

Age-of-Information Bandits

Kavya Bhandari, Santosh Fatale, Urvidh Narula, Sharayu Moharir, and Manjesh Kumar Hanawal
Indian Institute of Technology Bombay

Abstract—We consider a system with a single source that measures/tracks a time-varying quantity and periodically attempts to report these measurements to a monitoring station. Each update from the source has to be scheduled on one of K available communication channels. The probability of success of each attempted communication is a function of the channel used. This function is unknown to the scheduler.

The metric of interest is the Age-of-Information (AoI), formally defined as the time elapsed since the destination received the recent most update from the source. We model our scheduling problem as a variant of the multi-arm bandit problem with communication channels as arms. We characterize a lower bound on the AoI regret achievable by any policy and characterize the performance of UCB, Thompson Sampling, and their variants. In addition, we propose novel policies which, unlike UCB and Thompson Sampling, use the current AoI to make scheduling decisions. Via simulations, we show the proposed AoI-aware policies outperform existing AoI-agnostic policies.

I. INTRODUCTION

We consider a learning problem that focuses on the metric of Age of Information (AoI), introduced in [1]. AoI is formally defined as the time elapsed since the destination received the recent most update from the source. It follows that AoI is a measure of the freshness of the data available at the intended destination which makes it a suitable metric for time-sensitive systems like smart homes, smart cars, and other IoT based systems. Since its introduction, AoI has been used in areas like caching, scheduling, energy harvesting, and channel state information estimation. A comprehensive survey of AoI-based works is available in [2].

We focus on a system consisting of a single source that measures/tracks a time-varying quantity. The source updates a monitoring station by sending periodic updates using any one of K available communication channels at a given time (Figure 1). The probability of an attempted update succeeding is independent across communication channels and independent and identically distributed (i.i.d.) across time-slots for each channel. Channel statistics are unknown to the scheduler. AoI increases by one on each failed update and resets to one on each successful update. The goal is to determine which communication channel to use in each time-slot in order to minimize the cumulative AoI over a finite time-interval of T consecutive time-slots. We view our work as a key first step towards studying real IoT-type systems which have multiple sensors updating a central monitoring station via multiple communication channels.

Like the standard multi-arm bandit (MAB) problem and its numerous variants, our scheduling problem experiences a trade-off between exploring the various communication channels and exploiting the most promising communication




Fig. 1: A system consisting of a source, a monitoring station, and five communication channels. The source tracks a time-varying quantity and sends periodic updates to the monitoring station using any one of the five channels for each update.

channel, as observed from past observations. Henceforth, we refer to our problem as *AoI bandits*. The pseudo-regret of a policy at time T as the difference between the cumulative AoI in the first T time-slots under that policy and the cumulative AoI in the first T time-slots by the “genie” policy which uses the (statistically) best channel in each time-slot.

The key difference between our problem and the classical MAB problem is that since AoI is correlated across time-slots, the scheduling decision made in a time-slot has a cascading effect on the regret accumulated in future time-slots. Variants of the classical MAB problem like queuing bandits [3], [4] also exhibit this characteristic. This time-correlation has significant implications for both algorithm design and analysis.

Since the potential regret accumulated in a time-slot is a function of the current AoI, it is crucial to incorporate the current AoI in making scheduling decisions. We refer to policies that do this as *AoI-aware policies* and refer to policies that do not incorporate this information into their decision making as *AoI-agnostic policies*. Popular policies like UCB and Thompson Sampling are AoI-agnostic as they make decisions based only on the number of times each channel is used and the total number of successful transmissions on each channel. The performance analysis of policies for AoI bandits requires a novel approach where we upper bound the regret accumulated in all future time-slots as a result of using a sub-optimal channel in a time-slot.

A. Our Contributions

Lower bound on AoI regret: We show that the AoI regret of any α -consistent policy is $\Omega(K \log T)$.

Performance of AoI-agnostic policies: We show that the AoI regret of UCB [5] and Thompson Sampling [6] is $O(K \log^2 T)$ and the AoI regret of Q-UCB [3] and Q-Thompson Sampling [3] is $O(K \log^4 T)$.

New AoI-aware policies: We propose variants of UCB, Thompson Sampling, Q-UCB, and Q-Thompson Sampling which work in two phases. When AoI is “low”, the variants mimic the corresponding original policies and when AoI is “high”, the variants only exploit based on past observations. Via simulations, we show that the proposed variants outperform the original AoI-agnostic policies.

B. Related Work

In this section, we focus on AoI based work most relevant to our setting. Scheduling to minimize AoI has been explored in a variety of settings [7]–[12]. The key difference between these works and our work is that in these works, channel statistics and/or channel state information is assumed to be known, whereas we work in the setting where channel statistics are unknown and have to be learned. In addition, some of these works focus on the infinite time-horizon and evaluate the steady-state performance, whereas we provide finite-time guarantees.

A multi-arm bandit based approach to scheduling problems to minimize queue-length is the focus of [3], [4], [13]–[21]. The focus in [13]–[21] is on the infinite horizon problem, whereas [3], [4] focus on the finite horizon setting. The key difference between [3], [4] and our work is that their metric is queue-length regret whereas we focus on AoI regret. We evaluate the performance of policies proposed in [3] for our metric. The policies proposed in [4] cannot be applied in our setting due to the difference in the evolution of queue-length and AoI.

