

Structured Reinforcement Learning for Delay-Optimal Data Transmission in Dense mmWave Networks

Shufan Wang, Guojun Xiong, *Student Member, IEEE*, Shichen Zhang, Huacheng Zeng, *Senior Member, IEEE*, Jian Li, *Member, IEEE*, Shivendra Panwar, *Fellow, IEEE*

Abstract—We study the data packet transmission problem (mmDPT) in dense cell-free millimeter wave (mmWave) networks, i.e., users sending data packet requests to access points (APs) via uplinks and APs transmitting requested data packets to users via downlinks. Our objective is to minimize the average delay in the system due to APs’ limited service capacity and unreliable wireless channels between APs and users. This problem can be formulated as a restless multi-armed bandits problem with fairness constraint (RMAB-F). Since finding the optimal policy for RMAB-F is intractable, existing learning algorithms are computationally expensive and not suitable for practical dynamic dense mmWave networks. In this paper, we propose a structured reinforcement learning (RL) solution for mmDPT by exploiting the inherent structure encoded in RMAB-F. To achieve this, we first design a low-complexity and provably asymptotically optimal index policy for RMAB-F. Then, we leverage this structure information to develop a structured RL algorithm called mmDPT-TS, which provably achieves an $\tilde{O}(\sqrt{T})$ Bayesian regret. More importantly, mmDPT-TS is computation-efficient and thus amenable to practical implementation, as it fully exploits the structure of index policy for making decisions. Extensive emulation based on data collected in realistic mmWave networks demonstrate significant gains of mmDPT-TS over existing approaches.

Index Terms—Data Packet Transmission, Dense mmWave Networks, Structured Reinforcement Learning, Index Policy, Restless Multi-Armed Bandits

I. INTRODUCTION

MILLIMETER wave (mmWave) is a key technology for current 5G and beyond wireless networks [1]–[3]. It offers multi-GHz bandwidth of licensed and unlicensed spectrum for communications. As expected, it will play a crucial role in dealing with increased multimedia traffic, and emerging applications such as multi-user wireless virtual reality (VR) for

Manuscript received September 08, 2023; revised March 31, 2024 and June 11, 2024; accepted June 14, 2024. This work was supported in part by the National Science Foundation (NSF) grants 2148309 and 2315614; in part by the Office of the Under Secretary of Defense, Research and Engineering (OUSD R&E), National Institute of Standards and Technology (NIST), and industry partners as specified in the Resilient & Intelligent NextG Systems (RINGS) program; in part by the U.S. Army Research Office (ARO) under Grant W911NF-23-1-0072, and the U.S. Department of Energy (DOE) under Grant DE-EE0009341; and in part by the New York State Center for Advanced Technology in Telecommunications (CAT) and New York University (NYU) Wireless.

S. Wang, G. Xiong, and J. Li are with Stony Brook University, Stony Brook, NY, 11794. E-mail: {shufan.wang, guojun.xiong, jian.li.3}@stonybrook.edu

S. Zhang, H. Zeng are with Michigan State University, East Lansing, MI, 48824. E-mail: {sczhang, hzeng}@msu.edu

S. Panwar is with New York University, Brooklyn, NY, 11201. E-mail: sp1832@nyu.edu




Fig. 1: A dense mmWave network in a small conference room, where the dashed lines indicate communications between APs and users, i.e., users sending data packet requests to APs and APs transmitting real data packets to users. See Section V for more details on our mmWave testbed.

education, multi-player games and professional training, where high bandwidth data must be streamed to each user with low latency [4], [5].

Realizing this vision requires a dense deployment of many access points (APs) in a mmWave network and an efficient *data packet transmission policy*. Such a policy determines to send the data packet requests from users to the mmWave APs via uplink communication, which in turn transmit the requested data packets to users via downlink communication. Data packet transmission plays a pivotal role in enhancing load balancing, spectrum efficiency, energy efficiency of mmWave networks, and hence has gained much interest in recent years for the purpose of maximizing spectral [6]–[8] and energy efficiencies [9], [10]. Unfortunately, mmWave communication does not perform well in *dynamic environments* due to its vulnerability to blockage, sensitivity to mobility, and time-varying channel conditions. These factors lead to an intermittent link connectivity between a user and an AP, necessitating a dense deployment of APs to maintain communication reliability [11], [12].

In this paper, we consider such a dense, cell-free mmWave network where a set of APs serve a population of users in one area (e.g., a conference room, a concert hall, or a classroom). Suppose that all APs are reachable for all users. At each time, each user generates a data packet request, which is sent to one AP via the uplink communication. The corresponding AP then transmits the requested data packet to the user via the downlink communication. Since only the data packet request is sent from users to APs through uplinks while the real

data packets are transmitted from APs to users via downlinks, we assume that the uplink communication is reliable [11]–[13] and the downlink communication is unreliable. Then, an important problem is: *for each data packet request generated by a user, which AP should it be sent to so as to minimize the average delay due to the AP’s limited service capacity and the unreliable downlink communications via which the requested data packet is transmitted from AP to the user?*

Consider the example shown in Figure 1 of a conference room with 4 APs and 30 users. This will be used as our running motivation and our mmWave testbed environments in Section V, but our model and proposed solutions will not be limited to this scenario. We are interested in designing a data packet transmission policy to minimize *the average delay* in the system due to APs’ limited service capacity and unreliable wireless channels between APs and users. In addition to minimizing the average delay, ensuring fairness among users is also a key design concern in wireless networks [14]–[17]. To this end, we model the above data packet transmission problem (mmDPT) in a dense mmWave network as a *restless multi-armed bandits problem with fairness guarantee* (RMAB-F)¹, which is a generalization of the classical restless multi-armed bandits problem (RMAB) [18]. Our objective is to develop low-complexity reinforcement learning (RL) algorithms to solve this RMAB-F without the knowledge of system dynamics (e.g., the unknown data packet arrivals and time-varying mmWave channel qualities).

Limitations of Existing Methods. Although online RMAB has gained many efforts, existing solutions cannot be directly applied to our RMAB-F. A key challenge is that off-the-shelf RL algorithms, e.g., colored-UCRL2 [19] and Thompson sampling methods [20], [21], suffer from an exponential computational complexity, and their regret bounds grow exponentially with the size of state space. This is due to the fact that these algorithms need to repeatedly solve complicated Bellman equations for making decisions, and hence appear too slow for practical use, especially in highly dynamic mmWave environments. Many recent efforts have been devoted to developing low-complexity RL algorithms with order-of-optimal regret for online RMAB [22]–[28]; however, many challenges remain unsolved. For example, multi-timescale stochastic approximation algorithms [22]–[24] suffer from slow convergence and have no regret guarantee, and [25], [26] considered a finite-horizon setting while we focus on an infinite-horizon average-award setting in this paper. Exacerbating these limitations is the fact that none of them were designed with fairness constraints in mind. For example, [25]–[28] only focused on minimizing costs/delay in RMAB, and many existing RL or deep RL based policies for mmWave focused on maximizing throughput [29]–[32] with no finite-time performance analysis, while the controller in our RMAB-F faces a **new dilemma** on how to manage the balance between minimizing average delay and satisfying the fairness requirement. This adds a new layer of challenge to designing low-complexity RL algorithms for RMAB that is already quite challenging.

¹We refer to our mmDPT problem as a RMAB-F, and will interchangeably/equivalently use these two terms in the rest of this paper.




Fig. 2: Measurement setup and experiment scenario for data packet transmission in dense mmWave networks.

Structured RL for mmDPT. The lack of theoretical understanding on how to design efficient RL algorithms for RMAB-F or mmDPT motivates us to fill this gap by proposing *structured RL solutions* in this paper. Specifically, our structured RL solutions operate on a much smaller dimensional subspace by exploiting the inherent structure encoded in RMAB-F. This requires us to *first* design a low-complexity yet provably optimal index policy for RMAB-F, and *then* RL algorithms that leverage the structure of index policies for making decisions to reduce the high computational complexity and exponential factor in regret analysis. We summarize our contributions as follows:

- **Provably Optimal Index Policy.** We first develop a low-complexity index policy for RMAB-F to address the dimensional concerns when the system dynamics are known in Section III. Specially, we leverage a linear programming (LP) based approach to obtain a relaxed problem of RMAB-F, which is formulated as a LP using occupancy measures [33]. We then construct a mmDPT Index based on the occupancy measures obtained from the LP. Finally, we propose a low-complexity mmDPT Index Policy by carefully coupling the scheduling and fairness constraints to address the new dilemma via the above mmDPT Index. We offer a proof to show that mmDPT Index Policy is asymptotically optimal.
- **Structured RL Algorithm.** We further develop a low-complexity RL algorithm for RMAB-F without the knowledge of system dynamics in Section IV. Different from aforementioned off-the-shelf RL algorithms that either contend directly with an extremely large state space or do not incorporate the fairness constraint, we propose mmDPT-TS, a structured Thompson sampling (TS) method that learns to leverage the inherent structure in RMAB-F via our near-optimal mmDPT Index Policy for making decisions. We show that mmDPT-TS achieves an optimal sub-linear Bayesian regret with a low computational complexity, and hence can be easily implemented in realistic mmWave networks. To the best of our knowledge, our work is the first to develop a structured RL algorithm with low-complexity and order-of-optimal regret in the context of delay-optimal data packet transmission in dense, cell-free mmWave networks. We note that our proposed frameworks of designing low-complexity index policies and structured RL algorithms are very general, and can be applied to various large-scale combinatorial problems with fairness constraints.
- **Evaluations on 60GHz mmWave Testbed.** We build




Fig. 3: The measured error vector magnitude (EVM) of decoded signal constellations at users which receive data packet from APs (via downlinks) in our mmWave testbed (See Section V). The three curves correspond to the three transmissions in Figure 1.

a 60GHz mmWave testbed using software-defined radio (SDR) devices. The mmWave device is equipped with a planar antenna with 4×8 patch elements. Our evaluation is conducted in a conference room, and Figure 2 shows a photo of our testbed, see Section V for details. Experimental results using data collected from our mmWave testbed demonstrate that our mmDPT-TS produces significant performance gains over existing approaches.

II. SYSTEM MODEL AND PROBLEM FORMULATION

In this section, we present the system model and formulate the delay minimization problem of data packet transmission in dense, cell-free mmWave networks (mmDPT).

A. System Model

We consider a dense, cell-free mmWave network with a set of $\mathcal{N} = \{1, \dots, N\}$ mmWave APs serving one area (e.g., a conference room, a concert hall, or a classroom), where there is a set of $\mathcal{M} = \{1, \dots, M\}$ users. Due to the dense nature, all APs are reachable for all users [34]–[37]. Consider the example shown in Figure 1 of a conference room with 4 APs and 30 users. This will be used as our running motivation and our mmWave testbed environments in Figure 2, but our model and proposed solutions will not be limited to this scenario.

Time is divided into multiple units with each unit called a “slot”, which is indexed by $t \in \mathcal{T} = \{1, \dots, T\}$. At time slot t , user m generates a data packet request with probability p_m , which is sent to one AP for processing through uplinks available between APs and users. Since only requests are sent from users to APs, we assume that the uplink communication is reliable without delay. The rationality of this assumption is that only the request signal is sent via uplink communication, and the request indicator is very small-sized (typically in the order of a few bits to a few tens of bits). The exact size may depend on the network configuration or technology leverage. As a result, the delay due to uplink communication is negligible [11]–[13], [34], [35].

Without loss of generality (W.l.o.g.), we assume that requests generated by each user are independent of each other in each time slot. Upon receiving the request, the AP processes and transmits the requested data packet to that user through downlinks, which are often unreliable. A centralized controller is in charge of such a data packet transmission problem in the dense mmWave network in consideration. This is mainly due to the fact that we mainly focus on a dense mmWave

network such as a conference room as shown in Figure 1. Similar assumption applied to user association [34], [35] and beam alignment [36], [37] in a dense mmWave network.

