Decentralized Status Update for Age-of-Information Optimization in Wireless Multiaccess Channels

Zhiyuan Jiang[1], Bhaskar Krishnamachari[2], Xi Zheng[1], Sheng Zhou[1], Zhisheng Niu[1],
[1]{zhiyuan@, zhengx14@mails., sheng.zhou@, niuzhs@}tsinghua.edu.cn, Tsinghua University,
[2] bkrishna@usc.edu, University of Southern California

***Abstract*—We consider a system where multiple terminals transmit their randomly generated status updates to a base station (BS) sharing a wireless multiaccess uplink channel. The problem of interest, especially in massive Internet-of-Things systems, is that how to schedule the terminals to minimize the time-average age-of-information in a decentralized manner, namely terminals transmit autonomously without signalling exchange (overhead) with the BS or other terminals. Towards this end, the round-robin with one-packet buffers (the newest packet at each terminal only) policy (RR-ONE) is proposed and proved optimal among arrival-independent renewal (AIR) policies. In addition to its simple structure which is instrumental for decentralized implementation, RR-ONE is further proved asymptotically (massive terminals) optimal among all policies, including centralized and non-causal policies.**

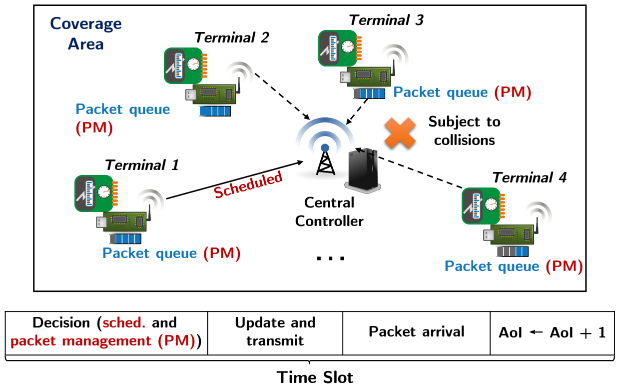

Fig. 1. Considered system architecture and status update procedure.

I. Introduction

Age-of-information (AoI) is a recently proposed metric specifically to quantify the timeliness of status information [1]; it characterizes the information monitoring latency at a destination node, or simply put, time elapsed since the last-updated packet's generation. This definition jointly accounts for the delay introduced by sampling and data communication, which distinguishes itself from the conventional end-to-end (e2e) communication (queuing and transmission) delay metric. The AoI only coincides with the conventional e2e delay at the time when a status update packet is successfully delivered; another distinct difference is that the e2e delay is defined for each packet, however, the AoI is a constantly evolving measurement at the destination. In systems where the timeliness of status information is critical for real-time applications, the AoI becomes an important metric to optimize.

The wireless communication system, which plays an integral part in status update systems, is therefore well motivated to optimize AoI, especially for future massive Internet-of-Things (IoT) enabled real-time applications. One of the fundamental restrictions of wireless communication systems is that the transmissions are subject to interference due to the broadcast nature of electromagnetic waves, which amounts to the fact that terminals cannot transmit simultaneously; otherwise collisions happen and transmissions fail with no data delivered. Therefore, terminals should be carefully *scheduled* to avoid such collisions. Meanwhile, in wireless uplinks with a large number of terminals, the scheduling design is faced with the overhead issue which is caused by requiring global state information for the scheduling decisions. The exchange of the state information entails significant signaling overhead which cannot be overlooked. In view of this, the scheduling policy is preferably *decentralized*, i.e., decisions are made autonomously at terminals, requiring only local information. For instance, the carrier-sensing-medium-access (CSMA) scheme is a widely-used and successful application of decentralized protocol in wireless networks. In particular, terminals transmit based on a contention protocol and scheduling decisions are made in a decentralized manner. However, the CSMA protocol is only designed for throughput maximization and may face severe challenges in status update systems.

