

Combinatorial Client-Master Multiagent Deep Reinforcement Learning for Task Offloading in Mobile Edge Computing

Tesfay Zemuy Gebrekidan
University of Southampton
Southampton, United Kingdom

Sebastian Stein
University of Southampton
Southampton, United Kingdom

Timothy J. Norman
University of Southampton
Southampton, United Kingdom

ABSTRACT

Recently, there has been an explosion of mobile applications that perform computationally intensive tasks such as video streaming, data mining, virtual reality, augmented reality, image processing, video processing, face recognition, and online gaming. However, user devices (UDs), such as tablets and smartphones, have a limited ability to perform the computation needs of the tasks. Mobile edge computing (MEC) has emerged as a promising technology to meet the increasing computing demands of UD. Task offloading in MEC is a strategy that achieves the demands of the UD by distributing the tasks between the UD and the MEC servers. Deep reinforcement learning (DRL) is gaining attention in task-offloading problems because it can adapt to dynamic changes and minimize online computational complexity. However, the various types of continuous and discrete resource constraints on the UD and MEC servers pose challenges to the design of an efficient DRL-based task-offloading strategy. Existing DRL-based task-offloading algorithms focus on the constraints of the UD, assuming the availability of enough storage resources on the server. Moreover, existing multiagent DRL (MADRL)-based task-offloading algorithms are homogeneous agents and consider homogeneous constraints as a penalty in their reward function. We proposed a novel combinatorial client-master MADRL (CCM_MADRL) algorithm for task offloading in MEC (CCM_MADRL_MEC) that enables UD to decide their resource requirements and the server to make a combinatorial decision based on the requirements of the UD. CCM_MADRL_MEC is the first MADRL in task offloading to consider server storage capacity in addition to the constraints in the UD. By taking advantage of the combinatorial action selection, CCM_MADRL_MEC has shown superior convergence over existing MADDPG and heuristic algorithms.

KEYWORDS

Multiagent Deep Reinforcement Learning; Combinatorial Action Selection; Mixed Constraints; Client-Master Multiagent Deep Reinforcement Learning

1 INTRODUCTION

Recently, there has been an explosion of mobile applications that perform computation-intensive tasks, such as video streaming, virtual reality, augmented reality, image processing, video processing, face recognition, and online gaming [1, 2, 10, 24]. However, UD, such as tablets and smartphones, have a limited ability to perform the

computation tasks of these applications. Mobile Cloud Computing (MCC) has been considered the key technology to meet the computation needs of UD by offloading their computation tasks to the cloud [14]. One of the challenges of MCC is latency caused by the distance of the MCC server from the UD [18]. MEC has emerged as a promising technology to address the challenges of MCC and the increasing computing demands of UD by providing the MCC service on the edge of the network.

Task offloading in MEC has become an attractive solution to cater to the diverse computing needs of UD [8] by distributing computational tasks between the UD and the MEC servers. Many existing task-offloading algorithms use traditional convex optimization methods for single-agent offloading scenarios [17]. DRL has recently gained attention in task offloading due to its advantage in reducing online computational complexity [16] and adapting to dynamic changes [11]. However, the existence of various types of resource constraints on UD and MEC servers, and the combination of discrete, continuous, and combinatorial action spaces pose challenges to the design of an efficient DRL-based task-offloading strategy. UD have limitations such as finite battery life and limited computational capabilities [10, 22], as well as quality of service (QoS) requirements such as latency. Similarly, the MEC servers come with storage constraints. DRL techniques, such as deep Q network (DQN), have yielded encouraging results by modeling the task-offloading problem as MDP with deep neural network (DNN) for function approximation [12]. However, due to the curse of dimensionality, DQN is insufficient for learning with large discrete action spaces [3] and a combination of continuous and discrete action spaces [23]. Although MADDPG-based task-offloading algorithms can handle continuous action spaces, the representation of discrete and continuous action spaces still poses a challenge [9, 23]. Despite the advances of MADRL in task offloading, such as cooperative offloading decisions [15] and mixed continuous and discrete action spaces [9, 23], most existing MADRL-based algorithms still formulate the constraints as a penalty in their reward function.

A comprehensive survey on task offloading by [8] has presented task offloading strategies in MEC from different perspectives, including the computational model, the decision-making entity, and the algorithm paradigm. Many algorithms have considered the wireless communication resource and the computing resource of the server. For example, the insufficient computing resource of the MEC server can be alleviated by using MEC-MCC collaboration or collaboration among multiple MEC servers [2]. Sub-channels are considered in the state and action spaces by some DRL-based task offloading algorithms [9, 15]. However, the storage constraint on the server is overlooked in existing DRL-based task-offloading algorithms.

The main contributions of this work are fourfold:

This preprint is a supplementary material for the extended abstract in [4] accepted in the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024).

- We proposed a novel CCM_MADRL algorithm for task offloading in MEC with various types of constraints at the UDs, the wireless network, and the server. Client agents are deployed at the UDs to decide their resource allocation and a master agent is deployed at the server to make combinatorial decisions based on the actions of the clients. The constraints of the UDs are considered as a penalty in the reward of the client agents whereas the channel and storage constraints are considered in the combinatorial decision of the master agent.
- By avoiding the number of sub-channels from the state and action spaces, and considering it as a constraint in the combinatorial action selection, we reduced the dimensionality
- This is the first DRL-based task offloading algorithm to consider combinations of continuous and discrete resource constraints on the UDs, the communication channel, and the storage capacity of the server.
- We develop different heuristic benchmarking methodologies and perform numerical analysis to determine the efficacy of the proposed algorithm.

2 SYSTEM MODEL

This section considers MEC for task offloading, which mainly includes a base station (BS), UDs, tasks, energy harvesting, and wireless networks. We consider a multi-user MEC scenario shown in Figure 1. In this scenario, there is a single wireless BS equipped with an MEC server that provides a computing and storage service and an SDN controller that controls communication between UDs and BS. The BS serves a set of $N = \{1, 2, 3, \dots, N\}$ UDs. A UD in the set N is denoted by n . For local processing, we consider that each UD n has a minimum and maximum computational resource allocation budgets denoted by f_n^{min} and f_n^{max} , respectively, in Gigahertz (GHz) cycles per second. Similarly, to offload its task to the server, we consider that each UD has a minimum and maximum transmission power allocation threshold denoted by p_n^{min} and p_n^{max} , respectively, in dBm. Furthermore, we consider a UD to have a minimum battery threshold of b_n^{min} and a maximum battery capacity of b_n^{max} in Mega-joules (MJ). The BS has multiple constraints and characteristics, such as the server storage constraint z_e in bits and the number of processing units on the server U_e each having an equal processing capacity of f_e in gigahertz cycles per second. Similar to [9, 15], we consider a wireless network of bandwidth of W in megahertz that is equally divided between K sub-channels. The list of notations and terms used in this work are presented in Table 1.

We consider a task-offloading problem for T time steps of τ_{max} length each. It is assumed that each UD n generates one task at each time step. If the processing of a task is not completed in τ_{max} , it is discarded before the next time step starts. The task model, the processing model, and the energy harvesting are described in the following sections.

The task offloading model is designed by combining the settings of different existing works. Since [15] has used a data set from Huawei Technologies, we adapted it disregarding the blockchain part. We considered the energy harvesting process in [23], and took advantage of the efficient estimation of the completion times of tasks by [21], which computes the completion time of tasks on the server based on the completion time of other tasks scheduled before

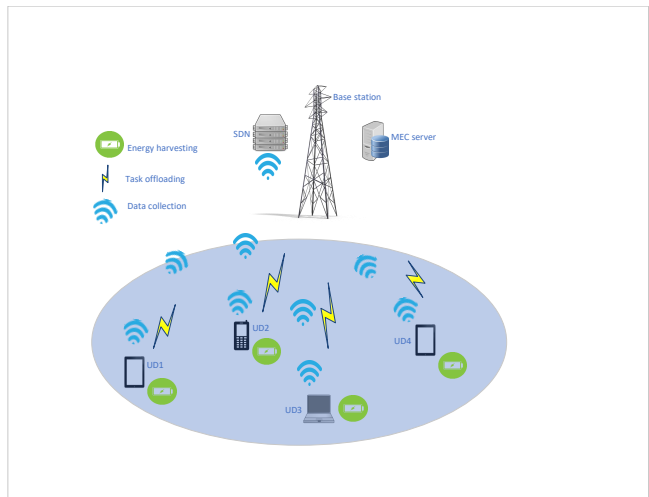


Figure 1: Network model

them. Similarly, we assume that the server processes the tasks in the order of their arrival on the server. The arrival times of the tasks are determined by their offloading times. In the following, we present the task and computing model for the task offloading problem.

