

Multi-Agent Reinforcement Learning for Cooperative Coded Caching via Homotopy Optimization

Xiongwei Wu, *Student Member, IEEE*, Jun Li, *Senior Member, IEEE*,
Ming Xiao, *Senior Member, IEEE*, P. C. Ching, *Life Fellow, IEEE*,
H. Vincent Poor, *Life Fellow, IEEE*

Abstract—Introducing cooperative coded caching into small cell networks is a promising approach to reducing traffic loads. By encoding content via maximum distance separable (MDS) codes, coded fragments can be collectively cached at small-cell base stations (SBSs) to enhance caching efficiency. However, content popularity is usually time-varying and unknown in practice. As a result, caching contents are anticipated to be intelligently updated by taking into account limited caching storage and interactive impacts among SBSs. In response to these challenges, we propose a multi-agent deep reinforcement learning (DRL) framework to intelligently update caching contents in dynamic environments. With the goal of minimizing long-term expected fronthaul traffic loads, we first model dynamic coded caching as a cooperative multi-agent Markov decision process. Owing to MDS coding, the resulting decision-making falls into a class of constrained reinforcement learning problems with continuous decision variables. To deal with this difficulty, we custom-build a novel DRL algorithm by embedding homotopy optimization into a deep deterministic policy gradient formalism. Next, to empower the caching framework with an effective trade-off between complexity and performance, we propose centralized, partially and fully decentralized caching controls by applying the derived DRL approach. Simulation results demonstrate the superior performance of the proposed multi-agent framework.

Index Terms—Small cell networks, MDS codes, homotopy optimization, deep multi-agent reinforcement learning

I. INTRODUCTION

In recent years, there has been a significant surge in mobile data traffic, which is anticipated to impose a heavy traffic burden on wireless networks for the foreseeable future [1], [2]. As a consequence, wireless networks could become very

This work was supported in part by the U.S. National Science Foundation under Grant CCF-1908308, in part by EU Marie Skłodowska-Curie Actions Project entitled “High-reliability Low-latency Communications with Network Coding”, in part by Swedish Foundation for International Cooperation in Research and Higher Education (STINT), project “Efficient and Secure Distributed Machine Learning with Gradient Descend”, and in part by National Natural Science Foundation of China under Grant 61872184. (*Corresponding author: Jun Li and Xiongwei Wu.*)

X. Wu and P. C. Ching are with the Department of Electronic Engineering, The Chinese University of Hong Kong, Shatin, Hong Kong SAR of China (e-mail: xwwu@ee.cuhk.edu.hk; pching@ee.cuhk.edu.hk).

J. Li is with the School of Electronic and Optical Engineering, Nanjing University of Science and Technology, Nanjing 210094, China. (e-mail: jun.li@njust.edu.cn).

M. Xiao is with the School of Electrical Engineering and Computer Science, Royal Institute of Technology (KTH), Stockholm, Sweden (email: mingx@kth.se).

H. V. Poor is with the Department of Electrical Engineering, Princeton University, Princeton, NJ 08544. (e-mail: poor@princeton.edu).

congested, and thus find it difficult to satisfy user requests with satisfactory quality of service. To cope with this challenge, edge caching has been proposed as a promising solution towards fifth generation (5G) communications and beyond [2]. By endowing caching units in wireless edge nodes, e.g., small-cell base stations (SBSs), popular content can be pre-fetched close to users. Subsequently, the caching content is able to be delivered to users without duplicated transmissions in fronthaul and backhaul links [2]. This process significantly decreases traffic loads, alleviates network congestion, reduces delay, and thus improves system performance [3]–[5].

In general, caching policies should be designed according to system features, e.g., users arrivals and content popularity, to better satisfy user demands. These features, in practice, usually exhibit unknown and temporal dynamics. For instance, content popularity is generally time-varying because the most popular content at the current epoch may not receive the highest attention in the future; and mobile users could change locations as time passes [6]. Thus, with a limited caching storage, it is crucial to learn how to reasonably update caching contents given the real-time observations of system features. Fortunately, by embedding deep learning into reinforcement learning (RL), deep RL (DRL) has emerged as an effective tool to address decision-making in dynamic environments [7]. This artificial intelligence technique can be leveraged to learn an optimal policy to maximize long-term performance criteria through interactions with environments [7], [8]. In this way, utilizing DRL is envisioned to empower “intelligent” caching, i.e., updating caching resources by tracking and adapting to dynamic features of wireless networks [9].

A. Related Work

Prior studies generally investigated the potentials of edge caching by optimizing average performance criteria. For instance, the studies in [10]–[12] examined effective caching strategies to alleviate traffic loads, and to reduce system cost and download latency. With the aid of caching resources at SBSs, the studies in [13], [14] investigated the joint design of SBS beamforming and clustering. Caching strategies in these studies were designed to store either entire content items or uncoded fragments, which are referred to as *uncoded caching*.

To further improve caching efficiency, *coded caching* has recently gained considerable research attention. The study in

[15] proposed a novel coded caching scheme, which provides a global caching gain relating to cumulative storage over all caching units. The research in [16], [17] investigated cooperative coded caching by utilizing maximum separable distance (MDS) codes to reduce traffic loads. MDS coded caching was also examined in [18], [19] to augment SBS collaboration and thus offers great advantages to lower latency and reduce power consumption compared with *uncoded caching*. The above-mentioned studies mainly investigated offline caching policies by assuming time-invariant content popularity distributions.

To exploit dynamic features in wireless networks, extant works have been devoted to designing caching policy by using RL. The study in [20] utilized Q-learning to find an optimal caching policy to minimize network cost. To counter the curse of dimensionality in conventional RL, DRL-based caching policies were advocated in [9], [21]–[23] by using deep neural networks as function approximators. Moreover, the study in [24] proposed a multi-agent DRL framework to maximize cache hit ratios for centralized and decentralized settings. The authors in [25] investigated cache placement by using cooperative multi-agent multi-armed bandit learning at small cell networks (SCN). A decentralized caching scheme was proposed in [26] by utilizing federated deep reinforcement learning. Nevertheless, the research in [9], [20]–[26] focused on *uncoded caching*. That is, each content item is entirely cached without exploring SBSs cooperatively fetching coded fragments of each content item.

B. Contributions

Indeed, cooperatively pre-fetching MDS coded fragments has been proven to significantly reduce traffic loads, download delay, and transmission cost over caching uncoded fragments at SCN [17]–[19]. By distributing MDS coded fragments of a content item to multiple SBSs, mobile users are allowed to access a cluster of SBSs simultaneously to download coded fragments of the desired content item. To date, very few works have investigated how to “intelligently” update MDS coded content items under dynamic environments (e.g., time-varying content popularity). Unlike uncoded caching schemes [7]–[9], [20]–[25] addressing binary decision-making, coded caching essentially entails continuous caching decisions that are subject to storage constraints. Therefore, the resulting decision-making for cooperative coded caching is a constrained RL with a continuous action space. The prior study in [27] adopted Q-learning to tackle this issue by quantizing continuous actions into discrete values. This method has been shown to probably impair performance compared with the continuous RL approaches such as deep deterministic policy gradient (DDPG) [27]. Another MDS caching study in [28] adopted the conventional DDPG for decision making, which may not be able to efficiently handle the storage constraints and result in suboptimal decisions. It is worth noting that, the centralized control in [27], [28] could lead to excessive communication overhead because the cloud processor (CP) needs frequent communications with SBSs to aggregate information and inform SBSs of their caching decisions. As the number of SBSs increases, the dimension of continuous states

and actions in a centralized control would be very large, and thus obtaining an optimal caching policy is computationally prohibitive. As a consequence, it is very challenging, but essential, to design efficient DRL algorithms for cooperative coded caching. To bridge the research gap identified above, in this work, we address the following fundamental issues of cooperative coded caching: **i) how to design efficient DRL algorithms for a constrained RL problem with continuous decision variables; and ii) how to develop a multi-agent DRL-based framework with different levels of controls to obtain an equitable trade-off between performance and complexity.**

The design recipe of the proposed DRL entails homotopy optimization, which refers to a class of optimization methods for non-convex problems by use of homotopy [29]. The general idea is to develop a continuous transformation of the original problem and gradually trace the path from the solution of the transformed problem to the solution of the original one. The transformed problem could be an easy one, which inaccurately approximates the original problem; or a hard one, e.g., by adding some regularizer which penalizes the solution close to the optimum. In this work, we take the latter fashion and first introduce a cumulative penalty regularizer to the corresponding RL problem of cooperative coded caching, which can penalize poor caching decisions that violate the optimal characterization. Such a transformed problem is found to admit a structure that still can be addressed by DDPG. Then, we propose a novel DRL by embedding homotopy optimization into the formalism of DDPG, which favorably leads to better performance. We also devise its decentralized adaptations to reduce signaling overhead.

Finally, the main contributions of this work are summarized as follows:

- To the best of our knowledge, this is the first work to investigate a multi-agent DRL framework for MDS coded caching under time-varying content popularity. Specifically, we model cache updating for MDS coded caching as a cooperative multi-agent Markov decision process (MDP). With the goal of minimizing long-term expected cumulative fronthaul traffic loads, we judiciously define the system state, local observations and action space of each agent, as well as caching reward. Our formulated problem is a continuous RL with action constraints. We also characterize optimal decisions in a closed form.
- As a core contribution, we reformulate a general constrained RL problem, whose action space is inefficient to be satisfied through designing neural networks, into a tractable form that can be dealt with by utilizing homotopy optimization. Then, we custom-build a novel DRL, i.e., homotopy deep deterministic policy gradient (HDDPG), through recasting the basic elements of RL and unfolding the iterative process of homotopy optimization. The novelty of this approach lies in introducing a reasonable cumulative penalty to the objective of RL, and then properly manipulating it by using homotopy optimization.
- To endow the proposed DRL caching framework with different levels of control, we generalize the proposed HDDPG from centralized control to partially and fully

decentralized controls. Specifically, in the centralized control, the CP coordinates SBSs to conduct cache updating by using global information. To reduce complexity and communication overhead, we then propose a partially decentralized control by allowing SBSs to make decisions locally. Whilst, the decentralized policies are learned in a centralized manner. In the fully decentralized control, each SBS works as an independent learner and trains its caching policy independently based on local observations. The proposed decentralized controls could obtain a desirable trade-off between complexity and performance, and thus have the potential to handle large-scale wireless networks.

The remainder of this paper is organized as follows. Sec. II presents the problem statement. Sec. III introduces the proposed DRL. Sec. IV develops a centralized cooperative coded caching design. Sec. V proposes a partially decentralized caching design, and Sec. VI proposes a fully decentralized caching design. Sec. VII presents performance evaluations, and Sec. VIII concludes the paper.

II. PROBLEM STATEMENT

A. MDS Coded Caching at SCN

As illustrated in Fig. 1, we consider a SCN, in which a total of B SBSs are densely deployed and thus are capable of cooperatively providing communication services for users. Each SBS is endowed with a cache unit, which can cache popular content from the CP through a capacity-limited fronthaul. The CP is further connected to the core network through a backhaul. Suppose that a catalog of F content items are available at the CP. For ease of discussion, all of the content items are of the same size s bits, and each cache unit has a storage of $L \times s$ bits. Let $\mathcal{B} = \{1, \dots, B\}$ and $\mathcal{F} = \{1, \dots, F\}$ denote the indices of SBSs and content items, respectively.

To reduce traffic loads on the capacity-limited fronthaul and provide better services for mobile users, SBSs can proactively cache popular content. By applying MDS codes, each content item of size s bits is able to be encoded into a sufficiently long sequence of parity bits, and any s parity bits are sufficient to reconstruct the original content item [17], [18]. Therefore, SBSs with limited caching storage can cooperatively cache these coded parity bits to satisfy user requests locally as much as possible. More precisely, we define the cache allocation matrix as $\mathbf{L} = [l_{f,b}] \in \mathbb{R}^{F \times B}$, where element $l_{f,b} \in [0, 1]$ denotes the proportion of parity bits encoding content item f that are stored at SBS b , $\forall b \in \mathcal{B}, f \in \mathcal{F}$. Owing to the storage limit at SBSs, cache allocation needs to satisfy $\sum_{f \in \mathcal{F}} l_{f,b} \leq L, \forall b$. It is worth noting that, the parity bits of a content item available at SBS b should be independent of that cached at other SBSs; thus, users can always download distinct coded content items from multiple SBSs [17]. To guarantee this caching diversity among SBSs, the encoded information sequence of parity bits of every content item should be sufficiently long, e.g., larger than Bs .

