

Demystifying 60GHz Outdoor Picocells

Yibo Zhu, Zengbin Zhang, Zhinus Marzi, Chris Nelson,
Upamanyu Madhow, Ben Y. Zhao and Haitao Zheng
University of California, Santa Barbara, CA 93106, USA

{yibo,zengbin,ravenben,htzheng}@cs.ucsb.edu, {zh.marzi,sciencectn}@gmail.com, madhow@ece.ucsb.edu

ABSTRACT

Mobile network traffic is set to explode in our near future, driven by the growth of bandwidth-hungry media applications. Current capacity solutions, including buying spectrum, WiFi offloading, and LTE picocells, are unlikely to supply the orders-of-magnitude bandwidth increase we need. In this paper, we explore a dramatically different alternative in the form of 60GHz mmwave picocells with highly directional links. While industry is investigating other mmwave bands (*e.g.* 28GHz to avoid oxygen absorption), we prefer the unlicensed 60GHz band with highly directional, short-range links ($\sim 100\text{m}$). 60GHz links truly reap the spatial reuse benefits of small cells while delivering high per-user data rates and leveraging efforts on indoor 60GHz PHY technology and standards. Using extensive measurements on off-the-shelf 60GHz radios and system-level simulations, we explore the feasibility of 60GHz picocells by characterizing range, attenuation due to reflections, sensitivity to movement and blockage, and interference in typical urban environments. Our results dispel some common myths, and show that there are no fundamental physical barriers to high-capacity 60GHz outdoor picocells. We conclude by identifying open challenges and associated research opportunities.

Categories and Subject Descriptors

C.2.1 [Network Architecture and Design]: Wireless communication

Keywords

60 GHz wireless, wireless beamforming, picocells

1. INTRODUCTION

We are all witness today to an explosion in mobile broadband usage. Thanks to bandwidth-hungry HD/3D video applications, industry research predicts that aggregate bandwidth requirements will increase by 1000-fold by 2020 [1]. This staggering growth is taxing mobile cellular networks past their limits and creating a capacity crisis [2]. Existing methods to increase capacity have reached their limits: spectral efficiency is slowly approaching the

Shannon limit [15]; and acquiring more spectrum provides limited capacity gains, while facing regulatory hurdles and high monetary costs [2].

To deliver the three orders of magnitude bandwidth growth required by mobile applications, we need to consider dramatically different approaches for delivering wireless content. A viable approach must meet several key requirements. *First*, it must provide total capacity that scales with the number of base stations in a given region. This means base stations must support dense deployments with minimal mutual interference. *Second*, it must increase capacity per base station by increasing amount of spectrum used and minimizing intra-cell interference between nearby mobile users. *Finally*, per-user data rates must be at least an order of magnitude higher than currently available rates, while providing robustness and reliability that match or exceed current standards in mobile and static scenarios.

Going Small. To minimize incremental deployment costs, cellular operators are exploring techniques to deploy smaller, denser cells for higher capacity, *e.g.* LTE picocells on lampposts or WiFi hotspots [3, 13], which augment existing LTE macrocells. Shrinking cell radii from several kilometers down to 100s (or even 10s) of meters is necessary to improve link quality and increase cell capacity. However, this approach is *fundamentally limited* by interference constraints for the carrier frequencies employed in today's cellular systems [30]. Specifically, at carrier frequencies of 1-5 GHz, physical form factor constraints put a hard limit on the number of antennas that can be accommodated on a picocellular basestation and a mobile device. Thus, even MIMO technologies can only mitigate interference to a limited degree, and picocell systems based on low carrier frequencies (LTE, WiFi, TV whitespaces) are unlikely to ever deliver the capacity increase we need [30].

The Promise of mmWave. But there is hope. Higher carrier frequencies, particularly millimeter-wave (mmWave) bands (30-300GHz), offer an attractive alternative. Using highly directional beams that drastically reduce interference, they enable small cells to truly deliver on their promise. Take 60GHz for example. First, the 60GHz band provides 7GHz of unlicensed spectrum for bandwidth hungry mobile applications, and is supported by the IEEE 802.11ad standard [9] (currently targeting indoor multiGigabit wireless networks) that defines multiple data rates between 385Mbps and 6.76Gbps. Second, since 60GHz has a much shorter carrier wavelength, *i.e.* 5mm, it is possible to pack large antenna arrays into relatively small form factors. A 100-element 60GHz array easily fits in no more than a square inch¹, which can be easily integrated into today's mobile devices. Third, the high density of ele-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MobiCom'14, September 7-11, 2014, Maui, Hawaii, USA.

Copyright 2014 ACM 978-1-4503-2783-1/14/09 ...\$15.00.

<http://dx.doi.org/10.1145/2639108.2639121>.

¹In theory, this is hundreds of times smaller than arrays with the same number of elements for frequencies at or below 5GHz.

ments allows very high aggregate bandwidth, while the narrow directional beams limit interference and provide potentially “unlimited” levels of spatial reuse. Finally, low-cost 60GHz radio chips are already available on the mass market [4, 5], e.g. Dell Latitude 6430u laptops can be ordered with multi-Gbps Wilocity 60GHz chips for an additional \$37.5.

Clearly, 60GHz links offer an attractive and practical solution that satisfies many of the requirements we mentioned. The goal of this paper is to show that, while most of the attention in academia and industry focuses on 60 GHz for *indoor* applications, it is also a strong candidate for delivering the 1000-fold capacity increase demanded for outdoor cellular networks in dense urban environments.

A 60GHz Picocell Architecture. We propose a picocell architecture using 60GHz links, as depicted in Figure 1. We envision picocellular base stations densely deployed on lampposts or building ledges, with each station mounting multiple antenna arrays on each of its “faces.” Properly tuned, each array can transmit independently to a different mobile user. To see the potential of this approach, consider a face which is 6x6 square-inches to a side. This can accommodate up to 36 60GHz 100-element arrays. If each array transmits at 2Gbps, the maximum data rate per face is 72 Gbps. Thus a single base station with 4 faces could support up to 288 Gbps downlink! Augmenting existing LTE cells with such a high-capacity picocell deployment provides a path to future-proofing our cellular networks.

Despite the immense potential of 60GHz outdoor picocells, so far, both industry and academia have *explicitly* precluded 60GHz from outdoor cellular designs [28, 34]. For example, Samsung’s recent outdoor mmWave attempts demonstrated a device prototype for transmission at 28GHz (a licensed band) [28]. Key concerns include propagation loss due to oxygen absorption, which peaks at 60GHz, and increase in propagation loss at higher carrier frequencies (unless highly directional antennas are correctly aligned at both endpoints). Another concern is blockage (diffraction around obstacles is more difficult at higher frequencies). These worries led to the limited use of outdoor 60GHz links in static backhauls, where manually adjusted high-gain antennas are positioned high above to avoid potential blockages.

Our Contributions. Our first contribution is to dispel common myths about 60GHz via extensive measurements and simulations, showing that the potential for 60GHz outdoor picocells can be realized without bending the laws of physics. Unlike prior efforts that focus on propagation models [22, 14, 25, 33], we experiment with off-the-shelf 60GHz radios, and study key questions such as transmission range, robustness to blockage and user motion, and spatial reuse. We explore how these issues are affected by future hardware upgrades (especially in terms of increasing the number of antenna elements). Unlike proposals for cellular systems in other mmWave bands [28], design and evaluation of our *outdoor* network can leverage extensive efforts already put into 60GHz PHY hardware and standards for *indoor* networks. Our second contribution is to point out the cross-layer architectural innovations required to truly deliver on this potential, and identifying an interdisciplinary research area spanning the PHY, MAC and networking layers.

We now provide a quick preview of the myths we dispel.

- **Myth #1: 60GHz’s range is too small for outdoor picocells.** Within FCC’s power regulations, and *taking into account oxygen absorption and heavy rain*, 60GHz offers coverage area that exceeds 130m for 385Mbps and 79m for 2+Gbps. A propagation loss of 36 dB/km due to oxygen absorption and heavy rain makes it difficult to achieve kilometer-long links. But these factors only produce a 3.6 dB loss in signal strength at 100 m.

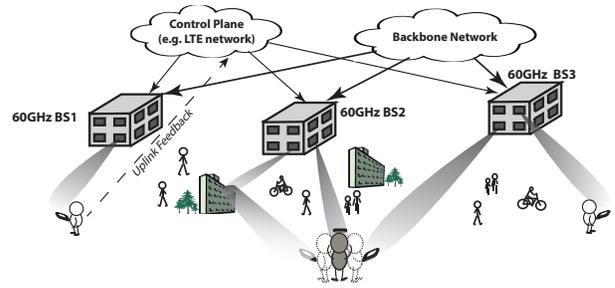


Figure 1: A potential 60GHz picocell architecture for outdoor mobile broadband augmenting existing LTE networks for significant increase in capacity. Each base station contains many antenna arrays, each containing a large number of elements. Leveraging an LTE uplink as the feedback loop, each base station can track and form highly directional beams simultaneously to multiple users. The beams follow users through direct or reflected paths, creating high-bandwidth downlink transmissions. Users can form beams in the reverse direction to build 60GHz uplinks, or use existing LTE uplinks.

