Data Freshness Oriented Sampling under Unknown Delay Statistics

an online learning approach

Haoyue Tang
Yale University

joint work with Yuchao Chen[1], Yin Sun[2], Leandros Tassiulas[3], Jintao Wang[1] and Pengkun Yang[1]
[1]Tsinghua University, [2]Auburn University, [3]Yale University

ITA Graduation Day, 2022

Table of contents

Background

Real-time service requires fresh data

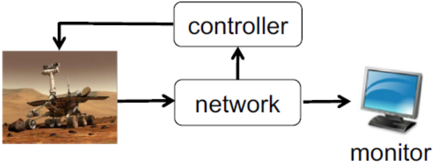

Background

Real-time service requires fresh data

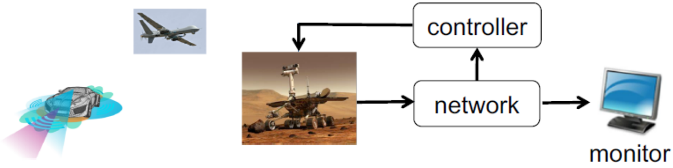

Real-time data analytics becomes important

Crowdsourcing

Online Learning
(ads bidding)

Social networks

Background

Real-time service requires fresh data

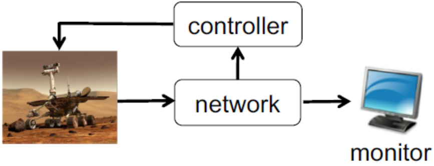

Real-time data analytics becomes important

Crowdsourcing

Online Learning
(ads bidding)

Social networks

Research Problem

How to measure and optimize information freshness?

Freshness Metric: Age of Information

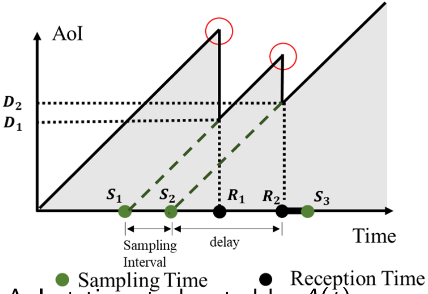

- By definition, the AoI at time t, denoted by $A(t)$

$$A(t) \triangleq t - S_{i(t)}, \tag{1}$$

where $i(t) := \arg\max\{i|R_i \leq t\}$ is the index of the recently received sample.

Freshness Metric: Age of Information

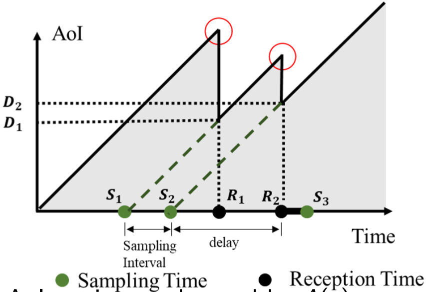

- By definition, the AoI at time t, denoted by $A(t)$

$$A(t) \triangleq t - S_{i(t)}, \tag{1}$$

where $i(t) := \arg\max\{i | R_i \leq t\}$ is the index of the recently received sample.
- No sample received: grows linearly
- New sample received: drops to packet delay

Freshness Metric: Age of Information

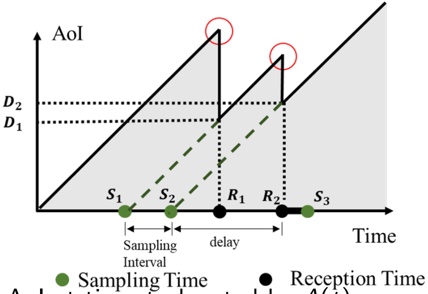

- By definition, the AoI at time t, denoted by $A(t)$

$$A(t) \triangleq t - S_{i(t)}, \tag{1}$$

where $i(t) := \arg\max\{i | R_i \leq t\}$ is the index of the recently received sample.

- No sample received: grows linearly
- New sample received: drops to packet delay

Previous Work

Minimizing AoI is different from max throughput/min delay

Challenges–Online Design

Existing work

- Unknown AoI penalty functions/AoI constraint in **known** environment Li (2021); Tripathi and Modiano (2021)
- min AoI in **unknown** environment: using RL/Bandits (Atay et al. (2021); Leng and Yener (2021), ...) Theoretic analysis are missing.

Challenges–Online Design

Existing work

- Unknown AoI penalty functions/AoI constraint in **known** environment Li (2021); Tripathi and Modiano (2021)
- min AoI in **unknown** environment: using RL/Bandits (Atay et al. (2021); Leng and Yener (2021), ...) Theoretic analysis are missing.

Age minimum sampling revisited–online algorithm

- Reformulate AoI minimization problem as a Renewal-Reward Process, then propose an online algorithm.
- Derive the convergence rate of the proposed online algorithm.
- Establish converse result for any causal sampling algorithm. Verify minimax order optimal.

