Cambridge University Press
978-1-107-05094-5 - Analytical Modeling of Heterogeneous Cellular Networks:
Geometry, Coverage, and Capacity
Sayandev Mukherjee
Excerpt
More information

1 Introduction

The rapid rate of increase in data traffic means that future wireless networks will have to support a large number of users with high data rates. A promising way to achieve this is by spectrum reuse through the deployment of cells with small range, such that the same time-frequency resources may be reused simultaneously in multiple cells. At the same time, the traditional *coverage* requirement for wireless users (supporting a modest rate at cell-edge users) is most economically met with cells having large range, i.e. the traditional *macrocellular* architecture. Thus the wireless cellular networks of the future are likely to be *heterogeneous*, i.e. have one or more *tiers* of small cells *overlaid* on the macrocellular tier.

Let us look at network design from the point of view of a service provider considering a deployment of a network in a certain region. Throughout this book, we only consider the *downlink*, i.e. the links from the BSs to the user terminals. The principal metrics we shall focus on are coverage and capacity.

(1) *Coverage* Intuitively, a user (or, more precisely, a user location) is *covered* if the communication link from the BS serving that user is sufficiently "good" that the user terminal can correctly receive both the control signaling and the data traffic at some minimum rate from the BS. Here, "good" means that (a) the received signal from the BS is "strong," i.e. the received signal power from the BS exceeds some threshold, and also that (b) the received *signal to interference plus noise ratio* (SINR) at the user exceeds some minimum value. In modern cellular networks such as LTE (3GPP 2010, Sec. 5.2.3.2), the threshold on received signal power is low enough to be comfortably exceeded by any modern receiver within a fairly large distance from the transmitting BS. Thus, for the purposes of analysis, the signal power requirement can be dropped without loss of generality, and the SINR is the key indicator of coverage.

(2) *Capacity* For the moment, the meaning of "capacity" is deliberately ambiguous, denoting both the maximum attainable data rate to a served user, and also the number of users that can be served by a given BS with some minimum data rate to each. We give more precise definitions of both coverage and capacity later, but for now it is sufficient to understand that the known relationship between achievable data rate on a link and the SINR on that link means that capacity in the sense used here is also dependent on SINR at the user terminals.

Since coverage and capacity are both related to the SINR at the user's receiver, we shall focus on the statistical properties of this SINR in the book. When we talk of the service provider *designing* a network to satisfy a certain coverage and/or capacity requirement, we mean the choice of deployment parameters for such a network, including:

(1) the number of tiers of the network;
(2) the densities of the BSs in the tiers; and
(3) the transmit powers of the BSs in the tiers.

Note that the service provider seeks to satisfy the coverage and/or capacity requirement with a deployment that is compatible with the service provider's targets for expected revenue, capital expenditure (CapEx), and operating expenditure (OpEx). In other words, there is a complicated *utility function* that depends on economic variables such as the pricing model for service, CapEx, OpEx, and expected revenue, in addition to the coverage criterion.

In this book, we focus only on the relationship between the deployment parameters of the network and the coverage and capacity. As we shall show, there are many sets of deployment parameters that are equivalent from the point of view of the SINR distribution. Thus, once an operating SINR distribution has been chosen, the service provider can choose one set of deployment parameters from the corresponding equivalence class based on its utility function, which takes into account the economic considerations.

When thermal noise is negligible, the SINR becomes the *signal to interference ratio*, or SIR. Before we can study the distribution of the SINR or SIR at a user, we need to define the wireless-channel model we shall be using in this book.

1.1 Wireless-channel model

There are two components to the so-called *link loss*, i.e. the difference between transmitted and received powers on a given wireless link: the distance-dependent *path loss* (which has no random component, and is completely known if the distance between transmitter and receiver is known), and the *fading* (which is random). We now discuss the models for these two components.

1.1.1 Path-loss model

We adopt the popular *slope-intercept model*, which says that the *path loss* (i.e. the ratio of the power transmitted over the link by the transmitter to the power received at the receiver) on the link between a BS in a given tier and an arbitrarily located user is given by

$$10\log_{10}\frac{\text{power transmitted (by transmitter)}}{\text{power received (at receiver)}} = 10\alpha\log_{10}(\text{distance in meters}) - 10\log_{10}K, \tag{1.1}$$

where the *slope* $\alpha > 2$ is the *path-loss exponent* and the *intercept* K is a constant accounting for the differences in height of the transmit and receive antennas, etc. It is important to bear in mind that α and K both depend upon the frequency band of operation.

