

Age Debt: A General Framework For Minimizing Age of Information

Vishrant Tripathi and Eytan Modiano

Laboratory for Information & Decision Systems, MIT

vishrant@mit.edu, modiano@mit.edu

Abstract—We consider the problem of minimizing age of information in general single-hop and multihop wireless networks. First, we formulate a way to convert AoI optimization problems into equivalent network stability problems. Then, we propose a heuristic low complexity approach for achieving stability that can handle general network topologies; unicast, multicast and broadcast flows; interference constraints; link reliabilities; and AoI cost functions. We provide numerical results to show that our proposed algorithms behave as well as the best known scheduling and routing schemes available in the literature for a wide variety of network settings.

I. INTRODUCTION

Many emerging applications require timely delivery of information updates over communication networks. Age of Information (AoI) is a metric that captures precisely this notion of timeliness of received information at a destination [1]–[3]. Unlike packet delay, AoI measures the lag in obtaining information at a destination node, and is therefore suited for applications involving time sensitive updates. Age of information, at a destination, is defined as the time that has elapsed since the last received information update was generated at the source. AoI, upon reception of a new update, drops to the time that has elapsed since generation of the update, and grows linearly otherwise. Over the past few years, there has been a rapidly growing body of work on using AoI as a metric for scheduling policies in wireless networks [4]–[9]. For detailed surveys of AoI literature see [10] and [11].

Scheduling to minimize AoI in single-hop wireless networks has been considered in [4]–[7]. These works prove constant factor optimality of three classes of policies - randomized, max-weight and Whittle index based; under both reliable and unreliable channels. Further, minimizing general cost functions of AoI in single-hop wireless networks has been considered in [9], [12].

AoI minimization over multi-hop wireless networks has also been considered in different settings. In [13], the authors developed stationary randomized policies to minimize weighted-sum-AoI of unicast flows with fixed paths in a multi-hop wireless network. AoI minimization in multihop networks with all-to-all broadcast flows was considered in [8], [14].

We observe that finding low complexity near optimal scheduling and routing schemes for AoI minimization which handle general network topologies, interference constraints,

cost functions, different types of flows and link reliabilities has remained an open problem.

In this work, we develop a unifying framework for making routing and scheduling decisions that optimize AoI in general multihop networks. We do this by transforming AoI-based network optimization problems into network stability problems. Instead of trying to solve for the best scheduling and routing policies directly, we assume that we have access to a set of target values which represent the average age cost for every flow in the network. These target values could be application specific freshness requirements provided by a network administrator, or they could be the solution to an optimization program that optimizes some utility function of the average age costs. Given these targets, we set up a virtual queuing network that is stable if and only if there exists a feasible network control policy that can achieve these targets. Then, we use Lyapunov drift based methods to stabilize this system of virtual queues and achieve the desired target age costs.

In Section II we describe our system model. Then, we introduce notions of age debt and debt-stable scheduling policies in Section III. We develop our heuristic scheduling and routing schemes using Lyapunov drift minimization techniques in Sections IV and V. Finally, we provide brief simulation results in Section VI.

II. SYSTEM MODEL

Consider a network with N nodes connected by a fixed undirected graph $G(V, E)$. An edge (i, j) means that nodes i and j can send packets to one another directly. We assume that at most one update can be sent over an edge in any given time-slot and takes exactly one time-slot to get delivered. We consider general interference constraints between the edges of the graph as well as unreliable transmissions. We further assume that all sources are active, i.e. they can generate fresh updates on demand and that there is no queuing at any node. Each node simply maintains a buffer for the freshest packet of each flow.

A flow is defined as a source node sending updates and a set of destination nodes that are interested in receiving updates from this source. Thus, each flow can be identified by its unique source node. Flows can be of three types depending on the the number of destination nodes:

- 1) **unicast**: the flow has a single destination node.
- 2) **multicast**: the flow has multiple destination nodes, which are a strict subset of the remaining nodes.

