A Diffusion Approximation for Stationary Distribution of Many-Server Queueing System In Halfin-Whitt Regime

Mohammadreza Aghajani
joint work with Kavita Ramanan

Brown University

APS Conference, Istanbul, Turkey
July 2015

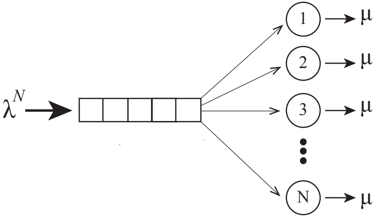

Where do they arise?

- Call Centers
- Health Care
- Data Centers

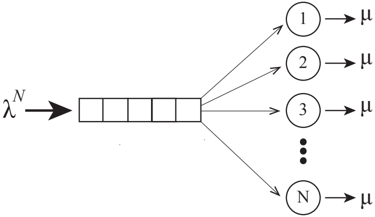

Relevant steady state performance measures:

- $\alpha_N = \mathbb{P}_{ss}\{\text{all } N \text{ servers are busy}\}$
- $\mathbb{P}_{ss}\{\text{wait} > t \text{ seconds}\}$

Asymptotic Analysis

Exact analysis for finite N is typically infeasible.

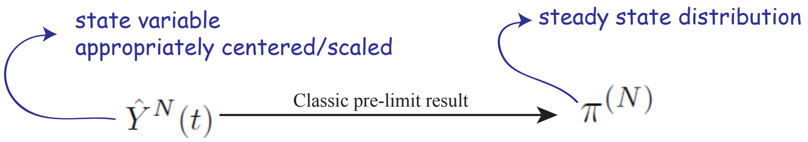

Asymptotic Analysis

Exact analysis for finite N is typically infeasible.

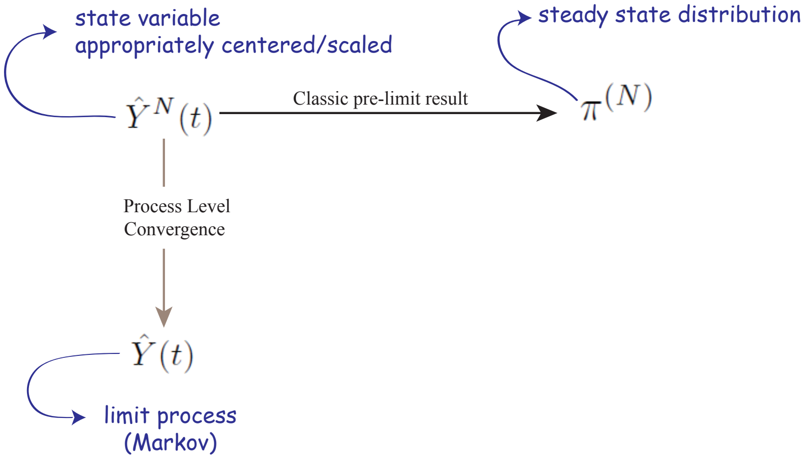

Asymptotic Analysis

Exact analysis for finite N is typically infeasible.

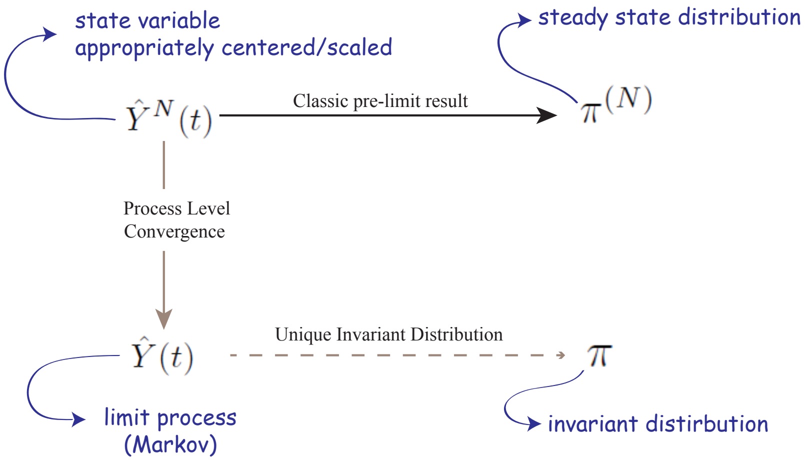

Asymptotic Analysis

Exact analysis for finite N is typically infeasible.

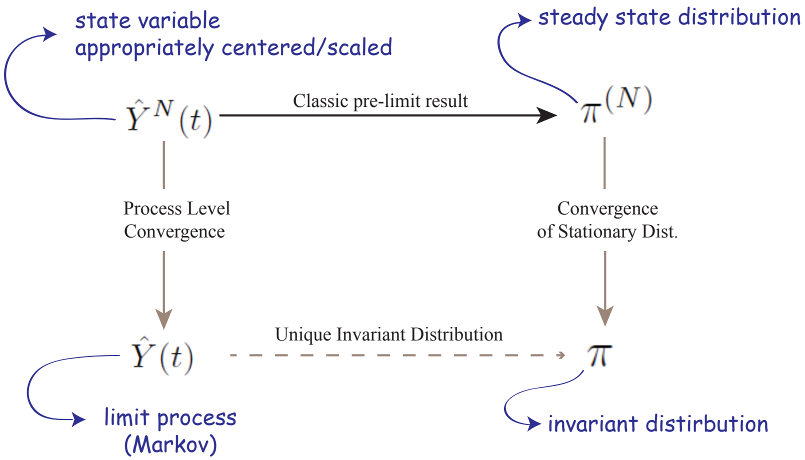

Outline

1. Recap on exponential service distribution

Outline

1. Recap on exponential service distribution
2. State representation for General service distribution

Outline

1. Recap on exponential service distribution
2. State representation for General service distribution
3. Characterization of the limit process

Outline

1. Recap on exponential service distribution
2. State representation for General service distribution
3. Characterization of the limit process
4. Proof of the main results

Outline

1. Recap on exponential service distribution
2. State representation for General service distribution
3. Characterization of the limit process
4. Proof of the main results
5. Ongoing work

1. Exponential Service Distribution

Halfin-Whitt Regime [Halfin-Whitt'81] for exponential service time

- Let $N \to \infty$, $\lambda^{(N)} = N\mu - \beta\sqrt{N} \to \infty$, $\rho^{(N)} = \lambda^{(N)}/N\mu \to 1$.
- Diffusion (CLT) scaling limit for $X_t^{(N)}$: # of customers in system.

