SIGCOMM: G: Many-to-Many Beam Alignment in Millimeter Wave Networks

Suraj Jog
PhD Advisor: Haitham Hassanieh
University of Illinois at Urbana-Champaign

Abstract — Millimeter Wave networks can deliver multi-Gbps wireless links that use extremely narrow directional beams. This provides us with a new opportunity to exploit spatial reuse where we can enable many wireless links to communicate in parallel in confined spaces, without interfering with each other. Exploiting such spatial reuse, however, requires aligning the beams of all nodes in the network. Aligning the beams is a difficult process which is complicated by indoor multipath, which can create interference, as well as by the inefficiency of interference detection schemes like carrier sense in directional networks. This paper presents BounceNet, the first many-to-many millimeter wave beam alignment protocol that can exploit dense spatial reuse and can scale the wireless network throughput with the number of clients. Results from three millimeter wave testbeds show that BounceNet can scale the throughput with the number of clients to deliver a total network data rate of more than 39 Gbps for 10 clients, which is up to 6.6× higher than current 802.11 mmWave standards.

1. INTRODUCTION

Millimeter wave (mmWave) is emerging as the de facto technology for next generation wireless networks [11, 17]. The abundance of bandwidth available in mmWave frequencies (above 24 GHz) has led to the design of wireless radios that can operate at several Gbps [2, 22], and as a result, mmWave will significantly change the future of wireless LANs by delivering links at fiber-like speed. This will enable new applications like multi-user wireless VR for education and professional training, where high bandwidth data must be streamed to each user in real-time [5, 13]. It will also enable large scale robotic factory automation, where many robots stream continuous real-time video back to servers that run AI algorithms to generate decisions to coordinate the robots and accomplish collaborative tasks [16, 23].

However, the above vision requires scaling mmWave networks from a single communication link to a network of many links without compromising the throughput per user. Fortunately, mmWave radios use very directional steerable narrow beams to focus their power. This presents a significant new opportunity for exploiting dense spatial reuse to enable many links to simultaneously communicate at multi-Gbps data rates. Consider the example in Fig. 1(a). In the current broadcast model for 802.11 WLANs, whenever a node is transmitting, all other nodes must stay silent to avoid interference. With more users, the throughput is divided since the entire medium is shared. In contrast, the use of very narrow beams in mmWave networks allows several Access Points (APs) and clients to communicate simultaneously on the same channel without interfering as shown in Fig. 1(b). Hence, mmWave can potentially scale the network throughput with the number of users by adding more APs.

Figure 1: Spatial reuse in traditional WiFi vs mmWave networks.

The directional nature of communication, however, brings its own new challenges. Millimeter wave APs and clients need to align their narrow beams towards each other in order to communicate at very high data rates. Past mmWave research focused on how to quickly find the best direction to align the beams for a single communication link [1, 10, 20, 21]. However, in a network with multiple links, selfishly choosing the best alignment for each AP-client pair independent of other APs and clients can create interference that severely harms the throughput of the interfering links. First, due to multipath reflections, even if two nodes are transmitting in completely different directions, their packets might still collide. Further, interference detection schemes like carrier sense (also known as *Listen-Before-Talk*) are ineffective in mmWave networks since the narrow directional beams prevent mmWave radios from hearing nearby transmissions, unless these transmissions are specifically directed towards them [3, 15]. As a result, we cannot directly apply single-link alignment schemes in larger networks by simply aligning the beams of each AP-client pair independently.

In this paper, we introduce BounceNet, the first many-to-many millimeter wave beam alignment protocol that efficiently aligns the beams of many APs and clients, allowing them to simultaneously communicate without interfering. To achieve this, we must address two key questions:

(1) How to align the beams of all the APs and clients in 3D space to densely pack as many links as possible while avoiding interference? The challenge arises from the fact that the choice of beam alignment at any node is intertwined with the choices at other APs and clients. To address this, BounceNet leverages the sparsity in the mmWave channel to perform *Physical Signal Routing*. Past work shows that mmWave signals travel along a small number of paths, e.g., 2 or 3 paths [19], between any two nodes in the network. BounceNet leverages this sparsity to reformulate the many-to-many alignment problem as a signal level routing problem

at the physical layer, where wireless signals are routed along different "air paths"[1] in a manner that maximizes network throughput. Specifically, BounceNet computes a combination of direct and reflected paths to route the wireless signals of the links in a manner that squeezes in as many links as possible to communicate simultaneously without interfering.