This work was supported by NSF Grants AST-1547331, CNS-1713725, and CNS-1701964, and by Army Research Office (ARO) grant number W911NF-17-1-0508.

- 3) **broadcast**: every node other than the source itself is a destination node.

We consider $K \leq N$ flows in this multi-hop network. For simplicity, we represent the source for the k th flow by k . The corresponding destination set for this flow is D_k . For every node j that is a destination for k , we maintain an age of information process $A_{kj}(t)$ which tracks how old the information is at node j about node k .

$$A_{kj}(t+1) = \begin{cases} \min(A_{kj}(t), t - t_g) + 1, & \text{if update generated} \\ & \text{at time } t_g \text{ is delivered at time } t. \\ A_{kj}(t) + 1, & \text{if no new delivery at time } t. \end{cases} \quad (1)$$

We associate a monotone increasing age cost function for each source k and corresponding destination $j \in D_k$ denoted by $f_{kj}(\cdot)$. Using these age cost functions, we maintain the effective age processes $B_{kj}(t) \triangleq f_{kj}(A_{kj}(t))$.

In general, a control policy needs to specify not only which links should be scheduled in each time-slot but also which flows should be transmitted along each link. We enumerate the set of all possible interference free choices of links and corresponding flows in the set \mathcal{S} . Thus, a member of set \mathcal{S} contains a subset of links and corresponding flows which can be sent on these links in a single time-slot without interference. A valid network control policy must choose an action that is a member of the set \mathcal{S} in every time-slot.

The typical goal of AoI-based scheduling and routing design for multihop networks is to minimize the time average of the expected age costs summed across flows:

$$\pi^* = \underset{\pi}{\operatorname{argmin}} \left(\lim_{T \rightarrow \infty} \mathbb{E} \left[\frac{1}{T} \sum_{t=1}^T \sum_{k=1}^K \sum_{j \in D_k} B_{kj}(t) \right] \right), \quad (2)$$

where $\pi(t) \in \mathcal{S}, \forall t, \pi$.

Next, we introduce the notions of age-achievability and age debt virtual queues and show how stabilizing this network of virtual queues leads to minimization of AoI.

III. AGE DEBT

We start by assuming that we have been given a target value of time average age cost for each source-destination pair; denoted by α_{kj} for the pair (k, j) . We aggregate the target values associated with each source-destination pair in the vector α . For any such target vector α , we define of age-achievability below.

Definition A vector α is **age-achievable** if there exists a feasible network control policy π such that

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T B_{kj}^\pi(t) \leq \alpha_{kj}, \forall j \in D_k, \forall k \text{ w.p. } 1. \quad (3)$$

In other words, a vector α is age-achievable if the time-average of the effective age process for *every* source-

destination pair (k, j) is upper bounded by the target value α_{kj} , under some feasible network control policy.

Note that the combination of general cost functions and achievability targets allows us to capture very general freshness requirements which might be useful in practical system specifications. For example, if an application requires that the empirical distribution of the age process $A_{kj}(t)$ should satisfy $\mathbb{P}(A_{kj}(t) \geq M) \leq \epsilon$, then we can capture this by setting the cost function $f_{kj}(h) = \mathbf{1}_{h \geq M}$ and the corresponding target to be $\alpha_{kj} = \epsilon$.

We now define a set of virtual queues called age-debt queues for every source-destination pair (k, j) . These queues measure how much the effective age process exceeds its target value α_{kj} , summed over time. Our definition of debt is motivated by the notion of throughput debt as introduced in [15].

Definition Given a target vector α , the **age debt queue** for source-destination pair kj at time t under a policy π , given by $Q_{kj}^\pi(t)$, evolves as

$$Q_{kj}^\pi(t+1) = \left[Q_{kj}^\pi(t) + B_{kj}^\pi(t+1) - \alpha_{kj} \right]^+, \forall j \in D_k, \quad \text{and } \forall k \in \{1, \dots, K\}. \quad (4)$$

To complete the definition, each age debt queue starts at zero, i.e. $Q_{kj}^\pi(0) = 0, \forall j, k$. Here $[x]^+ \triangleq \max(x, 0)$.