2.1 Task Model

The task model is based on the setting by [15] on a data set from Huawei Telecom. At each time step, each UD n generates a task denoted by its notation as n^1 that is represented by characteristics such as the size of the task z_n in bytes, the number of CPU cycles per bit required to process the task c_n , and the maximum deadline τ_n to which task processing is expected to finish.

Before processing the task, there are three decision variables: a binary decision of whether to process it locally or offload it to the MEC server x_n , a local resource allocation f_n , and a transmission power allocation p_n , which are described in detail in Section 4.2. So, we give a binary decision $X = \{x_n | n \in N\}$ to describe the processing mode, as seen in Equation (1).

$$x_n = \begin{cases} 1, & \text{MEC processing} \\ 0, & \text{Local processing} \end{cases} \quad (1)$$

Next, we present the local and MEC models.

2.2 Local Processing

A task is processed locally if one of the following two conditions happens as presented in Section 4.2: if the UD decides to process the task locally; if a UD proposes the task to the master agent but the master agent did not select it to be offloaded in the combinatorial action selection. Then, the UD processes the task using its local computational resource assigned to its task, which is restricted within its own resource allocation budget as $f_n | f_n^{min} \leq f_n \leq f_n^{max}$. The local computing latency to process the task is computed as $T_{loc_n} = \frac{z_n c_n}{f_n}$. The energy consumption in the local processing mode is calculated

¹Because a UD has one task at a time step, we use n to denote both the UD and its task to reduce the number of notations

Notation	Description
N	Set of UDs
T	Number of time steps
τ_{max}	Maximum length of a time step
n	A UD or a task of the UD
T_n	Total latency of processing task n
x_n	Binary indicator of local or offload for task n
E_n	Energy consumption of task n
E_{loc_n}	Energy consumption of task n in local processing
E_{off_n}	Energy consumption of offloading task n
T_{loc_n}	Computation time for the local processing of task n
T_{off_n}	Offloading time of task n
λ_1 & λ_2	weight coefficients of T_n and E_n
b_n	Battery level of US n
p_n	Transmission power allocation of the UD n
z_n	Size of task n
z_e	Storage capacity of the server
U_e	Number of processing units in the server
L_n	Cost of processing task n
R_n	Reward of processing task n
f_n	Resource allocation for local processing of task n
C_n	Number of CPU cycles to process one bit of task n
α_ϕ	Learning rage of the master agent
α_θ	Learning rage of client agent
J	Joules

Table 1: List of notations and terms

based on the size of the task and the allocation of resources for processing the task, as shown in Equation (2).

$$E_{loc_n} = \kappa z_n c_n f_n^2 \quad (2)$$

where κ is energy consumption coefficient [23].

2.3 MEC Processing

In this mode, the task is transferred to the MEC server to be processed by one of the processing units U_e of the server. The decision happens when the UD proposes the task to be offloaded and the master agent accepts it. To be processed on the server, the task needs transmission resources, which is a function of the transmission power p_n . The transmission power p_n is decided by the UD from its transmission power budget $p_n | p_n^{min} \leq p_n \leq p_n^{max}$ as discussed in Section 4.2. Then, the data transmission rate d_n in a single channel of the wireless network is calculated using Shannon's capacity as

$$d_n = \frac{W}{K} \log_2 (1 + p_n g_n) \quad (3)$$

where $g_n = h_n / \sigma^2$ is the normalized channel gain of the uplink channel between UD n and the BS, with channel gain h_n and the background noise variance σ^2 . The channel gain h_n is impacted by many factors, including distance. For simplicity, we assume that the UDs are stationary and have a stationary normalized channel gain depending on their distance from the BS. The variance of background noise σ^2 is also constant. We did not consider interference between multiple UDs because we assume that a channel is used by one task at a time.

Once the data transmission rate is determined, the transmission time T_{off_n} is computed as:

$$T_{off_n} = \frac{z_n}{d_n} \quad (4)$$

Then, the energy consumption of offloading task n to the server is calculated as $E_{off_n} = p_n T_{off_n}$. Like many works on task offloading [15, 20, 23], we assume that the communication resource required to return the information about the processed task to the UD is negligible.

Note that the energy consumption in task offloading is computed only for the UDs as they are battery-powered. However, the latency of processing the tasks on the server matters because the tasks have deadline constraints. Therefore, the total latency of processing a task on the server is determined by the transmission time, the earliest availability of the processing unit on the server, and the time required to process the task on the server. The processing time of task n in one of the processing units in the server is computed as $T_{ser_n} = \frac{z_n c_n}{f_e}$. However, the processing of the task on the server does not start as soon as the task has arrived at the server. The processing units on the server process one task at a time. Tasks offloaded to the server are processed in the order of their arrival at the server, which is determined by T_{off_n} . The tasks are assigned to the earliest free processing unit. Therefore, the start of processing task n depends on the earliest availability of a processing unit, which is determined by the number of processing units on the server U_e , and T_{ser_n} and T_{off_n} of other tasks that have shorter T_{off_n} than that of task n . Accordingly, the total latency of the offloading task n to the MEC server T_{MEC_n} is computed as $T_{MEC_n} = T_{ser_n} + \max(T_{off_n}, T_{ear_n})$ where T_{ear_n} is the estimated availability time of the first available processing unit U_e of the server after the arrival of task n and Max . ensures that the processing of the task starts when a free processing unit is found after the offloading of the task is finished. The T_{ear_n} is calculated based on the completion time of other accepted tasks on the server with the earliest offloading time than that of task n . This estimate is adapted from the work of [21].

2.4 Energy Harvesting

The energy harvesting process is adapted from the work in [23]. For simplicity, we assume that the UDs harvest e_n energy at the beginning of each time interval. Initially, each UD is full with its maximum battery capacity of b_n^{max} . The level of the battery in the next time interval depends on both energy consumption and harvesting, which evolves according to the following equation in the T time steps.

$$b_n(t+1) = \min(\max(b_n(t) - E_n(t) + e_n(t), 0), b_n^{max}) \quad (5)$$

where E_n is the energy consumption calculated based on E_{loc_n} and E_{off_n} as described in Section 3 and $\min(\cdot)$ and $\max(\cdot)$ ensure that the level of the battery cannot be negative and does not exceed the maximum capacity.

3 PROBLEM FORMULATION

The total processing latency is $T_n = (1 - x_n) T_{loc_n} + x_n T_{MEC_n}$ and the energy consumption is equal to $E_n = (1 - x_n) E_{loc_n} + x_n E_{off_n}$. Considering that the cost of processing a task is collectively determined

by its processing latency and energy consumption, the cost function for processing a task is specified as $L_n = \lambda_1 T_n + \lambda_2 E_n$ where λ_1 and λ_2 are weight coefficients. The CCM_MADRL_MEC is aimed to solve the optimization problem that can be formulated as the cost minimization for all UDs and T time steps while meeting the different constraints at the UDs and the server as follows:

$$\text{minimize } \sum_t \sum_{n \in N} L_n(t) \quad (6a)$$

$$\text{subject to } x_n \in \{0, 1\}, \quad \forall n \in N \quad (6b)$$

$$p_n^{\min} \leq p_n \leq p_n^{\max}, \quad \forall n \in N \quad (6c)$$

$$T_n \leq \tau_n, \quad \forall n \in N \quad (6d)$$

$$b_n \geq b_n^{\min}, \quad \forall n \in N \quad (6e)$$

$$f_n^{\min} \leq f_n \leq f_n^{\max}, \quad \forall n \in N \quad (6f)$$

$$\sum_{n \in N} x_n \leq K, \quad \forall n \in N \quad (6g)$$

$$\sum_{n \in N} x_n z_n \leq z_e, \quad \forall n \in N \quad (6h)$$

where Equation (6b) implies that a task is processed locally or uploaded to the MEC server, Equation (6c) indicates that the transmission power should be between the power allocation budget, Equation (6d) ensures that the processing time of each task cannot exceed its processing deadline, Equation (6e) guarantees that the battery level should not exceed the low battery level, Equation (6f) ensures that the local computational resource allocated to each task should be in the preset minimum and maximum values, Equation (6g) and ensures that the number of offloaded tasks does not exceed the number of sub-channels by ensuring that only one task uses a channel. It is used if and only if it is necessary to use only one channel for one user as used by [10]. Equation (6h) guarantees that the sum of the sizes of the off-loaded tasks does not exceed the storage capacity of the server.

