

Online Timely Status Updates with Erasures for Energy Harvesting Sensors

Ahmed Arafa¹, Jing Yang², Sennur Ulukus³, and H. Vincent Poor¹

¹Electrical Engineering Department, Princeton University

²School of Electrical Engineering and Computer Science, The Pennsylvania State University

³Department of Electrical and Computer Engineering, University of Maryland

Abstract—An energy harvesting sensor that is sending status updates to a destination through an *erasure* channel is considered, in which transmissions are prone to being erased with some probability q , independently from other transmissions. The sensor, however, is *unaware* of erasure events due to lack of feedback from the destination. Energy expenditure is normalized in the sense that one transmission consumes one unit of energy. The sensor is equipped with a *unit-sized* battery to save its incoming energy, which arrives according to a Poisson process of unit rate. The setting is *online*, in which energy arrival times are only revealed causally after being harvested, and the goal is to design transmission times such that the long term average age of information (AoI), defined as the time elapsed since the latest update has reached the destination *successfully*, is minimized. The optimal status update policy is first shown to have a *renewal* structure, in which the time instants at which the destination receives an update successfully constitute a renewal process. Then, for $q \leq \frac{1}{2}$, the optimal renewal policy is shown to have a *threshold* structure, in which a new status update is transmitted only if the AoI grows above a certain threshold, that is shown to be a decreasing function of q . While for $q > \frac{1}{2}$, the optimal renewal policy is shown to be *greedy*, in which a new status update is transmitted whenever energy is available.

I. INTRODUCTION

We consider an energy harvesting sensor monitoring some physical phenomenon and sending measurement status updates to a destination through an erasure channel, see Fig. 1. We design optimal policies that minimize the long term average AoI, with only causal knowledge of energy arrival times and with no erasure error feedback. The AoI is the time elapsed since the latest update has reached the destination successfully.

The AoI metric has been studied in the literature under various settings; mainly through modeling the update system as a queuing system and analyzing the long term average AoI, and through using optimization tools to characterize optimal status update policies, see, e.g., [1]–[11], and also the recent survey in [12]. AoI minimization in energy harvesting communications has recently gained attention in [13]–[27] under various service time (time for the update to take effect), battery capacity, and channel assumptions. With the exception of [16], an underlying assumption in these works is that energy expenditure is normalized, i.e., sending one status

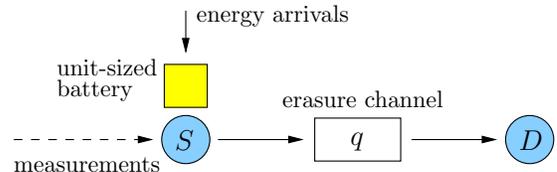


Fig. 1. Energy harvesting sensor with a unit-sized battery sending measurement status updates through an erasure channel.

update consumes one energy unit. Under a perfect channel model, optimality of online threshold policies, in which a new status update is transmitted only if the AoI grows above a certain threshold, is first shown in [18] for unit-sized batteries, and is later extended in [24]–[26], and independently and concurrently in [27], for general finite-sized batteries. Under an erasure channel model, references [19], [20] show the optimality of best-effort uniform updating policies, in which an update is sent every one time unit only if energy is available, for infinite batteries, with and without erasure error feedback. The use of coding to combat erasures along with best-effort and save-and-transmit online policies is studied in [21], extending the non-energy-harvesting work in [9].

In this paper, we extend the results of [19] to the case of *finite* batteries of unit size. Under Poisson energy arrivals with unit rate, with only causal knowledge of their arrival times, we design age-minimal status update policies for an *erasure* communication channel. With the sensor only knowing the erasure probability, and not the actual erasure state per transmission due to *lack of feedback* from the destination (receiver), we show that the optimal policy has a *renewal* structure. Specifically, the sensor should design its transmission times such that the times at which the destination receives an update *successfully* form a renewal process with i.i.d. inter-durations. Then, we show that the optimal renewal policy’s structure depends on the value of the erasure probability, q . Specifically, if $q \leq \frac{1}{2}$, the optimal policy is a *threshold* policy, in which a new update is transmitted only if the AoI grows above a certain threshold that is a decreasing function of q . On the other hand, if $q > \frac{1}{2}$, the optimal policy turns *greedy*, in which a new update is transmitted whenever energy is available. Intuitively, the higher the erasure probability, the more eager the sensor becomes to transmit new updates, so that when they are eventually received successfully the AoI is not too large.

This research was supported in part by the National Science Foundation under Grants CNS-1702808, ECCS-1650299, CNS-1526608 and ECCS-1807348.

II. SYSTEM MODEL AND PROBLEM FORMULATION

Consider an energy harvesting sensor that is sending status updates regarding measurements of some physical phenomenon to a destination. Updates are prone to erasures with some probability $q \in (0, 1)$, representing the case in which the destination is unable to decode the status update message, due to fading or noise contamination in the channel. Such erasure events occur independently of the transmitted status updates and are also mutually independent over time. However, when a status update does go through with no erasures, it arrives at the destination instantly as in [18], [26], [27]. Each status update carries a time stamp denoting when it was acquired at the sensor. From the destination's perspective, the AoI at time t , $a(t)$, is defined as the time elapsed since the latest update has been *successfully* received, i.e., with no erasures. This is mathematically given by

$$a(t) = t - u(t), \quad (1)$$

where $u(t)$ is the time stamp of the last update that has been successfully received prior to time t .

