Scheduling to Minimize Age of Information with Multiple Sources

Kumar Saurav
School of Technology and Computer Science
Tata Institute of Fundamental Research
Mumbai, India.
kumar.saurav@tifr.res.in

Rahul Vaze
School of Technology and Computer Science
Tata Institute of Fundamental Research
Mumbai, India.
rahul.vaze@gmail.com

***Abstract*—Finding an optimal/near-optimal scheduling algorithm to minimize the age of information (AoI) in a multi-source G/G/1 system is well-known to be a hard problem. In this paper, we consider this problem for the non-preemptive setting, where an algorithm is free to choose which update to transmit, but an update under transmission is not allowed to be preempted. For this problem, we propose a novel randomized scheduling algorithm and show that its competitive ratio is at most 3 plus the maximum of the ratio of the variance and the mean of the inter-arrival time distribution of sources. Notably, the competitive ratio is independent of the number of sources, or their service time distributions. For several common inter-arrival time distributions such as exponential, uniform and Rayleigh, the competitive ratio is at most 4.**

***Index Terms*—age of information, scheduling, competitive ratio.**

I. Introduction

With the advent of modern applications such as remote gaming, smart and connected cars, IoT, smart homes etc., information timeliness has become an important performance metric. Information timeliness refers to quick and periodic dissemination of information, for example, in networked cars, critical safety information needs to be updated quickly and often enough. In recent times, several metrics have been proposed for information timeliness that include the age of information (AoI) [1], the age of incorrect information [2], the age of incorrect estimates [3]. Because of its simplicity and elegance, AoI has become the de facto first choice for analysis, where the *age* at time t is defined as the time elapsed since the last received update was generated, and the AoI is the average of age across time.

Considered Problem: In this paper, we consider a scheduling problem in a G/G/1 system, where updates from different sources arrive to a single queue. The inter-generation time of updates for source i is assumed to be random with distribution $\mathcal{G}_i$, while the transmission (service) time distribution of an update from source i is $\mathcal{D}_i$. At any time, only one update from any source can be under transmission (service), and each transmission by a source incurs a fixed energy/transmission cost. We consider a cost function that is a linear combination of the sum of the AoI across all sources and the total transmission cost. The objective is to find a non-preemptive causal scheduling policy that minimizes the cost function. Discussion on *preemptive* causal scheduling policies can be found in the full version of this paper [4].

We acknowledge support of the Department of Atomic Energy, Government of India, under project no. RTI4001.

Prior Work: Early research on AoI considered M/M/1 [1], [5] and D/G/1 [6], [7] systems where updates arrive from a single source, and analyzed the AoI for different scheduling policies such as FCFS, LCFS, LCFS with preemption etc. Later, [8], [9] considered M/M/1 systems where updates arrive from multiple sources, and characterized the mean, the distribution of the AoI or the related performance metrics for each of the sources, for fixed scheduling policies such as FCFS, LCFS, etc. Distributional properties of AoI when the energy used to transmit updates is sourced from renewable sources has also been analyzed in [10], and the references therein.

An alternate AoI research direction has been towards finding a scheduling policy that minimizes the weighted sum of the AoI of the sources. For a single source continuous-time generate-at-will model, where the source can generate the update at any time, and the transmission delay/service time distribution is arbitrary, [11] derived an optimal causal scheduling policy. With multiple sources, finding an optimal causal policy has remained elusive, and prior works have restricted their analysis to simpler settings, where the goal is to find causal scheduling policies for which the competitive ratio (ratio of the cost of the causal scheduling policy and an optimal offline policy that is aware of the input in advance) is bounded [12]–[15]. In this direction, [12] considered a M/G/1 system with multiple sources, and using the memoryless property of the exponential update inter-generation time distribution, showed that a randomized policy has competitive ratio 3. For simpler settings of slotted-time model, [13]–[15] derived competitive ratio guarantees for scheduling policies, as discussed next.

In slotted-time model, in each slot, a centralized scheduler schedules at most one source to transmit its update to the monitor. If source i is chosen for transmission, the transmitted update is either successfully received at the monitor in the same slot with probability p_i, or gets lost with probability $1 - p_i$, independent across slots. When a new update is generated at each source in every slot, [13], [14] showed that for a simple randomized policy, the competitive ratio is at most 2. Similar results have been extended for richer models, e.g.

ISBN 978-3-903176-49-2 © 2022 IFIP

fading channel [16], multiple access channel [17], etc. When the updates are generated in each slot with a fixed probability, [15] derived a policy with competitive ratio 4.

Note that for AoI minimization, a source needs to transmit frequently. However, in real-world systems, transmissions incur cost (energy). Thus, there is an inherent tradeoff between minimizing the AoI and the transmission cost, that has largely been ignored in prior work such as [11]–[15]. In the work that consider AoI minimization with transmission cost, the AoI-transmission cost tradeoff has been modeled either by considering an objective function that is a combination of the AoI of the sources and the transmission cost [18], [19], or by considering an additional constraint on the transmission cost [16]. In either case, the model considered for update generation and service times have been far simpler (compared to the G/G/1 model in this paper) with slotted-time [19], single source [16], [18], zero service time [18], etc.

It is worth noting that there is a large body of work on optimal or near optimal algorithms for non-AoI scheduling problems such as the flow-time minimization [20], the completion time problem [21], the makespan problem [22], and the co-flow problem [23]. In terms of technical difficulty, AoI scheduling problems are fundamentally different than the non-AoI scheduling problems, primarily for two reasons: (i) for AoI scheduling problems, all the updates for each of the sources need not be transmitted/serviced (the problem is combinatorial in nature and need to choose the subset of updates to transmit), and (ii) AoI depends on the difference between the successive time instants when updates are transmitted (the scheduling decisions across time are correlated).

Our Contributions: To solve the considered problem, we propose a novel non-preemptive scheduling policy that consists of two randomized subroutines: (a) SR-PMS, which at each source ℓ, retains/discards each arriving update with a fixed probability, and (b) SR-NSS, that among all the sources, decides which source gets to transmit its latest (retained) update. Critically, if the selected source ℓ has no (new) update to transmit, the policy idles for a random period of time that is drawn independently from the service time distribution $\mathcal{D}_\ell$.

For the proposed policy, we show that its competitive ratio is upper bounded by $\max\{4, \ \ 3 + \max_\ell\{\sigma_\ell^2/\mu_\ell^2\}\}$, where σ_ℓ^2 and μ_ℓ are respectively the variance and the mean of the update inter-generation time for source ℓ. Notably, for many of the 'nice' distributions, e.g. exponential, uniform, Rayleigh, etc., the competitive ratio is upper bounded by 4. It is worth noting that the competitive ratio upper bound is independent of the service time distributions $\mathcal{D}_\ell$'s. As far as we know, this is first such result in the area of AoI scheduling, with general inter-generation time and service time distributions.

We also construct a 'tight' example with a single source to show that the dependence of the competitive ratio of the proposed policy on σ^2/μ^2 is unavoidable.

II. System Model

Consider a system consisting of N sources, where updates (henceforth, packets) are generated at each source ℓ intermittently, and the inter-generation time between the i^{th} and the $i+1^{st}$ packet of source ℓ is $X_{\ell i}$. We assume that $X_{\ell i}$ is independent and identically distributed according to some distribution $\mathcal{G}_\ell$, with mean $\mu_\ell < \infty$ and variance σ_ℓ^2. There is a single monitor, and all the N sources wish to send their updates to the monitor as soon as possible. At any time, at most one source can transmit its packet to the monitor, and packet i's transmission by source ℓ takes $d_{\ell i}$ time units (called transmission time) to complete (received at the monitor). We assume that the transmission time $d_{\ell i}$, for each packet i of source ℓ is independent and identically distributed according to some general distribution $\mathcal{D}_\ell$, with mean $\gamma_\ell < \infty$, independent of everything else.

Remark 1: For different sources, $\mathcal{G}_\ell$'s and $\mathcal{D}_\ell$'s may belong to different family of distributions. For example, we may have $\mathcal{G}_1$ a uniform distribution, $\mathcal{G}_2$ an exponential distribution, $\mathcal{D}_1$ a Rayleigh distribution, and $\mathcal{D}_2$ a log-normal distribution.

