + iy$ where $i^2 = -1$, and for $z \in \mathbb{C}$, $z^\flat \triangleq (\Re(z), \Im(z))$. The origin in $\mathbb{R}^2$ is denoted as $o = o^\flat$.

B. Three Circular Lemmas

Lemma 1 (Intersections of disks) *The intersection of n disks*

$$Q = \bigcap_{i \in [n]} b(x_i, r_i)$$

is either empty or has a boundary ∂Q formed by $1 \leq \check{n} \leq 2(n-1)$ arcs. If $\check{n} = 1$, it is a single disk, if $\check{n} = 2$, Q consists of two disk segments, and if $\check{n} > 2$, it consists of a convex polygon plus $\check{n}$ disk segments.

Proof: Follows from elementary geometry. ■

Lemma 2 (Equal distance ratio circle) *The set of locations $(x, y) \in \mathbb{R}^2$ such that the distance ratio to the two points $(0, 0)$ and $(b, 0)$ equals $\rho \neq 1$ describes a circle with center $(c(\rho), 0)$ where $c(\rho) = b/(1-\rho^2)$ and radius $r(\rho) = |(b/\rho)(\rho^2)^\ominus| = \rho|c(\rho)|$. For $\rho = 1$ (equidistance), the locations form the line $x = b/2$.*




Fig. 1: With the desired transmitter at the origin $\circ$ and an interferer at $\star$, the circles represent the locations where the line-of-sight SIR (squared distance ratio) has values 2, 3, 5, 10, 10 (left circles) and 1/2, 1/3, 1/5, 1/10 (right circles). The equidistant points to $\circ$ and $\star$ form the line $x = 1/2$, on which the SIR is 1.

Proof: The prescribed distance ratio implies

$$\frac{(x-b)^2 + y^2}{x^2 + y^2} = \rho^2,$$

which, for $\rho \neq 1$, can be rearranged to

$$\left(x - \frac{b}{1-\rho^2}\right)^2 + y^2 = \frac{b^2 \rho^2}{(1-\rho^2)^2}.$$

This is the equation for a circle with the given radius and center. Re-writing the expression for the center as $c(\rho) = (b/\rho^2)(\rho^2)^\ominus$ reveals that center and radius are related by a factor $\pm\rho$. For $\rho = 1$, the equidistant locations form the line $x = b/2$. ■

Remarks.

- If ρ is replaced by $1/\rho$, the radius stays the same, and the center point is reflected at $x = b/2$, i.e., $c(1/\rho) = b - c(\rho)$.
- The result is easily extended to general locations $\mathbf{x}$ and $\mathbf{x}'$ instead of $(0, 0)$ and $(b, 0)$, respectively, by setting $b = \|\mathbf{x}' - \mathbf{x}\|$, and placing the circle center at $\mathbf{x} + \frac{c(\rho)}{b}(\mathbf{x}' - \mathbf{x})$.
- The insight that the set of locations at which the distance ratio to two fixed points is constant is a circle dates back more than 2200 years—the Greek mathematician Apollonius defined a circle in this way [18, Ch. 2].

Fig. 1 shows the equal distance ratio circles for $b = 1$ and $\rho^2 = 2, 3, 5, 10, 20$. Interpreted in a wireless context, the circles represent the locations on the plane at which an SIR of exactly ρ^2 is achieved in a line-of-sight (LOS) environment if the desired transmitter is at $(0, 0)$ and the interferer at $(b, 0)$. The LOS propagation model assumed uses a path loss law of r^{-2} for a transmission over distance r and no fading. Letting $r_1 = \sqrt{x^2 + y^2}$ and $r_2 = \sqrt{(x-b)^2 + y^2}$ be the distances from (x, y) to $(0, 0)$ and $(b, 0)$, respectively, the SIR is given by the squared distance ratio

$$\text{SIR}_{(x,y)} = \frac{r_1^{-2}}{r_2^{-2}} = \frac{(x-b)^2 + y^2}{x^2 + y^2}.$$

Hence the locations at which $\text{SIR} = \rho^2$ form a circle as given by Lemma 2.

For $\rho > 1$, $\text{SIR} > \rho^2$ is achieved at locations inside the circle, while for $\rho < 1$, it is achieved at locations outside the circle. This fact is compactly represented using a Möbius transform on the complex plane.

Lemma 3 (Möbius transform characterization) *Let $z, z' \in \mathbb{C}$ denoted the locations of the desired transmitter and interferer in the complex plane. For $\rho \neq 1$, let*

$$f_\rho(z) \triangleq c + \rho|c| \frac{z - \rho c}{z + \rho c}, \quad z \in \mathbb{C}, \quad (1)$$

where $c = \|z' - z\|/(1 - \rho^2)$. For $\rho = 1$, $f_1(z) \triangleq z + \|z' - z\|/2$. Then, for $w \in \mathbb{C}$,

$$\frac{|w - z'|}{|w - z|} > \rho \Leftrightarrow w \in z + f_\rho(H^-) e^{i \arg(z' - z)}, \quad (2)$$

where $H^- \triangleq \{z \in \mathbb{C} : \Re(z) < 0\}$ is the open left half-plane of $\mathbb{C}$.

Proof: The modulus of the fraction in (1) is 1 on the imaginary axis while its argument ranges from $-\pi$ to π , so the axis is mapped onto the circle of radius $\rho|c|$ centered at c . The half-plane H^- is mapped inside the circle if $\rho > 1$, since $\Re(z) < 0$ decreases the modulus for $c < 0$. Conversely, $\rho < 1$, $\Re(z) < 0$ increases the modulus since $c > 0$. The rotation and translation in (2) place the disk in the correct position relative to z, z' . ■

Remarks.

- The transform (1) is a generalization of the bilinear transform that maps H^- to $|z| < 1$ commonly used in digital signal processing for the continuous time-to-discrete time mapping of system functions.
- Applying the transform to the imaginary axis, with $z = 0$ and $z' = b$, yields the circle in Lemma 2 for $\rho \neq 1$. For $\rho = 1$, the definition of f_1 ensures that this equivalence extends to the case of equidistant points to 0 and b , and that $b/(1 + \rho)$ lies in the image of the imaginary axis for any ρ .

III. COVERAGE AND STRINGENCY

A. Network Model

Let $\mathcal{X} = \{x_1, x_2, \dots\} \subset \mathbb{R}^2$ be a countable set of transmitter locations with at least two elements. For $y \in \mathbb{R}^2$, we let $\mathcal{X}_{-y} \triangleq \{x_i - y\}_{i \in \mathbb{N}}$ be the translated set, and we let r_i be the i -th smallest element of $\|\mathcal{X}_{-y}\|$, such that r_1 is the distance from y to the closest transmitter.

With nearest-transmitter association, a path loss law of the form $\ell(r) = r^{-\alpha}$ (where α is the path loss exponent), and unit transmit powers, the SIR at location y is

$$\text{SIR}_y = \frac{h_1 r_1^{-\alpha}}{\sum_{i>1} h_i r_i^{-\alpha}} = \frac{1}{\sum_{i>1} \frac{1}{H_i} \left(\frac{r_1}{r_i}\right)^\alpha}, \quad (3)$$

where the coefficients $H_i \triangleq h_1/h_i$, $i > 1$, are identically distributed with cdf F_H . All $(h_i)_{i \in \mathbb{N}}$ are independent with $\mathbb{E}h_i = 1$, and for $i > 1$, they are identically distributed. This formulation allows for different fading distributions for the desired transmitter and the interferers. In particular, we

consider Nakagami fading with parameter $p > 0$ for h_1 and parameter $q > 0$ for h_i , $i > 1$, which we refer to as Nakagami- (p, q) fading¹. In this case

$$F_H(x) = B_{p,q} \left(\frac{px}{px+q} \right) = B_{p,q}((x_q^L)^\oplus), \quad (4)$$

where

$$B_{p,q}(x) \triangleq \frac{\Gamma(p+q)}{\Gamma(p)\Gamma(q)} \int_0^x t^{p-1} (1-t)^{q-1} dt, \quad x \in [0, 1],$$

is the regularized incomplete beta function², i.e., the cdf of the beta distribution. It satisfies $B_{p,q}(1-x) \equiv 1 - B_{q,p}(x)$, which reflects the fact that the distribution of $1/H_i$ is related to that of H_i by flipping p and q . Its inverse³ is denoted as $B_{p,q}^{-1}$, i.e., $B_{p,q}(B_{p,q}^{-1}(u)) \equiv u$. This general fading model covers the entire range from severe fading (small p and q , modeling systems with no or little diversity) to fading with little variance (modeling systems with high diversity).