Energy expenditure is normalized in the sense that each status update transmission (and measurement/processing) consumes 1 energy unit. In this work, we investigate the case in which the sensor is equipped with a *unit-sized battery* to save its incoming energy. This special case models situations in which a single transmission (along with its processing requirements) from the sensor completely depletes its battery. Energy arrives (is harvested) according to a Poisson process with (normalized) rate 1. The setting is *online*, in which arrival times are revealed causally after the energy is harvested.

Let l_i denote the time at which the sensor measures and transmits its i th update, and denote by $\mathcal{E}(t)$ the energy available in the battery at time t . Then, we must have the following energy causality constraint:

$$\mathcal{E}(l_i^-) \geq 1, \quad \forall i. \quad (2)$$

Under any policy satisfying the above, the battery state evolves as follows:

$$\mathcal{E}(l_i^-) = \min\{\mathcal{E}(l_{i-1}^-) - 1 + A_i, 1\}, \quad \forall i, \quad (3)$$

where A_i denotes the amount of energy harvested in the time interval $[l_{i-1}, l_i)$, which has a Poisson distribution with parameter $l_i - l_{i-1}$. Without loss of generality, we assume $l_0 = 0$, $\mathcal{E}(0) = 1$, and that $a(0) = 0$. That is, the system starts at time 0 with *fresh* information.

In this work, we focus on the case in which there is *no erasure feedback* sent to the sensor from the destination. In other words, the sensor is oblivious to what occurs to its transmitted updates in the channel. However, the sensor is aware of the fact that updates are subject to erasures with probability q , and takes that probability knowledge into consideration while deciding on when to transmit. Let s_i denote the i th *actual* update time at the destination, i.e., the time at which the i th status update has successfully reached the destination. Then, we have $s_0 = l_0 = 0$ by assumption, and,

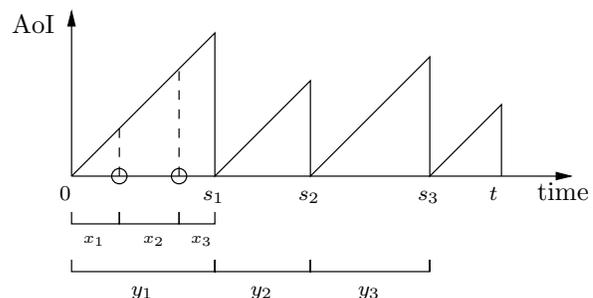


Fig. 2. Age evolution versus time with $n(t) = 3$ successful updates. Circles denote failed attempts. In this example, the first update is successfully received after three update attempts.

in general, $\{s_i\} \subseteq \{l_i\}$. Let $x_i \triangleq l_i - l_{i-1}$ and $y_i \triangleq s_i - s_{i-1}$ denote the inter-update attempt and the actual inter-update delays, respectively, and let $n(t)$ denote the number of updates that are successfully received by time t . We are interested in the average AoI given by the area under the age curve, see Fig. 2, which is given by

$$r(t) = \frac{1}{2} \sum_{i=1}^{n(t)} y_i^2 + \frac{1}{2} (t - s_{n(t)})^2. \quad (4)$$

The goal is to devise online feasible status update policies that minimize the long term average AoI through choosing the transmission times l_i 's (or equivalently the x_i 's). That is, to solve the following problem:

$$\begin{aligned} \min_{\{x_i\}} \quad & \limsup_{T \rightarrow \infty} \frac{1}{T} \mathbb{E}[r(T)] \\ \text{s.t.} \quad & (2) - (3), \end{aligned} \quad (5)$$

under causal knowledge of the energy arrival times, with no erasure feedback, and given the erasure probability q .

III. KEY CHARACTERISTIC OF THE OPTIMAL SOLUTION: RENEWAL STRUCTURE

In this section, we discuss a key characteristic of the optimal solution of problem (5). Specifically, we show that the optimal status update policy is a renewal policy, in which the *actual* inter-update times y_i 's are i.i.d. and that the *actual* update times s_i 's form a renewal process.

We first define some terminologies and notations. We use the term *epoch* to denote the time in between two consecutive successful updates. For instance, the i th epoch starts at time s_{i-1} and ends at s_i , and has a length of y_i time units. Note that an epoch may contain more than one update attempt, and the number of update attempts may vary from one epoch to another. We now slightly change notation to fit into our epoch definition and denote by $x_{i,k}$ the time in between the $(k-1)$ th and the k th update attempt in the i th epoch. Similarly, let $\tau_{i,k}$ denote the time until the k th energy arrival in the i th epoch *starting from the $(k-1)$ th update attempt*. For example, the first energy arrival in the i th epoch occurs at $s_{i-1} + \tau_{i,1}$, after which an update attempt occurs at $s_{i-1} + x_{i,1}$, with $x_{i,1} \geq \tau_{i,1}$ due to energy causality. Now say that this first update attempt

