Deep reinforcement learning for beam management in UAV relay mmWave networks

Dohyun Kim, *Member, IEEE*, Miguel R. Castellanos, *Member, IEEE*, and Robert W. Heath Jr., *Fellow, IEEE*

***Abstract*—Unmanned aerial vehicles (UAVs) offer a means to relay signals around obstacles in millimeter wave (mmWave) mobile ad hoc networks. Achieving these benefits, however, requires a dynamic beam management strategy that efficiently allocates resources for discovering, configuring, and exploiting communication links. Balancing these tasks is difficult due to the interplay between the overhead of beam acquisition and tracking and the resulting data rate over the link. In this magazine paper, we showcase how deep reinforcement learning (DRL) can jointly address the problems of blockage and mobility in mmWave ad hoc networks. We first summarize the problem of relay selection with realistic overhead penalties in which the beam management training time is characterized and minimized through a sequential decision-making approach. We then describe how hierarchical learning can be leveraged for choosing between distinct frequency bands for communication by addressing the issues posed by differing precoding training procedures. We conclude by overviewing how learning algorithms will be an important tool to overcome the challenges faced by future ad hoc networks.**

***Keywords*—mmWave MIMO, unmanned aerial vehicle, beam management, deep reinforcement learning**

I. Introduction

Modern wireless networks require gigabits-per-second data rates from mmWave communication to enable seamless sharing of raw or processed sensor data from various sources such as cameras, radars, and lidars [1]. Large throughputs are needed to aggregate shared sensor data and map a comprehensive “bird’s-eye view” perception of the surroundings that enhances the decision making capabilities of devices in the network. In commercial vehicular communication, such a map can be used to plan lane changing and acceleration/deceleration [2]. In tactical networks, precise real-time maps enhance defense efficiency against enemy ground units [3]. While the large bandwidths of the mmWave band can be leveraged to achieve these data rates, mmWave communication is sensitive to outages from blockages and beam misalignment from mobility. Both mobility and blockages are prevalent in highly dynamic scenarios, which means ad hoc networks should be designed to be especially robust to satisfy high data rate requirements.

UAVs have emerged as a key solution to building resilient mmWave networks [4]. UAVs exploit high operating altitudes to establish LOS ground-to-air channels and extend coverage. These aerial platforms prove especially valuable in ad hoc networks where infrastructure is sparse, as depicted in Fig. 1. The deployment of UAVs, however, introduces relay links that must be configured separately from the main link between the transmitter and receiver. As illustrated in Fig. 2, establishing a relay link typically includes the initial discovery of candidate UAV relays, followed by the selection of the best relay, and finally, beam configuration with respect to the selected relay. Link configuration can even include multi-band operation to increase channel diversity. For example, the sub-6 GHz band be used as an alternative to the mmWave band to leverage multi-path rich channels in cases of high blockage occurrences. Effective relay link configuration increases both coverage and spectral efficiency by choosing the best relay depending on the channel conditions [5].

Relay link configuration and maintenance incurs a notable amount of overhead from the pilot symbols needed for channel estimation and precoder training [6]. For example, in Fig. 2, the selected candidate relay may not be worth switching to because of the overhead associated with initial access. The main technical challenge in this context lies in finding the balance between insufficient beam alignment, leading to inaccurate relay link estimates due to fast-varying channels, and excessive beam alignment, resulting in a significant overhead from the pilot symbols. UAV relay networks face pronounced challenges, including high UAV mobility, complex 3D beam alignment, and channel volatility due to atmospheric conditions at mmWave frequencies [4]. State-of-the-art advancements on UAV networks have partially addressed such challenges, including trajectory

Dohyun Kim is with the Institute of Computer Technology, Seoul National University (SNU), Seoul 08826, South Korea (e-mail: dohyun.p.kim@snu.ac.kr). Miguel R. Castellanos and Robert W. Heath Jr. are with the Department of Electrical and Computer Engineering, North Carolina State University, 890 Oval Dr., Raleigh, NC 27606 USA (email: {mrcastel, rwheathjr}@ncsu.edu). This material is based upon work supported by the National Science Foundation under grant nos. NSF-ECCS-2153698, NSF-CCF-2225555, NSF-CNS- 2147955 and is supported in part by funds from federal agency and industry partners as specified in the Resilient & Intelligent NextG Systems (RINGS) program.