- **Myth #2: 60GHz is highly sensitive to blockage, hence robust outdoor connectivity is not feasible.** Indeed, 60 GHz links can be blocked easily by a human body (20-50 dB of loss), and we do not have enough transmit power to burn through obstacles. However, an electronically steerable antenna array can point its beam towards one of several reflected paths available off of outdoor surfaces. The additional reflection loss of 5-7 dB can be handled within our desired link budget. We show that this, together with the concept of “picoclouds” (multiple base stations serving a given mobile user), can provide robust connectivity even in obstacle-rich urban environments.
- **Myth #3: User motion will easily break 60GHz links.** Detailed measurements show that electronically steerable antennas in off-the-shelf radios can easily keep up with outdoor pedestrian motion, and only need to realign once every few seconds to maintain high rate connections. In essence, the time constant of electronic adaptation is small compared to the time constants of human mobility.
- **Myth #4: Interference will significantly limit network capacity for small cells.** While this is indeed true at lower carrier frequencies [30], 60GHz’s small interference footprint means that 60GHz picocells can be densely deployed with significant overlap but minimal interference, leading to a significant boost in spatial reuse. The reduced interference results from two factors: highly directional links that reduce interference between nearby users, and oxygen absorption and blockage from buildings that severely attenuate interference from distant cells.

In the rest of the paper, we elaborate on these results, and show that 60GHz outdoor picocells are indeed a feasible approach for delivering orders of magnitude increase in network capacity.

Limitations and Future Work. As an exploratory study, our work has several limitations that need to be addressed in future work. *First*, the preceding myths correspond to *physical* properties of 60GHz communication. Thus, dispelling them amounts to saying that, while the characteristics of 60 GHz are very different from those at lower carrier frequencies, there are no fundamental barriers in realizing its immense potential. However, as we discuss in §6, there are significant design challenges (and hence research opportunities) across a range of research disciplines, in-

cluding hardware, signal processing and cross-layer protocol design. *Second*, our measurement experiments are constrained by limitations of available 60GHz hardware, hence these results are supplemented by projections and simulation-based predictions (see §6 for a discussion on the role of such techniques for future research). *Third*, while picocellular deployments may be used to add capacity in a number of environments, our focus here for 60GHz picocells is restricted to dense urban environments with high availability of reflecting surfaces such as walls. These are also locations likely to first experience full utilization of LTE capacity. Some of our conclusions may not apply, for example, to picocells covering large open spaces, where lower carrier frequencies may be a better choice.

2. BACKGROUND AND METHODOLOGY

In this section we introduce background on 60GHz communications and describe our experimental methodology.

2.1 60GHz Background

Directional Transmissions. To overcome attenuation and reduce interference, 60GHz links use beamforming to concentrate transmission energy in a desired direction. Today, 60GHz beamforming radios are readily available and affordable, either using horn antennas or antenna arrays. Radios with horn antennas use a flaring metal horn to direct radio waves, but require mechanical rotators to achieve fine-grained directional control [42]. On the other hand, radios with antenna arrays implement directional transmissions using electronic beamsteering [26, 4], employing signal processing at RF and baseband to adjust beam directions with minimal delay (*e.g.* a few *ns*). Such instantaneous beam adaptation makes antenna arrays an ideal fit for picocells.

The performance of antenna arrays scales with the number of antenna elements. The more the elements, the stronger the beam and the less the signal leakage. Today’s low-cost 60GHz consumer products use 2×8 arrays (16 elements), while researchers have demonstrated 60GHz 16×16 array prototypes (total of 256 elements) [27]. Furthermore, recent breakthroughs make it computationally feasible to adapt even larger arrays (1000+ elements, but still fitting in palm-sized form factors!) [31], which puts us in a position to take advantage of future hardware advances to synthesize even narrower beams to further increase spatial reuse.

Oxygen Absorption and Rain Loss. At longer ranges, the advantages of directionality are offset by signal strength losses from oxygen absorption, which peaks at 60GHz, causing a loss of about 16 dB/km [6, 40]. While this makes realizing long-range (kilometers) 60 GHz link very difficult, it is actually an advantage for picocellular networks because it significantly enhances spatial reuse while having relatively small impact on the link budget at the link ranges of interest to us, as also confirmed by our testbed experiments in §3. For example, the net loss due to oxygen absorption and very heavy rain of 2 inches/hour is 36 dB/km, which works out to a modest 3.6 dB for a typical link range of 100m.

Impact of FCC Regulations. The emitted power at 60GHz is constrained by the Federal Communications Commission (FCC) as follows: the transmit power plus beamforming gain, termed Effective Isotropic Radiated Power (EIRP), is limited to 40 dBm. For antennas that are extremely directional (*e.g.*, using parabolic dishes), the FCC has relaxed its specifications to permit further increase of EIRP [8], but we are unable to take advantage of this in the operating regime of interest to us: our architecture is based on highly directional, electronically steerable beams, but the beamwidth is still too large for the new FCC regulations to permit an EIRP higher

than 40 dBm. Thus, increasing the number of elements in the transmit array does not improve link range if we already at an EIRP of 40 dBm, hence enhancements in link range must come from the receiver, by decreasing the receiver’s noise figure or by increasing its antenna directivity. Of course, increasing the number of transmit elements has a big impact on *network* performance, narrower transmit beams reduce interference and enhance spatial reuse.

Leveraging IEEE 802.11ad. The IEEE 802.11ad standard for 60GHz [9] defines a number of modulation and coding schemes to adapt to different range-rate tradeoffs. In our envisioned picocellular architecture, we leverage this 802.11ad *physical layer* on the downlink, and show that link ranges of the order of 100m are attainable (see Table 1). Of course, *our MAC and network layers will be different from that of 802.11ad*, given the need to adapt to a far more dynamic outdoor environment than the indoor networks the 802.11ad standard was currently intended for.

Data rate	RX sensitivity	4×4 RX	8×8 RX	10×10 RX
385 Mbps	-68 dBm	63m	114m	138m
1.155 Gbps	-64 dBm	41m	77m	94m
2.31 Gbps	-61 dBm	30m	57m	70m

Table 1: Outdoor range/rate tradeoffs with 802.11ad PHY

The link budget calculations resulting in these ranges use conservative assumptions (measurements from our testbed lead to more optimistic predictions). The receiver sensitivity for each data rate (*i.e.*, minimum received power required to sustain that data rate), obtained from the 802.11ad standard, already accounts for large values of implementation loss (5 dB) and noise figure (10 dB). The link budget accounts for the FCC EIRP restriction of 40 dBm and loss due to free space propagation and oxygen absorption. On top of that, we add a link margin of 15 dB, which is more than enough to account for reflection loss (in case we are using a bounce rather than a LoS path, which typically leads to at most 5-7 dB loss) and for rain (even for heavy rain, the additional loss on top of oxygen absorption is about 20 dB/km, hence the rain loss over 100m is 2 dB). That is, the receiver sensitivity as a function of range R (in meters) is computed as

$$P_{RX}(dBm) = EIRP(dBm) + G_{RX}(dBi) - L_{ploss}(R) - L_{margin} \quad (1)$$

where $G_{RX}(dBi) = 10 \log_{10} N_{RX}$ is the receiver antenna directivity as a function of the number of receive elements N_{RX} (we consider $N_{RX} = 16, 64, 100$ for the $4 \times 4, 8 \times 8$ and 10×10 arrays). The path loss L_{ploss} in dB is

$$L_{ploss}(R) = 10 \log_{10} \frac{16\pi^2 R^2}{\lambda^2} + \alpha R \quad (2)$$

with $\alpha = 0.016$ dB/m (16 dB/km) accounting for oxygen absorption, and $\lambda = 0.005$ m at 60 GHz.

2.2 Measurement Methodology

Unlike prior efforts that focus on propagation models using custom equipment [14, 22, 25, 33], we experiment with commercial off-the-shelf 60GHz radios. We study key issues regarding the feasibility of 60GHz picocells, including range, sensitivity to motion and blockage, and interference footprint.

Measurement Environments. We perform experiments in two types of environments.

- *Controlled Environments.* These are environments without any uncontrolled pedestrian. We use these measurements to evaluate the impact of individual factors in isolation. To emulate a picocell, we place a radio on a second floor balcony (6m high), representing a basestation fixed on a lamppost (Figure 2). A second radio represents the end user.

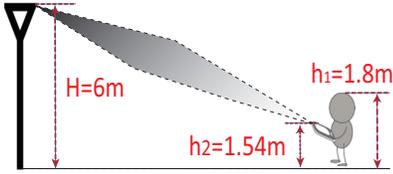


Figure 2: Controlled environment: a picocell system with 60GHz radios at both basestations and mobiles.



Figure 3: Real-life environment: a downtown street area with random pedestrians.

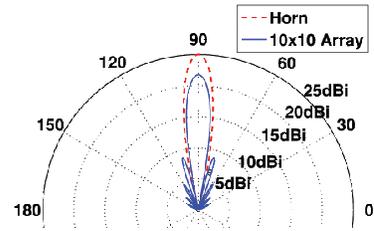


Figure 4: Radiation patterns of our testbed’s horn antenna and a 10x10 rectangular array.

