AoI-Penalty Minimization for Networked Control Systems with Packet Loss

Markus Klügel[*§], Mohammad H. Mamduhi[†‡], Sandra Hirche[†], Wolfgang Kellerer[§]

[*]Airbus, Ottobrunn/Munich, Germany
[†]Chair of Information-Oriented Control, Technical University of Munich
[‡]Division of Decision and Control Systems, KTH Royal Institute of Technology, Stockholm
[§]Chair of Communication Networks, Technical University of Munich
[*]markus.kluegel@airbus.de, [†]mh.mamduhi@tum.de, [‡]hirche@tum.de, [§]wolfgang.kellerer@tum.de

Abstract—In this work the impact of age of information (AoI) is studied from the perspective of networked control systems (NCS), i.e., control loops that are closed over networks. We formulate the estimation problem of a linear time invariant (LTI) system and show that related performance metrics can be optimized by minimizing age-penalty functions. From the variety of possible penalties that make sense from an NCS point of view, we derive a general age-penalty minimization problem. We characterize properties of penalty functions that are trivial or non-trivial to solve and show that for non-trivial age-penalties, the optimal transmission policy over a single link with packet loss is AoI-threshold based. Then, we propose an algorithm to find the optimal threshold. Simulation results verify that threshold policies with optimal threshold can serve to optimally solve a variety of NCS related estimation problems.

I. Introduction

The concept of age of information (AoI) has been recently introduced to capture application layer delay experience in a precise fashion [1], [2]. In contrast to per-packet delay, information age refers to the elapsed time since the generation of the latest received information at the receiver side and can incorporate the effect of not transmitting a packet or packet loss on end-to-end delay. As the metric is rather new and less understood than, e.g., delay, it has raised significant interest in the research community.

Since the seminal works on AoI [1], [2], research interest has been on the different related aspects, for example the impact of queuing on the average AoI [1]–[6], the impact of multi-hop transmissions [7], [8] and packet drops [9]. All together have led to a good general understanding of the metric and related options. More recently, the focus is shifting towards scheduling problems for AoI minimization [8], [10]–[12]. The most often seen approach here is to target a weighted-sum average AoI.

The use of AoI metric is typically motivated from the context of internet of things (IoT) and cyber physical systems (CPS), where timeliness of information at the receiver side is crucial. The timeliness requirement is often used to motivate that AoI, or a weighted average thereof, needs to be decreased. However, clean derivations that weighted AoI is an expressive metric for IoT or CPS are mostly missing. Alternatives to weighted average AoI do exist: In [13] and its journal version, [14], Sun et al. proposed the use of a general age-penalty

function, which captures the dissatisfaction of the communication sink with a certain AoI. For random networking delays with given distributions, the authors show that the optimal transmission times result in a threshold policy. Kosta et al. [15] proposed the use of a cost of update delay (CoUD) metric and derived closed form expressions for the expected CoUD through M/M/1 queues.

While the above works capture a general age penalty, which is assumed non-negative and non-decreasing, there is no comprehensive intuition on the exact type of age-function that make sense for IoT and CPS specific problem classes. This gap is partly closed by the works [16], [17] of Sun et al. and Sun and Cyr from different perspectives. In [16], the authors investigated sampling strategies for remote estimation of a Wiener process with random transmission delays. While they showed the estimation error is indeed a function of AoI, they also motivated that only in special cases, minimization of the estimation error is transferable to directly minimizing the average AoI. In [17], the authors investigated an alternative definition of information freshness, which is the mutual information among transmitter and receiver. They then showed that the mutual information can be expressed as function of AoI and showed the optimal sampling strategy to be threshold based. Interestingly, maximization of mutual information leads to minimization a non-positive and non-decreasing penalty function. This counters the intuitive assumption that AoI penalties should be non-negative.

A. Contribution of This Work

Concluding on the existing literature, we observe that while AoI has raised some research interests, the clear motivational link from the IoT and CPS scenario to the AoI metric is mostly missing, with exceptions being [16], [17]. We aim at this gap by re-motivating AoI from a networked control systems (NCS) perspective.

We establish an estimation problem from the perspective of control theory and show that related performance metrics can be expressed as nonlinear penalty functions of AoI. Due to the variety of possible penalty functions, we ask the question which impact the penalty function itself has on AoI optimization. We divide the resulting class of penalty functions into trivial functions, which we find to lead to the

solution of never sampling or never transmitting packets, and non-trivial functions. Then, we investigate optimal scheduling policies for non-trivial penalty minimization over a single link with packet loss. We find the optimal policy to be threshold based and establish a binary search algorithm for the optimal threshold. Finally, we show with simulations that by using appropriate penalty functions, different estimation problems can be optimally solved.

We note that compared with the existing scheduling works, our problem targets scheduling for AoI penalty instead of weighted average AoI, whereas compared to the existing works with AoI penalty, it targets scheduling instead of packet generation.

II. Estimation Problems and AoI

In this section, we first introduce the estimation problem of a single non-scalar stochastic linear time invariant (LTI) system and then express the estimation objective as a function of AoI. The targeted system is depicted in Figure 1. Consider a plant $\mathcal{P}$ that could refer to, e.g., an environment or process being monitored by sensors or a robot whose state needs to be estimated. $\mathcal{P}$'s state evolves over time according to the following stochastic discrete time dynamics

$$\mathcal{P} : \boldsymbol{x}_s[t+1] = \boldsymbol{A}\boldsymbol{x}_s[t] + \boldsymbol{w}_s[t], \qquad (1)$$

where $\boldsymbol{x}_s[t] \in \mathbb{R}^n$ represents the system state at the source at slotted time $t \in \mathbb{N}_+$, $\boldsymbol{w}_s[t] \in \mathbb{R}^n$ denotes an exogenous stochastic Gaussian disturbance with zero mean and covariance Σ acting on the state dynamics of the source, and $\boldsymbol{A} \in \mathbb{R}^{n \times n}$ is the system matrix. We assume the system matrix $\boldsymbol{A}$ can be selected freely, i.e. the system can be stable or unstable in open-loop. The system is equipped with a unit installed at the source including a sensor, to sample the state of the plant $\mathcal{P}$, and a transmission decision maker that decides at each time-step to either send the sample for transmission or not. The decision variable of the transmission decision maker at a time-step t, denoted by $\delta_s[t]$, is computed based on a transmission policy $\mu : \mathbb{R}^n \mapsto \{0, 1\}$ and can be stated as follows:

$$\delta_s[t] = \mu(\boldsymbol{x}_s[t]) = \begin{cases} 1, & \boldsymbol{x}_s[t] \text{ is sent for transmission}, \\ 0, & \text{otherwise.} \end{cases} \qquad (2)$$

The samples that are sent for transmission go through a time-slotted communication network. Each transmission has a cost for the system, denoted by D, and the success probability of receiving a sample by the destination is $p_s \in [0, 1]$. That is, the communication channel is modeled by an erasure channel with erasure probability $1 - p_s$. Therefore, we introduce the random variable $\gamma[t] \in \{0, 1\}$ as the delivery indicator such that $\gamma[t] = 1$ with probability p_s. That is, a data packet including $\boldsymbol{x}_s[t]$ will be received by the destination only if $\delta_s[t] = 1$ and $\gamma[t] = 1$.

