

Asynchronous Coded Caching With Uncoded Prefetching

Hooshang Ghasemi^{ID}, *Member, IEEE*, and Aditya Ramamoorthy^{ID}, *Senior Member, IEEE*

Abstract—Coded caching is a technique that promises huge reductions in network traffic in content-delivery networks. However, the original formulation and several subsequent contributions in the area, assume that the file requests from the users are synchronized, i.e., they arrive at the server at the same time. In this work, we formulate and study the coded caching problem when the file requests from the users arrive at different times. We assume that each user also has a prescribed deadline by which they want their request to be completed. In the offline case, we assume that the server knows the arrival times before starting transmission and in the online case, the user requests are revealed to the server over time. We present a linear programming formulation for the offline case that minimizes the overall transmission rate from the server subject to the constraint that each user meets his/her deadline. While the online case is much harder, we introduce a novel heuristic for it and show that under certain conditions, with high probability the request of each user can be satisfied with her/his deadline. Our simulation results indicate that in the presence of mild asynchronism, much of the benefit of coded caching can still be leveraged.

Index Terms—Coded caching, asynchronous, deadlines, linear programming.

I. INTRODUCTION

CACHING is a core component of solving the problem of large scale content delivery over the Internet. Conventional caching typically relies on placing popular content closer to end-users. Statistically, popular content is requested more frequently and the cache can be used to serve the user requests in this case. Contacting the central server that has all the content is not needed. This serves to reduce the induced network traffic.

In their pioneering work [1], Maddah-Ali and Niesen considered the usage of coding in the caching problem. In this “coded caching” setting, there is a server containing a library of N files. There are K users each with a cache that can store up to M files. The users are connected to the server via an error-free shared broadcast link (see Fig. 1). The system operates in two distinct phases. In the *placement phase* the content of the caches is populated by the server. This phase

does not depend on the future requests of the users which are assumed to be arbitrary. In the *delivery phase* each user makes a request and the server transmits potentially coded signals to satisfy the requests of the users. The work of [1] demonstrated that significant reductions in the network traffic were possible as compared to conventional caching. Crucially, these gains continue to hold even if the popularity of the files is not taken into account.

While this is a significant result, the original formulation of the coded caching problem assumes that the user requests are synchronized, i.e., all file requests from the users arrive at the server at the same time. Henceforth, we refer to this as the synchronous setting. From a practical perspective, it is important to consider the asynchronous setting where user requests arrive at different times. In this case, a simple strategy would be to wait for the last request to arrive and then apply the scheme of [1]. Such a strategy will be quite good in terms of the overall rate of transmission from the server. However, this may be quite bad for an end user’s experience, e.g., the delay experienced by the users will essentially be dominated by the arrival time of the last request.

In this work, we formulate and study the coded caching problem when the user requests arrive at different times. Each user has a specific deadline by which his/her demand needs to be satisfied. The goal is to schedule the transmission of packets so that each user is able to recover the requested file from the transmitted packets and his/her cache content within the prescribed deadline. We present algorithms for both the offline and online versions of this problem.

This paper is organized as follows. In Section II we discuss the background and related work and overview our main contributions. The problem formulation appears in Section III. Sections IV and V discuss our work on the offline and the online versions of the problem, respectively. We conclude the paper with a discussion of opportunities for future work in Section VII.

II. BACKGROUND, RELATED WORK AND SUMMARY OF CONTRIBUTIONS

A coded caching system contains a server with N files, denoted W_n , $n = 1, \dots, N$, each of size F subfiles, where a subfile is a basic unit of storage. These subfiles are indexed as $W_{n,j}$, $j = 1, \dots, F$. The system also contains K users each connected to the server through an error free, broadcast shared link. Each of the users is equipped with a local cache. The i -th cache can store the equivalent of $M_i F$ subfiles. We denote the cache content of user i by Z_i , where Z_i is a function of $W_1, \dots, W_N$. Our formulation supports users with different cache sizes. A block diagram of a coded caching system for $N = K = 3$ is depicted in Fig. 1.

Manuscript received July 8, 2019; revised March 13, 2020 and June 16, 2020; accepted June 17, 2020; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor B. Shrader. This work was supported in part by the National Science Foundation (NSF) under Grant CCF-1718470 and Grant CCF-1910840. This article was presented in part at the 2017 IEEE International Symposium on Information Theory and in part at the 2017 Asilomar Conference on Signals, Systems, and Computers. (Corresponding author: Hooshang Ghasemi.)

Hooshang Ghasemi was with the Electrical and Computer Engineering Department, Iowa State University, Ames, IA 50011 USA. He is now with Qualcomm Inc., San Diego, CA 92121 USA (e-mail: hghasemi@qti.qualcomm.com).

Aditya Ramamoorthy is with Iowa State University, Ames, IA 50011 USA (e-mail: adityar@iastate.edu).

Digital Object Identifier 10.1109/TNET.2020.3003907

1063-6692 © 2020 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission.
See <https://www.ieee.org/publications/rights/index.html> for more information.




Fig. 1. Block diagram of the coded caching system with $N = K = 3$ where cache of each user contains $1/3$ -rd fraction of each file as shown. Users 1, 2, and 3 request file A , B , and C respectively and the server transmits three subfiles over the shared link to satisfy their demands.

In general, a user might choose to store a coded combination of the subfiles, i.e., Z_i can be a non-trivial function of $\{W_{n,j}\}_{n=1,\dots,N,j=1,\dots,F}$. However, in this work, we assume that an uncoded placement scheme is being used by the coded caching system, i.e., user i caches at most a $M_i F$ -sized subset of all the subfiles at the server, i.e., Z_i is a subset of $\{W_{n,j}\}_{n=1,\dots,N,j=1,\dots,F}$. The uncoded placement scheme was shown in [1] to have excellent performance. It has also been considered in several follow-up works as well. It is well recognized that the delivery phase in the uncoded placement case, corresponds to an index coding problem [2]. While the optimal solution for an arbitrary index coding problem is known to be hard, techniques such as clique cover on the side information graph are well-recognized to have good performance [2]. In this case, each transmitted equation from the server is such that a certain number of users “benefit” from it simultaneously. Under this assumption, we formulate and study the asynchronous coded caching problem when the file requests arrive at the server at different times. Each user specifies a deadline by which he/she expects the request to be satisfied.¹ We assume that

- the delivery phase proceeds via a clique cover and
- transmitting a single packet over the shared link takes a certain number of time slots.

We study the rate gains of coded caching under this setup, i.e., among the class of strategies that allow the users to meet their deadlines, we attempt to determine those where the server transmits the fewest number of packets. Both the offline and online versions of the problem are studied. In the offline scenario, we assume that information about all request arrival times and deadlines are known to the server before transmission, whereas in the online scenario, the arrival times and deadlines are revealed to the server as time progresses.

A. Main Contributions

- *Linear programming (LP) formulation in the offline case.* We propose an LP in the offline scenario that determines a schedule for the equations that need to be transmitted from the server. A feasible point of the LP can be interpreted as a coding solution that can be used by the server, such that each user meets its deadlines.

¹It is not too hard to see that in the absence of deadlines the server can simply wait for enough user requests to arrive before starting transmission. Thus, the deadline-free case essentially reduces to the synchronous setting. Section IV.B of [3] has more details.

The computational complexity of solving this LP can be quite high for a large number of users. Accordingly, we develop a dual decomposition technique where the dual problem decouples into a set of independent minimum cost network flow problems that can be solved efficiently [4].

- *A novel online algorithm.* For the online problem, we demonstrate that, in general, coding within subfiles of the same file is essential. Interestingly, this is not needed in the synchronous case where the transmitted signals by the server are coded combination of subfiles belonging to different files. Furthermore, we propose a novel online algorithm that is inspired by recursively solving the offline LP and interpreting the corresponding output appropriately. Under certain conditions, we also show that the algorithm will result in a solution that satisfies the deadline constraints with high probability.

For both scenarios, we present exhaustive simulation results that corroborate our findings and demonstrate the superiority of our algorithm concerning prior work. Overall, our work indicates that under mild asynchronism, much of the benefits of coded caching can still be leveraged.

B. Related Work

The area of coded caching has seen a flurry of research activity along several dimensions in recent years. From a theoretical perspective, significant work has attempted to understand the fundamental rate limits of a coded caching system [5]–[7]. Extensions of the basic model to general networks have been examined in [8]–[10]. Issues related to subpacketization (i.e., the number of subfiles F) have been considered in [11]–[13]. A high subpacketization level can cause several issues in practical implementations. Coded caching ideas have also been used within the domain of distributed computing [14]–[16].

There are relatively few prior works that have considered asynchronism within the context of coded caching. To our best knowledge, it was first studied in [17]. They considered the decentralized coded caching model [18] and a situation where each subfile has a specific deadline. Only the online case was considered and heuristics for transmission from the server were proposed. The heuristics are found to have good performance. However, the transmission time for a packet was not considered in their formulation. Reference [19] also considers the asynchronous setting; again, they do not consider the transmission time of a transmitted packet. In that sense, their setting is closer to the work of [17] and can be viewed as a set of rules that the server should follow in the online case. [19] (Section III.C) also considers an offline setting for the centralized placement scheme of [1]. Our LP formulation can be viewed as a bound on the possible performance of any online scheme. Our proposed online algorithm has significantly better performance than the ones presented in [17].

Reference [18] (Section V.C) also discusses the issue of asynchronism within the context of decentralized coded caching, without considering deadlines or packet transmission times. They advocate a further subpacketization of each subfile (referred to as a segment in [18]). It is important to note that any system will need to commit to a certain subpacketization scheme before deployment. Given this subpacketization and

with user specified deadlines, the formalism of our work and our algorithms can be used to arrive at schemes that address asynchronous requests.

The work of [20] proposes an algorithm for the online scenario under the assumption of decentralized coded caching for reducing the worst-case load of fronthaul links in fog radio access networks (F-RANs); this is a different model than ours. Their work does not take transmission time into account and considers the scenario where each user has the same deadline.

The asynchronous setting has also been considered in [21] for video delivery by taking into account an appropriately defined audience retention rate. Their work considers a probabilistic arrival model and presents a decentralized coded caching scheme for it.

III. PROBLEM FORMULATION AND PRELIMINARIES

We assume that time $\tau \geq 0$ is slotted. Let $[n]$ denote the set $\{1, \dots, n\}$ and the symbol $\oplus$ represent the XOR operation. We assume that the server contains $N \geq K$ files² denoted by $W_n, n = 1, \dots, N$. The subfiles are denoted by $W_{n,f}$ so that $W_n = \{W_{n,f} : f \in [F]\}$ and the cache of user i by $Z_i \subseteq \{W_{n,f} : n \in [N], f \in [F]\}$. Z_i contains at most $M_i F$ subfiles. In the delivery phase, user i requests file W_{d_i} , where $d_i \in [N]$, from the server. We let $\Omega^{(i)}$ denote the indices of the subfiles that are not present in the i -th user's cache, i.e.,

$$\Omega^{(i)} \triangleq \{f : f \in [F], W_{d_i,f} \notin Z_i\}.$$

The equations in the delivery phase are assumed to be of the *all-but-one* type.

Definition 1 (All-But-One Equation): Consider an equation E such that

$$E \triangleq \bigoplus_{l=1}^{\ell} W_{d_l, f_l}.$$

We say that E is of the all-but-one type if for each $l \in [\ell]$, we have $W_{d_l, f_l} \notin Z_{i_l}$ and $W_{d_l, f_l} \in Z_{i_k}$ for all $k \in [\ell] \setminus \{l\}$. It is evident that an all-but-one equation transmitted from the server allows each of the users participating in the equation to recover a missing subfile that they need. The asynchronous coded caching problem can be formulated as follows.

Inputs:

- **User requests:** User i requests file W_{d_i} , with $d_i \in [N]$ at time T_i .
- **Deadlines:** The i -th user needs to be satisfied by time $T_i + \Delta_i$, where Δ_i is a positive integer.
- **Transmission Delay:** Each subfile needs r time-slots to be transmitted over the shared link, i.e., each subfile can be treated as equivalent to r packets, where each packet can be transmitted in one time slot.

