Stein's Method for Steady-State Diffusion Approximations

Anton Braverman
and
Jim Dai

School of Operations Research and Information Engineering
Cornell University

Applied Probability Frontiers, BIRS, June 4, 2015

Visionary Peter

Visionary Peter

- Numerical methods for Markov chains

Visionary Peter

- Numerical methods for Markov chains
- Impact on professional societies
  - INFORMS

Visionary Peter

- Numerical methods for Markov chains
- Impact on professional societies
  - INFORMS
  - APS

Visionary Peter

- Numerical methods for Markov chains
- Impact on professional societies
  - INFORMS
  - APS
  - IMS

Visionary Peter

- Numerical methods for Markov chains
- Impact on professional societies
  - INFORMS
  - APS
  - IMS
- Impact through editorship:
  - MOR

Visionary Peter

- Numerical methods for Markov chains
- Impact on professional societies
  - INFORMS
  - APS
  - IMS
- Impact through editorship:
  - MOR
  - SSy

Visionary Peter

- Numerical methods for Markov chains
- Impact on professional societies
  - INFORMS
  - APS
  - IMS
- Impact through editorship:
  - MOR
  - SSy
  - JAP

Outline

Convergence Rate in Diffusion approximations for Many-Server Queues

- Exponential service time distribution
- Phase type service time distribution
  Multi-dimensional piecewise Ornstein–Uhlenbeck (OU) processes.

Current status:

- There is a huge literature on stochastic process convergence.
- However, there is little work on rate of convergence for steady-state approximations.

$M/M/n+M$ queue

- Arrival rate λ.
- Service rate μ.
- Mean patience time $1/\alpha < \infty$.
- Steady-state number of customers in system $X(\infty)$.

$$\tilde{X}(\infty) = \frac{X(\infty) - x(\infty)}{\sqrt{\lambda}}.$$

Theorem 1 (a)

$\exists\ C_K = C_K(\alpha, \mu) > 0$ such that

$$\sup_{x\in\mathbb{R}} \Big| \mathbb{P}\big\{\tilde{X}(\infty) \le x\big\} - \mathbb{P}\big\{Y(\infty) \le x\big\} \Big| \le C_K \frac{1}{\sqrt{\lambda}} \quad \text{for all } \lambda \ge 1 \text{ and } n \ge 1,$$

$Y(\infty)$ has density $\kappa \exp\Big(- \int_0^x b(y)dy \Big)$ for some piecewise linear $b(\cdot)$.

Equilibrium system size $x(\infty)$ and drift $b(\cdot)$

- Flow balance equation

$$\lambda = \big(x(\infty) \wedge n\big)\mu + \big(x(\infty) - n\big)^+ \alpha,$$

- Solution

$$x(\infty) = \begin{cases} n + \frac{\lambda - n}{\alpha} & \text{if } \lambda > n\mu, \\ \lambda/\mu & \text{if } \lambda \leq n\mu \end{cases}$$

- Drift

$$b(x) = \big[(x+\zeta)^- - \zeta^-\big]\mu - \big[(x+\zeta)^+ - \zeta^+\big]\alpha \quad \text{for } x \in \mathbb{R}, \tag{1}$$

  where

$$\zeta = \frac{x(\infty) - n}{\sqrt{\lambda}}.$$

- Theorem 1 (a) says that

$$d_K(\tilde{X}(\infty), Y(\infty)) \leq C_K \frac{1}{\sqrt{\lambda}} \quad \text{for all } \lambda \geq 1 \text{ and } n \geq 1,$$

  where $d_K(\cdot,\cdot)$ denotes the Kolmogorov distance.
- *Not* a limit theorem

Piece-Wise Ornstein-Uhlenbeck (OU) Process

A piece-wise OU process in $\mathbb{R}$ is a diffusion process satisfying

$$Y(t) = Y(0) + \sigma B(t) + \theta t + \alpha_1 \int_0^t Y(s)^- ds - \alpha_2 \int_0^t Y(s)^+ ds.$$