4 COMBINATORIAL CLIENT-MASTER MADRL ALGORITHM FOR TASK OFFLOADING IN MEC

To solve the optimization problem of the cost minimization in Equation (6a), we convert the optimization problem into a reward maximization problem and apply CCM_MADRL_MEC. The states, client and master actions, and the formulation of the reward function are presented as follows.

4.1 State

The state $S(t)$ of the MEC environment at time t , which includes the set of states of the UDs, is described as $S(t) = \{S_n(t)\}, \quad \forall n \in N$. Constant values such as the number of sub-channels K , the number of processing units on the server U_e , the processing capacity f_e , and the storage capacity z_e of the server are excluded from the state

information. The state of a UD, $S_n(t)$, is characterized by five components: task state $S_n^{\text{task}}(t)$, normalized channel gain state $S_n^{\text{gain}}(t)$, power transmission budget $S_n^{\text{pow}}(t)$, local resource allocation budget $S_n^{\text{res}}(t)$, and battery state $S_n^{\text{battery}}(t)$ as defined in Equation (7):

$$S_n(t) = \{S_n^{\text{task}}(t), S_n^{\text{gain}}(t), S_n^{\text{pow}}(t), S_n^{\text{res}}(t), S_n^{\text{battery}}(t)\} \quad (7)$$

where $S_n^{\text{task}}(t) = [z_n(t), c_n(t), \tau_n(t)]$, $S_n^{\text{gain}}(t)$ is $g_n = h_n/\sigma^2$ as described in Section 3, $S_n^{\text{pow}}(t) = f_n^{\max}$, $S_n^{\text{res}}(t) = f_n^{\max}$, and $S_n^{\text{battery}}(t)$ is as described in Equation (5).

4.2 Action

At the beginning of each time step, the UDs make decisions about their resource allocations using client agents. Then, the SDN controller collects information about the state and action of the UDs and performs one of the following three procedures using a master agent: 1) for the UDs that decide to make a local processing, the server does not interfere. 2) if the number of UDs that are proposed to offload is greater than the number of sub-channels or if the sum of their sizes is greater than the capacity of the storage capacity of the server, the server makes a combinatorial decision on which of the requests of the UDs to approve and which of them to reject. 3) If the proposed requests are less than the constraints, the server accepts all of them. Finally, sub-channels are assigned to accepted UDs, and then the task offloading and processing process starts.

Existing DRL-based task offloading algorithms, such as [15] and [9], included the number of sub-channels in their state and action space. However, the sub-channels have equal transmission capacity from the perspective of a UD as seen in Equation (3). If we restrict that a channel is used by only one UD at a time, it does not matter which channel a UD uses. Therefore, the inclusion of sub-channels in the state and action space incurs a dimensionality problem without playing any significant role. We excluded channel information from the state and action space and considered them as a constraint in the combinatorial action selection. Note that a channel can be reused by multiple UDs one after the other, by lifting the constraint that one channel must be used by one channel. In such a case, only the storage capacity of the server becomes the constraint in the combinatorial action selection.

The actions of the client agents and master agent are as follows.

Client Actions. At each time step, each client agent produces three actions, which are all continuous value actions between $[0, 1]$ inclusive. The action space can be expressed as:

$$A_c(t) = x_{c,n}(t), p_{c,n}(t), f_{c,n}(t) \leftarrow \theta_n(S_n(t)), \quad \forall n \in N \quad (8)$$

where $S_n(t)$ is the state of UD n as described in Equation (7), and θ_n is the parametrized policy function of the client, $x_{c,n}(t)$ is the task offloading decision by client n (if $x_{c,n}(t) < 0.5$ then x_n in Equation (1) becomes 0 otherwise the task is proposed to be considered for the combinatorial decision), $p_{c,n}(t)$ is the client action that decides the transmission power using Equation (9), and $f_{c,n}(t)$ is the action that decides the local computational resources allocation using Equation (10).

The actions of the client agents determine p_n and f_n are follows:

$$p_n = \max \left(p_n^{\min}, p_{c,n}(t), p_n^{\max} \right) \quad (9)$$

$$f_n = \max \left(f_n^{\min}, f_{c,n}(t), f_n^{\max} \right) \quad (10)$$

Master Action. For the client actions with $x_{c,n}(t) \geq 0.5$, the master takes the combinations of states and actions of the clients and provides a binary output for the combinatorial decision on which of them should be allocated locally and which of them should be accepted for processing by the MEC server.

$$A_m(t) = x_{m,n}(t) \leftarrow \phi(S, A, S_c, A_c) \quad (11)$$

where S and A are the set of states and actions of all client agents, and S_c and A_c are the set of states and actions of the client agents whose $A_{c,n} \geq 0.5$ and ϕ is the policy of the master.

The combinatorial action selection in the master agent is built by modifying the critic in the classical multiagent deep deterministic policy gradient (MADDPG) algorithm [13] and incorporating the per-action DQN [7].

Figure 2 shows the interaction diagram of the CCM_MADRL algorithm and the MEC system. Client agents represent the policies of the UDs. The master agent represents the policy on the MEC server. The environment represents the allocation of resources in the UDs and on the server. After a client produces its output, it does the following as mentioned in Section 4.2: if $x_{c,n} < 0$ assigns $x_n = 0$ and starts the local allocation. Otherwise, it forwards $x_{c,n}, p_n, f_n$ to the master for the combinatorial decision. Then, the master agent produces the binary decision and applies it to the UDs and the server. Finally, a shared reward is computed and provided to the master agent to train its value function. The client agents are also trained using a TD error computed by the master agent as feedback.

4.3 System Reward Function

To compute the reward, we use the negative of the objective function and we compute a penalty function for the constraints of the tasks and UDs. The constraints of the server are already considered in the combinatorial action selection and do not need to be incorporated as a penalty. In [23], they included a dropoff penalty in their reward for running out of batteries. In ours, we start the penalty from a preset minimum battery threshold.

$$L'_n = \lambda_1 \min((\tau_n - T_n), 0) + \lambda_2 \min((b_n - b_n^{\min}), 0) \quad (12)$$

Since the design of our cooperative learning formulation is based on the cost minimization problem in Equation (6a), our system reward function is equal to the negative of the system cost function and the penalty function. Thus, we can formulate the system reward function as follows:

$$\bar{r}(S(t), A(t)) = -\frac{1}{|N|} \sum_{n \in N} L_n(t) - L'_n(t) \quad (13)$$

The power and resource allocation decisions made in the current step by the UDs affect their operational life in the next time steps by affecting energy consumption. Therefore, the DRL must consider




Figure 2: The interaction diagram of the agents and the MEC environment. Client agents output their actions $\{x_n, p_n, f_n\}$. Clients with $\{x_{c,n} < 0.5\}$ start local processing; and the others propose their tasks to the master agent, who makes the combinatorial decision on which of the proposed tasks should be offloaded and which of them should be designated for local processing

immediate reward and long-term return using the Bellman equation as shown in Equation (14)

$$Q(S, A, S_n, A_n | \phi) = (1 - \alpha_\phi) Q(S, A) + \alpha_\phi (R(S, A) + \gamma_n P(S' | S, A) \max_{S'_n, A'_n} Q(S', A' | \phi')) \quad (14)$$

where α_ϕ is the learning rate, γ is the discount factor, ϕ is the policy of the critic, ϕ' is the policy of the target critic, $S = \{S_1, \dots, S_N\}$ and $S' = \{S'_1, \dots, S'_N\}$ are combined current and next states, $A = \{A_1, \dots, A_N\}$ and $A' = \{A'_1, \dots, A'_N\}$ are combined current and next actions of the client agents, and S'_n and A'_n are the corresponding states and actions of the agents. The role of the four parameters in Q' is presented in the following section.

