







# Rate Adaptation Aware Positioning for Flying Gateways Using Reinforcement Learning

Gabriella Pantaleão<sup>(✉)</sup>, Rúben Queirós, Hélder Fontes,  
and Rui Campos

INESC TEC and Faculdade de Engenharia, Universidade do Porto, Porto, Portugal  
{gabriella.pantaleao,ruben.m.queiros,helder.m.fontes,  
rui.l.campos}@inesctec.pt

**Abstract.** With the growing connectivity demands, Unmanned Aerial Vehicles (UAVs) have emerged as a prominent component in the deployment of Next Generation On-demand Wireless Networks. However, current UAV positioning solutions typically neglect the impact of Rate Adaptation (RA) algorithms or simplify its effect by considering ideal and non-implementable RA algorithms. This work proposes the Rate Adaptation aware RL-based Flying Gateway Positioning (RARL) algorithm, a positioning method for Flying Gateways that applies Deep Q-Learning, accounting for the dynamic data rate imposed by the underlying RA algorithm. The RARL algorithm aims to maximize the throughput of the flying wireless links serving one or more Flying Access Points, which in turn serve ground terminals. The performance evaluation of the RARL algorithm demonstrates that it is capable of taking into account the effect of the underlying RA algorithm and achieve the maximum throughput in all analysed static and mobile scenarios.


**Keywords:** Aerial networks · Rate Adaptation · UAV positioning · Deep Reinforcement Learning

## 1 Introduction

Unmanned Aerial Vehicles (UAVs), commonly known as drones, have emerged as a promising technology for ensuring cost-effective on-demand wireless connectivity [6, 18]. The ability of UAVs to fly and navigate autonomously allow them to provide access to communications services where no infrastructure coverage exists, which can be particularly useful in remote or difficult-to-reach areas. UAVs are especially relevant in the context of sudden fluctuations in traffic demands that impair the effective allocation of radio resources, a scenario frequently seen in crowded events, for instance. Thus, UAVs are suitable platforms for delivering and enhancing connectivity in heterogeneous scenarios, carrying communications hardware to deploy Wi-Fi or cellular coverage. UAVs are highly

adaptable and can be quickly deployed and optimally positioned, in contrast to conventional ground-based solutions [8].

In flying networks, relays are often employed to improve network coverage and capacity, as represented in Fig. 1. The distance between UAVs directly influences the communications range between the flying nodes, as well as the Quality of Service (QoS) associated to the inter-UAV wireless links. Therefore, the use of intermediate UAVs, herein named Flying Gateways (FGWs), enables the extension of the communication range between the Backhaul network and the Flying Access Points (FAPs), while preserving the quality of the established connection [7]. By optimizing the positioning of the FGWs, there is an effective balance of the load in the network while improving coverage, resulting in a more efficient and reliable flying network. Nevertheless, the optimization of the positions of UAVs in flying networks according to specific link metrics is still a challenge, considering their dynamic nature.



**Fig. 1.** Relay network topology example, adapted from [12].

Machine Learning approaches can be employed to enhance various aspects of the performance and operation of flying networks, including delay, throughput, transmission power, and cache resource utilization [1, 3, 5, 11, 15]. Reinforcement Learning (RL) techniques are particularly useful, as they are intrinsically related to control theory, given its feedback-driven optimization process. RL represents a relevant approach to handle the continuous changing environment of flying networks, taking into account long-term, sequential and cumulative rewards, as it maps states into the effectiveness of the actions [10]. With RL, decisions are based on the current states of the environment, meaning that the agent can be trained using real-time network measurements instead of relying on approximations [1]. This allows the agent to have a more accurate response to the current environment conditions.

The main contribution of this paper is the Rate Adaptation aware RL-based Flying Gateway Positioning (RARL) algorithm, which addresses the problem of finding the position for the FGW that maximizes the throughput obtained in the FAPs, considering the influence of the underlying RA algorithm. The selected RA algorithm was the Minstrel-High Throughput (HT), designed as the evolution of the Minstrel algorithm for the IEEE 802.11n standard [13, 16, 17].

