Competitive Algorithms for Minimizing the Maximum Age-of-Information

Rajarshi Bhattacharjee
Indian Institute of Technology Madras
India
brajarshi91@gmail.com

Abhishek Sinha
Indian Institute of Technology Madras
India
abhishek.sinha@ee.iitm.ac.in *

ABSTRACT

In this short paper, we consider the problem of designing a near-optimal competitive scheduling policy to maximize the freshness of available information uniformly across N mobile users. Motivated by the unreliability and non-stationarity of the emerging 5G-mmWave channels for high-speed users, we forego of *any* statistical modeling assumptions of the wireless channels and user-mobility. Instead, we allow the channel states and the mobility patterns to be dictated by an omniscient adversary. It is not difficult to see that no competitive scheduling policy can exist for the corresponding throughput-maximization problem in this adversarial model. Surprisingly, we show that there exists a simple online distributed scheduling policy with a finite competitive ratio for maximizing the freshness of information in this model. We also prove that the proposed policy is competitively optimal up to an $O(\log N)$ factor.

1. INTRODUCTION

Apart from the throughput, maximizing the freshness of information available at the user-end is a principal design objective for the emerging 5G wireless standards. The *Age-of-Information* (AoI) is a newly-proposed metric that captures the information-freshness in a quantitative fashion [1]. However, the channel states and user-mobility are challenging to model and predict in high-speed non-stationary environments. This paper is concerned with the following question: Does there exist a scheduling policy that minimizes the maximum AoI across all users, irrespective of the channel dynamics and user-mobility patterns? The question is considerably general, as it does not make any assumptions on either the channel statistics or the user-mobility, both of which may be dictated by an omniscient adversary in the worst case. In this paper, we affirmatively answer the question by showing that a simple distributed scheduling policy is competitively optimal up to an $O(\log N)$ factor.

Closely related to this work, in a recent paper [2], we investigated the problem of minimizing the *average* AoI for N *static* users confined to a single cell. There we showed that the greedy Max-Age (MA) policy is competitively optimal up to a factor of $O(N)$ in the adversarial setting. In our previous paper [3], we proved that the MA policy is optimal for minimizing the maximum AoI for static users having stochastic binary erasure channels. Within the stochastic framework, the paper [4] proposes a Max-Weight policy, which is shown to be optimal up to a constant factor for minimizing the average AoI. Due to space constraints, we refer the reader to the book [1] for a comprehensive introduction to this active area of research and an extensive bibliography.

Contributions: In contrast with the existing works, this is the first paper to study the problem of minimizing the *maximum* AoI for *mobile* users in an *adversarial framework*. Our main results are summarized and contrasted with that of [2] in Table 1. On the technical side, compared to [2], the proof of achievability in Theorem 1 warrants the introduction of additional concepts of "Max-users" and "super-intervals". This is essential because, due to users' mobility, the simple round-robin structure of the scheduled users in [2] does not hold in this case. Furthermore, the proof of converse in Theorem 2 proceeds differently, making use of a maximal inequality similar to Massart's lemma.

2. SYSTEM MODEL

A set of N users move around in an area having a total of M Base Stations (BS). The coverage areas corresponding to each BS (*i.e.,* the cells) are disjoint. Time is slotted, and at every new slot, a user can either stay in its current cell or move to any other $M-1$ cells of its choice. Our mobility model is considerably general, as it does not make any statistical assumptions on the speed or user movement patterns. At each slot, all BSs receive a fresh update packet for each user from an external source (*e.g.,* a high-speed optical backbone). The fresh packets replace the stale packets in the BS buffers. Each BS can beamform and schedule a downlink fresh packet transmission to only one user under its coverage area at a slot. It is the scheduling decisions that we can control and optimize. The channel state for any user at any slot can be either Good or Bad. The BSs are assumed to be oblivious to the current channel state conditions (*i.e.,* no CSIT). If at any slot, a BS schedules a transmission to a user (under its coverage area) having Good channel, the user decodes the packet successfully. Otherwise, the packet is lost. In the worst case, the states of the N channels (corresponding to N users) and the user movements at every slot may be dictated by an *omniscient adversary* [5].