Symmetric case: $p = q$. In the symmetric case, all $(h_i)_{i \in \mathbb{N}}$ are independent and identically distributed (iid), and the fading model is the standard Nakagami- m model with $p = q = m$. For Rayleigh fading ($m = 1$), $F_H(x) = x^\oplus$, while for general m , the H_i are beta prime distributed with $F_H(x) = B_{m,m}(x^\oplus)$.

Special asymmetric cases: We consider two special asymmetric cases that can be used to upper- and lower-bound the performance if the exact fading distributions are unknown. The first ($p = 1$, $q \rightarrow \infty$) is pessimistic and assumes Rayleigh fading in the desired link and no fading in the interfering links. In this case, the H_i are exponential. The second ($p \rightarrow \infty$, $q = 1$) is optimistic and assumes no fading in the desired link and Rayleigh fading in the interfering links. In this case, the $1/H_i$ are exponential.

As an alternative to the SIR, it is often advantageous to use the *signal fraction* (SF) [17], defined as the desired signal power divided by the total received power, i.e., $\text{SF} = \text{SIR}^\oplus$ and

$$\{\text{SF}_y > t\} \Leftrightarrow \{\text{SIR}_y > t^\ominus\}, \quad t \in [0, 1]. \quad (5)$$

The inverse relationship $t = \theta^\oplus$ can also be expressed as $t = \theta$ MH, where MH refers to the Möbius homeomorphic scale that compresses $\mathbb{R}^+$ to $[0, 1)$. It offers advantages for integration and visualization⁴ [17].

B. Coverage

Definition 1 (Coverage) *A location $y \in \mathbb{R}^2$ is said to be covered if $\mathbb{P}(\text{SIR}_y > \theta) > u$, where $\theta \geq 0$ and $u \in [0, 1)$ are the QoS parameters. The set of all covered locations is the coverage manifold*

$$\mathcal{C} \triangleq \{y \in \mathbb{R}^2 : \mathbb{P}(\text{SIR}_y > \theta) > u\}.$$

The set of locations covered when x is the desired transmitter is denoted by $\mathcal{C}_x$, i.e., $\mathcal{C} = \bigcup_{x \in \mathcal{X}} \mathcal{C}_x$.

¹Nakagami- m refers to the symmetric case $p = q = m$.

²In Matlab, `betainc(x, p, q)`.

³In Matlab, `betaincinv(u, p, q)`.

⁴See <https://stogblog.net/2020/09/18/unmasking-distributions-with-infinite-support/>. The frequently used dB scale does the opposite—it extends $(0, 1]$ to $(-\infty, 0]$, which makes it impossible to draw a complete plot.

Frequently a different definition of “coverage” is employed in the literature, in which y is “covered” if $\text{SIR}_y > \theta$. As mentioned in Subsec. I-B, such a definition based on the instantaneous SIR would make coverage vary randomly over small spatial scales (as the SIR depends on the small-scale fading), and the covered locations would form a very complex, highly non-contiguous random set with wave-length level granularity, which would make the resulting maps non-reproducible and unusable for planning or for assessing user experience⁵. In contrast, coverage maps used in practice are deterministic and based on *local averaging*, where small-scale fading is averaged out [19], [20]. This implies that fading affects coverage only via its statistics, consistent with Def. 1.

The coverage manifold critically depends on a parameter called *stringency* that is a function of θ , u , and the fading distribution.

C. Stringency

Definition 2 (Stringency and QoS regimes) *The stringency of the QoS pair (θ, u) is defined as*

$$\sigma(\theta, u) \triangleq \frac{\theta}{F_H^{-1}(1-u)}. \quad (6)$$

If $\sigma < 1$, we call the QoS requirements lax; if $\sigma > 1$, they are stringent; for $\sigma = 1$, they are balanced.

The stringency is proportional to the SIR threshold parameter θ and monotonically increasing with u , with $\sigma(\theta, u) \rightarrow \infty$ as $u \rightarrow 1$. The next lemma gives concrete expressions for the stringency for different fading models.

Lemma 4 (Impact of fading on stringency) *For Nakagami- (p, q) fading,*

$$\sigma(\theta, u) = \frac{\theta}{\frac{q}{p}(B_{p,q}^{-1}(1-u))^{\ominus}}. \quad (7)$$

For symmetric Rayleigh fading ($p = q = 1$), this simplifies to $\sigma(\theta, u) = \theta u^{\ominus}$, and for $p = q = 1/2$,

$$\sigma(\theta, u) = \theta \tan^2(\pi u/2).$$

For increasingly severe Nakagami- m fading,

$$\sigma(\theta, u) \xrightarrow{m \rightarrow 0} \begin{cases} 0 & u \in [0, 1/2) \\ \theta & u = 1/2 \\ \infty & u > 1/2, \end{cases}$$

while, without fading ($m \rightarrow \infty$), $\sigma(\theta, u) = \theta$.

In the pessimistic asymmetric case ($p = 1$, $q \rightarrow \infty$),

$$\sigma(\theta, u) = \frac{\theta}{-\log u},$$

while in the optimistic asymmetric case ($p \rightarrow \infty$, $q = 1$),

$$\sigma(\theta, u) = -\theta \log(1-u).$$

Proof: For Nakagami- (p, q) fading, the quantile function

$$F_H^{-1}(u) = \frac{q}{p}(B_{p,q}^{-1}(u))^{\ominus} \quad (8)$$

⁵More details on the definitions of “coverage” can be found at <https://stogblog.net/2020/09/24/what-is-coverage/>.




Fig. 2: Stringency $\sigma(\theta, u)$ in dB for Nakagami-3 fading. To cover the full range $\theta \in \mathbb{R}^+$, the MH scale is used, i.e., $\theta \text{ MH} = \theta^{\oplus} \in [0, 1)$. The red line is the contour at 0 dB, on which $\theta \text{ MH} = B_{3,3}^{-1}(1-u)$.

is obtained by inverting (4). For $p = q = 1$, this reduces to $F_H^{-1}(u) = u^{\ominus}$, and the result follows from $1/(1-u)^{\ominus} \equiv u^{\ominus}$.

For $p = q = 1/2$, the underlying beta cdf is $B_{1/2,1/2}$, which is the arcsine distribution, hence $F_H^{-1}(u) = (\sin^2(\pi u/2))^{\ominus}$, and the result follows from $\tan^2 \varphi \equiv 1/(\cos^2(\varphi))^{\ominus}$.

For $m \rightarrow 0$, we start from the Gauss hypergeometric function representation [21, p. 263, (6.6.8)] to derive the asymptotic form

$$\begin{aligned} B_{m,m}(x) &= \frac{\Gamma(2m)}{\Gamma^2(m)} \frac{x^m}{m} {}_2F_1(m, 1-m; 1+m; x) \\ &\sim \frac{m}{2} \frac{x^m}{m} \left(1 + mx + \frac{m}{2}x^2 + \dots\right), \quad m \rightarrow 0 \\ &\sim \frac{x^m}{2}, \quad m \rightarrow 0, x \in [0, 1). \end{aligned}$$

Hence $F_H^{-1}(u) \sim ((2u)^{1/m})^{\ominus}$, $u < 1/2$, $F_H^{-1}(1/2) = 1$ (which holds for all $m > 0$), while $F_H^{-1}(u) \rightarrow \infty$ for $u > 1/2$, and it follows that

$$F_H^{-1}(u) \xrightarrow{m \rightarrow 0} \begin{cases} 0 & u \in [0, 1/2) \\ 1 & u = 1/2 \\ \infty & u > 1/2. \end{cases}$$

Without fading, $H_i \equiv 1$, so $F_H(x) = \mathbf{1}(x \geq 1)$, $F_H^{-1}(1-u) = \mathbf{1}(u \in [0, 1))$, and $\sigma(\theta, u) = \theta$.

In the asymmetric cases, the results follow from H_i or $1/H_i$ being exponential, respectively. ■

Fig. 2 shows the stringency $\sigma(\theta, u)$ for Nakagami-3 fading, using the MH scale for θ to display it in its entire range.

Remarks.

