Optimal Scheduling for Minimizing the Age-of-Information for Wireless Erasure Channels

Abhishek Sinha

Assistant Professor
EE, IIT Madras

† Joint work with

Igor Kadota, Eytan Modiano (MIT)

Arunabh Srivastava, Krishna Jagannathan (IIT Madras)

Friday 28th June, 2019

† INFOCOM 2018 (Best Paper Award), RAWNET 2019

Age Of Information (AoI)

What is AoI - A new metric to measure the freshness of information

Age Of Information (AoI)

What is AoI - A new metric to measure the freshness of information

Definition [KYG12]: The AoI $h(t)$ for a UE is the time elapsed since the UE received the latest packet. Mathematically,

$$h(t) = t - u(t),$$

where $u(t)$ is the timestamp of the *latest* received packet.

Age Of Information (AoI)

What is AoI - A new metric to measure the freshness of information

DEFINITION [KYG12]: The AoI $h(t)$ for a UE is the time elapsed since the UE received the latest packet. Mathematically,

$$h(t) = t - u(t),$$

where $u(t)$ is the timestamp of the *latest* received packet.

Saw-Tooth Variation of AoI with time

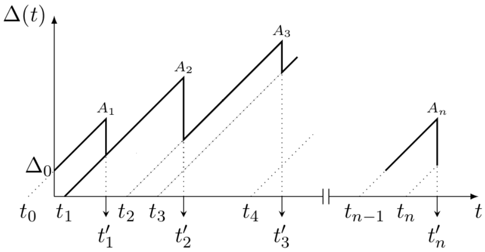

Use Case I - Self-Driving Car

- A Self-Driving Car uses many sensors to navigate through traffic on the road.
  - *e.g.*, Waymo by Google uses the LIDAR, eight laser sensors, cameras, GPS and radar systems

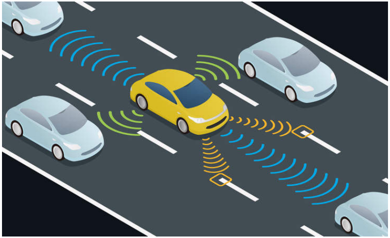

A Self-Driving Car

- The controller needs to obtain the *latest* readings from all sensors, and cannot ignore even one sensor for a long time

☞ Constraint: Due to wireless interference, can communicate with only a limited number of sensors per slot.

Use Case II- Automated Surveillance

- Automated intrusion detection in large areas requires a well-connected sensor network
- The central server requires live information from all sensors to detect intrusions
- Necessary to communicate with all sensors to identify the intruders with high accuracy

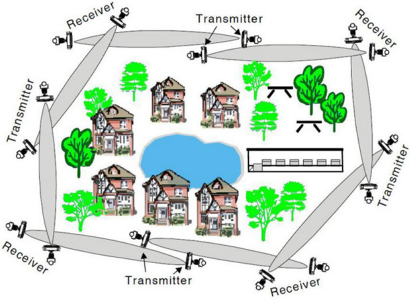

An Intrusion Detection System

Use Case II- Automated Surveillance

- Automated intrusion detection in large areas requires a well-connected sensor network
- The central server requires live information from all sensors to detect intrusions
- Necessary to communicate with all sensors to identify the intruders with high accuracy

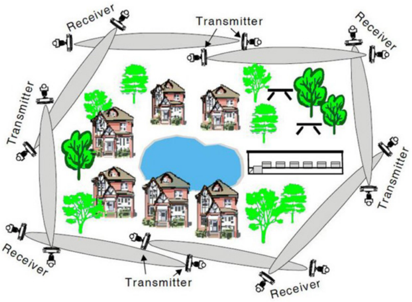

An Intrusion Detection System

☞ Constraint: Throughput constraints on the wireless links and wireless interference constraints

System Model

- A BS serves N UEs
- ARRIVAL: The BS receives one fresh packet per slot from a core network
- SCHEDULING: The BS can transmit the latest packet to only one UE per slot
- CHANNEL: The channel between the BS and the i^{th} UE is modelled by a `erasure channel` with erasure probability $1 - p_i$.

