Stochastic Approximation MCMC, Online Inference, and Applications in Optimization of Queueing Systems

Xiang Li[1], Jiadong Liang[1], Xinyun Chen[2], and Zhihua Zhang[1]

[1]Peking University

[2]The Chinese University of Hong Kong, Shenzhen

Problem Formulation

GI/GI/1 Service System

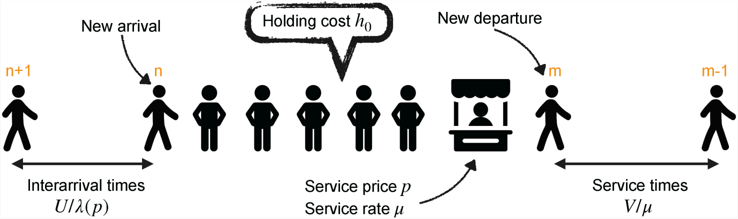

Figure 1: GI/GI/1 service

- **Service price**: p, **service rate**: μ, and **demand (customer) arrival rate** $\lambda(p)$.
- **Interarrival time**: $U_n/\lambda(p)$ with positive i.i.d random variable U_n and $\mathbb{E}U_n = 1$.
- **Service time**: V_n/μ with positive i.i.d random variable V_n and $\mathbb{E}V_n = 1$.
- **Holding cost per workload per unit time**: h_0.

Objective

Problem 1:

$$\min_{(\mu,p)\in\mathcal{B}} f(\mu,p) \equiv \underbrace{h_0\mathbb{E}[Q_\infty(\mu,p)]}_{\text{Holding cost}} + \underbrace{c(\mu)}_{\text{Staffing cost}} - \underbrace{p\lambda(p)}_{\text{Revenue}},\ \mathcal{B} \equiv [\underline{\mu},\bar{\mu}] \times [\underline{p},\bar{p}].$$

Objective

Problem 1:

$$\min_{(\mu,p)\in\mathcal{B}} f(\mu,p) \equiv \underbrace{h_0\mathbb{E}[Q_\infty(\mu,p)]}_{\text{Holding cost}} + \underbrace{c(\mu)}_{\text{Staffing cost}} - \underbrace{p\lambda(p)}_{\text{Revenue}},\ \mathcal{B} \equiv [\underline{\mu},\bar{\mu}] \times [\underline{p},\bar{p}].$$

This problem is equivalent to the **Problem 2:**

Find ϑ such that

$$h(\vartheta) := \int H(\vartheta,x)d\pi_\vartheta(x) = 0,\ \vartheta \in \mathcal{B}. \tag{1}$$

We denote the solution of (1) as ϑ^*

Equivalence

- First-order condition of Problem 1;
- Characterization of the steady state expectation's gradient by Chen et al. [2023].

Equivalence

- First-order condition of Problem 1;
- Characterization of the steady state expectation's gradient by Chen et al. [2023].

Combining these yields the form of $H(\vartheta, x)$ and π_ϑ with $\vartheta = (\mu, p)$ and $x = (w, y)$.

$$H(\vartheta, x) = \Bigg(\underbrace{-\lambda(p) - p\lambda'(p) + h_0\lambda'(p)\left(w + y + \frac{1}{\mu}\right)}_{\partial_p f(\mu, p)}, c'(\mu) - h_0\frac{\lambda(p)}{\mu}\left(w + y + \frac{1}{\mu}\right)\Bigg)$$

Equivalence

- First-order condition of Problem 1;
- Characterization of the steady state expectation's gradient by Chen et al. [2023].

Combining these yields the form of $H(\vartheta, x)$ and π_ϑ with $\vartheta = (\mu, p)$ and $x = (w, y)$.

$$H(\vartheta, x) = \left(-\lambda(p) - p\lambda'(p) + h_0\lambda'(p)\left(w + y + \tfrac{1}{\mu}\right),\ \underbrace{c'(\mu) - h_0\frac{\lambda(p)}{\mu}\left(w + y + \frac{1}{\mu}\right)}_{\partial_\mu f(\mu, p)}\right)$$

Equivalence

- First-order condition of Problem 1;
- Characterization of the steady state expectation's gradient by Chen et al. [2023].

Combining these yields the form of $H(\vartheta, x)$ and π_ϑ with $\vartheta = (\mu, p)$ and $x = (w, y)$.

$$H(\vartheta, x) = \left(-\lambda(p) - p\lambda'(p) + h_0\lambda'(p)\left(w + y + \tfrac{1}{\mu}\right), \underbrace{c'(\mu) - h_0\frac{\lambda(p)}{\mu}\left(w + y + \frac{1}{\mu}\right)}_{\partial_\mu f(\mu, p)}\right)$$

π_ϑ is the invariant distribution of (W_n, Y_n) under the choice (μ, p).

