

Demystifying Millimeter-Wave V2X: Towards Robust and Efficient Directional Connectivity Under High Mobility

Song Wang*

University of California, San Diego
sowang@ucsd.edu

Jingqi Huang*

University of California, San Diego
jih032@ucsd.edu

Xinyu Zhang

University of California, San Diego
xyzhang@ucsd.edu

ABSTRACT

Millimeter-wave (mmWave) networking represents a core technology to meet the demanding bandwidth requirements of emerging connected vehicles. However, the feasibility of mmWave vehicle-to-everything (V2X) connectivity has long been questioned. One major doubt lies in how the highly directional mmWave links can sustain under high mobility. In this paper, we present the first comprehensive reality check of mmWave V2X networks. We deploy an experimental testbed to mimic a typical mmWave V2X scenario, and customize a COTS mmWave radio to enable microscopic investigation of the channel and the link. We further construct a high-fidelity 3D ray-tracer to reproduce the mmWave characteristics at scale. With this toolset, we study the mmWave V2X coverage, mobility and blockage, codebook/beam management, and spatial multiplexing. Our measurement debunks some common misperceptions of mmWave V2X networks. In particular, due to the constrained roadway network structures, we find the beam management can be handled easily by the often-denounced beam scanning schemes, as long as the codebook is properly designed. Blockage can be almost eliminated through proper basestation deployment and cooperation. Highly effective spatial multiplexing can be realized even without sophisticated MIMO radios. Our work points to possible ways to realize efficient and reliable mmWave networks under high mobility, while maintaining the simplicity of standard network protocols.

CCS CONCEPTS

- Hardware → Wireless devices;
- Networks → Network measurement; Network simulations; Network mobility;

KEYWORDS

millimeter-wave networks, mmWave, V2X, beam management, 5G, MU-MIMO

ACM Reference format:

Song Wang, Jingqi Huang, and Xinyu Zhang. 2020. Demystifying Millimeter-Wave V2X: Towards Robust and Efficient Directional Connectivity Under High Mobility. In *Proceedings of The 26th Annual International Conference on Mobile Computing and Networking, London, United Kingdom, September 21–25, 2020 (MobiCom '20)*, 14 pages.
<https://doi.org/10.1145/3372224.3419208>

1 INTRODUCTION

Vehicle-to-Everything (V2X) connectivity represents a key vertical application of 3GPP 5G NR, which has specified 4 categories of use cases [1]: vehicles platooning, extended sensors, advanced driv-

*Both authors contributed equally to this work.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

MobiCom '20, September 21–25, 2020, London, United Kingdom

© 2020 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-7085-1/20/09.

<https://doi.org/10.1145/3372224.3419208>




Figure 1: Testbed deployment.

ing and remote driving. These use cases typically require high-throughput data synchronization between vehicles and the road infrastructures (henceforth referred to as *UE* and *basestation* following 3GPP terminology). For example, state-of-the-art self-driving cars can generate up to 750Mb of sensor data per second [2], from LIDAR, camera, and on-board diagnostic units. By fusing the data through real-time communication, each vehicle's perceptual range can be extended, creating an intelligent and safer transportation system [3–8].

Millimeter-wave (mmWave) networking is a niche technology to satisfy such demanding applications, due to abundant spectrum resources and high link capacity. However, the practicality of mmWave V2X has long been questioned [5, 9–14]. mmWave is known for its intrinsic limitations, *i.e.*, high attenuation loss, high directionality, vulnerable to blockage, *etc.* These properties seemingly contradict the high mobility and channel dynamics that are typical in V2X scenario. Consequently, although mmWave network protocols (*e.g.*, 802.11ad/ay, 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 most widely cited 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 the transmitter (Tx) and receiver (Rx) radios move, they may need to constantly search for the best-aligned beam pair out of hundreds to thousands of candidates. Standard beam searching protocols traverse many possible beams [15–17], and are often criticized for the large overhead even under human mobility [18–22]. Intuitively, with much higher moving speed, the beam alignment for V2X links needs to be done even more frequently, thus becoming intractable. This is often considered as the most critical challenge for mmWave V2X [9–12, 17, 23, 24].

(2) *Blockage recovery.* The mmWave directional beams are vulnerable to blockages from roadway obstacles including foliage, pedestrian, and tall vehicles [25]. Unlike indoor environment with rich multi-paths, the blocked link cannot be easily recovered as




Figure 2: Basestation and UE setup.

reflection paths are sparse in the V2X environment [17, 26–28].

(3) *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 SDM easily becomes unmanageable, especially under high mobility. *MIMO mmWave* [29–33] 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 paper, we demystify the above challenges using a mmWave V2X testbed together with large-scale 3D ray-tracing simulation. Our experimental facilities aim to represent a 3GPP vehicle-to-infrastructure (V2I) network comprised of *Roadside Unite (RSU)* basestations, which is expected to be the dominating deployment model of mmWave V2X. Our testbed runs in one typical urban/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 with realistic roadway traffic. Our experiments confirm the practicality of mmWave V2X even with standard protocols. Drawing on the experimental results, we identify unique opportunities to simplify the mmWave V2X operations, including beam management, blockage recovery and spatial reuse, without sacrificing performance.

Our key insights are summarized as follows:

(1) The exhaustive beam searching works well even under high mobility, as long as the codebook is properly designed. Prior work [19, 34, 35] revealed the sparsity of mmWave channel between a given pair of Tx/Rx locations (Sec. 7). Our measurement generalizes this property: we found the propagation paths in V2X scenarios show a surprisingly high sparsity and temporal persistence *even across geometric locations*. As a result, a small codebook can be used to achieve full coverage and the beam searching takes negligible time despite high mobility (Sec. 4). Our experiment shows a maximum beam searching latency of 2.4ms across various scenarios and vehicle velocities.

(2) Blockage happens infrequently in mmWave V2X and can be effectively mitigated through simple solutions. We observe an 8.5% blockage rate in urban scenarios with dense traffic, and less than 2% for other scenarios. A blocked link can be quickly recovered by steering its beams towards non-line-of-sight (NLoS) paths, thanks to the sparsity of paths. For links with no NLoS paths, blockage probability can be significantly reduced by raising basestation height, or leveraging the 3GPP NR multi-connectivity feature (Sec. 5).

(3) The spatial multiplexing schemes in mmWave V2X can be simplified due to the high directionality, but intelligent designs are still needed for some key problems. The inter-cell interference

can be significantly reduced by using Rx beamforming on the UE side. The use of multi-array radios, as proposed by 3GPP [36], may cause a destructive *co-phasing* effect that reduces link quality. Interestingly, co-phasing does not affect the beam selection and hence can be addressed in an orthogonal manner. In addition, 3GPP introduces a sophisticated hybrid beamforming scheme, which employs digital precoding across multiple phased arrays, so as to serve multiple UEs simultaneously without interference. However, we found that a simple approach without explicit precoding can achieve superior performance than hybrid beamforming, especially in scenarios with dense traffic.

Although our testbed is built from COTS 60 GHz radios, the above insights are transferable to 3GPP NR mmWave, because we mainly profile the relative impacts of channel/protocol/environmental factors (e.g., beamwidth, road way structure and traffic density) which are independent of the specific mmWave bands in use. Through this measurement campaign, we make the following contributions.

(i) We deploy a mmWave V2X network testbed and customize the radios to enable codebook redesign, beam switching, multi-array operations, and real-time channel quality measurement. The deployment follows the 3GPP guidelines [1], and can be easily reproduced due to its low cost. Furthermore, we build a high-fidelity 3D mmWave channel simulator, by integrating state-of-the-art 3D environment model, vehicular traffic generator, and mmWave ray-tracing engine. The simulation framework has been benchmarked against real measurement in terms of accuracy and can be reused for mmWave V2X research in general. (ii) We demystify the practicality of mmWave V2X networking under high mobility. Both of our testbed measurement and large-scale simulation converge to the somewhat surprising conclusion that mmWave V2X can achieve high performance even with a rudimentary beam management scheme, and blockage probability can be reduced to almost 0 through a high basestation deployment or multi-connectivity.

(iii) We quantitatively evaluate the potential multiplexing gains that can be achieved in mmWave V2X, through Tx/Rx beamforming, multi-array co-phasing, or hybrid beamforming. Our measurement reveals unique opportunities to simplify the multiplexing schemes in V2X, e.g., spatial orthogonality without digital precoding. Our dataset and code for the 3D ray-tracing are available at <http://m3.cs.berkeley.edu/mmwave-v2x-testbed/>.

2 METHODOLOGY

In this section, we introduce the two experimental facilities in our measurement campaign.

2.1 mmw V2X Testbed

Our mmWave V2X testbed is deployed along a 1km road segment with 2 vehicle lanes and 2 bike lanes, as illustrated in the map in Fig. 1. The nearby landscape 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.

2.1.1 *MmWave Radios.* The mmWave basestations are mounted on the roadside lampposts, 5m above the ground, with 80–100m separation. Each basestation covers up to 200m in LoS when the Tx beam with maximum gain is used and Rx turned to quasi-omni mode. The deployment complies with the definition of road-side unit (RSU) in the 3GPP TR 37.885 [1], more specifically the *UE-type RSU* with 100m inter-basestation distance and 5m antenna height. We show in Sec. 3 that RSUs contribute to 97% of the coverage in V2X. Hence our current testbed only comprises RSU basestations




Figure 3: Large-scale 3D ray-tracing overview.

(4 in total).

The mmWave basestation is designed to be lightweight, off-grid, and low cost. As illustrated in Fig. 2, each basestation consists of 2 COTS mmWave radios, an outdoor WiFi mesh router, a 60W solar panel, and a weather-proof enclosure with a high-capacity battery and various switches/wires. The total component costs less than \$2,000. Below we provide more details on the hardware modules.

We build the mmWave basestation from the Airfide 802.11ad 60 GHz radio [37]. 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 2-bit phase shifter per element. The 4 corner elements 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 weight 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 [38]. Note that this is lower than the outdoor EIRP limit for the 60 GHz band (83 dBm [38]). Although the resulting RSS may be lower than the maximum allowable values, our measurement insights in this paper 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 customize the firmware and driver to enable fine-grained channel measurement (Sec. 2.1.5).