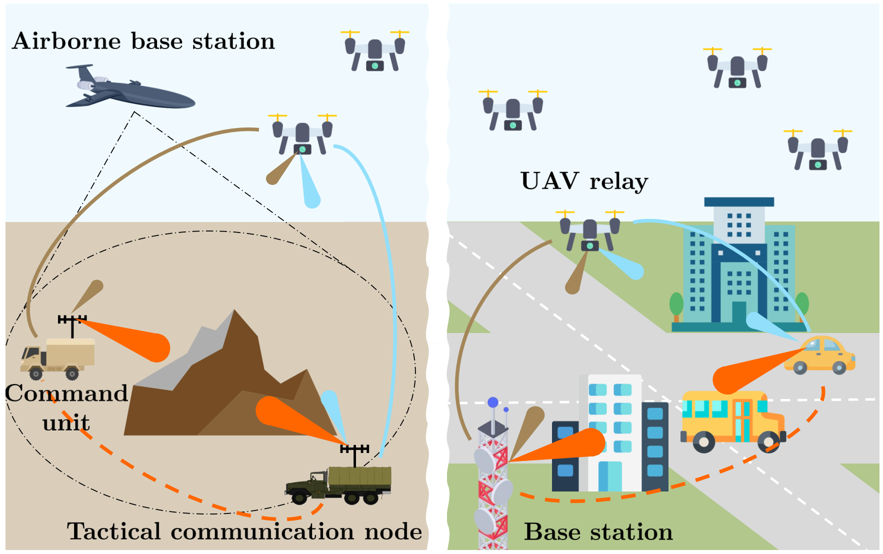

Fig. 1: Snapshot examples of UAV relay mmWave links in solid line, blocked mmWave links in dashed line, and sub-6 GHz coverage in dashdotted line. Rotorcraft UAVs can be distributed as relay nodes in tactical and urban networks, while fixed-wing UAVs may serve as aerial base stations supporting non-structured ad hoc infrastructure.

or placement optimization, yet less emphasis has been made minimizing beam management overhead [7], [8]. Although the position information from UAV trajectories can aid in tracking, it does not decrease the training time for switching links. Furthermore, minimizing the beam management overhead is complicated since it is intertwined with relay link evaluation.

DRL is a flexible framework that can efficiently solve the relay link selection and configuration problem by balancing the exploration-exploitation tradeoff. Unlike conventional signal processing methods, which may struggle with acquiring accurate channel models, DRL-based approaches excel with minimal requirements for online training data derived from system observations. The online data requirement is especially beneficial since the rapid dynamics of UAV relay networks can exacerbate overfitting concerns of offline learning approaches. DRL leverages computationally powerful neural networks to solve complicated sequential decision problems by continually adapting based on prior choices. The DRL framework has proven to be effective in similar dynamic settings such as network access, connectivity, and localization problems in multi-UAV wireless networks [8]. In the context of UAV relay networks, relay link selection and configuration involves a large number of decisions that are difficult to compare without accruing overhead. A learning-approach can acquire experience from past decisions to create better communication links in the future without needing to evaluate all choices. Building upon the fundamentals of DRL, incorporating variations such as hierarchical or transfer learning is the key to designing algorithms that can minimize link configuration time in UAV relay mmWave networks.

This article presents an overview on using DRL algorithms for beam management in UAV relay mmWave networks. For the scope of this article, we focus on data rate maximization as the objective assuming a link maintenance procedure depicted in Fig. 2. The procedure assumes a simplistic criterion for UAV trajectory/placement, relying on nearest neighbours. We first summarize the joint relay selection and beam management problem and the proposed DRL-based solution that picks the best relay and triggers beam realignment based on adaptive threshold learning. We then explore band assignment, as a complementary direction to relay selection addressing blockage in UAV relay mmWave networks, with a hierarchical learning approach. Lastly, we outline future research directions by pointing out prospective communication aspects of future UAV relay mmWave networks and corresponding learning algorithm developments.

