






User Dynamics-Aware Edge Caching and Computing for Mobile Virtual Reality

Mushu Li , Member, IEEE, Jie Gao , Senior Member, IEEE, Conghao Zhou , Member, IEEE, Xuemin Shen , Fellow, IEEE, and Weihua Zhuang , Fellow, IEEE

Abstract—In this article, we present a novel content caching and delivery approach for mobile virtual reality (VR) video streaming. The proposed approach aims to maximize VR video streaming performance, i.e., minimizing video frame missing rate, by proactively caching popular VR video chunks and adaptively scheduling computing resources at an edge server based on user and network dynamics. First, we design a scalable content placement scheme for deciding which video chunks to cache at the edge server based on tradeoffs between computing and caching resource consumption. Second, we propose a machine learning-assisted VR video delivery scheme, which allocates computing resources at the edge server to satisfy video delivery requests from multiple VR headsets. A Whittle index-based method is adopted to reduce the video frame missing rate by identifying network and user dynamics with low signaling overhead. Simulation results demonstrate that the proposed approach can significantly improve VR video streaming performance over conventional caching and computing resource scheduling strategies.

Index Terms—Virtual reality, deep reinforcement learning, caching, content delivery, resource scheduling.

I. INTRODUCTION

As a promising use case in sixth-generation (6G) communication networks, mobile virtual reality (VR) is anticipated to reshape the way people study, work, and entertain via the digital transformation of the physical world [1]. A myriad of novel applications can be enabled by mobile VR, such as immersive gaming, telesurgery, metaverses, etc. [2]. With mobile VR, users can watch 360° stereoscopic videos through VR headsets and access the digital world with a fully immersive experience.

Delivering VR videos to mobile VR headsets in real time can be challenging. VR videos are characterized by their ultra-high resolutions (up to 11520×6480 pixels for panoramic videos), while conventional videos typically have a resolution of 4K or

less [3]. From a communication perspective, VR video streaming consumes considerable bandwidth on both backhaul and wireless links for delivering high-resolution videos to VR headsets. In addition, VR headsets play 360° stereoscopic videos, i.e., two different videos for the left and right eyes, respectively, creating the perception of depth that people experience in the real world. To reduce storage resource usage, monoscopic videos, which consist of flat two-dimensional equirectangular video frames, can be stored at the cloud server and projected into stereoscopic videos in content delivery [4]. Extensive VR video processing requires powerful computing capabilities, and computing capabilities in VR headsets are usually insufficient for low-latency video processing [5].

The aforementioned challenges in VR video delivery can be addressed using two main approaches: tile-based content delivery and mobile edge computing. First, in tile-based content delivery, VR videos are divided spatio-temporally into small video chunks (VCs). Only the VCs that cover the user's current viewing area, i.e., the field of view (FoV), are sent to and played on the VR headset [6]. Such a tile-based solution can significantly reduce the data sizes of videos to be downloaded from the cloud server while satisfying users' requests. Second, mobile edge computing (MEC), as an innovation in the fifth-generation (5G) networks, leverages storage and computing resources at access points near VR headsets to reduce content delivery delay. Edge servers are located at access points to cache popular VCs and project monoscopic VCs (MVCs) into stereoscopic VCs (SVCs) for resource-limited VR headsets [7]. While MEC and tile-based content delivery make VR video streaming possible, the following questions need to be answered:

- Given limited storage space, which VCs should be cached at an edge server?
- How should computing resources at the edge server be scheduled to meet video delivery requests from multiple VR headsets?

Addressing the first question necessitates the effective selection of VCs cached at an edge server, referred to as content placement. Determining the optimal caching solution can be complicated [8]. The edge server can cache MVCs, which have smaller data sizes than the corresponding SVCs covering the same spatial area, to reduce caching resource usage given limited cache size. Alternatively, the edge server can cache SVCs, which are processed VCs and ready to send to VR headsets, to reduce video processing time in content delivery [9]. Furthermore, spatio-temporally partitioning a short video can

Manuscript received 15 January 2023; accepted 1 May 2023. Date of publication 16 May 2023; date of current version 16 November 2023. This work was supported by the Natural Sciences and Engineering Research Council (NSERC) of Canada. The guest editor coordinating the review of this manuscript and approving it for publication was Dr. Jing Li. (Corresponding author: Mushu Li.)

Mushu Li is with the Department of Electrical, Computer and Biomedical Engineering, Toronto Metropolitan University, Toronto M5B 2K3, ON, Canada (e-mail: mushu.li@iee.org).

Jie Gao is with the School of Information Technology, Carleton University, Ottawa K1S 5B6, ON, Canada (e-mail: jie.gao6@carleton.ca).

Conghao Zhou, Xuemin Shen, and Weihua Zhuang are with the Department of Electrical and Computer Engineering, University of Waterloo, Waterloo N2L 3G1, ON, Canada (e-mail: c89zhou@uwaterloo.ca; sshen@uwaterloo.ca; wzhuang@uwaterloo.ca).

Digital Object Identifier 10.1109/JSTSP.2023.3276595

produce a large number of VCs. These VCs may be spatially overlapped and stitched together to produce the video covering a user's FoV [6]. The vast number of potential caching options and coupling among VCs complicate content placement at the edge server.

Regarding the second question, managing limited computing resources for satisfying multiple content delivery requests is non-trivial due to the dynamics in VR video streaming. Millimeter wave (mmWave) communications are widely adopted to enable high-speed VR content delivery in many research works [10]. However, physical obstacles in the propagation environment, including the user's body, can disrupt communication links, resulting in a time-varying transmission rate. In addition, proactive content delivery can facilitate smooth VR video streaming by requesting and downloading VCs in advance based on predicted user viewing trajectories [11]. The spatial area in a video displayed to a user changes over time as the user's head orientation changes, and the viewing trajectories can be different for different videos depending on the content (e.g., sports or movies) and for different users due to the user viewing preference. The dynamic of user viewing trajectories introduces additional uncertainty in computing resource management at the edge server [12]. Allocating computing and communication resources for satisfying VR content delivery requests based on inaccurate viewing trajectory predictions may degrade resource utilization and result in a poor VR video streaming experience for all nearby VR headsets sharing the resources.

In this article, we study content caching and delivery for mobile VR video streaming. The research objective is to optimize VR video streaming performance, i.e., reducing the frame missing rate, for multiple VR users given limited communication, caching, and computing resources. First, we propose a content placement scheme to determine the optimal caching solution, given a content delivery delay threshold, channel quality and processing rate statistics, and VC data sizes. The proposed heuristic scheme uses a distributed optimization technique, i.e., alternating direction method of multipliers (ADMM), to facilitate parallel content placement. Using the caching solution, we propose a machine learning-based scheduling scheme to allocate limited computing resources at the edge server to reduce the frame missing rate in satisfying content delivery requests generated by multiple VR headsets. Finally, we study adaptive video quality adjustment to maximize the cached video quality at the edge server. This study covers the main aspects of VR video streaming, and our contributions are summarized as follows:

- We develop a caching placement and computing resource management approach to effectively improve video streaming performance for mobile VR.
- We propose a content placement scheme for caching VCs by trading off caching resource utilization and content delivery delay. The proposed scheme is proven to achieve a near-optimal caching solution.
- We propose a deep reinforcement learning (DRL)-based content delivery scheduling scheme for computing resource management. The scheme minimizes the frame missing rate in video streaming by adapting to network and user dynamics in a distributed manner.

The remainder of this article is organized as follows. Section II introduces related works. In Sections III and IV, we present the system model and problem formulation, respectively. In Sections V and VI, we propose content caching and content delivery scheduling schemes, respectively. We then discuss adaptive video quality adjustment in Section VI. Simulation results are presented in Section VII. Section VIII concludes this study.

II. RELATED WORKS

MEC can reduce content delivery and processing delay in VR video streaming by providing storage and computing capabilities near mobile VR headsets. An edge server can cache both MVCs and SVCs. A caching solution was evaluated in [9] to minimize average content delivery delay based on the tradeoffs among communication, computing, and caching resource usage. A similar problem was investigated in [13] considering caching capabilities at VR headsets. A mixed-integer problem was formulated and solved to determine whether and where to cache MVCs and SVCs. Additionally, some works take into account the correlations among VCs in order to reduce resource usage in caching and delivering videos covering users' FoVs, and the correlations vary depending on the granularity of partitioning a VR video. When a VR video is spatially partitioned into small segments, multiple VCs need to be stitched together to generate the video covering a user's FoV. In such a case, a tile-based content placement solution was proposed in [14]. A Choquet integral method was applied to identify the popularity of individual VCs based on user viewing behavior statistics, and popular VCs were cached at the edge server. The estimation of VC popularity was also studied in [15] and [16] for caching popular and high-quality VCs, while minimizing content delivery delay. When a video is coarsely partitioned in the spatial dimension, a single VC can cover a user's FoV, whereas multiple VCs may cover overlapped spatial areas. A study of content placement in such a case was presented in [6], where a group of cached VCs are synthesized to generate VCs that are not cached yet. The objective was to minimize content delivery delay, while maximizing caching resource utilization by properly selecting VCs to cache. From the existing works, VR content placement is a complex problem with different types of VCs (different quality levels, monoscopic or stereoscopic, etc.), and a scalable content placement solution is necessary for effectively caching VCs given a large number of potential caching options.

In addition to content placement, real-time content delivery scheduling and computing resource management can further improve video streaming performance by maximizing computing resource utilization at the edge server. To prevent users from feeling spatially disoriented and dizzy in VR video streaming, the time difference between the user's movement and corresponding display at the VR headset, referred to as motion-to-photon (MTP) delay, should be less than 20 ms [17]. In order to meet this strict requirement, a VR headset can predict user viewpoints based on historical viewpoint trajectories and request the VCs proactively [18], [19]. As a result, the VR video streaming experience depends on viewpoint prediction accuracy. In [20],




Fig. 1. Network model.

the viewpoint prediction window was determined to maximize users' VR experience, and computing resource allocation was jointly determined. Furthermore, machine learning algorithms were adopted in [21] and [22] to determine computing resource scheduling policies, given user viewpoint trajectories. In spite of the research efforts, designing a lightweight scheduling scheme remains an open issue when network dynamics and viewpoint prediction uncertainty are taken into account [12].

III. SYSTEM MODEL

The considered VR video caching and delivery scenario is shown in Fig. 1. An edge server, located at the access point, connects and serves U mobile VR users. Each user has a VR headset that plays 360° 3D stereoscopic VR videos at \hat{f} frames per second (fps). We refer to the time interval between the playback of two adjacent frames as a time slot, and the length of a time slot is $\delta = 1/\hat{f}$. In each time slot, each VR headset predicts the location of future user viewpoints, generates a content delivery request for the VC to be played, and sends the request to the edge server. Video processing, including stitching multiple VCs and projecting 2D monoscopic videos to 3D stereoscopic videos, can consume extensive computing resources and is conducted at the edge server.

The edge server is equipped with E computing units, and each computing unit can conduct video processing for only one content delivery request at any time. Moreover, the edge server has a cache with size C for storing popular VCs. For retrieving VCs that are not cached, the edge server connects to the cloud server via a wired link with data rate R_B . When transmitting the requested VC from the edge server to the corresponding VR headset, a mmWave band is used with high-speed rate R_H , and a sub-6 GHz band acts as a backup with low-speed rate R_L if the high-speed mmWave band is in outage [7]. We use a two-stage Markov chain to model the data rate between the edge server and a VR headset. The probabilities of the high-speed state transiting from and to the low-speed state are denoted by p_L and p_H , respectively. As a result, the average data rate from the edge server to a VR headset is $\bar{R}_E = (R_L + R_H)/(p_L + p_H)$.

