

Optimizing Age-of-Information in Adversarial and Stochastic Environments

Abhishek Sinha* and Rajarshi Bhattacharjee†

*School of Technology and Computer Science, Tata Institute of Fundamental Research, Mumbai 400 005, India

† College of Information and Computer Sciences, University of Massachusetts at Amherst

Email: *abhishek.sinha@tifr.res.in, †rbhattacharj@cs.umass.edu

Abstract—We design efficient online scheduling policies to maximize the freshness of information delivered to the users in a cellular network under both adversarial and stochastic channel and mobility assumptions. The information freshness achieved by a policy is investigated through the lens of a recently proposed metric - *Age-of-Information* (AoI). We show that a natural greedy scheduling policy is competitive against any optimal offline policy in minimizing the AoI in the adversarial setting. We also derive universal lower bounds to the competitive ratio achievable by any online policy in the adversarial framework. In the stochastic setting, we show that a simple index policy is near-optimal for minimizing the average AoI in two different mobility scenarios. Further, we prove that the greedy scheduling policy minimizes the peak AoI for static users in the stochastic setting. Simulation results show that the proposed policies perform well under realistic assumptions.

Index Terms—Age-of-Information, Competitive Analysis, Optimal Scheduling, Fundamental Limits

I. INTRODUCTION

QUALITY-of-Service (QoS) offered by data networks has been traditionally evaluated along three major dimensions, namely, throughput, latency, and energy/spectral efficiency. There is an extensive body of literature on cross-layer resource allocation algorithms to optimize the above metrics in various wired and wireless networks [4]–[8]. However, it has been widely argued that the standard QoS metrics are primarily geared towards quantifying the degree of utilization of system resources and less towards measuring the actual user experience [9]. With the explosive growth of hand-held mobile devices and the advent of the Internet of Things (IoT) and real-time AR and VR systems, future networks aim to optimize the Quality of Experience (QoE) for its end-users [10]. In order to incorporate the QoE objectives directly into the design decisions, a fundamentally new metric, called *Age-of-Information* (AoI), has been proposed recently for measuring the freshness of information available to the users [11], [12]. Informally, the AoI of a user is defined as the average length of time between its successive information updates. From the perspective of a network operator, minimizing either the average or the peak AoI of the users may be of interest [13]. For example, in the case of non-critical status updates, it makes sense to minimize the average AoI of all users. On the other hand, in applications

where the most outdated node is the bottleneck (*e.g.*, mission critical or industrial IoT applications), minimizing the peak AoI across all devices is a reasonable objective. Apart from uncertain wireless channels and bandwidth constraints, the mobility of the users poses significant challenges to optimize the AoI. It is well-known that mobility increases the capacity of wireless ad hoc networks [14]–[16]. Furthermore, the effect of mobility on the QoS, latency, and capacity of wireless cellular networks have been investigated in the papers [17]–[19]. However, to the best of our knowledge, the effect of user mobility on the freshness of information has not been previously explored. In this paper, we design near-optimal scheduling policies for minimizing the average and peak AoI in a cellular network under two widely different channel and mobility scenarios.

For analytical tractability, most of the existing papers on wireless communication work with stationary channel models [20]. In rapidly varying environments, such as high-speed trains and vehicle-to-vehicle communication, the stationarity assumption no longer holds in practice. This is particularly true with the emerging 5G mmWave technology, which needs to accurately beamform the wireless signal to mitigate severe attenuation loss at high frequencies [21], [22]. On the other hand, designing accurate and analytically tractable non-stationary channel models has remained an overarching challenge to the research community for decades [23], [24]. Furthermore, accurate channel estimation in rapidly-varying environments is often infeasible for applications requiring ultra-low latency. Responding to this challenge, we investigate the following question in the first half of the paper - is it possible to design a scheduling policy that minimizes the AoI irrespective of the channel dynamics and user-mobility patterns? The question is considerably general, as we do not make any assumption on either the channel statistics or the user-mobility, both of which may be dictated by an omniscient adversary in the worst case. The adversarial model is also useful for guaranteeing reliable communication in the presence of tactical jammers, where the interferers, in reality, may behave adversarially [25], [26]. A similar problem in the context of *stability* of wireless networks was considered in [27] under adversarial arrival and link transmission rates.

To formalize the problem, we first introduce an adversarial binary erasure channel model, which may be considered as an adversarial counterpart of the celebrated Binary Erasure Channel (BEC) model. However, unlike a similar adversarial model

This paper was presented in part at [1], [2], and [3]. The work by Rajarshi Bhattacharjee was completed while he was working with the first author as a project associate at the Indian Institute of Technology Madras.

considered in [28], we neither make any causality assumption nor impose any upper-bound on the fraction of erasures. Thus, our channel model is considerably more general. We propose a greedy scheduling policy that competitively minimizes the average and peak AoI in this model. See Table I for our main results for the adversarial framework.

In contrast with the non-stationary environment, minimizing the AoI in stationary stochastic environments has been extensively studied. In the paper [29], the authors study the average AoI minimization problem for static users associated with a single access point. The authors show that the greedy Max-Age (MA) policy is optimal for minimizing the average AoI in a symmetric static network. In the same paper, the authors also propose a 4-optimal Max-Weight scheduling policy (MW) for an arbitrary static network consisting of a single access point. The proposed MW policy has been reported to practically improve the information freshness in WiFi networks by two orders of magnitude [30]. The paper [31] extends the MW policy by taking into account additional throughput constraints. The paper [32] designs optimal stationary scheduling policies for minimizing the AoI in multi-hop networks with static users under general interference constraints. The paper [33] considers the problem of designing an AoI-optimal trajectory for a mobile agent to facilitate the information dissemination from a central station to a set of ground terminals. However, the problem of designing an AoI optimal scheduling policy for mobile users has not been considered before. In the second half of the paper we tackle this question and show that a natural multi-cell extension of the Max-Weight policy performs well under certain mobility assumptions.

Our contributions:

In this paper, we consider the AoI-optimal scheduling problem for mobile users in both adversarial and stochastic models. Our objective is to design simple scheduling policies that performs competitively in the adversarial environments and achieves near-optimality in the more benign stationary stochastic environments. The adversarial formulation of the problem is new and, to the best of our knowledge, has not been considered before. On the other hand, while the AoI-optimal scheduling problem for static users in the stochastic setting is now well-understood [29], [31], the multi-cell extension of the problem with mobile users is new. In particular, we make the following contributions in this paper:

- 1) Within the adversarial framework of Section II, we show that a greedy online scheduling policy is $2N^2$ -competitive for the Average AoI metric. Using Yao's minimax principle, we also establish a fundamental lower bound by showing that no online policy can have a competitive ratio smaller than $2N - 1$.
- 2) For the peak AoI metric, we show that the same greedy scheduling policy achieves a competitive ratio of $2N$ in the adversarial setting. Using Yao's minimax principle, we show that no online policy can have a competitive ratio better than $\Omega\left(\frac{N}{\ln(N)}\right)$. Table I summarizes our main results in the adversarial model.

- 3) In Section III, we consider the AoI-optimal multi-user scheduling problem for mobile users in a stationary environment. For the average-age metric, we design a 2-optimal scheduling policy for mobile users with i.i.d. uniform mobility. As a by-product of our analysis, we improve upon the best-known 4-approximation bound known for static users [29], [34].
- 4) To minimize the peak AoI metric in the stochastic setting, we show that the greedy policy is optimal for a single-cell static network. We give a short proof of this optimality result by exhibiting a solution to a countable-state average-cost MDP problem in the "closed-form", which might be of independent interest. This result supplements Theorem 5 of [29], which establishes the optimality of the greedy policy for the average AoI metric for *symmetric* static networks using stochastic dominance arguments. We further show that the greedy policy achieves the optimal large-deviation rate. Table II summarizes our main results in the stochastic model.

Discussion: Our motivation for studying the adversarial and the stochastic models together stems from the following two reasons - (1) to highlight the complementary strengths and weaknesses of each model, and (2) to show that, despite the differences, there exist simple scheduling policies that perform well in both settings. More specifically,

- 1) While the stochastic model makes strong assumptions on the environment, the adversarial framework makes virtually no assumptions. Hence, the adversarial model allows us to study the optimal scheduling problem in an arbitrary non-stationary environment, for which very few results exist in the literature. However, because of the difference in the generality of the assumptions, considerably stronger performance guarantees can be established in the stochastic setting compared to the adversarial environment. This brings us to the next point.
- 2) Given the difference in the performance guarantees, it is natural to ask whether there exist "universal" policies that perform near-optimally in the stochastic setting and simultaneously enjoy non-trivial performance guarantees in the adversarial environment. We answer the above question in the affirmative by exhibiting a simple greedy scheduling policy that achieves the above goal.

The rest of the paper is organized as follows. In Section II, we describe the adversarial model and prove upper and lower bounds on the competitive ratio achievable within this framework for minimizing AoI. In Section III, we describe the stochastic model and design near-optimal policies for minimizing the AoI in the stochastic setting. In Section IV, we compare the performance of the proposed scheduling policies via numerical simulations. Section V concludes the paper with some pointers to a few related open problems.

II. AOI MINIMIZATION IN ADVERSARIAL ENVIRONMENTS

We consider an optimal downlink scheduling problem in a wireless cellular network where N users roam around in a region having M Base Stations (BS). The environment, which is entirely specified by the channel states and the

TABLE I
SUMMARY OF THE RESULTS ON THE COMPETITIVE RATIOS IN THE ADVERSARIAL FRAMEWORK

| Metrics | Cost function | Mobility | Upper Bound | Achieving Policy | Lower Bound | Optimality gap |
|-------------|---|----------|-------------|------------------|----------------------------|----------------|
| Average AoI | $N^{-1} \sum_{t=1}^T \sum_{i=1}^N h_i(t)$ | Yes | $O(N^2)$ | CMA | $O(N)$ | $O(N)$ |
| Peak AoI | $\sum_{t=1}^T \max_{i=1}^N h_i(t)$ | Yes | $O(N)$ | CMA | $\Omega(\frac{N}{\ln(N)})$ | $O(\ln(N))$ |

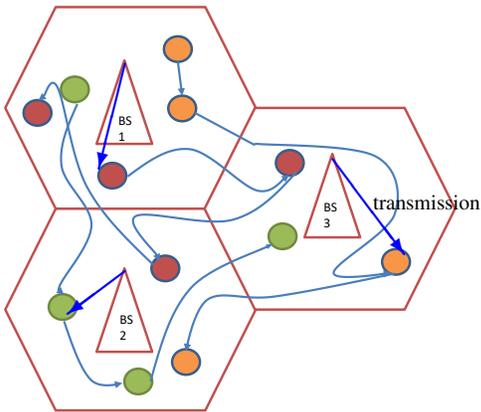


Fig. 1. Illustrating the movement of $N = 3$ users (each with a distinct color) in an area with $M = 3$ cells. The blue arrows indicate packet transmissions by the BS.

user mobility pattern, can evolve in an arbitrary fashion. Instead of trying to fit a complicated probabilistic model with multiple parameters [23], [24], we take a conservative view and model the environment using an adversarial framework. The adversarial viewpoint can be practically motivated by considering URLLC-type traffic, which requires extremely low latency with very high reliability [35]. Besides being analytically tractable, all achievability results in the adversarial model (Theorem 1) carry over to more benign stochastic environments. Moreover, as we will see in the sequel, policies having a good competitive ratio in the adversarial setting sometimes translate to optimal policies in the stochastic setting (Theorem 7).

A. Adversarial System Model

The main system components are described below.

a) Network model: The area covered by a Base Station (henceforth referred to as *BS*) is referred to as a *cell*. The cells are assumed to be spatially disjoint. Time is slotted, and at each time slot, a user can either stay in its current cell or move to any other $M - 1$ cells (the new cell need not be adjacent). The movement could be dictated by an omniscient adversary. Our mobility model is considerably general, as it does not make any assumptions (statistical or otherwise) on the speed or user movement patterns. See Figure 1 for a schematic.

b) Traffic Model: We consider a *saturated* traffic model where at the beginning of every slot, each of the M Base Stations receives a fresh update packet for each user from an external source (e.g., a high-speed optical core network). Since our objective is to maximize the freshness of information at the user-end, any stale packet at the BS buffers is replaced

by incoming fresh packets at each slot. Recall that a stale packet for a user at a slot is defined as any buffered packet(s) from the previous slot. Each BS can beamform and schedule a downlink packet transmission at each slot to only one user within its coverage area. Although to be specific, we consider downlink scheduling in this paper, an analogous problem can also be formulated for uplink transmissions, and all of our results apply to this case as well.

The saturated traffic model is standard in applications which involve continuous status updates [36], such as monitoring and surveillance with sensor networks [37], velocity and position updates for autonomous vehicles [38], command and control information exchange in mission-critical systems, disseminating stock-index updates and live game scores. The saturated traffic model has also been used earlier in connection with designing rate-optimal scheduling policies [39]–[41]. The advantage of this model is that it separates the arrival process from the scheduling policy, leading to a cleaner analysis in both adversarial and stochastic settings.

c) Channel states, Control, and Objective: The policy controls and optimizes user scheduling decisions by each BS. We consider a binary erasure channel model where the channel state for any user at any slot can be either **Good** or **Bad** (also referred to as **ON/OFF** in the literature [42], [43]). Here a *slot* is defined to be the duration of the entire packet transmission. In practice (e.g., in 5G systems), a packet is transmitted over several resource blocks (RBs) consisting of multiple OFDM symbols. Thus, if the aggregate of channel states remains poor over a packet transmission block so that the probability of a packet decoding error remains sufficiently large (e.g., an outage event), the channel for that slot is considered to be in **Bad** (or **OFF**) state. Else, the channel state for that slot is considered to be **Good** (or **ON**).

The schedulers are assumed to be oblivious to the current channel states (i.e., no CSIT). An online scheduling policy π first selects a user in each cell (if the cell contains at least one user), and then transmits the latest packet from the BS to the selected users over the wireless channel. If the corresponding channel is in **GOOD** state, the user decodes the packet successfully. Otherwise, the packet is lost. A lost packet is never retransmitted as the scheduler receives fresh packets at every slot. In the adversarial model, we posit that the channel states are dictated by an omniscient adversary [27]. In particular, we allow the situation where the adversary knows the scheduling policy in advance and chooses the channel realizations *after* the scheduling decisions have been made for a slot. On the other hand, the scheduling policy π is necessarily *online* and has no information about the channel states in the current or future slots. The set of all admissible scheduling

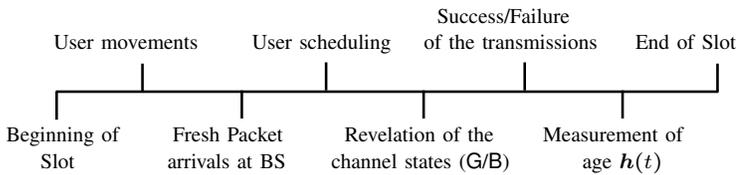


Fig. 2. Event sequence at each slot

policies is denoted by Π . See Figure 2 for the timeline of events taking place at every slot.