There is a single centralized scheduler, and at any time t, the scheduler has causal information of all the sources, and gets to decide which source should transmit when the channel becomes free (previous transmission is completed). With the non-preemptive restriction, packets can be transmitted in any order or discarded if their transmission has never started, however, a packet under transmission cannot be preempted or discarded. We relax the non-preemptive restriction in the full version of this paper [4].

Definition 1: At any time t, the channel is said to be *busy*, if a packet is already under transmission by some source. Otherwise, the channel is *free*. A source can begin transmission of a packet only when the channel is free.

At any time $t \geq 0$, age of source ℓ at the monitor is $\Delta_\ell(t) = t - \lambda_\ell(t)$, where $\lambda_\ell(t)$ denotes the generation time of the latest packet of source ℓ that has been received at the monitor (i.e., completely transmitted by source ℓ) until time t. Thus, the *age of information (AoI)* $\Delta_\ell^{av}(t)$ of source ℓ until time t is

$$\Delta_\ell^{av}(t) = \frac{1}{t}\int_0^t \Delta_\ell(i)di. \tag{1}$$

Each time source ℓ transmits a packet, it incurs a transmission cost of c_ℓ units, where $c_\ell \geq 0$ is finite constant, and includes the cost for channel usage, as well as the energy required to transmit the packet. Hence, the average transmission cost incurred by source ℓ until time t is given by

$$C_\ell^{av}(t) = \frac{c_\ell U_\ell(t)}{t}, \tag{2}$$

where, $U_\ell(t)$ denotes the number of packets transmitted by source ℓ until time t (including the packet under transmission).

Definition 2: A *causal scheduling policy* (in short, causal policy) refers to a centralized algorithm (employed by the scheduler) that at each time t when the channel is free, based only on the causal information of all the sources available at time t, schedules at most one source to transmit its packet. In this paper, all policies are non-preemptive.

The objective is to find a causal policy (Definition 2) that minimizes a linear combination of the AoI and the average transmission cost of all the sources (called the *weighted sum cost*). Formally, the weighted sum cost of policy π is

$$\Gamma(\pi) = \lim_{t\to\infty} \frac{1}{N} \sum_{\ell=1}^{N} (C^{av}_{\ell,\pi}(t) + \rho_\ell \Delta^{av}_{\ell,\pi}(t)), \tag{3}$$

where $\rho_\ell > 0$ is a finite constant (weight parameter) corresponding to the AoI of source ℓ, and $C^{av}_{\ell,\pi}(t)$ and $\Delta^{av}_{\ell,\pi}(t)$ respectively denote the average transmission cost and the AoI of source ℓ until time t, under policy π. The objective is formulated as the optimization problem

$$\min_{\pi\in\Pi} \ \Gamma(\pi), \tag{4}$$

where Π is the set of all causal policies π.

Remark 2: In (4), we consider the cost function (3) to be a linear combination of the AoI and the average transmission cost, similar to [18], [19]. An alternate formulation for (4) could be in which the cost function is AoI, and there is an additional constraint on the average transmission cost for each source. However, as is well understood from the theory of penalty functions [24], for an appropriate choice of the weights ρ_ℓ's (Lagrangian multipliers that can be found in an iterative way), the solution of the considered formulation corresponds to that of the alternate formulation.

Remark 3: When the cost per transmission $c_\ell = 0$ for each source ℓ, the objective (4) simplifies to an AoI minimization problem with multiple sources, where the packet inter-generation times, as well as the transmission time for packets may follow any general distribution. In past, such an AoI minimization problem has been considered under restricted settings, such as with single source [11], [18], generate-at-will model [11], [13], discrete-time model, [13], [15], zero transmission time [18], etc.

In this paper, to quantify the performance of a causal policy, we compare it against the performance of an optimal offline policy using the metric of competitive ratio (Definition 4).

Definition 3: A policy $\pi^\star_{OF}$ is called an optimal offline non-preemptive policy, if its weighted sum cost $\Gamma(\pi^\star_{OF})$ (3) is minimum among all non-preemptive policies that know the generation time of all the packets (at each source) in advance. For the reasons discussed in Remark 4, we assume that the transmission time for each packet is realized once the packet is transmitted, and it is not known to $\pi^\star_{OF}$ non-causally.

Definition 4: For a causal non-preemptive policy π, its competitive ratio CR_π is defined as the ratio of the expected weighted sum cost (3) for policy π and the expected weighted sum cost (3) for an optimal offline non-preemptive policy $\pi^\star_{OF}$ (Definition 3), where the expectation $\mathbb{E}[\cdot]$ is jointly with respect to the distributions $\mathcal{G}_\ell$ and $\mathcal{D}_\ell$ (for each source ℓ), and the corresponding scheduling policy (i.e., π or $\pi^\star_{OF}$). Mathematically, $\text{CR}_\pi = \mathbb{E}[\Gamma(\pi)]/\mathbb{E}[\Gamma(\pi^\star_{OF})]$.

Remark 4: Compared to a causal policy, an offline policy has access to more information, and that typically allows lower bounding of the cost of an optimal causal policy. However, more the extra information an offline policy has, larger is the gap between the performance of the optimal offline policy and a causal policy. Since the goal is to find an optimal/near-optimal causal policy, we ideally want an offline policy to have as little extra information as possible over the causal policy, while allowing analytical tractability. For the considered problem, we let an offline policy only know the inter-generation time of updates non-causally, but the transmission time/delay experienced by any packet that is transmitted is revealed to it causally.

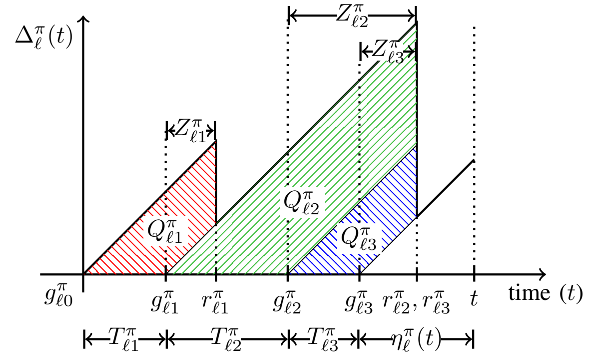

Fig. 1: Sample age plot of source ℓ under under policy π. Since update ℓ^π_3 is transmitted before update ℓ^π_2, we have $r^\pi_{\ell 2} = r^\pi_{\ell 3}$.

III. Preliminaries

In this section, we define certain quantities that determine the weighted sum cost (3) of any causal/offline policy. Subsequently, we derive a general expression for the weighted sum cost (3) of a policy in terms of these quantities.

For any policy π (causal or offline), consider a subset of packets $\mathcal{F}^\pi_\ell$ generated at source ℓ, such that the packets of source ℓ transmitted by π lie in $\mathcal{F}^\pi_\ell$. Let the packets in $\mathcal{F}^\pi_\ell$ be indexed as $\ell^\pi_1, \ell^\pi_2, \ldots$ in increasing order of their generation times, with the generation time of packet ℓ^π_i being $g^\pi_{\ell i}$.

Remark 5: We allow $\mathcal{F}^\pi_\ell$'s to be a *superset* of the set of transmitted packets to aid the analysis in Appendix C.

For each packet ℓ^π_i, define $s^\pi_{\ell i}$ and $r^\pi_{\ell i}$, respectively, as the earliest time instant when the transmission of a packet ℓ^π_j (where $j \geq i$) begins and completes, under policy π. By definition, $g^\pi_{\ell i} \leq s^\pi_{\ell i} \leq r^\pi_{\ell i}$. Also, as shown in Figure 1, $s^\pi_{\ell i}$'s and $r^\pi_{\ell i}$'s may be equal for successive packets in $\mathcal{F}^\pi_\ell$.