II. SETTING

A. Our System

We consider a system with a source and a monitoring station. The source tracks/measures a time-varying quantity and relays its measurements to the monitoring station via K communication channels as shown in Figure 1. We use C_i , $1 \leq i \leq K$ to denote the K channels. Time is divided into slots. In each time-slot, the source attempts to update the monitoring station by sending its current measurement via one of the K communication channels. Each attempted communication via C_i is successful with probability μ_i and unsuccessful otherwise, independent of all other channels and across time-slots. The values of the μ_i s are unknown to the scheduler. For a problem instance $\boldsymbol{\mu}$, let $\mu_{\min} = \min_{i=1:K} \mu_i$, $k^* = \arg \max_{k=1:K} \mu_k$, and $\mu^* = \mu_{k^*}$. We assume that $\mu_{\min} > 0$.

B. Metric: Age-of-Information Regret (AoI Regret)

The age-of-information is a metric that measures the freshness of information available at the monitoring station. It is formally defined as follows.

Definition 1 (Age-of-Information (AoI)). *Let $a(t)$ denote the AoI at the monitoring station in time-slot t and $u(t)$ denote the index of the time-slot in which the monitoring station received*

the latest update from the source before the beginning of time-slot t . Then,

$$a(t) = t - u(t).$$

By definition,

$$a(t) = \begin{cases} 1 & \text{if the update in slot } t-1 \text{ succeeds} \\ a(t-1) + 1 & \text{otherwise.} \end{cases}$$

Let $a_{\mathcal{P}}(t)$ be the AoI in time-slot t under a given policy \mathcal{P} , and let $a^*(t)$ be the corresponding AoI under the genie policy that always uses the optimal channel, i.e., C_{k^*} , where $k^* = \arg \max_{1 \leq k \leq K} \mu_k$. We define the AoI regret at time T as the cumulative difference in expected AoI for the two policies in time-slots 1 to T .

Definition 2 (Age-of-Information Regret (AoI Regret)). *AoI regret under policy \mathcal{P} is denoted by $R_{\mathcal{P}}(T)$ and*

$$R_{\mathcal{P}}(T) = \sum_{t=1}^T \mathbb{E}[a_{\mathcal{P}}(t) - a^*(t)]. \quad (1)$$

For concreteness and technical convenience, we make the following assumption on the initial state of the system.

Assumption 1 (Initial Conditions). *The system starts operating in time-slot $t = -\infty$ and the source sends an update to the monitoring station using one channel in each time-slot. Any candidate policy starts making scheduling decisions at $t = 1$. The policy does not use information from observations in time-slots $t \leq 0$ to make decisions in time-slots $t \geq 1$.*

The goal is to design a scheduling policy/algorithm¹ which minimizes AoI regret (Definition 2).

III. MAIN RESULTS AND DISCUSSION

In this section, we state and discuss our main results. A summary of our analytical results is provided in Table I. In addition to this, we propose new policies and compare the performance with known policies via simulations.

Algorithm	Regret
Any α -consistent policy	$\Omega(K \log T)$
UCB [5]	$O(K \log^2 T)$
Thompson Sampling [6]	$O(K \log^2 T)$
Q-UCB [3]	$O(K \log^4 T)$
Q-Thompson Sampling [3]	$O(K \log^4 T)$

TABLE I: Summary of our analytical results

A. Lower Bound on AoI Regret

We first provide a lower bound on the performance of any α -consistent policy defined as follows.

Definition 3. (α -consistent policies [22]) *Let $k(s)$ denote the index of the channel scheduled in time-slot s . A scheduling*

¹We use the terms policy and algorithm interchangeably.

policy is said to be an α -consistent policy for $\alpha \in (0, 1)$, if for any channel success probability vector $\boldsymbol{\mu}$, there exists a constant $C(\boldsymbol{\mu})$ such that

$$\mathbb{E} \left[\sum_{s=1}^t \mathbb{1}\{k(s) = k^*\} \right] \leq C(\boldsymbol{\mu})t^\alpha, \quad \forall k \neq k^*.$$

Theorem 1. (Lower Bound) For any α -consistent policy \mathcal{P} ,

$$R_{\mathcal{P}}(T) \geq \frac{(K-1)D(\boldsymbol{\mu})}{\mu^*} ((1-\alpha) \log T - \log(4KC)),$$

where $D(\boldsymbol{\mu}) = \frac{\Delta}{KL(\mu_{\min}, \frac{\mu^*+1}{2})}$, $\Delta = \mu^* - \max_{k \neq k^*} \mu_k$.

We thus conclude that the AoI regret of any α -consistent policy scales as $\Omega(K \log T)$.

B. AoI Regret of Popular AoI-agnostic Policies

Definition 4 (AoI-Agnostic Policies). A policy is AoI-agnostic if, given past scheduling decisions and the number of successful updates sent via each of the K channels in the past, it does not explicitly use age information in a time-slot to make scheduling decisions.

We now characterize the performance of four known AoI-agnostic policies, namely, UCB [5], Thompson Sampling (TS) [6], Q-UCB [3], and Q-Thompson Sampling (Q-TS) [3].