4.4 The master Agent with per-client DQN

In MADDPG, there is only a single Q-value for the combined state and action pair of all actors, which is calculated as:

$$Q(S, A) = (1 - \alpha_\phi) Q(S, A | \phi) + \alpha_\phi (R(S, A) + \gamma Q(S', A' | \phi')) \quad (15)$$

To customize the critic to select combinatorial actions, it should be able to provide a Q-value per each actor (or client in the case of CCM_MADRL). Therefore, the master agent with per-client DQN in the CCM_MADRL adapts the concept of per-action DQN where the state S_n and the action A_n of each client agent are appended to the combined state and the action of the client agents to calculate the relative Q value in the combination of the state and the action as seen in Equation (14). The combined rewards are given only to selected clients in the task offloading problem, as seen in Algorithm 3 so that they will have different Q values to distinguish them in action selection. The reward is used by the master agent to train a value

function for the selected clients using per-client DQN. Client agents are also trained using a TD error computed by the master agent as feedback.

The master agent applies the coalition action selection approach in [5]. However, for ease of benchmark comparison with MADDPG, it uses per-action DQN instead of a transformer neural network.

4.5 Algorithms

Since the master agent in the CCM_MADRL_MEC algorithm has two roles: providing feedback for training the clients, similar to the MADDPG, and participating in the combinatorial action selection of the clients, it follows different procedures for both. Therefore, the algorithm is presented below in three parts: a main algorithm, an action selection algorithm, and a training algorithm in the following sections.

Main Algorithm. The main algorithm runs the action selection algorithm, the training environment, and the evaluation environment. At each episode, it runs for T time steps as seen in lines 7 to 12 of Algorithm 1. At each step, the action selection algorithm is called, and the rewards are computed, and then the experience is recorded to replay memory. When the iteration over the steps is complete, the training algorithm is called and the trained policies of the client agents and the master agent are evaluated.

Action Selection Algorithm. Because the action selection algorithm has to use exploration in the training environment, we present the computation of ε , which is used for ε -greedy to determine whether to explore a new action or to exploit the learned knowledge in the master agent and to scale the noise in the client agents as seen in Algorithm 2. At each episode, ε is updated using Equation (16).

$$\varepsilon = \text{Min_Epsilon} + (\text{Max_Epsilon} - \text{Min_Epsilon}) \cdot e^{-\frac{\text{episode}}{\text{Max_Episodes}}} \quad (16)$$

where Min_Epsilon and Max_Epsilon are the minimum and maximum values of the decaying epsilon, episode , is the current episode, and Max_Episodes is the maximum number of episodes.

The action selection algorithm applies the exploration of actions for the client agents and the master agent as follows. Note that the evaluation flag is used to indicate whether the actions are running for the training environment or for the evaluation environment. In the evaluation, exploration is not needed. For client agents, the exploration is performed by adding noise to the actual output of the client agents, as seen in lines 5 to 8 of the Algorithm 2. After adding noise to the actual action, the values are clipped to $[-1, 1]$ so that they are within the activation function of the client agents, Tanh in this case. All actions, explored or exploited, are scaled to be between $[0, 1]$ before applying to compute the resource allocation in Section 4.2.

The master agent follows ε -greedy for exploration and exploitation as seen in lines 41 to 47. First, a random number is generated as seen in line 11 to decide whether to explore or exploit. If the number is less than ε , the master agent shuffles the proposed actions as seen in line 45 and follows the combinatorial action selection procedure described below. Otherwise, the master agent computes the Q value based on the states and actions of the proposed actions and appends

Algorithm 1 CCM_MADRL main Algorithm

```

1: Initialize  $\text{Max\_Episodes} = 2000$ ,  $\text{Min\_Epsilon} = 0.01$ ,
    $\text{Max\_Epsilon} = 1$ ,  $\gamma = 0.99$ 
2: Initialize client agents  $\theta_n \quad \forall n \in N$  and the master agent  $\phi$  with
   random weights
3: Initialize target client agents  $\theta'_n \leftarrow \theta_n$  and the target master agent
    $\phi' \leftarrow \phi$ ,  $\forall n \in N$ 
4: Initialize replay memory  $RM$ 
5: for episode = 1 to  $\text{Max\_Episodes}$  do
6:   Reset environment and get initial state  $S_n(t=1)$ ,  $\forall n \in N$ 
7:   for  $t = 1$  to  $T$  do
8:     Go to Algorithm 2 using evaluation = False flag to select
       client and master actions
9:     Execute actions and observe total reward  $\bar{r}(t)$  and next
       state  $S_n(t+1)$ ,  $\forall n \in N$ 
10:    Store transition  $(S_n(t), A_n(t), \bar{r}(t), S_n(t+1))$ ,  $\forall n \in N$ 
       in to  $RM$ 
11:    Update the state  $S_n(t) \leftarrow S_n(t+1)$ ,  $\forall n \in N$ 
12:  end for
13:  Go to Algorithm 3 for training
14:  for EvalEpisode in EvalEpisodes do
15:    Reset and seed episode to EvalEpisode and find state
        $S(t=1)$ 
16:    for  $t = 1$  to  $T$  do
17:      Go to Algorithm 2 using evaluation = True flag to select
        client and master actions
18:      Execute actions and observe total reward  $\bar{r}(t)$  and next
        state  $S_n(t+1)$ ,  $\forall n \in N$ 
19:      Update the state  $S_n(t) \leftarrow S_n(t+1)$ ,  $\forall n \in N$ 
20:    end for
21:  end for
22: end for

```

the Q value along with the identifiers of the tasks n to Qs and $Index$ and follows the combinatorial action selection procedure.

After the actions of the client agents are provided, the master agent follows one of the following three procedures in the exploitation mode as described in Section 4.2. If all client agents decide to process their tasks locally, as in line 18, the master agent does not intervene. Line 20 computes the Q values of the proposed tasks using per-client DQN, and appends them to Qs and $Index$ along with their identifiers n , to be considered in the combinatorial decision. If the number of proposed tasks or the sum of their sizes is less than the number of sub-channels and the storage constraint on the server, the server accepts all of them as seen in lines 24 to 26. If the number of proposed tasks is greater than the number of sub-channels or if the sum of their size is greater than the storage capacity of the server, the master agent uses the Q-values computed using the states and actions of the client agents that proposed to offload their tasks to make decisions. It starts to approve the proposed actions of the clients with the highest Q-values until the number of sub-channels or the storage constraint is met. The remaining agents are designated to process their tasks locally. The algorithm is provided in Algorithm 2. The procedure of the exploitation is provided similarly

to the procedure of exploitation except that the proposed actions are shuffled randomly rather than getting sorted by their Q-values.

Training Algorithm. The algorithm for training the client agents and the master agent is provided in Algorithm 3. Because the structure of the master agent is different from the critic of MADDPG as seen in Section 4.4, Algorithm 3 is significantly different from existing MADDPG training algorithms in that: It generates multiple Q-values rather than one combined Q-value, because the master agent has to make a combinatorial decision using the relative Q-values of the clients as seen in lines 7,17, and 33; The client agents are trained by computing the highest Q value from the tasks offloaded to the server as seen in lines 13 and 38; If all tasks are allocated locally by the clients the master agent uses *all_zeros* as a placeholder to hold the combined Q-value as seen in lines 11, 22, and 35. Note that after the master agent decides which of the tasks should be processed in the server as described in Section 4.2, the reward is applied only to the selected tasks when training. I.e. the reward is computed at a system level using Equation (13) but it is used only by the tasks offloaded to the server when training the algorithm so that the tasks are distinguished by their Q-values in the action selection algorithm. This technique is adapted from the coalition action selection in [5].