At the destination, an estimator receives the data samples and estimates the system state. Before deriving the state estimation and the resulting estimation error, let us now introduce the general definition of AoI. Consider a single communication process where $\mathcal{K} = \{t_1, ..., t_i, ...\}$, $t_{i+1} \geq t_i$,
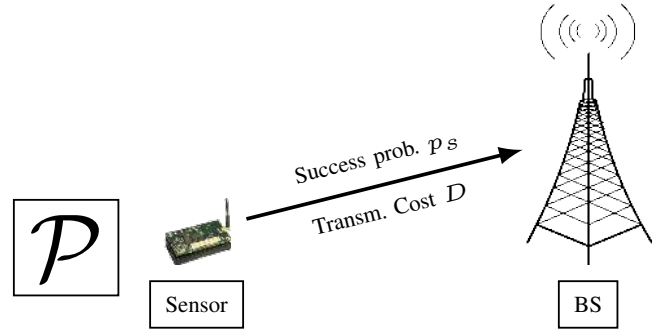

Fig. 1. A sensor observes the state $\boldsymbol{x}_s[t]$ of the plant $\mathcal{P}$, which can refer to, e.g., a robot or an environment process. The state is transmitted over a single link with success probability p_s and transmission cost D. At the BS, an estimate $\hat{\boldsymbol{x}}_s[t]$ of the plants' state is generated from the received samples.

represents a sequence of time instances when data packets including the state samples are generated at the source and $\mathcal{T} = \{t'_1, ..., t'_i, ...\}$, $t'_{i+1} \geq t'_i$, denoting a sequence of instants when those data packets are received at the destination. The instances t_i can be treated as arrival process to the network and t'_i as departure, where both can follow an arbitrary process satisfying $t'_i \geq t_i$, $\forall i$. We denote by $S_i \triangleq t'_i - t_i$ the total service time of the network, i.e., the time between packet generation and reception, including all queuing and transmission delays, and by $Y_i = t_i - t_{i-1}$ the inter arrival time. Define $\mathcal{K}(T) = \{t_i \in \mathcal{K} : t_i \leq T\}$ to be the set of sampling times within a considered interval $[0, T]$, and note that $|\mathcal{K}(T)|$ is the number of transmissions that fall into the observed interval. Moreover, $s[t] = \sup\{t_i \in \mathcal{K}(T) : t_i + S_i \leq t\}$ denotes the time stamp of the latest received data packet before time-step t. Using these definitions, the AoI at time t is defined as $\Delta[t] \triangleq t - s[t]$ [1], [2]. The most attractive property of $\Delta[t]$ is that as long as no packet is received, $s[t]$ remains constant and hence $\Delta[t]$ increases linearly. However, when a packet i is received at $t := t'_i$, $s[t]$ is instantaneously reset to t_i and $\Delta[t] = t'_i - t_i = S_i$ holds. Both effects together create a sawtooth curve evolution of $\Delta[t]$, that is depicted in most of the available literatures, e.g. [1], [2].

To fully characterize the estimation problem and the estimation objective, we need to determine the information available at the source, i.e., the transmission unit, and at the destination, i.e., the estimator. Denoted by $\mathcal{I}_s[t]$ and $\mathcal{I}_d[t]$, the information available at the source and destination are respectively as follows:

$$\mathcal{I}_s[t] = \{\boldsymbol{x}_s[t], \delta_s[0], \dots, \delta_s[t-1], \gamma[0], \dots, \gamma[t-1], \Delta[t-1]\}$$
$$\mathcal{I}_d[t] = \{\delta_s[t], \gamma[t], \Delta[t], \boldsymbol{x}_s[t - \Delta[t]]\}.$$

The estimator can then compute the state estimation given by $\hat{\boldsymbol{x}}_s[t] = \mathbb{E}[\boldsymbol{x}_s[t] | \mathcal{I}_d[t]]$.

Assuming that the latest received state information at the destination is $\boldsymbol{x}_s[t - \Delta[t]]$, we have from (1)

$$\boldsymbol{x}_s[t] = \boldsymbol{A}^{\Delta[t]} \boldsymbol{x}_s[t - \Delta[t]] + \sum_{r=1}^{\Delta[t]} \boldsymbol{A}^{r-1} \boldsymbol{w}_s[t - r], \qquad (3)$$

$$\hat{\boldsymbol{x}}_s[t] = \mathbb{E}\left[\boldsymbol{x}_s[t]|\mathcal{I}_d[t]\right] = \boldsymbol{A}^{\Delta[t]}\boldsymbol{x}_s[t - \Delta[t]]. \qquad (4)$$

The estimation error can then be computed straightforwardly as follows:

$$\boldsymbol{e}_s[t] = \sum_{r=1}^{\Delta[t]} \boldsymbol{A}^{r-1}\boldsymbol{w}_s[t - r]. \qquad (5)$$

It can be seen from the estimation error dynamics that its dependency on the AoI is tightly coupled with the plant's stability properties, i.e., if $\mathcal{P}$ is stable (all the eigenvalues of the $\boldsymbol{A}$ matrix lie within the unit circle) then the estimation error increases slower w.r.t. to the increase of the AoI and has an upper bounded norm in expectation, i.e., becomes insensitive to increasing AoI asymptoticly. This is straightforward to show noting that AoI appears as the exponent of the $\boldsymbol{A}$ matrix. Otherwise, if $\mathcal{P}$ is an unstable plant, i.e. if the spectral radius of its corresponding $\boldsymbol{A}$ matrix is larger than one, then increasing AoI results in an ever increasing estimation error. This equivalently means that as time goes forward, the error growth rate w.r.t. AoI increases for unstable plants, and decreases for stable ones where it asymptoticly stops to grow for $t \to \infty$. The conclusion is that the estimation error in (5) is an increasing convex function w.r.t. to AoI if the plant is unstable and an increasing concave function of the AoI if the plant is stable.

III. A PENALTY PERSPECTIVE ON AoI

From an estimation problem perspective, it makes sense to minimize the mean-square norm of $\boldsymbol{e}_s[t]$, which is given by

$$\mathbb{E}\{\|\boldsymbol{e}_s[t]\|_2^2\} = \mathbb{E}\{\|\sum_{r=1}^{\Delta[t]} \boldsymbol{A}^{r-1}\boldsymbol{w}_s[t - r]\|_2^2\} = f(\Delta[t]).$$

Observe that the middle term is an increasing function $f(\Delta[t])$ of AoI and can be treated as penalty to be minimized. However, it is not the only possible penalty function to be used. Another intuitive goal would be to limit the probability that the error norm surpasses a threshold θ. Because the error is a sum of multi-variate Gaussian random variables with zero mean, this probability is increasing with AoI and leads to an AoI penalty as $\Pr\{\|\boldsymbol{e}_s[t]\|_2 > \theta\} = f(\Delta[t])$. We can further constrain the excess probability, demanding $\Pr\{\|\boldsymbol{e}_s[t]\|_2 > \theta\} \leq \eta$. Because $\Pr\{\|\boldsymbol{e}_s[t]\|_2 > \theta\}$ is increasing w.r.t. AoI, the expected excess rate is a penalty function of the form $\mathbb{E}\{\mathbb{1}\{\Pr\{\|\boldsymbol{e}_s[t]\|_2 > \theta\} > \eta\}\} = \mathbb{E}\{\mathbb{1}\{\Delta[t] > K\}\}$ for an appropriately chosen K.