In our system model, we add a “request queue” to each AP n for each user m to store the number of data packet requests sent to AP n at time slot t . We denote the queue length as S_{mn}^t . The rationality of our model (i.e., each AP maintaining a request queue for each user) is that the number of data packet requests to one AP may be larger than its service capacity limited by both computation and communication resources due to the dense nature of mmWave networks, and the data packet requests may not be processed immediately and hence there will be a service delay for users. Another point is due to the fact that the wireless channels between APs and users are unreliable. As a result, the time that AP n takes to transmit the requested data packet to user m is a random variable, which heavily relies on the mmWave channel quality, denoted as q_{mn}^t , between AP n and user m at time slot t . For example, Figure 3 shows the error vector magnitude (EVM) measured at the receiver on our mmWave testbed, which vary significantly over time. Assume each processed packet is equally sized with Q bits. Let D_{mn}^t be the number of data packets that is successfully delivered from AP n to user m via downlinks at time slot t . We model D_{mn}^t as a random variable with probability distribution $\mathbb{P}(\cdot|q_{mn}^t)$ to reflect the randomness of wireless fading. Denote $C(q_{mn}^t)$ as the throughput of the wireless channel between AP n and user m at time slot t . Therefore, the distribution of D_{mn}^t can be formally given by

$$\mathbb{P}(D_{mn}^t = d|q_{mn}^t) = \mathbb{P}((d+1)Q > C(q_{mn}^t) \geq dQ). \quad (1)$$

B. MDP-based Problem Formulation

We formulate the delay minimization problem for mmDPT for the above model as a Markov decision process (MDP) [38].

State. We denote the queue length of data packet requests from user m at time slot t as $\mathbf{S}_m^t := (S_{m1}^t, \dots, S_{mN}^t)$, where S_{mn}^t is the number of data packet requests from user m sent to AP n at time slot t as described above. Let $\mathbf{S}^t := (\mathbf{S}_1^t, \dots, \mathbf{S}_M^t)$. W.l.o.g., we assume $S_{mn}^t \leq S_{\max}, \forall m, n, t$, where S_{\max} is the maximum number of data packet requests from a user sent to an AP, and can be arbitrarily large but bounded. For ease of readability, we denote the finite state space in our model as \mathcal{S} .

Action. Action $A_{mn}^t = 1$ means that the centralized controller determines to send the data packet request from user m to AP n via the uplink channel at time slot t ; and $A_{mn}^t = 0$, otherwise. Denote $\mathcal{A} = \{0, 1\}$ and let $\mathbf{A}_m^t := (A_{m1}^t, \dots, A_{mN}^t)$, $\mathbf{A}^t := (\mathbf{A}_1^t, \dots, \mathbf{A}_M^t)$. Since at most one data packet request can be sent from a user to an AP at each time slot, we have

$$\sum_{n \in \mathcal{N}} A_{mn}^t \leq 1, \quad \forall m \in \mathcal{M}, t \in \mathcal{T}. \quad (2)$$

In addition, we impose a *fairness* constraint among APs (e.g., due to resource constraints). Specifically, at most B data

packet requests can be simultaneously sent to any AP at any time slot, i.e.,

$$\sum_{m \in \mathcal{M}} A_{mn}^t \leq B, \quad \forall n \in \mathcal{N}, t \in \mathcal{T}. \quad (3)$$

A data packet transmission policy π in a dense mmWave network maps the states of all queues \mathbf{S}^t to transmission decisions \mathbf{A}^t , i.e., $\mathbf{A}^t = \pi(\mathbf{S}^t)$. Denote the set of all feasible policies as Π .

Controlled Transition Kernel. As aforementioned, when there are $S_{mn}^t = S$ data packet requests in the queue, AP n may process and successfully transmit $d \leq S$ packets to user m , which occurs with probability $\mathbb{P}(D_{mn}^t = d | q_{mn}^t)$ as defined in (1). If a new data packet request is generated, and sent to AP n at the same time, then the length of corresponding request queue becomes $S + 1 - d$. More precisely, for $\forall d \in [0, S_{mn}^t]$, we have $\mathbb{P}(S_{mn}^{t+1} = S + 1 - d | S_{mn}^t = S, A_{mn}^t = 1) = p_m \mathbb{P}(D_{mn}^t = d | q_{mn}^t)$. Otherwise, the queue length becomes $S - d$, i.e., $\mathbb{P}(S_{mn}^{t+1} = S - d | S_{mn}^t = S, A_{mn}^t = 0) = \mathbb{P}(D_{mn}^t = d | q_{mn}^t)$. Similarly, if AP n can successfully transmit more than S_{mn}^t packets, i.e., $d > S_{mn}^t$, then we have $\mathbb{P}(S_{mn}^{t+1} = 1 | S_{mn}^t = S, A_{mn}^t = 1) = p_m \left(1 - \sum_{d=0}^{S_{mn}^t} \mathbb{P}(D_{mn}^t = d | q_{mn}^t)\right)$, when a new data packet request from user m is sent to AP n at the same time; and otherwise $\mathbb{P}(S_{mn}^{t+1} = 0 | S_{mn}^t = S, A_{mn}^t = 0) = 1 - \sum_{d=0}^{S_{mn}^t} \mathbb{P}(D_{mn}^t = d | q_{mn}^t)$. The overall transition probability from user m to AP n is denoted as $P_{mn}(s|s', a')$.

Data Packet Transmission Problem. Our objective is to design a policy π that minimizes the *average delay* in a dense, cell-free mmWave network due to APs' limited service capacity and unreliable wireless channels between APs and users, while ensuring that each data packet request can only be sent to one AP, and no more than B data packet requests can be sent to any AP at any time slot. By Little's Law, the average delay minimization problem is equivalent to minimizing the average total number of requests in the system. Therefore, the data packet transmission problem in dense mmWave networks (mmDPT) can be formulated as the following MDP:

$$\text{mmDPT} : \min_{\pi \in \Pi} \limsup_{T \rightarrow \infty} \mathbb{E}_{\pi} \left[\frac{1}{T} \sum_{t=1}^T \left(\sum_{n=1}^N \sum_{m=1}^M S_{mn}^t \right) \right] \\ \text{subject to Constraints (2) and (3),} \quad (4)$$

where the subscript denotes the fact that the expectation is taken with respect to the measure induced by policy π . Problem mmDPT (4) is an example of the RMAB-F, and in theory it can be solved optimally as an infinite-horizon average cost per stage problem using relative value iteration [38]. However, this approach suffers from the curse of dimensionality, i.e., the computational complexity grows exponentially in the size of state space as a function of the number of user M , rendering such a solution impractical. In addition, this approach lacks insight for the solution structure. We overcome this difficulty by developing an index-based policy that is computationally appealing and provably optimal.

III. INDEX POLICY DESIGN AND ANALYSIS

We now propose an index policy for Problem mmDPT (4). We begin by introducing a so-called "relaxed problem", which

can be posed as a LP problem. The solution to this LP forms the building block of our proposed index policy, which we prove to be asymptotically optimal.

A. The Relaxed Problem

Following Whittle's approach [18], we first relax the instantaneous constraints in Problem mmDPT (4) to average constraints, and obtain the following "relaxed problem":

$$\min_{\pi \in \Pi} \limsup_{T \rightarrow \infty} \mathbb{E}_{\pi} \left[\frac{1}{T} \sum_{t=1}^T \left(\sum_{n=1}^N \sum_{m=1}^M S_{mn}^t \right) \right] \\ \text{s.t.} \quad \limsup_{T \rightarrow \infty} \mathbb{E}_{\pi} \left[\frac{1}{T} \sum_{t=1}^T \sum_{n=1}^N A_{mn}^t \right] \leq 1, \quad \forall m \in \mathcal{M}, \\ \limsup_{T \rightarrow \infty} \mathbb{E}_{\pi} \left[\frac{1}{T} \sum_{t=1}^T \sum_{m=1}^M A_{mn}^t \right] \leq B, \quad \forall n \in \mathcal{N}. \quad (5)$$

It is clear that the optimal value achieved by (5) is a lower bound of that achieved by Problem mmDPT (4). It is also known that the relaxed problem (5) can be reduced to an equivalent LP using occupancy measures [33].

Definition 1. The occupancy measure Ω_{π} of a stationary policy π for the infinite-horizon MDP is defined as the expected average number of visits to each state-action pair (s, a) , i.e.,

$$\Omega_{\pi} = \left\{ \omega_{mn}(s, a) \triangleq \lim_{T \rightarrow \infty} \frac{1}{T} \mathbb{E}_{\pi} \left(\sum_{t=1}^T \mathbb{1}(S_{mn}^t = s, A_{mn}^t = a) \right) : \forall m \in \mathcal{M}, n \in \mathcal{N}, s \in \mathcal{S}, a \in \mathcal{A} \right\}. \quad (6)$$

It can be easily checked that the occupancy measure satisfies $\sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \omega_{mn}(s, a) = 1$, and hence $\omega_{mn}, \forall m \in \mathcal{M}, n \in \mathcal{N}$ is a probability measure. Using this definition, the relaxed problem (5) can be equivalently reformulated as a LP [33]:

$$\min_{\omega_{mn} \in \Omega_{\pi}} \sum_{n=1}^N \sum_{m=1}^M \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \omega_{mn}(s, a) s \quad (7a)$$

$$\text{s.t.} \quad \sum_{n=1}^N \sum_{s \in \mathcal{S}} \omega_{mn}(s, 1) \leq 1, \quad \forall m \in \mathcal{M}, \quad (7b)$$

$$\sum_{m=1}^M \sum_{s \in \mathcal{S}} \omega_{mn}(s, 1) \leq B, \quad \forall n \in \mathcal{N}, \quad (7c)$$

$$\sum_{s' \in \mathcal{S}} \sum_{a \in \mathcal{A}} \omega_{mn}(s, a) P_{mn}(s' | s, a), \\ = \sum_{s' \in \mathcal{S}} \sum_{a \in \mathcal{A}} \omega_{mn}(s', a) P_{mn}(s | s', a), \quad \forall s \in \mathcal{S}, \quad (7d)$$

$$\sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \omega_{mn}(s, a) = 1, \quad \forall m \in \mathcal{M}, n \in \mathcal{N}, \quad (7e)$$

where (7b) and (7c) are restatements of constraints (2) and (3), respectively; (7d) represents the fluid transition of the occupancy measure, which holds due to the ergodic theorem for finite MDPs [38], [39], that under optimal solutions, the occupancy measure will be stable under the transition, where fluid in rate for a state-action occupancy measure equals to the

fluid out rate; and (7e) follows from the fact that the occupancy measure is a probability measure.

Let $\omega^* = \{\omega_{mn}^*(s, a) : \forall m \in \mathcal{M}, n \in \mathcal{N}, s \in \mathcal{S}, a \in \mathcal{A}\}$ be an optimal solution to the above LP (7a)-(7e). We now construct a Markovian stationary policy $\chi^* = \{\chi_{mn}^*(s, a), \forall m \in \mathcal{M}, n \in \mathcal{N}\}$ from ω^* as follows: if the number of requests in the queue from user m at AP n at time slot t is s , then $\chi_{mn}^*(s, a)$ chooses action a with a probability equal to

$$\chi_{mn}^*(s, a) := \frac{\omega_{mn}^*(s, a)}{\sum_{a \in \mathcal{A}} \omega_{mn}^*(s, a)}. \quad (8)$$

Unfortunately, the above policy (8) does not always provide a feasible solution to Problem mmDPT (4). This is due to the fact that both request transmission constraints in (4) must be strictly met at each time slot, instead of just in the average sense as in (5). Exacerbating this issue is that with a randomized policy, both relaxed constraints may be violated severely during each time slot, resulting in poor performance. To overcome this challenge, we next introduce a computationally appealing index policy for Problem mmDPT (4).

