1 Introduction

As old as I have become, many developments my eyes have seen. When I was a young man, changes used to take place from time to time, every now and then. Nowadays, they occur so often, life itself seems to change every day.

—A 97-year-old man commented, making reference to the increasing pace of the developments in the communications field through the last century

These comments, although exaggerated at first sight, may not be so, in light of the advancements seen in the telecommunication industry since Graham Bell carried out the first successful bidirectional telephone transmission in 1876 [1]. Since then, society has witnessed

- most long-distance communications having at least a wireless component, freed from wires and operated through air at the speed of light,
- wireless and mobile communications made available to over 8.6 billion connections across the globe [2],
- new types of communications and social interactions emerging through both the Internet [3] and social networking [4] and
- many other breakthroughs, which have certainly changed our everyday lives.

These developments, although of great importance to the 97-year-old man, are probably just small steps towards a new era – the era of digital and pervasive communications – which will continue to change the world we are inhabiting in unpredictable and fascinating manners.

As a matter of fact, today, we are on the brink of another significant societal change. While the network has mainly served humans up to now, this capability will increasingly be extended to machines in the near future too. By 2022, it is expected that there will be not only 8.4 billion handheld or personal mobile-ready devices, but also 3.9 billion machine-to-machine connections [2]. The emergence of this machine-originated data traffic will drive further the demand for network capacity, but also impose additional requirements on network performance, mainly in the areas of end-to-end latency and reliability. These are currently the major challenges for many new applications.

Nowadays, most of the data services reside on the Internet, far away from the user equipment (UE), where the speed of light becomes one of the main factors limiting end-to-end latency. To address this problem, processing will have to move closer to the UE, e.g. into a cloud computing infrastructure, which will extend – and act as a ramification of – the network. In addition, an intelligent and adaptive network management and a well-designed congestion control can also help to significantly enhance reliability, thus enabling new real-time applications, such as augmented reality or efficient machine communication.

With these new requirements and changes, communication networks are evolving to become our main interface with the virtual world, and increasingly also with the physical one. This future network will simplify and automate many aspects of life, allowing one to effectively "create time," by improving the efficiency in everything we do [5].

Making this vision of the future network a reality will require from a technical perspective both

- ultra-broadband wireless access, providing orders of magnitude improved throughput, delay and reliability as well as quality of service (QoS) control and
- a highly adaptable and remotely programmable cloud computing infrastructure located close to the edge of the network.

Throughout this book, we argue that small cells, and more specifically, ultra-dense deployments are one of the answers to the technological challenges of creating an ultra-broadband wireless access that connects mobile UEs, machines and objects to a processing cloud engine. In more detail, this book serves as a tool to shed new light on the fundamental understanding of ultra-dense networks.

As an introduction to the content of the book, the remainder of this chapter first depicts the current industry capacity challenge and follows with an overview of the small cell technology and its history, from both an industry and an academic perspective. Then, the individual parts and chapters of the book are introduced, as well as their relationship to various aspects of deploying and operating small cell networks. To conclude, some of the key nomenclature used in this book is presented together with a list of the most relevant publications in the field of ultra-dense networks by the authors of this book.

As an important disclaimer, let us note that this book is going to focus on the study of single antenna UEs and base stations (BSs), thus using single-input single-output transmission modes. As a result, this book does not consider either multiple-input multiple-output (MIMO) or multi-cell coordinated transmissions/receptions, and all the statements within are done accordingly.

1.1 The Capacity Challenge

Voice-based services, such as voice over internet protocol (VoIP), were the killer applications at the beginning of this century, demanding an average of tens of kilobits per second per UE for this type of connection [6], while the streaming of high-definition (HD) video is probably the most popular service today, requiring tens of megabits per second per video feed [7]. Future services, however, such as three-dimensional visualization, augmented and virtual reality, online gaming using multiple displays and the robot-to-robot exchange of HD laser imaging detection and ranging (LIDAR) maps will require much more capacity, with expected average throughputs per UE exceeding 1 Gbps [8] – and who knows what else tomorrow will bring?

With this enormous challenge of improving the average throughput per UE by orders of magnitude, before making any decision on any technology investment, which is likely to be costly, it is advised that network operators and service providers understand well the different dimensions that they have to improve wireless capacity.

In a simplified form, the Shannon–Hartley theorem [9],

$$C = B \cdot \log_2\left(1 + \frac{S}{N}\right), \tag{1.1}$$

provides an insight into which are the variables that influence the amount of information – capacity, C, in *bps* – that a transmitter can send to a receiver

- over a communication channel of a specified bandwidth, B, in Hertz
- with a received signal power, S, in Watts
- in the presence of an additive white Gaussian noise (AWGN) power, N, in Watts.

From this theorem, it can be inferred that the capacity, C, of a UE can be scaled up by increasing

- the bandwidth, B, per UE and/or
- the signal-to-noise ratio (SNR), $\frac{S}{N}$, of such UE, or more accurately, the signal-to-interference-plus-noise ratio (SINR), $\frac{S}{I+N}$, of the UE in a multi-cell multi-UE network, like the ones that will be studied in this book, where I stands for Gaussian interference in Watts.

Importantly, equation (1.1) also shows that to scale the capacity, C, of a UE, increasing the bandwidth, B, per UE is generally a more promising technique than increasing the SINR, $\frac{S}{I+N}$, of such a UE, since the former yields a linear scaling, while the latter only a logarithmic one.[1]

---

[1] Even though multi-antenna technology is not considered in this book, for completeness, it should be noted that multiple antennas can be used to either

- leverage spatial multiplexing through MIMO techniques or
- increase the SINR, $\frac{S}{I+N}$, of a UE via beamforming.

In particular, MIMO transmissions/receptions, can take advantage of the spatial resources, and linearly increase the capacity, C, of a UE with the number of spatial streams multiplexed. This can be treated as a “virtual” increase of the bandwidth, B, per UE. When taking a cell or a network perspective, it should also be noted that multi-user MIMO, coordinated beamforming and multi-cell coordinated transmissions/ receptions can be used to increase the spatial multiplexing in the cell or the network and/or the SINR, $\frac{S}{I+N}$, of a UE. Readers interested in related topics are referred to [10] and references therein.

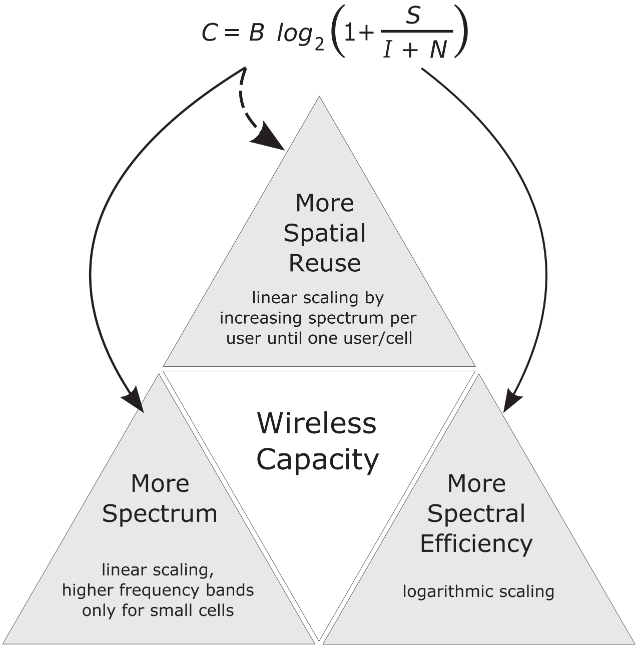

Figure 1.1 Dimensions for scaling the capacity of a UE in a wireless network [8].

