Minimizing the Age of Information in Broadcast Wireless Networks

Igor Kadota, Elif Uysal-Biyikoglu, Rahul Singh and Eytan Modiano

LIDS Student Conference, February 2, 2017

Outline

- Age of Information (AoI) and Network Model
- Symmetric Network and the Greedy Policy
- General Network and the Index Policy
- Numerical Results
- Max Throughput vs Min AoI

Age of Information

Example: Single Source

Single Destination

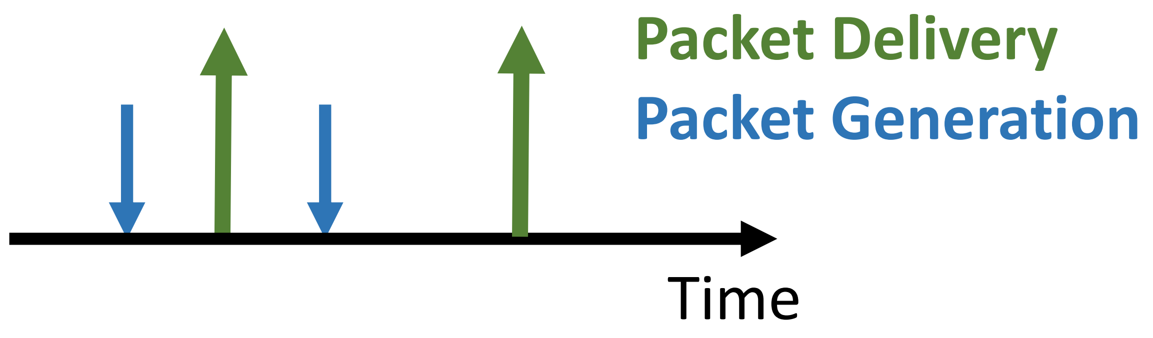

Measuring the freshness of the information

Age of Information

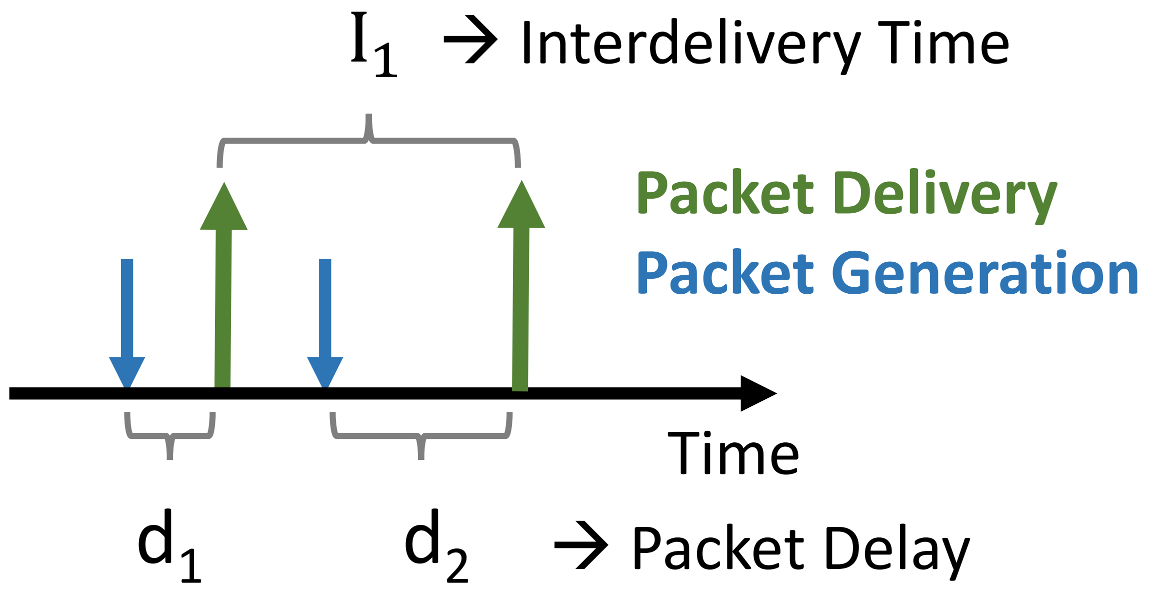

Interdelivery Time: time elapsed between consecutive packet deliveries.

Packet Delay: time elapsed from generation to delivery of a packet.

Age of Information

AoI: time elapsed since the most recently delivered packet was generated.

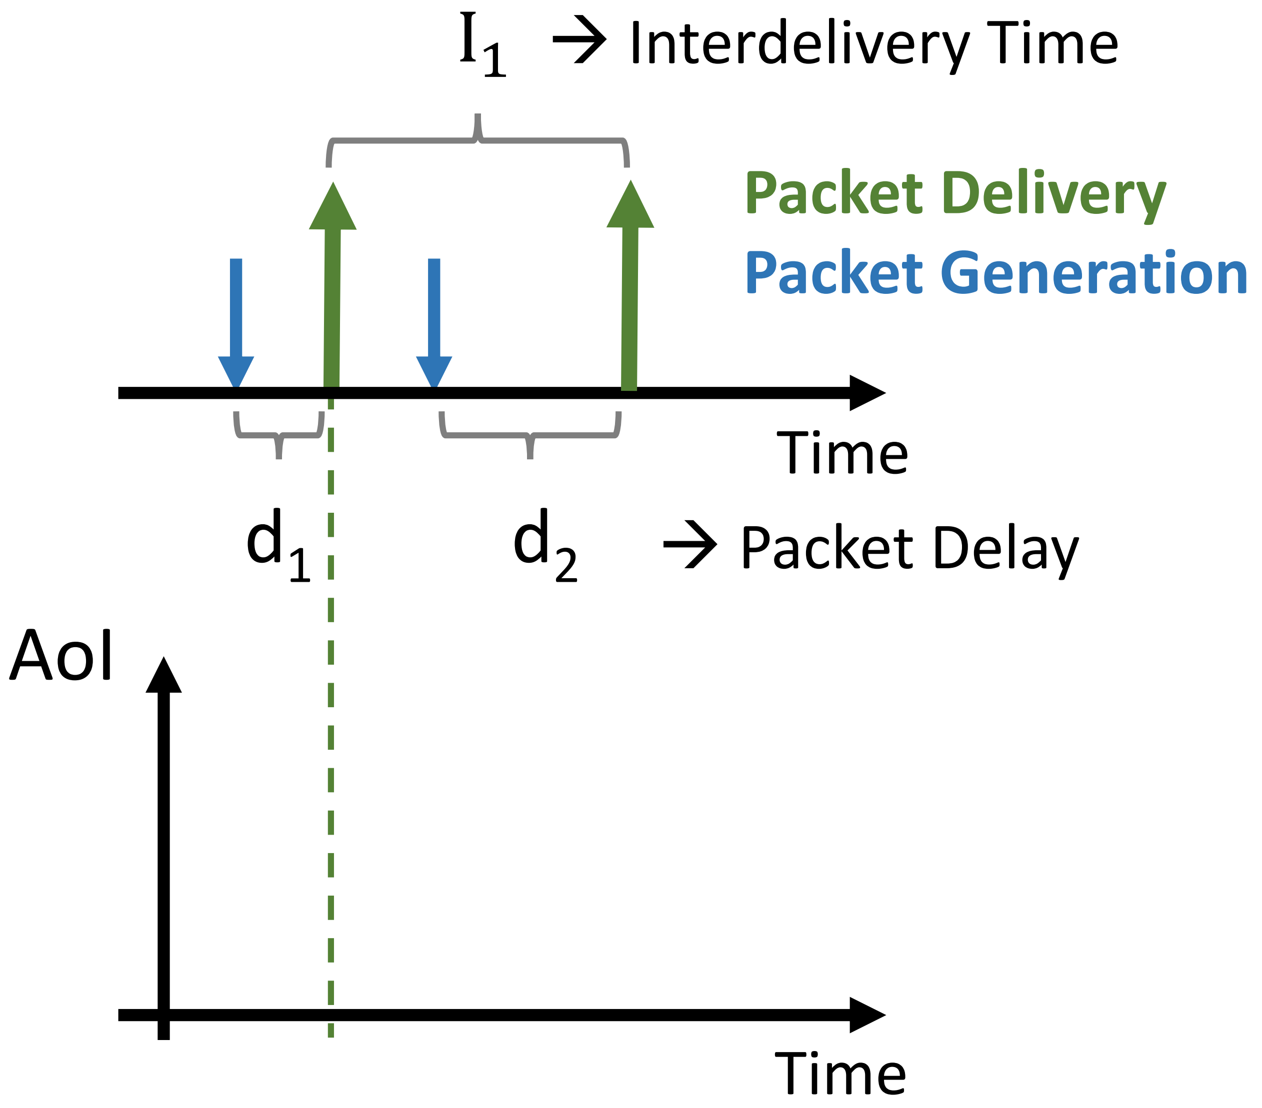

[1] S. Kaul, R. Yates, and M. Gruteser, “Real-time status: How often should one update?”, 2012.

Age of Information

AoI: time elapsed since the most recently delivered packet was generated.

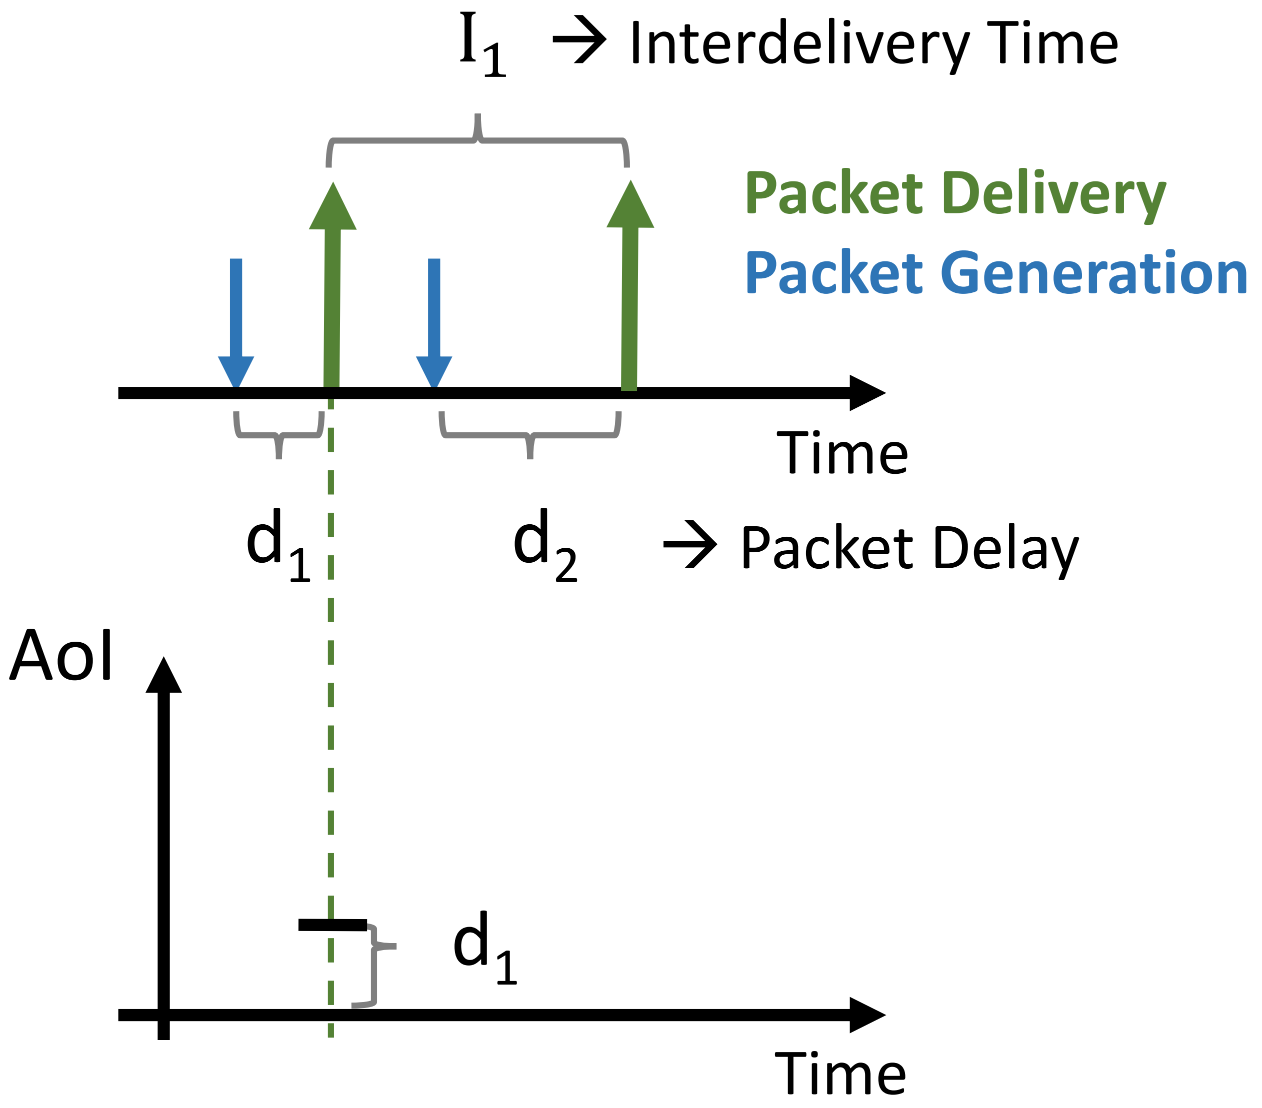

[1] S. Kaul, R. Yates, and M. Gruteser, “Real-time status: How often should one update?”, 2012.