However, simply maximizing the number of links that can operate concurrently might force some APs and clients to always communicate along reflected paths, which typically achieve lower data rates. Therefore, to ensure fairness between clients, BounceNet generates several combinations of direct and reflected path routings and switches between them (shown in Fig. 4(b)), such that each client gets a fair share to communicate on its highest data rate path.

(2) How does BounceNet quickly learn the paths and interference patterns in the network in order to adapt the beam alignment in dynamic and mobile environments? To deal with dynamic environments, today's networks periodically perform a beam search (typically every 100 millisecond) to learn the paths along which an AP-client pair can communicate. However, in a network with multiple links, it is not sufficient to learn the paths only between each AP and its corresponding client. In order to avoid interference, we also need to learn the paths between every other node pair in the network, since these are the paths that create interference. As a result, for a network of N APs and clients, the beam search process must be performed $O(N^2)$ times, and this results in prohibitively large overhead.

Therefore, instead of performing the search for each node pair independently, BounceNet redesigns the beam search protocol to jointly find all the paths between the nodes. By coordinating the APs and sharing measurements, BounceNet is able to amortize the cost of the beam search and reduce it to $O(N)$, as we describe in Section 3. Since the beam search is inherent to mmWave and is required to maintain connectivity between clients and APs, BounceNet's design does not introduce any additional overhead compared to current standards [12]. This allows BounceNet to quickly learn the paths and reconfigure the beam alignment to maintain high throughput.

We have designed BounceNet to be backward compatible with the current mmWave wireless LAN standard 802.11ad/ay making it easy to integrate into future standards. We have implemented BounceNet by using extensive real measurements from three indoor wireless testbeds. Our results demonstrate the first design of a wireless LAN that can deliver more than 39 Gbps for 10 clients, thus showcasing BounceNet's ability to scale the overall network data rate with the number of clients.

[1]Routing signals along an isolated "air path" is possible in mmWave due to 1) the sparsity in the mmWave channel, and 2) the lack of scattering effects at mmWave frequencies which ensures the signal reflects off obstacles and does not scatter in many directions [19].

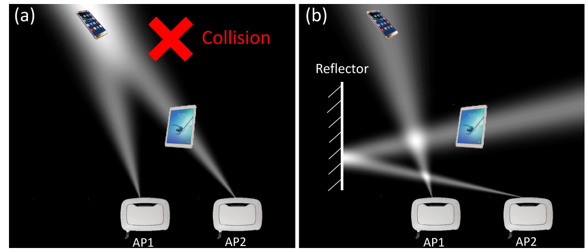

Figure 2: (a) Interference caused by naive alignment along direct paths (b) Routing signals along both direct and indirect paths to achieve higher network throughput.

2. BOUNCENET'S PHYSICAL SIGNAL ROUTING

BounceNet's objective is to route signals to/from clients in a manner that maximizes the number of AP-client pairs that can communicate simultaneously. The choice of routing will govern the many-to-many beam alignment. Past work aligns the beams of each AP-client pair (link) independently by selfishly choosing the best alignment for each link (typically the direct Line-of-sight alignment). However, this can create interference that is hard to deal with in directional networks [3, 15]. Consider the example in Fig. 2(a), where both links communicating along their Line-of-sight paths causes interference, and as a result, cuts the throughput of each link in half. However, routing the signal of AP2 along an indirect path by *bouncing* it off a reflector as shown in Fig. 2(b), can enable both links to operate simultaneously without interfering. Hence, we can get significantly higher network throughput by routing physical signals of the links along both direct and indirect paths, as opposed to each link naively aligning their beams along their respective direct paths.