II. Jointly addressing blockage and mobility in UAV relay networks

In this section, we focus on relay selection and band assignment for addressing blockage as a subproblem of candidate selection step of UAV relay link maintenance. We view the relay selection and band assignment problems as sequential decision making formulation. A Markov decision process (MDP) can then be leveraged to represent the learning model [9]. We explain the design of state, action, and reward that constitute an MDP and the corresponding learning algorithm.

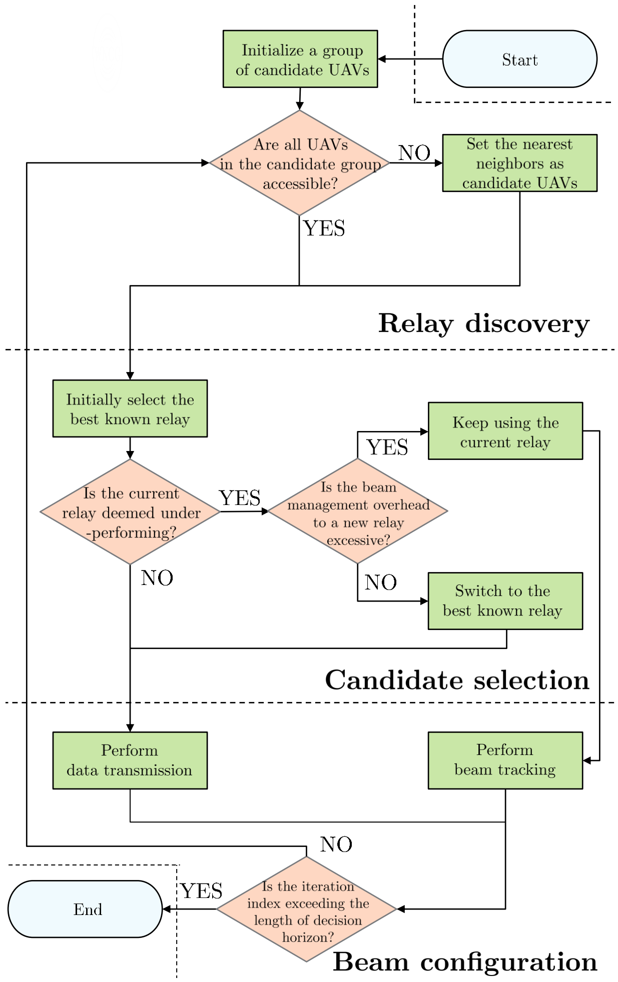

Fig. 2: The flow chart of UAV relay link maintenance procedure that consist of three steps: (1) discovering the available relays to construct a set of candidates, (2) selecting the candidate relay to use, (3) configuring beams with respect to the selected relay. The evaluation criterion in each step can be influenced by the associated beam management overhead.

A. *Joint relay selection and beam management*

Selecting the best relay is a nontrivial task due to the resources required to probe and establish a link. Even the act of checking whether a relay could result in a higher rate than the current link incurs a performance penalty. In general, there is a tradeoff between the rate increase from switching to a better link versus the rate decrease from training overhead. This introduces two opposing objectives for system design: reducing overhead and increasing link quality. An optimal relay selection method should include the penalty from training overhead to accurately depict how resources spent on link establishment deteriorate throughput.

Within the network, the objective of a transmitter communicating with a single receiver is to maximize the its cumulative data rate. The transmitter can either establish a direct link to the receiver or an indirect two-hop link via a relay. The transmitter can function in either of two transmission modes: data transmission, where it sends data symbols to the receiver, and beam alignment, where it either updates the current link or establishes a new link. At each time instant, the transmitter faces a sequence of decisions, weighing options of persisting in data transmission, enhancing the current link through additional training, or entirely switching to a new relay. Assuming no data is transmitted during beam alignment, the objective to maximize data rate will force the transmitter to find a balance between beam management and data transmission.