- **Waiting time** W_t: $W_{t+1} = \left(W_t - \frac{U_t}{\lambda(p)} + \frac{V_t}{\mu}\right)_+$;
- **Busy period** Y_t: $Y_{t+1} = \left(Y_t + \frac{U_t}{\lambda(p)}\right) \cdot \mathbb{1}\{W_{t+1} > 0\}$.

Batch Based Methods

Characteristic of $h(\vartheta) = \int H(\vartheta, x) d\pi_\vartheta(x)$:

- π_ϑ is dependent on ϑ;

Batch Based Methods

Characteristic of $h(\vartheta) = \int H(\vartheta, x) d\pi_\vartheta(x)$:

- π_ϑ is dependent on ϑ;
- We can't reach the sample from π_ϑ directly (often rely on some transition kernel $P_\vartheta(\cdot|\cdot)$).

Batch Based Methods

Characteristic of $h(\vartheta) = \int H(\vartheta, x) d\pi_\vartheta(x)$:

- π_ϑ is dependent on ϑ;
- We can't reach the sample from π_ϑ directly (often rely on some transition kernel $P_\vartheta(\cdot|\cdot)$).

Batch based algorithms:

- Fix ϑ_n, run the Markov chain $P_{\vartheta_n}(\cdot|\cdot)$ until mixing;

Batch Based Methods

Characteristic of $h(\vartheta) = \int H(\vartheta, x) d\pi_\vartheta(x)$:

- π_ϑ is dependent on ϑ;
- We can't reach the sample from π_ϑ directly (often rely on some transition kernel $P_\vartheta(\cdot|\cdot)$).

Batch based algorithms:

- Fix ϑ_n, run the Markov chain $P_{\vartheta_n}(\cdot|\cdot)$ until mixing;
- Collect a batch of data with size $\log n$, build an estimator H_n for $h(\vartheta_n)$;

Batch Based Methods

Characteristic of $h(\vartheta) = \int H(\vartheta, x) d\pi_\vartheta(x)$:

- π_ϑ is dependent on ϑ;
- We can't reach the sample from π_ϑ directly (often rely on some transition kernel $P_\vartheta(\cdot|\cdot)$).

Batch based algorithms:

- Fix ϑ_n, run the Markov chain $P_{\vartheta_n}(\cdot|\cdot)$ until mixing;
- Collect a batch of data with size $\log n$, build an estimator H_n for $h(\vartheta_n)$;
- Update $\vartheta_{n+1} = \Pi_{\mathcal{B}}(\vartheta_n - \eta_n H_n)$

Related Works: SGD for Steady State Optimization

- Pricing and capacity sizing in queues:
  - [Kim and Randhawa, 2018]:
  - [Chen et al., 2023]: optimizing optimal price and service rate by SGD whose batch size goes to infinity.
- Learning in queues:
  - [Suri and Zazanis, 1988, L'Ecuyer and Glynn, 1994]:
  - [Fu, 1990, Ravner and Snitkovsky, 2023]:
- Inventory models:
  - [Huh et al., 2009, Zhang et al., 2020, Yuan et al., 2021].

Gradient-based Online Learning in Queue (GOLiQ)[Chen et al., 2023]

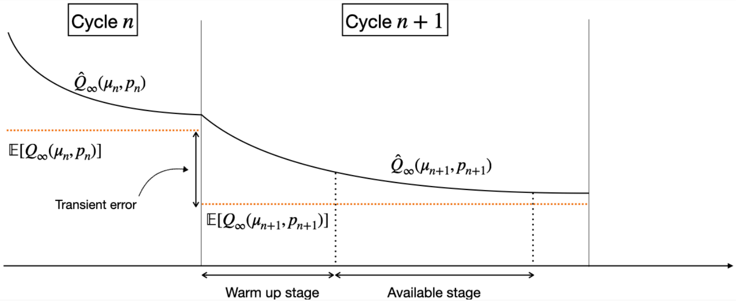

Figure 2: Update rule of GOLiQ

Related Works: GOLiQ

Gradient-based Online Learning in Queue (GOLiQ)[Chen et al., 2023]

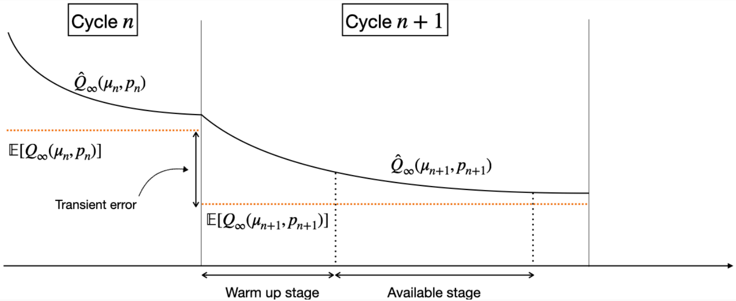

Figure 2: Update rule of GOLiQ

- Computational burden on the optimizer;

Related Works: GOLiQ

Gradient-based Online Learning in Queue (GOLiQ)[Chen et al., 2023]

