Contents lists available at ScienceDirect

Ad Hoc Networks



journal homepage: www.elsevier.com/locate/adhoc

Multi-agent reinforcement learning for network routing in integrated access backhaul networks

Check for updates

Shahaf Yamin^{*}, Haim H. Permuter

School of Electrical and Computer Engineering, Ben-Gurion University of the Negev, Be'er-Sheva, Israel

ARTICLE INFO

ABSTRACT

Keywords: Multi-agent reinforcement learning Integrated access backhaul Network routing

networks involving fiber-connected base stations, wireless base stations, and multiple users. Physical constraints prevent the use of a central controller, leaving base stations with limited access to real-time network conditions. These networks operate in a time-slotted regime, where base stations monitor network conditions and forward packets accordingly. Our objective is to maximize the arrival ratio of packets, while simultaneously minimizing their latency. To accomplish this, we formulate this problem as a multi-agent partially observed Markov Decision Process (POMDP). Moreover, we develop an algorithm that uses Multi-Agent Reinforcement Learning (MARL) combined with Advantage Actor Critic (A2C) to derive a joint routing policy on a distributed basis. Due to the importance of packet destinations for successful routing decisions, we utilize information about similar destinations as a basis for selecting specific-destination routing decisions. For portraying the similarity between those destinations, we rely on their relational base-station associations, i.e., which base station they are currently connected to. Therefore, the algorithm is referred to as Relational Advantage Actor Critic (Relational A2C). To the best of our knowledge, this is the first work that optimizes routing strategy for IAB networks. Further, we present three types of training paradigms for this algorithm in order to provide flexibility in terms of its performance and throughput. Through numerical experiments with different network scenarios, Relational A2C algorithms were demonstrated to be capable of achieving near-centralized performance even though they operate in a decentralized manner in the network of interest. Based on the results of those experiments, we compare Relational A2C to other reinforcement learning algorithms, like O-Routing and Hybrid Routing. This comparison illustrates that solving the joint optimization problem increases network efficiency and reduces selfish agent behavior.

In this study, we examine the problem of downlink wireless routing in integrated access backhaul (IAB)

1. Introduction

The increasing demand for wireless communication and the limitations of the electromagnetic spectrum have led to the development of more efficient methods for managing networks. To meet these needs, the 3rd Generation Partnership Project (3GPP) has established a standard, New Radio (NR), which includes novel designs and technologies to support fifth-generation (5G) networks [1]. One of the key features of this new protocol is the inclusion of new bands at millimeter wave (mmWave) frequencies. These frequencies offer the potential for increased data rates by exploiting the spatial diversity available in these bands, which are currently less congested compared to traditional bands. However, operating at mmWave frequencies also introduces new physical challenges, such as severe path and penetration losses. To overcome these challenges, network density can be increased and beam-forming methods can be used [2]. Although increasing network density has potential benefits, the deployment and operation of fiber between the Next-Generation Node Base Station (gNB) and the core network can be costly. Integrated Access and Backhaul (IAB) is a promising solution for successful 5G adoption as it allows for only a fraction of the gNBs to be connected to traditional fiber-like infrastructures, thus reducing redundant deployment and operational costs by utilizing spatial diversity [3]. The gNBs connected to fiber are called IAB donors, while the remaining gNBs are called IAB nodes and use multi-hop wireless connections for backhaul traffic. While IAB networks are cost-effective in terms of deployment and operation, ensuring reliable network performance remains a challenging research area due to their highly non-stationary nature. The dynamic nature of the topology, tight delay constraints, and limited information regarding the status of the network are some of the factors to be considered when supporting these requirements.

* Corresponding author. E-mail addresses: yamins@post.bgu.ac.il (S. Yamin), haimp@bgu.ac.il (H.H. Permuter).

https://doi.org/10.1016/j.adhoc.2023.103347

Received 6 April 2023; Received in revised form 30 June 2023; Accepted 1 November 2023 Available online 8 November 2023 1570-8705/© 2023 Elsevier B.V. All rights reserved.



Routing plays a crucial role in the context of network congestion control, where each destination may have multiple paths, and base stations monitor network conditions to make routing decisions. There are two main approaches for implementing routing algorithms in wireless networks: centralized and distributed. In a centralized approach, there is a central network processor that is responsible for path selection, while in a distributed approach, each node makes next-hop decisions based only on its own observations without knowledge of other nodes' decisions. In practical implementations, due to bandwidth limitations and multi-hop structures, information sharing is limited to the base station's neighborhood. This implies that base stations can only observe a part of the current network state, and enhance, when operating in a distributed manner next-hop transmission decisions are only based on partial observations.

This paper specifically focuses on the analysis of distributed routing algorithms developed for the downlink channel in integrated access backhaul (IAB)-based networks that exhibit physical limitations. We design a deep reinforcement-learning-based algorithm to achieve optimal routing policies. One unique aspect of our approach is that, given that successful routing decisions rely heavily on packet destinations, we consider that it is beneficial for the agent to use knowledge of similar destinations when making routing decisions. To portray the similarity between these destinations, we use their relational basestation associations, which refer to the base stations they are currently connected to. Our study differs from previous works that focused on designing routing policies to optimize packet paths to specific destinations without sharing information with other similar destinations. These previous algorithms are generally not suitable for a dynamic topology like the one in this study. We propose a novel algorithm for learning a decentralized routing policy using deep policy gradients. To make the learning efficient, we use the Advantage Actor Critic (A2C) [4] algorithm, which is a combination of policy gradient [5], temporal difference estimation [6], and deep neural networks, and enables learning from experience in an unknown environment with a large state space through interactions with the environment. Our proposed algorithm is called Relational A2C.

In the current study, we present numerical results that evaluate the performance of the Relational A2C algorithm in several network scenarios. We also compare the results obtained by Relational A2C with those from other methods such as [7–10]. Our results show that the proposed approach outperforms existing methods and is able to achieve performance comparable to that of centralized systems. To the best of our knowledge, this is the first work that addresses the routing problem in IAB networks using deep reinforcement learning (RL). Additionally, our algorithm formulates the routing problem as a joint optimization problem, which promotes agent cooperation and reduces selfish behavior, resulting in more efficient use of network resources.

1.1. Related work

Routing in networks has been and still is the subject of extensive research, as seen in [11,12]. There is a large body of literature on routing strategies for wireless networks, including earlier protocols such as DSR [13] and AODV [14] for ad-hoc networks, and various routing protocols for delay and disruption tolerant networks (DTNs) [15], as well as strategies for resource constrained wireless networks (such as sensor networks or internet-of-things networks) [16]. However, many existing routing protocols were designed for specific wireless network scenarios and may not be easily adaptable to other scenarios. For example, routing protocols for ad-hoc networks assume a connected network, while routing protocols for DTNs assume a disconnected network. In this study, we focus on routing in IAB networks that are characterized by dynamic topology changes and strict time constraints, and cannot be expected to be subject to these assumptions. This has

motivated the introduction of methods that can acquire a nearly optimal policy without requiring a-priori knowledge. A major technique that is capable of achieving this goal is RL, which is a class of machine learning algorithms that can learn an optimal policy via interaction with the environment without knowledge of the system dynamics (such algorithms are also known as model-free algorithms) [17, Ch. 1]. One of the most popular RL techniques is O-learning [18], which can learn the optimal policy online by estimating the optimal action-value function. Early works that applied Q-learning to network routing used the classical tabular Q-learning method [7,19,20]. This system enables each device to forward a limited number of packets in each time slot; for this it receives a reward over the ACK signal. However, it becomes computationally difficult to apply this method when the state space becomes large. This issue has motivated the combination of deep learning [21] with RL, giving rise to the deep RL class of algorithms. These algorithms have attracted much attention in recent years due to their ability to approximate the action-value function for large state and action spaces. Recently, the authors in [22] proposed a deep RLbased algorithm called deep Q-network (DQN), which combines deep neural networks and Q-learning. Recent studies that derived DRL-based algorithms for network routing problems can be found in [8,23-28]:

A DRL approach for routing was developed in [23] with the objective of minimizing the delay in the network. In this approach, a single controller finds all the paths of all source-destination pairs given the demand in the network, which represents the bandwidth request of each source-destination pair. However, this approach results in complex state and action spaces, and does not scale for large networks as it depends on a centralized controller. Moreover, in this approach, the state representation does not capture the network topology, which is highly dynamic in our scenario. Motivated by the high complexity of the state and action representations of the approaches proposed in [23,24], a feature engineering approach has recently been proposed in [25] that only considers some candidate end-to-end paths for each routing request. This proposed representation was shown to outperform the representation of the approaches proposed in [23,24] in some use-cases. In [26], the authors employed a deep recurrent Q-Network (DRQN), to determine the routing policy, this is a mixture of a DQN and a long short-term memory (LSTM) network. Using device-specific characteristics, such as the last k actions taken and the next m destinations of packets in queues, the algorithm is trained for each device. The LSTM layer in the DRON algorithm uses past observations for the prediction of the whole network state, which, in turn, allows the agent to select the next hop for a specific packet in the next time step. The work in [27] applied an algorithm called hierarchical-DQN (h-DQN) [29]. h-DQN facilitates exploration in complicated environments by decomposing the original problem into a hierarchy of sub-problems such that higherlevel tasks invoke lower levels as if they were primitive actions. In [8], the authors used another RL algorithm called Actor Critic algorithm, which is a policy-based RL algorithm where the policy learned directly via parametrization. They compared their results with that of the algorithm in [7,20] and showed that their proposed algorithm achieves better performance.

1.2. Paper structure and notations

The organization of this paper is as follows. In Section 2, we present the problem formulation and assumptions. In Section 3, we provide the mathematical background for our proposed solution, including Markov Decision Processes (MDPs), RL, and Multi-Agent Reinforcement Learning (MARL). Section 4 presents our mathematical formulation for the problem and explains the rationale behind the chosen MARL approach. In Section 5, we provide a review of existing routing algorithms. Section 6 describes our proposed algorithm, which is based on the A2C method and includes three different training paradigms, ranging from fully decentralized to centralized training. Simulation results, including a comparison with existing routing algorithms, are



Fig. 1. IAB network illustration with 1 IAB donor, 3 IAB nodes and 9 UEs.

presented in Section 7. Finally, in Section 8, we conclude this work and discuss future research directions.

Throughout this work, we use \mathbb{N} to denote natural numbers, bold letters (e.g., **X**) to denote vectors, and **X**_i denotes the *i*th element in the vector **X**, $i \ge 0$. Further, **X**_{t,i} denotes the *i*th element in the vector **X** at the *t*th time-step, $i \ge 0, t \in \mathbb{N}$. we use calligraphic letters to denote sets, e.g., \mathcal{X} , and the cardinality of a set is denoted by $|\cdot|$, e.g., $|\mathcal{X}|$ is the cardinality of the set \mathcal{X} . Lastly, $\mathbb{E}[\cdot]$ denotes the stochastic expectation.

2. Problem formulation

We consider a multi-hop IAB wireless network with an IAB donor, IAB nodes and User Equipments (UE) [3], as shown in Fig. 1. The IAB donor is wired to the core network, whereas IAB nodes use wireless communication to backhaul their traffic from the core network via a multi-hop connection. Our primary focus lies in analyzing the downlink channel within this network. Both the IAB donor and IAB nodes provide access and backhaul interfaces for the UE and IAB nodes, respectively. We model this network by an undirected weighted graph G = (N, L, d), where \mathcal{N} , \mathcal{L} denote the sets of nodes and wireless links, respectively, and d : $\mathcal{L} \rightarrow \mathbb{N}$ assigns a delay to each wireless link. There are three sets present in \mathcal{N} : a set \mathcal{D} of the IAB donor, a set \mathcal{B} of the IAB nodes and a set \mathcal{U} of the UEs, i.e., $\mathcal{N} = \mathcal{D} \cup \mathcal{B} \cup \mathcal{U}$. Each of the nodes $n \in D \cup B$ is equipped with an independent buffer queue, and a transceiver with beam-forming and routing capabilities. Each of the links $(n,m) \in \mathcal{L}$ is a bidirectional link between node n and node m, portraying a time-varying wireless channel.

We assume that time is slotted by $t \in \mathbb{N}$, and, for simplification, we assume that packets are constant in length and that transmission rates are limited to transmit integer numbers of packets per slot. As another simplification, we represent the wireless link's spectral efficiency as a delay between the two nodes of the graph using the mapping *d*. This assumption is based on the fact that links with varying degrees of spectral efficiency will require a different number of transmissions to transfer the same amount of data, so using a link with low spectral efficiency rather than a link with high spectral efficiency will result in an increased delay.

Once an IAB node or UE is activated, it is connected to an already active node, i.e., either an IAB donor or another IAB node that has a path to an IAB donor. Thus, we build the network topology in an iterative greedy fashion, similar to that in [30], where we set constraints on the maximum number of IAB parents (P_{parent}^{max}), IAB children ($C_{children}^{max}$), the number of users each base station has ($U_{children}^{max}$), and the number of associated base stations each user has (U_{parent}^{max}). It should be noted that by using the following topology generation scheme, we obtain a connected graph, i.e., there is a path from any base station to any other node in the network. We assume that all our nodes operate at mmWave bands for backhaul and access transmission and reception (inband backhauling) with beamforming capabilities. Therefore, similar to [31], we disregard the interference between non-assigned nodes since narrow beam mmWave frequencies have a power limit rather than an interference limit.