With this in mind, a fundamental question arises:

How can we increase the bandwidth, B, per UE?

In a network with multiple UEs, the bandwidth, B, per UE can be scaled up by either increasing

- the amount of frequency resources invested into the network and/or
- the network densification, and in turn, its associated spatial frequency reuse.

For the sake of clarity, it should be noted that the reduced cell size in a denser network results in an improved spatial reuse of the frequency resources, since there are less UEs in each cell to share the available cell bandwidth. As a result, each UE has access to more of such frequency resources.

Overall, as depicted in Figure 1.1, this leaves one with three main approaches to enhance the capacity, C, of the UE, i.e. using a wider bandwidth, deploying a denser network and improving the signal quality – the SINR of the UE, where the two first ones may be more appealing due to their intrinsic linear scaling of capacity.

To put things in context, and show how each of these three degrees of freedom have historically contributed to the increase of the capacity of practical networks, Webb [11] put together an interesting analysis, indicating that, between 1950 and 2000, such network capacity has increased around

- 2700× from densifying the network with smaller cells,

- 15× by using more bandwidth in the sub-6 GHz bands (from 150 MHz to 3 GHz) and
- 10× by improving the spectral efficiency (waveforms and multiple access techniques, modulation, coding and medium access control (MAC) methods such as scheduling, hybrid automatic repeat request (HARQ), etc.).

From this study, it is clear that – by far – the majority of the capacity gains in the past were achieved by increasing the spatial frequency reuse through densifying the network with smaller cells. This leaves one with the following question:

How much further can we increase the spatial reuse by reducing the cell size?

Answering the above question from a theoretical perspective is one of the primary goals of this book.

For completeness, and before proceeding any further, it should be noted here at this point that the operation at higher carrier frequencies, e.g. millimetre wave bands, also offers the possibility of accessing large amounts of spectrum and the associated very wide bandwidths, thus enabling extreme data rates.

Higher carrier frequencies, however, are also associated with higher radio frequency attenuations, which limit their network coverage. Although this can be compensated to some extent by means of multi-antenna technologies – and more in detail through beamforming – a substantial coverage disadvantage will always remain for a network operating at higher carrier frequencies [12].

Another challenge with the operation at higher frequency bands is the regulatory aspects. For non-technical reasons, the rules defining the allowed radiation may change in these higher frequency bands, from a specific absorption rate (SAR)-based limitation to a more effective isotropic radiated power (EIRP)-like limitation. These restrictions may also impose further coverage constraints [13].

More importantly, the millimetre wave technology in general – and the beamforming one in particular – are not mature as of today, or at least, are not cost-effective for ultra-dense deployments. The features required to deal with

- the beam alignment and tracking, possibly needed at both communication ends and
- the related issues arising from unexpected device rotation, blockage and mobility

are still of high complexity, and lack of robust QoS provisioning. This together with the large energy consumption of current millimetre wave access points make this solution still too expensive.

As a result, the more mature and developed sub-6 GHz technology still remains a front runner for ultra-dense deployments, especially due to its ability to satisfy communication requirements in non-line-of-sight (NLoS) and outdoor-to-indoor propagation conditions. For these reasons, we focus on low carrier frequencies in this book, and analyze in detail the impact of network densification and the increase of spatial reuse on network performance, as posed earlier. However, due to its potential, we leave the door open to the analysis of millimetre wave deployments for next editions of this book, and in particularly to the inter-working of sub-6 GHz and millimetre wave technology.

1.2 Network Densification

In a multi-cell multi-UE network, UEs being served within a cell share the available bandwidth. Thus, reducing the cell size – while deploying more cells to maintain the same level of coverage – also reduces the number of UEs per cell, and in turn, increases the bandwidth per UE. Through this approach, the bandwidth per UE can be increased, until each cell only serves a single UE. When densifying further, beyond this point, only the signal power – and potentially the SINR – of the UE can be improved by reducing the distance between the UE and its serving small cell BS.

Overall, by increasing the bandwidth, B, per UE, the capacity, C, scales up linearly, until the one UE per cell limit is reached, after which the scaling becomes only logarithmic through improvements on the SINRs of the UEs, as indicated by equation (1.1).

Figure 1.2 illustrates this capacity scaling behaviour, showing that with increasing cell densities, the capacity

- initially increases quickly due to the spatial frequency reuse, but then
- slows down, when the one UE per cell limit is reached, and the gains are mainly dominated by improvements on the SINRs of the UEs, through proximity gains [8].

From Figure 1.2, it is also important to note that the results were obtained assuming an active UE density of 300 BSs/km^2, typical in some dense urban scenarios, and that the one UE per cell limit is reached for an inter-site distance (ISD) of around 30–40 m. This indicates that there is still plenty of room for network densification in major cities like Manhattan and London, where the average ISD is around 200 m.

A second aspect of densification is that the required transmit power may reduce to an extent where its contribution to the total energy consumption becomes insignificant, and the processing power of the small cell BS becomes the dominant factor. Moreover, with reduced cell sizes, the required number of small cells to provide coverage increases, and as a result, many of them may not serve any UE for most of the time. However, they still consume energy and transmit unnecessary pilot signals, which may cause inter-cell interference. This issue can be addressed by introducing idle mode capabilities, where small cell BSs are only woken up to actively serve UEs. With efficiently controlled idle modes, the network energy consumption reduces and the SINR of the UE significantly improves.

The main challenge of network densification, however, is the issue of increasing costs for equipment, deployment and operational expenses. In this light, it is important to note that the cost of a small cell BS, estimated in 2015, only accounts for approximately 20% of the total deployment costs associated with outdoor small cells.

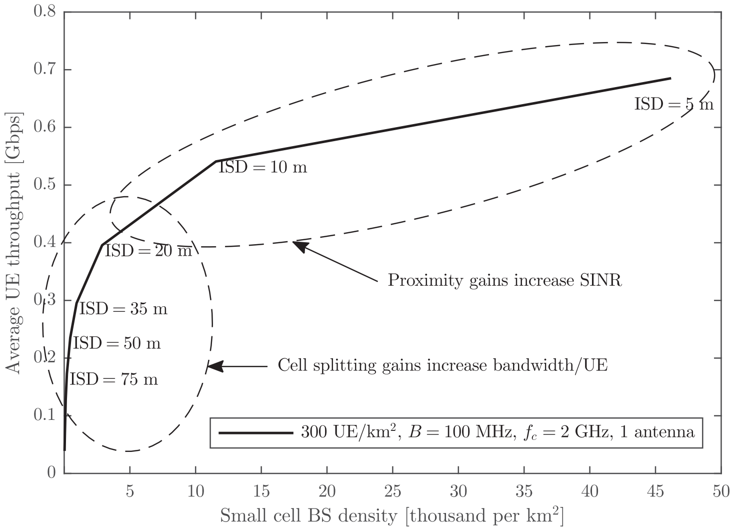

Figure 1.2 Capacity scaling with densification for different inter-site distances in a dense urban scenario considering a hexagonal BS deployment and a semi-clustered UE distribution. For more details on the scenario, models and results, the reader is referred to [8].

