

Shaping Rewards, Shaping Routes: On Multi-Agent Deep Q-Networks for Routing in Satellite Constellation Networks

Manuel M. H. Roth^{*1}, Anupama Hegde¹, Thomas Delamotte², Andreas Knopp²

¹*Institute of Communications and Navigation, German Aerospace Center (DLR), Oberpfaffenhofen, Germany*

²*Chair of Signal Processing, University of the Bundeswehr Munich, Neubiberg, Germany*

Effective routing in satellite mega-constellations has become crucial to facilitate the handling of increasing traffic loads, more complex network architectures, as well as the integration into 6G networks. To enhance adaptability as well as robustness to unpredictable traffic demands, and to solve dynamic routing environments efficiently, machine learning-based solutions are being considered. For network control problems, such as optimizing packet forwarding decisions according to Quality of Service requirements and maintaining network stability, deep reinforcement learning techniques have demonstrated promising results. For this reason, we investigate the viability of multi-agent deep Q-networks for routing in satellite constellation networks. We focus specifically on reward shaping and quantifying training convergence for joint optimization of latency and load balancing in static and dynamic scenarios. To address identified drawbacks, we propose a novel hybrid solution based on centralized learning and decentralized control.

1 Introduction

As larger and more complex Low Earth Orbit (LEO) Satellite Constellation Networks (SCNs) with Inter-satellite Links (ISLs) are proposed, effective routing for Non-Terrestrial Networks (NTNs) is becoming an increasingly important topic. To support broadband traffic, as well as the integration with terrestrial networks, novel and specifically tailored routing and network management techniques are required. Considerations for 6G networks include an adaptive, flexible NTN-integration, which highlights the need for new approaches to interact seamlessly with the AI-native network of networks [1]. The optimization goal at hand is finding the "best" routing configuration in a highly dynamic, physically large network with non-uniformly distributed traffic and its Quality of Service (QoS) requirements. In this context, Deep Reinforcement Learning (DRL)-based schemes have been identified to be particularly promising op-

timization methods for such network control problems [2]. The approaches combine the ability of reinforcement learning to solve control problems effectively with the capacity of Deep Neural Networks (DNNs) to approximate complex, high-dimensional functions. In this investigation, we look into the viability of such DRL schemes based on Deep Q-Networks (DQNs) for SCNs. Various Fully-Distributed Multi-Agent DRL (FD-MADRL) approaches have been proposed for SCNs in the literature [3, 4]. The idea is that each satellite acts as an individual agent and makes decisions based on its local observations. However, a careful design and reward shaping is required to avoid loops and to find coherent paths. Learning coordinated behaviors can pose a challenge due to the complex interactions which can arise from the policies of individual agents. Utilizing Centralized DRL (CDRL) techniques instead, enables a larger scope and thus facilitates the establishment of effective end-to-end routes. But, centralization is difficult due to the physical size of the network. To improve scalability, hierarchical strategies have been proposed [5]. A path is constructed by the assistance of different controllers (or group leaders) at different hierarchical levels.

We base our investigation on these proposed approaches. The main contribution of this work is the analysis of the FD-MADRL routing approach in different SCN scenarios, focusing on the joint optimization of latency and load balancing. Drawing from the results, we propose a hybrid routing solution based on centralized learning and decentralized control.

1.1 Routing in Satellite Constellation Networks

SCNs differ from terrestrial networks in several aspects. Most importantly, there are frequent ISL and Earth-Satellite Link (ESL) handover events due to the dynamic network topology. Furthermore, traffic requirements are non-uniformly distributed. As the con-

^{*}Corresponding author. E-Mail: manuel.roth@dlr.de

stellation moves relative to Earth, hot spots are geographical rather than topological. Moreover, the on-board processing is typically limited. For this investigation, we assume four ISLs on each satellite as in [6]. To evaluate the performance of a routing scheme, various QoS metrics can be considered, e.g. packet dropping rate, latency, throughput, and link utilization [6].

