Optimizing Age of Information in Wireless Networks with Throughput Constraints

Igor Kadota

SQUALL Seminar, Dec 5, 2017

ACK: Eytan Modiano, Elif Uysal-Biyikoglu, Abhishek Sinha, Rahul Singh
To appear in INFOCOM 2018.

Outline

- Age of Information (AoI)
- Network Model
- Scheduling Policies for AoI
- Discussion: Throughput versus AoI
- Scheduling Policies for AoI with Throughput Requirements
- Final Remarks

Age of Information (AoI)

Example:

Delivery of Packets to BS

Packet Generation at i

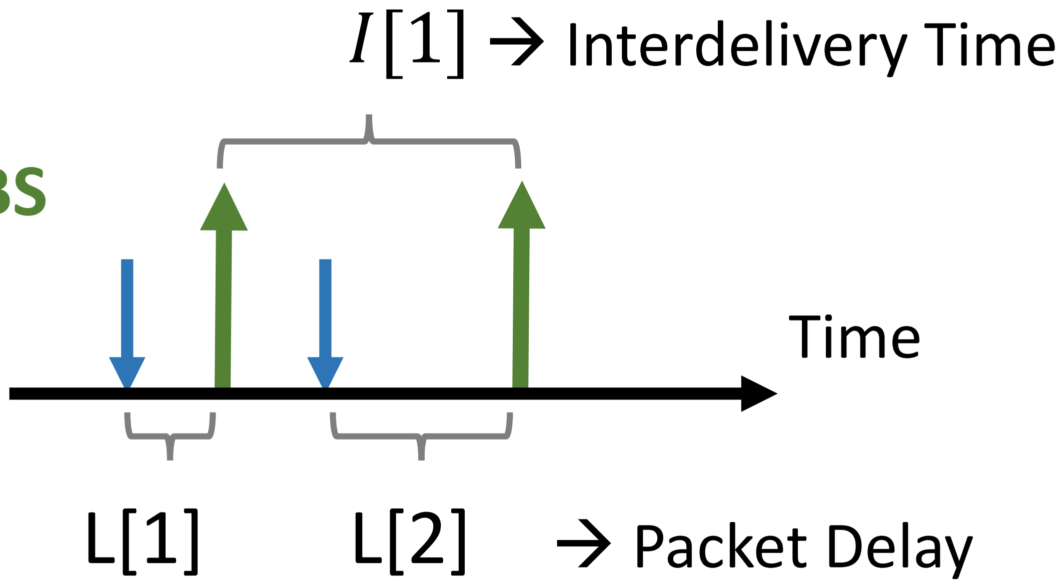

Single Source

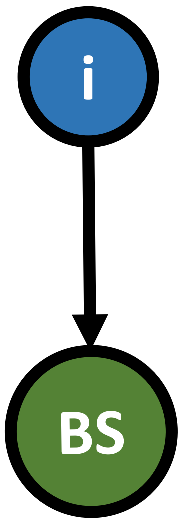

Single Destination

How fresh is the information at the destination?

Age of Information (AoI)

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

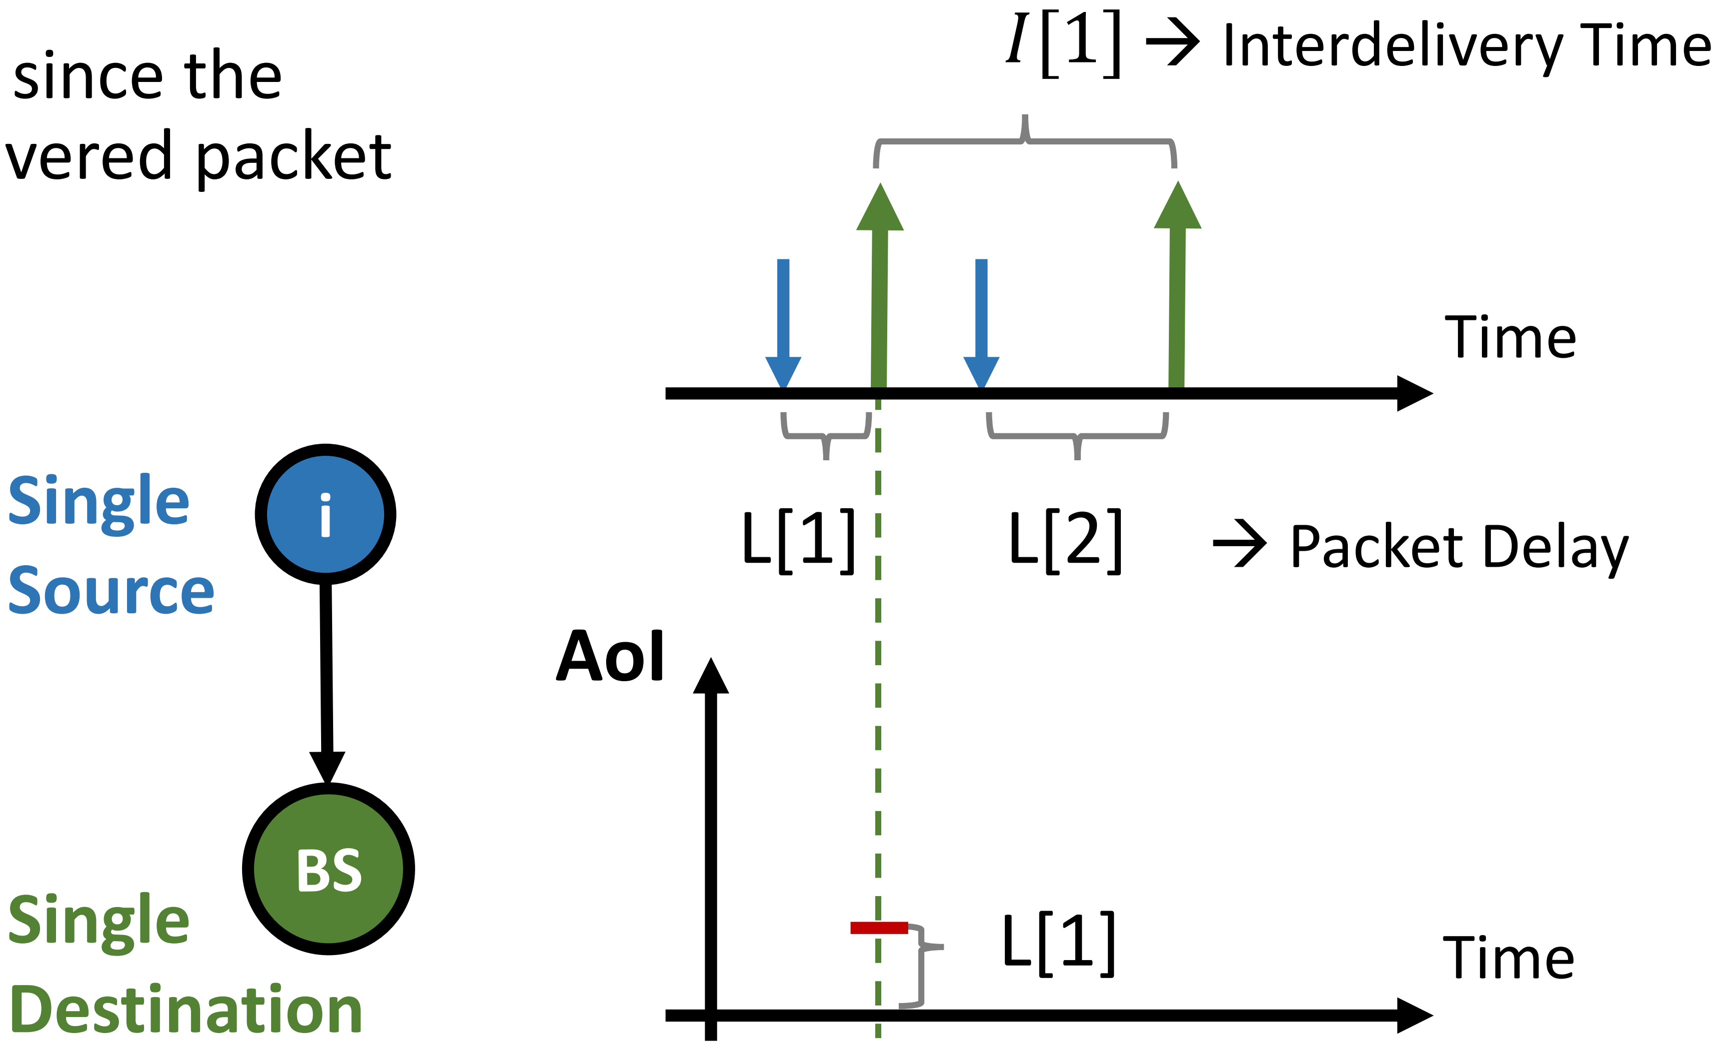

Age of Information (AoI)

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

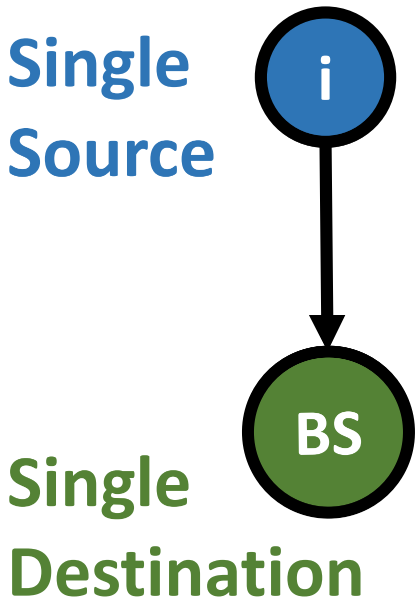

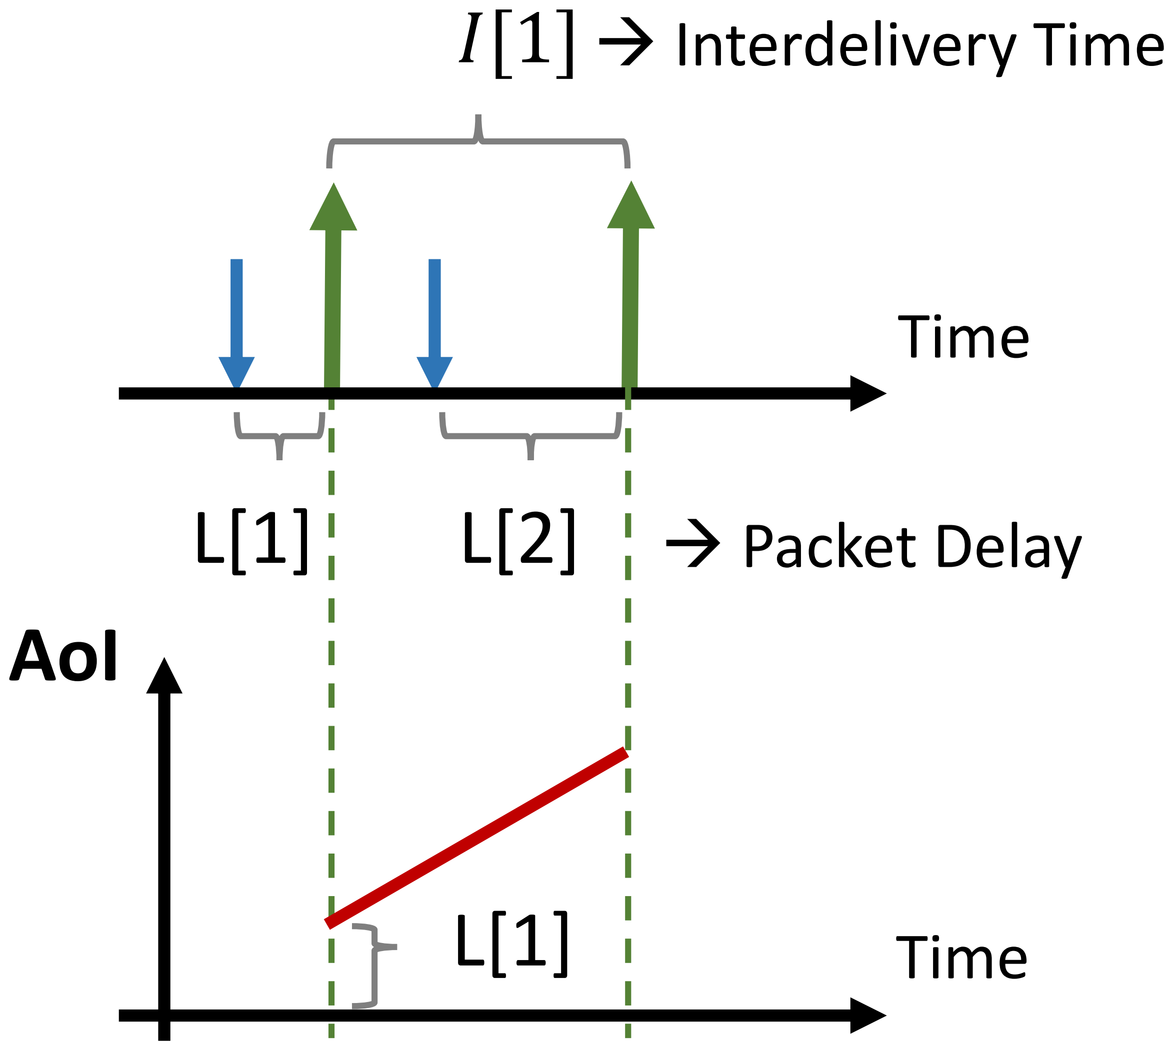

Age of Information (AoI)

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

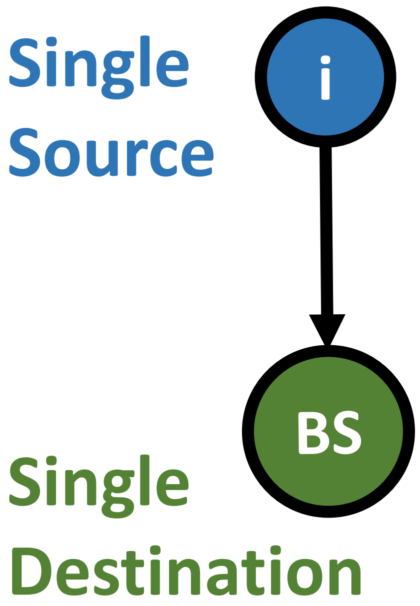

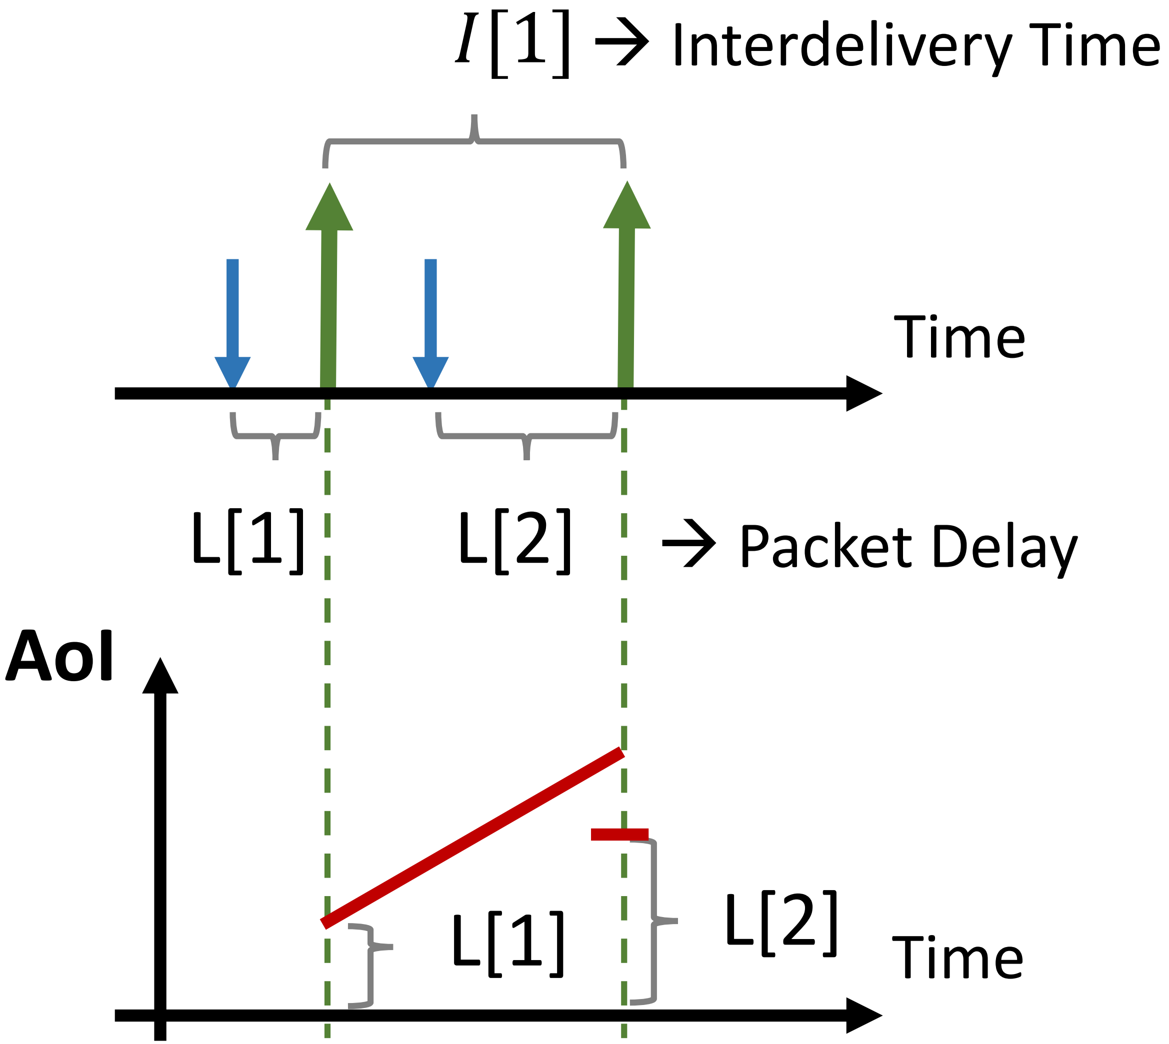

Age of Information (AoI)

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.

Relation between AoI, delay and interdelivery time?

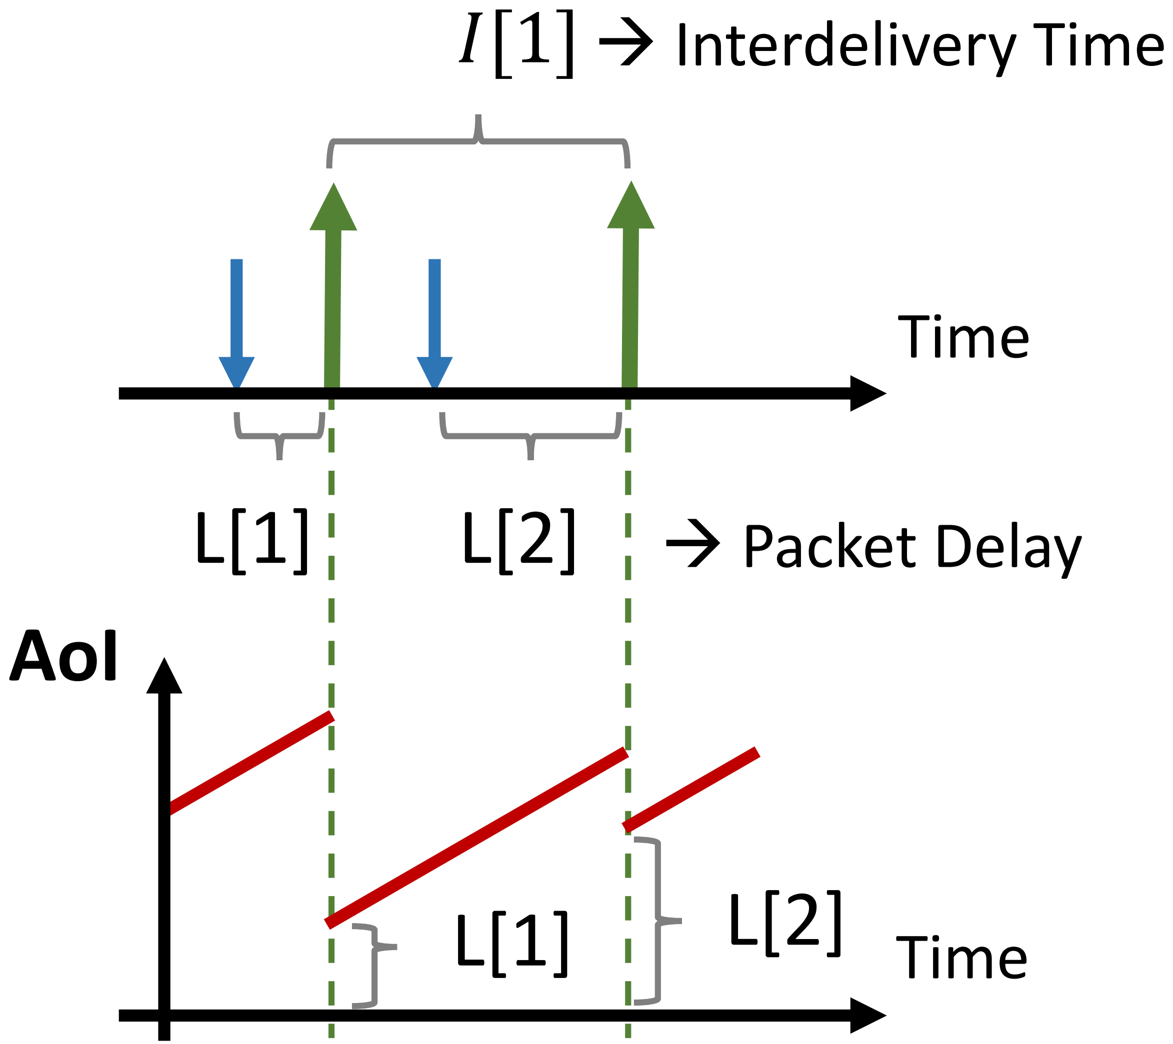

AoI, Delay and Interdelivery time

- 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	

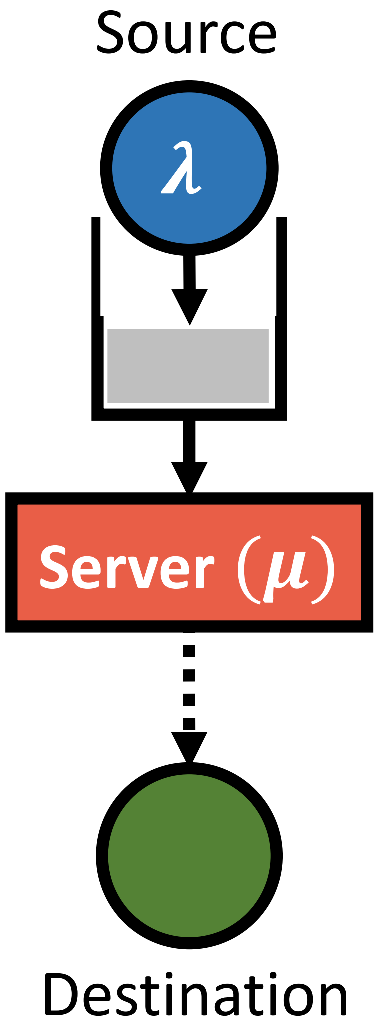

[1] S. Kaul, R. Yates, and M. Gruteser, “Real-time status: How often should one update?”, 2012.

