Extending Cell Tower Coverage through Drones

Ashutosh Dhekne
University of Illinois (UIUC)

Mahanth Gowda
University of Illinois (UIUC)

Romit Roy Choudhury
University of Illinois (UIUC)

ABSTRACT

This paper explores a future in which drones serve as extensions to cellular networks. Equipped with a WiFi interface and a (LTE/5G) backhaul link, we envision a drone to fly in and create a WiFi network in a desired region. Analogous to fire engines, these drones can offer on-demand network service, alleviating unpredictable problems such as sudden traffic hotspots, poor coverage, and natural disasters. While realizing such a vision would need various pieces to come together, we focus on the problem of "drone placement". We ask: when several scattered users demand cellular connectivity in a particular area, *where should the drone hover* so that the aggregate demands are optimally satisfied? This is essentially a search problem, i.e., the drone needs to determine a 3D location from which its SNR to all the clients is maximized. Given the unknown environmental conditions (such as multipath, wireless shadows, foliage, and absorption), it is not trivial to predict the best hovering location.

We explore the possibility of using RF ray tracing as a hint to narrow down the scope of search. Our key idea is to use 3D models from Google Earth to roughly model the terrain of the region, and then simulate how signals would scatter from the drone to various clients. While such simulations offer coarse-grained results, we find that they can still be valuable in broadly guiding the drone in the right direction. Once the drone has narrowed down the 3D search space, it can then physically move to quickly select the best hovering location. Measurement results from a WiFi mounted drone, communicating with 7 clients scattered in the UIUC campus, are encouraging. Our early prototype, *DroneNet*, reports 44% throughput gain with only 10% measurement overhead compared to a full scan of the entire region.

CCS Concepts

•**Networks** → **Wireless access points, base stations and infrastructure;** *Mobile networks;*

Keywords

Infrastructure Mobility; Drones; Ray Tracing

1. INTRODUCTION

Outdoor cellular network traffic is steadily on the rise. Video is already the dominant application for cellular networks [1,2]. In the near future, in-vehicle entertainment [3], uploads from numerous cameras, and various IoT applications (including smart cities, precision agriculture [4], and wearables [5]) will further add to the bandwidth pressure. Predictions indicate a 1000x increase in wireless data demand by 2020 [6,7].

While technological advances in MIMO, beamforming, spectrum sensing, and others have coped with this pressure thus far, there is wide agreement that such opportunities are saturating. Users are beginning to experience spatial or temporal degradations in the quality of service. For instance, areas with tall buildings are suffering from poor SNR due to wireless shadows [8]; flash crowds at political rallies, sports events, and other social occasions are creating sudden traffic spikes [9]; natural disasters are destroying local network infrastructure, warranting a quick and temporary replacement service. Solving these problems with additional tower installations does not scale—the cost of over-provisioning is becoming excessive, exacerbated by the difficulty in finding installation sites in dense urban regions. This paper takes an exploratory step and envisions drones as "elastic extenders" of cell towers. We call our system *DroneNet*.

DroneNet's model of operation bears similarity to cloud computing. Clouds leverage statistical aggregation opportunities, i.e., given that only a fraction of clients are requesting resources at a given time, the total resources in the cloud need not scale with the number of clients. Yet, any given client can still avail powerful resources from the cloud. Drones bring similar flexibility to the wireless world. Not every user experiences dead zones at a given time; neither are all users located in a flash crowd. Moreover, traffic demand exhibits a power law behavior; few users form the majority of demand at a time. Hence, a limited set of drones may be adequate to address all the dynamic needs, leading to a win-win situation for both the clients and network service providers.

Figure 1 illustrates a toy example. It shows multiple possible locations at which the drone can hover. From any of these locations, the drone connects to the clients with a WiFi link, while the backhaul operates over a 4G/LTE link back to the cell tower. As mentioned earlier, the core research question pertains to determining the best hovering position. Observe that moving closer to the ground improves client proximity, however, the multipath and shadowing effects get severely exacerbated. Moreover, the line of sight (LOS) to the cell tower also gets disrupted. Moving vertically higher offers better LOS to clients and the cell tower, but at the expense of longer distance to these clients, reducing data rates. Lateral movements also pose tradeoffs – for instance, the left and right-most positions in the Figure 1 both offer LOS paths to one of the clients but blocks the other. A combination of lateral and vertical movements can bring the drone to a position that maximizes a given function of SNR. Our goal is to efficiently determine this location.