B. The mmDT Index Policy

Conventional index policies including Whittle index policy [18] and many others [13], [25], [26], [28], [29], [39]–[44] simply schedule a request to the highest indexed AP, i.e., AP $n^* = \arg \max_n \chi_{mn}^*(s, 1)$ forms channel, through via a request from user m at time slot t is sent to AP n^* . Unfortunately, such a simple index policy will not work for Problem mmDPT (4) since it only accounts for constraint (2) but ignores the new dilemma faced by the controller in our RMAB-F, which is introduced by the fairness constraint (3), i.e., at most B requests can be sent to one AP at each time slot. Intuitively, to capture both constraints, the transmission decisions should involve some couplings between requests and APs. For simplicity, we denote

$$\phi_{mn}(s) = \chi_{mn}^*(s, 1), \quad (9)$$

and call it the mmDPT Index for requests from user m at AP n when its state is s . To address the aforementioned issue when applying existing index policies, our mmDPT Index Policy prioritizes the request from user m at time slot t to AP n according to a decreasing order of their mmDPT Index, and transmits requests to APs based on mmDPT Index as long as both constraints (2) and (3) are satisfied.

Specifically, at each time slot t , we first locate the users with new generated requests, and denote the set of these users as $\mathcal{M}^t := \{m | \text{user } m \text{ generates new requests}\}$. We construct the mmDPT Index set $\mathcal{I}(t) := \{\phi_{mn}(S_{mn}^t), \forall m \in \mathcal{M}^t, n \in \mathcal{N}\}$ with all elements in $\mathcal{I}(t)$ sorted in a decreasing order. Denote the largest index in $\mathcal{I}(t)$ as $\phi_{m'n'}(S_{m'n'}^t)$. Then we check if B requests have been sent to AP n' , which leads two cases: i) if not, AP n' activates channel with user m' and remove all indices related with m' , i.e., $\mathcal{I}(t) = \mathcal{I}(t) \setminus \{\phi_{m'n}(S_{m'n}^t), \forall n \in \mathcal{N}\}$; ii) otherwise, we remove all indices related with AP n' such that $\mathcal{I}(t) = \mathcal{I}(t) \setminus \{\phi_{mn'}(S_{mn'}^t), \forall m \in \mathcal{M}^t\}$. We repeat this process until $\mathcal{I}(t)$ is empty (i.e. all requests are satisfied).

Algorithm 1 mmDPT Index Policy

- 1: Construct LP (7a)-(7e) and solve the occupancy measure $\omega_{mn}^*(s, a), \forall m, n, s, a$;
 - 2: Compute $\chi_{mn}^*(s, a)$ according to (8) and construct mmDPT Index $\phi_{mn}(s) = \chi_{mn}^*(s, 1)$ in (9);
 - 3: **for** At each time slot t **do**
 - 4: Locate $\mathcal{M}^t := \{m | \text{user } m \text{ generates new requests}\}$;
 - 5: Construct the mmDPT Index set $\mathcal{I}(t) := \{\phi_{mn}(S_{mn}^t), \forall m \in \mathcal{M}^t, n \in \mathcal{N}\}$ with elements sorted in a decreasing order;
 - 6: **while** $\mathcal{I}(t)$ is non-empty **do**
 - 7: Find the largest index $\phi_{m'n'}(S_{m'n'}^t)$ in $\mathcal{I}(t)$;
 - 8: **if** Fewer than B requests transmitted to AP n' **then**
 - 9: AP n' activates channel and transmit request from user m' , and remove all indices related with m' , i.e., $\mathcal{I}(t) = \mathcal{I}(t) \setminus \{\phi_{m'n}(S_{m'n}^t), \forall n \in \mathcal{N}\}$;
 - 10: **else**
 - 11: Remove all indices related with n' such that $\mathcal{I}(t) = \mathcal{I}(t) \setminus \{\phi_{mn'}(S_{mn'}^t), \forall m \in \mathcal{M}^t\}$.
 - 12: **end if**
 - 13: **end while**
 - 14: **end for**
-

Remark 1. Unlike Whittle-based policies [18], [40], [45], [46], our mmDPT Index Policy does not require the indexability condition, which is often hard to establish [47]. Like Whittle-based policies, our mmDPT Index Policy is computationally efficient since it is merely based on solving a LP, which can be efficiently solved in polynomial time [48]–[50], and we leverage the Gurobi Optimizer [51] in our experiments. A line of works [25], [26], [41]–[43] designed index policies without indexability requirement for finite-horizon RMAB, and hence cannot be applied to our infinite-horizon average-cost formulation in Problem mmDPT (4). We note that the design of our index policy is largely inspired by the LP based approach in [28]. However, [28] only accounted for constraint (2), while our mmDPT Index Policy faces the new dilemma due to fairness constraint (3). Further distinguishing our work is that we propose a structured RL algorithm via Thompson sampling with a provably sub-linear Bayesian regret in Section IV.

C. Asymptotic Optimality

We now show that our mmDPT Index Policy is asymptotically optimal in the same asymptotic regime as that in Whittle [18] and many others [39], [40], [52]. With some abuse of notation, let the number of users and APs be ρM and ρN , and the resource constraint be ρB in the asymptotic regime with $\rho \rightarrow \infty$. In other words, we consider M classes of users with each class containing ρ , and similarly for the APs and fairness constraint. Denote $X_{mn}^\rho(\pi^*, s, a; t)$ as the number of requests from class- m users with the state at class- n APs being s and action a being taken at time slot t under mmDPT Index Policy π^* . We will be interested in the this fluid-scaling process with parameter ρ , and define the expected long-term average cost as $V_{mn}^\rho := \limsup_{T \rightarrow \infty} \frac{1}{T} \mathbb{E}_{\pi^*} \sum_{t=1}^T \sum_{n=1}^N \sum_{m=1}^M \sum_{(s,a)} s \frac{X_{mn}^\rho(\pi^*, s, a; t)}{\rho}$.

Our mmDPT Index Policy π^* is asymptotically optimal only when $V_{\pi^*}^\rho \leq V_\pi^\rho, \forall \pi \in \Pi$. W.l.o.g., we let π^{opt} denote the optimal policy for Problem mmDPT (4). Before presenting our main result in this section, we first state the following technical condition called “global attractor” [52].

Definition 1. *An equilibrium point $X^{\rho,*}/\rho$ under mmDPT Index Policy π^* is a global attractor for the process $X^\rho(\pi^*; t)/\rho$, if, for any initial point $X^\rho(\pi^*; 0)/\rho$, the process $X^\rho(\pi^*; t)/\rho$ converges to $X^{\rho,*}/\rho$.*

The global attractor indicates that all trajectories converge to $X^{\rho,*}$. Though it may be difficult to establish analytically that a fixed point is a global attractor for the process [39], such an assumption has been widely made in [39], [40], [45], [52] and is only verified numerically. Our experimental results in Section V show that such convergence indeed occurs for our mmDPT Index Policy π^* .

Theorem 1. *Our mmDPT Index Policy π^* is asymptotically optimal under Definition 1, i.e., $\lim_{\rho \rightarrow \infty} V_{\pi^*}^\rho - V_{\pi^{opt}}^\rho = 0$.*

IV. STRUCTURED REINFORCEMENT LEARNING

The computation of mmDPT Index Policy requires the knowledge of transition probabilities associated with MDPs (see Section II-B). Unfortunately, the mmWave environments are highly dynamic and these parameters are often unknown and time varying. Hence, we now consider to learn the mmDPT (i.e., RMAB-F) without the knowledge of system dynamics. Our goal is to develop a low-complexity structured RL algorithm and characterize its finite-time performance.

A. Structured RL Algorithm: mmDPT-TS

Algorithm Overview. We adapt the Thompson Sampling (TS) method to our problem. Specifically, we design a structured TS algorithm via mmDPT Index Policy awareness, entitled mmDPT-TS, as summarized in Algorithm 2. For ease of expression, denote the true transition kernel of the MDP associated with requests from user m at AP n , i.e., $P_{mn}(s'|s, a), \forall s, a$ as $\theta_{mn}^*, \forall m \in \mathcal{M}, n \in \mathcal{N}$, which is unknown to the controller. Let $h_{mn}^t = (S_{mn}^1, A_{mn}^1, S_{mn}^2, A_{mn}^2, \dots, S_{mn}^t, A_{mn}^t)$ be the history of states and actions up to time slot $t, \forall m \in \mathcal{M}, n \in \mathcal{N}$. We focus on a Bayesian framework, and denote μ_{mn}^1 as the prior distribution for $\theta_{mn}^*, \forall m \in \mathcal{M}, n \in \mathcal{N}$, i.e., $\mathbb{P}(\theta_{mn}^* \in \Theta) = \mu_{mn}^1(\Theta)$ for any arbitrary set Θ . mmDPT-TS operates in episodes and decomposes the total operating time T into K episodes (the value of K will be specified later). Let t_k be the start time of episode k and $T_k = t_{k+1} - t_k$ be the length of the episode, satisfying $T = \sum_k T_k$. W.l.o.g., we set $T_0 = 1$. Each episode consists of two phases: posterior updates and policy execution.

Posterior Updates. At time slot t , the posterior distribution $\mu_{mn}^t, \forall m, n$ can be computed based on the history h_{mn}^t , i.e., $\mu_{mn}^t(\Theta) = \mathbb{P}(\theta_{mn}^* \in \Theta | h_{mn}^t)$, for any set Θ . After applying action A_{mn}^t and observing the next state S_{mn}^{t+1} , the posterior distribution μ_{mn}^{t+1} can be updated using Bayes' rule as:

$$\mu_{mn}^{t+1}(d\theta) = \frac{\theta_{mn}(S_{mn}^{t+1} | S_{mn}^t, A_{mn}^t) \mu_{mn}^t(d\theta)}{\int \tilde{\theta}_{mn}(S_{mn}^{t+1} | S_{mn}^t, A_{mn}^t) \mu_{mn}^t(d\tilde{\theta})}, \forall m, n. \quad (10)$$

Algorithm 2 mmDPT-TS

Input: Prior distribution $\mu_{mn}^1, \forall m, n$;

- 1: Initialize $C_{mn}^1(s, a) = 0, \forall m, n, s, a; t = 1, T_0 = 0, t_1 = 1; \pi^{*,1}$ to be any policy;
- 2: **for** episodes $k = 1, 2, \dots$ **do**
- 3: **while** $t \leq t_k + T_{k-1}$ and $C_{mn}^t(s, a) \leq 2C_{mn}^{t_k}(s, a), \forall m, n, s, a$ **do**
- 4: Execute policy $\pi^{*,k}$ and observe new state S_{mn}^{t+1} ;
- 5: Update μ_{mn}^{t+1} according to (10);
- 6: $T_k \leftarrow t - t_k, t \leftarrow t + 1$;
- 7: **end while**
- 8: $t_{k+1} \leftarrow t$;
- 9: Sample θ_{mn}^{k+1} from $\mu_{mn}^{t_{k+1}}$, compute $\omega_{mn}^{*,k+1}$ via (11);
- 10: Construct the mmDPT Index Policy $\pi^{*,k+1}$ according to Algorithm 1 using $\omega_{mn}^{*,k+1}$.
- 11: **end for**

To execute the constructed mmDPT Index Policy, we need to determine when episode k terminates. Let $C_{mn}^t(s, a)$ be the number of visits to state-action pairs (s, a) until t for the MDP associated with requests from user m at AP n , satisfying $C_{mn}^t(s, a) = C_{mn}^{t-1}(s, a) + \mathbb{1}(S_{mn}^t = s, A_{mn}^t = a), \forall s, a, m, n$. Inspired by [53], episode k ends if its length is no less than that of episode $k - 1$, or the number of visits to some state-action pairs (s, a) satisfies $C_{mn}^{t_{k+1}}(s, a) > 2C_{mn}^{t_k}(s, a), \forall m, n$. Thus, $t_0 = 1$ and $t_{k+1}, k \geq 1$ is given by $t_{k+1} = \min\{t > t_k : t > t_k + T_{k-1} \text{ or } C_{mn}^{t_{k+1}}(s, a) > 2C_{mn}^{t_k}(s, a), \forall m, n, s, a\}$.

