

# On the Capacity of Picocellular Networks

Dinesh Ramasamy<sup>†</sup>, Radhakrishna Ganti<sup>††</sup>, Upamanyu Madhow<sup>†</sup>

**Abstract**—The orders of magnitude increase in projected demand for wireless cellular data require drastic increases in spatial reuse, with picocells with diameters of the order of 100-200 m supplementing existing macrocells whose diameters are of the order of kilometers. In this paper, we observe that the nature of interference changes fundamentally as we shrink cell size, with near line of sight interference from neighboring picocells seeing significantly smaller path loss exponents than interference in macrocellular environments. Using a propagation model proposed by Franceschetti, which compactly models increased interference in small cells, we show that the network capacity does not scale linearly with the reduction in cell size with standard frequency reuse strategies. Rather, more sophisticated resource sharing strategies based on beamforming and base station cooperation are required to realize the potential of small cells in providing high spectral efficiencies and quasi-deterministic guarantees on availability. Numerical results justifying these conclusions include Chernoff bounds on outage probability for random base station deployment (according to a spatial Poisson process), and simulations for deployment in a regular grid.

## I. INTRODUCTION

Cellular networks today face a capacity crisis due to the proliferation of data-hungry mobile devices, with projections of orders of magnitude growth in data demand from today’s baseline. The required increase in network capacity dictates drastic reductions in cell size, with picocells with diameters of the order of 100-200 m complementing the existing macrocellular infrastructure, where cells have diameter of the order of kilometers. In principle, for an interference-limited network, reduction of cell sizes by two orders of magnitude should translate directly to a corresponding increase in capacity. However, neighboring interferers for small cells can do far more damage than in a macrocellular environments: interference from neighboring picocells can often be near line of sight (LoS), whereas the interference seen by a macrocell from its neighbors may often be attenuated by buildings and other obstacles, and can experience destructive multipath interference. In this paper, we explore the quantitative consequence of this observation on the network capacity attainable with picocellular architectures. For simplicity, we concentrate on interference-limited performance, ignoring thermal noise, and focus on the downlink (which currently is, and is expected to remain, the bottleneck because of the asymmetric nature of data traffic).

Intuitively, in a picocellular environment, we expect near-LoS interference from adjacent cells to decay as  $1/r^2$ , but that from far-away cells to decay faster (a common approach is to model it as  $1/r^4$  or faster). However, instead of transitions between power laws as a function of range, however, a unified path loss model that captures the preceding intuition is  $e^{-\gamma r}/r^2$ , derived by Franceschetti et al [1] based on a “wandering photon” model (streams of photons emitted by the transmitter either get absorbed or scattered in a random direction on hitting an obstacle). The decay exponent  $\gamma$  depends on the density of obstacles and the elevation of the transmit/receive antennas. In our investigation, we set  $\gamma = 150 m^{-1}$  for picocellular environments in which base station heights are expected to be below rooftops (e.g., for a lamppost based deployment) and  $\gamma = 500 m^{-1}$  for a macrocellular environment with elevated base stations. Thus, scaling down from macrocells to picocells fundamentally changes the propagation model, and we can no longer expect linear scaling of network capacity with reduction in cell size. Another important difference from conventional capacity analysis is our desired regime of operation. While digital cellular networks supporting voice can operate at relatively low signal-to-interference ratio (SIR), our goal here is to provide high peak rates for bursty data traffic, operating at high SIR (and hence high spectral efficiency) with very low probability of outage (e.g., 0.1-0.01 %). This requires that neighboring base stations must coordinate their transmissions, and our goal is to explore the performance of coordination strategies at various levels of sophistication. Our performance analysis employs Chernoff bounds on outage probability for base stations deployed in the plane according to a Poisson point process (corresponding to irregular cell sizes), and simulation-based estimates of outage probability for base stations deployed according to a regular grid on the plane. The former provides a conservative design framework, while the latter may be a more accurate reflection of deployments in Manhattan-like urban environments.