Age of Information

AoI: time elapsed since the most recently delivered packet was generated.

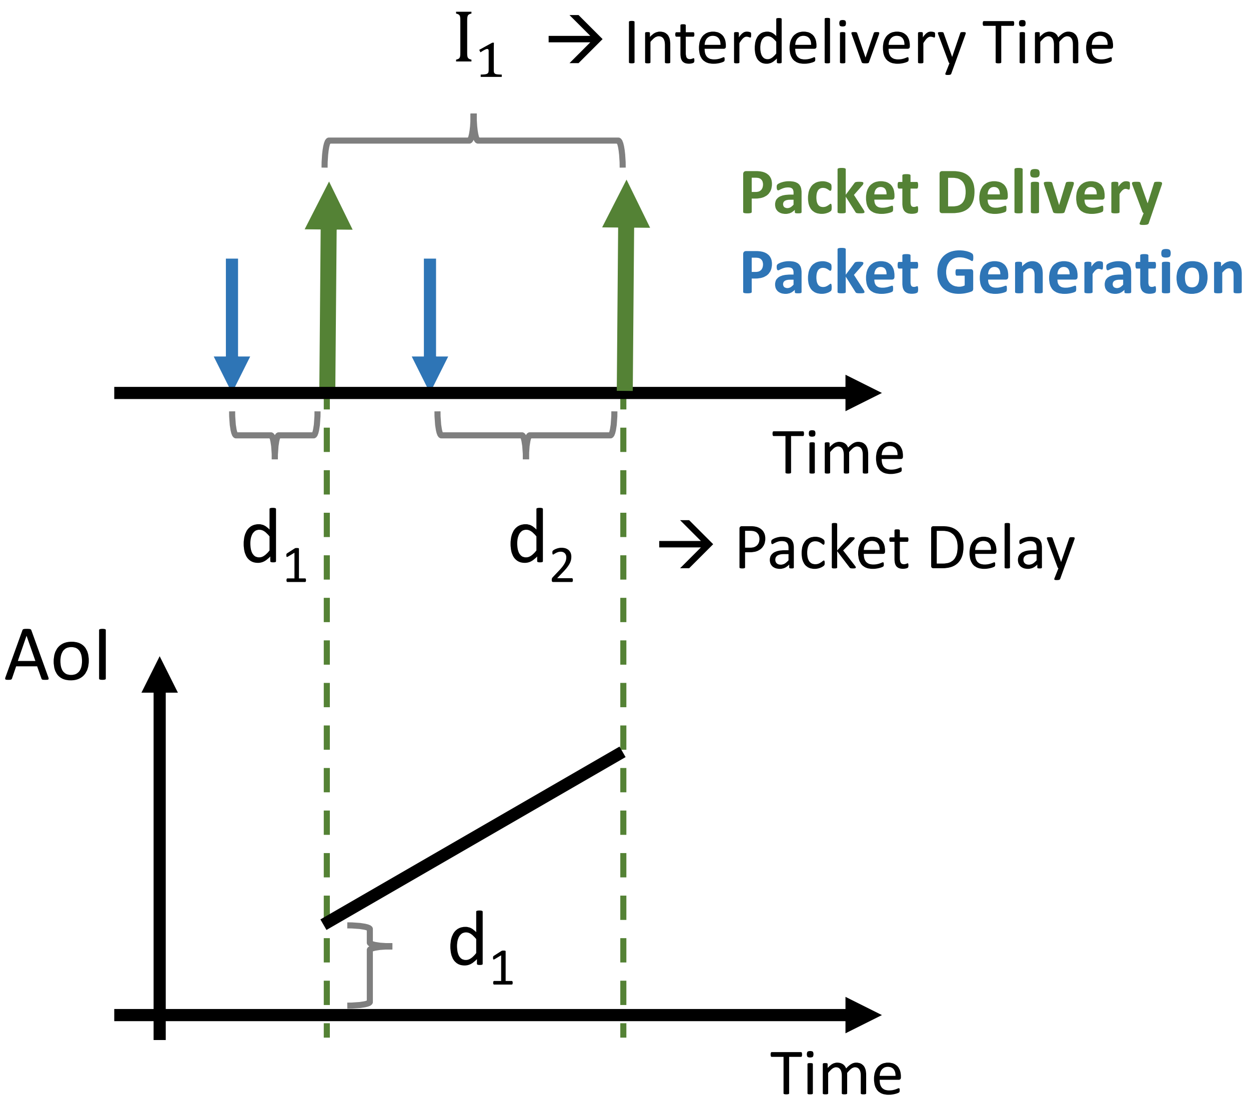

[1] S. Kaul, R. Yates, and M. Gruteser, “Real-time status: How often should one update?”, 2012.

Age of Information

AoI: time elapsed since the most recently delivered packet was generated.

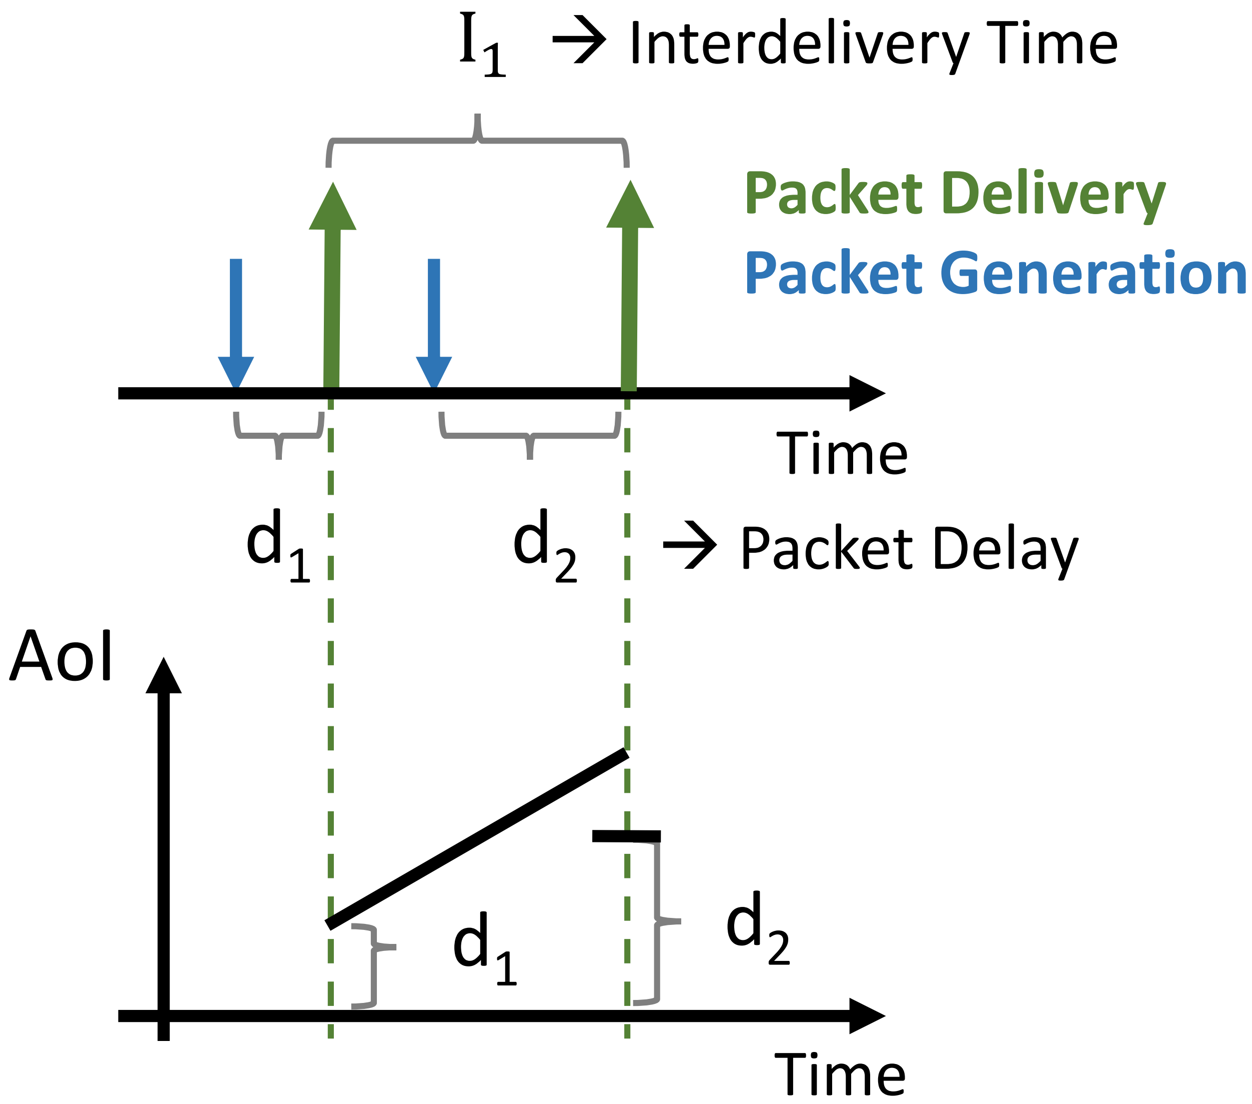

[1] S. Kaul, R. Yates, and M. Gruteser, "Real-time status: How often should one update?", 2012.

Age of Information

AoI: time elapsed since the most recently delivered packet was generated.

At time t: $\mathrm{AoI} = \mathrm{t} - \tau(t)$

$\tau(t)$ is the time stamp of the most recently delivered packet.

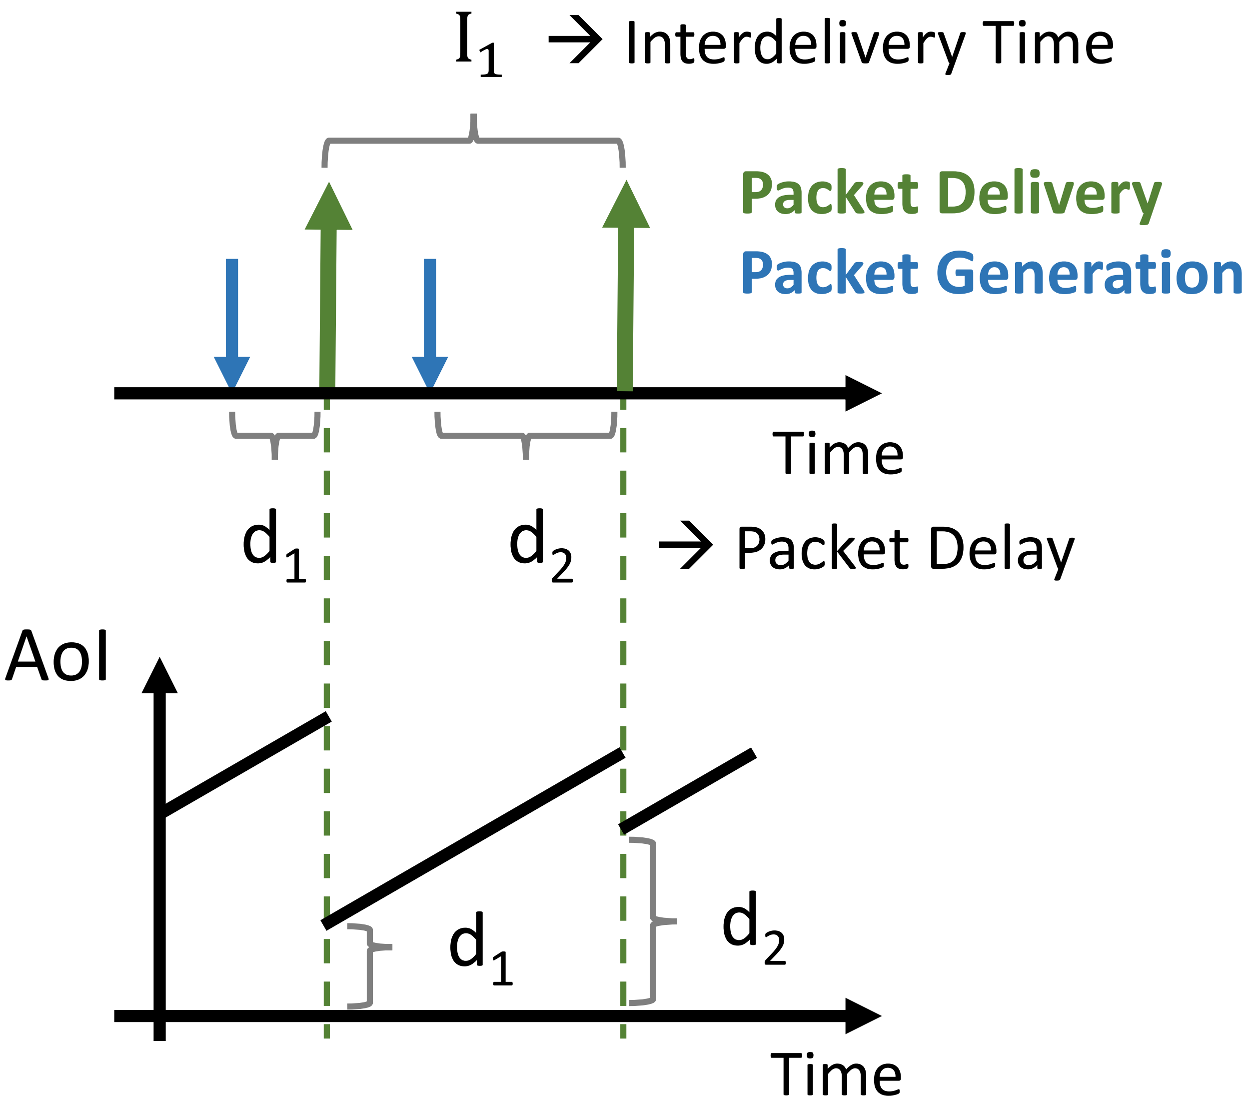

[1] S. Kaul, R. Yates, and M. Gruteser, “Real-time status: How often should one update?”, 2012.