DRL can be applied to solve the joint relay selection and beam management problem by learning a policy based on prior observations. Three concepts are necessary to specify an MDP and understand DRL algorithms: state, action, and reward. The state represents the quality of all available communication links, including beam pairs and achievable spectral efficiency. The action comprises the index of the selected relay and indicators for beam training or data transmission. The reward is determined by the cumulative data rate. DRL algorithms approach the sequential decision making formulation through trial-and-error. The aim of the algorithm is to learn a policy, which maps a state to an action, by executing actions per state and observing the ensuing reward. The reward is used to evaluate how good a policy is, which the algorithm can then use as feedback to guide the evolution of the policy. One important benefit of DRL algorithms is their ability to adapt to dynamic conditions. The optimal decision policy may change depending on external factors, like the blockage frequency and the number of candidate relays. A DRL algorithm is able to continually update the policy based on recent observations of the environment.

B. *Joint band assignment and beam management*

Modern wireless networks can leverage multiple operation over multiple frequency bands to ensure robust and reliable communication. While the aid of relays can help mitigate the effect of blockages, relays cannot completely eliminate blockages, for example, in severely congested scenarios. In these cases, however, systems can leverage the availability of multipath channels in the sub-6 GHz band to ensure transmission continues. These lower frequency channels cannot achieve the high data rates of mmWave communication due to their lower bandwidth, but their channel characteristics make them more resilient. The research challenge lies in effectively combining mmWave and sub-6 GHz systems to achieve high average throughput while maintaining a stable link.

Similar to the joint relay selection and beam management problem, the transmitter needs to select the operational band and a beam management mode. A distinctive research challenge emerges from the different signal processing architectures employed in the mmWave and sub-6 GHz band. MmWave systems typically use hybrid precoding architectures in which the digital and analog signal processing are separated. This functional split means that the digital and analog beams must be configured separately, which generally incurs a high amount of overhead relative to sub-6 GHz. Fortunately, both digital and analog precoding can be trained separately. At lower frequencies, fully-digital precoding is viable due to the manageable number of antennas. Beam training at lower frequencies typically involves estimating the channel using a quantized codebook and feeding back the information from the receiver to the transmitter. This procedure is less intensive than mmWave beam management, which means sub-6 GHz are generally easier to configure and manage.

The problem of distinctive beam management procedures in the mmWave and sub-6 GHz band can be approached by introducing a decision hierarchy in the DRL algorithm. In hierarchical reinforcement learning (HRL), as described throughout Algorithm 1, the aim of the agent is to learn two policies: the upper-level policy and lower-level policy. The upper-level policy makes high-level decisions that decompose a task into subtasks, and the lower-level policy chooses a subtask based on a goal [10]. Since the beam management procedure depends on the band of operation, it makes sense to assign the band assignment problem to the upper-level policy. Once a band is chosen, the lower-level policy can choose among the beam management subtasks, which can include initial access, tracking, or partial training. By separating the decision into two layers, the algorithm is able to use a divide-and-conquer approach to learn effective policies for each problem.

Algorithm 1 Joint band assignment and beam management strategy based on HRL

```
Input: Length of decision horizon, Boolean random
  variable RoundSkip
Randomly initialize online critic and actor network
  for upper-level and lower-level policies
for each decision iteration do
  if RoundSkip then
    Continue using upper-level action
  else
    Aggregate states of the latest consecutive
      skipped rounds
    Set goal as according to importance relabling
    Set reward as cumulative reward over the latest
      consecutive skipped rounds
    Store upper-level transition in upper-policy's
      experience replay
    Update actor and critic networks of the upper-
      level policy
    Update band of operation
  end if
  Deploy the lower-level policy's action
  Obtain the intrinsic reward from the upper-level
    policy
  Update beam management mode
  Store lower-level transition in lower-policy's ex-
    perience replay
  Update actor and critic networks of the lower-level
    policy
end for
```

C. Selected numerical results

In Fig. 3, we show a performance comparison between a DRL-based relay selection algorithm and a few baselines. The genie-aided policy has perfect knowledge of the channel information, the optimal threshold policy applies offline learning based on a fixed offline training dataset, and the direct policy uses the direct link and follows the genie-aided policy's beam management procedure. Further details on the baselines, exhaustive simulation parameters, and more results are available in [11] (and references therein). The DRL algorithm compares the current link rate to two adaptive thresholds that are learned over time. The thresholds dictate when the transmitter performs relay switching and/or beam alignment. Intuitively, these thresholds should change depending on the current quality of the network links. The results show that the DRL algorithm is able to outperform static policies that do not change over time. The DRL-based policy achieves a higher spectral efficiency that an offline-learning approach since the adaptive threshold can match the dynamic channel.

