Self-Tuning Sectorization: Deep Reinforcement Learning Meets Broadcast Beam Optimization

Rubayet Shafin[®], Hao Chen, Young-Han Nam, Sooyoung Hur[®], Jeongho Park, Jianzhong Zhang, *Fellow, IEEE*, Jeffrey H. Reed[®], and Lingjia Liu[®], *Senior Member, IEEE*

Abstract—Beamforming in multiple input multiple output (MIMO) systems is one of the key technologies for modern wireless communication. Creating appropriate sector-specific broadcast beams are essential for enhancing the coverage of cellular network and for improving the broadcast operation for control signals. However, in order to maximize the coverage, patterns for broadcast beams need to be adapted based on the users' distribution and movement over time. In this work, we present selftuning sectorization: a deep reinforcement learning framework to optimize MIMO broadcast beams autonomously and dynamically based on users' distribution in the network. Taking directly UE measurement results as input, deep reinforcement learning agent can track and predict the UE distribution pattern and come up with the best broadcast beams for each cell. Extensive simulation results show that the introduced framework can achieve the optimal coverage, and converge to the oracle solution for both single sector and multiple sectors environment, and for both periodic and Markov mobility patterns.

Index Terms—MIMO, sectorization, broadcast beams, coverage maximization, self-organizing networks, and deep reinforcement learning.

I. INTRODUCTION

CELLULAR data traffic has witnessed an exponential growth over the last few years primarily due to the widespread use of mobile devices and novel application services. Cisco Visual Networking Index (VNI) forecast predicts a threefold increase of global IP traffic from 122 exabyte (EB) in 2017 to 296 EB in 2022 [1]. In order to handle this massive data-flow and ensure superior quality of experience (QoE) to

Manuscript received April 4, 2019; revised December 26, 2019; accepted February 20, 2020. Date of publication March 19, 2020; date of current version June 10, 2020. The work of Rubayet Shafin and Lingjia Liu was supported in part by the U.S. National Science Foundation (NSF) under Grant ECCS-1811497 and Grant CCF-1937487. The work of Jeffrey H. Reed was supported by U.S. NSF under Award 1642873. The associate editor coordinating the review of this article and approving it for publication was E. de Carvalho. (*Corresponding author: Lingjia Liu.*)

Rubayet Shafin was with Samsung Research America, Plano, TX 75023 USA. He is now with Wireless@VT, Bradley Department of Electrical and Computer Engineering, Virginia Tech, Blacksburg, VA 24061 USA.

Hao Chen and Jianzhong Zhang are with the Standards and 5G Mobility Innovation Laboratory, Samsung Research America, Plano, TX 75023 USA. Young-Han Nam was with Samsung Research America, Plano, TX 75023

USA. He is now with Mavenir, Richardson, TX 75081 USA.

Sooyoung Hur and Jeongho Park are with the Air Technology Lab, Samsung Electronics, Suwon 443742, South Korea.

Jeffrey H. Reed and Lingjia Liu are with Wireless@VT, Bradley Department of Electrical and Computer Engineering, Virginia Tech, Blacksburg, VA 24061 USA (e-mail: ljliu@ieee.org).

Color versions of one or more of the figures in this article are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TWC.2020.2979446

the end users, wireless cellular networks are also becoming extremely complicated. With the coexistence of different types of networks, managing networks efficiently has become a critical issue for 5G [2]-[4] and beyond systems. In order to reduce the network management complexity and operational cost, self organizing network (SON) has been introduced in Third Generation Partnership Project (3GPP) as one of the enabling technologies for advanced mobile networks [5], [6]. SON aims to achieve autonomous functionalities within Radio Access Network (RAN). These self-X functionalities include self-configuration, self-optimization, and self-healing [7], [8]. Self-optimization within SON refers to the process of selftuning of network parameters for achieving optimum performance in terms of any predefined metric of interest. The idea is to dynamically update the cellular radio resource parameters based on the changes in propagation characteristics, traffic pattern or network deployment scenarios. User distribution in wireless cellular network changes dynamically over time. These changes are the result of users' mobility behavior. For instance, in the day time, users are more densely populated in the commercial area whereas at night, users are primarily clustered in residential areas. Users' large time-scale movement also depends on specific time within the week (workdays and weekends) or year (holidays). Accordingly, to maximize the overall throughput and coverage of the wireless networks, sector-specific cellular radio parameters should also be updated taking into account the changes in users' distribution.

Multiple Input Multiple Output (MIMO) system [9] is one of the back-bones for current and next generation cellular network. Massive MIMO [10], where a large number of antennas are deployed at the base stations (BS), is envisioned as a key enabler for 5G systems. Beamforming refers to a MIMO technique for coherently combining the signals generated by multiple antennas in the MIMO arrays. 3-dimensional (3D) massive MIMO/full-dimension (FD) MIMO [11]-[13] promises tremendous throughput gain by enabling simultaneous beamforming in both elevation and azimuth domain. With large antenna array, it is possible to create sharp narrow beams towards desired users, and hence reduce the interference significantly [14]; this beamforming is used to improve users' throughput and is therefore user-specific. Cellular networks, on the other hand, also require to create wide beams. In fact, sectorization can be viewed as a process of expansive beam generation where a separate wide beam is used to cover a separate sector belonging to the same cell-site.

1536-1276 © 2020 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information. These sector-specific broad beams are essential for connecting as many users as possible, essentially providing the coverage for cellular networks. Another important application for wide beams is the broadcast technologies for sending out the wireless control and access signals as prescribed by LTE and LTE-Advanced systems. As a result, generating the accurate wide broadcast beam patterns that cover the maximum number of users in the network is critical.

Unfortunately, most of the works in the MIMO literature focus on maximizing MIMO throughput or increasing the reliability of the data plane. Meanwhile, at present, broadcast beam parameters are set manually in modern cellular networks: a group of network engineers do the drive tests and physically visit each base station site to fix the parameters controlling the shape, tilt and beam-widths of these sectorspecific broadcast beams. Once fixed, these broadcast beam parameters are not changed until some major fault/complain is reported. In other words, the parameters remain unchanged for a long period of time– often years, and as a result, currently, these parameter cannot be updated based on users' movement or change in user distribution. Accordingly, this fixed parameter setup results in strictly suboptimal solution in terms of overall network coverage.

Reinforcement learning (RL) is shown to be a useful tool for dynamic spectrum access (DSA) as well as small cell networks. A Q-learning based framework has been introduced in [15] for managing cumulative interference, originated from multiple cognitive radios, at the primary users' receivers in wireless regional area networks (WRANs). The introduced RL system is shown to autonomously learn policy that handles the cumulative interference at the primary users and keeps interference level at the primary protection contour below a predefined threshold. An RL-based power control strategy has been developed in [16] for cognitive femtocell networks, and it has been shown that RL can enhance the capacity of femtocells while ensuring a minimum quality of service (QoS) to macrocells. In a similar setup, [17] proposes an RL framework for interference management in small-cell networks. The problem of dynamic channel assignment (DCA) has been addressed in [18] by utilizing a real-time RL-based approach. A Multirate transmission control (MTC) strategy has been proposed in [19] using Q-learning algorithm for wideband code division multiple access (CDMA) systems.

In cellular networks, users movement changes in a dynamic fashion. To maximize the coverage area, the broadcast beam parameters controlling cell-sectorization need to be dynamically updated based on user movement. However, selecting the best broadcast beams simultaneously for all sectors in the network and updating the beams autonomously based on users' movement or distribution is a challenging problem, primarily because of the large number of combinatorial possibilities for beam selection-this pertains to very large action space in RL framework. Recently deep reinforcement learning (DRL) [20], [21] has been proved to be capable of learning human-level control policies on a varieties of different Atari games [22]. DRL agents learn to estimate the Q-values of selecting the best possible actions from current state of the video games. However, compared to traditional Q-learning, in deep learning based Q-network, the Q-values are approximated using deep neural network instead of storing the Q-values for all state-action pairs in a tabular form. As a result, DRL has the ability to predict the correct Q-values even for very large state and action space. Our recent work [23] shows that DRL based resource allocation can help improve the network performance of a DSA network.