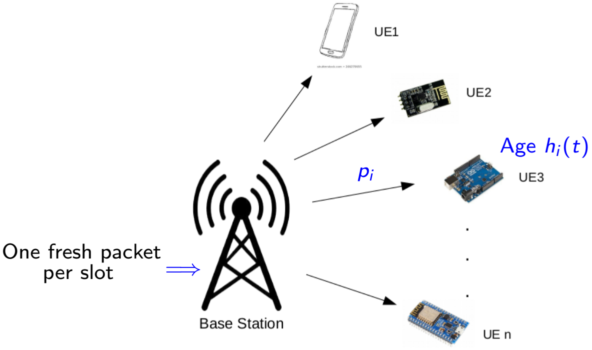

Problem Statement-I and Results

Objective: Design a UE scheduling policy to maximize the value of information.

Problem 1: Minimize the Average-AoI

Design a downlink scheduling policy which minimizes the long-term average-AoI (H_{avg}) of the UEs as defined below

$$H_{\text{avg}} \equiv \limsup_{T \to \infty} \frac{1}{T} \sum_{t=1}^{T} \left(\frac{1}{N} \sum_{i=1}^{N} \mathbb{E}(h_i(t)) \right)$$

Problem Statement-I and Results

Objective: Design a UE scheduling policy to maximize the value of information.

Problem 1: Minimize the Average-AoI

Design a downlink scheduling policy which minimizes the long-term average-AoI (H_{avg}) of the UEs as defined below

$$H_{\mathrm{avg}} \equiv \limsup_{T\to\infty} \frac{1}{T}\sum_{t=1}^{T}\left(\frac{1}{N}\sum_{i=1}^{N}\mathbb{E}(h_i(t))\right)$$

Our Results

1. Derivation of a universal lower-bound for H_{avg}
2. Designing a 4-approximation policy `MW`
3. Extension of `MW` with throughput-constraint

Converse

Theorem: Universal Lower Bound

For *any* UE scheduling policy π, we have

$$H^{\pi}_{\text{avg}} \geq \frac{1}{2N}\left(\sum_{i=1}^{N} \frac{1}{\sqrt{p_i}}\right)^2.$$

Converse

Theorem: Universal Lower Bound

For *any* UE scheduling policy π, we have

$$H^{\pi}_{\text{avg}} \geq \frac{1}{2N}\left(\sum_{i=1}^{N} \frac{1}{\sqrt{p_i}}\right)^2.$$

Proof Outline:

- The proof uses the fact that, irrespective of any policy π, a *maximum* of T *transmission attempts* can be made in T slots.
- This, along with the dynamics of age process, yields a lower-bound upon application of the Cauchy-Schwartz inequality.
- Finally, the proof concludes by using the SLLN and Fatou's Lemma.

Achievability

The Max-Weight Policy (MW)

At time slot t, the MW policy schedules the user $i^{\texttt{MW}}(t)$ having the highest index $p_i h_i^2(t)$, i.e.,

$$i^{\texttt{MW}}(t) \in \arg\max_i p_i h_i^2(t).$$

- The MW policy requires the knowledge of the channel statistics ($\boldsymbol{p}$).

Achievability

The Max-Weight Policy (MW)

At time slot t, the MW policy schedules the user $i^{\text{MW}}(t)$ having the highest index $p_i h_i^2(t)$, i.e.,

$$i^{\text{MW}}(t) \in \arg\max_i p_i h_i^2(t).$$

- The MW policy requires the knowledge of the channel statistics ($\boldsymbol{p}$).

Theorem: Performance of MW

The MW policy is a 4-approximation scheduling policy for the Problem 1.

Proof Outline

- The proof follows a Lyapunov-drift argument with a quadratic Lyapunov function.
- We compare the drift of MW with the drift of the "best" randomized policy π^*
  - With our methodology, the approximation guarantee of MW is essentially limited by that of π^*

Extension: Minimizing avg-AoI with Throughput-Constraints

We consider the above problem with the constraint that the UE_i has a throughput-requirement of $\alpha_i, \forall i$.

Lemma (Feasibility of $\boldsymbol{\alpha}$)

The throughput vector $\boldsymbol{\alpha}$ is feasible iff

$$\sum_i \frac{\alpha_i}{p_i} < 1.$$

Extension: Minimizing avg-AoI with Throughput-Constraints

We consider the above problem with the constraint that the UE_i has a throughput-requirement of $\alpha_i, \forall i$.

