Parallel-Server Stochastic Systems with Dynamic Affinity Scheduling and Load Balancing

Mark S. Squillante*
Mathematical Sciences Department
January 18, 2005

* Based mostly on joint work with R. Nelson

Problem Motivation

Parallel Computing Example

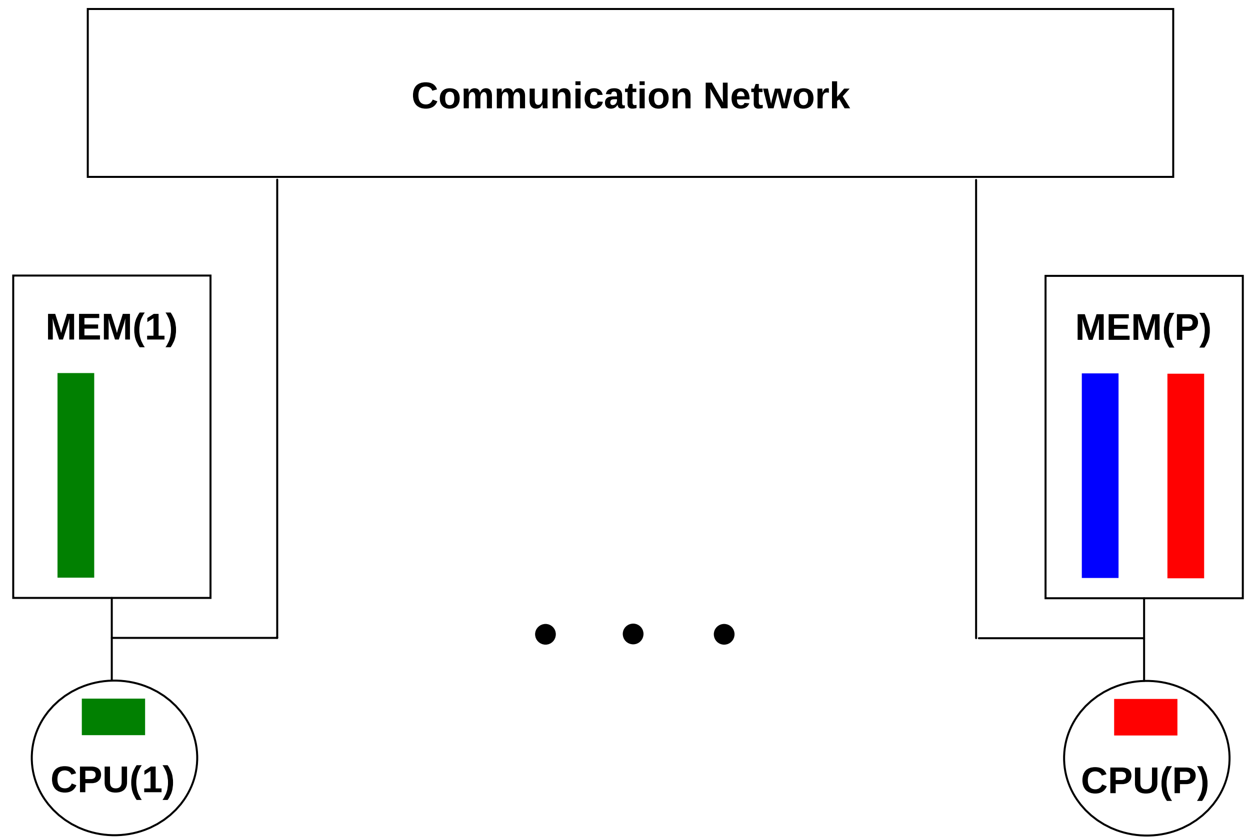

Problem Motivation

Parallel Computing Example

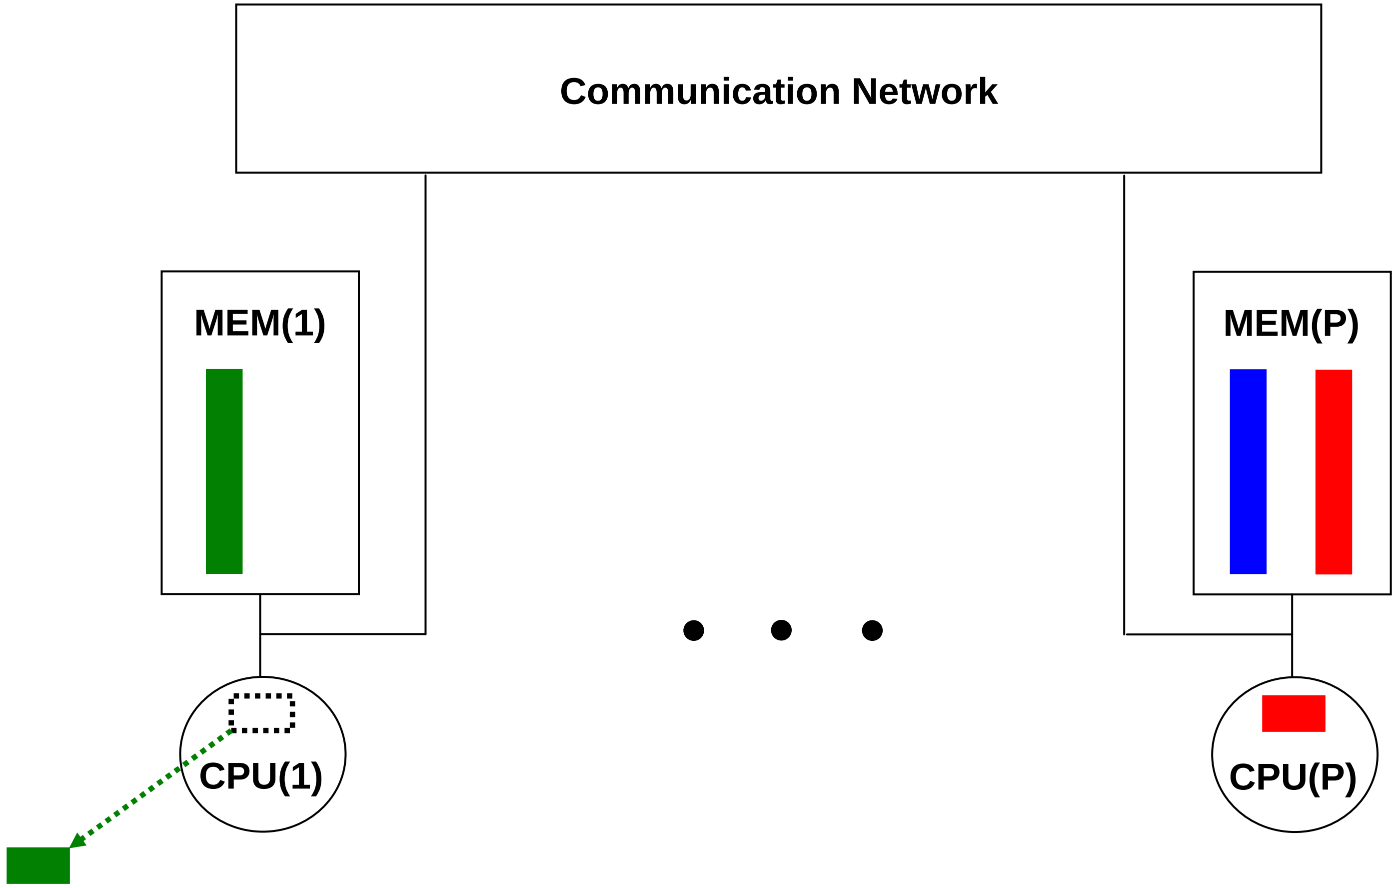

Problem Motivation

Dynamic Scheduling Tradeoff: Should we leave CPU(1) idle?

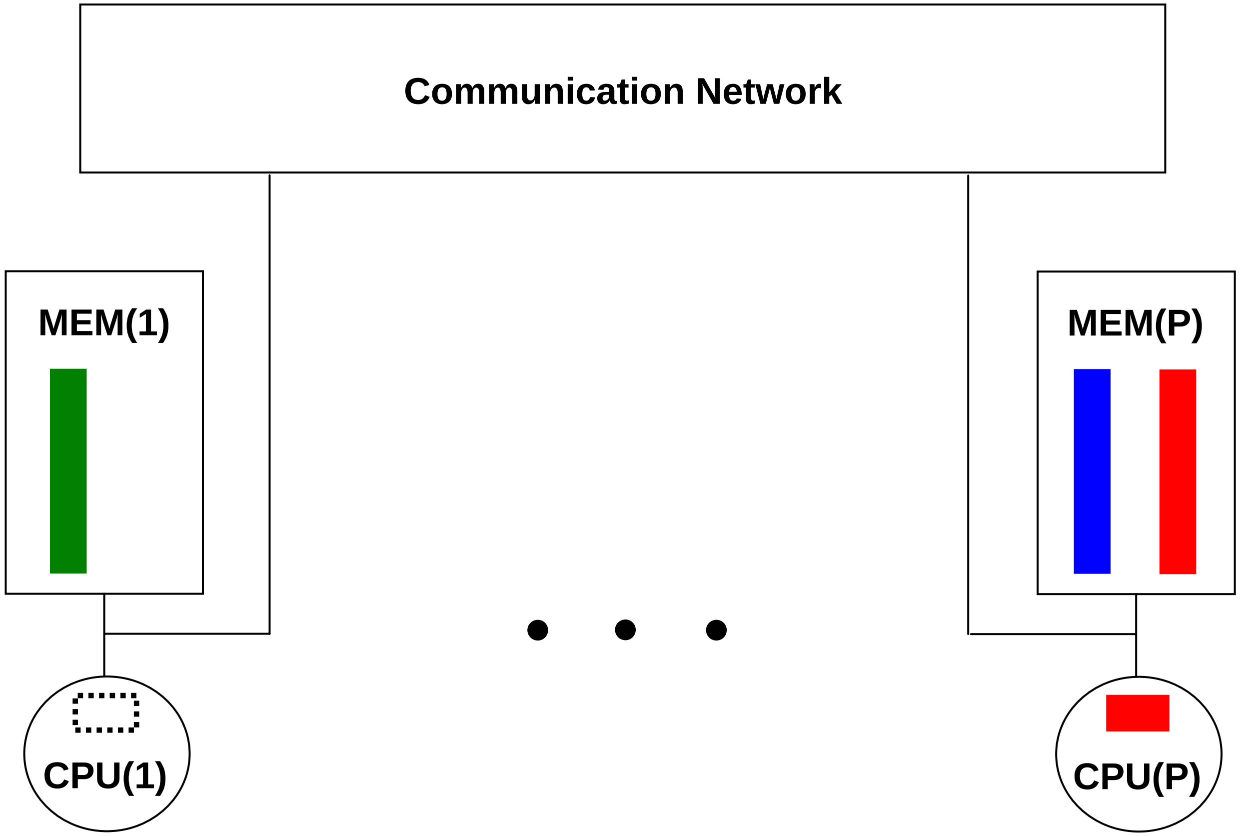

Problem Motivation

Dynamic Scheduling Tradeoff: Or serve Blue-Task at CPU(1)?