*Contributions:* We begin by considering omnidirectional transmission. We first show, via a lower bound on outage probability, that randomized resource reuse strategies (e.g., each base station randomly choosing one of  $K$  orthogonal channels), cannot provide the quasi-deterministic performance guarantees we require. We then consider a strategy in which each base station orthogonalizes its resources with  $S - 1$  nearest neighbors (we specify  $S$  later). We find that, for a particular value of outage probability, picocellular networks must employ a larger value of  $K$  than macrocellular networks, so that the network capacity increase is significantly smaller than the factor by which we shrink cell size. Next, we explore the extent to which interference can be reduced by employing

The authors are with <sup>†</sup>ECE Department, UCSB and <sup>††</sup>EE Department, IIT Madras. Email: {dineshr, madhow}@ece.ucsb.edu and rganti@ee.iitm.ac.in. This work was supported by the National Science Foundation through the grant CNS-0832154, and by the Institute for Collaborative Biotechnologies through the grant W911NF-09-0001 from the U.S. Army Research Office. The content of the information does not necessarily reflect the position or the policy of the Government, and no official endorsement should be inferred.

antenna arrays at the base station. Given the relatively small form factor required for picocellular networks, the number of antennas on a base stations operating in bands from 1-5 GHz cannot be too large. We consider base station arrays with 8 antenna elements, and find that beamforming on the downlink, together with resource orthogonalization with  $S-1$  nearest neighbors, does cut down significantly on interference. However, while beamforming significantly increases the mean SIR for a given value of channel reuse factor  $K$  (and therefore allows reduction in  $K$  for a desired SIR target), it does not provide the quasi-deterministic performance we desire in terms of very small outage rates. The final system we consider, therefore, is the use of Coordinated Multipoint (CoMP), where neighboring base stations (each equipped with an antenna array) form a larger distributed array for beamforming to mobiles in *virtual cells* defined by clusters of collaborating base stations. We find that this strategy provides enough directivity to allow aggressive spatial reuse (small channel reuse factor  $K$ ) while maintaining high SIR with quasi-deterministic guarantees on outage. Further, when this strategy is combined with Space Division Multiple Access (SDMA) on the downlink, it is able to provide orders of magnitude gain in network capacity relative to macrocellular networks, along with quasi-deterministic guarantees on outage. Our results indicate the importance of designing efficient techniques for tight base station coordination if the promise of increased spatial reuse from shrinking cell size is to be realized.

*Related work:* A spatial Poisson point process (PPP) to model irregular cells was proposed in [2], which computed coverage probabilities and ergodic rates for a conventional  $r^{-a}$  ( $a > 2$ ) model for path loss. While we differ from this work in our path loss model (and hence in our conclusions on performance with small cells) and consider a near-LOS rather than Rayleigh faded model for the desired user, we use similar stochastic geometry techniques for deriving our Chernoff bounds. Urban channel measurements supporting our  $e^{-\gamma r}/r^2$  path loss model appear in [1], [3], while a power law regime change from  $a = 2$  to  $a = 4$  is reported in the measurements in [4]. Shadowing due to randomly placed buildings is shown in [5] to also give rise to a  $e^{-\gamma r}/r^2$  path loss function.

## II. SYSTEM MODEL

We focus on the downlink, which is the bottleneck for mobile data.

### A. Propagation model

We employ the wandering photon propagation model proposed by Franceschetti et al [1] to model path loss, which is given by

$$\ell(r) = e^{-\gamma r}/r^2, \quad (1)$$

with the parameter  $\gamma$  determined by the propagation geometry.