1. Exponential Service Distribution

Halfin-Whitt Regime [Halfin-Whitt'81] for exponential service time

- Let $N \to \infty$, $\lambda^{(N)} = N\mu - \beta\sqrt{N} \to \infty$, $\rho^{(N)} = \lambda^{(N)}/N\mu \to 1$.
- Diffusion (CLT) scaling limit for $X_t^{(N)}$: # of customers in system.

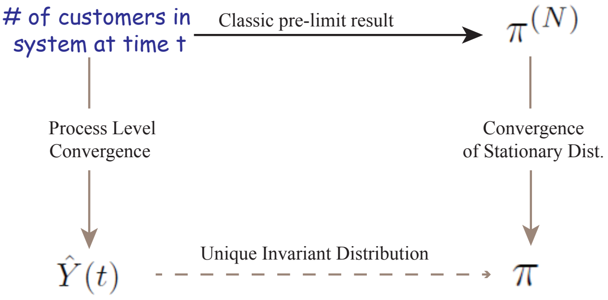

1. Exponential Service Distribution

Halfin-Whitt Regime [Halfin-Whitt'81] for exponential service time

- Let $N \to \infty$, $\lambda^{(N)} = N\mu - \beta\sqrt{N} \to \infty$, $\rho^{(N)} = \lambda^{(N)}/N\mu \to 1$.
- Diffusion (CLT) scaling limit for $X_t^{(N)}$: # of customers in system.

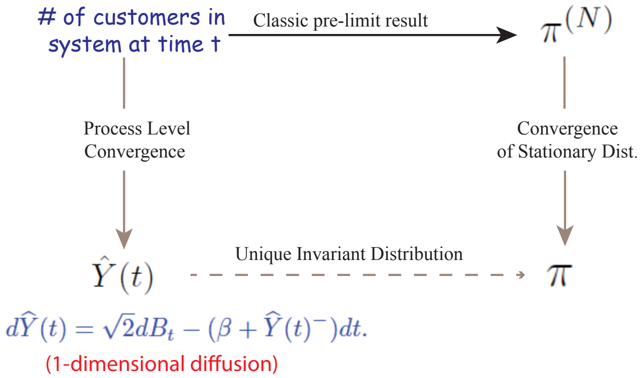

1. Exponential Service Distribution

Halfin-Whitt Regime [Halfin-Whitt'81] for exponential service time

- Let $N \to \infty$, $\lambda^{(N)} = N\mu - \beta\sqrt{N} \to \infty$, $\rho^{(N)} = \lambda^{(N)}/N\mu \to 1$.
- Diffusion (CLT) scaling limit for $X_t^{(N)}$: # of customers in system.

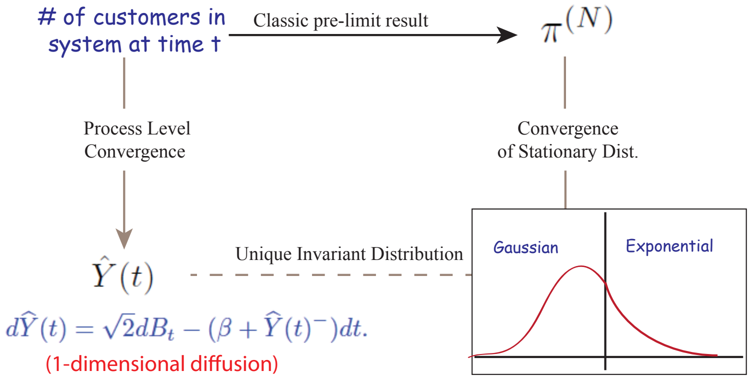

1. Exponential Service Distribution

Halfin-Whitt Regime [Halfin-Whitt'81] for exponential service time

- Let $N \to \infty$, $\lambda^{(N)} = N\mu - \beta\sqrt{N} \to \infty$, $\rho^{(N)} = \lambda^{(N)}/N\mu \to 1$.
- Diffusion (CLT) scaling limit for $X_t^{(N)}$: # of customers in system.

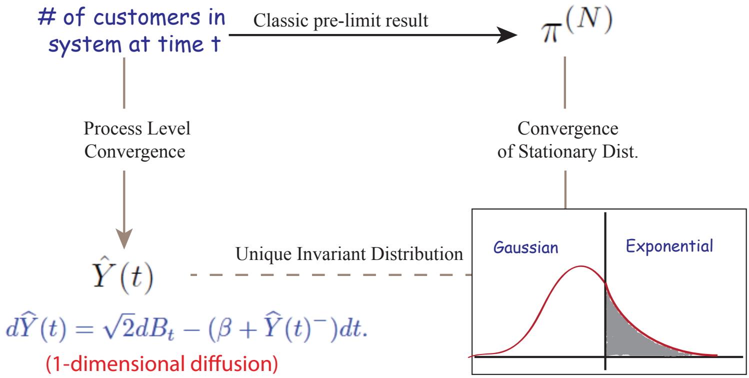

- $\mathbb{P}_{ss}$(all N servers are busy) $\to \pi([0,\infty)) \in (0,1)$.