HotMobile '17, February 21-22, 2017, Sonoma, CA, USA

DOI: http://dx.doi.org/10.1145/3032970.3032984

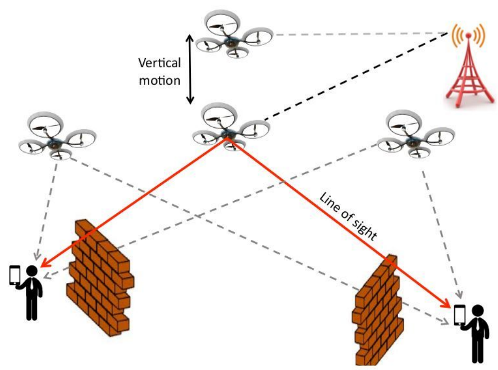

Figure 1: Drone locations present tradeoffs: Closer to the ground aggravates multipath and blocks line of sight (LOS) to the cell tower, while vertically higher placement increases distance to clients. Lateral movements can achieve LOS to some clients but get blocked from others.

A brute force solution would be to fly the drone and conduct SNR measurements to empirically search for the optimal location. However, a large 3D search space – say 2 or 3 city blocks in Chicago – makes this approach prohibitively time consuming. Movements in clients and changes in traffic patterns will occur at faster time scales, rendering this brute force search useless. Hence we require a solution that is lightweight and quick. Simple strategies like hovering at the centroid of a group of clients are unsuitable due to the non-monotonous relation between distance and SNR. Results from historical searches are also not useful since each new situation is somewhat unique in its placement of clients and the type of traffic demands.

We explore the possibility of using RF ray tracing as a hint to narrow down the scope of search. Our key idea is to model the dominant structures located in an area—such as the buildings and trees—to roughly model the terrain of the region, and then simulate how signals would bounce and scatter from the drone to the various clients. Of course, such simulations yield coarse-grained results since the simulated SNR is sensitive to centimeter-scale errors. Nonetheless, we find that these simulated results can still be valuable in broadly guiding the drone towards the right direction, i.e., towards areas where the SNR is relatively better. Once the drone arrives in this area, it physically conducts measurements to fine-tune its hovering location. Given a considerably smaller search space, the operation incurs far less time. The drone now hovers at this location offering connectivity to clients. If client positions or traffic changes substantially, the drone recomputes the ray-tracing results and finds a new hovering location.

Early prototype with 7 clients spread across an area of $160 \times 280m$ in the UIUC campus shows promise. We conduct flights with an octocopter, carrying an Almond WiFi AP and continuously transmitting packets to clients. The octocopter moves in a raster-scan over an area of $50 \times 68m$, and at three different heights at 15, 30, and $45m$, above ground level. We conduct ray tracing simulations using Remcom Wireless Insite [10]. We observe consistent correlation between the ray-tracing model and measurements – *DroneNet* exploits this to extend 44% throughput gain with only 10% of full measurement overhead.

We briefly summarize the contributions as follows:

(1) *Exploiting inaccurate ray tracing as an opportunity to reduce the search space for drone placement.* Fine tuning the SNR search through small scale physical movements of the drone.

(2) *Measurement based evaluation on a real drone, flying on the UIUC campus while communicating to 7 clients on the ground.* Significant engineering effort towards building this infrastructure, including payload optimization, battery management, data rate selection, power control, ground truthing, etc.

The rest of the paper expands on these ideas and experimentation effort. However, we first address some of the natural questions and discuss other alternatives to solve the problem.

2. NATURAL QUESTIONS

Does today's battery technology support longer flights? Flying a drone requires significant battery power. However, constant hovering may be unnecessary. Drones may hover only under extreme situations, while under less dire needs, they may fly and park at charging stations on top of buildings, lamp-posts, or fences. Gas/solar powered drones offering extended lifetime are also becoming available [11]. Alternative designs based on balloons consume relatively lesser energy [12]. We are aware of current battery limitations, however, we see opportunities emerging in the future that will make extended drone flight viable.

How practical is obtaining terrain model? Public data from Google warehouse is available for many buildings [13]. However, it is possible to create 3D models of an area by taking pictures from drones and using photogrammetry techniques [14]. We do not rely on material of the terrain.

How compelling are the gains? The gains are promising. Figure 2 shows that 16-18dB of SNR gain is possible from drone mobility at various heights (30m, 45m, and 60m). This translates into a throughput gain of 44%.