The Airfide radio follows the standard 802.11ad Sector Level Sweeping (SLS) beam searching protocol [15]. With SLS, an access point periodically broadcasts a burst of 52-byte frames called *Beam-forming (BF) frames* through each beam in the codebook. A client selects a beam based on the RSS of BF frames and then feeds the beam selection back to the access point. Similar periodic beam sweeping have been standardized in various mmWave systems, *e.g.*, the SLS in 802.11ay [16] and SS burst in 3GPP TS 38.331 [39].

2.1.2 Solar-Based Power System. Drawing power from the mains requires non-trivial modification to the incumbent lamp-post infrastructure. We thus build a simpler, self-contained power system for the basestation, 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 the WiFi router. The battery provides 12V 6A DC output, which is stepped up to 48V by a DC/DC converter to power the PoE switch. Using DC power instead of 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 to 3 hours of extra battery life and makes it smaller and safer. Both the switch

and battery are housed in the weather-proof enclosure.

The solar panel (Newpowa NPA50S-12H) 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 on a sunny day (with the basestation running). The battery supports ~ 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 four 5dBi dipole antennas with a 500m range when running on the 2.4GHz WiFi band. We flush the router with an OpenWrt 19.07.0 system [40] and repurpose it as a mesh router. The mesh interface uses 802.11ac IBSS (ad-hoc) mode, and operates on an L2 routing kernel module *batman-adv* [41]. The *batman-adv* controls the packet routing/forwarding and emulates a virtual switch for all participating nodes with a stable $\leq 5ms$ latency.

The WiFi routers from all 4 basestations form a multi-hop control network, and can be extended into a more sophisticated mesh topology as more nodes join. In order to send commands through the control network, we associate a controller PC to one of the routers through its 5GHz interface. This gateway router can then route 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. Through the control plane network, the controller PC can configure the radio parameters of each basestation, such as codebook, beacon interval, and beam pattern.

2.1.4 Vehicular UE. The testbed vehicle can collect fine-grained statistics of the mmWave channel, and link/network level performance. The vehicular UE 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's moving direction. A GPS receiver is used to log the vehicle UE's location ($\pm 1m$ error) every 0.1s.

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

2.1.5 Customizing the mmWave Radio Firmware. Per-beam RSS extraction. Per-beam RSS is crucial for evaluating beam management schemes. In COTS mmWave radios, per-beam RSS cannot be accessed from the userspace. To overcome this barrier, we disassemble the NIC's firmware, find the on-board memory address that stores the RSS measurement result (extracted from radio tap header of 802.11ad BF frames), use the Talon-tools [42] to inject a piece of code to the firmware, which then transfers the RSS measurement to a memory area accessible by the userspace. We then write a Python program to dump the RSS measurements repeatedly and transmit them to the host PC. A similar RSS extraction approach has been mentioned in recent mmWave measurement systems [43, 44].

Unlike previous mmWave measurement devices [43, 44] which run in *monitor mode* for channel sensing alone, our testbed requires seamless switch between per-beam RSS extraction and normal data transfer, which is only possible in the *managed mode* with a fully functional network stack. We found that, in managed mode, the per beam RSS is updated only when the UE proactively scans for a basestation. We thus identify the corresponding NIC command

Table 1: Major configurations of 3 experimental setups.

	Testbed measurement	Large-scale 3D ray-tracing	Scenario-specific 3D ray-tracing
Venue	mmWave V2X testbed (Sec. 2.1)	Wireless Insite 3D ray-tracing (Sec. 2.2)	Wireless Insite 3D ray-tracing (Sec. 2.2)
Scenarios	Mixture of urban and suburban	Mixture of urban and highway	Urban/Suburban/Highway
Basestations	2 Airfide COTS 802.11ad radios each with $8 \times 6 \times 6$ phased arrays	Simulated 60GHz mmWave radios Omni-coverage multi-panel arrays	Simulated 60GHz mmWave radios Single $6 \times 6 \times 24$ phased arrays or $3 \times 2 \times 6 \times 6 \times 24$ phased arrays for MU-MIMO (Sec. 6.3)
Deployment	3GPP RSU basestations (Sec. 2.1)	Co-siting with LTE eNBs (Sec. 3)	Scenario specific (Fig. 4)
Codebook	39-beam codebook (Sec. 4.1)	39-beam codebook 8-array 39-beam codebook (Sec. 6.2)	39-beam codebook Pruned codebook (Sec. 4.1) “Uneven” codebook (Sec. 4.3)
Vehicle UE	Vehicle UE passing by each basestation # of UE: 1 Average speed: 35mph	10min of SUMO generated traffic # of UEs: 100 Average speed: 35mph	10min of SUMO traffic for each scenario Traffic load: light/heavy (Sec. 6.3) # of UEs: 48/37/101 Average speed: 40/75mph (Sec. 4.3)
Recorded data	Per-beam RSS Vehicle location	Per-beam RSS Vehicle location	Per-beam RSS Vehicle location Channel information: AoA/AoD, CSI
Referred by	Mobility: Sec. 4.1, Sec. 4.2 Multiplexing: Sec. 6.1, Sec. 6.2	Large-scale coverage: Sec. 3	Mobility: Sec. 4.1, Sec. 4.2, Sec. 4.3 Blockage: Sec. 5 Multiplexing: Sec. 6.1, Sec. 6.3

(*WMI_START_SCAN_CMD*), and insert a command call before every RSS extraction cycle. Using this method, we are able to obtain 33 per-beam RSS measurements per second under the managed mode.

Basestation MAC address extraction. Unlike recent mmWave measurement systems [43, 44], our testbed comprises of multiple basestations, and the UE radio must associate the RSS samples to the corresponding basestation. Using the same technique as in RSS extraction, we identify the memory address for the entire BF frame, whose header contains the sender’s MAC address field. We modify our firmware patch so that the MAC address is extracted along with RSS.

Short beam searching periodicity. In 802.11ad, the BF frames are transmitted every Beacon Interval (BI), which is fixed to 100ms on our testbed radio, *i.e.*, only 10 per-beam RSS measurements per second. To improve the sampling granularity, we overwrite the memory address of the BI value in the firmware and reconfigure it to 30ms. Note that the BI can be set as low as 5ms for stable RSS measurement [43]. 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.

2.2 3D Ray Tracer for mmWave V2X

To represent a realistic V2X environment, our ray-tracing simulator takes the real-world 3D map as input, and places various moving vehicles on the roads based on the *SUMO* [45] vehicular traffic generator. We then use *Wireless Insite* [46], a commercial ray-tracing simulator, to track the propagation effects of the mmWave signals (including reflection, scattering, and penetration) and derive the channel/link characteristics. We elaborate on this process below.

2.2.1 Recreating large-scale environment in ray-tracing simulator. We select a $2.8 \times 2.4 \text{ km}^2$ campus area as our simulation field, which involves urban, suburban, rural like sites, and highway sections. The testbed site is included as part of the simulation, so that we can compare and validate the simulation results with real measurement. We export the 2D map of this area from OpenStreetMap (OSM) [47], which provides detailed information on road networks and building footprints. We then use *blender-osm* [48] to render the 2D map into 3D models. Fig. 3 showcases the 3D model for 3 environments versus their corresponding 3D Google maps. The building models

are simplified, *e.g.*, they may miss the glass windows and fine textures on walls, but this level of details are not likely to change the conclusions from our experiments, because the mmWave channel is known to be sparse-dominated by major reflecting objects and less affected by the details in the physical space [49–51].

Next we import the 3D models to the Wireless Insite ray-tracing simulator [46]. Wireless Insite exhaustively ray-traces all 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 pathloss higher than -172 dB assuming 82 dBm EIRP [38] and -90 dBm receiver sensitivity, which is consistent with our testbed radios.

Given the environment models, we deploy two sets of basestations for different experiments: (1) We place a mmWave basestation for every LTE basestation location (fetched from *OpenCellID* [52]) in the geographical area to study the coverage when co-siting 4G/5G. (2) In addition, we handpick typical urban, suburban, and highway locations following the 3GPP TR37.885 guidelines [1] to create environment diversity in simulation. Fig. 4 illustrates the basestation deployment in 3 scenarios.

2.2.2 Bridging traffic and ray-tracing simulator. To generate realistic vehicular traffic, we first feed the OSM 2D road map into SUMO. Then, following the V2X traffic parameters in 3GPP TR 37.885 [1, 53], we configure the dimension, acceleration, and occurrence probability of different types of vehicles. SUMO follows the configuration file and generates continuous traffic flows with 10ms sampling granularity. Then we extract the vehicles’ locations from SUMO’s API, specifically its Traffic Control Interface (TraCI) [54]. Finally, we place 3D models of the vehicles at each location in the corresponding simulation scenario within Wireless Insite. The 3D models of vehicles are treated as a part of the propagation environment in Wireless Insite, and reflection/scattering on the vehicles’ bodies are taken into account.

2.2.3 Reuse Channel Information to Emulate Beamforming. To simulate the channel between a pair of mmWave radios, Wireless Insite treats each antenna element separately. So the simulation time grows quadratically with the number of Tx and Rx antenna elements. Even on a powerful PC with dual Nvidia 1080Ti GPU, 4.7 GHz CPU, and 128 GB SSD RAM, it takes 30 minutes to simulate a 10ms snapshot of the channel between one pair of phased arrays




Figure 4: Scenario-specific 3D ray-tracing basestation deployment of (a) Urban, (b) Suburban, (c) Highway.




Figure 5: Measurement and 3D ray-tracing comparison of: (a) Best beam RSS; (b) Per beam RSS; (b) Best beam index.

with 32 elements each in a typical urban scenario. For mmWave MIMO, the simulation time is further multiplied by the number of Tx and Rx arrays. To overcome the computation barrier, we pre-calculate the magnitude/phase and AoA/AoD for all the propagation paths between the centers of a Tx array and Rx array. Then, assuming far-field links, the received signal from the given Tx array with a given beam G_b can be formulated as follows, without tracing the paths separately between every pair of antenna elements.

$$\mathbf{y} = \sum_i^P A_i G_b(\phi_i - \varphi_m, \theta_i - \vartheta_m) \quad (1)$$

Here A_i denotes the magnitude of the i_{th} path; G_b is the beam's directional gain pattern; ϕ_i and θ_i are the azimuth and elevation AoD of the i_{th} path; φ_m and ϑ_m are the azimuth and elevation facing angles of phased array m ; and P denotes the total number of paths. The beamforming pattern G_b is generated by the DFT method [55], which aligns the phases of all antenna elements for the desired direction. We assume the phased array faces perpendicular to the nearest road segment and parallel to the horizontal plane.