2. General Service Distribution

- Statistical data shows that service times are generally distributed (Lognormal, Pareto, etc. see e.g. [Brown et al. '05])

Goal: To extend the result for general service distribution

2. General Service Distribution

- Statistical data shows that service times are generally distributed (Lognormal, Pareto, etc. see e.g. [Brown et al. '05])

Goal: To extend the result for general service distribution

- **Challenges**
  - $X^{(N)}$ is no longer a Markov Process
  - need to keep track of residual times or ages of customers in service to make the process Markovian
  - Dimension of any finite-dim. Markovian representation grows with N

2. General Service Distribution

- Statistical data shows that service times are generally distributed (Lognormal, Pareto, etc. see e.g. [Brown et al. '05])

Goal: To extend the result for general service distribution

- **Challenges**
  - $X^{(N)}$ is no longer a Markov Process
  - need to keep track of residual times or ages of customers in service to make the process Markovian
  - Dimension of any finite-dim. Markovian representation grows with N
- **Prior Work**
  - Some particular service distributions [Jelenkovic-Mandelbaum], [Gamarnik-Momcilovic], [Puhalski-Reiman].
  - Results using X^N obtained by [Puhalskii-Reed], [Reed], [Mandelbaum-Momcilovic], [Dai-He] (with abandonment), etc.
  - However, there are not many results on stationary distribution.

2. General Service Distribution

- Statistical data shows that service times are generally distributed (Lognormal, Pareto, etc. see e.g. [Brown et al. '05])

Goal: To extend the result for general service distribution

- **Challenges**
  - $X^{(N)}$ is no longer a Markov Process
  - need to keep track of residual times or ages of customers in service to make the process Markovian
  - Dimension of any finite-dim. Markovian representation grows with N
- **Prior Work**
  - Some particular service distributions [Jelenkovic-Mandelbaum], [Gamarnik-Momcilovic], [Puhalski-Reiman].
  - Results using X^N obtained by [Puhalskii-Reed], [Reed], [Mandelbaum-Momcilovic], [Dai-He] (with abandonment), etc.
  - However, there are not many results on stationary distribution.

A way out: Common State Space (infinite-dimensional)

A Measure-valued Representation

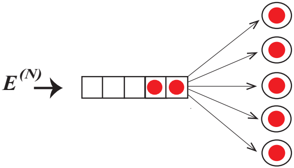

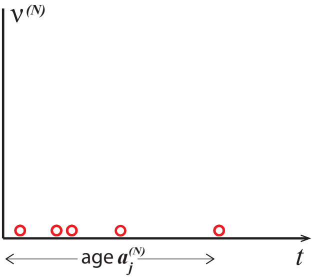

- $E^{(N)}$ represents the cumulative external arrivals
- $a_j^{(N)}$ represents age of the jth customer to enter service
- $\nu^{(N)}$ keeps track of the ages of all the customers in service

$$\nu_t^{(N)} = \sum_j \delta_{a_j^{(N)}}(t)$$

A Measure-valued Representation

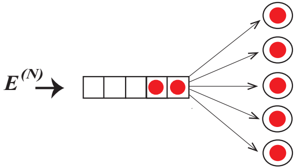

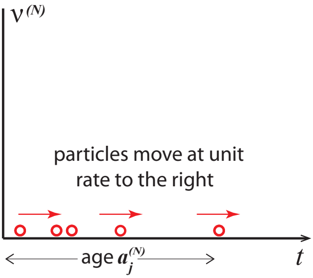

- $E^{(N)}$ represents the cumulative external arrivals
- $a_j^{(N)}$ represents age of the jth customer to enter service
- $\nu^{(N)}$ keeps track of the ages of all the customers in service

$$\nu_t^{(N)} = \sum_j \delta_{a_j^{(N)}(t)}$$

A Measure-valued Representation

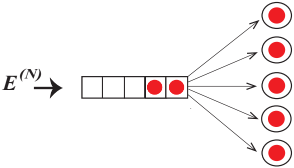

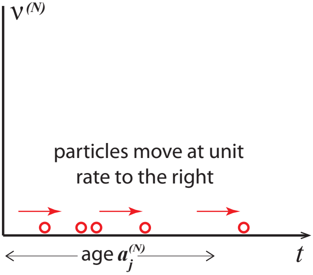

- $E^{(N)}$ represents the cumulative external arrivals
- $a_j^{(N)}$ represents age of the jth customer to enter service
- $\nu^{(N)}$ keeps track of the ages of all the customers in service

$$\nu_t^{(N)} = \sum_j \delta_{a_j^{(N)}}(t)$$

A Measure-valued Representation

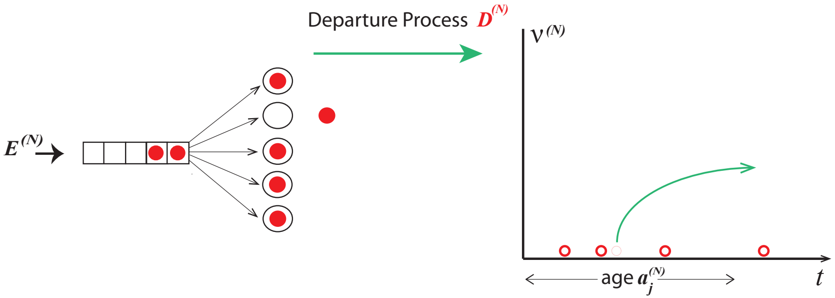

- $E^{(N)}$ represents the cumulative external arrivals
- $a_j^{(N)}$ represents age of the jth customer to enter service
- $\nu^{(N)}$ keeps track of the ages of all the customers in service
- $K^{(N)}$ cumulative entry into service
- $D^{(N)}$ cumulative departure process