In what follows, we introduce how SBSs cooperatively transmit coded content items to mobile users. The operation time of a SCN is divided into a series of epochs, indexed

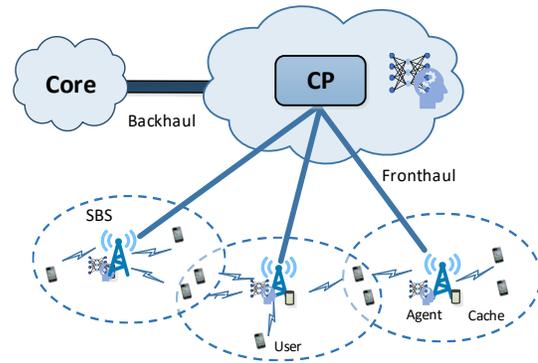


Fig. 1: A downlink cache-enabled SCN.

by $t = 0, 1, \dots$. For each epoch t , a number of active users \mathcal{K}^t are randomly distributed in the horizontal plane. We also assume that the duration of each epoch is relatively short, such that active users are considered to be quasi-static during a single epoch. Furthermore, each user $k \in \mathcal{K}^t$ is able to be served by the neighboring SBSs, which are specified by the communication radius [30]. Specifically, we set an indicator variable $e_{k,b}^t = 1$ to indicate that $d_{k,b}^t \leq r_0$, where r_0 denotes the communication radius of each SBS, and $d_{k,b}^t$ denotes the distance between user k and SBS b at epoch t ; otherwise, $e_{k,b}^t = 0$. Accordingly, the neighboring SBSs to serve user k are given by $\{b \in \mathcal{B} | e_{k,b}^t = 1\}$. For notational convenience, we collect all active users served by SBS b at epoch t as set $\mathcal{K}_b^t = \{k \in \mathcal{K}^t | e_{k,b}^t = 1\}$, for $\forall b \in \mathcal{B}$. The network connectivity at epoch t is denoted by $\mathcal{E}^t = \{e_{k,b}^t\}_{k \in \mathcal{K}^t, b \in \mathcal{B}}$. In addition, every user is assumed to request one content item at a single epoch. Evidently, if content item f is not fully stored at the neighboring SBSs associated with user k at epoch t , i.e.,

$$\sum_{b \in \mathcal{B}} l_{f,b}^t e_{k,b}^t < 1, \quad (1)$$

the missing part $(1 - \sum_{b \in \mathcal{B}} l_{f,b}^t e_{k,b}^t)s$ needs to be transmitted by the CP via fronthaul. This event is referred to as *cache miss*, which introduces additional fronthaul traffic loads. We summarize all of the key notations¹ in Table I.

TABLE I: Summary of main notations

| Notation | Description |
|----------------------|--|
| b, \mathcal{B} | SBS index, indices of all SBSs |
| k, \mathcal{K}_b^t | User index, active users observed by SBS b at epoch t |
| \mathcal{K}^t | All of the active users at epoch t |
| f | Index of a content item |
| f_k^t | Index of the content item requested by user k at epoch t |
| $l_{f,b}^t$ | Fraction of content item f stored at SBS b at epoch t |
| L | Normalized caching storage at each SBS |
| $e_{k,b}^t$ | Whether or not user k can access SBS b at epoch t |
| \mathcal{E}^t | Network connectivity at epoch t |

As a result, to mitigate traffic burden on fronthaul, optimized caching policies depend highly on knowledge of network connectivity, i.e., \mathcal{E}^t , and content popularity, i.e., $\{p_f, \forall f\}$,

¹Without further definition, in this paper, notation $(\cdot)^t$ denotes the value of (\cdot) taken at epoch t .

where p_f denotes the probability of content item f being requested. These elements generally exhibit unknown and time-varying dynamics in practice. Indeed, caching content needs to be temporally updated based on historical observations in order to provide better download services for future requests. We thereby introduce a dynamic cooperative coded caching problem in the following subsection.

B. Cooperative Multi-Agent MDP

In the cooperative coded caching, SBSs are anticipated to collaboratively cache the coded content items, which can be specified by optimizing continuous variables $\{l_{f,b}^t\}$. Therefore, we formulate the considered cooperative coded caching problem as a cooperative multi-agent MDP.

Definition 1 A cooperative multi-agent MDP is specified by a tuple $(\mathcal{B}, \mathcal{S}, \mathcal{A}, \mathcal{P}, R, \gamma)$, where \mathcal{B} denotes the set of agents; \mathcal{S} denotes the state space, which aggregates all of the agent local observations $\mathbf{S} = \cup_{b \in \mathcal{B}} \{\mathbf{S}_b\}$; and \mathcal{A} denotes the action space for a joint action $\mathbf{A} = \cup_{b \in \mathcal{B}} \{\mathbf{A}_b\}$. Let \mathcal{A}_b be the action space of local action \mathbf{A}_b , and then $\mathcal{A} = \cup_{b \in \mathcal{B}} \{\mathcal{A}_b\}$. \mathcal{P} collects all of the transition probability $Pr\{\mathbf{S}'|\mathbf{S}, \mathbf{A}\}$ for $\forall \mathbf{S}, \mathbf{S}' \in \mathcal{S}, \mathbf{A} \in \mathcal{A}$. All of the agents share a common reward R after they cooperatively take actions $\{\mathbf{A}_b\}_{b \in \mathcal{B}}$. $\gamma \in [0, 1)$ denotes a discount factor.

As aforementioned, user requests are expected to be satisfied by SBSs locally as much as possible; otherwise, the missing fragments could introduce additional traffic burden and transmission delay on the fronthaul. Therefore, in this paper, our goal is to minimize the expected fronthaul traffic loads. Accordingly, the basic elements in a cooperative multi-agent MDP are defined as follows.

State: We assume that a user request can be observed by his or her neighboring SBSs only. Consequently, SBS b has local observation of the environment, which is defined as follows:

$$\mathbf{S}_b^t = \left[\{f_k^t\}_{k \in \mathcal{K}_b^t}, \{\mathcal{E}_k^t\}_{k \in \mathcal{K}_b^t}, \{l_{f,b}^t\}_{f \in \mathcal{F}} \right], \quad (2)$$

where $f_k^t \in \mathcal{F}$ denotes the index of the content item requested by user k at epoch t ; and $\mathcal{E}_k^t = \{e_{k,b'}^t\}_{b' \in \mathcal{B}}$ implies the strategy of SBS collaboration in order to satisfy user k 's request, which can be inferred through acquiring user location. By aggregating observations of all SBSs, the system state is defined as:

$$\mathbf{S}^t = \left[\{f_k^t\}_{k \in \mathcal{K}^t}, \mathcal{E}^t, \mathbf{L}^t \right]. \quad (3)$$

Action: By the end of each epoch t , all SBSs need to update their cached content. Accordingly, we define the action of SBS b at the current epoch as $\mathbf{A}_b^t = [a_{f,b}^t, \forall f \in \mathcal{F}]$, where element $a_{f,b}^t = l_{f,b}^{t+1}$; and the corresponding action space is given by:

$$\mathcal{A}_b = \left\{ \mathbf{A}_b^t \mid 0 \leq a_{f,b}^t \leq 1, \forall f \in \mathcal{F}, \sum_{f \in \mathcal{F}} a_{f,b}^t \leq L \right\}. \quad (4)$$

As such, a joint action can be given by $\mathbf{A}^t = [a_{f,b}^t, \forall f \in \mathcal{F}, b \in \mathcal{B}]$.

Reward: After executing joint action \mathbf{A} , the system state turns into \mathbf{S}^{t+1} with transition probability $Pr\{\mathbf{S}^{t+1}|\mathbf{S}^t, \mathbf{A}^t\}$. In this cooperative task, all of the agents shall receive a common reward $R(\mathbf{S}^{t+1}, \mathbf{S}^t, \mathbf{A}^t)$, which indicates how good a joint action \mathbf{A}^t is. Therefore, it should be consistent with the goal of reducing fronthaul traffic loads. The total traffic loads for updating caching resources and satisfying user requests in the coming epoch are given by:

$$C^{t+1} = \sum_{f \in \mathcal{F}, b \in \mathcal{B}} \max \left\{ l_{f,b}^{t+1} - l_{f,b}^t, 0 \right\} + \sum_{k \in \mathcal{K}^{t+1}} \max \left\{ 0, 1 - \sum_{b \in \mathcal{B}} l_{f,b}^{t+1} e_{k,b}^{t+1} \Big|_{f=f_k^{t+1}} \right\}. \quad (5)$$

Accordingly, we design the reward as $R^{t+1} = R(\mathbf{S}^{t+1}, \mathbf{S}^t, \mathbf{A}^t) \triangleq -C^{t+1}/|\mathcal{K}^{t+1}|$, which indicates the normalized traffic loads to satisfy each content request at the coming epoch.

Toward this end, the goal of this study is to find a cooperative caching policy π^* , which maximizes the total expected cumulative caching reward, i.e.:

$$\pi^* = \arg \max_{\pi \in \Pi} \mathbb{E} [V|\pi], \quad (6)$$

and the cumulative reward is defined as:

$$V \triangleq \sum_{t=0}^{\infty} (\gamma)^t R^{t+1}, \quad (7)$$

where π denotes a mapping from state space to action space; Π denotes the set of feasible caching policies; and the expectation is over all of the rewards $\{R(\mathbf{S}^{t+1}, \mathbf{S}^t, \mathbf{A}^t)\}$. Furthermore, a characterization for optimal decisions is presented in the following proposition.

Proposition 1 Consider that all SBSs are fully loaded at the initial epoch, i.e., $\sum_{f \in \mathcal{F}} l_{f,b}^0 = L, \forall b \in \mathcal{B}$. There exists an optimal decision sequence $\{(\mathbf{A}^t)^*\}$ satisfying $\sum_{f \in \mathcal{F}} (a_{f,b}^t)^* = L, \forall b \in \mathcal{B}$, where $(\mathbf{A}^t)^* = [(a_{f,b}^t)^*]$, for any $t \geq 1$.

Proof. See Appendix A. ■

Remark 1 Proposition 1 implies that optimal caching decisions are likely to be the case where caching units are fully loaded. This result is reasonable and would provide further insight for algorithm design. Nevertheless, calculating an optimal cooperative caching policy offline depends on knowledge of network dynamics (e.g., transition probability $Pr\{\mathbf{S}'|\mathbf{S}, \mathbf{A}\}$), which is generally difficult to obtain in real applications. Even if this knowledge could be obtained, problem (6) is still intractable due to its no closed-form expression. In view of this, one can resort to DRL to handle this problem through utilizing historical experiences without knowing exact dynamic information. On the other hand, as previously mentioned, the decision variables $\{a_{f,b}^t\}$ for cooperative coded caching are continuous. Simply quantizing these variables into discrete values may lead to performance loss as well as an exponentially large number of actions.

III. A NOVEL HOMOTOPY DDPG

To develop an effective DRL algorithm for the considered problem, we first introduce a policy-based RL algorithm and identify the arising challenges. Then, we recast a general policy based RL problem with constraints into a tractable form, which is suitable to be addressed by leveraging homotopy optimization. Finally, we custom-build a novel DRL algorithm by embedding homotopy optimization into DDPG.

A. Fundamentals of DDPG

DDPG is one of the policy-based RL algorithms, which is widely used to handle continuous decision-making [31]. Built upon actor-critic architectures, this algorithm employs deep neural networks as function approximators to learn a deterministic policy that can map high-dimensional states into feasible continuous actions. Typically, a DDPG based RL framework consists of two networks, i.e., critic and actor, which are detailed as follows.

Actor: The actor network corresponds to a deterministic policy, which can generate an action \mathbf{A} under a given system state \mathbf{S} , i.e., $\mathbf{A} = \pi_{\theta}(\mathbf{S})$, and θ is the parameter of the associated deep neural networks. This parametrized policy $\pi_{\theta}(\cdot)$ aims to maximize the expected cumulative reward, i.e.:

$$J(\theta) = \mathbb{E} [V|\pi_{\theta}]. \quad (8)$$

Critic: The critic network $Q_{\phi}(\mathbf{S}, \mathbf{A})$ serves as an estimator to predict an action-value function (also termed as Q-function), i.e., $\mathbb{E}[V|\mathbf{S}, \mathbf{A}, \pi_{\theta}]$, and ϕ denotes the parameter of the associated deep neural networks. In general, the critic is designed to fine-tune the actor, which yields

$$\pi_{\theta}(\mathbf{S}) = \arg \max_{\mathbf{A} \in \mathcal{A}} Q_{\phi}(\mathbf{S}, \mathbf{A}). \quad (9)$$

By recalling (7), it leads to the following recursive equation:

$$Q_{\phi}(\mathbf{S}, \mathbf{A}) = \mathbb{E}_{\mathbf{S}', R|\mathbf{S}, \mathbf{A}} [R + \gamma Q_{\phi}(\mathbf{S}', \pi_{\theta}(\mathbf{S}'))], \quad (10)$$

where \mathbf{S}' denotes the subsequent state after taking \mathbf{A} under state \mathbf{S} ; R denotes the corresponding instant reward; and the expectation is over all of the possible occurrences of (\mathbf{S}', R) .

