Online Energy Minimization Under A Peak Age of Information Constraint

Kumar Saurav
School of Technology and Computer Science
Tata Institute of Fundamental Research
Mumbai, India.
kumar.saurav@tifr.res.in

Rahul Vaze
School of Technology and Computer Science
Tata Institute of Fundamental Research
Mumbai, India.
rahul.vaze@gmail.com

***Abstract*—We consider a node where packets of fixed size are generated at arbitrary intervals. The node is required to maintain the peak age of information (AoI) at the monitor below a threshold by transmitting potentially only a subset of the generated packets. At any time, depending on packet availability and current AoI, the node can choose the packet to transmit, and its transmission speed. We consider a power function (rate of energy consumption) that is increasing and convex in transmission speed, and the objective is to minimize the energy consumption so as to satisfy the peak AoI constraint at all times. For this problem, we propose a (customized) greedy policy and derive an upper bound on its competitive ratio (CR) that depends on the power function, but is independent of the packet generation times as well as the time horizon. We also derive a lower bound on the CR of all causal policies, and show that the dependence of the CR of the proposed greedy policy on the system parameters (such as packet size, peak AoI and power function) is similar to that of an optimal causal policy.**

***Index Terms*—Age of information, speed scaling, energy consumption**

I. INTRODUCTION

In the times of COVID-19 pandemic, real-time networked applications such as autonomous vehicles, immersive gaming, telemedicine and telesurgery played a vital role in enabling people to maintain social-distancing while minimizing the loss in lives and productivity [1]. However, such applications need to maintain data-freshness at the nodes, which requires minimizing the age of latest data-packet (formally called the *age of information* (AoI) [2]) at the nodes.

In practice, AoI minimization is a complicated task. Factors such as limited energy, unavailability of fresh data-packets, etc., often restrict the control options for minimizing AoI. For example, to maximize the operational life, devices with limited energy need to restrict the number of data-packets that they transmit, as well as their transmission rate (speed), thus leading to large AoI. Hence, there is an inherent AoI-energy tradeoff that any transmission policy must consider. In this paper, we consider a particular instance of AoI-energy tradeoff, where the objective is to minimize the energy consumption subject to a peak AoI constraint at all times, over a fixed horizon of time.

We acknowledge support of the Department of Atomic Energy, Government of India, under project no. RTI4001.

To be precise, we consider a node where data-packets (in short, packets) of fixed size (say, W bits) are generated at arbitrary time instants, with bounded inter-generation time. The node requires that by transmitting a subset of these packets, the peak AoI at the monitor is maintained below a threshold at all times, while consuming minimum energy. We consider a speed scaling model, where the node can choose its transmission speed, and operating at speed s consumes power $P(s)$ that is a convex and increasing function of s. Thus, for every generated packet, the node needs to decide whether to transmit it or not, and at what speed. At any time, the node can also preempt a packet being transmitted. If the node transmits a packet, the AoI drops to the age of the transmitted packet, at the instant the node finishes transmitting the entire W bits of the packet.

A. Prior Work

In past, in the context of AoI, the speed scaling model has been considered in [3], [4]. In particular, in [3], a node controls the number of bits of packet delivered to the monitor in a fixed time slot by adjusting the transmission power. In [4], a node can control the transmission delay of packets by adjusting the transmission power at the instant the packets are transmitted. The speed scaling model has also been considered in context of AoI-distortion tradeoff [5]–[7], where a node controls the number of bits it transmits for each packet. Transmitting fewer bits take less time, and is assumed sufficient for minimizing AoI, but this results in distorted information at the monitor, where the level of distortion increases with decrease (increase) in the number of bits transmitted (transmission time).

The closest work to the considered problem is in the non-AoI setting, for problem called as job scheduling with deadlines [8], [9], where jobs arrive at a server sequentially, and the server needs to process all the arriving jobs (with adaptive speed) before their corresponding deadlines, while consuming minimum energy. When power function is $P(s) = s^\alpha$ ($\alpha > 1$), [8] proposed a causal policy for which the **competitive ratio** (i.e., the ratio of the energy consumed by the causal policy to the energy consumed by an optimal offline policy that knows the generation time of all the packets in advance, maximized over all inputs) is at most α^α [9]. In the special case when the power function is $P(s) = 2^s - 1$, and all the packets

ISBN 978-3-903176-37-9 © 2021 IFIP

are of equal size, with a common deadline, [10], [11] gave a 3-competitive policy for the job scheduling problem.

The problem considered in [8], [9] differs from what we consider in this paper in one key aspect. In [8], [9] every job that arrives at the server needs to be processed completely, while in the considered problem, there is only a peak AoI constraint which could be satisfied by transmitting only a subset of the generated packets. Since the objective is to minimize total energy subject to peak AoI constraint, an additional challenge compared to [8], [9] (where only speed and the order of processing is to be decided) is to identify the subset of packets to transmit.

B. Our Contributions

1) We derive a lower bound on the competitive ratio of all causal policies, and thus characterize the dependence of the competitive ratio of an optimal causal policy on the system parameters such as packet size, peak AoI and power function. For example, for power function $P(s) = s^{\alpha}$ (where $\alpha > 1$, and s denotes speed), we show that the competitive ratio of any causal policy π is $\text{CR}_{\pi} \geq c^{\alpha}$ (where $c \geq 1.07$ is a constant). So, CR_{π} increases exponentially with increase in α.
2) We propose a simple causal greedy policy π^g that at any time t, transmits the latest available packet with a speed $s(t)$ (that depends on the AoI at time t), and show that for convex power function $P(s(t))$, the competitive ratio of the proposed policy has similar order of dependence on the system parameters as the lower bound discussed in point 1. For example, for $P(s) = s^{\alpha}$ $(\alpha > 1)$, we show that the competitive ratio of π^g is $\text{CR}_{\pi^g} \leq 2{\cdot}3^{\alpha}{+}1$.

II. System Model

Consider a node where data-packets (in short, packets), each of size W bits are generated intermittently[1]. In particular, the i^{th} packet at the node is generated at time t_i, where t_i is determined by external factors (possibly adversarial), with inter-generation time $X_i = t_i - t_{i-1}$. A packet i is said to be delivered (at the monitor) at time τ_i if the node finishes transmitting the W bits of packet i at time τ_i.