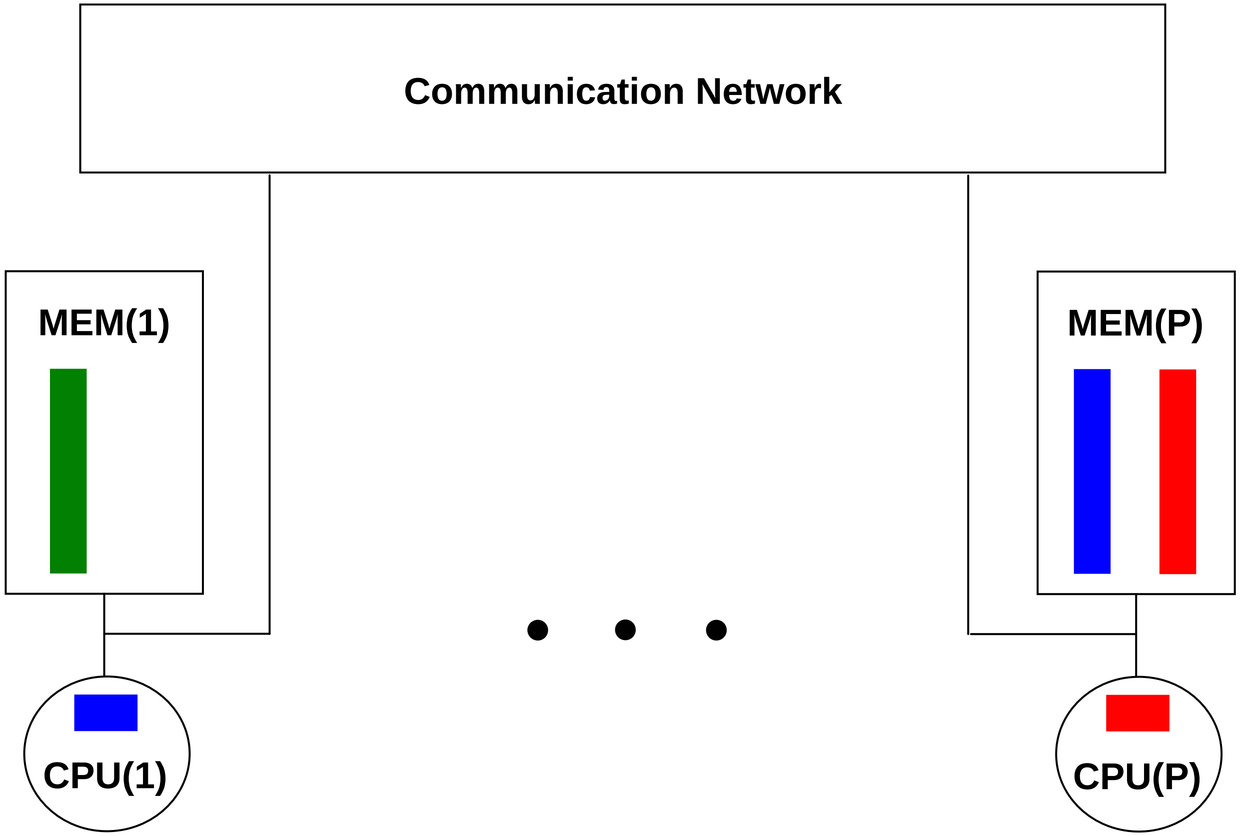

Problem Motivation

Dynamic Scheduling Tradeoff: But larger Blue-Task service time

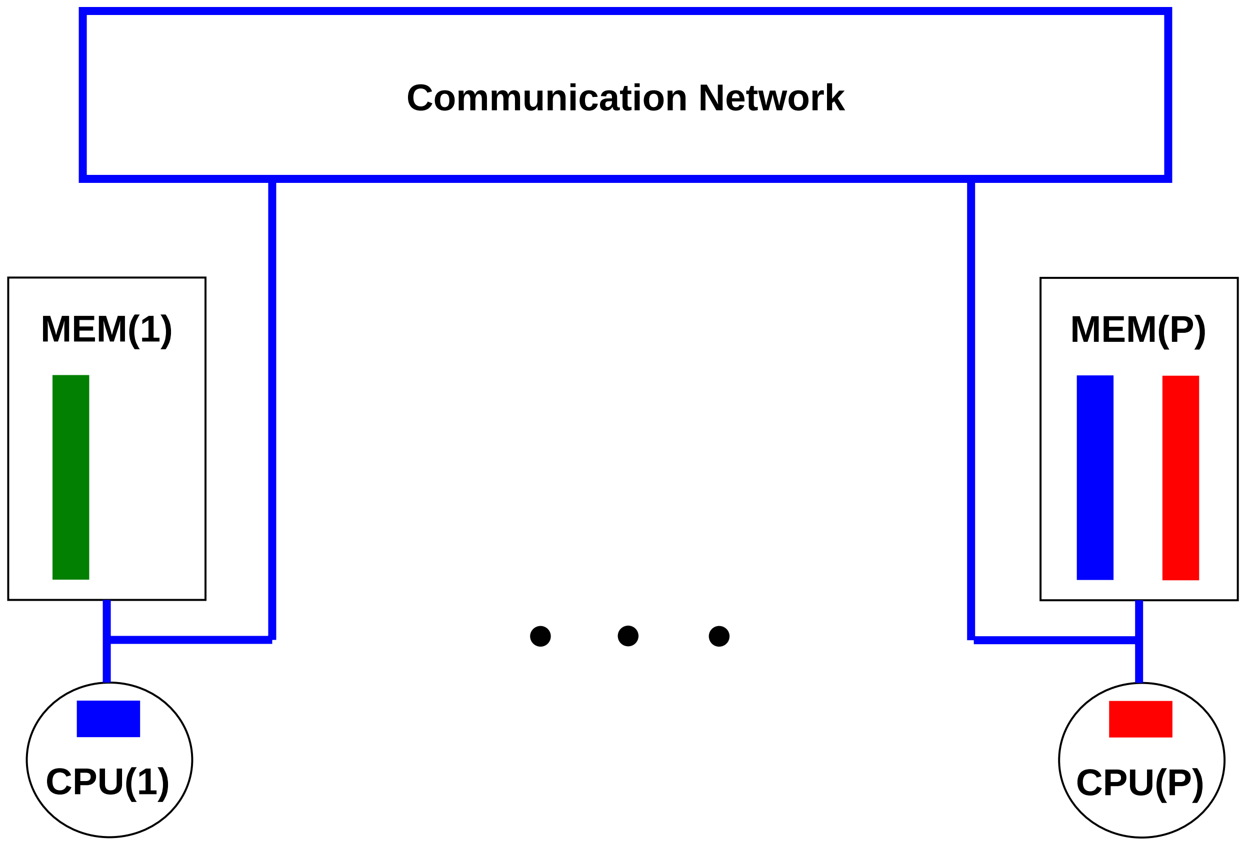

Problem Motivation

Cache Affinity [SquillanteLazowska89,93]