Fig. 4 depicts the episodic reward comparison between two HRL and a DRL approach with selected simulation parameters shown in Table I. Exhaustive simulation parameters, baselines, and further simulation results are available in [12]. The HRL approaches outperform the DRL counterpart both in how fast the episodic rate reward increases and how much the reward increases. As DRL algorithms are based on trial-and-error, being able to uniformly explore the environment is a key requirement for these methods to work well. Without hierarchy, the DRL algorithm is limited in exploration and is unable to fully explore the decision space. HRL explores more efficiently by letting the upper-level policy define exploration goals while the lower-level policy focuses on reaching those goals. The benefits of HRL can be further improved by leveraging off-policy corrections,

	MmWave band	Sub-6 GHz band
Array type	Fully-connected hybrid architecture	Fully digital architecture
Antenna number	32×16 system, 4 streams	8×8 system, 4 streams
Bandwidth (OFDM subcarriers)	850 MHz (256)	150 MHz (32)
Codebook type	Discrete Fourier transform	Type-I precoder matrix indicator

TABLE I: Selected simulation parameters of the joint band assignment and beam management problem.

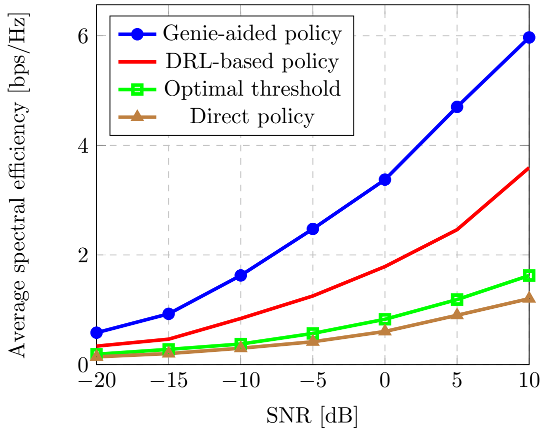

Fig. 3: Spectral efficiency comparison between genie-aided policy, proposed DRL-based policy, and baselines. The benchmarks include a direct policy, in which relays are not used, an optimal threshold policy, which decides to switch when the current link rate falls below a fixed threshold, and a genie-aided policy. The DRL algorithm adaptively adjusts its decision-making policy based on the current channel conditions.

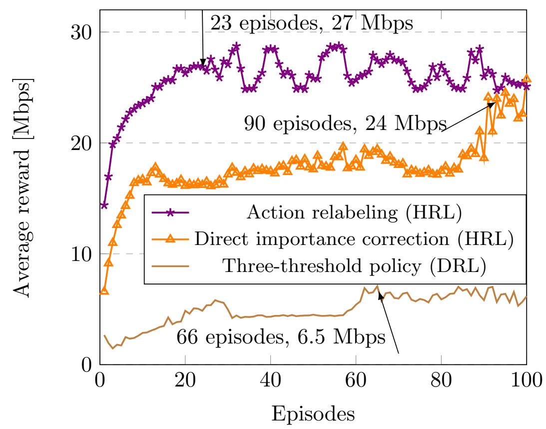

Fig. 4: Comparison of episodic reward of DRL algorithm to HRL algorithms. Compared to the DRL baseline, the HRL algorithms have a boost in episodic reward of more than 10 Mbps with efficient exploration. With off-policy correction based on action relabeling, HRL algorithms can reach peak episodic reward in less than half episodes than the DRL counterpart.

optimizing the update period of the upper-level policy, and employing action skipping in the upper-level policy training [12].

