

Deep Learning Predictive Band Switching in Wireless Networks

Faris B. Mismar, Ahmad AlAmmouri, Ahmed Alkhateeb, Jeffrey G. Andrews,
and Brian L. Evans

Abstract

In cellular systems, the user equipment (UE) can request a change in the frequency band when its rate drops below a threshold on the current band. The UE is then instructed by the base station (BS) to measure the quality of candidate bands, which requires a *measurement gap* in the data transmission, thus lowering the data rate. We propose an online-learning based band switching approach that does not require any measurement gap. Our proposed classifier-based band switching policy instead exploits spatial and spectral correlation between radio frequency signals in different bands based on knowledge of the UE location. We focus on switching between a lower (e.g., 3.5 GHz) band and a millimeter wave band (e.g., 28 GHz), and design and evaluate two classification models that are trained on a ray-tracing dataset. A key insight is that measurement gaps are overkill, in that only the relative order of the bands is necessary for band selection, rather than a full channel estimate. Our proposed machine learning-based policies achieve roughly 30% improvement in mean effective rates over those of the industry standard policy, while achieving misclassification errors well below 0.5% and maintaining resilience against blockage uncertainty.

I. INTRODUCTION

With each successive cellular standard using a rapidly increasing number of different frequency bands in different parts of the spectrum, the band selection problem has become ever more complicated. In particular, user equipment (UEs) would like to use the band or bands that

F. B. Mismar, A. AlAmmouri, J. G. Andrews, and B. L. Evans are with the Wireless Networking and Communications Group, Dept. of Electrical and Comp. Eng., The University of Texas at Austin, Austin, TX, 78712, USA. e-mail: {faris.mismar, alammouri}@utexas.edu and {jandrews, bevans}@ece.utexas.edu. A. Alkhateeb is with the School of Electrical, Computer and Energy Engineering at Arizona State University, Tempe, AZ 85287, USA. email: alkhateeb@asu.edu.

A preliminary version of this work was presented at the 2018 International Conference on Communications Workshops [1].

maximize their quality of experience (QoE), which is highly correlated to their achieved data rate. The choice of the optimal frequency band can be challenging. On the one hand, lower frequency bands generally have more benign propagation properties and thus produce higher signal to noise ratios (SNRs), but higher frequency bands such as millimeter wave (mmWave) offer much higher bandwidth as well as beamforming gains and will typically be more lightly loaded. So, if the SNR on a mmWave band is acceptable, it is likely to provide a much higher data rate than a lower band and a UE would usually benefit from being efficiently switched over to the mmWave band. Similarly, if coverage is lost on the mmWave band, the UE should be quickly switched back to the lower frequency band.

Despite its increasing importance, the signaling procedure for band switching has seen only incremental changes over the evolution of multiple successive 3GPP standards [2], [3]. This signaling (or *control plane*) procedure is shown in Fig. 1 and described as follows. If the received power at the UE drops below a certain threshold on its current frequency band, call it f_j , it requests a band switch from its serving base station (BS). This request is followed by a *measurement gap*, where the data (or *user plane*) flow is stopped to allow the user to tune its reception circuitry to the frequency of the target band, call it $f_{j'}, j' \neq j$, to measure the channel. The industry standards justify ceasing the data flow to preserve the UE battery. After obtaining the measurements, the user reports them back to the BS. The BS estimates, based on the measurements, whether the user would benefit from switching to $f_{j'}$ or not, and hence, grants or denies the request. A key issue with the aforementioned procedure is its dependence on the measurement gap which despite being a control plane procedure causes interruption in the flow in the data plane and reduces the user overall throughput. The 3GPP standards also introduced mobile load balancing (MLB) as a means of transferring traffic served by a congested BS to nearby BSs that have spare resources [4]. However, MLB requires periodic communication between BSs about their resources, introducing a significant overhead. MLB is also triggered by the BS desire to relieve its congestion, while the band switch is triggered when the UE desires to maximize its QoE.

It would be desirable to introduce a reliable method that can support the band switch procedure without interrupting the user data flow by measurement gaps. The aim of this paper is to propose a novel online-learning based gap-free algorithm for band switching that utilizes the spatial and spectral correlations over different frequency bands along with the previous band switching requests and decisions for nearby users. This online-learning based algorithm adapts to the

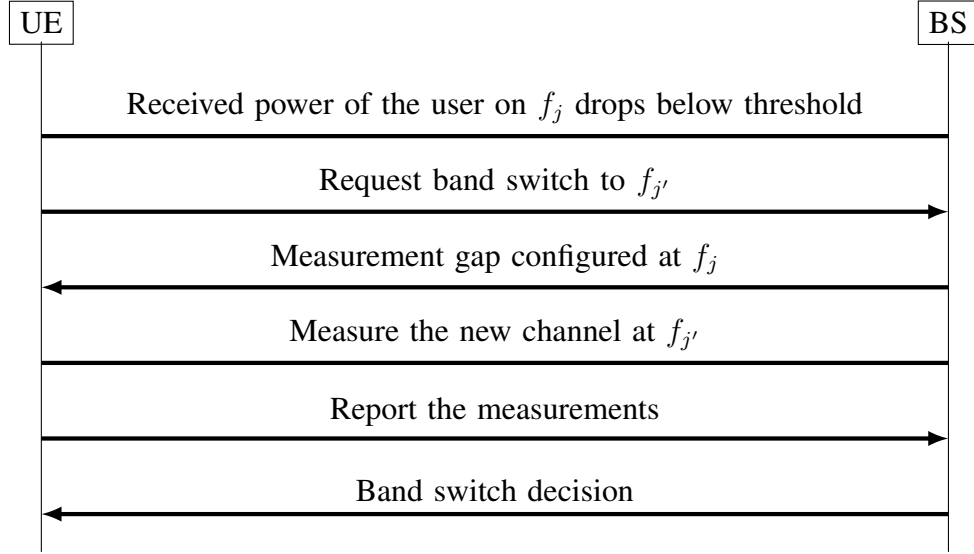


Fig. 1. The band switch procedure between frequency bands in one base station (BS) in the downlink direction.

changes in the environment as experienced by the users. Precisely, we propose a predictive algorithm, based on deep neural network (DNN) classifiers, which allow the BS to decide whether to grant or deny the band switching requests without the need for measurement gaps.

A. Related Work

Predicting the success of a band switch from one frequency band to another without explicitly measuring the channel at the target frequency band falls under the genre of problems commonly referred to as: “channel estimation using out-of-band information” [5]. In the simplest form of this problem, there are forward and backward (downlink and uplink) links occupying the same frequency bands at different time slots. In this case, we can use channel reciprocity [6] to estimate the channel of the backward link using the measurements on the forward link, or vice versa. Even with a separation of frequency duplex bands on the order of ten megahertz, a spatial correlation between the signals on the two frequency bands still exists due to the common propagation paths, blockages, and reflectors [6], [7]. Interestingly, the spatial correlation between two frequency bands that are separated by tens of gigahertz still exists [8]. However, it cannot be directly used to accurately estimate the channel on one frequency band by only using the measurements from the other, but it can be used to aid the channel estimation and reduce its complexity. For example, this correlation was exploited in [9]–[12] for cell discovery, channel covariance estimation, and beam selection in mmWave bands using sub-6 GHz measurements. In the case

of band switching, we are not interested in accurate channel estimation since the objective is not to use the estimate in decoding the messages, beam selection, or precoder/combiner design used in multiple-input multiple-output (MIMO) communications. Instead, our goal is much simpler: *ranking* the downlink channel quality of the two frequency bands or technologies.

The major challenge in exploiting the spatial correlation between frequency bands is the lack of accurate mathematical models that describe how the channel changes across these frequencies (or technologies). This challenge makes a data-driven or a machine learning (ML) approach more attractive to follow and implement. With more publicly available datasets that are based on field-measurements or sophisticated ray-tracing simulations [13], [14], we expect the interest in this approach to dramatically increase. Nevertheless, the applications involving dual-band ray-tracing datasets with ML classification to study channels is a nascent research area.

Although relevant to dual-band resource management, the work in [15] did not address the impact of measurement gaps on UE data rates. It focused on granting resources to users at mmWave first, while we allow granting resources to both mmWave and sub-6 GHz simultaneously without any specific preference. Furthermore, statistical path loss models were used for both mmWave and sub-6 GHz bands which may be privy to the spatial and spectral correlation of channels that we otherwise capture using ray-tracing datasets.

The work in [16] studied only one type of 3GPP dual-band handovers, which we call the “legacy” policy later in this paper. However, similar to [15], the use of statistical path loss models voids the opportunity to exploit the correlation across bands; therefore insights about the performance of the various algorithms, including the second type of 3GPP dual-band handover algorithms—the “blind” policy—could not be derived. Furthermore, the objective was to improve energy efficiency through handover avoidance, unlike our proposed algorithm the objective of which is to improve the UE data rates by eliminating measurement gaps.

In [17] dual connectivity was studied. Dual connectivity requires a local coordinator to manage the traffic between the cells, unlike band switch procedures. As a result, a backhaul latency constraint between the BSs was imposed. Furthermore, empirical pathloss models were used. Multiple BSs with a single UE were simulated while our focus on a single BS with dual band and multiple UEs. The use of a single UE may prevent the employment of ML techniques due to the limited number of learning observations—a problem we avoid altogether through the use of a ray-tracing dataset of RF propagation paths from UEs to a BS. Moreover, a band switch time-to-trigger mechanism was introduced in [17], whereby the band switch is only granted after

the band switch criterion is fulfilled for a period of time. This, unlike our proposed approach, introduces further latency to the band switch procedure [18].

