



Cache Updating Strategy Minimizing the Age of Information with Time-Varying Files' Popularities

Haoyue Tang, Philippe Ciblat, Jintao Wang, Michèle Wigger, Roy D Yates

► To cite this version:

Haoyue Tang, Philippe Ciblat, Jintao Wang, Michèle Wigger, Roy D Yates. Cache Updating Strategy Minimizing the Age of Information with Time-Varying Files' Popularities. IEEE Information Theory Workshop (ITW), Apr 2021, Riva de Garda, Italy. hal-03291067

HAL Id: hal-03291067

<https://hal.science/hal-03291067>

Submitted on 19 Jul 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Cache Updating Strategy Minimizing the Age of Information with Time-Varying Files' Popularities

Haoyue Tang, Philippe Ciblat, Jintao Wang, Michèle Wigger, Roy D. Yates

Abstract—We consider updating strategies for a local cache which downloads time-sensitive files from a remote server through a bandwidth-constrained link. The files are requested randomly from the cache by local users according to a popularity distribution which varies over time according to a Markov chain structure. We measure the freshness of the requested time-sensitive files through their Age of Information (AoI). The goal is then to minimize the average AoI of all requested files by appropriately designing the local cache's downloading strategy. To achieve this goal, the original problem is relaxed and cast into a Constrained Markov Decision Problem (CMDP), which we solve using a Lagrangian approach and Linear Programming. Inspired by this solution for the relaxed problem, we propose a practical cache updating strategy that meets all the constraints of the original problem. Under certain assumptions, the practical updating strategy is shown to be optimal for the original problem in the asymptotic regime of a large number of files. For a finite number of files, we show the gain of our practical updating strategy over the traditional square-root-law strategy (which is optimal for fixed non time-varying file popularities) through numerical simulations.

I. INTRODUCTION AND PROBLEM STATEMENT

Consider a local cache connected by a capacity-constrained link to a remote network server, as shown in Figure 1. The server stores N time-sensitive files that change in a continuous manner. The local cache maintains a copy of each file, and, upon request, sends the copy to a local user. By the capacity-constrained link from the local cache to the server, the cache cannot maintain the latest version of each item, and the copy it sends to the user can be outdated. The goal of our study is to measure the freshness of the copies sent to the users in terms of their *age of information (AoI)*, and to propose a cache updating strategy that minimizes the average AoI of the downloaded copies. Our main focus is on a setup where the popularities of the various files vary over time.

Existing works for cache updating with time-sensitive files and time-varying environments mainly aim at minimizing the average AoI when channel states change randomly [1], [2], energy arrivals are random [3], or packet arrivals are random [4]. A few works consider time-varying popularities for cache




Fig. 1. Local cache system.

updating, e.g., [5]–[7]. Similar to our work, [5], [6] propose policies for downloading files to a local cache over a capacity-limited link, but their figure of merit is not the AoI. Only [7] considers both time-varying popularities and AoI.

Nevertheless the setting in [7] differs from ours in that in [7]: i) the cache (as opposed to the server-cache link) is capacity-limited; ii) the goal is to minimize the missed-files probability in the cache (as opposed to the average AoI), iii) the current popularity value depends on the past requests weighted by their AoI (as opposed to being determined by a stationary Markov chain independent of the files' AoI as assumed in this paper); and finally iv) the considered AoI is different from ours since their age is defined as the request rate within a given time-interval.

Our paper can actually be regarded as an extension of our previous work [8] to time-varying file popularities. It is also related to [9], which extended [8] to variable update durations.

We now explain our model in more detail. Consider the system in Fig. 1, which comprises a remote server holding N files $n = 1, \dots, N$ that are subject to version updates, a local cache downloading the latest versions through a bandwidth limited network, and users requesting files from the local cache. In each time slot $t \in \{1, \dots, T\}$, the local cache can download the current version of no more than M files from the remote server due to the bandwidth constraint. For each file $n \in \{1, \dots, N\}$, let $\{u_{n,t}\} \in \{0, 1\}$ be the download decision at time t , i.e., $u_{n,t} = 1$, if file n is downloaded at time t , and $u_{n,t} = 0$ otherwise. The bandwidth constraint requires that

$$\sum_{n=1}^N u_{n,t} \leq M, \quad \forall t. \quad (1)$$

We denote by $X_{n,t}$ the AoI of file n in slot t , i.e., the number of slots that have passed since the local cache has downloaded file n . For convenience of exposition, we assume that all the files in the local cache are updated at time one so that $X_{n,1} = 1, \forall n$. Afterwards for $t > 1$, the AoI evolves as

$$X_{n,t+1} = \begin{cases} 1, & u_{n,t} = 1; \\ X_{n,t} + 1, & u_{n,t} = 0. \end{cases} \quad (2)$$