The effectiveness of the RL approach was validated through simulations using Network Simulator 3 (ns-3), which served as the platform for training, validating and testing the model.

The rest of the paper is organized as follows. Section 2 presents the related work on UAV positioning in flying networks. The design of the RARL algorithm is detailed in Sect. 3. Section 4 introduces the ns-3 simulation setup, as well as the simulation results on the performance of the RARL algorithm. Finally, Sect. 5 presents the conclusions and points out the future work.

## 2 Related Work

With the advancements made in the context of flying networks, many studies have been conducted to improve the positioning of the flying nodes. The work reported in [7] includes the implementation of centralized on-demand Gateway UAV positioning algorithm, relying on the awareness of the incoming traffic behaviour. Based on a mathematical approach, the goal is to minimize the capacity of the bidirectional wireless link between the FGW and the FAPs, while ensuring the bitrates required by each FAP. The results indicate that the position of the FGW is an essential aspect of the Backhaul network configuration, since constant FGW position updates can improve network performance. However, the authors consider an ideal underlying RA algorithm.

In [5], a solution for flying Base Stations, serving a region with numerous users, was proposed. This deployment relies on RL techniques for the UAV to be able to learn how to optimize its trajectory, contouring obstacles and reaching the intended service area. In this study, the UAV is assumed to return to its landing location when reaching the limit period of the flight. With a Q-Learning approach, the system was implemented based on making direct movement decisions, allowing the optimization of the sum rate of the transmission over the duration of the flight. With this solution, the algorithm needs no prior knowledge of the environment and has the ability to learn the structure of the network, resulting in improved performance. However, a limiting factor of this study is that it considers a static environment, with no changes in the course of the flight.

Overall, the main UAV positioning approaches include brute-force searching, mathematical optimizations [2, 7, 19], heuristics [9], and Machine Learning algorithms [1, 3, 5, 11, 15]. Despite their importance, optimizing UAV positions in a network to maximize one or more link metrics remains a challenging task, given the multiple factors to consider. From the literature review, Deep Reinforcement Learning (DRL) emerges as a promising approach for UAV positioning, supporting the choice of the Q-Learning approach in the implementation of the RARL algorithm.

It is worth noting that few solutions have considered the impact of the Backhaul network on the QoS experienced by ground terminals [7]. In addition, they are focussed on static and invariant scenarios, assuming mostly symmetric links. Moreover, none of the studies analysed included a realistic RA algorithm, which limits the applicability of these methods. Our work aims to address the

forementioned shortfalls, by considering a continuous analysis of the network state and accounting for the implementation of realistic and dynamic adjustments of the data rate in the network.

### 3 RARL Algorithm Design

The RARL algorithm relies on the implementation of Deep Q-Learning, a Deep RL method that has been proven to be highly applicable in continuous control tasks, making it well-suited for addressing the positioning of UAVs within flying networks. Three different scenarios were defined to test the RARL algorithm, as presented below:

1. **Asymmetric links scenario:** with three static nodes corresponding to the Backhaul, the FGW and the FAP, considering that the link between the FGW and the Backhaul node and the link between the FAP and the FGW are asymmetric, meaning that they have different values for the transmission power. This scenario is illustrated in Fig. 2a.
2. **Moving FAP scenario:** with the same network topology as the **asymmetric links** scenario, where the link between the FAP and the FGW and the link between the FGW and the Backhaul node are symmetric, but the FAP moves, as presented in Fig. 2b.
3. **Two FAPs scenario:** with four static nodes, corresponding to the Backhaul, the FGW and two FAPs, where the link between the FGW and the Backhaul node handles the traffic of the two symmetric links between the FGW and the FAPs, as shown in Fig. 2c.

Without loss of generality, in the two last scenarios we have considered symmetric wireless links in order to isolate the problems. The learning process was adapted for each of the three scenarios. For all the cases, the agent was instantiated in the FGW node, which is responsible for performing the different actions selected. Nevertheless, the algorithm modelling was tailored for each scenario. The action space, the observation space and the reward functions defined for each scenario are described below.

