

Update or Wait: How to Keep Your Data Fresh

Yin Sun[†], Elif Uysal-Biyikoglu[‡], Roy Yates^{*}, C. Emre Koksal[†], and Ness B. Shroff^{†§}

[†]Dept. of ECE, [§]Dept. of CSE, The Ohio State University, Columbus, OH

[‡]Dept. of EEE, Middle East Technical University, Ankara, Turkey

^{*}Dept. of ECE, Rutgers University, North Brunswick, NJ

Abstract—In this work we study how to manage the freshness of status updates sent from a source to a remote monitor via a network server. A proper metric of data freshness at the monitor is the *status update age*, which is defined as how old the freshest update is since the moment this update was generated at the source. A logical policy is the *zero wait policy*, i.e., the source submits a fresh update once the server is free, which achieves the maximum throughput and the minimum average delay. Surprisingly, this zero wait policy does not always minimize the average age. This motivates us to study how to optimally control the status updates to keep data fresh and to understand when the zero wait policy is optimal. We introduce a penalty function to characterize the level of “dissatisfaction” on data staleness, and formulate the average age penalty minimization problem as a constrained semi-Markov decision process (SMDP) with an uncountable state space. Despite of the difficulty of this problem, we develop efficient algorithms to find the optimal status update policy. We show that, in many scenarios, the optimal policy is to wait for a certain amount of time before submitting a new update. In particular, the zero wait policy can be far from the optimum if (i) the penalty function grows quickly with respect to the age, and (ii) the update service times are highly random and positive correlated. To the best of our knowledge, this is the first optimal control policy which is proven to minimize the age-of-information in status update systems.

I. INTRODUCTION

In recent years, the proliferation of mobile devices and applications has significantly boosted the need for real-time information updates, such as news, weather reports, email notifications, stock quotes, social updates, mobile ads, etc. Timely status updates are also critical in network-based monitoring and control systems, including sensor networks used in temperature and air pollution monitoring, surround monitoring in autonomous vehicles, and phasor data updates in power grid stabilization systems.

A common need in these applications is to maximize the freshness of the data at the monitor. In light of this, a metric called *status update age*, or simply *age*, was defined in [1]. At time t , if the freshest update at the monitor has a time stamp $U(t)$, the status update age is $\Delta(t) = t - U(t)$. Hence, the age is the time elapsed since the freshest packet was generated.

Most existing research on status update age focuses on an “enqueue-and-forward” model [1]–[7], where status updates are randomly generated or arrive at a source node. The source enqueues these updates and forwards them later to a remote monitor through a network cloud. It is worth noting that the goal of age minimization differs from those of throughput

This work was supported in part by DTRA grant HDTRA1-14-1-0058, NSF grants CNS-1446582, CNS-1409336, CNS-1518829, CNS-1514260, CNS-1422988, and CNS-1054738, ARO grant W911NF-14-1-0368, and ONR grant N00014-15-1-2166. E. Uysal-Biyikoglu was supported in part by TUBITAK and in part by a Science Academy BAGEP award.




Fig. 1. System model.

maximization and delay minimization: A high update frequency improves the system utilization, but may also induce a large waiting time in the queue which in turn increases the age; on the other hand, a low update frequency can reduce the queueing delay, but the monitor may end up having stale status information due to not enough updates [1], [5], [7]. In [6], it was found that a good policy is to discard the old packets waiting in the queue if a new sample arrives, which can greatly reduce the impact of queueing delay.

In this paper, we study a “generate-at-will” model depicted in Fig. 1. In this model, the source keeps monitoring the network server’s idle/busy state, and in contrast to [1]–[7], is able to generate status updates at any time by its own will. Hence, no updates need to be generated when the server is busy, which completely eliminates the waiting time in the queue and hence the queue in Fig. 1 is always empty. A simple *zero wait policy*, also known as the *work-conserving policy* in queueing theory, that submits a fresh update once the server becomes idle, achieves the maximum throughput and the minimum average delay. Surprisingly, this zero wait policy does **not** always minimize the average age of the information [8]. The following example reveals the reason behind this phenomenon:

Example: Suppose that the source submits a stream of update packets to a remote monitor. The service times of these updates form a periodic sequence

$$0, 0, 2, 2, 0, 0, 2, 2, 0, 0, 2, 2, \dots$$

Suppose that update 1 is generated and submitted at time t and delivered at time t . Under the zero wait policy, update 2 is also generated at time t and delivered at time t . However, despite of its short service time, update 2 has not brought any fresher information to the monitor after update 1 is delivered, because both updates are sampled at time t . Therefore, the potential benefit of the zero service time of update 2 is wasted! This issue occurs periodically over time: Whenever two consecutive updates have zero service time, the second update of the two is wasted. Therefore, 1/4 updates in this sequence are wasted in the zero wait policy!

For comparison, consider a policy that waits for 0.5 seconds after each update with zero service time, and does not wait




Fig. 2. Evolution of the status update age $\Delta(t)$.

after each update with a service time of 2 seconds. The average age of the zero wait policy is 2 seconds, and the average age of the second policy is 1.85 seconds.¹ Hence, the zero wait policy is not optimal.

This example points out a key difference between status update systems and data communication systems: In status update systems, an update packet is useful only if it carries fresh information to the monitor; however, in communication systems, all packets are equally important. While the theory of data communications is quite mature, the optimal control of status updates remains open.

For a source that can insert waiting times between updates, the aim of this paper is to answer the following questions: How to optimally submit update packets to maximize data freshness at the monitor? When is the zero wait policy optimal? To that end, the following are the key contributions of this paper:

- We generalize existing status update studies by introducing two new features: age penalty functions and non-*i.i.d.* service processes. We define an age penalty function $g(\Delta)$ to characterize the level of “dissatisfaction” for data staleness, where $g(\cdot)$ is *measurable, non-negative, and non-decreasing*, which is determined by the specific application. The update service process is modeled as a stationary ergodic Markov chain with an uncountable state space, which generalizes the *i.i.d.* service processes studied in previous work [1]–[8], [10].
- We formulate the average age penalty minimization problem as a constrained semi-Markov decision process (SMDP) with an uncountable state space. Despite of the difficulty of this problem, we manage to solve it by a divide-and-conquer approach: We first prove that there exists a stationary randomized policy that is optimal for this problem (Theorem 1). Further, we prove that there exists a stationary deterministic policy that is optimal for this problem (Theorem 2). Finally, we develop a low-complexity algorithm to find the optimal stationary deterministic policy (Theorem 3). To the best of our knowledge, this is the first optimal control policy which is proven to minimize the age-of-information (i.e., maximize data freshness) in status update systems.
- We further investigate when the zero wait policy is optimal. For the special case of proportional penalty function and *i.i.d.* service times, we devise a simpler solution to minimize the average age (Theorem 4). This solution explicitly characterizes when the zero wait policy

is optimal, and when it is not. For general age penalty functions and correlated service processes, sufficient conditions for the optimality of the zero wait policy are provided (Lemma 3).

- Our theoretical and simulation results demonstrate that, in many scenarios, the optimal policy is to wait for a certain amount of time before submitting a new update. In particular, the zero wait policy can be far from optimality if (i) the penalty function grows quickly with respect to the age, and (ii) the update service times are highly random and positive correlated.

II. SYSTEM MODEL AND PROBLEM FORMULATION

We consider a system depicted in Fig. 1, where a source generates status update packets and sends them to a remote monitor through a network server. This server provides First-Come First-Served (FCFS) service to the submitted update packets. The service of an update packet is considered complete, when it is successfully received by the monitor. After that, the server becomes available for sending the next packet.