Haoyue Tang and Jintao Wang are with Dept. of Electronic Engineering, Tsinghua University, Beijing, China ({thy17@mails, wangjintao@mail}.tsinghua.edu.cn). Philippe Ciblat and Michèle Wigger are with Telecom Paris, Institut Polytechnique de Paris, Palaiseau, France ({ciblat, wigger}@telecom-paris.fr). Roy D. Yates is with Rutgers University, New Brunswick, NJ, USA (ryates@rutgers.edu).

This work was supported by the ERC (Grant No. 715111), the National Key R&D Program of China (Grant No.2017YFE0112300) and Beijing National Research Center for Information Science and Technology under Grant BNR2019RC01014 and BNR2019TD01001.

In each slot t , the number of requests for file n depends on its current popularity mode $R_{n,t} \in \mathcal{R} := \{1, \dots, R\}$, where $R > 0$. The expected number of requests of file n is determined by a function $\omega_n : \mathcal{R} \rightarrow \mathbb{R}^+$, so that the expected number of requests for file n at time slot t is given by $\omega_n(R_{n,t}) > 0$. It is assumed that for each file n , the sequence $\{R_{n,t}\}_{t=1}^T$ evolves according to an R -state Markov chain with transition probabilities $P_{r,r'}^n := \Pr(R_{n,t+1} = r' | R_{n,t} = r)$ for $(r, r') \in \mathcal{R}^2$.

Let Π be a set of (cache updating) strategies, such that the design of the downloading decisions $\{u_{n,t}\}$ at slot t only depend on the current and past popularity modes $\{R_{n,\tau}\}_{\tau \leq t}$ and AoIs $\{X_{n,\tau}\}_{\tau \leq t}$ as well as on the statistics $\{P_{r,r'}^n\}$. The future popularity modes $\{R_{n,\tau}\}_{\tau > t}$ cannot be used. The goal in this article is to design a strategy $\pi \in \Pi$ that minimizes the expected total AoI of all requested files averaged over an infinite-time horizon. The corresponding optimization problem can be written as:

Problem 1 (Original problem):

$$\pi^* = \arg \min_{\pi \in \Pi} \lim_{T \rightarrow \infty} \mathbb{E}_\pi \left[\frac{1}{T} \sum_{t=1}^T \sum_{n=1}^N \omega_n(R_{n,t}) X_{n,t} \right], \quad (3a)$$

$$\text{s.t.} \quad \sum_{n=1}^N u_{n,t} \leq M, \quad \forall t. \quad (3b)$$

Problem 1 can be cast into a Markov Decision Process (MDP) where the state contains both the current AoIs $\mathbf{X}_t = [X_{1,t}, \dots, X_{N,t}]^T$ and popularities $\mathbf{R}_t = [R_{1,t}, \dots, R_{N,t}]^T$ of all the files, and Eq. (3b) can be cast as a constraint on the action space $\mathcal{A} := \{\mathbf{u}_t | \sum_{n=1}^N u_{n,t} \leq M\}$. The cardinality of the action space \mathcal{A} of this MDP grows exponentially in M , and thus even for moderate values of M its solution cannot be found using standard algorithms such as *Relative Value Iteration* [10].

We therefore slightly relax the hard bandwidth-constraint (3b) and derive an optimal solution for the relaxed optimization problem, see solution π_R^* to Problem 2 in Section II. Motivated by this solution, in Section II-D we propose a practical updating strategy $\hat{\pi}$ which satisfies the original constraint (3b). In Section III, we prove that, under mild conditions, the strategy $\hat{\pi}$ is optimal for Problem 1 when N goes to infinity for a fixed N/M . Numerical illustrations are provided in Section IV. Concluding remarks are drawn in Section V.

