

# Deep Reinforcement Learning for Uplink Multi-Carrier Non-Orthogonal Multiple Access Resource Allocation Using Buffer State Information

Eike-Manuel Bansbach, Yigit Kiyak and Laurent Schmalen

Communications Engineering Lab, Karlsruhe Institute of Technology, 76187 Karlsruhe, Germany  
(email: e.bansbach@kit.edu)

**Abstract**—For orthogonal multiple access (OMA) systems, the number of served user equipments (UEs) is limited to the number of available orthogonal resources. On the other hand, non-orthogonal multiple access (NOMA) schemes allow multiple UEs to use the same orthogonal resource. This extra degree of freedom introduces new challenges for resource allocation. Buffer state information (BSI), like the size and age of packets waiting for transmission, can be used to improve scheduling in OMA systems. In this paper, we investigate the impact of BSI on the performance of a centralized scheduler in an uplink multi-carrier NOMA scenario with UEs having various data rate and latency requirements. To handle the large combinatorial space of allocating UEs to the resources, we propose a novel scheduler based on actor-critic reinforcement learning incorporating BSI. Training and evaluation are carried out using Nokia’s “wireless suite”. We propose various novel techniques to both stabilize and speed up training. The proposed scheduler outperforms benchmark schedulers.

## I. INTRODUCTION

While in 5G multiple access (MA) is mainly realized using orthogonal MA (OMA), non-orthogonal MA (NOMA) is considered as a key enabling technology to improve the spectral efficiency of next generation’s mobile communication networks [1]. Since orthogonal frequency-division multiple access (OFDMA) hinders the straightforward sharing of a physical resource block (PRB) by multiple user equipments (UEs) [2], the number of orthogonal PRBs limits the number of served UEs [3]. In contrast, NOMA allows multiple UEs to use the same PRB [2] by cohabitation of UEs in the power domain at the transmitter side and successive interference cancellation (SIC) at the receiver side [4]. Distinctness of the superposed messages of different UEs can be achieved by either power control of the UE’s transmit power or by combining UEs with sufficiently distinct channel gains. Due to OFDMA’s high robustness against frequency-selective fading [5], multi-carrier NOMA (MC-NOMA) as combination of both techniques [6] is considered as an answer to the challenges of the next generation of communication networks [2].

For an OFDMA system with UEs having alike quality of service (QoS) and time-invariant data rate requirements, the

resource allocation of PRBs to the UEs can be described as a convex optimization problem. If the UEs belong to different QoS classes with varying guaranteed bit rates (GBRs), varying maximum packet delay budgets (PDBs) and time-variant data traffic, the optimization problem is non-convex [5]. One way to solve this optimization problem is by modelling it as a Markov decision process (MDP) and use deep reinforcement learning (DRL) [7]. With wireless networks becoming increasingly heterogeneous, DRL-based protocols tailored to specific applications can outperform general-purpose solutions [8]. We discussed several DRL-based approaches for scheduling in OFDMA systems in [9]. The approach proposed in [9] uses buffer state information (BSI), i.e., the size and age of packets inside the buffer waiting for transmission. It is able to handle up to 32 UEs belonging to different QoS classes and outperforms non-DRL-based benchmark algorithms. However, NOMA capabilities aren’t considered.

Since NOMA increases the number of possible combinations when allocating resources, a DRL algorithm that can handle a large action space is necessary. While [10] uses Q-learning for NOMA resource allocation, all devices belong to the same QoS class and have time-invariant data traffic. In [11], the deep deterministic policy gradient (DDPG) algorithm, capable of dealing with an infinite number of DRL actions, e.g., allocation options, is introduced. While [12] and [13] use Q-learning for PRB assignment and DDPG for power allocation, [6] and [14] directly use DDPG algorithm for PRB assignment. In [6], a flexible amount of users can be allocated to a PRB for a downlink MC-NOMA scenario. This DDPG-based approach outperforms various benchmarks. In [14], an uplink MC-NOMA system is investigated. For time-varying data and two QoS classes, the DDPG-based approach outperforms a Q-learning based reference. In [6] BSI is neglected, in [14] BSI is reduced to the buffer queue length.

In this work, we extend our work of [9] towards a DDPG algorithm for uplink MC-NOMA resource allocation with full BSI. All UEs have alike transmit power and we ensure discriminability by learning to combine UEs with sufficiently distinct channel gains. Our main contributions are the inclusion of BSI into the uplink MC-NOMA resource allocation as well as, to the best of our knowledge, the first DRL-based algorithm for uplink MC-NOMA resource allocation without power control of the UEs’ transmit powers.

This work has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No. 101001899). Parts of this work have been funded by the German Federal Ministry of Education and Research (BMBF) within the project Open6GHub (grant agreement 16KISK010).