Further, let $Z^\pi_{\ell i} = r^\pi_{\ell i} - g^\pi_{\ell i}$ denote the difference between the time instant when packet ℓ^π_i is generated, and the earliest time instant when the transmission of a packet ℓ^π_j $(j \geq i)$ is completed. By definition,

$$Z^\pi_{\ell i} = (s^\pi_{\ell i} - g^\pi_{\ell i}) + (r^\pi_{\ell i} - s^\pi_{\ell i}) = w^\pi_{\ell i} + d_{\ell i}, \tag{5}$$

where $w^\pi_{\ell i} = s^\pi_{\ell i} - g_{\ell i}$ (which we call as the waiting time for packet ℓ^π_i), and $d_{\ell i} = r^\pi_{\ell i} - s^\pi_{\ell i} \sim \mathcal{D}_\ell$ is the transmission time of the packet whose transmission began at time $s^\pi_{\ell i}$. Note that $d_{\ell i}$ is independent of policy π.

Next, with respect to $\mathcal{F}^\pi_\ell$, we define a **period** as the time interval between the generation time of two consecutive packets in $\mathcal{F}^\pi_\ell$. Thus, the interval $\mathcal{P}^\pi_{\ell i} = (g^\pi_{\ell(i-1)}, g^\pi_{\ell i}]$ represents the i^{th} period with respect to $\mathcal{F}^\pi_\ell$, and the length (duration) of period $\mathcal{P}^\pi_{\ell i}$ is $T^\pi_{\ell i} = g^\pi_{\ell i} - g^\pi_{\ell(i-1)}$. Note that $T^\pi_{\ell i}$ depends on the choice of $\mathcal{F}_\ell^\pi$, which further depends on both the packet generation process and the packets transmitted by policy π.

Remark 6: Without loss of generality, we assume that the age of all the sources at time $t = 0$ is 0 (i.e., $\Delta_\ell(0) = 0$, $\forall \ell \in \{1, \cdots, N\}$). In (4), since we are interested in the weighted sum cost over infinite time horizon, this assumption does not affect the final solution of (4), but simplifies the analysis, and for each source ℓ, allows us to assume $g_{\ell 0}^\pi = 0$ (i.e., the first period of every source starts at time 0).

Remark 7: As shown in Figure 1, when $g_{\ell 0}^\pi = 0$, with respect to $\mathcal{F}_\ell^\pi$, any time t can be written as $t = \sum_{i=1}^{R_\ell^\pi(t)} T_{\ell i}^\pi + \eta_\ell^\pi(t)$, where $R_\ell^\pi(t)$ denotes the number of packets in $\mathcal{F}_\ell^\pi$ generated until time t, while $\eta_\ell^\pi(t) = t - g_{\ell R_\ell^\pi(t)}^\pi$ denotes the length of the ongoing period of source ℓ at time t.

Assumption 1: All system parameters are finite such that there exists some causal/offline policy π for which the weighted sum cost $\Gamma(\pi)$ (3) is finite with probability 1.

Lemma 1: Under Assumption 1, for minimizing the weighted sum cost $\Gamma(\pi)$ (3), it is sufficient to consider only those policies π (causal or offline) and subset of packets $\mathcal{F}_\ell^\pi$, for which $T_{\ell i}^\pi$ ($\forall i$) and $\eta_\ell^\pi(t)$ are finite with probability 1.

Proof: Since there exists a policy with finite weighted sum cost (3) (Assumption 1), for the weighted sum cost minimization problem (4), it is sufficient to consider only those policies for which the weighted sum cost (3) is finite. Now note that by definition, source ℓ does not transmit any packet generated in intervals of lengths $T_{\ell i}^\pi$'s and $\eta_\ell^\pi(t)$ (under policy π). Therefore, if $T_{\ell i}^\pi$ or $\eta_\ell^\pi(t)$ is infinite, the age of the source grows to infinity, making the weighted sum cost (3) infinite. Hence, it is sufficient to consider policies π and subset of packets $\mathcal{F}_\ell^\pi$ for which $T_{\ell i}^\pi$'s and $\eta_\ell^\pi(t)$ are finite. ■

In view of Lemma 1, we restrict our attention to policies π and corresponding subsets $\mathcal{F}_\ell^\pi$, for which $T_{\ell i}^\pi$ ($\forall i$) and $\eta_\ell^\pi(t)$ are finite with probability 1, and respectively define Π_S and Π_{OF} as the set of all such causal and non-causal (offline) policies.

Remark 8: From Remark 7, we get $\sum_{i=1}^{R_\ell^\pi(t)} T_{\ell i}^\pi = t - \eta_\ell^\pi(t)$, where $\eta_\ell^\pi(t)$ is finite for all policies $\pi \in \Pi_S \cup \Pi_{OF}$ (by definition of Π_S and Π_{OF}). Therefore, for $\pi \in \Pi_S \cup \Pi_{OF}$ (i.e., the policies of interest), as $t \to \infty$, we have $\sum_{i=1}^{R_\ell^\pi(t)} T_{\ell i}^\pi = t - \eta_\ell^\pi(t) \approx t$. Hence for simplicity, in the rest of this paper, when $t \to \infty$, we consider $\sum_{i=1}^{R_\ell^\pi(t)} T_{\ell i}^\pi = t$, i.e., any large time t is equal to the sum of the length of periods of source ℓ (for any $\ell \in \{1, \cdots, N\}$) until time t.

Next, for all policies $\pi \in \Pi_S \cup \Pi_{OF}$, Lemma 2 provides a general expression for the weighted sum cost (3) in terms of the quantities defined so far.

Lemma 2: For any policy $\pi \in \Pi_S \cup \Pi_{OF}$, and the subset of packets $\mathcal{F}_\ell^\pi$, the weighted sum cost is $\Gamma(\pi)$

$$= \lim_{t\to\infty} \frac{1}{N} \sum_{\ell=1}^{N} \left(\frac{\rho_\ell \sum_{i=1}^{R_\ell^\pi(t)} (T_{\ell i}^\pi)^2}{2\sum_{i=1}^{R_\ell^\pi(t)} T_{\ell i}^\pi} + \frac{\rho_\ell \sum_{i=1}^{R_\ell^\pi(t)} T_{\ell i}^\pi w_{\ell i}^\pi}{\sum_{i=1}^{R_\ell^\pi(t)} T_{\ell i}^\pi} + \frac{\rho_\ell \sum_{i=1}^{R_\ell^\pi(t)} T_{\ell i}^\pi d_{\ell i}}{\sum_{i=1}^{R_\ell^\pi(t)} T_{\ell i}^\pi} + \frac{c_\ell U_\ell^\pi(t)}{\sum_{i=1}^{R_\ell^\pi(t)} T_{\ell i}^\pi} \right), \quad (6)$$

where $U_\ell^\pi(t) \leq R_\ell^\pi(t)$ denotes the number of packets transmitted by π (including that under transmission), until time t. Further, $U_\ell^\pi(t)$'s satisfy

$$\lim_{t\to\infty} \sum_{\ell=1}^{N} \frac{\gamma_\ell \cdot U_\ell^\pi(t)}{\sum_{i=1}^{R_\ell^\pi(t)} T_{\ell i}^\pi} \leq 1, \quad \text{(with probability 1)}. \quad (7)$$

Proof: See Appendix A. ■

Using this formalism, we next propose and analyze a novel randomized scheduling policy to solve problem (4).

IV. Stationary Randomized Causal Policy

Definition 5: At any time t, a packet ℓ_i^π is called **fresh**, if its generation time $g_{\ell i}^\pi$ is greater than the generation time of the latest generated packet of source ℓ that has been completely transmitted until time t, i.e., $g_{\ell i}^\pi > \lambda_\ell^\pi(t)$.

Consider a stationary randomized policy π_{sr} (Algorithm 1), that consists of following two subroutines: (i) SR-PMS (for packet management), and (ii) SR-NSS (for scheduling sources for transmission). At any time t, (i) if a packet is generated at source ℓ, then SR-PMS marks the generated packet with probability p_ℓ (and discards it otherwise), and (ii) if the channel becomes free (latest transmission is completed/received at the monitor), SR-NSS chooses source ℓ (among all the N sources) for transmission with probability

$$\hat{p}_\ell = \frac{(p_\ell/\mu_\ell)}{\sum_{i=1}^{N}(p_i/\mu_i)}, \quad (8)$$

where the probability vector $[p_1, p_2, ..., p_N]$ is obtained by solving the following *convex optimization problem*

$$\arg\min_{[p_1,...,p_N]} \sum_{\ell=1}^{N} \left(\frac{2\rho_\ell \mu_\ell}{p_\ell} + \frac{c_\ell p_\ell}{\mu_\ell} \right), \quad (9)$$

$$s.t. \quad \sum_{\ell=1}^{N} \frac{p_\ell \gamma_\ell}{\mu_\ell} \leq 1, \quad (10)$$

$$p_\ell \in [0, 1], \quad \forall \ell \in \{1, \cdots, N\}. \quad (11)$$

If the chosen source ℓ has at least one fresh marked packet (Definition 5), then its latest generated marked packet is transmitted. Else, SR-NSS waits for a random time duration, independently sampled from distribution $\mathcal{D}_\ell$, before choosing a source again.