Policy Execution. At the policy execution phase of each episode, mmDPT-TS constructs and executes mmDPT Index Policy. This is the key contribution and novelty of our proposed structured RL algorithm mmDPT-TS, which leverages our proposed near-optimal mmDPT Index Policy for making decisions, instead of contending directly with an extremely large state-action space (e.g., via solving complicated Bellman equations). These together contribute to the sub-linear Bayesian regret of mmDPT-TS with a low computational complexity, which will be discussed in detail later. Specifically, at the beginning of episode k , the parameters $\{\theta_{mn}^k, \forall m, n\}$ are sampled from the posterior distributions $\{\mu_{mn}^{t_k}, \forall m, n\}$. Using these samples, mmDPT-TS solves the following LP:

$$\begin{aligned} \min_{\{\omega_{mn}\}} & \sum_{n=1}^N \sum_{m=1}^M \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \omega_{mn}^k(s, a) s \\ \text{s.t.} & \sum_{s' \in \mathcal{S}} \sum_{a \in \mathcal{A}} \omega_{mn}^k(s, a) \theta_{mn}^k(s' | s, a) \\ & = \sum_{s' \in \mathcal{S}} \sum_{a \in \mathcal{A}} \omega_{mn}^k(s', a) \theta_{mn}^k(s | s', a), \forall s \in \mathcal{S}, \\ & \text{Constraints (7b), (7c), and (7e)}. \end{aligned} \quad (11)$$

We denote the optimal solution to the above LP (11) as $\{\omega_{mn}^{*,k}(s, a), \forall m, n, s, a\}$, using which mmDPT-TS computes the mmDPT Index in (9), and then constructs the mmDPT Index Policy according to Algorithm 1. We denote the policy as $\pi^{*,k}$, and then execute this policy in this episode. We summarize this process in Algorithm 2.

Remark 2. *mmDPT-TS leverages the low-complexity provably optimal mmDPT Index Policy for making decisions, and hence only needs to solve a LP (11) at each episode (in polynomial time [48]–[50]). This differentiates mmDPT-TS from state of the arts, which are often computationally expensive. For example, [53] proposed a TS method for MDPs and the optimal policy is approximated via solving complicated Bellman equations. [21] extended [53] to RMAB, however, the computation of Whittle index policy also relies on repeatedly solving Bellman equations. Another line of deep RL based approaches, e.g., [29]–[32] neither incorporate fairness constraint, nor have finite-time performance analysis.*

B. The Learning Regret

We characterize the finite-time performance of mmDPT-TS using the Bayesian regret. Specifically, the Bayesian regret of a learning policy π is defined as

$$R(T) = \mathbb{E}_{\pi, \mu^1} \left[\sum_{t=1}^T \sum_{n=1}^N \sum_{m=1}^M S_{mn}^t - T J(\theta^*) \right], \quad (12)$$

where $J(\theta^*)$ is performance of mmDPT Index Policy under the perfect knowledge of the true transition kernel $\theta^* := \{\theta_{mn}^*, \forall m, n\}$; and the expectation is taken with respect to prior distributions $\mu^1 := \{\mu_{mn}^1, \forall m, n\}$ and policy π . We follow RMAB literature, e.g., [21], [25] to define the Bayesian regret with respect to the mmDPT Index Policy, which is asymptotically optimal.

Assumption 1. *Let $J(\theta^k)$ be the average cost for mmDPT Index Policy under $\theta^k := \{\theta_{mn}^k, \forall m, n\}$. For $\forall \theta^k$, $J(\theta^k)$ does not depend on the initial state and satisfies the average cost Bellman equation $\forall t \in [t_k, t_{k+1}]$:*

$$C(\mathbf{S}^t, \mathbf{A}^t) = J(\theta^k) + V_{\theta^k}(\mathbf{S}^t) - \sum_{\mathbf{S}'} \theta^k(\mathbf{S}' | \mathbf{S}^t, \mathbf{A}^t) V_{\theta^k}(\mathbf{S}'), \quad (13)$$

where $V_{\theta^k}(\mathbf{S}^t)$ is the bias value function [38] and unique up to a constant.

Assumption 1 is standard in TS-based methods [21], [53], [54], which ensures that the average cost of mmDPT Index Policy is well defined. The span of the bias value function V under transition kernel θ^k is defined as [54]:

$$\text{Span}(V_{\theta^k}) := \max_{\mathbf{S} \in S^{MN \times 1}} V_{\theta^k}(\mathbf{S}) - \min_{\mathbf{S} \in S^{MN \times 1}} V_{\theta^k}(\mathbf{S}), \quad (14)$$

which is an essential factor for bounding the Bayesian regret of mmDPT-TS. Under Assumption 1 and the span definition, our main result in this section is stated as follows:

Theorem 2. *The Bayesian regret of mmDPT-TS satisfies*

$$R(T) = \mathcal{O}(S_{max}^3 M^2 N^2 \sqrt{T \log T}). \quad (15)$$

The high level idea of the proof is similar to [21], [53], but we provide an explicit upper bound on the span with respect to the dimension of state space S_{max} , the number of APs N , and the number of users M , by leveraging the structure encoded in our RMAB-F. This is one of main contributions in this work. For ease of readability, we present a proof outline below, and relegate the details to Appendix VII.

C. Proof Sketch of Theorem 2

Regret Decomposition. Let K_T be the number of episodes until time horizon T . Given the average cost Bellman equation (13), the Bayesian regret (12) can be decomposed as

$$\begin{aligned} R(T) &= \underbrace{\mathbb{E} \left[\sum_{k=1}^{K_T} T_k J(\theta^k) \right] - T \mathbb{E}[J(\theta^*)]}_{R_1: \text{regret due to Bayesian estimating error}} \\ &\quad + \underbrace{\mathbb{E} \left[\sum_{k=1}^{K_T} \sum_{t=t_k}^{t_{k+1}-1} [V_{\theta^k}(\mathbf{S}^t) - V_{\theta^k}(\mathbf{S}^{t+1})] \right]}_{R_2: \text{regret due to time-varying policies between episodes}} \\ &\quad + \underbrace{\mathbb{E} \left[\sum_{k=1}^{K_T} \sum_{t=t_k}^{t_{k+1}-1} [V_{\theta^k}(\mathbf{S}^{t+1}) - \sum_{\mathbf{S}'} \theta^k(\mathbf{S}' | \mathbf{S}^t, \mathbf{A}^t) V_{\theta^k}(\mathbf{S}')] \right]}_{R_3: \text{regret due to model mismatch}}. \end{aligned}$$

We then proceed to derive bounds on K_T and R_1, R_2, R_3 .

Bounding K_T : Since we consider dynamic episodes as inspired by [53], the number of episodes K_T can be upper bounded by $2\sqrt{S_{max} MNT \log T}$.

Bounding R_1 : Following the monotone convergence theorem, $R_1 \leq \sum_{k=1}^{\infty} \mathbb{E}[\mathbb{1}_{t_k \leq T} (T_{k-1} + 1) J(\theta^k)] - T \mathbb{E}[J(\theta^*)] = \mathbb{E}[\sum_{k=1}^{K_T} (T_{k-1} + 1) J(\theta^k)] - T \mathbb{E}[J(\theta^*)] \leq \mathbb{E}[K_T]$.

Bounding R_2 : Similar to [53], R_2 can be upper bounded by the value of span and K_T as $\mathbb{E}[\text{Span}(V) K_T]$.

Bounding R_3 : R_3 is the regret part due to *model mismatch*, which is one key contribution of our proof compared to existing results [21], [53]. We first construct a confidence ball in each episode which characterizes the distance between the true transition kernel and the sampled transition kernel. We show that with high probability the true transition kernel lies in the confidence ball and the complementary event is a rare event. Bounding these two regrets leads to the bound on $R_3 \leq \text{Span}(V) S_{max} MN \sqrt{T \log T}$.

Bounding $\text{Span}(V)$: The value of span plays a critical role in the Bayesian regret analysis, which is characterized in Lemma 1. This is another key contribution in this work.

Lemma 1. *The span of V under $\theta^k, \forall k$ is upper bounded by*

$$\text{Span}(V) := \text{Span}(V_{\theta^k}) \leq (S_{max}^2 + S_{max}) MN / 2.$$

Remark 3. *mmDPT-TS achieves a sub-linear Bayesian regret $\tilde{\mathcal{O}}(\sqrt{T \log T})$ as state-of-the-art TS-based methods [53] for MDPs and [21] for RMAB. Different from them, we provide an explicit upper bound on the span of the bias value function in Lemma 1. In contrast, the span is assumed to be upper bounded by a constant in [53], and the bound in [21] relies on an “ergodicity coefficient”, which is a unknown parameter and varies across different MDP realizations. This assumption is not necessary in our analysis since we leverage the underlying structure in RMAB-F to bound the span.*

V. EXPERIMENTS

In this section, we numerically evaluate the performance of our proposed mmDPT Index Policy and mmDPT-TS using both real traces collected from a 60GHz mmWave testbed and synthetic traces.


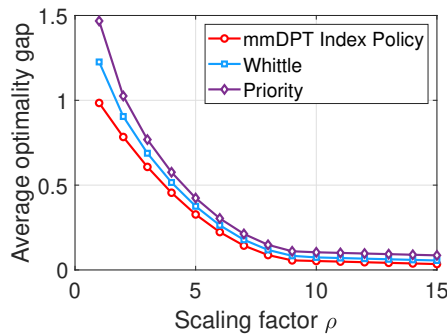



Fig. 4: Asymptotic optimality: 60GHz mmWave testbed.

Fig. 5: Asymptotic optimality: Synthetic data traces.

Fig. 6: Global attractor: Synthetic data traces.

User	AP			
	1	2	3	4
1	0.09,0.07,0.03,0.01	0.095,0.075,0.025,0.005	0.08,0.06,0.04,0.02	0.085,0.065,0.045,0.025
21	0.08,0.07,0.06,0.05	0.085,0.075,0.065,0.055	0.07,0.06,0.05,0.04	0.075,0.065,0.055,0.045
41	0.07,0.06,0.05,0.04	0.075,0.065,0.055,0.045	0.06,0.05,0.04,0.03	0.065,0.055,0.045,0.035
61	0.06,0.05,0.04,0.03	0.065,0.055,0.045,0.035	0.05,0.04,0.03,0.02	0.055,0.045,0.035,0.025
81	0.05,0.04,0.03,0.02	0.055,0.045,0.035,0.025	0.04,0.03,0.02,0.01	0.045,0.035,0.025,0.015

TABLE I: The probability of successfully delivering 1, 2, 3, 4 packets over frames in synthetic traces for some users.

A. Evaluation Setup

60GHz mmWave Testbed. Since commodity off-the-shelf (COTS) 802.11ad devices can generate the desirable data traces for our simulations, we build a 60GHz mmWave communication testbed using software-defined radio (SDR) devices. Specifically, the testbed is emulation based and measured repeatedly through one transmitter and one receiver, both of which are built using one computer, one ADI EVAL-HMC6300 Board (60GHz RF Frontend), one USRP X310, and one planar antenna. The planar antenna has 4×8 patch elements for beam steering to compensate the high path loss of mmWave signal propagation. We implement a simplified version of IEEE 802.11ad protocol [55] on the testbed for data packet transmission. The instantaneous bandwidth of signal transmission is 100MHz. The FFT size of OFDM symbols is 512, and the modulation scheme is QPSK. All signal processing modules are implemented in the computer using C++. Figure 2(c) shows a snapshot of received/decoded signal constellation at the receiver using the testbed in Figure 2(a).

We consider a dense, cell-free mmWave network in a conference room, which consists of 4 APs and 30 user devices as shown in Figure 2(a) with an example snapshot in Figure 1. The APs are placed on two-side walls, while 30 user devices are uniformly distributed over the whole conference room. Using this mmWave testbed, we conduct real-time data packet transmissions from each AP to each user devices to collect data traces for our simulation. We measure the error vector magnitude (EVM) of the decoded signal constellations at the user device (receiver) for every packet (0.128ms). A total of 468,750 EVM samples are recorded over 60 seconds for each AP-user pair. Figure 3 shows three instances of EVM traces, where we draw 1,000 samples for each instance out of the total samples with a step-size of 468 for ease of illustration. To the end, we collect 4×30 EVM traces for those 4 APs and 30 user devices. The measured EVM samples

are used to infer their corresponding packet error rate (PER) based on the selected modulation and coding scheme (MCS) and other parameters specified in the 802.11ad standard [55]. Specifically, the simulation time is divided into 60 frames, each of which consists of EVM samples recorded in 1 second. For simplicity, we use EVMs in each framework to infer PER and further the transition probabilities in Section II based on [56]. These inferred values over time are used as the input of our simulation to evaluate our proposed algorithms.