To extend this to a network with large number of links, we formulate the problem using a conflict graph $G(V,E)$. For every AP-client pair in the network, we define a supernode in the conflict graph as shown in Fig. 3(a). In our formulation, we represent each path that the AP-client pair can use to communicate, as a vertex $v \in V$ in the graph. As a result, each supernode comprises of one vertex corresponding to the direct path, and a vertex each for every indirect path between the AP-client pair. We represent interference in the network using edges $e \in E$. Therefore, an edge between two vertices in the conflict graph means that signals cannot be routed along the two corresponding paths simultaneously. Further, note that the conflict graph also has edges between vertices belonging to the same supernode, because each link can be routed only along one path at a time.

The best beam alignment that maximizes network throughput is equivalent to solving for the Maximum Independent Set on the conflict graph $G(V,E)^2$. However, computing the Maximum Independent Set belongs to the class of NP-Hard problems [7], and as the number of links grows, $G(V,E)$ becomes dense and the problem becomes very hard to solve. Therefore, in order to simplify the problem, we design our solution based on the following intuition: *Direct Paths achieve higher data rates as compared to Indirect*

[2]The Maximum Independent Set is defined as the largest set of vertices in the graph such that no two of the vertices have an edge.

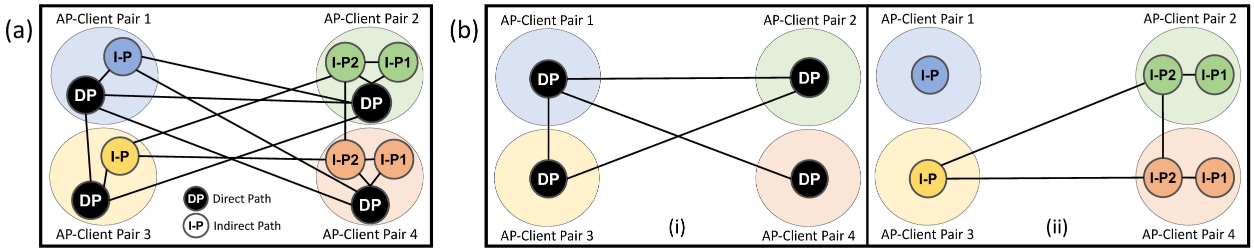

Figure 3: (a) Conflict graph $G(V,E)$ with supernode defined for each AP-Client pair; (b) Multi-Layer conflict graph by decomposing into (i) *Direct Path Conflict Graph:* Subgraph induced on vertices corresponding to direct paths in $G(V,E)$ (ii) *Indirect Path Conflict Graph:* Subgraph induced on vertices corresponding to indirect paths in $G(V,E)$.

Paths due to higher SNR, and therefore we must prioritize routing signals over the direct paths.

Towards this end, we decouple the routing along the direct and indirect paths by decomposing the conflict graph into a multi-layered graph construct as shown in Fig. 3(b). This decoupling step simplifies the problem substantially, since the resulting Direct Path and Indirect Path conflict graphs become very sparse. As a result, with very high probability these graphs are chordal. For such class of graphs, the Maximum Independent Set problem is no longer NP-hard and can be solved efficiently in linear time [7]. Hence, BounceNet's routing is performed in two phases:

1. *Direct Path Routing:* BounceNet routes as many links as possible along the direct paths first, by solving the Maximum Independent Set on the Direct Path Conflict Graph.
2. *Indirect Path Routing:* BounceNet then makes use of the Indirect Path Conflict Graph to squeeze in the remaining links by routing their signals along indirect paths, such that all links can communicate without interfering.

However, while this allows BounceNet to maximize the number of links operating concurrently, it may force some clients to always communicate along reflected paths which typically achieve lower data rates. Therefore, to ensure fairness between links, BounceNet reformulates the routing problem as a Weighted Maximum Independent Set problem on the multi-layered conflict graph, to generate multiple different many-to-many beam alignments. BounceNet switches between these alignments in different time slots, such that the scheduling ensures that each client gets a fair chance to transmit on its highest data rate path.

