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

Haoyue Tang*, Philippe Ciblat[#], Jintao Wang*, Michèle Wigger[#], Roy Yates[ʃ]
Tsinghua University *
Télécom Paris #
Rutgers University ʃ

ITW2020

Outline

- Problem Statement
  - Age of Information
  - Optimization Problem Formulation

- Optimization Problem Resolution
  - Markov Decision Process (MDP) Formulation
  - Constraint Modification and Constrained Markov Decision (CMDP) Reformulation
  - Single File Decomposition
  - Linear Programming (LP) Solution
  - Practical Scheduling Algorithm for Primal Problem

- Performance Evaluation
  - Asymptotic optimality
  - Numerical Simulations

System Model

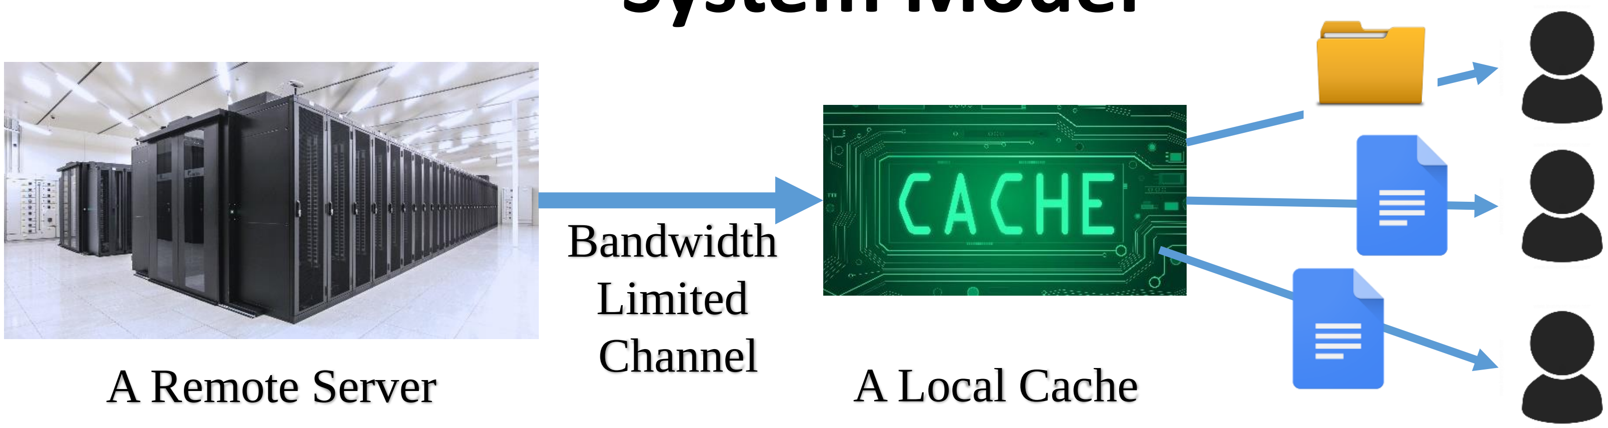

Remote server stores $\mathcal{N}$ files, each has the following three characteristics:

(1) time-sensitive (e.g., web crawling)

(2) has its own popularity

(3) the popularity (i.e., the number of requests each time) is time varying.

Users request files, then pick them up in **the local cache.**

Local Cache stores copies of all the files (but probably not the freshest version). The local cache pulls file updates through a bandwidth limited **channel**. Due to bandwidth constraint, only a subset of files can be downloaded simultaneously.

Question: How to design update strategies that adapt to time-varying popularities, so that users can receive fresh files?

Freshness Model – Age of Information (AoI)

- Def: **AoI** measures time elapsed since the latest file version has been updated at the local cache.

- Let $u_{n,t} = 1$ if file n is updated in slot t; $u_{n,t} = 0$ indicates file n is not updated in slot t.

- Let $X_{n,t}$ denote the AoI of file n at the beginning of slot t:
  - If file n is not updated, $u_{n,t} = 0$, then $X_{n,t+1} = X_{n,t} + 1$
  - If file n is updated, $u_{n,t} = 1$, then $X_{n,t+1} = 1$