B. Contributions

In this paper, we provide an answer to the question whether a reliable band switch method exists to maximize the users' achievable data rates. Specifically, this paper makes the following contributions:

- 1) Motivate the use of deep learning in ranking the downlink channel quality of the two frequency bands—a mathematically intractable problem and a requirement for the band switching procedure.
- 2) Offer several insights about the different band switch policies and their respective impact on performance. Furthermore, we show how the choice of the band switch threshold can have adverse impacts on the performance.
- 3) Motivate a data-driven approach to band switching, where we use a ray-tracing dataset in deep learning.
- 4) Create a unified framework to describe the band switch policies in a single equation and use this equation to explain the various band switching policies and their relevant performance.

II. SYSTEM MODEL

In this section, we describe the adopted network and channel models.

A. Network Model

We consider a radio network comprised of one BS serving single-antenna user equipment (UEs) in an arbitrary association area. The BS has two frequency bands; one in the sub-6 range and one at mmWave. Note that we assume that sub-6 and mmWave BSs are co-located to minimize the financial costs of deployment. Moreover, the BS utilizes a different number of antennas for each frequency band. Let $j \in \{\text{sub-6}, \text{mmWave}\}$ denote the frequency band and let $N^{(j)}$ denote the number of antennas on the j^{th} band, then the received signal at the i^{th} UE from the BS at the j^{th} frequency band is

$$r_{(i,j)} = P_{\text{TX}}^{(j)} \mathbf{h}_{(i,j)}^* \mathbf{f}_{(i,j)} s_{(i,j)} + n_{(i,j)}, \quad (1)$$

where $P_{\text{TX}}^{(j)}$ is the transmit power of BS on the j^{th} frequency band, $\mathbf{h}_{(i,j)} \in \mathbb{C}^{N^{(j)} \times 1}$ is the channel vector, $\mathbf{f}_{(i,j)}$ is the beamforming vector, $s_{(i,j)}$ is the transmitted signal, and $n_{(i,j)} \sim \text{Normal}(0, \sigma_j^2)$ is the thermal noise computed over the bandwidth $B^{(j)}$ including the UE noise figure. We focus on codebook-based analog beamforming, where the beamforming vector is chosen from a pre-defined codebook $\mathcal{F}^{(j)}$ [19]. In this case, the BS chooses the optimal beamforming vector \mathbf{f}^* that maximizes the receive SNR from the i^{th} user on the j^{th} frequency band from the codebook $\mathcal{F}^{(j)}$

$$\mathbf{f}_{(i,j)}^* := \arg \max_{\mathbf{f}_{(i,j)} \in \mathcal{F}^{(j)}} |\mathbf{h}_{(i,j)}^* \mathbf{f}_{(i,j)}|^2. \quad (2)$$

Let the codebook size be denoted by $N_{\text{CB}}^{(j)}$ and assume that all codewords are normalized, i.e., $\|\mathbf{f}_{(i,j)}\|^2 = 1$. Based on this, the received SNR at time step t at the i^{th} UE on the j^{th} frequency band is

$$\gamma^{(i,j)}[t] = \frac{P_{\text{TX}}^{(j)}[t]}{\sigma_j^2} |\mathbf{h}_{(i,j)}^*[t] \mathbf{f}_{(i,j)}^*[t]|^2, \quad (3)$$

and the instantaneous achievable rate is

$$R^{(i,j)}[t] = B^{(j)} \log_2(1 + \gamma^{(i,j)}[t]), \quad (4)$$

where $B^{(j)}$ is the available bandwidth at the j^{th} frequency. Note that the rate in (4) does not include the overhead of switching to a different frequency band nor the beam training overhead. These overheads cause a loss in throughput, which is typically related to the coherence time of the channel and the frame length.

B. Channel Model

Here we discuss the channel coherence time, the beam training time, the band switching overhead, and the effective throughput.

Channel coherence time: Let the coherence time for sub-6 GHz and mmWave frequency bands be denoted as $T_C^{\text{sub-6}}$ and T_C^{mmWave} , respectively. The exact values depend on the environment, the antenna configuration, and the user movement. Hence, to maintain the generality of the framework, we do not assume specific values for the channel coherence times and we discuss our choices of the coherence times in Section VI, which is only needed to numerically evaluate the performance of the different algorithms.

Beam training time: For the beam training overhead, we define the training penalty per beam as T_{beam} . Thus, the total beam training time, T_B , is related to the number of all possible beams, which is the size of the codebook $\mathcal{F}^{(j)}$ in our case (i.e., $T_B = T_{\text{beam}} N_{\text{CB}}^{(j)}$).

Band switching overhead: At the beginning of each radio frame, the UE can request a band switch operation from its serving BS if it is not satisfied by its current signal quality. We are using two time slots per frame, and a time slot does not necessarily have an integer duration that is compliant with the standards. The radio frame duration is not constant and is equal to the channel coherence time. These band switches can only happen at the beginning in the first time slot with zero delay by design. The BS uses a certain policy to determine whether the change to a different frequency band should be granted or denied. However, there is a time penalty for the band switch request, which is used by the BS to take a decision regarding the user request. We denote this overhead by T_H , which is determined by the algorithm or the policy used in the BS to respond to the band switch request and the existence or absence of a measurement gap. The exact values of T_H are given in Section III, where we present different band switch policies.

Effective throughput: Using the previous definitions for the channel coherence time T_C , the beam training time T_B , and the band switch overhead T_H , we can define the instantaneous weight $w^{(i,j,k)}[t]$ for the i^{th} UE that is connected to the BS on the j^{th} frequency band at the time step t as

$$w^{(i,j,k)}[t] := \max \left(0, 1 - \frac{T_B^{(j)} + T_H^{(k)}}{T_C^{(j)}} \right) \quad (5)$$

which accounts for the prolonged band switching overhead as a result of longer measurement gaps, since the throughput cannot be below zero. Here, the j^{th} band is the band *after* the band switch decision is made, which is a new frequency band if the band switch was granted and the old frequency band if the band switch was denied. This enables us to compute the effective throughput for the said UE as

$$R_E^{(i,j,k)}[t] = w^{(i,j,k)}[t] R^{(i,j)}[t]. \quad (6)$$

After discussing the system model and providing the necessary definitions, we present the current policies discussed in the industry standards [2] for the BS to make band switch decisions in the next section.

III. BAND SWITCH POLICIES

A band switch policy has to answer the following two questions: (i) when should the UE request this band switch? (ii) what is the information needed by the BS to make a decision for the band switch request and how? The first is typically solved by a pre-defined rate threshold




Fig. 2. Legacy band switch timing diagram. The shaded gray rectangles represent the measurement gaps.

$r_{\text{threshold}}$, such that if the UE rate is below this threshold, it requests a band switch. For the second question, the standards specify two policies today [2]. These policies are the measurement-based legacy approach and the blind approach. We also discuss the optimal policy as a benchmark. This policy is optimal in that it does not require the UE to use a measurement gap, and therefore the UE throughput is at the highest possible given the channel conditions, as we show later in the section. To provide a unified framework for the different policies, we define the following decision variables: $x_{\text{br}}, y \in \{0, 1\}$, where $x_{\text{br}} = 1$ if the UE requests a band switch, and $x_{\text{br}} = 0$ otherwise, and $y = 1$ if the BS grants the band switch and $y = 0$ otherwise. It is understood that y is only defined if $x_{\text{br}} = 1$. Further, the threshold $r_{\text{threshold}}$ is defined for all policies except the optimal policy. It is set based on how soon the UE should request the band switch from the BS.

A. Legacy Policy

The legacy policy, also known as the measurement-based policy, is shown in Fig. 2. When the user throughput is below the threshold $r_{\text{threshold}}$, it requests a band switch from the BS and it stops its transmission to measure the channel at the desired frequency band. After measuring the downlink channel, the user reports the measurements back to the BS, which decides whether to grant or deny the band switch request based on the measurements provided by the user. The measurement gap duration, denoted by T_G , is set to be a fraction of the coherence time [2]

$$T_G := \rho T_C, \quad 0 < \rho \leq 1. \quad (7)$$

Further, if we denote the overhead due to a band switch signaling request and its decision response as $\beta > 0$, then the band switch time overhead $T_H^{(\text{legacy})}$ is equal to $T_G + \beta$. By using this policy, the BS can make an informed decision regarding the band switch using the rates from both bands, which guarantees a certain QoE for the user. However, this comes at the expense

of the measurement gap, where the BS stops its transmission to the UE so it can measure the target channel, which causes an interruption in the data flow and reduces the UE throughput.

By employing this policy, one of three scenarios are possible at the beginning of each frame: (i) the UE does not request a band switch, which happens if its current rate is higher than the threshold; (ii) the UE requests a band switch and it is granted by the BS, which happens if the user's current rate is lower than the threshold, and the rate at the target band is higher than its current rate; or (iii) the UE requests a band switch, but it is denied, which happens if the UE current rate is lower than the threshold, and the rate at the target band is lower than its current rate. For the legacy policy, the decision variables are defined as

$$x_{\text{br}}^{(i)}[t] = \mathbb{1}[(R^{(i,j)}[t] < r_{\text{threshold}})], \quad \forall i, \quad (8)$$

$$y^{(i)}[t] = \mathbb{1}[(\hat{R}^{(i,j')}[t] > R^{(i,j)}[t])], \quad \forall i, \quad (9)$$

where j is the current serving BS, j' is the target BS, and $\hat{R}^{(i,j')}$ is the estimated rate the UE would get if the band switch were granted.

