Interference Management and Capacity Analysis for mm-Wave Picocells in Urban Canyons

Zhinus Marzi and Upamanyu Madhow, *Fellow, IEEE*

***Abstract*— Millimeter (mm) wave picocellular networks are a promising approach for delivering the 1000-fold capacity increase required to keep up with projected demand for wireless data: the available bandwidth is orders of magnitude larger than that in existing cellular systems, and the small carrier wavelength enables the realization of highly directive antenna arrays in compact form factors, thus drastically increasing spatial reuse. In this paper, we carry out an interference analysis for mm-wave picocells in an urban canyon with a dense deployment of base stations. Each base station sector can serve multiple simultaneous users, which implies that both intra- and inter-cell interference must be managed. We propose a *cross-layer* approach to interference management based on (i) suppressing interference at the physical layer and (ii) managing the residual interference at the medium access control layer. We provide an estimate of network capacity, and establish that 1000-fold increase relative to conventional LTE cellular networks is indeed feasible.**

***Index Terms*— mm-wave picocells, 60 GHz, interference management, cross-layer design, capacity analysis.**

I. Introduction

RECENT years have seen an explosion in cellular data demand due to bandwidth-hungry multimedia applications. This is projected to require a *1000-fold* capacity gain by 2020 [1]. In response to this demand, both industry and academic communities have converged upon the mm-wave frequency band (30-300 GHz) as the next frontier for cellular communication [2], [3]. This is because of two major reasons. First, this frequency band offers an enormous amount of bandwidth compared to existing cellular networks (for example, in the United States, 14 GHz of contiguous unlicensed spectrum is available in the 60 GHz band). Second, the short wavelength at this band ($\leq$ 10 mm) means that electronically large antenna arrays can be made physically small (for example, at 60 GHz, an 8×8 array occupies an area of less than a square inch, while a 32×32 array fits within 10 square inches). This enables a drastic increase in spatial reuse relative to existing systems, via a dense deployment of base stations with small form factor, with each base station capable of forming highly directive links.

There is a growing body of research on the feasibility of mm-wave small cells in terms of link budget and channel modeling [4]–[8]. There is also a recognition that the problem of beam discovery and user tracking is particularly important [9]–[13]: mobile users must be accurately tracked in order to form highly directive beams, and the relative ease of blockage of mm waves implies that an inventory of multiple feasible paths to a given user must be maintained in order to facilitate switching in the event of blockage. Providing an adequate backhaul for mm-wave picocells is another challenge, with mm-wave backhaul (possibly using a band different from that used for the access link) as one viable option [14]–[17]. In short, there are many challenges that must be addressed in order to realize the system concept driving the work reported here. In this paper, however, we sidestep these issues, assuming that such challenges will be eventually surmounted, and focus on estimating the capacity of the resulting system. In order to do so, we must characterize the interference in such a system, and provide sensible interference management strategies that are tailored to the unique characteristics and geometry of the system.

Manuscript received April 30, 2019; revised September 9, 2019; accepted October 10, 2019. Date of publication October 17, 2019; date of current version November 27, 2019. This work was supported in part by the National Science Foundation under Grant CNS-1518812. *(Corresponding author: Zhinus Marzi.)*

The authors are with the Department of Electrical and Computer Engineering, University of California at Santa Barbara, Santa Barbara, CA 93106 USA (e-mail: zhinus_marzi@ucsb.edu; madhow@ece.ucsb.edu).

Color versions of one or more of the figures in this article are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/JSAC.2019.2947819

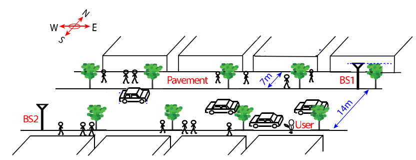

Fig. 1. Picocellular network deployed along an urban canyon. The basestations are mounted on lampposts on the pavement (7 meters from a canyon wall), at a height of 6 meters.

While the system design concepts presented here are of rather general applicability, our numerical results are for a particular setting that we feel has great promise, as also discussed in some of our prior publications [5], [9], [10], [18]. We propose to employ the 60 GHz unlicensed band for base station to mobile communication in outdoor picocells: More specifically, we consider picocellular base stations deployed on lampposts on each side of the street along an urban canyon (e.g. a typical street in New York City), as depicted in Figure 1 (discussed further in Section III). Each base station "face," or sector, could potentially support multiple simultaneous users. We currently assume that this is accomplished by employing multiple subarrays, each capable of RF beamforming to a different user. Alternatively, if and when digital beamforming becomes feasible for large mm-wave arrays, a single array could simultaneously form beams towards multiple users.

Contributions: Prior work at lower carrier frequencies shows that interference becomes a fundamental limiting factor in picocellular settings [19]. However, as we show here, using a capacity analysis taking both inter- and intra-cell interference into account, the narrow beams synthesized using large arrays at 60 GHz alleviate this problem.

In our previous work [18], we characterize *inter-cell* interference, using an analysis accounting for the geometry of the urban canyon. The approach involves studying the interference caused by the main beam and sidelobes separately, since they have distinct characteristics. Section IV provides a summary of the main results of this work that we draw on in our analysis here. While this prior work considers only one subarray per base station face, it extends naturally to the multiple subarray scenario considered here.

The key challenge addressed in this paper is to quantify the gain in spatial reuse by employing multiple subarrays per base station face. The effect of additional inter-cell interference caused by the increase in the aggregate number of transmitters in the system is characterized by adapting our prior analysis in [18]. However, the characterization and management of the *intra-cell interference* originating from the other transmitting subarrays on the same base station are challenging, and constitute the main thrust of this paper. A brief overview of our roadmap to estimate the capacity gain of mm-wave picocellular networks is as follows:

• We first characterize the LoS and NLoS elements of intra-cell interference

• We then propose a *cross-layer* approach to deal with the intra-cell interference, in which we combine techniques from two broad areas that have been studied in the literature: (a) downlink linear precoding and power control [20]–[24] (b) powerful optimization approaches recently developed for network-level resource allocation [25], [26]. Here is a brief description of our two-step method:

1) Given that a resource block is assigned to a given set of active users, we develop a PHY-layer building block which employs an optimal linear method (i.e., Linear Minimum Mean Square Error (LMMSE)) for beamforming, together with power allocation, in order to suppress the LoS intra-cell interference among the active users.
2) We then incorporate the PHY-layer block in designing a MAC-layer protocol which solves an optimization problem to determine the set of active users on each resource block.

• We evaluate our proposed scheme via simulations of picocells along an urban canyon, taking both inter- and intra-cell interference into account. Our cross-layer method enables us to push the limits of spatial reuse, and to estimate the capacity gains provided by cell/antenna densification.

• Finally, we compute the overall capacity per square kilometer for a typical region in Manhattan area, and demonstrate that dense mm-wave picocellular networks can indeed deliver the promised 1000-fold capacity increase over today's cellular networks.

II. Related Work

There are a number of prior papers that investigate the capacity of mm-wave networks in various architectures. Among those [27]–[31] study outdoor cellular network architectures.

The authors of [3], [27] show that spectral efficiency in mm-wave cellular systems can reach that of state-of-the-art LTE systems by employing highly directional antennas. They consider a 1-GHz bandwidth time-division duplex (TDD) mm-wave system which could easily provide a 20-fold increase in average cell throughput in comparison to a 20+20-MHz LTE system. Thus, the capacity gain essentially comes from the increase in bandwidth, and in contrast to the present work, the spectral efficiency improvement due to highly directional antennas is not explored. Moreover, [3] considers hexagonally shaped cells where the base stations are placed randomly, as opposed to our more structured scenario of regularly placed base stations in an urban canyon.

Similarly, [28] conducts system level simulations of the 60 GHz band for capacity evaluation in outdoor scenarios such as college campuses and urban environments. Despite their use of large 20×20 antenna arrays (compared to 8×8 in this paper), their overall capacity estimate is much smaller than ours (400 Gbps/km^2 vs. 2.7 Tbps/km^2 even for our least sparse scenario). This is because [28] does not employ any interference suppression schemes (other than conventional beamforming) or opportunistic resource allocation strategies. They instead apply a round-robin scheme to manage interference, which does not exploit the spatial diversity of users. This prohibits dense deployment of base stations, resulting in capacity saturation at a much lower level compared to ours.

Coverage and attainable data rates in outdoor mm-wave networks are investigated in [29], [30] using stochastic geometry models, with base stations, users and obstacles placed in the 2-D plane according to Poisson point processes, unlike our structured 3D model with regular base station placement. The focus of this work is to show that mm-wave networks can provide coverage comparable with that at lower carrier frequencies while benefiting from larger bandwidth, and the enhanced spatial reuse and interference suppression enabled by large mm-wave arrays is not considered.

There are a few other papers which study mm-wave network capacity in other architectures. For example, the authors of [32] investigate coexistence of device-to-device (D2D) mm-wave links with 4G cellular networks. They establish that the resource sharing optimization problem of this scenario is difficult to solve, and propose a heuristic approach which avoids the LoS interference. This leads to higher aggregate capacity (compared to 4G cellular networks) through a larger number of concurrent transmissions. The authors of [31], conduct extensive simulation for a complicated urban environment in Korea. They investigate a multilevel topology through wireless backhaul link and examine the effects of antenna configuration (arrangement, tilting angle, and spacing) on coverage and capacity.