We now introduce a notion of stability for these age debt queues. This is similar to how mean rate stability is typically defined in queueing networks [16].

Definition We say that the network of age debt queues is **stable** under a policy π and a given target vector α if the following condition holds:

$$\lim_{T \rightarrow \infty} \mathbb{E} \left[\sum_{k=1}^K \sum_{j \in D_k} \frac{Q_{kj}^\pi(T)}{T} \right] = 0, \quad (5)$$

where the expectation is taken over the randomness in the channel processes.

We also establish an equivalence relationship between age-achievability of a vector α and the stability of the corresponding network of age debt queues.

Lemma 1: A target vector α is age-achievable if and only if there exists a network control policy π , possibly dependent on α , that stabilizes the network of source-destination age debt queues.

Proof: See Appendix A in the technical report [17]. ■

Next, we define a debt-stable scheduling policy. Such a policy takes a target vector α as an input and stabilizes the network of corresponding age debt queues.

Definition A **debt-stable** scheduling policy π stabilizes the set of age-debt queues for any given target vector α that is age-achievable.

The notions introduced until now effectively allow us to convert the minimum age cost problem introduced in (2) into a network stability problem. Suppose π^* is a solution to the optimization problem (2). Further, suppose that expected time average of the k th effective age process under π^* is given by

$$\lim_{T \rightarrow \infty} \mathbb{E} \left[\frac{1}{T} \sum_{t=1}^T B_{kj}^{\pi^*}(t) \right] = \alpha_{kj}^*, \forall (k, j). \quad (6)$$

Clearly, if we have oracle access to an optimal age cost vector α^* and know how to design a debt-stable policy then we can perform minimum age cost scheduling. If the debt-stable policy is much lower in computational complexity than solving (2) directly, then we can also solve (2) at the same lower complexity (assuming oracle access to α^*). We now discuss a heuristic approach to designing debt-stable policies.

IV. LYAPUNOV DRIFT APPROACH

A. Single-Hop Broadcast

We first consider the special case of single-hop broadcast networks. This setting is easier to analyze since it only requires scheduling and no routing. It also highlights key structural properties of our proposed policy.

Consider a N node star network where each of the nodes $1, \dots, N-1$ has an edge to node N . These nodes wish to send packets to the central node N . Due to broadcast interference constraints, only one node can transmit in any given time-slot. Since the destination for every flow is N , we can drop the destination in our notation. The age evolution is given by

$$A_i^{\pi}(t+1) = \begin{cases} A_i^{\pi}(t) + 1, & \text{if } i \notin \pi(t) \text{ or } c_i(t) = 0 \\ 1, & \text{if } i \in \pi(t) \text{ and } c_i(t) = 1. \end{cases} \quad (7)$$

Here $\pi(t)$ is the source scheduled in time-slot t and $c_i(t)$ is an indicator variable denoting edge reliability between node i and node N at time t . Given an age-cost function $f_i(A_i(t))$ and a corresponding target value α_i , the debt queue evolution for node i is given by:

$$Q_i^{\pi}(t+1) = \left[Q_i^{\pi}(t) + f_i(A_i^{\pi}(t+1)) - \alpha_i \right]^+. \quad (8)$$

Given a target vector α , we will use a Lyapunov drift based scheduling scheme to try and achieve debt stability. To do so, we first define a Lyapunov function for our system of virtual queues:

$$L(t) \triangleq \sum_{i=1}^{N-1} Q_i^2(t). \quad (9)$$

Given the entire network state $\Phi(t)$ (all AoI and debt queue values), we define the *age debt scheduling policy* π^{AD} as:

$$\pi^{\text{AD}}(t) = \underset{a \in \mathcal{S}}{\operatorname{argmin}} \left(\mathbb{E}[L(t+1) - L(t) | \Phi(t)] \right), \quad (10)$$




Fig. 1

where the expectation is taken over the randomness in channel reliabilities $c(t)$. The following lemma describes what the drift minimizing actions look like in practice.