1.2 Reinforcement Learning Architectures

Markov Decision Processes (MDPs) provide a formal, mathematically idealized framework for modeling reinforcement learning problems [7]. In MDPs, an agent interacts at time step t with an environment by choosing an action a_t from a set of possible actions A . The interaction results in a new state s_{t+1} , and a reward r_t which the agent seeks to maximize. This results in a trade-off: agents must balance exploiting known rewards through greedy selection versus exploring unknown states and actions with potentially better long-term rewards. By decomposing the problem into immediate plus discounted future rewards using the discount factor γ , we can express the optimization by simpler, recursive sub-problems [7]. The estimated viability of each action for a given state, the Q-value, can thus be described by:

$$Q(s_t, a_t) = r_t + \gamma \max_{a_{t+1}} Q(s_{t+1}, a_{t+1}) \quad (1)$$

Building on this foundation, DQNs utilize DNNs to approximate the Q-function in complex routing environments. During training, agents interact with the environment and store their experiences in replay buffers for increased stability. While CDRL is based on a singular agent, in FD-MADRL, each satellite is assumed to act as an independent agent with its own model and only partial knowledge of the environment. While this is well-suited for dynamic environments, in which changes and disruptions have to be handled locally and quickly, it also introduces additional complexity. The end-to-end coherence and stability becomes questionable, the agents have to learn the behavior of other agents as well. Complex and unexpected interactions between agents can emerge.

For FD-MADRL, the state space consists of information about the adjacent link loads, the previous hop and the destination. The action space encompasses the available ISLs, so the available next hops.

1.3 Reward shaping

Defining suitable rewards is imperative to enforce the intended optimization behavior. For CDRL, it is possi-




Figure 1: Abstracted sub-network representation: satellites (green) with ISLs (green-dotted), and link load (color according to load level).

ble to formulate composite rewards which split feedback into global and more immediate local rewards [5]. At time step t , the global reward of an agent can capture relevant end-to-end characteristics. Specifically, the flow-specific latency $L(t)$, the maximum or average link utilization $U(t)$, and the resulting packet dropping rate $D(t)$ can be considered. Each metric is weighted by a factor, denoted as w_1 , w_2 , and w_3 . However, this reward is only received after a significant time lag, denoted by T , due to the multi-hop propagation delays. The global reward can thus be represented as:

$$r_{t+T}^G = w_1 U(t) - w_2 L(t) - w_3 D(t) \quad (2)$$

The local reward represents immediate feedback consisting of directly observable metrics in the scope of a node and its neighborhood. Thus, it is assumed to be received in the subsequent time step $t + 1$. In this context, local link thresholds $\delta_{th}^L(t)$, local delays $l^L(t)$, and packet losses $d^L(t)$ due to saturated links can be considered. As before, these metrics are weighted by factors, denoted as w_4 , w_5 , and w_6 . The local reward can thus be described by:

$$r_{t+1}^L = w_4 \delta_{th}^L(t) - w_5 l^L(t) - w_6 d^L(t) \quad (3)$$

In our investigation of FD-MADRL, we specifically focus on the effectiveness of such local rewards. For fully decentralized designs, composite rewards are typically not considered [3, 4]. The limited scope makes it difficult to evaluate the impact of individual actions of independent agents. Instead of using local delays, we utilize the estimated resulting decrease in distance towards the destination, as in [4]. We denote the value of this path reduction ψ . Most importantly, to form coherent paths, loops have to be punished, but only at the nodes which are actually causing them. In this case a punishment $-\Psi$ is received, which should be at least a magnitude higher than ψ . Including a similarly large reward, so $+\Psi$, for a successful path

improved training performance - even though it was not propagated to all nodes along the path.

To incentivise decentralized load balancing, a threshold-based approach proved effective. The utilized values of ξ are in the domain of ψ for balanced paths. Depending on which policy is more desired, the values for ψ and ξ can be varied. $-\Xi$ punishes the usage of saturated links, which results in packet drops. As we intend to strongly discourage this behaviour, $-\Xi$ is in the domain of $-\Psi$. The resulting local reward considered for FD-MADRL is thus the sum of:

$$r'_{t+1} = \begin{cases} +\Psi, & \text{if goal reached} \\ +\psi, & \text{if overall path reduction} \\ -\psi, & \text{if overall path extension} \\ -\Psi, & \text{if path not reduced \& loop} \\ +\xi_1, & \text{if link load} < 0.4 \\ -\xi_2, & \text{if } 0.4 < \text{link load} \leq 0.8 \\ -\xi_3, & \text{if link load} > 0.8 \\ -\Xi, & \text{if link saturated} \end{cases} \quad (4)$$

2 Results

2.1 Simulation Environment

To analyze the considered approaches, a specifically designed environment was created using the gym library [8]. The DQN functionality were implemented using Keras for Tensorflow [9, 10]. The resulting grid network represents a sub-network of an SCN, e.g. a cluster in a distributed architecture [6]. For the comparisons we use a cluster of 12 satellites as well as a cluster of 24 satellites with varied link loads. The larger cluster is shown in Fig. 1. While the loads are set randomly, they comply with observed link characteristics of plausible non-uniformly distributed traffic scenarios [6]. For this performance comparison, we focus on the latency on a per-hop basis, i.e. the hop count, as well as the link loads. Due to the limited scope of this manuscript, only the described reward design is investigated.