#### 3.1 Action Space

The actions of the FGW are based on discrete sequential movements, based on incrementing the position in the 2D venue in the selected direction by 25 m, a distance defined as a compromise between the training time and the impact of the FGW's movement on the link metrics, selected every 1 s, defined as the decision interval. Thus, the action space comprises five movement options, including the possibility of remaining in the same position: *Left*, *Right*, *Up*, *Down*, and *Same*.

#### 3.2 Observation Space

The observations must characterize the system so that the agent can recognize the best positions for maximizing and balancing the throughput values in the FGW and FAP nodes. For the **asymmetric links** and **moving FAP** scenarios, the observations are defined as:

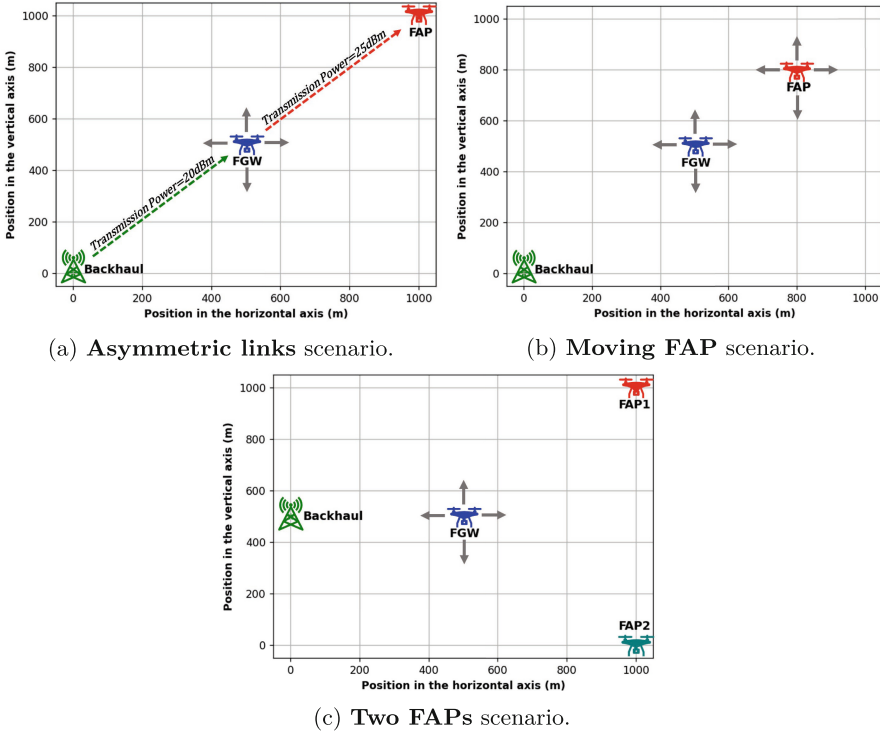


Fig. 2. Three scenarios analysed.

- the coordinates of the FGW;
- the distances between the FGW and the Backhaul node and between the FGW and the FAP;
- the throughput values measured in each link, calculated as the number of bytes received throughout the course of a decision interval, i.e. every second.

For the **two FAPs** scenario, the approach to sense the environment is different, given the existence of more than one FAP. In this case, the observations do not include the throughput values.


### 3.3 Reward Function

The underlying Minstrel-HT RA algorithm has a significant impact on the obtained throughput values, heavily influenced by the distances travelled by the FGW. This means that even if it travels to the same final position, the initial conditions imposed by the link metrics in the starting point impact the final throughput value obtained at the destination point. This means that the closer the UAV originally is to the final position, the fewer fluctuations are seen in the observed throughput. This is due to the fact that, when using Minstrel-HT,

the throughput variation has a slow response in cases where the link quality improves, as shown in [4, 13, 14, 16]. Hence, modelling the system becomes challenging due to the need to identify positions that optimize and balance the throughput values, which in turn requires addressing the peaks and variations introduced by the Minstrel-HT algorithm.