B. Blind Policy

Similar to the legacy policy, when the UE throughput is below the threshold, it requests a band switch from the BS. However, in this policy, the BS instructs the UE to band switch to a different band without any need for a measurement gap. Given the nature of this band switch approach, if the SNR is worse at the target frequency, the throughput drops significantly. Hence, although the measurement gap is eliminated, the BS cannot guarantee the user a higher throughput after the band switch, which causes a low QoE for the user. The decision variables in this case are as follows

$$x_{\text{br}}^{(i)}[t] = \mathbb{1}[(R^{(i,j)}[t] < r_{\text{threshold}})], \quad \forall i, \quad (10)$$

$$y^{(i)}[t] = 1, \quad \forall i, \quad (11)$$

since the band switch requests are always granted by the BS. Here, $T_H^{(\text{blind})} = \beta$, which is only for the signaling overhead since there is no measurement gap requirement.

C. Optimal

To define an upper bound for the various band switch policies, we define the optimal policy to be the one where the BS knows the instantaneous quality of the channels of the different

bands perfectly, so there is no need for a measurement gap. Hence, it asks the user to switch to a different band if the target rate is higher than its current rate. It also eliminates the need for a pre-defined rate threshold, since the band switch request and decision are combined and executed at the beginning of each frame by the BS. Based on this, the optimal effective throughput in this case is given by

$$R_E^{*(i)}[t] = \max_{j \in \{\text{sub-6, mmWave}\}} \left(1 - \frac{T_B^{(j)}}{T_C^{(j)}} \right) R^{(i,j)}[t]. \quad (12)$$

Finally, the decision variables can be written as

$$x_{\text{br}}^{(i)}[t] = 1, \quad \forall i, \quad (13)$$

$$y^{(i)}[t] = \mathbb{1}[(\hat{R}^{(i,j')}[t] > R^{(i,j)}[t])], \quad \forall i. \quad (14)$$

D. Overhead of Band Switching

Based on the previous discussion, besides the standards-imposed signaling overhead requirement of β , which is common across all policies, only the legacy policy causes band switch overhead. This overhead is equal to the measurement gap. Hence, $T_H^{(i,j,k)} = T_G^{(i,j,k)} + \beta$ for $k \in \{\text{legacy}\}$ and $T_H^{(i,j,k)} = \beta, k \in \{\text{blind, optimal}\}$. Note that deterioration in user throughput due to band switch overhead in the legacy policy drives the setting of the pre-defined threshold to lower values to avoid spending long times in measurement gaps. When this threshold is set low, the signal quality has to be bad for the user to request the band switch. This prevents the user from utilizing possibly better channels on other frequency bands or technologies. Moreover, with the introduction of mmWave frequency bands [20] in the fifth generation of wireless communications (5G) standard, the design of the band switch procedure becomes yet more critical since radio frequency signals at mmWave bands are more sensitive to blockages by various objects. For example, it was shown in [21] that the antenna gains on the mmWave bands can suffer from up to 25 dB attenuation due to the user hand grip on the mobile device and it varies with the different hand grips. Hence, under large blockage losses, the user would benefit from a fast transition to other frequency bands, and relying on the measurement gaps does not help.

The objective of this work is to propose a new band switch policy that eliminates the measurement gap, as in the blind policy, but ensures a certain QoE as in the legacy policy. Our policy relies on deep learning classification, which we introduce in the next section, before discussing the details of our algorithm.

TABLE I
DEEP NEURAL NETWORK CLASSIFIER LEARNING FEATURES

	Parameter	Type	Description
\mathbf{x}_0	Bias term	Integer	This is equal to unity.
\mathbf{x}_1	$R_E^{(\text{sub-6})}$	Float	Effective achievable rate in the sub-6 GHz band.
\mathbf{x}_2	$R_E^{(\text{mmWave})}$	Float	Effective achievable rate in the mmWave band.
\mathbf{x}_3	Source technology	Binary	(= 1 for sub-6 and = 0 for mmWave).
$(\mathbf{x}_4, \mathbf{x}_5, \mathbf{x}_6)$	Coordinates	Float	The coordinates of the UEs distance from the base station. based on the coordinates of i^{th} UE (i.e., $d_i[t]$).
\mathbf{x}_7	Band switch requested	Binary	UE requested band switch ($x_{\text{br}}^{(i)} = 1$)?
\mathbf{y}	Band switch decision	Binary	UE band switch request is granted ($y^{(i)} = 1$)?

IV. PROPOSED POLICY

For our proposed policy, we define \mathcal{N} as the set of users within the serving BS area. This set has a cardinality $|\mathcal{N}| := T_{\text{simulation}} N_{\text{UE}}$. Here, $T_{\text{simulation}}$ is the total simulation time and is equal to the number of stratified samples the total set \mathcal{N} is divided into (i.e., partitions). Also, N_{UE} is the number of active UEs within the BS serving area. Further, we define a set of users \mathcal{U} (also within the serving BS area) with an objective to improve the band switch performance using the proposed policy, but without a measurement gap.

To achieve this objective, we use the locations and measurements of the set $\mathcal{N} \setminus \mathcal{U}$, also served by this base station, such that $\mathcal{U} \subset \mathcal{N}$. This is achieved by exploiting the spatial and spectral correlation of the channels over different frequency bands and different locations.

We keep the minimum threshold criteria used in the legacy and blind policies, where the UE requests a band switch if its rate drops below a pre-defined threshold, $r_{\text{threshold}}$. Then the BS grants the band switch if the estimated rate at the target frequency is higher than the UE current rate. The difference is that the BS does not ask the UE to interrupt its transmission to measure the channel, but instead uses a machine learning approach to estimate the rate at the target frequency. Hence, the decision variables x_{br}, y are defined in the same way as in the legacy approach given in (8) and (9).

The major challenge in this approach, which relies on exploiting the spatial and spectral correlation between the channels, is the lack of accurate mathematical models that capture these correlations. Hence, we propose the use of DNN classifiers in the solution of our problem.

We use the classifiers to predict the band switch of other locations in the proximity of these

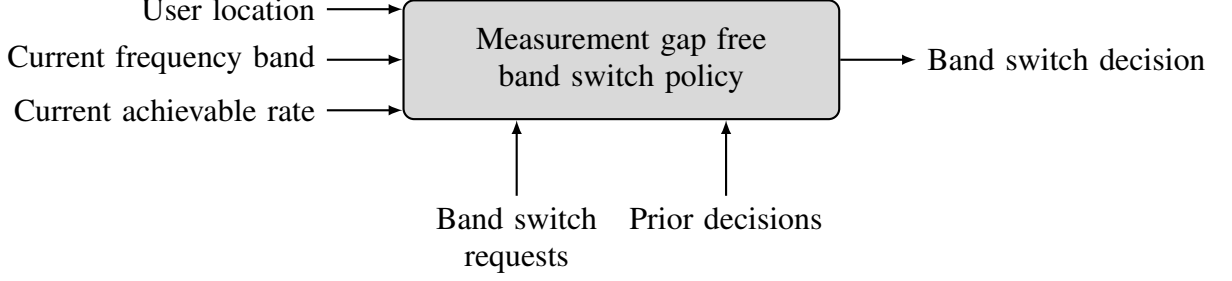


Fig. 3. Illustration of our proposed algorithm. The list of learning features is shown in Table I.

learned locations. Our algorithm is illustrated in Fig. 3 and is specified in Algorithm 1. The main steps of Algorithm 1 are as follows:

- Construct the dataset for the UEs within the serving BS area in the current time step t , which contains the rates from the spatially correlated wireless channels and band switch decisions and the stratified sets¹ $\mathcal{N}^{(t)}$ and $\mathcal{U}^{(t)}$ (i.e., partitions).
- Train the classifier using a randomly sampled subset of this dataset (or the *learning* set). The size of this subset is given by $|\mathcal{N}^{(t)} \setminus \mathcal{U}^{(t)}| := \lceil (1 - q_{\text{exploitation}}) N_{\text{UE}}^{(t)} \rceil$, where $q_{\text{exploitation}}$ is a fraction of the total dataset size and $N_{\text{UE}}^{(t)}$ is the stratified subset size (UE count).
- Use this classifier to predict the band switch for the set of UEs belonging to the subset $\mathcal{U}^{(t)}$.

Classifier choice: For the classifier we consider two options: DNN and extreme gradient boosting (XGBoost) trees. XGBoost is a scalable tree boosting system that achieves high classification performance [22]. DNN is a feed-forward architecture that uses layers of neurons of a given depth d and width w and can approximate arbitrary functions under assumptions on the activation functions [23]. An activation function defines the output of a neuron with respect to its inputs. A DNN optimizes a convex loss function through a learning rate $\eta > 0$. XGBoost optimizes an objective function containing a convex loss function (e.g., binary logistic loss) and a regularization term $\alpha \|\mathbf{w}\|_1 + \frac{1}{2} \lambda \|\mathbf{w}\|_2^2 + \gamma T$, where \mathbf{w} is the vector containing the leaf weights in the boosted tree, α and λ are the regularization terms for their respective norms, γ is the complexity control, and T is the number of leaves.