A. Main Contributions and Related Work

Concerning the aforementioned scenario and corresponding challenges, our main contributions include: (a) Among arrival-independent renewal (AIR) scheduling policies, whose decisions are independent with packet-arrival processes and hence suitable for decentralization, a round-robin policy with one-packet buffers (only retains the most up-to-date packet) at terminals (RR-ONE) is proved optimal. The proof technique leverages a generalized Poisson-arrival-see-time-average (PASTA) theorem which, as far as we know, has not been adopted in the related literature before. (b) RR-ONE is proved asymptotically optimal among *all* policies with a massive number of terminals. It is shown that the optimum time-average AoI is proportional to the number of terminals asymptotically; the optimum linear scaling factor is $\frac{1}{2}$; RR-ONE is proved to achieve the optimum scaling factor.

Among the recent progress on AoI optimization, e.g., [1]–[7], the study on multi-queue scheduling problems is the most related work [5]–[7]. Hsu *et al.* approach this problem considering the wireless broadcast channel where the scheduling decisions are centralized; they prove the optimal policy is age-threshold-based. The scheduling problem of multiple sources inside a finite-length transmission frame to minimize AoI is proved NP-hard [6]. The Whittle's index is leveraged by Kadota *et al.* [5] based on a restless multi-armed-bandit formulation; it is shown that an age-greedy policy is optimal in the symmetric case and the Whittle's index is derived for the asymmetric case. The major distinction between our formulation and existing work is that we assume: (a) status packets arrive randomly; (b) limited information is available for the decisions to facilitate decentralized implementation. As far as we know, no existing work on AoI has addressed a practical scenario involving both aspects.

II. System Model and Main Results

Consider a base station (BS), alternatively referred to as central controller or fusion center, which collects status packets from multiple terminals, as shown in Fig. 1. A time-slotted system is considered. The status data packets are generated and stored at terminal buffers. The number of packets generated at time t of terminal n is denoted by $L_n(t)$ and $L_n(t)$ is assumed to be a Bernoulli random variable with parameter λ_n; the arrival processes $\{L_n(t),\, t = 1, 2, ...\}$, $n \in \{1, ..., N\}$ are independent over terminals and time. The number of terminals is denoted by N. Let $U_{n,\pi}(t)$ denote the scheduling decision of terminal n at time t for a given policy π, i.e., $U_{n,\pi}(t) = 1$ if terminal n is scheduled and $U_{n,\pi}(t) = 0$ otherwise.

The time-average AoI is denoted by $\bar{h}_\pi^{(T,N)} \triangleq \frac{1}{TN}\sum_{t=1}^{T}\sum_{n=1}^{N} h_{n,\pi}(t)$, where the AoI at the t-th time slot for terminal n based on policy π is denoted by $h_{n,\pi}(t)$, and the time horizon is T. Denote time-average AoI over infinite time horizon as

$$\bar{h}_\pi^{(\infty,N)} \triangleq \lim_{T\to\infty} \bar{h}_\pi^{(T,N)}. \tag{1}$$

The evolution of AoI can be written as

$$h_{n,\pi}(t+1) = h_{n,\pi}(t) - U_{n,\pi}(t)\prod_{m\neq n}(1 - U_{m,\pi}(t))g_n(t) + 1, \tag{2}$$

where $g_n(t)$ denotes the AoI reduction with a successful update from terminal n. Consequently, we have $g_n(t) = 0$ when queue-n is empty at time t. The AoI for each terminal always increases by one after each time slot. Based on this definition (2), whenever a collision happens, i.e., more than one terminals transmit in the same time slot, no status is updated. Note that transmission failures only happen with collisions, otherwise the transmission is assumed successful; this corresponds to the interference-limited regime and therein failures due to noise are negligible. In addition, denote the average AoI of terminal-n under policy π as $\bar{h}_{n,\pi}^{(T)} \triangleq \frac{1}{T}\sum_{t=1}^{T} h_{n,\pi}(t)$.