Lemma 2: Suppose that the links between each source i and the destination N are i.i.d. Bernoulli w.p. p_i in every time-slot. Further, if each age cost function $f_i(\cdot)$ is upper bounded by a constant D , then the policy $\pi(t)$ below minimizes an upper bound on the Lyapunov drift in every time-slot.

$$\pi(t) = \underset{i \in 1, \dots, N-1}{\operatorname{argmax}} \left(p_i Q_i(t) (f_i(A_i(t) + 1) - f_i(1)) \right). \quad (11)$$

Proof: See Appendix B in the technical report [17]. ■

In other words, a drift minimizing policy chooses the source with the largest product of link reliability, current age debt and current age cost. While we use the boundedness assumption to establish this structural property of the drift minimizing action, it is *not required* to implement the age-debt policy in practice. Further, the boundedness assumption is mild, since the upper bound D can be chosen to be arbitrarily large. For example, it can be chosen to be in the order of years for a system where AoI is in the order of seconds or milliseconds.

Note that to compute $\pi^{\text{AD}}(t)$, the scheduler needs to iterate over the set of sources only once. So the per slot computational complexity of this policy grows linearly in N . This is similar to the complexity of the Whittle index policy proposed in [4], [9] and the max-weight policies proposed in [4], [6]. By contrast, a dynamic programming approach to solve (2) directly has per slot computational complexity that grows exponentially in N . This highlights the key strength of our approach. If the scheduler has some way to set the targets for each source optimally, then the age debt policy is a good low complexity heuristic for age minimization.

B. General Networks

The general multihop setting is more challenging. Simply using one-slot Lyapunov drift to try and achieve debt stability does not work directly. We highlight this with a simple example.

Consider the three node network described in Figure 1 with a single unicast flow from node 1 to node 3. The interference constraint enforces that only one of the two edges a and b can be activated in any time-slot. Suppose that we are interested in minimizing the time average of the age process $A_{13}(t)$. Given a target value α_{13} , we set up the age debt queue as follows:

$$Q_{13}^{\pi}(t+1) = \left[Q_{13}^{\pi}(t) + A_{13}^{\pi}(t+1) - \alpha_{13} \right]^+. \quad (12)$$

We will try to use the one slot Lyapunov drift minimizing policy to stabilize $Q_{13}(t)$ in this network. To do so, we solve the following optimization in every time-slot:

$$\pi^{\text{AD}}(t) = \underset{x \in \{a, b\}}{\operatorname{argmin}} \left(\mathbb{E}[Q_{13}^2(t+1) - Q_{13}^2(t) | \Phi(t)] \right). \quad (13)$$

At $t = 1$, activating either edge a or edge b has no effect on the debt $Q_{13}(2)$ since node 2 does not have any packet from node 1. If we break ties in favour of edge b , then it is activated but no new packet is delivered to node 3. At $t = 2$, since node 2 still does not have any new update from node 1, no action taken can affect the debt $Q_{13}(3)$. Using the same tie-break rule, we would again schedule edge b . This process keeps on repeating and the age debt queue Q_{13} blows up irrespective of the value of α_{13} , even though the age optimal policy in this setting is to simply alternate between a and b in every time-slot.

The example above illustrates why one-slot Lyapunov drift based techniques fail in stabilizing debt queues in multihop networks. The policy designer using Lyapunov drift is constrained to optimizing *only one time-step into the future*. So, if every possible scheduling and routing action has no effect on the age debt queues in the immediate next time-slot, the one step drift minimizing procedure does not provide any information on which action should be chosen to stabilize the network.

