Time-Critical Decisions with Real-Time Information Extraction

Dissertation Defense

Xingran Chen, ESE Department

Advisor: Prof. Shirin Saeedi Bidokhti

Introduction

- Backgrounds

The Real World

Information

Networks

Decisions

Actions

- Examples

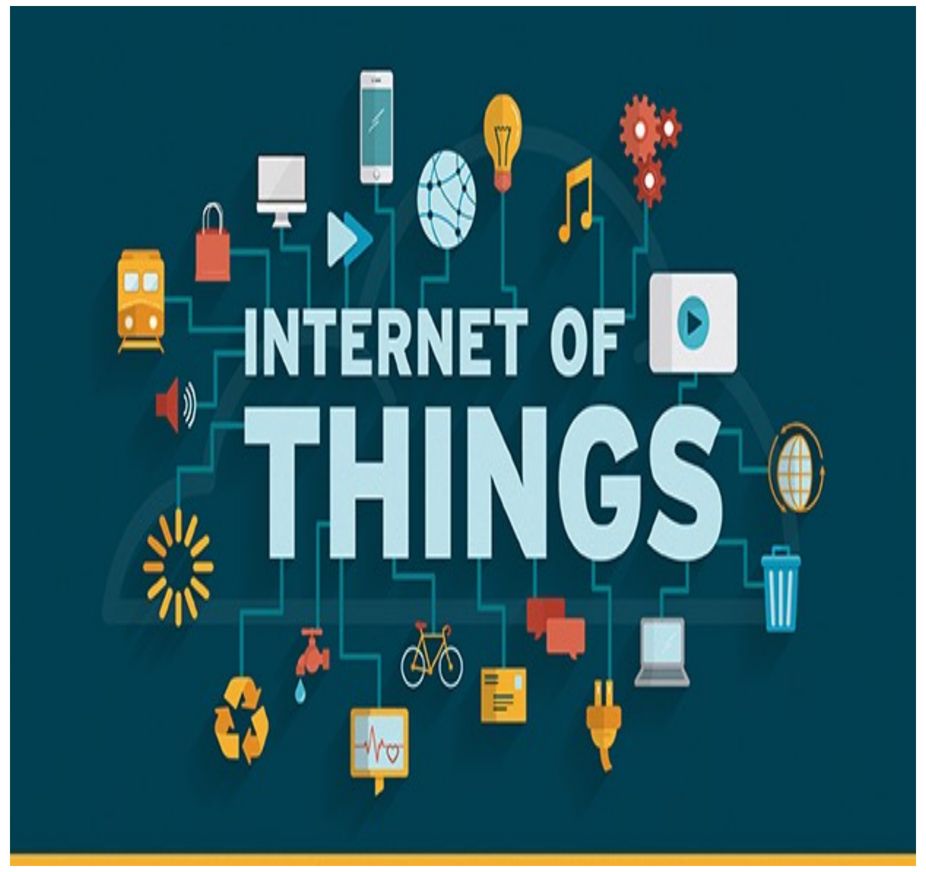

remote sensing, estimation, and control in wireless networks

testing and control strategies to track or contain the disease

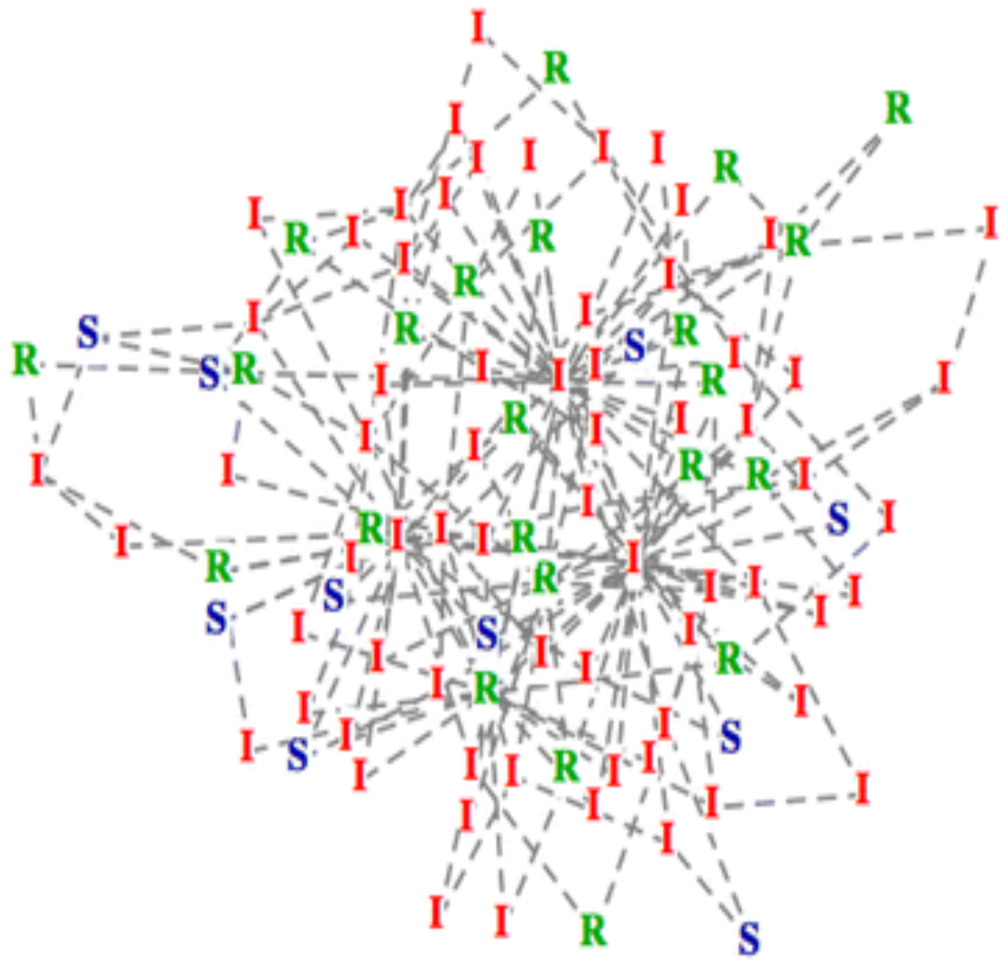

Introduction

- Universal challenges:
  - Sequential decision making, real-time (partial) observations.
  - Contrast between optimal and timely information extraction.

- Entire Goal:

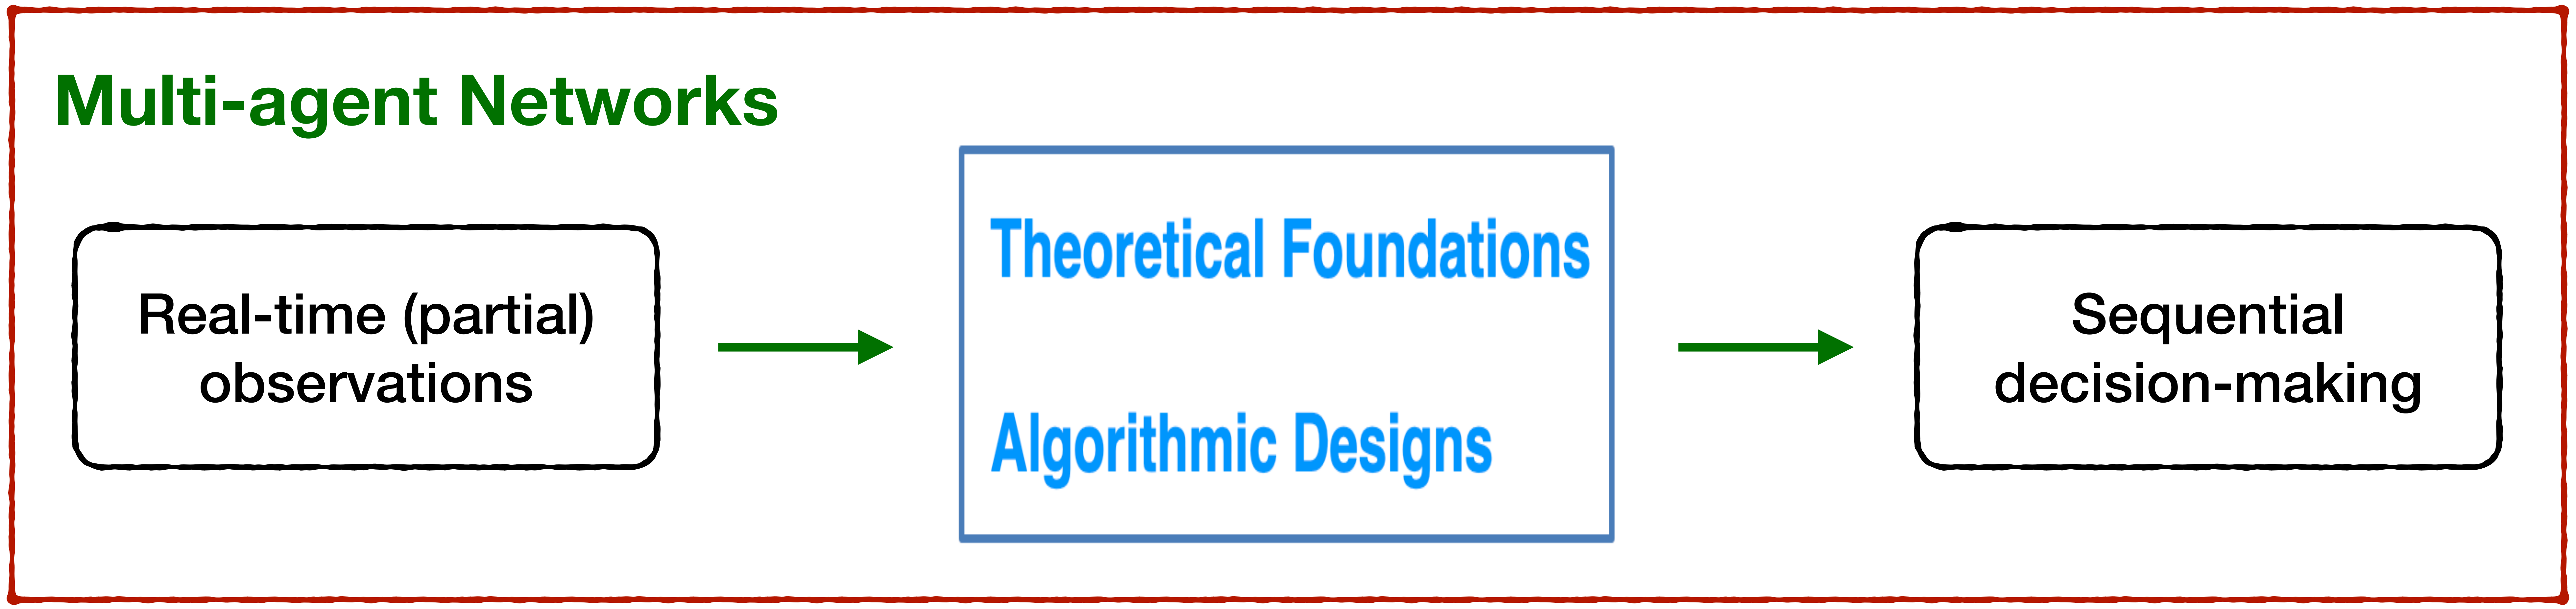

Introduction

- Outline of the Thesis:
  - Chapter 2 — Age of information (AoI): decentralized transmission policies, minimum AoI, random access channels.
  - Chapter 3 — Beyond AoI: decentralized sampling and transmission policies, minimum estimation error, random access channels.
  - Chapter 4 — Extension to Ad-hoc Networks: decentralized sampling and transmission policies, minimum AoI/estimation error, ad-hoc networks.
  - Chapter 5 — Tradeoffs between AoI and rate: broadcast transmission policies, AoI vs. communication rate.
  - Chapter 6 — From AoI to Public Health: testing and isolation policies, processes evolving temporally and spatially, containing of the spread of COVID-19.

Chapters 2 ~ 4

Backgrounds

- Quality of user experience > Quality of service [1].
- Traditional designs:

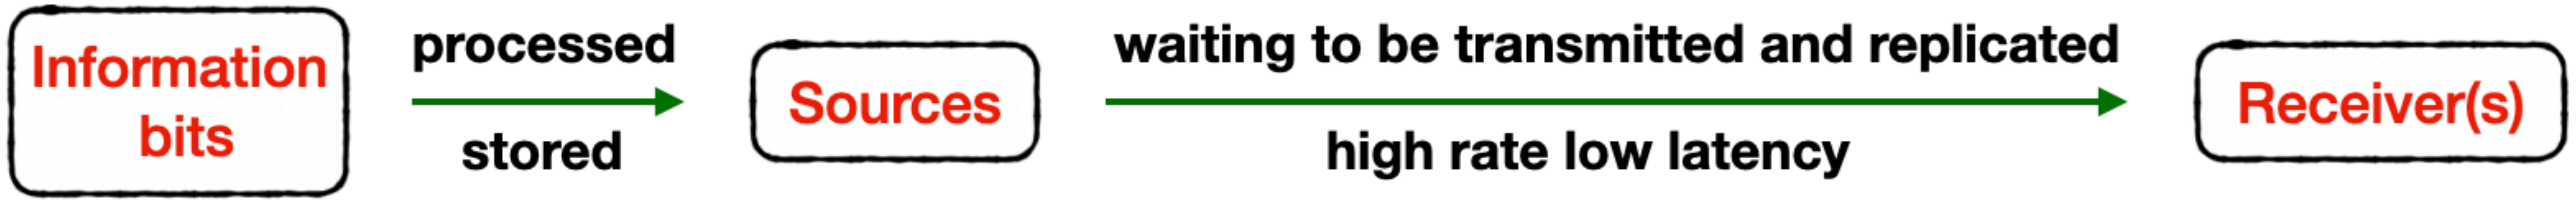

- Now, information is collected and communicated in real-time.
- Go beyond classical metrics: connectivity, rate, reliability, bit error, and latency.
- Age of Information, h_k, quantify the freshness of information [2, 3]
- u_k: generation time of the latest update, $h_k = k - u_k$.

[1] Banerjee-Ulukus, The freshness game: timely communications in the presence of an adversary, 2023.
[2] Kaul-Gruteser-Rai-Kenny, Minimizing age of information in vehicular networks, 2011.
[3] Kaul-Yates-Grusteser, On piggybacking in vehicular networks, 2011.

Chapters 2 ~ 4

Problem Formulation

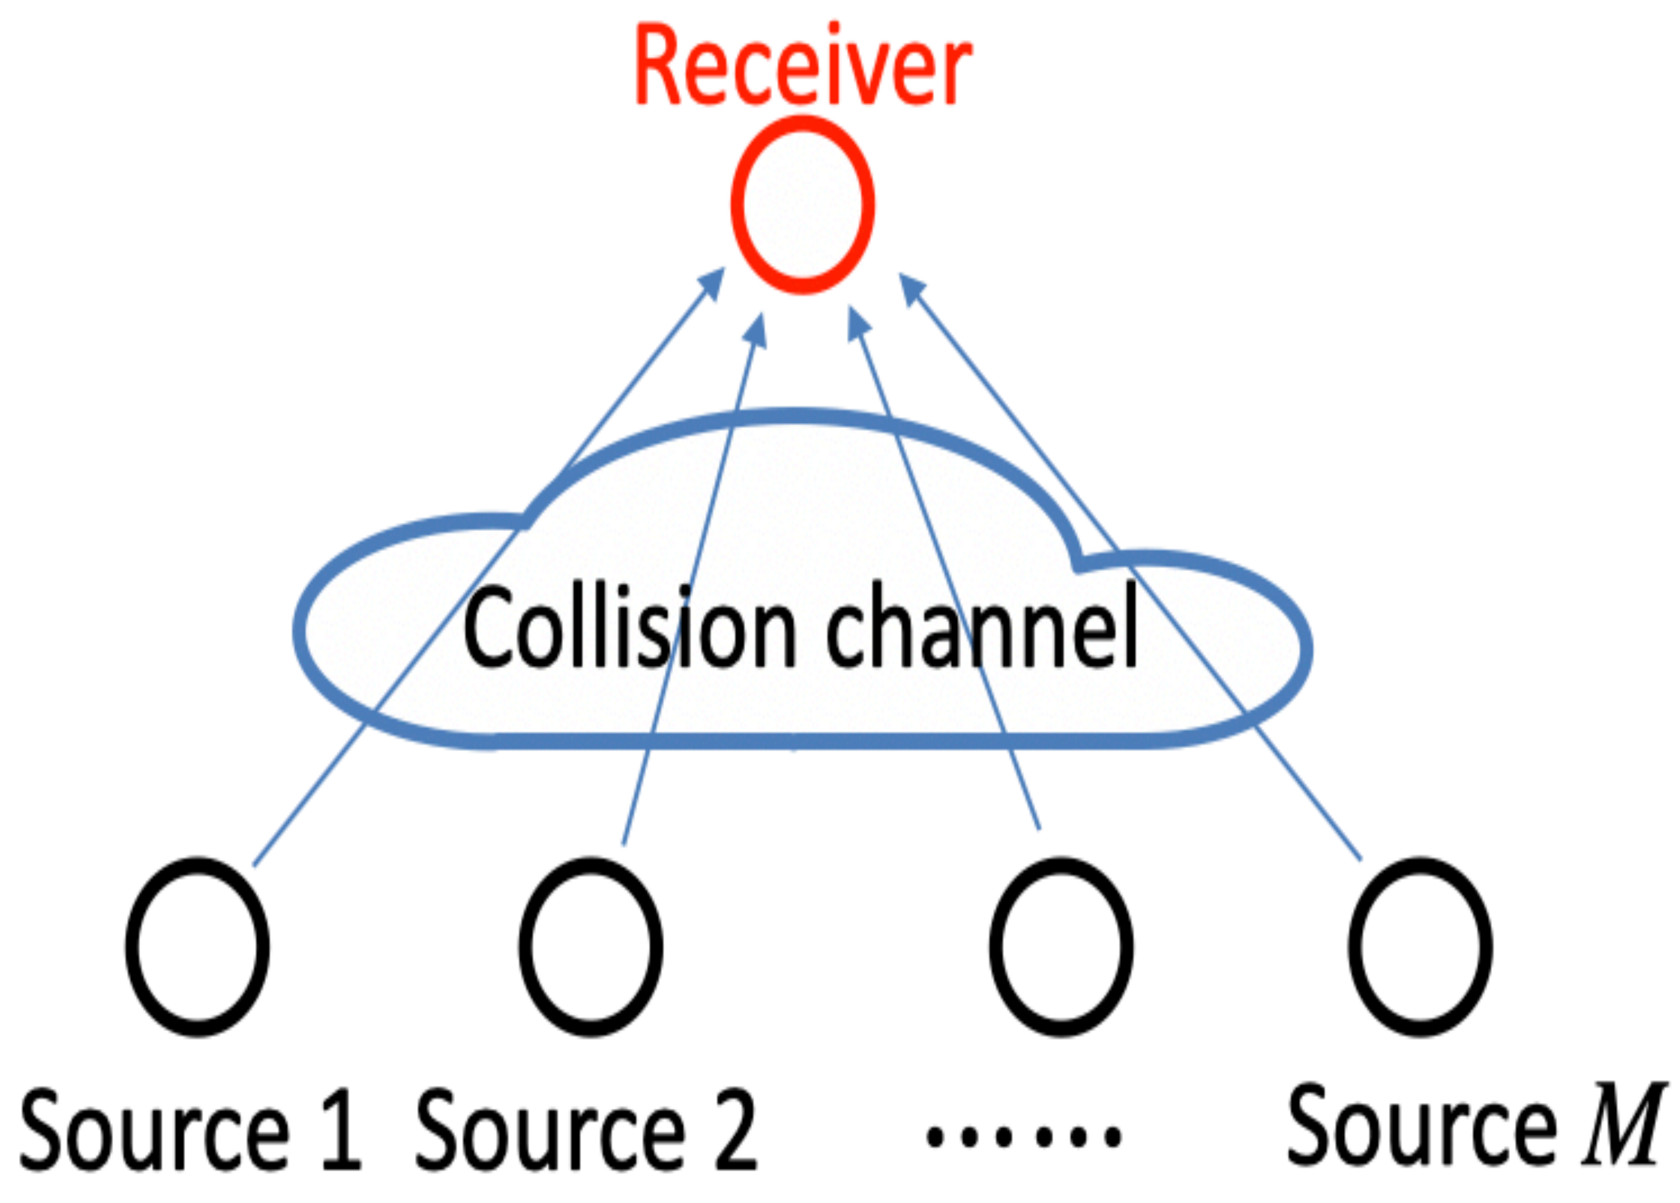

- M identical source nodes
- Physical process $X_{i,k+1} = X_{i,k} + \Lambda_{i,k}$, $\Lambda_{i,k} \sim \mathcal{N}(0,\sigma^2)$
- Collision channel, collision feedback
- Minimum mean square error estimator: $\hat{X}_{i,k}$
- Age of Information for source i: $h_{i,k}$
- Deciding in a decentralized manner.
- **Goal**: minimize normalized average estimation error (NAEE)

$$L(M) = \min_{\pi} \lim_{K\to\infty} \mathbb{E}\Big[\frac{1}{KM^2} \sum_{k=1}^{K} \sum_{i=1}^{M} \big(X_{i,k} - \hat{X}^{\pi}_{i,k}\big)^2\Big]$$

Chapters 2 ~ 4

Motivations & Literature Review

- Sampling: single-user systems [1,2,3,4,5] → Our works: multi-user systems
- Reliable vs. Timely Communication: centralized policies [6,7,8,9] → Our works: decentralized policies
- Distributed decision making: no collision feedback [10,11,12] → Our works: collision feedback
- Ad-hoc networks: centralized policies or decentralized stationary randomized policies [13,14] → Our works: general decentralized policies

[1] Rabi-Moustakides-Baras, Adaptive sampling for linear state estimation, 2012.
[2] Lipsa-Martins, Remote state estimation with communication costs for first-order LTI systems, 2011.
[3] Molin-Hirche, Event-triggered state estimation: an iterative algorithm and optimality properties, 2017.
[4] Nayyar-Basar-Teneketzis-Veeravalli, Communication scheduling and remote estimation with energy harvesting sensor, 2012.
[5] Chakravorty-Mahajan, Remote estimation over a packet-drop channel with Markovian state, 2020.
[6] Talak-Modiano, Age-delay tradeoffs in queueing systems, 2021.
[7] Sun-Polyanskiy-Uysal Biyikoglu, Remote estimation of the Wiener process over a channel with random delay, 2020.
[8] Kadota-Sinha-Modiano, Scheduling algorithms for optimizing age of information in wireless networks with throughput constraints, 2019.
[9] Kadota-Modiano, Minimizing the age of information in wireless networks with stochastic arrivals, 2019.
[10] Gatsis-Pajic-Ribeiro-Pappas, Opportunistic control over shared wireless channels, 2015.
[11] Taricco, Joint channel and data estimation for wireless sensor networks, 2012.
[12] Zhang-Vasconcelos-Cui-Mitra, Distributed remote estimation over the collision channel with and without local communication, 2022.
[13] Tripathi-Talak-Modiano, Information freshness in multihop wireless networks, 2023.
[14] Jones-Modiano, Minimizing age of information in spatially distributed random access wireless networks, 2022.

Chapters 2 ~ 4

Oblivious Policies, Non-oblivious Policies, and Age of Information

- Oblivious Policies: actions do not depend on samples ($X_{i,k}$).
- Non-oblivious Policies: actions depend on samples ($X_{i,k}$).

Lemma 1. In oblivious policies, for any node i, $\mathbb{E}\left[(X_{i,k} - \hat{X}_{i,k})^2\right] = \mathbb{E}[h_{i,k}]\sigma^2$.

- In oblivious policies, minimization of estimation error ⇔ minimization of AoI
- Minimize normalized average AoI (NAAoI)

$$J(M) = \min_{\pi} \lim_{K \to \infty} \mathbb{E}\left[\frac{1}{KM^2} \sum_{k=1}^{K} \sum_{i=1}^{M} h_{i,k}^{\pi}\right]$$

Chapters 2 ~ 4

Oblivious Policies and Age of Information

$$J(M) = \min_{\pi} \lim_{K\to\infty} \mathbb{E}\left[\frac{1}{KM^2}\sum_{k=1}^{K}\sum_{i=1}^{M} h_{i,k}^{\pi}\right]$$

- More general setting: generating packets by a Bernoulli process with θ
- When θ is small

Theorem 1. Suppose $\theta < 1/eM$ and define $\eta = \lim_{M\to\infty} M\theta$. Any slotted ALOHA scheme [1] achieves $\lim_{M\to\infty} J^{SA}(M) = 1/\eta$. → lower bound, optimality

- When θ is large: Select nodes through a thinning process
- age-gain [2] → age reduction when receiving a new packet
- Thinning process: age-gain > $\mathtt{T}(k)$ → active, slotted ALOHA
- $\mathtt{T}(k)$ → adaptive thinning, $\mathtt{T}^*$ → stationary thinning

[1] Bertsekas-Gallager, Data Networks, 2nd ed. Hoboken, NJ, USA: Prentice-Hall, 1992.
[2] Kadota-Modiano, Minimizing the age of information in wireless networks with stochastic arrivals, 2021.

