

Such situations result in unpredictable wireless propagation conditions ($\sim 73\%$ drop in prediction accuracy according to our experiments) that cannot be fully characterized without new RF information.

B. Multiverse at Edge for Beamforming in Unseen Scenarios

We propose the Multiverse a software-based paradigm that runs ray tracing to emulate the RF propagation patterns, in unseen environments. This has the potential to replace—and correct—erroneous DL predictions, while still avoiding the latency of exhaustive sweep, defined by the standards. In the Multiverse, several twins, each offering a different level of fidelity and associated computation cost, coexists and ray tracing emulation is performed for all of them and each unseen scenario. The emulation cost of each twin increases with its level of fidelity. Thus, the accuracy of emulating ‘reality’ is determined by the available computational resources at the edge, such as whether having access to a CPU unit, commercial grade GPUs, or a GPU farm. On the other hand, depending on wireless latency constraints, choosing a twin with lower complexity may be favored. Our framework is apt for exploring this accuracy-latency trade-off by selecting a twin from the ‘Multiverse’ and set of Top- K beam recommendations within the selected twin, according to the user defined latency constraint. Moreover, the ray tracing outputs from the Multiverse can also generate a labeled data point for continuous learning of the DL models at the vehicle. Ultimately, the prediction of the twin must be timely (i.e., the vehicle can use its prediction before it speeds away) and accurate (i.e., close to exhaustive search). Taken together, the main challenges involve (i) detecting seen and unseen scenarios locally at the vehicle, (ii) choosing one of the twins from the Multiverse in an unseen scenario based on the latency constraint, (iii) ensuring the complete cycle of query-compute-response can be completed in less time than performing exhaustive search locally.

C. Overview of the Proposed Framework

In our proposed framework, the vehicle first detects if it encountered a seen or unseen scenario. In a seen scenario, the vehicle predicts the beam locally using DL-based method and there is no need to invoke any twin from the Multiverse. In an unseen scenario, e.g., a blockage of certain dimensions that was not included in the training data for DL-based method, we use the Multiverse for beamforming instead. Upon triggering the Multiverse, if the unseen environment is not included in the Multiverse of twins, the vehicle relays the sensor data to setup a digital replica in the ray tracing tool Wireless InSite [10] that models the geographical location of the car, dimension and placement of the obstacles, twin-specific antenna models, and multipath emulation settings. At the end of this step, lookup tables are generated and added to the Multiverse of twins to be used by upcoming vehicles later. These look up tables include information about the locations of the receiver (Rx) and signal-to-noise-ratio (SNR) of all beams obtained by ray tracing. In the next step, a decision strategy at the Multiverse identifies (a) which of the twins from the Multiverse must be used, and (b) what is the optimum subset of beams for the selected twin, according to the latency constraint. The candidate top- K beams are then swept by the BS to identify the optimum beam and start the transmission. Finally, the ray tracing outputs from the Multiverse are paired

with the local sensor data to fine-tune the DL model at the vehicle for the future.

D. Summary of Contributions

Our main contributions are as follows:

- We propose the Multiverse paradigm, where different twins emulate the real world with varying levels of fidelity. We propose an optimization algorithm that takes into account the fidelity of the twins and user-defined latency constraint to automatically select the optimum twin from the Multiverse and associated top- K beams in a *case-by-case* basis. This algorithm is executed at the Multiverse with complexity $O(N\mathcal{B})$, where N is the number of twins in the Multiverse and \mathcal{B} is the number of beams in the codebook.
- We propose to leverage the ray tracing outputs from the Multiverse to generalize the real world DL models to unseen environments, in actual deployment conditions. Thus, the labels from the Multiverse and real world sensor data are paired to fine-tune the local DL model. We observe $> 96\%$ accuracy, while fine-tuning the local DL model with a regression loss.
- We rigorously evaluate the Multiverse paradigm and how well a twin’s predictions match with the ground-truth, using a publicly available real world dataset [11] for vehicle-to-infrastructure (V2I) communication. Our results reveal that the Multiverse offers up to 79.43% and 85.22% top-10 beam selection accuracy for Line-of-Sight (LOS) and non-Line-of-Sight (NLOS) scenarios, respectively. Finally, we demonstrate that the Multiverse decreases the beam selection overhead by 67.70–90.79% compared to exhaustive search method proposed by the state-of-the-art 802.11ad and 5G-NR standards.
- We publish the *first-of-its-kind* dataset and simulation code for the Multiverse with precise maps, experimentally measured antenna patterns, and building materials using the *Wireless InSite* [10] software in [12]. We provide APIs to interface the Wireless InSite software framework with our previously released FLASH dataset [11] obtained from a sensor-equipped Lincoln MKZ autonomous car.

II. RELATED WORK

In this section, we first survey the state-of-the-art deep learning methods for unseen scenarios. We then summarize the related work on digital twins within the wireless domain, as the main focus of this paper.

A. Deep Learning Methods for Unseen Scenarios

Invariant learning refers to a class of methods developed to address distribution shifts with the goal of obtaining a generalizable predictor over different environments or scenarios [13]. These methods posit that features are drawn from multiple distributions, but the relationship between label and features is invariant across environments [14], [15]. The incentive is that if a predictor is optimal across many environments seen during training, then it will generalize better to future unseen environments [16]. However, achieving true invariance (generalizing to any arbitrary distribution) is still a challenging open problem, and the state-of-the-art methods are often studied under restrictive assumptions such as linearity of the data generative process and access to many environments in the training phase. Another research area that comes close

to our setting is out-of-distribution (OOD) learning. In the machine learning literature, the distribution refers to “label distribution”, which means that OOD samples should not have overlapping labels with respect to the training data [17]. However, in our setting the labels are the same across different scenarios and the OOD setting is not applicable. Another possible approach is transfer learning, which we were extensively explored in our prior work [6]. However, these method requires training labels from the unseen scenarios to fine-tune and adapt the models, whereas in Multiverse we do not consider such an assumption. Similarly, Meta-learning, or “learning to learn” [18], aims to extract information from a set of observed tasks that enables fast adaptation to new tasks (thus also referred as multi-task learning), with only a small amount of samples and computation available for adaptation [19]. Meta-learning approaches span a variety of topics such as learning an initialization for gradient-based learning algorithms [20], [21], [22], [23] and finding a low-dimensional subspace to fine-tune model parameters [24], [25]. However, they either target generalization to new classes not seen in the training set or require training labels from unseen scenarios [26]; thus, out of scope for this work.

B. Digital Twin for Optimization

Dong et al. [27] train a DL model for base station and user association to optimize the energy consumption. The DL model is periodically updated in the digital twin, using the information flow from the real world. Similarly, Lu et al. [28] formulates how to set up a digital twin in a mobile network via reinforcement and transfer learning in order to ensure connectivity with low computation latency. Sun et al. [29] propose a lightweight digital twin framework within a network of UAVs, which performs energy-efficient UAV placement with the objective of establishing connectivity with ground units in emergency situations. Boas et al. [30] propose to use a two-stage deep learning scheme and digital twin for channel state information (CSI) estimation. The authors use untrained neural networks (UNNs) to learn the CSI, which is then fed into a conditional generative adversarial networks (cGAN) to create the corresponding digital world. After this step, their framework only needs the node locations to generate CSI.

C. Digital Twin for Beamforming

Li et al. [31] use digital twin to generate training data for channel estimation, which is important in beamforming. The emulated 5G complaint data is used to train a deep generative model that predicts the channel using only the precoder matrix indicator (PMI). Zeulin et al. [32] use digital twin to generate training data (channel map) through a ray tracing simulator in mmWave band. They then train a model to a predict subset of candidate AoA/AoD options. The success measure of their Machine Learning (ML) model is the probability of capturing half-power beamwidth and error in AoA. Cui et al. [33] utilize digital twins for beamforming with the assistance of Reconfigurable Intelligence Surfaces (RIS). They jointly optimize user association and power allocation by running a deep reinforcement learning (DRL) at the twin. The success of their eventual goal, i.e. maximizing sum-rate of uplink in the user-centric cell-free systems, is demonstrated in simulation.

TABLE I
NOTATION SUMMARY

Notation	Description	Section
C_{Tx}	Codebook of transmitter with B beams	Sec. III-A
$p_{t_b_i}$	Received signal strength for beam $t_b \in C_{Tx}$	Sec. III-A
S_l	Seen scenario	Sec. IV
S_u	Unseen scenario	Sec. IV
$X_{l,j}^{LiD}, X_{l,j}^{Img}, X_{l,j}^{CrD}$	LiDAR, image, and GPS data for sample j and l^{th} seen scenario	Sec. IV-A
$Y_{l,j}$	RF ground-truth for sample j and l^{th} seen scenario	Sec. IV-A
$X_{l,j}^{LiD}, X_{l,j}^{Img}, X_{l,j}^{CrD}$	Test samples from LiDAR, image and coordinate for sample j and l^{th} seen scenario	Sec. IV-A
n_l	Number of samples in l^{th} seen scenario	Sec. IV-A
f_θ	Learning model	Sec. IV-A
$\mathcal{L}(\theta; S_l)$	Loss function for seen scenario S_l	Sec. IV-A
$t_{l,j}^*$	Best beam for sample j and l^{th} seen scenario	Sec. IV-A
\mathcal{M}_u	Multiverse for unseen scenario u	Sec. IV-B
$\mathcal{T}_{u,i}$	i^{th} twin for unseen scenario u	Sec. IV-B
map_u	Imported OpenStreetMap for unseen scenario u	Sec. IV-B1
O_u	Present structures or obstacles for unseen scenario u	Sec. IV-B1
$C_{Tx}^{T_{u,i}}$	Transmitter codebook for twin $\mathcal{T}_{u,i}$ ($ C_{Tx}^{T_{u,i}} = B_i$)	Sec. IV-B1
ρ_i	Number of allowed reflections for i^{th} twin	Sec. IV-B1
N_u^{obs}	Number of obstacles in unseen scenario S_u	Sec. IV-B1
$L_{\mathcal{T}_{u,i}}$	Set of location to perform ray tracing for twin $\mathcal{T}_{u,i}$ ($ L_{\mathcal{T}_{u,i}} = n_{\mathcal{T}_{u,i}}$)	Sec. IV-B1
l_j	A single point from $L_{\mathcal{T}_{u,i}}$	Sec. IV-B1
$N_j^{m,i}$	Number of total propagation rays delivered at Rx at location l_j for the Tx beam t_m and twin $\mathcal{T}_{u,i}$	Sec. IV-B1
$P_j^{m,i}$	Power at location l_j for Tx beam t_m in twin $\mathcal{T}_{u,i}$	Sec. IV-B1
\mathcal{N}	Noise power	Sec. IV-B1
$SNR_j^{m,i}$	SNR at location l_j for Tx beam t_m in twin $\mathcal{T}_{u,i}$	Sec. IV-B1
$C_{Comp}^{\mathcal{T}_{u,i}}$	Computation cost of twin $\mathcal{T}_{u,i}$ including cost of generating map ($C_{map}^{\mathcal{T}_{u,i}}$) and look up tables ($C_{lookup}^{\mathcal{T}_{u,i}}$)	Sec. IV-B1
t	Number of transmissions through walls	Sec. IV-B1
ζ	Number of allowed diffraction	Sec. IV-B1
w	Number of reflective surfaces reached by the rays emitted from the Tx	Sec. IV-B1
$LT(\mathcal{T}_{u,i})$	Look up table for twin $\mathcal{T}_{u,i}$	Sec. IV-B1
$C^{5G-NR}(K)$	Communication cost in 5G-NR for K beams	Sec. IV-B2
$p(\cdot)$	Probability of inclusion	Sec. IV-B2
$\gamma(\cdot)$	Outlier detection algorithm	Sec. IV-D

Novelty of the Multiverse: In summary, the state-of-the-art work concerning digital twin in wireless applications is based on simulation alone. Also, to the best of our knowledge, current wireless digital twin models are applicable for stationary or pedestrian-speed mobility. Our work distinguishes itself over prior work by: (a) using real data for validation, (b) including vehicular-mobility components, (c) offering a Multiverse of twins to choose from, (d) proposing a pioneering work (to the best of our knowledge) on interactive digital twin based beam selection in the mmWave band. Also, unlike prior work, we pledge to release datasets and software APIs for rigorous validation by other researchers.

III. SYSTEM ARCHITECTURE

In this section, we first review classical beam initialization and formally present the beam selection problem. We then introduce the system architecture in our framework that exploits the DL-based method for seen and Multiverse for unseen scenarios (see Fig. 1). We summarize the notations in Tab. I.

A. Traditional Beam Initialization

The 802.11ad and 5G-NR standards use exhaustive search for beamforming wherein the transmitter (Tx) sends out probe frames in all beams sequentially and the receiver (Rx) listens to these frames with a quasi-omni-directional antenna setting [34]. This process is then repeated with the Tx-Rx roles reversed. Assume the Tx has a pre-defined codebook $C^{Tx} = \{t_1, \dots, t_B\}$, consisting of B elements. After transmitting B

probe frames, the beam with the maximum signal strength is chosen as optimal. Formally, the optimal beam at Tx is calculated as:

$$t^* = \arg \max_{1 \leq m \leq B} p_{t_m}, \quad (1)$$

where p_{t_m} is the received signal strength at the Rx, when the transmitter uses beam t_m . The exhaustive search methods are slow and particularly impractical in a V2X scenarios, where the optimal beam changes rapidly, due to mobility.

B. DL-Based Prediction at Vehicle and Multiverse at Edge

The RF propagation pattern in the mmWave band is affected by factors such as the positioning of the transmitter and receiver, atmospheric features, and presence of obstacles specially the ones that block the LOS. As a result, in dynamic scenarios such as vehicular networks, the RF profiles continuously change over time which makes the DL-based methods prone to faulty predictions in unseen scenarios. Thus, we propose an interactive Multiverse-based solution, where the vehicle: (a) recognizes the unseen scenarios locally; (b) invokes a twin in the Multiverse for beam selection; (c) uses emulation outputs for fine-tuning the local DL models. Our proposed frameworks runs by coordination between the DL-based beam selection and Multiverse and has three main thrusts as follows:

- **Beam Selection using DL-based framework:** While encountering a seen scenario, we use the DL-based beam selection that exploits the local sensor data at the vehicle to predict the optimum beam (see Sec. IV-A).
- **Beam Selection using Multiverse at Edge:** We use the Multiverse including twins with different levels of fidelity and associated computation cost at the edge, for unseen scenarios. We first create the Multiverse offline by emulating the propagation patterns using a ray tracing tool, where we consider different number of reflections and custom beam patterns. After this offline creation step, a lookup table is generated for each twin that contains the SNR of the beams. The vehicles use these look up tables to choose the optimum twin from the Multiverse and top- K beams locally, while considering the communication latency constraint (see Sec. IV-B).
- **Fine-Tuning with Multiverse Ground-Truth:** We incorporate the emulated beams from the Multiverse to label the sensor data at the vehicle. Thus, we fine-tune the local model to account for previously unseen scenarios in the future (see Sec. IV-C).

IV. MULTIVERSE-BASED BEAM SELECTION

In this section, we introduce our mmWave beam selection framework, consisting of the DL-based prediction at the vehicle and Multiverse at edge. We denote the set of seen (i.e. labeled) and unseen (i.e. unlabeled) scenarios by $\{\mathcal{S}_l\}_{l=1}^L$ and $\{\mathcal{S}_u\}_{u=1}^U$, respectively. Here, \mathcal{S}_l and \mathcal{S}_u are samples from a total of L and U seen and unseen scenarios, respectively. When a vehicle comes within the range of Multiverse-enabled mmWave BS, it (a) uses a local trained DL model that utilizes multimodal sensor data for seen scenario \mathcal{S}_l , and (b) invokes the Multiverse for unseen scenario \mathcal{S}_u . Fig. 2 shows the interactions of real world and Multiverse in our proposed framework.

