

Efficient Beam Alignment in Millimeter Wave Systems Using Contextual Bandits

Morteza Hashemi, Ashu Sabharwal, C. Emre Koksal, and Ness B. Shroff

Abstract—In this paper, we investigate the problem of beam alignment in millimeter wave (mmWave) systems, and design an optimal algorithm to reduce the overhead. Specifically, due to directional communications, the transmitter and receiver beams need to be aligned, which incurs high delay overhead since without a priori knowledge of the transmitter/receiver location, the search space spans the entire angular domain. This is further exacerbated under dynamic conditions (e.g., moving vehicles) where the access to the base station (access point) is highly dynamic with intermittent on-off periods, requiring more frequent beam alignment and signal training. To mitigate this issue, we consider an online stochastic optimization formulation where the goal is to maximize the directivity gain (i.e., received energy) of the beam alignment policy within a time period. We exploit the inherent correlation and unimodality properties of the model, and demonstrate that contextual information improves the performance. To this end, we propose an equivalent structured Multi-Armed Bandit model to optimally exploit the exploration-exploitation tradeoff. In contrast to the classical MAB models, the contextual information makes the lower bound on *regret* (i.e., performance loss compared with an oracle policy) independent of the number of beams. This is a crucial property since the number of all combinations of beam patterns can be large in transceiver antenna arrays, especially in massive MIMO systems. We further provide an asymptotically optimal beam alignment algorithm, and investigate its performance via simulations.

I. INTRODUCTION

To meet the exponentially growing demand in mobile data, the trend in wireless networks is migrating to higher frequencies combined with increasing number of antennas per device and per base station. For instance, it is envisioned that in 5G cellular systems certain portions of the millimeter wave (mmWave) band will be used, spanning the spectrum between 30GHz to 300GHz. However, propagation loss at mmWave frequencies is much higher due to a variety of factors including atmospheric absorption, basic Friis transmission-effect, and low penetration. When the users and/or surrounding objects are mobile, this effect is more pronounced such that different propagation paths become highly variable with intermittent on-off periods. Thus, unlike existing communication schemes, mmWave systems require highly directional communications to compensate for large channel losses. Thanks to recent advances in antenna technologies, large directional antenna arrays with much smaller form factors can be deployed in

relatively small chip areas. Such arrays have the potential to focus the signal energy toward a specific direction, “making up” for the channel losses.

In order to fully utilize directional communications, the transmitter and receiver beams need to be aligned. The experimental results in [1] demonstrate that in a system with a 7 degree beam width, a misalignment of 18 degree reduces the link budget by around 17 dB, which can reduce the maximum throughput by up to 6 Gbps or break the link entirely [2]. On the other hand, as a result of such “pencil-beams” at the transmitter and receiver, beam alignment incurs a large overhead that scales with device mobility and the product of the transmitter-receiver beam resolution. In exhaustive search methods, both users and base stations have a predefined codebook of beam directions that cover the entire angular space and are used sequentially to transmit and receive. Thus, the complexity of this exhaustive search is $O(N^2)$, where N is the number of possible beam directions. To improve the search efficiency, the transmitter and receiver steering is decoupled in the 802.11ad standard such that the transmitter starts with a quasi-omnidirectional beam, while the receiver scans the space for the best beam direction. The process is then reversed [3]. This approach reduces the search complexity to $O(N)$. Still, for a beam of a few degrees, the delay can be hundreds of milliseconds to seconds [4], which would easily stall real-time applications.

Dynamic conditions make the beam alignment more challenging since there is the need for frequent beam alignment. Under such scenarios, we pose the following question that *given the outcome of the past beam alignments, is it possible to extract some information and reduce the search space for the subsequent beam alignment procedures?* In particular, our work is based on the fact that successive beam alignments are stochastically correlated, and thus, outcome of the previous “beam matching” provides *contextual information* for the subsequent matchings, thus eliminating the need to search the entire angular domain. We exploit *correlation* and *unimodality* properties across various beam matching. Specifically, for a given beam matching, we call the difference between the transmitter and receiver direction as *misalignment*. Because of correlation if matching at a larger misalignment is successful (i.e., received energy is above a threshold τ), with a high probability a matching will be successful at a smaller misalignment as well. Furthermore, the directivity gain (or received energy) can be approximated as a unimodal function of the misalignment value. We exploit this contextual information in order to obtain a beam search scheme that quickly identifies the best beam direction and maximizes the directivity gain. We formulate the problem of finding the best beam pair as

M. Hashemi and C. E. Koksal are with the Department of Electrical and Computer Engineering, Ohio State University, Columbus, OH, 43210 USA, e-mail: hashemi.20@osu.edu, koksal.2@osu.edu

A. Sabharwal is with the Department of Electrical and Computer Engineering, Rice University, Houston, TX 77005, USA, e-mail: ashu@rice.edu.

N. B. Shroff is with the Department of Electrical and Computer Engineering and Department of Computer Science and Engineering, Ohio State University, Columbus, OH, 43210 USA, e-mail: shroff.11@osu.edu.

an online stochastic optimization where the objective is to maximize the expected amount of received energy within a given time period.

To find the optimal solution, we show that this problem can be considered as an instance of the Multi Armed Bandit (MAB) model in which *each transmit and receive beam pair is considered as a single arm*. Thus, the objective is to design a sequential arm selection (or, equivalently, beam alignment) strategy that maximizes the *expected reward* (received energy) over a given time horizon. Performance of MAB models is usually expressed in terms of *regret* that is defined as the total expected reward loss compared with an oracle policy. Regret of the best algorithm is in the form of $O(K \log(T))$ in which K is the number of arms and T is the time horizon [5]. In distinction with the classical MAB models, we exploit the contextual information of beam alignment that leads to a structured MAB model, and prove that *the regret does not scale with the number of arms that is equal to the number of beam matchings*. This is a crucial property in (massive) MIMO systems, and provides a fundamental performance limit satisfied by any exhaustive beam selection algorithm. This limit quantifies the inevitable performance loss due to the need to explore sub-optimal beam pairs. It also characterizes the performance gains that can be achieved by devising beam pair selection schemes that optimally exploit the correlations and the structural properties of the MAB problem. *Therefore, in contrast to the previous works that explore the sub-optimal beam pairs by heuristics, our method optimally explores the sub-optimal beam pairs*. The following example illustrates how we achieve this goal.