Chapters 2 ~ 4

Oblivious Policies and Age of Information

$$J(M) = \min_{\pi} \lim_{K\to\infty} \mathbb{E}\left[\frac{1}{KM^2}\sum_{k=1}^{K}\sum_{i=1}^{M} h_{i,k}^{\pi}\right]$$

Theorem 2. For stationary thinning, $\mathrm{T}^* = \lfloor eM - 1/\theta + 1 \rfloor$. For any $\theta = 1/o(M)$, $\lim_{M\to\infty} J(M) = e/2$.

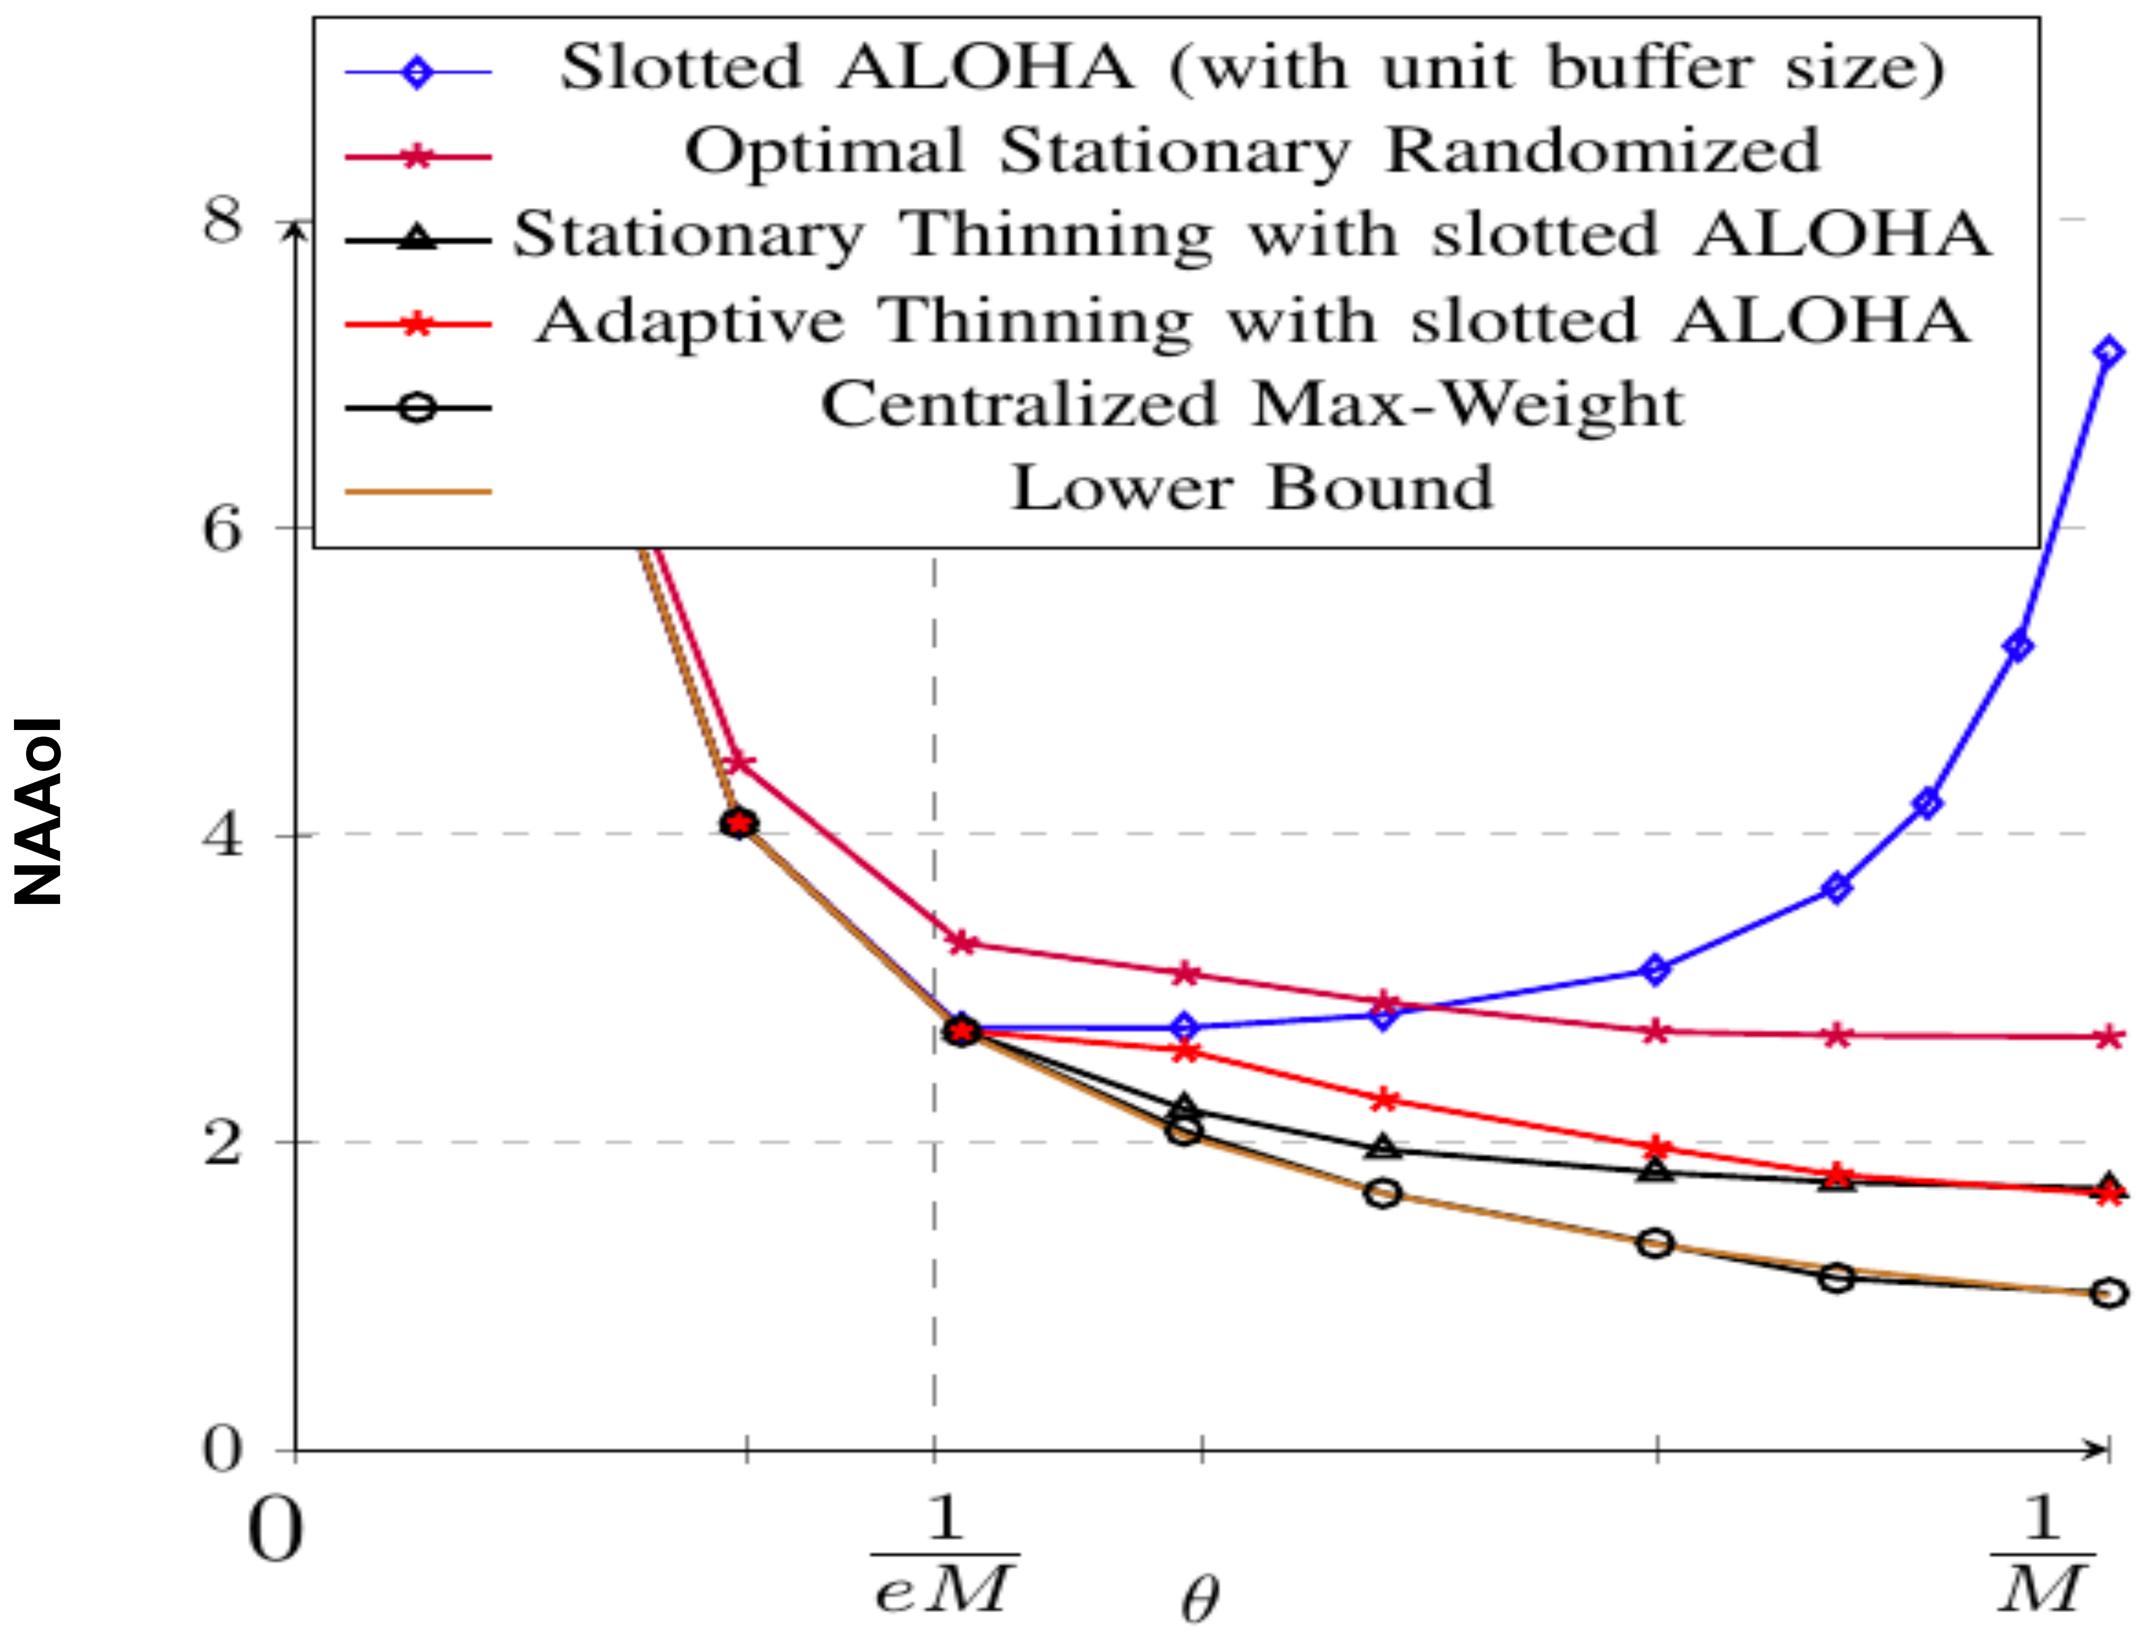

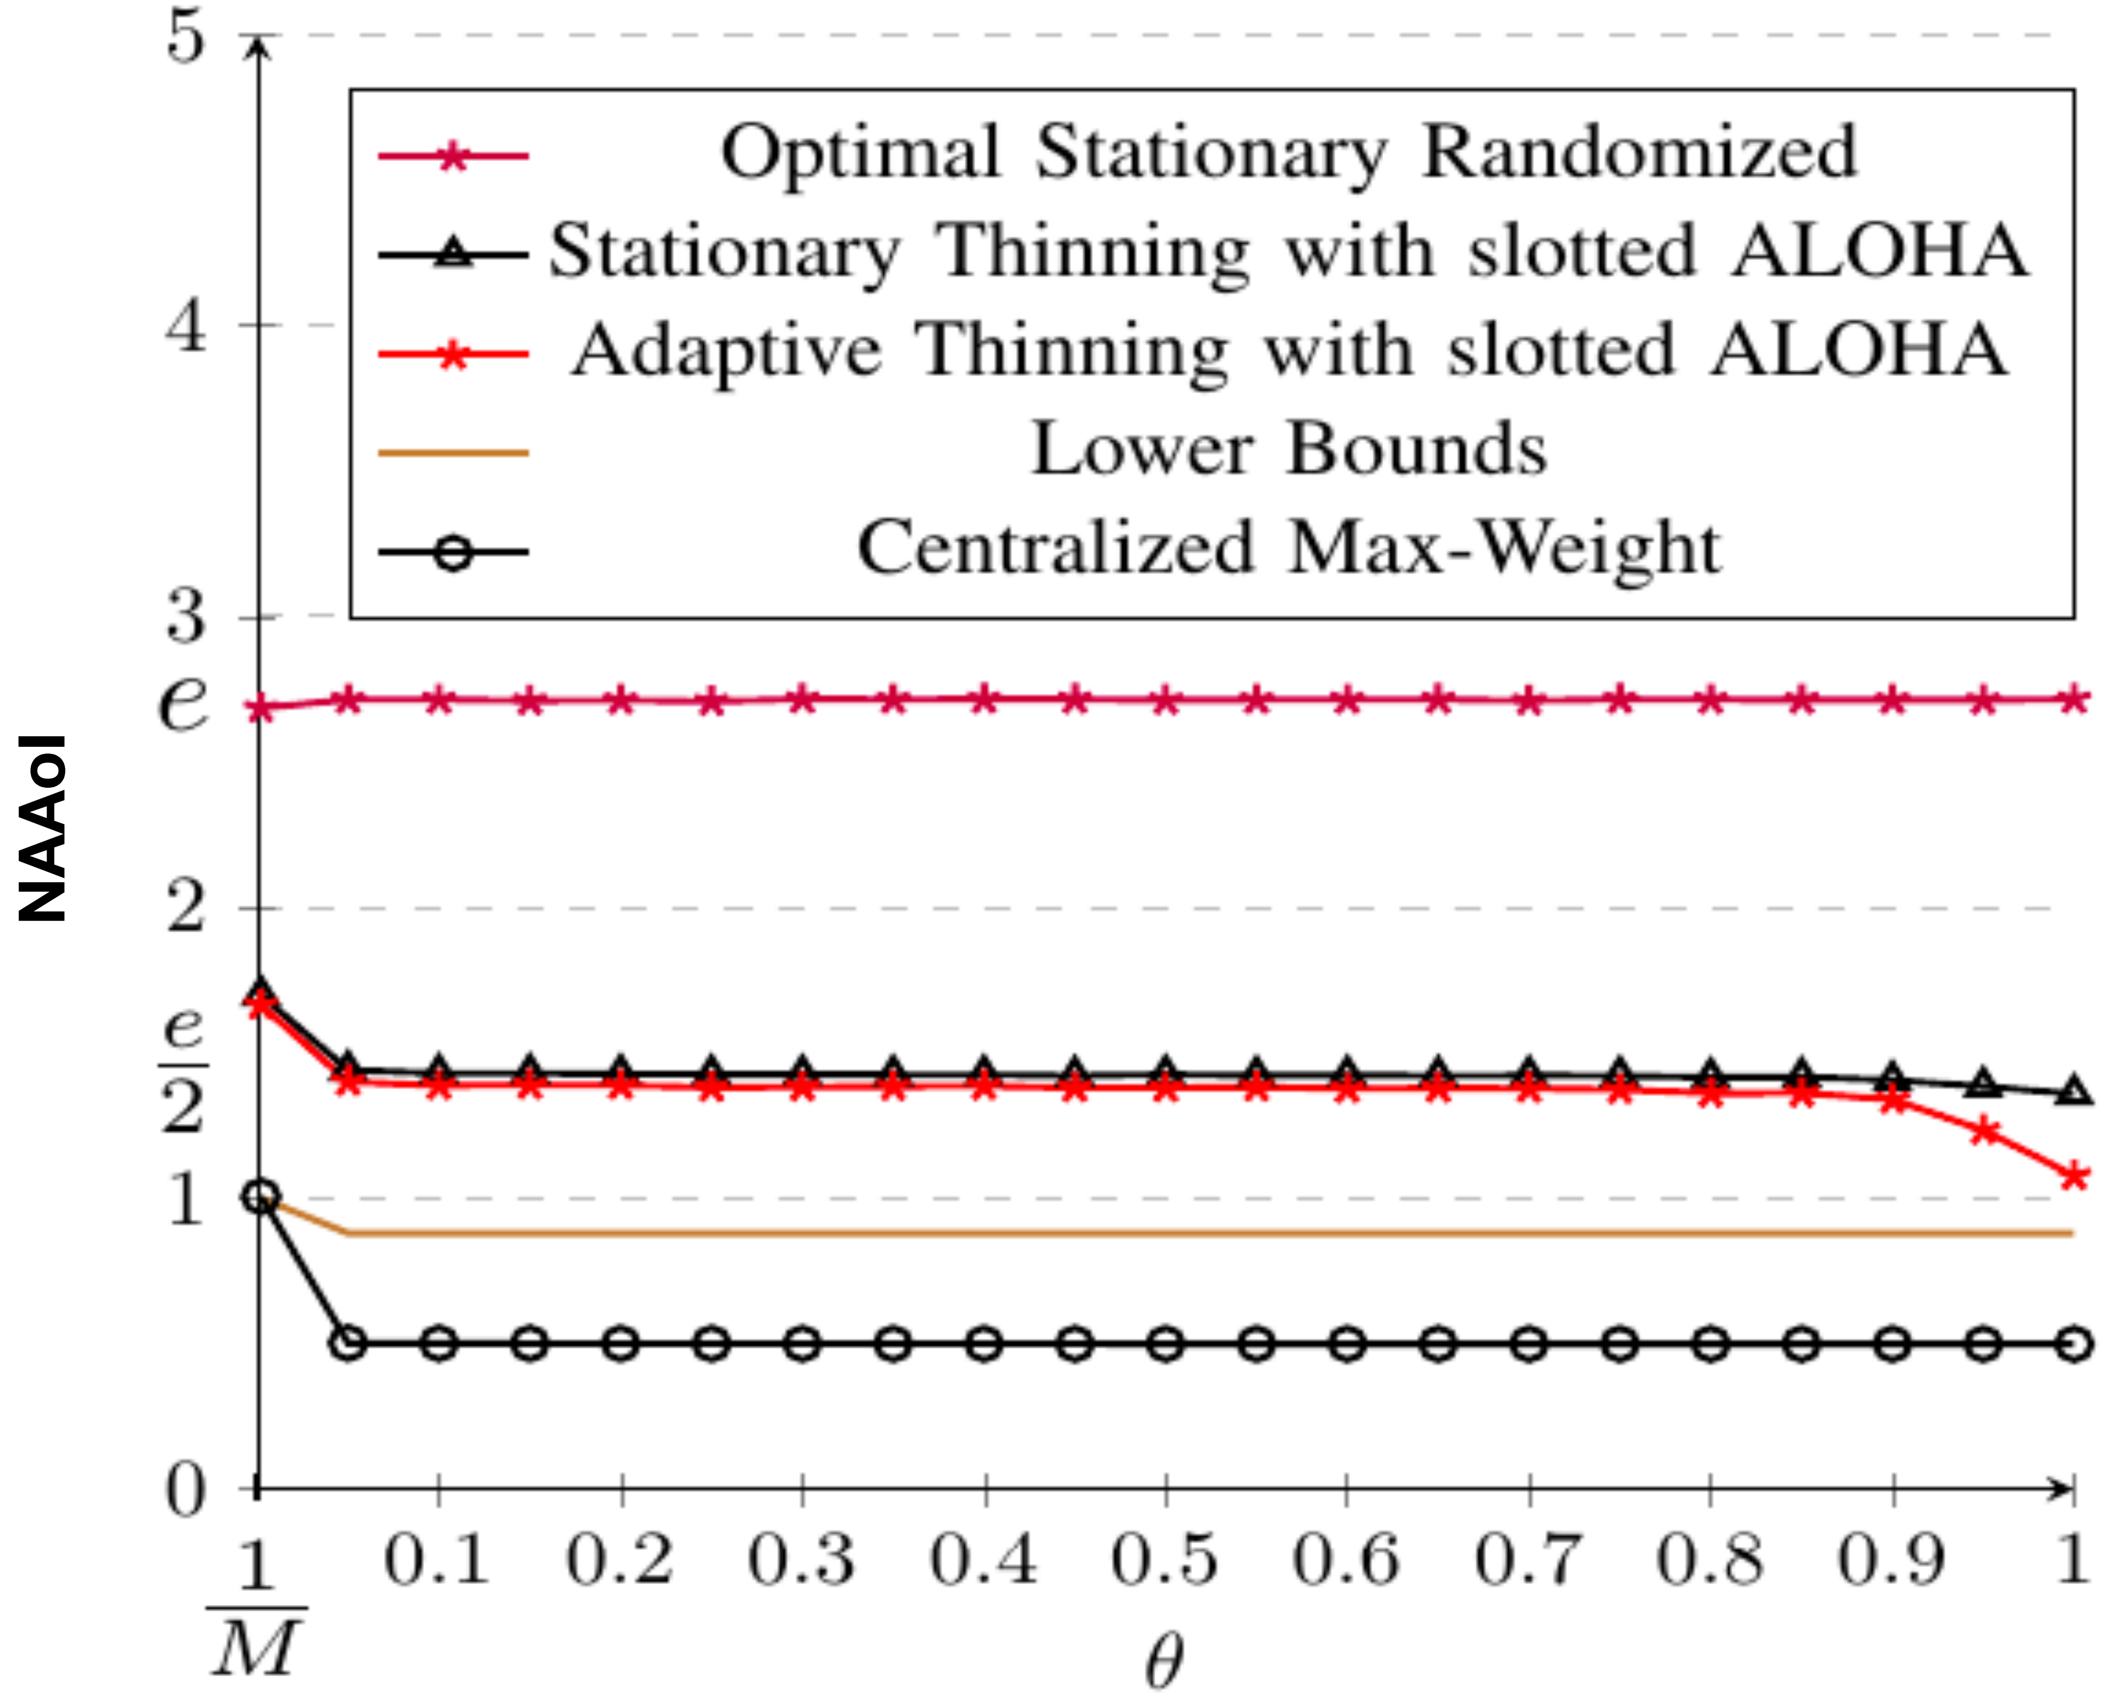

Chapters 2 ~ 4

Oblivious Policies and Age of Information

$$J(M) = \min_{\pi} \lim_{K\to\infty} \mathbb{E}\Big[\frac{1}{KM^2}\sum_{k=1}^{K}\sum_{i=1}^{M} h_{i,k}^{\pi}\Big]$$

- Stationary thinning can be applied to other transmission policies.
- Theoretical guarantees & simulations are provided.

- Achievements:

- [C1] X. Chen, K. Gatsis, H. Hassani and S. Saeedi-Bidokhti, Age of Information in Random Access Channels, IEEE ISIT, 2020.
- [J1] X. Chen, K. Gatsis, H. Hassani and S. Saeedi-Bidokhti, Age of Information in Random Access Channels, IEEE TIT, 2022.
- IEEE Communications Society & Information Theory Society Joint Paper Award 2023

Chapters 2 ~ 4

Non-oblivious Policies

$$L(M) = \min_{\pi} \lim_{K \to \infty} \mathbb{E}\Big[\frac{1}{KM^2} \sum_{k=1}^{K} \sum_{i=1}^{M} \big(X_{i,k} - \hat{X}^{\pi}_{i,k}\big)^2\Big]$$

- Error process: $\psi_i(k) = |X_i(k) - \hat{X}_i(k)|$
- Error-based Thinning: $\psi_i(k) \geq \beta \rightarrow$ active, slotted ALOHA.
- Minimize $L(M) \rightarrow$ find optimal β^*

Theorem 3. Let M be sufficiently large. The optimal β^* is approximately given by $\beta^* \approx \sigma\sqrt{eM}$, and $\hat{L} \approx e\sigma^2/6$.