A. Beam Selection Using DL-Based Framework

To avoid the costly exhaustive search in traditional beam selection, one approach is using DL models that predict the

best beam using non-RF multimodal data, such as LiDAR, camera images, and GPS coordinates [7], [8], [35]. In this method, a training set is available prior to deployment for seen environments $\{\mathcal{S}_l\}_{l=1}^L$. The training data corresponding to the l^{th} seen scenario includes the multimodal sensor data as well as the RF ground-truth $\{(X_{l,j}^{\text{LiD}}, X_{l,j}^{\text{img}}, X_{l,j}^{\text{CrD}}), Y_{l,j}\}_{j=1}^{n_l}$. Here, $X_{l,j}^{\text{LiD}}, X_{l,j}^{\text{img}}, X_{l,j}^{\text{CrD}}$ are LiDAR, image, and GPS coordinate samples and $Y_{l,j} \in \{0, 1\}^B$ is the corresponding label for the sample j and l^{th} seen scenario. Moreover, n_l is the total of training samples for seen scenario \mathcal{S}_l .

The learning model $f_\theta(\cdot)$ is a function parameterized by θ , i.e., a neural network with weights θ . The empirical loss of the model parameters θ on dataset for seen scenario \mathcal{S}_l is defined as $\mathcal{L}(\theta; \mathcal{S}_l) := \frac{1}{n_l} \sum_{j=1}^{n_l} [\ell(f_\theta(X_{l,j}^{\text{LiD}}, X_{l,j}^{\text{img}}, X_{l,j}^{\text{CrD}}), Y_{l,j})]$, where ℓ is a loss function measuring the discrepancy between predicted and true labels, cross entropy as an instance. The standard DL training approach finds a model that minimizes the loss across all of the training samples by solving: $\min_{\theta} \frac{1}{N_l} \sum_{l=1}^L n_l \mathcal{L}(\theta, \mathcal{S}_l)$, with $N_l = \sum_{l=1}^L n_l$. After the model training, the best beam is predicted as:

$$t_{l,j}^* = \sigma(f_\theta(\mathbb{X}_{l,j}^{\text{LiD}}, \mathbb{X}_{l,j}^{\text{img}}, \mathbb{X}_{l,j}^{\text{CrD}})), \quad (2)$$

where σ denotes the *softmax* activation, $\mathbb{X}_{l,j}^{\text{LiD}}, \mathbb{X}_{l,j}^{\text{img}}$, and $\mathbb{X}_{l,j}^{\text{CrD}}$ are test samples from LiDAR, image and coordinate sensors for sample j and l^{th} seen scenario, respectively.

B. Beam Selection Using Multiverse at Edge

The proposed DL-based beam selection method in Sec. IV-A exploits the sensor data to predict the optimum beam locally at the vehicles. However, it fails to provide accurate prediction for unseen scenarios $\{\mathcal{S}_u\}_{u=1}^U$. In the case of an unseen scenarios, the Multiverse is triggered to emulate the beam profiles instead, which operates at the edge. In the Multiverse, a collection of N digital twins are available with different levels of fidelity: $\mathcal{M}_u = \{\mathcal{T}_{u,i}\}_{i=1}^N$ for each unseen scenario \mathcal{S}_u . After running the ray tracing emulation, lookup tables are generated for each twin in the Multiverse, with per-beam SNRs over different locations on the road. In this section, we explain how we create the Multiverse of twins at the edge and present our strategy to determine the optimum twin from the Multiverse and set of top- K beams according to the communication latency constraint.

1) *Offline Multiverse Creation:* The Multiverse creation for each unseen scenario \mathcal{S}_u consists of three steps: (a) twin-world creation (map geometry), (b) modeling the propagation paths using a ray tracing tool, (c) creation of lookup table. Having a high fidelity twin of the real world RF propagation patterns is complex and time intensive. Thus, we assume that creating the twins is performed offline. Nevertheless, with enough computation resources, it can be completed in near real-time, as studies in this direction are under development [36].

Twin-world Creation. In the Multiverse, we consider properties with respect to map precision, transmitter codebook, and mmWave propagation properties. While we consider several key metrics in the Multiverse, our baseline can be extended in the future to twins that also incorporate weather patterns, such as the effect of rain/snow. We initialize the i^{th} twin for unseen scenario \mathcal{S}_u as $\mathcal{T}_{u,i} = f_{\text{win}}(\text{map}_u, O_u, C_i^{\text{Tx}}, \rho_i)$ in the Multiverse. Here, map_u and O_u denote the imported OpenStreetMap [37] and present structures or obstacles for unseen scenario \mathcal{S}_u . Moreover, C_i^{Tx} and ρ_i are transmitter

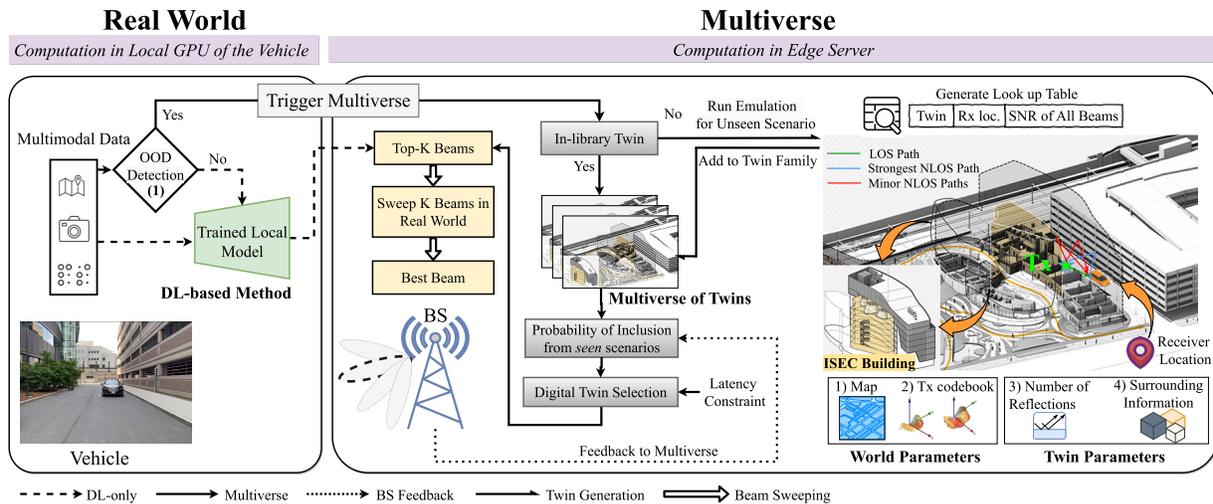


Fig. 2. The system model of the proposed Multiverse-based beam selection framework. In the offline Multiverse creation step, we run ray tracing emulation to obtain the beam profiles for each twin. We report the ray tracing output as a look up table including the, Rx location index, Twin index, and SNR values for all beams of twin $\mathcal{T}_{u,i}$ over all locations $l_j \in L_{\mathcal{T}_{u,i}}$. In the real-time Multiverse interaction, the Multiverse solves Eq. 7 to identify the optimum twin and subset of K beams from the selected twin.

codebook and number of allowed reflections for i^{th} twin. In our design, each twin can have a different codebook with \mathcal{B}_i beams for i^{th} twin ($|\mathcal{C}_i^{\text{Tx}}| = \mathcal{B}_i$). We characterize present objects in the twin with $O_u = \{(x_k, y_k, d_{0,k}, d_{1,k}, d_{2,k})\}_{k=1}^{N_u^{\text{obs}}}$, where N_u^{obs} is the number of structures or obstacles that are present in the unseen scenario \mathcal{S}_u and are replicated in the Multiverse. Moreover, (x_k, y_k) are the obstacle coordinates on the map, whereas $d_{0,k}, d_{1,k}, d_{2,k}$ represent the height, width, and length of the obstacle- k .

Modeling the Propagation Paths. To model the propagation paths in twin $\mathcal{T}_{u,i}$, we place the receiver at $L_{\mathcal{T}_{u,i}} = \{l_j\}_{j=1}^{n_{\mathcal{T}_{u,i}}}$ different locations on the road and send out mmWave rays from the BS, where the Tx is located, with ρ_i reflections for each location. We then perform the propagation study for each beam $t_m \in \mathcal{C}_i^{\text{Tx}}$ at each $l_j \in L_{\mathcal{T}_{u,i}}$ location.

We follow X3D [10] for modeling the mmWave propagation path that takes into account the phase information of the rays. We calculate the total received power, i.e. aggregated power from individual rays at the Rx, for the Tx beam t_m , twin $\mathcal{T}_{u,i}$, and location l_j as:

$$P_j^{m,i} = \frac{\lambda^2 \beta}{8\pi\eta_0} \left| \sum_{n=1}^{N_j^{m,i}} [E_{\theta,n} g_{\theta}(\theta_n, \phi_n) + E_{\phi,n} g_{\phi}(\theta_n, \phi_n)] \right|^2, \quad (3)$$

where $N_j^{m,i}$ is the number of total propagation rays delivered at Rx at location l_j for the Tx beam t_m and twin $\mathcal{T}_{u,i}$. Here, λ is the wavelength, and β is the ratio of the area under the overlapping frequency spectrum of the transmitted and received signals to the area under the frequency spectrum of the transmitted signal [38]. Moreover, Eq. 3, η_0 is the impedance of free space (where $\eta_0 = 377\Omega$), θ_n and ϕ_n represent direction of arrivals, $E_{\theta,n}$ and $E_{\phi,n}$ are the theta and phi components of the electric field of the n^{th} ray at the Rx. Further, $g_{\phi}(\theta, \phi) = \sqrt{|G_{\theta}(\theta, \phi)|} e^{j\psi_{\theta}}$, where G_{θ} is the theta component of the receiver antenna gain, and ψ_{θ} is the relative phase of the θ component of the far zone electric field. From Eq. 3, one can observe that the received power is square-inversely proportional to the carrier signal frequency,

f_c , given that $\lambda = c/f_c$, c being the speed of light and the exponent 2 is the path loss exponent for the X3D model, which is compliant with the latest mmWave measurements in urban canyon scenarios [39]. This means that the received power, as well as the range of the signal, depends on the operating frequency of the communication system, i.e. higher the f_c faster the attenuation is, which translates into smaller range and less received power. Nevertheless, high f_c signals bring a high data rate opportunity, due to the high signal bandwidth, introducing a trade-off for mobile operators, who can decide based on their business strategies. There are also studies with WI that involve higher [40] and lower frequencies [41] than our operating frequency, 60 GHz. Molecular absorption, which is defined as atmospheric absorption in WI, is another parameter that affects the signal propagation, hence the received power. We used the default atmospheric absorption setting in WI, adapted from the model in the National Telecommunications and Information Administration (NTIA) report [42], because WI and beam decisions show high accuracy as we detail in Sec. VI.

Once the received power is determined (in Watts), the power in dBm is calculated using $p_j^{m,i}(\text{dBm}) = 10 \log_{10}[p_j^{m,i}(\text{W})] - L_s(\text{dBm})$, where L_s is any loss in the system other than path loss, such as within cables and electronics at the Rx. We calculate the SNR values for the beam $t_m \in \mathcal{C}_i^{\text{Tx}}$ at location $l_j \in L_{\mathcal{T}_{u,i}}$ as $\text{SNR}_j^{m,i} = P_j^{m,i} - \mathcal{N}$, where \mathcal{N} is the noise power in dBm at location l_j for twin $\mathcal{T}_{u,i}$. On the other hand, generating a twin $\mathcal{T}_{u,i}$ has an associated computation cost that is the sum of: (a) cost of generating the twin-world and (b) cost of the generating propagation rays over locations on the road and all beams for creating the lookup table. Thus, the total cost for creating a twin is:

$$\mathcal{C}_{\text{Comp}}^{\mathcal{T}_{u,i}} = \mathcal{C}_{\text{map}}^{\mathcal{T}_{u,i}} + \mathcal{C}_{\text{lookup}}^{\mathcal{T}_{u,i}}, \quad (4)$$

where $\mathcal{C}_{\text{map}}^{\mathcal{T}_{u,i}}$ is the cost to generate/import the maps for creating the world, and $\mathcal{C}_{\text{lookup}}^{\mathcal{T}_{u,i}}$ is the cost of running ray tracing emulation to generate the lookup table, which is further broken down as $\mathcal{C}_{\text{lookup}}^{\mathcal{T}_{u,i}} = \mathcal{B}_i \times n_{\mathcal{T}_{u,i}} \times \mathcal{C}_{\text{prop}}(\rho_i)$. Here, $\mathcal{C}_{\text{prop}}(\rho_i)$ is the cost for calculating one propagation ray with ρ_i reflections that depends on the available compute resource. However, following the Wireless InSite manual [38], the cost

for propagation path calculation depends on four emulation parameters as:

$$C_{\text{prop}}(\rho_i) \propto \frac{(\rho_i + \iota + 1)!}{\iota! \rho_i!} \times w^{\zeta+1}, \quad (5)$$

where ρ_i denotes the number of reflections for twin $\mathcal{T}_{u,i}$. Moreover, and ι and ζ are the number of transmissions (penetration) through the walls and diffractions, respectively. Finally, w is the number of reflective surfaces (walls for example) that are reached by the rays emitted from the Tx and completely depends on the map geometry.

Creation of Lookup Table. After modeling all the propagation paths for different Rx location in each twin, we create a lookup table, which is stored at the Multiverse and consists of four parameters: (a) Rx location index $l_j \in L_{\mathcal{T}_{u,i}}$, (b) Twin index $\mathcal{T}_{u,i}$, and (c) SNR values $\text{SNR}_j^{m,i}$ for all beams $t_m \in C_i^{\text{Tx}}$ of twin $\mathcal{T}_{u,i}$ over all locations $l_j \in L_{\mathcal{T}_{u,i}}$, hence represented as a dictionary: $LT(\mathcal{T}_{u,i}) = \{(l_j, \mathcal{T}_{u,i}, \text{SNR}_j^{m,i})\}_{j=1}^{n_{\mathcal{T}_{u,i}}}$. In Multiverse, it is possible to run ray tracing emulations and generate the look up tables for different scenarios in advance. In this case, the Multiverse will have the look up tables ready, when a new vehicle encountered the same scenario. However, in this paper, we present Multiverse for mmWave vehicular networks, which are highly dynamic. Hence, it is not feasible to predict any possible combination of the Rx, Tx, and obstacles in advance. Thus, in Multiverse, we detect the changes in the environment, run the ray tracing emulation, and generate the look up tables as they occur in the environment, to account for unpredictability of the inherently dynamic vehicular environments.

2) *Real-Time Multiverse Interaction:* In Multiverse, multiple twins each with a different level of fidelity are available ($\mathcal{M}_u = \{\mathcal{T}_{u,i}\}_{i=1}^N$). Whenever an unseen scenario is encountered, the BS has N twins from which it can select the top- K beams for sweeping. The selection of top- K beams depends on the selected twin. For example, a low fidelity twin may require sweeping 10 beams, while a twin with higher fidelity may require only 4 beams, at a certain location. Thus, there is dependency between the selected twin and set of top- K beams. Moreover, our observations indicate that the fidelity of the twins depends on the location on the road in which the receiver is located. For example, in a specific region, a less complex twin might be superior to a twin with higher number of reflections. That is because the emulation parameters such as number of reflections are constant for each twin in Multiverse, and unnecessary constructive and destructive interference at the Rx, with a more complex twin, may result in discrepancy between the emulation and the real world. Thus, depending on the scenario and region on the road, a less complex twin might be closer to the real world. Finally, the permissible latency for each vehicle depends on its targeted application. For example a vehicle may need critical safety directions that must be sent with minimum latency, compared to sending media data that is not time sensitive. In this case, there is an upper bound on the number of allowed beams, which effects selecting the optimum twin and set of top- K beams. In this section, we present an optimization problem that jointly selects the twin and set of top- K beams for each region on the road, while a) providing the highest fidelity b) imposing the minimum communication latency c) accounting for the communication latency requirement.