Classifier training: We train the hyperparameters of the classifiers using grid search and K -fold cross-validation. The list of learning features is shown in Table I. Let the feature matrix

¹Given that the set \mathcal{N} is stratified, it should be clear that $\bigcup_{t=1}^{T_{\text{simulation}}} \mathcal{N}^{(t)} = \mathcal{N}$ since $\mathcal{N}^{(t)}$ is disjoint.

Algorithm 1: Measurement gap free band switch policy

Input: Parameters listed in Tables II and III. Instantaneous rates by UE location per frequency band (or technology) of a given BS association area. Simulation time $T_{\text{simulation}}$. Set of all UEs \mathcal{N} and target UEs $\mathcal{U} \subset \mathcal{N}$, defined by $q_{\text{exploitation}}$.

Output: A vector $\hat{\mathbf{y}}$ containing a prediction whether the measurement gap-free band switch should be granted or denied for the set \mathcal{U} in the same BS association area, a confusion matrix \mathbf{C} , and the area under the receiver operating characteristic (ROC) curve for the set.

```

1 for  $t := 1$  to  $T_{\text{simulation}}$  do
2   Stratified sampling with no replacement from  $\mathcal{N}$  into  $\mathcal{N}^{(t)}$ . Let  $N_{\text{UE}}^{(t)} := |\mathcal{N}^{(t)}|$ .
3    $N_{\text{learning}}^{(t)} := \lceil (1 - q_{\text{exploitation}}) \cdot N_{\text{UE}}^{(t)} \rceil$ 
4   At random, sample with no replacement a subset of users from  $\mathcal{N}^{(t)}$  into a set  $\mathcal{U}^{(t)}$ .
5   Build the learning dataset  $[\mathbf{X} | \mathbf{y}]$  for UEs  $\{1, 2, \dots, N_{\text{learning}}^{(t)}\}$ , where  $\mathbf{X}$  is in Table I
      and  $\mathbf{y}$  is based on (9), both for all  $N_{\text{learning}}^{(t)}$  UEs.
6   Randomly split the data  $[\mathbf{X} | \mathbf{y}]$  into a training and a test data (using  $q_{\text{training}}$  split ratio).
7   Train the DNN classifier balancing the weights of classes  $\mathbf{y}$  in the training data and use
      grid search on  $K$ -fold cross-validation to tune the hyperparameters based on the
      binary cross-entropy loss function [23].
8   forall  $u \in \mathcal{U}^{(t)}$  do
9     Use the DNN classifier to predict  $\hat{y}^{(u)}$  based on  $\mathbf{X}^{(u)}$ .
10  end
11  Obtain  $\xi$  the area under the ROC curve for this model using  $\hat{\mathbf{y}} := [\hat{y}^{(u)}]$ .
12  Build the confusion matrix  $\mathbf{C}$  by observing  $\mathbf{y}$  and  $\hat{\mathbf{y}}$ .
13  Invalidate the DNN classifier.
14 end

```

be denoted by $\mathbf{X} \in \mathbb{R}^{N_{\text{UE}} \times p}$, $p > 2$. The industry standards require two features for the band switch decision, as we showed in Section III. The supervisory label vector is a column vector $\mathbf{y} \in \{0, 1\}^{N_{\text{UE}}}$, where 0 means the band switch was denied and 1 means granted based on (8).

Our proposed approach is shown in comparison to the legacy approach in Fig. 4 and it operates in two phases: a) learning phase and b) exploitation phase.

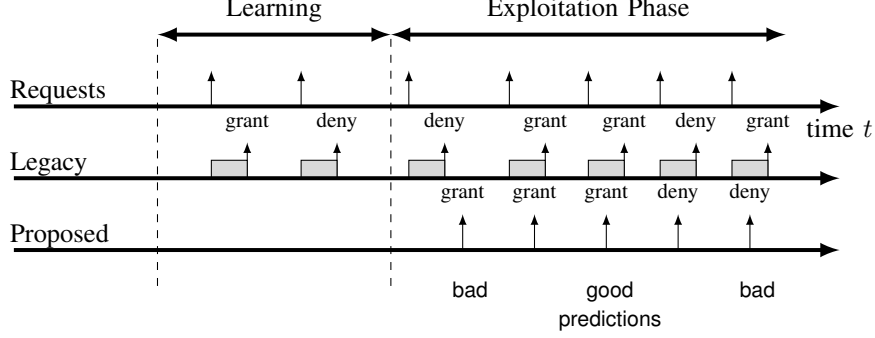


Fig. 4. Proposed band switch time diagram. The shaded gray rectangles represent the transmission gaps.

Learning phase: In the learning phase, the UE follows the legacy approach discussed earlier while the proposed algorithm stores the learning features \mathbf{X} and \mathbf{y} . Machine learning is then applied on this data to build a classifier that estimates band switch decisions but without the need for a measurement gap. During this phase, we let all UEs request band switches by setting \mathbf{x}_7 to unity (or $r_{\text{threshold}}$ to $+\infty$). This is in order for the classifier to learn the relationship between channels regardless of band switch requests. We use both DNN and XGBoost classification algorithms in the implementation of this phase and compute class weights for imbalanced the classes \mathbf{y} .

Exploitation phase: In the exploitation phase, the classifier uses prediction to eliminate the measurement gap for the set of UEs \mathcal{U} which were not used in the learning phase. The predicted decision either grants or denies the band switch from the j^{th} band. The exploitation phase essentially represents the *generalization capacity* of the classifier for the current radio frame.

Classifier invalidation: the invalidation (i.e., purging and retraining) of classifiers dealing with wireless channels in online-learning setting prevents changes in the channel state information from not being reflected onto the classifier [24]. Given that we only allow the band switching to take place in the beginning of the radio frame as stated earlier, this leaves the classifier with $(T_{\text{RF}} - 1)$ timeslots to realistically collect data and train. The number of measurements per radio frame is $|\mathcal{N} \setminus \mathcal{U}| = \lceil (1 - q_{\text{exploitation}}) \cdot |\mathcal{N}| \rceil$, after which the classifier must be invalidated.

V. PERFORMANCE MEASURES

In this section we describe our choices of the performance measures to benchmark our algorithm. These measures describe the performance of both the QoE of the users and the performance of the classifier.

A. Effective Achievable Rate

For the different policies discussed, we evaluate the effective rate of all the users in the network using (6). We are interested in the statistics of the effective achievable rates. Namely, the cumulative distribution function (CDF) and the mean.

B. Confusion Matrix

We define the misclassification count $E: 0 \leq E \leq n$ as the number of incorrectly predicted band switches during the exploitation phase. We build a confusion matrix $\mathbf{C} \in \mathbb{Z}_+^{2 \times 2}$ having the true and predicted band switch decision counts and write

$$E := \text{Tr}(\mathbf{J}\mathbf{C}^\top), \quad (15)$$

where \mathbf{J} is a 2×2 anti-diagonal identity matrix. The misclassification error μ can be derived by dividing E by $n := \lfloor q_{\text{exploitation}} N_{\text{UE}} \rfloor$.

C. ROC Area Under the Curve

The receiver operating characteristic (ROC) curve is a two-dimensional curve used to visualize classifiers based on the tradeoff between hit rates and false alarm rates. To compare the performance of classifiers, we reduce the ROC performance to a single scalar quantity known as the ROC area under the curve [25]. This area $\xi: 0 \leq \xi \leq 1.0$ where 0.5 means the classifier is as good as a random guess and 1.0 means it produces perfect prediction.

So far, we have discussed the different band switch policies, including our proposed policy. We have highlighted the desired performance measures we are interested in. In the next section, we discuss the data we use and how we construct in detail.

VI. DATASET CONSTRUCTION

To test the performance of the proposed algorithm, we rely on the DeepMIMO dataset [13]. The choice of this dataset is based on its use of accurate ray-tracing tools to generate spatially and spectrally correlated channels for specific scenarios. Hence, we avoid using oversimplified mathematical models that could lead to misleading results. A better choice would be to use a dataset that is based on actual field-measurements. However, such dataset is not available yet to the best of our knowledge, and is highly non-trivial to generate.

In the O1 outdoor scenario of the DeepMIMO dataset, the UEs are placed on a uniform grid on a main street for both the sub-6 GHz and mmWave frequency bands, where the BS uses OFDM and uniform planar array (UPA) antennas. The adopted O1 scenario is shown in Fig. 5.

A. Channel Coherence Time

The channel coherence time over which the channel remains highly correlated is known to be given by [26]

$$T_C(\alpha) \approx \frac{c}{f_c v_s \sin \alpha}, \quad (16)$$

where c is the speed of light, v_s is the speed of the UE, α is the angle between the direction of travel and the direction of the BS, $v_s \sin \alpha$ is the relative speed of the user with regards to the BS, and f_c is the center frequency. This equation has been widely used to measure the channel coherence time for the sub-6 GHz range, where omni-directional antennas are used. However, at mmWave, where directional antennas along with beamforming are employed to combat the high isotropic path loss, (16) does not accurately measure coherence time [27]. This is because by combining directional antenna arrays with beamforming, the signal power is focused on a beamwidth-defined angular space directed towards the UE location. Hence, only the variations in the channel within this angular space are relevant, which increases the channel coherence compared to (16). The coherence time of the beamformed channel, referred to as the beam coherence time, is given by [27]

$$T_C(\alpha) \approx \frac{D}{v_s \sin \alpha} \frac{\Theta}{2}, \quad (17)$$

where D is the Euclidean distance from the serving BS and Θ is the beamwidth of the beams used by serving BS (in radians). Since UEs are located at different locations with different distances to the BS, they have different coherence times. However, to maintain a fixed frame length for all users connected at the same band, we assume the cell-center beam coherence time (e.g., the 1st percentile). This conservative assumption is also motivated by the practical case where the BS may not have full knowledge of the UE parameters, such as their distance and accurate location. To sum up, we assume that the coherence time for sub-6 GHz and mmWave bands is given by

$$T_C^{\text{sub-6}} := \left(\frac{c}{f_c v_s \sin \alpha} \right)_{0.01}, \quad (18)$$

$$T_C^{\text{mmWave}} := \left(\frac{D}{v_s \sin \alpha} \frac{\Theta}{2} \right)_{0.01}, \quad (19)$$

respectively, where $(\mathcal{X})_{0.01}$ is the 1st percentile of the set \mathcal{X} . For the frame time, we set the frame duration to be equal to the channel coherence time for simplicity. Hence, the overheads for beam training and band switch, mentioned earlier, are related to a single parameter, which is the coherence time.