As the problem is symmetric with respect to users, w.l.o.g. we assume that $T_1 \leq T_2 \leq \dots \leq T_K$. Let $T_{\max} = \max_i (T_i + \Delta_i)$. Note that upon sorting the set of arrival times and deadlines, i.e., $\cup_{i=1}^K \{T_i, T_i + \Delta_i\}$, we can divide the interval $[T_1, T_{\max})$ into at most $2K - 1$ non-overlapping intervals. Let the integer β , where $1 \leq \beta \leq 2K - 1$ denote the number of intervals. Let $\Pi_1, \dots, \Pi_\beta$ represent the intervals where Π_i appears before Π_j if $i < j$; $|\Pi_\ell|$ denotes the length




Fig. 2. Offline solution corresponding to the Example 1 (system with $N = K = 3$). The double-headed arrows show the active time slots for each user. The transmitted equations are shown above the timeline.

of interval Π_ℓ . The intervals are left-closed and right-open. An easy to see but very useful property of the intervals that we have defined is that for a given i , either $[T_i, T_i + \Delta_i) \cap \Pi_\ell = \Pi_\ell$ or $[T_i, T_i + \Delta_i) \cap \Pi_\ell = \emptyset$. Fig. 2 shows an example when $K = 3$. We define $U_\ell \triangleq \{i \in [K] : [T_i, T_i + \Delta_i) \cap \Pi_\ell = \Pi_\ell\}$, and $D_\ell \triangleq \{d_i \in [N] : i \in U_\ell\}$. Thus, U_ℓ is the set of active users in time interval Π_ℓ and D_ℓ is the corresponding set of active file requests.

Outputs:

- **Transmissions at Each Time Slot:** If the problem is feasible, the schedule specifies which equations (of the all-but-one type) need to be transmitted at each time. The schedule is such that each user can recover all its missing subfiles within its deadline. The equations transmitted at time $\tau \in \Pi_\ell$ only depend on D_ℓ .

We consider two versions of the above problem.

- **Offline Version:** In the offline version, we assume that the server is aware of $\{T_i, \Delta_i, d_i\}_{i=1}^K$ at $\tau = 0$. However, at time $\tau \in \Pi_\ell$ the transmitted equation(s) will only depend on D_ℓ , i.e., the server cannot start sending missing subfiles for a given user until its request arrives.
- **Online Version:** In the online version, information about the file requests are revealed to the server as time progresses. At each time τ , the server only has information about $\{T_i, \Delta_i, d_i\}$ if $T_i \leq \tau$, i.e., the requests that have arrived by time τ .

We begin by defining some relevant sets; for convenience, a tabulated list of most of the items needed in the subsequent sections can be found in Table I. Consider a subset of users $U \subseteq [K]$. For each user $i \in U$ we let $\mathcal{F}_{\{i,U\}}$ denote the indices of all missing subfiles of the i -th user that have been stored in the cache of the other users in U , i.e.,

$$\mathcal{F}_{\{i,U\}} \triangleq \left\{ f \in \Omega^{(i)} : W_{d_i,f} \in Z_j \text{ for all } j \in U \setminus \{i\} \right\}.$$

Definition 2: User Group. A subset $U \subseteq [K]$ is said to be a user group if $\mathcal{F}_{\{i,U\}} \neq \emptyset$ for all users $i \in U$ so that there is at least one all-but-one type equation associated with U .

For a user group U there are $\prod_{i \in U} |\mathcal{F}_{\{i,U\}}|$ different all-but-one equations. This is because for any choice of $f_i \in \mathcal{F}_{\{i,U\}}$ for $i \in U$, we can construct the all-but-one equation $\bigoplus_{i \in U} W_{d_i, f_i}$. Thus, for each $i \in U$ there are $|\mathcal{F}_{\{i,U\}}|$ choices for f_i . Recall that U_ℓ is the set of active users in time interval Π_ℓ and D_ℓ represents their file requests. Let $\mathcal{U}_\ell$ be a subset of the power set of U_ℓ (i.e. the set of all subsets of U_ℓ) such that each element in $\mathcal{U}_\ell$ is a user group (cf. Definition 2). For any $U \subseteq [K]$, let $\mathcal{I}_U$ be the set of indices of all time intervals where the users in U are simultaneously active, i.e.,

$$\mathcal{I}_U \triangleq \left\{ \ell : [T_i, T_i + \Delta_i) \cap \Pi_\ell = \Pi_\ell, \forall i \in U \right\}.$$

²We assume that $N \geq K$ as it corresponds to the worst case rate (under most reasonable placement schemes) where each of the K users can request a different file. Furthermore, it is also the more practical scenario.

For each missing subfile $W_{\{d_i, f\}}$ (where $f \in \Omega^{(i)}$) we let $\mathcal{U}_{\{i, f\}}$ denote the set of user groups where it can be transmitted, i.e.,

$$\mathcal{U}_{\{i, f\}} \triangleq \left\{ U \in \cup_{\ell=1}^{\beta} \mathcal{U}_{\ell} : i \in U, f \in \mathcal{F}_{\{i, U\}} \right\}.$$

We note here that for a fixed i , there are potentially multiple indices $f_1, f_2, \dots, f_l \in \Omega^{(i)}$ such that $U \in \mathcal{U}_{\{i, f_j\}}$ for $j = 1, \dots, l$.

Example 1: Consider a system (shown in Fig. 2) with $N = 3$ files, W_1, W_2 , and W_3 where each file is divided into three subfiles, so that $F = 3$. There are $K = 3$ users with the following cache content, $Z_1 = \{W_{2,1}, W_{2,2}, W_{3,3}\}$, $Z_2 = \{W_{1,1}, W_{2,3}, W_{3,1}\}$, and $Z_3 = \{W_{2,2}, W_{3,2}, W_{2,1}\}$. Thus, $M_i = 1$ for $i \in [K]$. The arrival times are $T_1 = 0$, $T_2 = 1$, $T_3 = 3$, and deadlines are $\Delta_1 = 5$, $\Delta_2 = 4$, and $\Delta_3 = 2$. The i -th user requests file W_i , for $i = 1, \dots, 3$. Therefore, $\Omega^{(1)} = \{1, 2, 3\}$, $\Omega^{(2)} = \{1, 2\}$, and $\Omega^{(3)} = \{1, 3\}$.

In this system we have $\mathcal{F}_{\{1, \{1,2\}\}} = \{1\}$ as $W_{1,1} \in Z_2$, and $\mathcal{F}_{\{2, \{1,2\}\}} = \{1, 2\}$ as $W_{2,1}, W_{2,2} \in Z_1$. Therefore, $\{1, 2\}$ is an user group and the corresponding all-but-one equations are $W_{1,1} \oplus W_{2,1}$ and $W_{1,1} \oplus W_{2,2}$. However, $\mathcal{F}_{\{1, \{1,3\}\}} = \emptyset$ thus $\{1, 3\}$ is not an user group.

As $U = \{1, 2\}$ is an user group, we have $\mathcal{U}_2 = \{\{1\}, \{2\}, \{1, 2\}\}$. The set of time intervals where user group $\{1, 2\}$ is active is $\mathcal{I}_{\{1,2\}} = \{2, 3\}$. Finally, note that user group $U = \{2, 3\}$ is a member of $\mathcal{U}_{\{2,1\}}$ since $2 \in U$ and $1 \in \mathcal{F}_{\{2,U\}} = \{1, 2\}$. Similarly, $U \in \mathcal{U}_{\{2,2\}}$ as well since $2 \in U$ and $2 \in \mathcal{F}_{\{2,U\}}$.

IV. OFFLINE ASYNCHRONOUS CODED CACHING

In this section, we discuss the offline version of the problem where the server knows the arrival times/deadlines of all the requests at $\tau = 0$. The offline solution of the system in Example 1 is depicted in Fig. 2 where the transmitted equation in each time slot appears above the timeline. It can be verified that each user can recover the missing subfiles that they need. In what follows we argue that the offline setting can be cast as a linear programming problem.

A. Linear Programming Formulation

For each time interval Π_{ℓ} with $\ell = 1, \dots, \beta$ and for each $U \in \mathcal{U}_{\ell}$, we define variable $x_U(\ell) \in [0, |\Pi_{\ell}|]$ that represents the portion of time interval Π_{ℓ} that is allocated to an equation that benefits user group U . The actual equation will be determined shortly. For each missing subfile $W_{\{d_i, f\}}$ and each $U \in \mathcal{U}_{\{i, f\}}$, we define variable $y_{\{i, f\}}(U) \in [0, r]$ that represents the portion of the missing subfile $W_{\{d_i, f\}}$ transmitted within some or all of the equations associated with $x_U(\ell)$ for $\ell \in \mathcal{I}_U$. As pointed out before, for a fixed i , U can be used to transmit different missing subfiles needed by user i . However, a single equation can only help recover one missing subfile needed by i . Thus, $\sum_{\ell \in \mathcal{I}_U} x_U(\ell)$ must be shared between the appropriate $y_{\{i, f\}}(U)$'s. Accordingly, we need the following constraint for user i and a user group U which contains i .

$$\sum_{f \in \mathcal{F}_{\{i, U\}}} y_{\{i, f\}}(U) \leq \sum_{\ell \in \mathcal{I}_U} x_U(\ell).$$

TABLE I
LIST OF VARIABLES USED IN THE DESCRIPTION

Variable	Description
T_i	arrival time of user i
$T_i + \Delta_i$	deadline of user i
β	number of time intervals
Π_{ℓ}	time interval ℓ
$\mathcal{U}_{\ell}$	set of the active users in time interval ℓ
$\Omega^{(i)}$	set of the indices of missing subfiles of user i
$\mathcal{U}_{\ell}$	set of all subsets of $\mathcal{U}_{\ell}$ that are user groups
$\mathcal{I}_U$	set of the indices $\ell \in [\beta]$ so that $U \in \mathcal{U}_{\ell}$
$\mathcal{U}_{\{i, f\}}$	set of all user groups that $W_{d_i, f}$ can be transmitted within
$x_U(\ell)$	portion of Π_{ℓ} allocated to user group U
$y_{\{i, f\}}(U)$	portion of $W_{d_i, f}$ transmitted within user group U
$\mathcal{F}_{\{i, U\}}$	set of the indices of $f \in \Omega^{(i)}$ that can be transmitted within U

In addition, at time interval Π_{ℓ} at most $|\Pi_{\ell}|$ packets can be transmitted, so that $\sum_{U \subseteq \mathcal{U}_{\ell}} x_U(\ell) \leq |\Pi_{\ell}|$. To ensure that each missing subfile $W_{\{d_i, f\}}$ is transmitted in exactly r time slots, we have $\sum_{U \in \mathcal{U}_{\{i, f\}}} y_{\{i, f\}}(U) = r$.

The following LP minimizes the overall rate of transmission from the server while respecting all the deadline constraints of the users under the assumption that the server only transmits all-but-one equations. However, we point out that in general this may not be the information-theoretically optimal strategy for the server.

$$\begin{aligned}
& \min_{\{x_U(\ell), y_{\{i, f\}}(U)\}} \sum_{\ell=1}^{\beta} \sum_{U \in \mathcal{U}_{\ell}} x_U(\ell) \\
& \text{s.t.} \quad \sum_{U \in \mathcal{U}_{\ell}} x_U(\ell) \leq |\Pi_{\ell}|, \quad \text{for } \ell = 1, \dots, \beta, \\
& \quad \sum_{f \in \mathcal{F}_{\{i, U\}}} y_{\{i, f\}}(U) \leq \sum_{\ell \in \mathcal{I}_U} x_U(\ell), \quad \text{for } i \in U, \\
& \quad U \in \cup_{\ell=1}^{\beta} \mathcal{U}_{\ell}, \\
& \quad \sum_{U \in \mathcal{U}_{\{i, f\}}} y_{\{i, f\}}(U) = r, \quad \text{for } f \in \Omega^{(i)}, i \in [K], \\
& \quad x_U(\ell), y_{\{i, f\}}(U) \geq 0, \quad \text{for } \forall i \in [K], \ell \in [\beta], \\
& \quad U \in \cup_{\ell=1}^{\beta} \mathcal{U}_{\ell}.
\end{aligned} \tag{1}$$

Note that [17] considers the case when each missing subfile has a prescribed deadline. Our LP above can be modified in a straightforward manner to incorporate this aspect.

B. Interpretation of Feasible Point of (1) as a Coding Solution

We start by assigning time intervals to user groups. The time interval Π_{ℓ} , $\ell \in [\beta]$, will be arbitrarily assigned to user groups $U \in \mathcal{U}_{\ell}$ so that the time assigned to one user group does not overlap with another. The constraint $\sum_{U \in \mathcal{U}_{\ell}} x_U(\ell) \leq |\Pi_{\ell}|$ implies that such an assignment exists. For each user group U and each $i \in U$, suppose that $f_1, \dots, f_l \in \mathcal{F}_{\{i, U\}}$ are such that $y_{\{i, f_j\}}(U) \neq 0$ for $j = 1, \dots, l$. We assign $y_{\{i, f_j\}}(U)$ part of the total time allocated to user group U , i.e., $\sum_{\ell \in \mathcal{I}_U} x_U(\ell)$, to the missing subfile W_{d_i, f_j} for $j = 1, \dots, l$. The constraint $\sum_{f \in \mathcal{F}_{\{i, U\}}} y_{\{i, f\}}(U) \leq \sum_{\ell \in \mathcal{I}_U} x_U(\ell)$ ensures that such an assignment always exists, i.e., it is possible to assign




Fig. 3. Interpretation of feasible point in (1) for Example 1. For readability, only equations corresponding to user groups $\{1, 2\}$ and $\{2, 3\}$ are depicted.