To better understand the variation of the network metrics under the effect of Minstrel-HT in the ns-3 simulations for the reward function design, a preliminary study was carried out to understand the Signal-to-Noise Ratio (SNR) and the throughput. For this purpose, the Backhaul node was positioned in (0, 0) and the FAP in (1000, 1000), with the configuration showed in Fig. 2a. The FGW was set to move from the coordinates (25, 25) to (975, 975). The position of the FGW was increased incrementally, both horizontally and vertically, by 25 m every 1 s. This approach was employed to precisely examine the impact of distance on link metrics.

For the **asymmetric links** and **moving FAP** scenarios, since the SNR, shown in Fig. 3b, has greater sensitivity to changes in position when compared to the throughput, presented in Fig. 3a, the SNR was considered as a more suitable measure of the environment's state. This means that, the objective can be modelled according to the SNR, with the reward functions translating into maximizing the SNR values and minimizing the difference between the SNR obtained in both the link connecting the Backhaul node and the FGW and in the link between the FGW and the FAP.



(a) Throughput variation in nodes FAP and FGW. (b) SNR variation in nodes FAP and FGW.



**Fig. 3.** Link metrics comparison for the **asymmetric links** and **moving FAP** scenarios.

The goal is to maximize the throughput values and to minimize possible throughput imbalances between the two nodes acting as receivers in each wireless link. For this purpose, the objective is then to maximize the SNR value in the FGW,  $SNR_{FGW}$ , and the SNR value in FAP,  $SNR_{FAP}$ , and penalize the difference between them, by multiplying it by a weight of 2. This constant was chosen empirically, as it allows the penalty to be sufficient to impact the reward

value when the links are imbalanced, without impacting negatively the learning process. The reward function is defined in Eq. 1:

$$Reward = SNR_{FGW} + SNR_{FAP} - 2|SNR_{FGW} - SNR_{FAP}| \quad (1)$$

In the case of the **two FAPs** scenario, the objective remains the same: to maximize the throughput values while minimizing throughput imbalances between the two nodes. However, in this scenario, variations in throughput and SNR values differ due to differences in the network topology and to the need to route traffic through the FGW node. As the distance between the FGW and the Backhaul node increases, there is an increase in the Packet Loss Ratio, which implies a reduction of the link capacity. This is demonstrated by the results obtained from the same preliminary study conducted for this scenario, where the Backhaul node was positioned in (0, 500), the FAP1 in (1000, 1000) and the FAP2 in (1000, 0), as represented in Fig. 2c. The results of the throughput variation in Fig. 4a and the SNR variation in Fig. 4b show that the throughput and SNR are not proportional, due to the interdependency of the links. For instance, if we move the FGW closer to the Backhaul node, the links between the FGW and the FAPs get stretched, thus leading to a reduction in the throughput obtained in the link between the FGW and the Backhaul node.



(a) Throughput variation in FAPs and FGW.

(b) SNR variation in FAPs and FGW.

**Fig. 4.** Link metrics comparison for the **two FAPs** scenario.

Hence, in this scenario, the reward was designed to rely directly on the throughput values obtained in the FAPs. As in the first reward function, the reward was designed to penalize the possible throughput imbalances obtained in the FAPs and favour higher throughput values. Similarly, a factor of 2 was considered. The reward function for the **two FAPs** scenario is defined in Eq. 2, where  $T$  refers to the throughput.

$$Reward = T_{FAP_1} + T_{FAP_2} - 2|T_{FAP_1} - T_{FAP_2}| \quad (2)$$

## 4 Simulation Configuration and Results

The simulations performed to validate the proposed RARL algorithm were carried out using ns-3. This section presents the configuration of the simulation environment and the simulation results obtained for each scenario.

### 4.1 Simulation Configuration

The ns-3 simulator played a crucial role in simulating real-world behaviour in various scenarios. It served as the foundation for the DRL agent to gather relevant information from the environment, specifically the link metrics of interest. It is worth noting that the different connections were designed as independent links, and downlink traffic was considered. Table 1 summarizes the most relevant parameters applied in the simulation.