A Measure-valued Representation

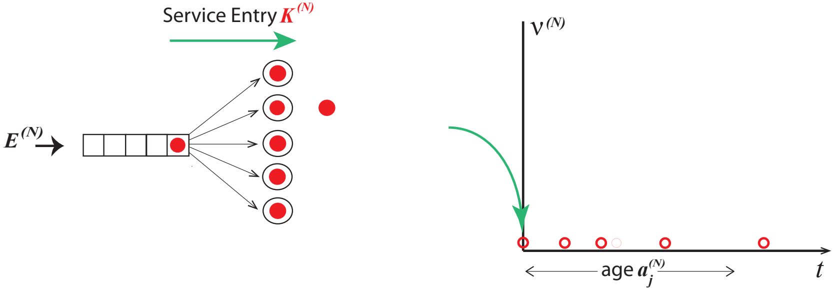

- $E^{(N)}$ represents the cumulative external arrivals
- $a_j^{(N)}$ represents age of the jth customer to enter service
- $\nu^{(N)}$ keeps track of the ages of all the customers in service
- $K^{(N)}$ cumulative entry into service
- $D^{(N)}$ cumulative departure process

A New Representation

- State descriptor $S_t^{(N)} = \left(X_t^{(N)}, \nu_t^{(N)}\right)$ is used in [Kaspi-Ramanan '11,'13] and [Kang- Ramanan '10, '12.]
  - Diffusion limit for $\nu^{(N)}$ is established in a distribution space $\mathbb{H}_{-2}$.
  - An extra component needs to be added for the limit process to be Markov.

A New Representation

- State descriptor $S_t^{(N)} = \left(X_t^{(N)}, \nu_t^{(N)}\right)$ is used in [Kaspi-Ramanan '11,'13] and [Kang- Ramanan '10, '12.]
  - Diffusion limit for $\nu^{(N)}$ is established in a distribution space $\mathbb{H}_{-2}$.
  - An extra component needs to be added for the limit process to be Markov.
- Instead of the whole measure ν, we define the functional

$$Z_t^{(N)}(r) \doteq \langle \frac{\overline{G}(\cdot + r)}{\overline{G}(\cdot)}, \nu_t^{(N)} \rangle = \sum_{j \text{ in service}} \frac{\overline{G}(a_j(t) + r)}{\overline{G}(a_j(t))}, \ r \geq 0,$$

which we call Frozen Departure Process.

A New Representation

- State descriptor $S_t^{(N)} = \left(X_t^{(N)}, \nu_t^{(N)}\right)$ is used in [Kaspi-Ramanan '11,'13] and [Kang- Ramanan '10, '12.]
  - Diffusion limit for $\nu^{(N)}$ is established in a distribution space $\mathbb{H}_{-2}$.
  - An extra component needs to be added for the limit process to be Markov.
- Instead of the whole measure ν, we define the functional

$$Z_t^{(N)}(r) \doteq \langle \frac{\overline{G}(\cdot + r)}{\overline{G}(\cdot)}, \nu_t^{(N)} \rangle = \sum_{j \text{ in service}} \frac{\overline{G}(a_j(t) + r)}{\overline{G}(a_j(t))}, \; r \geq 0,$$

which we call Frozen Departure Process.

- We use the state variable

$$Y_t^{(N)} = (X_t^{(N)}, Z_t^{(N)}) \in \mathbb{R} \times \mathbb{H}^1(0, \infty).$$

Main Results

Now we establish diffusion level "change of limits" for $Y^{(N)}(t)$.

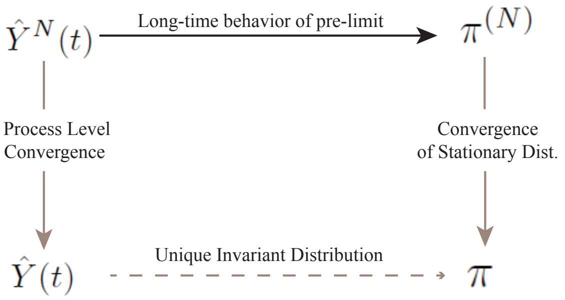

Main Contributions

Now we establish diffusion level "change of limits" for $Y^{(N)}(t)$.

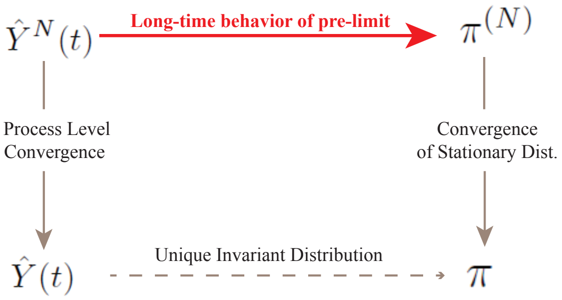

Main Contributions

Main Results

Now we establish diffusion level "change of limits" for $Y^{(N)}(t)$.

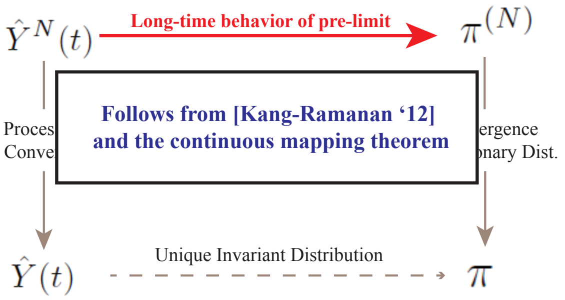

Main Contributions

Now we establish diffusion level "change of limits" for $Y^{(N)}(t)$.

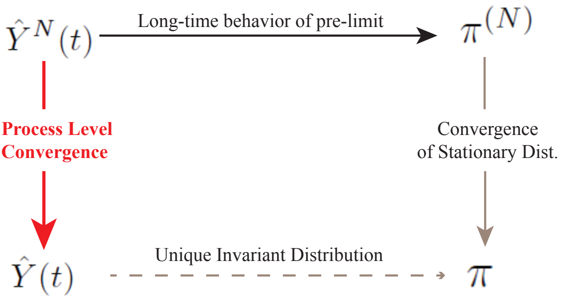

Main Contributions

Main Results

Now we establish diffusion level “change of limits” for $Y^{(N)}(t)$.