Modeling Fidelity (probability of inclusion). A simple way to model the fidelity is by benchmarking the performance of each twin with respect to the ground-truth from the so far seen scenarios. We define the probability of inclusion as the probability of the best beam from the ground-truth being in the top- K predictions from the Multiverse. Our observations indicate that the probability of inclusion for each twin is related to the sub-regions in the coverage area of the transmitter. Moreover, it is affected by the number of allowed beams K . To that end, we leverage the empirical observations to model the probability of inclusion as $p(K, LT(\mathcal{T}_{l,i}), \mathcal{S}_l, r)$, where parameters K , $LT(\mathcal{T}_{l,i})$, \mathcal{S}_l , and r denote that number of selected beams, set of lookup table entries for twin $\mathcal{T}_{l,i}$, seen scenarios \mathcal{S}_l , and the region in which the receiver is located, respectively. In other words, we use the *seen scenarios* to assess the fidelity of the twins and generate the probability of inclusion, $p(K, LT(\mathcal{T}_{l,i}), \mathcal{S}_l, r)$ at the Multiverse. We then use the generated probability of inclusion parameter to select the optimum twin and set of top- K beams in an *unseen scenarios*. For example, consider a setting, where we have the LOS and NLOS scenarios as seen and unseen scenarios, respectively. Having the LOS scenario as *seen* indicates that the ground-truth labels are already available for LOS scenario. In fact, these ground-truth labels are the ones that are used to train the DL model, as described in Sec. IV-A. On the other hand, we run the ray tracing emulation in the Multiverse for the seen LOS scenario. Given the ground-truth and emulation outputs for LOS scenario, we compute the probability of inclusion, $p(K, LT(\mathcal{T}_{l,i}), \mathcal{S}_l, r)$, for set of K beams ($1 \leq K \leq \mathcal{B}$), different twins ($\{\mathcal{T}_{l,i}\}_{i=1}^N$), and each region of the road (r). As a result, the dimensionality of the parameter probability of inclusion equals to $\mathcal{B} \times N \times n_{\mathcal{T}_{u,i}}$, which is available at the Multiverse.

Modeling Communication Cost. We use the state-of-the-art 5G-NR standard [43] to model the communication cost for sweeping the top- K beams as:

$$C^{5\text{G-NR}}(K) = T_p \times \left\lceil \frac{K-1}{32} \right\rceil + T_{ssb}, \quad (6)$$

In 5G-NR standard, synchronization signal (SS) blocks are sent in each beam and multiple SS blocks from different beams are grouped into one SS burst. The NR standard defines that the SS burst duration (T_{ssb}) is fixed to 5ms, which is transmitted with a periodicity (T_p) of 20ms [43]. Thus, within each SS burst the communication time increases linearly, and there is sudden 20ms increase when the number of beams exceeds the 32 beams that are allowed in one SS burst.

Optimization. Upon encountering an unseen scenario, we solve the following optimization problem at the Multiverse to obtain the optimum twin and set of top- K beams as:

$$\max_{K,i} p(K, LT(\mathcal{T}_{l,i}), \mathcal{S}_l, r) + \frac{\alpha}{2} \left(1 - \frac{C^{5\text{G-NR}}(K)}{C_{\text{Comm}}^{\text{Max}}}\right), \quad (7a)$$

$$\text{s.t. } C^{5\text{G-NR}}(K) < C_{\text{Comm}}, \quad (7b)$$

$$r^b \leq r \leq r^e, \quad (7c)$$

$$\mathcal{S}_l \in \{\mathcal{S}_l\}_{l=1}^L, \quad (7d)$$

$$\alpha > 0. \quad (7e)$$

In Eq. 7, the first term in objective forces the optimization problem to choose twins with higher fidelity (probability of inclusion). As opposed to this, the second term prevents the selection of unnecessarily high K values for time sensitive

applications, by modeling the communication cost to have an inverse relationship in the objective. Thus, maximizing the objective in Eq. 7 provides the solution that results in maximizing the probability of inclusion, while minimizing the communication cost. Moreover, we note that the maximum value for the probability of inclusion is equal to one. To ensure that the two objective elements are within the same range, we normalize the communication cost by its maximum (C_{Comm}^{Max}), which corresponds to sweeping all \mathcal{B} beams within the codebook. The control parameter α in Eq. 7 weights the importance between these two terms in the objective function. Moreover, the parameter \mathcal{C}_{Comm} is the user defined communication constraints at the vehicle, which applies an upper bound on the number of allowed beams for time sensitive applications.

To solve Eq. 7, the Multiverse requires four key information including, the probability of inclusion, communication latency constraint, and approximate location of the receiver. The probability of inclusion is computed at the Multiverse from the seen scenarios as previously described in this section. Moreover, the communication latency constraint is also known, since we are considering a downlink channel from the BS (where the Multiverse is located) to the vehicles. The approximate location of the receiver can also be estimated by the BS considering the speed of the vehicle or it can be sent in the uplink when triggering the Multiverse. The computational complexity for running this optimization is in the order of $O(N\mathcal{B})$ with N and \mathcal{B} denoting the number of twins and beams in the codebook. Whenever an unseen scenario is encountered, the vehicle triggers the Multiverse. The Multiverse solves Eq. 7 as the vehicle enters different regions of the road, or there is a change in the latency requirement.

After solving Eq. 7, the optimum twin and the associated top- K beams are determined. Thus, the Multiverse retrieves the set of top- K beams within the selected twin. The BS then sweeps through those suggested top- K beams in the real world, obtains the optimum beam, and starts the transmission. After identifying the optimal beam, the BS can send a feedback (see Fig. 2) to the Multiverse. In Multiverse, we envision the probability of inclusion as a cumulative metric that is computed over seen scenarios (up until the present in time) and continuously updated after receiving the feedback of the BS on the suggested top- K beams for unseen scenarios. In this regard, we first run the ray tracing emulation for the unseen scenarios and generate the look up tables. Given the feedback from the BS and the look up tables for the unseen scenario, we update probability of inclusion. Note that this operation is performed after solving Eq. 7 and returning the top- K beams to the BS. Thus, running the emulation for the unseen scenario does not halt the interaction between the real world and Multiverse. Moreover, the BS feedback can be stored at the Multiverse to update the probability of inclusion once the emulation outputs are ready. Overall, we use the ray tracing output for the unseen scenarios to i) update the probability of inclusion given the feedback from the BS ii) fine-tune the local DL model as we will discuss in Sec. IV-C.

C. Fine-Tuning With Multiverse Ground-Truth

When the vehicle uses its local DL models for beam selection, the local multimodal sensor data is used for inference (see Sec. IV-A). This results in near-real time prediction of the beam, due to locality of all actions. However, pure machine learning-based methods suffer from the adaptation to unseen environments, which undermines the accuracy of the model prediction. We propose to incorporate the emulation output

from the Multiverse to fine-tune the local DL model at the vehicle. Once a twin within the Multiverse is invoked to obtain the optimal beam, the ray tracing based emulation outputs become the ground-truth labels and are communicated back to the real world. We propose to pair these labels with the local sensor data at the vehicle to fine-tune the local model, periodically, few epochs of training for example. As a result, the local model can account for the previously unseen environment in the future, instead of continuous triggering of twins from the Multiverse. Following the notations of Sec. IV-A, the fine-tuned model parameters $\hat{\theta}$ is obtained by minimizing the loss $\mathcal{L}(\hat{\theta}; \mathcal{S}_u)$ on the unseen scenario \mathcal{S}_u , where $\mathcal{L}(\hat{\theta}; \mathcal{S}_u) := \frac{1}{n_u} \sum_{j=1}^{n_u} [\ell(f_{\hat{\theta}}(X_{u,j}^{LiD}, X_{u,j}^{Img}, X_{u,j}^{CrD}), Y_{u,j})]$. Here, $Y_{u,j}$ is the labels derived from twin $\mathcal{T}_{u,i}$ for unseen scenario \mathcal{S}_u ($|Y_{u,j}| = n_u$) and $X_{u,j}^{LiD}, X_{u,j}^{Img}, X_{u,j}^{CrD}$ are the real sensor data recorded at the vehicle. In our system model, Multiverse suggests a subset of top- K beams for sweeping to the BS. The BS sweeps through these top- K beams to identify the best one. In this case, the BS can operate as a safety mechanism and monitor the performance of the suggested beam by the Multiverse. For example, if the suggested beams from the Multiverse result in Gbps throughput, it is likely that the selected beam is the optimal one. However, if the throughput is low, it indicates that the emulations are not close to the real world. In this case, we can discard the emulation output from the Multiverse, or modify the emulation setting to obtain a twin with higher fidelity.

Training the DL model requires the sensor data as well as the labels. In our current vision of Multiverse, the labels are available after performing the ray tracing, but not the sensor data. Thus, we proposed to fine-tune the DL model at the vehicles by pairing the real world sensor data with the labels from the Multiverse. However, in a future work, one can run emulation to generate synthetic sensor data along with the ray tracing to train the DL models and deploy them at the vehicles for various types of scenarios in advance.

D. End-to-End System

Given a test sample, we first identify if it is either in-distribution or out-of-distribution, corresponding to seen and unseen scenarios, respectively. There are state-of-the-art works where kernel density-based [44], nearest neighbour-based [45], [46], or reconstruction-based [47] methods are proposed for identifying unseen scenarios. For example, one of our previous works [48] focuses on using statistical metrics recorded on the training data to detect unreliable predictions. This includes metrics that characterize the prediction probability confidence and ratio of correctly classified samples for each class, beams in our case. These statistics are generated at the training phase and are used as a threshold at inference to identify the outliers. We denote the outlier detection operation by $\gamma(\cdot)$.

Consider a test sample with $\mathbb{X}_m^{LiD}, \mathbb{X}_m^{Img}, \mathbb{X}_m^{CrD}$ as the LiDAR, image and coordinate samples, respectively. If the sample is detected as in-distribution, we use the DL-based method to locally predict the optimum beam using the sensor data. If the sample is out-of-distribution, the vehicle submits a request to get access to the Multiverse of twins and solves Eq. 7 to select the best twin and top- K beams. Formally,

$$t^* = \begin{cases} \sigma(f_{\hat{\theta}}(\mathbb{X}_{l,m}^{LiD}, \mathbb{X}_{l,m}^{Img}, \mathbb{X}_{l,m}^{CrD})) & \text{if } \gamma(\mathbb{X}_{l,m}^{\{LiD,Img,Crd\}}) = D_{ID} \\ \text{Solve Eq. 7} & \text{if } \gamma(\mathbb{X}_{l,m}^{\{LiD,Img,Crd\}}) = D_{OOD} \end{cases} \quad (8)$$

where σ denotes *softmax* activation, D_{ID} implies that the data sample belongs to the seen scenarios (in-distribution) and D_{OOD} denotes that $\gamma(\cdot)$ has identified that the sample is out-of-distribution (OOD), i.e. unseen scenario. For out-of-distribution samples, our framework revolves around (i) creating a Multiverse in the edge server of base station, (ii) solving Eq. 7 to obtain the optimum twin and set of top- K beams (see Fig. 2), and (iii) updating the probability of inclusion based on the feedback from the BS.

The *trigger Multiverse* module calls for the Multiverse \mathcal{M} from the edge, following Eq. 8. The edge has a few in-library twins within the Multiverse. This includes the twins that are encountered before and are already available at the edge. If the requested twin from the vehicle is already available at the edge, Multiverse directly moves to solving Eq. 7 and reports the selected top- K beams within the selected twin to the BS. If the requested twin is not available, the sensor data of the vehicle is sent in uplink to the edge, where the ray tracing tool runs to generate a look up table for the unseen environment. The output look up table is added to twin family in the Multiverse. Note that if the requested twin is already available at the edge, the Multiverse provides an output in competitive time compared to exhaustive search. If not, the computationally extensive ray tracing must be run at edge which might not provide an output in short contact times in vehicular network. However, the simulation output can be later used to account for the new vehicles, as they encounter the same scenario.

V. EXPERIMENT DESIGN

We consider the real world to be a V2I setting in an urban canyon of a metro city. We deploy the Multiverse using a ray tracing simulation tool, Wireless InSite (WI) by RemCom [10]. WI is capable of modeling complex electromagnetic propagation patterns in 3-dimensional (3D) urban, indoor, rural, and mixed path environments.

A. Real World: Creation of Ground-Truth

We use a publicly available mmWave beam selection dataset called FLASH [11] to generate the ground-truth from the real world. This dataset is captured within downtown Boston on an asphalt road flanked with high rising buildings (see Fig. 3a). The dataset includes synchronized sensor data from on-board GPS, a GoPro Hero4 camera, and a Velodyne LiDAR, all mounted on a 2017 Lincoln MKZ Hybrid autonomous car (see Fig. 3b). The ground-truth includes the received signal strength indicator (RSSI), which is acquired using the Talon AD7200 mmWave radio [49], with a pre-defined codebook of 34 beams. The dataset includes both LOS and NLOS scenarios with approximately 5K and 3K samples each, respectively. The LOS scenario corresponds to the vehicle (i.e. receiver) passing in front of the transmitter without any obstacle, while in NLOS scenario, a vehicle is located in front of the Tx, blocking LOS.

B. Multiverse Setup

In this section, we describe our approach for creating the scenarios and setting the RF parameters. The modular steps to create different twins is depicted in Fig. 4.

1) *Creating Scenarios in Wireless InSite*: The Wireless InSite ray tracing tool allows users to create a digital environment with 3D maps. We replicate the testbed of FLASH dataset in WI, using the closest materials and positioning for

the objects. A sample map visual for the digital world is provided in Fig. 3c.

Import the Map. In WI, a digital world can have various features, which are stored as layers, e.g. buildings, roads, water bodies, and foliage. Buildings are modeled as boxes that can have different material features, e.g. concrete, glass, metal, providing near-realistic ray reflection and penetration. The first step in creating a scenario is to obtain an appropriate map. OpenStreetMaps [37] provides such data, which we process in Blender [50] as geodata for WI. We confirm that the road and valley width as well as building dimensions in the digital world match with the ones in the real world.

Modeling and Placement of Obstacles. In WI, the dimension and coordinate of obstacles can be set, as desired. While a LOS scenario is ready for simulation runs after setting up the map and adding all the aforementioned details, there is one more step to finalize a NLOS scenario, modeling the obstacles. Following the FLASH testbed, where the obstacle is a 2018 Toyota Camry with the dimensions of $1.44\text{ m} \times 1.84\text{ m} \times 4.88\text{ m}$, height, width, and length, respectively, we model the obstacle in WI as a metal box with the same dimensions and place it 35 cm away from the Tx, shown as white box in Figs. 3d.

Placement of Tx and Rx. Following the FLASH testbed [8], the Tx antenna is placed on a cart with the height of 75 cm . Moreover, the height of Tx itself is 20 cm . Thus, we set the Tx height to be 95 cm in WI. The receiver car in the FLASH dataset has a height of 147.5 cm . Thus, we set the Rx height to be 167.5 cm in the simulation environment. In each simulation, we have one Tx location and one or more Rx locations (denoted as $L_{\mathcal{T}_{u,i}}$ for $\mathcal{T}_{u,i}$). We consider two set of scenarios, including LOS (\mathcal{S}_1) and NLOS (\mathcal{S}_2)

2) Setting RF Parameters in Wireless InSite:

Selection of Propagation Model. WI provides a number of propagation models. We use X3D because it provides high fidelity simulations, taking the following channel parameters into account: 3D ray tracing, indoor-outdoor suitability, atmospheric attenuation (temperature, pressure, humidity), reflection angle dependent reflection coefficients, and MIMO system analysis. Computation-wise X3D uses GPU acceleration.

Designing the Tx and Rx Antennas. In the real world experiments [11], the Talon AD7200 routers are used with Qualcomm's QCA9500 FullMAC IEEE 802.11ad Wi-Fi chip for both Tx and Rx. From the open source characterization of the antenna pattern [49], we retrieve the SNR measurements in 3D for 802.11ad standard (60GHz band) to create the Tx and Rx antenna patterns in WI. The 3D legacy azimuth (ϕ) and elevation (θ) measurements span $[-90^\circ, 90^\circ]$ and $[0^\circ, 32.4^\circ]$, respectively, having 101 and 10 sample points. We show the virtual representation of ϕ and θ in Fig. 7e. Moreover, we match the antenna orientations of WI and the ones in the FLASH dataset on the xy-plane, having the azimuth ($\phi = 0^\circ$) axis of Tx and Rx parallel (See Fig. 5). Steinmetzer et al. [49] indicate that they do not provide beamwidths or steering angles because of the strong variations in the antenna patterns. Thus, we have to depend on case by case basis visual inspection when it comes to beamwidth judgement. Nevertheless, by testing a few custom designed simple directional antennas, for example E and H-plane half-power beam widths of 15° , we confirmed that beamwidth changes correctly affect the received power in WI, e.g. higher received power through smaller beamwidth when the same transmit power is used [51].