This suggests that to be able to use one-slot drift minimizing techniques for stability there should be a virtual queue for every intermediate node that tracks both the current age debt at the destination and the potential reduction in debt at the destination upon forwarding a fresh packet. If we can set up such queues, then large values of debt at intermediate nodes would lead to fresh packets being sent to the next hops via one-slot drift minimizing actions, eventually reaching the destination and stabilizing the age debt queues.

Let $Q_{kj}^i(t)$ denote such a debt queue corresponding to flow (k, j) at an intermediate node i . These additional queues at every intermediate node combined with the original debt queues form our virtual network. The Lyapunov function that we use for scheduling and routing is given by:

$$L(t) \triangleq \sum_{k=1}^K \sum_{j \in D_k} \left(Q_{kj}^2(t) + \sum_{i \notin D_k, i \neq k} (Q_{kj}^i(t))^2 \right) \quad (14)$$

The Age Debt scheduling and routing policy is to choose the activation set and corresponding flows that minimizes the expected Lyapunov drift.

$$\pi^{\text{AD}}(t) = \underset{a \in \mathcal{S}}{\operatorname{argmin}} \left(\mathbb{E}[L(t+1) - L(t) | \Phi(t)] \right), \quad (15)$$

where the expectation is taken over the randomness in link reliabilities and $\Phi(t)$ represents the entire network state (all AoI and debt queue values).

C. Intermediate Debt Queues

We now discuss how to set up the age debt queues $Q_{kj}^i(t)$ for intermediate nodes to augment the original network of queues. Note that there are no intermediate nodes for broadcast flows since every node other than the source is a destination.

Consider a source-destination pair (k, j) for a unicast/multicast flow k and an intermediate node i that is not a destination for the flow originating at k . We want to set up the age debt queue $Q_{kj}^i(t)$ at i for the pair (k, j) . We maintain

an age process for flow k at node i , even though there is no associated cost or target value for this age process.

$$A_{ki}(t+1) = \begin{cases} \min(A_{ki}(t), t - t_g) + 1, & \text{if update generated} \\ & \text{at time } t_g \text{ is delivered at time } t. \\ A_{ki}(t) + 1, & \text{if no new delivery at time } t. \end{cases} \quad (16)$$

Here $A_{ki}(t)$ measures how old the information at node i is regarding node k . We split the debt queue's evolution into two cases.

Case 1: When node i forwards a flow k packet on a set of adjacent links L . Let h_{ij}^L be the minimum number of hops it takes to reach node j from node i , where the first hop can only include edges in the set L . Here, h_{ij}^L measures the minimum delay with which the packet that was forwarded by i gets delivered at j . The age debt queue $Q_{kj}^i(t)$, when node i is forwarding a flow k packet along the link set L , evolves as:

$$Q_{kj}^i(t+1) = \left[Q_{kj}^i(t) + f_{kj}(\min\{A_{ki}(t), A_{kj}(t)\} + h_{ij}^L) - \alpha_{kj} \right]^+ \quad (17)$$

This measures the most optimistic change in age debt possible at the destination using the current packet transmission from node i .

Case 2: When node i does not forward a packet from node k along any of its adjacent edges, then the age debt queue evolves as below.

$$Q_{kj}^i(t+1) = \left[Q_{kj}^i(t) + B_{kj}(t+1) - \alpha_{kj} \right]^+ \quad (18)$$

This means that the intermediate queue simply tracks the change in debt at the destination when it is not forwarding a relevant packet. If the destination is not receiving fresh packets from anywhere in the network then this would increase the intermediate debt queues.

Thus, the debt at an intermediate node i for a source-destination pair (k, j) blows up if either the destination has not received fresh packets for a long time and node i did not forward any packets from k (i.e. (18)) or if node i keeps forwarding stale packets from k (i.e. (17)). A drift minimizing policy will then try to ensure that either the destination debt queue is small, or node i forwards fresh packets of flow k towards the destination.