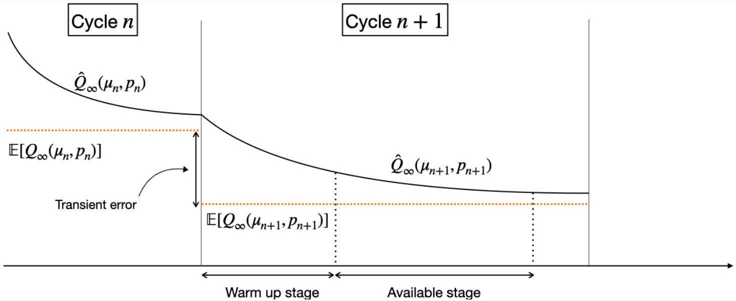

Figure 2: Update rule of GOLiQ

- Computational burden on the optimizer;
- Waste of data;

Related Works: GOLiQ

Gradient-based Online Learning in Queue (GOLiQ)[Chen et al., 2023]

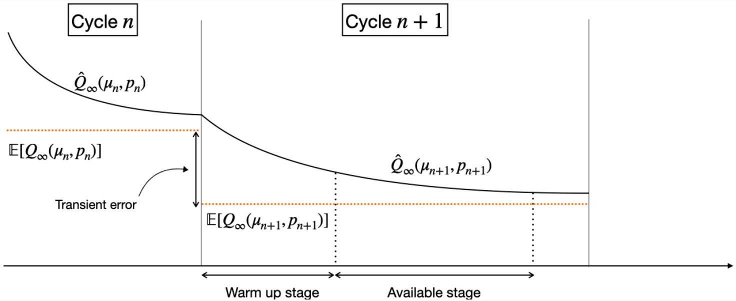

Figure 2: Update rule of GOLiQ

- Computational burden on the optimizer;
- Waste of data;
- No statistical inference.

Is there any method that finds the optimizer adaptively during GI/GI/1 queue service system is processing?

Stochastic Approximation MCMC

Recall the target function of **Problem 2:**

$$h(\vartheta) = \int H(\vartheta, x) d\pi(x)$$

Stochastic Approximation with Markov Chain Monte Carlo

Recall the target function of Problem 2:

$$h(\vartheta) = \int H(\vartheta, x) d\pi(x)$$

Batch based algorithm with batch size $\equiv 1$.

Stochastic Approximation with Markov Chain Monte Carlo

Recall the target function of Problem 2:

$$h(\vartheta) = \int H(\vartheta, x) d\pi(x)$$

Batch based algorithm with batch size $\equiv 1$.

Routine

- Draw **one** sample x_{k+1} with a Markov transition kernel P_{ϑ_k}, which admits π_{ϑ_k} as the invariant distribution;

Stochastic Approximation with Markov Chain Monte Carlo

Recall the target function of Problem 2:

$$h(\vartheta) = \int H(\vartheta, x) d\pi(x)$$

Batch based algorithm with batch size $\equiv 1$.

Routine

- Draw **one** sample x_{k+1} with a Markov transition kernel P_{ϑ_k}, which admits π_{ϑ_k} as the invariant distribution;
- Update $\vartheta_{k+1} := \Pi_{\mathcal{B}}(\vartheta_k - \eta_k H(\vartheta_k, x_{k+1}))$.

Related Works: State-Dependent Stochastic Approximation

- [Andrieu et al., 2005]: Stability and consistency of SAMCMC.
- [Liang, 2010]: Strong consistency and asymptotic normality of SAMCMC.
- [Atchadé et al., 2017]: Batch based state-dependent stochastic proximal algorithm.
- [Karimi et al., 2019, Li and Wai, 2022]: Nonasymptotic convergence rate of SAMCMC under various conditions.

Related Works: State-Dependent Stochastic Approximation

- [Andrieu et al., 2005]: Stability and consistency of SAMCMC.
- [Liang, 2010]: Strong consistency and asymptotic normality of SAMCMC.
- [Atchadé et al., 2017]: Batch based state-dependent stochastic proximal algorithm.
- [Karimi et al., 2019, Li and Wai, 2022]: Nonasymptotic convergence rate of SAMCMC under various conditions.

Limitations:

- Lack of convergence rate for strongly convex problems;

Related Works: State-Dependent Stochastic Approximation

- [Andrieu et al., 2005]: Stability and consistency of SAMCMC.
- [Liang, 2010]: Strong consistency and asymptotic normality of SAMCMC.
- [Atchadé et al., 2017]: Batch based state-dependent stochastic proximal algorithm.
- [Karimi et al., 2019, Li and Wai, 2022]: Nonasymptotic convergence rate of SAMCMC under various conditions.

Limitations:

- Lack of convergence rate for strongly convex problems;
- Lack of analysis in online learning scenario;

Related Works: State-Dependent Stochastic Approximation

- [Andrieu et al., 2005]: Stability and consistency of SAMCMC.
- [Liang, 2010]: Strong consistency and asymptotic normality of SAMCMC.
- [Atchadé et al., 2017]: Batch based state-dependent stochastic proximal algorithm.
- [Karimi et al., 2019, Li and Wai, 2022]: Nonasymptotic convergence rate of SAMCMC under various conditions.