The UCB and Thompson Sampling policies are known to perform well for MAB and Q-UCB, and Q-Thompson Sampling are variants of UCB and Thompson Sampling respectively, proposed in [3] for queueing bandits. The difference between the original policies and their ‘‘Q-’’ variants is that, in each time-slot, the variants force the policy to explore with a probability which decays with time, similar to the ϵ -greedy policy proposed in [5]. For the sake of completeness, we also provide a formal description of these policies.

Algorithm 1: UPPER CONFIDENCE BOUND (UCB)

1 **Initialise:** Set $\hat{\mu}_k = 0$ to be the estimated success probability of Channel k , $T_k(0) = 0 \forall k \in [K]$.
2 **while** $1 \leq t \leq K$ **do**
3 Schedule update on Channel $k(t) = t$
4 Receive rewards $X_{k(t)}(t) \sim \text{Ber}(\mu_{k(t)})$
5 $\hat{\mu}_{k(t)} = X_{k(t)}(t)$
6 $T_{k(t)}(t) = 1$
7 $t = t + 1$
8 **while** $t \geq K + 1$ **do**
9 Schedule update on Channel $k(t)$ such that

$$k(t) = \arg \max_{k \in [K]} \hat{\mu}_k(t) + \sqrt{\frac{8 \log t}{T_k(t-1)}}$$

10 Receive reward $X_{k(t)}(t) \sim \text{Ber}(\mu_{k(t)})$
11 $\hat{\mu}_{k(t)} =$

$$\frac{(\hat{\mu}_{k(t)} \cdot T_{k(t)}(t-1) + X_{k(t)}(t))}{(T_{k(t)}(t-1) + 1)}$$

12 $T_{k(t)}(t) = T_{k(t)}(t-1) + 1$
12 $t = t + 1$

In the following we bound regret of each of these policies. For notational convenience write $c = \frac{-1}{\log(1-\mu^*)}$.

Theorem 2. (Performance of UCB) Consider any problem instance $\boldsymbol{\mu}$. Under Assumption 1, we have

$$R_{UCB}(T) \leq \begin{cases} \frac{1}{\mu_{\min}} + \frac{c \log T}{\mu_{\min}} \left(1 + (K-1) \times \left(\frac{32 \log T}{\Delta^2} + 1 + \frac{\pi^2}{3} \right) \right), & \text{for } T > K \\ \left(\frac{1}{\mu_{\min}} - \frac{1}{\mu^*} \right) T, & \text{for } T \leq K. \end{cases}$$

We thus conclude that AoI regret of UCB scales as $O(K \log^2 T)$. The proof of Theorem 2 first characterizes the AoI regret as a function of the expected number of times a sub-optimal channel is scheduled under UCB. The result then follows using a known upper bound from [5] on this quantity.

Remark 1. There is a gap between the lower bound obtained in Theorem 1 and the bound obtained for UCB in Theorem 2. In subsequent results, we see that a gap exists for TS, Q-UCB, and Q-TS as well. Closing this gap, either through tighter analytical results or designing a policy that provably meets the lower bound obtained in Theorem 1 is an open problem.

Algorithm 2: THOMPSON SAMPLING (TS)

1 **Initialise:** Set $\hat{\mu}_k = 0$ to be the estimated success probability of Channel k , $T_k(0) = 0 \forall k \in [K]$.
2 **while** $t \geq 1$ **do**
3 $\alpha_k(t) = \hat{\mu}_k(t)T_k(t-1) + 1,$
4 $\beta_k(t) = (1 - \hat{\mu}_k(t))T_k(t-1) + 1,$
5 For each $k \in [K]$, pick a sample $\hat{\theta}_k(t)$ of distribution,

$$\hat{\theta}_k(t) \sim \text{Beta}(\alpha_k(t), \beta_k(t)).$$

6 Schedule update on Channel $k(t)$ such that

$$k(t) = \arg \max_{k \in [K]} \hat{\theta}_k(t)$$

7 Receive reward $X_{k(t)}(t) \sim \text{Ber}(\mu_{k(t)})$
8 $\hat{\mu}_{k(t)} =$

$$\frac{(\hat{\mu}_{k(t)} \cdot T_{k(t)}(t-1) + X_{k(t)}(t))}{(T_{k(t)}(t-1) + 1)}$$

9 $T_{k(t)}(t) = T_{k(t)}(t-1) + 1$
10 $t = t + 1$

Theorem 3. (Performance of Thompson Sampling) Consider any problem instance $\boldsymbol{\mu}$. Then, under Assumption 1, we have

$$R_{TS}(T) \leq \begin{cases} \frac{1}{\mu_{\min}} + \frac{c \log T}{\mu_{\min}} (1 + O(K \log T)), & \text{for } T > K \\ \left(\frac{1}{\mu_{\min}} - \frac{1}{\mu^*} \right) T, & \text{for } T \leq K. \end{cases}$$

We thus conclude that AoI regret of TS scales as $O(K \log^2 T)$. The proof of Theorem 3 follows on the same lines as that of Theorem 2 using known results for TS [23].