At any time t, the *age of information* (AoI) at the monitor is equal to $\Delta(t) = t - \mu(t)$, where $\mu(t)$ is the generation time of the latest packet that has been delivered to the monitor until time t. Further, in an interval $[0, T]$, peak AoI at the monitor is defined to be $\max_{t\in[0,T]} \Delta(t)$. We consider a peak AoI constraint, i.e., for any given time horizon T, the node requires that the peak AoI at the monitor in the interval $[0, T]$ is less than D, where D is a known constant.

Remark 1: If inter-generation time $X_i > D$ for any i, then the peak AoI constraint is infeasible. Therefore, for the problem to be meaningful, we need the inter-generation time of packets to be bounded, and in particular, less than D. Therefore, we assume that there exists some $\epsilon > 0$, such that $\forall i$, $X_i < D - \epsilon$. Further, without loss of generality, we assume that the first packet is generated at time $t = 0$, and the initial AoI $\Delta(0) < D - \epsilon$ (required for the problem to be feasible).

Remark 2: At any time t, the peak AoI constraint only requires that $\mu(t)$, i.e., the generation time of the latest packet delivered to the monitor until time t is less than D time units old. Therefore, to satisfy the peak AoI constraint, it is not necessary to transmit every generated packet. For example, at time t, if packets i and j (with generation times t_i and t_j respectively) are available at the node, and $t_i < t_j$, then for satisfying the peak AoI constraint, it is sufficient to transmit packet j only, without transmitting packet i.

Definition 1: A policy that at any time t, *can* interrupt an ongoing transmission of a packet, and begin transmitting a newly generated packet, is called an interruptive policy[2]. Note that the class of interruptive policies, by definition, includes all non-interruptive policies, that never interrupts transmission of any packet.

We consider a speed scaling model, where at any time t, the node can transmit a packet at speed $s(t) \geq 0$ (in bits/sec) adaptively, i.e., the node can choose $s(t)$ using causal information available at time t. Also, transmitting a packet at speed $s(t)$ consumes power $P(s(t))$, which is an increasing and convex function of speed $s(t)$, e.g., $P(s) = s^{\alpha}$ $(\alpha > 1)$, or $P(s) = 2^s - 1$, motivated by Shannon's rate function. Therefore, if the node transmits packets at high speed, it incurs low AoI, but consumes large amount of energy.

In this paper, we consider the problem of finding a causal policy that chooses the subset of packets to transmit (where interruption is allowed), the time interval over which the packets are transmitted, and their instantaneous transmission speed $s(t)$, so that the peak AoI is maintained below D at all times over a time horizon T, while consuming minimum energy. Formally, the objective can be stated as follows.

$$\min_{\pi\in\Pi} \; E_{\pi}(\sigma) = \int_{t=0}^{T} P(s(t))dt \tag{1a}$$

$$\text{s.t.} \;\; \Delta(t) < D, \;\; \forall t \in [0, T], \tag{1b}$$

where Π is the set of all causal policies for packet scheduling and the choice of speed $s(t)$, and $\sigma = \{t_1, t_2, ...\}$ is the sequence of packet generation times (we assume that the time horizon T is known in advance). Note that the choice of packet being transmitted by a policy π at any time t is inherently captured by (1a).

Remark 3: In this paper, the set of all causal policies implicitly refers to the set of all interruptive causal policies that may interrupt ongoing packet transmission (Definition 1).

Remark 4: For simplicity, we assume $P(0) = 0$, which appears as an offset term in (1a), and does not affect the causal policy that solves the optimization problem (1a).

Definition 2: A policy $\pi^{\star}$ is said to be offline optimal if it satisfies the peak AoI constraint (1b), and there exists no other policy $\tilde{\pi}$ that can simultaneously satisfy the peak

[1]The case of non-identical packet sizes has been considered in the companion paper [12].

[2]In literature (e.g., [13]), an interruptive policy is also known as a preemptive policy. However, a non-interruptive policy can discard a packet upon generation, whereas a non-preemptive policy transmits every generated packet.

ISBN 978-3-903176-37-9 © 2021 IFIP

AoI constraint (1b), and consume less energy than $\pi^\star$, even if $\tilde{\pi}$ knows the generation time of all the packets in advance. Optimal offline policies are useful as they provide a lower bound on the energy consumed by any causal policy π.

From prior work [8], [10], [11], it is known that finding an optimal causal policy for energy minimization problems under hard constraints (such as individual/common deadline for packets) is a challenging task. Hence, a usual approach is to find a causal policy π whose competitive ratio, defined as the ratio of the energy $E_\pi(\sigma)$ consumed by a causal policy π (to satisfy the peak AoI constraint) and the energy $E_{\pi^\star}(\sigma)$ consumed by an optimal offline policy $\pi^\star$ (Definition 2), maximized over all possible sequence of packet generation times σ, is small. Mathematically, the **competitive ratio** of policy π is

$$\text{CR}_\pi = \max_\sigma \frac{E_\pi(\sigma)}{E_{\pi^\star}(\sigma)}. \quad (2)$$

By definition, a policy with small competitive ratio is close to optimal offline policy (the strongest benchmark), as well as robust to input variations. In the rest of this paper, we will consider a particular non-interruptive causal policy, and show that its competitive ratio is at most $2P(3\hat{s})/P(\hat{s})+1$, where $\hat{s} = W/D$. We will also derive a lower bound on the competitive ratio of any causal policy π, and show that for different power functions of interest, the competitive ratio of the considered policy has similar characteristics (dependence on parameters) as the derived lower bound. For a non-interruptive causal policy which is much simpler to implement than a general interruptive causal policy, this is a significant result.

A. An Equivalent Deadline Constraint Problem

The peak AoI constraint (1b) can also be interpreted as a deadline constraint, where a deadline is defined as follows.

Definition 3: At any time t, if $\mu(t)$ is the generation time of the latest packet that has been delivered to the monitor until time t, then the deadline at time t is defined as $d(t) = \mu(t) + D$, which is the earliest time instant at which the peak AoI constraint (1b) will be violated if no packet is delivered to the monitor after time t. Equivalently, $d(t) = t + (D - \Delta(t))$.

Remark 5: Note that $d(t) = \mu(t) + D$ is a non-decreasing function of t. In fact, deadline $d(t)$ increases in steps whenever a fresh packet j ($t_j > \mu(t)$; see Definition 4 below) is delivered to the monitor. This happens because $\mu(t)$ (i.e., the generation time of the latest packet delivered until time t) increases discontinuously to the generation time t_j of packet j, at the instant packet j is delivered to the monitor.