- Comparing the two asymmetric cases, the stringency in the pessimistic case is significantly larger. The gap between the two diverges as $u \rightarrow 1$, and so does the ratio for both $u \rightarrow 0$ and $u \rightarrow 1$. For all u , the ratio is larger than 2.
- In the case of increasingly severe fading, reliabilities less than $1/2$ become trivially satisfied, while those above $1/2$ become impossible to achieve. The limit as $m \rightarrow 0$

can be expected to hold for general fading models whose variance goes to ∞ .

IV. BOUNDING THE COVERAGE MANIFOLD USING Q CELLS

A. Q Cells

Definition 3 (Q cell) The Q cell of transmitter $x \in \mathcal{X}$, denoted as Q_x , is defined as

$$Q_x(\rho) \triangleq \left\{ y \in \mathbb{R}^2 : \frac{\|\mathcal{X} \setminus \{x\} - y\|}{\|x - y\|} > \rho \right\}, \quad \rho > 0, \quad (9)$$

where $\|\mathcal{X} \setminus \{x\} - y\|$ is the distance to the interferer closest to y .

The next lemma shows that the Q cell can be analytically described using the Möbius transform put forth in Lemma 3.

Lemma 5 (Complex Q cell characterization) For $z \in \mathcal{X}^\sharp$, let

$$Z_z(\rho) \triangleq z + \bigcap_{z' \in \mathcal{X}^\sharp \setminus \{z\}} f_\rho(H^-) e^{j \arg(z' - z)}, \quad (10)$$

with f_ρ given in (1). Then $Q_x(\rho) \equiv Z_{x^\sharp}^b(\rho)$.

Proof: The right side of (10) consists of the locations at which the distance ratio $|z'|/|z| > \rho$ for each interferer z' . This is equivalent to the condition in (9) since if the condition is satisfied for the nearest interferer, it is satisfied for all interferers. ■

Remarks.

- For $\rho = 1$, the Q cell Q_x corresponds to the interior of the Voronoi cell V_x , which is defined as the intersection of the half planes delineated by the perpendicular bisectors of z and z' .
- For $\rho \neq 1$, the Q cell is the intersection of the regions whose boundaries are the equal distance ratio circles in Lemma 2.
 - For $\rho > 1$, the half planes get morphed into disks per f_ρ , and the Q cell is the intersection of these $|\mathcal{X}| - 1$ disks.
 - For $\rho < 1$, the Q cell consists of $\mathbb{R}^2$ (or $\mathbb{C}$) minus a union of exclusion disks for each interferer.

For a given stringency, increasing α reduces ρ if $\rho > 1$, resulting in larger Q cell. If $\rho < 1$, increasing α increases ρ , which enlarges the exclusions disks.

- Not every interferer contributes to the shaping of the Q cell. Those that do are termed *strong interferers*.

Definition 4 (Strong and dominant interferers) x' is a strong interferer to x for a given ρ if the removal of x' increases the Q cell.

A dominant interferer at location y is an interferer that by itself is strong enough to cause y to be uncovered.

Remarks.

- For $\rho < 1$, every interferer is strong.
- For $\rho > 1$, the intersection of the $|\mathcal{X}| - 1$ disks in the definition of Q_x can be reduced to the intersection of




Fig. 3: Example Q cell for the transmitter at x , shown in red. Its boundary is formed by the 4 black circles, at distance ratio $\rho = \sqrt{2}$ with respect to each strong interferer whose locations are marked with a red cross. The red lines are the edges of the Voronoi diagram, and the white lines inside the central Voronoi cell V_o mark the second-order Voronoi edges. The entire square region is colored according to $\mathbb{P}(\text{SIR}_y > 1)$ for symmetric Rayleigh fading and $\alpha = 4$, using the color map on the right. It is apparent that all locations y for which $\mathbb{P}(\text{SIR}_y > 1) > 0.8$ are inside the Q disk.

the disks of the strong interferers to x . By Lemma 1, their number determines the number of arcs $\tilde{n}_x$ that define the boundary of Q_x . $\tilde{n}_x$ equals the number of Delaunay neighbors of x . The number of dominant interferers at y equals the number of disks that y is outside of. For $y \in Q_x$, there are no dominant interferers.

- For $\rho \downarrow 1$, the number of strong interferers to x tends to the number of sides of the Voronoi cell of x . As $\rho \rightarrow \infty$, the number may decrease to 1.
- Denote by $V_{xx'}$ the second-order Voronoi cell of x and x' , i.e., the set of locations that have x and x' as their two closest points in $\mathcal{X}$ [22], [23]. If

$$y \in (V_x \cap V_{xx'}) \setminus Q_x,$$

x' is the dominant interferer at y .

- By construction, the corner points of the Q cell lie on an edge of the second-order Voronoi diagram. These are the locations that have exactly two dominant interferers at equal distance. This property will be exploited in Subsec. IV-C.

An example Q cell for $\rho = \sqrt{2}$ with the dominant interferers and their equal distance ratio circles is shown in Fig. 3. As stipulated, the corner points of the Q cell lie on edges of the second-order Voronoi diagram, which are shown in white.

Fig. 4 shows the Q cells for different values of ρ for a set $\mathcal{X}$ with 25 elements.

B. Coverage Outer Bounds

The key missing piece to bound the coverage manifold is the connection between θ and u and the radius parameter ρ of the Q cells. By the following theorem, the connection is established by equating ρ^α with the stringency σ .




Fig. 4: Q cells for $\rho = 1.25, 1.50, 1.75, 2.00, 2.50$ for a set of 25 transmitters. The blue lines are the Voronoi cell boundaries.

Theorem 1 (Outer bounds on coverage cells and manifold)
The coverage cell of transmitter x is bounded as

$$C_x \subseteq Q_x(\rho(\theta, u)), \quad (11)$$

where $\rho(\theta, u) = \sigma(\theta, u)^{1/\alpha}$, and the coverage manifold is bounded as

$$\mathcal{C} \subseteq \mathcal{Q}, \quad (12)$$

where

$$\mathcal{Q} \triangleq \bigcup_{x \in \mathcal{X}} Q_x(\rho(\theta, u))$$

is the union of all Q cells.

Proof: For a transmitter x and a location y , let $r_1 = \|x - y\|$ and $r_2 = \|\mathcal{X} \setminus \{x\} - y\|$ be the serving distance and nearest interferer distance, respectively. With $v \triangleq r_2/r_1$,

$$\begin{aligned} \mathbb{P}\left(\frac{h_1 r_1^{-\alpha}}{h_2 r_2^{-\alpha}} > \theta\right) &= \mathbb{P}(H v^\alpha > \theta) \\ &= 1 - F_H(\theta v^{-\alpha}). \end{aligned} \quad (13)$$

$1 - F_H(\theta v^{-\alpha}) > u$ implies

$$v^{-\alpha} < \frac{F_H^{-1}(1 - u)}{\theta} \iff v > \sigma(\theta, u)^{1/\alpha}.$$

Hence a distance ratio smaller than $\rho(\theta, u)$ to the nearest interferer is sufficient for y to be uncovered. The bound on the coverage manifold is a consequence of the bounds of the individual coverage cells. ■

Remarks.

- For $\rho > 1$, to achieve a reliability u , a user needs to be located inside all circles with distance ratio ρ (Lemma 2).
- For Rayleigh fading, if coverage is expressed in terms of the SF threshold t , $\rho(t^\ominus, u)^\alpha = t^\ominus u^\ominus$ by (5), which

shows a striking symmetry of ρ in the reliability u and the SF threshold t . The stringent regime is simply $t + u > 1$.

- If only the nearest interferer for each location y is considered, (11) is an equality.
- In the lax regime $\rho < 1$, $\mathcal{C} \subseteq \mathcal{Q}$ trivially holds since $\mathcal{Q} = \mathbb{R}^2$. For $\rho = 1$, $\mathcal{Q}$ is still essentially $\mathbb{R}^2$ (minus the Voronoi edges).
- Since the Q cells are disjoint in the stringent regime, a user cannot be handed over from one transmitter to the next without losing coverage. Only in the lax regime, seamless coverage is possible.

The Q cell Q_o in Fig. 3 is drawn for $\rho = \sqrt{2}$, which means it is an outer bound of the coverage cell C_o for all combinations of θ , u , and α for which $\sigma(\theta, u) = 2^{\alpha/2}$. For Rayleigh fading and $\alpha = 4$, all Q cells are the same if $u = (4/\theta)^\oplus$. The color map of the figure indicates $\mathbb{P}(\text{SIR}_y > 1)$ for each y in the square region, and the corresponding reliability is $u = 4^\oplus = 0.8$. Hence outside the Q cell, in $V_o \setminus Q_o$, it is guaranteed that $\mathbb{P}(\text{SIR}_y > 1) < 0.8$, which can be confirmed from the colors—none of these locations reach a shade of yellow.