Jointly optimizing relay selection, band assignment, and beam management can be a straightforward extension based on HRL algorithms. In the following section, we provide examples of novel challenges and possible learning approaches.

III. Beam Management Challenges in Future UAV Relay mmWave Networks

In this section, we look into prospective UAV relay mmWave networks, as illustrated in Fig. 5, focusing on the new array architectures, network topology, and their associated issues in terms of learning algorithm designs.

A. Generalized codebook designs

Regardless of the algorithm used to jointly address blockage and mobility, the codebook employed by mmWave networks can be a bottleneck of the achievable rate. Current 3GPP mmWave beam management standards [2] involve sweeping analog beams over the angular domain. Large arrays with narrow beam patterns can lead to prohibitively high overhead as beam alignment becomes more difficult. Furthermore, the use of frequency-flat phase-shifters in wideband systems can lead to a phenomenon known as beam squit, where the beamforming direction skews as the frequency gets far away from the center frequency. Beam squint leads to a significant loss of array gain across the entire bandwidth and reduced data rate. Overall, an innovation in the array architecture and codebook design is needed to enhance the rate performance of ad hoc networks.

Frequency-selective true time delay (TTD) architectures can be exploited to both mitigate beam and significantly reduce training overhead. TTD beamforming has been used in radar systems for decades as a way to combat beam squinting. Recent studies, however, have also proposed TTD-based multi-frequency probing in which a beam pattern with arbitrary direction and band range can be generated [13]. Such beam patterns can be designed to minimize the overhead induced from ini-

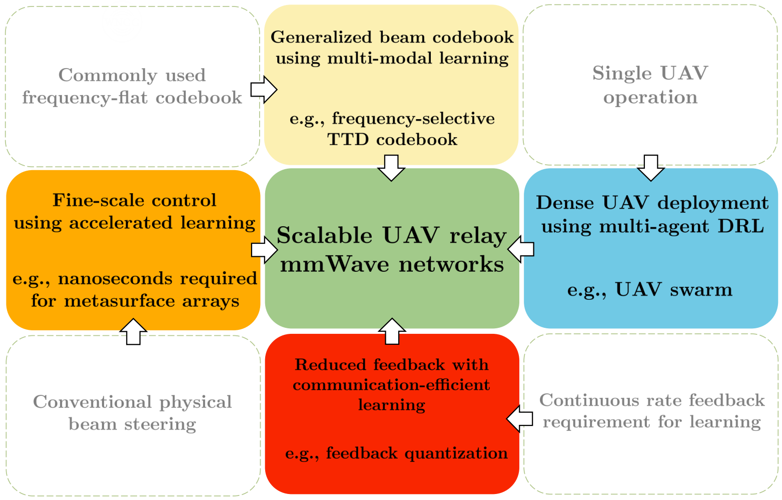

Fig. 5: Summary of four selected prospective aspects and learning solutions of future resilient and intelligent UAV relay mmWave networks. Conventional methods are denoted in dotted bubbles with arrows pointing to colored bubbles representing technical advancements needed for scalable UAV relay mmWave networks.

tial access and beam tracking while minimizing energy leak to undesired directions. Multi-modal data, such as images or videos, may be used with DRL algorithms to construct a TTD codebook, which consists of delay and phase, based on feedback that advocates data rate penalizes power efficiency loss.

B. Adaptive arrays

Advances in adaptive arrays have significantly reduced the operation time required for controlling and switching beams. One example is the lens array that electronically steers beams by switching on and off a subset of antennas. Another example is the metasurface reflective arrays that take input electrical control signals to produce a desired aperture field. The array examples share the common operating time ranging in few nanoseconds, which implies measurements will be streamed to the DRL algorithms in the same time scale. Typical DRL algorithms require extensive computation even for a single iteration of weight update and gradient computation that can be as long as several milliseconds. Aggregating the streamed data to compute a single iteration of DRL is possible, yet the resulting DRL algorithm may suffer from slow responsiveness to real-time changes in the environment. To apply DRL algorithms for the next-generation arrays, a breakthrough is needed to overcome the the computational bottleneck of the clock speed of processors stuck at several gigahertz.