Limitations:

- Lack of convergence rate for strongly convex problems;
- Lack of analysis in online learning scenario;
- Restricted and abstract condition (bounded noise, uniformly Lipschitz).

How to determine if SAMCMC is suitable for the GI/GI/1 service model? If applicable, does its performance surpass batch training-based algorithms?

Our Contributions

- We use the SAMCMC algorithm for the first time to address the pricing and capacity sizing problem in the GI/GI/1 service system.
- Optimal convergence rate on finding the best choice of $\vartheta = (\mu, p)$, $\mathbb{E} \left\| \vartheta_t - \vartheta^* \right\|^2 = \mathcal{O}(1/t)$.
- Optimal cumulative regret $\mathcal{O}(\log t)$.
- Efficient online statistical inference scheme based on the functional central limit theorem of ϑ_n.

SAMCMC with Queue

Recall the form of $H(\vartheta, x)$ in **Problem 2**.

$$H(\vartheta, x) = \left(-\lambda(p) - p\lambda'(p) + h_0\lambda'(p)\left(w + y + \frac{1}{\mu}\right), c'(\mu) - h_0\frac{\lambda(p)}{\mu}\left(w + y + \frac{1}{\mu}\right)\right)$$

SAMCMC with Queue

Recall the form of $H(\vartheta, x)$ in **Problem 2**.

$$H(\vartheta,x)=\left(-\lambda(p)-p\lambda'(p)+h_0\lambda'(p)\left(w+y+\frac{1}{\mu}\right),\, c'(\mu)-h_0\frac{\lambda(p)}{\mu}\left(w+y+\frac{1}{\mu}\right)\right)$$

Algorithm: SAMCMC for GI/GI/1 service system

1. Customer n gets service;
2. Compute recursively

$$W_{n+1}=(W_n-\frac{U_n}{\lambda(p)}+\frac{V_n}{\mu})_+;\ Y_{n+1}=(Y_n+\frac{U_n}{\lambda(p)})\mathbb{1}\{W_{n+1}>0\};$$

3. Construct $H(\vartheta_n, x_{n+1})$;
4. Update (μ_{n+1}, p_{n+1}).

Convergence and Regret

Assumptions

- Smooth and strongly convex objective function
- Light-tail service time
- Uniform stability
- Appropriate step size choice

Synchronous Coupling

Q: Why does GOLiQ need increasing batch size?

Synchronous Coupling

Q: Why does GOLiQ need increasing batch size?

A: Neglect of the smooth dependence of parameters (μ, p) on GI/GI/1 dynamics.

Synchronous Coupling

Q: Why does GOLiQ need increasing batch size?

A: Neglect of the smooth dependence of parameters (μ, p) on GI/GI/1 dynamics.

How to characterize it? Synchronous coupling !

Synchronous Coupling

Q: Why does GOLiQ need increasing batch size?

A: Neglect of the smooth dependence of parameters (μ, p) on GI/GI/1 dynamics.

How to characterize it? Synchronous coupling !

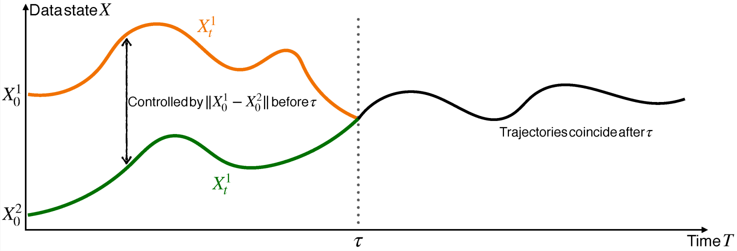

$$W_{t+1} = (W_t - \frac{U_t}{\lambda(p)} + \frac{V_t}{\mu})_+; \;\; Y_{t+1} = (Y_t + \frac{U_t}{\lambda(p)})\mathbb{1}\{W_{t+1} > 0\}$$

Poisson's Equation

Solution of **the Poisson's equation**:

$$H(\vartheta, x) - \mathbb{E}[H(\vartheta, X_\infty(\vartheta))] = U_H(\vartheta, x) - P_\vartheta U_H(\vartheta, x)$$

Poisson's Equation

Solution of **the Poisson's equation**:

$$H(\vartheta, x) - \mathbb{E}[H(\vartheta, X_\infty(\vartheta))] = U_H(\vartheta, x) - P_\vartheta U_H(\vartheta, x)$$

Lemma (Lipschitz continuity of $U_H(\cdot, x)$)

For any initial state $x = (w, y)^\top$ and two settings of parameter $\vartheta_i = (\mu_i, \lambda_i)^\top,\ i = 1, 2$, we have

$$\|U_H(\vartheta_1, x) - U_H(\vartheta_2, x)\| \leq C \|\vartheta_1 - \vartheta_2\| (\|x\| + 1).$$