Remark 9: The intuition for choosing p_ℓ's as the solution of (9)—(11) is that in Lemma 4, we show that the expression (9), with some additional constant terms, is an upper bound on the expected weighted sum cost for the proposed policy $\mathbb{E}[\Gamma(\pi_{sr})]$. Thus the proposed policy is choosing p_ℓ's to minimize an upper bound on its weighted sum cost, under constraint (10) that follows from (7) (that relates the number of packets that each source may transmit). Note that (9)—(11) is convex, and can be easily solved using tools such as the CVX in Matlab.

Remark 10: Note that if a source ℓ is chosen by SR-NSS which does not have a fresh marked packet, then SR-NSS waits for a random amount of time as per distribution $\mathcal{D}_\ell$, instead of choosing a different source to transmit (π_{sr} is non-work conserving). This is primarily to simplify the theoretical analysis of π_{sr} (SR-PMS + SR-NSS) as when SR-NSS waits for random time with distribution $\mathcal{D}_\ell$ (instead of choosing a different source), the time duration between two successive instants when any source is chosen by SR-NSS to transmit, does not depend on the availability of marked packets at the sources (i.e. on when SR-PMS marks the packets at the sources). Moreover, as shown in [11], when there is a single source with random transmission times, a non-work conserving policy is optimal for minimizing AoI. In Section V, we also show that when the transmission costs are large, π_{sr} has lower weighted sum cost (3) compared to when its waiting times are identically 0. Thus, it is not trivially wasteful to wait.

Algorithm 1 Stationary randomized policy π_{sr}.

/ SR-Packet Management Subroutine (SR-PMS) */*
for each packet generated at source $\ell \in \{1, \cdots, N\}$ **do**
  mark the packet with probability p_ℓ, and discard it with probability $1 - p_\ell$;
end for
/ SR-Node Scheduling Subroutine (SR-NSS) */*
for time $t \geq 0$, if channel is free **do**
  among the N sources, choose source ℓ with probability $\hat{p}_\ell$ (8);
  if source ℓ has at least one fresh marked packet **then**
    transmit the latest marked packet of source ℓ;
  else
    wait for random time $d_\ell \sim \mathcal{D}_\ell$;
  end if
end for

The main result of this paper is as follows.

Theorem 1: The stationary randomized policy π_{sr} (Algorithm 1) has competitive ratio (against non-preemptive optimal offline policy $\pi^\star_{OF}$)

$$\textsc{cr}_{\pi_{sr}} \leq \max\{4, \;\; 3 + \max_\ell\{\sigma_\ell^2/\mu_\ell^2\}\}, \tag{12}$$

where σ_ℓ^2 and μ_ℓ respectively denote the variance and the mean of packet inter-generation times at source ℓ.

Notably, Theorem 1 shows that the competitive ratio of π_{sr} is independent of the transmission time distribution $\mathcal{D}_\ell$. Intuitively, this is because the optimal offline policy $\pi^\star_{OF}$ also does not know the realization of the transmission times of packets non-causally (Definition 3), and hence, the impact of random transmission time on $\pi^\star_{OF}$ is similar to that on π_{sr}.

However, the competitive ratio of π_{sr} depends on the distribution of the inter-generation time of packets. In particular, it depends on σ_ℓ^2/μ_ℓ^2, which for several common distributions, is upper bounded by a constant.

Corollary 1: When $\mathcal{G}_\ell$'s are exponential, uniform or Rayleigh, the competitive ratio of π_{sr} is at most 4.

Proof: For exponential, uniform and Rayleigh distributions, the variance σ_ℓ^2 and mean μ_ℓ depend on same parameter, such that $\sigma_\ell^2/\mu_\ell^2 \leq 1$. Hence, (12) implies that $\textsc{cr}_{\pi_{sr}} \leq 4$. ■

Note that the dependence of π_{sr}'s competitive ratio on σ_ℓ^2/μ_ℓ^2 is unavoidable. This is because $\pi^\star_{OF}$ knows the generation time of packets in advance, and for certain distributions, $\pi^\star_{OF}$ can use this information and minimize the variance of its period lengths to 0 (irrespective of σ_ℓ^2), whereas the period length (AoI) of π_{sr} always depends on σ_ℓ^2. Next, we make this concrete via constructing a tight example.

Example 1: Consider a system with a single source ($N = 1$), where the packet inter-generation time is distributed as

$$X = \begin{cases} \alpha, & \text{with probability } 0.5, \\ \epsilon & \text{with probability } 0.5, \end{cases} \tag{13}$$

and $\epsilon \to 0^+$, while α is a large (but finite) positive constant. Thus, the mean and the variance of the packet inter-generation time X are $\mu = (\alpha+\epsilon)/2 \approx \alpha/2$ and $\sigma^2 = (\alpha-\epsilon)^2/4 \approx \alpha^2/4$, respectively. Also, let the cost per transmission $c = 0$, and the transmission time $d_i = 0$ for each packet (hence the expected transmission time $\gamma = 0$). For this example, let $\mathcal{F}^\pi_\ell$ to be the set of packets that are transmitted by policy π (thus for each source ℓ, $T^\pi_{\ell i}$'s $\forall i$, are the inter-generation time of packets that are transmitted by policy π).

Consider a threshold policy π_{tp}, that on generation of a packet at time t, transmits it immediately if $t - \lambda(t) \geq \alpha$ (where $\lambda(t)$ denotes the generation time of the latest packet that got transmitted until time t), and discards it otherwise. Let g_i^{tp} denote the generation time of the i^{th} packet transmitted under policy π_{tp}. Then the period lengths under policy π_{tp} are $T_i^{tp} = g_i^{tp} - g_{i-1}^{tp} = m_i\alpha + n_i\epsilon$, where m_i and n_i denote the number of packets generated in the i^{th} period with inter-generation time α and ϵ, respectively. Since the threshold for transmission of packets is α, any packet i with inter-generation time $X_i = \alpha$ is always transmitted. Therefore, either $m_i = 0$ and $n_i = \alpha/\epsilon$, or $m_i = 1$ and $n_i < \alpha/\epsilon$.

However, since $\alpha/\epsilon \to \infty$ (because $\epsilon \to 0^+$), and the inter-generation time of packets take values ϵ or α, each with probability 0.5, the probability that $m_i = 0$ and $n_i = \alpha/\epsilon \to \infty$ is 0. Hence, with probability 1, $m_i = 1$, and n_i is finite. Also, $n_i < \infty$ implies that as $\epsilon \to 0^+$, $n_i\epsilon \to 0$ as well. Therefore, the period lengths $T_i^{tp} = m_i\alpha + n_i\epsilon \approx \alpha = 2\mu$ (a constant). Thus, substituting $N = 1$, $c = 0$, $T_i^{tp} = 2\mu$ and $w_i^{tp} = d_i = 0$ $\forall i$ in (6), we get $\Gamma(\pi_{tp}) \to \rho\mu$.