The source generates and submits updates at successive times S_0, S_1, \dots . Update i , submitted at time S_i , is delivered at time $D_i = S_i + Y_i$, where $Y_i \geq 0$ is the service time of update i . Suppose that update 0 is submitted to an idle server at time $S_0 = -Y_0$ and delivered at $D_0 = 0$, as shown in Fig. 2. The source has access to the idle/busy state of the server and is able to generate updates at any time by its own will. Hence, the source should not generate update $i + 1$ when the server is busy processing update i , because this will incur an unnecessary waiting time in the queue. After update i is delivered at time D_i , the source may introduce a *waiting time* $Z_i \in [0, M]$ before submitting update $i + 1$ at time $S_{i+1} = D_i + Z_i$, where M represents the maximum amount of waiting time allowed by the system. The source can switch to a low-power sleep mode during $[D_i, S_{i+1})$. We assume that the service process (Y_0, Y_1, \dots) is a *stationary and ergodic Markov chain* with a possibly uncountable state space and a positive mean $0 < \mathbb{E}[Y_i] < \infty$, which generalizes the *i.i.d.* service processes in previous studies [1]–[8], [10]. The ergodicity of this Markov chain is assumed in the sense of ergodic theory [11], which allows the Markov chain to be periodic.² This Markovian service process model is introduced to study the impact of temporal-correlation on the optimality of the zero wait policy. In Section IV, we will see that zero wait policy is near-optimal when the service process is negative correlated; and can be far from optimality when the service process is positive correlated.


At any time t , the monitor’s most recently received update packet is time-stamped with $U(t) = \max\{S_i : D_i \leq t\}$. The *status update age* $\Delta(t)$ is defined as [1]

$$\Delta(t) = t - U(t), \quad (1)$$


which is also referred to as *age-of-information* or simply *age*. As shown in Fig. 2, the age $\Delta(t)$ is a stochastic process that increases linearly with t between updates, with downward jumps occurring when updates are delivered.

²The results in this paper can be readily extended to a more general model where the Markov chain (Y_0, Y_1, \dots) has a longer memory, i.e., the sequence (W_0, W_1, \dots) forms a Markov chain, with W_i defined as $W_i = (Y_i, Y_{i+1}, \dots, Y_{i+k})$ for some finite k .

¹To compute the average age, one need to use (8), (12), and refer to Fig. 2. Its matlab code is provided in Appendix A of [9].



(a) An exponential age penalty function $g_1(\Delta) = e^{0.2\Delta} - 1$.



(b) A stair-step age penalty function $g_2(\Delta) = \lfloor \Delta \rfloor$.

Fig. 3. Evolutions of two age penalty functions.

We introduce an age penalty function $g(\Delta)$ to represent the level of “dissatisfaction” for data staleness or the “need” for new information update, where the function $g : [0, \infty) \rightarrow [0, \infty)$ is assumed to be *measurable, non-negative, and non-decreasing*. Some examples of $g(\cdot)$ are power function $g(\Delta) = \Delta^a$, exponential function $g(\Delta) = e^{a\Delta}$, and the stair-step function $g(\Delta) = \lfloor a\Delta \rfloor$, where $a \geq 0$ and $\lfloor x \rfloor$ is the largest integer no greater than x . Two examples of age penalty functions are depicted in Figure 3. Note that this age penalty model is quite general, which allows $g(\cdot)$ to be discontinuous and non-convex.

To analyze the average age penalty, we decompose the area under the curve $g(\Delta(t))$ into a sum of disjoint components: Consider the time interval $[0, D_n]$, where $D_n = \sum_{i=0}^{n-1} (Y_i + Z_i)$. In this interval, the area under $g(\Delta(t))$ can be seen as the concatenation of the areas Q_i , $0 \leq i \leq n-1$, such that

$$\int_0^{D_n} g(\Delta(t)) dt = \sum_{i=0}^{n-1} Q_i,$$

where

$$Q_i = \int_{Y_i}^{Y_i + Y_i + Z_i} g(\tau) d\tau.$$

Let us define

$$q(y, z, y') = \int_y^{y+z+y'} g(\tau) d\tau. \quad (2)$$

Then, Q_i can be expressed as $Q_i = q(Y_i, Z_i, Y_{i+1})$. We assume that

$$\mathbb{E} [q(Y_i, M, Y_{i+1})] < \infty. \quad (3)$$

Our goal is to minimize the average age penalty by controlling the sequence of waiting times (Z_0, Z_1, \dots) . Let $\pi \triangleq$

(Z_0, Z_1, \dots) denote a status update policy. We consider the class of *causally feasible* policies, in which control decisions are made based on history and current information of the system, as well as the distribution of the service process (Y_0, Y_1, \dots) . Specifically, Z_i is determined based on the past realizations of (Y_0, Y_1, \dots, Y_i) , without using the realizations of future service times $(Y_{i+1}, Y_{i+2}, \dots)$; but the conditional distribution of $(Y_{i+1}, Y_{i+2}, \dots)$ based on (Y_0, Y_1, \dots, Y_i) is available. Let Π denote the set of all causally feasible policies satisfying $Z_i \in [0, M]$ for all i .

The average age penalty can be represented by

$$\limsup_{n \rightarrow \infty} \frac{\mathbb{E} \left[\int_0^{D_n} g(\Delta(t)) dt \right]}{\mathbb{E} [D_n]}.$$

Using this, the stochastic optimization problem for minimizing the average age penalty can be formulated as

$$\begin{aligned} \bar{g}_{\text{opt}} &= \min_{\pi \in \Pi} \limsup_{n \rightarrow \infty} \frac{\mathbb{E} \left[\sum_{i=0}^{n-1} q(Y_i, Z_i, Y_{i+1}) \right]}{\mathbb{E} \left[\sum_{i=0}^{n-1} (Y_i + Z_i) \right]} \\ \text{s.t. } &\liminf_{n \rightarrow \infty} \frac{1}{n} \mathbb{E} \left[\sum_{i=0}^{n-1} (Y_i + Z_i) \right] \geq T_{\min}, \end{aligned} \quad (4)$$

where \bar{g}_{opt} is the optimum objective value of Problem (4), the expectation \mathbb{E} is taken over the stochastic service process (Y_0, Y_1, \dots) for given policy π , and T_{\min} is the minimum average update period of the source due to hardware and physical constraints (e.g., limited power resource and cooling capacity). We assume $M > T_{\min}$ such that Problem (4) is feasible and $\bar{g}_{\text{opt}} < \infty$.

Problem (4) belongs to the class of constrained semi-Markov decision processes (SMDP) with a possibly uncountable state space, which is well-known for its difficulty. In this problem, Y_i is the state of the embedded Markov chain, Z_i is the control action taken after observing Y_i , $Y_i + Z_i$ is the update period, and $q(Y_i, Z_i, Y_{i+1})$ is the reward related to both stage i and $i+1$. The class of SMDPs include Markov decision problems (MDPs) [12], [13] and optimization problems of renewal processes [14] as special cases. Most existing studies on SMDPs deal with (i) unconstrained SMDPs, e.g., [13], [15]–[17], or (ii) constrained SMDPs with a countable state space, e.g., [18]–[21]. However, the optimality equations (e.g., Bellman’s equation) for solving unconstrained SMDPs are not applied to constrained SMDPs [22], and the studies for problems with a countable state space cannot be directly applied to Problem (4), which has an uncountable state space.

III. OPTIMAL STATUS UPDATE POLICY

In this section, we develop a chain of novel theoretical results to solve Problem (4): First, we prove that there exists a *stationary randomized* policy that is optimal for Problem (4). Further, we prove that there exists a *stationary deterministic* policy that is optimal for Problem (4). Finally, we develop a low-complexity algorithm to find the *optimal stationary deterministic* policy that solves Problem (4).