The status update procedure is described in Fig. 1. We assume the following sequence of events in each time slot. At the beginning of each time slot, scheduling decisions are made, including:

- *Terminal scheduling*: Decide which terminal updates and transmits in this time slot.
- *Packet management*: Once scheduled, the terminal can apply a packet management scheme, e.g., it can choose a packet from its queue to transmit, or drop several packets.

Based on the scheduling decision, the scheduled terminal transmits its update packet in the uplink (assuming one packet is transmitted in each time slot), and thereby the AoI is refreshed at the BS. Afterwards, packets arrive randomly at terminals (the age of newly arrived packets is zero) and then the age of all packets and AoIs of all terminals increase by one. This marks the end of a time slot. The AoI at the t-th time slot is defined as the AoI at the end of the time slot.

The objective is to minimize the infinite-horizon time-average AoI (1) over all feasible policies. As a first step, the following definition and Lemma 1 (cf. proof in [8]) enable us to only consider work-conserving non-collision (WCNC) policies without loss of optimality.

Definition 1 (WCNC policy): A WCNC policy is a policy that is never idle when there is at least one packet in terminal queues, nor schedules multiple terminals simultaneously. □

Lemma 1: For a non-WCNC policy, there exists at least one WCNC policy that achieves lower AoI. □

As discussed previously, for overhead concerns that decisions should be decentralized, we first consider AIR policies defined as follows. Denote the resultant scheduling interval process of terminal-n based on policy π as $X_{n,\pi}^{(k)}$, $k = 1, 2, ...$ where k is the scheduling interval index. Define $R_{n,\pi}(t)$ as the counting process of scheduling times before time t for terminal n, i.e., $R_{n,\pi}(t) \triangleq \sup\{r : \sum_{k=0}^{r} X_{n,\pi}^{(k)} \leq t\}$.

Definition 2 (AIR policy): A policy π is an AIR policy if the following conditions are both met.

(i) The scheduling interval processes $\{X_{n,\pi}^{(k)},\, n = 1, ..., N\}$ are independent with the packet arrival processes at terminals, with finite first and second raw moments denoted by m_n and v_n respectively.
(ii) The counting processes $R_{n,\pi}(t)$ are renewal processes. □

By definition, the set of AIR policies is essentially a subset of all policies. The condition (i) is in fact reflecting the practical perspective that the scheduling decisions are desired to be independent of the packet arrival processes to enable decentralized implementation and reduce signalling exchange overhead. The condition (ii) does enforce an additional constraint that the scheduling intervals are i.i.d.; however the distributions can be arbitrary as long as they have finite first and second moments. Note that, notwithstanding these conditions, it is found (Theorem 2) that the optimal AIR policy with proper packet management is asymptotically optimal among all policies in the massive IoT regime.

A. Main Results

Definition 3 (RR-ONE): RR-ONE, denoted by RR in the subscript, is defined as a policy that schedules the n_{RR}-th terminal at each time slot which satisfies $n_{\mathsf{RR}} =$

$\min\{n : \tau_n = \max_{m=1,...,N} \tau_m\}$, and only stores the latest-arrival packet at each terminal. The time since last update from terminal m is denoted by τ_m. □

Theorem 1: RR-ONE is the optimal AIR policy to minimize the time-average AoI, with

$$\bar{h}_{\mathsf{RR}}^{(\infty,N)} = \frac{1}{N}\sum_{n=1}^{N}\frac{1}{\lambda_n} + \frac{N-1}{2}. \square \tag{3}$$