We address the stochastic nature of packet arrivals using a Poisson process with parameter λ . The packets are distributed among the base stations based on available wireless bandwidth, with any remaining packets distributed uniformly. Each packet is assigned a Time To Live (TTL) as its time limit. To handle dropped packets and balance delay and loss optimization, we employ an iterative approach. When a packet's TTL expires, we inject a new packet with a delay equal to the expired TTL. If the replacement packet is dropped again, we introduce a new packet with a delay of twice the TTL, and so on. Base stations prioritize packets based on their TTL using an unlimited-sized prioritized queue. This ensures that packets with shorter TTLs receive higher priority for delivery, maintaining an efficient processing order.

We denote \mathcal{N}_i as the set of neighbors of node $i \in \mathcal{N}$, i.e., $(i, j) \in \mathcal{L}, \forall j \in \mathcal{N}_i$, and let $C, K \in \mathbb{N}$ be the number of wireless channels and activated base stations, respectively. In each time step, each base station extracts a set of packets from its queue. This is followed by deciding where to send each packet, which means that the *i*th base station has to choose one destination from \mathcal{N}_i for each packet. In our model, users may move or change their base-station associations between two consecutive time slots, which would be resolved in a change of the network topology. In addition, the edge delays are slowly varying around their initial values to modulate the changes in the wireless medium.



Fig. 2. Decision process framework.

3. Background

In the following section, we introduce the mathematical foundations, on which we will base our proposed solution. We begin by exploring MDPs and Partially Observe MDPs (POMDPs). Then, we describe the tools from the field of RL and MARL that we use as the basis for our method.

3.1. MDPs and Partially Observe MDPs preliminaries

We can define MDP as a tuple, $\langle S, A, \mathcal{R}, \mathcal{P} \rangle$, where S, A and \mathcal{R} are the sets of environment states, actions and rewards, respectively. In addition, let \mathcal{P} be the set of the probabilities $\{\Pr(s', r|s, a)\}_{s',s\in S, a\in A, r\in \mathcal{R}}$, whereas the probability $\Pr(s', r|s, a)$ represents the agent probability of observing next state s' and reward r after being at state s and performing action a. The agent's policy can be represented as the following mapping $\pi : S \times A \rightarrow [0, 1]$, which represents a mapping from the current state to the probability distribution on the action space.

A Partially Observe MDP (POMDP) is defined as a tuple $(\mathcal{O}, S, \mathcal{A}, \mathcal{R}, \mathcal{P})$ [32]. An agent interacting with the environment at state $s \in S$ observes an observation $o(s) \in \mathcal{O}$. After observing o, the agent selects an action $a \in \mathcal{A}$ based on this observation, which means that now the agent's policy is determined by its observations. That is, policy can be represented as the following mapping $\pi : \mathcal{O} \times \mathcal{A} \rightarrow [0, 1]$, which represents a mapping from the current observation to the probability distribution on the action space. After performing an action, the remainder of the decision process is the same as for the MDP case.

3.2. RL preliminaries

We examine a general model-free RL framework [17] applied to a specific MDP, in which the agent interacts with the environment and learns to accomplish the task at hand by a series of discrete actions. In this case, the agent does not assume any prior knowledge of the underlying environment statistics. Fig. 2 describes the decision making processes between the agent and the environment: at each discrete time step *t*, the agent observes the current state of the environment S_t and executes an action A_t according to its policy π . Then, the agent receives an immediate reward R_t and the environment transitions to a new state S_{t+1} , based on the transition kernel $Pr(S_{t+1}|S_t, A_t)$, to the next state S_{t+1} .

Under this framework, the agent's goal is to select actions that maximize the expected cumulative discounted reward G_t , where we define G_t as follows:

$$G_{t} = \sum_{n=0}^{\infty} \gamma^{n} R_{t+n+1}, \ \gamma \in [0,1).$$
(1)

The discount factor γ determines how much immediate rewards are favored over more distant rewards. For a general MDP framework, the action-value function $Q_{\pi}(S_i, A_i)$ represents the expected accumulated



Fig. 3. Multi-agent decision process framework.

discounted reward starting from state S_t , picking action A_t , and following policy π afterwards, whereas the value function $V_{\pi}(S_t)$ represents the expected accumulated discounted reward starting from state S_t , and following policy π [17, Ch. 3]:

 $Q_{\pi}(S_t, A_t) \triangleq E_{\pi}[G_t|S_t, A_t], V_{\pi}(S_t) \triangleq E_{\pi}[G_t|S_t].$

The optimal policy π^* is a policy that satisfies $Q_*(S_t, A_t) \triangleq Q_{\pi^*}(S_t, A_t) \ge Q_{\pi}(S_t, A_t)$ for any policy π and for every possible state–action pair. By continuously interacting with the environment, the RL agent aims to learn the optimal policy π^* .

3.3. MARL preliminaries

MARL also addresses sequential decision-making problems, but with more than one agent involved, as illustrated in Fig. 3. In particular, both the evolution of the system state and the reward received by each agent are influenced by the joint actions of all agents. More intriguingly, each agent has to optimize its own long-term reward, which now becomes a function of the policies of all other agents [33].

As we are interested in optimizing the network performance in this study, it is essential that our agents act cooperatively. In a fully cooperative setting, all agents will tend to share a common reward function, such as $R_1 = R_2 = \cdots = R_N = R$. We note that this model may also be referred to as a multi-agent MDP (MMDP) in the AI community [34]. In this model, the value function and Q-function are identical for all agents, which enables the application of single-agent RL algorithms, where all agents are coordinated as one decision maker [33].

Besides the common-reward model, another slightly more general and emerging model for cooperative MARL considers team-average rewards [33,35]. Specifically, agents are allowed to have different reward functions, which may be kept private for each agent, while the goal for cooperation is to optimize the long-term reward corresponding to the average reward $R \triangleq \frac{1}{N} \sum_{i=0}^{N-1} R_i$. The average-reward model, which allows more heterogeneity among agents, includes the model above as a special case. It also preserves privacy among agents, and facilitates the development of decentralized MARL algorithms [33,35].

4. RL for IAB network routing

In the following section, we formulate and evaluate the IAB network routing problem using a Multi-Agent POMDP. First, we propose and analyze an approach to solve this problem. Next, we proceed to formulate it mathematically as a Multi-Agent POMDP. We conclude our discussion by describing our evaluation metrics.

During the agents' training phase, their experience is formed from tuples of $\langle s, a, r \rangle$ that are chained together into time sequences by the

next state, s'. We can define the decision-making process in our network experience in two ways. One approach is to select the base station as the decision-maker, where each base station acts as an independent agent and makes its own decisions on how to forward packets. Alternatively, we can select the packet as the decision-maker, where every packet in the network becomes an agent and independently decides its action when it reaches the front of a base station's queue.

In our work, we have employed the packet-based approach to mathematically define the decision-making process. It is important to note that this approach serves as a convenient mathematical framework for the development of our algorithms. In practical implementation, the base stations, acting as agents, make decisions on behalf of the packets in their respective queues. This choice allows us to leverage the benefits of the packet-based approach and its associated mathematical convenience while ensuring a practical and realistic implementation. Considering that our network consists of multiple packets, the decisionmaking process can be seen as a multi-agent problem with increased complexity. This complexity arises from the simultaneous operation of multiple packets within the same network.

4.1. Formulating IAB routing using a Multi-Agent POMDP framework

This section outlines the mathematical formulation of our problem. Due to the presence of multiple agents with a joint goal, we formulate the problem of multi-agent routing as a Multi-Agent POMDP [32] with discrete actions. Each agent observes statistics relevant to itself and does not observe the entire global state. We denote the Multi-Agent POMDP by a tuple $(S, \mathcal{O}, \mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_N, \mathcal{P}, \mathcal{R}, N, \gamma)$, where N is the number of agents and S is the environment state space. Environment state $s \in S$ is not fully observable. Instead, agent *i* draws a private observation $\mathbf{o}_i \in \mathcal{O}$ that is correlated with s. \mathcal{A}_i is the action space of agent *i*, yielding a joint action space $\mathcal{A} = \mathcal{A}_1 \times \mathcal{A}_2 \times \cdots \times \mathcal{A}_N$. Pr(s', r|s, a) \in \mathcal{P} is the state-reward transition probability, where $r, \mathbf{s}, \mathbf{s}', \mathbf{a} \in \mathcal{R} \times S \times$ $S \times A$. \mathcal{R} , γ represent the available rewards set and the discount factor, respectively. Agent $i \in \{1, 2, ..., N\}$ uses a policy $\pi_i : \mathcal{O} \times \mathcal{A}_i \to [0, 1]$ to choose actions after drawing observation o_i . After all agents have taken actions, the joint action **a** triggers a state transition $s \rightarrow s'$ based on the state transition probability Pr(s'|s, a).

In Mulit-Agent POMDP, MARL can be used as a computational tool. Below, we specify each component using the MARL definitions.

- **Observations.** The agent can only observe information relevant to the packet it controls. Specifically, we consider the packet's current node, time to live (TTL), and queue delay in this work. Let $\mathbf{o}_i = (n, t, Queue Delay(i, n))$ be the agent *i* observation. The scope of *n* in this context is limited to only the base stations and the final destination, i.e., other users cannot relay messages.
- Actions. If agent *i* is authorized to perform a wireless hop based on $\mathbf{o}_i = (n, t)$, the allowable action set includes all wireless links that are available from node *n*. For example, $a_i = l_{n,n'}$, where $l_{n,n'} \in \mathcal{L}$ represents the link between nodes *n* and *n'*. Let $\mathbf{a} = (a_1, a_2, \dots, a_N)$ be the joint action. Whenever the agent does not have permission to conduct a wireless hop, its action will be defined as null.
- **Transitions.** Joint action **a** triggers a state transition $s \rightarrow s'$ with probability $\Pr(s'|s, \mathbf{a})$ which depends on the dynamics of the environment and on the frequency at which the agent is polled to provide an action. In this case, the frequency is dependent on the duration of each time slot in the system.
- Reward. Let D_i represent the immediate delay of the *i*th agent. To be more specific, D_i consists of two components that cause delay to the *i*th agent: the instant wireless link delay and the delay caused by waiting at the current base-station queue. We define the current agent's delay as the sum of these two terms.

If we define the observation as $\mathbf{o}_i = (n, t)$ and the action as $a_i = l_{n,n'}$, we can denote the immediate delay of the agent as $D_i = -(q_i^n + d((n, n')))$, where q_i^n represents the delay induced by the *n*th node queue to the *i*th agent. Accordingly, we define the agents' joint reward using the immediate delay representation as follows: $R \triangleq \frac{1}{N} \sum_{i \in \mathcal{I}} D_i \in \mathcal{R}$, where \mathcal{I} represents the set of active agents.

Let $\Pi \triangleq \times_{i=0}^{N-1} \pi_i$ represent the joint agent policy. Let $\mathbf{s}_i \in S$ denote the state of the network at time $t \in \mathbb{N}$. Our objective is to derive an RL-based algorithm to identify the set of policies $\{\pi_i\}_{i=0}^{N-1}$ that maximize the expected accumulated discounted reward over a finite time horizon, i.e,

$$\Pi^{\star} = \underset{\Pi}{\operatorname{argmax}} \left\{ \mathbb{E}_{\Pi} \left[\sum_{t=1}^{T} \gamma^{t-1} R_{t+1} \Big| \mathbf{s}_{1} \right] \right\}.$$
⁽²⁾

4.2. Evaluation metrics

A reactive routing scheme is used in a multi-hop IAB network to dynamically minimize packet delay while ensuring that packets reach their destination on time. There might be multiple hops, links with low spectral efficiency, and nodes with overloaded queues in a packet's path, all of which may cause delay. There are various metrics to measure or estimate the congestion within the network. In our scenario we define our congestion estimation using the following metrics:

- Average Packet latency The average time it takes for a packet to travel from its source to its destination.
- Arrival Ratio The percentage of packets that made it to their destination successfully.

This multi-objective problem aims to minimize the packet latency while simultaneously maximizing the arrival ratio. Therefore, it may suffer from a Pareto-front, which means that optimizing w.r.t. one objective, leads to a sub-optimal solution w.r.t. another objective [36]. Despite the fact that the network performance measurements are well defined, an individual agent does not necessarily have access to their signals. For example, the arrival ratio represents a metric which is dependent upon the entire network; due to its multi-hop structure, at each time slot an individual cannot even obtain a good estimation of this value.

5. Existing routing algorithms

The following section describes existing solutions for solving a routing problem. Throughout this section, we explore and modify those algorithms to solve the problem of routing in an IAB network.

5.1. Centralized-routing

In this policy, the algorithm has access to the full network state. During each time step, the next hop is determined by computing the shortest path to the packet's destination. By observing the full state while calculating the shortest path, this algorithm also accounts for delays caused by queues at other base stations along the packet's path. Even though this algorithm has a high complexity, it achieves network performance improvement with full knowledge of the network's state, and serve as a benchmark in comparison to a decentralized approach.

5.2. Minimum-Hop routing

In this policy, the algorithm has access only to the links' delay [9]. Based solely on those link delays, the next hop is determined for each time step by computing the shortest path to the packet's destination. In addition to serving as a benchmark, this algorithm is intended to analyze the influence of queue delay on the resulting delay.