System Model

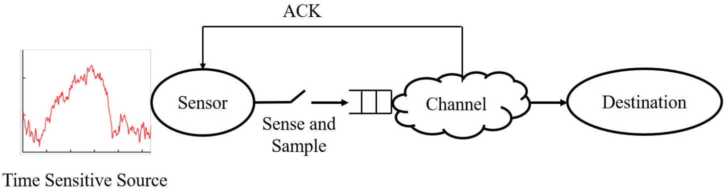

- Point-to-point link: **Sensor** senses and submits samples to the **Destination**
- **Channel**: FIFO queue with i.i.d. transmission times
- **Feedback**: zero-delay ACK
  - Busy/Idle state of the channel is known to the sensor.

Delay D_k of sample k in the channel is i.i.d following an absolutely distribution $\mathbb{P}_D$.

Problem Formulation

Minimizing the AoI under a sampling frequency constraint

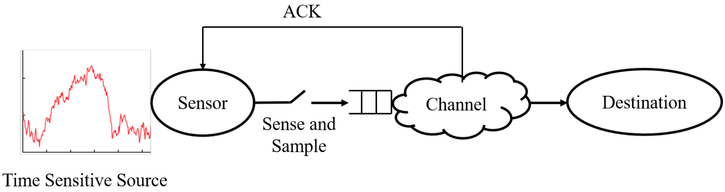

Optimization Problem

$$\inf_{\pi \in \Pi} \limsup_{T \to \infty} \frac{1}{T} \mathbb{E}\left[\int_{t=0}^{T} A(t) \mathrm{d}t\right], \tag{2a}$$

$$\text{s.t.} \quad \liminf_{K \to \infty} \frac{1}{K} \mathbb{E}\left[\sum_{k=1}^{K} (S_{k+1} - S_k)\right] \geq \frac{1}{f_{\max}}. \tag{2b}$$

Observation: Samples waiting in the queues are not longer fresh.
Solution: Focus on policy that waits for $W_k \geq 0$ to submit sample $(k+1)$ after the ACK of the k-th sample is received.

Problem Resolution

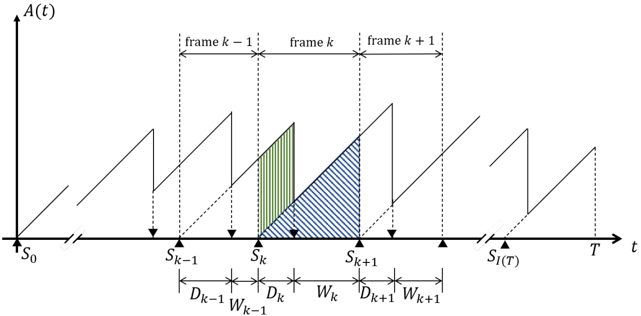

- Frame length k: $L_k = D_k + W_k$.
- Cumulative Age in frame k: $R_k = \frac{1}{2}(D_k + W_k)^2 + D_k(D_{k-1} + W_{k-1})$.

Problem Resolution

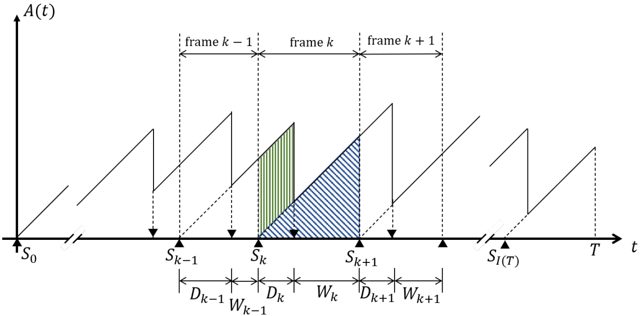

- Frame length k: $L_k = D_k + W_k$.
- Cumulative Age in frame k: $R_k = \frac{1}{2}(D_k + W_k)^2 + D_k(D_{k-1} + W_{k-1})$.

Delay D_k is i.i.d, for stationary policy π:

$$\overline{A}_\pi = \limsup_{K\to\infty} \frac{\mathbb{E}\left[\sum_{k=1}^K R_k\right]}{\mathbb{E}\left[\sum_{k=1}^K L_k\right]} = \limsup_{K\to\infty} \frac{\mathbb{E}\left[\sum_{k=1}^K \frac{1}{2}(D_k + W_k)^2\right]}{\mathbb{E}\left[\sum_{k=1}^K L_k\right]} + \overline{D}. \tag{3}$$

Problem Resolution

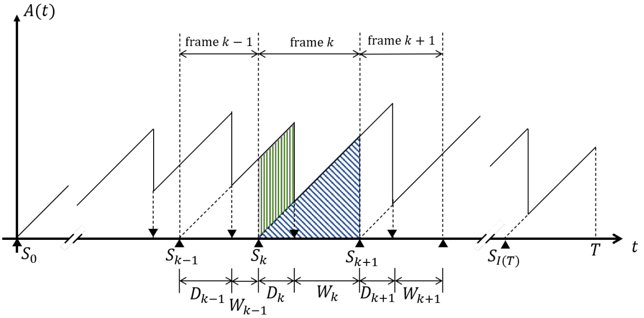

- Frame length k: $L_k = D_k + W_k$.
- Cumulative Age in frame k: $R_k = \frac{1}{2}(D_k + W_k)^2 + D_k(D_{k-1} + W_{k-1})$.