$y_{\{i,f\}}(U)$'s (for fixed i) to the available (strictly) positive $x_U(\ell)$'s, such that there is no overlap between them. This assignment is not unique in general. However, this is not a problem as any assignment can be used to determine the equations. This process is repeated for all users $i \in U$.

The equation transmitted on a particular interval is simply the XOR of the subfile indices that map to that interval. This equation is valid since the missing subfile $W_{d_i,f}$ with $f \in \mathcal{F}_{\{i,U\}}$ is in the cache of all the users in $U \setminus \{i\}$.

Finally, according to the constraint $\sum_{U \in \mathcal{U}_{\{i,f\}}} y_{\{i,f\}}(U) = r$, each missing subfile $W_{d_i,f}$ is transmitted in its entirety in some equations. The following example serves to illustrate the arguments above.

Example 2: Consider again the system in Example 1. Part of a feasible solution to the LP in (1), corresponding to user groups $U = \{2, 3\}$ and $U' = \{1, 2\}$, is presented below.

$$\begin{aligned} x_{\{1,2\}}(2) &= 0.5, & x_{\{1,2\}}(3) &= 0.5, & x_{\{2,3\}}(3) &= 1, \\ y_{\{1,1\}}(\{1,2\}) &= 1, & y_{\{2,1\}}(\{1,2\}) &= 0.5, \\ y_{\{2,2\}}(\{1,2\}) &= 0.5, \\ y_{\{2,1\}}(\{2,3\}) &= 0.5, & y_{\{2,2\}}(\{2,3\}) &= 0.5, \\ y_{\{3,1\}}(\{2,3\}) &= 1. \end{aligned}$$

According to the solution, $x_{\{2,3\}}(3) = 1$. Therefore, only one unit of Π_3 is assigned to U (though $|\Pi_3| = 2$). This is denoted by the light blue color line in Fig. 3. For user 3 $\in U$, there is only one missing subfile in $\mathcal{F}_{\{3,\{2,3\}\}}$, namely $W_{3,1}$. As $y_{\{3,1\}}(\{2,3\}) = 1$ it is assigned to $x_{\{2,3\}}(3)$ in its entirety. This is depicted by the gray line in Fig. 3. For user 2 in U we have $\mathcal{F}_{\{2,\{2,3\}\}} = \{1, 2\}$. The solution specifies $y_{\{2,1\}}(\{2,3\}) = y_{\{2,2\}}(\{2,3\}) = 0.5$. Thus, we assign the first half of $x_{\{2,3\}}(3)$ to missing subfile $W_{2,1}$ and the second half to $W_{2,2}$ (see the dark blue and dotted dark blue lines in Fig. 3). Accordingly, the server transmits equations such that the first half of the time interval assigned to user group U corresponds to the $E_2 = W_{2,1} \oplus W_{3,1}$ whereas the second half corresponds to $E_3 = W_{2,2} \oplus W_{3,1}$. The interpretation of the user group U' is similar (see Fig. 3).

Remark 1: The output of the above LP will typically result in a fractional solution for the variables. A fractional solution can be interpreted by assuming that each packet that is transmitted over the shared edge can be subdivided as finely as needed. Thus, in each time slot, we could transmit multiple equations that may serve potentially different subsets of users. This assumption is reasonable if the underlying subfiles and hence the packets are quite large. In any case, the above LP provides a lower bound on the performance of a solution where integrality constraints are enforced.

Remark 2: We note that for the offline solution, within a given time interval, the user groups can be assigned in any order according to the $x_U(\ell)$'s as long as they don't overlap. Moreover, the assignment of $y_{i,f}(U)$'s is also arbitrary as long as the constraints of the LP are respected. However, for the online case (cf. Section V), the order does matter since we make the best effort decision on each individual slot as we do not know the future arrivals.

C. Modified LP With Fixed User Group Assignments

Note that the LP in (1) includes the variables $x_U(\ell)$'s that determine the user groups in the different time-intervals. Suppose instead that at a certain time τ we are given the total time allocated to user group U thus far (denoted by $\tilde{z}_U$) and only need to determine the $y_{\{i,f\}}(U)$ values for each i . Let $\tilde{\mathcal{U}}$ be the set of user groups used until time τ . For user $i \in [K]$, this can be written as a related LP as shown below which returns the total number of missing packets that user i has obtained until time τ .

$$\begin{aligned} \max_{\{y_{i,f}(U)\}} & \sum_{U \in \tilde{\mathcal{U}}} \sum_{f \in \mathcal{F}_{i,U}} y_{i,f}(U) \\ \text{s.t.} & \sum_{f \in \mathcal{F}_{i,U}} y_{\{i,f\}}(U) \leq \tilde{z}_U, \text{ for } i \in U, U \in \tilde{\mathcal{U}}, \\ & \sum_{U \in \mathcal{U}_{\{i,f\}}} y_{\{i,f\}}(U) \leq r, \text{ for } f \in \Omega^{(i)}, \\ & y_{\{i,f\}}(U) \geq 0, \text{ for } U \in \tilde{\mathcal{U}}. \end{aligned} \quad (2)$$

This follows since given the z_U 's we only need to find an assignment of the $y_{\{i,f\}}(U)$'s to the corresponding z_U 's that respect the first constraint above. Moreover, since we are not considering the entire transmission time, each missing subfile may not be transmitted in its entirety.

For instance, in Example 1 suppose that at time $\tau = 4$, we have $z_{\{1\}} = 2, z_{\{1,2\}} = 1, z_{\{2,3\}} = 1$, then the LP above in (2) for user 1 has the optimum point $y_{\{1,1\}}(\{1\}) = y_{\{1,3\}}(\{1\}) = 1, y_{\{1,2\}}(\{1,2\}) = 1$. Likewise for user 2, the LP has the optimum point $y_{\{2,1\}}(\{1,2\}) = 1$ and $y_{\{2,2\}}(\{2,3\}) = 1$.

Remark 3: The complexity of our solution in (1) does not have any dependence on arrival times T_i 's and deadlines Δ_i 's. Our formulation of the LP in terms of the intervals allows us to circumvent this potential dependence.

Nevertheless, the complexity of solving the LP does grow quite quickly (cubic) in the problem parameters (number of constraints + number of variables) [22]. Next, we discuss a solution based on dual decomposition that is much faster.

D. Dual Decomposition Based LP Solution

As it stands, the LP in (1) cannot be interpreted as a network flow. Yet, intuitively one can view the missing subfiles from each user as flowing through the user groups and getting absorbed in sinks that correspond to their valid time intervals. However, the flows corresponding to different users can be shared as the all-but-one equations allow different users to benefit from the same equation. We note here that a similar sharing of flows also occurs in the problem of minimum cost multicast with network coding [23]. The LP in (1) can,

however, be modified slightly so that the corresponding dual function is such that it can be evaluated by solving a set of *decoupled minimum cost network flow optimizations*.

1) *Decoupling Procedure*: For each user $i \in U$ the variable $x_U^{(i)}(\ell)$ represents the amount of flow corresponding to user i outgoing from user group U to time interval Π_ℓ . Evidently, this amount can't be more than $x_U(\ell)$. Therefore, we have

$$x_U^{(i)}(\ell) \leq x_U(\ell),$$

which holds for all $i \in U$ and all $U \in \mathcal{U}_\ell$, $\ell \in [\beta]$. We define $\mathcal{U}_\ell^{(i)} \subseteq \mathcal{U}_\ell$ to be the subset of possible user groups at time interval Π_ℓ that include user i , i.e., $i \in U$ for all $U \in \mathcal{U}_\ell^{(i)}$. By the flow interpretation of $x_U^{(i)}(\ell)$, we have

$$\sum_{f \in \mathcal{F}_{\{i,U\}}} y_{\{i,f\}}(U) = \sum_{\ell \in \mathcal{I}_U} x_U^{(i)}(\ell),$$

for all $U \in \cup_{\ell=1}^\beta \mathcal{U}_\ell^{(i)}$. For $i = 1, \dots, K$, let $\mathcal{C}_i$ denote the following set of constraints.

$$\begin{aligned} \sum_{f \in \mathcal{F}_{\{i,U\}}} y_{\{i,f\}}(U) &= \sum_{\ell \in \mathcal{I}_U} x_U^{(i)}(\ell), \quad \text{for } U \in \cup_{\ell=1}^\beta \mathcal{U}_\ell^{(i)}, \\ \sum_{U \in \mathcal{U}_{\{i,f\}}} y_{\{i,f\}}(U) &= r, \quad \text{for } f \in \Omega^{(i)}, \\ x_U^{(i)}(\ell), y_{\{i,f\}}(U) &\geq 0, \quad \text{for } U \in \mathcal{U}_\ell^{(i)}, \ell \in [\beta], f \in \Omega^{(i)}. \end{aligned}$$

Then, the original LP can be compactly rewritten as

$$\begin{aligned} \min \quad & \sum_{\ell=1}^\beta \sum_{U \in \mathcal{U}_\ell} x_U(\ell) \\ \text{s.t.} \quad & x_U^{(i)}(\ell) \leq x_U(\ell) \quad \text{for } U \in \mathcal{U}_\ell^{(i)}, \ell \in [\beta], i \in [K], \\ & \sum_{U \in \mathcal{U}_\ell} x_U(\ell) \leq |\Pi_\ell|, \quad \text{for } \ell \in [\beta], \\ & \mathcal{C}_1, \mathcal{C}_2, \dots, \mathcal{C}_K. \end{aligned} \quad (3)$$

It is not too hard to see that the LPs in (1) and (3) are equivalent. The only difference with (1) is the introduction of variables $x_U^{(i)}(\ell)$ (for appropriate ranges of i , U and ℓ) such that the second set of inequality constraints in (1) are replaced by equality constraints. Moreover, the original constraints are maintained by setting $x_U^{(i)}(\ell) \leq x_U(\ell)$.

By the Slater's constraint qualification condition [24], we know that if the primal LP is feasible, then strong duality holds and the primal and dual optimal values are the same. Thus, we proceed by considering the dual of the LP in (3) with respect to the constraints that involve the variables $x_U(\ell)$. The Lagrangian $\mathcal{L}(\{x_U(\ell), x_U^{(i)}(\ell), \lambda_U^{(i)}(\ell)\}_{i \in U, U \in \mathcal{U}_\ell, \ell \in [\beta]}, \{\zeta_\ell\}_{\ell \in [\beta]})$ can be expressed as

$$\begin{aligned} \mathcal{L} \triangleq & \sum_{\ell=1}^\beta \sum_{U \in \mathcal{U}_\ell} x_U(\ell) + \sum_{\ell=1}^\beta \sum_{U \in \mathcal{U}_\ell} \sum_{i \in U} \lambda_U^{(i)}(\ell) (x_U^{(i)}(\ell) - x_U(\ell)) \\ & + \sum_{\ell=1}^\beta \zeta_\ell \left(\sum_{U \in \mathcal{U}_\ell} x_U(\ell) - |\Pi_\ell| \right) \end{aligned}$$

where $\lambda_U^{(i)}(\ell)$'s and ζ_ℓ 's are nonnegative dual variables. It turns out that minimizing the Lagrangian for fixed dual variables can be simplified by defining $\gamma_U^{(i)}(\ell) = \lambda_U^{(i)}(\ell)/(1 + \zeta_\ell)$ for




Fig. 4. Min-cost flow network associated with subproblem (5) corresponding to the second user, $\mathcal{N}_2(\Gamma_2, \zeta_2, \zeta_3)$. The constraints and costs are given in the text.