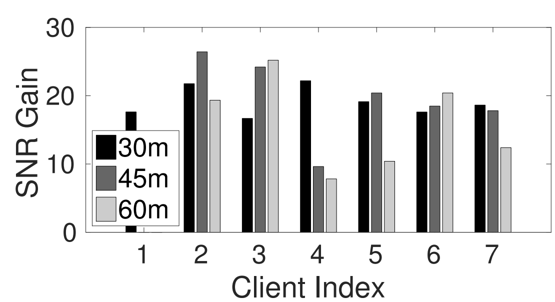

Figure 2: SNR gains for various heights across all clients

Are the gains emerging from moving closer to clients? Moving closer does not necessarily improve SNR due to multipath. To elaborate more, suppose that drone mobility is constrained in a region R. Let P_{close} denote a location within R that is closest to a client. Now, let P_{10} denote a subset of locations within R such that the SNR to the client from these locations is ranked in the top 10 percentile. Figure 3 shows the separation between P_{10} and various locations in P_{close}. Evidently many high SNR positions exist when the drone is 20-30 meters far from P_{close}. This suggests that the drone need not go closer to achieve a near-optimal SNR.

Figure 4 shows the best SNR in R for different heights of the drone at various client positions. SNR can both decrease or increase with height depending upon terrain structure, client location. SNR variation is non monotonous and going closer to the client by decreasing the height is not always beneficial.

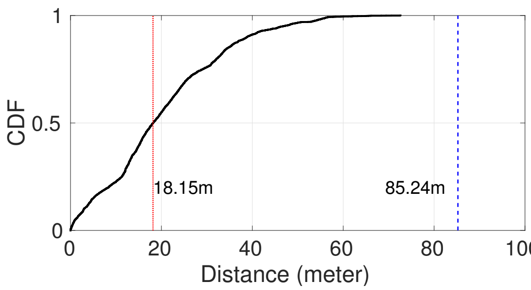

Figure 3: High SNR positions are scattered everywhere

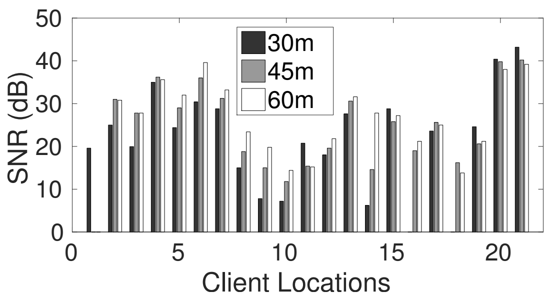

Figure 4: Best observed SNR for different drone heights

Why not use Femtocells or WiFi Hotspots? While beneficial, static infrastructure entails high density to simultaneously cover all shadows. Moreover, the height of deployment is limited to height of buildings. A drone on the other hand can fly higher and offer much higher coverage. The drone's motion allows on-demand mitigation of dead zones and hotspots.

3. SYSTEM DESIGN

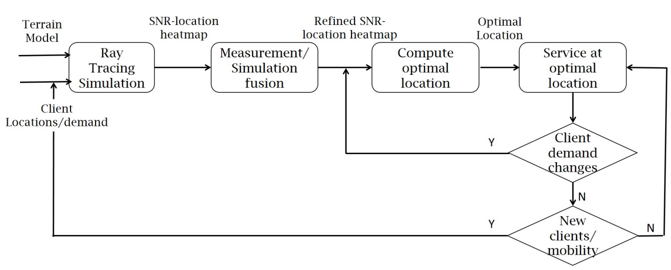

Figure 5: Flow of operations in *DroneNet*

Figure. 5 summarizes the flow of operations in *DroneNet*. Given a set of client locations and the terrain model, *DroneNet* first runs a low fidelity, light weight ray tracing simulation to compute SNR at each client as a function of drone location. A 3D heatmap showing this SNR for each client is obtained. The SNRs are translated into throughput using Shannon's equation[1]. The sum throughput over all clients is computed, resulting in a 3D heatmap of aggregate-throughput as a function of drone location. *DroneNet* then conducts a quick scan of physical measurements around the regions of high throughout in the heatmap. During the scan, the drone finds the position offering maximum aggregate-throughput across all clients and hovers there. Whenever the traffic demand of clients changes, or clients move, *DroneNet* should adapt accordingly. However, we leave addressing client mobility to future work. We now expand the two key modules enabling *DroneNet* design: (1) Ray tracing simulation, and (2) Fusion of simulation and measurements

[1]We consider throughputs of only client-drone links (without including the drone LTE backhaul) since they create bottleneck. The LTE backhaul is expected to have clear line of sight to the basestation, therefore of high quality. However, client-drone links are more challenging due to terrain shadowing.