5.3. Back-pressure routing

In this policy [10], each base-station stores a queue for each possible destination. The following procedure describes this algorithm. In each time slot, each node calculates the difference between the specific queue length and the corresponding queue length located at each of its neighboring nodes (this calculation occurs for each queue). By utilizing this difference for all available destinations, each base station is able to make two different decisions. First, it makes a greedy decision regarding which packets to extract from the queues based on the size of the differentiation (in other words, the destinations with the greatest degree of differentiation will be selected). Second, the node will determine the packet's next hop based on a greedy decision (the neighbor with the highest differentiation score).

5.4. Q-Routing

(

In this policy, each node uses an off policy iterative method termed Q-learning [7]. Next, Q-learning is first used to learn a representation of the network state in terms of Q-values and then these values are used to make routing decisions. Given this representation of the state, the action *a* at node *n* is to find the best neighbor node to deliver the a packet that results in lower latency for the packet to reach its destination. The following procedure describes this algorithm. As soon as node *n* extracts a packet destined to node *d* from his queue, it selects the next-hop decision based on the ϵ -greedy policy w.r.t. its Q-values, i.e.,

$$y = \begin{cases} UniformRandom(\mathcal{A}_n) & w.p. \ \epsilon, \\ \arg\max_{a \in \mathcal{A}_n} \hat{Q}^{(n)}(a, d) & w.p. \ 1 - \epsilon. \end{cases}$$
(3)

Once the packet has been sent from node *n* to neighboring node *y*, node *y* will send its best estimate back to node *n*, $\max_{a' \in A_y} \hat{Q}^{(y)}(a', d)$ for the destination *d* back to node *n* over the ACK signal. This value essentially estimates the remaining time in the journey of the packet. Following that, the Q-value is modified by the following formula:

$$\hat{Q}_{new}^{(n)}(y,d) = \hat{Q}_{old}^{(n)}(y,d) + \alpha \cdot \left(r + \gamma \cdot \max_{a' \in \mathcal{A}_y} \hat{Q}^{(y)}(a',d) - \hat{Q}_{old}^{(n)}(y,d)\right).$$
(4)

5.5. Full Echo Q-Routing

This algorithm is a variant of the Q-Routing algorithm described above, and addresses the well-known issue of exploration against exploitation [7]; similarly to Q-Routing, it utilizes an iterative off-policy method termed Q-Learning. In order to eliminate the problem of exploration against exploitation, this algorithm executes the following procedure. As soon as node *n* extracts a packet destined to node *d* from its queue, it selects the next-hop decision based on the greedy policy w.r.t. its Q-values, $y = \operatorname{argmax}_{a \in \mathcal{A}_n} \hat{Q}^{(n)}(a, d)$. Once the packet has been sent from node *n* to neighboring node *y*, node *n* receives information from all neighboring nodes, which transmit their respective best estimates for packets destined to node *d*, i.e., $\max_{a' \in \mathcal{A}_y} \hat{Q}^{(y)}(a', d), \forall y \in \mathcal{N}_n$. Using those estimations, node *n* modifies its Q-values for each neighbor by utilizing the following formula:

$$\hat{Q}_{new}^{(n)}(y,d) = \hat{Q}_{old}^{(n)}(y,d) + \alpha \cdot \left(r + \gamma \cdot \max_{a' \in \mathcal{A}_y} \hat{Q}^{(y)}(a',d) - \hat{Q}_{old}^{(n)}(y,d)\right), \forall y \in \mathcal{N}_n.$$
(5)

5.6. Hybrid routing

In this algorithm, a Q-Routing agent is trained simultaneously with an on policy iterative method [8]. In this case, Q-learning is used to learn a representation of the network state in terms of Q-values, and then Hybrid routing uses these values to update the agent's policy parameters by using the Actor-Critic method. As soon as node *n* extracts a packet destined to node *d* from its queue, it selects its next action based on sampling its policy distribution, $\pi_n(d, \cdot; \theta_n)$. Then, it sends the packet to one of its neighboring nodes *y*. The corresponding Q-value is updated based on a Q-Routing update rule, and then its policy parameters θ_n are updated according to the following formula:

$$\theta_{n} \leftarrow \theta_{n} + \alpha \cdot \nabla_{\theta_{n}} \log \pi_{n}(y, d; \theta_{n}) \cdot \left(r + \gamma \max_{a' \in \mathcal{A}_{y}} \hat{Q}^{(y)}(a', d) - \max_{a \in \mathcal{A}_{n}} \hat{Q}^{(n)}(a, d)\right).$$

6. Our proposed multi-agent relational A2C routing algorithm

In this study, we are primarily interested in exploring a decentralized solution to the problem of routing in an IAB network. In the following section, we present our novel solutions to this issue. First, we discuss the motivation behind our proposed solution and the challenges it seeks to address. Following this, we describe the main characteristics of our solution. Finally, we present three different training paradigms within our approach, ranging from fully decentralized training to centralized training.

As a first step, we attempted to solve the above task using traditional RL techniques such as Q-routing [7], Full-Echo Q-Routing [7], and Hybrid Routing [8]. These methods have not achieved a high degree of generalization due to the challenges posed by the above task, such as partial observability, a large state space, and multi-agent optimization. Essentially, these methods assume that each agent acts independently and does not share their experience with other agents, resulting in performance degradation as a result of insufficient correlation between reward signals, network state, and other agents' policies.

Our proposed solution addresses these issues by formulating this problem as a Mutli-Agent POMDP as described in Section 4.1. We define our algorithm objective in the same manner as described in Eq. (2) to encourage cooperation among the different agents. Furthermore, we leverage the homogeneity between destinations in order to support an invariant number of users, as we might encounter in a real-world scenario. To this end we categorize destinations into groups based on their relational base-station association, i.e., which base station they are currently connected to. Through this categorization, policy and value functions are shared by each group. As part of this process, each agent uses an iterative online on-policy method called Advantage Actor-Critic (A2C) [17, Ch. 13]. According to this scheme, the actor decides what action to take and the critic informs the actor of its effectiveness and how it should be adjusted. Based on the current observation, critic produces an estimated representation of the network state, and the actor uses this information to update its policy. Every agent in the network represents its own strategy through its actor. To contend with the issue of the large size of the state space, we propose using neural networks to approximate both the actor and the critic. Consequently, since we categorize agents according to their relational association with base stations, we refer to this algorithm as Relational A2C.

In this section, we present three algorithms based on our method, each in its own subsection. The first algorithm is Relational A2C, in which training is centralized. Specifically, all agents use the same global actor and global critic. Fig. 4(a) shows how these are updated based on information from all of the IAB stations. The second algorithm is the Dec-Relational A2C algorithm, in which each base station has a local actor and critic that are trained based on the local base station experience. Fig. 4(b) illustrates how the information from the base station is used to update these models. Lastly, the third algorithm is based on federated learning [37], which we refer to as Fed-Relational A2C. Most of the time, this algorithm operates in a decentralized manner, but once within a given period of time, it converts the weighted averages of these local model weights into a global model, which then is broadcast to the base stations.

6.1. Relational A2C

In this section, we present the general idea behind our method and explain our centralized training solution. As a first step, we unify all packets destined to the same destination in accordance with the same

Fig. 4. Illustration of centralized and decentralized training paradigms for our proposed method.

Fig. 5. An illustration of the categorization mapping *H* for a scenario with K = 4.

policy. In the following step, for any possible graph topology, denoted as $\mathcal{L}^* \triangleq \{\mathcal{L} | \mathcal{L} \text{ is a set of graph edges that forms network topology}\}$, we define a mapping $H : \mathcal{N} \times \mathcal{L}^* \to \{0,1\}^K$ that maps agents destined to nodes $n \in \mathcal{N}$ to a given group. In this case, K and $|\{0,1\}^K|$ represent the numbers of available base stations and different groups that our agents are divided into, respectively. Next, we define our categorization mapping, $H : \mathcal{N} \times \mathcal{L}^* \to \{0,1\}^K$ as follows:

Next, we present the centralized training paradigm, namely *Relational A2C* (Fig. 4(a)). In view of the fact that we are utilizing neural networks, we would like to prevent bias in our input and categorize

$$H(i, \mathcal{L})[j] \triangleq \begin{cases} 1, & \text{if } (i, j) \in \mathcal{L}, \\ 0, & \text{if } (i, j) \notin \mathcal{L}. \end{cases}$$

Fig. 6. Relational A2C neural architecture located at the *i*th base station.

agents in a meaningful way. Thus, we propose encoding the observations of the *n*th agent under the assumption that it is located at the *i*th node as follows:

$$\mathbf{o}_{n} \triangleq [oneHot(i), t, H(destination(n), \mathcal{L}), QueueDelay(n, i)], \tag{6}$$

where we define $oneHot : \mathcal{N} \to \{0,1\}^K$ as follows:

one
$$Hot(n)[j] \triangleq \begin{cases} 1, & j = n, \\ 0, & \text{otherwise.} \end{cases}$$

The last variables, t and QueueDelay(n, i), represent the TTL and the queue delay of agent n at the node i, respectively.

Let $\Pi(\cdot; \theta), \hat{V}_{\Pi}(\cdot; \mathbf{w}), \theta \in \mathbb{R}^{L_1}, \mathbf{w} \in \mathbb{R}^{L_2}, L_1, L_2 \ll |\mathcal{A}| \cdot |\mathcal{S}|$ be our actor and critic representations, respectively, as depicted in Fig. 6. In spite of the fact that we are only provided with one actor and one critic for all groups, we are still able to differentiate between them due to the mapping *H*. Considering this representation, and assuming the packet *n* is located at the *i*th base station, the packet strategy determines the next hop decision by sampling the corresponding actor distribution $\Pi(\mathbf{o}_n; \theta)$. Next, the *i*th base station sends the agent via a packet to its next hop decision. After that, the base station receives instantaneous feedback through the ACK signal that contains the critic's estimation for the next hop and the agent's instant delay, $\hat{V}_{\Pi}(\mathbf{o}'_n; \mathbf{w})$ and D_n , respectively. The critic's value represents the next-node estimate of remaining time in a packet's journey to the agent destination, when the agent follows policy Π .

Next, we consider $\hat{V}_{\Pi}(\mathbf{o}_{t,n}; \mathbf{w})$ as the estimation of the *n*th agent averaged delay given the current observation, i.e., $\mathbb{E}[\sum_{k=0}^{T-t-1} \gamma^k \cdot D_{n,t+1+k} | \mathbf{o}_{t,n}, \Pi]$. We derive the following equality based on the Bellman equation for a single agent scenario applied for the *n*th agent:

$$V_{\Pi}(\mathbf{o}_{t,n};\mathbf{w}) = \mathbb{E}[D_{n,t+1} + \gamma \cdot V_{\Pi}(\mathbf{o}_{t+1,n};\mathbf{w})|\mathbf{o}_{t,n},\Pi].$$

Hence, we propose to update the critic by minimizing the estimation error. This is achieved by minimizing the following objective T w.r.t. w:

$$\mathcal{T}(\Pi) = \mathbb{E}\left[\left(\hat{V}_{\Pi}(\mathbf{O}_n; \mathbf{w}) - \left(D_n + \gamma \cdot \hat{V}_{\Pi}(\mathbf{O}'_n; \mathbf{w})\right)\right)^2 \middle| \Pi \right],\tag{7}$$

where O_n , D_n and O'_n represent the random variables of observation, instant delay and next observation for a random agent. We update the critic's parameters using the gradient descent method:

for some learning rate $\alpha \in (0, 1)$. Next, we seek to identify a set of policy parameters θ that maximizes the expected cumulative discounted reward (refer to the definition in (2)). To do so, we derive the following lemma.

Lemma 1. The gradient of the objective $J(\Pi)$ w.r.t. θ is proportional to the following estimator:

$$\nabla_{\boldsymbol{\theta}} J(\boldsymbol{\Pi}) \propto \mathbb{E} \left[\sum_{n' \in \mathcal{I}_{t}} \nabla_{\boldsymbol{\theta}} \log \boldsymbol{\Pi}(\boldsymbol{a}_{t,n'} | \mathbf{o}_{t,n'}; \boldsymbol{\theta}) \left(\boldsymbol{Q}_{\boldsymbol{\Pi}}(\mathbf{s}_{t}, \mathbf{a}_{t}) - \frac{1}{|\mathcal{I}_{t}|} \sum_{n \in \mathcal{I}_{t}} \hat{V}_{\boldsymbol{\Pi}}(\mathbf{o}_{t,n}; \mathbf{w}) \right) \right],$$

where I_t represents the set of agents that were active at timestep t.