AoI, Delay and Interdelivery time

- <u>Example</u>: 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**

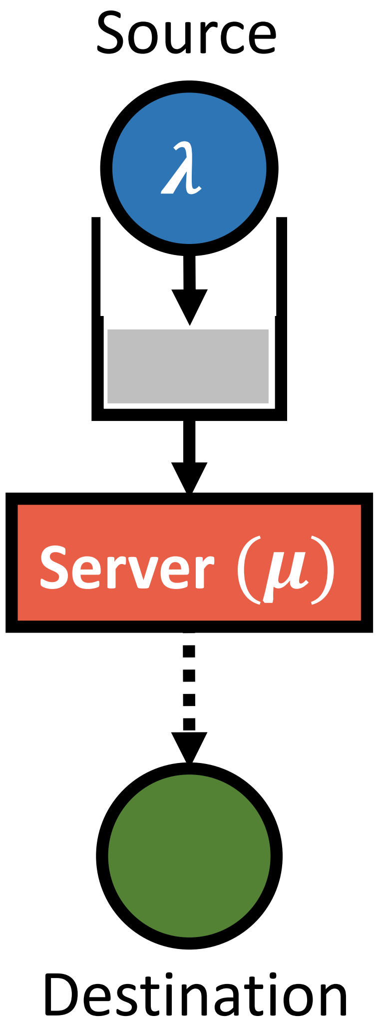

Low time-average AoI when packets with **<u>low delay</u>** are **<u>delivered regularly</u>**.

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

Network Model

Network - Example

Wireless Parking Sensor

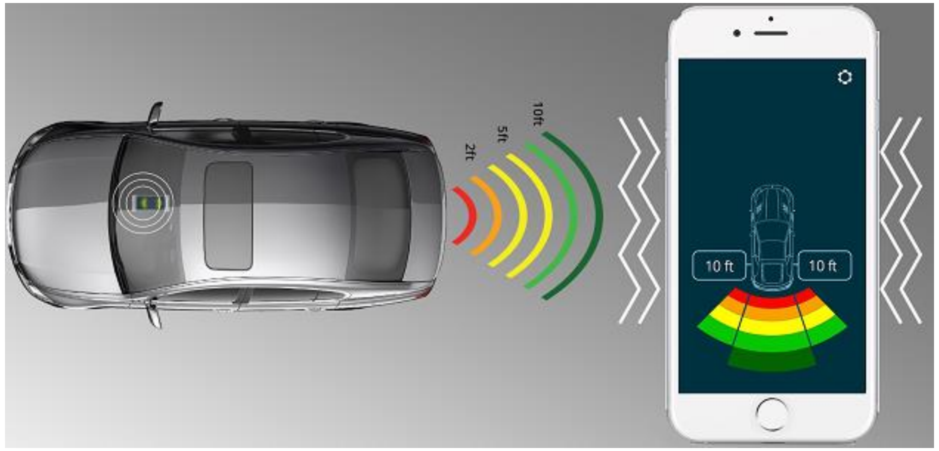

Wireless Rearview Camera

Wireless Tire-Pressure Monitoring System

Network - Description

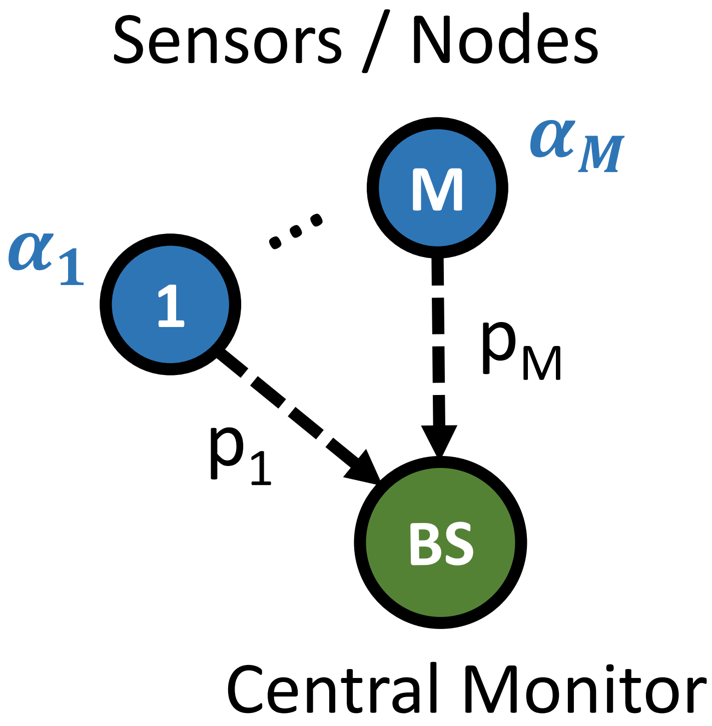

1) Central Monitor requires **fresh data** (low AoI)
2) Some sensors are more **important** than others (weights $\boldsymbol{\alpha_i}$)
3) Channel is **shared** and **unreliable** (probability of successful transmission $\boldsymbol{p_i}$)

Network - Scheduling Policy

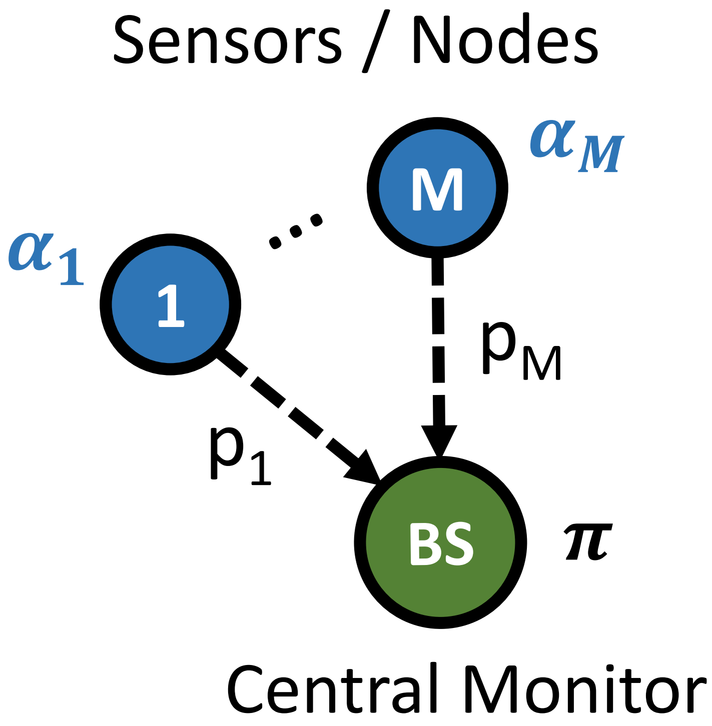

Scheduling decision during slot k:

1) **BS selects** a single node i
   $[u_i(k) = 1$ and $u_j(k) = 0, \forall j \neq i]$

2) Selected node **samples** new data and then **transmits** [packet delay = 1 slot]

3) Packet is **delivered** to the BS wp p_i
   $[d_i(k) = 1$ and $d_j(k) = 0, \forall j \neq i]$

Class of non-anticipative policies Π. **Arbitrary policy $\boldsymbol{\pi} \in \boldsymbol{\Pi}$.**

Network - Performance Metric

- Age of Information associated with node i at the beginning of slot k is given by $h_i(k)$.

- **Recall:** selected node samples new data and then transmits [packet delay = 1 slot]

- Evolution of AoI:

$$\mathbf{h_i}(k+1) = \begin{cases} 1, & \text{if } d_i(k) = 1 \\ h_i(k) + 1, & \text{otherwise} \end{cases}$$

Delivery of packets from sensor i to the BS

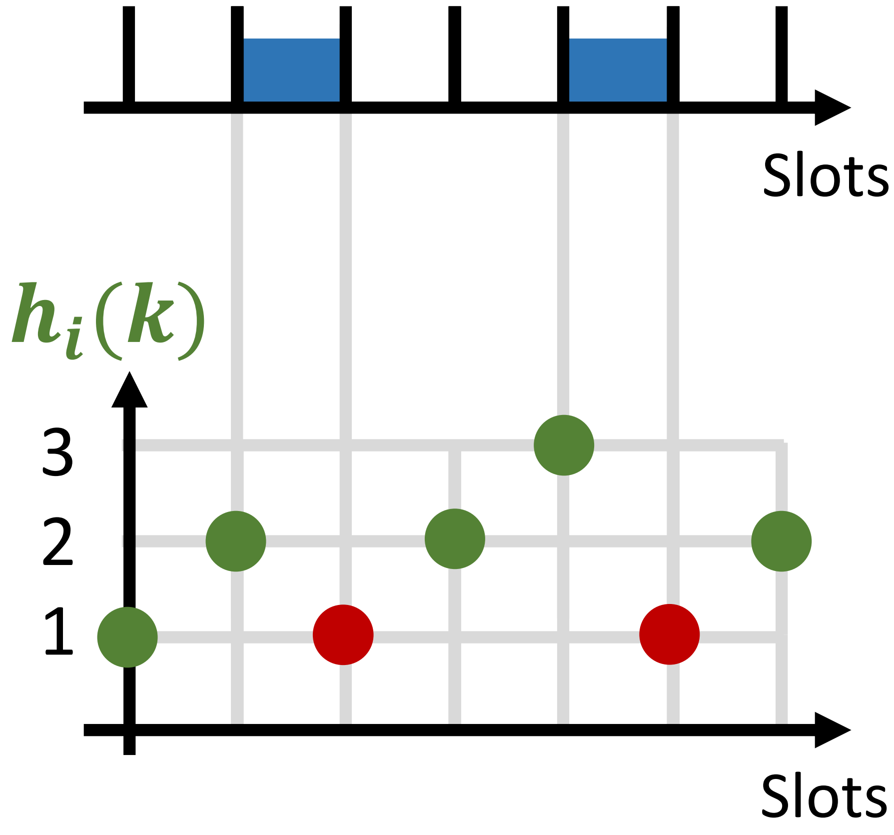

Network - Objective Function

- Expected Weighted Sum Age of Information when policy π is employed:

$$\mathbb{E}[\, J_K^{\pi} \,] = \frac{1}{KM}\,\mathbb{E}\left[\sum_{k=1}^{K}\sum_{i=1}^{M} \alpha_i h_i^{\pi}(k)\right]$$

, where $h_i^{\pi}(k)$ is the AoI of node i and α_i is the positive weight

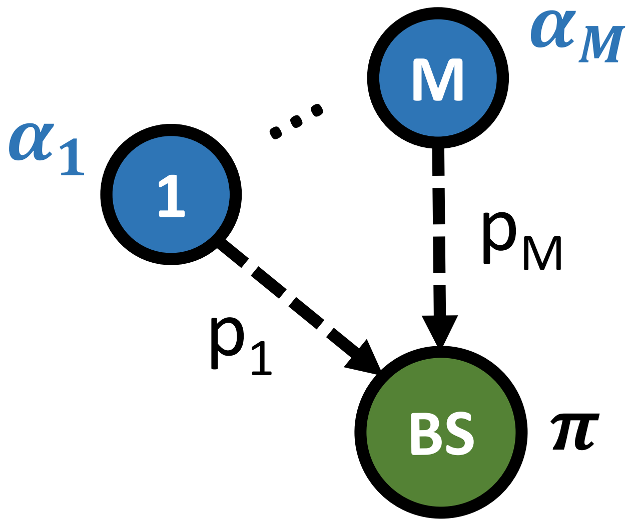

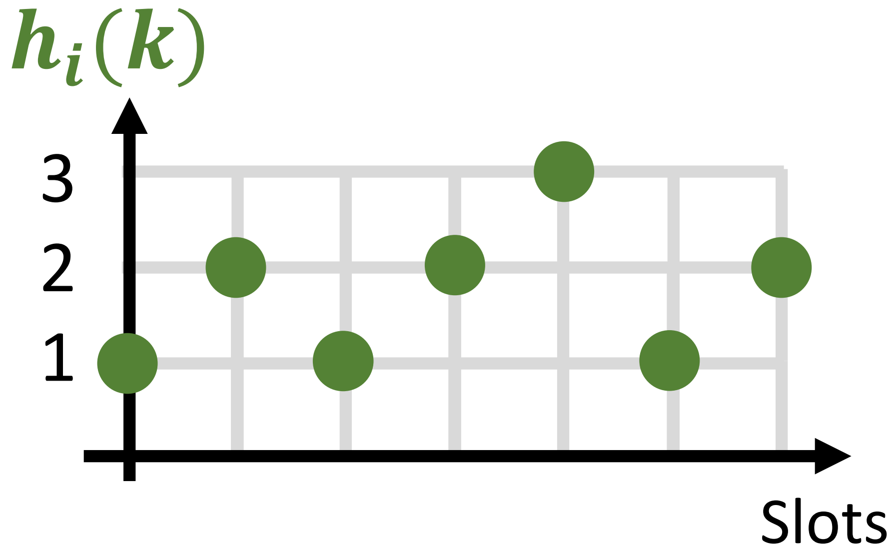

Network - Challenges

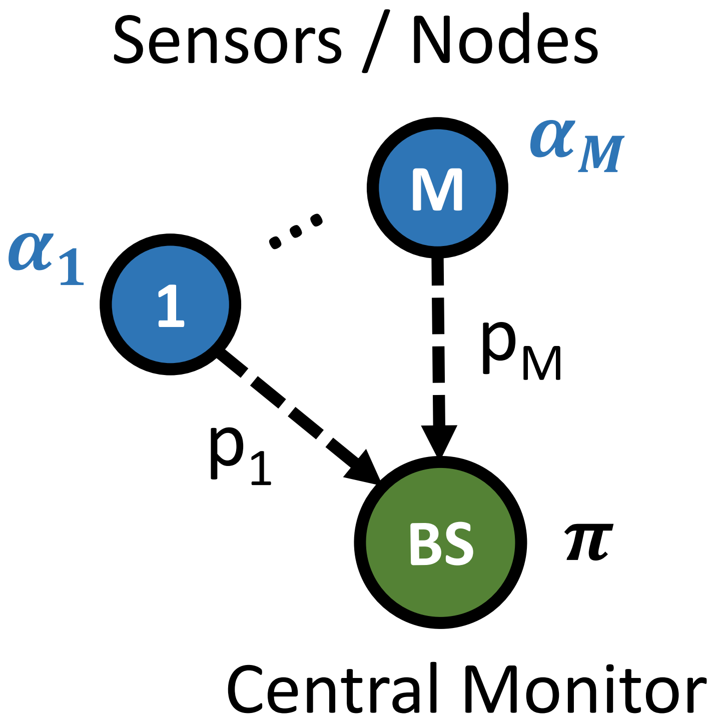

1) Scheduling policies must be **low-complexity** in order to be meaningful

2) **Evolution of AoI** is not simple

3) **Unreliability** of the wireless channel makes scheduling more challenging

Optimal Scheduling Policy for **Symmetric** Networks

Obs.: symmetry when $p_i = p \in (0,1]$ and $\alpha_i = \alpha > 0, \forall i.$

Optimality of Greedy

- **Greedy Policy**: in slot k, select the node with highest value of $\boldsymbol{h_i(k)}$.

Theorem: for **any** symmetric network with $p \in (0,1]$ and $\alpha > 0$, the Greedy policy attains the minimum expected sum AoI, i.e.

$$\text{Greedy} = \underset{\pi \in \Pi}{\operatorname{argmin}} \frac{\alpha}{KM} \mathbb{E}\left[\sum_{k=1}^{K} \sum_{\boldsymbol{i=1}}^{\boldsymbol{M}} \boldsymbol{h_i(k)}\right]$$

Intuition of the proof: ideal channels

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

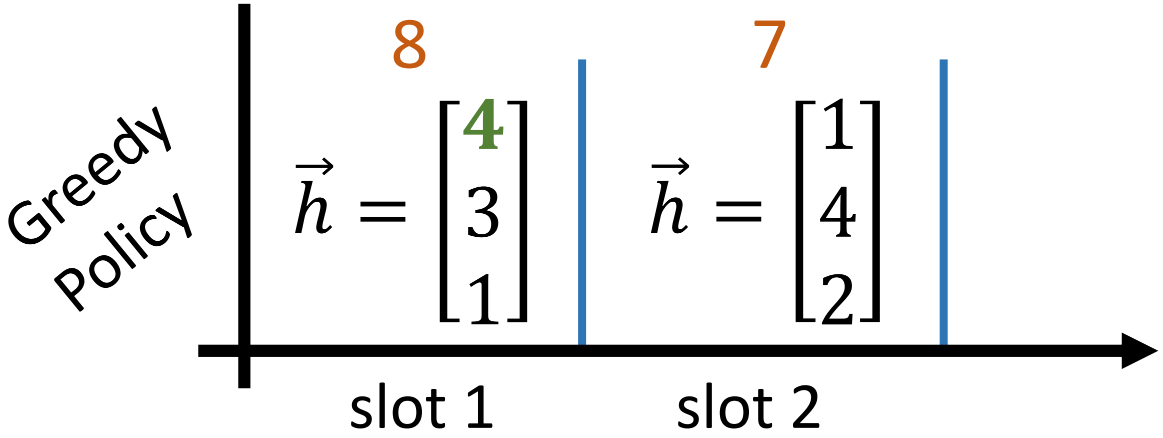

$$\vec{h} = \begin{bmatrix} h_{\mathbf{1}} \\ h_{\mathbf{2}} \\ h_{\mathbf{3}} \end{bmatrix}$$

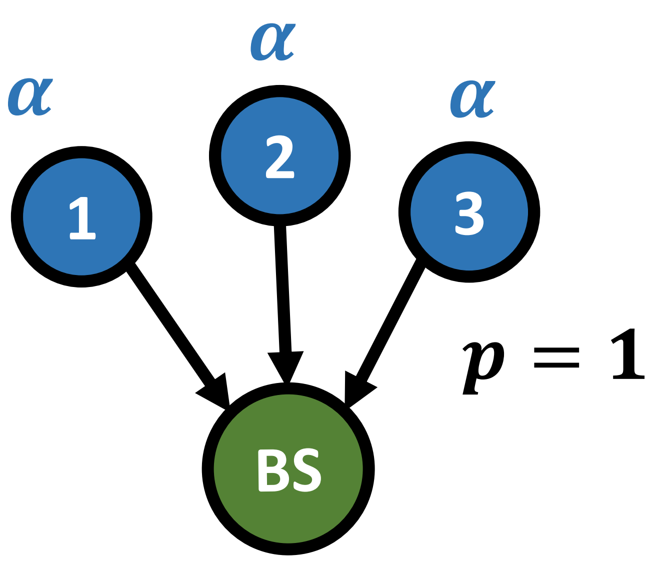

Greedy Policy

Intuition of the proof: ideal channels

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

Greedy Policy

	slot 1	slot 2	slot 3	slot 4	slot 5
$\sum h$	8	7	6	6	6
	$\vec{h} = \begin{bmatrix} \mathbf{4} \\ 3 \\ 1 \end{bmatrix}$	$\vec{h} = \begin{bmatrix} 1 \\ \mathbf{4} \\ 2 \end{bmatrix}$	$\vec{h} = \begin{bmatrix} 2 \\ 1 \\ \mathbf{3} \end{bmatrix}$	$\vec{h} = \begin{bmatrix} \mathbf{3} \\ 2 \\ 1 \end{bmatrix}$	$\vec{h} = \begin{bmatrix} 1 \\ \mathbf{3} \\ 2 \end{bmatrix}$

- Greedy achieves the lowest $\sum_{i=1}^{M} h_i(k)$ in **every** slot k → Greedy is optimal.