II. PROBLEM RESOLUTION

A. Relaxed problem description

Similarly to [2], [4], [8], [11], [12], we relax Constraint (3b) into an expected infinite-time horizon constraint. This leads to the following relaxed optimization problem:

Problem 2 (Relaxed problem):

$$\pi_R^* = \arg \min_{\pi \in \Pi} \lim_{T \rightarrow \infty} \mathbb{E}_\pi \left[\frac{1}{T} \sum_{t=1}^T \sum_{n=1}^N \omega_n(R_{n,t}) X_{n,t} \right], \quad (4a)$$

$$\text{s.t.} \quad \lim_{T \rightarrow \infty} \mathbb{E}_\pi \left[\frac{1}{T} \sum_{t=1}^T \sum_{n=1}^N u_{n,t} \right] \leq M. \quad (4b)$$

Problem 2 can be cast into the framework of Constrained Markov Decision Processes (CMDP) [13]. Its action space however is even larger than that of Problem 1. What makes Problem 2 tractable is that it can be decoupled into independent sub-problems with smaller action spaces, as we now explain.

The CMDP associated with Problem 2 is a countable-state CMDP with finite set of actions. Consequently, [14] asserts that Problem 2 can be solved, and that the optimal policy π_R^* can be determined by introducing the Lagrangian function

$$\mathcal{L}(\pi, W) = \lim_{T \rightarrow \infty} \mathbb{E}_\pi \left[\frac{1}{T} \sum_{t=1}^T \sum_{n=1}^N (\omega_n(R_{n,t}) X_{n,t} + W u_{n,t}) - W M \right]. \quad (5)$$

associated with the cost function (4a) and the constraint (4b).

In the remainder of this subsection and the next-following Subsection II-B, we minimize the Lagrangian for a fixed value of W . In Subsection II-C we then determine the appropriate value(s) of W which lead to the solution of Problem 2.

To minimize the Lagrangian for a fixed value W , we first notice that Eq. (5) is separable over the various files. This is easier to see after swapping the order of the summations over t and n in Eq. (5) (the order can be swapped because the sum over n is finite). As a consequence (see [15, Chapter 4] for more details), the minimizing policy $\pi_F^*(W)$ in (5) for given W factorizes as

$$\pi_F^*(W) = \bigotimes_{n=1}^N \pi_F^{n,*}(W), \quad (6)$$

where $\pi_F^{n,*}(W)$ denotes the solution to the following optimization problem for file n only:

Problem 3 (per-File relaxed problem):

$$\pi_F^{n,*}(W) \triangleq \arg \min_{\pi \in \Pi} \lim_{T \rightarrow \infty} \mathbb{E}_\pi \left[\frac{1}{T} \sum_{t=1}^T (\omega_n(R_{n,t}) X_{n,t} + W u_{n,t}) \right]. \quad (7)$$

B. An algorithmic solution for Problem 3

We solve the optimization problem (7) for a fixed file index n and Lagrange multiplier W . For simplicity, we omit the subscript n . W is also omitted except when it appears in a mapping expression.

Problem 3 can be cast into an MDP with a two-dimensional state (X_t, R_t) , action $u_t \in \{0, 1\}$, and instantaneous cost

$$C(X_t, R_t, u_t) = \omega(R_t) X_t + W u_t. \quad (8)$$

According to Eq. (2), if $u_t = 1$, the file is downloaded in slot t and the AoI drops to 1 in the next slot; otherwise, the AoI grows by 1. The state transition relationship thus is

$$\Pr((X_t, R_t) \rightarrow (1, r)) = P_{R_t, r}, \quad \text{if } u_t = 1; \quad (9a)$$

$$\Pr((X_t, R_t) \rightarrow (X_t + 1, r)) = P_{R_t, r}, \quad \text{if } u_t = 0. \quad (9b)$$

Definition 1: A policy π is called *stationary*, if for each time t and AoI-popularity-mode pair $(X_t = x, R_t = r)$, the action

$u_t = 1$ is chosen with a probability $\xi_{x,r}$ that only depends on the AoI-popularity-mode pair (x, r) but not on the time t .

