Coverage and Robustness of mm-Wave Urban Cellular Networks: Multi-Frequency HetNets are the 5G Future

Ljiljana Simić, Soumendra Panda, Janne Riihijärvi and Petri Mähönen
Institute for Networked Systems, RWTH Aachen University
Kackertstrasse 9, D-52072 Aachen, Germany
E-mail: lsi@inets.rwth-aachen.de

***Abstract*—Cellular deployments in the spectrum-rich mm-wave bands are anticipated to provide multi-Gbps connectivity in future 5G networks. Although measurements have demonstrated the feasibility of outdoor mm-wave links, the *network-level* opportunities and challenges of mm-wave cellular deployments are yet to be fully understood. In this paper we present a study of mm-wave urban network coverage and robustness - i.e. handover and beamsteering opportunities for supporting user mobility and resolving blockage of directional mm-wave links by moving obstacles. We use real 3D building data with ray tracing and investigate the impact of base station placement and density, building materials, and antenna directivity. Our results show that mm-wave networks can provide multi-Gbps connectivity *locally*, but that over 25% of the network area remains in outage even with dense pico-cellular deployments and that supporting mobility is very challenging. Moreover, primary link blockage can be resolved via handover or beamsteering to a secondary link at only up to 40% and 25% of high-throughput locations, respectively. Given these networking challenges, we argue that *multi-frequency heterogeneous networks* are the key to exploiting the great promise of mm-wave for future 5G deployments, i.e. mm-wave as a hotspot capacity booster, rather than a solution for comprehensive cellular coverage.**

***Index Terms*—Millimeter-wave, 60 GHz, urban outdoor cellular network, coverage, handover, beamsteering, 5G**

I. Introduction

The several GHz of spectrum resources in the millimeter-wave (mm-wave) bands have elicited intense interest in mm-wave networks as a solution to the "capacity crunch" in current microwave cellular bands, meeting the challenges of exponential data traffic growth in future 5G systems [1], [2]. Indeed, the results of early outdoor urban measurements [3]–[7] are promising, demonstrating that multi-Gbps mm-wave *links* are feasible. However, the *network-wide* properties of mm-wave cellular deployments are still poorly understood.

Several analytical studies have provided important insights into coverage trends in mm-wave networks [8]–[11]. However, these theoretical results are based on simplistic propagation models and abstract blockage models, that cannot fully capture the exact complex interactions of mm-wave signals with a real urban environment. Given the critical reliance of mm-wave connectivity on the availability of line-of-sight (LOS) or strong reflected non-LOS (NLOS) paths, the *exact geometry* and *building materials* of a given urban area must be explicitly taken into account to reliably evaluate the achievable coverage of a proposed mm-wave network deployment.

Moreover, these early studies focus almost exclusively on coverage probability, and cannot provide any insight into important networking aspects of directional mm-wave networks, i.e. handover and beamsteering opportunities, which are crucial for supporting (i) user mobility and (ii) resolving blockage of the primary LOS link by moving obstacles common in urban settings. In fact, evaluating this *robustness* of mm-wave network coverage has been altogether largely unaddressed in the existing literature.

In this paper we present a study of mm-wave cellular network coverage and robustness, based on detailed ray tracing simulations using real 3D urban building data from a major business and residential district in central Frankfurt, Germany. We investigate the impact of base station (BS) placement and density, building materials, and antenna directivity. To the best of our knowledge, our work is the first comprehensive system-level study of mm-wave urban cellular deployments in a real environment, investigating in detail both the achievable network performance *and* its robustness with respect to several key system design parameters.

Our results show that mm-wave cellular networks can provide very high throughput connectivity *locally*, but that *continuous* coverage could not be achieved even with dense pico-cellular deployments. However, we also show that strategic BS placement improves network coverage substantially, highlighting the importance of environment-specific site planning for future real-word deployments. Our results demonstrate a number of challenges for mm-wave cellular networks. Firstly, our results reveal that smooth cell-edge handovers would not be possible almost half of the time, posing a very serious problem to supporting mobility. Secondly, although mm-wave networks can provide multi-Gbps connectivity to most non-shadowed regions, they exhibit very poor handover robustness against blockage of those primary LOS links from the serving BS. Thirdly, our detailed network-wide study reveals more restricted beamsteering opportunities than reported by earlier measurement campaigns: the serving BS can resolve link blockage by switching to a secondary NLOS link only for under half of cell-edge receiver locations and for under a

quarter of high-throughput receiver locations over the network.

Importantly, our results thus highlight serious networking challenges that must be addressed before future 5G networks can exploit the immense potential of mm-wave spectrum. In light of these, we argue that mm-wave deployments are best suited to providing high-throughput hotspot data services within a multi-frequency heterogeneous network (HetNet) architecture, where handover to an underlying microwave network tier guarantees comprehensive coverage.

The rest of this paper is organized as follows. In Section II we discuss related work. Section III details our evaluation methodology. In Section IV we present and analyze our results. In Section V we conclude the paper and discuss the practical implications of our results.

II. Related Work

The seminal analytical studies in [8]–[11] provided initial mm-wave cellular network capacity estimates. However they are based on simplistic LOS/NLOS power-law propagation models, with [8], [9] applying abstract probabilistic models of a link being LOS/NLOS/in outage and [10], [11] using tools from stochastic geometry to model building blockage using so-called random shape theory (RST). Consequently, for the sake of analytical tractability these prior works sacrifice accuracy in estimating mm-wave coverage, not being able to capture the real exact interactions of mm-wave signals with the urban environment and predict the existence of e.g. a reflected NLOS link at a *specific* network location. These analytical models are thus also inherently unsuitable for studying other important aspects of mm-wave networks, such as coverage fragmentation and handover and beamsteering opportunities.

We note that we have compared in [12] the mm-wave coverage estimates obtained with our site-specific ray tracing simulations with the analytical models from [10], [11], showing that the RST model is overly optimistic for highly directional networks due to poor estimation of NLOS coverage, which critically depends on the actual urban layout.

We thus emphasize that the sort of detailed simulation-based methodology we adopt is the only approach that is both sufficiently *accurate* (*cf.* analytical studies as in [8]–[11]) and *scalable* (*cf.* measurement studies as in [3]–[7]) to yield reliable network-wide results on mm-wave deployments. We note especially that due to the directional nature of mm-wave links, measurement campaigns are currently prohibitively labour- and time-intensive for obtaining results at a comparably fine-grained spatial resolution over the network. Finally, we note that the existing ray tracing studies of outdoor mm-wave systems focus mainly on mm-wave channel modelling [13], [14] or consider rather elementary coverage evaluation [15].