Lemma (Feasibility of $\boldsymbol{\alpha}$)

The throughput vector $\boldsymbol{\alpha}$ is feasible iff

$$\sum_i \frac{\alpha_i}{p_i} < 1.$$

Proposition: Universal Lower-Bound with TPUT Constraint

The avg-AoI is lower-bounded by the value of the following program

$$\min \frac{1}{2N} \sum_i \frac{1}{\beta_i}$$

Subject to,

$$\beta_i \geq \alpha_i, \forall i$$
$$\sum_i \frac{\beta_i}{p_i} \leq 1$$
$$\beta_i \geq 0.$$

Approximately-Optimal MW Policy

For a scalar parameter $V > 0$, define the weight

$$W_i(t) = p_i h_i^2(t) + 2Vp_i q_i^+(t),$$

where $q_i^+(t)$ is the "debt-queue" for the UE_i having the dynamics

$$q_i^+(t+1) = \Big(q_i^+(t) - \mu_i(t)\Big)^+ + \alpha_i.$$

At time t, the MW-T policy schedules the UE_i having the largest value of the weight $W_i(t)$, i.e.,

$$i^{\texttt{MW-T}} \in \arg\max_i \Big(p_i h_i^2(t) + 2Vp_i q_i^+(t)\Big).$$

Approximately-Optimal MW Policy

For a scalar parameter $V > 0$, define the weight

$$W_i(t) = p_i h_i^2(t) + 2Vp_i q_i^+(t),$$

where $q_i^+(t)$ is the "debt-queue" for the UE_i having the dynamics

$$q_i^+(t+1) = \Big(q_i^+(t) - \mu_i(t)\Big)^+ + \alpha_i.$$

At time t, the MW-T policy schedules the UE_i having the largest value of the weight $W_i(t)$, i.e.,

$$i^{\texttt{MW-T}} \in \arg\max_i \Big(p_i h_i^2(t) + 2Vp_i q_i^+(t)\Big).$$

Optimality of MW-T

The MW-T policy is a 4-optimal scheduling policy in this setting for $0 < V \leq 2$.

Problem Statement-II and results

Emerging applications like URLLC and Cyber Physical Systems require a more stringent uniform control of AoI across all devices.

Problem 2: Minimize the Peak-AoI

Design a downlink scheduling policy which minimizes the long-term peak-AoI ($H_{\max}$) of the UEs as defined below

$$H_{\max} \equiv \limsup_{T \to \infty} \frac{1}{T} \sum_{t=1}^{T} \mathbb{E}(\max_i h_i(t))$$

Problem Statement-II and results

Emerging applications like URLLC and Cyber Physical Systems require a more stringent uniform control of AoI across all devices.

Problem 2: Minimize the Peak-AoI

Design a downlink scheduling policy which minimizes the long-term peak-AoI ($H_{\max}$) of the UEs as defined below

$$H_{\max} \equiv \limsup_{T\to\infty} \frac{1}{T}\sum_{t=1}^{T} \mathbb{E}(\max_i h_i(t))$$

Our Results

1. Derivation of an Optimal Policy - `Max-Age` (`MA`)
2. Large Deviation Optimality for `MA`
3. Extension of `MA` with throughput-constraint

Optimal Policy - Max Age (MA)

Max Age Policy (MA)

At time slot t, the MA policy schedules the user $i^{\mathtt{MA}}(t)$ having the highest instantaneous age, i.e.,

$$i^{\mathtt{MA}}(t) \in \arg\max_i h_i(t).$$

- Unlike MW, the MA policy is greedy and is oblivious to the channel statistics ($\boldsymbol{p}$).
    - Upshots: Easy to implement as it requires no channel estimations.

Optimal Policy - Max Age (MA)

Max Age Policy (MA)

At time slot t, the MA policy schedules the user $i^{\texttt{MA}}(t)$ having the highest instantaneous age, i.e.,

$$i^{\texttt{MA}}(t) \in \arg\max_i h_i(t).$$

- Unlike MW, the MA policy is greedy and is oblivious to the channel statistics ($\boldsymbol{p}$).
  - Upshots: Easy to implement as it requires no channel estimations.