Learning Algorithm: As a category of policy gradient approaches, actor parameter θ is updated by using stochastic gradient descent, where the gradient of the policy can be given by *Deterministic Policy Gradient Theorem* [31]; concerning the critic network, parameter ϕ is updated according to *Temporal Difference*. Readers are referred to [31] for greater details.

Although DDPG has achieved great success in addressing many continuous decision-making tasks, the action space in our problem (defined by (4)) could restrain it from being efficient. Specifically, to confine the output of the actor to be feasible, the authors in [28] propose to first use the activation function *SoftMax* to normalize the output of the last layer, which is then multiplied by the caching capacity (e.g., L). In other words, denoting the output of the neurons in the last layer by $\mathbf{z} = [z_{f,b}]_{f \in \mathcal{F}, b \in \mathcal{B}}$, one can construct the following activation function:

$$\sigma_{f,b}(\mathbf{z}) = L \times \frac{\exp(z_{f,b})}{\sum_{f' \in \mathcal{F}} \exp(z_{f',b})}, \forall f \in \mathcal{F}, \forall b \in \mathcal{B}. \quad (11)$$

However, the resulting elements could surpass 1 when $L \gg 1$; directly clipping it to 1 may lead to a very poor caching decision if an element $\sigma_{f,b}(\mathbf{z}) = L$ exists. This practice contradicts Proposition 1 and may degrade the performance. In the following subsections, we formally analyze this issue and propose an efficient approach to overcome this challenge.

B. A Homotopy Optimization Based Approach

For a class of RL problems, the corresponding action space \mathcal{A} could be some constraints inefficient to be directly satisfied through designing neural networks μ_{θ} . More specifically, we consider the following situation: let $\mu_{\theta}(\mathbf{S})$ be the output of neural networks μ_{θ} given state \mathbf{S} . There exists some state $\mathbf{S} \in \mathcal{S}$ such that $\mu_{\theta}(\mathbf{S}) \notin \mathcal{A}$. We thereby term $\mu_{\theta}(\mathbf{S})$ as a proto-action. To deal with this issue, a straightforward approach is to use a mapping function $\sigma_{\mathcal{A}}(\cdot)$, which can project a proto-action $\mu_{\theta}(\mathbf{S})$ into the action space, i.e., $\sigma_{\mathcal{A}}[\mu_{\theta}(\mathbf{S})] \in \mathcal{A}$. Thus, a feasible policy function can be given by $\pi_{\theta}(\cdot) = \sigma_{\mathcal{A}}[\mu_{\theta}(\cdot)]$. Accordingly, the associated policy-based RL problem is supposed to take the following form:

$$\max_{\theta} J(\theta|\sigma_{\mathcal{A}}) \triangleq \mathbb{E} [V|\mu_{\theta}, \sigma_{\mathcal{A}}]. \quad (12)$$

Nevertheless, for many constrained RL applications, poor actions are likely to be generated after projection. Like the example in the previous subsection, mapping a proto-action with a dominant element $a_{f,b}^t = L$ into an one-hot vector may lead to a very sparse caching vector; this case implies a very low caching resource available at SBSs. When frequently encountering this instance during training, using mapping methods may not guarantee network parameters to be efficiently updated. Thus, it will lead to a suboptimal policy.

To remedy this method, a natural idea is to seek a proper way to penalize the performance loss caused by mapping a proto-action $\mu(\mathbf{S}^t)$ into a feasible one given any state \mathbf{S}^t . Let $g(\cdot|\sigma_{\mathcal{A}})$ be a general penalty function, which needs to be designed according to the corresponding problem. In the coded caching problem, inspired by Proposition 1, a penalty function can be given by

$$g(\mathbf{S}|\sigma_{\mathcal{A}}) = BL - \|\text{vec}(\sigma_{\mathcal{A}}(\mathbf{A}))\|_1|_{\mathbf{A}=\mu_{\theta}(\mathbf{S})}, \quad (13)$$

where $\|\cdot\|_1$ denotes l_1 -norm; and this penalty indicates the remaining storage over all caching units after taking action $\sigma_{\mathcal{A}}(\mathbf{A})$.

The proposed approach is then built upon maximizing a homotopy function:

$$\max_{\theta} J_{\text{hom}}(\theta|\lambda, \sigma_{\mathcal{A}}) \triangleq J(\theta|\sigma_{\mathcal{A}}) + \lambda G, \quad (14)$$

where the discount cumulative penalty, i.e., $G = \mathbb{E} \left[\sum_{t=0}^{+\infty} (\gamma)^t g(\mathbf{S}^t|\sigma_{\mathcal{A}}) \right]$, is finite due to a discount factor $\gamma \in [0, 1)$; and $\lambda \leq 0$ is a homotopy variable, such that:

$$J_{\text{hom}}(\theta|\lambda, \sigma_{\mathcal{A}}) = \begin{cases} J(\theta|\sigma_{\mathcal{A}}), & \lambda = 0, \\ J_{\text{hom}}(\theta|\lambda_{\min}, \sigma_{\mathcal{A}}), & \lambda = \lambda_{\min}. \end{cases} \quad (15)$$

At this stage, we introduce the following lemma to address problem (12) through a typical homotopy optimization method [29].

Lemma 1 *On the basis of homotopy optimization, one can initialize a sequence of positive values, i.e., $\delta^i, i = 1, \dots, I$, subject to:*

$$\sum_{i=1}^I \delta^i = -\lambda_{\min}, \quad (16)$$

and also initialize a point (θ^0, λ^0) , where θ^0 denotes a (local) optimizer of $J_{\text{hom}}(\theta|\lambda^0, \sigma_{\mathcal{A}})$ and $\lambda^0 = \lambda_{\min}$; and then iterate the following update:

$$\lambda^i = \lambda^{i-1} + \delta^i, \quad (17)$$

and calculate a local optimizer θ^i of $J_{\text{hom}}(\theta|\lambda^i, \sigma_{\mathcal{A}})$ by using gradient descent starting from θ^{i-1} . Eventually, this homotopy approach is able to result in point $(\theta^I, 0)$, where θ^I is a local minimizer of problem (12) [29].

The motivation of the homotopy optimization approach is as follows: Starting with a sufficient small value of $\lambda_{\min} < 0$, a very large cumulative penalty $|\lambda_{\min} G|$ may penalize the corresponding policy (parametrized by θ^0) to generate desirable actions, e.g., caching decisions that fully exploit available caching storage in the considered problem. Thereafter, by using homotopy optimization, we attempt to carefully tune policy parameter θ^0 to a (local) optimizer θ^I of the original problem (12), which is likely to produce good decisions despite of applying mapping function $\sigma_{\mathcal{A}}(\cdot)$.

C. Proposed DRL Algorithm

In each iteration of homotopy optimization, computing an optimizer (e.g., θ^i) of problem (14) offline is somehow impractical under an unknown temporally evolving environment. For this reason, we custom-build HDDPG for problem (12) by recasting the basic elements of DRL and unfolding the iterative procedure of homotopy optimization introduced in Lemma 1.

Specifically, as an actor-critic approach, HDDPG maintains a parametrized critic $Q_{\text{hom},\phi}(\mathbf{S}, \mathbf{A})$ and actor $\mu_{\theta}(\mathbf{S})$ in addition to a mapping function $\sigma_{\mathcal{A}}(\cdot)$, where a feasible policy is given by $\pi_{\theta}(\mathbf{S}) = \sigma_{\mathcal{A}}[\mu_{\theta}(\mathbf{S})]$. Following the sketch of a plain DDPG,² we introduce the proposed algorithm as follows. Evidently, the objective $J_{\text{hom}}(\theta|\lambda, \sigma_{\mathcal{A}})$ can be equivalently reformulated as:

$$J_{\text{hom}}(\theta|\lambda, \sigma_{\mathcal{A}}) = \mathbb{E} \left[\sum_{t=0}^{+\infty} (\gamma)^t (R^{t+1} + \lambda g(\mathbf{S}^t|\sigma_{\mathcal{A}})) \right]. \quad (18)$$

Accordingly, we define the homotopy reward after taking action \mathbf{A}^t as:

$$R_{\text{hom}}^{t+1} = R_{\text{hom}}(\mathbf{S}^{t+1}, \mathbf{S}^t, \mathbf{A}^t) \triangleq R^{t+1} + \lambda g(\mathbf{S}^t|\sigma_{\mathcal{A}}), \quad (19)$$

which is known at epoch $t+1$. Then, the homotopy Q-function can be given by:

$$Q_{\text{hom}}(\mathbf{S}^t, \mathbf{A}^t) = \mathbb{E} \left[\sum_{\tau=0}^{+\infty} (\gamma)^{t+\tau} R_{\text{hom}}^{t+\tau+1} \middle| \mathbf{S}^t, \mathbf{A}^t, \mu_{\theta}, \sigma_{\mathcal{A}} \right],$$

which implies the discount cumulative homotopy reward after taking an action \mathbf{A}^t under state \mathbf{S}^t and thereafter following

²For clarity, we term the DRL algorithm proposed in [31] as plain DDPG.

policy $\pi_{\theta}(\cdot) = \sigma_{\mathcal{A}}[\mu_{\theta}(\cdot)]$. As a direct deduction of the Bellman optimality equation [32], we have the following *Homotopy Bellman Optimality Equation*.

Lemma 2 *An optimal $Q_{\text{hom}}^*(\mathbf{S}, \mathbf{A})$ satisfies the following recursive equality:*

$$Q_{\text{hom}}^*(\mathbf{S}, \mathbf{A}) = \mathbb{E}_{R_{\text{hom}}, \mathbf{S}'|\mathbf{S}, \mathbf{A}} \left[R_{\text{hom}} + \gamma \max_{\mathbf{A}' \in \mathcal{A}} Q_{\text{hom}}^*(\mathbf{S}', \mathbf{A}') \right],$$

where \mathbf{S}' denotes the subsequent state after taking an optimal action \mathbf{A} ; and R_{hom} denotes the associated homotopy reward.

Accordingly, to estimate an optimal homotopy Q-function, the critic $Q_{\text{hom},\phi}(\mathbf{S}, \mathbf{A})$ can be learned by using Lemma 2. Specifically, we update ϕ by minimizing the following loss function:

$$\text{Loss}(\phi) = \mathbb{E}_{\xi_{\text{hom}}} \left[(y - Q_{\text{hom},\phi}(\mathbf{S}, \mathbf{A}))^2 \right], \quad (20)$$

where $\xi_{\text{hom}} = (\mathbf{S}, \mathbf{A}, R_{\text{hom}}, \mathbf{S}')$; and y denotes the target value:

$$y = R_{\text{hom}} + \gamma Q_{\text{hom},\phi}(\mathbf{S}', \mathbf{A}') \big|_{\mathbf{A}' = \pi_{\theta}(\mathbf{S}')}. \quad (21)$$

Regarding the update of the actor, it depends on the gradient of the objective $J_{\text{hom}}(\theta|\lambda, \sigma_{\mathcal{A}})$, which brings us to the following *Deterministic Policy Gradient Theorem for HDDPG*.

Lemma 3 *Consider a homotopy deep deterministic policy with a continuous action space \mathcal{A} and a homotopy variable λ , as well as a mapping function $\sigma_{\mathcal{A}}$. Suppose that $\sigma_{\mathcal{A}}(\cdot)$ is continuous. Then, the deterministic policy gradient exists when $\nabla_{\theta} \mu_{\theta}(\mathbf{S})$ and $\nabla_{\mathbf{A}} Q_{\text{hom},\phi}(\mathbf{S}, \mathbf{A})$ exist, i.e.:*

$$\nabla_{\theta} J_{\text{hom}}(\theta|\lambda, \sigma_{\mathcal{A}}) \triangleq \mathbb{E}_{\mathbf{S}} \left[\nabla_{\mathbf{A}} Q_{\text{hom},\phi}(\mathbf{S}, \mathbf{A}) \big|_{\mathbf{A} = \sigma_{\mathcal{A}}[\mu_{\theta}(\mathbf{S})]} \nabla_{\mathbf{A}'} \sigma_{\mathcal{A}}(\mathbf{A}') \big|_{\mathbf{A}' = \mu_{\theta}(\mathbf{S})} \nabla_{\theta} \mu_{\theta}(\mathbf{S}) \right]. \quad (22)$$

Proof. See Appendix B. ■

Finally, we leverage inexact gradient descent methods to update $\{\theta, \phi, \lambda\}$ [33]. In particular, the updates of ϕ, θ occur at each epoch, i.e.:

$$\phi \leftarrow \phi - \alpha_c \nabla_{\phi} \text{Loss}(\phi), \quad (23)$$

$$\theta \leftarrow \theta + \alpha_a \nabla_{\theta} J_{\text{hom}}(\theta|\lambda, \sigma_{\mathcal{A}}), \quad (24)$$

where α_c and α_a are the learning rates of the critic and actor, respectively; and λ can be updated slowly, i.e., after every I_0 epochs, one can execute the following:

$$\lambda^{iI_0} \leftarrow \lambda^{(i-1)I_0} + \delta^i, \quad (25)$$

where sequence $\{\lambda^i\}_{i=1}^I$ should meet the equality in (16).