III. Evaluation Methodology

We give our mm-wave cellular network model in Section III-A and define directional network interference models in Section III-B. Section III-C details how we combined these with ray tracing simulations to obtain our results in Section IV.

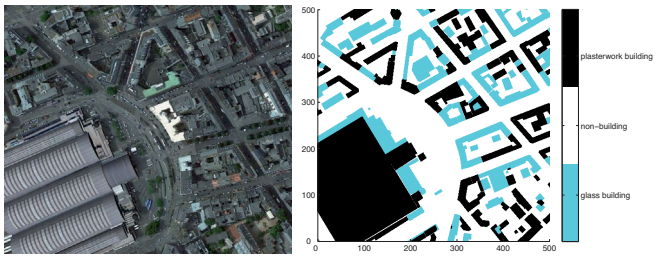

(a) (b)

Fig. 1. Frankfurt 500 m $\times$ 500 m outdoor network study area: (a) map and (b) diagram of the *mixed* material scenario.

A. Network Model & Study Scenarios

Fig. 1a shows our 500 m $\times$ 500 m network study area in central Frankfurt, a major business and residential region in Germany. Given the central railway station in the left bottom corner and surrounding city blocks in Fig. 1a, the mobile data traffic demand in our study area is likely to be high, making it a very good example of a location where mm-wave networks would be deployed in the first instance. We have obtained high-resolution 3D building data for this region, enabling the use of precise site-specific propagation models for computing mm-wave coverage. We note we have also conducted studies for areas of Seoul and New York, obtaining comparable qualitative results, which we omit here for the sake of brevity. BSs are placed over the network area in Fig. 1 with densities of 30, 60, or 90 BS/km^2, with locations of the deployment sites assumed either (i) to be a realization of a Poisson point process (PPP) – i.e. random deployment or (ii) to fall on a uniform grid – i.e. regular deployment.

We consider three distinct cases of building materials to study the impact on the availability of reflected NLOS links. As the boundary cases we consider: (i) the best-case scenario of all *glass* walled buildings, capturing most reflections enabled by the geometry of the urban layout, and (ii) the worst-case scenario of all *plasterwork* walled buildings, enabling practically no reflected paths. The reflection losses for these two surface materials are assumed to be 3 dB and 30 dB, respectively [16]. Thirdly, as an approximation of a more realistic heterogeneous urban environment, in the *mixed* material scenario we randomly assign either glass or plasterwork to individual buildings, as illustrated in Fig. 1b.

We consider two different antenna patterns at the transmitter and receiver: (i) 30° beamwidth and gain of $G_{tx,max} = G_{rx,max} = 15$ dBi in the main lobe, and (ii) 10° beamwidth with $G_{tx,max} = G_{rx,max} = 25$ dBi gain (in both cases we assume $G_{tx,min} = G_{rx,min} = -40$ dBi outside the main lobe). Throughout, we assume each BS operates at the carrier frequency of 60 GHz with a transmit power of $P_{tx} = 15$ dBm (i.e. EIRP within typical spectrum regulatory limits [17]); BS antennas are mounted on the building corner nearest to the deployment site; and the transmitter and receiver antenna heights are set to 6 m and 1.5 m, respectively. The contribution of noise is taken into account via a receiver noise figure of

6 dB and thermal noise over a channel bandwidth of 1 GHz.

B. Modelling Interference in Directional mm-Wave Networks

One of the challenges of evaluating the performance of mm-wave networks with highly directional antennas is that the SINR (signal-to-interference-and-noise ratio) at a given receiver depends not only on the orientations of its receive antenna and the transmit antenna of its serving BS, but also on the current orientations of *all* the other interfering BSs in the network[1]. In this section we define three general *interference models* that can be used (in either analytical or numerical approaches) to study the range of interference conditions that can arise in such directional networks. We first give the bounds on achievable SINR by considering the worst and best-case interference that a receiver can experience. The final model estimates the expected behavior, by considering the average interference for a typical receiver over the ensemble of possible antenna configurations at the interfering BSs.

Let us first formulate a general version of the Friis transmission equation, incorporating the effects of both antenna directionality and the directional propagation effects arising from the environment especially through reflected NLOS paths. Let $G_{rx}(\boldsymbol{\alpha}, \boldsymbol{\theta})$ denote the gain of the receiver antenna in direction $\boldsymbol{\theta}$ when the antenna main lobe is oriented in direction $\boldsymbol{\alpha}$, and define $G_{tx}(\boldsymbol{\beta}, \boldsymbol{\phi})$ for the transmitter antenna analogously. The total received power is then obtained by integrating over all combinations of directions through which the transmitter and the receiver can exchange energy, yielding

$$P_{rx}(\boldsymbol{\alpha}, \boldsymbol{\beta}) = \iint_{S^2 \times S^2} G_{rx}(\boldsymbol{\alpha}, \boldsymbol{\theta}) G_{tx}(\boldsymbol{\beta}, \boldsymbol{\phi}) \frac{P_{tx}}{L(\boldsymbol{\theta}, \boldsymbol{\phi})} \, d\boldsymbol{\theta} \, d\boldsymbol{\phi}, \tag{1}$$

where $L(\boldsymbol{\theta}, \boldsymbol{\phi})$ is the path loss between the transmitter and the receiver for power transmitted in direction $\boldsymbol{\phi}$ and received from direction $\boldsymbol{\theta}$. To simplify our notation, we define the *total gain* between the receiver and the transmitter with respective antenna orientations $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ as

$$T(\boldsymbol{\alpha}, \boldsymbol{\beta}) = \iint_{S^2 \times S^2} \frac{G_{rx}(\boldsymbol{\alpha}, \boldsymbol{\theta}) G_{tx}(\boldsymbol{\beta}, \boldsymbol{\phi})}{L(\boldsymbol{\theta}, \boldsymbol{\phi})} \, d\boldsymbol{\theta} \, d\boldsymbol{\phi}, \tag{2}$$

allowing us to re-express (1) as $P_{rx}(\boldsymbol{\alpha}, \boldsymbol{\beta}) = T(\boldsymbol{\alpha}, \boldsymbol{\beta}) P_{tx}$. Given we consider $\{1, \ldots, J\}$ BSs over the network, we denote the corresponding total gain between the receiver and a given BS j as $T_j(\boldsymbol{\alpha}, \boldsymbol{\beta}_j)$ and denote the transmit power of BS j as $P_{tx,j}$. For simplicity we assume that the receiver is always *optimally beamformed* towards its serving BS i (as in [10], [11]), i.e. they have the antenna orientations

$$(\boldsymbol{\alpha}_{\text{opt}}, \boldsymbol{\beta}_{i,\text{opt}}) = \arg \max_{\boldsymbol{\alpha}, \boldsymbol{\beta}_i} T_i(\boldsymbol{\alpha}, \boldsymbol{\beta}_i). \tag{3}$$

We define our *maximum interference* model by assuming that the interfering (non-serving) BSs also beamform optimally towards the receiver, giving

$$\text{SINR}_{\text{max}} = \frac{P_{tx,i} T_i(\boldsymbol{\alpha}_{\text{opt}}, \boldsymbol{\beta}_{i,\text{opt}})}{N + \sum_{j \neq i} P_{tx,j} \max_{\boldsymbol{\beta}_j} \{T_j(\boldsymbol{\alpha}_{\text{opt}}, \boldsymbol{\beta}_j)\}}, \tag{4}$$

[1]The aggregate mm-wave network interference in practice depends on instantaneous user density and mobility and network-specific policies on e.g. association, handover, scheduling, etc.

where N is the noise power (thermal noise plus receiver noise figure) and the sum in the denominator is over all interfering BSs (i.e. all BSs $\{1, \ldots, J\}$ except the serving BS i).