The present paper differs from the preceding body of work in two main aspects. First, capacity and interference analysis for the urban canyon model (which is well matched to big cities, where there is greatest demand for mobile capacity) and structured placement of base stations has not been considered in prior work, except for our own preliminary results reported in [18]. Second, we explore opportunities to improve spectral efficiency in mm-wave networks, while capacity gains attained in previous works are solely due to the larger bandwidth of mm-wave band. A key distinguishing feature of our work is that we propose and evaluate a cross-layer approach, by utilizing large antenna arrays to suppress interference, and employing novel scheduling approaches to handle residual interference induced in dense deployment of base stations.

The work with the closest perspective to ours is [33], which evaluates capacity for a mm-wave WLAN in a single room, with a 60 GHz access point in the center of the ceiling and users uniformly distributed in the room. A *heuristic static* predefined space time division multiple access (STDMA) algorithm that separates users in either space or time domain is employed. The room is partitioned in order to manage the level of interference the access point causes to other partitions when serving a user in a given partition. The authors then determine which partitions could be served simultaneously while attenuating mutual interference by nullforming. However, this static approach for fixed indoor environments is not directly applicable to more dynamic and complicated scenarios such as mobile users in urban canyons. The achievable spatial multiplexing gain in mm-wave WPAN networks is studied in [34]. An Exclusive Region (ER) is defined for each flow based on a simplified model of the antenna pattern in a 2-D setting, and concurrent transmission is allowed for users which are outside each other's ERs. This approach is difficult to extend to more complex urban canyon geometries, and is conservative compared to our dynamic cross-layer approach.

III. System Model

In this paper, we consider street canyons where base stations are placed in a zig-zag pattern, such that immediate neighbors are on opposite sides of the street. Each base station has two sets of antenna arrays placed on opposite faces, aligned such that one set faces east and the other faces west.

Figure 1 depicts a canyon segment between two neighboring base stations BS1 and BS2, separated by distance d. We term such a canyon segment a *picocell* of width d. Each user in the picocell could be served by either an eastward-facing antenna of BS2 or a westward-facing antenna of BS1. Thus, each picocell is covered by two sets of arrays, each belonging to a different BS.

In this paper, we have adopted a *spatially sparse* channel model motivated by comprehensive experimental results presented in [35]. Studies have shown that more recent channel models such as NYUSIM [36], which suggest sparser channels compared to legacy 3GPP models, are more realistic and are a better match to actual experimental results [37]. Sparse mm-wave channels can be accurately estimated by efficient algorithms proposed in the literature [9]. We assume that channel knowledge is available at both the base station and mobile users. Consider a base station bearing K antenna arrays on each face. Note that link distances are large enough that all transmitters installed on a face could be approximated as co-located from the users' point of view. Under this abstraction, our results apply to hybrid and digital beamforming as well, since we assume the same location and orientation for each subarray. Therefore, the channel matrix from any of these K transmitters on each face to the q-th user is the same and denoted by $\mathbf{H}_q$. The channel matrix $\mathbf{H}_q$ is of size $M \times N$ where M is the antenna size of the mobile user and N that of the transmitter and is characterized by the path loss and spatial frequencies between any of the K transmitters and the q-th mobile user. We assume $\mathbf{H}_q$ is known to all K transmitters as well as to the q-th mobile user. The parameters of the system model are listed in Table I.

TABLE I
System Model Parameters

Parameter	Definition	Default
B	Total bandwidth	2GHz
K	Num. of antenna arrays	-
Q	Num. of users	4
F	Frequency reuse factor	2
d	Dist. between adjacent BSs	50
N	Num. of ant. elements in Tx array	64 (8×8)
M	Num. of ant. elements in Rx array	16 (4×4)
d_e	Hor./Vert. dist. of ant. elements	$\lambda/2$=2.5mm
-	Street width	14m
-	Pavement width	7m
-	User height	$\sim\mathcal{U}$[1.5m,2m]
-	Num. of canyons in 1km^2	15

IV. Background on Inter-Cell Interference

While our focus in this paper is on intra-cell interference, our overall capacity analysis must, of course, account for inter-cell interference as well. For the latter, we employ the analysis in our previous work [18], which we briefly review in this section.

We define inter-cell interference as the interference induced by the transmitters on other basestations. We make two simplifying assumptions: (a) we ignore interference across parallel urban canyons, as well as interference which might leak from cross streets; (b) we do not consider potential reflections from horizontal ledges. While more detailed modeling is needed to refine the interference and capacity estimates provided here to account for such effects, we expect the qualitative conclusions to remain unchanged.

We investigate the inter-cell interference caused by the main lobe and side lobes separately, for they have different characteristics. Since we consider a large number of antenna elements, the main beam is narrow and is well modeled by a single ray. Side lobes are much weaker, but their directions are difficult to predict, hence we must be more careful in bounding their effect. In the following subsections, we elaborate on this by reviewing two theorems proved in [18].

A. Main Lobe Interference

We consider transmitters with a large number of elements forming a pencil beam towards the desired user. This "desired" beam can be along the LoS, or it can be a single bounce from a wall or the ground (e.g., when steering around an obstacle blocking the LoS). Given the highly directive nature of the beam and the limited diffraction at small wavelengths [38] we can use ray tracing to understand the interference that such a beam creates for neighboring basestations.

In [18], we show that the main beam escapes skyward after a few bounces (Figure 2), assuming that we can ignore the effect

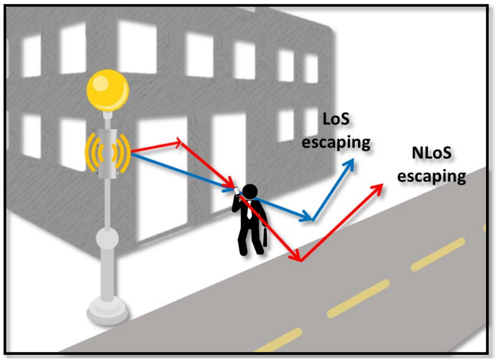

Fig. 2. Main lobe will escape skyward after a few bounces.

of potential reflections from horizontal ledges. Specifically, Theorem 1 bounds the number of neighboring cells that are affected by the main beam's interference, assuming that each face only creates interference in the direction it is facing.

Theorem 1: The maximum range over which the main beam can create interference is bounded by $\frac{H_{BS}+h_{max}}{H_{BS}-h_{max}}d$*. Thus, the main beam from a face creates interference for at most* $N_{max} = \lceil\frac{H_{BS}+h_{max}}{H_{BS}-h_{max}}\rceil$ *adjacent BSs in the direction it is facing. We denote by* h_{max} *the maximum height of users, by* H_{BS} *the height of a basestation, and by* d *the width of a picocell shared among two opposite facing antennas on adjacent basestations.*

For typical values of $H_{BS} = 6m$ and $h_{max} = 2m$ employed in our simulations, Theorem 1 implies that the main beam interferes with two adjacent basestations in the direction of the face producing the beam. We emphasize that this theorem determines the *presence* of inter-cell interference due to the main lobe using geometric modeling and ray tracing. It does not attempt to quantify the strength of the interference, and thus does not need to rely on a signal attenuation model.

B. Sidelobe Interference

While the main beam points towards a user inside the picocell, the emission direction of sidelobes is highly variable, hence it is not possible to limit side lobe interference to a finite number of adjacent picocells. However, as shown in [18] the cumulative sidelobe interference seen within a given picocell is bounded to a relatively small value. This is because the strength of the interference from a distant picocell decays geometrically with distance due to oxygen absorption and reflection losses, along with the quadratic decay due to path loss.

Specifically, for a user served by BS_0, Theorem 2 quantifies the interference from basestations [c, ∞) and (-∞,-c] (c $\geq$ 0).

Denote by P the smallest received power over the desired link, which is given by

$$P = P_{Tx}G_{Tx}G_{Rx}(\frac{\lambda}{4\pi L_{max}})^2 e^{-\beta L_{max}} \quad (1)$$

where P_{Tx}, G_{Tx} and G_{Rx} are the transmitter power and the gains of Tx and Rx antenna arrays, respectively. The parameters λ, β and L_{max} denote, respectively, the wavelength, oxygen absorption coefficient (16 dB/km) and maximum length of a link inside a picocell.

Theorem 2: For a user in cell 0, the sidelobe interference due to the BSs $[c,\infty)$ *and* $(-\infty, c]$ *is bounded by* $\alpha_c P$*, where* P *is the smallest received power over the desired link.*

$$\alpha_c = \frac{\sum_{n=c}^{\infty} I_n + \sum_{n=-\infty}^{-c} I_n}{P} \quad (2)$$

where α_c *decays geometrically with* c*. Here* I_n *denotes the sidelobe interference originated from the* n*th base station.*

In brief, by Theorem 1, main beam interference is induced by $\lceil\frac{H_{BS}+h_{max}}{H_{BS}-h_{max}}\rceil$ adjacent BSs and if we wish to avoid it for $H_{BS} = 6m$ and $h_{max} = 2m$, every 3 adjacent BSs have to coordinate.