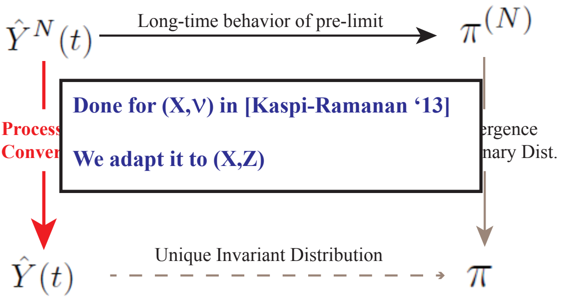

Main Contributions

Main Results

Now we establish diffusion level "change of limits" for $Y^{(N)}(t)$.

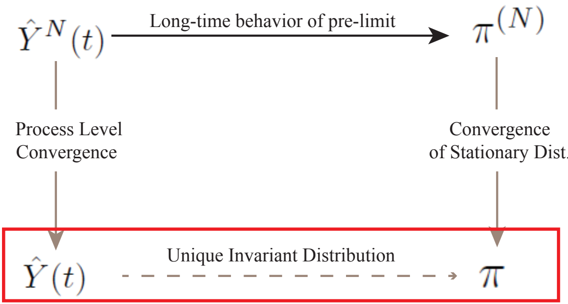

Main Contributions

- Characterization of the limit (X, Z) in terms of an SPDE in an appropriate space that makes it Markov

Main Results

Now we establish diffusion level "change of limits" for $Y^{(N)}(t)$.

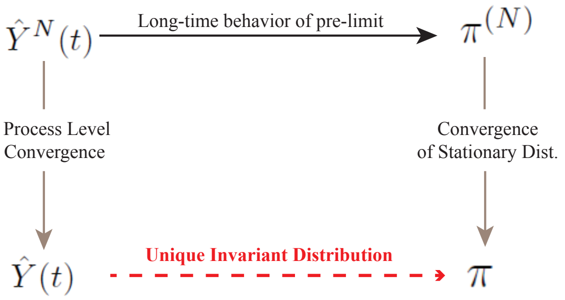

Main Contributions

- Characterization of the limit (X, Z) in terms of an SPDE in an appropriate space that makes it Markov
- Showing that (X, Z) has a unique invariant distribution

Now we establish diffusion level "change of limits" for $Y^{(N)}(t)$.

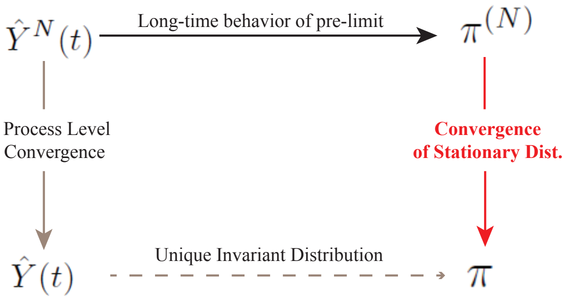

Main Contributions

- Characterization of the limit (X, Z) in terms of an SPDE in an appropriate space that makes it Markov
- Showing that (X, Z) has a unique invariant distribution
- Proving $\pi^{(N)} \mapsto \pi$, with partial characterization of π

Comments on Our Results:

- Previously, $\{\hat{X}_{\infty}^{(N)}\}$ (the X-marginal of $\pi^{(N)}$) was only shown to be tight [Gamarnik-Goldberg]. We proved the convergence.

Implications of our Results

Comments on Our Results:

- Previously, $\{\hat{X}_\infty^{(N)}\}$ (the X-marginal of $\pi^{(N)}$) was only shown to be tight [Gamarnik-Goldberg]. We proved the convergence.
- The limit π is now the invariant distribution of a Markov process. We can use basic adjoint relation type formulations to characterize it.

Implications of our Results

Comments on Our Results:

- Previously, $\{\hat{X}_\infty^{(N)}\}$ (the X-marginal of $\pi^{(N)}$) was only shown to be tight [Gamarnik-Goldberg]. We proved the convergence.
- The limit π is now the invariant distribution of a Markov process. We can use basic adjoint relation type formulations to characterize it.
- As the limit process (X, Z) is infinite dimensional, we use the newly developed method of asymptotic coupling to prove the uniqueness of invariant distribution.

3. Characterization of Limit Process

Consider the following "SPDE":

$$\begin{cases} dX_t = -d\mathcal{M}_t(1) + dB_t - \beta dt + Z_t'(0)dt, \\ dZ_t(r) = \left[Z_t'(r) - \bar{G}(r)Z_t'(0)\right] dt - d\mathcal{M}_t \left(\Phi_r 1 - \bar{G}(r)1\right) + \bar{G}(r)dZ_t(0) \end{cases}$$

with boundary condition $Z_t(0) = -X_t^-$, and initial condition Y_0.

B is a standard Brownian motion, $\mathcal{M}$ is an independent martingale measure.

3. Characterization of Limit Process

Consider the following “SPDE”:

$$\begin{cases} dX_t = -d\mathcal{M}_t(1) + dB_t - \beta dt + Z_t'(0)dt, \\ dZ_t(r) = \left[Z_t'(r) - \bar{G}(r)Z_t'(0)\right] dt - d\mathcal{M}_t\left(\Phi_r 1 - \bar{G}(r)1\right) + \bar{G}(r)dZ_t(0) \end{cases}$$

with boundary condition $Z_t(0) = -X_t^-$, and initial condition Y_0.

B is a standard Brownian motion, $\mathcal{M}$ is an independent martingale measure.

Assumptions: I. hazard rate function $h(x) \doteq g(x)/\overline{G}(x)$ is bounded;
II. G has finite $2+\epsilon$ moment for some $\epsilon > 0$;

Theorem

If Assumptions I. and II. hold, for every initial condition Y_0, the SPDEs above a unique continuous $\mathbb{R} \times \mathbb{H}^1(0,\infty)$-valued solution, which is a Markov process.