A. VR Video Model

Tile-based content caching and delivery are considered. As shown in Fig. 2, each VR video has two dimensions. In the spatial dimension, a video can be divided into I tiles, and we denote the

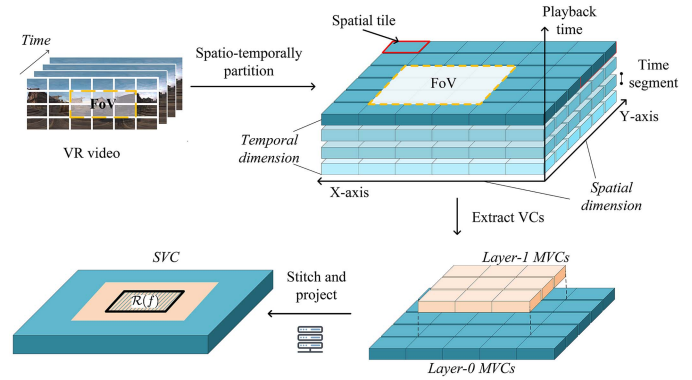


Fig. 2. VR video model.

number of tiles needed to cover an FoV by \hat{I} . In the temporal dimension, videos can be divided into J time segments with a constant duration in seconds¹ [23]. Accordingly, VR videos are divided into VCs along both dimensions. In the temporal dimension, the playback time of a VC is the duration of a time segment.

In the spatial dimension, there are two types of VCs: MVCs and SVCs. An MVC contains the monographic content of a spatial tile. For MVCs, layered encoding is adopted to differentiate the video quality [3]. There are two types of MVCs: layer-0 MVCs, which have a low resolution to satisfy the minimum video quality requirement, and layer-1 MVCs, which can be added on top of layer-0 MVCs for quality improvement. We denote the sets of layer-0 and layer-1 MVCs for all VR videos by \mathcal{T}_0 and \mathcal{T}_1 , respectively. Tuple $t = (i, j, y)$ represents the MVC corresponding to the i -th tile, j -th time segment, and y -th layer, where $i \in \{1, \dots, I\}$, $j \in \{1, \dots, J\}$, and $y \in \{1, 0\}$.

There are X MVCs, including both layer-0 and layer-1 MVCs, to be stitched and projected to an SVC, and the projection is facilitated by using auxiliary files (e.g., depth map). SVCs are delivered by the edge server to mobile VR headsets. To obtain the SVC centering on tile i , X_0 layer-0 MVCs and $X - X_0$ layer-1 MVCs, corresponding to the tiles closest to tile i , are stitched and projected. In the spatial dimension, the area covered by an SVC depends on the parameter X_0 , which must be at least \hat{I} to ensure that an SVC can fully cover the area of an FoV. The video quality of an SVC can be enhanced by encoding a large portion of layer-1 MVCs [24], [25]. On the other hand, a VR headset can play the videos covering multiple FoVs when an SVC is encoded with a large portion of layer-0 MVCs (corresponding to a large X_0), enabling the VR headset to adapt to frequent viewpoint changes [26].

Denote the set of all SVCs by \mathcal{T}_S , and denote the sets of layer-0 and layer-1 MVCs for generating SVC f by $\mathcal{B}(f, X_0)$ and $\mathcal{E}(f, X_0)$, respectively, where $\mathcal{B}(f, X_0) \subseteq \mathcal{T}_0$ and $\mathcal{E}(f, X_0) \subseteq \mathcal{T}_1$. Since an SVC contains videos for both eyes, the overall data size of all MVCs used to generate an SVC is smaller than the data size of the SVC. Therefore, caching SVCs consumes more cache spaces at the edge server but decreases video processing

¹As time slots are measured in milliseconds (e.g., for a 60 fps video, a time slot lasts 17 milliseconds), a time segment contains hundreds of time slots.




Fig. 3. Content delivery process diagram.

delay when delivering SVCs to VR headsets. In the considered scenario, the cloud server stores all MVCs, and the edge server can cache some MVCs and SVCs to reduce content delivery delay. The sets of MVCs and SVCs cached at the edge server are denoted by \mathcal{M} and \mathcal{S} , respectively.

B. Content Delivery Model

The VR content delivery process diagram is illustrated in Fig. 3. Each VR headset has a local buffer to store the downloaded SVCs. In each time slot, a VR headset tracks the user viewpoint and plays the SVC covering the user's FoV. Denote the viewpoint of user u in time slot k by $v_{u,k} = (i_{u,k}, j_{u,k})$, where $i_{u,k}$ and $j_{u,k}$ represent the spatial tile at the center of user u 's FoV in time slot k and the time segment of the video played by the VR headset in time slot k , respectively. When an SVC covers a larger area than an FoV, i.e., $X_0 > \hat{I}$, multiple viewpoints can be rendered by an SVC. Denote the set of viewpoints that can be rendered by SVC f by \mathcal{R}_f . In time slot k , if SVC f in the buffer of VR headset u can render the current viewpoint, i.e., $v_{u,k} \in \mathcal{R}_f$, the SVC is fetched and played by the VR headset. Otherwise, if no SVC in the buffer can render the viewpoint, the VR headset fails to play the correct video frame, and a frame missing event occurs. In addition, to identify which SVC will be requested proactively, in each time slot, each VR headset consecutively predicts user viewpoints to be rendered after the current slot, until it finds the next future viewpoint that cannot be rendered by any SVCs in the current local buffer. Such a predicted viewpoint is referred to as the desired viewpoint. Then, the VR headset generates a VR content delivery request proactively for obtaining the SVC to render the desired viewpoint. Denote the predicted viewpoint of time slot k for VR headset u by $\hat{v}_{u,k}$. The content delivery request from VR headset u in time slot k is denoted by $d_{u,k} = \hat{v}_{u,k'} |_{k' > k}$, where k' is the time slot to render desired viewpoint $\hat{v}_{u,k'}$. Since the historical information of the past viewpoint trajectories is updated after each time slot, a VR headset conducts predictions for future viewpoints in each time slot. Only the prediction result obtained in the current time slot is used to generate the content delivery request.

The edge server receives content delivery requests from multiple VR headsets in each time slot. A scheduler at the edge server identifies which request to respond first when a computing unit is available. If a computing unit is allocated to a request, i.e., the request is scheduled, there are two possible scenarios for fetching and processing the VCs to satisfy the request:

- 1) *Case 1*: The edge server caches an SVC that can satisfy the request, i.e., $\{\exists f | f \in \mathcal{S}, d_{u,k} \in \mathcal{R}_f\}$. In this case, the SVC is fetched from the cache and delivered to the corresponding VR headset².
- 2) *Case 2*: No SVC cached on the edge server can satisfy the request. In this case, the edge server fetches MVCs in $\mathcal{B}(f, X_0)$ and $\mathcal{E}(f, X_0)$ from the cache or the cloud server, processes the MVCs for generating SVC f , and delivers SVC f to the corresponding VR headset.

Regardless of whether the request from a VR headset is scheduled or not, the VR headset tracks and predicts the viewpoints in each time slot. If the viewpoint prediction result changes before the request is scheduled, the VR headset sends an updated content delivery request to the edge server, and the previously unscheduled request becomes obsolete and discarded.

C. Delay Model

A content delivery delay is the duration from the instant that a request is scheduled to the instant that the SVC rendering the desired viewpoint is downloaded to the VR headset. The delay includes three parts: the transmission delay for downloading MVCs that are not cached from the cloud server to the edge server, the computing delay for stitching and processing the MVCs, and the transmission delay for downloading the SVC from the edge server to the VR headset. When the request from VR headset u in time slot k , i.e., request $d_{u,k}$, is scheduled, a computing unit is allocated to process the request and obtain the SVC. The delay for transmitting the SVC from the edge server to the VR headset is

$$T^E(d_{u,k}; R) = w_f^S R^{-1}, d_{u,k} \in \mathcal{R}_f \quad (1)$$

where the value of R can be either R_L or R_H . Parameter w_f^S represents the data size of SVC f , which can be approximated by the overall data size of MVCs in set $\mathcal{C}(f) = \mathcal{B}(f, X_0) \cup \mathcal{E}(f, X_0)$. Let w_t^M denote the data size of MVC t . Then,

$$w_f^S = \sum_{t \in \mathcal{C}(f)} \alpha w_t^M \quad (2)$$

where parameter α represents the ratio of data size change from MVCs to the corresponding SVC. If no cached SVC can satisfy the request, i.e., Case 2 in Subsection III-B, the edge server needs to generate SVC f from MVCs. Firstly, the MVCs not in the cache are downloaded from the cloud server. The transmission delay is

$$T^B(d_{u,k}) = \sum_{t \in \{\mathcal{C}(f)\} \setminus \{\mathcal{C}(f) \cap \mathcal{M}\}} w_t^M R_B^{-1}, d_{u,k} \in \mathcal{R}_f. \quad (3)$$

When all MVCs in set $\mathcal{C}(f)$ are available at the edge server, the MVCs are stitched and projected to SVC f . The corresponding computing delay is

$$T^C(d_{u,k}) = \sum_{t \in \mathcal{C}(f)} w_t^M \chi^{-1}, d_{u,k} \in \mathcal{R}_f \quad (4)$$

²In Case 1, a computing unit is still needed to process the SVC to satisfy the video requirements of the corresponding VR headset [27], whereas the processing delay is negligible.

where parameter χ denotes the time to process each data bit in MVC stitching and monoscopic-to-stereoscopic projection. Therefore, the overall delay of content delivery for request $d_{u,k}$ is

$$T(d_{u,k}) = T^B(d_{u,k}) + T^C(d_{u,k}) + T^E(d_{u,k}, R). \quad (5)$$

IV. PROBLEM FORMULATION

Our main objective is to minimize the frame-missing rate when playing VR videos on multiple VR headsets. Achieving the objective requires that the edge server efficiently caches VCs to reduce the average content delivery delay and dynamically allocate computing resources to individual content delivery requests in real time. Therefore, we formulate two optimization problems for caching content placement and content delivery scheduling, respectively.

A. Content Placement

The objective of content placement is to maximize the probability that content delivery delay $T(d_{u,k})$ is lower than the delay requirement H , i.e., $P(T(d_{u,k}) < H)$, by selecting proper VCs to cache. The corresponding problem can be formulated as:

$$\max_{\{\mathcal{M}, \mathcal{S}\}} \sum_{d \in \mathcal{V}} p(d) P(T(d) < H) \quad (6a)$$

$$\text{s.t.} \quad \sum_{t \in \{\mathcal{M}\}} w_t^M + \sum_{f \in \mathcal{S}} w_f^S \leq C, \quad (6b)$$

$$\mathcal{M} \subseteq \{\mathcal{T}_0 \cup \mathcal{T}_1\}, \mathcal{S} \subseteq \mathcal{T}_S \quad (6c)$$

where set \mathcal{V} denotes all possible viewpoints. In (6a), $p(d)$ represents the probability that request d is generated by a VR headset. Such probability can be obtained from historical viewpoint popularity profiles, which reflect the average fraction of time that a user viewpoint falls onto a specific tile in a specific time segment [28]. The problem is a non-linear combinatorial optimization problem since multiple viewpoints may be rendered by an SVC. It is similar to the maximal coverage problem [29], which is proved to be NP-hard. Furthermore, the number of VCs is large, which makes the problem more complex. We propose a two-stage solution for solving the content caching problem. We first present a heuristic algorithm to solve (6) given a fixed cache size. Then, we adopt an iterative optimization method to enable parallel and scalable content placement.

B. Content Delivery

Due to limited computing resources, the scheduler at the edge server should allocate the computing resource to satisfy content delivery requests that can reduce the frame missing rate as much as possible. We refer to the probability that an SVC in the buffer matches the current viewpoint as the hit probability. A higher hit probability means a lower frame missing rate at a VR headset. Define scheduling variable $a_{u,k}$, where $a_{u,k} = 1$ if the request from VR headset u is scheduled in time slot k , and $a_{u,k} = 0$ otherwise. The scheduling variables are determined to maximize the hit probability, i.e., minimize the time-averaged number of events where SVCs in the local buffers of VR headsets can

render the viewpoints of the corresponding users. The objective function is

$$\max_{a_{u,k}, \forall u,k} \lim_{K \rightarrow \infty} \frac{1}{K} \mathbb{E} \left[\sum_{u=1}^U \sum_{k=1}^K \mathbf{1}[v_{u,k} \in \mathcal{U}_{u,k}(\mathbf{a}_{u,k-1})] \right] \quad (7)$$

where $\mathcal{U}_{u,k}(\mathbf{a}_{u,k-1})$ denotes the set of viewpoints that can be rendered by SVCs in VR headset u 's buffer in time slot k . Such a set depends on the scheduling decisions $\mathbf{a}_{u,k-1} = [a_{u,1}, \dots, a_{u,k-1}]$. Function $\mathbf{1}[x]$ is equal to one if x is true and zero otherwise. Due to the limited number of computing units, we formulate a stochastic constraint:

$$\lim_{K \rightarrow \infty} \frac{1}{K} \mathbb{E} \left[\sum_{u=1}^U \sum_{k=1}^K a_{u,k} \right] \leq \frac{E \times \delta}{\bar{T}} \quad (8)$$

where \bar{T} represents the average content delivery delay, i.e., $\bar{T} = \sum_{d \in \mathcal{V}} p(d) T(d)$. The average number of time slots for delivering an SVC is \bar{T}/δ . The constraint in (8) is a relaxed constraint imposing that, on average, at most $(E \times \delta)/\bar{T}$ computing units can be scheduled in a time slot.