The next theorem states that the optimum solution π_F^* is a stationary policy with a specific threshold structure.

Theorem 1: There exists an optimal stationary policy π_F^* and a set of thresholds $\{\tau_r\}_{r \in \mathcal{R}}$ such that π_F^* downloads the file with probability 1 in state r if $x > \tau_r$ and it keeps idle with probability 1 if $x < \tau_r$.

The proof is similar to [2, Lemma 1] but where instead of the AoI and the channel quality, the AoI and the popularity mode should be considered for the two-dimensional state. Many other papers prove the optimality of threshold policies, but usually with a one-dimensional state [4], [12].

By Theorem 1, there exists a stationary optimal policy such that the AoI is bounded as $X_{\max} = \max_{r \in \mathcal{R}} \tau_r$, where the maximum exists because \mathcal{R} is finite. Inspecting, (7), one sees in particular that there must exist a stationary optimal (threshold) policy with

$$X_{\max} \leq \max_{r \in \mathcal{R}} \frac{W + \omega(r)}{\omega(r)} =: X_{\text{ub}}. \quad (10)$$

The term X_{ub} is finite because $\omega(r) > 0$ and because \mathcal{R} is finite. Above inequality (10) is obtained by showing that for any policy violating (10), it is possible to find an improved policy that in popularity state $R = r$ updates file n whenever its age $X > \frac{W + \omega(r)}{\omega(r)}$. Notice that, in the following subsections, we will often write $X_{\text{ub}}^n(W)$ instead of X_{ub} to make the dependence on n and W explicit.

We can thus restrict to finding the optimal stationary policy that solves the CMDP with a restricted *finite* state space $\{1, \dots, X_{\text{ub}}\} \times \mathcal{R}$. Different algorithms for finding such a policy are described in [10]. In this paper, we resort to the *Linear Programming (LP)* approach developed in [13], and rewrite Problem 3 in terms of the two steady-state distributions $\{\mu_{x,r}\}$ and $\{\nu_{x,r}\}$. Here, $\mu_{x,r}$ stands for the steady-state probability of having AoI $X = x$ and popularity state $R = r$, whereas the so called *occupation measure* $\nu_{x,r}$ stands for the steady-state probability of simultaneously having $(X = x, R = r)$ and taking the “download action” $u = 1$. As a consequence, the probability of updating a file for a given AoI $X = x$ and popularity state $R = r$ equals $\xi_{x,r} = \nu_{x,r} / \mu_{x,r}$. We set $0/0 = 1$ by convention. But notice that this convention has no effect on the solution of the optimization problem because state (x, r) is reached with probability $\mu_{x,r} = 0$. With these definitions, and because the cost $\omega(r)x$ only depends on the state and the constraint u is nonzero only for a single action, we can apply [13, Theorem 4.3] to obtain:

Theorem 2 (Equivalent to Theorem 4.3 in [13]): Let X_{ub} be defined as in (10). The optimal stationary policy $\{\xi_{x,r}^*\}_{x,r}$ solving Problem 3 is given by

$$\xi_{x,r}^* = \frac{\nu_{x,r}^*}{\mu_{x,r}^*}, \quad (11)$$

where $\nu_{x,r}^*$ and $\mu_{x,r}^*$ are obtained by the following LP problem:

$$\{\mu_{x,r}^*, \nu_{x,r}^*\} = \arg \min_{\{\mu_{x,r}, \nu_{x,r}\}} \sum_{x=1}^{X_{\text{ub}}} \sum_{r=1}^R (\omega(r)x\mu_{x,r} + W\nu_{x,r}), \quad (12a)$$

$$\text{s.t. } \mu_{1,r} = \sum_{x=1}^{X_{\text{ub}}} \sum_{r'=1}^R \nu_{x,r'} P_{r',r}, \quad (12b)$$

$$\mu_{x,r} = \sum_{r'=1}^R (\mu_{x-1,r'} - \nu_{x-1,r'}) P_{r',r}, \forall x > 1, \quad (12c)$$

$$\sum_{x=1}^{X_{\text{ub}}} \sum_{r=1}^R \mu_{x,r} = 1, \quad (12d)$$

$$\nu_{x,r} \leq \mu_{x,r}, \quad (12e)$$

$$0 \leq \mu_{x,r}, 0 \leq \nu_{x,r}, \forall x, r. \quad (12f)$$

Before going further, we have some remarks:

- As mentioned in [13, Theorem 4.3], any solution of the LP described in Theorem 2 leads to a stationary optimal policy through Eq. (11). Conversely, any stationary optimal policy for Problem 3 is also a solution to Theorem 2.
- The set of constraints in the above-mentioned LP is just a straightforward application of [13, Theorem 4.3] except for Eq. (12e). This constraint has been added since the LP is written with respect to the occupation measure and the steady-state distribution (which is a sum of all the occupation measures). So, by construction, $\nu_{x,r} \leq \mu_{x,r}$.

C. An algorithmic solution for Problem 2

In Section II-B, we described an LP approach to obtain an optimal policy $\pi_F^{n,*}(W) = \{\xi_{x,r}^{n,*}\}$ (or equivalently, $\{\mu_{x,r}^{n,*}(W), \nu_{x,r}^{n,*}(W)\}_{x,r}$) for Problem 3 for any file n and Lagrange multiplier W . By Eq. (6), the product of these policies minimizes the Lagrangian function in Eq. (5) for the given multiplier W . So, at each time t , for given AoI vector $(X_{1,t}, \dots, X_{N,t}) = [x_1, \dots, x_N]$ and popularity vector $(R_{1,t}, \dots, R_{N,t}) = [r_1, \dots, r_N]$, the optimal policy $\pi_F^*(W)$ that minimizes the Lagrangian (5) for parameter W updates each file n independently of all the other files with probability $\xi_{x_n, r_n}^{n,*}(W) = \frac{\nu_{x_n, r_n}^{n,*}(W)}{\mu_{x_n, r_n}^{n,*}(W)}$. The average proportion of time spent on downloading files in this optimal policy is

$$d^*(W) := \sum_{n=1}^N \sum_{x=1}^{X_{\text{ub}}^n(W)} \sum_{r=1}^R \nu_{x,r}^{n,*}(W), \quad (13)$$

and the expected average AoI

$$a^*(W) \triangleq \sum_{n=1}^N \sum_{x=1}^{X_{\text{ub}}^n(W)} \sum_{r=1}^R \omega(r)x\mu_{x,r}^{n,*}(W). \quad (14)$$

Thus, for a given Lagrange multiplier W and the optimal stationary policy $\pi_F^*(W)$, the Lagrangian in (5) can be compactly written as:

$$\mathcal{L}(\pi_F^*(W), W) = a^*(W) + Wd^*(W) - MW. \quad (15)$$

It remains to find optimal value of W . As in [16], notice that for each W and any policy π :

$$a^*(W) + Wd^*(W) - MW \leq a_\pi + Wd_\pi - MW, \quad (16)$$

where a_π and d_π denote the average proportion of time spent on downloading files and the expected average AoI under policy π . By (16), if for some \tilde{W} the constraint $d^*(\tilde{W}) = M$, then it has the smallest expected AoI $a^*(\tilde{W})$ among all policies respecting constraint $d_\pi \leq M$. Since $d^*(W)$ is non-increasing in W [16, Lemma 3.3], if the desired value of \tilde{W} above exists, it coincides with

$$W^* \triangleq \inf\{W | d^*(W) \leq M\}. \quad (17)$$

In this case, the optimal policy for Problem 2 is given by $\pi_R^* = \pi_F^*(W^*)$.