The variety of options for AoI penalties leads us to the task of investigating the impact of such penalties on AoI optimization in general, which is studied in the following.

A. Expression for Expected AoI Penalty

Consider an age-penalty function $f(\Delta[t]) = f(t - s[t])$ that captures the level of dissatisfaction of the sink with respect to the current AoI. We assume that $f(\Delta[t])$ is a mapping $f : \mathbb{R} \mapsto \mathbb{R}$ with $f(t) = 0 \ \forall t < 0$, but apart from that we make no further restrictions on f at this point. Because $\Delta[t]$ has a sawtooth form, this greatly impacts $f(\Delta[t])$, potentially

leading to continuity breaks at all reception instances t_i'. Examples for penalty functions that we can imagine are shown in Figure 2, together with the time evolution of $f(\Delta[t])$ for a given sequence of t_i and t_i'.

We want to derive an expression for the infinite time-horizon average age-penalty. In analogy to [1], [2], this can be done from the geometric perspective shown in Figure 2 to be the time-averaged sum of values of $f(t)$. By using the time intervals denoted in the figure and considering a time interval between $[0, T]$, the average penalty can be obtained as

$$\overline{F}(T) = \frac{1}{T}\sum_{t=0}^{T-1} f(t) = \frac{Q_0 + \sum_{i=1}^{|\mathcal{K}(T)|} Q_i + Q'}{T}, \qquad (6)$$

where, Q_0 denotes the sum over the first truncated interval, and Q' is the sum over the last remaining interval. For $i \geq 1$, the aggregate sum Q_i can be written as:

$$
\begin{aligned}
Q_i &= \sum_{t=t_{i-1}'}^{t_i'-1} f(t - t_{i-1}) \\
&= \sum_{t=t_{i-1}}^{t_i'-1} f(t - t_{i-1}) - \sum_{t=t_{i-1}}^{t_{i-1}'-1} f(t - t_{i-1}) \\
&= \sum_{\tau=0}^{t_i'-t_{i-1}-1} f(\tau) - \sum_{0}^{t_{i-1}'-t_{i-1}-1} f(\tau) \\
&= F(t_i' - t_{i-1}) - F(t_{i-1}' - t_{i-1}), \qquad (7)
\end{aligned}
$$

where $F(t) = \sum_{t=0}^{t-1} f(t)$. Now, by using that $S_i = t_i' - t_i$ and $Y_i = t_i - t_{i-1}$, we have

$$Q_i = F(t_i' - t_i + t_i - t_{i-1}) - F(t_{i-1}' - t_{i-1}) \qquad (8)$$
$$= F(Y_i + S_i) - F(S_{i-1}). \qquad (9)$$

Finally, following the approach of [2], by letting $T \to \infty$ and using the notion of effective arrival rate $\lambda = \lim_{T\to\infty}|\mathcal{K}(T)|/T$, we get that

$$
\begin{aligned}
\overline{F} = \lim_{T\to\infty} \overline{F}(T) &= \lim_{T\to\infty} \frac{|\mathcal{K}(T)|}{T} \frac{1}{|\mathcal{K}(T)|} \sum_{i=1}^{|\mathcal{K}(T)|} Q_i \\
&= \lambda \mathbb{E}\{Q_i\} = \lambda \mathbb{E}\{F(Y + S)\} - \lambda \mathbb{E}\{F(S)\}, \qquad (10)
\end{aligned}
$$

where Y and S are random variables with probability distributions matching those of the inter-arrival times Y_i and the network service times S_i, respectively. We can now use the general result of queuing theory that $\lambda = \frac{1}{\mathbb{E}\{Y\}}$. Plugging this into (10) yields:

$$\overline{F} = \frac{\mathbb{E}\{F(Y + S)\} - \mathbb{E}\{F(S)\}}{\mathbb{E}\{Y\}}.$$

B. Trivial and Non-Trivial Age-Cost Functions

In many cases, the sampling times t_i and reception times t_i' can be influenced by implementing sampling policies for t_i or scheduling and queuing strategies for t_i'. Indirectly, this can be used to influence Y and T, leading to optimization problems either for the source or the network. In this section,
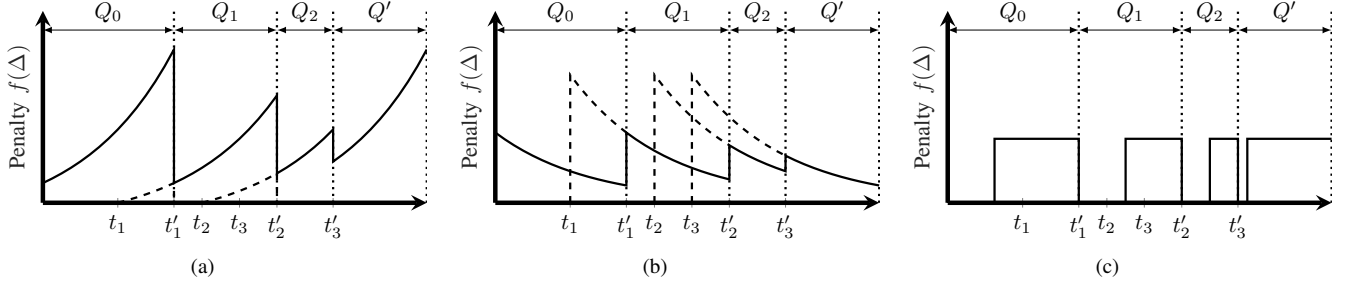
Fig. 2. Possible penalty functions: (a) $f(\Delta) = e^{\alpha\Delta}$, (b) $f(\Delta) = e^{-\alpha\Delta}$ and (c) $f(\Delta) = \mathbb{1}\{\Delta \geq K\}$. Given in the figure are examples of the packet generation times t_i, reception times t_i' and the evolution of $f(t - t_i)$ in for discrete times. The vertical, dashed lines indicate the borders of intervals used in the derivation flow.

we identify age-penalty functions that are trivial to optimize and, in reaction, conclude on properties of non-trivial age-penalty functions, which are of course the cases of interest. At this point, it is worth noting that in literature, the penalty is typically assumed non-decreasing and non-negative [13]–[15].

Let's assume for an instant that we have full control over all time instances t_i and t_i' and no limitations are made except for $t_i \leq t_i'$ $\forall i$. Then we make the following observations:

Proposition 1 (Trends for non-decreasing $f(\tau)$): If $f(\tau)$ is non-decreasing $\forall \tau$, then the optimal choice for sampling times t_i is $t_i \rightarrow t_i'$, $\forall i$ and the optimal choice for reception times is such that $t_i' \rightarrow t_{i-1}'$, $\forall i$.

Proof: The proof is given in Appendix A. ∎

Proposition 2 (Trends for non-increasing $f(\tau)$): If $f(\tau)$ is non-increasing $\forall \tau$, then the optimal choice for reception times t_i' is $t_i' \rightarrow \infty$, $\forall i$, while the optimal choice of transmission times is $t_i \rightarrow -\infty$, $\forall i$.