The majority of the costs are site leasing (26%), backhaul (26%), planning (12%) and installation (8%) [14]. The good piece of news is that this challenge can be addressed by changing the deployment model from an operator deployment to a “drop and forget” end-user deployment, and reusing the existing power and backhaul infrastructure. In this model, the end-user simply connects the small cell BS to the power and the backhaul, which then triggers a fully automatic configuration, and a continuous self-optimization process during operation. This end-user deployment model is feasible for both the residential and the enterprise markets.

Due to cost and performance reasons, it is also important to highlight that it becomes increasingly important to deploy the small cell BSs wherever the UEs are, since small cells cannot compensate for misplacement as well as larger cells do. If the small cell BSs are not deployed in an intelligent manner following the distribution of UEs, a larger number of small cell BSs will be required to achieve the one UE per cell limit. However, accurate UE demand distributions are hard to derive today, because of the limited accuracy of conventional localization techniques in cellular networks, such as triangulation. One may think of using more accurate techniques such as the global positioning system (GPS) for this purpose, but its performance is poor indoors, where 80% of the traffic demand is located [15]. Advanced planning tools for small cell deployment are still an open challenge.

In summary, densification continues to have a high potential to increase capacity until reaching the one UE per cell limit. To maintain high performance and energy efficiency, idle mode capabilities that switch off small cells when they are not serving

UEs are necessary. Transitioning to a “drop and forget” deployment by the end-user has a high potential to reduce the deployment and operational costs.

1.3 A Brief History of Small Cells

In this section, we provide an overview of the small cell technology and its history, first from an industry perspective, and subsequently from a theoretical one. This last part of the chapter serves as an introduction to the rest of the book, which will be formalized by the outline presented in the next section.

1.3.1 From an Idea to a Market Product

The idea of the small cell has been around for over three decades [16]. Initially, “small cell” was the term used to describe the cell size in a metropolitan area, where a macrocell – with a cell diameter on the order of kilometres – would be split into a larger number of smaller cells with reduced transmit power, known today as metropolitan macrocells or microcells. These small cells had a cell diameter of a few hundreds of metres.

In the 1990s, picocells appeared with even a smaller cell size, between a hundred metres to around a few tens of metres [17]. These “more traditional” small cells were used for coverage and capacity infill, that is, where macrocell penetration was insufficient to give a good connection or where the macrocell was at its capacity limit. These types of small cell BSs were essentially a smaller version of the macrocell BSs, which also had to be planned, managed and interfaced with the network. This last point is probably the most important reason why small cells – other than metropolitan macrocells or microcells – have not gained much popularity for quite some time. Essentially, the costs associated with deploying and running a large number of small cells outweighed the performance advantage that this kind of cellular topology provided.

In the 2000s, new thinking on the deployment and configuration of cellular systems began to address the cost and the operational aspects of small cell deployments, which enabled the cost-effective deployment of even smaller cells [18]. Such thinking crystallized in the home BS concept first [19] and the femtocell one later [20]. A femtocell is a low-cost cellular BS with advanced auto-configuration and self-optimization capabilities, which allows the end-user – without any operator involvement – to deploy this small form factor BS in a plug-and-play manner within the home. Femtocells use a broadband internet connection as backhaul, and connect to the cellular network through dedicated gateways, which enables a better scaling to millions of femtocell BSs. Early results on the performance of 3G universal mobile telecommunication system (UMTS) femtocells were presented in [21–23], which were shortly afterward extended with a bulk of studies on self-optimization and offloading strategies, multiple antenna techniques and energy management methods [20, 24–30]. Soon after, results on fourth-generation (4G) wireless interoperability for microwave access (WiMAX) and long-term evolution (LTE) small cells followed too [31–36]. Femtocells also emerged as the first step towards a heterogeneous network deployment model [37–42].

Following this early research and development in the femtocell area, in 2007, several industry players advocating for small cell technology formed the Femto Forum – rebranded as the Small Cell Forum in 2012 – to create a venue for promotion, standardization and regulation. Moreover, governments started funding research projects on femtocells, e.g. the European Union ICT-4-248523 BeFEMTO project, which focused on the analysis and the development of 4G LTE-compliant femtocell technologies [43]. To highlight the success of the femtocell in the research fora, it is worth highlighting that the number of publications registered in the Institute of Electrical and Electronics Engineers (IEEE) Xplore digital library [44] including the word “femtocell” or “femtocells” increased from 3 in 2007 to 11 (2008), 52 (2009), 117 (2010), 1088 (2015) and 3178 (2019).

The first commercial deployments of residential femtocells started in 2008 when Sprint launched a nationwide service in the United States, followed in 2009 by Vodafone in Europe and Softbank in Japan. Since then, small cell technology has quickly proliferated. The number of deployed small cells for the first time exceeded those of macrocells in 2011 [45], and over 77 operators used small cells worldwide in 2015 [46]. The business impact of small cells took off with 4G LTE, with this small cell type being – by far – the most widely deployed today, where multimode cells with additional 3G UMTS and/or Wi-Fi capabilities are also widely available [47].

Due to its success, the scope of femtocells was then extended from residential deployments to public indoor spaces, which generate most of the cellular traffic. By 2015, 71 operators had deployed in-building small cells, known as pico- or microcells, in enterprise or public buildings [46], and from 2020 to 2025, it is expected that the number of such in-building small cells will grow from 10.7 to 208.6 million [48]. Their design is aimed at reducing planning and deployment costs, decreasing the need for large customer support teams and eliminating the need for massive reprovisioning. The generally good availability of internet protocol internet protocol (IP) backhaul, such as Ethernet, in enterprise and public buildings is an important deployment advantage. However, the overall system configuration and the overlap with outdoor macrocells and microcells must be monitored and well managed. To this end, in-building small cell deployments are equipped with per-call and QoS analytics, as well as self-organizing network (SON) features [49].

The expansion of femtocells to the outdoor space is more difficult due to challenges such as cost, site rental, backhaul availability, network provisioning and management as well as monetization issues. However, even in this area, small cells have proved to be a viable approach in the form of metrocells – smaller and more flexible versions of metropolitan macrocells or microcells, with whom they share many hardware and software features, most notably the support of a high number of simultaneous UEs. Metrocells, however, shine for their SON capabilities, providing self-configuration and self-optimization of, e.g. neighbour relation management, inter-cell interference mitigation and handover parameter configuration, among others. Outdoor small cells are mainly serving operator-deployed, public networks in urban, suburban or rural environments, although a number of outdoor small cell deployments are already dedicated to particular businesses and enterprises (e.g. oil drilling rigs or power stations). In the fifth-generation (5G) era, when more small cells will be related to industrial and internet of things (IoT) services, there will be more vertical-specific small cells deployed outdoors, managed by enterprise specialists. Recent surveys indicate that in the outdoor environment, the growth rate of small cell deployments will double that of indoor ones from 2020 to 2025, 41% versus 20% [48]. By 2025, outdoor small cell deployments will reach 2.76 million, with urban scenarios reaching a total installed base of 11.2 million small cells [48].