The notations of the client and master actions are changed in the training algorithm because of the subscript i for the minibatch which is used to iterate over the entries of the minibatch of size M . Unlike MADDPG, which computes the Q-values of the minibatch as a batch, the Q-values of the minibatch in the CCM_MADRL are computed individually because they are processed conditionally as seen with many if clauses in the algorithm. The training algorithm starts by selecting a minibatch of size M from the replay memory. Each entry of the minibatch includes the combined state S and action A of all client agents S , the set of binary actions of the tasks by the master agent A^{mas} , the total reward of the tasks \bar{r} , the combined next state S' , and a flag that indicates whether the episode was terminated or not *done*.

The master agent is trained by lines 2 up to 28. Line 2 computes the target action for every client and every entry of the minibatch using their next state. The target action is to be used to compute the target Q-value using the master agent. Then, lines 3 and 4 concatenate the target actions of the client agents because the master agent accepts a combined state and action of all clients as input as seen in line 7. Lines 5 and 6 check if the client agents have decided to process the tasks locally or propose them to the master agent. For each task that was proposed to the master, a relative Q-value is computed in line 7 and the maximum Q-value will be computed in line 13. Line 9 concatenates the Q-values of the offloaded tasks in the same entry of a minibatch. If no task was offloaded, a Q-value will be computed using a placeholder to train the master agent so that it is used to give feedback like classical MADDPG. Then, the combined reward is provided to the actions selected to offload their tasks as seen in lines 16 to 20. The current and target Q-values in lines 17 and 17 are used to compute the TD error in line 26. As seen in lines 21 to 25, if all of the client agents, at any entry of the minibatch, decide to process their tasks locally, the master agent concatenates the state and action of all agents and appends *all_zeros* as a placeholder to learn the Q value when all tasks are processed

Algorithm 2 The action selection algorithm for the client agents and the master agent

```

1: Input: state  $S_n$  for each client agent  $n$ ,  $\{K, z_e\}$ , and Evaluation flag
2: Output: client actions  $A_n$  for each  $n$  with  $x_n$  decided by collaboration with the master agent
3: Get action  $A_n \leftarrow \pi_n(S_n, \theta_n)$ ,  $\forall n \in N$ 
4: if Evaluation == False then
5:   Compute  $\varepsilon$  using Equation (16)
6:    $noise = \text{random}(|N| \text{ by } |A_n|) * \varepsilon$ 
7:    $A_n = A_n + noise_n$ ,  $\forall n \in N$ 
8:   Clip  $A_n$  to  $[-1, 1]$ 
9: end if
10: Scale  $A_n$  to  $[0, 1]$  using  $\frac{A_n}{2} + 0.5$ ,  $\forall n \in N$ 
11: Generate a random number
12:  $Qs = []$ ,  $Index = []$ 
13: if  $random < \varepsilon$  or Evaluation == True then
14:    $S = \{S_n\}$ ,  $A = \{A_n\}$ ,  $\forall n \in N$ 
15:   for  $n \in N$  do
16:     Get  $x_{c,n}$  from  $A_n$  as described in Section 4.2
17:     if  $x_{c,n} < 0.5$  then
18:        $x_n = 0$ 
19:     else
20:        $Append(Qs, Q(S, A, S_n, A_n, \phi))$ ,  $Append(Index, n)$ 
21:     end if
22:   end for
23:   if  $Length(Index) \leq K$  and  $Sum(z_n \forall n \in N \text{ and } x_{c,n} \geq 0.5) \leq z_e$  then
24:     for  $\forall n \in N$  and  $x_{c,n} \geq 0.5$  do
25:        $x_n = 1$ 
26:     end for
27:   else
28:     Sort  $Qs$ , and adjust  $Index$  accordingly
29:      $TotalSizeOfAccepted = 0$ 
30:     while  $Length(Index) > K$  do
31:        $n = Pop(Index)$ 
32:        $x_n = 1$ 
33:       if  $TotalSizeOfAccepted + z_n \leq z_e$  then
34:          $TotalSizeOfAccepted = TotalSizeOfAccepted + z_n$ 
35:       else
36:          $x_n = 0$ 
37:       end if
38:     end while
39:   end if
40: else
41:   Collect the index  $n$  of the tasks with  $x_{c,n} \geq 0.5 \forall n \in N$  to  $Index$ 
42:   if  $Length(Index) \leq K$  and  $Sum(z_n \forall n \in N \text{ and } x_{c,n} \geq 0.5) \leq z_e$  then
43:     Execute lines 24 to 26
44:   else
45:     shuffle  $Index$  in to random order
46:     Execute lines 29 to 38
47:   end if
48: end if

```

Algorithm 3 The training algorithm for the client agents and the master agent

```

1: Sample a random minibatch of transitions
   ( $S, A, A^{mas}, \bar{r}, S', done$ ) of size  $M$  from  $RM$ 
2: Set target actions  $A'_{i,n} \leftarrow \pi_n(S'_{i,n}, \theta'_n)$ ,  $\forall n \in N$ , and for  $i = 1$ 
   to  $M$ 
3:  $S'_i = \{S'_{i,n}\}$ ,  $\forall n \in N$ , and for  $i = 1$  to  $M$ 
4:  $A'_i = \{A'_{i,n}\}$ ,  $\forall n \in N$ , and for  $i = 1$  to  $M$ 
5: Get  $x_{i,n}$  from  $A'_{i,n}$  as described in Section 4.2,  $\forall n \in N$ , and for
    $i = 1$  to  $M$ 
6: if  $x'_{i,n} \geq 0.5$ ,  $\forall x'_{i,n} \in a'_{i,n}$ ,  $\forall n \in N$ , and for  $i = 1$  to  $M$  then
7:   Append ( $Q'_N, Q(S'_i, A'_i, S'_{i,n}, A'_{i,n}, \phi')$ )
8: end if
9:  $Q'_i = Q'_N$ , for  $i = 1$  to  $M$ 
10: if Length( $Q'_i$ ) is 0 for any  $i$  then
11:   next $Q_i = Q(S'_i, A'_i, all\_zeros, all\_zeros, \phi')$ 
12: else
13:   next $Q_i = Max(Q'_i)$ 
14: end if
15:  $y = []$ ,  $Qs = []$ 
16: if  $A^{mas}_{i,n} = 1 \exists n \in N \exists i \in M$  then
17:   Append ( $Qs, Q(S_i, A_i, S_{i,n}, A_{i,n}, \phi)$ )
18:   target $Q = \bar{r}_i + \gamma nextQ_i * (1 - done_i)$ 
19:   append( $y, targetQ$ )
20: end if
21: if  $A^{mas}_{i,n} = 0$ ,  $\forall n \in N \exists i \in M$  then
22:   Append ( $Qs, Q(S_i, A_i, all\_zeros, all\_zeros, \phi)$ )
23:   target $Q = \bar{r}_i + \gamma nextQ_i * (1 - done_i)$ 
24:   append ( $y, targetQ$ )
25: end if
26: Compute the TD error:  $\delta = \frac{1}{Length(y)} \sum_{j=1}^{Length(y)} (y_j - Qs_j)^2$ 
27: Update parameters of master agent  $\phi$ :  $\phi \leftarrow \phi + \alpha_\phi \cdot \nabla_\phi \delta$ 
28: Update target master network  $\phi' \leftarrow \phi$ 
29: for each client  $n$  do
30:    $Q_i^N = []$  for  $i = 1$  to  $M$ , tar $Q = []$ 
31:   Set new actions  $A^{new}_{i,n} \leftarrow \pi_n(S_{i,n}, \theta_n)$ ,  $\forall n \in N$ , and for  $i = 1$ 
     to  $M$ 
32:    $A^{new}_i = \{A^{new}_{i,n}\}$ ,  $\forall n \in N$ , and for  $i = 1$  to  $M$ 
33:   Append ( $Q_i^N, Q(S_i, A^{new}_i, S_{i,n}, A^{new}_{i,n}, \phi)$ ),  $\forall n \in N$  with
      $A^{mas}_{i,n} = 1$ , and for  $i = 1$  to  $M$ 
34:   if Length( $Q_i^N$ ) is 0 for any  $i$  then
35:      $Qloc = Q(S_i, A^{new}_i, all\_zeros, all\_zeros, \phi)$ 
36:     Append (tar $Q, Qloc$ )
37:   else
38:     Append (tar $Q, Max(Q_i^N)$ )
39:   end if
40:   Compute the gradient for the client:  $\nabla_{\phi_n} J \phi_n \leftarrow$ 
      $-\frac{1}{Length(tarQ)} \sum_{j=1}^{Length(tarQ)} \nabla_{\phi_n} tarQ_j$ 
41:   Update the client parameters  $\theta_n$ :
42:      $\theta_n \leftarrow \theta_n - \alpha_\theta \nabla_{\theta_n} J(\theta_n)$ 
43:   Update target client networks  $\theta'_n \leftarrow \theta_n$ 
44: end for