Proof. The full proof is attached in Appendix B. \Box

According to its definition, the action-value function $Q_{II}(\mathbf{s}_t, \mathbf{a}_t)$ represents the average delay of all packets while the network is at state \mathbf{s}_t , uses action \mathbf{a}_t , and follows the policy Π . Our proposal is to estimate the average network delay $\hat{Q}_{II}(\mathbf{s}_t, \mathbf{a}_t)$ using our active agents in order to accommodate the partially observed issue, i.e.,

$$\hat{Q}_{\Pi}(\mathbf{s}_t, \mathbf{a}_t) = \frac{1}{|\mathcal{I}_t|} \sum_{n \in \mathcal{I}_t} \hat{Q}_{\Pi}(\mathbf{o}_{t,n}, a_{t,n}) = \frac{1}{|\mathcal{I}_t|} \sum_{n \in \mathcal{I}_t} D_{n,t+1} + \gamma \cdot \hat{V}_{\Pi}(\mathbf{o}_{t+1,n}; \mathbf{w}).$$

Following that, we can estimate the gradient of the objective $J(\Pi)$ w.r.t. θ at the *t*th time-step using the following sample:

$$\nabla_{\boldsymbol{\theta}} \hat{\boldsymbol{J}}(\boldsymbol{\Pi}) = \sum_{i \in \mathcal{I}_{t}} \left(\frac{1}{|\mathcal{I}_{t}|} \sum_{n \in \mathcal{I}_{t}} D_{n,t+1} + \gamma \cdot \hat{V}_{\boldsymbol{\Pi}}(\mathbf{o}_{t+1,n}; \mathbf{w}) - \hat{V}_{\boldsymbol{\Pi}}(\mathbf{o}_{t,n}; \mathbf{w}) \right) \\ \cdot \nabla_{\boldsymbol{\theta}} \log \boldsymbol{\Pi}(\boldsymbol{a}_{t,i} | \mathbf{o}_{t,i}; \boldsymbol{\theta}).$$
(8)

This estimation allows us to apply the stochastic gradient ascent method in order to find the local optimal solution. That is, at each time step t, the parameter θ is updated by

$$\theta \leftarrow \theta + \eta \nabla_{\theta} \hat{J}(\Pi)$$

for some learning rate $\eta \in (0, 1)$. By using this method, we expect to increase the cooperation between the different agents through the optimization of the joint objective (Fig. 4(a) illustrates this procedure).

The steps of the proposed Relational A2C algorithm are summarized in Algorithm 1 below:

Algorithm 1 The Relational A2C Algorithm for Simultaneously Optimize Routing Strategy

1: Initialize Actor and Critic weights θ , w and learning rates η , α . 2: **for** time step t = 1, 2, ..., T **do** 3: Extract active agents \mathcal{I}_{t} . for Agent $n \in I_t$ do 4: 5: Observe $\mathbf{o}_{t,n}$. $\mathbf{a}_{n,t} = a, \text{w.p. } \Pi(a | \mathbf{o}_{t,n}; \theta), \forall a \in \mathcal{A}_{n,t},$ 6: end for 7: 8: Execute actions. for Agent $n \in I_t$ do Q٠ Obtain the delay $D_{n,t+1}$ and next observation $\mathbf{o}_{t+1,n}$ associ-10: ated with the n^{th} agent. Set Temporal Difference Error. 11: $\delta_{t,n} = D_{n,t+1} + \gamma \cdot \hat{V}(\mathbf{o}_{t+1,n}; \mathbf{w}) - \hat{V}(\mathbf{o}_{t,n}; \mathbf{w}).$ 12: 13: end for Update critic and actor parameters based on Eq. (7) and Eq. (8). 14: $\boldsymbol{\theta}_{t+1} \leftarrow \boldsymbol{\theta}_t + \eta \cdot \left(\sum_{n \in \mathcal{I}_t} \nabla_{\boldsymbol{\theta}_t} \log \left(\boldsymbol{\Pi}(\boldsymbol{a}_{t,n} | \mathbf{o}_{t,n}; \boldsymbol{\theta}) \right) \cdot \left(\frac{1}{|\mathcal{I}_t|} \sum_{n' \in \mathcal{I}_t} \delta_{t,n'} \right) \right).$ 15: $\boldsymbol{w}_{t+1} \leftarrow \boldsymbol{w}_t - \alpha \cdot \nabla_{\boldsymbol{w}} \left(\sum_{n \in \mathcal{I}_t} (\delta_{t,n})^2 \right).$ 16: 17: end for

6.2. Dec-Relational A2C

In this section we extend our method to support a decentralized training paradigm, namely, Dec-Relational A2C. For decentralized training, we propose decoupling the neural network across the different base stations so that each base station uses its own set of weights for both actor and critic, as depicted in Fig. 6. With a dedicated neural network for each base station, the base station ID in the observation becomes redundant, so we drop it, leading to a new observation $\mathbf{o}_n = [t, H(destination(n), \mathcal{L}), QueueDelay(n)],$ where QueueDelay represents the amount of time the *n*th agent waited at the current base station queue. Let $\hat{\Pi}_k(\cdot; \boldsymbol{\theta}_k), \hat{V}_{\Pi}(\cdot; \mathbf{w}_k), \boldsymbol{\theta}_k \in \mathbb{R}^{L_1}, \mathbf{w}_k \in \mathbb{R}^{L_2}$, be our actor and critic representations at the kth base-station, respectively. In addition, let $k, n, k', \mathbf{o}_n, \mathbf{o}'_n$ be our current node, agent, next node decision, observation and next step observation, respectively. Following that, based on the Bellman equation we propose to modify the temporal difference estimation as follows:

$$\delta_{k,n,k'} = D_n + \gamma \cdot \hat{V}_{\Pi}(\mathbf{o}'_n; \mathbf{w}_{k'}) - \hat{V}_{\Pi}(\mathbf{o}_n; \mathbf{w}_k), \tag{9}$$

such that there is a mutual update between the different base stations through the ACK signal. Essentially, by using this modification we hope that the next base-station estimation combined with the instant delay will represent the estimated path delay from the current base station for this agent. Next, let I_k represent the set of agents who took a next-hop decision from base station k. The critic is then updated at each base station *k* by minimizing the following objective T_k w.r.t. \mathbf{w}_k :

$$\mathcal{T}_{k}(\hat{\Pi}_{k}) = \mathbb{E}\bigg[\sum_{n\in\mathcal{I}_{k}}\delta_{k,n,k_{n}'}^{2}\Big|\hat{\Pi}_{k}\bigg],\tag{10}$$

where we update critic's parameters using the gradient descent method:

 $\mathbf{w}_k \leftarrow \mathbf{w}_k - \alpha \cdot \nabla_{\mathbf{w}_k} \mathcal{T}_k(\hat{\Pi}_k),$

for some learning rate $\alpha \in (0, 1)$. As a next step, we aim to approximate the objective gradient w.r.t. θ_k .

Lemma 2. The gradient of the objective $J(\Pi)$ w.r.t. θ_k is proportional to the following estimator:

 $\nabla_{\theta_k} J(\Pi)$

$$\mathbb{E}\left[\sum_{n\in\mathcal{I}_{k,t}}\nabla_{\boldsymbol{\theta}_k}\log\Pi(\boldsymbol{a}_{t,n}|\mathbf{o}_{t,n};\boldsymbol{\theta}_k)\big(\boldsymbol{Q}_{\Pi}(\mathbf{s}_t,\mathbf{a}_t)-\frac{1}{|\mathcal{I}_{k,t}|}\sum_{n'\in\mathcal{I}_{k,t}}\hat{V}_{\Pi}(\mathbf{o}_{t,n};\mathbf{w}_k)\big)\right],$$

where $I_{k,t}$ represents the set of agents that were active at the kth base station at timestep t.

Proof. The full proof is attached in Appendix C.

To support decentralized training, we first proposed a local estimation of average network delay, $\hat{Q}_{II}^{(k)}(\mathbf{s}_t, \mathbf{a}_t) = \frac{1}{|I_{k,l}|} \sum_{n \in I_{k,t}} D_{n,t+1} +$ $\gamma \cdot \hat{V}_{\Pi}(\mathbf{o}_{t+1,n}; \mathbf{w}_{a_{t,n}})$. Using this estimation and following Lemma 2, we can estimate the objective $J(\Pi)$ gradient w.r.t. θ_k using the following sample

$$\nabla_{\boldsymbol{\theta}_{k}} \hat{J}(\hat{\boldsymbol{\Pi}}_{k}) = \sum_{n \in \mathcal{I}_{k}} \nabla_{\boldsymbol{\theta}_{n}} \log \left(\hat{\boldsymbol{\Pi}}_{k}(\boldsymbol{a}_{n} | \boldsymbol{o}_{n}; \boldsymbol{\theta}_{k}) \right) \left(\frac{1}{|\mathcal{I}_{k}|} \sum_{n' \in \mathcal{I}_{k}} \delta_{k,n', \boldsymbol{a}_{n'}} \right), \tag{11}$$

where we neglect the time indexing for notational simplicity. Afterwards, we update our agents policies using the gradient ascent method:

$$\theta_k \leftarrow \theta_k + \eta \nabla_{\theta_k} \hat{J}(\hat{\Pi}_k)$$

0

for some learning rate $\eta \in (0, 1)$. We term this method as *Dec-Relational* A2C. By following this method we are able achieve fully decentralized training of our network (Fig. 4(b) illustrates this procedure). The steps of this algorithm are summarized in Algorithm 2 below:

Algorithm 2 The Dec-Relational A2C Algorithm for Simultaneously **Optimize Routing Strategy**

1: Initialize Learning Rates η , α . 2: for Base Station k = 0, 1, ..., K - 1 do 3: Initialize Actor and Critic weights θ_k , \mathbf{w}_k . 4: end for 5: **for** time step t = 1, 2, ..., T **do** for Base Station $k = 0, 1, \dots, K - 1$ do 6: Extract active agents $\mathcal{I}_{k,t}$. 7: for Agent $n \in \mathcal{I}_{k,t}$ do 8: 9. Observe $\mathbf{o}_{t,n}$. $a_{t,n} = a, \text{w.p. } \hat{\Pi}_k(a | \mathbf{o}_{t,n}; \boldsymbol{\theta}_k), \forall a \in \mathcal{A}_{n,t},$ 10: 11: end for 12: end for 13: Execute actions. 14: for Base Station $k = 0, 1, \dots, K - 1$ do 15: for Agent $n \in I_{k,t}$ do 16: Obtain the delay $D_{n,t+1}$ and next observation $\mathbf{o}_{t+1,n}$ associated with the n^{th} agent. 17: Set Temporal Difference Error. $\delta_{t,n} = D_{n,t+1} + \gamma \cdot \hat{V}(\mathbf{o}_{t+1,n}; \mathbf{w}_{a_{t,n}}) - \hat{V}(\mathbf{o}_{t,n}; \mathbf{w}_k).$ 18: 19: end for 20: end for 21: Update critics' and actors' parameters based on Eq. (10) and Eq. (11).

22: **for** Base Station
$$k = 0, 1, ..., K - 1$$
 do

23:
$$\theta_{t+1,k} \leftarrow \theta_{t,k} + \eta \cdot \left(\sum_{n \in I_{k,t}} \nabla_{\theta_k} \log(\hat{\Pi}_k(a_{t,n} | \mathbf{o}_{t,n}; \theta_{t,k})) \right)$$

$$\left(\frac{1}{|I_{k,t}|}\sum_{n'\in I_{k,t}}\delta_{t,n'}\right)\right).$$
24: $\boldsymbol{w}_{t+1,k} \leftarrow \boldsymbol{w}_{t,k} - \alpha \cdot \nabla_{\boldsymbol{w}_k}\left(\frac{1}{|I_{k,t}|}\sum_{n\in I_{k,t}}(\delta_{t,n})^2\right)$
25: end for
26: end for

6.3. Fed-Relational A2C

To conclude, we propose another version that combines the features of the previous versions. As part of this approach, the weights of the network are updated using a federated learning approach [37]. Therefore, we refer to this method as Fed-Relational A2C. In this approach

2

2

Table 1

0.1	··· · · · · · · · · · · · · · · · · ·	1	
Similiation	nerwork	nvper-parameters	

$ \mathcal{D} $	$ \mathcal{B} $	$ \mathcal{U} $	Ν	P ^{max} parent	$C_{children}^{max}$	$U_{children}^{max}$	U_{parent}^{max}	TTL	User's speed $\left[\frac{m}{sec}\right]$	С
1	9	100	300 000	3	3	35	2	50	3	1

Table 2 Algorithm's hy	per-parameters.				
γ	α	η	ϵ_D	ϵ_{min}	Federated update frequency [Time-slots]
0.995	0.0001	0.0001	0.9999	0.01	1000

the agents constantly update their weights in a decentralized manner, similar to the decentralized approach, while the agents periodically report their weights to a central controller; based on the weights they submit, the controller calculates a new set of shared weights. The final step is to relay the updated weights to the base stations. Thus, we find this method to be somewhere in the middle between the previous centralized and decentralized approaches. Additionally, in this approach, weights are shared periodically among the different base stations, resulting in the same observations as in the relational approach. The steps of this algorithm are summarized in Algorithm 3 below:

Algorithm 3 The Federated-Relational Advantage Actor Critic (Fed-Relational A2C) Algorithm for Simultaneously Optimize Routing Strategy

- 1: Initialize Actor and Critic weights and learning rates θ , w, η , α .
- 2: Initialize federated update period τ .
- 3: for BS k = 0, 1, ..., K 1 do \triangleright Broadcast Actor and Critic weights. 4: $\mathbf{w}_k \leftarrow \mathbf{w}, \ \theta_k \leftarrow \theta$.
- 5: end for
- 6: **for** time step t = 1, 2, ..., T **do**
- 7: **if** *t* MOD $\tau == 0$ **then**
- 8: Collect Actor and Critic weights from each base-station $\{\mathbf{w}_k\}_{k=0}^{K-1}, \{\theta_k\}_{k=0}^{K-1}$ with the number of updates since last global update $\{n_k\}_{k=0}^{K-1}$.