Poisson's Equation

Solution of **the Poisson's equation**:

$$H(\vartheta, x) - \mathbb{E}[H(\vartheta, X_\infty(\vartheta))] = U_H(\vartheta, x) - P_\vartheta U_H(\vartheta, x)$$

Lemma (Lipschitz continuity of $U_H(\cdot, x)$)

For any initial state $x = (w, y)^\top$ and two settings of parameter $\vartheta_i = (\mu_i, \lambda_i)^\top,\ i = 1, 2$, we have

$$\|U_H(\vartheta_1, x) - U_H(\vartheta_2, x)\| \leq C \|\vartheta_1 - \vartheta_2\| (\|x\| + 1).$$

Insure the decomposition of random oracle $H(\vartheta_t, X_{t+1})$ with **diminishing residual terms**,

$$H(\vartheta_t, X_{t+1}) = h(\vartheta_n) + \text{zero mean noise} + o(1).$$

Convergence Rate

Theorem ($\mathcal{L}_2$ Convergence)

If we use SAMCMC to estimate $\vartheta = (\mu, p)$ in a stream style, then given the above mild assumptions, we have

$$\mathbb{E}\left\|\vartheta_t - \vartheta^*\right\|^2 = \mathcal{O}(\eta_t)$$

Minimax Optimal: When $\eta_t = \gamma t^{-1}$, $\mathbb{E}\left\|\vartheta_t - \vartheta^*\right\|^2$ could achieve the optimal order of t, i.e., $\mathcal{O}(t^{-1})$.

Regret Analysis

Upon SAMCMC finishes T iterations;

- I_t: The duration of iteration t.
- $\sum_{t=1}^{T} I_t$: Total time elapsed.

Regret Analysis

Upon SAMCMC finishes T iterations;

- l_t: The duration of iteration t.
- $\sum_{t=1}^{T} l_t$: Total time elapsed.
- The cumulated regret in T iterations

$$R(T) := \mathbb{E}\left[\sum_{t=2}^{T+1} \left(h_0(w_t + S_t) - p_t + c(\mu_t) l_t\right) - f(\mu^\star, p^\star) \sum_{t=2}^{T+1} l_t\right]$$

Regret Analysis

Upon SAMCMC finishes T iterations;

- I_t: The duration of iteration t.
- $\sum_{t=1}^{T} I_t$: Total time elapsed.
- The cumulated regret in T iterations

$$R(T) := \mathbb{E}\left[\sum_{t=2}^{T+1} \left(h_0(w_t + S_t) - p_t + c(\mu_t)I_t\right) - f(\mu^\star, p^\star)\sum_{t=2}^{T+1} I_t\right]$$

Theorem (Optimal Regret)

For the cumulative regret defined above, under regular assumptions, when the step size is chosen as $\eta_t = \frac{\gamma}{t}$, then we have that

$$\mathbb{E}R(T) = \mathcal{O}(\log T).$$

Online Statistical Inference

Standardized Time Series

Proposed by Glynn and Iglehart [1990].

For a discrete-time stochastic process x_t, if

$$\bar{x}_T(r) := \frac{1}{\sqrt{T}} \sum_{t=1}^{[Tr]} (x_t - \mu) \xrightarrow{d.} \underbrace{\sigma}_{\text{nuisance}} \cdot \underbrace{\chi(r)}_{\text{known}}$$

Standardized Time Series

Proposed by Glynn and Iglehart [1990].

For a discrete-time stochastic process x_t, if

$$\bar{x}_T(r) := \frac{1}{\sqrt{T}} \sum_{t=1}^{[Tr]} (x_t - \mu) \xrightarrow{d.} \underbrace{\sigma}_{\text{nuisance}} \cdot \underbrace{\chi(r)}_{\text{known}}$$

Consider a **positive continuous** functional $g : C[0,1] \to \mathbb{R}$ such that

- **Homogeneity:** $g(\alpha x) = \alpha g(x),\ \forall \alpha > 0$ and $x \in C[0,1]$;
- **Location invariance:** $g(x - \beta e) = g(x)$ for $\beta \in \mathbb{R}$ and $x \in C[0,1]$, where $e(t) = t$.

Standardized Time Series

Proposed by Glynn and Iglehart [1990].

For a discrete-time stochastic process x_t, if

$$\bar{x}_T(r) := \frac{1}{\sqrt{T}} \sum_{t=1}^{[Tr]} (x_t - \mu) \xrightarrow{d.} \underbrace{\sigma}_{\text{nuisance}} \cdot \underbrace{\chi(r)}_{\text{known}}$$

Consider a **positive continuous** functional $g : C[0,1] \to \mathbb{R}$ such that

- **Homogeneity:** $g(\alpha x) = \alpha g(x), \ \forall \alpha > 0$ and $x \in C[0,1]$;
- **Location invariance:** $g(x - \beta e) = g(x)$ for $\beta \in \mathbb{R}$ and $x \in C[0,1]$, where $e(t) = t$.

$$\frac{\bar{x}_T(1) - \mu}{g(\bar{x}_T)} \xrightarrow{d.} \frac{\chi(1)}{g(\chi)}$$

Standardized Time Series

Proposed by Glynn and Iglehart [1990].