```

locally, and it is only used to provide feedback in training the client agents as seen in line 35.

The training of client agents is seen from lines 29 to 44. They are trained similarly to the training of actors in classical MADDPG except that the feedback is computed differently as seen in lines 33 to 39, because the Q value is provided for the client agents that offloaded their task to the server. Therefore, if one or more clients were offloaded their task, the feedback for training the clients is computed from the Q value of one of the offloaded tasks as they are trained with the same rewards. The maximum Q value of the offloaded tasks is considered for consistency. The computation of the maximum Q-value is the same as that of the training for the master agent.

We used prioritized experience replay [19] in the CCM_MADRL algorithms for better efficiency in the training.

5 EXPERIMENTAL EVALUATION

To evaluate the merits of the combinatorial action selection by introducing a master agent to MADDPG, in task offloading problem with various constraints, we compare our algorithm with other benchmarks and heuristic algorithms as follows.

5.1 Benchmark Comparison

The main benchmark for our algorithm is MADDPG for two reasons. First, most existing task-offloading algorithms use DDPG and MADDPG. Second, our algorithm is an extension of MADDPG. However, we also developed different heuristic benchmarks. The heuristics differ from the proposed CCM_MADRL_MEC in that, instead of training a master agent to make combinatorial decisions about the clients, a stationary algorithm is used to make a decision on which of the clients to approve for the MEC server based on some ordering mechanism. The benchmark algorithms are discussed below.

- MADDPG: This benchmark uses actor agents to make decisions. Its difference from CCM_MADRL_MEC is in the procedure 2) of the action selection in Section 4.2, where the SDN allocates the tasks to the sub-channels according to their order of offloading rather than making combinatorial decisions. Tasks that are not assigned to any channel are dropped. The actor agent in the UD will assign τ_{max} to its T_{MEC_n} for the dropped tasks as a penalty. Because the penalty can be unfair for benchmark comparison, the following heuristics are developed to have equivalent combinatorial decisions with the CCM_MADRL_MEC for the tasks that are not accepted.
- MADDPG with the shortest offloading time first heuristic: This is similar to MADDPG, with the distinction that tasks not assigned to sub-channels or storage are designated for local processing.
- MADDPG with deadline/size first heuristic: This differs from the MADDPG with the shortest offloading time first heuristic in that it uses the increasing order of deadline/size as a priority rather than offloading time.

Even if the way tasks are accepted by the MEC server differs between the benchmarks and the CCM_MADRL_MEC, the order of processing of the accepted tasks is always in the order of arrival at the MEC using Equation (4)

5.2 Experimental Settings

The experimental setting is provided in Table 2. Note that all UDs have the same minimum battery, power, and resource allocation threshold but their maximum budget is generated from uniform distribution. As described in Section 2, the experimental setting is customized from the settings in [15] and in [23]. The typical storage capacity of modern servers is GBs and TBs. However, because we chose a small experimental setting due to computational resources, we considered a storage constraint of 400 MB so that the task of-flooding problem is combinatorial to the server. We used a seed of 37 for the reproducibility of the simulation environment. Evaluation episodes are seeded with their index. A discount factor of 0.99 is used in this experiment.

Table 2: Experimental parameters of CCM_MADRL_MEC

Param	Value	Param	Value
$ N $	50	P^{max}	24 dBm
K	10	P^{min}	1 dBm
τ_n	[0.1-0.9] s	P_n^{max}	$[P^{min} - P^{max}]$ dBm
W	40 MHz	P_n^{min}	P^{min} dBm
f_{max}	1.5 GHz	f_{min}	0.4 GHz
f_n^{max}	$[f_{min} - f_{max}]$ GHz	b_n^{max}	3.2 MJ
f_n^{min}	f_{min} GHz	b_n^{min}	0.5 MJ
g_n	[5-14] dB	f_e	4 GHz
z_n	[1-50] MB	z_e	400 MB
c_n	[300 - 737.5] cycles	U_e	8
κ	5×10^{-27}	λ_1, λ_2	0.5, 0.5
e_n	0.001 J	b_n^{max}	$[b^{min} - b^{max}]$ MJ
b_n^{min}	b^{min} MJ		

5.3 Generalizability

Because convergence is affected by the initialization of the weights of the DNNs and the exploration and exploitation sequence, we evaluate the algorithms in a different evaluation environment. At every episode of the training environment, the DRL is evaluated with 50 evaluation episodes.

5.4 Results

First, we run an experiment with 10 steps per episode for 2000 training episodes. We ran 10 experiments using different initialization of DNN weights and different sequences of exploration and exploitation for each run and plotted the result with a 95% confidence interval as seen in Figure 3 (A). The CCM_MADRL algorithm has performed better than the others because once the clients choose their action, the master agent also makes a combinatorial decision on the action of the clients. On the other hand, MADDPG-based heuristic and benchmark algorithms use only actors to provide actions. The critic is only used to provide feedback to the actors. Similarly, Figure 3 (B) shows CCM_MADRL has a smaller percentage of tasks whose deadline has expired before finishing processing of the tasks compared to the other algorithms. Note that Figure 3 (A) is the performance based on the combined reward and combined penalty for the time and energy consumption in Equation (13) but Figure 3

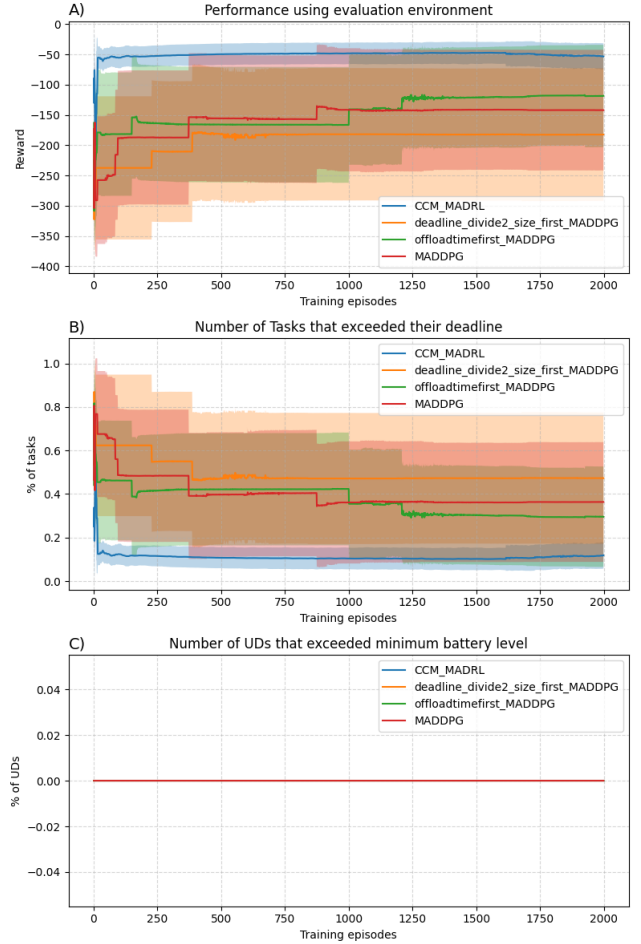


Figure 3: Comparison of the CCM_MADRL in the evaluation environment with the heuristic and MADDPG algorithms with a learning rate of 0.0001 and 0.001 for the clients and master respectively

(B) is a percentage of tasks whose deadline expired before finishing to the total number of tasks generated in the episode. Figure 3 (C) shows that none of the UDs exceed the minimum battery threshold in an episode. This is because the experiments are run for 10 steps per episode.