9: FEDERATED UPDATE $(\{\mathbf{w}_k\}, \{\boldsymbol{\theta}_k\}_{k=0}^{K-1}, K, \{n_k\}_{k=0}^{K-1}).$

- 10: end if
- Follow decision and training phases of Decentralized Algorithm
 Lines 6-25.
- 12: end for
- 13: **procedure** FEDERATEDUPDATE(CriticWeights,ActorWeights,K,Number OfUpdates)
- 14: $\{\mathbf{w}_k\}_{k=0}^{K-1} \leftarrow$ CriticWeights, \triangleright Unpack the Critic Weights.15: $\{\theta_k\}_{k=0}^{K-1} \leftarrow$ ActorWeights, \triangleright Unpack the Actor Weights.16: $\{n_k\}_{k=0}^{K-1} \leftarrow$ NumberOfUpdates
number. \triangleright Unpack each BS's update
- 17: $\mathbf{w} \leftarrow \sum_{k=0}^{K-1} \frac{n_k}{\sum_{k=0}^{K-1} n_k} \mathbf{w}_k, \ \boldsymbol{\theta} \leftarrow \sum_{k=0}^{K-1} \frac{n_k}{\sum_{k=0}^{K-1} n_k} \boldsymbol{\theta}_k, \quad \triangleright \text{ Apply FedAvg Update Rule.}$
- 18: **for** BS k = 0, 1, ..., K 1 **do** \triangleright Broadcast the new calculated weights to the agents
- 19: $\mathbf{w}_k \leftarrow \mathbf{w}, \, \boldsymbol{\theta}_k \leftarrow \boldsymbol{\theta}.$
- 20: end for
- 21: end procedure

It should be mentioned that in a general POMDP setting solved using MARL techniques, as considered here, a common problem is multiple local optimum points [39] within the joint policy space, which may be resolved with convergence to a less desirable, local optima strategy solution [39]. As a result, convergence to the optimal policy is not guaranteed theoretically. In practice, however, it achieves very good performance even in various POMDP models with infinitely large state space. For example, the work in [40] developed an Actor Critic algorithm for teaching multiple agents how to play Starcraft games directly from screen images, and achieved very good performance at various stages.

Fig. 7. Network topology illustration.

7. Experiments and insights

In this section, we describe our main research insights and their associated experiments. First, we evaluate several connectivity scenarios in order to demonstrate the importance of network routing in an IAB network. Next, we study the impact of network load as well as the impact of mobility. This analysis shows that high loads affect network performance significantly in the case of mobility for most of the routing algorithms. We then analyze how changes in online traffic patterns and node failure affect the network. As a result of this experiment, we gained valuable insights into how the proposed routing can adapt to and recover from online changes. Lastly, we analyzed the different routing convergence times

As part of these experiments, we study and evaluate various routing methods within an IAB network, as discussed in Section 5. In particular, the performance of Relational A2C is compared to six other algorithms: Centralized Routing, Minimum-Hop Routing, Back-pressure Routing, Q-Routing, Full Echo Q-Routing, and Hybrid Routing. We refer the reader to Section 5 for a detailed explanation of each algorithm.

To conduct those experiments, we have developed a gym-based simulated IAB environment [41]. The simulation takes place over a 2-dimensional grid.¹ Tables 1,2 describe network and algorithms hyper parameters, respectively. Furthermore, Relational A2C was implemented as described in Algorithms 1,2,3 in Section 6, with all methods based on the A2C network architecture shown in Fig. 6. In the following, all the metrics we mentioned in Section 4 are used as the figure-of-merit for evaluating the performance of the different algorithms.

Fig. 7. illustrates a network composed with 1 IAB Donor, 9 IAB Nodes and 100 users.

¹ For reproducing our results, we refer the reader to run our code, https: //github.com/Shahaf-Yamin/Routing-In-IAB-Networks.

Fig. 8. Visualization of different connectivity scenarios.

Fig. 9. Performance illustration of different network's connectivity cases.

7.1. The importance of routing in an IAB network

In this section, we examine the importance of routing in an IAB network by measuring network performance in various scenarios of connectivity. Our study emphasizes the importance of routing algorithms by illustrating that for high connectivity, where routing is needed, performance is much higher than for low connectivity, where routing is not needed due to the limited number of paths. In order to change the network's connectivity we have modified the constraints that dictate the number of parents each IAB node/User may have and the number of devices (IAB children/Users) that each IAB can support.

Using the results illustrated in Fig. 9, we can conclude that using a single path topology (Fig. 8.a.) yielded the lowest arrival ratio and the longest delay. We can also deduce that higher network connectivity is highly recommended to support the expected 5G's requirements of high load and low latency. Moreover, while higher connectivity may have some benefits, we must also consider that a higher level of connectivity may result in more interference between nearby base stations, which may contradict our assumption regarding interference. In light of this conclusion, the purpose of this study is to examine how using routing can improve network performance, while only taking partial connectivity into account.

7.2. The influence of network load

This section examines how network load affects the performance of different algorithms in various mobility scenarios. That is, we evaluate each algorithm's resilience under different loads for various mobility scenarios. Following this, we conclude that an increase in load/mobility affects network performance significantly for most of the proposed routing algorithms. As a next step, we will describe in detail our experiments and their corresponding empirical results.

To change the network load, the parameter λ of the Poisson distribution has been modified. This parameter indicates the average number of packets generated in each time-slot by the IABs. To modify λ , we scanned various loads successively from bottom-to-top, and then from top-to-bottom. The results presented are an average of 10 different measurements for each load across five different network topologies.

First, we evaluate the performance of our algorithms over a static topology, which means that users cannot change their base-station association or location. Next, we increase the user's speed to $3_{\frac{m}{sec}}$, which means that users can change their location and change their association with a base station at any time. We then refer to this scenario as a dynamic topology and evaluate our algorithms' performance. In addition, we investigate how the speed of UEs affects the performance of the different algorithms. Furthermore, due to the lack of further information provided by the arrival ratio, we present only the average delay metric. The following sections provide results and insights from those experiments.

7.2.1. Static topology

This experiment evaluated the algorithm's performance under a static network topology and changing load. In Fig. 10(a), we present the performance of the algorithms. Based on these results, we can conclude that all three versions of the Relational A2C algorithm outperformed other algorithms, despite acting in a decentralized manner. Moreover, these results also revealed that using exploration in conjunction with exploitation (Full-Echo Q-Routing) greatly improved performance at low to medium loads when compared with Q-Routing, whilst they achieved similar performances for higher loads.

As a further interesting phenomenon, we observe that Hybrid Routing improves performance when compared with traditional Q-Routing under low loads, but degrades performance under higher loads. We

Fig. 10. Static topology and dynamic topology — average delay performance for different routing algorithms under different loads.

explain this phenomenon by stating that higher loads require better coordination among agents, as making an incorrect routing decision while the network is already congested will probably result in packet loss. Our conclusion is that since each agent learns an independent stochastic policy as part of the algorithm, it is less effective at handling non-stationary problems than traditional Q-Routing algorithms. In our next experiment, we allow UEs to travel and change their base-station association, which will further increase the non-stationarity of the problem.

Additionally, as network load increases, we observed performance degradation of the Shortest-Path algorithm. Our explanation for this phenomenon is that queue delay does not have a significant influence on performance for low loads, but as the load increases, it becomes more significant.

7.2.2. Dynamic topology

This experiment evaluated the algorithm's performance under a dynamic network topology and changing load. Following this, based on the results illustrated in Fig. 10 we can determine that although acting in a decentralized manner, all versions of Relational A2C's algorithm managed to achieve superior performance than the other algorithms for this scenario. Further, Hybrid Routing does not perform well in this non-stationary environment, as we observe further degradation in performance when compared to a static scenario, and at low load it is even less effective than Q-Routing.

Moreover, when compared with the static scenario, Full-Echo Q-Routing exhibits performance degradation at medium to high loads. It performs worse in this range than the traditional Q-Routing. Our explanation for this phenomenon is that when using Full Echo Q-Routing, frequent topology changes may cause multiple modifications to the routing policy, which will result in increased instability and longer routing times as loads increase. In addition, this phenomenon coincides with the phenomena observed in [7], in which the authors applied those algorithms to a grid topology.

Furthermore, it is worth noting that the Shortest Path Routing and Centralized Routing algorithms, operating in a centralized manner with access to complete information about user equipment (UE) connections, demonstrates robustness and consistency in its performance despite UE mobility. By considering the entire network topology, the algorithm can effectively adapt to changes and reroute traffic accordingly. Similarly, the Maximum Backpressure algorithm aims to optimize queue lengths between base stations to maximize throughput. Although user mobility can indirectly influence queue lengths due to changes in traffic patterns, the algorithm's primary objective is to manage the dynamics of the queue system and enhance throughput. As a result, the performance of the Maximum Backpressure algorithm is less directly affected by user mobility compared to Full-Echo Q-Routing.

The significant performance gap between all versions of Relational A2C and the traditional hybrid algorithm supports our claim that by following our proposed algorithms, agents are able to cooperate more effectively and increase the stability of our routing policy, proving that agents that cooperate achieve better outcomes than agents who act selfishly. Thus, these results demonstrate that optimizing the joint goal and solving the coordination problem between agents significantly improves performance for both static and dynamic topologies.

When comparing the performance of static versus dynamic topologies as illustrated in Fig. 10, we observe that all tabular RL baselines suffer from performance degradation when users are permitted to move and change their base-station association. Based on these results, which indicate a significant difference between static and dynamic topologies, we propose the following experiment to examine the effect of UE movement on algorithm performance.

7.2.3. Results from examining the influence of dynamic topology changes

In this experiment, we evaluated the algorithm's robustness to dynamic topology changes under constant loads. For this purpose, we scanned various UE speeds and measured the arrival ratio and average delay with the different algorithms. We evaluate the algorithms' performance for two possible loads, medium load ($\lambda = 3$) and high load ($\lambda = 5$), presented in Figs. 11(a) and 11(b), respectively.

Figs. 11(a) and 11(b) demonstrate that increasing the UE speed will not have a significant impact on the Relational A2C algorithm's performance for varying network loads. Also, in medium loads, Full Echo Q-Routing achieves superior performance to Q-Routing, with both suffering from similar performance degradation as a result of their different UE speeds. Furthermore, we observe that the combination of higher load with increased speed results in more severe performance degradation for Full Echo Q-Routing compared to traditional Q-Routing. This observation is consistent with the results of our previous experiments. Moreover, Hybrid Routing shows the greatest degradation among all Q-Routing algorithms, further demonstrating that it is less able to cope with this non-stationary setting than traditional Q-Routing algorithms. Aside from this, we observe that all other RL-based algorithms suffer from performance degradation when UE speed is increased, which serves as an additional indication of our proposed approach's superiority.

7.3. Experiment results for online changes

The purpose of this section is to examine how online changes impact algorithm performance in a variety of scenarios. Our study indicates

Fig. 11. Average network delay for different routing algorithms under different UE speed.

Fig. 12. Static topology and dynamic topology — average delay performance for different routing algorithms under bursts of traffic. low load — $\lambda = 2$, high load — $\lambda = 5$.

that Relational A2C-based approaches are superior for handling such online changes when compared to other algorithms. As a next step, we will describe the online scenarios that we have studied. In the first case, we analyzed the algorithm's response to bursts of traffic. In the second case, we investigate how the algorithm responds to node failures and recovery situations. Also, in these experiments we capture those online changes by measuring the algorithm's performance using a sliding window with a length of 100 time slots.

7.3.1. Experiment results for bursts of traffic

During this experiment we measure the algorithm's response to bursts of traffic. This was accomplished by evaluating the algorithm performance while changing the network load. Our experiment involves changing the network load rapidly from low ($\lambda = 2$) to high ($\lambda = 5$) and then back to low. The results presented here are an average of 10 measurements across five different topologies of the network. Furthermore, we conducted this experiment for both static and dynamic topologies, as depicted in Figs. 12(a) and 12(b), respectively. Further, due to the absence of further information offered by the arrival ratio, we present only the average delay metric measured through time for this experiment.

In both static and dynamic scenarios, it is evident that the Relational A2C versions achieve superior performance than the traditional RL-algorithms. Additionally, all algorithms exhibit similar reaction times when measuring the impact of changes in load from low to high

and from high to low. When comparing between static and dynamic scenarios, we are able to see that the Relational A2C algorithms did not suffer from any performance degradation, while the tabular methods greatly suffer.

Further, we observe that Full Echo Q-Routing is superior to traditional Q-Routing in static scenarios, while they achieve similar performance in dynamic scenarios. Additionally, Hybrid Routing suffers from the greatest degradation of all Q-Routing algorithms when working with high loads, emphasizing its inability to cope with such a nonstationary scenario. Our conclusions are consistent with the results of our previous experiments, which further verifies the results.

7.3.2. Experiment results for node failure

In this experiment, we evaluate how the algorithms respond to a scenario of node failure and recovery. To achieve this, the algorithm performance was evaluated while removing a random base station from the network for a specified period of time. Since we insert dynamic changes through node failures, we now consider only static topologies, in which users cannot move. The following results are based on an average of 30 consecutive experiments with random base-station failure in each experiment.

This experiment was conducted for a high load as shown in Fig. 13. To illustrate the arrival ratio metric in such an online setting, we measured the number of packets that dropped within our sliding window. From these results, it is evident that all algorithms are able to recover

(a) Average Delay Through Time.

(b) Dropped Packets Through Time.

Fig. 13. Performance for different routing algorithms in case of node failure for high load $\lambda = 5$.

Fig. 14. Illustration of the performance of different routing algorithms through their training procedures.

from the node failure situation, although we observe that the average delay does not return to its initial value after the node has recovered. In our opinion, this is due to the fact that the other base stations' queues are already congested at this point. This results in a more congested network than before the node failure. In conjunction with the fact that the arrival rate process does not change over time, performance is degraded under higher loads.