Characterization of Limit Process

Given initial condition $y_0 = (x_0, z_0)$, we can "explicitly" solve the SPDE:

- X is a solution to a non-linear Volterra equation ([Reed], [Puhalskii-Reed],[Kaspi-Ramanan])

Characterization of Limit Process

Given initial condition $y_0 = (x_0, z_0)$, we can "explicitly" solve the SPDE:

- X is a solution to a non-linear Volterra equation ([Reed], [Puhalskii-Reed],[Kaspi-Ramanan])
- The service entry process K satisfies

$$K(t) = B_t - \beta t - X^+(t) + x_0^+.$$

Characterization of Limit Process

Given initial condition $y_0 = (x_0, z_0)$, we can "explicitly" solve the SPDE:

- X is a solution to a non-linear Volterra equation ([Reed], [Puhalskii-Reed],[Kaspi-Ramanan])
- The service entry process K satisfies

$$K(t) = B_t - \beta t - X^+(t) + x_0^+.$$

- Given X (and hence K), the equation for Z is a transport equation.

$$Z_t(r) = z_0(t+r) - \mathcal{M}_t(\Psi_{t+r}1) + \big(\Gamma_t K\big)(r).$$

$\{\Psi_t; t \geq 0\}$ and $\{\Gamma_t; t \geq 0\}$ are certain family of mappings on continuous functions.

Invariant Distribution of the Limit Process

Existence: "Standard." Follows from Krylov-Bogoliubov Theorem.

Invariant Distribution of the Limit Process

Existence: "Standard." Follows from Krylov-Bogoliubov Theorem.

Uniqueness:

- Key challenge: State Space $\mathcal{Y} \doteq \mathbb{R} \times \mathbb{H}^1$ is infinite dimensional
- Traditional recurrence methods are not easily applicable.
- In some cases, traditional methods fail: the stochastic delay differential equation example in [Hairer et. al.'11].

Invariant Distribution of the Limit Process

Existence: "Standard." Follows from Krylov-Bogoliubov Theorem.

Uniqueness:

- Key challenge: State Space $\mathcal{Y} \doteq \mathbb{R} \times \mathbb{H}^1$ is infinite dimensional
- Traditional recurrence methods are not easily applicable.
- In some cases, traditional methods fail: the stochastic delay differential equation example in [Hairer et. al.'11].
- We invoke the asymptotic coupling method (Hairer, Mattingly, Sheutzow, Bakhtin, et al.)

Invariant Dist. of the Limit Process: Uniqueness

Theorem (Hairer et. al'11, continuous version)

Assume there exists a measurable set $A \subseteq \mathcal{Y}$ with following properties:

(I) $\mu(A) > 0$ *for any invariant probability measure* μ *of* $\mathcal{P}_t$.

(II) *For every* $y, \tilde{y} \in A$, *there exists a measurable map* $\Gamma_{y,\tilde{y}} : A \times A \to \tilde{\mathcal{C}}(\mathcal{P}_{[0,\infty)}\delta_y, \mathcal{P}_{[0,\infty)}\delta_{\tilde{y}})$, *such that* $\Gamma_{y,\tilde{y}}(\mathcal{D}) > 0$.

Then $\{\mathcal{P}_t\}$ *has at most one invariant probability measure.*

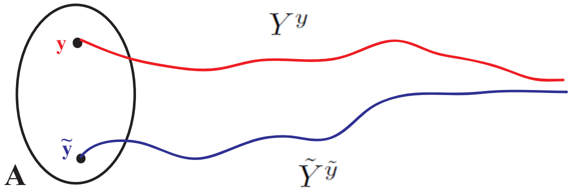

Invariant Dist. of the Limit Process: Uniqueness

Theorem (Hairer et. al'11, continuous version)

Assume there exists a measurable set $A \subseteq \mathcal{Y}$ with following properties:

(I) $\mu(A) > 0$ *for any invariant probability measure* μ *of* $\mathcal{P}_t$.

(II) *For every* $y, \tilde{y} \in A$, *there exists a measurable map* $\Gamma_{y,\tilde{y}} : A \times A \to \tilde{\mathcal{C}}(\mathcal{P}_{[0,\infty)}\delta_y, \mathcal{P}_{[0,\infty)}\delta_{\tilde{y}})$, *such that* $\Gamma_{y,\tilde{y}}(\mathcal{D}) > 0$.

Then $\{\mathcal{P}_t\}$ *has at most one invariant probability measure.*

To prove the uniqueness of the inv. dist. for a Markov kernel $\mathcal{P}$:

- Specify the subset A.
- For $y, \tilde{y} \in A$, construct $(Y^y, \tilde{Y}^{\tilde{y}})$ on a common probability space:
  - verify the marginals of Y^y and $\tilde{Y}^{\tilde{y}}$.
  - show the asymptotic convergence: $\mathbb{P}\left\{d(Y^y(t), \tilde{Y}^{\tilde{y}}(t)) \to 0\right\} > 0$.

Then $\Gamma_{y,\tilde{y}} = \text{Law}(Y^y, \tilde{Y}^{\tilde{y}})$ is a legitimate asymptotic coupling.

Theorem

Under assumptions I, II and IV, the limit process has at most one invariant distribution.