AoI, Delay and Interdelivery time

- Example: M/M/1 queue

  Controllable arrival rate λ and fixed service rate $\mu = 1$ packet per second.

λ	$\mathbb{E}[delay]$	$\mathbb{E}[interdel.]$	Average AoI
0.01			
0.53			
0.99			

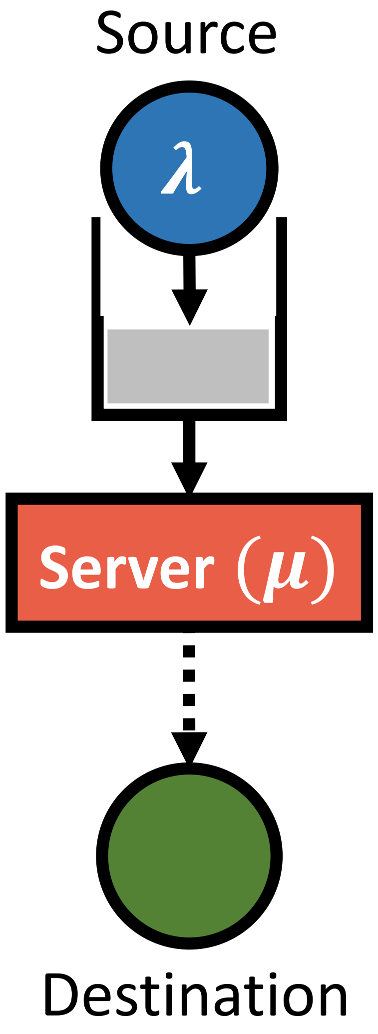

AoI, Delay and Interdelivery time

- Example: M/M/1 queue

  Controllable arrival rate λ and fixed service rate $\mu = 1$ packet per second.

λ	$\mathbb{E}[delay]$	$\mathbb{E}[interdel.]$	Average AoI
0.01	1.01	**100.00**	
0.53	2.13	1.89	
0.99	**100.00**	1.01	

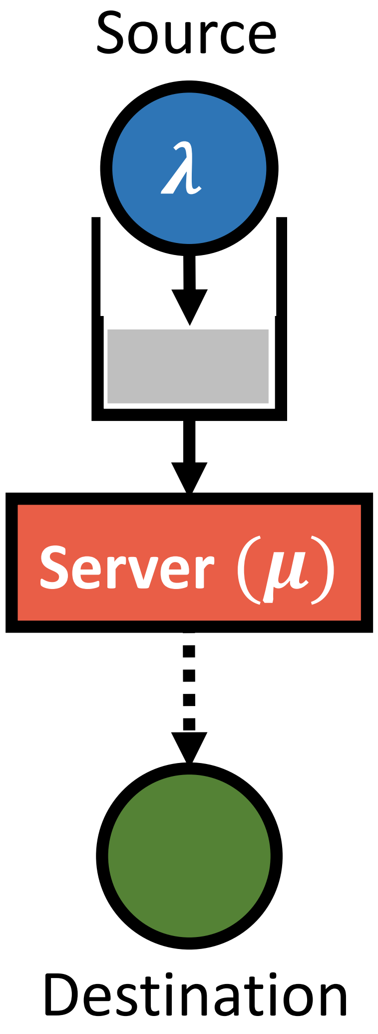

AoI, Delay and Interdelivery time

- Example: M/M/1 queue

Controllable arrival rate λ and fixed service rate $\mu = 1$ packet per second.

λ	$\mathbb{E}[delay]$	$\mathbb{E}[interdel.]$	Average AoI
0.01	1.01	**100.00**	**101.00**
0.53	2.13	1.89	3.48
0.99	**100.00**	1.01	**100.02**

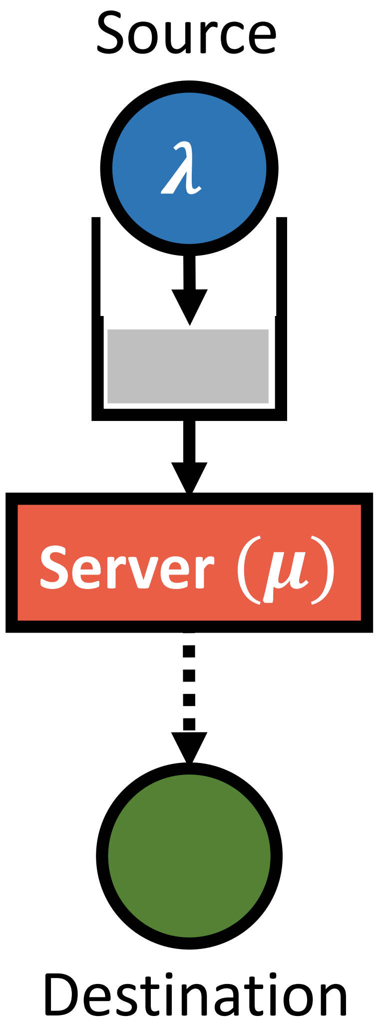

A low AoI is achieved when packets with low delay are delivered regularly.

Minimum throughput requirement DOES NOT guarantee regular deliveries.

[1] S. Kaul, R. Yates, and M. Gruteser, “Real-time status: How often should one update?”, 2012.

Network Model

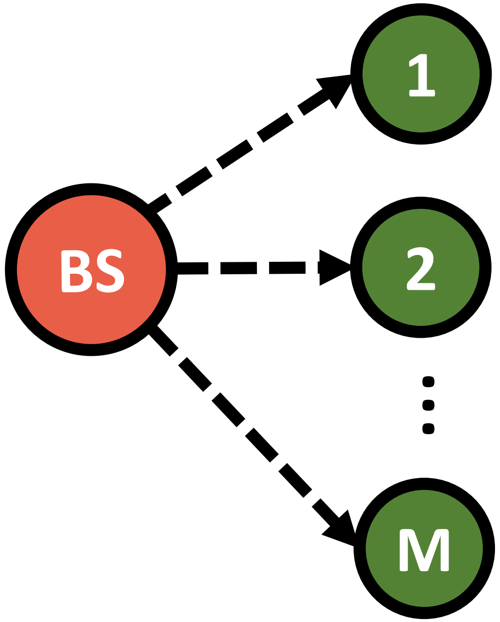

Network Model

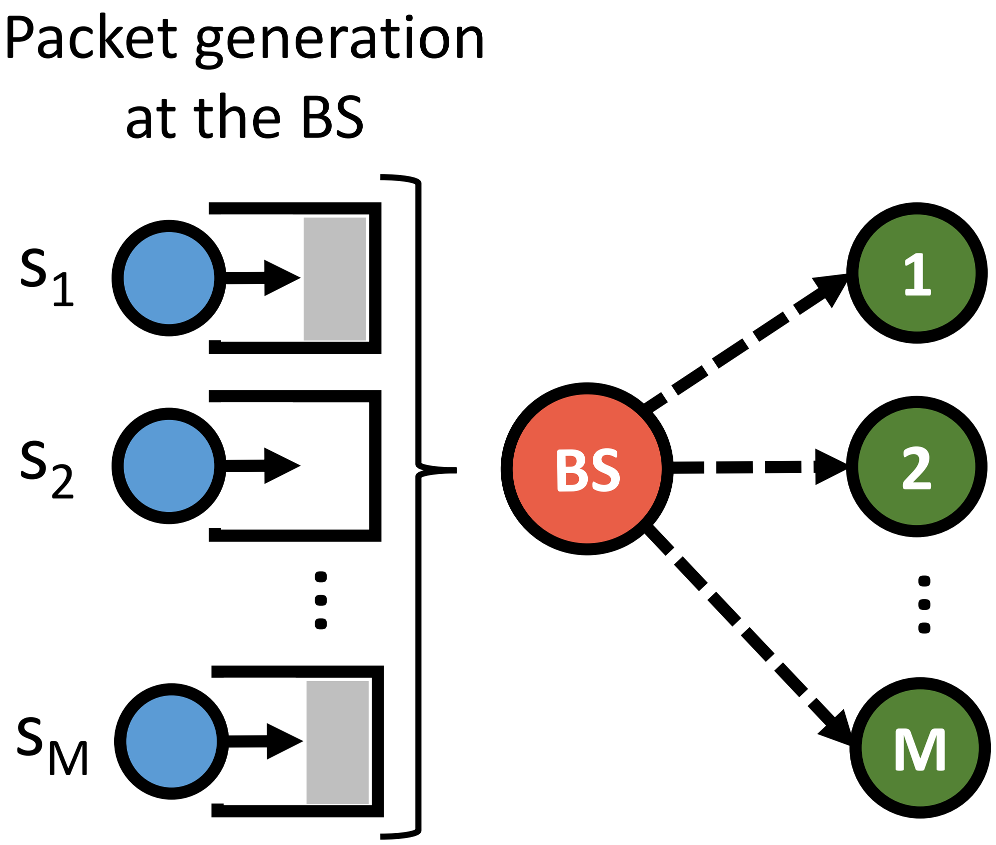

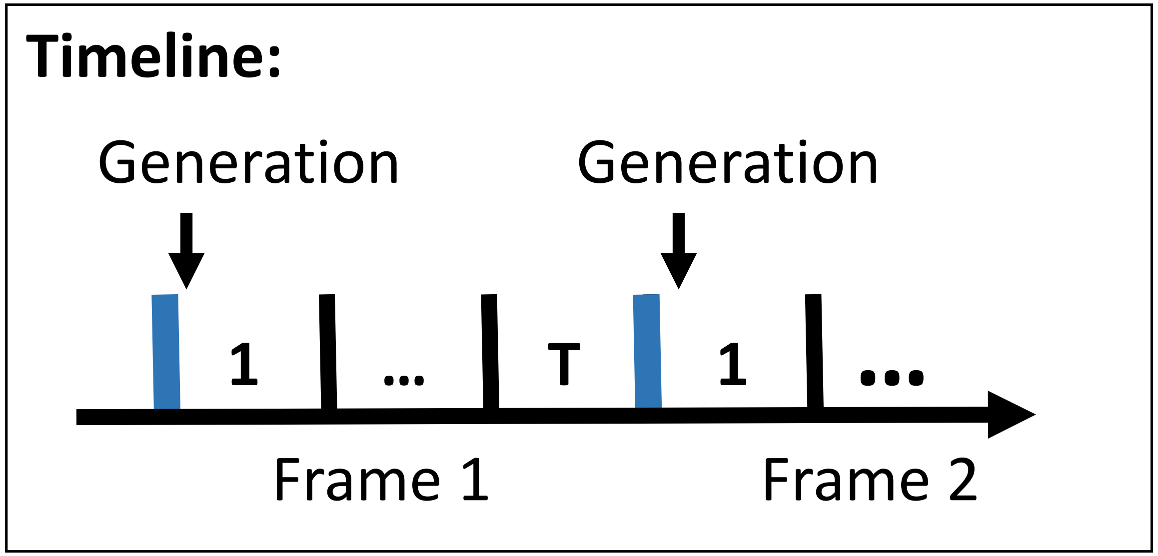

New packets replace undelivered packets from the previous frame.