A. Optimality of Stationary Randomized Policies

A policy $\pi \in \Pi$ is said to be a *stationary randomized* policy, if it observes Y_i and then chooses a waiting time $Z_i \in [0, M]$ based only on the observed value of Y_i . In this

case, Z_i is determined according to a conditional probability measure $p(y, A) \triangleq \Pr[Z_i \in A | Y_i = y]$ that is invariant for all $i = 0, 1, \dots$. We use Π_{SR} ($\Pi_{\text{SR}} \subseteq \Pi$) to denote the set of stationary randomized policies such that

$\Pi_{\text{SR}} = \{\pi : \text{After observing } Y_i = y_i, Z_i \in [0, M] \text{ is chosen according to probability measure } p(y_i, A), i = 0, 1, \dots\}$.

Note that (Y_i, Z_i, Y_{i+1}) is stationary and ergodic for all stationary randomized policies. In the sequel, when we refer to the stationary distribution of a stationary randomized policy $\pi \in \Pi_{\text{SR}}$, we will remove subscript i . In particular, the random variables (Y_i, Z_i, Y_{i+1}) are replaced by (Y, Z, Y') , where Z is chosen based on the conditional probability measure $\Pr[Z \in A | Y = y] = p(y, A)$ after observing $Y = y$, and (Y, Y') have the same joint distribution as (Y_0, Y_1) . The first key result of this paper is stated as follows:

Theorem 1: (Optimality of Stationary Randomized Policies)

If $g(\cdot)$ is measurable and non-negative, (Y_0, Y_1, \dots) is a stationary ergodic Markov chain with $Y_i \geq 0$ and $0 < \mathbb{E}[Y_i] < \infty$, condition (3) is satisfied, then there exists a stationary randomized policy that is optimal for Problem (4).

Proof sketch of Theorem 1: For any policy $\pi \in \Pi$, define finite horizon average occupations

$$a_{n,\pi} \triangleq \frac{1}{n} \mathbb{E} \left[\sum_{i=0}^{n-1} q(Y_i, Z_i, Y_{i+1}) \right] - \frac{\bar{g}_{\text{opt}}}{n} \mathbb{E} \left[\sum_{i=0}^{n-1} (Y_i + Z_i) \right], \quad (5)$$

$$b_{n,\pi} \triangleq \frac{1}{n} \mathbb{E} \left[\sum_{i=0}^{n-1} (Y_i + Z_i) \right]. \quad (6)$$

Let Γ_{SR} be the set of limit points of sequences $((a_{n,\pi}, b_{n,\pi}), n = 1, 2, \dots)$ associated with stationary randomized policies in Π_{SR} . We first prove that Γ_{SR} is convex and compact. Then, we show that there exists an optimal policy π_{opt} of Problem (4), such that the sequence $((a_{n,\pi_{\text{opt}}}, b_{n,\pi_{\text{opt}}}), n = 1, 2, \dots)$ associated with policy π_{opt} has a limit point (a^*, b^*) satisfying $(a^*, b^*) \in \Gamma_{\text{SR}}$, $a^* \leq 0$, and $b^* \geq T_{\min}$. Since $(a^*, b^*) \in \Gamma_{\text{SR}}$, there exists a stationary randomized policy π^* achieving this limit point (a^*, b^*) . Finally, we show that policy π^* is optimal for Problem (4), which completes the proof. The details are available in Appendix A. ■

The convexity and compactness properties of the set of occupation measures are essential in the study of constrained MDPs [23, Sec. 1.5], which dates back to Derman's monograph in 1970 [24]. Recently, it was used in stochastic optimization for discrete-time queueing systems and renewal processes, e.g., [14], [25]. The techniques in these studies cannot directly handle constrained SMDPs with an uncountable state space, like Problem (4). One crucial novel idea in our proof is to introduce \bar{g}_{opt} in the definition of average occupation in (5), which turns out to be essential in later steps for showing the optimality of the stationary randomized policy π^* .

By Theorem 1, we only need to consider the class of stationary randomized policies Π_{SR} . Therefore, Problem (4) can be simplified to the following functional optimization problem (as shown in Appendix A):

$$\begin{aligned} \min_{p(y,A)} & \frac{\mathbb{E}[q(Y, Z, Y')]}{\mathbb{E}[Y + Z]} \\ \text{s.t.} & \mathbb{E}[Y + Z] \geq T_{\min} \\ & 0 \leq Z \leq M \end{aligned} \quad (7)$$

where $p(y, A) = \Pr[Z \in A | Y = y]$ is the conditional probability measure of some stationary randomized policy, and (Y, Y') have the same distribution as (Y_0, Y_1) .

B. Optimality of Stationary Deterministic Policies

A policy $\pi \in \Pi_{\text{SR}}$ is said to be a *stationary deterministic policy* if $Z_i = z(Y_i)$ for all $i = 0, 1, \dots$, where $z : [0, \infty) \rightarrow [0, M]$ is a deterministic function. We use Π_{SD} ($\Pi_{\text{SD}} \subseteq \Pi_{\text{SR}}$) to denote the set of stationary deterministic policies such that

$$\Pi_{\text{SD}} = \{\pi : Z_i = z(Y_i) \text{ for all } i, 0 \leq z(y) \leq M, \forall y \geq 0\}.$$

Theorem 2: (Optimality of Stationary Deterministic Policies) If $g(\cdot)$ is measurable and non-decreasing, then there exists a stationary deterministic policy that is optimal for Problem (7).

Proof sketch of Theorem 2: Since $g(\cdot)$ is non-decreasing, $q(y, \cdot, y')$ is convex. Using Jensen's inequality, we can show that for any feasible stationary randomized policy $\pi_1 \in \Pi_{\text{SR}}$, there is a feasible stationary deterministic policy that is no worse than policy π_1 . The details are provided in Appendix B of [9]. ■

Let μ_Y be the probability measure of Y_i , then any bounded measurable function $z : [0, \infty) \rightarrow [0, M]$ belongs to the Lebesgue space $L^2(\mu_Y)$ [26, Section 3], because

$$\int_0^\infty |z(y)|^2 d\mu_Y(y) \leq \int_0^\infty M^2 d\mu_Y(y) = M^2 < \infty.$$

By Theorems 1 and 2, we only need to consider the class of stationary deterministic policies Π_{SD} and Problem (4) is simplified to the following functional optimization problem:

$$\min_{z(\cdot) \in L^2(\mu_Y)} \frac{\mathbb{E}[q(Y, z(Y), Y')]}{\mathbb{E}[Y + z(Y)]} \quad (8)$$

$$\begin{aligned} \text{s.t.} & \mathbb{E}[Y + z(Y)] \geq T_{\min} \\ & 0 \leq z(y) \leq M, \forall y \geq 0, \end{aligned} \quad (9)$$

where $z(\cdot)$ is the function associated with a stationary deterministic policy $\pi \in \Pi_{\text{SD}}$. The optimum objective value of Problem (8) is equal to \bar{g}_{opt} .

C. A Low Complexity Solution to Problem (8)

Lemma 1: If $g(\cdot)$ is measurable, non-negative, and non-decreasing, then the functional $h : L^2(\mu_Y) \rightarrow [0, \infty)$ defined by

$$h(z) = \frac{\mathbb{E}[q(Y, z(Y), Y')]}{\mathbb{E}[Y + z(Y)]}$$

is quasi-convex.