2.2 Performance Evaluation

2.2.1 Static Path Finding

Firstly, we investigate a static case. The goal is to find paths with suitable latency and link load characteristics. To this end, we compare the total rewards received during training in each episode for the 12-node and 24-node clusters to gain insights into scalability. As shown in Fig. 2, more than twice as many episodes are required to achieve maximum rewards consistently. When analyzing the results, we observed that




Figure 2: Smoothed rewards over training for different DQN-based approaches in static scenarios, highlighting differences in convergence.




Figure 3: Comparison of FD-MADRL (in green) and a rule-based multi-cost approach (in red) in terms of latency (number of hops, on the x-axis) and the resulting maximum link load on the chosen path (y-axis).

oftentimes sub-optimal decisions were made in the last hop. For instance, in the 24-node cluster shown in Fig. 1, to reach node 23 from node 4, going up to node 8 was learned, as the adjacent link has a lower load. This however can result in a route going over node 22 which then takes the almost saturated link to 23. As assumed, the limited scope of individual nodes can result in costly end-to-end decisions.

The performance of the approach in comparison to a dynamic Shortest Path First (SPF) scheme using a dynamic multi-cost approach based on Dijkstra’s algorithm [11] is shown in 3. The rule-based algorithm is designed to favor short paths with minimum load, and represents a typical state-of-the-art approach [6]. Links with high link loads ($> 80\%$ of capacity used) are actively avoided. The results show that the current reward design achieves short paths: for most routes FD-MADRL requires fewer hops. However, the approach tends to saturate links. While this static example may not reveal significant issues, it’s crucial to consider potential consequences in dynamic scenarios where a lack of foresight can result in decreased overall system performance.

2.2.2 State Evolution: Dynamic Link Loads

If we reflect the impact of previous decisions in the link loads, differences in training can be observed, as




Figure 4: Smoothed rewards over training for FD-MADRL with dynamic, mutable link loads.

shown in Fig. 4. The idea is to test the robustness of the approach by evolving the state space according to previously made decisions. This corresponds more to an actual SCN scenario, where routing decisions impact future states and link loads change dynamically. In a sense, we investigate subsequent snapshots of the constellation. To this end, we change the network to be affected by the previous action trajectory: an additional load of 20% is added to each utilized link of the previous episode. These additional loads are only episode-specific, to avoid saturation on all links.

In comparison to Fig. 2, the achieved rewards are lower and less stable in the given training setup. While this can be explained by higher link loads resulting in more negative rewards, the current setup is also not able to robustly respond to the more varied state space. While a high performance can be achieved for the 12-node cluster nonetheless, the results are unstable for the 24-node cluster. In this larger network, more scenarios can arise which lead to significant negative rewards: either by selecting saturated links or by inadvertently introducing loops. Moreover, learning to proactively avoid potential bottlenecks is difficult with fully decentralized agents. These results highlight the necessity for adjustments such as information exchange between nodes, or schemes specifically designed to enable cooperation between agents.

3 Discussion

The results show that FD-MADRL is capable of converging to correct routing policies for the considered scenario, in terms of hop count and load balancing. However, with increasing cluster size, and state complexity, the training is significantly more difficult. More varied QoS requirements and network dynamics will require even more complex state representations and reward structures. Since DRL-based approaches are expected to outperform conventional, rule-based

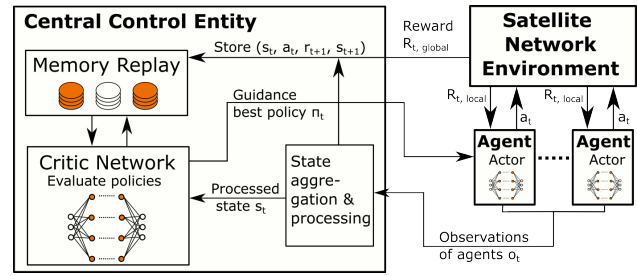


Figure 5: Proposed extended architecture: actor-critic approach using Centralized Learning and Decentralized Control (CL-DC). As in FD-MADRL actors make decisions locally, but receive central guidance from critic network.

algorithms in such complex scenarios, they are essential for the evaluation of the proposed schemes. For FD-MADRL, we have shown that the limited scope of each agent negatively impacts the end-to-end routes, in both static and dynamic scenarios. Moreover, with increasingly complex policies, the agents have to learn the behavior of other agents, which is counterproductive for scalability. A CDRL-based scheme on the other hand, loses the flexibility of decentralization, introduces a single point of failure, and needs additional signalling.

