Politecnico di Milano - DEI
Milan, Italy

The JMT Simulator for Performance Evaluation of Non-Product-Form Queueing Networks

Marco Bertoli, Giuliano Casale, Giuseppe Serazzi

Speaker: Giuliano Casale

March 26th, 2007
40th ANSS 2007 Symp. – Norfolk, VA, US

Outline

- Introduction
- The JMT Simulator
  - Generalities
  - Statistical Analysis of Simulation Results
  - Non-Product-Form Features
- Case Study

Capacity Planning & Performance

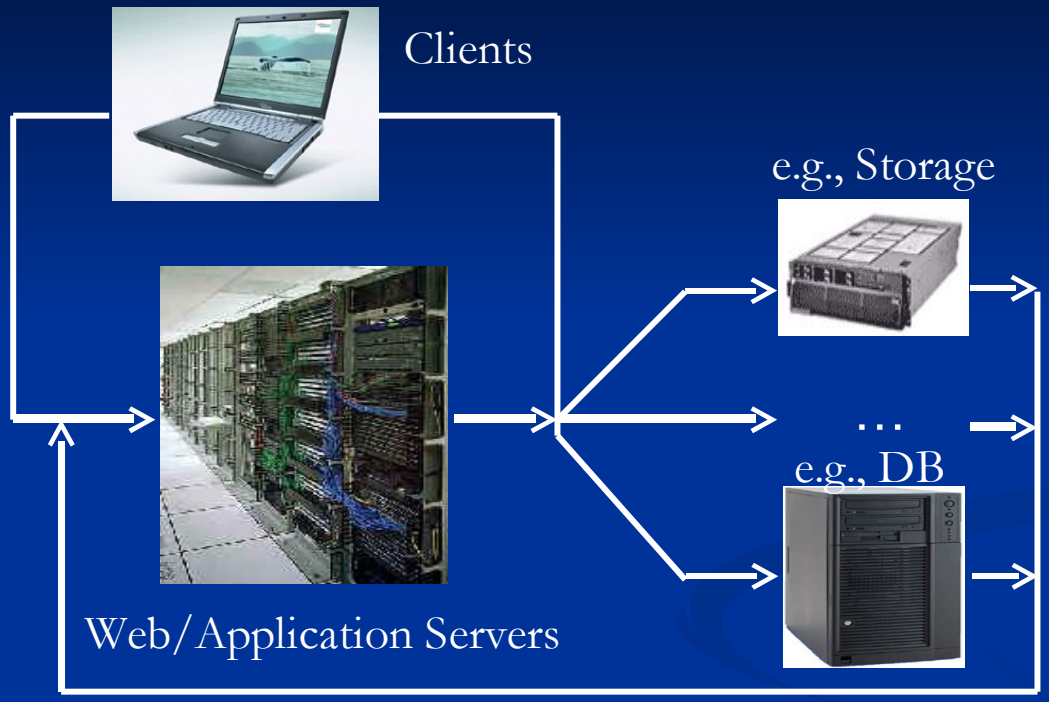

- What will be the worst-case quality of service?
- What-if I upgrade my hardware?
- What-if I consolidate servers using virtualization?

Capacity Planning Example

Select Best Design/Configuration (Routing, Num of Servers/Reqs)