- For oblivious policy, $J^{\text{oblivious}} \approx e/2$, $L^{\text{oblivious}} \approx e\sigma^2/2$

Proposition 1. For large M, $L^{\text{oblivious}}/\hat{L} \approx 3$.

Chapters 2 ~ 4

Non-oblivious Policies

$$L(M) = \min_{\pi} \lim_{K\to\infty} \mathbb{E}\Big[\frac{1}{KM^2}\sum_{k=1}^{K}\sum_{i=1}^{M}\big(X_{i,k}-\hat{X}^{\pi}_{i,k}\big)^2\Big]$$

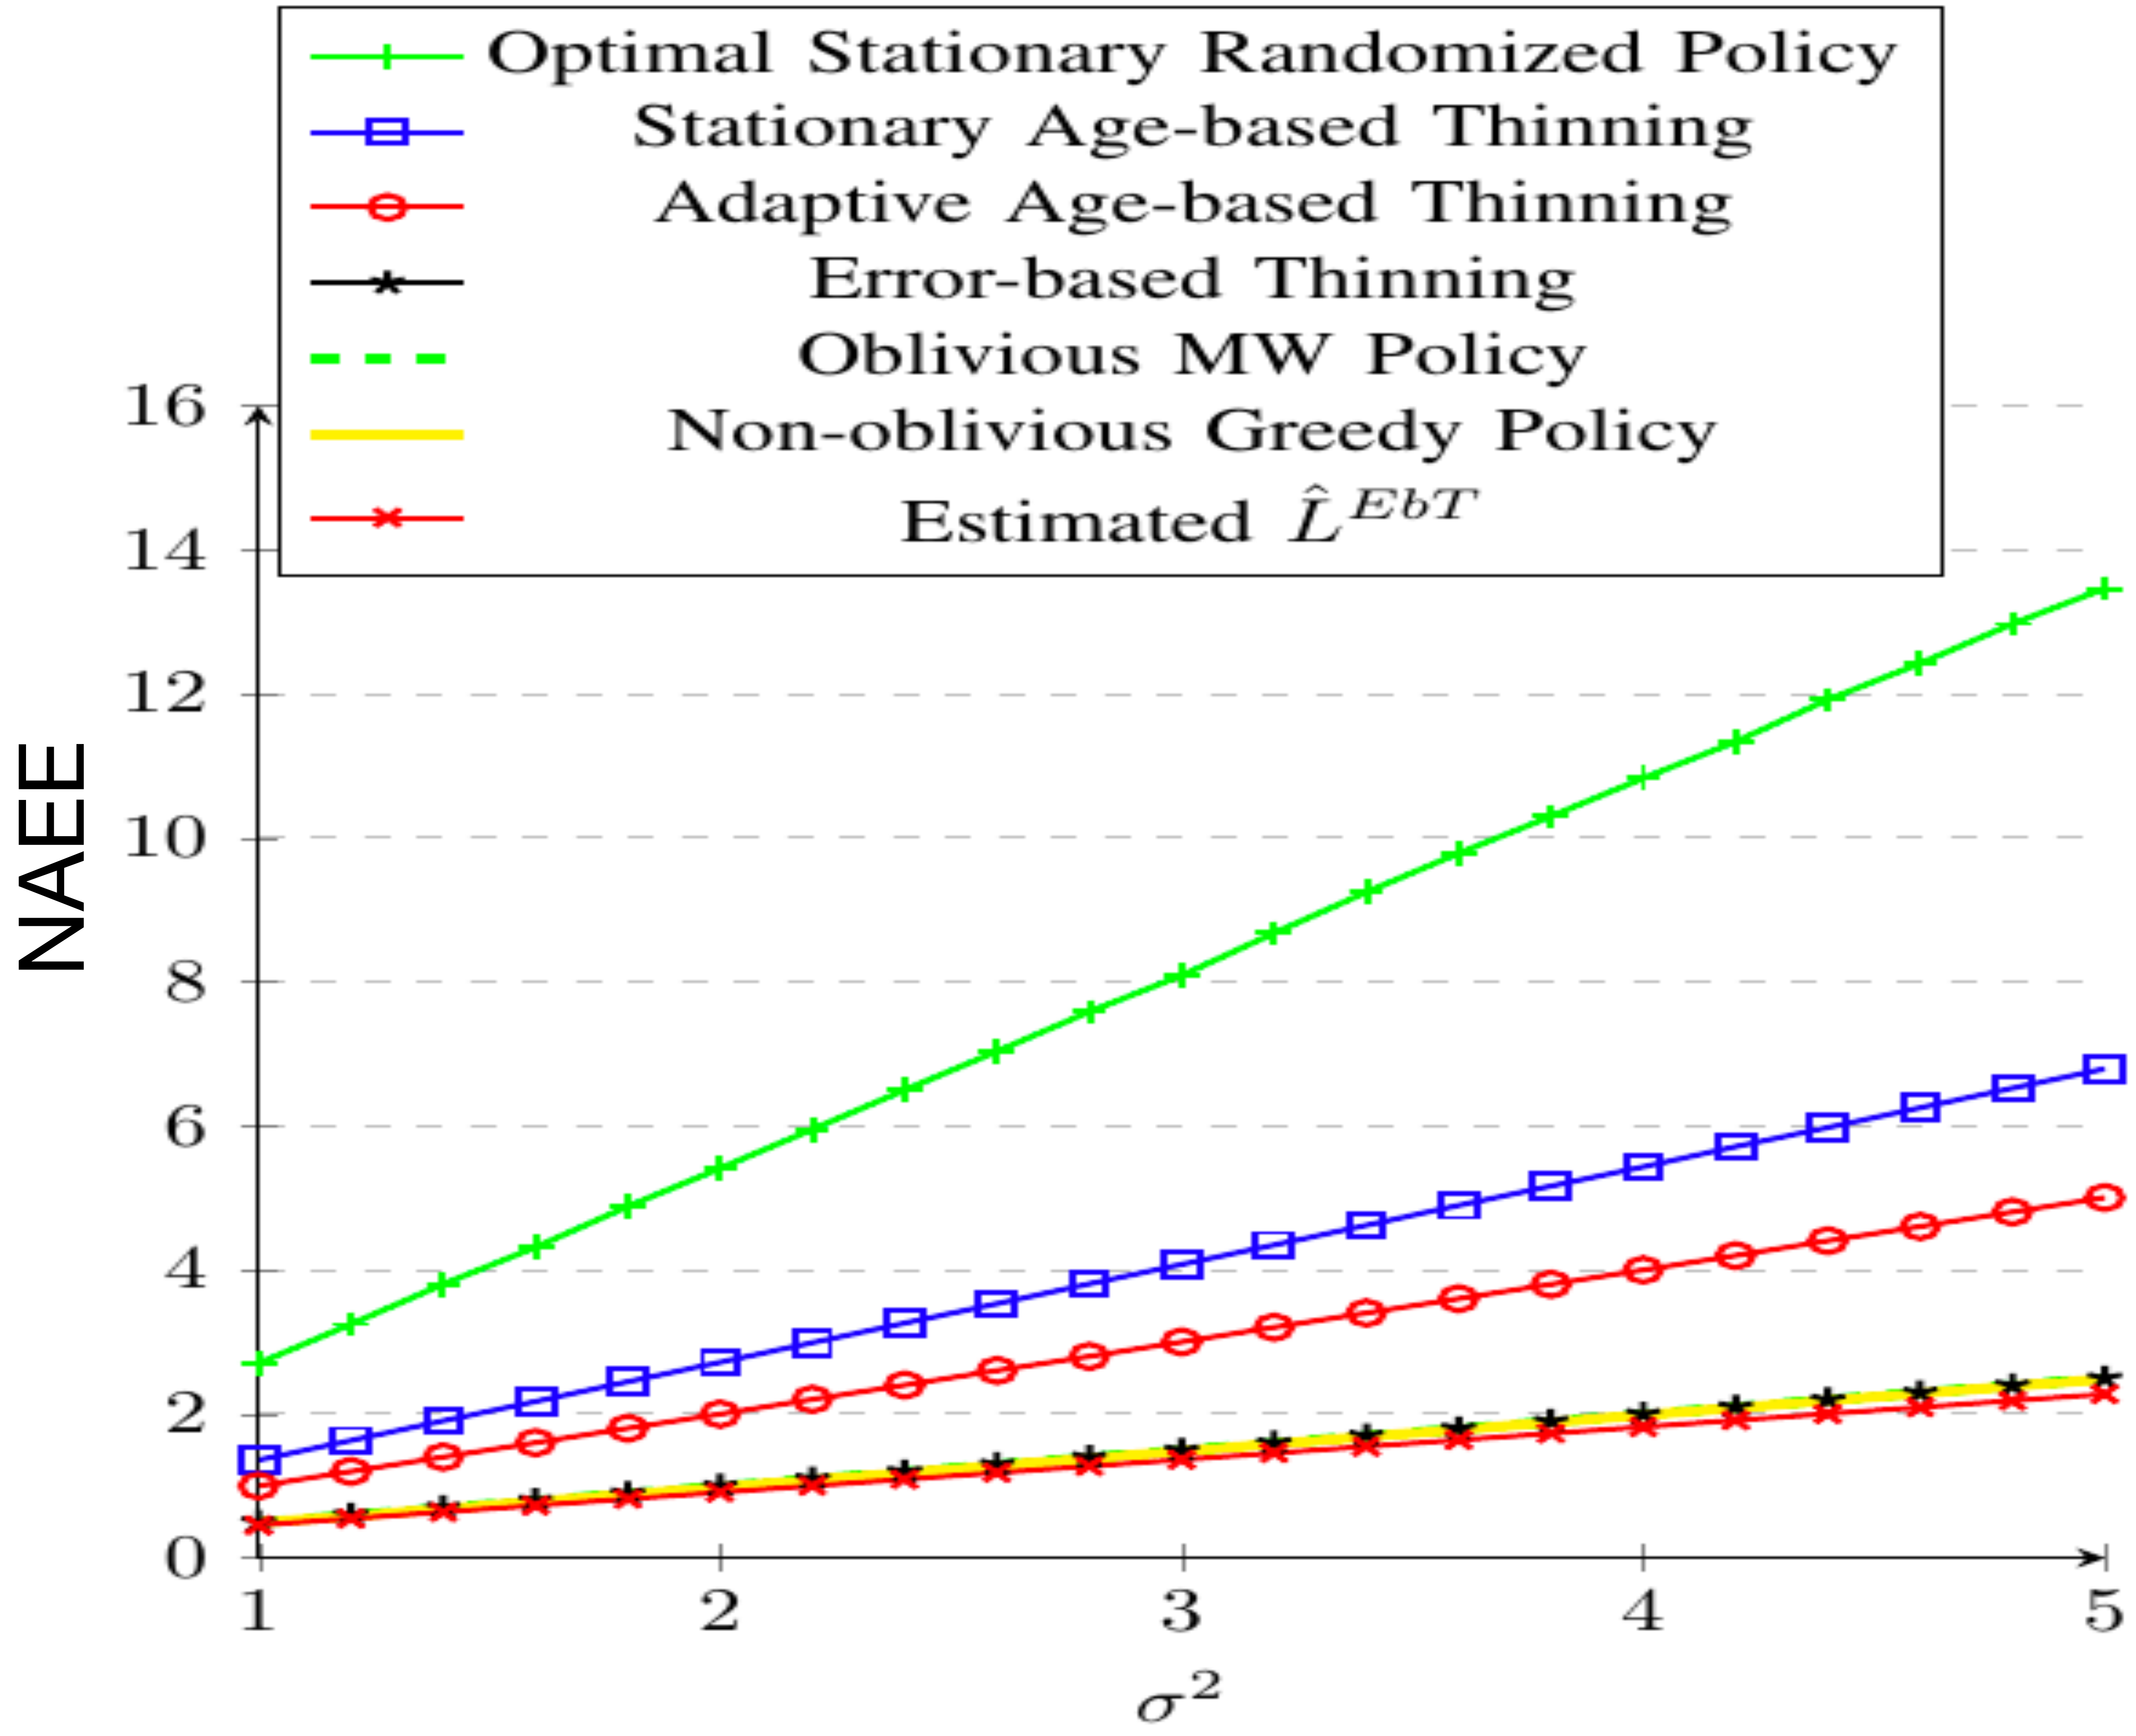

Chapters 2 ~ 4

Non-oblivious Policies

$$L(M) = \min_{\pi} \lim_{K\to\infty} \mathbb{E}\Big[\frac{1}{KM^2}\sum_{k=1}^{K}\sum_{i=1}^{M}\big(X_{i,k} - \hat{X}^{\pi}_{i,k}\big)^2\Big]$$

- The framework can be extended to two general settings:
- Extension 1: $X_{i,k+1} = \gamma X_{i,k} + \Lambda_{i,k}$, $\Lambda_{i,k} \sim \mathcal{N}(0,\sigma^2)$, $\gamma > 0$.
- Extension 2: unreliable channels, packets are erased with an erasure probability.

- Achievements:

- [C2] X. Chen, X. Liao, and S. Saeedi-Bidokhti, Real-time Sampling and Estimation on Random Access Channels: Age of Information and Beyond, IEEE INFOCOM, 2021.
- [J2] X. Chen, X. Liao, and S. Saeedi-Bidokhti, Beyond AoI: Real-time Sampling and Estimation on Reliable and Unreliable Random Access Channels, IEEE/ACM ToN, submitted.
- IEEE INFOCOM Student Conference Award

Chapters 2 ~ 4

General Setting in Ad-hoc Networks

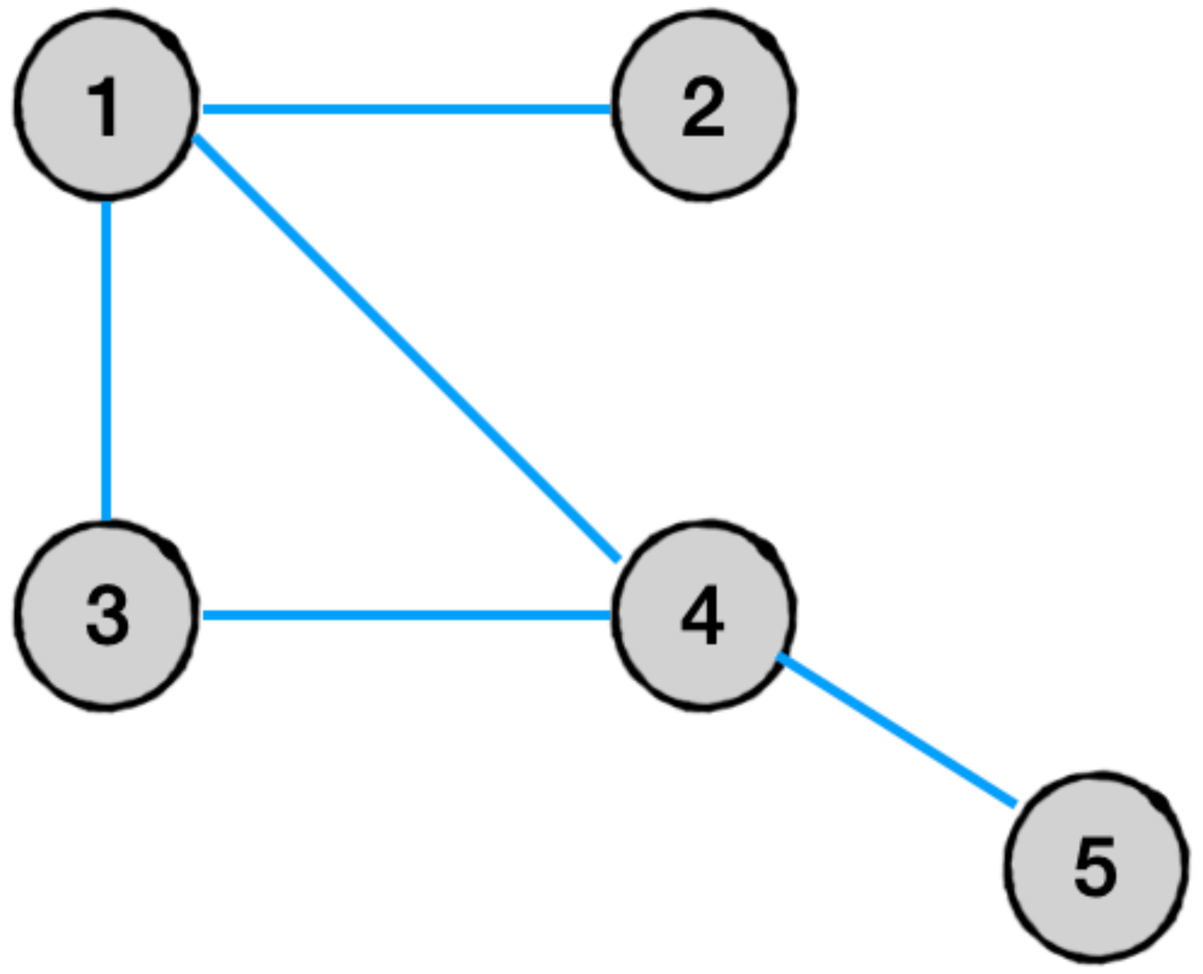

- Consider an ad-hoc (connected) network with M sources.
- Physical process: $X_{i,k+1} = X_{i,k} + \Lambda_{i,k}$, $\Lambda_{i,k} \sim \mathcal{N}(0, \sigma^2)$
- Each source can be either a sender or a receiver.
- Collision channel, collision feedback
- Every source estimates the processes for every other sources.
- Source i: estimate $X_{j,k}$ for source j ($\hat{X}^j_{i,k}$), calculate the AoI of source j ($h^j_{i,k}$).
- Every source decides (i) when to sample, (ii) who to communicate with, and (iii) what to transmit.

Chapters 2 ~ 4

Challenges

- Two main challenges → Multi-agent Reinforcement Learning

  (i) increased dimensions of decision making,

  (ii) network topologies.

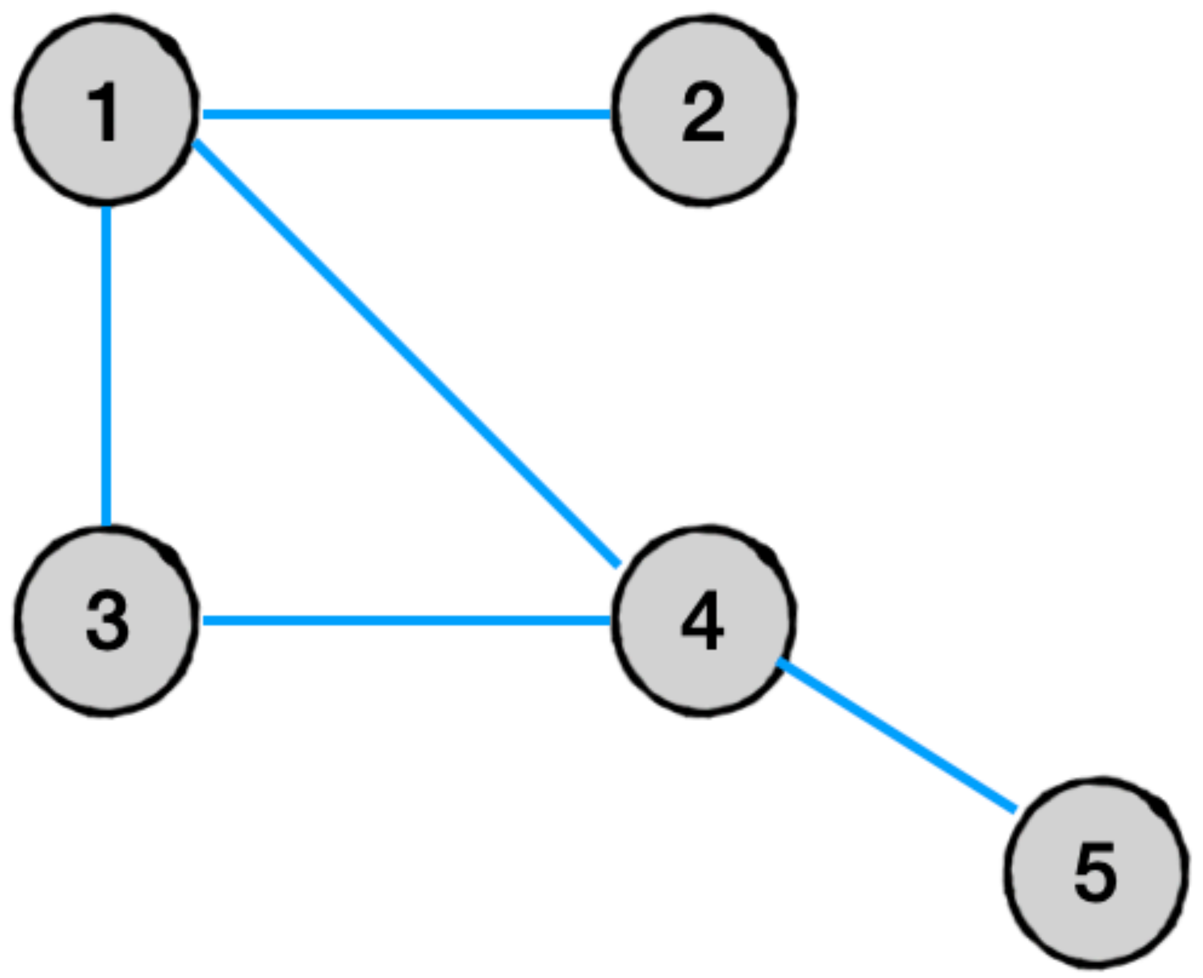

- **Goal**: minimize the average estimation error,

$$L(M) = \min_{\pi} \lim_{K \to \infty} \frac{1}{M^2 K} \mathbb{E}[\sum_{k=1}^{K} \sum_{i=1}^{M} \sum_{j=1}^{M} \left(X_{j,k} - X_{i,k}^{j,\pi}\right)^2]$$

- In oblivious policies,

$$J(M) = \min_{\pi} \lim_{K \to \infty} \frac{1}{KM^2} \mathbb{E}[\sum_{k=1}^{K} \sum_{i=1}^{M} \sum_{j=1}^{M} h_{i,k}^{j,\pi}]$$

Chapters 2 ~ 4

Graphical Reinforcement Learning

- Classical Actor-Critic Framework [1]