Next, to see the impact on the battery level, we repeat the above experiment by changing the number of steps per episode to 100 and the b_{max} to b_{min} 1J. Figure 4 (C) shows that CCM_MADRL_MEC has more UDs running below the minimum battery threshold than the benchmark and heuristic algorithms. However, as explained above for Figure 3, the algorithm is by computing a scalar reward, which is a sum of energy and time consumption, as seen in Equation (13) but the percentage of UDs exceeding their battery threshold is from the ration of UDs that exceed battery threshold to the total number of UDs. Therefore, it is affected by the scales given to the deadline penalty and energy penalty. We used $\lambda_1 = \lambda_2 = 0.5$ as weight coefficients of processing time and energy consumption. We

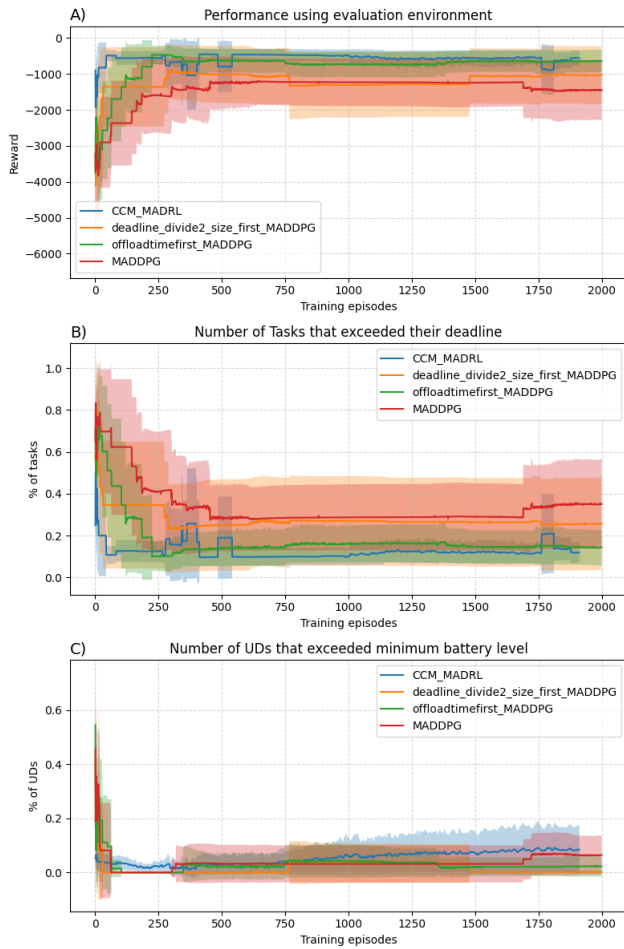


Figure 4: Comparison of the CCM_MADRL with the heuristic and MADDPG algorithms on the evaluation environment with learning rates of 0.0001 and 0.001 and $b_{max} = b_{min} 1J$ for the clients and master respectively

repeated the algorithm for $lambda_1 = 1$ and $lambda_2 = 5$ as seen in Figure 5.

We ran the experiment further with $lambda_1 = 1$ and $lambda_2 = 5$ because Figure 4 (B) looks like it is the exact inverse of Figure 4 (B) means $lambda_1 = lambda_2 = 0.5$ has not given enough balance for the energy consumption, which is a very low number, and the time consumption, which is relatively larger. The result is plotted in Figure 5 with a 95% confidence interval of 40 runs, unlike the above two experiments which are plotted with a 95% confidence interval of 10 runs. However, the plots do not show a significant difference with 10 runs and 40 runs. Subplot (B) and subplot (A) are now varied but subplot (C) shows that CCM_MADRL_MEC has still more UDs running below the minimum battery threshold than the benchmark and heuristic algorithms. Nonetheless, CCM_MADRL_MEC has performed better than the other algorithms in the combined performance and shows more advantage in the number of tasks whose

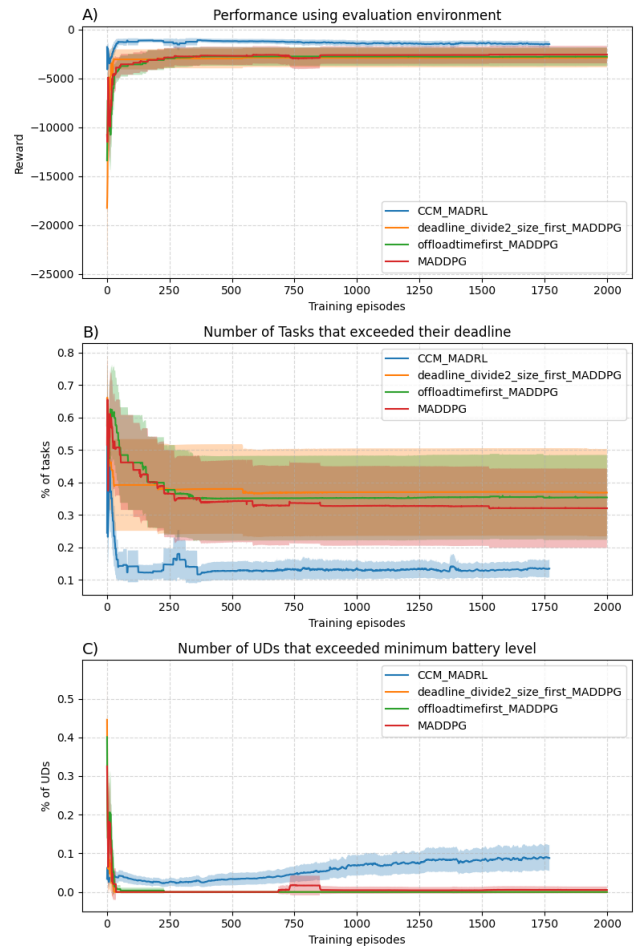


Figure 5: Comparison of the CCM_MADRL with the heuristic and MADDPG algorithms on the evaluation environment with learning rates of 0.0001 and 0.001 and $b_{max} = b_{min} 1J$ for the clients and master respectively

deadlines exceeded. Using scalar reward for multi-objective functions does not represent the underlying problem but it is not the scope of this work. The work in [6] has studied the essence and techniques of multi-objective reinforcement learning.

It can be seen that the benchmark and heuristic algorithms demonstrated a closer performance to the CCM_MADRL_MEC on 100 steps per episode than on 10 steps per episode. This is because the training is performed at the end of each episode. This caused the training after 10 steps to overfit to the early episodes, whereas the training after every 100 episodes to generalize. The benchmark and heuristic algorithms are impacted by overfitting more than CCM_MADRL_MEC because they use only their actors to select action while the CCM_MADRL_MEC uses the advantage of both clients and master to mitigate overfitting and sticking to local optimal.

Note that the CCM_MADRL_MEC in Figure 4 and Figure 5 are not plotted until the last episode. The experiment is run on

Iridis², an HPC cluster at the University of Southampton. The experiments for 60 hours each. All of the experiments for heuristic and benchmark algorithms were finished earlier, but some of the runs for the CCM_MADRL_MEC ran out of time before reaching the last episode. For convenience in plotting with the 95% confidence interval, all runs of the CCM_MADRL_MEC are clipped after the run with the least number of episodes. Note that the benchmark and heuristic algorithms have only one Q value in the critic for a combination of state and actions of the actors. On the other hand, the number of Q values to train in the CCM_MADRL_MEC is equal to the number of offloaded tasks or 1 if all of them are allocated locally.