## II. UPLINK MC-NOMA SYSTEM MODEL

The used system model is the *NomaULTimeFreq-ResourceAllocation-v0* environment provided by Nokia's "Wireless Suite" problem collection [15]. It simulates an uplink MC-NOMA scenario and provides benchmarks. The set of  $K$  UEs is given by  $\mathcal{K} = \{\text{UE}_1, \dots, \text{UE}_K\}$  and the set of  $N$  PRBs by  $\mathcal{N} = \{\text{PRB}_1, \dots, \text{PRB}_N\}$ . Each UE has a buffer with  $L$  slots, which contain packets waiting for transmission. The UEs belong to one of four QoS classes, identified by their corresponding QoS identifier (QI)  $q_k \in \{1, 2, 3, 4\}$ . The QoS classes are given as GBR services, like conversational voice, conversational video and delay critical GBR, and as non-GBR services, like web browsing. Upon initialization of an environment, UEs are randomly spread over a  $1 \text{ km}^2$  squared area and assigned a QI. The area is an empty Euclidean space with a transceiver base station (BS) at its center. The UEs roam around at rectilinear trajectories with random speeds [16]. The speeds are independently sampled from a normal distribution according to [17]. At the edges of the area, the UEs bounce off at specular angles [16].

The environment is described by Alg. 1. First, the UEs are placed randomly across the area (initialization). At the beginning of every environment step, the scheduler receives the channel quality indicator (CQI)  $c_k \in \{0, \dots, 15\}$  of every UE as well as the ages and sizes of the packets inside a UE's buffer. The scheduler assigns up to  $M$  UEs to each PRB, where  $M$  is the number of NOMA capabilities of a single PRB. Thus, up to  $M$  UEs can do NOMA using a single PRB. The possible actions for each NOMA resource of a PRB are either assigning one of the  $K$  UEs or leave it empty, resulting in  $K+1$  possible actions. For every chosen UE, the achievable rate is calculated by

$$R = \frac{B}{N} \log_2(1 + \text{SINR}),$$

with the overall system bandwidth  $B$ . The UE's signal-to-interference-plus-noise-ratio SINR is calculated by

$$\text{SINR} = \frac{P_{\text{rx}}}{\sigma_n^2 + \zeta + \xi},$$

where  $P_{\text{rx}}$  is the UE's power received at the BS,  $\sigma_n^2$  the power of additive white Gaussian noise,  $\zeta \hat{=} -105 \text{ dBm}$  a constant interference power throughout the coverage area and  $\xi$  the interference caused by other UEs occupying the current PRB.  $P_{\text{rx}}$  is calculated using the distance  $d$  between the UE and the BS, applying the free-space loss to the transmit power  $P_{\text{tx}} = 13 \text{ dBm}$ , which is identical for all UEs. After each PRB is assigned and the transmitted bits are deleted from the buffers, it is checked whether the remaining packets exceed their latency requirements, given by their PDB. The PDB depends on the UE's QoS class. If there are packets that exceed their PDB, the environment returns a negative reward (penalty) by summing up the bits of packets exceeding their PDB [16]. Afterwards, the UEs move and new packets are generated. While SIC demodulates the signals of UEs in the order of decreasing received power and eliminates

---

### Algorithm 1 Pseudocode of *NomaULTimeFreqResourceAllocation-v0* environment [15]

---

```

Initialize:  $K$  UEs with random QI,  $N$  PRBs,  $M$  UEs per
PRB,  $L$  buffer slots per UE, maximum runtime  $T$ 
Randomly place UEs, initialize speed, generate packets
for  $t = 1, \dots, T_{\text{max}}$  do
  Get CQI and BSI of all UEs
  for  $n = 1, \dots, N$  do
    Choose up to  $M$  UEs to do NOMA on  $\text{PRB}_n$ 
    Set of chosen UEs:  $\mathcal{D} \subseteq \mathcal{K}, |\mathcal{D}| \leq M$ 
    NOMA Interference:  $\xi = 0$ 
    for  $\text{UE} \in \mathcal{D}$  do
      Calculate receive power  $P_{\text{rx}}$  and achievable rate  $R$ 

      Transmit packets, delete from UE's buffer
       $\xi \leftarrow \xi + P_{\text{rx}}$ 
    end for
    Update BSI
  end for
  Calculate penalty by checking packet ages [16]
  Move UEs, generate new packets
end for

```

---

the interfering waveforms one at a time [18, Sec. 16.3-4], the *NomaULTimeFreq-ResourceAllocation-v0* environment emulates SIC by stepwise adding interference. The reward function is predefined by the environment [16] and is not modified within this paper.

While the OFDMA resource allocation problem in [9] focuses on scheduling by the packet's urgency, the MC-NOMA problem extends the OFDMA problem by combining different UEs on the same PRB, while mitigating the interference among the UEs.

