UNIVERSITY OF CALIFORNIA SAN DIEGO

Demystifying Millimeter-Wave V2X: Testbed Measurement and Large-Scale Simulation

A thesis submitted in partial satisfaction of the
requirements for the degree
Master of Science

in

Electrical Engineering (Communication Theory and Systems)

by

Jingqi Huang

Committee in charge:

Professor Xinyu Zhang, Chair
Professor Sujit Dey
Professor Dinesh Bharadia

2020

The thesis of Jingqi Huang is approved, and it is acceptable in quality and form for publication on microfilm and electronically:

_____

_____

_____

Chair

University of California San Diego

2020

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGEMENTS

I sincerely thank my supervisor Professor Xinyu Zhang and my colleague Song Wang for their support in not only this thesis, but also in my research during my master study. I would also like to thank Professor Sujit Dey and Professor Dinesh Bharadia for serving as my thesis defense committee.

The thesis is currently being prepared for submission for publication. Song Wang; Jingqi Huang; Xinyu Zhang. The dissertation/thesis author was the primary investigator and author of this material.

ABSTRACT OF THE THESIS

Demystifying Millimeter-Wave V2X: Testbed Measurement and Large-Scale Simulation

by

Jingqi Huang

Master of Science in Electrical Engineering (Communication Theory and Systems)

University of California San Diego, 2020

Professor Xinyu Zhang, Chair

Recent year has witnessed an explosive growth in the demand for intelligent vehicles. In particular, multi-gigabit throughput and high reliability of vehicle communications are the key requirements to enable intelligent vehicles. Although mmWave is believed to be the right track to meet the extreme throughput requirement, its reliability for vehicle communications is still unclear due to the directional beams and limited field-of-view (FoV). In this thesis, we present the first comprehensive reality check on millimeter-wave vehicle-to-everything (mmWave V2X) communications. We leverage both real-world measurements using COTS devices and large scale simulation using high-fidelity 3D ray-tracer to explore the feasibility of mmWave V2X. Specifically, we study the basic performance, impact of vehicle mobility and obstacle

blockage, the effect of codebook/beam management and multi-array basestaion, and the potential solutions to MU-MIMO and interference management. By extensive evaluations, we correct some common misunderstandings and show that mmWave V2X is feasible and its major concerns can be addressed by simple solutions.

Chapter 1

Introduction

Vehicle-to-Everything (V2X) connectivity represents a key vertical application of 5G. The 5G standard has identified 4 categories of V2X use cases [3GP18]: vehicles platooning, extended sensors, advanced driving and remote driving. All these cases require high-throughput data synchronization between connected vehicles and road infrastructures. For example, state-of-the-art self-driving cars can generate up to 1000 megabytes of sensor data per second [3GP17c], from LIDAR, camera, and other onboard units. By sharing the data in real-time, each vehicle's perceptual range can be extended, creating a smoother and safer transportation system [SVBH18, AQE+20, YLD+19, QAB+18]. 5G millimeter-wave (mmWave) networking is a niche technology to meet such demanding applications, due to its multi Gbps of bit-rate.

However, the practicality of mmWave V2X has long been questioned [VSBHJ16, SVBH18, VCS+17, GPMRH16, AG18a]. mmWave is known for its intrinsic limitations, *i.e.*, high attenuation loss, high directionality, prune to blockage, *etc.* These properties seemingly contradict the high mobility nature of V2X. Consequently, although mmWave network protocols (*e.g.*, 802.11ad and 802.15.3c) and commercial devices have existed for more than one decade, the real-world use cases are limited to quasi-stationary point-to-point scenarios, *e.g.*, cordless HDMI and wireless backhaul. The specific challenges for mmWave V2X include the following.

(1) Intractable beam management under high mobility. To combat high attenuation, mmWave radios have to use phased array antennas to form highly directional beams. As a vehicle user equipment (UE) moves, the basestation has to constantly search for the best beam that aligns with the UE's. Standard beam searching protocols traverse all possible beams [IEE12, IEE18, GPR+18] and often criticized for the large overhead for even under human mobility [ZWX+18, SVZR15, HAR+18, WZ17]. Intuitively, since the vehicle UE moves much faster, the beam alignment needs to be done even more frequently, thus becoming intractable. This is often considered as the most significant challenges for mmWave V2X [VSBHJ16, GPMRH16, VCS+17, AG18a, VZH15, AMS+18, GPR+18].

(2) Blockage recovery. The mmWave directional beams are vulnerable to blockage by roadway obstacles including foliage, pedestrian, and tall vehicles [WVMH17]. Unlike indoor environments with rich multi-paths, the blocked link cannot be easily recovered as reflection paths are sparse in V2X environment [GPR+18, EANH16, AMB17, AG18b].

(3) Limited coverage. Due to high attenuation/penetration loss, 5G mmWave basestations have a much shorter range compared with 4G LTE. Nonetheless, a quantitative evaluation is still lacking, with respect to the exact deployment density needed to achieve reasonable coverage.

(4) High complexity in implementing multiplexing mechanisms. The use of directional beams opens up new opportunities to achieve ultra-fine-grained *spatial division multiplexing (SDM)*. However, such extreme SDM requires building an interference map between every pair of basestation and UE, and for each beam they use. The complexity of such SDM easily becomes unmanageable, especially under high mobility. *MIMO mmWave* [AEALH14, AMGPH14] is another way to achieve multiplexing gain, by using an array of phased arrays to serve multiple UEs simultaneously. Similar to SDM, it suffers from a curse of dimension, and the overhead may easily outweigh the capacity gain.

In this thesis, we demystify the above challenges on a mmWave V2X testbed together with large-scale 3D ray-tracing simulation. Our experimental testbed runs in one typical suburban

scenario and characterizes both the channel dynamics and higher layer behaviors. Leveraging existing 3D environment models, our simulation framework reproduces a large-scale $6.7km^2$ site with a mix of urban/highway scenarios and synthetic roadway traffic. Our experiments confirm the practicality of mmWave V2X with a little to none modification to the existing exhaustive beam searching mechanism. Drawing on the experimental results, we identify unique opportunities to simplify the mmWave V2X operations, including spatial reuse, MU-MIMO and interference management, without sacrificing performance.

Our key conclusions are as follow:

(1) mmWave V2X works, even with heuristic beam management. Our real measurement shows over 700Mbps of stable throughput under 50mph vehicle velocity. Our result shows less than 5% of beam searching overheads for various scenarios and vehicle mobilities.

(2) mmWave V2X links have an overall blockage probability of 4%. During link blockage, the basestation can quickly steer beams to recovery paths with less than 5% overhead. The blockage can be further alleviated by raising basestation or multi-connectivity.

(3) Spatial reuse, multiplexing and interference management can be addressed by heuristic ways without much performance loss.

The rest of this thesis is organized as follows. We introduce our evaluation methodology in Chapter 2, and then present the link throughput and coverage elevation in Chapter 3. We evaluate mobility in Chapter 4, which provides guidance for codebook/beam management design in Chapter 5. Blockage is also evaluated and discussed in Chapter 6. For multi-array beasesations and multiplexing and inter-cell interference, we derive design hints from the simulations and further verify them in Chapter 7, Chapter 8 and Chapter 9. We discuss related work in Chapter 10 and conclude the thesis in Chapter 11.

Chapter 2

Methodology

Our results are derived from two different sets of experiments: a real-life mmWave V2X testbed, and a large scale high-fidelity ray-tracing simulation. Our mmWave V2X testbed is deployed on a university campus with a mixture of typical urban and suburban environments. With the COTS mmWave radios equipped, the testbed can capture wireless channel dynamics and characterize higher layer behaviors. We introduce the details of the testbed in Sec.2.1.

Limited by the deployment and hardware, the real-life testbed only runs in one typical scenario and does not provide CSI measurement, which is crucial to MU-MIMO analysis. Hence we create large scale ray-tracing simulations that cover a 2.8x2.4 km^2 area with diverse environments. With professional traffic generator and real-life 3D environment models, we can investigate the channel dynamic under different scenarios (*e.g.* highway, urban, suburban, *etc.*) with different settings (*e.g.* AP height, traffic density, *etc.*) and perform statically analysis. We introduce the details of the simulation in Sec.2.2.

2.1 mmw V2X Testbed

Our mmWave V2X testbed is deployed on a 1km long road segment, consisting of 2 vehicle lanes and 2 bike lanes. Fig. 2.1 shows a map of the deployment. The nearby landscape

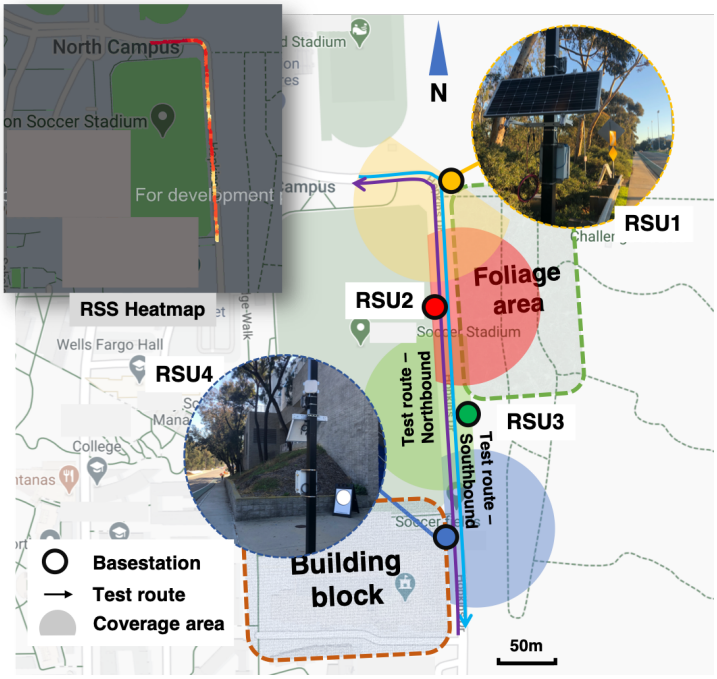

Figure 2.1: mmw V2X AP deployment.

represents a mixture of typical urban scenarios with several 5-story buildings on the south, and suburban scenarios with foliage and football court on the north.

The mmWave basestations are mounted on the lampposts along the road, 3-4m above the ground. Two neighboring basestations are on the opposite sides of the road, separated by 80-100m. Each basestation covers up to 200m in LoS when the Tx beam with maximum gain is used and Rx beam under quasi-omni mode. The deployment of mmWave basestations follows the Urban Micro (UMi) scenario in 3GPP TR 38.901 channel model [3GP17a] except for the height which is constrained by the lampposts (basestation height is 10m in TR 38.901).

The mmWave basesation is designed to be lightweight, off-grid, and low cost. As illustrated in Fig. 2.2, it consists of 2 COTS mmWave radios, an outdoor WiFi mesh router, a 60W solar panel, and a weather-proof enclosure hosting a high-capacity backup battery and various switches/wires. The basestation is self-contained and can operate on solar power alone. The total cost of all its components is less than $2000. Below we provide more details on the hardware

components.

2.1.1 mmWave Radios

The mmWave basestation is built from the Airfide 802.11ad multi-panel 60 GHz radio [air19]. Each radio has 8 phased array panels arranged in a 2×4 layout. Each panel consists of 6×6 quasi-omni antenna elements with 6 dBi gain and a 2-bit phase shifter. The 4 elements on the corners are disabled, so only 32 elements are usable in effect. The phase shifters can be reconfigured to form up to 128 different beam patterns, each corresponding to a beamforming weights vector in a predefined codebook. The maximal EIRP of beams over all 8 phased arrays is 42 dBm, *i.e.*, the indoor EIRP limit set by FCC. Note that this is lower than the outdoor EIRP limit for the 60 GHz band (83 dBm [Com20]). Although the resulting RSS may be lower than the maximum allowable values, our measurement insights in this thesis are independent of the absolute RSS.

The Airfide radio comprises an 802.11ad NIC (with Qualcomm QCA9500 FullMAC WiGig chip and QCA6335 baseband) and an embedded Linux system running OpenWrt, which accepts command line configurations through SSH. We customized the firmware and driver to enable fine-grained channel measurement (details in Sec.2.1.5).

2.1.2 Solar-Based Power System

Drawing power from the mains requires non-trivial modification of the lamppost infrastructure. We thus build a simpler, self-contained power system comprised of a solar power generator and an energy storage device, *i.e.*, a 256Wh battery. An 802.3at PoE switch is used to simultaneously power and connect the mmWave radios and backhaul router. The battery provides 12V6A DC output which is converted to 48V by an DC/DC converter to power the PoE swithc. Using DC power instead factory AC/DC converter skips the AC/DC rectifier, saves extra space

and energy waste from the rectifier. The DC power system gives the basestation $2\tilde{3}$ hours more of battery life and make it smaller and safer. Both the switch and battery are housed in the waterproof enclosure.

The Newpowa NPA50S-12H solar panel has a footprint of $23.1 \times 19.9 \times 1.2$ inches. It provides 60W peak power and fully charges the battery in 12 hours in a sunny day (with the basestation running). The battery supports $\sim$10 hours of the basestation's normal operations. The entire system can be put into sleep mode through a remote switch.

2.1.3 Low-frequency Control Plane

To enable reliable access and control of the mmWave radios, each basestation co-sits with an EnGenius 802.11ac outdoor router. The router has 4 5dBi dipole antennas with 500m range when running on the 2.4GHz WiFi band. We flush the router with an OpenWrt 19.07.0 system [ope20a] and repurpose it as a mesh router. The mesh interface uses 802.11ac IBSS (ad-hoc) mode, and operates on a L2 routing kernel module *batman-adv* [bat20]. The batman-adv controls the packet routing/forwarding and emulates a virtual switch for all participating nodes with a stable $\leq 5ms$ latency.

