Drifting Towards Progress in **Multiserver Scheduling**

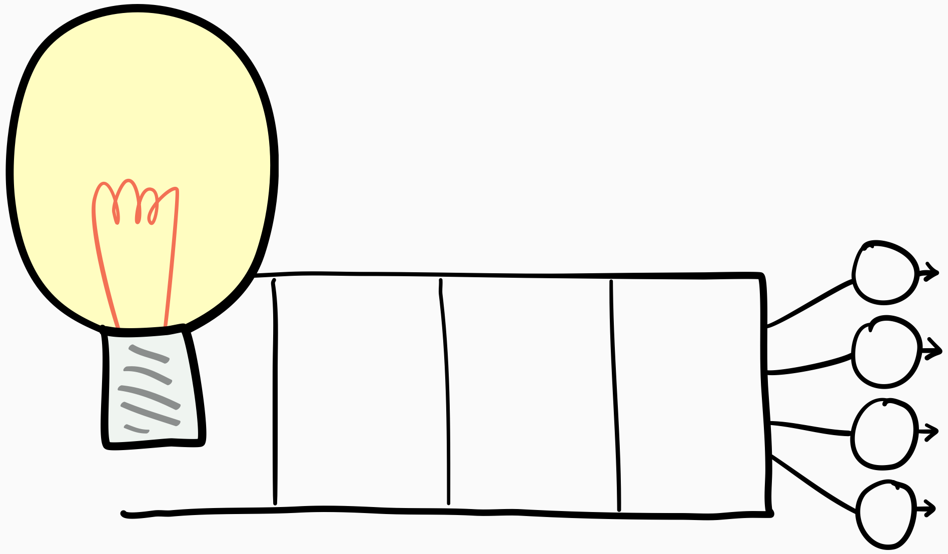

Ziv Scully

CMU (now) → UC Berkeley → MIT/Harvard → Cornell (Fall 2023)

Multiserver systems are ubiquitous but hard to analyze

Multiserver systems are ubiquitous but hard to analyze

- E.g. FCFS M/G/k hard to approximate [GHDZ'10]

Multiserver systems are ubiquitous but hard to analyze

- E.g. FCFS M/G/k hard to approximate [GHDZ'10]
- Even harder: complex *architectures*

Multiserver systems are ubiquitous but hard to analyze

- E.g. FCFS M/G/k hard to approximate [GHDZ'10]
- Even harder: complex *architectures*
- Even harder: complex *scheduling*

Multiserver systems are ubiquitous but hard to analyze

- E.g. FCFS M/G/*k* hard to approximate [GHDZ'10]
- Even harder: complex *architectures*
- Even harder: complex *scheduling*

I think there's hope!

Multiserver systems are ubiquitous but hard to analyze

- E.g. FCFS M/G/*k* hard to approximate [GHDZ'10]
- Even harder: complex *architectures*
- Even harder: complex *scheduling*

I think there's hope!

- Progress on complex *architectures*:
  dispatching [ES'12, BBLM'16, Y'17], generalized switches [HM'20], networks [BDM'17], redundancy [GHS'17], fork-join [WHJSS'19]

Multiserver systems are ubiquitous but hard to analyze

- E.g. FCFS M/G/*k* hard to approximate [GHDZ'10]
- Even harder: complex *architectures*
- Even harder: complex *scheduling*

I think there's hope!

- Progress on complex *architectures*:
  dispatching [ES'12, BBLM'16, Y'17], generalized switches [HM'20], networks [BDM'17], redundancy [GHS'17], fork-join [WHJSS'19]
- Progress on complex *scheduling*:
  SRPT [GSH'18, GSH'19], Gittins policy [SGH'20]

Techniques I won’t do justice to

Techniques I won't do justice to

- Product-form Markov chains
  - Redundancy models and more [GR'20]
  - **Limitation:** what about general job sizes?

Techniques I won't do justice to

- Product-form Markov chains
  - Redundancy models and more [GR'20]
  - **Limitation:** what about general job sizes?
- Process convergence and limit interchange
  - JSQ and power-of-*d* dispatching [MBLW'16] and more
  - Dispatching to SRPT servers [DW'06]
  - **Limitation:** what do we learn about pre-limit system?