3.1 Ray tracing simulations

Using ray tracing simulations, *DroneNet* predicts SNR heatmap without undergoing the overhead of empirical measurements. A basic "shoot and bounce" ray tracing [15] simulation is conducted using Remcom Wireless Insite software [10]. We describe the steps involved using Figure 6 as an illustration. It shows a few buildings, one client, and an area where the drone can potentially fly. The goal is to predict SNR from ray tracing.

1. Many rays pointing in all directions are generated from the client position. Figure 6 depicts rays emitting from *Client-1*
2. Each ray advances, and hits objects in the environment such as buildings and trees.
3. This causes deflections—reflections and diffractions. While each reflection creates one reflected ray, diffraction creates multiple rays. The deflection of rays generated from *Client 1* is observable in Figure 6.
4. Appropriate amplitude is set for deflected rays based on propagation delay and the absorption, reflection and refraction coefficients of incident material. We assume a single material for all structures, resulting in some inaccuracies.
5. Rays with weak signal strength are eliminated. This includes: (1) Rays undergoing more than 4 successive reflections. (2) Rays undergoing more than 3 successive diffractions. (3) Rays traveling longer than a certain threshold.
6. Finally, all rays arriving at the drone location are noted. The SNR is computed as a function of amplitude and delay of the arriving rays. Figure 6 shows the set of rays converging at the drone location *L1* using which the SNR at *L1* can be computed.

Similarly, SNR of all other possible client-drone links are computed to obtain their own SNR heatmaps.

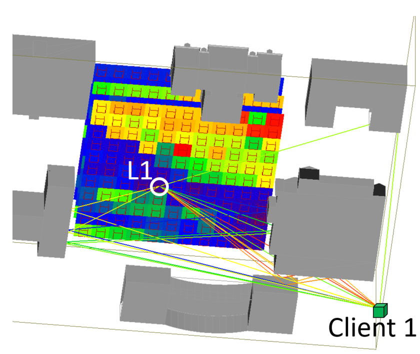

Figure 6: 3D SNR heatmap

Figure 7 shows the comparison between measured SNR and ray tracing predicted SNR. While ray tracing simulations can model the reality well (Figure 7(a)), it can also be less accurate for some clients (Figure 7(b)).

The simulations can be offloaded to cloud. Real time results [16, 17] are achievable through parallelization and tradeoffs with accuracy. In this paper, however, we conduct simulations offline.

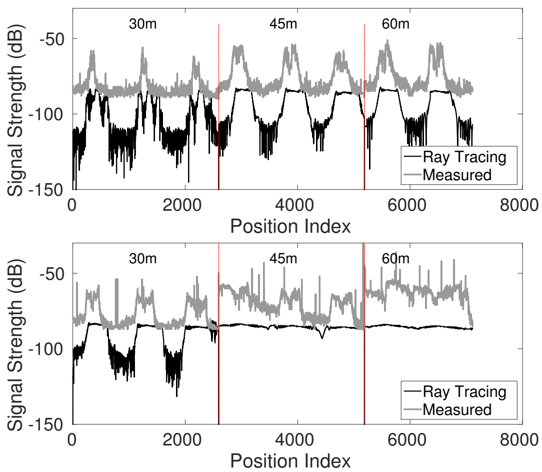

Figure 7: (a) Ray tracing can capture variation in SNR well. (b) Ray tracing can be inaccurate in some cases

3.2 Fusion of ray tracing and measurements

While ray-tracing predictions are promising, hovering a drone entirely based on predictions gave poor performance. We define *SNR gain* of a candidate position for drone hovering as the difference between the measured SNR of that position and the median of measured SNRs of all candidate positions (more details in Section 4). We take top 10%-ile high SNR drone positions from ray-tracing predictions and plot a CDF of their *SNR gain* in Figure 8. Evidently, the median is close to 0 suggesting that ray-tracing alone is not enough. The low accuracy arises from inexact terrain modeling, imprecise material absorptions, unmodeled foliage and limited number of rays/reflections/diffractions used to reduce complexity. *Oracle* is an imaginary system which can magically predict measurements to 100% accuracy. We observe in Figure 8 that Oracle has a substantial *SNR gain*. *DroneNet* achieves 57% of these gains by fusing ray-tracing predictions with sparse measurements, as elaborated below.

