Age-Optimal Causal Labeling of Memoryless Processes

Yunus İnan, Emre Telatar
EPFL, Lausanne, Switzerland
Email: {yunus.inan, emre.telatar}@epfl.ch

***Abstract*—We consider a problem of labeling a memoryless temporal point process. The labelings have to be done causally. We study the tradeoff between the rate and age of the labeled process and characterize the optimal tradeoff curve. We show that the optimal curve can be achieved with rather simple labeling procedures which repeatedly do the following: Wait for T time units and label the next arrival.**

***Index Terms*—Age of Information, Causal Labeling, Memoryless Process, Markov Decision Process, Policy Updates**

I. Introduction

Timely data transmission requirements on large-scale networks have raised the importance of data transmission protocols that classify and convey the freshest data to keep the receivers up-to-date. *Age-of-Information* (AoI), a metric proposed by Kaul et al. [1], is well-suited for evaluation of such protocols, and studied extensively in recent work — see [2], [3] for detailed surveys. Now, think of a component in the network, which receives data at a higher rate than it is allowed to send. In this case, one may ask (i) what and (ii) when to send with an aim to optimize AoI? Studying the first question may need classification of data according to their importance, which might be based on a distortion metric [4]–[6]; or one may resort to packet management techniques [7] in the absence of such metrics. The second question is addressed for example in [8]–[10], and in [11]–[16] when the throughput is limited due to energy constraints. All of these studies assume *a priori* knowledge of the network dynamics, e.g., the data is known to be conveyed through a single-server queue [4], [5], or through a Multiple Access Channel [6], etc. In this work, we assume that the sender is oblivious of the network dynamics except the limit on the output data rate.

Our work focuses on arrivals modeled as a memoryless point process with an arrival rate greater than the limited output rate. The sender thus needs to filter out some of the arrivals. We call this filtering operation a 'labeling procedure', where an arrival is passed through the network as soon as it is labeled. We consider causal labelings, i.e., a data can be labeled only after it has arrived. These labeling procedures relate to both questions (i) and (ii) above. As the sender is oblivious of the network dynamics, we study the tradeoff between the rate and the age of the labeled process, which are to be defined in Section II.

The outline of this work is as follows: In Section II, we provide the problem definition in a discrete-time setting. The rate-age tradeoff is related to an appropriate Markov Decision Problem (MDP) formulation. In Section III, we study a finite-state approximation which allows to characterize the optimal labeling procedures for the original one. The optimal procedures turn out to be rather simple: Wait T time slots and label the next arrival, where T is tuned to match the output rate. In Section IV, we extend the results to a continuous-time model, where the arrivals are modeled as a Poisson process.

II. Problem Definition

Consider a discrete-time memoryless point process, where the interarrival times $Z_1, Z_2, \dots$ are independent and identically distributed (i.i.d.) with Geometric distribution. Suppose their success probability is p, which is also equal to the rate of the arrival process. The arrivals can also be modeled as i.i.d. Bernoulli random variables X_t with success probability p, with the natural filtration $\{\mathcal{F}_t\}$ where $\mathcal{F}_t := \sigma(X_1, \dots, X_t)$. For convenience assume $X_0 = 1$. We define a causal labeling on the point process as a sequence of (possibly random) labeling functions adapted to the filtration $\{\mathcal{F}_t\}$ as $S_t : \{0,1\}^t \to \{1, \dots, t\} \cup \{?\}$, with the following restriction that if $S_t \neq ?$, then (i) X_{S_t} must be 1 and (ii) $S_t \neq S_\tau$ for all $\tau < t$. Arrival X_{S_t} is labeled at time t, and 'phantom' labelings and multiple labelings of a single point are not allowed. We assume $S_0 = 0$.

Given a procedure $\{S_t\}$, we define the rate R as the expected long term average of the number of labelings. More precisely,

$$R^{(S)} := E\left[\limsup_{t\to\infty} \frac{1}{t}\sum_{\tau=1}^{t} \mathbb{1}\{S_t \neq ?\}\right]. \tag{1}$$

Define the most recent labeling at time t as $M_t := \max\{S_\tau : 0 \le \tau < t, S_\tau \neq ?\}$ with $M_0 = 0$. The instantaneous age $\Delta_t^{(S)}$ and average age $\Delta^{(S)}$ are then defined as

$$\Delta_t^{(S)} := t - M_t, \quad \Delta^{(S)} := E\left[\limsup_{t\to\infty} \frac{1}{t}\sum_{\tau=1}^{t} \Delta_\tau^{(S)}\right]. \tag{2}$$

Sometimes we omit the superscript (S) from the above expressions for brevity. At this point, it might be useful to give the following examples of different labeling procedures to demonstrate what they resemble in practice.

Example 1. *Label every point upon arrival with probability α. This labeling procedure constitutes a renewal process whose interarrival times W are Geometric random variables with parameter αp. Hence, the rate and age of this renewal process are $R = \alpha p$ and $\Delta = \frac{E[W(W+1)]}{2E[W]} = \frac{1}{\alpha p} = \frac{1}{R}$.*

Example 2. *Label every k^{th} point upon arrival. Likewise, this procedure yields a renewal process whose interarrival times are sum of k Geometric random variables with success probability p. Hence, $R = \frac{p}{k}$, and $\Delta = \frac{k+1}{2p} = \frac{1+1/k}{2}\frac{1}{R}$, which is strictly smaller than the age resulting from the labeling procedure in Example 1 for $k > 1$. This labeling procedure can be extended to cover the rates $R = pr$ for rational r and* *the resulting age can be shown to be smaller than the one in Example 1 as well, see Appendix A.*

We are interested in finding the achievable region of possible (R, Δ) pairs with causal labeling procedures. More precisely, we are interested in finding the boundary curve of such pairs. Given the large class of possible labelings, this search seems difficult at first sight. However, we can eliminate some of the labeling procedures to make the search tractable.