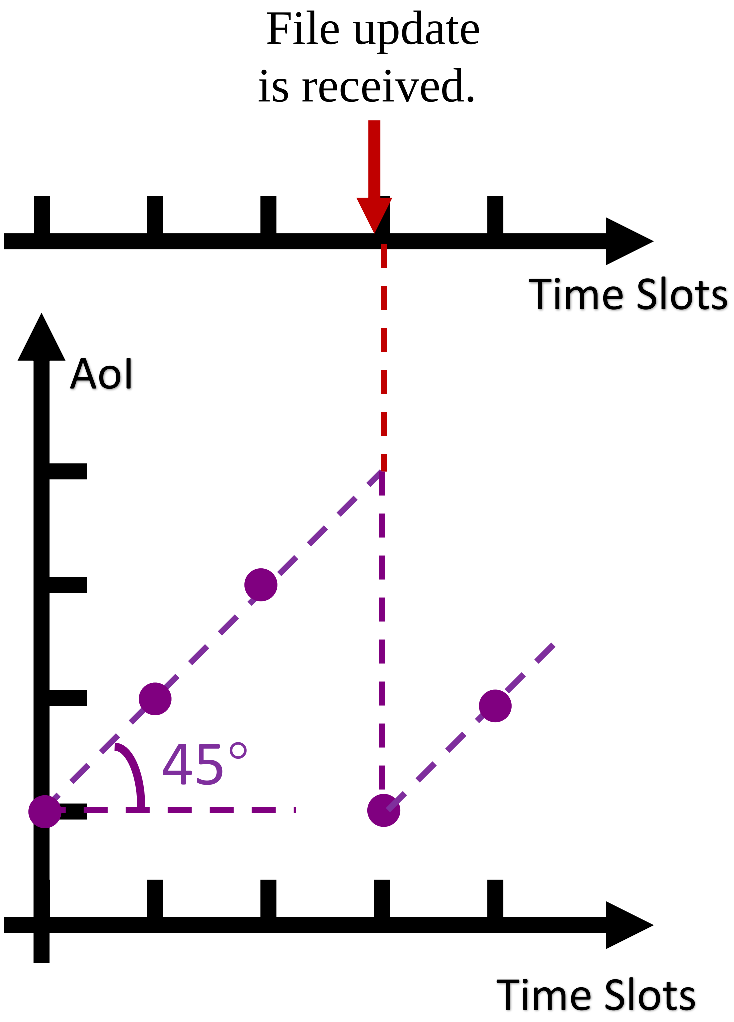

Popularity model and channel model

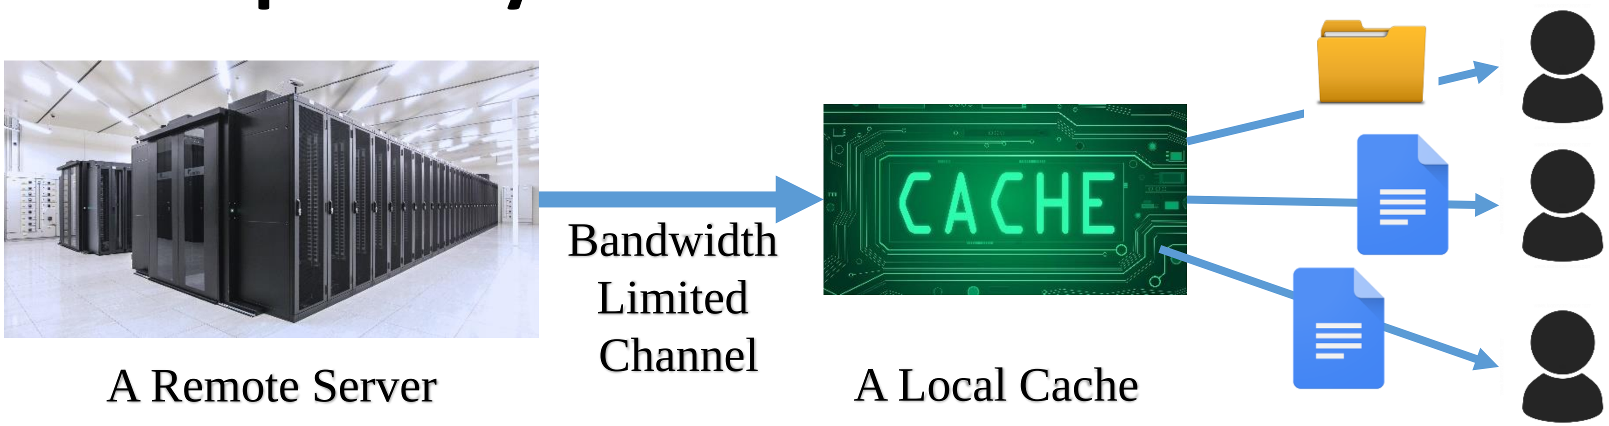

- Time Varying File Popularities:
  - ➢We characterize the popularity of file n by using **popularity mode** $R_{n,t}$
  - ➢The number of requests for file n in slot t depends on the $R_{n,t}$, which can be characterize by functions $\omega_n(R_{n,t})$.
  - ➢The popularity mode $R_{n,t}$ evolves like a Markov Chain, i.e.,

$$P^n_{r,r'} := \Pr\left(R_{n,t+1} = r' \middle| R_{n,t} = r\right), \forall (r, r'), t$$

- Bandwidth Limited Channel:
  - ➢In each slot, no more than M files can be updated.