Network Model

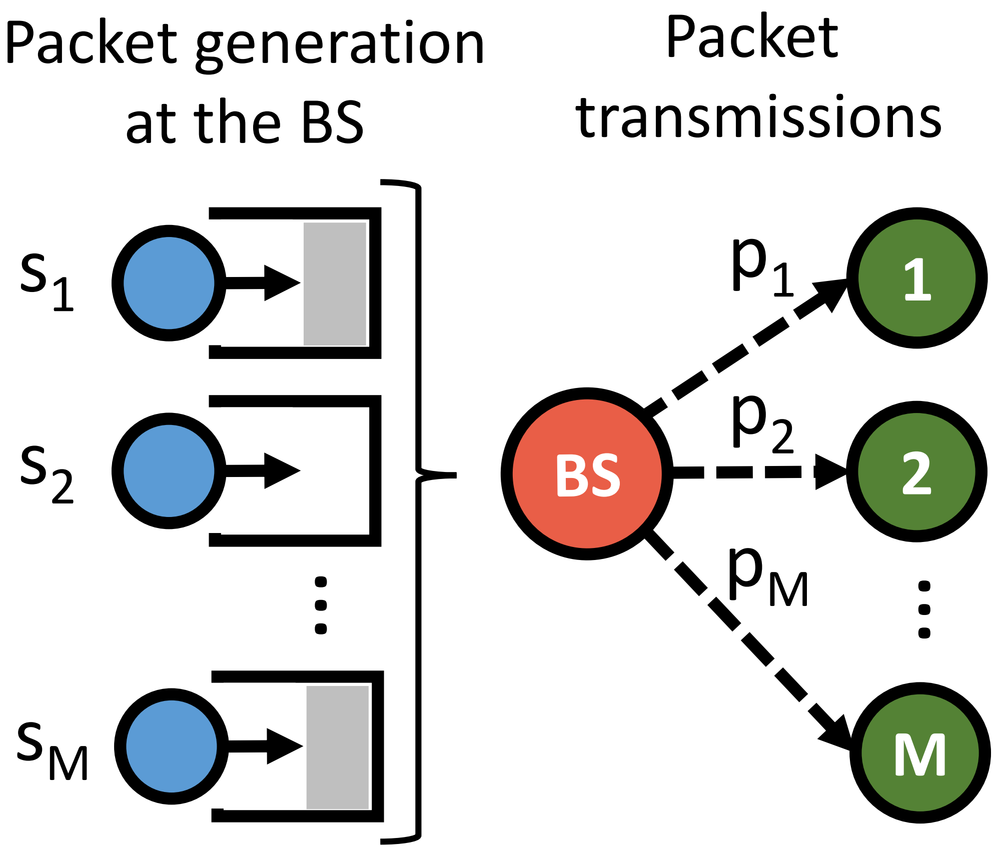

Timeline:

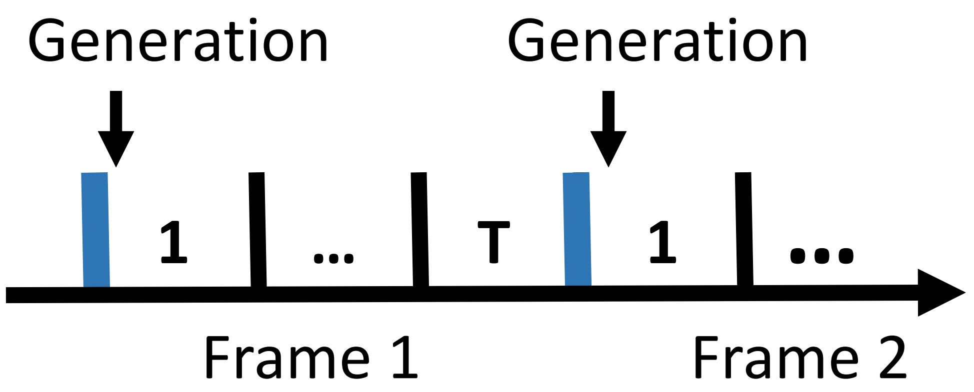

Packet transmission:

- In a slot, a packet is transmitted to a selected client i.
- p_i is the prob. of a successful transm.
- Instantaneous feedback

Network Model

- Goal: design a scheduling policy that provides fresh information to the clients.

- Objective function is the Expected Weighted Sum AoI:

$$\mathrm{EWSAoI} = \frac{1}{\mathrm{KT}} \mathbb{E}\left\{ \sum_{\mathrm{i}=1}^{\mathrm{M}} \boldsymbol{\alpha_i} \, \mathbf{AoI}_i \right\}$$

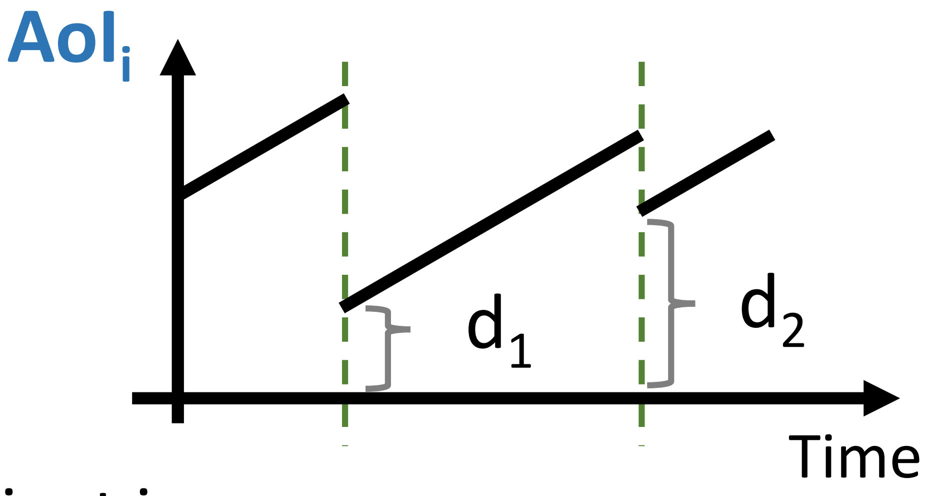

where $\boldsymbol{\alpha_i}$ is the client`s weight and

$\mathbf{AoI}_i$ is the area under the AoI curve for client i

Network Model

- Goal: design a scheduling policy that provides fresh information to the clients.

- Equivalent Objective function:

$$J_K^{\pi^*} = \min_{\pi \epsilon \Pi} \mathbb{E} \left\{ \sum_{k=1}^{K} \sum_{i=1}^{M} \boldsymbol{\alpha_i} \boldsymbol{h_{\mathrm{k},i}} \right\}$$

where $\boldsymbol{\alpha_i}$ is the client`s weight and

$\boldsymbol{h_{\mathrm{k},i}}$ is the number of frames since the last delivery from client i

Π is the class of non-anticipatory policies and π^* is the optimal policy.

Symmetric Network

- Network with symmetric clients: $\alpha_i = \alpha$ and $p_i = p, \forall i \in \{1, \ldots, M\}$

- Greedy Policy (G): in each slot, select the client with undelivered packet and highest value of $h_{k,i}$.

Theorem 1: Optimality of the Greedy Policy.

Among the class of admissible policies Π, G attains the minimum time average sum AoI.

General Network

- Network with clients having (possibly) different α_i and p_i
- Objective Function: $J_K^{\pi^*} = \min_{\pi \epsilon \Pi} \mathbb{E} \left\{ \sum_{k=1}^{K} \sum_{i=1}^{M} \alpha_i \mathrm{h}_{\mathrm{k,i}} \right\}$
- Policy π is a mapping from all possible states in each possible slot to the associated scheduling choice. In general, computing π is complex.
- Index Policy: in each slot, select the client with undelivered packet and *highest value of* $C_i(h_{k,i})$.
- The Index Policy is a low-complexity heuristic that is extensively used in the literature for its strong performance [K. Liu, 2010; R. Weber, 1990; et al.]

General Network: Whittle Index

- For designing the Index Policy, we use the RMAB framework in [16].
  - We relax our problem to the case of a single client, $M = 1$, and add a cost per transmission, $C > 0$.

- The solution to this relaxed problem yields:
  - Condition for indexability;
  - Expression for the Whittle Index, $C_i(h_{k,i})$.

- Challenges:
  - Indexability is often hard to establish
  - Indexable problems might not have closed-form solutions for the Whittle Index.

[16] P. Whittle, “Restless bandits: Activity allocation in a changing world”, 1988.

General Network: Definitions

- Indexability:
  - Consider the relaxed problem with a single client and cost per transmission.
  - Let $\mathcal{P}(C)$ be the set of states for which it is optimal to idle when the cost for transmission is C.
  - The problem is indexable if $\mathcal{P}(C)$ increases monotonically from Ø to the entire state space as C increases from 0 to +∞.
  - The condition checks if the problem is suited for an Index Policy.

- Whittle Index:
  - Given indexability, $C(h)$ is the infimum cost C that makes both scheduling decisions equally desirable in state h.
  - $C(h)$ represents how valuable is to transmit a client in state h.

[16] P. Whittle, "Restless bandits: Activity allocation in a changing world", 1988.

General Network: Index Policy

- We establish that the problem is indexable and find a **closed-form** solution for the Whittle Index.
- Index Policy: in each slot, select the client with undelivered packet and *highest value of* $C_i(h_{k,i})$, where:

$$C_i(h_i) = \frac{T\alpha_i}{2} p_i h_i \left[h_i + \frac{1+(1-p_i)^T}{1-(1-p_i)^T} \right]$$

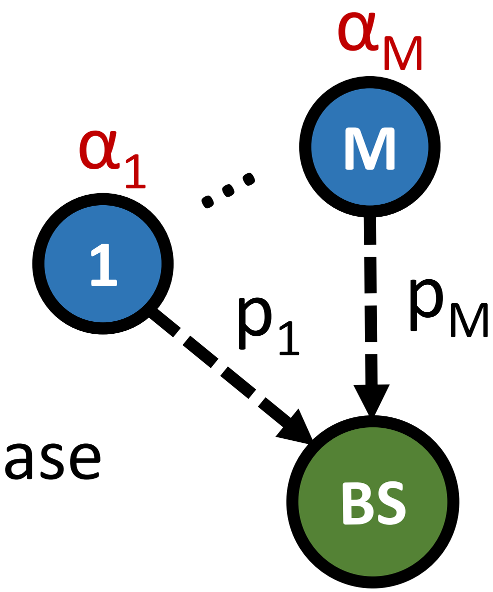

- Observe that the client ordering imposed by the Index Policy is the same as the one imposed by the Greedy Policy for the case of symmetric networks. Thus, the Index Policy is OPTIMAL.

Numerical Results

- Metric is the normalized AoI: $\mathrm{EWSAoI} / \mathrm{MT}$
- Comparison:
  - **Greedy Policy** Simulation (each point is an average over 1k runs)
  - **Index Policy** Simulation (each point is an average over 1k runs)
  - **Optimal Policy** Computation (using Dynamic Programming)
- Two settings:
  - Symmetric Network
  - General Network