Proof: See Appendix C of [9]. ■

Therefore, Problem (8) is a functional quasi-convex optimization problem. In order to solve Problem (8), we consider the following functional convex optimization problem with a parameter c :

$$\begin{aligned} f(c) = \min_{z(\cdot) \in L^2(\mu_Y)} & \mathbb{E}[q(Y, z(Y), Y')] - c\mathbb{E}[Y + z(Y)] \\ \text{s.t.} & \mathbb{E}[Y + z(Y)] \geq T_{\min} \\ & 0 \leq z(y) \leq M, \forall y \geq 0. \end{aligned} \quad (10)$$

It is easy to show that $\bar{g}_{\text{opt}} \leq c$ if and only if $f(c) \leq 0$ [27]. Therefore, we can solve Problem (8) by a two-layer nested algorithm, such as Algorithm 1. In the inner layer, we

Algorithm 1 Two-layer bisection method for Problem (8)

given $l = 0$, sufficiently large $u > \bar{g}_{SD}$, tolerance ϵ_1 .
repeat
 $c := (l + u)/2$.
given $\zeta_l = 0$, sufficiently large $\zeta_u > 0$, tolerance ϵ_2 .
 $\zeta := \zeta_l$, $\nu := \zeta + c$.
 Compute $z_\nu(\cdot)$ in (11).
if $\mathbb{E}[z_\nu(Y)] + \mathbb{E}[Y] < T_{\min}$ **then**
 repeat
 $\zeta := (\zeta_l + \zeta_u)/2$, $\nu := \zeta + c$.
 Compute $z_\nu(\cdot)$ in (11).
 if $\mathbb{E}[z_\nu(Y)] + \mathbb{E}[Y] \geq T_{\min}$, $\zeta_u := \zeta$; **else**, $\zeta_l := \zeta$.
 until $\zeta_u - \zeta_l \leq \epsilon_2$.
end if
if $f(c) \leq 0$, $u := c$; **else**, $l := c$.
until $u - l \leq \epsilon_1$.
return $z(\cdot) := z_\nu(\cdot)$.

use *bisection* to solve Problem (10) for given parameter c ; in the outer layer, we employ *bisection* again to search for a c^* such that $f(c^*) = 0$ and thus $\bar{g}_{\text{opt}} = c^*$. Algorithm 1 has low complexity. It requires at most $\lceil \log_2((u - l)/\epsilon_1) \rceil \times \lceil \log_2((\zeta_u - \zeta_l)/\epsilon_2) \rceil$ iterations to terminate. Each iteration involves computing $\mathbb{E}[z_\nu(Y)]$ based on (11). The optimality of Algorithm 1 is guaranteed by the following theorem:

Theorem 3: If $g(\cdot)$ is measurable, non-negative, and non-decreasing, then an optimal solution π_{opt} to Problem (8) is obtained by Algorithm 1, where the function $z_\nu(\cdot)$ is determined by

$$z_\nu(y) = \sup\{z \in [0, M] : \mathbb{E}[g(y + z + Y')|Y = y] \leq \nu\}. \quad (11)$$

Proof Sketch of Theorem 3: We use Lagrangian duality theory to solve Problem (8). Different from traditional finite dimensional optimization problems [27], Problem (8) is an infinite dimensional functional optimization problem. Therefore, the Karush-Kuhn-Tucker (KKT) theorem for infinite dimensional space [28], [29] and the calculus of variations are required in the analysis. In particular, since the Lagrangian may not be strictly convex for some penalty functions, one-sided Gâteaux derivative (similar to sub-gradient in finite dimensional space) is used to solve the KKT conditions in Lebesgue space $L^2(\mu_Y)$. The proof details are provided in Appendix D of [9]. ■

IV. WHEN IS IT BETTER TO WAIT THAN TO UPDATE?

When $T_{\min} \leq \mathbb{E}[Y]$, a logical policy is the *zero wait* policy: the source submits a fresh update once the prior update completes service, i.e., $\pi_{\text{zero wait}} = (0, 0, \dots)$. As mentioned before, this zero wait policy is not always optimal to keep data fresh. When $T_{\min} > \mathbb{E}[Y]$, due to the constraint (9), the minimum possible average waiting time is $\mathbb{E}[z(Y)] = T_{\min} - \mathbb{E}[Y]$. However, even in this case, the optimal policy may have additional waiting time such that $\mathbb{E}[z(Y)] > T_{\min} - \mathbb{E}[Y]$. In this section, we will study when it is optimal to submit updates with the minimum wait and when it is not.

Algorithm 2 Bisection method for solving Problem (13)

given $l = 0$, sufficiently large u , tolerance ϵ .
repeat
 $\beta := (l + u)/2$.
 $o := \mathbb{E}[(\beta)_Y^{M+Y}] - \max\left(T_{\min}, \frac{\mathbb{E}[(\beta)_Y^{M+Y}]^2}{2\beta}\right)$, where
 $(x)_a^b = \min[\max[x, a], b]$.
if $o \geq 0$, $u := \beta$; **else**, $l := \beta$.
until $u - l \leq \epsilon$.
 Compute $z(\cdot)$ by (15).
return $z(\cdot)$.

A. A Special Case of $g(\Delta) = \Delta$ with i.i.d. Service Times

Consider the case that $g(\Delta) = \Delta$ and the Y_i 's are *i.i.d.*. In this case, Problem (8) has a simpler solution than that provided by Algorithm 1. Interestingly, this solution explicitly characterizes whether the optimal control $z(\cdot)$ can have minimum wait such that $\mathbb{E}[z(Y)] = T_{\min} - \mathbb{E}[Y]$.

As shown in Fig. 2, $q(y, z, y')$ is the area of a trapezoid, computed as

$$q(y, z, y') = \frac{1}{2}(2y + z + y')(z + y').$$

Because the Y_i 's are *i.i.d.*, Y and Y' in Problem (8) are also *i.i.d.* Using this, we can obtain

$$\begin{aligned} & \mathbb{E}[q(Y, z(Y), Y')] \\ &= \mathbb{E}\left[\frac{1}{2}(Y + z(Y))^2 + (Y + z(Y))Y'\right] \\ &= \frac{1}{2}\mathbb{E}[(Y + z(Y))^2] + \mathbb{E}[Y + z(Y)]\mathbb{E}[Y'], \end{aligned} \quad (12)$$

where in (12) we have used that $\mathbb{E}[Y^2] = \mathbb{E}[Y'^2]$. Hence, Problem (8) can be reformulated as

$$\min_{z \in L^2(\mu_Y)} \frac{\mathbb{E}[(Y + z(Y))^2]}{2\mathbb{E}[Y + z(Y)]} + \mathbb{E}[Y] \quad (13)$$

$$\begin{aligned} \text{s.t. } & \mathbb{E}[Y + z(Y)] \geq T_{\min} \\ & 0 \leq z(y) \leq M, \forall y \geq 0. \end{aligned} \quad (14)$$

The following lemma tells us that Problem (13) is a functional convex optimization problem.

Lemma 2: The functional $h_1 : L^2(\mu_Y) \rightarrow \mathbb{R}$ defined by

$$h_1(z) = \frac{\mathbb{E}[(Y + z(Y))^2]}{\mathbb{E}[Y + z(Y)]}$$

is convex on the domain

$$\text{dom } h_1 = \{z : z(y) \in [0, M], \forall y \geq 0, z \in L^2(\mu_Y)\}.$$

Proof: See Appendix E of [9]. ■

Using the KKT theorem for infinite dimensional space and the calculus of variations, we can obtain

Theorem 4: The optimal solution to Problem (13) is

$$z(y) = (\beta - y)_0^M, \quad (15)$$

where $(x)_a^b \triangleq \min[\max[x, a], b]$ and $\beta > 0$ is the root of the

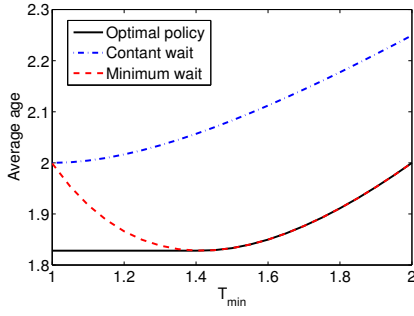


Fig. 4. Average age vs. T_{\min} with *i.i.d.* discrete service times.