Proof idea. Let $y = (x_0, z_0)$ and $\tilde{y} = (\tilde{x}_0, \tilde{y}_0)$. Recall

$$\begin{cases} X_t = x_0 - \mathcal{M}_t(1) + B_t - \beta t + \int_0^t Z_s'(0)ds, & t \geq 0, \\ Z_t(r) = z_0(t+r) - \mathcal{M}_t(\Psi_{t+r}1) + \big(\Gamma_t K\big)(r), & r \geq 0. \end{cases}$$

Now define

$$\begin{cases} \tilde{X}_t = \tilde{x}_0 - \mathcal{M}_t(1) + \tilde{B}_t - \beta t + \int_0^t \tilde{Z}_s'(0)ds, & t \geq 0, \\ \tilde{Z}_t(r) = \tilde{z}_0(t+r) - \mathcal{M}_t(\Psi_{t+r}1) + \big(\Gamma_t \tilde{K}\big)(r), & r \geq 0. \end{cases}$$

where

$$\boxed{\tilde{B}_t = B_t + \int_0^t \left(\Delta Z_s'(0) - \lambda \Delta X_s\right) ds.}$$

Invariant Dist. of the Limit Process: Uniqueness

Define $A = \{(x, z) \in \mathcal{Y}; x \geq 0\}$.

- For every invariant distribution μ of $\mathcal{P}$, $\mu(A) > 0$.

Invariant Dist. of the Limit Process: Uniqueness

Define $A = \{(x, z) \in \mathcal{Y}; x \geq 0\}$.

- For every invariant distribution μ of $\mathcal{P}$, $\mu(A) > 0$.

Asymptotic Convergence:

- $\Delta X_t = \Delta x_0 e^{-\lambda t} \quad \Rightarrow \quad \Delta X_t \to 0.$

Invariant Dist. of the Limit Process: Uniqueness

Define $A = \{(x,z) \in \mathcal{Y}; x \geq 0\}$.

- For every invariant distribution μ of $\mathcal{P}$, $\mu(A) > 0$.

Asymptotic Convergence:

- $\Delta X_t = \Delta x_0 e^{-\lambda t} \quad \Rightarrow \quad \Delta X_t \to 0.$

Lemma (2)

When $y, \tilde{y} \in A$, we have $\Delta Z'_{\cdot}(0) \in \mathbb{L}^2$

- Using Lemma 2, $\Delta Z_t \to 0$ in $\mathbb{H}^1(0,\infty)$.

$$\Delta Z_t(r) = \Delta z_0(t+r) + \bar{G}(r)\Delta X_t^- + \int_0^t \Delta X_s^- g(t+r-s)ds - \int_0^t \Delta Z'_s(0)\bar{G}(t+r-s)ds.$$

Invariant Dist. of the Limit Process: Uniqueness

Define $A = \{(x, z) \in \mathcal{Y}; x \geq 0\}$.

- For every invariant distribution μ of $\mathcal{P}$, $\mu(A) > 0$.

Asymptotic Convergence:

- $\Delta X_t = \Delta x_0 e^{-\lambda t} \quad \Rightarrow \quad \Delta X_t \to 0.$

Lemma (2)

When $y, \tilde{y} \in A$, we have $\Delta Z'_{\cdot}(0) \in \mathbb{L}^2$

- Using Lemma 2, $\Delta Z_t \to 0$ in $\mathbb{H}^1(0, \infty)$.

$$\Delta Z_t(r) = \Delta z_0(t+r) + \bar{G}(r)\Delta X_t^- + \int_0^t \Delta X_s^- g(t+r-s)ds - \int_0^t \Delta Z'_s(0)\bar{G}(t+r-s)ds.$$

Distribution of $\tilde{Y}$:

- By Girsanov Theorem, the distribution of $\tilde{B}$ is equivalent to a Brownian motion. Novikov condition follows from Lemma 2.

$$\tilde{Y} \sim \mathcal{P}_{\lfloor \infty \rfloor} \delta_{\tilde{y}}.$$

4. Convergence of Steady-State Distributions

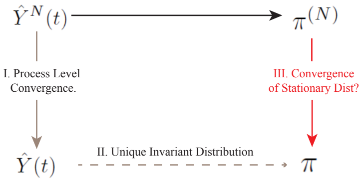

Further Assumptions:

III. $\varrho \doteq \sup\{u \in [0, \infty), \; g = 0 \text{ a.e. on } [a, a+u] \text{ for some a } \in [0, \infty)\} < \infty.$

IV. g has a density g' and $h_2(x) \doteq \frac{g'(x)}{\bar{G}(x)}$ is bounded.

Theorem (Aghajani and 'R'13)

Under assumptions I-IV and if G has a finite $3+\epsilon$ moment, the sequence $\{\pi^{(N)}\}$ converges weakly to the unique invariant distribution π of Y.

Convergence of Steady-State Distributions

Proof sketch.

Step 1.

Under assumptions on G, the sequence $\{\pi^{(N)}\}$ of steady state distributions of pre-limit processes is tight in $\mathbb{R} \times \mathbb{H}^1(0, \infty)$.

Proof idea: establish uniform bounds on $(X^{(N)}, Z^{(N)})$ in N, t, using results in [Gamarnik and Goldberg'13].

Convergence of Steady-State Distributions

Proof sketch.

Step 1.

Under assumptions on G, the sequence $\{\pi^{(N)}\}$ of steady state distributions of pre-limit processes is tight in $\mathbb{R} \times \mathbb{H}^1(0, \infty)$.

Proof idea: establish uniform bounds on $(X^{(N)}, Z^{(N)})$ in N, t, using results in [Gamarnik and Goldberg'13].

Step 2.

Every subsequential limit of $\{\pi^{(N)}\}$ is an invariant distribution for the limit process Y.

Convergence of Steady-State Distributions

Proof sketch.

Step 1.

Under assumptions on G, the sequence $\{\pi^{(N)}\}$ of steady state distributions of pre-limit processes is tight in $\mathbb{R} \times \mathbb{H}^1(0,\infty)$.

Proof idea: establish uniform bounds on $(X^{(N)}, Z^{(N)})$ in N, t, using results in [Gamarnik and Goldberg'13].

Step 2.

Every subsequential limit of $\{\pi^{(N)}\}$ is an invariant distribution for the limit process Y.