- $B = \{B(t), t \geq 0\}$ is the one-dimensional standard Brownian motion.
- When $\alpha_1 = \alpha_2 = \alpha$, Y becomes a $(\sigma^2, \theta, \alpha)$-OU process whose stationary distribution is normal
$$N\Big(\theta/\alpha, \sigma^2/(2\alpha)\Big).$$
- The generator for an OU process is

$$Gf(x) = \frac{1}{2}\sigma^2 f''(x) + (\theta - \alpha x) f'(x) \quad \text{for } f \in C^2(\mathbb{R}).$$

Universal approximation

- Overloaded: $\lambda >> n\mu$, $\zeta > 0$

$$Y(\infty) \sim f(x) \sim \begin{cases} N(-\zeta, 1/(2\mu)) & \text{if } x < -\zeta, \\ N(0, 1/(2\alpha)) & \text{if } x \geq -\zeta. \end{cases}$$

- Underloaded: $\lambda << n\mu$, $\zeta < 0$

$$Y(\infty) \sim f(x) \sim \begin{cases} N(0, 1/(2\mu)) & \text{if } x < |\zeta|, \\ N\big((1-\mu/\alpha)\,|\zeta|\,, 1/(2\alpha)\big) & \text{if } x \geq -\zeta. \end{cases}$$

- Gurvich, Mandelbaum, Huang (2014), *Mathematics of Operations Research*
- Glynn, Ward (2003), *Queueing Systems*

Other Versions of Theorem 1

Theorem 1 (b) (Wasserstein metric)

$\exists\ C_W = C_W(\alpha, \mu) > 0$ such that

$$\sup_{h \in \mathrm{Lip}(1)} \left| \mathbb{E}h(\tilde{X}(\infty)) - \mathbb{E}h(Y(\infty)) \right| \le C_W \delta \quad \text{for all } \lambda \ge 1 \text{ and } n \ge 1,$$

$$\mathrm{Lip}(1) = \{h : \mathbb{R} \to \mathbb{R}, |h(x) - h(y)| \le |x - y|\}.$$

Theorem 1 (c)

$\exists\ C_P = C_P(m, \alpha, \mu) > 0$ such that

$$\sup_{h \in \mathcal{P}_m} \left| \mathbb{E}h(\tilde{X}(\infty)) - \mathbb{E}h(Y(\infty)) \right| \le C_P \frac{1}{\sqrt{\lambda}} \quad \text{for all } \lambda \ge 1 \text{ and } n \ge 1,$$

where

$$\mathcal{P}_m = \{h \in C_1(\mathbb{R}) : |h(x)| \le |x|^m \text{ for all } x \in \mathbb{R}\}.$$

Phase-type service time distributions

An $M/Ph/n+M$ System

- Phase-type service times with mean service time $1/\mu$.
- For example, an H_2 (hyper-exponential) random variable S has the following representation:

$$S = \begin{cases} S_1 & \text{with probability } p_1, \\ S_2 & \text{with probability } p_2, \end{cases} \quad \text{and} \quad \begin{array}{l} p_1 + p_2 = 1, \\ S_i \sim \text{Exponential}(\nu_i), \end{array}$$

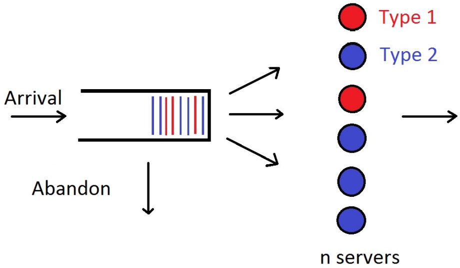

Phase-Type Random Variables

Definition (Neuts 1981)

A phase-type random variable corresponds to the hitting time of a continuous time Markov chain (CTMC) to an absorbing state. Inputs: (p, P, ν).

- For example, an H_2 (hyper-exponential) random variable S has the following representation:

$$S = \begin{cases} S_1 & \text{with probability } p_1, \\ S_2 & \text{with probability } p_2, \end{cases} \quad \text{and} \quad \begin{array}{l} p_1 + p_2 = 1, \\ S_i \sim \text{ Exponential}(\nu_i), \end{array}$$

$$\text{mean service time } = \frac{1}{\mu} = p_1 \frac{1}{\nu_1} + p_2 \frac{1}{\nu_2},$$

$$p = \begin{pmatrix} p_1 \\ p_2 \end{pmatrix}, \quad P = \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}, \quad \nu = \begin{pmatrix} \nu_1 \\ \nu_2 \end{pmatrix}.$$

An $M/Ph/n+M$ System (cont.)