V. CHOOSING TARGET VECTORS

In the preceding sections, we have developed a general framework of age achievability where given a target average age cost for every source-destination pair, we formulate a corresponding network stability problem and attempt to solve it via one slot Lyapunov drift minimization. In this section, we discuss how to choose the right target vectors, such that they lead to minimum sum age cost.

In the absence of an optimization oracle that provides access to α^* or a system administrator who specifies average age cost targets based on the underlying application requirements, we

develop a simple heuristic to dynamically update α in order to optimize utility based on the state of the underlying debt queues.

The following optimization problem needs to be solved to find the best target vector α^* .

$$\begin{aligned} \underset{\alpha}{\operatorname{argmin}} \quad & \left(\sum_{k=1}^K \sum_{j \in D_k} \alpha_{kj} \right), \\ \text{s.t. } & \alpha \text{ is age-achievable.} \end{aligned} \quad (19)$$

Note that this problem has the same optimal value as (2). One way to dynamically set the target vectors is to take a flow control approach for solving the optimization problem (19), similar to [18]. Algorithm 1 describes the details.

Algorithm 1: Age Debt - Flow Control

Input : parameter $V > 0$, upper bound α_{\max} ,
initialization $\alpha^{(1)}$

```

1 while  $t \in 1, \dots, T$  do
2   Use  $\alpha^{(t)}$  to update debt queue values at time  $t$ 
3   Update  $\alpha$  by solving the optimization below:
4

$$\alpha^{(t+1)} = \underset{\alpha}{\operatorname{argmin}} \left( \sum_{k=1}^K \sum_{j \in D_k} V \alpha_{kj} - \alpha_{kj} Q_{kj}(t) \right),$$


$$\text{s.t. } \alpha \geq 1, \alpha \leq \alpha_{\max}.$$

5   Use  $\alpha^{(t+1)}$  to compute the scheduling and routing
    decision that minimizes drift:

$$\pi(t) = \underset{\alpha \in S}{\operatorname{argmin}} \left( \mathbb{E}[L(t+1) - L(t) | \Phi(t)] \right)$$

6 end
```

The flow control based age debt policy tries to tradeoff between the stability of the queueing network and the optimization of targets using a parameter $V > 0$. In every time-slot, the flow control optimization sets the target α for the next time-slot and then the scheduling and routing decisions are computed by minimizing Lyapunov drift.

The update optimization in step 4 of Algorithm 1 can be simplified to the rule below:

$$\alpha_{kj}^{(t+1)} = \begin{cases} \alpha_{\max}, & \text{if } Q_{kj}(t) > V \\ 1, & \text{if } Q_{kj}(t) \leq V, \end{cases} \forall (k, j). \quad (20)$$

When current debts are high, future targets are set to be high pushing the debts lower. Similarly, when the current debts are low, future targets are also set low, pushing the debts higher. The parameter V decides the threshold between high and low values of the debt queues.

We also develop a separate gradient descent based approach to dynamically update the target vector α . Details of this algorithm are included in our technical report [17].

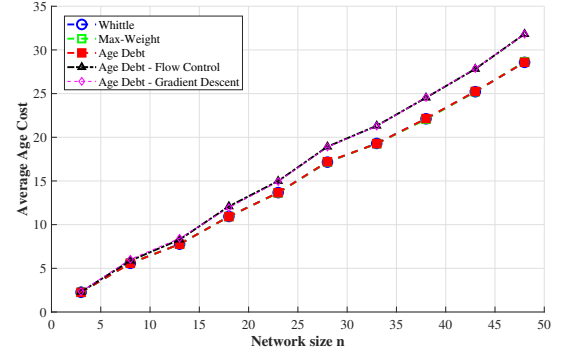


Fig. 2: Weighted-sum AoI minimization in broadcast networks with unreliable channels