Proof: The proof can be argued analog to that of Proposition 1. ∎

Of course, in reality the t_i and t_i' can not be chosen arbitrarily, i.e., we do not have full control over them and sometimes even must assume them given according to an exogenous process. However, Propositions 1 and 2 give the view that if $f(\tau)$ is non-decreasing, packet sampling should be as late as possible prior to packet reception and reception should be with as large frequency as possible. On the other hand, if $f(\tau)$ is non-increasing, then packets may be infinitely delayed, which can be easily realized, e.g., by dropping the packets at the sink. We conclude that the core property that leads to non-trivial solution is a non-decreasing property of $f(\tau)$. While this seems intuitive at first glance, it slightly opposes the penalty functions found in literature, which are *non-decreasing and non-negative* [13]–[15]. Our core result however is that non-negativity of $f(\tau)$ does not significantly change the problem properties, as the complicating trends are created already by assuming monotonic increasing functions, while monotonic decreasing functions suffice to make solution trivial.

We further remark that the results of Propositions 1 and 2 can be easily adapted when restricting $t_i \in [a_i, b_i]$ and $t_i' \in [a_i', b_i']$, i.e. to time-allowed intervals, respectively. This means that penalty functions with changing monotonicity can

be divided into a sequence of intervals with fixed monotonicity, leading to the particular result that as long as the AoI is in an interval of decreasing monotonicity, there is no need for transmission.

IV. AGE-COST MINIMIZATION FOR SINGLE LINKS

We now return to the estimation problem of Section II and map it to an AoI penalty minimization problem for a non-decreasing AoI penalty $f(\Delta[t])$, which may comply with any of the penalties we have derived above. We assume that the source samples a new packet upon every transmission, which yields an age-optimal sampling according to Proposition 1. As the source has full knowledge, we have $\Delta_s[t] = 0$ $\forall t$. If it decides to transmit and the transmission succeeds, then the AoI at the sink, i.e. the estimator, reduces to $\Delta_d[t] := \Delta_s[t-1] + 1 = 1$, to account for the one-slot transmission delay. If the transmission fails or the node does not transmit, the AoI at the sink increases by one, i.e., $\Delta_d[t] = \Delta_d[t-1]+1$.

A. Problem Formulation

Given this set-up, the target is to minimize the infinite time-horizon age penalty together with communication cost, i.e.,

$$U = \lim_{T\to\infty} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}\{f(\Delta_d[t]) + \delta_s[t]D\}. \quad (11)$$

The problem is formulated as average cost per stage problem [18]. For this, define the network state to be $\Delta_d[t]$, the possible actions $\delta_s[t] \in \{0, 1\}$ and the cost per state $g(\Delta_d[t]) = f(\Delta_d[t]) + \delta_s[t]D$, respectively. The state evolves according to the equation $\Delta_d[t+1] = \psi(\Delta_d[t], \delta_s[t], \gamma[t])$, where

$$\psi(\Delta_d[t], \delta_s[t], \gamma[t]) = \begin{cases} 1, & \text{if } \delta_s[t] = 1 \text{ and } \gamma[t] = 1 \\ \Delta_d[t] + 1, & \text{otherwise.} \end{cases}$$

The goal is now to find a transmission policy, i.e., a mapping $\mu : \mathbb{N}_+ \mapsto \{0, 1\}$ that defines a law for $\delta_s[t] = \mu(\Delta_d[t])$ and that solves the problem

$$\min_{\mu \in \Pi} \lim_{T\to\infty} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}\{f(\Delta_d[t]) + \mu(\Delta_d[t])D\}, \quad (12a)$$

$$\Delta_d[t+1] = \psi(\Delta_d[t], \mu(\Delta_d[t]), \gamma[t]), \quad (12b)$$

$$\gamma[t] \sim B(p_s) \;\; \forall t \in \mathbb{N}, \;\; \Delta_d[0] = \Delta_0, \quad (12c)$$

where Π is the set of all admissible policies and $B(p_s)$ denotes the Bernoulli distribution with success probability p_s.

B. Optimal Policy via Dynamic Programming

Problem (12) can be solved with dynamic programming [18]. With this technique, each state is associated with a value, given by the value function $J(\Delta_d)$. It is shown that if $J(\Delta_d) = J^*(\Delta_d)$ is chosen correctly, then the optimal policy μ^* satisfies the Bellman-Equation [18]:

$$J^*(\Delta_d) = g(\Delta_d) + \mathbb{E}\{J^*(\psi(\Delta_d, \mu^*(\Delta_d)))\}, \quad \forall \Delta_d. \quad (13)$$

As shown in [18], the optimal value function can be derived by the iterative procedure called *value iteration*. In this, a sequence of value functions $J_k(\Delta_d)$ is created according to

$$J_{k+1}(\Delta_d) := \min_{\mu \in \Pi}\{g(\Delta_d) + \mathbb{E}\{J_k(\psi(\Delta_d, \mu(\Delta_d)))\}\} \quad, \forall \Delta_d, \quad (14)$$

which converges as $\lim_{k\to\infty} J_k = J^*$ for any J_0. The optimal policy is found as by-product from the argument minimizing (14) for $k \to \infty$. Equation (14) can be stated explicitly in different forms as:

$$J_{k+1}(\Delta_d) = \min_{\mu \in \Pi}\{f(\Delta_d) + \mathbb{E}\{J_k(\psi(\Delta_d, \mu(\Delta_d)))\}\}$$

$$\overset{(1)}{=} f(\Delta_d) + \min_{\delta_s \in \{0,1\}}\{(1 - \delta_s)J_k(\Delta_d + 1)$$
$$+ \delta_s(p_s J_k(1) + (1 - p_s)J_k(\Delta_d + 1) + D)\}$$

$$= f(\Delta_d) + \min\{J_k(\Delta_d + 1),$$
$$p_s J_k(1) + (1 - p_s)J_k(\Delta_d + 1) + D\} \quad (15)$$

$$= f(\Delta_d) + J_k(\Delta_d + 1)$$
$$+ \min_{\delta_s \in \{0,1\}}\{\delta_s[p_s(J_k(1) - J_k(\Delta_d + 1)) + D]\}. \quad (16)$$

In step (1) we use the fact that μ maps each state to a $\delta_s \in \{0, 1\}$ and for each, the expectation is explicitly written out. We now start by establishing a monotonicity property of J^*:

Lemma 1: It holds that $J^*(\Delta_d + n) \geq J^*(\Delta_d) \ \forall n, \Delta_d \geq 0$.

Proof: We prove the claim by induction. Choose $J_0(\Delta_d) = f(\Delta_d)$ and observe that $J_0(\Delta_d + n) \geq J_0(\Delta_d)$ $\forall n, \Delta_d \geq 0$ holds because $f(\Delta_d)$ was assumed non-decreasing. Now assume that, for any k, $J_k(\Delta_d + n) \geq J_k(\Delta_d)$ $\forall n, \Delta_d \geq 0$ is true. Then all terms in the minimum of (15) are non-decreasing with Δ_d, so the minimum operation itself produces a non-decreasing function w.r.t. Δ_d. As $f(\Delta_d)$ is also assumed non-decreasing, $J_{k+1}(\Delta_d)$ is the sum of non-decreasing functions and hence non-decreasing itself. By taking the limit $k \to \infty$, the property transfers to J^*. ∎

We can now narrow down the form of the optimal policy:

Proposition 3: The optimal policy μ^* is a stationary threshold policy of the form $\mu^*(\Delta_d) = \mathbb{1}\{\Delta_d \geq \overline{\Delta}^*\}$ for an optimal age-threshold $\overline{\Delta}^* \in \mathbb{N}_+$.