In this work, we present a DRL-based framework for MIMO broadcast beam optimization for optimal cell-sectorization in order to maximize the coverage of cellular network. This will be an important step towards realizing the potential of SON. Our detailed contributions in this paper are summarized below:

- We introduce a double deep Q-network (DQN)-based framework [24] for dynamically optimizing sectorspecific MIMO broadcast beams for cellular network. The learning-based algorithm can autonomously update or self-tune the beam parameters based on users' mobility patterns or changes in user distribution.
- 2) We introduce self-tuning sectorization algorithms for both single sector and multiple sector environments. For multiple sector case, we have introduced a novel neural network architecture for computing the Q-values corresponding to different broadcast beam selection, with only linear increase in complexity as the number of BSs increases.
- 3) Finally, we present extensive simulation work using raytracing data for validating our solution. We consider both periodic and Markov mobility patterns, and show that the introduced DRL-based algorithm can achieve perfect convergence with *Oracle* for both single sector and multiple sector environment and for any user distribution.

The rest of the paper is organized as follows: Section II presents the network model and problem statement; Section III presents the beam learning framework; Section IV introduces the DRL-based optimization strategies for both single cell and multiple cell environments; Section V presents the simulation work before we conclude the paper in Section VI.

II. NETWORK MODEL AND PROBLEM STATEMENT

We consider a cellular network consisting of G BSs and K UEs. We assume the BSs can have one or multiple sectors, and there are total M sectors in the network, where $M \geq G$. Each sector is equipped with a two dimensional (2D) antenna array whose phases can be configured so that different arraybeam widths (in both elevation and azimuth domain) and elevation tilt (e-tilt) angle can be updated. Placing 2D antenna array enables the BSs to beamform in both elevation and azimuth directions, and this is essentially the setup for full dimension (FD) MIMO systems [11], [25]. The elevation beam-width, ϕ , azimuth beam-width, ψ , and e-tilt angle, ζ , constitute the parameter set in constructing the broadcast beams for each sector. In this work, we focus on optimizing the broadcast beams/sector-wide beams for cellular network. Let us denote the number of antenna elements in elevation and azimuth directions by N_1 and N_2 , respectively. Hence, total $N = N_1 N_2$ number of antenna weights need to be tuned for generating the FD-MIMO broadcast beams. We can represent the $N_1 \times N_2$ antenna weight matrix into a $N \times 1$ weight

TABLE I NOTATION FOR SYSTEM VARIABLES

Variable	Notation
No. of BSs	G
No. of Sectors	M
No. of UEs	K
Elevation beam-width	ϕ
Azimuth beam-width	ψ
E-tilt angle	ζ
No. of antennas at the BSs in elevation direction	N_1
No. of antennas at the BSs in azimuth direction	N_2
Total no. of antennas	N
Broadcast signal from <i>m</i> -th BS	x_m
Broadcast beamforming vector for <i>m</i> -th BS	\mathbf{f}_m
Received signal at k -th UE	y_k
Channel between m -th BS and k -th UE	$\mathbf{h}_{m,k}$
Beam-pool	\mathcal{W}
No. of possible beams in beam-pool	J
j-th beam-weight vector in beam-pool	\mathbf{w}^{j}
n-th antenna weight in j -th beam	w_n^j
UEs' SINR threshold for connectivity	T

vector, w, following a vectorization operation. Each choice of weight vector, w, in fact, consists of a specific choice of ϕ , ψ , and ζ . A collection of notations used in this paper is summarized in Table I.

Assuming each UE has a single antenna, the downlink broadcast received signal at k-th UE under m-th cell-sector can be written as

$$y_{k} = \mathbf{h}_{m,k}^{T} \mathbf{f}_{m} x_{m} + \sum_{\substack{m'=1\\m' \neq m}}^{M} \mathbf{h}_{m',k}^{T} \mathbf{f}_{m'} x_{m'} + n_{k}, \qquad (1)$$

where $\mathbf{h}_{m,k}$ is the $N \times 1$ channel vector for the channel between *m*-th sector and the *k*-th UE, x_m is the broadcast signal from *m*-th sector, and \mathbf{f}_m is the corresponding $N \times 1$ broadcast precoding vector for *m*-th sector. It can be clearly observed from (1) that broadcast beams from one sector interfere with the beams from other sectors. Hence, in order to maximize the network coverage, selecting the appropriate broadcast beams for all the sectors is critical.

In this work, we adopt a DRL-based approach where an agent is responsible for selecting the proper antenna configurations for all sectors. Each BS, for its sectors, has the same pool of possible antenna weight vectors available, W: $\{\mathbf{w}^1, \mathbf{w}^2, \dots, \mathbf{w}^J\}$, where J is the total number of beamweight vectors in the pool; $\mathbf{w}^j = [w_1^j, w_2^j, \dots, w_N^j]$ is the *j*-th vector in the beam pool, and w_n^q is the antenna weight for the *n*-th antenna element corresponding to *q*-th weight vector. Accordingly, each sector chooses its precoder, f, from the beam pool, i.e., $\mathbf{f}_m \in \mathcal{W}$. It is to be noted here again that each of the weight vector in the pool corresponds to a particular choice of elevation and azimuth beam-widths and e-tilt angle. The agent selects one out of J beam patterns for each sector based on users' distribution/mobility patterns. This selection behavior is referred to as actions in reinforcement learning.

All BS in the network transmit sector-specific signals using the wide broadcast beams selected by the agent. UEs collect measurement results such as Reference Signal Received Power (RSRP) or Reference Signal Received Quality (RSRQ), and report them to the agent as observation of the mobile environment. Assuming k-th UE in the network is associated with m-th sector, from (1), the received signal-to-interference-plusnoise ratio (SINR) for k-th user can be expressed as:

$$\operatorname{SINR}_{k} = \frac{\left|\mathbf{h}_{m,k}^{T}\mathbf{f}_{m}\right|^{2}}{\sum_{\substack{m'=1\\m'\neq m}}^{M}\left|\mathbf{h}_{m',k}^{T}\mathbf{f}_{m'}\right|^{2} + \sigma^{2}},$$
(2)

where σ^2 is the noise variance. In this work, we use the number of connected UEs as a metric to measure the cell coverage. Number of connected UEs in the network can be defined as the number of UEs whose received signal-to-interference-plus-noise ratio (SINR) are above a predefined threshold, *T*. For any user distribution, the objective, hence, is to select the optimal beam pattern indices for all the sectors under all BSs that maximize the coverage or total number of connected UEs in the network. The problem can formally be written as:

$$\max_{\mathbf{f}_1, \mathbf{f}_2, \dots, \mathbf{f}_M} \sum_{k=1}^K \mathbb{1}_{\mathrm{SINR}_k > T}$$
(3)

s.t.
$$\mathbf{f}_m \in \mathcal{W}, \quad 1 \le m \le M,$$
 (4)

where the indicator function, $\mathbb{1}_{x>T}$, is defined as

$$\mathbb{1}_{x>T} := \begin{cases}
1, & \text{if } x > T \\
0, & \text{if } x \le T.
\end{cases}$$
(5)

The user distribution changes over time, and hence optimal beam patterns that maximize the number of connected UEs at time t_1 may not be the same as that at time t_2 , where $t_1 \neq t_2$. The agent, therefore, has to be able to identify users' mobility pattern, and then dynamically and autonomously select optimal beams for all the sectors in order to maximize network coverage. It is to be noted here that we are not using users' location information to optimize the beam patterns. In order to minimize the feedback from the network, the agent will be merely using users' RSRP values to for the optimization.