Definition 1 (Strictly Increasing Procedures)**.** *A labeling procedure $\{S_t\}$ is strictly increasing if its subsequence $\{S_t : S_t \neq ?\}$ is strictly increasing.*

Lemma 1. *For any labeling procedure $\{S_t\}$, there exists a strictly increasing modification $\{\tilde{S}_t\}$ such that $R^{(\tilde{S})} \leq R^{(S)}$ and $\Delta^{(\tilde{S})} = \Delta^{(S)}$.*

Proof: Take $\tilde{S}_t = S_t$ if $S_t = M_{t+1}$; otherwise $\tilde{S}_t = ?$. Then, $\Delta_t^{(\tilde{S})} = \Delta_t^{(S)}$ and $R_t^{(\tilde{S})} \leq R_t^{(S)}$, where $R_t^{(S)} = \sum_{\tau=1}^{t} \mathbb{1}\{S_\tau \neq ?\}$. Thus, $\tilde{R} \leq R$ and $\tilde{\Delta} = \Delta$. ∎

Lemma 1 can be interpreted as follows: At time t, a strictly increasing modification of $\{S_t\}$ will consider arrivals after the most recent labeling M_t, and the (R, Δ) pair pertaining to this modification will be closer to the boundary curve we are trying to find. Hence we may focus on strictly increasing labelings that omit arrivals before and including M_t.

We further restrict the space of labeling procedures we are interested in by introducing the lemma below.

Lemma 2. *For any $t_1 < t_2 \leq \tau$ such that $X_{t_1} = X_{t_2} = 1$, and for every strictly increasing $\{S_t\}$ with $S_\tau = t_1$, there exists a strictly increasing modification $\{\tilde{S}_t\}$ with $\tilde{S}_\tau = t_2$ and $\Delta_t^{(\tilde{S})} \leq \Delta_t^{(S)}$, $R_t^{(\tilde{S})} \leq R_t^{(S)}$. Consequently, $R^{(\tilde{S})} \leq R^{(S)}$ and $\Delta^{(\tilde{S})} \leq \Delta^{(S)}$.*

Proof: Define $\{\tilde{S}_t\}$ such that $\tilde{S}_t = S_t$ for $t < \tau$, $\tilde{S}_\tau = t_2$ and for all $t > \tau$,

$$\tilde{S}_t = \begin{cases} S_t, & S_t > t_2 \\ ?, & \text{else} \end{cases}. \tag{3}$$

In words, whenever an arrival later than t_2 is labeled by $\{S_t\}$, it is also labeled by $\{\tilde{S}_t\}$. Observe $\tilde{M}_t \geq M_t$ and $R_t^{(\tilde{S})} \leq R_t^{(S)}$, thus $\Delta^{(\tilde{S})} \leq \Delta^{(S)}$ and $R^{(\tilde{S})} \leq R^{(S)}$. ∎

Set the index of the freshest arrival as $I_t := \max\{\tau \leq t : X_\tau = 1\}$. Observe that Lemmas 1 and 2 imply

Corollary 1. *For any $\{S_t\}$, there exists a strictly increasing modification $\{\tilde{S}_t\}$ such that $\tilde{S}_t = I_t$ or $\tilde{S}_t = ?$ and $R^{(\tilde{S})} \leq R^{(S)}$, $\Delta^{(\tilde{S})} \leq \Delta^{(S)}$.*

Corollary 1 tells that one should examine the procedures that label only the freshest arrival. Let $\mathcal{S}_F$ be the space of such labeling procedures. The following theorem gives a lower bound to $CR + \Delta$, $C > 0$ for a specific class of labeling functions and hence gives a lower bound to the boundary curve of feasible (R, Δ) region.

Theorem 1. *For $\{S_t\} \in \mathcal{S}_F$ such that $\sup_t E[\Delta_t^2] < \infty$,*

$$CR + \Delta \geq \limsup_{t\to\infty} \frac{1}{t} \sum_{\tau=1}^{t} E\big[C\mathbb{1}\{S_\tau \neq ?\} + \Delta_\tau\big]. \tag{4}$$

Proof: Since $\frac{1}{t}\sum_{\tau=1}^{t} \mathbb{1}\{S_t = ?\}$ is bounded for all t, we directly apply Reverse Fatou's lemma to the first term on the left-hand side and obtain $R \geq \limsup_{t\to\infty} \frac{1}{t}\sum_{\tau=1}^{t} \Pr\{S_t = ?\}$. Since $\sup_t E[\Delta_t^2]$ is finite, $\sup_t E\left[\left(\frac{1}{t}\sum_{\tau=1}^{t} \Delta_t\right)^2\right]$ is also finite and constitutes a uniformly integrable family; allowing the use of Reverse Fatou's lemma [17]. Thus, $\Delta \geq \limsup_{t\to\infty} \frac{1}{t}\sum_{\tau=1}^{t} E[\Delta_\tau]$. Lastly, we observe

$$\begin{aligned} CR + \Delta &\geq \limsup_{t\to\infty} \frac{C}{t} \sum_{\tau=1}^{t} \Pr\{S_t = ?\} + \limsup_{t\to\infty} \frac{1}{t} \sum_{\tau=1}^{t} E[\Delta_\tau] \\ &\geq \limsup_{t\to\infty} \frac{1}{t} \sum_{\tau=1}^{t} E\big[C\mathbb{1}\{S_\tau \neq ?\} + \Delta_\tau\big]. \end{aligned}$$ ∎

The expression on the right-hand side of (4) contains a summation whose τ^{th} term is $\mathcal{F}_\tau$-measurable. This tells that for any $\{S_t\}$, this expression is the average reward (cost, in our case) of a Markov Reward Process with state space $\mathcal{B} := \{0,1\}^*$ and the problem of choosing an appropriate labeling procedure $\{S_t\}$ is a Markov Decision Problem (MDP), which is formulated as finding the infimal limsup average cost

$$\lambda^* := \inf_{\{S_t\}\in\mathcal{S}_F} J^{\{S_t\}}(C), \text{ where}$$

$$J^{\{S_t\}}(C) := \limsup_{t} \frac{1}{t} \sum_{\tau=1}^{t} E\big[C\mathbb{1}\{S_\tau \neq ?\} + \Delta_\tau\big].$$

In the current problem formulation, and given the definition of the labeling procedures at the beginning of this manuscript, the states and actions seem to be complicated. However, Corollary 1 tells that it is sufficient to consider only two actions: (i) label the freshest arrival or (ii) wait, which we denote by l and w respectively. Since all arrivals before I_t are ignored at time t, we can reduce the state space to binary strings $\boldsymbol{B}_t$ of length $t - M_t$, with its $(I_t - M_t)^{\text{th}}$ element being 1 and its other elements being 0, e.g., [000100].