We are interested in competitively optimizing the information freshness for all users. Formally, our objective is to design a decentralized scheduling policy that minimizes some measure of the aggregate Age-of-Information of the users as defined next. For any slot $t \geq 1$, let $t_i(t) < t$ denote the last time prior to time t at which the i^{th} user successfully received a packet from some BS. The Age-of-Information (AoI) of the user at time t , denoted by $h_i(t)$, is defined as:

$$h_i(t) \equiv t - t_i(t).$$

In other words, $h_i(t)$ denotes the length of the time elapsed since the i^{th} user received its last update packet before time t . Thus, the process $\{h_i(t)\}_{t \geq 1}$ quantifies the *staleness* (or age) of the information available to the i^{th} user. Accordingly, we define an N -dimensional state-vector $\mathbf{h}(t)$, whose i^{th} component, $h_i(t)$, denotes the AoI of the i^{th} user at time t . Clearly, the plot of $h_i(t)$ vs. time t has a saw-tooth shape that increases linearly with unit-slope until a fresh packet is received. Upon the reception of a fresh packet, the AoI $h_i(t)$ instantaneously drops to 1. From that point onwards, $h_i(t)$ again increases linearly, repeating the saw-tooth pattern [29]. See Figure 3 for an illustration. In this paper, we consider optimizing the following two different aggregate AoI metrics:

a) **Average AoI:** The time-averaged cost corresponding to the average AoI for N users up to time T is defined as:

$$\text{AoI}_{\text{avg}}(T) = \frac{1}{NT} \sum_{t=1}^T \left(\sum_{i=1}^N h_i(t) \right). \quad (1)$$

b) **Peak AoI:** The instantaneous peak AoI at a slot is defined as the maximum age among all users. The time-averaged cost corresponding to the peak AoI for a time-horizon of length T is defined as:

$$\text{AoI}_{\text{peak}}(T) = \frac{1}{T} \sum_{t=1}^T \max_{i=1}^N h_i(t) \quad (2)$$

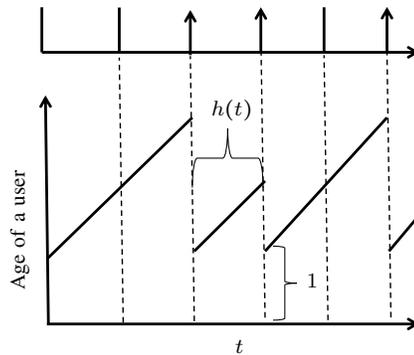


Fig. 3. Time-evolution of the Age-of-Information for a user. Solid arrows denote successful transmissions.

Performance Metric: As standard in the literature on online algorithms [44], [45], we gauge the performance of any online scheduling policy \mathcal{A} using the notion of *competitive ratio* (denoted by $\eta^{\mathcal{A}}$). Recall that the competitive ratio is defined as the worst-case ratio of the cost incurred by an online policy to that of an optimal *offline* policy OPT. The OPT policy is assumed to be equipped with hindsight knowledge. Specializing to our context, let σ be a sequence representing the channel states and the user's locations for the entire time-horizon. Then, the competitive ratio of a policy \mathcal{A} is defined as [45]:

$$\eta^{\mathcal{A}} = \sup_{\sigma} \left(\frac{\text{Cost of the policy } \mathcal{A} \text{ on } \sigma}{\text{Cost of OPT on } \sigma} \right). \quad (3)$$

In the above definition, the supremum is taken over *all* finite-length admissible sequences σ . Depending on the objective, the cost function in the definition (3) can be taken to be either Eqn. (1) or Eqn. (2). We emphasize that, while the online policy \mathcal{A} has only causal information (*i.e.*, knows only the subsequence σ_1^{t-1} at time t), the policy OPT is assumed to be equipped with non-causal knowledge of the entire sequence σ_1^T right at the beginning. Our objective is to design an online scheduling policy \mathcal{A} with a small competitive ratio so that it performs close to the OPT policy.

Discussion: In the sum-rate maximization problem, the objective is to maximize the total rate of successful packet transmissions to the users [4]. It is interesting to note that finding a scheduling policy with a small competitive ratio for the sum-throughput objective is too strong a requirement, as all deterministic policies suffer from unbounded competitive ratios. This can be understood from the following simple example. Assume that two stationary users are located in a single cell. If an online policy \mathcal{A} schedules the first user at any slot, the adversary can set the channel corresponding to the first user to **Bad** and the second user's channel to **Good** and vice versa. Hence, all transmissions by policy \mathcal{A} are unsuccessful. On the other hand, at any slot, the optimal policy schedules the user that has the **Good** channel state for that slot. Hence, the OPT policy achieves unit throughput, resulting in an unbounded competitive ratio. In the following, we show that, somewhat surprisingly, there exists a policy that achieves a finite competitive ratio for the AoI minimization problem.

B. Achievability

For our achievability results, we consider the following online scheduling policy:

Cellular Max-Age (CMA): At every slot, each BS j schedules a downlink packet transmission to the user with the highest age among all users in BS j 's cell at that slot (ties are broken in an arbitrary but fixed order).

Clearly, the CMA policy is decentralized as the schedulers at each BS need to know the state (AoI) of the users in their local cells only. In the following, we upper bound the competitive ratios of the CMA policy for the Average AoI objective (Eqn. (1)) and the Peak AoI objective (Eqn. (2)). Surprisingly, it turns out that the bounds are independent of the total number of Base Stations M . In the context of the CMA policy, we now state a few definitions which facilitate the achievability proofs.

Max-user: For any slot t , we define the (global) *Max-user* as the user having the highest age among all N users under the CMA policy (ties are broken similarly as in the CMA policy). Clearly, the identity of the Max-user changes with time. Observe that, by definition, the CMA policy continues to schedule the current Max-user *irrespective* of its location until it successfully receives a packet. In the following slot, a different user assumes the role of the Max-user, and the process continues.

Super-interval: Under the CMA policy, the time interval between two consecutive successful transmissions to the current Max-user is called a *super-interval*. Throughout a super-interval, the identity of the Max-user remains invariant. The Max-user corresponding to the i^{th} super-interval is denoted by M_i . Note that the super-intervals are contiguous and disjoint. Let T_i be the index of the time slot at which the i^{th} super-interval ends. Thus, $\Delta_i \equiv T_i - T_{i-1}$ denotes the length of the i^{th} super-interval. See Figure 4 for a schematic. Note that there could be more than one successful transmission within a super-interval to users other than the Max-user (by different Base Stations). For notational consistency, we define $T_j \equiv 0$, and $\Delta_j \equiv 0, \forall j \leq 0$. The above definitions and observations lead to the following key result:

Proposition 1: At the k^{th} slot of the i^{th} super-interval, the age of the Max-user M_i under the CMA policy is upper bounded by $k + \sum_{j=1}^{N-1} \Delta_{i-j}$.

Proof: We claim that the Max-user M_i corresponding to the i^{th} super-interval must have had a successful transmission within the last $N - 1$ super-intervals. If not, since there are a total of N users, 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 j^{th} user would have had less age than M_i (which was not scheduled at all in the last $N - 1$ super-intervals) when the j^{th} user became the Max-user for the second time. Thus, at the beginning of the i^{th} super-interval, the age of the new

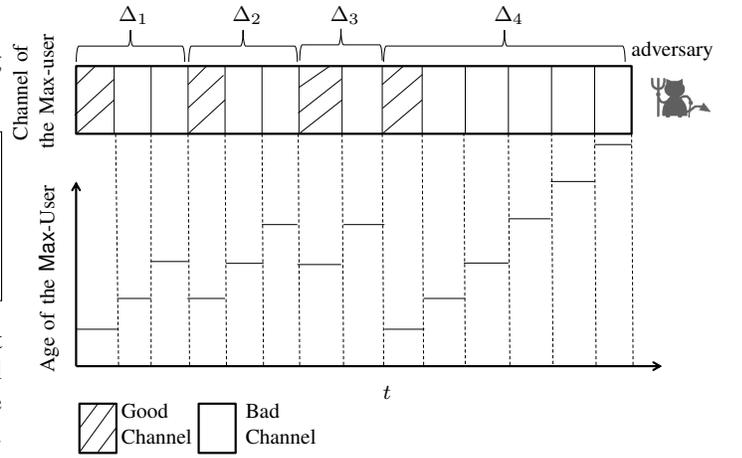


Fig. 4. Illustrating Super-intervals under the CMA policy

Max-user is upper bounded by $\sum_{j=1}^{N-1} \Delta_{i-j}$. The proposition follows from this fact. ■

The previous proposition leads to the following achievability result for the CMA policy:

Theorem 1 (Upper bounds): The competitive ratios of the CMA policy for the Average AoI and Peak AoI metrics can be upper bounded as follows:

$$\eta_{\text{avg}}^{\text{CMA}} \leq 2N^2, \text{ and } \eta_{\text{peak}}^{\text{CMA}} \leq 2N.$$

Discussion on the proof technique: We prove the bounds for the Average AoI and the Peak AoI metrics separately. The achievability proofs proceed by upper bounding the cost incurred by the CMA policy and lower bounding the cost incurred by the OPT policy in each super-interval generated by the CMA policy. Note that the quantities, such as the super-intervals $\{\Delta_i\}_{i \geq 1}$, the number of super-intervals K , and the Max-User M_i appearing in the proofs below are defined with reference to the CMA policy only. See Figure 4 for an illustration. Hence, these quantities are determined exclusively by the input channel state and user location sequence σ , which remains the same for all policies, including OPT, while computing the competitive ratio via Eqn. (3). In other words, we first run the CMA policy on the input sequence σ , determine the super-intervals $\{\Delta_i\}_{i=1}^K$ and the corresponding Max-users $\{M_i\}_{i=1}^K$, and then use these quantities to analyze the performance of OPT. We emphasize that if a user j corresponds to the Max-user in the i^{th} super-interval, then the user j need not have the highest age among all users at the same slot under the operation of a different scheduling policy (e.g., OPT).

Proof:

a) *Average-AoI:* Note that, at any slot of the i^{th} super-interval, the age of every user is upper bounded by that of the Max-user M_i . Hence, using Proposition 1, the total cost incurred by the CMA policy during the i^{th} super-interval may

be upper bounded as:

$$\begin{aligned}
C_i^{\text{CMA}} &\equiv \sum_{t \in i^{\text{th}} \text{ super-interval}} \sum_{t=1}^N h_i(t) \\
&\leq \sum_{k=1}^{\Delta_i} N \left(k + \left(\sum_{j=1}^{N-1} \Delta_{i-j} \right) \right) \\
&= N \left(\frac{\Delta_i(\Delta_i + 1)}{2} + \sum_{j=1}^{N-1} \Delta_i \Delta_{i-j} \right) \\
&\leq \frac{N}{2} \left(N\Delta_i^2 + \Delta_i + \sum_{j=1}^{N-1} \Delta_{i-j}^2 \right),
\end{aligned}$$

where in the last step, we have used the AM-GM inequality to obtain $\Delta_i \Delta_{i-j} \leq \frac{1}{2}(\Delta_i^2 + \Delta_{i-j}^2)$, $1 \leq j \leq N-1$.

Let K be the total number of super-intervals in the entire time-horizon of length T . The total cost incurred by the CMA policy for the entire time horizon can be upper bounded as:

$$\begin{aligned}
\text{Cost}^{\text{CMA}}(T) &= \sum_{i=1}^K C_i^{\text{CMA}} \\
&\leq \frac{N}{2} \sum_{i=1}^K \left(N\Delta_i^2 + \Delta_i + \sum_{j=1}^{N-1} \Delta_{i-j}^2 \right) \\
&\leq \frac{N}{2} \sum_{i=1}^K \left(2N\Delta_i^2 + \Delta_i \right). \tag{4}
\end{aligned}$$

On the other hand, observe that the user M_i must experience Bad channels throughout the i^{th} super-interval generated by the CMA policy. This is true, as otherwise, the user M_i would have successfully received a packet under the CMA policy, which, by design, always schedules a packet transmission to the Max-user irrespective of its location. Hence, the cost incurred by the OPT policy during the i^{th} super-interval may be lower bounded by the cost of the user M_i as follows:

$$\begin{aligned}
C_i^{\text{OPT}} &\geq \underbrace{(N-1) \sum_{k=1}^{\Delta_i} 1}_{\text{lower-bound to the cost incurred by other users}} + \\
&\quad \underbrace{\sum_{k=1}^{\Delta_i} (1+k)}_{\text{lower-bound to the cost incurred by } M_i} \\
&\geq \frac{1}{2}\Delta_i^2 + N\Delta_i. \tag{5}
\end{aligned}$$

In the above, the lower-bound to the cost incurred by M_i is obtained using the fact that, due to successive Bad channels, its age increases *linearly* under any policy (including OPT) throughout the i^{th} super interval, starting from the minimum age of at least one. Finally, the total cost incurred during the entire horizon of length T is obtained by summing up the cost incurred in the constituent super-intervals. Hence, from Eqns. (4) and (5), the competitive ratio η^{CMA} of the CMA policy

may be upper bounded as:

$$\begin{aligned}
\eta_{\text{avg}}^{\text{CMA}} &= \frac{\frac{1}{NT} \sum_{i=1}^K C_i^{\text{CMA}}}{\frac{1}{NT} \sum_{i=1}^K C_i^{\text{OPT}}} \\
&\stackrel{(a)}{\leq} \frac{\frac{N}{2} \sum_{i=1}^K \left(2N\Delta_i^2 + \Delta_i \right)}{\sum_{i=1}^K \left(\frac{1}{2}\Delta_i^2 + N\Delta_i \right)} \\
&\leq 2N^2.
\end{aligned}$$

b) *Peak AoI*: The proof proceeds essentially in the same way as the Average Age case. Using proposition 1, the total cost C_i^{CMA} incurred by the CMA policy during the i^{th} super-interval may be upper bounded as:

$$\begin{aligned}
C_i^{\text{CMA}} &= \sum_{t \in i^{\text{th}} \text{ super interval}} \max_{i=1}^N h_i(t) \\
&\leq \sum_{k=1}^{\Delta_i} \left(k + \sum_{j=1}^{N-1} \Delta_{i-j} \right) = \frac{1}{2}(\Delta_i^2 + \Delta_i) + \sum_{j=1}^{N-1} \Delta_i \Delta_{i-j} \\
&\stackrel{(a)}{\leq} \frac{1}{2}(\Delta_i^2 + \Delta_i) + \frac{1}{2} \sum_{j=1}^{N-1} (\Delta_i^2 + \Delta_{i-j}^2) \\
&= \frac{N}{2}\Delta_i^2 + \frac{1}{2}\Delta_i + \frac{1}{2} \sum_{j=1}^{N-1} \Delta_{i-j}^2.
\end{aligned}$$

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 K super-intervals in the time-horizon of length T , the total cost incurred by the CMA policy over the entire time horizon may be upper bounded as:

$$\text{Cost}^{\text{CMA}}(T) = \sum_{i=1}^K C_i^{\text{CMA}} \leq \frac{1}{2} \sum_{i=1}^K \left(2N\Delta_i^2 + \Delta_i \right). \tag{6}$$

On the other hand, the cost incurred by the OPT policy during the i^{th} super-interval is trivially lower bounded by the cost of the user M_i during the i^{th} super-interval. Note that, as in the previous proof, the user M_i experiences successive Bad channels throughout the i^{th} super-interval. Hence,

$$C_i^{\text{OPT}} \geq \sum_{k=1}^{\Delta_i} (1+k) = \frac{1}{2}\Delta_i^2 + \frac{3}{2}\Delta_i. \tag{7}$$

Finally, the cost of the entire horizon of length T may be obtained by summing up the cost incurred in each super-intervals. Noting that $\Delta_0 = 0$, using Eqns. (6) and (7), the competitive ratio η^{CMA} of the CMA policy may be upper bounded as:

$$\eta_{\text{peak}}^{\text{CMA}} = \frac{\frac{1}{T} \sum_{i=1}^K C_i^{\text{CMA}}}{\frac{1}{T} \sum_{i=1}^K C_i^{\text{OPT}}} \stackrel{(a)}{\leq} \frac{\frac{1}{2} \sum_{i=1}^K \left(2N\Delta_i^2 + \Delta_i \right)}{\sum_{i=1}^K \left(\frac{1}{2}\Delta_i^2 + \frac{3}{2}\Delta_i \right)} \leq 2N.$$

Tightness of the Bounds

Next, we show that our analysis of the CMA policy is tight and the bounds given in Theorem 1 cannot be improved by more than a constant factor. ■

Theorem 2: In the adversarial setting described above, we have

$$\eta_{\text{avg}}^{\text{CMA}} \geq N^2 \text{ and } \eta_{\text{peak}}^{\text{CMA}} \geq 2N - 1.$$

We defer the proof of the above theorem to Appendix VI-A. The proof proceeds by constructing a suitable channel state sequence for N static users in a single cell running the CMA policy.