The routers from all 4 nodes form a multi-hop control network and can be extended into a more sophisticated mesh topology as more nodes join. We configure the routers' 5GHz interfaces into infrastructure mode. In order to send control commands through the control network, we simply associate a controller PC to one of the routers, which routes commands to the desired basestation. The mesh interfaces and the 5GHz interfaces are bridged so any router in the control network can serve as a gateway for the controller PC to access. More specifically, the controller can configure the radio parameters of each basestation, such as codebook, beacon interval, and beam pattern.

2.1.4 Vehicular Client

The testbed vehicle aims to collect fine-grained statistics of the mmWave channel, the link and network-level performance, along with the coordinates of the vehicle. The vehicular client is equipped with the same Airfide mmWave radio as the basestations. The radio is installed on a car roof rack and can be adjusted to face towards or against the vehicle moving direction. A GPS receiver is used to log the radio location ($\pm 1m$ error) with 10Hz refresh rate.

The client radio can measure the channel statistics (*e.g.*, per beam RSS from the basestation's beacon header frame), as well as the end-to-end performance (*e.g.*, MCS, TCP throughput, and packet loss rate). Due to firmware limitations, the radio has to use a quasi-omni beam in receiving mode. During the data collection, a host PC on the vehicle connects to the radios via its on-board WiFi and the GPS receiver via Bluetooth.

2.1.5 Customizing the mmWave Radio Firmware

Per beam RSS extraction

In real-life measurement, per beam RSS is crucial for evaluating and comparing beams' performance. In COTS mmWave radio, per beam RSS is hidden from user space access by the NIC. We disassemble the NIC's firmware, find the on-board memory address for the RSS measurement and inject a piece of code to the firmware using Talon-tools [SWH17] that transfers the RSS measurement to a memory area accessible by the user space program. We then write a Python program that dumps the RSS measurements repeatedly and transmits it to the controller PC over TCP. Similar RSS extraction approach has been mentioned in [WHZ+20, PSL+18].

Different from previous works [WHZ+20, PSL+18] which run in monitor mode, our testbed requires seamless switch between Per beam RSS extraction and data transfer, which is only possible in managed mode. Unlike monitor mode, in managed mode, the per beam RSS information (extract from radio tap header of 802.11ad BF frames) only updated when

the radio proactively scans for a basestation. We find the correspond NIC command called for proactive scan (*WMI_START_SCAN_CMD*) and insert a command call before every Per beam RSS extraction cycle. Using this RSS extraction method, we can get 33 per beam RSS measurements per second under managed mode.

Basestation MAC address extraction

Previous Per beam RSS extractions [WHZ+20, PSL+18] focus on single AP case. Since our testbed comprises of multiple basestations, the client radios may receive RSS measurements from different basestations and cannot differentiate them. Using the same technique of finding RSS memory address, we find memory address for the entire BF frame, whose header contains a sender's mac address field. We modify our firmware patch so that the Per Beam RSS is extracted along with the MAC address of the corresponding basestation. We then distinguish the per beam RSS measurements by their mac addresses.

Short beam searching periodicity

In 802.11ad, BF frames are transmitted every Beacon Interval (BI). In our radios, BI value is fixed to 100ms, meaning we can only have 10 Per beam RSS measurement per second. To obtain the fine-grained Per beam RSS measurement, we overwrite the memory address of BI value in the firmware and fix it to 30ms. Note that as mentioned in X-Array [WHZ+20], BI value can be set as low as 5ms for stable RSS measurement. However, smaller BI means larger beam searching overhead. In addition, the measurement granularity is already limited by our Per beam RSS extraction method (33 samples/sec). Hence, we choose 30ms to balance the granularity and overhead.

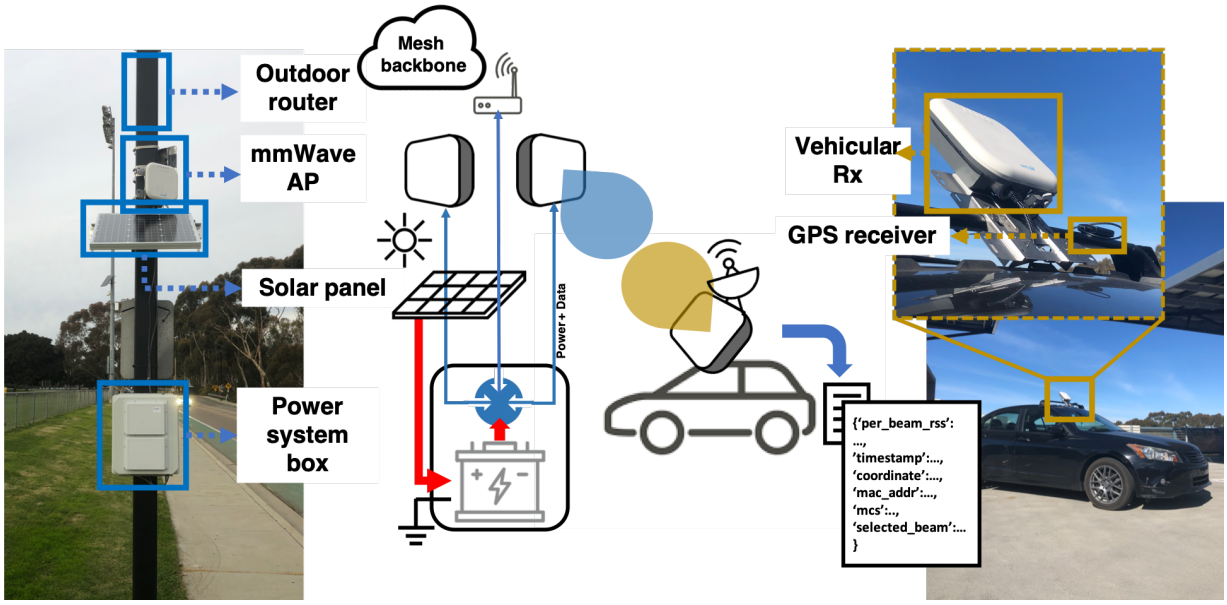

Figure 2.2: mmw V2X AP and client setup.

2.2 Large Scale Simulation

In this section, we discuss how we build a high fidelity simulation framework to reproduce the mmWave channel and network dynamics in a realistic environment. At a high level, our mmWave V2X simulator takes the real-world 3D map as input and places various vehicles on the roads based on the Simulation of Urban MObility (SUMO) [(DL19a] vehicular traffic simulator. We then use *Wireless Insite* [Rem20], a commercial ray-tracing simulator, to track the propagation paths of the mmWave signals and derive the channel characteristics.

2.2.1 Recreating the large-scale environment in ray-tracing simulator

We select a $2.8 \times 2.4 km^2$ campus area as our simulation field. This area includes various of typical environments, *e.g.* urban, suburban, rural, highway, *etc.*. We also include our testbed area in the simulation so that we can compare and validate the simulation result with real measurement. We export the 2D map of this area from OpenStreetMap [Ope20b], which provides detailed information on road networks and building footprints. We then use *blender-osm* [Eli20] to

render the 2D map into 3D model. Fig. 2.3 shows the overview of the 3D model and 3 detailed environments v.s. their corresponding Google 3D map. Note that the 3D model does not contain details of the buildings. But the quasi-optical propagation property of mmWave makes it blind to the details [ITU13]. Similar detail-less ray-tracing simulations are widely used in mmWave simulations and proven to be reliable [HWG+19, KBGP+18].

Next, we import the 3D models to the ray-tracing simulator - Wireless Insite [Rem20] and create simulation environments. Wireless Insite exhaustively ray-traces all possible propagation paths for each pair of Tx/Rx and provides rich path information, *e.g.* RSS, CSI, AoA/AoD, *etc.*. We filter out the paths with path loss higher than $-172dB$ assuming $82dBm$ EIRP [Com20] and $-90dBm$ receiver sensitivity.

With the environment models created, next we add basestations to the simulation. We deploy two sets of basestations for different experiments: (1) We deploy a mmWave basestations for every LTE basestation in the whole simulation area to study the coverage achieved by co-locating mmWave basetations with LTE basestations. The LTE basestation deployment plan (2) We handpicked typical environments for urban, suburban, and highway following the guideline from 3GPP TR38.901 above 6GHz channel study [3GP17a] and deploy the basestations accordingly.

2.2.2 Bridging traffic and ray-tracing simulator

Real-life traffic is crucial in client mobility evaluation and MU-MIMO user selection analysis. We feed the OSM 2D road map into SUMO to generate the road network. Then we define the traffic follow the traffic model in [KBGP+18, 3GP18]. We specify the dimension, accelerate, and occurrence probability of different types of vehicles according to the traffic model and write it into a .xml configuration file. SUMO follows the configuration file and generates 10 minutes of continuous traffic flows with 10ms granularity. Then we use a python program to extract the vehicles' locations from SUMO's Traffic Control Interface (TraCI) [(DL19b]. The

vehicles' locations for each 10ms snapshot are stored in a .mat file for later creations of ray-tracing simulations.

Finally, we add vehicle 3D models/receivers to the simulation. Vehicle objects in Wireless Insite are stored as an ".object" file. We first copy the ".object" files to the simulation environment with numbers and types according to the SUMO generated traffic model. Then we move the coordinates of vehicles according to SUMO generated .mat file.

2.2.3 Reuse Channel Information to Emulate Beamforming

To compensate for the intrinsic limitation of short wavelength and hence high attenuation loss, mmWave radios have to communicate through directional beams, formed by many-element phased array antennas. To simulate the channel between a pair of mmWave radios, Wireless Insite treats each antenna element separately. As a result, the simulation time grows quadratically with the number of transmitter and receiver antenna elements. Even on a powerful PC with dual Nvidia 1080Ti GPU, 4.7 GHz CPU, and 128 GB SSD RAM, it takes half an hour to simulate an 10ms snapshot of the channel between one 32-element transmitter and 32-element receiver phased array in a typical urban scenario. Since our simulation often involves different large size phased array for the same environment, computation overhead can be huge. To overcome the computation barrier, we set one omni-directional antenna element with 0dB gain for each basestation and receiver for each scenario, to just capture the paths' magnitude/phase and AoA/AoD. Assuming far-field assumption, the received signal from the given transmit array can be formulated as:

$$\mathbf{y} = \sum_{i}^{P} A_i G(\phi_i, \theta_i) \tag{2.1}$$

where A_i and G_i denote the magnitude, beamforming gain, and phase of the i_{th} path and P is total number of paths. The beam pattern $G(\phi, \theta)$ is generated by the DFT method [Dud83], *i.e.* align the phases of all antenna elements for the desired direction.

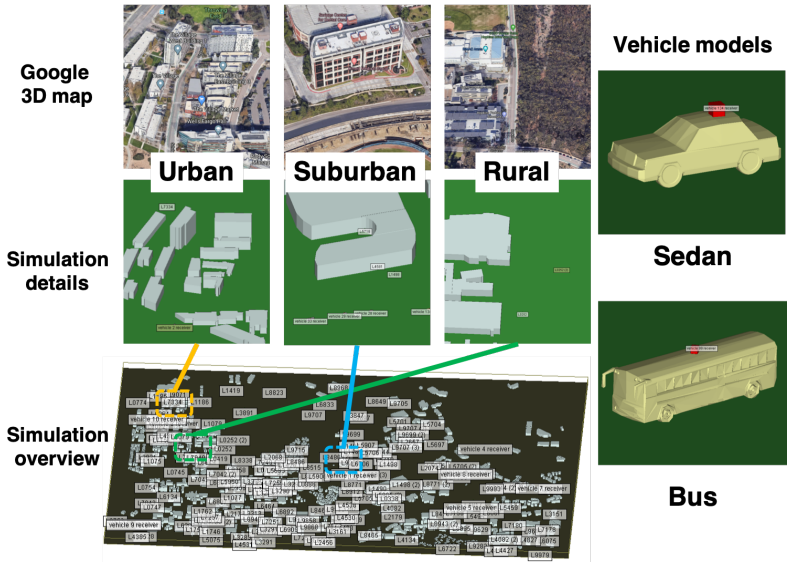

Figure 2.3: Simulation overview.

For our MU-MIMO evaluation, we compute the received signals from additional transmit array with a displacement of d from the reference array with the following equation,

$$\mathbf{y}^{'} = \sum_{i} A_i G(\phi_i, \theta_i) e^{j\left(2\pi \frac{mod(dcos\theta_i, \lambda)}{\lambda}\right)} \tag{2.2}$$

where the $e^{j[.]}$ term represents the relative difference of the current array to the reference array. In this way, we can reuse the paths' information in the same environment to emulate different phased array sizes. The emulation saved $N_{Tx} \times N_{Rx}$ times of computation overhead for every Tx-Rx pair, where N_{Tx} and N_{Rx} are the number of the phased array at TX and Rx side.

2.2.4 Validation

To verify whether the 3D ray tracing can faithfully reproduce the real mmWave channel, we feed a 3D model of the testbed deployment site into the ray tracer, and compare the simulation with real measurements. Our simulation configuration attempts to maximize the consistency between the simulated and actual radio hardware (phased array geometries, beamforming codebook, deployment location/height, output power, *etc.*).