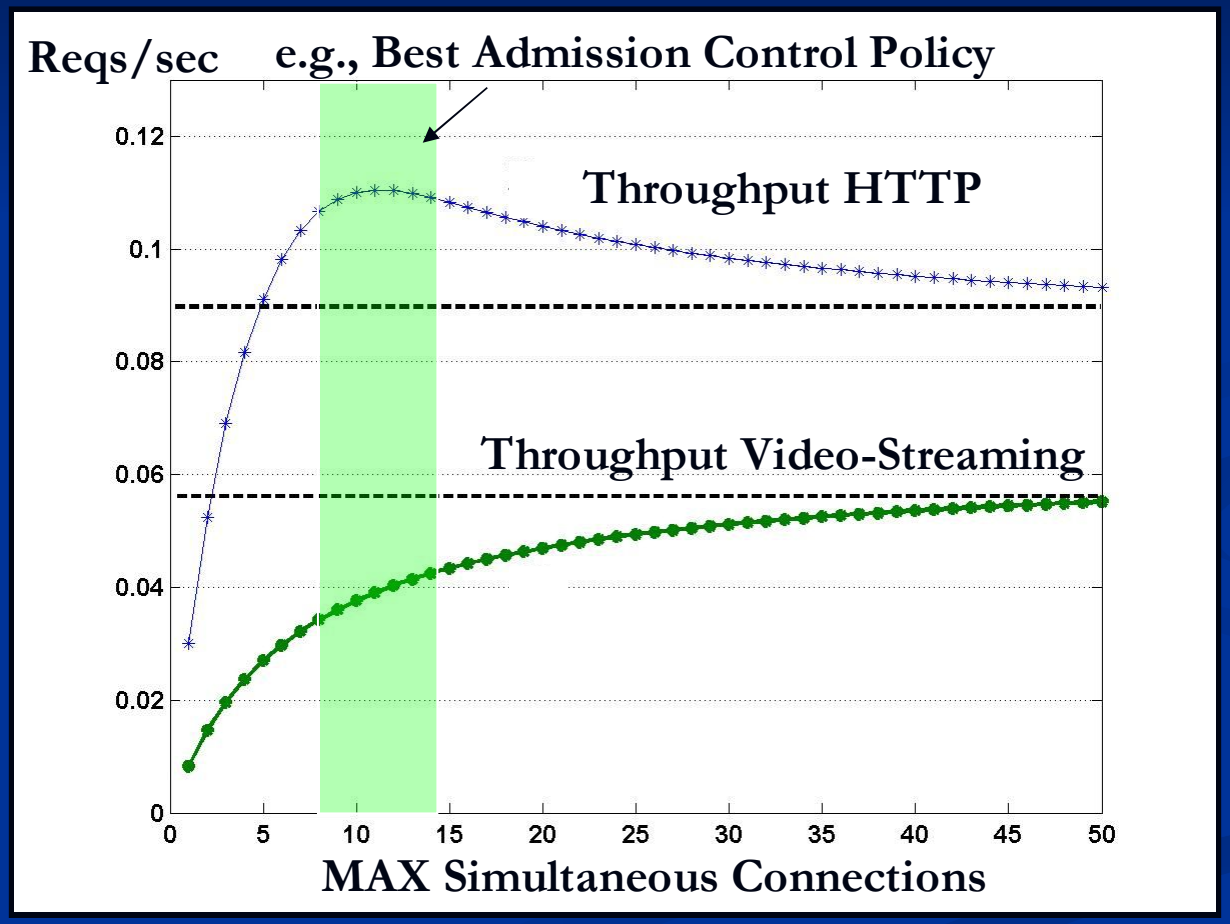

Restrictive Assumptions of PF Theory

- Models are analytically tractable if:
  - State independent service
  - Non-idling service policies
  - No Blocking
  - No Finite buffers
  - No Priorities
  - FIFO = Exponential service
  - Static Routing
  - ...
- Strong restrictions for real applications
- PF networks still valuable bounds/approximations