In this work, we consider both single cell and multiple cells network scenarios. In the single cell case, the agent optimizes the broadcast beam for one cell-this represents a noise-limited environment. In this case, the DRL only needs to learn the optimal beam according to the cell environment including UE mobility pattern. On the other hand, in the multiple-cell case, the broadcast beams for all the cells need to be updated simultaneously- this represents an interferencelimited environment. We are addressing the challenges of these two scenarios where UEs are assumed to be moving according to some mobility pattern; first, the periodic case, where users' movement change in a periodic fashion, and second, the Markov case, where users' mobility is determined by following a transition probability matrix.

III. LEARNING FRAMEWORK

In this section, we present learning framework for MIMO broadcast beam optimization using DRL as a self-tuning sectorization mechanism. We first briefly describe the background of DRL which will set up the foundation for broadcast beamlearning strategy developed in subsequent sections.

A. Beam Learning Framework

Appropriate MIMO broadcast beam selection for cellsectorization is critical for wireless network performance optimization. Our objective here is to build a mechanism that automatically facilitates the selection of best beams for all the sectors. Moreover, we would need the sectors autonomously update their beam parameters based on different scenarios or user distributions, and realize self-tuning sectorization. Towards this, our learning framework is described as follows:

1) Specification of Design Parameters: First of all, network designer needs to decide on the objective function that needs to be optimized [26]. For broadcast beam optimization, an important objective function is the network coverage or total number of connected UEs in the network. The optimization parameters in this problem are the beam weights for each sector antenna element. It is necessary to select the optimal beam for each BS from a set of possible beams. Next, the system designer needs to decide on what input, such as RSRP or RSRQ, are required from the UEs in order to learn their mobility behavior and optimize the beams. Finally, in order to avoid random broadcast beams during the deployment stage, a simulation platform based on ray-tracing data is built to train the DRL agent offline.

2) Learning Engine: An agent or learning engine has the task of learning the UE mobility pattern and selecting the best beam parameters for each scenario. It takes feedbacks from UEs as inputs, and suggests the optimal beam vectors for all sectors. Updating the beams based on user distribution by autonomously identifying underlying mobility pattern requires training. However, online training is often not desirable because of stringent network management requirements from the operators. Hence, the training needs to be done offline, and the training environment has to be close to the real cellular environment as much as possible so that the optimal beams in the training stage—the procedure is presented in details in the next subsection.

3) Online Deployment and Occasional Re-Training: Once the learning engine is trained offline, the learned agent is deployed for real-time operation. It will enable the BSs to choose the optimal beams and update the selections based on users mobility pattern. Since users' mobility pattern in the network don't change too frequently, the beam parameters learned offline can remain unchanged for a long period of time– on the order of weeks or months. Whenever, there is a need to support new scenarios or any change in mobility patterns is identified, the learning engine would need to be re-trained offline based on recent data. The newly learned beam parameters will then be pushed to the respective BSs for updated operation.

B. Offline Training

Dynamically updating the broadcast beam patterns according to the cellular environment and user distribution for all cells in real time is intrinsically a difficult problem. Directly deploying a DRL agent and training it online is not only slow but also costly. During the online training stage of the DRL,

the agent may output some random beams according to the greedy exploration algorithm. Some of these random beams may not be acceptable to operators because of degraded network performance. In order to address this issue, we develop an offline training mechanism using ray-tracing data to train the DRL network before real deployment. By providing azimuth angle of arrival, elevation angle of arrival, azimuth angle of departure, elevation angle of departure, and path loss value of each path for each location in a cell, raytracing can well-capture the cellular environment so that the learned beam in the offline training platform could be the same as the online deployment case. The offline training is focused on learning the UE distribution pattern from users' location history data. The location data includes UEs' location and the corresponding time stamp. The location history data contains the UEs' mobility pattern information. Together with ray-tracing data, which contains the information about signal propagation environment, UEs' location history data are used to train the DRL network so that the DRL agent could learn the best broadcast beam according to both the cellular environment and UE distribution pattern. It is to be noted that for each training time step, the BSs select one set of actions and it throws the agent to a new state, upon which the new reward is computed. Hence, for each training step, the agent needs to access one ray-tracing data. After offline training, it will be deployed to provide real-time broadcast beam selection results for all the BSs in the cellular network. In the following, we describe the detailed steps of offline training.

According to 3GPP standard on minimization of drive test (MDT), a BS could configure its UEs to report measurement results, time stamp, and location information [27]. Therefore, we assume that UE location history information is available for a cellular network. During one training step, a batch of time stamps are selected from the location history data, and the corresponding UEs' location information is incorporated to ray-tracing data for every time stamp. Therefore, the UE distribution at the selected timestamp is combined with ray-tracing data. We call the ray-tracing data with UE distribution information as scenario-specific ray-tracing data and the UEs who report their measurement information during the timestamp as selected UEs. Based on the current BSs' broadcast beam and scenario-specific ray-tracing data, the receive power for the selected UEs could be calculated and accordingly the network coverage. A reward could be provided to the DRL agent based on the coverage and the DRL agent could accordingly update its selection of broadcast beams based on selected optimizer. These offline training steps could be repeated many times until the DRL agent converges. After the DRL agent converges, it could be deployed in the cellular network for real-time broadcast beam selection. Details on the DRL agent design is discussed in next section. The entire offline training process is pictorially depicted in Fig. 1 and Algorithm 1.

IV. DRL FOR BROADCAST BEAM OPTIMIZATION

In this section, details on the design of DRL framework for self-tuning sectorization are presented. The DRL network is utilized in order to track optimal beams during both the offline training and online deployment. To be specific, a deep

Algorithm 1 Offline Training

Input:

1: UE location history data, ray-tracing data of a cellular network

Output:

2: trained DRL agent for broadcast beam selection

STEP 1: Initialization

3: Define a pool of candidate antenna patterns;

STEP 2: Learning Best Beams

- 4: while algorithm doesn't achieve convergence do
- 5: Select a batch of UE location at different timestamps;
- 6: incorporate UE location distribution to ray-tracing data to create scenario-specific ray-tracing data;
- 7: calculate the received power for each UE in the scenariospecific ray-tracing data based on the current BSs' broadcast beam;
- 8: calculate the network coverage, and calculate a total reward as a function of network coverage;
- 9: DRL updates its neural weights based on the learning algorithm and reward

10: end while



Fig. 1. Offline training.

Q-network (DQN)-based architecture has been introduced to select MIMO broadcast beams for all sectors in a dynamic environment. For better stability of the results, we use DQN with *experience replay* [21], [28]. The agent (decision maker) interacts with the environment by selecting the best broadcast beam parameters. The DRL has three main components: state, action and reward. The dynamics between state, action and reward are shown in Fig. 2. Agent interacts with environment by observing the state of the network, and taking action that maximizes the reward or network performance metric.

A. Background of DRL

We consider a reinforcement learning framework where an agent or controller dynamically interacts with an unknown environment, \mathcal{E} , by taking sequential decisions or actions in discrete time steps. At each time step, t, the agent interacting with the environment observes a state, $s_t \in S$, selects an action, a_t , from a set of allowable actions, \mathcal{A} , and receives an immediate scalar reward, $r_t \in \mathcal{R}(s_t, a_t)$. Based on agent's current action, agent enters into new state, s_{t+1} . The cumulative discounted reward, R_t , at time step, t, is defined as

$$R_t = \sum_{k=0}^{\infty} \gamma^k r_{t+k},\tag{6}$$



Fig. 2. Reinforcement learning framework for beam optimization.

where $\gamma \in (0, 1]$ is the reward discount factor, which balances between the impact of recent rewards and earlier rewards. The learning objective is to maximize the expected cumulative reward at each state, s_t . The Q-value, $Q^{\pi}(s, a)$, for state-action pair, (s, a), is defined as the expected cumulative discounted reward for taking action, a, in state, s, and following a policy, π , onward, i.e.,

$$Q^{\pi}(s,a) = \mathbb{E}[R_t|s,a],\tag{7}$$

where $\mathbb{E}[\cdot]$ denotes expectation. Q-learning adopts a value iteration approach to find the Q-values for each state-action pair, and optimal value function $Q^*(s, a)$ is the one which provides maximum action value for state, s, and action, a, achievable by following any policy:

$$Q^*(s,a) = \max_{\pi} Q^{\pi}(s,a)$$
. (8)

Using Bellman equation [20], the optimal value function in (7) can be expressed as

$$Q^{*}(s,a) = \mathbb{E}_{s'}\left[r_t + \gamma \max_{a'} Q^{*}(s',a')|s,a\right].$$
 (9)

The value iteration algorithm can solve the Bellman equation, and the update rule is given by

$$Q_{i+1}(s,a) \leftarrow \mathbb{E}_{s'}\left[r_t + \gamma \max_{a'} Q_i(s',a')|s,a.\right].$$
(10)

In deep Q-learning, the value functions are approximated by deep neural network parameterized by the weights, ζ :

$$Q(s, a, \zeta) \approx Q^{\pi}(s, a). \tag{11}$$

This helps to estimate the Q-values even for very large stateaction space, and reduces the computational complexity. Next, we describe each of these components in details, and explain how we model the state, action, and reward in DRL-based MIMO broadcast beam optimization problem.