Intuition of the proof: coupling argument [3]

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

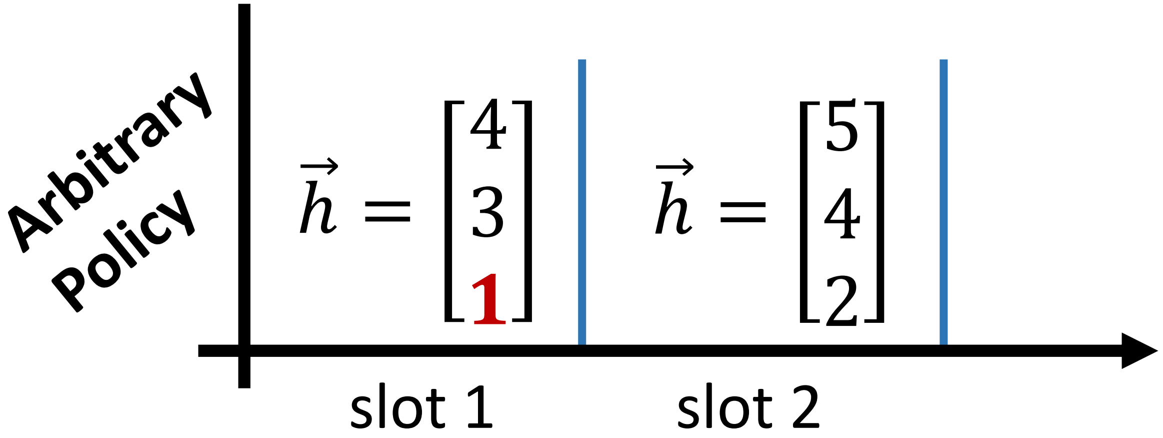

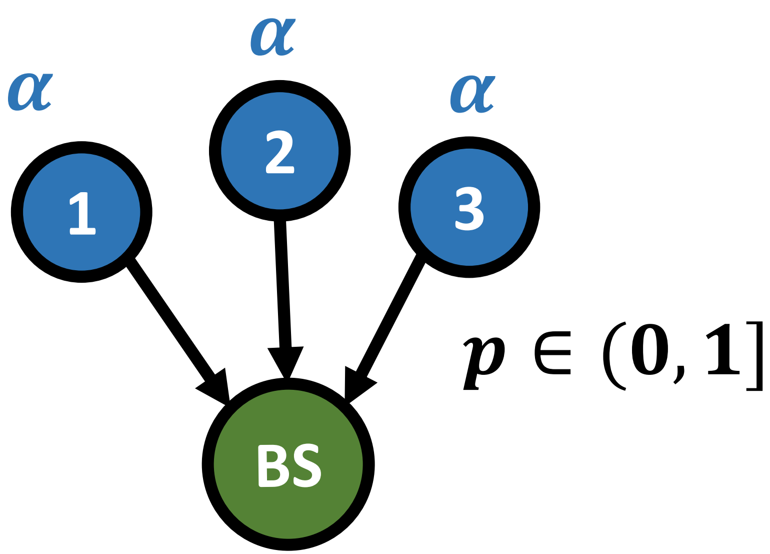

Arbitrary Policy

[3] D. Stoyan, **Comparison Methods for Queues and other Stochastic Models,** 1983

Intuition of the proof: coupling argument [3]

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

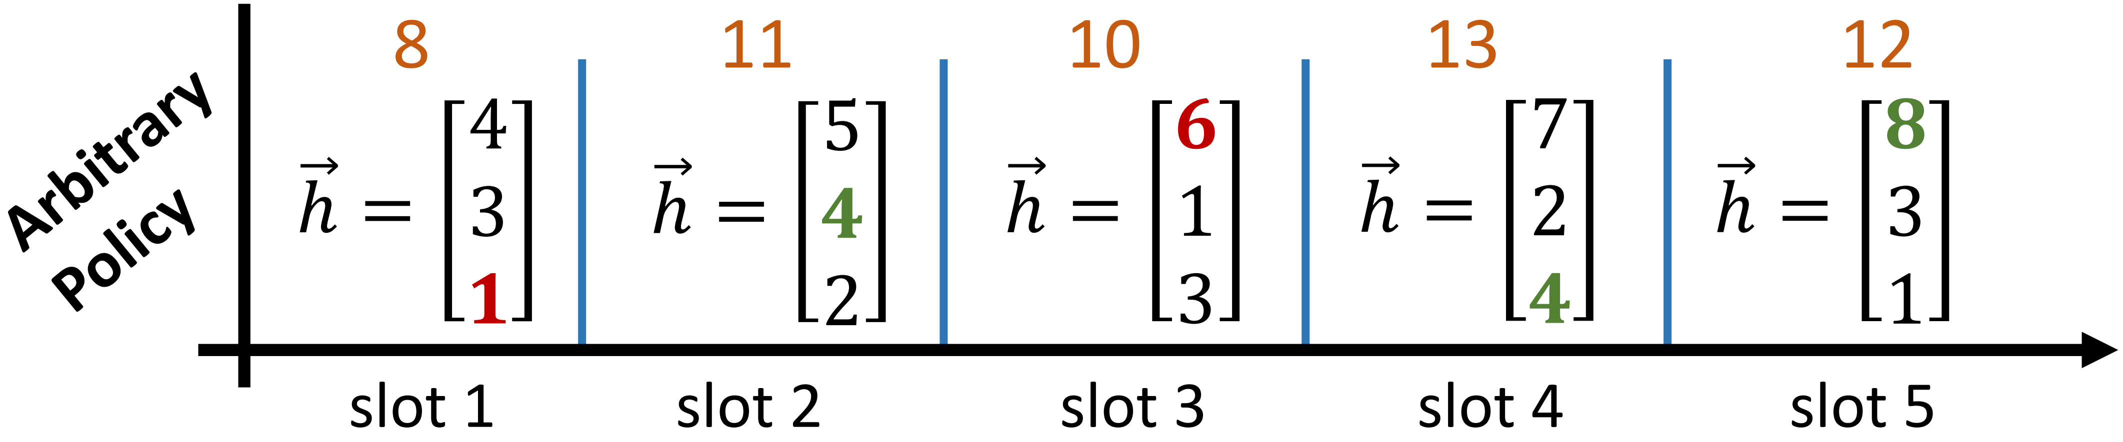

- Goal is to show that Greedy achieves the lowest $\sum_{i=1}^{M} h_i(k)$ in **every** slot k.

[3] D. Stoyan, **Comparison Methods for Queues and other Stochastic Models,** 1983

Intuition of the proof: coupling argument [3]

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

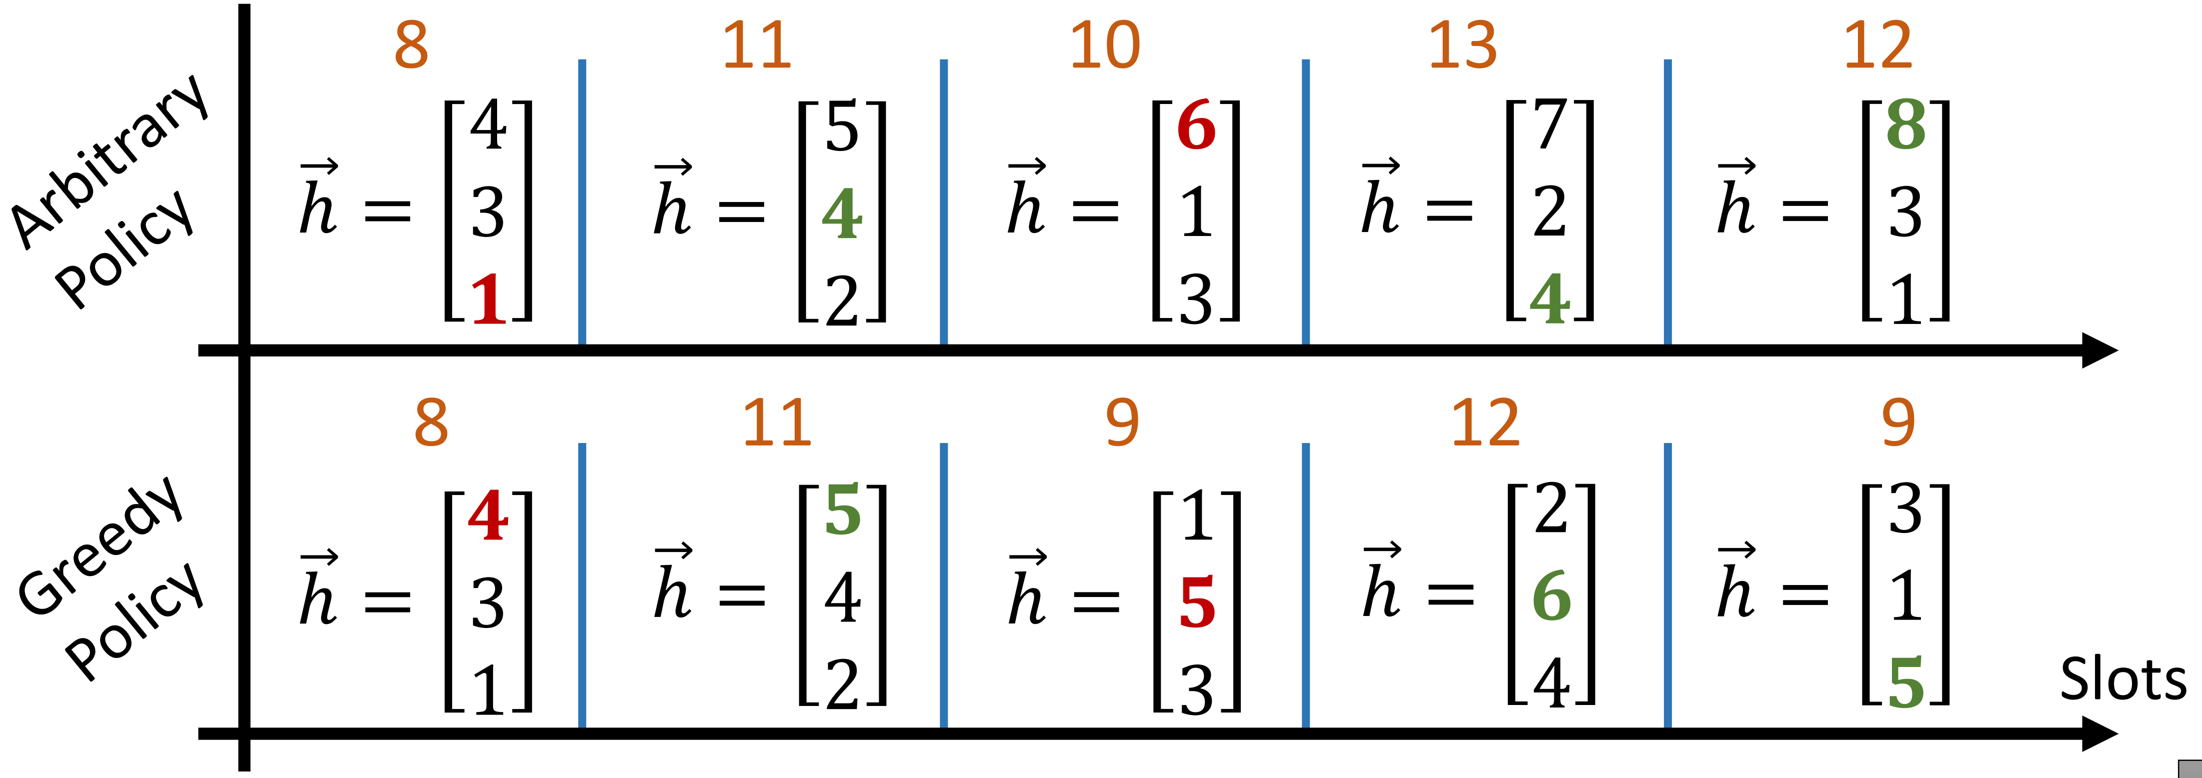

Scheduling Policies for **General** Networks

AoI from a different perspective

Lemma:

$$\lim_{K\to\infty} J_K^{\pi} \triangleq \lim_{K\to\infty} \frac{1}{KM}\sum_{k=1}^{K}\sum_{i=1}^{M} \boldsymbol{\alpha_i h_i(k)} = \frac{1}{M}\sum_{i=1}^{M} \frac{\boldsymbol{\alpha_i}}{2}\left[\frac{\overline{\mathbb{M}}[\boldsymbol{I_i^2}]}{\overline{\mathbb{M}}[\boldsymbol{I_i}]} + 1\right], \text{wp1}$$

where $I_i[m]$ is the inter-delivery time of node i and $\overline{\mathbb{M}}[I_i]$ is the **sample mean** of $I_i[m]$.

$$\lim_{K\to\infty} \frac{1}{K}\sum_{k=1}^{K} \boldsymbol{h_i(k)} = \frac{1}{2}\left[\frac{\overline{\mathbb{M}}[\boldsymbol{I_i^2}]}{\overline{\mathbb{M}}[\boldsymbol{I_i}]} + 1\right], \text{wp1}$$

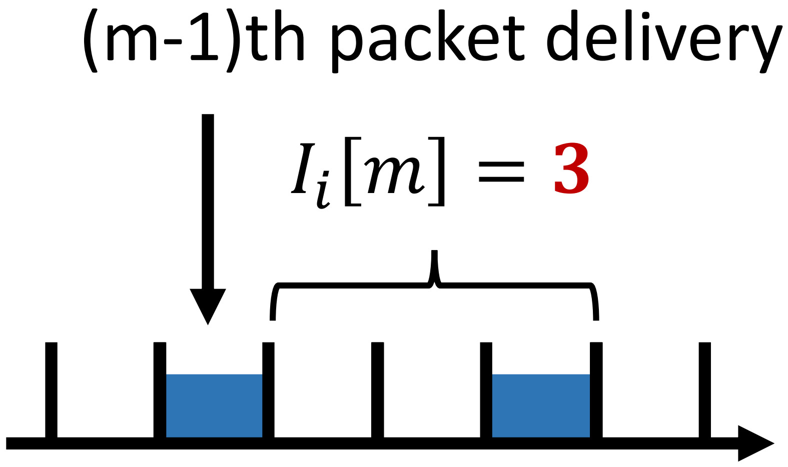

AoI from a different perspective

Lemma:

$$\lim_{K\to\infty} J_K^{\pi} \triangleq \lim_{K\to\infty} \frac{1}{KM}\sum_{k=1}^{K}\sum_{i=1}^{M} \boldsymbol{\alpha_i h_i(k)} = \frac{1}{M}\sum_{i=1}^{M}\frac{\boldsymbol{\alpha_i}}{2}\left[\frac{\overline{\mathbb{M}}[\boldsymbol{I_i^2}]}{\overline{\mathbb{M}}[\boldsymbol{I_i}]}+1\right], \mathrm{wp1}$$

where $I_i[m]$ is the inter-delivery time of node i

and $\overline{\mathbb{M}}[I_i]$ is the **sample mean** of $I_i[m]$.

$$\lim_{K\to\infty}\frac{1}{K}\sum_{k=1}^{K}\boldsymbol{h_i(k)} = \frac{1}{2}\left[\frac{\overline{\mathbb{M}}[\boldsymbol{I_i^2}]}{\overline{\mathbb{M}}[\boldsymbol{I_i}]}+1\right], \mathrm{wp1}$$

Notice:

$$\frac{\mathbb{E}[\boldsymbol{I_i^2}]}{\mathbb{E}[\boldsymbol{I_i}]} = \frac{\mathbb{V}ar[I_i]}{\mathbb{E}[I_i]} + \mathbb{E}[I_i]$$

(delivery regularity)

Intuition of the proof:

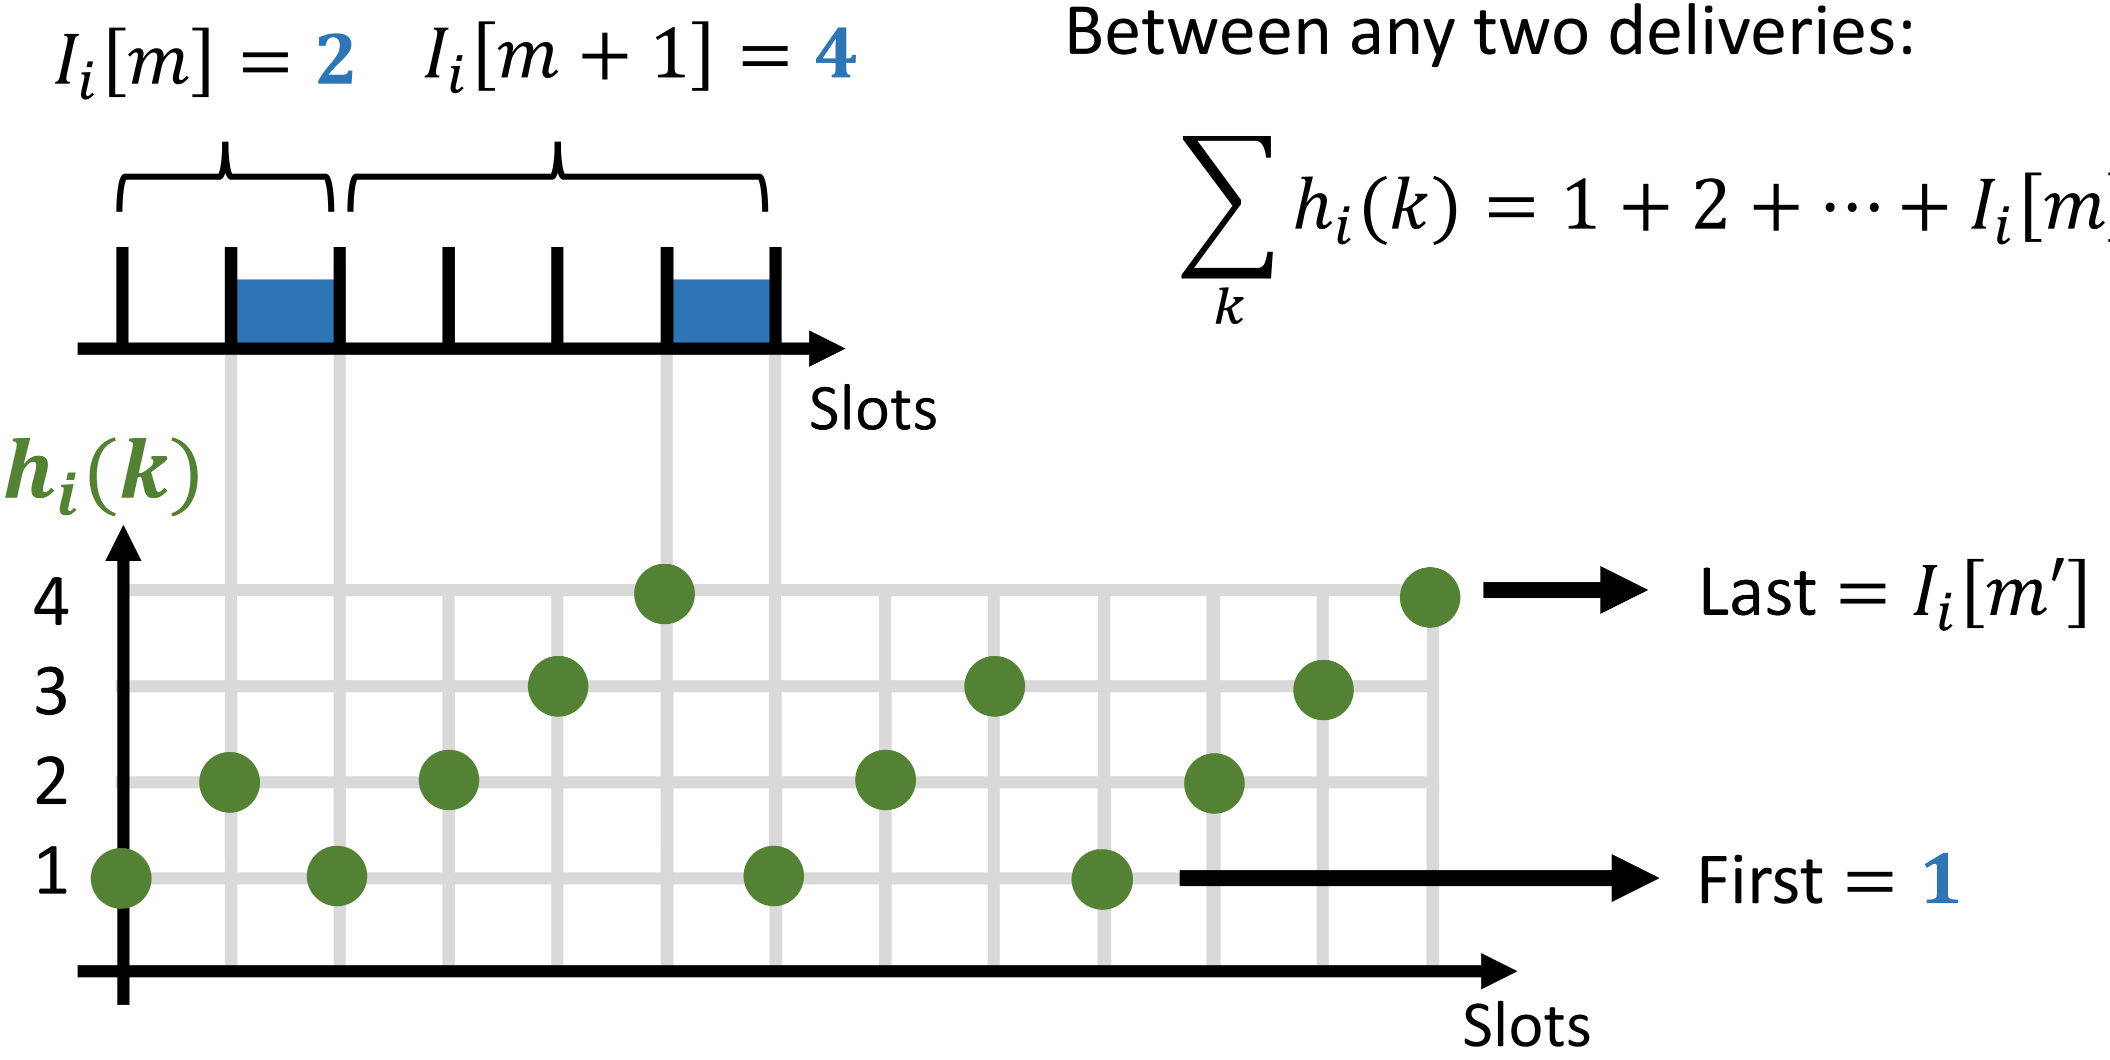

Between any two deliveries:

$$\sum_k h_i(k) = 1 + 2 + \cdots + I_i[m]$$

Intuition of the proof:

