Towards Joint Learning of Optimal MAC Signaling and Wireless Channel Access

Alvaro Valcarce, *Senior Member, IEEE,* and Jakob Hoydis, *Senior Member, IEEE*

Abstract—**Communication protocols are the languages used by network nodes. Before a user equipment (UE) exchanges data with a base station (BS), it must first negotiate the conditions and parameters for that transmission. This negotiation is supported by signaling messages at all layers of the protocol stack. Each year, the telecoms industry defines and standardizes these messages, which are designed by humans during lengthy technical (and often political) debates. Following this standardization effort, the development phase begins, wherein the industry interprets and implements the resulting standards. But is this massive development undertaking the only way to implement a given protocol? We address the question of whether radios can learn a pre-given target protocol as an intermediate step towards evolving their own. Furthermore, we train cellular radios to emerge a channel access policy that performs optimally under the constraints of the target protocol. We show that multi-agent reinforcement learning (MARL) and learning-to-communicate (L2C) techniques achieve this goal with gains over expert systems. Finally, we provide insight into the transferability of these results to scenarios never seen during training.**

Index Terms—**communication system signaling, learning systems, mobile communication**

Fig. 1. MAC protocol constituent parts. This paper trains UEs to jointly learn the policies highlighted as dashed blocks.

I. INTRODUCTION

CURRENT medium access control (MAC) protocols obey fixed rules given by industrial standards (e.g., [1]). These standards are designed by humans with competing commercial interests. Despite the standardization process' high costs, the ensuing protocols are often ambiguous and not necessarily optimal for a given task. This ambiguity increases the costs of testing, validation and implementation, specially in cross-vendor systems like cellular networks. In this paper, we study if there might be an alternative approach to MAC protocol implementation based on reinforcement learning.

To implement a MAC protocol for a wireless application, vendors must respect the standardized
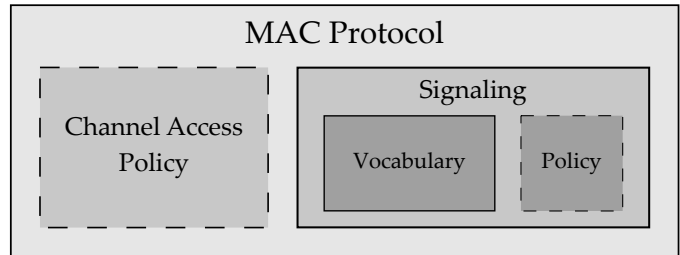
A. Valcarce and J. Hoydis are with Nokia Bell-Labs France, Route de Villejust, 91620 Nozay, France. Email: {alvaro.valcarce_rial, jakob.hoydis}@nokia-bell-labs.com

signaling (i.e., the structure of the protocol data units (PDUs), the control procedures, etc) and the channel access policy (listen before talk (LBT), FDMA, etc). These are the main building blocks of a MAC protocol (see Fig 1) and their implementation ultimately defines the performance that users will experience. The signaling defines what control information is available to the radio nodes and when it is available. In doing so, it constraints the actions that a channel access policy can take and puts an upper bound to its attainable performance.

As wireless technologies evolve and new radio environments emerge (e.g., industrial IoT, beyond 6 GHz, etc), new protocols are needed. The optimal MAC protocol may not be strictly contention-based or coordinated as shown in Table I. For instance, 5GNR supports grant-free scheduling by means of a Configured Grant to support ultra reliable low latency communications (URLLC). But are this acces policy and its associated signaling jointly optimal? We are interested in the question of which MAC protocol is optimal for a given use case and whether it can be learned through experience.

A. Related work

Current research on the problem of emerging MAC protocols with machine learning focuses mainly on new channel access policies. Within this body of research, contention-based policies domi-

TABLE I
Channel access schemes, MAC protocols and the technologies that use them

		Channel Access Schemes					
			frequency division multiple access (FDMA)				
		TDMA	Time-invariant	Time-variant		CDMA	Spatial
				Non-OFDMA	OFDMA		
MAC protocols	Contention based	Aloha CSMA Cognitive Radio	N.A.	Cognitive Radio			
	Coordinated	GSM Token Ring Bluetooth	FM/AM radio DVB-T	POTS	LTE 5G New Radio (5GNR)	3G GPS	Satcoms MIMO

nate (see [2], [3], [4] or [5]), although work on coordinated protocols also exists (e.g., [6], [7], [8]). Other approaches such as [9] propose learning to enable/disable existing MAC features.

Given a channel access scheme (e.g., time division multiple access (TDMA)), [2] asks whether an agent can learn a channel access policy in an environment where other agents use heterogeneous policies (e.g., q-ALOHA, etc). Whereas this is interesting, it solely focuses on contention-based access and it ignores the signaling needed to support it. Instead, we focus on the more ambitious goal of jointly emerging a channel access policy and its associated signaling.

Dynamic spectrum sharing is a similar problem, where high performance has recently been achieved (see [10], [11]) by subdividing the task into the smaller sub-problems of channel selection, admission control and scheduling. However, these studies focus exclusively on maximizing channel-access efficiency under the constraints of a pre-given signaling. None focuses on jointly optimizing the control signaling and channel access policy.

The field of Emergent Communication has been growing since 2016 [12]. Advances in deep learning have enabled the application of MARL to the study of how languages emerge and to teaching natural languages to artificial agents (see [13] or [14]). Due to the multi-agent nature of cellular networks and to the control-plane/user-plane traffic separation, these techniques generalize well to the development of machine-type languages (i.e., signaling). In cellular systems, we interpret the MAC signaling as the language spoken by UEs and the BS to coordinate while pursuing the goal of delivering traffic across a network.

B. Contribution

For the reasons mentioned above, we believe that machine-learned MAC protocols have the potential to outperform their human-built counterparts in certain scenarios. This idea has been recently used for emerging new digital modulations (see, e.g., [15]). Protocols built this way are freed from human intuitions and may be able to optimize control-plane traffic and channel access in yet unseen ways.

The first step towards this goal is to train an intelligent software agent to learn an existing MAC protocol. Our agents are trained tabula rasa with no previous knowledge or logic about the target protocol. We show that this is possible in a simplified wireless scenario with MARL, self-play and tabular learning techniques, and lay the groundwork for scaling this further with deep learning. In addition, we measure the influence of signaling on the achievable channel access performance. Finally, we present results on the extension of the learned signaling and channel access policy to other scenarios.

The rest of this paper is organized as follows. Section II formalizes the concepts of channel access, protocol and signaling. It then formulates the joint learning problem and defines the research target. Section III describes the multi-agent learning framework. Section IV illustrates the achieved performance, provides an example of the learned signaling and discusses the transferability of these results. A final discussion is provided in section V.

II. Problem definition

A. Definitions

We distinguish the concepts of *channel access scheme* and *MAC protocol* as follows:
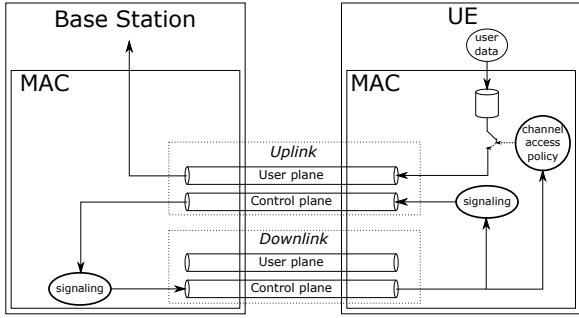
Fig. 2. MAC protocol abstraction with UL data traffic only.