Example: *Let us consider a scenario where the transmitter beam direction is fixed at 75° angle with respect to the receiver. For the sake of exposition, we assume a 2D setting. Assume that there are 16 possible directions at the receiver, as shown in Fig. 1. Using the exhaustive beam selection scheme, each of 16 directions will be examined one at a time, and the direction with the largest received energy (from beacon messages) is picked. However, under dynamic conditions (e.g., with mobility), the optimal beam direction can potentially change within a short period of time. In this case, we consider maximizing the received energy within a given period of time. Using our proposed scheme, the receiver assigns an index to each beam direction, and the beam with the highest index will be selected. The important point is that due to the correlation and unimodality properties, the search space will be limited to the neighborhood of the beam with maximum index. As a result, it prevents the need of a uniform exploration over the entire angular domain, thus mitigating the overhead of beam alignment when the number of beam directions becomes large.*

In summary, our contributions are as follows:

- We consider the beam alignment problem, and investigate the fundamental performance limits of the search-based beam alignment between the transmitter and receiver antennas. We model the problem of finding the best beam alignment as an online stochastic optimization problem.
- We exploit contextual information of the problem and formulate an equivalent structured Multi-Armed Bandit

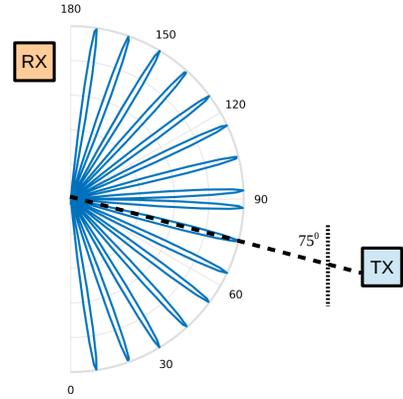


Fig. 1. Beam alignment of the transmitter and receiver with 16 beams.

model. We experimentally demonstrate that the received power (approximately) follows a unimodal pattern.

- We derive a lower bound on the regret of any search-based algorithm, and demonstrate that the regret does not scale with the number of transmission and receive beams, thanks to the underlying structure of the problem. Finally, we propose Unimodal Beam Alignment (UBA) algorithm that is shown to be asymptotically optimal.

II. RELATED WORK

A. Beam Alignment

The authors in [6] perform initial access for clustered mmWave small cells using the power delay profile. In this case, base stations are coordinated in clusters, and communicate through a backhaul network. Base stations share their measurement reports obtained from the mobile devices, and location of the mobile is estimated based on the shared measurements. This will enable the base stations to point at the estimated mobile location. Although this method is limited to line-of-sight scenarios, the probability of having at least three line-of-sight links (needed for mobile localization) increases by assuming larger cluster sizes. In another line of research, the authors in [7] proposed a fast-discovery hierarchical search method, while [8] exploits the sparse multipath structure of the mmWave channel to optimize the choice of beamforming directions. A cell discovery method is proposed in [9] in which the base station periodically transmits synchronization signals to scan the entire angular space in time-varying random directions. In [10], a beam alignment technique is designed based on adaptive subspace sampling and hierarchical beam codebooks. The authors in [11] use spatial information extracted at sub-6 GHz to help estimate the best beam pairs at the transmitter and receiver at mmWave frequencies. In [12], a beam alignment scheme based on scanning several directions by one-shot is proposed. In contrast to the previous works, we focus on exploiting contextual information in standalone mmWave systems in order to reduce the search space, and thus the overhead of beam alignment operation. Note that there are other related work to reduce the overhead of beam search in integrated sub-6 GHz-mmWave systems [13, 14].

B. Multi-Armed Bandit

Multi-Armed Bandit (MAB) framework formulates sequential decision problems where an agent (i.e., decision maker) has to strike an optimal trade-off between exploitation and exploration by sequentially selecting an action (or an arm), and observing the corresponding reward. Rewards of a given arm are random variables with unknown distributions. The objective is to maximize the expected reward over a given time horizon by selecting the optimal arm at each time slot. Most of the existing works focus on *unstructured MAB* problems in which the reward associated with different arms are not related [5, 15]. In contrast to the unstructured MAB models, when the average rewards are structured, deriving the optimal regret bound and designing of optimal decision algorithms is more challenging [16]. Unimodal bandits are specific instances of bandit models in which the average reward of arms are correlated. In [17], unimodal bandits with a continuous set of arms are studied, and the authors show that the regret of the order of $O(\sqrt{T})$ is achievable under some strong regularity assumptions on the reward functions. For the same problem, the authors in [18] provide an algorithm that achieves $O(\sqrt{T} \log(T))$ regret under relaxed regularity assumptions. In this paper, we cast the problem of mmWave beam alignment as a unimodal bandit model, and derive the regret bound.

III. MODEL AND OBJECTIVE

A. System Model

In mmWave systems, once the connection is lost, there are two options for connection establishment and subsequent beamforming: digital or analog. Digital beamforming is highly efficient in delay such that with the observations from all receive antennas, beamforming can be done by one-shot processing of the observed beacons. However, to achieve digital beamforming, there is the need for a separate analog-to-digital converter (ADC) for each antenna, which may not even be feasible for even a small to mid-sized antenna array due to high energy consumption. On the other hand, while analog beamforming requires only one ADC, it can focus on one direction at a time, making the search process costly in delay. Given the fragility of the mmWave channel, the need to scan the entire space leads to the loss of opportunities to utilize the mmWave channel upon its availability. In order to avoid high energy consumption by mmWave components, we focus on analog beamforming in which a single RF chain is deployed at the transmitter and receiver. Other implementations (e.g., hybrid architectures) may look at combinations of directions, which is out of scope of this work. Figure 2 depicts a schematic of the analog beamforming, and directional beams at the transmitter and receiver. We assume that the transmitter and receiver are equipped with phased array antennas with M_t and M_r identical antennas respectively, equally spaced by a distance d along an axis. For the sake of exposition, we only consider the receiver side, while a similar argument is held for the transmitter. Due to the use of analog architecture, the signal at the input of the decoder is a scalar, identical to a

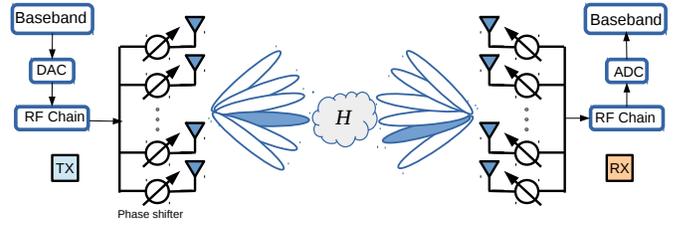


Fig. 2. Beam alignment using analog beamforming

weighted combination of signal x_m across all antennas. Thus, the received signal at the mmWave receiver can be written as:

$$y_m = \mathbf{w}_r^\dagger \mathbf{H} \mathbf{w}_t \cdot x_m + n_m, \quad (1)$$

where \mathbf{w}_r and \mathbf{w}_t are the beamforming vectors. The white Gaussian noise n_m is normalized to have unit variance. If the transmitter uses N_t training precoding vectors \mathbf{w}_t , and the receiver uses N_r training combining vectors \mathbf{w}_r , then the collected signals (divided through by the training signal) is given by:

$$\mathbf{Y} = \mathbf{W}_r^\dagger \mathbf{H} \mathbf{W}_t + \mathbf{N},$$

where $\mathbf{W}_r = [\mathbf{w}_{r_1}, \dots, \mathbf{w}_{r_{N_r}}]$ is the $M_r \times N_r$ combining matrix, and $\mathbf{W}_t = [\mathbf{w}_{t_1}, \dots, \mathbf{w}_{t_{N_t}}]$ is the $M_t \times N_t$ precoding matrix. Furthermore, \mathbf{N} is the $N_r \times N_t$ post-processing noise matrix. Hence, at each time slot, the problem of finding the best beam pair boils down to finding the largest value of matrix \mathbf{Y} . In the exhaustive search, one should examine all $N_r \times N_t$ elements of \mathbf{Y} to find the largest index, which determines the optimal beam index at the transmitter and receiver. The authors in [11] use spatial information extracted at sub-6 GHz to help estimate the largest index of \mathbf{Y} . Applying the same framework, our beam alignment method exploits correlation across the elements of \mathbf{Y} in order to reduce the search space to sub-matrices of \mathbf{Y} .