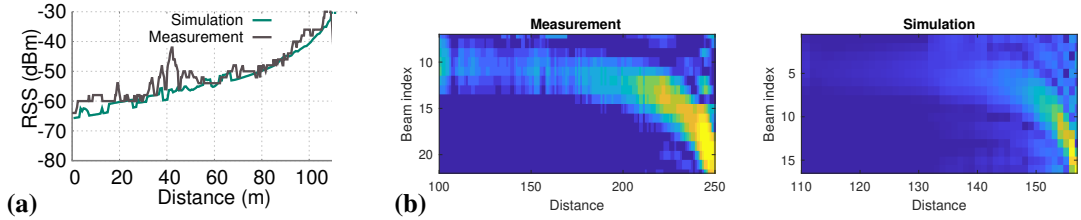

Figure 2.4: (a) Best beam RSS comparison; (b) Per beam RSS comparison.

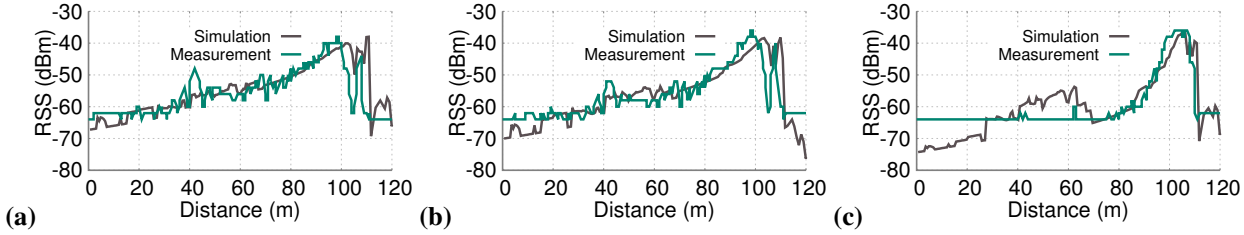

Figure 2.5: Per beam RSS showcases: (a) Beam 12; (b) Beam Beam 15; (c) Beam 19.

During the measurement, we turn on basestation RMi1 (shown in Fig. 2.1) and drive along the road segment nearby. The vehicle client logs the per-beam RSS. Fig. 2.4 compares the testbed measurement and simulation. We can see that both the RSS and the best beam index in simulation share similar patterns with the measurement result. Fig. 2.5 further shows that the RSS variation patterns in simulation match the measurement well. Therefore, the general insights derived from our simulation would also work in the real mmWave V2X.

Chapter 3

Link Throughput and Coverage

3.1 Link Level Throughput

We first evaluate the TCP throughput between a mmWave basestation and the moving vehicle. The vehicle client starts association to the basestation at 200m south of basestation *RMi2*.

We then initiate an iPerf3 TCP stream from the AP to the client while quickly accelerate to a certain speed and cruise. The traffic is captured by tcpdump on the client side. Note that the AP operates standard 802.11ad SLS beam searching mechanism, which exhaustively probes all beams in codebook every Beacon Interval (BI=100ms in our case). It represents the simplest beam searching solution and often criticized for large overhead and poor reactiveness [HAR+18].

Fig.3.1 shows 3 traces of TCP throughput and packet loss under different vehicle speeds and AP codebooks. Surprisingly, all three traces exhibit fair throughput. Even for the latter 2 traces where the vehicle is moving at 50mph, the average throughput is still over 700 Mbps. In addition, there is rarely packet loss, indicating the links are stable. The result implies that despite the seemingly challenging scenarios, mmWave V2X is achievable. In fact, at a link level, even a simple beam searching design is able to achieve a stable link with fair throughput under high mobility. This surprising simplicity of mmWave V2X comes from vehicles' predictable mobility

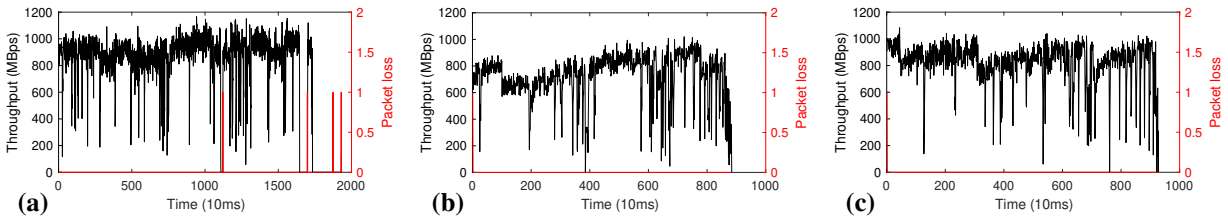

Figure 3.1: TCP performance (a) 20mph, "full coverage" codebook. (b) 50mph, "full coverage" codebook. (c) 50mph, "uneven" codebook.

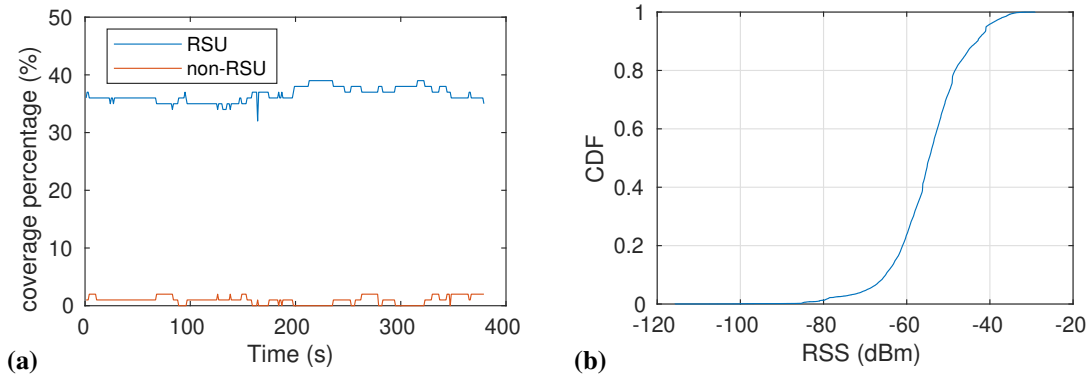

Figure 3.2: Large scale coverage showcase (a) Coverage over time, (b) RSS CDF.

and resulting AoD sparsity and persistence, which we will discuss in Sec.4.

The TCP traces also show the non-trivial impact of mobility on TCP performance. The first two traces show the performance difference of the same codebook under 20mph and 50mph. The 50mph trace averages 717.4Mbps with some sections lower than 600Mbps, comparing with 20mph trace which averages 872.0Mbps and almost always over 800Mbps. The effect of vehicle speed is due to the vehicle moving out of the current beam's coverage within two beam sweepings. Intuitively, one can avoid this effect by shortening beam sweeping intervals. However, frequent beam sweeping leads to high overhead, which reduces throughput. In Sec.5, we show how one can design the codebook to alleviate this effect without overhead.

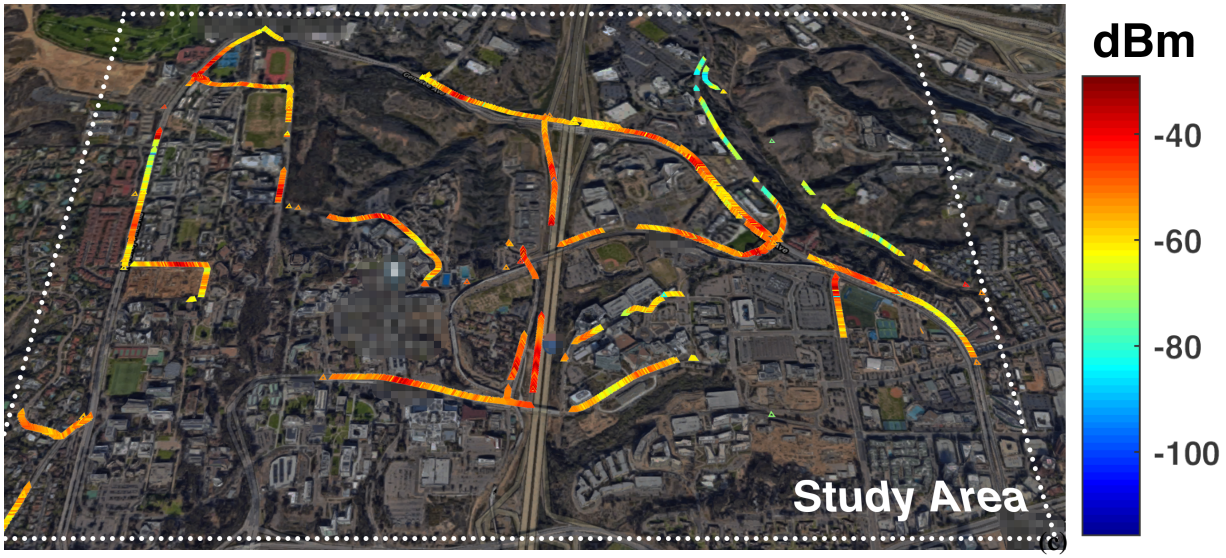

Figure 3.3: Large scale coverage heatmap.

3.2 Coverage at a Large Scale

Existing research [Qua19] has demonstrated that mmWave basestaions can achieve 75% coverage in a dense urban scenario by co-locating with the LTE basestations. But this result may not hold for all scenarios, due to the various density and the placement of LTE basestations in different environments.

To investigate the potential coverage for mmWave V2X basestation co-locating with LTE basestations, we co-locate mmWave V2X basestaions with the LTE basestations in the large scale simulation and examine the coverage. The coordinates of all 170 LTE basestations within this area are collected from *OpenCellID* [Lab20]. We assume mmWave V2X basestations have omni-coverage. Fig. 3.3 shows the coverage heatmap of vehicle clients and their RSS. The coverage is only around 37% at anytime and the average RSS is $-54dBm$, which translates to 3080Mbps. The result suggests the mmWave V2X can be poorly covered by co-locating with LTE basestations.

The discrepancy between our result and [Qua19] originates from basestastion deployment. [Qua19] select a typical urban scenario where basestations are densely deployed and near the

roads. Our large scale simulation runs in a more generic environment where basestations are evenly distributed over the field. The basestations far from road hardly contribute to the coverage. To verify this effect, we define ”roadside” type basestations as the ones within 20m of the road and not blocked by the building. We plot the CDF of percentage of the coverage achieved by the ”roadside” and ”non-roadside” basestations in Fig. 3.2 (a). The result shows that although comprising only 39% of all basestations, the ”roadside” type basestations contribute 97.43% of the overall coverage. This implies the ”roadside” basestations contribute almost all coverage in mmWave V2X.

Since co-locating with LTE basestation only provides 37% of the coverage, theoretically 2 times extra of the ”roadside” basestaions need to be deployed to get mmWave V2X full coverage. That means that we need to deploy 20 more ”roadside” basestations per km^2. Such numerous basestations can be hard to deploy in suburban and rural scenarios due to the facility limitation. This motivates the lightweight and self-sufficient basestation deployment.

Chapter 4

Mobility

In this section, we investigate the impact of vehicle mobility on the spatial profile of the mmWave V2X channel, along with the implications for beam management. The vehicle mobility changes the propagation paths' Angle-of-Departure/Angle-of-Arrival (AoD/AoA), while an optimal beamforming scheme should direct the signal power towards the dominant AoD/AoA. Therefore, our measurement mainly focuses on the AoD/AoA dynamics. For clarity, we first investigate the AoD dynamics from the basestation's perspectives and assume omnidirectional Rx beamforming. Later in Sec. 4.3 we show the similarity of AoA and AoD and how they can be decoupled to reduce beamforming complexity.

4.1 AoD Sparsity Across Geographical Locations

It is well known that the mmWave channel is sparse [SVZR15, NCF+14]. Between a given pair of Tx and Rx, there usually only exist a small set of dominant propagation paths and AoAs/AoDs, created by LoS and a few strong reflectors. However, for mobile Tx or Rx, the dominant AoAs/AoDs keep changing with the Tx/Rx locations, and are generally considered to be densely distributed across a wide angular range [PDDW17, ZWX+18]. As a result, practical phased arrays need to generate evenly distributed beams, either quasi-omni or directional, to

cover the entire field of view [IEE12, IEE18]. The key challenge of mmWave management essentially lies in efficiently selecting the best beam(s) for arbitrarily located Tx/Rx. Prior research generally demoted the exhaustive beam scanning method, and instead attempted to reduce the beam search overhead through efficient beam searching or specialized hardware capabilities [GPMRH16, HAR+18, VCS+17], which inevitably adds to system complexity. However, we find the mmWave V2X channel exhibits sparsity even across locations. With the right codebook design, even the simple exhaustive beam scanning can efficiently identify the best beams. Fig. 4.2 shows the histogram of best beams' index, when one basestation is serving one vehicle moving from 130m away. The basestation uses a 39-beam codebook, which uniformly covers the $70^\circ - 90^\circ$ azimuth/ $40^\circ - 140^\circ$ elevation angles. We see that only a small subset of beams between index 30 to 35 are repeatedly selected and most are unused.