$i \in U$, $U \in \mathcal{U}_\ell$, and $\ell \in [\beta]$. We define $\Gamma^{(i)} = \{\gamma_U^{(i)}(\ell), \ell \in \mathcal{I}_U, U \in \mathcal{U}_\ell^{(i)}\}$, $\mathbf{x} = \{x_U(\ell), U \in \mathcal{U}_\ell, \ell \in [\beta]\}$, and $\mathbf{x}^{(i)} = \{x_U^{(i)}(\ell), \ell \in \mathcal{I}_U, U \in \mathcal{U}_\ell^{(i)}\}$. The dual function $g(\Gamma^{(1)}, \dots, \Gamma^{(K)}, \{\zeta_\ell\}_{\ell \in [\beta]})$ is obtained by solving for

$$\begin{aligned} \min_{\mathbf{x}, \mathbf{x}^{(1)}, \dots, \mathbf{x}^{(K)}} \quad & \mathcal{L} \\ \text{s.t.} \quad & \mathcal{C}_1, \mathcal{C}_2, \dots, \mathcal{C}_K. \end{aligned}$$

It is evident that the dual function $g(\Gamma^{(1)}, \dots, \Gamma^{(K)}, \{\zeta_\ell\}_{\ell \in [\beta]})$ takes a nontrivial value only if

$$\sum_{i \in U} \gamma_U^{(i)}(\ell) = 1, \quad \forall U \in \mathcal{U}_\ell, \ell \in [\beta].$$

The evaluation of $g(\Gamma^{(1)}, \dots, \Gamma^{(K)}, \{\zeta_\ell\}_{\ell \in [\beta]})$ at a fixed set of dual variables $\Gamma^{(i)}$'s and ζ_ℓ 's can therefore be written as

$$\begin{aligned} \min_{\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(K)}} \quad & \sum_{\ell=1}^\beta \sum_{U \in \mathcal{U}_\ell} \sum_{i \in U} (1 + \zeta_\ell) \gamma_U^{(i)}(\ell) x_U^{(i)}(\ell) - \sum_{\ell=1}^\beta \zeta_\ell |\Pi_\ell| \\ \text{s.t.} \quad & \mathcal{C}_1, \mathcal{C}_2, \dots, \mathcal{C}_K. \end{aligned} \quad (4)$$

We emphasize that (4) is still a convex problem and that $\gamma_U^{(i)}(\ell), \zeta_\ell \geq 0$. Let $h_i(\Gamma_i, \{\zeta_\ell\}_{\ell \in [\beta]})$, $i \in [K]$ be

$$\begin{aligned} h_i(\Gamma_i, \{\zeta_\ell\}_{\ell \in [\beta]}) &\triangleq \min_{\mathbf{x}^{(i)}} \sum_{\ell=1}^\beta \sum_{U \in \mathcal{U}_\ell^{(i)}} (1 + \zeta_\ell) \gamma_U^{(i)}(\ell) x_U^{(i)}(\ell), \\ \text{s.t.} \quad & \mathcal{C}_i. \end{aligned} \quad (5)$$

Then, the dual function becomes

$$\begin{aligned} g(\Gamma^{(1)}, \dots, \Gamma^{(K)}, \{\zeta_\ell\}_{\ell \in [\beta]}) &= \sum_{i=1}^K h_i(\Gamma_i, \{\zeta_\ell\}_{\ell \in [\beta]}) - \sum_{\ell=1}^\beta \zeta_\ell |\Pi_\ell|, \end{aligned} \quad (6)$$

if $\sum_{i \in U} \gamma_U^{(i)}(\ell) = 1$ for all $U \in \mathcal{U}_\ell$, $\ell \in [\beta]$. We present an approach to maximize the dual function in (6) shortly.

The sub-problem in (5) for fixed Γ_i and $\{\zeta_\ell\}_{\ell \in [\beta]}$, is a standard minimum-cost flow problem. The associated flow network corresponding to user i , $i \in [K]$, depends on Γ_i and $\{\zeta_\ell\}_{\ell \in [\beta]}$ and we denote it by $\mathcal{N}_i(\Gamma_i, \{\zeta_\ell\}_{\ell \in [\beta]})$. It contains a source node s and three intermediate layers followed by a terminal node t (see Fig. 4 for an example). The nodes in the first, second, and third layer correspond to missing subfiles in $\Omega^{(i)}$, user groups in $\cup_{\ell \in [\beta]} \mathcal{U}_\ell^{(i)}$, and time intervals

$\{\Pi_\ell : \ell \in [\beta] \text{ and } i \in U_\ell\}$ respectively. The edges in $\mathcal{N}_i(\Gamma_i, \{\zeta_\ell\}_{\ell \in [\beta]})$ can be expressed as follows. There are $|\Omega^{(i)}|$ edges going from source node s to each of missing subfiles in $\Omega^{(i)}$. Also, for each $f \in \mathcal{F}_{\{i,U\}}$ there is an edge going from missing subfile node f to user group node U . Furthermore, there is an edge going from user group $U \in \cup_{\ell \in [\beta]} \mathcal{U}_\ell^{(i)}$ to time interval Π_ℓ for each $\ell \in \mathcal{I}_U$. Finally, corresponding to each time interval in $\{\Pi_\ell : \ell \in [\beta] \text{ and } i \in U_\ell\}$ there is an edge going from this time interval to the terminal node t .

In flow network $\mathcal{N}_i(\Gamma_i, \{\zeta_\ell\}_{\ell \in [\beta]})$, $i \in [K]$, a zero cost is assigned to all edges except those from the user group nodes to the time intervals. The cost of the edge between user group U and time interval Π_ℓ is $(1 + \zeta_\ell)\gamma_U^{(i)}(\ell)$. The edge between time interval Π_ℓ and the terminal node has a capacity constraint of $|\Pi_\ell|$ and the edge between the source node and a missing subfile has a capacity constraint of r ; the other edges have no capacity constraint. The variable $x_U^{(i)}(\ell)$ is the amount of flow carried by the edge from user group U to time interval Π_ℓ . The source injects a flow of value $|\Omega^{(i)}|r$ which needs to be absorbed in the terminal.

We emphasize that minimum cost network flow algorithms have been subject of much investigation [4] within the optimization literature and large scale instances can be solved very quickly. For our work, we leverage Capacity Scaling algorithms within the open-source LEMON package [25].

2) *Maximizing the Dual Function*: The dual function in (6) is concave (as it can be expressed as the pointwise infimum of a family of affine functions of the dual variables [24]). We exploit the projected subgradient method to maximize the dual function iteratively. Let $x_U^{(i)}(\ell, n-1)$ for all $i \in [K], U \in \mathcal{U}_\ell$ denote the optimal point of (5) when solved for $i \in [K]$ at the $n-1$ iteration. Let $\{\gamma_U^{(i)}(\ell, n-1), \zeta_\ell(n-1), \forall U \in \mathcal{U}_\ell^{(i)}, \ell \in [\beta], i \in [K]\}$ denote a dual feasible point of (6) at the $(n-1)$ -th iteration.

According to the subgradient method, at the n -th iteration, for $i \in [K]$, we first compute

$$\begin{aligned} \tilde{\gamma}_U^{(i)}(\ell, n) &= \gamma_U^{(i)}(\ell, n-1) + \theta_n x_U^{(i)}(\ell, n)(1 + \zeta_\ell(n-1)), \\ \tilde{\zeta}_\ell(n) &= \zeta_\ell(n-1) + \theta_n \left(\sum_{U \in \mathcal{U}_\ell} \sum_{i \in U} \gamma_U^{(i)} \right. \\ &\quad \left. \times (\ell, n) x_U^{(i)}(\ell, n) - |\Pi_\ell| \right), \end{aligned}$$

where θ_n is the step size. These intermediate variables are projected onto the feasible set and primal recovery is performed by the method of [26]. The details can be found in the Appendix A. Numerical results appear in Section VI.

V. ONLINE ASYNCHRONOUS CODED CACHING

In the online scenario, at time τ only information about the already arrived requests are known to the server, i.e., it only knows T_i, d_i and Δ_i for $i \in [K]$ such that $T_i \leq \tau$. Ideally, one would want to design an online algorithm that is guaranteed to be feasible whenever the corresponding offline version is feasible. However, this appears to be a hard problem. Specifically, routinely used algorithms such as earliest-deadline-first (EDF) do not have this property [27]. In the upcoming subsection we demonstrate that the online solution requires additional ideas from a coding standpoint.




Fig. 5. Online solution corresponding to the Example 3. Note that the server is forced to transmit $W_{1,2} \oplus W_{1,3}$ at $\tau = 1$.

A. Necessity of Coding Across Missing Subfiles of a User

Example 3: Consider a system with $N = K = 3$ and $M_i = 1$ with $Z_i = \{W_{n,i} : n \in [N]\}$ for $i \in [K]$ (also depicted in Fig. 1). The arrival times and deadlines of the users are $T_i = i$, and $\Delta_i = 2$ for $i \in [K]$ (as shown in Fig. 5). We assume that user i is interested in files W_i for $i \in [K]$ and that transmitting a subfile takes a single time slot, i.e., $r = 1$.

Suppose that the server does not code across any user's missing subfiles. At $\tau = 1$, it has the choice to transmit either $W_{1,2}$ or $W_{1,3}$. We emphasize that it has to transmit either of these as the deadline for user 1 is $T_1 + \Delta_1 = 3$. If the server transmits $W_{1,3}$, then consider the scenario where $(T_3, \Delta_3) = (2, 2)$ and $(T_2, \Delta_2) = (3, 2)$, i.e., the third user's request comes at $\tau = 2$ and the second user's request comes at $\tau = 3$. In this case, the server is forced to transmit $W_{1,2}$ at $\tau = 2$, which implies that user 3 misses its deadline. On the other hand, if the server transmits $W_{1,2}$ at $\tau = 1$, then $(T_2, \Delta_2) = (2, 2)$ and $(T_3, \Delta_3) = (3, 2)$ will cause user 2 to miss its deadline.

This issue can be circumvented if we transmit a linear combination of both $W_{1,2}$ and $W_{1,3}$ in the first time slot as shown in Fig 5. Intuitively, this is the correct strategy since transmitting $W_{1,3} \oplus W_{1,2}$ allows the server to hedge its bets against the identity of the next request arrival. This example demonstrates that coding across missing subfiles of user 1 is strictly better than the alternative. We emphasize that the synchronized model of [1] and the offline scenario do not require this.

Accordingly, for the online scenario we treat each missing subfile $W_{d_i,f}$ as an element of a large enough finite field $\mathbb{F}$. This allows us to consider linear combinations of the missing subfiles over $\mathbb{F}$. Note that any equation of the form

$$\bigoplus_{i \in U} \bigoplus_{f \in \mathcal{F}_{\{i,U\}}} \alpha_{\{i,f\}} W_{\{d_i,f\}}, \quad (7)$$

where the coefficients $\alpha_{\{i,f\}}$ belong to the field $\mathbb{F}$ and $\bigoplus$ represents $\mathbb{F}$ -addition is also an all-but-one equation from which user i can recover $\bigoplus_{f \in \mathcal{F}_{\{i,U\}}} \alpha_{\{i,f\}} W_{\{d_i,f\}}$.

B. Recursive LP Based Algorithm

In the online scenario at time τ our only decision is to transmit an equation in the time slot $[\tau, \tau + 1)$. In particular, it is possible that a request arrives at $\tau + 1$ and that can change the situation drastically. It makes intuitive sense to transmit equations that benefit a large number of users. However, we also need to take into account the deadline constraints of each user. These requirements need to be balanced. At the top level, our approach can be summarized as follows.

- *Solving a Linear Program When a User Request Arrives*: We solve an LP which is similar to (1) each time a new




Fig. 6. An illustration of arrival times and user groups associated with the already submitted equations upon time $\tau = 8$ in Example 4. Associated with each user group in each time slot, an equation has been submitted by the server at the same time slot.

user request comes into the system. This specifies a set of $x_U(\ell)$ and $y_{i,f}(U)$ variables. However, in the offline case, the ordering of the $x_U(\ell)$'s within an interval does not matter (*cf.* Remark 2). In the online case, this is no longer true. As we have no knowledge of future arrivals, it becomes important to choose the “best” user group for the time slot in which transmission needs to take place.

- *Deciding Which User Group to Pick:* Based on the $x_U(\ell)$ variables we first decide a candidate list of feasible user groups that can be chosen for transmission at each time slot. Suppose user group U is a candidate. We calculate a metric for U depending upon (i) the benefit of this equation to the participating users within U , and (ii) the stringency of the deadlines of the users in U . Our measure of stringency for user i is the ratio of the remaining number of missing packets of user i to the number of remaining time-slots for user i . If the calculated metric for U is above a pre-defined threshold then an equation corresponding to U (of the form in (7)) is transmitted.
- *Updating Variables and Continuing Recursively:* Following this, we update certain variables and the process continues for each time slot thereafter. When the next user request arrives into the system, the history of the variable assignments is used to solve a new LP (similar to (1)), and the process continues recursively.

1) Measuring the Benefit of User Group U to User i :

As we need to commit to a user group at each time instant in the online case, we first discuss how we can measure the benefit to a user $i \in U$ if user group U is chosen for transmission at a given time slot. In particular, if U has been used for transmission in the past, then the current transmission may be less beneficial to some of the users or of no benefit. We demonstrate this by means of the following example.