Techniques I won't do justice to

- Product-form Markov chains
  - Redundancy models and more [GR'20]
  - **Limitation:** what about general job sizes?
- Process convergence and limit interchange
  - JSQ and power-of-*d* dispatching [MBLW'16] and more
  - Dispatching to SRPT servers [DW'06]
  - **Limitation:** what do we learn about pre-limit system?
- Many techniques for G/G/*k*
  - Keifer-Wolfowitz vector, random walks, and more [LG'17]
  - **Limitation:** what about non-FCFS scheduling?

Today's focus:
drift methods

Today's focus: *drift methods*

1. What are drift methods?

Today's focus: *drift methods*

1. What are drift methods?

2. How do we adapt drift methods to handle complex scheduling?

Today's focus: *drift methods*

1. What are drift methods?

2. How do we adapt drift methods to handle complex scheduling?

3. What's next for drift methods?

What is drift?

Let X_t be a Markov process with

- stationary distribution X_∞
- *generator* D_X:

What is drift?

Let X_t be a Markov process with

- stationary distribution X_∞
- *generator* D_X:

$$D_X f(x) = \lim_{\delta \to 0} \frac{\mathbf{E}[f(X_\delta) \mid X_0 = 0] - f(x)}{\delta}$$

What is drift?

Let X_t be a Markov process with

- stationary distribution X_∞
- *generator* D_X:

$$D_X f(x) = \lim_{\delta \to 0} \frac{\mathbf{E}[f(X_\delta) \mid X_0 = 0] - f(x)}{\delta}$$

D_X operator gives "expected drift" of value over time

What are drift methods?

Compositions of “ingredients” involving generator D_X

What are drift methods?

Compositions of “ingredients” involving generator D_X

Example: JSQ with M/M arrivals

What are drift methods?

Compositions of “ingredients” involving generator D_X

Example: JSQ with M/M arrivals

N = number in system

$$\mathbf{E}[N_{\text{JSQ}}] \leq \mathbf{E}[N_{\text{M/M/1}}] + ???$$

What are drift methods?

Compositions of “ingredients” involving generator D_X

Example: JSQ with M/M arrivals

N = number in system

$$\mathbf{E}[N_{\text{JSQ}}] \le \mathbf{E}[N_{\text{M/M/1}}] + ???$$

$$??? \le o\left(\frac{1}{1-\rho}\right)$$

ρ = load

What are drift methods?

Compositions of "ingredients" involving generator D_X

Example: JSQ with M/M arrivals

N = number in system

Basic Adjoint Relationship ⇒

$$\mathbf{E}[N_{\mathrm{JSQ}}] \le \mathbf{E}[N_{\mathrm{M/M/1}}] + ???$$

$$??? \le o\left(\frac{1}{1-\rho}\right)$$

ρ = load

What are drift methods?

Compositions of “ingredients” involving generator D_X

stationary drift is zero

Basic Adjoint Relationship ⇨

Example: JSQ with M/M arrivals

N = number in system

$$\mathbf{E}[N_{\text{JSQ}}] \leq \mathbf{E}[N_{\text{M/M/1}}] + ???$$

$$??? \leq o\left(\frac{1}{1-\rho}\right)$$

ρ = load

What are drift methods?

Compositions of “ingredients” involving generator D_X

Example: JSQ with M/M arrivals

stationary drift is zero

N = number in system

Basic Adjoint Relationship

$$\mathbf{E}[N_{\text{JSQ}}] \leq \mathbf{E}[N_{\text{M/M/1}}] + ???$$

State Space Collapse

$$??? \leq o\left(\frac{1}{1-\rho}\right)$$

ρ = load

What are drift methods?