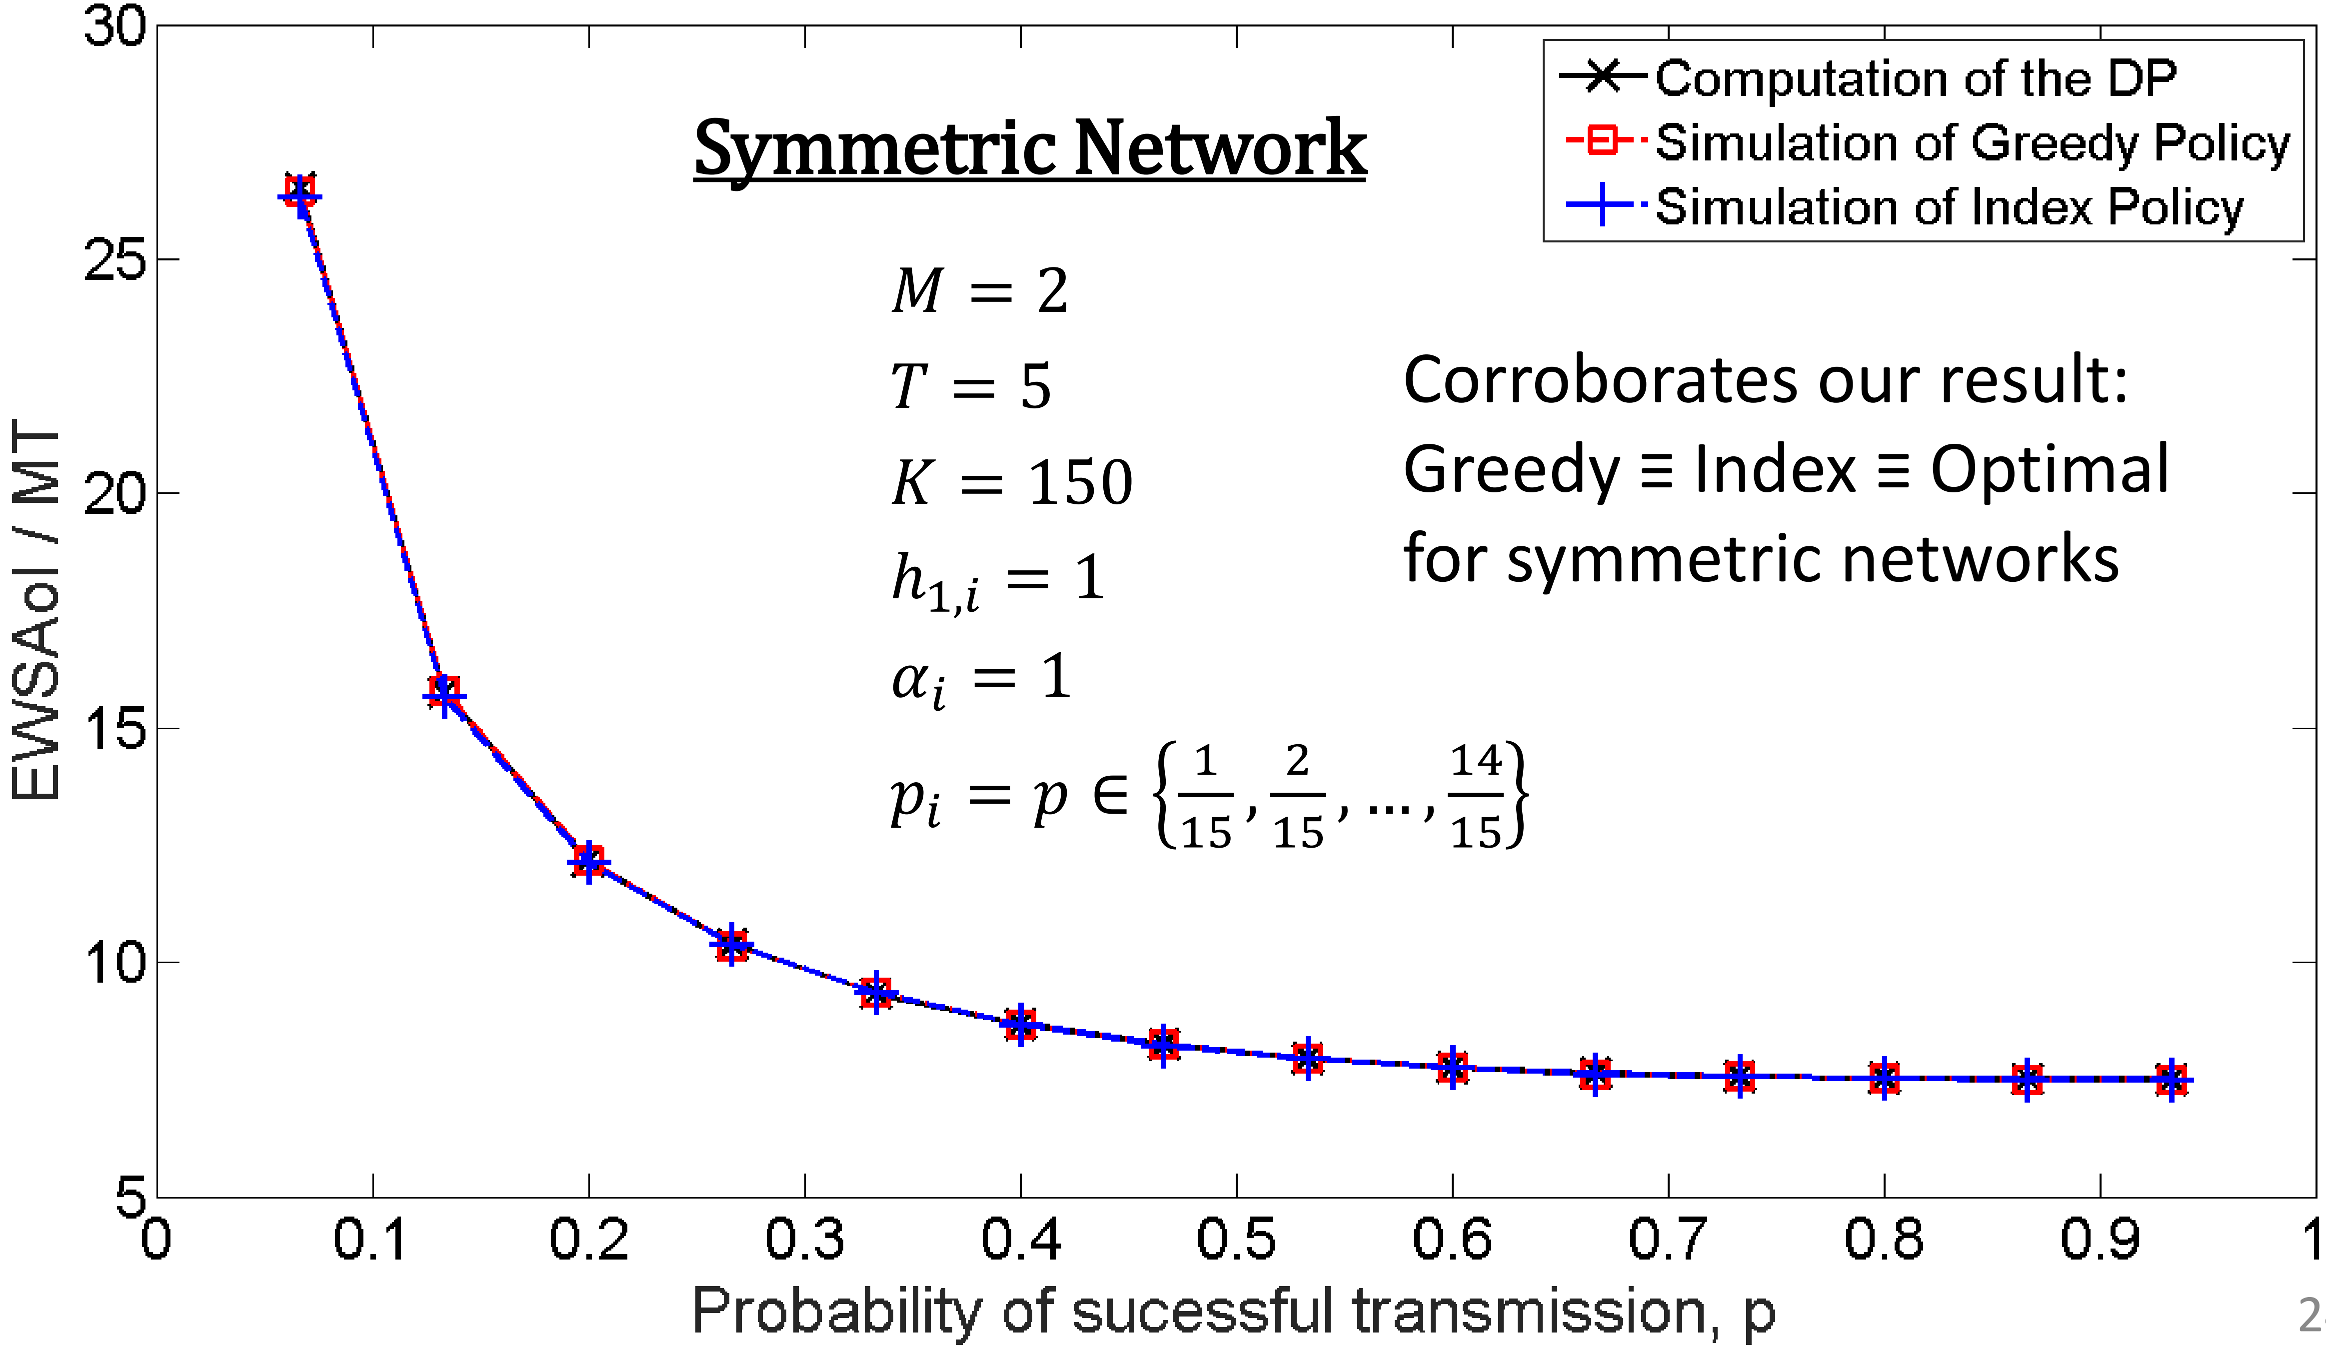
Symmetric Network
$M = 2$
$T = 5$
$K = 150$
$h_{1,i} = 1$
$\alpha_i = 1$
$p_i = p \in \left\{\frac{1}{15}, \frac{2}{15}, \ldots, \frac{14}{15}\right\}$
Corroborates our result:
Greedy ≡ Index ≡ Optimal
for symmetric networks
Computation of the DP
Simulation of Greedy Policy
Simulation of Index Policy
EWSAoI / MT
Probability of sucessful transmission, p
30
25
20
15
10
5
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1

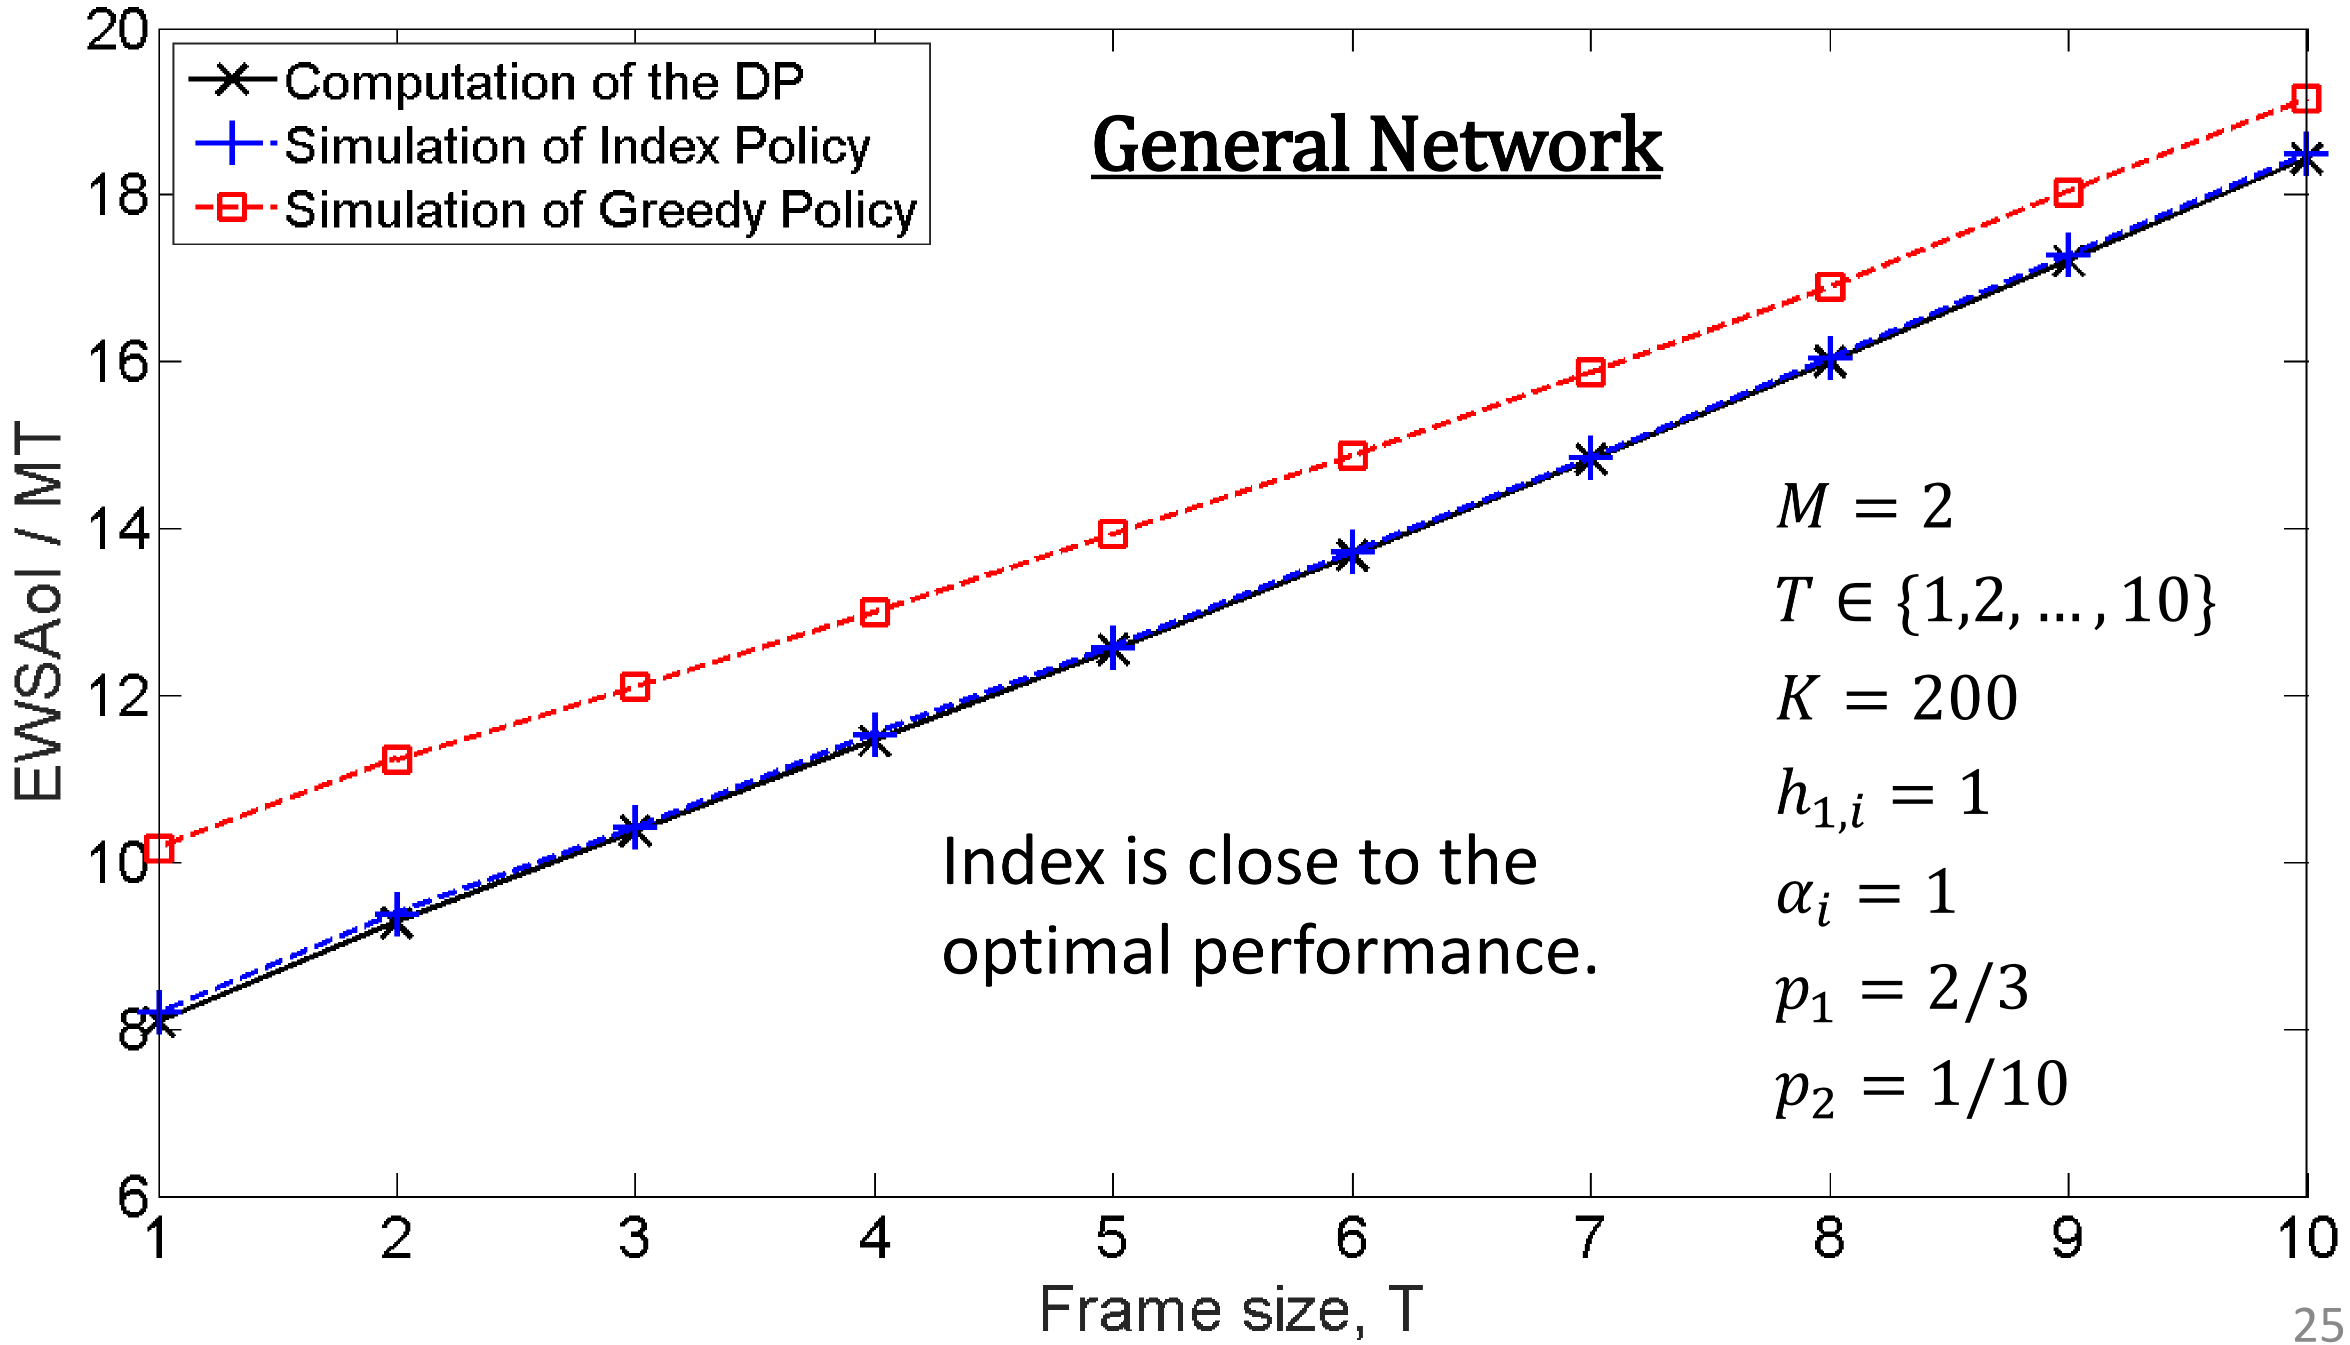
General Network
Computation of the DP
Simulation of Index Policy
Simulation of Greedy Policy
EWSAoI / MT
Frame size, T
20
18
16
14
12
10
8
6
1
2
3
4
5
6
7
8
9
10
Index is close to the optimal performance.
$M = 2$
$T \in \{1,2,\dots,10\}$
$K = 200$
$h_{1,i} = 1$
$\alpha_i = 1$
$p_1 = 2/3$
$p_2 = 1/10$