Example 4: Consider a system $N = K = 5$, $M_i = 2$ for all users $i \in [K]$, and $r = 1$. The placement scheme is the same as [1] so that each file is divided to $F = 10$ subfiles and each user misses 6 subfiles. The cache content and missing subfiles are specified in Table II. We assume that the current time is $\tau = 8$ and that the request of users $1, \dots, 4$ have arrived to the server. More specifically, we have $T_1 = 0, T_2 = 2, T_3 = 3, T_4 = 6$ with deadlines $\Delta_i = 15$ for all users $i \in [K]$. The server has already transmitted eight equations corresponding to the user groups depicted in Fig. 6.

Suppose that the server considers scheduling the user group $\{2, 3, 5\}$ at $\tau = 8$. We note here that users 2 and 3 have already participated in previous transmissions. Thus, it needs to be determined how beneficial this equation is to each user. Considering user 2, as shown in Table II, it can recover a linear combination of $\{W_{\{d_2,f\}}, f = 2, 8, 9\}$ (from user group $\{2, 3\}$), $W_{\{d_2,2\}}$ (from user group $\{1, 2, 3\}$) and $W_{\{d_2,8\}}$ (from user group $\{2, 3, 4\}$) from the prior equations; this in turn implies that it can also recover $W_{\{d_2,9\}}$ by solving linear

TABLE II

CACHE CONTENT AND MISSING SUBFILES OF EXAMPLE 4. CACHE CONTENTS $Z_i = \{W_{\{n,f\}}, n \in [5], f \in \text{SECOND COLUMN}\}$ AND MISSING SUBFILES $\{W_{\{d_i,f\}}, f \in \text{THIRD COLUMN}\}$

User	Cache content indices	Missing subfiles indices
1	1, 2, 3, 4	5, 6, 7, 8, 9, 10
2	1, 5, 6, 7	2, 3, 4, 8, 9, 10
3	2, 5, 8, 9	1, 3, 4, 6, 7, 10
4	3, 6, 8, 10	1, 2, 4, 5, 7, 9
5	4, 7, 9, 10	1, 2, 3, 5, 6, 8

equations. However, note the the user group $\{2, 3, 5\}$ also results in user 2 recovering $W_{\{d_2,9\}}$. Thus, from user 2's perspective, this equation is of no benefit. A similar argument shows that this user group does not benefit user 3 either.

We observe at this point that the modified LP in Section IV-C, can be used here to determine a given user's benefit. In particular, the past history of the transmission contains information on the total time allocated to a given user group U . At time τ , if we consider transmitting user group $\hat{U}$, we can add it to the set of user groups and update its time allocation. Following this, for each user the LP in (2) can be used to measure the number of subfiles that can be transmitted for it, were $\hat{U}$ to be chosen.

Let $\mathcal{U}_{\text{sent}}(\tau)$ denote the set of user groups chosen for time $\leq \tau$, $\hat{U}$ be a user group under consideration at time τ , and let $\tilde{\mathcal{U}}_{\text{sent}}(\tau) = \mathcal{U}_{\text{sent}}(\tau) \cup \{\hat{U}\}$. Let $\tilde{z}_U(\tau)$ denote the time allocated to user group $U \in \tilde{\mathcal{U}}_{\text{sent}}(\tau)$, where $\tilde{z}_{\hat{U}}(\tau)$ is incremented by one if $\hat{U} \in \tilde{\mathcal{U}}_{\text{sent}}(\tau)$, otherwise it is set to 1. Consider the following LP.

$$\begin{aligned}
 & \max_{\{\tilde{y}_{\{i,f\}}(U)\}} \sum_{U \in \tilde{\mathcal{U}}_{\text{sent}}(\tau), U \ni i} \sum_{f \in \mathcal{F}_{\{i,U\}}} \tilde{y}_{\{i,f\}}(U) \\
 \text{s.t.} \quad & \sum_{f \in \mathcal{F}_{\{i,U\}}} \tilde{y}_{\{i,f\}}(U) \leq \tilde{z}_U(\tau) \text{ for } U \in \tilde{\mathcal{U}}_{\text{sent}}(\tau), U \ni i \\
 & \sum_{U \in \mathcal{U}_{\{i,f\}} \cap \tilde{\mathcal{U}}_{\text{sent}}(\tau)} \tilde{y}_{\{i,f\}}(U) \leq r \text{ for } f \in \Omega^{(i)}, \\
 & \tilde{y}_{\{i,f\}}(U) \geq 0.
 \end{aligned} \tag{8}$$

Now suppose that we have already tracked the number of useful packets for user i until time τ . Then the above LP can be used to determine the benefit of transmitting user group $\hat{U}$ at time τ .

Remark 4: The LP in (8) can also be expressed as a maximum flow problem. The associated flow network consists of a source node s , a node for each $f \in \Omega^{(i)}$, a node for each user group $U \in \tilde{\mathcal{U}}_{\text{sent}}(\tau)$, and a terminal node t . There are edges with capacity r going from s to each $f \in \Omega^{(i)}$ and edges from $f \in \Omega^{(i)}$ to node $U \in \tilde{\mathcal{U}}_{\text{sent}}(\tau)$ if $f \in \mathcal{F}_{\{i,U\}}$. The flow on such an edge is $\tilde{y}_{\{i,f\}}(U)$. Moreover, from each node $U \in \tilde{\mathcal{U}}_{\text{sent}}(\tau)$ to t there exist an edge of capacity $\tilde{z}_U(\tau)$. These capacity constraints model the first two inequality constraints in (8). Fig. 7 illustrates an example of this network. It is well-known that if all capacities in a flow network are integers, there exists an integral maximum flow ([28], Chapter 7). Therefore, there exists an integral solution for $\tilde{y}_{\{i,f\}}(U)$'s in (8) if $\tilde{z}_U(\tau)$'s are integers.

2) *Solving LP Upon User Arrival:* Consider a time $\tau = T_k$ when the request of the k -th user arrives at the server. We




Fig. 7. Max flow network associated with LP in (8).

let $\mathcal{U}_{\text{sent}}(\tau)$ be the set of user groups associated with the previously transmitted equations. We also let $z_U(\tau)$ be the total time allocated to equations corresponding to user group U prior to time τ . Thus, if in time interval $[\tau, \tau+1)$ the server transmits an equation that exclusively benefit users in U then $z_U(\tau+1) = z_U(\tau) + 1$ otherwise $z_U(\tau+1) = z_U(\tau)$. Time intervals $\Pi_{1,k}, \dots, \Pi_{\beta_k, k}$ are formed by the set of times in

$$\{T_k\} \cup \{T_i + \Delta_i : i \in [k], T_i + \Delta_i > T_k\}.$$

As in the offline case in (1), the sets of active users $U_{\ell, k}$, user groups $\mathcal{U}_{\ell, k}$ and $\mathcal{I}_U^{(k)}$ are defined corresponding to these time intervals, e.g., $U_{\ell, k}$ is the set of active users in $\Pi_{\ell, k}$. Moreover, $\mathcal{V}_k$ is a set of user groups that either already have been transmitted or might be transmitted after $\tau = T_k$. That is $\mathcal{V}_k = \mathcal{U}_{\text{sent}}(\tau) \cup \{\mathcal{U}_{\ell, k} : \ell \in [\beta_k]\}$. The variables $x_U(\ell)$'s and $y_{\{i, f\}}(U)$'s have the same interpretation as the offline case. With these variables, the server solves the following LP.

$$\begin{aligned} \min_{\{x_U(\ell), y_{\{i, f\}}(U)\}} & \sum_{\ell=1}^{\beta_k} \sum_{U \in \mathcal{U}_{\ell, k}} x_U(\ell) \\ \text{s.t.} & \sum_{U \in \mathcal{U}_{\ell, k}} x_U(\ell) \leq |\Pi_{\ell, k}|, \quad \text{for } \ell = 1, \dots, \beta_k \\ & \sum_{f \in \mathcal{F}_{\{i, U\}}} y_{\{i, f\}}(U) \leq \sum_{\ell \in \mathcal{I}_U^{(k)}} x_U(\ell) + z_U(T_k) \\ & \text{for } i \in U, U \in \mathcal{V}_k, \\ & \sum_{U \in \mathcal{U}_{\{i, f\}}} y_{\{i, f\}}(U) = r, \quad \text{for } f \in \Omega^{(i)}, \forall i \in \cup_{\ell=1}^{\beta_k} U_{\ell, k}, \\ & x_U(\ell), y_{\{i, f\}}(U) \geq 0, \quad \text{for } i \in [k], \ell \in [\beta], U \in \mathcal{V}_k. \end{aligned} \quad (9)$$

An important feature of time intervals $\Pi_{1, k}, \dots, \Pi_{\beta_k, k}$ is that these time intervals end at a deadline and except the first time interval $\Pi_{1, k}$ that starts with arrival time T_k , the other time intervals start with a deadline. Thus, we have $U_{\ell+1, k} \subset U_{\ell, k}$, i.e., the set of active users in interval $\Pi_{\ell+1, k}$ is a subset of the active users in interval $\Pi_{\ell, k}$ for the range of ℓ .

Next, the server creates a list of candidate user groups. Let $\{x_U^*(\ell), \forall U \in \mathcal{U}_{\ell, k}, \ell = 1, \dots, \beta_k\}$ be the solution of (9) and let $\mathcal{X}^* = \{x_U^*(\ell) : x_U^*(\ell) \geq 1\}$. The elements of $\mathcal{X}^*$ are first ordered based on time intervals. Then, among the elements with the same time interval, they are ordered based on length of user group. Therefore, for two elements $x_U^*(\ell), x_{U'}^*(\ell') \in \mathcal{X}^*$ we say $x_U^*(\ell)$ is before $x_{U'}^*(\ell')$ if $\ell < \ell'$, or if $\ell = \ell'$ and $|U| \geq |U'|$. We let $\mathcal{X}_{\text{sorted}}^*$ denote the sorted version of $\mathcal{X}^*$ using this procedure. Let $v_i(\tau)$ be the number of missing

Algorithm 1 Recursive LP Algorithm

Input: Caches Z_i for $i \in [K]$, η_0 , $\{T_i, \Delta_i\}$, for $i \in [K]$.

- 1: **Initialization:**
- 2: set $\mathcal{U}_{\text{sent}}(0) \leftarrow \emptyset$, $\mathcal{X}_{\text{off}} \leftarrow \emptyset$, $\ell_{\text{off}} \leftarrow 0$, $m \leftarrow 1$ and $k \leftarrow 1$.
- 3: set $\mathcal{M}_i = \emptyset$, and $v_i(0) = 0$ for $i = 1, \dots, K$.
- 4: **for** $\tau = 0, 1, 2, \dots, T_{\max}$ **do**
- 5: **if** $\tau = T_i + \Delta_i$ and $v_i(\tau) < r|\Omega^{(i)}|$ for some i **then**
- 6: return INFEASIBLE.
- 7: **end if**
- 8: **if** $\tau = T_i$ (a new user makes request) for some i **then**
- 9: $k = \arg \max_{i \in [K]} \tau = T_i$
- 10: Solve LP (9). Form $\mathcal{X}^*$ and then $\mathcal{X}_{\text{sorted}}^*$.
- 11: **end if**
- 12: **If** $\tau = T_i$ or $\tau = T_i + \Delta_i$ for some i **then** $\ell_{\text{off}} \leftarrow \ell_{\text{off}} + 1$.
- 13: **if** $\mathcal{X}_{\text{sorted}}^* \neq \emptyset$ **then**
- 14: Pick first in order $x_{U^*}^*(\ell) \in \mathcal{X}_{\text{sorted}}^*$ with $\eta_{U^*}(\tau) \geq \eta_0$.
- 15: Randomly select $\alpha_{\{i, f, j, m\}}$'s from $\mathbb{F}$ and send

$$\bigoplus_{i \in U} \bigoplus_{f \in \mathcal{F}_{\{i, U\}}} \bigoplus_{j=1}^r \alpha_{\{i, f, j, m\}} W_{\{d_i, f, j\}},$$
- 16: **If** $U^* \in \mathcal{U}_{\text{sent}}(\tau)$ **then** $z_{U^*}(\tau+1) \leftarrow z_{U^*}(\tau) + 1$, **otherwise** $z_{U^*}(\tau+1) = 1$ and $\mathcal{U}_{\text{sent}}(\tau+1) \leftarrow \mathcal{U}_{\text{sent}}(\tau) \cup \{U^*\}$
- 17: **If** $\tilde{x}_{U^*}(\ell_{\text{off}}) \in \mathcal{X}_{\text{off}}$ **then** $\tilde{x}_{U^*}(\ell_{\text{off}}) \leftarrow \tilde{x}_{U^*}(\ell_{\text{off}}) + 1$, **otherwise** $\tilde{x}_{U^*}(\ell_{\text{off}}) = 1$ and $\mathcal{X}_{\text{off}} \leftarrow \mathcal{X}_{\text{off}} \cup \{\tilde{x}_{U^*}(\ell_{\text{off}})\}$
- 18: Set $x_{U^*}^*(\ell) \leftarrow x_{U^*}^*(\ell) - 1$, if $x_{U^*}^*(\ell) < 1$ remove it from $\mathcal{X}_{\text{sorted}}^*$.
- 19: For all $i \in U^*$, set $v_i(\tau+1) \leftarrow w_{\{i, U^*\}}(\tau)$, set $\mathcal{M}_i \leftarrow \mathcal{M}_i \cup \{m\}$, then $m \leftarrow m + 1$
- 20: for all $U \in \mathcal{U}_{\text{sent}}(\tau) \setminus \{U^*\}$ set $z_U(\tau+1) \leftarrow z_U(\tau)$
- 21: for all $i \in [K] \setminus U^*$ set $v_i(\tau+1) \leftarrow v_i(\tau)$.
- 22: **end if**
- 23: **end for**

packets (subfiles when $r = 1$) that have been transmitted for user i until time τ ; this value is tracked in Algorithm 1.