To efficiently simulate the mmWave MIMO setup involving multiple phased arrays, we first pick a reference Tx array, and then compute the received signals from another Tx array with displacement $D = [d_x, d_y, d_z]$ as:

$$\mathbf{y}' = \sum_i^P A_i G_b(\phi_i - \varphi_m, \theta_i - \vartheta_m) e^{j\left(2\pi \frac{\text{mod}(R_i D_i, \lambda)}{\lambda}\right)} \quad (2)$$

where the $e^{j[\cdot]}$ term represents the relative phase difference between the selected Tx array and the reference array; R_i is the rotation matrix $[\sin\theta_i \cos\phi_i, \cos\theta_i, \sin\theta_i \sin\phi_i]^T$, and λ is the carrier wavelength. We then repeat Eq. (2) for each pair of Tx array and Rx array with their corresponding channel information. In this way, we can reuse the pre-calculated paths to emulate phased arrays with different sizes and displacement. The emulation method reduces the ray-tracing computation time by $N_{Tx} \times N_{Rx}$ times, where N_{Tx} and N_{Rx} are the number of phased arrays at Tx and Rx side.

2.2.4 From RSS to Bitrate. To convert the RSS value into bitrate, We use an empirical noise floor setting to calculate SNR from RSS, and then follow existing the IEEE 802.11 ad and 802.11ay [15, 16] to map the SNR value to MCS with a certain bitrate. We calibrate the noise floor so that the MCS variation from the simulation matches the testbed measurement. For interference analysis (Sec. 6.1), the interference is aggregated on top of the noise floor.

2.2.5 Validation of the 3D V2X Channel Simulation. To verify whether the 3D ray tracing can faithfully reproduce the real mmWave channel, we feed a 3D model of our testbed site into the ray tracer, and compare the simulation with the real measurement. The simulation parameters are configured to maximize consistency with the actual radio hardware (phased array geometries, beamforming codebook, deployment location/height, output power, etc.).

Fig. 5 compares the measured and simulated channel as a UE moves along the road till it reaches closest to one basestation. We see both the best-beam RSS and per beam RSS (Fig. 5(a) and (b)) match well across the trajectory. The best beam index (Fig. 5(c)) deviates by less than ± 1 on average. Occasionally, there exist 5-10 dB differences between measurement and simulation because the Automatic Gain Controller (AGC) module in our vehicle UE misreads the link RSS measurement and adds too much Rx gain. Although the testbed radio does not allow disabling the AGC, our results on beam management remain unaffected, since the AGC scales all beam sweeping signal measurements at the same time and the anomalies only happen for a very short period of time.

Note that for mmWave phased array radios, the channel state information including phase and magnitude can be directly derived from the per-beam RSS [44]. Meanwhile, the simulation and measurement match well in terms of per-beam RSS and best beam indexes (i.e., AoA), which are the key factors in making important protocol-level decisions, such as beam switching, spatial reuse, blockage, etc. These benchmarks confirm that our 3D ray-tracing simulation can accurately capture the mmWave channel profile.

2.3 Experimental Scenarios and Setup

Our experiments are performed under 3 types of setups: (1) *Testbed measurement*. (2) *Large-scale 3D ray-tracing*. (3) *Scenario-specific 3D ray-tracing*. We list the major configurations of each setup in Table 1. For ease of exposition, we reuse each set of configurations in the following sections, and refer to them by the type of setup.

3 LARGE SCALE COVERAGE

Recent simulation study [56] demonstrated that, by simply co-locating mmWave basestations with existing LTE basestations, 85% of an urban area can be covered, indicating potential deployment cost reduction by infrastructure sharing. But V2X UEs have uniquely structure geometric distribution, which means that the prior observation [56] deserves a cross-validation. We thus rerun the co-siting deployment in the *large-scale 3D ray-tracing* setup




Figure 6: 3D ray-tracing - Large scale coverage heatmap of co-locating mmWave basestations with LTE basestations.




Figure 7: 3D ray-tracing - Large scale coverage: (a) Coverage contribution, (b) Link RSS.

in Table 1, where the coordinates and the height of all LTE basestations (173 in total, 25.4 per km^2) are obtained from *OpenCellID* [52]. The mmWave basestations are all configured to use 60 GHz multi-panel phased arrays [36] which can achieve omni-coverage [43] with 83dBm EIRP (FCC limit [38]). The heatmap in Fig. 6 shows the RSS spatial distribution of the best beam. Although some major roads are covered with strong RSS, most receive no or poor signals. *Only around 37% of the locations of all the road segments are covered.*

We further categorize the basestations into 2 types: *RSUs*, according to 3GPP TR 37.885 [1], are the ones located within 20m away from a road and not blocked by nearby buildings; *non-RSUs* are the rest. The RSUs comprise 39% of all basestations on our simulated site. Fig. 7 (a) shows the percentage of links served by the two types of basestations and the overall RSS distribution. We see that, among all the locations covered by the basestations, over 97% are owing to the RSUs. Fig. 7 (b) shows the average link RSS is -54 dBm, which translates to a non-trivial bitrate of $3.08Gbps$. On the other hand, the majority of the non-RSUs, though not too far from the road (within 20m-50m), are either within a building or blocked by nearby buildings. The result shows *only the RSUs can effectively provide sufficient coverage for mmWave V2X*. Since the co-siting deployment only provides 37% of the coverage, theoretically *at least 2x extra RSUs need to be deployed to achieve full coverage*. Given the current RSU density ($173 \times 39\% \div 6.72km^2 \approx 10$ per km^2), it means 20 more RSUs per km^2 . The resulting infrastructure and operational cost (tower, backhaul, power, and radio hardware, etc.) will surge accordingly. A lightweight, self-sufficient basestation solution with green energy, wireless backhaul, and cloud RAN, may be needed to contain the cost.

4 BEAM MANAGEMENT UNDER HIGH MOBILITY

In this section, we investigate the impact of mobility on the spatial profile of the mmWave V2X channel along with the implications for beam management. We focus our studies on the standard RSU basestation setting in 3GPP TS 38.201 [57] with 100-200m of range covering a relatively simple road segment. The vehicle mobility changes the propagation paths' AoD/AoA, and an optimal beamforming scheme should direct the signal power towards the dominant AoD/AoA directions. Thus our measurement mainly focuses on the AoD/AoA dynamics. Without loss of generality, we mainly

investigate the AoD dynamics from the basestations' perspectives.

4.1 AoD Sparsity Across Locations

It is well known that the mmWave channel is sparse [19, 51, 58]. *Between a given pair of Tx and Rx, there usually only exists 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* [18, 59]. As a result, practical phased arrays need to generate evenly distributed beams to cover the entire angular Field-of-View (FoV) [17, 51]. The key challenge of mmWave V2X beam management thus lies in efficiently selecting the best beam(s) for arbitrarily located Tx/Rx. Prior research generally demoted the exhaustive beam scanning method, and attempted to reduce the beam search overhead through compressive searching or specialized hardware capabilities [10, 11, 20, 60, 61], which inevitably adds to system complexity. Plus, the formidable computation time of the algorithms may not justify the savings in beam searching time.

In contrast, we find the *mmWave V2X channel exhibits sparsity even across locations*. We made this observation first through the *testbed measurement* experiment in Table 1. In this experiment, the basestation radios are loaded with a 39-beam codebook, which covers $\{70^\circ, 80^\circ, 90^\circ\}$ azimuth and $\{40^\circ, 48^\circ, 56^\circ, 64^\circ, 72^\circ, 80^\circ, 88^\circ, 96^\circ, 104^\circ, 112^\circ, 120^\circ, 128^\circ, 136^\circ\}$ elevation angles with 19.1° beamwidth. Then we drive the vehicle UE along the 200m road segment near each of the basestations at 35mph and record the best beams (the beam with the largest RSS) along the route. Fig. 8 shows the measured histogram of best beams' index. We see that only a small subset of beams between index 30 to 35 are repeatedly selected and most of the rest are unused.

We further examine the underlying channel characteristics leading to this phenomenon. Our testbed cannot provide accurate AoA/AoD estimation since that requires CSI input [62]. Hence we examine the propagation paths from the *Scenario-specific 3D ray-tracing setup* (Table 1). We create urban, suburban, and highway scenarios involving 48, 37, and 101 moving vehicles respectively, and run a 10-minute simulation for each. We accumulate path data from all the vehicle UEs within the 10-minute trial and scatter plot the AoD of all LoS/NLoS paths from basestations in urban and highway scenarios in Fig. 9. Each dot represents a path from a certain UE to the basestation, with its azimuth AoD shown on x axis and elevation on y axis. Clearly, the AoD scatters are clustered into sparse "curves", each corresponding to the vehicles on one road lane. *Unlike indoor mobile UEs whose mobility is mainly induced by 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 outliers due to reflections.* Note that the highway basestation manifests two sets of curves distributed at $-90^\circ \rightarrow 90^\circ$ and $90^\circ \rightarrow -90^\circ$, because it is placed between two roads with reverse directions. This implies that for basestations serving complicated road structures, the AoD curves can be decomposed to simpler curves, each corresponding to a simple road segment. Moreover, multi-panel phased array [36, 43] will be widely used in 5G to expand the basestations' FoV. Each panel covers one or a few simple road segments, and the AoD sparsity still holds for each panel.

Pruning beamforming codebook to reduce the complexity of V2X beam management. Because of the strong correlation between path angles and the road geometries, one can prune the beamforming codebook by selecting only those candidate beams within the AoD range, thereby reducing beam searching overhead




Figure 8: Measurement - Best beam histogram of 4 testbed basestations: (a)RSU1, (b)RSU2, (3)RSU3, (4)RSU4.




Figure 9: 3D ray-tracing - AoD distribution showcases (a) Urban basestation 1. (b) Urban basestation 2. (c) Highway basestation 1. (d) Highway basestation 2.