Fig. 3. Left: Venn diagram of all possible signaling messages $\mathcal{M} = \mathcal{M}_{DL} \cup \mathcal{M}_{UL}$ of size (l, m) bits. The set M^* of messages for the optimal signaling S^*, as well as the set M^k of messages for the k^{th} arbitrary signaling S^k are also shown. Right: An optimum signaling S^* and the k^{th} arbitrary signaling S^k within the abstract space $\mathcal{S}$ of all possible signaling (i.e., the signaling corpus).

- **Channel access scheme**: Method used by multiple radios to share a communication channel (e.g., a wireless medium). The channel access scheme is implemented and constrained by the physical layer (PHY). Sample channel access schemes are frequency division multiplexing (FDM), time division multiplexing (TDM), etc.
- **MAC protocol**: Combination of a *channel access policy* and a *signaling* with which a channel access scheme is used. Sample channel access policies are LBT, dynamic scheduling, etc. Signaling is the vocabulary and rules (i.e., signaling policy) used by radios to coordinate and is described by the PDU structure, the sub-headers, etc. The channel access policy decides when to send data through the user-plane pipe, and the signaling rules decide when to send what through the control-plane pipe (see Fig. 2).

Table I illustrates these definitions with examples of technologies that use these mechanisms.

B. Target MAC signaling

A complete commercial MAC layer provides numerous services to the upper layers (channel access, multiplexing of service data units (SDUs), hybrid automatic repeat request (HARQ), etc). Our goal is to replace the MAC layer in a mobile UE by a learning agent that can perform all these functions and their associated signaling. However for simplicity, this paper targets a leaner MAC layer with two main functions: wireless channel access, and automatic repeat request (ARQ). Let this agent be denoted as the MAC *learner*. This differs from an *expert* MAC implemented through a traditional design-build-test-and-validate approach. Several learners are then concurrently trained in a mobile cell to jointly learn
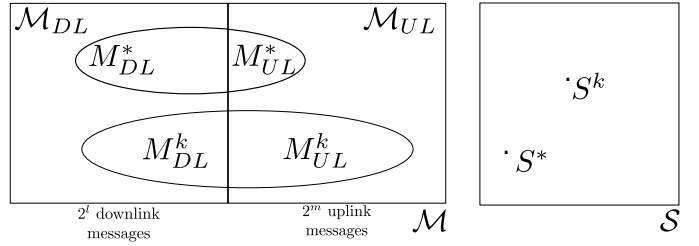
a *channel access policy* and the MAC *signaling* needed to coordinate channel access with the BS. The BS uses an expert MAC that is not learned.

Let $\mathcal{S}$ be the set of all possible MAC signaling that UEs and a BS may ever use to communicate (see Fig. 3). We formalize a signaling as a vocabulary with downlink (DL) and UL messages, plus mappings from observations to these messages. Since different signalings are possible, let us denote the k^{th} signaling $S^k \in \mathcal{S}$ as:

$$S^k = [M^k_{DL}, M^k_{UL},$$
$$O^{BS} \rightarrow M^k_{DL}, \quad (1)$$
$$O^{UE} \rightarrow M^k_{UL}]$$

where $M^k_{DL} \subseteq \mathcal{M}_{DL}$ and $M^k_{UL} \subseteq \mathcal{M}_{UL}$ are the sets of DL and UL messages of signaling S^k. O^{BS} and O^{UE} are generic observations obtained at the BS and the UE respectively, and can include internal states, local measurements, etc. In other words, a signaling defines the messages that can be exchanged and the rules under which they can be exchanged. These rules give hence meaning to the messages.

In the above definitions, $|M^k_{DL}|$ and $|M^k_{UL}|$ are the sizes of the DL and UL signaling vocabularies, which implicitly define the amount of control data a single message can carry, i.e., the richness of the control vocabulary. Messages from non-compositional protocols with larger vocabularies can therefore feed more control information to the channel access policy. Although this comes at the expense of a higher signaling overhead, the richer context available to the radio nodes can enable more sophisticated algorithms for channel access. The size of the signaling vocabulary is hence an important hyper-parameter in emergent protocols

to balance the trade-off between channel access performance and signaling cost.

The mappings from observations to messages define when to send what. A MAC signaling policy π_S describes one possible way of implementing this mapping. This *signaling policy* shall not be confused with the *channel access policy*, which describes when and how to transmit data.

The BS is an expert system with full knowledge of a standard MAC signaling $S^{BS} \in \mathcal{S}$. Our first objective is to enable the UEs to communicate with the BS by learning to understand and use its signaling. Out of the many signaling UEs could learn, S^{BS} is the learning target. Note that S^{BS} is not necessarily the optimal signaling $S^* \in \mathcal{S}$, which depends on a chosen metric, such as throughput, latency, etc.

The ideas presented here are generalizable to vocabularies of any size. For simplicity, we have reduced the size of the target signaling vocabulary to the minimum number of messages needed to support both uncoordinated and coordinated MAC protocols. In this paper, S^{BS} has the following DL messages that the MAC learners need to interpret:

- scheduling grants (SGs)
- acknowledgments (ACKs).

Similarly, the UL messages that the MAC learners need to learn to use are:

- scheduling requests (SRs).

Other messages such as buffer status reports can be added to accommodate for larger vocabularies.

C. Channel access policy

The experiments described in this paper use a TDMA scheme, which was chosen for simplicity. The learning framework proposed here is nevertheless equally applicable to other channel access schemes, such as those listed in Table I.

Unlike for the MAC signaling, we impose no a-priori known policy with which to use the access scheme. This is to allow UEs to explore the full spectrum of channel access policies between contention-based and coordinated. For example, a channel access policy might follow a logic that ignores the available MAC signaling. Other channel access policies may take it into consideration and implement a coordinated access scheme based on a sequence of SRs, SGs and ACKs.

Hence, the second objective of this research is for the UEs to learn an UL channel access policy π_P that leverages the available signaling to perform optimally (according to some chosen channel access performance metric). The channel access policy is denoted with the subscript P to highlight that the MAC layer controls channel access by steering the underlying physical layer. The MAC layer commands the PHY through an application programming interface (API) and is therefore constrained by the services it offers. The MAC has no means of influencing the wireless shared channel without a PHY API (e.g., [16]). In our experiments, these services are limited to the in-order delivery of MAC PDUs through a packet erasure channel.

D. Channel model

All learners share the same UL frequency channel for transmitting their UL MAC PDUs and they access this shared UL data channel through their respective PHY layers. From the viewpoint of the MAC learner, the PHY is thus considered part of the shared channel. The channel accessed by each MAC learner is a *packet erasure channel*, where a transmitted PDU is lost with a certain block error rate (BLER). For simplicity, we assume that the BLER is the same for all data links and abstract away all other PHY features. This comes without loss of generality to the higher-layer protocol analysis.

Since we want to study the effects of the control signaling on the performance of the shared data channel, we assume that the UL and DL control channels are error free, costless and dedicated to each user without any contention or collisions.