Remark 1. *At this point, we have not imposed that the labeling strategies depend only on the current state $\boldsymbol{B}_t$. The procedures can depend on the whole past (also called history dependent). Thus, one may argue that by reducing the state space, the history may not be recovered. However, this is not true as one is able to construct $X_1, \ldots, X_t$ from $\boldsymbol{B}_1, \ldots, \boldsymbol{B}_t$.*

Note that $\boldsymbol{B}_t$'s are binary strings which contain at most a single 1. We can represent such a string by a pair (m, n) of non-negative integers, where m is the number of elements from the beginning of the string until and including the 1, and n is the remaining number of zeros. For instance, the buffer content [000100] becomes $(4, 2)$; and [000] becomes $(0, 3)$.

The current problem is classified as *countable-state average cost MDP* [18]. This class of problems is in general difficult to work with and it is not guaranteed that an optimal stationary policy exists. Even when it exists, it can be difficult to find such an optimal strategy. In the next section, we give a finite-state approximation to the problem with an aim to use the methods for finite-state problems; and as we will see, the finite-state formulation fortunately allows us to characterize the optimal strategies for the countable-state model as well.

III. Finite-state approximation

Assume 'phantom' arrivals are generated whenever $m+n = L$, with no sending cost, i.e., do not contribute to R. With a similar proof as in Lemma 2, one can show that if the length of the buffer reaches L, procedures that label the phantom arrival will have a smaller Δ compared to the procedures that wait instead. Hence, the boundary curve of pairs pertaining to such procedures will lie under the original (R, Δ) curve. This truncated problem is finite state. Furthermore, since the state $(0,L)$ is recurrent under any policy — because the Geometric distribution has infinite tail and thus the buffer length reaches L infinitely often — the problem is unichain, i.e., every policy induces a Markov Chain with a single recurrent class [18]. Define

$$\lambda_L^* := \inf_{\{S_t\}\in\mathcal{S}_{F,L}} J^{\{S_t\}}(C)$$

where $\mathcal{S}_{F,L}$ is the set of labeling procedures that only label the freshest arrival and always label the phantom arrival. Observe that λ_L^* is non-decreasing with L. Moreover, the sequence $\{\lambda_L^*\}$ has a limit. To see this, take the strategy 'label every point upon arrival' in the untruncated problem. The average cost corresponding to this strategy will be $\frac{C}{E[Z]} + \frac{E[Z(Z+1)]}{2E[Z]} = Cp + 1/p$. Then, $\lambda_L^* \le Cp + 1/p$ for all L. Since $\{\lambda_L^*\}$ is bounded from above and is non-decreasing, it has a limit which we denote by λ_∞^*.

At this moment, the problem has become finite state and unichain; and it is known that there exists an optimal stationary policy for such problems. One may therefore focus on stationary policies and their evaluation methods. A stationary policy $s : \mathbb{N}\times\mathbb{N} \to \{l, w\}$ in our problem is evaluated by solving for λ, $\boldsymbol{h}$ in the system of linear equations below. [18]

$$h(m,n)+\lambda = \begin{cases} m+n+C + ph(n+1,0) + qh(0,n+1), & s(m,n)=l,\ m+n<L,\ m\ge 1 \\ m+n + ph(m+n+1,0) + qh(m,n+1), & s(m,n)=w,\ m+n<L \\ L+ph(1,0) + qh(0,1), & m+n=L \end{cases} \tag{5}$$

where $q := 1-p$. We choose the state $(1,0)$, i.e., the buffer content [1], as the reference state and set $h(1,0)=0$. λ gives the average cost of the unichain stationary labeling policy and $h(m,n)$ is called relative value of the state (m,n) [18].

Remark 2. *$h(m,n)+\lambda$ is equal to the one-step cost plus the expected relative value of the next state depending on the action. E.g., if $s(m,n)=l$, then the one-step cost is $m+n+C$ and the next state will be $(n+1,0)$ with probability p and $(0,n+1)$ with probability q. Thus for any policy, the state transitions are inferred from* (5).

Let us provide a brief summary on policy updates. Choose a stationary policy $s(m,n)$, evaluate it by solving (5) and obtain the relative values $h(m,n)$, $m+n \le L$ together with the average cost λ. Given the relative values, take a state (m_0,n_0), $m_0+n_0 < L$, $m_0 \ge 1$ and consider the policy

$$s'(m_0,n_0) = \begin{cases} l, & C+ph(n_0+1,0)+qh(0,n_0+1) \le ph(m_0+n_0+1,0)+qh(m_0,1) \\ w, & \text{else} \end{cases} \tag{6}$$

and $s(m,n)=s'(m,n)$ for all other states. Now, solve (5) with respect to $s'(m,n)$ to obtain λ'. It is known that $\lambda' \le \lambda$ [18], i.e., the policy s is updated to a better policy s'. In fact, if one does the above procedure not only for (m_0,n_0), but also for every possible state, and then solves (5), the procedure is the well-known policy iteration algorithm, see [18] for example.

We aim to characterize the optimal strategies by using policy updates as a tool. Instead of direct application of the generic policy iteration algorithm, we consider the procedure described above where at k^{th} step we choose a single state (m_k,n_k), update $s(m_k,n_k)$ and solve (5). Denote the average cost at the end of k^{th} step as $\lambda_L^{(k)}$, denote the updated policy and the relative values as $\boldsymbol{s}^{(k)}$ and $\boldsymbol{h}^{(k)}$ respectively.