B. Band-Selective Blockage

Further, we choose to have occasional blockage in the mmWave frequency band. To generate this blockage using DeepMIMO, we generate two channels: one with blockage and the other without blockage. We further combine the mmWave channels into one by introducing a Bernoulli random variable for the i^{th} UE:

$$b_i \sim \text{Bernoulli}(p), \quad i = 1, 2, \dots, N_{\text{UE}} \quad (20)$$

$$\mathbf{h}^{(i)}[t] = b_i \mathbf{h}_b^{(i)}[t] + (1 - b_i) \mathbf{h}_{nb}^{(i)}[t], \quad (21)$$

where p is the blockage probability, $\mathbf{h}_b^{(i)}$ is the channel with blockage on the first mmWave path, and $\mathbf{h}_{nb}^{(i)}$ is the channel with no blockage, all for the i^{th} UE. Hence, some locations along the street are assumed to be blocked from the BS (non-line of sight), while others have a line of sight. To study the behavior of the proposed algorithm against uncertainty, we vary the blockage probability p in the simulation during the exploitation phase.

C. Analog Beamforming

We adopt a multi-antenna setup, where the BS employs a UPA of $M_y^{(j)}$ and $M_z^{(j)}$ antennas in the elevation and azimuth directions respectively at the j^{th} band. Therefore, we write the channel in (1) as $\mathbf{h} \in \mathbb{C}^{M_y^{(j)} M_z^{(j)} \times 1}$ in a vectorized form. In our implementation of analog beamforming, we focus on discrete Fourier transform (DFT) codebooks. We focus on DFT codebooks as they are a common practice for effective codebook design when the channels are spatially correlated [28]. Let the $M \times N_{\text{CB}}^{(j)}$ matrix $\mathbf{F}^{(j)}$ be the concatenation of M beamforming vectors in the codebook $\mathcal{F}^{(j)}$, then the matrix $\mathbf{F}^{(j)}$ is constructed as

$$\mathbf{F}^{(j)} = \mathbf{F}_z^{(j)} \otimes \mathbf{F}_y^{(j)} \quad (22)$$

where $\mathbf{F}_y^{(j)} \in \mathbb{C}^{M_y^{(j)} \times M_y^{(j)}}$ and $\mathbf{F}_z^{(j)} \in \mathbb{C}^{M_z^{(j)} \times M_z^{(j)}}$ concatenate the DFT codebook vectors in the y and z directions for the j^{th} frequency band. In the next section, we use this dataset to evaluate the performance of the proposed algorithm and compare it with the other algorithms discussed in Section III.



Fig. 5. Scenario O1 of the DeepMIMO dataset [13]. We use base station (BS) 3 and users on User Grid 1.

TABLE II
RADIO ENVIRONMENT PARAMETERS

Parameter	Value
Subcarrier bandwidth (sub-6, mmWave)	(180, 1800) kHz
Center frequency	(3.5, 28) GHz
UE noise figure	7 dB
DeepMIMO Scenario O1 Base Station	3
DeepMIMO Scenario O1 number of antennas (M_x, M_y, M_z)	(1, 64, 4)
DeepMIMO Scenario O1 OFDM limit	64
Band switch threshold rate for sub-6 GHz $r_{\text{threshold}}^{\text{sub-6}}$	1.72 Mbps
Band switch threshold rate for mmWave $r_{\text{threshold}}^{\text{mmWave}}$	7.00 Mbps
Measurement gap fraction of coherence time ρ	0.6

VII. SIMULATION RESULTS

In this section, we evaluate the proposed algorithm using the DeepMIMO dataset described in Section VI using the performance measures outlined in Section V.

A. Setup

The DNN and XGBoost classifier hyperparameters are both shown in Table III. As mentioned earlier, the users are placed on a uniform grid on a main street in the association area of this

TABLE III
HYPERPARAMETERS OF CLASSIFIERS USED IN THE IMPLEMENTATION OF OUR ALGORITHM

DNN		XGBoost	
Parameter	Value	Parameter	Value
Exploitation split $q_{\text{exploitation}}$	0.8	Exploitation split $q_{\text{exploitation}}$	0.8
K -fold cross-validation K	2	K -fold cross-validation K	2
Optimizer	[29]	ℓ_1 regularization term α	{0,1}
Learning rate η	0.05	ℓ_2 regularization term λ	{0,1}
Activation function	sigmoid	Complexity control term γ	{0,0.02,0.04}
Depth of neural network d	{1,3,5}	Sample weights	{0.5,0.7}
Width of the hidden layer w	{3,5,10}	Child weights	{0,10}

co-sited BS such that the i^{th} UE has the Cartesian coordinate (x_i, y_i, z_i) . This grid has an area of 550 m and a width of 35 m for an area of 19,250 square meters. The spacing between every two adjacent users in this uniform grid is 0.2 m. The height of all UEs $z_i = 2$ m is constant throughout the simulation. We set the mmWave channel blockage probability in (20) to $p = 0.4$ during the *learning* phase. Given that we choose not to perform oversampling in beamforming, $N_{\text{CB}} := M_y$. We set the beam training time $T_{\text{beam}} := 1 \mu\text{s}$ [27]. Further, in (18) and (19), we set $\alpha \sim \text{Uniform}(0, \pi)$ and $\Theta := 102/M_y$ following [27]. The users move at a vehicular speed $v_s = 50$ km/h within the BS association area every discrete time step t . There is a total of 54,480 users in the simulation divided into partitions of 5,448 users, hence $T_{\text{simulation}} = 10$ radio frames. In an attempt to find the *absolute* training dataset size that best maximizes the ROC area, we choose from $q_{\text{training}} \in \{1 \times 10^{-3}, 5 \times 10^{-3}, 7 \times 10^{-3}, 1 \times 10^{-2}, 3 \times 10^{-2}, 5 \times 10^{-2}, 7 \times 10^{-2}, 1 \times 10^{-1}, 3 \times 10^{-1}, 4 \times 10^{-1}, 5 \times 10^{-1}, 7 \times 10^{-1}\}$. With 12 OFDM subcarriers per physical resource block (PRB) and a subcarrier spacing of 15 kHz, we have the bandwidth $B = 180$ kHz per PRB. We use one PRB for the sub-6 GHz frequency band and ten PRBs for the mmWave frequency band. In other words, $B^{\text{sub-6}} = 180$ kHz and $B^{\text{mmWave}} = 1800$ kHz. Initially, the coherence times for sub-6 GHz and mmWave based on (18) and (19) are 6.17 and 19.16 ms respectively. This is justified due to the increased beamforming gain at mmWave, which slows down the time fluctuation of the channel [27], [30]. We choose the transmit energy of 0.1 W/Hz at mmWave and set the transmit energy at sub-6 GHz to 1 W/Hz. We set the measurement gap fraction ρ to 0.6 aligned with the industry standards of the gap duration per frame [31]. Further, we set the band switch thresholds as 1.72 Mbps and 7.00 Mbps for sub-6 GHz and

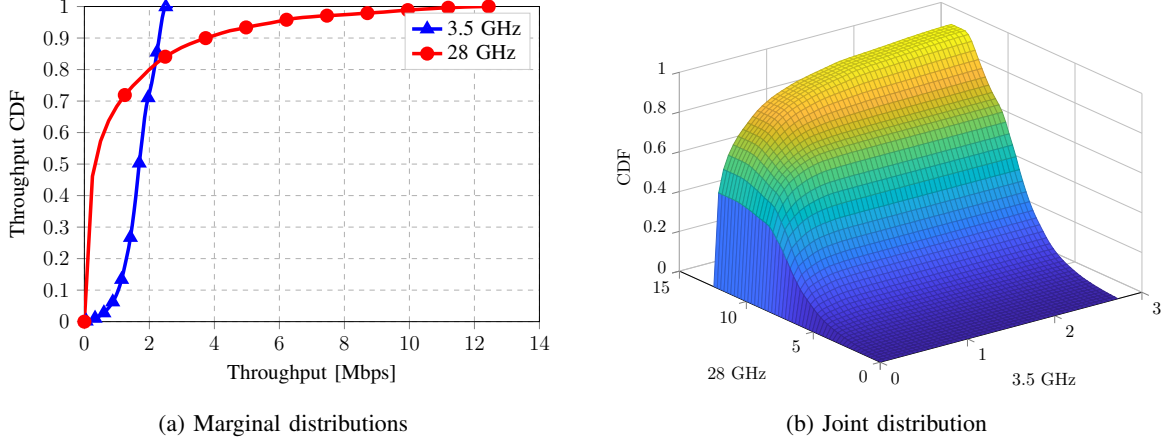


Fig. 6. Instantaneous throughput distributions for both sub-6 and mmWave frequency bands.

mmWave as the empirical means of the throughput distributions. The exploitation ratio $q_{\text{exploitation}}$ of 0.8 means that the total number of UEs that will use the trained model for band switching is $0.8 \times 54,480 = 43,584$ UEs. In the learning phase, we set the band switch request threshold to $+\infty$. This allows us to capture all the available spatial correlation information between the channels without any omission.