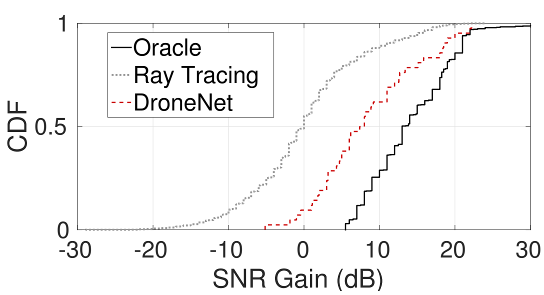

Figure 8: Ray tracing predictions fail to capture SNR gains

1. Perform ray tracing simulations to obtain sum-throughput heatmap as outlined in Section 3.1
2. Select the top 10%ile drone locations in the sum-throughput heatmap. Let this set of points be denoted as ht
3. Divide the entire region into chunks of size $25m^2$. Find the chunk that contains the largest number of positions from ht.
4. Scan the chunk determined from the above step to conduct physical throughput measurements. The position with maximum throughput in the chunk based on measurements is selected for hovering.

After determining the optimal location, the drone needs to relocate there. The relocation accuracy allows some tolerance because there will be many spatially contiguous points with similar throughput around the optimal position.

4. EVALUATION

4.1 Experimental Setup

Our experimental platform consists of an X8 quadcopter from 3D Robotics [18]. We placed an Almond WiFi AP [19], powered by a LiPo battery, on the quadcopter and set it to transmit at a frequency of $2.437GHz$, ($20MHz$ bandwidth), with a transmit power of 28dBm. Whereas we use WiFi for the client-drone link, it can also be a cellular link with a femtocell (instead of the WiFi AP) operating in a designated LTE band. An Android phone was additionally used to time-stamp and location stamp the packets sent out by the Almond AP. 7 clients running linux on Raspberry Pis were spread in an area ($160 \times 280m$) around the Engineering Quad at the University of Illinois, as shown in Figure 9(a).

The Raspberry Pis were connected to an Atheros WiFi dongle. The AP transmitted 400 packets per second, which were captured by the clients from which SNR was computed. During the packet transmissions, the drone flew in a raster scan at a speed of 1m/s covering the area ($50 \times 68m$) highlighted in Figure 9(a). Three different heights—30, 45 and 60 m—were used for drone flights above this area. This provides us a platform to measure the spatial heatmap of SNR variation.[2]

The 3D terrain model of most buildings around the Quad were obtained from Google 3D Warehouse [13]. Some of the unavailable models were manually created using AutoCAD softwares by borrowing relevant design diagrams from the architecture department. The top and perspective view of the buildings from Figure 9(a) is modeled in Figure 9 (b, c). The model includes 3D terrain structures, client locations and drone flight path. The 3D model was input to Remcom Wireless Insite [10], producing another spatial heatmap of SNR variations used for predictions.

Let the *Oracle* be a system whose ray-tracing predictions match exactly with measurements. We define *Random* as a system where the drone randomly chooses a position to hover. We define *Ray Tracing* as a system where the drone hovers at the best location predicted by ray tracing simulations. We characterize performance gains of systems – *Ray Tracing*, *DroneNet*, and *Oracle* over *Random*. We use throughput as the performance metric and translate SNR into throughput using Shannon's equation.

4.2 Performance Results

Single client throughput: Figure 10 quantifies the gain in throughput. *DroneNet* gain varies from $1.2x$ to $4.2x$ over various clients. While *RayTrace* gains has a significant difference from *Oracle* gains, *DroneNet* captures reasonable *Oracle* gains. **Multi client throughput:** Figure 11 shows the pattern of gain across multiple gains. While the gains decrease with higher clients, we need not optimize for all clients simultaneously.

[2]Experiments were conducted as per FAA regulations. Operator was trained by a general aviation pilot and flight permissions were obtained from campus police.