C. Minimax Lower Bounds for the Competitive Ratios

We now use Yao's minimax principle to obtain lower bounds to the achievable competitive ratio for minimizing the Average AoI and the Peak AoI metrics. These lower bounds are universal and hold for any online policy. For this, we first recall Yao's minimax principle, which is an immediate consequence of Von Neumann's Minimax Theorem in game theory:

Theorem 3 (Yao's Minimax principle [45]): The competitive ratio of the best randomized online algorithm against an oblivious adversary is equal to the competitive ratio of the best deterministic online algorithm under a worst-case input distribution.

From the above principle, it is clear that a lower bound to the competitive ratio under any deterministic online algorithm for a given input distribution \mathbf{p} yields a lower bound to the competitive ratio in the adversarial setting, *i.e.*,

$$\eta \geq \frac{\mathbb{E}_{\sigma \sim \mathbf{p}}(\text{Cost of the Best Deterministic Online Policy})}{\mathbb{E}_{\sigma \sim \mathbf{p}}(\text{Cost of OPT})}. \quad (8)$$

Towards this end, we construct a suitable channel-state distribution \mathbf{p} for N static users located at a single BS. Then, we upper bound the expected cost incurred by OPT and lower bound the expected cost incurred by the CMA policy under the distribution \mathbf{p} to lower bound η .

Theorem 4 (Minimax Lower Bounds): The competitive ratio of any online policy $\pi \in \Pi$ with N static users located at a single cell is lower bounded as follows:

$$\eta_{\text{avg}}^{\pi} \geq \frac{N}{2} + \frac{1}{2N}, \quad \eta_{\text{peak}}^{\pi} \geq \Omega(N/\ln(N)).$$

Proof: As stated above, to apply Yao's minimax principle (Eqn. (8)), we need to (a) lower bound the expected cost incurred by the online policy π , and (b) upper bound the expected cost incurred by the optimal offline policy for some suitably chosen channel state distribution \mathbf{p} . Selecting a channel state distribution \mathbf{p} , which simultaneously yields a tight lower bound and leads to a tractable analysis, is a non-trivial task. Towards this, we consider the following channel state distribution for N static users located at a single cell.

Distribution of Channel States \mathbf{p} : At every 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.

The rationale behind the above choice of the channel state distribution will become clear when we evaluate OPT's expected cost. In general, the optimal offline policy's cost is obtained by solving a Dynamic Program, which is challenging to analyze. However, with the selected channel state distribution \mathbf{p} , only one user's channel is in Good state at any time. Hence, in this case, the OPT policy schedules the user that has Good channel state for that slot. This leads to a tractable analysis of OPT's expected cost for both the Average AoI and Peak AoI objectives.

Case I: Average AoI

1) **Computation of OPT's cost:** Let the random variable $C_i(T)$ denote the total cost incurred by the i^{th} user up to time T , *i.e.*,

$$C_i(T) = \sum_{t=1}^T h_i(t).$$

Hence, the limiting time-averaged expected cost incurred by the OPT policy may be expressed as:

$$\bar{c}(\text{OPT}) \equiv \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{i=1}^N \mathbb{E}(C_i(T)) = \sum_{i=1}^N \lim_{T \rightarrow \infty} \frac{\mathbb{E}(C_i(T))}{T}, \quad (9)$$

In the following, we will show that the above limits exist for the channel state distribution \mathbf{p} . We now use the Renewal Reward Theorem [46] to evaluate the RHS of Eqn. (9). As stated before, since only one channel is in Good state at a slot, the optimal policy OPT simply schedules the user having Good channel. It can be verified that, under the OPT policy, for each user i , the sequence of age random variables $\{h_i(t)\}_{t \geq 1}$ constitute a renewal process. Clearly, the time instants of scheduling the i^{th} user constitutes renewal instants. A generic renewal interval of length τ for the i^{th} user consists of two parts - a sequence of Good channels of length τ_{G} , followed by a sequence of Bad channels of length τ_{B} . Hence, the AoI cost $c_i(\tau)$ incurred by the user i in any generic renewal cycle may be written as the sum of the costs incurred in two parts:

$$\begin{aligned} c_i(\tau) &= c_i(\tau_{\text{G}}) + c_i(\tau_{\text{B}}) \\ &= \sum_{t=1}^{\tau_{\text{G}}} 1 + \sum_{t=1}^{\tau_{\text{B}}} (1+t) \\ &= \tau_{\text{G}} + \frac{3}{2}\tau_{\text{B}} + \frac{1}{2}\tau_{\text{B}}^2. \end{aligned}$$

Let $q \equiv \frac{1}{N}$ be the probability that the channel is Good for the i^{th} user at any slot. Hence, from our construction, the random variables τ_{G} and τ_{B} follows a Geometric distribution with the following p.m.f.

$$\begin{aligned} \mathbb{P}(\tau_{\text{G}} = k) &= q^{k-1}(1-q), \quad k \geq 1. \\ \mathbb{P}(\tau_{\text{B}} = k) &= q(1-q)^{k-1}, \quad k \geq 1. \end{aligned}$$

Thus, the expected cost incurred by the i^{th} user at any renewal cycle is given by:

$$\mathbb{E}(c_i(\tau)) = \frac{1}{1-q} + \frac{3}{2q} + \frac{2-q}{2q^2} = \frac{1}{q^2(1-q)}. \quad (10)$$

Moreover, the expected length of any renewal cycle can be computed to be:

$$\mathbb{E}(\tau) = \mathbb{E}(\tau_G) + \mathbb{E}(\tau_B) = \frac{1}{q(1-q)}. \quad (11)$$

Using Renewal Reward Theorem [46], we have

$$\lim_{T \rightarrow \infty} \frac{\mathbb{E}(C_i(T))}{T} = \frac{\mathbb{E}(c_i(\tau))}{\mathbb{E}(\tau)} = \frac{1}{q} = N, \quad \forall i.$$

Hence, from Eqn. (9), we conclude that the limiting time-averaged total expected cost incurred by OPT is given by

$$\bar{C}(\text{OPT}) = N^2. \quad (12)$$

2) *Lower bound to the cost of policy π* : In order to lower bound the expected cost incurred by any online policy π under the distribution \mathbf{p} , we appeal to a special case of Theorem 5, stated later in Section III-A. Theorem 5 gives a lower bound to the average cost incurred by any scheduling policy in a stochastic setting when the channels are modeled as binary erasure channels (BEC) with fixed probabilities of success. Since the proof of Theorem 5 does not require the channels to be mutually independent, it is applicable to the channel state distribution \mathbf{p} as well. By plugging in $p_i = \frac{1}{N}$, $\forall i$, and $M = 1$ in Theorem 5, we conclude that under the channel state distribution \mathbf{p} , the time-averaged expected cost for any online scheduling policy π is lower bounded as:

$$\bar{C}(\pi) = \limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{i=1}^N \mathbb{E}(C_i(T)) \geq \frac{N^3 + N}{2}. \quad (13)$$

Finally, using Yao's minimax principle in conjunction with Eqns. (12) and (13), we conclude that the competitive ratio η_{avg}^π of any online scheduling policy π is lower bounded as:

$$\eta_{\text{avg}}^\pi \geq \sup_T \frac{C_T(\pi)}{C_T(\text{OPT})} \geq \limsup_{T \rightarrow \infty} \frac{C_T(\pi)/T}{C_T(\text{OPT})/T} \geq \frac{N}{2} + \frac{1}{2N}. \quad (14)$$

We point out that the lower bound in Eqn. (14) can be further improved in the case of $N = 2$ users using a more careful analysis. The following result shows that the lower bound for average AoI for $N = 2$ users may be improved to 1.5 from 1.25 as given by Theorem 4.

Proposition 2: For the case of $N = 2$ users, we have the following improved bound: $\eta_{\text{avg}}^\pi \geq 1.5$.

Please refer to Appendix VI-B for the proof. Proposition 2 suggests that the lower bound on the competitive ratio for the average age metric is likely to be loose and may be improved upon further using a more refined analysis.

Next, we consider the Peak-AoI objective (2) and derive a minimax lower bound for this cost metric.

Case II- Peak AoI

3) *Upper bound to OPT's cost*: We use the same channel state distribution \mathbf{p} as before. Recall that, under the distribution \mathbf{p} , at any slot t :

$$\begin{aligned} \mathbb{P}(\text{user } i\text{'s channel is Good at slot } t) &= 1/N, \\ \mathbb{P}(\text{user } i\text{'s channel is Bad at slot } t) &= 1 - 1/N, \end{aligned}$$

independent of everything else. The OPT policy, with non-causal channel state information, schedules the user having a **Good** channel at every slot. Thus, the limiting distribution of the age of any user is Geometric ($\frac{1}{N}$), i.e.,

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

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:

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

Let the random variable H_{max} denote limiting peak-age of the users. For any $-\log(1 - 1/N) > \lambda > 0$, we have:

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

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

$$\mathbb{E}(H_{\text{max}}) \leq 1 - \frac{1}{\lambda} \log(1 - e^\lambda(1 - 1/N)). \quad (15)$$

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

$$e^x \leq \frac{1}{1-x}, \quad \forall x < 1. \quad (16)$$

As a result,

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

Next, for upper bounding the RHS of Eqn. (15), we start with the simple analytical fact that for any $0 < \alpha < 1$,

$$\inf_{0 < x < 1} \frac{1 - (1-x)e^{\alpha x}}{x} = 1 - \alpha. \quad (17)$$

This result can be verified by using Eqn. (16) to conclude that for any $0 < x < 1$, we have

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

where the infimum is achieved when $x \rightarrow 0^+$. Substituting $x = \frac{1}{N}$ in the inequality (17), we have

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

Hence, using Eqn. (15), 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:

$$\mathbb{E}(H_{\max}) \leq N \ln N + o(N \ln N). \quad (18)$$

4) Lower Bound to the expected cost of any online policy π :

To establish a lower bound to the expected cost of any online policy π , we use Theorem 7, established in Section III-B. Theorem 7 gives the minimum time-averaged peak-AoI cost in the stationary environment when all channels can be modeled as BEC. As in the average AoI case, it can be verified that the theorem continues to hold under the channel state distribution \mathbf{p} . Taking $p_i = \frac{1}{N}$, $\forall i$, and $M = 1$ in Theorem 7, we conclude that under the distribution \mathbf{p} , the time-averaged expected cost for any online scheduling policy π is lower bounded as:

$$\liminf_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \text{Cost}^\pi(T) = \liminf_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \mathbb{E}(\max_i h_i^\pi(t)) \geq N^2. \quad (19)$$

Combining Eqns. (18) and (19) with Eqn. (8) and using Cesàro's summation formula, we have for any online policy $\pi \in \Pi$:

$$\begin{aligned} \eta_{\text{peak}}^\pi &\geq \sup_{T>0} \frac{\mathbb{E} \text{Cost}^\pi(T)}{\mathbb{E} \text{Cost}^{\text{OPT}}(T)} \\ &\geq \limsup_{T \rightarrow \infty} \frac{\mathbb{E} \text{Cost}^\pi(T)/T}{\mathbb{E} \text{Cost}^{\text{OPT}}(T)/T} \\ &\geq \Omega\left(\frac{N}{\ln N}\right). \end{aligned}$$

Discussion and open problems: Combining the above results, we see that for the average AoI metric, there is a factor of $O(N)$ gap between the competitive ratio of the CMA policy and the corresponding lower bound. Theorem 2 shows that the upper bound to the competitive ratio of the CMA policy, given by Theorem 1, is tight and cannot be improved further (up to a constant factor). Hence, either there exists a different online policy with a smaller competitive ratio, or the current lower bound for average AoI can be improved with a more involved analysis. Reducing the current optimality gap for the average AoI metric is an interesting open problem. ■

III. AOI MINIMIZATION IN STOCHASTIC ENVIRONMENTS

In this section, we study the problem of AoI-optimal multi-user scheduling in a cellular wireless network when the channel and user mobility can be modeled as a stationary stochastic process. In the following, we highlight the major differences between the adversarial model in Section II and the stochastic model in this section. Our main results for the stochastic model are summarized in Table II.

Stochastic System Model

a) *Channel model:* As in the adversarial model, we consider a cellular system where N users roam in an area having M Base Stations. However, contrary to the adversarial model, the wireless link between the i^{th} user and the associated BS is modeled using a stationary binary erasure channel (BEC) with the probability of successful reception of a transmitted packet being p_i , $0 < p_i \leq 1$, $\forall i$. Hence, when the associated BS schedules a downlink packet transmission to the i^{th} user, the packet is either successfully received with probability p_i (if the channel is in Good state) or is lost otherwise (if the channel is in Bad state). Due to the power control mechanisms employed by Base Stations, the success probabilities (*i.e.*, the parameter p_i 's) may vary among the users [47], [48]. The channels are i.i.d. with respect to the time but need not be independent across the users.

b) *Mobility model:* Contrary to the adversarial setting, where we allow arbitrary mobility patterns, in the stochastic setting, we model the user mobility using a stationary ergodic process. Formally, let the random variable $C_i(t) \in \{1, 2, \dots, M\}$ denote the index of the cell to which the i^{th} user is associated with at time t . We assume that the stochastic process $\{C_i(t)\}_{t \geq 1}$ is a stationary ergodic process such that $\mathbb{P}(C_i(t) = j) = \psi_{ij}$, $\forall i \in [N]$, $j \in [M]$, $t \geq 1$. The probability measure ψ denotes the time-invariant occupancy distribution of the cells by the users. The mobility of the users may be correlated or independent of each other. Many different stochastic mobility models proposed in the literature fall under the above general scheme, including the i.i.d. mobility model, random walk model, and the random waypoint model [49]–[52].

c) *Packet arrival model and Policy Space:* We use the same saturated traffic model for packet arrivals as the adversarial setup. The policy space Π is also identical to the adversarial framework. In particular, we study online scheduling policies that know the channel parameters $\{p_i\}_{i=1}^N$ but do not know the current or future realizations of random channel states.

d) *Performance metrics:* For a given policy $\pi \in \Pi$, its long-term Average AoI metric is defined as:

$$\text{AoI}_{\text{avg}}^\pi = \limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \frac{1}{N} \left(\sum_{i=1}^N \mathbb{E}^\pi(h_i(t)) \right). \quad (20)$$

Similarly, the long-term Peak AoI metric achieved by a policy $\pi \in \Pi$ is defined as:

$$\text{AoI}_{\text{peak}}^\pi = \limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \mathbb{E}(\max_i h_i(t)). \quad (21)$$

Our goal is to design scheduling policies that minimize the long-term expected Average AoI and the long-term expected Peak AoI.