## III. REINFORCEMENT LEARNING

### A. Reinforcement Learning Problem

In RL, an agent tries to skillfully map actions to observed system states in order to maximize a numerical reward [9]. Starting at an observed state, the agent takes an action, receives a reward and follow-up state, takes another action, et cetera. This sequential decision making can be formalized as an MDP [19, Chap. 3]. Let  $\mathcal{S}$  denote a finite set of states,  $\mathcal{A}$  a finite set of actions and  $S_t \in \mathcal{S}$ ,  $A_t \in \mathcal{A}$  as well as  $W_t \in \mathbb{R}$  random variables describing the state, action and reward at time  $t$ . Assuming the Markov property is fulfilled, the dynamics of an MDP can be fully described by the state-transition probabilities  $p(S_{t+1} = s' | S_t = s, A_t = a) : \mathcal{S} \times \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$  with  $s', s \in \mathcal{S}$  and  $a \in \mathcal{A}$ , as well as the expected reward  $w(S_t = s, A_t = a, S_{t+1} = s') : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}$  when taking action  $A_t = a$  at state  $S_t = s$  with follow-up state  $S_{t+1} = s'$  [19, Sec. 3.1].

With  $G_t := \sum_{k=0}^{\infty} \gamma^k W_{t+k+1}$ ,  $\gamma \in [0, 1]$  as the discounted return, the cumulative reward in the long run is calculated [19, Sec. 3.3]. A policy  $\pi(a|s)$  defines probabilities of taking action  $A_t = a$  given state  $S_t = s$  in order to maximize  $G_t$  [19,

Sec 3.4]. Following a policy  $\pi$ , the action-value function  $q_\pi(s,a)$  denotes the expected return for choosing an action  $A_t = a$  at state  $S_t = s$  [19, Sec. 3.5]:

$$q_\pi(s,a) := \mathbb{E}_\pi \{G_t | S_t = s, A_t = a\} \quad \forall s \in \mathcal{S}, \forall a \in \mathcal{A}.$$

A policy  $\pi$  is called an optimal policy  $\pi^*$ , if its decisions maximize the action-value function

$$q^*(s,a) := \max_\pi q_\pi(s,a) \quad \forall s \in \mathcal{S}, \forall a \in \mathcal{A},$$

where  $q^*(s,a)$  is called the optimal action-value function [19, Sec. 3.6]. Similarly, the value function  $v_\pi(s)$  as the expected return when in  $S_t = s$  and following  $\pi$  can be defined by  $v_\pi(s) := \mathbb{E}_\pi \{G_t | S_t = s\} \forall s \in \mathcal{S}$  [19, Sec. 3.5]. The RL problem can be solved using value-based or policy-based methods. While the former try to learn  $q^*(s,a)$ , the latter directly optimize the policy  $\pi(a|s)$  [19, Chap. 13].

### B. Value-based Methods

The objective of Q-learning is to use temporal difference learning to learn a function  $Q(s,a)$  that approximates  $q^*(s,a)$  [20]. Performing action  $A_t = a$  at state  $S_t = s$  and observing the reward  $W_{t+1} = w$  and follow-up state  $S_{t+1} = s'$  yields the tuple  $(s,a,w,s')$ . The approximated Q-function  $Q(s,a)$  can be updated using

$$Q(s,a) \leftarrow (1 - \alpha)Q(s,a) + \alpha[w + \gamma \max_{a'} Q(s',a')],$$

with the step-size  $\alpha \in [0,1]$  [20]. Using a more accurate estimate incorporating the observed reward  $w$  for state-action pair  $(s,a)$ ,  $Q(s,a)$  gets updated. Q-learning converges to the optimal action-value function  $q^*(s,a)$  [20]. The optimal action  $A_t = a^*$  given state  $S_t = s$  is chosen greedily by  $a^* = \arg \max_{a \in \mathcal{A}} Q(s,a)$  [20].

### C. Policy-based Methods

Instead of learning  $q^*(s,a)$ , policy-based methods directly optimize the policy parameters  $\theta \in \mathbb{R}^d$  of a parametrized policy  $\pi_\theta := \pi(a|s,\theta)$ . With the performance measure  $J(\theta) := v_{\pi_\theta}(s_0)$  as the expected return when starting at state  $S_0 = s_0$  and following policy  $\pi_\theta$ , the policy can be updated using *gradient ascent* [19, Chap. 13]. Following the stochastic policy gradient theorem (PGT), it can be shown that [21]

$$\nabla J(\theta) \propto \mathbb{E}_\pi \left[ \sum_a q_\pi(S_t, a) \nabla_\theta \pi(a|S_t, \theta) \right].$$

$\nabla J(\theta)$  is proportional to the expectation of the gradient of the probability choosing action  $a$  given state  $S_t$ , weighted with its expected return  $q_\pi(s,a)$ . Hence, the computation of the performance gradient is reduced to a simple expectation [22]. Using stochastic gradient ascent,  $\theta$  can be updated by

$$\theta \leftarrow \theta + \epsilon \sum_a \hat{q}(S_t, a) \nabla_\theta \pi(a|S_t, \theta),$$

where  $\epsilon$  is a step-size and  $\hat{q}(S_t, a)$  is a learned approximation of  $q_\pi(s,a)$  [19, Sec. 13.2]. Hence, policy-gradient methods require an estimate of the action-value function  $q_\pi(S_t, a)$  [22].