We further examine the underlying channel characteristics leading to this phenomenon. Our testbed cannot provide accurate AoA/AoD since AoD/AoD recovery algorithms require CSI as input [GMD04]. Hence we examine the propagation paths from simulated scenarios. The scatter plot in Fig. 4.1 shows the AoD of all paths (including reflection paths) from basestations in a typical simulated urban and highway scenarios. The path data is accumulated from all the vehicle clients (around 40) passing by within 10 minutes. Each dot in this figure represents a path from a certain vehicle to the basstation, with its azimuth AoD shown on x axis and elevation AoD on y axis. Clearly, the AoD scatters are clustered into sparse "curves" in the figure. Each curve corresponds to the vehicles on a road lane. Unlike human activities, the vehicle movement is confined within fixed routes, and hence the AoDs are limited to only the ones pointing at the lanes, and a few ones outside of the curves due to reflections. Note that the highway basestation has two set of curves distributed at $-90^\circ - 90^\circ$ and $90^\circ - -90^\circ$. This is because the highway basestation is placed in between two roads going different directions. The similar AoD distributions could also happen to the basestations with many different roads in front. In this case, we can simply split the basestations into many sectors with each one serving a road, since the AoD distribution

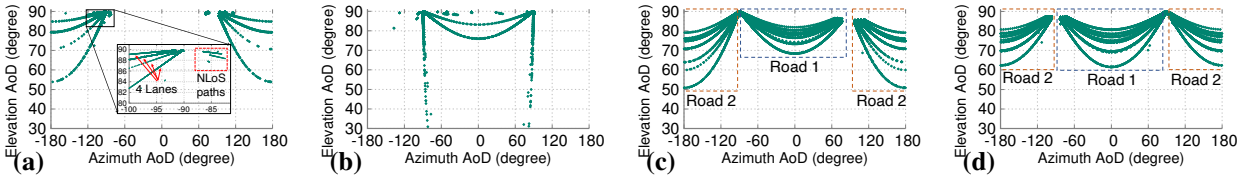

Figure 4.1: AoD distribution showcases (a) Urban basestation 1. (b) Urban basestation 2. (c) Highway basestation 1. (c) Highway basestation 2.

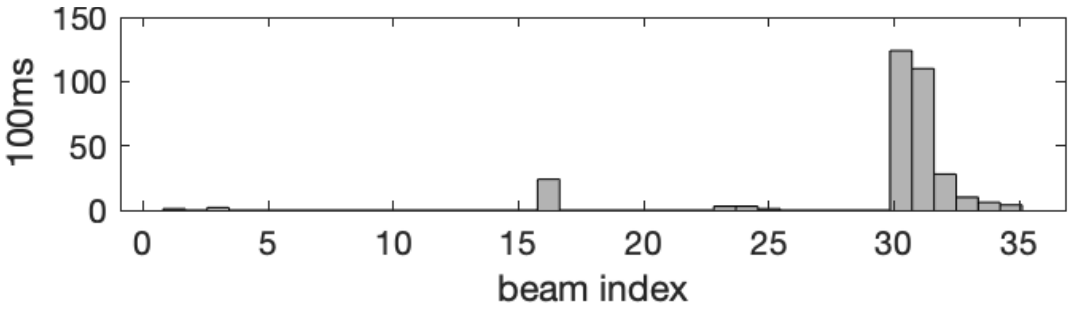

Figure 4.2: Best beam histogram.

on one road is sparse.

Because of this strong correlation between signal propagation angles and the road geometries, one can prune the beamforming codebook by selecting only those candidate beams within the AoD range. Consequently, the beam searching overhead can be substantially reduced. Table. 4.1 compares the overhead of the pruned and full codebook under different V2X scenarios and phased array sizes. We assume omnidirectional client adopts 5G NR standard to calculate beam searching overhead, which takes 17.84μs to scan 1 beam and has beam sweeping periodicity of $T_s = 20ms$. Since 5G NR limits the number of beams searched within T_s interval to 64 beams, the full coverage codebook cannot even finish a full scan in one T_s. The pruned codebooks, on the other hand, all fit within one T_s with overhead $\leq 1ms$, which translate to only $1ms/20ms = 5\%$ overhead under default $T_s = 20ms$ in 5G NR [3GP17b]. Due to more driving lanes, the highway has the most beams but still has 1.1ms low overhead. The suburban reduces to the smallest codebook since there are not too many lanes nor many reflections.

Table 4.1: Beam searching overhead.

	Full coverage codebook		Fit to AoD range	
Setting	# beams	overhead	# beams	overhead
Urban 6x6	133	1.1ms×3	33	0.6ms
Urban 12x24	481	1.1ms×8	45	0.8ms
Suburban 6x6	133	1.1ms×3	28	0.5ms
Suburban 12x24	481	1.1ms×8	39	0.7ms
Highway 6x6	133	1.1ms×3	47	0.8ms
Highway 12x24	481	1.1ms×8	61	1.1ms

4.2 AoD Persistence

It is generally believed that high mobility necessitates frequent beam searching, which leads to huge overhead [VSBHJ16, GPMRH16, NCF+14]. This is often considered the most significant challenge for mmWave V2X. However, contrary to this intuition, we found the AoDs' changing rates are nearly static for most of the regions, and beam searching only needs to be called occasionally. Fig. 4.3 shows the best beam index time series from the aforementioned experiment. When the vehicle is more than 20m away from the AP, the best beam is either 30 or 31. Beam 30 and 31 in our codebook point at 70° and 80° respectively, implying the AoD changes in this 100m area are less than 10°. Only when the vehicle is very close to the basestation (within 20m), the AoD starts to change more quickly over time.

The AoD persistence in far regions originates from geometric relations between the vehicle client and basestation. The AoD's changing rate ω_v, *i.e.*, vehicle's angular velocity relative to the basestation, are a fraction of vehicle's linear velocity v_v by link distance d.

$$\omega_v = \frac{v_v}{d} \tag{4.1}$$

In a typical roadside basestation setup, the link distance as denominator is often large (*e.g.* 20-100m) compared to the vehicle's linear velocity (*e.g.*, 22m/s 50m/s) for far regions. As a

Table 4.2: RSS loss (comparing with $T_s = 5ms$) in the near regions.

	6x6 Array		12x24 Array	
T_s	40mph	75mph	40mph	75mph
20ms	-0.26dB	-0.56dB	-2.46dB	**-6.34dB**
40ms	-0.64dB	-2.36dB	**-7.88dB**	**-7.21dB**
80ms	-2.70dB	**-12.11dB**	**-8.28dB**	**-12.61dB**
160ms	**-14.85dB**	**-12.75dB**	**-10.42dB**	**-16.3dB**

result, the AoD changes very slowly in the far regions. Fig. 4.4 (a) shows the scatter plot of AoD change rate versus the vehicle's horizontal distance to the basestation in simulated scenarios. Apart from a few NLoS paths in the urban scenario, most of AoD changes over $20^\circ/s$ happen when vehicle is within 20m of AP. The highway basestation has less large AoD changes due to the larger link distance induced by higher basestation hanging position.

Since the far regions dominate a vehicle's trajectory, an infrequent beam searching would not sacrifice link performance for most of time. Fig.4.4 (b) shows the histogram of AoD changes of all traces from the urban and highway scenarios. We mark the AoD change thresholds to cause 1dB and 3dB loss with $T_s = 100ms$ standard beam searching interval in 802.11ad and $T_s = 160ms$ maximum allowable beam searching interval in 5G NR [3GP17b]. The two beam searching intervals cause $1.1ms/100ms = 1.10\%$ and $1.1ms/160ms = 0.69\%$ overhead and even with such negligible overhead, exhaustive beam searching can still catch most of AoD changes without more than $1dB$ of loss.

4.3 Receiver beamforming

It is reasonable to expect receivers with beamforming capability are widely used in V2X for range extension and interference avoidance (Chapter9). However, Rx beamforming quadratically increases the beam searching overhead. For each Rx beam, the entire Tx beams need to be searched, *i.e.* $O(mn)$ time complexity where n is Tx beam numbers and m is Rx beam

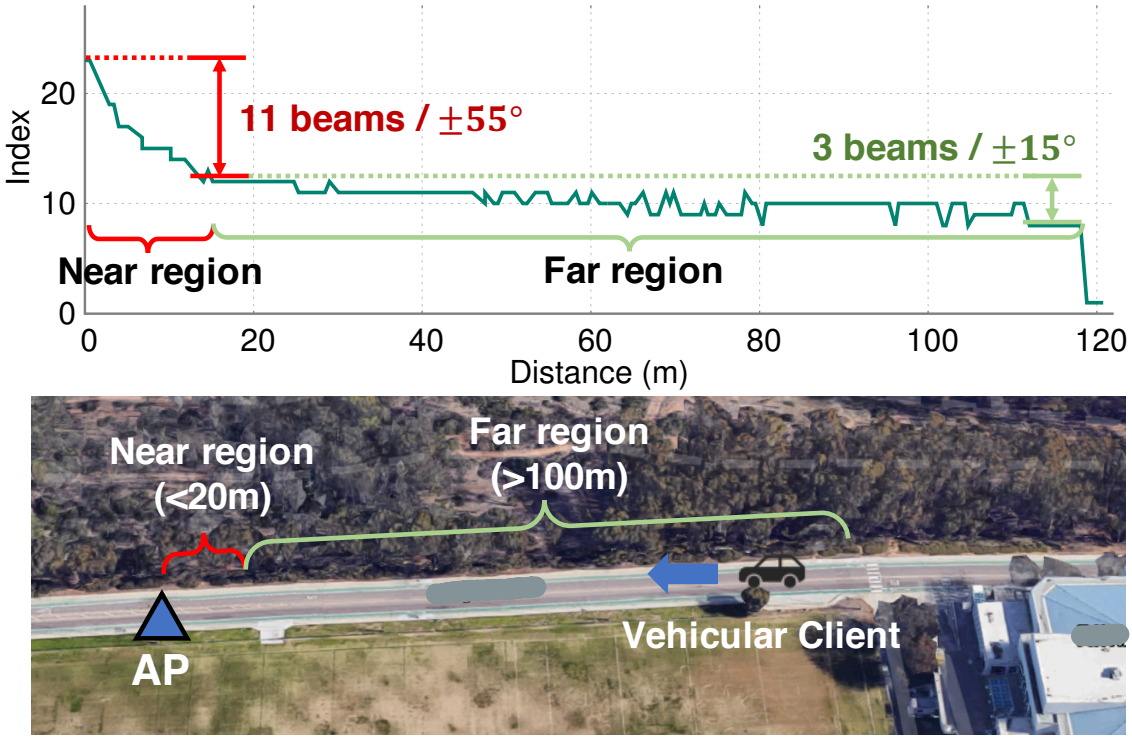

Figure 4.3: Best beam dynamic in near/far regions.

numbers. This means for an urban 6×6 array basestation with 0.6ms overhead (Table. 4.1), even a 4-beam Rx codebook will bring the total overhead to 2.4ms and exceed the limit of one beam search period. It is unclear if the extra overhead of Rx beamforming can justify its benefits. In the existing 802.11ad/ay mmWave system, it is a common practice to alleviate beam sweeping overhead by *single sided Tx/Rx beam searching*, *i.e.* assuming channel reciprocity, the Tx beams scan its beams with Rx in omnidirectional mode (TXSS) and vice versa[IEE12]. The beam sweeping overhead reduces from $O(mn)$ to $O(n)$. For outdoor mmWave networks especially the far region clients, this method is not applicable since the omnidirectional Rx significantly limits the range and as a result, the clients may not be able to receive the basestations beam searching signals.

For mmWave V2X, it turns out we can leverage AoA's properties to perform similar single-sided Tx beam searching. In Sec. 4, we show a basestation can serve far region clients

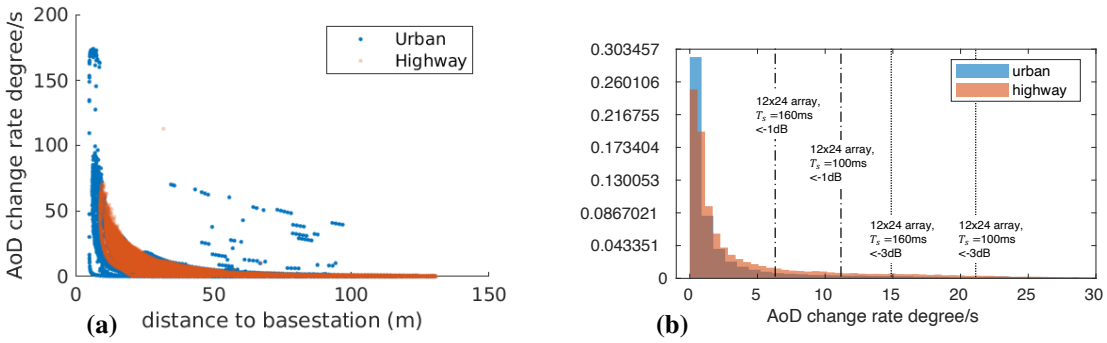

Figure 4.4: (a) AoD changes v.s. distance to AP; (b) AoD changes histogram.

with just a few beams with less frequent beam searching. Since AoD and AoA is complementary for LoS paths, the AoA inherits the same sparsity and persistence, implying the far region clients should also be able to reach the basestation with a few beams given LoS present. To verify AoA's sparsity, we utilize the same simulated highway setting from Fig. 4.1 and examine the distribution of AoAs of links with low RSS ($< -90dBm$, Rx sensitivity of our mmWave radio [air19]) under omni-directional Rx. We only consider highway scenarios since in urban/suburban scenarios, the inter basestation distance for 5G deployment is typically small ($\sim$100m [3GP17a]). As a result, even cell edge users can still get $-49dBm$ RSS. In such cases, omnidirectional Rx can easily pick up the basestation's beam searching signal. Fig. 4.5 shows the scatter plot AoA distributions. The AoA scatters are clustered and only spans 6° on azimuth and 1° on elevation. Recall that even a large 6×6 phased array forms beams with 19°. It implies that a far region client can easily use one beam to receive the Tx beam searching signal. Note that in reality, the outdoor receiver may have even higher sensitivity than our indoor radios, hence the client can use and the AoA distribution is more confined.

To verify the effectiveness of using single-sided Tx beam searching on far region clients, we reuse the same simulated highway scenarios in previous section of this chapter and apply single-sided Tx beam searching to the vehicle clients. Without loss of generality, we assume the client's Rx beam points towards the moving direction. Table. 4.3 shows the percentage of links that can be received by using single-sided Tx searching. We see even for a large 12×12 Rx

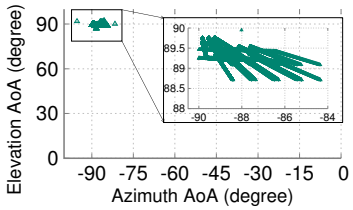

Figure 4.5: AoA distribution of far region clients in highway scenarios.