Discussion: Our stochastic model is simplistic in the sense that it does not account for time-varying path losses. A more realistic model would take into account the time-varying nature of successful packet reception probabilities as a function of the users' current distances from the associated BS. However, such a model is challenging to analyze and not expected to shed much additional insight into the design

TABLE II
SUMMARY OF THE RESULTS FOR THE STOCHASTIC SETTING

| Metrics | Cost function | Mobility | Upper Bound | Attaining Policy | Lower Bound | Approx. |
|-------------|--|----------|--------------------------------|------------------|--|---------|
| Average AoI | $\limsup_{T \rightarrow \infty} \frac{1}{NT} \sum_{t=1}^T \left(\sum_{i=1}^N \mathbb{E}^\pi (h_i(t)) \right)$ | Yes | $\frac{N}{Mp(1-(1-M^{-1})^N)}$ | MMW | $\frac{1}{2Ng(\psi)} \left(\sum_{i=1}^N \sqrt{\frac{1}{p_i}} \right)^2 + \frac{1}{2}$ | 2 |
| Peak AoI | $\limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \mathbb{E}(\max_i h_i(t))$ | No | $\sum_{i=1}^N \frac{1}{p_i}$ | CMA | $\sum_{i=1}^N \frac{1}{p_i}$ | 1 |

of good scheduling policies. Our choice of the above simple model is motivated by the following two reasons:

- 1) In practice, dynamic power control mechanisms employed in cellular networks tightly regulate the transmission power from the BS to each user so that the received SNR levels remain almost constant [48], [53], [54]. This keeps the probability of successful packet receptions for each user roughly invariant irrespective of their locations and thus, mitigates the cell-edge effects.
- 2) Although our stochastic model is simple, it sheds valuable insights into the design of good scheduling policies even in non-stationary environments. In Section IV, we show that the effect of time-varying channel parameters can be heuristically incorporated into the proposed policies, which lead to improved performance.

Proof Techniques: The problem of finding an optimal policy for minimizing the AoI in the stochastic setting (Eqns. (20) and (21)) reduces to solving an infinite state, unbounded cost Markov Decision Process (MDP) problem with an average cost objective. Such problems are notoriously difficult to tackle, and at present, there is no comprehensive theory (see Section 4.6 of [55]). Moreover, standard numerical approximation schemes for infinite-state MDPs do not typically provide explicit performance guarantees [56]. As a result, solving the average cost MDP problems with infinite states require problem-specific techniques.

For the problem of minimizing the peak AoI in Section III-B, we directly establish the optimality of a greedy index policy in Theorem 7 by making use of a clever guess for the differential value function in Bellman's equation. However, a similar direct method could not be found for the problem of minimizing the average AoI metric. Hence, for this problem, we resort to an indirect approach. In Theorem 5, we establish a lower bound to the average AoI incurred by *any* admissible scheduling policy. Following this result, in Theorem 6, we establish that the proposed MMW policy achieves an average AoI, which is at most twice as large as the above lower bound under certain modelling assumptions.

A. Minimizing the Average AoI in the Stochastic Setting

1) *Lower bound:* The following theorem gives a universal lower bound to the average AoI metric attained by any admissible policy.

Theorem 5 (Converse): In the stochastic setup, the optimal average AoI is lower bounded as:

$$\text{AoI}_{\text{avg}}^* \geq \frac{1}{2Ng(\psi)} \left(\sum_{i=1}^N \sqrt{\frac{1}{p_i}} \right)^2 + \frac{1}{2}, \quad (22)$$

where the function $g(\psi)$ denotes the expected number of cells having at least one user under the stationary occupancy distribution ψ .

Since $g(\psi) \leq \min\{M, N\}$, we also have the following weakening of the above bound, which is agnostic of the user mobility statistics:

$$\text{AoI}_{\text{avg}}^* \geq \frac{1}{2N \min\{M, N\}} \left(\sum_{i=1}^N \sqrt{\frac{1}{p_i}} \right)^2 + \frac{1}{2}.$$

Proof: We use a sample-path-based argument to obtain an almost sure lower bound to the average AoI. We then use Fatou's lemma [57] to convert the almost sure bound to a bound in expectation.

Consider a sample path under the action of any arbitrary admissible scheduling policy $\pi \in \Pi$. Let the random variable $N_i(T)$ denote the number of packets received by the i^{th} user up to time T . Also, let the random variable T_{ij} denote the time interval between receiving the $(j-1)^{\text{th}}$ packet and the j^{th} packet, and the random variable D_i denote the time interval between receiving the last (*i.e.*, $N_i(T)^{\text{th}}$) packet and the time-horizon T for the i^{th} user. Hence, we can write

$$T = \sum_{j=1}^{N_i(T)} T_{ij} + D_i. \quad (23)$$

Since the AoI of any user increases by one at each slot until a new packet is received and then it drops to one again (Figure 3), the average AoI up to time T may be lower bounded as:

$$\begin{aligned} \overline{\text{AoI}}_T &\equiv \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T h_i(t) \\ &= \frac{1}{NT} \sum_{i=1}^N \left(\sum_{j=1}^{N_i(T)} \frac{1}{2} T_{ij} (T_{ij} + 1) + \frac{1}{2} D_i (D_i + 1) \right) \\ &\stackrel{(a)}{=} \frac{1}{2NT} \sum_{i=1}^N \left(N_i(T) \left(\frac{1}{N_i(T)} \sum_{j=1}^{N_i(T)} T_{ij}^2 \right) + D_i^2 \right) + \frac{1}{2} \\ &\stackrel{(b)}{\geq} \frac{1}{2NT} \sum_{i=1}^N \left(N_i(T) \bar{T}_i^2 + D_i^2 \right) + \frac{1}{2}, \end{aligned} \quad (24)$$

where in Eqn. (a) we have used (23), and in Eqn. (b) we have defined $\bar{T}_i \equiv \frac{1}{N_i(T)} \sum_{j=1}^{N_i(T)} T_{ij}$ and used Jensen's inequality.

Rearranging Eqn. (23), we can express the random variable \bar{T}_i as:

$$\bar{T}_i = \frac{T - D_i}{N_i(T)}.$$

With this substitution, the term within the bracket in Equation (24) simplifies to

$$N_i(T)\bar{T}_i^2 + D_i^2 = \frac{(T - D_i)^2}{N_i(T)} + D_i^2 \geq \frac{T^2}{N_i(T) + 1}, \quad (25)$$

where the last inequality is obtained by minimizing the middle expression by viewing it as a quadratic in the variable D_i . Hence, from Eqns. (24) and (25), we obtain the following lower bound to the average AoI under the action of any admissible scheduling policy:

$$\overline{\text{AoI}_T} \geq \frac{T}{2N} \sum_{i=1}^N \frac{1}{N_i(T) + 1} + \frac{1}{2}. \quad (26)$$

Next, we incorporate the scheduling constraints to control the RHS of inequality (26). Let the random variable $A_i(T)$ denote the total number of transmission attempts made to the i^{th} user by all Base Stations up to time T . Also, let the random variable $g_j(T)$ denote the fraction of time that BS_j contained at least one user in its coverage area. Since a BS can attempt a downlink transmission only when there is at least one user in its coverage area, the total number of transmission attempts to all users by the Base Stations can be upper bounded by the following global balance condition:

$$\sum_{i=1}^N A_i(T) \leq T \sum_{j=1}^M g_j(T) \equiv Tg(T), \quad (27)$$

where $g(T) \equiv \sum_j g_j(T)$. Using inequality (27), we can lower bound the RHS of inequality (26) as:

$$\overline{\text{AoI}_T} \geq \frac{1}{2Ng(T)} \left(\sum_{i=1}^N A_i(T) \right) \left(\sum_{i=1}^N \frac{1}{N_i(T) + 1} \right) + \frac{1}{2}. \quad (28)$$

An application of the Cauchy-Schwartz inequality on the RHS of the above inequality yields:

$$\overline{\text{AoI}_T} \geq \frac{1}{2Ng(T)} \left(\sum_{i=1}^N \sqrt{\frac{A_i(T)}{N_i(T) + 1}} \right)^2 + \frac{1}{2}. \quad (29)$$

Note that, the i^{th} user successfully received $N_i(T)$ packets out of a total of $A_i(T)$ packet transmission attempts made by the Base Stations over the wireless erasure channels with success probability p_i . Without any loss of generality, we may fix our attention on those scheduling policies only for which $\lim_{T \rightarrow \infty} A_i(T) = \infty, \forall i$ almost surely. Otherwise, at least one of the users will receive only a finite number of packets, resulting in infinite average AoI. Hence, using the Strong law of large numbers [57], we obtain:

$$\lim_{T \rightarrow \infty} \frac{N_i(T)}{A_i(T)} = p_i, \quad \forall i \text{ w.p. } 1. \quad (30)$$

Moreover, using the ergodicity property of the user mobility, we conclude that almost surely:

$$\lim_{T \rightarrow \infty} g_j(T) = \mathbb{P}_{\psi}(\text{BS}_j \text{ contains at least one user}),$$

where we recall that ψ denotes the stationary cell occupancy distribution. Thus, we have almost surely

$$\begin{aligned} \lim_{T \rightarrow \infty} g(T) &= \lim_{T \rightarrow \infty} \sum_j g_j(T) \\ &= \sum_{j=1}^M \mathbb{P}_{\psi}(\text{BS}_j \text{ contains at least one user}) \\ &\equiv g(\psi), \end{aligned} \quad (31)$$

where the function $g(\psi)$ denotes the expected number of non-empty cells under the stationary occupancy distribution ψ . Hence, combining equations (30) and (31) together with the lower bound in equation (29), we have almost surely:

$$\liminf_{T \rightarrow \infty} \overline{\text{AoI}_T} \geq \frac{1}{2Ng(\psi)} \left(\sum_i \sqrt{\frac{1}{p_i}} \right)^2 + \frac{1}{2}. \quad (32)$$

Finally,

$$\begin{aligned} \text{AoI}_{\text{avg}}^* &\geq \liminf_{T \rightarrow \infty} \mathbb{E}(\text{AoI}_T) \\ &\stackrel{(a)}{\geq} \mathbb{E}(\liminf_{T \rightarrow \infty} \text{AoI}_T) \\ &\geq \frac{1}{2Ng(\psi)} \left(\sum_i \sqrt{\frac{1}{p_i}} \right)^2 + \frac{1}{2}, \end{aligned}$$

where the inequality (a) follows from Fatou's lemma. This concludes the proof of Theorem 5. Note that the proof continues to hold even when the mobility of the users is not independent across the users. ■

Discussion: Theorem 5 suggests that the user mobility statistics affects the lower bound only through the stationary cell-occupancy distribution ψ . Hence, given the stationary distribution ψ , the lower bound (22) is agnostic of the specifics of the mobility model. A similar result was obtained earlier in connection with the capacity region of wireless networks (see [58], Corollary 5, p. 88). These observations can be intuitively understood as follows. Note that, we are interested in the asymptotic AoI of the users averaged over an arbitrarily long time-horizon (*viz.* Eqn. (20)). For many "regular" stochastic processes, *e.g.*, positive recurrent Markov Chains or general Ergodic processes, it is well-known that the long-term behaviour of the process is entirely determined by its limiting steady-state distribution (see, *e.g.*, Theorem 7.2.1 of [59]). In practice, we are interested in scheduling policies that are "well-behaved" in the above sense. Hence, it is not surprising that the lower bound (22), which is used to upper bound the approximation ratio of the scheduling policies, depends only on the induced steady-state distribution and not on the details of the transition probabilities of the specific mobility model. The appearance of the quantity $g(\psi)$ in the lower bound should not be surprising either as it denotes the *typical* number of non-empty cells at a time in the long run. Since a Base Station can transmit a packet only if at least one user is present in its coverage area, the quantity $g(\psi)$, in some sense, represents the multi-user diversity of the system.

Closed-form expression for $g(\psi)$: To get a sense of the bound (22), we now derive a closed-form expression for $g(\psi)$ under certain assumptions. This expression will be used later

in Theorem 6 to establish the 2-optimality guarantee of the MMW policy discussed in the following section.

Using the linearity of expectation, we have

$$\begin{aligned} g(\psi) &= \mathbb{E}_{\psi} \sum_{j=1}^M \mathbb{1}(\text{BS}_j \text{ contains at least one user}) \\ &= \sum_{j=1}^M \mathbb{P}_{\psi}(\text{BS}_j \text{ contains at least one user}). \end{aligned} \quad (33)$$

Since the cells are disjoint, we readily conclude from Eqn. (33) that $g(\psi) \leq \min\{M, N\}$. Recall that ψ_{ij} denotes the probability that the i^{th} user is in BS_j . If the mobility of the users is independent of each other, the expected number of non-empty cells $g(\psi)$ in Eqn. (33) simplifies to:

$$g(\psi) = \sum_{j=1}^M \left(1 - \prod_{i=1}^N (1 - \psi_{ij})\right). \quad (34)$$

We now evaluate the above expression for the case when the limiting occupancy distribution of each user is uniform across all cells, i.e., $\psi_{ij} = \frac{1}{M}, \forall i, j$. The uniform stationary distribution arises, for example, when the user mobility can be modelled as a random walk on a regular graph [60]. In this case, Eqn. (34) simplifies to

$$g(\psi^{\text{unif}}) = M \left(1 - \left(1 - \frac{1}{M}\right)^N\right). \quad (35)$$

For $M = 1$, we have $g(\psi) = 1$. For $M \geq 2$, we have the following bounds:

$$e^{-\frac{\beta}{M}} \stackrel{(a)}{\leq} \left(1 - \frac{1}{M}\right) \stackrel{(b)}{\leq} e^{-\frac{1}{M}}, \quad (36)$$

where $\beta \equiv \log(4) \leq 1.387$. The inequality (b) is standard. To prove the inequality (a), consider the concave function

$$f(x) = 1 - x - e^{-\beta x}, 0 \leq x \leq \frac{1}{2},$$

for some $\beta > 0$. Since a concave function of a real variable defined on an interval attains its minima at one of the end points of the closed interval, and since $f(0) = 0$, we have $f(x) \geq 0, \forall x \in [0, \frac{1}{2}]$, if $f(\frac{1}{2}) \geq 0$, i.e., $e^{\beta/2} \geq 2$, i.e., $\beta \geq \ln(4)$. Thus, the inequality (a) holds for $M \geq 2$ with $\beta = \ln(4)$. The inequality (36) directly leads to the following bounds for $M \geq 2$:

$$M \left(1 - e^{-\frac{N}{M}}\right) \leq g(\psi^{\text{unif}}) \leq M \left(1 - e^{-1.387 \frac{N}{M}}\right). \quad (37)$$

Please refer to Figure 5 for the plot of the normalized $g(\cdot)$ function.

2) *Achievability*: We now propose an online scheduling policy, called ‘‘Multi-cell Max-Weight’’ or π^{MMW} , that approximately minimizes the long-term average AoI.

The policy π^{MMW} : At each slot t , each BS transmits a packet to a user in its cell that has the highest index among all other users in its cell. The index $I_i(t)$ of the i^{th} user at time t is defined as $I_i(t) \equiv p_i h_i^2(t), \forall i \in [N]$.

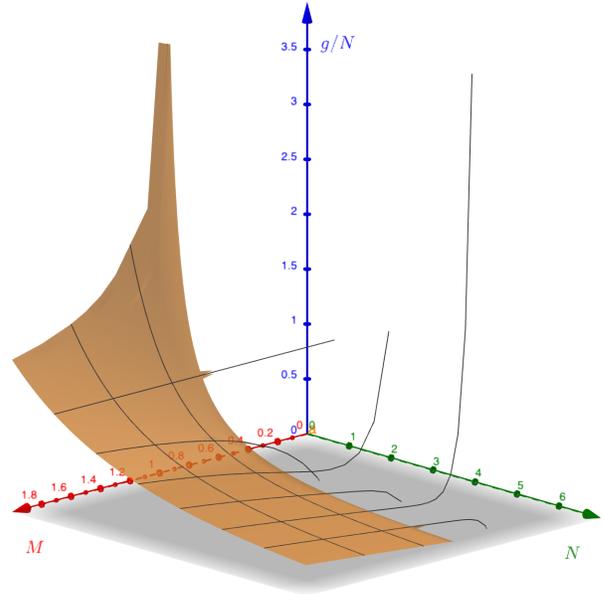


Fig. 5. Plot of $g(\psi)/N$ as a function of M and N . Note that the average AoI is inversely proportional to $g(\psi)/N$ in the case when all users are identical (i.e., $p_i = p, \forall i$).