- Let $X_1^n(t)$ be the number of type 1 customers in system at time t.
- Let $X_2^n(t)$ be the number of type 2 customers in system at time t.
- Denote

$$X^n(\infty) = \left(X_1^n(\infty), X_2^n(\infty)\right)$$

  to be the random vector having the stationary distribution.
- The computation of the distribution of $X^n(\infty)$ can be expensive or unrealistic.

An $M/H_2/500 + M$ System

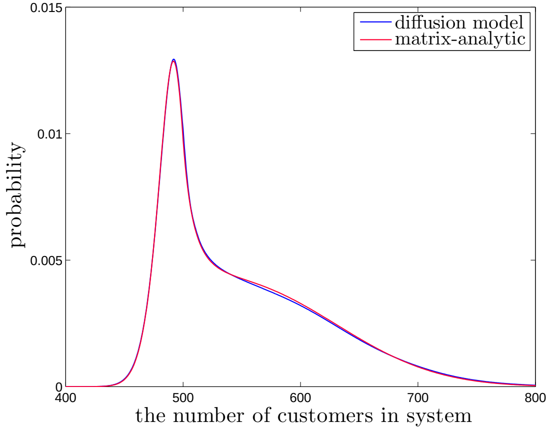

Figure : $\lambda = 522$, $p = (0.9351, 0.0649)$, $1/\nu = (0.1069, 13.89)$, mean patience time $= 2$.

Fig. 1 of D-He (2013), *Stochastic Systems*

Fig. 2 of Dai-He (2013): an $M/H_2/20 + M$ System

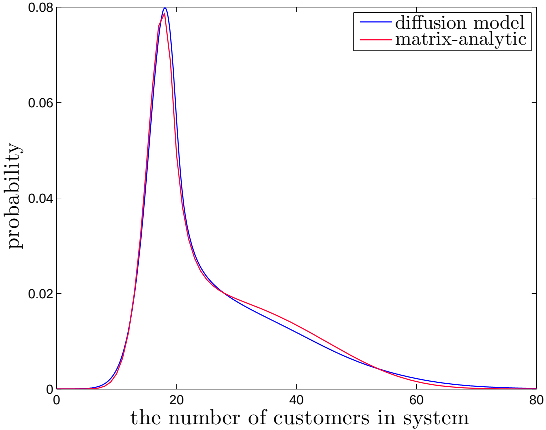

Figure : $\lambda = 22.2$, $p = (0.9351, 0.0649)$, $1/\nu = (0.1069, 13.89)$, mean patience time $= 2$.

$M/Ph/n/+M$ Results

- Let $\beta \in \mathbb{R}$ be fixed.
- Assume the number of server n follows the square-root safety staffing rule:

$$n\mu = \lambda_n + \beta\sqrt{\lambda_n}, \tag{2}$$

where λ_n is the arrival rate.
- The sequence of systems is in the Quality- and Efficiency-Driven (QED) regime, also known as the Halfin-Whitt (1981) regime.

Theorem 2 (a)

$\exists\ C = C(\alpha, \beta, p, P, \nu)$ such that

$$\sup_{h \in \mathrm{Lip}(1)} \left| \mathbb{E}h(\tilde{X}^n(\infty)) - \mathbb{E}h(Y(\infty)) \right| \leq C\lambda_n^{-1/4}, \quad \forall n \geq 1,$$

where

$$\tilde{X}^n(t) = \frac{1}{\sqrt{\lambda_n}}\Big(X^n(t) - \gamma n\Big), \qquad \gamma_i = \frac{p_i/\nu_i}{p_1/\nu_1 + p_2/\nu_2}.$$

$M/Ph/n+M$ results (cont.)