Cost Metric: We are concerned with competitively optimizing the information freshness for all users. Accordingly, we define the N-dimensional state-vector $\boldsymbol{h}(t)$, where $h_i(t)$ denotes the length of the time interval prior to time t before which the i^{th} user successfully received its most

*This work is supported by the grant IND-417880 from Qualcomm, USA.

MAMA 2020, Boston, MA, USA

Table 1: Summary of the results on the Competitive Ratios (η) in the adversarial framework

Metrics	Cost function	Mobility	Upper Bound	Lower Bound	Optimality gap
Average Age [2]	$N^{-1}\sum_{t=1}^{T}\sum_{i=1}^{N} h_i(t)$	No	$O(N^2)$	$O(N)$	$O(N)$
Maximum Age (This paper)	$\sum_{t=1}^{T}\max_{i=1}^{N} h_i(t)$	Yes	$O(N)$	$\Omega(\frac{N}{\ln(N)})$	$O(\ln(N))$

recent packet. The variable $h_i(t)$ is known as the *Age-of-Information* of the i^{th} user at time t [1]. Clearly, the plot of $h_i(t)$ has a saw-tooth shape that increases linearly with unit-slope (*i.e.,* $h_i(t+1) = h_i(t)+1$) until the i^{th} user receives a new packet. Upon reception of a new packet, $h_i(t)$ drops to 1. From that point onwards, $h_i(t)$ again continues increasing and repeats the saw-tooth pattern [4]. The cost $C(t)$ at time t is defined to be the maximum age among all users, *i.e.,* $C(t) = \max_{i=1}^{N} h_i(t)$. The cumulative cost incurred over a time-horizon of length T is defined as: $\mathsf{Cost}(T) = \sum_{t=1}^{T} C(t)$.

Performance index: As standard in the literature on online algorithms, we compare the performance of any online scheduling algorithm $\mathcal{A}$ against that of an optimal *offline* scheduling algorithm OPT using the notion of competitive ratio $\eta^{\mathcal{A}}$, defined as follows:

$$\eta^{\mathcal{A}} = \sup_{\boldsymbol{\sigma}} \left(\frac{\text{Cost of the online policy } \mathcal{A} \text{ on } \boldsymbol{\sigma}}{\text{Cost of the offline } \mathsf{OPT} \text{ policy on } \boldsymbol{\sigma}}\right). \tag{1}$$

In the above definition, the supremum is taken over all finite-length sequences $\boldsymbol{\sigma}$ that encodes the channel states and user locations at every slot. Note that, while the online policy $\mathcal{A}$ has only causal information, the policy OPT is assumed to be equipped with non-causal knowledge of the entire sequence $\boldsymbol{\sigma}$ right at the beginning. Our objective is to design an online scheduling policy $\mathcal{A}$ with a small competitive ratio.

3. ACHIEVABILITY

Consider the following distributed scheduling policy:
Cellular Max-Age (CMA): At every slot, each BS j schedules a downlink transmission to the user that has the highest age among all users located in BS j's coverage area at that slot (ties are broken in an arbitrary but fixed order).
Theorem 1 below states a performance guarantee for CMA, which is, quite surprisingly, independent of the number of BSs M.

THEOREM 1. $\eta^{\mathsf{CMA}} \leq 2N$.

PROOF: For any slot t, define the global *Max-user* that has the highest age among all N users (ties are broken identically as in the CMA policy). The identity of the Max-user changes with time. However, by definition, the CMA policy continues to schedule the user corresponding to the current Max-user *irrespective* of its location until the transmission to it is successful. In the subsequent slot, a different user assumes the role of the Max-user, and the process continues.