The challenge of accelerating DRL is a complex task requiring technical contributions from both the hardware and software layers. On one hand, new computing hardware architectures, such as multicore processors and field-programmable gate arrays, needs dedicated design and implementation on UAVs to speed up DRL up to the nanosecond processing time scales [14]. On the other hand, reducing computational complexity of fundamental operations using approximate matrix multiplication and reducing neural network size with pruning/quantization seeks attention.

C. UAV swarm communication

UAV swarm communication introduces unique challenges to form a cohesive system while maintaing communication between the nodes [8]. The computational complexity will not only scale with the size of the network, but the candidate selection step will also necessitate more sophisticated methods, including routing and interference management. Notably, trajectory/placement control will become increasingly crucial to ensure synchronization among UAVs, and power constraints on each UAV will become more stringent, especially in multidirectional communication scenarios such as data dissemination. In case of a decentralized swarm scenario, multi-agent DRL can allow each UAV to distributively execute its policy based on data exchangement between neighbors. The number of neighbor nodes should balance the tradeoff between increase in sum reward from representative cooperation data versus the overhead from data exchangement between nodes. Conventional DRL algorithms may be used in UAV swarm scenarios with centralized control and communication, where simplified deep learning architectures like binary neural networks can compensate the higher computational demands.

D. Limited feedback

DRL algorithms deployed in UAV relay mmWave networks operate based on reward in the form of feedback from receiving nodes. This feedback should be encoded in bits and sent via a feedback channel, implying that a tradeoff exists between the feedback accuracy and overhead. DRL algorithms should incorporate several feedback parameters in the action to efficiently minimize feedback overhead while maintaining an acceptable accuracy. The feedback parameters may include but not be limited to: how often the feedback is sent, what type of feedback (e.g. channel information or spectral efficiency) will be used, and the number of quantization bits used in the feedback channel. DRL should adopt the methods studied in communication-efficient learning, such as distributed learning, federated learning, and split learning, to select the best beam under the feedback constraints [15].

IV. Conclusions

In this article, we have discussed the benefits of DRL in approaching the complicated problems that can arise in dynamic UAV relay mmWave networks. Sequential decision-making formulations and DRL algorithms are useful to identify the main performance bottleneck of beam management overhead and minimize the overhead. In cases where the beam management procedure can vary over decisions, incorporating hierarchical structure in learning enhances data rate performance. DRL will become a foundational technology in ensuring scalable UAV-networks with high data rate, potentially using TTD codebook construction on nanosecond time scale.

References

[1] J. Choi, V. Va, N. Gonzalez-Prelcic, R. Daniels, C. R. Bhat, and R. W. Heath, "Millimeter-wave vehicular communication to support massive automotive sensing," *IEEE Commun. Mag.*, vol. 54, no. 12, pp. 160–167, Dec. 2016.

[2] M. H. C. Garcia, A. Molina-Galan, M. Boban, J. Gozalvez, B. Coll-Perales, T. Şahin, and A. Kousaridas, "A Tutorial on 5G NR V2X Communications," *IEEE Commun. Surveys Tuts.*, 3rd Quart. 2021.

[3] J. F. Harvey, M. B. Steer, and T. S. Rappaport, "Exploiting high millimeter wave bands for military communications, applications, and design," *IEEE Access*, vol. 7, pp. 52 350–52 359, Apr. 2019.

[4] M. M. Azari, S. Solanki, S. Chatzinotas, and M. Bennis, "THz-empowered UAVs in 6G: Opportunities, challenges, and trade-offs," *IEEE Commun. Mag.*, vol. 60, no. 5, pp. 24–30, May 2022.

[5] N. Zhao, W. Lu, M. Sheng, Y. Chen, J. Tang, F. R. Yu, and K.-K. Wong, "UAV-assisted emergency networks in disasters," *IEEE Wireless Commun.,*, vol. 26, no. 1, pp. 45–51, Feb. 2019.