Note that the path-loss model only applies to the attenuation over the air, i.e. after the transmission has left the transmit antenna, and just before entering the receive antenna. The final expression for received power can account for factors such as antenna gains as decrements to the above path loss.

1.1.2 Fading model

We show in the subsequent sections that the calculation of both the joint and marginal *complementary CDFs* (CCDFs) of the downlink SINR at an arbitrarily located user in the network are examples of a class of problems that are made tractable by pairing the PPP model for BS locations with a particular class of *independent identically distributed* (i.i.d.) fading processes on all links, namely a mixture of Nakagami-m (Mallik, 2010) fading processes with positive integer-valued m (see Appendix A for details on this and other distributions used in the book). The convention in wireless communications is to refer to the fading process by the name of the distribution (e.g. Rayleigh, Nakagami, etc.) of the *magnitude* of the equivalent complex baseband signal. The *fade attenuation* on a link is the square of this magnitude. For the Nakagami mixture fading process, the distributions of the fade attenuations on the links are given by a mixture of *Erlang* random variables, i.e. a mixture of gamma random variables with positive integer-valued shape parameters (see Appendix A). Moreover, the *probability density function* (PDF) of an arbitrary positive-valued continuous random variable can be approximated uniformly to any desired accuracy by a mixture of Erlang PDFs (see Lemma 2.6). The distribution of received power on a wireless fading channel with gamma-distribution fast fading and lognormal shadowing can be modeled by a single gamma random variable (of non-integer order), based on matching the first and second moments of the two distributions, as shown, for example, in Heath, Kountouris & Bai (2013). A similar approach (Atapattu, Tellambura & Jiang, 2011) shows how the distribution of received power in many wireless fading channel with lognormal shadowing can be approximated using mixtures of gamma PDFs (with non-integer order). Finally, Almhana *et al.* (2006, eqns. (16)–(17)) shows how to approximate such mixtures of general gamma PDFs by mixtures of Erlang PDFs.

Note that although it is customary in the literature to assume lognormal shadow fading on wireless links, we do not work with the lognormal distribution in our analysis. The reason is that the lognormal distribution is just not very tractable, which explains the plethora of approximations in the literature. Further, it has been observed, at least in a single-tier network, that lognormal shadowing has almost no effect on the (reciprocal of the) SIR (Błaszczyszyn, Karray & Klepper, 2010) when the serving BS is that received most strongly at the user.

The simplest example of a mixture of Erlang PDFs is the exponential PDF, corresponding to the Rayleigh fading process. We shall show that the joint or marginal

SINR CCDF for an arbitrary mixture of Erlang PDFs can be derived from the result for Rayleigh fading. For this reason, we begin by assuming *independent identically distributed* (i.i.d.) Rayleigh fading on all links from all BSs to any user location. Recall that, here, "Rayleigh" refers to the distribution of the magnitude of the equivalent complex baseband signal at the receiver. The fade attenuation (i.e. the power "gain" on the link to the receiver) is distributed as the square of a Rayleigh random variable, i.e. it is an exponential random variable (see Appendix A).

1.2 Distribution of the SINR at an arbitrary user

In this section, we shall make our first attempt at calculating, analytically, the distribution of the SINR at an arbitrary user location in the network. We begin with a snapshot of the wireless network *at a particular moment in time*. Further, let us assume there is just one tier of BSs, as in the currently deployed macrocellular networks, and that each BS transmits with fixed power, assumed to be the same at all BSs. Consider an arbitrarily located user, and suppose it is located at a distance R_0 from its serving BS (where the serving BS was selected according to some criterion – we discuss such criteria later). BSs in the network other than the serving BS are then *interferers*. Suppose there are M such interfering BSs, labeled $1, \dots, M$, with their distances from the user being $R_1, \dots, R_M$, respectively. Further, suppose that all BSs transmit with the same power P^{tx}, and that the fades on all links between the $M+1$ BSs (the serving BS, labeled 0, and the M interfering BSs $1, \dots, M$) and the user are i.i.d. Rayleigh. Then the received power at the user from BS k ($k = 0, \dots, M$) is given by

$$Y_k = \frac{PH_k}{R_k^\alpha}, \; k = 0, \dots, M, \quad \{H_k\}_{k=0}^M \text{ i.i.d. Exp}(1), \quad P = KP^{\text{tx}}, \tag{1.2}$$

where the exponential distribution for H_k, $k = 0, \dots, M$, follows from the i.i.d. Rayleigh fading assumption, and K and α are the intercept and slope, respectively, of the slope-intercept path-loss model (1.1) describing links from the BSs to the user. Let us further simplify by assuming that thermal noise power at the user's receiver is negligible compared to interference power. Thus the SINR (actually SIR) at the user is given by

$$\Gamma = \frac{Y_0}{\sum_{k=1}^M Y_k} = \frac{H_0/R_0^\alpha}{\sum_{k=1}^M H_k/R_k^\alpha}. \tag{1.3}$$