The policy π^{MMW} is a multi-cell generalization of the 4-approximation single cell scheduling policy MW proposed in [29]. In Theorem 6, we show that π^{MMW} is a 2-approximation policy for statistically identical users with i.i.d. uniform mobility. The policy π^{MMW} is also a 2-approximation policy for static users located in a single cell (possibly with varying transmission success probabilities). Hence, our result strictly improves upon the 4-approximation guarantee of the MW policy for static users [29]. Our result also complements Theorem 5 of [29], where the authors showed that the MW policy is exactly optimal for statistically identical static users located in a common cell.

Theorem 6 (Achievability): π^{MMW} is a 2-approximation scheduling policy for the following two scenarios:

- 1) static users located in a common cell
- 2) statistically identical users with i.i.d. uniform mobility.

Note that scenarios (1) and (2) represent two extreme ends of the user mobility landscape. Scenario (1) corresponds to zero mobility, whereas scenario (2) corresponds to an infinite mobility model in the limit. In our notations, the second scenario is characterized by $p_i = p, \forall 1 \leq i \leq N$, $\psi_{ij} = \frac{1}{M}, \forall 1 \leq i \leq N, 1 \leq j \leq M$, such that the cell occupancy random variables $C_i(t)$'s are independent across time. The i.i.d. mobility model has been considered earlier in [16] in connection with characterizing the capacity and delay tradeoffs in wireless networks.

Proof: Let the scheduling decisions at slot t be denoted by the binary control vector $\mu(t) \in \{0, 1\}^N$, where $\mu_i(t) = 1$ if

and only if the following two conditions hold simultaneously: (1) $C_i(t) = j$, *i.e.*, the i^{th} user is within the coverage area of the j^{th} BS at slot t , for some $1 \leq j \leq M$, and (2) BS $_j$ schedules a packet transmission to the i^{th} user at time t . Since a BS can schedule only one transmission per slot to a user in its coverage area, the control vector must satisfy the following constraint:

$$\sum_{i:C_i(t)=j} \mu_i(t) \leq 1, \quad \forall j, t.$$

For performance analysis, we consider the following linear Lyapunov function:

$$L(\mathbf{h}(t)) = \sum_{i=1}^N \frac{h_i(t)}{\sqrt{p_i}}. \quad (38)$$

The above linear Lyapunov function should be contrasted with the quadratic Lyapunov function used in [29]. The conditional transition probabilities for the age of the i^{th} user may be expressed as follows (see Figure 2):

$$\begin{aligned} \mathbb{P}(h_i(t+1) = 1 | \mathbf{h}(t), \boldsymbol{\mu}(t+1), \mathbf{C}(t+1)) &= \\ \mu_i(t+1)p_i, \\ \mathbb{P}(h_i(t+1) = h_i(t) + 1 | \mathbf{h}(t), \boldsymbol{\mu}(t+1), \mathbf{C}(t+1)) &= \\ 1 - \mu_i(t+1)p_i, \end{aligned}$$

where the first conditional probability corresponds to the event $\mathcal{E}_i(t)$ where the i^{th} user was scheduled at slot t and the packet transmission was successful, and the second equation corresponds to the complement of the previous event. Hence, for each user $i \in [N]$, we have:

$$\begin{aligned} \mathbb{E}(h_i(t+1) | \mathbf{h}(t), \boldsymbol{\mu}(t+1), \mathbf{C}(t+1)) \\ = h_i(t) - \mu_i(t+1)p_i h_i(t) + 1. \end{aligned} \quad (39)$$

From the above equation, we can evaluate the one-step conditional drift of the Lyapunov function $L(\cdot)$ as follows:

$$\begin{aligned} \mathbb{E}(L(\mathbf{h}(t+1)) - L(\mathbf{h}(t)) | \mathbf{h}(t), \boldsymbol{\mu}(t+1), \mathbf{C}(t+1)) \\ = - \sum_{i=1}^N \mu_i(t+1) \sqrt{p_i} h_i(t) + \sum_{i=1}^N \frac{1}{\sqrt{p_i}}. \end{aligned} \quad (40)$$

Now consider the policy π^{MMW} , under which, each Base Station schedules the user i having the highest weight $\sqrt{p_i} h_i(t)$ among all users in its cell. To analyze the performance of the π^{MMW} policy, we define an auxiliary stationary randomized scheduling policy **RAND**, under which every BS randomly schedules a user in its cell with probability $\mu_i^{\text{RAND}}(t+1) \propto 1/\sqrt{p_i}$ ¹. Using the basic fact that maximum of a set of real numbers is at least as large as any convex combination of the numbers, we conclude that:

$$\begin{aligned} \mathbb{E}\left(\sum_{i=1}^N \mu_i^{\text{MMW}}(t+1) \sqrt{p_i} h_i(t) | \mathbf{h}(t), \boldsymbol{\mu}(t+1), \mathbf{C}(t+1)\right) \\ \geq \sum_{j=1}^M \frac{\sum_{i:C_i(t+1)=j} h_i(t)}{\sum_{i:C_i(t+1)=j} \frac{1}{\sqrt{p_i}}}. \end{aligned}$$

¹We use the usual convention that summation over an empty set is zero.

Hence, combining the above inequality with Eqn. (40), we have the following upper bound on the one-step expected drift of the Lyapunov function under the action of the policy π^{MMW} :

$$\begin{aligned} \mathbb{E}^{\text{MMW}}(L(\mathbf{h}(t+1)) - L(\mathbf{h}(t)) | \mathbf{h}(t), \mathbf{C}(t+1)) \\ \leq - \sum_{j=1}^M \frac{\sum_{i:C_i(t+1)=j} h_i(t)}{\sum_{i:C_i(t+1)=j} \frac{1}{\sqrt{p_i}}} + \sum_{i=1}^N \frac{1}{\sqrt{p_i}}. \end{aligned} \quad (41)$$

Taking expectation of the above inequality w.r.t. the random cell-occupancy vector $\mathbf{C}(t+1)$, we have:

$$\begin{aligned} \mathbb{E}^{\text{MMW}}(L(\mathbf{h}(t+1)) - L(\mathbf{h}(t)) | \mathbf{h}(t)) \\ \leq - \sum_{j=1}^M \mathbb{E}(Z_j(t) | \mathbf{h}(t)) + \sum_{i=1}^N \frac{1}{\sqrt{p_i}}, \end{aligned} \quad (42)$$

where we define the random variable $Z_j(t) \equiv \frac{\sum_{i:C_i(t+1)=j} h_i(t)}{\sum_{i:C_i(t+1)=j} \frac{1}{\sqrt{p_i}}}$, $\forall j$, and let $0/0 \equiv 0$. This convention is consistent, because if a cell is empty, the corresponding term is not present in the summation (42).

The inequality (42) holds for any arbitrary user mobility model. To make progress, we now analyze two special cases as stated in the statement of the theorem.

CASE I: STATIC USERS LOCATED IN A COMMON CELL

In this case, we have $M = 1$, and hence,

$$\mathbb{E}^{\text{MMW}}(L(\mathbf{h}(t+1)) - L(\mathbf{h}(t)) | \mathbf{h}(t)) \leq - \frac{\sum_{i=1}^N h_i(t)}{\sum_{i=1}^N \frac{1}{\sqrt{p_i}}} + \sum_{i=1}^N \frac{1}{\sqrt{p_i}}.$$

Taking expectations of both sides w.r.t. $\mathbf{h}(t)$, we obtain:

$$\mathbb{E}(L(\mathbf{h}(t+1)) - L(\mathbf{h}(t))) \leq - \frac{\sum_{i=1}^N \mathbb{E}(h_i(t))}{\sum_{i=1}^N \frac{1}{\sqrt{p_i}}} + \sum_{i=1}^N \frac{1}{\sqrt{p_i}}.$$

Summing up the above inequalities for $t = 1, 2, \dots, T$, dividing both sides by T and then taking the limit as $T \rightarrow \infty$, we obtain:

$$\text{AoI}_{\text{avg}}^{\text{MMW}} = \limsup_{T \rightarrow \infty} \frac{1}{NT} \sum_{t=1}^T \sum_{i=1}^N \mathbb{E}(h_i(t)) \leq \frac{1}{N} \left(\sum_{i=1}^N \frac{1}{\sqrt{p_i}} \right)^2.$$

Comparing the above with the lower bound in Eqn. (22) and realizing that $g(\psi) = 1$ in this case, we have

$$\text{AoI}_{\text{avg}}^{\text{MMW}} \leq 2\text{AoI}_{\text{avg}}^*.$$

CASE II: STATISTICALLY IDENTICAL USERS WITH I.I.D. UNIFORM MOBILITY

In the case of i.i.d. uniform mobility, the users move to any one of the M Base Stations chosen uniformly at random independent of everything else. We begin our analysis from the drift upper bound given in equation (42). Note that, given the age vector $\mathbf{h}(t)$, under this mobility model, the r.v.s $Z_j(t)$'s are identically distributed. Hence,

$$\sum_{j=1}^M \mathbb{E}(Z_j(t) | \mathbf{h}(t)) = M \mathbb{E}(Z_1(t) | \mathbf{h}(t)).$$

The RHS of the above can be expressed as the expectation of the ratio of two correlated random variables:

$$\mathbb{E}(Z_1(t) | \mathbf{h}(t)) = \mathbb{E}_{\mathbf{h}(t)} \left[\frac{\sum_{i=1}^N h_i(t) W_i}{\sum_{i=1}^N \frac{1}{\sqrt{p_i}} W_i} \right], \quad (43)$$

where the indicator r.v.s $\{W_i\}_{i=1}^N$ are i.i.d. such that

$$\mathbb{P}(W_i = 1) = \frac{1}{M} = 1 - \mathbb{P}(W_i = 0).$$

In general, it is non-trivial to obtain a tight lower bound to the expression (43), which is defined as the expectation of the ratio of two correlated random variables. In the following, we evaluate the expectation exactly in the special case of statistically identical users with $p_i = p, \forall i \in [N]$.

Note that, we can express the random variable $\sum_{j=1}^M Z_j(t)$ as

$$\sum_{j=1}^M Z_j(t) = \sum_{i=1}^N h_i(t) Y_i(t),$$

where $Y_i(t) = \left(\frac{1}{\sqrt{p_i}} + \sum_{k \neq i} \frac{1}{\sqrt{p_k}} \mathbb{1}(C_i(t) = C_k(t)) \right)^{-1}$. Given our i.i.d. mobility assumption, the r.v. $\mathbf{C}(t+1)$ is independent of $\mathbf{h}(t)$. Hence,

$$\mathbb{E}(Y_i(t) | \mathbf{h}(t)) = \sum_{n=0}^{N-1} \sum_{S: i \notin S, |S|=n} \left(\frac{1}{\sqrt{p_i}} + \sum_{k \in S} \frac{1}{\sqrt{p_k}} \right)^{-1} \times \frac{1}{M^n} \left(1 - \frac{1}{M} \right)^{N-n-1}. \quad (44)$$

Since the users are statistically identical, i.e., $p_i = p, \forall i$, the summation (44) has a closed-form expression. Clearly, for all $0 \leq n \leq N-1$, we have:

$$Y_i(t) = \frac{\sqrt{p}}{n+1}, \quad \text{w.p.} \binom{N-1}{n} \frac{1}{M^n} \left(1 - \frac{1}{M} \right)^{N-n-1}.$$

To evaluate the expectation of $Y_i(t)$, we integrate the binomial expansion of $(1+x)^{N-1}$ in the range $[0, \beta]$ to obtain the identity:

$$\frac{1}{N} \left((1+\beta)^N - 1 \right) = \beta \sum_{n=0}^{N-1} \frac{1}{n+1} \binom{N-1}{n} \beta^n.$$

Substituting $\beta = \frac{1}{M-1}$ in the above, we obtain

$$\mathbb{E}(Y_i(t) | \mathbf{h}(t)) = \sqrt{p} \frac{M}{N} \left(1 - \left(1 - \frac{1}{M} \right)^N \right) \equiv Y^* \text{ (say)}. \quad (45)$$

From Eqn. (42) and (45), we have

$$\mathbb{E}^{\text{MMW}}(L(\mathbf{h}(t+1)) - L(\mathbf{h}(t)) | \mathbf{h}(t)) \leq -Y^* \sum_i h_i(t) + \frac{N}{\sqrt{p}}.$$

Taking expectation of both sides, we have

$$\mathbb{E}^{\text{MMW}}(L(\mathbf{h}(t+1)) - L(\mathbf{h}(t))) \leq -Y^* \sum_i \mathbb{E} h_i(t) + \frac{N}{\sqrt{p}}.$$

Summing up the above inequalities for $t = 1, 2, \dots, T$, dividing both sides by T and then taking limit as $T \rightarrow \infty$, we obtain

$$\begin{aligned} \text{AoI}_{\text{avg}}^{\text{MMW}} &= \limsup_{T \rightarrow \infty} \frac{1}{NT} \sum_{t=1}^T \sum_i \mathbb{E} h_i(t) \\ &\leq \frac{N}{Y^* \sqrt{p}} = \frac{N}{Mp \left(1 - \left(1 - \frac{1}{M} \right)^N \right)}. \end{aligned} \quad (46)$$

On the other hand, the lower bound from Theorem 5, specialized to the case of identical users, yields:

$$\text{AoI}_{\text{avg}}^* \geq \frac{N}{2Mp \left(1 - \left(1 - \frac{1}{M} \right)^N \right)}. \quad (47)$$

Eqns. (46) and (47), we have

$$\text{AoI}_{\text{avg}}^{\text{MMW}} \leq 2 \text{AoI}_{\text{avg}}^*.$$

The above inequality shows that the MMW scheduling policy is 2-optimal for statistically identical users with uniform i.i.d. mobility. ■

Discussion: In this Section, we have presented a universal lower bound (Theorem 5) and near-optimal achievability results for two important special cases (Theorem 6) for the average AoI metric in the stochastic setting. It will be interesting to establish similar achievability results for more general user mobility (e.g., random walk) and channel models (e.g., Gilbert-Elliott model [61]).

B. Minimizing the Peak AoI in the Stochastic Framework

In this section, we consider the problem of minimizing the long-term *peak*-AoI metric (21), for static users in a single cell. By directly solving the associated countable state Bellman equation, we prove that the greedy CMA scheduling policy is optimal for minimizing the Peak AoI. Furthermore, we also establish the large-deviation optimality of the CMA policy.

Theorem 7 (Optimal Policy for minimizing the Peak AoI): The greedy CMA policy is optimal for minimizing the peak AoI objective given in Eqn. (21). Moreover, the optimal value of the peak AoI is given by $\text{AoI}_{\text{peak}}^* = \sum_{i=1}^N \frac{1}{p_i}$.

We prove Theorem 7 by proposing a closed-form solution to the Bellman's equation corresponding to the associated average-cost MDP and then verifying that the candidate solution indeed satisfies the Bellman's equation.

Proof: The stochastic control problem under investigation is an instance of a countable-state average-cost MDP with a finite action space with the following components. The state of the system at a slot t given by the instantaneous AoI of all users: $\mathbf{h}(t) \equiv (h_1(t), h_2(t), \dots, h_N(t))$. The per-stage cost at time t is $\max_{i=1}^N h_i(t)$, which is unbounded, in general. Finally, the finite action space $\mathcal{A} = \{1, 2, \dots, N\}$ corresponds to the user scheduled at a given slot.