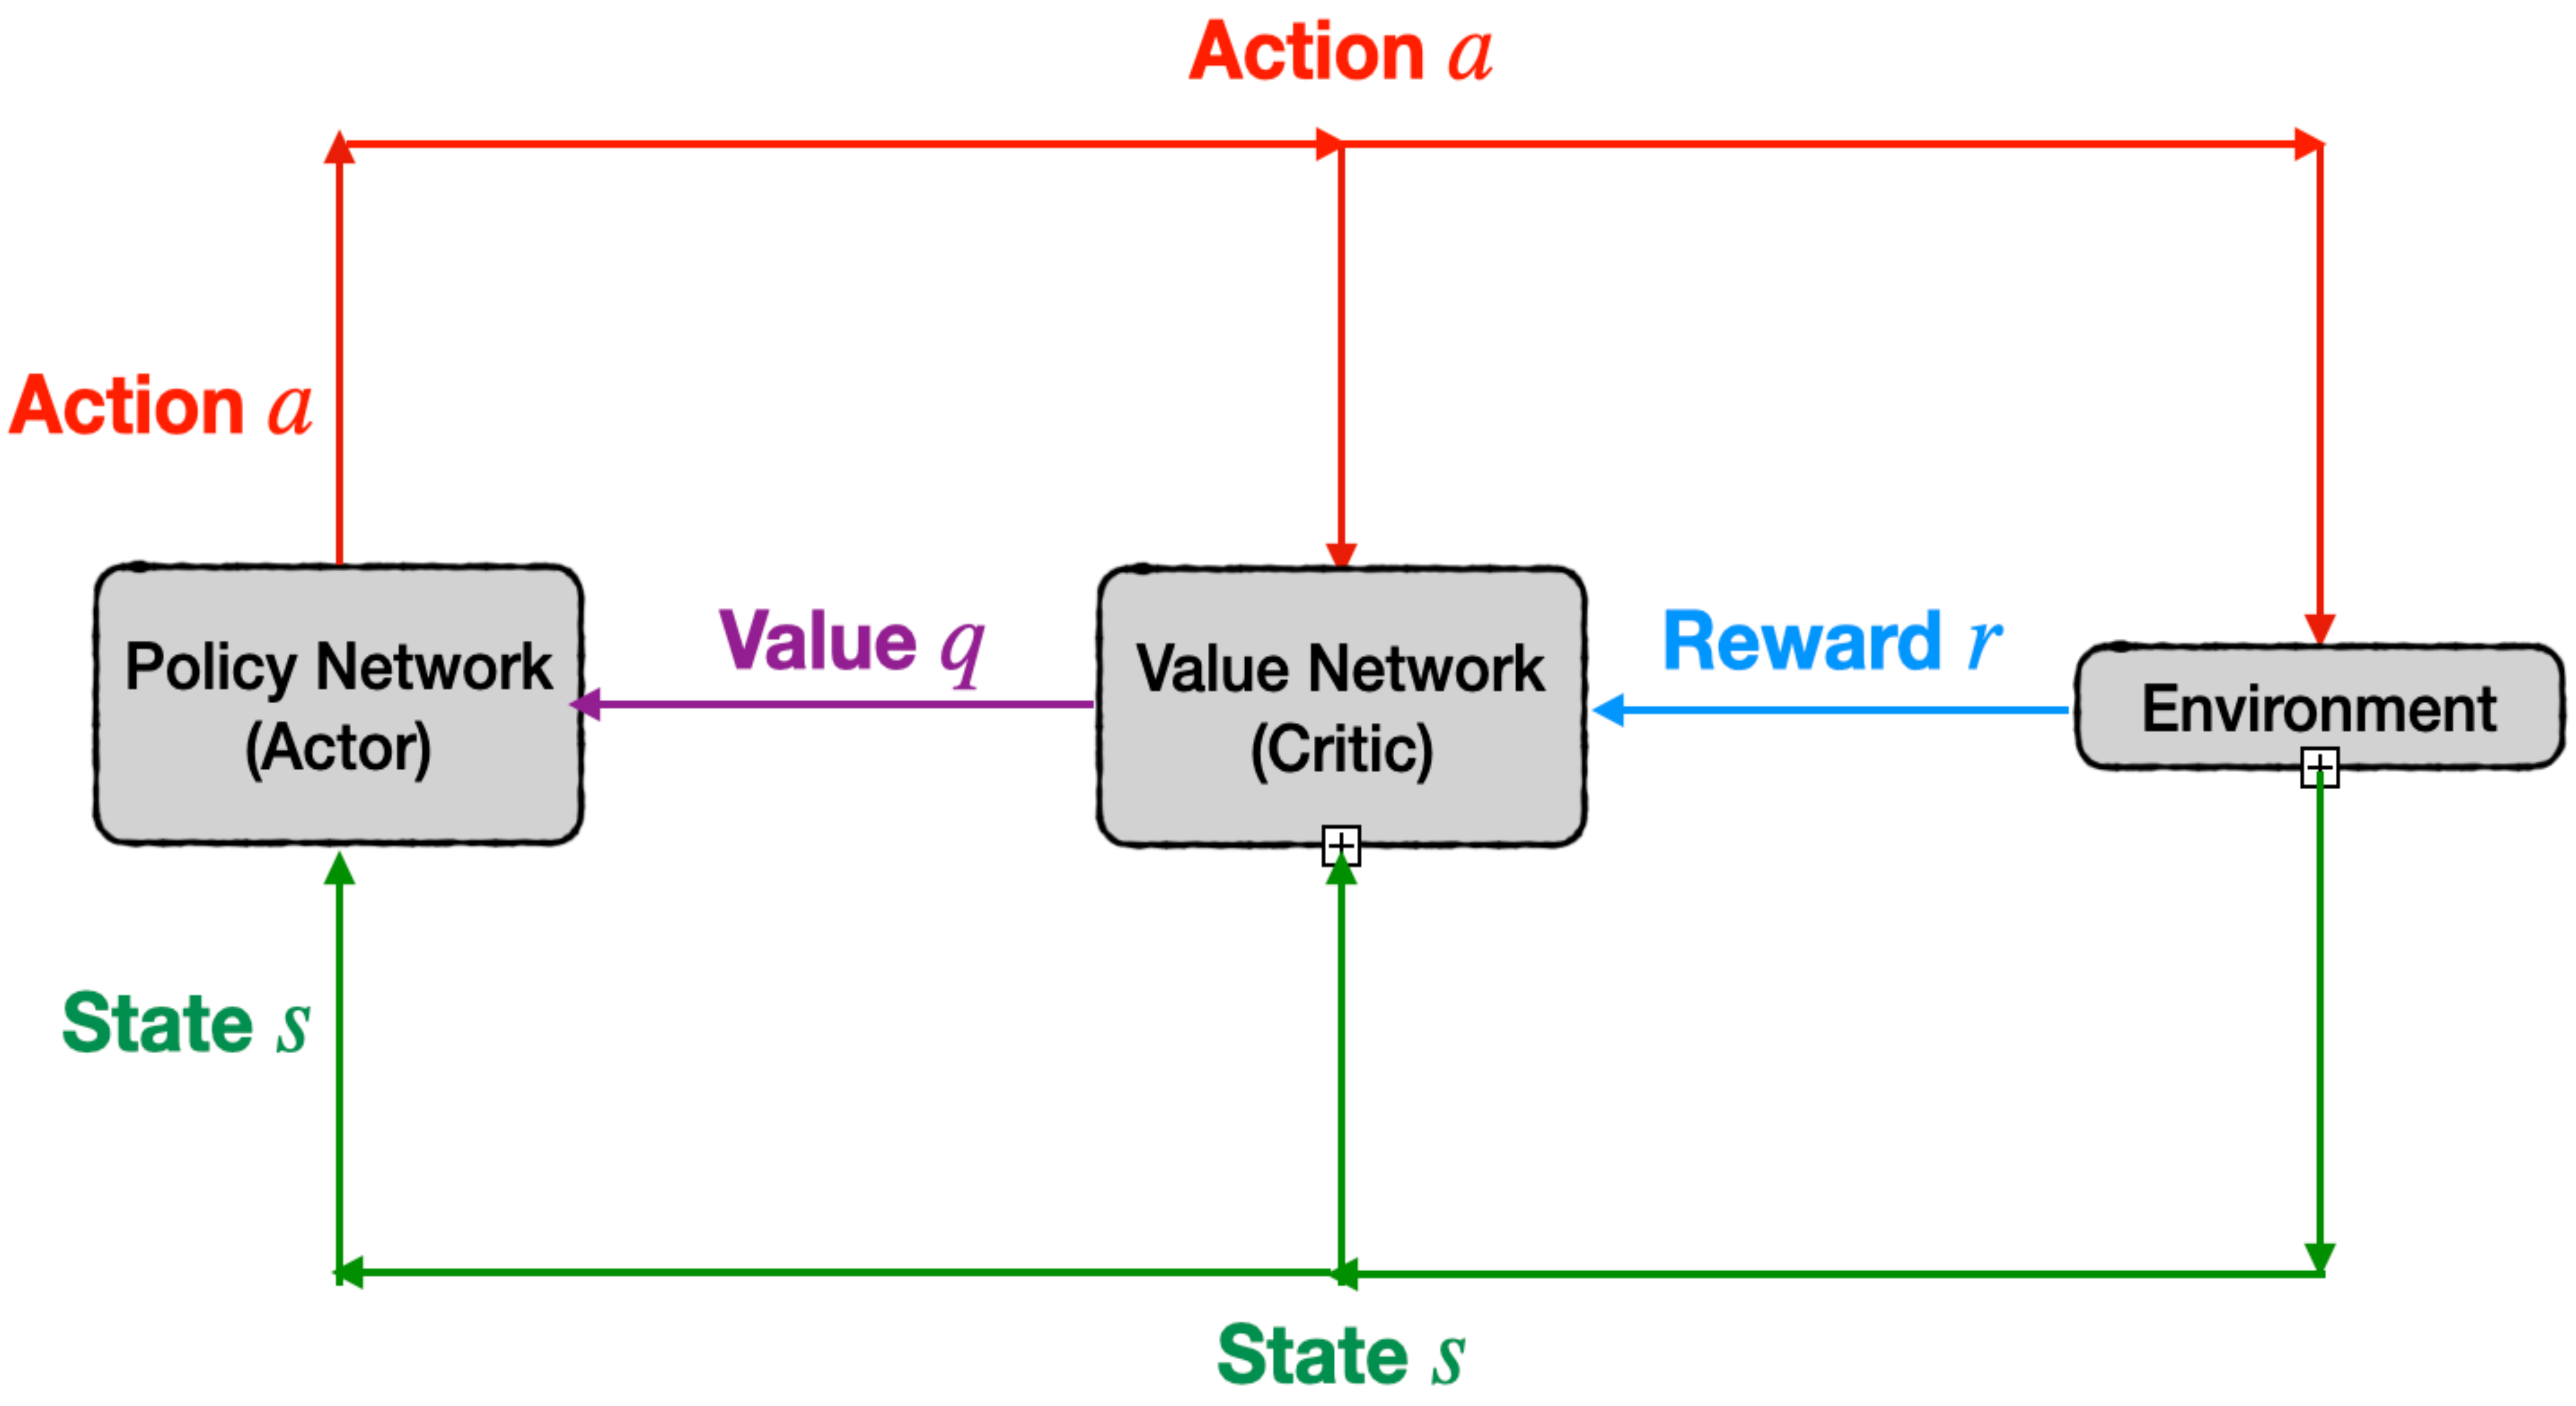

[1] Mnih-Badia-Mirza-Graves-Harley-Lillicrap-Silver-Kavukcuoglu, Asynchronous Methods for Deep Reinforcement Learning, 2016.

Chapters 2 ~ 4

Graphical Reinforcement Learning

- Critic ← Graph Neural Networks
- Actor ← Graph Recurrent Neural Networks
- State ← Observation
- $A \sim F_{\text{Softmax}}(HBH')$: # of parameters in B is independent of # of sources. Transferability

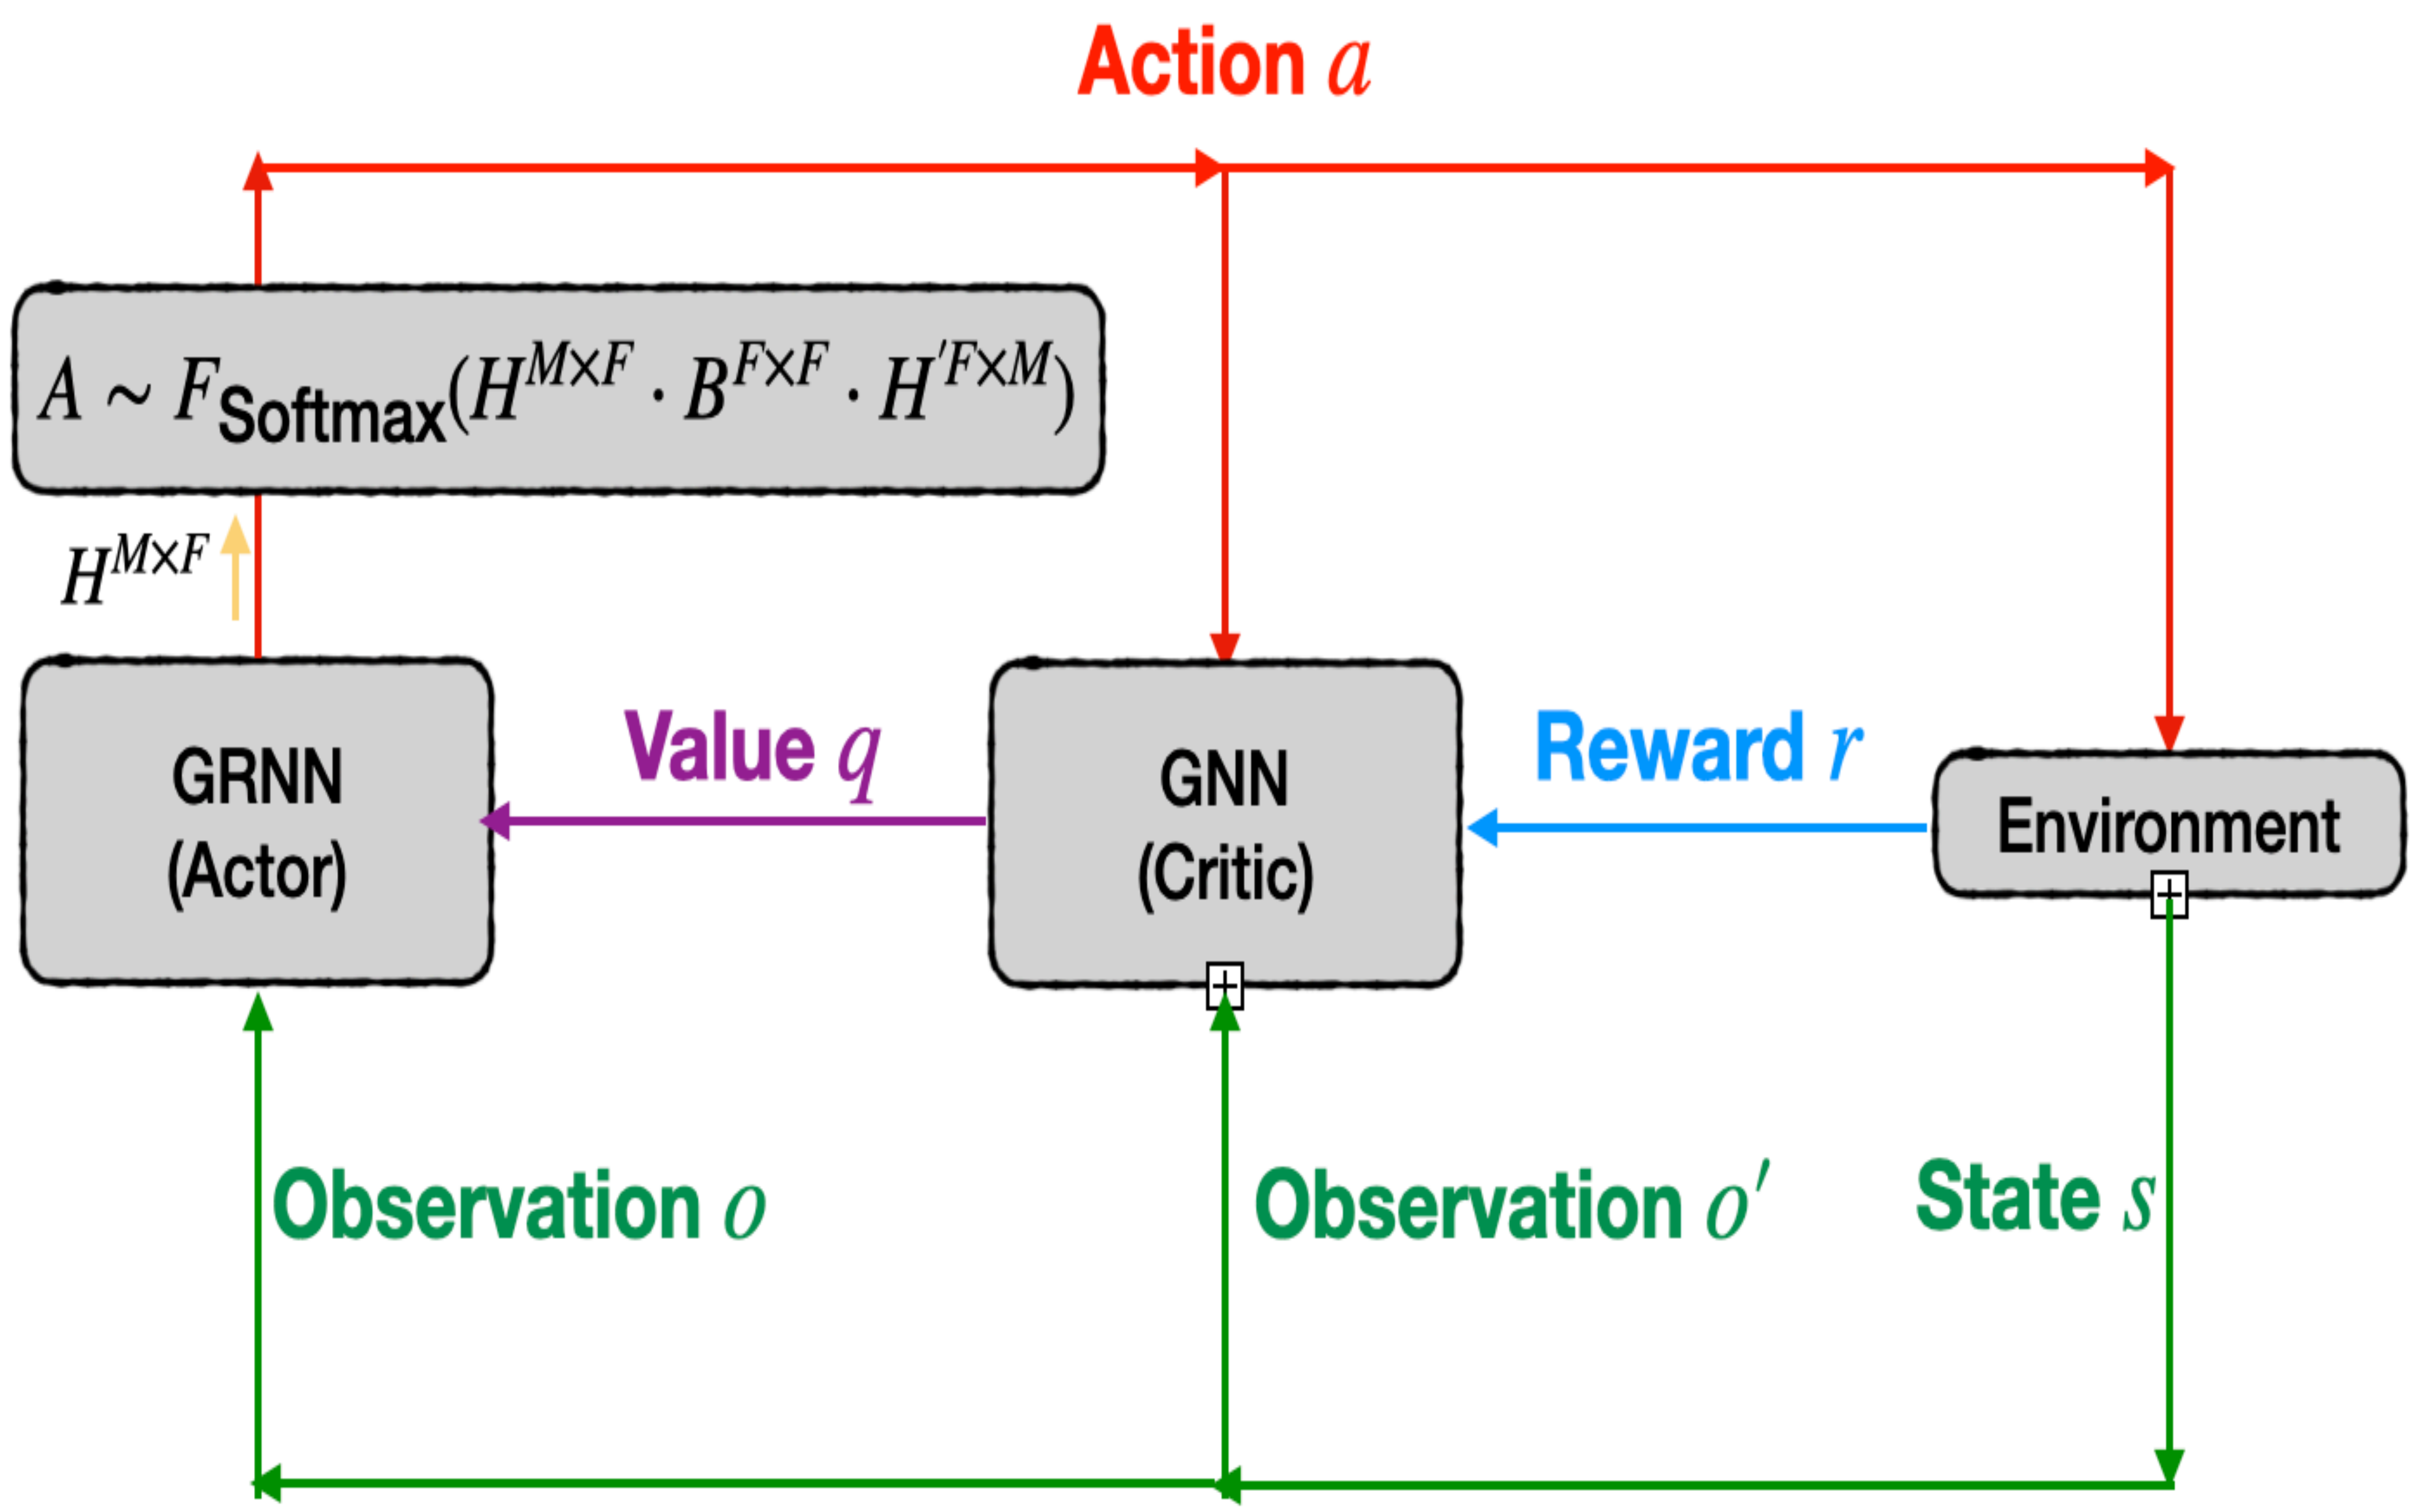

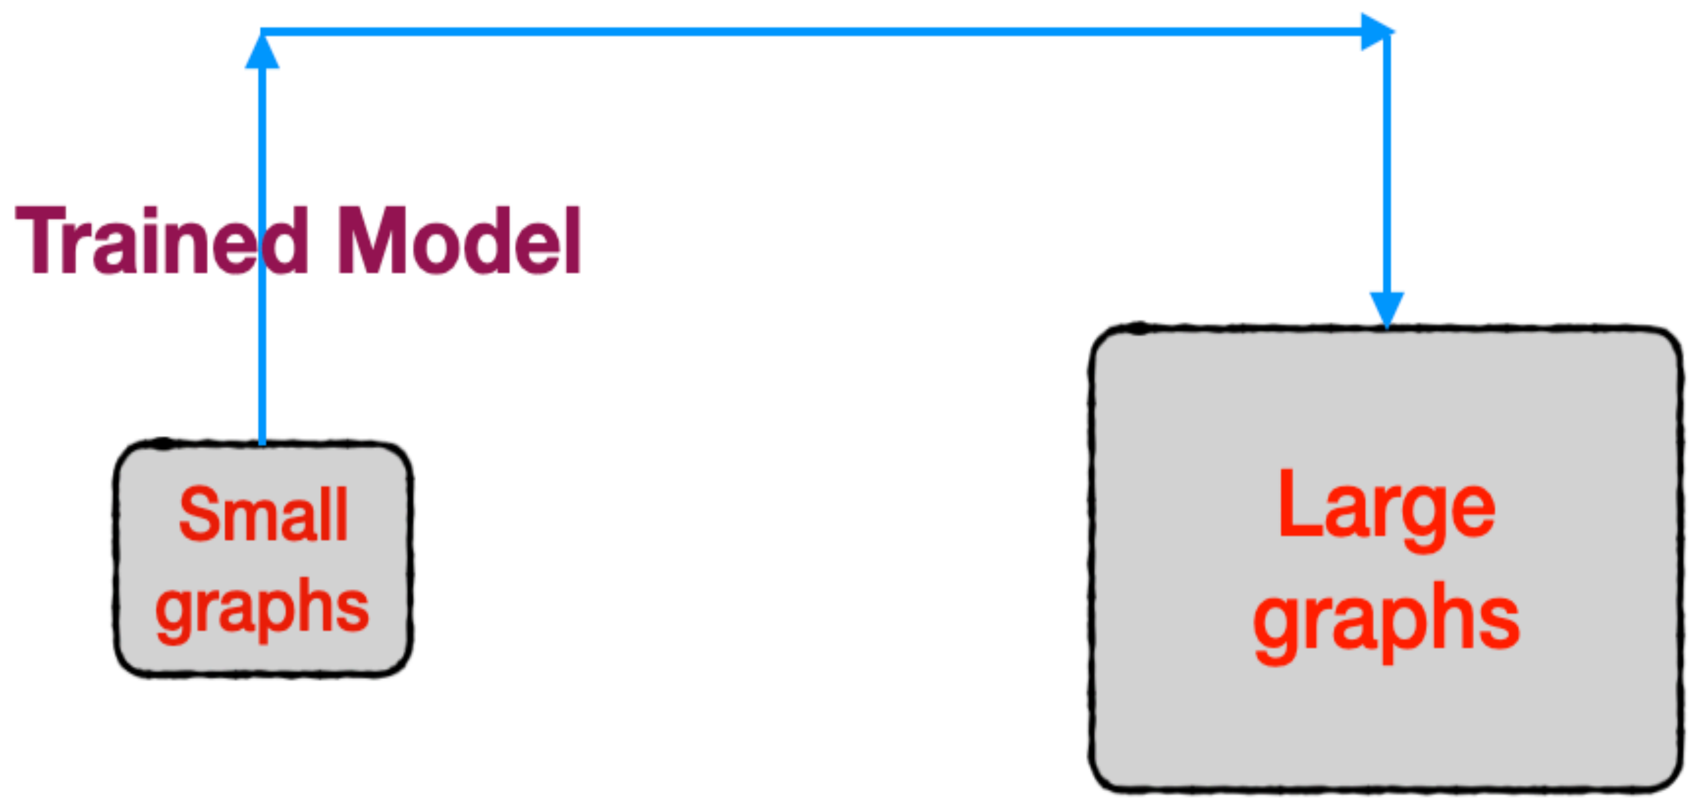

Chapters 2 ~ 4

Transferability

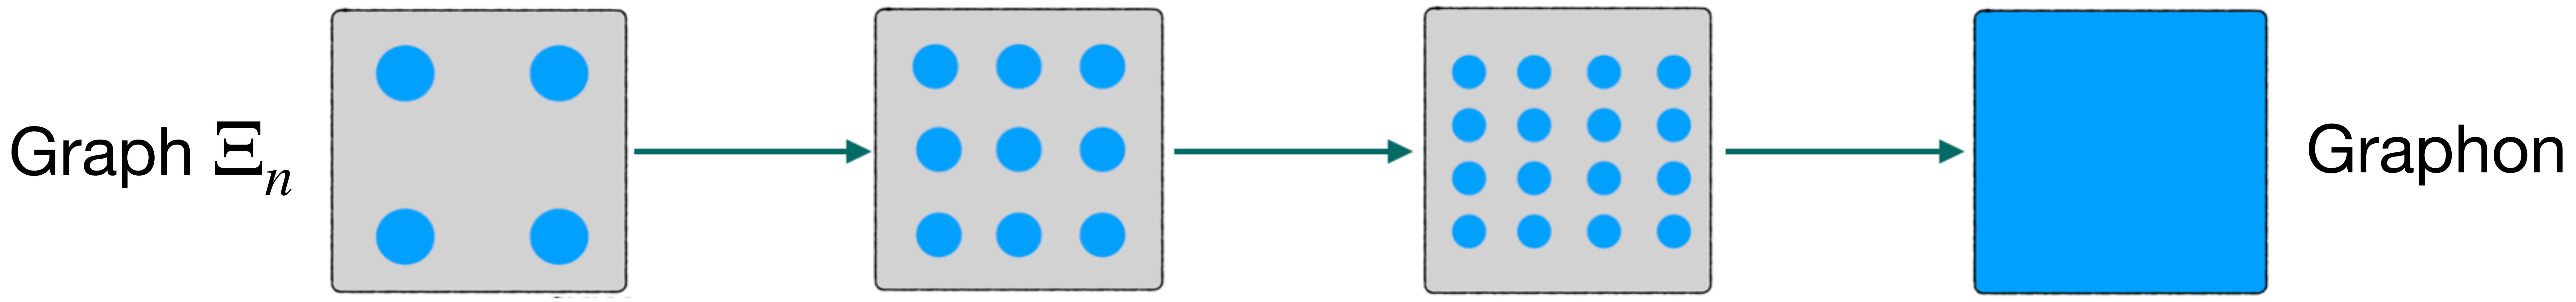

- Graphon [1]: a limit of sequence of convergent graphs. $W : [0,1]^2 \to [0,1]$, bounded, measurable, and symmetric.