E. BS signaling policy

Each time step, the BS receives zero or more SRs from the UEs. It then chooses one of the requesting UEs at random and a SG is sent in response. An exception occurs if the UE had made a successful data transmission concurrently with the SR. In this case, instead of an SG, an ACK is sent to the UE and another UE is then scheduled at random. This is because only one UE can be scheduled each time slot. More complex scheduling algorithms could be used here, but random scheduling has been chosen for simplicity.
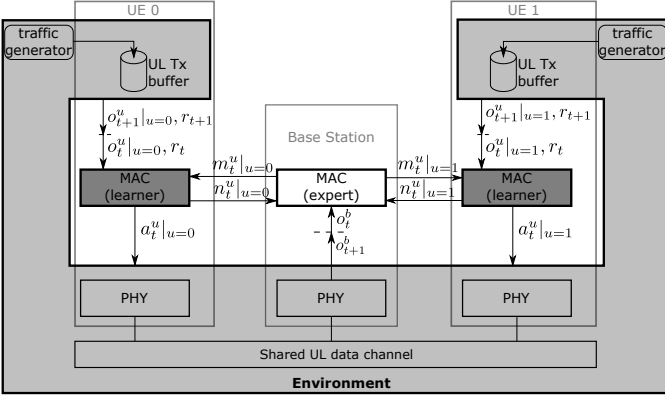
Fig. 4. System model with one BS and two UEs

F. Multi-Agent Reinforcement Learning formulation

For any given UE, the sequential decision-making nature of channel access lets us model it as a Markov Decision Process (MDP), which can then be solved using the tools of reinforcement learning (RL). Using RL to train multiple simultaneous learners (i.e., UEs) constitutes what is known as a MARL setup. If the observations received by each learner differ from those of other learners, the problem becomes a partially observable Markov decision process (POMDP).

We have formulated this POMDP as a cooperative Markov game (see [17]), where all learners receive exactly the same reward from the environment. Learners are trained cooperatively to maximize the sum R of rewards:

$$R \triangleq \sum_{t=0}^{t_{max}} r_t \qquad (2)$$

where t_{max} is the maximum number of time steps in an episode. This design decision reflects the objective of optimizing the performance of the whole cell, rather than that of any individual UE. An alternative approach that delivers different rewards to different UEs depending on their radio conditions could perhaps yield higher network-wide performance. However, this would require the design of multiple reward functions, which is known to be difficult (see [18]), and left to future investigations.

1) Time dynamics: The MARL architecture is illustrated in Figure 4. U is the set of all MAC learners. Then, on each time step t, each MAC learner $u \in [0, |U|)$ invokes an action a_t^u on the environment, and it receives a reward r_{t+1} and an observation o_{t+1}^u. The actions of all learners are aggregated into a joint action vector $\mathbf{a}_t$. The environment then executes this joint action and delivers the same reward r_{t+1} plus independent observations to all learners. The BS also receives a scalar observation o_{t+1}^b following the execution of the joint learner action.

The environment studied in this paper demands that each MAC learner delivers a total of P UL MAC SDUs to the BS. We performed experiments with the following simple SDU traffic models:

- **Full buffer start**: The UL Tx buffer is filled with P SDUs at $t = 0$.
- **Empty fuffer start**: The UL Tx buffer is empty at $t = 0$. Then, it is filled with probability 0.5 with one new SDU each time step until a maximum of P SDUs have been generated.

The learners must also indicate awareness that the SDUs have been successfully delivered by removing them from their UL transmit buffer. An episode ends when the P SDUs of each and all of the $|U|$ MAC learners have successfully reached the BS and the learners have removed them from their buffers.

2) Observation space: Each time step t, the environment delivers a scalar observation $o_t^u \in O^{\mathcal{U}} = [0, L]$ to each learner u describing the number of SDUs that remain in the learner's UL transmit buffer. Here, L is the transmit buffer capacity and all SDUs are presumed to be of the same size. For example, the environment observation $o_t^u|_{u=5,t=2} = 3$ indicates to learner 5 that, at time $t = 2$, three SDUs are yet to be transmitted.

Similarly, at each time step the BS receives a scalar observation $o_t^b \in O^{\mathcal{B}} = [0, |U| + 1]$ from the environment. This observation can take the following meanings:

- $o_t^b = 0$: UL channel is idle
- $o_t^b \in [1, |U|]$: Collision-free UL transmission received from UE $o_t^b - 1$
- $o_t^b = |U| + 1$: Collision in UL channel.

For example, the environment observation $o_t^b|_{t=4} = 3$ indicates that, at time $t = 4$, the BS successfully received a MAC PDU from UE 2.

3) Channel access action space: Learners follow a channel access policy by executing actions $a_t^u \in \mathcal{A}_P$ every time step. These actions have an effect on the environment by steering the physical layer and are $\mathcal{A}_P = \{0, 1, 2\}$, which are interpreted by the environment as follows:

- $a_t^u = 0$: Do nothing

- $a_t^u = 1$: Transmit the oldest SDU in the buffer. Invoking this action when the UL buffer is empty is equivalent to invoking $a_t^u = 0$.
- $a_t^u = 2$: Delete[1] the oldest SDU in the buffer.

The channel access actions from all learners are aggregated into a joint action vector $a_t = [a_t^0, a_t^1, ..., a_t^{|U|-1}]$, which is then executed on the environment at once. For example, invoking action $a_2 = [1, 2, 0]$ on the environment indicates that, at time $t = 2$, UE 0 attempts an SDU transmission while UE 1 deletes a SDU from its buffer, and UE 2 remains idle.

4) Uplink signaling action space: Following the approach introduced in [12], in each time step, MAC learners can also select an UL signaling action $n_t^u \in \mathcal{A}_S$. This action maps to a message received by the BS MAC, and it exerts no direct effect onto the environment. These messages are thus transmitted through a dedicated UL control channel that is separate from the shared UL data channel modeled by the environment (see Fig. 2). The UL signaling actions available to the MAC learners are $\mathcal{M}_{UL} = \{0, 1\}$, which are interpreted by the BS MAC expert as follows:

- $n_t^u = 0$: Null uplink message
- $n_t^u = 1$: Scheduling Request.

5) Downlink messages: Finally, the BS MAC expert can, each time step, select a DL signaling action m_t^u towards each MAC learner. These messages have no direct effect on the environment and are transmitted through a dedicated DL control channel. The actions available to the BS MAC expert are $\mathcal{M}_{DL} = \{0, 1, 2\}$ and have the following meanings:

- $m_t^u = 0$: Null downlink message
- $m_t^u = 1$: Scheduling Grant for next time step
- $m_t^u = 2$: ACK corresponding to the last SDU received in the uplink.

Note that the MAC learners are unaware of the meanings of these messages. They must learn them from experience. Larger sizes of the DL vocabulary $\mathcal{M}_{DL}$ would let the BS communicate more complex schedules to the UEs (e.g., by granting radio resources further into the future). This would increase training duration significantly and is therefore out of the scope of this paper but left for future investigation.

---

[1]A third action space with buffer-management actions could have also been defined. For simplicity, we only use this buffer related action and include it in $\mathcal{A}_P$

6) Learners' memory: We endow each MAC learner with an internal memory $h_t^u \in \mathcal{H}^{\mathcal{U}}$ to store the past history of observations, PHY and UL signaling actions, as well as the DL messages received from the BS. N denotes the size, in number of past time steps, of this internal memory, which at time step t takes the form:

$$h_t^u = [m_{t-N}^u, a_{t-N}^u, n_{t-N}^u, o_{t-N}^u, ..., \\ m_{t-1}^u, a_{t-1}^u, n_{t-1}^u, o_{t-1}^u]. \quad (3)$$

For example, if $N = 1$, the internal memory h_t at time t contains only the observation, actions, and messages from time $t - 1$. The motivation for this memory is the need to disambiguate instantaneous observations that may seem equal to a given MAC learner, but emanate from the un-observed actions concurrently taken by other learners. In short, memory addresses the problem of partial observability by the learners. This is loosely based on the idea of *fingerprinting* introduced in [19].

The current state of the memory h_t^u and the current observation o_t^u constitute the two main inputs considered by the learners during action selection: $(a_t^u, n_t^u) = f(o_t^u, h_t^u)$. For example, a learner u with current memory state $h_t^u = [1, 0, 0, 1, 0, 0, 1, 1]$ and current observation $o_t^u = 1$ describes a situation where the learner received a SG in response to a previous SR due to a non-empty Tx buffer. The learner's current policy will then take this context into consideration for deciding the next channel access action and UL signaling message (if any).