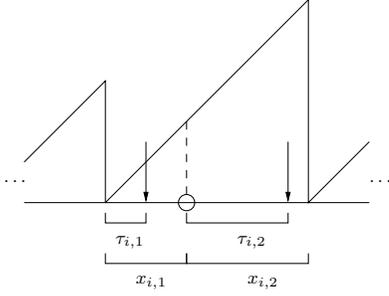


Fig. 3. AoI in the i th epoch with two update attempts. Arrows represent energy arrivals, and the circle denotes a failed update attempt.

has failed. Then, the sensor waits for the second energy arrival in the epoch occurring at $s_{i-1} + x_{i,1} + \tau_{i,2}$, after which the second update attempt occurs at $s_{i-1} + x_{i,1} + x_{i,2}$, and so on. Note that according to the definition of $\tau_{i,k}$'s, they do not necessarily represent the energy inter-arrival times, since $x_{i,k}$ can be strictly larger than $\tau_{i,k}$ (see Fig. 3).

Observe that transmission attempts occurring in the i th epoch may depend, in principal, on the history of events (transmission attempts and energy arrivals) that had occurred before the epoch started, which we denote by \mathcal{H}_{i-1} . Theorem 1 below shows, under some regularity conditions, that this is not the case; events in an epoch are independent of the history of events in previous epochs. Before we make that statement precise, we focus on the following special case of online policies, which are also the focus in [18], [26]:

Definition 1 (Uniformly Bounded Policy) *An online policy whose inter-update times have a bounded second moment.*

Intuitively, one would expect practical status update policies to be uniformly bounded as per Definition 1, so that the inter-update delays do not grow infinitely large (in expectation). We now state the main result of this section in Theorem 1 below. The proof of the theorem, which is fully presented in the Appendix, is similar in essence to the proofs of [26, Theorems 1 and 2] albeit some notable differences.

Theorem 1 *In the optimal solution of problem (5), any uniformly bounded policy is outperformed by a renewal policy in which the epoch lengths, y_i 's, are i.i.d.*

IV. OPTIMAL RENEWAL POLICIES: THRESHOLD POLICIES

We now analyze the best renewal policy and show that it has a threshold structure. Theorem 1 shows that epoch starting times, s_i 's, at which the system resets by making both the sensor's battery and the AoI drop to 0 simultaneously, constitute a renewal process. Since epoch lengths are i.i.d., we drop the subscript i from all random variables and denote the epoch duration by y and the inter-update attempt duration by x . Observe that we do not differentiate between different update attempts in a single epoch since the sensor is unaware of this information due to lack of erasure feedback. From the sensor's

point of view, it only designs a single inter-update attempt duration x . However, it takes the value of q into account while doing so as we show in the sequel. Let us (re)define τ as the time elapsed until the next energy arrival starting from the previous update attempt. Given that the sensor is unaware of erasure events, the sensor is ignorant of when the epoch ends. Therefore, inter-update attempt times in the epoch are functions of only the most recent energy arrival time. That is, x is only a function of τ .

By the strong law of large numbers for renewal processes (the renewal reward theorem) [28], problem (5) now reduces to an optimization over a single epoch as follows:

$$\begin{aligned} \min_{x(\cdot)} \quad & \frac{\mathbb{E}[R]}{\mathbb{E}[y]} \\ \text{s.t.} \quad & x(\tau) \geq \tau, \quad \forall \tau, \end{aligned} \quad (6)$$

where the random variable R denotes the area under the age curve in the epoch. We now introduce the following auxiliary problem to solve the above:

$$\begin{aligned} p(\lambda) \triangleq \min_{x(\cdot)} \quad & \mathbb{E}[R] - \lambda \mathbb{E}[y] \\ \text{s.t.} \quad & x(\tau) \geq \tau, \quad \forall \tau, \end{aligned} \quad (7)$$

for some $\lambda \geq 0$. One can show that the optimal solution of problem (6) is given by λ^* that solves $p(\lambda^*) = 0$ and that such λ^* is unique since $p(\lambda)$ is decreasing in λ [29]. Before solving problem (7), we first evaluate the terms $\mathbb{E}[y]$ and $\mathbb{E}[R]$ in what follows.

The expected epoch length can be found using iterated expectations by conditioning on how many erasure events occurred in it. We now write the following:

$$\begin{aligned} \mathbb{E}[y] &= (1-q)\mathbb{E}[x(\tau_1)] \\ &\quad + q(1-q)(\mathbb{E}[x(\tau_1)] + \mathbb{E}[x(\tau_2)]) \\ &\quad + q^2(1-q)(\mathbb{E}[x(\tau_1)] + \mathbb{E}[x(\tau_2)] + \mathbb{E}[x(\tau_3)]) \\ &\quad + \dots \quad (8) \\ &= \mathbb{E}[x(\tau)](1 + q + q^2 + \dots) \quad (9) \\ &= \frac{\mathbb{E}[x(\tau)]}{1-q}, \quad (10) \end{aligned}$$

where $\tau \sim \exp(1)$, and the second equality follows since τ_j 's are i.i.d. $\exp(1)$ random variables by the memoryless property of the exponential distribution. The expected area under the age curve in a single epoch can be found similarly as follows:

$$\begin{aligned} \mathbb{E}[R] &= (1-q)\frac{1}{2}\mathbb{E}[x^2(\tau_1)] \\ &\quad + q(1-q)\frac{1}{2}\mathbb{E}[(x(\tau_1) + x(\tau_2))^2] \\ &\quad + q^2(1-q)\frac{1}{2}\mathbb{E}[(x(\tau_1) + x(\tau_2) + x(\tau_3))^2] \\ &\quad + \dots \quad (11) \end{aligned}$$

$$\begin{aligned} &= \frac{1}{2}\mathbb{E}[x^2(\tau)](1 + q + q^2 + \dots) \\ &\quad + (\mathbb{E}[x(\tau)])^2(q + 2q^2 + 3q^3 + \dots) \quad (12) \end{aligned}$$

$$= \frac{\frac{1}{2}\mathbb{E}[x^2(\tau)]}{1-q} + \frac{q(\mathbb{E}[x(\tau)])^2}{(1-q)^2}, \quad (13)$$

where the second equality again follows since τ_j 's are i.i.d., and after some algebraic manipulations.

Using (10) and (13), one can write the following Lagrangian [30] for problem (7):

$$\begin{aligned} \mathcal{L} = & \frac{\frac{1}{2}\mathbb{E}[x^2(\tau)]}{1-q} + \frac{q(\mathbb{E}[x(\tau)])^2}{(1-q)^2} - \lambda \frac{\mathbb{E}[x(\tau)]}{1-q} \\ & - \int_0^\infty (x(\tau) - \tau) \eta(\tau) d\tau, \end{aligned} \quad (14)$$

where η is a Lagrange multiplier. Taking (the functional) derivative with respect to $x(t)$ and equating to 0, we get that the optimal x satisfies

$$x(t) = \lambda - \frac{2q}{1-q} \mathbb{E}[x(\tau)] + \frac{\eta(t)}{e^{-t}/1-q}. \quad (15)$$

Now let us define

$$\lambda' \triangleq \lambda - \frac{2q}{1-q} \mathbb{E}[x(\tau)]. \quad (16)$$

The sign of λ' has a major implication on the optimal policy's structure, which we discuss in detail next.

If $\lambda' < 0$ then we must have $\eta(t) > 0, \forall t$, to maintain positivity of $x(t)$. By complementary slackness [30] this further implies that $x(t) = t, \forall t$, i.e., a *greedy zero-wait policy* is optimal in this case, in which energy is used to send an update whenever it arrives. This case occurs for relatively high values of q which we specify precisely towards the end of this section. The value of $p(\lambda)$ in this case can be computed by plugging in $x(\tau) = \tau$ with $\mathbb{E}[x(\tau)] = 1$ and $\mathbb{E}[x^2(\tau)] = 2$ to get after some direct manipulations that

$$p(\lambda) = \frac{1 - \lambda(1-q)}{(1-q)^2}, \quad (17)$$

which admits an optimal long term average AoI, λ^* , of

$$\lambda^* = \frac{1}{1-q}. \quad (18)$$

Note that such greedy policy is always feasible and therefore (18) can generally serve as an upper bound on λ^* .

Now if $\lambda' \geq 0$, then by complementary slackness [30] we get that (see [18] and [26])

$$x(t) = \begin{cases} \lambda', & t < \lambda' \\ t, & t \geq \lambda' \end{cases}. \quad (19)$$

That is, the optimal status update policy is a λ' -*threshold policy*. Using this, one can directly compute $\mathbb{E}[x(\tau)] = \lambda' + e^{-\lambda'}$ and substitute back in (16) to get that

$$\frac{1+q}{1-q} \lambda' + \frac{2q}{1-q} e^{-\lambda'} = \lambda. \quad (20)$$

Direct first derivative analysis shows that the left hand side above is increasing in λ' for $\lambda' \geq 0$, and therefore, since its value at $\lambda' = 0$ is $2q/(1-q)$, (20) has a unique solution in λ' for every given $\lambda \geq 2q/(1-q)$, i.e., $2q/(1-q)$ is

the best achievable long term average AoI if $\lambda' \geq 0$. Now observe that for $q > 1/2$, the greedy zero-wait policy achieves a lower long term average AoI than that, given by $1/(1-q)$. We therefore conclude that in the optimal policy, λ' can only be non-negative if $q \leq 1/2$. Continuing with this assumption, we use (19), and some algebraic manipulations, to get

$$p(\lambda') = \frac{(1-q) \left(e^{-\lambda'} - \frac{1}{2} (\lambda')^2 \right) - q \left(\lambda' + e^{-\lambda'} \right)^2}{(1-q)^2}, \quad (21)$$

with λ' as defined in (20). Now observe that solving $p(\lambda') = 0$ for $\lambda' \geq 0$ is tantamount to having $p(0) \geq 0$ (since $p(\lambda)$ is monotonically decreasing [29] in λ , and λ is an increasing function of λ' from (20)). In other words, we must have

$$p(0) = \frac{1-2q}{(1-q)^2} \geq 0 \iff q \leq \frac{1}{2} \quad (22)$$

as assumed before.

In conclusion, the optimal policy's structure depends on the value of the erasure probability, q . If $q > \frac{1}{2}$ then (21) does not admit a positive λ' solution for $p(\lambda') = 0$, and therefore it holds that $\lambda' < 0$, and the greedy zero-wait policy is optimal. While if $q \leq \frac{1}{2}$ then the optimal policy is a λ' -threshold policy as in (19), with the optimal λ' solving $p(\lambda') = 0$.