Table 2: 3D ray-tracing - Codebook size and beam searching latency comparison of full codebook and pruned codebook.

Setting	Full codebook	Pruned codebook		
Setting	# beams	Latency	# beams	Latency
Urban 6 × 6	133	40.3ms	33	1.3ms
Urban 12 × 24	481	141.3ms	45	1.8ms
Suburban 6 × 6	133	40.3ms	28	1.1ms
Suburban 12 × 24	481	141.3ms	39	1.5ms
Highway 6 × 6	133	40.3ms	47	1.8ms
Highway 12 × 24	481	141.3ms	61	2.4ms

substantially. To demonstrate the effectiveness, we examine the results of the *Scenario-specific 3D ray-tracing setup* with the pruned codebook. The beams in the pruned codebook share the same width as the full codebook and their pointing directions are less than half beam width away from the basestation’s AoD curves. For example, for the urban basestation 2 in Fig. 9(b), the beam pointing at 60° azimuth, 70° elevation is pruned since its nearest point on the AoD curve, 46° azimuth, 81° elevation, is 17.8° away from its pointing direction which is larger than half beamwidth.

We follow the 3GPP NR specification to calculate the beam searching latency [17], which is the time required for searching all beams in a codebook, and overhead, which is the time/frequency resources allocated to beam searching over the 20ms T_s and the 400MHz maximum bandwidth. 3GPP NR uses a *SS block* to scan a beam, which takes 17.84 μ s in time and 57.6MHz (without considering frequency repetition [17]) in frequency, and has a sweeping periodicity of $T_s = 20ms$. Table 2 compares the beam searching latency of the pruned and full codebook. Since 3GPP NR [63] limits the number of beams searched within T_s interval to 64 beams, the full coverage codebook cannot finish a full scan in one T_s . The pruned codebooks, on the other hand, all fit within one T_s with the latency $\leq 2.4ms$, which translates to only $\frac{2.4ms \times 57.6MHz}{20ms \times 400MHz} = 1.7\%$ overhead under the default $T_s = 20ms$ in NR [63]. Due to more driving lanes, the highway has the most beams after pruning, but still has a low latency of 2.4ms. The suburban reduces to the smallest codebook since there are not too many lanes nor NLoS reflections.

4.2 AoD Persistence over Time

It is generally observed that, for conventional indoor mmWave systems, higher mobility necessitates more frequent beam searching, which leads to huge overhead [19, 64]. This is also often considered the most significant challenge for mmWave V2X [9, 11, 51] given the vehicles’ high speed. However, contrary to this intuition, we find through the *testbed measurement* that the AoD changes slowly

for most of regions, and beam searching only needs to be called occasionally. We load the basestation RSU2 in our testbed with the aforementioned 39-beam codebook and drive the vehicle UE along the 120m road segment on one side of RSU2 at 35mph. Fig. 10 shows the best beam index versus the UE’s distance to the basestation. When the vehicle is more than 20m away from the basestation, the best beam is either 30 or 31, which point at 70° and 80° respectively, implying the AoD change in this 100m area is 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 the geometric relations between the vehicle UE and basestation. The AoD’s changing rate ω_v , *i.e.*, vehicle’s angular velocity relative to the basestation, is a fraction of the vehicle’s linear velocity v_v by link distance d , *i.e.*, $\omega_v = \frac{v_v}{d}$. In a typical RSU setup, d as denominator is often large (e.g. 20-100m) compared to the v_v (e.g., 22m/s-50m/s) for far regions. So the AoD changes very slowly in these far regions. Fig. 11 shows the scatter plot of the 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 AoD changes over 20°/s happen when the vehicle is within 20m of the basestation. We emphasize that the slow change of AoD shows up not only in testbed measurement (Fig. 10), but also simulation which contains multiple scenarios (Fig. 11 and Fig. 12). This differs from existing observations of indoor mmWave networks [18–20]. It holds true for mmWave V2X because the high mmWave basestation (5-35m above the ground [1]) promises a longer link distance. Also, the vehicle UEs would not experience sudden orientation changes, unlike indoor handheld mobile devices [21]. These unique characteristics of vehicle UEs result in a slower AoD change over time.

Since the far regions dominate a vehicle’s trajectory, an infrequent beam searching would not sacrifice link performance for most of the time. Fig. 12 shows the histogram of AoD changes of all traces from the urban and highway scenarios. We mark the AoD change thresholds to the values that cause 1 dB and 3 dB link RSS loss (comparing to the oracle) with beam searching interval $T_s = 100ms$, which is consistent with our testbed setup, and falls between the largest and second largest T_s values (160ms and 80ms) proposed by 3GPP NR [63]. The 100ms beam searching interval causes an overhead of $\frac{2.4ms \times 57.6MHz}{100ms \times 400MHz} = 0.35\%$. We observe that, with a large 100ms beam searching interval, only a small fraction of links (<5%) suffer from RSS loss of over 1dB, and even smaller (<1.2%) for 3dB loss. Therefore, for the majority of the time, frequent beam




Figure 10: *Measurement - Best beam index dynamics in near/far regions of a testbed basestation.*




Figure 11: *3D ray-tracing - AoD changes v.s. distance to basestation.*

Figure 12: *3D ray-tracing - AoD changes histogram.*

Table 3: *3D ray-tracing - RSS loss (comparing with $T_s = 5ms$ in the near regions.*

T_s	6x6 Array		12x24 Array	
	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

searching is not necessary.

4.3 Beam Management and Codebook Optimization

Hierarchical beam refinement offers marginal gain. Certain mmWave protocols like 802.11ad further run a semi-hierarchical beam searching scheme after the exhaustive beam searching, called *beam refinement* [15, 16], where a narrower beam within the FoV of the current beam is selected to improve the link gain. Similar to beam refinement, recent work [20, 65] proposed intelligent hierarchical beam searching algorithms to further reduce the overhead down to $O(\log n)$, n being the number of beams. In mmWave V2X, such hierarchical beam refinement may not be needed for typical RSU basestations. In Sec. 4.1 and 4.2, we showed that even with the narrowest beam, the standard exhaustive beam searching overhead is still less than 1.7%. On the other hand, most hierarchical beam refinement algorithms involve significant computation overhead that cannot be justified by the 1.7% of saving.

Note that, our foregoing discussion about exhaustive beam scanning/searching assumes the Tx (basestation) scans its beams while the Rx (UE) operates in quasi-omni mode. This is consistent with the beam searching procedure in 3GPP NR, or the sector-level sweeping (SLS) in 802.11ad. These standard protocols also involve an Rx-side beam scanning, which is the reciprocal procedure. Due to reciprocity of the AoD/AoA along the LoS path, the AoA should inherit the same sparsity and persistence properties. So the UE codebook can be pruned in the same way as the basestation's. Accordingly, the Rx-side beam scanning would involve marginal overhead just

as the Tx. More generally, after pruning, the Tx plus Rx scanning overhead is sub-linear with respect to the number of beams in the original codebook.

Beam searching in near basestation regions. Whereas a small codebook can achieve optimal performance thanks to AoD sparsity and persistence, the beam selection in the near-basestation regions is still non-trivial. As mentioned above, the AoD changes more quickly here, requiring shorter beam scanning intervals. So even a smaller codebook needs to make a tradeoff between beam searching overhead and RSS loss. To demonstrate this effect, we examine the results of the *scenario-specific 3D ray-tracing* with $T_s = \{20ms, 40ms, 80ms, 160ms\}$ (proposed as 3GPP NR options for T_s [66]), under two different phased array sizes (6×6 and 12×24), and two vehicle UE speeds (40mph and 75mph). Table 3 shows the mean RSS loss of different T_s settings comparing to the minimal $T_s = 5ms$. We see over 6dB of loss for $T_s > 40ms$ on a 12×24 array and over 3dB loss for $T_s > 80ms$ on a smaller 6×6 array. The result implies that, for vehicles with high mobility, the UE either chooses the minimal $T_s = 5ms$ beam searching interval and raises the overhead up to $\frac{2.4ms \times 57.6MHz}{5ms \times 400MHz} = 6.9\%$, or chooses a larger T_s and suffers up to 16.3dB RSS loss. Next, we show that a proper codebook design can make a much better tradeoff between overhead and RSS loss.

Reducing overhead with uneven codebook. An intuitive way of reducing beam sweeping overhead is to “widen” the beams pointing to the areas near the basestation. Albeit having lower gain, wide beams cover larger angular range and hence require less frequent beam sweeping. For example, a 27.5dBi and a 20.7dBi beam generated by the DFT method [55] have a main-lobe width of 2.4° and 11.4°, respectively. For a vehicle UE 10m away from a basestation moving at 75mph, according to the angular speed equation in Sec. 4.1, it takes $T_s = 6.2ms$ for the 27.5dBi beam to keep track of the client and $T_s = 29.7ms$ for the 20.7dBi beam. The wide beam reduces the beam management overhead from 5.6% to 1.1%. On the other hand, the higher gain of the narrow beam does not necessarily translate to higher throughput. In near-basestation regions, the links usually have a large SNR margin. So a wide beam may lead to the same link bit-rate despite the lower directional gain. Therefore, *an “uneven” codebook with high directionality for far regions and wide beams for near regions can potentially reduce the beam management overhead without sacrificing link quality.*

To verify the advantage of the uneven codebook, we revisit the urban scenario in the *scenario-specific 3D ray-tracing* with the “Fit-to-AoD” codebook and the “Uneven” codebook. The “Fit-to-AoD” codebook, whose beam patterns are shown in Fig. 13 (a), is the same as the *pruned codebook* for this basestation (Sec. 4.1). The “Uneven” codebook, as shown in Fig. 13, has 4 4.7° DFT beams pointing at the basestation’s far region and 17 11.4° DFT beams pointing at the near region. We examine all near region links’ *beam coherent time* (i.e., the duration within which a certain beam stays as the best beam), and bitrate loss of using two codebooks under different T_s compared to using the optimal best beam. From the result in Fig. 14, we see that the uneven codebook exhibits much larger beam coherent time. Only 5.2% of the uneven codebook links have a beam coherent time less than 20ms, comparing to 29.0% of even codebook links. This indicates under $T_s = 20ms$, 94.8% of the links always keep the same best beam under uneven codebook, in contrast to 71.0% of links with even codebook. Fig. 15 shows the bitrate loss. When $T_s = 100ms$, the uneven codebook improves the number of links with 0 loss from 75% to 85%; when $T_s = 20ms$, the uneven codebook closes the margin of 0 loss links from 94% to 98%.