One possibility, which turns out to be wasteful, is to orthogonalize transmissions among such sets of 3 basestations (i.e., with a frequency reuse of 3). Moreover, from the computations associated with Theorem 2 shown in [18], the cumulative interference caused by sidelobes from basestations beyond this set (c $\geq$ 3) is at least 40dB weaker than the desired received power. Thus, a frequency reuse of 3 leads to very large SINR, which (a) can lead to only logarithmic gains in capacity, which do not compensate for the up-front loss of 1/3 of the degrees of freedom, on which the capacity depends linearly; (b) hardware constraints limit the size of the constellation, making it difficult to fully utilize even the logarithmic gains due to enhanced SINR. For example, if we limit the highest spectral efficiency supported by our coded modulation strategy to $s_M = 6$ bps/Hz (e.g., uncoded 64-QAM or an even larger constellation with nontrivial channel coding), the numerical results in [18] show that it is this limit, rather than the inter-cell interference, which determines network capacity.

Thus, given the interference reduction due to narrow beams, orthogonalization is wasteful, and much larger network capacity can be obtained by lightly coordinating among neighboring base stations while keeping full spatial reuse, which is the regime suggested in [18].

V. Intra-Cell Interference

In addition to cell densification, one can attain further spatial reuse *within* the cell by increasing the number of subarrays on each base station (or equivalently, employing digital beamforming to support multiple users with a single array). However, this benefit comes with the pitfall of *intra-cell interference,* where a transmitter interferes with receivers in the *same* cell that it does not target. If not carefully managed, this could significantly reduce the spectral efficiency of spatially correlated users.

In this section, we consider K subarrays placed on each face of a basestation (Fig. 3). We first characterize intra-cell interference in our system model and then propose a *cross-layer* approach to deal with it. To this end, we employ a two-step method briefly described below:

1) Given that a resource block is assigned to a pre-defined set of users, we develop a building block at the PHY-layer, which employs LMMSE beamforming

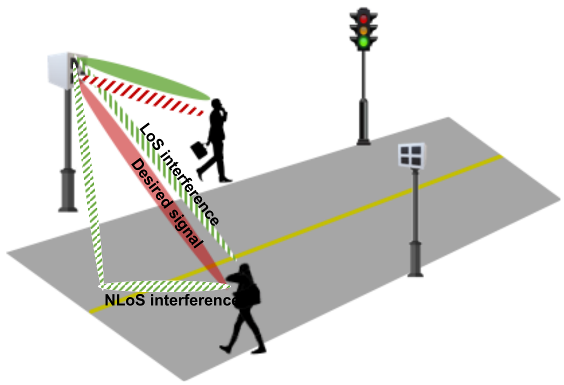

Fig. 3. Multiple subarrays placed on each face of a basestation which leads to intra-cell interference.

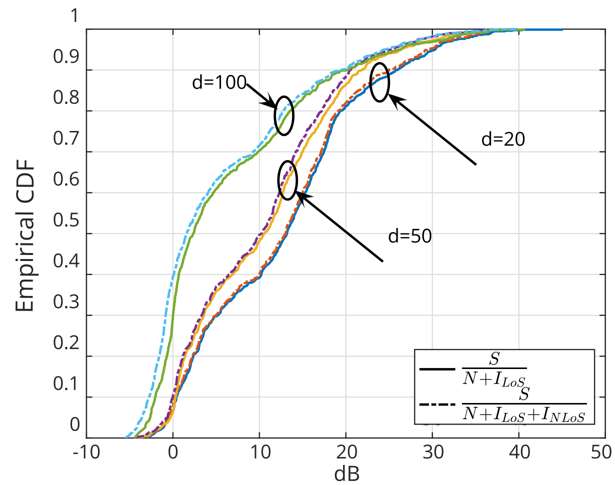

Fig. 4. CDF of signal-to-intracellular interference shows that NLoS interference can be neglected. Solid lines correspond to ignoring NLoS interference, while dashed lines include NLoS interference.

(which is SINR-optimal among linear interference suppression techniques), and associated optimal power allocation to suppress and manage LoS intra-cell interference.

2) We then incorporate the PHY-layer block in designing the MAC-layer protocol, which determines the set of active users on each of the resource blocks.

We evaluate our proposed scheme via comprehensive simulations of picocells along an urban canyon in which both inter- and intra-cell interference are taken into account. Our simulation results demonstrate that, as we shrink cells (down to the cell width of 20m), users' spectral efficiency is mostly ($\geq$ 97%) limited by the hardware limitations. Comparison with our previous results with a single subarray per face [18] indicates that we are able to increase the capacity by a factor of K (at least for small number of subarrays per face, e.g., $K = 2$) in small cells. Larger picocells are more prone to interference and do not benefit as much from multiple subarrays. However, even here our proposed scheme provides users with sufficient spectral efficiency to attain large network capacity gain. Lastly, we compute the overall capacity per square kilometer for a typical region in Manhattan area and demonstrate that dense mm-wave picocellular networks can actually deliver the promised 1000-fold capacity increase over the conventional LTE networks.

A. Intra-Cell Interference Characterization

As with inter-cell interference, intra-cell interference is composed of LoS and NLoS components (depicted in Fig. 3). However, with our assumption that users are served through the LoS path, LoS interference component is expected to be dominant, for the following reasons:

1) The receiver's main lobe is unlikely to encompass the NLoS components of interference. The LoS component, in contrast, gets *amplified* by the same amount as the desired signal.
2) NLoS components are subject to higher path loss.
3) NLoS components suffer from reflection loss induced by reflecting surfaces.

Our simulation results for the same urban canyon scenario also validate this assumption (depicted in Fig. 4)

We therefore assume that intra-cell interference can be sufficiently alleviated by suppressing the LoS component alone. For the rest of this section, by the term interference, we refer to the LoS component of intra-cell interference.

B. PHY Layer Design: Power Allocation and Beamforming

Mitigation of co-channel interference in multiuser MIMO has been extensively studied in the literature [20]–[24]. Different approaches such as precoding, transmitter or/and receiver beamforming, and power adaptation have been explored. In this section, we restrict ourselves to RF beamforming and power control.

In the context of power control and beamforming, there are two classical optimization problems: (a) sum-rate maximization and (b) minimum-rate maximization, subject to the power constraint(s). The former is often studied in the context of information-theoretic capacity and does not guarantee fair sharing of resources among users. We therefore focus on the latter, which guarantees a minimum level of QoS (Quality of Service) for each of the streams.

Minimum-rate optimization can be translated to the following problem:

$$\mathcal{S}(P_T) = \begin{cases} \max_{\{\boldsymbol{\omega}_1,\boldsymbol{\omega}_2,\cdots,\boldsymbol{\omega}_K\}} & \min_i \text{SINR}_i \\ \text{s.t.} & \sum_{k=1}^{K} \|\boldsymbol{\omega}_k\|_2^2 \leq P_T \end{cases} \tag{3}$$

where $\boldsymbol{\omega}_k \in \mathbb{C}^N$ is the transmit beamforming vector aimed at the k-th user, $\|\boldsymbol{\omega}_k\|_2^2$ is the power consumed by the k-th subarray, and SINR_k is the signal to interference ratio at k-th receiver

$$\text{SINR}_k = \frac{\left|\boldsymbol{\omega}_k^H \mathbf{h_k}\right|^2}{\sum_{\substack{i=1 \\ i\neq k}}^{K} \left|\boldsymbol{\omega}_i^H \mathbf{h_k}\right|^2 + \sigma_k^2}$$

A straightforward argument shows that (3) results in the same SINR for all the users, hence the maximum index of fairness is guaranteed.