Start the procedure with the initial policy 'label every point upon arrival'. That is, $s^{(0)}(m,n)=l$ if $m>0$ and $n=0$; otherwise it is equal to w. Note that $s^{(0)}(0,0)=w$ and although the state $(0,0)$ will never be encountered according to our formulation, it provides convenience in the description of the procedure. The average cost corresponding to this policy is $\lambda_L^{(0)} \le Cp+1/p$.

At k^{th} step, choose $(m_k,n_k)=(k,0)$. Apply the update rule in (6) to obtain

$$s^{(k)}(k,0) = \begin{cases} l, & C \le d^{(k-1)}(k,0) \text{ or } k=L \\ w, & \text{else} \end{cases} \tag{7}$$

where $d^{(k-1)}(k,0) := \big(ph^{(k-1)}(k+1,0)+qh^{(k-1)}(k,1)\big) - \big(ph^{(k-1)}(1,0)+qh^{(k-1)}(0,1)\big)$. In fact, without solving the linear system in (5), it is possible to calculate $d^{(k-1)}(k,0)$. Repeated application of (5) with $h^{(k-1)}(1,0)=0$ yields

$$d^{(k-1)}(k,0) = E[G^{(k,0)}] - \lambda_L^{(k-1)} E[T^{(k,0)}]$$

where $T^{(k,0)}$ is the time until return to the reference state under policy $\boldsymbol{s}^{(k-1)}$ if we start from $(k,0)$ and opt not to label; and $G^{(k,0)}$ is the accumulated cost until the return. $T^{(k,0)}$ has a truncated Geometric distribution, i.e., $T^{(k,0)} = \min\{T, L-k\}$ for a geometrically distributed T with parameter p. Observe that given $T^{(k,0)} = t^{(k,0)}$, the accumulated cost $G^{(k,0)}$ will be equal to $\sum_{\tau=k+1}^{t^{(k,0)}+k} \tau + \mathbb{1}\{t^{(k,0)} < L-k\}C$. Therefore, the expectations above can be calculated straightforwardly and we obtain

$$d^{(k-1)}(k,0) = \frac{k}{p} + \frac{1}{p^2} - \frac{\epsilon_k}{p^2} - \frac{L\epsilon_k}{p} + (1-\epsilon_{k+1})C - \frac{\lambda_L^{(k-1)}}{p}(1-\epsilon_k)$$

where $\epsilon_k := q^{L-k}$. The update rule (7) is then equivalent to

$$\lambda_L^{(k-1)}(1-\epsilon_k) \underset{w}{\overset{l}{\lessgtr}} k + \frac{1}{p} - \frac{\epsilon_k}{p} - L\epsilon_k - Cp\epsilon_{k+1}.$$

Continue the procedure until $s^{(K_L)}(K_L,0) \ne s^{(K_L+1)}(K_L+1,0)$ for the first time for some $K_L \ge 0$. We are interested in large L as we want to make the 'phantom' arrivals as rare as possible. Thus, the limiting behavior of K_L is of interest.

Lemma 3. *$\lim_{L\to\infty} K_L =: K$ exists and is finite. Furthermore,*

$$K = \left\lceil \frac{-1+\sqrt{1+8C+\frac{4q}{p^2}}}{2} - \frac{1}{p} \right\rceil. \tag{8}$$

Proof: See Appendix B ■

Since K_L's are integer, Lemma 3 implies the existence of an L_1 such that for all $L > L_1$, $K_L = K$. From now on, assume $L > L_1$ so that the termination time of our procedure is K. The next step is to show that for $m > K$ and $n = 0$, the policy will not updated if L is large enough.

Lemma 4. *There exists an L_2 such that for all $L > L_2$, one obtains a worse policy by the modification $\tilde{s}^{(K)}(m,n) = w$ for any (m,n) such that $m > K$ and $n = 0$.*

Proof: See Appendix C ■

At this point, we have covered the states given in the region (a) of Figure 1. It only remains to find the optimal actions for states (m,n) with $n > 0$, i.e., regions (b) and (c) in Figure 1. Our procedure did not modify the actions of these states. Hence, $s^{(K)}(m,n) = w$ for $n > 0$. Now we show that the actions of region (b) should remain unchanged.

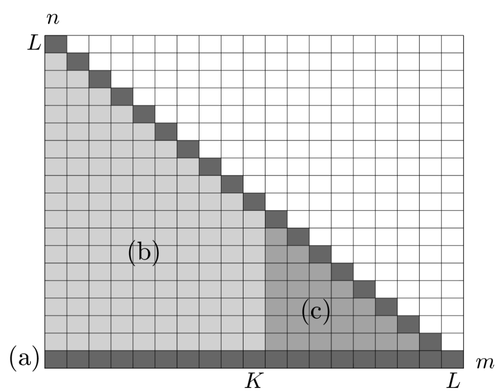

Fig. 1. States as (m,n) pairs covered until different steps our analysis. (a) corresponds to the darkest shaded region and includes the states $n = 0$ and $m+n = L$, which are covered with Lemma 4. (b) corresponds to the slightly shaded region including the states $n > 0$, $m \le K$, $m+n < L$, which are covered with Lemma 5. The moderately shaded region (c) corresponds to the remaining states and these states turn out to be transient according to the Markov Chain induced by $\boldsymbol{s}^{(K)}$.

Lemma 5. *There exists an L_3 such that for all $L > L_3$, one obtains a worse policy by the modification $\tilde{s}^{(K)}(m,n) = l$ for any (m,n) such that $m \le K$ and $n > 0$.*

Proof: See Appendix D ■

Lemma 5 points out an unintuitive result: If an arrival is not labeled upon arrival, it will remain unlabeled. For instance, suppose an arrival occured at $K-1$. The optimal procedure skips it and waits for the next arrival even though it occurs very late.