7) Reward function: The environment delivers a reward of $r_t = -1$ on all time steps. This motivates MAC learners to finish the episode as quickly as possible (i.e., in the smallest number of time steps). The conditions for the termination of an episode were described in Section II-F1.

III. METHODS

A. Tabular Q-learning

Given the definition of reward r_t of section II-F1, the return G_t at time t is defined as (see [20]):

$$G_t \triangleq \sum_{k=0}^{\infty} \gamma^k r_{t+k+1} \quad (4)$$

where $\gamma \in [0, 1]$ is a discount factor. The action-value function, under an arbitrary policy π, for a given observation-action pair (o_t, a_t) is then defined
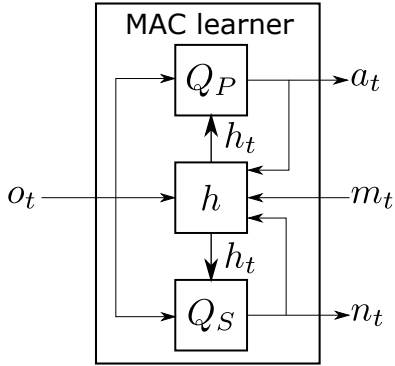
Fig. 5. Internal structure of a MAC learner

as the expected return when that observation is made and the action taken:

$$Q_\pi(o_t, a_t) = E_\pi[G_t | a_t, o_t]. \tag{5}$$

Q-learning [21] is a well known technique for finding the optimal action-value function, regardless of the policy being followed. It stores the value function $Q(o, a)$ in two-dimensional tables, whose cells are updated iteratively with:

$$Q(o_t, a_t) \leftarrow Q(o_t, a_t) + \\ \alpha[r_{t+1} + \gamma \max_a Q(o_{t+1}, a) - Q(o_t, a_t)]. \tag{6}$$

Independently of how the table is initialized, Q will converge to the optimal action-value function as long as sufficient exploration is provided.

Independent Q-learners are known to violate the Q-learning convergence requirements (see [22]). However, in practice, good results have been obtained when combining it with memory-capable agents and decentralized training. These two techniques are essential to address the non-stationarity of multi-agent Q-learning problems and have been successfully applied in [12] on an L2C setting. Our MAC learner architecture and training procedure are a tabular adaptation of the RIAL architecture described in [12], whose success is one of the main motivations for using Q-learning here.

Each MAC learner is composed of two Q tables (see Fig. 5), which contain action value estimates Q_P for the physical-layer actions and action value estimates Q_S for the signaling actions. The reduced size of the problem described in Section II allows for the application of a traditional Q-learning approach based on Q tables (see [20]). This is less data hungry than deep learning approaches, i.e., it needs less samples to converge and is therefore

appropriate as a proof of concept. However, tabular Q-learning scales poorly and is certainly not suitable in larger problems (e.g., $|U| >> 2$, more SDUs, larger memories, etc.). In such cases, deep learning approaches are needed (the tables in Fig. 5 can be easily replaced by deep neural networks).

Then, every time step t, each MAC learner follows two different policies π_P and π_S to act:

$$\pi_P : O^{\mathcal{U}} \times \mathcal{H}^{\mathcal{U}} \longrightarrow \mathcal{A}_P \\ \pi_S : O^{\mathcal{U}} \times \mathcal{H}^{\mathcal{U}} \longrightarrow \mathcal{M}_{UL}. \tag{7}$$

Although $a_t \in \mathcal{A}_P$ and $n_t \in \mathcal{M}_{UL}$ actions are chosen independently, both policies are synchronous due to the conditioning of both Q tables on the same current observation and state of the internal memory:

$$a_t = \arg\max_a Q_P(a | o_t, h_t) \tag{8}$$

$$n_t = \arg\max_n Q_S(n | o_t, h_t). \tag{9}$$

B. Training procedure

MAC learners are trained using self-play [23], which is known to reach a Nash equilibrium of the game when it converges. This yields policies that can be copied to all UEs and used during deployment. Training is centralized, with one central copy of the Q_P and Q_S tables being updated regularly as experience from all MAC learners is collected. In our experiments, we update the tables after each time step, although other schedules are possible. Then, each time step, all learners choose their actions based on the same version of the trained tables (i.e., decentralized execution) combined with an ϵ-greedy policy. The exploration probability ϵ is annealed between training episodes with $\epsilon_e = max(\epsilon_{e-1} * F_\epsilon^e, 0.01)$, where $\epsilon_0 = 1$, F_ϵ is the exploration decay factor, and $e > 0$ is the training episode number. Asymptotically, ϵ-greedy guarantees that learners try all actions an infinite number of times. This ensures that $Q(o_t, a_t)$ converges to the optimal $Q^*(o_t, a_t)$. However, training is never run indefinitely, so the theoretical guarantee looses relevance and leaves the door open to alternative exploration methods such as e.g., optimistic initial action values (see [20]). This *centralized training with decentralized execution* approach has been chosen to address the non-stationarity pathology typical of setups with independent learners

[24]. Deploying the same policies to all learners reduces the variance of observations during learning and helps convergence. This training procedure can be performed in a server farm, where thousands of cross-vendor UEs (cabled and/or over-the-air (OTA)) would contribute with experiences to the training of a central UE MAC model. From a practical viewpoint, centralized training also spares UEs from executing costly training algorithms. This way, most of the computational workload is shifted towards a central server. This avoids impacting the UEs' battery performance and escapes constraints due to UE mobility. It is precisely this procedure which could replace the more traditional design-build-test-and-validate approach to MAC layer development in future protocol stacks.

Different BS models can also be incorporated to the training testbed to increase the environment variance. This can generalize the learned policies and improve performance when deployed in mobile networks different from the ones used during training. Along these lines, zero-shot coordination techniques (see [25]) could also help.

The system was trained for a fixed number N_{tr} of consecutive training episodes, followed by a fixed number N_{eval} of consecutive evaluation episodes. This sequence of $N_{tr} + N_{eval}$ episodes is called a training session. The Q tables were only updated during training episodes, but not during evaluation episodes. The ϵ-greedy policy with ϵ decay was followed only during training episodes. Evaluation episodes follow strictly the learned policy and are thus free of exploration variance. For each configuration (set of hyper-parameters), training sessions were repeated N_{rep} times with different random seeds. The convergence time grows with the problem size, specially with hyper-parameters P and $|U|$. Fig. 9 illustrates how even a small scenario with $|U| = 2$ UEs and $P = 2$ SDUs requires several million training episodes to converge.

IV. RESULTS

The training procedure described above was tested in simulation, and its performance R measured as the reward collected per episode (see (2)). The optimal hyper-parameters have been chosen via grid search in the following discretized sets:

- Discount factor: $[0, 0.5, 1]$
- Exploration decay factor: $0.99991, 0.999991, 0.9999991$

TABLE II
SCENARIO SETTINGS AND SIMULATION HYPER-PARAMETERS

Symbol	Name	Value
BLER	block error rate	$0, 10^{-4}, ..., 10^{-1}$
SB	start buffer	full, empty
P	number of MAC SDUs	$1, 2$
$|U|$	number of UEs	$1, 2$
t_{max}	maximum time steps per episode	$4, 8, 16, 32$
N_{tr}	number of training episodes	$2^{13}, 2^{14}, ...$
N_{eval}	number of evaluation episodes	128
N_{rep}	training session repetitions	$4, 8$
γ	discount factor	1
F_{ϵ}	exploration decay factor	$0.999991, 0.9999991$
α	learning rate	0.3
N	memory length of MAC learner	$1, 2, 3, 4$

- Learning rate: $[0.05, 0.06, 0.1, 0.2, 0.3, 0.4]$

Table II collects the best-performing parameters.