1) State: State in the introduced RL framework is designed as to reflect the network coverage situation which can be obtained from UE measurements. To be specific, we can design the state as the connection indicators of UEs in the network (a vector of 1/0s). Each UE reports its status to its attached BS. If a UE's SINR falls below a predefined threshold, T, a zero is placed at the element of the vector corresponding to that UE. Otherwise, a one is placed. Accordingly, a '0' in the state vector will represent that the corresponding UE has poor connection, and a '1' will indicate that the UE has good connection. The DRL state representation adopted in this work is pictorially depicted in Fig. 3.



Fig. 3. DRL state representation for beam optimization problem.

2) Action: An action of the agent is defined as the selection of beam index from a pool of candidate beam patterns. Agent observes the states and the corresponding reward, and takes the best possible action that maximizes the cumulative discounted future reward for the next time step. At the beginning of training, agent explores different actions in an attempt to learn the best beams for different user distribution. However, once the training phase is complete, agent exploits the learned information and only selects the best known actions that maximize the cumulative reward for each user distribution. We would like to highlight the fact that continuous beam space would not be feasible for cellular network coverage optimization since it can produce many beams which may not be practically realizable at BS arrays. Hence, beam pool should consist of discrete set of beams and need to be judiciously selected based on the particular network under consideration.

3) Reward: A reward in this work refers to any network performance metric. One way to design the reward can be the total number of connected UEs in the network based on the state and action taken in the previous state. Another approach to design the reward can be the function of the measurement results, for example, a function of the SINR or RSRP vector. In this work, we adopt the first approach for designing the reward. It is to be noted here that maximizing total number of connected UEs in the network is equivalent to maximizing the coverage of the cellular network.

The agent's goal is to maximize the cumulative discounted future reward. The agent gathers its experiences as tuples, (s_t, a_t, r_t, s_{t+1}) , where s_t denotes current UE connection state, a_t denotes the action taken at state, s_t ; r_t is the instantaneous reward obtained from state, s_t and by taking action, a_t ; and s_{t+1} is the next state. The agent stores history of its experiences in a memory called experience replay memory [28], and replay memory stores the tuples, (s_t, a_t, r_t, s_{t+1}) , for all time steps. The DRL agent randomly samples mini-batches of experience from the replay memory for training, and selects an action based on ϵ -greedy policy, i.e., with probability ϵ , it tries a random action, and with probability $(1 - \epsilon)$, the agent selects the best known action so far. The optimum action in a particular state is selected based on maximum Q-values [20] corresponding to that state. In DQN-based reinforcement learning, the Q-values are predicted using deep neural network. Input to the neural network is the UEs' connection vector representing the state of the

RL environment, and output is the Q-values corresponding to all the possible actions, i.e., beam indices from the beam-pool. In the following subsections, we detail the broadcast beam optimization strategies for both single cell and multiple cell scenarios.

B. Broadcast Beam Optimization for Dynamic Environment

In this subsection, we present the framework for dynamically optimizing MIMO broadcast beams, where the RL agent needs to simultaneously control the beam parameters for all the sectors based on different user distributions. For the single cell case, beam parameters corresponding to only one sector need to be optimized. This could serve as an example where a legacy LTE sector is replaced with one massive MIMO unit. The goal is to maximize the number of connected UEs for different dynamic user distributions. The agent keeps a single replay memory containing the agent's experience tuples, and randomly samples from it-this random sampling from experience replay memory helps to decorrelate the data [22]. However, for multiple sectors case, there needs to be some significant changes on the RL framework compared to that for single sector beam optimization. In the multiple sector environment, each sector has its own pool of beams or action sets. Each sector can hence independently select its own beam parameters. The setup is similar to that of multi agent system [29], [30]. The goal remains the same-to maximize the overall network coverage. This is a challenging problem in terms of computational tractability. For an illustration, let us consider that there are m sectors in the network, and each sector has j possible beam patterns (actions) to select from. Hence, total number of actions, i.e., all possible combinations of sectors' beam patterns, becomes j^m , which increases exponentially with total number of sectors. If there are 40 base stations, and each has 5 possible actions to choose from, total possible combination of beam patterns becomes 5^{40} , which is an extraordinarily large number, making it difficult to achieve optimal solution within reasonable time.

One way to find the appropriate broadcast beams for multiple sectors simultaneously is to use a single large neural network with large number of output nodes that can be used to predict the Q-values for all possible j^m actions. However, total number of training samples needed to train such neural network would be extremely large, which may not be feasible at all for any practical purposes. In other words, the learning algorithm can almost never achieve convergence with this architecture for even moderate size cellular network.

To address this issue, we introduce a novel low-complexity algorithm for optimizing the broadcast beams for multiple sectors where the action space grows only linearly, instead of exponentially, with total number of sectors in the network. Let us again assume that there are m sectors, and each sector has j possible actions (beam-weight set) to choose from. Unlike the single cell case, for multiple cell environment, we assume the agent preserves different replay memories for different sectors. Moreover we use m different neural networks for independently computing the Q-values for j sectors. Each neural network is responsible to predict the optimum action for the corresponding sector only. With this architecture, number of



Fig. 4. Replay buffer architecture for multiple sector case.

actions increases only linearly, but we can still achieve perfect convergence with reasonably short computation time, which demonstrated through extensive simulation in Section V. It is to be noted that deep Q learning algorithm proposed in [21] is designed to create a single NN-based agent that can learn to play Atari games where the number of valid actions for the player is quite limited. On the other hand, in terms of training methodology, the architecture presented in this section for multiple sector scenario is scalable with growing action space, and in this sense, it provides a architectural generalization of the work in [21]. The details of the architectures for replay memory and neural networks for multiple sector broadcast beam optimization are briefly described next.

1) Replay Memory Architecture: The replay memory architecture for multiple sectors broadcast beam optimization is shown in figure 4. There are separate buffers for each sector. The same current state, reward, and the next state are stored in all the replay memories/buffers for the sectors. However, the replay memories differ in the actions taken (beam indices chosen) by the each sector. While all the sectors observe the same current state, s_t , reward, r_t , and next state, s_{t+1} , the action stored are different–BS 1's action is stored in buffer 1, BS 2's action is stored in buffer 2, and so on. The rationale behind this buffer architecture is that states and rewards are network specific, and same states and rewards are observed by all sectors. On the other hand, each sector takes its own action, and their joint actions regulate the overall network state and the corresponding reward.

2) Neural Network Architecture: For Q-value prediction, a deep convolutional neural network is used in this work. For the suitability of computing the Q-values using convolutional neural network, we transform the $(K \times 1)$ UE connection vector into an $(\frac{K}{100} \times 100)$ frame. Four such frames are stacked together, and fed as the input to the neural network for computing the Q-values. We used three convolutional layers– all with rectified linear unit (ReLU) activation function. First convolution layer has 32 (8 × 8) filters. Second and third convolution layers have 64 (4 × 4) filters and 64 (3 × 3) filters, respectively. Finally, a dense layer with linear activation function is used as the output layer.