In case $W \mapsto d^*(W)$ is discontinuous around W^* , then we need to define two intermediate policies as follows. We recall that policy $\pi_F^*(W)$ is a finite product of policies $\pi_{F,ls}^{n,*}(W)$, each defined by the finite set $\{\xi_{x,r}^{n,*}(W)\}_{1 \leq x \leq X_{ub}^n(W), r \in \mathcal{R}}$. When we consider $W \in [0, W^*]$, according to (10), we can define $X_{ub,ls}^n = \max_{W \in [0, W^*]} X_{ub}^n(W)$. We now consider the zero-padded set $\{\xi_{x,r}^{n,*}(W)\}_{1 \leq x \leq X_{ub,ls}^n, r \in \mathcal{R}}$. As each $\xi_{x,r}^{n,*}(W)$ is bounded by 1, there exists a subsequence $\{W_{ls,\ell}\}_{\ell \geq 1}$ left-converging to W^* s.t. $\xi_{x,r,ls}^{n,*} := \lim_{\ell \rightarrow \infty} \xi_{x,r}^{n,*}(W_{ls,\ell})$ exists. Similarly, by considering \bar{W} larger than W^* and $X_{ub,rs}^n = \max_{W \in [W^*, \bar{W}]} X_{ub}^n(W)$, there exists a subsequence $\{W_{rs,\ell}\}_{\ell \geq 1}$ right-converging to W^* s.t. $\xi_{x,r,rs}^{n,*} := \lim_{\ell \rightarrow \infty} \xi_{x,r}^{n,*}(W_{rs,\ell})$ exists. Consequently, with an abuse of notation (we do not mention the selected subsequences anymore as well as the integer index ℓ), we can define

$$\pi_{F,ls}^* := \lim_{W \uparrow W^*} \pi_F^*(W) \quad (18)$$

and

$$\pi_{F,rs}^* := \lim_{W \downarrow W^*} \pi_F^*(W). \quad (19)$$

Notice that $\pi_{F,ls}^*$ and $\pi_{F,rs}^*$ may depend on the selected subsequence but this has no impact on the final result of this subsection.

As the mapping $\pi \mapsto d_\pi$ is continuous, the following limits exist:

$$d_{ls} = \lim_{W \uparrow W^*} d_{\pi_F^*}(W) \quad (20)$$

and

$$d_{rs} = \lim_{W \downarrow W^*} d_{\pi_F^*}(W). \quad (21)$$

According to (16), see [16, Theorem 4.4] and [14] for details, the following mixed policy is optimal

$$\pi_R^* = \lambda \pi_{F,ls}^* + (1 - \lambda) \pi_{F,rs}^*, \quad (22)$$

with

$$\lambda := \frac{M - d_{rs}}{d_{ls} - d_{rs}}. \quad (23)$$

Notice that the mixed policy in Eq. (22) means that at time $t = 0$, each of the two pure policies $\pi_{F,ls}^*$ and $\pi_{F,rs}^*$ is chosen with probability λ and $1 - \lambda$. The selected policy is then played

until the end of the process in slot $t = T$. Such a policy is obviously not stationary, and moreover, Constraint (4b) is not necessarily satisfied for all realizations of the decision process. An optimal *stationary* policy for Problem 2 can however easily be found based on π_R^* given above. The idea is to define the new steady-state probabilities and occupation measures

$$\bar{\mu}_{x,r}^{n,*} \triangleq \lambda \mu_{x,r,ls}^{n,*} + (1 - \lambda) \mu_{x,r,rs}^{n,*}, \quad (24)$$

$$\bar{\nu}_{x,r}^{n,*} \triangleq \lambda \nu_{x,r,ls}^{n,*} + (1 - \lambda) \nu_{x,r,rs}^{n,*}, \quad (25)$$

where $\{\mu_{x,r,ls}^{n,*}, \nu_{x,r,ls}^{n,*}\}$ and $\{\mu_{x,r,rs}^{n,*}, \nu_{x,r,rs}^{n,*}\}$ denote the steady-state probabilities and the occupation measures associated with the policies $\pi_{F,ls}^*$ and $\pi_{F,rs}^*$, respectively. Consider now the corresponding stationary policy $\bar{\pi}_R^*$, which updates at each time t the file n independently of all other files with a probability $\bar{\xi}_{x,r}^{n,*} \triangleq \frac{\bar{\nu}_{x,r}^{n,*}}{\bar{\mu}_{x,r}^{n,*}}$, if $X_{n,t} = x$ and $R_{n,t} = r$. By setting $\bar{X}_{ub}^n = \max(X_{ub,ls}^n, X_{ub,rs}^n)$, one can show this policy achieves the following average AoI

$$\sum_{n=1}^N \sum_{x=1}^{\bar{X}_{ub}^n} \sum_{r=1}^R \omega(r) x [\lambda \mu_{x,r,ls}^{n,*} + (1 - \lambda) \mu_{x,r,rs}^{n,*}] \quad (26)$$

and the following average downloading time

$$\sum_{n=1}^N \sum_{x=1}^{\bar{X}_{ub}^n} \sum_{r=1}^R [\lambda \nu_{x,r,ls}^{n,*} + (1 - \lambda) \nu_{x,r,rs}^{n,*}] = M. \quad (27)$$

Consequently, this policy offers the same AoI as the mixed policy π_R^* and satisfies the constraint, and is thus optimal.