The AoD sparsity and uneven codebook greatly reduce code-




Figure 13: Simulated Beam patterns of (a) Fit-to-AoD even codebook; (b) Uneven codebook;




Figure 14: 3D ray-tracing - Figure 15: 3D ray-tracing - Bi-Beam coherent time com-plateau loss comparison of two codebooks.

book/beam management complexity. However, since the spatial channel profile depends on the basestation deployment and road geometries, the uneven codebook needs to be tailored in a *site-specific* way so that the beams' pointing direction fits to the AoD curves and the width of near region beams is optimized based on the vehicles' typical speed.

5 BLOCKAGE: IMPACTS AND REMEDIES

To investigate the severity of blockage in mmWave V2X, we examine the channel measurement in the *scenario-specific* 3D ray-tracing (Table 1). We categorize blockages into two types: *recoverable blockage* which can be saved by strong NLoS reflection paths, and *fatal blockage* where the link breaks once the LoS is blocked. Fig. 17 shows the blockage probability in the 3 scenarios. Overall, only 8.5% of the V2X links in the urban scenario suffer from blockage, and 82.4% of the blockages are recoverable. The suburban and highway links suffer even less from blockage, 1% and 2% respectively, due to the lack of building obstacles. This seemingly counter-intuitive low blockage probability originates from the roadside deployment of the basestations. Following the 3GPP TR37.885 [1], the RSU basestations are deployed near the roads and 5-10m above the ground, rarely with any buildings or trees in sight. Even tall trucks can barely block the LoS paths of other nearby UEs.

When a recoverable blockage occurs, the mmWave link needs to quickly identify alternative NLoS reflection paths. Due to sparsity of the AoDs across locations (Sec. 4.1), we expect a certain correlation exists between the strongest NLoS path and the LoS path right before blockage. To verify this hypothesis, Fig. 16 plots the distribution of azimuth AoD gap between the two paths in the foregoing experiments. We see the AoD gaps all fall below 20° with only a few outliers. In particular, in the suburban and highway scenarios, the AoD gaps fall below 10° for 50% and 99% of the recoverable links, respectively. The result implies that *to recover from blockage, a mmWave V2X link only needs to prioritize and search for alternative beams adjacent to the latest LoS beam, in order to reduce the overhead*.

On the other hand, to eliminate the fatal blockages, one simple solution is to raise the basestation's height. This way, the LoS path bears a larger elevation angle relative to the ground, and can overcome even tall vehicle obstacles. To verify its effectiveness, we

increase the basestation height by 10m in the previous simulation. The first group of bars in Fig. 18 shows the percentage of fatal blockage links that appear in low basestation scenarios but not in high basestation scenarios. We see that, *over 99% of the fatal blockages are saved by raising the basestations, regardless of the deployment scenarios*.

In case when raising the basestation height is infeasible due to cost or space constraints, an alternative solution is *Multi-Connectivity* (MC)—a 3GPP NR feature that allows UE devices to connect to more than one basestations simultaneously to improve robustness [67]. To verify its effectiveness, we repeat the above simulation with MC enabled. Since the availability of MC links may be affected by the traffic density, we generate *dense* traffic for both urban and highway scenarios in SUMO with all the traffic parameters (including vehicle speed and the distance between adjacent vehicles and etc.) following the classic traffic model [68, 69], which has been proven to faithfully represent real roadway traffic patterns. The traffic in the previous experiment serves as *normal traffic*. The number of the vehicles in the same area of dense traffic is roughly 2× over the normal traffic. The right group of bars in Fig. 18 shows that MC can save 95% of the fatal blockages under normal traffic and 40% for dense traffic. MC performs worse in dense traffic because the MC links have a higher chance to suffer blockage from more vehicles in dense traffic. Many existing studies [9, 70, 71] have explored the possibility of overcoming blockage using surrounding reflection objects (ground, buildings, other vehicles, etc.). But these solutions are opportunistic by nature and environment-dependent.

6 SPATIAL MULTIPLEXING

6.1 Inter-cell Spatial Reuse

Whereas the high directionality of mmWave radios reduces interference footprint, the dense deployment of V2X basestations may still exacerbate inter-cell interference. To characterize the inter-cell spatial reuse, we first measure the interference between a pair of basestations in the *testbed measurement* setup as a small-scale showcase. Then we examine all basestations under the urban scenario in the *scenario-specific* 3D ray-tracing, and analyze the impact of worst-case interference on all UEs. Without loss of generality, we focus on downlink only. To isolate the impacts of frequency reuse, we configure all basestations to share one frequency band for both the testbed experiment and the simulation. For the testbed experiment, we drive the vehicle UE starting from RSU 3 towards RSU 2 while simultaneously measuring the per-beam RSS from both basestations. We assume handoff occurs at the middle point in between. For each UE location, the serving basestation selects the best beam, whereas the SINR is calculated assuming the strongest interfering beam is used by the other basestation. Fig. 19 shows the measurement setup and interference versus the vehicle UE's distance to RSU 3. The interference intensifies as the distance increases. Some of the cell edge (long-distance) links have SINR below 0 dB, rendering the link unusable. Meanwhile, the UE, in this case a quasi-omni Rx, usually falls in the range of two adjacent basestations. Due to the dense deployment, even for cell center UEs (*i.e.* link distance is around 10m), the average SINR is merely 33dB, which translates to a 37.5% throughput loss comparing to the interference-free case.

Receiver beamforming can substantially reduce inter-cell interference. An intuitive way to avoid interference is using Rx beams to amplify the desired basestation's signal while suppressing others. Since our testbed radio does not support Rx beamforming, we emulate Rx beamforming by placing the Rx to face towards the desired basestation (the radio cannot receive the signal from




Figure 16: 3D ray-tracing - Distribution of angle gaps between recoverable paths AoD and last LoS AoD.




Figure 17: 3D ray-tracing - Percentage of links under recoverable blockages and fatal blockages.




Figure 18: 3D ray-tracing - Percentages of fatal blockages saved by high basestation and multi-connectivity.

the backside), and repeat the measurement mentioned earlier in this section. The interference level with Rx beamforming is shown in Fig. 19. We see that the inter-cell interference becomes negligible. We further cross-validate the results in the simulated scenario. Fig. 20 shows the SINR distribution w/ and w/o Rx beamforming. We see *Rx beamforming effectively reduces the interference, improving SINR by more than 30 dB on average*. Even the worst link SNR improves by 35 dB. However, Rx beamforming does not eradicate all interference. This is because: (i) The sidelobes due to imperfect beamforming may still pick up the interference signal from undesired directions, and (ii) Rx beamforming can only isolate interfering signals with a large enough angle separation.

Transmit beamforming needs to sacrifice link quality for interference avoidance. 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 UE is served with a high-quality beam as well. To investigate this tradeoff, we revisit the links under the urban scenario in the *scenario-specific 3D ray-tracing*. For each UE location in the simulation, we exhaustively apply every Tx beam to serve the link, and record the link RSS and the strongest interference this Tx beam causes to other basestations’ UEs. Fig. 21 shows a scatter plot of a link’s RSS versus the strongest interference. In general, *the interference increases as the link RSS increases*. We mark the areas with link RSS over -50dBm and interference level less than -15dB in the figure. The links falling in this area have large RSS and low interference and are good candidates for interference avoidance. We see there are very few links in this area, indicating that *transmit beamforming rarely offers any opportunities to achieve high link RSS while maintaining a low interference*.

In summary, mmWave V2X can achieve spatial reuse for most UEs with Rx beamforming, while UEs without Rx beamforming capabilities should be served by multiplexing (e.g., time/frequency) across neighboring basestations rather than Tx beamforming. The Tx/Rx beamforming capability is widely adopted in mmWave systems, especially for outdoor scenarios [2]. Hence this observation applies to general V2X networks regardless of the specific mmWave frequency bands being used.

6.2 Multi-Array basestations

3GPP NR basestations adopt “Array-of-Phased-Arrays” (APA) [36, 43], *i.e.*, multiple phased array panels form a large planar or circular layout for gain or coverage improvement. Recent work [43] pointed out that such APA structure may suffer from a “co-phasing” effect, *i.e.*, when the different phased arrays beamform simultaneously to the same UE, their signals may combine incoherently, thus degrading link quality. To verify whether co-phasing impacts mmWave V2X, we repeat the *testbed measurement* experiment (Table 1) with 8 phased arrays activated on each basestation. All the 8 phased arrays are facing the same direction and are loaded with the same 39-beam codebook. Note that loading the same codebook on multiple arrays is a common practice among multi-array mmWave

radios, because the hardware register space is limited and the standard mmWave protocols constrain the total number of beams a radio can scan [43]. Then we compare the result with the single-array measurement in Sec. 4.1. We scale the RSS value from the single-array trace by 8 (linear scale) to emulate the ideal coherently combined RSS. Fig. 22 shows the strongest beams’ RSS. We see in the multi-array trace, at 87m, 115, and 120m, the co-phasing effect manifests in the form of RSS “notches” (marked in red), leading to up to 8dB loss compared to the coherent trace.

The co-phasing effect creates significant fluctuations in link quality and will be harmful to upper layer applications. However, we find *the co-phasing effect does not affect beam selection*. Fig. 23 shows that the best beam indices in the above two experiments are nearly identical. This is because, for the same UE location, co-phasing creates the same RSS loss across all beams. Thus, the best beam index remains unchanged after co-phasing. Therefore, *the beam selection scheme can keep its simplicity and does not need to be changed to accommodate co-phasing*.