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

Problem Formulation

- The goal is to design a policy π that designs update decisions $\{u_{n,t}\}$ for all the files, so that the expected AoI of the requested files can be minimized:

Problem 1: (Primal Problem)

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

$$\sum_{n=1}^{N} u_{n,t} \le M, \forall t.$$

AoI

This problem can be formulated as an MDP:

➢State: the AoI and popularity of all the files $\{X_{n,t}, R_{n,t}\}_{n=1}^{N}$

➢Action: all the update decisions that satisfy bandwidth constraint $\sum_{n=1}^{N} u_{n,t} \le M$.

➢Cost: the total AoI of all the requested files $\sum_{n=1}^{N}\omega_n\left(R_{n,t}\right)X_{n,t}$

Infinite State Space, Action space has cardinality $\mathcal{O}(N^M)$!
Value/Policy Iteration have high computational complexity.

We need scalable reformulations

Problem Resolution (1)- CMDP reformulation

Challenge: high computational complexity of multi-file MDP

Solution: Decouple into multiple single file MDP, then solve them separately

- Step 1: Bandwidth constraint relaxation

In each slot → **On average across all slots**, no more than M files can be updated.

Problem 2: (Relaxed Problem)

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

$$\sum_{n=1}^{N} u_n(t) \leq M, \forall t \Rightarrow \lim_{T\to\infty} \mathbb{E}_\pi\left[\frac{1}{T}\sum_{t=1}^{T}\sum_{n=1}^{N} u_{n,t}\right] \leq M.$$

➢Action space becomes larger: $\{u_{n,t}\}_{n=1}^{N} \in \{0,1\}^N$.

Still a hard CMDP with 2^N possible actions!

Problem Resolution (2)-Decomposition

- What is decoupled?
  - Action space: $u_{n,t} \in \{0, 1\}, \forall n, t$
  - Cost function: $a(\pi) = \sum_{n=1}^{N} a_n(\pi) = \sum_{n=1}^{N} \lim_{T\to\infty} \mathbb{E}_\pi \left[\frac{1}{T}\sum_{t=1}^{T} \omega_n\left(R_{n,t}\right)X_{n,t}\right]$
  - Probability transition function: $\left(X_{n,t}, R_{n,t}\right)$ depends only on $\left(X_{n,t-1}, R_{n,t-1}\right)$ and action $u_{n,t}$
- Only the constraint is not decoupled.
- Solution: Place the constraint into objective function using Lagrange multiplier:

$$\pi^*(W) = \arg\min L(\pi, W)$$

$$= \lim_{T\to\infty} \mathbb{E}_\pi \left[\frac{1}{T}\sum_{t=1}^{T}\sum_{n=1}^{N}\left(\omega_n\left(R_{n,t}\right)X_{n,t} + W u_{n,t}\right) - WM\right]$$

- For fixed W, per-File decomposition for min $L(\pi, W)$

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

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

Problem Resolution (3)-Solving Decoupled Problem

Problem 3: (Per-File decoupled problem)

$$\pi_F^{*,n}(W) = \arg\min_{\pi} \lim_{T\to\infty} \mathbb{E}_\pi \left[\frac{1}{T}\sum_{t=1}^{T} \omega_n(R_{n,t})X_{n,t} + Wu_{n,t}\right]$$

- The problem is also an Markov Decision Process (MDP) with infinite space. For each of the decoupled problem (we omit superscript n) we have the following results:

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

Result 2: The maximum threshold

$$X_{\max} = \max_r \tau_r \leq \max_r \left\lceil \frac{W + \omega(r)}{\omega(r)} \right\rceil := X_{ub}$$

Problem Resolution (3)-Solving Decoupled Problem

- For fixed W, we can search over all possible stationary policies π:
  - We introduce two auxiliary variables to represent π:
  - $\mu_{x,r}^{n}$ the steady state probability file n in state (x, r)
  - $\nu_{x,r}^{n}$ the steady state probability file n in state (x, r) and is updated
  - Rewrite expected AoI as a function of $\{\mu_{x,r}^{n}, \nu_{x,r}^{n}\}_{x=1}^{X_{ub}}$
- Linear Programming objectives:

$$\lim_{T\to\infty} \mathbb{E}_{\pi}\left[\frac{1}{T}\sum_{t=1}^{T} \omega_n(R_{n,t})X_{n,t} + W u_{n,t}\right] = \sum_{x=1}^{X_{ub}}\sum_{r=1}^{R}\left(\omega(r)x\mu_{x,r}^{n} + W\nu_{x,r}^{n}\right)$$