Let T_i be the time slot at which a total of i successful packet transmissions have just been made exclusively by the Max-users up to time T_i. Let $\Delta_i \equiv T_i - T_{i-1}$ denote the length of the i^{th} *super-interval* - defined as the time interval between the i^{th} and $i-1^{\text{th}}$ successful transmissions by the corresponding Max-user. The super-intervals are contiguous and disjoint. Let us denote the Max-user corresponding to the i^{th} super-interval by M_i. As argued above, the user M_i gets scheduled by the CMA policy persistently during the entire i^{th} super-interval of length Δ_i, irrespective of its locations during that period. Note that, unlike the case of static users [2], there could be more than one successful transmissions within a super-interval by users other than the Max-user.

We now claim that the Max-user M_i corresponding to the i^{th} super-interval must have a successful transmission by the beginning of the last $N-1$ super-intervals. If not, by the pigeonhole principle, some other user $j \neq M_i$ must become the Max-user at least twice in the previous N super-intervals. However, this cannot be true as then the user j would have less age than M_i when the user j became the Max-user for the second time in the previous N super-intervals.

Hence, at the k^{th} slot of the i^{th} super-interval, the age of the Max-user M_i is upper bounded by $k+\sum_{j=1}^{N-1}\Delta_{i-j}$, where for notational consistency, we have defined $T_j \equiv 0$, and $\Delta_j \equiv 0, \forall j \leq 0$. Thus, the cost C_i^{CMA} incurred by the CMA policy during the i^{th} interval may be upper-bounded as:

$$\begin{aligned} C_i^{\mathsf{CMA}} &\leq \sum_{k=1}^{\Delta_i}\left(k+\sum_{j=1}^{N-1}\Delta_{i-j}\right) = \frac{1}{2}(\Delta_i^2+\Delta_i) + \sum_{j=1}^{N-1}\Delta_i\Delta_{i-j} \\ &\stackrel{(a)}{\leq} \frac{1}{2}(\Delta_i^2+\Delta_i) + \frac{1}{2}\sum_{j=1}^{N-1}(\Delta_i^2+\Delta_{i-j}^2) \\ &= \frac{N}{2}\Delta_i^2 + \frac{1}{2}\Delta_i + \frac{1}{2}\sum_{j=1}^{N-1}\Delta_{i-j}^2. \end{aligned}$$

where in (a), we have used the AM-GM inequality to conclude $\Delta_i\Delta_{i-j} \leq \frac{1}{2}(\Delta_i^2+\Delta_{i-j}^2), 1 \leq j \leq N-1$. Hence, assuming that there are a total of K super-intervals in a time-horizon of length T, the total cost incurred by the CMA policy over the entire time horizon is upper bounded as:

$$\mathsf{Cost}^{\mathsf{CMA}}(T) = \sum_{i=1}^{K} C_i^{\mathsf{CMA}} \leq \frac{1}{2}\sum_{i=1}^{K}\left(2N\Delta_i^2+\Delta_i\right). \tag{2}$$

On the other hand, the cost incurred by the OPT policy during the i^{th} super-interval is trivially lower bounded by the sum of the ages of the user M_i during the super-interval. Note that, a Max-user consistently experience Bad channels throughout the corresponding super-interval. Hence,

$$C_i^{\mathsf{OPT}} \geq \sum_{k=1}^{\Delta_i}(1+k) = \frac{1}{2}\Delta_i^2 + \frac{3}{2}\Delta_i, \tag{3}$$

Finally, the cost of a entire horizon of length T may be obtained by summing up the cost incurred in the constituent super-intervals. Noting that $\Delta_0 = 0$, using Eqns. (2) and (3), the competitive ratio η^{CMA} of the CMA policy may be upper bounded as follows:

$$\eta^{\mathsf{CMA}} = \frac{\sum_{i=1}^{K} C_i^{\mathsf{CMA}}}{\sum_{i=1}^{K} C_i^{\mathsf{OPT}}} \stackrel{(a)}{\leq} \frac{\frac{1}{2}\sum_{i=1}^{K}\left(2N\Delta_i^2+\Delta_i\right)}{\sum_{i=1}^{K}\left(\frac{1}{2}\Delta_i^2+\frac{3}{2}\Delta_i\right)} \leq 2N. \quad \blacksquare$$

4. CONVERSE

THEOREM 2. *For any online policy $\mathcal{A}$, $\eta^{\mathcal{A}} \geq \Omega(\frac{N}{\ln N})$.*