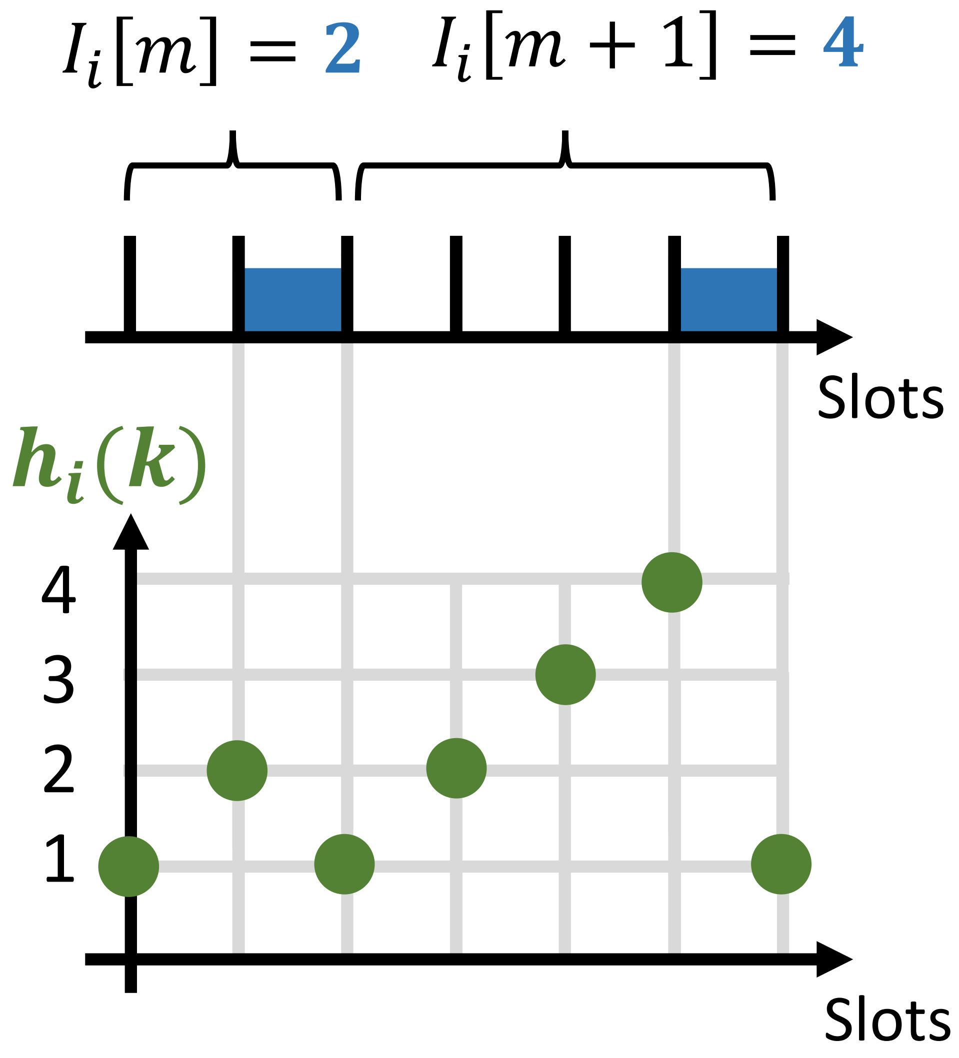

Between any two deliveries:

$$\sum_k h_i(k) = \frac{I_i[m]^2 + I_i[m]}{2}$$

Hence, the **time-average AoI** is:

$$\frac{1}{K}\sum_{k=1}^{K} h_i(k) \approx \frac{1}{\bar{\mathbb{M}}[I_i]}\left[\frac{\bar{\mathbb{M}}\left[I_i^2\right] + \bar{\mathbb{M}}[I_i]}{2}\right]$$

$$\frac{1}{K}\sum_{k=1}^{K} h_i(k) \approx \frac{1}{2}\left[\frac{\bar{\mathbb{M}}\left[I_i^2\right]}{\bar{\mathbb{M}}[I_i]} + 1\right]$$

Lower Bound

Lemma:

$$\lim_{K\to\infty} J_K^\pi \triangleq \lim_{K\to\infty} \frac{1}{KM}\sum_{k=1}^{K}\sum_{i=1}^{M} \boldsymbol{\alpha_i h_i(k)} = \frac{1}{M}\sum_{i=1}^{M}\frac{\boldsymbol{\alpha_i}}{2}\left[\frac{\overline{\mathbb{M}}\left[\boldsymbol{I_i^2}\right]}{\overline{\mathbb{M}}[\boldsymbol{I_i}]}+1\right], \text{wp1}$$

Theorem:

$$\lim_{K\to\infty} \mathbb{E}[J_K^\pi] \geq L_B, \text{where } \mathrm{L_B} = \frac{1}{2M}\min_{\pi\in\Pi}\left\{\sum_{i=1}^{M}\boldsymbol{\alpha_i}\mathbb{E}\left[\overline{\mathbb{M}}[I_i^\pi]\right]\right\}+\frac{1}{2M}\sum_{i=1}^{M}\boldsymbol{\alpha_i}$$

Proof: applying Fatou's Lemma to the non-negative sequence J_K^π and then the generalized mean inequality $\overline{\mathbb{M}}\left[I_i^2\right] \geq (\overline{\mathbb{M}}[I_i])^2$.

Transmission Scheduling Policies

- Next, we develop three different **scheduling policies**:

Scheduling Policy	Technique	Optimality Ratio	Simulation Result
Optimal Stationary Randomized Policy	Renewal Theory	2-optimal	~ 2-optimal
Max-Weight Policy	Lyapunov Optimization	2-optimal	close to optimal
Whittle's Index Policy	RMAB Framework	8-optimal	close to optimal

$$L_B \leq \lim_{K\to\infty} \mathbb{E}\left[J_K^{\boldsymbol{R}}\right] \leq 2L_B$$

Randomized Policies

- **Stationary Randomized Policy:**
  in slot k, select node i with probability $\boldsymbol{\mu_i}$.
- Clearly, $d_i(k) \sim Ber(\boldsymbol{\mu_i p_i})$ and $I_i[m] \sim Geo(\boldsymbol{\mu_i p_i})$ iid for node i
- Sequence of packet deliveries from node i is a **renewal process:**

$$I_i[m] = \mathbf{4}$$

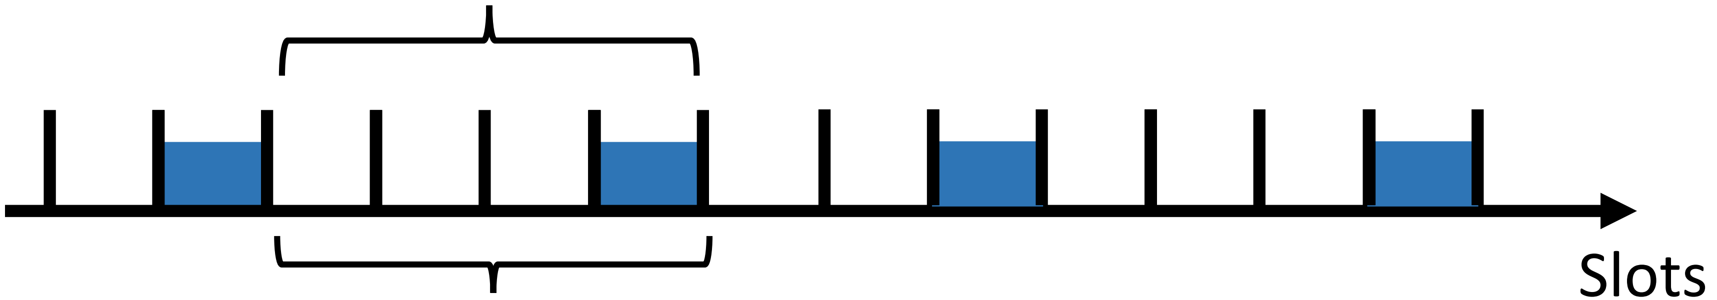

$$\sum_k h_i(k) = \frac{I_i[m]^2 + I_i[m]}{2}$$

Randomized Policies

- **Stationary Randomized Policy:**

  in slot k, select node i with probability $\boldsymbol{\mu_i}$.

- Clearly, $d_i(k) \sim Ber(\boldsymbol{\mu_i p_i})$ and $I_i[m] \sim Geo(\boldsymbol{\mu_i p_i})$ iid for node i

- Sum of $h_i(k)$ is a **renewal-reward process:**
  - Period length $\mathbb{E}[I_i] = (\boldsymbol{\mu_i p_i})^{-1}$
  - Reward in a period is $\mathbb{E}\left[I_i^2 + I_i\right]/2 = (\boldsymbol{\mu_i p_i})^{-2}$

Hence, by the elementary renewal theorem:

$$\lim_{K\to\infty} \mathbb{E}[J_K^R] = \frac{1}{M}\sum_{i=1}^{M} \boldsymbol{\alpha_i} \frac{\mathbb{E}[reward]}{\mathbb{E}[period]} = \frac{1}{M}\sum_{i=1}^{M} \frac{\boldsymbol{\alpha_i}}{\boldsymbol{\mu_i p_i}} =$$

Randomized Policies

- **Stationary Randomized Policy:**

  in slot k, select node i with probability $\boldsymbol{\mu_i}$.

- Clearly, $d_i(k) \sim Ber(\boldsymbol{\mu_i p_i})$ and $I_i[m] \sim Geo(\boldsymbol{\mu_i p_i})$ iid for node i

- Sum of $h_i(k)$ is a **renewal-reward process:**
  - Period length $\mathbb{E}[I_i] = (\boldsymbol{\mu_i p_i})^{-1}$
  - Reward in a period is $\mathbb{E}\left[I_i^2 + I_i\right]/2 = (\boldsymbol{\mu_i p_i})^{-2}$

Hence, by the elementary renewal theorem:

$$\lim_{K\to\infty} \mathbb{E}[J_K^R] = \frac{1}{M}\sum_{i=1}^{M} \boldsymbol{\alpha_i} \frac{\mathbb{E}[reward]}{\mathbb{E}[period]} = \frac{1}{M}\sum_{i=1}^{M} \frac{\boldsymbol{\alpha_i}}{\boldsymbol{\mu_i p_i}} = \frac{1}{M}\sum_{i=1}^{M} \boldsymbol{\alpha_i}\mathbb{E}[I_i^{\boldsymbol{R}}]$$

Randomized Policies

Theorem: there exists **R** which is **2-optimal** for **any** network configuration.

- Proof:

Recall that: $$L_B = \frac{1}{2M} \min_{\pi \in \Pi} \left\{ \sum_{i=1}^{M} \alpha_i \mathbb{E}\left[\bar{\mathbb{M}}[I_i^{\pi}]\right] \right\} + \frac{1}{2M} \sum_{i=1}^{M} \alpha_i$$

Policy that solves the minimization above yields some value of $\mathbb{E}\left[\bar{\mathbb{M}}[I_i^{\pi}]\right]$.

There exists **R** such that $\mathbb{E}\left[I_i^{R}\right] = \mathbb{E}\left[\bar{\mathbb{M}}[I_i^{\pi}]\right]$ for all nodes. For this particular **R**, we have:

$$\lim_{K \to \infty} \mathbb{E}\left[J_K^{R}\right] = \frac{1}{M} \sum_{i=1}^{M} \alpha_i \mathbb{E}\left[\bar{\mathbb{M}}[I_i^{\pi}]\right] \leq 2L_B$$

Randomized Policies

Theorem: there exists **R** which is **2-optimal** for **any** network configuration.

- **Optimal Stationary Randomized Policy:** $$\mu_i^* = \frac{1}{M} \underset{\mu_i, \forall i}{\operatorname{argmin}} \left\{ \sum_{i=1}^{M} \frac{\alpha_i}{\mu_i p_i} \right\}$$

  The optimal solution is $\mu_i^* \propto \sqrt{\alpha_i / p_i}$

Corollary: Optimal **R*** is 2-optimal.

Summary (so far…)

- Objective Function: $$\min_{\pi\in\Pi}\lim_{K\to\infty}\mathbb{E}[\,J_K^\pi\,] = \min_{\pi\in\Pi}\lim_{K\to\infty}\frac{1}{KM}\mathbb{E}\left[\sum_{k=1}^{K}\sum_{i=1}^{M}\alpha_i h_i^\pi(k)\right]$$

- Network Model:

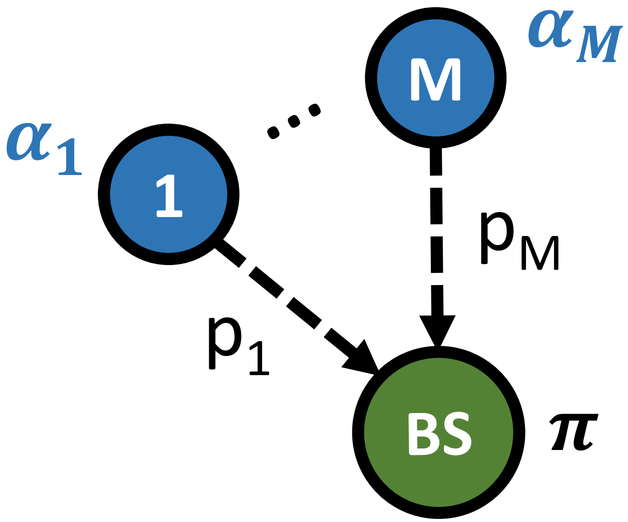

Policy	Decision in slot k	Performance
Greedy	highest $h_i(k)$	optimal when symmetric
Randomized	node i wp $\propto \sqrt{\alpha_i/p_i}$	2-optimal
Max-Weight		
Whittle's Index		

Max-Weight Policy

- Max-Weight is designed to minimize the **Lyapunov Drift.**

- Evolution of AoI:

$$\mathbf{h_i}(\boldsymbol{k}+\mathbf{1}) = \begin{cases} 1, & \text{if } d_i(k) = 1 \\ \boldsymbol{h_i}(\boldsymbol{k}) + 1, & \text{otherwise} \end{cases}$$

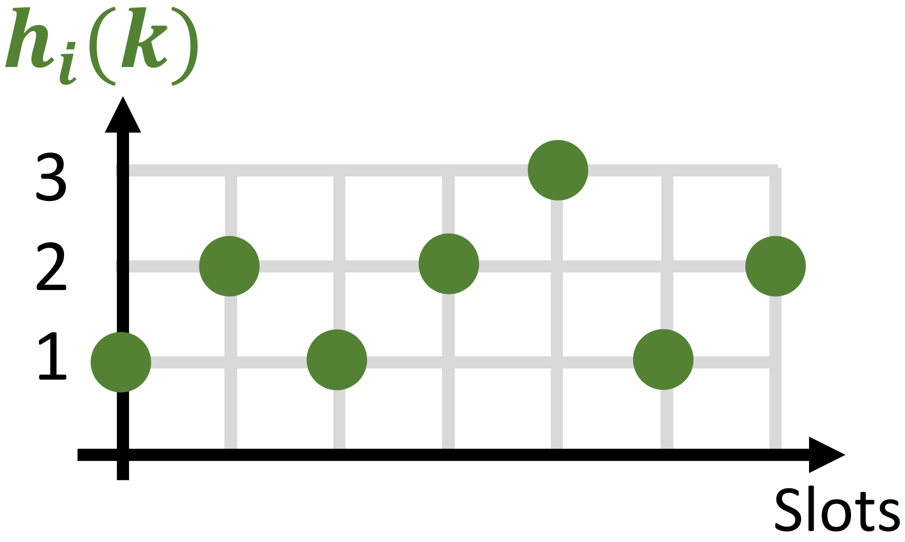

- Equivalently: $h_i(k+1) = h_i(k)[1 - \boldsymbol{d_i}(\boldsymbol{k})] + 1$

$$h_i(k+1) = h_i(k)[1 - \boldsymbol{c_i}(\boldsymbol{k})\boldsymbol{u_i}(\boldsymbol{k})] + 1$$

Max-Weight Policy

- Lyapunov Function: $L(k) = \frac{1}{M}\sum_{i=1}^{M} \beta_i h_i(k)$, where $\beta_i > 0$ is a constant
- Lyapunov Drift: $\Delta(k) = \mathbb{E}\{\, L(k+1) - L(k) \mid h_i(k) \,\}$

Max-Weight Policy

- Lyapunov Function: $L(k) = \frac{1}{M}\sum_{i=1}^{M} \beta_i h_i(k)\,,\text{where }\; \beta_i > 0$ is a constant
- Lyapunov Drift: $\Delta(k) = \mathbb{E}\{\, L(k+1) - L(k) \mid h_i(k)\,\}$
- Recall that: $h_i(k+1) = h_i(k)[1 - \boldsymbol{c_i(k)}\boldsymbol{u_i(k)}] + 1$

$$\Delta(k) = -\frac{1}{M}\sum_{i=1}^{M} \beta_i h_i(k)\boldsymbol{p_i}\mathbb{E}\{\boldsymbol{u_i(k)} \mid h_i(k)\,\} + \frac{1}{M}\sum_{i=1}^{M} \beta_i$$

- **MW policy:** in slot k, schedule node with highest value of $\beta_i \boldsymbol{p_i} h_i(k)$

Max-Weight Policy

Theorem: MW with $\beta_i = \boldsymbol{\alpha_i}/\boldsymbol{\mu_i^* p_i}$ is **2-optimal** for **any** network configuration

- **MW policy** with $\beta_i = \boldsymbol{\alpha_i}/\boldsymbol{\mu_i^* p_i}$:

  in slot k, schedule node with highest value of $\sqrt{\boldsymbol{\alpha_i p_i}}\, h_i(k) \equiv \boldsymbol{\alpha_i p_i} h_i^2(k)$

Max-Weight Policy

Theorem: MW with $\beta_i = \boldsymbol{\alpha_i}/\boldsymbol{\mu_i^*}\boldsymbol{p_i}$ is **2-optimal** for **any** network configuration

- **MW policy** with $\beta_i = \boldsymbol{\alpha_i}/\boldsymbol{\mu_i^*}\boldsymbol{p_i}$:

  in slot k, schedule node with highest value of $\sqrt{\boldsymbol{\alpha_i}\boldsymbol{p_i}}h_i(k) \equiv \boldsymbol{\alpha_i}\boldsymbol{p_i}h_i^2(k)$

- Proof outline:
  - MW minimizes drift while Optimal R* does not. Thus:

    $$\Delta^{\boldsymbol{MW}}(k) \leq \Delta^{\boldsymbol{R}^*}(k)$$

  - Manipulating the expression and substituting $\beta_i = \boldsymbol{\alpha_i}/\boldsymbol{\mu_i^*}\boldsymbol{p_i}$, gives:

    $$\lim_{K\to\infty} \mathbb{E}\left[J_K^{\boldsymbol{MW}}\right] \leq \lim_{K\to\infty} \mathbb{E}\left[J_K^{\boldsymbol{R}^*}\right]$$

  - Hence, Max-Weight is 2-optimal.

Summary (so far...)