Proof: It is comprehensively discussed in numerous mathematical literature that the optimal policy for the average cost problem is stationary under three mild conditions [19, Theorem 3]; 1) the problem should be well-defined, i.e. not trivial, 2) the optimal policy should exist, and 3) the state-space transition between every two states should be possible in finite
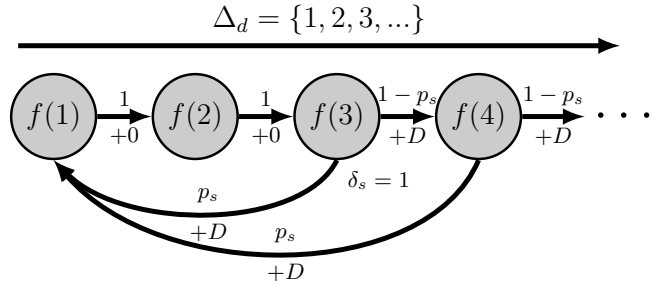

Fig. 3. State transitions and cost per state of stochastic shortest path problem with a threshold policy $\mu(\Delta_d) = \mathbb{1}\{\Delta_d \geq 3\}$. Each node corresponds to an AoI state. Reaching a node induces the cost of $f(\Delta)$ as shown in the nodes. Traversing an edge corresponds to switching a state, which happens with the probability as given above the edge, and induces a cost shown below the edge.

expected time and cost. All these conditions are satisfied in the problem (12) because first it is defined on a measurable space, and second it can be shown that the estimation error (5) has a continuous probability transition function over its uncountable state space, is irreducible, and positive recurrent, resulting in asymptotic converges for the estimation error to a bounded set under the transmission policy (2), (Detail derivations of these results are, however, out of scope of this paper.). Therefore, the optimal policy μ^* is stationary. By (16), $\mu^*(\Delta_d) = 1$ if and only if $p_s(J_k(1) - J^*(\Delta_d + 1)) + D \leq 0$. Assume that $\mu^*(\Delta_d) = 1$, then because $J^*(\Delta_d)$ is non-decreasing, it follows that $p_s(J_k(1) - J^*(\Delta_d + 1 + n)) + D \leq 0$ and hence $\mu^*(\Delta_d + n) = 1 \ \forall n \geq 0$. Now assume that $\mu^*(\Delta_d) = 0$, which induces $p_s(J_k(1) - J^*(\Delta_d + 1)) + D \geq 0$. Again because $J^*(\Delta_d)$ is non-decreasing, it must hold that $\mu^*(\Delta_d - n) = 0$ $\forall 0 \leq n \leq \Delta_d - 1$. As result, the optimal policy must be a threshold policy as claimed. ∎

C. Properties of the Optimal Threshold

For the remaining parts of this work, we denote by $U(\overline{\Delta})$ the expected time-average cost for a threshold policy with threshold $\overline{\Delta}$. It remains to find the optimal threshold $\overline{\Delta}^*$. To do so, we leverage the result of [18], according to which an average cost-per-stage problem can be cast to a stochastic shortest path problem if there is a recurrent state, i.e., a state such that at any point in time, the probability that it will be (re-)visited at some time instant in the future is one. Threshold policies on our given problem induce a specific pattern how the system moves through the state-space, which is depicted in Figure 3. It can be seen that under the assumption of a threshold policy as stated in Proposition 3, $\Delta_d = 1$ is recurrent for any $p_s > 0$, as is visited after each successful transmission

with probability one. It can be shown that [18]:

$$U(\overline{\Delta}) = \lim_{T \to \infty} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}\{f(\Delta_d[t]) + \mu(\Delta_d[t])D\} = \frac{\mathbb{E}\{\overline{C}_\mu\}}{\mathbb{E}\{\overline{N}_\mu\}},$$

where, $\mathbb{E}\{\overline{C}_\mu\} = \sum_{t=t'_i}^{t'_{i+1}-1} \mathbb{E}\{f(\Delta_d[t]) + \mu(\Delta_d[t])D \mid \mu, p_s\},$

$$\mathbb{E}\{\overline{N}_\mu\} = \mathbb{E}\{t'_{i+1} - t'_i \mid \mu, p_s\}.$$

In fact, $\mathbb{E}\{\overline{N}_\mu\}$ denotes the expected time between two consecutive visits of $\Delta_d = 1$, and $\mathbb{E}\{\overline{C}_\mu\}$ is the expected aggregated cost in this interval. Considering the state evolution shown in Figure (3), we can derive explicit expressions $\mathbb{E}\{\overline{N}_\mu\}$ and $\mathbb{E}\{\overline{C}_\mu\}$ given $\overline{\Delta}$, as follows

$$\mathbb{E}\{\overline{C}_\mu \mid \overline{\Delta}\} = \sum_{\Delta=1}^{\overline{\Delta}-1} f(\Delta) + \sum_{r=0}^{\infty}(1-p_s)^r \left[f(\overline{\Delta}+r) + D\right],$$

$$\mathbb{E}\{\overline{N}_\mu \mid \overline{\Delta}\} = \overline{\Delta} - 1 + \sum_{r=0}^{\infty}(1-p_s)^r = \overline{\Delta} - 1 + \frac{1}{p_s},$$

where we define $0^0 = 1$ for $p_s = 1$, $r = 0$, respectively. Using these statements, we make the following observations:

Proposition 4: Given $p_s > 0$, $U(\overline{\Delta}) < \infty$ is bounded $\forall \overline{\Delta} \geq 1$ if and only if

$$\sum_{r=0}^{\infty}(1-p_s)^r f(r) < \infty. \tag{17}$$

Proof: For $p_s > 0$, the denominator $\mathbb{E}\{\overline{N}_\mu\}$ is finite for any $\overline{\Delta} \geq 1$, so $U(\overline{\Delta}) < \infty$ holds if, and only if $\mathbb{E}\{\overline{C}_\mu \mid \overline{\Delta}\}$ is finite. This is the case if, and only if, $f(\Delta) < \infty$ $\forall \Delta$ and $\sum_{r=0}^{\infty}(1-p_s)^r f(\overline{\Delta}+r)$ is finite. Due to monotonicity of $f(\Delta)$, for any threshold $1 \leq \overline{\Delta} \leq \infty$ it holds that

$$\sum_{r=0}^{\infty}(1-p_s)^r f(\overline{\Delta}+r) \geq \sum_{r=0}^{\infty}(1-p_s)^r f(r) = \tag{18}$$

$$\sum_{r=0}^{\overline{\Delta}-1}(1-p_s)^r f(r) + (1-p_s)^{\overline{\Delta}} \sum_{r=0}^{\infty}(1-p_s)^r f(\overline{\Delta}+r). \tag{19}$$

Thus, if we assume that $\sum_{r=0}^{\infty}(1-p_s)^r f(r)$ is infinite, the same holds for $\sum_{r=0}^{\infty}(1-p_s)^r f(\overline{\Delta}+r)$ by (18) and hence for $\mathbb{E}\{\overline{C}_\mu \mid \overline{\Delta}\}$. On the other hand, if it is finite, the same must hold for $\sum_{r=0}^{\infty}(1-p_s)^r f(\overline{\Delta}+r)$, otherwise a contradiction would be formed from the second term of (19). If $f(\Delta) = \infty$ for any Δ, the above terms are all unbounded, respectively. ∎

Proposition 4 induces that, for some age-penalty functions $f(\Delta)$ and p_s, the expected cost can actually become unbounded. This is particularly the case for $f(\Delta) = \exp(\alpha\Delta)$ if α is too large. In contrast, the expected AoI can be seen to be bounded at any times.