Definition 4: At any time t, a packet i is defined to be fresh if its generation time $t_i \leq t$ is greater than $\mu(t)$, i.e. $t_i > \mu(t)$. Otherwise, the packet is stale.

Note that the peak AoI constraint (1b) is satisfied if and only if $D - \Delta(t) > 0$, for all $t \in [0, T]$. Also, at any time t, $d(t) = t + (D - \Delta(t))$ is greater than t if and only if $D - \Delta(t) > 0$. Therefore, the peak AoI constraint (1b) is equivalent to the following deadline constraint:

$$d(t) > t, \quad \forall t \in [0, T]. \quad (3)$$

In other words, the peak AoI constraint (1b) is equivalent to the constraint that at any time $t \in [0, T]$, the current deadline $d(t)$ must be in future. Hence, hereafter we consider the deadline constraint (3) (instead of peak AoI constraint (1b)) while minimizing the objective function (1a).

Remark 6: As we show in Section V, considering peak AoI constraint (1b) as deadline constraint (3) reveals several key properties of an optimal offline policy $\pi^\star$, and simplifies the overall analysis in this paper.

Definition 5: A policy π is defined to be feasible if it satisfies the deadline constraint (3) at all times $t \in [0, T]$.

Proposition 1: At any time $t \in [0, T]$, if $d(t) \leq T$, then for any feasible policy π (Definition 5), at least one fresh packet has to be delivered to the monitor in interval $[t, d(t))$.

Proof: If the deadline at time $t \in [0, T]$ is $d(t) \leq T$, and no fresh packet is delivered to the monitor in interval $[t, d(t))$, then the deadline constraint (3) will be violated at time $d(t)$ (follows from the definition of $d(t)$; Definition 3). Hence, a feasible policy must deliver at least one fresh packet to the monitor in interval $[t, d(t))$, for all $t \in [0, T]$, if $d(t) \leq T$. ■

Remark 7: In order to satisfy the deadline constraint (3), note that only fresh packets are needed/useful. Hence, in rest of the paper, we only consider packets that are fresh, and at any time, the term 'packet' implicitly means a fresh packet.

Definition 6: For each packet i generated at time t_i, we define $d_i = t_i + D$.

B. Property of Convex Power Function

Convexity of power function implies the following results.

Lemma 1: Energy consumed in transmitting w bits in a fixed time interval $[p, q)$ is minimum if for $t \in [p, q)$, the bits are transmitted at a constant speed $s_w(t) = w/(q-p)$. Also, the minimum energy consumed for transmitting the w bits in interval $[p, q)$ is $P(w/(q-p))(q-p)$.

Corollary 1: For fixed w, $P(w/y)y$ decreases with increase in y.

III. Lower Bound on the Competitive Ratio of all Causal Policies

Before we discuss a particular causal policy for minimizing the energy consumption (1a) (under the deadline constraint (3)), it is important to note the fundamental limitations of any causal policy π. Towards that end, Theorem 1 provides a lower bound on the competitive ratio of any causal policy π, and shows (in Corollary 2) that for certain power functions $P(\cdot)$, the competitive ratio of any causal policy π is an increasing function of W/D.

Theorem 1: For any causal policy π, its competitive ratio

$$\text{CR}_\pi \geq \frac{c_1 P(c_2 W/D)}{P(c_3 W/D)}, \quad (4)$$

where c_1, c_2 and c_3 are finite positive constants, $c_2 - c_3 \geq 0.14$, and $c_2/c_3 \geq 1.07$.

Proof Sketch: Recall that competitive ratio of a causal policy π is defined to be the worst case ratio of the energy consumed by the policy π to the energy consumed by an

ISBN 978-3-903176-37-9 © 2021 IFIP

optimal offline policy $\pi^\star$ over all inputs. Thus, to derive the lower bound (4) on the competitive ratio of all causal policies (i.e., to prove Theorem 1), we consider a particular scenario where the AoI at the monitor at time $t = 0$ is $\Delta(0) = D/2$, the time horizon $T = 3D/2 - \delta$ (for $\delta \to 0^+$), and the packets are generated according to one of the two instances of packet generation times σ: (i) $\sigma_1 = \{0, D/4, D/2\}$, and (ii) $\sigma_2 = \{0, D/4, 5D/6\}$. In this scenario, for indexing all causal policies, we consider different cases based on the packet(s) that any causal policy may transmit in interval $[0, D/2)$, and for each case, we compute the competitive ratio (2), where the maximization is with respect to $\sigma \in \{\sigma_1, \sigma_2\}$. Finally, we take the minimum over the competitive ratio obtained for different cases considered above, and obtain (4), where c_1, c_2 and c_3 are finite positive constants, $c_2 - c_3 \geq 0.14$, and $c_2/c_3 \geq 1.07$. For the detailed proof, see Technical Report [12]. ■

Corollary 2: For any causal policy π, (i) if $P(s) = s^\alpha$ ($\alpha > 1$), the competitive ratio $\text{CR}_\pi \geq c_1(c_2/c_3)^\alpha \geq c_1 1.07^\alpha$ increases exponentially with increase in α, while (ii) for $P(s) = 2^s - 1$, the competitive ratio $\text{CR}_\pi \geq 2^{(c_2-c_3)W/D} \geq 2^{0.14W/D}$ increases exponentially with increase in W/D.

Although the lower bound (4) in Theorem 1 is for all interruptive policies, in the next section, we propose a simple non-interruptive causal greedy policy π^g, and show that its competitive ratio (2) is upper bounded by $c_1' P(c_2' W/D)/P(c_3' W/D) + 1$ (where $c_1' = 2$, $c_2' = 3$, and $c_3' = 1$). Thus, we show that the dependence of the competitive ratio of π^g on the system parameters is similar to the policy-independent lower bound (4).

IV. A Greedy Policy π^g

Consider a greedy policy π^g (Algorithm 1) that at any time t, if the node is idle (i.e., not transmitting any packet) and the deadline $d(t) < T$, transmits the latest available (fresh) packet with *constant* speed $s^g(t)$ (5), starting at time t, throughout until the W bits of the chosen packet are delivered to the monitor (and waits otherwise),

$$s^g(t) = \max\left\{\frac{W}{d(t) - t}, \frac{W}{D/3}\right\}. \tag{5}$$

Remark 8: Note that the greedy policy π^g never interrupts any ongoing packet transmission. However, this is not a constraint, and in general, a policy is allowed to interrupt ongoing packet transmission to solve (1a).