A. Expert baseline

Algorithm 1 Expert channel access policy

Require: o_t, m_{t-1}, a_{t-1}
  if $m_{t-1} = SG$ **and** $o_t > 0$ **and** $a_{t-1} \neq 1$ **then**
    Transmit oldest SDU in the Tx buffer ($a_t = 1$)
  else if $m_{t-1} = ACK$ **and** $o_t > 0$ **then**
    Delete oldest SDU in the Tx buffer ($a_t = 2$)
  else
    Do nothing ($a_t = 0$)
  end if

Algorithm 2 Expert signaling policy

Require: $o_t, m_{t-1},$
  if $o_t > 0$ **then**
    if $m_{t-1} = 0$ **or** ($o_t > 1$ **and** $m_{t-1} = ACK$) **then**
      Send an SR to the BS ($n_t = 1$)
    else
      Do nothing ($n_t = 0$)
    end if
  else
    Do nothing ($n_t = 0$)
  end if

For comparison purposes, the performance obtained by a population of expert (i.e., non-learner) UEs is also shown. The expert UE has complete knowledge about the BS signaling semantics and only transmits following the reception of an SG. Similarly, it only deletes an SDU from its buffer

following reception of an ACK. These UEs follow a fully coordinated channel access policy and do not profit from potential gains due to stochastic contention. Pseudo code for the channel access and signaling policies implemented by the expert UE is provided in the Algorithms 1 and 2, respectively.

B. Optimal policy

The optimal policy $\pi^* = (\pi_P^*, \pi_S^*)$ for a single MAC learner that needs to transmit a single SDU (i.e., $|U| = 1$ and $P = 1$) can be intuitively deducted and its performance modeled analytically. Considering that the total reward that can be collected in a given episode depends on t_{max}, there may be different optimal policies for different values of this parameter. Indeed, the optimal channel access policy for low t_{max} consists in transmitting an SDU at $t = 0$ and immediately removing it from the Tx buffer in the next time step. UEs using this policy ignore the ACKs because on average, waiting for the ACK before removing the SDU takes too long. Let this optimal policy be $\pi^{(1)}$ with expected performance:

$$R^{(1)} = b \cdot (2 - t_{max}) - 2 \qquad (10)$$

where b denotes the BLER. When the SDU transmission at $t = 1$ succeeds, this leads to a sum reward of -2. If the transmission fails, the UE will get, under policy $\pi^{(1)}$, the mininum reward of $-tmax$. The value of $tmax$ for which this policy is optimal depends hence on the BLER.

For higher t_{max}, the MAC learner is encouraged to wait for an ACK because the risk of ignoring it is too high due to the long episode duration. Let this optimal policy for $|U| = 1$ be $\pi^{(2)}$. One can easily derive the expected performance analytically as:

$$R^{(2)} = \begin{cases} -t_{max}, & t_{max} < 4 \\ -(b+3), & t_{max} = 4 \\ \\ (b-1)(3 + \sum_{i=4}^{t_{max}-1} ib^{i-3}) & t_{max} > 4 \\ -t_{max} \cdot b^{t_{max}-3} \end{cases} \qquad (11)$$

The expected optimum performance in this scenario ($|U| = 1$ and $P = 1$) is shown in Fig. 6 and is:

$$R^* = max(R^{(1)}, R^{(2)}). \qquad (12)$$

It is worth noting that the MAC learners' behavior is an artifact of the reward function. The
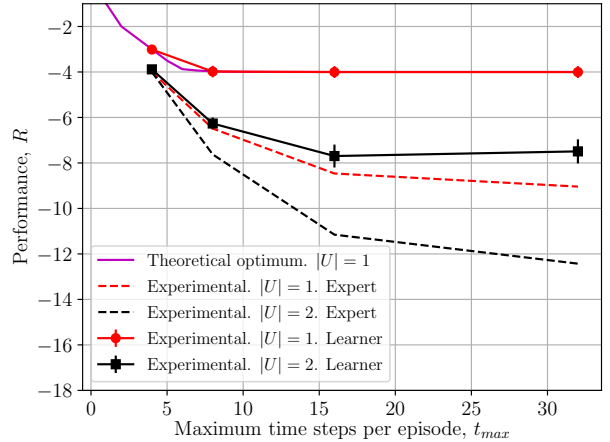

Fig. 6. Average performance over eight independent training sessions. The performance is evaluated over $N_{eval} = 128$ episodes with different seeds. The environment was configured with $P = 1$, $BLER = 0.5$ and start buffer full. MAC learners were trained with a learning rate of 0.05. Learners on the $|U| = 1$ environment were trained with an internal memory length of $N = 1$ for $N_{tr} = 2^{13} = 8192$ episodes, while learners on the $|U| = 2$ environment were trained with an internal memory length of $N = 3$ for $N_{tr} = 2^{18} = 262144$ episodes. No exploration was used (i.e., $\epsilon = 0$). The length of the error bars equal the standard error of the mean performance (calculated over the 8 independent training sessions).

reward described in section II-F7 encourages them to accomplish a task as fast as possible. Alternatively, one could provide a reward of 1 (or 1 for each message) as soon as the messages have been successfully delivered and a reward of -1 if the maximum duration has passed.

Fig. 6 shows how MAC learners use $\pi^{(1)}$ or $\pi^{(2)}$ depending on the maximum episode length. Both of these policies ignore SGs and are therefore useless in scenarios with more than $|U| = 1$ UE, where collisions impose the need for MAC learners to respect the scheduling allocation decided by the BS.

Fig. 6 also shows experimental results for the $|U| = 2$ scenario to illustrate how performance scales with increasing numbers of UEs. Similarly to the single-UE case, ACKs are largely ignored in the low t_{max} regime when $|U| = 2$ UEs are present. In this larger scenario, the vast size of the signaling solution space makes it difficult to find the optimal policy, much less a closed-form expression for its performance. Nevertheless, a measure of gain can still be obtained by comparing the performance of the learned policies against that of a known expert (see Fig. 6 for the gains above the expert).

C. BLER impact on MAC training

In the absence of block errors, MAC learners learn to ignore ACKs. On the other hand, unexpected PDU losses motivate MAC learners to interpret the DL ACKs before deleting transmitted SDUs from the Tx buffer. This is illustrated in the learned MAC signaling trace of Fig. 7, which shows the MAC learners removing SDUs from the UL Tx buffer at time steps immediately following the reception of an ACK (deletions at $t = 3$ and $t = 6$ for MAC 1, and at $t = 5$ for MAC 0).

Interestingly, MAC 0 also removed a SDU at $t = 7$ before it had time to process the received ACK. The reason is because, unlike in all other transmissions, transmission at $t = 6$ was preceded by a SG at $t = 5$, which guarantees it to be collision-free. In this scenario, BLER was so low that, on average, removing the SDU from the buffer before waiting for an ACK yields a higher average reward. This suggests that MAC protocols may perform better by skipping ACKs during dynamic scheduling in low BLER regimes. This is clearly something not typically done in human-designed protocols.

As expected, having to wait for the ACKs before deleting SDUs from the Tx buffer reduces performance, since episodes take longer to complete (see Fig. 8). The presence of BLER in the data channel also slows down training due to a larger number of state transitions (see Fig. 9).

