Age-of-Information-Aware Mobile Crowdsensing for Uncertain Event Capture

IEEE INFOCOM 2026

May 21, 2026

Jinrui Zhou[1], Yin Xu[1], Haotian Xu[1], Mingjun Xiao[1], Jie Wu[2]

1 School of Computer Science and Technology & Suzhou Institute for Advanced Study, University of Science and Technology of China

2 Department of Computer and Information Sciences, Temple University, USA

CONTENTS

CONTENTS

Introduction

Mobile CrowdSensing (MCS)

➢ Mobile CrowdSensing (MCS) is a crowdsourcing-based sensing paradigm that a platform can recruit a crowd of mobile users (a.k.a., workers) to collect data from some Points of Interest (PoIs) with carried smart devices.

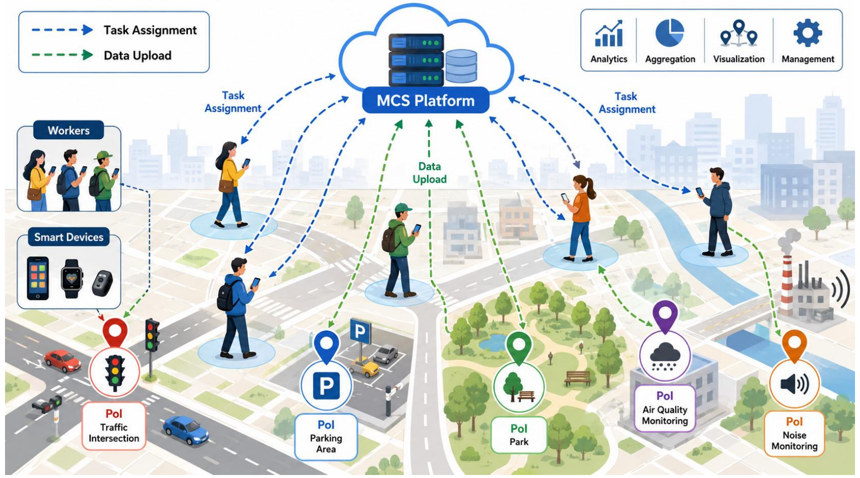

Age of Information (AoI)

➢ Age of Information (AoI), defined as the elapsed time since the latest update or generation of the data collected or transmitted from some PoI, is proposed to measure the freshness of data.

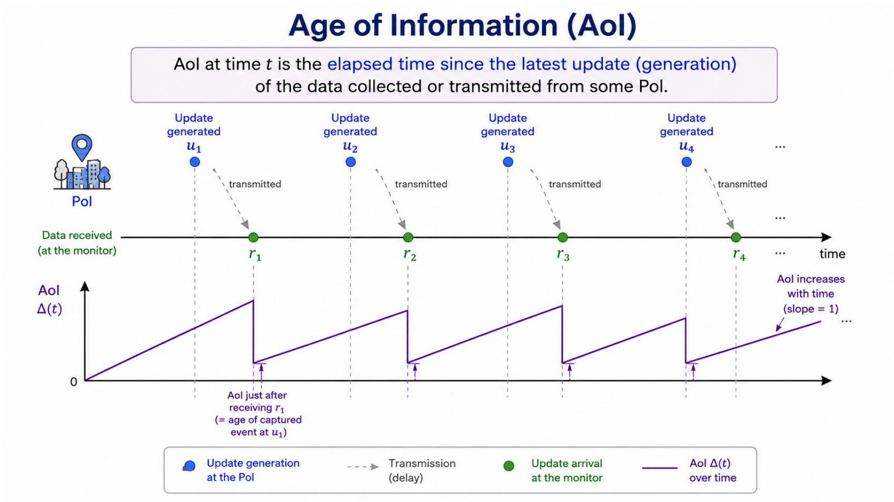

Introduction

Motivation

- ✓ The platform does not know **whether** an event has occurred at a particular location or **when** it will happen. Sending workers blindly (e.g., P3) may waste effort and result in stale data.

- ✓ With a **limited budget**, overspending on distant PoIs (e.g., P2) will deprive other PoIs of adequate exploration and updates, ultimately failing to minimize the average AoI values across the system.