Theorem 2: RR-ONE is asymptotically optimal among all policies in the massive IoT regime; it achieves the optimum asymptotic scaling factors, i.e.,

$$\lim_{N\to\infty}\frac{\bar{h}_{\mathsf{RR}}^{(\infty,N)}}{N} = \lim_{N\to\infty}\frac{\bar{h}_{\text{opt}}^{(\infty,N)}}{N} = \frac{1}{2},\ \forall\lambda_i, \tag{4}$$

$$\lim_{\frac{1}{\lambda_i}\to\infty}\lambda_i\bar{h}_{\mathsf{RR}}^{(\infty,N)} = \lim_{\frac{1}{\lambda_i}\to\infty}\lambda_i\bar{h}_{\text{opt}}^{(\infty,N)} = \frac{1}{N}, \quad \forall\lambda_{j,j\neq i} \text{ and } N, \tag{5}$$

where $\bar{h}_{\text{opt}}^{(\infty,N)}$ denotes the minimum time-average AoI. □

Corollary 1: The minimum time-average AoI scales linearly with the number of terminals asymptotically, and the optimum scaling factor is $\frac{1}{2}$. □

The above results mainly establish the optimality of RR-ONE; another important note is that RR-ONE can be easily adapted for decentralized implementation. Roughly illustrated, each terminal is assigned a unique time slot to transmit, in a frame of length N, and only retains the most up-to-date packet in its buffer. A detailed protocol which accounts for variable N, i.e., random terminal appearances, is described in [8].

III. Proof of Theorem 1: Optimal AIR Policy

The quest for the optimal policy, among all policies with any given N, to minimize the time-average AoI seems elusive, because the problem can be essentially viewed as a restless multi-armed bandit problem with time- and arm-correlated reward functions. Besides, there is a strong probability that the optimal policy requires global information exchange and hence decentralization-unfriendly. Therefore, in this section, we resolve to derive the optimal AIR policy to minimize the time-average AoI in (1) following a generalized PASTA theorem, i.e., the arrival-see-time-average (ASTA) property with a Markov state process as the observed process and an independent outside observer [9].

First, consider the queue evolution of terminal-n based on AIR policies; it is similar with an $M/G/1$ queue given the definition of AIR policies, with a subtle, but important, difference that the service (in this case the service time is the scheduling interval) begins immediately after a packet departure, even if there is no packet waiting in the queue. In the case that there is no packet is in the queue, the service proceeds independently till the end, i.e., scheduled, during which period two possible circumstances can occur: 1) there are (at least one) packet arrivals and thereby one of the packets is updated under a certain packet management policy; 2) there is no packet arrival and consequently no packet is updated. It is clear that under this queue model, the optimal packet management, under arbitrary scheduling policy, is to always update the most up-to-date packet, i.e., the packet that arrives the last; the resultant queue is equivalent to having a buffer size of one and storing only the latest arrival packet. Note that this packet management policy is not necessarily optimal with preemptive service model due to service interruption [4]. However, the service (update of a terminal) in this paper is assumed instant, and hence, without loss of optimality, we only consider this packet management policy henceforth. The age of the packet in queue-n (buffer size is one) is denoted by $A_n(t)$, $t = 1, 2, ...$, a sample path of which is shown in the left of Fig. 2. Upon a packet arrival, e.g., a_i in Fig. 2, the age $A_n(t)$ drops to one (measured at the end of the time slot). When terminal-n is scheduled at the time of s_i, the AoI at the BS is updated to the age of the packet at terminal n, i.e., $A_n(s_i)$. Note that we prescribe a generalized age of $A_n(s_i)$ that between each update and next packet arrival, e.g., between s_1 and a_1, $A_n(t)$ equals the AoI of terminal n at the BS although there is no packet in the queue during the time. By doing this, we make $A_n(t)$ evolve *independently* with $h_{n,\pi}(t)$ while not affecting the AoI update procedure; this is crucial for the ASTA property to apply.