Fig. 1. Schematic overview of an actor-critic algorithm adapted from [23, Fig. 1]. The dotted lines indicate that the critic is responsible for updating the actor and itself.

### D. Actor-critic Methods

While policy-based methods are able to handle large or continuous action spaces, value-based methods have a lower variance in the estimates of expected returns. Actor-critic methods combine the advantages of both [23]. A parametrized *actor*  $\pi_\theta(s)$  with parameters  $\theta$  is defined. Using, e.g., Q-learning, a *critic* with parameters  $\Phi$  estimates the action-value function  $Q(s,a|\Phi) \approx q_\pi(s,a)$  of the actor's policy  $\pi_\theta(s)$ . As depicted in Fig. 1, both the actor and the critic receive a state  $S_t$ . The actor chooses an action  $A_t$  following  $\pi_\theta(s)$ . Using the reward  $W_{t+1}$ , the critic updates its estimator  $Q(s,a|\Phi)$ . Afterwards,  $Q(s,a|\Phi)$  is used to update the actor according to the stochastic PGT update step [24].

### E. Deterministic Policy Gradient

Now consider a deterministic parametrized policy  $\rho_\theta(s)$  with parameters  $\theta$ , which maps exactly one action to every state. Given its performance measure  $J(\theta) := v_{\rho_\theta}(s_0)$  when starting at state  $S_0 = s_0$ , the deterministic PGT is given by [22]

$$\nabla J(\theta) = \mathbb{E}_\rho [\nabla_\theta \rho_\theta(S_t) \nabla_a q_\rho(S_t, a)|_{a=\rho_\theta(S_t)}],$$

where  $q_\rho(s,a)$  is the action-value function following policy  $\rho_\theta$ . Compared to the stochastic PGT, deterministic PGT updates only for the taken action  $a = \rho_\theta(S_t)$ . Moreover, instead of  $q_\rho(s,a)$ , only its gradient  $\nabla_a q_\rho(s,a)$  is taken into account, which avoids the estimation of  $q_\rho(s,a)$  and is computationally attractive. The stochastic policy gradient converges to the deterministic policy gradient for  $\sigma^2 \rightarrow 0$ , where  $\sigma^2$  is the output variance of the stochastic policy [22].

### F. Deep Deterministic Policy Gradient Algorithm

In [11], the DDPG algorithm, an actor-critic setup using deterministic PGT, is introduced. Since feedforward neural networks define a parametrized mapping  $\mathbf{y} = f(\mathbf{x}, \theta)$ , RL algorithms can be implemented using deep neural networks (DNNs) [19, Sec. 9.7]. However, there are two challenges when using DNNs [11]: First, most DNN optimization algorithms assume that the samples used for optimization are i.i.d.. Storing observed transitions  $(s_t, a_t, s_{t+1}, r_{t+1})$  in a finite replay memory and randomly sampling from it ensures learning using (approximately) independent transitions. Second, using only




Fig. 2. Sketch of the proposed decision process of scheduling the NOMA resources of one PRB. The decision matrix is initialized as  $\mathbf{X} = \mathbf{0} \in \mathbb{R}^{M \times (K+1)}$  and is updated row-wise using the action probabilities  $\mathbf{x}_m \in \mathbb{R}^{(K+1)}$  of NOMA-resource  $m$ .

a single DNN for, e.g., Q-learning, updating the DNN while using it for calculating the target  $Q(s', a')$  destabilizes learning. Using a copy of the networks, the actual network is updated and the copy, called target, is used to estimate  $Q(s', a')$ . The target network gets periodically synchronized with the updated network. The full algorithm is given by [11, Alg. 1].