B. Problem Statement

In order to establish a mmWave link, the transmitter selects a beam direction that determines the phase shifter weights to steer the beam in a certain direction. Similarly, the receiver selects a receive beam index to receive the signals in a certain direction. To obtain a high beamforming gain, the transmitter and receiver beams should be well aligned with each other. We let $D_{i,j} = (i, j)$ to denote a pair of beam direction in which i is the beam index at the transmitter, and j denotes the receiver beam index. There are N_t and N_r beams at the transmitter and receiver, respectively. Further, we define $\mathcal{D} = \{(i, j) : 1 \leq i \leq N_t, 1 \leq j \leq N_r\}$ as the set of all possible beam pairs such that there exists $K = N_t \times N_r$ distinct *matching* between the transmit and receive beams. For each pair of transmit and receive beams, *misalignment* is defined as follows.

Definition. (Misalignment) Given the pair $(i, j) \in \mathcal{D}$ of transmitter and receiver beams, the misalignment $\delta_{i,j}$ captures the angular mismatch between the i -th transmitter beam and j -th receiver beam.

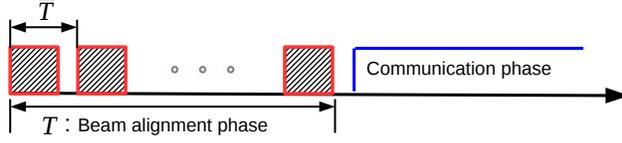


Fig. 3. Sequence of beam alignment operations followed by a data communication phase.

Set $\mathcal{A} = \{\delta_{i,j} : 1 \leq i \leq N_t, 1 \leq j \leq N_r\}$ contains all possible values of the misalignment values such that \mathcal{A} is partially ordered. With an abuse of notation, we let δ_k to denote the k^{th} ($k \leq K$) misalignment value.

In order to detect the transmitted beacon signals, the received energy level should lie above a certain threshold τ , which is determined based on the quality of service (QoS) needed. Once the received energy is larger than τ , we call it a *successful matching*. For a fixed transmit power, the received signal energy can be expressed as a function of the misalignment value, i.e., $p(\delta_k)$ such that $p(\cdot)$ is non-increasing. For the sake of notations, we use $p_k := p(\delta_k)$. The probability of success for the misalignment δ_k is given by $\theta_k = \mathbb{P}(p_k \geq \tau)$. We assume a fixed situation (e.g., LOS) such that from one run of the alignment algorithm to another, the situation remains fixed and only the orientations and distances change, i.e., the success probabilities θ_k 's are time-invariant. Furthermore, we let a binary random variable X_k represent the success ($X_k = 1$) or failure ($X_k = 0$) of the matching. The optimal matching is unique, i.e., there exists k^* such that $p_{k^*}\theta_{k^*} \geq p_k\theta_k$, for all $k \neq k^*$.

Problem formulation: Time is slotted, and we let T to denote the length of beam alignment phase followed by the data communication phase. Further, \bar{T} denotes the length of pilot signal by which each beam is measured. This value captures the amount of time that it takes to examine a single pair of beams, as shown in Fig. 3. We formulate the problem of finding the best beam pair as an online stochastic optimization problem such that an optimal beam selection policy π maximizes the expected amount of energy received from beacon messages up to a certain finite time T . In this case, we let $s_k^\pi(T)$ denote the number of times that the misalignment k is selected under policy π and within the time period T . Therefore, the optimal beam selection policy solves the following optimization problem:

$$\max_{\pi} \sum_k \mathbb{E}[s_k^\pi(T)] p_k \theta_k \quad (2a)$$

$$\text{s.t. } \bar{T} \sum_k s_k^\pi(T) \leq T \text{ and } s_k^\pi(T) \in \mathbb{N}. \quad (2b)$$

In this formulation, small misalignment and large probability of success is desirable. In addition, given that examining each beam matching takes \bar{T} on average, the total number of beam examinations is upper-bounded by $\frac{T}{\bar{T}}$, which is reflected in the first constraint. The second constraint implies that the number of each matching examination should be an integer.

IV. EQUIVALENT MULTI-ARMED BANDIT MODEL

The beam alignment formulation implies an MAB model such that each combination of the transmitter and receiver beam is considered as an arm, which leads to $K = N_t \times N_r$ total arms. In this work, we use the terms ‘‘arm’’ and ‘‘beam direction’’ interchangeably. In this case, $s_k^\pi(T)$ denotes the number of times that arm k has been selected under policy π . Moreover, the reward of arm k has Bernoulli distribution with parameter θ_k such that it is p_k with probability θ_k and 0 with probability $1 - \theta_k$. The average reward of arm k is denoted by $\mu_k := p_k \theta_k$.

A. Contextual Information

We explore a new type of contextual information that correlates the misalignment and the received energy. In particular, due to physics of signal propagation, if matching at a larger misalignment is successful, a matching at a smaller misalignment will be successful with a high probability. On the other hand, if matching at a smaller misalignment fails, then a matching at a larger misalignment will fail with a high probability as well. Hence, we have:

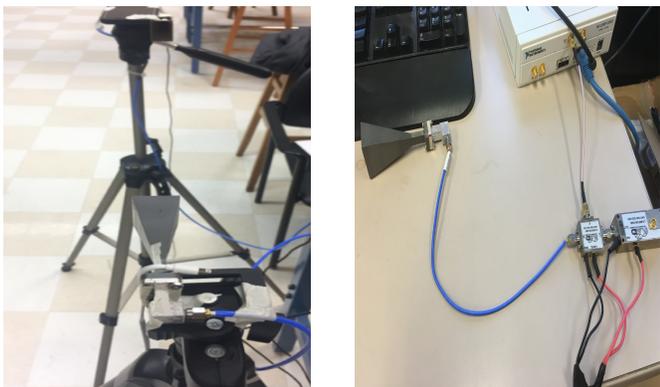
$$(\delta_m > \delta_n \text{ and } X_n = 0) \Rightarrow X_m = 0, \quad (3)$$

$$(\delta_m < \delta_n \text{ and } X_n = 1) \Rightarrow X_m = 1. \quad (4)$$

Equivalently, we can define the vector of success probabilities to satisfy the following condition: $\boldsymbol{\theta} = (\theta_1, \theta_2, \dots, \theta_K) \in \mathcal{T}$, where $\mathcal{T} = \{\boldsymbol{\theta} \in [0, 1]^K : \theta_1 \geq \theta_2 \geq \dots \geq \theta_K\}$. In addition to the correlation property, we note that amount of energy received can be approximated as a unimodal function of misalignment. In this case, $\boldsymbol{\theta} \in \mathcal{U}$ such that $\mathcal{U} = \{\boldsymbol{\theta} \in [0, 1]^K : \exists k^*, p_1\theta_1 < \dots < p_{k^*}\theta_{k^*}, p_{k^*}\theta_{k^*} > p_{k^*+1}\theta_{k^*+1} > \dots > p_K\theta_K\}$.