- Objective Function: $$\min_{\pi\in\Pi}\lim_{K\to\infty}\mathbb{E}[J_K^\pi] = \min_{\pi\in\Pi}\lim_{K\to\infty}\frac{1}{KM}\mathbb{E}\left[\sum_{k=1}^{K}\sum_{i=1}^{M}\alpha_i h_i^\pi(k)\right]$$

- Network Model:

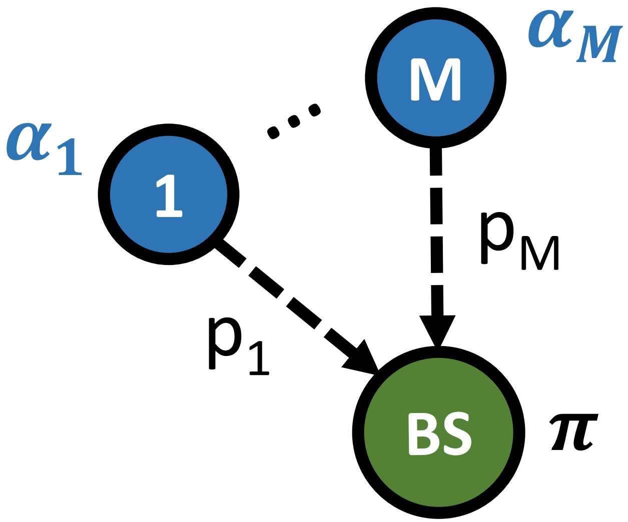

Policy	Decision in slot k	Performance
Greedy	highest $h_i(k)$	optimal when symmetric
Randomized	node i wp $\propto \sqrt{\alpha_i/p_i}$	2-optimal
Max-Weight	highest $\alpha_i p_i h_i^2(k)$	2-optimal

Whittle's Index Policy

A Markov Bandit is characterized by a MDP in which:

- $u(k) = 0$ freezes the process and gives no reward;
- $u(k) = 1$ continues the process with $\mathbb{P}$ and gives reward $r(h(k))$

- In contrast, when the bandit is **restless:**
  - $u(k) = 0$ continues the process with $\mathbb{P}_0$ and gives reward $r_0(h(k))$
  - $u(k) = 1$ continues the process with $\mathbb{P}_1$ and gives reward $r_1(h(k))$

- The AoI of each node evolves as a **restless bandit**.
  Hence, we can use the RMAB framework [2] to design an Index Policy.

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

Whittle's Index Policy

- For designing the Index Policy, we use the RMAB framework in [2].
  - We relax our problem to the case of a single node i , $M = 1$, and add a cost per transmission, $C > 0$:

$$\min_{\pi \in \Pi} \lim_{K \to \infty} \frac{1}{K} \mathbb{E}\left[\sum_{k=1}^{K} \boldsymbol{\alpha_i h_i(k)} + C\boldsymbol{u_i(k)}\right]$$

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

- Challenges

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

Whittle's Index Policy

- Consider the relaxed problem with a single node and cost per transmission C.

$$\min_{\pi \in \Pi} \lim_{K \to \infty} \frac{1}{K} \mathbb{E}\left[\sum_{k=1}^{K} \boldsymbol{\alpha_i h_i(k)} + C\boldsymbol{u_i(k)}\right]$$

- Condition for Indexability:
  - Let $\mathcal{P}(C)$ be the set of states $h_i(k)$ 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's Index Policy

- Consider the relaxed problem with a single node and cost per transmission C.

$$\min_{\pi\in\Pi}\lim_{K\to\infty}\frac{1}{K}\mathbb{E}\left[\sum_{k=1}^{K}\boldsymbol{\alpha_i h_i(k)}+C\boldsymbol{u_i(k)}\right]$$

- 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 node in state h.

Whittle's Index Policy

- We establish that the problem is **indexable** and find a **closed-form** solution for the Whittle Index.
- **Index Policy**: in slot k, select the node with *highest value of* $C_i(h_i(k))$, where:

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

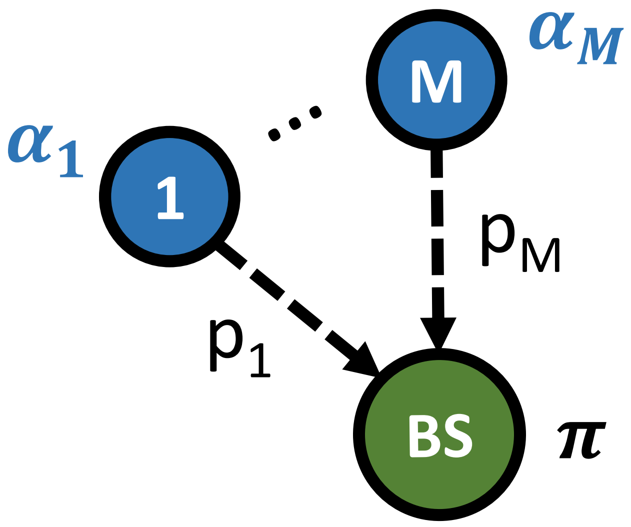

Whittle's Index Policy

- We establish that the problem is **indexable** and find a **closed-form** solution for the Whittle Index.

- **Index Policy**: in slot k, select the node with *highest value of* $C_i(h_i(k))$, where:

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

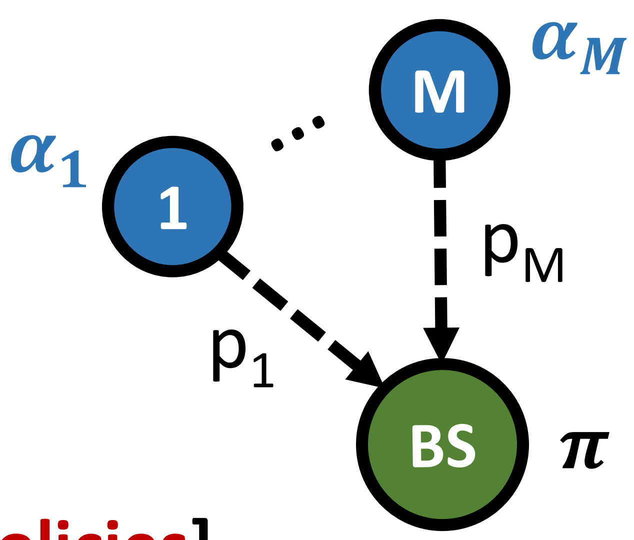

- Whittle's Index is similar to Max-Weight: $\boldsymbol{\alpha_i p_i} h_i^2(k)$

- For Symmetric Networks:

Whittle's ≡ Max-Weight ≡ Greedy. [All optimal policies]

Summary (so far…)

- Objective Function: $\min_{\pi\in\Pi} \lim_{K\to\infty} \mathbb{E}[J_K^\pi] = \min_{\pi\in\Pi} \lim_{K\to\infty} \frac{1}{KM}\mathbb{E}\left[\sum_{k=1}^{K}\sum_{i=1}^{M} \alpha_i h_i^\pi(k)\right]$

- Network Model:

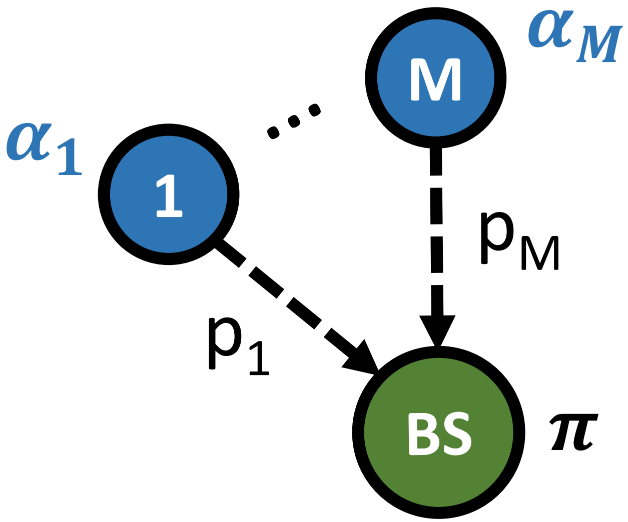

Policy	Decision in slot k	Performance
Greedy	highest $h_i(k)$	optimal when symmetric
Randomized	node i wp $\propto \sqrt{\alpha_i/p_i}$	2-optimal
Max-Weight	highest $\alpha_i p_i h_i^2(k)$	2-optimal
Whittle's Index	highest $C_i(h_i(k))$	8-optimal

Numerical Results

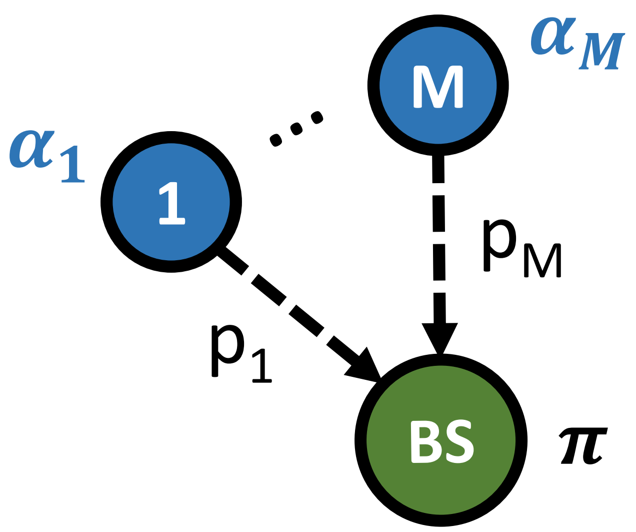

- Metric:
  - Expected Weighted Sum AoI : $\mathbb{E}[\, J_K^{\pi} \,]$

- Network Setup with M nodes. Node i has:
  - weight $\boldsymbol{\alpha_i} = (M + 1 - \boldsymbol{i})/M$ [decreasing with i]
  - channel reliability $\boldsymbol{p_i} = \boldsymbol{i}/M$ [increasing with i]

- Each simulation runs for $K = M \times 10^6$ slots
- Each data point is an average over 10 simulations

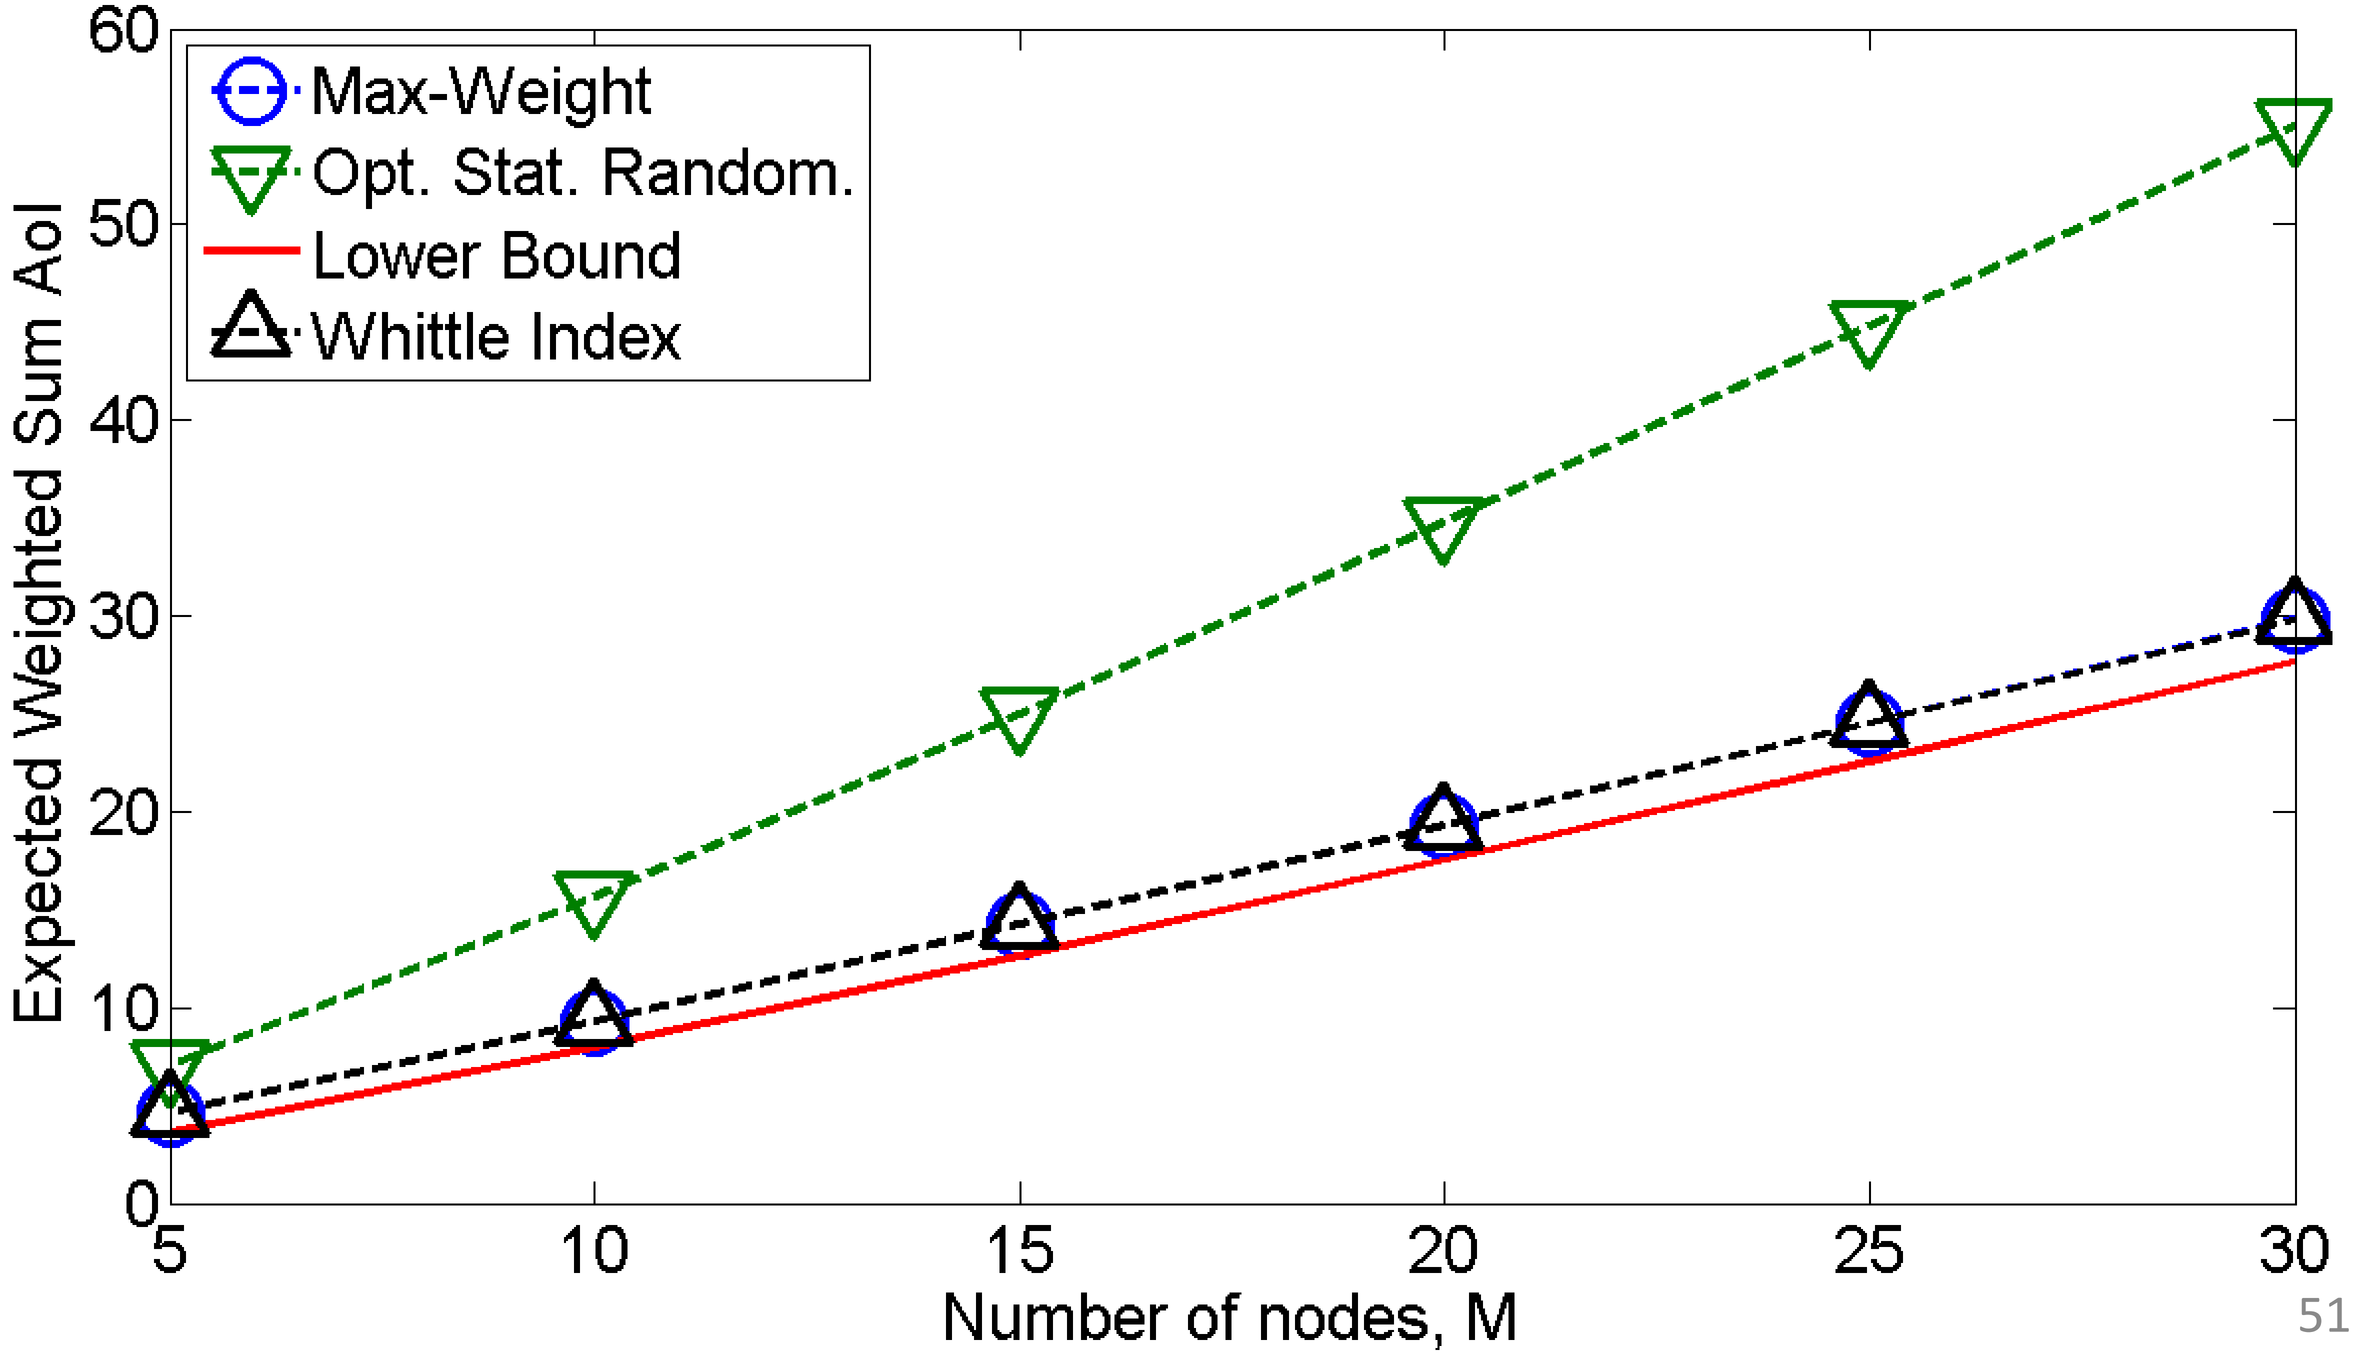

Max-Weight
Opt. Stat. Random.
Lower Bound
Whittle Index
Expected Weighted Sum AoI
Number of nodes, M
0
10
20
30
40
50
60
5
10
15
20
25
30

Remarks (so far…)

- We developed and evaluated low-complexity scheduling policies:
  - Greedy, Optimal Stationary Randomized, Max-Weight and Whittle's Index
  - Obtained optimality ratios and validated their performance using Numerical Results
- Main ideas:
  - Age of Information is a measure of information freshness
  - Randomized Policy is the simplest possible policy and it is 2-optimal.
  - Max-Weight and Whittle's Index have near optimal AoI-performance.
- **Next:** Throughput versus Age of Information

Long-term Throughput vs Age of Information

Throughput

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

$$\mathbb{E}[\,T_K^{\pi}\,] = \frac{1}{KM}\mathbb{E}\left[\sum_{k=1}^{K}\sum_{i=1}^{M}\boldsymbol{\alpha_i d_i^{\pi}(k)}\right],$$ where $\boldsymbol{d_i^{\pi}(k)}$ indicates a delivery and $\boldsymbol{\alpha_i}$ is the positive weight

- Goal is to find the scheduling policy that $\max\limits_{\pi\in\Pi} \mathbb{E}[\,\boldsymbol{T_K^{\pi}}\,]$.

$$\mathbb{E}[\,T_K^{\pi}\,] = \frac{1}{KM}\sum_{k=1}^{K}\sum_{i=1}^{M}\boldsymbol{\alpha_i p_i}\mathbb{E}[\boldsymbol{u_i^{\pi}(k)}]$$

- **Optimal policy**: in slot k, select node with highest value of $\boldsymbol{\alpha_i p_i}$

Throughput vs Age of Information

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

$$\mathbb{E}[\, T_K^{\pi} \,] = \frac{1}{KM}\mathbb{E}\left[\sum_{k=1}^{K}\sum_{i=1}^{M}\alpha_i \boldsymbol{d}_i^{\pi}(\boldsymbol{k})\right]$$

, where $\boldsymbol{d}_i^{\pi}(\boldsymbol{k})$ indicates a delivery

- The **AoI minimization** problem is given by:

$$\mathbb{E}[\, J_K^{\pi} \,] = \frac{1}{KM}\mathbb{E}\left[\sum_{k=1}^{K}\sum_{i=1}^{M}\alpha_i \boldsymbol{h}_i^{\pi}(\boldsymbol{k})\right]$$

, where $\boldsymbol{h}_i^{\pi}(\boldsymbol{k})$ is the AoI

- We want to compare the scheduling policies that solve each problem.