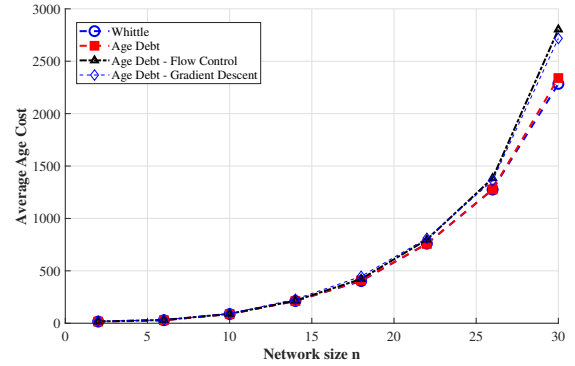


Fig. 3: Functions of Age minimization in broadcast networks with reliable channels

VI. SIMULATIONS

First, we consider the weighted-sum AoI problem in broadcast networks with unreliable channels. There are N nodes in the network and the weight of the i th node w_i is set to i/N . Link connection probabilities are chosen uniformly from the set $[0.6, 1]$. Figure 2 plots the performance of the age debt policy along with the max-weight and Whittle index policies proposed in [4] against the system size. We observe that when the age debt policy is provided the max-weight average cost as the target vector, it replicates near optimal performance of max-weight. Further, the flow control and gradient descent versions of age debt only have a small gap to the max-weight/Whittle policies despite not having access to α beforehand.

Next, we consider general functions of age minimization in the single-hop wireless broadcast setting. There are N nodes in the network and the cost of AoI for each node is chosen from the set of functions $\{15A(t), e^{A(t)}, (A(t))^2 \text{ and } (A(t))^3\}$. Figure 3 plots the performance of the age-debt policy and its variants along with the Whittle index policy proposed in [9] as the system size increases. As for the linear AoI case, we observe that age debt is able to replicate the Whittle policy's performance when provided its average cost as the target vector. The flow control and gradient descent variants are also only a small gap away in performance without knowing α beforehand.

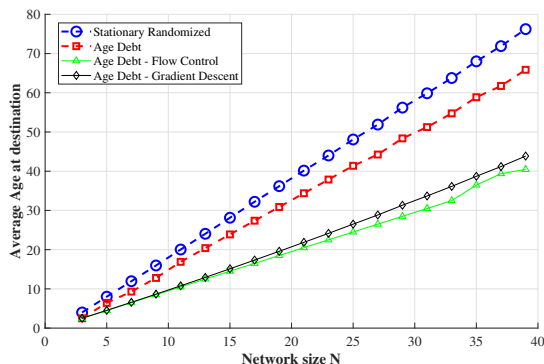


Fig. 4: Age minimization of a single unicast flow in line networks (neighboring nodes interfere)




Fig. 5: Age minimization of broadcast flows in multihop networks with 5 nodes

Next, we consider scheduling for a single unicast flow on the line network, as studied in [13]. Consider N nodes arranged in a line network from 1 to N . Node 1 wants to sent packets to node N , however not all nodes can transmit simultaneously. We consider a simple interference constraint - in any given time-slot either all even numbered nodes or all odd numbered nodes can forward packets. This ensures that no two adjacent nodes send interfering transmissions. Figure 4 plots the performance of age-debt and its flow-control and gradient-descent variations along with the optimal stationary randomized policy proposed in [13]. We observe that age-debt outperforms the stationary randomized policy despite using its average costs as the target vector. The dynamic variants of age-debt significantly outperform the stationary randomized policy. We also note that the gap in performance would increase in settings where multiple paths are available which age-debt can utilize for routing, unlike the stationary randomized approach.

Finally, we consider average age minimization for all-to-all broadcast flows in multihop networks similar to [8]. We consider all possible connected network topologies with 5 or 6 nodes (a total of 133 graphs). Figure 5 plots the performance of the age-debt policy and its variants along with the near optimal minimum connected dominating set (MCDS) based scheme proposed in [8] for each of these networks. The x-axis represents the graph labels numbered from 1 to 133, sorted according to the average age achieved by the MCDS scheme.