6.3 MU-MIMO: Channel Orthogonality or Spatial Orthogonality?

Huge overhead in mmWave MU-MIMO V2X. MmWave MU-MIMO is a major feature in 3GPP NR and the upcoming 802.11ay [16, 57]. Unlike the aforementioned multi-array setup, mmWave MU-MIMO requires that the basestation has not only multiple phased arrays, but also multiple RF chains, so that it can simultaneously send multiple data streams to different UEs. A hybrid beamforming algorithm is generally used to cancel the inter-UE interference:

$$\hat{y} = HWP\bar{x} \quad (3)$$

where H is the channel matrix from the basestation to a selected UE group, W and P are the beamforming weight and precoding matrix, and \bar{x} is the desired signals targeting different UEs.

The feasibility and effectiveness of hybrid beamforming strongly depends on whether the channels between the basestation and different UEs are orthogonal. Poor orthogonality may make the precoding matrix P unsolvable. Orthogonality can be quantified by the *channel condition number*, which should be close to 1 in order for all the UEs to be uncorrelated [72]. To investigate the dynamics of the channel orthogonality, we evaluate the coherent time of the channel condition number. We simulate a basestation with $3 \times 12 \times 24$ arrays in the *scenario-specific 3D ray-tracing* experiment (Table 1), which can serve 3 UEs simultaneously. We exhaustively traverse all UE group combinations and record all the time durations with condition numbers consistently below a certain threshold. Here we define the threshold as the channel condition number which translates to an average bitrate of 2480Mbps ($\frac{1}{4}$ of the maximum in the *scenario-specific 3D ray-tracing* results). Fig. 24 shows the CDF of the durations. We see that in all 3 scenarios, there is a non-trivial fraction of UE groups with extremely short coherent time. At least 40% of the UE groups’ coherent time durations are less than 10ms




Figure 19: **Measurement - Interference w/ and w/o receiver beamforming.**




Figure 20: **3D ray-tracing - Interference SINR distribution.**




Figure 21: **3D ray-tracing - Beam used freq v.s. resulting interference.**




Figure 22: **Measurement - Co-phasing effect showcase.**




Figure 23: **Measurement - Best beam index dynamics.**

and 50% less than 20ms. Therefore, to ensure a reasonably high link quality, a proper UE grouping algorithm needs to be invoked at least every 10ms for 40% of time. This may result in formidable overhead, since evaluating the orthogonality requires all the UEs within the basestation's coverage to feedback their CSI.

Harnessing spatial orthogonality to simplify MU-MIMO.

An alternative and simpler way of harnessing MU-MIMO gain is to directly steer different phased arrays' beams to different UEs:

$$\bar{y} = H[w_1 w_2 \cdots w_N]^T I \bar{x} \quad (4)$$

where I is identity precoding matrix, and the beamforming weights w_1, w_2, \dots, w_N steer N phased arrays to N different UEs. We refer to the approach as *spatial-orthogonal multiplexing*, and the aforementioned hybrid beamforming method as *channel orthogonal multiplexing*. Spatial-orthogonal multiplexing is often considered inferior to hybrid-beamforming since it only uses analog beamforming to mitigate inter-UE interference. Nonetheless, due to the sparse multipath in mmWave V2X, the channel orthogonality mainly comes from the angular separation of the links [73]. Hence, as long as the different basestation-to-UE LoS paths have large angular separation, the mutual interference may become minimal. One obvious advantage of this approach lies in its *extremely low feedback overhead*. Only a single value, *i.e.*, AoD, needs to be fed back from each UE to the basestation, and the AoD can be estimated by using the beacon broadcast from the basestation [43]. In contrast, *hybrid beamforming requires CSI feedback, whose overhead is orders of magnitude higher*—comprising a multiplication of the number of UEs, subcarriers, and beams.

To verify the effectiveness of spatial-orthogonal multiplexing, we compare it with hybrid beamforming within the *scenario-specific 3D ray-tracing setup* (Table 1), and use 12×24 arrays in both methods. For the former, we first identify a UE with the highest RSS, and then group it with other UEs with largest AoD separations. For the latter, we exhaustively search for the UE groups with the best channel condition and then apply zero-forcing precoding. To penalize hybrid beamforming for its CSI feedback overhead, we multiple its throughput by $\frac{3T_o}{T_s}$ where T_o is each UE's feedback overhead in 3GPP NR (Eq. 16 in [17]). The box plot in Fig. 25 shows the 90%, 75%, median, 25%, and 10% throughput of each UE. *Spatial orthogonal multiplexing demonstrates a median throughput gain of 85%, 107%, and -3% under urban, suburban and highway scenarios, respectively.* Note that spatial orthogonal multiplexing achieves lower throughput gain (-3%) in the highway scenario. This is because the vehicles on the highway have larger separation, and channel-orthogonal user groups often have strong orthogonality which leads to high throughput. On the other hand, the absolute median throughput on the high way is larger than the other two scenarios. This is caused by the sparser basestation deployment and more lanes in the highway scenario where one basestation's angular coverage is larger and the possibility of finding a set of UEs with large angular separation is higher. Despite the lower gain in the highway scenario, spatial-orthogonal multiplexing shows much less throughput standard deviation, *i.e.*, 0.83 Gbps which is only 23% of channel orthogonal multiplexing in the highway scenario. This implies that spatial-orthogonal multiplexing provides not only good but also stable throughput performance.

The surprising advantage of spatial orthogonality mainly comes from its superior UE selection. We observe that when there is not enough high RSS UEs to form an orthogonal group, the condition number based UE selection in hybrid beamforming would rather choose a set of UEs with low RSS than mixing high RSS UEs with low RSS UEs. In other words, the condition number based user selection scheme *prioritizes equal power allocation instead of achievable throughput* [74]. This becomes a problem in mmWave V2X because of the uneven distribution of RSS caused by large spatial separation of UEs. On the other hand, the spatial-orthogonal multiplexing chooses UEs based on their individual RSS, which accurately reflects the throughput given good spatial orthogonality.

In practice, the road traffic density and the phased array size may affect the opportunity of finding UEs with sufficient angular separations. To quantify such opportunities for spatial-orthogonal multiplexing, we vary the phased array size in the above simulation scenario. Fig. 26 shows the percentage of vehicles which can be included in a group with angular separations larger than the beam width, which we denote as *candidate group*. The percentage is heavily affected by the deployment scenario and the phased array size. In the highway scenario, around 35% of vehicles can always form a candidate group regardless of the basestations' phased array size. However, the urban and suburban scenarios require the basestation use large phased arrays to form candidate groups. In particular, for the urban scenario, no vehicle can form a candidate group over 40% of the time using 6×6 arrays and 90% of the time using 2×2 . For the suburban, over 80% of the time, no valid vehicle group for spatial-orthogonal multiplexing exists when using a phased array smaller than 12×24 . To summarize, *the simple spatial-orthogonal multiplexing approach has a higher chance to benefit the highway scenario or the basestations with large phased arrays. Yet hybrid beamforming is still needed to handle the cases where inter-UE*




Figure 24: 3D ray-tracing - Coherent time CDF of UEs served by hybrid beamforming (channel orthogonal multiplexing).




Figure 25: 3D ray-tracing - Comparison between channel and spatial-orthogonal multiplexing.




Figure 26: 3D ray-tracing - Opportunities for spatial-orthogonal multiplexing.

interference is unavoidable.

7 RELATED WORK

mmWave V2X simulation and measurement. Recent research used simulation extensively to examine the mmWave V2X channel, and reveal the unique challenges imposed by a mix of high mobility and directionality. Anjinappa and He *et al.* [12, 75] studied the temporal and angular characteristics of sub-6 GHz and mmWave bands along an urban trajectory, and provided guidelines for effective beam tracking. Both works are limited to the beam tracking problem only. Tassi *et al.* [76] characterized the downlink performance of a mmWave network along a highway, and analyzed the link outage probability and coverage probability in the presence of truck blockage. Antonescu *et al.* [27] provided instructions for modeling the channel propagation, blockage behaviour and multipath fading in the simulated environment. Unlike our work, the above studies are based on statistical channel simulation alone. It is unclear to what extent the results hold in real scenarios.

Due to 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 channel profiling without real-time beam management. For example, Sato *et al.* [77, 78] measured the shadowing effect and attenuation loss of a truck placed between a pair of mmWave radios. Keusgen *et al.* [79] conducted measurement in a specific street canyon environment. They pointed out that reflected paths do not contribute significantly to the received power compared with LoS, but may still save the link when 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.

Network protocols 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. Following this idea, Va *et al.* [23, 80] proposed to use location feedback from the vehicles to aid beam prediction. However, their simulation experiments assume perfect GPS information. González-Prelcic *et al.* [11] leveraged radar sensing information to help maintain the mmWave link. Other recent works [81, 82] made use of multi-connectivity and ambient reflectors 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.* [24] proposed a simpler machine learning mechanism to adapt beam selection to the V2X environment. In contrast, our work shows that a small set of beams can cover most of the possible best AoA/AoD due to their sparsity and persistence. By adopting a smaller codebook, even with the heuristic brute-force beam scanning, the overhead is less than 5% in typical scenarios. Thus complicated beam management mechanisms may not be necessary.

General solutions for efficient and reliable mmWave networks. The AoA/AoD sparsity observed in this paper extends the insights from existing measurement profiling of mmWave channels [19, 34, 35]. However, the existing studies only considered indoor human mobility where sparsity exists for every single link but does not hold 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 the V2X scenario exists across locations along the whole route.

The potential of spatial reuse in mmWave networks has been widely studied [83–85]. Recent research investigated hybrid beamforming [86, 87] and user selection for mmWave MU-MIMO [88], and proposed to find best analog configurations for multi-stream beams [89]. In contrast, our work discusses the feasibility of mmWave MU-MIMO using channel and spatial orthogonal multiplexing specifically for V2X scenarios. We not only reveal a surprising fact that spatial-orthogonal multiplexing may perform better than the channel-orthogonal multiplexing, but also find that the effectiveness of spatial-orthogonal multiplexing can be affected by traffic density and basestation phased array sizes.

The directionality nature of mmWave beam is also considered to be used in interference reduction [90]. Recent studies on mmWave WPAN [91], mesh network [92], and data center networks [93] tried to leverage directional mmWave beams and minimize interference. They concluded that Tx/Rx beamforming can effectively reduce inter-link interference. But as mentioned Sec. 6.1, Rx beamforming cannot fully eliminate the interference, and Rx beam steering is limited due to the AoD sparsity. Hence, it is non-trivial to leverage spatial reuse between adjacent basestations in mmWave V2X.

