On Minimizing the Maximum Age-of-Information for Wireless Erasure Channels

Arunabh Srivastava

Joint work with

Abhishek Sinha (advisor), Krishna Jagannathan

Department of Electrical Engineering

IIT Madras

RAWNET 2019, Avignon, France

May 24, 2019

Age Of Information (AoI)

What is AoI - A new metric to evaluate the freshness of information at UE

- DEFINITION: The AoI $h(t)$ at time t for a UE is defined as the time elapsed since the UE received the last packet prior to time t. Mathematically,

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

where u(t) is the timestamp of the last received packet by the UE

Age Of Information (AoI)

What is AoI - A new metric to evaluate the freshness of information at UE

- DEFINITION: The AoI $h(t)$ at time t for a UE is defined as the time elapsed since the UE received the last packet prior to time t. Mathematically,

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

where u(t) is the timestamp of the last received packet by the UE

Saw-Tooth Variation of AoI with time

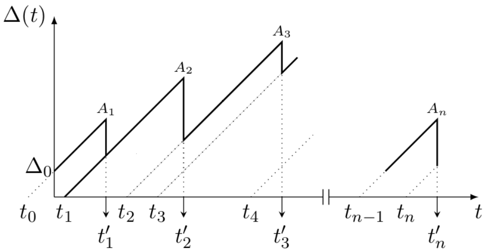

The jumps correspond to reception of new packets by the UEs.

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

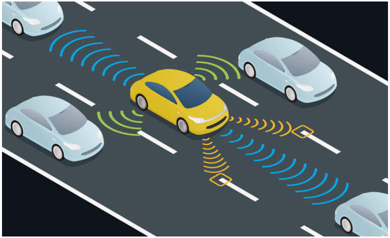

A Self-Driving Car

- The controllers need to obtain the *latest* readings from all sensors and cannot ignore even one sensor for a long 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

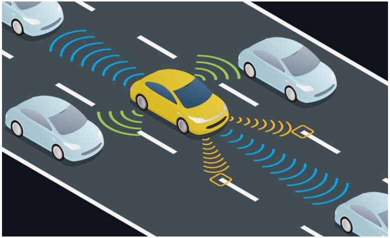

A Self-Driving Car

- The controllers need 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- Intrusion Detection

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

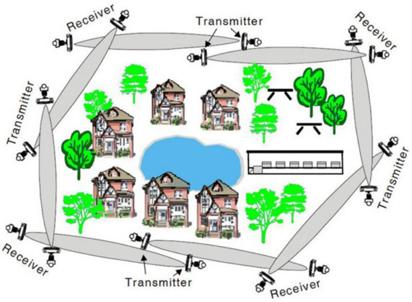

An Intrusion Detection System

Use case II- Intrusion Detection

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

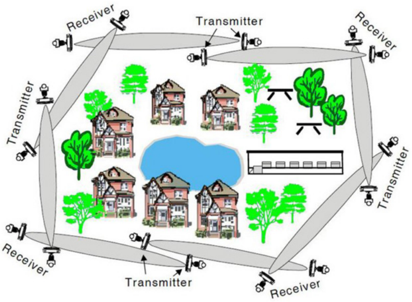

An Intrusion Detection System

☞ 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 a packet to only one UE per slot
- CHANNEL: The channel between the BS and the i^{th} UE is modelled by a binary erasure channel (BEC) with erasure probability $1 - p_i$.

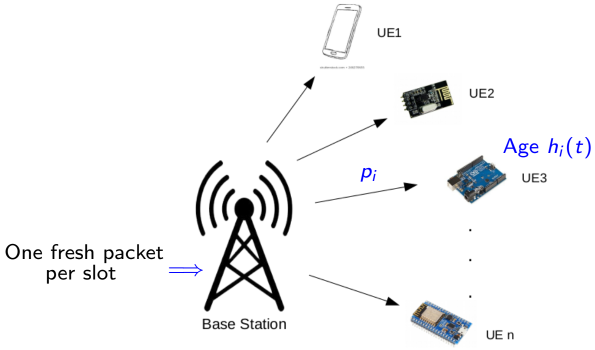

Problem Statement and Results

Objective: Design an optimal scheduling policy to maximize the value of information.

Problem 1: 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 and Results

Objective: Design an optimal scheduling policy to maximize the value of information.

Problem 1: 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 the 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 $i^{\texttt{MA}}(t)$ having the highest instantaneous age, i.e.,

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

- 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.,

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

- 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!).

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 of Age-vectors $\{\boldsymbol{h}(t)\}_{t\geq 1}$ is Positive Recurrent under the action of the MA Policy.

The above theorem implies that the age of *each UE* reaches the lowest value 1 infinitely often with probability 1.

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

Large Deviation Optimality for 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 for 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 universal upper-bound (applicable to all scheduling policies) 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 the eMBB 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

- We do not have the exact optimal policy to Problem 2 yet.
- Inspired by the optimality of MA, we propose the MATP policy which approximately solves the associated Bellman Equation.