We conclude this section by stating a few remarks and presenting some numerical examples to further illustrate the results of this paper. First, observe that for the case of no erasures, i.e., $q = 0$, (originally considered in [18]) we get from (20) and (21) that $\lambda' = \lambda$ and $p(\lambda) = e^{-\lambda} - \frac{1}{2} \lambda^2$, respectively, coinciding with the optimal solution in [18]. Second, for a given $\lambda \geq 0$, (20) shows that $\lambda' \leq \lambda$ with equality if and only if $q = 0$. This shows that the problem with erasures does *not* have the recurring property shown in [18], [26], [27] that the optimal long term average AoI equals the optimal threshold; they are only equal if $q = 0$.

As for the numerical examples, in Fig. 4 we plot the optimal AoI versus the erasure probability, along with the corresponding optimal threshold. For the case $q \leq \frac{1}{2}$, we basically start with a large enough value of λ' that makes $p(\lambda') < 0$, and then use a bisection search (in between 0 and that large enough value) to find λ' that solves $p(\lambda') = 0$. We then use (20) to find the optimal long term average AoI, λ^* . We also plot the optimal long term average AoI for the case $B = \infty$, which is shown in [19] to be equal to $\frac{1+q}{2(1-q)}$. Clearly, the solution for the $B = \infty$ case serves as a lower bound for the solution for the $B = 1$ case. From the figure, we also see that, quite intuitively, the larger the erasure probability, the larger the AoI, i.e., λ^* is monotonically increasing in q . In addition, we see that the optimal threshold λ' is monotonically decreasing in q . This is quite intuitive, since the sensor should be more eager to send new updates if the erasure probability is high, so that when the update is eventually received successfully the AoI would not be large.

V. CONCLUSION

The long term average age-minimal online policy has been derived for an energy harvesting sensor with unit-sized battery,

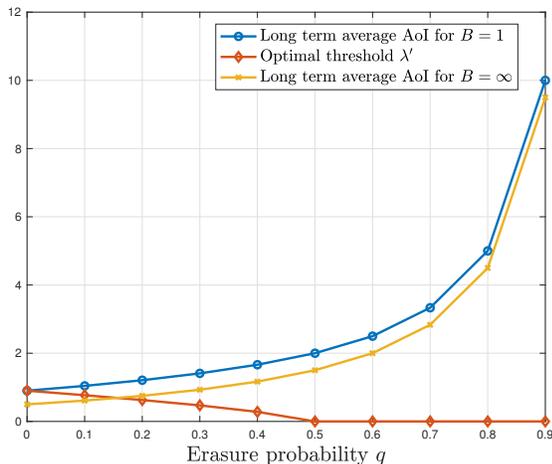


Fig. 4. Comparison of optimal AoI for $B = 1$ and $B = \infty$, along with λ' , versus q .

communicating with a destination over an erasure channel, without erasure error feedback information. Firstly, a key structure of the optimal policy has been shown, namely, that transmission times should be such that updates reach the destination successfully at times that constitute a renewal process. Then, it has been shown that the optimal renewal-type policy has an erasure probability-dependent threshold structure, in which a new update is transmitted only if the AoI grows above a certain threshold value. The optimal threshold has been shown to be decreasing with the erasure probability, reaching exactly 0 when the erasure probability is $\frac{1}{2}$, and staying at 0 afterwards. Such 0-threshold policies are interpreted as greedy policies, in which a new update is sent whenever an energy arrives. The rationale is that with higher erasure probability the sensor should be more aggressive in transmitting new updates, so that when they eventually reach the destination successfully the AoI would not be too large.

APPENDIX PROOF OF THEOREM 1

We follow an indirect approach to prove the theorem. Basically, we derive an achievable lower bound on the long term average AoI using renewal-type policies in a *genie-aided* system in which there exists a genie that informs the sensor when updates are successful, i.e., the epochs' start times. However, we enforce a constraint on the sensor *not to use* the lack of this piece of information to infer that its update is unsuccessful and act accordingly to change its policy within the same epoch. This seemingly unintuitive constraint simplifies the proof as we will see later on. Now observe that such genie-aided system cannot perform worse than the original system that we consider in this paper, and hence, a lower bound on this genie-aided system is also a lower bound on the original one. We then conclude the proof by showing that such lower bound is also achievable in the original system by showing that the optimal renewal-type policy does not

actually need the information provided by the genie, thereby proving optimality of renewal-type policies in the original system as well. Next, we provide the details.

In the genie-aided system, consider any online feasible uniformly bounded policy $\{x_{i,k}\}$. Focusing on the i th epoch, let us denote by $R_{i,m}$ the area under the age curve in the i th epoch given that it went through m update attempts, and by R_i the area under the age curve in it *irrespective* of how many update attempts. Let us also denote by $e_{i,k}$ the event that the k th update attempt in the i th epoch gets erased. We can now write the following:

$$R_{i,m} = \frac{1}{2} (x_{i,1} + x_{i,2} + \dots + x_{i,m})^2, \quad (23)$$

$$R_i = \sum_{m=1}^{\infty} R_{i,m} \cdot \prod_{k=1}^{m-1} \mathbb{1}(e_{i,k}) \mathbb{1}(e_{i,m}^c), \quad (24)$$

where $\mathbb{1}(\cdot)$ is the indicator function, and the superscript c denotes the complement of an event.