### B. BS placement model

We employ the following basestation (BS) placement models in our analysis. Let  $\lambda$  denote the spatial density of BSs. *PPP model:* The BS locations are modeled by a spatial Poisson

Point Process (PPP) of intensity  $\lambda$  [2]. *Square grid model:* The BSs are arranged at the vertices of a square tessellation of side  $1/\sqrt{\lambda}$ . We investigate two cell size regimes: (i) *Small cells:* The density of BSs  $\lambda$  is  $100 \text{ km}^{-2}$  and the path loss parameter  $\gamma$  in (1) is set at  $150^{-1} \text{ m}^{-1}$ . (ii) *Large cells* are modeled by a path loss parameter of  $\gamma = 500^{-1} \text{ m}^{-1}$  and a BS density of  $\lambda = 1 \text{ km}^{-2}$ . Small cell BSs are typically mounted at lower heights (say 4 m) than their large cell counterparts (say 13 m) and the values of  $\gamma$  used are consistent with predictions in [4] for these antenna heights.

### C. Performance measure

We assume that all BSs transmit with the same transmit power, and that the noise floor is easily cleared. Thus, the metric of interest is the Signal-to-Interference Ratio, SIR. Without loss of generality, consider a “desired” mobile user at the origin  $o$ . Let  $x$  denote the location of the closest BS to the user at the origin. Let  $\Phi$  denote the set of locations of the BSs in the network which use the same resource block as the BS at  $x$ . Then the downlink SIR for the desired user, assuming omnidirectional transmission, is given by

$$\text{SIR} = \frac{\ell(\|x\|)}{\sum_{z \in \Phi \setminus \{x\}} \ell(\|z\|)}. \quad (2)$$

Henceforth, we abuse notation and denote  $\ell(\|x\|)$  by  $\ell(x)$ . The outage probability for an SIR target  $\theta$  is given by  $\mathbb{P}[\text{SIR} < \theta]$ . We assume that the network designer expects to cater to high peak rate demands with overwhelming probability (essentially in a “quasi-deterministic” manner), hence we are interested in large SIRs (corresponding to high spectral utilization) and very small probability of outage. In the forthcoming sections, we compute outage probability versus SIR target curves for a number of different transmission and resource sharing policies. We analyze three systems (with increasing levels of sophistication in transmission and coordination) in Sections III-V, and briefly discuss design implications in Section VI.

## III. OMNIDIRECTIONAL TRANSMISSION

For omnidirectional transmission from the base stations, we consider two resource allocation schemes for the PPP model: (i) Random reuse, where each BS randomly chooses one in  $K$  available resource blocks, irrespective of the choices of the neighboring BSs (no coordination) and (ii) Local orthogonalization using  $K$  resource blocks, using approximate abstractions for the resource sharing model motivated by the square grid, for which regular reuse is possible. We derive bounds on outage probability for the PPP model, and employ simulations for the square grid.

### A. Randomized reuse does not work

We first present a lower bound on outage probability for random reuse using  $K$  resource blocks. The BSs are distributed in space as a PPP( $\Phi$ ) of intensity  $\lambda$  and they each choose one of  $K$  available resource blocks at random with equal probability. The user is assumed to be at the origin  $o$ . Let  $R$  denote the distance to the nearest BS. This is the BS that




Fig. 1. Lower bound on the probability of outage as a function of SIR for small cells with random resource allocation from  $K = 4, 9$  resource blocks for the small cell scenario in a PPP deployment.




Fig. 2. Left: Square grid deployment of BSs colored using regular reuse  $K = 4$ . Right: Square grid deployment of BSs forming “virtual cells” used for CoMP, and these virtual cells are colored using regular reuse  $K = 4$ .

serves the user. Let  $f_R(r)$  denote the PDF of  $R$ . For a spatial PPP, the distribution of  $R$  is  $f_R(r) = 2\pi\lambda r \exp(-\pi\lambda r^2)$ . From the strong Markov property and independent thinning property [6], the BSs outside the disc  $B(o, R)$  that use the same resource block as the BS closest to the origin is a PPP  $\Phi'$  with density  $\lambda/K$ . Let us now consider the SIR at this user:  $\text{SIR} = \ell(r) / \sum_{x \in \Phi'} \ell(x)$ . The probability of outage (for any  $\theta > 0$  dB) can be lower bounded in the following manner (protocol model for interference, which follows from the inequality  $\sum_{x \in \Phi'} \ell(x) \geq \max_{x \in \Phi'} \ell(x)$ ):

