Imperial College London

Fluid Approximation of Closed Queueing Networks with Discriminatory Processor Sharing

Lulai Zhu[1], Giuliano Casale[1] and Iker Perez[2]
[1] Imperial College London, UK
[2] University of Nottingham, UK

DPS Discipline

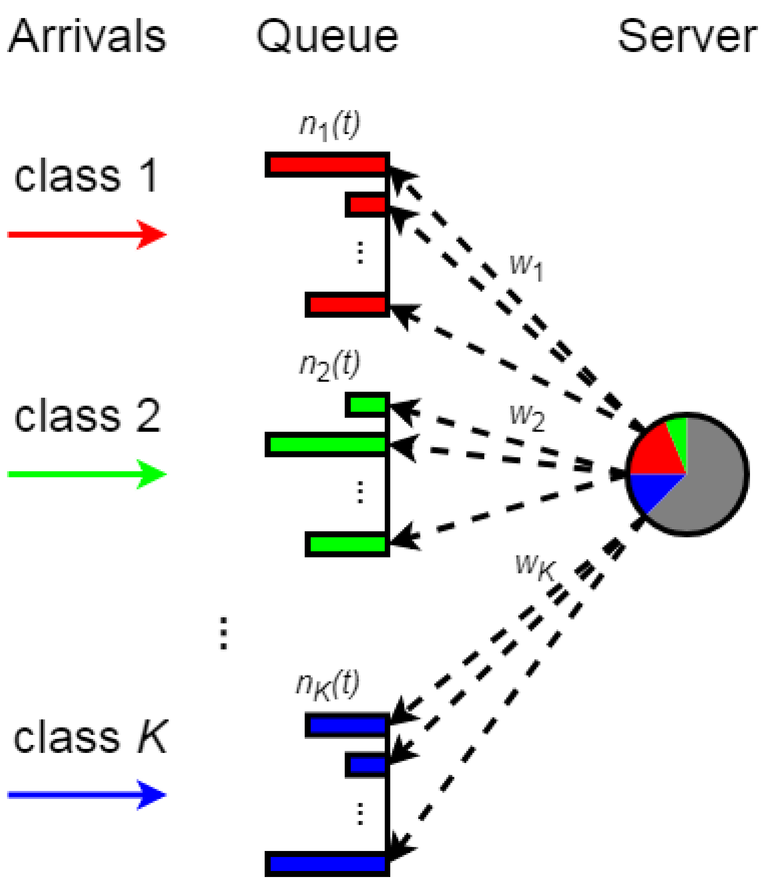

DPS Queueing System

Instantaneous fraction of the service capacity allotted to a class-k job:

$$\frac{w_k}{\sum_{l=1}^{K} w_l n_l(t)}$$

Imperial College London

Related Work

- **Heavy-Traffic Regime:**
  - *A Closed Network with a Discriminatory Processor-Sharing Server* [Mitra & Weiss, 1989]

$$\frac{dx_k(t)}{dt} = \lambda_k\big(b_k - x_k(t)\big) - \frac{\lambda_k g_k w_k x_k(t)}{\sum_{l=1}^{K} w_l x_l(t)}$$

- **Moderately Heavy-Traffic Regime:**
  - *Asymptotic analysis of a large closed queueing network with discriminatory processor sharing* [Morrison, 1991]