Two such identical neural networks are used in predicting the Q-values. One is used for computing the running Q-values—this neural network is called the *evaluation network*. The other neural network, called the *target network*, is held fixed for some training duration, say for P episodes, and every P episodes, the weights of the evaluation neural network is

Algorithm 2 Broadcast Beam Optimization for Multiple Sectors

Input:

- 1: RSRP measurements from the UEs in the network **Output:**
- 2: Optimum broadcast beam patterns for all sectors that maximizes the number of connected UEs

STEP 1: Initialization

- 3: Define a pool of candidate antenna pattern;
- 4: Define the maximum exploration rate, ϵ_{max} , minimum exploration rate, ϵ_{min} , exploration decay rate, optimizer's learning rate, α , and reward discount factor, γ ;
- 5: Initialize the replay memory, \mathcal{D} .

STEP 2: Optimization of Beam Weights

- 6: **for** episode = 1, 2, ..., Z, **do**
- 7: Initialize the state vector at time step 1 as s_1 ;
- 8: **for** t = 1, 2, ..., T', **do**
- 9: Sample c from Uniform (0, 1)
- 10: **if** $c \leq \epsilon$ **then**

11: Select an action (choose a beam index) for each sector randomly from the beam pool

12: else

13: **for** m = 1, 2, ..., M **do**

14: Select the action for *m*-th BS,
$$a_t^m = \arg\max_{a^m} Q_m^*(\mathbf{s}_t, a^m; \theta_m)$$

15: end for

- 16: **end if**
- 17: Apply the selected beam patterns on the antenna arrays of the corresponding BSs
- 18: Observe the resulting RL state, s_{t+1} , the UE connection vector.
- Pre-process the state vector into a frame before feeding to Neural Network
- 20: Compute the reward, r_t , which is the number of connected UEs.
- 21: **for** m = 1, 2, ..., M **do**

22: Store the experience tuple for *m*-th sector, $e_t^m = (\mathbf{s}_t, a_t^m, r_t, \mathbf{s}_{t+1})$, in *m*-th replay memory, \mathcal{D}^m .

23: Sample random mini-batches of experience $(\mathbf{s}_j, a_j^m, r_j, \mathbf{s}_{j+1})$, from \mathcal{D}^m

24: if
$$s_{j+1}$$
 is a terminal state then
25: Set $y_j^m = r_j$
26: else
27: Set $y_j^m = r_j + \gamma \max_{a_m} Q_m(\mathbf{s}_t, a_m; \theta)$
28: end if
29: Perform a gradient descend on
 $(y_j^m - Q(s_j, a_j^m; \theta))^2$
30: end for
31: end for
32: end for

transferred to the target neural network. It has been shown that this two neural network-based approach for Q-value prediction provides better stability of results at convergence [22].

The neural network architecture for predicting the Q-values for multiple sectors are shown in Fig. 5. The depiction



Fig. 5. Neural network architecture for multiple sector case.



Fig. 6. Periodic change in scenarios.

is presented for M sectors case, where M separate neural networks are used for predicting the Q-values for M sectors. Input to all neural networks are the same state vectors. Neural networks are identical, and the number of output for each neural network is J. Hence, size of action space is JM, instead of J^M , i.e., total number of actions grows only linearly with number of sectors. The optimal action predicted by the Q-values of neural network 1 is stored in Buffer 1, which corresponds to sector 1. Similarly, the action predicted by the Q-values of neural network 2 is stored in Buffer 2, which corresponds to sector 2, and so on. It is to be noted that neural networks do not share any weight information during training, and each neural network independently predicts the optimal actions for the corresponding sectors. Hence, there is no additional signaling overhead among the neural networks. The beam learning procedure for multiple BS environment is presented in Algorithm 2.

V. SIMULATION RESULTS AND PERFORMANCE ANALYSIS

In this section, we present the simulation results and performance evaluation for self-tuning sectorization mechanism through DRL-based MIMO broadcast beam optimization. We first present the results for single sector environment followed by multiple sectors case. Both periodic and Markov mobility patterns have been considered for the evaluation.

A. Results for Single Sector Dynamic Environment:

In this sub-section, we present the performance evaluation for our algorithm for single sector dynamic environment. The sector is equipped with a two dimensional (2D) antenna array with 4 antenna elements in both elevation and azimuth directions. The horizontal distance between BS antenna elements is 0.5 wave-length and the vertical distance between antenna elements is 1.48 wave-length. We first consider two scenarios or user distributions, and assume that users switch between Scenario-1 and Scenario-2 periodically every 8 time steps (see Fig. 6). The BS is located at a height of 35 m from ground, and users are distributed randomly in the cell. Based on users' X-, Y-, and Z-coordinates, two scenarios are defined

IADLE II	
SIMULATION PARAMETERS	
RL Parameter	Specification
Reward Discount Factor, γ	0.0001
Learning Rate, α	0.001

	0.000
Learning Rate, α	0.001
Initial exploration probability, ϵ_{max}	1.0
Final exploration probability, ϵ_{\min}	0.000001
Training batch size	32
Optimizer	Adam
Network Parameter	Specification
Antenna array at BSs	4×4
Antenna separation in azimuth domain	1.48λ
Antenna separation in elevation domain	0.5λ
UEs' SINR threshold for connectivit, T	-6 dB
BSs height from the ground	35m

as follows: Scenario-1: $X \ge 2600$ m, $Z \ge 10$ m; Scenario-2: $X \le 2700$ m, $Z \le 12$ m. For simulation, this partition is used as users' mobility pattern. The received power of each UE is calculated based on ray-tracing data. Noise level is set as -95 dBm, and SINR threshold level is kept at 0.1 dB. For a particular user, if the received SINR is above this threshold, we consider the user to be connected; otherwise, we consider it to be not-connected. A set of simulation parameters used in this work is summarized in Table II. The general rationale for selecting the hyper-parameters are described below:

1) Initial Exploration Rate: At the beginning of training, agent needs to gather experiences, and explore as much as possible. Accordingly, the initial exploration rate, ϵ_{max} , is set to 1, which corresponds to complete exploration and no exploitation.

2) Final Exploration Rate: Towards the end of training, agent should have acquired enough knowledge about the environment and the underlying user distributions. In this phase, rather than exploration, the agent should focus on exploitation by taking the already known best actions for different sectors. Hence, final exploration rate should be close to zero. However, in order to avoid the situation where two rewards are very close to each other and the agent is stuck with the slightly lesser reward, the final exploration, ϵ_{\min} , in practice, is not set at exactly 0. In this work, we set ϵ_{\min} at a very small number, 0.000001, which correspond to very high exploitation phase.

3) Exploration Decay Rate: Exploration decay rate is set based on the total number of training samples available and the number of training samples used for initial exploration phase. Usually, the exploration rate is decreased in regular interval. Denoting the number of training samples dedicated for initial observation as $T_{\rm obs}$, in this work, we followed the algorithm below for decaying the exploration rate at time step, t.

if $t \leq T_{obs}$ then
$\epsilon_t = \epsilon_{\max}$
else if $t > T_{obs}$ and $\epsilon_t > \epsilon_{min}$ then
$\epsilon_t = \epsilon_{t-1} - \frac{(\epsilon_{\max} - \epsilon_{\min})}{T_{\exp}}$
else
$\epsilon_t = \epsilon_{\min}$
end if



Fig. 8. Results for periodic mobility pattern in a single sector dynamic environment: (a) average squared difference (ASD) between reward achieved by DRL agent and the reward obtained by oracle; (b) average mismatch (AM) between actions taken by the DRL agents and the oracle.



Fig. 9. Users' distribution patterns for 2 scenarios.

Here, T_{expl} denotes a parameter > 1 controlling the speed of decay. In our work, we set $T_{\text{expl}} = 5000$, and $T_{\text{obs}} = 1000$.