Remark 2 *In contrast with plain DDPG, which constitutes a special case of the proposed HDDPG, i.e., $\lambda = 0$, properly introducing a penalty term into the objective function assists to infer which actions should be better to take and avoid becoming stuck in suboptimal solutions. More importantly, we unfold the homotopy optimization approach in Lemma 1 into a DRL, which can be done through interacting with environments.*

In the ensuing sections, we will apply HDDPG to the cooperative coded caching problem and propose a centralized caching design, and further generalize HDDPG in decentralized settings to reduce complexity and communication cost.

IV. CENTRALIZED HDDPG BASED COOPERATIVE CODED CACHING

In this section, we introduce a centralized HDDPG (C-HDDPG) design for multi-agent cooperative coded caching. As illustrated in Fig. 2(a), the system operation is at the level of centralized control. The CP serves as a centralized agent and coordinates the cooperative caching policies for all SBSs based on global information. To realize this, the CP maintains a (centralized) critic network $Q_\phi(\mathbf{S}, \mathbf{A})$ and a (centralized) actor network $\mu_\theta(\mathbf{S})$, as well as a mapping function $\sigma_{\mathcal{A}}(\cdot)$. In what follows, we first introduce a detailed implementation of the proposed centralized design, and then analyze its communication overhead and complexity.

A. Proposed Centralized HDDPG based Design

The system operation includes two procedures, i.e., network training and network evaluation. In general, during network evaluation, the CP simply leverages the actor and mapping function to make caching decisions, while the critic is only necessary during training procedure to fine-tune the actor. The details of network design and training procedure are introduced as follows.

Network Design: In general, both networks, i.e., critic and actor, can be implemented by fully connected neural networks where each hidden layer has a batch of neurons and an activation function to perform nonlinear transformations [34]. The output of the critic should be a scalar, which corresponds to the estimated value of the Q-function. To generate feasible actions, we elaborate on how to design the actor network μ_θ and mapping function $\sigma_{\mathcal{A}}$. It is evident that the number of neurons in the output layer of μ_θ should match the dimension of a joint action, i.e., $F \times B$ (and these neurons output a long vector $\mathbf{z} = [z_{1,1}, z_{2,1}, \dots, z_{f,b}, \dots, z_{F,B}]$). Then, we use the activation function (e.g., realized by *Scaling* and *SoftMax* in (11)) to refine \mathbf{z} , i.e., $\sigma_{f,b}(z)$, which thereafter is filtered by a mapping function $\sigma_{\mathcal{A}}(\cdot) = \min\{\mathbf{1}_{F \times B}, \cdot\}$. Accordingly, any proto-action $\mu_\theta(\mathbf{S})$ can be mapped into a feasible action, i.e., $\min\{\mathbf{1}_{F \times B}, \mu_\theta(\mathbf{S})\}$.

Update: To proceed, the technique of *Replay Buffer (RB)* Ξ is introduced to store historical experiences $\xi^t = (\mathbf{S}^t, \mathbf{A}^t, R_{\text{hom}}^{t+1}, \mathbf{S}^{t+1})$, which serves as the data set for network training. The buffer size $|\Xi|$ is usually finite, and thus the most outdated experience should be replaced by the current one as long as Ξ is fully loaded. Subsequently, at each epoch, we can randomly sample a mini-batch of N experiences (e.g., set Ξ_N) from RB to update parameters of the critic and the actor networks. More concretely, parameter ϕ of the critic network can be updated by minimizing the following loss function:

$$Loss(\phi) = \mathbb{E}_{\xi^t \sim \Xi_N} \left[\left(y_{\phi^-}^t - Q_{\text{hom},\phi}(\mathbf{S}^t, \mathbf{A}^t) \right)^2 \right], \quad (26)$$

³For instance, $\min\{\mathbf{1}, \mathbf{X}\}$ is an element-wise operator that executes $\max\{1, x_{f,b}\}$ for any element $x_{f,b}$ of \mathbf{X} .

where the expectation is over all of the sampled experiences; $y_{\phi^-}^t$ denotes the target value:

$$R_{\text{hom}}^{t+1} + \gamma Q_{\text{hom},\phi^-}(\mathbf{S}^{t+1}, \sigma_{\mathcal{A}}[\mu_{\theta^-}(\mathbf{S}^{t+1})]), \quad (27)$$

and $Q_{\text{hom},\phi^-}(\mathbf{S}, \mathbf{A})$ and $\mu_{\theta^-}(\mathbf{S})$ denote the target critic and the target actor with parameters ϕ^- and θ^- , respectively. To stabilize training [31], target networks should be slowly updated, i.e.:

$$\phi^- \leftarrow \tau \phi + (1 - \tau) \phi^-, \quad (28)$$

$$\theta^- \leftarrow \tau \theta + (1 - \tau) \theta^-, \quad (29)$$

where τ is a very small step size. With regard to updating parameter θ , the corresponding homotopy deterministic policy gradient $\nabla_{\theta} J_{\text{hom}}$ can be estimated by (22). In addition, the homotopy variable should be updated according to (25).

Exploration: To avoid being stuck in suboptimal decisions, exploration is usually needed during network training. The purpose of this process is to gather sufficient experiences, which then are used to infer what actions should be adopted under different states. In continuous decision-making applications, a typical method is to add Ornstein-Uhlenbeck random noise to the action generated by the actor [31], i.e.:

$$\pi_{\theta}^{\text{explore}}(\mathbf{S}^t) = \sigma'_{\mathcal{A}}(\pi_{\theta}(\mathbf{S}^t) + \beta^t \Delta^t), \quad (30)$$

where $\sigma'_{\mathcal{A}}$ is a simple mapping function that maps a noise-perturbed action into a feasible one; $\Delta^t = [\delta_{f,b}^t] \in \mathbb{R}^{F \times B}$; and each element $\delta_{f,b}^t$ denotes a sample drawn from a continuous Ornstein-Uhlenbeck process [31]; and $\beta^t \geq 0$ is a diminishing parameter.

To this end, an entire implementation of this centralized control is shown in Algorithm 1.

B. Fronthaul Communication Complexity

In the proposed centralized caching design, the CP needs to frequently communicate with SBSs during network training and evaluation. Herein, we briefly analyze fronthaul communication complexity of this centralized control, which is described by the total dimension of variables that are transmitted between the CP and SBSs. We consider the worst case, in which each SBS is fully loaded and serves a maximum number of users, e.g., $|\mathcal{K}_b^t| = K$. Specifically, during network training, the CP needs to obtain information about the system state $[\{f_k^t\}, \mathcal{E}^t, \mathbf{L}^t]$ at each epoch. The dimension of user requests should be BK . Network connectivity \mathcal{E}^t can be computed by knowing the coordinates of the active users; by denoting the coordinates as two dimensional vectors, the total dimension of user positions is $2BK$. Clearly, the CP has the exact information about cache allocation \mathbf{L}^t , which is termed as \mathbf{A}^{t-1} in its RB ; thus, no fronthaul cost is involved. Afterwards, the CP uses fronthaul links to inform caching decisions $\{\mathbf{A}_b^t\}$; the total dimension of the involved variables is given by BF . Regarding reward R^{t+1} , it can be inferred from state \mathbf{S}^{t+1} . Hence, the overall fronthaul communication complexity during network training is $\mathcal{O}(3BK + BF)$. When the system runs in an evaluation procedure, the CP again needs to know the system state and inform each SBS of its

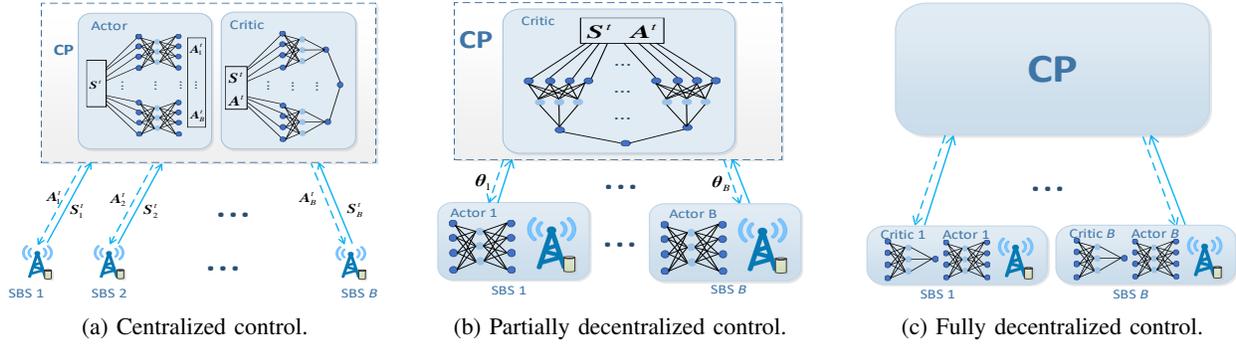


Fig. 2: Proposed multi-agent DRL-based cooperative coded caching framework.

Algorithm 1 Proposed C-HDDPG based Cooperative Coded Caching

- 1: Initialize $\tau, \alpha_c, \alpha_a, \gamma, N, I_0, \lambda = \lambda_{\min}, i = 0$
- 2: Initialize parameter ϕ for critic network and θ for actor network
- 3: Initialize parameters $\phi^- \leftarrow \phi, \theta^- \leftarrow \theta$ for target critic network and target actor network
- 4: Initialize $RB \Xi$ and mapping function σ_A
- 5: Initialize $\delta^1, \delta^2, \dots, \delta^I$
- 6: **for** $t = 0, 1, 2, \dots$ **do**
- 7: Input S^t to actor and output $A^t = \pi_{\theta}^{\text{explore}}(S^t)$
- 8: Take action A^t and observe S^{t+1}, R^{t+1}
- 9: Calculate R_{hom}^{t+1} by (19)
- 10: Store $\xi^t = (S^t, A^t, R_{\text{hom}}^{t+1}, S^{t+1})$ into RB
- 11: **procedure** TRAINHDDPG
- 12: Randomly sample a mini-batch of N experiences from relay buffer as Ξ_N
- 13: Update ϕ, θ by (23) and (24), respectively
- 14: Update ϕ^-, θ^- by (28) and (29), respectively
- 15: **if** $t == i \times I_0$ **then**
- 16: $\lambda \leftarrow \lambda + \delta^i, i \leftarrow i + 1$
- 17: **end if**
- 18: **end procedure**
- 19: **end for**

caching decision. Consequently, the corresponding fronthaul communication complexity is $\mathcal{O}(3BK + BF)$.

Moreover, the critic and actor are built upon system states and joint actions, i.e., (S^t, A^t) , which is in the order of $\mathcal{O}(B^2)$ of local observations and actions, i.e., (S_b^t, A_b^t) . For this reason, the computational complexity would be excessively high as the number of agents increases for a continuous RL problem [35]. To address this issue, we now focus on developing efficient decentralized algorithms in following sections.

V. PARTIALLY DECENTRALIZED HDDPG BASED COOPERATIVE CODED CACHING

In this section, to circumvent excessive communication cost and high complexity in the centralized design, we develop a partially decentralized (PD)-HDDPG based cooperative coded caching design. This scheme operates at the level of partially decentralized control, in the sense that a (centralized) critic

Algorithm 2 Proposed PD-HDDPG based Cooperative Coded Caching

- 1: Initialize $\tau, \alpha_c, \alpha_a, \gamma, N, I_0, \lambda = \lambda_{\min}$
- 2: Initialize parameter ϕ for the critic and $\theta = \{\theta_b, \forall b \in \mathcal{B}\}$ for actors
- 3: Initialize parameters $\phi^- \leftarrow \phi, \theta^- \leftarrow \theta$ for target critic and target actors
- 4: Initialize $RB \Xi$ and mapping functions $\{\sigma_b, \forall b \in \mathcal{B}\}$
- 5: Initialize $\delta^1, \delta^2, \dots, \delta^I$
- 6: **for** $t = 0, 1, 2, \dots$ **do**
- 7: **for** $b \in \mathcal{B}$ **do**
- 8: Observe S_b^t and compute action A_b^t through π_b with proper exploration
- 9: **end for**
- 10: Execute $\{A_b^t, \forall b\}$ in a real environment and observe $S^{t+1} = \{S_1^{t+1}, \dots, S_B^{t+1}\}, R^{t+1}$
- 11: Let $A^t = \{A_1^t, \dots, A_B^t\}$ and calculate R_{hom}^{t+1} by (19)
- 12: The CP pushes $\xi^t = (S^t, A^t, R_{\text{hom}}^{t+1}, S^{t+1})$ into RB
- 13: **procedure** TRAINHDDPG
- 14: Randomly sample a mini-batch of experiences Ξ_N
- 15: Calculate $\nabla_{\phi} \text{Loss}(\phi)$ and $\phi \leftarrow \phi - \alpha_c \nabla_{\phi} \text{Loss}(\phi)$
- 16: **for** $b \in \mathcal{B}$ **do**
- 17: $\theta_b \leftarrow \theta_b + \alpha_a \nabla_{\theta_b} J_{\text{hom}}$
- 18: **end for**
- 19: $\phi^- \leftarrow (1 - \tau)\phi^- + \tau\phi, \theta^- \leftarrow (1 - \tau)\theta^- + \tau\theta$
- 20: **if** $t == i \times I_0$ **then**
- 21: $\lambda \leftarrow \lambda + \delta^i, i \leftarrow i + 1$
- 22: **end if**
- 23: **end procedure**
- 24: **end for**

is used to train (local) actors that separately approximate the caching policy of each SBS.