Together with Lemma 5 we have covered the regions (a) and (b). Now observe that for the policy $\boldsymbol{s}^{(K)}$, the states in region (c) are transient since these states cannot be reached from any other state — check the state transitions in light of Remark 2 to see that transitions from regions (a) and (b) are to (a) and (b). Therefore, no matter what action is taken at a transient state in region (c), the average cost remains the same. This completes the proof that an optimal strategy is indeed $\boldsymbol{s}^{(K)}$ as any of its modification results in a higher cost.

Let us summarize what we have shown so far: There exists an $L' = \max\{L_1, L_2, L_3\}$ such that for all $L > L'$, the optimal strategy is

$$s^{(K)}(m,n) = \begin{cases} l, & m > K,\ n = 0 \\ w, & \text{else} \end{cases}$$

to which we shall refer as 'wait K, label next' strategy. One can easily calculate λ_L^* for $L > L'$ as

$$\lambda_L^* = \frac{\frac{K(K+1)}{2} + \frac{K}{p} - (\frac{L}{p} + \frac{1}{p^2})\epsilon_K + \frac{1}{p^2} + C(1-\epsilon_{K+1})}{K + \frac{1}{p}(1-\epsilon_K)}$$

and thus (recall $\epsilon_k = q^{L-k}$)

$$\lambda_\infty^* = \frac{K^2 + (2/p+1)K + 2/p^2 + 2C}{2(K+1/p)}.$$

Since $\lambda_L^* \le \lambda^*$ for all L, $\lambda_\infty^* \le \lambda^*$. Moreover, λ_∞^* can be achieved via 'wait K, label next' strategy in the untruncated problem. Therefore $\lambda_\infty^* \ge \lambda^*$, and thus $\lambda_\infty^* = \lambda^*$.

The 'wait K, label next' strategy satisfies the square integrability condition in Theorem 1, i.e., $\sup_t E[\Delta_t^2] < \infty$. To see this, observe that at time t, the previous labeled arrival must have arrived at most $K+Z$ time slots ago. Hence, we obtain $E[\Delta_t^2] \le E[(K+Z)^2] = (K+\frac{1}{p})^2 + \frac{q}{p^2} < \infty$ for all t. Furthermore, this strategy is stationary and constitutes a renewal process, which implies $R = \lim_t \frac{1}{t}\sum_{\tau=1}^{t} E[C\mathbb{1}\{S_\tau \neq ?\}]$ and $\Delta = \lim_t \frac{1}{t}\sum_{\tau=1}^{t} E[\Delta(\tau)]$, thus $CR + \Delta = \lambda^*$. This finally proves

Theorem 2. *The boundary of the feasible region of (R,Δ) pairs is characterized as the interpolation of the set of pairs $\{(R_k,\Delta_k)\}_{k\in\mathbb{N}}$ where*

$$\Delta_k = \frac{k^2 + (2/p+1)k + 2/p^2}{2(k+1/p)}$$
$$R_k = \frac{1}{k+1/p}.$$

As a final remark, we point out that the pairs $\{(R_k,\Delta_k)\}_{k\in\mathbb{N}}$ are achievable with 'wait k, label next' strategies, and the remaining pairs on the boundary curve are achievable with time sharing between at most two of such strategies, e.g., for $0 < \rho < 1$, do 'wait k, label next' ρ fraction of the time and do 'wait $k+1$, label next' $1-\rho$ fraction of the time.

IV. Extension To Poisson Processes

In this section, we extend the results obtained for the discrete-time problem to a continuous-time problem: The arrivals are modeled as a Poisson process. Let $\mathcal{N}(t)$ be the counting process associated with a stationary Poisson process of intensity ν, with its natural filtration $\{\mathcal{F}_t\}_{t\ge 0}$. Similar to the discrete case, the causal labelings for the continuous model are defined as the collection of functions $\{S_t\}_{t\in\mathbb{R}^+}$, such that every S_t is $\mathcal{F}_t$-measurable and $S_t : \mathcal{P}^{[0,t]} \to [0,t] \cup \{?\}$

where $\mathcal{P}^{[0,t]}$ denotes the space of the sample paths of $\mathcal{N}(t)$ on the interval $[0,t]$. Observe that the process that tracks the number of labelings until t constitutes another counting process. Denote this process by $\mathcal{R}(t)$. Note that $\mathcal{R}(t) \leq \mathcal{N}(t)$.

The rate and the average age are defined analogously to the discrete-time case — see (1) and (2) — namely

$$\mathcal{R} := E\left[\limsup_{t\to\infty} \frac{1}{t}\mathcal{R}(t)\right] \text{ and}$$

$$\delta := E\left[\limsup_{t\to\infty} \frac{1}{t}\int_{\tau=0}^{t} \Delta_\tau d\tau\right].$$

Recall that the proofs of Lemma 1 and Lemma 2 involve path-wise coupling arguments and can be extended to continuous-time straightforwardly. Therefore Lemma 1, Lemma 2 and thus Corollary 1 hold for the continuous-time model as well. Once more, this restricts the class of labelings to the set $\mathcal{S}_F$: the procedures that only label the freshest arrival. In the following theorem, we extend the results of the discrete model and show that 'wait T, label next' type of strategies are optimal for the continuous-time problem as well; which concludes our work.

Theorem 3. *For $\{S_t\}_{t\in\mathbb{R}^+} \in \mathcal{S}_F$ such that $\sup_{t\in\mathbb{R}^+} E[\Delta_t^2] < \infty$, the boundary curve of achievable $(\mathcal{R},\delta)$ pairs lie on the curve*

$$\delta(\mathcal{R}) = \frac{1}{2\mathcal{R}} + \frac{\mathcal{R}}{2\nu^2}, \quad \mathcal{R} \leq \nu$$

which are achieved with 'wait T, label next' strategies with $T = \frac{1}{\mathcal{R}} - \frac{1}{\nu}$.