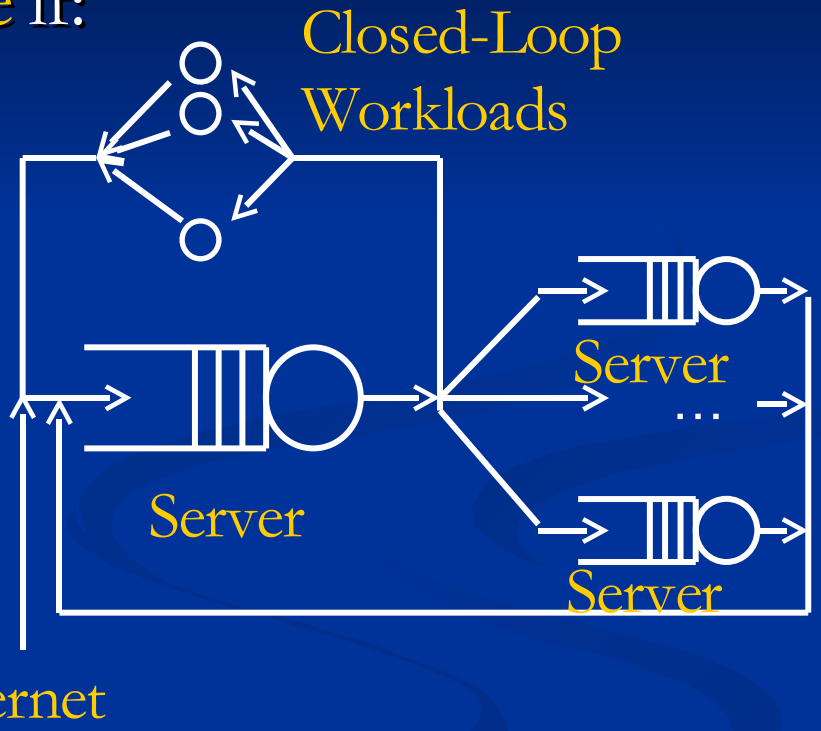

Non-Product Form (NPF) Models

- NPF models hard to approximate
  - Multiple NPF features make the model 'impossible'
  - No theory for multiclass NPF models
  - Simulation typically required/preferred (easier…)
- Java Modelling Tools project (2006)
  - Open source set of analytical and simulative tools
  - http://jmt.sourceforge.net
  - Strong diffusion for both research and teaching
  - Maximum portability (Java)

JSIM: NPF models simulator

- Core simulation module of the JMT suite
- Comes with two graphical interfaces:

JSIMgraph

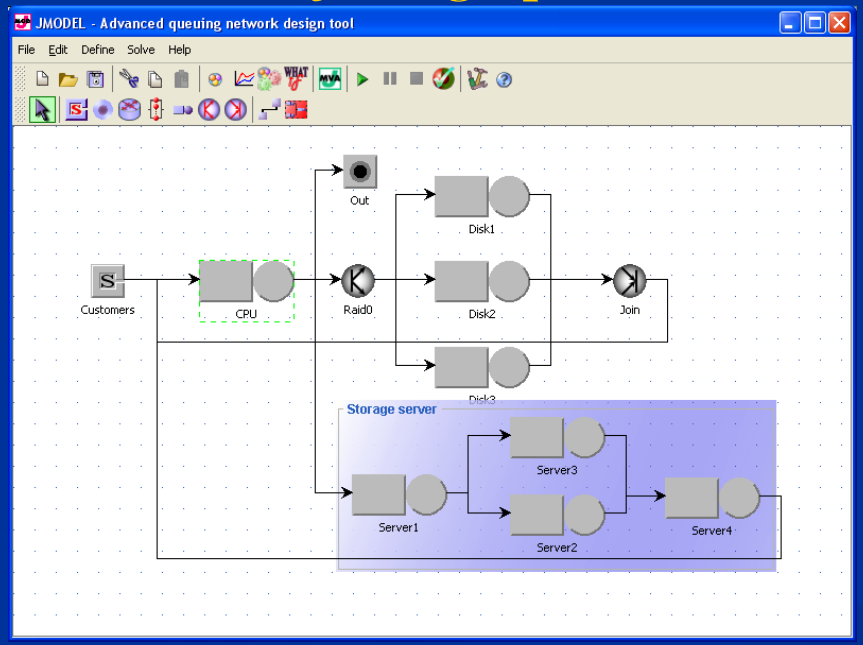

JSIMwiz

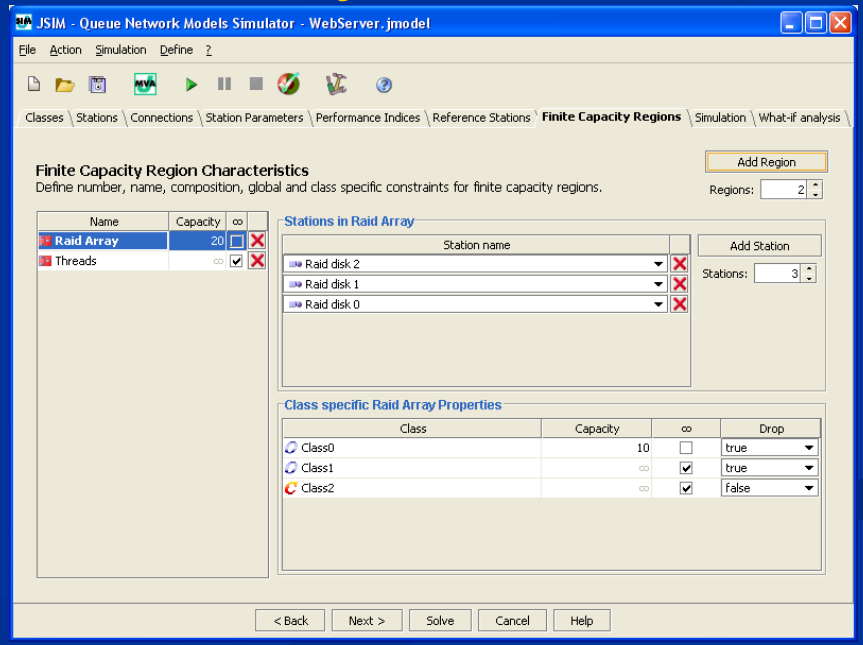

JSIM

- Architecture and Design Choices
- Statistical Analysis
- Non-Product-Form Modeling Features

JMT & JSIM: Architecture

- "Model-View-Controller"-like pattern
  - Better reuse and isolation of components