A. Partially Decentralized Multi-Agent HDDPG

In a PD multi-agent framework, each agent maintains an actor and mapping function to produce its actions. To augment collaboration among multiple agents, these actors are trained with the aid of a (centralized) critic. Specifically, agent b

has an actor μ_b (parametrized by θ_b) and mapping function⁴ σ_b , which are able to map a (local) proto-action $\mu_b(\mathbf{S}_b)$ into the corresponding action space \mathcal{A}_b . Accordingly, the policy function for agent b can be expressed by $\pi_b = \sigma_b[\mu_b(\cdot)]$.

On the basis of homotopy optimization, all agents cooperatively seek policies to jointly maximize the following homotopy function:

$$J_{\text{hom}}(\theta_1, \dots, \theta_B | \lambda, \sigma_{\mathcal{A}}) = \mathbb{E} \left[\sum_{t=0}^{\infty} (\gamma)^t R_{\text{hom}}^{t+1} | \mu_{\theta}, \sigma_{\mathcal{A}} \right], \quad (31)$$

where we define $\mu_{\theta} \triangleq \{\mu_1, \dots, \mu_B\}$ and $\sigma_{\mathcal{A}} \triangleq \{\sigma_1, \dots, \sigma_B\}$; and the homotopy reward R_{hom}^t can be given by (19). Next, a (centralized) critic $Q_{\text{hom},\phi}(\mathbf{S}^t, \mathbf{A}_1^t, \mathbf{A}_2^t, \dots, \mathbf{A}_B^t)$ is leveraged to estimate $\mathbb{E} \left[\sum_{\tau=0}^{\infty} (\gamma)^{t+\tau} R_{\text{hom}}^{t+\tau+1} | \mathbf{S}^t, \mathbf{A}_b^t, \mu_b, \sigma_b, \forall b \in \mathcal{B} \right]$. Similar to C-HDDPG, parameter ϕ can be learned by minimizing the following loss function:

$$\text{Loss}(\phi) = \mathbb{E}_{\xi_t \sim \Xi_N} \left[\left(y_{\phi^-}^t - Q_{\text{hom},\phi}(\mathbf{S}^t, \mathbf{A}_1^t, \dots, \mathbf{A}_B^t) \right)^2 \right], \quad (32)$$

where $y_{\phi^-}^t$ denotes the target value, i.e., $R_{\text{hom}}^{t+1} + \lambda Q_{\text{hom},\phi^-}(\mathbf{S}^{t+1}, \mathbf{A}_1, \dots, \mathbf{A}_B) |_{\mathbf{A}_b = \sigma_b[\mu_b^-(\mathbf{S}_b^{t+1})]}$, and θ_b^-, ϕ^- are parameters of target actor μ_b^- and target critic Q_{hom,ϕ^-} , respectively. Furthermore, the gradient $\nabla_{\theta_b} J_{\text{hom}}(\theta_1, \dots, \theta_B | \lambda, \sigma_{\mathcal{A}})$ can be approximated by:

$$\mathbb{E}_{\xi_t \sim \Xi_N} \left[\nabla_{\mathbf{A}_b} Q_{\text{hom},\phi}(\mathbf{S}^t, \mathbf{A}_1, \dots, \mathbf{A}_B) |_{\mathbf{A}_b = \pi_b(\mathbf{S}_b^t)} \nabla_{\theta_b} \pi_b(\mathbf{S}_b^t) \right]. \quad (33)$$

Similarly, homotopy variable λ should be updated in accordance with (25).

B. Implementation

As depicted in Fig. 2 (b), we propose a PD-HDDPG based cooperative coded caching design. Specifically, the CP maintains a centralized (critic), while each SBS has a local actor and mapping function. In addition, the actor and mapping function are designed in the same manner as the centralized scheme to ensure that their outputs are feasible to (4). Particularly, $\sigma_b(\cdot) = \min\{\mathbf{1}_F, \cdot\}, \forall b \in \mathcal{B}$. During the training procedure, the critic and (local) actors should be learned in the CP. We again adopt the techniques of exploration and RB , and the entire procedure is similar to what we have presented in Algorithm 1. The detailed implementation is shown in Algorithm 2. Notably, after all actors are fine-tuned, the CP needs to send actor parameters (e.g., $\theta_b, \forall b$) to SBSs, which thereafter can locally compute actions.

C. Fronthaul Communication Complexity

During the training procedure, fronthaul communication complexity is the same as that of C-HDDPG, i.e., $\mathcal{O}(3BK + BF)$. When the system runs in an evaluation procedure,

⁴For simplicity of notation, μ_b and σ_b are abbreviations of μ_{θ_b} and $\sigma_{\mathcal{A}_b}$, respectively.

Algorithm 3 Proposed FD-HDDPG based Cooperative Coded Caching

- 1: Initialize $\tau, \alpha_c, \alpha_a, \gamma, I_0, \lambda = \lambda_{\min}$
- 2: Initialize parameters $\phi = \{\phi_b, \forall b \in \mathcal{B}\}$ for critics and $\theta = \{\theta_b, \forall b \in \mathcal{B}\}$ for actors
- 3: Initialize parameters $\phi^- \leftarrow \phi, \theta^- \leftarrow \theta$ for target critics and target actors
- 4: Initialize $\Xi_b, N_b, \forall b \in \mathcal{B}$ and $\{\sigma_b, \forall b \in \mathcal{B}\}$
- 5: Initialize $\delta^1, \dots, \delta^I$
- 6: **for** $t = 0, 1, 2, \dots$ **do**
- 7: **for** $b \in \mathcal{B}$ **do**
- 8: Execute $\mathbf{A}_b^t = \pi_b(\mathbf{S}_b^t)$ with proper exploration and calculate $g_b(\mathbf{S}_b^t) \leftarrow L - \|\mathbf{A}_b^t\|_1$
- 9: **end for**
- 10: Observe $\mathbf{S}_b^{t+1}, \forall b \in \mathcal{B}$, and R^{t+1}
- 11: Calculate R_{hom}^{t+1} by (34)
- 12: **for** $b \in \mathcal{B}$ **do**
- 13: SBS b stores $\xi^t = (\mathbf{S}_b^t, \mathbf{A}_b^t, R_{\text{hom}}^{t+1}, \mathbf{S}_b^{t+1})$ into Ξ_b
- 14: **end for**
- 15: **procedure** TRAINHDDPG
- 16: **for** $b \in \mathcal{B}$ **do**
- 17: Randomly sample a mini-batch of experiences $\Xi_{b,N}$
- 18: $\phi_b \leftarrow \phi_b - \alpha_c \nabla_{\phi_b} \text{Loss}(\phi_b)$
- 19: $\theta_b \leftarrow \theta_b + \alpha_a \nabla_{\theta_b} J_{\text{hom}}$
- 20: $\phi_b^- \leftarrow (1-\tau)\phi_b^- + \tau\phi_b, \theta_b^- \leftarrow (1-\tau)\theta_b^- + \tau\theta_b$
- 21: **if** $t == i \times I_0$ **then**
- 22: $\lambda \leftarrow \lambda + \delta^i, i \leftarrow i + 1$
- 23: **end if**
- 24: **end for**
- 25: **end procedure**
- 26: **end for**

each SBS computes its action locally; obviously, no fronthaul communication is incurred when observing content requests and positions of local users.

VI. FULLY DECENTRALIZED HDDPG BASED COOPERATIVE CODED CACHING

To further reduce complexity and fronthaul signaling, we propose a fully decentralized (FD) control for cooperative coded caching. Particularly, each SBS serves as an independent learner to locally train its caching policy. Hereunder, we first present the FD-HDDPG based caching design, and then briefly summarize the complexity of all of the proposed designs.

A. Fully Decentralized Cooperative Coded Caching Design

As shown in Fig 2(c), each SBS has a set of critic, actor, and mapping functions. These basic elements are designed in the same manner as that of C-HDDPG, but built upon local observations. More precisely, with an actor $\mu_b(\cdot)$ and mapping function σ_b , SBS b can obtain a feasible action by mapping a proto-action $\mu_b(\mathbf{S}_b)$ into the action space \mathcal{A}_b , i.e., $\sigma_b[\mu_b(\mathbf{S}_b)]$. SBSs are encouraged to cooperate with each other and receive a common reward R^t from the environment as the performance criterion to evaluate their policies. On the basis of homotopy

optimization, a (local) critic $Q_{\text{hom}, \phi_b}(\mathbf{S}_b^t, \mathbf{A}_b^t)$ is designed for SBS b to estimate $\mathbb{E} \left[\sum_{\tau=0}^{+\infty} (\gamma)^{t+\tau} R_{\text{hom}}^{t+\tau+1} | \mathbf{S}_b^t, \mathbf{A}_b^t, \mu_b, \sigma_b \right]$, where R_{hom}^t is defined as:

$$R_{\text{hom}}^t \triangleq R^t + \lambda \sum_{b \in \mathcal{B}} g_b(\mathbf{S}_b^{t-1} | \sigma_b), \quad (34)$$

and $g_b(\mathbf{S}_b^{t-1} | \sigma_b) \triangleq L - \|\sigma_b(\mu_b(\mathbf{S}_b^{t-1}))\|_1, \forall b$. Subsequently, each agent is envisioned to independently train its critic and actor. In addition, the training procedure should follow the same workflow as C-HDDPG, which is presented in Algorithm 3 in greater detail.

B. Fronthaul Communication Complexity

During network training, although fronthaul communications are not necessary for SBSs to obtain local observations, each SBS still needs to know the homotopy reward. SBSs first locally computes $\{g_b(\mathbf{S}_b^t | \sigma_b)\}$, which are then aggregated by the CP and subsequently sent back to each SBS. Therefore, fronthaul communication complexity during training is given by $\mathcal{O}(2B)$. For network evaluation, SBSs can directly calculate local actions according to their observations; thus, no fronthaul communication is needed. Finally, we show the fronthaul communication complexity of all proposed algorithms in Table II.

TABLE II: Fronthaul communication complexity

| | C-HDDPG | PD-HDDPG | FD-HDDPG |
|------------|-------------------------|-------------------------|-------------------|
| Training | $\mathcal{O}(3BK + BF)$ | $\mathcal{O}(3BK + BF)$ | $\mathcal{O}(2B)$ |
| Evaluation | $\mathcal{O}(3BK + BF)$ | 0 | 0 |

Remark 3 As a comparison, we summarize the characteristics of all algorithms. Operating in the centralized control, C-HDDPG incurs the highest fronthaul communication complexity in either training or evaluation procedure. In this algorithm, the CP needs to coordinate SBSs by aggregating local information and distributing individual actions. To reduce fronthaul communication complexity, PD-HDDPG is then devised with a partially decentralized manner. That is, a centralized critic is applied to train local policies for each SBS, which is capable to produce actions based on local observation. In this way, fronthaul communication complexity during the evaluation procedure is as low as that of FD-HDDPG. In the fully decentralized control, FD-HDDPG has the lowest fronthaul communication complexity during training and evaluation procedures, which might compromise performance. The reason being that each SBS serves as an independent decision-maker, which trains caching policies based on local information without the coordination of the CP. By enabling different levels of control, the proposed caching framework is envisioned to possess advantages of superior performance, as well as scalability to large-scale systems.