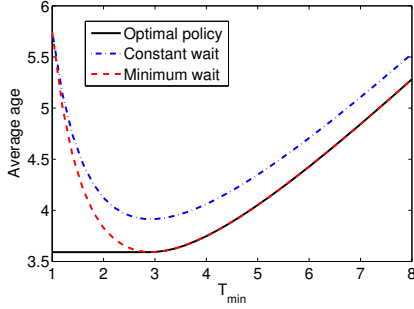


Fig. 5. Average age vs. T_{\min} with *i.i.d.* log-normal distributed service times.

following equation:

$$\mathbb{E}[(\beta)_Y^{M+Y}] = \max\left(T_{\min}, \frac{\mathbb{E}[(\beta)_Y^{M+Y}]^2}{2\beta}\right). \quad (16)$$

Proof: See Appendix F of [9]. ■

Equation (15) has the form of a water-filling solution, where the water-level β is given by the root of equation (16). One can observe that (11) reduces to (15) if $g(\Delta) = \Delta$, the Y_i 's are *i.i.d.*, and τ is replaced by $\beta + \mathbb{E}[Y]$. The root β of equation (16) can be simply solved by the bisection search method in Algorithm 2. We note that Algorithm 2 has a lower complexity than Algorithm 1 in the special case of $g(\Delta) = \Delta$ and *i.i.d.* service process, while Algorithm 1 can obtain the optimal policy in more general scenarios.

Interestingly, Theorem 4 provides a closed-form criterion on whether the optimal $z(\cdot)$ satisfies $\mathbb{E}[z(Y)] = T_{\min} - \mathbb{E}[Y]$. Specifically, (15) implies $(\beta)_Y^{M+Y} = Y + z(Y)$. This and (16) tell us that if $T_{\min} \geq \frac{\mathbb{E}[(Y+z(Y))^2]}{2\beta}$, then the optimal control $z(\cdot)$ satisfies

$$\mathbb{E}[Y + z(Y)] = T_{\min} \geq \frac{\mathbb{E}[(Y + z(Y))^2]}{2\beta}, \quad (17)$$

such that the optimal policy has the minimum average waiting time $\mathbb{E}[z(Y)] = T_{\min} - \mathbb{E}[Y]$; otherwise, if $T_{\min} < \frac{\mathbb{E}[(Y+z(Y))^2]}{2\beta}$, the optimal control $z(\cdot)$ satisfies

$$\mathbb{E}[Y + z(Y)] = \frac{\mathbb{E}[(Y + z(Y))^2]}{2\beta} > T_{\min}, \quad (18)$$

such that the average waiting time $\mathbb{E}[z(Y)]$ of the optimal policy is larger than $T_{\min} - \mathbb{E}[Y]$.

Furthermore, we consider the case $T_{\min} = 0$, where the constraint (14) is always satisfied and can be removed. Note




Fig. 6. Average age vs. the correlation coefficient ρ between Y_i and Y_{i+1} for discrete service times, where $T_{\min} \leq \mathbb{E}[Y]$.




Fig. 7. Average age vs. the correlation coefficient ρ between Y_i and Y_{i+1} for log-normal distributed service times, where $T_{\min} \leq \mathbb{E}[Y]$.

that we have the same problem for all $T_{\min} \leq \mathbb{E}[Y]$ and hence we can pick $T_{\min} = \mathbb{E}[Y]$. By substituting $T_{\min} = \mathbb{E}[Y]$ into (17) and (18), we can obtain the criterion on whether the zero waiting policy is optimal.

1) *Simulation Results:* We use “Optimal policy” to refer to the policy provided in Theorem 3 (or its special case in Theorem 4), and compare it with two reference policies:

- “Constant wait”: Each update is followed by a constant wait $Z_i = \text{const}$ before submitting the next update with $\text{const} = T_{\min} - \mathbb{E}[Y]$.
- “Minimum wait”: The update waiting time is given by a deterministic function $Z_i = z(Y_i)$, where $z(\cdot)$ is given by (15) and β is chosen to satisfy $\mathbb{E}[z(Y)] = T_{\min} - \mathbb{E}[Y]$.³

When $\mathbb{E}[Y] = T_{\min}$, both the constant wait and minimum wait policies reduce to the zero wait policy. Two models of the service processes are considered: The first one is a discrete Markov chain with a probability mass function $\Pr[Y_i = 0] = \Pr[Y_i = 2] = 0.5$ and a transition matrix

$$P = \begin{bmatrix} p & 1-p \\ 1-p & p \end{bmatrix}.$$

Hence, the Y_i 's are *i.i.d.* when $p = 0.5$, and the correlation coefficient between Y_i and Y_{i+1} is $\rho = 2p - 1$. The second one is a log-normal distributed Markov chain, where $Y_i = e^{\sigma X_i} / \mathbb{E}[e^{\sigma X_i}]$ and (X_0, X_1, \dots) is a Gaussian Markov process satisfying the first-order AR equation

$$X_{i+1} = \eta X_i + \sqrt{1 - \eta^2} W_i,$$

where $\sigma > 0$ is the scale parameter of log-normal distribution, $\eta \in [-1, 1]$ is the parameter of the AR model, and the W_i 's are *i.i.d.* Gaussian random variables with zero mean

³This policy was called “ β -minimum” in [8].




Fig. 8. Average age vs. the distribution parameter σ of log-normal distributed service times, where $T_{\min} \leq \mathbb{E}[Y]$ and $\rho = (e^{0.5} - 1)/(e - 1)$.




Fig. 9. Average age penalty vs. the parameter α of stair-step penalty functions with discrete service times, where $T_{\min} \leq \mathbb{E}[Y]$, $g(\Delta) = \lfloor \alpha \Delta \rfloor$, and $\rho = 0.4$.

and unit variance. The log-normal distributed Markov chain is normalized such that $\mathbb{E}[Y_i] = 1$. According to the properties of log-normal distribution, the correlation coefficient between Y_i and Y_{i+1} is $\rho = (e^\eta - 1)/(e - 1)$. Then, the Y_i 's are *i.i.d.* when $\eta = 0$. The value of M is set to be 10.

Figures 4 and 5 illustrate the average age vs. T_{\min} for *i.i.d.* discrete and log-normal distributed service times, respectively. In both figures, one can observe that the constant wait policy always incurs a larger average age than the optimal policy. In addition, as expected from (17) and (18), as T_{\min} exceeds a certain threshold, the optimal policy meets the constraint (14) with equality. For smaller values of T_{\min} , the constraint is not active for the optimal solution. Consequently, the average age achieved by the minimum wait policy deviates from the optimum for small values of T_{\min} .

B. General Age Penalties and Correlated Service Processes

For general age penalties and correlated service processes, it is essentially difficult to find closed-form characterization on whether the optimal control $z(\cdot)$ can have minimum wait such that $\mathbb{E}[z(Y)] = T_{\min} - \mathbb{E}[Y]$. Therefore, we focus on the case of $T_{\min} \leq \mathbb{E}[Y]$ and study when the zero wait policy minimizes the average age penalty. Sufficient conditions for the optimality of the zero wait policy are provided as follows:

Lemma 3: Suppose that $T_{\min} \leq \mathbb{E}[Y]$, $g(\cdot)$ is measurable, non-negative, and non-decreasing. The zero wait policy is optimal for Problem (8) if one of the following is satisfied:

- 1). The correlation coefficient between Y_i and Y_{i+1} is -1 ;
- 2). The Y_i 's are equal to a constant value;
- 3). $g(\cdot)$ is a constant function.