Graph representation: In order to demonstrate the implications of contextual information, we can utilize a graph representation to capture the order of arms (beam pairs) with respect to each other. In this model, each arm corresponds to a node of a graph and each edge is associated with a relationship specifying which node of the edge gives the largest expected reward, thus providing a partial ordering over the arm space. Furthermore, from any node there is a path leading to the unique node with the maximum expected reward along which the expected reward is monotonically increasing. Under the assumption of unimodal expected reward, we can move from low expected rewards to high ones just by climbing them in the graph, preventing the need of a uniform exploration over all the graph nodes. This assumption reduces the complexity in the search for the optimal arm, since the optimal policy can avoid pulling the arms corresponding to some subset of non-optimal nodes.

Experimental observations: For the purpose of illustration, we provide experimental results to observe the typical propagation pattern as a function of misalignment. In particular, we consider the case of clear line of sight, and run a set of experiments in which two horn antennas placed on tripods and facing each other symmetrically. The transmitter antenna is set to be the stationary antenna, while the receiver antenna is



(a) Two horn antennas at 40 GHz placed on tripods (b) Connection of antenna with the oscillator, mixer and USRP

Fig. 4. Experimental setup connection and equipment

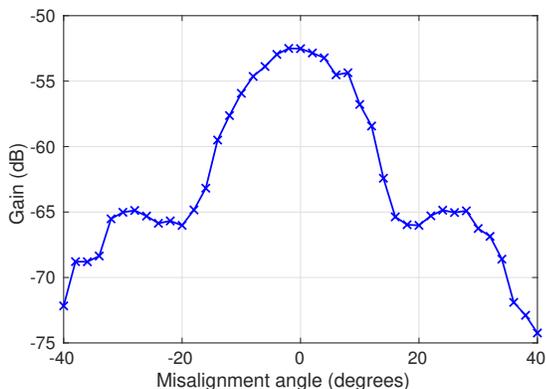


Fig. 5. Received power at distance of 1 meter vs. misalignment.

rotated throughout the experiment. Two software defined radio (NI USRP-2901) are set up as the transmitter and receiver at 4 GHz. The transmit antenna is connected to an up-converter with the output at 40 GHz. The receiver antenna takes the 40 GHz carrier and sends it to the down-converter with output of 4 GHz. A power spectrum of gain (dB) vs. frequency is displayed in real time for data collection, and then we record the average peak value of gain. The receive horn antenna sweeps 2 degrees incrementally on both clockwise and counter-clockwise directions until the gain is indistinguishable from the thermal noise floor (about 85 dB). Our experimental setup is shown in Fig. 4, and Fig. 5 demonstrates the received power as a function of misalignment angle between the transmitter and receiver antennas. We observe that received power approximately (due to the sidelobes) follows the unimodal pattern. In this work, we assume that the transmitter and receiver deploy highly directional antenna arrays in which the effect of sidelobes is negligible and thus the received power can be approximated as a unimodal function. Moreover, although we would expect the misalignment of 0 degree provides the highest signal energy, under blockage and reflection scenarios a misalignment of $\delta \neq 0$ may provide a larger gain. Our model on contextual information is general and captures these conditions as well.

V. PERFORMANCE ANALYSIS AND OPTIMAL ALGORITHM

A. Regret Analysis

In order to assess the performance of policy π for beam alignment, or equivalently arm selection, we consider *regret* as the performance metric, defined as follows:

$$\mathcal{R}^\pi(T) = p_{k^*} \theta_{k^*} T - \sum_k \mathbb{E}[s_k^\pi(T)] p_k \theta_k. \quad (5)$$

From this definition, *regret measures the expected reward loss (over a time period of T) compared with an oracle policy that would know everything*. In the case that the expected reward of the various arms are not correlated, regret of the best algorithm is in the form of $O(K \log(T))$ in which K is the number of arms [5]. As a result, regret scales linearly with the number of arms. In beam alignment, for a beam of a few degrees the total number of arms (i.e., all combinations of transmitter and receiver beam pairs) becomes very large. Therefore, in order to avoid the scaling factor, we exploit the structural properties of the arms and their reward functions, and show that due to contextual information, the scaling factor is constant and does not scale with the number of beam matchings (i.e., size of the decision space). To this end, for any arm k , we denote the set of its neighbors by:

$$N(k) = \{j \in \{k-1, k+1\} : p_k \theta_k \leq p_j\}.$$

Furthermore, given that the arms reward follow Bernoulli distribution, the Kullback-Leibler (KL) divergence of two Bernoulli distributions with respective parameters θ and θ^* is defined as: $I(\theta, \theta^*) = \theta \log \frac{\theta}{\theta^*} + (1-\theta) \log \frac{1-\theta}{1-\theta^*}$. It has been shown in [19] that the problem of learning in a unimodal bandit setting presents a lower bound over the regret of the following form:

Theorem 1. *For any beam alignment algorithm π , the lower bound on the regret is given by:*

$$\limsup_{T \rightarrow \infty} \frac{\mathcal{R}^\pi(T)}{\log(T)} \geq c(\theta),$$

in which $c(\theta)$ is a function of arms reward, and is given by:

$$c(\theta) = \sum_{k \in N(k^*)} \frac{p_{k^*} \theta_{k^*} - p_k \theta_k}{I(\theta_k, \theta_{k^*})}. \quad (6)$$

From (6), we observe that $c(\theta)$ is equal to a summation over a constant number of terms (i.e., independent of K). On the other hand, in the case that the structural properties is not exploited, the regret is lower bounded by:

$$c'(\theta) = \sum_{k \neq k^*} \frac{p_{k^*} \theta_{k^*} - p_k \theta_k}{I(\theta_k, \theta_{k^*})},$$

where $c'(\theta)$ linearly increases with the number of possible arms, or equivalently, number of beam pairs at the transmitter and receiver.

B. Unimodal Beam Alignment (UBA) Algorithm

Next, we present an algorithm whose regret matches the lower bound given in Theorem 1. Our algorithm is asymptotically optimal, and is based on UCB algorithm that uses the KL divergence as an index for arm. In particular, each arm is attached an index that resembles the KL-UCB index, but the arm selected at a given time is the arm with maximal index within the *neighborhood* of the arm that yields the highest empirical reward. Let $k(t)$ be the arm selected at time t , and $s_k(t)$ denote the number of times arm k has been selected up to time t . The empirical reward of arm k at time t is:

$$\hat{\mu}_k(t) = \begin{cases} \frac{\sum_{n=1}^t \mathbb{1}\{k(n)=k\} p_k X_k(n)}{s_k(t)} & \text{if } s_k(t) \neq 0, \\ 0 & \text{otherwise.} \end{cases} \quad (7)$$

At any time slot t , we denote by $L(t) = \arg \max_{1 \leq k \leq K} \hat{\mu}_k(t)$ the index of the arm with the highest empirical reward. $L(t)$ is referred to as the *leader* at time t . Further, we define $l_k(t) = \sum_{n=1}^t \mathbb{1}\{L(n) = k\}$ the number of times that arm k has been the leader up to time t . Now, the index of arm k at time t is defined as:

$$b_k(t) = \sup \left\{ q \in [0, p_k] : I \left(\frac{\hat{\mu}_k(t)}{p_k}, \frac{q}{p_k} \right) \leq \frac{f(t)}{s_k(t)} \right\}, \quad (8)$$

in which $f(t) = \log(l_{L(t)}(t)) + c \log(\log(l_{L(t)}(t)))$ and c is a positive number. At any time slot, the algorithm selects the arm “close” to arm $L(t)$ and with the maximum index. The pseudocode is provided in Algorithm 1 where $\gamma = 2$ is the maximum degree of the graph representing the relation between arms. Next, we provide the finite time analysis of this algorithm.