Algorithm 1 Greedy Policy π^g.

$t \leftarrow$ current time;
if $d(t) \leq T$ **and** node idle **and** packet available **then**
  transmit the latest packet with constant speed $s^g(t)$ (5) throughout the time interval $[t, t + W/s^g(t))$;
end if

Theorem 2: Let $\hat{s} = W/D$. Then, the competitive ratio (CR_{π^g}) of greedy policy π^g is bounded as follows.

$$\text{CR}_{\pi^g} \leq \frac{2P(3\hat{s})}{P(\hat{s})} + 1. \tag{6}$$

Remark 9: In (6), if $P(s) = s^\alpha$ ($\alpha > 1$), $\text{CR}_{\pi^g} \leq 2\cdot 3^\alpha + 1$, while if $P(s) = 2^s - 1$, $\text{CR}_{\pi^g} \leq 2(2^{3W/D} - 1)/(2^{W/D} - 1) + 1$. Note that this dependence of CR_{π^g} on the parameters is similar to that of the lower bound in Theorem 1 for all causal policies.

Although greedy policy π^g may appear natural, the speed $s^g(t)$ (5) has been chosen carefully such that (i) π^g is feasible (since $s^g(t) \geq W/(d(t) - t)$, if π^g begins to transmit a packet at time t, the packet will be delivered to the monitor before deadline $d(t)$), and (ii) speed $s^g(t)$ cannot be arbitrarily large (in fact, $s^g(t)$ cannot be greater than $3W/D$), unless a particular event happens (defined in Proposition 2) with regard to the sequence of packet generation times, and for which we can lower bound the energy consumed by an optimal offline policy $\pi^\star$ (Lemma 4 in Appendix C).

Proposition 2: If π^g begins to transmit a packet j at time t, then the speed $s^g(t) > 3W/D$ only if no packet was generated in interval $[t - 2D/3, t)$.

Proof: See Appendix A. ■

We next provide some intuition for the choice of speed $s^g(t)$ (5) used by the greedy policy π^g, that is at least equal to $3W/D$. For a greedy policy such as Algorithm 1, it is critical to have $s^g(t) \geq 3W/D$, because as shown in Example 1 below, a smaller value of speed such as $2W/D$ may require π^g to transmit some packets at much higher speed (compared to an optimal offline policy $\pi^\star$), thus consuming large amount of energy (compared to the offline optimal policy $\pi^\star$).

Example 1: Let at time $t = 0$, $\Delta(0) = 0$, $D = 2$, and $T = 3 + \delta/2$, where $\delta \to 0^+$. Also, let three packets be generated at time $t = 0, t = \delta$, and $t = 1 + \delta$, respectively. Then, an optimal offline policy $\pi^\star$ will only transmit the third packet, with constant speed $W/(1-\delta)$ for $(1-\delta)$ time units, whereas the greedy policy π^g (Algorithm 1) with speed $s^g(t) = \max\{W/(d(t) - t), 3W/D\}$ will transmit all three packets with constant speed $3W/D$ (each for $D/3$ time units).

However, if the speed choice $s^g(t)$ (5) for π^g is replaced with $\hat{s}^g(t) = \max\{W/(d(t)-t), 2W/D\}$ (where t is the time when transmission of packet begins), then π^g will still transmit all three packets, but the third packet will be transmitted with a constant speed $W/\delta \to \infty$.

In the rest of this paper, we prove Theorem 2 in two steps. In step 1 (Section V), we derive some structural results for an optimal offline policy $\pi^\star$, and then in step 2 (Section VI), we use the derived results for the optimal offline policy $\pi^\star$ to derive the upper bound (6) on the competitive ratio CR_{π^g}.

V. Properties of an Optimal Offline Policy $\pi^\star$

Consider an optimal offline policy $\pi^\star$. In this section, without loss of generality, we only consider the packets that are transmitted by $\pi^\star$ in interval $[0, T]$, and index them as $1, 2, 3, \ldots$ in ascending order of their generation times (i.e., $t_1 < t_2 < t_3 < \ldots$). Therefore, between the generation time of (transmitted) packets $i-1$ and i, many other packets might have been generated, however, they are not transmitted by $\pi^\star$.

Lemma 2: If $\pi^\star$ chooses to transmit packet i, and $d_i \leq T$ (where $d_i = t_i + D$), then in interval $[t_i, d_i)$, $\pi^\star$ transmits at least two packets completely (i.e., $2W$ bits).

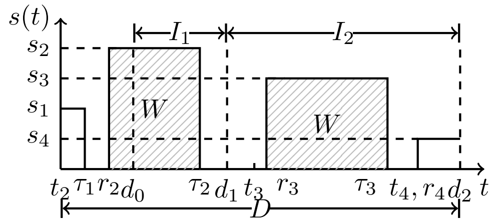

(a) Typical speed profile in a period. Here, a packet generated at t_i is scheduled for transmission at r_i, and delivered to the monitor at τ_i.

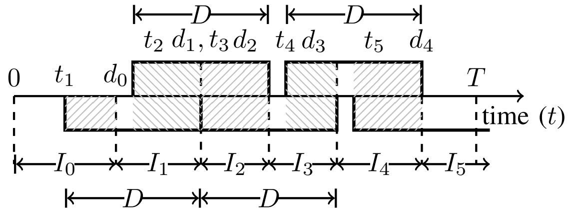

(b) Time-axis as union of disjoint intervals (frames).

Fig. 1: Periods and Frames.

Proof: If $\pi^\star$ delivers packet i at time τ_i, then the deadline at time τ_i is $d(\tau_i) = d_i$ (where $d_i = t_i + D$). Since $d_i \le T$, from Proposition 1, it follows that $\pi^\star$ delivers packet $i+1$ in interval $[\tau_i, d_i)$ (i.e., $\tau_{i+1} \in [\tau_i, d_i)$. Thus, in interval $[t_i, d_i)$, $\pi^\star$ transmits at least two complete packets i and $i+1$ (i.e., $2W$ bits), as shown in Figure 1(a) (for $i = 2$). ∎

Thus, Lemma 2 implies that for $\pi^\star$, the intervals $[t_i, d_i)$ are special, and hence, we call them *periods*, defined next.

Definition 7: With respect to $\pi^\star$, the interval $\chi_i^{\pi^\star} = [t_i, d_i)$ (where $d_i = t_i + D$) is called **period** i.