Next, consider the performance of the stationary randomized policy π_{sr} (Algorithm 1) for the same input. Since $\gamma = 0$, $c = 0$ and $N = 1$, from the convex optimization problem (9), it is immediate that $p = 1$ (i.e., π_{sr} marks every generated packet). Also, since the transmission time $d_i = 0$ for every packet, the channel is always free, and every marked packet gets immediately transmitted. Therefore, $T_i^{sr} = X_i$ and $w_i^{sr} = d_i = 0$. Further $\mathbb{E}[X_i] = \mu < \infty$, which implies that $T_i^{sr} = X_i$ is finite with probability 1, i.e., $\pi_{sr} \in \Pi_S$. Thus, substituting $N = 1$, $c = 0$, $T_i^{sr} = X_i$ and $w_i^{sr} = d_i = 0$ ($\forall i$) in (6), and using the renewal reward theorem [25], we get $\Gamma(\pi_{sr}) = \rho\mathbb{E}[X^2]/(2\mathbb{E}[X]) = (\rho/2)(\sigma^2 + \mu^2)/\mu$. Hence, $\frac{\Gamma(\pi_{sr})}{\Gamma(\pi_{tp})} = \frac{1}{2}\left(\frac{\sigma^2}{\mu^2} + 1\right)$. Since $\pi^\star_{OF}$ is at least as good as π_{tp}, we get that the competitive ratio of π_{sr} is proportional to σ^2/μ^2.

Next, we present the proof of Theorem 1 in two steps. First, we derive a lower bound on the weighted sum cost for an optimal offline policy $\pi^\star_{OF}$ (Definition 3), as follows.

Lemma 3: Let $h_\ell(t)$ denote the number of packets generated at source ℓ until time t, and an optimal offline policy $\pi^\star_{OF}$ transmits $U^\star_\ell(t)$ number of these packets. Then, the expected weighted sum cost for policy $\pi^\star_{OF}$ is

$$\mathbb{E}[\Gamma(\pi^\star_{OF})] \geq \frac{1}{N}\sum_{\ell=1}^{N}\left(\frac{\rho_\ell\mu_\ell}{2f^\star_\ell} + \rho_\ell\gamma_\ell + \frac{c_\ell f^\star_\ell}{\mu_\ell}\right), \tag{14}$$

where $f^\star_\ell = \lim_{t\to\infty} U^\star_\ell(t)/h_\ell(t) \in [0,1]$. Further, $\sum_{\ell=1}^{N}\gamma_\ell f^\star_\ell/\mu_\ell \leq 1$.

Next, we compute an upper bound on the expected weighted sum cost of policy π_{sr} (Algorithm 1), described as follows.

Lemma 4: The expected weighted sum cost for policy π_{sr} (Algorithm 1) is $\mathbb{E}[\Gamma(\pi_{sr})]$

$$\leq \frac{1}{N}\sum_{\ell=1}^{N}\left(\frac{2\rho_\ell\mu_\ell}{p_\ell} + \frac{c_\ell p_\ell}{\mu_\ell} + \rho_\ell\gamma_\ell - \frac{\rho_\ell\mu_\ell\theta_\ell}{2}\right), \tag{15}$$

where $\theta_\ell = 1 - \sigma^2_\ell/\mu^2_\ell$ (σ^2_ℓ and μ_ℓ respectively denotes the variance and the mean of packet inter-generation times at source ℓ), and p_ℓ is as defined in (9)—(11) for π_{sr}.

Proof of Theorem 1: Using Lemma 3 and 4, we complete the proof of Theorem 1 as follows. Recall that $f^\star_\ell$'s (defined in Lemma 3) satisfy the constraints (10) and (11). Also, under the same constraints, p_ℓ's minimize (9). Therefore,

$$\sum_{\ell=1}^{N}\left(\frac{2\rho_\ell\mu_\ell}{p_\ell} + \frac{c_\ell p_\ell}{\mu_\ell}\right) \leq \sum_{\ell=1}^{N}\left(\frac{2\rho_\ell\mu_\ell}{f^\star_\ell} + \frac{c_\ell f^\star_\ell}{\mu_\ell}\right). \tag{16}$$

From (15) and (16), we get $\mathbb{E}[\Gamma(\pi_{sr})]$

$$\leq \frac{1}{N}\sum_{\ell=1}^{N}\left(\frac{2\rho_\ell\mu_\ell}{f^\star_\ell} + \frac{c_\ell f^\star_\ell}{\mu_\ell} + \rho_\ell\gamma_\ell - \frac{\rho_\ell\mu_\ell\theta_\ell}{2}\right). \tag{17}$$

Since competitive ratio (Definition 4) of π_{sr} is $\text{CR}_{\pi_{sr}} = \mathbb{E}[\Gamma(\pi_{sr})]/\mathbb{E}[\Gamma(\pi^\star_{OF})]$, using (14) and (17), we get

$$\begin{aligned}\text{CR}_{\pi_{sr}} &\leq \frac{\frac{1}{N}\sum_{\ell=1}^{N}\left(\frac{\rho_\ell\mu_\ell}{2f^\star_\ell}\left(4 - f^\star_\ell\theta_\ell\right) + \rho_\ell\gamma_\ell + \frac{c_\ell f^\star_\ell}{\mu_\ell}\right)}{\frac{1}{N}\sum_{\ell=1}^{N}\left(\frac{\rho_\ell\mu_\ell}{2f^\star_\ell} + \rho_\ell\gamma_\ell + \frac{c_\ell f^\star_\ell}{\mu_\ell}\right)}, \\ &\leq \max_\ell\{4 - f^\star_\ell\theta_\ell\}, &(18)\\ &\leq \max\left\{4,\ \ 3 + \max_\ell\{\sigma^2_\ell/\mu^2_\ell\}\right\}, &(19)\end{aligned}$$

where we get (19) by substituting $\theta_\ell = 1 - \sigma^2_\ell/\mu^2_\ell$, and maximizing the R.H.S. of (18) with respect to $f^\star_\ell \in [0,1]$. ■

Proof sketches for Lemma 3 and 4 are provided in Appendix B and C, respectively. The detailed proofs are provided in the full version of this paper [4].

V. Numerical Results

To analyze the stationary randomized policy π_{sr} (Algorithm 1), we perform its parametric and comparative analysis using numerical simulations. This includes analyzing the effect of the system parameters (such as the number of sources N, transmission cost c_ℓ, and distributions $\mathcal{G}_\ell$ and $\mathcal{D}_\ell$) on the weighted sum cost $\Gamma(\pi_{sr})$, as well as considering relevant settings for AoI minimization from prior work, and comparing the performance of π_{sr} with other state-of-the-art policies. In this section, we discuss some key observations, and provide the detailed results in the full version of this paper [4].

Based on the parameters used in [15], we consider a system with $N = 4$ sources, with weight $[4,4,1,1]$, mean packet inter-generation time $\mu \cdot [1,(4/3),2,4]$, mean transmission time for packets $\gamma \cdot [4,2,(4/3),1]$, and cost per transmission $c \cdot [2,1,1,2]$, where μ, γ and c are parameters that we specify later for each simulation. Also, to understand the effect of non-work conserving property of SR-NSS (in π_{sr}; Remark 10) on the weighted sum cost $\Gamma(\pi_{sr})$, we consider a work-conserving policy π^{wc}_{sr} which is identical to π_{sr}, except that whenever π^{wc}_{sr} chooses a source to transmit which does not have a fresh marked packet, another source is chosen immediately. We denote the upper bound (15) on $\mathbb{E}[\Gamma(\pi_{sr})]$, and the lower bound (14) on $\mathbb{E}[\Gamma(\pi^\star_{OF})]$ by UB_{sr} and LB respectively.

To analyze the effect of the mean packet inter-generation time on $\Gamma(\pi_{sr})$, we simulate the considered system in Figure 2, by fixing $c = 1$, $\gamma = 2$, and varying μ. Also, we assume $\mathcal{G}_\ell$'s and $\mathcal{D}_\ell$'s (for all ℓ) to be the exponential distribution (because of which $\sigma^2_\ell/\mu^2_\ell = 1, \forall \ell$). Figure 2 shows that when $\theta_\ell = \sigma^2_\ell/\mu^2_\ell$ are fixed for each source ℓ, $\Gamma(\pi_{sr})$ increases with increase in μ (i.e. $\mu_\ell, \forall\ell$). This is as expected because when μ_ℓ's are large, sources need to wait longer for fresh packets to get generated, and hence, they cannot transmit at optimal time instants that would minimize $\Gamma(\pi_{sr})$, even if the channel is free. Further, in Figure 2, note that initially when μ is small, the rate of increase in $\Gamma(\pi_{sr})$ with respect to μ is small (almost 0), compared to when μ is large. This is because when μ is small, the time instants when the sources get to transmit is mainly restricted by their transmission times.