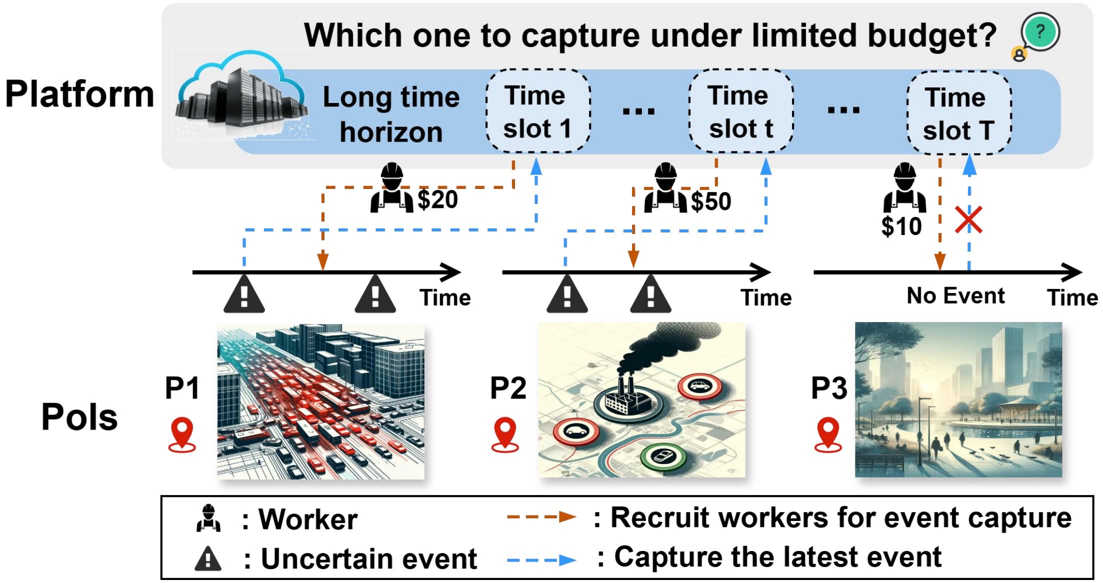

Goal: An efficient event capture strategy under uncertainty and budget limits.

Introduction

Key Challenges

Challenge 1: Dual Uncertainty

- Event generation and capture are probabilistic
- Platform doesn't know if or when an event occurs at a PoI
- AoI **may not reset** even after a capture

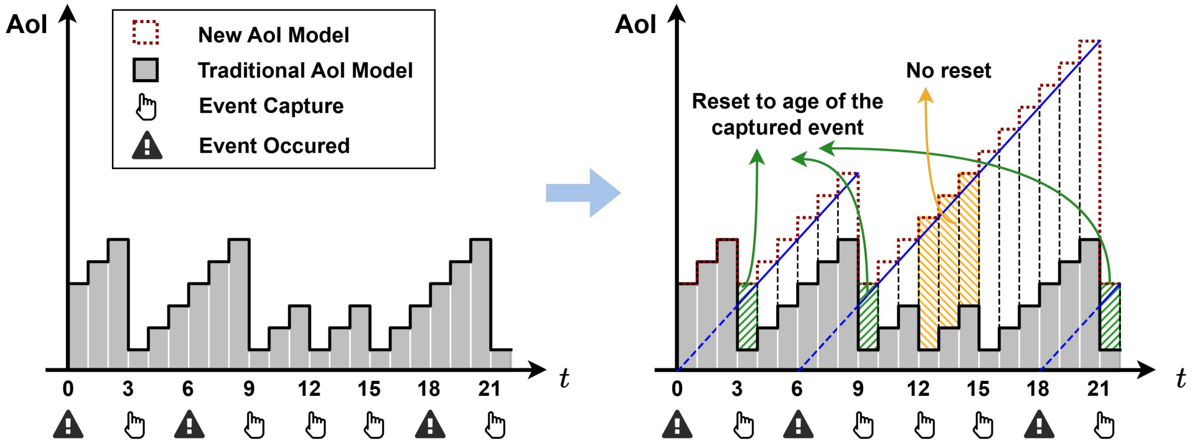

Challenge 2: Temporal Coupling

- Each capture affects current and **future** AoI
- AoIs of other PoIs keep increasing
- Budget must be allocated carefully over time