Theorem 2. *Let fix $\theta \in \mathcal{T} \cap \mathcal{U}$. For all ϵ , the regret under the proposed UBA policy and at time T is bounded by:*

$$\mathcal{R}(T) \leq (1 + \epsilon) \sum_{k \in N(k^*)} \frac{p_{k^*} \theta_{k^*} - p_k \theta_k}{I(\theta_k, \theta_{k^*})} \log(T).$$

Proof. Proof is provided in Appendix A. \square

C. Transmitter-Receiver Beam Alignment

By using the proposed UBA algorithm, the transmitter and receiver are able to refine the search space through successive rounds of beam alignment. In order to guarantee a finite time running of this algorithm, we add an additional termination condition to Algorithm 1 and continue with the data communication phase thereafter. We use the **peak-to-average ratio** as the termination condition in order to detect when the best beam direction is found. Therefore, as the UBA algorithm proceeds, we evaluate the ratio of the received energy to the average of previously received signals energy. If the ratio is higher than a threshold Ψ , we terminate the UBA algorithm and declare the beam as the best beam direction. Specifically, at time slot t , we calculate the peak-to-average ratio as follows:

$$\psi(t) = \frac{p_{k(t)}}{\frac{1}{t} \sum_{n=1}^t p_{k(n)}},$$

in which $p_{k(t)}$ denotes the energy level of beam direction selected at time t . Therefore, when the condition $\psi(t) \geq \Psi$ is

Algorithm 1 Unimodal Beam Alignment (UBA)

At time slot $t \geq 1$, select the beam direction with index $k(t)$ where:

$$k(t) = \begin{cases} L(t) & \text{if } \frac{l_{L(t)}(t)-1}{\gamma+1} \in \mathbb{N}, \\ \arg \max_{k \in N(L(t))} b_k(t) & \text{otherwise.} \end{cases} \quad (9)$$

satisfied, we declare the beam with index $k(t)$ as the optimal direction, and the UBA algorithm stops at time t . The authors in [1] have experimentally evaluated the peak-to-average ratio for LOS and NLOS situations such that $\Psi = 4$ is acceptable for detecting LOS. It should be noted that the proposed UBA scheme does not rely on the existence of LOS scenarios, while the threshold Ψ can be different under various environmental conditions. In particular, environmental conditions (e.g., blockage or reflection) alter the success probability of beam matchings, while the proposed UBA is oblivious to the underlying “physical layer” condition. In fact, based on the past observations, the UBA biases the search space towards the best beam direction.

In order to provide a complete beam alignment algorithm, similar to the IEEE 802.11ad standard, we decouple the transmitter and receiver steering such that the transmitter starts with a quasi-omnidirectional beam, while the receiver uses the UBA algorithm (instead of exhaustive search) to find the best beam direction. The process is then reversed to have the transmitter scan the space while keeping the receiver quasi-omnidirectional. As a result, we enhance the 802.11ad standard beam alignment by using the UBA algorithm instead of exhaustive search.

VI. NUMERICAL RESULTS

A. Setup

We compare the performance of the UBA algorithm with the exhaustive search scheme in which the receiver scans all different directions and samples the beam in all directions. The combination of transmitter and receiver beams that delivered the maximum power is picked as the direction of the signal. We perform the comparison of the UBA algorithm with the exhaustive search method under two different scenarios: *directional* and *quasi-directional*. Under directional conditions, probability of success is either very high or very low (e.g., in LOS scenarios). On the other hand, quasi-directional scenario occurs when the variance of success probabilities is smaller than directional situation (e.g., NLOS conditions). We evaluate the performance of beam alignment in terms of regret that measures the performance loss compared with the optimal alignment (i.e., an oracle policy). In this case, a lower regret implies a higher amount of received energy, and thus a higher accuracy in beam alignment. We also compare the beam alignment accuracy and delay overhead when the termination condition of peak-to-average ratio is used. In simulations, we fix the transmitter beam direction, and the receiver scans the angular domain to find the optimal beam direction.

B. Regret Performance

We set the vector of success probabilities as follows:

$$\theta_{\text{Directional}} = (0.99, 0.98, 0.96, 0.93, 0.9, 0.1, 0.06, 0.04);$$

$$\theta_{\text{Quasi-directional}} = (0.95, 0.9, 0.8, 0.65, 0.45, 0.25, 0.15, 0.1).$$

Figure 6 demonstrates the regret of the UBA method compared with the exhaustive beam sampling method under the directional and quasi-directional scenarios with 8 beam directions. From the results, we observe that the regret increases over time since compared with an oracle policy, the **total performance loss** keeps increasing. However, the regret curve is concave and its rate of increase, decreases with time (i.e., error decreases). In addition, exploiting the structural properties using the UBA algorithm greatly reduces the regret that is equivalent to a higher amount of received energy. This implies a higher beam alignment accuracy that is proportional with the received energy. In addition, both methods achieve a lower regret under the the directional scenario, as expected.

C. Scaling with the Size of System

Due to recent advances in antenna technologies, large directional antenna arrays with much smaller form factors can be deployed in relatively small chip areas. As a result, spatial resolution and number of the beams can be very large at the transmitter and receiver. Within this context, we investigate the effect of number of beam pairs on the performance of UBA. Figure 7 demonstrates the regret metric for $K = 8$ and $K = 16$ beam pairs. From the results, we observe that the performance of UBA scheme does not degrade with the number of beams that is a function of the number of antennas at the transmitter and receiver. This is a crucial property in massive antenna systems. Similar to Fig. 6, UBA scheme achieves a better performance compared with the exhaustive beam sampling method for both $K = 8$ and $K = 16$ beam pairs.

D. Beam Alignment Accuracy and Delay Overhead

Next, we investigate the accuracy and delay overhead of the proposed UBA algorithm combined with the peak-to-average ratio termination condition. We set the number of beams to be equal to 8 beams at the receiver, and the goal is to find the best beam (e.g., misalignment angle of zero). Using the exhaustive search method, 8 time slots is needed to examine all beams and pick the one with the highest received energy. This method is deterministic in a sense that the output is correct with the guaranteed delay overhead of 8 slots. On the other hand, our method finds the optimal beam direction with a high probability while its delay overhead is smaller than the exhaustive method. We set $\Psi = 4$ [1], and consider a scenario in which beam alignment success probability for the optimal beam is relatively small, i.e., $\theta = (0.8, 0.5, 0.35, 0.3, 0.25, 0.2, 0.15, 0.1)$. In this case, Fig. 8(a) reports the CDF of the optimal beam detection. From the results, we observe that in more than 85% of iterations, we correctly predict the optimal beam direction. The important point, however, is that our method significantly reduces the delay overhead. Figure 8(b) depicts the scatter

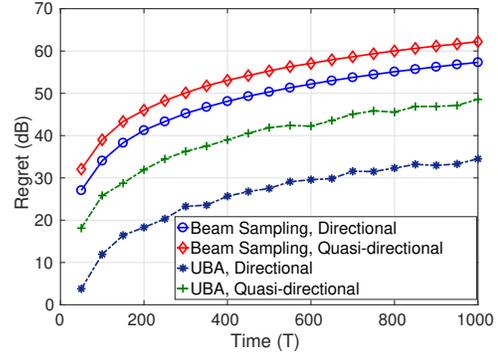


Fig. 6. Regret of UBA and exhaustive beam search with 8 beams.