Algorithm 3: Q-UPPER CONFIDENCE BOUND (Q-UCB)

```
1 Initialise: Set  $\hat{\mu}_k = 0$  to be the estimated success
  probability of Channel  $k$ ,  $T_k(0) = 0 \forall k \in [K]$ .
2 while  $t \geq 1$  do
3   let  $E(t) \sim \text{Ber}\left(\min\left\{1, 3K \frac{\log^2 t}{t}\right\}\right)$ 
4   if  $E(t) = 1$  then
5     Explore: Schedule update on a channel chosen
      uniformly at random
6   else
7     Exploit: Schedule update on Channel  $k(t)$  such
      that
      
$$k(t) = \arg \max_{k \in [K]} \hat{\mu}_k(t) + \sqrt{\frac{\log^2 t}{2T_k(t-1)}}$$

8   Receive reward  $X_{k(t)}(t) \sim \text{Ber}(\mu_{k(t)})$ 
9    $\hat{\mu}_{k(t)} =$ 
      $(\hat{\mu}_{k(t)} \cdot T_{k(t)}(t-1) + X_{k(t)}(t)) / (T_{k(t)}(t-1) + 1)$ 
10   $T_{k(t)}(t) = T_{k(t)}(t-1) + 1$ 
11   $t = t + 1$ 
```

Theorem 4. (Performance of Q-UCB) Consider any problem instance μ . Under Assumption 1, there exists a constant t_0 such that

$$R_{Q\text{-UCB}}(T) \leq \begin{cases} \frac{c \log T + 1 + cK \log^4 T + O\left(\frac{K}{T^2}\right)}{\left(\frac{1}{\mu_{\min}} - \frac{1}{\mu^*}\right)^{\mu_{\min}}} & \text{for } T > t_0 \\ \left(\frac{1}{\mu_{\min}} - \frac{1}{\mu^*}\right)^{\mu_{\min}} T, & \text{for } T \leq t_0. \end{cases}$$

We thus conclude that AoI regret of Q-UCB scales as $O(K \log^4 T)$. The proof of Theorem 4 first characterizes the AoI regret as a function of the expected number of times a sub-optimal channel is scheduled under Q-UCB. The result then follows using results in [3].

Theorem 5. (Performance of Q-TS) Consider any problem instance μ . There exists a constant t_0 such that

$$R_{Q\text{-TS}}(T) \leq \begin{cases} \frac{c \log T + 1 + cK \log^4 T + O\left(\frac{K}{T^2}\right)}{\left(\frac{1}{\mu_{\min}} - \frac{1}{\mu^*}\right)^{\mu_{\min}}} & \text{for } T > t_0 \\ \left(\frac{1}{\mu_{\min}} - \frac{1}{\mu^*}\right)^{\mu_{\min}} T, & \text{for } T \leq t_0. \end{cases}$$

We conclude that AoI regret of Q-Thompson Sampling scales as $O(K \log^4 T)$. The proof of Theorem 5 follows on the same lines as that of Theorem 4.

C. Our AoI-aware Policies

In this section, we propose AoI-aware variants of the four policies discussed in the previous section.

In the classical MAB with Bernoulli rewards, the contribution of a time-slot to the overall regret is upper bounded by one, but in AoI bandits it can be more than one. Also, unlike the MAB, for AoI bandits, the difference between AoIs under a candidate policy and the genie policy in a time-slot can be unbounded. This motivates the need to take the current AoI value into account when making scheduling decisions.

Algorithm 4: Q-THOMPSON SAMPLING (Q-TS)

```
1 Initialise: Set  $\hat{\mu}_k = 0$  to be the estimated success
  probability of Channel  $k$ ,  $T_k(0) = 0 \forall k \in [K]$ .
2 while  $t \geq 1$  do
3   let  $E(t) \sim \text{Ber}\left(\min\left\{1, 3K \frac{\log^2 t}{t}\right\}\right)$ 
4   if  $E(t) = 1$  then
5     Explore: Schedule update on a channel chosen
      uniformly at random
6   else
7     Exploit:
8      $\alpha_k(t) = \hat{\mu}_k(t)T_k(t-1) + 1,$ 
9      $\beta_k(t) = (1 - \hat{\mu}_k(t))T_k(t-1) + 1,$ 
10    For each  $k \in [K]$ , pick a sample  $\hat{\theta}_k(t)$  of
      distribution,
      
$$\hat{\theta}_k(t) \sim \text{Beta}(\alpha_k(t), \beta_k(t)).$$

      Schedule update on Channel  $k(t)$  such that
      
$$k(t) = \arg \max_{k \in [K]} \hat{\theta}_k(t)$$

11   Receive reward  $X_{k(t)}(t) \sim \text{Ber}(\mu_{k(t)})$ 
12    $\hat{\mu}_{k(t)} =$ 
      $(\hat{\mu}_{k(t)} \cdot T_{k(t)}(t-1) + X_{k(t)}(t)) / (T_{k(t)}(t-1) + 1)$ 
13    $T_{k(t)}(t) = T_{k(t)}(t-1) + 1$ 
14    $t = t + 1$ 
```

Intuitively, it makes sense to explore when AoI is low and exploit when AoI is high since the cost of making a mistake is much higher when AoI is high. We use this intuition to design AoI-aware policies. The key idea behind these policies is that they mimic the original policies when AoI is below a threshold and exploit when AoI is equal to or above a threshold, for an appropriately chosen threshold.