A further, but equally significant, conclusion we can draw from these measurements is that while Relational A2C has a lower delay than Centralized Routing in Fig. 13(a), its packet loss rate is higher in Fig. 13(b). This demonstrates the importance of using both arrival rate and average delay in our analysis.

7.4. Algorithm convergence

In this section, we examine the algorithm's convergence time under a constant load, starting with a random initialization point. According to our results, all versions of our algorithm were able to converge to a similar solution, as opposed to the other baselines that suffer from performance degradation. Thus, we conclude that all Relational A2Cbased algorithms were able to achieve a better routing solution than traditional algorithms. The results presented in Fig. 14 are the average of those measurements.

Additionally, we found that centralized training achieved the fastest convergence among the different training paradigms of our proposed method. In fact, this convergence gap arises because the different agents can share their experience implicitly through the mutual updates of both the actor and the critic among all base stations. Furthermore, we conclude that Fed-Relational A2C performs better than Dec-Relational A2C in terms of convergence. In general, this difference can be understood intuitively, since federated learning approaches manage the trade-off between a fully decentralized training paradigm and a centralized training paradigm. Another interesting phenomenon is that Full Echo Q-Routing appears to be the fastest algorithm that converges to a stable solution (an approximated point of equilibrium). This rapid convergence can be explained by the fact that Full Echo Q-Routing receives all of the rewards available to it, regardless of the chosen action. As a result, it is able to reduce the number of interactions required for convergence with the environment.

7.5. Algorithm scalability

In this section, we examine the algorithm's performance for different number of base-stations for two possible loads, medium load ($\lambda = 3$) and high load ($\lambda = 5$), presented in Figs. 15(a) and 15(b), respectively. The following results are based on an average of 5 consecutive experiments with different topology in each experiment. Figs. 15(a) and 15(b) provide insightful evidence regarding the impact of increasing the number of base-stations on network performance. The results reveal that for medium loads, a higher number of base-stations may not necessarily lead to improved network performance. However, under higher loads, increasing the number of base-stations can result in a significant improvement in performance. These findings suggest a nonlinear relationship between the number of base-stations and network performance, with the effectiveness depending on the specific load conditions.

Upon comparing the performance of various algorithms, several key observations have emerged. Firstly, both the Fed-Relational A2C and Dec-Relational A2C algorithms exhibit similar performances across all scenarios. However, there is a slight degradation in their performance compared to the Relational A2C algorithm, particularly when using 20 base-stations. This can be attributed to the fact that in the Relational A2C algorithm, agents share the same actor and critic, facilitating the exchange of knowledge through mutual gradient updates.

Furthermore, we observed that the Hybrid routing algorithm experiences the most significant performance degradation as the number of base-stations increases. This degradation can be attributed to the dynamic nature of the topology, where user equipment (UE) mobility leads to frequent changes in base-station associations. The dynamic

Relational-A2C

Dec-Belational-A2C

(a) Average Delay Through Time For Medium Load $\lambda = 3$.

(b) Average Delay Through Time For High Load $\lambda = 5$.

Fig. 15. Performance for different routing algorithms in case of different number of base-stations.

120

changes introduced by increasing the network size pose challenges in updating the routing policy effectively in an online manner, thus impacting performance.

Additionally, both the Q-Routing and Full-Echo Q-Routing algorithms demonstrate high variance and a strong dependency on the number of base-stations. For example, at high loads with only 5 basestations, Full-Echo Q-Routing exhibits the worst performance. Although increasing the number of base-stations improves its performance, Q-Routing fails to achieve the performance levels attained by Relational A2C and similar algorithms at medium loads.

These findings provide insights into the strengths and weaknesses of each algorithm and emphasize the effectiveness and robustness of the Relational A2C algorithm in various scenarios. By optimizing routing strategies in IAB networks, these insights contribute to the ongoing efforts of enhancing network performance and efficiency.

8. Conclusion

In this paper, we have investigated the problem of routing in an IAB network, where multiple IAB nodes operate concurrently to efficiently route packets while avoiding network congestion. Our objective was to develop a joint routing policy that maximizes the network's arrival ratio and minimizes the average packet delay. To achieve this, we proposed the novel Relational A2C algorithm, which utilizes online learning and observations from IABs to determine the optimal joint routing strategy.

Through extensive evaluations and comparisons with other algorithms, we demonstrated that the Relational A2C algorithm outperforms baseline approaches in various scenarios. Its performance is comparable to a centralized approach, showcasing its ability to learn nearcentralized policies. This highlights the effectiveness and superiority of our proposed approach over existing methods.

Moreover, our study emphasized the crucial role of network routing in mitigating congestion and maximizing the utilization of network resources in IAB networks. By developing an efficient joint routing policy, we can significantly enhance the overall network performance and improve the delivery of packets to their destinations.

In summary, this paper contributes to the field of IAB network routing by introducing the Relational A2C algorithm and showcasing its effectiveness. By optimizing joint routing policies, we can improve network performance, and future research can focus on incorporating robust security measures for practical deployment.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

Acknowledgments

This work was supported in part by the German Research Foundation (DFG) via the German Israeli Project Cooperation (DIP), by the Israeli Science Foundation (ISF) under Grant 899/21, by the Minister of Defense, and by Israeli Innovation Authority as part of the Worldwide Innovative Networking (WIN) Consortium.

Appendix A. MDP formulation

Under the assumption that each packet is an individual agent, this section introduces the proof that for the multiple packets scenario, we can formulate the problem of IAB network routing as an MDP with multiple agents. Table A.3. summarizes the Multi-Agent MDP formulation, while Lemmas 3, 4 prove that it is indeed a Multi-Agent MDP. To formulate the problem as a Multi-Agent MDP we shall define the following tuple: $\langle S, A, R, P, \gamma, N \rangle$, where N represents the maximal number of packets (agents) our system supports.

Under this formulation, as illustrated in Fig. 16, each agent has four available modes. The following explains the meaning of each mode,

- Of f In this mode, the packet is not a part of the network at the moment, either because it has already been delivered to its destination, or suffered from a TTL expiration event, or because it is awaiting injection into the network.
- *Queue* In this mode, the packet is waiting in a specific basestation queue.
- *Decision* In this mode, the packet is currently required to make a next-hop decision.
- *Air* In this mode, the packet is transferred over the air from one node to another in the network.

We define our reward as follows:

$$r_t(\mathbf{s}_t, \mathbf{a}_t) = -\sum_{i=0}^{N-1} \left(ElapsedTime_{t,i} + d((n_{t,i}, a_{t,i})) \right) \cdot \mathbb{1}_{[s_{t,i,3} = Decision]},$$
(12)

specifically, $ElapsedTime_{t,i}$ represents the period of agent *i* waiting at node $n_{t,i}$ queue before transmission and $d((n_{t,i}, a_{t,i}))$ represents the transmission delay between node $n_{t,i}$ and node $a_{t,i}$ at time step $t \in \mathbb{N}$. The last term, $\mathbb{1}_{\{s_{t,i}\}==Decision\}}$, represents if agent *i* has permission to conduct a wireless hop or not. Next, we provide the full proof that this formulation is indeed a Multi-Agent MDP.

MDP formulation.	
Туре	Representation
Action Space of agent i at time t	$\mathcal{A}_{i}^{t} = \left((\mathbf{s}_{t,i,3} = Decision) \rightarrow \left\{ a \forall a \in \mathcal{N}, (n_{t,i}, a) \in \mathcal{L} \right\} \right) \land \emptyset$
Action Space at time t	$\mathcal{A}_{t} = X_{i=0}^{N-1} \mathcal{A}_{t,i}$
State	$\mathbf{s}_{t} = \bigcup_{i=0}^{N-1} \mathbf{s}_{t,i},$ where $\mathbf{s}_{t,i} = \{n_{t,i}, ttl_{t,i}, ElapsedTime_{t,i}, AgentMode_{t,i}\}$
Action	$\mathbf{a}_t \in \mathcal{A}^t$
Reward	$-\sum_{i=0}^{N-1} \left(ElapsedTime_{t,i} + d((n_{t,i}, a_{t,i})) \right) \cdot \mathbb{1}_{[s_{t_{i,j,3}}=1]}$

Fig. 16. Different modes of each agent.

Lemma 3. Given the history \mathcal{H}_t the reward distribution depends only on s_t and \mathbf{a}_t . Thus, let $\Pr : S \times \mathcal{A} \times \mathbb{Z} \to [0, 1]$ be the agent's reward distribution.

T-11- A O

Proof.

$$\Pr(r_t|\mathcal{H}_t) \stackrel{(a)}{=} \Pr(r_t|\mathbf{s}^t, \mathbf{a}^t, r^{t-1}) \stackrel{(b)}{=} \Pr(r_t|\mathbf{s}_t, \mathbf{a}_t).$$
(13)

(*a*) is true straight from the definition of history. Further, it follows that (*b*) is true since the reward does not depend on the previous history once the action and state are determined. This claim is true since it is composed of $ElapsedTime_{t,i}$, which is a part of the current state, and $d((n_{t,i}, a_{t,i}))$, which is a random variable that is dependent only on the current state and action. \Box

Lemma 4. Given the history \mathcal{H}_t the transition probability matrix \Pr depends only on \mathbf{s}_t and \mathbf{a}_t . Thus, let $\Pr : S \times A \times S \rightarrow [0, 1]$ be the transition probability matrix.

Proof. We begin by defining the following probabilities, Whenever an agent reaches a TTL expiration event, its mode will transition to Off:

$$\Pr\left(\mathbf{s}_{t+1,i} = (\cdot, \cdot, \cdot, Of f) | \mathbf{s}_t, \mathbf{a}_t, s_{t,i,1} = 0\right) = 1.$$
(14)

Next, we deal with the Air mode transition probabilities:

$$\Pr\left(\mathbf{s}_{t+1,i} = (a_t, ttl_{t,i} - 1, d((n_{t,i}, a_{t,i})), Air) | \mathbf{s}_t, \mathbf{a}_t, s_{t,i,3} = Decision\right) = 1.$$
(15)

$$\Pr\left(\mathbf{s}_{t+1,i} = (n_{t,i}, ttl_{t,i} - 1, ElapsedTime_{t,i} - 1, Air)|\mathbf{s}_{t}, \mathbf{a}_{t}, s_{t,i,3} = Air, s_{t,i,2} > 0\right)$$

= 1. (16)

$$\Pr\left(\mathbf{s}_{t+1,i} = (n_{t,i}, ttl_{t,i} - 1, 1, Queue) | \mathbf{s}_t, \mathbf{a}_t, s_{t,i,3} = Air, s_{t,i,2} = 0\right) = 1.$$
(17)

Next, we explain why the current state combined with agent action provides sufficient information regarding queue transition probabilities, so the history condition does not matter. In the current state, the agent knows both the TTL values and the locations of all the agents. A prioritized queue, in combination with its state and action, enables the agent to determine which agents will transition from queue to decision mode next.

Last, the injection process of new agents in our system, which is based over sampling an i.i.d. random process, is independent at the agent history. Thus, we can define our transition probabilities as follows:

$$\Pr(\mathbf{s}_{t+1}|\mathcal{H}_t) \stackrel{(a)}{=} \Pr(\mathbf{s}_{t+1}|\mathbf{s}^t, \mathbf{a}^t, r^{t-1}) \stackrel{(b)}{=} \Pr(\mathbf{s}_{t+1}|\mathbf{s}_t, \mathbf{a}_t).$$

Equality (*a*) is from the definition of history. Equality (*b*) is due to the fact that given that you are at state s_t and you have chose action a_t , the previous states that you have visited have no influence over your next transition probabilities, as we have just shown for each agent mode. \Box

Appendix B. Proof of Lemma 1

Proof. Let $\tau(T) \triangleq (\mathbf{s}_0, \mathbf{o}_0, \mathbf{a}_0, r_1, \dots, \mathbf{s}_{T-1}, \mathbf{o}_{T-1}, \mathbf{a}_n, r_T, \mathbf{s}_T, \mathbf{o}_T)$ represent a possible trajectory with length *T*. In addition, let $\Pr(\tau(T)) \triangleq \Pr(\mathbf{s}_0, \mathbf{o}_0, \mathbf{a}_0, r_1, \dots, \mathbf{s}_{T-1}, \mathbf{o}_{T-1}, \mathbf{a}_{T-1}, r_T, \mathbf{s}_T, \mathbf{o}_T)$, by applying the chain rule and Lemmas 3, 4, we receive $\Pr(\tau(T)) = \Pr(\mathbf{s}_0)\Pi_{t=0}^{T-1} (\Pi_{n=0}^{N-1}\Pi(a_{t,n}|\mathbf{o}_{t,n}; \theta))$ $\Pr(\mathbf{o}_t|\mathbf{s}_t)\Pr(\mathbf{s}_{t+1}, r_{t+1}|\mathbf{s}_t, \mathbf{a}_t)$. Thus, we get

$$\nabla_{\boldsymbol{\theta}} \log \Pr(\tau(T)) = \sum_{t=0}^{T-1} \sum_{n=0}^{N-1} \nabla_{\boldsymbol{\theta}} \log \Pi(\boldsymbol{a}_{t,n} | \boldsymbol{o}_{t,n}; \boldsymbol{\theta}).$$
(18)