Based on the renewal process condition of AIR policies, and following the same arguments in, e.g., [4], the time-average AoI can be readily calculated by the sum of the geometric areas $Q_{k,n}$ in Fig. 2:

$$\bar{h}_{n,\pi}^{(\infty,N)} = \lim_{T\to\infty}\frac{K}{T}\frac{1}{K}\sum_{k=1}^{K}Q_{k,n} = \frac{\mathbb{E}[Q_{k,n}]}{m_n}. \tag{6}$$

The last equality is based on the elementary renewal theorem.

$$\begin{aligned}
\bar{h}_{n,\pi}^{(\infty,N)} &= \frac{1}{m_n}\mathbb{E}\left[X_{n,\pi}^{(k)}A_n(s_k) + \left(X_{n,\pi}^{(k)}-1\right)\frac{X_{n,\pi}^{(k)}}{2}\right] \\
&\overset{(a)}{=} \frac{1}{m_n}\left(\mathbb{E}\left[X_{n,\pi}^{(k)}\right]\mathbb{E}\left[A_n(s_k)\right] + \frac{1}{2}(v_n - m_n)\right) \\
&= \mathbb{E}\left[A_n(s_k)\right] + \frac{v_n - m_n}{2m_n} \\
&\overset{(b)}{\geq} \mathbb{E}\left[A_n(s_k)\right] + \frac{m_n - 1}{2},
\end{aligned} \tag{7}$$

where the equality (a) is based on the arrival-independent condition of AIR policies, and the inequality (b) follows from $v_n \geq m_n^2$; note that the equality holds when the scheduling interval is a constant m_n.

It is now clear that the main challenge is to calculate $\mathbb{E}\left[A_n(s_k)\right]$. First we have Lemma 2 (cf. proof in [8]) which shows that $\{A_n(t), t = 1, 2...\}$ is a Markov state process and its steady-state stationary distribution is given below.

Lemma 2: $\{A_n(t), t = 1, 2...\}$ is a Markov state process with the steady-state stationary distribution given as $\mu_n(j) = \lambda_n(1-\lambda_n)^{j-1}$, where $\mu_n(j)$ denotes the probability that the steady-state of terminal n is state j (age of packet at terminal-n equals j). □

Then the challenge of calculating $\mathbb{E}\left[A_n(s_k)\right]$ is tackled by treating $\mathbb{E}\left[A_n(s_k)\right]$ as the average state value of a Markov state process by an independent outside observer. Armed with

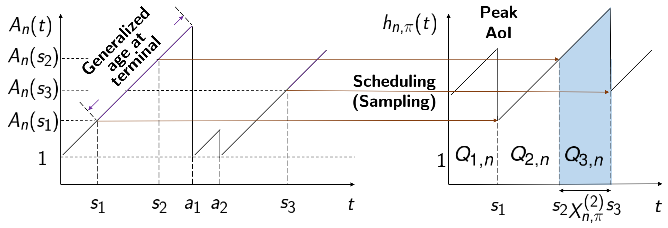

Fig. 2. Age of the packet at terminal-n assuming one packet buffer (left) and AoI at the BS (right).

this, we invoke the ASTA property [9, Theorem 3.14] which can be seen as a generalization of the well-known PASTA theorem to non-Poisson observers. Using the notations in [9, Theorem 3.14], let U be $\{A_n(t),\, t = 1, 2...\}$, and N be the counting process of the number of scheduling times before time t. Then based on the AIR policy conditions, U and N are independent and U is a Markov state process. The continuous condition in [9, Theorem 3.14] follows by design of update sequence described in Fig. 1. Therefore, we obtain

$$\mathbb{E}\left[A_n(s_k)\right] = \lim_{K\to\infty} \frac{1}{K}\sum_{k=1}^{K} A_n(s_k) = \mathbb{E}\left[A_n(t)\right] = \frac{1}{\lambda_n}. \quad (8)$$