- *Real-Life Environments.* We also performed experiments in a downtown street walkway (Figure 3) and a campus plaza during busy hours with uncontrolled pedestrian. The basestation is placed on a second floor balcony (6m high).

60GHz Radio Platforms. We consider two types of 60GHz radios with antenna arrays, representing today’s off-the-shelf hardware and high-performance radios for future deployments. By comparing their results, we also seek to identify issues that can be addressed by reasonable hardware upgrades, and those that are fundamental to 60GHz picocells and hence require careful network design. For both platforms, each radio operates on a single 60GHz channel of width 2.16GHz.

- *Wilocity 2x8 Antenna Arrays.* We use pairs of Dell 6430u laptops (as users) and D5000 docks (as basestations), designed for indoor usage. Each has a Wilocity 60GHz radio chipset with a 2x8 rectangular antenna array, operating according to IEEE 802.11ad. Because Wilocity does not provide any information on the beamwidth, we apply the antenna array theory [18] to compute the lower bound on the beamwidth: the horizontal 3dB beamwidth is 13° and the vertical 3dB beamwidth is 60°. Each radio adapts link rate between 385Mbps and 3.85Gbps, and maintains directional connectivity over time (e.g., if the LoS path is blocked, it seeks a reflected path and beamforms along it). Currently, the radio does not provide external control of link directions, hence we rely on its built-in algorithm to manage link directionality. The radio does not report signal strength values but link rates, so we monitor link performance using fine-grained reports on TCP/UDP throughput and link rate.

- *Emulating 10x10 Antenna Arrays.* There are no suitable 60GHz radios with large antenna arrays on the market. Instead, we emulate the performance of a large antenna array using a horn antenna. For experiments that do not consider interference, this can provide accurate results. This is because 60GHz link performance is largely determined by directionality and signal patterns of the main beam lobe². In this regard, large antenna arrays and horn antennas are very close except for differences in antenna gain.

Specifically, we use two HXI Gigalink 6451 radios with 0dBm transmit power and 1Gbps link rate, operating on a proprietary (non-802.11ad) configuration. Each radio has a horn antenna of 10° 3dB beamwidth and 25dBi gain. Its main lobe pattern closely align with that of a 10x10 array with 1dBi elements with 21dBi gain (Figure 4). We set each radio on a mechanical rotator and adjust its beam direction in units of 0.5°. Each measurement records the received signal power level (RSS), repeated once every 50ms. The measured noise level is -80dBm.

²Antenna arrays produce side lobes (*i.e.* unwanted transmissions in undirected directions) aside from the main lobe. The side lobes are at least 13.26dB weaker than the main lobe and thus not responsible for data communication [20].

Since HXI radios use a proprietary physical layer and only support a single data rate, they cannot fully capture the performance of a picocell operating based on 802.11ad. However, we can extrapolate the attainable range to 802.11ad settings: increasing the EIRP to 40 dBm from the 25 dBm in our testbed leads to a 15–4=11dB increase in RSS (counting the 4dB gain difference of horn antenna and 10x10 arrays). This provides us with a practical means of using our testbed to verify link budget calculations as in eq. (1), and also to predict the supported data rate for each RSS measurement following the 802.11ad receiver sensitivity table. Specifically, we use the 802.11ad receiver sensitivity table to derive the required SNR for each data rate, and use it to derive the data rate supported by each RSS measured by our testbed. A similar methodology has been used in prior work on 60GHz systems [42].

3. SINGLE-USER FEASIBILITY STUDY

Using testbed measurements, we perform initial evaluation on the feasibility of 60GHz picocell systems, focusing on dispelling the four myths discussed in §1.

Assumptions. Due to hardware limitations, our measurement study makes a few assumptions. *First*, our experiments emulate a single array forming a beam to serve a single user, using either 2x8 WiloCity arrays at both ends or 10x10 arrays at both the base station and the user. Our envisioned picocellular base stations would have multiple such arrays, one for each user, as discussed in §1. *Second*, we assume the target user is stationary (sitting or standing), although her body may rotate or move a short distance due to natural posture changes. *Finally*, basestations do not coordinate in these measurements. We study the benefits of coordination via simulations in §4.

3.1 Measurements in Controlled Environments

We begin by verifying the framework for the link budget calculations in §2.1. Our measurements use the HXI radios since they report detailed receiver RSS values. Figure 5(a) plots the RSS values at different link distances under clear weather, along with the well-known Friis model [16] on 60GHz free space propagation with and without oxygen absorption (16dB loss per km). The results show that 60GHz outdoor LoS propagation does follow the classical Friis model with oxygen absorption loss, *i.e.* eq. (2).

3.1.1 Range & Throughput

Our first task is to evaluate the range and throughput of 60GHz transmissions in the context of outdoor picocells.

Picocell LoS Range. Figure 5(b) lists the measured LoS range of the Wilocity 2x8 array under different link rate targets. On a clear day, it can reach 23m at 385Mbps, and 9m at 2.31Gbps. While this is clearly insufficient for 60GHz outdoor picocells, we can easily extrapolate beyond the hardware limitations of the Wilocity platform, which is designed for indoor short-range usage and has an

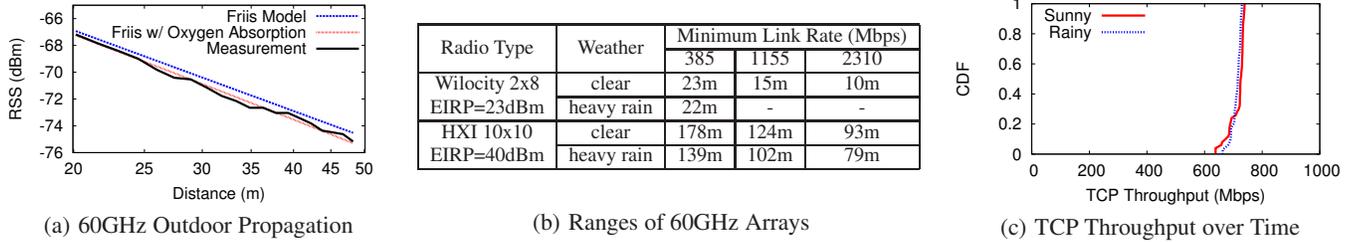


Figure 5: Measurements on 60GHz range and throughput. (a) The RSS results collected by HXI radios show that 60GHz outdoor propagation follows the Free-Space Friis model with 16dB/km oxygen absorption loss. (b) Downlink LoS ranges when 60GHz radios operate in the 802.11ad single carrier modulation (SC) setting, with 2×8 and 10×10 antenna arrays. The Wilocity 2×8 results are based on actual measurements, while the HXI 10×10 results are projected by increasing EIRP to 40dBm and mapping a data rate to each measured RSS value based on the 802.11ad specification. (c) CDF of per-minute TCP throughput measured over 10 days using the Wilocity radios.

EIRP of 23dBm. When we increase the EIRP from 23dBm to 40dBm, the path loss formula in (2) can be used to translate the 17 dB increase in signal strength to a corresponding increase in range. Because of the mixture of free space loss of $1/R^2$ and exponential decay of $e^{-\alpha R}$, the actual range increase depends on the initial value of range. While this is easy enough to calculate numerically, it is instructive to separately address these two effects via an example. Specifically, can we get a data rate of 385 Mbps at 100 m using 2×8 arrays as in the Wilocity radios in heavy rain (e.g., 2 in/hr, where the loss due to rain and oxygen absorption is about 36 dB/km)? The exponential loss is 3.6 dB at 100 m, and ignoring the baseline loss under the clear conditions for our measurements at the starting range of 23 m, we still have $17 - 3.6 = 13.4$ dB signal strength increase to compensate for free space loss, leading to $13.4/2 = 6.7$ dB, or 4.67X, increase in range relative to the baseline of 23 m. This yields a range of 100+ m. The point of this calculation is to show that, for the moderate ranges of interest, oxygen absorption and rain have a relatively small effect on the link budget.

In order to quantify the effect of using larger arrays, we project the measurements using our HXI radios, which we using as a proxy for 10×10 arrays, by adding 15 dB to the EIRP (going from 25 dBm to 40 dBm), and subtracting 4 dB in receiver directivity (25dBi horn replaced by 10×10 receive array with 1 dBi elements). If we keep the EIRP fixed, then increasing the number of transmit elements does not help the link budget (it only decreases the emitted power per element), hence the source of performance gain is the increase in receiver directivity. Based on this projection, the range in clear weather is as high as 178 m at 385 Mbps.

Figure 5(b) also lists transmission ranges under heavy rain (0.38in/h). The measured range of Wilocity radios remains the same while that of 10×10 arrays, projected from our HXI measurements, drops to 139m at 385Mbps and 102m at 1.15Gbps. This is as expected, since the additional loss due to rain impacts longer links. However, the range (100+m) is still sufficient to support picocell operations.