Proposition 5: Consider two thresholds $\overline{\Delta}$ and $\overline{\Delta}+1$. The expected cost $U(\overline{\Delta}) \geq U(\overline{\Delta}+1)$ if, and only if,

$$\frac{1}{\overline{\Delta}}\left[\sum_{\Delta=1}^{\overline{\Delta}} f(\Delta) + \frac{D}{p_s}\right] \geq p_s \sum_{r=0}^{\infty}(1-p_s)^r f(\overline{\Delta}+1+r). \tag{20}$$

Proof: The proof is given in Appendix B. ∎

Note that the term $D/p_s = \sum_{r=0}^{\infty}(1-p_s)^r D$ is the expected transmission cost due to packet losses and does not change independent of the value of $\overline{\Delta}$. So equation (20) can be interpreted such that the left hand side contains the fixed, average cost per stage up to stage $\overline{\Delta}$, whereas the right hand side denotes the expected cost from all visits of stages larger than $\overline{\Delta}$. If the latter is lower than the former, then $U(\overline{\Delta}) \geq U(\overline{\Delta}+1)$ and it makes sense to increase the threshold. Otherwise, it makes sense to decrease it.

Proposition 6: Let $\overline{\Delta}^*$ be the optimal threshold. For any $\overline{\Delta} \leq \overline{\Delta}^*$, it holds that $U(\max\{\overline{\Delta}-1,1\}) \geq U(\overline{\Delta})$, whereas for any $\overline{\Delta} > \overline{\Delta}^*$, it holds that $U(\overline{\Delta}) \geq U(\overline{\Delta}-1)$.

Proof: The restriction to $\max\{\overline{\Delta}-1,1\}$ is used solely to ensure that $U(\overline{\Delta})$ is defined. By re-formulating (20), we obtain that $U(\overline{\Delta}) \geq U(\overline{\Delta}+1)$ if, and only if,

$$\Upsilon(\overline{\Delta}) - \frac{D}{p_s} \leq 0, \tag{21}$$

where $\Upsilon(\overline{\Delta}) = \overline{\Delta}p_s \sum_{r=0}^{\infty}(1-p_s)^r f(\overline{\Delta}+1+r) - \sum_{\Delta=1}^{\overline{\Delta}} f(\Delta)$. We can bound the terms $\Upsilon(\overline{\Delta})$ as:

$$\Upsilon(\overline{\Delta}) = \overline{\Delta}p_s \sum_{r=0}^{\infty}(1-p_s)^r f(\overline{\Delta}+1+r) - \sum_{\Delta=1}^{\overline{\Delta}} f(\Delta)$$

$$\overset{(1)}{\geq} \overline{\Delta}p_s \sum_{r=0}^{\infty}(1-p_s)^r f(\overline{\Delta}+r) - \sum_{\Delta=1}^{\overline{\Delta}} f(\Delta)$$

$$\overset{(2)}{\geq} (\overline{\Delta}-1)p_s \sum_{r=0}^{\infty}(1-p_s)^r f(\overline{\Delta}+r) + f(\overline{\Delta}) - \sum_{\Delta=1}^{\overline{\Delta}} f(\Delta)$$

$$= (\overline{\Delta}-1)p_s \sum_{r=0}^{\infty}(1-p_s)^r f(\overline{\Delta}+r) - \sum_{\Delta=1}^{\overline{\Delta}-1} f(\Delta) = \Upsilon(\overline{\Delta}-1).$$

Step (1) follows from monotonicity of $f(\Delta)$ and (2) from the equality $\sum_{r=0}^{\infty}(1-p_s)^r f(\overline{\Delta}) = f(\overline{\Delta})/p_s$, in combination with $f(\Delta+r) \geq f(\Delta)$ $\forall r \geq 0$. The result is that $\Upsilon(\overline{\Delta})$ is non-decreasing in $\overline{\Delta}$. Note that it is also non-negative, because

$$\overline{\Delta}p_s \sum_{r=0}^{\infty}(1-p_s)^r f(\overline{\Delta}+r) \geq \overline{\Delta} f(\overline{\Delta}+1) \geq \sum_{\Delta=1}^{\overline{\Delta}} f(\Delta).$$

The only reasons for $\Upsilon(\overline{\Delta}) - D/p_s$ to be negative thus are a large enough D or low enough p_s. On the other hand, the terms D/p_s are constant, such that the left hand side of (21) is non-decreasing. So if there is a finite, optimal $\overline{\Delta}^* \geq 1$, then by Proposition 5, $U(\overline{\Delta})$ is non-increasing $\forall \overline{\Delta} \leq \overline{\Delta}^*$ and non-decreasing $\forall \overline{\Delta} \geq \overline{\Delta}^*$. ∎

We can use Proposition 6 to decide whether a threshold is above or below the optimal one. In particular, a corollary is that we can identify $\overline{\Delta}^*$ as

$$\overline{\Delta}^* = \inf\{\overline{\Delta} \geq 1 : \Upsilon(\overline{\Delta}) \geq D/p_s\}. \tag{22}$$

This result leads to intuitive conclusions: For example, if $D = 0$ the optimal threshold must always be $\overline{\Delta}^* = 1$,

Algorithm 1 Binary Search for $\overline{\Delta}^*$

1: Set LB = UB = 1
2: **while** $\Upsilon(\text{UB}) < D/p_s$ **do**
3:     LB := UB
4:     UB := $2 * \text{UB}$
5: **end while**
6: $\overline{\Delta}^* = \lceil(\text{LB} + \text{UB})/2\rceil$
7: **while** $\overline{\Delta}^* < \text{UB}$ **do**
8:     **if** $\Upsilon(\overline{\Delta}^*) < D/p_s$ **then**
9:         LB := $\overline{\Delta}^*$
10:     **else**
11:         UB := $\overline{\Delta}^*$
12:     **end if**
13:     $\overline{\Delta}^* = \lceil(\text{LB} + \text{UB})/2\rceil$
14: **end while**
15: Output: $\overline{\Delta}^*$

Fig. 4. Expected infinite-time averaged cost for different thresholds $\overline{\Delta}$ and cost functions $f(\Delta)$. The values $D = 10$, $p_s = 0.9$, $A = 0.9$, $\Sigma = 1$, $\theta = 0.6$, $\alpha = 11$ and $K = 4$ were used.

while if $f(\Delta) = 0\ \forall \Delta \geq 1$ but $D > 0$, it is unbounded, i.e., $\overline{\Delta}^* = \infty$. In some cases we can use (22) to derive $\overline{\Delta}^*$ analytically, however this often leads to complex formulas that are out of the scope of this work. Here, we thus propose to use a binary search as shown in Algorithm 1. The algorithm comprises of two loops: In the first one, the value of UB is increased exponentially, until it is ensured to be above the optimal threshold. Then, a binary search is realized in the second loop, which stops when the optimal point is found.