In other words, the time-average of random sampling (s_k) of the Markov process $\{A_n(t),\, t = 1, 2...\}$ equals the steady-state average. Combining with (7), the time-average AoI is

$$\bar{h}_{\pi}^{(\infty,N)} = \frac{1}{N}\sum_{n=1}^{N} \bar{h}_{n,\pi}^{(\infty,N)} \geq \frac{1}{N}\sum_{n=1}^{N}\left(\mathbb{E}\left[A_n(s_k)\right] + \frac{m_n - 1}{2}\right)$$
$$= \frac{1}{N}\sum_{n=1}^{N}\left(\frac{1}{\lambda_n} + \frac{m_n - 1}{2}\right). \quad (9)$$

The scheduling rate of all terminals equals one for WCNC policies; therefore, according to the elementary renewal theorem, $\sum_{n=1}^{N} \frac{1}{m_n} = 1$. It follows from the arithmetic-harmonic-mean inequality (equality holds when $m_n = N$, $\forall n = 1, ..., N$) and (9) that

$$\bar{h}_{\pi}^{(\infty,N)} \geq \frac{1}{N}\sum_{n=1}^{N}\frac{1}{\lambda_n} + \frac{N-1}{2}. \quad (10)$$

The equality holds in (10) under two conditions: 1) $m_n = N, \forall n = 1, ..., N$; 2) The scheduling interval is a constant m_n. These two conditions can be both satisfied with RR-ONE. For a sanity check, RR-ONE is indeed an AIR policy. With this, we can conclude the proof of Theorem 1.

IV. Proof of Theorem 2: Asymptotic Optimality

In future massive IoT systems, the main challenge is to accommodate a large number of terminals while maintaining timely status updates. Hence, it is of particular interest to consider the problem in the asymptotic regime where the number of terminals grows large. Towards this end, it will be shown in the following that RR-ONE, given its simple structure, is asymptotically optimal among all policies with arbitrary information and even non-causal packet arrival knowledge. First, we obtain two AoI lower bounds and compare these with the achievable AoI by RR-ONE; the conclusion follows by showing that they have identical asymptotic scaling factors.

First, we introduce two lower bounds of the time-average AoI of any feasible policies in Lemma 3 and Lemma 4. The proofs are both straightforward by considering two genie-aided systems wherein the AoI is always updated to one after each scheduling, and transmissions are no longer subject to collisions, respectively. For detailed proofs, please see [8].

Lemma 3: The time-average AoI in (1) is no less than $\frac{N+1}{2}$, i.e.,

$$\bar{h}_{\pi}^{(\infty,N)} \geq \frac{N+1}{2},\ \forall N = 1, 2, ...,\ \lambda_n \in [0, 1], n \in \{1, ..., N\}. \square \quad (11)$$

Lemma 4: The time-average AoI in (1) is no less than $\frac{1}{N}\sum_{n=1}^{N}\frac{1}{\lambda_n}$, i.e.,

$$\bar{h}_{\pi}^{(\infty,N)} \geq \frac{1}{N}\sum_{n=1}^{N}\frac{1}{\lambda_n}, \forall N = 1, 2, ...,\ \lambda_n \in [0, 1],$$
$$n \in \{1, ..., N\}. \square \quad (12)$$

It follows that the minimum time-average AoI, denoted by $\bar{h}_{\text{opt}}^{(\infty,N)}$, cannot be less than either bound, i.e.,

$$\bar{h}_{\text{opt}}^{(\infty,N)} \geq \max\left[\frac{N+1}{2}, \frac{1}{N}\sum_{n=1}^{N}\frac{1}{\lambda_n}\right]. \quad (13)$$

After obtaining two lower bounds in Lemma 3 and 4, combining with the achievable AoI by RR-ONE derived in (10), we can conclude the asymptotic optimality of RR-ONE, and the optimum scaling result follows immediately. Based on Lemma 3, Lemma 4 and Theorem 2, it follows that $\forall N, \lambda_1, ..., \lambda_N$,

$$\max\left[\frac{N+1}{2}, \frac{1}{N}\sum_{n=1}^{N}\frac{1}{\lambda_n}\right] \leq \bar{h}_{\text{opt}}^{(\infty,N)} \leq \frac{1}{N}\sum_{n=1}^{N}\frac{1}{\lambda_n} + \frac{N-1}{2}. \quad (14)$$