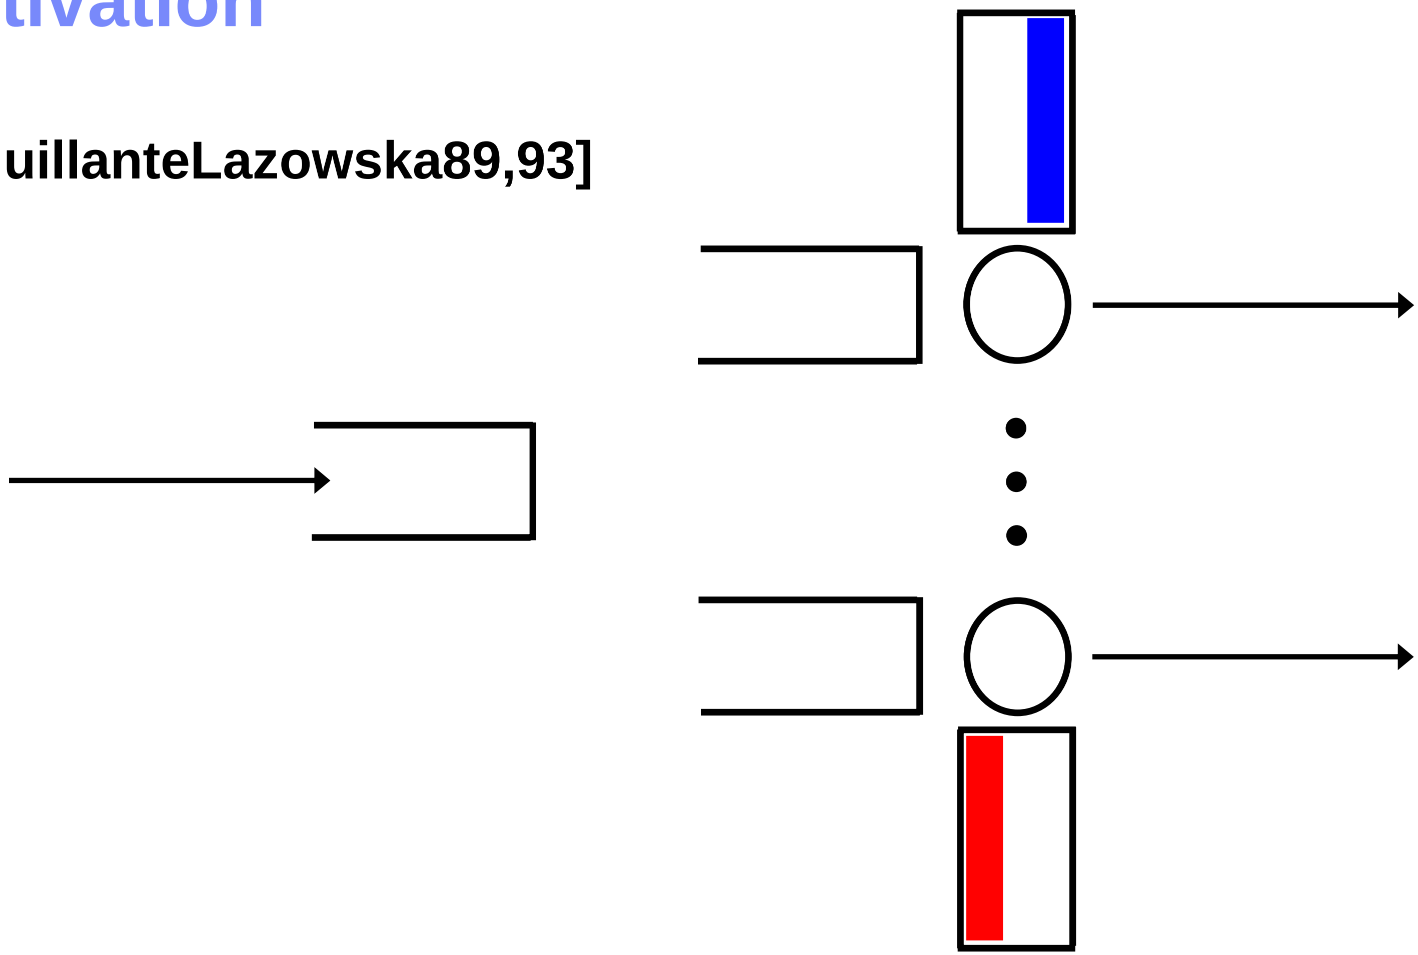

Key Points of Fundamental Tradeoff

- Customers can be served on any server of a parallel-server queueing system
- Each customer is served most efficiently on one the servers
- Load imbalance among queues occurs due to stochastic properties of system

General Overview

- Optimal Scheduling Policy
  - Fluid control problem: $c\mu$-type scheduling policy
  - Brownian control problem: dynamic threshold-type scheduling policy
  - Optimal threshold settings for dynamic scheduling policy depend upon stochastic properties and traffic intensity of system
- Analysis of Dynamic Threshold Scheduling
  - Consider generalized threshold scheduling policy
  - Matrix-analytic analysis and fix-point solution, asymptotically exact
  - Numerical experiments
  - Optimal settings of dynamic scheduling policy thresholds
- Stochastic Derivative-Free Optimization

Scheduling Policy

$$\beta \equiv (T_r^\ell, T_s^\ell, T_r^u, T_s^u), \;\; 0 \le \{T_r^\ell, T_s^\ell\} \le \{T_r^u, T_s^u\} < \infty$$
$$K \equiv \max\{T_r^\ell, T_s^\ell\}, \;\; \hat{T} \equiv \max\{T_r^u, T_s^u\}$$

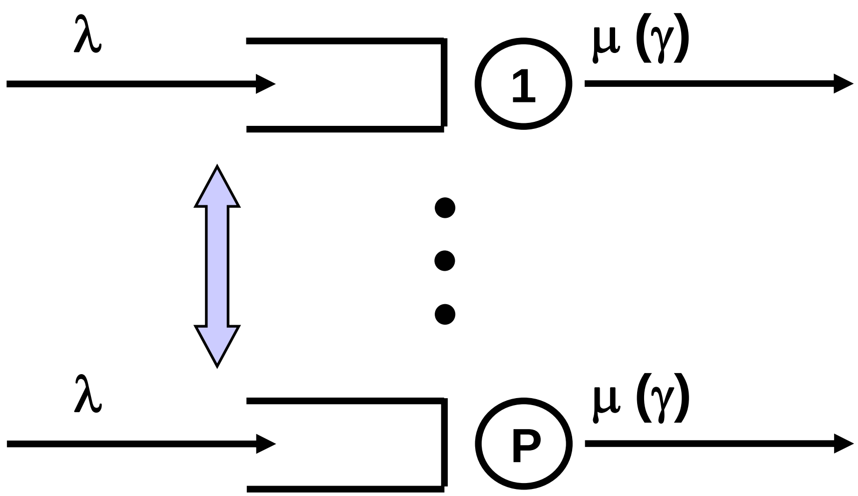

Scheduling Model

$$\beta \equiv (T_r^\ell, T_s^\ell, T_r^u, T_s^u), \;\; 0 \leq \{T_r^\ell, T_s^\ell\} \leq \{T_r^u, T_s^u\} < \infty$$

$$K \equiv \max\{T_r^\ell, T_s^\ell\}, \;\; \hat{T} \equiv \max\{T_r^u, T_s^u\}$$

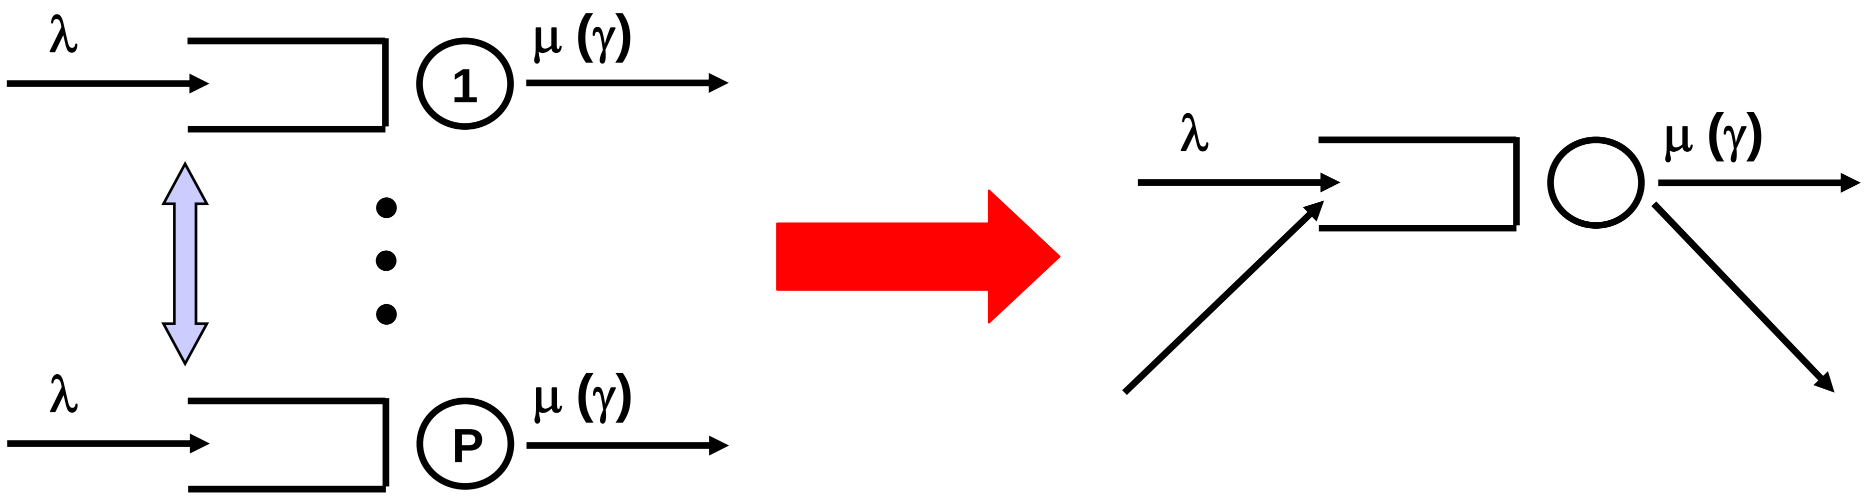

Background on Matrix-Analytic Methods

BD process

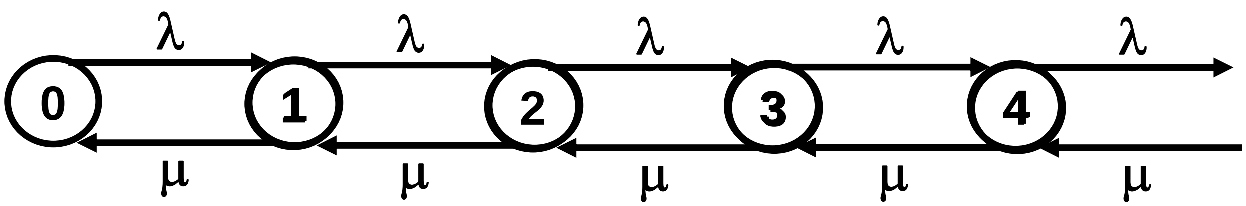

Define $\pi_n \equiv \lim_{t\to\infty} \mathrm{P}[N(t)=n], \quad \boldsymbol{\pi} \equiv (\pi_0, \pi_1, \pi_2, \ldots)$

$$\mathbf{Q} = \begin{bmatrix} -\lambda & \lambda & 0 & 0 & 0 & \ldots \\ \mu & -(\lambda+\mu) & \lambda & 0 & 0 & \ldots \\ 0 & \mu & -(\lambda+\mu) & \lambda & 0 & \ldots \\ 0 & 0 & \mu & -(\lambda+\mu) & \lambda & \ldots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix}$$

Background on Matrix-Analytic Methods

BD process

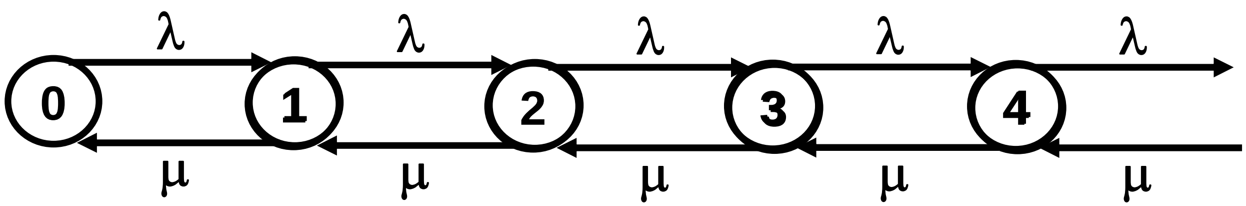

Then solution of $\boldsymbol{\pi}\mathbf{Q} = 0$ and $\boldsymbol{\pi}\mathbf{e} = 1$ given by

$$\pi_i = \pi_0 \rho^i$$

where

$$0 = \lambda - \rho(\lambda + \mu) + \rho^2\mu, \ \rho < 1 \ \Rightarrow \ \rho = \lambda/\mu < 1$$

$$\boldsymbol{\pi}\mathbf{e} = 1 \ \Rightarrow \ \pi_0 \sum_{i=0}^{\infty} \rho^i = 1 \ \Rightarrow \ \pi_0 = 1 - \rho$$