$$\sum_{l=1}^{K} \frac{b_k}{g_k} = 1 - \frac{a}{\sqrt{N}}$$

where $a = O(1)$ as $N \to \infty$

Model Definition

M: Number of stations

K: Number of classes

N: Total population of jobs

$\mathcal{D}$: Set of delay stations

$\mathcal{Q}$: Set of queueing stations

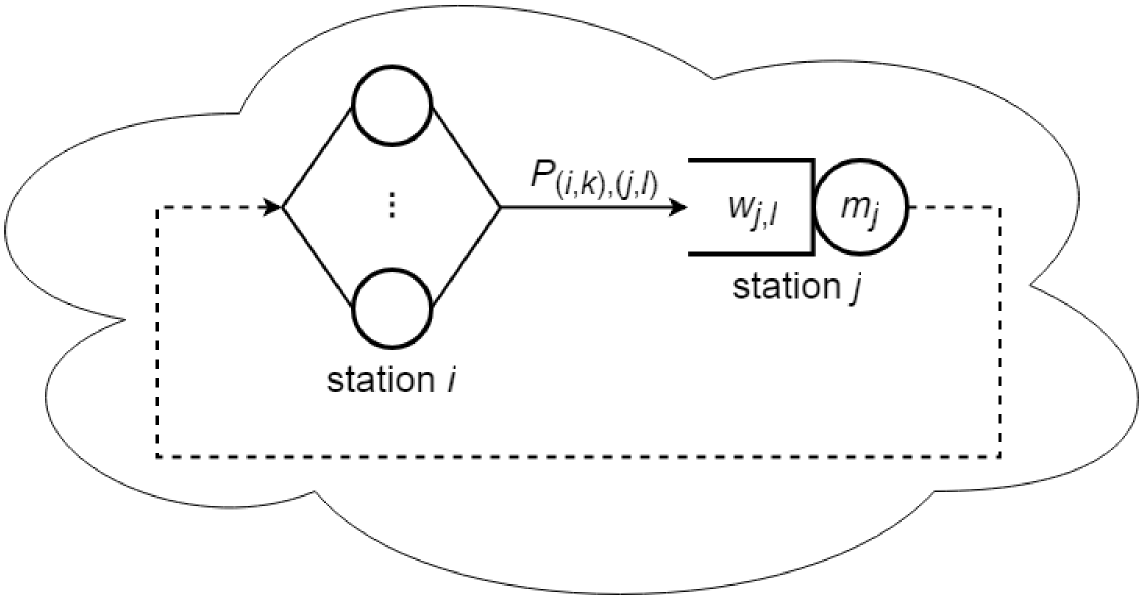

Reference Model

Imperial College London

Model Definition

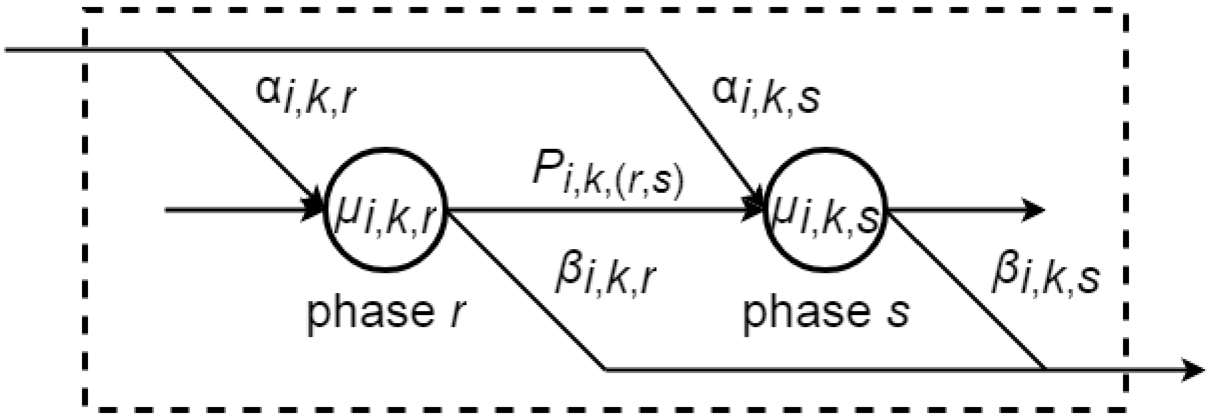

Service Time Distribution

- Probability that a class-k job in phase r in station i proceeds to station j as a class-l job in phase s:

$$P_{(i,k,r),(j,l,s)} = \begin{cases} P_{i,k,(r,s)} + \beta_{i,k,r} P_{(i,k)(i,k)}\, \alpha_{i,k,s} & \text{if } i = j, k = l \\ \beta_{i,k,r} P_{(i,k),(j,l)}\, \alpha_{j,l,r} & \text{otherwise} \end{cases}$$

where

$$\beta_{i,k,r} = 1 - \sum_{s=1}^{R_{i,k}} P_{i,k,(r,s)}$$

Imperial College London

Fluid Approximation

- **Standard Markovian Technique:** State space explosion
- **Fluid Approximation:** Discontinuity of the ODE system
  - Preload a self-looping job into every queueing station:

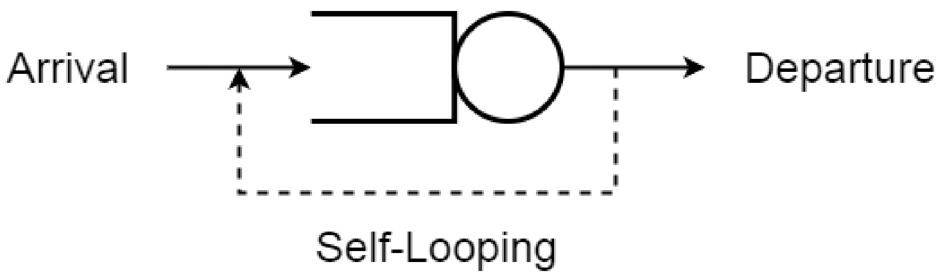

  - Assign the self-looping job in each queueing station i a service weight:

$$w_{i,0} = \frac{\sum_{k=1}^{K} w_{i,k}}{K}$$

Fluid Approximation

- Constraints on the service weight of the self-looping job:

$$\begin{cases} w_{i,0} \geq 0 & \text{if } \forall k \in \mathbb{N}^*_{\leq K}, w_{i,k} \geq 0 \\ w_{i,0} = a & \text{if } \forall k \in \mathbb{N}^*_{\leq K}, w_{i,k} = a \end{cases}$$

  where $a \in \mathbb{R}_{\geq 0}$ is arbitrary
- Assume the model to be a multi-chain QN, where no transient classes would possibly exist

Fluid Approximation

- State of the model:

$$\boldsymbol{n}(t) = \left(n_{i,k,r}(t)\right) \in \left\{\boldsymbol{n} \in \mathbb{N}^R : \sum_{i=1}^{M}\sum_{k=1}^{K}\sum_{r=1}^{R_{i,k}} n_{i,k,r} = N\right\}$$

where

$$R = \sum_{i=1}^{M}\sum_{k=1}^{K} R_{i,k}$$

- Set of possible state changes in the model:

$$L = \left\{-\boldsymbol{e}_{i,k,r} + \boldsymbol{e}_{j,l,s} : i,j \in \mathbb{N}^*_{\leq M}; k,l \in \mathbb{N}^*_{\leq K}; r \in \mathbb{N}^*_{\leq R_{i,k}}; s \in \mathbb{N}^*_{\leq R_{j,l}}\right\}$$

where $\boldsymbol{e}_{i,k,r}$ is a vector of all zeros except for a one in the (i,k,r)-th entry and $\boldsymbol{e}_{j,l,s}$ is similarly defined