[6] M. Giordani, M. Polese, A. Roy, D. Castor, and M. Zorzi, "A tutorial on beam management for 3GPP NR at mmWave frequencies," *IEEE Commun. Surveys Tuts.*, vol. 21, no. 1, pp. 173–196, 1st Quart. 2019.

[7] P. A. Karegar and A. Al-Anbuky, "UAV-assisted data gathering from a sparse wireless sensor adaptive networks," *Wireless Netw.*, vol. 29, no. 3, pp. 1367–1384, Dec. 2022.

[8] Y. Bai, H. Zhao, X. Zhang, Z. Chang, R. Jäntti, and K. Yang, "Towards autonomous multi-UAV wireless network: A survey of reinforcement learning-based approaches," *IEEE Commun. Surveys Tuts.*, 4th Quart. 2023.

[9] M. L. Puterman, *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. New York, NY, USA: Wiley, 2014.

[10] M. Eppe, C. Gumbsch, M. Kerzel, P. D. Nguyen, M. V. Butz, and S. Wermter, "Intelligent problem-solving as integrated hierarchical reinforcement learning," *Nat. Mach. Intell.*, vol. 4, no. 1, pp. 11–20, Jan. 2022.

[11] J. Ye and X. Ge, "Beam management optimization for V2V communications based on deep reinforcement learning," *Sci. Rep.*, vol. 13, no. 1, p. 20440, Nov. 2023.

[12] D. Kim, M. R. Castellanos, and R. W. Heath, "Joint band assignment and beam management using hierarchical reinforcement learning for multi-band communication," *arXiv preprint arXiv:2308.13202*, 2023.

[13] I. K. Jain, R. R. Vennam, R. Subbaraman, and D. Bharadia, "mmFlexible: Flexible directional frequency multiplexing for multi-user mmWave networks," in *Proc. IEEE Int. Conf. Comput. Commun.*, May 2023, pp. 1–10.

[14] M. Rothmann and M. Porrmann, "A survey of domain-specific architectures for reinforcement learning," *IEEE Access*, vol. 10, pp. 13 753–13 767, Jan. 2022.

[15] J. Park, S. Samarakoon, A. Elgabli, J. Kim, M. Bennis, S.-L. Kim, and M. Debbah, "Communication-efficient and distributed learning over wireless networks: Principles and applications," *Proc. IEEE*, vol. 109, no. 5, pp. 796–819, May 2021.

Dohyun Kim *dohyun.p.kim@snu.ac.kr* is a postdoctoral researcher in the Institute of Computer Technology at Seoul National University. Dr. Kim's PhD work was focused on developing deep reinforcement learning algorithms for jointly addressing blockage and mobility in mmWave networks. His current research interests includes reconfigurable antennas and approximate computing to envision cross-layer optimized machine learning applications for next-generation wireless systems. Along with his coauthors, he was the recipient of the 2022 Fred W. Ellersick MILCOM Best Paper Award and the 2022 Vanu Bose Best Paper Award.

Miguel R. Castellanos *mrcastel@ncsu.edu* is a postdoctoral researcher in the Department of Electrical and Computer Engineering at North Carolina State University. Part of his research has focused on incorporating concepts from electromagnetics into wireless communication to address problems such as user radiation exposure, depolarization, near-field effects, and mutual coupling. His recent work analyzes reconfigurable antennas and revisits electromagnetic-based models for the array manifold.

Robert W. Heath Jr. *rwheathjr@ncsu.edu* is the Lampe Distinguished Professor in the Department of Electrical and Computer Engineering at the North Carolina State University and also codirector of 6GNC. He has authored or coauthored several books, including *Introduction to Wireless Digital Communication* (Prentice Hall, 2017) and Foundations of MIMO Communication (Cambridge University Press, 2018). He is the recipient or corecipient of several awards, including the 2019 IEEE Kiyo Tomiyasu Award, 2020 IEEE Signal Processing Society Donald G. Fink Overview Paper Award, 2020 North Carolina State University Innovator of the Year Award, 2021 IEEE Vehicular Technology Society James Evans Avant Garde Award, and 2022 IEEE Vehicular Technology Society Best Vehicular Electronics Paper Award.