We define our *minimum interference* model by assuming that all the interfering BSs have their antennas optimally steered away from the receiver, giving

$$\text{SINR}_{\text{min}} = \frac{P_{tx,i} T_i(\boldsymbol{\alpha}_{\text{opt}}, \boldsymbol{\beta}_{i,\text{opt}})}{N + \sum_{j \neq i} P_{tx,j} \min_{\boldsymbol{\beta}_j} \{T_j(\boldsymbol{\alpha}_{\text{opt}}, \boldsymbol{\beta}_j)\}}. \tag{5}$$

Finally, we define our *mean interference* model by averaging the interference power from non-serving BSs, giving

$$\text{SINR}_{\text{mean}} = \frac{P_{tx,i} T_i(\boldsymbol{\alpha}_{\text{opt}}, \boldsymbol{\beta}_{i,\text{opt}})}{N + \mathbb{E}\{\sum_{j \neq i} P_{tx,j} T_j(\boldsymbol{\alpha}_{\text{opt}}, \boldsymbol{\beta}_j)\}}, \tag{6}$$

where the expectation $\mathbb{E}\{\cdot\}$ is taken over all the antenna orientations $\boldsymbol{\beta}_j$ of the interfering BSs, assumed to be uniformly distributed on the set of allowable directions on the sphere.

These three interference models clearly form a partial stochastic order in the sense that for the three SINR distributions we have $\text{SINR}_{\text{max}} \leq \text{SINR}_{\text{mean}} \leq \text{SINR}_{\text{min}}$ in probability.

C. Simulation Methodology

To study the mm-wave cellular network scenarios in Section III-A, we firstly conducted extensive ray tracing simulations using WinProp [18] to obtain accurate, site-specific estimates of mm-wave signal propagation with respect to our 3D urban study area. The resolution of receiver (user) location sampling was set to 5 m. The ray tracing data of the received signal strength (RSS) over the network study area was then post-processed in MATLAB to obtain estimates of SINR by applying the interference models defined in Section III-B. We emphasize that we use ray tracing solely to obtain *average* signal and interference powers at each location, which suffice to determine coverage distributions once fast fading is averaged out (i.e. we do not consider time-domain characteristics of the mm-wave channel). For studying the achievable coverage in Section IV-A, for a given user location the serving BS is defined as the one yielding the highest RSS (in its optimal beamsteering direction), over all BSs in the network (over all their antenna orientations). For studying network robustness in Section IV-B, other BSs and beamsteering directions were considered in turn as potentially serving a user location. We also geometrically distinguished between LOS and NLOS BSs for a given user location, to help study the role of reflections in mm-wave coverage.

1) Ray Tracing Simulations: Conducting ray tracing simulations of mm-wave networks is extremely time consuming due to the large number of possible configurations that must be simulated, and the high computational load that each simulation run entails when using a high-resolution system model. The key challenge arises from the use of directional antennas to overcome the high path loss inherent at mm-wave, resulting in a very large number of possible antenna orientations that need to be explored through dedicated ray tracing runs. Specifically, for each BS in a given network realization, ray tracing is performed for all distinct antenna

orientations, whereby the BS antenna is steered in increments of the beamwidth through all azimuthal angles and selected downtilt angles (i.e. $\{-5°, -35°\}$ for the $30°$ beamwidth and $\{-5°, -15°, -25°, -35°\}$ for the $10°$ beamwidth). This results in 24 distinct orientations in the $30°$ beamwidth case, and 144 orientations in the $10°$ beamwidth case, each of which entails a dedicated ray tracing simulation run *for each BS*. Further, studying the impact of BS deployment patterns and densities, building materials, and antenna patterns on network performance entails a dedicated ray tracing run for each parameter combination. Finally, in the case of PPP distributed BSs, all of these simulations were further repeated over 40 and 100 random network realizations for the $10°$ and $30°$ antenna cases, respectively, to obtain statistically meaningful results.

2) Estimating Interference Models from Ray Tracing Data: Like other currently available commercial propagation and network planning tools, originally designed for analyzing primarily microwave cellular networks, WinProp only supports antenna directionality at the transmitter. Therefore, we finally estimate the user SINR by approximating the interference models in Section III-B from our ray tracing data as follows. This task is simplified by our focus on mm-wave networks, where, in the absence of diffraction effects, NLOS links typically carry a small fraction of power compared to LOS links. First, for a given receiver we find the optimal beamforming configuration with respect to its serving BS i by solving (3) numerically from the ray tracer data via an exhaustive search over all considered BS antenna orientations. (Similarly, we select as the serving BS i the one with the maximum RSS, by searching over *all* BSs and all their antenna orientations.) Second, as per our antenna model in Section III-A, we add to the computed RSS from BS i the main lobe gain $G_{rx,max}$, thus approximating the term $P_{tx,i}T_i(\boldsymbol{\alpha}_{\text{opt}}, \boldsymbol{\beta}_{i,\text{opt}})$. Finally, we compute estimates of the interference term from non-serving BSs in (4)-(6) by maximizing, minimizing, or averaging, respectively, the RSS from each interfering BS j via search over all its orientations $\boldsymbol{\beta}_j$ from the ray tracer data. We then add to this $G_{rx,max}$ if BS j is in the main lobe of the receiver antenna (whose direction is fixed to $\boldsymbol{\alpha}_{\text{opt}}$ as the direction to its serving BS i), and the back lobe gain $G_{rx,min}$ otherwise.