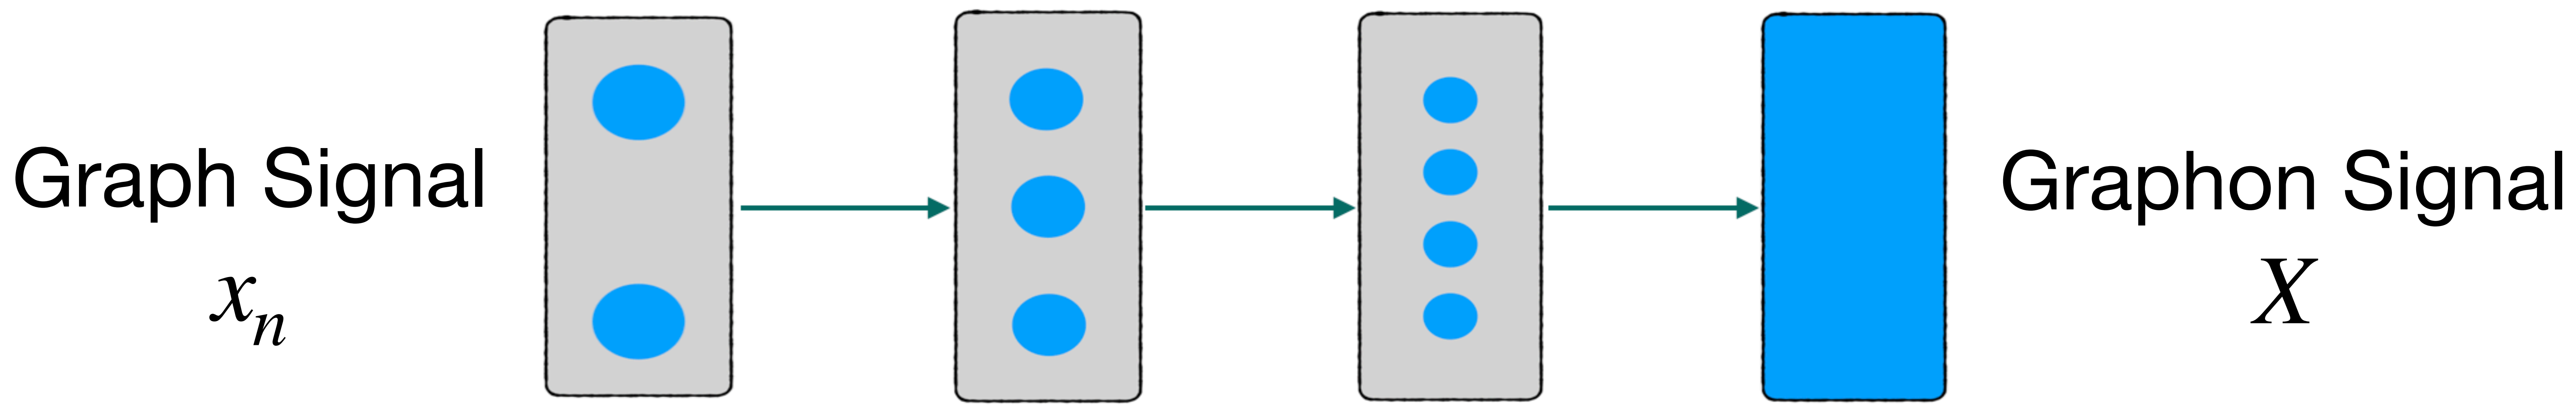

- Graphon signal [1]: $X \in L^2([0,1])$
- Output & operator: $Y(v) = (T_W X)(v) = \int_0^1 W(u,v)X(u)du$

[1] Ruiz-Chamon-Ribeiro, Transferability properties of graph neural networks, 2022.

Chapters 2 ~ 4

Transferability

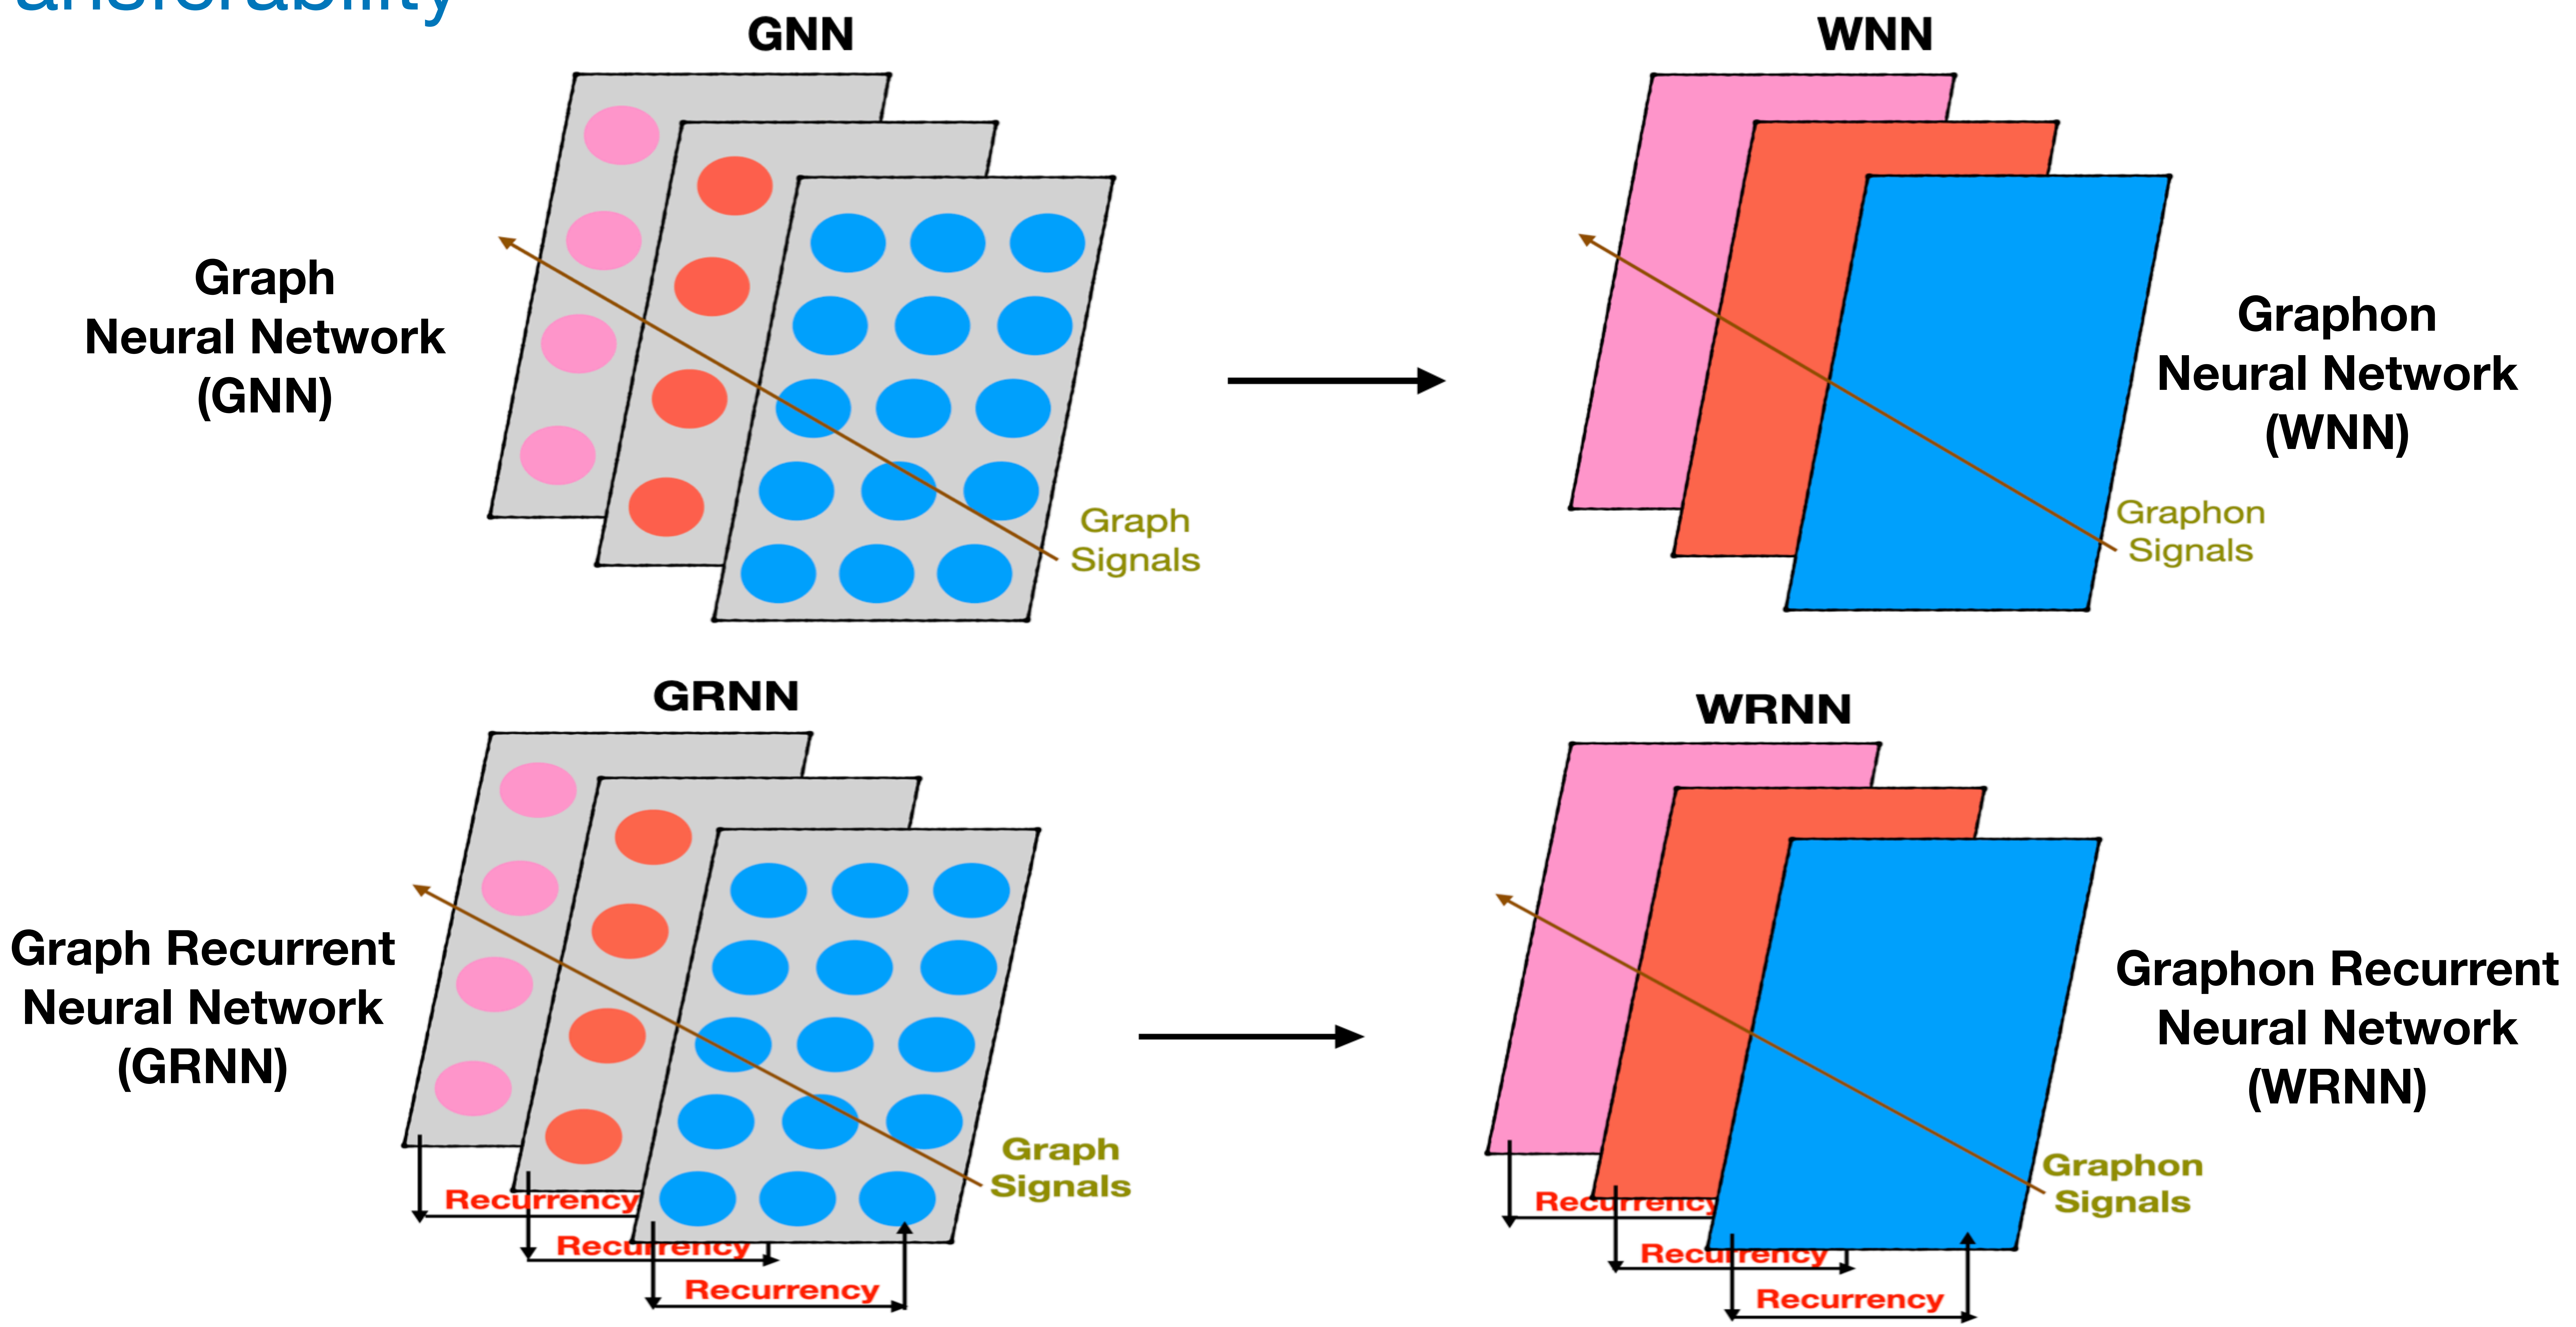

Chapters 2 ~ 4

Transferability in GRNN

- Given (W, X), a (Ξ_n, x_n) with dimension n can be induced by (W, X)
- Given (Ξ_n, x_n), a (W_{Ξ_n}, X_n) is induced from (Ξ_n, x_n).

Assumption 1. The spectral response $h(\lambda)$ of a convolutional filter $T_{H,W}$ is L-Lipschitz in $[-1, -\epsilon] \cup [\epsilon, 1]$, and ℓ-Lipschitz in $(-\epsilon, \epsilon)$ with $\ell < L$. Moreover, $|h(\lambda)| < 1$.

Assumption 2. The activation functions satisfy $|\rho(x) - \rho(y)| \leq |x - y|$, and $\rho(0) = 0$.

Theorem 4. $Y = \Phi(X; W)$ is a WRNN, the convolutional layers of $\Phi(X; W)$ satisfying Assumptions 1 & 2. Let (Ξ_n, x_n), (W_{Ξ_n}, X_n) be defined above, $Y_n = \Phi(X_n; W_{\Xi_n})$. Then, $||Y - Y_n|| \leq \Theta_n ||X|| + c||X - X_n||$, where $\Theta_n \to 0$ as $n \to \infty$.

Remark: $||Y - Y_n|| \to 0$ as $n \to \infty$.

Chapters 2 ~ 4

Transferability in Action Distribution

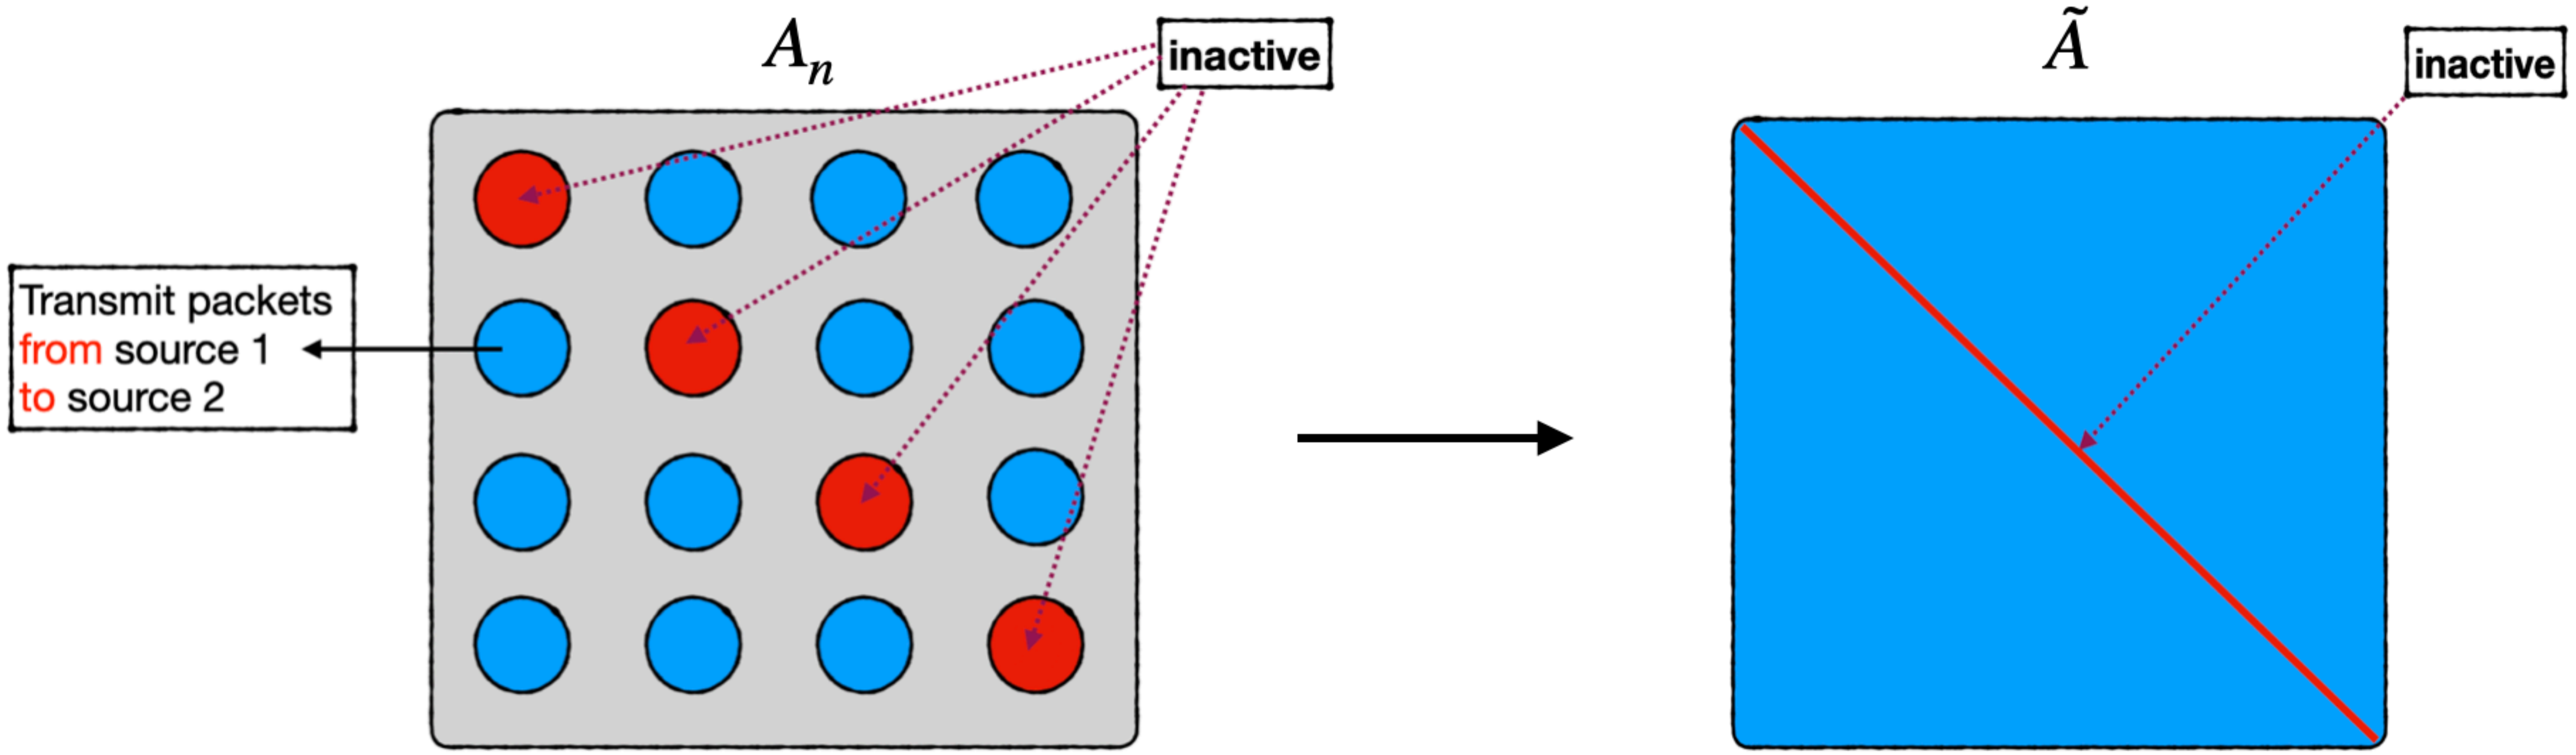

- Given $A_n \rightarrow$ obtain $\tilde{A}_n$. [refer to the Thesis]

Theorem 5. Let $Y = \Phi(X; W)$ is a WRNN satisfying Assumptions 1 & 2, then $\|\tilde{A} - \tilde{A}_n\| \rightarrow 0$ as $n \rightarrow \infty$.