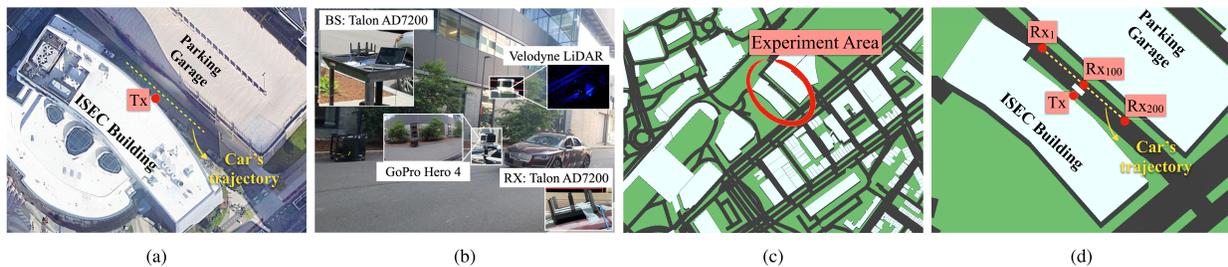


Fig. 3. (a) Experiment location in real world; (b) FLASH experiment setup [8]; (c) Experiment area in the virtual world (*twin*), showing the map beyond the area of interest in downtown Boston; (d) Simulation location in *twin*, showing the first, middle, and last sample points, in order. The 3D box between Tx and Rx₁₀₀ models the car obstacle in the NLOS scenario.

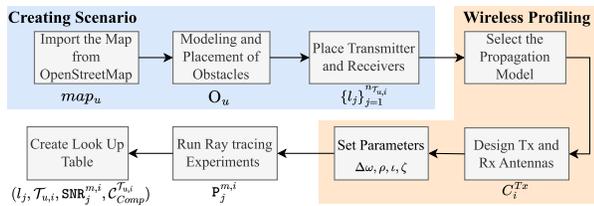


Fig. 4. A typical workflow in Wireless InSite (WI) for creating a twin (suppose $\mathcal{T}_{u,i}$ for scenario \mathcal{S}_u) in the Multiverse.

Selection of $\Delta\omega$, ρ , l , and ζ . WI allows users to select ray spacing $\Delta\omega$ (angle between two adjacent emitted rays), maximum number of reflections ρ , maximum number of transmissions l , and maximum number of diffractions ζ . In our experiments, we select ray spacing to be $\Delta\omega = 0.25^\circ$. We set ρ to be either 1 or 3, because we observe that given the geometry of the environment and the effect of reflections on the received power, this number of reflections is enough to deliver signals to the receivers. We choose the l to be 0, because the data collection environment (urban canyon) is mostly concrete on the sides and the floor is asphalt, which typically do not allow mmWave signals to penetrate through. Finally, ζ is set to be 1, because after 2 diffractions, mmWave signal strength is negligible.

3) *Ray Tracing in Wireless InSite*: For each scenario, we perform ray tracing over Rx locations and 34 Tx beams for each twin. In the experiments, the angle between two rays is set as $\Delta\omega = 0.25^\circ$ and rays are transmitted at all directions in the Euclidean space, following the antenna pattern specifications. Since the rays are emitted in discrete angles and spread as they travel through space, it is not guaranteed that rays would be received at exact Rx location. Thus, in order to ensure that received rays are correctly estimated, WI puts a sphere around the receiver point, with the radius r . Depending on how large that r is, several rays might pass through the sphere. There are two steps to identify which ray to select: i) rays are sorted according to the geometry faces they interact with and the similar ones are eliminated in order to prevent over-predicting the delivered energy at the receiver and ii) the closest ray to the actual Rx point is selected. According to WI manual, $r = \Delta\omega \times D_{max}$, where D_{max} is the coverage range of Tx (≈ 20 m in FLASH). Thus, we estimate $r \approx 8.73$ cm. The X3D model then applies Exact Path Calculations (EPC), which adjusts the reflection and diffraction points within the sphere so that the selected rays pass through the exact receiver points. We consider the highest received power value of the received ray from each Tx beam at each of the $L_{\mathcal{T}_{u,i}}$ Rx locations for twin $\mathcal{T}_{u,i}$.

TABLE II
MULTIVERSE PARAMETERS

Materials	Buildings: Concrete Road: Asphalt Foliage: Grass
Antenna orientation	T _x and R _x facing opposite directions
Waveform	$f_c = 60$ GHz, BW = 2.16 GHz
Tx power	24 dBm [53]
Tx and Rx pattern source	Talon antenna measurements [54]
Tx height	On a cart, 0.95 m
Rx height	On a car, 1.645 m
Noise power (\mathcal{N})	-100.99 dBm
Antenna sensitivity	-250 dBm
Ray spacing ($\Delta\omega$)	0.25°
Number of allowed diffraction (ζ)	1
Number of transmissions through walls (l)	0
Scenarios	LOS (\mathcal{S}_1) and NLOS (\mathcal{S}_2)

TABLE III
DIFFERENCES BETWEEN TWINS

Twins	Baseline ($\mathcal{T}_{u,1}$)	1-Reflection ($\mathcal{T}_{u,2}$)	3-Reflection ($\mathcal{T}_{u,3}$)
Environment	Open space	Urban canyon	Urban canyon
Imported map (map_u)	Random	Boston	Boston
Number of reflections (ρ_i)	3	1	3
$C_{lookup}^{\mathcal{T}_{u,i}}$ (hour)	~ 0.4	~ 50	> 100
Number of Rx locations	$n_{\mathcal{T}_{u,1}} = 1$	$n_{\mathcal{T}_{u,2}} = 200$	$n_{\mathcal{T}_{u,3}} = 200$

C. Multiverse Setup: Twins in Multiverse

We create the Multiverse with three twins each having different cost and precision. All twins have the same materials for the surrounding buildings and foliage, antenna orientations, transmitted waveforms, transmission power, Tx and Rx antenna patterns, coordinates for the Rx sample points, noise power, and antenna sensitivity. To create differences between *twins*, we change (i) environment used in the twin world, (ii) number of allowed reflections (ρ) in simulations, and (iii) number of Rx locations, $L_{\mathcal{T}_{u,i}}$. The common parameters in the digital world creation are given in Tab. II, whereas the differences between twins are highlighted in Tab. III.

1) *Baseline Twin $\mathcal{T}_{u,1}$* : Creating a precise map is expensive and may not be available all the time. Thus, in order to show the gain of precise geometry (which we create for the next twins) and to form a baseline for our evaluations, we use an *open area* environment as the Baseline twin (excluding the buildings present in the FLASH testbed). The *Open area* environment is not a free-space, as reflections from the far surrounding buildings in a city is still allowed. We set the Tx and only one Rx around center region of the *open area*, locating them 4.33 m away, corresponding to the minimum Tx-Rx distance in the FLASH testbed, and record the SNR for 34 beams. We use only one Rx location, because in the *open area* experiments, different Rx locations are negligibly affected by the surroundings, due to high attenuation. Such an experiment setting significantly cuts both time and computation in

ray tracing and forms a basis for our next *twin* evaluations. For this twin, we consider three reflections in our simulations. We design Baseline twin for both the LOS (\mathcal{S}_1) and NLOS (\mathcal{S}_2) scenarios, with $B_1 = 34$. For the \mathcal{S}_2 , we place the obstacle between Tx and Rx, 35 cm away from the Tx, parallel to the situation in the FLASH experiments. To account for different regions on the road, we first compute the angle between the Tx and Rx locations in FLASH. We then multiply the output from open area experiment with the antenna patterns from Talon router at the identified angle, for each beam.

Specific characteristics. Following the notations of Sec. IV, the Baseline twin ($\mathcal{T}_{u,1}$) has: (a) $map_1 = open\ area$, (b) $O_1 =$ retrieved from FLASH [11], (c) $\rho_1 = 3$, and (d) $|L_{\mathcal{T}_{u,1}}| = 1$.

2) *1-Reflection Twin* $\mathcal{T}_{u,2}$: For designing this twin, we switch the simulation environment to the location where the FLASH experiments took place. Creating the realistic environment and using antenna patterns from real life measurements brings a trade-off between significantly increased computation time and precision in beam selection and received power. We set the Tx point in WI by precisely measuring the Tx location in the FLASH experiment location. Moreover, we import the buildings to WI with the same geometry as of the real world. In WI simulations, we collect received beam power at 200 consecutive points, which are uniformly distributed over a 40 meter trajectory, Tx being 4.33 m away from the trajectory's middle point (see Fig 3d). We allow one reflection in the simulations to save on computation time and realize the effects of simpler settings that mostly allow LOS communication. For 1-Reflection twin, we have both the LOS (\mathcal{S}_1) and NLOS (\mathcal{S}_2) scenarios, with $B_2 = 34$. For \mathcal{S}_2 , we again place the obstacle between Tx and Rx, 35 cm away from the Tx, parallel to the situation in the FLASH experiments.

Specific characteristics. The 1-reflection twin ($\mathcal{T}_{u,2}$) has: (a) $map_2 = Boston$, (b) $O_2 =$ retrieved from FLASH [11], (c) $\rho_2 = 1$, and (d) $|L_{\mathcal{T}_{u,2}}| = 200$.

3) *3-Reflection Twin* $\mathcal{T}_{u,3}$: This twin creates a more comprehensive setting in the digital world in order to provide a more precise beam profiling, even in NLOS cases. We consider three reflections in the simulations to cope with potentially dense reflective environments. Like $\mathcal{T}_{u,2}$ in Sec. V-C.2, we run the simulations at the same place as the FLASH dataset was collected. We use 1 Tx, 200 Rx locations, and the same obstacle settings as in $\mathcal{T}_{u,2}$ for both the \mathcal{S}_1 and \mathcal{S}_2 scenarios.

Specific characteristics. Finally, for 3-reflection twin ($\mathcal{T}_{u,3}$): (a) $map_3 = Boston$, (b) $O_3 =$ retrieved from FLASH [11], (c) $\rho_3 = 3$, and (d) $|L_{\mathcal{T}_{u,3}}| = 200$.

D. Observation From Ray Tracing Experiments

Representative Ray Tracing. A collection of outputs from the representative raytracing analysis in the WI simulator for $\mathcal{T}_{u,1}$, $\mathcal{T}_{u,2}$, and $\mathcal{T}_{u,3}$ are presented in Fig. 5 for LOS (\mathcal{S}_1) and NLOS (\mathcal{S}_2) scenarios. We denote different Rx locations in $\mathcal{T}_{u,2}$, and $\mathcal{T}_{u,3}$ with Rx_i , where $1 \leq i \leq |L_{\mathcal{T}_{u,2}}|, |L_{\mathcal{T}_{u,3}}|$, respectively, and $|L_{\mathcal{T}_{u,2}}| = |L_{\mathcal{T}_{u,3}}| = 200$, according to Tab. III. For $\mathcal{T}_{u,1}$, the only Rx location is denoted as 'Rx' ($|L_{\mathcal{T}_{u,1}}| = 1$). In Fig. 5, we present the ray between the Tx and Rx_{130} for $\mathcal{T}_{u,2}$, and $\mathcal{T}_{u,3}$. Moreover, we show Rx_{100} (middle sample point) for location reference. In all examples, Antenna-10 (t_{10}) is used at Tx. We enlarge the corresponding antenna patterns for visual purposes, but we keep their location precise. In each sub-figure, red ray delivers the highest power

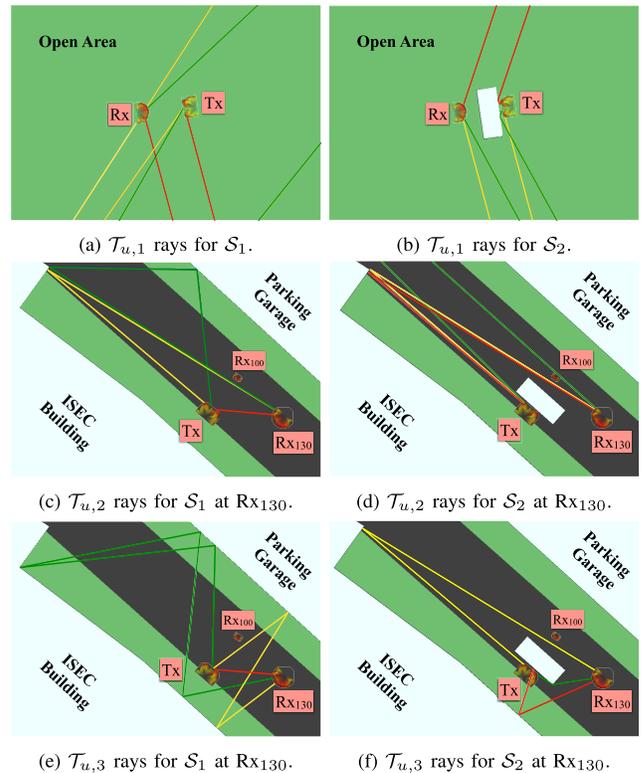


Fig. 5. A set of examples for the first, second, and fifth strongest received rays at Rx_{130} , using antenna-10 (t_{10}) at the Tx in the Multiverse \mathcal{M} for $\mathcal{T}_{u,1}$, $\mathcal{T}_{u,2}$, and $\mathcal{T}_{u,3}$ and both \mathcal{S}_1 and \mathcal{S}_2 scenarios. Red ray is the strongest one, while the ray power decreases towards colder color (red→yellow→green). Antenna locations are correct, their patterns are not in scale, enlarged for visual purposes. Due to the mismatch of Tx and Rx antenna height, there was no direct path for LOS rays when the antennas are directly in front of each other.

and ray strength decreases in the order of red→yellow→green, where they represent first, second, and fifth strongest rays.

We observe that when the antennas' $\phi = 0^\circ$ axes are aligned and the antennas are close, no direct ray is possible. This is due to the height difference between Tx and Rx, which comes from the precise heights we set for the antennas following the FLASH experiment setting [8], and the antenna patterns from Talon [49]. Thus, the rays in $\mathcal{T}_{u,1}$, given in Figs. 5a and 5b, are delivered by reflection and diffraction from distant buildings. As a result, the Baseline twin, does not feature the reflections from the building in FLASH testbed, making results less reliable. This suggests the idea that a carefully created twin at the same experiment location would yield more robust results, which in fact, we confirm in Sec. VI. As the antennas become more distant, the height difference is compensated by longer ray travel distance, making a direct ray between Tx and Rx possible (see red line 5c and 5e). In $\mathcal{T}_{u,2}$ and $\mathcal{T}_{u,3}$, direct paths are identical, having the $p_{130}^{10,2} = p_{130}^{10,3} = -69.65dBm$. The rest of the LOS rays tend to be drastically different due to the allowed number of reflections ($\rho_2 = 1$ and $\rho_3 = 3$). In Fig. 5c yellow and green rays are delivered to the Rx with at most one reflection, whereas in Fig. 5e these rays are able to undergo multiple reflections.

In NLOS cases direct path is prevented. Thus, rays have to undergo reflections and/or diffractions in order to reach to the Rx. This is where we see the difference between $p_{130}^{10,2}$ and $p_{130}^{10,3}$. The rays in $\mathcal{T}_{u,2}$ are only allowed to reflect once compared to $\mathcal{T}_{u,3}$, where $\rho_i = 3$. Thus, in $\mathcal{T}_{u,2}$, rays are

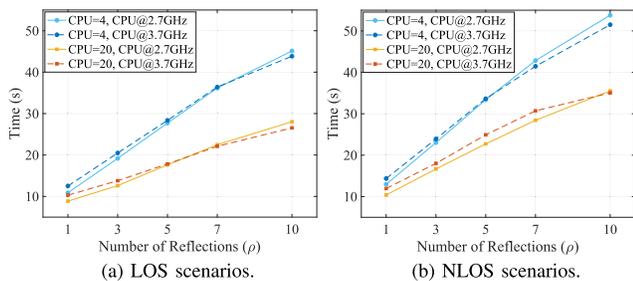


Fig. 6. Comparing computation time to run a single antenna experiment on WI for different number of reflections and CPU usage on computers with Intel i9-11900 processor with 32GB RAM (CPU@2.7GHz) and Intel i9-10900K processor with 48GB RAM (CPU@3.7GHz). Overall, NLOS requires more time to complete simulations compared to the LOS case, because the obstacle in the NLOS case creates more surface for rays to reflect, increasing the calculation time. These plots suggest that with more computation power, fast ray tracing and beam selection operations could be performed at a plausible fidelity level.