Synthetic Traces. We simulate a dense mmWave network with 4 APs and 100 users. The request arrival probability is drawn from a Poisson process. The mass function is defined as $f(1000; 2000)$, and normalized with an average of 0.5. To model the dynamic nature of mmWave channels, we divide the simulation time into frames, each of which consists of 10^4 time slots, and the distribution over the number of packets successfully delivered from an AP to one user (defined in (1)) is fixed in one frame but varies across frames. We assume that at most $d = 4$ packets can be transmitted and the corresponding distribution of selected user 1, 21, 41, 61, 81 is presented in Table I. The distribution of user with index number between them are arithmetic sequences. For example, (0.09, 0.07, 0.03, 0.01) corresponds to the probabilities of (1, 2, 3, 4) packets transmitted from AP 1 to user 1 in frame 1, respectively. Hence the remaining 0.8 probability corresponds to no packet delivery. We set the maximum queue size as $S_{\max} = 15$ and the fairness constraint as $B = 20$.

Baselines. We compare our mmDPT Index Policy with (a) Whittle index based (Whittle) [13]; and (b) priority index based (Priority) [29]. We note that none of these policies can be directly applied to Problem mmDPT (4) since their problems were cast as a RMAB without the fairness constraint (3). To this end, we augment them with the sorting step as in the design of mmDPT Index Policy (see Section III-B), and refer to the resulting algorithms as Whittle and Priority, respectively.

Correspondingly, when the system dynamics are unknown, we compare our mmDPT-TS with (a) a TS method [21] to learn the above Whittle policy (TS-Whittle); (b) the above priority index enabled learning policy (IDEA) for packet scheduling in mmWave networks [29]; (c) Deep Q-network (DQN) based packet scheduling policy [32]; and (d) soft actor-critic (SAC) based scheduling policy [31]. Again, the design of these (deep) RL based scheduling policies did not incorporate


the fairness constraint (3). For sake of fair comparison, we augment them in the same manner as aforementioned, and the prior μ_{mn}^t is sampled from a Dirichlet distribution for mmDPT-TS.

B. Evaluation Results

Asymptotic Optimality. We first validate the asymptotic optimality of mmDPT Index Policy (see Theorem 1). We compare the cumulative cost (measured in average delay) suffered by all users under different policies, with that obtained from the theoretical lower bound obtained via solving the LP (7a)-(7e). We call this difference *the optimality gap*. The average optimality gap, which is the ratio of the optimality gap and the scaling parameter ρ (see Section III-C), is presented in Figure 4 using the real traces collected from our mmWave testbed and Figure 5 using synthetic traces. We observe that the average optimality gap decreases significantly and closes to zero as ρ increases. This verifies the asymptotic optimality in Theorem 1. An interesting observation is that though Whittle and Priority scheduling policies did not incorporate the fairness constraint, when augmented with our proposed sorting step as aforementioned, their asymptotic performance can be guaranteed. This further validates the independent interest of our proposed framework for designing index policies.

Global Attractor. The asymptotic optimality of mmDPT Index Policy is under the definition of global attractor. For ease of illustration, we randomly pick three state-action pairs: (5, 1) for requests from user 1 at AP 3, (5, 0) for requests from user 20 at AP 2, and (13, 0) for requests from user 70 at AP 4, all in frame 1 with 10^4 time slots. As shown in Figures 6, the occupancy measure of requests from user 1 at AP 3 for state-action pair (5, 1) indeed converges using synthetic traces. Similar observations can be made for the other two cases. Therefore, the convergence indeed occurs for mmDPT Index Policy and hence we verify the global attractor condition.


Learning Regret and Running Time. The learning regret of mmDPT-TS and other baselines under real and synthetic traces in any particular frame are shown in Figure 7(a) and (b), respectively, where we use the Monte Carlo simulation with 2,000 independent trails of a single-threaded program on Ryzen 7 7800X3D desktop with 32 GB RAM. We observe that mmDPT-TS consistently achieves a much smaller regret compared to other baselines. The corresponding running time is shown in Figure 8(a) and (b), respectively, where the error bars are drawn based on the standard deviation. Note that although TS-Whittle is also an index-aware TS based method, there is often no explicit expression for its intrinsic index policy, i.e., the Whittle index policy, which is often computed through numeral methods [13], [57]. In particular, we use value iteration to compute the Whittle index for TS-Whittle in our experiments. We observe that the running time of mmDPT-TS is similar to that of SAC and outperforms all others. However, SAC has a much larger regret than mmDPT-TS as shown in Figure 7. These observations are consistent with our motivation that existing learning polices either do not incorporate fairness constraint and hence cannot be directly applied to mmDPT, or do not have a finite-time



(a) 60GHz mmWave testbed.

(b) Synthetic data traces.

Fig. 7: Accumulated regret.



(a) 60GHz mmWave testbed.

(b) Synthetic data traces.

Fig. 8: Average running time.

(regret) performance guarantee, or are computationally expensive, while our mmDPT-TS achieves all at once.

VI. CONCLUSION

We studied the data packet transmission problem (mmDPT) in a dense, cell-free mmWave network to minimize the average delay experienced by all users in the system. We proposed a low-complexity structured RL solution mmDPT-TS for mmDPT by exploiting the inherent problem structure. We proved that mmDPT-TS achieved a sub-linear Bayesian regret. Experimental results based on the data collected from realistic mmWave networks corroborate our theoretical analysis.

VII. APPENDIX

A. Proof of Theorem 1

Since $\lim_{\rho \rightarrow \infty} V_{\pi^*}^\rho - V_{\pi_{opt}}^\rho$ is non-negative, $\forall \pi^*$ from Algorithm 1, we only need to show that it is non-positive. Let $X_{mn}^\rho(\pi^*, s, a)$ be the average number of class- m users with state at class- n APs being s with action a taken under π^* , which is the global attractor based on Definition 1. The key is then to show [39]

$$\lim_{\rho \rightarrow \infty} X_{mn}^\rho(\pi^*, s, a) / \rho = \omega_{mn}(s, a), \forall n, m.$$

We denote $A_{mn}^{\pi^*}(s)$ as the set of all combinations (m', j) , $m' \in \mathcal{N}$, $j \in \mathcal{S}$ such that class- m' users with state at class- n APs being j have larger indices than those of class- m users with state at class- n APs being s under the mmDPT Index index policy π^* . The transition rates of the process $X_{mn}^\rho(\pi^*, t) / \rho$ are then defined as

$$x \rightarrow x - \frac{e_{mn,s}}{\rho} + \frac{e_{mn,s'}}{\rho} \quad (16)$$

at rate $\sum_a P_{mn}(s'|s, a)x_{mn}^\rho(s, a)$, where $x_{mn}^\rho(s, 1) = \min\left(\rho B - \sum_{(m', j) \in A_{mn}^*(s)} x_{m'n}^\rho(j, 1), 0\right)$ and $e_{mn, s} \in \mathbb{R}^{S \times 1}$ is unit vector with the s -th position being 1. It follows from [58] that there exists a continuous function $f_\ell(x)$ to model the transition rate of the process $X_{mn}^\rho(\pi^*; t)$ from state x to $x + \ell/\rho, \forall \ell \in \mathcal{L}$ according to (16), with \mathcal{L} being the set composed of a finite number of vectors in \mathbb{N}^{SN} . Hence, the process $X_{mn}^\rho(\pi^*; t)/\rho$ is a density dependent population processes as in [39], [58]. Note that the process $X_{mn}^\rho(\pi^*; t)$ can be expressed as

$$dX_{mn}^\rho(\pi^*; t)/dt = F(X_{mn}^\rho(\pi^*; t)),$$

with $F(\cdot)$ being Lipschitz continuous and satisfying $F(X_{mn}^\rho(\pi^*; t)) = \sum_{\ell \in \mathcal{L}} \ell f_\ell(X_{mn}^\rho(\pi^*; t))$. Under the condition that the considered MDP is unichain, the process $\frac{X_{mn}^\rho(\pi^*; t)}{\rho}$ has a unique invariant probability distribution $\zeta_{\pi^*}^\rho$, which is tight [39]. Thus, we have $\zeta_{\pi^*}^\rho \left(\frac{X_{mn}^\rho(\pi^*; t)}{\rho}\right)$ converge to the Dirac measure in $X_{mn}^{\rho, *}/\rho$ when $\rho \rightarrow \infty$, which is a *global attractor* of $\frac{X_{mn}^\rho(\pi^*; t)}{\rho}$, i.e., $\omega_{mn}(s, a), \forall s \in \mathcal{S}, a \in \mathcal{A}$. Combing these together, we have

$$\begin{aligned} \lim_{\rho \rightarrow \infty} V_{\pi^*}^\rho &\stackrel{(a)}{=} \lim_{\rho \rightarrow \infty} \sum_n \sum_m \sum_{(s, a)} s D_{mn}^{\rho, *}(s, a) / \rho \\ &\stackrel{(b)}{=} \sum_n \sum_m \sum_{(s, a)} s \omega_{mn}^*(s, a) \\ &\stackrel{(c)}{\leq} \lim_{\rho \rightarrow \infty} V_{\pi^{opt}}^\rho, \end{aligned}$$

where (a) is from the definition of $D_{mn}^{\rho, *}(s, a)$, (b) holds since $\lim_{\rho \rightarrow \infty} D_{mn}^{\rho, *}(s, a) / \rho = \omega_{mn}(s, a)$, and (c) is because the optimal value of (7a)-(7e) is a lower bound of that of (4).

B. Proof of Lemma 1.

Since $V_{\theta^k}, \forall k$ is unique up to a constant δ , hence $V_{\theta^k} + \delta$ also satisfies (13) [38]. Define the stationary distribution over \mathbf{S} under transition kernel θ^k with mmDPT Index Policy as η_{θ^k} . W.l.o.g, we assume that V_{θ^k} satisfies $\eta_{\theta^k}^\top V_{\theta^k} = J(\theta^k) - S_{max}MN$. Following [38], V_{θ^k} is computed as the asymptotic bias of policy $\pi^{*, k}$ under θ^k as

$$V_{\theta^k}(\mathbf{S}) = \sum_{t=0}^{\infty} \sum_{\mathbf{S}'} \{\theta^k(\mathbf{S}^{t+1} = \mathbf{S}' | \mathbf{S}^t, \mathbf{A}^t) [C(\mathbf{S}^{t+1}, \mathbf{A}^{t+1}) - S_{max}MN] | \mathbf{S}^0 = \mathbf{S}\}.$$

Since at each time slot, we have that

$$\sum_{\mathbf{S}'} \theta^k(\mathbf{S}^{t+1} = \mathbf{S}' | \mathbf{S}^t, \mathbf{A}^t) [C(\mathbf{S}^{t+1}, \mathbf{A}^{t+1}) - S_{max}MN] \leq 0,$$

it leads to $V_{\theta^k}(\mathbf{S}) \leq 0, \forall \mathbf{S}$. The equality holds when the current state is $\mathbf{S} = S_{max}[1, 1, \dots]^{1 \times MN}$ and the stationary distribution is $\eta_{\theta^k}(\mathbf{S}) = 1$ when $\mathbf{S} = S_{max}[1, 1, \dots]^{1 \times MN}$ and 0 otherwise. By leveraging $\eta_{\theta^k}^\top V_{\theta^k} = J(\theta^k) - S_{max}MN$, we still consider the MDP under θ^k with $J(\theta^k) := S_{max}MN$, which implies that $\eta_{\theta^k}(\mathbf{S}) = 1$ when $\mathbf{S} = S_{max}[1, 1, \dots]^{1 \times MN}$ and 0 otherwise. This leads to the fact that $V_{\theta^k}(\mathbf{S}) \geq -(S_{max}^2 + S_{max})MN/2, \forall \mathbf{S}$. Hence, $\text{Span}(V_{\theta^k}) = \max_{\mathbf{S}} V_{\theta^k}(\mathbf{S}) - \min_{\mathbf{S}} V_{\theta^k}(\mathbf{S}) \leq (S_{max}^2 + S_{max})MN/2$.