In this study, the simulations utilized the IEEE 802.11n Wi-Fi standard, commonly referred to as Wi-Fi 4. Moreover, the antenna gain was defined to have 0 dBi, meaning that for simplicity the antenna is considered to be isotropic. When it comes to the propagation loss model, the Friis Free Space Propagation Model was applied as it offers an approximation of the line-of-sight path loss that occurs in free space environments. Finally, to avoid the overhead associated with the robustness of Transmission Control Protocol, the traffic protocol chosen was the User Datagram Protocol (UDP), with a fix packet size of 1400 bytes. Moreover, in all the defined scenarios, the UDP data rate was set to be 70 Mbit/s to ensure that there was always traffic being generated, meaning that the links were saturated, maintaining a consistent load.

**Table 1.** Simulation parameters considered for the three scenarios.

ns-3.37 Simulator Parameters	
Wi-Fi Standard	IEEE 802.11n
Channel Bandwidth	20 MHz
Antenna Gain	0 dBi
Propagation Loss Model	Friis
Rate Adaptation Algorithm	Minstrel-HT
Application Traffic	UDP
UDP Data Rate	70 Mbit/s
Packet Size	1400 bytes

## 4.2 Simulation Results

The initial positions of the FGW were defined as extreme locations, tailored to each specific scenario. The results are presented below according to the scenario.

**Asymmetric Links Scenario.** Regarding the **asymmetric links** scenario, the aim was to assess the impact of the asymmetric links in the optimal positioning of the FGW. Thus, the Backhaul and the FAP nodes were stationary, located at the coordinates (0, 0) and (1000, 1000), respectively. Furthermore, the initial position for the FGW was at the coordinates (500, 500), to show that the model was able to recognize the asymmetry between the links and find the optimal position in the new scenario, following the trajectory shown in Fig. 5a, as depicted below:


- Initial position of the FGW: (500, 500);
- Final position of the FGW: (175, 525);
- Final distance between FGW and FAP: 950 m;
- Final distance between FGW and Backhaul: 550 m.

Regarding the **asymmetric links** scenario, it is possible to see that the FGW recognizes the asymmetry of the links, as the final position is significantly closer to the Backhaul node, given that its transmission power was lower than the transmission power of the FAP. After a few iterations, the throughput in the FAP and in the FGW reached the value of around 17 Mbit/s, as seen in Fig. 5b. Nevertheless, the instabilities with the Minstrel-HT RA algorithm were evident, with fluctuations occurring during the Modulation and Coding Scheme data rate changes.


Figure 5a also presents the baseline solution, with the optimal trajectory to the optimal final position, determined as the point that ensures the SNR values in both links are the same, following Eq. 3, where  $P$  refers to the transmission power,  $G$  to the antenna gain,  $D$  to the distance between nodes and  $f$  to the frequency of operation. The throughput variation for the baseline solution, shown in Fig. 5c, follows a pattern similar to the one achieved with the RARL algorithm. Still, RARL converges faster (cf. Fig. 5b), validating the RA aware implementation of the RARL algorithm, as it was able to quickly meet the objective of throughput convergence and maximization in the different links.

$$P_{Rx} = P_{Tx} + G_{Tx} + G_{Rx} + 20 \log \left( \frac{c}{4\pi Df} \right) \quad (3)$$


**Moving FAP Scenario.** The FAP movement in the **moving FAP** scenario relied on the use of the Waypoint Mobility Model available in ns-3. For the evaluation presented herein, the movement was defined to evidence the behaviour regarding the FGW positioning when the FAP moves closer and further away. The following movement was defined:



(a) Trajectory of the FGW to final position.



(b) Throughput values evolution for the RARL algorithm.




(c) Throughput values evolution for the baseline model.