Link Throughput over Time. We also measured 60GHz link throughput over 10 days (February 2014) to examine its long-term performance in outdoor environments. The weather condition for the first 9 days was clear and the last day we had heavy rain (0.38in/h). We used the Wilocity radios, set the link distance to 20m, close to its maximum range of 23m, and recorded per-minute TCP throughput using Iperf. Figure 5(c) plots the CDF of the per-minute TCP throughput of the first 9 days and that of the last 1 day. We see that the Wilocity 60GHz link offers 638+Mbps TCP throughput continuously over time (and 700 Mbps 80% of the time). The minor rate fluctuations (638-700 Mbps) are randomly distributed

across the measurement period, and may be the result of occasional structure vibrations caused by wind and/or pedestrians walking on the balcony. We see no difference in throughput between clear weather and heavy rain, which confirms that, at short distances, 60GHz is highly robust to weather conditions. We also confirm this observation using the HXI radio measurements. At 50m, heavy rain (0.38in/h) only introduces 0.86dB RSS loss, which did not introduce any degradation to the link.

3.1.2 Robustness to Blockage

Outdoor 60GHz transmissions experience blockage from various obstacles including buildings, trees, pedestrians and vehicles. Our experiments focus on pedestrians, because they are moving obstacles commonly seen in outdoor settings. We also examine the effectiveness of switching to reflected paths on addressing blockage.

Impact of Pedestrian Blockage. Our experiments show that when holding a 60GHz radio and facing away from a basestation, the user herself will block the LoS signal. This observation triggers another important question: “will the connection to a user, either LoS or via reflection, be constantly blocked by nearby pedestrians?” To answer this question, we measure for both radios the LoS peer impact region, *i.e.* the area around a user where a peer of similar height blocks the LoS path between the basestation and the user. For Wilocity radios, we placed RF absorber around the radio to prevent it from switching path. We also derive the impact region of reflection paths following the same method. Figure 6 shows the LoS impact regions for the HXI radio at different link distances (the Wilocity radio result is similar and thus omitted for brevity).

We see that the observed impact regions are all small areas with a narrow cone shape, and the width (y-axis) is significantly smaller than its length (x-axis). The region also enlarges linearly with the link distance. This is because the basestation is placed in a high location (6m in our experiments, representing a lamppost placement) and the peer must be quite close to the user to block the signal.

Overall, our results imply that in outdoor environments, surrounding static peers can potentially block paths to a user. The impact, however, is limited because the peer has to be in a small region near the user. On the other hand, our measurements in busy walkways (§3.2) show that groups of moving pedestrians have a heavier impact on connectivity. However, such blockage can be effectively addressed by either switching to another reflection path (discussed next) or connecting to another basestation (see §3.2).

Handling Blockage via Reflection. 60GHz signals are reflected by most surfaces with losses of the order of 5-7 dB [21]. Thus, when the LoS path (or an existing path) to a user is blocked, the

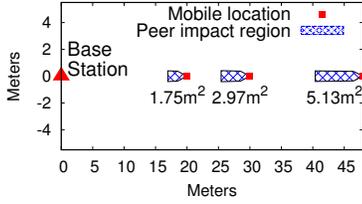


Figure 6: Impact of human blockage in terms of peer impact region for HXI radio. The base station is positioned at (0,0), and receiver is at 20m, 30m and 48m away.

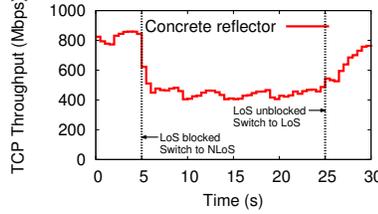


Figure 7: Measured TCP throughput as a Wilocity 60GHz link switches between different paths.

Radio	Path Length	Flat steel, Aluminum, Plastic board	Exterior Building Walls		
			Brick, Paint, Whitewash	Coarse paint	Coarse concrete
Wilocity (2x8)	10m	0%	0%	16%	16%
	15m	0%	17%	33%	33%
HXI (10x10)	30m	0%	17%	33%	46%
	50m	0%	19%	25%	25%
RSS Loss		0dB	3dB	4.5dB	5dB

Table 2: The reflection loss of eight different materials measured using both Wilocity and HXI radios.

basestation can adjust link directionality to create another reflection path. In fact, our measurements show that Wilocity radios can already switch to reflected paths in real time upon LoS blockage (see Figure 7).

Next, we study the impact of reflection on link performance. We experimented with eight surfaces, commonly seen in outdoor street environments. These include three flat surfaces and five exterior building surfaces. We repeated our experiments with different incident angles (10° - 80°), and found that the angle has a very minor impact on the result³. For brevity we only show the results of 45° . Table 2 shows the normalized throughput loss (compared to a LoS path of the same length) for both 2×8 and 10×10 arrays at different path lengths. The HXI radio experiments also report the amount of RSS loss due to reflection. For both radios, the first three flat (and smooth) materials are perfect reflectors with negligible signal loss. The five building surfaces introduce 3-5dB signal loss and up to 46% throughput loss. The throughput loss varies with the link distance while the signal loss (in dB) remains constant. We noticed that in our results with the HXI radios the measured reflection losses were slightly smaller (*i.e.* better) than those reported in [21]. This might be due to the subtle difference in the reflection materials. Overall, these results show that reflection paths are good alternatives to LoS paths, offering ample opportunities for picocell coverage.

3.1.3 Robustness to User Motion

In practice even a “stationary” user will naturally make local movements, such as change her posture, or move a small distance (*e.g.* up to a few meters). These minor movements could change 60GHz link directionality and thus performance. We now present some initial results on their impact on link performance and implication on picocell design.

Specifically, we seek to examine whether 60GHz radios can realign their arrays to maintain a LoS connection. To do so, we place absorber around the transmitter radio to prevent the use of any reflection path, and measure TCP (or UDP) throughput continuously during the motion to study its impact. We consider four types of user motion: rotating the body, moving closer/further to the basestation⁴, moving perpendicularly to the basestation, and moving in a circular motion. Each experiment starts with the basestation and user radio directly facing each other and separated by 10m (Wilocity) or 40m (HXI). While the absolute results depend on factors like link distance and user body shape, our goal is to understand

the general trend and its implication, as well as how the two 60GHz radios react to these motions.

Wilocity 2×8 Array Results. Since 802.11ad already enables radio realignment via beam tracking, we expect that the Wilocity radios can realign with each other to a certain degree. Thus the goal of our measurements is to understand the speed of realignment and the robustness against user motion.

Because Wilocity radios do not expose any information on beam directionality, we first performed a set of UDP experiments to understand their array alignment behaviors. In each experiment, we first let the basestation and user radios face each other and record UDP throughput, and then move/rotate the user radio locally for 1-3s and stop for 15s (to allow re-alignment).

From both UDP and link rate reports, we see that the low-cost Wilocity radios already adapt beam direction on-the-fly to cope with changes in link directionality. But the adaptation has two limitations. *First*, the adaptation is not instantaneous: in our experiments, it takes 2-7s for UDP/link rate to fully recover. Such delay can come from the time for adjusting the array and that for coordinating both ends. *Second*, the amount of beam adjustment is limited, *i.e.* the array cannot tune to all the directions, possibly due to the hardware enclosure and/or low-cost array design.

With these observations in mind, we now study the impact of user motion on TCP performance. Each experiment starts from the two radios being perfectly aligned with TCP throughput at 900Mbps. *First*, we experimented on user rotation at different rotation speeds ($40^\circ/s$ - $100^\circ/s$). We found that until the rotation angle reaches 85° , the TCP throughput remains at 600Mbps. After that, the link breaks down.

Next, for the other three patterns, Figure 8 shows the TCP throughput during motion, as a function of the receiver’s deviation distance from the original position. We make two observations. *First*, a user moving further/closer or in a small circle at moderate speeds (Figure 8(a)-(b)) lead to a minor degradation in TCP throughput. *Second*, moving perpendicular to the LoS has more impact since it introduces larger/faster directionality changes. During the movement, TCP throughput fluctuates and gradually degrades to 400Mbps as the deviation reaches 8m. Specifically, when the user moves at 1.5m/h, the throughput first drops due to array misalignment, then rises as both radios align their arrays, and repeats the same pattern. By examining the patterns of peak/valley across all the speeds tested (1.5-5mph), we see that the radio alignment frequency aligns with the above UDP results, *e.g.* once every 2-7s.

In summary, our experiments on Wilocity radios show that today’s low-cost 60GHz radios already offer online array alignment to handle moderate user motions. However, these radios are designed for quasi-static indoor environments, and the speed and amount

³Our results align with existing results: for thick materials (≥ 3 cm), the reflection loss decreases as the incident angle increases [21].

⁴Since the basestation is placed higher than the user radio, moving further/closer to the basestation does change link directionality.