- Transferability $\longleftarrow$ graph filters [1]

[1] Ruiz-Chamon-Ribeiro, Transferability properties of graph neural networks, 2022.

Chapters 2 ~ 4

Simulations

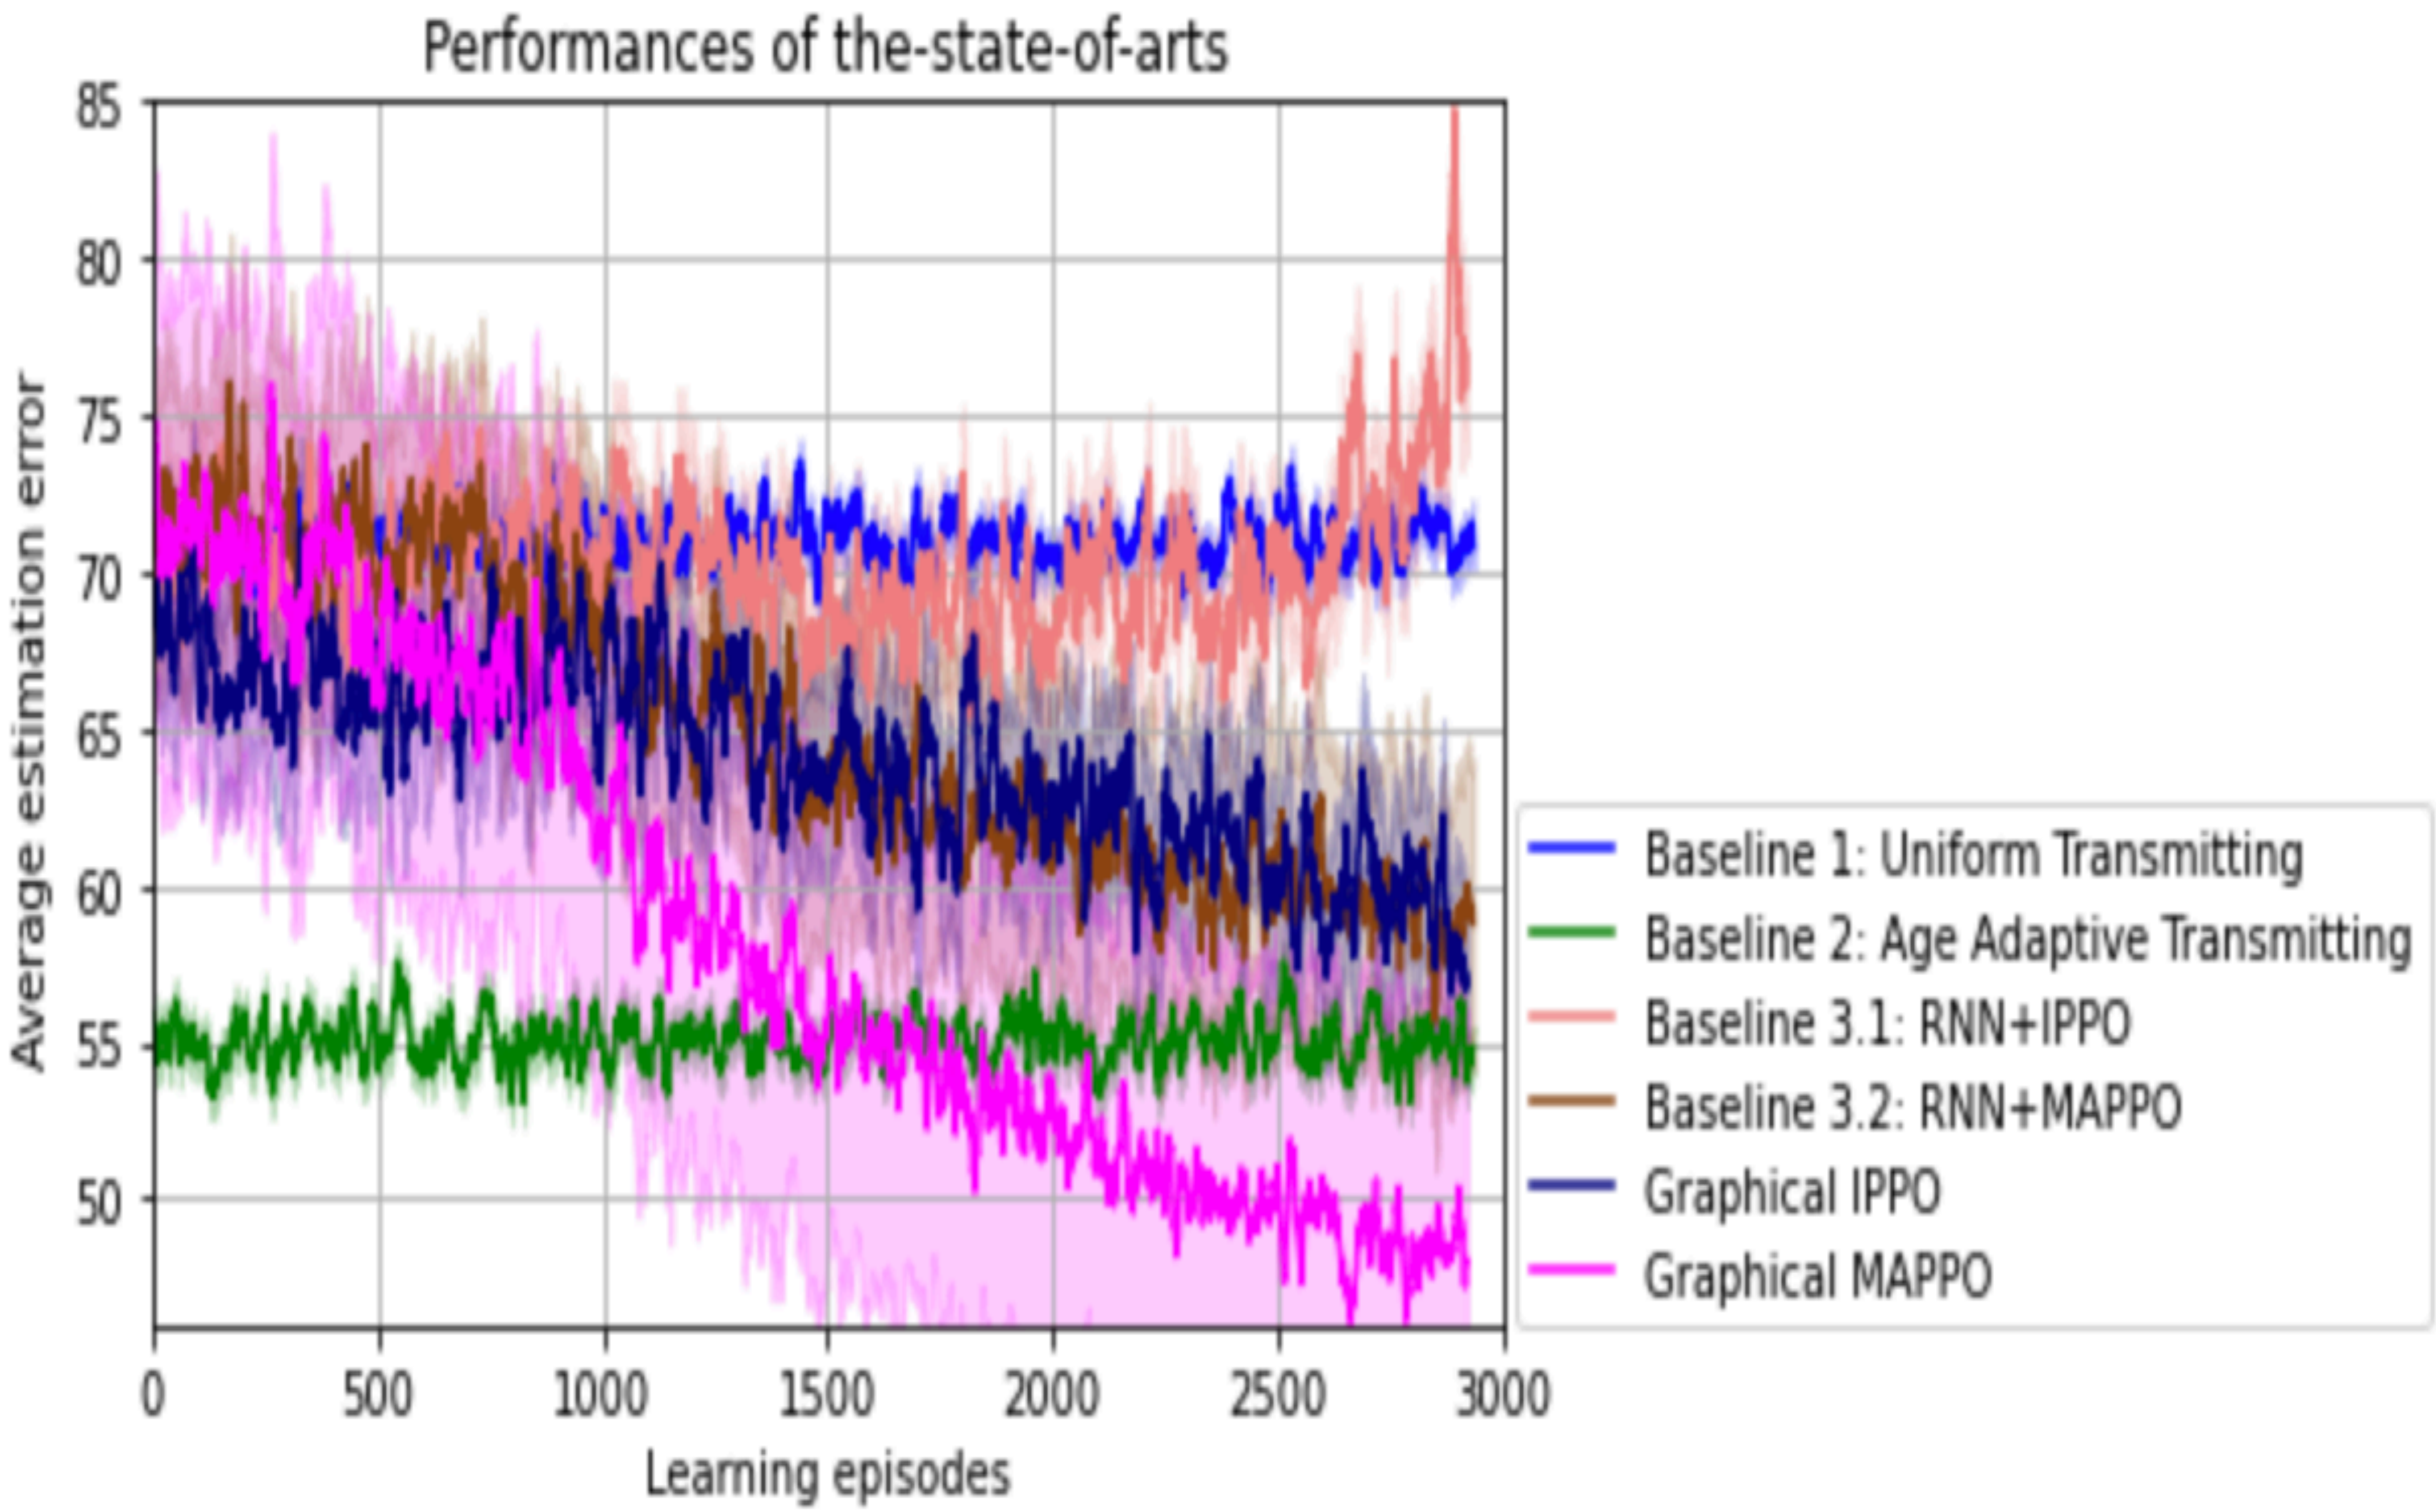

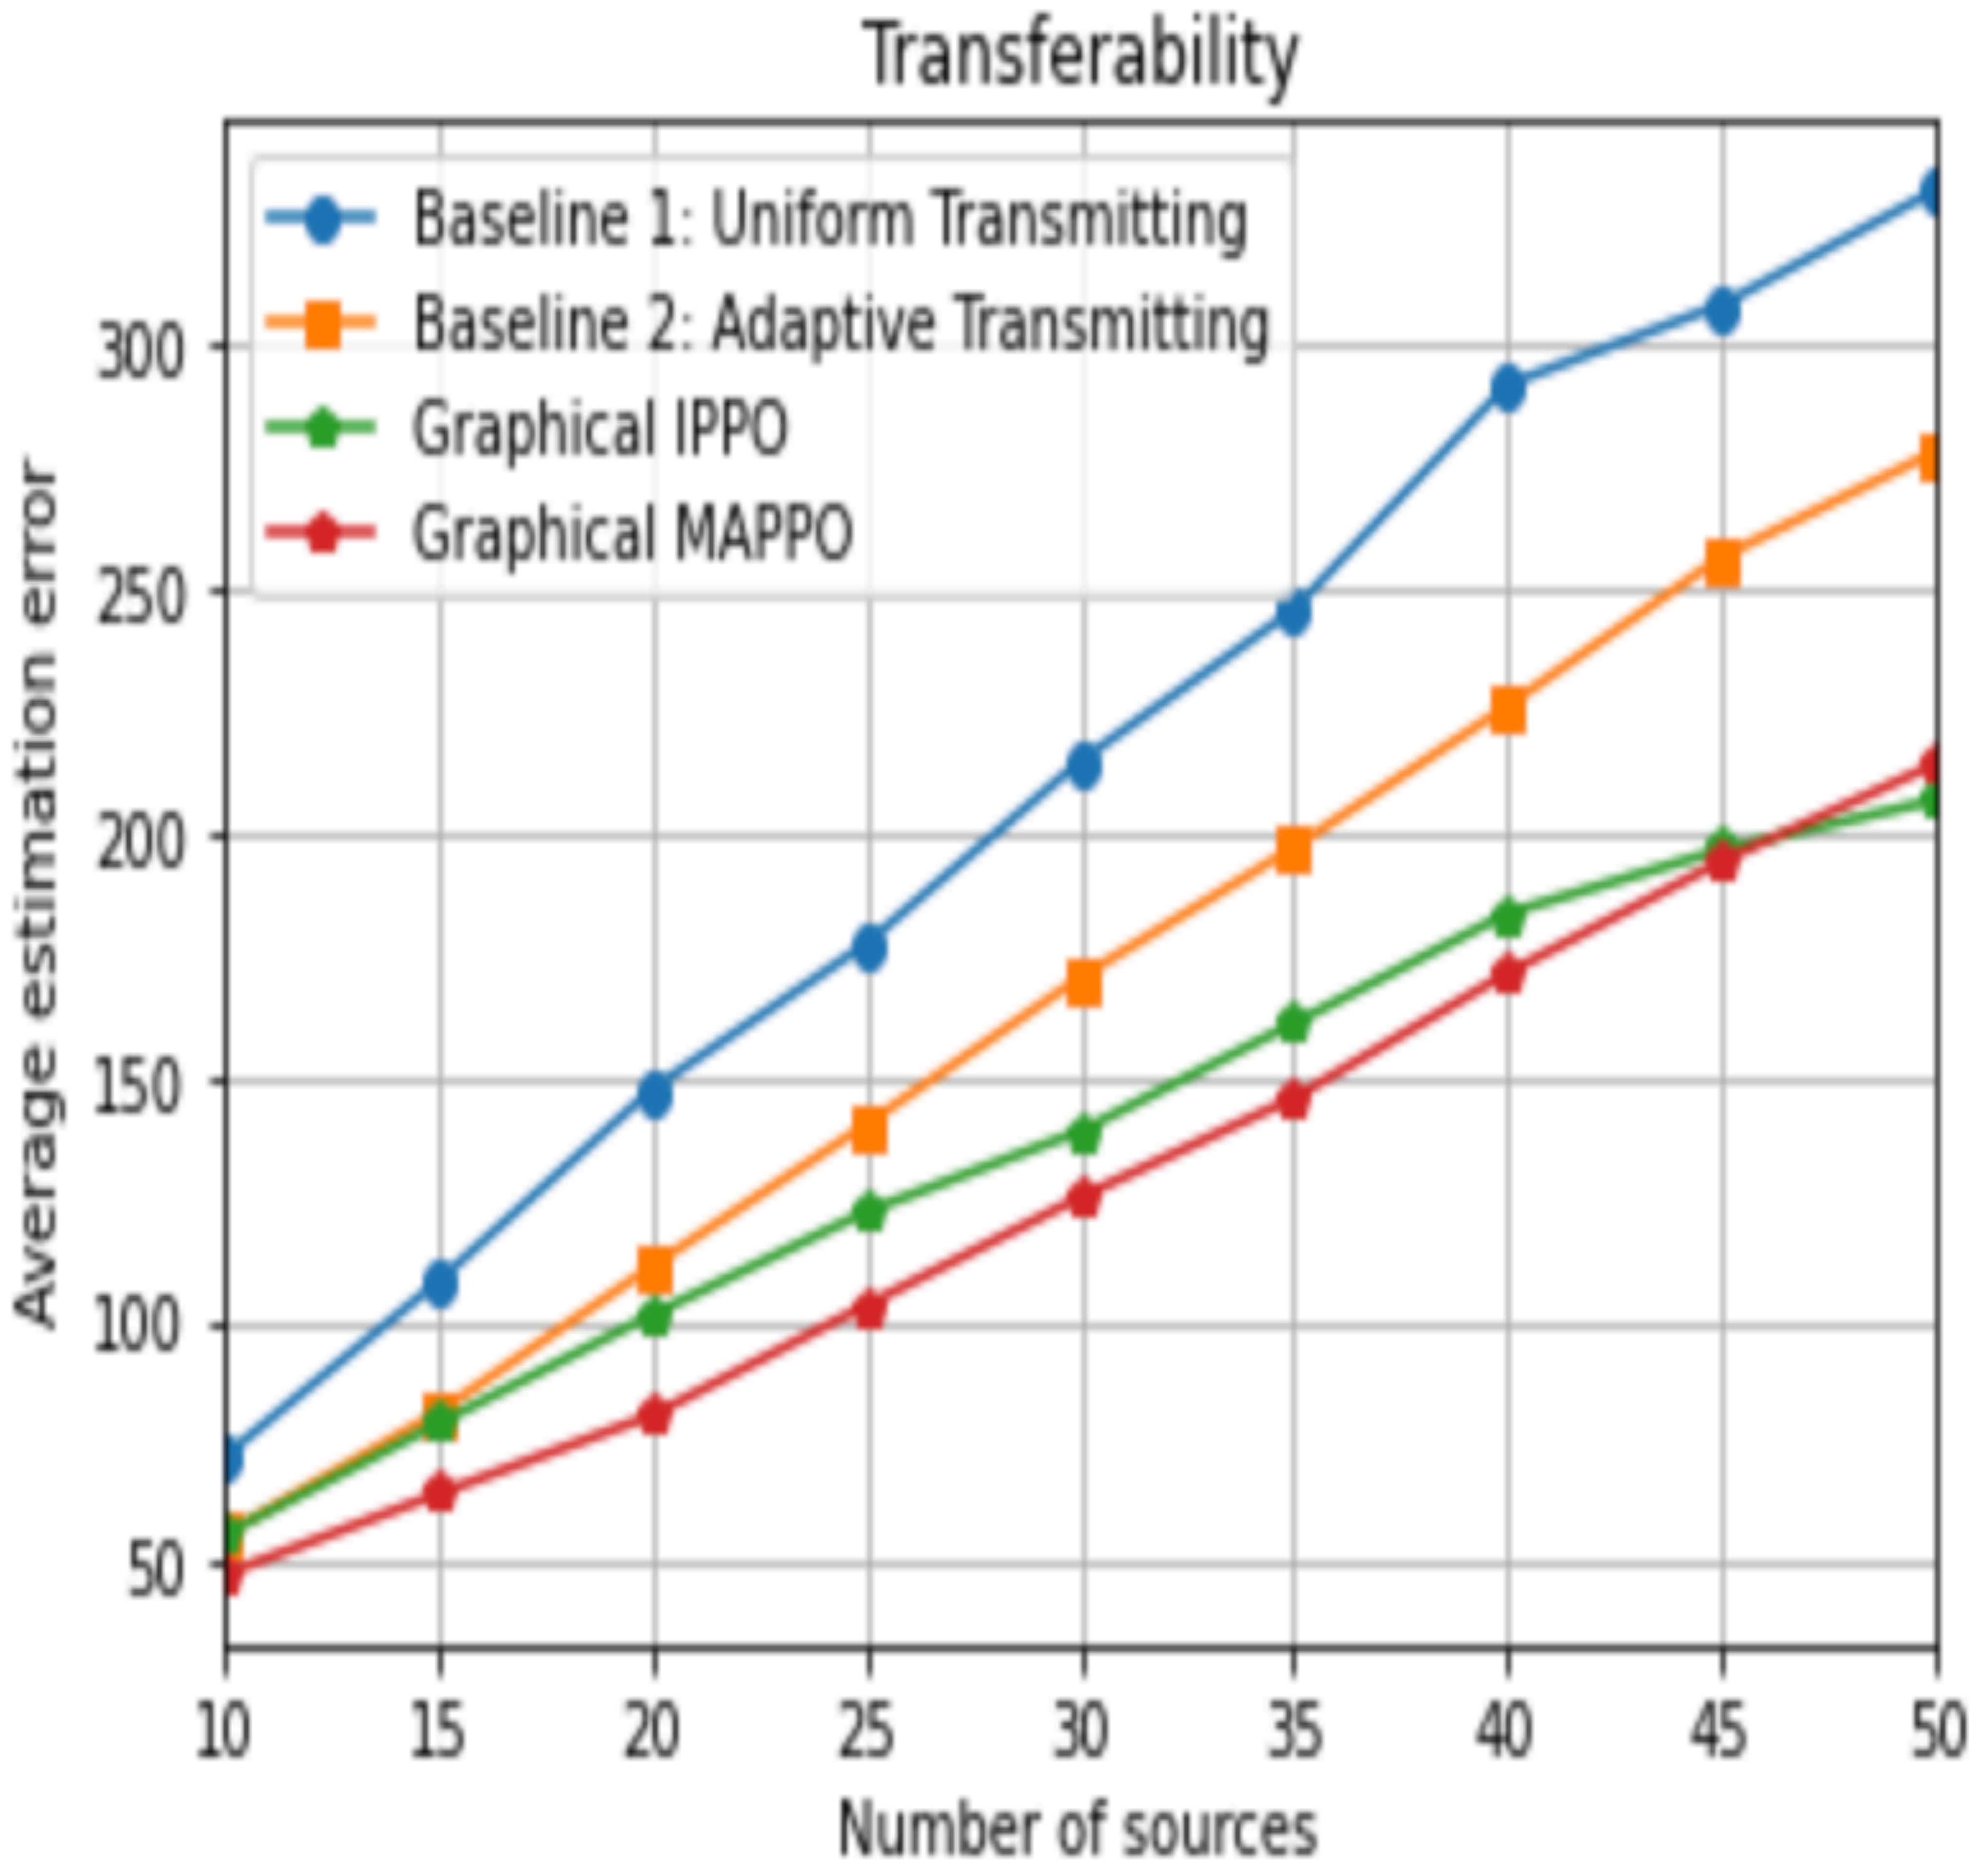

- Achievements: this work is ready to submit.