We simulate the radio environment (given in Table II) using three different scenarios:

- *Scenario A*: All users start in sub-6 and attempt to change band to mmWave.
- *Scenario B*: All users start in mmWave and attempt to change band to sub-6 GHz.
- *Scenario C*: 70% are in sub-6 and 30% are in mmWave.

We refer to the source code for the details of the implementation of this simulation [32]. Before presenting the results for these scenarios, we start with analyzing the raw data from the dataset. In Fig. 6, we show the distribution of the effective throughput over all users for the 3.5 GHz and the 28 GHz bands. From the marginal distributions, we can see that the effective throughput for the mmWave bands goes up to 12 Mbps, while it only goes up to 3 Mbps for the 3.5 GHz band. This is due to the large bandwidth that is available in the mmWave band. However, due to blockage, these high rates only occur with a small probability, since the two CDF curves cross at 0.8. Overall, the figure shows that the effective rate at mmWave can be very high, due to the large bandwidth, or very low due to the blockage. This wide range of rates motivates the optimized design of the band switch policy, since ineffective design can cause significant deterioration in the UE throughput, or can prevent the user from harvesting a high




Fig. 7. The distribution of the absolute difference between the rate on 3.5 GHz and 28 GHz.

rate from mmWave bands.

From the joint distribution, the general trend is that a higher throughput on 3.5 GHz means a higher throughput on 28 GHz. This is due to the correlation between the channels caused by common propagation paths, reflectors, and obstructions. By simple computations using the joint distributions, one can see that the 25% of the users can get higher throughput on the 28 GHz band. Intuitively, these are the users who do not suffer from blockage and are at a short distance from the BS. These users would benefit from operating at mmWave. To quantify this gain, we plot Fig. 7, which shows the distribution of the absolute difference between the throughput at 3.5 GHz and 28 GHz, $|\Delta R_E| := |R_E^{\text{sub-6}} - R_E^{\text{mmWave}}|$. Next, we analyze the performance of the different band switch policies discussed earlier.

B. Band Switch Policies

1) *Scenario A:* We start with Scenario A, where all the users start in 3.5 GHz band and can request a band switch to the 28 GHz band. The results are shown in Fig. 8, where we show the distribution of the effective throughput under different band switch policies. Starting with the legacy approach, the effect of the measurement gap is clear in the figure and results in a performance gap compared to the optimal policy, especially in the low rate regime, where the negative impact of the measurement gap is detrimental. Precisely, the effect is more dominant for users suffering from a low throughput at sub-6 GHz and requested a band switch from the

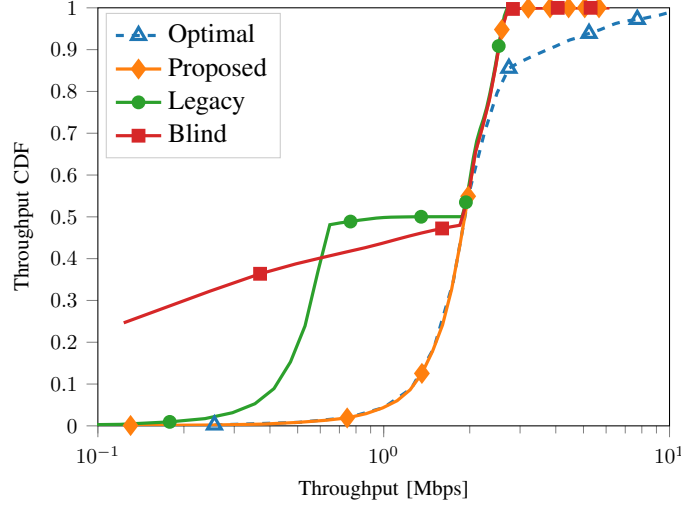


Fig. 8. The distribution of the effective throughput for Scenario A under different band switch polices.

BS, but their request got denied because the throughput at mmWave is also low, possibly due to blockage. Hence, having a measurement gap makes their throughput even worse. This can be observed from the curve for the small throughput regime. Further, we observe that there are more users in the small throughput range (i.e., less than 0.5 Mbps) in the blind policy than the legacy policy. This can be justified as follows: for the same class of users suffering from extremely low throughput in the sub-6 GHz band, but would not benefit from switching to the mmWave, the legacy policy instructs them to stay in the sub-6 GHz band at the expense of a measurement gap, while the blind policy switches these users to the mmWave band, which deteriorates their throughput even more. However, there is a point where the throughput at mmWave is around the same as the sub-6 GHz. At this point, the blind policy is more efficient since it has the advantage of not requiring a measurement gap.

For the proposed algorithm, Fig. 8 also shows that it has the best performance compared to the previous two; it is identical to the optimal in the low rate regime and identical to the other policies in the high rate regime. This is due to: (i) the elimination of the measurement gap, hence users with low throughput do not suffer more if their band switch request got denied as in the legacy approach, and (ii) the accurate band switch decisions, which prevents switching users to a band with low throughput as in the blind policy. Note that there is a performance gap between the three policies and the optimal in the high rate regime. This performance gap is due to the band switch threshold introduced in these policies, but missing from the optimal. Hence,




Fig. 9. The distribution of the effective throughput for Scenario B under different band switch policies.

TABLE IV
NORMALIZED MEAN EFFECTIVE THROUGHPUT FOR DIFFERENT SCENARIOS

Scenario	Normalized mean effective throughput R_E			
	Legacy	Blind	Proposed	Optimal
Scenario A	0.55	0.54	0.75	1.00
Scenario B	0.43	0.88	1.00	1.00
Scenario C	0.35	0.76	1.00	1.00

users with high throughput in the sub-6 GHz band do not benefit from the higher throughput in the mmWave bands following these policies. But this is not the case for the optimal algorithm, since the BS picks the band with the maximum throughput each frame without a threshold. However, as we will show later, our proposed algorithm can overcome this issue by increasing the band switch threshold, without losing its accuracy. Finally, to quantify the gains provided by the different band switch policies, we list the mean effective throughput in Table IV. Based on the values for Scenario A, the proposed algorithm provides a gain in the mean effective throughput of 39% and 37% over the blind and the legacy policies, respectively, and just 20% behind the optimal algorithm. Overall, the results for this scenario are promising and show the effectiveness of the proposed algorithm.

2) *Scenario B*: For the second scenario, the results are shown in Fig. 14. There are a few differences between this scenario and Scenario A. Firstly, all policies achieve the same performance in the high rate regime, which is due to the assumption that all users start in the

TABLE V
BAND SWITCH GRANTS AND REQUESTS BASED ON THE BAND SWITCH POLICY AND THE SCENARIO

Scenario	Policy	Band switches requested	Band switches granted
Scenario A	Legacy	22,458	2,751
	Blind	22,458	22,458
	Proposed	22,458	2,724
Scenario B	Legacy	41,609	32,558
	Blind	41,609	41,609
	Proposed	41,609	32,569
Scenario C	Legacy	43,033	21,514
	Blind	43,033	43,033
	Proposed	43,033	21,619

mmWave band. To be precise, the high throughput regime (above 4 Mbps) can only be achieved on mmWave bands as shown in Fig. 6. Hence, given that the users start in the mmWave band and this range is above the band switch threshold, these users remain in mmWave regardless of the policy, which justifies the identical performance of the different policies in the high throughput regime. Secondly, the blind policy achieves an identical performance to the optimal policy for the low throughput region. To justify this, we need the data in Table V, which show the number of band switch requests and the number of the granted ones for each policy and each scenario. From this table, we can see that up to around 70% of the band switch requests in this scenario are granted (assuming the optimal policy). Hence, the blind policy is identical to the optimal policy 70% of the time. Among this 70% of the users are the users who suffer from extremely low throughput at mmWave, mostly due to blockage. Hence, the blind algorithm results in the optimal decision for these users which justifies the identical performance for the low throughput regime. However, the blind policy also makes the wrong decision 30% of the time, which results in a gap between this policy and the optimal policy in the medium throughput regime. Finally, the legacy is always the worst in this scenario, which is due the measurement gap and the fact that the blind policy is accurate 70% of the time without having a measurement gap. The averages of the effective throughput are also shown in Table IV. The proposed policy achieves 130% gain in the throughput compared to the legacy policy and 13% compared to the blind policy. Also, the effective throughput for the proposed is almost identical to the optimal policy.




Fig. 10. The distribution of the effective throughput for Scenario C under different band switch policies.

3) *Scenario C*: The results for this scenario are presented in Fig. 10. As expected, the results lie between the previous two, and all the curves can be justified using the same arguments we have above. The reason we include this scenario is to have an idea on the gains we might observe in practice, since part of the users will be using mmWave and the others using sub-6 GHz. The mean gains are also presented in Table V.