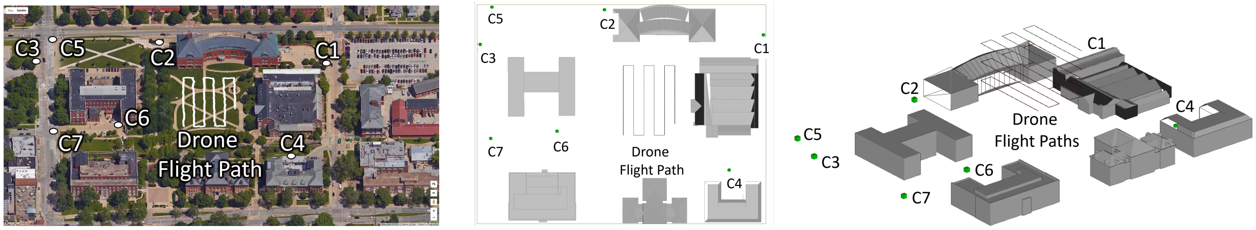

Figure 9: (a) Client locations and drone flight path (b) 3D Terrain Model – top view (c) 3D Model – perspective view

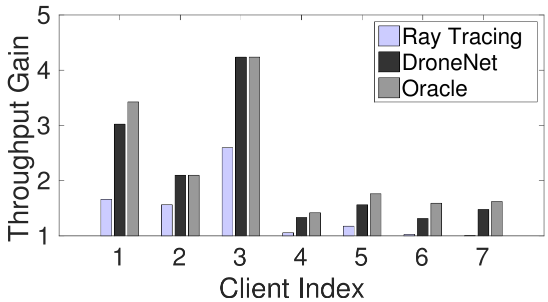

Figure 10: *DroneNet* gains are comparable with the *Oracle*

Power law [20] indicates that 50% traffic demand comes from 1% of the users who from the bottleneck. Addressing the bottleneck will benefit all other users.

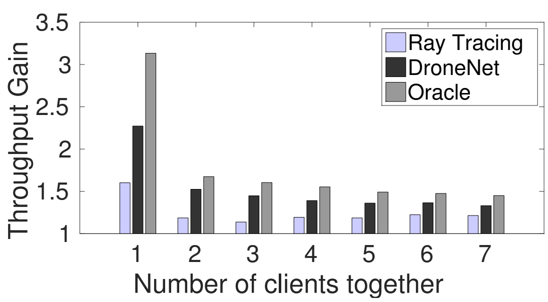

Figure 11: *DroneNet* provides up to 44% gains for 3 clients and 36% for 5 clients

Gains with constrained motion: Figure 12 computes the gains over various spatial constraints of drone motion. The experimental area, A, is partitioned into 4, 6, and 8 regions shown as $A/4$, $A/6$ and $A/8$ on the x-axis. Even with constrained motion, there is a graceful degradation in gains, suggesting that enough diversity from shadowing exists. This also suggests that even constrained mobility can offer gains. Figure 13 zooms in to show the CDF of gains for $A/6$. On an average, *DroneNet* achieves 57% of oracle gains with only 10% of measurement overhead.

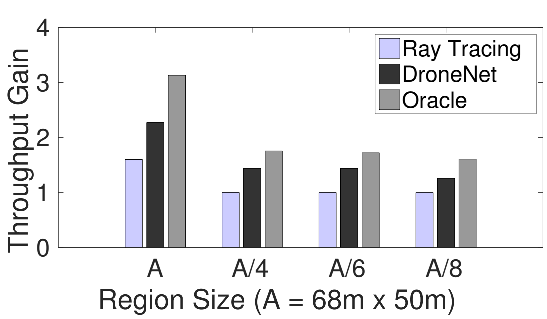

Figure 12: 25 $\sim$ 45% gains, with constrained mobility

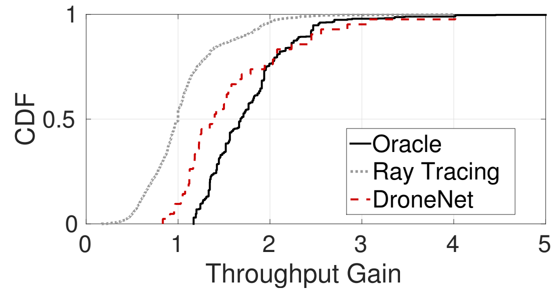

Figure 13: CDF of gains for A/6 region size

5. RELATED WORK

5.1 Mobility

Beamforming: While Beamforming techniques [21] help addressing similar problems as *DroneNet* by steering the radiation pattern based on demand, we believe combining beamforming with physical mobility can offer higher flexibility.
Mobile Infrastructure: *DroneNet*'s notion of mobile infrastructure is similar to Google Loon [22] and Facebook drones [23]. *DroneNet*, however, attempts to increase spectrum efficiency and address dynamic demands even in high density urban networks and not just in remote areas. DARPA Landroids [24] use autonomous robots for offering communication in challenging terrains for military. We believe that ideas in *DroneNet* will be useful for Landroids. Cellular towers are moved on vans [25] during sporting events. While useful, it viable only during large gatherings. A very early version of our work was presented as a Mobicom poster [26].