Throughput vs Age of Information

- For a fixed vector of weights $\vec{\alpha}$, we have:

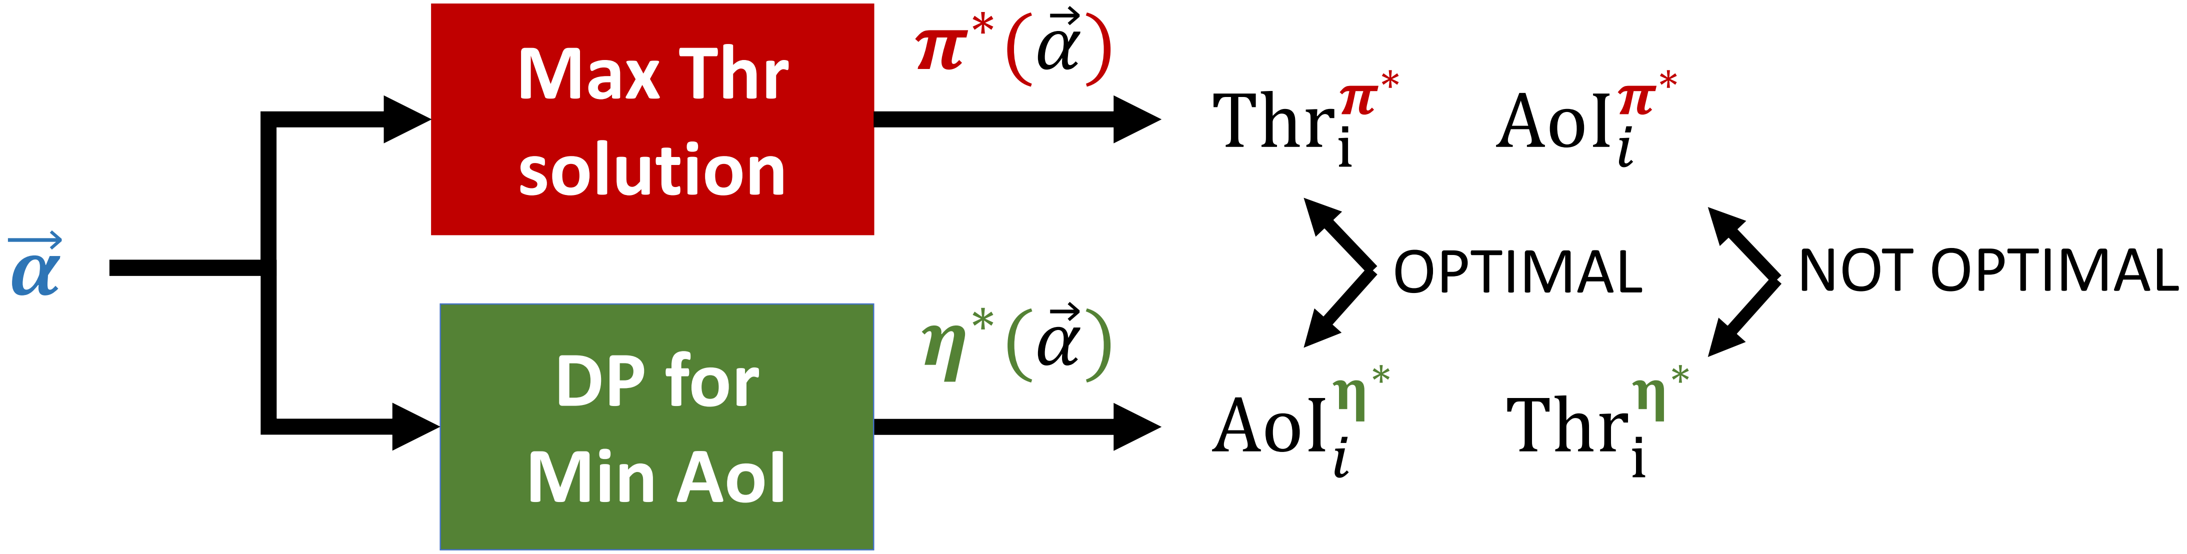

- We sweep $\vec{\alpha}$ and plot the results next:
  - **Red** is for metrics associated with the Throughput policy $\pi^*(\vec{\alpha})$.
  - **Green** is associated with the AoI policy $\eta^*(\vec{\alpha})$.