Table 4.3: % of links covered by single-sided Tx beamforming.

Rx Array Size	4x4	6x6	12x12
% of links	100%	100%	98.90%

phased array, there are 98.9% of links covered by one beam. The result shows the effectiveness of single-sided Tx beam searching.

4.4 Beam Searching Problem in Near basestation regions

While a smaller codebook can achieve optimal performance due to AoD sparsity and persistence, the beam selection in the near regions is still non-trivial. Due to its proximity to AP, the AoD changing rates in these areas are typically high, requiring short beam scanning interval. As a result, even a smaller codebook still needs to make a trade-off between beam sweeping overhead and RSS loss. To demonstrate this effect, we select one typical urban AP and alters T_s along with the vehicle speed. We then examine all paths of clients in the near AP region, *i.e.* link distance less than 20m. Table.4.2 shows the mean RSS loss of $T_s = \{10ms, 20ms, 40ms, 80ms, 160ms\}$, as suggested by 5G NR [3GP17a], comparing to minimal $T_s = 5ms$. We see over $-6dB$ of loss for $T_s > 40ms$ on 12x24 array and over $-3dB$ loss for $T_s > 80ms$ on smaller 6x6 array. This implies for vehicle with high mobility, it either chooses the minimal $T_s = 5ms$ beam searching interval and loses up to $1.1ms/5ms = 22\%$ to overhead, or chooses a larger T_s and suffers up to $-16.3dB$ RSS loss. We show in Chapter. 5 how to avoid overhead without major performance loss via proper codebook design.

Chapter 5

Codebook/Beam management

An intuitive way of reducing beam sweeping overhead is to "widen" the beams pointing to the areas near the basestation. "Widening" beams usually means reduce the maximal gain. However, the wide beam covers larger angular range and hence requires less frequent beam sweeping. For example, a $27.5dBi$ beam generated by DFT method has a main-lobe width of $2.4°$; a $20.7dBi$ beam has a main-lobe width of $11.4°$. For a vehicle client 10m away from a basestation moving at 75mph, it requires $T_s = 6.2ms$ for $27.5dBi$ beams to keep track of the client and $T_s = 29.7ms$ for $20.7dBi$ beams. The wide beam reduces the overhead from 16% to 3%. However, the higher gain of narrow beams does not translate to throughput. In near regions, links usually have a large SNR margin. So a wide beam may lead to the same link bit-rate despite the lower directional gain. Assuming 82dBm EIRP, the 2 beams lead to $-30.4dBm$ and $-37.2dBm$ of RSS respectively. According to the RSS to bitrate mapping in 802.11ad [IEE12, IEE18], both of these two beams lead to 8.09Gbps (MCS 12). Therefore, using an "uneven" codebook with high directionality beams for further away area and wide beams for a nearby area can potentially reduce the beam management overhead without sacrificing link quality.

To verify the advantage of the uneven codebook, we create two codebooks, one with 39 beams evenly distributed in $30-150°$ azimuth and $110-90°$ elevation angles. Fig. 5.1 (a) and

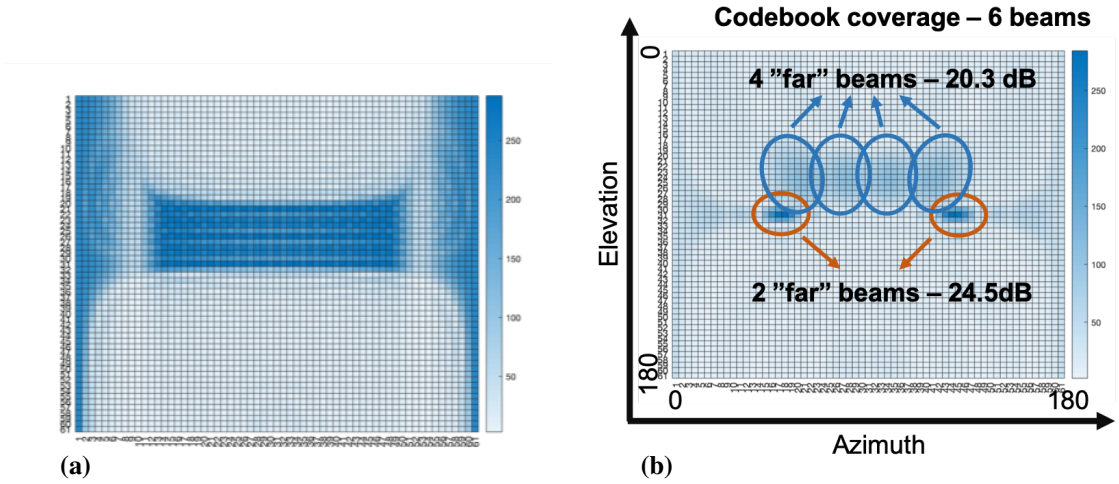

Figure 5.1: (a) Full coverage codebook; (b) Uneven codebook.

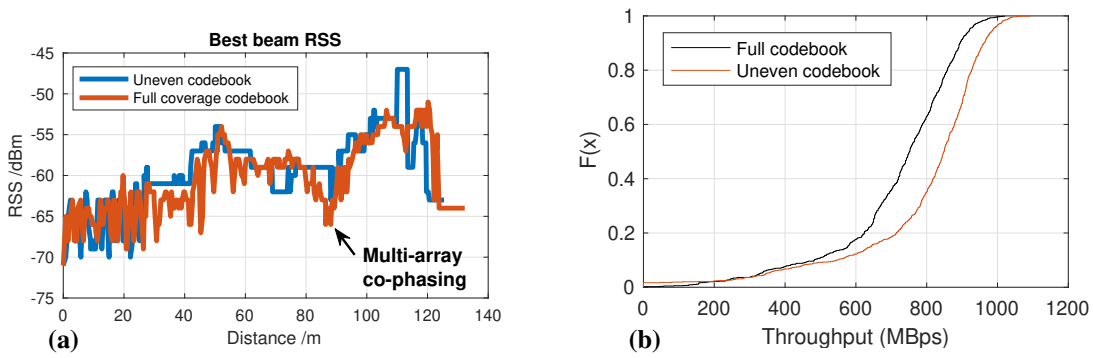

Figure 5.2: (a) RSS traces of two codebooks; (b) TCP throughput of two codebooks under 50mph.

(b) shows the simulated beam pattern for even and uneven codebooks, respectively. Fig.5.2 (a) shows the RSS trace, and Fig. 5.2 (b) shows the throughput CDF of uneven and full coverage codebook. We can clearly see that the two codebooks exhibit nearly identical performance, but the full coverage codebook suffers co-phasing effect due to the displacement of multiple arrays. This will be mentioned in Chapter 7. We load the two code into one of our deployed basestation, run multiple trials and get the CDF of throughput in Fig. 5.2 (b), where we can see the uneven codebook can achieve around 100 Mbps higher throughput than the full coverage codebook.

Need for site-specific beamforming codebook. The AoD sparsity and uneven codebook

greatly reduce codebook/beam management complexity. However, since the spatial channel profile depends on the basestation deployment and road geometries, the codebook needs to be tailored in a site-specific way, *i.e.*, which direction should the far/near beam pointing to achieve optimally coverage, how wide should the near beam be to accommodate vehicles' speed. Since the AoD shows a similar curve pattern across all roadside basestations, we believe the environment adaption can be achieved with simple mechanisms.

Chapter 6

Blockage

One major concern for mmWave V2X lies in the link outages caused by roadway obstacles including foliage, pedestrian, and tall vehicles [VSBHJ16]. To investigate the severity of the blockage in mmWave V2X, we examine the blockage occurrence probability in different scenarios for 10 minutes. We identify the blockage by the absence of the LoS path. We categorize blockage into two types: recoverable blockage where blockage can be saved by NLoS, and fatal blockage where there is no viable path. Fig. 6.1 (a) shows the average blockage probability in different environments in large scale simulations. The result indicates that *overall* blockage is not severe, the severest urban scenario only has around 8.5% blocked probability. Without densely building

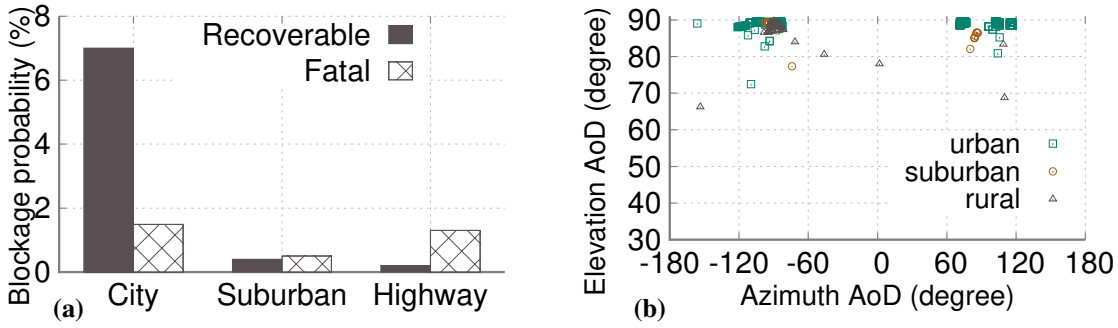

Figure 6.1: (a) Recoverable blockage v.s. fatal blockage; (b) Recoverable path angular distribution.

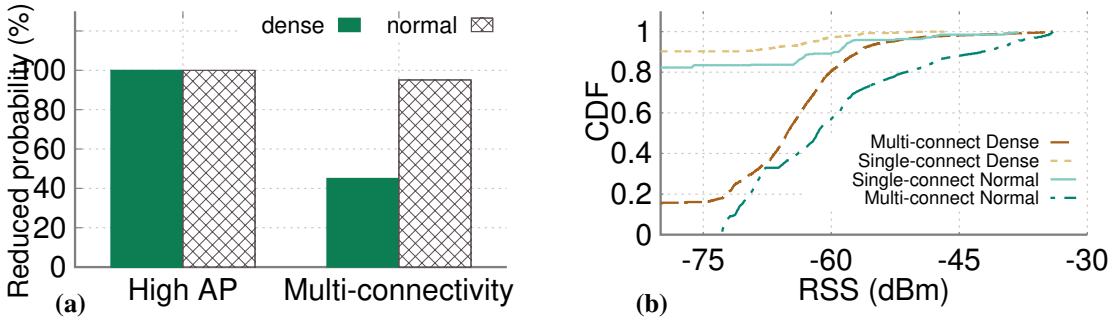

Figure 6.2: (a) Blockage saved by high-basestation setup and multi-connectivity in highway; (b) RSS distribution under multi-connectivity in highway.

and vehicle obstacle, the overall blockage possibility of suburban and highway is only 1% and 2%. The low blockage probability happens because the basestation is hanged 3-4m higher than the rooftop of the vehicles, and most of the obstacles are not high enough to block the LoS. Note that even for the urban scenario, 86% of the blockage can be recovered by strong NLoS due to the reflections provided by numerous surrounding buildings.

Although blockage happens infrequently, it may cause temporary blackout, unless a strong NLOS reflection path is available for recovery. In Fig. 6.1 (a), although the total blockage probability are low as mentioned before, the blockage are mostly fatal blockage. This implies that the most of the blockages leads to complete link outages, which may have severe impact to upper layer performance.

For those recoverable blockage cases, the mmWave link needs to quickly identify alternative NLoS reflection paths. Due to the sparsity of overall AoD distribution at basestation side, we expect the same sparsity exist in the AoD distribution of the strongest NLoS path. We collect the azimuth AoD and elevation AoD of the best path under recoverable blockages from the previous experiments. Fig. 6.1 (b) plots the AoD distribution for all scenarios. We see the AoDs are clustered around -100° and 100° in azimuth, and around 92° in elevation in all scenarios. This shows the sparsity of the recovery paths' AoD. Similar to the AoD sparsity induced small codebook size in Sec. 4, this result implies that while recoverable blockage happens, the base

station can save recoverable blockage by using a small set of beams covering these angles.

To reduce fatal blockages, one simple solution is to directly increase the basestation height. In this way, the LoS path bears a larger elevation angle relative to the ground, and hence cannot be easily blocked. The first group of bars in Fig. 6.2 (a) shows the effectiveness of a high basestation. It is observed that over 99% of the fatal blockages are eliminated when the basestation height increases from 5m to 15m.

When improving basestation height is infeasible (*e.g.*, due to cost or tower limit), an alternative solution is *Multi-Connectivity* (MC)–a 5G feature that allows client devices to connect to multipole basestations simultaneously to improve robustness [PGM+17].

To verify the impact of MC, we implement MC in the simulations of different scenarios and see how much blockage can be saved by MC. Since the mechanism of MC between mmWave basestations and users has not been finalized, we intuitively allow vehicles to connect with the nearest available basestaion in different scenarios and run the simulation for 10 minutes. Intuitively, traffic density also plays a role in the effectiveness of MC. Thus we categorize the traffic into two types: dense and normal. Normal traffic uses the same traffic model in Sec. 2.2, and dense traffic set the vehicle density to be $2\times$ over the traffic model in normal traffic. The right group of the bars in Fig. 6.2 (a) shows the percentage of blockages saved by MC. It is shown that MC can save 40% of the fatal blockage for dense traffic and 95% for normal traffic. MC works better under normal traffic because MC links have lower chance to be blocked by the vehicles obstacles in normal traffic than in dense traffic. The result shows that MC can significantly reduce the fatal blockage in normal traffic, but can only reduce near half of the fatal blockage in dense traffic.