The first two policies (Algorithm 5 and Algorithm 6 in Appendix A of [24]) are variants of Thompson Sampling and UCB respectively. These policies maintain an estimate of the success probability of the best arm, denoted by $\hat{\mu}^*$. When AoI is not more than $\frac{1}{\hat{\mu}^*}$, the two policies mimic UCB and Thompson Sampling respectively, and exploit the “best” arm (based on past observations) otherwise. The third and fourth policies are variants of Q-UCB and Q-Thompson Sampling. When AoI is one, the two policies mimic Q-UCB and Q-Thompson Sampling respectively and exploit the “best” arm (based on past observations) otherwise. These policies are formally defined in Appendix A of in [24] (Algorithms 7 and 8).

In the next section, we compare the performance of all eight policies via simulations.

IV. SIMULATIONS

We present two sets of simulation results. In the first set of results (Figures 2, 3, and 4), we consider the first five parameter settings in Table II. We fix the number of

Algorithm 5: AOI-AWARE THOMPSON SAMPLING (AA-TS)

```

1 Initialise: Set  $\hat{\mu}_k = 0$  to be the estimated success
  probability of Channel  $k$ ,  $T_k(0) = 0 \forall k \in [K]$ .
2 while  $t \geq 1$  do
3    $\alpha_k(t) = \hat{\mu}_k(t)T_k(t-1) + 1$ ,
4    $\beta_k(t) = (1 - \hat{\mu}_k(t))T_k(t-1) + 1$ ,
5   Let  $\text{limit}(t) = \min_{k \in [K]} \frac{\alpha_k(t) + \beta_k(t)}{\alpha_k(t)}$ 
6   if  $a(t-1) > \text{limit}(t)$  then
7     Exploit: Select channel with highest estimated
      success probability
8   else
9     Explore:
10    For each  $k \in [K]$ , pick a sample  $\hat{\theta}_k(t)$  of
      distribution,
      
$$\hat{\theta}_k(t) \sim \text{Beta}(\alpha_k(t), \beta_k(t)).$$

      Schedule update on Channel  $k(t)$  such that
      
$$k(t) = \arg \max_{k \in [K]} \hat{\theta}_k(t)$$

11    Receive reward  $X_{k(t)}(t) \sim \text{Ber}(\mu_{k(t)})$ 
12     $\hat{\mu}_{k(t)} =$ 
      
$$(\hat{\mu}_{k(t)} \cdot T_{k(t)}(t-1) + X_{k(t)}(t)) / (T_{k(t)}(t-1) + 1)$$

13     $T_{k(t)}(t) = T_{k(t)}(t-1) + 1$ 
14     $t = t + 1$ 

```

arms to five and vary the range of success probabilities of these five arms. The success probability of the five arms is equally spaced in this range, for example, if the range is 0.1 to 0.3, the success probabilities for the five arms are $\{0.1, 0.15, 0.2, 0.25, 0.3\}$. In the second set of results (Figures 5, 6, and 7), we consider the last five parameter settings in Table II. We fix the range of success probabilities and vary the number of arms. As in the first set of simulations, the success probability of the arms is equally spaced in the specified range. Each reported data-point is the average value of 1000 independent iterations.

Setting	Range	Number of Arms (K)
1.a	[0.1; 0.3]	5
1.b	[0.1; 0.4]	5
1.c	[0.1; 0.5]	5
1.d	[0.1; 0.6]	5
1.e	[0.1; 0.7]	5
2.a	[0.05; 0.9]	2
2.b	[0.05; 0.9]	4
2.c	[0.05; 0.9]	6
2.d	[0.05; 0.9]	8
2.e	[0.05; 0.9]	10

TABLE II: Simulation parameters settings. The success probability of the arms is equally spaced in the specified range.

Due to space constraints, we show the time-evolution of regret for two of the five settings in each set. In addition, we show the regret at $T = 10000$ for all five settings in each set.

Consistent with expectations, AoI-agnostic policies are outperformed by their AoI-aware versions across all settings. The most notable observation is that AA-TS consistently outperforms all other policies, followed closely by TS. Also, the performance of AA-TS improves significantly relative to TS as the uniform gap between the success probabilities decreases, i.e., the optimal channel becomes harder to find. Notably, TS and Q-TS always outperform UCB and Q-UCB respectively, for all settings considered. Q-TS performs significantly worse than TS, but the same does not always hold true for Q-UCB and UCB respectively, which only seem to follow this trend for a sufficiently high success probability of the optimal channel.




Fig. 2: AoI regret as a function of time for Setting 1.b

V. PROOFS

In this section, we discuss the proofs of the results presented in Section III. Some details have been relegated to Appendix B of [24] due to space constraints.

A. Proof of Theorem 1

To prove this theorem, we construct an alternative service process described in [3], such that under any scheduling policy, the AoI evolution for this system has the same distribution as that for the original system. The service process is constructed as follows: let $\{U(t)\}_{t \geq 1}$ be i.i.d random variables distributed uniformly in $(0, 1)$. Let the service process for Channel k be given by $R_k(t) = \mathbb{1}\{U(t) \leq \mu_k\}$ for all t . Note that $\mathbb{E}[R_k(t)] = \mu_k$, i.e., the marginals of the service offered by each channel under this construction is the same as that in the original system.