We note that our approximation of (4)-(6) is highly accurate for LOS BSs, but somewhat less accurate for NLOS BSs. This is due to the fact that we are unable to determine from WinProp ray tracing data the direction at the receiver of interference from NLOS BSs, and thus cannot reliably assign receiver antenna gains based on whether the BS location is geometrically in the receiver main lobe. To keep the computational complexity of our simulations feasible, we treat NLOS BSs the same as the LOS case. Nonetheless, our approach is statistically equivalent to assuming that interference power from a NLOS BS is received from a *random* direction uniformly distributed over the admissible receiver antenna orientations, especially considering the resulting SINR distribution over the network is aggregated over many receiver locations and network realizations. Namely, locations of BSs and receivers are approximately independent, implying due to isotropy of the BS locations that the probability of interference from a NLOS BS to be assigned the main lobe gain of the receiver antenna is beamwidth/$360°$, in line with the random interference direction model. Consequently, we expect the approximation error in our SINR distribution results arising from this estimation of NLOS BS interference to be negligible.

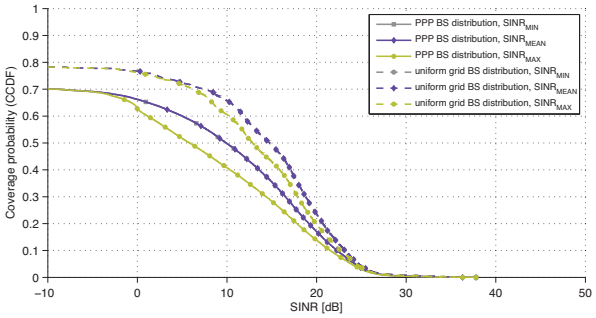

Fig. 2. Coverage probability vs. SINR for the different directional interference models in (4)-(6) (90 BS/km^2, $30°$ beamwidth, glass).

Finally, to briefly illustrate the impact of choosing among the three interference models, in Fig. 2 we show the corresponding SINR distributions obtained from our simulations, for a BS density of 90 BS/km^2, $30°$ antenna beamwidth, and glass building materials (as the spatial likelihood of interference is highest in this case). Fig. 2 shows that the SINR_{min} and $\text{SINR}_{\text{mean}}$ curves are almost identical, indicating that typical interference is negligible; this is also consistent with the fact that the 1 GHz channel bandwidth makes our mm-wave network noise-limited [8], [9]. Fig. 2 also shows that the discrepancy between our SINR_{max} and $\text{SINR}_{\text{mean}}$ models is within 5 dB (for lower BS densities, the difference is within $1-2$ dB). In simulation results presented in the remainder of this paper in Section IV, we adopt the $\text{SINR}_{\text{mean}}$ model in (6) to incorporate the average expected effects of interference in a mm-wave network. Nonetheless, we note that it is in general worthwhile to consider the SINR_{max} model as a bound for potential instantaneous interference. Fig. 2 indicates that the reduction in coverage probability for a given SINR due to worst-case directional interference conditions is up to 5% and 10%, for the uniform grid and PPP BS distributions, respectively (higher in PPP due to clustered neighbour BSs).

IV. Results & Analysis

In this section we present and analyze the results of our study in terms of the achievable coverage (Section IV-A) and network robustness (Section IV-B) of the mm-wave deployments in our urban study area.

A. Achievable Coverage

1) Impact of BS Density: The seminal outdoor mm-wave cellular measurements conducted by Rappaport *et al.* [6] demonstrated that a LOS coverage range of 200 m is feasible; these encouraging early results appear to have been over-interpreted by a part of research community to mean that 5G mm-wave cellular networks will provide coverage with

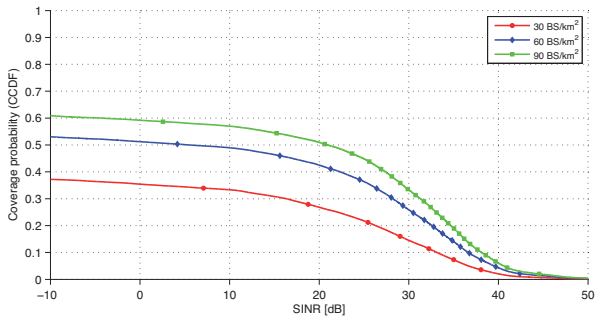

Fig. 3. Coverage probability vs. SINR for different BS densities (mixed materials, 10° beamwidth, PPP BS distribution).

this order of cell radius [1], [2], [19]. In order to provide more accurate and reliable estimates, in Fig. 3 we present the coverage probability vs. SINR in our real urban study area for different BS densities of 30, 60, and 90 BS/km^2; the equivalent average cell radius is around 100, 70, and 55 m, respectively. The results in Fig. 3 are for the case of mixed building materials, a 10° antenna beamwidth, and PPP BS placement. We note that in practice the minimum cell-edge SINR is around -5 dB (e.g. a representative LTE auto-rate function is defined by 3GPP over the range of -6.5 dB to 17 dB [20]). Fig. 3 shows that the cell-edge (-5 dB SINR) coverage probability is only around 48%, 52%, and 60%, for a BS density of 30, 60, and 90 BS/km^2, respectively. Namely, even a pico-cellular BS density of 90 BS/km^2 results in 40% of our study area being in outage. This demonstrates that even an equivalent cell radius which is around a quarter of the often-cited 200 m range is far from sufficient to ensure adequate mm-wave coverage.

Our results thus highlight that the LOS range of several hundred meters demonstrated in early mm-wave measurements does not represent a typical cell radius that can be expected in a normal city environment. This is owing to the fact that, unlike in conventional microwave cellular networks, diffraction does not play a significant role in mm-wave propagation [21], and so the feasible cell size to achieve coverage in a real urban environment is much lower than the mm-wave range in LOS conditions. The key limiting factors are: (i) blockage of LOS paths by impenetrable buildings and (ii) lack of feasible NLOS paths due to geometry of the urban environment and/or insufficiently reflecting building materials. To further examine these factors, in the remainder of this paper, we consider only the highest BS density of 90 BS/km^2, and we turn our attention to considering the impact of BS placement, antenna directionality, and building material on the achievable coverage and robustness of mm-wave networks.