Apart from its effectiveness in reducing blockage probability in normal traffic, MC also tends to provide higher quality links under blockage. Fig. 6.2 (b) shows the RSS distribution comparison of the **blocked** link under single and multiple connectivity. With MC, over 50% of the links reach an RSS of over -62 dBm under blockage in light traffic and -65 dBm in dense traffic,

while over 80% of the links are totally blocked in both light and dense traffic in single-connectivity case.

Many studies mentioned that multi-connectivity in mmWave faces great challenges in mmWave multi-connectivity. Due to the directional beam in mmWave, multiple links has to be discovered and maintained by beam scanning, which will cause huge overhead [GMRZ16, PGM+17]. But previous sections (Sec. 4 and Sec. 6 show that with a small codebook covering major AoD directions, even the naive beam scanning mechanism induce less than 5% overhead.

Chapter 7

Multi-Array Basestations

To improve basestations' coverage and range, 5G NR adopts "Array-of-Phased-Arrays" (APA) structure [WHZ+20]. An APA basestation use a series of phased arrays to form a large planar or circular array for gain or coverage improvement. Recent work [WHZ+20] points out such APA structure may suffer from "co-phasing" effect across phased arrays, *i.e.* the signals combine incoherently because of AoA/AoD and cause RSS loss. We find such "co-phasing" effect also happens for V2X. We select one of our basestation and first turn on all phased arrays can collect RSS trace. Each phased array is loaded with the same codebook. Then we turn on only one phased array and repeat. We multiply the RSS value from the one array trace by 8 ($+9.03dB$) to emulate the coherent combined RSS for all 8 phased arrays. Fig. 7.1 (a) shows the strongest beams' RSS comparison between the two traces. We see in the multi-array trace, at $87m$, 115, and $120m$, the co-phasing effect manifests in the form of "notches" (marked red), leading to up to $8dB$ loss compare to the coherent trace.

The co-phasing effect creates significant fluctuations in RSS trace and is even more harmful for upper-layer applications. However, we find the co-phasing effect does not affect beam selection. Fig. 7.1 (b) shows the best beam index dynamic for the above two traces. Two best beam index dynamics are near identical across two traces. This is because the co-phasing

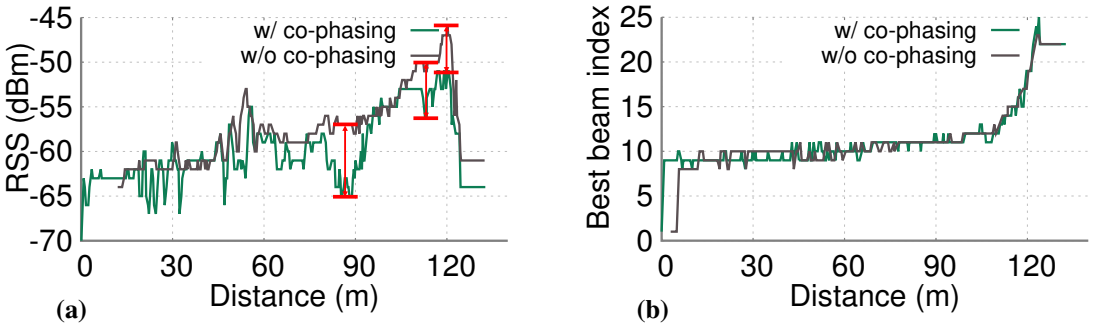

Figure 7.1: (a) Co-phasing effect showcase; (b) Best beam index dynamic w/ and w/o co-phasing.

creates the same RSS loss across all beams. Thus the best beam after co-phasing is still the best beam. It implies that the co-phasing problem and beam selection can be addressed separately.

Chapter 8

Is mmWave Hybrid Beamforming Necessary?

MmWave MU-MIMO is a major feature in 5G NR and the upcoming 802.11ay [IEE18]. It relies on multiple RF chains and phased arrays to simultaneously send multiple data streams to different users. Since the phased arrays' beams may partially overlap, a hybrid beamforming algorithm is generally used to cancel the inter-user interference.

$$y = H^{M\times N} W^{N\times B} P^{B\times B} \bar{x} \tag{8.1}$$

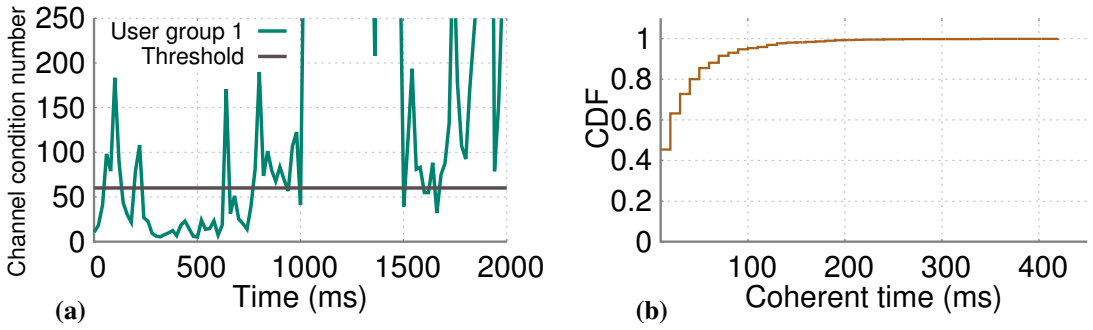

Figure 8.1: (a) Channel conditional number dynamic; (b) CDF of channel orthogonality coherent time.

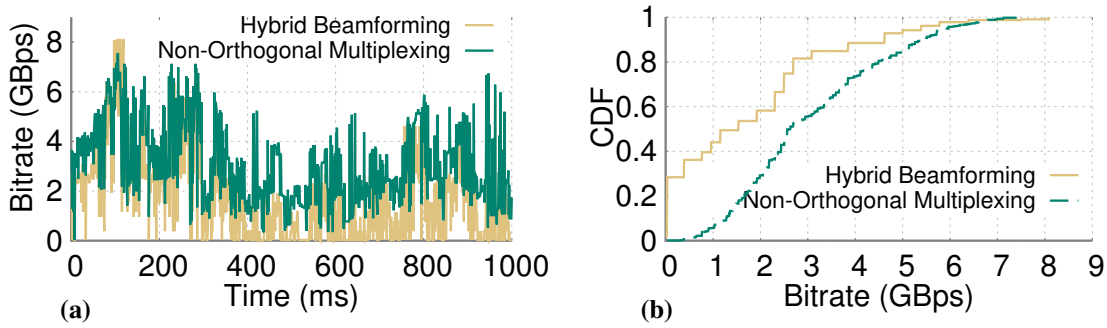

Figure 8.2: (a) Bitrate variation in 10s for hybrid beamforming/non-orthogonal multiplexing; (b) Bitrate CDF comparison.

where H is the channel of selected users, W and P is the beamforming weight matrix and precoding matrix and $\bar{x}$ is the desired signals targeting different users.

The feasibility and effectiveness of hybrid beamforming strongly depends on whether the channels between the basestation and different users are orthogonal. Poor orthogonality may make the precoding matrix P unsolvable. Orthogonality can be reflected by the *channel condition number* [TV05], which should be close to 1 in order for all the users to be uncorrelated. To showcase the dynamics of the orthogonality in V2X scenarios, we conduct simulation in an urban environment. We select the vehicle groups with the strongest orthogonality at the first snapshot and track its channel conditional number change along time. In Fig. 8.1 (a). We can see the channel condition number fluctuates very frequently. This indicates the good orthogonality cannot be kept for a significant duration.

We further check the coherence time of the orthogonality, *i.e.*, duration within which the channel condition number of a selected set of vehicles remains below a certain threshold. Here we define the threshold as channel condition number of 60, which achievable bitrate is $\frac{1}{4}$ of the maximum bitrate. The CDF result in Fig. 8.1 (a) shows that the selected vehicle group cannot keep orthogonal for even one beacon interval (100 ms) over 95% of the time. Therefore, to ensure the effectiveness of mmWave MU-MIMO, a proper orthogonal user grouping algorithm needs to

be invoked frequently. This may result in substantial overhead, since evaluating the orthogonality requires all the users to feedback their CSI to the basestation. An alternative, simpler way of harnessing the multiplexing gain is to directly steer different phased arrays' beams to different users–a method often considered inferior to MU-MIMO due to concerns of inter-user interference. Nonetheless, as long as the users have large angular separation, mutual interference may become minimal. We refer to the approach as non-orthogonal multiplexing. One obvious advantage of this approach lies its low feedback overhead. Only a single value, *i.e.*, AoD, needs to be fed back from each user to the basestation, and the AoD can be estimated by simple using the beacon broadcast [PSL$^+$18, WHZ$^+$20]. In contrast, the MU-MIMO requires CSI feedback, whose overhead comprises a multiplication of the number of users, subcarriers, and beams. To verify the effectiveness of this approach, we conduct a simulation experiment in the urban environment with dense traffic. For MU-MIMO, we select the vehicles with the strongest orthogonality of every timestamp and calculate their average throughput. Fig. 8.2 (a) shows the throughput of a random subset of vehicles, averaged over 100s. Surprisingly, the simple non-orthogonal multiplexing approach achieves even higher performance than MU-MIMO. Fig. 8.2 (b) further plots the CDF of average throughput comparison under a 10-minute experiment in the urban environment with dense traffic. The result demonstrates a 1.3 Gbps throughput gap on average between these two solutions. The results verify the absolute advantage of the non-orthogonal multiplexing, even when the savings in feedback overhead are not counted.

In practice, the road traffic density and the phased array size may affect the opportunity of finding users with sufficient angular separation. To investigate the opportunity, we conduct a 10-minute simulation in city and highway scenarios and vary the phased array sizes. Fig. 8.3 shows the percentage of vehicles which can be included in a group with an angle separation larger than the beam width, which we denote as *candidate group*. The result shows that the percentage of candidate vehicles is heavily affected by the scenario and the phased array size. In highway scenario, there are always around 35% of vehicles which are able to form a candidate

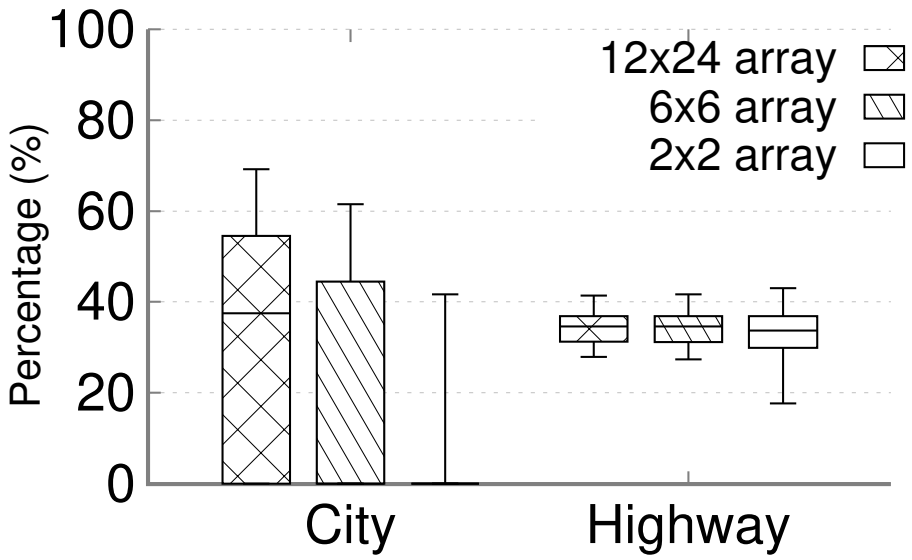

Figure 8.3: Multiplexing on phased array opportunity

group regardless of the phased array size of the basestations. However, vehicles in city scenario requires basestation to use large phased array to form candidate groups. Especially, there is no vehicle can form a candidate group in over 60% of the time using 6×6 arrays and 90% using 2×2 arrays. Thus the opportunity of forming candidate vehicle group in city scenario is more sensitive to the phased array size. Because vehicles in city scenario are located denser, larger phased arrays is needed to form narrower beam and leverage spatial multiplexity. This implies that the simple non-orthogonal multiplexing approach has a higher chance to be used in highway scenario. Although hybrid beamforming demonstrates a disadvantage compared the non-orthogonal multiplexing approach, we still need it implementing MU-MIMO in city scenario due to lack of the opportunity of the simple approach.

Chapter 9

Inter-cell Interference and Spatial Reuse

While the high directionality of mmWave radios reduces interference footprint, the dense deployment of V2X cells may still exacerbate inter-cell interference. To characterize the inter-cell spatial reuse, we select a basestation in the simulated urban scenario, and investigate all links within its cell edges. Without loss of generality, we focus on downlink only. For each link, the basestation selects the best beam, whereas the SINR is calculated assuming the strongest interfering beams are used from the two adjacent basestations. Fig. 9.1 (a) shows the scatter plot of link SINR and vehicle's distance to the basestation. The interference intensifies as the distance increases. Some of the cell edge (long-distance) links have SINR below 0dB, rendering the link unusable. This is because the distances to the desired/interfering basestations are similar. Meanwhile, the receiver, in this case a quasi-omni one, can receive from both ends. Due to the dense AP deployment, even for cell center links (*i.e.* distance is 10m) the average SINR is merely 33dB, which translates to a bitrate 5005MBps, over 3000MBps throughput loss comparing to the interference-free case.