4) Learning Rate: Learning rate, α , determines how fast information acquired from recent experiences overrides the

information from prior experiences. In practice, α is set between 0 and 1. A learning rate of 0 implies the Q-values are never updated, and hence no learning takes place. On the other extreme, a learning rate of 1 means the agent only considers



Fig. 10. Results for periodic mobility pattern in a multiple sector dynamic environment: (a) average squared difference (ASD) between reward achieved by DRL agent and the reward obtained by oracle; (b) average mismatch (AM) between actions taken by DRL agents for each sector and the corresponding oracles.



Fig. 11. Instantaneous rewards (a) and instantaneous actions (b) at convergence for multiple sectors environment and periodic user-mobility pattern.

the information from the most recent experience, and ignores any information previously acquired. In this work, we start training with initial learnig rate, $\alpha = 0.001$, and every 20000 training steps, we reduce our learning rate by a factor of 10.

5) Reward Discount Factor: Reward discount factor, γ , indicates how the agent values the future reward. In practice, the value of γ is set between 0 and 1. γ close to zero indicates that immediate rewards are more valued than the distant future rewards. On the other hand, γ close to 1 implies that long term cumulative future rewards are more important than the current reward. Based on our DRL environment, we set the reward discount factor at 0.0001.

At each time step, the RL agent has 10 actions to choose from, i.e., there are 10 different beam weight vectors available for the agent. Each of the actions corresponds to a unique beam pattern. As an illustration, one such beam pattern and the associated elevation and azimuth cuts are shown in Fig. 7. Based on the change in user distribution, the agent adaptively selects the beam that maximizes the total number of connected UEs. Figure 8a shows the average squared difference (ASD) between the reward (total number of connected UEs) obtained by the DRL agent and the reward predicted by Oracle:

$$ASD = \frac{1}{N'} \sum_{n=1}^{N'} \left(R_n^{\text{Agent}} - R_n^{\text{Oracle}} \right)^2, \qquad (12)$$

where R_n^{Agent} and R_n^{Oracle} denote instantaneous reward at *n*-th time step obtained by DRL agent and Oracle, respectively; N' represent the number of time steps used for averaging. *Oracle* is defined as an entity which has the complete and perfect knowledge of the environment and user



Fig. 12. Results for Global solution for periodic mobility pattern in a multiple sector dynamic environment: (a) average squared difference (ASD) between reward achieved by DRL agent and the reward obtained by oracle; (b) average mismatch (AM) between actions taken by DRL agents for each sector and the corresponding oracles.

distribution; it is essentially an exhaustive search method in order to compute the maximum attainable reward at any given scenario. Each point in Fig. 8a represent ASD over N' = 200 time steps. In Fig. 8a, we have also shown the shaded error bar, which represent the maximum difference from the mean value within every N' time steps. It can be observed that at the beginning of training, ASD between rewards obtained by the RL agent and the Oracle is quite high. However, as time goes by, ASD gradually decreases, and finally, at the completion of training, rewards from RL agent converges completely with that from Oracle. This is due to the fact that at the beginning of training, the agent explores different actions and collects the memory. During the exploration phase, the agent tries out all available actions, and attempts to learn the optimal beam weights for different user distributions. Over time, this exploration rate decreases, and exploitation increases, i.e., agent tends to choose more frequently the best known actions so far that maximize the reward.

Fig. 8b shows the results for average mismatch (AM) in actions (selected beam pattern) taken by the DRL agent and the Oracle, respectively, where AM is defined as

$$AM = \frac{1}{N'} \sum_{n=1}^{N'} \mathbb{1}_{\left(A_n^{\text{Agent}} \neq A_n^{\text{Oracle}}\right)},\tag{13}$$

where A_n^{Agent} and A_n^{Oracle} denote the actions selected for *n*-th time step by the DRL agent and the Oracle, respectively, and the indicator function, $\mathbb{1}_{(A_n^{\text{Agent}} \neq A_n^{\text{Oracle}})}$, is defined as

$$\mathbb{1}_{\left(A_{n}^{\text{Agent}}\neq A_{n}^{\text{Oracle}}\right)} = \begin{cases} 1, & \text{if } A_{n}^{\text{Agent}}\neq A_{n}^{\text{Oracle}}\\ 0, & \text{if } A_{n}^{\text{Agent}}=A_{n}^{\text{Oracle}}. \end{cases}$$
(14)

It can be observed that action mismatch is quite large at the start of the training because of high exploration rate. However, at the end of training phase, actions taken by the DRL agent and the Oracle converge completely, and average mismatch reduces to zero. It is to be noted that the introduced DRL-based self-sectorization method is applicable for any discrete number of actions. However, as the number actions grows large, the difference between optimal number of connected users corresponding to different best beam combinations becomes smaller. Further increasing the number of action would not provide much gain, however, may potentially cause longer training time, especially for the multiple sector scenarios.

B. Results for Multiple Sector Dynamic Environment

In this sub-section, we present the simulation results for multiple sector dynamic environment. We consider two sectors, each at a height of 35 m from ground. Each sector has two possible beam patterns to choose from. Two scenarios are considered similar to single sector case in the previous sub-section. The scenarios with line of sight (LoS) and nonline of sight (NLoS) UEs are shown in Fig. 9. We assume the scenarios periodically change every 8 time steps. The agent is responsible for simultaneously selecting the optimal beam patterns for both sectors for maximizing the number of connected UEs in the network. The average squared difference in rewards achieved by the agent and the oracle for multiple sectors scenario is shown in Fig. 10a. Similarly to single cell case, as training increases, overall rewards attained by the agent and the oracle converge completely. In other words, the agent is able to dynamically optimize the beam patterns for both sectors simultaneously in the interference environment, and maximize the overall rewards from the network in all scenarios or user distributions. In Fig. 10b, we show the average action mismatch for both sectors. It can be observed that towards the end of exploration phase, average action mismatches between the sectors and the corresponding Oracles reduce to zero. The instantaneous rewards and actions at



Fig. 13. Results for average squared difference (ASD) in reward between the DRL agent and the oracle for periodic mobility pattern in a single sector dynamic environment. ASDs for different size of action space have been plotted in figures (a) - (e).

convergence of the algorithm are shown in Fig. 11, where, for clarity, we zoom in for time steps between 4000 and 4030. We can observe that scenarios change every 8 time steps

and maximum number of connected UEs are different for the two scenarios. Optimal strategy for sector-1 is to select action 1 while in scenario 1, and select action 2 while in



Fig. 14. Results for ASD in reward between DRL agent and the oracle for periodic mobility pattern in a multiple sector dynamic environment. ASDs for different size of action space have been plotted in figures (a) and (b).

scenario 2. On the other hand, optimal strategy for sector-2 is to select action 2 for both scenarios. In reinforcement learning, it is, in general, difficult to obtain convergence if the reward values are too close. However, we can observe from Fig. 11 that the DRL agent can completely converge with the oracle and take the corresponding best actions even when the reward values for scenario 1 and scenario 2 differ. This indicates the accuracy of self-tuning sectorization strategy developed in this work.

Fig. 10 and Fig. 11 are based on our introduced neural network architecture, where Q-values corresponding to each sector is predicted by a separate neural network. For comparison, in Fig. 12, we present the global solution, where a single neural network is responsible for predicting the Q-values for all sectors. Hence, if there are 4 actions available for each sector, for a 2-sectors environment, the neural network needs to predict Q-values for $4^2 = 16$ actions. We can observe from Fig. 12 that for the single NN-based architecture, it requires more than 7500 time steps for the DRL to converge with Oracle. In comparison, for the introduced NN-architecture in Fig. 10, the DRL agent can converge with Oracle within only about 4000 time steps. These results demonstrate the training advantage of the introduced neural network architecture, especially for large action space, over the traditional state of the art DRL training method [21].