Delay D_k is i.i.d, for stationary policy π:

$$\overline{A}_\pi = \limsup_{K\to\infty} \frac{\mathbb{E}\left[\sum_{k=1}^K R_k\right]}{\mathbb{E}\left[\sum_{k=1}^K L_k\right]} = \limsup_{K\to\infty} \frac{\mathbb{E}\left[\sum_{k=1}^K \frac{1}{2}(D_k + W_k)^2\right]}{\mathbb{E}\left[\sum_{k=1}^K L_k\right]} + \overline{D}. \quad (3)$$

$Q_k := \frac{1}{2}(W_k + D_k)^2$ and L_k are i.i.d $\Rightarrow$ **Renewal-Reward Process Optimization**

Algorithm Design (1): Offline Policy

Assuming $\mathbb{P}_D$ is known, computing $\pi^\star$:

$$\gamma^\star := \inf_{\pi} \limsup_{K\to\infty} \frac{\mathbb{E}\left[\sum_{k=1}^{K} \frac{1}{2}(D_k + W_k)^2\right]}{\mathbb{E}\left[\sum_{k=1}^{K}(D_k + W_k)\right]}, \text{s.t., } \mathbb{E}\left[\frac{1}{K}\sum_{k=1}^{K}(D_k + W_k)\right] \geq \frac{1}{f_{\max}}.$$

Algorithm Design (1): Offline Policy

Assuming $\mathbb{P}_D$ is known, computing $\pi^\star$:

$$\gamma^\star := \inf_{\pi} \limsup_{K\to\infty} \frac{\mathbb{E}\left[\sum_{k=1}^{K} \frac{1}{2}(D_k + W_k)^2\right]}{\mathbb{E}\left[\sum_{k=1}^{K}(D_k + W_k)\right]}, \text{s.t., } \mathbb{E}\left[\frac{1}{K}\sum_{k=1}^{K}(D_k + W_k)\right] \geq \frac{1}{f_{\max}}.$$

Step 1: Fractional Programming Reformulation (Assume $\gamma^\star$ is known)

$$\pi^\star = \arg\min_{\pi} \left(\underbrace{\frac{1}{2}(D + \pi(D))^2}_{\text{Reward}} - \gamma^\star \underbrace{(D + \pi(D))}_{\text{Length}} \right), \text{s.t.} \mathbb{E}[D + \pi(D)] \geq \frac{1}{f_{\max}}. \quad (4)$$

Algorithm Design (1): Offline Policy

Assuming $\mathbb{P}_D$ is known, computing $\pi^\star$:

$$\gamma^\star := \inf_{\pi} \limsup_{K\to\infty} \frac{\mathbb{E}\left[\sum_{k=1}^{K} \frac{1}{2}(D_k + W_k)^2\right]}{\mathbb{E}\left[\sum_{k=1}^{K}(D_k + W_k)\right]}, \text{s.t., } \mathbb{E}\left[\frac{1}{K}\sum_{k=1}^{K}(D_k + W_k)\right] \geq \frac{1}{f_{\max}}.$$

Step 1: Fractional Programming Reformulation (Assume $\gamma^\star$ is known)

$$\pi^\star = \arg\min_{\pi} \left(\underbrace{\frac{1}{2}(D + \pi(D))^2}_{\text{Reward}} - \gamma^\star \underbrace{(D + \pi(D))}_{\text{Length}} \right), \text{s.t.} \mathbb{E}[D + \pi(D)] \geq \frac{1}{f_{\max}}. \quad (4)$$

Step 2: Lagrange Reformulation (for a given γ)

$$\mathcal{L}(\pi, \gamma, \nu) := \frac{1}{2}(D + \pi(D))^2 - \gamma(D + \pi(D)) + \underbrace{\nu\left(\frac{1}{f_{\max}} - D - \pi(D)\right)}_{\text{freq cons}, \nu \geq 0} \quad (5)$$

Algorithm Design (1): Offline Policy

Assuming $\mathbb{P}_D$ is known, computing $\pi^\star$:

$$\gamma^\star := \inf_{\pi} \limsup_{K\to\infty} \frac{\mathbb{E}\left[\sum_{k=1}^{K}\frac{1}{2}(D_k+W_k)^2\right]}{\mathbb{E}\left[\sum_{k=1}^{K}(D_k+W_k)\right]}, \text{s.t., } \mathbb{E}\left[\frac{1}{K}\sum_{k=1}^{K}(D_k+W_k)\right] \geq \frac{1}{f_{\max}}.$$

Step 1: Fractional Programming Reformulation (Assume $\gamma^\star$ is known)

$$\pi^\star = \arg\min_{\pi}\left(\underbrace{\frac{1}{2}(D+\pi(D))^2}_{\text{Reward}} - \gamma^\star \underbrace{(D+\pi(D))}_{\text{Length}}\right), \text{s.t.} \mathbb{E}[D+\pi(D)] \geq \frac{1}{f_{\max}}. \quad (4)$$