By definition, a period starts at the generation time of a packet transmitted by $\pi^\star$, and in each period i, $\pi^\star$ transmits at least two packets (from Lemma 2). Thus, consecutive periods overlap as shown in Figure 1(b). Hence, it is difficult to generalize Lemma 2 directly to the whole interval $[0, T]$. Therefore, we further define *frames* (that are non-overlapping) with respect to $\pi^\star$ as follows.

Definition 8: With respect to $\pi^\star$, the interval $I_i = [d_{i-1}, d_i)$ (where $d_{i-1} = t_{i-1} + D$ and $d_i = t_i + D$) is called **frame** i.

Remark 10: As shown in Figure 1(b), consecutive frames partition the time-axis between $[0, T]$ (assuming the initial frame I_0 starts at time $t = 0$). Therefore, $\forall i \neq j$, $I_i \cap I_j = \phi$, and there exists consecutive frames $0, 1, ..., m$, such that T lies in frame m, and $[0, T] \subseteq \cup_{i=1}^m I_i$. So, the properties of $\pi^\star$ in a frame can be easily generalized to the entire interval $[0, T]$.

Remark 11: In Definition 7 and Definition 8, note that periods and frames are defined with respect to the packets transmitted by an optimal offline policy $\pi^\star$. Also, frame i (interval $[d_{i-1}, d_i)$) is always a proper subset of period i (interval $[t_i, d_i)$), with $d_{i-1} > t_i$. Hence, length of frame i is always less than $d_i - t_i = D$.

Figure 1(a) shows a typical relation between periods and frames, where at time t_2, the deadline is d_0, and packet 1 is delivered to the monitor at time $\tau_1 < d_0$. Thus the deadline is updated at time τ_1 to d_1. Then, the interval $I_1 = [d_0, d_1)$ is called frame 1. Similarly, within interval I_1, packet 2 (with generation time t_2) is delivered to the monitor at time $\tau_2 < d_1$. Therefore, at time τ_2, the deadline gets updated to d_2, and the interval $I_2 = [d_1, d_2)$ is called frame 2.

Remark 12: Although Figure 1(a) shows typical properties of frames and periods defined with respect to the packets transmitted by $\pi^\star$, the speed profile shown in Figure 1(a) is not necessarily the speed chosen by $\pi^\star$, since that $(\pi^\star)$ is unknown. We show in Proposition 3 that within a frame, speed of $\pi^\star$ exhibits several structural properties.

Proposition 3: The optimal offline policy $\pi^\star$

1) transmits the W bits of a packet with constant speed,
2) never interrupts an ongoing packet transmission,
3) delivers packet $i+1$ in frame i ($\forall i$), and
4) never decreases the transmission speed within a frame.

Proof: See Appendix B. ∎

VI. Proof of Theorem 2

Consider an arbitrary sequence of packet generation times σ. From Remark 10, we know that the time axis can be partitioned into frames defined with respect to $\pi^\star$ (Definition 8). Therefore, consider the consecutive frames $0, 1, ..., m$ such that the total time horizon interval $[0, T]$ is a subset of the union of frames 0 to m, and time horizon T lies in frame m (as shown in Figure 1(b) for $m = 5$). Note that if $m = 0$ (i.e., T is less than the initial deadline $d(0)$), then neither π^g, nor $\pi^\star$ transmits any packet because the deadline constraint (3) is trivially satisfied in the interval $[0, T]$. Hence, we only consider the case where $m \ge 1$.

Since length of a frame is always less than D (Remark 11), the time horizon $T < (m+1)D$. Therefore, in interval $[0, T]$, if π^g transmits $x^g \ge 0$ number of packets with speed $3W/D$ consuming E_x^g units of energy, then $E_x^g < P(3W/D)(m+1)D$ (product of power consumption and upper bound on the length of time interval $[0, T]$). Let $y^g \ge 0$ denote the number of packets that π^g transmits with speed greater than $3W/D$ (π^g transmits an entire packet at a constant speed; either equal to $3W/D$, or greater than $3W/D$), and E_y^g denotes the total energy consumed by π^g in transmitting these y^g packets. Then, the total energy consumed by π^g in interval $[0, T]$ is

$$E_{\pi^g} = E_x^g + E_y^g \le (m+1)P(3W/D)D + E_y^g. \quad (7)$$

Remark 13: Since π^g transmits each of the y^g packets at a constant speed greater than $3W/D$, the energy consumed by π^g in transmitting the y^g number of packets is $E_y^g > y^g P(3W/D)D/3$.

From Proposition 3 (Property 3), it follows that $\pi^\star$ delivers exactly one packet in each frame $i = 0, 1, 2, ..., m-1$. Thus, $\pi^\star$ transmits m complete packets in interval $[0, T]$. Also, the length of each frame is less than D (length of a period). Hence, the energy consumed by $\pi^\star$ in transmitting each of these m packets is at least $P(W/D)D$. Next, we show in Lemma 3 that out of these m packets that $\pi^\star$ transmits completely, there exists a subset $\mathcal{Z}$ consisting of y^g number of packets such that $\pi^\star$ consumes at least E_y^g units of energy in transmitting the packets in $\mathcal{Z}$ (where y^g and E_y^g are defined as in Remark 13).

Lemma 3: There exists a subset $\mathcal{Z}$ consisting of y^g number of packets such that $\pi^\star$ transmits all the packets in $\mathcal{Z}$, and

ISBN 978-3-903176-37-9 © 2021 IFIP

consumes at least E_y^g units of energy while transmitting the packets in $\mathcal{Z}$.

Proof: Note that for $y^g = 0$, the claim is trivially satisfied. So, for the rest of the proof, we assume $y^g \geq 1$. Recall that π^g transmits y^g number of packets at speed greater than $3W/D$, consuming E_y^g units of energy. Without loss of generality, let these packets be indexed as $1, 2, ..., y^g$. Also, let π^g transmit a packet $j \in \{1, 2, ..., y^g\}$ over a contiguous time interval U_j (π^g always transmits packets over contiguous intervals), and consumes energy e_j (in transmitting packet j). Since π^g transmits only one packet at a time, $U_i \cap U_j = \phi$ for $i \neq j$. Therefore, to prove Lemma 3, it is sufficient to show that in each interval U_j (where $j \in \{1, 2, ..., y^g\}$), $\pi^\star$ transmits at least one packet completely (entire W bits), consuming at least e_j units of energy. Hence, to conclude the proof, we show in Lemma 4 (in Appendix C), that for each packet j that π^g transmits with speed greater than $3W/D$, $\pi^\star$ transmits at least one packet $\hat{j}$ (where packet j and $\hat{j}$ may be same) completely during the time interval when π^g transmits packet j, at a constant speed that is at least equal to the constant speed with which π^g transmits packet j. ∎

Remark 14: $\pi^\star$ transmits a total of m packets in interval $[0, T]$ (exactly one packet in each of the frames $0, 1, 2, ..., m-1$). Also, from Lemma 3, it follows that $\pi^\star$ transmits at least y^g number of packets. Therefore, $y^g \leq m$.