D. An algorithmic solution for Problem 1

For our original Problem 1, we propose a (sub-optimal) policy $\hat{\pi}$ that behaves as the policy $\bar{\pi}_R^*$ derived above, except that at time-instances t where $\bar{\pi}_R^*$ downloads more than M files, $\hat{\pi}$ randomly choose M files among those that were to be updated by strategy $\bar{\pi}_R^*$.

III. OPTIMALITY IN AN ASYMPTOTIC REGIME

As shown in the following theorem, for N/M fixed and $N \rightarrow \infty$, the AoI of the practical policy $\hat{\pi}$ defined in Section II-D converges to the AoI of the optimal policy π^* that solves the original Problem 1.

Theorem 3: Assume that the N files have the same popularity statistics $P_{r,r'}$ and the same expected number of requests given a popularity mode $\omega(r)$. Assume also a fixed ratio $\theta := N/M$ independent of N . Assume that the Markov chain induced by policy $\bar{\pi}_R^*$ is ergodic. The policy $\hat{\pi}$ defined in Section II-D is then asymptotically optimal in the sense

$$\lim_{\substack{N \rightarrow \infty \\ \theta = N/M}} \frac{a_{\hat{\pi}} - a_{\pi^*}}{a_{\pi^*}} = 0, \quad (28)$$

where π^* is the optimal policy for Problem 1.

Proof: Omitted for space limitation. See [17]. ■

Notice that the condition of ergodicity for the Markov chain is mild here since the randomness of the popularity mode may lead to aperiodic and positive recurrent states for the Markov chain induced by $\bar{\pi}_R^*$.

Theorem 3 can easily be extended to the case of a finite number of classes of users, where a class is formed by all the users with the same popularity statistics and the same expected number of requests given a popularity mode $\omega(r)$. The setup in Theorem 3 corresponds to a single class.

IV. NUMERICAL RESULTS

We consider two popularity modes $\mathcal{R} = \{1, 2\}$ such that for some $q \in (0, 1)$ all files n have following transition matrix

$$\mathbf{P}^n = \begin{bmatrix} q & 1-q \\ 1-q & q \end{bmatrix}, \quad \forall n \in \{1, \dots, N\}. \quad (29)$$

The expected number of requests of file n in the two states is $\omega_n(1) = 0.2\bar{\omega}_n$ and $\omega_n(2) = 1.8\bar{\omega}_n$, where $\bar{\omega}_n \propto 1/n^\alpha$ follows a Zipf distribution with coefficient $\alpha = 1.5$. Due to Eq. (29), the steady-state probability for both popularities are identical, so $\bar{\omega}_n$ is the average number of requests for file n .

In Fig. 2, we plot the average AoI versus q with different M and $N = 64$ for the proposed policy $\hat{\pi}$, the relaxed policy $\bar{\pi}_R^*$, and the square-root law [8] designed with the average popularity. Notice that the AoI obtained with the relaxed policy $\bar{\pi}_R^*$ is a lower bound. We also recall that the square-root law does not take into account the time-varying characteristic of the popularity mode.




Fig. 2. AoI vs q for different M and $N = 64$.

We first observe that the proposed policy achieves a better AoI than the square-root law for any $q \neq 0.5$. For $q = 0.5$ the current popularity mode does not provide any information on the next one, therefore the proposed policy provides marginal gain. When q deviates more from 0.5, the gain becomes larger since the future popularity mode is better predicted and it is thus more important to take it into account. Moreover, the square-root law does not depend on q since the average number of requests over infinite horizon does not depend on q in our simulations. We further observe that when M increases, the AoI decreases since more updates can be done. Finally, the gap between the proposed policy and the relaxed one (which does not satisfy the hard constraint (1)) slightly depend on q .