VII. PERFORMANCE EVALUATIONS

A. Simulation Setup

Unless stated otherwise, we consider the following default settings: a SCN covers a square area of $[0, 1] \text{ km} \times [0, 1]$

km; four SBSs are uniformly deployed in the region, each of which has a communication radius of $r_0 = 500 \text{ m}$ and can provide service for a maximum of 100 mobile users during each epoch; and mobile users are randomly distributed by following a Poisson point process (PPP) with density 9.5×10^{-5} during each epoch. Moreover, user preferences towards content are considered to have multiple patterns, i.e., each preference pattern follows a Zipf distribution, i.e., $p_f = \zeta_f^{-\kappa} / \sum_{f' \in \mathcal{F}} \zeta_{f'}^{-\kappa}$, where ζ_f denotes the popularity rank of content item f , which is temporally and randomly evolving as time passes; κ denotes a skewness factor and randomly takes a value from $\{0.5, 1, 1.5, 2\}$; a catalog of 20 content items are encoded by MDS codes; and each SBS has a fractional caching capacity $L/F = 0.2$, which indicates that each SBS can fetch 20% of the total content.

B. Convergence Behavior

To analyze the proposed caching framework, we consider the following baselines:

- **Centralized Optimization-Cache Update (CO-CU):** This is a centralized optimization-based design, which is performed at the CP [17], [18]. Specifically, one can first estimate the probability of each content item that could be requested by users under the coverage of SBS b , i.e., $p_{f,b} = N_{f,b} / \sum_{f' \in \mathcal{F}} N_{f',b}, \forall f$, where $N_{f,b}$ denotes how many requests of content item f that SBS b receives at the current epoch. Similar to the cooperative caching problem formulated in [17], the joint caching decision (e.g. \mathbf{A}^t) can be optimized by minimizing the expected fronthaul traffic loads together with cache updating cost.
- **Local Optimization-Cache Update (LO-CU):** This scheme works at a level of decentralized control. Instead of optimizing joint action \mathbf{A}^t via CP in CO-CU, each SBS separately calculates its caching decision (e.g. \mathbf{A}_b^t) by minimizing the expected fronthaul loads corresponding to local user requests from its communication range.
- **Random Cache Update (RCU):** At each epoch, every SBS randomly updates its caching resource until it reaches storage limits.
- **Plain DDPG based Schemes:** To assess the effectiveness of the proposed DRL, we consider to implement plain DDPG in different levels of controls, i.e., **C-DDPG, PD-DDPG and FD-DDPG**, each of which follows a similar idea to the proposed caching approach.

To implement the proposed algorithms, every critic has three hidden layers, each of which contains 512 neurons. Each actor consists of three hidden layers with 256, 128, and 64 neurons, respectively. All of the networks are trained by the Adam optimizer with a polynomial learning rate (e.g., readers are referred to [36] for additional details), where we set initial learning rates for actors and critics as 0.01 and 0.001, respectively, and the power factor for decay as 0.9. A mini-batch of 100 experiences are randomly sampled every time from RB that is capable of storing 5000 past experiences. Every target critic or target actor is updated by a step size $\tau = 0.001$, and the discount factor $\gamma = 0.99$. To perform policy explorations, we use an Ornstein-Uhlenbeck process

with mean 0 and variance 1; the associated diminishing parameter β is initialized as 0.9, and decreased at a rate of 0.995 every epoch until it reaches 0.0001. Finally, we initialize the following sequence to update the homotopy variable λ , i.e., $\{\delta^i = -0.1 \times \lambda_{\min}, i = 1, 2, \dots, 10\}$, where $\lambda_{\min} = -0.005$ and we update λ every $I_0 = 1000$ epochs.

As shown in Fig. 3, we first illustrate the learning curves of the proposed algorithm under centralized control. Particularly, we vary parameter λ_{\min} to investigate the impact of penalty. Each result is averaged over $N = 5000$ epochs, i.e., $\sum_{\tau=t-N+1}^t R^\tau/N$. It can be observed that in the first 10^4 epochs, the curves of HDDPG based algorithms rise markedly, and notable gaps can be observed between plain DDPG and HDDPG based algorithms. In subsequent epochs, the learning curves of these DRL algorithms increase gradually until convergence. Clearly, when λ_{\min} is -0.005 or -0.015, HDDPG based algorithms achieve higher caching rewards than those of plain DDPG; when λ_{\min} goes down to -0.1, the curve increases fairly slowly and converges to a level that is very close to plain DDPG. Therefore, if $\lambda_{\min}G$ is significantly large compared with the objective, it could dominant the actual objective, eventually leading to a suboptimal policy. These observations demonstrate that properly introducing the penalty term to RL (e.g., $\lambda_{\min}G$ in (14)) could assist agents to infer better actions and speed-up convergence behavior. Furthermore, we propose to implement HDDPG (with $\lambda_{\min} = -0.005$) by initializing RB with 10% warm-up experiences via optimization baselines (e.g., CU-CO) rather than the conventional exploration method used in plain DDPG, i.e., only utilizing Ornstein-Uhlenbeck random noise to explore action spaces. In this way, the proposed implementation can further improve performance compared with the conventional exploration under the same $\lambda_{\min} = -0.005$. This result implies that taking advantage of a good baseline improves the efficiency of exploration in DRL.

Fig. 4 and 5 show the learning curves of the proposed algorithms under partially and fully decentralized scenarios. As anticipated, the proposed PD-HDDPG and FD-HDDPG respectively outperform plain DDPG in both scenarios. It is worth noting that, the shaded region around each learning curve shows reward deviations, which measure the robustness of each policy. Obviously, the DRL-based designs exhibit more centered results while the rewards, achieved by the optimization baseline, spread out over quite a broad range. This observation demonstrates the effectiveness of HDDPG to track and adapt to dynamic features of wireless networks.

C. Impact of System Parameters

In this subsection, we study the impacts of system parameters on the proposed caching framework. All of the results are obtained by averaging over 10^4 epochs. We first investigate the impact of caching capacity under a larger content catalog size (e.g., $F = 50$). Clearly, as shown in Fig. 6, C-HDDPG is always superior to other algorithms. When fractional caching capacity (e.g., L/F) is 10%, C-HDDPG achieves the lowest fronthaul traffic loads, e.g., 0.47, in contrast with PD-HDDPG and FD-HDDPG, e.g., 0.52. The superiority of C-HDDPG demonstrates the effectiveness of using global information to

enhance SBS collaboration. As fractional caching capacity grows larger, the gap between PD-HDDPG and FD-HDDPG becomes bigger. Indeed, with the aid of a centralized critic to train local policies, PD-HDDPG can allow SBSs to tightly collaborate in comparison to the fully decentralized scheme. Although FD-HDDPG depends on local observations only, it still outperforms CO-CU and LO-CU by 6.11% and 11.51% respectively, under the scenarios being studied. This observation demonstrates the remarkable advantages of using DRL algorithms to learn policies under dynamic environments over conventional optimization-based algorithms.

Hereunder, we conduct experiments to investigate the impacts of content popularity by varying the skewness factor of Zipf distribution. Moreover, under each scenario being investigated, the corresponding skewness factor is fixed as a constant, in which a larger skewness factor indicates a more concentrated content popularity. As can be seen, fronthaul traffic loads decrease as the skewness factor becomes larger for all of the algorithms except for RCU. The reason for this is that user requests are more likely to be accessed in local SBSs if their preferences are more centered. Furthermore, PD-HDDPG achieves comparable fronthaul traffic loads in contrast with C-HDDPG when the skewness factor is smaller than 1; after that, using centralized control only produces a marginal performance gain over decentralized control, yet with a significant implementation cost. This finding demonstrates that PD-HDDPG can efficiently obtain a satisfactory trade-off between complexity and performance.

To investigate the scalability of the proposed algorithms, we carry out experiments by varying the content catalog size. As depicted in Fig. 8, PD-HDDPG obtains a comparable performance to C-HDDPG when the content catalog size is smaller than 50; as more content items are considered, C-HDDPG achieves better performance due to utilization of global information. It is worth noting that, vast gaps can be observed between the proposed algorithms and baselines under either centralized or decentralized scenarios. More specifically, over the entire horizontal axis, C-HDDPG and PD-HDDPG can decrease fronthaul traffic loads by 10.54% and 7.68% respectively in comparison to CO-CU; whereas FD-HDDPG can reduce fronthaul traffic loads by 10.40% compared with LO-CU. All of these results corroborate the scalability of the proposed algorithms. Notably, the curve of RCU increases greatly and eventually surpasses 1 as the content catalog size grows larger. This is because the cache updating cost introduced by RCU could overtake fronthaul traffic loads arising from satisfying user requests under large scenarios.

We further investigate how the number of agents (i.e., SBSs) impacts the proposed multi-agent algorithms. In these settings, we vary the number of SBSs, and set the distance between two adjacent SBSs as 300 m. As shown in Fig 9, traffic loads exhibit a decreasing trend as more SBSs are available to participate in cooperative coded caching. In addition, all of the curves decrease relatively slowly when more than 15 SBSs are deployed, which implies that most users might already be able to access a maximum number of local SBSs, that is usually limited by communication coverage. More importantly, when the number of agents becomes large, PD-HDDPG always

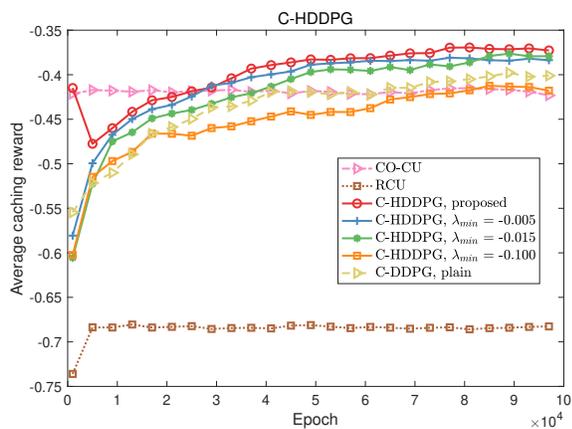


Fig. 3: Learning curves of C-HDDPG.

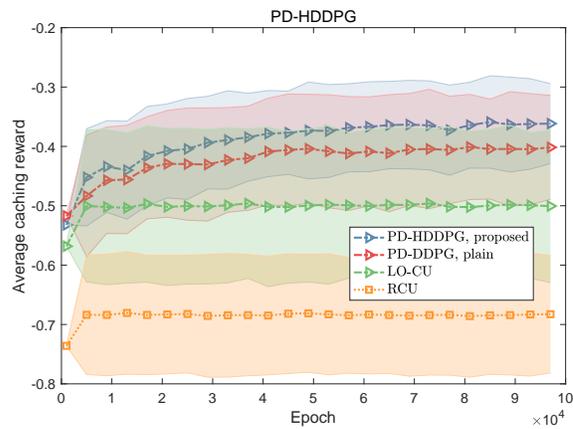


Fig. 4: Learning curves of PD-HDDPG.

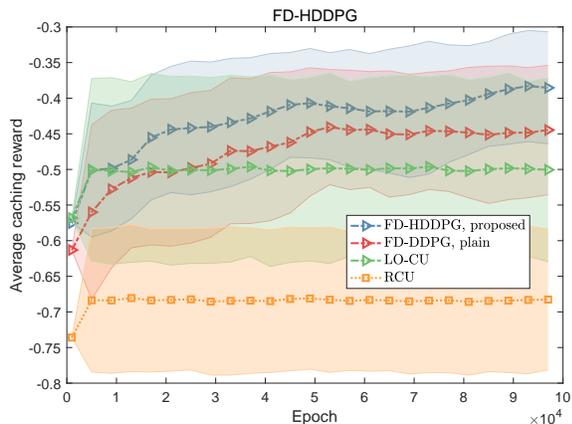


Fig. 5: Learning curves of FD-HDDPG.

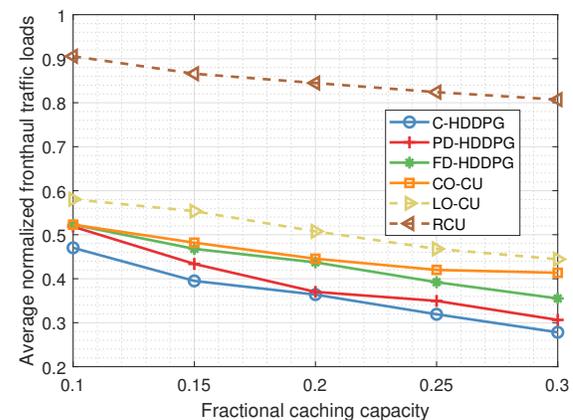


Fig. 6: Impact of fractional caching capacity.

achieves comparable results to C-HDDPG with significant reductions of signaling overhead and complexity. Concerning FD-HDDPG, it achieves slightly larger traffic loads than those of C-HDDPG and PD-HDDPG by at most 3.23% and 2.34%, respectively. These results demonstrate the potentials of utilizing decentralized controls as a large number of SBSs are deployed. Finally, Fig 10 illustrates the impact of SBS geometry on fronthaul traffic loads. We consider six SBSs and vary the distance of any two adjacent SBSs. As anticipated, the average traffic fronthaul loads become higher when SBSs are installed farther from each other. This is because, given a fixed communication range at SBSs, a larger SBS distance scenario provides less opportunity for SBS collaboration. Moreover, we can observe that PD-HDDPG generally performs pretty close to the centralized control, which again corroborates the effectiveness of utilizing the centralized critic. Indeed, PD-HDDPG operates with the least signaling overhead yet at the cost of slight performance degradation.