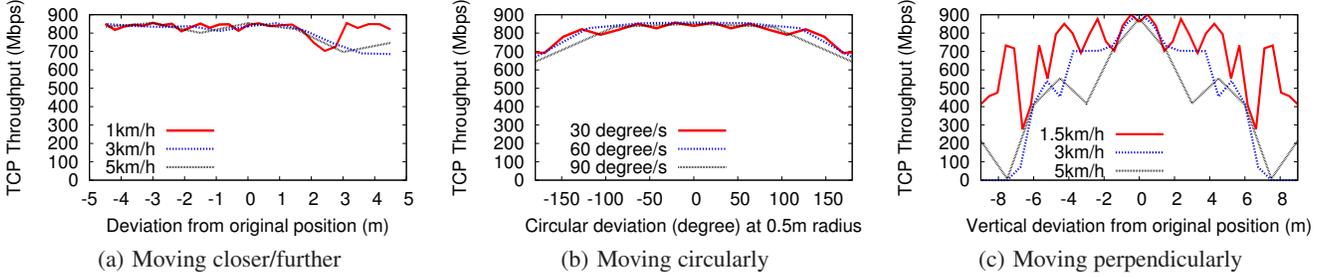


Figure 8: Measured throughput using our Wilocity testbed (2x8) when the receiver moves in three different motion patterns.

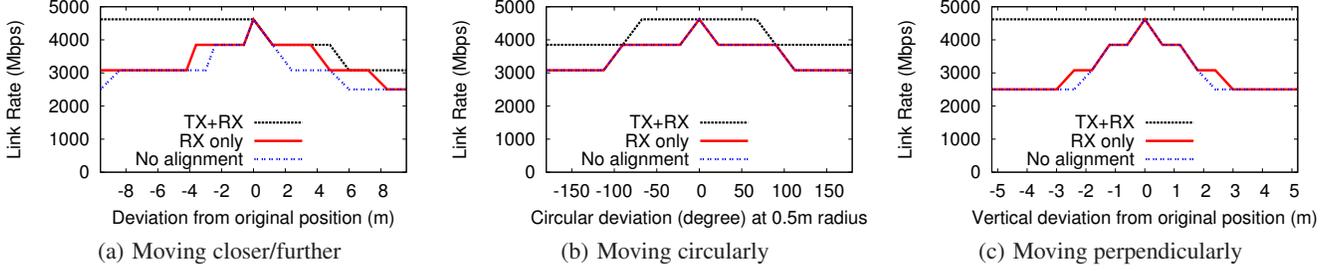


Figure 9: Measured RSS values using our HXI testbed (10x10) when the receiver moves in three different motion patterns. We consider three types of alignment methods: no alignment, partial alignment (RX only) and full alignment involving both TX and RX (TX+RX).

of their beam adjustments are not enough to handle certain motion patterns that introduce fast/large changes in link directionality.

Alignment Requirements for 10x10 Arrays. The above discussion triggers another question: “As the link becomes more directional by using larger arrays, how fast should array alignment take place to handle user motion?” As already discussed, the horn antennas on the HXI radios serve as a good proxy for 10x10 arrays. Hence, to answer this question, we mount the HXI radios on mechanical rotators (with a precision of 0.5°) and manually adjust beam directions as the user moves. Each experiment starts with the two radios perfectly aligned. As the user moves, we measure the RSS values of three scenarios: adjusting both TX and RX arrays (full alignment), only adjusting RX array (partial alignment), and no alignment. From these results we estimate the alignment speed required to support typical user motion. Figure 9 plots the link rate projected from the RSS values measured during motion. We make two key observations.

- *Full alignment is required to handle user motion*, which is particularly true for movements towards and perpendicularly to the basestation. It involves both ends and requires a coordination channel between the basestation and the user.
- *Re-alignment once every second is sufficient for local movements.* Again perpendicular movements lead to the fastest signal degradation. Without alignment, the transmitter and receiver beams deviate from each other gradually, e.g. 3.6° when the user moves 2.5m. The RSS drops by 2.47dB and the link throughput to half. At a moderate motion speed of 3-5mph, moving 2.5m takes 1.8s-3s. Thus at a link distance of 40m, re-aligning once every second should maintain quality connections during local movements. We also study the re-alignment frequency at other link distance values, and obtain similar findings.

3.1.4 Interference and Spatial Reuse

We now study the spatial reuse provided by 60GHz links. While averaged performance over large-scale deployments is required to

accurately quantify spatial reuse, we use testbed measurements to explore several worst-case interference scenarios when 60GHz picocells are deployed in a street environment. We also examine how the choice of arrays affects the degree of spatial reuse.

Interference Scenarios. Intuitively, two directional links interfere when their directions are too similar to each other. Hence, in the context of 60GHz picocells, we consider in our experiments three example “worst-case” downlink interference scenarios, shown in Figure 10. All three scenarios involve 2 basestations (BS_a , BS_b) deployed on lampposts at the edge of the sidewalk next to the street, each serving a user (a and b) on the sidewalk, using the same 60GHz channel. In each scenario, the interference condition depends heavily on the angular separation between the two links, which is defined by the distance between the two basestations (D), user a 's horizontal distance to BS_a (d) in scenario 1, and the sidewalk depth (l) in scenario 2 and 3. In our experiments, we consider $l = 7m$, modeling major urban streets like those in Manhattan.

Since these interference measurements require at least two links, we used Wilocity radios in our measurements (we only have two HXI radios or one link due to cost constraints). To estimate the performance of larger arrays (thus narrower beams), we mount a metal horn-shaped case on the Wilocity device to reduce its beamwidth to roughly approximate a 10x10 array. However this also eliminates side-beams so our larger-array results are rather optimistic. Later in §4 we use network simulations to further study this issue. Finally, for all the experiments, we place radio absorbers to ensure that both links do not change their path directions.

Findings. We plot in Figure 10(d-f) the normalized TCP throughput loss due to interference. We make two key observations. *First*, the low-cost 2x8 arrays have wide beams that result in large interference. The two basestations must be widely separated to avoid inter-basestation interference: 20-22m for scenario 1 and 2, and $\gg 22m$ for scenario 3. But upgrading to a larger array can potentially reduce basestation separation to 12-14m for scenario 1 and 2, and 22m for scenario 3, enabling dense basestation deployment.

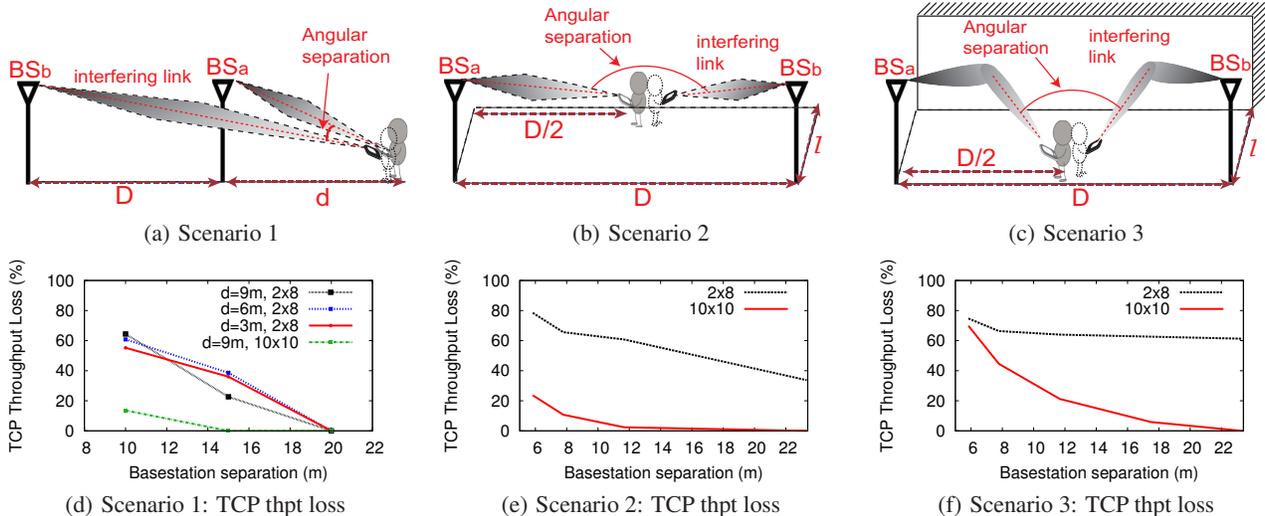


Figure 10: Three sample “worst-case” interference scenarios for a sidewalk environment and the TCP throughput loss due to interference. Two 60GHz basestations are deployed on lampposts (6m in height) on the edge of the sidewalk, each serving a user on the sidewalk.

Second, the use of reflected paths (for overcoming blockage) creates diverse interference conditions beyond those of LoS. In certain scenarios, even using a narrow beam array may not eliminate interference. This calls for *coordinated scheduling*, where basestations coordinate user scheduling to proactively avoid heavy interference.

3.2 Measurements in Real-life Environments

Our experiments took place in two sites: a downtown street walkway area and a campus plaza area, both during busy hours with many moving pedestrians. These “in the wild” experiments differ from the “controlled” experiments as follows. *First*, the user is surrounded by *random moving pedestrians* that create unpredictable blockage. *Second*, we use Wilocity radios and leverage its built-in array adaptation mechanisms to maintain link connectivity. *Third*, we deploy multiple basestations around the area to study whether switching between basestations can effectively improve connectivity. At both sites, we placed basestations on a 2nd floor balcony of 6 meters in height, and experimented at 10 user locations (6 at downtown, 4 at campus) and 2 basestation locations. At each location, we let the user face the basestation(s), thus her own body did not block the LoS path. We then recorded TCP throughput and link layer rate (reported by the radio driver) for 1 hour per location.