We observe that age-debt achieves the same performance

as the MCDS scheme when provided its average cost as the target vector. Further, age-debt with flow control achieves performance that is very close to that of the MCDS scheme without requiring knowledge of α .

Note that the complexity of implementing the flow-control scheme is polynomial in the network size per time-slot. This suggests that age-debt and its variants are a good candidate for low complexity near optimal age scheduling in general networks.

Interesting directions of future work involve proving performance bounds on the age debt policy, implementing it in a distributed fashion, and considering stochastic arrivals and time-varying topologies in the underlying network.

REFERENCES

- [1] S. Kaul, R. Yates, and M. Gruteser, "Real-time status: How often should one update?," in *Proc. IEEE INFOCOM*, pp. 2731–2735, 2012.
- [2] C. Kam, S. Kompella, and A. Ephremides, "Age of information under random updates," in *Proc. IEEE Int. Symp. Information Theory (ISIT)*, pp. 66–70, 2013.
- [3] 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 Trans. Information Theory*, vol. 63, pp. 7492–7508, Nov. 2017.
- [4] 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 Trans. Netw.*, vol. 26, no. 6, pp. 2637–2650, 2018.
- [5] I. Kadota, A. Sinha, and E. Modiano, "Scheduling algorithms for optimizing age of information in wireless networks with throughput constraints," *IEEE/ACM Trans. Netw.*, vol. 27, no. 4, pp. 1359–1372, 2019.
- [6] R. Talak, S. Karaman, and E. Modiano, "Optimizing information freshness in wireless networks under general interference constraints," in *Proc. ACM Int. Symp. Mobile Ad Hoc Netw. Comput. (MobiHoc)*, pp. 61–70, 2018.
- [7] A. Maatouk, S. Kriouile, M. Assaad, and A. Ephremides, "On the optimality of the whittle's index policy for minimizing the age of information," *arXiv preprint arXiv:2001.03096*, 2020.
- [8] S. Farazi, A. G. Klein, J. A. McNeill, and D. R. Brown, "On the age of information in multi-source multi-hop wireless status update networks," in *Proc. IEEE Int. Workshop Signal Process. Adv. Wireless Commun. (SPAWC)*, pp. 1–5, 2018.
- [9] V. Tripathi and E. Modiano, "A whittle index approach to minimizing functions of age of information," in *Proc. 57th Allerton Conf. Commun. Control Comput.*, pp. 1160–1167, IEEE, 2019.
- [10] 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.
- [11] Y. Sun, I. Kadota, R. Talak, and E. Modiano, "Age of information: A new metric for information freshness," *Synthesis Lectures on Communication Networks*, vol. 12, no. 2, pp. 1–224, 2019.
- [12] P. R. Jhunjunwala and S. Moharir, "Age-of-information aware scheduling," in *Proc. IEEE SPCOM*, 2018.
- [13] R. Talak, S. Karaman, and E. Modiano, "Minimizing age-of-information in multi-hop wireless networks," in *Proc. 55th Allerton Conf. Commun. Control Comput.*, pp. 486–493, IEEE, 2017.
- [14] S. Farazi, A. G. Klein, and D. R. Brown, "Fundamental bounds on the age of information in general multi-hop interference networks," in *Proc. IEEE INFOCOM AoI Workshop*, pp. 96–101, 2019.
- [15] I. Hou, V. Borkar, and P. Kumar, "A theory of qos for wireless," in *Proc. Infocom*, pp. 486–494, IEEE, 2009.
- [16] M. J. Neely, "Stability and capacity regions or discrete time queueing networks," *arXiv preprint arXiv:1003.3396*, 2010.
- [17] V. Tripathi and E. Modiano, "Age debt: A general framework for minimizing age of information," technical report, 2020. <https://arxiv.org/abs/2101.10225>.
- [18] L. Georgiadis, M. J. Neely, and L. Tassiulas, *Resource allocation and cross-layer control in wireless networks*. Now Publishers Inc, 2006.