Fig. 3: AoI regret as a function of time for Setting 1.e




Fig. 5: AoI regret as a function of time for Setting 2.b




Fig. 4: AoI regret at $T = 10000$ for Settings 1.a – 1.e




Fig. 6: AoI regret as a function of time for Setting 2.d

The proof of the claim that for any scheduling policy, the AoI evolution for this system with coupled service processes across channels has the same distribution as that for the original system follows using arguments from Section 8.1 in [3]. We use the following Lemma from [3] to prove Theorem 1.

Lemma 1 (Corollary 20, [3]). *Let $T_k(t)$ be the number of time-slots in which Channel k is used in the time-interval 1 to $t - 1$. For a problem instance $\boldsymbol{\mu}$, let $\mu_{\min} = \min_{i=1:K} \mu_i > 0$ and $\mu^* = \max_{i=1:K} \mu_i$. For any α -consistent policy \mathcal{P} , there exist constants τ and C , s.t. for any $t > \tau$,*

$$\Delta \sum_{k \neq k^*} \mathbb{E}[T_k(t+1)]$$

$$\geq (K-1)D(\boldsymbol{\mu})((1-\alpha)\log t - \log(4KC)),$$

$$\text{where } D(\boldsymbol{\mu}) = \frac{\Delta}{KL(\mu_{\min}, \frac{\mu^*+1}{2})}, \text{ and } \Delta = \mu^* - \max_{k \neq k^*} \mu_k.$$

Proof of Theorem 1. Let the AoI in time-slot t , under an α -consistent policy and the genie policy be denoted by $a(t)$ and $a^*(t)$ respectively. Let $S(t)$ and $S^*(t)$ be indicator random variables denoting successful updates in time-slot t by an α -consistent policy and the genie policy respectively. By




Fig. 7: AoI regret at $T = 10000$ for Settings 2.a – 2.e

definition,

$$\begin{aligned} a(t) &= (1 - S(t))(a(t-1) + 1) + S(t), \\ a^*(t) &= (1 - S^*(t))(a^*(t-1) + 1) + S^*(t). \end{aligned}$$

It follows that

$$\begin{aligned} a(t) - a^*(t) &= (1 - S(t))(a(t-1) + 1) + S(t) \\ &\quad - (1 - S^*(t))(a^*(t-1) + 1) - S^*(t). \end{aligned}$$

In the coupled system, $a^*(t) \leq a(t)$, for all t . Therefore,

$$a(t) - a^*(t) \geq (S^*(t) - S(t))(a^*(t-1)).$$

Taking expectations, it follows that

$$\mathbb{E}[a(t) - a^*(t)] \geq \mathbb{E}[S^*(t) - S(t)] \cdot \mathbb{E}[a^*(t-1)],$$

since $a^*(t-1)$ is independent of $S^*(t)$ and $S(t)$. Since the genie policy always uses the best channel, $a^*(t)$ is a geometric random variable with parameter μ^* . It follows that $\mathbb{E}[a^*(t)] = \frac{1}{\mu^*}$, and therefore,

$$R_{\mathcal{P}}(T) \geq \frac{1}{\mu^*} \sum_{t=1}^T \mathbb{E}[S^*(t) - S(t)]. \quad (2)$$

Let $Y_k(t)$ be an indicator random variable denoting if an update sent on Channel k in time-slot t will be successful. Let $Y^*(t)$ be an indicator random variable denoting if an update sent on the optimal channel in time-slot t will be successful. Let $k(t)$ be the index of the channel used by the α -consistent policy in time-slot t . It follows that

$$S^*(t) = Y^*(t) \text{ and } S(t) = \sum_{k=1}^K \mathbb{1}\{k(t) = k\} Y_k(t).$$

Therefore,

$$\mathbb{E}[S^*(t) - S(t)] = \mathbb{E} \left[\sum_{k \neq k^*} \mathbb{1}\{k(t) = k\} Y_k(t) \right]$$

$$\begin{aligned} &= \sum_{k \neq k^*} (\mathbb{P}(\mathbb{1}\{k(t) = k\} = 1) \mathbb{P}(\mu_k < U(t) \leq \mu^*)) \\ &= \sum_{k \neq k^*} (\mu^* - \mu_k) \mathbb{P}(\mathbb{1}\{k(t) = k\} = 1) \\ &\geq \Delta \sum_{k \neq k^*} \mathbb{P}(\mathbb{1}\{k(t) = k\} = 1). \end{aligned} \quad (3)$$

From (2) and (3),

$$\begin{aligned} R_{\mathcal{P}}(T) &\geq \frac{\Delta}{\mu^*} \sum_{t=1}^T \sum_{k \neq k^*} \mathbb{P}(\mathbb{1}\{k(t) = k\} = 1) \\ &= \frac{\Delta}{\mu^*} \sum_{k \neq k^*} \mathbb{E}[T_k(T+1)]. \end{aligned} \quad (4)$$

By Lemma 1 and (4),

$$R_{\mathcal{P}}(T) \geq \frac{(K-1)D(\boldsymbol{\mu})}{\mu^*} ((1-\alpha) \log T - \log(4KC)).$$

■

B. Proofs of Theorems 2 and 3

We use the following lemmas to prove Theorems 2 and 3. The proofs of these lemmas are available in Appendix B of [24]. In the first lemma, we bound AoI regret as a function of the number of times a sub-optimal channel is used.