#### IV. ACTOR-CRITIC METHODS FOR NOMA-OFDMA UPLINK RESOURCE ALLOCATION

##### A. Sequential NOMA Allocation

While the general DRL framework introduced in Sec. III receives  $S_t$ , takes  $A_t$  and immediately receives  $W_{t+1}$  and  $S_{t+1}$ , the uplink MC-NOMA scenario in Alg. 1 involves two major challenges. First, the BSI and, therefore, the state  $S_t$  is updated after all NOMA-resources of a PRB are allocated. Thus, multiple sequential allocation actions, allocating the NOMA-resources of a PRB, need to be taken without new state information. If, e.g., UE  $k$  is allowed to use the first NOMA-resource and is able to transmit all the packets inside its buffer, allocating the second NOMA resource to UE  $k$  does not make sense. However, the state containing information about the occupancy of all UEs' buffers is not updated and still assumes that UE  $k$  has packets to transmit. To combat this issue, we introduce a sequential decision making structure, shown in Fig. 2, which is inspired by [25]. The rows  $\mathbf{x}_m \in \mathbb{R}^{(K+1)}$ ,  $m = \{1, \dots, M\}$ , of the matrix  $\mathbf{X} \in \mathbb{R}^{M \times (K+1)}$  contain the probabilities of the  $K + 1$  actions to be chosen for NOMA resource  $m$ . So,  $\mathbf{X} = (\mathbf{x}_1 \dots \mathbf{x}_M)^T$  with  $\mathbf{x}_m = (x_{m,1} \dots x_{m,K+1})$  and  $x_{m,k} \in [0,1]$ . We initialize  $\mathbf{X} = \mathbf{0}$  and iteratively decide for allocation, replacing the  $m^{\text{th}}$  row of  $\mathbf{X}$  with the actor output  $\mathbf{x}_m$ . The updated  $\mathbf{X}$  is used to enable prediction of the changes applied to the state, e.g., which UE's buffers may have been emptied using already allocated NOMA-capabilities. After the actor decided for  $\mathbf{x}_M$ , the actions  $\mathbf{a} = (a_1 \dots a_M)$  are sampled from the probability distributions provided by  $\mathbf{x}_m$ .

##### B. Early Termination

The second issue we face in the uplink MC-NOMA scenario are sparse rewards. While the state is updated after each PRB is processed, the rewards are calculated after all PRBs are processed. Thus,  $M \cdot N$  actions have been carried out. Not getting rewards immediately on single actions, but on a group

of actions, prolongs training. We combat sparse rewards by introducing early termination for training. Especially at the beginning of training, the actor decides for suboptimal actions, leading to full UE buffers and, therefore, to situations without a chance for choosing beneficial actions. We assume that a sufficiently trained agent is able to avoid such ill-conditioned situations. Therefore, we terminate a training episode early if the training reward  $W_{\text{train}} < W_{\text{cap}}$ , where  $W_{\text{cap}} \in \mathbb{R}$  is an empirically determined value. Since the reward penalizes packets that exceeded their PDB (punishment), the reward is solely negative  $W_{t+1} \leq 0$  (penalty-only) and it is  $W_{\text{train}} < 0$ . With early termination, we save computation time in ill-conditioned situations and are able to combat the prolonged training due to sparse rewards.

##### C. Traffic-based Masking

To avoid the actor choosing UEs with empty buffers, we introduce traffic-based masking. By defining a Boolean mask  $\mathbf{h} \in \{0,1\}^{(K+1)}$ , indicating whether a UE has packets in its buffer or not, the probability vector  $\mathbf{x}_m$  is updated by an elementwise multiplication  $\mathbf{x}_m \leftarrow \mathbf{x}_m \odot \mathbf{h}$ . We apply normalization to ensure that after masking  $\mathbf{x}_m$  is a probability distribution.

##### D. Architecture of Actor and Critic Networks

In [9], we introduced encoder neural networks (ENNs) for effective state space compression. The state  $s_k$  of every UE  $k$ , containing the size and ages of packets residing in a UE's buffer as well as the UE's CQI  $c_k$  and QI  $q_k$ , is compressed to a vector  $s'_k \in \mathbb{R}^3$  using ENNs. To remedy the issue of learning a bias towards a specific UE, we randomly shuffle the order of UEs by  $\mathbf{s} \leftarrow \mathbf{P}_{\text{rand}} \mathbf{s}$ , where  $\mathbf{P}_{\text{rand}}$  is a random permutation matrix. Figure 3 shows the setup of the actor network. The compressed and randomly permuted UE information  $\mathbf{s}_k, k \in \mathcal{K}$ ,  $\mathbf{X}$  as well as  $n_{\text{PRB}} \in \mathcal{N}$ , denoting the current PRB to allocate, are fed to a DNN with a softmax-function at its output layer. Using  $\mathbf{P}_{\text{rand}}^{-1}$ , the order of UEs is restored and the masking is applied. To ensure exploration during training, parameter space noise [26] is applied to the weights and biases of the output layer. To avoid overfitting of the DNN, we use dropout layers.

We furthermore use age capping for handling packets that exceeded their PDB [9]. The critic's architecture is similar in terms of input information processing, however, the DNN's output layer only has a single neuron, since solely the expected return for executing actions  $\mathbf{X}$  given state  $S_t = \{s_1, \dots, s_K, n_{\text{PRB}}\}$  needs to be estimated.