Proof: We discretize the time axis by dividing it into small intervals of length h. Then, for any strategy $\{S_t\} \in \mathcal{S}_F$, consider a modification that only makes decisions at times that are multiples of h. Suppose the continuous-time strategy $\{S_t\}$ has made k labelings at times $l_1,\ldots,l_k$ on the interval $((n-1)h, nh]$; and Δ_t drops down to values $a_1,\ldots,a_k$ respectively for each labeling. We can infer that the labeled arrivals have occured at times $t_1 = l_1 - a_1, \ldots, t_k = l_k - a_k$. Observe that only the first labeled arrival can belong to a past interval; otherwise it would not be the freshest arrival. This implies that $l_i - t_i \leq h$ for $i > 1$. Suppose the discrete-time modification labels only the k^{th} arrival at time nh and its instantaneous age $\tilde{\Delta}_n$ drops down to $\lfloor \frac{nh-(l_k-a_k)}{h} \rfloor h$. If $k > 1$, $\tilde{\Delta}_n = 0$, hence smaller than Δ_t. If $k = 1$, then $\tilde{\Delta}_n = \lfloor \frac{nh-(l_1-a_1)}{h} \rfloor h \leq \lfloor \frac{h+a_1}{h} \rfloor h \leq a_1 + h \leq \Delta_t + h$ for all $t \in ((n-1)h, nh]$. Hence, the discrete age can at most be h higher than the continuous age for all t, yielding

$$\frac{1}{t}\int_{\tau=0}^{t} \Delta_\tau d\tau \geq \frac{1}{(N+1)}\sum_{n=1}^{N}(\tilde{\Delta}_n - h)$$

where $N = \lfloor t/h \rfloor$. Then, taking limsup on both sides, we have

$$\limsup_{t\to\infty} \frac{1}{t}\int_{\tau=0}^{t} \Delta_\tau d\tau \geq \limsup_{N\to\infty} \frac{1}{N}\sum_{n=1}^{N} \tilde{\Delta}_n - h.$$

Also observe that the number of labelings made by the discrete-time modification is always smaller than $\mathcal{R}(t)$ — as it only labels the k^{th} arrival. Proceeding similarly as above, we have

$$\limsup_{t\to\infty} \frac{1}{t}\mathcal{R}(t) \geq \limsup_{N\to\infty} \frac{1}{Nh}\sum_{n=1}^{N} R_n^{(h)}$$

where $R_n^{(h)} := \mathbb{1}\{\exists t \in ((n-1)h, nh] : S_t \neq ?\}$. Using Reverse Fatou's lemma once more as in Theorem 1, we obtain

$$C\mathcal{R} + \delta \geq \limsup_{N\to\infty} \frac{1}{Nh}\sum_{n=1}^{N} E[\tilde{\Delta}_n h + CR_n^{(h)}] - h.$$

The limsup expression above can be lower bounded by the average cost of a very similar problem that we have considered in discrete-time. Recall that the discrete-time modification is able to label from an interval as long as there exists at least one arrival. This implies that the interarrival times of the discrete model are Geometrically distributed with parameter $p = 1 - e^{-\nu h}$, which is equal to the probability that at least one arrival occurs on an interval of length h. Furthermore, the discrete-time modification is able to label from the same interval more than once. Hence, comparing with our discrete-time formulation in Section II, we see that $R_n^{(h)} \geq \mathbb{1}\{S_n \neq ?\}$, where S_n is defined as in Section II. The expression above is then lower bounded by h times the optimal average cost of an MDP with

$$\limsup_{N\to\infty} \frac{1}{N}\sum_{n=1}^{N} E[\Delta_{n+1} - 1 + \frac{C}{h^2}\mathbb{1}\{S_n \neq ?\}]. \tag{9}$$

where Δ_n is defined as in our discrete-time formulation in Section II. The crucial difference is the first term, which is a time shifted version of the original problem. With similar truncation arguments as in Section III, we can analyze the modified problem and study its optimal stationary policies. However, observe that any stationary and unichain policy of the modified problem can be simulated with our original problem — a time shift of instantaneous ages does not change the long-term average. Hence, the optimal policy will exactly be the optimal policy of the original problem. Now write down the optimal average cost and do the appropriate scaling to obtain

$$\delta + C\mathcal{R} \geq \lim_{h\to 0} h\frac{K^2 + (2/p+1)K + 2/p^2 + 2C/h^2}{2(K+1/p)} - 2h$$

$$= \lim_{h\to 0} \frac{T^2 + 2T/\nu + 2/\nu^2 + Th + 2C}{2(T+1/\nu)}$$

where $T := hK(h)$. We have written $K(h)$ to emphasize its dependence in h. Substituting C/h^2 and $p = 1 - e^{-\nu h}$ in (8), we see that T tends to $\sqrt{2C + \frac{1}{\nu^2}} - \frac{1}{\nu}$ as $h \to 0$, which is a constant. Hence,

$$\delta + C\mathcal{R} \geq \frac{T^2 + 2T/\nu + 2/\nu^2 + 2C}{2(T+1/\nu)}.$$

Similar to the end of the proof of Theorem 2, we argue that 'wait T and label' strategy is stationary, satisfies the square integrability condition and attains the average cost above; concluding that the above is an equality. ■

References

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

[2] R. D. Yates, Y. Sun, D. R. Brown III, S. K. Kaul, E. Modiano, and S. Ulukus, "Age of information: An introduction and survey," 2020.

[3] A. Kosta, N. Pappas, and V. Angelakis, *Age of Information: A New Concept, Metric, and Tool*, 2017.

[4] Y. Sun, Y. Polyanskiy, and E. Uysal, "Sampling of the wiener process for remote estimation over a channel with random delay," *IEEE Transactions on Information Theory*, vol. 66, no. 2, pp. 1118–1135, 2020.

[5] T. Z. Ornee and Y. Sun, "Sampling and remote estimation for the ornstein-uhlenbeck process through queues: Age of information and beyond," *IEEE/ACM Transactions on Networking*, vol. 29, no. 5, pp. 1962–1975, 2021.