$$\begin{aligned} \Pr[\text{SIR} < \theta \mid R = r] &\geq \Pr[\exists x \in \Phi' : \|x\| < \ell^{-1}(\theta^{-1}\ell(r))] \\ &= 1 - e^{-\pi(\lambda/K)((\ell^{-1}(\theta^{-1}\ell(r)))^2 - r^2)}, \end{aligned}$$

where  $\ell^{-1}(x)$  is the inverse of  $\ell(x)$ . Averaging over  $R$ , the distance from the user to the nearest BS, we get a lower bound on outage probabilities for different SIR thresholds  $\theta$ . Figure 1 plots these for randomized reuse with  $K = 4$  and  $K = 9$  resource blocks for the small cell scenario. Even at a moderate SIR of 5 dB, the outage probability is at least 10%, so that randomized reuse is not a viable scheme for small cells. Thus, BSs must collaborate in order to attain the desired large SIR/low outage regime.

### B. Orthogonalization with neighbors

We will now examine reuse schemes which aim to orthogonalize resources among neighboring BSs.

1) *Square grid deployment:* The BSs are placed on the square grid of density  $\lambda$  given by  $\Phi = \{(m, n)/\sqrt{\lambda} : m, n \in \mathbb{Z}\}$ . The  $k^2$ -regular reuse




Fig. 3. Left: Probability of outage as a function of SIR for square grids with regular reuse  $K = 4, 9$  for both the small and large cell scenarios. Right: Upper bound on the probability of outage as a function of SIR corresponding to orthogonalization for the small and large cell scenarios in a PPP deployment. The parameter pair  $(K, S)$  was set to  $(4, 9)$  and  $(9, 25)$ .

policy on this grid is the following: All BSs in the set  $\Phi_{i,j} = \{(k \times (m, n) + (i, j)) / \sqrt{\lambda} : m, n \in \mathbb{Z}\}$  use the same resource (reuse rate  $K = k^2$  as there are  $k^2$  such sets corresponding to  $i, j = 0, \dots, k - 1$ ). This regular reuse policy for  $K = 4$  is shown in Figure 2 (left). Figure 3 (left) plots simulation-based outage curves for reuse rates  $K = 4$  and  $K = 9$  for small and large cells.

2) *PPP deployment:* Resource allocation schemes for random deployments that aim to coordinate with neighboring BSs to reduce out of cell interference are not easy to analyze. We therefore propose an analytically tractable, tunable abstraction for such coordination schemes. Let  $K$  denote the number of resource blocks available in the network. Consider the  $S$  BSs closest to the mobile user. The network ensures that from this set only the BS which is the closest to the mobile uses the resource block intended for the mobile. All BSs apart from the nearest  $S$  are assumed to choose a resource block at random from the  $K$  available resource blocks. We now give an upper bound on the probability of outage for this abstraction of coordinated resource reuse.

Consider a user at the origin  $o$ . Let  $R$  denote the distance to the nearest BS. Let  $X$  denote the distance to the  $S$ -th nearest BS from the mobile user. In our model, only BSs beyond this distance  $X$  from the mobile can interfere with the user. The BSs are deployed as a PPP( $\Phi$ ) of intensity  $\lambda$  and therefore the joint pdf of  $(R, X)$  is given by:

$$p(r, x) = \frac{(2\pi\lambda)^2 r x}{(S-2)!} (\lambda\pi(x^2 - r^2))^{S-2} e^{-\lambda\pi x^2} \mathbf{1}(x \geq r). \quad (3)$$

We observe that the interfering BSs lie outside the disc  $B(o, X)$  and form a PPP  $\Phi'$  of density  $\lambda/K$ . Let  $I$  denote the interference term  $\sum_{z \in \Phi'} \ell(z)$ . We upper bound the outage probability using a Chernoff bound on  $I$ , as follows:

$$\begin{aligned} \Pr[\text{SIR} < \theta \mid (R, X) = (r, x)] &= \Pr[I\theta > \ell(r)] \\ &\stackrel{s \geq 0}{\leq} \Pr[e^{sI\theta} > e^{s\ell(r)}] \leq \frac{\mathbb{E}e^{sI\theta}}{e^{s\ell(r)}}. \end{aligned}$$