We can argue that Algorithm 1 terminates within a finite number of steps for finite optimal thresholds and has logarithmic complexity with respect to the optimal threshold value: For any finite, optimal threshold $\overline{\Delta}^*$ the first loop terminates after a finite number of iterations because UB is increased in each iteration and $\Upsilon(\text{UB}) \geq D/p_s$ whenever $\text{UB} \geq \overline{\Delta}^*$. The second loop must terminate because the difference $\text{UB} - \text{LB}$ is decreased in each iteration as long as $\text{UB} > \text{LB} + 1$, while $\overline{\Delta}^* = \text{UB}$ if $\text{UB} = \text{LB} + 1$. Finally, assume that $\overline{\Delta}^*$ is the optimal threshold, then the first while loop will terminate after $K_1 := \lceil\log_2(\overline{\Delta}^*)\rceil$ steps, leading to $\text{UB} = 2^{K_1}$ and $\text{LB} = 2^{K_1-1}$, and the second after

$$K_2 := \lceil\log_2(2^{K_1} - 2^{K_1-1})\rceil = \lceil\log_2(2^{K_1-1}(2-1))\rceil = K_1 - 1.$$

Let N_1 be the number of operations required to operate the first while loop, N_2 the ones for line 6 and N_3 that for operating the second loop, then the algorithm will terminate after $N_1 \lceil\log_2(\overline{\Delta}^*)\rceil + N_2 + N_3(\lceil\log_2(\overline{\Delta}^*)\rceil - 1)$ operations. That is, the complexity of Algorithm 1 is $\Omega(\log_2(\overline{\Delta}^*))$.

Algorithm 1 relies only on being able to evaluate the function $\Upsilon(\overline{\Delta})$. While this could be challenging due to the infinite sum involved, it leads to expressions that can be identified relatively easily for the cases we are interested in. Also, the sum can be reasonably approximated with a finite number of summands, as these converge to zero for bounded functions.
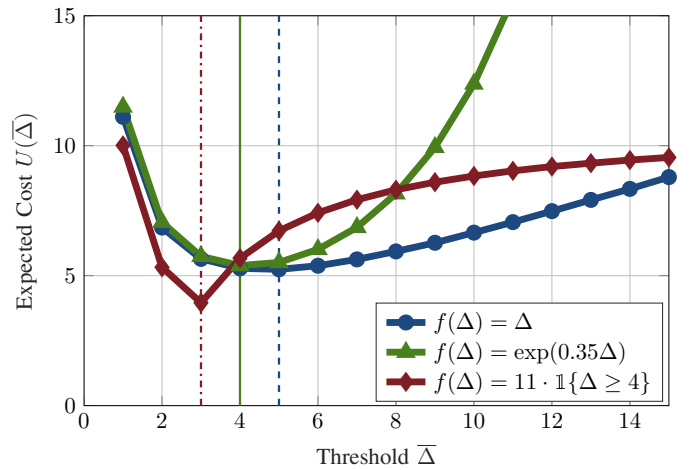
V. SIMULATION RESULTS

We now present some simulation results. The system described in Figure 1 is considered, with $D = 10$, $p_s = 0.9$ and a scalar LTI system with parameters $A = 0.9$, $\Sigma = 1$. Threshold policies are simulated, where $\overline{\Delta}$ is increased and for each value, transmissions over a single link with defined packet success rate are simulated according to the respective law.

Figure 4 shows the results for four different penalty functions, which are: Linear AoI, mean square error and error excess probability of an LTI system, as well as an AoI excess indicator. Each point in the figure is a time-average cost over 10^5 time steps and represents the expected time-average total cost $U(\overline{\Delta})$. We also calculate the optimal thresholds according to Algorithm 1, which are shown with dashed lines in the figure. It can be seen that the cost curves have the properties derived in Section IV-C, i.e., there is a unique optimum and the cost is decreasing for smaller and increasing for larger thresholds, and that the proposed algorithm indeed finds the optimal threshold. Further, it can be seen that the optimal threshold strongly depends on the used penalty. This is particularly the case for the error excess probability shown in yellow, which is monotonically decreasing for the given setting, i.e., has an infinite optimal threshold. The reason is that in this case $f(\Delta)$ is upper bounded and the right hand side of (20) is smaller than D/p_s, or intuitively, that the expected cost of allowing larger AoI is always lower than the expected cost for transmitting. From these results it becomes clear that while the targeted NCS scenarios in principle require timely transmission, considering the AoI itself as metric only implicitly transfers to actually solving the problem at hand. However, using AoI-penalties can help to solve the respective problems optimally.

VI. Conclusion

In this work, we investigated the age of information (AoI) metric from the perspective of networked control systems (NCS). We showed that estimation problems of linear time invariant (LTI) systems can be expressed in terms of an AoI-penalty for different minimization goals. Due to the variety of possible AoI penalties, we investigated what makes a penalty non-trivial to optimize and found the main property of those functions to be positive monotonicity. Finally, we investigated an average age-penalty minimization problem for a single link with packet losses and found the optimal policy to be threshold based. For the optimal threshold we proposed a binary search algorithm and proved its optimality both analytically and by means of simulations.

Appendix A
Proof of Proosition 1

Proof: Consider the sequence $\mathcal{K} = \{t_1, ..., t_i, ...\}$. We can choose a sequence $\{c_1, ..., c_i, ...\}$ of constants such that $0 \le c_i < t_i' - t_i$, $\forall i$, to generate a new sequence $\mathcal{K}_c = \{t_1 + c_1, ..., t_i + c_i, ...\}$ while $\mathcal{T}$ is unchanged. Note that $t_i + c_i < t_i'$ and that as the number of sampled packets remains the same, $\lambda = \mathbb{E}\{Y\}$ is unchanged. The sums over Q_i change to Q_i' as

$$Q_i' = \sum_{t=t_{i-1}'}^{t_i'-1} f(t - t_{i-1} - c_{i-1}) \le \sum_{t=t_{i-1}'}^{t_i'-1} f(t - t_{i-1}) = Q_i$$

because $f(\tau)$ is non-decreasing, which consequently leads to

$$\overline{F}' = \lim_{T \to \infty} \frac{\sum_{i=1}^{|\mathcal{K}(T)|} Q_i'}{\mathbb{E}\{Y\}|\mathcal{K}(T)|} \le \lim_{T \to \infty} \frac{\sum_{i=1}^{|\mathcal{K}(T)|} Q_i}{\mathbb{E}\{Y\}|\mathcal{K}(T)|} = \overline{F}. \quad (23)$$