For a discrete-time stochastic process x_t, if

$$\bar{x}_T(r) := \frac{1}{\sqrt{T}} \sum_{t=1}^{[Tr]} (x_t - \mu) \xrightarrow{d.} \underbrace{\sigma}_{\text{nuisance}} \cdot \underbrace{\chi(r)}_{\text{known}}$$

Consider a **positive continuous** functional $g : C[0,1] \to \mathbb{R}$ such that

- **Homogeneity:** $g(\alpha x) = \alpha g(x),\ \forall \alpha > 0$ and $x \in C[0,1]$;
- **Location invariance:** $g(x - \beta e) = g(x)$ for $\beta \in \mathbb{R}$ and $x \in C[0,1]$, where $e(t) = t$.

$$\frac{\bar{x}_T(1) - \mu}{g(\bar{x}_T)} \xrightarrow{d.} \frac{\chi(1)}{g(\chi)} \qquad \textbf{Cancel out } \sigma\textbf{ !}$$

Functional Central Limit Theorem

Local linearity assumption: There exists a positive definite matrix G and constant L_G, δ_G such that

$$\|g(\vartheta) - G(\vartheta - \vartheta^*)\| \leq L_G \cdot \|\vartheta - \vartheta^*\|^2 \text{ for any } \|\vartheta - \vartheta^*\| \leq \delta_G.$$

.

Theorem (Functional Central Limit Theorem)

Under the above all assumptions, we have

$$\phi_T(r) := \frac{1}{\sqrt{T}} \sum_{t=1}^{\lfloor Tr \rfloor} (\vartheta_t - \vartheta^\star) \xrightarrow{d} \psi(r) := G^{-1} S^{1/2} W(r)$$

within the topology $\mathcal{D}([0,1]; \|\cdot\|_\infty)$.

S is the asymptotic covariance related to the Markovian dynamics of GI/GI/1 system.

Construction of Confidence Intervals: Random Scaling Method

- 'Cancel out function': $g(\phi) = \sqrt{\int_0^1 (\phi_r - r\phi_1)^2 dr}$;
- Inference statistics: $h(\phi_T) = v^\top \bar{\vartheta}_T \Big/ \left(\frac{1}{T} \sqrt{\sum_{i=1}^{T} i^2 \left| v^\top (\bar{\vartheta}_i - \bar{\vartheta}_T) \right|^2} \right)$
- Computed in an online manner. Simple and flexible.

Construction of Confidence Intervals: Random Scaling Method

- **'Cancel out function':** $g(\phi) = \sqrt{\int_0^1 (\phi_r - r\phi_1)^2 dr}$;
- **Inference statistics:** $h(\boldsymbol{\phi}_T) = v^\top \bar{\vartheta}_T \Big/ \left(\frac{1}{T} \sqrt{\sum_{i=1}^{T} i^2 \left| v^\top (\bar{\vartheta}_i - \bar{\vartheta}_T) \right|^2} \right)$
- Computed in an online manner. Simple and flexible.

Batch means: [Chen et al., 2020, Zhu and Dong, 2021]

- $h(\boldsymbol{\phi}_T) = \boldsymbol{\phi}_T(1) \Big/ \sqrt{\frac{1}{m(m-1)} \sum_{i=1}^{m} \left(\frac{\boldsymbol{\phi}_T(e_i) - \boldsymbol{\phi}_T(s_i)}{e_i - s_i} - \boldsymbol{\phi}_T(1) \right)^2}$;

Construction of Confidence Intervals: Random Scaling Method

- ‘Cancel out function’: $g(\phi) = \sqrt{\int_0^1 (\phi_r - r\phi_1)^2 dr}$;
- Inference statistics: $h(\boldsymbol{\phi}_T) = v^\top \bar{\vartheta}_T \Big/ \left(\frac{1}{T} \sqrt{\sum_{i=1}^{T} i^2 \left| v^\top (\bar{\vartheta}_i - \bar{\vartheta}_T) \right|^2} \right)$
- Computed in an online manner. Simple and flexible.

Batch means: [Chen et al., 2020, Zhu and Dong, 2021]

- $h(\boldsymbol{\phi}_T) = \boldsymbol{\phi}_T(1) \Big/ \sqrt{\frac{1}{m(m-1)} \sum_{i=1}^{m} \left(\frac{\boldsymbol{\phi}_T(e_i) - \boldsymbol{\phi}_T(s_i)}{e_i - s_i} - \boldsymbol{\phi}_T(1) \right)^2}$;
- **Complicate:** should adjust (s_i, e_i);
- **Inconvenient:** often choose increasing batch size.