[6] Y. İnan, R. Inovan, and E. Telatar, "Optimal policies for age and distortion in a discrete-time model," in *2021 IEEE Information Theory Workshop (ITW)*, 2021, pp. 1–6.

[7] M. Costa, M. Codreanu, and A. Ephremides, "On the age of information in status update systems with packet management," *IEEE Transactions on Information Theory*, vol. 62, no. 4, pp. 1897–1910, 2016.

[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 Transactions on Information Theory*, vol. 63, no. 11, pp. 7492–7508, 2017.

[9] B. Wang, S. Feng, and J. Yang, "To skip or to switch? minimizing age of information under link capacity constraint," in *2018 IEEE 19th International Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*, 2018, pp. 1–5.

[10] E. T. Ceran, D. Gündüz, and A. György, "Average age of information with hybrid arq under a resource constraint," *IEEE Transactions on Wireless Communications*, vol. 18, no. 3, pp. 1900–1913, 2019.

[11] B. T. Bacinoglu, Y. Sun, E. Uysal, and V. Mutlu, "Optimal status updating with a finite-battery energy harvesting source," *Journal of Communications and Networks*, vol. 21, no. 3, pp. 280–294, 2019.

[12] B. T. Bacinoglu, E. T. Ceran, and E. Uysal-Biyikoglu, "Age of information under energy replenishment constraints," in *2015 Information Theory and Applications Workshop (ITA)*, 2015, pp. 25–31.

[13] R. D. Yates, "Lazy is timely: Status updates by an energy harvesting source," in *2015 IEEE International Symposium on Information Theory (ISIT)*, 2015, pp. 3008–3012.

[14] A. Arafa, J. Yang, S. Ulukus, and H. V. Poor, "Age-minimal transmission for energy harvesting sensors with finite batteries: Online policies," *IEEE Transactions on Information Theory*, vol. 66, no. 1, pp. 534–556, 2020.

[15] S. Farazi, A. G. Klein, and D. R. Brown, "Age of information in energy harvesting status update systems: When to preempt in service?" in *2018 IEEE International Symposium on Information Theory (ISIT)*, 2018, pp. 2436–2440.

[16] O. Ozel, "Timely status updating through intermittent sensing and transmission," in *2020 IEEE International Symposium on Information Theory (ISIT)*, 2020, pp. 1788–1793.

[17] S. Resnick, *A Probability Path*, ser. Modern Birkhäuser Classics. Birkhäuser Boston, 2003. [Online]. Available: https://books.google.ch/books?id=T9-PMDSVDNsC

[18] D. P. Bertsekas, *Dynamic Programming and Optimal Control, Vol. II*, 3rd ed. Athena Scientific, 2007.

APPENDIX

A. Extension of Example 2 to Rational r:

Let $r = \frac{m}{n}$ with $m \leq n$. Write $n = m\alpha + \beta$, where $\alpha \geq 1$ is the quotient and β is the remainder. Note that $\beta < m$. Write r as

$$r = \frac{m}{(\alpha+1)\beta + \alpha(m-\beta)}.$$

Consider the following strategy: (i) Label every $\alpha + 1$th arrival, do this β times. Then (ii) label every αth arrival, do this $m - \beta$ times. Then repeat (i) and (ii) consecutively. By renewal theory, the average rate will be pr. Also note that the interarrival times are sum of β Geometric random variables with parameter $\frac{p}{\alpha+1}$ plus sum of $m - \beta$ Geometric random variables with parameter $\frac{p}{\alpha}$. With some algebra, one can obtain

$$\Delta = \frac{r(\alpha+1)}{2R}\left(1 + \frac{\beta r}{m}\right).$$

To check if this expression is smaller than $1/R$, observe

$$\begin{aligned}
\frac{r(\alpha+1)}{2R}\left(1 + \frac{\beta r}{m}\right) &= \frac{1}{2R}\frac{m}{n}(\alpha+1)(1+\beta/n) \\
&= \frac{1}{2R}\frac{n-\beta+m}{n}(1+\beta/n) \\
&= \frac{1}{2R}(1 - \frac{\beta}{n} + \frac{m}{n})(1 + \frac{\beta}{n}) \\
&= \frac{1}{2R}(1 - \frac{\beta^2}{n^2} + \frac{m}{n} + r\frac{\beta}{n}) \\
&\leq \frac{1}{2R}(1 + \frac{m + r\beta}{n}) \leq \frac{1}{2R}(1+1) = \frac{1}{R}
\end{aligned}$$

as $n = m\alpha + \beta$ and $\alpha \geq 1$ and $r \leq 1$.

B. Proof of Lemma 3

K_L is defined as the minimum non-negative integer satisfying $\lambda_L(k) < g_L(k)$, where

$$\lambda_L(k) = \frac{\frac{k(k+1)}{2} + \frac{k}{p} - (\frac{L}{p} + \frac{1}{p^2})\epsilon_k + \frac{1}{p^2} + C(1 - \epsilon_{k+1})}{k + \frac{1}{p}(1-\epsilon_k)}$$

is the average cost pertaining to the 'wait k, label next' strategy and

$$g_L(k) := \frac{k + 1 + \frac{1}{p} - \frac{\epsilon_{k+1}}{p} - L\epsilon_{k+1} - Cp\epsilon_{k+2}}{(1 - \epsilon_{k+1})}.$$

Define $g_L(L-1) = \infty$ for convenience. This ensures that $K_L \leq L - 1$. Observe that for any finite $k \geq 0$ and small ϵ, we can find large enough L such that a sufficient condition for the inequality $\lambda_L(k) < g_L(k)$ can be obtained as

$$k^2 + (\tfrac{2}{p} + 1)k + \tfrac{2}{p^2} + 2C < 2(k + \tfrac{1}{p})(k + 1 + \tfrac{1}{p}) - \epsilon$$

which is equivalent to

$$\epsilon < (k + \tfrac{1}{p})^2 + (k + \tfrac{1}{p}) - 2C - \tfrac{1}{p^2} + \tfrac{1}{p} =: f_L(k).$$