Theorem (Optimality of MA)

The MA policy is an optimal policy for Problem 1. Moreover, the optimal long term peak AoI is given by

$$H^*_{\max} = \sum_{i=1}^{N} \frac{1}{p_i}.$$

Proof Outline

- Problem 1 is an instance of a countable-state average-cost MDP with a finite action space.
    - Very hard to solve exactly, due to infinite state-space (VI, PI do not work!).
- Our proof starts by writing down the associated Bellman Equation (BE):

$$\lambda^* + V(\mathbf{h}) = \min_i \Big(p_i V(1, h_{-i} + 1) + (1 - p_i) V(\mathbf{h+1}) \Big) + \max_i h_i \quad (1)$$

- Note that, (1) is a system of infinitely many non-linear equations.

Proof Outline

- Problem 1 is an instance of a countable-state average-cost MDP with a finite action space.
  - Very hard to solve exactly, due to infinite state-space (VI, PI do not work!).
- Our proof starts by writing down the associated Bellman Equation (BE):

$$\lambda^* + V(\mathbf{h}) = \min_i \Big(p_i V(1, h_{-i} + 1) + (1 - p_i) V(\mathbf{h+1}) \Big) + \max_i h_i \quad (1)$$

  - Note that, (1) is a system of infinitely many non-linear equations.
- We next propose the following linear candidate solution to the BE:

$$V(\mathbf{h}) = \sum_j \frac{h_j}{p_j}, \quad \lambda^* = \sum_j \frac{1}{p_j} \quad (2)$$

Proof Outline

- Problem 1 is an instance of a countable-state average-cost MDP with a finite action space.
  - Very hard to solve exactly, due to infinite state-space (VI, PI do not work!).
- Our proof starts by writing down the associated Bellman Equation (BE):

$$\lambda^* + V(\mathbf{h}) = \min_i \Big(p_i V(1, h_{-i} + 1) + (1 - p_i) V(\mathbf{h+1}) \Big) + \max_i h_i \quad (1)$$

  - Note that, (1) is a system of infinitely many non-linear equations.
- We next propose the following linear candidate solution to the BE:

$$V(\mathbf{h}) = \sum_j \frac{h_j}{p_j}, \quad \lambda^* = \sum_j \frac{1}{p_j} \quad (2)$$

- Finally, we show that (2) satisfies the BE under MA.

Stability of the Age Process

We next show that, under the MA policy, the age-process is stable.

Theorem

The Markov Chain $\{h(t)\}_{t\geq 1}$ is Positive Recurrent under the action of the MA Policy.

Positive recurrence of the age-process implies

- Each UE is served infinitely often w.p. 1.
- The expected time between two consecutive service of a UE has a finite expected value.

Proof Outline: The proof follows a Lyapunov-drift approach with a Linear Lyapunov function. Details in the paper.

Large Deviation Optimality of MA

A more refined performance measure of a scheduler is its large-deviation exponent I defined below

$$I = -\lim_{k\to\infty}\lim_{t\to\infty}\frac{1}{k}\log\mathbb{P}(\max_i h_i(t)\geq k).$$

- ☞ The larger the value of I, the (exponentially) smaller the probability of age exceeding a threshold.

Large Deviation Optimality of MA

A more refined performance measure of a scheduler is its large-deviation exponent I defined below

$$I = -\lim_{k\to\infty}\lim_{t\to\infty}\frac{1}{k}\log\mathbb{P}(\max_i h_i(t)\geq k).$$

- ☞ The larger the value of I, the (exponentially) smaller the probability of age exceeding a threshold.

Theorem (MA is LD-Optimal)

The MA policy maximizes the Large-Deviation exponent and the value of the optimal exponent is given by

$$I^{MA} = \max I = -\log(1-p_{min}).$$

Proof Outline: The proof proceeds by deriving a converse (universal upper-bound) and a matching lower-bound for the MA policy. Details in the paper.

Extension: Minimizing Age with Throughput Constraints

As an extension, we consider a scenario, where UE_1 is throughput-constrained and the rest of the UEs are delay-constrained.

Problem 2: Minimize Age with TPUT Constraint

Find an optimal scheduling policy which minimizes the long-term max-age of all UEs subject to the throughput-constraint of one UE.

- By relaxing the throughput constraint, we obtain the following relaxed objective:

$$\lambda^{**} = \inf_{\pi \in \Pi} \limsup_{T \to \infty} \frac{1}{T} \sum_{t=1}^{T} \mathbb{E}(\max_i h_i(t) + \beta \bar{a}_1(t)),$$

where $\bar{a}_1(t) = \mathbb{1}(\text{UE}_1 \text{ did not successfully receive a packet in slot } t)$, and $\beta \geq 0$ is a scalar Lagrangian coefficient.