Fluid Approximation

- Transition rate of the model:

$$q_{\boldsymbol{n},\boldsymbol{n}+\boldsymbol{l}} = f(\boldsymbol{n},\boldsymbol{l})$$

where $f : \mathbb{R}^R \times \mathbb{Z}^R \to \mathbb{R}_{\geq 0}$ is a function defined by

$$f(\boldsymbol{x},\boldsymbol{l}) = \begin{cases} \mu_{i,k,r} P_{(i,k,r),(j,l,s)} x_{i,k,r} & \text{if } \boldsymbol{l} \in L, i \in \mathcal{D} \\ \mu_{i,k,r} P_{(i,k,r),(j,l,s)} \dfrac{m_i w_{i,k} x_{i,k,r}}{\sum_{k'=1}^{K} w_{i,k'} \left(\sum_{r'=1}^{R_{i,k'}} x_{i,k',r'} + \frac{1}{K} \right)} & \text{if } \boldsymbol{l} \in L, i \in \mathcal{Q} \\ 0 & \text{otherwise} \end{cases}$$

- The function $f(\boldsymbol{x},\boldsymbol{l})$ is continuous, bounded and Lipschitz continuous on any non-empty compact and convex subset of its global domain

$$D = \left\{ \boldsymbol{x} \in \mathbb{R}^R : \forall i \in \mathcal{Q}, \sum_{k'=1}^{K} w_{i,k'} \left(\sum_{r'=1}^{R_{i,k'}} x_{i,k',r'} + 1/K \right) \neq 0 \right\}$$

Fluid Approximation

Lemma 1. Let $\{\boldsymbol{n}^{[v]}(t) : v \in \mathbb{N}^*\}$ be the sample paths of an infinite sequence of QN models whose v-th element is the original model scaled with a total population of vN jobs as well as vm_i servers and v self-looping jobs in each queueing station i. The sequence $\{\boldsymbol{n}^{[v]}(t) : v \in \mathbb{N}^*\}$ consistutes a density-dependent family of CTMCs.

Fluid Approximation

Lemma 2. Define a vector field $\boldsymbol{F} : \mathbb{R}^R \rightarrow \mathbb{R}^R$ by

$$\boldsymbol{F}(\boldsymbol{x}) = \sum_{l} \boldsymbol{l} f(\boldsymbol{x}, \boldsymbol{l})$$

Specifically,

$$\begin{aligned} &F_{i,k,r}(\boldsymbol{x}) \\ &= \sum_{j=1}^{M} \sum_{l=1}^{K} \sum_{s=1}^{R_{j,l}} \Big(f\big(\boldsymbol{x}, -\boldsymbol{e}_{j,l,s} + \boldsymbol{e}_{i,k,r}\big) \\ &- f\big(\boldsymbol{x}, -\boldsymbol{e}_{i,k,r} + \boldsymbol{e}_{j,l,s}\big)\Big) \end{aligned}$$

The initial value problem (IVP)

$$\begin{cases} \dfrac{d\boldsymbol{x}(t)}{dt} = \boldsymbol{F}\big(\boldsymbol{x}(t)\big) \\ \boldsymbol{x}(0) = \boldsymbol{n}(0) \end{cases}$$

has a unique solution $\boldsymbol{x}(t)$ in any finite time horizon *T*.

Fluid Approximation

Theorem 1. Given the same initial state

$$\frac{\boldsymbol{n}^{[v]}(0)}{v} = \boldsymbol{n}(0)$$

the normalized sequence $\left\{\boldsymbol{n}^{[v]}(t)/v : v \in \mathbb{N}^*\right\}$ converges to the solution $\boldsymbol{x}(t)$ to the IVP

$$\begin{cases} \dfrac{d\boldsymbol{x}(t)}{dt} = \boldsymbol{F}\big(\boldsymbol{x}(t)\big) \\ \boldsymbol{x}(0) = \boldsymbol{n}(0) \end{cases}$$

in the sense that

$$\lim_{v \to \infty} \mathbb{P}\left(\sup_{t \in [0,T]} \left\| \frac{\boldsymbol{n}^{[v]}(t)}{v} - \boldsymbol{x}(t) \right\|_1 > \delta \right) = 0 \quad \text{for } \delta \in \mathbb{R}_{>0}$$

Fluid Approximation

Corollary 1. Suppose that the condition of Theorem 1 is satisfied. The expectation of the normalized sequence $\{\boldsymbol{n}^{[v]}(t)/v : v \in \mathbb{N}^*\}$ converges to the solution $\boldsymbol{x}(t)$ to the IVP

$$\begin{cases} \dfrac{d\boldsymbol{x}(t)}{dt} = \boldsymbol{F}\big(\boldsymbol{x}(t)\big) \\ \boldsymbol{x}(0) = \boldsymbol{n}(0) \end{cases}$$

in any finite time horizon T:

$$\lim_{v\to\infty} \mathbb{E}\left(\frac{\boldsymbol{n}^{[v]}(t)}{v}\right) = \boldsymbol{x}(t) \quad \text{for } t \in [0, T]$$