Note that user i needs to recover $r|\Omega^{(i)}| - v_i(\tau)$ missing packets and it has $T_i + \Delta_i - \tau$ time slots to obtain them. We use the ratio of these quantities as a measure of the stringency of the deadline of user i . Let $w_{\{i, \hat{U}\}}$ denote the total number of missing packets of user i that can be communicated by the user groups chosen thus far and by picking $\hat{U}$ at time τ . The LP in (8) allows us to compute $w_{\{i, \hat{U}\}}$ and in turn $w_{\{i, \hat{U}\}}(\tau) - v_i(\tau)$. Therefore the metric $\eta_U(\tau)$ is obtained by the following weighted sum.

$$\eta_U(\tau) \triangleq \sum_{i \in U} \frac{(r|\Omega^{(i)}| - v_i(\tau))}{T_i + \Delta_i - \tau} (w_{\{i, U\}}(\tau) - v_i(\tau)).$$

At time $\tau = T_k$, the server picks the first element $x_{U^*}^*(\ell) \in \mathcal{X}_{\text{sorted}}^*$ such that $\eta_{U^*}(\tau) \geq \eta_0$ for some threshold η_0 and transmits an equation corresponding to it. Unlike the synchronous case, we choose a random linear combination of all missing packets of user i that can be transmitted by user group U .

When $r > 1$, we subdivide a missing subfile into r packets that are denoted $W_{\{d_i, f, j\}}$ for $j = 1, \dots, r$. Thus, the server

transmits

$$\bigoplus_{i \in U} \bigoplus_{f \in \mathcal{F}_{\{i,U\}}} \bigoplus_{j=1}^r \alpha_{\{i,f,j,m\}} W_{\{d_i,f,j\}},$$

at time interval $[\tau, \tau + 1)$ where m denotes the m -th equation transmitted by the server and $\alpha_{\{i,f,j,m\}}$ are chosen independently and uniformly at random from the finite field $\mathbb{F}$. If none of the elements in $x_U^*(\ell) \in \mathcal{X}_{\text{sorted}}^*$ satisfy $\eta_U(\ell) \geq \eta_0$ then nothing will be transmitted at this time interval.

If a new user request does not come at time $\tau + 1$, then the server updates the user group values and then solves (8) again to decide the user group for the time slot $[\tau + 1, \tau + 2)$. The process continues this way until the next user request comes when the LP in (9) is solved. The complete details are provided in Algorithm 1.

In general, there is no guarantee that Algorithm 1 will return a feasible schedule if the corresponding offline schedule is feasible. In that sense, Algorithm 1 can be viewed as a heuristic with good experimental performance. However, if Algorithm 1 does not return “INFEASIBLE”, we can show that a feasible solution for the corresponding offline LP can be identified. This fact coupled with usage of the Schwartz-Zippel Lemma allows us to conclude that our algorithm works with high probability if it does not return “INFEASIBLE”. The proofs of the following claim and lemma appear in Appendix B and C respectively.

Claim 1: For user requests, $\{T_i, \Delta_i, d_i\}$, where $i \in [K]$, if Algorithm 1 does not return “INFEASIBLE” then there exists a feasible integral solution for the offline LP in (1). The following lemma shows that if Algorithm 1 does not return “INFEASIBLE” then with high probability each user recovers all its missing subfiles from the transmitted equations.

Lemma 1: If Algorithm 1 does not return “INFEASIBLE” then with probability at least $\left(1 - \frac{1}{|\mathbb{F}|}\right)^{rKF}$ all requests will be satisfied within their deadline.

Thus, by choosing the field size $|\mathbb{F}|$ large enough, we can make the probability of success as large as we want. We point out that increasing the field size results in a corresponding increase in the computational requirements at the server and the user nodes.

VI. SIMULATION RESULTS AND COMPARISONS WITH PRIOR WORK

In this section we present simulation results for both the proposed offline and the online algorithms (software are available in [29]). Prior work in this area is primarily the work of [17] that presents heuristics for the online scenario. However, we note that [17] works with deadlines for subfiles and does not take into account the time required to transmit a packet. It uses intuitively plausible rules to decide the equations transmitted by the server depending on the deadlines of the users.

For both scenarios, the request arrival times $\{T_i, i \in [K]\}$ are generated according to a Poisson process with parameter λF . The arrival time is quantized to the nearest time slot. The deadlines $\Delta_i, i \in [K]$ are generated uniformly at random from the range $[\Delta_{\min}, \Delta_{\max}]$ (these values will be specified for each setting below).

TABLE III

EXECUTION TIME FOR SOLVING THE LP USING OUR APPROACH; WE RUN 1000 ITERATIONS OF SUBGRADIENT ASCENT. COLUMNS 2 & 3 INDICATE THE SIZE OF THE ASSOCIATED FLOW NETWORK. THE TABLE IS ORDERED BY THE NUMBER OF NODES IN THE FLOW NETWORK

(K, t)	No. of nodes	No. of edges	Exec. time (min)	Exec. time Orig. (min)
(100, 2)	986,161	17,643,986	8,026	—
(20, 4)	178,542	1,778,703	1065	—
(40, 2)	61,959	567,780	63	—
(20, 2)	7,542	43,507	1.9	21.9
(10, 4)	3,917	29,369	0.8	5.33
(10, 2)	915	3,866	0.08	0.03




Fig. 8. Convergence of primal recovery to the optimal solution for a system with $N = K = 20$, $r = 1$, and $t = 2$. Dashed line is the optimal value obtained by solving (1).

A. Offline Scenario Simulation

In the first set of simulations we examine the execution time of our approach for various values of (K, t) where $t = KM/N$ is an integer; the placement scheme in [1] was used. In these simulations we set $r = 1$, $\lambda = 0.4$, $F = \binom{K}{t}$, $\Delta_{\min} = \binom{K-1}{t}$, and $\Delta_{\max} = \binom{K}{t+1}$. Table III shows the details of the overall execution time and the size of the corresponding flow networks for the various instances. The last column of the table corresponds to the execution time (in MATLAB) of the LP in (1), while the second-last column corresponds to the execution time of the proposed approach above. It is evident that the proposed approach is significantly faster. In fact, memory requirements make it infeasible to even formulate the problems corresponding to the first three rows in MATLAB. Fig. 8 shows the convergence of the primal recovery procedure to the actual rate for a system with $N = K = 20$, $t = 2$, and $r = 1$. It can be observed that there is a clear convergence of the solution to the optimal value.

B. Online Scenario Simulation

For the online scenario we consider both centralized [1] and decentralized [18] placement schemes for a system with $N = K = 6$ and $M = 2$ with $\Delta_{\min} = (KM/N)F$ and $\Delta_{\max} = KF$. For each experiment we run 200 trials for generating the arrivals. For the centralized case, we use the placement scheme of [1] and the placement is fixed during each experiment. In the




Fig. 9. Centralized Placement in [1]: (a) average coding gain over all feasible offline problem instances, (b) feasibility probability of the online algorithm conditioned on feasibility of the offline problem. The placement has been fixed for all trials and at each trial a new arrival time and deadline is generated. In this simulation, we set $\eta_0 = 0.4 - \frac{0.5}{\lambda}$ and $\eta_0 = 0.8 - \frac{0.2}{\lambda}$ in Case I and II respectively.

decentralized scheme, at each trial the cache content of each user is independently and uniformly chosen as well.

For each set of generated arrivals, we first run the offline LP to check whether it is feasible. The online algorithm is run only if the offline LP is feasible. The online algorithm requires a threshold η_0 (see Section V-B). We run simulations with a low threshold (case I) and a high threshold (case II). The coding gain is defined as the ratio of the uncoded rate³ to the rate achieved by the system. Fig. 9 (a) and Fig. 10 (a) depict plots of the coding gain vs. $1/(F\lambda)$ in centralized and decentralized cases, respectively. As λ decreases, the arrivals are spaced further apart on average, and the coding gain of any scheme is expected to reduce.

The coding gain is computed by taking an average over all instances where a given scheme is feasible. For the offline scheme, this means that we take the average of all instances where it is feasible. For the online algorithm, some of the arrival patterns may result in infeasibility; these instances were not taken into account when computing the average coding gain. This explains why the coding gain of the case II sometimes appears to be higher than the offline algorithm. However, the coding gain of the case I is significantly lower, because of its low threshold.

³The uncoded rate is simply the total number of missing subfiles of all users normalized by F .




Fig. 10. Decentralized placement scheme for $N = K = 6$, $M = 2$, and $F = 100$: (a) average coding gain over all feasible offline problem instances, (b) feasibility probability of the online algorithm conditioned on feasibility of the offline problem. At each trial cache content of each user is placed randomly and uniformly. In this simulation, we set $\eta_0 = 0.4 - \frac{0.5}{\lambda}$ and $\eta_0 = 0.8 - \frac{0.2}{\lambda}$ in Case I and Case II respectively.

The feasibility probability of a scheme vs. the arrival rate is plotted in Fig. 9 (b) and Fig. 10 (b) for the centralized and decentralized placement schemes respectively. As expected the low threshold online algorithm has a very high feasibility probability ≈ 1 for a range of arrival parameters, while the high threshold algorithm has a lower feasibility probability. Note that the high threshold algorithm (when compared to the low threshold case) only transmits an equation when a large enough number of users benefit from the transmission. Thus, its feasibility probability is lower, but when it is feasible, its coding gain is much higher than the low threshold case.

For both plots, we also include the results of [17]. In this scheme feasibility and coding gain can be traded off by setting a threshold for the defined misfit function (Section III in [17]). We use this scheme by setting the threshold to zero; this is the so-called First-Fit Rule in [17]. The First-Fit rule prefers feasibility over coding gain. The setting in [17] considers a scenario where each subfile has a deadline. We have adapted their algorithm for our case. It can be observed that the feasibility probability of [17] is quite poor. Accordingly, we also plot the fraction of subfiles that meet the deadline; this is somewhat better. The coding gain numbers for [17] are also quite unreliable as the algorithm is infeasible in most cases. Thus, we do not plot it.




Fig. 11. Decentralized placement scheme with deadlines for individual subfiles for $K = N = 6$, $M = 2$, and $F = 20$: (a) average coding gain over all feasible offline problem instances, (b) feasibility probability of the online algorithm conditioned on feasibility of the offline problem. For the scheme in [17] two probabilities are reported. The first one is the probability that all requests are satisfied, and the second one is the probability that a fixed request is satisfied (lines with circle and diamond marks respectively). At each trial, the cache content of each user is populated randomly and uniformly. In this simulation, we set $\eta_0 = 0.4 - \frac{0.5}{\lambda}$ and $\eta_0 = 0.8 - \frac{0.2}{\lambda}$ in Case I and Case II respectively.

C. Scenario Where Individual Subfiles Have Deadlines

The work of [17] considers a situation where each subfile has its own deadline. This is inspired by applications such as video delivery over the Internet. We emphasize that this setting can be captured by our techniques. In particular, suppose that each user requests a set of subfiles from the server where the subfile requests arrive at different times and each subfile has a different deadline. In this case, we can treat each subfile request of user $i \in [K]$ as corresponding to a distinct virtual user whose cache content is the same as user i . However, the requests of the users are different. In this situation, each virtual user has precisely one missing subfile. Thus, the issue of coding over the corresponding subfiles does not arise.