Overall, the results for the different scenarios show the superiority of the proposed policy compared to the legacy and the blind policies; up to 130% improvement in the effective throughput depending on the considered scenario. It also justifies the use of a machine learning approach to solve this problem. Next, we provide more technical discussions on the accuracy of the proposed algorithm, more insights, and possible extensions to this work.

C. Discussion

We start with discussing the predictive accuracy of the proposed algorithm. Fig. 11 shows that our proposed algorithm usually made the right decisions; only a very few times did it deny the band switch when it was supposed to grant it (and vice versa).

To compare the performance of the ML classifiers², we show the performance of XGBoost alongside DNN. In Fig. 11 and Fig. 12, we show the confusion matrix for the three considered

²These comparisons may not be general for any ML algorithm, but are valid for these important ML classification algorithms.

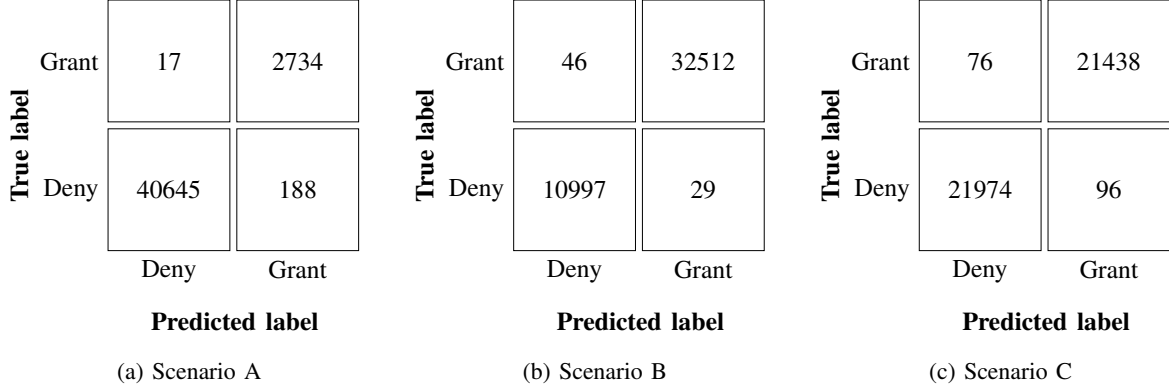


Fig. 11. The confusion matrix \mathbf{C} for the three scenarios using DNN.

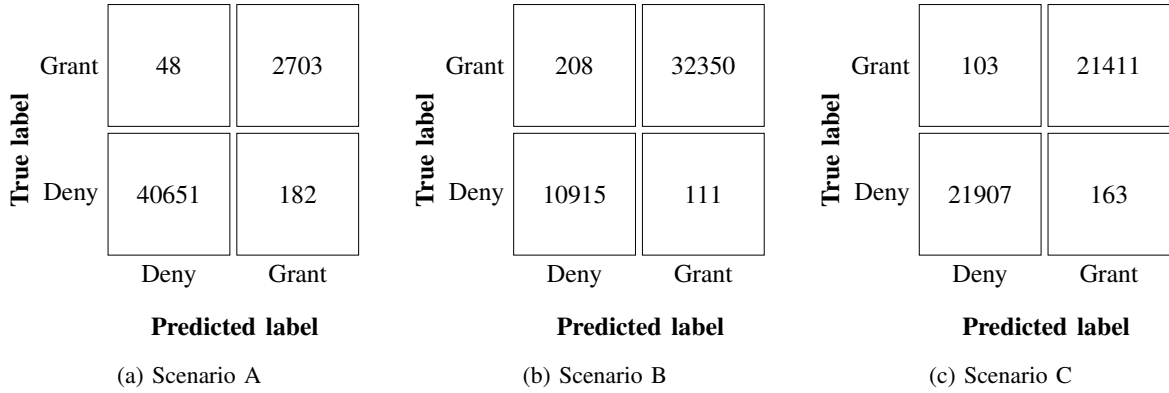


Fig. 12. The confusion matrix \mathbf{C} for the three scenarios using XGBoost.

scenarios using DNN and XGBoost respectively. Precisely, the misclassification error (μ) using DNN (XGBoost) is 0.47% (0.53%), 0.17% (0.73%), and 0.39% (0.61%) for Scenarios A, B, and C, respectively. The run-time complexity of XGBoost using the hyperparameters in Table III has an upper bound in $\mathcal{O}(n + n \log n) = \mathcal{O}(n \log n)$ [22], where $n := N_{\text{learning}}$. However, for DNN this complexity is super-linear since training DNNs requires matrix multiplications [23]. Matrix multiplications have a run-time complexity between $\mathcal{O}(n^2)$ and $\mathcal{O}(n^3)$. Hence, the classifier choice between DNN and XGBoost is a trade-off between decision speed and accuracy: if accuracy is desired, then choose DNN, but if less run-time complexity is desired for decision speed, XGBoost is a more attractive choice.

The second point we highlight here is the amount of training data that we require to have an accurate prediction. In Fig. 13, we show the ROC area (ξ) and the misclassification error (μ) for different training data sizes. In particular, the figure shows that training using only 1,362

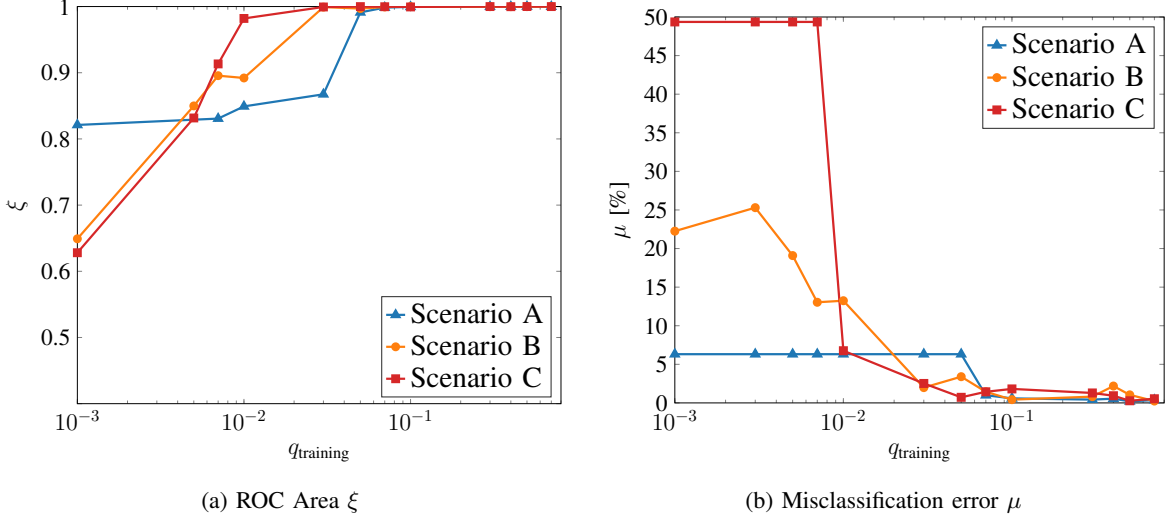


Fig. 13. The classification performance of the proposed algorithm for different training data sizes and different scenarios.

measurements (i.e., $1/40$ of the data for a grid of an area of 19,250 square meters) is enough to have an excellent performance—less than 2% misclassification error. In other words, having knowledge about the previous band switch decisions of 7 random samples per 100 square meters is enough to predict the band switch decisions for the rest of the locations. This absolute number depends on the spatial correlation between the channels on different locations, as well as the hyperparameters and the choice of the machine learning algorithm. Further, this insight should be understood alongside the other considerations, such as the user grid size, the collection period, and the blockage probabilities, as discussed in this section.

Note that our presented results so far are for a single band switch threshold value. However, we claim that the performance gap between the optimal algorithm and the proposed one can be reduced by increasing the threshold. To verify this claim, we show the mean effective throughput for different band switch thresholds in Table VI. In Scenario A, we observe that as we increase the band switch threshold $r_{\text{threshold}}$, the performance gap between the mean effective throughputs of the proposed and the optimal rates shrink considerably. While both the legacy and blind rates also get better, their performance is not close to the optimal: the legacy because of the measurement gap and the blind because of the undesired band switch. However, in Scenario B, both the blind and legacy rates deteriorate as we increase the band switch thresholds. In the legacy policy, it is also due to the measurement gaps, and for the blind policy, it is because users who were getting up to 10 Mbps on mmWave are now getting 3 Mbps at best as shown




Fig. 14. The distribution of the effective throughput for Scenario B under different band switch policies.

TABLE VI
IMPACT OF DIFFERENT BAND SWITCH THRESHOLDS FOR SCENARIOS A AND B

	$r_{\text{threshold}}$	Normalized mean effective throughput R_E			
		Legacy	Blind	Proposed	Optimal
Scenario A	1.72	0.55	0.54	0.75	1.00
	2.00	0.45	0.46	0.77	1.00
	2.60	0.34	0.60	1.00	1.00
Scenario B	2.00	0.43	0.88	1.00	1.00
	9.00	0.39	0.84	1.00	1.00
	12.50	0.33	0.76	1.00	1.00

in Fig. 6. As expected, we do not see much of a change in the proposed rate as we increase the band switch threshold, aligned with the CDFs in Fig. 7 and Fig. 14.