Heuristic Policy - MATP

Let g_i denote the expected cost when UE_1 did not receive a packet successfully, i.e., $g_i = \beta - \beta p_1 \mathbb{1}(i = 1)$.

The MATP Policy

At any slot t, the MATP policy serves the user $i^{\texttt{MATP}}(t)$ having highest value of $h_i(t) - g_i$, i.e.,

$$i^{\texttt{MATP}} \in \arg\max_i h_i(t) - g_i.$$

Heuristic Policy - MATP

Let g_i denote the expected cost when UE_1 did not receive a packet successfully, i.e., $g_i = \beta - \beta p_1 \mathbb{1}(i = 1)$.

The MATP Policy

At any slot t, the MATP policy serves the user $i^{\text{MATP}}(t)$ having highest value of $h_i(t) - g_i$, i.e.,

$$i^{\text{MATP}} \in \arg\max_i h_i(t) - g_i.$$

Proposition: Approximate Optimality of MATP

There exists a value function $V(\cdot)$, such that, under the MATP policy, we have

$$||V - TV||_\infty \leq \beta p_1,$$

where $T(\cdot)$ is the associated Bellman Operator.

Minimize the Peak-Age

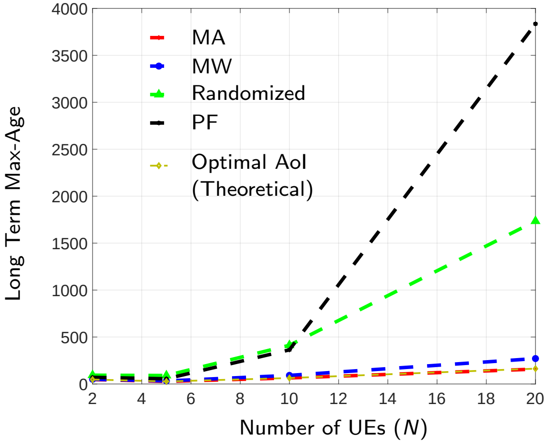

Problem 1: Performance of the Max-Age (MA) policy with three other Scheduling Policies for different number of UEs.

Age vs Throughput Variation of MATP with the β Parameter

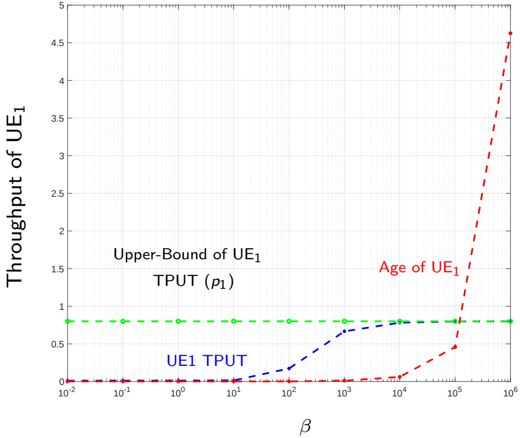

PROBLEM 2: Variation of Throughput of UE_1 with the parameter β.

Minimize the Peak-Age with TPUT Constraint

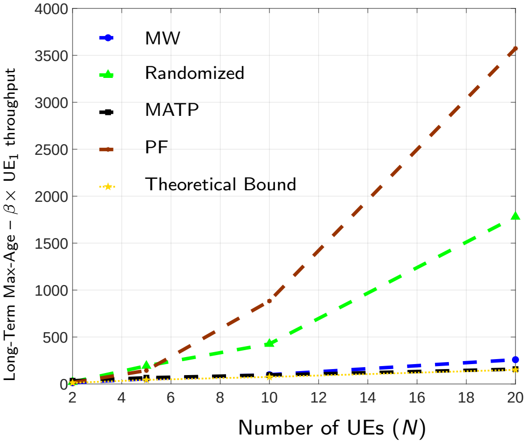

PROBLEM 2: Comparative Performance of the Proposed MATP Policy with other well-known scheduling policies.

Conclusion

- We formulated the problem of minimizing the average-age and peak-age in the single-hop setting
- Derived an approximately optimal policy for the former and an optimal policy for the latter
- Also Established large-deviation optimality of `MA` and Positive Recurrence of the Age process under `MA`.
- Future work will be on deriving an exactly optimal policy for the throughput-constrained case

Thank You