Several factors affect the scheduling decision. Generally, the edge server should first schedule the request that needs to be satisfied urgently. However, if the channel quality is poor or receiving the SVC from the edge server takes a long time, the VR headset may not receive the requested SVC on time. In addition, high viewpoint movement dynamics pose a challenge to accurate viewpoint prediction. As a result, downloaded SVCs may not be played, resulting in a waste of computing and communication resources. Both network and user viewpoint movement dynamics should be taken into account in computing resource management.

V. MOBILE VR CONTENT CACHING

In this section, we first propose a content placement scheme to solve Problem (6) given the predefined video quality setting X_0 , while taking into account the trade-offs among multiple resources. Then, we extend the scheme for caching a large number of VCs and solve Problem (6) in a parallel manner.

A. Content Placement Scheme

We first identify the properties of content placement for different types of VCs. Let $\mathcal{F} = \mathcal{M} \cup \mathcal{S}$. Let function $L(\mathcal{F})$ represent the overall probability of delay requirement satisfaction as follows:

$$L(\mathcal{F}) = \sum_{d \in \mathcal{R}(\mathcal{F})} p(d) P(T(d) < H) \quad (9)$$

where $\mathcal{R}(\mathcal{F})$ represents the set of viewpoints that can be rendered by the VCs in set \mathcal{F} , and $\mathcal{R}(\mathcal{F}) = \{\cup_{f \in \mathcal{S}} \mathcal{R}_f\} \cup \{\cup_{f \in \{\mathcal{C}(f) \subseteq \mathcal{M}\}} \mathcal{R}_f\}$. The following lemmas can be obtained.

Lemma 1: If the edge server caches only SVCs, i.e., $\mathcal{F} = \mathcal{S}$, the real-valued set function $L(\mathcal{F})$ is a non-decreasing and submodular function, i.e.,

$$L(\mathcal{S} \cup \{f\}) - L(\mathcal{S}) \geq L(\mathcal{S}' \cup \{f\}) - L(\mathcal{S}') \quad (10)$$

Algorithm 1: Content Placement Scheme (Part 1).

```

% Placing MVCs to fill the cache
1: Initialize set  $\mathcal{M} \leftarrow \emptyset$ .
2: while  $w(\mathcal{M}) \leq C$  do
3:   Find MVCs in set  $\mathcal{C}(f^*)$ ,  $f^* \in \mathcal{T}_S$ , where:
      
$$f^* = \operatorname{argmax}_f l(\mathcal{C}(f)|\mathcal{M}).$$

4:   if  $w(\mathcal{C}(f^*) \cup \mathcal{M}) \leq C$  then
5:      $\mathcal{M} \leftarrow \mathcal{C}(f^*) \cup \mathcal{M}$ .
6:   else
7:     Break
8:   end if
9: end while
10: Return set  $\mathcal{M}$ .

```

where $\mathcal{S} \subseteq \mathcal{S}' \subseteq \mathcal{T}_S$ and $f \in \mathcal{T}_S \setminus \mathcal{S}'$.

Proof: See Appendix A. \square

Lemma 2: If the edge server caches only MVCs, i.e., $\mathcal{F} = \mathcal{M}$, the real-valued set function $L(\mathcal{F})$ is a non-decreasing and supermodular function, i.e.,

$$L(\mathcal{M} \cup \{t\}) - L(\mathcal{M}) \leq L(\mathcal{M}' \cup \{t\}) - L(\mathcal{M}') \quad (11)$$

where $\mathcal{M} \subseteq \mathcal{M}' \subseteq \{\mathcal{T}_0 \cup \mathcal{T}_1\}$ and $t \in \{\mathcal{T}_0 \cup \mathcal{T}_1\} \setminus \mathcal{M}'$.

Proof: See Appendix B. \square

Problem (6) is a maximal coverage problem when only SVCs are cached. A greedy algorithm can be applied to obtain a near-optimal solution to the maximal coverage problem [29], [30]. However, when caching both MVCs and SVCs, the submodularity of function $L(\mathcal{F})$ no longer holds. Therefore, we design a heuristic algorithm to cache different types of VCs. The main idea is to successively increase the value of $L(\mathcal{F})$ by updating the sets of VCs cached at the edge server, i.e., \mathcal{M} and \mathcal{S} . The proposed content placement algorithm has three parts, as presented in Algorithms 1 to 3, respectively. The three parts aim to trade off content delivery delay reduction, by caching SVCs, and cache resource usage minimization, by caching MVCs. In the algorithms, we introduce an index as a basis for placing the VCs in the cache, which is defined by

$$l(\mathcal{C}|\mathcal{F}) = L(\mathcal{C} \cup \mathcal{F}) - L(\mathcal{F}).$$

The index, $l(\mathcal{C}|\mathcal{F})$, represents the increased probability of delay requirement satisfaction by caching the VCs in set \mathcal{C} , given that the VCs in set \mathcal{F} are already cached. In addition, the overall data size of the VCs in set \mathcal{F} is defined by

$$w(\mathcal{F}) = \sum_{t \in \mathcal{M}} w_t^M + \sum_{f \in \mathcal{S}} w_f^S.$$

First, in Algorithm 1, MVCs are cached to maximize the probability of delay requirement satisfaction. The algorithm provides an initial caching solution in which the edge server caches as many VCs as possible, since MVCs have smaller data sizes than the corresponding SVCs. Specifically, the algorithm selects MVCs with the highest indexes to increase $L(\mathcal{F})$ greedily. For simplicity, in each iteration, a group of MVCs are cached based on the SVC that they can be stitched and projected to. MVCs

Algorithm 2: Content Placement Scheme (Part 2).

```

% Replace MVCs with SVCs
1: Initialize set  $\mathcal{D} \leftarrow \emptyset$ .
2: Run Algorithm 1, obtain  $\mathcal{M}$ , and  $\mathcal{F} \leftarrow \mathcal{M}$ .
3: while  $L(\mathcal{F}') \geq L(\mathcal{F})$  do
4:    $\mathcal{F} \leftarrow \mathcal{F}'$ .
5:   Find SVC  $f_1^*$  or MVCs  $\mathcal{C}(f_1^*)$ , where:
      
$$\{f_1^*, f_2^*\} = \operatorname{argmax}_{\{f_1, f_2\}} \begin{cases} l(f_1|\mathcal{F} \setminus \mathcal{C}(f_2)), & \text{where } f_1 \notin \mathcal{S}, \\ l(\mathcal{C}(f_1)|\mathcal{F} \setminus \mathcal{C}(f_2)) \end{cases}$$

      and  $\mathcal{C}(f_2) \subseteq \mathcal{M}$ .
6:    $\mathcal{D} \leftarrow \mathcal{C}(f_2^*)$ .
7:   while  $w(\{f_1^*\} \cup \{\mathcal{F} \setminus \mathcal{D}\}) \geq C$  do
8:      $\mathcal{D} \leftarrow \mathcal{C}(f_3)$ , where  $f_3$  can be obtained as follows:
          
$$f_3 = \operatorname{argmax}_f L(\mathcal{F} \setminus \{\mathcal{D} \cup \{\mathcal{C}(f)\}\}), \mathcal{C}(f) \subseteq \mathcal{M}.$$

9:   end while
10:  Let  $\mathcal{F}' \leftarrow \{f_1^*\} \cup \{\mathcal{F} \setminus \mathcal{D}\}$ , if an SVC was selected in Line 4, or  $\mathcal{F}' \leftarrow \{\mathcal{C}(f_1^*)\} \cup \{\mathcal{F} \setminus \mathcal{D}\}$ , if MVCs were selected in Line 4.
11: end while
12: Return set  $\mathcal{F}_A \leftarrow \mathcal{F}$  and  $\Omega_A = L(\mathcal{F})$ .

```

Algorithm 3: Content Placement Scheme (Part 3).

```

% Fill the cache by SVCs to avoid local optimum
1: Run Algorithm 2, obtain  $\mathcal{F}_A$ .
2: Delete all MVCs in set  $\mathcal{F}_A$ , and the result is denoted by  $\mathcal{F}_B$ .
3: while  $w(\mathcal{F}_B) \leq C$  do
4:   Find stereoscopic video  $f^*$ , where:
      
$$f^* = \operatorname{argmax}_f l(f|\mathcal{F}_B).$$

5:   if  $w(\{f^*\} \cup \mathcal{F}_B) \leq C$  then
6:      $\mathcal{F}_B \leftarrow \{f^*\} \cup \mathcal{F}_B$ .
7:   else
8:     Break
9:   end if
10: end while
11: if  $L\{\mathcal{F}_B\} < \Omega_A$  then
12:    $\mathcal{F} \leftarrow \mathcal{F}_A$ .
13: else
14:    $\mathcal{F} \leftarrow \mathcal{F}_B$ .
15: end if
16: Return set  $\mathcal{F}$  as the final caching solution, and update  $\mathcal{M}$  and  $\mathcal{S}$  accordingly.

```

with the highest indexes are sequentially cached until the cache is full. Then, based on the initial caching solution provided by Algorithm 1, Algorithm 2 replaces cached MVCs with SVCs and other MVCs that are not cached. The objective is to further increase the probability of delay requirement satisfaction by reducing the computing delay until no MVCs can be replaced to further improve the probability. At the end of Algorithm 2, a local

optimum of Problem (6) can be obtained. In order to prevent the algorithm from being trapped into a local optimum, Algorithm 3 evaluates the probability of delay requirement satisfaction when only SVCs are cached. We select the caching solution, i.e., \mathcal{F} , which has the highest probability of delay requirement satisfaction from two cases: jointly caching SVCs and MVCs (i.e., \mathcal{F}_A in Algorithm 2) and caching SVCs only (i.e., \mathcal{F}_B in Algorithm 3).

The reason for comparing the two cases is given in the following proposition. Let $\omega \in [0, 1]$ denote the percentage of the cache space assigned for caching SVCs, i.e., $\sum_{f \in \mathcal{S}} w_f^S = \omega C$. The probability of delay requirement satisfaction by caching SVCs with size ωC and MVCs with size $(1 - \omega)C$ through Algorithm 2 is denoted by $L(\omega)$. Let $L^M(1 - \omega)$ and $L^S(\omega)$ represent the probabilities of delay requirement satisfaction by caching MVCs and SVCs, respectively, where $L(\omega) = L^M(1 - \omega) + L^S(\omega)$.

Proposition 1: The following assumptions are made:

- The same type of VCs has the same data size.
- The ratio ω can be any real value between 0 and 1, and $L^M(1 - \omega)$, and $L^S(\omega)$ are differentiable.

If the cache is initially filled by MVCs, and the MVCs are gradually replaced by the SVCs, i.e., increasing ω , there are no more than two local maximizers of ω for Problem (6). If two local maximizers exist, one of them is either $\omega = 1$ or $\omega = 0$.

Proof: See Appendix C. \square

As shown in Proposition 1, Algorithm 2 may find a local optimum under mild conditions. Therefore, we need to compare it with the probability of delay requirement satisfaction when $\omega = 1$ in Algorithm 3. The output of Algorithm 3 provides a near-optimal caching solution given the limited cache size.