Step 3.

Combine Steps 1 and 2. By uniqueness of invariant distribution for the limit process Y, we have our final result. □

Makes key use of the fact that Y is Markovian.

Summary and Conclusion

Some subtleties

- Finding a more tractable representation
  - conserved the Markov property of the diffusion limit
  - been able to remove the problematic ν component
- Prove the uniqueness of invariant distribution for the inf. dim. limit process

*Our proposed asymptotic coupling scheme does not work.

Summary and Conclusion

Some subtleties

- Finding a more tractable representation
  - conserved the Markov property of the diffusion limit
  - been able to remove the problematic ν component
- Prove the uniqueness of invariant distribution for the inf. dim. limit process
- **Key Challenge** Choosing the right space for Z

Space	Markov Property	SPDE Charac.	Uniqueness of Stat. Dist.
$\mathbb{C}[0,\infty)$	Yes	No	Unknown*
$\mathbb{C}^1[0,\infty)$	Yes	Yes	Unknown
$\mathbb{L}^2(0,\infty)$	Unknown	No	Yes
$\mathbb{H}^1(0,\infty)$	Yes	Yes	Yes

*Our proposed asymptotic coupling scheme does not work.

Summary and Conclusion

Some subtleties

- Finding a more tractable representation
  - conserved the Markov property of the diffusion limit
  - been able to remove the problematic ν component
- Prove the uniqueness of invariant distribution for the inf. dim. limit process
- **Key Challenge** Choosing the right space for Z

Space	Markov Property	SPDE Charac.	Uniqueness of Stat. Dist.
$\mathbb{C}[0,\infty)$	Yes	No	Unknown*
$\mathbb{C}^1[0,\infty)$	Yes	Yes	Unknown
$\mathbb{L}^2(0,\infty)$	Unknown	No	Yes
$\mathbb{H}^1(0,\infty)$	Yes	Yes	Yes

- In our construction, $A \neq \mathcal{Y}$ and therefore, the continuous-time version of Asymptotic Coupling theorem does not immediately follow from the discrete-time version.

*Our proposed asymptotic coupling scheme does not work.

5. What Else Can This be Used For?

- Seems to be a useful framework to do diffusion control (fluid version is done in [Atar-Kaspi-Shimkin '12])
- Use generator to get error bounds for finite N ([Braverman-Dai] in finite dimension.)
- Characterization of invariant distribution using infinitesimal generator of the limit process and basic adjoint relation.

Characterization of the generator $\mathcal{L}$ of the diffusion process Y.

- for $f(x,z) = \tilde{f}(x, z(r_1), ..., z(r_n))$ with $\tilde{f} \in \mathbb{C}_c^2(\mathbb{R}^{n+1})$:

$$\mathcal{L}f = \begin{cases} \mathcal{L}_+ f & \text{if } x > 0, \\ \mathcal{L}_- f & \text{if } x < 0. \end{cases}$$

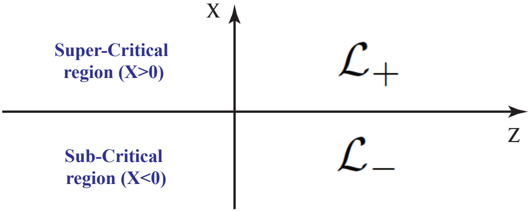

Characterization of Invariant Distribution

Characterization of the generator $\mathcal{L}$ of the diffusion process Y.

- for $f(x,z) = \tilde{f}(x, z(r_1), ..., z(r_n))$ with $\tilde{f} \in \mathbb{C}_c^2(\mathbb{R}^{n+1})$:

$$\mathcal{L}f = \begin{cases} \mathcal{L}_+ f & \text{if } x > 0, \\ \mathcal{L}_- f & \text{if } x < 0. \end{cases}$$

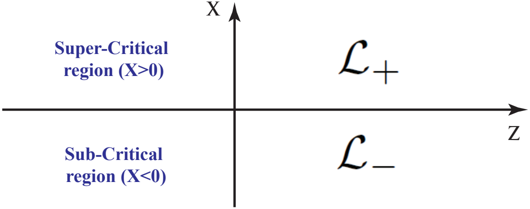

- $\mathcal{L}_+$ and $\mathcal{L}_-$ are second order differential operators, whose explicit forms are known.
- $\mathcal{L}_-$ is the generator of an "infinite-server" queue.
- $\mathcal{L}_+$ is the generator of the limit of a system composed of N decoupled closed queues.

Characterization of Invariant Distribution

An Idea: analyze sub-critical and super-critical systems and identify φ_+ and φ_- which satisfy $\mathcal{L}^*_-\varphi = 0$ and $\mathcal{L}^*_-\varphi = 0$, respectively, then glue them together such that φ is smooth at the boundary.

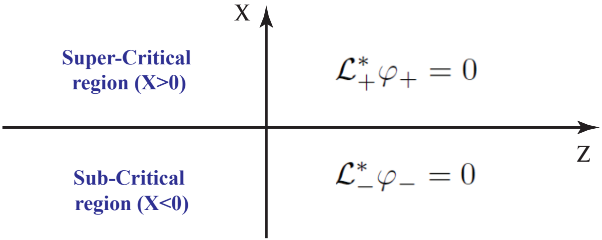

Summary and Conclusions:

- Introduced a more tractable SPDE framework for the study of diffusion limits of many-server queues
- Use of the asymptotic coupling method (as opposed to Lyapunov function methods) to establishing stability properties of queueing networks: more suitable for infinite-dimensional processes
- Strengthened the Gamarnik-Goldberg tightness result to convergence of the X-marginal
- A wide range of service distributions satisfy our assumptions, including Log-Normal, Pareto (for certain parameters), Gamma, Phase-Type, etc. Weibull does not.

Future challenges:

- Complete the characterization of the stationary distribution of the limit Markovian process.
- Extensions to more general systems