C. Proof of Theorem 2

We will first decompose the regret into three terms, corresponding to sampling error, time-varying policy and model mismatch. Define $K_T = \text{argmax}\{k : t_k \leq T\}$ be the number of episodes of mmDPT-TS until time T. Recall t_t as the start time of episode k , for $t_k \leq t < t_{k+1}$, the Bellman equation (Assumption 1) holds:

$$C(\mathbf{S}^t, \mathbf{A}^t) = J(\theta^k) + V_{\theta^k}(\mathbf{S}^t) - \sum_{\mathbf{S}'} \theta^k(\mathbf{S}' | \mathbf{S}^t, \mathbf{A}^t) V_{\theta^k}(\mathbf{S}').$$

We use $\mathbf{S}^t, \mathbf{A}^t$ to represent the state and action matrix of time t and also simplify $\theta_{m, n}^k$ to θ^k and V is the value function. By rearranging the Bellman equation we can have:

$$\begin{aligned} R(T, \pi) &= \mathbb{E} \left[\sum_{t=1}^T [C(\mathbf{S}^t, \mathbf{A}^t) - J(\theta^*)] \right] \\ &= \mathbb{E} \left[\sum_{k=1}^{K_T} \sum_{t=t_k}^{t_{k+1}-1} C(\mathbf{S}^t, \mathbf{A}^t) \right] - T\mathbb{E}[J(\theta^*)] \\ &= \mathbb{E} \left[\sum_{k=1}^{K_T} T_k J(\theta^k) \right] - T\mathbb{E}[J(\theta^*)] \\ &\quad + \mathbb{E} \left[\sum_{k=1}^{K_T} \sum_{t=t_k}^{t_{k+1}-1} [V_{\theta^k}(\mathbf{S}^t) - \sum_{\mathbf{S}'} \theta^k(\mathbf{S}' | \mathbf{S}^t, \mathbf{A}^t) V_{\theta^k}(\mathbf{S}')] \right] \\ &= \underbrace{\mathbb{E} \left[\sum_{k=1}^{K_T} T_k J(\theta^k) \right] - T\mathbb{E}[J(\theta^*)]}_{R_1} \\ &\quad + \underbrace{\mathbb{E} \left[\sum_{k=1}^{K_T} \sum_{t=t_k}^{t_{k+1}-1} [V_{\theta^k}(\mathbf{S}^t) - V_{\theta^k}(\mathbf{S}^{t+1})] \right]}_{R_2} \\ &\quad + \underbrace{\mathbb{E} \left[\sum_{k=1}^{K_T} \sum_{t=t_k}^{t_{k+1}-1} [V_{\theta^k}(\mathbf{S}^{t+1}) - \sum_{\mathbf{S}'} \theta^k(\mathbf{S}' | \mathbf{S}^t, \mathbf{A}^t) V_{\theta^k}(\mathbf{S}')] \right]}_{R_3}, \end{aligned}$$

where R_1, R_2 and R_3 corresponds to sampling error, time-varying policy and model mismatch, respectively. We will first bound the number of episodes (K_T) then analyze them individually.

1) *Bound on number of episodes K_T* : Note that K_T is a random variable because the number of visits $C_{m, n}^t(s, a)$ depends on the dynamical state trajectory. We provide an upper bound on K_T as follows.

Lemma 2.

$$K_T \leq 2\sqrt{S_{max}MNT \log T}.$$

Proof. Define macro episodes with start times $t_{n_i}, i = 1, 2, \dots$ where $t_{n_1} = t_1$ and

$$t_{n_{i+1}} = \min\{t_k > t_{n_i} : C_{m, n}^{t_k}(s, a) > 2C_{m, n}^{t_{k-1}}(s, a) \text{ for any } (m, n, s, a)\}.$$

This condition is related to the second stopping criterion. Let γ be the number of macro episodes until time T and define $n_{\gamma+1} = K_T + 1$.

Let $\tilde{T}_i = \sum_{k=n_i}^{n_{i+1}-1} T_k$ be the length of the i -th macro episode. By the definition of macro episodes, any episode except the last one in a macro episode must be triggered by the first stopping criterion. Therefore, within the i -th macro episode, $T_k = T_{k-1} + 1$ for all $k = n_i, n_i + 1, \dots, n_{i+1} - 2$. Hence,

$$\begin{aligned} \tilde{T}_i &= \sum_{k=n_i}^{n_{i+1}-1} T_k = \sum_{j=1}^{n_{i+1}-n_i-1} (T_{n_{i-1}+j} + T_{n_{i+1}-1}) \\ &\geq \sum_{j=1}^{n_{i+1}-n_i-1} (j+1) + 1 = \frac{(n_{i+1} - n_i)(n_{i+1} - n_i + 1)}{2}. \end{aligned}$$

We then obtain

$$K_T = n_{\gamma+1} - 1 \leq \sum_{i=1}^{\gamma} \sqrt{2\tilde{T}_i},$$

with the fact that $\sum_{i=1}^{\gamma} \tilde{T}_i = T$, we have

$$K_T \leq \sum_{i=1}^{\gamma} \sqrt{2\tilde{T}_i} \leq \sqrt{\gamma \sum_{i=1}^{\gamma} 2\tilde{T}_i} = \sqrt{2\gamma T}, \quad (17)$$

where the second inequality is Cauchy-Schwarz. Next we will bound the number of episodes of γ . The start of micro episodes can be expressed as:

$$\{t_{n_1}\} \cup \left(\bigcup_{(s,a) \in \mathbf{S} \times \mathbf{A}, m \in \mathcal{M}, n \in \mathcal{N}} \{t_k : k \in \Gamma_{m,n}^{(s,a)}\} \right),$$

where $\Gamma_{m,n}^{(s,a)} = \{k \leq K_T : C_{m,n}^{t_k}(s,a) > 2C_{m,n}^{t_k-1}(s,a)\}$. The size of $\Gamma_{m,n}^{(s,a)}$ satisfies $|\Gamma_{m,n}^{(s,a)}| \leq \log C_{m,n}^{T+1}(s,a)$. This can be proved by contradiction. Define $t_{K_T+1} = T + 1$, if $|\Gamma_{m,n}^{(s,a)}| \geq \log C_{m,n}^{T+1}(s,a) + 1$,

$$\begin{aligned} C_{m,n}^{t_{K_T}}(s,a) &= \prod_{k \leq K_T, C_{m,n}^{t_{k-1}}(s,a) \geq 1} \frac{C_{m,n}^{t_k}(s,a)}{C_{m,n}^{t_{k-1}}(s,a)} \\ &> \prod_{k \in \Gamma_{m,n}^{(s,a)}, C_{m,n}^{t_{k-1}}(s,a) \geq 1} 2 \geq C_{m,n}^{T+1}(s,a), \end{aligned}$$

which contradicts with $C_{m,n}^{t_{K_T}}(s,a) \leq C_{m,n}^{T+1}(s,a)$. Therefore $|\Gamma_{m,n}^{(s,a)}| \leq \log C_{m,n}^{T+1}(s,a)$. We can bound γ with:

$$\begin{aligned} \gamma &\leq \sum_{M \in \mathcal{M}, N \in \mathcal{N}} \sum_{s \in \mathcal{S}, a \in \mathcal{A}} |\Gamma_{m,n}^{(s,a)}| \\ &\leq \sum_{M \in \mathcal{M}, N \in \mathcal{N}} \sum_{s \in \mathcal{S}, a \in \mathcal{A}} \log C_{m,n}^{T+1}(s,a) \\ &\leq 2S_{max}MN \log \sum_{s \in \mathcal{S}, a \in \mathcal{A}} \frac{C_{m,n}^{T+1}(s,a)}{2S_{max}} \\ &\leq 2S_{max}MN \log \sum_{s \in \mathcal{S}, a \in \mathcal{A}} C_{m,n}^{T+1}(s,a) \\ &= 2S_{max}MN \log T. \end{aligned} \quad (18)$$

Combine 18 and 17 we complete the proof of Lemma 2. \square

2) *Bound on R_1* : One key property of Thomson Sampling is $\mathbb{E}[f(\boldsymbol{\theta}^k, X)] = \mathbb{E}[f(\boldsymbol{\theta}^*, X)]$, but this is different in dynamic episode version, we provide the following lemma.

Lemma 3. *Under mmDPT-TS, t_k is a stopping time for any episode k . Then for any measurable function f and any $\delta(h_{t_k})$ -measurable random variable X , we have*

$$\mathbb{E}[f(\boldsymbol{\theta}^k, X)] = \mathbb{E}[f(\boldsymbol{\theta}^*, X)].$$

Proof. The only randomness in $f(\boldsymbol{\theta}^k, X)$ is the random sampling in the algorithm, which gives the following equation:

$$\begin{aligned} \mathbb{E}[f(\boldsymbol{\theta}^k, X)|h_{t_k}] &= \mathbb{E}[f(\boldsymbol{\theta}^k, X)|h_{t_k}, t_k, \mu_k] \\ &= \mathbb{E}[f(\boldsymbol{\theta}^*, X)|h_{t_k}], \end{aligned}$$

\square

The result follows by taking the expectation for both sides. From monotone convergence theorem we have:

$$\begin{aligned} R_1 &= \mathbb{E} \left[\sum_{k=1}^{K_T} T_k J(\boldsymbol{\theta}^k) \right] - T \mathbb{E}[J(\boldsymbol{\theta}^*)] \\ &= \mathbb{E} \left[\sum_{k=1}^{\infty} \mathbf{1}_{t_k \leq T} T_k J(\boldsymbol{\theta}^k) \right] - T \mathbb{E}[J(\boldsymbol{\theta}^*)] \\ &= \sum_{k=1}^{\infty} \mathbb{E} \left[\mathbf{1}_{t_k \leq T} T_k J(\boldsymbol{\theta}^k) \right] - T \mathbb{E}[J(\boldsymbol{\theta}^*)] \\ &\leq \sum_{k=1}^{\infty} \mathbb{E} \left[\mathbf{1}_{t_k \leq T} (T_{k-1} + 1) J(\boldsymbol{\theta}^k) \right] - T \mathbb{E}[J(\boldsymbol{\theta}^*)]. \end{aligned}$$

From Lemma 3, we have

$$\mathbb{E} \left[\mathbf{1}_{t_k \leq T} T_k J(\boldsymbol{\theta}^k) \right] = \mathbb{E} \left[\mathbf{1}_{t_k \leq T} T_k J(\boldsymbol{\theta}^*) \right],$$

therefore we obtain

$$\begin{aligned} R_1 &\leq \sum_{k=1}^{\infty} \mathbb{E} \left[\mathbf{1}_{t_k \leq T} (T_{k-1} + 1) J(\boldsymbol{\theta}^k) \right] - T \mathbb{E}[J(\boldsymbol{\theta}^*)] \\ &= \mathbb{E} \left[\sum_{k=1}^{K_T} T_{k-1} J(\boldsymbol{\theta}^*) \right] - T \mathbb{E}[J(\boldsymbol{\theta}^*)] \\ &\leq E[K_T]. \end{aligned}$$

3) *Bound on R_2* : R_2 can be simplified as follows:

$$\begin{aligned} R_2 &= \mathbb{E} \left[\sum_{k=1}^{K_T} \sum_{t=t_k}^{t_{k+1}-1} [V_{\boldsymbol{\theta}^k}(\mathbf{S}^t) - V_{\boldsymbol{\theta}^k}(\mathbf{S}^{t+1})] \right] \\ &= \mathbb{E} \left[\sum_{k=1}^{K_T} [v(S^{t_k}, \boldsymbol{\theta}^k) - v(S^{t_{k+1}}, \boldsymbol{\theta}^k)] \right] \\ &\leq \mathbb{E} \left[\text{Span}(V) K_T \right]. \end{aligned}$$