Background on Matrix-Analytic Methods

QBD process

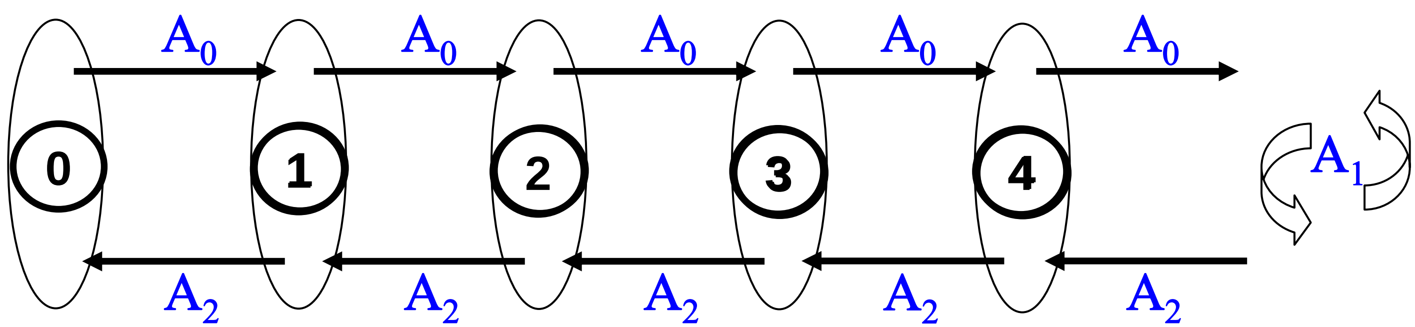

Define $\boldsymbol{\pi}_n \equiv (\pi_{n,1}, \ldots, \pi_{n,m}), \quad \boldsymbol{\pi} \equiv (\boldsymbol{\pi}_0, \boldsymbol{\pi}_1, \boldsymbol{\pi}_2, \ldots)$

$$\mathbf{Q} = \begin{bmatrix} \mathbf{B}_0 & \mathbf{A}_0 & \mathbf{0} & \mathbf{0} & \mathbf{0} & \cdots \\ \mathbf{A}_2 & \mathbf{A}_1 & \mathbf{A}_0 & \mathbf{0} & \mathbf{0} & \cdots \\ \mathbf{0} & \mathbf{A}_2 & \mathbf{A}_1 & \mathbf{A}_0 & \mathbf{0} & \cdots \\ \mathbf{0} & \mathbf{0} & \mathbf{A}_2 & \mathbf{A}_1 & \mathbf{A}_0 & \cdots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix}$$

Background on Matrix-Analytic Methods

QBD process

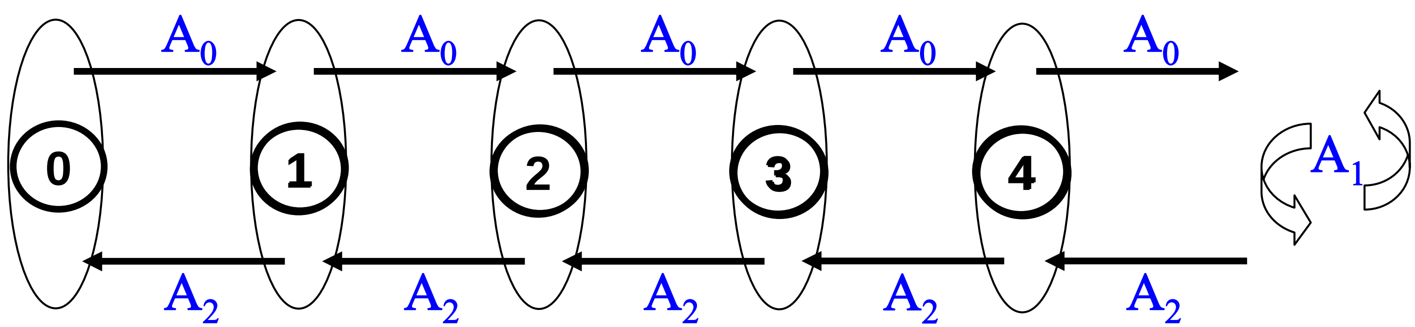

Then solution of $\pi Q = 0$ and $\pi e = 1$ given by

$$\pi_i = \pi_0 R^i$$

where

$$0 = A_0 + R A_1 + R^2 A_2, \quad \mathrm{sp}(R) < 1$$
$$0 = \pi_0 (B_0 + R A_2), \quad \pi_0 (I - R)^{-1} e = 1$$

Mathematical Analysis

$$T_r^\ell = T_s^\ell = K = 1,\ \ T_r^u < T_s^u = \widehat{T},\ \ \exp(\lambda),\ \exp(\mu),\ \exp(\gamma)$$

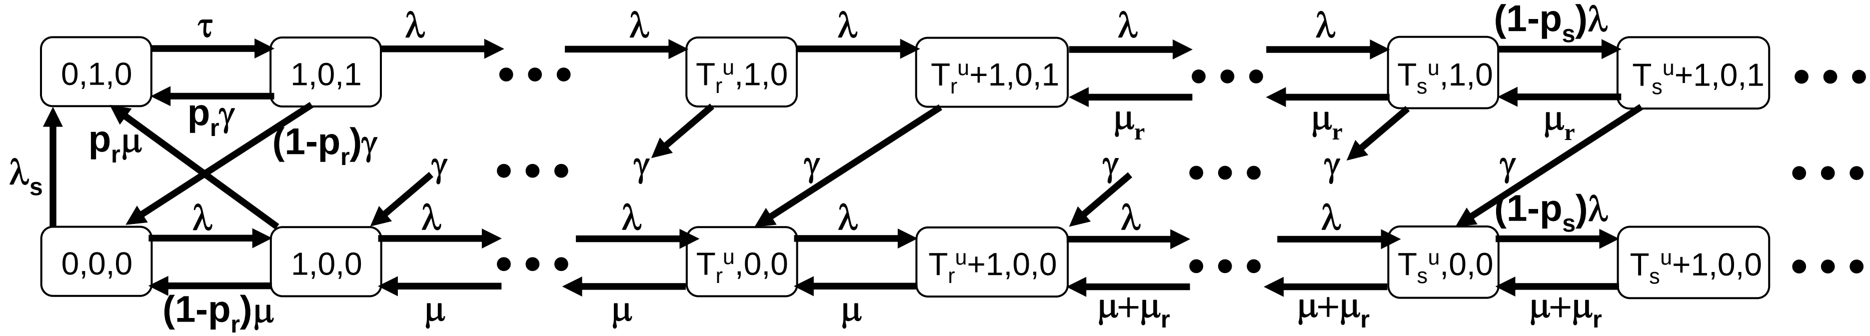

State Vector (i, j, v):

- i: total number of customers waiting or receiving service at the server of interest
- j: number of customers in the process of being migrated to the server of interest
- v: K-bit binary vector denoting customer type of up to the first K customers at the server

Mathematical Analysis

$$T_r^\ell = T_s^\ell = K = 1, \quad T_r^u < T_s^u = \hat{T}, \quad \exp(\lambda),\ \exp(\mu),\ \exp(\gamma)$$

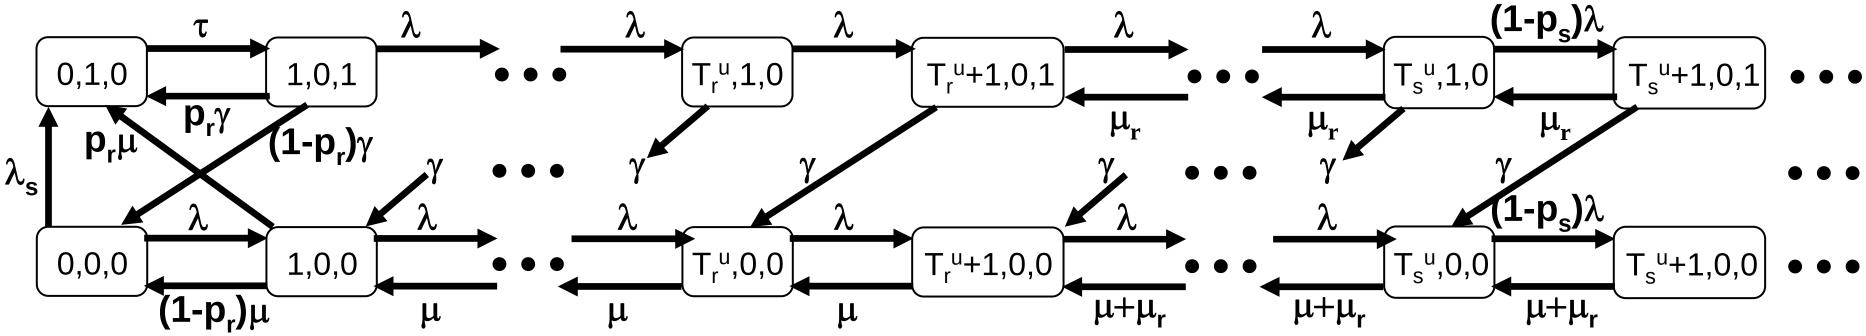

Define $\boldsymbol{\pi}_0 \equiv (\pi_{0,0,0}, \pi_{0,1,0}), \quad \boldsymbol{\pi}_n \equiv (\pi_{n,0,0}, \pi_{n,0,1}),\ n > 0$

$$\mathbf{Q} = \begin{bmatrix} \mathbf{B}_{00} & \mathbf{B}_{01} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \cdots \\ \mathbf{B}_{10} & \mathbf{B}_{11} & \mathbf{A}_0 & \mathbf{0} & \mathbf{0} & \cdots \\ \mathbf{0} & \mathbf{A}_2 & \mathbf{A}_1 & \mathbf{A}_0 & \mathbf{0} & \cdots \\ \mathbf{0} & \mathbf{0} & \mathbf{A}_2 & \mathbf{A}_1 & \mathbf{A}_0 & \cdots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix}$$

Mathematical Analysis

$$T_r^\ell = T_s^\ell = K = 1, \quad T_r^u < T_s^u = \hat{T}, \quad \exp(\lambda), \; \exp(\mu), \; \exp(\gamma)$$