In general, if the action space grows large, more training time is required for the algorithm to converge. However, exact training time required can be determined through experimentation. For example, Fig. 13 shows the results on average squared difference (ASD) in reward (number of connected UEs) obtained by the DRL agent and the Oracle for a single sector. In Fig. 13, for a fixed set of hyper-parameters, we provide a comparison on how size of action space affects the convergence time, where we vary the number of available actions (possible beam patterns) from 2 to 10. We can observe that for the single sector case, training for the action sizes from 2 to 10 all can converge within about 3000 time steps.



Fig. 15. State transition diagram for Markov mobility.

On the other hand, Fig. 14 presents ASD between the DRL agent and the Oracle for two-cells dynamic environment where the comparison on convergence time is shown for the numbers of actions 2 and 4. Unlike the performance on single cell cases in Fig. 13, for the multiple sectors case, we can observe that as the number of actions doubles, from 2 to 4, it requires more time for the DRL agent to converge with the Oracle. However, from these experiments, it is notable that even though the action space increases linearly, the required convergence time doesn't need to be doubled with doubling the action space.

C. Multi-Sectors Environment With Markovian Mobility Pattern

In this sub-section, we present the performance analysis for DRL-based self-tuning beamforming in multiple sector environment and for the case where user distributions alternate between two scenarios following a Markovian mobility pattern. It is to be noted here that, in general, users' mobility pattern has some intrinsic regularity. For example, users can be clustered more in the commercial area during day time while they move to residential are in the evening. Hence, periodic mobility patterns considered in previous two subsections rather closely depict the actual mobility pattern in cellular network. Nevertheless, in this sub-section, we consider



Fig. 16. Results for Markov mobility pattern in a multiple sector dynamic environment: (a) average squared difference (ASD) between reward achieved by DRL agent and the reward obtained by oracle; (b) average mismatch (AM) between actions taken by DRL agent for each sector and the corresponding oracles.



Fig. 17. Instantaneous reward (a) and instanteneous actions (b) at convergence for multiple sectors environment and Markov user mobility pattern.

the Markovian mobility in order to verify the robustness of the developed self-tuning sectorization algorithm for the extreme case when users' mobility pattern doesn't have any regularity and users move between different scenarios in random fashion.

We consider two scenarios defined similarly to the ones in Section V-A, and assume the users' locations switch between these two scenarios with transition probability governed by the state transition diagram shown in Fig. 15. Moreover, we consider two sectors each having two possible beam patterns to choose from for each scenario. Fig. 16a shows the average squared difference for rewards attained by the RL agent and the oracle for Markov mobility pattern. We can observe that similarly to the periodic cases presented in previous two subsections, RL agent does converge with the oracle even for probabilistic mobility, and ASD goes to zero after the training phase. Average mismatch in actions between the sectors and the corresponding oracles are shown in Fig. 16b. It can be seen that average mismatch in actions for both sectors reduce to zero at the end of the training phase. Finally, the instantaneous rewards achieved and the actions taken by the sectors at convergence of the algorithm are shown in Fig. 17, which, again, indicates perfect convergence for Markov mobility pattern in multiple cell environment.

VI. CONCLUSION

In this work, we have developed a framework for self-tuning cell sectorization through MIMO broadcast beam optimization using deep reinforcement learning. To be specific, we have introduced learning strategies for both single sector and multiple sectors environment with dynamic user distribution. The introduced solutions can autonomously and adaptively update the RF parameters based on the changes in user distributions. Simulation results show that the DRL-based method completely converges with the Oracle-suggested optimal solutions for both periodic and Markovian user mobility patterns.

REFERENCES

- Cisco Visual Networking Index, "Global mobile data traffic forecast update, 2015–2020," Cisco Syst., San Jose, CA, USA, White Paper, Nov. 2018, p. 9.
- [2] J. G. Andrews *et al.*, "What will 5G be?" *IEEE J. Sel. Areas Commun.*, vol. 32, no. 6, pp. 1065–1082, Jun. 2014.
- [3] F. Boccardi, R. W. Heath, A. Lozano, T. L. Marzetta, and P. Popovski, "Five disruptive technology directions for 5G," *IEEE Commun. Mag.*, vol. 52, no. 2, pp. 74–80, Feb. 2014.
- [4] C.-X. Wang *et al.*, "Cellular architecture and key technologies for 5G wireless communication networks," *IEEE Commun. Mag.*, vol. 52, no. 2, pp. 122–130, Feb. 2014.
- [5] S. Hämäläinen, H. Sanneck, and C. Sartori, *LTE Self-Organising Networks (SON): Network Management Automation for Operational Efficiency*. Hoboken, NJ, USA: Wiley, 2012.
- [6] M. Peng, D. Liang, Y. Wei, J. Li, and H.-H. Chen, "Self-configuration and self-optimization in LTE-advanced heterogeneous networks," *IEEE Commun. Mag.*, vol. 51, no. 5, pp. 36–45, May 2013.
- [7] O. Sallent *et al.*, "A roadmap from UMTS optimization to LTE selfoptimization," *IEEE Commun. Mag.*, vol. 49, no. 6, pp. 172–182, Jun. 2011.
- [8] H. Hu, J. Zhang, X. Zheng, Y. Yang, and P. Wu, "Self-configuration and self-optimization for LTE networks," *IEEE Commun. Mag.*, vol. 48, no. 2, pp. 94–100, Feb. 2010.
- [9] L. Liu, R. Chen, S. Geirhofer, K. Sayana, Z. Shi, and Y. Zhou, "Downlink MIMO in LTE-advanced: SU-MIMO vs. MU-MIMO," *IEEE Commun. Mag.*, vol. 50, no. 2, pp. 140–147, Feb. 2012.
- [10] T. L. Marzetta, "Noncooperative cellular wireless with unlimited numbers of base station antennas," *IEEE Trans. Wireless Commun.*, vol. 9, no. 11, pp. 3590–3600, Nov. 2010.
- [11] Y. Kim *et al.*, "Full dimension MIMO (FD-MIMO): The next evolution of MIMO in LTE systems," *IEEE Wireless Commun.*, vol. 21, no. 2, pp. 26–33, Apr. 2014.
- [12] R. Shafin, L. Liu, J. Zhang, and Y.-C. Wu, "DoA estimation and capacity analysis for 3-D millimeter wave Massive-MIMO/FD-MIMO OFDM systems," *IEEE Trans. Wireless Commun.*, vol. 15, no. 10, pp. 6963–6978, Oct. 2016.
- [13] R. Shafin and L. Liu, "Multi-cell multi-user massive FD-MIMO: Downlink precoding and throughput analysis," *IEEE Trans. Wireless Commun.*, vol. 18, no. 1, pp. 487–502, Jan. 2019.
- [14] R. Shafin, L. Liu, Y. Li, A. Wang, and J. Zhang, "Angle and delay estimation for 3-D massive MIMO/FD-MIMO systems based on parametric channel modeling," *IEEE Trans. Wireless Commun.*, vol. 16, no. 8, pp. 5370–5383, Aug. 2017.
- [15] A. Galindo-Serrano and L. Giupponi, "Distributed Q-learning for aggregated interference control in cognitive radio networks," *IEEE Trans. Veh. Technol.*, vol. 59, no. 4, pp. 1823–1834, May 2010.
- [16] H. Saad, A. Mohamed, and T. ElBatt, "Distributed cooperative Q-learning for power allocation in cognitive femtocell networks," in *Proc. IEEE Veh. Technol. Conf. (VTC Fall)*, Sep. 2012, pp. 1–5.
- [17] M. Bennis, S. M. Perlaza, P. Blasco, Z. Han, and H. V. Poor, "Selforganization in small cell networks: A reinforcement learning approach," *IEEE Trans. Wireless Commun.*, vol. 12, no. 7, pp. 3202–3212, Jul. 2013.
- [18] J. Nie and S. Haykin, "A Q-learning-based dynamic channel assignment technique for mobile communication systems," *IEEE Trans. Veh. Technol.*, vol. 48, no. 5, pp. 1676–1687, 1999.
- [19] Y.-S. Chen, C.-J. Chang, and F.-C. Ren, "Q-learning-based multirate transmission control scheme for RRM in multimedia WCDMA systems," *IEEE Trans. Veh. Technol.*, vol. 53, no. 1, pp. 38–48, Jan. 2004.
- [20] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. Cambridge, MA, USA: MIT Press, 2018.
- [21] V. Mnih et al., "Playing atari with deep reinforcement learning," 2013, arXiv:1312.5602. [Online]. Available: http://arxiv.org/abs/1312.5602
- [22] V. Mnih et al., "Human-level control through deep reinforcement learning," *Nature*, vol. 518, no. 7540, pp. 529–533, Feb. 2015.
- [23] H.-H. Chang, H. Song, Y. Yi, J. Zhang, H. He, and L. Liu, "Distributive dynamic spectrum access through deep reinforcement learning: A reservoir computing-based approach," *IEEE Internet Things J.*, vol. 6, no. 2, pp. 1938–1948, Apr. 2019.
- [24] H. Van Hasselt, A. Guez, and D. Silver, "Deep reinforcement learning with double Q-learning," in *Proc. AAAI*, Phoenix, AZ, USA, vol. 2, 2016, p. 5.
- [25] Y.-H. Nam *et al.*, "Full-dimension MIMO (FD-MIMO) for next generation cellular technology," *IEEE Commun. Mag.*, vol. 51, no. 6, pp. 172–179, Jun. 2013.