Hence, the total energy consumed by $\pi^\star$ in interval $[0, T]$ is

$$E_{\pi^\star} \geq (m - y^g)P(W/D)D + E_y^g. \tag{8}$$

From (7) and (8), we obtain an upper bound on the competitive ratio (2) for π^g as follows.

$$\begin{aligned}
\text{CR}_{\pi^g} &\leq \frac{(m+1)P(3W/D)D + E_y^g}{(m - y^g)P(W/D)D + E_y^g}, \\
&\overset{(a)}{\leq} \frac{(m+1)P(3W/D)D}{(m - y^g)P(W/D)D + y^g P(3W/D)D/3} + 1, \\
&\overset{(b)}{\leq} \frac{(m+1)P(3W/D)}{mP(W/D)} + 1, \qquad (9) \\
&\overset{(c)}{\leq} \frac{2P(3W/D)}{P(W/D)} + 1, \qquad (10)
\end{aligned}$$

where in (a), we have used the fact that $E_y^g > y^g P(3W/D)D/3$ (Remark 13), (b) follows because $P(3W/D)D/3 \geq P(W/D)D$ (transmitting W bits at speed $3W/D$ consumes more energy than transmitting W bits at speed W/D), and we get (c) by maximizing the R.H.S. of (9) with respect to $m \geq 1$. With $\hat{s} = W/D$, we get the result.

VII. Numerical Results

Figure 2 shows the AoI plot for π^g when size of packets is $W = 1$ Mbit, peak AoI is $D = 3$ msec, and inter-generation time X of packets follow uniform distribution with values in interval $(0, 2.5)$. Also, the stem plot in Figure 2 shows the generation time of packets, and the speed with which they were transmitted by π^g. The transmission speed of a packet is 0 if it is not transmitted, otherwise the speed is at least

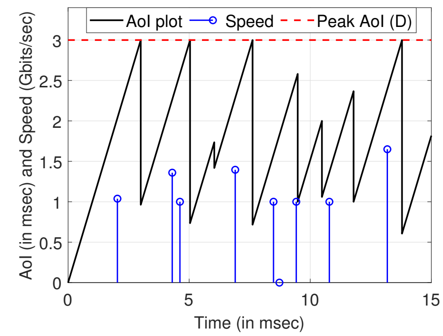

Fig. 2: AoI plot and transmission speed of packets.

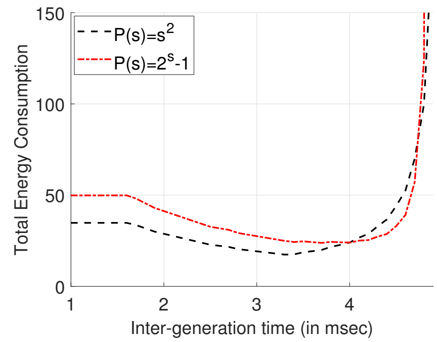

Fig. 3: Energy consumption as a function of inter-generation time.

$3W/D = 1$ Gbit/sec. Note that the speed is greater than 1 Gbit/sec only if inter-generation time of packets is greater than $2D/3 = 2$ msec.

To further understand the effect of inter-generation time of packets on energy consumption, Figure 3 plots the energy consumed by π^g as a function of inter-generation time X (where X is deterministic, packets are of size $W = 1$ Mbit, and peak AoI and time horizon are $D = 5$ msec and $T = 100$ msec respectively). When $X \leq D/3 \approx 1.7$ msec, π^g always has a fresh packet to transmit, and hence, remains busy throughout the interval $[0, T]$ transmitting packets with speed $3W/D = 3/5$ Gbits/sec. So, as long as $X \leq D/3$, energy consumption remains constant. When $X \in (D/3, 2D/3]$, π^g transmits fewer packets, but with same speed $3W/D$, and hence, consumes lesser energy. However, when $X > 2D/3$, π^g transmits fewer packets, but at speed larger than $3W/D$. So, energy consumption starts to increase with increase in X, and becomes unbounded at values of X close to D.

Next, for any policy π (causal or offline), consider the universal lower bound (ULB) (11) on E_π (total energy consumption in interval $[0, T]$ under policy π), derived in [12].

$$E_\pi \geq \max\{0, P(2W/D)(T - D)\}. \tag{11}$$

From (11), it follows that the energy consumption of any causal policy increases with increase in W/D. So, to visualize this numerically, Figure 4 plots the energy consumption of π^g against values of W/D (when initial AoI $\Delta(0) = 0$, time horizon $T = 100$ msec, and inter-generation time of packets is uniformly distributed in interval $[0, 4.5]$ msec.). Note that for varying W/D, we fixed $D = 5$ msec, and varied W.

ISBN 978-3-903176-37-9 © 2021 IFIP

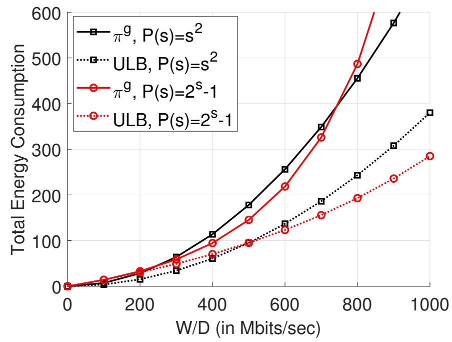

Fig. 4: Dependence of energy consumption of π^g on the ratio W/D, and the universal lower bound (ULB) (11).

Clearly, the energy consumption of π^g increases with increase in W/D, and it grows exponentially when the power function is $P(s) = 2^s - 1$.