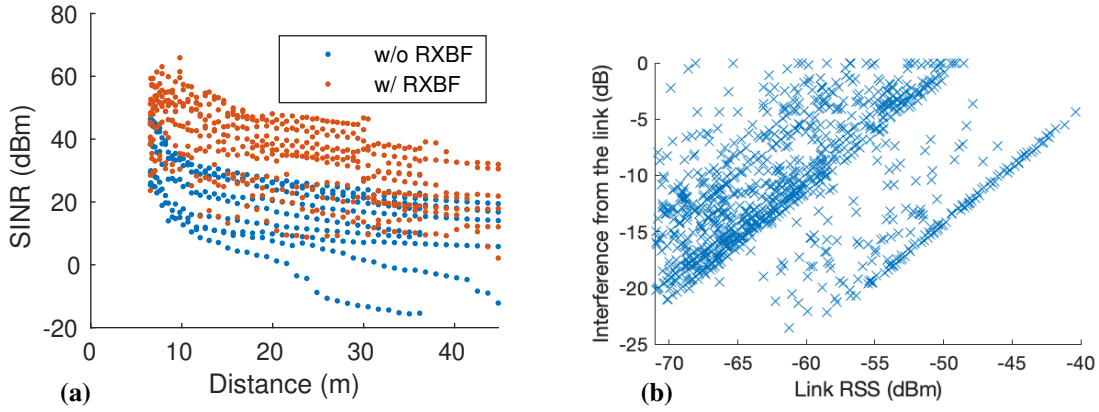

Figure 9.1: (a) Interference SINR scatter; (b) Beam used freq v.s. how much interference it causes on average.

9.1 Interference Avoidance through Receiver Beamforming

An intuitive way to avoid interference is using Rx beams to amplify desired basestation's signal while suppressing others. The orange dots in Fig.9.1 (a) shows the link SINR/distance with Rx beamforming enabled. The Rx beamforming effectively reduces the interference, improving SINR by 9dB to 20dB for cell center and cell edge links, respectively. The worst link on the cell edge is now 10dB. However, even with Rx beamforming, the average link SINR is still at a mediocre level of 15dB. This is because (i) the sidelobes generated by imperfect beamforming may still pick up the interference signal from undesired directions, and (ii) The Rx beamforming can only separate interfering signals with a large angle separation.

9.2 Interference Avoidance through Transmit Beamforming

Interference can also be avoided by steering the interferer's beam away from the receiver. However, the interfering basestation itself needs to ensure the desired user is served with a high-quality beam as well. To investigate the trade-off between serving desired users and avoiding interference, we examine the links from above urban basestation. For each link, we exhaustively apply every Tx beam to this link and record the link RSS and the strongest interference this Tx

beam causes on other basestations' clients. Fig. 9.1 (b) shows the scatter plot of a link's RSS versus the strongest interference. In general, the interference increases as the link RSS increases. We mark the area with link RSS over $-50dBm$ and inference level less than $-15dB$ in the figure. The links falling in this area have large link RSS and low interference and are good candidates for interference avoidance. We see there are few links in this area, indicating there interference avoidance opportunities by transmitting beamforming is rare.

Overall, the spatial reuse efficiency can be significantly improved by Rx beamforming. However, the inter-cell interference cannot be fully eliminated and there are non-trivial improvement margins for future studies.

Chapter 10

Related Work

10.1 mmWave V2X simulation and measurement

Recent research used simulation extensively to examine the underlying mmWave V2X channel, and reveal the unique challenges imposed by a mix of high mobility and directionality. Anjinappa and He *et al.* [AG18a, HWG+19] study the temporal and angular characteristics of sub-6 GHz and mmWave bands along an urban trajectory, and provided guidelines for effective beam tracking. Both of them only consider link-level performance, while our work aimS specifically at the design and validation of mmWave V2X. Tassi *et al.* [TEPN17] characterizes the downlink performance of a mmWave network along a highway, and analyzed the link outage probability and coverage probability amid truck blockage. Abdulla *et al.* [AW17] indicates that the standard SINR-based average performance metrics is not sufficient for characterizing V2V communications since it may obscure the performance for specific node configuration (referred as *fine-grained reliability*), and characterized the probability of packet delivery for suburban and urban intersections based on the fine-grained reliability. Unlike our work, each of the above studies provides detailed investigation only in the simulation without any measurement validation. So it is not sure whether all of their results hold for the real life scenario.

Due to the lack of a low-cost programmable testbed and the difficulty of large scale measurement for mmWave V2X, real measurement studies are quite limited and mostly focus on a single link. For example, Sato *et al.* [SFT+01, SVK+15] measure the shadowing effect and attenuation loss of a truck, placed between a pair of mmWave radios. Keusgen *et al.* [KWP+] conduct measurement in a street canyon environment. They confirmed that reflected paths do not contribute significantly to the received power under the present of LoS, but they still strong enough to maintain a link if the LoS is blocked.

In contrast to the above simulation/measurement studies, our work represents the first large scale profiling of mmWave V2X network using a high-fidelity 3D ray-tracing simulation, combined with a reconfigurable V2X network testbed.

10.2 Solutions for mmWave V2X

A key observation from prior research in mmWave V2X beam management is that, as long as one can tame the beam searching overhead, the V2X link can be as efficient/stable as a static one. Follow this idea, Va *et al.* [VSBH16, VZH15] propose to use location feedback from the vehicles to aid beam prediction. However, their simulation experiments assume perfect GPS information. González-Prelcic *et al.* [GPMRH16] leverage the radar sensing information to help maintain the mmWave link. Other recent work [GMRZ16, CVGP+16] make use of multi-connectivity and neighboring environments to overcome blockage, and reduce beam tracking overhead with a DSRC/LTE side channel. The need for extra hardware or RF channel limits the practical use of these approaches. Asadi *et al.* [AMS+18] proposes a simpler machine learning mechanism to adapt beam selection to the V2X environment. However, our work has shown that a small set of beams can cover most of the possible best AoA/AoD due to their sparsity and persistence. By adopting small codebook, even with the heuristic brute-force beam scanning, the overhead is less than 5% in various scenarios. Thus complicated beam management mechanisms

may not be necessary.

10.3 Solution for related components in mmWave

The AoA/AoD sparsity that we observe in this work shares similarity with other measurement profiling of mmWave channels [WZZ17, SVZR15, Smu02]. By leveraging such sparsity, the beam searching scope can be narrowed down substantially [SZRC16]. However, these work only considered human activities where sparsity exists in only every single location but does not holds across different locations. Our work shows that AoA/AoD sparsity in V2X scenario mainly comes from the fixed-route of the vehicle movement, since most of the AoA/AoDs are pointing from/to the lanes. So the AoA/AoD sparsity in V2X scenario is across locations along the whole route which is not mentioned in previous work.

The potential of spatial reuse in mmWave networks has been widely studied [CPYC12, NSG+15, NLJ+15]. Other work focus on hybrid beamforming [SGR15, HLL+17] and user selection for user group in mmWave MU-MIMO [LS17, LS17]. Recent work [GHC+18] proposes to find best analog configuration for multi-stream beam steering to a given target user-group. In contrast, our work discusses the feasibility of mmWave MU-MIMO using hybrid beamforming and analog beam steering specifically for V2X scenario under different environments and basestation phased array size.

The directionality nature of mmWave beam is also considered to be used in interference reduction [ZZM+14]. Some recent study for mmWave WPAN [NGL+15], mesh network [ZNL+16], and data center system [VTR+10] try to leverage directional mmWave beams and minimize interference. They concluded that Tx/Rx beamforming can effectively reduce interlink interference. But as mentioned Chapter. 9, receiver beamforming cannot fully eliminate the interference and transmitter/basestation beam steering is limited due to the AoD sparsity. Consequently, It is hard to leverage spatial multiplexity between adjacent basestation in mmWave

V2X scenario.

Although these components are well studied in mmWave network, they are not designed specifically for V2X scenario. Thus the performance of these aspects in mmWave V2X scenario still remain unknown. To fulfill this gap, our work is exploring and unveiling mmWave V2X by comprehensive profiling the unique characteristics of these components as mentioned in previous paragraphs.

Chapter 11

Conclusion

In this thesis, we have presented a real-world measurement and large scale simulations of 5G mmWave V2X using a COTS platform and high fidelity 3D ray-tracer. The evaluations dispel some common myths, discover some unknown issues, and show that 5G mmWave V2X works even with simple solutions. Our findings provide hints toward solving significant challenges that remain. Our key insights can be summarised as follows:

(1) The legacy exhaustive beam searching works well even under high mobility, with proper codebook designs. Despite the vehicles' high mobility, the propagation paths show a surprising sparsity and persistence across geometric locations. As a result, small codebooks can be used to achieve full coverages and the beam searching can be done less frequently. Hence the beam management for V2X is well tractable (Chapter 4).

(2) Blockage happens infrequently in mmWave V2X and can be further alleviated. During blockage, the link can be quickly recovered by steering the beams towards the NLoS paths, thanks to the sparsity of these NLoS paths. There is also a non-trivial proportion of cases where there is no NLoS path for link recovery. For those cases, the blockage probability can be significantly reduced by raising basestation height and multi-connectivity. (Chapter 6).

(3) Spatial reuse, multiplexing and interference management can be addressed by heuristic

ways without much performance loss.

Finally, a few points are worth elaborating on, regarding the future work of the evaluation and mmWave V2X application:

(1) This thesis focuses on V2I evaluation to provide insights for the general V2X case, which includes V2V, V2P and etc. There is no doubt that the specific design principals for other V2X cases may still remain unknown and need to be further demystified.

(2) As mentioned in Chapter 3, dense deployment may be required to get full coverage of mmWave V2X basestations. The dense deployment of the mmWave bastations needs a huge amount of infrastructure. Reduction of the basestation infrastructure, such as integrated access and backhaul (IAB), is an interesting and well-motivated research area.

(3) As mentioned in Chapter 8, although non-orthogonal multiplexing achieves better performance than hybrid beamforming in MU-MIMO, it does not apply to all the scenarios. Thus a solution to boost up the performance of MU-MIMO in general scenarios still needs to be discovered.

(4) As mentioned in Chapter 9, Rx beamforming can not fully eliminate the inter-cell interference, and Tx beamforming cannot solely consider avoiding interference. Consequently, an advanced solution to eliminate the inter-cell interference is still unknown.

We believe that our work paves the road towards the high speed and robust mmWave V2X system.

The thesis is currently being prepared for submission for publication. Song Wang; Jingqi Huang; Xinyu Zhang. The dissertation/thesis author was the primary investigator and author of this material.

Bibliography

[3GP17a] 3GPP. 3gpp tr 38.901. https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=2991, 2017.

[3GP17b] 3GPP. Nr; radio resource control (rrc); protocol specification. https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3197, 2017.

[3GP17c] 3GPP. Service requirements for enhanced v2x scenarios. https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3180, 2017.

[3GP18] 3GPP. Study on evaluation methodology of new vehicle-to- everything v2x use cases for lte and nr; (release 15). TR 37.885 V15.0.0, 2018.

[AEALH14] Ahmed Alkhateeb, Omar El Ayach, Geert Leus, and Robert W Heath. Channel estimation and hybrid precoding for millimeter wave cellular systems. *IEEE Journal of Selected Topics in Signal Processing*, 8(5):831–846, 2014.

[AG18a] C. K. Anjinappa and I. Guvenc. Millimeter-wave v2x channels: Propagation statistics, beamforming, and blockage. In *IEEE Vehicular Technology Conference (VTC-Fall)*, 2018.

[AG18b] Chethan Kumar Anjinappa and Ismail Guvenc. Millimeter-wave v2x channels: Propagation statistics, beamforming, and blockage. In *2018 IEEE 88th Vehicular Technology Conference (VTC-Fall)*, pages 1–6. IEEE, 2018.

[air19] Airfide inc. http://airfidenet.com, 2019.

[AMB17] Bogdan Antonescu, Miead Tehrani Moayyed, and Stefano Basagni. mmwave channel propagation modeling for v2x communication systems. In *2017 IEEE 28th Annual International Symposium on Personal, Indoor, and Mobile Radio Communications (PIMRC)*, pages 1–6. IEEE, 2017.

[AMGPH14] Ahmed Alkhateeb, Jianhua Mo, Nuria Gonzalez-Prelcic, and Robert W Heath. Mimo precoding and combining solutions for millimeter-wave systems. *IEEE Communications Magazine*, 52(12):122–131, 2014.

[AMS+18] Arash Asadi, Sabrina Müller, Gek Hong Sim, Anja Klein, and Matthias Hollick. Fml: Fast machine learning for 5g mmwave vehicular communications. In *IEEE INFOCOM 2018-IEEE Conference on Computer Communications*, pages 1961–1969. IEEE, 2018.

[AQE+20] Fawad Ahmad, Hang Qiu, Ray Eells, Fan Bai, and Ramesh Govindan. Carmap-fast 3d feature map updates for automobiles. In *17th USENIX Symposium on Networked Systems Design and Implementation (NSDI 20)*, Santa Clara, CA, 2020. USENIX Association.

[AW17] Mouhamed Abdulla and Henk Wymeersch. Fine-grained reliability for v2v communications around suburban and urban intersections. *arXiv preprint arXiv:1706.10011*, 2017.

[bat20] Wiki - batman-adv - open mesh. `https://www.open-mesh.org/projects/batman-adv/wiki/Wiki`, 2020.

[Com20] Federal Communications Commission. Fcc 15.255. `https://www.ecfr.gov/cgi-bin/text-idx?SID=7aee7a7d5d4a0dab75e7c6519b898a01&mc=true&node=se47.1.15_1255&rgn=div8`, 2020.

[CPYC12] Qian Chen, Xiaoming Peng, Juan Yang, and Francois Chin. Spatial reuse strategy in mmwave wpans with directional antennas. In *2012 IEEE Global Communications Conference (GLOBECOM)*, pages 5392–5397. IEEE, 2012.