Proof: See Appendix G of [9]. ■

1) *Simulation Results:* We now provide some simulation results for general age penalties and/or correlated service processes. Figures 6 and 7 depict the average age vs. the correlation coefficient ρ between Y_i and Y_{i+1} for discrete and log-normal distributed service times, respectively. In Fig. 6, the regime of ρ is $[-1, 1]$. We observe that the zero wait policy is optimal when $\rho \in [-1, -0.5]$, and the performance gap between the optimal policy and the zero wait policy grows with ρ when $\rho \geq -0.5$. This is in accordance with the example in the introduction: As ρ grows, the occurrence of two consecutive zero service times (i.e., $(Y_i, Y_{i+1}) = (0, 0)$) increases. Therefore, more and more updates are wasted in the zero wait policy, leading to a larger gap from the optimum. In Fig. 7, the regime of ρ is $[(e^{-1} - 1)/(e - 1), 1]$. In this case, the sub-optimality gap of the zero wait policy also increases with ρ . The point $\rho = 1$ is not plotted in these figures because the corresponding Markov chains are not ergodic.

Figure 8 considers the average age vs. the parameter σ of log-normal distributed service times, where $\rho = (e^{0.5} - 1)/(e - 1)$. We observe that the zero wait policy is optimal for small σ and is not optimal for large σ . Note that when $\sigma = 0$, the service times are constant, i.e., $Y_i = 1$ for all i .

Figures 9 shows the average age penalty vs. the parameter α of stair-step penalty functions, where $g(\Delta) = \lfloor \alpha \Delta \rfloor$. The correlation coefficient is $\rho = 0.4$ for discrete service times. We find that the zero wait policy is optimal if $\alpha = 0$, in which case $g(\Delta)$ is a constant function. When $\alpha > 0$, the zero wait policy may not be optimal. Due to space limitations, additional simulation results of power and exponential penalty functions are provided in [9].

These simulation results suggest that the conditions in Lemma 3 are sufficient but not necessary.

V. CONCLUSION

We studied the optimal control of status updates from a source to a remote monitor via a FCFS network server. We generalized the existing studies on status update age to cover general age penalty functions and non-*i.i.d.* service processes. We developed efficient algorithms to find the optimal status update policy for minimizing the average age penalty. We showed that, surprisingly, in many scenarios, the optimal policy is to wait for a certain time before submitting a new update. In particular, the widely-adopted zero wait policy can be far from the optimum if (i) the penalty function grows quickly with respect to the age, and (ii) the update service times are highly random and positive correlated.

APPENDIX A PROOF OF THEOREM 1

A. An Upper Bound of \bar{g}_{opt}

By restricting Π in Problem (4) to Π_{SR} , we obtain the following problem:

$$\begin{aligned} \bar{g}_{SR} = \min_{\pi \in \Pi_{SR}} \limsup_{n \rightarrow \infty} & \frac{\mathbb{E} \left[\sum_{i=0}^{n-1} q(Y_i, Z_i, Y_{i+1}) \right]}{\mathbb{E} \left[\sum_{i=0}^{n-1} (Y_i + Z_i) \right]} \\ \text{s.t. } \liminf_{n \rightarrow \infty} & \frac{1}{n} \mathbb{E} \left[\sum_{i=0}^{n-1} (Y_i + Z_i) \right] \geq T_{\min}, \end{aligned} \quad (19)$$

where \bar{g}_{SR} is the optimum objective value of Problem (19). Since $\Pi_{\text{SR}} \subseteq \Pi$, we can obtain

$$\bar{g}_{\text{SR}} \geq \bar{g}_{\text{opt}}. \quad (20)$$

It is easy to show that the (Y_i, Z_i, Y_{i+1}) 's are stationary and ergodic for all stationary randomized policies. This, together with the condition that $g(\cdot)$ is measurable, tells us that $q(Y_i, Z_i, Y_{i+1})$ is stationary and ergodic [11, Theorems 7.1.1 and 7.1.3]. For any stationary randomized policy $\pi = (Z_0, Z_1, \dots) \in \Pi_{\text{SR}}$, we obtain

$$\frac{1}{n} \mathbb{E} \left[\sum_{i=0}^{n-1} q(Y_i, Z_i, Y_{i+1}) \right] = \mathbb{E}[q(Y_0, Z_0, Y_1)], \quad (21)$$

$$\frac{1}{n} \mathbb{E} \left[\sum_{i=0}^{n-1} (Y_i + Z_i) \right] = \mathbb{E}[Y_0 + Z_0]. \quad (22)$$

Hence, Problem (19) can be reformulated as Problem (7).

B. The Upper Bound of \bar{g}_{opt} is Tight, i.e., $\bar{g}_{\text{SR}} = \bar{g}_{\text{opt}}$

We will show $\bar{g}_{\text{SR}} = \bar{g}_{\text{opt}}$ in 4 steps. The following definitions are needed: Since \bar{g}_{opt} is finite, for each causally feasible policy $\pi = (Z_0, Z_1, \dots) \in \Pi$ we can define $a_{n,\pi}$ and $b_{n,\pi}$ as in (5) and (6), respectively.

Further, define Γ_{SR} as the set of limit points of sequences $((a_{n,\pi}, b_{n,\pi}), n = 1, 2, \dots)$ associated with all stationary randomized policies $\pi \in \Pi_{\text{SR}}$. Because the renewal reward $q(Y_i, Z_i, Y_{i+1})$ and renewal interval $Y_i + Z_i$ are stationary and ergodic for all stationary randomized policies $\pi \in \Pi_{\text{SR}}$, the sequence $(a_{n,\pi}, b_{n,\pi})$ has a unique limit point in the form of

$$(\mathbb{E}[q(Y, Z, Y')] - \bar{g}_{\text{opt}} \mathbb{E}[Y + Z], \mathbb{E}[Y + Z]). \quad (23)$$

Hence, Γ_{SR} is the set of all points $(\mathbb{E}[q(Y, Z, Y')] - \bar{g}_{\text{opt}} \mathbb{E}[Y + Z], \mathbb{E}[Y + Z])$, where each point is associated with a conditional probability measure $p(y, A) = \Pr[Z \in A | Y = y]$, and the measure of (Y, Y') is the same as that of (Y_0, Y_1) . Note that $\mathbb{E}[Y] = \mathbb{E}[Y']$.

Step 1: We will show that Γ_{SR} is a convex and compact set.

It is easy to show that Γ_{SR} is convex by considering a stationary randomized policy that is a mixture of two stationary randomized policies.

For compactness, let $((d_j, e_j), j = 1, 2, \dots)$ be any sequence of points in Γ_{SR} , we need to show that there is a convergent subsequence (d_{j_k}, e_{j_k}) whose limit is also in Γ_{SR} . Since $(d_j, e_j) \in \Gamma_{\text{SR}}$, there must exist $(Y, Z_{(j)}, Y')$ with conditional probability $p_j(y, A) = \Pr[Z_{(j)} \in A | Y = y]$, such that $d_j = \mathbb{E}[q(Y, Z_{(j)}, Y')] - \bar{g}_{\text{opt}} \mathbb{E}[Y + Z_{(j)}]$, $e_j = \mathbb{E}[Y + Z_{(j)}]$. Let μ_j be the joint probability measure of $(Y, Z_{(j)}, Y')$, then (d_j, e_j) is uniquely determined by μ_j . For any L satisfying $L \geq M$, we can obtain

$$\begin{aligned} & \mu_j(Y \leq L, Z_{(j)} \leq L, Y' \leq L) \\ &= \Pr(Y \leq L, Y' \leq L) \\ &\geq \Pr(Y + Y' \leq L) \\ &\geq 1 - \frac{\mathbb{E}[Y + Y']}{L}, \quad \forall j, \end{aligned}$$

where the equality is due to the fact that $Z_{(j)} \leq M \leq L$ and the last inequality is due to Markov's inequality. Therefore,

for any ϵ , there is an L such that

$$\liminf_{j \rightarrow \infty} \mu_j(|Y| \leq L, |Z_{(j)}| \leq L, |Y'| \leq L) \geq 1 - \epsilon.$$

Hence, the sequence of measures μ_j is tight. By Helly's selection theorem [11, Theorem 3.9.2], there is a subsequence of measures μ_{j_k} that converges weakly to a limit measure μ_∞ .