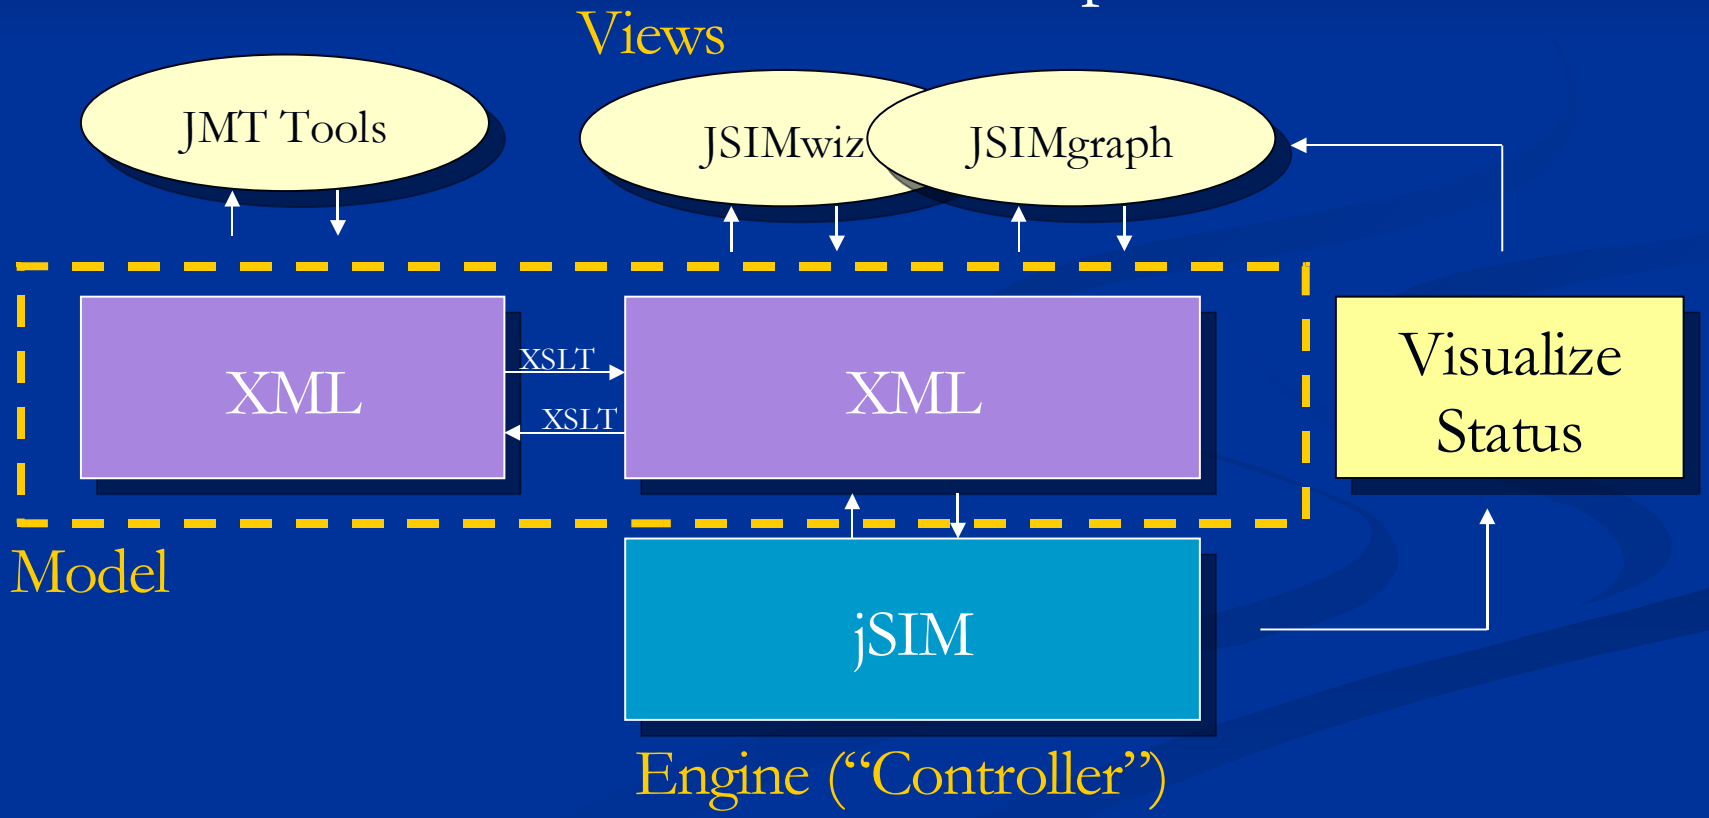

DES Engine & Simulation Entities

- Discrete Event Calendar for queue activity
- Simulation entities are compound objects
  - Input Section
  - Service Section
  - Output Section

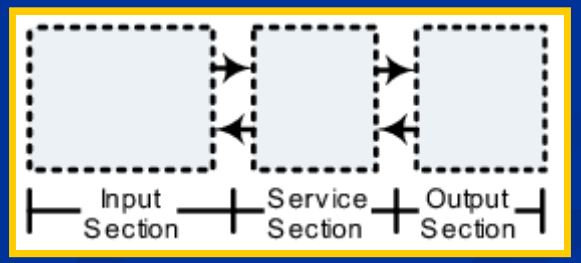

- Examples

Exogenous arrivals

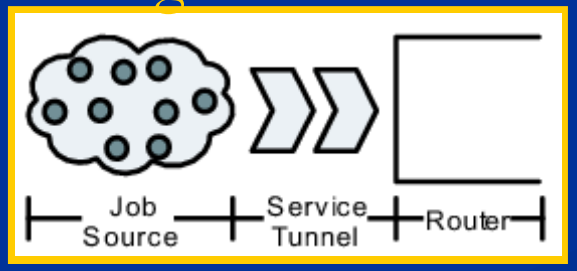

Single Server Queue

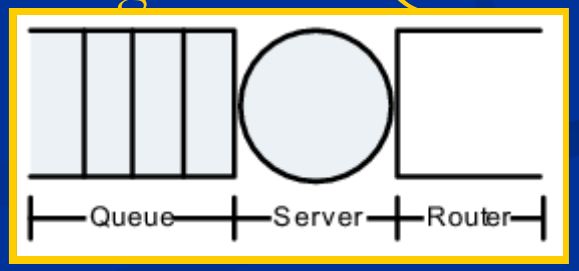

Simulation Coordination

- Strong messaging system
  - Complete separation between sections
  - External contributors need only to implement the correct messaging behavior