Let us recall our goal, which is that we would like to solve the following optimization problem:

$$\Pi^{\star} = \arg \max_{\Pi} J(\Pi) = \arg \max_{\Pi} \underbrace{\mathbb{E}[\sum_{t=0}^{T-1} \gamma^{t} \cdot r_{t+1} | \Pi]}_{J(\Pi)}.$$

Following that, we can now calculate the gradient of the objective $J(\Pi)$ w.r.t. the policy parameters θ :

$$\begin{split} \nabla_{\theta} J(\theta) \\ &= \nabla_{\theta} \mathbb{E} \bigg[\frac{1}{N} \sum_{n=0}^{N-1} G_n \bigg] \stackrel{(a)}{=} \sum_{\tau(T)} \nabla_{\theta} \operatorname{Pr}(\tau(T)) \cdot \frac{1}{N} \sum_{t=0}^{T-1} \gamma^t \cdot r_{t+1} \\ &= \sum_{\tau(T)} \nabla_{\theta} \operatorname{Pr}(\tau(T)) \cdot \frac{\operatorname{Pr}(\tau(T))}{\operatorname{Pr}(\tau(T))} \cdot \sum_{t=0}^{T-1} \gamma^t \cdot r_{t+1} \\ &= \sum_{\tau(T)} \operatorname{Pr}(\tau(T)) \nabla_{\theta} \log \operatorname{Pr}(\tau(T)) \cdot \sum_{t=0}^{T-1} \gamma^t \cdot r_{t+1} \end{split}$$

$$\begin{split} \stackrel{(b)}{=} & \sum_{\tau(T)} \Pr(\tau(T)) \bigg(\sum_{l=0}^{T-1} \sum_{n'=0}^{N-1} \nabla_{\theta} \log \Pi(a_{t,n'} | \mathbf{o}_{t,n'}; \theta) \bigg) \cdot \sum_{l'=0}^{T-1} \gamma^{l'} \cdot r_{l'+1} \\ \stackrel{(c)}{=} & \mathbb{E} \bigg[\sum_{t=0}^{T-1} \bigg(\sum_{n' \in I_{t}}^{N-1} \nabla_{\theta} \log \Pi(a_{t,n'} | \mathbf{o}_{t,n'}; \theta) \bigg) \cdot \bigg(\sum_{l'=0}^{T-1} \gamma^{l'} \cdot r_{l'+1} \bigg) \bigg] \\ \stackrel{(d)}{=} & \mathbb{E} \bigg[\sum_{t=0}^{T-1} \bigg(\sum_{n' \in I_{t}}^{N-1} \nabla_{\theta} \log \Pi(a_{t,n'} | \mathbf{o}_{t,n'}; \theta) \bigg) \cdot \bigg(\sum_{l'=t}^{T-1} \gamma^{l'} \cdot r_{l'+1} \bigg) \bigg] \\ \stackrel{(e)}{\cong} & \mathbb{E} \bigg[\bigg(\sum_{n' \in I_{t}}^{T-1} \nabla_{\theta} \log \Pi(a_{t,n'} | \mathbf{o}_{t,n'}; \theta) \bigg) \cdot \bigg(\sum_{l'=t}^{T-1} \gamma^{l'} \cdot r_{l'+1} \bigg) \bigg] \\ \stackrel{(f)}{=} & \mathbb{E}_{s_{t} \sim \mu(s), \mathbf{o}_{t} \sim \Pr(\cdot | \mathbf{s}_{t}), \mathbf{a}_{t} \sim \Pi(\cdot | \mathbf{o}_{t})} \bigg[\mathbb{E} \bigg[\bigg(\sum_{n' \in I_{t}}^{T-1} \nabla_{\theta} \log \Pi(a_{t,n'} | \mathbf{o}_{t,n'}; \theta) \bigg) \\ \cdot \sum_{l'=t}^{T-1} \gamma^{l'} \cdot r_{l'+1} \bigg| \mathbf{s}_{t}, \mathbf{o}_{t}, \mathbf{a}_{t} \bigg] \bigg] \\ \stackrel{(g)}{\cong} & \mathbb{E}_{s_{t} \sim \mu(s), \mathbf{o}_{t} \sim \Pr(\cdot | \mathbf{s}_{t}), \mathbf{a}_{t} \sim \Pi(\cdot | \mathbf{o}_{t})} \bigg[\mathbb{E} \bigg[\sum_{l'=t}^{T-1} \gamma^{l'} \cdot r_{l'+1} \bigg| \mathbf{s}_{t}, \mathbf{a}_{t} \bigg] \bigg] \\ \stackrel{(g)}{\cong} & \mathbb{E}_{s_{t} \sim \mu(s), \mathbf{o}_{t} \sim \Pr(\cdot | \mathbf{s}_{t}), \mathbf{a}_{t} \sim \Pi(\cdot | \mathbf{o}_{t})} \bigg] \\ \stackrel{(g)}{=} & \mathbb{E}_{s_{t} \sim \mu(s), \mathbf{o}_{t} \sim \Pr(\cdot | \mathbf{s}_{t}), \mathbf{a}_{t} \sim \Pi(\cdot | \mathbf{o}_{t})} \bigg] \underbrace{ \sum_{Q_{II}(s_{t}, \mathbf{a}_{t})} \bigg] \\ \stackrel{(h)}{\longrightarrow} & \mathbb{E}_{s_{t} \sim \mu(s), \mathbf{o}_{t} \sim \Pr(\cdot | \mathbf{s}_{t}), \mathbf{a}_{t} \sim \Pi(\cdot | \mathbf{o}_{t})} \bigg] \\ \stackrel{(h)}{=} & \mathbb{E}_{s_{t} \sim \mu(s), \mathbf{o}_{t} \sim \Pr(\cdot | \mathbf{s}_{t}), \mathbf{a}_{t} \sim \Pi(\cdot | \mathbf{o}_{t})} \bigg] \\ \stackrel{(h)}{\longrightarrow} & \mathbb{E}_{s_{t} \sim \mu(s), \mathbf{o}_{t} \sim \Pr(\cdot | \mathbf{s}_{t}), \mathbf{a}_{t} \sim \Pi(\cdot | \mathbf{o}_{t})} \bigg] \\ \stackrel{(h)}{\longrightarrow} & \mathbb{E}_{s_{t} \sim \mu(s), \mathbf{o}_{t} \sim \Pr(\cdot | \mathbf{s}_{t}), \mathbf{a}_{t} \sim \Pi(\cdot | \mathbf{o}_{t})} \bigg] \\ \stackrel{(h)}{\longrightarrow} & \mathbb{E}_{s_{t} \sim \mu(s), \mathbf{o}_{t} \sim \Pr(\cdot | \mathbf{s}_{t}), \mathbf{a}_{t} \sim \Pi(\cdot | \mathbf{o}_{t})} \bigg] \\ \stackrel{(h)}{\longrightarrow} & \mathbb{E}_{s_{t} \sim \mu(s), \mathbf{o}_{t} \sim \Pr(\cdot | \mathbf{s}_{t}), \mathbf{a}_{t} \sim \Pi(\cdot | \mathbf{o}_{t})} \bigg] \\ \stackrel{(h)}{\longrightarrow} & \mathbb{E}_{s_{t} \sim \mu(s), \mathbf{o}_{t} \sim \Pr(\cdot | \mathbf{s}_{t}), \mathbf{a}_{t} \sim \Pi(\cdot | \mathbf{o}_{t})} \bigg] \\ \stackrel{(h)}{\longrightarrow} & \mathbb{E}_{s_{t} \sim \mu(s), \mathbf{o}_{t} \sim \Pr(\cdot | \mathbf{s}_{t}), \mathbf{a}_{t} \sim \Pi(\cdot | \mathbf{o}_{t})} \bigg]$$

Equality (a) follows from the definition of expectation combined with the linearity of both expectation and derivative (under the assumption that the trajectories set is independent of the policy parameters). Equality (*b*) follows from Eq. (18), and equality (*c*) follows from the definition of expectation. Next, equality (d) holds since the policy of agents that are not active at a certain time slot $t \in (0, ..., T-1)$ is independent of the policy parameters, and therefore, the corresponding derivative is equal to zero. The proportion relation (e) holds under the assumption of an MDP with a stationary distribution, which we denote with $\mu(s), s \in S$. Equality (f) holds according to the law of total expectation. Equality (g) holds since the derivative of the agent's policy is deterministic when considering the current state, observation and action, and therefore, it can be moved out of the inner expectation. Equality (h) follows from the definition of action-value function and from the fact that given the current state and action, the current observation is irrelevant to the expected return. Last, equality (i) is due to the fact that reducing a baseline function that is independent of the policy actions does not introduce any bias under the expectation [17, Ch. 13].

Appendix C. Proof of Lemma 2

Proof. Let $\tau(T) \triangleq (\mathbf{s}_0, \mathbf{o}_0, \mathbf{a}_0, r_1, \dots, \mathbf{s}_{T-1}, \mathbf{o}_{T-1}, \mathbf{a}_n, r_T, \mathbf{s}_T, \mathbf{o}_T)$ represent a possible trajectory with length *T*. In addition, we denote the joint parameters θ as follows: $\theta \triangleq \bigcup_{k=0}^{K-1} \theta_k$. Next, let $\Pr(\tau(T)) \triangleq \Pr(\mathbf{s}_0, \mathbf{o}_0, \mathbf{a}_0, r_1, \dots, \mathbf{s}_{T-1}, \mathbf{o}_{T-1}, \mathbf{a}_{T-1}, r_T, \mathbf{s}_T, \mathbf{o}_T)$, by applying the chain rule and Lemmas 3, 4; thus we receive $\Pr(\tau(T)) = \Pr(\mathbf{s}_0)\Pi_{t=0}^{T-1}(\Pi_{n=0}^{N-1}\Pi(a_{t,n}|\mathbf{o}_{t,n};\theta))$ $\Pr(\mathbf{o}_t|\mathbf{s}_t)\Pr(\mathbf{s}_{t+1}, r_{t+1}|\mathbf{s}_t, \mathbf{a}_t)$. Thus, we get

$$\nabla_{\theta_k} \log \Pr(\tau(T)) = \sum_{t=0}^{T-1} \sum_{n=0}^{N-1} \nabla_{\theta_k} \log \Pi(a_{t,n} | \mathbf{o}_{t,n}; \theta) \text{ for } \forall k \in (0, \dots, K-1).$$
(19)

Next, we derive the gradient of the joint objective $J(\Pi)$ w.r.t. $\theta_k.$

ν

$$\begin{split} \hat{f}_{\theta_{k}} J(\theta) &= \nabla_{\theta} \mathbb{E} \left[\frac{1}{N} \sum_{n=0}^{N-1} G_{n} \right]^{(\theta)} = \sum_{\tau(T)} \nabla_{\theta_{k}} \Pr(\tau(T)) \cdot \sum_{t=0}^{T-1} \gamma^{t} \cdot r_{t+1} \\ &= \sum_{\tau(T)} \nabla_{\theta_{k}} \Pr(\tau(T)) \cdot \frac{\Pr(\tau(T))}{\Pr(\tau(T))} \cdot \sum_{t=0}^{T-1} \gamma^{t} \cdot r_{t+1} \\ &= \sum_{\tau(T)} \Pr(\tau(T)) \nabla_{\theta_{k}} \log \Pr(\tau(T)) \cdot \sum_{t=0}^{T-1} \gamma^{t} \cdot r_{t+1} \\ &= \sum_{\tau(T)} \Pr(\tau(T)) \left(\sum_{t=0}^{T-1} \sum_{n'=0}^{N-1} \nabla_{\theta_{k}} \log \Pi(a_{t,n'} | \mathbf{o}_{t,n'}; \theta) \right) \sum_{t'=0}^{T-1} \gamma^{t'} \cdot r_{t'+1} \\ \begin{pmatrix} \psi \\ = \sum_{\tau(T)} \Pr(\tau(T)) \left(\sum_{t=0}^{T-1} \sum_{n'=0}^{N-1} \nabla_{\theta_{k}} \log \Pi(a_{t,n'} | \mathbf{o}_{t,n'}; \theta) \right) \cdot \left(\sum_{t'=0}^{T-1} \gamma^{t'} \cdot r_{t'+1} \right) \\ \begin{pmatrix} \psi \\ = \sum_{\tau(T)} \sum_{n'=0}^{T-1} \left(\sum_{n'\in I_{k,t}} \nabla_{\theta_{k}} \log \Pi(a_{t,n'} | \mathbf{o}_{t,n'}; \theta_{k}) \right) \cdot \left(\sum_{t'=t}^{T-1} \gamma^{t'} \cdot r_{t'+1} \right) \\ \end{pmatrix} \\ \begin{pmatrix} \psi \\ = \mathbb{E} \left[\left(\sum_{n'\in I_{k,t}} \nabla_{\theta_{k}} \log \Pi(a_{t,n'} | \mathbf{o}_{t,n'}; \theta_{k}) \right) \cdot \left(\sum_{t'=t}^{T-1} \gamma^{t'} \cdot r_{t'+1} \right) \right] \\ \begin{pmatrix} \psi \\ = \mathbb{E} \left[\left(\sum_{n'\in I_{k,t}} \nabla_{\theta_{k}} \log \Pi(a_{t,n'} | \mathbf{o}_{t,n'}; \theta_{k}) \right) \cdot \sum_{t'=t}^{T-1} \gamma^{t'} \cdot r_{t'+1} \\ \\ \end{pmatrix} \right] \\ \begin{pmatrix} \psi \\ = \mathbb{E} \left[\left(\sum_{n'\in I_{k,t}} \nabla_{\theta_{k}} \log \Pi(a_{t,n'} | \mathbf{o}_{t,n'}; \theta_{k}) \right) \left(\sum_{t'=t}^{T-1} \gamma^{t'} \cdot r_{t'+1} | \mathbf{s}_{t}, \mathbf{o}_{t}, \mathbf{a}_{t} \right) \right] \\ \begin{pmatrix} \psi \\ = \mathbb{E} \left[\sum_{t'\in I_{k,t}} \nabla_{\theta_{k}} \log \Pi(a_{t,n'} | \mathbf{o}_{t,n'}; \theta_{k}) \right] \\ \\ \begin{pmatrix} \psi \\ = \mathbb{E} \left[\sum_{n'\in I_{k,t}} \nabla_{\theta_{k}} \log \Pi(a_{t,n'} | \mathbf{o}_{t,n'}; \theta_{k}) \right] \left(\sum_{t'=t}^{T-1} \gamma^{t'} \cdot r_{t'+1} | \mathbf{s}_{t}, \mathbf{a}_{t} \right) \right] \\ \begin{pmatrix} \psi \\ = \mathbb{E} \left[\left(\sum_{n'\in I_{k,t}} \nabla_{\theta_{k}} \log \Pi(a_{t,n'} | \mathbf{o}_{t,n'}; \theta_{k}) \cdot Q_{\Pi}(\mathbf{s}_{t}, \mathbf{a}_{t} \right) \right] \\ \\ \begin{pmatrix} \psi \\ = \mathbb{E} \left[\left(\sum_{n'\in I_{k,t}} \nabla_{\theta_{k}} \log \Pi(a_{t,n'} | \mathbf{o}_{t,n'}; \theta_{k}) \cdot \left(Q_{\Pi}(\mathbf{s}_{t}, \mathbf{a}_{t} \right) \right] \\ \\ \\ \end{pmatrix} \\ \\ \end{pmatrix} \\ \end{pmatrix} \\ \end{pmatrix} \\ \end{pmatrix} \\ \end{pmatrix} \\ \end{pmatrix}$$