3. Learning Paths & Interference

To enable a practical system, BounceNet must be able to quickly learn the paths and interference patterns in the network in real-time, in order to construct the conflict graph and compute the optimal signal routing. Today's networks periodically perform a beam search every 100 millisecond (ms) to learn the directions of the paths along which an AP-client pair can communicate. At a high level, the beam search process works as follows. The AP starts by transmitting announcement frames in all directions by sequentially sweeping its narrow beam along different sectors. During this phase, the client listens to the channel in all directions using a quasi-omnidirectional beam pattern, so that the client can discover all directions which the AP can use to communicate with it. The above process is repeated with the client and AP switching roles, so as to discover the directions of the paths from the client side.

However, in a network with multiple links, it is not sufficient to learn only the paths between each AP and its corresponding client. We also need to learn the paths between the other node pairs in the network, since these paths can create interference between links. Past work has leveraged sparsity to speed up the beam search for a single communication link and reduce it to a millisecond [10]. However, for a network of N APs and clients, the search needs to be performed $O(N^2)$ times. For $N = 10$, even with fast algorithms like [10], the overhead is 100 ms which is prohibitively high.

To address this, BounceNet redesigns the beam search protocol to jointly find the paths between all node pairs, instead of performing the search for each node pair separately. Specifically, BounceNet ensures that whenever a node m is performing its sequential scan, every other node in the network (both APs and clients) listen in the quasi-omnidirectional mode to discover the directions that m can use to communicate with them. This process continues with every node in the network performing the scan one after the other while the remaining nodes listen in the quasi-omnidirectional mode. After all nodes finish their scans, the information is aggregated at a lead AP, which can then compute the paths between every node pair in the network, since it knows the directions that any node can use to communicate with any other node.

Therefore, BounceNet can now learn the paths between every AP-client pair, as well as the paths that lead to interference between the different communication links. As a result, BounceNet can construct the conflict graph and compute the optimal many-to-many alignment as described in Section 2. Note that BounceNet needs to perform only $O(N)$ beam scans to learn all the paths, since every node does the sequential scan only once. By coordinating the beam scan process, BounceNet is able to amortize the cost of the search and reduce the overhead from $O(N^2)$ to $O(N)$.

4. BounceNet's Architecture

In this section, we will piece together the components of BounceNet and present the overall flow of the system. BounceNet is designed to be backward compatible with the standard, and consequently follows the same high level structure as the standard. Time is divided into transmission

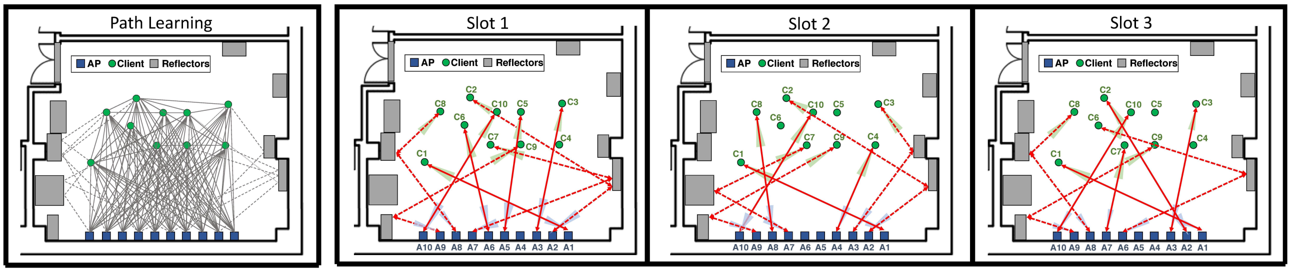

(a) Association Phase (b) Data Transmission Phase

Figure 4: Beacon Interval consisting of (a) Association Phase and (b) Data Transmission Phase.

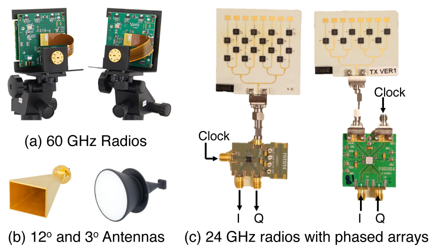

Figure 5: Experimental hardware used to evaluate BounceNet.

cycles (typically referred to as Beacon Intervals), with each cycle consisting of an Association Phase, followed by a Data Transmission Phase as shown in Fig. 4.