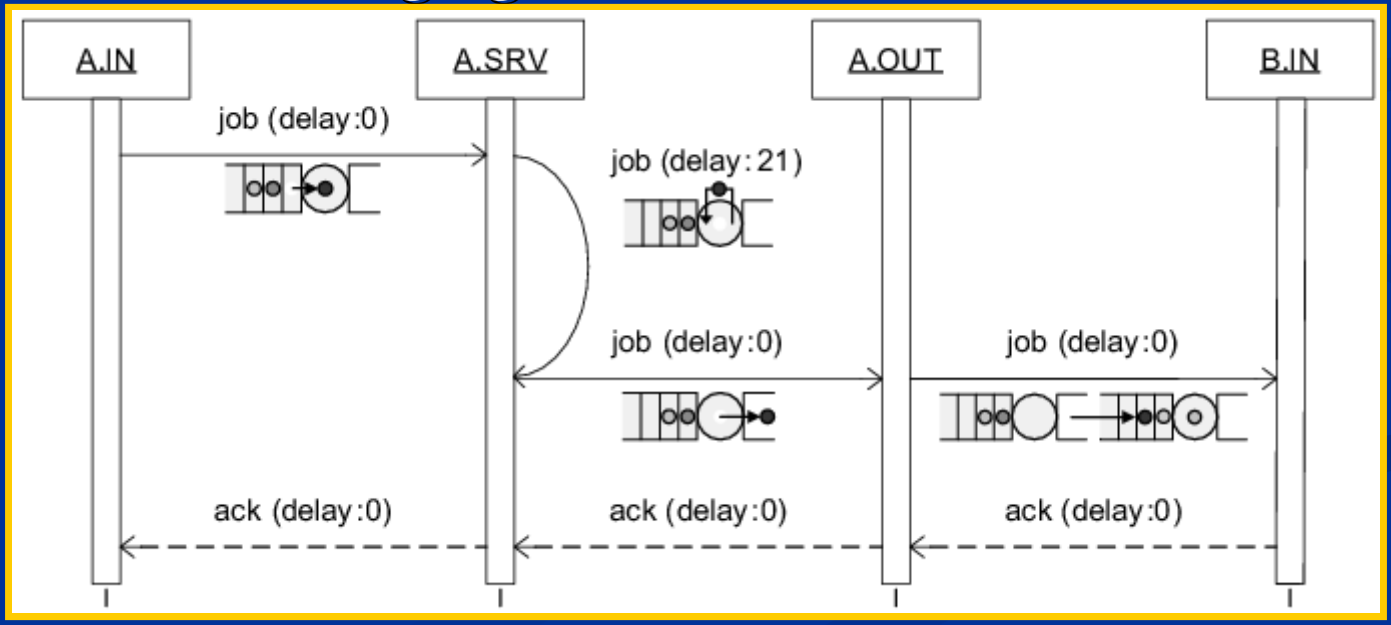

Control of Simulation Experiments

- Simplification of Experiment control
  - Maximum Relative Error [Pawlikowski , 1990]
    - Ratio Half-width marginal C.I. / Estimated Mean