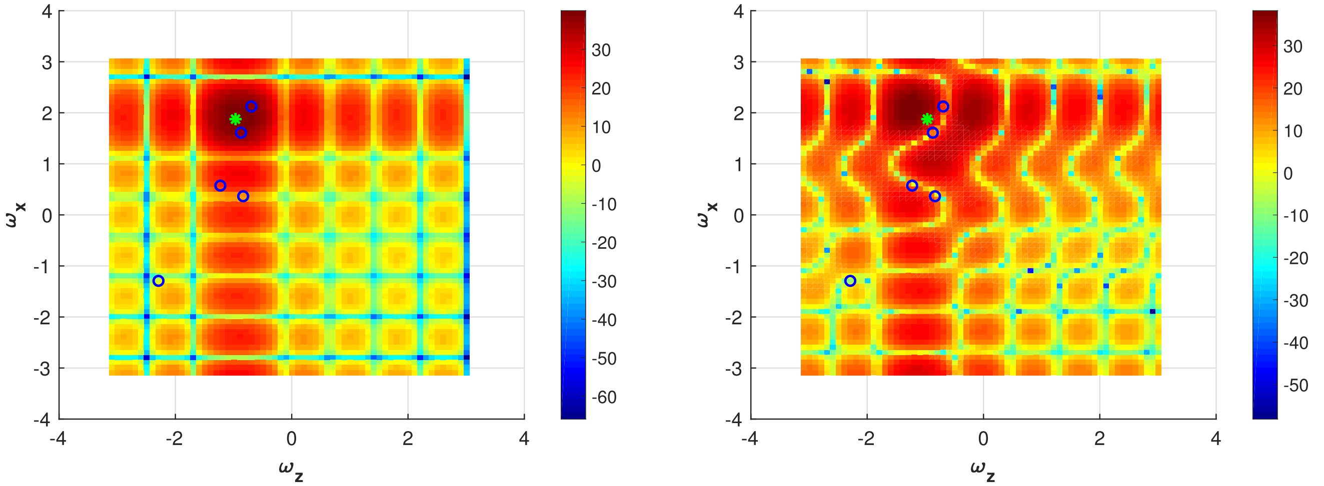

Fig. 5. Transmit antenna patterns causing intra-cell interference (left) and the new antenna patterns after employing interference suppression via Algorithm 1 (right). The spatial frequency of the target user is marked by a green star, and the remaining users are marked by blue circles. Employing Algorithm 1 (right) aligns the null directions with the non-targeted users.

Our solution to problem (3) builds on previous work in [20], [21]. We start with the related power optimization problem

$$\mathcal{P}(\gamma) = \begin{cases} \min_{\{\boldsymbol{\omega}_1, \boldsymbol{\omega}_2, \cdots, \boldsymbol{\omega}_K\}} \sum_{k=1}^{K} \|\boldsymbol{\omega}_k\|_2^2 \\ \text{s.t.} \qquad \min_i \text{SINR}_i \geq \gamma \end{cases} \tag{4}$$

It was shown in [21] that (3) and (4) are inverse problems, meaning that $\mathcal{S}(\mathcal{P}(\gamma_0)) = \gamma_0$ and $\mathcal{P}(\mathcal{S}(P_T)) = P_T$. Furthermore, (4) has an iterative solution [20]. We leverage these observations to formulate Algorithm 1, which iteratively solves (4) for increasing values of γ until the power constraint in (3) is saturated. The solution to (4) employs LMMSE to estimate the transmit beamforming vector (lines 8-15 in Algorithm 1), followed by power allocation to enforce the minimum SINR constraints (line 16 in Algorithm 1).

Figure 5 illustrate how the algorithm distorts the transmitter antenna pattern by pushing nulls toward the users that the transmitter does not target. This improves SINR but might cause SNR degradation due to sidelobe enhancement.

Algorithm 1 PHY Layer Design

1: **Input:** $\{p_i^0, \mathbf{h}_i\}$ $\forall i$, γ, $\Delta\gamma$
2: **Output:** $\{\boldsymbol{\omega}_i\}$ $\forall i$, γ
3: **procedure** BEAMFORMING AND POWER ADAPTATION
4: Compute normalized channels: $\tilde{h}_k = h_k/\sigma_k^2$ $\forall k$
5: **while** $G_{max}\|\boldsymbol{\omega}_k\|_2^2 \leq EIRP$, $\forall k$ **do**
6: $\gamma = \gamma + \Delta\gamma$
7: $n \leftarrow 0$
8: **repeat**
9: **for** $k \in \{1, 2, \cdots, K\}$ **do**
10: $\hat{\boldsymbol{\omega}}_k^n = \text{argmin}_{\boldsymbol{\omega}_k} \sum_{j=1, j\neq k}^{K} p_j^n |\boldsymbol{\omega}_k^H \tilde{h}_j|^2 + \|\boldsymbol{\omega}_k\|_2^2$, s.t. $\boldsymbol{\omega}_k^H \tilde{h}_k = 1$
11: $p_k^{n+1} = \gamma \sum_{j=1, j\neq k}^{K} p_j^n |(\hat{\boldsymbol{\omega}}_k^n)^H \tilde{h}_j|^2 + \gamma \|\hat{\boldsymbol{\omega}}_k^n\|_2^2$
12: $\tilde{p}_k^{n+1} = \gamma \sum_{j=1, j\neq k}^{K} \tilde{p}_j^n |(\hat{\boldsymbol{\omega}}_j^n)^H \tilde{h}_k|^2 + \gamma$
13: **end for**
14: $n \leftarrow n + 1$
15: **until** convergence
16: $\boldsymbol{\omega}_k = \sqrt{\tilde{p}_k}\hat{\boldsymbol{\omega}}_k$ $\forall k$
17: **end while**
18: **end procedure**

Remarks on the Algorithm

- Intuitively, the goal of the optimization problems in (3) and (4) is to manipulate the transmitter's antenna pattern to minimize the induced interference toward the non-targeted users while maintaining constant gain along the desired direction. In addition, power adaptation is employed to cope with variations in link distance.
- In practice, we have individual power constraints on the Equivalent Isotropically Radiated Power (EIRP), which impose the following constraint:

$$G_{max}\|\boldsymbol{\omega}_k\|_2^2 \leq EIRP \quad \forall k \in \{1, 2, \cdots, K\}$$

where G_{max} is the maximum array gain provided by the antenna and EIRP is the limit established by FCC (Federal Communications Commission) for different frequencies (e.g., EIRP=40 dBm at 60 GHz). Our iterative solution allows us to impose the individual power constraints by setting the stopping criteria as when any of the transmit powers has reached the threshold (line 5 in Algorithm 1).
- We have omitted the effect of the receiver antenna array in our formulation. Specifically, the channel matrix $\mathbf{H}_k$ has been replaced by a vector $\mathbf{h}_k$. This is for two reasons:
  1) In order to limit the complexity of mobile receivers, interference suppression is employed at the base station alone.
  2) For intra-cell interference, the receiver antenna provides an array gain of M for both signal and interference. Thus, it does not affect performance in an interference-limited scenario.

However, we take the receiver arrays back into account for our simulation results, given their role in enhancing signal-to-noise ratio (Section VI).

C. MAC Layer Design: Resource Allocation

The preceding PHY layer optimization is for sharing a single resource block among a pre-defined set of users. In this section, we consider interference management in the MAC layer, where resources are divided into blocks (resource granularity) and only certain users allowed to operate in each block (user selection). Intuitively, these additional degrees of freedom can be exploited in the following manner: by selecting spatially separated users to operate in the same block, we can mitigate interference and increase spectral efficiency.

Preliminaries

Consider a cell with Q users sharing frequency band B over a frame of duration T. We make two assumptions:

1) The frame duration T is small enough that mobile users can be considered to be quasi-stationary over a frame. For example, for a frame length of 1 ms as proposed in [39], a pedestrian speed of 5 km/h and a vehicular speed of 90 km/h would result in displacements of only 1.4 mm and 2.5 cm, respectively, during a frame. Even for the narrow beams considered here, it is easy to see that the spatial channel is well modeled as stationary across multiple such frames for typical ranges.
2) The directive antenna arrays employed on both transmitter and receiver suppress multipath fading sufficiently that we may approximate the channel as frequency non-selective.

We consider resource allocation via time division, so that at every point in time each active user utilizes the entire bandwidth B. For simplicity we allow an infinite time granularity.

We need to allocate each time slot (small portions of a frame) to a subset of users. Denoting by $\mathcal{Q}$ the set of all users, we define $\mathcal{P}_{\leq K}(\mathcal{Q})$ as the set of all possible subsets of users (*configurations*) that can be served simultaneously by (up to) K antenna arrays:

$$\mathcal{P}_{\leq K}(\mathcal{Q}) = \{U_c \subset \mathcal{Q} \mid |U_c| \leq K\}$$

We wish to find the fraction of a frame that should be allocated to each of these configurations in order to maximize sum (or minimum) spectral efficiency. More specifically, let x_c represent the portion of the time frame allocated to the c-th configuration. We want to find policy $\mathbf{x} = [x_1, x_2, \cdots, x_C]^T$ where $C = \sum_{k=0}^{K} \binom{Q}{k}$ is the cardinality of $\mathcal{P}_{\leq K}(\mathcal{Q})$.

The spectral efficiency for the q-th user under policy $\mathbf{x}$ is then defined as

$$r_q = \sum_{c=1}^{C} x_c \ \log(1 + \gamma_c^q) \quad \text{(bits/sec/Hz)} \tag{5}$$

where γ_c^q is the SINR of the q-th user under c-th configuration (where U_c is the set of active users.) Clearly, we set $\gamma_c^q = 0$, for $q \notin U_c$).