D. The importance of signaling

A major question that arises in *learning-to-communicate* settings is whether gains are due to either optimized action policies or better communication (i.e., signaling in our case), or both. In our MARL formulation, MAC learners have two action spaces and are trained to learn two distinct policies (i.e., a channel access and a signaling policy). Hence, in an extreme case, it is conceivable for the learners to ignore all DL signaling and learn an optimized channel access policy instead. In this case, no MAC protocol signaling would be needed.

To address the previous question, we have calculated the Instantaneous Coordination (*IC*) metric proposed in [26]. For the u^{th} UE, IC^u is defined in this scenario as the mutual information between the
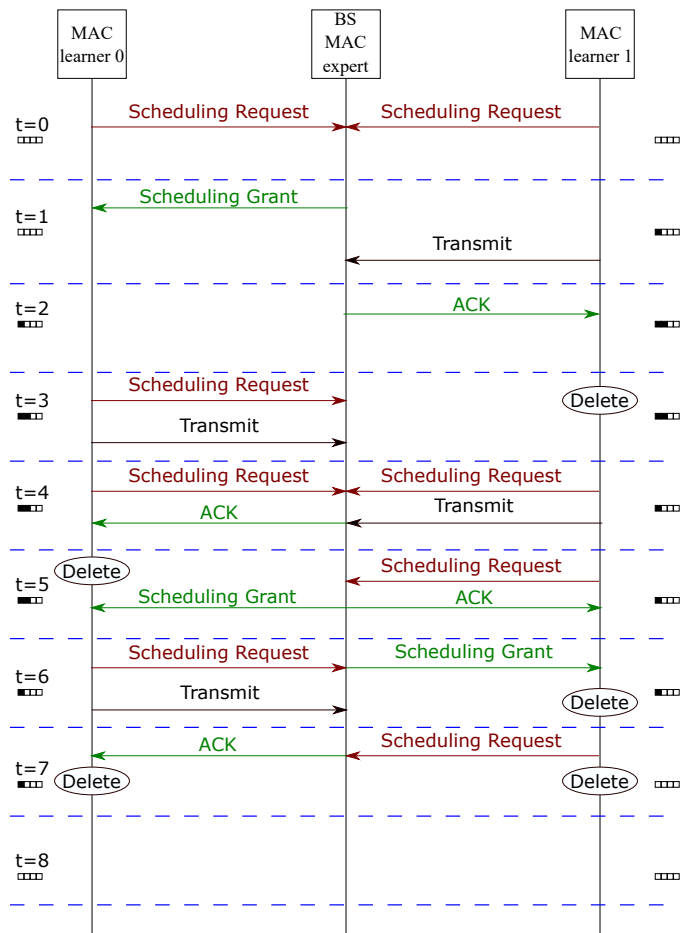

Fig. 7. Best performing evaluation episode in a scenario with $|U| = 2$, $P = 2$, $N = 4$, $t_{max} = 32$, empty buffer start and $BLER = 10^{-3}$. The small 4-cells grid depicts the UL Tx buffer, with one dark cell per available SDU.

BS's DL MAC messages received at time t and the UE MAC's channel access actions at time $t + 1$:

$$IC^u = I(m_t^u; a_{t+1}^u)$$

$$= \sum_{a_{t+1}^u \in \mathcal{A}_\mathcal{P}} \sum_{m_t^u \in \mathcal{M}_{DL}} p(m_t^u, a_{t+1}^u) \log \frac{p(m_t^u, a_{t+1}^u)}{p(m_t^u) p(a_{t+1}^u)}.$$

$$(13)$$

The marginal probabilities $p(m_t^u)$ and $p(a_{t+1}^u)$, and the joint probabilities $p(m_t^u, a_{t+1}^u)$ can be obtained by averaging *DL message* and *channel access action* occurrences in simulation episodes after training.

Fig. 10 shows that performance R and Instantaneous Coordination IC are clearly correlated (i.e., higher performance is concurrently achieved with higher IC). The Pearson correlation coefficient $\rho_{IC,R} = 0.91$ can be calculated as:

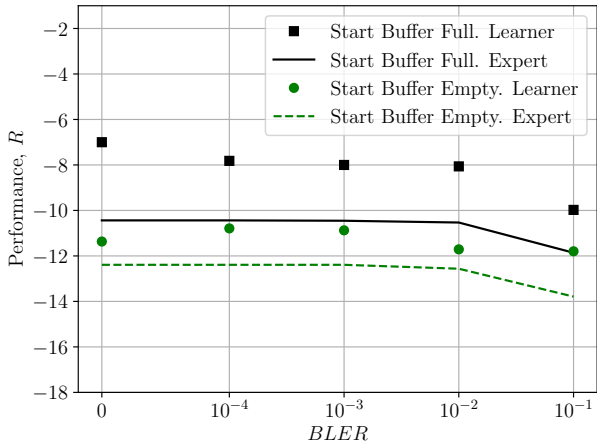$$\rho_{IC,R} = \frac{cov(IC, R)}{\sigma_{IC} \sigma_R} \qquad (14)$$

Fig. 8. Best performance out of 4 independent training sessions. The performance is evaluated over N_{eval} = 128 episodes with different seeds. The environment was configured with $|U|$ = 2, P = 2 and t_{max} = 32. MAC learners were trained with learning rates between 0.1 and 0.3 and an internal memory length of N = 4. Start buffer full learners were trained for $N_{tr} = 2^{20} \approx 1$ million episodes, while start buffer empty learners were trained for $N_{tr} \leq 4$ million episodes. ϵ-greedy exploration with ϵ decay from 1 to 0.01 was used.

Fig. 10. Relationship between the Instantaneous Coordination (IC) and performance R across various training sessions. IC is averaged across all UEs. The length of the vertical error bars equal the standard error of the mean IC, while the length of the horizontal error bars equal the standard error of the mean performance (calculated over 4 independent training sessions). The environment was configured with $|U|$ = 2, P = 2, t_{max} = 32 and empty start buffer. MAC learners were trained with a learning rate of 0.3 and an internal memory length of N = 4 for N_{tr} = 3M episodes. ϵ-greedy exploration with ϵ decay from 1 to 0.01 was used.
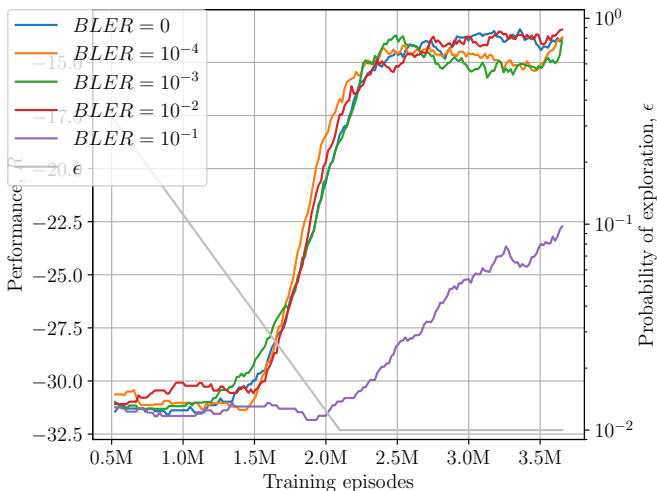

Fig. 9. Average learning progress (over 4 independent training sessions) under various BLER conditions. The environment was configured with $|U|$ = 2, P = 2, t_{max} = 32 and empty start buffer. MAC learners were trained with a learning rate of 0.3 and an internal memory length of N = 4 for $N_{tr} = 2^{22} \approx 4$ million episodes. ϵ-greedy exploration with ϵ decay from 1 to 0.01 was used.