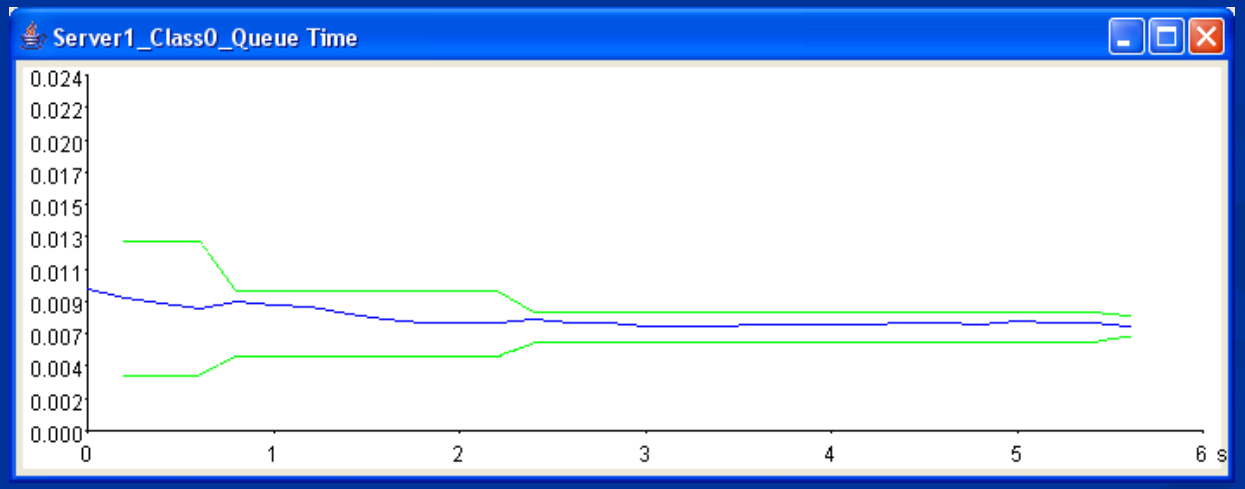

  - Maximum Number of samples (Long run analysis)
  - Maximum Simulation Time

What-if Analysis

Simulation Results...

Queue Length | Utilization

Queue Length

Average number of customers for each chosen class at each chosen station.

Description

Station Name:	CPU1
Class Name:	All classes
Conf.Int/Max Rel.Err (0-1):	0.9 / 0.1
Samples:	11
X min:	0
X max:	1
Y min:	0.05
Y max:	10.409

Plot

Show confidence interval range

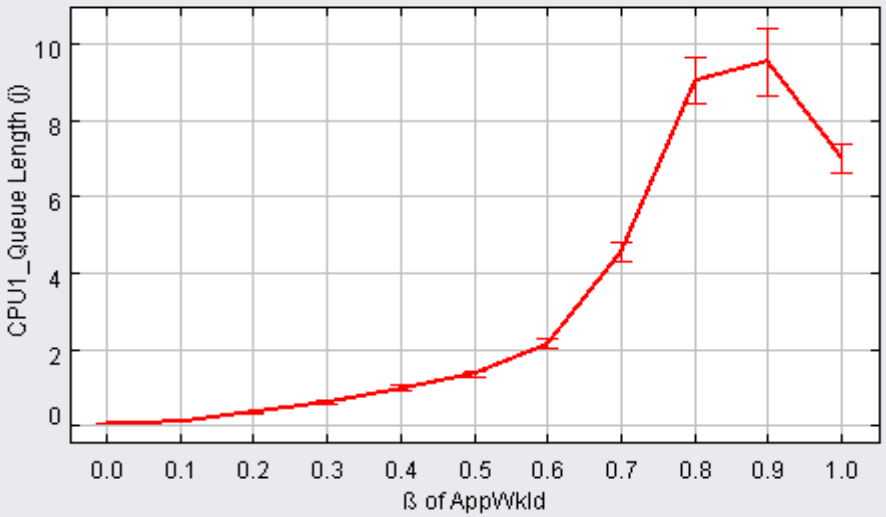

Simulation results table

*	ß = 0.000	ß = 0.100	ß = 0.200	ß = 0.300	ß = 0.400	ß = 0.500	ß = 0.600	ß = 0.700	ß = 0.80
Mean value (j)	0.050	0.135	0.341	0.615	0.985	1.355	2.142	4.549	9.034
Max (j)	0.051	0.140	0.369	0.650	1.063	1.448	2.279	4.814	9.628
Min (j)	0.050	0.131	0.313	0.580	0.906	1.263	2.006	4.284	8.439

JSIM

- Architecture and Design Choices
- Statistical Analysis of Simulation Results
- Non-Product-Form Modeling Features

Statistical Analysis

- Automatic removal of the initial bias
  - R-5 Heuristic
  - MSER-5 Rule (Marginal Standard Error Rule)
- C.I. generation using spectral methods
  - Spectral Analysis [Heidelberger & Welch, 1981]
  - Used also for run-length control