## V. RESULTS

##### A. Training, Evaluation and Test Setup

We employ the environment described in Sec. II. For a small environment (SE) with  $K = 20$  UEs,  $N = 10$  PRBs and  $W_{\text{cap}} = -80\,000$ , we show the success of masking. For a large environment (LE) with  $K = 32$  UEs,  $N = 25$  PRBs and  $W_{\text{cap}} = -150\,000$ , we show the necessity of dropout layers. All setups have  $L = 8$  buffer slots per UE

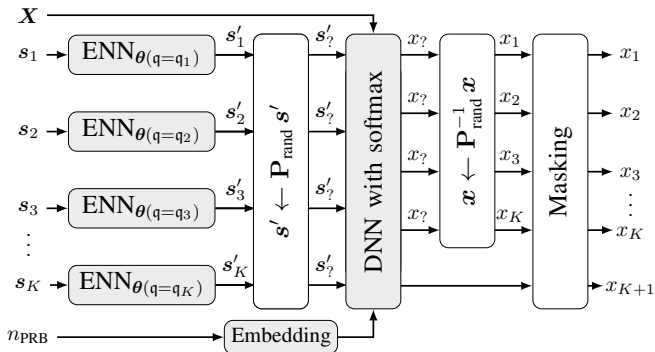


Fig. 3. Structure of the actor network, modified from [9]. ENNs, UE shuffling and PRB embedding is carried out as described in [9]. Learnable segments are highlighted by gray shading. Probabilities of taking actions  $x_k$  are returned, where  $x_{k+1}$  denotes the action of leaving the resource empty. The probability matrix of past action probabilities  $\mathbf{X} = [x_1 \dots x_M]$  is given as input.

TABLE I  
PARAMETERS OF THE DNNs FOR THE SE AND LE

	input width	hidden width	output width	hidden layers	activation functions
SE	105	603	21	3	ReLU
LE	165	963	33	3	ReLU

and  $M = 2$ . The parameters of the embedding and ENNs are given by [9, Tab. I]. Depending on SE and LE, the DNN architecture differs according to Tab. I. The different agents we train and test are summarized in Tab. II. We benchmark against the *NOMA uplink proportional fair channel aware* (NPFCA) agent, provided by [15]. Given the infinite set of environment realizations  $\mathfrak{S}$ , we choose  $|\mathfrak{S}_{\text{eval}}| = 4$  realizations for the evaluation and  $|\mathfrak{S}_{\text{test}}| = 100$  for the test set, with  $\mathfrak{S}_{\text{eval}} \cap \mathfrak{S}_{\text{test}} = \emptyset$ . Furthermore,  $\mathfrak{S}_{\text{train}} \subset \mathfrak{S} \setminus (\mathfrak{S}_{\text{eval}} \cup \mathfrak{S}_{\text{test}})$  denotes the training set, where the training realizations are sampled from. During training, an environment is stopped after  $t_{\text{stop}}$ , either after  $T_{\text{max}} = 600$  time steps are executed or by early termination, i.e.,  $t_{\text{stop}} \leq 600$  time steps. After five training episodes, the agent is evaluated. To save computing time for LE, agents are only evaluated if they achieve a training reward better than the mean reward of the benchmark agent, which is  $W_{\text{eval}} = -1552$ . The evaluation and tests are limited to 65536 allocation steps, resulting in  $T_{\text{max,test}} = 6553$  time steps for SE and  $T_{\text{max,test}} = 2621$  time steps for LE.

### B. Training and Evaluation Results

In Fig. 4, the training performance of the different agents is shown. If an agent is evaluated, its mean evaluation reward over  $\mathfrak{S}_{\text{eval}}$  is plotted as a function of the training episode.

TABLE II  
OVERVIEW OF THE AGENTS

	Environment	$K$	$N$	Masking	Dropout
S-def	SE	20	10	—	—
S-mask	SE	20	10	✓	—
L-mask	LE	32	25	✓	—
L-drop	LE	32	25	✓	✓




Fig. 4. Mean evaluation reward of the agents introduced in Tab. II.

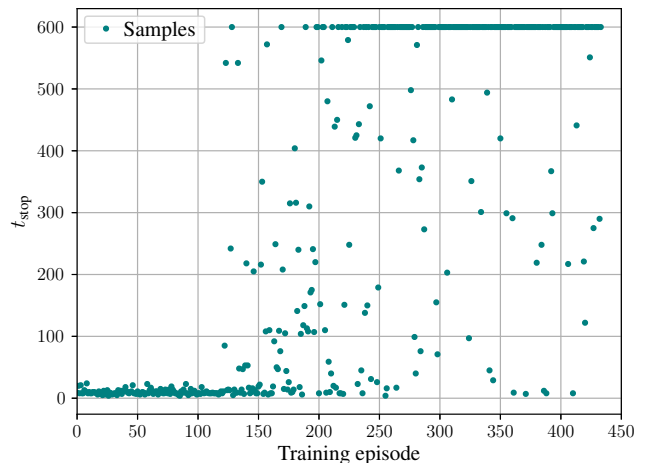


Fig. 5. Maximum environment runtime  $t_{\text{stop}}$  of the training environments of the L-drop agent when early termination is applied.