C. Cutting Corners

Here we focus on the stringent regime $\rho > 1$ and show that it is possible to improve on the outer coverage bounds by “cutting the corners” of the Q cells. The reason why the corner points of the Q cells are of special interest is that they are equidistant to the two nearest interferers, while the definition of the Q cell only considers a single nearest interferer. By taking into account the joint interference from the two equidistant interferers, we can obtain a better outer bound of the coverage provided by each transmitter, i.e., a *refined Q cell*.

First we define a version of a Q cell that always considers the two nearest interferers.

Definition 5 (Second-Order Q Cell) For $|\mathcal{X}| > 2$, the second-order Q cell $Q_x^{(2)}$ is the set of location that are covered if the two nearest interferers to each location are considered. Formally,

$$\begin{aligned} Q_x^{(2)} &\triangleq \left\{ y \in \mathbb{R}^2 : \mathbb{P}\left(\frac{h_1 r_1^{-\alpha}}{h_2 r_2^{-\alpha} + h_3 r_3^{-\alpha}} > \theta\right) > u \right\} \\ &= \left\{ y \in \mathbb{R}^2 : \mathbb{P}\left(\frac{1}{H_2 v_2^\alpha} + \frac{1}{H_3 v_3^\alpha} < \frac{1}{\theta}\right) > u \right\}, \end{aligned} \quad (14)$$

where $(r_i)_{i \in [3]}$ are the ordered distances to y and $v_i = r_i/r_1$ and $H_i = h_1/h_i$, as before.

Remarks.

- By construction, $C_x \subseteq Q_x^{(2)} \subseteq Q_x$.
- This definition could be extended easily to account for more than 2 interferers by adding more terms $1/(H_i \rho_i)$ to (14), but the higher-order Q cells are increasingly intractable.

The key purpose of introducing the second-order cell is the definition of a refined version of the Q cell if $\tilde{n}_x > 1$, i.e., if the Q cell has corner points. Before doing so, we need to introduce another variant of the Q cell.

Definition 6 (Interference-cloned Q Cell) Let

$$\rho^*(\theta, u) \triangleq \left(\frac{\theta}{F_{H^*}^{-1}(1-u)} \right)^{1/\alpha}, \quad (15)$$

where $F_{H^*}^{-1}$ is the quantile function of $H^* \triangleq h_1/(h_2 + h_3) = H_2 H_3 / (H_2 + H_3)$. The Q cell $Q_x(\rho^*)$ is the interference-cloned version of the standard Q cell $Q_x(\rho)$.

The interference-cloned Q cell is the Q cell obtained if every interferer is replaced by two interferers at the same distance, each with independent fading or, equivalently, if every interferer has two antennas, sufficiently spaced for the fading to be independent and each transmitting at unit power. The ratio

$$\frac{\rho^*}{\rho} = \left(\frac{F_H^{-1}(1-u)}{F_{H^*}^{-1}(1-u)} \right)^{1/\alpha}$$

is larger than one since $F_{H^*}(x) > F_H(x)$ and thus $F_H^{-1}(1-u) > F_{H^*}^{-1}(1-u)$. For Nakagami- (p, q) fading,

$$F_{H^*}(x) = B_{p,2q} \left(\frac{px}{px+q} \right) = B_{p,2q}((x \frac{p}{q})^\oplus), \quad (16)$$

which is a slight modification of (4). For $p = q = 1$, $F_{H^*}(x) = x(x+2)/(x+1)^2$ and thus $F_{H^*}^{-1}(u) = (1-u)^{-1/2} - 1$. It follows that $(\rho^*)^\alpha = \theta \sqrt{u}^\ominus$ and $\rho^*/\rho = \left(\frac{\sqrt{u}^\ominus}{u^\ominus} \right)^{1/\alpha}$.

The next lemma establishes a connection between $Q^{(2)}(\rho)$ and $Q(\rho^*)$.

Lemma 6 (Corner points of the interference-cloned cell)

The corner points of $Q_x(\rho^*)$ (if any) lie on the boundary $\partial Q_x^{(2)}$ of the second-order cell.

Proof: For the interference-cloned cell,

$$\begin{aligned} \mathbb{P} \left(\frac{h_1 r_1^{-\alpha}}{(h_2 + h_3) r_2^{-\alpha}} > \theta \right) &= \mathbb{P}(H^* v^\alpha > \theta) \\ &= 1 - F_{H^*}(\theta v^{-\alpha}), \end{aligned} \quad (17)$$

as in the proof of Thm. 1 but with H replaced by H^* . Hence a distance ratio of $(\theta/F_{H^*}^{-1}(1-u))^{1/\alpha}$ demarcates the boundary of the interference-cloned cell $Q_x(\rho^*)$.

Now, at the corner points $Q_x(\rho^*)$, there are two actual equidistant interferers (without cloning), so at these locations, cloning has the same effect as taking the two nearest interferers in account, hence they lie on $\partial Q^{(2)}(\rho)$ and have $\mathbb{P}(\text{SIR}_y > \theta) \leq u$, with equality if there are only two interferers. ■

The remaining question is how to “interpolate” between $Q_x(\rho)$ and $Q_x(\rho^*)$ to define a refined Q cell.

Definition 7 (Refined Q Cell) Let $\rho^*(\theta, u)$ be given as in (15) and denote by $\mathbf{q}_k$, $k \in [\tilde{n}_x]$, the corner points of $Q_x(\rho^*)$ and by φ_k and φ'_k the angles of the tangents to the two arcs meeting at $\mathbf{q}_k$. Letting $\bar{\varphi}_k = (\varphi_k + \varphi'_k)/2$, we define

$$\mathcal{H}_x \triangleq \bigcap_{k=1}^{\tilde{n}_x} \mathbf{q}_k^\# + H^- e^{i(\bar{\varphi}_k - \pi/2)}$$

to be the intersection of the half-planes whose boundary is the line through $\mathbf{q}_k$ at angle $\bar{\varphi}_k - \pi/2$, such that $\mathbf{x}^\# \in \mathcal{H}_x$. If




Fig. 5: Example Q cell $Q(\rho)$ with $\rho = \sqrt{3}$ (outermost cell, including the beige parts) and corresponding interference-cloned cell $Q(\rho^*)$, where $\rho^* = 2.07$, (innermost cell, green) and refined Q cell $Q^*(\rho)$ (blue boundary), and the simulated boundary of the second-order cell $Q^{(2)}$ (red line). The corners of $Q(\rho^*)$, marked with red crosses, lie on the boundary of $Q^{(2)}$. The black line segments mark the boundaries of the half-planes touching the corners of $Q(\rho^*)$, effectively “cutting the corners” off of $Q(\rho)$. In this example, the refined Q cell is 8% smaller than the original Q cell.




Fig. 6: Reliability on the boundary of the Q cells in Fig. 5, with the colors matching those of the corresponding cells in Fig. 5. The boundary is represented in polar coordinates relative to the transmitter $\mathbf{x}$. The black line indicates the target reliability $u = 0.9$, corresponding to $\rho = (u^\ominus)^{1/4} = \sqrt{3}$ for Rayleigh fading, $\theta = 1$, and $\alpha = 4$.

$\tilde{n}_x = 0$, $\mathcal{H} = \mathbb{C}$, per the usual convention that the intersection of 0 sets is the universal set.

The refined Q cell is

$$Q_x^*(\rho) \triangleq Q_x(\rho) \cap \mathcal{H}_x^b.$$

As before we set $Q^* \triangleq \bigcup_{\mathbf{x} \in \mathcal{X}} Q_x^*$.

Fig. 5 shows an example of a Q cell $Q(\rho)$, its interference-cloned version $Q(\rho^*)$, and the refined Q cell $Q^*(\rho)$, together with the simulated boundary of the $Q^{(2)}$ cell. It is apparent that the corners of $Q^*(\rho)$ lie on the boundary of $Q^{(2)}$, hence the corners of $Q(\rho)$ can be cut up to those inner corners.

Theorem 2 (Improved outer bound) The refined Q cell Q_x^*

is a tighter outer bound on C_x than Q_x , i.e.,

$$C_x \subseteq Q_x^* \subseteq Q_x \quad \forall x \in \mathcal{X}, \quad (18)$$

and $\mathcal{C} \subseteq \mathcal{Q}^* \subseteq \mathcal{Q}$.