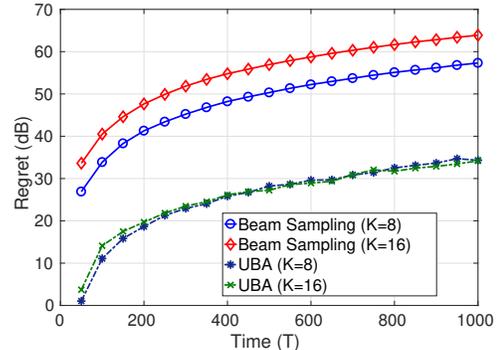


Fig. 7. Regret of UBA and the exhaustive beam search with 8 and 16 beams.

plot for detecting each beam as the optimal vs. the amount of time it takes. Size of each scatter point represents density of data. From the results, we observe that most of the beam alignment operations lead to beam 1 (i.e., high accuracy) with delay of less than 5 time slots (i.e., low overhead). Figure 8(c) also shows the CDF of delay overhead in detecting beam 1 as the optimal beam. We extend the simulation to 128 receiver beams. From the results shown in Fig. 9, we observe that delay overhead is significantly improved at the cost of some error in detecting the optimal beam direction. In fact, the delay overhead of 128 time slots (to examine each direction) is reduced to 12.5 time slots (averaged over 1000 iterations).

VII. CONCLUSION

In this paper, we investigated the beam alignment problem in mmWave systems where the transmit and receive antenna arrays require to frequently find the optimal beam pair that maximizes the received energy from beacon messages. In order to reduce the overhead of exhaustive search methods, we investigated an online stochastic optimization problem and proposed an equivalent structured Multi-Armed bandit model. In this case, the problem of finding the best beam pair is reduced to finding the optimal arm at each time slot such that the overall regret is minimized. We exploit the contextual information in order to reduce the search space, and thus the overhead of exhaustive beam selection. Thanks to the structural properties, we demonstrated that the regret bound does not depend on the size of decision space that is equal to the the number of transmit and receive beams multiplied. This

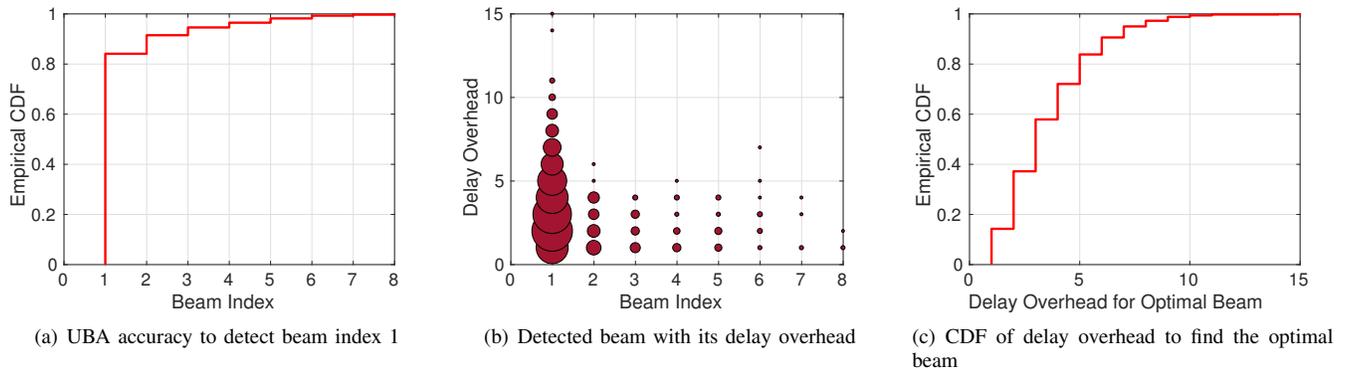


Fig. 8. Accuracy and delay overhead of the proposed UBA algorithm with the detection threshold $\Psi = 4$ and **8 beams**.

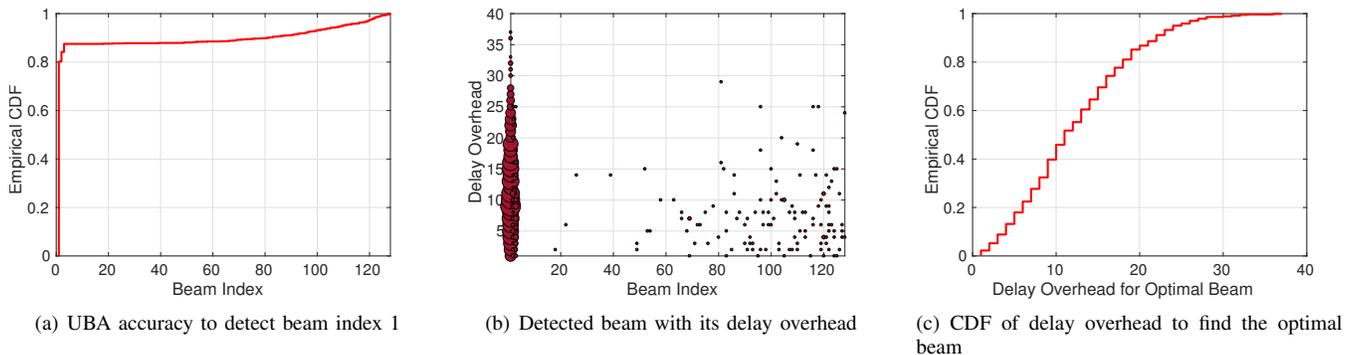


Fig. 9. Accuracy and delay overhead of the proposed UBA algorithm with the detection threshold $\Psi = 4$ and **128 beams**.

is a crucial property in MIMO settings in which the number of all combinations of transmit and receive beams grows quickly. We further proposed an asymptotically optimal algorithm for the beam alignment problem and demonstrated its performance via simulations.

ACKNOWLEDGMENT

This work was supported in part by the following NSF grants: CNS-1518916, CNS-1314822, CNS-1618566, CNS-1514260, CNS-1518829, and ONR grants: N00014-15-1-2166 and N00014-17-1-2417.

The authors would like to thank Hongliang Si and Nathan Weirich for performing the experiments.