Conditioned on $R_k = r_k$, $k = 0, \dots, M$, the CCDF of the SIR at the user can be written as follows (see Problem 1.1):

$$\mathbb{P}\left\{\Gamma > \gamma \mid R_0 = r_0, \dots, R_M = r_M\right\} = \prod_{k=1}^M \mathbb{E}\left[\exp\left(-\gamma r_0^\alpha \frac{H_k}{r_k^\alpha}\right)\right]. \tag{1.4}$$

It is important to recognize that the conditional CCDF of the SINR is a *Laplace transform* of a sum of independent random variables:

$$\mathbb{P}\left\{\Gamma > \gamma \;\middle|\; R_0 = r_0, \ldots, R_M = r_M\right\} = \mathbb{E}\, \mathrm{e}^{-sW}\Big|_{s=\gamma r_0^\alpha}, \quad W = \sum_{k=1}^{M} \frac{H_k}{r_k^\alpha}. \tag{1.5}$$

Remark 1.1 The above calculation can be extended (see Problem 1.2) to the case where the common distribution of $H_0, \ldots, H_M$ is not exponential but Erlang (Torrieri & Valenti, 2012), i.e. the fading process is not Rayleigh but Nakagami-m with integer-valued m (Mallik, 2010). As we shall show later (see Lemma 2.6), this is sufficient to compute the conditional CCDF of the SINR to any desired degree of accuracy for any *arbitrary* fading process.

Remark 1.2 If M is infinite, the random variable W defined in (1.5) exists if and only if $\sum_{k=1}^{\infty} 1/r_k^\alpha < \infty$ (Haenggi & Ganti, 2009, p. 8).

Unfortunately, (1.5) is not quite the analytic *coup* it may seem. The reason is that (1.4) is the CCDF of the SIR at the user *conditioned* on the location of the user relative to the BSs $0, 1, \ldots, M$. In other words, (1.4) is the CCDF of the SIR at *one specific* user location. This is not very useful to a network operator, who wants to know either (a) the distribution of the SIR at an *arbitrary* user location, which is the *unconditional* CCDF of the SIR, or, equivalently, the expectation of (1.4) with respect to the joint distribution of $(R_0, R_1, \ldots, R_M)$; or (b) the CCDF of the SIR at an arbitrary user location at a distance of $R_0 = r_0$ from its serving BS, which is the expectation of (1.4) with respect to the joint distribution of $(R_1, \ldots, R_M)$ given $R_0 = r_0$. Now, for a *fixed* deployment of BSs in a network and an arbitrary user location in that network, analytic expressions for the joint distribution of $(R_0, R_1, \ldots, R_M)$ and the conditional joint distribution of $(R_1, \ldots, R_M)$ given $R_0 = r_0$ are unknown. Thus the expectation of (1.4) with respect to either of these distributions cannot be computed analytically.

Problems

1.1 Derive (1.4). *Hint*: Use (1.3) and the fact that $H_0 \sim \mathrm{Exp}(1)$.

1.2 Show that (1.4) can be generalized to the case where $H_0, H_1, \ldots, H_M$ are i.i.d. Erlang(m, c) with shape parameter $m \in \{1, 2, \ldots\}$ and rate parameter $c > 0$:

$$f_{H_k}(x) = \frac{c^m}{(m-1)!} x^{m-1} \exp(-cx), \quad x \geq 0, \quad k = 0, 1, \ldots, M.$$

Hint: Observe that

$$f_{H_k}(x) = \frac{(-1)^{m-1}}{(m-1)!} c^m \frac{\partial^{m-1}}{\partial c^{m-1}} \exp(-cx), \quad x \geq 0,$$

i.e. the Erlang PDF is the $(m-1)$th derivative of the PDF of the Exp$(1/c)$ distribution with respect to the parameter c. Thus the Laplace transform of the Erlang PDF is given by the $(m-1)$th derivative with respect to c of the Laplace transform of the Exp$(1/c)$ PDF.