Overall, it is expected that the installed base of small cells will reach 70.2 million in 2025 as operators seek to densify their networks. This growth will likely be led by Asia-Pacific and North America, with Europe lagging [48]. While residential deployments will continue to rise, the non-residential market is expected to grow far more quickly, at an annual growth of 36%, accounting for 75% of annual deployments and 55% of the total installed base in 2025 led by the urban public and enterprise public sectors. The total installed base of 5G new radio (NR) or multimode small cells in 2025 is predicted to be 13.1 million, over one-third of the total in use [48]. From 2025, 5G NR small cells are expected to overtake 4G LTE-only as well as combined 4G LTE and 5G NR models [48].

1.3.2 The Evolution of the Small Cell Theoretical Understanding: A Brief Summary of This Book

As indicated in the previous sections, densification has a high potential to increase capacity, until reaching the one UE per cell limit, and the market is already heading down this path, deploying denser and denser networks in dense urban scenarios, as well as indoors, in enterprises and factories. However, there are fundamental questions around whether ultra-dense networks – which do not exist out there yet – will behave similarly as today's more sparse ones, or whether they follow different fundamentals, which may impact performance. For example:

> *Will the larger number of small cells create an inter-cell interference overload that will render any communication impossible?*

Answering this and other fundamental questions related to ultra-dense networks is the objective of this book.

The Old Understanding

When it comes to small cell deployment and network performance analysis, the theoretical work of M. Haenggi, J. G. Andrews, F. Baccelli, O. Dousse and M. Franceschetti stands out. In their seminal work [50] and references therein, the authors created a mathematical framework based on stochastic geometry to analyze the performance of random networks in a tractable manner.

In a nutshell, this mathematical stochastic geometry framework allows one to theoretically calculate, sometimes even in a closed-form expression, the coverage probability of a typical UE, which is defined as the probability that the SINR, γ, of the typical UE is larger than an SINR threshold, γ_0, i.e. $\Pr\left[\text{fl} > \gamma_0\right]$. Based on this coverage probability – also known as success probability – the SINR-dependent area spectral efficiency (ASE) in bps/Hz/km^2 can also be investigated, among other metrics.

This framework has become the de facto tool for the theoretical performance analysis of small cell networks in the entire wireless community. Good tutorials and more references on the fundamentals of this mathematical tool can be found in [51–54] and the references therein. Chapter 2 of this book and in more detail Section 2.3 will also provide a more detailed introduction to the topic.

Many efforts have been made since 2009 to extend the capabilities of this stochastic geometry framework to improve the understanding of small cell networks. M. Haenggi et al. further developed the framework to account for different stochastic processes different that the basic homogeneous Poisson point process (HPPP) [55], as well as distinct performance metrics, such as the typicality of the typical UE [56] and the transmission delay [57], among others. T. D. Novlan et al. further extended the framework to study uplink transmissions, calculating the aggregated inter-cell interference, using the probability generating functional of the HPPP [58]. M. Di Renzo et al. also did a good number of extensions, by considering more detailed wireless channel characteristics in the modelling, such as other non-HPPP distributions, building obstructions, shadow fading and non-Rayleigh multi-path fast fading, of course, at the expense of tractability [59].

When it comes to the analysis of different wireless network technologies and features using stochastic geometry, it is worth highlighting the extensive work of J. G. Andrews, V. Chandrasekar, H. S. Dhillon et al., which touches on spectrum allocation [37], sectorization [38], power control [39], small cell-only networks [60], multi-tier heterogeneous networks [61], MIMO [62], load-balancing [63], device-to-device communications [64], content caching [65], IoT networks [66] and unmanned aerial vehicle (UAV) communications [67], to cite a few.

Regarding higher frequency bands and the massive use of antennas, the studies of R. W. Heath, T. Bai et al. stand out, for example, those on massive MIMO [68] and millimetre wave [69] performance analysis, random blockage [70], millimetre wave ad hoc networks [71] and secure communications [72], shared millimetre wave spectrum [73] as well as wireless power systems [74].

For further reference, and with regard to the analysis of other relevant network aspects, it is worth pointing out the research of G. Nigam et al. on coordinated multipoint joint transmission [75], H. Sun et al. on dynamic time division duplex [76] and Y. S. Soh et al. on energy efficiency [77]. Many other analyses studying different types of stochastic processes, performance metrics, wireless characteristics and network features can be found in the literature, which show both the generality and the strength of stochastic geometry framework.

Among all the mentioned results, one of the most important theoretical findings is that by J. G. Andrews and H. S. Dhillon et al., concluding that the fears of an inter-cell interference overload in small cell networks are not well-grounded, neither in a small cell-only network [60] nor in a heterogeneous one [61]. Instead, their results showed that the increase in the inter-cell interference power due to the larger number of interfering small cell BSs in a dense network is exactly counterbalanced by the increase in the signal power due to the closer proximity of transmitters and receivers. This conclusion is powerful, meaning that an operator can continually densify its network – no problem – and expect that

- the spectral efficiency in each cell stays roughly constant, and as a result that
- the network capacity – or in more technical words, the ASE – linearly grows with the number of deployed cells.

This behaviour – or capacity scaling law with the small cell BS density – is referred hereafter in this book as *the linear capacity scaling law*.

This exciting message created big hype in the community, and also in the industry, presenting the small cell BS as the ultimate mechanism in providing a superior broadband experience. Consequently, this raised a new fundamental question:

> *Can we infinitely reduce the cell size to achieve an infinite spatial reuse, and in turn, an infinitely large capacity?*

Unfortunately, a few important caveats to realize such a linear scale of capacity in an ultra-dense network were quickly found. Among them, it is worth highlighting

- the need for an open-access operation [78] and
- the impact of the transition of a large number of interfering links from NLoS to line-of-sight (LoS) [79],

which will be discussed in the following paragraphs.

The Effect of the Access Method

Closed-access operation provides an experience comparable to Wi-Fi access point, in which the owner of the small cell BS can select which UEs can associate to it. This is an appealing model for the small cell owner but prevents the UE to connect to the strongest cell. This degrades the SINR of the UE and breaks the linear scale of capacity observed in both [60] and [61]. This is because the inter-cell interference power can now grow much faster than the signal power. This is particularly true when the UE, let us say in a block of apartments, moves away from its closed-access small cell – to which it can connect – and gets closer to a neighbouring one – which it cannot access.

Open-access operation has been widely adopted in small cell BS products to address this issue and restore the linear capacity scaling law. Since the performance impact of closed- and open-access operation is intuitive and well understood [78], this topic is not theoretically treated in this book.