V. CONCLUSION

The paper formulates a cache updating problem that aims to minimize the average AoI of the requested files under a bandwidth constraint on the server-cache link and assuming time-varying file popularities. We relaxed the hard bandwidth constraint and formulated the problem as a Constrained Markov Decision Process, which we then decoupled and solved through Linear Programming. Inspired by this approach, we proposed a practical updating strategy that satisfies the hard bandwidth constraint, and we showed that the proposed strategy is asymptotically optimal for a large number of files and under certain assumptions. Numerical results showed that for a fixed number of files the proposed strategy outperforms in many configurations the square-root law policy previously proposed for fixed popularities.

REFERENCES

- [1] Y. Sun, E. Uysal-Biyikoglu, R. D. Yates, C. E. Koksal, and N. B. Shroff, "Update or wait: How to keep your data fresh," *IEEE Transactions on Information Theory*, vol. 63, pp. 7492–7508, Nov 2017.
- [2] H. Tang, J. Wang, L. Song, and J. Song, "Minimizing age of information with power constraints: Opportunistic scheduling in multi-state time-varying networks," *IEEE Journal on Selected Areas in Communications*, 2020.
- [3] A. Arafat, J. Yang, S. Ulukus, and H. V. Poor, "Age-minimal transmission for energy harvesting sensors with finite batteries: Online policies," *IEEE Transactions on Information Theory*, vol. 66, pp. 534–556, Jan 2020.
- [4] Y. Hsu, E. Modiano, and L. Duan, "Scheduling algorithms for minimizing age of information in wireless broadcast networks with random arrivals," *IEEE Transactions on Mobile Computing*, pp. 1–1, 2019.
- [5] B. Bharath, K. Nagananda, D. Gunduz, and H. Poor, "Caching with time-varying popularity profiles: a learning-theoretic perspective," *IEEE Transactions on Communications*, vol. 66, pp. 3837–3847, Sep 2018.
- [6] J. Gao, S. Zhang, L. Zhao, and X. Shen, "The design of dynamic probabilistic caching with time-varying content popularity," *IEEE Transactions on Mobile Computing*, 2020.
- [7] C. Kam, S. Kompella, G. D. Nguyen, J. E. Wieselthier, and A. Ephremides, "Information freshness and popularity in mobile caching," in *IEEE International Symposium on Information Theory (ISIT)*, pp. 136–140, June 2017.
- [8] R. Yates, P. Ciblat, A. Yener, and M. Wigger, "Age-optimal constrained cache updating," in *IEEE International Symposium on Information Theory (ISIT)*, pp. 141–145, June 2017.
- [9] H. Tang, P. Ciblat, J. Wang, M. Wigger, and R. Yates, "Age of information aware cache updating with file- and age-dependent update durations," in *International Symposium on Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks (WiOpt)*, June 2020.
- [10] M. Puterman, *Markov Decision Processes*. Wiley, 1994.
- [11] P. Whittle, "Restless bandits: activity allocation in a changing world," *Journal of Applied Probability*, pp. 287–298, 1988.
- [12] I. Kadota, A. Sinha, E. Uysal-Biyikoglu, R. Singh, and E. Modiano, "Scheduling policies for minimizing age of information in broadcast wireless networks," *IEEE/ACM Transactions on Networking*, vol. 26, pp. 2637–2650, Dec 2018.
- [13] E. Altman, *Constrained Markov decision processes*. CRC Press, 1999.
- [14] L. Sennott, "Constrained Average Cost Markov Decision Chains," *Probability in the Engineering and Information Sciences*, vol. 7, pp. 69–83, 1993.
- [15] S. Stefanov, *Separable Programming: Theory and Methods*, vol. 53 of *Applied Optimization*. SpringerScience, 2001.
- [16] F. Beutler and K. Ross, "Optimal policies for controlled markov chains with a constraint," *Journal of Mathematical Analysis and Applications*, vol. 112, no. 1, pp. 236 – 252, 1985.
- [17] H. Tang, P. Ciblat, J. Wang, M. Wigger, and R. Yates, "Cache updating strategy minimizing the age of information with time-varying files' popularities (supplementary files)," [arXiv.org/abs/2010.04787](https://arxiv.org/abs/2010.04787), Sep 2020.