- Bandwidth consumptions:

$$d_n^*(W) = \lim_{T\to\infty} \mathbb{E}_{\pi}\left[\frac{1}{T}\sum_{t=1}^{T} u_{n,t}\right] = \sum_{x=1}^{X_{ub}}\sum_{r=1}^{R} \nu_{x,r}^{n,*}(W)$$

Solve the LP for each file n and obtain the optimum solution $\{\mu_{x,r}^{n,*}(W), \nu_{x,r}^{n,*}(W)\}$.

The optimum policy $\pi_F^*(W)$ updates file n in (x, r) with probabity $\xi_{x,r}^{n,*}(W) = \frac{\nu_{x,r}^{n,*}(W)}{\mu_{x,r}^{n,*}(W)}$

$$W\uparrow, d^*(W) = \sum_{n=1}^{N} d_n^*(W) \text{ (bandwidth consumption)}\downarrow$$

Problem Resolution (4)-Solving Relaxed Problem

Back to the relaxed constraint problem

$$\lim_{T\to\infty} \mathbb{E}_\pi \left[\frac{1}{T}\sum_{t=1}^{T}\sum_{n=1}^{N} u_n(t)\right] \le M$$

- We search for policy π_R^* by adjusting W to satisfy the relaxed constraint.
- Let $W^* = \inf_W \{W | d^*(W) \le M\}$

Case 1: $d^*(W^*) = M$

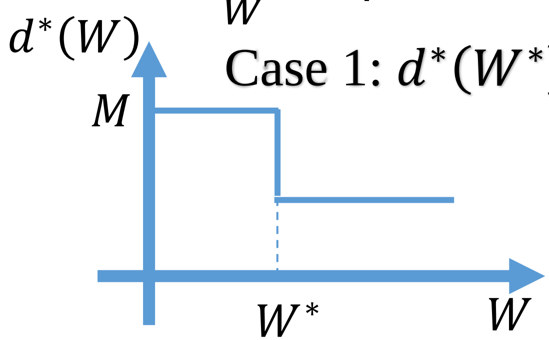

$\overline{\xi}_{x,r}^{n,*} = \xi_{x,r}^{n,*}(W^*)$

Case 2: $d^*(W^*) \neq M$,

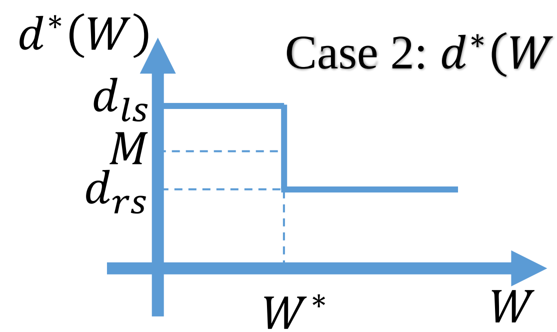

Let $d_{ls} = \lim_{W\uparrow W^*} d^*(W)$ and $d_{rs} = \lim_{W\downarrow W^*} d^*(W)$

Compute $\{\overline{\mu}_{x,r}^{n,*}, \overline{\nu}_{x,r}^{n,*}\}$ from $\{\mu_{x,r}^{n,ls}, \nu_{x,r}^{n,ls}\}$ and $\{\mu_{x,r}^{n,rs}, \nu_{x,r}^{n,rs}\}$

Finally obtain $\overline{\xi}_{x,r}^{n,*} = \frac{\overline{\nu}_{x,r}^{n,*}}{\overline{\mu}_{x,r}^{n,*}}$

π_R^* : In slot t, suppose file n is in state (x, r), update file n with probability $\overline{\xi}_{x,r}^{n,*}$.

Problem Resolution (5)-Solving Primal Problem

Back to the original problem: minimize the expected average AoI $a(\pi)$ under constraint:

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

- We derive policy $\hat{\pi}$ that satisfy the hard constraint based on π_R^* for the relaxed problem:
  - If π_R^* updates less or equal M files, $\hat{\pi}$ updates all of them
  - If π_R^* updates more than M files, $\hat{\pi}$ chooses M randomly to update.