The resource allocation problem: Like the optimization problems for beamforming and power adaptation, the resource allocation problem could also be formulated to maximize either the sum-rate or the min-rate. In order to provide fairness among users, we focus on the min-rate version, which can be formulated as follows:

$$\begin{aligned} &\max_{\mathbf{x}} \min_{q} \ r_q \\ &\text{s.t. } S^T \mathbf{x} = \mathbf{r} \\ &\quad \mathbb{1}^T \mathbf{x} = 1 \\ &\quad \mathbf{x} \succeq \mathbb{0} \end{aligned} \tag{6}$$

In the first constraint, we have rewritten the equations in (5) in a matrix form by defining $S_{C \times Q} = [s_{cq}]$ where $s_{cq} = \log_2(1 + \gamma_c^q)$ is the spectral efficiency of the q-th user under the c-th configuration and $\mathbf{r} = [r_1, r_2, \cdots, r_Q]^T$ is the vector of resultant spectral efficiency over a unit time frame. The last two conditions ensure that the sum of the portions allocated to different configurations adds up to one and neither of them can be negative.

In theory, allocation policies resulting from (6) should maximize the min-rate among users. However, in practice we might not be able to attain the theoretical rate due to hardware constraints on constellation size. If s_M is the hardware-constrained spectral efficiency limit, the maximum min-rate will be bounded by $(K/Q)\, s_M$. This corresponds to the *saturation* point where all transmitters operate at their highest modulation rate, s_M.

Figure 6 shows the empirical CCDF of maximum min-rate for different cell sizes, along with the saturation point imposed by the various modulations (i.e., $(K/Q)\, s_M$). As depicted in Figure 6, for smaller picocells with larger number of users ($d \leq$ 20m and $Q \geq K$) spectral efficiency is limited to the saturation point imposed by 64-QAM modulation ($s_M = 6$ bps/Hz) and hence constrained by hardware rather than noise or interference. This is because smaller cells have (almost) vertically aligned beams which will lead to more diverse spatial frequencies as compared to less slanted beams at larger cells. As a result, our interference suppression algorithm performs more effectively in smaller cells. Furthermore, a larger number of users could increase the attainable spectral efficiency by enabling us to utilize multiuser diversity for avoiding interference.

Remarks

- The optimization problem in (6) maximizes the worst users' spectral efficiency and therefore will result in equal rate for all users in $\mathcal{Q}$. Its performance is therefore inherently bounded by that of the worst user. However, there are certain scenarios where we can maximize the sum-rate as well; for example, when we have surplus resources after providing all users with some minimum required spectral efficiency, r_{min}.

  Therefore, if the resultant min-rate provided by the allocation policy in (6) is greater than r_{min}, we employ the following optimization problem to maximize the sum-rate by utilizing multiuser diversity.

$$\begin{aligned} &\max_{\mathbf{x}} \mathbb{1}^T S^T \mathbf{x} \\ &\text{s.t. } S^T \mathbf{x} \succeq r_{min} \mathbb{1} \\ &\quad \mathbb{1}^T \mathbf{x} = 1 \\ &\quad \mathbf{x} \succeq \mathbb{0} \end{aligned} \tag{7}$$

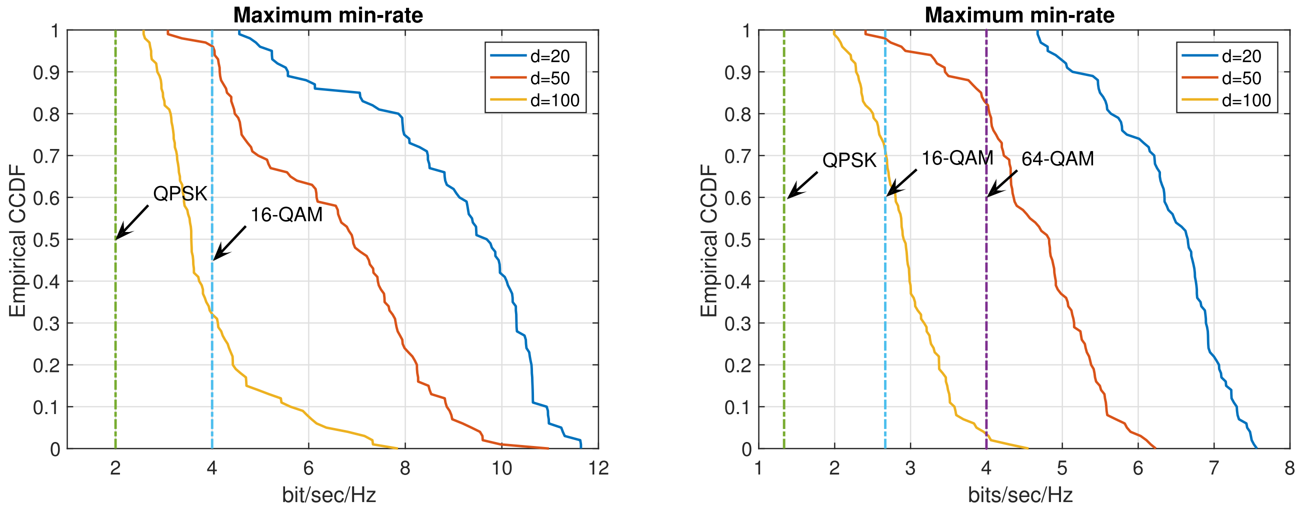

Fig. 6. Empirical CCDF of the maximum min-rate for (a) Q=4, K=4 (b) Q=6, K=4.

- An important observation is that an optimal allocation policy typically allocates more resource blocks to configurations with a larger number of users. This is because the overall data rate is linearly proportional to the number of simultaneous users, whereas the dependence on SINR is logarithmic. However, there are settings in which time multiplexing leads to a higher data rate than spatial multiplexing (for example, when users are highly spatially correlated such that by eliminating their mutual interference, higher data rates can be attained even over smaller portion of a resource block).
- Figure 7 demonstrates this phenomenon by showing a few examples for the solution to the resource allocation problem. The optimal solution tends towards serving maximum number of users simultaneously (blue portions) unless the induced interference is so large that only a subset of them are served (green or red portions).
- Our proposed recipe is designed toward achieving the maximum capacity, ignoring complexity considerations. An important topic for future work is to tune such PHY/MAC algorithms to meet specific overhead, computational complexity, and throughput requirements.

VI. Capacity Estimation

We now obtain numerical estimates of capacity using simulations, which show that dense mm-wave picocells provide a significant gain in capacity over conventional LTE cellular networks. Our goal in this section is to estimate the overall capacity over 1km^2; we do this by carrying out Monte Carlo simulations for users in a single picocell in the middle of a canyon and then roughly extrapolate our results to approximate per km^2 capacity using an example of a real-world urban area (Manhattan area). The system model and PHY/MAC layer parameters are summarized in Tables I and II, respectively.

[1]We have set this to largest possible spectral efficiency that can be provided to each of the Q users.

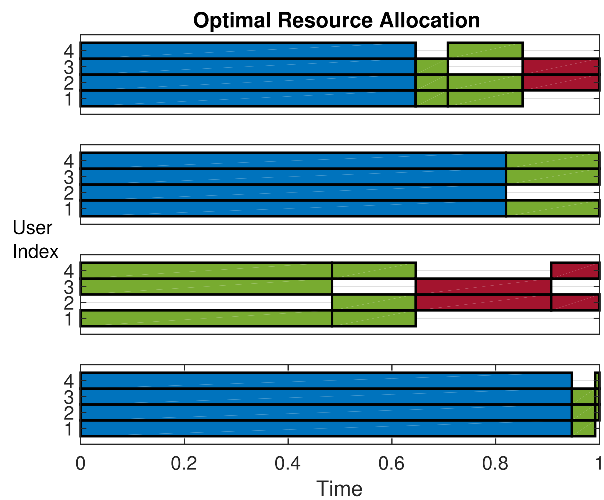

Fig. 7. Optimal solution of the resource allocation problem for different realizations of mobile users. The picocell parameters are d=50m and K=Q=4. Optimal allocation policies tend to serve the largest possible number of users (blue portions) while in some cases it is better to turn off a subset of subarrays, i.e., green/red portions.

A. Single Canyon Simulations

Our interference analysis in the preceding sections is partially geometry dependent and specifically tailored for cells along an urban canyon. Hence, for our simulations, we consider an urban canyon of length 1 km and investigate a picocell in the middle of this canyon, where users would see the most interference (Fig. 8).

We consider 8 × 8 basestation TX arrays and 4 × 4 mobile RX arrays. These values are chosen because they are close to the current state of the art (32 element arrays are already deployed in commercial 60 GHz products), and it turns out that they suffice to provide high spectral efficiency even as we scale down cell sizes.