4) *Bound on R_3* :

$$\begin{aligned} R_3 &= \mathbb{E} \left[\sum_{k=1}^{K_T} \sum_{t=t_k}^{t_{k+1}-1} [V_{\boldsymbol{\theta}^k}(\mathbf{S}^{t+1}) - \sum_{\mathbf{S}'} \boldsymbol{\theta}^k(\mathbf{S}'|\mathbf{S}^t, \mathbf{A}^t) V_{\boldsymbol{\theta}^k}(\mathbf{S}')] \right] \\ &= \mathbb{E} \left[\sum_{k=1}^{K_T} \sum_{t=t_k}^{t_{k+1}-1} [\sum_{\mathbf{S}'} (\boldsymbol{\theta}^*(\mathbf{S}'|\mathbf{S}^t, \mathbf{A}^t) - \boldsymbol{\theta}^k(\mathbf{S}'|\mathbf{S}^t, \mathbf{A}^t)) V_{\boldsymbol{\theta}^k}(\mathbf{S}')] \right], \end{aligned}$$

where inner summation is bounded by:

$$\begin{aligned}
 & \sum_{S'} (\boldsymbol{\theta}^*(S'|S^t, \mathbf{A}^t) - \boldsymbol{\theta}^k(S'|S^t, \mathbf{A}^t)) V_{\boldsymbol{\theta}^k}(S') \\
 & \leq \text{Span}(V) \sum_{S'} |\boldsymbol{\theta}^*(S'|S^t, \mathbf{A}^t) - \boldsymbol{\theta}^k(S'|S^t, \mathbf{A}^t)| \\
 & \leq \text{Span}(V) \sum_{S'} |\boldsymbol{\theta}^*(S'|S^t, \mathbf{A}^t) - \hat{\boldsymbol{\theta}}^k(S'|S^t, \mathbf{A}^t)| \\
 & \quad + \text{Span}(V) \sum_{S'} |\boldsymbol{\theta}^k(S'|S^t, \mathbf{A}^t) - \hat{\boldsymbol{\theta}}^k(S'|S^t, \mathbf{A}^t)|.
 \end{aligned}$$

Define confidence set

$$\begin{aligned}
 B_{mn}^k &= \{\boldsymbol{\theta} : \sum_{S'_{mn}} |\boldsymbol{\theta}(S'_{mn}|S_{mn}, S_{mn}) - \hat{\boldsymbol{\theta}}^k(S'_{mn}|S_{mn}, A_{mn})| \\
 & \leq \beta_{mn}^k(S_{mn}, A_{mn})\},
 \end{aligned}$$

where $\beta_{mn}^k(S_{mn}, A_{mn}) = \sqrt{\frac{14S_{max} \log 4t_k T}{\max(1, C_{mn}^{t_k}(S_{mn}, A_{mn}))}}$. Therefore we have

$$\begin{aligned}
 R_3 & \leq \underbrace{2\text{Span}(V) \mathbb{E} \left[\sum_{k=1}^{K_T} \sum_{t=t_k}^{t_{k+1}-1} \sum_{m,n} \beta_{mn}^k(S_{mn}^t, A_{mn}^t) \right]}_{\text{Term 1}} \\
 & \quad + \underbrace{2\text{Span}(V) \mathbb{E} \left[\sum_{k=1}^{K_T} \sum_{m,n} T_k (\mathbb{1}_{\boldsymbol{\theta}_{mn}^* \notin B_{mn}^k} + \mathbb{1}_{\boldsymbol{\theta}_{mn}^k \notin B_{mn}^k}) \right]}_{\text{Term 2}}.
 \end{aligned}$$

For Term 1, by the above definition of $\beta_{mn}^k(S_{mn}^t, A_{mn}^t)$, we have

$$\begin{aligned}
 & \sum_{k=1}^{K_T} \sum_{t=t_k}^{t_{k+1}-1} \sum_{m,n} \beta_{mn}^k(S_{mn}^t, A_{mn}^t) \\
 & = \sum_{k=1}^{K_T} \sum_{t=t_k}^{t_{k+1}-1} \sum_{m,n} \sqrt{\frac{14S_{max} \log 4t_k T}{\max(1, C_{mn}^{t_k}(S_{mn}, A_{mn}))}} \\
 & \leq \sum_{k=1}^{K_T} \sum_{t=t_k}^{t_{k+1}-1} \sum_{m,n} \sqrt{\frac{28S_{max} \log 4t_k T}{\max(1, C_{mn}^t(S_{mn}, A_{mn}))}} \\
 & = \sum_{t=1}^T \sum_{m,n} \sqrt{\frac{28S_{max} \log 4t_k T}{\max(1, C_{mn}^t(S_{mn}, A_{mn}))}} \\
 & \leq \sqrt{56S_{max} \log(T)} \sum_{t=1}^T \sum_{m,n} \frac{1}{\max(1, C_{mn}^t(S_{mn}, A_{mn}))} \\
 & \leq \sqrt{56S_{max} \log(T)} 3MN \sqrt{2S_{max} T} \\
 & = 3MNS_{max} \sqrt{112T \log(T)} \\
 & \leq 33MNS_{max} \sqrt{T \log(T)},
 \end{aligned}$$

where the first inequality is due to the fact that $C_{mn}^t(S_{mn}^t, A_{mn}^t) \leq 2C_{mn}^{t_k}(S_{mn}^t, A_{mn}^t)$ for all t in the k -th episode.

For Term 2,

$$\begin{aligned}
 & \sum_{k=1}^{K_T} \sum_{m,n} T_k (\mathbb{1}_{\boldsymbol{\theta}_{mn}^* \notin B_{mn}^k} + \mathbb{1}_{\boldsymbol{\theta}_{mn}^k \notin B_{mn}^k}) \\
 & = 2\mathbb{1}_{\boldsymbol{\theta}_{mn}^k \notin B_{mn}^k} = 2\mathbb{P}(\boldsymbol{\theta}_{mn}^k \notin B_{mn}^k).
 \end{aligned}$$

By the definition of confidence ball, we have

$$\mathbb{P}(\boldsymbol{\theta}_{mn}^k \notin B_{mn}^k) \leq \frac{1}{15Tt_k^6},$$

thus we get

$$\begin{aligned}
 & 2\text{Span}(V) \mathbb{E} \left[\sum_{k=1}^{K_T} \sum_{m,n} T_k (\mathbb{1}_{\boldsymbol{\theta}_{mn}^* \notin B_{mn}^k} + \mathbb{1}_{\boldsymbol{\theta}_{mn}^k \notin B_{mn}^k}) \right] \\
 & \leq \frac{4}{15} \text{Span}(V) \sum_{k=1}^{\infty} t_k^{-6} \leq \text{Span}(V).
 \end{aligned}$$

Combine the above results we have

$$\begin{aligned}
 R_3 & \leq 2\text{Span}(V) \mathbb{E} \left[\sum_{k=1}^{K_T} \sum_{t=t_k}^{t_{k+1}-1} \sum_{m,n} \beta_{mn}^k(S_{mn}^t, A_{mn}^t) \right] \\
 & \quad + 2\text{Span}(V) \mathbb{E} \left[\sum_{k=1}^{K_T} \sum_{m,n} T_k (\mathbb{1}_{\boldsymbol{\theta}_{mn}^* \notin B_{mn}^k} + \mathbb{1}_{\boldsymbol{\theta}_{mn}^k \notin B_{mn}^k}) \right] \\
 & \leq 66\text{Span}(V) MNS_{max} \sqrt{T \log(T)} + \text{Span}(V).
 \end{aligned}$$

The total regret then follows

$$\begin{aligned}
 R(T, \pi) &= R_1 + R_2 + R_3 \\
 & \leq E[K_T] + E[\text{Span}(V)K_T] \\
 & \quad + 66\text{Span}(V) MNS_{max} \sqrt{T \log(T)} + \text{Span}(V) \\
 & = 2\sqrt{S_{max} MNT \log T} (\text{Span}(V) + 1) \\
 & \quad + 66\text{Span}(V) MNS_{max} \sqrt{T \log(T)} + \text{Span}(V) \\
 & = \mathcal{O}(S_{max}^3 M^2 N^2 \sqrt{T \log T}).
 \end{aligned}$$

where the last equation holds from Lemma 1.

REFERENCES

- [1] T. S. Rappaport, S. Sun, R. Mayzou, H. Zhao, Y. Azar, K. Wang, G. N. Wong, J. K. Schulz, M. Samimi, and F. Gutierrez, "Millimeter wave mobile communications for 5g cellular: It will work!" *IEEE Access*, vol. 1, pp. 335–349, 2013.
- [2] J. G. Andrews, S. Buzzi, W. Choi, S. V. Hanly, A. Lozano, A. C. Soong, and J. C. Zhang, "What will 5g be?" *IEEE Journal on selected areas in communications*, vol. 32, no. 6, pp. 1065–1082, 2014.
- [3] M. Agiwal, A. Roy, and N. Saxena, "Next generation 5g wireless networks: A comprehensive survey," *IEEE Communications Surveys & Tutorials*, vol. 18, no. 3, pp. 1617–1655, 2016.
- [4] S. Jog, J. Wang, H. Hassanieh, and R. R. Choudhury, "Enabling dense spatial reuse in mmwave networks," in *Proc. of ACM SIGCOMM Posters and Demos*, 2018.
- [5] O. Abari, D. Bharadia, A. Duffield, and D. Katabi, "Enabling high-quality untethered virtual reality," in *Proc. of USENIX NSDI*, 2017.
- [6] Y. Niu, Y. Li, D. Jin, L. Su, and A. V. Vasilakos, "A survey of millimeter wave communications (mmwave) for 5g: opportunities and challenges," *Wireless networks*, vol. 21, no. 8, pp. 2657–2676, 2015.
- [7] D. Liu, L. Wang, Y. Chen, M. El-kashlan, K.-K. Wong, R. Schober, and L. Hanzo, "User association in 5g networks: A survey and an outlook," *IEEE Communications Surveys & Tutorials*, vol. 18, no. 2, 2016.
- [8] M. Mezzavilla, M. Zhang, M. Polese, R. Ford, S. Dutta, S. Rangan, and M. Zorzi, "End-to-end simulation of 5g mmwave networks," *IEEE Communications Surveys & Tutorials*, vol. 20, no. 3, pp. 2237–2263, 2018.
- [9] X. Wang, L. Kong, F. Kong, F. Qiu, M. Xia, S. Arnon, and G. Chen, "Millimeter wave communication: A comprehensive survey," *IEEE Communications Surveys & Tutorials*, vol. 20, no. 3, pp. 1616–1653, 2018.
- [10] X. Li, T. N. Guo, and A. B. Mackenzie, "Multi-agent reinforcement learning with measured difference reward for multi-association in ultra-dense mmwave network," *IEEE Access*, vol. 10, pp. 118747–118758, 2022.