**Fig. 5.** Analysis of the RARL algorithm performance for the **asymmetric links** scenario.

1. The FAP is stationary during the first Waypoint, at (600, 600), for the initial 20s, as observed in Fig. 6a. The final results are:
  - Initial position of the FGW: (400, 400);
  - Final position of the FGW: oscillation between the coordinates (250, 450) and (275, 450);
  - Final distance between FGW and FAP: 380 m;
  - Final distance between FGW and Backhaul node: 500 m.
2. When it comes to the second Waypoint, represented in Fig. 6b, the movement of the FAP from (600, 600) to (1000, 1000). The FGW behaviour is described below:
  - Initial position of the FGW: (275, 450);
  - Final position of the FGW: the FGW adapts progressively to the FAPs position, finally oscillating between (550, 450) and (575, 450). The final position of this trajectory validates the RARL algorithm, given the proximity to the geometric centre between the FAP and the Backhaul node;
  - Final distance between FGW and FAP: 700 m;


- Final distance between FGW and Backhaul node: 730 m.
- 3. Finally, in the third part of the movement, displayed in Fig. 6c, the FAP moves to the coordinates (700, 300). The FGW behaviour is described below:
  - Initial position of the FGW: (575, 450);
  - Final position of the FGW: the final position of the FGW resulted in the oscillation between the coordinates (300, 450) and (325, 450).
  - Final distance between FGW and FAP: 400 m;
  - Final distance between FGW and Backhaul node: 550 m.



(a) First part of the movement.



(b) Second part of the movement.




(c) Third part of the movement.

**Fig. 6.** Trajectory of the FGW to final position for the **moving FAP** scenario.

When it comes to the evolution of the throughput values throughout the simulation in the **moving FAP** scenario, it is possible to observe in Fig. 7a that the throughput in the FGW node remained overall constant, having a uniform behaviour. On the other hand, the throughput measured in the FAP suffered multiple variations. Given that the FAP movement begins at 20s, the fluctuations occurred just in the initial moments. As soon as the FGW is able to adjust the position to the movement, the throughputs converge, reaching around 17 Mbit/s. Due to the convergence occurring during the movement of the FAP, it

can be concluded that the FGW was able to successfully follow the trajectory of the FAP.

A baseline model was also tested to evaluate the RARL algorithm performance. In the baseline model, the FGW defines a trajectory that follows the movement of the FAP node, maintaining the same distance between the FAP and the Backhaul node. The throughputs in the receiving nodes are presented in Fig. 7b and lead to the conclusion that the RARL was able to learn how to converge and maximize the throughput values even when the environment is constantly changing. As an RA aware algorithm, the RARL algorithm was able to detect the throughput imbalances and ensure the convergence of the throughputs throughout the FAP, outperform the baseline solution.



(a) Throughput values evolution for the RARL algorithm. (b) Throughput values evolution for the baseline model.

**Fig. 7.** Throughput values evolution for the **moving FAP** scenario.



**Two FAPs Scenario.** When it comes to the **two FAPs** scenario, given the presence of an additional FAP node, the distribution of the nodes through the venue included the Backhaul node at (0, 500), FAP1 at (1000, 1000) and FAP2 at (1000, 0). The analysis of the implementation with the Minstrel-HT, shown in Fig. 8a, was based on placing the FGW at a considerable distance to all the nodes in this configuration, having the following results:

- Initial position of the FGW: (25, 25);
- Final position of the FGW: (325, 375);
- Final distance between FGW and FAP1: 920 m;
- Final distance between FGW and FAP2: 770 m;
- Final distance between FGW and Backhaul: 350 m.

Overall, in the **two FAPs** scenario, the first link between the Backhaul node and the FGW can act as a bottleneck, limiting the maximum throughput achievable in the FAPs. The throughput variation in the different nodes is shown in Fig. 8b, showcasing the effect of the applied underlying RA algorithm with the instabilities associated with data rate transitions. It is possible to observe


that the throughput in the FGW is constantly higher, showing values around 23 Mbit/s. However, around the moment when there is a transition to a higher data rate in FAP1, there is a significant fluctuation in all the values. Finally, the throughput values in the FAPs converge to approximately 11 Mbit/s, accomplishing the objective of reaching higher values.