Proof: First inclusion: Since $Q_x^{(2)}$ is convex, the intersection with a tangential half-plane at any point of its boundary leaves it unchanged. So $Q_x^{(2)} \subseteq Q_x^*$, which implies $C_x \subseteq Q_x^*$. Second inclusion: Since the corner points of Q_x^* lie inside Q_x , the intersection with $\mathcal{H}_x^b$ makes Q_x^* smaller than Q_x —unless there are no corner points. ■

If there are more than two interferers, the first inclusion in (18) becomes strict, and if $Q_x(\rho^*)$ has at least two corner points, the second inclusion becomes strict.

Fig. 6 shows the probabilities $\mathbb{P}(\text{SIR} > 1)$ on the boundary of the Q cells in Fig. 5. The boundary is represented in polar coordinates relative to the transmitter o . The Q cells are constructed for a target probability $u = 0.9$. It is apparent that the refined Q cell has a significantly higher reliability at the angles of the corner points of the standard Q cell.

V. COVERAGE ESTIMATION IN INFINITE NETWORKS

In this section, we focus on the stringent QoS regime and assume that $\mathcal{X}$ is countably infinite, in such a way that it forms a *proper point pattern* as per [24, Def. 1]. This means its convex hull spans $\mathbb{R}^2$ and all Voronoi cells are finite. If $\mathcal{X}$ is random, i.e., a point process, these properties are assumed to hold in an almost sure sense.

A. Covered Area Fraction

A key quantity in infinite networks is the covered area fraction.

Definition 8 (Area fraction covered by a set) For $B \subset \mathbb{R}^2$,

$$\eta_B \triangleq \lim_{r \rightarrow \infty} \frac{|B \cap b(o, r)|}{\pi r^2} \in [0, 1]$$

is the area fraction covered by B .

An immediate consequence of Theorems 1 and 2 is $\eta_{\mathcal{Q}} \geq \eta_{\mathcal{Q}^*} \geq \eta_{\mathcal{C}}$. In the next two subsections, we quantify $\eta_{\mathcal{Q}}$ for different types of deployments $\mathcal{X}$.

For the asymptotic scaling of $\eta_{\mathcal{Q}}$, we have the following lemma. We define $\delta \triangleq 2/\alpha \in (0, 1)$.

Lemma 7 (Asymptotic coverage of $\eta_{\mathcal{Q}}$) For all stationary and ergodic point processes, $\eta_{\mathcal{Q}}(\rho) = \Theta(\rho^{-2}) = \Theta(\theta^{-\delta})$, $\rho \rightarrow \infty$.

Proof: As $\rho \rightarrow \infty$, all equal distance ratio circles of $x \in \mathcal{X}$ become concentric at x , and their radii are proportional to the distance from x to the interfering transmitters and decreasing with ρ^{-1} . Hence asymptotically only transmitters with minimum distance, given by $\arg \min_{x' \in \mathcal{X} \setminus \{x\}} \|x - x'\|$ characterize the Q cell. If there is a single interferer at the minimum distance, the Q cell is a disk with area proportional to $\rho^2/(\rho^2 - 1)^2 = \Theta(\rho^{-2})$ by Lemma 2. If there are multiple interferers at the minimum distance, i.e., if $\mathcal{X}$ forms a lattice,

the Q cell is the intersection of a fixed number of disks with equal radius, whose area also scales with $\Theta(\rho^{-2})$. This behavior applies to all Q cells, hence the covered area fraction scales in the same way. Lastly, since $\rho^\alpha \propto \theta$ by Thm. 1, $\rho^{-2} \propto \theta^{-\delta}$. ■

B. Lattices

In regular lattices, all Voronoi cells V_x are congruent, and so are all Q cells Q_x . Hence their area fractions can be calculated by considering just a single cell.

Lemma 8 (Q cell area fraction in lattices) For the square lattice,

$$\eta_{\mathcal{Q}^\square}(\rho) = \frac{4\rho^2 \arctan\left(\frac{\sqrt{2\rho^2-1}-1}{\sqrt{2\rho^2-1}+1}\right) - 2\sqrt{2\rho^2-1} + 2}{(\rho^2 - 1)^2}. \quad (19)$$

For the triangular lattice,

$$\eta_{\mathcal{Q}^\triangle}(\rho) = \frac{4\sqrt{3}\rho^2 \arctan\left(\frac{r-\sqrt{3}}{\sqrt{3}r+1}\right) - \sqrt{3}r + 3}{(\rho^2 - 1)^2}, \quad (20)$$

where $r = \sqrt{4\rho^2 - 1}$. For both square and triangular lattices,

$$\eta_{\mathcal{Q}}(\rho) < \frac{4}{(1 + \rho)^2}. \quad (21)$$

Proof: (19) and (20) are obtained from straightforward calculation of the areas of the disk segments and the inner square and hexagon, respectively. For the square lattice, the tightest square outer bounding the Q cell has side length $2/(1 + \rho)$; for the triangular lattice, the tightest hexagon outer bounding the Q cell has inner radius $a = \sqrt{2}/(3^{1/4}(1 + \rho))$, hence its area $6a^2/\sqrt{3} = 4/(1 + \rho)^2$ is the same. ■

Since the arctan in (19) tends to $\pi/4$, $\eta_{\mathcal{Q}^\square} \sim \pi\rho^{-2}$, $\rho \rightarrow \infty$. For the triangular lattice, the arctan tends to $\pi/6$, so $\eta_{\mathcal{Q}^\triangle} \sim (2\pi/\sqrt{3})\rho^{-2}$, which is 15% larger. The upper bound (21) is asymptotically larger by a factor $2\sqrt{3}/\pi \approx 1.10$ than the Q cells for the triangular lattice, so the bound is at most 10% off.

C. Point Processes

If $\mathcal{X}$ is a stationary and ergodic point process of density λ ,

$$\eta_{\mathcal{Q}} = \lambda \mathbb{E}_o(|Q_o(\rho)|),$$

where $\mathbb{E}_o$ is the expectation with respect to the Palm measure of $\mathcal{X}$ at o [25, Ch. 8]. In words, the area fraction of $\mathcal{Q}$ is the mean area of the typical Q cell divided by the mean area $1/\lambda$ of the typical Voronoi cell.

The following lemma gives the area fraction of $\mathcal{Q}$ for the Poisson point process (PPP).

Lemma 9 (Q cell area fraction for PPP) For the PPP, $\eta_{\mathcal{Q}} = \rho^{-2}$.

Proof: The area fraction only depends on ρ , so the result follows from re-writing the MD for the nearest interferer case in [26, Cor. 3] in terms of ρ . ■




Fig. 7: Area fraction η_Q for the triangular lattice (20), the square lattice (19), the PPP ρ^{-2} (from Lemma 9), and a noisy square lattice with variance 1/16 (simulated). The upper bound is from (21), and the approximation (22) is also shown.

point process	$\rho \downarrow 1$	$\rho \rightarrow \infty$
triangular lattice	$1 - \frac{10}{9}(\rho - 1)$	$\frac{2\pi}{\sqrt{3}}\rho^{-2}$
square lattice	$1 - \frac{4}{3}(\rho - 1)$	$\pi\rho^{-2}$
PPP	$1 - 2(\rho - 1)$	ρ^{-2}

TABLE I: Asymptotics of $\eta_Q(\rho)$ for three types of deployments from Lemmas 8 and 9.

This lemma generalizes [26, Eqn. (8)] to arbitrary fading, including asymmetric cases, and it establishes which combinations of θ , u , and α result in the same Q cell coverage.

Fig. 7 shows the area fractions for the square lattice, the noisy square lattice where a Gaussian perturbation with variance 1/16 is added independently to both coordinates of each lattice point, and the PPP. Table I shows the asymptotic behavior of $\eta_Q(\rho)$ for three point processes.

Since practical deployments often are about in the middle of lattices and PPPs in terms of regularity, a good general approximation is $\eta_Q(\rho) \approx f_Q(\rho)$ for

$$f_Q(\rho) = \frac{3}{2}\rho^{-2} - \frac{1}{2}e^{-3(\rho-1)}. \quad (22)$$

This function has slope $-3/2$ at $\rho = 1$ and decays as $f_Q(\rho) \sim \frac{3}{2}\rho^{-2}$ as $\rho \rightarrow \infty$. It is apparent from Fig. 7 that it matches the area fraction of the noisy lattice quite exactly.