Let $(Y, Z_{(\infty)}, Y')$ and $p_\infty(y, A) = \Pr[Z_\infty \in A | Y = y]$ denote the random vector and conditional probability corresponding to the limit measure μ_∞ , respectively. We can define $d_\infty = \mathbb{E}[q(Y, Z_{(\infty)}, Y')] - \bar{g}_{\text{opt}} \mathbb{E}[Y + Z_{(\infty)}]$, $e_\infty = \mathbb{E}[Y + Z_{(\infty)}]$. Since the function $q(y, z, y')$ is in the form of an integral, it is continuous and thus measurable. Using the continuous mapping theorem [11, Theorem 3.2.4], we can obtain that $q(Y, Z_{(j_k)}, Y')$ converges weakly to $q(Y, Z_{(\infty)}, Y')$. Then, using the condition (3), together with the dominated convergence theorem (Theorem 1.6.7 of [11]) and Theorem 3.2.2 of [11], we can obtain $\lim_{k \rightarrow \infty} (d_{j_k}, e_{j_k}) = (d_\infty, e_\infty)$. Hence, $((d_j, e_j), j = 1, 2, \dots)$ has a convergent subsequence. Further, we can generate a stationary randomized policy $\pi_{\infty, \text{SR}}$ by using the conditional probability $p_\infty(y, A)$ corresponding to μ_∞ . Then, (d_∞, e_∞) is the limit point generated by the stationary randomized policy $\pi_{\infty, \text{SR}}$, which implies $(d_\infty, e_\infty) \in \Gamma_{\text{SR}}$. In summary, any sequence (d_j, e_j) in Γ_{SR} has a convergent subsequence (d_{j_k}, e_{j_k}) whose limit (d_∞, e_∞) is also in Γ_{SR} . Therefore, Γ_{SR} is a compact set.

Step 2: We will show that there exists an optimal policy $\pi_{\text{opt}} \in \Pi$ of Problem (4) such that the sequence $(a_{n,\pi_{\text{opt}}}, b_{n,\pi_{\text{opt}}})$ associated with policy π_{opt} has at least one limit point in Γ_{SR} .

Since the sequence (Y_0, Y_1, \dots) is a Markov chain, the observation Y_{i+1} depends only on the immediately preceding state Y_i and not on the history state and control $Y_0, \dots, Y_{i-1}, Z_0, \dots, Z_{i-1}$. Therefore, Y_i is the *sufficient statistic* [13, p. 252] for solving Problem (4). This tells us that there exists an optimal policy $\pi_{\text{opt}} = (Z_0, Z_1, \dots) \in \Pi$ of Problem (4) in which the control action Z_i is determined based on only Y_i , but not the history state and control $Y_0, \dots, Y_{i-1}, Z_0, \dots, Z_{i-1}$ [13]. We will show that the sequence $(a_{n,\pi_{\text{opt}}}, b_{n,\pi_{\text{opt}}})$ associated with this policy π_{opt} has at least one limit point in Γ_{SR} .

It is known that Z_i takes values in the standard Borel space $(\mathbb{R}, \mathcal{R})$, where \mathcal{R} is the Borel σ -field. According to [11, Theorem 5.1.9], for each i there exists a conditional probability measure $p'_i(y, A)$ such that $p'_i(y, A) = \Pr(Z_i \in A | Y_i = y)$ for almost all y . One can use this conditional probability $p'_i(y, A)$ to generate a stationary randomized policy $\pi'_{i, \text{SR}} \in \Pi_{\text{SR}}$. Then, the one-stage expectation $(\mathbb{E}[q(Y_i, Z_i, Y_{i+1})] - \bar{g}_{\text{opt}} \mathbb{E}[Y_i + Z_i], \mathbb{E}[Y_i + Z_i])$ is exactly the limit point generated by the stationary randomized policy $\pi'_{i, \text{SR}}$. Thus, $(\mathbb{E}[q(Y_i, Z_i, Y_{i+1})] - \bar{g}_{\text{opt}} \mathbb{E}[Y_i + Z_i], \mathbb{E}[Y_i + Z_i]) \in \Gamma_{\text{SR}}$ for all $i = 0, 1, 2, \dots$. Using (5), (6), and the fact that Γ_{SR} is convex, we can obtain $(a_{n,\pi_{\text{opt}}}, b_{n,\pi_{\text{opt}}}) \in \Gamma_{\text{SR}}$ for all $n = 1, 2, 3, \dots$. In other words, the sequence $(a_{n,\pi_{\text{opt}}}, b_{n,\pi_{\text{opt}}})$ is within Γ_{SR} . Since Γ_{SR} is a compact set, the sequence $(a_{n,\pi_{\text{opt}}}, b_{n,\pi_{\text{opt}}})$ must have a convergent subsequence, whose limit is in Γ_{SR} .

Step 3: Let $(a^, b^*) \in \Gamma_{\text{SR}}$ be one limit point of the sequence $(a_{n,\pi_{\text{opt}}}, b_{n,\pi_{\text{opt}}})$ associated with policy π_{opt} . We will show that $a^* \leq 0$ and $b^* \geq T_{\min}$.*

Policy π_{opt} is feasible for Problem (4) and meanwhile

achieves the optimum objective value \bar{g}_{opt} . Hence,

$$\limsup_{n \rightarrow \infty} \frac{c_{n, \pi_{\text{opt}}}}{b_{n, \pi_{\text{opt}}}} = \bar{g}_{\text{opt}}, \quad (24)$$

$$\liminf_{n \rightarrow \infty} b_{n, \pi_{\text{opt}}} \geq T_{\min}, \quad (25)$$

where

$$c_{n, \pi_{\text{opt}}} \triangleq \frac{1}{n} \mathbb{E} \left[\sum_{i=0}^{n-1} q(Y_i, Z_i, Y_{i+1}) \right].$$

By (6), $b_{n, \pi_{\text{opt}}}$ is upper bounded by

$$b_{n, \pi_{\text{opt}}} \leq \frac{1}{n} \mathbb{E} \left[\sum_{i=0}^{n-1} (M + Y_{i+1}) \right] = M + \mathbb{E}[Y] < \infty.$$

Hence, by (5), we have

$$\begin{aligned} a_{n, \pi_{\text{opt}}} &= c_{n, \pi_{\text{opt}}} - \bar{g}_{\text{opt}} b_{n, \pi_{\text{opt}}} \\ &\leq \max\{c_{n, \pi_{\text{opt}}} - \bar{g}_{\text{opt}} b_{n, \pi_{\text{opt}}}, 0\} \\ &= \max\left\{\frac{c_{n, \pi_{\text{opt}}}}{b_{n, \pi_{\text{opt}}}} - \bar{g}_{\text{opt}}, 0\right\} b_{n, \pi_{\text{opt}}} \\ &\leq \max\left\{\frac{c_{n, \pi_{\text{opt}}}}{b_{n, \pi_{\text{opt}}}} - \bar{g}_{\text{opt}}, 0\right\} (M + \mathbb{E}[Y]). \end{aligned}$$

Taking the lim sup in this inequality and using (24), yields

$$\limsup_{n \rightarrow \infty} a_{n, \pi_{\text{opt}}} \leq 0. \quad (26)$$

Because (a^*, b^*) is one limit point of $(a_{n, \pi_{\text{opt}}}, b_{n, \pi_{\text{opt}}})$, we have

$$a^* \leq \limsup_{n \rightarrow \infty} a_{n, \pi_{\text{opt}}}, \quad b^* \geq \liminf_{n \rightarrow \infty} b_{n, \pi_{\text{opt}}}. \quad (27)$$

By (25)-(27), we have $a^* \leq 0$ and $b^* \geq T_{\min}$.