A baseline optimal trajectory leading to the geometric centre of the three nodes is presented in Fig. 8a. This implementation was able to converge the throughput values throughout most of the simulation time. Nevertheless, the FGW eventually gets to a position where its throughput decreases, consequently decreasing the throughputs achieved at the FAPs. This evidences that the final position was not optimal, showcasing that the RARL algorithm indeed maximized the throughput values in the nodes, while minimizing imbalances.



(a) Trajectory of the FGW to the final position.

(b) Throughput values evolution for the RARL algorithm.



(c) Throughput values evolution for the baseline model.

**Fig. 8.** Analysis of the RARL algorithm performance for the **two FAPs** scenario.

## 5 Conclusions

Given the great impact on the network performance, the optimal positioning of the FGW is a critical element in the flying network design. This paper proposes the RARL algorithm, a DRL-based FGW positioning approach that considers the effect of two aspects that have been overlooked in the state of the art: the influence of underlying RA algorithms and the impact of the Backhaul network configuration.

The evaluation of the performance of the RARL algorithm was carried out in ns-3. It was possible to observe that the fluctuations and instabilities, associated with the influence of Minstrel-HT in the link metrics, were overcome. This is supported by the trajectory of the FGW leading to a maximization of the defined reward functions, despite potential interference caused by the underlying RA algorithm in the throughput measured throughout the displacement of the FGW.

The comparisons of the RARL algorithm with the baseline models demonstrate its capability of converging and maximizing the throughput values in the nodes in all the three scenarios studied, meaning that this work endorses the possibility of implementing an RA aware positioning algorithm for real-world deployments. As future work, we intend to enhance the trajectory performed by the FGW, as it was overall not fully optimized, in order to reduce resource consumption and the time to converge to the optimal position. Moreover, given that this work is an initial evaluation of the position optimization of UAVs considering the influence of RA algorithms, after performing the necessary steps into scaling the complexity of the network configuration and topology, a real-world implementation should be carried out. For this, after the model being trained in the simulation environment, it will be implemented in the FGW through transfer learning.

**Acknowledgements.** This work is co-financed by Component 5 - Capitalization and Business Innovation, integrated in the Resilience Dimension of the Recovery and Resilience Plan within the scope of the Recovery and Resilience Mechanism (MRR) of the European Union (EU), framed in the Next Generation EU, for the period 2021 - 2026, within project Produtech\_R3, with reference 60. The second author thanks the funding from FCT, Portugal under the PhD grant 2022.10093.BD.

## References

1. Almeida, E.N., Campos, R., Ricardo, M.: Traffic-aware UAV placement using a generalizable deep reinforcement learning methodology. In: 2022 IEEE Symposium on Computers and Communications (ISCC), pp. 1–6 (2022). <https://doi.org/10.1109/ISCC55528.2022.9912770>
2. Alzenad, M., El-Keyi, A., Yanikomeroglu, H.: 3-D placement of an unmanned aerial vehicle base station for maximum coverage of users with different QoS requirements. *IEEE Wirel. Commun. Lett.* **7**(1), 38–41 (2018). <https://doi.org/10.1109/LWC.2017.2752161>
3. Anokye, S., Ayepah-Mensah, D., Seid, A.M., Boateng, G.O., Sun, G.: Deep reinforcement learning-based mobility-aware UAV content caching and placement in