Chapter 5

- Tradeoffs between AoI and communication rate in broadcast networks

- Result 1: Coding is beneficial to the AoI, and the benefits increases with # of users

- Result 2: Tradeoffs exist → the system has to sacrifice AoI to achieve a higher rate

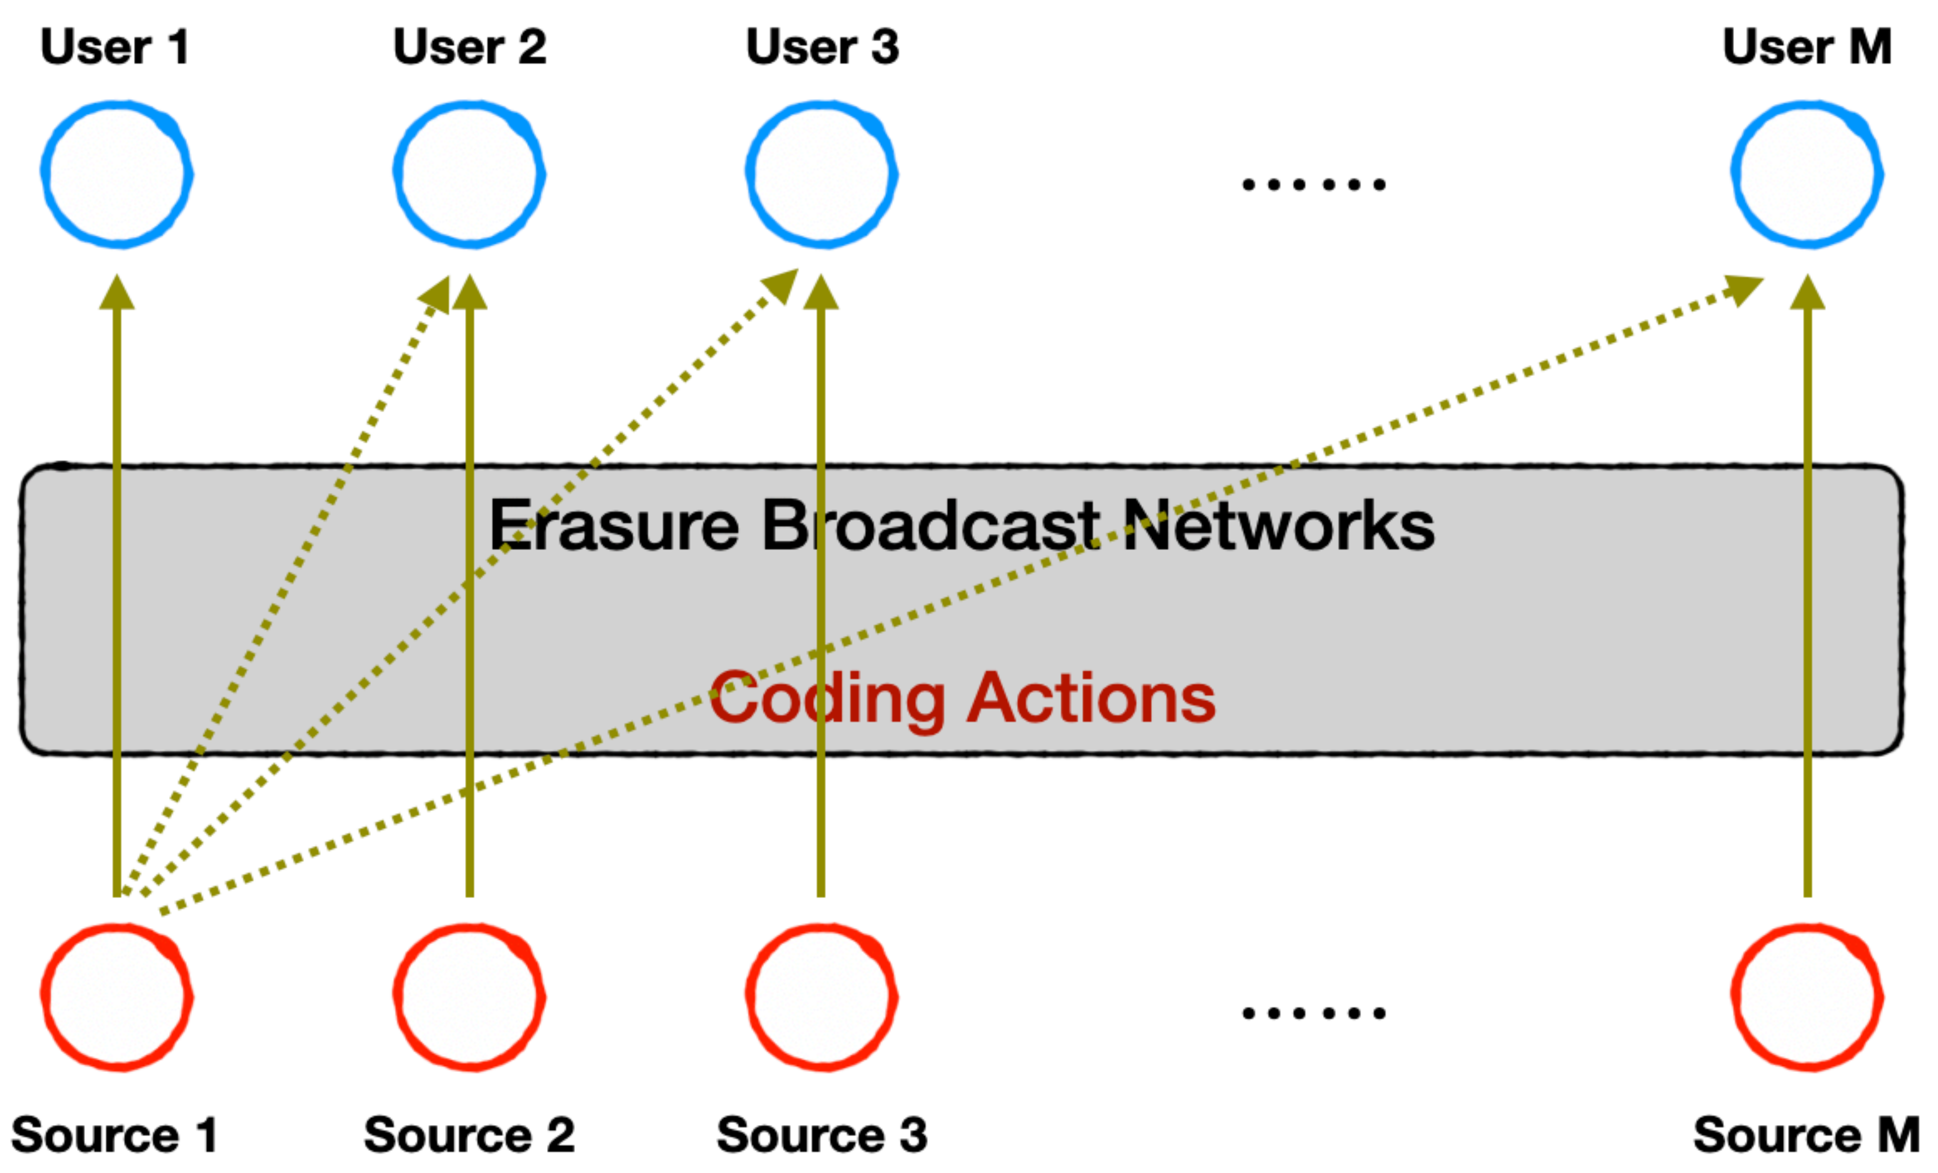

- Achievements:

- [C3] X. Chen and S. Saeedi-Bidokhti, Benefits of Coding on Age of Information in Broadcast Networks, IEEE ITW, 2019.

- [C4] X. Chen, R. Liu, S. Wang, and S. Saeedi-Bidokhti, Timely Broadcasting in Erasure Networks: Age-Rate Tradeoffs, IEEE ISIT, 2021.

- [J3] X. Chen and S. Saeedi-Bidokhti, Timely Broadcasting Mechanisms in Erasure Networks: Age-Rate Tradeoffs, IEEE TWC, submitted.

Chapter 6

Backgrounds

- Timely inference and detection for processes that evolve both temporally and spatially. COVID-19 in contact networks
- How to contain the spread as soon as possible? Sequentially policy.
- Testing has a dual role: (i) detect infected nodes, and (ii) learn the spread.
- Paradigm I: contact tracing [1, 2] —— pure exploitation
- Paradigm II: random testing —— pure exploration
- Silent spread: an undetected individual may infect its neighbors
- Tradeoffs between exploitation of knowledge and exploration of the unknown.

[1] Kojaku,-Dufresne-Mones-et al, The effectiveness of backward contact tracing in networks, 2021.
[2] Ou-Sinha-Suen-et al. Who and when to screen: Multi-round active screening for network recurrent infectious diseases under uncertainty, 2020.

Chapter 6

Motivations & Literature Review

- Estimation & Prediction: SIR and variants [1,2,3] → Our work: testing and isolation policy.
- Differential Equation Approximations: no heterogeneity [4,5] → Our work: heterogeneity & spread
- Comparing to other RL:
  - Multi-armed Bandit [6,7] → Our work: time-variant actions
  - Active Search [8,9] → Our work: dynamic target
  - POMDP [10] → Our work: general setting
- Novel Exploitation-Exploration Tradeoffs

[1] Bastani-Drakopoulos-Gupta-et al., Efficient and targeted COVID-19 border testing via reinforcement learning, 2021.
[2] Ramos-Ferrandez-Perez-et al., A simple but complex enough θ-SIR type model to be used with COVID-19 real data: application to the case Italy, 2021.
[3] Hu-Geng, Heterogeneity learning for SIRS model: an application to the COVID-19, 2021.
[4] Tanaka-Kuga-Tanimoto, Pair approximation model for the vaccination game: predicting the dynamic process of epidemic spread and individual actions against contagion, 2021.
[5] Kabir-Tanimoto, Evolutionary vaccination game approach in metapopulation migration model with information spreading on different graphs, 2019.
[6] Auer-Bianchi-Fischer, Finite-time Analysis of the Multiarmed Bandit Problem, 2002.
[7] Agrawal-Goyal, Regret analysis of stochastic and nonstochastic multi-armed bandit problems, 2012.
[8] Bilgic-Mihalkova-Getoor, Active learning for networked data, 2010
[9] Wang-Garnett-Schneider, Active search on graphs, 2013.
[10] Singh-Liu-Shroff, A Partially Observable MDP Approach for Sequential Testing for Infectious Diseases such as COVID-19, 2020.

Chapter 6

Problem Formulation

- Susceptible (S), Latent (L), Infectious (I), and Recovered(R)
- β: transmission rate, λ: ill-being rate, γ: recovery rate

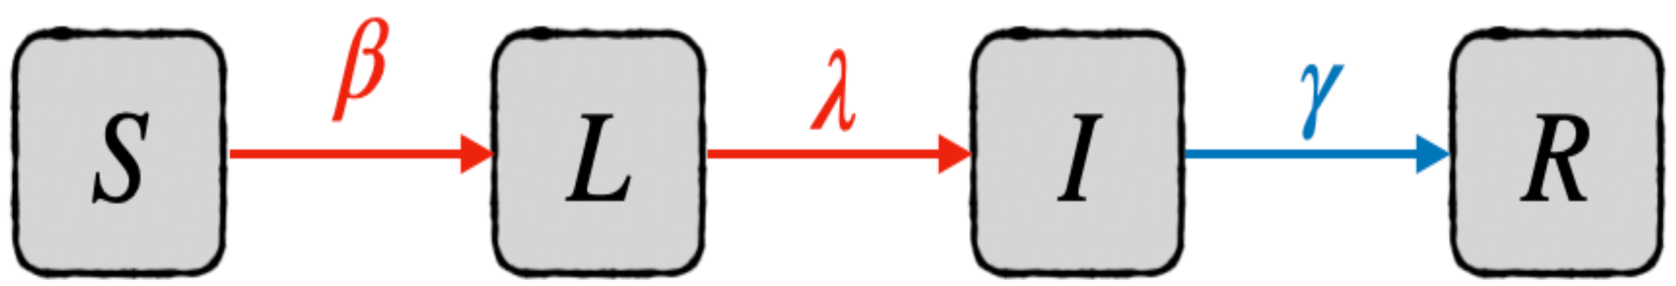

- An individual tested positive will be isolated immediately.
- A recovered individual can not be infected again.
- Only infectious individuals can infect others.
- $B(t)$: testing budget; $\mathscr{K}^{\pi}(t)$: the set of tests; $C^{\pi}(T)$: the cumulative infections
- **Goal**: minimize the cumulative infections under budget constraints

$$\min_{\pi:|\mathscr{K}^{\pi}(t)|\leq B(t),0\leq t\leq T-1} \mathbb{E}[C^{\pi}(T)]$$

Chapter 6

Supermodularity

- $S(\mathscr{D}; t)$: the expected number of newly infectious nodes incurred by nodes in $\mathscr{D}$ on day t.

Lemma 2. $\mathbb{E}[C^\pi(t+1) - C^\pi(t)] = S(\mathscr{V}(t) \backslash \mathscr{K}^\pi(t); t)$.

Remark $\min_{\pi : |\mathscr{K}^\pi(t)| \leq B(t), 0 \leq t \leq T-1} \mathbb{E}[C^\pi(T)] \rightarrow \min_{|\mathscr{K}^\pi(t)| \leq B(t)} S(\mathscr{V}(t) \backslash \mathscr{K}^\pi(t); t)$

Theorem 6. $S(\mathscr{K}^\pi(t); t)$ is a supermodular [1] and increasing monotone function on $\mathscr{K}^\pi(t)$.

- By Algorithm A in [1], $S(\mathscr{V}(t) \backslash \tilde{\mathscr{K}}(t); t) \leq (1 + \epsilon_t)$OPT,

$$\epsilon_t = \max_{a \in \mathscr{V}(t)} \frac{S(\mathscr{V}(t); t) - S(\mathscr{V}(t) \backslash \{a\}; t) - S(a(t); t)}{S(\mathscr{V}(t); t) - S(\mathscr{V}(t) \backslash \{a\}; t)}$$

[1] Ilev, An approximation guarantee of the greedy descent algorithm for minimizing a supermodular set function, 2001.

Chapter 6

Exploitation and Exploration

- Reward: $r_i(t) = S(\{i\}; t) \rightarrow$ the expected number of newly infectious nodes incurred by node i on day t.

Lemma 3. $S(\mathscr{V}(t) \backslash \mathscr{K}^\pi(t); t) \leq S(\mathscr{V}(t); t) - \sum_{i \in \mathscr{K}^\pi(t)} r_i(t)$.

Remark $\min_{|\mathscr{K}^\pi(t)| \leq B(t)} S(\mathscr{V}(t) \backslash \mathscr{K}^\pi(t); t) \rightarrow \max_{|\mathscr{K}^\pi(t)| \leq B(t)} \sum_{i \in \mathscr{K}^\pi(t)} r_i(t)$

- Exploitation: Re-arrange $\{r_i(t)\}_i$ in descending order, and test the first $B(t)$ nodes.
- Exploration: Test node i with probability $\min\{1, B(t) r_i(t) / \sum r_i(t)\}$.
- Question: How to estimate $\{r_i(t)\}_i$?

Chapter 6

Message-Passing Framework

- $\underline{u}_i(t)$: the prior probability vector of the true probability vector of node i
- $\underline{w}_i(t)$: the posterior probability vector of the true probability vector of node i

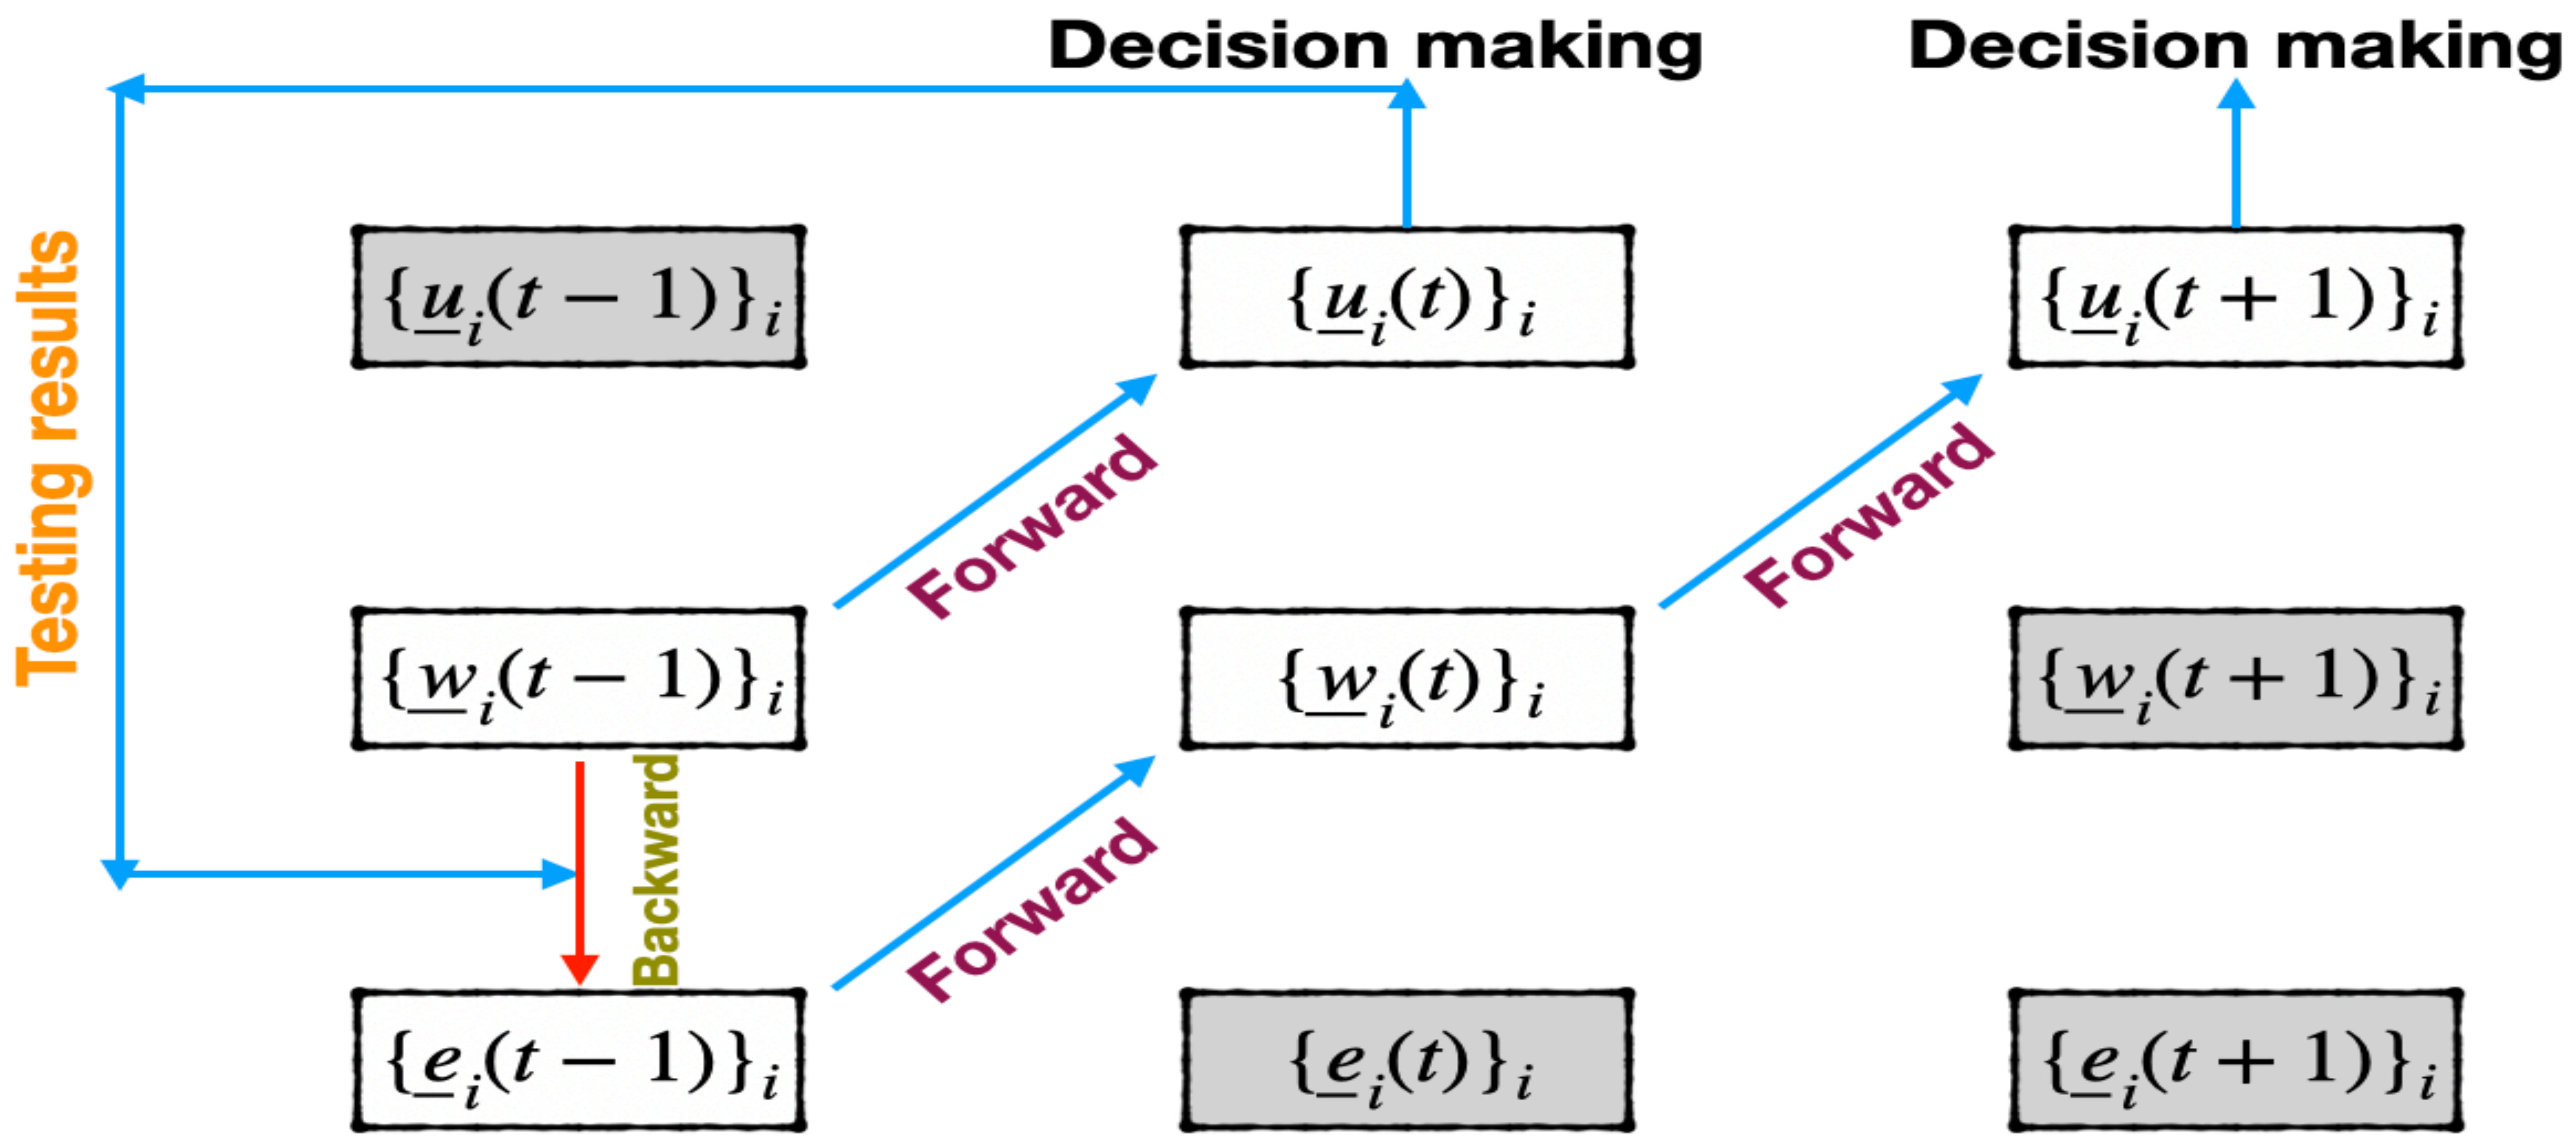

Chapter 6

Backward Updating Is Necessary

- Example 1: A line network with N nodes. Set $\beta = 1, \lambda = 0, \gamma = 0.$ $\rightarrow$ No L and R states. $B(t) = 1$. On the initial day, each node is infected with probability $1/N$. No isolation policy.

Theorem 7. Without the backward updating, for any testing policy that sequentially computes $\{\underline{u}_i(t)\}_i$, with probability $1/e$, $\sum_i ||\underline{v}_i(t) - \underline{u}_i(t)|| \rightarrow \Theta(N)$ as $t \rightarrow \infty$ for large N. With the backward updating, there exists a policy, such that $\sum_i ||\underline{v}_i(t) - \underline{u}_i(t)|| = 0$ for $t \geq 2N$.