Deviation between the two AoI models:

- AoI keeps increasing if no updated data is captured (e.g., 3rd & 4th captures)
- AoI resets to the captured event's age, not zero, if data is obtained (e.g., 2nd & 5th captures)

CONTENTS

System, Modeling, and Problem

MCS System and AoI Models

System settings:

- ✓ PoIs: $\mathcal{N} = \{1, \ldots, N\}$
- ✓ Time slots: $\mathcal{T} = \{1, 2, \ldots, T\}$
- ✓ Episode number: m
- ✓ Episode length: L

$$T = mL$$

- ✓ Probabilities: $\theta = (\theta_1, \theta_2, \ldots, \theta_N)$
- ✓ Event capture strategy:

$$\pi(t) = (\pi_1(t), \ldots, \pi_N(t)) \in \{0, 1\}^N$$

- ✓ Cost and budget: c_i, B

AoI of event:

- Definition: the AoI in the i -th PoI at time t : $X_i(t;\theta)$
- Dynamic:

$$Pr[X_i(t+1;\theta) = 0] = \theta_i,$$
$$Pr[X_i(t+1;\theta) = X_i(t;\theta) + 1] = 1 - \theta_i,$$

AoI of event copy:

- Definition: the AoI of event copy in the platform, denoted as $Y_i(t;\theta,\pi)$
- Dynamic:

$$Y_i(t;\theta,\pi) = \begin{cases} X_i(t-1;\theta) + 1 & \text{if } \pi_i(t) = 1 \\ Y_i(t-1;\theta,\pi) + 1 & \text{if } \pi_i(t) = 0 \end{cases}$$

System, Modeling, and Problem

Problem Formulation

Our goal is to determine an event capture strategy that minimizes the cumulative weighted average AoI value under the budget constraint.

$$\textbf{P1:} \quad \min_\pi \quad \frac{1}{T}\sum_{l=1}^{m}\sum_{t=1}^{L}\mathbb{E}\Big(\sum_{i=1}^{N}\gamma_i Y_i(t)\Big),$$

← Expected cumulative weighted average AoI.

$$s.t. \quad \sum_{i=1}^{N}\pi_i(t) = 1, \text{for any } t \in \mathcal{T},$$

← Only one PoI will be selected in each time slot

$$\frac{1}{T}\sum_{t=1}^{T}\sum_{i=1}^{N}\pi_i(t)c_i \leq B,$$

← Budget constraint

$$Eq.\ (2) \sim Eq.\ (4).$$

← AoI updating processes

System, Modeling, and Problem

Problem Transformation

Constrained restless bandit:

- ✓ Learner: platform
- ✓ Arms: PoIs
- ✓ Action: choose a PoI A_t
- ✓ Arm's state: AoI $X_i(t;\theta)$
- ✓ State transition:

  unknown AoI dynamic

- ✓ Penalty: weighted AoI of all events

$$r(t)=\sum_{i=1}^{N}\gamma_i Y_i(t)=\sum_{i=1}^{N}\gamma_i Y_i(t-1)+1+\gamma_{A_t}(X_{A_t}(t)-Y_{A_t}(t-1)-1).$$

- ✓ Cost: $c(t)=\sum_{i=1}^{N}\pi_i(t)c_i.$

Our goal is to select arms sequentially so as to minimize cumulative penalty while respecting the overall budget constraint.

CONTENTS

Algorithm Design and Theoretical Analysis

Belief-DPP Bandit Policy

Assume θ is known
→ handle temporal coupling & budget

Belief state — estimate of current AoI

$$X_i(t;\theta) \quad \rightarrow \quad X_i^b(t)$$

get the approximation of penalty

$$\widetilde{Y}_i(t) = \pi_i(t)X_i^b(t) + (1 - \pi_i(t))(Y_i(t-1) + 1).$$

Virtual queue — track budget usage

$$Z(t) = \max\{0, Z(t-1) + \sum_{i=1}^{N} \pi_i(t-1)c_i - B\}.$$

Belief-DPP bandit policy — balances

immediate AoI reduction vs. budget

$$\pi^*(t) = \arg\min V\sum_{i=1}^{N}\gamma_i\widetilde{Y}_i(t) + Z(t)(\sum_{i=1}^{N}\pi_i(t)c_i - B)$$

Algorithm Design and Theoretical Analysis

Thompson Sampling

Estimate unknown event probabilities θ for each PoI by **posterior distribution**.