TABLE IV
COMPARING TALON AND WI METRICS

	ϕ	ϕ_δ	θ	θ_δ	SNR
Talon [50]	$[-90^\circ, 90^\circ]$	1.8°	$[0^\circ, 32.4^\circ]$	3.6°	same
WI	$[-100^\circ, 100^\circ]$	2°	$[0^\circ, 36^\circ]$	4°	same

more likely to travel longer to arrive at Rx when there is no direct path. This way, signals in $\mathcal{T}_{u,2}$ attenuate more, causing lower received power for the strongest beam, e.g., $p_{130}^{10,2} = -108.4 \text{ dBm}$ and $p_{130}^{10,3} = -78.16 \text{ dBm}$ in Figs. 5d and 5f, respectively. Again, due to different number of reflections, we observe yellow and green rays follow different paths in NLOS cases. Overall, we tend to observe $p_j^{m,3} \geq p_j^{m,2}$. Moreover, we expect more accurate beam matching performance for $\mathcal{T}_{u,3}$, as we also confirm in the Sec. VI.

Cost for Creating the Twins. The cost of creating $\mathcal{T}_{u,1}$, $\mathcal{T}_{u,2}$, and $\mathcal{T}_{u,3}$ is determined by Eq. 4. In Eq. 4, the $\mathcal{C}_{\text{map}}^{\mathcal{T}_{u,i}}$ consists of one-time cost for: (i) finding appropriate maps and confirming building/object dimensions with their real life counterparts, (ii) antenna pattern design, (iii) precise antenna locations and heights, (iv) obstacle locations and dimensions, (v) deciding on appropriate waveforms and materials, and (vi) selecting the most representative propagation model. We observe $\mathcal{C}_{\text{map}}^{\mathcal{T}_{u,1}} < \mathcal{C}_{\text{map}}^{\mathcal{T}_{u,2}} = \mathcal{C}_{\text{map}}^{\mathcal{T}_{u,3}}$, because $\mathcal{T}_{u,1}$ is an open area without any building details, whereas $\mathcal{T}_{u,2}$ and $\mathcal{T}_{u,3}$ are exact map replica of the real world. On the other hand, $\mathcal{C}_{\text{lookup}}^{\mathcal{T}_{u,i}}$ is an offline computation cost. In order to create the lookup tables, we use a Dell XPS computer with Intel i9 processor and 32GB RAM following Sec. IV-B.1. Detailed analysis of computation time with respect to the CPU usage is given in Fig. 6, where we observe a semi-linear relation between the computation time and the number of reflections (ρ). The $\mathcal{C}_{\text{lookup}}^{\mathcal{T}_{u,i}}$ for $\mathcal{T}_{u,1}$, $\mathcal{T}_{u,2}$, and $\mathcal{T}_{u,3}$ are given in Tab. III.

Antenna Similarity (visual representation). In Fig. 7a-7d, we showcase several antenna pattern examples that we have re-created in WI ray tracing software and provide their comparisons with the Talon antenna patterns for visual similarity. The antenna pattern in Fig. 7a belongs to the 24th beam (element t_{24}), whereas Fig. 7b is the Rx in WI. In Fig. 7c and 7d, we provide an example comparison between the Talon antenna patterns (solid blue line) and the re-created beams in WI (shaded area) for the 24th beam (element t_{24}) in 2D azimuth ($\theta = 0^\circ$) and elevation ($\phi = 0^\circ$), respectively. The slight discrepancy comes from the fact that WI's user defined antenna patterns only accept sample points with

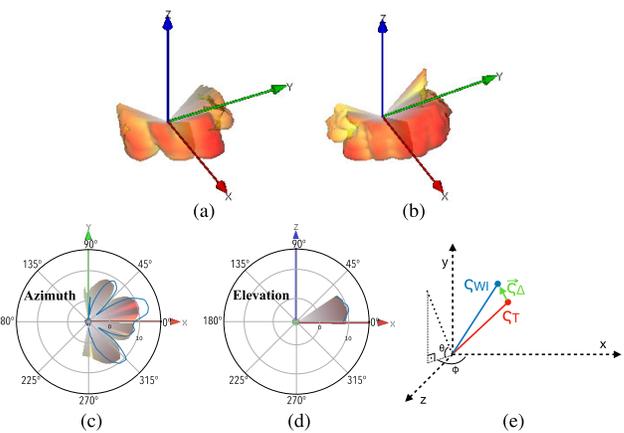


Fig. 7. Re-created antenna pattern examples in Wireless InSite: (a) Tx Antenna-24 (t_{24}) and (b) the Rx antenna. Antenna pattern comparisons between the Talon antenna patterns (solid blue line) and the re-created beams in WI (shaded area): (c) Azimuth, (d) Elevation. (e) Antenna pattern similarity: The definition of ϕ and θ on the Cartesian coordinates and vector discrepancy calculation, where ζ_T and ζ_{WI} are points in experimental and simulation antenna pattern, respectively, and ζ_Δ is the vector difference between the ζ_T and ζ_{WI} .

integer increments ($\phi_\delta, \theta_\delta$). To preserve the complete SNR values, we distribute the experimental SNR values over the closest azimuth and elevation regions, while keeping $\phi = 0$ as the reference point; thus, $\phi \in [-100^\circ, 100^\circ]$ and $\theta \in [0^\circ, 36^\circ]$. The complete metrics for antenna pattern sample comparison is given in Tab. IV.

Antenna Similarity (quantitative measures). In addition to a visual comparison, we also provide a quantitative measure for the antenna similarity between Talon and WI. We measure the antenna discrepancy with an average score from each of these 34 beams, given by the magnitude of the vector differences between experimental and simulation antenna patterns. A visual vector representation of antenna pattern samples is given in Fig. 7e, where ζ_T and ζ_{WI} represent points in experimental and simulation antenna pattern, respectively, which were created with corresponding SNR values, ϕ , and θ . The discrepancy score for a single antenna, given in Eq. 9, is calculated by averaging the normalized magnitudes of difference vectors, $\vec{\zeta}_\Delta$, where the normalization factor is the magnitude of the corresponding antenna pattern sample ($|\vec{\zeta}_{WI}| = |\vec{\zeta}_T|$).

$$\Delta_{\text{single}} = \frac{1}{Q} \sum_{i=1}^Q \frac{|\vec{\zeta}_{WI_i} - \vec{\zeta}_{T_i}|}{|\vec{\zeta}_{WI_i}|}, \quad (9)$$

where $Q = 1010$ (representing the number of sample points in the antenna patterns), and $0 \leq \Delta_{\text{single}} \leq 2$. The final average antenna pattern similarity score, averaged over 34 beams, is found to be $\Delta = 0.0931$ for both scenarios, \mathcal{S}_1 and \mathcal{S}_2 . Smaller the Δ , more similar the antenna patterns are. The designed set of antenna patterns for the Multiverse using WI simulation are available in [12].

VI. PERFORMANCE OF THE MULTIVERSE

We validate our proposed framework using the FLASH dataset, presented in Sec. V-A. We consider two set of scenarios, including LOS (\mathcal{S}_1) and NLOS (\mathcal{S}_2). For each scenario, the Multiverse consists of three twins each, $\mathcal{M}_1 = \{\mathcal{T}_{1,1}, \mathcal{T}_{1,2}, \mathcal{T}_{1,3}\}$ for LOS scenario as an example.

TABLE V

THE $Acc_{(1,0)}$ WHILE BEING EXPOSED TO THE UNSEEN ENVIRONMENTS. THE DL-BASED METHODS EXPERIENCE EXTENSIVE DROP IN ACCURACY DUE NON-ADAPTABILITY TO THE CHANGES IN THE ENVIRONMENT

Train Scenario	Train $Acc_{(1,0)}$	Test $Acc_{(1,0)}$		Drop in $Acc_{(1,0)}$
		S_1	S_2	
S_1 (LOS)	83.9	66.40	10.30	56.1
S_2 (NLOS)	90.97	5.96	79.38	73.42

A. Evaluation Metrics

Given the ground-truth measurements from FLASH dataset, $G \in \mathbb{R}^{|\mathcal{B}|}$, and the SNRs $S \in \mathbb{R}^{|\mathcal{B}|}$ for \mathcal{B} beams (either obtained from the Multiverse or predicted by DL-based method), we evaluate the beam selection performance as:

$$Acc_{(K,T)} = \frac{1}{V} \sum_{l=1}^V \mathbb{1}_{(\exists g \subset G | g \in S_k)}, \quad (10)$$

where V denotes the number of samples and τ is a Boolean predicate, with $\mathbb{1}_\tau$ to be 1 if τ is true, and 0 otherwise. In Eq. 10, the set g denotes a subset of beams in G (ground-truth from FLASH dataset) such that the observed SNR of beams in g is within the $T - dB$ threshold of the optimum beam, i.e., $g = \{t_m \in \mathcal{B} | SNR_{t_m} \geq SNR_{t^*} - T\}$, where SNR_{t^*} denotes the SNR of the optimum beam from FLASH ground-truth. Moreover, S_K denotes the top- K beams obtained by either the Multiverse or DL-based method, defined as:

$$S_K = \{s | \arg \max_{s \subset \{1, \dots, |\mathcal{B}|\}, |s|=K} \sum_{m \in s} SNR_{t_m}\}, \quad (11)$$

where SNR_{t_m} denotes the inferred SNR for beam t_m . This is intuitive as our observation indicates that there are closely competitive beams in Talon radio [49], used for collecting the FLASH dataset. Thus, selecting any of them results in a favorable observed SNR. For $T = 0$ and $K = 1$, we get the conventional top-1 accuracy, where only the best beam in FLASH is compared with the best inferred beam. Throughout our experiments, we mark the results based on the seen scenario, unless explicitly mentioned otherwise. For example, the results for ‘‘LOS’’ scenario correspond to a setting where the LOS scenario is considered as seen, and NLOS scenario is considered as unseen. Similarly, the experiments for ‘‘NLOS’’ scenario correspond to having NLOS and LOS scenarios as seen and unseen, respectively.

B. Motivation for the Multiverse

We consider a scenario where the mmWave beam initialization is performed by only relying on multimodal data and DL-based method proposed in Sec. IV-A. We follow the models released publicly along with the FLASH dataset to train a CNN [8]. The proposed architecture is inspired by ResNet [54] and exploits feature level fusion to reinforce the prediction, by incorporating the sensor data from all modalities. In Tab. V, we benchmark the performance of the DL-based method when the CNN is trained on one scenario and tested on another (unseen during training), e.g. trained on LOS scenario S_1 and tested on NLOS scenario S_2 . We observe that the maximum drop in accuracy ($Acc_{(1,0)}$) is 73.42%, while making a transition from S_2 to S_1 scenario. In Fig. 8, we present the distribution of the beams in LOS and NLOS scenarios. In the mmWave beamforming, the codebooks are fixed; thus, out-of-distribution classes (i.e., beams) are not observed. However,

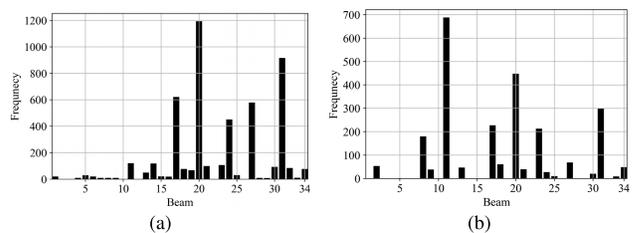


Fig. 8. Distribution beams for (a) LOS and (b) NLOS scenario in FLASH dataset. Presence of obstacles significantly affect the optimal beam in the mmWave band.

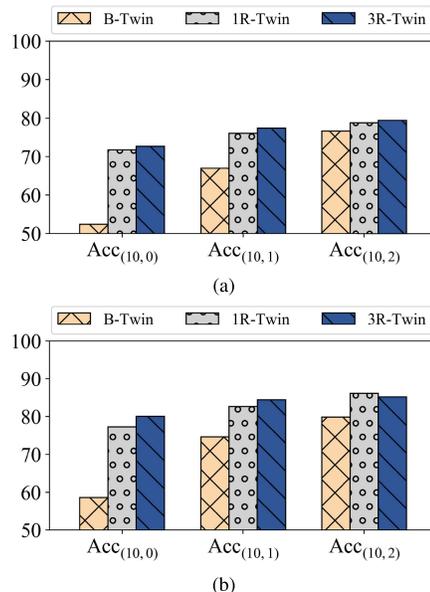


Fig. 9. Benchmarking the performance of three twins in the Multiverse against the real-world measurements for (a) LOS S_1 and (b) NLOS S_2 scenarios. The accuracy for $K = 10$ and SNR thresholds of 0, 1, 2 dB indicate the fidelity of the Multiverse to the real world setting. The 3-Reflection twin and 1-Reflection twin exhibits more fidelity over the Baseline twin.

the presence of obstacles can affect the distributions of the beam significantly (for example the dominant optimal beam shifts from 20 to 11 while transitioning from LOS and NLOS scenario), which cannot be handled with the DL-only method.

Observation 1: From Tab. V, we observe that if a DL-only method is used for beam selection, the accuracy severely degrades, when unseen scenarios are encountered.

C. Validation of the Multiverse Concept

In Fig. 9, we compare the performance of the Multiverse against the ground-truth measurements from the FLASH dataset to validate the fidelity of emulation outputs. We report the metric $Acc_{(K,T)}$ in Eq. 10 and set the parameter K as 10 and gradually relax the SNR threshold T with $T = \{0 dB, 1 dB, 2 dB\}$. Note that $T = 0$ is the most extreme case, where only the optimum beam in FLASH is considered for justification. From Fig. 9, we observe that for the LOS scenario (S_1), Baseline twin $T_{1,1}$ provides the $Acc_{(10,0)}$ of 52.38%, while 1-Reflection and 3-Reflection twins $T_{1,2}$ and $T_{1,3}$ perform closely with accuracy of 71.74% and 72.69%, respectively. The $Acc_{(10,0)}$ increases as we relax the threshold on SNR and ranges between 76-79% across all three twins. For NLOS scenario, S_2 , we observe that 3-Reflection twin $T_{2,3}$ exhibits superiority in $Acc_{(10,0)}$ against $T_{2,1}$ and $T_{2,2}$

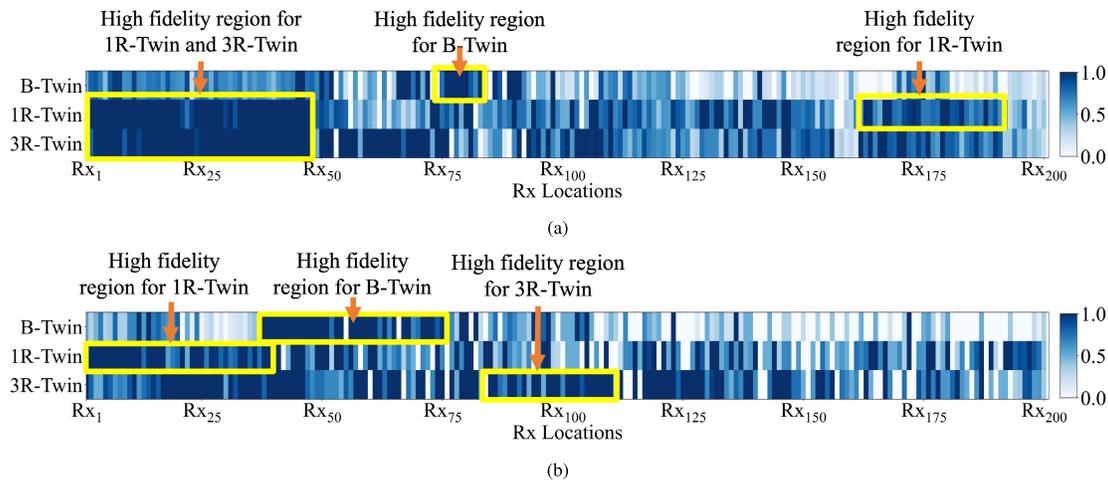


Fig. 10. Probability of inclusion for three twins for (a) LOS \mathcal{S}_1 and (b) NLOS \mathcal{S}_2 scenarios. Each column shows a region on the road and each row depicts the probability of inclusion for one of the twins. While the performance on twins varies in different regions, the 3-Reflection twin outperforms the Baseline and 1-Reflection twins on average. All twins offer higher probability of inclusion in \mathcal{S}_1 compared to \mathcal{S}_2 . A few sample regions where one twin has more fidelity than other, is highlighted in both (a) and (b).

by 21.49% and 2.78%, respectively. Overall, the designed Multiverse depicts the accuracy ranging between 52.38 – 80.08%, 66.92 – 84.41% and 76.63 – 85.22% for SNR thresholds of 0 dB, 1 dB, 2 dB, respectively, for both LOS (\mathcal{S}_1) and NLOS (\mathcal{S}_2) scenarios across all twins.