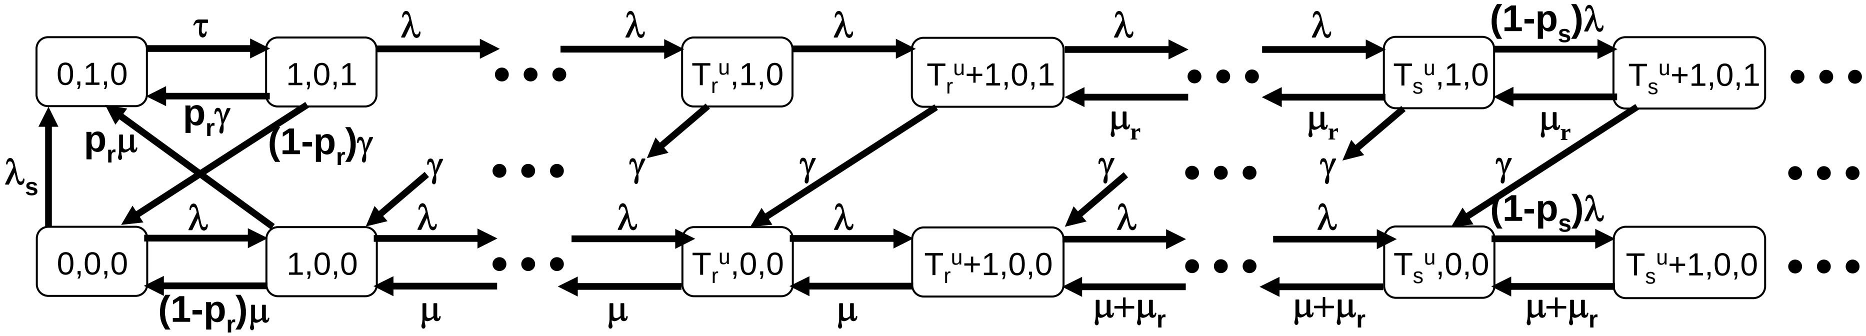

Then solution of $\boldsymbol{\pi}\mathbf{Q} = 0, \; \boldsymbol{\pi}\mathbf{e} = 1$ is given by

$$\boldsymbol{\pi}_{\hat{T}+k} = \boldsymbol{\pi}_{\hat{T}} \mathbf{R}^k, \quad k \geq 0$$

$$\mathbf{0} = \mathbf{A}_0 + \mathbf{R}\mathbf{A}_1 + \mathbf{R}^2\mathbf{A}_2, \quad \mathrm{sp}(\mathbf{R}) < 1$$

$$\mathbf{0} = (\boldsymbol{\pi}_0, \boldsymbol{\pi}_1, \ldots, \boldsymbol{\pi}_{\hat{T}}) \begin{bmatrix} \mathbf{B}_{00} & \mathbf{B}_{01} \\ \mathbf{B}_{10} & \mathbf{B}_{11} + \mathbf{R}\mathbf{A}_2 \end{bmatrix}$$

$$1 = (\boldsymbol{\pi}_0, \boldsymbol{\pi}_1, \ldots, \boldsymbol{\pi}_{\hat{T}-1})\mathbf{e} + \boldsymbol{\pi}_{\hat{T}}(\mathbf{I} - \mathbf{R})^{-1}\mathbf{e}$$

Mathematical Analysis

$$T_r^\ell = T_s^\ell = K = 1, \quad T_r^u < T_s^u = \hat{T}, \quad \exp(\lambda),\ \exp(\mu),\ \exp(\gamma)$$

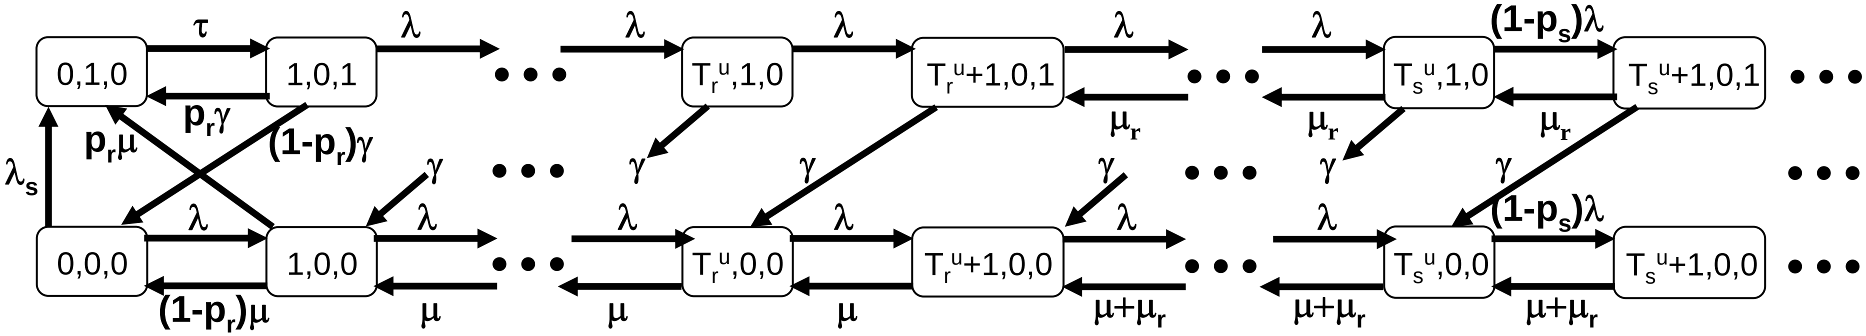

Probabilistic Interpretation of $\mathbf{R}$:

$$r_{00} = \frac{\lambda(1-p_s)}{\mu+\mu_r}, \qquad r_{01} = 0$$

$$r_{11} = \frac{\lambda(1-p_s)+\gamma+\mu_r - ((\lambda(1-p_s)+\gamma+\mu_r)^2 - 4\lambda(1-p_s)\mu_r)^{1/2}}{2\mu_r}$$

$$r_{10} = \frac{{r_{11}}^2\gamma}{(\mu+\mu_r)(1-r_{11})}$$

Mathematical Analysis

$$T_r^{\ell} = T_s^{\ell} = K = 1, \quad T_r^u < T_s^u = \hat{T}, \quad \exp(\lambda), \ \exp(\mu), \ \exp(\gamma)$$

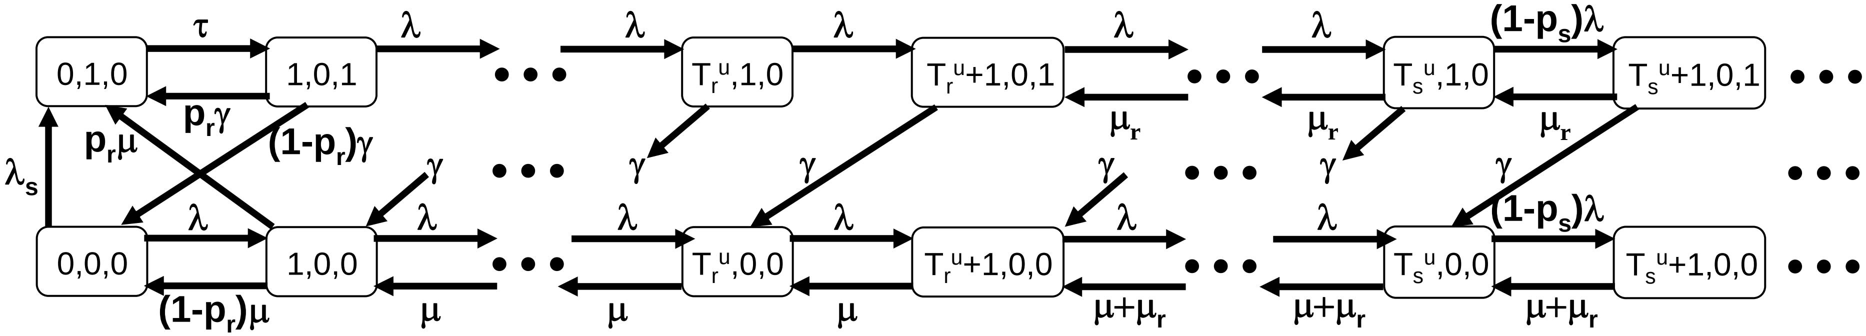

Let I, J, $\underline{V}$ be generic r.v.s of the process

$$\begin{aligned} p_s &= \mathsf{P}\left[\text{Find Queue with } I + J < T_s^{\ell} \text{ in } \leq L_p \text{ Random Probes}\right] \\ &= 1 - \left(1 - \mathsf{P}\left[I + J < T_s^{\ell}\right]\right)^{L_p} \end{aligned}$$

Equating inflow and outflow of SI migrated customers:

$$\lambda_s \mathsf{P}\left[I + J < T_s^{\ell}\right] = \lambda p_s \mathsf{P}\left[I \geq T_s^u\right]$$

Mathematical Analysis

$$T_r^\ell = T_s^\ell = K = 1, \quad T_r^u < T_s^u = \widehat{T}, \quad \exp(\lambda), \ \exp(\mu), \ \exp(\gamma)$$

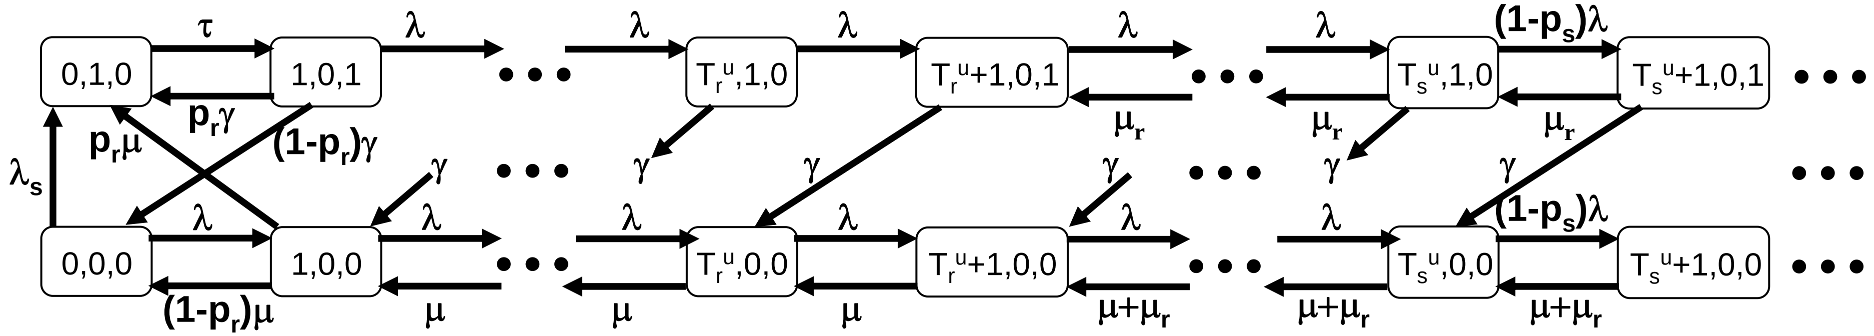

Let I, J, $\underline{V}$ be generic r.v.s of the process

$$\begin{aligned} p_r &= \mathsf{P}\,[\,\text{Find Queue with } I > T_r^u \text{ in } \le L_p \text{ Random Probes}\,] \\ &= 1 - (1 - \mathsf{P}\,[\,I > T_r^u\,])^{L_p} \end{aligned}$$