APPENDIX A PROOF OF THEOREM 2

Proof. Similar to [20], we split the T rounds in two sets: those rounds in which the best arm k^* is the leader, i.e., $L(t) = k^*$, and those in which the leader is another arm, i.e., $L(t) \neq k^*$. Therefore:

$$\begin{aligned} \mathcal{R}(T) &= \sum_{k \neq k^*} (\mu_{k^*} - \mu_k) \mathbb{E} \left[\sum_{t=1}^T \mathbf{1}\{k(t) = k\} \right] \\ &= \sum_{k \neq k^*} (\mu_{k^*} - \mu_k) \mathbb{E} \left[\sum_{t=1}^T \mathbf{1}\{L(t) = k^* \text{ and } k(t) = k\} \right] \\ &\quad + \sum_{k \neq k^*} (\mu_{k^*} - \mu_k) \mathbb{E} \left[\sum_{t=1}^T \mathbf{1}\{L(t) \neq k^* \text{ and } k(t) = k\} \right]. \end{aligned}$$

If we consider the first term, the proposed algorithm behaves like the UCB algorithm restricted to the optimal arm and its neighborhood, and the regret upper bound is the one presented in [21], i.e., for every $\epsilon > 0$:

$$\mathcal{R}_1(T) \leq (1 + \epsilon) \sum_{k \in N(k^*)} \frac{\mu_{k^*} - \mu_k}{I(\theta_k, \theta_{k^*})} [\log(T) + \log(\log(T))] + C,$$

where C is a constant. For the second part, we have:

$$\mathcal{R}_2(T) = \sum_{k \neq k^*} (\mu_{k^*} - \mu_k) \mathbb{E} \left[\sum_{t=1}^T \mathbf{1}\{L(t) \neq k^* \text{ and } k(t) = k\} \right],$$

or $\mathcal{R}_2(T) \leq \sum_{k \neq k^*} \mathbb{E}[l_k(T)]$. Next, we provide an upper bound on the number of times that arm k has been the leader, i.e., $l_k(T)$, with $\hat{l}_k(T)$ that is the number of rounds spent with arm k as leader in the case only its neighborhood is considered during the whole time horizon T . Therefore, we have:

$$\begin{aligned} \mathcal{R}_2(T) &\leq \sum_{k \neq k^*} \mathbb{E}[l_k(T)] \leq \sum_{k \neq k^*} \mathbb{E}[\hat{l}_k(T)] = \sum_{k \neq k^*} \sum_{t=1}^T \mathbb{E}[\mathbf{1}\{L(t) = k\}] \\ &= \sum_{k \neq k^*} \sum_{t=1}^T \mathbb{E}[\mathbf{1}\{\hat{\mu}_k(t) = \max_{j \in N(k)} \hat{\mu}_j(t)\}], \end{aligned} \quad (10)$$

where, with an abuse of notations, $L(t)$ denotes the leader at round t in this modified problem where only $N(k)$ is considered. Since arm k is the leader, its empirical mean is the maximum in its neighborhood, i.e., $\hat{\mu}_k(t) \geq \hat{\mu}_{k'}(t)$ in which $k' = \arg \max_{i \in N(k)} \mu_i$. Thus,

we have: $\mathcal{R}_2(T) \leq \sum_{k \neq k^*} \sum_{t=1}^T \mathbb{E}[\mathbf{1}\{\hat{\mu}_k(t) \geq \hat{\mu}_{k'}(t)\}] = \sum_{k \neq k^*} \sum_{t=1}^T \mathbb{P}(\hat{\mu}_k(t) \geq \hat{\mu}_{k'}(t))$. Defining $\Delta_k = \max_{k' \in N(k)} \hat{\mu}_{k'} - \mu_k$ as the expected loss incurred in choosing arm k instead of its best adjacent one k' , we have:

$$\begin{aligned} \mathcal{R}_2(T) &\leq \sum_{k \neq k^*} \sum_{t=1}^T \mathbb{P}(\hat{\mu}_k(t) - \mu_k - \frac{\Delta_k}{2} - \hat{\mu}_{k'}(t) + \mu_{k'} - \frac{\Delta_k}{2} \geq 0) \\ &\leq \sum_{k \neq k^*} \left[\underbrace{\sum_{t=1}^T \mathbb{P}(\hat{\mu}_k(t) - \mu_k - \frac{\Delta_k}{2} \geq 0)}_{\mathcal{R}_{2,1}^k(T)} + \underbrace{\sum_{t=1}^T \mathbb{P}(\hat{\mu}_{k'}(t) - \mu_{k'} + \frac{\Delta_k}{2} \leq 0)}_{\mathcal{R}_{2,2}^k(T)} \right] \end{aligned}$$

For the first term, we have:

$$\begin{aligned} \mathcal{R}_{2,1}^k(T) &= \sum_{t=1}^T \mathbb{P}(\hat{\mu}_k(t) \geq \mu_k + \frac{\Delta_k}{2}) \\ &= \sum_{t=1}^T \sum_{h=1}^t \mathbb{P}(s_k(t) = h | \hat{\mu}_k(t) \geq \mu_k + \frac{\Delta_k}{2}) \mathbb{P}(\hat{\mu}_k(t) \geq \mu_k + \frac{\Delta_k}{2}) \\ &\leq \sum_{t=1}^T \sum_{h=1}^t \mathbb{P}(s_k(t) = h | \hat{\mu}_k(t) \geq \mu_k + \frac{\Delta_k}{2}) e^{-\frac{h\Delta_k^2}{2}}, \quad (11) \end{aligned}$$

where the last inequality is due to the Chernoff-Hoeffding inequality expressed as follows:

Lemma 1. (Chernoff-Hoeffding inequality) *Let X_1, \dots, X_n be random variables with common range $[0, 1]$ and such that $\mathbb{E}[X_t | X_1, \dots, X_{t-1}] = \mu$. Let $S_n = X_1 + \dots + X_n$. Then for all $a \geq 0$, we have: $\mathbb{P}(S_n \geq n\mu + a) \leq e^{-2a^2/n}$.*

Therefore, we have:

$$\begin{aligned} \mathcal{R}_{2,1}^k(T) &\leq \sum_{t=1}^T \left(\sum_{h=1}^{x_0} \mathbb{P}(s_k(t) = h | \hat{\mu}_k(t) \geq \mu_k + \frac{\Delta_k}{2}) e^{-\frac{h\Delta_k^2}{2}} \right. \\ &\quad \left. + \frac{2}{\Delta_k^2} e^{-\frac{x_0\Delta_k^2}{2}} \right), \end{aligned}$$

since $\sum_{t=x+1}^{\infty} e^{-kt} \leq \frac{1}{k} e^{-kx}$. Then, we have: $\mathcal{R}_{2,1}^k(T) \leq \sum_{t=1}^T \left(x_0 \mathbb{P}(s_k(t) \leq x_0) e^{-\frac{h\Delta_k^2}{2}} + \frac{2}{\Delta_k^2} e^{-\frac{x_0\Delta_k^2}{2}} \right)$, where we dropped the conditioning. Since the expected number of times that the non-optimal arm has been played is bounded and its variance is bounded as well, using Bernstein's inequality (provided below), we have: $\mathbb{P}(s_k(t) \leq x_0) \leq e^{-x_0/5}$, and since x_0 is lower bounded, we conclude that $\mathcal{R}_{2,1}^k(T)$ is bounded by a constant, i.e., $\mathcal{R}_{2,1}^k(T) \leq C'$.