For any fixed λ_n, $n = 1, ..., N$, divide both sides of (14) by N, and let N goes to infinity, we obtain

$$\lim_{N\to\infty} \frac{\max\left[\frac{N+1}{2}, \frac{1}{N}\sum_{n=1}^{N}\frac{1}{\lambda_n}\right]}{N} = \frac{1}{2},$$
$$\lim_{N\to\infty} \frac{\frac{1}{N}\sum_{n=1}^{N}\frac{1}{\lambda_n} + \frac{N-1}{2}}{N} = \frac{1}{2}, \quad (15)$$

and therefore (4) follows. Based on the same arguments, (5) follows, concluding the proof of Theorem 2.

Corollary 2: The time-average AoI achieved by a uniformly random scheduling policy with one-packet buffers (UN-ONE) is at least $\bar{h}_{\text{UN}}^{(\infty,N)} \geq N$. □

Proof: The proof (cf. [8]) is based on Lemma 3. ■

Remark 1: Based on Corollary 2, the UN-ONE policy, which in fact can be seen as a performance bound of CSMA scheme without considering the contention time overhead, has a much larger AoI compared with RR-ONE. In particular, when the number of terminals grows large, the UN-ONE policy does not achieve the optimum scaling factor and thus is arbitrarily worse than RR-ONE. □

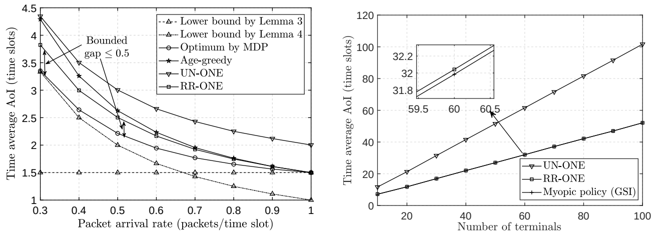

Fig. 3. Performance comparisons with 2 terminals having identical packet arrival rate (x-axis). Performance evaluations with a large number of terminals (right); the arrival rates are uniform-randomly generated in $[0, 1]$.

V. Simulation Results

In this section, computer simulation based experiments are conducted to evaluate the AoI performance of scheduling policies. Relative value iterations for average cost function enable us to obtain the optimum performance numerically by formulating the problem as an Markov decision process (MDP). Note that, similar with most practical applications, the MDP based approach suffers from the curse of dimensionality and hence only small-scale problems can be solved thereby. Nevertheless, we obtain the minimum time-average AoI of a 2-terminal case and compare its performance with RR-ONE. The performance of RR-ONE is obtained by running RR-ONE for 10^5 time slots and calculating the time-average AoI. In addition, we also simulate UN-ONE which schedules a terminal uniformly random at each time slot, and an age-greedy policy which chooses the terminal with the largest AoI. In fact, UN-ONE can be regarded as the CSMA scheme which is shown to be optimal to maximize throughput with greedy sources. The age-greedy policy is found optimal without considering random packet arrivals [5]. It is observed from Fig 3(a) that the performance gap between RR-ONE and the optimum given by numerically solving the MDP is larger with lower packet arrival rates; on the other hand, RR-ONE achieves the optimum when λ approaches one which can be concluded from Lemma 3. Given this gap characterization, since RR-ONE is proved optimal among AIR policies, the optimal scheduling policy with low packet arrival rates must be a non-AIR policy. Specifically, the following intuition explains this. Suppose that the probability of both terminals having arrival packets in the same time slot is negligible when arrival rates are sufficiently low; then the optimal policy is immediately obvious that it should schedule the terminal with packet arrival in each time slot; note that this policy is not an AIR policy because the scheduling decision depends on packet arrivals and hence terminals have to report their queue status. The optimum AoI in this case is also obvious: it should be the same with what is shown in Lemma 4, i.e., completely determined by the inter-arrival time and this can be observed from Fig 3(a). Nevertheless, it is noted that the performance gap is bounded (within 0.5 time slots) in this 2-terminal case by observing the RR-ONE performance and Lemma 4.