Theorem 3 (Uncovered users) *For any independent stationary and (jointly) ergodic point processes of transmitters and users, the fraction μ of uncovered users is lower bounded by*

$$\mu > 1 - \frac{4}{(1 + \rho)^2}. \quad (23)$$

If the transmitters form a PPP, the fraction of uncovered users is larger than $1 - \rho^{-2}$.

Proof: The triangular lattice is the optimum deployment, for which the Q cell area is upper bounded by (21). The result for the PPP follows from Lemma 9. ■

The theorem establishes that for $\rho = 4/\sqrt{3} - 1 \approx 1.31$, at least 25% of the users are uncovered, and for $\rho = 2\sqrt{2} - 1 \approx 1.83$, at least 50% of the users are uncovered. For the PPP, the two thresholds are $\rho = 2/\sqrt{3} \approx 1.155$ and $\rho = \sqrt{2}$, respectively.

D. Connection to the SIR Meta Distribution

We first state an important fact about the SIR meta distribution (MD) and the covered area fraction η_C .

Fact 1 (Meta distribution and η_C) *If $\mathcal{X}$ is a stationary and ergodic point process,*

$$\eta_C(\rho(\theta, u)) = \bar{F}_{[\text{SIR}]}(\theta, u),$$

where $\bar{F}_{[\text{SIR}]}(\theta, u) \triangleq \mathbb{P}(\mathbb{P}(\text{SIR}_o > \theta \mid \mathcal{X}) > u)$ is the meta distribution of the SIR [14], [15].

Hence, the SIR MD can be either interpreted as η_C or as the fraction of users who achieve an SIR of θ with probability at least u for a stationary and ergodic point process of users.

Consequently, η_Q and η_{Q^*} are upper bounds on the meta distribution. In particular, (23) is a universal upper bound for all point processes and all values of θ and u , as stated in the next corollary.

Corollary 1 (Upper bounds on MD) *For any independent stationary and ergodic point processes of transmitters and users, the MD is bounded as*

$$\bar{F}_{[\text{SIR}]}(\theta, u) < \frac{4}{(1 + \rho)^2} = \frac{4}{\left(1 + \left(\frac{\theta}{F_H^{-1}(1-u)}\right)^{1/\alpha}\right)^2}. \quad (24)$$

For the PPP,

$$\bar{F}_{[\text{SIR}]}(\theta, u) < (F_H^{-1}(1-u))^\delta \theta^{-\delta}. \quad (25)$$

Proof: Follows from Thm. 3 and Fact 1. ■

Remarks.

- For Rayleigh fading, the upper bound (25) simplifies to $(\theta u^\ominus)^{-\delta}$.
- If the MD exceeds the right side of (25), the transmitters form a more regular point process than the PPP.
- If only the nearest interferer is considered in the MD and coverage definition, $\bar{F}_{[\text{SIR}]}(\theta, u) = \eta_C = \eta_Q(\rho(\theta, u))$. Lemma 9 is a consequence of this equality—the result follows from the nearest interferer-only version of the MD.

E. Separability in Poisson Networks

If $\mathcal{X}$ is a stationary PPP, the stringent regime corresponds to the *separable regime* where, for iid fading, $\bar{F}_{[\text{SIR}]}(\theta, u) = g(u)\theta^{-\delta}$ [16, Thm. 2]. “Separable” refers to the fact that the MD can be written as a product of a function of u and $\theta^{-\delta}$. The bound (25) establishes that $g(u) < (F_H^{-1}(1-u))^\delta$. For Nakagami- m fading, the stringent/separable regime is $\{(\theta, u) : u > 1 - B_{m,m}(\theta^\oplus)\}$.

For Rayleigh fading, $\bar{F}_{[\text{SIR}]}(\theta, u) \sim \text{sinc } \delta(\theta u^\ominus)^{-\delta}$, $u \rightarrow 1$, [16, Cor. 5], which suggests that

$$\eta_C \approx \text{sinc } \delta \eta_Q. \quad (26)$$

The asymptotic proportionality of the MD to $\theta^{-\delta}$ as $\theta \rightarrow \infty$ holds for all stationary point processes with Rayleigh fading [16, Thm. 3]. For η_Q , it holds for all types of fading.

Only for the PPP, η_Q is separable. For all point processes, *asymptotic separability* holds in the sense that $\eta_Q(\rho) \sim \tilde{g}(u)\theta^{-\delta}$, $\theta \rightarrow \infty$, where $\tilde{g}(u) = \lim_{\theta \rightarrow \infty} \eta_Q(\rho)\theta^\delta$, which follows from Lemma 7. It is expected that asymptotic separability also holds for the MD.

F. Scaled Q Cells for Coverage Manifold Estimation

Q cells only account for one interferer, while refined Q cells account for two in specific locations. Shrinking the Q cells can approximately incorporate the effect of the remaining interferers. If the MD (*i.e.*, η_C) is available, the Q cells can be shrunk to match the area fraction. Applied to the refined cells, letting $\nu = \eta_C/\eta_{Q^*}$, each cell is scaled as

$$\hat{Q}_x^* \triangleq x + \sqrt{\nu}(Q_x^* - x),$$

and we let $\hat{Q}^* \triangleq \bigcup_{x \in \mathcal{X}} \hat{Q}_x^*$. This way, $\eta_{\hat{Q}^*} = \eta_C$.

To assess the accuracy of Q_x^* and $\hat{Q}_x^*$ as an outer bound and approximation of C_x , respectively, we use the probabilities $\mathbb{P}(\text{SIR} > \theta)$ averaged over the Q cell boundaries. Formally,

$$U_x \triangleq \frac{1}{|\partial Q_x^*|} \int_{y \in \partial Q_x^*} \mathbb{P}(\text{SIR}_y > \theta) dy, \quad (27)$$

$$\hat{U}_x \triangleq \frac{1}{|\partial \hat{Q}_x^*|} \int_{y \in \partial \hat{Q}_x^*} \mathbb{P}(\text{SIR}_y > \theta) dy. \quad (28)$$

For random $\mathcal{X}$, the cdfs of SIR_y are conditioned on $\mathcal{X}$, and U and $\hat{U}$ are random variables. Without subscript, U and $\hat{U}$ refer to the average probabilities on the boundary of the typical cells Q_o^* and $\hat{Q}_o^*$, respectively.

Along the boundary ∂C_x , the reliability is u by definition, and, since $C_x \subset Q_x^*$, $U_x < u$. Hence the proximity of U_x to u and the deviation of $\hat{U}_x$ from u are sensible indicators for the accuracy of the Q cells.

The effect of the scaling can be viewed as an increase in ρ and, in turn, an increase of the reliabilities (or the threshold θ) on the cell boundaries. Using $\eta_{Q^*} \approx \rho^{-2}$, which holds in general, we have $\hat{\rho}^{-2} = \nu \rho^{-2} \approx \eta_C$, so $Q_x^*(\hat{\rho}) \approx \hat{Q}_x^*(\rho)$ for $\hat{\rho} = \rho/\sqrt{\nu}$. Then, from (13),

$$\hat{u} \triangleq 1 - F_H(\theta \hat{\rho}^{-\alpha}) = 1 - F_H(\theta \nu^{1/\delta} \rho^{-\alpha}) \quad (29)$$

is an estimate of the maximum reliability on the boundary of $\hat{Q}_x^*$. For Rayleigh fading, $\hat{u} = 1 - (\theta \nu^{1/\delta} \rho^{-\alpha})^\oplus = (\nu^{-1/\delta} u^\ominus)^\oplus$, where u is the target reliability of the unscaled Q cells.

1) *Square lattice*: Fig. 8 displays the refined Q cell Q_o^* , the scaled version $\hat{Q}_o^*$, and the true coverage cell C_o for the square lattice for $\rho = \sqrt{2}$. Fig. 9 shows the probabilities $\mathbb{P}(\text{SIR} > 1)$ along the boundary $\partial \hat{Q}_o^*$ for targets $u \in \{0.6, 0.7, 0.8, 0.9, 0.95\}$. It is apparent that with increasing u , the probabilities are increasingly well concentrated around u , from 0.6 ± 0.05 to 0.95 ± 0.005 . The averages over the




Fig. 8: Refined Q cell (outer cell, including blue part) and scaled refined Q cell (inner yellow cell) and boundary of coverage cell (red line) for $\rho = \sqrt{2}$ for the square lattice. The area scaling factor from the refined Q cell Q_o^* to the actual cell C_o is $\nu = 0.817$.