Theorem 1. ***(Posterior Update)*** *Given a prior distribution* $Q_i^l(\theta_i) = Beta(\alpha_i^l, \beta_i^l)$ *and a sequence of feedback* $\mathcal{S}_i = \{(t_j^i, s_j^i)\}_{j=0}^{|\Gamma_i|}$. *Then, the posterior distribution is given by* $Q_i^{l+1}(\theta_i|\mathcal{S}_i) = Beta(\alpha_i^{l+1}, \beta_i^{l+1})$, *where*

$$\alpha_i^{l+1} = \alpha_i^l + \sum_{j=1}^{|\Gamma_i|} I_j, \ \beta_i^{l+1} = \beta_i^l + \sum_{j=1}^{|\Gamma_i|} s_j^i - (1 - I_j) s_{j-1}^i. \quad (14)$$

Here, I_j *is an indicator of whether new data is obtained in the* j*-th selection, formally,* $I_j = \mathbb{I}(s_j^i < t_j^i - t_{j-1}^i)$.

Algorithm Design and Theoretical Analysis

Algorithm 1: TS-DPP for Uncertain Event Capture

input : prior Q^0, episode length L;

1 Initialize: posterior $Q^1 = Q^0$;
2 **for** *episodes* $l = 1, \ldots, m$ **do**
3   Initialize: $\mathcal{H}_0 = \emptyset$;
4   **for** $i = 1, \ldots, N$ **do**
5   Draw a parameter $\theta_i^l \sim Q_i^l$;
6   **for** $t = 1, \ldots, L$ **do**
7   Observe $\omega^o(t); Z(t)$;
8   **for** $i = 1, \ldots, N$ **do**
9   Calculate belief state $X_i^b(t)$ according to Eq. (16);
10   Select arm A_t according to Eq. (13);
11   Observe feedback $(A_t, X_{A_t}(t))$;
12   Update virtual queue $Z(t+1)$ according to Eq. (12);
13   Update record M_{t+1};
14   Update history $\mathcal{H}_t = \mathcal{H}_{t-1} \cup (A_t, X_{t,A_t})$;
15   **for** $i = 1, \ldots, N$ **do**
16   Update posterior Q_i^{l+1} according to Eq. (14);

Inner loop (lines 6-14):
Fix the parameter estimation, select the arm according to belief-DPP policy.

Outer loop:
Update the posterior distribution of each arm based on the history we get in the inner loop.

Algorithm Design and Theoretical Analysis

Theoretical Analysis

- TS-DPP analyzed via **Bayesian regret**
- Compared to oracle with known transition probabilities
- **Value function**: expected cumulative AoI penalty over an episode

$$J^{\theta}_{\pi,t}(\mathcal{H}) = \mathbb{E}_{\theta,\pi}\left[\sum_{j=t}^{L} r(j)|\mathcal{H}\right]$$

- Bayesian regret definition:

$$BR(T) = \mathbb{E}_{\theta^*\sim Q}\left[\sum_{l=1}^{m} \mathbb{E}_{\theta^l\sim Q^l} J^{\theta^*}_{\pi^l,1}(\emptyset) - mJ^{\theta^*}_{\pi^*,1}(\emptyset)\right]$$

Theorem 2. ***(Bayesian Regret Bound)*** *The Bayesian regret of TS-DPP algorithm satisfies the following bound:*

$$BR(T) = \mathcal{O}\left(\sqrt{L^3 T \log T}\right). \tag{19}$$

Implication:
sublinear regret → algorithm quickly approaches near-optimal

CONTENTS

Experimental Evaluation

Evaluation Setup

✓ **Datasets:**

Real-world: 3 event types, 1.65M occurrences, 1-second time slots, 100 episodes

Synthetic: N = 4, 8 PoIs, 50 episodes, Bernoulli events

✓ **Baselines:**

- Offline oracle (belief-DPP) → upper bound
- Lyapunov index method
- Whittle index approach