Step 2: Lagrange Reformulation (for a given γ)

$$\mathcal{L}(\pi,\gamma,\nu) := \frac{1}{2}(D+\pi(D))^2 - \gamma(D+\pi(D)) + \underbrace{\nu\left(\frac{1}{f_{\max}} - D - \pi(D)\right)}_{\text{freq cons},\nu\geq 0} \quad (5)$$

Step 3:

$$\pi^\star_{\gamma,\nu}(d) = \min \mathcal{L}(\pi,\gamma,\nu) = (\gamma+\nu-d)^+. \quad (6)$$

Offline Design: Compute $\gamma^\star+\nu^\star$ via Bi-section search Sun et al. (2017)

Algorithm Design (2): Online Policy

How to obtain $\gamma^\star, \nu^\star$ when $\mathbb{P}_D$ is unknown?

- $\gamma^\star = \frac{\mathbb{E}[\frac{1}{2}((\gamma^\star+\nu^\star-D)^+)^2]}{\mathbb{E}[(\gamma^\star+\nu^\star-D)^+]} \Rightarrow \overline{Q}(\gamma^\star) - \gamma^\star \overline{L}(\gamma^\star) = 0$
  Finding root of an equation: Robbins-Monro (SGD) Neely (2021)
- Dual optimizer $\nu^\star$: Virtual queue to satisfy sampling frequency constraint.

Algorithm Design (2): Online Policy

How to obtain $\gamma^\star, \nu^\star$ when $\mathbb{P}_D$ is unknown?

- $\gamma^\star = \frac{\mathbb{E}[\frac{1}{2}((\gamma^\star+\nu^\star-D)^+)^2]}{\mathbb{E}[(\gamma^\star+\nu^\star-D)^+]} \Rightarrow \overline{Q}(\gamma^\star) - \gamma^\star\overline{L}(\gamma^\star) = 0$
  Finding root of an equation: Robbins-Monro (SGD) Neely (2021)
- Dual optimizer $\nu^\star$: Virtual queue to satisfy sampling frequency constraint.

In frame k, observe transmission delay D_k and then:

- **Generate Sample/Batch k:** Wait for $W_k = (\gamma_k + \nu_k - D_k)^+$ to take sample $k+1$. The reward/length of frame k:

$$Q_k = \frac{1}{2}(D_k + W_k)^2, L_k = D_k + W_k. \tag{7}$$

Algorithm Design (2): Online Policy

How to obtain $\gamma^\star, \nu^\star$ when $\mathbb{P}_D$ is unknown?

- $\gamma^\star = \frac{\mathbb{E}[\frac{1}{2}((\gamma^\star+\nu^\star-D)^+)^2]}{\mathbb{E}[(\gamma^\star+\nu^\star-D)^+]} \Rightarrow \overline{Q}(\gamma^\star) - \gamma^\star \overline{L}(\gamma^\star) = 0$
  Finding root of an equation: Robbins-Monro (SGD) Neely (2021)
- Dual optimizer $\nu^\star$: Virtual queue to satisfy sampling frequency constraint.

In frame k, observe transmission delay D_k and then:

- **Generate Sample/Batch k:** Wait for $W_k = (\gamma_k + \nu_k - D_k)^+$ to take sample $k+1$. The reward/length of frame k:

$$Q_k = \frac{1}{2}(D_k + W_k)^2, L_k = D_k + W_k. \tag{7}$$

- **Approximate γ_k via SGD:** $\overline{g}_{\nu_k}(\gamma) := \overline{Q}(\gamma) - \gamma \overline{L}(\gamma)$, using noisy observation $\tilde{g}_{\nu_k}(\gamma_k) = Q_k - \gamma_k L_k$ for gradient descent:

$$\gamma_{k+1} = (\gamma_k + \eta_k \tilde{g}_{\nu_k}(\gamma_k))^+ = (\gamma_k + \eta_k (Q_k - \gamma_k L_k))^+ . \tag{8}$$

($\eta_k = \frac{1}{k\overline{D}_{\text{lb}}}$, diminishing step-sizes in SGD)

Algorithm Design (2): Online Policy

How to obtain $\gamma^\star, \nu^\star$ when $\mathbb{P}_D$ is unknown?

- $\gamma^\star = \frac{\mathbb{E}[\frac{1}{2}((\gamma^\star+\nu^\star-D)^+)^2]}{\mathbb{E}[(\gamma^\star+\nu^\star-D)^+]} \Rightarrow \overline{Q}(\gamma^\star) - \gamma^\star \overline{L}(\gamma^\star) = 0$
  Finding root of an equation: Robbins-Monro (SGD) Neely (2021)
- Dual optimizer $\nu^\star$: Virtual queue to satisfy sampling frequency constraint.