Theorem 2 (b)

For each integer $m > 0$, $\exists\ C = C(m, \alpha, \beta, p, P, \nu)$ such that such that if $h(x) : \mathbb{R}^d \longrightarrow \mathbb{R}$ satisfies

$$|h(x)| \le |x|^m, \quad \text{for } x \in \mathbb{R}^d,$$

then

$$\left|\mathbb{E}h(\tilde{X}^n(\infty)) - \mathbb{E}h(Y(\infty))\right| \le C_m \lambda_n^{-1/4}, \quad \forall n \ge 1.$$

- $Y(\infty)$ is the stationary distribution of a d-dimensional piecewise OU process

$$Y(t) = Y(0) + \sqrt{\Sigma}B(t) + \int_0^t b(Y(s))ds,$$

  Dieker and Gao (2013), *Annals of Applied Probability*
- The drift vector

$$b(x) = -\beta p - R(x - p(e'x)) - \alpha p(e'x)^+, \quad x \in \mathbb{R}^d.$$

- $e' = (1, \dots, 1)$ and $R = (I - P')\text{diag}(\nu)$.

$M/Ph/n+M$ results (cont.)

- The optimal rate should be $1/\sqrt{\lambda}$, not $1/\lambda^{1/4}$.
- The approximation should be universal: $M/H_2/100+M$

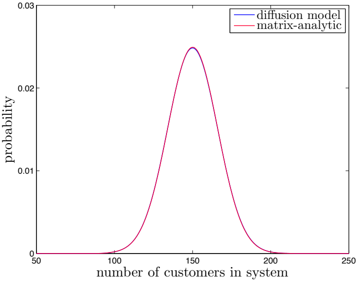

$\lambda = 150$

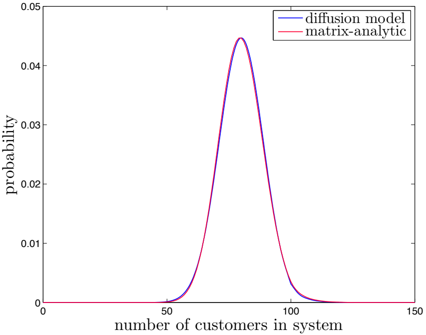

$\lambda = 80$

Stein Framework for Proofs

- Poisson (Stein) equation and gradient bounds
- Basic Adjoint Relationship (BAR) and the generator coupling
- State space collapse (SSC)
- Moment bounds

- Stein
- Louis Chen
- Andrew Barbour
- Chen, Goldstein, and Shao (2011)

Gurvich (2014), *Annals of Applied Probability*

Blanchet (2006)

An Outline of Proof for Theorem 1

Poisson Equation

- Generator of a piecewise OU process $Y = \{Y(t), t \geq 0\}$

$$G_Y f(x) = f''(x) + b(x) f'(x)$$

- Given an $h \in \mathrm{Lip}(1)$, solve $f = f_h$ from the Poisson equation

$$G_Y f(x) = h(x) - \mathbb{E}[h(Y(\infty))], \quad x \in \mathbb{R}.$$

- On each sample path of $\tilde{X}(\infty)$

$$G_Y f(\tilde{X}(\infty)) = h(\tilde{X}(\infty)) - \mathbb{E}[h(Y(\infty))].$$

- Key identify

$$\mathbb{E}[h(\tilde{X}(\infty))] - \mathbb{E}[h(Y(\infty))] = \mathbb{E}[G_Y f(\tilde{X}(\infty))]$$

Gradient Bounds

Assume $\mu = 1$. For any $h \in \mathrm{Lip}(1)$,

$$\|f_h'\| \le C_1(\alpha)\|h'\|, \quad \|f_h''\| \le C_2(\alpha)\|h'\|, \quad \text{and} \quad \|f_h'''\| \le C_3(\alpha)\|h'\|.$$

The constants $C_1(\alpha)$, $C_2(\alpha)$ and $C_3(\alpha)$ are given by

$$\begin{aligned} C_1(\alpha) &= \left(1 + 2\left(\frac{1}{\alpha} \vee 1\right) + \left(\frac{\alpha \vee 1}{\alpha \wedge 1} + 2\right)\right) / (1 \vee \alpha), \\ C_2(\alpha) &= \left(1 + 2\left(\frac{1}{\alpha} \vee 1\right) + \left(\frac{\alpha \vee 1}{\alpha \wedge 1} + 2\right)\right)\left(\sqrt{\frac{\pi}{2}} + \sqrt{\frac{\pi}{2\alpha}}\right), \\ C_3(\alpha) &= 3\left(1 + 2\left(\frac{1}{\alpha} \vee 1\right) + \left(\frac{\alpha \vee 1}{\alpha \wedge 1} + 2\right)\right). \end{aligned}$$