Lemma 2. Let $k(t)$ denote the index of the communication channel used in time-slot t and k^* be the index of the optimal channel. Let $\mathbf{K}(T) = \{k(1), k(2), \dots, k(T)\}$ be the sequence of channels used in time-slots 1 to T and

$$N(\mathbf{K}(T)) = \sum_{t=1}^T \mathbb{1}_{k(t) \neq k^*},$$

denote the number of time-slots in which a sub-optimal channel is used. Then, under Assumption 1, for $c = \frac{-1}{\log(1-\mu^*)}$,

$$\sum_{t=1}^T \mathbb{E}[a(t)] \leq \frac{T}{\mu^*} + \frac{1}{\mu_{\min}} + \frac{c \log T}{\mu_{\min}} (1 + \mathbb{E}[N(\mathbf{K}(T))]).$$

The next lemma summarizes the results from Theorem 1 in [5] and Theorem 2 in [23] to provide upper bounds on the number of time-slots in which a sub-optimal channel is picked by UCB and Thompson Sampling.

Lemma 3. Let $k(t)$ denote the index of the communication channel used in time-slot t and k^* be the index of the optimal channel. Let $\mathbb{E}_{UCB}[N(\mathbf{K}(T))]$ and $\mathbb{E}_{TS}[N(\mathbf{K}(T))]$ denote the expected number of time-slots in which a sub-optimal channel is picked in time-slots 1 to T by UCB and Thompson Sampling respectively. Then, for $t > K$,

$$\mathbb{E}_{UCB}[N(\mathbf{K}(T))] \leq (K-1) \left(\frac{32 \log T}{\Delta^2} + 1 + \frac{\pi^2}{3} \right),$$

$$\mathbb{E}_{TS}[N(\mathbf{K}(T))] \leq O(K \log T),$$

where $\Delta = \mu^* - \max_{k \neq k^*} \mu_k$.

We now use Lemmas 2 and 3 to prove Theorems 2 and 3.

Proof. (Proof of Theorems 2 and 3) Note that by Assumption

$$1, \sum_{t=1}^T \mathbb{E}[a^*(t)] = \frac{T}{\mu^*}. \text{ From Lemma 2, we have that,}$$

$$\sum_{t=1}^T \mathbb{E}[a(t)] = \frac{T}{\mu^*} + \frac{1}{\mu_{\min}} + \frac{c \log T}{\mu_{\min}} (1 + \mathbb{E}[N(\mathbf{K}(T))]).$$

The results then follow by using Lemma 3. \blacksquare

C. Proof of Theorems 4 and 5

We use the following lemmas to prove Theorems 4 and 5. The proofs are available in Appendix B in [24].

Lemma 4. *Let $k(t)$ denote the index of the communication channel used in time-slot t and k^* be the index of the optimal channel. Let $\mathbf{K}(T) = \{k(1), k(2), \dots, k(T)\}$ be the sequence of channels used in time-slots 1 to T and E_t be the event that $k(\tau) = k^*$ for $t - c \log T + 1 \leq \tau \leq t$. Then, for $c = \frac{-1}{\log(1-\mu^*)}$,*

$$\sum_{t=1}^T \mathbb{E}[a(t)] \leq \frac{T}{\mu^*} + \frac{c \log T + 1}{\mu_{\min}} + \frac{1}{\mu_{\min}} \mathbb{E} \left[\sum_{t=c \log T + 1}^T \mathbb{1}_{E_t^c} \right].$$

Lemma 5. *Let E_t be the event that $k(\tau) = k^*$ for $t - c \log T + 1 \leq \tau \leq t$. Let $\mathbb{E}_{Q\text{-UCB}}[\cdot]$ and $\mathbb{E}_{Q\text{-TS}}[\cdot]$ denote expectation under the Q-UCB and Q-TS policies. Then,*

$$\mathbb{E}_{Q\text{-UCB}} \left[\sum_{t=c \log T + 1}^T \mathbb{1}_{E_t^c} \right] \leq cK \log^4 T + O\left(\frac{K}{T^2}\right),$$

$$\mathbb{E}_{Q\text{-TS}} \left[\sum_{t=c \log T + 1}^T \mathbb{1}_{E_t^c} \right] \leq cK \log^4 T + O\left(\frac{K}{T^2}\right).$$

Proof of Theorems 4 and 5. Follows by Lemmas 4 and 5 using the same arguments as the proof of Theorem 2. \blacksquare

VI. CONCLUSIONS

We consider a variant of MAB, called AoI bandits. We first characterize a lower bound on the regret achievable by any policy for AoI bandits. Next, we analyze the performance of popular policies, namely UCB and Thompson Sampling for our setting. In addition, we analyze the performance of two policies, namely, Q-UCB and Q-Thompson Sampling proposed in [3]. The commonality between these four policies is that they are AoI-agnostic, i.e., conditioned on the number of times each channel is used in the past and the number of successful communications on each channel, these policies make decisions independent of the current AoI. We then propose four AoI-aware policies, which also take the current value of AoI into account while making decisions. Via simulations, we observe that the AoI-aware policies outperform the AoI-agnostic policies.