Compositions of "ingredients" involving generator D_X

Example: JSQ with M/M arrivals

stationary drift is zero

Basic Adjoint Relationship

N = number in system

$$\mathbf{E}[N_{\text{JSQ}}] \le \mathbf{E}[N_{\text{M/M/1}}] + ???$$

State Space Collapse

$$??? \le o\left(\frac{1}{1-\rho}\right)$$

ρ = load

show state drifts to a subspace

Drift method ingredients

Basic Adjoint Relationship (BAR)
stationary drift is zero

State Space Collapse (SSC)
show state drifts to a subspace

Drift method ingredients

Basic Adjoint Relationship (BAR)
stationary drift is zero

State Space Collapse (SSC)
show state drifts to a subspace

$D_X f$ for which f?

Drift method ingredients

Basic Adjoint Relationship (BAR)
stationary drift is zero

State Space Collapse (SSC)
show state drifts to a subspace

$D_X f$ for which f?

Poisson EQuation (PEQ)
find a function with useful drift

Drift method ingredients

Basic Adjoint Relationship (BAR)
stationary drift is zero

State Space Collapse (SSC)
show state drifts to a subspace

$D_X f$ for which f?

Poisson EQuation (PEQ)
find a function with useful drift

Metric eXtraction (MX)
compute metric we care about

Drift method ingredients

Basic Adjoint Relationship (BAR)
stationary drift is zero

State Space Collapse (SSC)
show state drifts to a subspace

$D_X f$ for which f?

Poisson EQuation (PEQ)
find a function with useful drift

Metric eXtraction (MX)
compute metric we care about

crucial for scheduling

Drift method recipes

Ingredients:

- **Basic Adjoint Relationship (BAR)**
- **State Space Collapse (SSC)**
- **Poisson EQuation (PEQ)**
- **Metric eXtraction (MX)**

Drift method recipes

Ingredients:

- **Basic Adjoint Relationship (BAR)**
- **State Space Collapse (SSC)**
- **Poisson EQuation (PEQ)**
- **Metric eXtraction (MX)**

BAR approach [BDM'17]: **BAR**

Drift method recipes

Ingredients:

- **Basic Adjoint Relationship (BAR)**
- **State Space Collapse (SSC)**
- **Poisson EQuation (PEQ)**
- **Metric eXtraction (MX)**

BAR approach [BDM'17]: **BAR**

Drift method [ES'12, MS'16]: **BAR** + **SSC**

Drift method recipes

Ingredients:

- **Basic Adjoint Relationship (BAR)**
- **State Space Collapse (SSC)**
- **Poisson EQuation (PEQ)**
- **Metric eXtraction (MX)**

BAR approach [BDM'17]: **BAR**

Drift method [ES'12, MS'16]: **BAR** + **SSC**

Stein framework [BDF'16]: **BAR** + **PEQ** + **MX**

Drift method recipes

Ingredients:

- **Basic Adjoint Relationship (BAR)**
- **State Space Collapse (SSC)**
- **Poisson EQuation (PEQ)**
- **Metric eXtraction (MX)**

BAR approach [BDM'17]: **BAR**

Drift method [ES'12, MS'16]: **BAR** + **SSC**

Stein framework [BDF'16]: **BAR** + **PEQ** + **MX**

SRPT and Gittins in M/G/k [SGH'20]: **BAR** + **MX**

Drift method recipes

Ingredients:

- **Basic Adjoint Relationship (BAR)**
- **State Space Collapse (SSC)**
- **Poisson EQuation (PEQ)**
- **Metric eXtraction (MX)**