The Impact of the NLoS to LoS Inter-Cell Interference Transition

A more fundamental problem than the access method was found in [79], which showed that, even if open-access operation is adopted, the inter-cell interference power can grow faster than the signal power, with the consequent degradation of the SINR of the UE, when the small cell BS density grows ultra-dense.

To understand this phenomenon, it is important to note that the linear scale of capacity presented in both [60] and [61] was obtained with the assumption of a single-slope path loss model, meaning that both the inter-cell interference and the signal powers decay at the same pace, $d^{-\alpha}$, over a given distance, d, where α is the path loss exponent.

Although simplistic, when the path loss exponent is "fine-tuned," this single-slope path loss model is applicable to sparse networks, such as metropolitan macrocell and microcell ones. However, this model may be inaccurate for denser networks, where the small cell BSs are deployed below the clutter of man-made structures. This is because the probability of the received signal strength abruptly changing due to a change in the LoS condition of the interfering and/or serving links is much larger in a dense network, where the small cell BSs are deployed at street level.

To model this critical channel characteristic, whether the UE is in NLoS or LoS with a small cell BS, the use of

- a multi-slope path loss model considering NLoS and LoS transmissions and
- a probabilistic function governing the switch between them

was proposed in [79], and implemented over the theoretical analysis framework presented in [60] and [61].

Intuitively speaking, the key difference between the single-slope and this new multi-slope path loss model is that

- a UE always associates to the nearest small cell BS in the former, while
- a UE may be connected with a further but stronger small cell BS in the latter.

This probabilistic model introduces randomness and renders decreasing distances less useful. The results of this new analysis showed an important fact. There is a small cell BS density region where

- the strongest interfering links transit from NLoS to LoS, while
- the signal ones stay LoS dominated due to the close proximity between the UEs and their serving small cell BSs.

As a result, the SINR of the UE – and the spectral efficiency in the small cell – do not stay constant, and the network capacity does not grow linearly with the small cell BS density anymore.

In this book, we will refer to this important phenomenon – the loss of linearity in the scale of capacity due to the transition of a large number of interfering links from NLoS to LoS – as *the ASE Crawl*.

Importantly, it should be noted that the ASE Crawl is not the result of a mathematical artefact, and that its impact was shown by the real-world experiments in [80], in which a densification factor of 100× led to a network capacity increase of 40× – clearly not a linear increase.

This new theoretical finding showed that the small cell BS density matters, and that it should not be taken lightly during the planning of a small cell deployment, as the ASE Crawl has a counterproductive effect. For completeness, however, it is fair to note that once the network density is ultra-dense, and the strongest interfering links transit from NLoS to LoS, both the inter-cell interference and the signal powers will again grow at a similar pace, as they are all LoS dominated, and the path loss decays at a similar rate. This restores the linear scale of capacity with the small cell BS density, but at a lower rate, as the path loss exponent in LoS is generally smaller than that in NLoS.

At this point in our story telling, it is worth noting that

- the need for an open-access operation and
- the impact of the transition of a large number of interfering links from NLoS to LoS

served as a wake-up call to the theoretical research community, which began to realize the importance of an accurate network and channel modelling, and started to review their understanding of ultra-dense networks. Some asked themselves whether some other important details were overlooked. Details that could change the performance trends expected for ultra-dense networks until then. This brought back again the original question:

Will the network capacity linearly grow with the small cell BS density or not?

Two frameworks should be highlighted in this quest, which will be discussed in the following:

- the implications of considering the near-field transmissions [81]; and
- the effect of the antenna height difference between the UEs and the small cell BSs [82].

The Myth of the Near-Field Effect

While looking at a more accurate channel model that could reveal new findings, the research in [81] presented a reasonable conjecture, indicating that the path loss exponent should be an increasing function of the distance, and proposed to capture this in a multi-slope path loss model similar to that presented in [79]. To illustrate the thinking behind it, the authors provided the following argument: *"There could easily be three distinct regimes in a practical environment:*

- *a first distance-independent 'near-field',* where $\alpha_1 = 0$,
- *second, a free-space like regime where* $\alpha_2 = 2$, *and*
- *finally, some heavily-attenuated regime where* $\alpha_3 > 3$,"

and then posed the following question:

What happens if densification pushes many BSs into the near-field regime?

The mathematical results derived in [81] provided an answer to such a question and concluded that the inter-cell interference power can grow faster than the signal power when the network is ultra-dense, even if both the inter-cell interference and the signal powers are LoS dominated. The intuition behind this is that when the UE enters the near-field range,

- the signal power is bounded, as the path loss becomes independent of the distance between such UE and its serving BS, i.e. $\alpha_1 = 0$, while
- the inter-cell interference power continues to grow, since more and more interfering small cell BSs approach the UE from every direction, when the network marches into the ultra-dense regime.

As a result, once the signal power enters the near-field range, the SINR of the UE cannot be kept constant, and will monotonically decrease with the small cell BS density.

This finding raised the alarm again, as it indicates that the near-field effect could lead to a void of capacity in an extreme densification case due to the overwhelming inter-cell interference.

Subsequent results on the topic, based on measurements, however, have shown that this alarm was unfounded [83]. The measurements, shown in figure 1 of [83], indicate that the near-field effect only takes place at sub-metre distances in practical ultra-dense networks with a carrier frequency of around 2 GHz and an antenna aperture of a few wavelengths. This is in line with the near-field effect theory, which indicates that the near-field is that part of the radiated field, where the distance from the source, an antenna of aperture, D, is shorter than the Fraunhofer distance, $d_f = 2D^2/\lambda$ [84], where λ is the wavelength. As a result, a BS density of around 10^6 BSs/km^2 would be needed for this near-field effect to be an issue, and this is unlikely to be seen in practice – at least as of today – as it means having one small cell BS every square metre.

Such result renders the near-field effect issue negligible in practical deployments. For this reason, despite being interesting and relevant, this topic is not theoretically further treated in this book.

The Challenge of the Small Cell Base Station Antenna Height

In parallel with the work done to shed new light on the performance impact of the near-field effect, new investigations presented yet another reason why the inter-cell interference power could grow faster than the signal power in an ultra-dense network, i.e. the antenna height difference, L, between the UEs and the small cell BSs [82].

By considering the antenna heights of both the UEs and the small cell BSs in the multi-slope path loss model presented in [79], and carrying a theoretical performance analysis on a fully loaded network, it was shown that

- the distance between a UE and its interfering small cell BSs decreases faster than the distance between such UEs and its serving BS when densifying the network.

This is because the UE can never get closer than a distance, L, to its serving small cell BS, since the UE cannot climb up towards it and that

- the inter-cell interference power, as a consequence, increases faster than the signal power at such a UE in a dense network.

As a result, this faster increase of the inter-cell interference power results in

- a decline of the SINR of the UE, which can be fast with the increase of the small cell BSs density due to the sheer number of interfering small cell BSs in an ultra-dense network, and in turn,
- a potential total network outage in the ultra-dense regime,

thus putting again a question mark on the benefits of network densification.