Minimum Delivery Ratio Constraint

- In our network setting, undelivered packets are replaced.
- Consider the problem of finding the scheduling policy $\eta \in \Pi$ that satisfies:

$$\mathbb{P}\left(\hat{q}_i^\eta \geq q_i\right) = 1, \forall i$$

where q_i is the minimum delivery ratio requirement of client i and $\hat{q}_i^\eta$ is:

$$\hat{q}_i^\eta \triangleq \liminf_{K \to \infty} \frac{1}{K} \sum_{k=1}^{K} \mathrm{e_i(k)} \quad \text{, where } \mathrm{e_i(k)} = \begin{cases} 1, & \text{if delivery (k, i)} \\ 0, & \text{otherwise} \end{cases}$$

- An equivalent problem is to find the η that maximizes the Expected Weighted Sum Throughput (next slide).

[8] I.-H. Hou, V. Borkar, and P. R. Kumar, "A theory of QoS for wireless," 2009.

Max Throughput vs Min AoI

- The **Throughput maximization** metric is given by:

$$\mathrm{EWST} = \frac{1}{K}\mathbb{E}\left\{\sum_{k=1}^{K}\sum_{i=1}^{M}\alpha_i\, e_i(k)\right\}, \text{ where } e_i(k) = \begin{cases} 1, & \text{if delivery } (k,i) \\ 0, & \text{otherwise} \end{cases}$$

- The **AoI minimization** metric is:

$$\mathrm{EWSAoI} = \frac{1}{K}\mathbb{E}\left\{\sum_{k=1}^{K}\sum_{i=1}^{M}\alpha_i\left(\frac{T^2}{2} + T^2 h_{k,i}\right)\right\}, \text{ where } h_{k+1,i} = \begin{cases} 1, & \text{if delivery } (k,i) \\ h_{k,i} + 1, & \text{otherwise} \end{cases}$$

- For comparing the scheduling policies that result from each problem, we consider their DP solutions.

Max Throughput vs Min AoI

- For a fixed vector of client weights $\vec{\alpha}$, the Dynamic Programs yield:

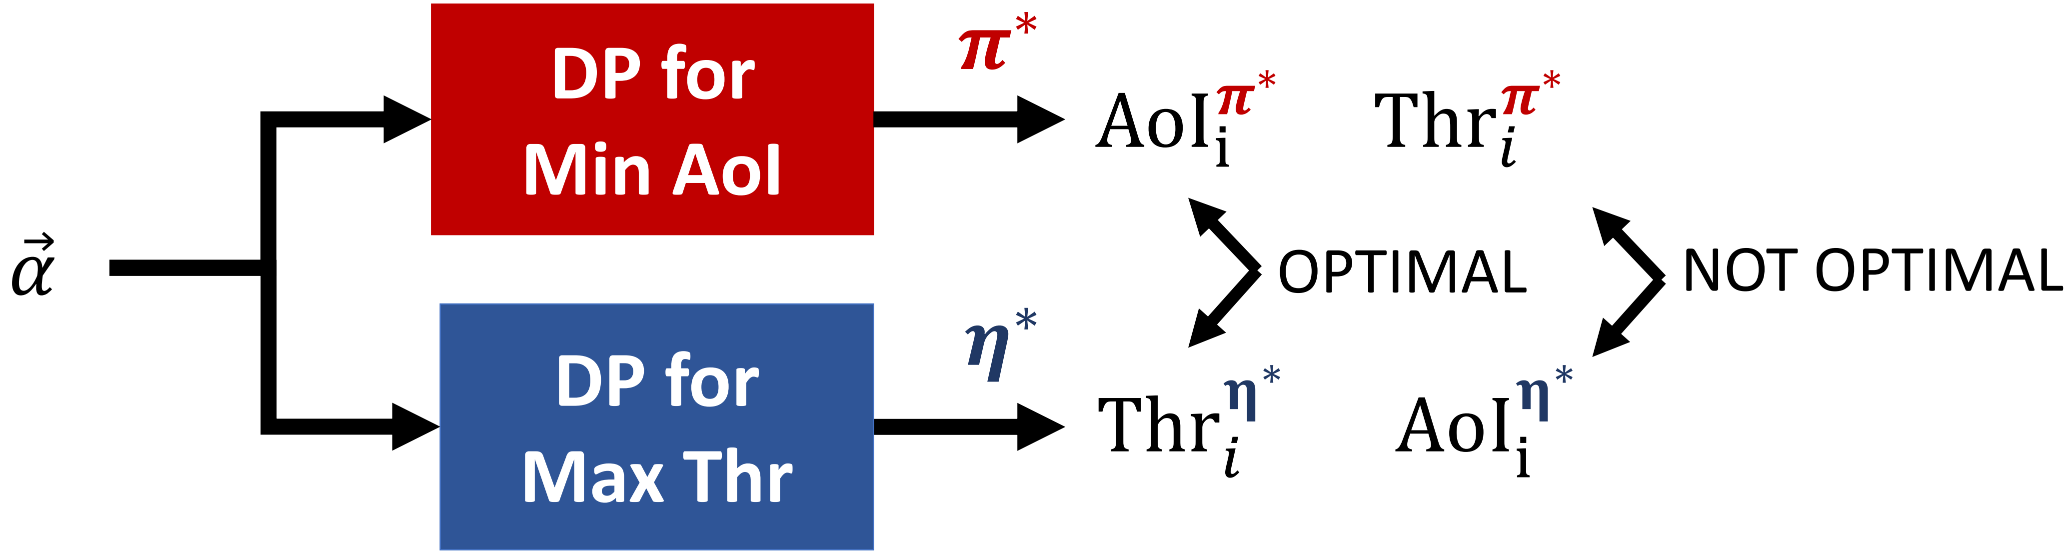

- We sweep $\vec{\alpha}$ and plot the results next:
  - **Red** is for metrics associated with $\boldsymbol{\pi}^*$.
  - **Blue** is associated with $\boldsymbol{\eta}^*$.