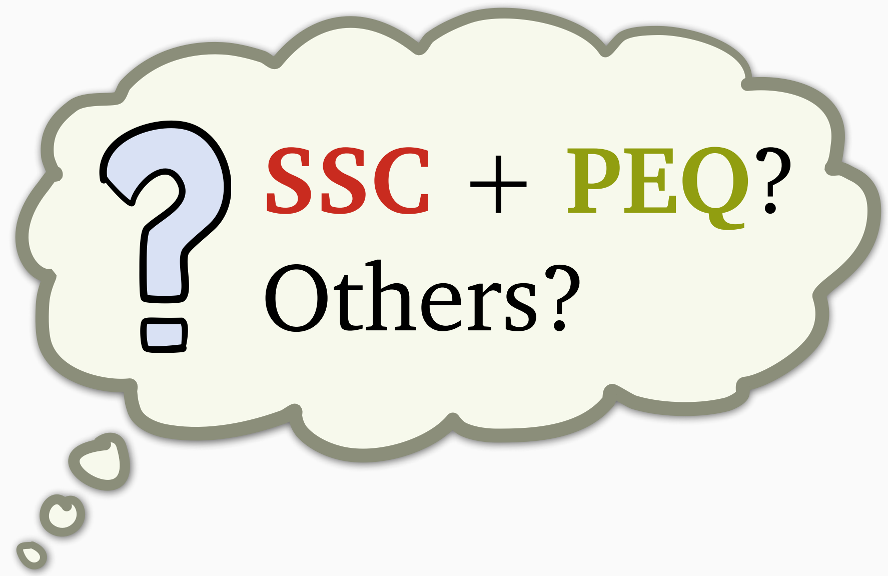

BAR approach [BDM'17]: **BAR**

Drift method [ES'12, MS'16]: **BAR** + **SSC**

Stein framework [BDF'16]: **BAR** + **PEQ** + **MX**

SRPT and Gittins in M/G/k [SGH'20]: **BAR** + **MX**

Drift method ingredients

Basic Adjoint Relationship (BAR)
stationary drift is zero

State Space Collapse (SSC)
show state drifts to a subspace

Poisson EQuation (PEQ)
find a function with useful drift

Metric eXtraction (MX)
compute metric we care about

Drift method ingredients

Today:
dispatching,
M/M arrivals

Basic Adjoint Relationship (BAR)
stationary drift is zero

State Space Collapse (SSC)
show state drifts to a subspace

Poisson EQuation (PEQ)
find a function with useful drift

Metric eXtraction (MX)
compute metric we care about

Basic Adjoint Relationship (BAR)

BAR: For all f that don't grow too fast,

$$\mathbf{E}[D_X f(X_\infty)] = 0$$

Basic Adjoint Relationship (BAR)

BAR: For all f that don't grow too fast,

$$\mathbf{E}[D_X f(X_\infty)] = 0$$

Example: M/M dispatching, x_i = # jobs at server i

Basic Adjoint Relationship (BAR)

BAR: For all f that don't grow too fast,

$$\mathbf{E}[D_X f(X_\infty)] = 0$$

Example: M/M dispatching, $x_i = \#$ jobs at server i

Let $f(x) = \#\text{ jobs} = \sum_{i=1}^{k} x_i$

Basic Adjoint Relationship (BAR)

BAR: For all f that don't grow too fast,

$$\mathbf{E}[D_X f(X_\infty)] = 0$$

Example: M/M dispatching, $x_i = \#$ jobs at server i

Let $f(x) = \#\text{ jobs} = \sum_{i=1}^{k} x_i \Rightarrow \mathbf{E}[I] = k(1-\rho)$

idle servers

Basic Adjoint Relationship (BAR)

BAR: For all f that don't grow too fast,

$$\mathbf{E}[D_X f(X_\infty)] = 0$$

Example: M/M dispatching, $x_i = \#$ jobs at server i

\# idle servers

$$\text{Let } f(x) = \#\text{ jobs} = \sum_{i=1}^{k} x_i \quad \Rightarrow \quad \mathbf{E}[I] = k(1-\rho)$$

Let $f(x) = (\#\text{ jobs})^2$

Basic Adjoint Relationship (BAR)

BAR: For all f that don't grow too fast,

$$\mathbf{E}[D_X f(X_\infty)] = 0$$

Example: M/M dispatching, x_i = # jobs at server i

idle servers

$$\text{Let } f(x) = \#\text{ jobs} = \sum_{i=1}^{k} x_i \implies \mathbf{E}[I] = k(1-\rho)$$

$$\text{Let } f(x) = (\#\text{ jobs})^2 \implies \mathbf{E}[N] = \mathbf{E}[N_{\mathrm{M/M/1}}] + \frac{\mathbf{E}[IN]}{1-\rho}$$