Let the optimal cost for the problem be denoted by λ^* and the differential cost-to-go from the state \mathbf{h} be denoted by $V(\mathbf{h})$ [62]. Then, following the standard theory of average cost countable state MDP (Proposition 4.6.1 of [62]), we set up the following Bellman Eqn.

$$\begin{aligned} &\lambda^* + V(\mathbf{h}) \\ &= \min_i \{ p_i V(\underbrace{\mathbf{1}}_{i^{\text{th}} \text{ coordinate}}, \mathbf{h}_{-i} + \mathbf{1}) + (1-p_i) V(\mathbf{h} + \mathbf{1}) \} \\ &+ \max_i h_i, \end{aligned} \quad (48)$$

where the vector \mathbf{h}_{-i} denotes the $N - 1$ dimensional vector of all coordinates excepting the i^{th} coordinate and $\mathbf{1}$ is an all-one vector of an appropriate dimension.

Discussion: The Bellman Equation (48) may be derived as follows. Suppose that the current age of the users is given by the vector \mathbf{h} . If the policy schedules a packet transmission to the i^{th} user, the transmission is successful with probability p_i and is unsuccessful with probability $1 - p_i$. If the transmission is successful, the AoI of all users, excepting the i^{th} user, is incremented by 1, and the AoI of the i^{th} user is reduced to 1. This explains the first term. On the other hand, if the transmission to the i^{th} user is unsuccessful, the AoI of all users are incremented by 1. This explains the second term within the bracket. Finally, the term $\max_i h_i$ denotes the stage cost.

Solution to the Bellman Equation (48): We now verify that the following constitutes a solution to the Bellman Equation (48):

$$V(\mathbf{h}) = \sum_j \frac{h_j}{p_j}, \quad \lambda^* = \sum_j \frac{1}{p_j}. \quad (49)$$

To verify the above solution, we start with the RHS of Eqn. (48). Upon substitution from Eqn. (49), the expression corresponding to the i^{th} user inside the min operator of Eqn. (48) simplifies to:

$$\begin{aligned} & p_i V(\mathbf{1}, \mathbf{h}_{-i} + \mathbf{1}) + (1 - p_i) V(\mathbf{h} + \mathbf{1}) \\ &= p_i \sum_{j \neq i} \frac{h_j + 1}{p_j} + 1 + (1 - p_i) \sum_{j=1}^N \frac{h_j + 1}{p_j} \\ &= p_i \sum_{j=1}^N \frac{h_j + 1}{p_j} - p_i \frac{h_i + 1}{p_i} + 1 + (1 - p_i) \sum_{j=1}^N \frac{h_j + 1}{p_j} \\ &= \sum_{j=1}^N \frac{h_j}{p_j} - h_i + \sum_{j=1}^N \frac{1}{p_j}. \end{aligned} \quad (50)$$

Hence,

$$\begin{aligned} \text{RHS} &= \min_i \{ p_i V(\mathbf{1}, \mathbf{h}_{-i} + \mathbf{1}) + (1 - p_i) V(\mathbf{h} + \mathbf{1}) \} + \max_i h_i \\ &= \sum_{j=1}^N \frac{1}{p_j} + \sum_{j=1}^N \frac{h_j}{p_j} - \max_i h_i + \max_i h_i \\ &= \lambda^* + V(\mathbf{h}) \\ &= \text{LHS}. \end{aligned}$$

Finally, to verify the regularity condition (Eqn. 4.122 of [62]), note that

$$\begin{aligned} \frac{\mathbb{E}(V(\mathbf{h}(t)) | \mathbf{h}(1))}{t} &= \frac{1}{t} \mathbb{E} \left[\sum_{j=1}^N \frac{h_j(t)}{p_j} | \mathbf{h}(1) \right] \\ &\leq \mathbb{E} \left[\frac{\max_{i=1}^N h_i(t)}{t} | \mathbf{h}(1) \right] \left(\sum_i \frac{1}{p_i} \right). \end{aligned}$$

Hence, for any scheduling policy π for which the regularity condition is violated, *i.e.*, $\limsup_{t \rightarrow \infty} \frac{\mathbb{E}^\pi(V(\mathbf{h}(t)) | \mathbf{h}(1))}{t} > 0$, for some $\mathbf{h}(1)$ with positive probability, the above bound implies that

$$\limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \mathbb{E}^\pi \left[\max_{i=1}^N h_i(t) \right] = \infty.$$

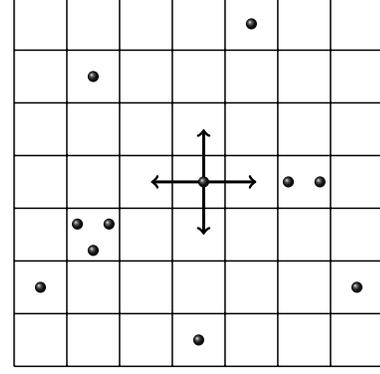


Fig. 6. Schematic of a grid network used in the simulations in the stationary regime. The black balls represent the users and the small squares denote the cells. There is one Base Station at every cell (not shown in this schematic). At every step, each user randomly moves to any one of the adjacent cells with equal probability.

Hence, without any loss of optimality, we may confine our attention to those policies for which the regularity condition holds. The proof now follows from Proposition 4.6.1 of [62]. ■

Large-Deviation Optimality: Theorem 7 establishes that the greedy CMA scheduling policy is optimal for minimizing the long-term *expected* peak AoI. However, for some applications, the scheduling policy is additionally required to ensure that the peak AoI metric stays within limit with a high probability after a sufficiently long time. An analogous problem was considered in connection with the queue stability in a similar multi-user single cell setting in [43]. The previous requirement leads to the following problem statement: Design a scheduling policy $\pi \in \Pi$ that maximizes the large deviation exponent, *i.e.*,

$$\max_{\pi \in \Pi} \left[- \lim_{k \rightarrow \infty} \lim_{t \rightarrow \infty} \frac{1}{k} \log \mathbb{P}^\pi \left(\max_{i=1}^N h_i(t) \geq k \right) \right]. \quad (51)$$

The following theorem shows that the CMA policy is also optimal in the above sense.

Theorem 8: The CMA policy achieves the optimal large-deviation exponent whose value is given by

$$- \lim_{k \rightarrow \infty} \lim_{t \rightarrow \infty} \frac{1}{k} \log \mathbb{P}^{\text{CMA}} \left(\max_i h_i(t) \geq k \right) = - \log \left(1 - \min_{i=1}^N p_i \right).$$

See Appendix VI-C for the proof.

IV. NUMERICAL EXPERIMENTS

A. Stationary regime

In our numerical experiments, we first simulate the CMA and the MMW policies in the stationary regime and compare their performances for minimizing the average AoI and peak AoI. We consider a grid network with $M = 100$ base stations arranged in the form of a 10×10 square grid. N users are initially placed uniformly at random on any of the M Base Stations. Each user executes an independent 2-dimensional

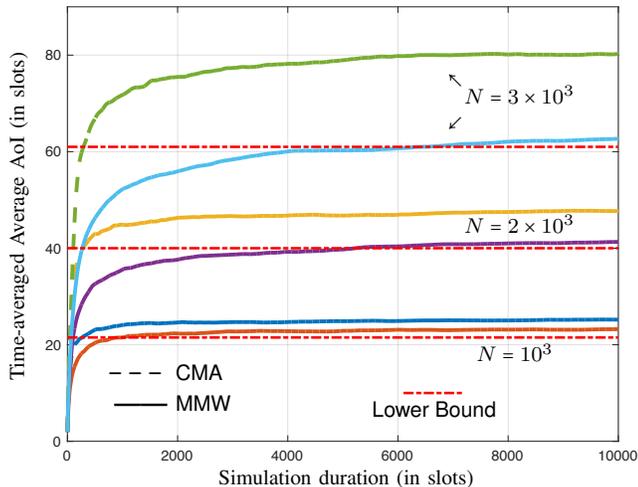


Fig. 7. Comparison between the CMA and MMW policy for statistically non-identical users for the Average-age metric in the stationary regime. The users are assumed to execute independent 2-D random walks on a 10×10 grid. The transmission success probability of each user is sampled uniformly at random from the interval $[0, 1]$. From the plots, it follows that the MMW policy outperforms the CMA policy for minimizing the average AoI.

random walk on the square grid at every time slot and moves to any one of the adjacent cells with equal probability. See Figure 6 for a schematic. As in our system model, we assume that each Base Station can schedule a new packet transmission to only one user located in its cell per slot. In the stationary regime, the transmission success probabilities p of each of the users are sampled independently and uniformly at random from the interval $I \subseteq [0, 1]$ and remains fixed throughout the simulation duration.

Results: Figure 7 shows the time variation of the average AoI for three different scenarios $N = 1000, 2000, \text{ and } 3000$. Not surprisingly, the average AoI increases as the number of users N in the network is increased. More interestingly, by numerically evaluating the lower bound (22) using the values of the simulation parameters, we see that the simulated long-term AoI under the MMW policy is very close (within $\sim 1\%$) to the theoretical lower bound. This result indicates that (1) the lower bound is tight, and (2) the performance of the MMW policy is near-optimal. We also see that the greedy CMA policy is outperformed by the MMW policy, which takes into account the transmission success probabilities of different users. However, as the plot 8 shows, when it comes to maximizing the peak AoI, the greedy CMA policy consistently outperforms the MMW policy. Hence, the choice of the particular AoI metric plays an important role in determining the performance of different scheduling policies.

B. Non-stationary regime

Next, we numerically evaluate the performance of the proposed scheduling policies in a non-stationary setup where the probability of successful transmission to each user is allowed to vary with time. Specifically, in this set of simulations, the users' movement is modeled using the *Lévy Mobility model*, which is known to capture human mobility accurately [63]. We simulate a discrete time version of this model, where at each

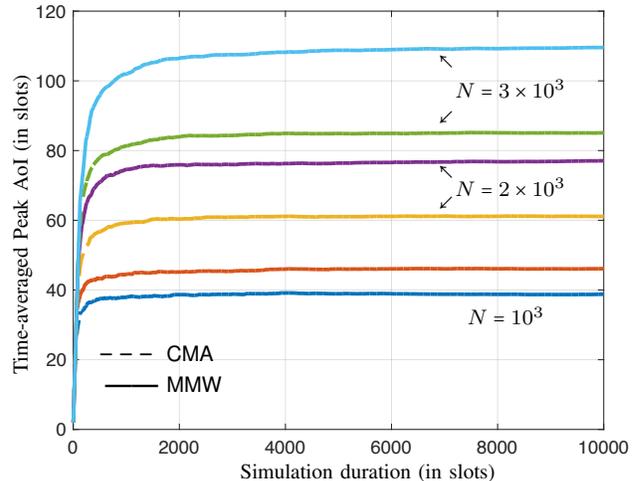


Fig. 8. Comparison between the CMA and MMW policy for statistically non-identical users for the Peak-Age metric in the stationary regime. The users are assumed to execute independent 2-D random walks on a 10×10 grid. The transmission success probability of each user is sampled uniformly at random from the interval $[0.2, 0.8]$. From the plots, it follows that the CMA policy outperforms the MMW policy for minimizing the peak AoI.

slot, each user moves independently at random by a distance ρ at an angle θ from its current location. The distance ρ and the angle θ are sampled uniformly at random from the intervals $[0, v_{\max}]$ and $[0, 2\pi]$ respectively, where v_{\max} is a constant determined by the maximum velocity of the users. We assume that a packet loss occurs when the received SNR falls below a certain threshold (*i.e.*, when a channel *outage* occurs). Under this assumption, we use the following approximate expression for the probability of a successful packet reception over a slow-fading wireless channel for a user i located at a distance $d_i(t)$ from its associated BS [48]:

$$p_i(t) = 1 - \left(\frac{1}{\underbrace{1 + G \cdot L(d_i(t))^{-1} \cdot \text{SNR}_i}_{\text{received SNR}}} \right)^\eta \quad (52)$$

In the above expressions, G is the receive antenna gain, $L(\cdot)$ denotes the path loss function, η is the diversity order of the system, and SNR_i represents the transmit Signal-to-Noise ratio for the i^{th} user. Note that the received SNR for a user is obtained by multiplying the transmit SNR, receive-antenna gain, and the reciprocal of the path loss function. In our simulations, we use the following mean alpha-beta-gamma (ABG) 5G path loss model as proposed in [64]:

$$L^{\text{ABG}}(d)[\text{dB}] = 10\alpha \log_{10} \left(\frac{d}{1 \text{ m}} \right) + \beta + 10\gamma \log_{10} \left(\frac{f_c}{1 \text{ GHz}} \right). \quad (53)$$

In the path loss formula above, α corresponds to the path-loss exponent, β is an optimized offset value for the path loss in dB and the parameter γ gives the dependence of the path loss on the carrier frequency f_c .

Simulation setup: We consider a 5G mmWave cellular wireless system operating at a carrier frequency of $f_c = 28$ GHz. Assume that $M = 100$ base stations are arranged in the form of a 10×10 uniform square grid such that the Inter-Site Distance (ISD) between any two adjacent base stations is 100 m. Base stations periodically receive fresh update packets

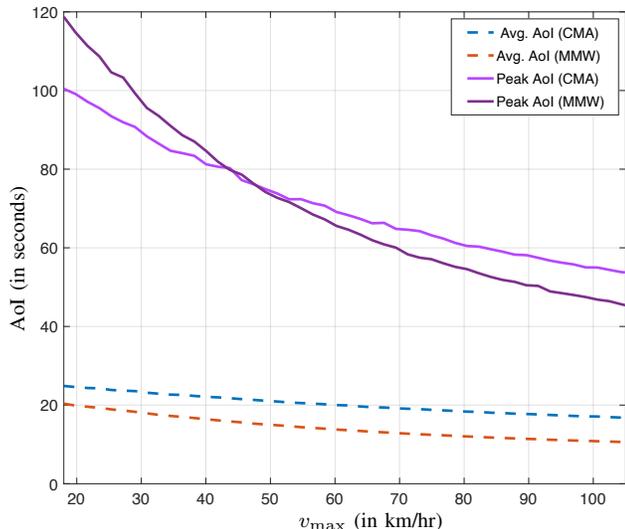


Fig. 9. Performance comparison between the CMA and MMW policies in the non-stationary regime. The propagation environment corresponds to an NLOS Urban Microcell environment consisting of 100 Base Stations and $N = 1000$ users. From the plots, we see that AoI decreases as the mobility of the users increases. Furthermore, the MMW policy out-performs the CMA policy in the high-mobility regime.

every second from external sources, and each BS schedules a downlink transmission immediately upon the reception of a fresh packet. We assume that $N = 1000$ users are roaming in the area at the maximum speed of v_{\max} km/hr according to the Lévy mobility model described above. At the beginning of the simulations, the users are placed uniformly at random over the entire region. The path-loss parameters are taken to be $\alpha = 3.5, \beta = 24.4$ dB, $\gamma = 1.9$, which correspond to an NLOS Urban Microcell environment [64]. The diversity order of the system is taken to be $\eta = 5$, the transmit Signal-to-Noise ratio (SNR) is taken to be 90 dB, and the gain of the receive antenna is taken to be $G = 20$ dB.

Results and Discussions: We numerically compare the performance of the following two policies in the non-stationary regime - (1) the CMA policy, and (2) the MMW policy. The MMW policy updates the probability of successful transmission to each user according to its distance from the associated BS (viz. Eqns. (52) and (53)). The average and the peak AoI metrics attained by these two policies are plotted in Figure 9 as a function of the maximum speed of the users v_{\max} . From the plots, we see that, in general, the peak and average AoI decrease as the maximum speed of the users increases. Furthermore, the MMW policy performs better than the CMA policy in minimizing the average AoI metric. On the other hand, the CMA policy is more effective in reducing the peak AoI metric up to some critical speed, beyond which the MMW policy performs better.