To balance these concepts, future research may consider a hybrid approach based on Centralized Learning - Decentralized Control (CL-DC). Such a scheme can improve coherence and end-to-end performance while maintaining the flexibility, robustness, and scalability of FD-MADRL. In this setup, a central controller provides guidance, while decisions are made locally by the agents. Local actions and rewards are forwarded to the central entity for learning. The centralized guidance improves stability by facilitating coordination between agents. Based on local conditions and constraints, the decentralized agents can make informed local decisions. As the Q-function generally requires the same information at training and execution time, future research may also focus on policy gradient techniques for multi-agent actor-critic architectures [12]. State aggregation and pre-processing may be considered to facilitate training. This hybrid approach combining aspects from FD-MADRL and actor-critic methods is depicted in Fig. 5.

In conclusion, while DQN-based architectures represent promising solutions for routing and network control in complex SCNs, they still have practical limitations. The investigated FD-MADRL approaches face scaling challenges, particularly in mastering complex scenarios where they are expected to surpass state-of-the-art methods. In-depth investigations are required to fully evaluate their actual viability.

Acknowledgment

This work has been supported by the Space for 5G & 6G programme of the European Space Agency (ESA), activity code 5A.079, <https://connectivity.esa.int/projects/aicoms>. Responsibility for the contents of this publication rests with the authors.

References

1. Muscinelli, E., Shinde, S. S. & Tarchi, D. Overview of Distributed Machine Learning Techniques for 6G Networks. en. *Algorithms* **15**, 210. issn: 1999-4893. <https://www.mdpi.com/1999-4893/15/6/210> (2023) (June 2022).
2. Zhang, C., Patras, P. & Haddadi, H. Deep Learning in Mobile and Wireless Networking: A Survey. *IEEE Communications Surveys & Tutorials* **21**, 2224–2287. issn: 1553-877X (2019).
3. You, X. *et al.* Toward Packet Routing With Fully Distributed Multiagent Deep Reinforcement Learning. *IEEE Transactions on Systems, Man, and Cybernetics: Systems* **52**, 855–868. issn: 2168-2232 (Feb. 2022).
4. Soret, B., Leyva-Mayorga, I., Lozano-Cuadra, F. & Thorsager, M. D. *Q-learning for distributed routing in LEO satellite constellations* tech. rep. arXiv:2306.01346 (arXiv, June 2023). <http://arxiv.org/abs/2306.01346> (2023).
5. Ali, R. E., Erman, B., Baştuğ, E. & Cilli, B. *Hierarchical Deep Double Q-Routing* tech. rep. arXiv:1910.04041 (arXiv, Mar. 2020). <http://arxiv.org/abs/1910.04041> (2022).
6. Roth, M. M. H., Brandt, H. & Bischl, H. *Distributed SDN-based Load-balanced Routing for Low Earth Orbit Satellite Constellation Networks* in *2022 11th Advanced Satellite Multimedia Systems Conference and the 17th Signal Processing for Space Communications Workshop (ASMS/SPSC)* (Sept. 2022), 1–8.
7. Sutton, R. S. & Barto, A. G. *Reinforcement learning: An introduction* (MIT press, 2018).
8. Brockman, G. *et al.* *OpenAI Gym* 2016. eprint: arXiv:1606.01540.
9. Chollet, F. *et al.* *Keras* <https://keras.io>. 2015.
10. Martín Abadi *et al.* *TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems* 2015. <https://www.tensorflow.org/>.
11. Dijkstra, E. W. A note on two problems in connexion with graphs. *Numerische mathematik* **1**, 269–271 (1959).
12. Lowe, R. *et al.* *Multi-Agent Actor-Critic for Mixed Cooperative-Competitive Environments* in *Advances in Neural Information Processing Systems* **30** (Curran Associates, Inc., 2017). <https://proceedings.neurips.cc/paper/2017/hash/68a9750337a418a86fe06c1991a1d64c-Abstract.html> (2023).