Single user, single basestation. We first consider the scenario where a single user is being served by a single basestation. To illustrate the impact of pedestrian blockage on link connectivity, we plot in Figure 11(a)-(b) the TCP and link rate traces at two sample user locations at the edge of the walkway, with a roughly 12m link distance. In Figure 11(a), the connection experienced outage (both link rate and TCP rate drop to 0) several times during the first 250s. During this time period, we observed a continuous crowd of pedestrians passing by, creating prolonged blockage to the user. After that, the link rate fluctuates (due to transient pedestrian blockage and subsequent rate adaptation), but the TCP throughput remains stable⁵. In fact, when a small group of 2-3 pedestrians pass by the user, they only create transient blockage. The link rate drops largely but quickly recovers.

⁵The Wilocity TCP throughput is capped to 1Gbps, due to the 1Gbps Ethernet interface that connects the radio to the laptop.

In Figure 11(b), between minute 13 and 23, there is a 2-minute outage followed by a period of rate fluctuations. This 10-minute period was in between two class sessions, thus large groups of students passed by our user location. Different from the last example, the link was able to recover and sustain high TCP throughput after the two-minute outage, despite the prolonged blockage. This is because the radios switched to a reflected path (via a nearby building wall) to avoid blockage, and the frequent fluctuation was the result of rate adaptation to mitigate reflection loss. As a result, the link didn’t fail repeatedly like in the previous example. In the rest of the hour, the pedestrian density was moderate and only caused transient blockage thus zero outage.

Figure 11(c) shows the link availability, *i.e.* (1-outage), over the busiest 300s across all the measured user locations (thus the period with the lowest availability), for the two basestation locations. The availability varies largely across user locations, and can be as low as 50%. The low availability scenarios are those where the user was on a crowded walkway and blocked by many moving pedestrians.

Single user, multiple basestations. We then examine scenarios where two basestations coordinate to serve a single user. To do so, we place the second basestation 10m-20m away from the first basestation⁶. Because Wilocity radios do not allow clients to switch between basestations on-the-fly, we emulate real-time switching by connecting a second receiver (co-located with the first receiver) to the second basestation on a different 60GHz channel. This allows us to simultaneously measure for each location its link availability to each basestation, and calculate the link availability when enabling real-time switching. Results in Figure 11(c) confirm that serving a user with a picocloud consisting of multiple basestation can effectively boost link availability.

3.3 Summary of Observations

Our detailed measurements dispel the four common myths about 60GHz outdoor picocells listed in §1. We now summarize our findings and their implications on picocell design.

- **Range:** While today’s low-cost radios have limited outdoor range due to its low EIRP and small array size (designed for indoor sce-

⁶We did not observe any notable impact of the separation distance.

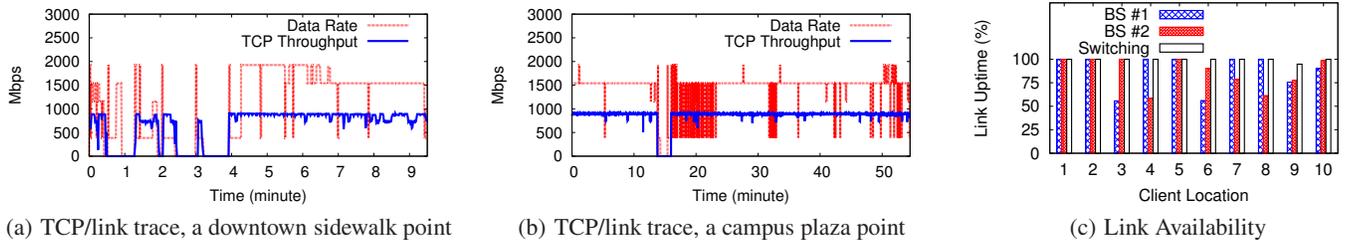


Figure 11: The impact of pedestrian blockage in real-life environments. (a) A 10-minute segment of TCP/ link rate traces at a downtown sidewalk location. (b) A full 1-hour trace at a campus plaza location. (c) Link availability at 10 locations w/ and w/o basestation switching.

narios), larger antenna arrays, *e.g.* 10×10 , can support a range up to 100+m at 385Mbps even in heavy rain conditions.

- **Blockage:** The user body, other pedestrians, and other obstacles can all block connections to a user. Hence it is critical to maintain a rich path inventory for each user by continuously tracking both LoS and various reflection paths. When necessary, users can be served using a picocloud comprising of multiple basestations (supported by results in §4).
- **User Motion:** Local user motion changes 60GHz link directionality and degrades performance, unless both endpoints realign their arrays in real time. Since they were designed for quasi-static indoor settings, realignment on Wilocity radios can take up to 7 seconds, making them less robust against certain motion patterns. For larger arrays (10×10) with narrower bandwidth, re-alignment frequency needs to be at least once per second.
- **Spatial Reuse:** Today’s low-cost 2×8 arrays have wide beams that limits the degree of spatial reuse. Using larger arrays up to 10×10 can significantly reduce interference footprints, but the large variety of path directions can still lead to considerable interference in certain scenarios. Handling these requires intelligent user scheduling and basestation coordination. While we will provide initial insights on these issues in §4, far more research is needed in this area.

4. SIMULATION RESULTS

In this section, we perform detailed simulations to further examine 60GHz picocells in terms of blockage and interference.

Simulation Setup. We consider a street in a typical urban environment: two basestations are deployed on opposite sides of the street (at lamppost height of 6m), forming a picocloud to serve users on both sidewalks. Figure 12 plots the geometry of a typical Manhattan street observed from the Google map. Our simulation considers pedestrians, cars and trees as obstacles that can potentially block or attenuate 60GHz links. We model each pedestrian as a cylinder of radius 0.3m with height uniformly distributed between 1.5m and 2m, and trees as cylinders with radius 0.4m and height 3m. We model cars as rectangular cuboids of size $2m \times 4.5m \times 1.5m$. We consider a moderate pedestrian density of $0.08person/m^2$ [23].

The link budget in our simulations is as in eq. (1)-(2) defined in §2.1, also consistent with the measurement-based feasibility analysis in §3. We consider regular square arrays with standard radiation patterns [18], varying the size of the transmit and receive arrays around default settings of 10×10 at each end. We fix the EIRP to 40dBm. Thus, the size of the transmit array does not impact the link budget, but can have a significant effect on interference. The receive array size impacts both link budget and interference. The 60 GHz channel is well characterized by a small number of dominant

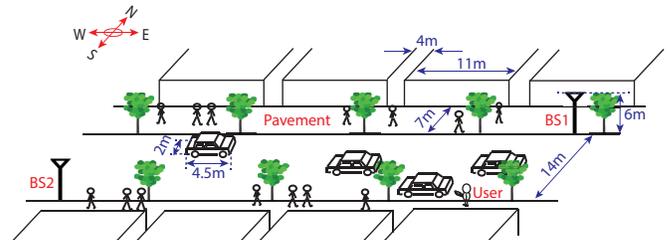


Figure 12: Simulation environment: a typical street in New York City (observed from the Google Map), where two basestations serve users on both sidewalks along the street.

rays (see [41] and references therein), with path strengths decaying rapidly with multiple bounces. Hence, in our model, we consider the LoS ray and one-bounce paths.

4.1 Link Availability

We consider several factors that may affect link availability. The first is user orientation, *e.g.* when facing basestation 1, the user’s body blocks the LoS path from basestation 2. We study its impact by comparing link availability among four possible orientations: north (facing far wall), south (facing near wall), east (facing basestation 1), and west (facing basestation 2). In general, the wall-facing orientations do not have LoS blockage, and the LoS availability is impaired mainly for the east-west orientations. We omit orientation-specific results for brevity; our later results are averaged over these four orientations. The second factor affecting link availability is blockage by other obstacles, *e.g.* other pedestrians, trees, and cars. We consider a moderate pedestrian density of 0.08 per square meter⁷. We find that, in general, cars are not tall enough to block LoS and wall reflections to users on sidewalks, but may disrupt ground reflections. Hence our simulations do not utilize any ground reflections from the street or sidewalk. Finally, the third factor is the presence of gaps between buildings (see Figure 12), which may eliminate possible reflected paths.

Taking into account these factors, we run simulations to estimate the distribution of RSS at different receiver locations, accounting for different receiver orientations and locations of other pedestrians and cars. We use this to compute the *availability rate* for three 802.11ad data rates listed in Table 1 (a data rate is “available” if the RSS exceeds its required receiver sensitivity).

Benefits of Picoclouds and Reflected Paths. Table 3 summarizes the availability rates for four scenarios: a mobile is served by (i) only the LoS path from basestation 1, (ii) a path among the

⁷We obtain similar results for much higher densities since it turns out that blockage of user body is the key source of link impairment.

Data rate (Mbps)	385	1155	2310
LoS (BS1)	66.0%	66.0%	66.0%
LoS + 2 NLoS (BS1)	94.2%	93.8%	93.5%
LoS (BS1 & BS2)	94.5%	94.5%	94.5%
6 paths (BS1 & BS2)	100%	99.95%	99.95%

Table 3: Effect of picocloud and reflection paths on availability, 10×10 RX array, 100m separation between the two basestations.