Drift

$$b(x) = \left[(x+\zeta)^- - \zeta^-\right]\mu - \left[(x+\zeta)^+ - \zeta^+\right]\alpha \quad \text{for } x \in \mathbb{R},$$

¡where

$$\zeta = \frac{x(\infty) - n}{\sqrt{\lambda}}.$$

Basic Adjoint Relationship (BAR)

Lemma

Random vector $X(\infty) \in S$ *has the stationary distribution of a Markov process* $X = \{X(t), t \geq 0\}$ *if and only if the following basic adjoint relationship (BAR) holds:*

$$\mathbb{E}\big[G_X f(X(\infty)\big] = 0 \quad \text{for all "good" } f : S \to \mathbb{R}. \tag{3}$$

- Echeverria (1982): Markov processes without boundary.
- Weiss (1981): Markov processes with boundaries.
- Harrison and Williams (1987), D-Kurtz (1994), semimartingale reflecting Brownian motions (SRBMs).
- Kang-Ramanan (2014), reflecting diffusion
- Glynn and Zeevi (2008, *Kurtz Festschrift*) provides sufficient conditions on f for (3) to hold for Markov chains.

Generator Coupling

- From the Stein equation,

$$\begin{aligned}\mathbb{E}\big[h(\tilde{X}(\infty))\big] - \mathbb{E}\big[h(Y(\infty))\big] &= \mathbb{E}\big[G_Y f_h(\tilde{X}(\infty))\big] \\ &= \mathbb{E}\big[G_Y f_h(\tilde{X}(\infty))\big] - \mathbb{E}\big[G_{\tilde{X}} f_h(\tilde{X}(\infty))\big] \\ &= \mathbb{E}\Big[G_Y f_h(\tilde{X}(\infty)) - G_{\tilde{X}} f_h(\tilde{X}(\infty))\Big].\end{aligned}$$

- Doing Taylor expansion on $G_{\tilde{X}} f_h(x)$ to bound

$$\big|G_Y f_h(x) - G_{\tilde{X}} f_h(x)\big| \quad \text{for } x = \delta(i - x(\infty)) \text{ with } i \in \mathbb{Z}_+,$$

where

$$\delta = \frac{1}{\sqrt{\lambda}}.$$

Pre-limit generator

- Set $x = \delta(i - x(\infty))$ for $i \in \mathbb{Z}_+$. The generator of birth-death process $\tilde{X}$ is

$$\begin{aligned} G_{\tilde{X}} f(x) &= \lambda\Big(f(x+\delta) - f(x)\Big) \\ &\quad + \big(\mu(i \wedge n) + \alpha(i-n)^+\big)\Big(f(x-\delta) - f(x)\Big). \end{aligned}$$

- The generator of Y is

$$G_Y f(x) = b(x) f'(x) + f''(x).$$

- Note that

$$\begin{aligned} \lambda - n\mu &= \big(x(\infty) - n\big)^+ \alpha - \big(x(\infty) - n\big)^- \mu, \\ (i \wedge n)\mu + (i-n)^+\alpha &= \big(n - (i-n)^-\big)\mu + (i-n)^+\alpha \\ &= n\mu + (i-n)^+\alpha - (i-n)^-\mu. \end{aligned}$$

- Therefore,

$$\big(\mu(i \wedge n) + \alpha(i-n)^+\big) = \lambda - b(x)/\delta.$$

Taylor Expansion

- Using Taylor expansion, we can write

$$\begin{aligned}
& \left| G_{\tilde{X}} f_h(x) - G_Y f_h(x) \right| \\
= \; & \left| -\frac{1}{2}\delta b(x) f_h''(x) + \frac{\delta^2}{2}\lambda(f_h''(\xi) - f_h''(x)) + \frac{\delta}{2}(\delta\lambda - b(x))(f_h''(\eta) - f_h''(x)) \right| \\
\le \; & \frac{\delta}{2}\Big[\, |b(x)| \, |f_h''(x)| + |f'''(\xi_1)| + (1 + \delta\,|b(x)|)\,|f_h'''(\eta_1)| \, \Big] \\
\le \; & \frac{\delta}{2}\Big[2\|f_h'''\| + \big(\|f_h''\| + \delta\|f_h''''\|\big)|b(x)| \Big]
\end{aligned}$$

where $\xi \in [x, x+\delta]$, $\eta \in [x-\delta, x]$, $\xi_1 \in [x, \xi]$, $\eta_1 \in [\eta, x]$.