Numerical Experiment

Convergence and Regret Results

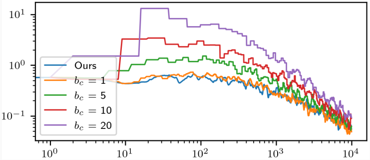

(a) Convergence rate under M/M/1

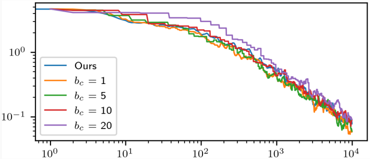

(b) Convergence rate under M/GI/1

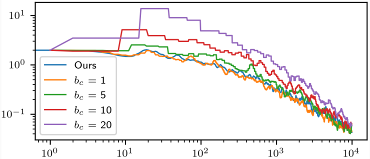

(c) Convergence rate under GI/M/1

b_c indicates that the batch size in the t-th cycle is $b_c \log t$ when we use GOLiQ.

Convergence and Regret Results

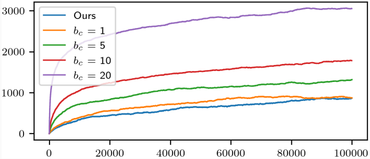

(d) Regret bound under M/M/1

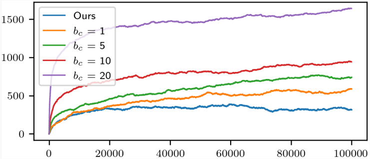

(e) Regret bound under M/GI/1

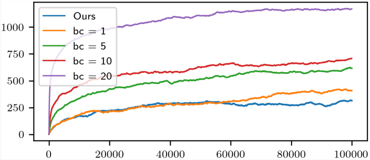

(f) Regret bound under GI/M/1

b_c indicates that the batch size in the t-th cycle is $b_c \log t$ when we use GOLiQ.

Inference

Table 1: Comparison of empirical coverage rates between random scaling and the batch mean. ES and IBS represent even-splitting and increasing batch size respectively.

Para.	Settings	Methods	$n = T/2^4$	$n = T/2^3$	$n = T/2^2$	$n = T/2$	$n = T$
p	M/M/1	Ours	89.2(1.39)	90.4(1.32)	93.2(1.13)	92.4(1.19)	94.8(0.99)
		ES	80.6(1.77)	86.2(1.54)	90.2(1.33)	91.6(1.24)	92.4(1.19)
		IBS	84.8(1.61)	92.4(1.19)	94.0(1.06)	95.6(0.92)	96.4(0.83)
μ	M/M/1	Ours	97.0(0.76)	97.0(0.76)	95.8(0.9)	94.0(1.06)	94.4(1.03)
		ES	97.8(0.66)	97.4(0.71)	96.6(0.81)	93.8(1.08)	93.6(1.09)
		IBS	95.4(0.94)	94.6(1.01)	90.8(1.29)	89.2(1.39)	90.4(1.32)

Summary

- Introduce GI/GI/1 service problem and batch-based method;
- Usage of SAMCMC;
- Optimal convergence rate and regret bound;
- FCLT and online inference procedure;
- Persuasive numerical experiments.

Questions?

Wasserstein Geometrical Ergodicity

Recall the update rule of **waiting time - busy period pair**:

$$W_{n+1} = (W_n - \frac{U_n}{\lambda(p)} + \frac{V_n}{\mu})_+; \; Y_{n+1} = (Y_n + \frac{U_n}{\lambda(p)})\mathbb{1}\{W_{n+1} > 0\}$$

Lemma (Wasserstein Geometrical Ergodicity)

For two sequences $\{(W_n^i, Y_n^i)\}_{n=0}^{1}$ with $i = 1, 2$ whose initial state are separately (W_0^1, Y_0^1) and (W_0^2, Y_0^2). Then under identical coupling, we have

$$\mathbb{E}\left[\left\|W_n^1 - W_n^2\right\| \mid W_0^1, W_0^2\right] = \mathcal{O}\left(\left\|W_0^1 - W_0^2\right\| e^{-\gamma n}\right).$$

$$\mathbb{E}\left[\left\|Y_n^1 - Y_n^2\right\| \mid \{W_0^i, Y_0^i\}_{i=1}^2\right] = \mathcal{O}\left((\left\|Y_0^1 - Y_0^2\right\| + \left\|W_0^1 - W_0^2\right\|)e^{-\gamma n}\right).$$

Online Computation of Confidence Intervals

The denominator of $h(\phi_T(r))$ can be computed in an online manner.