Fig. 9: Probability $\mathbb{P}(\text{SIR} > 1)$ along the boundary of the scaled refined cell $\hat{Q}_o^*$ in the square lattice for $\rho = \sigma(1, u)^{1/4} = (u^\ominus)^{1/4}$, $u \in \{0.6, 0.7, 0.9, 0.95\}$, with Rayleigh fading. The values of u are indicated using dashed lines.

boundary are 0.606, 0.703, 0.801, 0.949, which are very close to the target values of u .

2) *Poisson point process*: Fig. 10 shows refined Q cells Q_x^* , scaled refined Q cells $\hat{Q}_x^*$, and coverage cells C_x for a realization of a PPP of intensity 1 for Rayleigh fading.

From $\hat{Q}_x^* \approx C_x$ follows that $\mathbb{E}(\hat{U}) \approx u$, and, from (29), $\mathbb{E}(\hat{U}) \approx (\nu^{-1/\delta} (\mathbb{E}(U))^\ominus)^\oplus$, which can be inverted to yield

$$\mathbb{E}(U) \approx (\nu^{1/\delta} u^\ominus)^\oplus. \quad (30)$$

Fig. 11 shows the empirical cdfs $\mathbb{P}(U \geq x)$ and $\mathbb{P}(\hat{U} \geq x)$ for $\theta = 1$, estimated from 1000 Q cells with target reliability $u = 0.8$ ($\rho = \sqrt{2}$) in a realization of a PPP. The means are $\mathbb{E}(\hat{U}) = 0.821$, which is close to u as stipulated, and $\mathbb{E}(U) = 0.654$, which is close to $(\nu^{1/\delta} u^\ominus)^\oplus = 0.646$, in agreement with (30). The empirical variances are $\text{var}(U) \approx 0.0149$ and $\text{var}(\hat{U}) \approx 0.0069$, so the values on $\partial \hat{Q}_o^*$ are well concentrated in the interval $[(\nu^{1/\delta} u^\ominus)^\oplus, (\nu^{-1/\delta} u^\ominus)^\oplus]$, which is the interval between the dashed and dash-dotted lines in Fig. 11.

3) *Area scaling without MD*: If the MD is not available, the scaling factor ν can be set heuristically as

$$\nu(\delta, g) = (\text{sinc } \delta)^{1-g(1-\delta)}, \quad (31)$$




Fig. 10: Refined Q cells (outer cells, including blue part) and scaled refined Q cells (inner yellow cells) for $\rho = \sqrt{2}$ and boundaries of the coverage cells C (red lines) for $\theta = 1$, $u = 0.8$, $\delta = 1/2$, and Rayleigh fading for a realization of a PPP. The area scaling factor from the refined to the actual cell is $\nu = 0.676$.




Fig. 11: Empirical cdfs of $\mathbb{P}(\text{SIR} > 1 \mid \mathcal{X})$ for the PPP on the boundaries ∂Q_o^* and $\partial \hat{Q}_o^*$. The vertical lines indicate the target $x = u = 0.8$ (solid line) and $(\nu^{1/\delta} u^\ominus)^\oplus = 0.646$ (dashed line). The dashed-dotted line at $x = 0.898$ is the estimate (29) of the upper bound of the reliabilities. $\mathbb{E}(\hat{U}) = 0.821$ and $\mathbb{E}(U) = 0.654$.

where g quantifies the regularity of the point process. For a lattice, $g = 1$, while for the PPP, $g = 0.3$. For perturbed lattices with variance v , $g = \max\{0.3, 1 - 3.5v\}$.

The rationale behind (31) is the following:

- Since both η_C and η_{Q^*} scale similarly with ρ , ρ does not need to be included as an explicit parameter in an estimate of ν .
- If $\delta \rightarrow 0$, $Q_x^* \rightarrow C_x$ for all x , i.e., $\nu(0) = 1$. If $\delta \rightarrow 1$, the interference becomes overwhelming, and $|C_x| \rightarrow 0$, i.e., $\nu(1) = 0$. For intermediate δ , $\nu(\delta)$ is monotonically




Fig. 12: Area scaling factor $\nu(\delta) = \eta_C/\eta_{Q^*}$ for different deployments. The markers are obtained from simulations, the solid lines from (31).

decreasing.

- The base sinc δ is taken from (26), which is raised to an exponent that depends on the regularity g and δ and takes into account that (26) is the approximate scaling factor of the standard Q cells, not of the refined ones. For higher regularity, the exponent is smaller since the refined Q cell is more accurate.
- For $\delta \ll 1$ ($Q^* \approx C$ in all cases) and $1 - \delta \ll 1$ ($|C| \approx 0$ in all cases), the regularity of the point process has little influence. To reflect this, g is multiplied by $1 - \delta$.

Fig. 12 compares this simple heuristic with simulation results for the square lattice, perturbed lattices, and the PPP.

In the example of $\delta = 1/2$ used in Fig. 10, $(\text{sinc}(1/2))^{0.85} \approx 0.68$, which is close to the true scaling of 0.676 (see the red line and the red $\times$ in Fig. 12 at $\delta = 1/2$).

We summarize the steps to obtain the scaled refined Q cells in the following recipe.

Recipe 1 (Scaled refined Q cells)

- (1) For given θ and u , calculate the stringency $\sigma(\theta, u)$ and set $\rho = \sigma^{1/\alpha}$.
- (2) Determine the refined Q cells $Q_x^*(\rho)$ for each $x \in \mathcal{X}$.
- (3a) If the MD η_C is available, set $\nu = \eta_C/\eta_{Q^*}$ and scale each Q_x^* to $\hat{Q}_x^*$ so that $\eta_{\hat{Q}^*} = \eta_C$.
- (3b) If the MD is not available, set ν according to (31) and scale each Q_x^* to $\hat{Q}_x^*$ so that $\eta_{\hat{Q}^*} \approx \eta_C$.

VI. EXTENSIONS

Here we briefly discuss two extensions of the network model—to the case of different transmit power and to the case of more general path loss laws.

A. Different Transmit Powers

With different transmit powers, the natural association model is that a location y is served by the strongest transmitter (including transmit power), instead of always by the nearest transmitter. Similarly, the strongest interferer may not be the second-nearest transmitter.

Let P_i be the transmit power of transmitter x_i and, for a location $y \in \mathbb{R}^2$, let $\tilde{r}_i$ be the i -th smallest element of $\|x_i - y\|/P_i^{1/\alpha}$ and set $r_i \triangleq \tilde{r}_i P_i^{1/\alpha}$. This way, $\tilde{r}_i^{-\alpha}$ is the power received from the i -th strongest transmitter and r_i is the distance to the i -th strongest transmitter. The SIR is the same as (3) but with r_i replaced by $\tilde{r}_i$. The reliability condition (13) needs to be augmented by the transmit power ratio of the second-strongest and strongest transmitters, denoted by γ :

$$\mathbb{P}\left(\frac{h_1 r_1^{-\alpha}}{\gamma h_2 r_2^{-\alpha}} > \theta\right) = 1 - F_H(\theta \gamma v^{-\alpha})$$

Hence $\gamma > 1$ indicates that the strongest interferer has a higher transmit power than the serving transmitter, and its effect is a scaling of the threshold parameter θ , which decreases the disks radii. The resulting definition of the Q cell is

$$Q_x \triangleq \left\{y \in \mathbb{R}^2 : \frac{r_2}{r_1} > \sigma(\theta \gamma, u)^{1/\alpha}\right\}.$$

Due to the presence of the power ratio γ , the Q cells can no longer be expressed as a function of a global parameter ρ .

B. General Path Loss Laws

With a general (monotonically decreasing) path loss law $\ell(r)$, the condition for coverage (see the proof of Thm. 1) is

$$\frac{\ell(r_2)}{\ell(r_1)} < \frac{F_H^{-1}(1-u)}{\theta},$$

which in general cannot be expressed as a condition on the distance ratio r_2/r_1 . Consequently, the Q cell is not delimited by arcs but more complicated shapes, which depend on the concrete form of ℓ .

For the bounded law $\ell(r) = (c + r^\alpha)^{-1}$ with $c > 0$ and $r_2 > r_1$ (which is necessary for stringencies above 1),

$$\frac{\ell(r_2)}{\ell(r_1)} = \frac{c + r_1^\alpha}{c + r_2^\alpha} < \frac{r_1^\alpha}{r_2^\alpha},$$

which shows that the resulting Q cell is a subset of that for $c = 0$. The larger c , the smaller the Q cell, since the derivative of the ratio $\ell(r_2)/\ell(r_1)$ with respect to c is negative for stringencies larger than 1.