B. Parallel Content Placement

Spatio-temporally dividing VR videos that are several minutes in length may generate thousands of VCs. Although Algorithm 3 can provide a near-optimal caching solution, its complexity can be significant for caching a large number of VCs. Observing the independence among VCs in different time segments, we propose a parallel content placement scheme. Specifically, we first divide the whole VC catalog $\mathcal{A} = \{\mathcal{T}_0, \mathcal{T}_1, \mathcal{T}_S\}$ into G VC subsets, in which the g -th subset is denoted by \mathcal{A}_g . Each subset contains VCs in one or multiple time segments. Then, we iteratively determine the cache size allocated for caching VCs in each subset at the edge server. The allocated cache size for VCs in subset \mathcal{A}_g is denoted by C_g . Based on C_g , Algorithm 3 can be used to place VCs in a parallel manner. The optimization problem for determining the optimal cache sizes for the subsets is

$$\max_{\{C_g, \forall g\}} \sum_{g=1}^G P_g(C_g) \quad (12a)$$

$$\text{s.t.} \quad \sum_{g=1}^G C_g \leq C, \quad (12b)$$

$$C_g \geq 0, \forall s \quad (12c)$$

where function $P_g(C_g)$ represents the probability of delay requirement satisfaction for caching VCs in subset \mathcal{A}_g with cache size C_g . The function is defined by

$$P_g(C_g) = \sum_{d \in \mathcal{V}_g} p(d) P(T(d) < H)$$

where \mathcal{V}_g denotes the set of viewpoints to be rendered by the VCs in subset \mathcal{A}_g . The function is quasi-convex and monotonically increases as cache size C_g increases. While the closed-form formulation of $P_g(C_g)$ is intractable, we can solve Problem (12) via an optimization method.

We use a relaxed heavy ball ADMM technique to decouple Problem (12) into sub-problems, which maximize the probability of delay requirement satisfaction for caching VCs in individual subsets, i.e., $P_g(C_g)$, $\forall g$. The problem decomposition is presented in Appendix D. The algorithm for placing VCs in multiple subsets is Algorithm 4. In Lines 3 and 4, the auxiliary variables (provided in Appendix D) and cache sizes allocated for subsets are determined iteratively. The cache size for the g -th subset, i.e., C_g , obtained in the m -th iteration is denoted by C_g^m . From (21a), obtaining the value of C_g^{m+1} requires the knowledge of the gradient of $P_g(C_g^{m+1})$. We estimate the function through a piece-wise linear function $\hat{P}_g(C_g^{m+1})$. In each iteration, when C_g^{m+1} is updated, we use Algorithm 3 to obtain the probability of delay requirement satisfaction, as shown in Lines 7 to 9. Accordingly, function $\hat{P}_g(C_g)$ is updated by (13) shown at the bottom of the next page, in Line 10, where C_g^- represents the maximum cache size for subset g obtained in a previous iteration which is lower than C_g^{m+1} . Note that C_g^+ represents the minimum cache size for subset g obtained in a previous iteration, which is higher than C_g^{m+1} . As the number of iterations increases, the estimated function $\hat{P}_g(C_g^{m+1})$ approaches the actual function $P_g(C_g^{m+1})$. For exploring the actual value of $P_g(C_g)$ in the approximation, we apply a zero-mean variance-attenuated Gaussian noise on C_g^{m+1} in Line 5. The noise, following a normal distribution, $\mathcal{N}(0, \sigma^2)$, has variance σ^2 attenuated in each iteration, as shown in Line 11. The algorithm is stopped when the Lagrangian of the problem, as shown in (20), converges and the exploration on $\hat{P}_g(C_g)$ completes.

The aforementioned problem decomposition and parallel computing significantly reduce the decision space as compared to solving Problem (12) directly. Without parallel content placement, the worst time complexity of Algorithm 3 is $(|\mathcal{V}|(M_1 + M_3) + 2|\mathcal{V}|^2 M_2)$, where M_1 , M_2 , and M_3 are the iteration numbers in Algorithms 1, 2, and 3, respectively. With an increase in the number of viewpoints, the time complexity increases quadratically. In Algorithm 4, the content placement for different video subsets is in parallel. If the popularity functions $\{P_g(C_g), \forall g\}$ are known in advance, the number of iterations for running the relaxed heavy ball ADMM algorithm can be low, and the worst time complexity of content placement for a subset becomes $(|\mathcal{V}_g|(M_1 + M_3) + 2|\mathcal{V}_g|^2 M_2)$.

VI. DRL-BASED CONTENT DELIVERY

The content placement scheme proposed in Section V maximizes the probability of delay requirement satisfaction through

Algorithm 4: Parallel Content Placement Scheme.

```

1: Initialize estimation function  $\hat{P}_g^0(C_g), m = 0.$ 
2: while 1 do
3:   Calculate  $C_g^{m+1}\{\hat{P}_g^k(C_g)\}$  by (21a).
4:   Update variables in set  $\Gamma$  by (21b)-(21h).
5:    $C_g^{m+1} \leftarrow C_g^{m+1} + \mathcal{N}(0, \sigma^2), \forall g.$ 
6:   for  $g \in \mathcal{G}$  do
7:     Place VCs in subset  $g$  given cache size  $C_g^{m+1}$  by
       Algorithm 3.
8:      $\hat{P}_g^{m+1}(C_g^{m+1}) = P_g(C_g^{m+1}).$ 
9:   end for
10:  Update  $\hat{P}_g^{m+1}(C_g)$  by (13).
11:   $\sigma \leftarrow \alpha_{attn}\sigma, m = m + 1.$ 
12:  if  $|\mathcal{L}(\Gamma)^{m+1} - \mathcal{L}(\Gamma)|$  and  $\sigma$  are less than thresholds
       then
13:    Break.
14:  end if
15: end while
16: Return the caching solution with cache capacity  $C_g$  for
    all VC subsets.

```

caching VCs at the edge server. However, a viewpoint may not be always rendered in time, especially when unpopular videos are requested. Therefore, we propose a content delivery scheduling scheme to solve Problem (7) in this section, given the caching solution and predefined video quality setting X_0 . Specifically, we design a Whittle index (WI) based scheduling policy in restless multi-armed bandit (RMAB) and a DRL-based scheme to evaluate the WI in the scheduling policy.

A. Whittle Index Formulation

Problem (7) with stochastic constraint (8) is an RMAB problem. For U VR headsets, we define U corresponding controlled Markov chains. The state of the Markov chain for VR headset u in time slot k is denoted by $s_{u,k}$, and the state space for VR headset u is denoted by \mathcal{Y}^u . State $s_{u,k}$ consists of 1) the historical viewpoint trajectory of VR headset u , i.e., $\{v_{u,k^*}, k^* = k - P, \dots, k\}$; 2) the SVCs in the local buffer of VR headset u ;

3) the request $d_{u,k}$; and 4) the sets of VCs cached at the edge server.

Denote the binary control variable for the Markov chain by $a_{u,k}$, where $a_{u,k} = 1$ if the request of headset u is scheduled in time slot k , and $a_{u,k} = 0$ otherwise. Consider the average time length for a state transition as an epoch, which is denoted by ϕ . The average number of time slots in an epoch is $\phi = \bar{T}/(E \times \delta)$, which is utilized to find tractable state transition probabilities. The probability of transiting from state $s_{u,k}$ to $s_{u,k+\phi}$ is approximated by $P(s_{u,k+\phi}|s_{u,k}, a_{u,k})$, where $\sum_{s_{u,k+\phi} \in \mathcal{Y}^u} P(s_{u,k+\phi}|s_{u,k}, a_{u,k}) = 1$. The reward of scheduling a request is set as the number of future viewpoints that can be rendered from granting the request, given by (14), shown at the bottom of the this page, The discounted reward model is applied in (14), where κ is a discount factor. Problem (7) with constraint (8) is now an RMAB problem. Whenever a computing unit is available, we select one request from one of the U VR headsets to schedule with unknown transition probabilities.

Our method of solving the RMAB problem is based on the WI, which is a heuristic solution with excellent empirical performance [31]. Specifically, the WI-based method defines a subsidy, i.e., the WI, for each VR headset if its current request cannot be scheduled by the edge server immediately according to the current state of the VR headset. The edge server can receive a higher long-term reward by scheduling the requests with a higher subsidy, i.e., higher WI. Therefore, the edge server can allocate computing resources to requests by comparing the WIs of requests. The WI of state $s_{u,k}$ can be obtained from (15). In (15), $Q(s_{u,k}, a_{u,k})$ is the state-action value, i.e., the Q value, for state $s_{u,k}$, and action $a_{u,k}$. Eq (15) shown at the bottom of the this page, presents the WI of VR headset u in time slot k .

In time slot k , the edge server collects the WIs $\{\lambda(s_{u,k}), \forall u\}$ from all VR headsets and schedules the request of the VR headset with the highest WI value if a computing unit is available. As shown in Appendix E, the RMAB problem is indexable if the content delivery delay for all requests is the same. Even though the content delivery delay can be different in practice, the standard deviation of the delay is usually small as compared to

$$\hat{P}_g^{m+1}(C_g) = \begin{cases} \frac{\hat{P}_g^{m+1}(C_g^{m+1}) - \hat{P}_g^m(C_g^-)}{C_g^{m+1} - C_g^-} (C_g - C_g^-) + \hat{P}_g^m(C_g^-), & \text{for } C_g^- \leq C_g < C_g^{m+1} \\ \frac{\hat{P}_g^m(C_g^+) - \hat{P}_g^{m+1}(C_g^{m+1})}{C_g^+ - C_g^{m+1}} (C_g^+ - C_g) + \hat{P}_g^{m+1}(C_g^{m+1}), & \text{for } C_g^{m+1} \leq C_g < C_g^+ \\ \hat{P}_g^m(C_g), & \text{Otherwise.} \end{cases} \quad (13)$$

$$R(s_{u,k}, a_{u,k}) = \begin{cases} \sum_{k'=k+\phi}^{\infty} \kappa^{k'-k} f(v; \mathcal{B}_{u,k}) \mathbf{1}[v \in \mathcal{U}_{u,k}(\mathbf{a}_{u,k-1})], & \text{if } a_{u,k} = 1; \\ 0, & \text{if } a_{u,k} = 0. \end{cases} \quad (14)$$

$$\lambda(s_{u,k}) = R(s_{u,k}, 1) + \kappa^\phi \left\{ \sum_{s_{u,k+\phi} \in \mathcal{Y}^u} P(s_{u,k+\phi}|s_{u,k}, 1) \max_{a_{u,k+\phi}} \{Q(s_{u,k+\phi}, a_{u,k+\phi})\} \right. \\ \left. - \sum_{s_{u,k+\phi} \in \mathcal{Y}^u} P(s_{u,k+\phi}|s_{u,k}, 0) \max_{a_{u,k+\phi}} \{Q(s_{u,k+\phi}, a_{u,k+\phi})\} \right\}. \quad (15)$$




Fig. 4. DRL-based content delivery scheduling scheme diagram.

the average content delivery delay³, and the proposed WI-based solution can still be applied.

B. DRL-Based Scheduling

Since the state-action transition probability is unknown for the RMAB problem, the value of the WI in (15) cannot be obtained straightforwardly. We adopt the DRL approach to approximate the WI value for each VR headset in each time slot. Two neural networks are constructed for each VR headset. One neural network approximates the Q value, and the other neural network approximates the WI value. The weight vectors of the two neural networks corresponding to the Q value and the WI value are denoted by θ_Q and θ_W , respectively.

The proposed DRL-based content delivery scheduling scheme is illustrated in Fig. 4. At the beginning of time slot k , VR headset u observes its state $s_{u,k}$, predicts the future user viewpoints using viewpoint prediction techniques [19], and obtains content delivery requests. The Q value and WI evaluation networks approximate the set of Q values and WI values regarding state $s_{u,k}$, denoted by $Q(s_{u,k}, a_{u,k} | \theta_Q)$ and $\lambda(s_{u,k}, a_{u,k} | \theta_W)$, respectively. With the set of requests, the approximated WI values are sent to the edge server, and the edge server schedules the request with the highest WI when a computing unit is available. Furthermore, at each time slot, the weight vectors θ_Q and θ_W are updated according to the reward obtained in content delivery scheduling. In particular, if the request of VR headset u is scheduled, the VR headset observes the reward as given by (14). Otherwise, we use the approximated WI $\lambda(s_{u,k} | \theta_W)$ as the reward, which represents the difference of Q values between two successive states when a request is not satisfied. According to [32], weight vector θ_Q is updated by minimizing the loss function for evaluating the Q value, given by

$$L(\theta_Q) = [Q(s_{u,k}, a_{u,k} | \theta_Q) - Q(s_{u,k+\phi}, a_{u,k} | \theta_Q) \\ \times \lambda(s_{u,k} | \theta_W)(1 - a_{u,k}) - R(s_{u,k}, a_{u,k})a_{u,k}]^2. \quad (16)$$

Furthermore, we generate a reference WI value based on the outputs of the Q value and WI value neural networks, where

$$\hat{\lambda}(s_{u,k}) = \lambda(s_{u,k} | \theta_W) - \varphi [Q(s_{u,k}, 0 | \theta_Q) - Q(s_{u,k}, 1 | \theta_Q)]. \quad (17)$$

In (17), φ is the step size for approximating the WI value. From (17), if the approximation of $\lambda(s_{u,k})$ is accurate,

³The standard deviation of the content delivery delay is bounded by half of the computing delay, which is only a portion of the content delivery delay under mild conditions.