The S-def agent has only a small window where its training performance is good enough to get evaluated. With increasing training episodes, the S-mask agent is steadily improving. We conclude that masking stabilizes and improves training significantly. However, for the LE, the L-mask agent returns volatile evaluation rewards and its performance collapses after approximately 650 training episodes. We assume that this is attributed to overfitting of the DNN. Adding dropout layers with an outage probability of  $p = 0.2$  stabilizes training, shown by the L-drop agent. Due to sufficient evaluation performance and high computation effort, training of the L-drop agent was stopped after 440 training episodes. Figure 5 shows the impact of early termination on the training of the L-drop agent. With increasing training, the runtime  $t_{\text{stop}}$  of the environment realizations increases as well. Especially in the beginning of training, early termination significantly speeds up training. The L-drop agent gets a performance boost at approximately 200 episodes, which, compared to Fig. 4, is the start of evaluation of the L-drop agent.

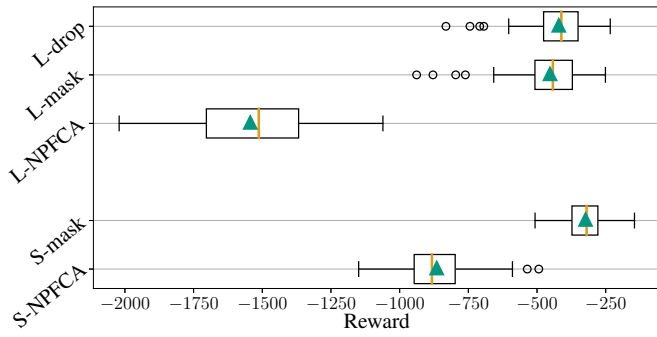


Fig. 6. Test rewards of the agents introduced in Tab. II and the benchmark agents. The lower two agents are specialized to the SE, the upper three agents to the LE.

### C. Test Results

The improvement of our proposed agents compared to the benchmark agent is shown in Fig. 6. The agents are tested using the same 100 environment realizations and we plot the obtained rewards. Green triangles indicate the mean evaluation reward, vertical orange lines inside the box, limited by the lower and upper quartile, depict the median reward over all environments. Open circles are outliers. Our agents outperform the benchmark agent by getting merely 37% and 27% of the benchmarks penalty, for SE and LE respectively.

### D. Discussion

In [9], we have shown that incorporating ENNs compressing BSI as well as age capping enable the design of an agent that outperforms benchmark agents without BSI. For a detailed discussion of the benefits of ENNs and age capping, we refer the interested reader to [9]. Due to the large action space of the MC-NOMA system, the Q-learning approach from [9] is not feasible anymore. By changing the DRL technique of [9] to an actor-critic approach and adding traffic-based masking, we can extend our previous work of [9] and design an agent for resource allocation in an uplink MC-NOMA system using BSI. The proposed scheme assumes to train and test on a fixed number of UEs. The generalization of the agent to a variable number of UEs, e.g., training for  $K$  UEs and testing for  $K_{\text{test}} < K$ , is ongoing.

## VI. CONCLUSION

In this work, we have proposed a centralized DRL agent for an uplink MC-NOMA resource allocation problem using BSI. We proposed to use a DDPG-based approach with a stochastic policy and combine it with the methods introduced in [9]. To enable the decision for multiple actions per PRB, we proposed feedback of the probability matrix of past action to the actor. To speed up training, we introduced early termination, which interrupts training in ill-conditioned situations. Furthermore, we have shown that a traffic-based masking of actions as well as dropout layers stabilize training of the agent. For  $K = 20$  and  $K = 32$  UEs, we significantly outperform the benchmark agent. Thus, we enabled the use of BSI for an uplink MC-NOMA resource allocation problem, which improved the performance of the scheduler.