VIII. CONCLUSION

We have proposed a deep multi-agent reinforcement learning framework for dynamic cooperative coded caching at small cell networks. Particularly, we have developed a novel deep reinforcement learning algorithm, i.e., homotopy DDPG, to address the challenges arising from the resultant continuous

decision-making with constraints. From an engineering perspective, we have proposed centralized, partially decentralized, and fully decentralized controls to balance complexity and performance. Simulation results have confirmed that the proposed DRL outperforms plain DDPG under different levels of controls; and the proposed decentralized designs also achieve satisfactory performance compared with the centralized design.

APPENDIX

A. Proof of Proposition 1

Consider an optimal decision sequence $\{\mathbf{A}^t\}$, which results in an optimal value $J^* = \sum_f (\gamma)^t R^{t+1}$. Suppose that there exists t', b' such that $\sum_f a_{f,b'}^t = L'$ and $L' < L$; in addition, the corresponding decision $\mathbf{A}^{t'}$ is anticipated to impact rewards $R^{t'+1}$ and $R^{t'+2}$. To proceed, we first denote $\mathcal{F}' = \{f | a_{f,b'}^t - l_{f,b'}^t > 0\}$. Accordingly, one can create a sequence $\{o_f, \forall f \in \mathcal{F}\}$ where $0 \leq o_f \leq l_{f,b'}^t - a_{f,b'}^t, \forall f \in \mathcal{F} \setminus \mathcal{F}'$, and $o_f = 0, \forall f \in \mathcal{F}'$, such that $\sum_{f \in \mathcal{F}} o_f = L - L'$. Then, we consider a new decision $\underline{\mathbf{A}}^{t'} = [\underline{a}_{f,b}^{t'}]$ for epoch t' , where $\underline{a}_{f,b}^{t'} = a_{f,b}^{t'}$ for $\forall b \neq b'$, and $\underline{a}_{f,b'}^{t'} = a_{f,b'}^{t'} + o_f$; clearly, this gives rise to $\sum_{f \in \mathcal{F}} \underline{a}_{f,b'}^{t'} = L$. By checking the traffic loads in (5), one can verify that $\underline{R}^{t'+1} \geq R^{t'+1}$

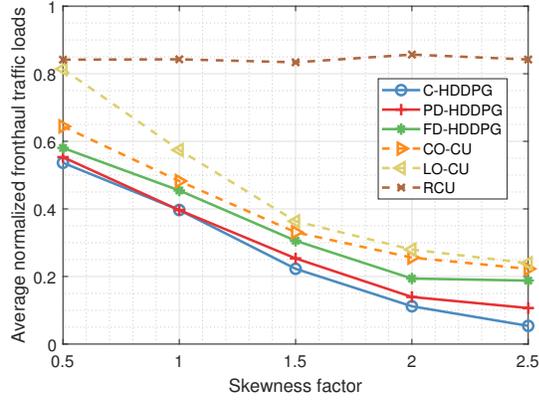


Fig. 7: Impact of content popularity.

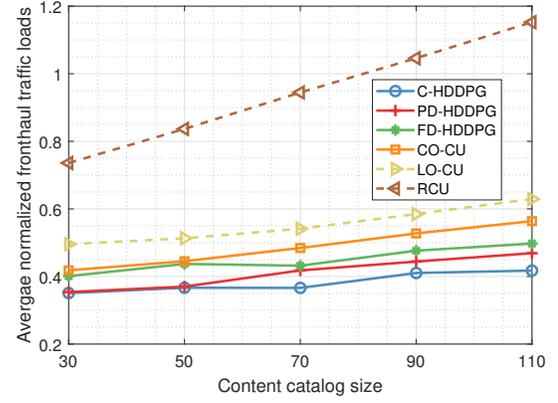


Fig. 8: Impact of content catalog size.

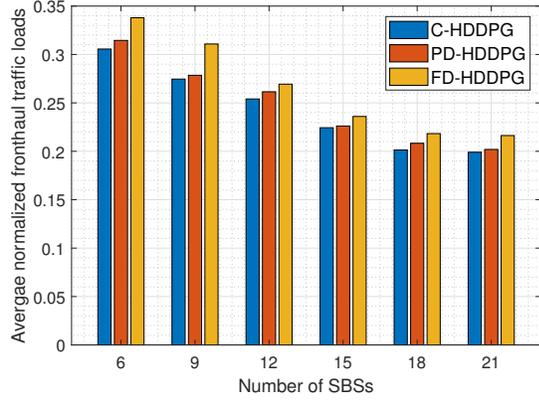


Fig. 9: Impact of number of SBSs.

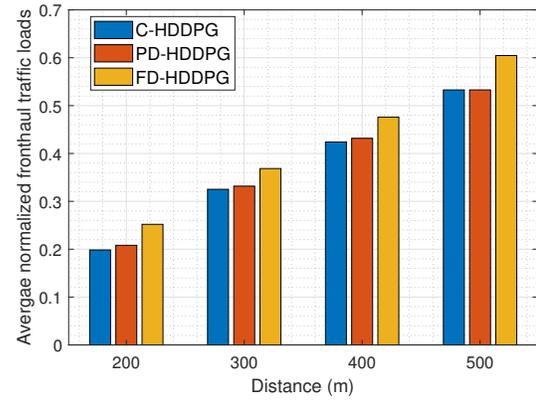


Fig. 10: Impact of SBS geometry.

and $\underline{R}^{t'+2} \geq R^{t'+2}$. We thereby claim that $\{\underline{\mathbf{A}}^t\}$ is also an optimal decision sequence where $\underline{\mathbf{A}}^t = \mathbf{A}^t$ for $\forall t \neq t'$. Hence, Proposition 1 holds.

B. Proof of Lemma 3

This proof follows similar procedures to the *Deterministic Policy Gradient Theorem* in [37]. Accordingly, $\nabla_{\theta} J_{\text{hom},\theta} = \nabla_{\theta} \int_{\mathcal{S}} p^0(\mathbf{S}) Q_{\text{hom},\phi}(\mathbf{S}, \mathbf{A})|_{\mathbf{A}=\sigma_{\mathcal{A}}[\mu_{\theta}(\mathbf{S})]} d\mathbf{S}$, where $p^0(\cdot)$ denotes probability density of state at epoch 0. Then, following the standard steps in [37] yields the following

$$\begin{aligned} & \nabla_{\theta} Q_{\text{hom},\phi}(\mathbf{S}, \mathbf{A})|_{\mathbf{A}=\sigma_{\mathcal{A}}[\mu_{\theta}(\mathbf{S})]} \\ &= \nabla_{\mathbf{A}} Q_{\text{hom},\phi}(\mathbf{S}, \mathbf{A})|_{\mathbf{A}=\sigma_{\mathcal{A}}[\mu_{\theta}(\mathbf{S})]} \nabla_{\mathbf{A}'} \sigma_{\mathcal{A}}(\mathbf{A}')|_{\mathbf{A}'=\mu_{\theta}(\mathbf{S})} \nabla_{\theta} \mu_{\theta}(\mathbf{S}) \\ &+ \int_{\mathcal{S}} \gamma p(\mathbf{S} \rightarrow \hat{\mathbf{S}}, 1) \nabla_{\theta} Q_{\text{hom},\phi}(\hat{\mathbf{S}}, \hat{\mathbf{A}})|_{\hat{\mathbf{A}}=\sigma_{\mathcal{A}}[\mu_{\theta}(\hat{\mathbf{S}})]} d\hat{\mathbf{S}}, \end{aligned} \quad (35)$$

where $p(\mathbf{S} \rightarrow \hat{\mathbf{S}}, t_0)$ denotes the probability density of state \mathbf{S} transiting to state $\hat{\mathbf{S}}$ after t_0 epochs; thereafter, one can continue to unfold (35), resulting in

$$\begin{aligned} & \nabla_{\theta} Q_{\text{hom},\phi}(\mathbf{S}, \mathbf{A})|_{\mathbf{A}=\sigma_{\mathcal{A}}[\mu_{\theta}(\mathbf{S})]} \\ &= \int_{\mathcal{S}} \sum_{t=0}^{\infty} \gamma^t p(\mathbf{S} \rightarrow \hat{\mathbf{S}}, t) \nabla_{\mathbf{A}} Q_{\text{hom},\phi}(\hat{\mathbf{S}}, \hat{\mathbf{A}})|_{\hat{\mathbf{A}}=\sigma_{\mathcal{A}}[\mu_{\theta}(\hat{\mathbf{S}})]} \\ & \cdot \nabla_{\mathbf{A}'} \sigma_{\mathcal{A}}(\mathbf{A}')|_{\mathbf{A}'=\mu_{\theta}(\hat{\mathbf{S}})} \nabla_{\theta} \mu_{\theta}(\hat{\mathbf{S}}) d\hat{\mathbf{S}}. \end{aligned} \quad (36)$$

Now, by using (36), we can obtain

$$\begin{aligned} \nabla_{\theta} J_{\text{hom},\theta} &= \int_{\mathcal{S}} \rho(\mathbf{S}) \nabla_{\mathbf{A}} Q_{\text{hom},\phi}(\mathbf{S}, \mathbf{A})|_{\mathbf{A}=\sigma_{\mathcal{A}}[\mu_{\theta}(\mathbf{S})]} \\ & \nabla_{\mathbf{A}'} \sigma_{\mathcal{A}}(\mathbf{A}')|_{\mathbf{A}'=\mu_{\theta}(\mathbf{S})} \nabla_{\theta} \mu_{\theta}(\mathbf{S}) d\mathbf{S}, \end{aligned} \quad (37)$$

where $\rho(\mathbf{S})$ denotes the discounted state distribution [37]. This completes the proof.

REFERENCES

- [1] C. V. N. Index, "Cisco visual networking index: Global mobile data traffic forecast update, 2016–2021," link: <http://goo.gl/yTuVx>, 2017.
- [2] E. Bastug, M. Bennis, and M. Debbah, "Living on the edge: The role of proactive caching in 5G wireless networks," *IEEE Commun. Mag.*, vol. 52, no. 8, pp. 82–89, Aug. 2014.
- [3] D. Liu, B. Chen, C. Yang, and A. F. Molisch, "Caching at the wireless edge: design aspects, challenges, and future directions," *IEEE Commun. Mag.*, vol. 54, no. 9, pp. 22–28, Sept. 2016.
- [4] J. Li, Y. Chen, Z. Lin, W. Chen, B. Vucetic, and L. Hanzo, "Distributed caching for data dissemination in the downlink of heterogeneous networks," *IEEE Trans. Commun.*, vol. 63, no. 10, pp. 3553–3568, Oct. 2015.
- [5] J. Li, H. Chen, Y. Chen, Z. Lin, B. Vucetic, and L. Hanzo, "Pricing and resource allocation via game theory for a small-cell video caching system," *IEEE J. Sel. Areas Commun.*, vol. 34, no. 8, pp. 2115–2129, Oct., 2016.
- [6] C. Bettstetter, "Mobility modeling in wireless networks: categorization, smooth movement, and border effects," *ACM SIGMOBILE Mob. Comput. Commun. Rev.*, vol. 5, no. 3, pp. 55–66, Jan. 2001.
- [7] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski et al., "Human-level control through deep reinforcement learning," *Nature*, vol. 518, no. 7540, p. 529, Feb. 2015.