Equality (a) follows from the definition of expectation combined with the linearity of both the expectation and derivative (Under the assumption that the trajectories set is independent of the policy parameters). Equality (b) follows from Eq. (19), and equality (c) follows from the definition of expectation. Next, equality (d) holds since the policy of agents that are active in a certain time slot $t \in (0, ..., T - 1)$ is dependent on the policy parameters of the kth base-station only if they are acting within these base stations, and therefore, the remaining agent's derivative is equal to zero (the ones that are active outside of base-station k or the ones that are not active at all). The proportion relation (e) holds under the assumption of a MDP with a stationary distribution, which we denote with $\mu(s), s \in S$. The equality (*f*) holds according to the law of total expectation. Equality (g) holds since the derivative of the agent's policy is deterministic when considering the current state, observation and action, and therefore, it can be moved out of the inner expectation. Equality (h) follows from the definition of action-value function and from the fact that given the current state and action, the current observation is irrelevant to the expected return. Last, equality (i) is due to the fact that reducing a baseline function that is independent with policy actions does not introduce any bias under the expectation [17, Ch. 13].

S. Yamin and H.H. Permuter

References

- 3GPP, NR; NR and NG-RAN Overall Description; Stage 2, Technical Specification (TS), Vol. 38 no. 300, 2019.
- [2] B.D. Van Veen, K.M. Buckley, Beamforming: A versatile approach to spatial filtering, IEEE Assp Mag. 5 (2) (1988) 4–24.
- [3] 3GPP, NR; Study on Integrated Access and Backhaul, Technical Specification (TS), Vol. 38, no. 3874, 2019.
- [4] V. Konda, J. Tsitsiklis, Actor-critic algorithms, in: Advances in Neural Information Processing Systems, Vol. 12, 1999.
- R.J. Williams, Simple statistical gradient-following algorithms for connectionist reinforcement learning, Mach. Learn. 8 (3–4) (1992) 229–256, http://dx.doi.org/ 10.1007/BF00992696.
- [6] G. Tesauro, et al., Temporal difference learning and TD-gammon, Commun. ACM 38 (3) (1995) 58–68.
- [7] J.A. Boyan, M.L. Littman, Packet routing in dynamically changing networks: A reinforcement learning approach, in: Advances in Neural Information Processing Systems, 1994, pp. 671–678.
- [8] S. Zeng, X. Xu, Y. Chen, Multi-agent reinforcement learning for adaptive routing: A hybrid method using eligibility traces, in: 2020 IEEE 16th International Conference on Control & Automation, ICCA, 2020, pp. 1332–1339, http://dx. doi.org/10.1109/ICCA51439.2020.9264518.
- [9] R. Bellman, On a routing problem, Q. Appl. Math. 16 (1) (1958) 87-90.
- [10] L. Tassiulas, A. Ephremides, Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks, IEEE Trans. Automat. Control 37 (12) (1992) 1936–1948, http://dx.doi.org/10. 1109/9.182479.
- [11] K. Akkaya, M. Younis, A survey on routing protocols for wireless sensor networks, Ad Hoc Netw. 3 (3) (2005) 325–349.
- [12] Z. Mammeri, Reinforcement learning based routing in networks: Review and classification of approaches, Ieee Access 7 (2019) 55916–55950.
- [13] D.B. Johnson, D.A. Maltz, Dynamic source routing in ad hoc wireless networks, in: Mobile Computing, Springer, 1996, pp. 153–181.
- [14] C. Perkins, E. Royer, Ad-hoc on-demand distance vector routing, in: Proceedings WMCSA'99. Second IEEE Workshop on Mobile Computing Systems and Applications, 1999, pp. 90–100, http://dx.doi.org/10.1109/MCSA.1999.749281.
- [15] S. Jain, K. Fall, R. Patra, Routing in a delay tolerant network, in: Proceedings of the 2004 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications, 2004, pp. 145–158.
- [16] K. Sha, J. Gehlot, R. Greve, Multipath routing techniques in wireless sensor networks: A survey, Wirel. Personal Commun. 70 (2) (2013) 807–829.
- [17] R.S. Sutton, A.G. Barto, Reinforcement Learning: An Introduction, MIT Press, 2018.
- [18] C.J. Watkins, P. Dayan, Q-learning, Mach. Learn. 8 (3-4) (1992) 279-292.
- [19] S. Choi, D.-Y. Yeung, Predictive Q-routing: A memory-based reinforcement learning approach to adaptive traffic control, Adv. Neural Inf. Process. Syst. 8 (1995).
- [20] S. Kumar, R. Miikkulainen, Confidence-based q-routing: An on-line adaptive network routing algorithm, in: Proceedings of Artificial Neural Networks in Engineering, Citeseer, 1998.
- [21] Y. LeCun, Y. Bengio, G. Hinton, Deep learning, Nature 521 (7553) (2015) 436-444.
- [22] V. Mnih, K. Kavukcuoglu, D. Silver, A.A. Rusu, J. Veness, M.G. Bellemare, A. Graves, M. Riedmiller, A.K. Fidjeland, G. Ostrovski, et al., Human-level control through deep reinforcement learning, Nature 518 (7540) (2015) 529–533.
- [23] G. Stampa, M. Arias, D. Sánchez-Charles, V. Muntés-Mulero, A. Cabellos, A deep-reinforcement learning approach for software-defined networking routing optimization, 2017, arXiv preprint arXiv:1709.07080.
- [24] A. Valadarsky, M. Schapira, D. Shahaf, A. Tamar, Learning to route, in: Proceedings of the 16th ACM Workshop on Hot Topics in Networks, 2017, pp. 185–191.
- [25] J. Suarez-Varela, A. Mestres, J. Yu, L. Kuang, H. Feng, P. Barlet-Ros, A. Cabellos-Aparicio, Feature engineering for deep reinforcement learning based routing, in: ICC 2019 - 2019 IEEE International Conference on Communications, ICC, 2019, pp. 1–6, http://dx.doi.org/10.1109/ICC.2019.8761276.
- [26] X. You, X. Li, Y. Xu, H. Feng, J. Zhao, H. Yan, Toward packet routing with fully distributed multiagent deep reinforcement learning, IEEE Trans. Syst. Man Cybern.: Syst. (2020).
- [27] R.E. Ali, B. Erman, E. Baştuğ, B. Cilli, Hierarchical deep double Q-routing, in: ICC 2020 - 2020 IEEE International Conference on Communications, ICC, 2020, pp. 1–7, http://dx.doi.org/10.1109/ICC40277.2020.9149287.
- [28] V. Manfredi, A.P. Wolfe, B. Wang, X. Zhang, Relational deep reinforcement learning for routing in wireless networks, in: 2021 IEEE 22nd International Symposium on a World of Wireless, Mobile and Multimedia Networks, WOWMOM, 2021, pp. 159–168, http://dx.doi.org/10.1109/WoWMoM51794.2021.00029.
- [29] T.D. Kulkarni, K. Narasimhan, A. Saeedi, J. Tenenbaum, Hierarchical deep reinforcement learning: Integrating temporal abstraction and intrinsic motivation, in: Advances in Neural Information Processing Systems, Vol. 29, 2016, pp. 3675–3683.

- Ad Hoc Networks 153 (2024) 103347
- [30] M. Polese, M. Giordani, A. Roy, D. Castor, M. Zorzi, Distributed path selection strategies for integrated access and backhaul at mmwaves, in: 2018 IEEE Global Communications Conference, GLOBECOM, 2018, pp. 1–7, http://dx.doi.org/10. 1109/GLOCOM.2018.8647977.
- [31] M. Simsek, O. Orhan, M. Nassar, O. Elibol, H. Nikopour, IAB topology design: A graph embedding and deep reinforcement learning approach, IEEE Commun. Lett. 25 (2) (2020) 489–493.
- [32] M.L. Puterman, Markov Decision Processes: Discrete Stochastic Dynamic Programming, John Wiley & Sons, 2014.
- [33] K. Zhang, Z. Yang, T. Basar, Multi-agent reinforcement learning: A selective overview of theories and algorithms, 2019, CoRR abs/1911.10635, arXiv:1911. 10635.
- [34] C. Boutilier, Planning, learning and coordination in multiagent decision processes, in: TARK, Vol. 96, Citeseer, 1996, pp. 195–210.
- [35] K. Zhang, Z. Yang, H. Liu, T. Zhang, T. Basar, Fully decentralized multi-agent reinforcement learning with networked agents, in: International Conference on Machine Learning, PMLR, 2018, pp. 5872–5881.
- [36] P. Ngatchou, A. Zarei, A. El-Sharkawi, Pareto multi objective optimization, in: Proceedings of the 13th International Conference On, Intelligent Systems Application to Power Systems, 2005, pp. 84–91, http://dx.doi.org/10.1109/ISAP. 2005.1599245.
- [37] B. McMahan, E. Moore, D. Ramage, S. Hampson, B.A. y Arcas, Communicationefficient learning of deep networks from decentralized data, in: Artificial Intelligence and Statistics, PMLR, 2017, pp. 1273–1282.
- [38] T. Mai, H. Yao, Z. Xiong, S. Guo, D.T. Niyato, Multi-agent actor-critic reinforcement learning based in-network load balance, in: GLOBECOM 2020 - 2020 IEEE Global Communications Conference, 2020, pp. 1–6, http://dx.doi.org/10.1109/ GLOBECOM42002.2020.9322277.
- [39] J. Hu, M.P. Wellman, Nash Q-learning for general-sum stochastic games, J. Mach. Learn. Res. 4 (Nov) (2003) 1039–1069.
- [40] J. Foerster, G. Farquhar, T. Afouras, N. Nardelli, S. Whiteson, Counterfactual multi-agent policy gradients, in: Proceedings of the AAAI Conference on Artificial Intelligence, Vol. 32, no. 1, 2018.
- [41] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, W. Zaremba, Openai gym, 2016, arXiv preprint arXiv:1606.01540.

Shahaf Yamin is a distinguished electrical engineer with a Bachelor's degree from Ben Gurion University (2020) and is currently pursuing a Master's degree in electrical engineering. His research primarily centers around machine learning, specifically multi-agent reinforcement learning, applied to wireless networks. Shahaf's strong analytical and problem-solving skills are dedicated to advancing the field and developing innovative approaches to optimize network performance. He actively participates in conferences, workshops, and extracurricular activities, including mentoring students and collaborating with other researchers.

Haim H. Permuter (M'08-SM'13) received his B.Sc. (summa cum laude) and M.Sc. (summa cum laude) degrees in Electrical and Computer Engineering from the Ben-Gurion University, Israel, in 1997 and 2003, respectively, and the Ph.D. degree in Electrical Engineering from Stanford University, California in 2008. Between 1997 and 2004, he was an officer at a research and development unit of the Israeli Defense Forces. Since 2009 he is with the department of Electrical and Computer Engineering at Ben-Gurion University where he is currently a professor, Luck-Hille Chair in Electrical Engineering. Haim also serves as head of the communication, cyber/ and information track in his department. Prof. Permuter is a recipient of several awards, among them the Fullbright Fellowship, the Stanford Graduate Fellowship (SGF), Allon Fellowship, and the U.S.Israel Binational Science Foundation Bergmann Memorial Award. Haim served on the editorial boards of the IEEE Transactions on Information Theory in 2013-2016 and has been reappointed again in 2023.