where σ_{IC} and σ_R are the standard deviations of the Instantaneous Coordination and performance samples respectively. This suggests that high performance may only be achievable under high influence (i.e., high IC) from BS onto UEs. The signaling vocabulary (see Fig. 1) plays thus a major role in the performance a MAC protocol can achieve.
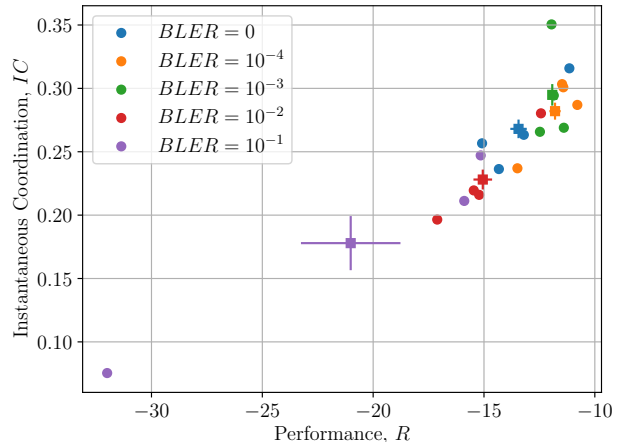
E. Generalization

How well can the learned signaling and channel access policy perform in conditions never seen during training? This question is about the generalization capacity of the learning algorithm (i.e., its robustness against environment variations, see [27]). We address this by first training the MAC learner in a default environment. We then confront the trained learners against environments that differ in a single hyperparameter from the default parameters.

From Figs. 11 and 12, generalization seems robust across environment variations of BLER and $|U|$. MAC learners can therefore be quickly trained in low-load scenarios with mild BLER conditions. Then, these trained systems can still be expected to perform well in more challenging channels. However, Figure 13 shows that performance degrades rapidly with the number of SDUs to be transmitted. In this case, the trained learners have overfit to the number of SDUs to be transmitted. Training in environments with a larger number of SDUs may alleviate this problem, although convergence in that case was elusive in our experiments.
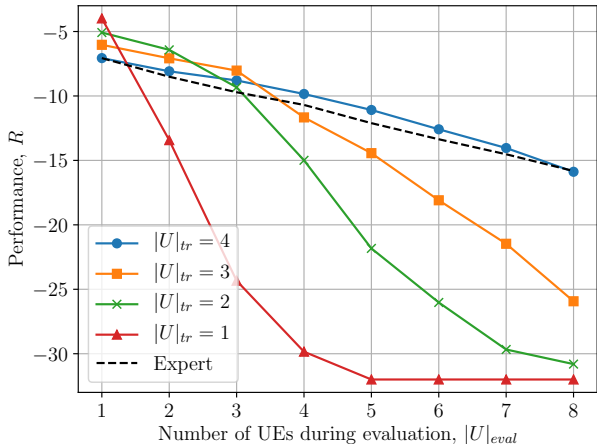
Fig. 11. Best learner's performance (over 16 independent training sessions). Training proceeds with $|U|_{tr}$ MAC learners but performance is evaluated with $|U|_{eval}$ MAC learners. The environment was configured with $P = 1$, $BLER = 10^{-4}$, $t_{max} = 32$ and an empty start buffer. MAC learners were trained with a learning rate of 0.06 and an internal memory length of $N = 4$ for $N_{tr} = 2^{24} = 16M$ episodes. ϵ-greedy exploration with ϵ decay from 1 to 0.01 was used.
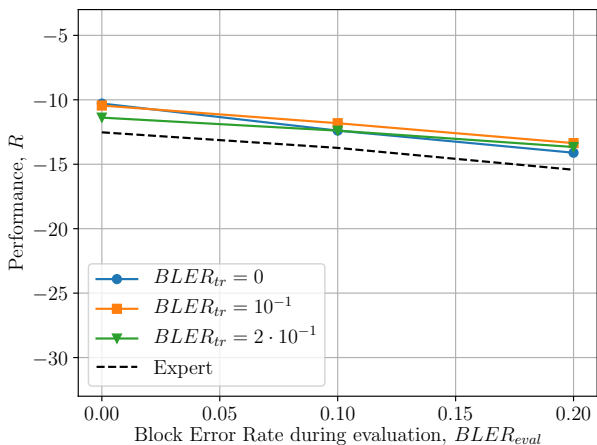

Fig. 12. Best learner's performance (over 32 independent training sessions). Training proceeds with $BLER_{tr}$ but performance is evaluated with $BLER_{eval}$. The environment was configured with $|U| = 2$, $P = 2$, $t_{max} = 32$ and empty start buffer. MAC learners used a learning rate of 0.1 and an internal memory length of $N = 4$. ϵ-greedy exploration with ϵ decay from 1 to 0.01 was used. Learners on the $BLER_{tr} = 0, 10^{-1}, 10^{-2}$ environments were trained respectively for $N_{tr} = 2^{23}, 2^{24}, 2^{25}$ ($\approx 8M, 16M, 32M$) episodes.

Fig. 13. Best learner's performance and (over 4 independent training sessions). Training proceeds with P_{tr} SDUs, but performance is evaluated with P_{eval} SDUs.

The environment was configured with $|U| = 2$, $BLER = 10^{-4}$, $t_{max} = 32$ and an empty start buffer. MAC learners were trained with a learning rate of 0.06 and an internal memory length of $N = 4$ for $N_{tr} = 2^{24} \approx 16M$ episodes. ϵ-greedy exploration with ϵ decay from 1 to 0.01 was used.

that deep learning approaches may be able to scale these techniques to larger and more practical scenarios. The paper has also shown that these agents can achieve the optimal performance when trained *tabula rasa*, i.e., without any built-in knowledge of the target MAC signaling or channel access policy.

A main experimental observation of these experiments is that the trained agents understand but do not always comply with the DL signaling from the BS. This yields MAC protocols that can not be classified as contention-based or coordinated, but fluidly in between. This is in line with 5GNR's grant-free scheduling mechanism, although a major difference is that the protocols learned by our agents are not BS-controlled. A key advantage of protocols learned this way is that they can co-exist with the pre-existing human-designed protocol, while exploiting the optimizations uncovered during training.

Finally, this research has shown that the learned protocols depend on the deployment scenario. This is one reason why they might yield higher performance than non-trainable protocols, whose signaling is static and independent of these parameters.

V. CONCLUSIONS

This paper has demonstrated that RL agents can be jointly trained to learn a given simple MAC signaling and a wireless channel access policy. It has done so using a purely tabular RL approach in a simplified wireless environment. It is expected
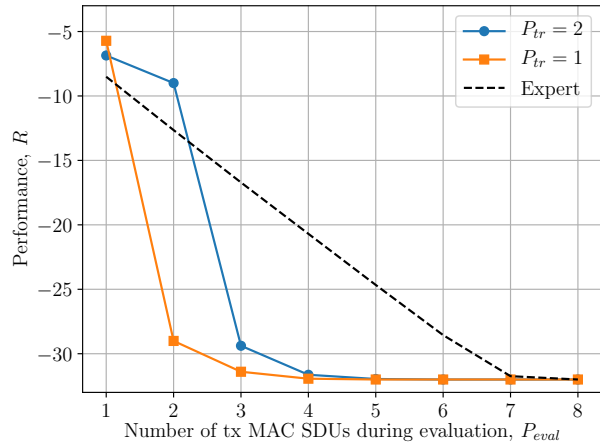
A. Future work

For the methodology presented here to be practical, a number of obstacles must be overcome. The immediate next steps include a study on the sensitivity of the learned policies to the training environment. This should respond to whether it is possible to train MAC learners in networks with a reduced number of UEs and traffic volume, and then deploy them in larger networks. A more detailed scalability analysis is also necessary (can these agents be trained in a reasonable amount of time in larger environments?). The mirror problem of training a BS to learn an expert UE signaling may also illuminate the difficulties of protocol learning.