Step 4: We will show that there exists a stationary randomized policy that is optimal for Problems (4) and (7), and thus $\bar{g}_{\text{SR}} = \bar{g}_{\text{opt}}$. By the definition of Γ_{SR} , $(a^, b^*) \in \Gamma_{\text{SR}}$ must be the limit point generated by a stationary randomized policy $\pi^* \in \Pi_{\text{SR}}$. Let (Y, Z^*, Y') be a random vector with the stationary distribution of policy π^* . Then, (23) implies*

$$(a^*, b^*) = (\mathbb{E}[q(Y, Z^*, Y')] - \bar{g}_{\text{opt}} \mathbb{E}[Y + Z^*], \mathbb{E}[Y + Z^*]).$$

Using $a^* \leq 0$ and $b^* \geq T_{\min}$, we can obtain

$$\mathbb{E}[q(Y, Z^*, Y')] - \mathbb{E}[Y + Z^*] \bar{g}_{\text{opt}} \leq 0, \quad (28)$$

$$\mathbb{E}[Y + Z^*] \geq T_{\min}. \quad (29)$$

By (28) and $\mathbb{E}[Y + Z^*] > 0$, we have

$$\frac{\mathbb{E}[q(Y, Z^*, Y')]}{\mathbb{E}[Y + Z^*]} \leq \bar{g}_{\text{opt}}.$$

Further, the inequality (29) suggests that the stationary randomized policy π^* is feasible for Problem (7). Hence,

$$\frac{\mathbb{E}[q(Y, Z^*, Y')]}{\mathbb{E}[Y + Z^*]} \geq \bar{g}_{\text{SR}}.$$

Therefore, $\bar{g}_{\text{SR}} \leq \bar{g}_{\text{opt}}$. This and (20) suggest that

$$\frac{\mathbb{E}[q(Y, Z^*, Y')]}{\mathbb{E}[Y + Z^*]} = \bar{g}_{\text{SR}} = \bar{g}_{\text{opt}}.$$

This completes the proof.

REFERENCES

- [1] S. Kaul, R. Yates, and M. Gruteser, "Real-time status: How often should one update?" in *Proc. INFOCOM*, Apr. 2012, pp. 2731–2735.
- [2] —, "Status updates through queues," in *Conf. on Info. Sciences and Systems*, Mar. 2012.
- [3] R. Yates and S. Kaul, "Real-time status updating: Multiple sources," in *Proc. IEEE Int. Symp. Inform. Theory*, Jul. 2012.
- [4] C. Kam, S. Kompella, and A. Ephremides, "Age of information under random updates," in *Proc. IEEE Int. Symp. Inform. Theory*, 2013.
- [5] —, "Effect of message transmission diversity on status age," in *Proc. IEEE Int. Symp. Inform. Theory*, June 2014, pp. 2411–2415.
- [6] M. Costa, M. Codreanu, and A. Ephremides, "Age of information with packet management," in *Proc. IEEE Int. Symp. Inform. Theory*, June 2014, pp. 1583–1587.
- [7] L. Huang and E. Modiano, "Optimizing age-of-information in a multi-class queueing system," in *Proc. IEEE Int. Symp. Inform. Theory*, 2015.
- [8] R. Yates, "Lazy is timely: Status updates by an energy harvesting source," in *Proc. IEEE Int. Symp. Inform. Theory*, 2015.
- [9] Y. Sun, E. Uysal-Biyikoglu, R. Yates, C. E. Koksall, and N. B. Shroff, "Update or wait: How to keep your data fresh," technical report. [Online]. Available: <http://arxiv.org/abs/1601.02284>
- [10] T. Bacinoglu, E. T. Ceran, and E. Uysal-Biyikoglu, "Age of information under energy replenishment constraints," in *Proc. Info. Theory and Appl. Workshop*, Feb. 2015.
- [11] R. Durrett, *Probability: Theory and Examples*, 4th ed. Cambridge University Press, 2010.
- [12] R. Bellman, *Dynamic Programming*. Princeton University Press, 1957.
- [13] D. P. Bertsekas, *Dynamic Programming and Optimal Control*, 3rd ed. Belmont, MA: Athena Scientific, 2005, vol. 1.
- [14] M. J. Neely, "Dynamic optimization and learning for renewal systems," *IEEE Trans. Auto. Control*, vol. 58, no. 1, pp. 32–46, Jan 2013.
- [15] S. M. Ross, *Applies Probability Models with Optimization Applications*. San Francisco, CA: Holden-Day, 1970.
- [16] A. Federgruen, P. J. Schweitzer, and H. C. Tijms, "Denumerable undiscounted semi-Markov decision processes with unbounded rewards," *Mathematics of Operations Research*, vol. 8, no. 2, pp. 298–313, 1983.
- [17] D. Klabjan and D. Adelman, "Existence of optimal policies for semi-Markov decision processes using duality for infinite linear programming," *SIAM J. Control Optim.*, vol. 44, no. 6, pp. 2104–2122, Jan. 2006.
- [18] F. J. Beutler and K. W. Ross, "Optimal policies for controlled Markov chains with a constraint," *Journal of Mathematical Analysis and Applications*, vol. 112, no. 1, pp. 236 – 252, 1985.
- [19] —, "Time-average optimal constrained semi-Markov decision processes," *Adv. in Appl. Probab.*, vol. 18, no. 2, pp. 341–359, 1986.
- [20] E. Feinberg, "Constrained semi-Markov decision processes with average rewards," *Zeitschrift für Operations Research*, vol. 39, no. 3, pp. 257–288, 1994.
- [21] M. Baykal-Gürsoy and K. Gürsoy, "Semi-Markov decision processes: nonstandard criteria," *Probability in the Engineering and Informational Sciences*, vol. 21, pp. 635–657, Oct. 2007.
- [22] E. A. Feinberg, "Constrained discounted semi-Markov decision processes," in *Markov Processes and Controlled Markov Chains*, Z. Hou, J. A. Filar, and A. Chen, Eds. Kluwer Academic Publishers, 2002.
- [23] E. Altman, *Constrained Markov decision processes*. Boca Raton, FL, USA: Chapman & Hall/CRC, 1999.
- [24] C. Derman, *Finite State Markovian Decision Processes*. Orlando, FL, USA: Academic Press, Inc., 1970.
- [25] M. J. Neely, *Stochastic Network Optimization with Application to Communication and Queueing Systems*. Morgan & Claypool, 2010.
- [26] W. Rudin, *Real and Complex Analysis*, 3rd ed. New York: McGraw-Hill Book Co., 1987.
- [27] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, U.K.: Cambridge University Press, 2004.
- [28] D. Butnariu and A. N. Iusem, *Totally Convex Functions for Fixed Points Computation and Infinite Dimensional Optimization*. Norwell, MA: Kluwer Academic Publisher, 2000.
- [29] G. G. Magaril-Il'yayev and V. M. Tikhomirov, *Convex Analysis: Theory and Applications*. American Mathematical Society, 2003.