- mobile edge networks. *IEEE Syst. J.* **16**(1), 275–286 (2022). <https://doi.org/10.1109/JSYST.2021.3082837>
4. Arif, T.Y., Munadi, R., et al.: Evaluation of the minstrel-HT rate adaptation algorithm in IEEE 802.11 n WLANs. *Int. J. Simul. Syst. Sci. Technol.* **18**(1) (2017)
  5. Bayerlein, H., De Kerret, P., Gesbert, D.: Trajectory optimization for autonomous flying base station via reinforcement learning. In: 2018 IEEE 19th International Workshop on Signal Processing Advances in Wireless Communications (SPAWC), pp. 1–5 (2018). <https://doi.org/10.1109/SPAWC.2018.8445768>
  6. Cicek, C.T., Gultekin, H., Tavli, B., Yanikomeroglu, H.: UAV base station location optimization for next generation wireless networks: overview and future research directions. In: 2019 1st International Conference on Unmanned Vehicle Systems-Oman (UVS), pp. 1–6 (2019). <https://doi.org/10.1109/UVS.2019.8658363>
  7. Coelho, A., Fontes, H., Campos, R., Ricardo, M.: Traffic-aware gateway placement for high-capacity flying networks. In: 2021 IEEE 93rd Vehicular Technology Conference (VTC2021-Spring), pp. 1–6 (2021). <https://doi.org/10.1109/VTC2021-Spring51267.2021.9448966>
  8. Fotouhi, A., et al.: Survey on UAV cellular communications: practical aspects, standardization advancements, regulation, and security challenges. *IEEE Commun. Surv. Tutor.* **21**(4), 3417–3442 (2019). <https://doi.org/10.1109/COMST.2019.2906228>
  9. Lagum, F., Bor-Yaliniz, I., Yanikomeroglu, H.: Strategic densification with UAV-BSs in cellular networks. *IEEE Wirel. Commun. Lett.* **7**(3), 384–387 (2018). <https://doi.org/10.1109/LWC.2017.2779483>
  10. Li, Y.: Deep reinforcement learning: an overview. arXiv preprint [arXiv:1701.07274](https://arxiv.org/abs/1701.07274) (2017)
  11. Liu, X., Liu, Y., Chen, Y.: Reinforcement learning in multiple-UAV networks: deployment and movement design. *IEEE Trans. Veh. Technol.* **68**(8), 8036–8049 (2019). <https://doi.org/10.1109/TVT.2019.2922849>
  12. Queiros, R., Ruela, J., Fontes, H., Campos, R.: Trajectory-aware rate adaptation for flying networks. arXiv preprint [arXiv:2307.06905](https://arxiv.org/abs/2307.06905) (2023)
  13. Queirós, R., Almeida, E.N., Fontes, H., Ruela, J., Campos, R.: Wi-Fi rate adaptation using a simple deep reinforcement learning approach. In: 2022 IEEE Symposium on Computers and Communications (ISCC), pp. 1–3 (2022). <https://doi.org/10.1109/ISCC55528.2022.9912784>
  14. Sammour, I., Chalhoub, G.: Evaluation of rate adaptation algorithms in IEEE 802.11 networks. *Electronics* **9**(9), 1436 (2020)
  15. Saxena, V., Jaldén, J., Klessig, H.: Optimal UAV base station trajectories using flow-level models for reinforcement learning. *IEEE Trans. Cogn. Commun. Netw.* **5**(4), 1101–1112 (2019). <https://doi.org/10.1109/TCCN.2019.2948324>
  16. Xia, D., Hart, J., Fu, Q.: Evaluation of the minstrel rate adaptation algorithm in IEEE 802.11g WLANs. In: 2013 IEEE International Conference on Communications (ICC), pp. 2223–2228 (2013). <https://doi.org/10.1109/ICC.2013.6654858>
  17. Yin, W., Hu, P., Indulskaja, J., Portmann, M., Mao, Y.: Mac-layer rate control for 802.11 networks: a survey. *Wirel. Netw.* **26**, 3793–3830 (2020)
  18. Zeng, Y., Zhang, R., Lim, T.J.: Wireless communications with unmanned aerial vehicles: opportunities and challenges. *IEEE Commun. Mag.* **54**(5), 36–42 (2016). <https://doi.org/10.1109/MCOM.2016.7470933>
  19. Zhang, S., Zhang, H., He, Q., Bian, K., Song, L.: Joint trajectory and power optimization for UAV relay networks. *IEEE Commun. Lett.* **22**(1), 161–164 (2018). <https://doi.org/10.1109/LCOMM.2017.2763135>