As this holds for any sequence of c_i according to the given bounds, $\overline{F}$ is non-increasing with c_i. An optimal choice of t_i thus is the maximum possible value, which is as close to t_i' as possible. In an analog argument, it can be seen that for any $0 \le c_i < t_i' - t_{i-1}'$

$$Q_i' = \sum_{t=t_{i-1}'}^{t_i'-c_i-1} f(t - t_{i-1}) \le \sum_{t=t_{i-1}'}^{t_i'-1} f(t - t_{i-1}) = Q_i,$$

leading to $t_i' \to t_{i-1}'$, $\forall i$. ∎

Appendix B
Proof of Proposition 5

Proof: The result of Proposition 5 follows mainly from mathematical reformulations, for which we only give some intermediate steps due to the large formulas involved. Observe that $U(\overline{\Delta}) \ge U(\overline{\Delta} + 1)$ transfers to:

$$\frac{\mathbb{E}\{\overline{C}_\mu \mid \overline{\Delta}\}}{\mathbb{E}\{\overline{N}_\mu \mid \overline{\Delta}\}} \ge \frac{\mathbb{E}\{\overline{C}_\mu \mid \overline{\Delta} + 1\}}{\mathbb{E}\{\overline{N}_\mu \mid \overline{\Delta} + 1\}} \quad (24)$$

By plugging in the explicit expressions, this leads to

$$\frac{\sum_{\Delta=1}^{\overline{\Delta}-1} f(\Delta) + \sum_{r=0}^{\infty}(1-p_s)^r \left[f(\overline{\Delta}+r)+D\right]}{\overline{\Delta}-1+\frac{1}{p_s}}$$

$$\ge \frac{\sum_{\Delta=1}^{\overline{\Delta}} f(\Delta) + \sum_{r=0}^{\infty}(1-p_s)^r \left[f(\overline{\Delta}+1+r)+D\right]}{\overline{\Delta}+\frac{1}{p_s}}$$

$$\frac{\sum_{\Delta=1}^{\overline{\Delta}} f(\Delta) + D + (1-p_s)\sum_{r=0}^{\infty}(1-p_s)^r \left[f(\overline{\Delta}+r)+D\right]}{\overline{\Delta}-1+\frac{1}{p_s}}$$

$$\ge \frac{\sum_{\Delta=1}^{\overline{\Delta}} f(\Delta) + \sum_{r=0}^{\infty}(1-p_s)^r \left[f(\overline{\Delta}+1+r)+D\right]}{\overline{\Delta}+\frac{1}{p_s}}$$

$$\quad (25)$$

$$\sum_{\Delta=1}^{\overline{\Delta}} f(\Delta) + \sum_{r=0}^{\infty}(1-p_s)^r \left[f(\overline{\Delta}+1+r)+D\right] + D\left(\overline{\Delta}+\frac{1}{p_s}\right)$$

$$\ge p_s\left(\overline{\Delta}+\frac{1}{p_s}\right)\left[\sum_{r=0}^{\infty}(1-p_s)^r \left[f(\overline{\Delta}+1+r)+D\right]\right]$$

$$\sum_{\Delta=1}^{\overline{\Delta}} f(\Delta) + D\left(\overline{\Delta}+\frac{1}{p_s}\right)$$

$$\ge p_s\overline{\Delta}\left[\sum_{r=0}^{\infty}(1-p_s)^r \left[f(\overline{\Delta}+1+r)+D\right]\right]$$

$$\frac{1}{\overline{\Delta}}\left[\sum_{\Delta=1}^{\overline{\Delta}} f(\Delta) + \frac{D}{p_s}\right] \ge p_s\sum_{r=0}^{\infty}(1-p_s)^r f(\overline{\Delta}+1+r).$$

∎

References

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

[2] S. K. Kaul, R. D. Yates, and M. Gruteser, "Status updates through queues," *2012 46th Annual Conference on Information Sciences and Systems, CISS 2012*, pp. 1–6, 2012.

[3] A. M. Bedewy, Y. Sun, and N. B. Shroff, "Age-Optimal Information Updates in Multihop Networks," *IEEE International Symposium on Information Theory - Proceedings*, pp. 576–580, jan.

[4] C. Kam, S. Kompella, G. D. Nguyen, J. E. Wieselthier, and A. Ephremides, "Controlling the age of information: Buffer size, deadline, and packet replacement," in *MILCOM 2016 - 2016 IEEE Military Communications Conference*. IEEE, nov, pp. 301–306.

[5] A. M. Bedewy, Y. Sun, and N. B. Shroff, "Optimizing data freshness, throughput, and delay in multi-server information-update systems," in *2016 IEEE International Symposium on Information Theory (ISIT)*. IEEE, jul, pp. 2569–2573.

[6] L. Huang and E. Modiano, "Optimizing age-of-information in a multi-class queueing system," in *2015 IEEE International Symposium on Information Theory (ISIT)*. IEEE, jun, pp. 1681–1685.

[7] C. Kam, S. Kompella, G. D. Nguyen, and A. Ephremides, "Effect of Message Transmission Path Diversity on Status Age," *IEEE Transactions on Information Theory*, no. 3, pp. 1360–1374, mar.

[8] R. Talak, S. Karaman, and E. Modiano, "Minimizing age-of-information in multi-hop wireless networks," in *2017 55th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*. IEEE, oct, pp. 486–493.

[9] C. Kam, S. Kompella, G. D. Nguyen, J. E. Wieselthier, and A. Ephremides, "Age of information with a packet deadline," in *2016 IEEE International Symposium on Information Theory (ISIT)*. IEEE, jul, pp. 2564–2568.

[10] I. Kadota, E. Uysal-Biyikoglu, R. Singh, and E. Modiano, "Minimizing the Age of Information in broadcast wireless networks," in *2016 54th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*. IEEE, sep, pp. 844–851.

[11] Y.-P. Hsu, E. Modiano, and L. Duan, "Scheduling Algorithms for Minimizing Age of Information in Wireless Broadcast Networks with Random Arrivals," pp. 1–13, dec.

[12] Y.-p. Hsu, E. Modiano, and L. Duan, "Age of information: Design and analysis of optimal scheduling algorithms," in *2017 IEEE International Symposium on Information Theory (ISIT)*. IEEE, jun, pp. 561–565.

[13] Y. Sun, E. Uysal-biyikoglu, R. Yates, C. E. Koksal, and N. B. Shroff, "Update or Wait : How to Keep Your Data Fresh," *IEEE INFOCOM 2016 - The 35th Annual IEEE International Conference on Computer Communications*, pp. 1–9, 2016.

[14] 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*, no. 11, pp. 7492–7508, nov.

[15] A. Kosta, N. Pappas, A. Ephremides, and V. Angelakis, "Age and value of information: Non-linear age case," in *IEEE International Symposium on Information Theory - Proceedings*. IEEE, jun, pp. 326–330.

[16] Y. Sun, Y. Polyanskiy, and E. Uysal-Biyikoglu, "Remote estimation of the Wiener process over a channel with random delay," in *2017 IEEE International Symposium on Information Theory (ISIT)*. IEEE, jun, pp. 321–325.

[17] Y. Sun and B. Cyr, "Information Aging through Queues : A Mutual Information Perspective," pp. 1–5, 2018.

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

[19] A. F. Eugene, P. O. Kasyanov, and N. V. Zadoianchuk, "Average cost markov decision processes with weakly continuous transition probabilities," *Mathematics of Operations Research*, vol. 37, no. 4, pp. 591–607, 2012.