VII. ACKNOWLEDGEMENTS

Sharayu Moharir's work was supported in part by an Indo-French grant on "Machine Learning for Network Analytics" and a seed grant from IIT Bombay.

REFERENCES

- [1] S. Ioannidis, A. Chaintreau, and L. Massoulie, "Optimal and scalable distribution of content updates over a mobile social network," in *IEEE INFOCOM 2009*. IEEE, 2009, pp. 1422–1430.
- [2] A. Kosta, N. Pappas, V. Angelakis *et al.*, "Age of information: A new concept, metric, and tool," *Foundations and Trends® in Networking*, vol. 12, no. 3, pp. 162–259, 2017.
- [3] S. Krishnasamy, R. Sen, R. Johari, and S. Shakkottai, "Regret of queueing bandits," in *Advances in Neural Information Processing Systems*, 2016, pp. 1669–1677.
- [4] T. Stahlbuhk, B. Shrader, and E. Modiano, "Learning algorithms for minimizing queue length regret," in *2018 IEEE International Symposium on Information Theory (ISIT)*. IEEE, 2018, pp. 1001–1005.
- [5] P. Auer, N. Cesa-Bianchi, and P. Fischer, "Finite-time analysis of the multiarmed bandit problem," *Machine learning*, vol. 47, no. 2-3, pp. 235–256, 2002.
- [6] W. R. Thompson, "On the likelihood that one unknown probability exceeds another in view of the evidence of two samples," *Biometrika*, vol. 25, no. 3/4, pp. 285–294, 1933.
- [7] B. Sombabu and S. Moharir, "Age-of-Information Aware Scheduling for Heterogeneous Sources," in *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking*, ser. MobiCom '18. New York, NY, USA: ACM, 2018, pp. 696–698. [Online]. Available: <http://doi.acm.org/10.1145/3241539.3267734>
- [8] V. Tripathi and S. Moharir, "Age of Information in Multi-Source Systems," in *GLOBECOM 2017-2017 IEEE Global Communications Conference*. IEEE, 2017, pp. 1–6.
- [9] V. Tripathi and E. Modiano, "A Whittle Index Approach to Minimizing Functions of Age of Information," *arXiv preprint arXiv:1908.10438*, 2019.
- [10] P. R. Jhunjhunwala and S. Moharir, "Age-of-Information aware scheduling," *SPCOM*, 2018.
- [11] Y.-P. Hsu, E. Modiano, and L. Duan, "Scheduling Algorithms for Minimizing Age of Information in Wireless Broadcast Networks with Random Arrivals: The No-Buffer Case," *arXiv preprint arXiv:1712.07419*, 2017.
- [12] I. Kadota, A. Sinha, and E. Modiano, "Optimizing age of information in wireless networks with throughput constraints," in *Proc. INFOCOM*, 2018.
- [13] D. R. Cox and W. Smith, *Queues*.
- [14] C. Buyukkoc, P. Varaiya, and J. Walrand, "The $c\mu$ rule revisited," *Advances in applied probability*, vol. 17, no. 1, pp. 237–238, 1985.
- [15] J. A. Van Mieghem, "Dynamic scheduling with convex delay costs: The generalized $c-\mu$ rule," *The Annals of Applied Probability*, pp. 809–833, 1995.
- [16] C. Lott and D. Teneketzis, "On the optimality of an index rule in multichannel allocation for single-hop mobile networks with multiple service classes," *Probability in the Engineering and Informational Sciences*, vol. 14, no. 3, pp. 259–297, 2000.
- [17] U. Ayesta, P. Jacko, and V. Novak, "Scheduling of multi-class multi-server queueing systems with abandonments," *Journal of Scheduling*, vol. 20, no. 2, pp. 129–145, 2017.
- [18] J. Niño-Mora, "Dynamic priority allocation via restless bandit marginal productivity indices," *Top*, vol. 15, no. 2, pp. 161–198, 2007.
- [19] A. Mahajan and D. Teneketzis, "Multi-armed bandit problems," in *Foundations and applications of sensor management*. Springer, 2008, pp. 121–151.
- [20] M. Larnaaga, U. Ayesta, and I. M. Verloop, "Dynamic control of birth-and-death restless bandits: Application to resource-allocation problems," *IEEE/ACM Transactions on Networking*, vol. 24, no. 6, pp. 3812–3825, 2016.
- [21] J. Gittins, K. Glazebrook, and R. Weber, *Multi-armed bandit allocation indices*. John Wiley & Sons, 2011.
- [22] T. L. Lai and H. Robbins, "Asymptotically efficient adaptive allocation rules," *Advances in applied mathematics*, vol. 6, no. 1, pp. 4–22, 1985.
- [23] E. Kaufmann, N. Korda, and R. Munos, "Thompson sampling: An asymptotically optimal finite-time analysis," in *International conference on algorithmic learning theory*. Springer, 2012, pp. 199–213.
- [24] K. Bhandari, S. Fatale, U. Naruala, S. Moharir, and M. K. Hanawal, "Age-of-Information Bandits." [Online]. Available: <https://arxiv.org/abs/2001.09317>