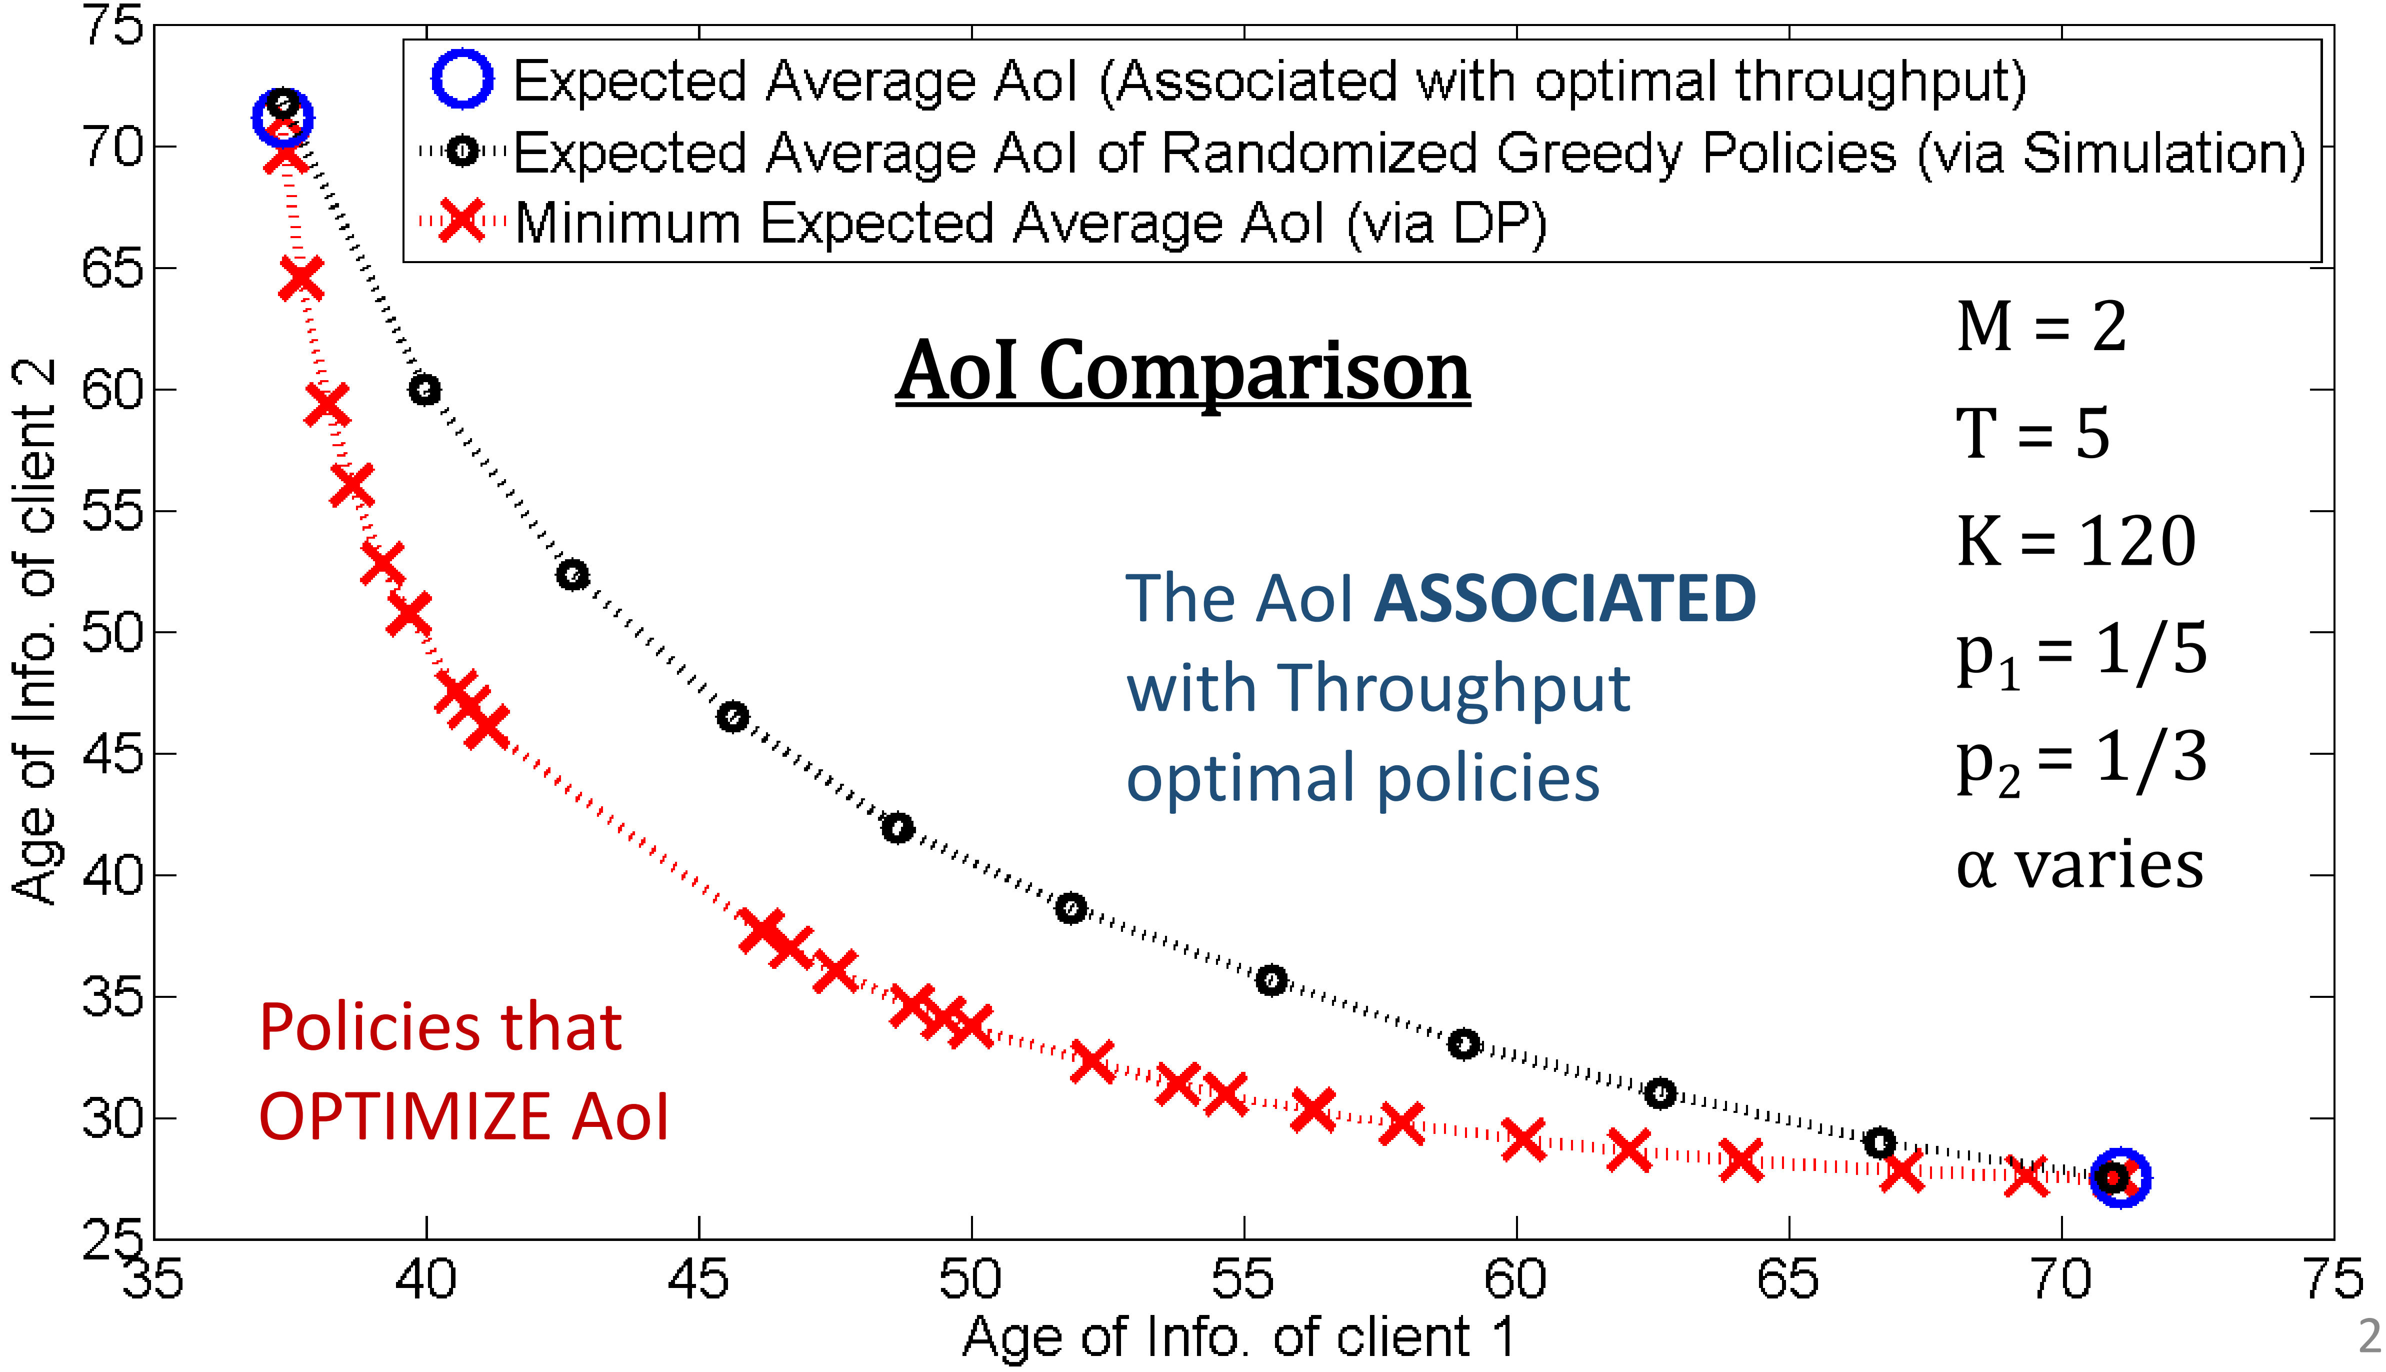
AoI Comparison
Expected Average AoI (Associated with optimal throughput)
Expected Average AoI of Randomized Greedy Policies (via Simulation)
Minimum Expected Average AoI (via DP)
M = 2
T = 5
K = 120
p1 = 1/5
p2 = 1/3
α varies
The AoI ASSOCIATED with Throughput optimal policies
Policies that OPTIMIZE AoI
Age of Info. of client 2
Age of Info. of client 1
75
70
65
60
55
50
45
40
35
30
25
35
40
45
50
55
60
65
70
75