6 CONCLUSION

In this paper, we propose a combinatorial client-master MADRL algorithm for task offloading in MEC, that considers various constraints at the UD, sub-channels, and server. By combining the advantages of both policy gradient and value functions to output continuous and combinatorial actions, CCM_MADRL provides better convergence than existing homogeneous MADRL algorithms because the master agent applies combinatorial decisions on the actions proposed by the clients.

In the future, we plan to extend CCM_MADRL_MEC to multi-server MEC where multiple servers cooperate to make combinatorial decisions.

ACKNOWLEDGMENTS

This research was sponsored by the U.S. Army Research Laboratory and the U.K. Ministry of Defence under Agreement Number W911NF-16-3-0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Research Laboratory, the U.S. Government, the U.K. Ministry of Defence or the U.K. Government. The U.S. and U.K. Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon. This work was also supported by an EPSRC Turing AI Acceleration Fellowship (EP/V022067/1).

REFERENCES

- [1] Seble Birhanu Engidayehu, Tahira Mahboob, and Min Young Chung. 2022. Deep Reinforcement Learning-based Task Offloading and Resource Allocation in MEC-enabled Wireless Networks. In *2022 27th Asia Pacific Conference on Communications (APCC)*. 226–230.
- [2] Guang Chen, Yueyun Chen, Zhiyuan Mai, Conghui Hao, Meijie Yang, and Liping Du. 2023. Incentive-Based Distributed Resource Allocation for Task Offloading and Collaborative Computing in MEC-Enabled Networks. *IEEE Internet of Things Journal* 10, 10 (2023), 9077–9091.
- [3] Gabriel Dulac-Arnold, Richard Evans, Hado van Hasselt, Peter Sunehag, Timothy Lillicrap, Jonathan Hunt, Timothy Mann, Theophane Weber, Thomas Degris, and Ben Coppin. 2015. Deep reinforcement learning in large discrete action spaces. *arXiv preprint arXiv:1512.07679* (2015).
- [4] Tesfay Zemuy Gebrekidan, Sebastian Stein, and Timothy J. Norman. 2024. Combinatorial Client-Master Multiagent Deep Reinforcement Learning for Task Offloading in Mobile Edge Computing: Extended Abstract. In press. Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024).
- [5] Tesfay Zemuy Gebrekidan, Sebastian Stein, and Timothy J. Norman. 2024. Deep Reinforcement Learning with Coalition Action Selection for Online Combinatorial Resource Allocation with Arbitrary Action Space. In press. Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024).
- [6] Conor F Hayes, Roxana Rădulescu, Eugenio Bargiacchi, Johan Källström, Matthew Macfarlane, Mathieu Reymond, Timothy Verstraeten, Luisa M Zintgraf, Richard Dazeley, Fredrik Heintz, et al. 2022. A practical guide to multi-objective reinforcement learning and planning. *Autonomous Agents and Multi-Agent Systems* 36, 1 (2022), 26.
- [7] Ji He, Jianshu Chen, Xiaodong He, Jianfeng Gao, Lihong Li, Li Deng, and Mari Ostendorf. 2015. Deep reinforcement learning with a natural language action space. *arXiv preprint arXiv:1511.04636* (2015).
- [8] Akhirl Islam, Arindam Debnath, Manojit Ghose, and Suchetana Chakraborty. 2021. A survey on task offloading in multi-access edge computing. *Journal of Systems Architecture* 118 (2021), 102225.
- [9] Wei Jiang, Daquan Feng, Yao Sun, Gang Feng, Zhenzhong Wang, and Xiang-Gen Xia. 2023. Joint Computation Offloading and Resource Allocation for D2D-Assisted Mobile Edge Computing. *IEEE Transactions on Services Computing* 16, 3 (2023), 1949–1963.
- [10] Te-Yi Kan, Yao Chiang, and Hung-Yu Wei. 2018. Task offloading and resource allocation in mobile-edge computing system. In *2018 27th Wireless and Optical Communication Conference (WOCC)*. 1–4.
- [11] Xiaowei Liu, Shuwen Jiang, and Yi Wu. 2022. A novel deep reinforcement learning approach for task offloading in MEC systems. *Applied Sciences* 12, 21 (2022), 11260.
- [12] Xiaowei Liu, Shuwen Jiang, and Yi Wu. 2022. A Novel Deep Reinforcement Learning Approach for Task Offloading in MEC Systems. *Applied Sciences* 12, 21 (2022).
- [13] Ryan Lowe, Yi Wu, Aviv Tamar, Jean Harb, Pieter Abbeel, and Igor Mordatch. 2017. Multi-agent actor-critic for mixed cooperative-competitive environments. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*. 6382–6393.
- [14] Michael Pendo John Mahenge, Chunlin Li, and Camilius A Sanga. 2022. Energy-efficient task offloading strategy in mobile edge computing for resource-intensive mobile applications. *Digital Communications and Networks* 8, 6 (2022), 1048–1058.
- [15] D. C. Nguyen, M. Ding, P. N. Pathirana, A. Seneviratne, J. Li, and H. Poor. 2023. Cooperative Task Offloading and Block Mining in Blockchain-Based Edge Computing With Multi-Agent Deep Reinforcement Learning. *IEEE Transactions on Mobile Computing* 22, 04 (2023), 2021–2037.
- [16] Lili Nie, Huiqiang Wang, Guangsheng Feng, Jiayu Sun, Hongwu Lv, and Hang Cui. 2023. A deep reinforcement learning assisted task offloading and resource allocation approach towards self-driving object detection. *Journal of Cloud Computing* 12, 1 (2023), 131.
- [17] Kuanishbay Sadatdiyrov, Laizhong Cui, Lei Zhang, Joshua Zhexue Huang, Salman Salloum, and Mohammad Sultan Mahmud. 2023. A review of optimization methods for computation offloading in edge computing networks. *Digital Communications and Networks* 9, 2 (2023), 450–461.
- [18] Dileep Kumar Sajjani, Abdul Rasheed Mahesar, Abdullah Lakhani, and Irfan Ali Jamali. 2018. Latency aware and service delay with task scheduling in mobile edge computing. *Communications and Network* 10, 04 (2018), 127–141.
- [19] Tom Schaul, John Quan, Ioannis Antonoglou, and David Silver. 2016. Prioritized experience replay. (2016).
- [20] Suzhen Wang, Zhongbo Hu, Yongchen Deng, and Lisha Hu. 2022. Game-Theory-Based Task Offloading and Resource Scheduling in Cloud-Edge Collaborative Systems. *Applied Sciences* 12, 12 (2022).
- [21] Jianyu Xiong, Peng Guo, Yi Wang, Xiangyin Meng, Jian Zhang, Linmao Qian, and Zhenglin Yu. 2023. Multi-agent deep reinforcement learning for task offloading in group distributed manufacturing systems. *Engineering Applications of Artificial Intelligence* 118 (2023), 105710.
- [22] Jia Yan, Suzhi Bi, and Ying-Jun Angela Zhang. 2018. Optimal Offloading and Resource Allocation in Mobile-Edge Computing with Inter-User Task Dependency. In *2018 IEEE Global Communications Conference (GLOBECOM)*. 1–8.
- [23] Jing Zhang, Jun Du, Yuan Shen, and Jian Wang. 2020. Dynamic Computation Offloading With Energy Harvesting Devices: A Hybrid-Decision-Based Deep Reinforcement Learning Approach. *IEEE Internet of Things Journal* 7, 10 (2020), 9303–9317.
- [24] Weiwen Zhang, Yonggang Wen, Jun Wu, and Hui Li. 2013. Toward a unified elastic computing platform for smartphones with cloud support. *IEEE Network* 27, 5 (2013), 34–40.

²<https://www.southampton.ac.uk/isolutions/staff/iridis.page>