<u>Age of Info.</u>

M = 2

K = 120

$p_1 = 1/3$

$p_2 = 1/3$

α varies

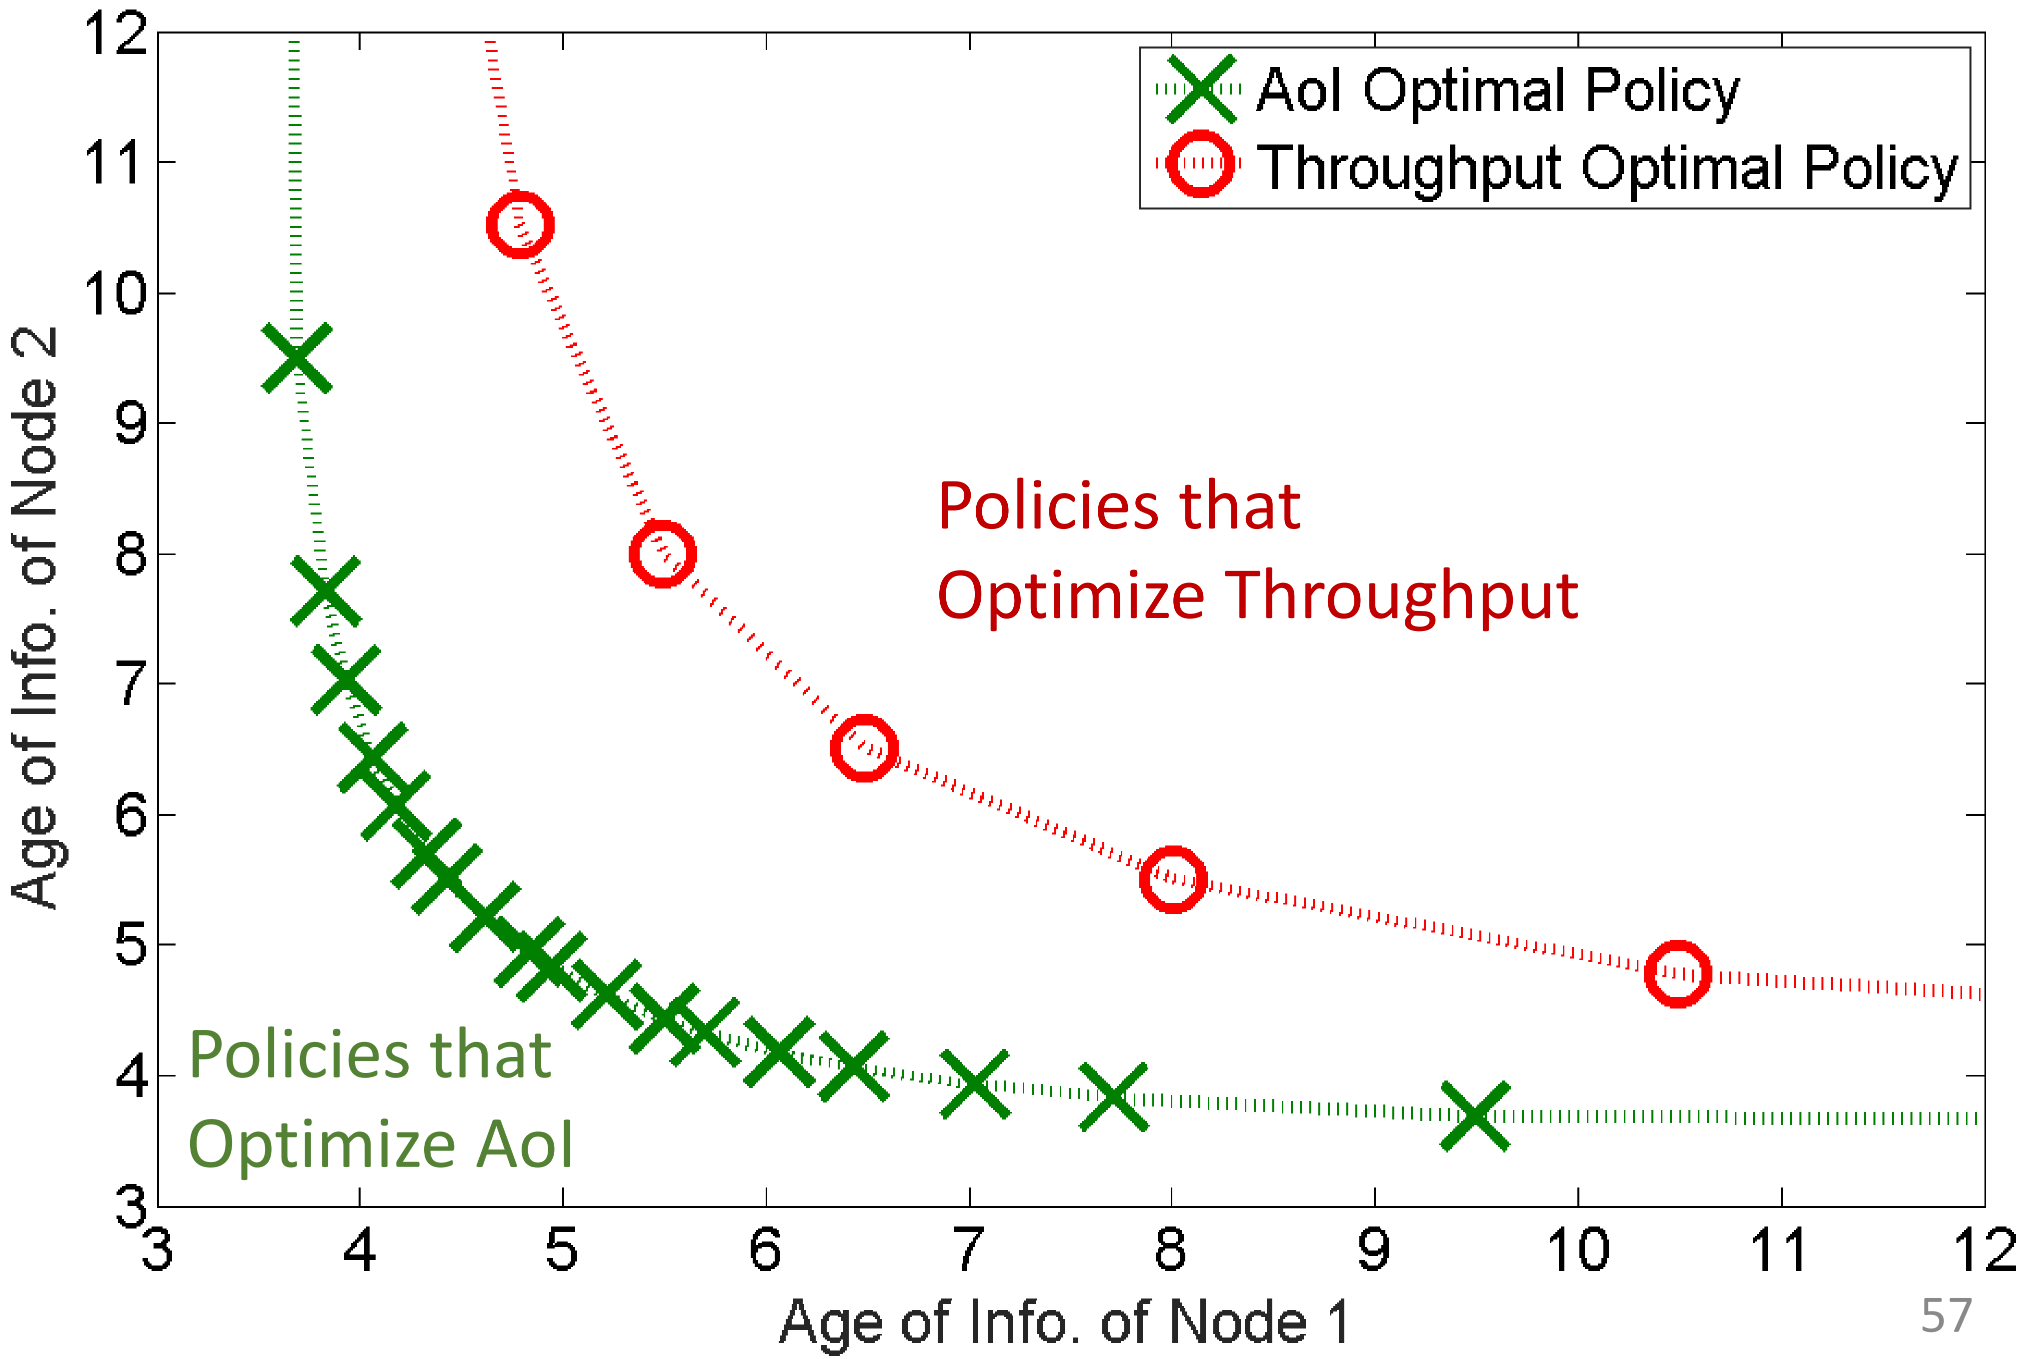

Throughput

M = 2

K = 120

$p_1 = 1/3$

$p_2 = 1/3$

α varies

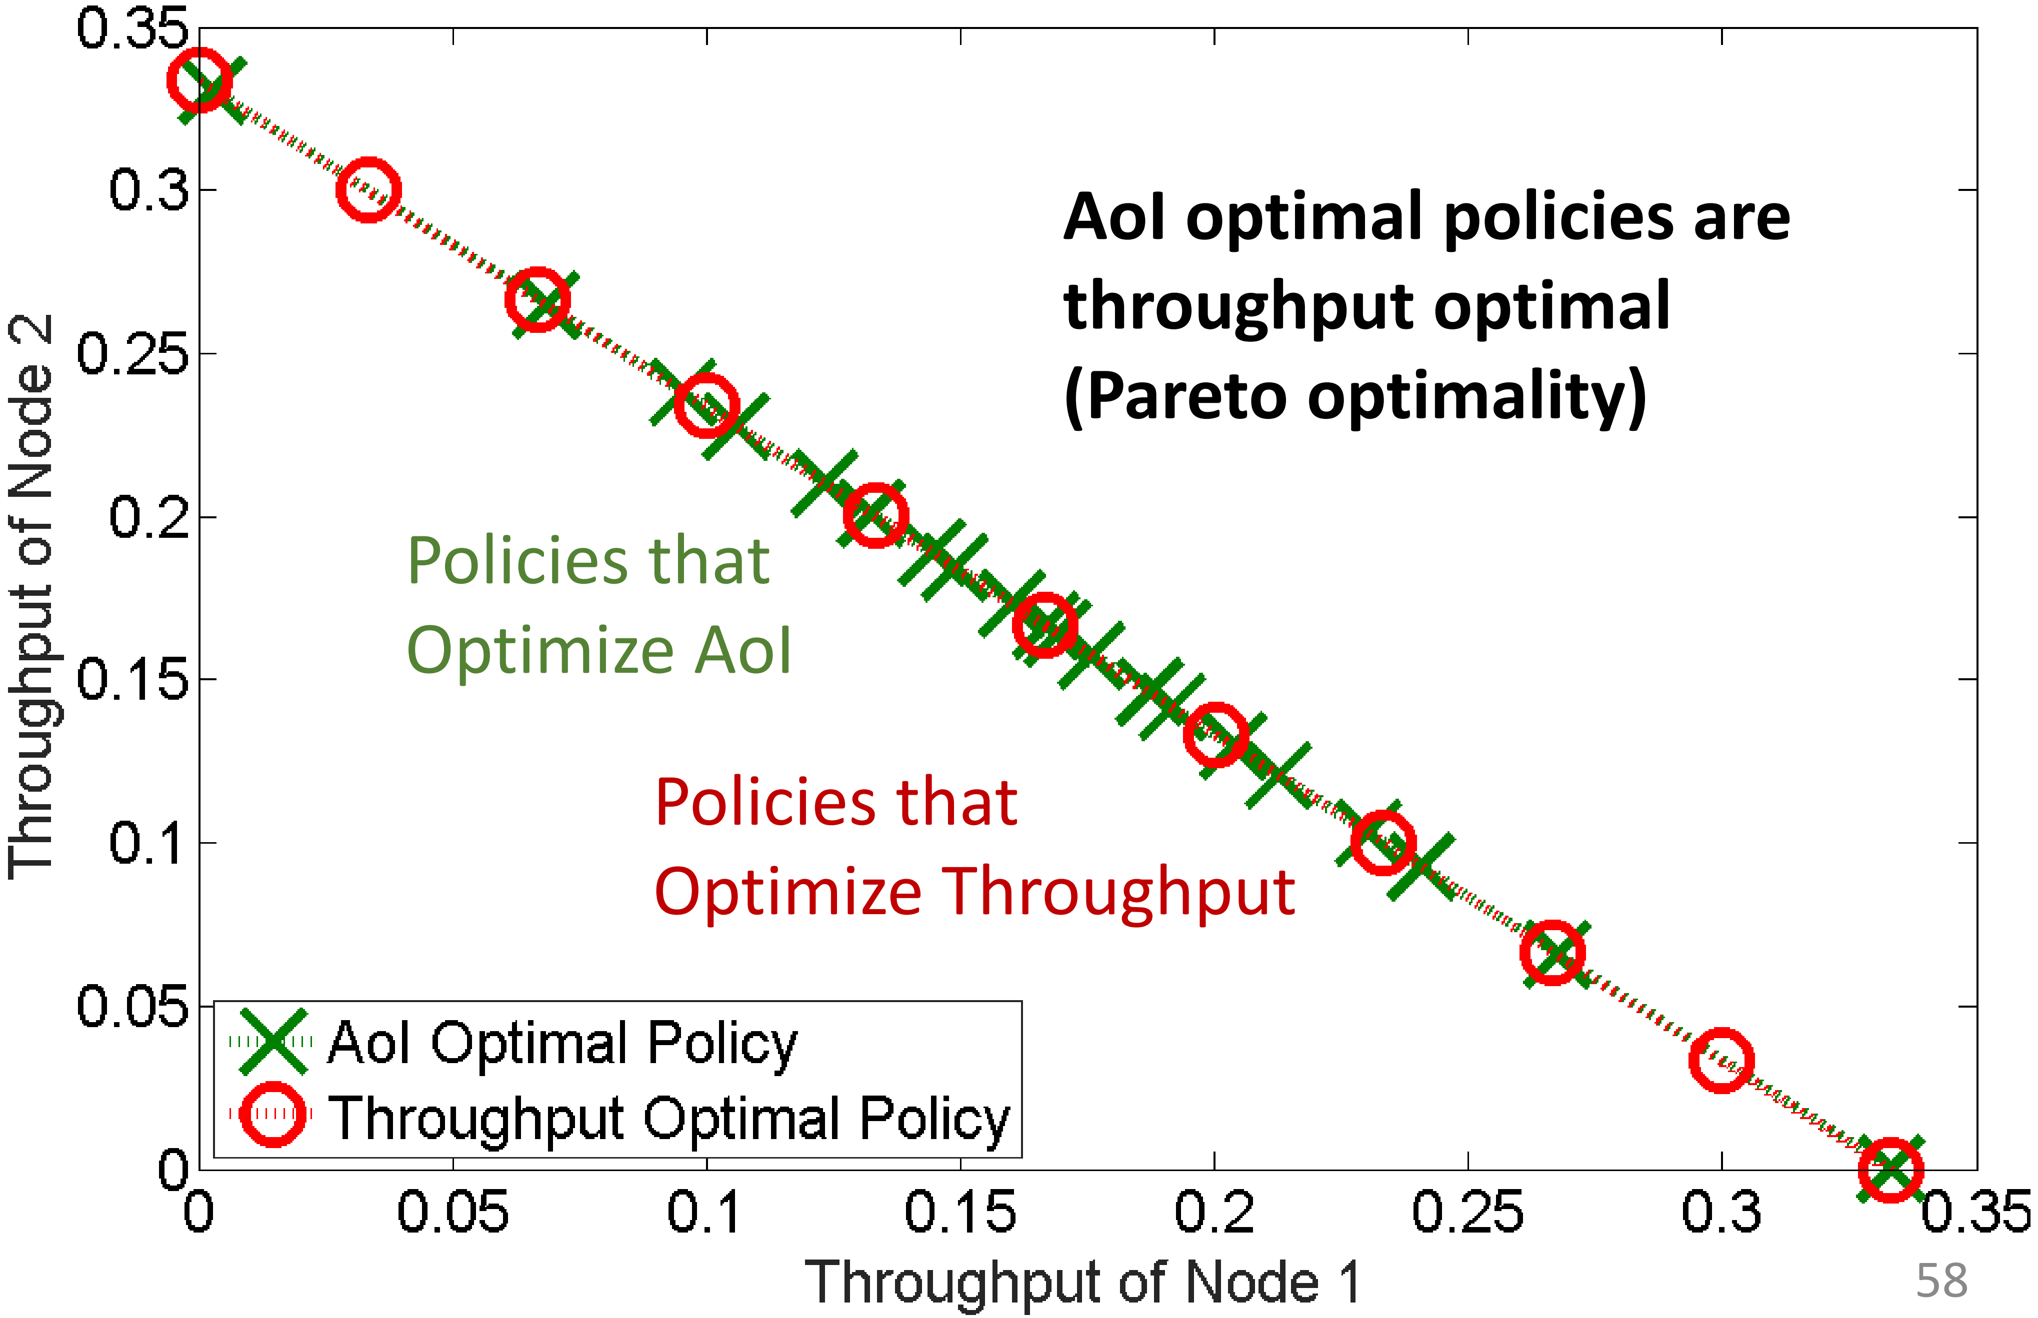

<u>Throughput</u>

M = 2

K = 120

$p_1 = 1/3$

$p_2 = 1/3$

α varies

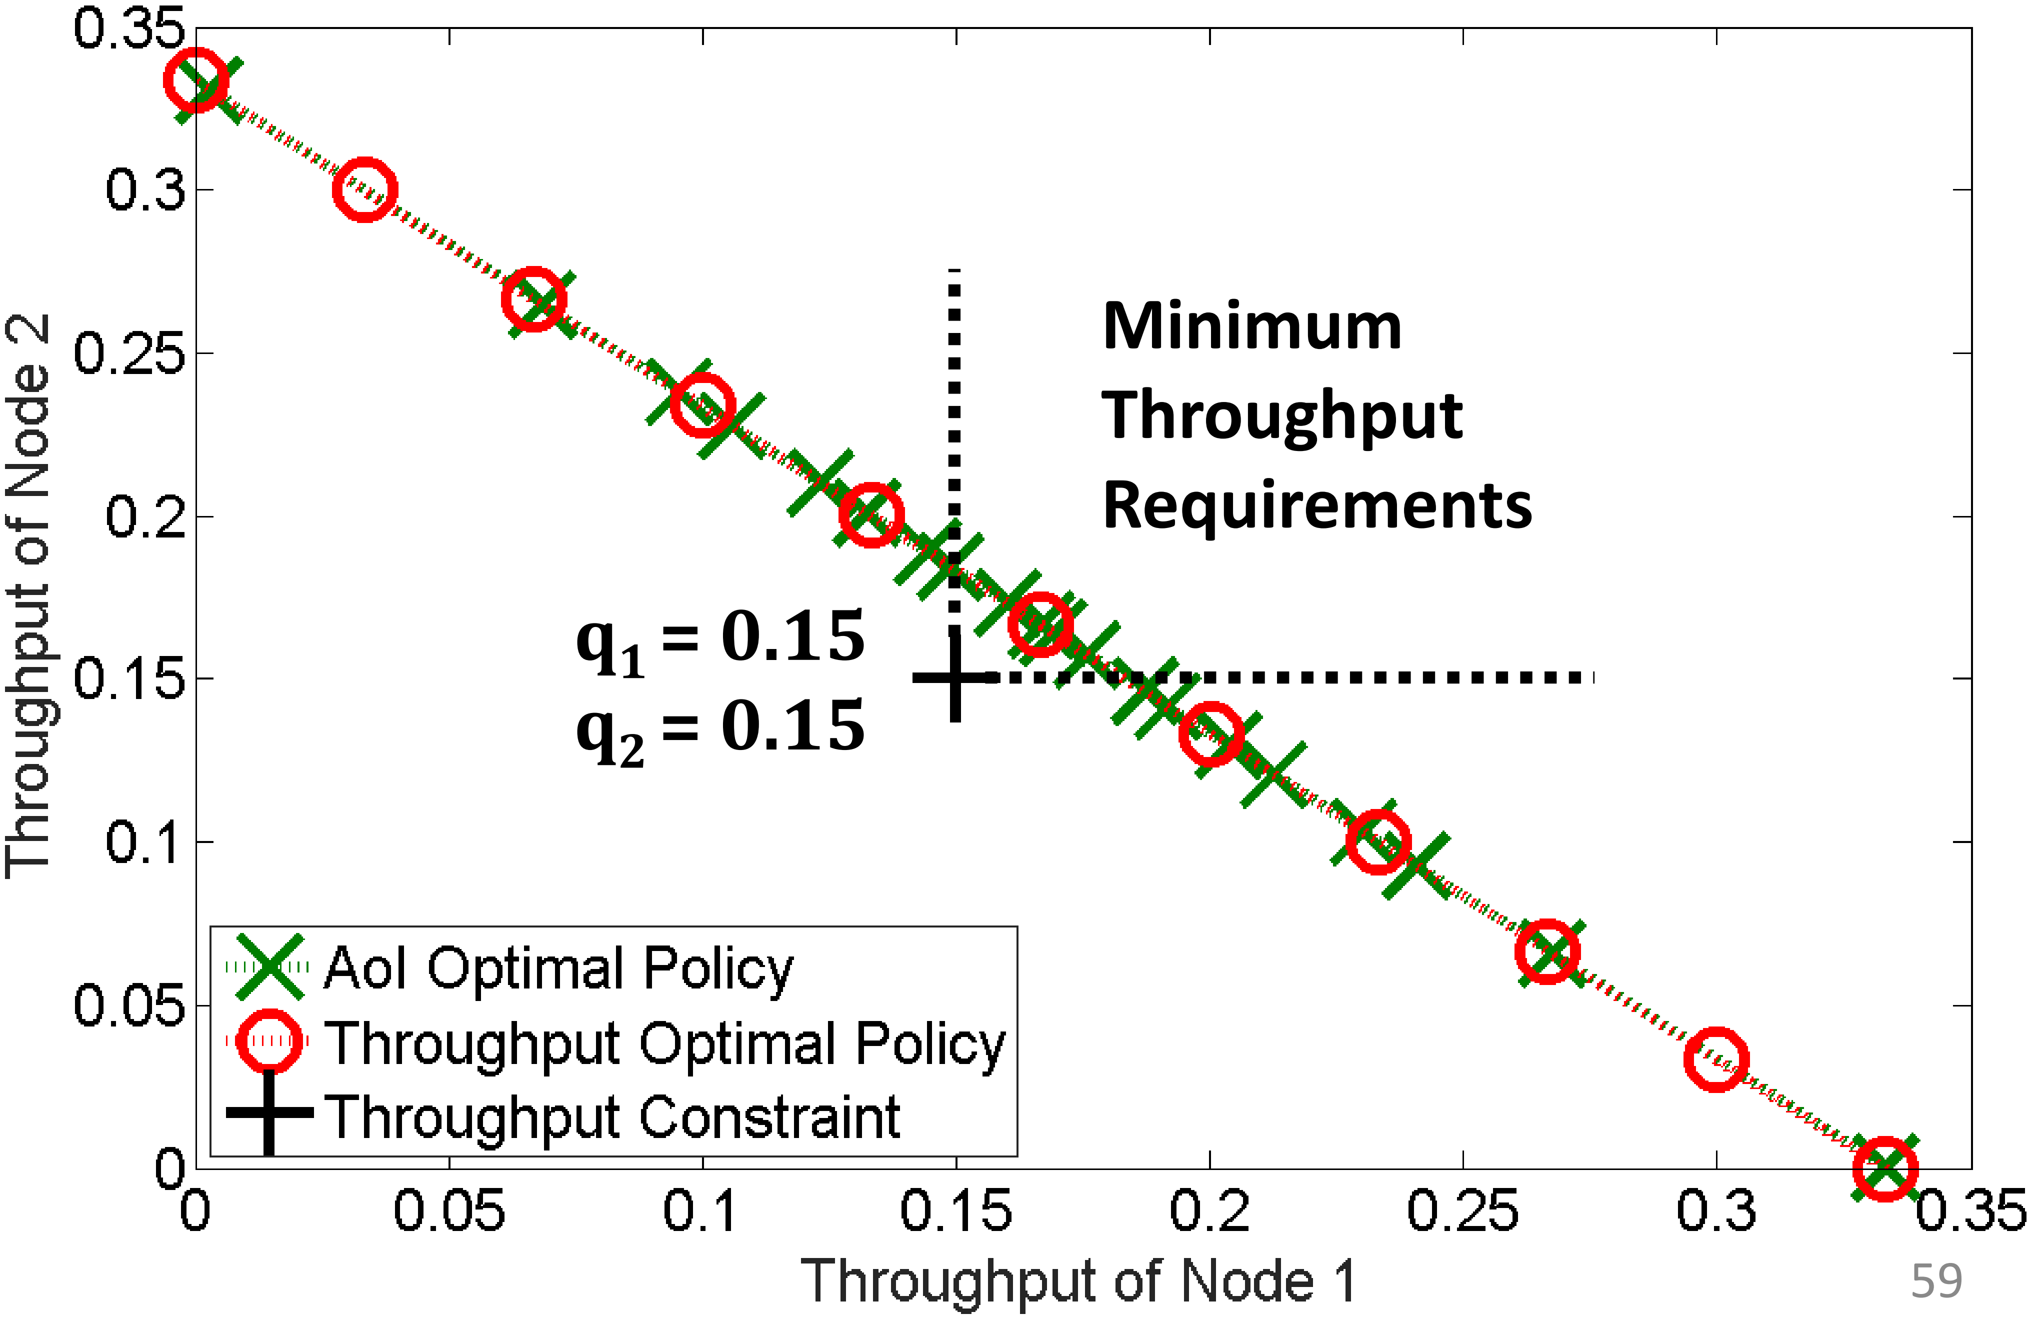

Minimizing Age of Information subject to Minimum Throughput Requirements

Minimum Throughput Requirement

- Long-term throughput of node i when policy π is employed is defined as:

$$\hat{q}_i^{\pi} := \lim_{K\to\infty} \frac{1}{K}\mathbb{E}\left[\sum_{k=1}^{K} d_i^{\pi}(k)\right]$$

- Minimum Throughput Requirements:

$$\hat{q}_i^{\pi} \geq \boldsymbol{q_i}\,, \forall i$$

we assume that the set $\{\boldsymbol{q_i}\}_{i=1}^{M}$ is feasible.

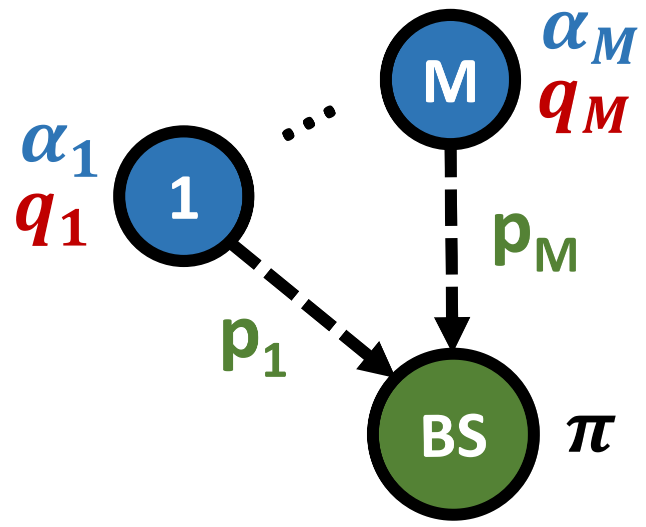

Optimization Problem

AoI Optimization

$$\mathrm{OPT}^* = \min_{\pi \in \Pi} \left\{ \lim_{K \to \infty} \frac{1}{KM} \mathbb{E} \left[\sum_{k=1}^{K} \sum_{i=1}^{M} \alpha_i h_i(k) \right] \right\} \tag{8a}$$

(Age of Information)

$$\text{s.t. } \hat{q}_i^{\pi} \geq q_i \ , \forall i \ ; \tag{8b}$$

(Minimum Throughput)

$$\textstyle\sum_{i=1}^{M} u_i(k) \leq 1 \ , \forall k \ . \tag{8c}$$

(Channel Interference)

- Next, we consider:
  - Stationary Randomized Policy;
  - Max-Weight Policy.

Randomized Policies

- **Stationary Randomized Policy:**

  in slot k, select node i with probability μ_i.

Optimization over Randomized policies

$$\mathrm{OPT}_{R^*} = \min_{R \in \Pi_R} \left\{ \frac{1}{M} \sum_{i=1}^{M} \frac{\alpha_i}{p_i \mu_i} \right\} \tag{25a}$$

$$\text{s.t.} \quad p_i \mu_i \geq q_i \;, \forall i \;; \tag{25b}$$

$$\textstyle\sum_{i=1}^{M} \mu_i \leq 1 \;, \forall k \;. \tag{25c}$$

- Solution given by KKT Conditions. Optimal **R*** is **2-optimal.**

Max Weight Policy

- Throughput **debt** at the beginning of slot k: $\boldsymbol{x_i(k)} = (k-1)\boldsymbol{q_i} - \sum_{t=1}^{k-1} d_i(t)$

- **MW policy**: in slot k, schedule node with highest value of

$$\boldsymbol{p_i} \max\{\boldsymbol{x_i(k)}, 0\} + V(\boldsymbol{\alpha_i}/\boldsymbol{\mu_i^*})h_i(k)$$

Theorem: MW policy is $\left(2+\frac{\epsilon}{V}\right)$-optimal in terms of AoI and **satisfies any feasible** throughput constraints $\boldsymbol{q_i}$.

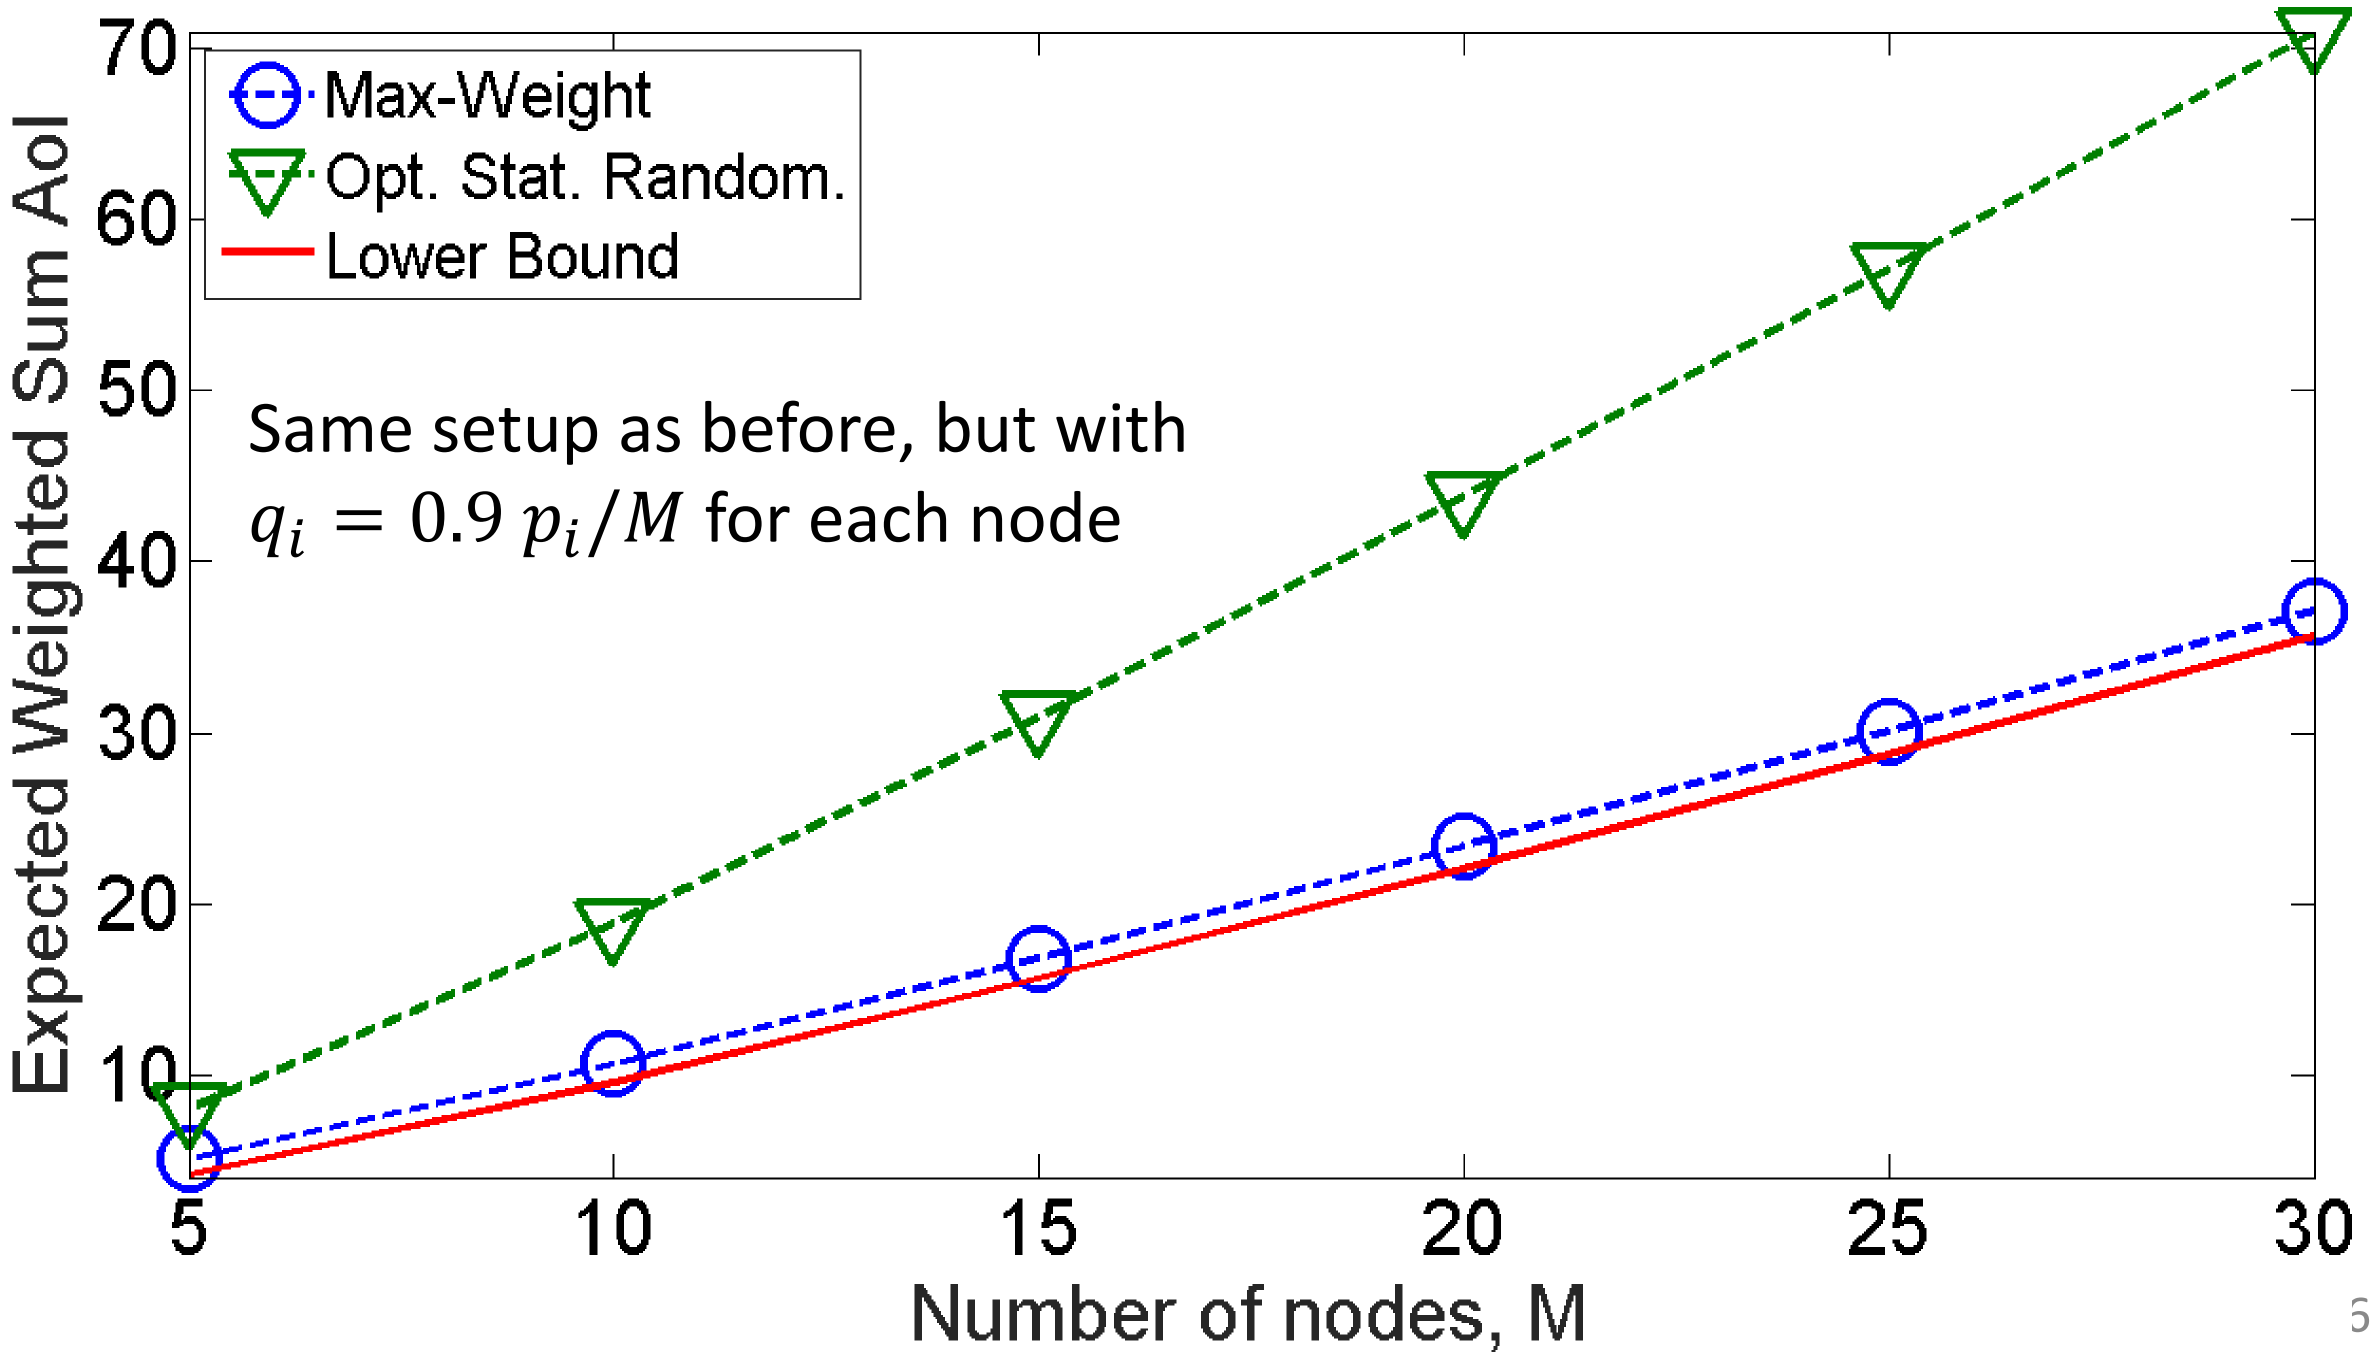
70
60
50
40
30
20
10
5
10
15
20
25
30
Expected Weighted Sum AoI
Number of nodes, M
Max-Weight
Opt. Stat. Random.
Lower Bound
Same setup as before, but with
$q_i = 0.9\ p_i/M$ for each node