Since a user in the target picocell can be served by one of two basestations on two different sides, it is unlikely for her body to block both LoS paths. Furthermore, as we shrink the

TABLE II
ALGORITHM 1 AND PHY LAYER PARAMETERS

Parameter	Definition	Default
$\mathbf{H}_q/h_q$	Channel matrix/vector toward q-th user	-
$\boldsymbol{\omega}_q$	Transmit beamforming vector toward q-th user	-
p_q^0	initial power allocation (dBm)	40-10log(N^2)
γ	initial SINR	0dB
$\Delta\gamma$	SINR step size	1dB
-	Maximum EIRP allowed by FCC at 60GHz	40dBm
-	Noise figure	6dB
-	Reflection loss	5dB
s_M	hardware-constrained spectral efficiency limit (64QAM)	6 bps/Hz
r_{min}	pre-defined minimum required spectral efficiency for each user (if possible)	$K \times S_M/Q$ [1]

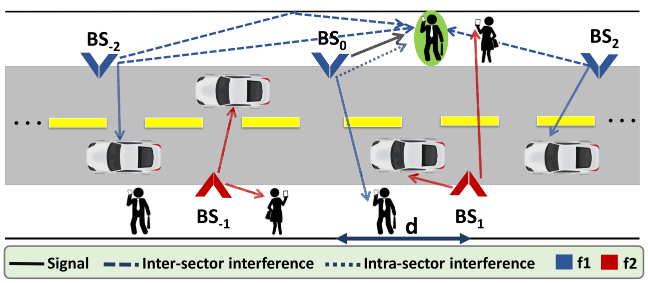

Fig. 8. Simulation scenario ($F = 2$).

picocell width, the LoS path slants more steeply downward, hence it is difficult for other obstacles (e.g., pedestrians, cars) to block it. Thus, in our computations, we assume for simplicity that at least one LoS path is available to every user. Of course, both LoS and first order NLoS paths are accounted for when computing interference from other subarrays. As noted in [18], interference from higher order reflections is negligible in comparison. In our numerical results, we have considered an average reflection loss of 5dB for NLoS paths.

By virtue of Theorems 1 and 2 from Section IV, we ignore interference coming from outside the 1 km segment of the canyon. Moreover, for a typical user served by BS_0, the interference induced by the base stations further than 2d away from BS_0 is negligible ($N_{max} = 2$). This is also verified by simulation results in our previous work [18]. Specifically, in the scenario depicted in Figure 8, the following sources would interfere with the shaded user served with one of K eastward facing antenna arrays of BS_0:

1) *inter-cell* interference from K eastward facing antenna arrays on BS_{-2}
2) *inter-cell* interference from K eastward facing antenna arrays on BS_{-1}
3) *intra-cell* interference from K-1 eastward facing antenna arrays on BS_0
4) *inter-cell* interference from K westward facing antenna arrays on BS_1
5) *inter-cell* interference from K westward facing antenna arrays on BS_2

Each of these is composed of LoS and multiple NLoS components.

As we have noted in Section IV, frequency reuse of one is a good design choice for a system with a single subarray per face. As we increase the number of subarrays per face, however, a frequency reuse of one incurs too much interference. Hence, in our simulations, we also employ frequency reuse of two which automatically eliminates items 2 and 4 above.

Moreover, in scenarios with more than one subarray per face ($K > 1$), we attenuate the LoS intra-cell interference (item 3) by employing Algorithm 1 in Section V-B, which takes the estimated spatial channels for candidate users as input, and determines the beamforming and power allocation scheme.

Employing the aforementioned interference management schemes, we compute the overall spectral efficiency, $\log_2(1+\text{SINR})$, for the users served by BS_0; taking into account the *residual* intra-cell interference from item 3 as well as inter-cell interference from applicable items (items 1, 2, 4, and 5 for frequency reuse of one and items 1 and 5 for frequency reuse of two). The resultant matrix S is then fed into the optimization problem (6) to obtain the maximum min-rate obtained by the optimal time allocation.

Figure 9 shows the empirical CCDF of the maximum min-rate provided for a typical user served by BS_0. Note that the hardware saturation points corresponding to QPSK, 16-QAM and 64-QAM modulations are $(K/Q)\, s_M = 1, 2$ and 3 respectively for the case $K = 2$ and $Q = 4$.

B. Capacity Calculations For Real-World Urban Area

We now use the preceding calculations to extrapolate the overall capacity per square kilometer in an urban area. We consider a 1km^2 region in Manhattan area (Fig. 10), which encompasses 15 urban canyons. Thus, we can get a rough estimate of the overall capacity per square kilometer of our approach via the following computations:

$$\begin{aligned}\text{Capacity (bps/km}^2) &= \text{Maximum min-rate (bps/Hz/user)}\\ &\quad\times \frac{B}{F}(\text{Hz}) \times 2Q\,(\text{Num. users / cell})\\ &\quad\times n_c\,(\text{Num. cells / km}^2) \qquad (8)\end{aligned}$$

where B, F and n_c are the total bandwidth, the frequency reuse factor and the number of picocells per square kilometer respectively. Note that $2Q$ in (8) refers to the number of users served within the picocell [2] which are covered by either eastward facing antennas of BS_0 or westward facing antennas of BS_1. In our example of a 1km^2 region in Manhattan shown in Figure 10, there are a total of 15 street canyons of length 1km (in both directions), each of which encompasses 1km$/d$ cells. Hence, we get $n_c \approx 150, 300$ and 750 for picocell widths of $d = 100$, 50 and 20 meters respectively.

We have summarized the preceding results in Table III specifying the overall attainable capacity for different scenarios. Note that the Maximum min-rate in equation (8) is replaced with the attained rate in Fig. 9 truncated to the hardware

[2]This requires $2Q \times n_c = 9000$ users/km^2 in our most extreme case: d=20m and Q=6 which is still much smaller than the population density of Manhattan area: 27,826 persons/km^2 [40].

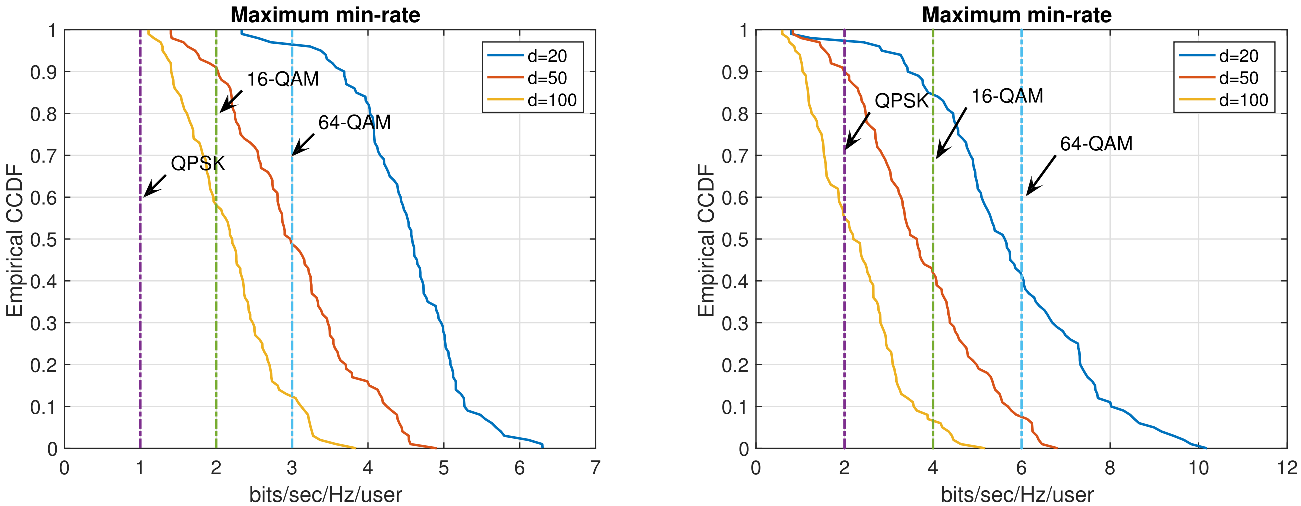

Fig. 9. Empirical CCDF of the maximum min-rate for users in a picocell in the middle of an urban canyon as depicted in Fig. 8, for (a) $K = 2$, $Q = 4$ and (b) $K = 4$, $Q = 4$.

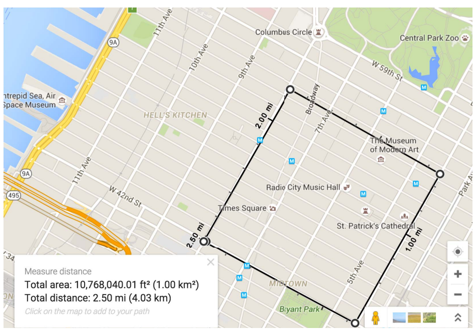

Fig. 10. 1 km^2 in Manhattan area, encompassing 15 street canyons.

TABLE III
CAPACITY (Tbps/km^2) OVER A TOTAL BANDWIDTH OF 2GHz FOR AN AREA IN NEW YORK EMPLOYING 8 × 8 AND 4 × 4 ANTENNA ARRAYS AS TRANSMITTER AND RECEIVERS

Capacity (Tbps/km^2)	$K = 1$	$K = 2$	$K = 1$	$K = 2$	$K = 4$
	$F = 1$		$F = 2$		
d=100 m	1.3	1.8	1.6	2.6	2.7
d=50 m	5.3	8.9	3.3	6.4	8.9
d=20 m	17.6	33.1	8.9	17.8	30.9

saturation point imposed by 64-QAM which is $(K/Q)\, s_M$ for $s_M = 6$ bps/Hz. Moreover, the first column in Table III corresponds to our previous results [18] for a single subarray per face.