## REFERENCES

- [1] Y. Yuan *et al.*, "Noma for next-generation massive IoT: Performance potential and technology directions," *IEEE Commun. Mag.*, vol. 59, no. 7, pp. 115–121, 2021.
- [2] W. Jiang, B. Han, M. A. Habibi, and H. D. Schotten, "The road towards 6G: A comprehensive survey," *IEEE Open J. Commun. Soc.*, vol. 2, pp. 334–366, 2021.
- [3] H. Tabassum, M. S. Ali, E. Hossain, M. J. Hossain, and D. I. Kim, "Uplink vs. downlink NOMA in cellular networks: Challenges and research directions," in *Proc. VTC Spring*, Sydney, Australia, 2017.
- [4] M.-R. Hojiej, J. Farah, C. A. Nour, and C. Douillard, "Resource allocation in downlink non-orthogonal multiple access (NOMA) for future radio access," in *Proc. VTC Spring*, 2015.
- [5] F. Shams, G. Bacci, and M. Luise, "A survey on resource allocation techniques in OFDM(A) networks," *Comput. Netw.*, vol. 65, pp. 129–150, June 2014.
- [6] S. Wang, T. Lv, W. Ni, N. C. Beaulieu, and Y. J. Guo, "Joint resource management for MC-NOMA: A deep reinforcement learning approach," *IEEE Trans. Wirel. Commun.*, vol. 20, no. 9, pp. 5672–5688, 2021.
- [7] J. Wang, C. Xu, Y. Huangfu, R. Li, Y. Ge, and J. Wang, "Deep reinforcement learning for scheduling in cellular networks," in *Proc. IEEE WCSP*, Xi'an, China, Oct. 2019.
- [8] M. P. Mota, A. Valcarce, J.-M. Gorce, and J. Hoydis, "The emergence of wireless MAC protocols with multi-agent reinforcement learning," in *IEEE GLOBECOM Workshops*, 2021.
- [9] E.-M. Bansbach, V. Eliachevitch, and L. Schmalen, "Deep reinforcement learning for wireless resource allocation using buffer state information," in *Proc. IEEE GLOBECOM*, Madrid, Spain, Dec. 2021.
- [10] M. V. da Silva, R. D. Souza, H. Alves, and T. Abrão, "A NOMA-based Q-learning random access method for machine type communications," *IEEE Wireless Commun. Lett.*, vol. 9, no. 10, pp. 1720–1724, 2020.
- [11] T. P. Lillicrap *et al.*, "Continuous control with deep reinforcement learning," in *Proc. ICLR*, San Juan, Puerto Rico, May 2016.
- [12] M. He, Y. Li, X. Wang, and Z. Liu, "NOMA resource allocation method in IoV based on prioritized DQN-DDPG network," *EURASIP J. Adv. Signal Process.*, vol. 120, 2021.
- [13] S. Wang, X. Wang, Y. Zhang, and Y. Xu, "Resource allocation in multi-cell NOMA systems with multi-agent deep reinforcement learning," in *Proc. WCNC*, 2021.
- [14] Y.-H. Xu, C.-C. Yang, M. Hua, and W. Zhou, "Deep deterministic policy gradient (DDPG)-based resource allocation scheme for NOMA vehicular communications," *IEEE Access*, vol. 8, pp. 18 797–18 807, 2020.
- [15] "Wireless-suite," Nokia, 2021, (accessed on: 11.04.2022). [Online]. Available: <https://github.com/nokia/wireless-suite>
- [16] A. Valcarce, "The TimeFreqResourceAllocation-v0 environment," 2020, (accessed on: 11.04.2022). [Online]. Available: <https://github.com/nokia/wireless-suite/blob/master/wireless/doc/TimeFreqResourceAllocation-v0>
- [17] S. Chandra and A. K. Bharti, "Speed distribution curves for pedestrians during walking and crossing," *Procedia - Social and Behavioral Sciences*, vol. 104, pp. 660–667, 2013.
- [18] J. Proakis and M. Salehi, *Digital Communications 5th Edition*. McGraw Hill, 2007.
- [19] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*, 2nd ed. The MIT Press, 2018.
- [20] C. J. C. H. Watkins and P. Dayan, "Q-learning," *Mach Learn*, vol. 8, no. 3, pp. 279–292, 1992.
- [21] R. S. Sutton, D. McAllester, S. Singh, and Y. Mansour, "Policy gradient methods for reinforcement learning with function approximation," in *Adv. Neural Inf. Process. Syst.*, vol. 12. MIT Press, 1999.
- [22] D. Silver, G. Lever, N. Heess, T. Degris, W. Daan, and M. Reidmiller, "Deterministic policy gradient algorithms," in *Proc. ICML*, Beijing, China, June 2014.
- [23] I. Grondman *et al.*, "A survey of actor-critic reinforcement learning: Standard and natural policy gradients," *IEEE Trans. Syst. Man Cybern., Part C*, vol. 42, no. 6, pp. 1291–1307, 2012.
- [24] V. Konda and J. Tsitsiklis, "Actor-critic algorithms," in *Adv. Neural Inf. Process. Syst.*, vol. 12. MIT Press, 1999.
- [25] Y. Zhang, Q. H. Vuong, K. Song, X.-Y. Gong, and K. W. Ross, "Efficient entropy for policy gradient with multi-dimensional action space," in *Proc. ICLR*, Vancouver, Canada, May 2018.
- [26] M. Plappert *et al.*, "Parameter space noise for exploration," in *Proc. ICLR*, Vancouver, Canada, May 2018.