Equating inflow and outflow of RI migrated customers:

$$\begin{aligned} \mu_r \mathsf{P}\,[\,I > T_r^u\,] &= \mu p_r \mathsf{P}\left[\,I + J \le T_r^\ell,\ I > 0,\ v_1 = 0\,\right] + \\ &\quad \gamma p_r \mathsf{P}\left[\,I + J \le T_r^\ell,\ I > 0,\ v_1 = 1\,\right] \end{aligned}$$

Mathematical Analysis

$$T_r^{\ell} = T_s^{\ell} = K = 1, \quad T_r^{u} < T_s^{u} = \hat{T}, \quad \exp(\lambda), \; \exp(\mu), \; \exp(\gamma)$$

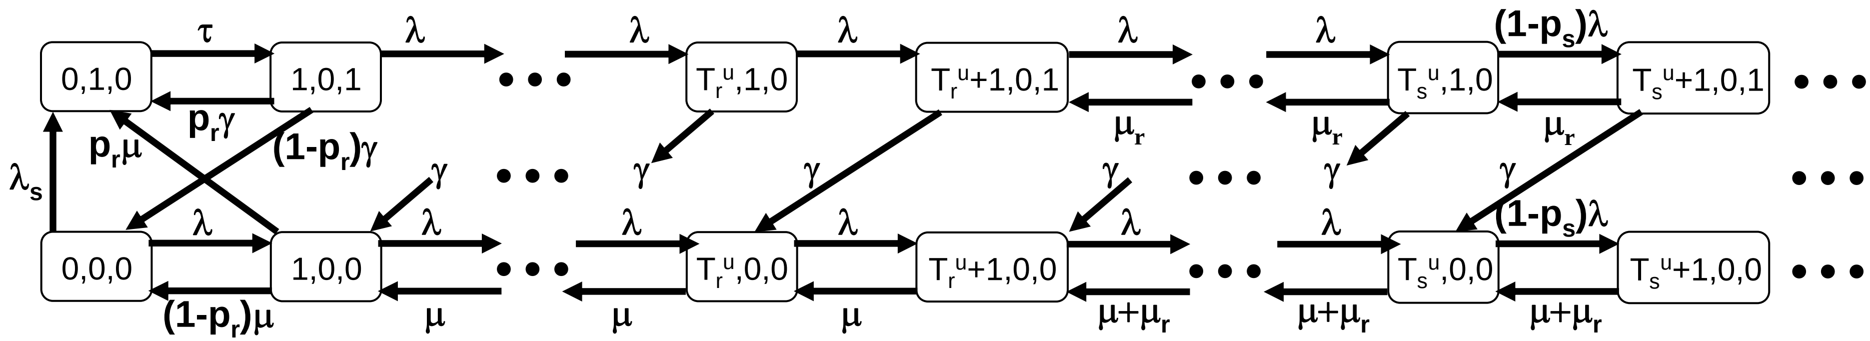

Final Solution Obtained via Fix-Point Iteration

1. Initialize $(\mathbf{p_s}, \lambda_s, \mathbf{p_r}, \mu_r)$
2. Compute stationary probability vector in terms of $(\mathbf{p_s}, \lambda_s, \mathbf{p_r}, \mu_r)$
3. Compute new values of $(\mathbf{p_s}, \lambda_s, \mathbf{p_r}, \mu_r)$ in terms of stationary vector
4. Goto 2 until differences between iteration values are arbitrarily small

Mathematical Analysis

$$T_r^\ell = T_s^\ell = K = 1, \quad T_r^u < T_s^u = \widehat{T}, \quad \mathsf{exp}(\lambda),\ \mathsf{exp}(\mu),\ \mathsf{exp}(\gamma)$$

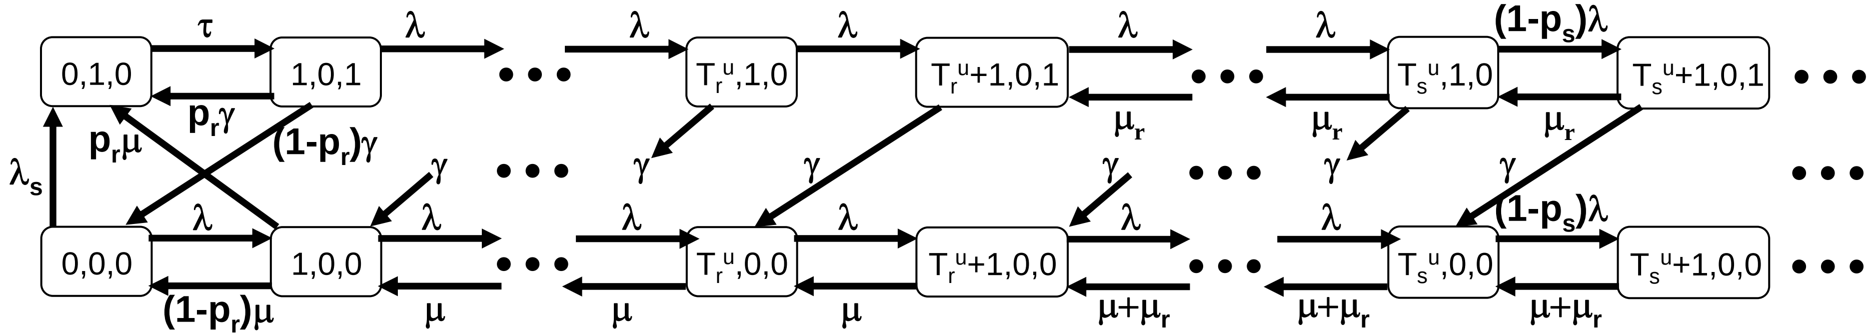

Stability criterion: $\mathsf{sp}(\mathbf{R}) = r_{00} < 1, \ \mathbf{A} \equiv \mathbf{A}_0 + \mathbf{A}_1 + \mathbf{A}_2$ reducible

$$\mathsf{E}[\mathcal{T}] = \lambda^{-1}\left\{\sum_{k=1}^{\widehat{T}-1} k\boldsymbol{\pi}_k\mathbf{e} + \widehat{T}\boldsymbol{\pi}_{\widehat{T}}(\mathbf{I}-\mathbf{R})^{-1}\mathbf{e} + \boldsymbol{\pi}_{\widehat{T}}(\mathbf{I}-\mathbf{R})^{-2}\mathbf{R}\mathbf{e}\right\}$$

$$p_{mt} = (\boldsymbol{\pi}_1,\ldots,\boldsymbol{\pi}_K)\nu + \sum_{k=K+1}^{\widehat{T}-1}\boldsymbol{\pi}_k\nu' + \boldsymbol{\pi}_{\widehat{T}}(\mathbf{I}-\mathbf{R})^{-1}\nu'$$

Mathematical Analysis

$$T_r^\ell = T_s^\ell = K = 1, \quad T_r^u < T_s^u = \widehat{T}, \quad \exp(\lambda),\ \exp(\mu),\ \exp(\gamma)$$

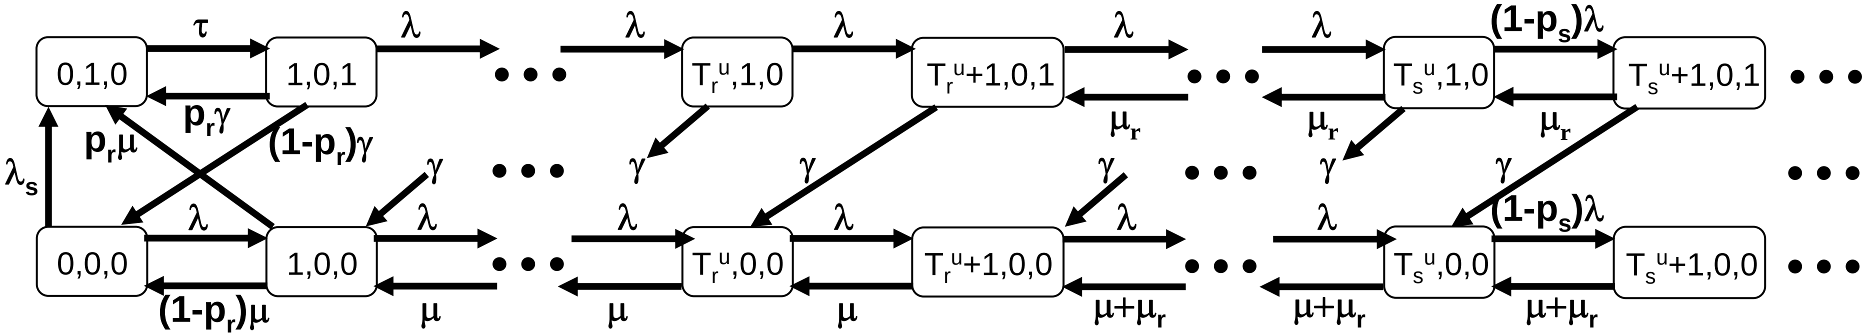

Tail asymptotics:

$$\mathsf{P}[Q > x] = \frac{\pi_{\widehat{T}}\mathbf{v}}{1-\eta}\eta^x + o(\eta^x), \quad \text{as } x \to \infty,$$

$$\mathsf{P}[Q > x] \sim \frac{\pi_{\widehat{T}}\mathbf{v}}{1-\eta}\eta^x, \quad \text{as } x \to \infty,$$

where $\mathbf{u}$ and $\mathbf{v}$ are the left and right eigenvectors corresponding to $\eta = r_{00}$ normalized by $\mathbf{ue} = 1$ and $\mathbf{uv} = 1$