Theorem 2 (asymptotic optimality of $\hat{\pi}$) Assume $\frac{N}{M} = \theta$ fixed as a constant independent of N, policy $\hat{\pi}$ is asymptotic optimal in the sense:

$$\lim_{N\to\infty} \frac{a(\hat{\pi}) - a(\pi^*)}{a(\pi^*)} = 0$$

Simulations: Algorithm Comparisons

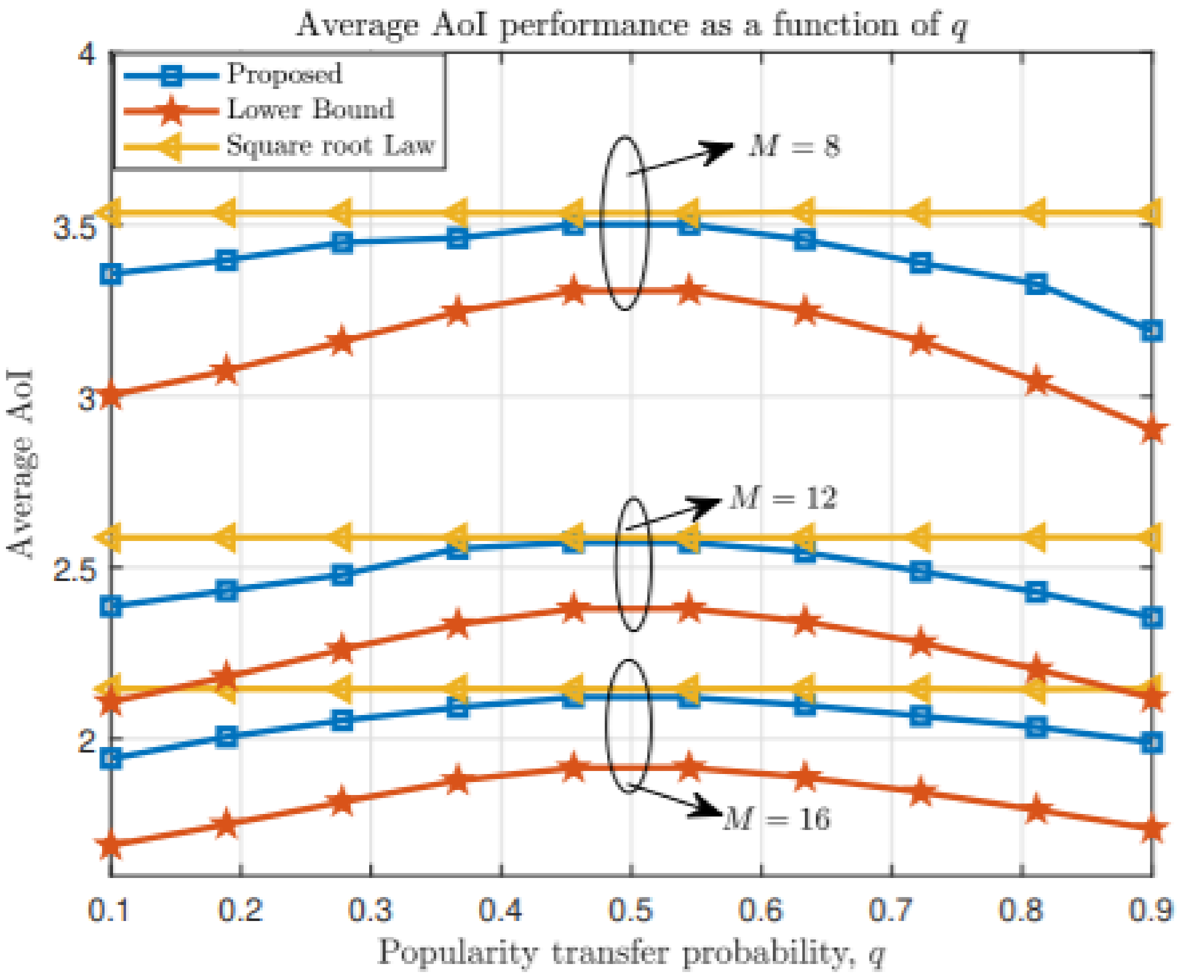

Assumption:
Each file has two popularity modes with transition matrix.

$$\mathbf{P}^n = \begin{bmatrix} q & 1-q \\ 1-q & q \end{bmatrix}$$

Square Root Law (sqrt) minimizes the expected AoI when the popularity of each file does not vary with time.

The expected AoI of proposed algorithm is close to the lower bound.

- When $q = 0.5$, popularity evolution do not depend on previous popularity mode: **sqrt and proposed algorithm have similar performance.**
- When $q \to 1$, $q \to 0$, popularity on the current slot highly depends previous state: **proposed algorithm has smaller AoI.**

Thank you! Q&A