In indoor spaces, *iMob* [27, 28] shows that Roomba robots can enhance throughput by exploiting multipath opportunities. *DroneNet* operates outdoors and whereas *iMob* directly uses empirical measurements for finding the optimal AP position, for a large outdoor 3D search space, such measurements become costly. *DroneNet* uses ray-tracing to reduce this cost.

A swarm of drones is used in [29] to demonstrate multihop connectivity. In contrast, we focus on predicting optimal placement of the drones utilizing ray-tracing techniques.

5.2 Spatial SNR Modeling

Path loss Models: Path loss models characterize RF attenuation as a power of propagation distance. The exponent and variation properties are determined empirically and used for predictions [30–32]. However, its accuracy is limited and doesn't capture sudden SNR changes at building boundaries.
Ray Tracing: Ray tracing has been studied extensively in the context of cellular network planning. Works in [33, 34] show that ray tracing models can predict the SNR in urban and semi-urban areas. Works in [35, 36] attempt to decrease the complexity for Manhattan style urban areas by transforming 3d ray tracing into 2d ray tracing. Similarly [37] proposes an angular partitioning technique to decrease complexity. *DroneNet* benefits from these techniques and, moreover, incorporates selective measurements for refining the model.

6. DISCUSSION AND ON-GOING WORK

Client Mobility: We do not address mobile clients in this paper. Plethora of research exists in vehicular networking community for exploiting the predictive nature of mobility [38]. We believe it would be an interesting approach to combine ray tracing with mobility models to address the challenges. *DroneNet*'s ray-tracing models can be generalized to a path instead of a point, however, we leave this for future work.

OFDM/MIMO: SNR is not a good indicator of throughput for multi-carrier OFDM systems. While direct throughput assessment (with Iperf) was infeasible because of flying constraints, it is possible to extend the ray-tracing simulations to OFDM by observing the channel frequency response (CFR). Evaluation can be conducted with platforms that export CFR [39]. Perhaps, more opportunities at the subcarrier/antenna level from OFDM/MIMO can be used for optimizations. We treat them as separate problems and leave them for future work.

Location Requirement: *DroneNet*'s ray-tracing module requires client locations to be known with an accuracy of couple of meters. Since we attempt to mitigate large shadows caused by terrain profile, we do not need wavelength level accuracy. GPS location sharing may have privacy concerns. We do not address these concerns in this paper, our focus is to explore the possibilities.

7. CONCLUSION

Demand for cellular traffic continues to increase while the spectrum resources are limited. Cellular users also suffer from intermittent regions of poor connectivity due to wireless shadows and dead zones caused by buildings. We envision a system of drones that can extend cell towers and mitigate the dynamic connectivity issues while also facilitating efficient use of the scarce spectrum. Our system *DroneNet* explores a key problem of determining optimal drone placements. While we scratch the surface, much more remains to be done. Early prototype with real flights offers sufficient promise for a long term research engagement.

8. REFERENCES

[1] "Are mobile tv and video still killer apps?." http://www.mobileindustryreview.com/2014/09/mobile-video-still-killer-app.html.

[2] "Mobile video is the new killer app for mobile operators." http://v-net.tv/2013/03/27/mobile-video-is-the-new-killer-app-for-mobile-operators/.

[3] "How self-driving cars could make streaming tv the next radio." https://goo.gl/Eg4RCu/.

[4] "Why iot, big data and smart farming is the future of agriculture." http://read.bi/2dSdFV2.

[5] "Wearable and big data: Potential challenges, potential rewards." https://goo.gl/UYiwLN/.

[6] "2020: Network for the networked society – an industry beyond smartphones." goo.gl/DdZk4n.

[7] "Mobile data traffic to increase 1000 times beyond 2020." http://channeleye.co.uk/mobile-data-traffic-to-increase-1000-times-beyond-2020/.

[8] "How to fix reception problems in cell phone dead zones." http://lsh.re/1BA7U.

[9] "Texas a & m fans set all-time record for data usage for at&t." goo.gl/7O8YRk.