State Space Collapse (SSC)

SSC: If there is "norm-like" f, constant c, and set A such that for all $x \notin A$,

$$D_X f(x) \leq -c f(x),$$

then distance from X_∞ to A has all finite moments

State Space Collapse (SSC)

SSC: If there is "norm-like" f, constant c, and set A such that for all $x \notin A$,

$$D_X f(x) \leq -cf(x),$$

then distance from X_∞ to A has all finite moments

Example: M/M dispatching, $x_i = \#$ jobs at server i

State Space Collapse (SSC)

SSC: If there is "norm-like" f, constant c, and set A such that for all $x \notin A$,

$$D_X f(x) \leq -c f(x),$$

then distance from X_∞ to A has all finite moments

Example: M/M dispatching, $x_i = \#$ jobs at server i

Let $f(x) = \|x\|_\perp$

State Space Collapse (SSC)

SSC: If there is “norm-like” f, constant c, and set A such that for all $x \notin A$,

$$D_X f(x) \leq -cf(x),$$

then distance from X_∞ to A has all finite moments

Example: M/M dispatching, $x_i = \#$ jobs at server i

Let $f(x) = \|x\|_\perp$ ⇒ Under JSQ dispatching, $\mathbf{E}[IN] = o\left((1-\rho)^\varepsilon\right)$

State Space Collapse (SSC)

SSC: If there is “norm-like” f, constant c, and set A such that for all $x \notin A$,

$$D_X f(x) \leq -cf(x),$$

then distance from X_∞ to A has all finite moments

Example: M/M dispatching, $x_i = \#$ jobs at server i

Let $f(x) = \|x\|_\perp$ ⇒ Under JSQ dispatching, $\mathbf{E}[IN] = o\left((1-\rho)^\varepsilon\right)$

In heavy traffic: $\mathbf{E}[N] = \mathbf{E}[N_{\mathrm{M/M/1}}] + \dfrac{\mathbf{E}[IN]}{1-\rho} \approx \mathbf{E}[N_{\mathrm{M/M/1}}]$

Poisson Equation (PEQ)

PEQ: Given holding cost h, solve for average cost $h_\infty = \mathbf{E}[h(X_\infty)]$ and potential function f:

$$D_X f(x) = h(x) - h_\infty$$

Poisson Equation (PEQ)

PEQ: Given holding cost h, solve for average cost $h_\infty = \mathbf{E}[h(X_\infty)]$ and potential function f:

$$D_X f(x) = h(x) - h_\infty$$

Why do we care? Because then

$$\mathbf{E}[D_X f(Y_\infty)] = \mathbf{E}[h(Y_\infty)] - \mathbf{E}[h(X_\infty)]$$

Poisson Equation (PEQ)

PEQ: Given holding cost h, solve for average cost $h_\infty = \mathbf{E}[h(X_\infty)]$ and potential function f:

$$D_X f(x) = h(x) - h_\infty$$

Why do we care? Because then

$$\mathbf{E}[D_X f(Y_\infty)] = \mathbf{E}[h(Y_\infty)] - \mathbf{E}[h(X_\infty)]$$
$$= \mathbf{E}[(D_X - D_Y) f(Y_\infty)]$$

BAR

Poisson Equation (PEQ)

PEQ: Given holding cost h, solve for average cost $h_\infty = \mathbf{E}[h(X_\infty)]$ and potential function f:

$$D_X f(x) = h(x) - h_\infty$$

Why do we care? Because then

$$\begin{aligned}\mathbf{E}[D_X f(Y_\infty)] &= \mathbf{E}[h(Y_\infty)] - \mathbf{E}[h(X_\infty)] \\ &= \mathbf{E}[(D_X - D_Y) f(Y_\infty)]\end{aligned}$$

BAR

simple

complicated

Poisson Equation (PEQ)

PEQ: Given holding cost h, solve for average cost $h_\infty = \mathbf{E}[h(X_\infty)]$ and potential function f:

$$D_X f(x) = h(x) - h_\infty$$

Why do we care? Because then

$$\mathbf{E}[D_X f(Y_\infty)] = \mathbf{E}[h(Y_\infty)] - \mathbf{E}[h(X_\infty)]$$
$$= \mathbf{E}[(D_X - D_Y)f(Y_\infty)]$$