2) Impact of Building Materials: To examine the effect of building materials, in Fig. 5 we present the coverage probability vs. SINR if different materials are assumed for the buildings in our study area (for a 10° antenna beamwidth and PPP BS placement). An example of the underlying ray tracing results of RSS (i.e. P_{rx}) from the strongest LOS or NLOS BS (with the optimal antenna orientation in each case) is illustrated in Fig. 4 for a single PPP network realization. To

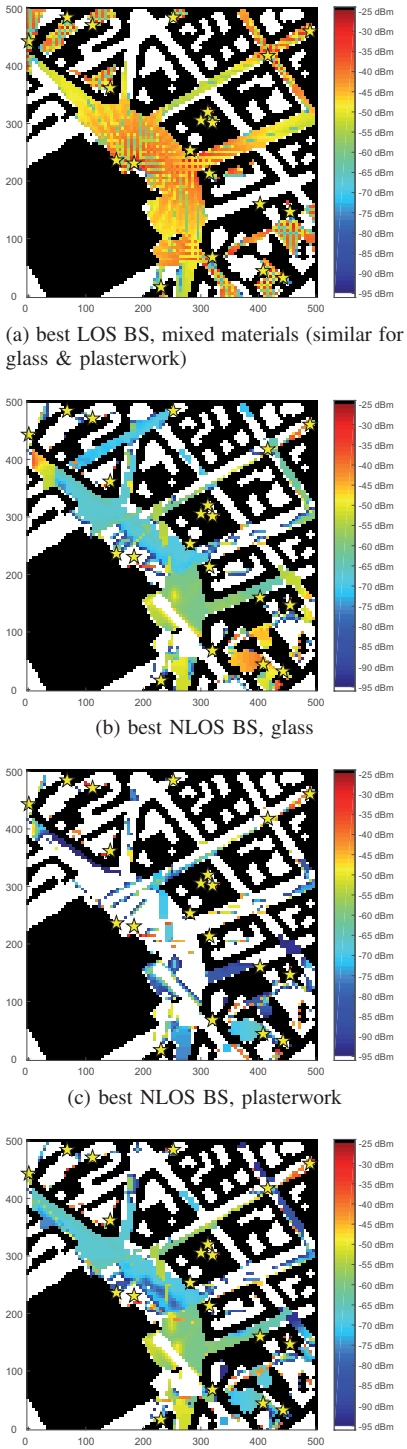

(a) best LOS BS, mixed materials (similar for glass & plasterwork)

(b) best NLOS BS, glass

(c) best NLOS BS, plasterwork

(d) best NLOS BS, mixed materials

Fig. 4. Heatmap of RSS (dBm) for an example PPP BS deployment for different building materials, from the best LOS or NLOS BS (90 BS/km^2, 10° beamwidth). Buildings are shown in black and BS locations as yellow stars; a receiver located in the coloured areas obtains the corresponding RSS from its best LOS/NLOS BS, whereas white indicates areas where RSS is below -95 dBm.

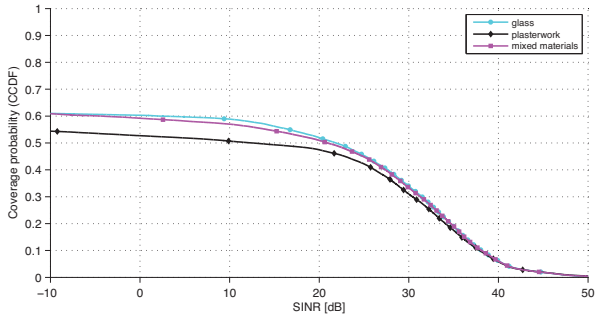

Fig. 5. Coverage probability vs. SINR for different building materials (90 BS/km^2, 10° beamwidth, PPP BS distribution).

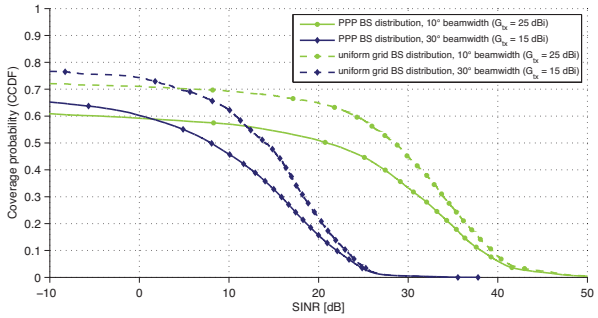

Fig. 6. Coverage probability vs. SINR for different antenna beamwidths and BS distributions (90 BS/km^2, mixed materials).

help relate Figs. 5 and 4, we note that, ignoring interference, an SINR of -5 dB and 10 dB in Fig. 5 corresponds to an RSS in Fig. 4 of -83 dBm and -68 dBm, respectively.

Comparing Figs. 4a and 4b affords an interesting observation: the majority of feasible strong NLOS paths in Fig. 4b occur in the same network areas that are already covered by LOS paths in Fig. 4a. In other words, the contribution of NLOS paths to improving aggregate coverage in the mm-wave network is relatively modest. Indeed, Fig. 5 shows that only an additional 10% of the study area can be covered by NLOS BSs in the highly reflective *glass* environment, compared to the *plasterwork* environment, where virtually no reflected paths are present, as illustrated in Fig. 4c. Importantly, this suggests that the majority of outage areas in a mm-wave network are due to shadowing of LOS paths by buildings, and that NLOS paths provide limited additional coverage in such LOS-shadowed regions. Nonetheless, we emphasize that Fig. 4 also illustrates the arguably more important networking role of NLOS paths: providing secondary NLOS links which can be used when the strongest LOS path becomes blocked by a moving obstacle (such as a pedestrian or car). The corresponding impact of building materials on beamsteering and handover opportunities in the mm-wave network will be further explored in Section IV-B. Finally, we note that Fig. 4 also demonstrates that for our *mixed* material configuration (*cf.* Fig. 1b), there are a similar number of reflected NLOS paths present as for the glass environment, resulting in only marginally lower coverage probability in Fig. 5.

3) Impact of BS Distribution & Antenna Directivity: Having established that coverage in mm-wave networks is very strongly dependent on the specific geometry of the local urban environment, we now consider to what extent a more regular BS placement and wider antenna beamwidth might improve overall coverage. Fig. 6 shows the coverage probability vs. SINR for two different BS deployment strategies, i.e. PPP and uniform grid distribution, and for two different values of antenna directionality, i.e. 10° beamwidth with $G_{tx,max} = G_{rx,max} = 25$ dBi and 30° beamwidth with $G_{tx,max} = G_{rx,max} = 15$ dBi (all assuming mixed building materials). Fig. 7 shows the corresponding heatmap of RSS from the serving BS (i.e. strongest BS with its best antenna orientation) for an example PPP realization of the BS placement (Figs. 7a and 7b) and the uniform grid BS placement (Figs. 7c and 7d), for the 10° and 30° antenna beamwidth cases, respectively.

Firstly, it is evident from Fig. 7 that a uniform grid BS placement improves coverage substantially, irrespective of antenna beamwidth, eliminating many previously LOS-shadowed outage areas. This is consistent with the grid BS placement more closely following the regular street layout of our urban study area. Fig. 6 shows that the uniform grid BS placement improves the cell-edge (-5 dB SINR) coverage by around 10% of the area, and by almost 15% of the area for high SINR coverage at e.g. 10 dB. This is an important result which suggests that simply placing BSs at strategic locations can improve mm-wave network coverage at no extra equipment or site cost. This in turn indicates that careful environment-specific planning of BS sites will be crucial for real-world mm-wave cellular deployments.