Experimental Evaluation

Comparing to Baselines

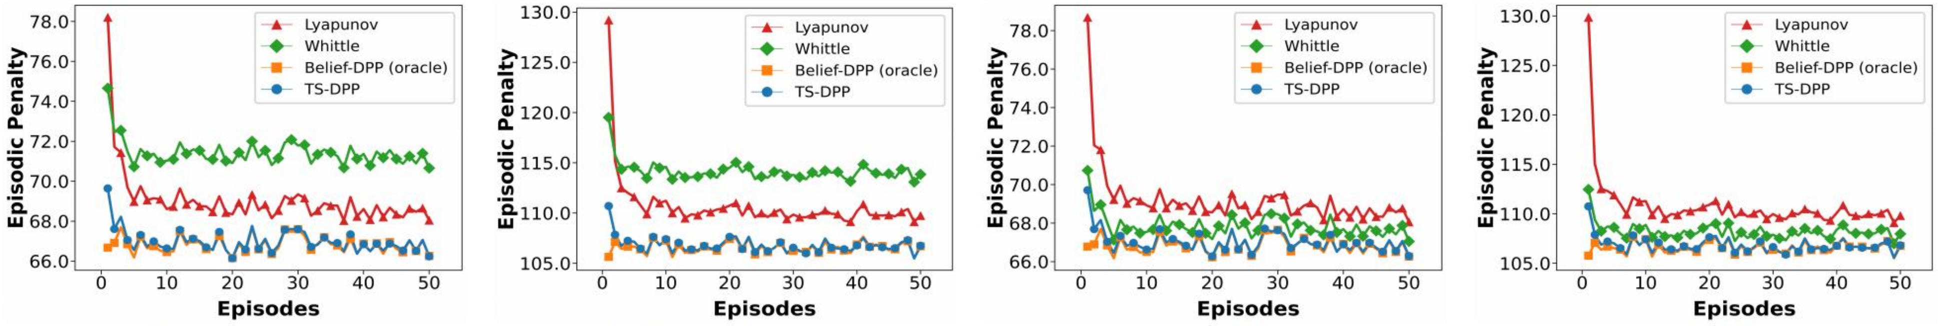

(a) $L = 20, V = 0.2$ (b) $L = 30, V = 0.2$ (c) $L = 20, V = 0.5$ (d) $L = 30, V = 0.5$

Fig. 3. Episodic penalty trends for each algorithm under parameter settings: $N = 4$, $L = \{20, 30\}$, $V = \{0.2, 0.5\}$.

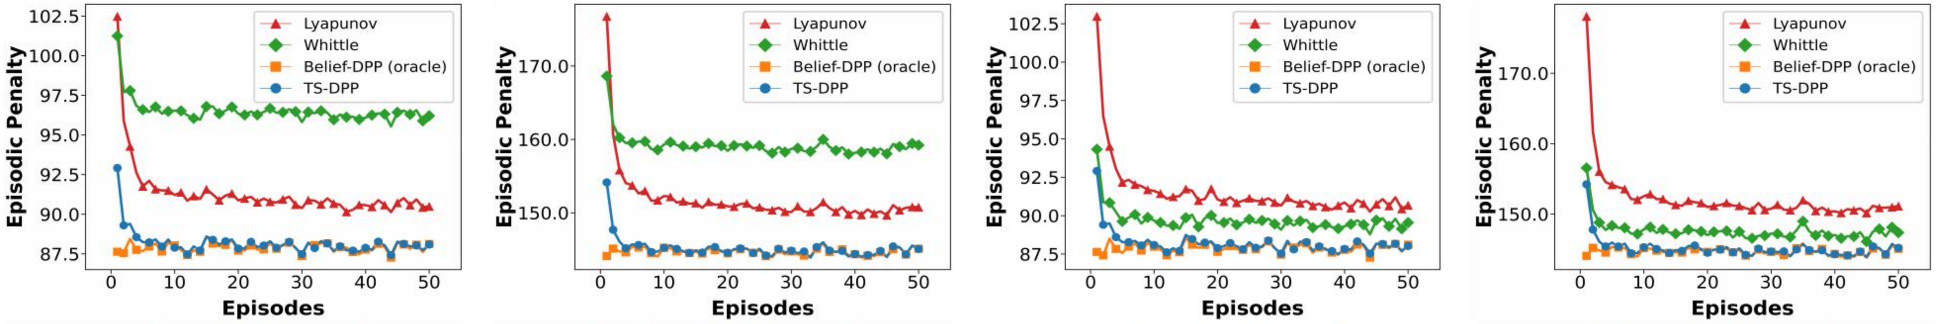

(a) $L = 20, V = 0.2$ (b) $L = 30, V = 0.2$ (c) $L = 20, V = 0.5$ (d) $L = 30, V = 0.5$

Fig. 4. Episodic penalty trends for each algorithm under parameter settings: $N = 8$, $L = \{20, 30\}$, $V = \{0.2, 0.5\}$.