1.3 Why SINR distributions are usually found via simulation

For a given deployment of BSs, although the expectation of (1.4) with respect to either the joint distribution of $(R_0, R_1, \ldots, R_M)$ or the conditional joint distribution of $(R_1, \ldots, R_M)$ given $R_0 = r_0$ cannot be computed analytically, it can always be determined numerically via Monte Carlo methods. However, this is equivalent to *simulating* the network deployment. Further, a simulation of the network can be done on a much more detailed, much less simplified model of the wireless channel and its impairments while still keeping the computational complexity within reasonable bounds. This is the reason why simulation is so popular among industry practitioners – in addition to being just about the sole method to obtain these SIR distributions, it also allows one to probe the sensitivity of this SIR distribution to any parameter of the deployment or wireless-channel model in any desired depth of detail.

Moreover, for a single-tier macrocellular network with "regular" placement of BSs in the centers of hexagonal "cells" on the familiar hexagonal lattice, and for the usual values of path-loss exponents for macrocellular wireless links, it has been observed that the total downlink interference power at an arbitrary user location is essentially that due to the two "rings" of BSs (6 BSs in the first ring, 12 BSs in the second ring) around the cell containing the user. This permits the simulation to be efficient with regard to memory requirements and run time by restricting the deployment region to the "19-cell wraparound region" (Chen *et al.*, 2011), where *wraparound* means that the upper and lower, and the right and left, boundaries of the region are assumed contiguous, so there are no edge effects in the simulation, and statistics on the SINR may be drawn from all locations in the region without introducing bias.

It is important, however, to recall that the 19-cell wraparound model for simulation was proposed for the study of macrocellular networks. Let us examine the assumptions behind this model.

(1) *Regular placement of BSs* While this assumption was never applicable to any real-world deployment, it could at least be claimed to be more or less accurate for macrocells, because of careful planning of cell sites by operators. However, it is increasingly inapplicable as the number of cells increases and their size decreases, as is the case for future networks. The reason for this is that careful planning of cell sites is often impossible if the number of sites is large, and the deployment map then takes on a more "random" appearance.

(2) *Interference limited to two rings of BSs* With the placement of BSs no longer regular, the interference is effectively that due to all BSs within some distance from the user, instead of those BSs located in certain "rings." However, the value of this effective range depends upon the parameters of the path-loss model and the transmit powers of the BSs, which are different for the smaller cell sizes of future networks. In particular, for a macro-femto overlay network, these ranges are different for the two tiers, and a suitably sized region would need to be considered such that edge effects due to both tiers can be eliminated. This is likely to enlarge the deployment region to be simulated, increase memory requirements, and reduce the speed of the simulator.

To summarize, in future heterogeneous networks with smaller cell sizes, where the BS locations are more irregular, the simulations need to be done over larger deployment regions and will become more time-consuming.

A single-tier network has only two deployment parameters: the transmit powers of the BSs and their density (number of BSs per unit area). Thus the total number of scenarios required to be simulated for a single-tier network is not large, especially since the choices for transmit power and density of the BSs required for system operation are tightly constrained. However, each additional tier in the network multiplies the number of possibilities for the overall combination of transmit powers and densities of the various tiers. In other words, if our goal is exhaustive simulation of all feasible operating scenarios for a multi-tier heterogeneous network, the number of scenarios to be simulated rises exponentially with the number of tiers. If one of the tiers is that of user-owned (not operator-owned) femtocells, the number of scenarios is further augmented by the fact that some femtocells may be switched on and off by the user depending on the need for coverage at that user's premises. In other words, the layout of BSs may even become dynamic. Clearly, an exhaustive simulation-based investigation of all possible operating scenarios for a multi-tier heterogeneous network is challenging at best, and infeasible at worst.

1.4 The role of analytic modeling

The analytic-modeling-based investigation of deployment scenarios has two phases. In the first phase, we use probabilistic models for the locations of the BSs to determine analytic expressions for the CCDF of the SINR in the deployment region. In other words, the use of a stochastic model (Poisson point process, or PPP) for the locations of the BSs allows us to write an analytic expression for the expectation of (1.4) with respect to either the joint distribution of $(R_0, R_1, \ldots, R_M)$ or the conditional joint distribution of $(R_1, \ldots, R_M)$ given $R_0 = r_0$. Further, these results can be extended to arbitrary fading distributions and arbitrary numbers of tiers of BSs.