Using the probability generating functional (PGFL) [6] of a

PPP, we evaluate  $\mathbb{E}e^{sI\theta}$ :

$$\mathbb{E} \prod_{z \in \Phi'} e^{s\theta\ell(z)} = \exp \left( -\frac{2\pi\lambda}{K} \int_{y=x}^{y=\infty} y \left( 1 - e^{s\theta\ell(y)} \right) dy \right).$$

A Chernoff bound on the outage probability is therefore given by  $\Pr [\text{SIR} < \theta \mid (R, X) = (r, x)] \leq b(r, x)$ , where,

$$b(r, x) = \min_{s \geq 0} e^{-s\ell(r)} \exp \left( -\frac{2\pi\lambda}{K} \int_{y=x}^{y=\infty} y \left( 1 - e^{s\theta\ell(y)} \right) dy \right).$$

Averaging this over the joint distribution of  $(R, X)$ , we have an upper bound on the outage probability:

$$\Pr [\text{SIR} < \theta] \leq \int_{r=0}^{r=\infty} \int_{x=r}^{x=\infty} p(r, x) b(r, x) dr dx.$$

*Choice of parameter  $S$  for reuse rate  $K$ :* The parameter  $S - 1$  in our abstraction of a coordinated reuse strategy is the number of neighboring BSs which are not in conflict with the user. For a given reuse rate  $K$ , we choose  $S$  by using regular reuse on the square grid as our guide. For e.g., for reuse 4, in the square grid this number corresponding to a user near the cell center is  $S - 1 = 8$  (refer figure 2 (left)), hence we set  $S = 9$  for reuse  $K = 4$ . Similarly, for reuse  $K = 9$ , we set  $S = 25$ . Figure 3 (right) plots these upper bounds for small and large cells.

#### IV. TRANSMIT BEAMFORMING

We now assume that each BS steers its beam towards its receiving node. The interfering BSs beamform towards their intended receivers, hence the orientation of their beams is modeled as random relative to the tagged user under investigation. We now quantify the improvement in outage probability due to the resulting attenuation in interference. The expression for SIR in (2) is modified to:  $\text{SIR} = \ell(x) / \sum_{z \in \Phi' \setminus \{x\}} \mathbf{h}_z \ell(z)$ , where  $0 \leq \mathbf{h}_z \leq 1$  is the gain of the directional transmit antenna at the interfering BS placed at  $z$  along the link between this BS and the tagged mobile receiver. We assume a steerable array with 8 elements placed uniformly on a circle of diameter equal to the carrier wavelength 15 cm. We evaluate the outage probability by including the effect of the array in the square grid simulations. For the random deployment model, we assume that all  $\mathbf{h}_z$  are i.i.d. and independent of the PPP. The distribution of  $\mathbf{h}_z$  depends on the antenna array used at the BSs and we compute the necessary statistics of  $\mathbf{h}_z$  by simulation.

*Square grid deployment:* We use the same regular reuse scheme described in Section III-B. The variables  $\mathbf{h}_z$  are evaluated using the beam patterns at the interfering BSs (which beamform toward their intended receivers). The probability of outage curves obtained by simulation for the small cell scenario for reuse rates  $K = 4$  and  $K = 9$  are plotted in Figure 4 (Left).