[CVGP+16] Junil Choi, Vutha Va, Nuria Gonzalez-Prelcic, Robert Daniels, Chandra R Bhat, and Robert W Heath. Millimeter-wave vehicular communication to support massive automotive sensing. *IEEE Communications Magazine*, 54(12), 2016.

[(DL19a] German Aerospace Center (DLR). Sumo at a glance - sumo documentation. `http://sumo.sourceforge.net/userdoc/Sumo_at_a_Glance.html`, 2019.

[(DL19b] German Aerospace Center (DLR). Traci - sumo documentation. `http://sumo.sourceforge.net/userdoc/TraCI.html`, 2019.

[Dud83] Dan E Dudgeon. Multidimensional digital signal processing. *Engewood Cliffs*, 1983.

[EANH16] Mohammed E Eltayeb, Tareq Y Al-Naffouri, and Robert W Heath. Compressive sensing for blockage detection in vehicular millimeter wave antenna arrays. In *2016 IEEE Global Communications Conference (GLOBECOM)*, pages 1–6. IEEE, 2016.

[Eli20] Vladimir Elistratov. blender-osm: Openstreetmap and terrain for blender. `https://github.com/vvoovv/blender-osm`, 2020.

[GHC+18] Yasaman Ghasempour, Muhammad K Haider, Carlos Cordeiro, Dimitrios Koutsonikolas, and Edward Knightly. Multi-stream beam-training for mmwave mimo networks. In *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking*, pages 225–239, 2018.

[GMD04] Fred K Gruber, Edwin A Marengo, and Anthony J Devaney. Time-reversal imaging with multiple signal classification considering multiple scattering between the targets. *The Journal of the Acoustical Society of America*, 115(6):3042–3047, 2004.

[GMRZ16] Marco Giordani, Marco Mezzavilla, Sundeep Rangan, and Michele Zorzi. Multi-connectivity in 5g mmwave cellular networks. In *Mediterranean Ad Hoc Networking Workshop (Med-Hoc-Net)*, pages 1–7. IEEE, 2016.

[GPMRH16] Nuria González-Prelcic, Roi Méndez-Rial, and Robert W Heath. Radar aided beam alignment in mmwave v2i communications supporting antenna diversity. In *Information Theory and Applications Workshop (ITA)*, 2016.

[GPR+18] Marco Giordani, Michele Polese, Arnab Roy, Douglas Castor, and Michele Zorzi. A tutorial on beam management for 3gpp nr at mmwave frequencies. *IEEE Communications Surveys & Tutorials*, 21(1):173–196, 2018.

[HAR+18] Haitham Hassanieh, Omid Abari, Michael Rodriguez, Mohammed Abdelghany, Dina Katabi, and Piotr Indyk. Fast millimeter wave beam alignment. In *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication*, 2018.

[HLL+17] Chen Hu, Jian Liu, Xiangbai Liao, Yingzhuang Liu, and Jun Wang. A novel equivalent baseband channel of hybrid beamforming in massive multiuser mimo systems. *IEEE Communications Letters*, 22(4):764–767, 2017.

[HWG+19] Danping He, Longhe Wang, Ke Guan, Bo Ai, Junhyeong Kim, and Zhangdui Zhong. Channel characterization for mmwave vehicle-to-infrastructure communications in urban street environment. In *European Conference on Antennas and Propagation (EuCAP)*, 2019.

[IEE12] IEEE Standards Association. IEEE Standards 802.11ad-2012: Enhancements for Very High Throughput in the 60 GHz Band, 2012.

[IEE18] IEEE 802.11ay Task Group. Status of Project IEEE 802.11ay. http://www.ieee802.org/11/Reports/tgay_update.html, 2018.

[ITU13] ITU. P.1410 : Propagation data and prediction methods required for the design of terrestrial broadband radio access systems operating in a frequency range from 3 to 60 ghz. `https://www.itu.int/rec/R-REC-P.1410-5-201202-I/en`, 2013.

[KBGP+18] Aldebaro Klautau, Pedro Batista, Nuria Gonzalez-Prelcic, Yuyang Wang, and Robert W. Heath Jr. 5G MIMO data for machine learning: Application to beam-selection using deep learning. In *2018 Information Theory and Applications Workshop, San Diego*, pages 1–1, 2018.

[KWP+] Wilhelm Keusgen, Richard J Weiler, Michael Peter, Mike Wisotzki, and Bariş Göktepe. Propagation measurements and simulations for millimeter-wave mobile access in a busy urban environment. In *2014 39th International Conference on Infrared, Millimeter, and Terahertz waves (IRMMW-THz)*.

[Lab20] Unwired Labs. Opencellid. `https://opencellid.org/#zoom=16&lat=37.77889&lon=-122.41942`, 2020.

[LS17] Gilwon Lee and Youngchul Sung. A new approach to user scheduling in massive multi-user mimo broadcast channels. *IEEE Transactions on Communications*, pages 1481–1495, 2017.

[NCF+14] Thomas Nitsche, Carlos Cordeiro, Adriana B Flores, Edward W Knightly, Eldad Perahia, and Joerg C Widmer. Ieee 802.11 ad: directional 60 ghz communication for multi-gigabit-per-second wi-fi. *IEEE Communications Magazine*, 52(12):132–141, 2014.

[NGL+15] Yong Niu, Chuhan Gao, Yong Li, Depeng Jin, Li Su, and Dapeng Wu. Boosting spatial reuse via multiple-path multihop scheduling for directional mmwave wpans. *IEEE Transactions on Vehicular Technology*, 65(8):6614–6627, 2015.

[NLJ+15] Yong Niu, Yong Li, Depeng Jin, Li Su, and Athanasios V Vasilakos. A survey of millimeter wave communications (mmwave) for 5g: opportunities and challenges. *Wireless networks*, 21(8):2657–2676, 2015.

[NSG+15] Yong Niu, Li Su, Chuhan Gao, Yong Li, Depeng Jin, and Zhu Han. Exploiting device-to-device communications to enhance spatial reuse for popular content downloading in directional mmwave small cells. *IEEE Transactions on Vehicular Technology*, 65(7):5538–5550, 2015.

[ope20a] Techdata: Engenius ens620ext v1.0.0. `https://openwrt.org/toh/hwdata/engenius/engenius_ens620ext`, 2020.

[Ope20b] OpenStreetMap. Openstreetmap. `https://www.openstreetmap.org/#map=5/38.007/-95.844`, 2020.

[PDDW17] Joan Palacios, Danilo De Donno, and Joerg Widmer. Tracking mm-wave channel dynamics: Fast beam training strategies under mobility. In *IEEE INFOCOM 2017-IEEE Conference on Computer Communications*, pages 1–9. IEEE, 2017.

[PGM+17] Michele Polese, Marco Giordani, Marco Mezzavilla, Sundeep Rangan, and Michele Zorzi. Improved handover through dual connectivity in 5g mmwave mobile networks. *IEEE Journal on Selected Areas in Communications*, 35(9):2069–2084, 2017.

[PSL+18] Joan Palacios, Daniel Steinmetzer, Adrian Loch, Matthias Hollick, and Joerg Widmer. Adaptive codebook optimization for beam training on off-the-shelf ieee 802.11 ad devices. In *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking*, pages 241–255, 2018.

[QAB+18] Hang Qiu, Fawad Ahmad, Fan Bai, Marco Gruteser, and Ramesh Govindan. Avr: Augmented vehicular reality. In *Proceedings of the 16th Annual International Conference on Mobile Systems, Applications, and Services*, MobiSys 2018, page 8195, New York, NY, USA, 2018. Association for Computing Machinery.

[Qua19] Qualcomm. 5g nr mmwave outdoor and indoor deployment strategy. `https://www.qualcomm.com/media/documents/files/deploying-5g-nr-mmwave-for-indoor-outdoor.pdf`, 2019.

[Rem20] Remcom. Wireless em propagation software - wireless insite. `https://www.remcom.com/wireless-insite-em-propagation-software/`, 2020.

[SFT+01] Katsuyoshi Sato, M Fujise, R Tachita, E Hase, and T Nose. Propagation in rof road-vehicle communication system using millimeter wave. In *Proceedings of the IEEE International Vehicle Electronics Conference*, 2001.

[SGR15] Richard A Stirling-Gallacher and Md Saifur Rahman. Multi-user mimo strategies for a millimeter wave communication system using hybrid beam-forming. In *2015 IEEE International Conference on Communications (ICC)*, pages 2437–2443. IEEE, 2015.

[Smu02] Peter Smulders. Exploiting the 60 ghz band for local wireless multimedia access: Prospects and future directions. *IEEE communications magazine*, 40(1):140–147, 2002.

[SVBH18] Takayuki Shimizu, Vutha Va, Gaurav Bansal, and Robert W Heath. Millimeter wave v2x communications: Use cases and design considerations of beam management. In *2018 Asia-Pacific Microwave Conference (APMC)*, pages 183–185. IEEE, 2018.

[SVK+15] Vasilii Semkin, Usman Virk, Aki Karttunen, Katsuyuki Haneda, and Antti V Räisänen. E-band propagation channel measurements in an urban street canyon. In *2015 9th European Conference on Antennas and Propagation (EuCAP)*, 2015.

[SVZR15] Sanjib Sur, Vignesh Venkateswaran, Xinyu Zhang, and Parmesh Ramanathan. 60 ghz indoor networking through flexible beams: A link-level profiling. In *Proceedings of the 2015 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems*, pages 71–84, 2015.

[SWH17] Daniel Steinmetzer, Daniel Wegemer, and Matthias Hollick. Talon tools: The framework for practical ieee 802.11ad research, 2017.

[SZRC16] Sanjib Sur, Xinyu Zhang, Parmesh Ramanathan, and Ranveer Chandra. Beamspy: enabling robust 60 ghz links under blockage. In *13th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 16)*, pages 193–206, 2016.

[TEPN17] Andrea Tassi, Malcolm Egan, Robert J Piechocki, and Andrew Nix. Modeling and design of millimeter-wave networks for highway vehicular communication. *IEEE Transactions on Vehicular Technology*, 66(12), 2017.

[TV05] David Tse and Pramod Viswanath. *Fundamentals of Wireless Communication*. Cambridge University Press, 2005.

[VCS+17] V Va, J Choi, T Shimizu, G Bansal, and RW Heath Jr. Inverse fingerprinting for millimeter wave v2i beam alignment. *ArXiv preprint*, 2017.

[VSBH16] Vutha Va, Takayuki Shimizu, Gaurav Bansal, and Robert W Heath. Beam design for beam switching based millimeter wave vehicle-to-infrastructure communications. In *IEEE International Conference on Communications (ICC)*, 2016.

[VSBHJ16] Vutha Va, Takayuki Shimizu, Gaurav Bansal, and Robert W Heath Jr. Millimeter wave vehicular communications: A survey. *Foundations and Trends® in Networking*, 10(1):1–113, 2016.

[VTR+10] Hars Vardhan, Navine Thomas, Seong-Ryong Ryu, Bhaskar Banerjee, and Ravi Prakash. Wireless data center with millimeter wave network. In *2010 IEEE global telecommunications conference GLOBECOM 2010*, 2010.

[VZH15] Vutha Va, Xinchen Zhang, and Robert W Heath. Beam switching for millimeter wave communication to support high speed trains. In *2015 IEEE 82nd Vehicular Technology Conference (VTC2015-Fall)*, 2015.

[WHZ+20] Song Wang, Jingqi Huang, Xinyu Zhang, Hyoil Kim, and Sujit Dey. X-Array: Approximating Omnidirectional Millimeter-WaveCoverage Using an Array of Phased Arrays. In *ACM MobiCom*, 2020.

[WVMH17] Y. Wang, K. Venugopal, A. F. Molisch, and R. W. Heath. Blockage and coverage analysis with mmwave cross street bss near urban intersections. In *2017 IEEE International Conference on Communications (ICC)*, pages 1–6, May 2017.

[WZ17] Teng Wei and Xinyu Zhang. Pose information assisted 60 ghz networks: Towards seamless coverage and mobility support. In *Proceedings of the 23rd Annual International Conference on Mobile Computing and Networking*, pages 42–55, 2017.

[WZZ17] Teng Wei, Anfu Zhou, and Xinyu Zhang. Facilitating robust 60 ghz network deployment by sensing ambient reflectors. In *14th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 17)*, pages 213–226, 2017.

[YLD+19] Wenqiang Yi, Yuanwei Liu, Yansha Deng, Arumugam Nallanathan, and Robert W Heath. Modeling and analysis of mmwave v2x networks with vehicular platoon systems. *IEEE Journal on Selected Areas in Communications*, 2019.

[ZNL+16] Yun Zhu, Yong Niu, Jiade Li, Dapeng Oliver Wu, Yong Li, and Depeng Jin. Qos-aware scheduling for small cell millimeter wave mesh backhaul. In *2016 IEEE International Conference on Communications (ICC)*, pages 1–6. IEEE, 2016.

[ZWX+18] Anfu Zhou, Leilei Wu, Shaoqing Xu, Huadong Ma, Teng Wei, and Xinyu Zhang. Following the shadow: Agile 3-d beam-steering for 60 ghz wireless networks. In *IEEE INFOCOM 2018-IEEE Conference on Computer Communications*, pages 2375–2383. IEEE, 2018.

[ZZM+14] Yibo Zhu, Zengbin Zhang, Zhinus Marzi, Chris Nelson, Upamanyu Madhow, Ben Y Zhao, and Haitao Zheng. Demystifying 60ghz outdoor picocells. In *Proceedings of the 20th annual international conference on Mobile computing and networking*, pages 5–16, 2014.