In frame k, observe transmission delay D_k and then:

- **Generate Sample/Batch k:** Wait for $W_k = (\gamma_k + \nu_k - D_k)^+$ to take sample $k+1$. The reward/length of frame k:

$$Q_k = \frac{1}{2}(D_k + W_k)^2, L_k = D_k + W_k. \tag{7}$$

- **Approximate γ_k via SGD:** $\overline{g}_{\nu_k}(\gamma) := \overline{Q}(\gamma) - \gamma \overline{L}(\gamma)$, using noisy observation $\tilde{g}_{\nu_k}(\gamma_k) = Q_k - \gamma_k L_k$ for gradient descent:

$$\gamma_{k+1} = (\gamma_k + \eta_k \tilde{g}_{\nu_k}(\gamma_k))^+ = (\gamma_k + \eta_k(Q_k - \gamma_k L_k))^+ . \tag{8}$$

($\eta_k = \frac{1}{k\overline{D}_{\text{lb}}}$, diminishing step-sizes in SGD)

- **Update dual optimizer ν_k:**

$$\nu_{k+1} = \left(\nu_k + \frac{1}{V}(L_k - \frac{1}{f_{\max}})\right)^+ . \tag{9}$$

Theoretic Analysis (1)

Constraint Satisfaction

U_k records the sampling constraint violation up to the k-th sample. The proposed algorithm satisfies the sampling frequency constraint in the sense that:

$$\liminf_{K\to\infty} \frac{1}{K}\mathbb{E}\left[\sum_{k=1}^{K} U_k\right] < \infty. \tag{10}$$

Proof: Lyapunov Drift Plus Penalty

Theoretic Analysis (2)–Convergence Analysis

Convergence of proposed algorithm

The time-average AoI of the proposed algorithm converges to $\overline{A}_{\pi^\star}$ almost surely,

$$\lim_{K\to\infty} \frac{\int_{t=0}^{S_{K+1}} A(t)\mathrm{d}t}{S_{K+1}} = \overline{A}_{\pi^\star}, \text{w.p.1}, \tag{11}$$

and the cumulative AoI regret up to frame K:

$$\mathbb{E}\left[\int_{t=0}^{S_{K+1}} A(t)\mathrm{d}t\right] - \overline{A}_{\pi^\star}\mathbb{E}[S_{K+1}] = \mathcal{O}\left(\ln K\right). \tag{12}$$

General idea: view $\overline{A}_t$ as an ODE, show its convergence to $\overline{A}_{\pi^\star}$.
Perturbed ODE approach requires perturbations γ_k within a closed set.

$$\gamma_{k+1} = \left(\gamma_k + \eta_k\left(Q_k - \gamma_k L_k\right)\right)^+ \Rightarrow \text{queue} = (\text{queue} + \text{arrival} - \text{departure})^+$$

$\gamma_k \in [0, \infty)$, novel analysis using **Drift Method from Heavy-Traffic Analysis**.

Theoretic Analysis (3)–Converse Result

Converse Result

Let $\mathcal{P}_w$ be the set of delay distribution, if $D \sim \mathbb{P} \in \mathcal{P}_w$, the age optimum sampling policy is not zero-wait, for any online policy π:

$$\inf_{\pi} \sup_{\mathbb{P}\in\mathcal{P}_w} \left(\mathbb{E}\left[\int_{t=0}^{S_{K+1}} A(t)\mathrm{d}t \right] - \overline{A}_{\pi^\star}\mathbb{E}[S_{K+1}] \right) = \Omega\left(\ln K\right). \tag{13}$$

Theoretic Analysis (3)–Converse Result

Converse Result

Let $\mathcal{P}_w$ be the set of delay distribution, if $D \sim \mathbb{P} \in \mathcal{P}_w$, the age optimum sampling policy is not zero-wait, for any online policy π:

$$\inf_{\pi} \sup_{\mathbb{P}\in\mathcal{P}_w} \left(\mathbb{E}\left[\int_{t=0}^{S_{K+1}} A(t)\mathrm{d}t \right] - \overline{A}_{\pi^\star}\mathbb{E}[S_{K+1}] \right) = \Omega\left(\ln K\right). \tag{13}$$

- Step 1: AoI regret lower bound⇒Estimation lower bound of $\mathbb{E}[(\hat{\gamma} - \gamma^\star)^2]$ given k samples.

Theoretic Analysis (3)–Converse Result

Converse Result

Let $\mathcal{P}_w$ be the set of delay distribution, if $D \sim \mathbb{P} \in \mathcal{P}_w$, the age optimum sampling policy is not zero-wait, for any online policy π:

$$\inf_{\pi} \sup_{\mathbb{P}\in\mathcal{P}_w} \left(\mathbb{E}\left[\int_{t=0}^{S_{K+1}} A(t)\mathrm{d}t \right] - \overline{A}_{\pi^\star}\mathbb{E}[S_{K+1}] \right) = \Omega\left(\ln K\right). \quad (13)$$

- Step 1: AoI regret lower bound⇒Estimation lower bound of $\mathbb{E}[(\hat{\gamma}-\gamma^\star)^2]$ given k samples.
- Step 2: Le Cam's two point method: for any distribution $\mathbb{P}_1$ and $\mathbb{P}_2$

$$\inf_{\hat{\gamma}} \sup_{\mathbb{P}} \mathbb{E}\left[(\hat{\gamma}-\gamma)^2\right] \geq (\gamma_1^\star - \gamma_2^\star)^2 \cdot \left(\int \min\{\mathbb{P}_1^{\otimes k}(\mathrm{d}x), \mathbb{P}_2^{\otimes k}(\mathrm{d}x)\} \right). \quad (14)$$

Construct two probabilities $\mathbb{P}_1$ and $\mathbb{P}_2$ that $(\gamma_1^\star - \gamma_2^\star)^2$ is large but $D_{\mathrm{KL}}(\mathbb{P}_1||\mathbb{P}_2)$ is small.