PROOF: We establish a slightly stronger result by proving the lower bound for the particular case when all N users remain static at a single cell throughout the entire time-horizon. Using Yao's minimax principle, it is well-known that a lower bound to the competitive ratios of *all* deterministic online algorithms under *any* input channel state distribution $\boldsymbol{p}$ yields a lower bound to the competitive ratio in the adversarial framework, *i.e.,*

$$\eta^{\mathcal{A}} \geq \frac{\mathbb{E}_{\boldsymbol{\sigma}\sim\boldsymbol{p}}(\mathsf{Cost}\text{ of the best deterministic online policy})}{\mathbb{E}_{\boldsymbol{\sigma}\sim\boldsymbol{p}}(\mathsf{Cost}\text{ of }\mathsf{OPT})}. \quad (4)$$

To apply Yao's principle in our setting, we construct the following channel state distribution $\boldsymbol{p}$: at any slot t, a user is chosen independently and uniformly at random and assigned a Good channel. The rest of the $N-1$ users are assigned Bad channels for that slot. Hence, at any slot: $\mathbb{P}(\text{user } i\text{'s channel is } \mathsf{Good}) = 1/N$, and is Bad otherwise. The rationale behind the above choice of the channel state distribution will be clear when we compute OPT's expected cost. With our chosen channel distribution $\boldsymbol{p}$, we see that only one user's channel is in Good state at any slot. This greatly simplifies the computation of OPT's expected cost. We lower bound the competitive ratio using Eqn. (4) by lower-bounding the numerator and upper-bounding the denominator for the above channel state distribution.

An Upper bound to OPT's cost: The OPT policy, with non-causal channel state information, schedules the user having a Good channel at a slot. Thus, the limiting distribution of the age of any user is Geometric ($\frac{1}{N}$), *i.e.,*

$$\lim_{t\to\infty} \mathbb{P}\big(h_i(t)=k\big) = \frac{1}{N}\left(1-\frac{1}{N}\right)^{k-1}, \quad k\geq 1, \forall i.$$

Hence, for upper-bounding the time-averaged cost of OPT using Cesàro's summation formula, we are required to upper-bound the expected value of maximum of N *dependent* and identically Geometrically distributed random variables. The MGF of the Geometric distribution G is given by:

$$\mathbb{E}(\exp(\lambda G)) = \begin{cases} \frac{e^{\lambda}/N}{1-e^{\lambda}(1-1/N)}, & \text{if } \lambda < -\log(1-1/N) \\ \infty \quad \text{o.w.} \end{cases}$$

Let the r.v. $H_{\max}$ denote limiting max-age of the users. We proceed as in the proof of Massart's lemma for upper bounding $\mathbb{E}(H_{\max})$. For any $-\log(1-1/N) > \lambda > 0$, we have:

$$\exp\big(\lambda \mathbb{E}(H_{\max})\big) \stackrel{(a)}{\leq} \mathbb{E}(\exp(\lambda H_{\max})) \leq \sum_{i=1}^{N} \mathbb{E}(\exp(\lambda G_i)) \leq \frac{e^{\lambda}}{1-e^{\lambda}(1-\frac{1}{N})},$$

where the inequality (a) follows from Jensen's inequality. Taking natural logarithm of both sides, we get

$$\mathbb{E}(H_{\max}) \leq 1 - \frac{1}{\lambda}\log\big(1-e^{\lambda}(1-1/N)\big). \quad (5)$$

Now, let us choose $\lambda = \frac{\alpha}{N}$, for some fixed α $(0<\alpha<1)$ that will be determined later. First, we verify that, with this choice for λ, we always have $\lambda < -\log(1-\frac{1}{N})$. Using the convexity of the function e^x, we can write

$$1 = e^0 \geq e^x + (0-x)e^x = (1-x)e^x \Longrightarrow e^x \leq \frac{1}{1-x}, \forall x<1. \quad (6)$$