The ultimate goal is to jointly train UEs and the BS to emerge a fully new MAC protocol (i.e., to learn, not only the signaling and channel access policies, but also the signaling vocabulary of Fig. 1). Recent research (e.g., [26], [13]) suggests this is hard when all agents (i.e., the UEs and the BS) begin with no previous protocol knowledge. Consequently, scheduled training that alternates between supervised learning and self-play, as suggested in [14], seems promising to emerge fully new protocols.

REFERENCES

[1] 3GPP, "3GPP TS 38.321. Medium Access Control (MAC) protocol specification," 2020.

[2] Y. Yu, T. Wang, and S. C. Liew, "Deep-Reinforcement Learning Multiple Access for Heterogeneous Wireless Networks," in *IEEE International Conference on Communications (ICC)*, Kansas City, MO, 2018.

[3] Y. Yu, S. C. Liew, and T. Wang, "Non-Uniform Time-Step Deep Q-Network for Carrier-Sense Multiple Access in Heterogeneous Wireless Networks," *CoRR*, oct 2019. [Online]. Available: http://arxiv.org/abs/1910.05221

[4] A. Destounis, D. Tsilimantos, M. Debbah, and G. S. Paschos, "Learn2MAC: Online Learning Multiple Access for URLLC Applications," *CoRR*, apr 2019. [Online]. Available: https://arxiv.org/abs/1904.00665

[5] Y. Yu, S. C. Liew, and T. Wang, "Multi-Agent Deep Reinforcement Learning Multiple Access for Heterogeneous Wireless Networks with Imperfect Channels," *CoRR*, mar 2020. [Online]. Available: https://arxiv.org/abs/2003.11210

[6] N. Naderializadeh, J. Sydir, M. Simsek, H. Nikopour, and S. Talwar, "When Multiple Agents Learn to Schedule: A Distributed Radio Resource Management Framework," *unpublished*, jun 2019. [Online]. Available: http://arxiv.org/abs/1906.08792

[7] N. Naderializadeh, J. Sydir, M. Simsek, and H. Nikopour, "Resource Management in Wireless Networks via Multi-Agent Deep Reinforcement Learning," *CoRR*, feb 2020. [Online]. Available: http://arxiv.org/abs/2002.06215

[8] F. AL-Tam, N. Correia, and J. Rodriguez, "Learn to Schedule (LEASCH): A Deep reinforcement learning approach for radio resource scheduling in the 5G MAC layer," *unpublished*, mar 2020. [Online]. Available: http://arxiv.org/abs/2003.11003

[9] H. B. Pasandi and T. Nadeem, "MAC Protocol Design Optimization Using Deep Learning," *CoRR*, feb 2020. [Online]. Available: http://arxiv.org/abs/2002.02075

[10] C. Bowyer, D. Greene, T. Ward, M. Menendez, J. Shea, and T. Wong, "Reinforcement Learning for Mixed Cooperative/Competitive Dynamic Spectrum Access," *2019 IEEE International Symposium on Dynamic Spectrum Access Networks, DySPAN 2019*, pp. 1–6, 2019.

[11] T. F. Wong, T. Ward, J. M. Shea, M. Menendez, D. Greene, and C. Bowyer, "A Dynamic Spectrum Sharing Design in the DARPA Spectrum Collaboration Challenge," in *GOMACTech*, San Diego, 2020, pp. 1–6.

[12] J. N. Foerster, Y. M. Assael, N. de Freitas, and S. Whiteson, "Learning to Communicate with Deep Multi-Agent Reinforcement Learning," in *NIPS'16: Proceedings of the 30th International Conference on Neural Information Processing Systems*, vol. abs/1605.0, Barcelona, Spain, 2016, pp. 2145–2153. [Online]. Available: http://arxiv.org/abs/1605.06676

[13] R. Lowe, J. Foerster, Y.-L. Boureau, J. Pineau, and Y. Dauphin, "On the Pitfalls of Measuring Emergent Communication," *CoRR*, vol. abs/1903.0, no. Aamas, 2019. [Online]. Available: http://arxiv.org/abs/1903.05168

[14] R. Lowe, A. Gupta, J. Foerster, D. Kiela, and J. Pineau, "On the interaction between supervision and self-play in emergent communication," in *International Conference on Learning Representations (submitted)*, 2020, pp. 1–13. [Online]. Available: http://arxiv.org/abs/2002.01093

[15] F. Ait Aoudia and J. Hoydis, "Model-free Training of End-to-end Communication Systems," *IEEE Journal on Selected Areas in Communications*, vol. 37, no. 11, pp. 2503 – 2516, 2019. [Online]. Available: https://ieeexplore.ieee.org/document/8792076

[16] Small Cell Forum, "5G FAPI: PHY API Specification," Small Cell Forum, Tech. Rep. June, 2019.

[17] M. L. Littman, "Markov games as a framework for multi-agent reinforcement learning," in *11th International Conference on Machine Learning (ICML)*, 1994, pp. 157–163.

[18] D. H. Wolpert and K. Tumer, "Optimal Payoff Functions for Members of Collectives," *Advances in Complex Systems*, vol. 4, no. 02n03, pp. 265–279, 2001.

[19] J. Foerster, N. Nardelli, G. Farquhar, T. Afouras, P. H. S. Torr, P. Kohli, and S. Whiteson, "Stabilising Experience Replay for Deep Multi-Agent Reinforcement Learning," in *34th International Conferecen on Machine Learning (ICML)*, feb 2017. [Online]. Available: https://arxiv.org/abs/1702.08887

[20] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction (2nd Edition, in preparation)*, 2nd ed. London, England: The MIT Press, 2018. [Online]. Available: http://ir.obihiro.ac.jp/dspace/handle/10322/3933

[21] C. Watkins, "Learning from delayed rewards," Ph.D. dissertation, King's College London, 1989.

[22] M. Tan, "Multi-Agent Reinforcement Learning: Independent vs. Cooperative Agents," in *Proceedings of the Tenth International Conference on Machine Learning*, 1993, pp. 330—337.

[23] G. Tesauro, "TD-Gammon, a Self-Teaching Backgammon Program, Achieves Master-Level Play," *Neural Computation*, vol. 6, no. 2, pp. 215–219, 1994.

[24] P. Hernandez-Leal, B. Kartal, and M. E. Taylor, "A Survey and Critique of Multiagent Deep Reinforcement Learning," *Autonomous Agents and Multi-Agent Systems*, vol. 33, no. 6, pp. 750–797, oct 2019. [Online]. Available: http://arxiv.org/abs/1810.05587http://dx.doi.org/10.1007/s10458-019-09421-1

[25] H. Hu, A. Lerer, A. Peysakhovich, and J. Foerster, ""Other-Play" for Zero-Shot Coordination," in *37th International Conference on Machine Learning (ICML)*, 2020, pp. 4399–4410.

[26] N. Jaques, "Social Influence as Intrinsic Motivation for Multi-Agent Deep Reinforcement Learning," in *36th International Conference on Machine Learning*, R. Chaudhuri, Kamalika and Salakhutdinov, Ed. Long Beach, California, USA: PMLR, 2019, pp. 3040—-3049. [Online]. Available: http://arxiv.org/abs/1810.08647

[27] C. Packer, K. Gao, J. Kos, P. Krähenbühl, V. Koltun, and D. Song, "Assessing Generalization in Deep Reinforcement Learning," *CoRR*, vol. abs/1810.1, 2018. [Online]. Available: http://arxiv.org/abs/1810.12282