Steady-State Validation

- A large set of 4914 model instances was obtained by setting
  - Topology: serial, parallel
  - Service time distribution: exponential (insensitive to phase-type distributions)
  - Number of stations: $M = 3, 5, 9$ (including a delay station)
  - Number of classes: $K = 2, 4, 8$
  - Total population of jobs: $N = 128, 256, 512$ ($N_k = N/K$ without class switching)
  - Service rates: balanced, unbalanced (randomly generated with 30 seeds)
  - Service weights: balanced, unbalanced (randomly generated with 30 seeds)

Steady-State Validation

- The proposed approach was compared with an AMVA algorithm for DPS presented in *Quantitative System Performance - Computer System Analysis Using Queueing Network Models* [Lazowska et al., 1984]:

$$R_{i,k}(\boldsymbol{N}) \approx \begin{cases} D_{i,k} & \text{if } i \in \mathcal{Q} \\ D_{i,k}\left(1 + \dfrac{\sum_{l=1}^{K} w_{i,l} Q_{i,l}(\boldsymbol{N} - \boldsymbol{e}_k)}{w_{i,k}}\right) & \text{otherwise} \end{cases}$$

- The results were evaluated against discrete-event simulation using a tolerance metric found in *Linearizer - A Heuristic Algorithm for Queueing Network Models of Computing Systems* [Chandy & Neuse, 1982]:

$$\text{Tolerance} = \max_{i \in \mathbb{N}^*_{\leq M}, k \in \mathbb{N}^*_{\leq K}} \frac{\left|Q^{\mathrm{C}}_{i,k} - Q^{\mathrm{R}}_{i,k}\right|}{N_k}$$