VII. CONCLUSIONS

Q cell analysis (QCA) is a powerful method to bound and approximate the coverage manifold in wireless networks. Standard Q cells constitute simple yet fairly accurate outer bounds on the coverage region of each transmitter. They depend on the QoS parameters θ and u only through the scalar stringency σ and can be calculated very efficiently as the intersection of a modest number of disks whose boundaries are formed by equal distance ratio circles with ratio $\rho = \sigma^{1/\alpha}$. As a generalization of Voronoi cells, they are defined as the intersection of Möbius transformed half-planes instead of just the intersection of half-planes. In the stringent regime, the complement of their union is guaranteed to be uncovered. Cutting corner points to account for equidistant interferers results in refined Q cells that tightly outer bound the coverage regions.

The stringency has a simple analytical expression for general Nakagami- (p, q) fading, which allows for different fading types in the serving and interfering links. It establishes that for all combinations of θ , u , and α that have the same distance ratio $\rho = \sigma^{1/\alpha}$, the coverage performance will be similar. Accordingly, it is unnecessary to explore many combinations of the three parameters—the coverage outer bound only depends on the single scalar parameter ρ .

Applied to infinite networks, Q cell coverage is much more tractable than the actual coverage (or SIR meta distribution) but closely related. It allows for the derivation of simple universal bounds. For instance, in any infinite network with $\rho = 2\sqrt{2} - 1$, at least 50% of the users are uncovered. In Poisson networks, more than 50% are uncovered at $\rho = \sqrt{2}$.

Compared to simulations or ML-based approaches, QCA is orders of magnitude faster, and the powerful bounding property of the Q cells (any location outside Q is guaranteed to be uncovered) is hard to achieve in simulations or ML. In addition, simple analytical results such as that in Thm. 3 could not be obtained by other techniques.

We expect that QCA will be useful to quickly explore the effect of advanced transmission schemes such as coordinated multipoint (CoMP) or non-orthogonal multiple access (NOMA), and the effects of resource allocation, scheduling and load balancing—and to facilitate their implementation. The basic Q cells (where one transmitter is the serving one and all others interfere) can be used to quickly determine which users reside outside Q and thus suffer from insufficient QoS (sometimes referred to as “edge users”). To meet QoS constraints, these users require additional resources, such as additional bandwidth (carrier aggregation or channel bonding), reduced interference (base station silencing), or increased signal power (CoMP).

REFERENCES

- [1] F. Baccelli and B. Błaszczyszyn, *Stochastic Geometry and Wireless Networks*. Foundations and Trends in Networking, NOW, 2009.
- [2] B. Błaszczyszyn, M. Haenggi, P. Keeler, and S. Mukherjee, *Stochastic Geometry Analysis of Cellular Networks*. Cambridge University Press, 2018.
- [3] A. D. Wyner, “Shannon-Theoretic Approach to a Gaussian Cellular Multiple-Access Channel,” *IEEE Transactions on Information Theory*, vol. 40, pp. 1713–1727, Nov. 1994.
- [4] S. Weber, J. G. Andrews, and N. Jindal, “An Overview of the Transmission Capacity of Wireless Networks,” *IEEE Transactions on Communications*, vol. 58, pp. 3593–3604, Dec. 2010.
- [5] S. Kusaladharma and C. Tellambura, “On Approximating the Cognitive Radio Aggregate Interference,” *IEEE Wireless Communications Letters*, vol. 2, pp. 58–61, Feb. 2013.
- [6] A. Shafie, N. Yang, S. Durrani, X. Zhou, C. Han, and M. Juntti, “Coverage Analysis for 3D Terahertz Communication Systems,” *IEEE Journal on Selected Areas in Communications*, vol. 39, pp. 1817–1832, June 2021.
- [7] V. Mordachev and S. Loyka, “On Node Density-Outage Probability Tradeoff in Wireless Networks,” *IEEE Journal on Selected Areas of Telecommunications*, vol. 27, pp. 1120–1131, July 2009.
- [8] J. Lee, N. Lee, and F. Baccelli, “Scaling Laws for Ergodic Spectral Efficiency in MIMO Poisson Networks,” *IEEE Transactions on Information Theory*, vol. 64, pp. 2791–2804, Apr. 2018.
- [9] Y. Qin, M. A. Kishk, and M.-S. Alouini, “A Dominant Interferer Plus Mean Field-Based Approximation for SINR Meta Distribution in Wireless Networks,” *IEEE Transactions on Communications*, vol. 71, pp. 3663–3678, June 2023.

- [10] V. V. Chetlur and H. S. Dhillon, "Downlink Coverage Analysis for a Finite 3-D Wireless Network of Unmanned Aerial Vehicles," *IEEE Transactions on Communications*, vol. 65, pp. 4543–4558, Oct. 2017.
- [11] M. Haenggi, "Simalysis: Symbiosis of simulation and analysis." Blog on Stochastic Geometry, <https://stogblog.net/2020/10/15/simalysis-symbiosis-of-simulation-and-analysis/>, Oct. 2020.
- [12] W. U. Mondal, P. Mankar, G. Das, V. Aggarwal, and S. V. Ukkusuri, "Deep Learning-Based Coverage and Rate Manifold Estimation in Cellular Networks," *IEEE Transactions on Cognitive Communications and Networking*, vol. 8, pp. 1706–1715, Aug. 2022.
- [13] W. U. Mondal, V. Goyal, S. V. Ukkusuri, G. Das, D. Wang, M.-S. Alouini, and V. Aggarwal, "Near-Perfect Coverage Manifold Estimation in Cellular Networks with Conditional GAN," *IEEE Networking Letters*, vol. 6, pp. 97–100, Feb. 2024.
- [14] M. Haenggi, "The Meta Distribution of the SIR in Poisson Bipolar and Cellular Networks," *IEEE Transactions on Wireless Communications*, vol. 15, pp. 2577–2589, Apr. 2016.
- [15] M. Haenggi, "Meta Distributions—Part 1: Definition and Examples," *IEEE Communications Letters*, vol. 25, pp. 2089–2093, July 2021.
- [16] K. Feng and M. Haenggi, "Separability, Asymptotics, and Applications of the SIR Meta Distribution in Cellular Networks," *IEEE Transactions on Wireless Communications*, vol. 19, pp. 4806–4816, July 2020.
- [17] M. Haenggi, "SIR Analysis via Signal Fractions," *IEEE Communications Letters*, vol. 24, pp. 1358–1362, July 2020.
- [18] C. S. Ogilvy, *Excursions in Geometry*. Oxford University Press, 1969.
- [19] D. de la Vega, S. López, J. M. Matías, U. Gil, I. Peña, M. M. Vélez, J. L. Ordiales, and P. Angueira, "Generalization of the Lee Method for the Analysis of the Signal Variability," *IEEE Transactions on Vehicular Technology*, vol. 58, pp. 506–516, Feb. 2009.
- [20] G. George, R. K. Mungara, A. Lozano, and M. Haenggi, "Ergodic Spectral Efficiency in MIMO Cellular Networks," *IEEE Transactions on Wireless Communications*, vol. 16, pp. 2835–2849, May 2017.
- [21] M. Abramowitz and I. A. Stegun, *Handbook of Mathematical Functions*. National Bureau of Standards, 1964. 2nd Ed.
- [22] D.-T. Lee, "On k -Nearest Neighbor Voronoi Diagrams in the Plane," *IEEE Transactions on Computers*, vol. C-31, pp. 478–487, June 1982.
- [23] F. Baccelli and A. Giovanidis, "A Stochastic Geometry Framework for Analyzing Pairwise-Cooperative Cellular Networks," *IEEE Transactions on Wireless Communications*, vol. 14, pp. 794–808, Feb. 2015.
- [24] M. Haenggi, "The Stable Point Process and its Applications to Wireless Networks," *IEEE Wireless Communications Letters*, vol. 13, pp. 1394–1398, May 2024.
- [25] M. Haenggi, *Stochastic Geometry for Wireless Networks*. Cambridge University Press, 2012.
- [26] M. Haenggi, "Meta Distributions—Part 2: Properties and Interpretations," *IEEE Communications Letters*, vol. 25, pp. 2094–2098, July 2021.