BAR

simple

complicated

Example: M/M/1, $x = \#$ jobs

Let $h(x) = \#$ jobs $= x$

Poisson Equation (PEQ)

PEQ: Given holding cost h, solve for average cost $h_\infty = \mathbf{E}[h(X_\infty)]$ and potential function f:

$$D_X f(x) = h(x) - h_\infty$$

Why do we care? Because then

$$\mathbf{E}[D_X f(Y_\infty)] = \mathbf{E}[h(Y_\infty)] - \mathbf{E}[h(X_\infty)]$$

$$= \mathbf{E}[(D_X - D_Y) f(Y_\infty)]$$

BAR

simple

complicated

Example: M/M/1, $x = \#$ jobs

Let $h(x) = \#\text{ jobs} = x$ ⇒ Solution: $f(x) = (\#\text{ jobs})^2 = x^2$

Metric eXtraction (MX)

Can understand $\mathbf{E}[h(X_\infty)]$ for some cost functions h

or $\mathbf{E}[h(Y_\infty)] - \mathbf{E}[h(X_\infty)]$

Metric eXtraction (MX)

Can understand $\mathbf{E}[h(X_\infty)]$ for some cost functions h

or $\mathbf{E}[h(Y_\infty)] - \mathbf{E}[h(X_\infty)]$

Question: what functions h are useful?

Metric eXtraction (MX)

Can understand $\mathbf{E}[h(X_\infty)]$ for some cost functions h

or $\mathbf{E}[h(Y_\infty)] - \mathbf{E}[h(X_\infty)]$

Question: what functions h are useful?

- Directly useful: number of jobs, functions thereof

Metric eXtraction (MX)

Can understand $\mathbf{E}[h(X_\infty)]$ for some cost functions h

or $\mathbf{E}[h(Y_\infty)] - \mathbf{E}[h(X_\infty)]$

Question: what functions h are useful?

- Directly useful: number of jobs, functions thereof
- Bound for all indicators → total variation distance

Metric eXtraction (MX)

Can understand $\mathbf{E}[h(X_\infty)]$ for some cost functions h

or $\mathbf{E}[h(Y_\infty)] - \mathbf{E}[h(X_\infty)]$

Question: what functions h are useful?

- Directly useful: number of jobs, functions thereof
- Bound for all indicators → total variation distance
- Bound for all Lipschitz → Wasserstein distance

Metric eXtraction (MX)

Can understand $\mathbf{E}[h(X_\infty)]$ for some cost functions h

or $\mathbf{E}[h(Y_\infty)] - \mathbf{E}[h(X_\infty)]$

Question: what functions h are useful?

- Directly useful: number of jobs, functions thereof
- Bound for all indicators → total variation distance
- Bound for all Lipschitz → Wasserstein distance
- **Problem:** for M/G arrivals, number of jobs is hard to deal with

Metric eXtraction (MX)

Can understand $\mathbf{E}[h(X_\infty)]$ for some cost functions h

or $\mathbf{E}[h(Y_\infty)] - \mathbf{E}[h(X_\infty)]$

Question: what functions h are useful?

- Directly useful: number of jobs, functions thereof
- Bound for all indicators → total variation distance
- Bound for all Lipschitz → Wasserstein distance
- **Problem:** for M/G arrivals, number of jobs is hard to deal with
  - **Solution:** look at *work* instead

Metric eXtraction (MX)

Can understand $\mathbf{E}[h(X_\infty)]$ for some cost functions h

or $\mathbf{E}[h(Y_\infty)] - \mathbf{E}[h(X_\infty)]$

Question: what functions h are useful?