Our setting is again one where $K = N = 6$, $M = 2$. Each file is subdivided into $F = 20$ subfiles. Arrival times and deadlines are generated similar to the previous simulations with Poisson parameters $1/(\lambda F)$ and the deadlines are randomly chosen uniformly from $[\Delta_{\min}, \Delta_{\max}]$ with $\Delta_{\min} = KMF/N$ and $\Delta_{\max} = KF$. Similar to the previous experiments we run 200 trials and at each trial, the cache content of each user is populated randomly and uniformly among all placement

schemes with cache of size MF subfiles. Thus, different users might request different number of subfiles from the server. The only difference is that here each requested subfile has its own arrival time and deadline. The results are illustrated in Fig. 11. It can be observed that our proposed approach provides significantly superior coding gain and feasibility probability as compared to the work of [17].

VII. CONCLUSIONS AND FUTURE WORK

In this work, we considered the asynchronous coded caching problem where user requests (with deadlines) arrive at the main server at different times. We considered both offline and online versions of this problem. We demonstrated that under the assumption of all-but-one equations, the offline scenario can be optimally solved by a linear program (LP). Moreover, we presented a low-complexity solution to this LP based on dual decomposition. In contrast to the synchronous case and the offline scenario, we show that the online scenario requires coding across missing subfiles of a given user. Furthermore, we present an online algorithm that leverages offline LP in a recursive fashion. Extensive simulation results indicate that our proposed algorithm significantly outperforms prior algorithms.

Our online algorithm considers the situation where there is no knowledge about future request arrival times and file identities; this corresponds to a worst-case scenario. It would be interesting to consider cases where there is statistical information available on the arrival times and file popularity and/or algorithms where these can be learned and to investigate how this knowledge can be used to further improve the performance of the algorithm. For instance, it may be possible to design better placement schemes under this knowledge. Throughout the paper, we implicitly considered the case when different users request different files. It may be interesting to adapt our techniques for the case of repeated requests.

APPENDIX

A. Quadratic Projection and Primal Recovery in Dual Decomposition

For the projection of $\tilde{\gamma}_U^{(i)}(\ell, n)$ and $\tilde{\zeta}_\ell(n)$ to the constraint space we simply set $\zeta_\ell(n) = \max(\tilde{\zeta}_\ell(n), 0)$ and $\{\gamma_U^{(i)}(\ell, n), \forall i \in U\}$ is obtained via the following quadratic optimization.

$$\min_{\{v_U^{(i)}: \forall i \in U\}} \sum_{i \in U} \left(v_U^{(i)} - \tilde{\gamma}_U^{(i)}(\ell, n) \right)^2 \quad \text{s.t.} \quad \sum_{i \in U} v_U^{(i)} = 1. \quad (10)$$

In [23, Appendix I] an algorithm has been proposed to solve (10). This solution can be explained as follows. For fixed $\ell \in [\beta]$ and for each $U \in \mathcal{U}_\ell$, we sort $\tilde{\gamma}_U^{(i)}(\ell, n)$ so that $\tilde{\gamma}_U^{(i_1)}(\ell, n) \geq \dots \geq \tilde{\gamma}_U^{(i_{|U|})}(\ell, n)$. We take $\hat{k}$ to be the minimum k such that

$$\frac{1}{k} \left(1 - \sum_{j=1}^k \tilde{\gamma}_U^{(i_j)}(\ell, n) \right) \leq -\tilde{\gamma}_U^{(i_{k+1})}(\ell, n)$$

or let $\hat{k} = |U|$ if such a k doesn't exist. Then $\gamma_U^{(i_j)}(\ell, n) = \tilde{\gamma}_U^{(i_j)}(\ell, n) + \frac{1 - \sum_{l=1}^{\hat{k}} \tilde{\gamma}_U^{(i_l)}(\ell, n)}{\hat{k}}$ if $j \in [\hat{k}]$ and zero otherwise.

The initial setting for the dual variables is chosen as $\gamma_U^{(i)}(\ell, 0) = 1/|U|$, for $i \in U$, $U \in \mathcal{U}_\ell$, $\ell \in [\beta]$, and $\zeta_\ell = 0$ for $\ell \in [\beta]$.

Primal Recovery: After solving the dual problem, the primal variables, i.e., $x_U(\ell, n)$'s, are recovered by the method of [26] whereby

$$x_U(\ell, n) = \sum_{l=1}^n \mu_l(n) \left(\max_{i \in U} x_U^{(i)}(\ell, l) \right) \quad (11)$$

where $\mu_l(n)$'s are sequence of convex combination weights for each non-negative integer n , i.e. $\sum_{l=1}^n \mu_l(n) = 1$ and $\mu_l(n) \geq 0$ for all $l = 1, \dots, n$. In [23], it has been shown that if the step size θ_n and convex combination weights $\mu_l(n)$ are chosen so that

- $\eta_{l,n} \geq \eta_{l-1,n}$ for all $l = 2, \dots, n$ and $n = 0, 1, \dots$,
- $\Delta_{\eta_n}^{\max} \rightarrow 0$ as $n \leftarrow \infty$, and
- $\eta_{1,n} \rightarrow 0$ as $n \leftarrow \infty$ and $\eta_{n,n} \leq \delta$ for all $n = 0, 1, \dots$ for some $\delta > 0$,

then $\{x_U(\ell, n)\}_{\ell \in [\beta], U \in \mathcal{U}_\ell}$ is an optimal primal solution. Here $\eta_{l,n} = \frac{\mu_l(n)}{\theta_n}$ and $\Delta_{\eta_n}^{\max} = \max_{l=2, \dots, n} \{\eta_{l,n} - \eta_{l-1,n}\}$. Some sequences for θ_n and $\mu_l(n)$ that satisfy the above conditions has been proposed by [23]. Among them we choose $\mu_l(n) = \frac{1}{n}$ and $\theta_n = n^{-\alpha}$ where $0 < \alpha < 1$. Then, the primal solution will be updated as,

$$x_U(\ell, n+1) = \frac{n}{n+1} x_U(\ell, n) + \frac{\max_{i \in U} x_U^{(i)}(\ell, n)}{n+1}. \quad (12)$$

B. Proof of Claim 1

Proof: For simplicity, we prove the claim for $r = 1$ and the proof for the general case follows directly. We will construct $x_U(\ell)$ and $y_{\{i,f\}}(U)$ variables for the offline LP from the decisions made in Algorithm 1. Note that we update the set $\mathcal{X}_{\text{off}}$ with the user groups chosen in Algorithm 1. It is not difficult to verify that for any $\tilde{x}_U(\ell) \in \mathcal{X}_{\text{off}}$ user group U is a member of $\mathcal{U}_\ell$. Moreover, the algorithm assigns integer values to $\tilde{x}_U(\ell)$. Now, for any $U \in \mathcal{U}_\ell$ in (1), we set $x_U(\ell) = \tilde{x}_U(\ell)$ if $\tilde{x}_U(\ell) \in \mathcal{X}_{\text{off}}$ and $x_U(\ell) = 0$ otherwise. Therefore, $x_U(\ell)$'s take integer values. Since at each time only one equation is transmitted in Algorithm 1, the first condition $\sum_{U \in \mathcal{U}_\ell} x_U(\ell) \leq |\Pi_\ell|$ holds for all $\ell \in [\beta]$.

For each $i \in [K]$ we define $[\tau_i, \tau_i + 1)$ to be the last time slot that user i benefits from the equation transmitted by the server. Clearly we have that $v_i(\tau_i + 1) \geq |\Omega^{(i)}|$ otherwise Algorithm 1 will be infeasible at $\tau = T_i + \Delta_i$. We let $U_{i,\text{last}}$ to be the user group associated with this equation where $i \in U_{i,\text{last}}$.

Note that Algorithm 1 tracks a set $\mathcal{U}_{\text{sent}}(\tau)$ that contains all the user groups that have been used by the algorithm before time τ . We let $\tilde{y}_{\{i,f\}}(U)$, $f \in \mathcal{F}_{\{i,U\}}$ and $U \in \mathcal{U}_{\text{sent}}(\tau_i)$ with $U \ni i$ be the solution of (8) when solving it for $w_{\{i,U_{i,\text{last}}\}}(\tau_i)$. Then, for each $U \in \cup_{\ell=1}^\beta \mathcal{U}_\ell$ with $U \ni i$ and for each $f \in \mathcal{F}_{\{i,U\}}$ we assign $y_{\{i,f\}}(U) = \tilde{y}_{\{i,f\}}(U)$ if $U \in \mathcal{U}_{\text{sent}}(\tau_i)$ and $y_{\{i,f\}}(U) = 0$ otherwise. We apply this assignment for all $i \in [K]$. Algorithm 1 assigns integer values to $z_U(\tau)$'s. From Remark 4 it follows that there exists an integral solution for $\tilde{y}_{\{i,f\}}(U)$'s and consequently the $y_{\{i,f\}}(U)$'s as well. With these assignments, we now demonstrate that the second and third conditions in (1) hold.

For the second condition we note that if $U \notin \mathcal{U}_{\text{sent}}(\tau_i)$ then $y_{\{i,f\}}(U) = 0$ and we have nothing to show. For $U \in \mathcal{U}_{\text{sent}}(\tau_i)$ we have that $y_{\{i,f\}}(U) = \tilde{y}_{\{i,f\}}(U)$. Recall that $\tilde{y}_{\{i,f\}}(U)$ is the solution of (8) at time $\tau = \tau_i$. By the way that $z_U(\tau)$ has been updated in Algorithm 1, we have $z_U(\tau) \leq z_U(T_{\max})$. Therefore, we have $z_U(\tau_i + 1) \leq z_U(T_{\max}) = \sum_{\ell \in \mathcal{I}_U} \tilde{x}_U(\ell)$ and from (8) for $w_{\{i,U_{i,\text{last}}\}}(\tau_i)$,

$$\begin{aligned} \sum_{f \in \mathcal{F}_{\{i,U\}}} y_{\{i,f\}}(U) &= \sum_{f \in \mathcal{F}_{\{i,U\}}} \tilde{y}_{\{i,f\}}(U) \leq \tilde{z}_U(\tau_i) = z_U(\tau_i + 1) \\ &\leq \sum_{\ell \in \mathcal{I}_U} \tilde{x}_U(\ell) = \sum_{\ell \in \mathcal{I}_U} x_U(\ell). \end{aligned}$$

For the third condition, consider any user $i \in [K]$ and any $f \in \Omega^{(i)}$. Recalling the definition of τ_i and $w_{\{i,U_{i,\text{last}}\}}(\tau_i)$, we know that $w_{\{i,U_{i,\text{last}}\}}(\tau_i) = v_i(\tau_i + 1) \geq |\Omega^{(i)}|$ which implies that in (8), we have

$$\begin{aligned} |\Omega^{(i)}| &\leq \sum_{U \in \mathcal{U}_{\text{sent}}(\tau_i), U \ni i} \sum_{f \in \mathcal{F}_{\{i,U\}}} \tilde{y}_{\{i,f\}}(U) \\ &= \sum_{f \in \Omega^{(i)}} \sum_{U \in \mathcal{U}_{\{i,f\}} \cap \mathcal{U}_{\text{sent}}(\tau_i)} \tilde{y}_{\{i,f\}}(U) \leq \sum_{f \in \Omega^{(i)}} 1 = |\Omega^{(i)}|, \end{aligned}$$

where the last inequality comes from the second constraint in (8). The middle equality holds by counting arguments for missing subfiles $f \in \Omega^{(i)}$ and user groups in $U \in \mathcal{U}_{\text{sent}}(\tau_i)$. To verify this, consider a bipartite graph in which the left and right nodes correspond to $f \in \Omega^{(i)}$ and $U \in \mathcal{U}_{\text{sent}}(\tau_i)$ with $U \ni i$ respectively. There is an edge between nodes corresponding to f and U if and only if $f \in \mathcal{F}_{\{i,U\}}$. We let $\tilde{y}_{\{i,f\}}(U)$ to be the label of this edge. By the definition of $\mathcal{U}_{\{i,f\}}$ we know that $f \in \mathcal{F}_{\{i,U\}}$ implies $U \in \mathcal{U}_{\{i,f\}}$. Therefore, outgoing edges from the node corresponding to f are the edges between f and the nodes $U \in \mathcal{U}_{\{i,f\}} \cap \mathcal{U}_{\text{sent}}(\tau_i)$. Similarly, the outgoing edges between node $U \in \mathcal{U}_{\text{sent}}(\tau_i)$ with $U \ni i$ are the edges between U and $f \in \mathcal{F}_{\{i,U\}}$. By counting $\tilde{y}_{\{i,f\}}(U)$ in two ways, from the left and right nodes, we have the required equality. Therefore, we have that $\sum_{U \in \mathcal{U}_{\{i,f\}} \cap \mathcal{U}_{\text{sent}}(\tau_i)} \tilde{y}_{\{i,f\}}(U) = 1$ for any $f \in \Omega^{(i)}$. This further implies that $\sum_{U \in \mathcal{U}_{\{i,f\}}} y_{\{i,f\}}(U) = 1$ for all $f \in \Omega^{(i)}$ and completes the proof. ■

C. Proof of Lemma 1

Proof: For simplicity, in the discussion below we assume that $r = 1$. The proof for $r > 1$ follows in straightforward manner. By the way that $\mathcal{M}_i$ and $v_i(\tau)$ are updated in Algorithm 1, we have $|\mathcal{M}_i| = v_i(\tau)$ at each time τ . Furthermore, $v_i(T_{\max}) = |\Omega^{(i)}|$ for all $i \in [K]$. Therefore, each user $i \in [K]$ benefits from $|\Omega^{(i)}|$ equations. For a $m \in \mathcal{M}_i$, let $\bigoplus_{i \in U} \bigoplus_{f \in \mathcal{F}_{\{i,U\}}} \alpha_{\{i,f,m\}} W_{\{d_i,f\}}$ represent the m -th equation (the dependence on index j is suppressed since we assume that $r = 1$). User $i \in U$ can recover $\bigoplus_{f \in \mathcal{F}_{\{i,U\}}} \alpha_{\{i,f,m\}} W_{\{d_i,f\}}$ from this equation since the missing subfiles $W_{\{d_j,f'\}}$, for $f' \in \mathcal{F}_{\{j,U\}}$ and $j \in U \setminus \{i\}$, exist in the cache of user i .