Theoretic Analysis (3)–Converse Result

Converse Result

Let $\mathcal{P}_w$ be the set of delay distribution, if $D \sim \mathbb{P} \in \mathcal{P}_w$, the age optimum sampling policy is not zero-wait, for any online policy π:

$$\inf_{\pi} \sup_{\mathbb{P}\in\mathcal{P}_w} \left(\mathbb{E}\left[\int_{t=0}^{S_{K+1}} A(t)\mathrm{d}t\right] - \overline{A}_{\pi^\star}\mathbb{E}[S_{K+1}]\right) = \Omega\left(\ln K\right). \tag{13}$$

- Step 1: AoI regret lower bound⇒Estimation lower bound of $\mathbb{E}[(\hat{\gamma}-\gamma^\star)^2]$ given k samples.
- Step 2: Le Cam's two point method: for any distribution $\mathbb{P}_1$ and $\mathbb{P}_2$

$$\inf_{\hat{\gamma}} \sup_{\mathbb{P}} \mathbb{E}\left[(\hat{\gamma}-\gamma)^2\right] \geq (\gamma_1^\star - \gamma_2^\star)^2 \cdot \left(\int \min\{\mathbb{P}_1^{\otimes k}(\mathrm{d}x), \mathbb{P}_2^{\otimes k}(\mathrm{d}x)\}\right). \tag{14}$$

  Construct two probabilities $\mathbb{P}_1$ and $\mathbb{P}_2$ that $(\gamma_1^\star - \gamma_2^\star)^2$ is large but $D_{\mathrm{KL}}(\mathbb{P}_1||\mathbb{P}_2)$ is small.
- Step 3: Estimate $\gamma^\star : g(\gamma^\star) = 0 \Rightarrow$ Non-Parametric Estimation of Hölder smooth function $g(\gamma)$ at a given point γ.

Achievability $\mathcal{O}(\ln K)$, minimax order optimal.

Simulations (1)

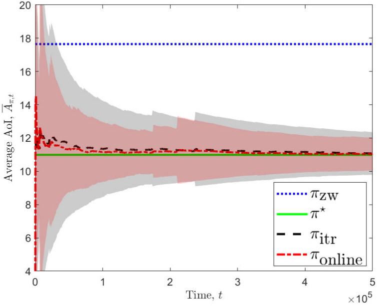

Figure: AoI evolution with time (red denotes the confidence interval).

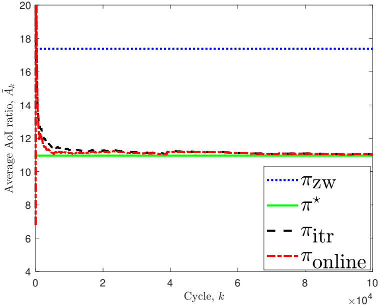

Figure: AoI evolution with frame $\mathbb{E}[\int_{t=0}^{S_{K+1}} A(t)\mathrm{d}t]/\mathbb{E}[S_{K+1}]$

The proposed algorithm learns the AoI minimum sampling policy adaptively when $T \to \infty$, the learning rate is faster than previous method.

Simulations (2)

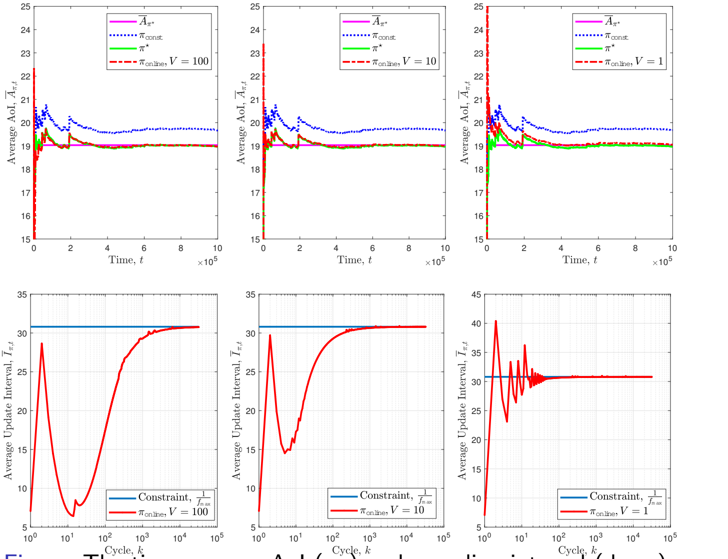

Figure: The time average AoI (up) and sampling interval (down).

π_{online} can satisfy the sampling frequency constraint, smaller V obeys better.

Challenges–Metric Constraint

AoI: **Time difference** between data generation and data usage.