Observation 2: From Fig. 9, we observe that the accuracy over the three twins is related as $\mathcal{T}_{u,1} < \mathcal{T}_{u,2} < \mathcal{T}_{u,3}$, regardless of the scenarios. We observe the maximum accuracy of 79.43% and 85.22% with 3-Reflection twin for LOS and NLOS scenarios, respectively, which corroborates the fidelity of the Multiverse against the ground-truth measurements from the FLASH dataset.

D. Comparing the Performance of Twins

In Fig. 10, we compare the probability of inclusion $p(K, LT(\mathcal{T}_{l,i}), \mathcal{S}_l, r)$ for $K = 10$, three twins ($i = \{1, 2, 3\}$), and both LOS and NLOS scenarios ($l = \{1, 2\}$). In this figure, each blue shaded rectangle represents the location on the road (i.e., r) for which the probability of inclusion is calculated, with 20cm spacing, as explained in Sec. V. From this figure we conclude three observations. First, the performance of the twins depends on the location of the receiver on the road. This is intuitive as the number of reflections in real world might vary in each location, unlike the Multiverse where it is constant for each twin. Thus, each twin might be closer to the real world scenario on a case by case basis. Second, we observe that 3-Reflection twin performs better than Baseline and 1-Reflection twins with the average probability of inclusion of 74.87 – 78.68% compared to that of 43.90 – 47.48% and 66.43 – 76.60% for LOS (\mathcal{S}_1) and NLOS (\mathcal{S}_2) scenarios, respectively. Third, the Multiverse provides higher probability of inclusion in \mathcal{S}_1 compared to \mathcal{S}_2 by 4 – 8% margin.

Observation 3: From Fig. 10, we observe that the performance of the twins varies over different regions of the road (i.e. there is not a twin that is superior at all time). Thus, the region wise selection of twins (as per Eq. 7) ensures that the optimal twins is selected on a case-by-case basis. The average probably of inclusion over all regions of the road is higher with 3-Reflection twin, which has a more complex modeling of the propagation pattern.

E. Digital Twin Selection From the Multiverse

Effect of Control Parameter. The proposed optimization problem in Eq. 7 selects the optimum twin and set of top- K beams in different regions of the road. In Fig. 11a, we show the effect of control parameter α on the optimization problem in Eq. 7 for LOS (\mathcal{S}_1) and NLOS (\mathcal{S}_2) scenarios, where there is no constraint on the communication latency. We observe that by increasing the control parameter α , both probability of inclusion and beam selection time decrease. Note that by increasing α , the optimization problem in Eq. 7 is enforced to decrease K to lower the communication latency cost, which results in reduced probability of inclusion and beam selection time, consequently. In particular, we observe that when $\alpha = 0$, the beam selection time for LOS and NLOS scenarios are 3.43ms and 3.60ms, respectively, and the probability of inclusion is 1 for both. Moreover, our experiments indicate that when $\alpha = 0$, 3-Reflection twin is the dominantly selected by ratio of 99.5% and 93% in LOS and NLOS scenarios, respectively. This is intuitive as when $\alpha = 0$ and there is no constraints on the communication latency, the algorithm is enabled to select the twin with the highest fidelity and as high values of K to ensure beam alignment. Moreover, we observe that, when $\alpha > 15$ the beam selection time is lower than 0.46ms, which corresponds to sweeping up to 3 beams according to Eq. 6, imposed due to weighting the second term in Eq. 7.

Observation 4: Increasing the control parameter α encourages the optimization problem in Eq. 7 to select a lower K and decrease the communication cost. From Fig. 11a, we observe that both probability of inclusion and beam selection time decrease as α increases (as a consequence of having a more strict bound on K).

Effect of Latency Constraint. In Tab. VI, we report the percentage of the times that each of the twins were selected under three communication latency constraints, $\mathcal{C}_{Comm_1} = 0.62ms$, $\mathcal{C}_{Comm_2} = 1.56ms$ and $\mathcal{C}_{Comm_3} = 3.12ms$, corresponding to the maximum allowed number of beams being 4, 10, and 20 according to the 5G-NR standard (Eq. 6), respectively. In Tab. VI, the summation of each column is equal to 100%, and each row depicts the usage of specific twins as the the communication latency constraint is relaxed,

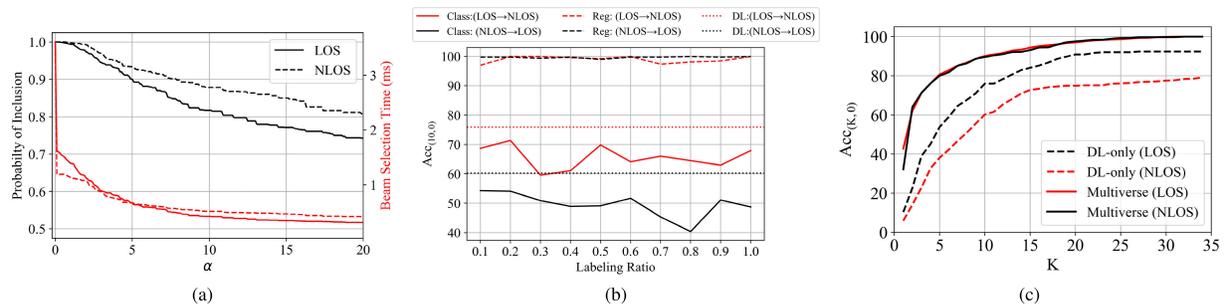


Fig. 11. (a) Analysis of probability of inclusion and beam selection time for different α values in Eq. 7. The probability of inclusion decreases as the optimization problem is weighted on the second term in Eq. 7 that encourages minimizing the communication cost. (b) The $Acc_{(10,0)}$ of DL-based method when fine-tuned by the labels obtained by 3-Reflection twin. We observe significant improvement by fine-tuning with regression loss. The classification method, on the other hand, shows lower performance than the case without fine-tuning. (c) Comparing the accuracy ($Acc_{(K,0)}$) of DL-only method with Multiverse for different values of K . Multiverse is superior to DL-only method for all values of K and both LOS/NLOS scenarios.

from the most strict (\mathcal{C}_{Comm_1}) to the least strict (\mathcal{C}_{Comm_3}). We report the selection percentage for $\alpha = \{0, 2\}$. For $\alpha = 0$, we observe, as the constraint on communication latency is relaxed, the percentage of the times that the twin with higher complexity (and higher probability of inclusion) is selected increases. In particular, the usage for 3-Reflection twin increases from 42.5% and 62% to 74% and 68.5% for LOS and NLOS scenarios, respectively. On the other hand, the usage for 1-Reflection twin decrease by 31% and 4.5% for each scenario and the usage for the Baseline twin stay within 5–6.5% range. This is intuitive as relaxing the communication latency constraint indicates that the users are inclined to have a more robust than fast communication; thus, the twin with higher complexity and higher probability of inclusion is more likely to be chosen.

When $\alpha = 2$, the second term in Eq. 7 is weighted as twice. In this case, we observe that the usage for Baseline twin increases from 7% and 5.5% to 12% and 10% for LOS and NLOS scenarios, respectively, as the communication latency is relaxed. The 1-Reflection twin exhibits 11.5% decrease in usage in LOS scenario and 5.5% increase in NLOS scenario. Finally, the percentage of the times that 3-Reflection twin is used decreases. By relaxing the communication latency constraint, the number of allowed beams increases. As a result, less complex twins such as Baseline and 1-Reflection twins can appear as competitive method to the 3-Reflection twin.

Observation 5: From Tab. VI, we observe that relaxing the communication latency constraint favors selecting the twin with higher complexity (and higher probability of inclusion). However, increasing the control parameter α in Eq. 7 increase the usage of Baseline and 1-Reflection twins as the communication constraint is relaxed.

F. Fine-Tuning With the Multiverse Ground-Truth

We consider a scenario in which a percentage of labels from the Multiverse (ray tracing outputs), which we refer to as *labeling ratio*, is paired with the local multimodal sensor data to fine-tune the local model at the vehicle. For testing, we only use the labels from FLASH dataset that is our ground-truth. We generate the labels from 3-Reflection twin in the Multiverse, which offers the highest average probability of inclusion. In one experiment, we use the multimodal sensor data and labels from FLASH ground-truth for \mathcal{S}_1 . We label the sensor data within \mathcal{S}_2 using the ray tracing outputs from the Multiverse (not FLASH) and fine-tune the model (LOS→NLOS). We then repeat the same experiment with \mathcal{S}_1 and \mathcal{S}_2 role-reversed (NLOS→LOS). As in Sec. VI-B,

we use the model architecture released by FLASH framework [8]. In Fig. 11b, we compare the $Acc_{(10,0)}$ of DL-only method without fine-tuning against fine-tuning with classification and regression loss. The $Acc_{(10,0)}$ without fine-tuning are 75.95% and 60.23% for LOS and NLOS scenarios (see dotted lines in Fig. 11b), respectively. We note that the $Acc_{(1,0)}$ and $Acc_{(10,0)}$ for LOS (\mathcal{S}_1) and NLOS (\mathcal{S}_2) scenarios in the Multiverse range between 41 – 72% and 47 – 80% according to Fig. 9, respectively. On the other hand, we observe that $Acc_{(10,0)}$ with fine-tuning ranges between 40.35 – 54.27% and 59.54 – 69.84% for \mathcal{S}_1 and \mathcal{S}_2 , with classification. Interestingly, the $Acc_{(10,0)}$ with fine-tuning is within the $Acc_{(1,0)}$ and $Acc_{(10,0)}$ from the Multiverse and it is even lower than the performance of DL method without fine-tuning, compare red solid and dotted lines for (LOS→NLOS) and black solid and dotted lines for (NLOS→LOS). We note that the labeling strategy in FLASH framework in one-hot encoding. This significantly limits the performance of fine-tuning with labels from Multiverse, since the ray tracing profiles from all beams except for the one with the highest power is lost. Thus, we go one step further and use the same FLASH framework with a regression loss. Here, we use the Mean Absolute Error (MAE) for training the model and replace the *softmax* layer with *relu*; however, when it comes to evaluation, we report the $Acc_{(10,0)}$ to benchmark the performance against two other competing methods (DL-only without fine-tuning and classification). Interestingly, we observe the fine-tuning with a regression model significantly improves the $Acc_{(10,0)}$. In particular, we observe that the $Acc_{(10,0)}$ with regression is higher than 96.95% for all labeling ratios, which is superior to the case without fine-tuning by 21 – 36.72% and classification by 24.62 – 55.64%.

Observation 6: From Fig. 11b, we observe > 96% accuracy, while fine-tuning the local DL model with a regression loss. With classification and one-hot encoding, the ray tracing information from all the beams, except for the one with the highest power, is lost. Thus, a regression based fine-tuning is more suitable in our use case.

G. Multiverse Against DL-Only Method

In Fig. 11c, we benchmark the performance of Multiverse against the DL-only method for LOS and NLOS scenarios. For DL-only method, we only use the DL models from Sec. IV-A and report the $Acc_{(K,0)}$ for 34 beams in FLASH dataset. For Multiverse, we report the accuracy, when the communication latency constraint is set as each K in x-axis and the control parameter α is set as zero. Note that the communication

TABLE VI

THE BREAKDOWN OF THE USAGE (IN %) FOR THREE TWINS IN THE MULTIVERSE FOR LOS (\mathcal{S}_1) AND NLOS (\mathcal{S}_2) SCENARIOS WHILE IMPOSED WITH THREE COMMUNICATION LATENCY CONSTRAINTS AND $\alpha=\{0, 2\}$. WHEN $\alpha = 0$, RELAXING THE COMMUNICATION CONSTRAINT RESULTS IN AN INCREASE IN CHOOSING 3-REFLECTION TWIN AND DECREASE IN CHOOSING BASELINE AND 1-REFLECTION TWINS. HOWEVER, WHEN $\alpha \neq 0$, THE USAGE OF BASELINE AND 1-REFLECTION TWINS INCREASES AS THE COMMUNICATION CONSTRAINT IS RELAXED

Twin	LOS \mathcal{S}_1 ($\alpha=0$)			LOS \mathcal{S}_1 ($\alpha=2$)		
	\mathcal{C}_{Comm_1}	\mathcal{C}_{Comm_2}	\mathcal{C}_{Comm_3}	\mathcal{C}_{Comm_1}	\mathcal{C}_{Comm_2}	\mathcal{C}_{Comm_3}
B-Twin ($\mathcal{T}_{1,1}$)	6.5	6	6	7	8	12
1R-Twin ($\mathcal{T}_{1,2}$)	51	31.5	20	51.5	46	40
3R-Twin ($\mathcal{T}_{1,3}$)	42.5	62.5	74	41.5	46	48
Twin	NLOS \mathcal{S}_2 ($\alpha=0$)			NLOS \mathcal{S}_2 ($\alpha=2$)		
	\mathcal{C}_{Comm_1}	\mathcal{C}_{Comm_2}	\mathcal{C}_{Comm_3}	\mathcal{C}_{Comm_1}	\mathcal{C}_{Comm_2}	\mathcal{C}_{Comm_3}
B-Twin ($\mathcal{T}_{2,1}$)	5	6.5	3	5.5	7.5	10
1R-Twin ($\mathcal{T}_{2,2}$)	33	29.5	28.5	36	35	41.5
3R-Twin ($\mathcal{T}_{2,3}$)	62	64	68.5	58.5	57.5	48.5

TABLE VII

BENCHMARKING MULTIVERSE AGAINST USING SINGLE DIGITAL TWINS AND DL-ONLY METHODS. THE MULTIVERSE RESULTS IN THE HIGHEST ACCURACY BY SELECTING THE APPROPRIATE TWIN AND SET OF TOP- K BEAMS WITH MODERATE BEAM SELECTION TIME AND COMPUTATION COST. IN THIS TABLE, “ \times ” DENOTES THE ONE UNIT OF TIME THAT SCALES ACCORDING TO THE AVAILABLE COMPUTATION RESOURCES

Methodology	$Acc_{(10,0)}$ (%)		Beam Selection Time (ms)		Computation Cost
	LOS	NLOS	LOS	NLOS	
DL-only	75.95	60.23	1.56	1.56	0.6ms
Single B-Twin	47.48	43.90	0.90	0.75	\times
Single 1R-Twin	76.60	66.43	0.74	0.50	100 \times
Single 3R-Twin	78.68	74.87	0.85	0.52	200 \times
Multiverse	90.15	89.56	0.89	0.62	157 \times

latency constraint in Eq. 7, imposes an upper bound on the number of allowed beams, and the number of selected beams can be less for each K in x-axis with Multiverse. From this figure, we observe that Multiverse is outperforming the DL-only method for both LOS and NLOS scenarios and over all K values. In particular, the maximum achieved accuracy with DL-only method are 92.43% and 79.08% and are achieved when $K > 26$ and $K > 33$ for LOS and NLOS scenarios, respectively. On the other hand, the Multiverse achieves a similar accuracy when $K > 12$ and $K > 4$. As a result, we conclude that for the same target K the Multiverse has higher accuracy and for the same target accuracy it results in the lower overhead. We provide quantitative comparison with respect to the beam selection time in Sec. VI-I.

Observation 7: From Fig. 11c, we observe that Multiverse is superior to DL-only method for all values of K and both LOS and NLOS scenarios.