Mathematical Analysis

General K, Pure Sender-Initiated Policy

Probabilistic Interpretation of $\mathbf{R}$:

$$
\begin{aligned}
r_{00} &= \lambda(1-p_s)/\mu \\
r_{nm} &= 0, \qquad 0 \le n < m \le 2^K - 1 \qquad \text{(lower triangular)} \\
r^{\mathbf{0}}_{nn} &= \Phi(n), \qquad 1 \le n \le 2^K - 1 \\
r^{\mathbf{0}}_{nm} &= C_k \Phi(n)^{k+1} \Psi(n)^k, \qquad (*) \\
& \qquad n > m > 0 \text{ s.t. } b_m = \mathcal{D}_k(b_n),\ 1 \le k < K \\
r_{n0} &= \frac{\gamma \sum_{k=1}^{n} r_{nk} r_{k1} + \mu \sum_{k=1}^{n-1} r_{nk} r_{k0}}{\lambda(1-p_s) + \mu(1 - r_{00} - r_{nn})}, \\
& \qquad 1 \le n \le 2^K - 1
\end{aligned}
$$

Mathematical Analysis

General K

Departure Transitions from State $(i, j, \underline{v})$, $i > 0$, $j \le (K-i)^+$	
To State	Rate
$(i-1, j, \mathcal{D}_1(\underline{v}))$, $i + j \le T_r^\ell \le K$	$(1-p_r)e(\underline{v})$
$(i-1, j+1, \mathcal{D}_1(\underline{v}))$, $i + j \le T_r^\ell \le K$	$p_r e(\underline{v})$
$(i-1, j, \mathcal{D}_1(\underline{v}))$, $T_r^\ell < i + j \le T_r^u$	$e(\underline{v})$
$(i-1, j, \underline{v})$, $\underline{v} \neq 0$, $i > T_r^u$	μ_r
$(i-1, j, \mathcal{D}_1(\underline{v})), \underline{v} \neq 0$, $i > T_r^u$	$e(\underline{v})$
$(i-1, j, \mathcal{D}_1(\underline{v}))$, $\underline{v} = 0$, $i > T_r^u$	$e(\underline{v}) + \mu_r$

$$\mathcal{D}_k(\underline{v}) = (v_{k+1}, \ldots, v_K, \underbrace{0, \ldots, 0}_{k \text{ times}})$$

Mathematical Analysis

General K

Arrival Transitions from State $(i, j, \underline{v}),\ i \geq 0,\ j \leq (K-i)^+$	
To State	Rate
$(i+1, j-1, \mathcal{A}_{i+1}(\underline{v})),\ j > 0$	$j\theta$
$(i, j+1, \underline{v}),\ i+j < T_s^{\ell} \leq K$	λ_s
$(i+1, j, \underline{v}),\ i < K,\ i+j < K$	λ
$(i+1, j, \underline{v}),\ K \leq i < T_s^u$	λ
$(i+1, j, \underline{v}),\ i \geq T_s^u$	$\lambda(1-p_s)$

$$\mathcal{A}_i(\underline{v}) = (v_1, v_2, \ldots, v_{i-1}, 1, \underbrace{0, 0, \ldots, 0}_{K-i \text{ times}}).$$

Numerical Results

Optimal Threshold Values												
	$1/\gamma = 1.6$				$1/\gamma = 2.4$				$1/\gamma = 3.2$			
λ	$\widehat{T}_r^\ell$	$\widehat{T}_s^\ell$	$\widehat{T}_r^u$	$\widehat{T}_s^u$	$\widehat{T}_r^\ell$	$\widehat{T}_s^\ell$	$\widehat{T}_r^u$	$\widehat{T}_s^u$	$\widehat{T}_r^\ell$	$\widehat{T}_s^\ell$	$\widehat{T}_r^u$	$\widehat{T}_s^u$
0.04	1	1	1	1	0	1	-	2	0	1	-	3
0.08	1	1	1	1	1	1	2	2	1	1	3	3
0.12	1	1	1	1	1	1	2	2	1	1	4	4
0.16	1	1	1	1	1	1	2	2	1	1	4	4
0.20	1	1	1	1	1	1	3	3	1	1	4	4
0.24	1	1	1	1	1	1	3	3	1	1	5	5
0.28	1	1	1	1	1	1	3	3	1	1	5	5
0.32	1	1	1	1	1	1	3	3	1	1	5	5
0.36	1	1	1	1	1	1	3	3	1	1	5	5
0.40	1	1	1	1	1	1	3	3	1	1	6	6
0.44	1	1	1	1	1	1	3	3	1	1	6	6
0.48	1	1	1	1	1	1	4	4	1	1	6	6
0.52	1	1	1	1	1	1	4	4	1	1	7	7
0.56	1	1	1	1	1	1	4	4	1	1	7	7
0.60	1	1	1	1	1	1	4	4	1	1	8	8
0.64	1	1	1	1	1	1	4	4	1	1	8	8
0.68	1	1	2	2	1	1	4	4	1	1	8	8
0.72	1	1	2	2	1	1	5	5	1	1	9	9
0.76	1	1	2	2	1	1	5	5	1	1	9	9
0.80	1	1	2	2	1	1	5	5	1	1	10	10
0.84	1	1	3	3	1	1	6	6	1	1	10	10
0.88	1	1	3	3	1	1	7	7	1	1	12	12
0.92	1	1	4	4	1	1	9	9	1	1	12	12
0.96	1	1	12	12	1	1	12	12	1	1	12	12
0.98	1	1	12	12	1	1	12	12	1	1	12	12

Numerical Results

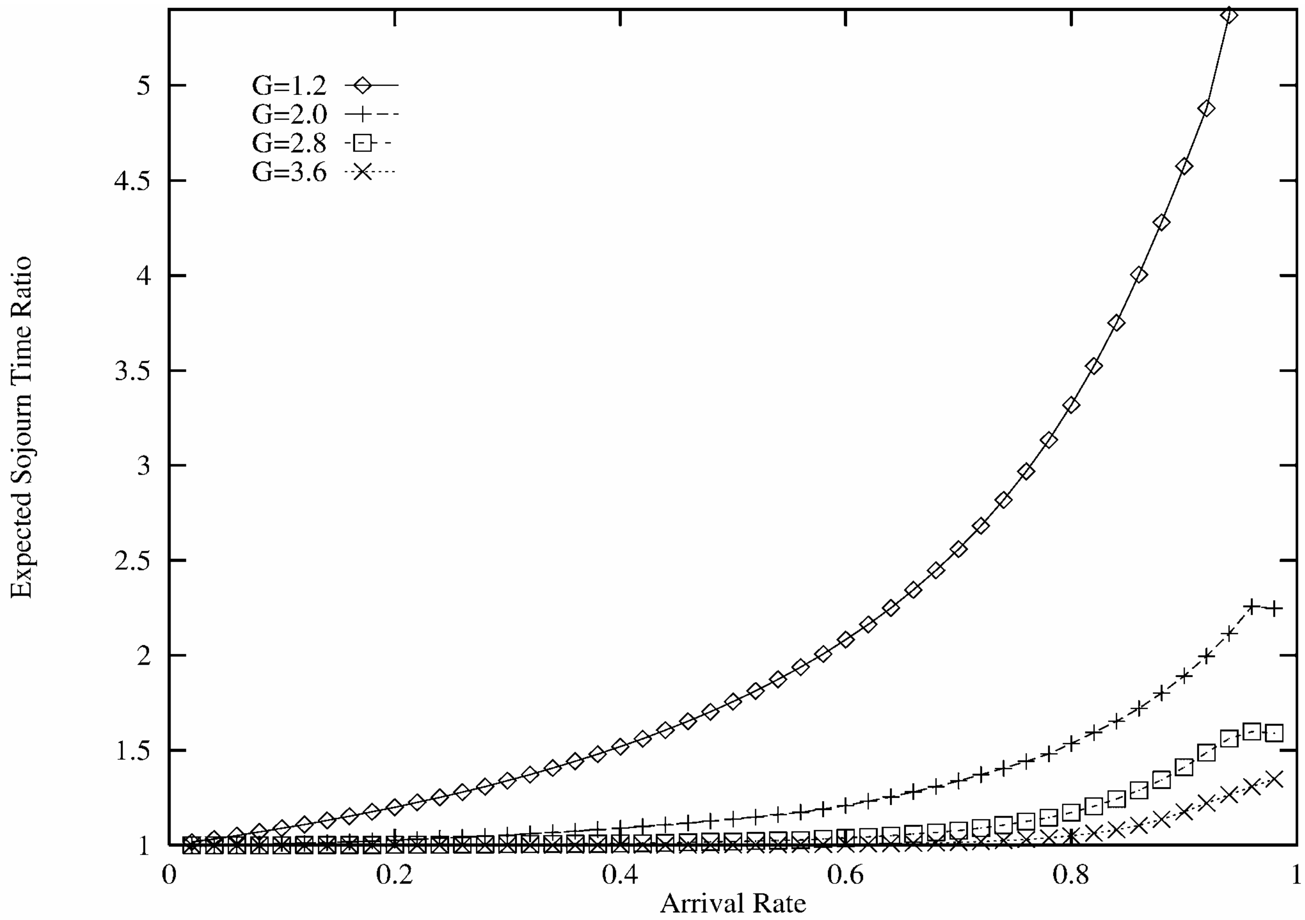

Relative Performance Benefits of the Migratory Scheduling Policy with Optimal Thresholds over the No-Migration Policy