*Random deployment:* We consider the same abstraction for local coordination as in Section III-B. Following the analysis in Section III-B, we see that, conditioned on  $(R, X) = (r, x)$ ,  $x \geq r$ , an upper bound on the probability of outage is given by:  $\Pr [\text{SIR} < \theta \mid (R, X) = (r, x)] \leq e^{-s\ell(r)} \mathbb{E}e^{sI\theta}$  for any  $s > 0$ ,




Fig. 4. Transmit beamforming in the small cell regime for reuse  $K = 4, 9$ . Left: Probability of outage versus SIR for the square grid deployment with regular reuse. Right: Upper bounds on probability of outage as a function of SIR.

where the expression for interference  $I$  is  $\sum_{z \in \Phi'} \mathbf{h}_z \ell(z)$ . Remembering that  $\{\mathbf{h}_z : z \in \Phi'\}$  are independent of the PPP( $\Phi'$ ) and of each other, we have that

$$\mathbb{E}e^{sI\theta} = \mathbb{E} \prod_{z \in \Phi'} e^{s\theta\mathbf{h}_z\ell(z)} = \mathbb{E} \prod_{z \in \Phi'} G(s\theta\ell(z)),$$

where  $G(t)$  is the moment generating function of the random variable  $\mathbf{h}_z$ , i.e.,  $G(t) = \mathbb{E}e^{t\mathbf{h}_z}$ . Proceeding as before we get the following upper bound on the conditional outage probability:  $\Pr [\text{SIR} < \theta \mid (R, X) = (r, x)] \leq c(r, x)$ , where

$$c(r, x) = \min_{s \geq 0} e^{-s\ell(r)} \exp \left( -\frac{2\pi\lambda}{K} \int_{y=x}^{y=\infty} y \left( 1 - G(s\theta\ell(y)) \right) dy \right),$$

which, when averaged over the joint pdf (3) of  $(R, X)$  gives an upper bound on the outage probability. We choose the parameter  $S$  as in Section III-B. The distribution of  $\mathbf{h}_z$  for the antenna array used was found to be well-approximated by  $\Pr [\mathbf{h}_z \leq x] = (1 - e^{-\beta x}) / (1 - e^{-\beta})$ ,  $\beta = 4.63$  in simulations, and this truncated Rayleigh distribution was used to obtain the upper bounds on the outage probabilities and these are plotted in Figure 4 (right) for the small cell scenario.

#### V. CoMP

We now examine, via simulations for the square grid model, the benefits of coordinated multipoint (CoMP), where neighboring base stations can transmit a common data packet in synchrony to a desired mobile. (Analysis of CoMP for the PPP model is difficult, and is a topic of ongoing research.) Instead of partitioning space based on the Voronoi tessellation induced by the BSs, as shown in Figure 2 (left), we partition space into “virtual cells”, in which a group of four base stations serve a virtual cell, as shown in Figure 2 (right). The tessellation of virtual cells is obtained by translating the standard Voronoi tessellation by  $(1/2\sqrt{\lambda}, 1/2\sqrt{\lambda})$ . Resource allocation is achieved by means of a regular reuse policy over these virtual cells (the coloring in Figure 2 (right) corresponds to regular reuse  $K = 4$ ).

We specify two CoMP architectures: (i) *CoMP beamforming*, where the 4 BSs that form the boundary of each virtual cell cooperate and *coherently* beamform as a single virtual array to a receiver inside the cell. (ii) *CoMP spatial multiplexing* (SDMA), where this 4 BS cluster transmits data




Fig. 5. Probability of outage for square grid deployment in the small cell scenario as a function of SIR with regular reuse  $K = 4, 9$  for CoMP beamforming on the left and CoMP multiplexing on the right.




Fig. 6. Tight containment of the spatial distribution of received power (in dB scale) due to CoMP beamforming.

simultaneously to two users in the virtual cell using *zero-forcing*. This doubles the number of users in the network relative to (i). Each BS is equipped with the 8 element steerable antenna array described in Section IV, so that the 4 BS cluster effectively has a 32 element distributed array, enabling far better containment of the transmitted energy within the cell (see Figure 6), and hence significantly reduced interference. Figure 5 plots outage probabilities for CoMP beamforming and multiplexing (reuse  $K = 4, 9$ ), modeling the positions of users within the virtual cells as random.