H. Multiverse Against Single Digital Twin

We consider a scenario in which only a single digital twin is available for beam selection, while encountering an unseen scenario. In Tab. VII, we report the accuracy ($Acc_{(10,0)}$), beam selection time, and computation cost for four competing methods. In the DL-only method, we use the deep learning models that are trained on the seen scenarios (see Sec. IV-A). In single B-Twin setting, we consider Baseline twin as the only option available and use the look up tables from this twin to solve Eq. 7. Note that, in this case, solving Eq. 7 corresponds to selecting the set of top- K beams only, whereas when multiple twins are available the optimization problem in Eq. 7 jointly optimizes the selection of the twin and set of top- K beams. The same logic applies to single 1R-Twin

and single 3R-Twin settings as well. Finally, we compare the above methods against Multiverse, where three twins each with a different level of fidelity are available. For all experiments, we set the control parameter α as zero in Eq. 7 to compare the twins with respect to their best achievable accuracy. Nevertheless, we also report the beam selection time to compare these four methods in the communication latency. From Tab. VII, we observe that Multiverse provides the highest accuracy compared to all three single twins (by 11.47–45.66% margin) as well as the DL-only (by 14.2 – 29.33% margin) for both LOS and NLOS scenarios. Interestingly, we observe that the Multiverse outperforms the single 3R-Twin, which shows the highest average accuracy according to Fig. 9. This is aligned with the experiments in Fig. 10, where we observe that there is not a dominant twin that outperforms others over all locations on the road. This corroborates the idea of region wise selection of twins in Eq. 7, where we select the optimum twin from Multiverse and top- K beams on a case-by-case basis.

From the beam selection time and computation cost perspective, the DL-only methods requires training a model for 100 epochs, which is costly; however, at the inference phase, it predicts the optimum beam in $\sim 0.6 ms$ (based on measurements on a NVIDIA V100 GPU). Considering the delay of sweeping top-10 beams (1.56 ms as per Eq. 6), the entire computation overhead with DL-only methods is $\sim 2.16 ms$. For single twins, we observe that the B-Twin has the highest average beam selection time over LOS and NLOS scenarios. This is because the B-Twin has the lowest average level of fidelity compared to two other twins. Thus, the optimization problem in Eq. 7 is encouraged to increase the number of selected top- K beams to ensure capturing the optimum beam, which results in the increased beam selection time. On the other hand, the Multiverse results in slightly higher average beam selection time compared to single 1R-Twin and 3R-Twin twins, but higher accuracy according to Tab. VII. As described in Sec. IV-B.1, the computation cost of the generating the look up tables for single twins is proportional to the number of reflections (ρ_i), Rx locations ($n_{\mathcal{T}_{u,i}}$), and beams (\mathcal{B}_i). Overall, we observe that the computation overhead is related as described in the last column of Tab. VII for different twins, and scales (with factor “ \times ”) depending on the available computation resources for running the ray tracing emulation. For Multiverse, different twins might be selected at different regions of the road. We identify the ratio of selecting each of the twins as 0.06, 0.315, and 0.625 for B-Twin, 1R-Twin, and 3R-Twin for LOS scenario, and 0.065, 0.295, and 0.64 for NLOS scenario, respectively. This results in the average relative computation cost of 157 \times . Note that in

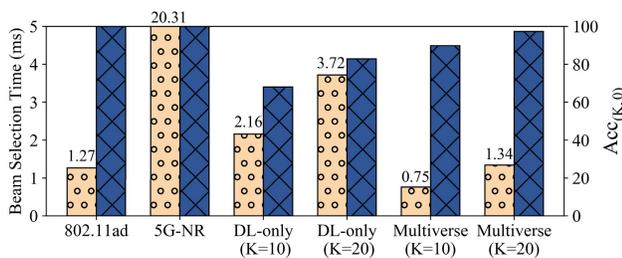


Fig. 12. The beam selection time and accuracy for two exhaustive search based methods (802.11ad and 5G-NR) and DL-only and Multiverse for $K = \{10, 20\}$. The beam selection time and accuracy for DL-only and Multiverse are averaged over LOS and NLOS scenarios for DL-only and Multiverse. The actual beam selection time of 20.31 ms for 5G-NR is scaled here, for better visibility and comparison purpose.

our system design, the look up tables are generated offline (see Sec. IV-B.1). Thus, the computation cost of running the ray tracing emulation is not relevant in the real-time Multiverse interaction step (see Sec. IV-B.2).

Observation 8: From Tab. VII, we observe that Multiverse is superior to DL-only as well as using single digital twins in accuracy, with relatively close beam selection time and moderate computation cost.

I. Benchmarking Multiverse in Accuracy and Overhead

In Fig. 12, we report the beam selection time and $Acc_{(K,0)}$ for two exhaustive search based methods (802.11ad and 5G-NR standards), DL-only and Multiverse for $K = \{10, 20\}$. We use Eq. 6 to model the beam selection time for different number of beams with 5G-NR, DL-only, and Multiverse. Later, we provide quantitative comparisons with 802.11ad standard as well. With 5G-NR standard, it takes 20.31 ms to exhaustively sweep 34 beams as per Eq. 6, and it provides 100% accuracy. The DL-only method, on the other hand, reduces the beam selection time to 2.16 ms and 3.72 ms for 10 and 20 beams, while providing the accuracy of 68.09% and 82.90%, respectively. With Multiverse, whenever the DL-based method identified as faulty, the BS exploits the look up tables and runs the optimization problem in Eq. 7 to obtain the optimum twin and associated top- K beams according to the communication latency constraints. The local execution of Eq. 7 takes $\sim 23\text{ }\mu\text{s}$ on a commercial laptop. Thus, given the the probability of inclusion from the so far seen scenario, the BS can determine the optimum twin and set of top- K beams in near real-time. From Fig. 12, we observe that the Multiverse is the most successful method with 0.75 ms and 1.34 ms beam selection time and 89.86% and 97.35% accuracy for $K = 10$ and $K = 20$, respectively. Interestingly, we observe that the Multiverse with 10 beams outperforms the DL-only method with 20 beams in both beam selection time and accuracy, despite of being imposed of more strict constraints. Moreover, we present the beam selection time with another exhaustive search based method, namely 802.11ad standard. This standard is a contention based channel access method and the minimum beam selection time with this standard is reported as 1.27 ms [49] for 34 beams. We note that the selected top- K beams in the Multiverse are bounded by 12 while targeting the inclusion probability of 92% (there is no need to sweep more than 12 beams with the Multiverse). Sweeping 12 beams ($K = 12$) imposes the latency of 0.4482 ms and 1.87 ms with 802.11ad and 5G-NR standards. Thus, we conclude that the Multiverse results

TABLE VIII

THE BEAM SELECTION TIME WITH EXHAUSTIVE SEARCH AND MULTIVERSE FOR THREE MMWAVE RADIOS WITH 34, 64, AND 256 BEAMS. IN THIS TABLE, THE FIRST AND SECOND NUMBERS DENOTES THE BEAM SELECTION TIME WITH 802.11AD AND 5G-NR STANDARDS, RESPECTIVELY. THE EFFECT OF IMPROVEMENT IN BEAM SELECTION TIME IS MORE PRONOUNCED WHEN THE CODEBOOK HAS MORE BEAMS

Radio	# of beams	Overhead with exhaustive search	Overhead with Multiverse
Talon (FLASH dataset)	34	1.27ms/20.31ms	0.4482ms/1.87ms
Talon (entire codebook)	64	2.54ms/24.992 ms	0.8964ms/3.74ms
NI radio [56]	256	9.56ms/144.992s	3.36ms/44.056s

TABLE IX

COMPARISON OF THE PROPOSED MULTIVERSE WITH THE RELEVANT STATE-OF-THE-ART MMWAVE BEAM SELECTION METHODS

Papers	Top-1 Acc. (%)	Modalities	Method	Validation data
Klautau <i>et al.</i> [58]	30	LiDAR	DL-based	Synthetic
Dias <i>et al.</i> [59]	20	LiDAR	DL-based	Synthetic
Xu <i>et al.</i> [60]	60	Camera	DL-based	Synthetic
Multiverse at the Edge (this paper)	66	Camera, LiDAR, GPS	DL-based Digital twin	Real

in 64.70% and 90.79% improvement in beam selection time compared to the 802.11ad standard and 5G-NR standards, respectively.

Note that the COTS mmWave radios typically have a larger codebook (than 34 beams), which results in a higher overhead. For example, the Talon AD7200 radio that are used for data collection in FLASH dataset actually support 64 beams. Another example is the National Instruments mmWave radio [55] that has a codebook with 16 beams, which also enables beamforming at both Tx and Rx for maximum efficiency; thus, consisting of total of 256 possible beam combinations. On the other hand, the Multiverse diminishes the number of beam by factor of 0.35 in FLASH dataset. In Tab. VIII we compare the beam selection time with exhaustive search and Multiverse for Talon and NI radios, considering the same scaling factor of 0.35. From this table, we conclude, with higher number of beam, the effect of reduction in the beam selection time is more pronounced. On the other hand, given that mmWaves provide high rates of up to 4.6 Gbps [56], a small improvement in the beam selection time (0.44 ms with 802.11ad and 18.44 ms with 5G-NR standard) may result in loss of up to $2.02 - 84.82\text{ Mb}$ of data per second, which in return may halt critical applications, such as video streaming and autonomous driving.

Observation 9: From Fig. 12, we observe that, by using Multiverse, the beam selection time is decreased by 67.70 – 90.79% compared to the 802.11ad and 5G-NR standards, respectively, that employ exhaustive search. The effect of reduction in beam selection time is more pronounced in larger codebooks. Moreover, we observe that the Multiverse with 10 beams outperforms the DL-only method with 20 beams in both beam selection time and accuracy, despite of being imposed of a more strict constraint.

J. Multiverse Against State-of-the-Art Multimodal Methods

In Tab. IX, we compare the performance of the Multiverse with the state-of-the-art mmWave beam selection methods by

Klautau et al. [57], Dias et al. [58], and Xu et al. [59]. We report the average top-1 accuracy for LOS and NLOS scenarios. For the Multiverse, we use the accuracies from 3-Reflection twin and consider the 2 dB threshold ($Acc_{(1,2)}$ in Eq. 10). We observe that the Multiverse outperforms the competing methods by 6-46% in accuracy. Unlike our work, other methods are only validated on synthetic data. The first two works use LiDAR sensors on the synthetically-generated Raymobtime dataset [60], and Xu et al. [59] generate images using Blender software [50] and use them to construct a LiDAR-like point cloud.

Observation 10: From Tab. IX, we observe that the Multiverse is superior to the state-of-the-art multimodal mmWave beam selection methods by 6 – 46%.

K. Beam Coherence Time

In vehicular networks, the maximum delay threshold in the mmWave band is characterized by the “beam coherence time” that is the duration of the time for which a beam is valid. In the FLASH dataset, the BS-vehicle communication coverage spans ~ 40 m [8]. We use the analysis from [61] to estimate the contact time (the total time that the vehicle is in the coverage range of the BS) as $T_{contact} = \frac{2h \tan(\frac{\phi_{BS}}{2})}{v_l}$. Here, ϕ_{BS} and h denote the coverage angle and height of the BS, 168.57° and 0.95 m, respectively, and v_l is the velocity of the vehicle. From this equation, we observe that a vehicle with higher speed has a lower beam coherence time. We estimate the minimum contact time between the vehicles and BS as 2.12 s, by taking into account the maximum speed of 20 mph in the FLASH dataset. Considering the 34 beams in the codebook, we estimate the beam coherence time as $2.12/34 = 62.4$ ms. As a result, each beam at the BS is approximately valid for 62.4 ms, which is the delay threshold for beam selection. Interestingly, we observe that the 5G-NR standard (which exploit the exhaustive sweeping as per Eq. 1) spends 32.55% of the coherence time for beam initialization. On the other hand, Multiverse occupies maximum 3% of the coherence time and the remaining time can be used for data transmission. The coherence time is even more strict, when the vehicle speed is higher, or the codebook has more beams. In the mmWave band, due to Gbps transmission rate, even small fractions of time result in an immense difference in the amount of exchanged data. Thus, although the traditional methods fit in the beam coherence time, they result in higher throughput loss than the Multiverse.

Observation 11: We observe that the 5G-NR standard spends 32.55% of the coherence time for beam initialization. On the other hand, Multiverse occupies 3% of the coherence time, and the remaining time can be used for the data transmission.

VII. FUTURE DIRECTIONS AND CHALLENGES

In this paper, we present a pioneering work on the digital twin field that offers different virtual world choices based on fidelity needs for mmWave beamforming. In the following, we provide future challenges and directions we plan on pursuing:

- **Computation Resources:** With ordinary computers today, ray tracing takes a long time to keep up with fast-pacing vehicular traffic. Thus, in today’s standards expensive computing sources are needed. However, as the computation power

for devices and the availability of such devices are in steady rise [62], more powerful computations will be ubiquitous in the future, making ray tracing possible in the order *ms*.

- **Re-ray Tracing:** As indicated previously, ray tracing is the bottleneck of the digital twin in terms of computation and timing, because each time a new environment is entered, the ray tracing needs to be run again. In this paper, we propose addressing this by creating beam dictionaries by constantly running the twin so that when an appropriate beam is requested, the twin can provide an optimal answer. However, in case of an environment that the digital twin has not seen or generated before, ray tracing needs to be run from scratch. In order to cut the ray tracing time, we will develop more efficient computation strategies so that the ray tracing from scratch is not needed.

- **Reconfigurable Networks:** Currently, possible Tx beam options depend on preset beam patterns and directions, which brings suboptimal beam selections, necessitating alternative ways to deliver signals to the Rx. Given that real-time fast beam steering are at a development stage [63], we will explore the computation feasibilities of beam generation that eliminates the dictionary dependency at Tx using digital twins. Additionally, reconfigurable intelligent surfaces (RIS), which aims to (i) redirect and (ii) change the phase of incoming rays so that they constructively aggregate at a target Rx, gains attention for digital twins in recent years [36], enabling reconfigurable networks. Accordingly, by constantly running the real world replica and updating the deployed ML models, the digital twins calculates optimal wireless parameters, such as standards, available bandwidth, beam features, beam directions, and RIS parameters. In future work, we will explore feasibility of reconfigurable networks using digital twins by extending our preliminary work [51].

- **Computation Load and Data Granularity Trade-off:** As we discussed in Sec. VI-C and VI-D, accuracy and computation load in twins are inversely proportional. Additionally, the level of detail in multimodal dataset transmitted to digital twins affects the beam selection performance by altering the virtual world precision, hence ray tracing outcome. In future work, we will explore the impact of data granularity on the beam selection accuracy and computation load, for which we will investigate alternative methods to mitigate, such as transfer learning.

- **Multiverse for Big Data Generation and Generalizable Learning Agents:** Thanks to the capabilities of WI [10] and the expansive details of the digital twin options in the Multiverse, in the future Multiverse could be used as a data generation factory, which will allow training deep learning models for smarter wireless decisions for various applications. For example, angle of ray arrivals, ray bouncing points, and associated power values could be used for optimal RIS placement; power delay profile, delay spread, RSSI values, and the features of reflection, diffraction, and transmission could be used for channel modeling without the hassle of setting complicated and expensive experiments; bit-error-rate (BER), RSSI, H-matrix, and SNR values could be used for MIMO and physical layer applications.

VIII. CONCLUSION

In this paper, we propose to expand the digital twin concept towards a Multiverse of twins and demonstrate its application for mmWave beam selection in V2X scenarios. The Multiverse

play a vital role when local DL methods are insufficient due to environment changes, and the exhaustive beam search is difficult due to mobility. Each twin in the Multiverse, created in Wireless InSite simulator with different settings, captures the real world with different cost/fidelity trade-off. Accordingly, each twin prepares a beam selection dictionary for a quick reference when the environment change is detected. We validate the beam selection decisions by the Multiverse through the experimental dataset, FLASH. Our evaluations show that this *Multiverse at the Edge* correctly predicts the top- K beams with upto 85.22% accuracy. Our Multiverse-based method yields upto 90.79% improvement in beam selection time compared to 802.11ad standard.