8 CONCLUSION

In this paper, we have presented a real-world measurement and large-scale 3D ray-tracing based characterization of mmWave V2X networks. Our experiments dispel some common myths, discover some unknown issues, and show that mmWave V2X works even with simple solutions. Our findings provide hints that can help accelerate the deployment of mmWave V2X. Our experimental facilities aim to reproduce typical 5G V2X scenarios following the 3GPP NR guidelines. Due to hardware capabilities and resource constraints, our current testbed is limited in scale and has to be complemented by 3D ray tracing simulation. As future work, we will expand the testbed deployment, and use state-of-the-art mmWave MIMO software radios [94] to conduct a more in-depth cross-layer evaluation of mmWave V2X.

ACKNOWLEDGEMENT

We sincerely thank our anonymous shepherd and reviewers for their insightful feedback. This work was supported in part by the US National Science Foundation through NSF CNS-1506657, CNS-1617321, CNS-1854472, CNS-1925767, CNS-1952942.

REFERENCES

[1] 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.

[2] 3GPP, "Service requirements for enhanced v2x scenarios," <https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3180>, 2017.

[3] H. Cao, S. Gangakhedkar, A. R. Ali, M. Gharba, and J. Eichinger, "A 5g v2x testbed for cooperative automated driving," in *2016 IEEE Vehicular Networking Conference (VNC)*. IEEE, 2016, pp. 1–4.

[4] K. Sakaguchi and R. Fukatsu, "Cooperative perception realized by millimeter-wave v2v for safe automated driving," in *2018 Asia-Pacific Microwave Conference (APMC)*. IEEE, 2018, pp. 180–182.

[5] T. Shimizu, V. Va, G. Bansal, and R. W. Heath, "Millimeter wave v2x communications: Use cases and design considerations of beam management," in *2018 Asia-Pacific Microwave Conference (APMC)*. IEEE, 2018, pp. 183–185.

[6] F. Ahmad, H. Qiu, R. Eells, F. Bai, and R. Govindan, "Carmap-fast 3d feature map updates for automobiles," in *USENIX Symposium on Networked Systems Design and Implementation*, 2020.

[7] W. Yi, Y. Liu, Y. Deng, A. Nallanathan, and R. W. Heath, "Modeling and analysis of mmwave v2x networks with vehicular platoon systems," *IEEE Journal on Selected Areas in Communications*, 2019.

[8] H. Qiu, F. Ahmad, F. Bai, M. Gruteser, and R. Govindan, "Avr: Augmented vehicular reality," in *Proceedings of the 16th Annual International Conference on Mobile Systems, Applications, and Services*, ser. MobiSys 2018. New York, NY, USA: Association for Computing Machinery, 2018.

[9] V. Va, T. Shimizu, G. Bansal, R. W. Heath Jr *et al.*, "Millimeter wave vehicular communications: A survey," *Foundations and Trends® in Networking*, vol. 10, no. 1, pp. 1–113, 2016.

[10] V. Va, J. Choi, T. Shimizu, G. Bansal, and R. Heath Jr, "Inverse fingerprinting for millimeter wave v2i beam alignment," *ArXiv preprint*, 2017.

[11] N. González-Prelcic, R. Méndez-Rial, and R. W. Heath, "Radar aided beam alignment in mmwave v2i communications supporting antenna diversity," in *Information Theory and Applications Workshop (ITA)*, 2016.

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

[13] D. Maamari, N. Devroye, and D. Tuninetti, "Coverage in mmwave cellular networks with base station co-operation," *IEEE transactions on Wireless Communications*, vol. 15, no. 4, pp. 2981–2994, 2016.

[14] J. Deng, O. Tirkkonen, R. Freij-Hollanti, T. Chen, and N. Nikaein, "Resource allocation and interference management for opportunistic relaying in integrated mmwave/sub-6 ghz 5g networks," *IEEE Communications Magazine*, vol. 55, no. 6, pp. 94–101, 2017.

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

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

[17] M. Giordani, M. Polesse, A. Roy, D. Castor, and M. Zorzi, "A tutorial on beam management for 3gpp nr at mmwave frequencies," *IEEE Communications Surveys & Tutorials*, vol. 21, no. 1, pp. 173–196, 2018.

[18] 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," in *IEEE INFOCOM 2018-IEEE Conference on Computer Communications*. IEEE, 2018, pp. 2375–2383.

[19] S. Sur, V. Venkateswaran, X. Zhang, and P. 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*, 2015, pp. 71–84.

[20] H. Hassanieh, O. Abari, M. Rodriguez, M. Abdelghany, D. Katabi, and P. Indyk, "Fast millimeter wave beam alignment," in *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication*, 2018.

[21] T. Wei and X. 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*, 2017, pp. 42–55.

[22] T. Nitsche, A. B. Flores, E. W. Knightly, and J. Widmer, "Steering with eyes closed: mm-wave beam steering without in-band measurement," in *2015 IEEE Conference on Computer Communications (INFOCOM)*. IEEE, 2015, pp. 2416–2424.

[23] V. Va, X. Zhang, and R. W. Heath, "Beam switching for millimeter wave communication to support high speed trains," in *2015 IEEE 82nd Vehicular Technology Conference (VTC2015-Fall)*, 2015.

[24] A. Asadi, S. Müller, G. H. Sim, A. Klein, and M. Hollick, "Fml: Fast machine learning for 5g mmwave vehicular communications," in *IEEE INFOCOM 2018-IEEE Conference on Computer Communications*. IEEE, 2018, pp. 1961–1969.

[25] 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)*, May 2017, pp. 1–6.

[26] M. E. Eltayeb, T. Y. Al-Naffouri, and R. W. Heath, "Compressive sensing for blockage detection in vehicular millimeter wave antenna arrays," in *2016 IEEE Global Communications Conference (GLOBECOM)*. IEEE, 2016, pp. 1–6.

[27] B. Antonescu, M. T. Moayyed, and S. Basagni, "mmwave channel propagation modeling for v2x communication systems," in *2017 IEEE 28th Annual International Symposium on Personal, Indoor, and Mobile Radio Communications (PIMRC)*. IEEE, 2017, pp. 1–6.

[28] C. K. Anjinappa and I. Guvenc, "Millimeter-wave v2x channels: Propagation statistics, beamforming, and blockage," in *2018 IEEE 88th Vehicular Technology Conference (VTC-Fall)*. IEEE, 2018, pp. 1–6.

[29] A. Alkhateeb, O. El Ayach, G. Leus, and R. W. Heath, "Channel estimation and hybrid precoding for millimeter wave cellular systems," *IEEE Journal of Selected Topics in Signal Processing*, vol. 8, no. 5, pp. 831–846, 2014.

[30] A. Alkhateeb, J. Mo, N. Gonzalez-Prelcic, and R. W. Heath, "Mimo precoding and combining solutions for millimeter-wave systems," *IEEE Communications Magazine*, vol. 52, no. 12, pp. 122–131, 2014.

[31] X. Gao, L. Dai, S. Han, I. Chih-Lin, and R. W. Heath, "Energy-efficient hybrid analog and digital precoding for mmwave mimo systems with large antenna arrays," *IEEE Journal on Selected Areas in Communications*, vol. 34, no. 4, pp. 998–1009, 2016.

[32] J. P. González-Coma, J. Rodríguez-Fernandez, N. González-Prelcic, L. Castedo, and R. W. Heath, "Channel estimation and hybrid precoding for frequency selective multiuser mmwave mimo systems," *IEEE Journal of Selected Topics in Signal Processing*, vol. 12, no. 2, pp. 353–367, 2018.

[33] L. He, J. Wang, and J. Song, "Spatial modulation for more spatial multiplexing: Rf-chain-limited generalized spatial modulation aided mm-wave mimo with hybrid precoding," *IEEE Transactions on Communications*, vol. 66, no. 3, pp. 986–998, 2017.

[34] T. Wei, A. Zhou, and X. Zhang, "Facilitating robust 60 ghz network deployment by sensing ambient reflectors," in *14th USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, 2017.

[35] P. Smulders, "Exploiting the 60 ghz band for local wireless multimedia access: Prospects and future directions," *IEEE communications magazine*, vol. 40, no. 1, pp. 140–147, 2002.

[36] Y. Huang, Y. Li, H. Ren, J. Lu, and W. Zhang, "Multi-panel MIMO in 5G," *IEEE Communications Magazine*, vol. 56, no. 3, 2018.

[37] "Airfide inc." <http://airfidenet.com>, 2019.

[38] F. C. Commission, "Fcc 15.255," https://www.ecfr.gov/cgi-bin/text-idx?SID=7aeef7a7d5d4a0dab75e7c6519b898a01&mc=true&node=se47.1.15_1255&rgn=div8, 2020.

[39] 3GPP, "Nr: radio resource control (rrc); protocol specification; (release 15)," TS 38.331, 2018.

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

[41] "Wiki - batman-adv - open mesh" <https://www.open-mesh.org/projects/batman-adv/wiki/Wiki>, 2020.

[42] D. Steinmetz, D. Wegemer, and M. Hollick. (2017) Talon tools: The framework for practical ieee 802.11ad research. [Online]. Available: <https://seemoo.de/talon-tools>

[43] S. Wang, J. Huang, X. Zhang, H. Kim, and S. Dey, "X-Array: Approximating Omnidirectional Millimeter-WaveCoverage Using an Array of Phased Arrays," in *ACM MobiCom*, 2020.

[44] J. Palacios, D. Steinmetz, A. Loch, M. Hollick, and J. 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*, 2018, pp. 241–255.

[45] G. A. C. (DLR), "Sumo at a glance - sumo documentation," http://sumo.sourceforge.net/userdoc/Sumo_at_a_Glance.html, 2019.

[46] Remcom, "Wireless em propagation software - wireless insite," <https://www.remcom.com/wireless-insite-em-propagation-software/>, 2020.

[47] OpenStreetMap, "Openstreetmap," <https://www.openstreetmap.org/#map=5/38.007/-95.844,2020>.

[48] V. Elistratov, "blender-osm: Openstreetmap and terrain for blender," <https://github.com/vvoovv/blender-osm>, 2020.