- Using Lipschitz continuity,

$$|b(x)| \le (\alpha \vee \mu)\,|x|.$$

Moment Bound

- Using quadratic Lyapunov function $V(x) = x^2$, one can prove moment bound

$$B \equiv \sup_n \mathbb{E}\left|\tilde{X}^n(\infty)\right| < \infty.$$

- Therefore, for any Lipschitz continuous h, one has

$$\left|\mathbb{E}h(\tilde{X}^n(\infty)) - \mathbb{E}h(Y(\infty))\right| = \left|\mathbb{E}G_X f_h(\tilde{X}^n(\infty)) - \mathbb{E}G_Y f_h(\tilde{X}^n(\infty))\right|$$
$$\leq \frac{\delta}{2}\Big[2\|f_h'''\| + \big(\|f_h''\| + \delta\|f_h''''\|\big)(\alpha \vee \mu) B\Big].$$

An Outline of the Proof for Theorem 2

Poisson equation

- Let

$$G_Y f(x) = \frac{1}{2} \sum_{i,j=1}^{d} \Sigma_{ij} \partial_{ij} f(x) + \sum_{i=1}^{d} \partial_i f(x) b_i(x) \quad \text{for } x \in \mathbb{R}^d,$$

  be the generator of the piecewise OU process $Y = \{Y(t) \in \mathbb{R}^d, t \geq 0\}$.
- For $h : \mathbb{R}^d \to \mathbb{R}$, find a solution f_h to the Poisson equation

$$G_Y f_h(x) = h(x) - \mathbb{E}h(Y(\infty)). \tag{4}$$

- Then, the following key identity holds

$$\mathbb{E}\big[h(\tilde{X}^n(\infty))\big] - \mathbb{E}\big[h(Y(\infty))\big] = \mathbb{E}\big[G_Y f_h(\tilde{X}^n(\infty))\big]. \tag{5}$$

Multi-dimensional Gradient Bounds

Lemma (Gurvich (2014))

Suppose $|h(x)| \leq |x|^{2m}$ for some $m > 0$, then the solution to Poisson equation

$$G_Y f_h(x) = h(x) - \mathbb{E}h(Y(\infty))$$

satisfies

$$|f(x)| \leq C_m(1 + |x|^2)^m,$$

$$|Df(x)| \leq C_m(1 + |x|^2)^m(1 + |x|),$$

$$\left|D^2 f(x)\right| \leq C_m(1 + |x|^2)^m(1 + |x|)^2,$$

$$\sup_{|y-x|<1, y\neq x} \frac{\left|D^2 f(x) - D^2 f(y)\right|}{|x - y|} \leq C_m(1 + |x|^2)^m(1 + |x|)^3.$$

A Markov Chain Representation

The system can be modeled as a continuous time Markov chain (CTMC)

$$U^n = \Big\{ U^n(t) \in \big(\{1, 2, \ldots, d\}\big)^{\infty}, t \geq 0 \Big\}.$$

- An example of the state u is given by

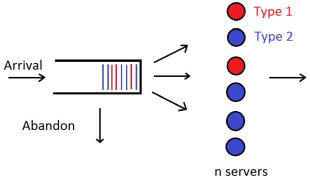

$$u = (1, 2, 1, 2, 2, 2, |, 2, 1, 2, 2, 1, 1, 2, 2)$$

- Because of customer abandonment, the CTMC U^n is positive recurrent.

Generator Coupling

- Recall that $f_h : \mathbb{R}^2 \to \mathbb{R}$ is a solution to the Poisson equation.
- $\mathbb{E}\big[G_{U^n} f_h(\tilde{X}^n(\infty))\big]$ is not well defined.
- With a general phase-type service distribution, system size process $\{(X_1^n(t), X_2^n(t)), t \geq 0\}$ is no longer a CTMC.
- U^n is a CTMC living on state space $\mathcal{U} = \{1, 2\}^\infty$. Its generator G_{U^n} acts on functions $F : \mathcal{U} \to \mathbb{R}$.
- BAR gives $\mathbb{E}\big[G_{U^n} A f_h(U^n(\infty))\big] = 0$.