Steady-State Validation

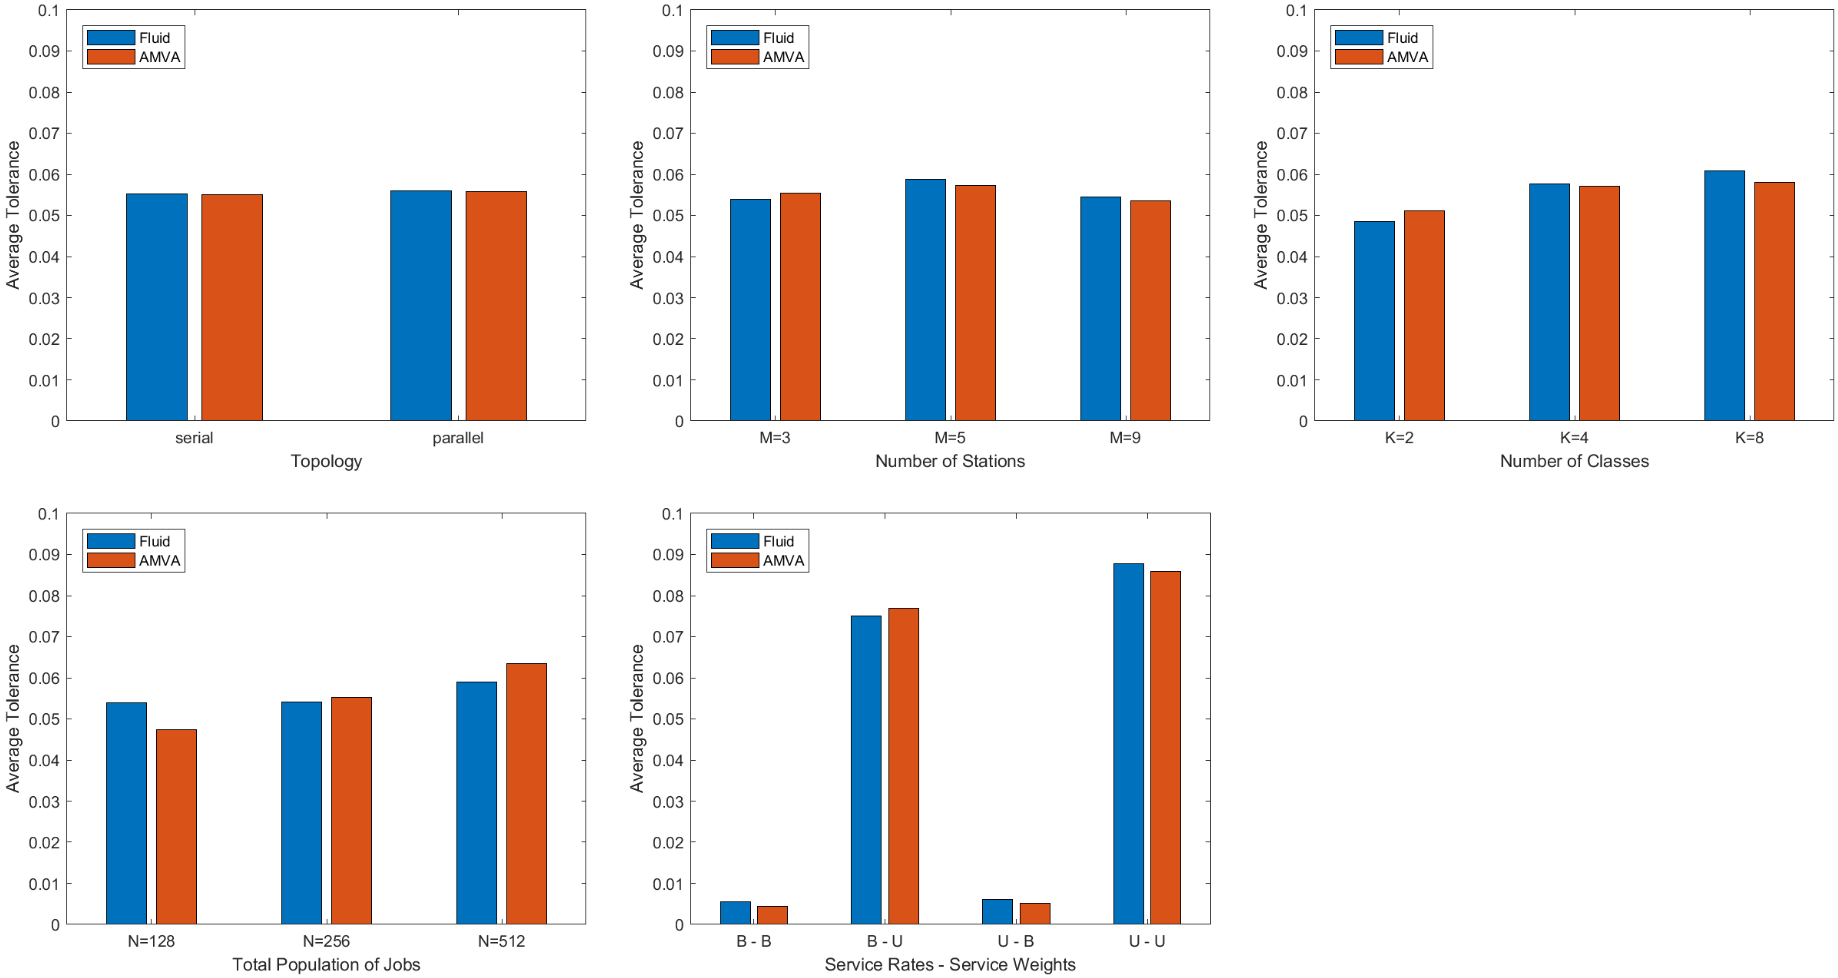

Transient Validation

- 4 typical model instances were obtained by setting
  - Topology: serial, parallel
  - Service time distribution: exponential, 2-phase coxian with a CV of 2 (sensitive to phase-type distributions)
  - Number of stations: $M = 5$ (including a delay station)
  - Number of classes: $K = 4$
  - Total population of jobs: $N = 256$ ($N_k = N/K$ without class switching)
  - Service rates: unbalanced (randomly generated with 1 seed)
  - Service weights: unbalanced (randomly generated with 1 seed)