Secondly, Fig. 6 reveals an interesting tradeoff between the 10° and 30° antenna beamwidth cases. In the cell-edge, low SINR region, the 30° antenna results in coverage of an additional 3% of the area, owing to its wider beamwidth which suffers less from building shadowing. We again note that although this coverage improvement is rather modest, the practical significance of this result lies in enabling more robust handovers in the mm-wave network, as we will discuss further in Section IV-B. However, due to its associated lower gain, the 30° antenna also achieves significantly lower coverage in the high SINR region compared to the 10° antenna: e.g. around 10% less of the area is covered by an SINR of 10 dB and up to 30% less of the area is covered by an SINR of 17 dB (corresponding to LTE data rates of 2.6 Gbps and 4.3 Gbps [20], respectively). Fig. 7 also clearly illustrates this tradeoff, where the 10° case in Figs. 7a and 7c shows the network having more orange (high data rate) regions but also more white regions in complete outage, compared to the 30° case in Figs. 7b and 7d. This result emphasizes that antenna directionality is also an important design parameter that should be optimized for a given mm-wave network deployment scenario.

Finally, we observe in Fig. 6 that even with our considered best case of a uniform grid BS deployment and the pico-

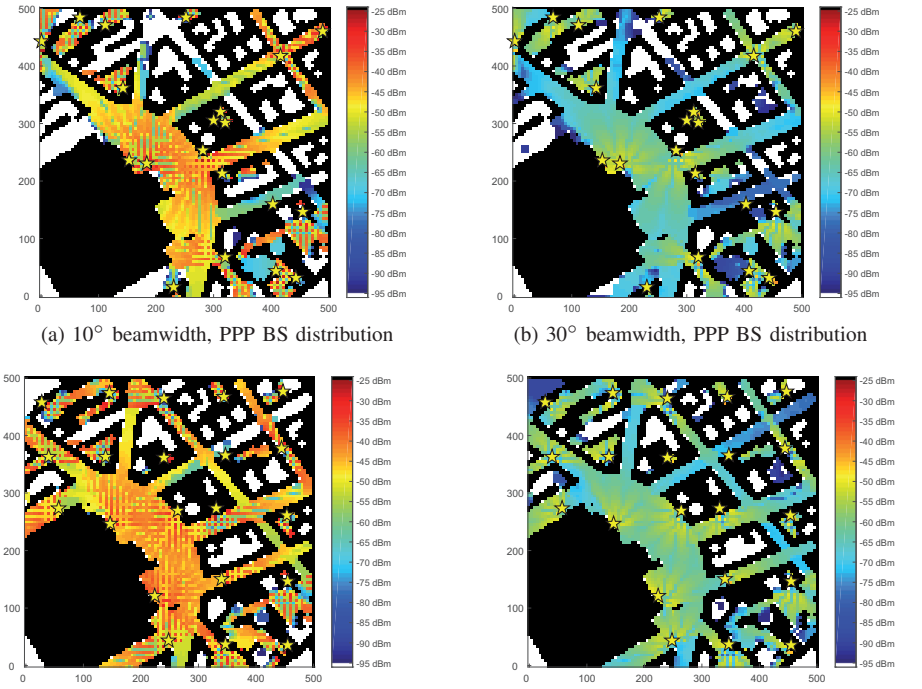

(a) 10° beamwidth, PPP BS distribution

(b) 30° beamwidth, PPP BS distribution

(c) 10° beamwidth, uniform grid BS distribution

(d) 30° beamwidth, uniform grid BS distribution

Fig. 7. Heatmap of RSS (dBm) for the uniform grid BS deployment and an example PPP BS deployment, for different antenna beamwidths (90 BS/km^2, mixed materials). Buildings are shown in black and BS locations as yellow stars; a receiver located in the coloured areas obtains the corresponding RSS from its serving BS, whereas white indicates areas where RSS is below -95 dBm.

cellular BS density of 90 BS/km^2, the maximum cell-edge coverage probability achieved is only 76%. This indicates that a combination of further network densification and more careful BS site planning is unavoidable for achieving *comprehensive* coverage in mm-wave cellular networks. On the other hand, in our best considered case, a user can be served at a very high data rate of 2.6 Gbps (SINR of 10 dB) in 70% of the network area. This is a very important result, which demonstrates that a standalone mm-wave cellular network can easily provide very high throughput connectivity *locally*, but that *continuous* coverage is very difficult (and expensive) to achieve.

B. Network Robustness

In Section IV-A we studied the achievable coverage in our network from the *serving* BS and its *optimal* beamsteering direction. Namely, to compute the coverage probability for different SINR thresholds, we considered for each user location in the network area what is the best achievable mm-wave connectivity, by selecting the combination of the BS and its antenna orientation that provides the highest possible RSS to that user over all BSs (and over all their antenna orientations) in the network. However, due to the demonstrated reliance of mm-wave connectivity on strong LOS links, and the sensitivity of those links to blockage by moving obstacles, this maximum achievable coverage is only one consideration when designing a mm-wave network. In this section, we focus on investigating the *robustness* of the mm-wave coverage in our network, in terms of the handover and beamsteering opportunities for supporting seamless connectivity in the presence of user mobility and moving obstacles (e.g. pedestrians and vehicles).

1) Handover Opportunities: Fig. 8 shows the distribution (taken only over *covered* user locations) of the number of available BSs for serving a given user at a -5 dB and 10 dB SINR threshold, representing cell-edge and high-throughput connectivity, respectively. The two subfigures show the results for the 10° and 30° antenna beamwidth cases, each in turn showing the effect of building material. For the sake of brevity we omit results for the uniform grid BS deployment, as they are qualitatively similar to the PPP case. We emphasize that the existence of multiple BSs that can serve a given user location is important in mm-wave networks for two reasons. First, in the low SINR region (-5 dB), it means the network can perform cell-edge handovers to support user mobility. Second, in the high SINR region (10 dB), it means the network can perform a handover from the serving BS to resolve link blockage of the best LOS path due to a moving obstacle.