Here, we should also highlight that the impact of this cap on the received signal power at the UEs due to the antenna height difference, L, between the UEs and the small cell BSs is not a mathematical construct. It was confirmed in [85] using the measurement data of [86], where an antenna height difference, $L = 4.5$ m, was considered. More importantly, and in contrast to the near-field effect, it should be also noted that this phenomenon occurs at more realistic and practical small cell BS densities of around 10^4 BSs/km^2, with small cell BS antenna heights of 10 m, thus being a more realistic threat.

Hereafter in this book, we will refer to this phenomenon – the continuous decrease in network capacity due to the antenna height difference between the UEs and the small cell BSs – as *the ASE Crash*.

Exploiting a Surplus of Small Cell Base Stations

When analyzing the previous results and trying to understand their implications, it is important to note that they all follow a traditional, macrocell-centric modelling assumption, in which the network is fully loaded, i.e. the number of UEs is always much larger than the number of BSs, and thus it is always safe to assume that there is at least one UE in the coverage area of every BS considered in the study. This assumption fits with macrocell as well as sparse small cell scenarios. Moreover, it is quite handy, as it allows to derive the capacity of the network at full load, and also makes the theoretical analysis of the network more tractable. However, in a dense deployment with many relatively small cell BSs, things are different. The probability of a small cell not serving any UE at a given time can be significantly high in some scenarios, and thus the always-on control signals have two negative impacts [13]:

- they impose an upper limit on the achievable network energy performance; and
- they cause inter-cell interference, thereby reducing the achievable data rates.

To address this, a number of mechanisms have been introduced for switching on and off small cell BSs – or at least their always-on signals – in the last years. For example, the 3rd-Generation Partnership Project (3GPP) LTE Release 12 [87] introduced mechanisms for turning on and off individual small cell BSs as a function of the traffic load to reduce the power consumption and the inter-cell interference. Moreover, sophisticated “lean carrier” approaches to allow a more dynamic on and off operation have been developed in 3GPP NR [12].

Embracing these practical considerations, the work in [88] revisited the ultra-dense network system model, and made a leap in theoretical performance analysis, accounting for finite active UE densities and idle mode capabilities. This new modelling brought new viewpoints, more industry aligned, which have led to a significantly different understanding of ultra-dense networks. In a nutshell, this work theoretically demonstrated the true benefits of an ultra-dense network, in which the number of small cell BSs is much larger than the number of UEs. This allows to reach the one UE per cell limit, where every UE can simultaneously reuse the entire spectrum managed by its serving small cell BS, without sharing it with other UEs. More importantly, it also showed how UEs can benefit from an improved performance in an ultra-dense network of this nature, because every small cell BS can

- tune its transmit power to the lowest possible one, just to cover its small intended range,
- and switch off its wireless transmissions through its idle mode capability, if there is no UE in its coverage area.

This results in both energy savings and a mitigated inter-cell interference; the latter because the control signals usually transmitted by active – but not empty – small cell BSs are switched off now, and do not interfere neighbouring transmitting cells.

From a theoretical perspective, and since multi-cell coordination is not considered here, when all UEs are served in an ultra-dense network, and the number of active small cell BSs is equal to the number of UEs, it is important to note that the number of interfering small cell BSs is automatically bounded, and thus so is the inter-cell interference. Since every active small cell BS serves a UE, no more active small cell BSs are needed. As a result, this bounded inter-cell interference power leads to an increasing SINR of the UE when densifying the network beyond this point, as the signal power continues to grow due to the closer proximity between a UE and its serving small cell BS. This leads to an enhanced overall network performance.

In the sequel of this book, we will refer to this phenomenon – the continuous increase in network capacity due to the surplus of small cell BSs with respect to UE and their idle mode capabilities – as *the ASE Climb*.

Channel-Dependent Scheduling and Multi-User Diversity

While the much larger number of small cell BSs with respect to that of UEs allows one to reach the one UE per cell limit, and in turn, a larger bandwidth, B, per UE, it also brings about a disadvantage. The lower number of UEs per small cell BS leads to a reduced multi-user diversity. In other words, a small cell BS has less UEs to choose from during its scheduling process, and thus, it is increasingly harder to opportunistically take advantage of potential constructive multi-path fast fading gains. Not only that, when the network goes ultra-dense, due to the closer proximity between a UE and its serving small cell BS, and the resulting higher probability of LoS transmissions, the radio-channel variations on a given time-frequency resource also become smaller and smaller. This also leads to a reduced multi-user diversity, given that the scheduler finds it increasingly harder to opportunistically find large multi-path fast fading gains.

The research in [89] studied multi-user diversity, and showed how the widely used channel-dependent schedulers indeed lose their ability to select a better UE for each scheduled time-frequency resource in each scheduling decision period, with the resulting loss in small cell capacity. It should be noted that the one UE per cell limit is the extreme case, where channel-dependent scheduling does not have any degree of freedom to select UEs during the scheduling process. This fact advocates for the use of simpler schedulers at small cell BS in ultra-dense networks, such as a simple round robin (RR) policy, to save on hardware processing complexity.

A New Capacity Scaling Law

Considering the above introduced theoretical findings, i.e. the ASE Crawl, the ASE Crash and the ASE Climb, new fundamental are questions raised, which are summarized in the following:

> *Will the negative impact of the ASE Crawl and Crash overweight the positive one of the ASE Climb, or the other way around?*
>
> *Which is the resulting capacity scaling law that best characterizes an ultra-dense network when considering all these characteristics?*

A new theoretical performance analysis came to answer these questions, using stochastic geometry, and proposed a new capacity scaling law for ultra-dense networks [90]. This is probably the most complete model and comprehensive capacity scaling law in the literature, as of the time of writing this book and up to the authors' knowledge. In short, this new study considered for the first time a model able to capture the combined effects and interactions of

- the transition of a large number of interfering links from NLoS to LoS,
- the antenna height difference between the UEs and the small cell BSs,
- a finite UE density, and the surplus of small cell BSs with respect to UE, as well as
- the idle mode capability at the small cell BSs,

and derived a new capacity scaling law, indicating that both the coverage probability and the ASE will asymptotically reach a maximum constant value in the ultra-dense regime.

Theoretically speaking, this research showed that in a densifying network

- the signal power caps because of the antenna height difference between the UEs and the small cell BSs, while
- the inter-cell interference power becomes bounded due to the finite UE density as well as the idle mode capability at the small cell BSs.

This results in a constant SINR of the UE in a densifying ultra-dense network of the above characteristics, leading to the mentioned asymptotic behaviour with the increase of the small cell BS density – *a constant capacity scaling law*. From this new capacity scaling law, it can be concluded that, for a given UE density, the network densification should not be abused indefinitely, but should be instead stopped at a given level, because any network densification beyond such point is a waste of both invested money and energy consumption.

At this point, to recap, it is important to note that this new constant capacity scaling law in the ultra-dense regime is significantly different from

- the initial linear scale of capacity introduced in both [60] and [61],
- the pessimistic ASE Crawl and ASE Crash, leading to a disastrous network performance with the densification, presented in [79] and [82], respectively, and
- the optimistic ASE Climb, discussed in [88],

and shows how the theoretical performance analysis and the understanding of ultra-dense networks have both improved up to now.