BounceNet uses the Association Phase to perform the path learning protocol from Section 3. Following the algorithm from Section 3, BounceNet can learn both the direct and indirect paths between every node pair in the network (shown in Fig. 4(a)) with only $O(N)$ scans. For a network with $N = 10$, the resulting overhead is only about 10 ms, as opposed to hundreds of ms required by the standard. With information about the paths, BounceNet can compute the signal routing as described in Section 2. In order to ensure fairness between links, BounceNet computes multiple many-to-many alignments and schedules them in different time slots during the Data Transmission Phase (Fig. 4(b)). The entire process repeats every Beacon Interval (every 100 ms). Hence, BounceNet is able to adapt to client mobility since it obtains fresh measurements from the beam scanning process and can re-learn the paths every 100 ms.

5. Evaluation and Results

We evaluated BounceNet using extensive real measurements from three indoor wireless testbeds (Fig. 5):

- A 60 GHz testbed with 3° beam directional antennas.
- A 60 GHz testbed with 12° beam directional antennas.
- A 24 GHz testbed with 8-element phased arrays.

We evaluate BounceNet in a 860 sq.ft. lab space, with the floor plan shown in Fig. 4. The APs were deployed along the wall of the lab with the clients scattered across the room.

I. BounceNet's Signal Routing in Practice: In Fig. 4(b), we show an example of BounceNet's Physical Signal Routing in practice from our 12° beam 60 GHz testbed for the case of 10 APs and clients in the network. We pick a random client configuration and plot the beam alignments computed by BounceNet for the first three time slots of the Data Transmission Phase. We can see that BounceNet leverages both direct and indirect paths to squeeze in as many links as possible for communication during the time slot. Further, over the three time slots, BounceNet schedules the direct paths for different clients, thus ensuring fairness among clients.

II. Total Network Data Rates Achieved: We also evaluate the aggregate data rate of all clients in the network. We vary the number of APs and clients from 1 to 10. We test on a total of 5000 random configurations of client locations in the lab space. We compare BounceNet against two schemes:

1. **802.11ad:** The standard does not leverage multipath reflections to align beams, thus resulting in sub-optimal alignments. Further, the standard uses a greedy mechanism for exploiting spatial reuse by measuring pairwise mutual interference between links [12]. This results in an overhead of $O(N^2)$ measurements, and hence, the standard is unable to exploit spatial reuse effectively.
2. **Baseline:** We also compare against a baseline where the beams of each AP and client are aligned independently along their respective best paths, and carrier sense is used for interference avoidance.

Fig. 6(i) shows the total network data rate as a function of the number of APs and clients in the network for each testbed. As the number of clients increases, BounceNet is able to scale the network data rate to deliver a total data rate of 39.2 Gbps and 32.8 Gbps for 10 clients using 60 GHz with 3° and 12° beams respectively. For 24 GHz, BounceNet achieves 18.2 Gbps for 10 clients. This is expected as side-lobe leakage from phased array antennas creates more interference in the network which limits spatial reuse.

802.11ad, on the other hand, is unable to properly exploit spatial reuse and shows limited gains. Specifically, for the case of 10 clients, BounceNet achieves 6.6×, 5×, and 3.1× gain in network throughput as compared to 802.11ad for 3° beam, 12° beam, and the phased array respectively. The baseline can exploit spatial reuse for 3° beam since the interference in this case is very limited. Hence, for 10 clients with 3° beam, BounceNet only achieves 1.27× gain over the baseline. This gain, however, increases to 2.7× and 3.4× for 12° beam and the phased array respectively, where there is more interference, and carrier sense is unable to avoid packet collisions. In fact, the baseline is unable to exploit spatial reuse and scale network throughput in such cases.