## VI. COMPARISON OF SCHEMES

Suppose the network designer wishes to operate in the regime where the outage probability is below 0.1%. Let us examine the attainable SIR targets (and hence spectral efficiency) for which such a guarantee can be given. From Figure 3, we see that reduction in cell size significantly reduces spectral efficiency. For example, for the square grid, the attainable SIR is 11dB for large cells (corresponding to a peak data rate in a cell of 18 Mbps for the typical 20MHz LTE slice). When one moves to a small cell deployment, the SIR at 0.1% outage drops to 4.5dB at  $K = 4$ , and we must increase the reuse factor to  $K = 9$  to get back to 11 dB SIR. Thus, while we have reduced cell size by 100X, this is offset by an increase in reuse factor from  $K = 4$  to  $K = 9$ , so that the network throughput only improves by 44X.

Now consider the table in Figure 7, which shows SIR at 0.1% outage for reuse  $K = 4$  for a number of schemes. With

Scheme	0.1% outage SIR (dB)	0.1% outage rate (Mbps)	Median rate (Mbps)
Omni	4.5	9	17
Tx BF	9	15	32
CoMP BF	16	26	35
CoMP SDMA	10	16.5	29
<b>Large cell</b>	<b>11</b>	<b>18</b>	<b>32</b>

Fig. 7. Comparison of median and three nines (0.1% outage) peak *per user* rates for the different schemes discussed using reuse  $K = 4$  in the small cell regime for square deployments. The benchmark numbers for reuse 4 and large cells (with omnidirectional antennas) is given in red. Rates are computed assuming a 20 MHz LTE block.

steerable arrays, the 0.1% outage SIR for small cells improves to 9dB, which is still smaller than the 11dB benchmark for large cells with omnidirectional antennas. For CoMP beamforming, the 0.1% outage SIR is 16dB, which well exceeds the large cell benchmark. The per user peak rate improves by a factor of 1.5X (to 26Mbps for a 20MHz LTE slice), and therefore the network capacity improves 150 fold. The 0.1% outage SIR for CoMP spatial multiplexing is 10dB, which is just shy of the 11dB large cell benchmark. Therefore the per user peak rates are essentially the same as the large cell, while the network capacity sees a 200X increase because the network can serve twice the number of users per virtual cell (and thus 200 times the number of users when compared to the large cell).

## VII. CONCLUSIONS

Our work shows the importance of accounting for the “amplification” of interference as we shrink cell size, and highlights the need for tight coordination among neighboring base stations for exploiting the increased spatial reuse from picocells in order to attain high spectral efficiency at low outage. An important area for future work is the development of efficient protocols for such coordination, including adaptation to bursty demands, mobility and irregular network geometry. It is of particular interest to devise approaches for significantly reducing the overhead for the synchronization and tight coordination required for CoMP, given its critical importance in realizing the potential capacity gains from small cells.

## REFERENCES

- [1] M. Franceschetti, J. Bruck, and L. Schulman, “A random walk model of wave propagation,” *Antennas and Propagation, IEEE Transactions on*, 2004.
- [2] J. Andrews, F. Baccelli, and R. Ganti, “A tractable approach to coverage and rate in cellular networks,” *Communications, IEEE Transactions on*, 2011.
- [3] N. Blaunstein, R. Giladi, and M. Levin, “Characteristics’ prediction in urban and suburban environments,” *Vehicular Technology, IEEE Transactions on*, 1998.
- [4] M. Feuerstein, K. Blackard, T. Rappaport, S. Seidel, and H. Xia, “Path loss, delay spread, and outage models as functions of antenna height for microcellular system design,” *Vehicular Technology, IEEE Transactions on*, 1994.
- [5] T. Bai, R. Vaze, and R. Heath, “Using random shape theory to model blockage in random cellular networks,” in *Signal Processing and Communications (SPCOM), 2012 International Conference on*, 2012.
- [6] D. Stoyan, W. Kendall, and J. Mecke, “Stochastic geometry and its applications.” 1995.