Transient Validation

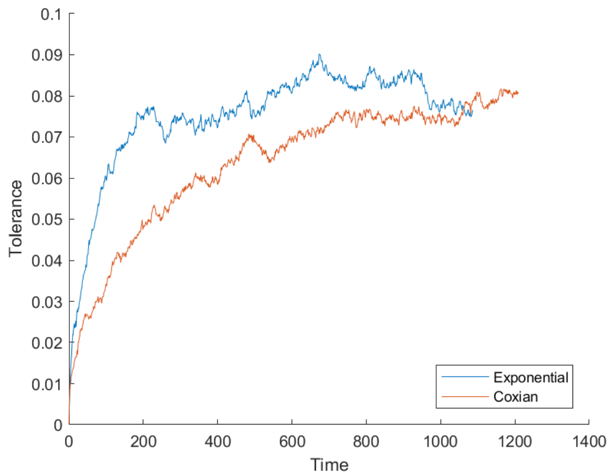

(a) Serial

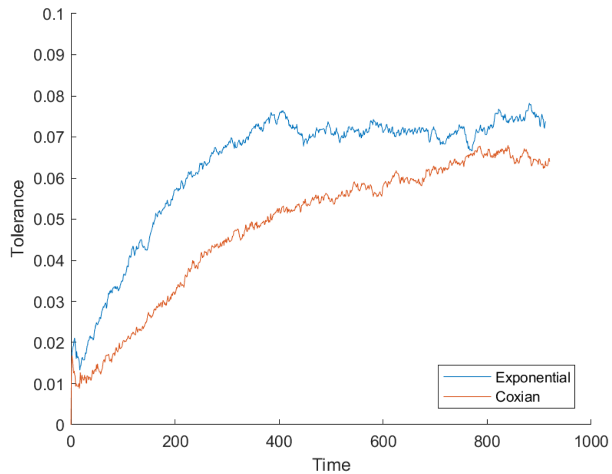

(b) Parallel

Transient Validation

(a) Serial - Exponential

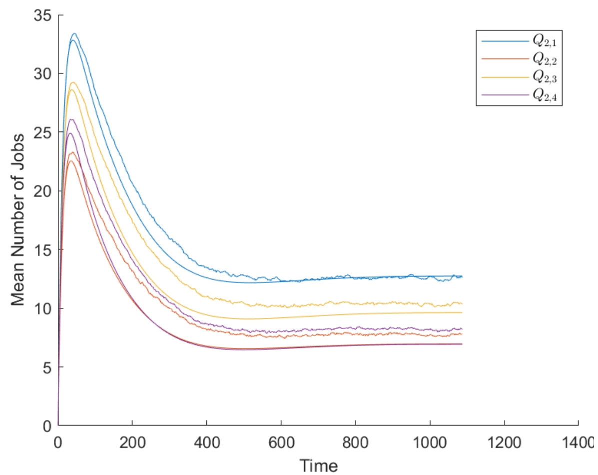

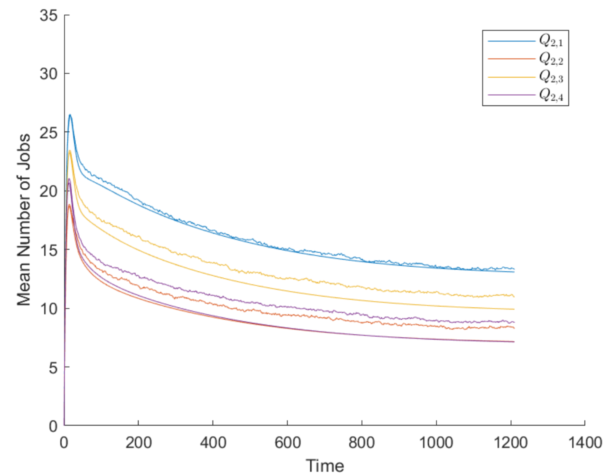

(b) Serial - Coxian

(c) Parallel - Exponential

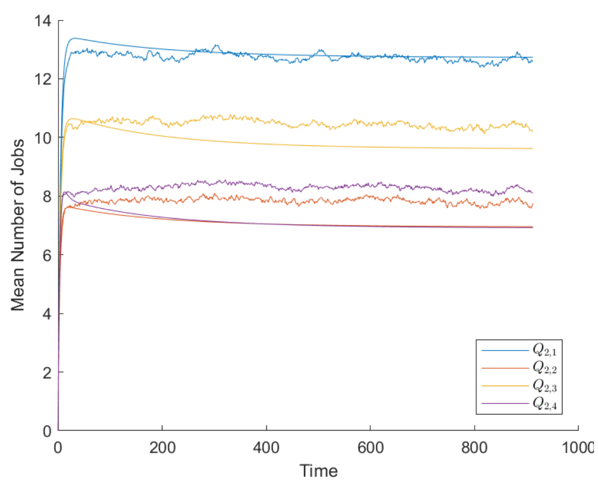

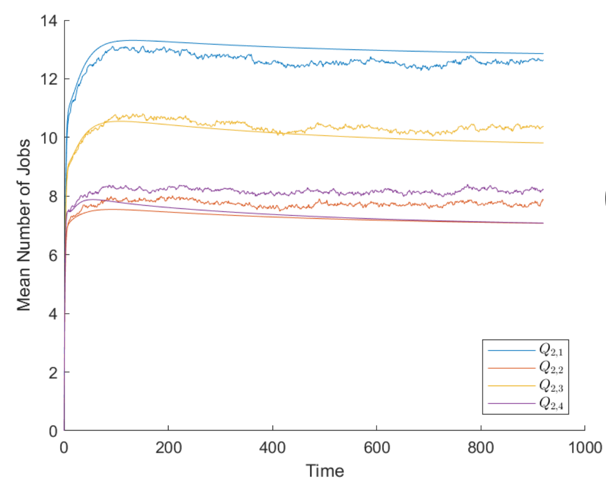

(c) Parallel - Coxian

Applications

- **Dynamic Analysis of Random Environments:**
  - *Blending randomness in closed queueing network models* [Casale, Tribastone & Harrison, 2014]
- **Approximation of Response Time Distributions:**
  - *LINE: Evaluating Software Applications in Unreliable Environments* [Casale, Tribastone & Harrison, 2014]
  - Not accurate for non-exponential service times
  - Not applicable at the system level

Imperial College London

Concept of Chains

- The routing probability matrix $\left[P_{(i,k),(j,l)}\right] \in \mathbb{R}^{MK \times MK}$ can be considered as defining a DTMC whose state is a station-class pair (i, k)
- For a multi-chain QN, the routing DTMC is assumed to be decomposable into ergodic subchains
- If we represent the routing DTMC as a weighted direct graph, then each chain manifests itself as a strongly connected subgraph
- Because all the chains are completely isolated from one another, it is impossible for a job to switch between different chains

Response Time CDF

- $\mathcal{V}_q$: Set of station-class pairs in chain q
- $\mathcal{E}_q$: Set of routes connecting station-class pairs in chain q
- Population of jobs in chain q:

$$N_q^* = \sum_{(i,k)\in \mathcal{V}_q} \sum_{r=1}^{R_{i,k}} n_{i,k,r}$$

- Set of classes that may be visited a chain-q job at station i:

$$\mathcal{C}_{i,q} = \{k : (i,k) \in \mathcal{V}_q\}$$

- Set of classes that are visited by chain-q jobs across all the stations:

$$\mathcal{C}_q = \bigcup_{i=1}^{M} \mathcal{C}_{i,q}$$

Response Time CDF

- Evaluate the mean service time $S_{i,k}$ of a class-k job at station i by applying the moment formula for phase-type distributions
- Modify the service time distribution of class k at station i so that any class-k arrival at the station starts in an additional phase 0 with probability $\alpha_{i,k,0} = 1$ and shifts from that phase to phase r with probability $P_{i,k,(0,r)} = \alpha_{i,k,r}$
  - Phase $(i, k, 0)$ therefore gathers up all class-k jobs having just arrived at station i
  - To limit the side effects of the modification, the service rate of phase $(i, k, 0)$ to $\mu_{i,k,r} = 1/\eta S_{i,k}$, where η is a small positive constant