Throughput

$M = 2$

$K = 120$

$p_1 = 1/3$

$p_2 = 1/3$

α varies

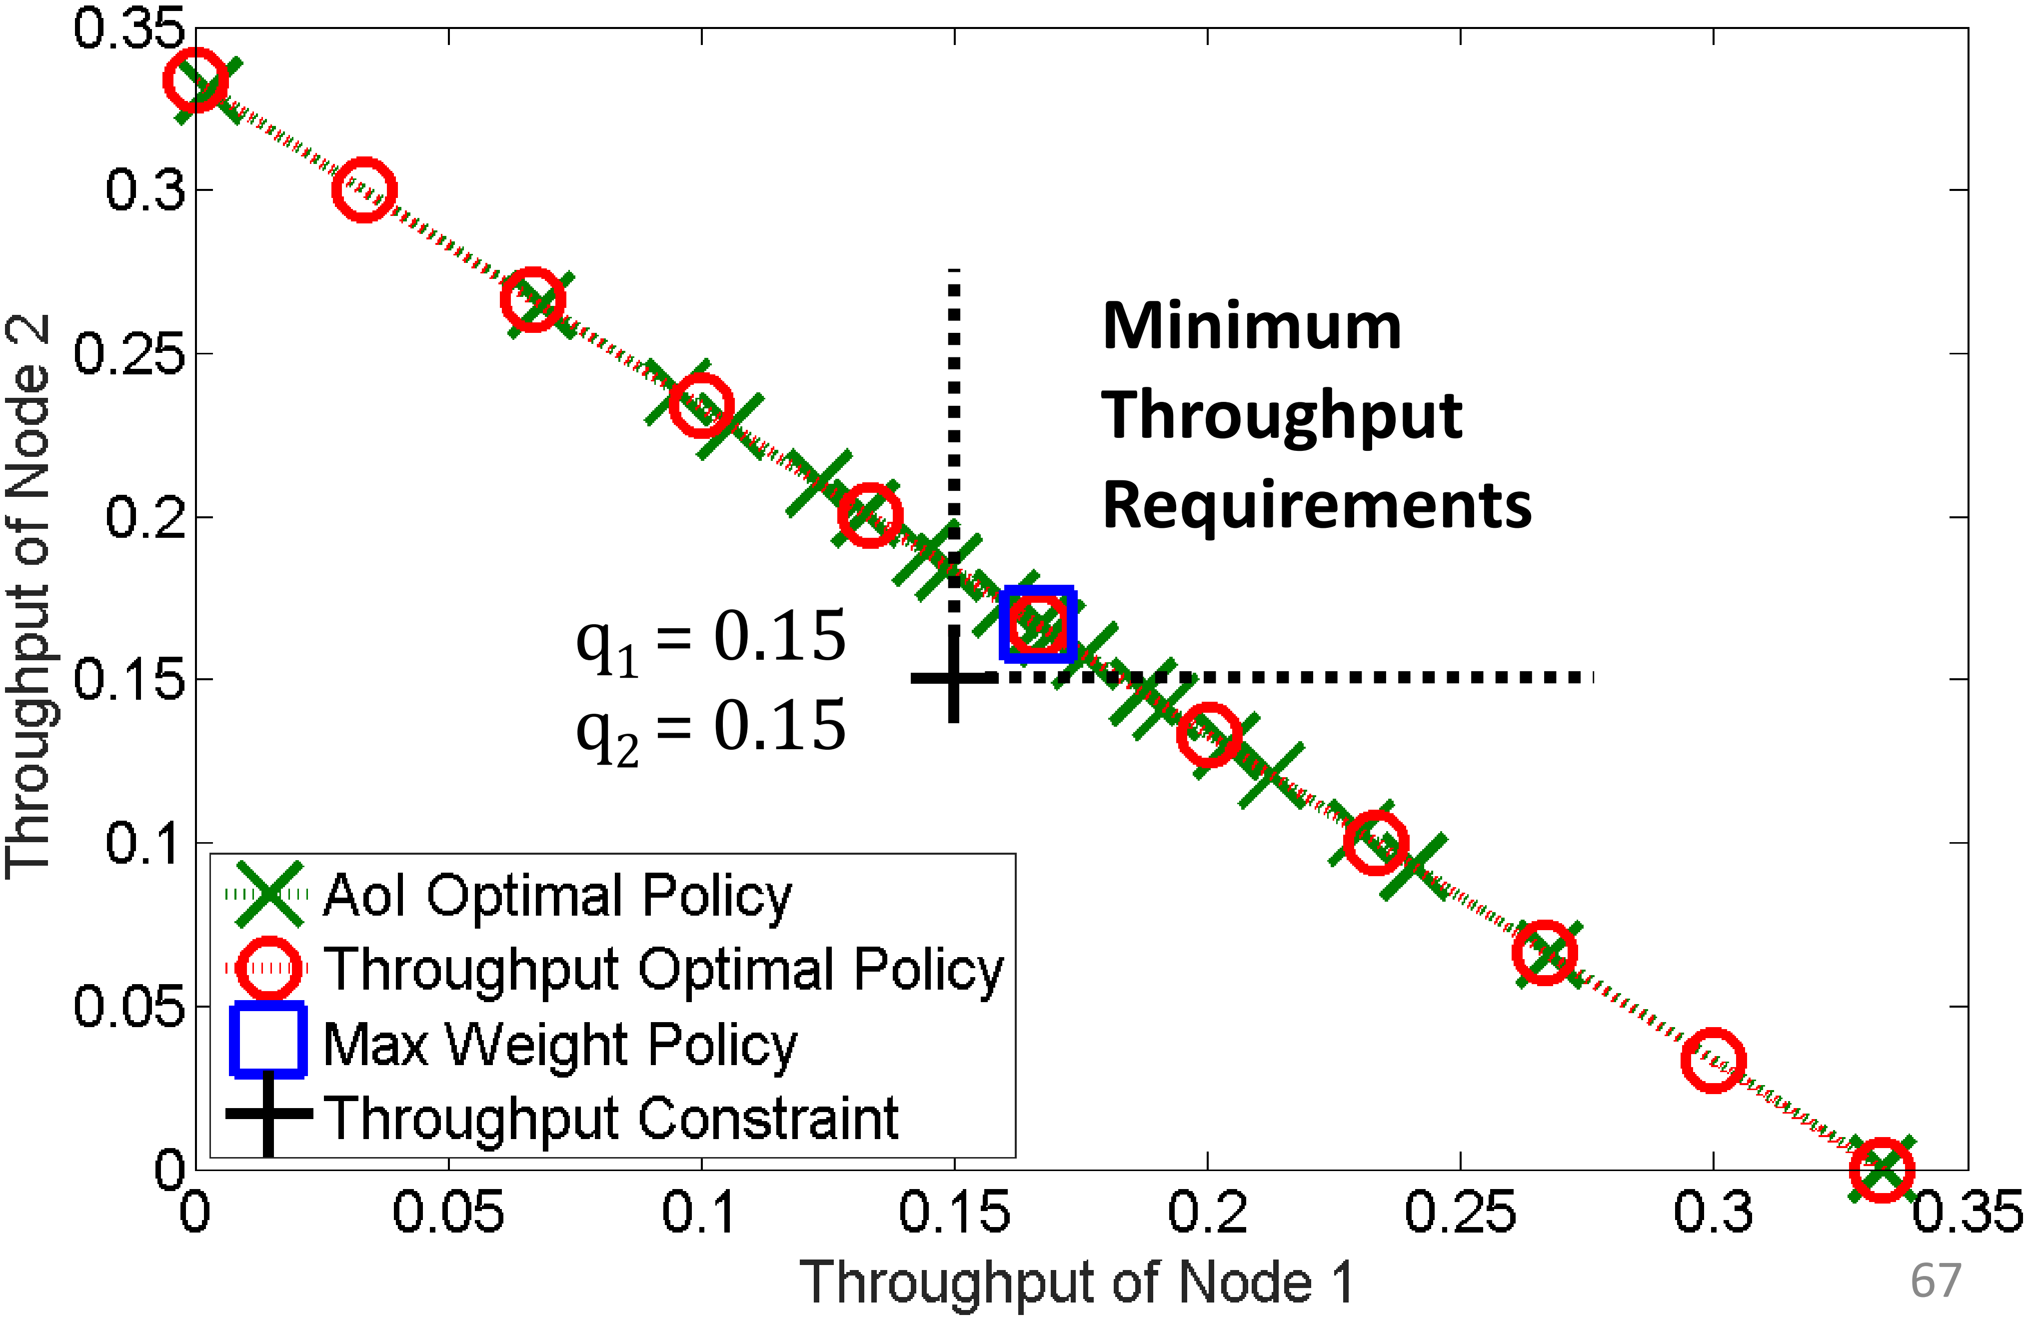

Age of Info.

M = 2

K = 120

$p_1 = 1/3$

$p_2 = 1/3$

α varies

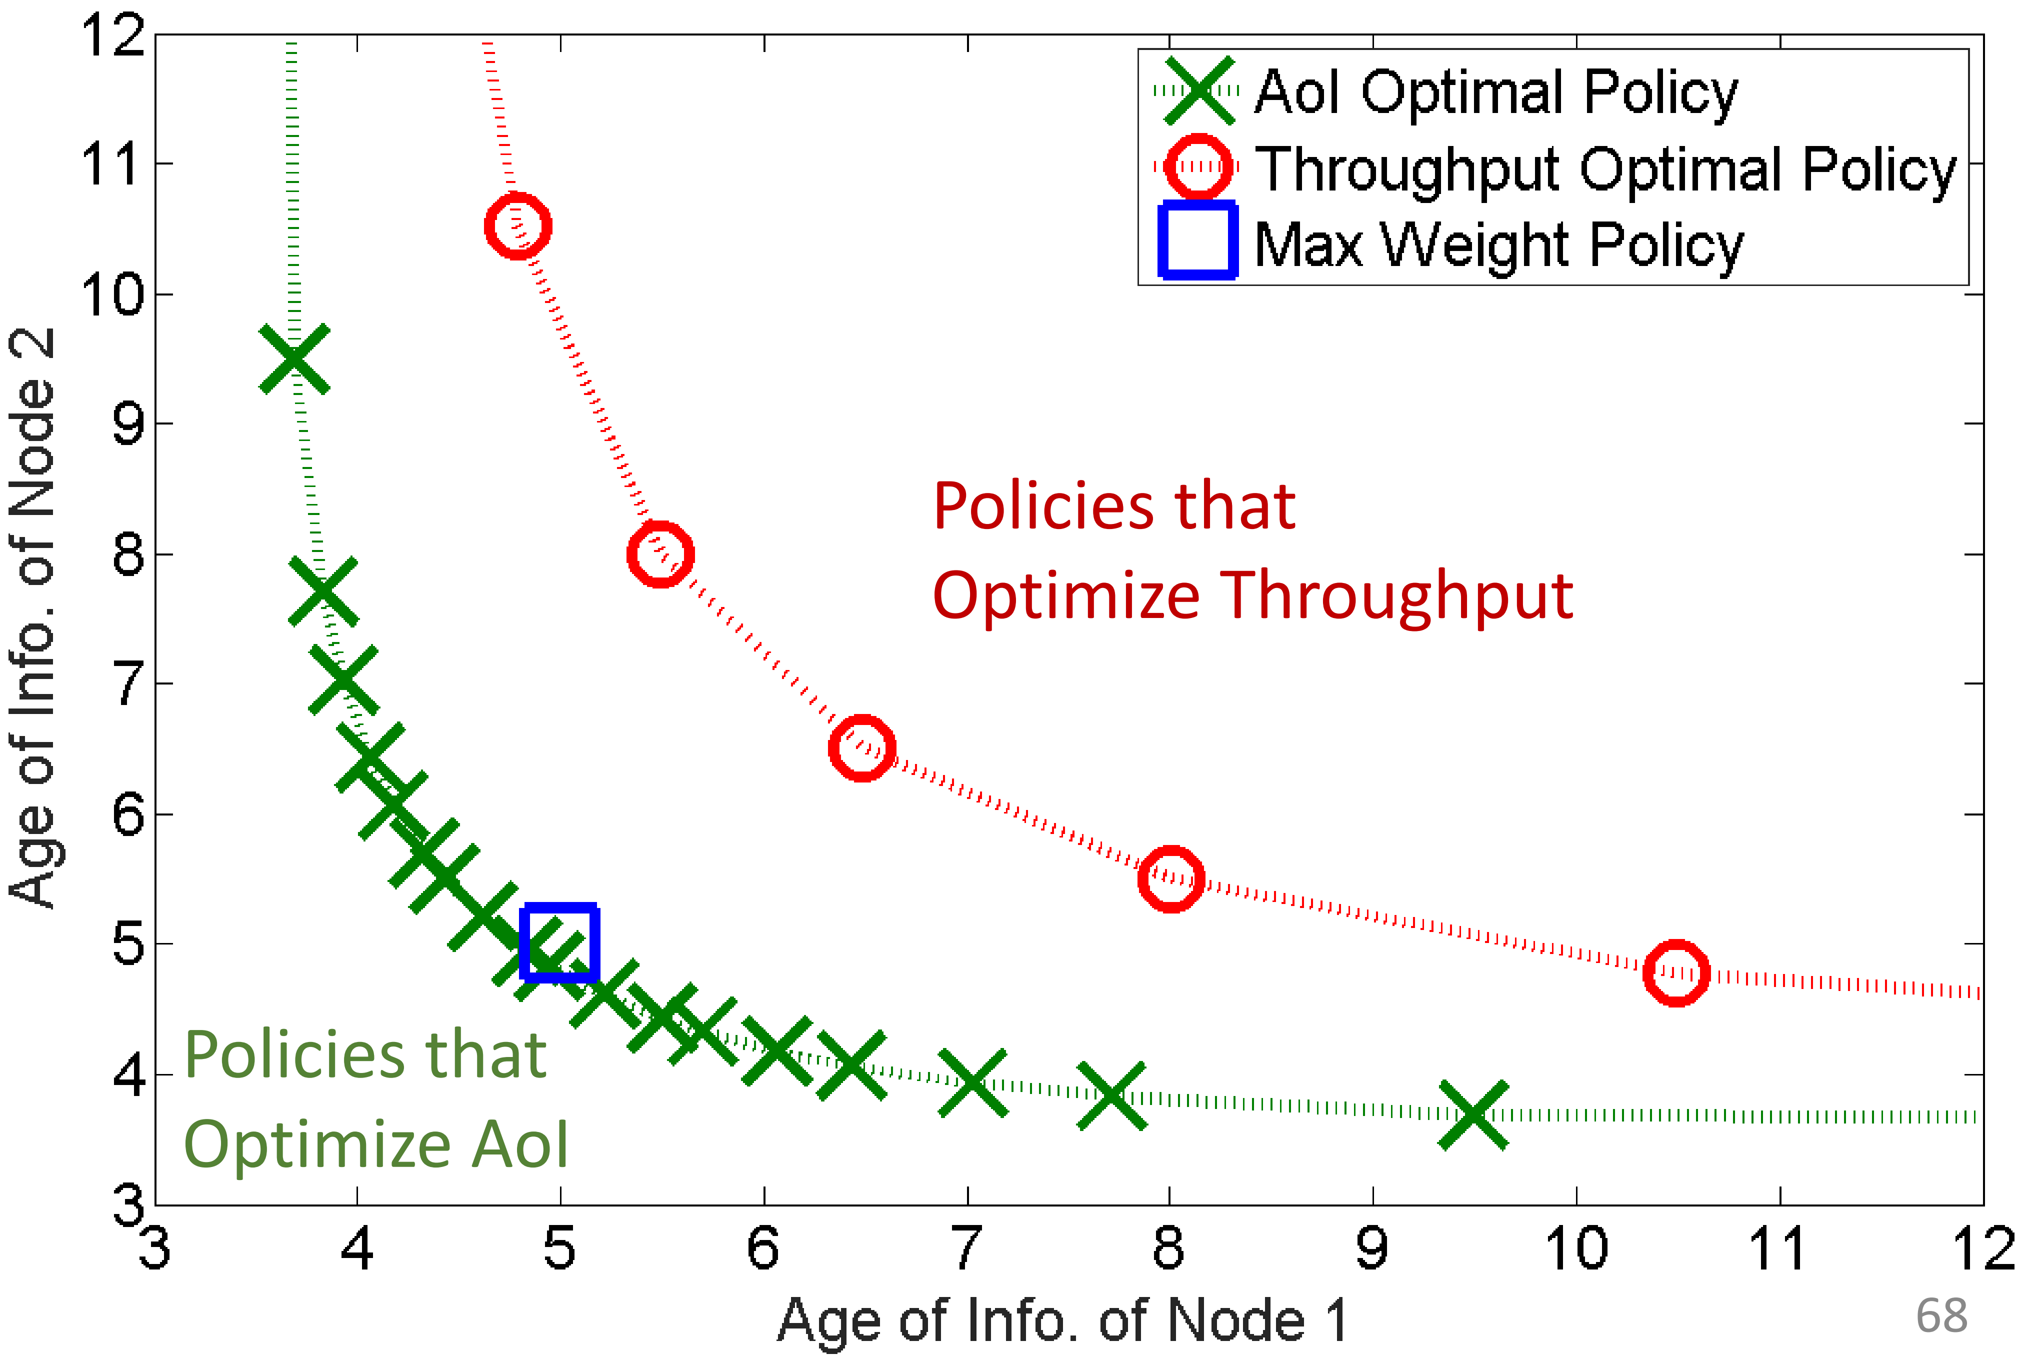

Final Remarks

- We developed and evaluated low-complexity scheduling policies:
  - Greedy, Optimal Stationary Randomized, Max-Weight and Whittle's Index
  - Obtained optimality ratios and validated their performance using Numerical Results

- Takeaways:
  - Age of Information is a measure of information freshness
  - Randomized Policy is the simplest possible policy and it is 2-optimal.
  - Max-Weight has near optimal AoI-performance and it satisfies any feasible throughput constraints.

THANK YOU!