Remarks

- Smaller picocells are less prone to interference, since the beams aiming towards their target users are slanted more steeply, and hence illuminate (and induce interference for) a smaller area around them. Moreover, almost vertical beams at smaller cells result in larger spacing between the spatial frequencies for different users. This makes it easier to isolate users with our proposed interference suppression algorithm, and hence it possible to gain more from employing additional subarrays per face. This feature, along with the increased spatial reuse attained with smaller cell sizes, leads to massive estimated capacity of up to 30.9 Tbps/km^2.
- Larger picocells are inherently more prone to interference due to their less slanted beams, which cause significant interference to users in a larger neighborhood around the target user. They also do not gain as much from employing additional subarrays per face (Table III). This is because almost horizontally aligned beams in wide cells lead to smaller spacing between spatial frequencies for different users, so that interference suppression is not as effective. Possible approaches to solve this problem are (a) increasing the number of antenna elements, which provides more degrees of freedom for employing interference suppression; or (b) increasing basestation height, which will separate users more in the spatial frequency domain.
- Employing a larger frequency reuse factor is a wasteful approach to deal with interference for smaller cells, and only leads to marginal improvement for larger cells. It is difficult to make up for the loss of degrees of freedom (on which the capacity depends linearly) with the gain in SINR (in which the capacity grows logarithmically). The gain in SINR is further limited by potential limits on constellation size due to hardware constraints

In addition to the preceding capacity estimates, we have also designed simulation scenarios in order to obtain insight into the effect of the MAC and PHY modules, by assigning various combinations of module status (ON/OFF) to the pair of (MAC,PHY) modules, for $K = 2, 4$. We assume the following natural default behaviors when the modules are turned off.

• **MAC = OFF** : The MAC module guarantees the max-min rate while utilizing the remaining resources to opportunistically maximize the overall network throughput. The default behavior, when the MAC module is disabled, is that each subarray selects a user randomly and commits to serving it

TABLE IV
DESIGN SPACE EXPLORATION OF CAPACITY (Tbps/km^2) WITH RESPECT TO INTER- AND INTRA-CELL MODULE STATUS ($d = 50$m, $Q = 4$)

			Intra-cell Module			
		MAC module	OFF		ON	
		PHY module	OFF	ON	OFF	ON
Inter-cell Module	F=1	K=2	5.6	4.9	8.6	8.9
		K=4	6.9	6.9	8.5	8.1
	F=2	K=2	3.2	4.2	5.8	6.4
		K=4	4.2	6.2	5.5	8.9

over the entire resource block (time/frequency). Note that such a simple scheme would not be able to guarantee a pre-defined minimum rate for all Q users (as our MAC module is designed to do), and the K selected users would potentially interfere with each other over the entire resource block.

• **PHY** $=$ **OFF** : Our PHY module, employs LMMSE beamforming and power allocation to suppress the LoS intra-cell interference among the active users sharing a resource block. The default behavior, when the PHY module is disabled, is that each subarray employs transmit beamforming toward its own target user while ignoring interference to other users.

The simulation results are presented in Table IV in terms of the average capacity (Tbps/km^2), employing the same simulation setup as in Table III. Some key observations from the simulation results shown in this table are as follows:

• Our PHY layer mechanism suppresses interference by pushing nulls toward non-targeted users. This reduces interference toward other users but might lead to SNR loss for the target user. As shown in the table, in cases with severe interference ($F = 1$ and $K = 4$) the degradation in SNR loss might break even with, or even outweigh, the benefits of interference suppression. In such cases, the MAC layer becomes more critical, since it is able to orthogonalize spatially correlated users in another domain (time/frequency).

• The MAC layer method is always beneficial and leads to higher capacity when compared to the corresponding scenarios with MAC layer OFF. The improvement due to the MAC layer becomes more significant when the number of users is larger than the available antennas ($Q >> K$), which makes it more likely that we can find a subset of size K in which the users are less spatially correlated. This makes it more likely that we can serve the maximum number K of concurrent users without incurring excessive intra-cell interference.

VII. DISCUSSION AND CONCLUSION

In this paper, we investigate the attainable downlink capacity of mm-wave picocellular networks deployed in dense urban environments. We assume that it is possible to transmit to multiple users simultaneously in a given base station sector. Our nominal configuration for enabling the latter capability is to deploy multiple arrays, each capable of RF beamforming, but the abstraction used here also applies to hybrid or digital beamforming using a single array.

TABLE V
COMPARING CONVENTION LTE AND mm-WAVE CELLULAR NETWORKS

	LTE	mm-wave		Gain
		d=20	d=100	
Capacity	1.2Gbps	30.9 Tbps	2.7 Tbps	$\geq$ 2250X
Bandwidth	500 MHz	2GHz		4X
Spatial reuse	–	–		$\geq$ 550X

We find that as we shrink the cell size (down to a cell width of 20m), the per-user spectral efficiency is mostly ($\geq 97\%$ of the time) bounded by hardware limitations (the bound we use is $s_M = 6$ bps/Hz, corresponding to uncoded 64QAM). Larger cells are more prone to interference, but our proposed scheme provides users with sufficient spectral efficiency for supporting smaller constellations such as QPSK.

We now provide a rough estimate of the capacity gains attained relative to conventional LTE networks. The downlink capacity of LTE network is estimated as 0.6 Gbps/km^2 over a total bandwidth of 255 MHz in [41]. However, the available bandwidth for downlink cellular networks is 500 MHz, hence the total capacity could be further increased by adding more channels per base station. Therefore, we estimate the total downlink capacity of LTE networks as 1.2 Gbps/km^2.

Table V compares the resultant capacity for mm-wave picocells computed via simulations with the benchmark capacity of LTE networks. We see that the targeted 1000-fold capacity increase is reachable even with the largest picocell size ($d = 100$m) considered here. Excluding the 4X gain from the larger bandwidth of 2GHz employed in our system (which is still a small fraction of the 14GHz of available bandwidth at 60GHz), the remaining gain ($\geq$ 550X) is attained through the larger spatial reuse from small cells and pencil beams. Of course, as mentioned in the introduction, many implementation challenges must be surmounted in order to attain these potential gains. Our results provide a compelling motivation for a sustained effort in addressing these challenges.

It is worth emphasizing again the contrast between our results and those at lower frequencies. As we increase cell density, interference can become a fundamental limiting factor at lower carrier frequencies [19]. Our analysis shows that this is not the case for mm-wave frequencies: the narrow beams yield large gains in spatial reuse, which translate to orders of magnitude capacity increases.

REFERENCES

[1] I. Chih-Lin, S. Han, Z. Xu, Q. Sun, and Z. Pan, "5G: Rethink mobile communications for 2020+," *Philos. Trans. Roy. Soc. A, Math., Phys. Eng. Sci.*, vol. 374, no. 2062, Mar. 2016, Art. no. 20140432.

[2] N. Bhushan *et al.*, "Industry perspective," *IEEE Wireless Commun.*, vol. 24, no. 5, pp. 6–8, Oct. 2017.

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

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

[5] Y. Zhu *et al.*, "Demystifying 60ghz outdoor picocells," in *Proc. 20th Annu. Int. Conf. Mobile Comput. Netw.*, Sep. 2014, pp. 5–16.

[6] T. S. Rappaport, F. Gutierrez, Jr., E. Ben-Dor, J. N. Murdock, Y. Qiao, and J. I. Tamir, "Broadband millimeter-wave propagation measurements and models using adaptive-beam antennas for outdoor urban cellular communications," *IEEE Trans. Antennas Propag.*, vol. 61, no. 4, pp. 1850–1859, Apr. 2013.

[7] H. Zhang, S. Venkateswaran, and U. Madhow, "Channel modeling and MIMO capacity for outdoor millimeter wave links," in *Proc. IEEE Wireless Commun. Netw. Conf.*, Apr. 2010, pp. 1–6.

[8] J. Karjalainen, M. Nekovee, H. Benn, W. Kim, J. Park, and H. Sungsoo, "Challenges and opportunities of mm-wave communication in 5G networks," in *Proc. 9th Int. Conf. Cognit. Radio Oriented Wireless Netw. Commun. (CROWNCOM)*, Jun. 2014, pp. 372–376.

[9] Z. Marzi, D. Ramasamy, and U. Madhow, "Compressive channel estimation and tracking for large arrays in mm-wave picocells," *IEEE J. Sel. Topics Signal Process.*, vol. 10, no. 3, pp. 514–527, Apr. 2016.

[10] D. Ramasamy, S. Venkateswaran, and U. Madhow, "Compressive tracking with 1000-element arrays: A framework for multi-gbps mm wave cellular downlinks," in *Proc. 50th Annu. Allerton Conf. Commun., Control, Comput. (Allerton)*, Oct. 2012, pp. 690–697.

[11] D. Ramasamy, S. Venkateswaran, and U. Madhow, "Compressive adaptation of large steerable arrays," in *Proc. ITA*, Feb. 2012, pp. 234–239.