V. CONCLUDING REMARKS

In this paper, we investigated the fundamental limits of Age-of-Information for mobile users in adversarial and stochastic environments. We also proposed efficient scheduling policies that come close to achieving the limits. In particular, we

showed that a greedy scheduling policy (CMA) is near-optimal for minimizing the peak AoI in the adversarial setting. The competitive ratio of the same policy is shown to be within a factor of $O(N)$ of the optimal value for minimizing the average AoI. In the stochastic setting, we proved that a Max-Weight-type scheduling policy (MMW) attains 2-approximation for minimizing the average AoI in two extreme mobility scenarios. Furthermore, the CMA policy is shown to be optimal for minimizing the peak AoI for static users in a single cell. Although, in this paper, we made some progress in the AoI-optimal scheduling problem for cellular networks, many interesting questions are still open. In the adversarial setting, an interesting problem is to reduce the current $O(N)$ gap to optimality for minimizing the average AoI metric. The problem of designing an optimal policy for minimizing the peak AoI for mobile users in the stochastic environment is open and can be investigated in the future. Extending the achievability result in the stochastic setting to more general mobility and channel state models (e.g., Gilbert-Elliott model [61]) is an important research direction. Furthermore, designing optimal scheduling algorithms in the adversarial setting when some estimate of the channel state for the immediate future is known, constitutes an interesting problem from a practical point-of-view [65].

ACKNOWLEDGEMENTS

This work is partially supported by the grant IND-417880 from Qualcomm, USA and a research grant from the Govt. of India under the Institutes of Eminence (IoE) initiative. The first author would like to thank his former students Subhankar Banerjee and Arunabh Srivastava for some useful discussions during the early stages of this work.

VI. APPENDIX

A. Proof of Theorem 2

Proof: We prove the theorem by exhibiting a channel state sequence for which the CMA policy achieves a competitive ratio of at least N^2 for the average AoI objective and at least $2N - 1$ for the peak AoI objective.

Consider a single-cell scenario where a BS serves N stationary users. In this case, it is easy to see that the CMA policy reduces to a persistent round-robin policy - the users are scheduled in a round-robin fashion such that each user is scheduled continuously until its transmission is successful. Consider the following channel-state sequence: all super-intervals are of constant length $\Delta \gg N$, where we take $\Delta \equiv 1 \pmod{N-1}$. During each super-interval, all users, apart from the Max-user, have Good channels at every slot. Hence, under the CMA policy, the time interval between two consecutive successful packet transmissions to any user is $N\Delta$ slots. Thus, at the beginning of any super-interval (apart from the first $N-1$ super-intervals) the ages of the users under the CMA policy in *ascending* order is given by $1, \Delta+1, 2\Delta+1, \dots, (N-1)\Delta+1$. Next, we consider the following two objectives.

a) *Average AoI*: The total cost incurred by the CMA policy in any interval (apart from the first $N-1$ super-intervals) is given by:

$$C_i^{\text{CMA}} = \sum_{j=1}^N \sum_{k=1}^{\Delta} ((j-1)\Delta + k) = \frac{1}{2}(N^2\Delta^2 + N\Delta).$$

Next, we upper-bound the cost incurred by the offline optimal policy OPT by comparing it to another (potentially sub-optimal) offline policy \mathcal{P} . The policy \mathcal{P} serves each of the $N-1$ users other than the Max user in a round-robin fashion in each super-interval and finally, it serves the Max user at the last time slot of each super-interval (See Figure 4). Clearly, under the action of the policy \mathcal{P} , the set of ages of the users at the beginning of every super-interval is given by $\{1, 2, \dots, N\}$. Thus the total cost incurred by the policy \mathcal{P} during the i^{th} super-interval is bounded as:

$$\begin{aligned} C_i^{\mathcal{P}} &\leq \underbrace{\sum_{k=1}^{\Delta} k}_{\text{Cost due to the Max user}} \\ &\quad + \underbrace{(N-1) \left(\left(\frac{\Delta-1}{N-1} - 1 \right) \frac{N(N-1)}{2} + \frac{3N(N-1)}{2} \right)}_{\text{Cost due to the other } N-1 \text{ users}} \\ &\leq \frac{\Delta^2}{2} + \frac{\Delta N^2}{2} + N^3. \end{aligned} \quad (54)$$

Let K be the number of super-intervals in the time-horizon T . We have

$$\text{Cost}^{\text{OPT}}(T) \leq \sum_{i=1}^K C_i^{\mathcal{P}}.$$

Hence, for large enough K , the competitive ratio of the CMA policy is lower bounded as:

$$\eta_{\text{avg}}^{\text{CMA}} = \frac{\sum_{i=1}^K C_i^{\text{CMA}}}{\sum_{i=1}^K C_i^{\text{OPT}}} \geq \frac{\frac{1}{2}(N^2\Delta^2 + N\Delta)}{\frac{\Delta^2}{2} + \frac{\Delta N^2}{2} + N^3}.$$

By taking Δ to be arbitrarily large, it follows from the above expression that

$$\eta_{\text{avg}}^{\text{CMA}} \geq N^2.$$

b) *Peak AoI*: We compute the competitive ratio for the same channel state sequence as before for the peak AoI objective. The cost incurred by the CMA policy during any super-interval is:

$$C_i^{\text{CMA}} = \sum_{k=1}^{\Delta} ((N-1)\Delta + k) = (N-1)\Delta^2 + \frac{\Delta(\Delta+1)}{2}. \quad (55)$$

Similar to the Average AoI case, we use the policy \mathcal{P} to upper-bound the cost incurred by the OPT policy for the given channel state sequence. As before, the total cost incurred by the policy \mathcal{P} during the i^{th} super-interval is bounded as:

$$C_i^{\mathcal{P}} \leq \sum_{k=1}^{\Delta} (N+k) = N\Delta + \frac{\Delta(\Delta+1)}{2}.$$

Thus, summing over all super-intervals, we obtain

$$\eta_{\text{peak}}^{\text{CMA}} = \frac{\sum_{i=1}^K C_i^{\text{CMA}}}{\sum_{i=1}^K C_i^{\text{OPT}}} \geq \frac{(N-1)\Delta^2 + \frac{\Delta(\Delta+1)}{2}}{N\Delta + \frac{\Delta(\Delta+1)}{2}}.$$

The result now follows by letting the length of the sub-intervals Δ to be arbitrarily large. \blacksquare

B. Proof of Proposition 2

Proof: We use the same proof technique as in Theorem 4 with the channel states for the two user being i.i.d. $\{\text{Good}, \text{Bad}\}$ or $\{\text{Bad}, \text{Good}\}$ w.p. $1/2$ each. Let $\vec{\mu}(t) \in \{(1, 0), (0, 1)\}$ denote the scheduling decision of the policy π at slot t . Define $\mathcal{F}_{t-1} \equiv \sigma(\vec{h}(k), \vec{\mu}(k), 1 \leq k \leq t-1)$ to be the sigma-algebra generated by the age and scheduling decision r.v.s up to time $t-1$. For any online policy, the scheduling decision $\vec{\mu}(t)$ at time t must be measurable in \mathcal{F}_{t-1} , $\forall t \geq 1$. Let $H_{\text{sum}}(t) \equiv \mathbb{E}^{\pi}(h_1(t)) + \mathbb{E}^{\pi}(h_2(t))$ be the expected sum of the ages of the users at time t . Let $B_t \in \mathcal{F}_t$ be the event for which the user 1 is scheduled under the policy π . Then, we can write

$$\begin{aligned} &\mathbb{E}^{\pi}(h_1(t+1)|\mathcal{F}_t) \\ &= \left(1 + \frac{1}{2}h_1(t)\right)\mathbb{1}(B_t) + \left(1 + h_1(t)\right)\mathbb{1}(B_t^c) \\ &= 1 + \frac{1}{2}h_1(t) + \frac{1}{2}h_1(t)\mathbb{1}(B_t^c) \\ &\stackrel{(a)}{\geq} 1 + \frac{1}{2}h_1(t) + \frac{1}{2}\min\{h_1(t), h_2(t)\}\mathbb{1}(B_t^c), \end{aligned} \quad (56)$$

Similarly, we can also write

$$\begin{aligned} &\mathbb{E}^{\pi}(h_2(t+1)|\mathcal{F}_t) \geq \\ &1 + \frac{1}{2}h_2(t) + \frac{1}{2}\min\{h_1(t), h_2(t)\}\mathbb{1}(B_t). \end{aligned} \quad (58)$$

Since $\mathbb{1}(B_t) + \mathbb{1}(B_t^c) = 1$, from the equations (56) and (58), we have

$$\begin{aligned} &\mathbb{E}^{\pi}(h_1(t+1) + h_2(t+1)|\mathcal{F}_t) \geq \\ &2 + \frac{1}{2}(h_1(t) + h_2(t)) + \frac{1}{2}\min\{h_1(t), h_2(t)\}. \end{aligned}$$

Taking expectations of both sides of the above equation, we get

$$H_{\text{sum}}(t+1) \geq 2 + \frac{1}{2}H_{\text{sum}}(t) + \frac{1}{2}\mathbb{E}\left(\min\{h_1(t), h_2(t)\}\right). \quad (59)$$

Let the random variable $S(t)$ denote the time elapsed since the last successful transmission (by any user) before time t . Clearly,

$$\min\{h_1(t), h_2(t)\} \geq S(t)$$

(the above inequality holds with equality for the two user case). Hence, the above inequality implies

$$H_{\text{sum}}(t+1) \geq 2 + \frac{1}{2}H_{\text{sum}}(t) + \frac{1}{2}\mathbb{E}(S(t)).$$

Summing up the above inequalities for $t = 1, 2, \dots, T$, and dividing both sides by T , we obtain

$$2\frac{H_{\text{sum}}(T+1)}{T} + \frac{1}{T}\sum_{t=1}^T H_{\text{sum}}(t) \geq 4 + \frac{1}{T}\sum_{t=1}^T \mathbb{E}(S(t)). \quad (60)$$

It is to be noted that $\{S(t)\}_{t \geq 1}$ is a renewal process with the time-stamp of successful transmissions constituting the

renewal instants. Let the random variable τ denote the length of any generic renewal cycle. Hence, using the renewal reward theorem [46] [66], it follows that

$$\begin{aligned} \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \mathbb{E}(S(t)) &= \frac{\mathbb{E}(\int_0^\tau S(t) dt)}{\mathbb{E}(\tau)} \\ &= \frac{\mathbb{E}(1 + 2 + \dots + \tau)}{\mathbb{E}(\tau)} \\ &= \frac{\mathbb{E}(\tau^2) + \mathbb{E}(\tau)}{2\mathbb{E}(\tau)} \\ &= 2, \end{aligned}$$

where the last inequality follows from the fact that the renewal cycle lengths T are distributed geometrically with the parameter $p = 1/2$. Thus, the limit of the RHS of Eqn. (60) exists and the limiting value is equal to 6. Next, we consider two possible cases.

Case I: $\liminf_{T \rightarrow \infty} \frac{H_{\text{sum}}(T+1)}{T} = 0$: In this case, consider a subsequence $\{T_k\}_{k \geq 1}$ along which $\lim_{k \rightarrow \infty} \frac{H_{\text{sum}}(T_k+1)}{T_k} = 0$.

For this subsequence, we have from Eqn. (60):

$$2 \frac{H_{\text{sum}}(T_k + 1)}{T_k} + \frac{1}{T_k} \sum_{t=1}^{T_k} H_{\text{sum}}(t) \geq 4 + \frac{1}{T_k} \sum_{t=1}^{T_k} \mathbb{E}(S(t)).$$

Taking $k \rightarrow \infty$, we conclude that

$$\limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T H_{\text{sum}}(t) \geq 6. \quad (61)$$

Case II: $\liminf_{T \rightarrow \infty} \frac{H_{\text{sum}}(T+1)}{T} = \alpha > 0$: From the definition of \liminf , it follows that there exists a finite T_0 such that, for all $T \geq T_0$, we have

$$\frac{H_{\text{sum}}(T+1)}{T} \geq \frac{\alpha}{2}. \quad (62)$$

Thus, for any $T \geq T_0$, we can write

$$\frac{1}{T} \sum_{t=1}^T H_{\text{sum}}(t) \geq \frac{1}{T} \sum_{t=T_0+1}^T H_{\text{sum}}(t) \stackrel{(a)}{\geq} \frac{\alpha}{2T} \sum_{t=T_0}^{T-1} t = \Omega(T).$$

Hence, in this case, we have

$$\limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T H_{\text{sum}}(t) = \infty.$$

Hence, from Eqns. (61) and (63), we conclude that, in either case, we have

$$\limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T H_{\text{sum}}(t) \geq 6. \quad (63)$$

Thus, In the case when $N = 2$, using the result of Proposition 2, the competitive ratio is lower bounded by

$$\eta \geq \frac{6}{2^2} = 1.5.$$

C. Proof of Theorem 8

Proof: Let $i^* = \arg \min_i p_i$. Now, under the action of any arbitrary policy π , at any slot $t \geq k$ and for all $k \geq 1$, we have

$$\mathbb{P}^\pi(\max_i h_i(t) \geq k) \geq \mathbb{P}^\pi(h_{i^*}(t) \geq k) \stackrel{(a)}{\geq} (1 - p_{\min})^k, \quad (64)$$

where the inequality (a) follows from the fact that consecutive k erasures just prior to time t for the i^* th user (which occurs with probability $(1 - p_{\min})^k$) ensures that the age of the user at time t is at least k .

Next, we analyze the large-deviation exponent under the action of the CMA policy. Using the union bound, we have

$$\mathbb{P}(\max_i h_i(t) \geq k) = \mathbb{P}(\bigcup_{i=1}^N h_i(t) \geq k) \leq \sum_{i=1}^N \mathbb{P}(h_i(t) \geq k). \quad (65)$$

Now, for any user i , the event $h_i(t) \geq k$ occurs if and only if it has been at least k slots since the i th user received a packet successfully before time t . Define $p_{\max} \equiv \max_i p_i$ and $p_{\min} \equiv \min_i p_i$. Since the CMA policy transmits other users successfully exactly once between two consecutive successful transmission to the user i (due to its round-robin nature), it follows that, during the last k slots prior to time t , at most $N-1$ users have successfully received a packet. Thus, we have the following bound:

$$\begin{aligned} \mathbb{P}(h_i(t) \geq k) &\leq \sum_{j=0}^{N-1} \binom{k}{j} p_{\max}^j (1 - p_{\min})^{k-j} \\ &\leq \binom{k}{N-1} (1 - p_{\min})^k \times \\ &\quad \frac{1 - p_{\min}}{(p_{\max} + p_{\min}) - 1} \left(\left(\frac{p_{\max}}{1 - p_{\min}} \right)^N - 1 \right) \\ &\leq c'(N, \mathbf{p}) k^N (1 - p_{\min})^k, \end{aligned}$$

where we have used the bound $\binom{k}{N-1} \leq \frac{k^N}{(N-1)!}$ and defined $c'(N, \mathbf{p}) \equiv \frac{1 - p_{\min}}{(N-1)! (p_{\max} + p_{\min}) - 1} \left(\left(\frac{p_{\max}}{1 - p_{\min}} \right)^N - 1 \right)^2$. Combining the above bound with Eqns. (65) and (64), we conclude that the CMA policy is optimal for the problem (51) and

$$-\lim_{k \rightarrow \infty} \lim_{t \rightarrow \infty} \frac{1}{k} \log \mathbb{P}^{\text{CMA}}(\max_i h_i(t) \geq k) = -\log(1 - \min_{i=1}^N p_i).$$

REFERENCES

- [1] A. Srivastava, A. Sinha, and K. Jagannathan, "On minimizing the maximum age-of-information for wireless erasure channels," in *2019 International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOPT)*, 2019, pp. 1–6.
- [2] S. Banerjee, R. Bhattacharjee, and A. Sinha, "Fundamental limits of age-of-information in stationary and non-stationary environments," in *2020 IEEE International Symposium on Information Theory (ISIT)*, 2020, pp. 1741–1746.
- [3] R. Bhattacharjee and A. Sinha, "Competitive algorithms for minimizing the maximum age-of-information," *SIGMETRICS Perform. Eval. Rev.*, vol. 48, no. 2, p. 6–8, Nov. 2020. [Online]. Available: <https://doi.org/10.1145/3439602.3439606>
- [4] L. Tassiulas and A. Ephremides, "Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks," *Automatic Control, IEEE Transactions on*, vol. 37, no. 12, pp. 1936–1948, 1992.