For each user $i \in [K]$ we define matrix $\mathbf{B}_i \in \mathbb{F}^{|\Omega^{(i)}| \times |\Omega^{(i)}|}$ whose rows and columns correspond to equation numbers in $\mathcal{M}_i$ and missing subfiles in $\Omega^{(i)}$ respectively. For $m \in \mathcal{M}_i$, assume that m -th equation is associated with user group U , where $i \in U$. Then, the entry of $\mathbf{B}_i$ for the row and column

corresponding to $m \in \mathcal{M}_i$ and $f \in \Omega^{(i)}$ is $\alpha_{\{i,f,m\}}$ if $f \in \mathcal{F}_{\{i,U\}}$ and zero otherwise. Therefore, if matrix $\mathbf{B}_i$ is invertible then user i can recover all the missing subfiles $W_{\{d_i,f\}}$, for $f \in \Omega^{(i)}$, from equations $\sum_{f \in \mathcal{F}_{\{i,U\}}} \alpha_{\{i,f,m\}} W_{\{d_i,f\}}$ for $m \in \mathcal{M}_i$. Thus, we need to show that the determinant of $\mathbf{B}_i$ is nonzero for all $i \in [K]$ with high probability.

Towards this end, let $h_i(\{\alpha_{\{i,f,m\}}, f \in \Omega^{(i)}, m \in \mathcal{M}_i\})$ denote the determinant of $\mathbf{B}_i$; we treat the $\{\alpha_{\{i,f,m\}}, f \in \Omega^{(i)}, m \in \mathcal{M}_i\}$ as indeterminates at this point. Note that since Algorithm 1 did not return “INFEASIBLE”, we have a feasible integral solution for the corresponding offline LP (cf. Claim 1). Thus, there exists an interpretation of this solution (cf. Section IV-B) such that in each time slot, only one equation is transmitted, i.e., unlike a fractional solution, we do not need to potentially transmit multiple equations in the same time slot. This in turn implies that there is a setting for coefficients $\alpha_{\{i,f,m\}}$ with $\alpha_{\{i,f,m\}} \in \{0,1\}$ such that the multivariate polynomial h_i evaluates to a non-zero value over $\mathbb{F}$, i.e., h_i is not identically zero. This further implies that $h = \prod_{i \in [K]} h_i$ is not identically zero. Now, since each $\alpha_{\{i,f,m\}}$ appears only once in $\mathbf{B}_i$ thus its degree in polynomial h_i is one. Also, h_i is a polynomial of degree $|\Omega^{(i)}| \leq F$ thus h is a polynomial of degree at most KF . Therefore, we can use Lemma 4 in [30] to show that by choosing $\alpha_{\{i,f,m\}}$ ’s independently and uniformly at random from $\mathbb{F}$, the determinants of $\mathbf{B}_i$ ’s, $i \in [K]$, are nonzero with probability at least $(1 - \frac{1}{|\mathbb{F}|})^{KF}$.

When $r > 1$ we will need to split a missing subfile $W_{\{d_i,f\}}$ into r packets and code over these as well. Thus, the corresponding system of equations will be of size $\mathbb{F}^{r|\Omega^{(i)}| \times r|\Omega^{(i)}|}$ leading to the bound $(1 - \frac{1}{|\mathbb{F}|})^{rKF}$. ■

REFERENCES

- [1] M. A. Maddah-Ali and U. Niesen, “Fundamental limits of caching,” *IEEE Trans. Inf. Theory*, vol. 60, no. 5, pp. 2856–2867, May 2014.
- [2] F. Arbabjolfaei et al., “Fundamentals of index coding,” *Found. Trends Commun. Inf. Theory*, vol. 14, nos. 3–4, pp. 163–346, 2018.
- [3] H. Ghasemi and A. Ramamoorthy, “Asynchronous coded caching,” in *Proc. IEEE Int. Symp. Inf. Theory*, Jun. 2017, pp. 2438–2442.
- [4] R. K. Ahuja, T. L. Maganti, and J. B. Orlin, *Network Flows: Theory, Algorithms and Applications*. Upper Saddle River, NJ, USA: Prentice-Hall, 1993.
- [5] H. Ghasemi and A. Ramamoorthy, “Improved lower bounds for coded caching,” *IEEE Trans. Inf. Theory*, vol. 63, no. 7, pp. 4388–4413, Jul. 2017.
- [6] Q. Yu, M. A. Maddah-Ali, and A. S. Avestimehr, “The exact rate-memory tradeoff for caching with uncoded prefetching,” *IEEE Trans. Inf. Theory*, vol. 64, no. 2, pp. 1281–1296, 2017.
- [7] Q. Yu, M. A. Maddah-Ali, and A. S. Avestimehr, “Characterizing the rate-memory tradeoff in cache networks within a factor of 2,” *IEEE Trans. Inf. Theory*, vol. 65, no. 1, pp. 647–663, Jan. 2019.
- [8] L. Tang and A. Ramamoorthy, “Coded caching for networks with the resolvability property,” in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Jul. 2016, pp. 420–424.
- [9] K. Wan, M. Ji, P. Piantanida, and D. Tuninetti, “Caching in combination networks: Novel multicast message generation and delivery by leveraging the network topology,” in *Proc. IEEE Int. Conf. Commun. (ICC)*, May 2018, pp. 1–6.
- [10] N. Naderializadeh, M. A. Maddah-Ali, and A. S. Avestimehr, “Fundamental limits of cache-aided interference management,” *IEEE Trans. Inf. Theory*, vol. 63, no. 5, pp. 3092–3107, May 2017.
- [11] Q. Yan, M. Cheng, X. Tang, and Q. Chen, “On the placement delivery array design for centralized coded caching scheme,” *IEEE Trans. Inf. Theory*, vol. 63, no. 9, pp. 5821–5833, Sep. 2017.
- [12] L. Tang and A. Ramamoorthy, “Coded caching schemes with reduced subpacketization from linear block codes,” *IEEE Trans. Inf. Theory*, vol. 64, no. 4, pp. 3099–3120, Apr. 2018.
- [13] E. Lampsiris and P. Elia, “Adding transmitters dramatically boosts coded-caching gains for finite file sizes,” *IEEE J. Sel. Areas Commun.*, vol. 36, no. 6, pp. 1176–1188, Jun. 2018.
- [14] M. Kiamari, C. Wang, and A. S. Avestimehr, “On heterogeneous coded distributed computing,” in *Proc. IEEE Global Telecommun. Conf. (GLOBECOM)*, 2017.
- [15] S. Li, M. A. Maddah-Ali, Q. Yu, and A. S. Avestimehr, “A fundamental tradeoff between computation and communication in distributed computing,” *IEEE Trans. Inf. Theory*, vol. 64, no. 1, pp. 109–128, 2017.
- [16] K. Konstantinidis and A. Ramamoorthy, “Resolvable designs for speeding up distributed computing,” *IEEE/ACM Trans. Netw.*, early access, May 29, 2020, doi: 10.1109/TNET.2020.2992989.
- [17] U. Niesen and M. A. Maddah-Ali, “Coded caching for delay-sensitive content,” in *Proc. IEEE Int. Conf. Commun. (ICC)*, Jun. 2015, pp. 5559–5564.
- [18] M. A. Maddah-Ali and U. Niesen, “Decentralized coded caching attains order-optimal memory-rate tradeoff,” *IEEE/ACM Trans. Netw.*, vol. 23, no. 4, pp. 1029–1040, Aug. 2015.
- [19] Y. Lu, W. Chen, and H. V. Poor, “Coded joint pushing and caching with asynchronous user requests,” *IEEE J. Sel. Areas Commun.*, vol. 36, no. 8, pp. 1843–1856, Aug. 2018.
- [20] Y. Jiang, W. Huang, M. Bennis, and F.-C. Zheng, “Decentralized asynchronous coded caching design and performance analysis in fog radio access networks,” *IEEE Trans. Mobile Comput.*, vol. 19, no. 3, pp. 540–551, Mar. 2020.
- [21] Q. Yang, M. M. Amiri, and D. Gündüz, “Audience-retention-rate-aware caching and coded video delivery with asynchronous demands,” 2018, *arXiv:1808.04835*. [Online]. Available: <http://arxiv.org/abs/1808.04835>
- [22] P. M. Vaidya, “An algorithm for linear programming which requires $O((m+n)n^2 + (m+n)^{1.5}nL)$ arithmetic operations,” *Math. Prog.*, vol. 47, pp. 175–201, May 1990.
- [23] D. S. Lun et al., “Minimum-cost multicast over coded packet networks,” *IEEE Trans. Inf. Theory*, vol. 52, no. 6, pp. 2608–2623, Jun. 2006.
- [24] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, U.K.: Cambridge Univ. Press, 2004.
- [25] LEMON. *Library for Efficient Modeling and Optimization in Networks*. [Online]. Available: <http://lemon.cs.elte.hu>
- [26] H. D. Sherali and G. Choi, “Recovery of primal solutions when using subgradient optimization methods to solve Lagrangian duals of linear programs,” *Oper. Res. Lett.*, vol. 19, no. 3, pp. 105–113, Sep. 1996.
- [27] H. Ghasemi, “Coded caching: Information theoretic bounds and asynchronism,” Ph.D. dissertation, Dept. Elect. Comput. Eng., Iowa State Univ., Ames, IA, USA, 2019. [Online]. Available: <http://www.ece.iastate.edu/~ghasemi/publication>
- [28] J. Kleinberg and E. Tardos, *Algorithm Design*. New Delhi, India: Pearson Education, 2006.
- [29] (2019). *Coded Caching*. [Online]. Available: <https://hooshanggh.github.io/Coded-Caching/>
- [30] T. Ho et al., “A random linear network coding approach to multicast,” *IEEE Trans. Inf. Theory*, vol. 52, no. 10, pp. 4413–4430, Oct. 2006.

Hooshang Ghasemi (Member, IEEE) received the Ph.D. degree from Iowa State University, in 2019. He was with the Electrical and Computer Engineering Department, Iowa State University. He is currently with Qualcomm Inc., San Diego, CA, USA.

Aditya Ramamoorthy (Senior Member, IEEE) received the B. Tech. degree in electrical engineering from the Indian Institute of Technology, Delhi, in 1999, and the M.S. and Ph.D. degrees from the University of California, Los Angeles (UCLA), in 2002 and 2005, respectively. He is currently a Professor of Electrical and Computer Engineering and (by courtesy) of Mathematics at Iowa State University. His research interests are in the areas of information theory and coding techniques with applications to networks, distributed storage and distributed computation. He served as an editor for the IEEE TRANSACTIONS ON INFORMATION THEORY, from 2016 to 2019, and for the IEEE TRANSACTIONS ON COMMUNICATIONS, from 2011 to 2015. He was the recipient of the 2020 Mid-Career Achievement in Research Award and the 2012 Early Career Engineering Faculty Research Award from Iowa State University, the 2012 NSF CAREER award, and the Harpole-Pentair professorship in 2009 and 2010.