- [11] W. Feng, Y. Wang, D. Lin, N. Ge, J. Lu, and S. Li, "When mmwave communications meet network densification: A scalable interference coordination perspective," *IEEE Journal on Selected Areas in Communications*, vol. 35, no. 7, pp. 1459–1471, 2017.
- [12] G. Yang, M. Xiao, and H. V. Poor, "Low-latency millimeter-wave communications: Traffic dispersion or network densification?" *IEEE Transactions on Communications*, vol. 66, no. 8, pp. 3526–3539, 2018.
- [13] S. K. Singh, V. S. Borkar, and G. Kasbekar, "User association in dense mmwave networks as restless bandits," *IEEE Transactions on Vehicular Technology*, 2022.
- [14] X. Liu, E. K. Chong, and N. B. Shroff, "A framework for opportunistic scheduling in wireless networks," *Computer networks*, vol. 41, no. 4, pp. 451–474, 2003.
- [15] I.-H. Hou, V. Borkar, and P. Kumar, "A theory of qos for wireless," in *Proc. of IEEE INFOCOM*, 2009.
- [16] T. Lan, D. Kao, M. Chiang, and A. Sabharwal, "An axiomatic theory of fairness in network resource allocation," in *Proc. of IEEE INFOCOM*, 2010.
- [17] F. Li, J. Liu, and B. Ji, "Combinatorial sleeping bandits with fairness constraints," *IEEE Transactions on Network Science and Engineering*, vol. 7, no. 3, pp. 1799–1813, 2019.
- [18] P. Whittle, "Restless Bandits: Activity Allocation in A Changing World," *Journal of Applied Probability*, pp. 287–298, 1988.
- [19] R. Ortner, D. Ryabko, P. Auer, and R. Munos, "Regret Bounds for Restless Markov Bandits," in *Proc. of ALT*, 2012.
- [20] Y. H. Jung and A. Tewari, "Regret Bounds for Thompson Sampling in Episodic Restless Bandit Problems," *Proc. of NeurIPS*, 2019.
- [21] N. Akbarzadeh and A. Mahajan, "On learning whittle index policy for restless bandits with scalable regret," *arXiv preprint arXiv:2202.03463*, 2022.
- [22] J. Fu, Y. Nazarathy, S. Moka, and P. G. Taylor, "Towards Q-Learning the Whittle Index for Restless Bandits," in *Proc. of Australian & New Zealand Control Conference*, 2019.
- [23] K. E. Avrachenkov and V. S. Borkar, "Whittle index based q-learning for restless bandits with average reward," *Automatica*, vol. 139, 2022.
- [24] J. A. Killian, A. Biswas, S. Shah, and M. Tambe, "Q-learning lagrange policies for multi-action restless bandits," in *Proc. ACM SIGKDD*, 2021.
- [25] G. Xiong, J. Li, and R. Singh, "Reinforcement Learning Augmented Asymptotically Optimal Index Policies for Finite-Horizon Restless Bandits," in *Proc. of AAAI*, 2022.
- [26] G. Xiong, X. Qin, B. Li, R. Singh, and J. Li, "Index-aware reinforcement learning for adaptive video streaming at the wireless edge," in *Proc. of ACM MobiHoc*, 2022.
- [27] S. Wang, L. Huang, and J. Lui, "Restless-UCB, an Efficient and Low-complexity Algorithm for Online Restless Bandits," in *Proc. of NeurIPS*, 2020.
- [28] G. Xiong, S. Wang, and J. Li, "Learning infinite-horizon average-reward restless multi-action bandits via index awareness," in *Proc. of NeurIPS*, 2022.
- [29] Y. Sun, G. Feng, S. Qin, and S. Sun, "Cell association with user behavior awareness in heterogeneous cellular networks," *IEEE Transactions on Vehicular Technology*, vol. 67, no. 5, pp. 4589–4601, 2018.
- [30] X. Zhang, S. Sarkar, A. Bhuyan, S. K. Kasera, and M. Ji, "A q-learning-based approach for distributed beam scheduling in mmwave networks," in *Proc. of IEEE DySPAN*, 2021.
- [31] M. G. Dogan, Y. H. Ezzeldin, C. Fragouli, and A. W. Bohannon, "A reinforcement learning approach for scheduling in mmwave networks," in *Proc. of IEEE MILCOM*, 2021.
- [32] T. H. L. Dinh, M. Kaneko, K. Wakao, K. Kawamura, T. Moriyama, H. Abeysekera, and Y. Takatori, "Deep reinforcement learning-based user association in sub6ghz/mmwave integrated networks," in *Proc. of IEEE CCNC*, 2021.
- [33] E. Altman, *Constrained Markov decision processes*. CRC Press, 1999, vol. 7.
- [34] G. Yao, M. Hashemi, R. Singh, and N. B. Shroff, "Delay-optimal scheduling for integrated mmwave sub-6 ghz systems with markovian blockage model," *IEEE Transactions on Mobile Computing*, 2022.
- [35] Y. Cao, B. Sun, and D. H. Tsang, "Delay-aware scheduling over mmwave/sub-6 dual interfaces: A reinforcement learning approach," in *2020 IEEE International Conference on Communications Workshops (ICC Workshops)*. IEEE, 2020, pp. 1–6.
- [36] D. C. Araújo and A. L. de Almeida, "Beam management solution using q-learning framework," in *2019 IEEE 8th International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP)*. IEEE, 2019, pp. 594–598.
- [37] Y.-N. R. Li, B. Gao, X. Zhang, and K. Huang, "Beam management in millimeter-wave communications for 5g and beyond," *IEEE Access*, vol. 8, pp. 13 282–13 293, 2020.
- [38] M. L. Puterman, *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, 1994.
- [39] I. M. Verloop, "Asymptotically Optimal Priority Policies for Indexable and Nonindexable Restless Bandits," *The Annals of Applied Probability*, vol. 26, no. 4, pp. 1947–1995, 2016.
- [40] Y. Zou, K. T. Kim, X. Lin, and M. Chiang, "Minimizing Age-of-Information in Heterogeneous Multi-Channel Systems: A New Partial-Index Approach," in *Proc. of ACM MobiHoc*, 2021.
- [41] W. Hu and P. Frazier, "An Asymptotically Optimal Index Policy for Finite-Horizon Restless Bandits," *arXiv preprint arXiv:1707.00205*, 2017.
- [42] G. Zayas-Cabán, S. Jasin, and G. Wang, "An Asymptotically Optimal Heuristic for General Nonstationary Finite-Horizon Restless Multi-Armed, Multi-Action Bandits," *Advances in Applied Probability*, vol. 51, no. 3, pp. 745–772, 2019.
- [43] X. Zhang and P. I. Frazier, "Restless bandits with many arms: Beating the central limit theorem," *arXiv preprint arXiv:2107.11911*, 2021.
- [44] G. Xiong, S. Wang, G. Yan, and J. Li, "Reinforcement learning for dynamic dimensioning of cloud caches: A restless bandit approach," in *Proc. of IEEE INFOCOM*, 2022.
- [45] D. J. Hodge and K. D. Glazebrook, "On the asymptotic optimality of greedy index heuristics for multi-action restless bandits," *Advances in Applied Probability*, vol. 47, no. 3, pp. 652–667, 2015.
- [46] K. D. Glazebrook, D. J. Hodge, and C. Kirkbride, "General notions of indexability for queueing control and asset management," *The Annals of Applied Probability*, vol. 21, no. 3, pp. 876–907, 2011.
- [47] J. Niño-Mora, "Dynamic Priority Allocation via Restless Bandit Marginal Productivity Indices," *Top*, vol. 15, no. 2, pp. 161–198, 2007.
- [48] N. Karmarkar, "A new polynomial-time algorithm for linear programming," in *Proc. of ACM STOC*, 1984.
- [49] J. Dunagan and S. Vempala, "A simple polynomial-time rescaling algorithm for solving linear programs," in *Proc. of ACM STOC*, 2004.
- [50] J. A. Kelner and D. A. Spielman, "A randomized polynomial-time simplex algorithm for linear programming," in *Proc. of ACM STOC*, 2006.
- [51] Gurobi Optimizer, <https://www.gurobi.com/solutions/gurobi-optimizer/>.
- [52] R. R. Weber and G. Weiss, "On An Index Policy for Restless Bandits," *Journal of Applied Probability*, pp. 637–648, 1990.
- [53] Y. Ouyang, M. Gagrani, A. Nayyar, and R. Jain, "Learning unknown markov decision processes: A thompson sampling approach," *Advances in neural information processing systems*, vol. 30, 2017.
- [54] P. L. Bartlett and A. Tewari, "Regal: A regularization based algorithm for reinforcement learning in weakly communicating mdps," *arXiv preprint arXiv:1205.2661*, 2012.
- [55] "IEEE 802.11ad-2012," <https://standards.ieee.org/ieee/802.11ad/45271/>, Accessed:08-July-2022.
- [56] C. R. da Silva, A. Lomayev, C. Chen, and C. Cordeiro, "Analysis and simulation of the ieee 802.11 ay single-carrier phy," in *Proc. of IEEE ICC*, 2018.
- [57] V. S. Borkar and S. Pattathil, "Whittle indexability in egalitarian processor sharing systems," *Annals of Operations Research*, vol. 317, no. 2, pp. 417–437, 2022.
- [58] N. Gast and G. Bruno, "A mean field model of work stealing in large-scale systems," *ACM SIGMETRICS Performance Evaluation Review*, vol. 38, no. 1, pp. 13–24, 2010.




Shufan Wang is currently pursuing the Ph.D. degree at Stony Brook University. He received the B.S. degree from Nanjing University in June 2017, and the M.S. degree from Binghamton University in September 2022. His research interests include online algorithms and reinforcement learning.



Guojun Xiong (Student Member, IEEE) is currently pursuing the Ph.D. degree at Stony Brook University. He received the B.S. degree in information science and technology from the Sun Yat-Sen University in 2015, and the M.S. degree in Electrical Engineering and Computer Science from University of Kansas in 2020. His research interests include wireless communication, networking, optimization and control, and reinforcement learning.




Shichen Zhang is currently a Ph.D. student in the Department of Computer Science and Engineering at Michigan State University (MSU), East Lansing, MI. He received his B.Eng. degree in Automation from Beijing University of Technology, Beijing, China, in 2018 and M.Eng. degree in Electrical and Computer Engineering from Cornell University, Ithaca, NY, in 2019. His current research interests focus on wireless networks, sensing systems, and deep learning.



Huacheng Zeng (Senior Member, IEEE) received a Ph.D. degree in Computer Engineering from Virginia Polytechnic Institute and State University (Virginia Tech), Blacksburg, VA. He is currently an Assistant Professor in the Department of Computer Science and Engineering at Michigan State University (MSU), East Lansing, MI. Prior to joining MSU, he was an Assistant Professor of Electrical and Computer Engineering at the University of Louisville, Louisville, KY, and a Senior System Engineer at Marvell Semiconductor, Santa Clara, CA. He is a

recipient of the NSF CAREER Award. He received the Best Paper Award from IEEE SECON 2021 and ACM WUWNet 2014. His research interest is broadly in wireless networking and sensing systems.



Jian Li (Member, IEEE) is an Assistant Professor of Data Science at Stony Brook University. Before that he was an Assistant Professor at Binghamton University, and a postdoc at University of Massachusetts Amherst. He received the Ph.D. degree in Computer Engineering from Texas A&M University in December 2016, and the B.E. degree from Shanghai Jiao Tong University in June 2012. His current research interests lie at the intersection of algorithms for reinforcement learning, federated learning, and stochastic optimization and control, with applications to next-generation networked systems.



Shivendra S. Panwar (Fellow, IEEE) received the B.Tech. degree in electrical engineering from the Indian Institute of Technology Kanpur, and the Ph.D. degree in electrical and computer engineering from the University of Massachusetts, Amherst, MA, USA. He is currently a Professor with the Electrical and Computer Engineering Department, NYU Tandon School of Engineering. He is the Director of the New York State Center for Advanced Technology in Telecommunications, the Co-Founder of New York City Media Lab, and a member of NYU Wireless.

He has been a Visiting Scientist at the IBM T.J. Watson Research Center and a consultant to Bell Laboratories. He has co-authored a textbook titled TCP/IP Essentials: A Lab-Based Approach (Cambridge University Press). His research interests cover the performance analysis and design of networks. His current research focuses on cross-layer research issues in wireless networks and multimedia transport over networks. He was a co-winner of the IEEE Communication Society's Leonard Abraham Prize for 2004, the ICC Best Paper Award in 2016, the IEEE HPSR Best Paper Award in 2023, and the Sony Research Award. He was a recipient of the Best Paper in 2011 Multimedia Communications Award. He has served as the Secretary for the Technical Affairs Council of the IEEE Communications Society. He is a Fellow of the National Academy of Inventors.