While the baseline provides higher total network data rate

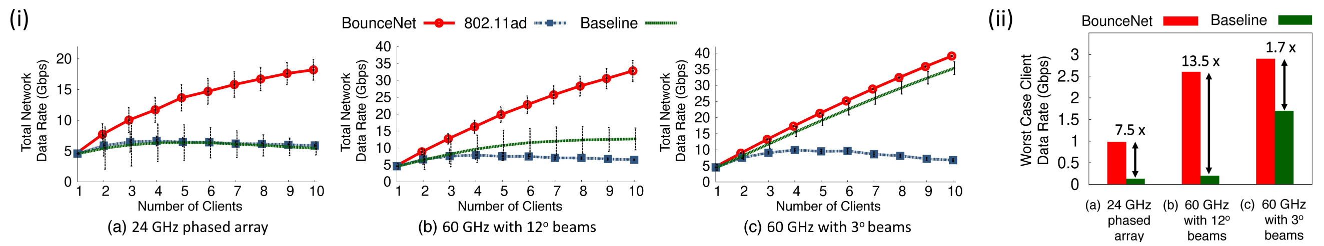

Figure 6: (i) Total network data rates for BounceNet, 802.11ad and Baseline, (ii) Median data rate of worst-case client in testbed for BounceNet and Baseline.

as compared to the standard, it does so at the cost of fairness. In Fig. 6(ii), we show the median value of the data rate achieved by the worst case client in the network, for BounceNet and the baseline. BounceNet improves the data rate of the worst case client by 7.5× for phased arrays and 13.5× for 12° beam as compared to the baseline. This is because the baseline does not try to avoid interference, and instead simply aligns each link independently. As a result, clients that interfere suffer from repeated packet collisions since interference avoidance schemes like carrier sense do not work in directional networks. BounceNet on the other hand, significantly benefits the worst case clients since it explicitly accounts for fairness in the network.

6. RELATED WORK

Millimeter Wave Networks: BounceNet is related to recent work on increasing the speed and robustness of beam alignment in mmWave networks to enable mobility [8, 10, 20, 27] and avoid blockage [1, 14, 25] All this work, however, focuses on a single communication link. BounceNet is the first to demonstrate many-to-many beam alignment and scale the network to multiple links. It is complementary to these systems and can benefit from faster beam search.

Our work is also related to recent mmWave work that deploys multiple APs to deal with blockage [24, 26]. However, [24] requires brute-force training to map all reflectors in the environment, a process that needs to be repeated every time the environment changes. [26] addresses blockage by having multiple APs jointly transmit the same signal to the client. However, it requires tight phase and frequency synchronization, which is hard to achieve in practice [18]. BounceNet on the other hand does not require phase, frequency or packet level synchronization. BounceNet is also able to learn the reflectors in the environment in real-time without any pretraining. There is also past work that uses 60 GHz wireless links in data centers [6, 9]. However, data centers have static topologies with predictable interference models, which does not apply to 802.11 LANs which have mobile clients.

WLANs with Directional Antennas: Past work has designed protocols for mobile ad-hoc networks and WLANs with directional antennas [3, 4]. However, these works assume the locations of the clients are known apriori, and ignore multipath effects. [15] leverages directional antennas at 2.4 GHz to increase spatial reuse. However, [15] assumes only APs to have directional antennas which simplifies the problem since the clients can easily perform interference detection in the omnidirectional mode.

7. CONCLUSION

In this paper, we introduced BounceNet, the first many-to-many mmWave beam alignment system that can efficiently align the beams of many APs and clients allowing them to simultaneously communicate without interfering. We evaluated BounceNet using three wireless testbeds and demonstrated that it can enable extremely dense spatial reuse and scale the network throughput with the number of clients.