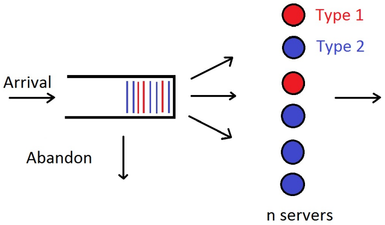

Applying Gradient Bounds

- Recall $\delta = \frac{1}{\sqrt{\lambda_n}}$.
- For $u = (1, 2, 1, 2, 2, 2, |, 2, 1, 2, 2, 1, 1, 2, 2)$, $(z_1, z_2, q_1, q_2) = (2, 4, 3, 5)$ and

$$\begin{aligned} G_U A f_h(u) = {} & \lambda p_1 f_h(x_1 + \delta, x_2) + \lambda p_2 f_h(x_1, x_2 + \delta) \\ & + \alpha q_1 f_h(x_1 - \delta, x_2) + \alpha q_2 f_h(x_1, x_2 - \delta) \\ & + z_1 \nu_1 f_h(x_1, x_2) + z_1 \nu_1 f_h(x_1, x_2 - \delta) \\ & - (\lambda + \alpha q + z_1 \nu_1 + z_2 \nu_2) f_h(x_1, x_2). \end{aligned}$$

Lemma

There exists a constant $C(m) = C(m, \beta, \alpha, p, \nu, P)$ such that for any $u \in \mathcal{U}$,

$$\begin{aligned} |G_U A f_h(u) - G_Y f_h(x)| \quad \leq \quad & C(m)(1 + |x|^2)^m (1 + |x|) \left|\delta q - p(e^T x)^+\right| \\ & + \delta\, C(m)(1 + |x|^2)^m (1 + |x|)^4. \end{aligned}$$

State Space Collapse in $M/Ph/n+M$ Case

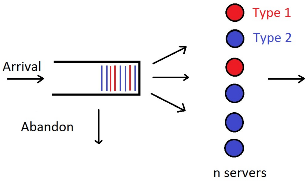

- The total queue size $\left(X_1^n(\infty) + X_2^n(\infty) - n\right)^+ = (e'X^n(\infty) - n)^+$.

Lemma (State-Space Collapse)

There exists $C(m) > 0$ such that $\forall n \geq 1$,

$$\mathbb{E}\left|\delta\left(Q_i^n(\infty) - p_i(e'X^n(\infty) - n)^+\right)\right|^{2m} \leq C(m)\delta^m \mathbb{E}[(e'\tilde{X}^n(\infty))^+]^m \quad \text{for } i = 1, 2.$$

Current Method: Limit Interchange

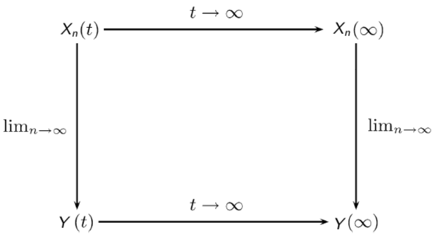

- Prove process convergence $X^n(\cdot) \Rightarrow Y(\cdot)$ see, for example, Reiman (1984), Peterson (1991), Bramson (1998), and Williams (1998).
- Process convergence does not imply $X^n(\infty) \Rightarrow Y(\infty)$.
- Need to justify the limit interchange

$$\lim_{n\to\infty} \lim_{t\to\infty} X^n(t) = \lim_{t\to\infty} \lim_{n\to\infty} X^n(t).$$

Limit Interchange Justifications

- Networks of single-server queues
  - Gamarnik & Zeevi (2006)
  - Budhiraja & Lee (2009)
  - Zhang & Zwart (2008)
  - Katsuda (2010, 2011)
  - Yao & Ye (2012)
  - Gurvich (MOR, 2014)
- Many-server systems
  - Tezcan (2008)
  - Gamarnik & Stolyar (2012)
  - D., Dieker & Gao (2014)
- No convergence rates

Summary

Use Stein Framework

for justifying steady-state approximations