Next, for a fixed time T , denote by N_T the number of epochs that have already *started* by time T . Given the history before the i th epoch, \mathcal{H}_{i-1} , and the number of update attempts in the i th epoch, m , let us define the vector $\tau_i^{(m)} \triangleq [\tau_{i,1}, \tau_{i,2}, \dots, \tau_{i,m}]$, and define the following statistical average of the area under the age curve in the i th epoch with m update attempts:

$$\hat{R}_{i,m}(\gamma^{(m)}, \mathcal{H}_{i-1}) \triangleq \mathbb{E} [R_{i,m} | \tau_i^{(m)} = \gamma^{(m)}, \mathcal{H}_{i-1}]. \quad (25)$$

Therefore, it holds that

$$\begin{aligned} & \mathbb{E} [R_{i,m} \mathbb{1}(i \leq N_T)] \\ &= \mathbb{E}_{\mathcal{H}_{i-1}} [\mathbb{E}_{\tau_i^{(m)}} [\hat{R}_{i,m}(\gamma^{(m)}, \mathcal{H}_{i-1})] \mathbb{1}(i \leq N_T) | \mathcal{H}_{i-1}] \end{aligned} \quad (26)$$

since $\mathbb{1}(i \leq N_T)$ is independent of $\tau_i^{(m)}$ given \mathcal{H}_{i-1} . We can similarly define the following statistical average of the i th epoch length with m update attempts:

$$\hat{y}_{i,m}(\gamma^{(m)}, \mathcal{H}_{i-1}) \triangleq \mathbb{E} [y_{i,m} | \tau_i^{(m)} = \gamma^{(m)}, \mathcal{H}_{i-1}]. \quad (27)$$

Next, observe that by (4) the following holds:

$$\frac{1}{T} \sum_{i=1}^{\infty} R_i \mathbb{1}(i \leq N_T - 1) \leq \frac{r(T)}{T} \leq \frac{1}{T} \sum_{i=1}^{\infty} R_i \mathbb{1}(i \leq N_T). \quad (28)$$

Following similar analysis as in [18, Appendix C-1], one can show that the term $\mathbb{E}[R_{N_T}]/T \rightarrow 0$ as $T \rightarrow \infty$, making the upper and lower bounds in (28) equal as $T \rightarrow \infty$. Therefore, we proceed by deriving a lower bound on $\frac{1}{T} \mathbb{E} [\sum_{i=1}^{\infty} R_i \mathbb{1}(i \leq N_T)]$ and conclude that it shall also serve as a lower bound on $\frac{1}{T} \mathbb{E} [r(T)]$ as $T \rightarrow \infty$ (the objective function of problem (5)). We do so through a series of inequalities at the top of the next page. There, (29) follows since, by definition of N_T , it holds that $\mathbb{E} [\sum_{i=1}^{\infty} y_i \mathbb{1}(i \leq N_T)] \geq T$; (30) follows by (24); (31) follows by the monotone convergence theorem, and the fact that erasure events are mutually independent and are independent of transmissions; (32) follows by (26); (33) follows again by the monotone