As a result, we have

$$e^{\lambda} \equiv e^{\frac{\alpha}{N}} \leq \frac{1}{1-\frac{\alpha}{N}} < \frac{1}{1-\frac{1}{N}}; \text{ i.e., } \lambda < -\log(1-\frac{1}{N}).$$

Next, for upper-bounding the RHS of Eqn. (5), we start with the simple analytical fact that for any $0<\alpha<1$,

$$\inf_{0<x<1} \frac{1-(1-x)e^{\alpha x}}{x} = 1-\alpha. \quad (7)$$

This result can be verified by using Eqn. (6) to conclude that for any $0<x<1$, we have

$$\frac{1-(1-x)e^{\alpha x}}{x} \geq \frac{1}{x}\big(1-\frac{1-x}{1-\alpha x}\big) = \frac{1-\alpha}{1-\alpha x} \geq 1-\alpha,$$

where the infimum is achieved when $x \to 0^+$. Substituting $x=\frac{1}{N}$ in the inequality (7), we have the following bound

$$1-e^{\alpha/N}(1-1/N) \geq \frac{1-\alpha}{N}.$$

Hence, using Eqn. (5), we have the following upper bound to the expected Max-age under OPT:

$$\mathbb{E}(H_{\max}) \leq 1 + \frac{N}{\alpha}\ln\frac{N}{1-\alpha}, \text{ for some } 0<\alpha<1.$$

Setting $\alpha = 1-\frac{1}{\ln N}$ yields the following asymptotic bound:

$$\mathbb{E}(H_{\max}) \leq N\ln N + o(N\ln N). \quad (8)$$

A Lower Bound to the cost of any online policy $\mathcal{A}$: Using Theorem 1 of [3], that gives the cost of the best online policy, with the parameters $p_i = \frac{1}{N}, \forall i$, we have:

$$\liminf_{T\to\infty} \frac{1}{T}\mathbb{E}\mathsf{Cost}^{\mathcal{A}}(T) = \liminf_{T\to\infty} \frac{1}{T}\sum_{t=1}^{T}\mathbb{E}(\max_i h_i^{\mathcal{A}}(t)) \geq N^2, \forall \mathcal{A}. \quad (9)$$

Combining Eqns. (8) and (9) with Eqn. (4) and using Cesàro's summation formula, we have for any online policy $\mathcal{A}$:

$$\eta^{\mathcal{A}} \geq \sup_{T>0} \frac{\mathbb{E}\mathsf{Cost}^{\mathcal{A}}(T)}{\mathbb{E}\mathsf{Cost}^{\mathsf{OPT}}(T)} \geq \frac{\liminf_{T\to\infty} \mathbb{E}\mathsf{Cost}^{\mathcal{A}}(T)/T}{\limsup_{T\to\infty} \mathbb{E}\mathsf{Cost}^{\mathsf{OPT}}(T)/T} \geq \Omega(\frac{N}{\ln N}). \quad \blacksquare$$

5. REFERENCES

[1] Yin Sun, Igor Kadota, Rajat Talak, and Eytan Modiano. Age of information: A new metric for information freshness. *Synthesis Lectures on Communication Networks*, 12(2):1–224, 2019.

[2] Subhankar Banerjee, Rajarshi Bhattacharjee, and Abhishek Sinha. Fundamental limits of age-of-information in stationary and non-stationary environments. *arXiv preprint arXiv:2001.05471 (to appear in ISIT 2020, LA, USA)*.

[3] Arunabh Srivastava, Abhishek Sinha, and Krishna Jagannathan. On minimizing the maximum age-of-information for wireless erasure channels. *arXiv:1904.00647 (appeared in RAWNET 2019)*.

[4] Igor Kadota, Abhishek Sinha, Elif Uysal-Biyikoglu, Rahul Singh, and Eytan Modiano. Scheduling policies for minimizing age of information in broadcast wireless networks. *IEEE/ACM Transactions on Networking*, 26(6):2637–2650, 2018.

[5] Matthew Andrews and Lisa Zhang. Routing and scheduling in multihop wireless networks with time-varying channels. *ACM Transactions on Algorithms (TALG)*, 3(3):33–es, 2007.