Response Time CDF

- Build a fluid ODE system that incorporates the modified service time distribution, and compute its steady-state solution $\tilde{\boldsymbol{x}}' \in \mathbb{R}^{R'}$ under the initial condition $\boldsymbol{x}'(0) \in \mathbb{R}^{R'}$, where $R' = R + 1$
  - The initial condition $\boldsymbol{x}'(0)$ derives from the given initial condition $\boldsymbol{x}(0)$ by inserting zero entries
- Create an auxiliary class k' for tagging class-k jobs
  - Class k' has the same service time distribution and the same service weight as class k at station i
  - The probability of a class-k' job at station i being routed to station j as a class-l job is set as $P_{(i,k'),(j,l)} = P_{(i,k),(j,l)}$
  - A tagged class-k job is thereby restored immediately after departure from station i

Response Time CDF

- Build another fluid ODE system that incorporates the auxiliary class, and compute its transient solution $\boldsymbol{x}''(t) \in \mathbb{R}^{R''}$ under the initial condition $\boldsymbol{x}''(0) \in \mathbb{R}^{R''}$, where $R'' = R + R_{i,k} + 2$
  - The initial condition $\boldsymbol{x}''(0)$ arises from the resultant steady-state solution $\widetilde{\boldsymbol{x}}'$ by inserting zero entries except that $x''_{i,k,0}(0) = 0$ and $x''_{i,k',0}(0) = \tilde{x}'_{i,k,0}$
- Response time CDF of class k at station i:

$$\Phi_{i,k}(t) \approx 1 - \frac{\sum_{r=0}^{R_{i,k'}} x''_{i,k',0}(t)}{x''_{i,k',0}(0)}$$

System Response Time CDF

- Add an end phase in the service time distribution of each class k in the set $\mathcal{C}_{D,q}$ at station D to tag chain-q jobs that are about to depart from the reference station and restore them upon next arrival.
- System response time CDF of chain q with respect to station D:

$$\Phi_{i,k}(t) \approx 1 - \frac{\sum_{k\in\mathcal{C}_{D,q}} x''_{D,k',R_{D,k'}+1}(t) + \sum_{(i,k)\in \mathcal{V}_q,\, i\neq D} \sum_{r=1}^{R_{i,k'}} x''_{i,k',r}(t)}{\sum_{k\in\mathcal{C}_{D,q}} x''_{D,k',R_{D,k'}+1}(0)}$$

A Simple Example

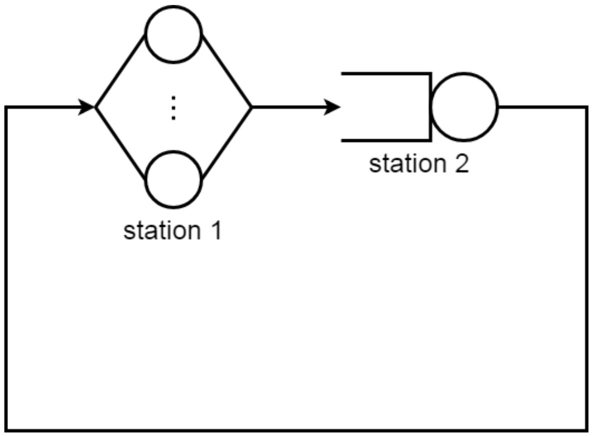

Topology

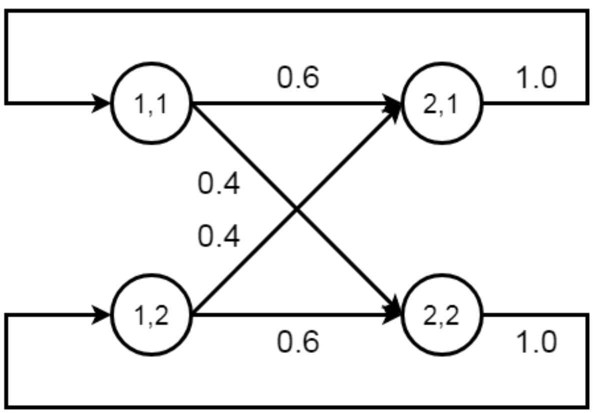

Routing CTMC

	Class 1	Class 2
Population	5	5
Think time distributions	$Exp(0.2)$	$Exp(0.5)$
Service time distributions	$HyperExp$ $([0.5, 0.5], [0.2, 1.0])$	$Coxian$ $([0.7, 1.0], [0.5, 0.1])$
Service weights	1.0	2.0
Routing probability matrix	$\begin{bmatrix} 0.0, 0.0, 0.6, 0.4 \\ 0.0, 0.0, 0.4, 0.6 \\ 1.0, 0.0, 0.0, 0.0 \\ 0.0, 1.0, 0.0, 0.0 \end{bmatrix}$	