Dynamic Time Division Duplex to Make the Most of Ultra-Dense Networks

As a by-product of the lower number of UEs per small cell BS in a densifying network, it is also important to note yet another impact. The per-small cell aggregated downlink and uplink traffic demands become highly dynamic. Sometimes a small cell BS may have much more downlink than uplink traffic, while the opposite may be true at a different time instant or for another small cell BS in a neighbouring location. To address such scenarios, dynamic time division duplexing (TDD), i.e. the possibility for dynamic assignment and reassignment of time resources between the downlink and uplink transmission directions, was developed in 3GPP LTE Release 12 [87], and is a key 3GPP NR technology component [12].

Contrary to a static or a semi-dynamic TDD system, where the number of time resources devoted to downlink and uplink transmissions are preconfigured, and may not match the instantaneous traffic demands of a small cell, when embracing dynamic TDD, each small cell BS can provide a tailored configuration of downlink and uplink time resources, e.g. subframes, to meet its instantaneous downlink and uplink traffic requests. In other words, the transmission direction can be dynamically changed on short time periods in each cell. However, it is important to note that such flexibility does not come free, but at the expense of introducing inter-cell interlink interference. For example, the downlink transmission of a small cell may interfere with the uplink reception in a neighbouring small cell (downlink-to-uplink inter-cell interference), and vice versa, the uplink transmission of a UE in a small cell may interfere with the downlink reception of another UE in a neighbouring small cell (uplink-to-downlink inter-cell interference).

As a consequence, dynamic TDD may present a trade-off that needs to be well understood for its proper operation in certain cases, as it may

- improve the efficiency in the usage of the time resources at the MAC layer, but
- degrade the performance of the physical (PHY) layer, introducing inter-cell interlink interference, which in turn, decreases the SINR of the UE, and as a result, the cell performance.

This performance degradation, if it takes place, may be particularly severe for the uplink reception due to strong downlink-to-uplink inter-cell interference. This is because the transmit power and the antenna gain of a small cell BS is usually larger than that of a UE. This calls for the implementation of downlink-to-uplink inter-cell interference mitigation techniques.

The work in [91] presents a system-level simulation analysis on these dynamic TDD trade-offs. Importantly, a new theoretical performance analysis based on the stochastic geometry framework presented in the previous sections of this "*brief history of small cells*" has been developed in [92] to study the dynamic TDD technology. Both MAC and PHY layer aspects are covered, exploring the mentioned trade-offs and the benefits of inter-cell interference cancellation.

1.4 Outline of This Book

In this section and following the description in the previous "*brief history of small cells*," we present the outline of this book, aimed at answering fundamental questions about network densification, and shedding new light on ultra-dense deployments.

The structure of the book is designed to show the fundamental differences between traditional sparse or dense small cell networks and ultra-dense ones, while the content is meant to teach readers the basis of theoretical performance analysis and empower them with the knowledge to develop their own frameworks.

Chapter 2 introduces the main building blocks of any performance analysis tool and describes the basic concepts of the system-level simulation and the theoretical performance analysis frameworks used in this book, paying particular attention to stochastic geometry.

Chapter 3 summarizes the modelling, derivations and results of probably one of the most important works on small cell theoretical performance analysis, that of J. G. Andrews et al. [60], which concluded that the fears of an inter-cell interference overload in small cell networks were not well grounded, and that the network capacity – or in more technical words, the ASE – linearly grows with the number of deployed small cells.

Chapter 4 analyzes in detail – from a theoretical perspective – the first caveat towards such linear growth of capacity in the ultra-dense regime, that of the impact of the transition of a large number of interfering links from NLoS to LoS. This chapter shows that the theoretical tools used to analyze traditional sparse or dense small cell networks do not apply to ultra-dense ones, and neither do their conclusions. The modelling used, the derivations done and the results obtained are carefully presented and discussed in this book chapter for the better understanding of the reader.

Chapter 5 studies in detail – and also in a theoretical manner – yet another and more important caveat towards a satisfactory network performance in the ultra-dense regime, that of the impact of the antenna height difference between the UEs and the small cell BSs. Similarly, as in the previous chapter, the modelling used, the derivations done and the results obtained are carefully presented and discussed in this book chapter for the better understanding of the reader. Moreover, some deployment guidelines are provided to mitigate such fundamental issue.

Chapter 6 brings the attention to an important feature of ultra-dense networks, the surplus of small cell BSs with respect to the UEs. Building on this fact, the ability of next generation small cell BSs to go into idle mode, transmit no signalling meanwhile, and thus mitigate inter-cell interference is presented and shown in this chapter, as a key tool to enhance ultra-dense network performance, and combat the previous presented caveats. Special attention is paid to the modelling and analysis of the idle mode capability at the small cell BS.

Chapter 7 investigates the impact of ultra-dense networks on multi-user diversity. A denser network reduces the number of UEs per small cell in a significant manner, and thus can significantly reduce – and potentially neglect – the gains of channel-dependent scheduling techniques. These performance gain degradations are theoretically analyzed in this book chapter, and the performance of a proportional fair (PF) scheduler is compared to that of an RR one.

Chapter 8, standing on the shoulders of all previous chapters, presents a new capacity scaling law for ultra-dense networks. Interestingly, the signal and the inter-cell interference powers become bounded in the ultra-dense regime due to the antenna height difference between the UEs and the small cell BSs and the finite UE density as well as the idle mode capability at the small cell BSs, respectively. This leads to a constant SINR of the UE, and thus to an asymptotic capacity behaviour in such a regime. From this new capacity scaling law, it can be concluded that, for a given UE density, the network densification should not be abused indefinitely, and instead, should be stopped at a given level. Network densification beyond such point is a waste of both invested money and energy consumption.

Chapter 9, using the new capacity scaling law presented in the previous chapter, explores three relevant network optimization problems: (i) the small cell BS deployment/activation problem, (ii) the network-wide UE admission/scheduling problem and (iii) the spatial spectrum reuse problem. These problems are formally presented, and exemplary solutions provided, with the corresponding discussion on the intuition behind the solutions.

Chapter 10, in contrast to all previous chapters of this book, which focused on the performance of the downlink, analyzes the performance of the uplink of an ultra-dense network. Importantly, this book chapter shows that the phenomena presented in – and the conclusions derived from – all the previous chapters also apply to the uplink, despite its different features, e.g. uplink transmit power control, inter-cell interference source distribution. System-level simulations are used in this book chapter to conduct the study.

Chapter 11 shows the benefits of dynamic TDD with respect to a more static TDD assignment of time resources in an ultra-dense network. As mentioned before, the number of UEs per small cell reduces in a significant manner in a denser network. As a result, a dynamic assignment of time resources to the downlink and the uplink according to the load in each small cell can avoid resource waste, and significantly enhance its capacity. Dynamic TDD is modelled and analyzed through system-level simulations in this book chapter, and its performance carefully explained.