REFERENCES

- [1] H. X. Nguyen, R. Trestian, D. To, and M. Tatipamula, "Digital twin for 5G and beyond," *IEEE Commun. Mag.*, vol. 59, no. 2, pp. 10–15, Feb. 2021.
- [2] L. U. Khan, Z. Han, W. Saad, E. Hossain, M. Guizani, and C. S. Hong, "Digital twin of wireless systems: Overview, taxonomy, challenges, and opportunities," *IEEE Commun. Surveys Tuts.*, vol. 24, no. 4, pp. 2230–2254, 4th Quart., 2022.
- [3] L. U. Khan, W. Saad, D. Niyato, Z. Han, and C. S. Hong, "Digital-twin-enabled 6G: Vision, architectural trends, and future directions," *IEEE Commun. Mag.*, vol. 60, no. 1, pp. 74–80, Jan. 2022.
- [4] D. Roy et al., "Going beyond RF: A survey on how AI-enabled multimodal beamforming will shape the NextG standard," *Comput. Netw.*, vol. 228, Jun. 2023, Art. no. 109729.
- [5] J. Gu, B. Salehi, D. Roy, and K. R. Chowdhury, "Multimodality in mmWave MIMO beam selection using deep learning: Datasets and challenges," *IEEE Commun. Mag.*, vol. 60, no. 11, pp. 36–41, Nov. 2022.
- [6] J. Gu, B. Salehi, S. Pimple, D. Roy, and K. R. Chowdhury, "TUNE: Transfer learning in unseen environments for V2X mmWave beam selection," in *Proc. ICC IEEE Int. Conf. Commun.*, May 2023, pp. 1658–1663.
- [7] B. Salehihikouei et al., "Deep learning on multimodal sensor data at the wireless edge for vehicular network," *IEEE Trans. Veh. Technol.*, vol. 71, no. 7, pp. 7639–7655, Jul. 2022.
- [8] B. Salehi, J. Gu, D. Roy, and K. Chowdhury, "FLASH: Federated learning for automated selection of high-band mmWave sectors," in *Proc. IEEE INFOCOM IEEE Conf. Comput. Commun.*, May 2022, pp. 1719–1728.
- [9] G. Pang, C. Shen, L. Cao, and A. V. D. Hengel, "Deep learning for anomaly detection: A review," *ACM Comput. Surveys*, vol. 54, no. 2, pp. 1–38, 2021.
- [10] *Wireless InSite, 3D Wireless Prediction Software*. Accessed: Oct. 10, 2022. [Online]. Available: <https://www.remcom.com/wireless-insite-em-propagation-software>
- [11] *FLASH Dataset*. [Online]. Available: <https://genesys-lab.org/multimodal-fusion-nextg-v2x-communications>
- [12] (2024). *GitHub Repository of Multiverse is Available At*. [Online]. Available: <https://github.com/batoolsalehi/Multiverse-TNET2024>
- [13] M. Arjovsky, L. Bottou, I. Gulrajani, and D. Lopez-Paz, "Invariant risk minimization," 2019, *arXiv:1907.02893*.
- [14] C. Shi, C. Zheng, K. Vafa, A. Feder, and D. Blei, "An invariant learning characterization of controlled text generation," in *Proc. NeurIPS Workshop Distribution Shifts: Connecting Methods Appl.*, 2022.
- [15] E. Rosenfeld, P. Ravikumar, and A. Risteski, "The risks of invariant risk minimization," 2020, *arXiv:2010.05761*.
- [16] D. Krueger et al., "Out-of-distribution generalization via risk extrapolation (REX)," in *Proc. Int. Conf. Mach. Learn.*, 2021, pp. 5815–5826.
- [17] J. Yang, K. Zhou, Y. Li, and Z. Liu, "Generalized out-of-distribution detection: A survey," 2021, *arXiv:2110.11334*.
- [18] S. Thrun and L. Pratt, "Learning to learn: Introduction and overview," in *Learning to Learn*. Boston, MA, USA: Springer, 1998, pp. 3–17.
- [19] J. Gu et al., "Meta-learning for image-guided millimeter-wave beam selection in unseen environments," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, 2023, pp. 1–5.
- [20] C. Finn, P. Abbeel, and S. Levine, "Model-agnostic meta-learning for fast adaptation of deep networks," in *Proc. Int. Conf. Mach. Learn.*, 2017, pp. 1126–1135.
- [21] A. Nichol, J. Achiam, and J. Schulman, "On first-order meta-learning algorithms," 2018, *arXiv:1803.02999*.
- [22] M. Andrychowicz et al., "Learning to learn by gradient descent by gradient descent," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 29, 2016, pp. 1–9.
- [23] M. Khodak, M.-F. F. Balcan, and A. S. Talwalkar, "Adaptive gradient-based meta-learning methods," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 32, 2019, pp. 1–12.
- [24] Y. Lee and S. Choi, "Gradient-based meta-learning with learned layer-wise metric and subspace," in *Proc. Int. Conf. Mach. Learn.*, Jul. 2018, pp. 2927–2936.
- [25] A. A. Rusu et al., "Meta-learning with latent embedding optimization," 2018, *arXiv:1807.05960*.
- [26] J. Snell, K. Swersky, and R. Zemel, "Prototypical networks for few-shot learning," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 30, 2017, pp. 1–11.
- [27] R. Dong, C. She, W. Hardjawana, Y. Li, and B. Vucetic, "Deep learning for hybrid 5G services in mobile edge computing systems: Learn from a digital twin," *IEEE Trans. Wireless Commun.*, vol. 18, no. 10, pp. 4692–4707, Oct. 2019.
- [28] Y. Lu, S. Maharjan, and Y. Zhang, "Adaptive edge association for wireless digital twin networks in 6G," *IEEE Internet Things J.*, vol. 8, no. 22, pp. 16219–16230, Nov. 2021.
- [29] W. Sun, S. Lian, H. Zhang, and Y. Zhang, "Lightweight digital twin and federated learning with distributed incentive in air-ground 6G networks," *IEEE Trans. Netw. Sci. Eng.*, vol. 10, no. 3, pp. 1214–1227, May 2023.
- [30] B. Vilas Boas, W. Zirwas, and M. Haardt, "Machine learning for CSI recreation in the digital twin based on prior knowledge," *IEEE Open J. Commun. Soc.*, vol. 3, pp. 1578–1591, 2022.
- [31] Y. Li, K. Li, L. Cheng, Q. Shi, and Z.-Q. Luo, "Digital twin-aided learning to enable robust beamforming: Limited feedback meets deep generative models," in *Proc. IEEE 22nd Int. Workshop Signal Process. Adv. Wireless Commun. (SPAWC)*, Sep. 2021, pp. 26–30.
- [32] N. Zeulin, A. Ponomarenko-Timofeev, O. Galinina, and S. Andreev, "ML-assisted beam selection via digital twins for time-sensitive industrial IoT," *IEEE Internet Things Mag.*, vol. 5, no. 1, pp. 36–40, Mar. 2022.
- [33] Y. Cui, T. Lv, W. Ni, and A. Jamalipour, "Digital twin-aided learning for managing reconfigurable intelligent surface-assisted, uplink, user-centric cell-free systems," *IEEE J. Sel. Areas Commun.*, vol. 41, no. 10, pp. 3175–3190, Oct. 2023.
- [34] T. Nitsche, C. Cordeiro, A. Flores, E. Knightly, E. Perahia, and J. Widmer, "IEEE 802.11ad: Directional 60 GHz communication for multi-gigabit-per-second Wi-Fi," *IEEE Commun. Mag.*, vol. 52, no. 12, pp. 132–141, Dec. 2014.
- [35] B. Salehi, M. Belgiovine, S. G. Sanchez, J. Dy, S. Ioannidis, and K. Chowdhury, "Machine learning on camera images for fast mmWave beamforming," in *Proc. IEEE 17th Int. Conf. Mobile Ad Hoc Sensor Syst. (MASS)*, Dec. 2020, pp. 338–346.
- [36] J. Hoydis et al., "Sionna RT: Differentiable ray tracing for radio propagation modeling," 2023, *arXiv:2303.11103*.
- [37] *OpenStreetMap*. Accessed: Jul. 30, 2022. [Online]. Available: <https://www.openstreetmap.org>
- [38] *Wireless InSite References*. Accessed: Nov. 2022. [Online]. Available: <https://www.remcom.com/wireless-insite-references>
- [39] L. Ahumada, E. Carreño, A. Anglès, and D. Schkolnik, "Shadowing correlation: Empirical results for mm-wave wireless links in urban street canyons," *IEEE Antennas Wireless Propag. Lett.*, vol. 17, no. 4, pp. 543–546, Apr. 2018.
- [40] F. Aghaei, H. B. Eldeeb, L. Bariah, S. Muhaidat, and M. Uysal, "Comparative characterization of indoor VLC and MMW communications via ray tracing simulations," *IEEE Access*, vol. 11, pp. 90345–90357, 2023.
- [41] C. B. Barneto et al., "Millimeter-wave mobile sensing and environment mapping: Models, algorithms and validation," *IEEE Trans. Veh. Technol.*, vol. 71, no. 4, pp. 3900–3916, Apr. 2022.
- [42] H. J. Liebe and D. H. Layton, "Millimeter-wave properties of the atmosphere: Laboratory studies and propagation modeling," NASA STI/Recon, Washington, DC, USA, Tech. Rep. 88, 1987, p. 21387.
- [43] C. N. Barati, S. Dutta, S. Rangan, and A. Sabharwal, "Energy and latency of beamforming architectures for initial access in mmWave wireless networks," *J. Indian Inst. Sci.*, vol. 100, no. 2, pp. 281–302, Apr. 2020.
- [44] S. Subramaniam, T. Palpanas, D. Papadopoulos, V. Kalogeraki, and D. Gunopulos, "Online outlier detection in sensor data using non-parametric models," in *Proc. 32nd Int. Conf. Very Large Data Bases*, 2006, pp. 187–198.

- [45] F. Angiulli and C. Pizzuti, "Fast outlier detection in high dimensional spaces," in *Proc. Eur. Conf. Princ. Data Mining Knowl. Discovery*. Berlin, Germany: Springer, 2002.
- [46] J. Zhang and H. Wang, "Detecting outlying subspaces for high-dimensional data: The new task, algorithms, and performance," *Knowl. Inf. Syst.*, vol. 10, no. 3, pp. 333–355, Oct. 2006.
- [47] M. Markou and S. Singh, "Novelty detection: A review—Part 2: Neural network based approaches," *Signal Process.*, vol. 83, no. 12, pp. 2499–2521, 2003.
- [48] A. Gritsenko, Z. Wang, T. Jian, J. Dy, K. Chowdhury, and S. Ioannidis, "Finding a 'new' needle in the haystack: Unseen radio detection in large populations using deep learning," in *Proc. IEEE Int. Symp. Dyn. Spectr. Access Netw. (DySPAN)*, 2019, pp. 1–10.
- [49] D. Steinmetzer, D. Wegemer, M. Schulz, J. Widmer, and M. Hollick, "Compressive millimeter-wave sector selection in off-the-shelf IEEE 802.11ad devices," in *Proc. Int. Conf. Emerg. Netw. Experiments Technol. (CoNEXT)*, 2017, pp. 414–425.
- [50] *Blender*. Accessed: Jul. 30, 2022. [Online]. Available: <https://www.blender.org>
- [51] U. Demir, S. Pradhan, R. Kumahia, D. Roy, S. Ioannidis, and K. Chowdhury, "Digital twins for maintaining QoS in programmable vehicular networks," *IEEE Netw.*, vol. 37, no. 4, pp. 208–214, Jul. 2023.
- [52] *WiFi Routers | AD7200*. Accessed: Nov. 11, 2022. [Online]. Available: <https://www.tp-link.com/us/home-networking/wifi-router/ad7200/>
- [53] *Legacy Measurements*. Accessed: Nov. 11, 2022. [Online]. Available: https://github.com/seemoo-lab/talon-sector-patterns/tree/master/legacy_measurements
- [54] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 770–778.
- [55] *Ni Radio*. [Online]. Available: <https://www.ni.com/en/shop/wireless-design-test/what-is-mmwave-transceiver-system.html>
- [56] *802.11ad*. Accessed: Oct. 3, 2023. [Online]. Available: <https://www.qualcomm.com/products/features/80211ad>
- [57] A. Klautau, N. Gonzalez-Prelcic, and R. W. Heath, "LiDAR data for deep learning-based mmWave beam-selection," *IEEE Wireless Commun. Lett.*, vol. 8, no. 3, pp. 909–912, Jun. 2019.
- [58] M. Dias, A. Klautau, N. González-Prelcic, and R. W. Heath, "Position and LiDAR-aided mmWave beam selection using deep learning," in *Proc. IEEE 20th Int. Workshop Signal Process. Adv. Wireless Commun. (SPAWC)*, Jul. 2019, pp. 1–5.
- [59] W. Xu, F. Gao, S. Jin, and A. Alkhateeb, "3D scene-based beam selection for mmWave communications," *IEEE Wireless Commun. Lett.*, vol. 9, no. 11, pp. 1850–1854, Nov. 2020.
- [60] A. Klautau, P. Batista, N. González-Prelcic, Y. Wang, and R. W. Heath, "5G MIMO data for machine learning: Application to beam-selection using deep learning," in *Proc. Inf. Theory Appl. Workshop (ITA)*, Feb. 2018, pp. 1–9.
- [61] G. R. Muns, K. V. Mishra, C. B. Guerra, Y. C. Eldar, and K. R. Chowdhury, "Beam alignment and tracking for autonomous vehicular communication using IEEE 802.11 ad-based radar," in *Proc. IEEE Conf. Commun. Workshops (INFOCOM WKSHPS)*, 2019, pp. 535–540.
- [62] Q. Guo, F. Tang, T. K. Rodrigues, and N. Kato, "Five disruptive technologies in 6G to support digital twin networks," *IEEE Wireless Commun.*, vol. 31, no. 1, pp. 149–155, Feb. 2023.
- [63] C. U. Bas et al., "Real-time millimeter-wave MIMO channel sounder for dynamic directional measurements," *IEEE Trans. Veh. Technol.*, vol. 68, no. 9, pp. 8775–8789, Sep. 2019.



Batool Salehi received the M.S. degree in electrical engineering from the University of Tehran, Iran, in 2019. She is currently pursuing the Ph.D. degree in computer engineering with Northeastern University under the supervision of Prof. K. Chowdhury. Her current research focuses on mmWave beamforming, the Internet of Things, and the application of machine learning in the domain of wireless communication.



Utku Demir received the B.S. and M.S. degrees in electrical and electronics engineering from Koç University, Istanbul, working with Sinem Çöleri in 2012 and 2014, respectively, and the M.S. and Ph.D. degrees in electrical engineering from the University of Rochester, working with Wendi Heinzelman, in 2017 and 2020, respectively. He was an EAI Post-Doctoral Research Fellow with the GENESYS Laboratory, Department of Electrical and Computer Engineering, Northeastern University, directed by Kaushik Chowdhury. In 2024, he joined Nexcepta Inc., as a Research Scientist. His research interests lie in the areas of wireless communications, mobile networks, signal processing, and machine learning.



Debashri Roy (Senior Member, IEEE) received the M.S. and Ph.D. degrees in computer science from the University of Central Florida, USA, in 2018 and 2020, respectively. She is currently an Assistant Professor with The University of Texas Arlington. Her research interests are in the areas of AI/ML enabled technologies in wireless communication, multimodal data fusion, nextG networks, and networked systems.



Suyash Pradhan is currently pursuing the M.S. degree in electrical and computer engineering with Northeastern University. He is also a Research Assistant with the GENESYS Laboratory under the supervision of Prof. Kaushik Chowdhury. His research interests revolve around multimodal data fusion, computer vision, and machine learning for enabling Internet of Things (IoT) applications.



Jennifer Dy (Member, IEEE) is currently a Professor with the Department of Electrical and Computer Engineering, Northeastern University, Boston, MA, USA, where she first joined as a Faculty Member in 2002. Her research spans both fundamental research in machine learning and their application to biomedical imaging, health, science, and engineering, with research contributions in unsupervised learning, dimensionality reduction, feature selection, learning from uncertain experts, active learning, Bayesian models, and deep representations.



Stratis Ioannidis (Member, IEEE) is currently a Professor with the Electrical and Computer Engineering Department, Northeastern University, Boston, MA, USA, where he also holds a courtesy appointment with the Khoury College of Computer Sciences. Prior to joining Northeastern University, he was a Research Scientist with the Technicolor Research Centers, Paris, France, and Palo Alto, CA, USA, and the Yahoo Laboratories, Sunnyvale, CA, USA. His research interests span machine learning, distributed systems, networking, optimization, and privacy.



Kaushik Chowdhury (Fellow, IEEE) received the M.S. degree from the University of Cincinnati, Cincinnati, OH, USA, in 2006, and the Ph.D. degree from Georgia Institute of Technology, Atlanta, GA, USA, in 2009. He is currently a Professor with the Department of Electrical and Computer Engineering, Northeastern University, Boston, MA, USA. His current research interests include dynamic spectrum access, applied machine learning for wireless, and networked robotics. He was a recipient of the Presidential Early Career Award for Scientists and Engineers (PECASE) in 2017, ONR Director of Research Early Career Award in 2016, and NSF CAREER Award in 2015.