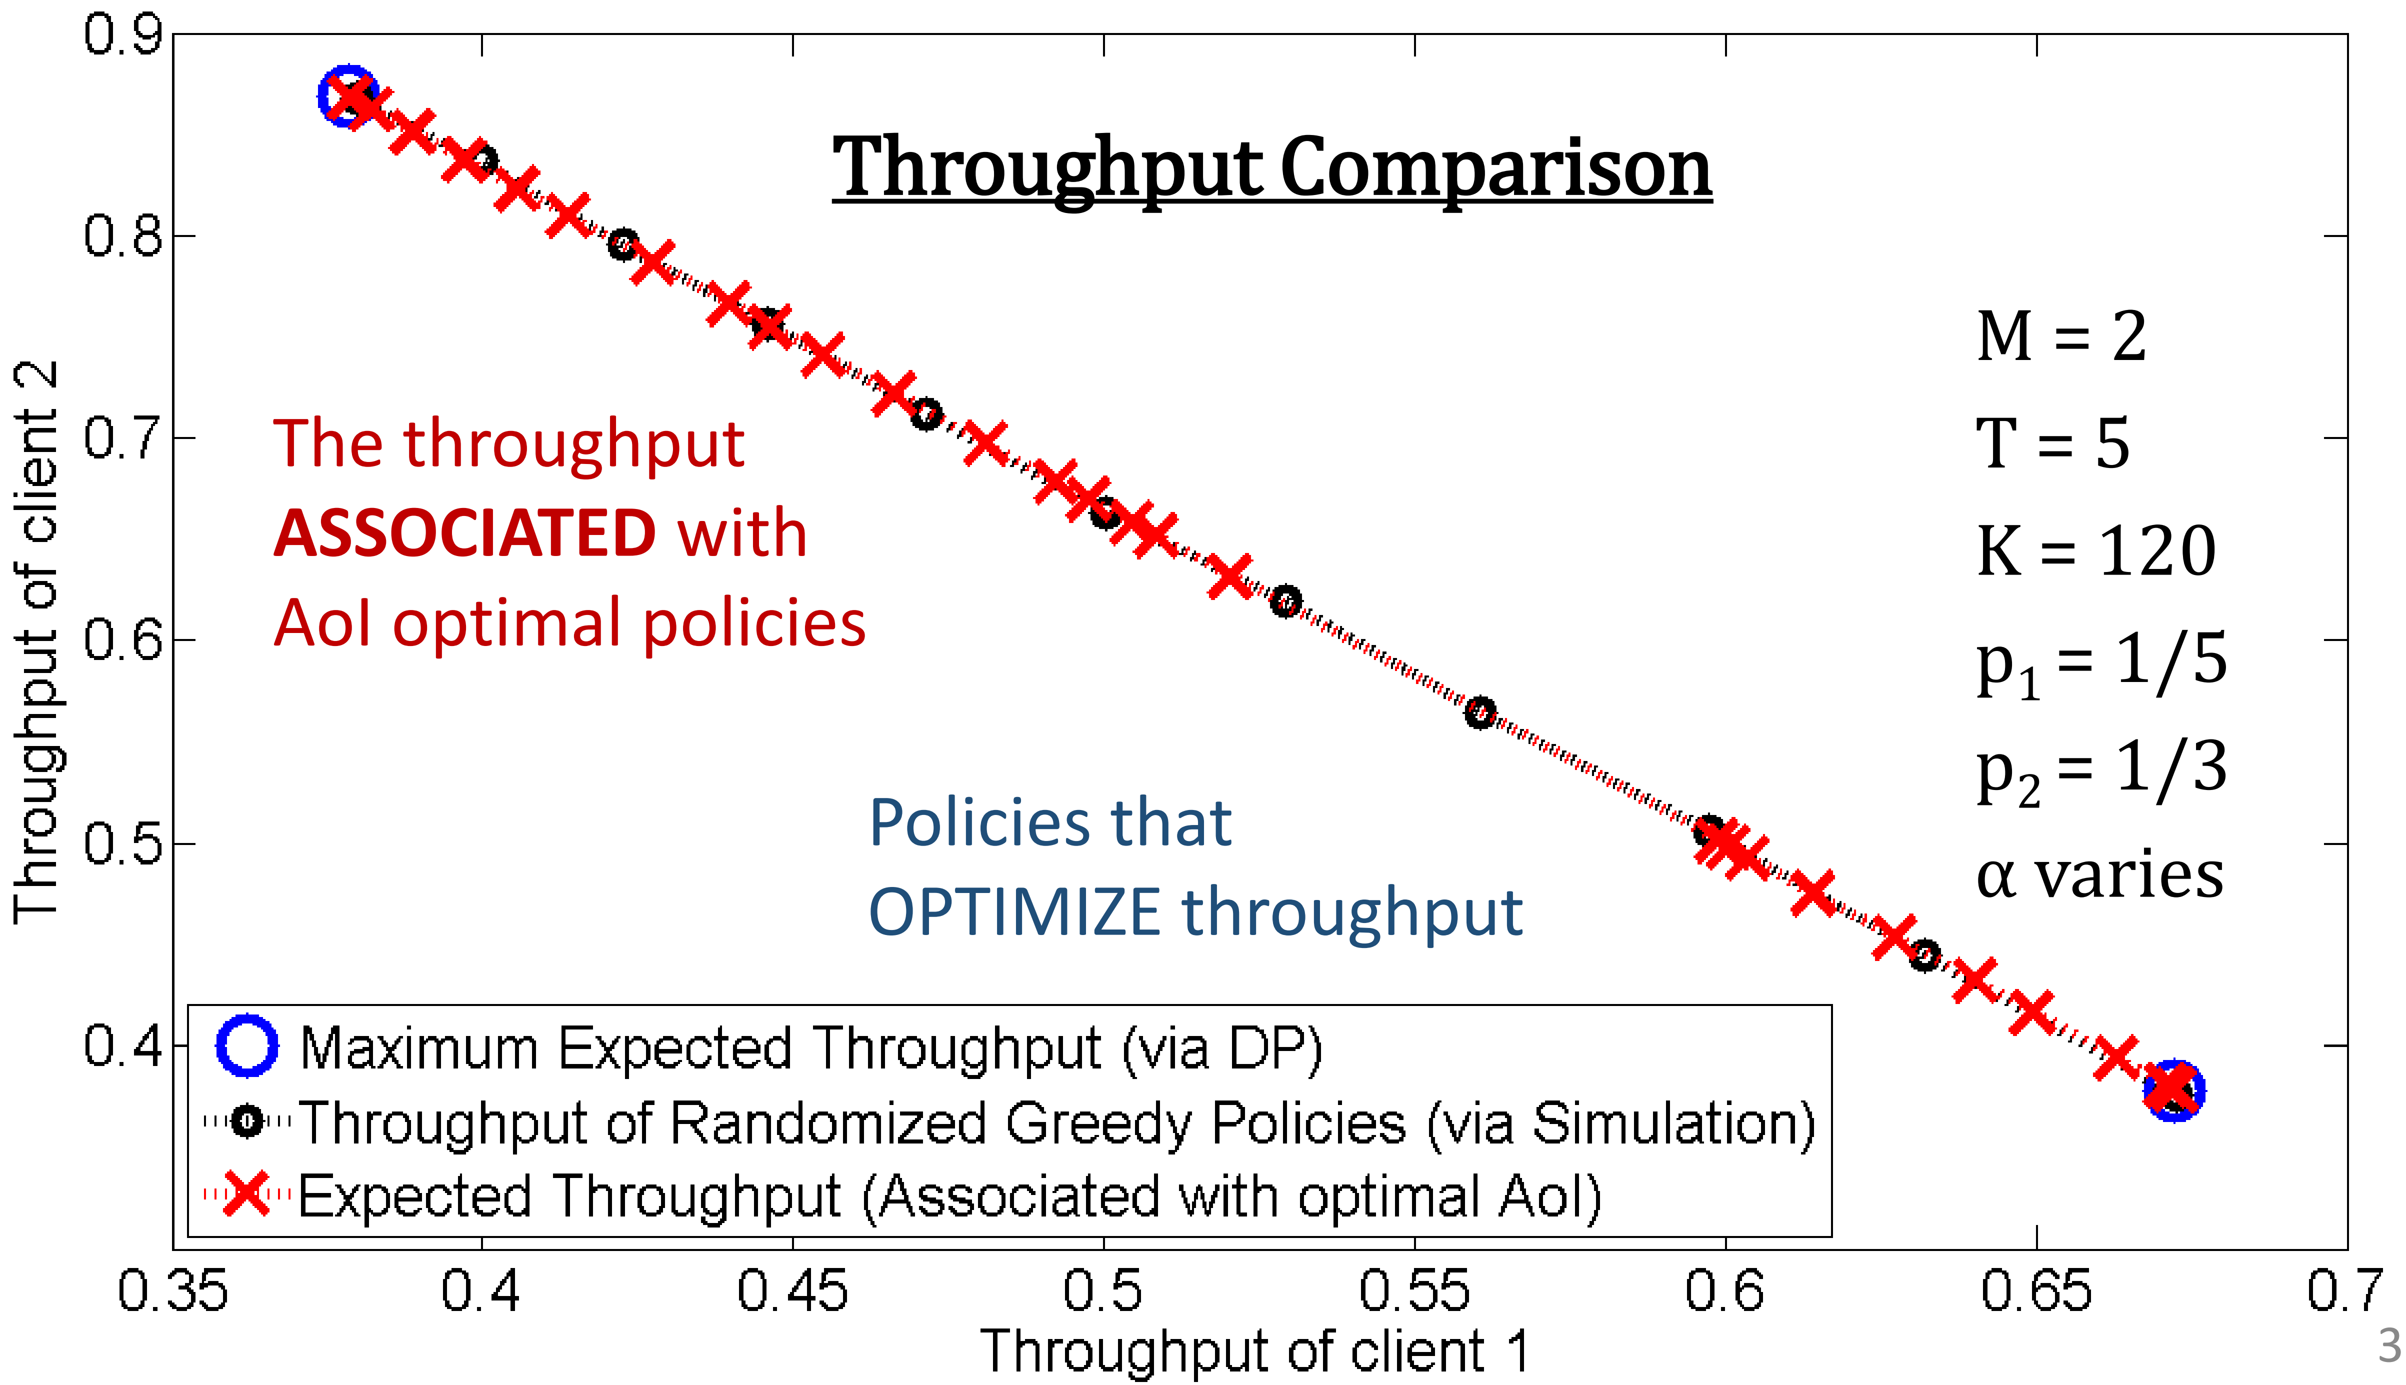

Throughput Comparison
M = 2
T = 5
K = 120
p1 = 1/5
p2 = 1/3
α varies
The throughput
ASSOCIATED with
AoI optimal policies
Policies that
OPTIMIZE throughput
Maximum Expected Throughput (via DP)
Throughput of Randomized Greedy Policies (via Simulation)
Expected Throughput (Associated with optimal AoI)
Throughput of client 2
Throughput of client 1
0.9
0.8
0.7
0.6
0.5
0.4
0.35
0.4
0.45
0.5
0.55
0.6
0.65
0.7

Max Throughput vs Min AoI

- The conclusion illustrated by the numerical results holds in general:

Thr Optimal Policies $\nRightarrow$ AoI Optimal
AoI Optimal Policies $\Rightarrow$ Thr. Optimal[1]

[1] Pareto optimal

- By minimizing the AoI, we are assured to achieve maximum throughput.

- However, it is still not known how to design a scheduling policy that achieves a **given throughput** with minimum AoI.

AoI, Throughput and Interdelivery times

- **Proposition**: Consider the network over an infinite horizon, namely $K \to \infty$, and assume that the steady-state distribution of the underlying MC exists when the stationary policy π is employed. Then, it follows

$$\mathrm{EWSAoI} = \mathrm{T}^2 \sum_{\mathrm{i}=1}^{\mathrm{M}} \alpha_{\mathrm{i}} + \frac{\mathrm{T}^2}{2} \sum_{\mathrm{i}=1}^{\mathrm{M}} \alpha_{\mathrm{i}} \left(\frac{\mathbb{V}ar[I_i]}{\mathbb{E}[I_i]} + \mathbb{E}[I_i] \right)$$

where I_i is the r.v. that represents the number of frames in the interval between two packet deliveries from client i, i.e. the interdelivery time.

Consider the AoI optimal policy π^* and the associated throughput performance. Under the conditions of the Proposition, we know that from all policies with the same throughput, policy π^* achieves the lowest value of $\mathbb{V}ar[I_i]$.

[11] B.Li, et al., “Throughput-optimal scheduling design with regular service guarantees in wireless networks,” 2015.

Outline / Contributions

- Age of Information (AoI) and Network Model
- Symmetric Network and the **OPTIMALITY** of the Greedy Policy
- General Network and the **DESIGN** of the Index Policy
- **VALIDATION** of the policies via Numerical Results
- **COMPARISON** of the Min AoI problem and the Max Throughput problem.

Supplementary Slides

Intuition of the proof: ideal channels

- Consider $M = 3, T = 1$ and $p = 1$ (ideal channels)
- Employ GREEDY policy. **Deliveries** are in green.

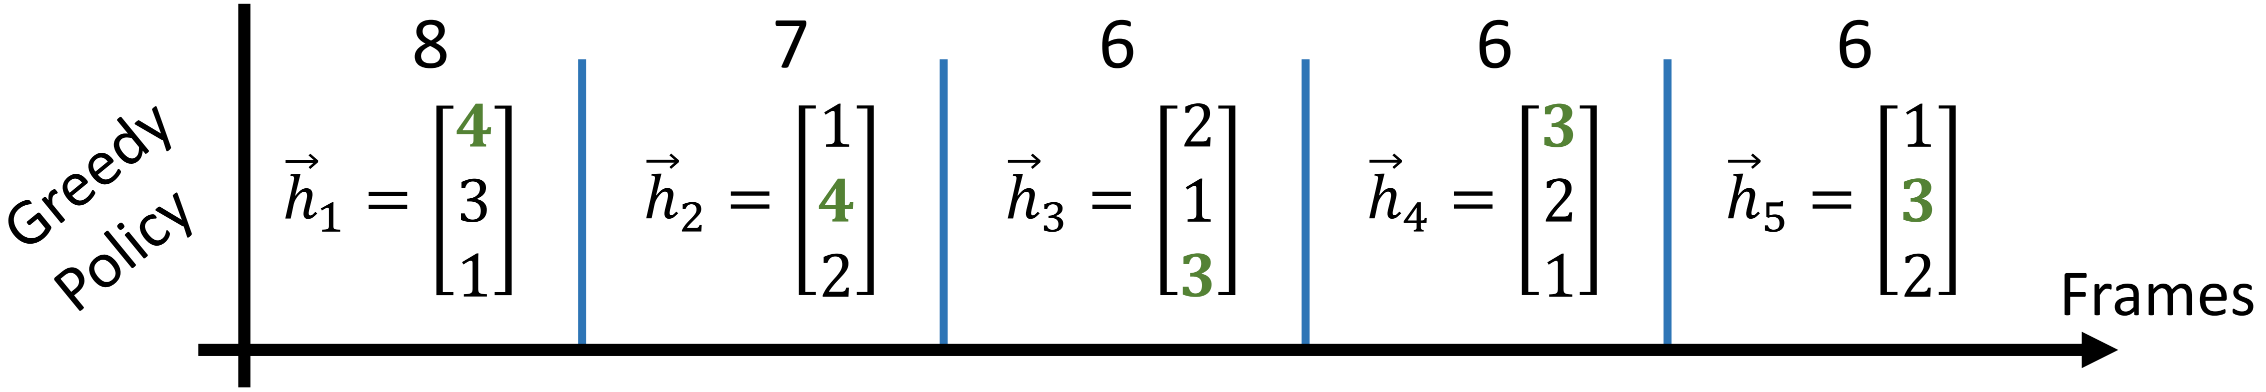

- One packet is scheduled and delivered in each frame ($T = 1$ slot).
- Greedy achieves the lowest $\sum_{i=1}^{M} \mathrm{h}_{\mathrm{k,i}}$ in every frame k → Greedy is optimal.
- Note that $\mathrm{h}_{\mathrm{k,1}} + \mathrm{h}_{\mathrm{k,2}} + \mathrm{h}_{\mathrm{k,3}} = 6, \forall k \geq 3$ (steady-state)

Intuition of the proof: coupling argument

- Consider $M = 3, T = 1$ and $p \in (0,1]$ **(unreliable channels)**
- Employ **ARBITRARY policy**. **Deliveries** are green. **Failed** transmissions are red.

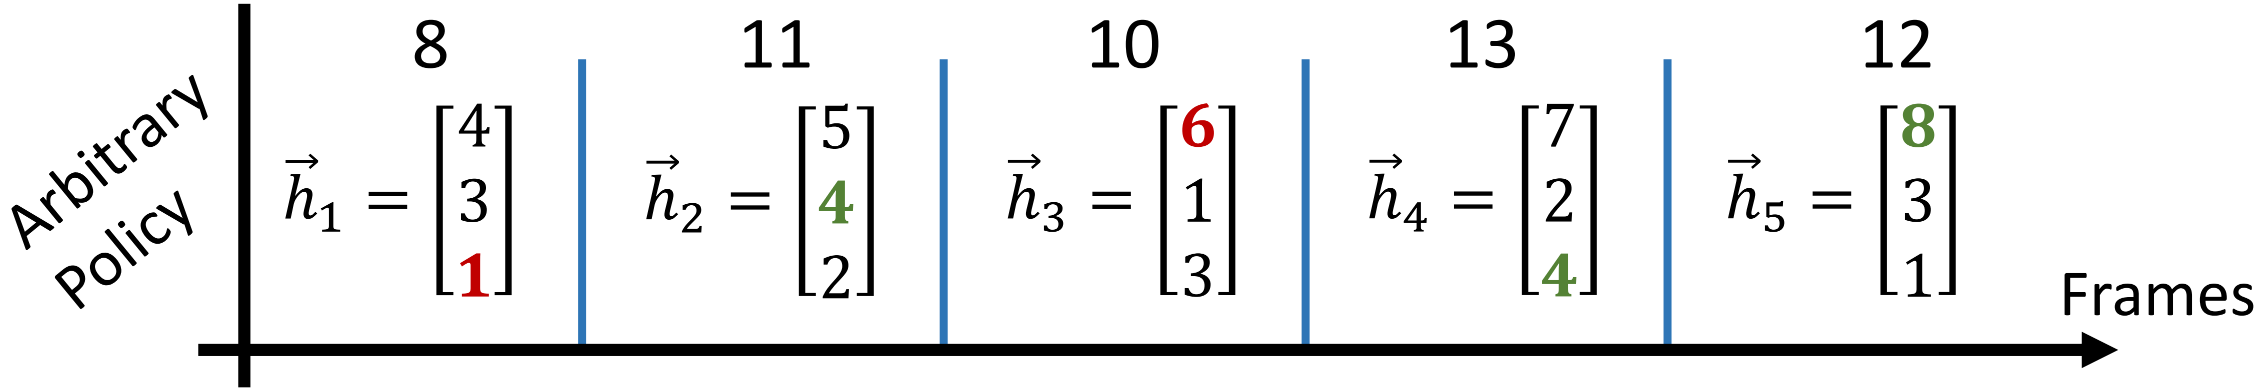

- Fix any sample path for the state of the active channel:
  - channel is OFF: no room for improvement. All policies are equivalent.
  - channel is ON ≡ ideal channels: the best policy is Greedy. (example next)

Intuition of the proof: coupling argument

- Consider $M = 3, T = 1$ and $p \in (0,1]$ **(unreliable channels)**
- Employ **ARBITRARY policy**. **Deliveries** are green. **Failed** transmissions are red.