$$\begin{aligned}\sum_{i=1}^{t} i^2 |v^\top(\bar{\vartheta}_i - \bar{\vartheta}_t)|^2 &= \sum_{i=1}^{t} i^2 \left((v^\top \bar{\vartheta}_i)^2 - 2v^\top \bar{\vartheta}_i v^\top \bar{\vartheta}_t + (v^\top \bar{\vartheta}_t)^2\right) \\ &= \sum_{i=1}^{t} i^2((v^\top \bar{\vartheta}_i)^2 - 2\sum_{i=1}^{t} i^2 v^\top \bar{\vartheta}_i v^\top \bar{\vartheta}_t + \sum_{i=1}^{t} i^2 (v^\top \bar{\vartheta}_t)^2 \\ &= O_{t,1} - 2O_{t,2} \cdot (v^\top \bar{\vartheta}_t) + O_{t,3} \cdot (v^\top \bar{\vartheta}_t)^2,\end{aligned}$$

where

$$O_{t,1} = \sum_{i=1}^{t} i^2((v^\top \bar{\vartheta}_i)^2, \ O_{t,2} = \sum_{i=1}^{t} i^2 v^\top \bar{\vartheta}_i, \ \text{and } O_{t,3} = \sum_{i=1}^{t} i^2.$$

Online Computation of Confidence Intervals

The denominator of $h(\phi_T(r))$ can be computed in an online manner.

$$O_{t,1} = \sum_{i=1}^{t} i^2((v^\top \bar{\vartheta}_i)^2, \; O_{t,2} = \sum_{i=1}^{t} i^2 v^\top \bar{\vartheta}_i, \text{ and } O_{t,3} = \sum_{i=1}^{t} i^2.$$

Each of them has their own update relation.

$$\begin{aligned} O_{t+1,1} &= O_{t,1} + (t+1)^2 (v^\top \bar{\vartheta}_{t+1})^2 \\ O_{t+1,2} &= O_{t,2} + (t+1)^2 (v^\top \bar{\vartheta}_{t+1}) \\ O_{t+1,3} &= O_{t,3} + (t+1)^2. \end{aligned}$$

References

Christophe Andrieu, Éric Moulines, and Pierre Priouret. Stability of stochastic approximation under verifiable conditions. *SIAM Journal on control and optimization*, 44(1):283–312, 2005.

Yves F Atchadé, Gersende Fort, and Eric Moulines. On perturbed proximal gradient algorithms. *The Journal of Machine Learning Research*, 18(1):310–342, 2017.

Xi Chen, Jason D Lee, Xin T Tong, and Yichen Zhang. Statistical inference for model parameters in stochastic gradient descent. *Annals of Statistics*, 48(1):251–273, 2020.

Xinyun Chen, Yunan Liu, and Guiyu Hong. An online learning approach to dynamic pricing and capacity sizing in service systems. *Operations Research*, 2023.

Michael C Fu. Convergence of a stochastic approximation algorithm for the gi/g/1 queue using infinitesimal perturbation analysis. *Journal of Optimization Theory and Applications*, 65:149–160, 1990.

Peter W Glynn and Donald L Iglehart. Simulation output analysis using standardized time series. *Mathematics of Operations Research*, 15(1):1–16, 1990.

Woonghee Tim Huh, Ganesh Janakiraman, John A Muckstadt, and Paat Rusmevichientong. An adaptive algorithm for finding the optimal base-stock policy in lost sales inventory systems with censored demand. *Mathematics of Operations Research*, 34(2):397–416, 2009.

References iii

Belhal Karimi, Blazej Miasojedow, Eric Moulines, and Hoi-To Wai. Non-asymptotic analysis of biased stochastic approximation scheme. In *Conference on Learning Theory*, pages 1944–1974. PMLR, 2019.

Jeunghyun Kim and Ramandeep S Randhawa. The value of dynamic pricing in large queueing systems. *Operations Research*, 66(2): 409–425, 2018.

Pierre L'Ecuyer and Peter W Glynn. Stochastic optimization by simulation: Convergence proofs for the gi/g/1 queue in steady-state. *Management Science*, 40(11):1562–1578, 1994.

Qiang Li and Hoi-To Wai. State dependent performative prediction with stochastic approximation. In *International Conference on Artificial Intelligence and Statistics*, pages 3164–3186. PMLR, 2022.

Faming Liang. Trajectory averaging for stochastic approximation mcmc algorithms1. *The Annals of Statistics*, 38(5):2823–2856, 2010.

Liron Ravner and Ran I Snitkovsky. Stochastic approximation of symmetric nash equilibria in queueing games. *Operations Research*, 2023.

Rajan Suri and Michael A Zazanis. Perturbation analysis gives strongly consistent sensitivity estimates for the m/g/1 queue. *Management Science*, 34(1):39–64, 1988.

Hao Yuan, Qi Luo, and Cong Shi. Marrying stochastic gradient descent with bandits: Learning algorithms for inventory systems with fixed costs. *Management Science*, 67(10):6089–6115, 2021.

Huanan Zhang, Xiuli Chao, and Cong Shi. Closing the gap: A learning algorithm for lost-sales inventory systems with lead times. *Management Science*, 66(5):1962–1980, 2020.

Yi Zhu and Jing Dong. On constructing confidence region for model parameters in stochastic gradient descent via batch means. In *2021 Winter Simulation Conference (WSC)*, pages 1–12. IEEE, 2021.