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.,

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

Heuristic Policy - MATP

- We do not have the exact optimal policy to Problem 2 yet.
- Inspired by the optimality of MA, we propose the MATP policy which approximately solves the associated Bellman Equation.

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.,

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

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.

Benchmark Policies

An Index Policy π schedules a UE at slot t maximizing an index function $I^{\pi}(t)$.

- Index Policies:
  - MA Max Age: $I^{\mathtt{MA}}(t) = \max_i h_i(t)$
  - MW Max Weight: $I^{\mathtt{MW}}(t) = p_i h_i^2(t)$
  - PF Proportional Fair: $I^{\mathtt{PF}}(t) = p_i / R_i(t)$, where $R_i(t)$ is the average rate for UE_i
  - MATP Max-Age with Throughput Constraints: $I^{\mathtt{MATP}} = \max_i (h_i(t) - g_i)$
- Non-Index Policies:
  - Rand Randomized Policy: Schedule a UE uniformly at random

Benchmark Policies

An Index Policy π schedules a UE at slot t maximizing an index function $I^{\pi}(t)$.

- Index Policies:

  MA Max Age: $I^{\mathtt{MA}}(t) = \max_i h_i(t)$

  MW Max Weight: $I^{\mathtt{MW}}(t) = p_i h_i^2(t)$

  PF Proportional Fair: $I^{\mathtt{PF}}(t) = p_i/R_i(t)$, where $R_i(t)$ is the average rate for UE_i

  MATP Max-Age with Throughput Constraints: $I^{\mathtt{MATP}} = \max_i(h_i(t) - g_i)$

- Non-Index Policies:

  Rand Randomized Policy: Schedule a UE uniformly at random

Theorem (Kadota, Sinha, Modiano, 2018)

The MW policy is a 2-optimal policy for the AVERAGE-AGE metric.

Long Term Peak Age

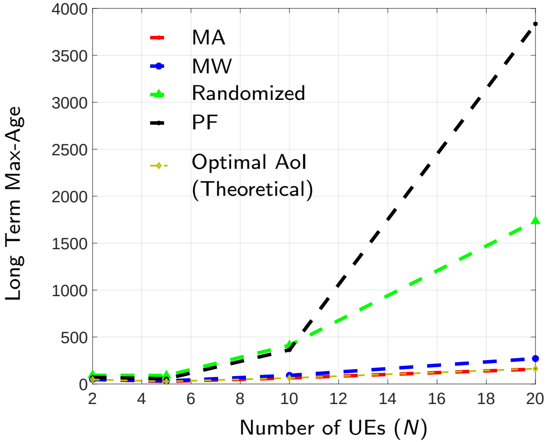

PROBLEM 1: Performance of the Max-Age (MA) policy with three other Scheduling Policies for different number of UEs.

Throughput Variation of MATP with the β Parameter

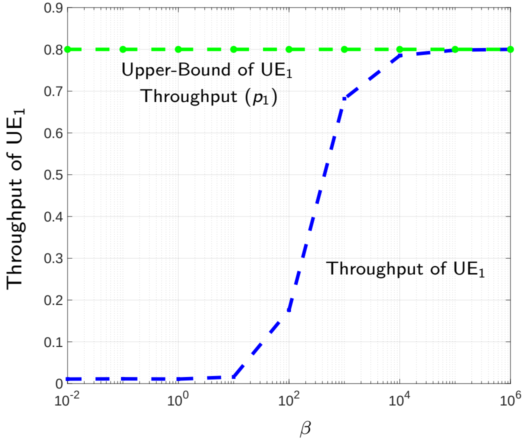

PROBLEM 2: Variation of Throughput of UE_1 with the parameter β.

Comparison of Policies

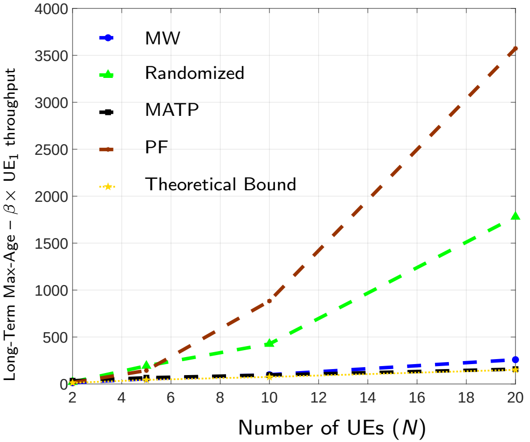

PROBLEM 2: Comparative Performance of the Proposed MATP Policy with other well-known scheduling policies.

Conclusion

- We formulated the problem of minimizing the long-term peak-age for a single-hop downlink communication setting
- We derived an optimal scheduling policy MA
- We 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-constraint case

Thank You