$$\begin{aligned}
& \frac{1}{T} \mathbb{E} \left[\sum_{i=1}^{\infty} R_i \mathbb{1}(i \leq N_T) \right] \\
& \geq \frac{\mathbb{E} \left[\sum_{i=1}^{\infty} R_i \mathbb{1}(i \leq N_T) \right]}{\mathbb{E} \left[\sum_{i=1}^{\infty} y_i \mathbb{1}(i \leq N_T) \right]} \tag{29} \\
& = \frac{\mathbb{E} \left[\sum_{i=1}^{\infty} \sum_{m=1}^{\infty} R_{i,m} \prod_{k=1}^{m-1} \mathbb{1}(e_{i,k}) \mathbb{1}(e_{i,m}^c) \mathbb{1}(i \leq N_T) \right]}{\mathbb{E} \left[\sum_{i=1}^{\infty} y_i \mathbb{1}(i \leq N_T) \right]} \tag{30} \\
& = \frac{\sum_{i=1}^{\infty} \sum_{m=1}^{\infty} q^{m-1} (1-q) \mathbb{E} [R_{i,m} \mathbb{1}(i \leq N_T)]}{\sum_{i=1}^{\infty} \mathbb{E} [y_i \mathbb{1}(i \leq N_T)]} \tag{31} \\
& = \frac{\sum_{i=1}^{\infty} \sum_{m=1}^{\infty} q^{m-1} (1-q) \mathbb{E}_{\mathcal{H}_{i-1}} \left[\mathbb{E}_{\tau_i^{(m)}} \left[\hat{R}_{i,m}(\gamma^{(m)}, \mathcal{H}_{i-1}) \right] \mathbb{1}(i \leq N_T) \middle| \mathcal{H}_{i-1} \right]}{\sum_{i=1}^{\infty} \mathbb{E} [y_i \mathbb{1}(i \leq N_T)]} \tag{32} \\
& = \frac{\sum_{i=1}^{\infty} \mathbb{E}_{\mathcal{H}_{i-1}} \left[\sum_{m=1}^{\infty} q^{m-1} (1-q) \mathbb{E}_{\tau_i^{(m)}} \left[\hat{R}_{i,m}(\gamma^{(m)}, \mathcal{H}_{i-1}) \right] \mathbb{1}(i \leq N_T) \middle| \mathcal{H}_{i-1} \right]}{\sum_{i=1}^{\infty} \mathbb{E} [y_i \mathbb{1}(i \leq N_T)]} \tag{33} \\
& = \frac{\sum_{i=1}^{\infty} \mathbb{E}_{\mathcal{H}_{i-1}} \left[\sum_{m=1}^{\infty} q^{m-1} (1-q) \mathbb{E}_{\tau_i^{(m)}} \left[\hat{y}_{i,m}(\gamma^{(m)}, \mathcal{H}_{i-1}) \right] \frac{\sum_{m=1}^{\infty} q^{m-1} (1-q) \mathbb{E}_{\tau_i^{(m)}} \left[\hat{R}_{i,m}(\gamma^{(m)}, \mathcal{H}_{i-1}) \right]}{\sum_{m=1}^{\infty} q^{m-1} (1-q) \mathbb{E}_{\tau_i^{(m)}} \left[\hat{y}_{i,m}(\gamma^{(m)}, \mathcal{H}_{i-1}) \right]} \mathbb{1}(i \leq N_T) \middle| \mathcal{H}_{i-1} \right]}{\sum_{i=1}^{\infty} \mathbb{E} [y_i \mathbb{1}(i \leq N_T)]} \tag{34} \\
& \geq \frac{\sum_{i=1}^{\infty} \mathbb{E}_{\mathcal{H}_{i-1}} \left[\sum_{m=1}^{\infty} q^{m-1} (1-q) \mathbb{E}_{\tau_i^{(m)}} \left[\hat{y}_{i,m}(\gamma^{(m)}, \mathcal{H}_{i-1}) \right] R^*(\mathcal{H}_{i-1}) \mathbb{1}(i \leq N_T) \middle| \mathcal{H}_{i-1} \right]}{\sum_{i=1}^{\infty} \mathbb{E} [y_i \mathbb{1}(i \leq N_T)]} \tag{35} \\
& \geq \frac{\sum_{i=1}^{\infty} \sum_{m=1}^{\infty} q^{m-1} (1-q) \mathbb{E}_{\mathcal{H}_{i-1}} \left[\mathbb{E}_{\tau_i^{(m)}} \left[\hat{y}_{i,m}(\gamma^{(m)}, \mathcal{H}_{i-1}) \right] \mathbb{1}(i \leq N_T) \middle| \mathcal{H}_{i-1} \right]}{\sum_{i=1}^{\infty} \mathbb{E} [y_i \mathbb{1}(i \leq N_T)]} R_{\min} \tag{36} \\
& = \frac{\sum_{i=1}^{\infty} \sum_{m=1}^{\infty} q^{m-1} (1-q) \mathbb{E} [y_{i,m} \mathbb{1}(i \leq N_T)]}{\sum_{i=1}^{\infty} \mathbb{E} [y_i \mathbb{1}(i \leq N_T)]} R_{\min} \tag{37} \\
& = R_{\min}. \tag{38}
\end{aligned}$$

convergence theorem; $R^*(\mathcal{H}_{i-1})$ in (35) denotes the minimum value of $\frac{\sum_{m=1}^{\infty} q^{m-1} (1-q) \mathbb{E}_{\tau_i^{(m)}} [\hat{R}_{i,m}(\gamma^{(m)}, \mathcal{H}_{i-1})]}{\sum_{m=1}^{\infty} q^{m-1} (1-q) \mathbb{E}_{\tau_i^{(m)}} [\hat{y}_{i,m}(\gamma^{(m)}, \mathcal{H}_{i-1})]}$; R_{\min} in (36) denotes the minimum value of $R^*(\mathcal{H}_{i-1})$ over all epochs and their corresponding histories, i.e., the minimum over all i and \mathcal{H}_{i-1} ; and (37) and (38) follow by the relationships between $\hat{y}_{i,m}$, $y_{i,m}$, and y_i which are the same as those between $\hat{R}_{i,m}$, $R_{i,m}$, and R_i that got us from (29) to (32).

Note that the online policy achieving $R^*(\mathcal{H}_{i-1})$ is only a function of the energy arrivals in the i th epoch, since the history \mathcal{H}_{i-1} is fixed. Now observe that by the memoryless property of exponential distribution, $\tau_{i,k}$'s are i.i.d. $\sim \exp(1)$. Therefore, if one repeats the policy that achieves R_{\min} over all epochs, which is possible since the genie provides information about epochs' start times, then one gets a renewal policy in which y_i 's are i.i.d.

We now argue that the best renewal policy does not depend on the genie's provided information. First, it is clear that when an epoch starts, the sensor's next inter-update attempt becomes independent of the past and only a function of the energy arrivals in the epoch, in particular the first arrival time. If the sensor receives an information from the genie that its first

update was successful, then this means a new epoch started and the process is repeated. On the other hand, if it does not hear from the genie, it is not allowed to act upon that information according to our enforced constraint that we stated at the beginning of the proof. Hence, it repeats the same policy, otherwise the constraint would be violated. Therefore, the policy does not change whether the genie sends its information or not. Finally, observe that this policy is achievable in the original system considered in this paper, i.e., the system with no genie. This completes the proof.