VIII. Conclusion

In this paper, we considered a node where packets of fixed size are generated with arbitrary (but bounded) inter-generation time. The node is required to maintain peak age of information (AoI) at the monitor below a given threshold (throughout a given interval of time) by transmitting these packets and controlling their transmission speed, and the objective is to minimize the total energy consumption. We proposed a (customized) greedy policy, and bounded its competitive ratio (CR) by comparing it against an optimal offline policy by deriving some structural results. Importantly, for polynomial power functions, the CR upper bound is independent of the system parameters (such as packet size, peak AoI, time horizon, and number of packets generated). For exponential power functions, we showed that the CR of any causal policy grows exponentially with increase in the ratio of the packet size and the peak AoI, and showed that the proposed greedy policy has CR of similar order as a lower bound on the CR of all causal policies.

References

[1] Y. Xiao and Z. Fan, "10 technology trends to watch in the covid-19 pandemic," https://www.weforum.org/agenda/2020/04/10-technology-trends-coronavirus-covid19-pandemic-robotics-telehealth/, April 27 2020.

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

[3] R. V. Bhat, R. Vaze, and M. Motani, "Throughput maximization with an average age of information constraint in fading channels," *IEEE Transactions on Wireless Communications*, vol. 20, no. 1, pp. 481–494, 2020.

[4] A. Arafa and S. Ulukus, "Age minimization in energy harvesting communications: Energy-controlled delays," in *2017 51st Asilomar Conference on Signals, Systems, and Computers*. IEEE, 2017, pp. 1801–1805.

[5] M. Bastopcu and S. Ulukus, "Age of information for updates with distortion: Constant and age-dependent distortion constraints," *arXiv preprint arXiv:1912.13493*, 2019.

[6] ——, "Partial updates: Losing information for freshness," *arXiv preprint arXiv:2001.11014*, 2020.

[7] N. Rajaraman, R. Vaze, and G. Reddy, "Not just age but age and quality of information," *IEEE Journal on Selected Areas in Communications*, vol. 39, no. 5, pp. 1325–1338, 2021.

[8] F. Yao, A. Demers, and S. Shenker, "A scheduling model for reduced cpu energy," in *Proceedings of IEEE 36th annual foundations of computer science*. IEEE, 1995, pp. 374–382.

[9] N. Bansal, T. Kimbrel, and K. Pruhs, "Speed scaling to manage energy and temperature," *Journal of the ACM (JACM)*, vol. 54, no. 1, pp. 1–39, 2007.

[10] A. Deshmukh and R. Vaze, "Online energy-efficient packet scheduling for a common deadline with and without energy harvesting," *IEEE Journal on Selected Areas in Communications*, vol. 34, no. 12, pp. 3661–3674, 2016.

[11] H. Rahman, A. Vinayachandran, S. R. B. Pillai, K. Appaiah, R. Vaze, and N. Kashyap, "Deadline constrained packet scheduling in the presence of an energy harvesting jammer," *IEEE Transactions on Green Communications and Networking*, 2020.

[12] K. Saurav and R. Vaze, "Online energy minimization under a peak age of information constraint," *arXiv preprint arXiv:2103.16425*, 2021.

[13] V. Kavitha, E. Altman, and I. Saha, "Controlling packet drops to improve freshness of information," *arXiv preprint arXiv:1807.09325*, 2018.

Appendix A
Proof of Proposition 2

Consider time t, when π^g begins to transmit a packet j with speed $s^g(t) > 3W/D$. To prove Proposition 2, we need to show that no packet must have been generated in interval $[t - 2W/D, t)$. For this, let at least one packet be generated in the interval $[t - 2D/3, t)$. Then we must have one of the following two cases:

a) There exists a packet $i \neq j$ generated in interval $[t - 2D/3, t)$ that π^g begins to transmit at time $r_i < t$: Recall that π^g is a non-interruptive policy (Remark 8). Hence, π^g must have delivered packet i (to the monitor) until time t. Also, by hypothesis, π^g begins to transmit packet j at time t with speed $s^g(t) > 3W/D$. Thus, from (5), it follows that $d(t) - t < D/3$. This implies that $d(t) < t + D/3$, which is possible only if no packet (including packet i) that was generated in interval $[t - 2D/3, t)$ was delivered to the monitor until time t. This contradicts the fact that packet i was generated in interval $[t - 2D/3, t)$, and delivered until time t. Hence, there cannot be a packet generated in interval $[t - 2D/3, t)$ that is transmitted by π^g starting at time $r_i < t$.

b) No packet generated in interval $[t - 2D/3, t)$ is transmitted by π^g starting at any time $\hat{t} < t$: Since π^g never idles if a fresh packet is available, this case is possible only if π^g remains busy in the entire interval $[t - 2D/3, t)$, transmitting packets that were generated before time $t - 2D/3$. However, such an event cannot happen. To show this, let packet i be the latest packet generated before time $t - 2D/3$ (say, at time $t_i = t - 2D/3 - \delta$, where $\delta > 0$). Because π^g only transmits a fresh packet, in interval $[t - 2D/3, t)$, π^g may transmit at most two packets generated before $t - 2D/3$: (i) a packet $\hat{i}$ that was being transmitted by π^g when packet i was generated at time t_i, and (ii) packet i itself. Since π^g transmits a packet with speed at least $3W/D$ (Eq.(5)), it takes at most $D/3$ time units to completely transmit (deliver) a packet. Therefore, π^g would finish transmitting both packet $\hat{i}$ and packet i before $t_i + 2D/3 = t - \delta$, and hence, must begin to transmit a packet generated in interval $[t - 2D/3, t)$ in sub-interval $[t - \delta, t)$, thus proving that this case is not possible.

ISBN 978-3-903176-37-9 © 2021 IFIP

Thus, we conclude that none of the above two cases are possible. This contradicts the assumption that a packet was generated in interval $[t-2D/3,t)$.