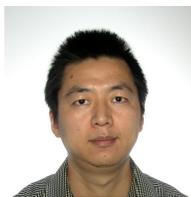
- [8] N. C. Luong, D. T. Hoang, S. Gong, D. Niyato, P. Wang, Y.-C. Liang, and D. I. Kim, "Applications of deep reinforcement learning in communications and networking: A survey," *IEEE Commun. Surveys Tuts.*, vol. 21, no. 4, pp. 3133–3174, Fourthquarter 2019.
- [9] A. Sadeghi, G. Wang, and G. B. Giannakis, "Deep reinforcement learning for adaptive caching in hierarchical content delivery networks," *IEEE Trans. Cogn. Commun. Netw.*, vol. 5, no. 4, pp. 1024–1033, Dec. 2019.
- [10] B. Blaszczyszyn and A. Giovanidis, "Optimal geographic caching in cellular networks," in *Proc. IEEE Int. Conf. Commun.*, June 2015, pp. 3358–3363.
- [11] X. Li, X. Wang, S. Xiao, and V. C. Leung, "Delay performance analysis of cooperative cell caching in future mobile networks," in *Proc. IEEE Int. Conf. Commun.*, June 2015, pp. 1–6.
- [12] X. Li, X. Wang, and V. C. Leung, "Weighted network traffic offloading in cache-enabled heterogeneous networks," in *Proc. IEEE Int. Conf. Commun.*, May 2016, pp. 1–6.
- [13] M. Tao, E. Chen, H. Zhou, and W. Yu, "Content-centric sparse multicast beamforming for cache-enabled cloud RAN," *IEEE Trans. Wireless Commun.*, vol. 15, no. 9, pp. 6118–6131, Sept. 2016.
- [14] X. Wu, X. Li, Q. Li, V. C. Leung, and P. Ching, "Latency driven fronthaul bandwidth allocation and cooperative beamforming for cache-enabled cloud-based small cell networks," in *Proc. IEEE Int. Conf. Acoust., Speech, and Signal Process.*, May, 2019, pp. 4594–4598.
- [15] M. A. Maddah-Ali and U. Niesen, "Fundamental limits of caching," *IEEE Trans. Inf. Theory*, vol. 60, no. 5, pp. 2856–2867, May 2014.
- [16] V. Bioglio, F. Gabry, and I. Land, "Optimizing MDS codes for caching at the edge," in *Proc. IEEE Global Commun. Conf.*, Dec. 2015, pp. 1–6.
- [17] J. Liao, K.-K. Wong, Y. Zhang, Z. Zheng, and K. Yang, "Coding, multicast, and cooperation for cache-enabled heterogeneous small cell networks," *IEEE Trans. Wireless Commun.*, vol. 16, no. 10, pp. 6838–6853, Oct. 2017.
- [18] X. Wu, Q. Li, X. Li, V. C. Leung, and P. Ching, "Joint long-term cache updating and short-term content delivery in cloud-based small cell networks," *IEEE Trans. Commun.*, vol. 68, no. 5, pp. 3173 – 3186, May 2020.
- [19] X. Wu, Q. Li, V. C. Leung, and P. Ching, "Joint fronthaul multicast and cooperative beamforming for cache-enabled cloud-based small cell networks: An MDS codes-aided approach," *IEEE Trans. Wireless Commun.*, vol. 18, no. 10, pp. 4970–4982, Oct. 2019.
- [20] A. Sadeghi, F. Sheikholeslami, and G. B. Giannakis, "Optimal and scalable caching for 5G using reinforcement learning of space-time popularities," *IEEE J. Sel. Topics Signal Process.*, vol. 12, no. 1, pp. 180–190, Feb. 2018.
- [21] Y. He, Z. Zhang, F. R. Yu, N. Zhao, H. Yin, V. C. Leung, and Y. Zhang, "Deep-reinforcement-learning-based optimization for cache-enabled opportunistic interference alignment wireless networks," *IEEE Trans. Veh. Technol.*, vol. 66, no. 11, pp. 10433–10445, Nov. 2017.
- [22] Y. Wei, F. R. Yu, M. Song, and Z. Han, "Joint optimization of caching, computing, and radio resources for fog-enabled iot using natural actor-critic deep reinforcement learning," *IEEE Internet Things J.*, Apr. 2018.
- [23] P. Wu, J. Li, L. Shi, M. Ding, K. Cai, and F. Yang, "Dynamic content update for wireless edge caching via deep reinforcement learning," *IEEE Commun. Lett.*, vol. 23, no. 10, pp. 1773–1777, Oct. 2019.
- [24] C. Zhong, M. C. Gursoy, and S. Velipasalar, "Deep reinforcement learning based edge caching in wireless networks," *IEEE Trans. Cogn. Commun. Netw.*, Mar. 2020.
- [25] X. Xu, M. Tao, and C. Shen, "Collaborative multi-agent multi-armed bandit learning for small-cell caching," *arXiv preprint arXiv:2001.03835*, 2020.
- [26] X. Wang, C. Wang, X. Li, V. C. Leung, and T. Taleb, "Federated deep reinforcement learning for internet of things with decentralized cooperative edge caching," *IEEE Internet Things J.*, Apr., 2020.
- [27] S. Gao, P. Dong, Z. Pan, and G. Y. Li, "Reinforcement learning based cooperative coded caching under dynamic popularities in ultra-dense networks," *IEEE Trans. Veh. Technol.*, vol. 69, no. 5, pp. 5442–5456, May 2020.
- [28] Z. Zhang and M. Tao, "Accelerated deep reinforcement learning for wireless coded caching," in *Proc. IEEE/CIC Int. Conf. Commun. in China*, Sept. 2019, pp. 249–254.
- [29] D. M. Dunlavy and D. P. O'Leary, "Homotopy optimization methods for global optimization," *Sandia Nat. Lab., Albuquerque, NM, USA, Tech. Rep. SAND2005-7495*, 2005., 2005.
- [30] B. Bharath, K. G. Nagananda, and H. V. Poor, "A learning-based approach to caching in heterogenous small cell networks," *IEEE Trans. Commun.*, vol. 64, no. 4, pp. 1674–1686, Apr. 2016.
- [31] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. M. O. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra, "Continuous control with deep reinforcement learning," *CoRR*, vol. abs/1509.02971, 2015.
- [32] R. S. Sutton, A. G. Barto *et al.*, *Introduction to Reinforcement Learning*. MIT Press Cambridge, 1998, vol. 135.
- [33] A. Cassioli, D. Di Lorenzo, and M. Sciandrone, "On the convergence of inexact block coordinate descent methods for constrained optimization," *Eur. J. Oper. Res.*, vol. 231, no. 2, pp. 274–281, 2013.
- [34] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, 2015.
- [35] R. Lowe, Y. Wu, A. Tamar, J. Harb, O. P. Abbeel, and I. Mordatch, "Multi-agent actor-critic for mixed cooperative-competitive environments," in *Adv. Neural Inf. Process. Syst.*, 2017, pp. 6379–6390.
- [36] P. Mishra and K. Sarawadekar, "Polynomial learning rate policy with warm restart for deep neural network," in *IEEE TENCON*, Oct. 2019, pp. 2087–2092.
- [37] D. Silver, G. Lever, N. Heess, T. Degris, D. Wierstra, and M. Riedmiller, "Deterministic policy gradient algorithms," in *Proc. Int. Conf. Machine Learning*, June 2014, pp. 387–395.



Xiongwei Wu (S'18) received the B.Eng. in Electronic Information Engineering from University of Electronic Science and Technology of China, Chengdu, China, in 2016. He is currently working toward the Ph.D. degree with the Chinese University of Hong Kong (CUHK), Shatin, Hong Kong SAR, China. From August 2018 to December 2018, he was a Visiting International Research Student with the University of British Columbia (UBC), Vancouver, BC, Canada. From July 2019 to January 2020, he was a Visiting Student Research Collaborator with Princeton University, Princeton, NJ, USA. His research interests include signal processing and resource allocation in wireless networks, decentralized optimization, and machine learning.



Jun Li (M'09-SM'16) received Ph. D degree in Electronic Engineering from Shanghai Jiao Tong University, Shanghai, P. R. China in 2009. From January 2009 to June 2009, he worked in the Department of Research and Innovation, Alcatel Lucent Shanghai Bell as a Research Scientist. From June 2009 to April 2012, he was a Postdoctoral Fellow at the School of Electrical Engineering and Telecommunications, the University of New South Wales, Australia. From April 2012 to June 2015, he was a Research Fellow at the School of Electrical Engineering, the University of Sydney, Australia. From June 2015 to now, he is a Professor at the School of Electronic and Optical Engineering, Nanjing University of Science and Technology, Nanjing, China. He was a visiting professor at Princeton University from 2018 to 2019. His research interests include network information theory, game theory, distributed intelligence, multiple agent reinforcement learning, and their applications in ultra-dense wireless networks, mobile edge computing, network privacy and security, and industrial Internet of things. He has co-authored more than 200 papers in IEEE journals and conferences, and holds 1 US patents and more than 10 Chinese patents in these areas. He was serving as an editor of IEEE Communication Letters and TPC member for several flagship IEEE conferences. He received Exemplary Reviewer of IEEE Transactions on Communications in 2018, and best paper award from IEEE International Conference on 5G for Future Wireless Networks in 2017.



Ming Xiao (S'2002-M'2007-SM'2012) received Bachelor and Master degrees in Engineering from the University of Electronic Science and Technology of China, Chengdu in 1997 and 2002, respectively. He received Ph.D degree from Chalmers University of technology, Sweden in November 2007. From 1997 to 1999, he worked as a network and software engineer in ChinaTelecom. From 2000 to 2002, he also held a position in the SiChuan communications administration. From November 2007 to now, he has been in the department of information science

and engineering, school of electrical engineering and computer science, Royal Institute of Technology, Sweden, where he is currently an Associate Professor. He was an Editor for IEEE Transactions on Communications (2012-2017), IEEE Communications Letters (Senior Editor Since January 2015) and IEEE Wireless Communications Letters (2012-2016), and has been an Editor for IEEE Transactions on Wireless Communications since 2018. He was the lead Guest Editor for IEEE JSAC Special issue on Millimeter Wave Communications for future mobile networks in 2017. He has been an area editor for IEEE Open Journal of the Communication Society since 2019.



H. Vincent Poor (S'72, M'77, SM'82, F'87) received the Ph.D. degree in EECS from Princeton University in 1977. From 1977 until 1990, he was on the faculty of the University of Illinois at Urbana-Champaign. Since 1990 he has been on the faculty at Princeton, where he is currently the Michael Henry Strater University Professor. During 2006 to 2016, he served as the dean of Princetons School of Engineering and Applied Science. He has also held visiting appointments at several other universities, including most recently at Berkeley and Cambridge.

His research interests are in the areas of information theory, machine learning and network science, and their applications in wireless networks, energy systems and related fields. Among his publications in these areas is the forthcoming book *Machine Learning and Wireless Communications*. (Cambridge University Press, 2021).

Dr. Poor is a member of the National Academy of Engineering and the National Academy of Sciences and is a foreign member of the Chinese Academy of Sciences, the Royal Society, and other national and international academies. Recent recognition of his work includes the 2017 IEEE Alexander Graham Bell Medal and a D.Eng. honoris causa from the University of Waterloo awarded in 2019.



P. C. Ching (F'11) received the B. Eng. (1st Class Honors) and Ph.D. degrees from the University of Liverpool, UK, in 1977 and 1981 respectively. From 1981 to 1982, he was Research Officer at the University of Bath, UK. In 1982, Prof. Ching returned to Hong Kong and joined the then Hong Kong Polytechnic as a lecturer. Since 1984, he has been with the Department of Electronic Engineering of the Chinese University of Hong Kong (CUHK), where he is currently Choh-Ming Li Professor of Electronic Engineering. He was Department Chair-

man from 1995 to 1997, Dean of Engineering from 1997 to 2003 and Head of Shaw College from 2004 to 2008. He became Director of the Shun Hing Institute of Advanced Engineering in 2004. From 2006 till end of 2014, Prof Ching was appointed as Pro-Vice-Chancellor/Vice-President of CUHK. Between 2013 to 2014, Prof. Ching also took up the Directorship of the CUHK Shenzhen Research Institute. Prof. Ching is very active in promoting professional activities, both in Hong Kong and overseas. He was a council member of the Institution of Electrical Engineers (IEE), past chairman of the IEEE Hong Kong Section, an associate editor of the IEEE Transactions on Signal Processing from 1997 to 2000 and IEEE Signal Processing Letters from 2001 to 2003. He was also a member of the Technical Committee of the IEEE Signal Processing Society from 1996 to 2004. He was appointed Editor-in-Chief of the HKIE Transactions between 2001 and 2004. He has been an Honorary Member of the editorial committee for Journal of Data Acquisition and Processing since 2000. Prof. Ching has been instrumental in organizing many international conferences in Hong Kong including the 1997 IEEE International Symposium on Circuits and Systems where he was the Vice-Chairman. He also served as Technical Program Co-Chair of the 2003 and 2016 IEEE International Conference on Acoustics, Speech and Signal Processing. Prof Ching was awarded the IEEE Third Millennium Award in 2000 and the HKIE Hall of Fame in 2010. In addition, Prof. Ching also plays an active role in community services. He was awarded the Silver Bauhinia Star (SBS) and the Bronze Bauhinia Star (BBS) by the HKSAR Government in 2017 and 2010, respectively, in recognition of his long and distinguished public and community services. He is presently Chairman of the Board of Directors of the Nano and Advanced Materials Institute Limited, Council Member of the Shaw Prize Foundation, as well as a Member of the Museum Advisory Committee (MAC) and the Chairperson of its Science Sub-committee. He is elected as President of the Hong Kong Academy of Engineering Sciences (HKAES) in 2018.