Fig. 8b reveals that even in the best case of all-glass buildings, only 60% of covered receiver locations have a second available BS at the -5 dB SINR threshold. This means that even in covered regions, smooth cell-edge handovers would not be possible almost *half* of the time, which poses a very serious problem for supporting user mobility – especially in our mm-wave network where frequent handovers are strictly necessary due to its dense pico-cellular deployment. Fig. 8b also illustrates the importance of reflective building materials for supporting handovers through NLOS links: for the plasterwork scenario, where NLOS links are rare, a handover from the serving to a second available BS can be performed at only

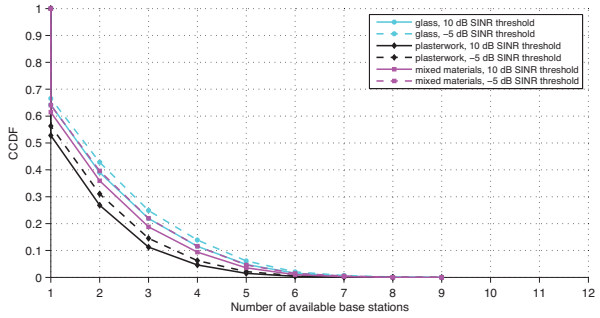

(a) 10° beamwidth

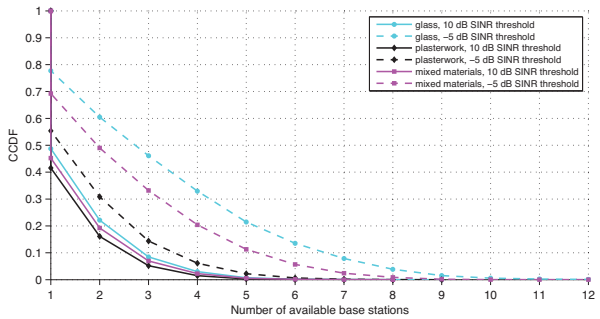

(b) 30° beamwidth

Fig. 8. Network-wide handover opportunity: distribution of the number of available BSs for receiver locations covered by SINR thresholds of -5 dB and 10 dB, for different building materials and antenna beamwidths (90 BS/km^2, PPP BS distribution).

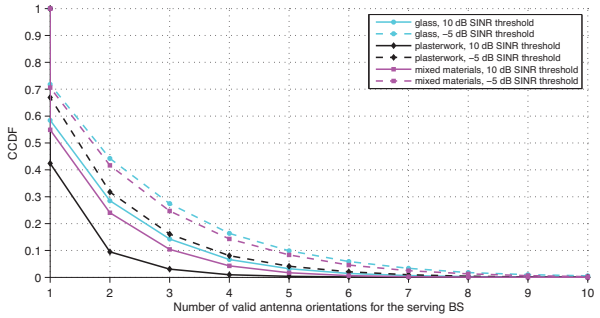

(a) 10° beamwidth

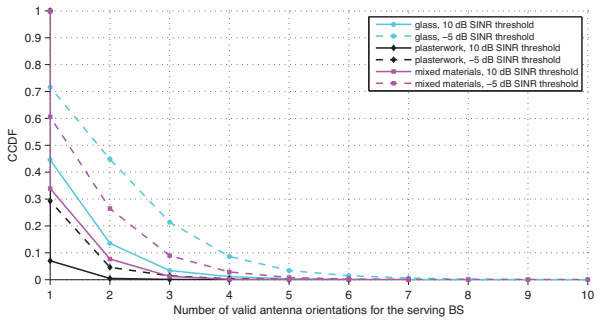

(b) 30° beamwidth

Fig. 9. Beamsteering degrees of freedom: distribution of the number of valid antenna orientations at the serving BS, for receiver locations covered by SINR thresholds of -5 dB and 10 dB, for different building materials and antenna beamwidths (90 BS/km^2, PPP BSs).

30% of cell-edge receiver locations. This effect is also evident, though somewhat less pronounced, for the 10° antenna case in Fig. 8a. Finally, comparing Figs. 8a and 8b for the mixed material scenario demonstrates the advantage of using less directional antennas to support cell-edge handovers: the 30° antenna in Fig. 8b supports handovers for an extra 10% of cell-edge receiver locations compared to the 10° case in Fig. 8a.

Considering the 10 dB SINR threshold curves in Fig. 8a shows that, even in the best case of all-glass buildings, only 40% of high-throughput receiver locations have a second available BS. In Fig. 8b this drops to 20%, due to the insufficient link budget associated with the lower gain of the 30° antennas. Therefore, our results indicate that, although the mm-wave networks can locally provide multi-Gbps connectivity to most covered (non-shadowed) areas (*cf.* Fig. 7), they *exhibit very poor robustness* against blockage by an obstacle of those primary high-throughput links from the serving BS.

2) Beamsteering Degrees of Freedom: In Fig. 9 we consider the possibility of instead resolving primary link blockage by beamsteering to a weaker NLOS link at the serving BS itself. Specifically, Fig. 9 shows the distribution of the number of valid antenna orientations at the serving BS for a user covered by a -5 dB and 10 dB SINR threshold, for the 10° and 30° antenna beamwidth cases and different building materials (for the PPP BS distribution). Fig. 9 reveals that, even in the best case of all-glass buildings, the serving BS has a secondary NLOS link at only 45% of cell-edge (-5 dB SINR) locations.

The availability of back-up NLOS links at the serving BS is even worse for high-throughput receiver locations: for the mixed material scenario in Fig. 9a, the serving BS can switch to a second or third NLOS link at only 25% and 10% of locations covered by the 10 dB SINR, respectively. This is a very interesting result, as it is much less optimistic about resolving blockage of the primary mm-wave link via beamsteering than existing outdoor mm-wave measurement results suggest [7], [22]. For instance, our recent measurements in [7] showed that several NLOS links were found in most of the seven considered receiver locations. However, our measurement locations were selected specifically to study the availability of NLOS paths in an environment where reflections would be likely; we would argue that most other outdoor mm-wave measurements reported thus far are also inherently somewhat biased in this sense. This highlights that, although real measurement data is incredibly important, it should not be taken in isolation, especially as the time-consuming nature of such campaigns means they cannot provide good *fine-grained* spatial statistics over all possible locations in the network area. We thus argue that performing detailed large-scale simulation studies such as the one reported here, to complement real data from measurement campaigns, is crucial for advancing our understanding of mm-wave networking challenges.

Finally, comparing Figs. 9a and 9b shows that the 10° antenna has a slightly higher number of antenna orientations resulting in valid secondary NLOS links than the 30° antenna,

which is consistent with the lower link budget afforded by the 30° antenna. However, we note that the 30° case has a much higher *percentage* of valid beamsteering directions over its 24 distinct antenna orientations (compared to 144 in the 10° case). This suggests that using the 30° antenna would require significantly less precision beamsteering effort and may simplify mm-wave networking. This result is further evidence that antenna directivity is a key parameter in mm-wave networking, which should be carefully optimized, not only for coverage, but also robustness and low control overhead.