Parameters

A Simple Example

For Response Time CDF of Class 1 at Queueing Station

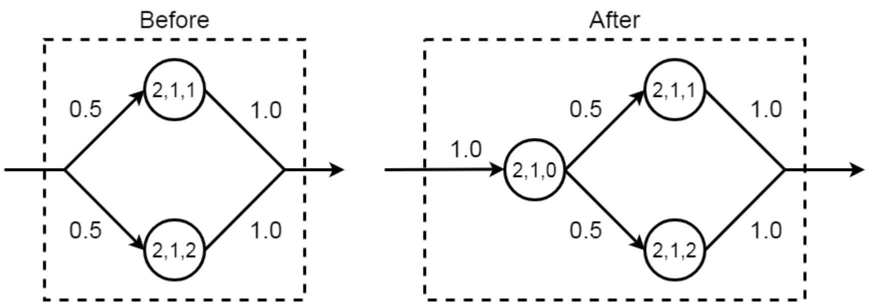

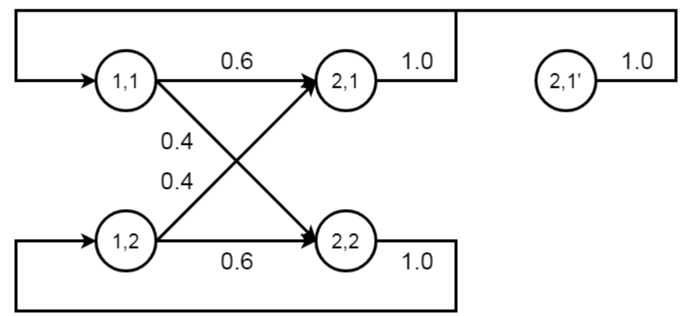

Service Time Distribution

Routing CTMC

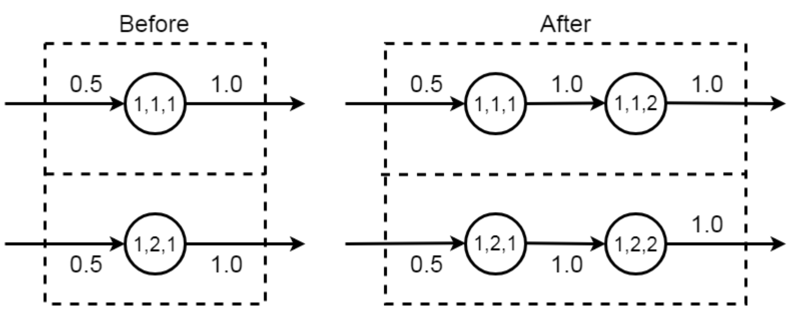

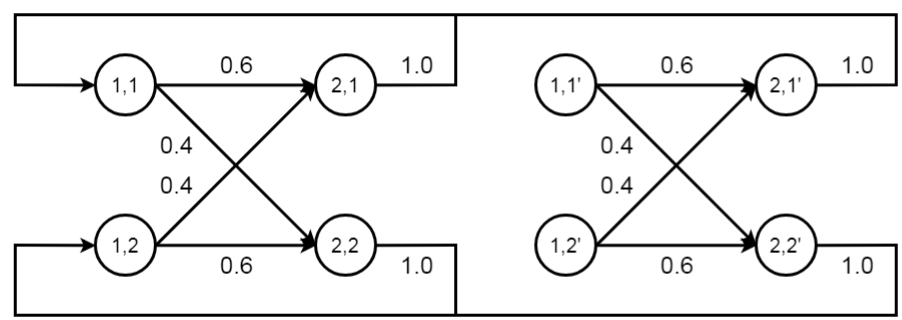

For System Response Time CDF with Respect to Delay Station

Imperial College London

A Simple Example

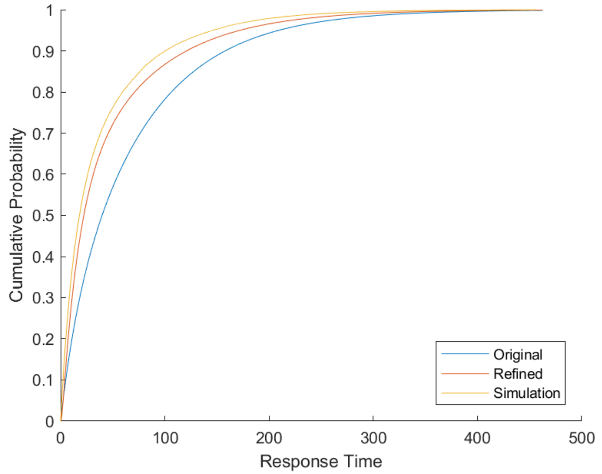

Response Time CDF of Class 1 at Queueing Station

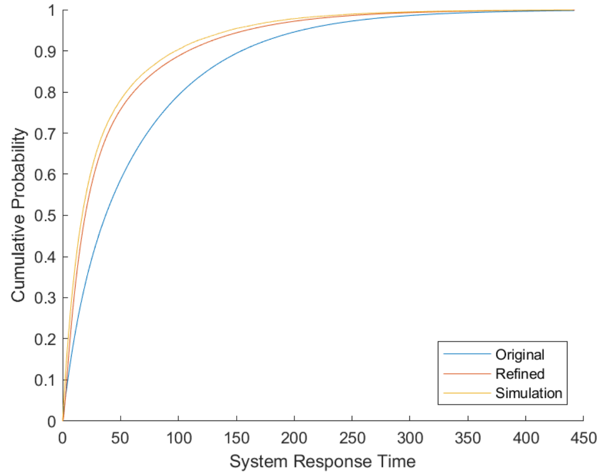

System Response Time CDF with Respect to Delay Station

Conclusion

- A fluid approach is proposed for approximating the transient and steady-state behavior of closed QNs with DPS
- Our reference model is defined generically, featuring
  - Arbitrary topology
  - Phase-type service times
  - Class switching
- The proposed approach has been verified against discrete-event simulation for both transient and steady-state analysis
- A refined method and its extension are introduced for approximating response time distributions at the station and system levels

Thank you!