Lemma 2. (Bernstein inequality) *Let X_1, \dots, X_n be random variables with range in $[0, 1]$ and $\sum_{t=1}^n \text{Var}[X_t | X_{t-1}, \dots, X_1] = \sigma^2$. Let $S_n = X_1 + X_2 + \dots + X_n$. Then for all $a \leq 0$, we have: $\mathbb{P}(S_n - \mathbb{E}[S_n] \geq a) \leq e^{-\frac{a^2/2}{\sigma^2 + a/2}}$.*

For the $\mathcal{R}_{2,2}^k(T)$ term, we have:

$$\begin{aligned} \mathcal{R}_{2,2}^k(T) &= \sum_{t=1}^T \mathbb{P}(\hat{\mu}_{k'}(t) \leq \mu_{k'} - \frac{\Delta_k}{2}) \\ &= \sum_{t=1}^T \sum_{h=1}^t \mathbb{P}(s_{k'}(t) = h \text{ and } \hat{\mu}_{k'}(t) \leq \mu_{k'} - \frac{\Delta_k}{2}) \\ &= \sum_{t=1}^T \sum_{h=1}^{t_b} \mathbb{P}(s_{k'}(t) = h \text{ and } \hat{\mu}_{k'}(t) \leq \mu_{k'} - \frac{\Delta_k}{2}) \\ &\quad + \sum_{t=1}^T \sum_{h=t_b+1}^t \mathbb{P}(s_{k'}(t) = h | \hat{\mu}_{k'}(t) \leq \mu_{k'} - \frac{\Delta_k}{2}) \mathbb{P}(\hat{\mu}_{k'}(t) \leq \mu_{k'} - \frac{\Delta_k}{2}) \\ &\leq \sum_{t=1}^{\infty} \mathbb{E}[\mathbf{1}\{s_{k'}(t) \leq t_b\}] + \sum_{t=1}^T \sum_{h=t_b+1}^t \mathbb{P}(\hat{\mu}_{k'}(t) \leq \mu_{k'} - \frac{\Delta_k}{2}) \\ &\leq C + \sum_{t=1}^T \sum_{h=t_b+1}^t e^{-\frac{t\Delta_k^2}{2}}. \end{aligned}$$

It is straightforward to show that $\mathcal{R}_{2,2}^k(T)$ is bounded by a constant as well. By considering the three partial results on \mathcal{R}_1 , $\mathcal{R}_{2,1}^k$, and $\mathcal{R}_{2,2}^k$, we have:

$$\mathcal{R}(T) \leq \mathcal{R}_1(T) + \sum_{k \neq k^*} \mathcal{R}_{2,1}^k(T) + \mathcal{R}_{2,2}^k(T),$$

and the theorem statement follows. \square

REFERENCES

- [1] T. Nitsche, A. B. Flores, E. W. Knightly, and J. Widmer, "Steering with eyes closed: mm-wave beam steering without in-band measurement," in *Computer Communications (INFOCOM), IEEE Conference on*. IEEE, 2015, pp. 2416–2424.
- [2] IEEE, "IEEE 802.11ad, amendment 3: Enhancements for very high throughput in the 60 GHz band," IEEE 802.11 Working Group, 2012.
- [3] L. Zhou and Y. Ohashi, "Efficient codebook-based MIMO beamforming for millimeter-wave WLANs," in *Personal Indoor and Mobile Radio Communications (PIMRC), 2012 IEEE 23rd International Symposium on*. IEEE, 2012, pp. 1885–1889.
- [4] Y. Zhu, Z. Zhang, Z. Marzi, C. Nelson, U. Madhow, B. Y. Zhao, and H. Zheng, "Demystifying 60GHz outdoor picocells," in *Proceedings of the 20th annual international conference on Mobile computing and networking*. ACM, 2014, pp. 5–16.
- [5] T. L. Lai and H. Robbins, "Asymptotically efficient adaptive allocation rules," *Advances in applied mathematics*, vol. 6, no. 1, pp. 4–22, 1985.
- [6] Y. Qi and M. Nekovee, "Coordinated initial access in millimeter wave standalone networks," *arXiv preprint arXiv:1605.03337v1*, 2016.
- [7] V. Desai, L. Krzymien, P. Sartori, W. Xiao, A. Soong, and A. Alkhateeb, "Initial beamforming for mmwave communications," in *Signals, Systems and Computers, 2014 48th Asilomar Conference on*. IEEE, 2014, pp. 1926–1930.
- [8] J. Singh and S. Ramakrishna, "On the feasibility of codebook-based beamforming in millimeter wave systems with multiple antenna arrays," *IEEE transactions on Wireless Communications*, vol. 14, no. 5, pp. 2670–2683, 2015.
- [9] C. N. Barati, S. A. Hosseini, S. Rangan, P. Liu, T. Korakis, S. S. Panwar, and T. S. Rappaport, "Directional cell discovery in millimeter wave cellular networks," *IEEE Transactions on Wireless Communications*, vol. 14, no. 12, pp. 6664–6678, 2015.
- [10] S. Hur, T. Kim, D. J. Love, J. V. Krogmeier, T. A. Thomas, and A. Ghosh, "Millimeter wave beamforming for wireless backhaul and access in small cell networks," *Communications, IEEE Transactions on*, vol. 61, no. 10, pp. 4391–4403, 2013.
- [11] A. Ali, N. González-Prelcic, and R. W. Heath Jr, "Millimeter wave beam-selection using out-of-band spatial information," *arXiv preprint arXiv:1702.08574*, 2017.
- [12] H. Hassanieh, O. Abari, M. Rodriguez, M. Abdelghany, D. Katabi, and P. Indyk, "Agile millimeter wave networks with provable guarantees," *arXiv preprint arXiv:1706.06935*, 2017.

- [13] M. Hashemi, C. E. Koksai, and N. B. Shroff, "Energy-efficient power and bandwidth allocation in an integrated sub-6 GHz-millimeter wave system," *arXiv preprint arXiv:1710.00980*, 2017.
- [14] —, "Out-of-band millimeter wave beamforming and communications to achieve low latency and high energy efficiency in 5G systems," *IEEE Transactions on Communications*, 2017.
- [15] H. Robbins, "Some aspects of the sequential design of experiments," in *Herbert Robbins Selected Papers*. Springer, 1985, pp. 169–177.
- [16] S. Bubeck, N. Cesa-Bianchi *et al.*, "Regret analysis of stochastic and nonstochastic multi-armed bandit problems," *Foundations and Trends® in Machine Learning*, vol. 5, no. 1, pp. 1–122, 2012.
- [17] E. W. Cope, "Regret and convergence bounds for a class of continuum-armed bandit problems," *IEEE Transactions on Automatic Control*, vol. 54, no. 6, pp. 1243–1253, 2009.
- [18] Y. Y. Jia and S. Mannor, "Unimodal bandits," in *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, 2011, pp. 41–48.
- [19] R. Combes and A. Proutiere, "Unimodal bandits: Regret lower bounds and optimal algorithms," in *International Conference on Machine Learning*, 2014, pp. 521–529.
- [20] S. Paladino, F. Trovò, M. Restelli, and N. Gatti, "Unimodal thompson sampling for graph-structured arms," in *AAAI*, 2017, pp. 2457–2463.
- [21] O. Cappé, A. Garivier, O.-A. Maillard, R. Munos, G. Stoltz *et al.*, "Kullback–leibler upper confidence bounds for optimal sequential allocation," *The Annals of Statistics*, vol. 41, no. 3, pp. 1516–1541, 2013.