As we shall see, this has the benefit of providing insights into the *combinations* of deployment parameters that affect the CCDF of the SINR, and therefore the different sets of deployment parameters that are *equivalent* in that they yield the same CCDF of the SINR. This analytic phase allows us to sift through the large space of combinations of deployment parameters to settle quickly on certain equivalence classes of deployment parameters, each class corresponding to some desired CCDF of the SINR. The service provider may then choose a set of deployment parameters from one of these equivalence classes based on its economic utility function.

In the next phase of the network design, the shortlist of deployment scenarios (as defined by the deployment parameters) chosen in the first phase may be investigated in depth via simulation. This effectively uses the power of detailed simulation, incorporating all relevant aspects whose behavior and impact on performance is to be investigated, for a few selected deployment scenarios.

2 Structure of the SINR calculation problem

We begin with the simple SINR calculation problem of Chapter 1 and generalize it to a deployment with more than one tier of BSs. We then examine the features of the problem that permit us to evaluate the CCDF of the SIR in terms of Laplace transforms of fading coefficients on the links to the user location from the BSs in the tiers. Then we abstract the problem formulation slightly in order to define a general probability calculation of a vector of random variables, which we call the *canonical problem*. We derive general results and conditions under which the canonical probability may be expressed in terms of Laplace transforms of certain random variables. This is one half of the mathematical core of the book. The other half is the study of stochastic models that yield tractable analytic expressions for the Laplace transforms in the canonical probability calculation, and we discuss that topic in Chapter 3.

2.1 Statement of the SINR calculation problem

Let us return to the problem of computing the distribution of the SINR at an arbitrarily located user somewhere in a network with one or more tiers of BSs. We begin by defining the *candidate serving BSs*, the criteria for their selection, the criterion for choosing the BS that will serve the user, and some notation to represent received power from the candidate serving BSs and from all interferers.

2.1.1 Candidate serving BSs and the serving BS

Consider a snapshot of the wireless network at a particular moment in time. This corresponds to the model for a single resource element (a single subcarrier over one transmission interval) in the long-term evolution (LTE) standard, or to a single transmission interval for a frequency non-selective channel in the HSPA standard. Consider a user located anywhere in this network. Label the tiers of the network $1, \ldots, n_{\text{tier}}$, and assume that the user is only allowed to access the BSs in tiers $1, \ldots, n_{\text{open}}$. For example, a macro-femto HCN with a mix of open access (OA) and closed subscriber group (CSG) femtocells would be represented as an HCN with $n_{\text{tier}} = 3$ and $n_{\text{open}} = 2$, with tier 1 representing the macrocells, tier 2 the OA femtocells, and tier 3 the CSG femtocells.

Next, we assume that each BS in a tier transmits with the maximum power allowed for BSs in that tier. This immediately models the case of reference symbols (LTE) or pilot channels (HSPA), but also covers the case of data channels if we assume that the cells are all fully loaded. For any given user, a single *candidate serving BS* is chosen from each tier according to some criterion, and the BS that actually serves the user is chosen from among these candidate serving BSs according to some criterion. In Chapter 5, we study some commonly used criteria for selection of candidate serving BSs from the tiers, and for the overall serving BS from among the candidate serving BSs.

2.1.2 Basic definitions

Let Φ_i denote the set of locations of the BSs in tier i, $i = 1, \ldots, n_{\text{tier}}$. A BS in tier i is identified by its location $b \in \Phi_i$. In other words, we often write "the BS $b \in \Phi_i$" to mean "the BS in tier i located at b." Note that, in this notation, $b \in \mathbb{R}^2$.

Instantaneous received power at the user

A BS $b \in \Phi_i$ at distance R_b from the user transmits with power P_i^{tx} and is received at the user with power

$$(\forall b \in \Phi_i)\ Y_b = \frac{P_i H_b}{R_b^{\alpha_i}}, \quad \{H_b\}_{b \in \Phi_i} \text{ i.i.d.}, \quad P_i = K_i P_i^{\text{tx}}, \quad i = 1, \ldots, n_{\text{tier}}, \tag{2.1}$$

where α_i and K_i are, respectively, the slope and intercept of the path-loss model (1.1) that describes links between BSs in tier i and the user location, and we assume that the fading processes on all links from all BSs to the user are *independent*. The fade attenuations $\{H_b\}_{b \in \Phi}$ on the links to the user from the BSs in tier i are therefore independent. Further, the fade attenuations on the links to the user from all BSs in tier i (except possibly from the candidate serving BS in tier i, if $i \in \{1, \ldots, n_{\text{open}}\}$ – see below) are assumed *identically distributed* with common *continuous* distribution described by the PDF $f_{H_i}(\cdot)$. Finally, let us assume that the thermal noise power at the user receiver (measured over the same bandwidth as the received powers) is N_0.