²In the case $p_{\max}/(1 - p_{\min}) = 1$, we take $c' \equiv N/(N-1)!$.

- [5] A. Mandelbaum and A. L. Stolyar, "Scheduling flexible servers with convex delay costs: Heavy-traffic optimality of the generalized $c\mu$ -rule," *Operations Research*, vol. 52, no. 6, pp. 836–855, 2004.
- [6] A. Sinha and E. Modiano, "Optimal control for generalized network-flow problems," *IEEE/ACM Transactions on Networking*, vol. 26, no. 1, pp. 506–519, Feb 2018.
- [7] M. J. Neely, "Stochastic network optimization with application to communication and queueing systems," *Synthesis Lectures on Communication Networks*, vol. 3, no. 1, pp. 1–211, 2010.
- [8] U. C. Kozat, I. Koutsopoulos, and L. Tassiulas, "A framework for cross-layer design of energy-efficient communication with qos provisioning in multi-hop wireless networks," in *IEEE INFOCOM 2004*, vol. 2. IEEE, 2004, pp. 1446–1456.
- [9] A. Gurijala and C. Molina, "Defining and monitoring qos metrics in the next generation wireless networks," in *2004 IEE Telecommunications Quality of Services: The Business of Success QoS 2004*, March 2004, pp. 37–42.
- [10] N. Banović-Ćurguz and D. Ilišević, "Mapping of qos/qoe in 5g networks," in *2019 42nd International Convention on Information and Communication Technology, Electronics and Microelectronics (MIPRO)*, May 2019, pp. 404–408.
- [11] S. Kaul, R. Yates, and M. Gruteser, "Real-time status: How often should one update?" in *INFOCOM, 2012 Proceedings IEEE*. IEEE, 2012, pp. 2731–2735.
- [12] 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.
- [13] S. Farazi, A. G. Klein, and D. R. Brown III, "Fundamental bounds on the age of information in multi-hop global status update networks," *Journal of Communications and Networks*, vol. 21, no. 3, pp. 268–279, 2019.
- [14] M. Grossglauser and D. N. Tse, "Mobility increases the capacity of ad hoc wireless networks," *IEEE/ACM Transactions on Networking (ToN)*, vol. 10, no. 4, pp. 477–486, 2002.
- [15] A. E. Gamal, J. Mammen, B. Prabhakar, and D. Shah, "Throughput-delay trade-off in wireless networks," in *IEEE INFOCOM 2004*, vol. 1. IEEE, 2004.
- [16] M. J. Neely and E. Modiano, "Capacity and delay tradeoffs for ad hoc mobile networks," *IEEE Transactions on Information Theory*, vol. 51, no. 6, pp. 1917–1937, 2005.
- [17] B. Baynat, R.-M. Indre, N. Nya, P. Olivier, and A. Simonian, "Impact of mobility in dense lte-a networks with small cells," in *2015 IEEE 81st Vehicular Technology Conference (VTC Spring)*. IEEE, 2015, pp. 1–5.
- [18] E. Anton, U. Ayesta, and F. Simatos, "On the impact of mobility in cellular networks," in *2019 International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOPT)*. IEEE, 2019, pp. 1–8.
- [19] A. Baiocchi, F. Sestini, and F. D. Priscoli, "Effects of user mobility on the capacity of a cdma cellular network," *European transactions on telecommunications*, vol. 7, no. 4, pp. 305–314, 1996.
- [20] W.-J. Hsu, T. Spyropoulos, K. Psounis, and A. Helmy, "Modeling time-variant user mobility in wireless mobile networks," in *IEEE INFOCOM 2007-26th IEEE International Conference on Computer Communications*. IEEE, 2007, pp. 758–766.
- [21] A. Ghazal, Y. Yuan, C. Wang, Y. Zhang, Q. Yao, H. Zhou, and W. Duan, "A non-stationary imt-advanced mimo channel model for high-mobility wireless communication systems," *IEEE Transactions on Wireless Communications*, vol. 16, no. 4, pp. 2057–2068, April 2017.
- [22] S. Wu, C.-X. Wang, M. M. Alwakeel, X. You *et al.*, "A general 3-d non-stationary 5g wireless channel model," *IEEE Transactions on Communications*, vol. 66, no. 7, pp. 3065–3078, 2017.
- [23] J. Bian, J. Sun, C. Wang, R. Feng, J. Huang, Y. Yang, and M. Zhang, "A winner+ based 3-d non-stationary wideband mimo channel model," *IEEE Transactions on Wireless Communications*, vol. 17, no. 3, pp. 1755–1767, March 2018.
- [24] A. Ghazal, C. Wang, Y. Liu, P. Fan, and M. K. Chahine, "A generic non-stationary mimo channel model for different high-speed train scenarios," in *2015 IEEE/CIC International Conference on Communications in China (ICC)*, Nov 2015, pp. 1–6.
- [25] R. Poisel, *Modern communications jamming principles and techniques*. Artech House, 2011.
- [26] A. Mpitzopoulos, D. Gavalas, C. Konstantopoulos, and G. Pantziou, "A survey on jamming attacks and countermeasures in wsns," *IEEE Communications Surveys & Tutorials*, vol. 11, no. 4, pp. 42–56, 2009.
- [27] M. Andrews and L. Zhang, "Routing and scheduling in multihop wireless networks with time-varying channels," *ACM Transactions on Algorithms (TALG)*, vol. 3, no. 3, pp. 33–es, 2007.
- [28] R. Bassily and A. Smith, "Causal erasure channels," in *Proceedings of the twenty-fifth annual ACM-SIAM symposium on Discrete algorithms*. SIAM, 2014, pp. 1844–1857.
- [29] 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 (TON)*, vol. 26, no. 6, pp. 2637–2650, 2018.
- [30] I. Kadota, M. S. Rahman, and E. Modiano, "Age of information in wireless networks: From theory to implementation," in *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking*, ser. MobiCom '20. New York, NY, USA: Association for Computing Machinery, 2020. [Online]. Available: <https://doi.org/10.1145/3372224.3418171>
- [31] I. Kadota, A. Sinha, and E. Modiano, "Optimizing age of information in wireless networks with throughput constraints," in *IEEE INFOCOM 2018-IEEE Conference on Computer Communications*. IEEE, 2018, pp. 1844–1852.
- [32] R. Talak, S. Karaman, and E. Modiano, "Minimizing age-of-information in multi-hop wireless networks," in *2017 55th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*. IEEE, 2017, pp. 486–493.
- [33] V. Tripathi, R. Talak, and E. Modiano, "Age optimal information gathering and dissemination on graphs," in *IEEE INFOCOM 2019-IEEE Conference on Computer Communications*. IEEE, 2019, pp. 2422–2430.
- [34] I. Kadota, A. Sinha, and E. Modiano, "Scheduling algorithms for optimizing age of information in wireless networks with throughput constraints," *IEEE/ACM Transactions on Networking*, 2019.
- [35] P. Popovski, C. Stefanović, J. J. Nielsen, E. de Carvalho, M. Angjelichinoski, K. F. Trillingsgaard, and A. Bana, "Wireless access in ultra-reliable low-latency communication (urllc)," *IEEE Transactions on Communications*, vol. 67, no. 8, pp. 5783–5801, 2019.
- [36] M. Costa, M. Codreanu, and A. Ephremides, "On the age of information in status update systems with packet management," *IEEE Transactions on Information Theory*, vol. 62, no. 4, pp. 1897–1910, 2016.
- [37] A. Javani, M. Zorghi, and Z. Wang, "Age of information in multiple sensing," in *2020 Information Theory and Applications Workshop (ITA)*. IEEE, 2020, pp. 1–10.
- [38] S. Kaul, M. Gruteser, V. Rai, and J. Kenney, "Minimizing age of information in vehicular networks," in *2011 8th Annual IEEE Communications Society Conference on Sensor, Mesh and Ad Hoc Communications and Networks*. IEEE, 2011, pp. 350–358.
- [39] S. Borst and P. Whiting, "Dynamic rate control algorithms for hdr throughput optimization," in *Proceedings IEEE INFOCOM 2001. Conference on Computer Communications. Twentieth Annual Joint Conference of the IEEE Computer and Communications Society (Cat. No. 01CH37213)*, vol. 2. IEEE, 2001, pp. 976–985.
- [40] H. J. Kushner and P. A. Whiting, "Convergence of proportional-fair sharing algorithms under general conditions," *IEEE transactions on wireless communications*, vol. 3, no. 4, pp. 1250–1259, 2004.
- [41] J. M. Holtzman, "Asymptotic analysis of proportional fair algorithm," in *12th IEEE International Symposium on Personal, Indoor and Mobile Radio Communications. PIMRC 2001. Proceedings (Cat. No. 01TH8598)*, vol. 2. IEEE, 2001, pp. F–F.
- [42] A. Ganti, E. Modiano, and J. N. Tsitsiklis, "Optimal transmission scheduling in symmetric communication models with intermittent connectivity," *IEEE Transactions on Information Theory*, vol. 53, no. 3, pp. 998–1008, 2007.
- [43] L. Ying, R. Srikant, A. Eryilmaz, and G. E. Dullerud, "A large deviations analysis of scheduling in wireless networks," *IEEE Transactions on Information Theory*, vol. 52, no. 11, pp. 5088–5098, 2006.
- [44] A. Fiat and G. J. Woeginger, *Online algorithms: The state of the art*. Springer, 1998, vol. 1442.
- [45] S. Albers, *Competitive online algorithms*. Citeseer, 1996.
- [46] R. G. Gallager, *Discrete stochastic processes*. Springer Science & Business Media, 2012, vol. 321.
- [47] F. de Sousa Chaves, M. Abbas-Turki, H. Abou-Kandil, and J. M. T. Romano, "Transmission power control for opportunistic qos provision in wireless networks," *IEEE Transactions on Control Systems Technology*, vol. 21, no. 2, pp. 315–331, 2013.
- [48] D. Tse and P. Viswanath, *Fundamentals of wireless communication*. Cambridge university press, 2005.
- [49] X. Ge, J. Ye, Y. Yang, and Q. Li, "User mobility evaluation for 5g small cell networks based on individual mobility model," *IEEE Journal on Selected Areas in Communications*, vol. 34, no. 3, pp. 528–541, 2016.

- [50] I. F. Akyildiz, Y.-B. Lin, W.-R. Lai, and R.-J. Chen, "A new random walk model for pcs networks," *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 7, pp. 1254–1260, 2000.
- [51] D. B. Johnson and D. A. Maltz, "Dynamic source routing in ad hoc wireless networks," in *Mobile computing*. Springer, 1996, pp. 153–181.
- [52] F. Bai and A. Helmy, "A survey of mobility models," *Wireless Adhoc Networks. University of Southern California, USA*, vol. 206, p. 147, 2004.
- [53] M. Chiang, P. Hande, and T. Lan, *Power control in wireless cellular networks*. Now Publishers Inc, 2008.
- [54] A. Zappone, L. Sanguinetti, G. Bacci, E. Jorswieck, and M. Debbah, "Energy-efficient power control: A look at 5g wireless technologies," *IEEE Transactions on Signal Processing*, vol. 64, no. 7, pp. 1668–1683, 2015.
- [55] D. P. Bertsekas, *Dynamic Programming and Optimal Control, Vol. II*, 3rd ed. Athena Scientific, 2007.
- [56] N. Ferns, P. Panangaden, and D. Precup, "Metrics for markov decision processes with infinite state spaces," in *Proceedings of the Twenty-First Conference on Uncertainty in Artificial Intelligence*, ser. UAI'05. Arlington, Virginia, USA: AUAI Press, 2005, p. 201–208.
- [57] D. Williams, *Probability with martingales*. Cambridge university press, 1991.
- [58] M. J. Neely, "Dynamic power allocation and routing for satellite and wireless networks with time varying channels," Ph.D. dissertation, Massachusetts Institute of Technology, 2003.
- [59] R. Durrett, *Probability: theory and examples*. Cambridge university press, 2010.
- [60] L. Lovász *et al.*, "Random walks on graphs: A survey," *Combinatorics, Paul erdos is eighty*, vol. 2, no. 1, pp. 1–46, 1993.
- [61] B. Buyukates and S. Ulukus, "Age of information with gilbert-elliott servers and samplers," in *2020 54th Annual Conference on Information Sciences and Systems (CISS)*, 2020, pp. 1–6.
- [62] D. P. Bertsekas, *Dynamic programming and optimal control*. Athena scientific Belmont, MA, 2005, vol. 2, no. 3.
- [63] I. Rhee, M. Shin, S. Hong, K. Lee, S. J. Kim, and S. Chong, "On the levy-walk nature of human mobility," *IEEE/ACM transactions on networking*, vol. 19, no. 3, pp. 630–643, 2011.
- [64] S. Sun, T. S. Rappaport, S. Rangan, T. A. Thomas, A. Ghosh, I. Z. Kovacs, I. Rodriguez, O. Koymen, A. Partyka, and J. Jarvelainen, "Propagation path loss models for 5g urban micro-and macro-cellular scenarios," in *2016 IEEE 83rd Vehicular Technology Conference (VTC Spring)*. IEEE, 2016, pp. 1–6.
- [65] A. Mandal, R. Bhattacharjee, and A. Sinha, "Optimizing age-of-information in adversarial environments with channel state information," in *2022 14th International Conference on COMMunication Systems NETWORKS (COMSNETS)*, 2022, pp. 522–530.
- [66] R. G. Gallager, *Stochastic processes: theory for applications*. Cambridge University Press, 2013.



Rajarshi Bhattacharjee arjshi Bhattacharjee-ajarshi Bhattacharjee received a B.E. degree in mechanical engineering from Jadavpur University, Kolkata, India, and an M.Tech. degree in computer science from Indian Statistical Institute, Kolkata, India. He is currently pursuing a Ph.D. degree in computer science at the University of Massachusetts Amherst, USA. From 2019 to 2020, he was a research assistant at the Indian Institute of Technology Madras, India, under Prof. Abhishek Sinha, where he worked on online caching algorithms and on optimizing age-of-information in communication networks. His areas of interest include randomized algorithms, numerical linear algebra, online algorithms, and machine learning.



Abhishek Sinha is currently a faculty member in the School of Technology and Computer Science at the Tata Institute of Fundamental Research, Mumbai, India. Prior to joining TIFR, he had been with the Dept. of Electrical Engineering at the Indian Institute of Technology Madras as an Assistant Professor. He received his Ph.D. from the Massachusetts Institute of Technology, where he was affiliated with the Laboratory for Information and Decision Systems. Thereafter, Abhishek worked as a senior engineer at Qualcomm Research, San Diego, in the 5G standardization group. He obtained his M.E. degree in Telecommunication Engg. from the Indian Institute of Science, Bangalore, and B.E. degree in Electronics and Telecommunication Engg. from Jadavpur University, Kolkata, India. He is a recipient of the INSA Medal for Young Scientists (2021), Best Paper Awards in INFOCOM 2018 and MobiHoc 2016, and Jagadis Bose National Science Talent Search (JBNSTS) scholarship, Kolkata, India. His areas of interest include theoretical machine learning, networks, and information theory.