APPENDIX B
PROOF OF PROPOSITION 3

a) Proof of Property 1: Let $\pi^\star$ transmits the W bits of a packet i with speed that varies with time. Now, consider another policy π', identical to $\pi^\star$, except that it (π') transmits the W bits of packet i with constant speed over the same time-interval where $\pi^\star$ transmits packet i. But, due to convexity of power function $P(\cdot)$ (Lemma 1), we know that over any given interval of time, transmitting a packet with constant speed consumes minimum energy. Therefore, π' consumes less energy than $\pi^\star$. But this cannot be true, because $\pi^\star$ is an optimal offline policy. Hence, $\pi^\star$ must transmit the W bits of each packet i with constant speed.

b) Proof of Property 2: Since $\pi^\star$ knows the generation time of all the packets in advance, it never transmits any packet partially (because transmitting a packet partially consumes energy, without meeting the deadline constraint (3)). Also, $\pi^\star$ only transmits fresh packets (Remark 7). Therefore, it never interrupts a packet's transmission to transmit it later. Hence, $\pi^\star$ never interrupts an ongoing packet transmission.

c) Proof of Property 3: Note that when $\pi^\star$ begins to transmit packet $i+1$ at time r_{i+1}, the deadline is d_i (the latest delivered packet at r_{i+1} is packet i). Therefore, $\tau_{i+1} < d_i$, where τ_{i+1} is the time when packet $i+1$ is delivered. Therefore, packet $i+1$ is delivered in one of the frames 0 to i. Hence, for $i=0$, the only possibility is that packet 1 is delivered in frame 0. Next, using induction, we show that for all i, π^g delivers packet $i+1$ in frame i. Let packet i is delivered in frame $i-1$. Then, there are two possible cases: (i) packet $i+1$ is delivered in frame i, and (ii) packet $i+1$ is delivered in frame $i-1$ (since packet i is delivered in frame $i-1$, packet $i+1$ cannot have been delivered in a frame previous to frame $i-1$). Note that if packet $i+1$ is delivered in frame $i-1$, then there would be two packets (packet i and packet $i+1$) that are delivered in frame $i-1$, at time τ_i and τ_{i+1} respectively. However, in this case, transmission of packet i will be redundant because at the start of frame $i-1$ (i.e., time d_{i-2}), the deadline is d_{i-1}, and due to Proposition 1, we know that delivery of a single packet is sufficient in interval $[d_{i-2}, d_{i-1})$. Hence, $\pi^\star$ being an optimal offline policy, will not waste energy delivering both packets i and $i+1$ in frame $i-1$. Thus, for all i, $\pi^\star$ delivers packet $i+1$ in frame i.

d) Proof of Property 4: From Property 3, we know that exactly one packet (packet $i+1$) is delivered in frame i. Also, packet $i+1$ is transmitted with constant speed (Property 1). Therefore, transmission speed may decrease only after packet $i+1$ is delivered. However, $\pi^\star$ being an optimal offline policy, instead of delivering packet $i+1$ before the deadline d_i (end of frame i), and decreasing the speed, it ($\pi^\star$) would transmit packet $i+1$ itself at lesser speed, over a larger interval. This follows due to convexity of power function $P(\cdot)$ (Lemma 1).

APPENDIX C

In this section, let the packets be indexed in the order they are generated, irrespective of whether they are transmitted by $\pi^\star$ or not. Also, in this section, let the deadline $d(t)$ at any time t be defined according to π^g.

Let π^g transmits a packet j with constant speed greater than $3W/D$. Proposition 4, shows that π^g must have begun to transmit packet j at time t_j (where t_j is the generation time of packet j), and got the transmission completed at time $d(t_j)$.

Proposition 4: If π^g transmits a packet j with speed greater than $3W/D$, then it must have begun to transmit packet j immediately after it was generated at time t_j, and completed the transmission at time $d(t_j)$.

Proof: Let transmission of packet j begins at time t. So, $t_j \le t$, where t_j is the generation time of packet j. From Proposition 2, we know that no packet is generated in interval $[t-2D/3,t)$. Also, it has been shown in proof of Proposition 2 that the transmission of packets generated before $t-2D/3$ (that are transmitted by π^g) gets completed before time t, and hence, t_j cannot be less than $t-2D/3$ (because transmission of packet j begins at time t). Hence, we must have $t_j = t$.

Further, π^g transmits packet j with constant speed $W/(d(t)-t) = W/(d(t_j)-t_j)$ starting time $t=t_j$. So, transmission of packet j completes at time $d(t_j)$. ■

Therefore, in interval $[t_j, d(t_j))$, π^g transmits packet j with constant speed greater than $3W/D$, which is precisely equal to $W/(d(t_j)-t_j)$. Lemma 4 shows that the optimal offline policy $\pi^\star$ also transmits an entire packet (of size W bits) in interval $[t_j, d(t_j))$, with constant speed at least equal to $W/(d(t_j)-t_j)$.

Lemma 4: For each packet j (W bits) transmitted by π^g in interval $[t_j, d(t_j))$ with constant speed $s^g(t) = W/(d(t_j)-t_j) > 3W/D$, there exists a packet $\hat{j}$ such that $\pi^\star$ transmits the W bits of packet $\hat{j}$ in interval $[t_j, d(t_j))$, with constant speed at least equal to $W/(d(t_j)-t_j)$.

Proof: Let π^g transmits a packet j with speed greater than $3W/D$. From Proposition 2, it follows that no packet is generated in interval $[t_j-2D/3, t_j)$, where t_j is the generation time of packet j. Without loss of generality, let packet ℓ be the latest packet that is generated before time $t_j-2D/3$. Since no packet is generated in interval (t_ℓ, t_j), packet ℓ remains fresh until time t_j. Also, note that π^g takes at most $D/3$ time units to deliver W bits (because $s^g(t) \ge 3W/D$), and $t_j - t_\ell > 2D/3$. Therefore, even if π^g was transmitting a previous packet when packet ℓ was generated at time t_ℓ, it (π^g) delivers packet ℓ before time t_j. So, the deadline at time t_j for π^g is $d(t_j) = d_\ell$.

Since no packet is generated in interval (t_ℓ, t_j), the generation time of latest packet delivered by $\pi^\star$ until time t_j can at most be t_ℓ. Hence, the deadline for $\pi^\star$ at time t_j is at most equal to d_ℓ. Due to Proposition 1, $\pi^\star$ must deliver a packet (i.e., transmit W bits) in interval $[t_j, d_\ell)$. Since $d_\ell = d(t_j)$, this implies that $\pi^\star$ must transmit at least W bits of some packet $\hat{j}$ in interval $[t_j, d(t_j))$. From Proposition 3 (Property 1), we know that that $\pi^\star$ transmits the W bits of a packet with constant speed. Therefore, $\pi^\star$ transmits packet $\hat{j}$ within some sub-interval of interval $[t_j, d(t_j))$, with constant speed at least equal to $W/(d(t_j)-t_j)$. ■

ISBN 978-3-903176-37-9 © 2021 IFIP