Algorithm 5: DRL-Based Content Delivery Scheduling Scheme.

- 1: Initialize weight vectors θ_W and θ_Q for all VR headsets. $\epsilon = 1$.
- 2: **for** time slot $k = 1, \dots, K$ **do**
- 3: **for** VR headset $u = 1 : U$ **do**
- 4: Request the new SVC and report WI set $\{\lambda(s_{u,k}; \theta_W)\}$.
- 5: **end for**
- 6: **if** any computing unit is available **then**
- 7: With probability ϵ , randomly schedule a request; otherwise, schedule request d^* from the VR headset with the highest WI.
- 8: **end if**
- 9: **for** VR headset $u = 1 : U$ **do**
- 10: Update the Q value evaluation network by $\theta_Q \leftarrow \theta_Q - \varphi_Q \nabla L(\theta_Q)$.
- 11: Update the WI evaluation network by $\theta_W \leftarrow \theta_W - \varphi_W \nabla L(\theta_W)$.
- 12: **end for**
- 13: **if** $\epsilon > \epsilon_{min}$, $\epsilon \leftarrow \beta_{attn}\epsilon$.
- 14: **end for**

$Q(s_{u,k}, 0 | \theta_Q)$ should be equal to $Q(s_{u,k}, 1 | \theta_Q)$. Thus, weight vector θ_Q is updated by minimizing the loss function for evaluating the WI value, given by

$$L(\theta_W) = [\varphi(Q(s_{u,k}, 0 | \theta_Q) - Q(s_{u,k}, 1 | \theta_Q))]^2. \quad (18)$$

The proposed DRL-based content delivery scheduling scheme is summarized in Algorithm 5, where φ_Q and φ_W are the learning rates for evaluating weight vectors θ_Q and θ_W , respectively. At each time slot, a VR headset calculates WI, which is a float number (4 bytes), in a distributed manner and uploads a WI for content delivery scheduling. Compared to sending the full VR headset state to the edge server for scheduling, the proposed scheme can significantly reduce communication overhead.

VII. ADAPTIVE VIDEO QUALITY ADJUSTMENT

The edge caching and content delivery scheduling schemes proposed in Sections V and VI are based on a predefined video quality setting, i.e., a given X_0 . In this section, we discuss how to optimize video quality in content delivery while meeting frame missing rate requirements. For simplicity, we assume that the data sizes of different MVCs are similar [23], and the data size of an SVC does not depend on X_0 . The frame missing rate is jointly determined by the following three factors, as illustrated in Fig. 5:

- 1) *Request waiting time*: The request waiting time can be obtained by the time difference between the instant when a content delivery request is generated and the next instant when a request is scheduled by the edge server. The average request waiting time is denoted by W .
- 2) *Content delivery delay*: When a request is scheduled by the edge server, a computing unit is occupied to process and deliver the corresponding SVC. The content delivery




Fig. 5. Illustration of content delivery processes for a VR headset.

TABLE I
METADATA OF VR VIDEOS [28]

No	Video length	Content	Category	Popularity
1	2'44"	Conan360-Sandwich	Performance	0.6
2	3'20"	Freestyle Skiing	Sport	0.2
3	4'48"	Google Spotlight-Help	Film	0.2

delay for request $d_{u,k}$, i.e., $T(d_{u,k})$, is evaluated as in Section III.C.

- 3) *Request deadline*: The request deadline is the time slot when the requested video will be played on the VR headset. In the example shown in Fig. 5, slot k' is the request deadline of the request generated in slot k .

The request waiting time depends on the content delivery scheduling scheme and the number of VR headsets in the system. The request deadline depends on the viewpoint dynamics of the corresponding user and parameter X_0 (since X_0 impacts the number of viewpoints that can be rendered by the VR headset using the SVCs in the local buffer). Let random variable $N(X_0)$ denote the number of frame missing events between two successive content delivery for a VR headset given X_0 , with mean denoted by $\bar{N}(X_0)$. The relation between $\bar{N}(X_0)$ and the video quality setting, i.e., X_0 , is given by the following lemma:

Lemma 3: $\bar{N}(X_0)$ monotonically decreases as X_0 increases, under the assumption that the data sizes of MVCs are constant.

Proof: See Appendix F. \square

Based on Lemma 3, parameter X_0 can be adjusted tentatively to maximize video quality, while satisfying the frame missing rate requirement. Specifically, a VR headset can monitor the real-time frame missing rate and upload it to the edge server. If the average frame missing rate is higher than the frame missing rate requirement, X_0 should be increased to reduce the frame missing rate by degrading video quality. Otherwise, X_0 can be decreased to improve the video quality, until the frame missing rate requirement is no longer satisfied.

VIII. SIMULATION RESULTS

A. Parameter Settings

In the simulation, we use a viewpoint-tracking data set composed of 48 users watching three spherical VR videos [28]. The details of the VR videos are provided in Table I. The $360^\circ \times 180^\circ$ equirectangular video is divided into 24×12 spatial tiles. A user's FoV spans 7×5 tiles. The length of a time segment, i.e., the playback duration of a VC, is 4 seconds. For each time segment, we map the spatial locations of viewpoints

TABLE II
NETWORK STRUCTURE FOR EVALUATING THE WI VALUE

Layers	Number of neurons	Activation function
CONV1	$(3 \times 3 \times 10)$	elu
POOL1	(2×2)	elu
CONV2	$(3 \times 3 \times 10)$	elu
LSTM	50	tanh
Fully Connected 1	125	elu
Fully Connected 2	64	elu
Fully Connected 3	64	elu
Fully Connected 4	24	elu
Output	1	linear

TABLE III
NETWORK STRUCTURE FOR EVALUATING THE Q VALUE

Layers	Number of neurons	Activation function
CONV1	$(3 \times 3 \times 10)$	elu
POOL1	(2×2)	elu
CONV2	$(3 \times 3 \times 10)$	elu
LSTM	50	tanh
Fully Connected 1	125	elu
Fully Connected 2	64	elu
Output	1	linear

into tiles to obtain the probability that request d is generated, i.e., $p(d)$. The data size of MVCs is randomly generated based on a Gaussian distribution, in which the mean is 30 Kbits and the standard deviation is 10 Kbits, and the value of α is 1.3. The content delivery delay threshold H is 85 ms. The wired transmission rate from the cloud server to the edge server is 700 Mbits/s. In parallel content placement, VCs corresponding to the same time segment, i.e., the playback duration of a VC, form a VC subset, and there are 163 subsets for the three videos. In the parallel content placement, the parameters of the relaxed heavy ball ADMM algorithm are set to $\varepsilon_1 = 0.8$, $\varepsilon_2 = 0.85$, and $\rho_1 = \rho_2 = 0.3$.

In the content delivery, there are 10 VR users with VR headsets. The length of a time slot is 33 ms. For each user, the viewpoint trajectory is randomly selected from the viewpoint-tracking data set. For processing the content delivery requests from users, the edge server has one computing unit, i.e., $E = 1$. The transition probabilities of the communication channel between the edge server and users, i.e., p_H and p_L , are 0.6 and 0.3, respectively. In the DRL algorithm for content delivery scheduling, the learning rates of the networks for evaluating the Q value and the WI value are $5e-4$ and $1e-4$, respectively. Neural network structures used to evaluate the Q value and the WI value are presented in Tables II and III, respectively. At each VR headset, a long short-term memory (LSTM) network with 50 neurons predicts the user's future viewpoint with an accuracy of around 94%.

B. Performance of the Content Placement Scheme

We first evaluate the performance of the proposed content placement scheme (i.e., Algorithms 1-3) for caching VCs within a VC subset. We compare our proposed scheme with three benchmark schemes: *MVC only* and *SVC only*, in which only




Fig. 6. The probability of delay requirement satisfaction with various computing rates, where $R_L = 50$ MB/s.




Fig. 7. The probability of delay requirement satisfaction with various channel transition probabilities, where $R_L = 50$ MB/s.

MVCs or SVCs are cached, respectively; *brute-force*, which searches the best cache size for caching SVCs, i.e., the best ω , from set $\{0, 0.1, \dots, 1\}$. The probability of delay requirement satisfaction with various computing rates χ is shown in Fig. 6, and the probability of delay requirement satisfaction with various channel transition probabilities, i.e., p_H , is shown in Fig. 7. In comparison with the *MVC-only* and *SVC-only* schemes, the proposed content placement scheme can significantly improve the delay requirement satisfaction probability. This is because the proposed scheme can find the near-optimal caching solution by balancing computing delay and cache resource usage. Furthermore, the proposed scheme outperforms the *brute-force* scheme, which is a computationally intensive scheme based on exhaustive search, in satisfying the delay requirement. In addition, in both Figs. 6 and 7, the performance of the *MVC only* scheme improves significantly as the computing rate increases, whereas the performance improvement of the *SVC only* scheme is not significant, especially when the transmission rate R_H is low. This is because the edge server may frequently download MVCs from the cloud server, and the performance bottleneck in the *SVC only* scheme is on the transmission delay for delivering

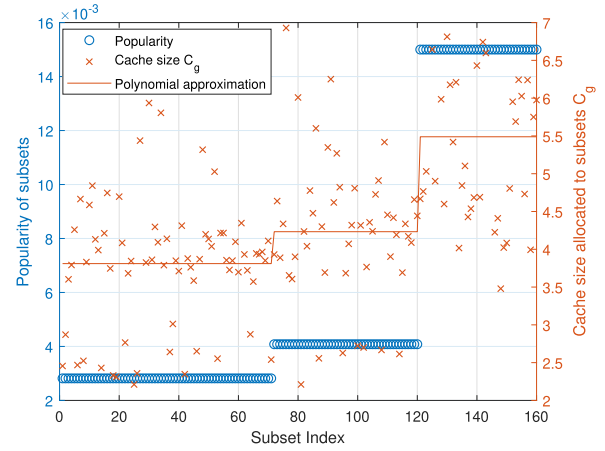


Fig. 8. The correlation between the popularity of VC subsets, and the allocated cache sizes for VC subsets.

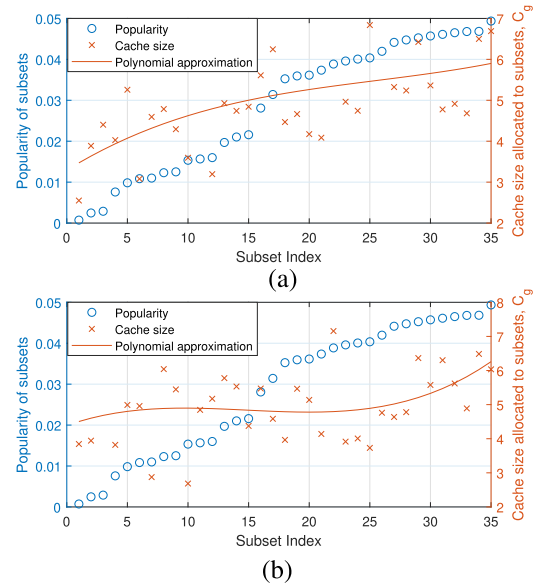


Fig. 9. The correlation between the popularity of VC subsets in two different videos and the allocated cache sizes for VC subsets. (a) video 1; (b) video 2.

VCs in both wireless and wired links. When the edge server has sufficient computing resources, caching more MVCs can increase the probability of delay requirement satisfaction.

The performance of the parallel content placement scheme (Algorithm 4) for caching VCs associated with multiple time segments is presented in Figs. 8 and 9, where $R_H = 88$ MB/s and $R_L = 50$ MB/s. Fig. 8 shows the correlation between the popularity of VC subsets, i.e., $\{\sum_{d \in V_g} p(d), \forall g\}$, and the allocated cache sizes for the subsets, i.e., $\{C_g, \forall g\}$, in which all 163 subsets are taken into account. Popular VC subsets are allocated with larger cache sizes than less popular VC subsets. We observe that even if the popularity of different VC subsets is the same, the allocated cache sizes can be different since viewpoint popularity distributions are different among subsets. The correlation between the popularity of VC subsets in two different videos and the allocated cache sizes is shown in Fig. 9.