Constraints

- The source sometimes changes fast, sometimes changes slow.

If prior knowledge on how the source evolves can be obtained, AoI may **not** be a good freshness metric.

Challenges–Metric Constraint

AoI: **Time difference** between data generation and data usage.

Constraints

- The source sometimes changes fast, sometimes changes slow.

If prior knowledge on how the source evolves can be obtained, AoI may **not** be a good freshness metric.

Addressing the change of the source

Online Sampling of a Wiener Source.

System Model

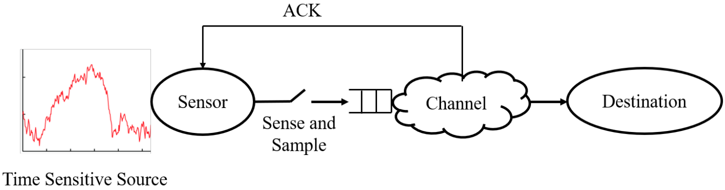

- Point-to-point link: **Sensor** senses and submits samples to the **Destination**
- **Channel**: FIFO queue with i.i.d. transmission times
- **Feedback**: zero-delay ACK
  - Busy/Idle state of the channel is known to the sensor.

System Model

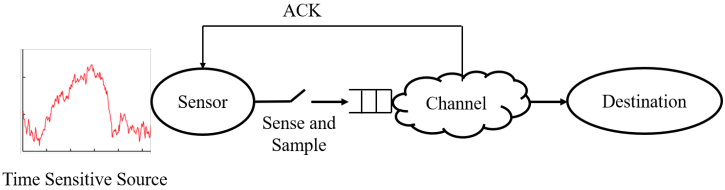

- Point-to-point link: **Sensor** senses and submits samples to the **Destination**
- **Channel**: FIFO queue with i.i.d. transmission times
- **Feedback**: zero-delay ACK
  - Busy/Idle state of the channel is known to the sensor.

Similar to the previous one, except we have some prior knownledge on source evolution:

The source X_t is now a Wiener process.

Problem Formulation

Design sampling strategy to minimize MSE

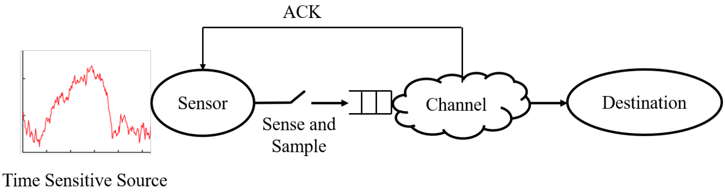

Optimization Problem

$$\inf_{\pi \in \Pi} \limsup_{T \to \infty} \frac{1}{T} \mathbb{E}\left[\int_{t=0}^{T} (X_t - \hat{X}_t)^2 \mathrm{d}t \right], \quad (15a)$$

$$\text{s.t.} \quad \liminf_{K \to \infty} \frac{1}{K} \mathbb{E}\left[\sum_{k=1}^{K} (S_{k+1} - S_k) \right] \geq \frac{1}{f_{\max}}. \quad (15b)$$

Problem Formulation

Design sampling strategy to minimize MSE

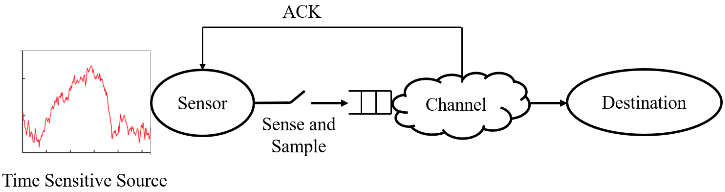

Optimization Problem

$$\inf_{\pi\in\Pi} \limsup_{T\to\infty} \frac{1}{T}\mathbb{E}\left[\int_{t=0}^{T}(X_t-\hat{X}_t)^2\mathrm{d}t\right], \tag{15a}$$

$$\text{s.t.} \quad \liminf_{K\to\infty} \frac{1}{K}\mathbb{E}\left[\sum_{k=1}^{K}(S_{k+1}-S_k)\right] \geq \frac{1}{f_{\max}}. \tag{15b}$$

MMSE Estimator:

$$\hat{X}_t = X_{i(t)}, i(t) := \arg\min_i\{R_i \leq t\}. \tag{16}$$

Similarly, consider policies wait for W_k to take the next sample.

Problem Resolution

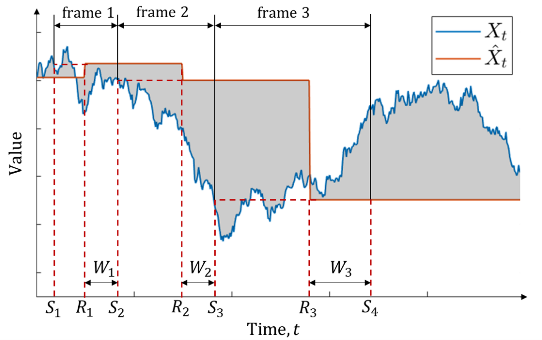

- Frame length k: $L_k = D_k + W_k$.
- Cumulative Estimation Error in frame k: $E_k = \int_{t=S_k}^{S_k+D_k}(X_t - X_{S_{t-1}})^2\mathrm{d}t + \int_{t=S_k+D_k}^{S_{k+1}}(X_t - X_{S_t})^2\mathrm{d}t$.