[49] S. Sun and T. S. Rappaport, "Antenna diversity combining and beamforming at millimeter wave frequencies," Ph.D. dissertation, Masterázs thesis, 2014.

[50] Y. Wang, K. Venugopal, A. F. Molisch, and R. W. Heath, "Analysis of urban millimeter wave microcellular networks," in *2016 IEEE 84th Vehicular Technology Conference (VTC-Fall)*. IEEE, 2016, pp. 1–5.

[51] T. Nitsche, C. Cordeiro, A. B. Flores, E. W. Knightly, E. Perahia, and J. C. Widmer, "Ieee 802.11 ad: directional 60 ghz communication for multi-gigabit-per-second wi-fi," *IEEE Communications Magazine*, vol. 52, no. 12, pp. 132–141, 2014.

[52] U. Labs, "Opencellid," <https://opencellid.org/#zoom=16&lat=37.77889&lon=-122.41942,2020>.

[53] A. Klautau, P. Batista, N. Gonzalez-Prelcic, Y. Wang, and R. 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*, 2018, pp. 1–1. [Online]. Available: http://ita.ucsd.edu/workshop/18/files/paper/paper_3313.pdf

[54] G. A. C. (DLR), "Traci - sumo documentation," <http://sumo.sourceforge.net/userdoc/TraCI.html>, 2019.

[55] D. E. Dudgeon, "Multidimensional digital signal processing," *Engewood Cliffs*, 1983.

[56] 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.

[57] 3GPP, "nr physical layer: General description."

[58] T. S. Rappaport, F. Gutierrez, E. Ben-Dor, J. N. Murdoch, Y. Qiao, and J. I. Tamir, "Broadband Millimeter-Wave Propagation Measurements and Models Using Adaptive-Beam Antennas for Outdoor Urban Cellular Communications," *IEEE Transactions on Antennas and Propagation*, vol. 61, no. 4, 2013.

[59] J. Palacios, D. De Donno, and J. Widmer, "Tracking mm-wave channel dynamics: Fast beam training strategies under mobility," in *IEEE INFOCOM 2017-IEEE Conference on Computer Communications*. IEEE, 2017, pp. 1–9.

[60] S. Sur, I. Pefkianakis, X. Zhang, and K.-H. Kim, "Towards Scalable and Ubiquitous Millimeter-Wave Wireless Networks," in *Proceedings of ACM Annual International Conference on Mobile Computing and Networking (MobiCom)*, 2018.

[61] A. Zhou, X. Zhang, and H. Ma, "Beam-Forecast: Facilitating Mobile 60 GHz Networks via Model-Driven Beam Steering," in *IEEE Conference on Computer Communications (INFOCOM)*, 2017, pp. 1–9.

[62] F. K. Gruber, E. A. Marengo, and A. J. Devaney, "Time-reversal imaging with multiple signal classification considering multiple scattering between the targets," *The Journal of the Acoustical Society of America*, vol. 115, no. 6, pp. 3042–3047, 2004.

[63] 3GPP, "Nr; radio resource control (rrc); protocol specification," <https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3197>, 2017.

[64] S. Sur, X. Zhang, P. Ramanathan, and R. Chandra, "BeamSpy: Enabling Robust 60 GHz Links Under Blockage," in *USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, 2016.

[65] S.-E. Chiu, N. Ronquillo, and T. Javidi, "Active learning and csi acquisition for mmwave initial alignment," *IEEE Journal on Selected Areas in Communications*, vol. 37, no. 11, pp. 2474–2489, 2019.

[66] 3GPP, "3gpp tr 38.901," <https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=2991>, 2017.

[67] A. Ravanshidi, P. Rost, D. S. Michalopoulos, V. V. Phan, H. Bakker, D. Aziz, S. Tayade, H. D. Schotten, S. Wong, and O. Holland, "Multi-connectivity functional architectures in 5g," in *2016 IEEE international conference on communications workshops (ICC)*. IEEE, 2016, pp. 187–192.

[68] J. L. Zambrano-Martinez, C. T. Calafate, D. Soler, and J.-C. Cano, "Towards realistic urban traffic experiments using dfrouter: Heuristic, validation and extensions," *Sensors*, vol. 17, no. 12, p. 2921, 2017.

[69] N. Akhtar, S. C. Ergen, and O. Ozkasap, "Vehicle mobility and communication channel models for realistic and efficient highway vanet simulation," *IEEE Transactions on Vehicular Technology*, vol. 64, no. 1, pp. 248–262, 2014.

[70] T. Bai and R. W. Heath, "Coverage analysis for millimeter wave cellular networks with blockage effects," in *2013 IEEE Global Conference on Signal and Information Processing*. IEEE, 2013, pp. 727–730.

[71] Q. Zhang, W. Saad, and M. Bennis, "Millimeter wave communications with an intelligent reflector: Performance optimization and distributional reinforcement learning," *arXiv preprint arXiv:2002.10572*, 2020.

[72] V. Sundarapandian, *Numerical linear algebra*. PHI Learning Pvt. Ltd., 2008.

[73] Y. Ghasempour and E. W. Knightly, "Decoupling beam steering and user selection for scaling multi-user 60 ghz wlans," in *Proceedings of the 18th ACM International Symposium on Mobile Ad Hoc Networking and Computing*, 2017, pp. 1–10.

[74] J. H. Wilkinson, F. L. Bauer, and C. Reinsch, *Linear algebra*. Springer, 2013, vol. 2.

[75] D. He, L. Wang, K. Guan, B. Ai, J. Kim, and Z. Zhong, "Channel characterization for mmwave vehicle-to-infrastructure communications in urban street environment," in *European Conference on Antennas and Propagation (EuCAP)*, 2019.

[76] A. Tassi, M. Egan, R. J. Piechocki, and A. Nix, "Modeling and design of millimeter-wave networks for highway vehicular communication," *IEEE Transactions on Vehicular Technology*, vol. 66, no. 12, 2017.

[77] K. 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.

[78] V. Semkin, U. Virk, A. Karttunen, K. Haneda, and A. V. Räisänen, "E-band propagation channel measurements in an urban street canyon," in *2015 9th European Conference on Antennas and Propagation (EuCAP)*, 2015.

[79] W. Keusgen, R. J. Weiler, M. Peter, M. Wisotzki, and B. 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)*, 2014.

[80] V. Va, T. Shimizu, G. Bansal, and R. W. Heath, "Beam design for beam switching based millimeter wave vehicle-to-infrastructure communications," in *IEEE International Conference on Communications (ICC)*, 2016.

[81] M. Giordani, M. Mezzavilla, S. Rangan, and M. Zorzi, "Multi-connectivity in 5g mmwave cellular networks," in *Mediterranean Ad Hoc Networking Workshop (Med-Hoc-Net)*. IEEE, 2016, pp. 1–7.

[82] J. Choi, V. Va, N. Gonzalez-Prelcic, R. Daniels, C. R. Bhat, and R. W. Heath, "Millimeter-wave vehicular communication to support massive automotive sensing," *IEEE Communications Magazine*, vol. 54, no. 12, 2016.

[83] Q. Chen, X. Peng, J. Yang, and F. Chin, "Spatial reuse strategy in mmwave wpans with directional antennas," in *2012 IEEE Global Communications Conference (GLOBECOM)*. IEEE, 2012, pp. 5392–5397.

[84] Y. Niu, L. Su, C. Gao, Y. Li, D. Jin, and Z. Han, "Exploiting device-to-device communications to enhance spatial reuse for popular content downloading in directional mmwave small cells," *IEEE Transactions on Vehicular Technology*, vol. 65, no. 7, pp. 5538–5550, 2015.

[85] Y. Niu, Y. Li, D. Jin, L. Su, and A. V. Vasilakos, "A survey of millimeter wave communications (mmwave) for 5g: opportunities and challenges," *Wireless networks*, vol. 21, no. 8, pp. 2657–2676, 2015.

[86] R. A. Stirling-Gallacher and M. S. Rahman, "Multi-user mimo strategies for a millimeter wave communication system using hybrid beam-forming," in *2015 IEEE International Conference on Communications (ICC)*. IEEE, 2015, pp. 2437–2443.

[87] C. Hu, J. Liu, X. Liao, Y. Liu, and J. Wang, "A novel equivalent baseband channel of hybrid beamforming in massive multiuser mimo systems," *IEEE Communications Letters*, vol. 22, no. 4, pp. 764–767, 2017.

[88] G. Lee and Y. Sung, "A new approach to user scheduling in massive multi-user mimo broadcast channels," *IEEE Transactions on Communications*, pp. 1481–1495, 2017.

[89] Y. Ghasempour, M. K. Haider, C. Cordeiro, D. Koutsonikolas, and E. Knightly, "Multi-stream beam-training for mmwave mimo networks," in *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking*, 2018, pp. 225–239.

[90] Y. Zhu, Z. Zhang, Z. Marzi, C. Nelson, U. Madhow, B. Y. Zhao, and H. Zheng, "Demystifying 60ghz outdoor picocells," in *Proceedings of the 20th annual international conference on Mobile computing and networking*, 2014, pp. 5–16.

[91] Y. Niu, C. Gao, Y. Li, D. Jin, L. Su, and D. Wu, "Boosting spatial reuse via multiple-path multihop scheduling for directional mmwave wpans," *IEEE Transactions on Vehicular Technology*, vol. 65, no. 8, pp. 6614–6627, 2015.

[92] Y. Zhu, Y. Niu, J. Li, D. O. Wu, Y. Li, and D. Jin, "Qos-aware scheduling for small cell millimeter wave mesh backhaul," in *2016 IEEE International Conference on Communications (ICC)*. IEEE, 2016, pp. 1–6.

[93] H. Vardhan, N. Thomas, S.-R. Ryu, B. Banerjee, and R. Prakash, "Wireless data center with millimeter wave network," in *2010 IEEE global telecommunications conference GLOBECOM 2010*, 2010.

[94] R. Zhao, T. Woodford, T. Wei, K. Qian, and X. Zhang, "M-Cube: A Millimeter-Wave Massive MIMO Software Radio," in *Proceedings of the ACM Annual International Conference on Mobile Computing and Networking (MobiCom)*, 2020.