For any tier $i = 1, \ldots, n_{\text{tier}}$, the total received power at the user from all BSs in tier i is given by

$$W_i \equiv \sum_{b \in \Phi_i} Y_b = \sum_{b \in \Phi_i} \frac{H_b}{R_b^{\alpha_i}}, \quad i = 1, \ldots, n_{\text{tier}}.$$

For each *open* or *accessible* tier $i = 1, \ldots, n_{\text{open}}$, let us denote the *candidate serving* BS from tier i by B_i. The received power at the user from the candidate serving BS in accessible tier i is therefore $U_i \equiv Y_{B_i}$, while the interference from all other BSs in the accessible tier i is

$$V_i \equiv \sum_{b \in \Phi_i \setminus \{B_i\}} Y_b = \sum_{b \in \Phi_i \setminus \{B_i\}} \frac{H_b}{R_b^{\alpha_i}} = W_i - U_i, \quad U_i \equiv Y_{B_i}, \quad i = 1, \ldots, n_{\text{open}}. \tag{2.2}$$

Serving tier, serving BS, and instantaneous SINR

If there are multiple accessible tiers in the network, one of the *candidate serving* BSs is chosen (based on any of several criteria to be defined later) as the *serving* BS. If there is only one accessible tier, the candidate serving BS for this tier (chosen according to some criterion) is also the serving BS. Let us denote the index of the *serving tier* by I. Then I is a random variable taking values in $\{1, \ldots, n_{\text{open}}\}$, and the serving BS is B_I. If $I = i$, the SINR at the user is given by

$$\Gamma_i \equiv \frac{U_i}{\sum_{\substack{j=1 \\ j\neq i}}^{n_{\text{open}}} U_j + \left[\sum_{k=1}^{n_{\text{open}}} V_k + \sum_{k=n_{\text{open}}+1}^{n_{\text{tier}}} W_k + N_0\right]}, \quad i = 1, \ldots, n_{\text{open}}. \tag{2.3}$$

Special case: single tier

Consider a deployment with a single tier: $n_{\text{tier}} = n_{\text{open}} = 1$. Then $I = 1$ with probability 1. Without loss of generality, denote the distance of the serving BS B_1 from the user by R_*, and the corresponding fading attenuation by H_* (instead of H_{B_1}). Dropping the subscript $i = 1$ for convenience, the SINR at the user is given from (2.3) by

$$\Gamma = \frac{U}{V + N_0}, \quad U = \frac{PH_*}{R_*^{\alpha}}, \quad P = KP^{\text{tx}}, \tag{2.4}$$

and V is given by (2.2) with $i = 1$. Suppose now that the common distribution of all the fading attenuations on all links to the user location from all BSs (including H_*, on the link from the serving BS to the user) is the exponential distribution with unit mean, Exp(1). It follows that the CCDF of Γ is given by

$$\mathbb{P}\{\Gamma > \gamma \mid R_* = r_*\} = \exp\left(-\frac{\gamma r_*^{\alpha}}{P} N_0\right) \mathbb{E}\left[\exp\left(-\frac{\gamma r_*^{\alpha}}{P} V\right) \mid R_* = r_*\right], \tag{2.5}$$

which may be seen as the counterpart of the probability (1.4), defined by conditioning only on the distance to the serving BS.

2.2 SINR distributions

The *cumulative distribution function* (CDF) of the wideband SINR at an arbitrary user is often referred to as the *geometry* in the literature. It is a representation of the "environment" that an arbitrarily located user in the network will encounter.

Throughout the book, we are interested in two kinds of joint SINR CCDFs, which we define below.

2.2.1 Joint CCDF of SINRs from candidate serving BSs

The first is the joint CCDF of the SINRs at the user when receiving from the candidate serving BSs of (all or a subset of) the open tiers:

$$\mathbb{P}\{\Gamma_{j_1} > \gamma_1, \ldots, \Gamma_{j_k} > \gamma_k\}, \quad \gamma_1 > 0, \ldots, \gamma_k > 0, \tag{2.6}$$

for all $k \leq n_{\text{open}}$ and distinct indices $(j_1, \ldots, j_k)$ such that $1 \leq j_1 < j_2 < \cdots < j_k \leq n_{\text{open}}$.