We increase the number of terminals and enter the massive IoT regime in Fig. 3(b). The MDP-based optimal solution is computationally intractable in this regime and hence we adopt the myopic policy with global state information (GSI) as an approximation of the optimum. The myopic policy leverages all the global information (though no future knowledge) to make a scheduling decision that minimizes the one-step expected AoI cost in the MDP formulation; by comparing it with RR-ONE helps us to understand how much GSI benefits the AoI performance. Based on our findings, it is even conjectured that the myopic policy is close to optimal. Based on Fig. 3(b), it is shown that the myopic policy with GSI outperforms RR-ONE only slightly, due to the reason that the packet management of using one-packet buffers eliminates most of the randomness of packet arrivals; most packets are dropped by packet management due to staleness and hence their randomness has no effect. The performance of UN-ONE is also shown; it has been proved in Corollary 2 that its linear scaling factor is 1 compared with $1/2$ for RR-ONE; this can be observed in the figure.

VI. Conclusions

In this paper, it is found that with a number of terminals sharing a common wireless uplink and randomly generated status packets at each terminal, the optimal AIR policy to minimize the time-average AoI is RR-ONE, i.e., scheduling terminals in a round-robin fashion and each terminal only retains the most up-to-date packet. In the asymptotic regime where the number of terminals is large, the optimum (among all policies) time-average AoI is proved to scale linearly with the number of terminals. The optimum scaling factor is $\frac{1}{2}$, and RR-ONE achieves it. Moreover, RR-ONE is suitable for decentralized implementation which significantly alleviates the signalling overhead of massive IoT systems.

Acknowledgment

This work is sponsored by NSFC (61701275, 91638204, 61621091) and Hitachi Ltd.

References

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

[2] L. Huang and E. Modiano, “Optimizing age-of-information in a multi-class queueing system,” in *IEEE Int’l Symp. Info. Theory*, pp. 1681–1685, Jun 2015.

[3] Y. Sun, E. Uysal-Biyikoglu, R. Yates, C. E. Koksal, and N. B. Shroff, “Update or wait: How to keep your data fresh,” in *IEEE INFOCOM*, pp. 1–9, April 2016.

[4] E. Najm, R. Yates, and E. Soljanin, “Status updates through M/G/1/1 queues with HARQ,” in *IEEE Int’l Symp. Info. Theory*, pp. 131–135, Jun 2017.

[5] I. Kadota, E. Uysal-Biyikoglu, R. Singh, and E. Modiano, “Minimizing the age of information in broadcast wireless networks,” in *Annu. Allerton Conf. Commun., Control, Comput.*, pp. 844–851, Sep 2016.

[6] Q. He, D. Yuan, and A. Ephremides, “Optimal link scheduling for age minimization in wireless systems,” *IEEE Trans. Inform. Theory*, in press.

[7] Y. P. Hsu, E. Modiano, and L. Duan, “Age of information: Design and analysis of optimal scheduling algorithms,” in *IEEE Int’l Symp. Info. Theory*, pp. 561–565, Jun 2017.

[8] Z. Jiang, B. Krishnamachari, X. Zheng, S. Zhou, and Z. Niu, “Timely status update in massive IoT systems: Decentralized scheduling for wireless uplinks,” *arXiv preprint*, 2018.

[9] B. Melamed and D. Yao, “The ASTA property,” *Advances in Queueing: Theory, Methods and Open Problems*, pp. 195–224, 1995.