Fig. 10. The average hit probability with LSTM-based prediction and dynamic channel.

We use 35 VC subsets for both videos 1 and 3 in Table I and assign random popularity for the VC subsets. Compared to video 1 (performance video), the popularity and the allocated cache size for a subset in video 3 (film video) are not strongly correlated, especially for the subsets with high popularity. The reason is, compared to video 1, user viewpoints in video 3 are more concentrated. In such a case, increasing the cache sizes for popular VC subsets may not significantly improve the overall probability of delay requirement satisfaction.

C. Performance of the Content Delivery Scheduling Scheme

The performance of the proposed WI-based content delivery scheduling scheme is compared with three benchmark schemes: *urgent-request-first (URF)*, in which the edge server always schedules the request for the SVC that will be played in the nearest future first; *round-robin*, in which the edge server schedules requests in a circular order; and *random*, in which the edge server schedules requests randomly. We generate two different video selection and viewpoint movement profiles based on the viewpoint-tracking data set: one for training the neural networks; another for testing the learning-based scheme performance. For each profile, we generate requests for 2,000 seconds. With the viewpoint movement prediction by LSTM, the performance of the average hit probability for the four schemes is shown in Fig. 10, in which the test profile is applied. The results in Fig. 10 are the moving average of the hit probability in 1,000 preceding decision epochs (around 8,000 time slots) for each decision epoch. As shown in the figure, the proposed WI-based content delivery scheduling can significantly improve the hit probability. When the transmission rate from the edge server to a headset is low, the proposed WI-based scheme can improve the hit probability by up to 30% compared to *URF* and *round-robin* schemes. Our proposed WI-based scheme improves performance because, compared to the benchmark schemes, our proposed WI-based scheme evaluates viewpoint and network dynamics in making content delivery scheduling decisions. The relation between video quality and average hit probability is shown in Fig. 11. As the portion of layer-0 MVCs in an SVC,




Fig. 11. The average hit probability versus various video quality settings.

i.e., X_0 , decreases, the hit probability decreases accordingly, while the video quality is improved. The proportional relation is consistent with Lemma 3. When implementing content placement and scheduling schemes, parameter X_0 should be adapted to achieve the optimal balance between the video quality and the hit probability.

IX. CONCLUSION

In this article, we have studied content caching and delivery schemes for mobile VR video streaming. Given limited network resources, a scalable content placement scheme has been designed to obtain caching solutions for VR videos by trading off caching and computing resources. In addition, we have developed a novel distributed learning-based content delivery scheduling scheme to coordinate real-time content delivery for multiple VR headsets. The proposed content delivery solution identifies the priority of content delivery among the requests from multiple VR headsets based on the network and user viewpoint dynamics, thereby maximizing resource utilization. Simulation results have demonstrated that the proposed solution outperforms benchmark content caching and delivery schemes. In the future, we will investigate efficient computing resource allocation and caching schemes for interactive VR applications, where the dynamics of VR user interactions will be evaluated.

APPENDIX A PROOF OF LEMMA 1

It can be directly proved by using (9) and the definition of submodularity. Let \mathcal{S}_1 and \mathcal{S}_2 be the subsets of \mathcal{S} satisfying $\mathcal{S}_1 \subseteq \mathcal{S}_2$. For any $f \in \mathcal{T}_S \setminus \mathcal{S}_2$, we have

$$L(\mathcal{S}_1 \cup \{f\}) - L(\mathcal{S}_1) = L(\mathcal{X}_1)$$

where

$$\mathcal{X}_1 = \{\cup_{f' \in \mathcal{S}_1 \cup \{f\}} \mathcal{R}(f')\} \cap \{\cup_{f' \in \mathcal{S}_1} \mathcal{R}(f')\}.$$

Similarly,

$$\mathcal{X}_2 = \{\cup_{f' \in \mathcal{S}_2 \cup \{f\}} \mathcal{R}(f')\} \cap \{\cup_{f' \in \mathcal{S}_2} \mathcal{R}(f')\}$$

and $L(\mathcal{S}_2 \cup \{f\}) - L(\mathcal{S}_2) = L(\mathcal{X}_2)$. Since $\mathcal{S}_1 \subseteq \mathcal{S}_2$, some viewpoints that can be rendered by the VR headset using SVC f and cannot be satisfied by SVCs in set \mathcal{S}_1 may be satisfied

by SVCs in set \mathcal{S}_2 . Therefore, $\mathcal{X}_2 \subseteq \mathcal{X}_1$, and $L(\mathcal{S}_1 \cup \{f\}) - L(\mathcal{S}_1) \geq L(\mathcal{S}_2 \cup \{f\}) - L(\mathcal{S}_2)$.

APPENDIX B PROOF OF LEMMA 2

It can be proved similarly to Lemma 1. Let \mathcal{M}_1 and \mathcal{M}_2 be the subsets of \mathcal{M} satisfying $\mathcal{M}_1 \subseteq \mathcal{M}_2$, and $L(\mathcal{M}_1 \cup \{t\}) - L(\mathcal{M}_1) = L(\mathcal{X}_1)$. In this case, \mathcal{X}_1 can be defined by

$$\mathcal{X}_1 = \cup_{f' \in \{f | \mathcal{C}(f) \subseteq \{\mathcal{M}_1 \cup \{t\}\}, t \in \mathcal{C}(f)\}} \mathcal{R}(f').$$

Similarly, let $L(\mathcal{M}_2 \cup \{t\}) - L(\mathcal{M}_2) = L(\mathcal{X}_2)$, where

$$\mathcal{X}_2 = \cup_{f' \in \{f | \mathcal{C}(f) \subseteq \{\mathcal{M}_2 \cup \{t\}\}, t \in \mathcal{C}(f)\}} \mathcal{R}(f').$$

Since $\mathcal{M}_1 \subseteq \mathcal{M}_2$, by MVC t , more SVCs can be stitched when MVCs in \mathcal{M}_2 are cached as compared to when MVCs in \mathcal{M}_1 are cached. Therefore, $\mathcal{X}_1 \subseteq \mathcal{X}_2$, and $L(\mathcal{M}_1 \cup \{t\}) - L(\mathcal{M}_1) \leq L(\mathcal{M}_2 \cup \{t\}) - L(\mathcal{M}_2)$.

APPENDIX C PROOF OF PROPOSITION 1

Let functions $L^M(\omega|x)$ and $L^S(\omega|x)$ denote the probabilities of delay requirement satisfaction by caching MVCs and SVCs with data size ωC , respectively, given that the same type of videos with data size x has been cached. If both $L^M(1-\omega)$ and $L^S(\omega)$ are differentiable, $L(\omega)$ can be differentiated as follows:

$$\begin{aligned} \nabla_{\omega} L(\omega) &= \nabla_{\omega} L^S(\omega) + \nabla_{\omega} [L^M(1) - L^M(\omega|1-\omega)] \\ &= \nabla_{\omega} [L^S(\omega) - L^M(\omega|1-\omega)] \\ &= \lim_{\Delta \rightarrow 0} \frac{1}{\Delta} \{L^S(\omega + \Delta) - L^S(\omega) \\ &\quad + L^M(\omega + \Delta|1-\omega-\Delta) - L^M(\omega|1-\omega)\} \\ &= \lim_{\Delta \rightarrow 0} \frac{1}{\Delta} \{L^S(\Delta|\omega) - L^M(\Delta|1-\omega-\Delta)\}. \end{aligned}$$

Assuming that all SVCs have the same data size, $L^S(\Delta|\omega)$ monotonically decreases as ω increases due to submodularity proved in Lemma 1. Moreover, it can be directly proved that $L^S(\Delta|\omega)$ is convex on ω . This is because, when caching an additional SVC, with fewer SVCs cached at the edge server, viewpoints that can be rendered by VR headsets using the additional SVC are less likely to intersect with viewpoints rendered by the cached SVCs, i.e.,

$$L^S(\Delta|\omega + \delta) - L^S(\Delta|\omega) \geq L^S(\Delta|\omega + 2\delta) - L^S(\Delta|\omega + \delta).$$

Similarly, it can be shown that function $L^M(\Delta|1-\omega-\Delta)$ monotonically decreases as ω increases due to supermodularity proved in Lemma 2, and the function is concave on ω .

If a local optimal ω^* exists, $\nabla_{\omega} L(\omega^*) = 0$. In such case, $L^S(\Delta|\omega) = L^M(\Delta|1-\omega-\Delta)$ holds for any Δ . Therefore, as shown in Fig. 12, there are three possible cases in finding the local optimal ω^* :

- *Case 1*: No intersection between $L^S(\Delta|\omega)$ and $L^M(\Delta|1-\omega-\Delta)$. Thus, no local optimum between the region $\omega \in (0, 1)$. The global maximizer of $L(\omega)$ is located either in $\omega = 0$ or $\omega = 1$.




Fig. 12. Illustration of the proof of Proposition 1.

- *Case 2*: Two intersections between $L^S(\Delta|\omega)$ and $L^M(\Delta|1-\omega-\Delta)$ (i.e., points A and B in Fig. 12). Thus, there are two local optimums when $\omega \in (0, 1)$, and only one of the local optimums is the local maximizer, denoted by ω^* . If $\omega^* \in (0, 1)$ is not the global maximizer, the global maximizer is located at $\omega = 1$. This is because, as ω increases from zero to the first intersection (i.e., intersection A in Fig. 12), $\nabla_{\omega} L(\omega)$ has to be greater than zero. Therefore, the value of $L(\omega)$ at intersection A has to be greater than $L(\omega)$.
- *Case 3*: Only one intersection between $L^S(\Delta|\omega)$ and $L^M(\Delta|1-\omega-\Delta)$ (i.e., point C in Fig. 12). There is at most one local maximizer when $\omega \in (0, 1)$, and the other local maximizer is either $\omega = 1$ or $\omega = 0$.

In summary, based on the three cases, there are mostly two local maximizers of ratio ω for Problem (6).

APPENDIX D PROBLEM DECOMPOSITION FOR (12)

The first step of parallel content placement is to decouple Problem (12) into subproblems for caching the VCs in individual subsets. Introducing auxiliary variables z_1 and $\{z_{2,g}, \forall g\}$, we reformulate constraints (12b) and (12c) as

$$\bar{C} - \frac{C_g}{|\mathcal{G}|} - z_1 = 0, \quad (19a)$$

$$C_g - z_{2,g} = 0, \forall g \quad (19b)$$

respectively, where \bar{C} represents the mean value of $\{C_g, \forall g\}$. We formulate the augmented Lagrangian of the problem as

$$\begin{aligned} \mathcal{L}(\Gamma) &= - \sum_{g \in \mathcal{G}} P_g(C_g) + u_1 \left(\bar{C} - \frac{C_g}{|\mathcal{G}|} - z_1 \right) \\ &\quad + \sum_g u_{2,g} (C_g - z_{2,g}) \\ &\quad + \frac{\rho_1}{2} \left(\bar{C} - \frac{C}{|\mathcal{G}|} - z_1 \right)^2 + \frac{\rho_2}{2} \sum_g (C_g - z_{2,g})^2 \quad (20) \end{aligned}$$

where u_1 and $\{u_{2,g}, \forall g\}$ are the Lagrangian multipliers for Constraints (19a) and (19b), respectively. Parameters ρ_1 and ρ_2 are penalty factors. The term Γ represents the set of variables, i.e., $\Gamma = \{C_g, z_{2,g}, u_{2,g}, \forall g, z_1, u_1\}$. We then adopt relaxed heavy ball ADMM [33] to update the variables in an iterative and parallel manner, where decoupled subproblems for iteration