Further, to show the behavior of the proposed policy against uncertainty due to the change of the blockage probability as in (20), we simulate Scenario C using $p \in \{0.2, 0.4, 0.6, 0.8\}$ in the exploitation data, while the learning data is fixed at $p = 0.4$. Then, we compute the mean effective throughput as shown in Fig. 15. The classifier remains resilient against uncertainty of band-selective blockage contrary to the other policies (with a significance up to the third decimal). This is due to the ability of the classifier to learn from the spatial relationship of the channels even with blockage as a result of: 1) coordinates being part of the learning features and 2) the relaxation of $r_{\text{threshold}}$ value in the learning phase.




Fig. 15. Normalized mean effective throughput for different blockage probabilities p in the *exploitation* phase

VIII. CONCLUSIONS

In this paper, we used both deep neural networks and XGBoost classifiers to rank the downlink channel quality of the frequency bands prior to the band switch, which is a mathematically intractable problem. The use of classifiers in an online learning setting eliminates the dependence on measurement gaps during a band switch in a dual-band base station. We exploited the spatial and spectral relationships in both the sub-6 GHz and mmWave bands through the use of a ray-tracing dataset. This brought forward two benefits: 1) reduces link latency by removing the need for a measurement gap and 2) reduces complexity in the UE and BS because channel estimation in the other frequency band is not required. We revealed insights as to why the deep learning classification method was needed and why it worked. We simulated one dual-band base station with many UEs in its association area and varied the blockage probability. In this simulation, our method improved downlink throughput by up to 1.3x compared with the legacy policy over different scenarios with a misclassification error less than 0.3%. The observed improvement is due to the classifier ability to exploit the spatial correlation of channels across the different frequency bands and thus accurately predict the effective achievable rate on the target frequency without the dependency on a measurement gap. This band selection method is better suited for 5G and beyond where maintaining high data rates is desired without interrupting the data flow. We focused on the case where the BS has only two bands: one centered at 3.5 GHz and the other at 28 GHz, since the dataset we use supports these two bands. An interesting extension is

for multiple bands, or when a handoff between multiple BSs is required due to mobility, when subsequent band switches are required (also known as the “ping-pong” effect), or when the BS uses different radio units per frequency band as in [33].

REFERENCES

- [1] F. B. Mismar and B. L. Evans, “Partially Blind Handovers for mmWave New Radio Aided by Sub-6 GHz LTE Signaling,” in *Proc. IEEE International Conference on Communications Workshops*, May 2018, pp. 1–5.
- [2] 3GPP TS 36.331 V15.5.1, “Evolved universal terrestrial radio access (E-UTRA); radio resource control (RRC); protocol specification (Release 15),” Apr. 2019.
- [3] 3GPP TS 36.331 V1.0.0, “Evolved universal terrestrial radio access (E-UTRA); radio resource control (RRC); protocol specification (Release 8),” Nov. 2007.
- [4] 3GPP TR 32.860 V14.0.0, “Technical Specification Group Services and System Aspects; Telecommunication management; Study on enhancements of OAM aspects of distributed Self-Organizing Networks (SON) functions (Release 14),” Jun. 2016.
- [5] N. Gonzalez-Prelcic, A. Ali, V. Va, and R. W. Heath Jr., “Millimeter-wave communication with out-of-band information,” *IEEE Communications Magazine*, vol. 55, no. 12, pp. 140–146, Dec. 2017.
- [6] K. Hugl, K. Kalliola, and J. Laurila, “Spatial reciprocity of uplink and downlink radio channels in FDD systems,” in *Proc. European Cooperation in the Field of Scientific and Technical Research*, vol. 273, no. 2. Citeseer, May 2002, p. 066.
- [7] M. Alrabeiah and A. Alkhateeb, “Deep learning for TDD and FDD massive MIMO: Mapping channels in space and frequency,” *arXiv preprint arXiv:1905.03761*, 2019.
- [8] T. Nitsche, A. B. Flores, E. W. Knightly, and J. Widmer, “Steering with eyes closed: Mm-Wave beam steering without in-band measurement,” in *Proc. IEEE INFOCOM*, Apr. 2015, pp. 2416–2424.
- [9] A. Ali, N. Gonzalez-Prelcic, and R. W. Heath, “Millimeter wave beam-selection using out-of-band spatial information,” *IEEE Trans. on Wireless Communications*, vol. 17, no. 2, pp. 1038–1052, Feb. 2018.
- [10] Y. Xiu, W. Wang, and Z. Zhang, “A message passing approach to acquire mm-Wave channel state information based on out-of-band data,” *IEEE Access*, vol. 6, pp. 45 665–45 680, Jul. 2018.
- [11] Y. Xiu, J. Wu, C. Xiu, and Z. Zhang, “Millimeter wave cell discovery based on out-of-band information and design of beamforming,” *IEEE Access*, vol. 7, pp. 23 076–23 088, Feb. 2019.
- [12] A. Ali, N. González-Prelcic, and R. W. Heath Jr, “Spatial covariance estimation for millimeter wave hybrid systems using out-of-band information,” *IEEE Trans. on Wireless Communications*, 2019, early access.
- [13] A. Alkhateeb, “DeepMIMO: A generic deep learning dataset for millimeter wave and massive MIMO applications,” in *Proc. of Information Theory and Applications Workshop*, Feb. 2019.
- [14] A. Alkhateeb. DeepMIMO. [Online]. Available: <https://github.com/DeepMIMO/DeepMIMO-codes>
- [15] O. Semiari, W. Saad, and M. Bennis, “Joint Millimeter Wave and Microwave Resources Allocation in Cellular Networks With Dual-Mode Base Stations,” *IEEE Trans. on Wireless Communications*, vol. 16, no. 7, pp. 4802–4816, Jul. 2017.
- [16] O. Semiari, W. Saad, M. Bennis, and B. Maham, “Caching meets millimeter wave communications for enhanced mobility management in 5G networks,” *IEEE Trans. on Wireless Communications*, vol. 17, no. 2, pp. 779–793, Feb. 2018.
- [17] M. Polese, M. Giordani, M. Mezzavilla, S. Rangan, and M. Zorzi, “Improved Handover Through Dual Connectivity in 5G mmWave Mobile Networks,” *IEEE Journal on Sel. Areas in Communications*, vol. 35, no. 9, pp. 2069–2084, Sep. 2017.
- [18] J. G. Andrews, S. Buzzi, W. Choi, S. V. Hanly, A. Lozano, A. C. K. Soong, and J. C. Zhang, “What will 5G be?” *IEEE Journal on Sel. Areas in Communications*, vol. 32, no. 6, pp. 1065–1082, Jun. 2014.

- [19] J. Mo, B. L. Ng, P. Huang, M. Kulkarni, A. AlAmmouri, J. C. Zhang, S. Chang, J. Lee, and W.-J. Choi, "Beam codebook design for 5G mmWave terminals," *IEEE Access*, vol. 7, pp. 98 387–98 404, Jul. 2019.
- [20] T. S. Rappaport, S. Sun, R. Mayzus, H. Zhao, Y. Azar, K. Wang, G. N. Wong, J. K. Schulz, M. Samimi, and F. Gutierrez Jr, "Millimeter wave mobile communications for 5G cellular: It will work!" *IEEE Access*, vol. 1, no. 1, pp. 335–349, May 2013.
- [21] A. AlAmmouri, J. Mo, B. L. Ng, J. C. Zhang, and J. G. Andrews, "Hand grip impact on 5G mmWave mobile devices," *IEEE Access*, vol. 7, pp. 60 532–60 544, May 2019.
- [22] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in *Proc. of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, Aug. 2016, pp. 785–794.
- [23] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*, 1st ed. Cambridge, MA, USA: The MIT Press, 2016.
- [24] F. B. Mismar and B. L. Evans, "Deep Learning in Downlink Coordinated Multipoint in New Radio Heterogeneous Networks," *IEEE Wireless Communications Letters*, vol. 8, no. 4, pp. 1040–1043, Mar. 2019.
- [25] T. Fawcett, "An introduction to ROC analysis," *Pattern Recognition Letters*, vol. 27, no. 8, pp. 861–874, Jun. 2006.
- [26] A. Ghosh, J. Zhang, J. G. Andrews, and R. Muhamed, *Fundamentals of LTE*. Prentice-Hall, 2010.
- [27] V. Va, J. Choi, and R. W. Heath Jr., "The Impact of Beamwidth on Temporal Channel Variation in Vehicular Channels and Its Implications," *IEEE Trans. on Veh. Technology*, vol. 66, no. 6, pp. 5014–5029, Jun. 2017.
- [28] D. J. Love and R. W. Heath Jr., "Equal gain transmission in multiple-input multiple-output wireless systems," *IEEE Transactions on Communications*, vol. 51, no. 7, pp. 1102–1110, Jul. 2003.
- [29] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. International Conference on Learning Representations ICLR*, May 2015.
- [30] D. Chizhik, "Slowing the time-fluctuating MIMO channel by beam forming," *IEEE Transactions on Wireless Communications*, vol. 3, no. 5, pp. 1554–1565, Sep. 2004.
- [31] 3GPP TS 36.133 V13.3.0, "Evolved universal terrestrial radio access (E-UTRA); requirements for support of radio resource management; protocol specification (Release 13)," May 2016.
- [32] F. B. Mismar. Source code. [Online]. Available: <https://github.com/farismismar/Bandswitch-DeepMIMO>
- [33] O. Semiari, W. Saad, M. Bennis, and M. Debbah, "Integrated Millimeter Wave and Sub-6 GHz Wireless Networks: A Roadmap for Joint Mobile Broadband and Ultra-Reliable Low-Latency Communications," *IEEE Wireless Communications*, vol. 26, no. 2, pp. 109–115, Apr. 2019.