A critical property of π_{sr} is that its competitive ratio (12) is independent of the transmission time distribution $\mathcal{D}_\ell$ of the sources. To verify this fact, we fix $c = 1$ and $\mu = 16$, and for each source ℓ, we choose $\mathcal{G}_\ell$ to be the exponential distribution, and $\mathcal{D}_\ell$ to be the log-normal distribution. For each source ℓ, defining the variance of $\mathcal{D}_\ell$ to be ν^2, we simulate the system with different values of parameters ν^2 and γ. As shown in Figure 3, the weighted sum cost $\Gamma(\pi_{sr})$ is less than the upper bound UB_{sr}, and UB_{sr} as well as the lower bound LB increases with γ. Hence, the ratio of UB_{sr} and LB is a constant, less than the competitive ratio (12). Note that for different values of ν^2, the plots of $\Gamma(\pi_{sr})$ overlap.

Further, we fix $\mu = \gamma = 1$, and for each source ℓ, we choose $\mathcal{D}_\ell$ to be the exponential distribution, and $\mathcal{G}_\ell$ to be the log-normal distribution (with variance $\sigma^2_\ell = 1$). We simulate the system for different values of c, and find that when the cost per transmission is large, the weighted sum cost for π^{wc}_{sr} exceeds the weighted sum cost for π_{sr}, as shown in Figure 4.

VI. Conclusions

In this paper, we have considered the scheduling problem of finding an optimal non-preemptive policy to minimize the sum of the AoI and the transmission cost, in the presence of multiple sources, and where the inter-generation time of updates and

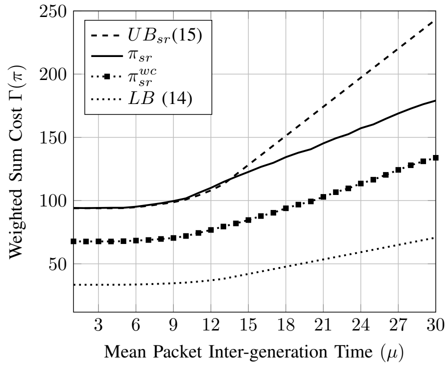

Fig. 2: $\Gamma(\pi)$ versus μ.

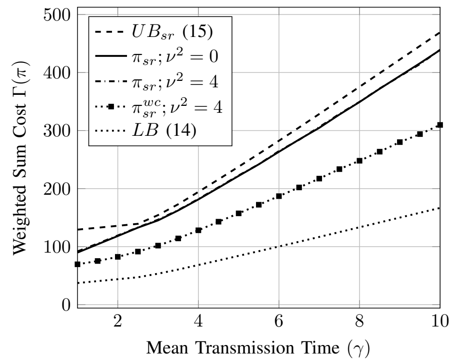

Fig. 3: $\Gamma(\pi)$ versus γ.

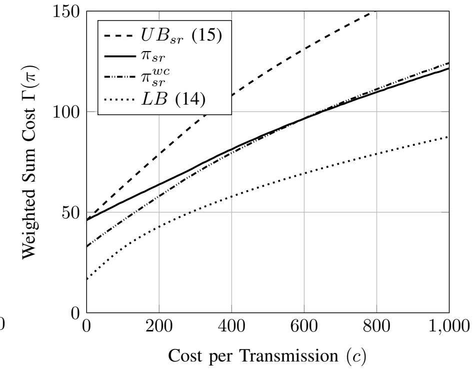

Fig. 4: $\Gamma(\pi)$ versus c.

the transmission time/delay for each update follow a general distribution. Instead of directly finding the optimal scheduling policy, we propose a randomized scheduling policy and upper bound its competitive ratio (by comparing against an offline optimal policy) by the ratio of the variance and the squared mean of the inter-generation time of updates. Notably the competitive ratio is independent of the transmission time/delay distributions, and is upper bounded by 4 for exponential, uniform, and Rayleigh inter-generation time distributions. In addition to the upper bound, we also presented a tight example to show that the competitive ratio of the considered algorithm has to depend on the ratio of the variance and squared mean of the inter-generation time of updates. Obvious question that remains open: are there policies that have constant competitive ratios, i.e., independent of the distribution of inter-generation time of updates?

References

[1] S. Kaul, R. Yates, and M. Gruteser, "Real-time status: How often should one update?" in *2012 Proceedings IEEE INFOCOM*. IEEE, 2012, pp. 2731–2735.

[2] C. Kam, S. Kompella, and A. Ephremides, "Age of incorrect information for remote estimation of a binary markov source," in *IEEE INFOCOM 2020-IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*. IEEE, 2020, pp. 1–6.

[3] B. N. B. B. Joshi, R. V. Bhat and R. Vaze, "Minimization of age of incorrect estimates of autoregressive markov processes," in *2020 18th International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOPT)*. IFIP, 2021.

[4] K. Saurav and R. Vaze, "Scheduling to minimize age of information with multiple sources," 2022. [Online]. Available: https://arxiv.org/abs/2204.02953

[5] S. K. Kaul, R. D. Yates, and M. Gruteser, "Status updates through queues," in *2012 46th Annual Conference on Information Sciences and Systems (CISS)*. IEEE, 2012, pp. 1–6.

[6] J. P. Champati, H. Al-Zubaidy, and J. Gross, "Statistical guarantee optimization for age of information for the d/g/1 queue," in *IEEE INFOCOM 2018-IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*. IEEE, 2018, pp. 130–135.

[7] Y. Inoue, H. Masuyama, T. Takine, and T. Tanaka, "A general formula for the stationary distribution of the age of information and its application to single-server queues," *IEEE Transactions on Information Theory*, vol. 65, no. 12, pp. 8305–8324, 2019.

[8] R. D. Yates and S. K. Kaul, "The age of information: Real-time status updating by multiple sources," *IEEE Transactions on Information Theory*, vol. 65, no. 3, pp. 1807–1827, 2018.

[9] M. Moltafet, M. Leinonen, and M. Codreanu, "Average age of information for a multi-source m/m/1 queueing model with packet management," in *2020 IEEE International Symposium on Information Theory (ISIT)*, 2020, pp. 1765–1769.

[10] M. A. Abd-Elmagid and H. S. Dhillon, "Age of information in multi-source updating systems powered by energy harvesting," *arXiv preprint arXiv:2109.07605*, 2021.

[11] Y. Sun, E. Uysal-Biyikoglu, R. D. Yates, C. E. Koksal, and N. B. Shroff, "Update or wait: How to keep your data fresh," *IEEE Transactions on Information Theory*, vol. 63, no. 11, pp. 7492–7508, 2017.

[12] K. Saurav, "3-competitive policy for minimizing age of information in multi-source m/g/1 queuing model," *arXiv preprint arXiv:2201.03502*, 2022.

[13] I. Kadota, A. Sinha, E. Uysal-Biyikoglu, R. Singh, and E. Modiano, "Scheduling policies for minimizing age of information in broadcast wireless networks," *IEEE/ACM Transactions on Networking*, vol. 26, no. 6, pp. 2637–2650, 2018.

[14] 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*, vol. 27, no. 4, pp. 1359–1372, 2019.

[15] I. Kadota and E. Modiano, "Minimizing the age of information in wireless networks with stochastic arrivals," *IEEE Transactions on Mobile Computing*, vol. 20, no. 3, pp. 1173–1185, 2021.

[16] R. V. Bhat, R. Vaze, and M. Motani, "Throughput maximization with an average age of information constraint in fading channels," *IEEE Transactions on Wireless Communications*, vol. 20, no. 1, pp. 481–494, 2020.

[17] ——, "Minimization of age of information in fading multiple access channels," *IEEE Journal on Selected Areas in Communications*, vol. 39, no. 5, pp. 1471–1484, 2021.

[18] K. Saurav and R. Vaze, "Minimizing the sum of age of information and transmission cost under stochastic arrival model," in *IEEE INFOCOM 2021 - IEEE Conference on Computer Communications*, 2021, pp. 1–10.

[19] N. Rajaraman, R. Vaze, and G. Reddy, "Not just age but age and quality of information," *IEEE Journal on Selected Areas in Communications*, vol. 39, no. 5, pp. 1325–1338, 2021.