V. Conclusions & Discussion

We presented a comprehensive system-level evaluation study of mm-wave cellular network coverage and robustness in a real 3D urban environment, investigating the impact of BS placement and density, building materials, and antenna directivity. Our results show that mm-wave network performance is very strongly dependent on the geometry and materials of the urban environment. Although multi-Gbps connectivity can be provided *locally* to most non-shadowed regions, even dense pico-cellular deployments cannot achieve *continuous* coverage over the network: in the best case, 24% of the network area remains in outage, whereas BS handovers to support user mobility are not possible at 40% of covered cell-edge locations. Moreover, our results reveal that the mm-wave networks exhibit *poor handover and beamsteering robustness* against blockage of the primary link by moving obstacles in the urban environment: the serving BS can maintain connectivity via handover to a second BS or beamsteering to a secondary NLOS link at only up to 40% and 25% of high-throughput locations, respectively.

Therefore, our work highlights several difficult networking challenges to exploiting the great promise of mm-wave spectrum for 5G cellular networks. It demonstrates that very careful, environment-specific network planning and optimization will be absolutely crucial for future real-world deployments, especially if mm-wave 5G systems are seen as a standalone cellular solution. An alternative interpretation of our results is that mm-wave technology should not be seen as a standalone, or even main, component of 5G cellular data access and coverage. Instead, we argue that the highly fragmented nature of mm-wave network coverage revealed by our study strongly indicates that mm-wave deployments should be integrated within a 5G HetNet architecture with spectrum agility. Namely, small-cell mm-wave deployments are ideally suited for providing ultra high speed capacity locally, with smooth handover to an underlying microwave network tier for ensuring comprehensive coverage and mobility management. We note that somewhat surprisingly, whereas a multi-frequency HetNet architecture was also proposed by early LMDS research on sub-60 GHz mm-wave systems [23], in the context of 5G it has been minimally discussed thus far and typically only with respect to integration of 5G mm-wave with legacy 4G/3G microwave deployments [19].

Our ongoing work is focusing on network resource allocation for mm-wave cellular deployments, in particular handover and beamsteering algorithms to support user mobility in the context of multi-frequency HetNets.

References

[1] J. G. Andrews *et al.*, "What will 5G be?" *IEEE J. Select. Areas Commun.*, vol. 32, no. 6, pp. 1065–1082, June 2014.

[2] T. S. Rappaport *et al.*, "Millimeter wave mobile communications for 5G cellular: It will work!" *IEEE Access*, vol. 1, pp. 335–349, May 2013.

[3] T. Rappaport, E. Ben-Dor, J. Murdock, and Y. Qiao, "38 GHz and 60 GHz angle-dependent propagation for cellular & peer-to-peer wireless communications," in *Proc. IEEE ICC*, Ottawa, 2012.

[4] G. MacCartney and T. Rappaport, "73 GHz millimeter wave propagation measurements for outdoor urban mobile and backhaul communications in New York City," in *Proc. IEEE ICC*, Sydney, 2014.

[5] R. J. Weiler *et al.*, "Outdoor millimeter-wave access for heterogeneous networks – path loss and system performance," in *Proc. IEEE PIMRC*, Washington, 2014.

[6] T. Rappaport, G. MacCartney, M. Samimi, and S. Sun, "Wideband millimeter-wave propagation measurements and channel models for future wireless communication system design," *IEEE Trans. Commun.*, vol. 63, no. 9, pp. 3029–3056, Sept. 2015.

[7] L. Simić, N. Perpinias, and M. Petrova, "60 GHz Outdoor Urban Measurement Study of the Feasibility of Multi-Gbps mm-Wave Cellular Networks," in *Proc. mmNet Workshop in IEEE INFOCOM*, San Franciso, 2016.

[8] M. R. Akdeniz *et al.*, "Millimeter wave channel modeling and cellular capacity evaluation," *IEEE Journal on Selected Areas in Communications*, vol. 32, no. 6, pp. 1164–1179, June 2014.

[9] M. N. Kulkarni, S. Singh, and J. G. Andrews, "Coverage and rate trends in dense urban mmwave cellular networks," in *Proc. IEEE GLOBECOM*, Austin, 2014.

[10] T. Bai and R. W. Heath, "Coverage and rate analysis for millimeter-wave cellular networks," *IEEE Trans. Wireless Commun.*, vol. 14, no. 2, pp. 1100–1114, Feb. 2015.

[11] T. Bai, A. Alkhateeb, and R. W. Heath, "Coverage and capacity of millimeter-wave cellular networks," *IEEE Commun. Mag.*, vol. 52, no. 9, pp. 70–77, Sept. 2014.

[12] L. Simić *et al.*, "Comparative study of coverage prediction using random shape theory and ray tracing for mm-wave cellular networks," *to be submitted to IEEE Wireless Commun. Lett.*, 2017.

[13] S. G. Larew, T. A. Thomas, M. Cudak, and A. Ghosh, "Air interface design and ray tracing study for 5G millimeter wave communications," in *Proc. IEEE GLOBECOM Workshops*, Atlanta, 2013.

[14] Z. Zhang, J. Ryu, S. Subramanian, and A. Sampath, "Coverage and channel characteristics of millimeter wave band using ray tracing," in *Proc. IEEE ICC*, London, 2015.

[15] M. Dong *et al.*, "Simulation study on millimeter wave 3D beamforming systems in urban outdoor multi-cell scenarios using 3D ray tracing," in *Proc. IEEE PIMRC*, Hong Kong, 2015.

[16] B. Langen, G. Lober, and W. Herzig, "Reflection and transmission behaviour of building materials at 60 GHz," in *Proc. IEEE PIMRC*, Hague, 1994.

[17] FCC, "In the Matter of Revision of Part 15 of the Commission's Rules Regarding Operation in the 57-64 GHz Band: Report and Order," Tech. Rep. 13-112, 2013.

[18] AWE Communications, *Winprop simulation suite*. [Online]. Available: http://awe-communications.com/

[19] S. Rangan, T. S. Rappaport, and E. Erkip, "Millimeter-wave cellular wireless networks: Potentials and challenges," *Proceedings of the IEEE*, vol. 102, no. 3, pp. 366–385, March 2014.

[20] 3GPP, "LTE; E-UTRA," Tech. Rep. 36 942 V12.0.0, 2014.

[21] A. F. Molisch and F. Tufvesson, "Propagation channel models for next-generation wireless communications systems," *IEICE Transactions on Communications*, vol. E97-B, no. 10, pp. 2022–2034, Oct. 2014.

[22] M. Samimi *et al.*, "28 GHz angle of arrival and angle of departure analysis for outdoor cellular communications using steerable beam antennas in New York City," in *Proc. IEEE VTC Spring*, Dresden, 2013.

[23] P. Mähönen, T. Saarinen, Z. Shelby, and L. Muñoz, "Wireless Internet over LMDS: Architecture and experimental implementation," *IEEE Commun. Mag.*, vol. 39, no. 5, pp. 126–132, May 2001.