$$\overline{\mathcal{E}}_\pi = \limsup_{K\to\infty} \frac{\mathbb{E}\left[\sum_{k=1}^K E_k\right]}{\mathbb{E}\left[\sum_{k=1}^K L_k\right]} = \limsup_{K\to\infty} \frac{\mathbb{E}\left[\sum_{k=1}^K \frac{1}{6}(X_{S_{k+1}} - X_{S_k})^4\right]}{\mathbb{E}\left[\sum_{k=1}^K (S_{k+1} - S_k)\right]} + \overline{D}. \qquad (17)$$

Optimal Policy: $S_{k+1} = \inf\{t \geq R_k | |X_t - X_{S_k}| \geq \sqrt{3(\gamma^\star + \nu^\star)}\}, \gamma^\star = \frac{\mathbb{E}[\delta X^4/6]}{\mathbb{E}[L]}$.

Simulations (1)

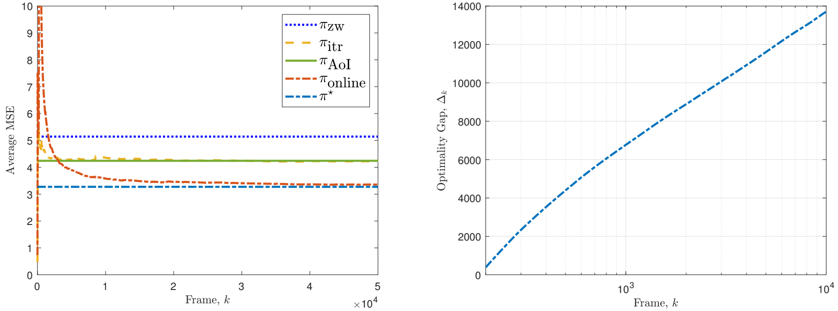

Figure: MSE evolution with frame $\mathbb{E}[\int_{t=0}^{S_{K+1}}(X_t - \hat{X}_t)^2 \mathrm{d}t]/\mathbb{E}[S_{K+1}]$ (left); Regret Growth Rate $\Delta_k := \mathbb{E}\left[\int_0^{S_{K+1}}(\hat{X}_t - X_t)^2 \mathrm{d}t\right] - \overline{\mathcal{E}}_{\pi^\star}\mathbb{E}[S_{k+1}]$. (Right)

Signal-aware optimum sampling is much better than signal-ignorant AoI optimal sampling; Regret growth rate of $\mathcal{O}(\ln k)$ is verified.

Conclusions

- Contribution:
  - The first to use Robbins-Monro to AoI related problems Neely (2021).
  - Develop a new method for proving convergence rate of stochastic approximation algorithm in an **open** set.
  - Converse bound for online algorithm using non-parametric statistics.
- Research Output:
  H. Tang, Y. Chen, J. Wang, J. Sun, J. Song: Sending Timely Status Updates through Channel with Random Delay via Online Learning, Infocom2022
  H. Tang, Y. Chen, J. Wang, P. Yang and L. Tassiulas, "Age Optimal Sampling under Unknown Delay Statistics", submitted to Trans on IT https://arxiv.org/abs/2202.13367
  H. Tang, Y. Sun and L. Tassiulas, "Sampling of the Wiener Process for Remote Estimation over a Channel with Unknown Delay Statistics", submitted to Mobihoc2022

Thank you! Questions?

E. U. Atay, I. Kadota, and E. H. Modiano. Aging wireless bandits: Regret analysis and order-optimal learning algorithm. In *WiOpt*, pages 57–64. IFIP, 2021.

S. Leng and A. Yener. An actor-critic reinforcement learning approach to minimum age of information scheduling in energy harvesting networks. In *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8128–8132, 2021. doi: 10.1109/ICASSP39728.2021.9415110.

B. Li. Efficient learning-based scheduling for information freshness in wireless networks. In *IEEE INFOCOM 2021 - IEEE Conference on Computer Communications*, pages 1–10, 2021. doi: 10.1109/INFOCOM42981.2021.9488709.

M. J. Neely. Fast learning for renewal optimization in online task scheduling. *Journal of Machine Learning Research*, 22(279):1–44, 2021.

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*, 63(11):7492–7508, 2017.

Y. Sun, Y. Polyanskiy, and E. Uysal. Sampling of the wiener process for remote estimation over a channel with random delay. *IEEE Transactions on Information Theory*, 66(2):1118–1135, 2020.

V. Tripathi and E. Modiano. An online learning approach to optimizing time-varying costs of aoi. In *Proceedings of the Twenty-Second International Symposium on Theory, Algorithmic Foundations, and Protocol Design for Mobile Networks and Mobile Computing*, MobiHoc '21, page 241–250, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450385589. doi: 10.1145/3466772.3467053. URL https://doi.org/10.1145/3466772.3467053.