- Directly useful: number of jobs, functions thereof
- Bound for all indicators → total variation distance
- Bound for all Lipschitz →
- **Problem:** for M/G arriva rd to deal with
  - **Solution:** look at *work* instead

WINE: get # jobs from work [SGH'20]

Drift method ingredients

Basic Adjoint Relationship (BAR)
stationary drift is zero

State Space Collapse (SSC)
show state drifts to a subspace

Poisson EQuation (PEQ)
find a function with useful drift

Metric eXtraction (MX)
compute metric we care about

Drift method ingredients

Basic Adjoint Relationship (BAR)
stationary drift is zero

State Space Collapse (SSC)
show state drifts to a subspace

Poisson EQuation (PEQ)
find a function with useful drift

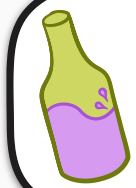

WINE
[SGH'20]

Metric eXtraction (MX)
compute metric we care about

References: drift methods

[BDF'16] BAR + PEQ + MX — Anton Braverman, J. G. Dai, & Jiekun Feng (2016). "Stein's method for steady-state diffusion approximations: an introduction through the Erlang-A and Erlang-C models". doi:10.1287/15-SSY212.

[BDM'17] BAR — Anton Braverman, J. G. Dai, & Masakiyo Miyazawa. "Heavy traffic approximation for the stationary distribution of a generalized Jackson network: the BAR approach". doi:10.1214/15-SSY199.

[ES'12] BAR + SSC — Atilla Eryilmaz & R. Srikant. "Asymptotically tight steady-state queue length bounds implied by drift conditions". doi:10.1007/s11134-012-9305-y.

[HM'20] BAR + SSC — Daniela Hurtado-Lange & Siva Theja Maguluri. "Transform methods for heavy-traffic analysis". doi:10.1287/stsy.2019.0056.

[MS'16] BAR + SSC — Siva Theja Maguluri & R. Srikant. "Heavy traffic queue length behavior in a switch under the MaxWeight algorithm". doi:10.1287/15-SSY193.

[SGH'20] BAR + MX — Ziv Scully, Isaac Grosof, & Mor Harchol-Balter. "The Gittins policy is nearly optimal in the M/G/k under extremely general conditions". doi:10.1145/3428328.

[Y'17] BAR + PEQ + MX — Lei Ying (2017). "Stein's method for mean field approximations in light and heavy traffic regimes". doi:10.1145/3084449.

References: other approaches

[BBLM'16] Mark van der Boor, Sem C. Borst, Johan S. H. van Leeuwaarden, & Debankur Mukherjee. "Scalable load balancing in networked systems: a survey of recent advances". arXiv:1806.05444.

[DW'06] Douglas Down & Wu. "Multi-layered round robin routing for parallel servers". doi:10.1007/s11134-006-7419-9.

[GHDZ'10] Varun Gupta, Mor Harchol-Balter, J. G. Dai, & Bert Zwart. "On the inapproximability of M/G/K: why two moments of job size distribution are not enough". doi:10.1007/s11134-009-9133-x.

[GHS'17] Kristen Gardner, Mor Harchol-Balter, & Alan Scheller-Wolf. "A better model for job redundancy: decoupling server slowdown and job size". doi:10.1109/TNET.2017.2744607.

[GR'20] Kristen Gardner & Rhonda Righter. "Product forms for FCFS queueing models with arbitrary server-job compatibilities: an overview". doi:10.1007/s11134-020-09668-6.

[GSH'18] Isaac Grosof, Ziv Scully, & Mor Harchol-Balter. "SRPT for multiserver systems". doi:10.1016/j.peva.2018.10.001.

[GSH'19] Isaac Grosof, Ziv Scully, & Mor Harchol-Balter. "Load balancing guardrails: keeping your heavy traffic on the road to low response times". doi:10.1145/3341617.3326157.

[LG'17] Yuan Li & David Goldberg (2017). "Simple and explicit bounds for multi-server queues with universal $1/(1-\rho)$ scaling". arXiv:1706.04628.

[WHJSS'19] Weina Wang, Mor Harchol-Balter, Haotian Jiang, Alan Scheller-Wolf, & R. Srikant. "Delay asymptotics and bounds for multitask parallel jobs". doi:10.1007/s11134-018-09597-5.