$e + 1$ are as follows:

$$\begin{aligned} C_g^{e+1} = & \underset{C_g^{e+1}}{\operatorname{argmin}} \left\{ -P_g(C_g^{e+1}) + u_1^e \left(\bar{C}^e - C_g^e + C_g^{e+1} - \frac{C}{|G|} - z_1^e \right) \right. \\ & + u_{2,g}^e (C_g^{e+1} - z_{2,g}^e) + \frac{\rho_1}{2} \\ & \times \left(\bar{C}^e - C_g^e + C_g^{e+1} - \frac{C}{|G|} - z_1^e \right)^2 \\ & \left. + \frac{\rho_2}{2} (C_g^{e+1} - z_{2,g}^e)^2 \right\}, \end{aligned} \quad (21a)$$

$$z_{2,g}^{e+1} = \max \left\{ \varepsilon_1 C_g^{e+1} + (1 - \varepsilon_1) \hat{z}_{2,g}^e + \frac{1}{\rho_2} \hat{u}_{2,g}^e, 0 \right\}, \forall g, \quad (21b)$$

$$u_1^{e+1} = \hat{u}_1^e + \rho_1 \left[\varepsilon_1 \left(\bar{C}^{e+1} - \frac{C}{|G|} \right) + (1 - \varepsilon_1) \hat{z}_1^e - z_1^{e+1} \right], \quad (21c)$$

$$u_{2,g}^{e+1} = \hat{u}_{2,g}^e + \rho_1 \left[\varepsilon_1 C_g^{e+1} + (1 - \varepsilon_1) \hat{z}_{2,g}^e - z_{2,g}^{e+1} \right], \forall g, \quad (21d)$$

$$\hat{u}_1^{e+1} = u_1^{e+1} + \varepsilon_2 (u_1^{e+1} - \hat{u}_1^e), \quad (21e)$$

$$\hat{u}_{2,g}^{e+1} = u_{2,g}^{e+1} + \varepsilon_2 (u_{2,g}^{e+1} - \hat{u}_{2,g}^e), \forall g, \quad (21f)$$

$$\hat{z}_1^{e+1} = z_1^{e+1} + \varepsilon_2 (z_1^{e+1} - \hat{z}_1^e), \quad (21g)$$

$$\hat{z}_{2,g}^{e+1} = z_{2,g}^{e+1} + \varepsilon_2 (z_{2,g}^{e+1} - \hat{z}_{2,g}^e), \forall g \quad (21h)$$

where $0 \leq \varepsilon_1 \leq 1$ and $0 \leq \varepsilon_2 \leq 1$.

APPENDIX E

INDEXABILITY OF THE RMAB PROBLEM

Let $\mathcal{P}(\lambda^*)$ denote the set of state \mathcal{Y} for which the VR headset stays passive, i.e., $a_{u,k} = 0$, according to subsidy $\lambda^*(s_{u,k})$, $s_{u,k} \in \mathcal{Y}$. If the RMAB problem is indexable, the following condition should be satisfied:

- Set $\mathcal{P}(\lambda^*)$ of the corresponding single-armed bandit process increases monotonically from the empty set to the whole state space as subsidy λ^* increases from $-\infty$ to $+\infty$ [34].

To prove the indexability, we assume that the content delivery delay for all requests is the same, which is denoted by ϕ . If the following condition holds, the problem is indexable [34]:

$$\frac{dQ(s_{u,k}, 0)}{d\lambda} \geq \frac{dQ(s_{u,k}, 1)}{d\lambda}. \quad (22)$$

Define function $B_\lambda(s_{u,k}) = \frac{dV(s_{u,k})}{d\lambda}$, where $V(s_{u,k}) = \max\{Q(s_{u,k}, 0), Q(s_{u,k}, 1)\}$. Denote the SVC that is scheduled in time slot k by b . The condition in (22) is equivalent to

$$1 + \kappa^\phi B_\lambda(s_{u,k+\phi}) \geq \kappa^\phi B_\lambda(s_{u,k+\phi} \cup \{b\}). \quad (23)$$

To prove the above relation, we first assume that b is the last SVC that can be downloaded, i.e., all SVCs have been downloaded in the buffer except b . The left-hand side (LHS) of (23) can be rewritten as

$$\begin{aligned} \text{LHS} &= 1 + \kappa^\phi \max[1 + \kappa^\phi B_\lambda(s_{u,k+\phi}), \kappa^\phi B_\lambda(s_{u,k+\phi} \cup \{b\})] \\ &\geq 1 + \kappa^{2\phi} B_\lambda(s_{u,k+2\phi} \cup \{b\}). \end{aligned} \quad (24)$$

The right-hand side (RHS) of (23) can be rewritten as

$$\text{RHS} = \kappa^\phi (1 + \kappa^\phi B_\lambda(s_{u,k+2\phi} \cup \{b\})). \quad (25)$$

It can be observed that the value of the LHS is greater than the value of the RHS since $\kappa < 1$.

Moreover, we assume that b is the last h -th SVCs to be downloaded from the edge until the SVCs are fully downloaded in the buffer, and (23) holds in such a case. In time slot k' , where $n' < n$, the last $(h+1)$ -th SVC is requested by the VR headset, which is denoted by b' . The LHS of (23) for skipping to download b' in time slot k' can be represented as

$$\begin{aligned} \text{LHS} &= 1 + \kappa^\phi \max \\ &[1 + \kappa^\phi B_\lambda(s_{u,k'+2\phi}), \kappa^\phi B_\lambda(s_{u,k'+2\phi} \cup \{b^*\})] \end{aligned} \quad (26)$$

where b^* can be either b' or b . Moreover, the RHS of (23) for downloading b' in time slot n' can be represented as

$$\begin{aligned} \text{RHS} &= \kappa^\phi \max \\ &[1 + \kappa^\phi B_\lambda(s_{u,k'+2\phi} \cup \{b'\}), \kappa^\phi B_\lambda(s_{u,k'+2\phi} \cup \{b', b\})]. \end{aligned} \quad (27)$$

Since b is the last h -th SVCs to be downloaded, the following condition holds:

$$1 + \kappa^\phi B_\lambda(s_{u,k'+2\phi} \cup \{b'\}) \geq \kappa^\phi B_\lambda(s_{u,k'+2\phi} \cup \{b', b\}) \quad (28)$$

where the RHS is equal to $\kappa^\phi [1 + \kappa^\phi B_\lambda(s_{u,k'+2\phi} \cup \{b'\})]$. If $b^* = b'$, it is straightforward that the LHS is greater than the RHS. If $b^* = b$, and the SVC b' is never requested again in the subsequent time slots, state $s_{u,k'+2\phi} \cup \{b'\}$ is equivalent to $s_{u,k'+2\phi}$. The condition that the value of the LHS is not less than the value of RHS still holds. If SVC b' will be requested again in the subsequent time slots, for both state $s_{u,k'+2\phi} \cup \{b\}$ in the LHS and state $s_{u,k'+2\phi} \cup \{b'\}$ in the RHS, there are h SVCs left to be downloaded, and $B_\lambda(s_{u,k'+2\phi} \cup \{b\}) = B_\lambda(s_{u,k'+2\phi} \cup \{b'\})$. In such a case, the value of the LHS is no less than the value of the RHS. Therefore, when the content delivery delay is the same for different requests, the indexability of the problem holds.

APPENDIX F

PROOF OF LEMMA 3

The distribution of request deadline $N(X_0)$ for a VR headset can be obtained by a Markov chain on viewpoint movement, as shown in Fig. 13. Denote the probability that the spatial location of a viewpoint moves out of a tile along a direction within a time slot by p , as shown in Fig. 13(a). The viewpoint can move either up, down, left, or right from the current tile. The user's watching behavior and the video they are watching are factors that influence this probability. Furthermore, denote the probability that the spatial location of a viewpoint moves into the tiles and the viewpoint can be rendered by downloading SVC by q , where the probability depends on the number of videos downloaded in the buffer. To simplify the notations, we define index ζ as a tier in the spatial area of an SVC, as shown in Fig. 13(a), and the number of tiles in tier ζ is denoted by $M(\zeta)$. Then, we define parameter F as the maximal tier of tiles where the located viewpoints can be rendered by an SVC. Then, we

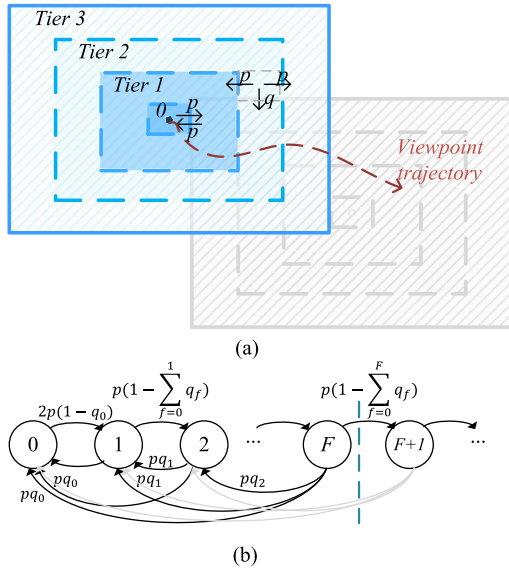


Fig. 13. (a) Illustration of viewpoint movement, where shadowed areas area are the video spatial area covered by two different videos; (b) Markov chain on viewpoint movement.

model the viewpoint movement across tires as a Markov chain as shown in Fig. 13(b), where state $V(k)$ represents the minimum tier where the viewpoint is located and the tiles are covered by downloaded SVCs in time slot k . The transition probabilities between states are defined by

$$P(1|0) = 4p(1 - q_0); P(0|0) = 1 - 4p(1 - q_0); \quad (29a)$$

$$P(V(k)+1|V(k)) = p \left(1 - \sum_{t=0}^{\beta} q_{\zeta} \right), V(k) > 0; \quad (29b)$$

$$P(V(k)-1|V(k)) = p \left(1 - \sum_{t=0}^{\beta-1} q_{\zeta} + 4q_{V(k)-1} \right), V(k) > 0; \quad (29c)$$

$$P(\zeta|V(k)) = 4pq_{\zeta}, 1 < \zeta \leq \beta, V(k) > 0; \quad (29d)$$


$$P(V(k)|V(k)) = 1 - 2p \left(1 + \sum_{\zeta=0}^{\beta-1} q_{\zeta} \right) + pq_{V(k)}, V(k) > 0 \quad (29e)$$

where $q_{\zeta} = qM(\zeta) / \sum_{\zeta'=0}^F M(\zeta')$, and $\beta = \min\{V(k), F\}$. The probability mass function of $N(X_0)$ can be obtained from evaluating the number of times that state $V(k) > F$ within $(W + H)$ time slots, given $V(0) = 0$. As shown in (29e), states are more likely to transit to the states with higher values than F when viewpoints move faster, i.e., higher p . When X_0 increases, F increases since more viewpoints can be rendered by the same SVCs, and $\bar{N}(X_0)$ decreases accordingly. We can tentatively decrease X_0 whenever the frame missing rate is higher than a requirement.