RX array	BS spacing	Clear			Rainy		
		Data rate (Mbps)					
		385	1155	2310	385	1155	2310
4×4	100m	99.9%	99.5%	99%	99.8%	99.4%	98%
	200m	99.4%	94.6%	85%	98.3%	89.2%	75.7%
8×8	100m	99.9%	99.75%	99.6%	100%	99.75%	99.7%
	200m	99.5%	99.45%	98.9%	99.5%	99%	97.2%
10×10	100m	100%	99.95%	99.95%	99.9%	99.9%	99.8%
	200m	99.5%	99.5%	99.3%	99.5%	99.5%	98.8%

Table 4: Availability rates (blockage scenario).

LoS and 2 wall-reflected paths (if feasible) from basestation 1, (iii) only the LoS paths from basestation 1 and 2, and (iv) a path among the total 6 possible paths from basestation 1 and 2. This result assumes a basestation spacing of 100m, but we observe similar conclusions for 150m and 200m. This result shows that using LoS path from a single basestation cannot provide robust connectivity. But when both basestations are involved, the link availability jumps from 66% to 94.5% at 385Mbps. Adding just two wall-reflected paths further increases the availability rate to 99.95-100%.

Basestation Spacing and RX Array Size. Assuming a picocloud of two basestations and the use of both LoS and reflected paths, Table 4 reports the availability rate for different basestation spacing and RX array size. We see that even heavy rain (50mm/hr) barely has any impact. For the 10×10 RX array and a BS spacing of 200m, we have high availability even at 2 Gbps (99.3% in clear, and 98.8% in rainy conditions). This implies that even with a sparse deployment, the network offers robust connectivity and coverage.

4.2 Interference and Spatial Reuse

We report on some preliminary results that indicate the massive spatial reuse achievable in our architecture. To this end, we lower the basestation spacing to 20m, and consider mutual interference between two users, each served by one of two basestations forming a picocloud. In addition to using different receive arrays, we also consider 10×10 and 32×32 transmit arrays. The EIRP is fixed, hence the transmit array size does not impact the link budget, but can have a significant effect on interference.

We first consider a scenario with “multiuser diversity”: there are five users, out of which two are picked at any time. We consider two strategies (neither of these is optimal, but they suffice to illustrate the impact of interference). In the first, termed without coordination (WoC), the basestations do not account for interference in scheduling: the first basestation serves the user to which it has the best path, and the second chooses the user to which it has the best path among the remaining 4 users. In the second strategy, termed With Coordination (WC), the first basestation chooses the user for which it has the best path (as before), but the second basestation chooses a user to maximize the minimum SINR among the two chosen users. Our simulations of availability are as before, except that we now also account for the presence of interference in determining receiver sensitivities. We find that the availability is 100% for all antenna combinations, with or without coordination! Basically, noise is not a limiting factor at this shorter range, and a user

RX array	Mode	TX array 10×10			TX array 32×32		
		Data rate (Mbps)					
		385	1155	2310	385	1155	2310
4×4	WoC	70%	60%	55%	89.5%	79.5%	72%
	WC	75%	68%	62%	89.8%	82.6%	76%
10×10	WoC	94.3%	91%	81%	98%	96.8%	93%
	WC	94.7%	91%	84%	98%	96.8%	94.8%

Table 5: Availability rates with blockage and interference, where the user pair are co-located in between the two basestations.

RX array	Mode	TX array 10×10			TX array 32×32		
		Data rate (Mbps)					
		385	1155	2310	385	1155	2310
4×4	WoC	59%	48%	46%	73%	58%	52%
	WC	66%	59%	56%	77%	66%	61%
10×10	WoC	90%	81%	75%	92.5%	90%	85%
	WC	91%	85%	81%	93%	91%	88%

Table 6: Availability rates in the offset geometry.

configuration that is “bad enough” (to produce large interference despite beamforming at both ends) is a rare event. We plan to develop an analytical framework for understanding the probabilities of such rare events in future work.

Next, we consider a “worst-case” scenario where two users next to each other are being served by the two base stations. Table 5 shows availability rates, averaging over the location of the user pair between the basestations, and their individual orientations. Here the transmit array size has a large effect. Using the larger array drops the outage rate from 30% to 10.5% for the smallest 4×4 receive array. For the 10×10 receive array, the main effect of using a larger transmit array is to increase the availability rate for the 2 Gbps rate.

The performance gets slightly worse in the “offset” geometry depicted in Figure 10(a), where the user pair is no longer between the two basestations. The availability in Table 6 drops to 93% even for the largest transmit and receive arrays considered. Of course, these worst-case settings are unlikely for randomly placed users, and can be avoided if we have multiuser diversity (so the picocloud does not have to simultaneously transmit to pairs of nearby users).

These results, despite their preliminary nature, indicate the promise of massive spatial reuse provided by the use of narrow beams at both transmit and receiver, and motivate more detailed analysis and simulations for characterizing the nature of interference in both average and worst cases.

5. RELATED WORK

mmWave Cellular Networks. Picocellular networks in the mm wave bands have attracted significant attention recently. Much recent effort has focused on trying to attain ranges of the order of 200m or even more, thus considered bands such as 28GHz, 38GHz and 73GHz to avoid oxygen absorption that peaks at 60GHz (see [29] and a recent survey [32]). Recent work also proposed to increase range by combining signals from many NLoS paths incoherently [12].

Our proposed 60GHz picocell design differs from existing efforts by taking a different perspective. Aiming for short links (50-100m)⁸, the oxygen absorption at 60GHz barely dents the link budget while still cutting down on interference from far-away basestations and hence significantly increasing spatial reuse. Together, the

⁸Mobility management becomes difficult at short ranges. This can be potentially addressed by the picocloud architecture, but still remains an open challenge that demands further research study.

large per-user data rate and large spatial reuse can potentially enable the 1000-fold increase in network capacity. Furthermore, our proposal leverages current and ongoing 60GHz hardware technologies and the IEEE 802.11ad standard developed for indoor applications, while the unlicensed nature of 60GHz bands lowers the entry barrier and promotes innovation.

Outdoor 60GHz Systems. The outdoor use of 60GHz is not new: there are point-to-point Gigabit links on the market today (e.g., [7]), attaining link ranges of kilometers using carefully installed parabolic dishes. These fixed beam antennas have significantly higher directivities than the electronically steerable antenna arrays envisioned in our architecture. Recent work has examined 60GHz mesh networks with electronically steerable links for backhaul [35, 36]. Finally, the impact of blockage on short-range outdoor 60 GHz links has been explored independently in [10, 11] using ray-tracing simulations. Unlike our work, this study is not grounded in measurements, and does not consider the detailed tradeoffs in the picocloud architecture.

60GHz Measurements. There have been significant efforts in 60 GHz channel measurement and modeling both for indoor [24, 17] and outdoor settings [37, 39, 21, 33], confirming the accuracy of quasi-optical propagation models on 60GHz propagation. Most measurements used channel sounding and probing hardware, while others have cleverly adapted commercial radios to indoor measurements [38, 42].

While prior measurement campaigns focus on link level characterization and propagation modeling, to the best of our knowledge, our work is the first measurement-centric study aimed specifically at design and validation of outdoor 60GHz picocell architectures.

6. CONCLUSION

Our measurements and simulations dispel some common myths on 60 GHz outdoor mobile communication, and show that outdoor 60 GHz picocells augmenting existing cellular networks indeed have the potential to deliver orders of magnitude increase in network capacity. Our findings lead us to identify the need for new, novel architectures and system designs such as “picoclouds,” as part of a new research area on 60GHz outdoor networks. We conclude by discussing some of the significant challenges that remain.

Cross-layer Modeling. In our experiments, off-the-shelf hardware provides insight into propagation, blockage and interference in simple scenarios. However, today’s hardware limitations imply that future research involving more complex scenarios must rely on simulations until more capable platforms (e.g., with externally programmable beamsteering) become available. Therefore, significant effort is necessary to develop models at different layers accurate enough to capture the cross-layer interactions unique to 60 GHz communication.

Specifically, the lowest layer involves tying the physics of propagation and diffraction to link-level models. This includes link availability and data rate as a function of pedestrian configurations, for group motion as well as independent individual motion. Another layer involves signal processing models and algorithms that account for hardware constraints, e.g., RF beamforming with a large number of antenna elements. We can then build to higher layers, including abstractions for user tracking and basestation coordination, and network-level resource management for seamless support of TCP flows.

User Tracking. To establish/maintain links, basestations must track users continuously and build a rich path inventory. GPS location works only for LoS paths, but its localization error is too large to enable highly directive beams. Brute-force beam scans in-

cur large overhead, which prevent online adaptation. A potential solution is compressive adaptation [31], where basestations periodically send 60GHz beacons in random beam directions, and then leverage user feedback to estimate directions of available paths to each user. However, significant effort is required to translate such ideas to practical system designs.

Control Plane. A robust control plane is needed for coordination between basestations and users during bootstrapping, handoff and array (re)alignment. A potential candidate is to use existing LTE services (Figure 1) for this purpose. The resulting control channel is reliable, but the design must be robust to latencies as high as 80ms [19], especially if LTE bands are congested.