Chapter 6

Exploration Is Necessary

- Example 2: A line network with N nodes. Set $\beta = 1, \lambda = 0, \gamma = 0.$ $\rightarrow$ No L and R states. $B(t) = 10$. Node 1 is infected. A slightly wrong initial estimate.

- A specific exploration: 1 (out of 10) test is done randomly, and the other 9 tests are done following exploitation.

Theorem 8. With probability $p_0 \geq 99/100$, on day T, $\dfrac{C^{\text{Exploitation}}}{C^{\text{Exploration}}} \geq c(N, p_0)$, where $c(N, p_0)$ is a constant depending on N and p_0, and $c(N, p_0)$ can be arbitrarily large.

Chapter 6

Simulations

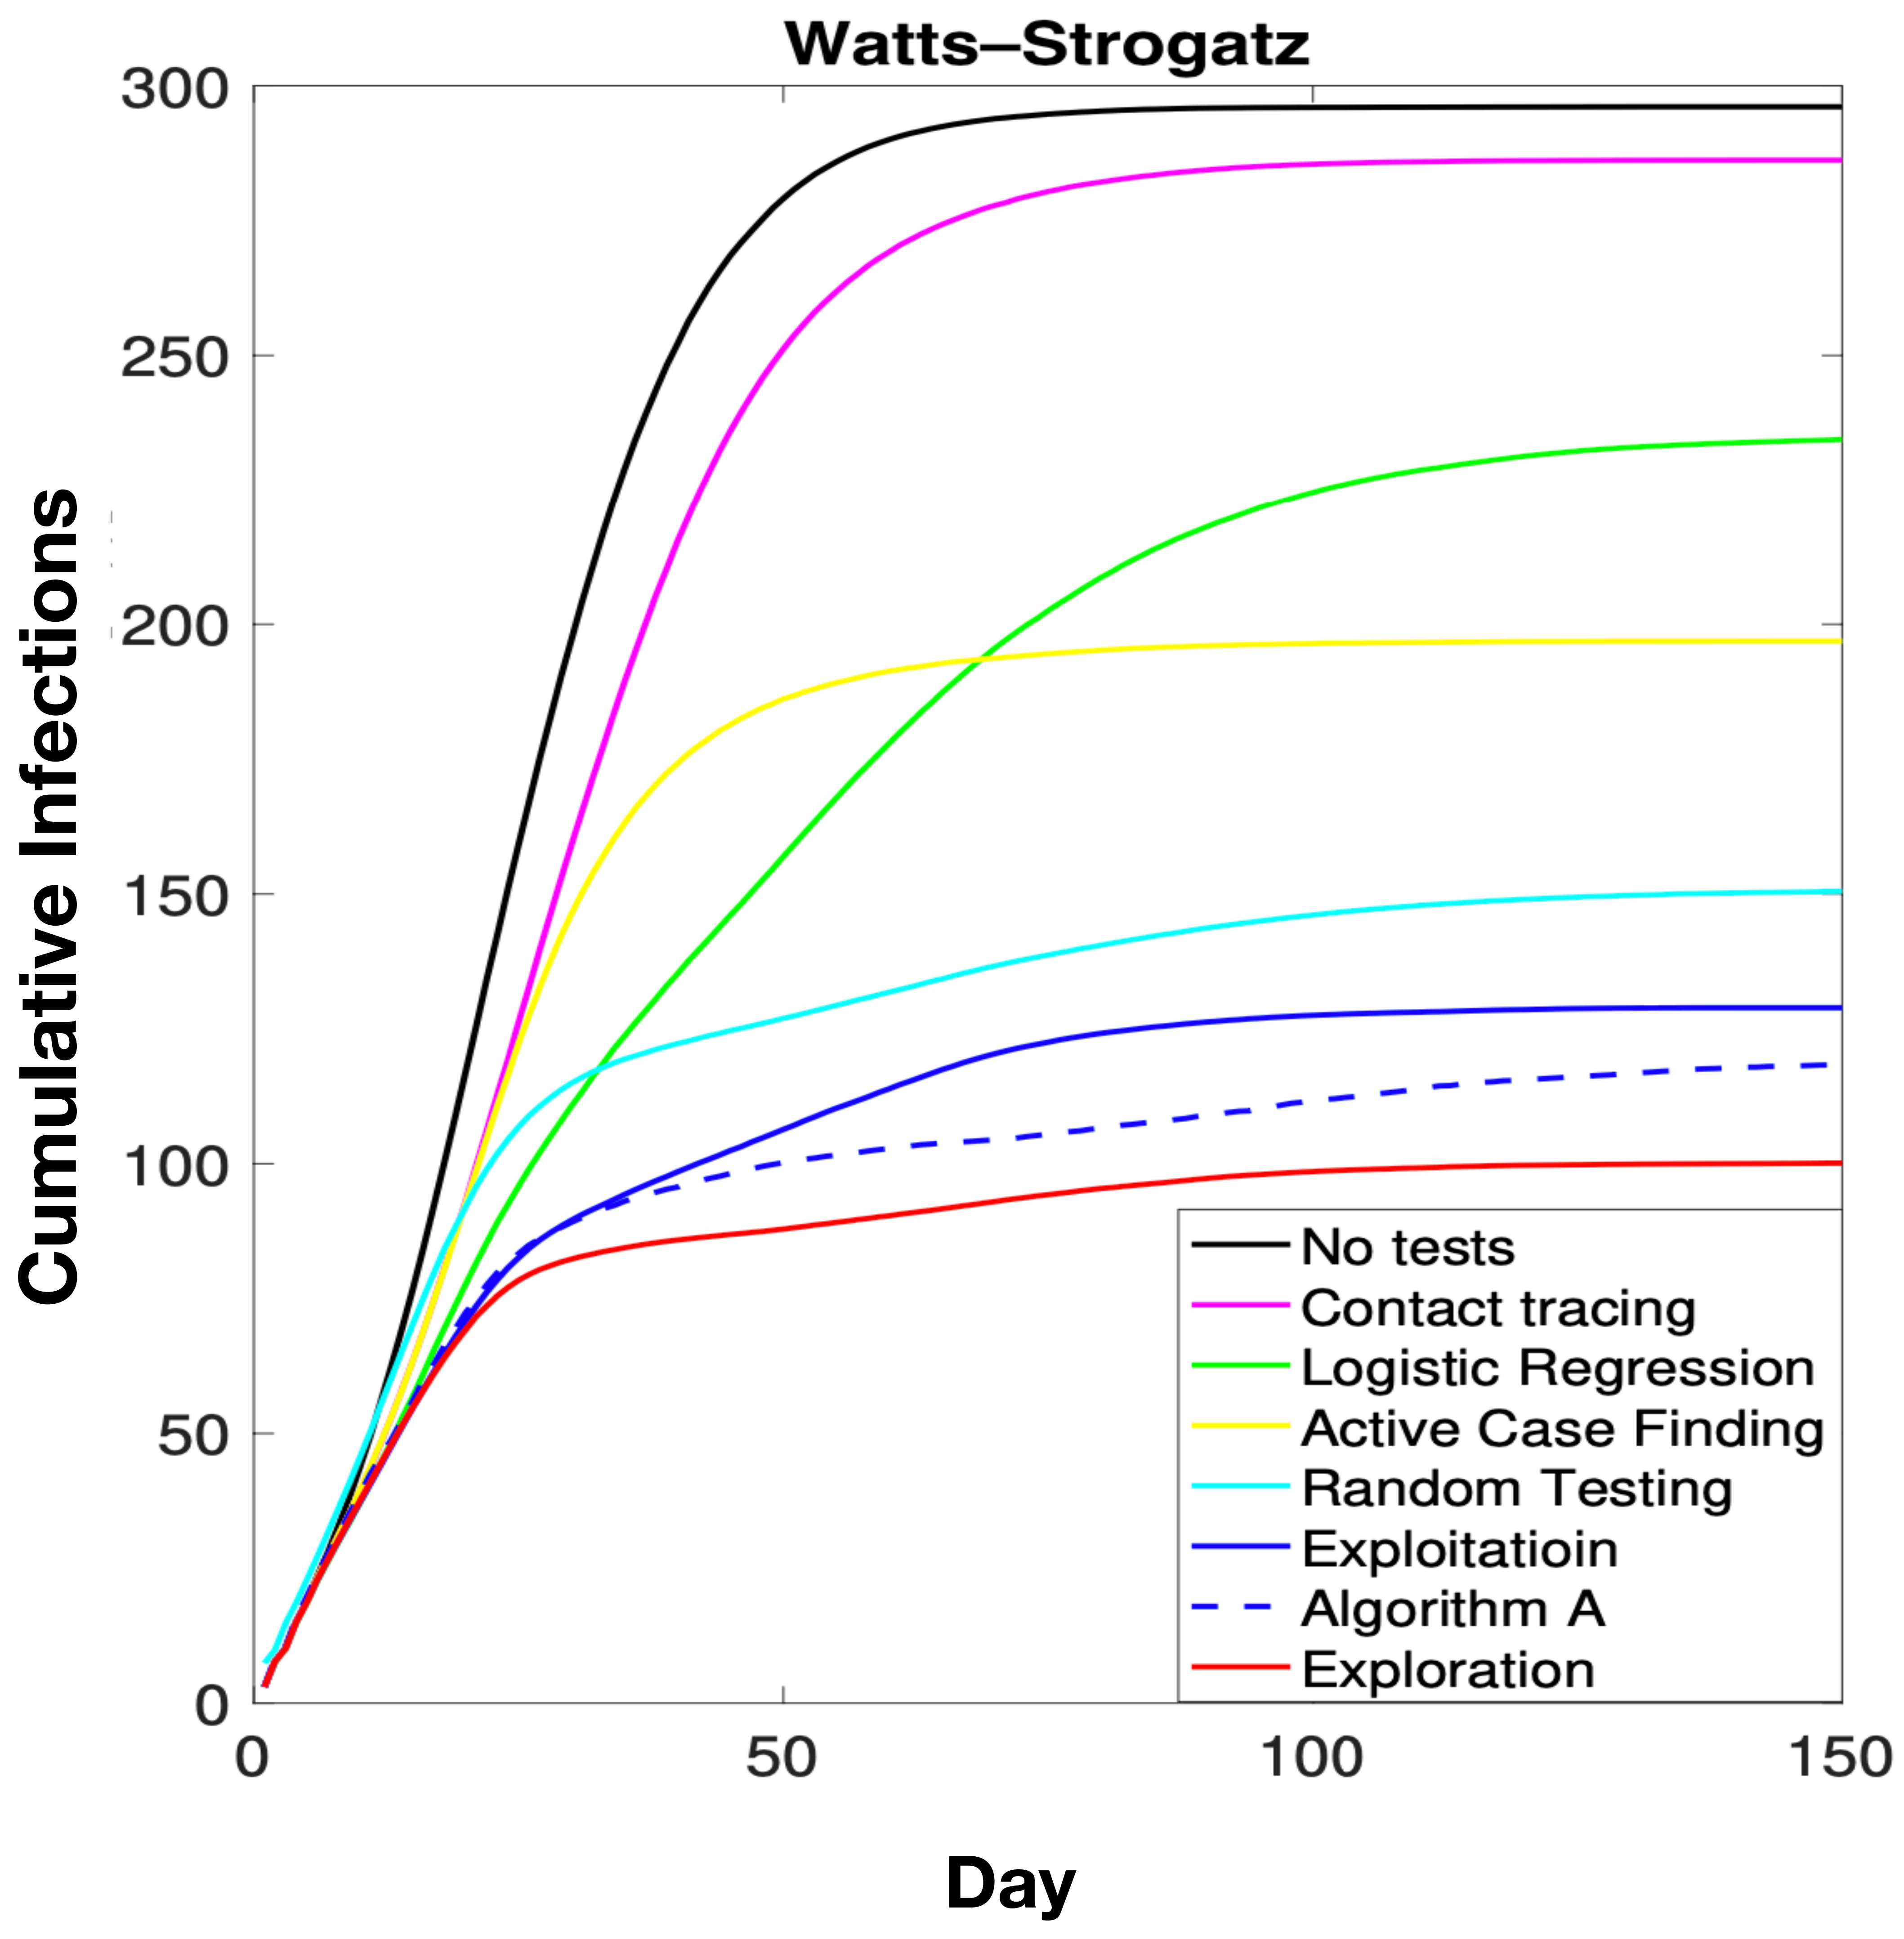

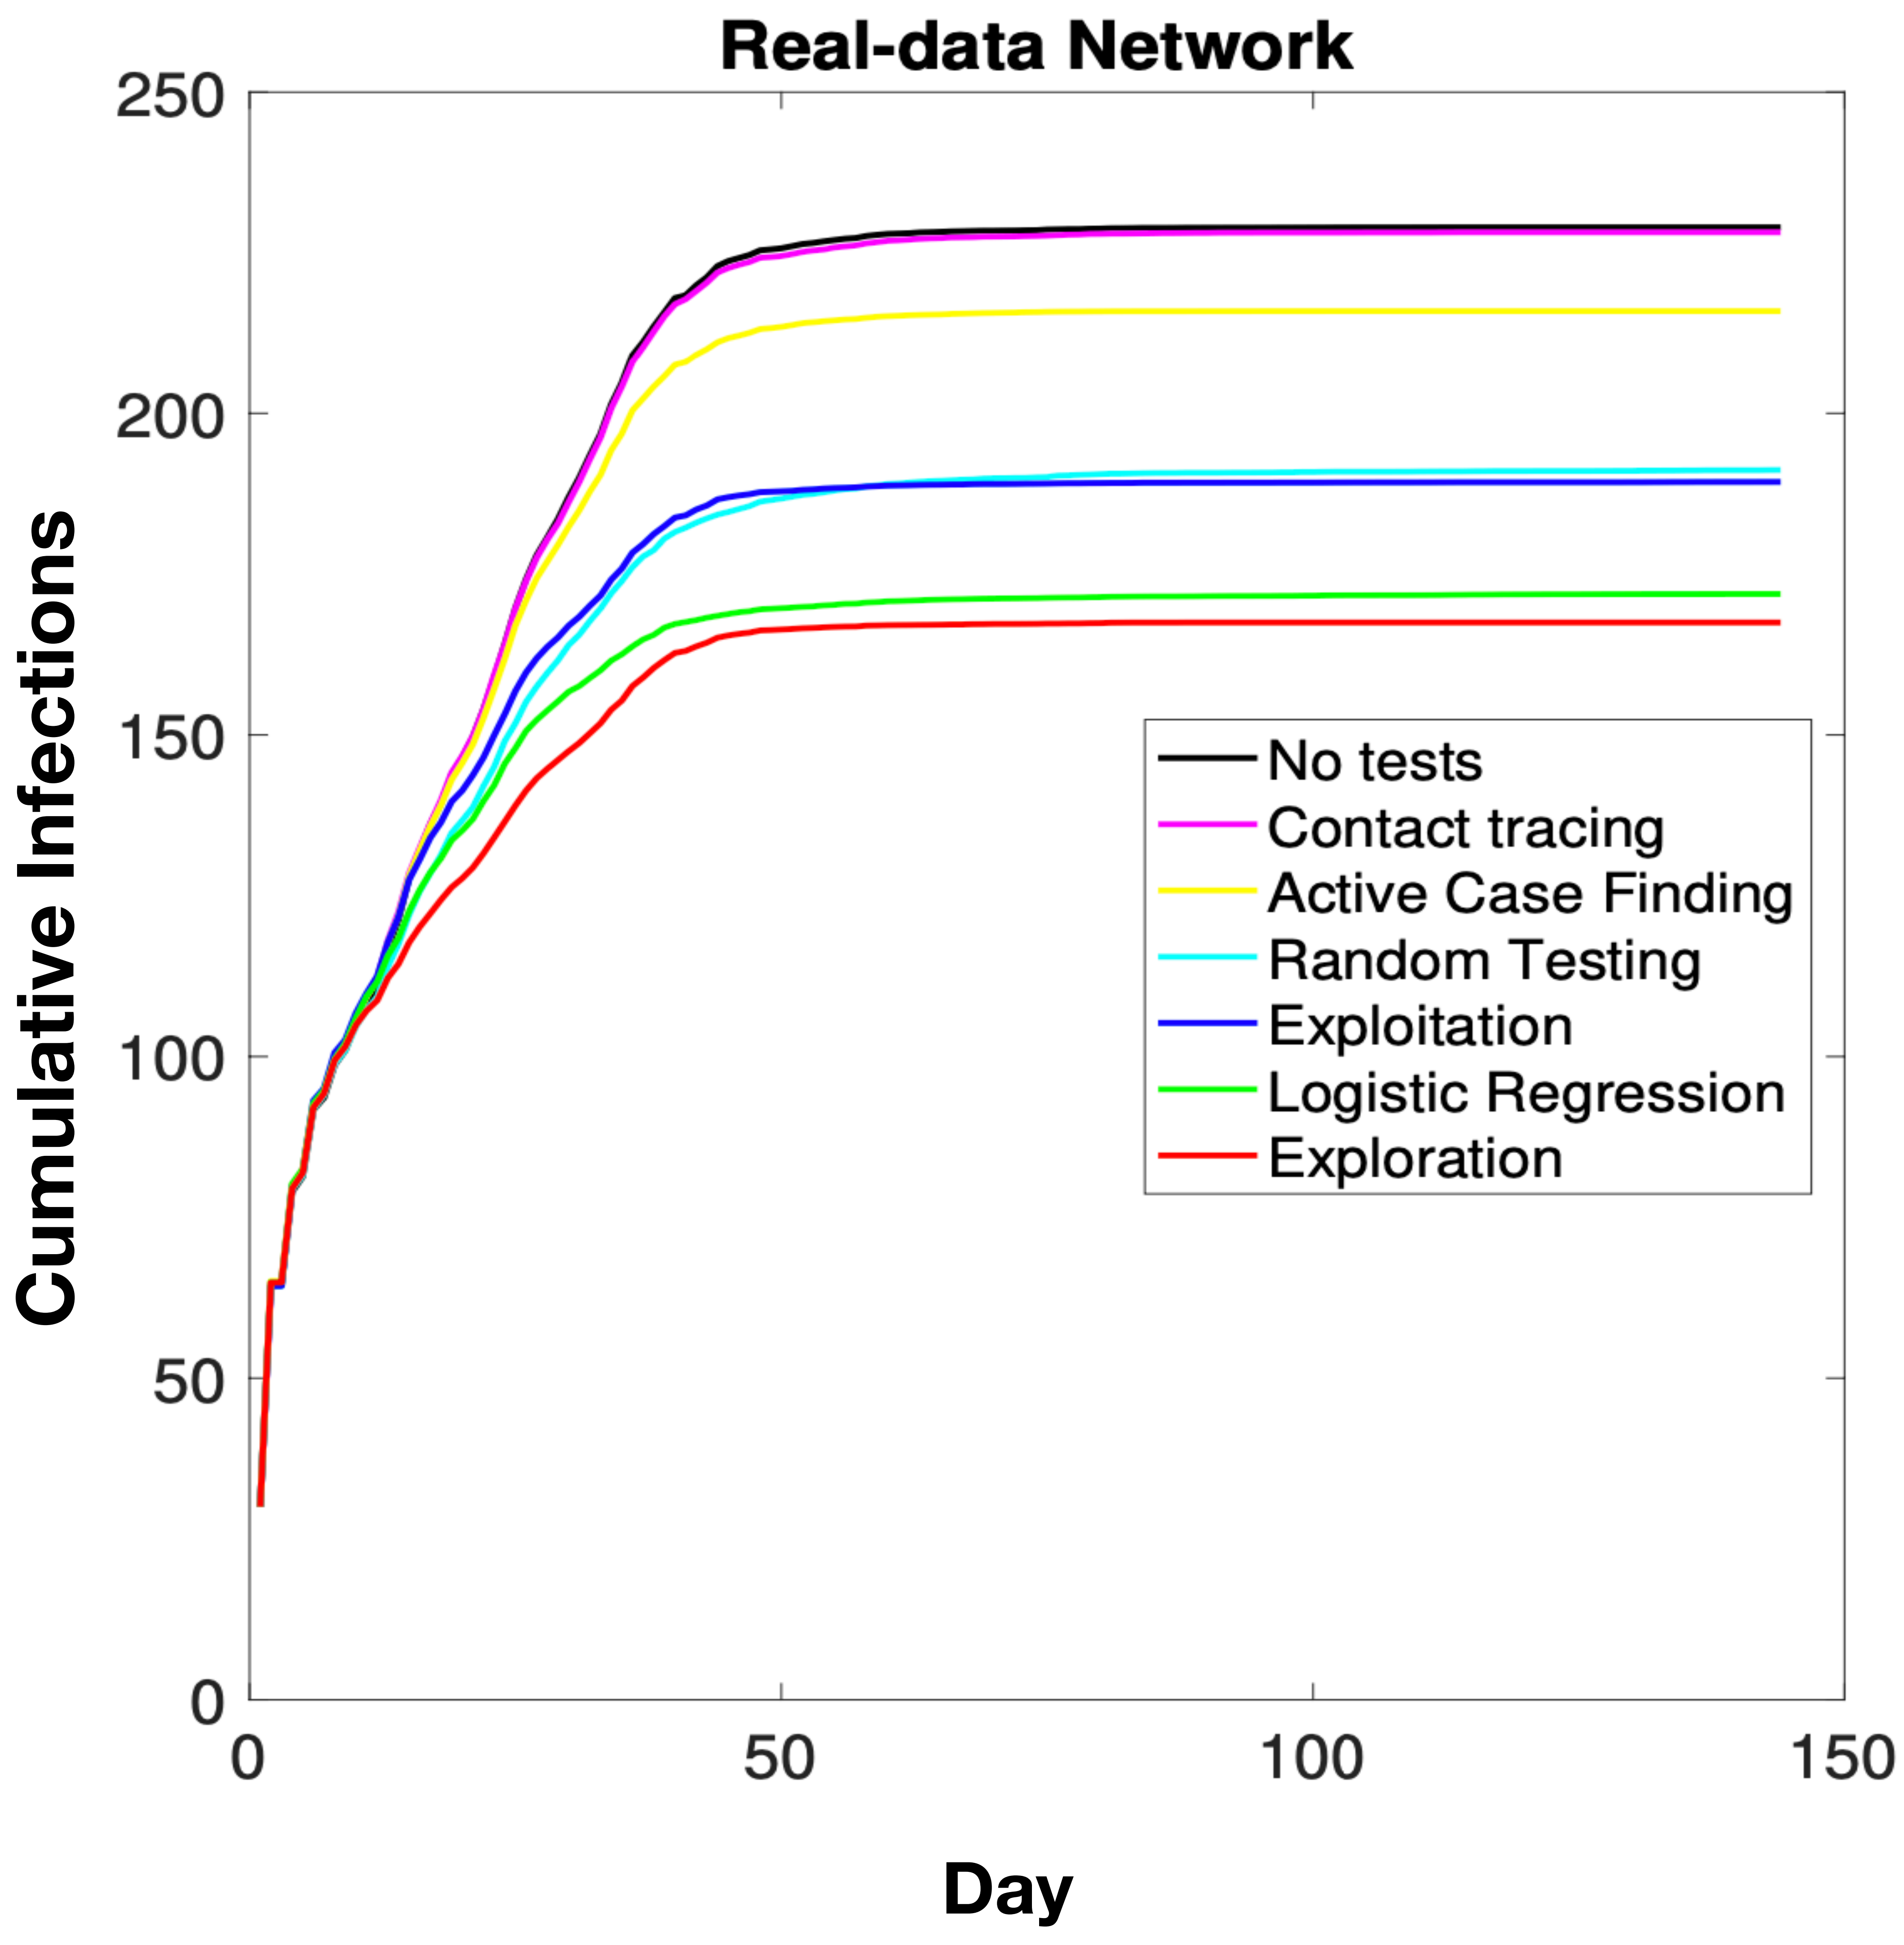

Chapter 6

Simulations

- Unregulated delay: initial start → the first time intervention start
- Clustering coefficient: nodes in a graph tend to cluster together
- Shortest path-length: the average shortest distance between every pairs
- **Conclusion**: When the above parameters increase, exploration becomes more beneficial as it provides better estimates of nodes' probabilities of infection.
- Relationship to previous chapters: (1) Temporal processes → Temporal and spatial processes; (2) Timeliness of nodes → Timeliness of networks
- Achievements:

- [J4] X. Chen, H. Nikpey, J. Kim, S. Sarkar, and S. Saeedi-Bidokhti, Containing a spread through sequential learning: to exploit or to explore? TMLR, 2023.

Acknowledgement

- Prof. Shirin Saeedi Bidokhti
- Prof. Saswati Sarkar, Prof. Hamed Hassani, Dr. Navid NaderiAlizadeh
- Collaborators
- Colleagues & ESE Department
- Thesis Committee

Thank you !