REFERENCES

- [1] S. Kaul, R. D. Yates, and M. Gruteser. Real-time status: How often should one update? In *Proc. IEEE Infocom*, March 2012.
- [2] R. D. Yates and S. Kaul. Real-time status updating: Multiple sources. In *Proc. IEEE ISIT*, July 2012.
- [3] C. Kam, S. Kompella, and A. Ephremides. Age of information under random updates. In *Proc. IEEE ISIT*, July 2013.
- [4] A. Kosta, N. Pappas, A. Ephremides, and V. Angelakis. Age and value of information: Non-linear age case. In *Proc. IEEE ISIT*, June 2017.
- [5] A. M. Bedewy, Y. Sun, and N. B. Shroff. Age-optimal information updates in multihop networks. In *Proc. IEEE ISIT*, June, 2017.
- [6] Y. Hsu, E. Modiano, and L. Duan. Age of information: Design and analysis of optimal scheduling algorithms. In *Proc. IEEE ISIT*, June 2017.

- [7] E. Najm and E. Telatar. Status updates in a multi-stream M/G/1/1 preemptive queue. In *Proc. IEEE Infocom*, April 2018.
- [8] 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. Inf. Theory*, 63(11):7492–7508, November 2017.
- [9] R. D. Yates, E. Najm, E. Soljanin, and J. Zhong. Timely updates over an erasure channel. In *Proc. IEEE ISIT*, June 2017.
- [10] R. D. Yates, P. Ciblat, A. Yener, and M. A. Wigger. Age-optimal constrained cache updating. In *Proc. IEEE ISIT*, June 2017.
- [11] Y. Sun, Y. Polyanskiy, and E. Uysal-Biyikoglu. Remote estimation of the Wiener process over a channel with random delay. In *Proc. IEEE ISIT*, June 2017. Longer version available: arXiv:1701.06734.
- [12] A. Kosta, N. Pappas, and V. Angelakis. Age of information: A new concept, metric, and tool. *Foundations and Trends in Networking*, 12(3):162–259, 2017.
- [13] R. D. Yates. Lazy is timely: Status updates by an energy harvesting source. In *Proc. IEEE ISIT*, June 2015.
- [14] B. T. Bacinoglu, E. T. Ceran, and E. Uysal-Biyikoglu. Age of information under energy replenishment constraints. In *Proc. ITA*, February 2015.
- [15] A. Arafa and S. Ulukus. Age-minimal transmission in energy harvesting two-hop networks. In *Proc. IEEE Globecom*, December 2017.
- [16] A. Arafa and S. Ulukus. Age minimization in energy harvesting communications: Energy-controlled delays. In *Proc. Asilomar*, October 2017.
- [17] B. T. Bacinoglu and E. Uysal-Biyikoglu. Scheduling status updates to minimize age of information with an energy harvesting sensor. In *Proc. IEEE ISIT*, June 2017.
- [18] X. Wu, J. Yang, and J. Wu. Optimal status update for age of information minimization with an energy harvesting source. *IEEE Trans. Green Commun. Netw.*, 2(1):193–204, March 2018.
- [19] S. Feng and J. Yang. Optimal status updating for an energy harvesting sensor with a noisy channel. In *Proc. IEEE Infocom*, April 2018.
- [20] S. Feng and J. Yang. Minimizing age of information for an energy harvesting source with updating failures. In *Proc. IEEE ISIT*, June 2018.
- [21] A. Baknina and S. Ulukus. Coded status updates in an energy harvesting erasure channel. In *Proc. CISS*, March 2018.
- [22] S. Farazi, A. G. Klein, and D. R. Brown III. Average age of information for status update systems with an energy harvesting server. In *Proc. IEEE Infocom*, April 2018.
- [23] A. Baknina, O. Ozel, J. Yang, S. Ulukus, and A. Yener. Sending information through status updates. In *Proc. IEEE ISIT*, June 2018.
- [24] A. Arafa, J. Yang, S. Ulukus, and H. V. Poor. Age-minimal online policies for energy harvesting sensors with incremental battery recharges. In *Proc. ITA*, February 2018.
- [25] A. Arafa, J. Yang, and S. Ulukus. Age-minimal online policies for energy harvesting sensors with random battery recharges. In *Proc. IEEE ICC*, May 2018.
- [26] A. Arafa, J. Yang, S. Ulukus, and H. V. Poor. Age-minimal transmission for energy harvesting sensors with finite batteries: Online policies. Available Online: arXiv:1806.07271.
- [27] B. T. Bacinoglu, Y. Sun, E. Uysal-Biyikoglu, and V. Mutlu. Achieving the age-energy tradeoff with a finite-battery energy harvesting source. In *Proc. IEEE ISIT*, June 2018.
- [28] S. M. Ross. *Stochastic Processes*. Wiley, 1996.
- [29] W. Dinkelbach. On nonlinear fractional programming. *Management Science*, 13(7):492–498, 1967.
- [30] S. P. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.