Picocloud Architecture & Protocol Design. In the picocloud architecture, basestations must coordinate to maintain user connectivity and minimize interference. This requires new, cross-layer design spanning the PHY, MAC and network layers. Frequent switching between basestations to handle blockage and mobility, together with rate adaptation switching between different paths, requires careful coordination to support applications such as TCP. Given the high frequency of blockages, it is particularly important to avoid naive handoff strategies that could lead to ping-pong effects and race conditions. Such considerations, together with the large per-user rate, call for careful designs that account for backhaul and storage constraints.

Hardware Design. Advances in system-on-chip design enable packing multiple radios into mobile devices. The integration of 60GHz radios poses little difficulty in terms of interference with existing radios, since the frequency band is so far from current WiFi and cellular frequencies. However, another consideration is the integration of (possibly multiple) antenna arrays to receive signals effectively regardless of the device orientation. Limiting radio energy cost for mobile devices is another. While signal processing and RF design for high-speed mmwave links are challenging, our advantage lies in leveraging existing and ongoing hardware development for IEEE 802.11ad.

7. ACKNOWLEDGMENTS

The authors thank the anonymous reviewers for their insightful feedback. This work is supported in part by NSF grant CNS-1317153. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the NSF.

8. REFERENCES

- [1] <http://www.techjournal.org/2011/09/mobile-broadband-useage-is-set-to-explode-infographic>.
- [2] http://money.cnn.com/2012/02/24/technology/spectrum_crunch_solutions/index.htm.
- [3] <http://www.dailywireless.org/2011/05/06/att-moves-to-hotzones-and-picocells>.
- [4] <http://www.siliconimage.com/news/releasedetails.aspx?id=649>.
- [5] http://news.cnet.com/8301-30685_3-57326718-264/wilocity-60ghz-wireless-revolution-begins-at-ces/.
- [6] http://transition.fcc.gov/Bureaus/Engineering_Technology/Documents/bulletins/oet70/oet70a.pdf.

- [7] Cisco Small Cell Wireless Backhaul Ecosystem: A Flexible and Proven Deployment Toolkit. http://www.cisco.com/c/en/us/solutions/collateral/service-provider/small-cell-solutions/brochure_c02-728436.html.
- [8] Fcc 13-112. http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-13-112A1.pdf.
- [9] IEEE 802.11 Task Group AD. http://www.ieee802.org/11/Reports/tgad_update.htm.
- [10] ABOUELSEoud, M., AND CHARLTON, G. The effect of human blockage on the performance of millimeter-wave access link for outdoor coverage. In *Proc. of VTC* (2013).
- [11] ABOUELSEoud, M., AND CHARLTON, G. System level performance of millimeter-wave access link for outdoor coverage. In *Proc. of WCNC* (2013).
- [12] AKDENIZ, M. R., LIU, Y., RANGAN, S., AND ERKIP, E. Millimeter wave picocellular system evaluation for urban deployments. *arXiv:1304.3963* (2013).
- [13] BALASUBRAMANIAN, A., MAHAJAN, R., AND VENKATARAMANI, A. Augmenting mobile 3G using WiFi. In *Proc. of MobiSys* (2010).
- [14] CORREIA, L., AND REIS, J. Wideband characterisation of the propagation channel for outdoors at 60 ghz. In *Proc. of PIMRC* (1996).
- [15] COSTELLO, D., AND FORNEY, G.D., J. Channel coding: The road to channel capacity. *Proc. of the IEEE* 95, 6 (2007), 1150–1177.
- [16] FRIIS, H. T. A note on a simple transmission formula. *Proc. of IRE* (1946).
- [17] GENG, S., KIVINEN, J., ZHAO, X., AND VAINIKAINEN, P. Millimeter-wave propagation channel characterization for short-range wireless communications. *IEEE TVT* 58, 1 (2009), 3–13.
- [18] HANSEN, R. C. *Phased array antennas*, 2nd ed. John Wiley & Sons, 2009.
- [19] HUANG, J., QIAN, F., GERBER, A., MAO, Z. M., SEN, S., AND SPATSHECK, O. A close examination of performance and power characteristics of 4G LTE networks. In *Proc. of MobiSys* (2012).
- [20] KELLEHER, K. *The Microwave Engineers' Handbook and Buyers' Guide*, 5th ed. New York: Horizon Press, 1964.
- [21] LANGEN, B., LOBER, G., AND HERZIG, W. Reflection and transmission behavior of building materials at 60ghz. In *Proc. of PIMRC* (1994).
- [22] LOVNES, G., REIS, J., AND RAEKKEN, R. Channel sounding measurements at 59 ghz in city streets. In *Proc. of PIMRC* (1994).
- [23] MAEDA, K., ET AL. Urban pedestrian mobility for mobile wireless network simulation. *Ad Hoc Networks* 7, 1 (2009), 153–170.
- [24] MALTSEV, A., MASLENNIKOV, R., SEVASTYANOV, A., KHORYAEV, A., AND LOMAYEV, A. Experimental investigations of 60 GHz WLAN systems in office environment. *IEEE JSAC* 27, 8 (2009).
- [25] MATIC, D., HARADA, H., AND PRASAD, R. Indoor and outdoor frequency measurements for mm-waves in the range of 60 ghz. In *Proc. of VTC* (1998).
- [26] NATARAJAN, A., REYNOLDS, S. K., TSAI, M.-D., NICOLSON, S. T., ZHAN, J.-H., KAM, D. G., LIU, D., HUANG, Y.-L., VALDES-GARCIA, A., AND FLOYD, B. A. A fully-integrated 16-element phased-array receiver in sigebicmos for 60-ghz communications. *IEEE JSSC* 46, 5 (2011), 1059–1075.
- [27] NISHI, S., AND TOKUDA, K. Development of mm-wave video transmission system-development of antenna. In *Proc. of APMC* (2001).
- [28] PI, Z., AND KHAN, F. An introduction to millimeter-wave mobile broadband systems. *IEEE Comm. Magazine* 49 (2009), 101–107.
- [29] PI, Z., AND KHAN, F. System design and network architecture for a millimeter-wave mobile broadband (mmb) system. In *Proc. of Sarnoff Symposium* (2011).
- [30] RAMASAMY, D., GANTI, R., AND MADHOW, U. On the capacity of picocellular networks. In *Proc. of ISIT* (2013).
- [31] RAMASAMY, D., VENKATESWARAN, S., AND MADHOW, U. Compressive tracking with 1000-element arrays: A framework for multi-Gbps mm wave cellular downlinks. In *Proc. of Allerton* (2012).
- [32] RANGAN, S., RAPPAPORT, T., AND ERKIP, E. Millimeter-wave cellular wireless networks: Potentials and challenges. *Proc. of the IEEE* 102, 3 (March 2014), 366–385.
- [33] RAPPAPORT, T. S., BEN-DOR, E., MURDOCK, J. N., AND QIAO, Y. 38GHz and 60 GHz angle-dependent propagation for cellular & peer-to-peer wireless communications. In *Proc. of ICC* (2012).
- [34] RAPPAPORT, T. S., GUTIERREZ, F., BEN-DOR, E., MURDOCK, J. N., QIAO, Y., AND TAMIR, J. I. Broadband millimeter-wave propagation measurements and models using adaptive-beam antennas for outdoor urban cellular communications. *IEEE TAP* 61, 4 (2013), 1850–1859.
- [35] SINGH, S., MUDUMBAI, R., AND MADHOW, U. Distributed coordination with deaf neighbors: efficient medium access for 60 GHz mesh networks. In *Proc. of Infocom* (2010).
- [36] SINGH, S., MUDUMBAI, R., AND MADHOW, U. Interference analysis for highly directional 60-ghz mesh networks: The case for rethinking medium access control. *IEEE/ACM ToN* 19, 5 (oct. 2011).
- [37] SMULDERS, P., AND CORREIA, L. Characterisation of propagation in 60 ghz radio channels. *Electronics & communication engineering journal* 9, 2 (1997), 73–80.
- [38] TIE, X., RAMACHANDRAN, K., AND MAHINDRA, R. On 60 ghz wireless link performance in indoor environments. In *Proc. of PAM* (2012).
- [39] VIOLETTE, E., ESPELAND, R., AND HAND, G. R. Millimeter-wave urban and suburban propagation measurements using narrow and wide bandwidth channel probes. *NASA STI/Recon Technical Report* 86 (1985), 25683.
- [40] WILTSE, J. Corrections to published curves for atmospheric attenuation in the 10 to 1000 ghz region. In *Proc. of ISAP* (1997).
- [41] ZHANG, H., VENKATESWARAN, S., AND MADHOW, U. Channel modeling and mimo capacity for outdoor millimeter wave links. In *Proc. of WCNC* (2010).
- [42] ZHOU, X., ZHANG, Z., ZHU, Y., LI, Y., KUMAR, S., VAHDAT, A., ZHAO, B. Y., AND ZHENG, H. Mirror mirror on the ceiling: Flexible wireless links for data centers. In *Proc. of SIGCOMM* (2012).