[20] L. Schrage, "Letter to the editor- a proof of the optimality of the shortest remaining processing time discipline," *Operations Research*, vol. 16, no. 3, pp. 687–690, 1968.

[21] L. A. Hall, A. S. Schulz, D. B. Shmoys, and J. Wein, "Scheduling to minimize average completion time: Off-line and on-line approximation algorithms," *Mathematics of operations research*, vol. 22, no. 3, pp. 513–544, 1997.

[22] R. L. Graham, "Bounds for certain multiprocessing anomalies," *Bell system technical journal*, vol. 45, no. 9, pp. 1563–1581, 1966.

[23] A. Bhimaraju, D. Nayak, and R. Vaze, "Non-clairvoyant scheduling of coflows," in *2020 18th International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOPT)*. IFIP, 2020, pp. 1–8.

[24] D. G. Luenberger, Y. Ye *et al.*, *Linear and nonlinear programming*. Springer, 1984, vol. 2.

[25] S. M. Ross, *Introduction to probability models*. Academic press, 2014.

Appendix A

Proof of Lemma 2: Figure 1 shows a general age plot for source ℓ in terms of the quantities defined in this section so far. Corresponding to each period $\mathcal{P}_{\ell i}^{\pi}$ until time t, the age cost is $Q^\pi_{\ell i} = (T^\pi_{\ell i})^2/2 + T^\pi_{\ell i} Z^\pi_{\ell i}$. Thus, the AoI for source ℓ satisfies $\lim_{t\to\infty} \Delta^{av}_{\ell,\pi}(t) = \lim_{t\to\infty} \left(\frac{\sum_{i=1}^{R^\pi_\ell(t)} \left((T^\pi_{\ell i})^2/2 + T^\pi_{\ell i} Z^\pi_{\ell i} \right)}{t} + \frac{(\eta^\pi_\ell(t))^2}{2t} \right)$.

Note that in the fraction $(\eta^\pi_\ell(t))^2/(2t)$, the numerator $(\eta^\pi_\ell(t))^2$ is finite with probability 1 (from Lemma 1), but in the denominator, $t \to \infty$. Thus,

$$\lim_{t\to\infty} \Delta^{av}_{\ell,\pi}(t) = \lim_{t\to\infty} \frac{\sum_{i=1}^{R^\pi_\ell(t)} \left(\frac{(T^\pi_{\ell i})^2}{2} + T^\pi_{\ell i} Z^\pi_{\ell i} \right)}{t}. \tag{20}$$

Substituting (20) and (2) into (3) and substituting $Z^\pi_{\ell i} = w^\pi_{\ell i} + d_{\ell i}$ (5) and $t = \sum_{i=1}^{R^\pi_\ell(t)} T^\pi_{\ell i}$ (Remark 8) in the denominator of the resulting expression, we get (6).

To obtain (7), note that at any time at most one packet can be under transmission, and the transmission of packet ℓ^π_i takes $d_{\ell i}$ time units, where $d_{\ell i}$'s are independent and identically distributed random variables with mean γ_ℓ. Therefore, assuming (without loss of generality) that policy π transmits updates $\ell^\pi_1, \cdots, \ell^\pi_{U^\pi_\ell(t)}$, we get $\sum_{\ell=1}^N \sum_{i=1}^{U^\pi_\ell(t)} d_{\ell i} \leq t$. Dividing both sides by t, and taking limit as $t \to \infty$, we get $1 \geq \lim_{t\to\infty} \sum_{\ell=1}^N \left(\frac{\sum_{i=1}^{U^\pi_\ell(t)} d_{\ell i}}{U^\pi_\ell(t)} \cdot \frac{U^\pi_\ell(t)}{t} \right)$. Since the limit of a product is equal to the product of the limits (when the limits exists, as in the above case), we get $1 \geq \sum_{\ell=1}^N \left(\lim_{t\to\infty} \frac{\sum_{i=1}^{U^\pi_\ell(t)} d_{\ell i}}{U^\pi_\ell(t)} \cdot \lim_{t\to\infty} \frac{U^\pi_\ell(t)}{t} \right) \overset{(a)}{=} \lim_{t\to\infty} \sum_{\ell=1}^N \frac{\gamma_\ell U^\pi_\ell(t)}{\sum_{i=1}^{R^\pi_\ell(t)} T^\pi_{\ell i}}$ (with probability 1), where (a) follows from the strong law of large numbers (and substituting $t = \sum_{i=1}^{R^\pi_\ell(t)} T^\pi_{\ell i}$ from Remark 8). Note that here we could use the strong law of large numbers because (i) $d_{\ell i}$'s (for all i) are independent and identically distributed with mean γ_ℓ, and (ii) as $t \to \infty$, $U^\pi_\ell(t) \to \infty$ as well (from Corollary 2 below, for $\mathcal{F}^\pi_\ell$ equal to the set of completely transmitted packets, in which case $U^\pi_\ell(t) = R^\pi_\ell(t) \to \infty$).

Corollary 2: For any policy $\pi \in \Pi_S \cup \Pi_{OF}$ and subset $\mathcal{F}^\pi_\ell$, when $t \to \infty$, the number of packets in $\mathcal{F}^\pi_\ell$ is $R^\pi_\ell(t) \to \infty$.

Proof: Since $t = \sum_{i=1}^{R^\pi_\ell(t)} T^\pi_{\ell i}$, where $T^\pi_{\ell i}$'s are finite, we get that as $t \to \infty$, $R^\pi_\ell(t) \to \infty$ as well. ∎

Appendix B

Proof Sketch for Lemma 3: Let $\mathcal{F}^\pi_\ell$ be equal to the set of packets of source ℓ that are transmitted by policy π. Then, $R^\pi_\ell(t) = U^\pi_\ell(t)$. Since $w^\pi_{\ell i} = s^\pi_{\ell i} - g^\pi_{\ell i} \geq 0$, for a lower bound on the weighted sum cost $\Gamma(\pi)$ for policy $\pi \in \Pi_{OF}$, in (6), we substitute $U^\pi_\ell(t) = R^\pi_\ell(t)$ and $w^\pi_{\ell i} = 0$, and get $\Gamma(\pi) = \lim_{t\to\infty} \frac{1}{N} \sum_{\ell=1}^N \left(\frac{\rho_\ell \sum_{i=1}^{R^\pi_\ell(t)} (T^\pi_{\ell i})^2}{2\sum_{i=1}^{R^\pi_\ell(t)} T^\pi_{\ell i}} + \frac{\rho_\ell \sum_{i=1}^{R^\pi_\ell(t)} T^\pi_{\ell i} d_{\ell i}}{\sum_{i=1}^{R^\pi_\ell(t)} T^\pi_{\ell i}} + \frac{c_\ell R^\pi_\ell(t)}{\sum_{i=1}^{R^\pi_\ell(t)} T^\pi_{\ell i}} \right)$. Using Jensen's inequality, we get that $\frac{\rho_\ell \sum_{i=1}^{R^\pi_\ell(t)} (T^\pi_{\ell i})^2}{2\sum_{i=1}^{R^\pi_\ell(t)} T^\pi_{\ell i}} \geq \frac{\rho_\ell \sum_{i=1}^{R^\pi_\ell(t)} T^\pi_{\ell i}}{2R^\pi_\ell(t)}$. Further, using the fact that $d_{\ell i}$'s $(\forall i)$ are independent and identically distributed as per distribution $\mathcal{D}_\ell$ (with mean γ_ℓ), and for each i, $d_{\ell i}$ is independent of the inter-generation time $T^\pi_{\ell i}$ of transmitted packets, we show that $\mathbb{E}\left[\frac{\rho_\ell \sum_{i=1}^{R^\pi_\ell(t)} T^\pi_{\ell i} d_{\ell i}}{\sum_{i=1}^{R^\pi_\ell(t)} T^\pi_{\ell i}} \right] = \rho_\ell \gamma_\ell$. Thus, $\mathbb{E}[\Gamma(\pi)] \geq \lim_{t\to\infty} \frac{1}{N} \sum_{\ell=1}^N \left(\mathbb{E}\left[\frac{\rho_\ell \sum_{i=1}^{R^\pi_\ell(t)} T^\pi_{\ell i}}{2R^\pi_\ell(t)} \right] + \rho_\ell \gamma_\ell + \mathbb{E}\left[\frac{c_\ell R^\pi_\ell(t)}{\sum_{i=1}^{R^\pi_\ell(t)} T^\pi_{\ell i}} \right] \right)$.