1.5 Definitions

In this section, we provide a number of definitions that may be handy for the reader to provide clarity and avoid confusion on some of the widely used concepts in this book. It is important to note that the definitions presented in the following are provided for guidance and are not exhaustive, and that a more complete definition of some of them – including modelling details – will be given in Chapter 2 of this book.

Subscriber The subscriber is the customer who has a contract with the service provider, in other words, the person named on the bill for the telephone line or internet connection, or the person who owns the subscriber identity module (SIM) card on a pay-as-you-go mobile contract. This may be an individual or an organization.

End-user The end-user is the individual actually using the phone or the internet connection. This will not always be the same person as the subscriber for example, he or she might be the subscriber’s employee, a customer, a family member or a friend.

User equipment A UE is any device used directly by an end-user to communicate. It can be a handheld telephone, a laptop computer equipped with a mobile broadband adapter or any other device.

Base station A BS is a specialized radio transmitter/receiver which connects the UE to a central network hub, the core network, and allows the connection to a network.

Cell site A cell site is the geographical location where the equipment that conforms the BS is deployed.

Cell A cell is the physical geographical area covered by the BS.

Deployment A deployment refers to a collection of two or more BSs which provides access to the network to the UEs in a geographical area, which is generally larger than what a single BS can cover.

Orthogonal deployment Two BSs are said to be deployed in an orthogonal manner when they use a different carrier frequency to communicate. For example, a BS operating at 2 GHz and another one functioning at 3.5 GHz are orthogonally deployed, and do not interfere with each other.

Co-channel deployment Two BSs are said to be deployed in a co-channel manner when they use the same carrier frequency to communicate. For example, two BSs operating at 2 GHz are co-channel deployed, and interfere with each other.

Antenna A rod, wire or other device used to transmit or receive radio signals.

Antenna radiation pattern An antenna radiation pattern is a diagrammatical representation of the distribution of the antenna radiated energy into space, as a function of direction. The most energy is radiated through the main lobe. The other parts of the pattern where the radiation is distributed sidewards are known as side lobes. These are the areas where the power is wasted. There is another lobe, which is exactly opposite to the direction of the main lobe. It is known as the back lobe.

Channel Broadly speaking, the channel is the physical or logical link that connects a data source, e.g. the UE, to a data sink, e.g. the small cell BS. The wireless channel is characterized by a large number of parameters, such as its carrier frequency and bandwidth, to cite a few. Understanding the variations of the wireless channel over frequency and time is of importance to analyze a system performance. Such variations can be roughly categorized in the following two groups, where some of the mentioned concepts are further developed in the posterior definitions [93]:

- Large-scale fading, due to (i) the path loss of the signal as a function of the distance and (ii) the shadowing by large obstructing objects, such as buildings and hills. This happens as the UE moves through distances of the order of such large obstructing objects and is typically frequency independent.
- Small-scale fading, due to the constructive and destructive interference of the multiple signal paths between the transmitter and receiver. This occurs at the spatial scale of the order of the wavelength of the carrier frequency and is frequency dependent.

Path loss The path loss is the attenuation in power density of an electromagnetic wave as it propagates through space. The path loss is influenced by the environment (dense urban, urban or rural, vegetation and foliage), the propagation medium (dry or moist air), the distance between the transmitter and the receiver, the height and the location of the antennas, as well as other phenomena such as refraction, diffraction and reflection.

Shadow fading A radio signal will typically experience obstructions caused by objects in its propagation path, thus generating random fluctuations of the received signal strength at the receiver, referred to as shadow fading. The number, locations, sizes and dielectric properties of the obstructing objects, as well as those of the reflecting surfaces and scattering obstacles are usually not known, or very hard to predict. Due to such unknown variables, statistical models are generally used to model shadow fading, where the received power due to the shadow fading may vary significantly (e.g. in tens of dB) over distances of the order of such obstructing objects, surfaces and obstacles.

Multi-path fast fading The obstructing objects in the propagation path from the transmitter to the receiver may also produce reflected, diffracted and scattered copies of the radio signal, resulting in multi-path components (MPCs). The MPCs may arrive at the receiver attenuated in power, delayed in time and shifted in frequency (and/or phase) with respect to the first and strongest MPC, usually the LoS component, thus adding up constructively or destructively. As a consequence, the received signal strength at the receiver may vary significantly over very small distances of the order of a few wavelengths.

Table 1.1. Author’s research articles on ultra-dense wireless deployments – foundation of this book

Factors	DL	UL	LoS	antH	IMC	PF	dynTDD	Rician	Shadow	Energy	HetNet	D2D	Downtilt	MIMO	UAV	LAA	caching	CoMP	DNA-GA
[94]	√		√	√	√			√											
[95]	√		√	√	√														
[96]	√		√					√											
[97]	√		√		√	√						√							
[98]	√		√	√									√						
[92]	√	√	√		√		√												
[99]	√		√		√						√								
[100]	√		√	√	√	√	√												
[101]	√		√					√	√	√	√								
[102]	√		√	√	√														
[103]	√		√					√	√										
[104]		√	√		√									√					
[105]		√	√	√	√														
[106]	√		√		√						√								
[107]	√		√	√				√											
[108]	√		√		√									√					
[109]	√	√					√												
[89]	√		√		√	√													
[110]	√		√	√															
[88]	√		√		√														
[111]	√	√	√	√	√	√	√	√	√		√		√	√					
[112]	√		√					√	√										
[113]		√	√																
[114]	√		√	√															
[115]	√		√		√														
[116]	√		√								√								
[117]		√	√																
[79]	√		√																
[118]	√		√																
[119]	√					√													

Table 1.1. (Cont)

[120]	√	√	√	√						√	√		√		√				
[121]	√		√	√											√				
[122]	√		√	√	√										√				
[123]		√	√	√											√				
[124]	√		√	√	√										√				
[125]	√		√	√											√				
[126]	√	√			√											√			
[127]	√	√			√											√			
[128]	√	√			√											√			
[129]	√	√			√											√			
[130]	√	√			√											√			
[131]	√		√								√				√		√		
[132]	√		√								√						√		
[133]	√											√					√		
[134]	√				√												√		
[135]	√		√		√												√		
[136]	√											√					√		
[137]		√						√	√				√	√					√
[138]		√						√	√				√	√					√
[139]		√						√	√										√
[140]		√						√	√										√
[8]	√		√		√	√		√	√	√	√		√	√					
[91]	√	√			√		√				√								
[141]	√	√			√		√												
[142]	√	√			√		√												
[143]	√	√			√		√				√								
[144]	√		√															√	
[145]	√		√	√	√			√	√									√	

Noise power The noise power is the measured total noise in a given bandwidth at the antenna of the receiving device when the signal is not present, i.e. the integral of noise spectral density over the bandwidth.

1.6 Related Publications

In this section, and through Table 1.1, we provide a comprehensive list of all the research papers published by the authors of this book on the fundamental understanding of ultra-dense networks. The content of this book builds upon such publications and references therein. The authors advise to read the papers to gain a deeper understanding of some of the concepts presented in this book.