Experimental Evaluation

Bayesian Regret

- TS-DPP achieves time-averaged regret that rapidly converges toward zero, closely following theoretical bounds.

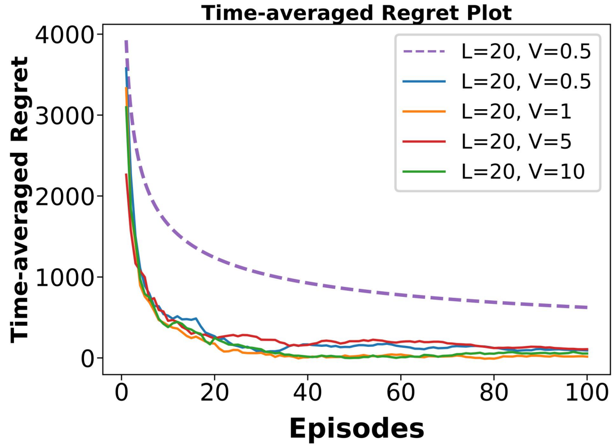

Efficiency of Thompson Sampling

- Thompson Sampling accurately learns event probabilities over episodes, moving from exploration to exploitation.

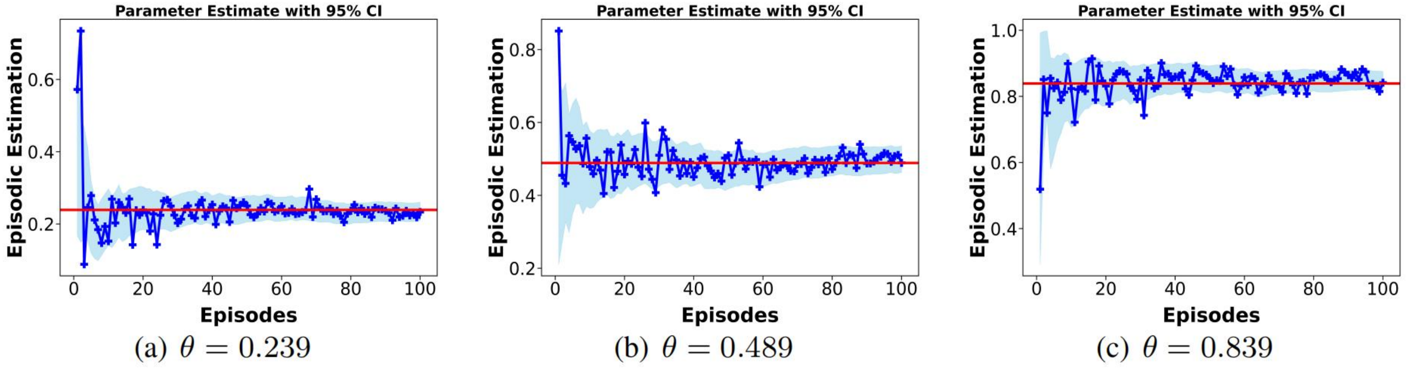

Fig. 7. The parameter estimates trend plots of three categories of probabilities.

CONTENTS

Conclusion

- Studied AoI-aware MCS for uncertain event capture
- Modeled as **constrained episodic restless bandit** with **unknown** transitions
- Proposed **Belief-DPP** and **TS-DPP** algorithms
- Theoretical guarantee: **sublinear Bayesian regret** $\mathcal{O}(\sqrt{L^3 T \log T})$
- Extensive simulations validate performance and robustness

INFOCOM 2026

Thank you for your attention!

Jinrui Zhou, Yin Xu, Haotian Xu, Mingjun Xiao, Jie Wu

zzkevin@mail.ustc.edu.cn