[10] "Wireless insite." http://www.remcom.com/.

[11] "Hycopter drone uses hydrogen gas to fly for 4 hours." https://goo.gl/4Cayi8/.

[12] "Is a drone-balloon hybrid a match made in heaven?." http://nofilmschool.com/2016/07/why-use-panasonic-balloncam-over-pure-drone.

[13] "3d warehouse." https://3dwarehouse.sketchup.com/.

[14] P. R. Wolf and B. A. Dewitt, *Elements of Photogrammetry: with applications in GIS*, vol. 3. McGraw-Hill, 2000.

[15] Y. Dama *et al.*, "Mimo indoor propagation prediction using 3d shoot-and-bounce ray (sbr) tracing technique for 2.4 ghz and 5 ghz," in *EUCAP*, 2011.

[16] Z. Yun and M. F. Iskander, "Ray tracing for radio propagation modeling: principles and applications," *IEEE Access*, 2015.

[17] Y. Tao, H. Lin, and H. Bao, "Gpu-based shooting and bouncing ray method for fast rcs prediction," *IEEE Transactions on Antennas and Propagation*, 2010.

[18] "3d robotics - x8+ drones." https://3dr.com/wp-content/uploads/2016/02/X8-Operation-Manual-vC.pdf.

[19] "Securifi almond." https://www.securifi.com.

[20] "Top 1% of mobile users consume half of worlds bandwidth, and gap is growing." https://nyti.ms/2ioEoOp.

[21] F. Rashid-Farrokhi, K. R. Liu, and L. Tassiulas, "Transmit beamforming and power control for cellular wireless systems," *IEEE JSAC*, 1998.

[22] "Google loon." https://www.solveforx.com/loon/.

[23] "Facebook internet drones." https://www.wired.com/2016/07/facebooks-giant-internet-beaming-drone-finally-takes-flight/.

[24] "Darpa landroids." https://defense-update.com/products/l/landroids.htm.

[25] "Vehicle mounted antenna mast." https://goo.gl/wrx4c2.

[26] A. Dhekne, M. Gowda, *et al.*, "Cell tower extension through drones: Poster," in *Poster, Mobicom*, 2016.

[27] M. Gowda, N. Roy, and R. R. Choudhury, "Infrastructure mobility: A what-if analysis," in *ACM Hotnets*, 2014.

[28] M. Gowda, A. Dhekne, and R. Roy Choudhury, "The case for robotic wireless networks," in *WWW*, 2016.

[29] A. Y. Chung, J. Jung, *et al.*, "Poster: Swarming drones can connect you to the network," in *Mobisys*, 2015.

[30] G. Durgin *et al.*, "Measurements and models for radio path loss and penetration loss in and around homes and trees at 5.85 ghz," *IEEE Transactions on Communications*, 1998.

[31] L. J. Greenstein *et al.*, "A new path-gain/delay-spread propagation model for digital cellular channels," *IEEE Transactions on Vehicular Technology*, 1997.

[32] V. Erceg, L. J. Greenstein, S. Y. Tjandra, S. R. Parkoff, *et al.*, "An empirically based path loss model for wireless channels in suburban environments,"

[33] K. R. Schaubach, N. Davis, *et al.*, "A ray tracing method for predicting path loss and delay spread in microcellular environments," in *IEEE VTC*, 1992.

[34] T. Rautiainen *et al.*, "Verifying path loss and delay spread predictions of a 3d ray tracing propagation model in urban environment," in *IEEE VTC*, 2002.

[35] G. Liang and H. L. Bertoni, "A new approach to 3-d ray tracing for propagation prediction in cities," *IEEE Transactions on Antennas and Propagation*, 1998.

[36] K. Rizk *et al.*, "Two-dimensional ray-tracing modeling for propagation prediction in microcellular environments," *IEEE Trans on Vehicular Tech.*, 1997.

[37] M. Catedra *et al.*, "Efficient ray-tracing techniques for three-dimensional analyses of propagation in mobile communications: application to picocell and microcell scenarios," *IEEE Antennas and Propagation Magazine*.

[38] V. Namboodiri and L. Gao, "Prediction-based routing for vehicular ad hoc networks," *IEEE Transactions on Vehicular Technology*, 2007.

[39] D. Halperin, W. Hu, A. Sheth, and D. Wetherall, "Predictable 802.11 packet delivery from wireless channel measurements," *ACM SIGCOMM*, 2010.