[12] M. E. Rasekh, Z. Marzi, Y. Zhu, U. Madhow, and H. Zheng, "Noncoherent mmWave path tracking," in *Proc. 18th Int. Workshop Mobile Comput. Syst. Appl.*, Feb. 2017, pp. 13–18.

[13] H. Hassanieh, O. Abari, M. Rodriguez, M. Abdelghany, D. Katabi, and P. Indyk, "Fast millimeter wave beam alignment," in *Proc. Conf. ACM Special Interest Group Data Commun.*, New York, NY, USA, Aug. 2018, pp. 432–445.

[14] M. N. Islam, S. Subramanian, and A. Sampath, "Integrated access backhaul in millimeter wave networks," in *Proc. IEEE Wireless Commun. Netw. Conf. (WCNC)*, Mar. 2017, pp. 1–6.

[15] M. N. Kulkarni, A. Ghosh, and J. G. Andrews, "Max-min rates in self-backhauled millimeter wave cellular networks," May 2018, *arXiv:1805.01040*. [Online]. Available: https://arxiv.org/abs/1805.01040

[16] M. E. Rasekh, D. Guo, and U. Madhow, "Interference-aware routing and spectrum allocation for millimeter wave backhaul in urban picocells," in *Proc. 53rd Annu. Allerton Conf. Commun., Control, Comput. (Allerton)*, Sep./Oct. 2015, pp. 1–7.

[17] M. E. Rasekh, D. Guo, and U. Madhow, "Joint routing and resource allocation for millimeter wave picocellular backhaul," Sep. 2018, *arXiv:1809.07470*. [Online]. Available: https://arxiv.org/abs/1405.6636

[18] Z. Marzi, U. Madhow, and H. Zheng, "Interference analysis for mm-Wave picocells," in *Proc. IEEE Global Commun. Conf. (GLOBECOM)*, Dec. 2015, pp. 1–6.

[19] D. Ramasamy, R. Ganti, and U. Madhow, "On the capacity of picocellular networks," in *Proc. IEEE Int. Symp. Inf. Theory*, Jul. 2013, pp. 241–245.

[20] E. Visotsky and U. Madhow, "Optimum beamforming using transmit antenna arrays," in *Proc. IEEE 49th Veh. Technol. Conf.*, vol. 1, May 1999, pp. 851–856.

[21] A. Wiesel, Y. C. Eldar, and S. Shamai (Shitz), "Linear precoding via conic optimization for fixed MIMO receivers," *IEEE Trans. Signal Process.*, vol. 54, no. 1, pp. 161–176, Dec. 2005.

[22] S. Ulukus and R. D. Yates, "Adaptive power control and MMSE interference suppression," *Wireless Netw.*, vol. 4, no. 6, pp. 489–496, Nov. 1998.

[23] F. Rashid-Farrokhi, K. J. R. Liu, and L. Tassiulas, "Transmit beamforming and power control for cellular wireless systems," *IEEE J. Sel. Areas Commun.*, vol. 16, no. 8, pp. 1437–1450, Oct. 1998.

[24] F. Rashid-Farrokhi, L. Tassiulas, and K. J. R. Liu, "Joint optimal power control and beamforming in wireless networks using antenna arrays," *IEEE Trans. Commun.*, vol. 46, no. 10, pp. 1313–1324, Oct. 1998.

[25] B. Zhuang, D. Guo, and M. L. Honig, "Traffic driven resource allocation in heterogenous wireless networks," May 2014, *arXiv:1405.6636*. [Online]. Available: https://arxiv.org/abs/1405.6636

[26] B. Zhuang, D. Guo, E. Wei, and M. Honig, "Large-scale spectrum allocation for cellular networks via sparse optimization," *IEEE Trans. Signal Process.*, vol. 66, no. 20, pp. 5470–5483, Oct. 2018.

[27] M. R. Akdeniz *et al.*, "Millimeter wave channel modeling and cellular capacity evaluation," *IEEE J. Sel. Areas Commun.*, vol. 32, no. 6, pp. 1164–1179, Jun. 2014.

[28] M. Abouelseoud and G. Charlton, "System level performance of millimeter-wave access link for outdoor coverage," in *Proc. IEEE Wireless Commun. Netw. Conf. (WCNC)*, Apr. 2013, pp. 4146–4151.

[29] T. Bai and R. W. Heath, Jr., "Coverage and rate analysis for millimeter-wave cellular networks," *IEEE Trans. Wireless Commun.*, vol. 14, no. 2, pp. 1100–1114, Feb. 2014.

[30] S. Akoum, O. El Ayach, and R. W. Heath, Jr., "Coverage and capacity in mmWave cellular systems," in *Proc. 46th Asilomar Conf. Signals, Syst. Comput. (ASILOMAR)*, Nov. 2012, pp. 688–692.

[31] J. S. Kim, J. S. Shin, S.-M. Oh, A.-S. Park, and M. Y. Chung, "System coverage and capacity analysis on millimeter-wave band for 5g mobile communication systems with massive antenna structure," *Int. J. Antennas Propag.*, vol. 2014, Jul. 2014, Art. no. 139063.

[32] J. Qiao, X. Shen, J. Mark, Q. Shen, Y. He, and L. Lei, "Enabling device-to-device communications in millimeter-wave 5G cellular networks," *IEEE Commun. Mag.*, vol. 53, no. 1, pp. 209–215, Jan. 2015.

[33] C. Yiu and S. Singh, "Empirical capacity of mmWave WLANs," *IEEE J. Sel. Areas Commun.*, vol. 27, no. 8, pp. 1479–1487, Oct. 2009.

[34] L. X. Cai, L. Cai, X. Shen, and J. W. Mark, "Spatial multiplexing capacity analysis of mmWave WPANs with directional antennae," in *Proc. IEEE Global Telecommun. Conf.*, Nov. 20078, pp. 4744–4748.

[35] M. K. Samimi and T. S. Rappaport, "3-D millimeter-wave statistical channel model for 5G wireless system design," *IEEE Trans. Microw. Theory Techn.*, vol. 64, no. 7, pp. 2207–2225, Jul. 2016.

[36] S. Sun, G. R. MacCartney, and T. S. Rappaport, "A novel millimeter-wave channel simulator and applications for 5g wireless communications," in *Proc. IEEE Int. Conf. Commun. (ICC)*, May 2017, pp. 1–7.

[37] T. S. Rappaport, S. Sun, and M. Shafi, "Investigation and comparison of 3GPP and NYUSIM channel models for 5G wireless communications," in *Proc. IEEE 86th Veh. Technol. Conf. (VTC-Fall)*, Sep. 2017, pp. 1–5.

[38] M. E. Rasekh, A. A. Shishegar, and F. Farzaneh, "A study of the effect of diffraction and rough surface scatter modeling on ray tracing results in an urban environment at 60 GHz," in *Proc. 1st Conf. Millim.-Wave THz Technol. (MMWaTT)*, Dec. 2009, pp. 27–31.

[39] M. Mezzavilla *et al.*, "End-to-end simulation of 5G mmWave networks," *IEEE Commun. Surveys Tuts.*, vol. 20, no. 3, pp. 2237–2263, 3rd Quart., 2018.

[40] Wikipedia Contributors. (2018). *Demographics of New York City—Wikipedia, the Free Encyclopedia*. Accessed: Sep. 11, 2018. [Online]. Available: https://en.wikipedia.org/w/index.php?title=Demographics_of_New_York_City&oldid=858444014

[41] M. Sauter, *3G, 4G and Beyond: Bringing Networks, Devices and the Web Together*. Hoboken, NJ, USA: Wiley, 2013.

Zhinus Marzi received the B.Sc. degree in electrical engineering from the Iran University of Science and Technology, Tehran, Iran, in 2010, the M.Sc. degree in electrical engineering from the Sharif University of Technology, Tehran, in 2012, and the Ph.D. degree in communications and signal processing from the University of California at Santa Barbara (UCSB) in 2019. She is currently an Research and Development Engineer with Broadcom Inc., where she has been involved in high-performance wireless communication system design. Her research interest includes millimeter-wave communications.

Upamanyu Madhow is Distinguished Professor of Electrical and Computer Engineering at the University of California, Santa Barbara. His current research interests focus on next generation communication, sensing and inference infrastructures centered around millimeter wave systems, and on robust machine learning. He received his bachelor's degree in electrical engineering from the Indian Institute of Technology, Kanpur, in 1985, and his Ph. D. degree in electrical engineering from the University of Illinois, Urbana-Champaign in 1990. He has worked as a research scientist at Bell Communications Research, Morristown, NJ, and as a faculty at the University of Illinois, Urbana- Champaign. Dr. Madhow is a recipient of the 1996 NSF CAREER award, and co-recipient of the 2012 IEEE Marconi prize paper award in wireless communications. He has served as Associate Editor for the IEEE TRANSACTIONS ON COMMUNICATIONS, the IEEE TRANSACTIONS ON INFORMATION THEORY, and the IEEE TRANSACTIONS ON INFORMATION FORENSICS AND SECURITY. He is the author of two textbooks published by Cambridge University Press, *Fundamentals of Digital Communication* (2008) and *Introduction to Communication Systems* (2014).