Since K_L is the smallest integer satisfying the above, we must have $\lambda_L(k-1) \geq g_L(k-1)$. Again, for large enough L, an equivalent sufficient condition will be

$$f_L(k-1) \leq -\epsilon.$$

$f(k)$ is a quadratic function of k and it always has two distinct real roots. The smaller root is always negative, hence take the larger root

$$\tilde{k} = \frac{-1 + \sqrt{1 + 8C + \frac{4q}{p^2}}}{2} - \frac{1}{p}$$

and note that for small enough ϵ both sufficient conditions are satisfied at $k = \lceil \tilde{k} \rceil$. This shows that $K_L \leq \lceil \tilde{k} \rceil$ for large enough L, which implies $\limsup_L K_L \leq \lceil \tilde{k} \rceil$. Furthermore, for $0 < C \leq \frac{1}{p}$, $-1 < \tilde{k} \leq 0$ and thus $\lceil \tilde{k} \rceil = 0$, so $\lim_L K_L = 0$ for such C. For $C > \frac{1}{p}$, with a similar argument, obtain necessary conditions by choosing large enough L and small enough ϵ; which are

$$f(k) > -\epsilon \text{ and } f(k-1) \leq \epsilon.$$

Observe that $f(\lceil \tilde{k} \rceil - 1)$ is negative and bounded away from zero — as $\tilde{k} > 0$ and the smaller root is always negative. Hence the necessary conditions will not be satisfied for $\lceil \tilde{k} \rceil - 1$ and $K_L > \lceil \tilde{k} \rceil - 1$ eventually. This concludes that $\liminf_L K_L \geq \lceil \tilde{k} \rceil$ and therefore $\lim_L K_L = \lceil \tilde{k} \rceil =: K$. We remark that $K \leq \lfloor Cp \rfloor$. This fact will be important when proving Lemma 4 and Lemma 5.

C. Proof of Lemma 4

This is equivalent to showing that the update condition

$$\lambda_L^{(K)}(1-\epsilon_m) \geq m + \frac{1}{p} - \frac{\epsilon_m}{p} - L\epsilon_m - \tfrac{Cp}{q}\epsilon_m.$$

is violated for all $K < m < L$. It is sufficient to check if the minimum of the function

$$f(t) = L + \tfrac{1}{p} + \tfrac{\log t}{\log \frac{1}{q}} - \lambda_L^{(K)}(1-t) - (L + \tfrac{1}{p} + \tfrac{Cp}{q})t$$

on $t \in [q^{L-K-1}, q]$ is greater than zero. f is concave, therefore the minimum occurs at boundaries. Observe $f(q) = Lp - Cp - p\lambda_L^{(K)} \geq Lp - Cp - p(Cp + \frac{1}{p})$ therefore greater than zero for $L \geq L_2 := C + (Cp + \frac{1}{p^2})$. The other boundary corresponds to $m = K+1$, and from the definition of K as being the termination time of the procedure — check the definition in the beginning of Appendix B — we must have $f(q^{L-K-1}) \geq 0$. Thus, the condition above is violated for all $m > K$.

D. Proof of Lemma 5

The condition in (6) implies that if we alter $s^{(K)}(m,n)$ to l, the policy does not improve if

$$C + ph(n+1,0) + qh(0,n+1) \geq ph(m+n+1,0) + qh(m,n+1).$$

Some calculation reveals that this condition is equivalent to (10) — at the bottom of this page.

Now, consider the three cases (for which $m \leq K$ according to the main statement of the Lemma):

(i) $n > K$

Then (10) is equivalent to

$$\frac{Cp - m}{q^{L-m-n} - q^{L-n}} + L + \frac{1}{p} + Cp \geq \lambda_L^{(K)}.$$

Recall that $Cp \geq K$ — see the end of Appendix B — hence the first term on the left-hand side is positive. We also know that $\lambda_L^{(K)} \leq Cp + 1/p$ as the latter is the average cost of the untruncated problem with 'label every point upon its arrival' policy. Hence the inequality always holds.

(ii) $n \leq K$, $m + n > K$

Observe that the (m,n) pairs lie on a bounded set that does not grow with L. Hence, for large enough L, the condition will be equivalent to

$$C + (K-n)(K+n+1)/2 - (m+n-K)/p \geq \lambda(K-n)$$

the left-hand side is minimized at $m = K$, which yields

$$C + (K-n)(K+n+1)/2 - n/p \geq \lambda(K-n).$$

Since the left-hand side is a concave quadratic function of n, it is minimized at either $n = 0$, or $n = K$. For $n = 0$, the state corresponds to $(K, 0)$ and we know the inequality holds from definition of K. For $n = K$, we have

$$C - K/p \geq 0$$

which we know is true.

(iii) $n \leq K$, $m + n \leq K$. Again, the possible pairs lie in a bounded set that does not grow with L. For large L, rewrite the condition as

$$C \geq \lambda m - m(m + 2n + 1)/2.$$

The right-hand side is maximized at $n = 0$, and similar to (ii) the result is immediate from the definition of K.

As K, m, n are integers, there also exists an L_3 such that the conditions in (ii), (iii) are exact for $L > L_3$.

$$\begin{aligned} C + \sum_{\tau=n+1}^{K} \tau + \sum_{\tau=K\vee n}^{L-1} (\tau+1) q^{\tau - K\vee n} + C(1 - q^{L-K\vee n}) - \lambda\Bigg(\sum_{\tau=n+1}^{K} 1 + \sum_{\tau=K\vee n}^{L-1} q^{\tau-K\vee n} \Bigg) \\ \geq \sum_{\tau=m+n+1}^{K} \tau + \sum_{\tau=K\vee(m+n)}^{L-1} (\tau+1) q^{\tau-K\vee(m+n)} + C(1 - q^{L-K\vee(m+n)}) - \lambda\Bigg(\sum_{\tau=m+n+1}^{K} 1 + \sum_{\tau=K\vee(m+n)}^{L-1} q^{\tau-K\vee(m+n)} \Bigg) \end{aligned} \tag{10}$$