- [26] S. Joseph, R. Misra, and S. Katti, "Towards self-driving radios: Physicallayer control using deep reinforcement learning," in *Proc. 20th Int. Workshop Mobile Comput. Syst. Appl. (HotMobile).* New York, NY, USA: ACM, 2019, pp. 69–74.
- [27] Radio Measurement Collection for Minimization of Drive Tests (MDT), document TS 37.320 V14.0.0, 3rd Generation Partnership Project, Mar. 2017.
- [28] L.-J. Lin, "Reinforcement learning for robots using neural networks," School Comput. Sci., Carnegie Mellon Univ., Pittsburgh, PA, USA, Tech. Rep. CMU-CS-93-103, 1993.
- [29] K. P. Sycara, "Multiagent systems," AI Mag., vol. 19, no. 2, p. 79, 1998.
- [30] M. Wooldridge, An Introduction to Multiagent Systems. Hoboken, NJ, USA: Wiley, 2009.



Rubayet Shafin received the B.S. degree in electrical and electronic engineering from the Bangladesh University of Engineering and Technology, Dhaka, Bangladesh, and the M.S. degree in electrical engineering from the University of Kansas in 2013 and 2017, respectively. He is currently pursuing the Ph.D. degree in electrical engineering with Virginia Tech. His research interests include MIMO systems, mmWave communications, and application of machine learning/deep learning in wireless communication. He was a recipient of the 2018 IEEE TAOS

Best Paper Award and the 2018 IEEE TCGCC Best Conference Paper Award.



Hao Chen received the B.S. and M.S. degrees in information engineering from Xi'an Jiaotong University, Shanxi, in 2010 and 2013, respectively, and the Ph.D. degree in electrical engineering from Kansas University, Lawrence, KS, USA, in 2017. Since 2016, he has been a Research Engineer with the Standards and Mobility Innovation Laboratory, Samsung Research America. His research interests include network optimization, machine learning, and 5G cellular systems.



Young-Han Nam received the B.S. and M.S. degrees from Seoul National University, South Korea, in 1998 and 2002, respectively, and the Ph.D. degree in electrical engineering from The Ohio State University, Columbus, OH, USA, in 2008. He was a Principal Engineer with Samsung Research America, Plano, TX, USA. He was engaged in base station algorithm development, standardization, design and analysis of the 3GPP LTE, LTE-Advanced, and 5G NR, from 2008 to 2019. In 3GPP standardization meetings in 2017, he served as a study item rappor-

teur of the 5G channel models to edit 3GPP TR38.900 and TR38.901. He is currently the Senior Director of Mavenir, Richardson, TX, USA, and also leading O-RAN Product Research and Development. His research interests include artificial intelligence, open RAN, MIMO, and channel modeling.



Sooyoung Hur received the B.S. degree in electrical engineering from Sogang University, Seoul, South Korea, in 2005, the M.S. degree in electrical engineering from the Korea Advanced Institute of Science and Technology, Daejeon, South Korea, in 2007, and the Ph.D. degree in electrical and computer engineering from Purdue University, West Lafayette, IN, USA, in 2013. Since 2013, he has been with Samsung Electronics. His research interests cover channel modeling and wireless system design for 5G and future wireless communication.



Jeongho Park received the B.S., M.S., and Ph.D. degrees in electronic engineering from Yonsei University in 1997, 2000, and 2005, respectively. Since 2005, he has been with Samsung Electronics, where he has been engaged in research and development of wireless communications systems. He is currently the Vice President of the Network Business Unit Research and Development Team, Samsung Electronics, and takes the lead of research and development of break-through air technologies in 5G and beyond 5G area. His current interests include

massive MIMO, millimeter-wave, load balancing, and SON technologies.



Jeffrey H. Reed is currently the Founder of Wireless@Virginia Tech, where he served as the Director until 2014. He is also the Founding Faculty Member of the Ted and Karyn Hume Center for National Security and Technology and served as an Interim Director when founded in 2010. He is also a Co-Founder of Cognitive Radio Technologies (CRT), a company commercializing of the cognitive radio technologies; Federated Wireless, a company developing spectrum sharing technologies; and PFP Cybersecurity, a company specializing in security for

embedded systems. His book *Software Radio: A Modern Approach to Radio Design* (Prentice Hall) and his latest textbook *Cellular Communications: A Comprehensive and Practical Guide* (Wiley-IEEE, 2014). In 2005, he became a Fellow to the IEEE for contributions to software radio and communications signal processing and for leadership in engineering education. He is a Past Member of CSMAC a group that provides advice to the NTIA on spectrum issues. In 2013, he was awarded the International Achievement Award by the Wireless Innovations Forum. In 2012, he served on the President's Council of Advisors of Science and Technology Advisory Group that examined ways to transition federal spectrum for commercial use.



Lingjia Liu (Senior Member, IEEE) received the B.S. degree in electronic engineering from Shanghai Jiao Tong University and the Ph.D. degree in electrical and computer engineering from Texas A&M University. Prior to joining the ECE Department at Virginia Tech (VT), he was an Associate Professor with the EECS Department, University of Kansas (KU). He spent more than four years working in the Mitsubishi Electric Research Laboratory (MERL) and the Standards and Mobility Innovation Lab, Samsung Research America (SRA), where he

received Global Samsung Best Paper Award in 2008 and 2010. He was leading Samsung's efforts on multiuser MIMO, CoMP, and HetNets in LTE/LTE-Advanced standards. His general research interests mainly lie in emerging technologies for 5G cellular networks including machine learning for wireless networks, massive MIMO, massive MTC communications, and mmWave communications.

He received the Air Force Summer Faculty Fellow, from 2013 to 2017, Miller Scholar at KU, 2014, Miller Professional Development Award for Distinguished Research at KU, 2015, the 2016 IEEE GLOBECOM Best Paper Award, the 2018 IEEE ISQED Best Paper Award, and the 2018 IEEE TAOS Best Paper Award.



Jianzhong (Charlie) Zhang (Fellow, IEEE) received the Ph.D. degree from the University of Wisconsin–Madison, Madison. From 2009 to 2013, he served as the Vice Chairman for the 3GPP RAN1 working group and led development of LTE and LTE-Advanced technologies such as 3D channel modeling, UL-MIMO, CoMP, and Carrier Aggregation for TD-LTE. He is currently the Vice President and the Head of the Standards and Mobility Innovation Lab, Samsung Research America, where he leads research, prototyping, and standards for 5G and future multimedia networks.