REFERENCES

- [1] Q. Zhang, J. Wei, S. Wang, S. Ma, and W. Gao, "RealVR: Efficient, economical, and quality-of-experience-driven VR video system based on MPEG OMAF," *IEEE Trans. Multimedia*, early access, Jul. 14, 2022, doi: 10.1109/TMM.2022.3190697.
- [2] X. Shen et al., "Toward immersive communications in 6G," *Front. Comput. Sci.*, vol. 4, pp. 1–45, 2023.
- [3] P. Maniatis, E. Bourtsoulatzis, and N. Thomos, "Tile-based joint caching and delivery of 360° videos in heterogeneous networks," *IEEE Trans. Multimedia*, vol. 22, pp. 2382–2395, 2020.
- [4] J. Xie, R. Girshick, and A. Farhadi, "Deep3D: Fully automatic 2D-to-3D video conversion with deep convolutional neural networks," in *Proc. 14th Eur. Conf. Comput. Vis.*, 2016, pp. 842–857.
- [5] K. Boos, D. Chu, and E. Cuervo, "FlashBack: Immersive virtual reality on mobile devices via rendering memoization," in *Proc. 14th Annu. Int. Conf. Mobile Syst., Appl., Serv.*, 2016, pp. 291–304.
- [6] J. Dai, Z. Zhang, S. Mao, and D. Liu, "A view synthesis-based 360° caching system over MEC-enabled C-RAN," *IEEE Tran. Circuits Syst. Video Technol.*, vol. 30, no. 10, pp. 3843–3855, Oct. 2020.
- [7] Y. Liu, J. Liu, A. Argyriou, and S. Ci, "MEC-assisted panoramic VR video streaming over millimeter wave mobile networks," *IEEE Trans. Multimedia*, vol. 21, pp. 1302–1316, 2019.
- [8] J. Gao, S. Zhang, L. Zhao, and X. Shen, "The design of dynamic probabilistic caching with time-varying content popularity," *IEEE Trans. Mob. Comput.*, vol. 20, no. 4, pp. 1672–1684, Apr. 2021.
- [9] Y. Sun, Z. Chen, M. Tao, and H. Liu, "Communications, caching, and computing for mobile virtual reality: Modeling and tradeoff," *IEEE Trans. Commun.*, vol. 67, no. 11, pp. 7573–7586, Nov. 2019.
- [10] P. Yang, T. Q. Quek, J. Chen, C. You, and X. Cao, "Feeling of presence maximization: MmWave-enabled virtual reality meets deep reinforcement learning," *IEEE Trans. Wireless Commun.*, vol. 21, no. 11, pp. 10005–10019, Nov. 2022.
- [11] J. Zou, C. Li, C. Liu, Q. Yang, H. Xiong, and E. Steinbach, "Probabilistic tile visibility-based server-side rate adaptation for adaptive 360-degree video streaming," *IEEE J. Sel. Topics Signal Process.*, vol. 14, no. 1, pp. 161–176, Jan. 2020.
- [12] M. Li, J. Gao, C. Zhou, X. Shen, and W. Zhuang, "Adaptive mobile VR content delivery for industrial 5.0," in *Proc. 1st Workshop Digit. Twin Edge AI Ind. IoT*, 2022, pp. 1–6.
- [13] T. Dang and M. Peng, "Joint radio communication, caching, and computing design for mobile virtual reality delivery in fog radio access networks," *IEEE J. Sel. Areas Commun.*, vol. 37, no. 7, pp. 1594–1607, Jul. 2019.
- [14] Y. Liu, J. Liu, A. Argyriou, L. Wang, and Z. Xu, "Rendering-aware VR video caching over multi-cell MEC networks," *IEEE Trans. Veh. Technol.*, vol. 70, no. 3, pp. 2728–2742, Mar. 2021.
- [15] P. Maniatis and N. Thomos, "Tile-based edge caching for 360° live video streaming," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 31, no. 12, pp. 4938–4950, Dec. 2021.
- [16] G. Papaioannou and I. Koutsopoulos, "Tile-based caching optimization for 360 videos," in *Proc. 20th ACM Int. Symp. Mobile Ad Hoc Netw. Comput.*, 2019, pp. 171–180.
- [17] R. Yao, T. Heath, A. Davies, T. Forsyth, N. Mitchell, and P. Hoberman, "Oculus VR best practices guide," 2017. Accessed: Jul. 17, 2022. [Online]. Available: <http://developer.oculusvr.com/best-practices>
- [18] J. He, M. A. Qureshi, L. Qiu, J. Li, F. Li, and L. Han, "Rubiks: Practical 360-degree streaming for smartphones," in *Proc. 16th Annu. Int. Conf. Mobile Syst., Appl., Serv.*, 2018, pp. 482–494.
- [19] A. T. Nasrabadi, A. Samiei, and R. Prakash, "Viewport prediction for 360° videos: A clustering approach," in *Proc. 30th ACM Workshop Netw. Operating Syst. Support Dig. Audio Video*, 2020, pp. 34–39.
- [20] X. Wei, C. Yang, and S. Han, "Prediction, communication, and computing duration optimization for VR video streaming," *IEEE Trans. Commun.*, vol. 69, no. 3, pp. 1947–1959, Mar. 2021.
- [21] C. Perfecto, M. S. Elbamby, J. D. Ser, and M. Bennis, "Taming the latency in multi-user VR 360: A QoE-aware deep learning-aided multicast framework," *IEEE Trans. Commun.*, vol. 68, no. 4, pp. 2491–2508, Apr. 2020.
- [22] X. Liu and Y. Deng, "Learning-based prediction, rendering and association optimization for MEC-enabled wireless virtual reality (VR) networks," *IEEE Trans. Wireless Commun.*, vol. 20, no. 10, pp. 6356–6370, Oct. 2021.
- [23] J. V. d.M. T. HooftS. VegaT. PetrangeliWauters, and F. D. Turck, "Tile-based adaptive streaming for virtual reality video," *ACM Trans. Multimedia Comput., Commun., Appl.*, vol. 15, no. 4, pp. 1–24, 2019.

- [24] M. Hu, L. Wang, B. Tan, and S. Jin, "Two-tier 360-degree video delivery control in multiuser immersive communications systems," *IEEE Trans. Veh. Technol.*, vol. 72, no. 3, pp. 4119–4123, Mar. 2023.
- [25] H. Yuan, S. Zhao, J. Hou, X. Wei, and S. Kwong, "Spatial and temporal consistency-aware dynamic adaptive streaming for 360-degree videos," *IEEE J. Sel. Topics Signal Process.*, vol. 14, no. 1, pp. 177–193, Jan. 2020.
- [26] S. Shi, V. Gupta, M. Hwang, and R. Jana, "Mobile VR on edge cloud: A latency-driven design," in *Proc. 10th ACM Multimedia Syst. Conf.*, 2019, pp. 222–231.
- [27] Y. Zhou, L. Tian, C. Zhu, X. Jin, and Y. Sun, "Video coding optimization for virtual reality 360-degree source," *IEEE J. Sel. Topics Signal Process.*, vol. 14, no. 1, pp. 118–129, Jan. 2020.
- [28] C. Wu, Z. Tan, Z. Wang, and S. Yang, "A dataset for exploring user behaviors in VR spherical video streaming," in *Proc. 8th ACM Multimedia Syst. Conf.*, 2017, pp. 193–198.
- [29] S. Khuller, A. Moss, and J. S. Naor, "The budgeted maximum coverage problem," *Inf. Process. Lett.*, vol. 70, no. 1, pp. 39–45, 1999.
- [30] K. Shanmugam, N. Golrezaei, A. G. Dimakis, A. F. Molisch, and G. Caire, "FemtoCaching: Wireless content delivery through distributed caching helpers," *IEEE Trans. Inf. Theory*, vol. 59, no. 12, pp. 8402–8413, Dec. 2013.
- [31] P. Whittle, "Restless bandits: Activity allocation in a changing world," *J. Appl. probability*, vol. 25, pp. 287–298, 1988.
- [32] K. Avrachenkov and V. Borkar, "Whittle index based Q-learning for restless bandits with average reward," *Automatica*, vol. 139, 2022, Art. no. 110186.
- [33] G. França, D. P. Robinson, and R. Vidal, "A nonsmooth dynamical systems perspective on accelerated extensions of ADMM," *IEEE Trans. Autom. Control* vol. 68, no. 5, pp. 2966–2978, May 2018.
- [34] K. Liu and Q. Zhao, "Indexability of restless bandit problems and optimality of whittle index for dynamic multichannel access," *IEEE Trans. Inf. Theory*, vol. 56, no. 11, pp. 5547–5567, Nov. 2010.



Mushu Li (Member, IEEE) received the Ph.D. degree in electrical and computer engineering from the University of Waterloo, Waterloo, ON, Canada, in 2021. She is currently a Postdoctoral Fellow with Toronto Metropolitan University, Toronto, ON. From 2021 to 2022, she was a Postdoctoral Fellow with the University of Waterloo. Her research interests include mobile edge computing, the system optimization in wireless networks, and machine learning-assisted network management. She was the recipient of the Natural Science and Engineering Research Council of Canada

Postdoctoral Fellowship, in 2022, NSERC Canada Graduate Scholarship, in 2018, and Ontario Graduate Scholarship, in 2015 and 2016.



Jie Gao (Senior Member, IEEE) received the M.Sc. and Ph.D. degrees in electrical engineering from the University of Alberta, Edmonton, AB, Canada, in 2009 and 2014, respectively. He was a Postdoctoral Fellow with Ryerson University, Toronto, ON, Canada, from 2017 to 2019 and a Research Associate with the University of Waterloo, Waterloo, ON, from 2019 to 2020. From 2020 to 2022, he was an Assistant Professor with the Department of Electrical and Computer Engineering, Marquette University, Milwaukee, WI, USA. He is currently an Assistant Professor with

the School of Information Technology, Carleton University, Ottawa, ON. His research interests include machine learning for communications and networking, cloud and multi-access edge computing, Internet of Things and industrial IoT solutions, and B5G/6G networks in general.



Conghao Zhou (Member, IEEE) received the B.Eng. degree from Northeastern University, Shenyang, China, in 2017, the M.Sc. degree from University of Illinois Chicago, Chicago, IL, USA, in 2018, and the Ph.D. degree in electrical and computer engineering from the University of Waterloo, Waterloo, ON, Canada, in 2022. He is currently a Postdoctoral Fellow with the University of Waterloo. His research interests include space-air-ground integrated networks, multi-tier computing, and machine learning in wireless networks.



Xuemin (Sherman) Shen (Fellow, IEEE) received the Ph.D. degree in electrical engineering from Rutgers University, New Brunswick, NJ, USA, in 1990. He is currently a University Professor with the Department of Electrical and Computer Engineering, University of Waterloo, Waterloo, ON, Canada. His research interests include network resource management, wireless network security, Internet of Things, 5G and beyond, and vehicular networks. Dr. Shen is a registered Professional Engineer of Ontario, Canada, an Engineering Institute of Canada Fellow, a Canadian Academy of Engineering Fellow, a Royal Society of Canada Fellow, a Chinese Academy of Engineering Foreign Member, and a Distinguished Lecturer of the IEEE Vehicular Technology Society and Communications Society.

He was the recipient of the Canadian Award for Telecommunications Research from the Canadian Society of Information Theory in 2021, the R.A. Fessenden Award in 2019 from IEEE, Canada, Award of Merit from the Federation of Chinese Canadian Professionals (Ontario) in 2019, James Evans Avant Garde Award in 2018 from IEEE Vehicular Technology Society, Joseph LoCicero Award in 2015, Education Award in 2017 from IEEE Communications Society, and Technical Recognition Award from Wireless Communications Technical Committee in 2019, and AHSN Technical Committee in 2013. He was also the recipient of the Excellent Graduate Supervision Award in 2006 from the University of Waterloo and the Premier's Research Excellence Award in 2003 from the Province of Ontario, Canada. He was the Technical Program Committee Chair/Co-Chair of IEEE Globecom'16, IEEE Infocom'14, IEEE VTC'10 Fall, IEEE Globecom'07, and the Chair of IEEE Communications Society Technical Committee on Wireless Communications. He is also the President of IEEE Communications Society. He was the Vice President for Technical & Educational Activities and Publications, Member-at-Large on the Board of Governors, Chair of the Distinguished Lecturer Selection Committee, Member of IEEE Fellow Selection Committee of the ComSoc. He was the Editor-in-Chief of the IEEE INTERNET OF THINGS JOURNAL, IEEE NETWORK, and *IET Communications*.



Weihua Zhuang (Fellow, IEEE) received the B.Sc. and M.Sc. degrees in electrical engineering from Dalian Marine University, Dalian, China, and the Ph.D. degree in electrical engineering from the University of New Brunswick, Fredericton, NB, Canada. She is currently a University Professor and a Tier I Canada Research Chair of Wireless Communication Networks with the University of Waterloo, Waterloo, ON, Canada. Her research interest include network architecture, algorithms and protocols, and service provisioning in future communication systems. Dr.

Zhuang was the recipient of the 2021 Women's Distinguished Career Award from IEEE Vehicular Technology Society, 2021 Technical Contribution Award in Cognitive Networks from IEEE Communications Society, 2021 R.A. Fessenden Award from IEEE Canada, and 2021 Award of Merit from the Federation of Chinese Canadian Professionals in Ontario. She was the Editor-in-Chief of the IEEE TRANSACTIONS ON VEHICULAR TECHNOLOGY from 2007 to 2013, General Co-Chair of 2021 IEEE/CIC International Conference on Communications in China, Technical Program Chair/Co-Chair of 2017/2016 IEEE VTC Fall, Technical Program Symposia Chair of 2011 IEEE Globecom, and an IEEE Communications Society Distinguished Lecturer from 2008 to 2011. She is an Elected Member of the Board of Governors and the President of the IEEE Vehicular Technology Society. She is a Fellow of the Royal Society of Canada, Canadian Academy of Engineering, and Engineering Institute of Canada.