REFERENCES

[1] O. Abari, D. Bharadia, A. Duffield, and D. Katabi. Enabling high-quality untethered virtual reality. In *NSDI*, 2017.
[2] Acer. TravelMate P6 TMP658-M-70S3 Laptop.
[3] R. R. Choudhury and N. H. Vaidya. Deafness: A MAC problem in ad hoc networks when using directional antennas. In *ICNP*, 2004.
[4] R. R. Choudhury, X. Yang, R. Ramanathan, and N. H. Vaidya. On designing MAC protocols for wireless networks using directional antennas. *IEEE Transactions on Mobile Computing*, 5(5):477–491, 2006.
[5] A. Connor-Simons. Enabling wireless virtual reality, MIT News, 2016.
[6] Y. Cui, S. Xiao, X. Wang, Z. Yang, C. Zhu, X. Li, L. Yang, and N. Ge. Diamond: Nesting the data center network with wireless rings in 3D space. In *NSDI*, 2016.
[7] A. Frank. Some polynomial algorithms for certain graphs and hypergraphs. In *Proceedings of the Fifth British Combinatorial Conference*, 1975.
[8] M. Haider, Y. Ghasempour, D. Koutsonikolas, and E. Knightly. Listeer: mmWave beam acquisition and steering by tracking indicator LEDs on wireless APs. In *ACM MobiCom*, 2018.
[9] D. Halperin, S. Kandula, J. Padhye, P. Bahl, and D. Wetherall. Augmenting Data Center Networks with Multi-Gigabit Wireless Links. In *ACM SIGCOMM*, 2011.
[10] H. Hassanieh, O. Abari, M. Rodriguez, M. Abdelghany, D. Katabi, and P. Indyk. Fast millimeter wave beam alignment. In *ACM SIGCOMM*, 2018.
[11] K. Hill. A look at Verizon's fixed millimeter wave testing. RCR Wireless News, May 2017. Press Release.
[12] IEEE Standards Association. IEEE Standards 802.11ad-2012: Enhancements for Very High Throughput in the 60 GHz Band, 2012.
[13] S. Jog, J. Wang, H. Hassanieh, and R. R. Choudhury. Enabling Dense Spatial Reuse in mmWave Networks. In *ACM SIGCOMM*, 2018.
[14] S. Kwon and J. Widmer. Multi-beam power allocation for mmwave communications under random blockage. In *IEEE VTC Spring*, 2018.
[15] X. Liu, A. Sheth, M. Kaminsky, K. Papagiannaki, S. Seshan, and P. Steenkiste. DIRC: Increasing indoor wireless capacity using directional antennas. *ACM SIGCOMM Computer Communication Review*, 39(4):171–182, 2009.
[16] B. Manz. 5G Cellular Networks Are the Future of Robotics, Mouser Electronics, 2016.
[17] Markets and Markets. Millimeter Wave Technology Market worth 4,632.8 Million USD by 2022, Press Release, 2017.
[18] H. S. Rahul, S. Kumar, and D. Katabi. JMB: scaling wireless capacity with user demands. In *ACM SIGCOMM*, 2012.
[19] S. Rangan, T. S. Rappaport, and E. Erkip. Millimeter-wave cellular wireless networks: Potentials and challenges. *IEEE*, 2014.
[20] M. E. Rasekh, Z. Marzi, Y. Zhu, U. Madhow, and H. Zheng. Noncoherent mmwave path tracking. In *ACM HotMobile*, 2017.
[21] S. Sur, V. Venkateswaran, X. Zhang, and P. Ramanathan. 60 GHz Indoor networking through flexible beams: A link-level profiling. In *SIGMETRICS*, 2015.
[22] TP-Link. Talon AD7200 Multi-Band Wi-Fi Router.
[23] V. Turk. These Supermarket Warehouse Robots Have Their Own Mobile Network, Vice Motherboard, 2016.
[24] T. Wei and X. Zhang. Pose Information Assisted 60 GHz Networks: Towards Seamless Coverage and Mobility Support. In *MobiCom'17*, 2017.
[25] T. Wei, A. Zhou, and X. Zhang. Facilitating Robust 60 GHz Network Deployment By Sensing Ambient Reflectors. In *NSDI*, 2017.
[26] D. Zhang, M. Garude, and P. H. Pathak. mmChoir: Exploiting joint transmissions for reliable 60 GHz mmwave WLANs. In *MobiHoc*, 2018.
[27] A. Zhou, L. Wu, S. Xu, H. Ma, T. Wei, and X. Zhang. Following the shadow: Agile 3-d beam-steering for 60 GHz wireless networks. *IEEE INFOCOM*, 2018.