Numerical Results

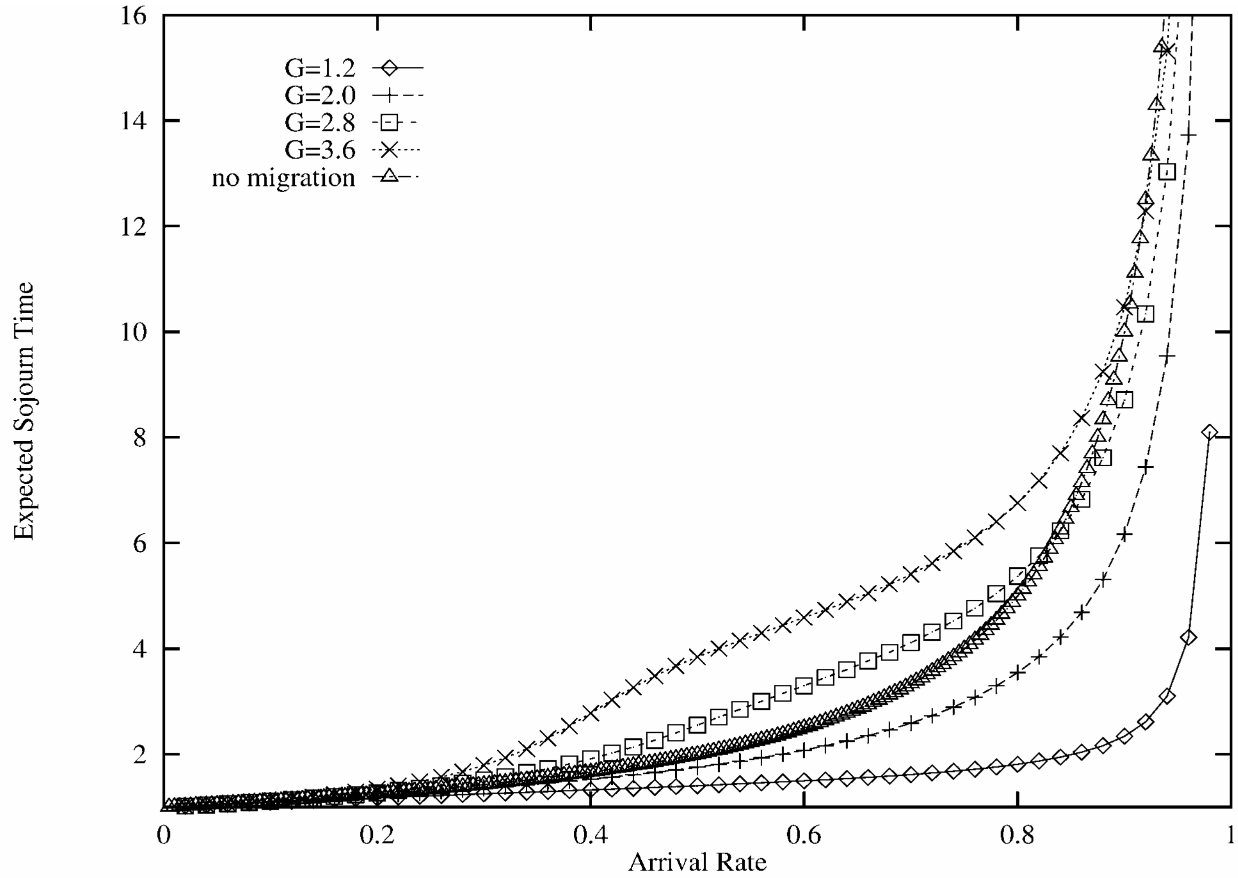

Expected Sojourn Times for the No-Migration Policy and the Migratory Scheduling Policy with Thresholds $\beta = (T_r^\ell, T_s^\ell, T_r^u, T_s^u) = (1, 1, 2, 2)$

Numerical Results

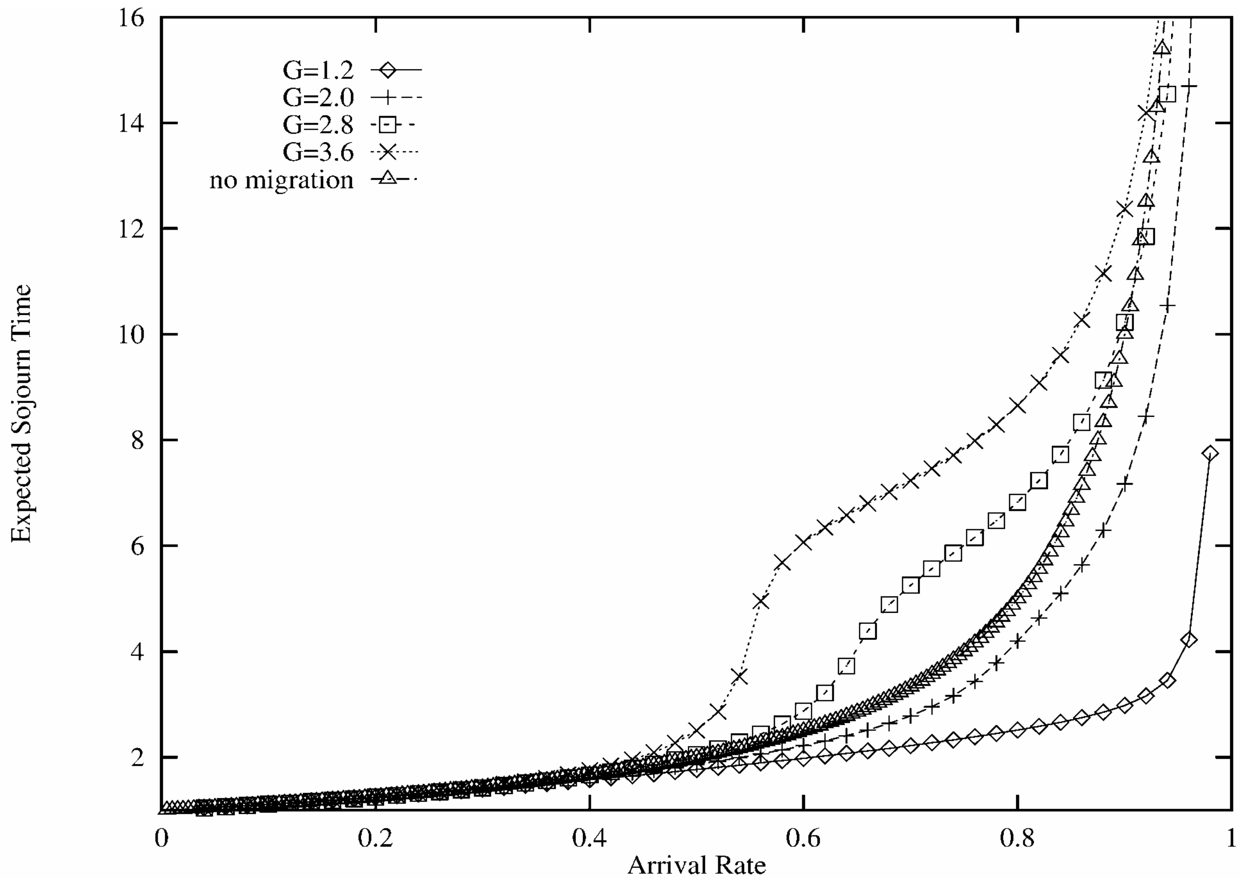

Expected Sojourn Times for the No-Migration Policy and the Migratory Scheduling Policy with Thresholds $\beta = (T_r^{\ell}, T_s^{\ell}, T_r^{u}, T_s^{u}) = (2, 2, 4, 4)$

Numerical Results

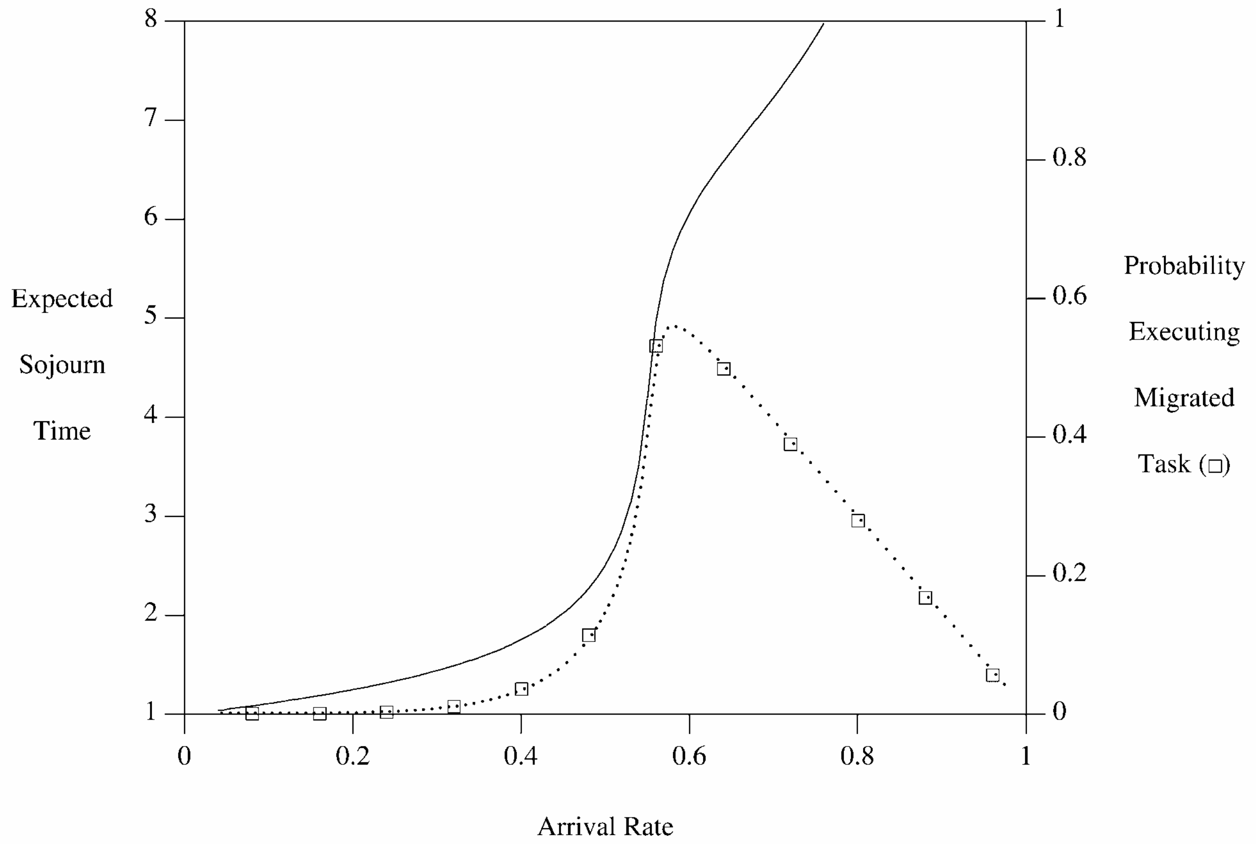

Expected Sojourn Time, and Probability of Executing a Migrated Task, for the Migratory Scheduling Policy with Thresholds $\beta = (T_r^\ell, T_s^\ell, T_r^u, T_s^u) = (2, 2, 4, 4)$

Stochastic Derivative-Free Optimization

- Internal Model
- External Model
- Trust Region

General Overview

- Optimal Scheduling Policy
  - Fluid control problem: $c\mu$ type scheduling policy
  - Brownian control problem: dynamic threshold-type scheduling policy
- Analysis of Dynamic Threshold Scheduling
  - Consider generalized threshold scheduling policy
  - Matrix-analytic analysis and fix-point solution, asymptotically exact
  - Numerical experiments
  - Optimal settings of dynamic scheduling policy thresholds
- Stochastic Derivative-Free Optimization