Finally, to prove (14), we show that for $\pi = \pi^\star_{OF}$, $\frac{R^\pi_\ell(t)}{\sum_{i=1}^{R^\pi_\ell(t)} T^\pi_{\ell i}} = \frac{f^\star_\ell}{\mu_\ell}$ is the number of packets of source ℓ that are transmitted by $\pi^\star_{OF}$ per unit time. Further, substituting $\frac{U^\pi_\ell(t)}{\sum_{i=1}^{R^\pi_\ell(t)} T^\pi_{\ell i}} = \frac{R^\pi_\ell(t)}{\sum_{i=1}^{R^\pi_\ell(t)} T^\pi_{\ell i}} = \frac{f^\star_\ell}{\mu_\ell}$ in (7), we get $\sum_{\ell=1}^N \gamma_\ell \frac{f^\star_\ell}{\mu_\ell} \leq 1$.

Appendix C

Proof Sketch for Lemma 4: Let $\mathcal{F}^{sr}_\ell$ be equal to the set of packets that are *marked* by π_{sr} (SR-PMS). Since the set of packets that are transmitted by π_{sr} is a subset of the marked packets, we have $U^{sr}_\ell(t) \leq R^{sr}_\ell(t)$. Therefore, to upper bound $\Gamma(\pi_{sr})$, in (6), we upper bound $U^{sr}_\ell(t)$ by $R^{sr}_\ell(t)$. Also, since packets of each source ℓ are marked independently with probability p_ℓ, and $d_{\ell i} \sim \mathcal{D}_\ell$ (with mean γ_ℓ), we get that $T^{sr}_{\ell i}$'s and $d_{\ell i}$'s $(\forall i)$ are independent. Using this fact, we show that $\mathbb{E}\left[\frac{\rho_\ell \sum_{i=1}^{R^{sr}_\ell(t)} T^{sr}_{\ell i} d_{\ell i}}{\sum_{i=1}^{R^{sr}_\ell(t)} T^{sr}_{\ell i}} \right] = \rho_\ell \gamma_\ell$. Then for $\pi = \pi_{sr}$, from (6) we get $\mathbb{E}[\Gamma(\pi_{sr})] = \lim_{t\to\infty} \frac{1}{N} \sum_{\ell=1}^N \left(\mathbb{E}\left[\frac{\rho_\ell \sum_{i=1}^{R^{sr}_\ell(t)} (T^{sr}_{\ell i})^2}{2\sum_{i=1}^{R^{sr}_\ell(t)} T^{sr}_{\ell i}} \right] + \mathbb{E}\left[\frac{\rho_\ell \sum_{i=1}^{R^{sr}_\ell(t)} T^{sr}_{\ell i} w^{sr}_{\ell i}}{\sum_{i=1}^{R^{sr}_\ell(t)} T^{sr}_{\ell i}} \right] + \rho_\ell \gamma_\ell + \mathbb{E}\left[\frac{c_\ell R^{sr}_\ell(t)}{\sum_{i=1}^{R^{sr}_\ell(t)} T^{sr}_{\ell i}} \right] \right)$.

Let $\hat{w}^{sr}_{\ell i}$ denote the difference between the successive time instants when source ℓ is chosen to transmit by π_{sr} (SR-NSS). Then by definition, $w^{sr}_{\ell i} \leq \hat{w}^{sr}_{\ell i}$. Thus, $\mathbb{E}\left[\frac{\rho_\ell \sum_{i=1}^{R^{sr}_\ell(t)} T^{sr}_{\ell i} w^{sr}_{\ell i}}{\sum_{i=1}^{R^{sr}_\ell(t)} T^{sr}_{\ell i}} \right] \leq \mathbb{E}\left[\frac{\rho_\ell \sum_{i=1}^{R^{sr}_\ell(t)} T^{sr}_{\ell i} \hat{w}^{sr}_{\ell i}}{\sum_{i=1}^{R^{sr}_\ell(t)} T^{sr}_{\ell i}} \right]$. By design, π_{sr} is such that the packet marking (SR-PMS) and source selection (SR-NSS) for transmission are independent. Thus, $T^{sr}_{\ell i}$ and $\hat{w}^{sr}_{\ell i}$ are mutually independent. Also, $\mathbb{E}[\hat{w}^{sr}_{\ell i}] \leq \mathbb{E}[T^{sr}_{\ell i}] = \mu_\ell/p_\ell$. Hence, $\mathbb{E}[T^{sr}_{\ell i} \hat{w}^{sr}_{\ell i}] = \mathbb{E}[T^{sr}_{\ell i}]\mathbb{E}[\hat{w}^{sr}_{\ell i}] \leq (\mu_\ell/p_\ell)^2$. Using this fact, we show that $\mathbb{E}\left[\frac{\rho_\ell \sum_{i=1}^{R^{sr}_\ell(t)} T^{sr}_{\ell i} \hat{w}^{sr}_{\ell i}}{\sum_{i=1}^{R^{sr}_\ell(t)} T^{sr}_{\ell i}} \right] \leq \frac{\rho_\ell (\mu_\ell/p_\ell)^2}{\mu_\ell/p_\ell} = \frac{\rho_\ell \mu_\ell}{p_\ell}$. Hence, $\mathbb{E}[\Gamma(\pi_{sr})] \leq \lim_{t\to\infty} \frac{1}{N} \sum_{\ell=1}^N \left(\mathbb{E}\left[\frac{\rho_\ell \sum_{i=1}^{R^{sr}_\ell(t)} (T^{sr}_{\ell i})^2}{2\sum_{i=1}^{R^{sr}_\ell(t)} T^{sr}_{\ell i}} \right] + \frac{\rho_\ell \mu_\ell}{p_\ell} + \rho_\ell \gamma_\ell + \mathbb{E}\left[\frac{c_\ell R^{sr}_\ell(t)}{\sum_{i=1}^{R^{sr}_\ell(t)} T^{sr}_{\ell i}} \right] \right)$.

Since the inter-generation time of packets at source ℓ are independent and identically distributed as per distribution $\mathcal{G}_\ell$ (with mean μ_ℓ), and every generated packet is marked with fixed probability p_ℓ, we get that the inter-generation time of the marked packets, i.e. $T^{sr}_{\ell i}$'s $(\forall i)$ are independent and identically distributed such that $\mathbb{E}[T^{sr}_{\ell i}] = \mu_\ell/p_\ell$ and $\mathbb{E}[(T^{sr}_{\ell i})^2] = (\sigma^2_\ell/p_\ell) + (2 - p_\ell)(\mu^2_\ell/p^2_\ell)$, where σ^2_ℓ is the variance of the inter-generation time of packets at source ℓ. Thus, using the renewal reward theorem [25], we get $\mathbb{E}\left[\frac{\rho_\ell \sum_{i=1}^{R^{sr}_\ell(t)} (T^{sr}_{\ell i})^2}{2\sum_{i=1}^{R^{sr}_\ell(t)} T^{sr}_{\ell i}} \right] = \frac{\rho_\ell}{2} \left(\frac{\sigma^2_\ell}{\mu_\ell} + \frac{\mu_\ell}{p_\ell}(2 - p_\ell) \right)$, and using the strong law of large numbers, we get $\mathbb{E}\left[\frac{c_\ell R^{sr}_\ell(t)}{\sum_{i=1}^{R^{sr}_\ell(t)} T^{sr}_{\ell i}} \right] = \frac{c_\ell p_\ell}{\mu_\ell}$. Therefore, $\mathbb{E}[\Gamma(\pi_{sr})] \leq \lim_{t\to\infty} \frac{1}{N} \sum_{\ell=1}^N \left(\frac{\rho_\ell}{2} \left(\frac{\sigma^2_\ell}{\mu_\ell} + \frac{\mu_\ell}{p_\ell}(2 - p_\ell) \right) + \frac{\rho_\ell \mu_\ell}{p_\ell} + \rho_\ell \gamma_\ell + \frac{c_\ell p_\ell}{\mu_\ell} \right)$. Rearranging the terms, we get (15).