Transient Filtering

- Superposition of several rules to improve effectiveness

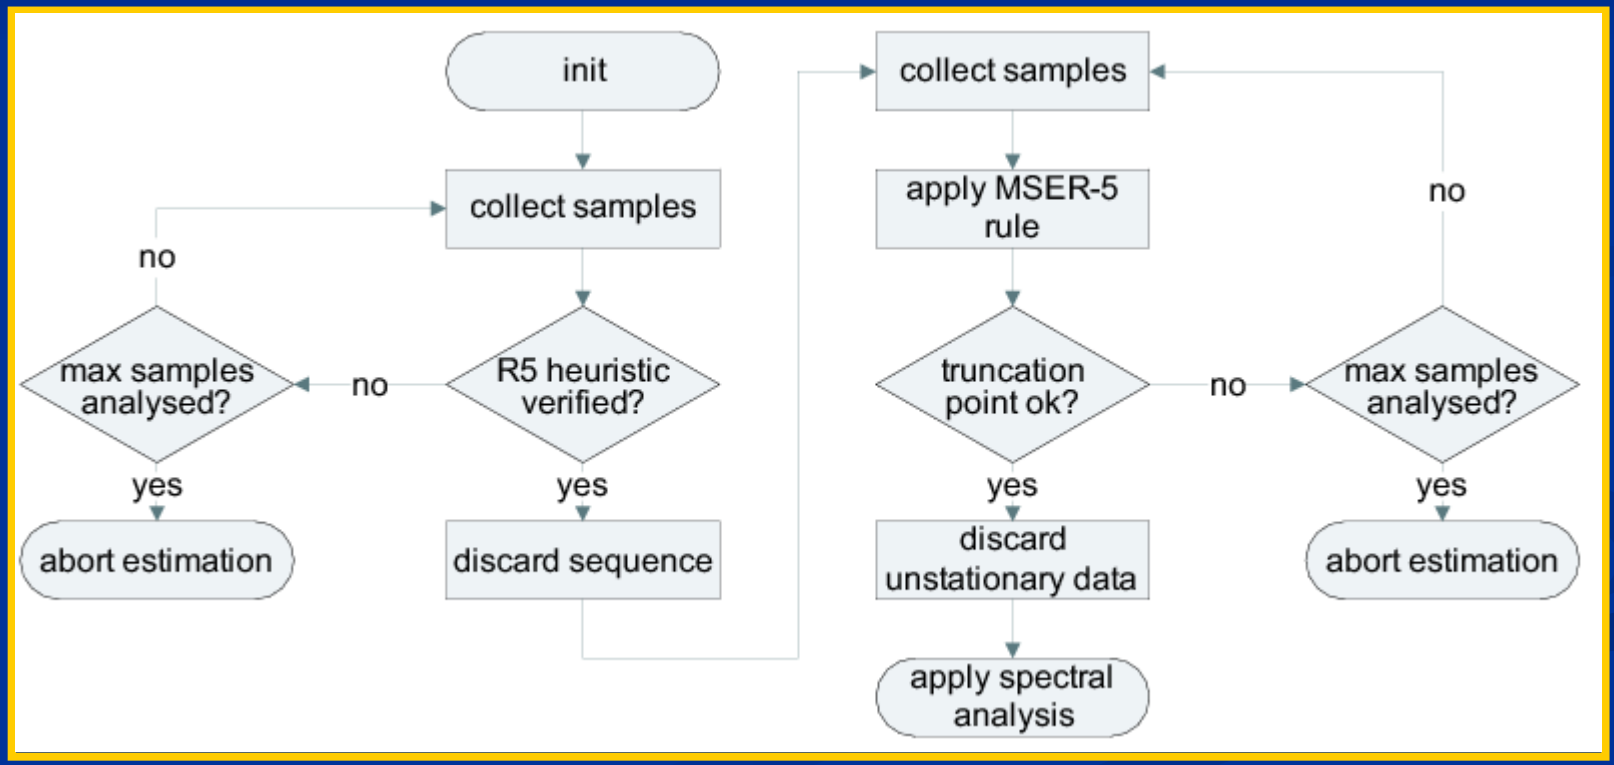

Transient Filtering (II)

- R5 Heuristic
  - Implemented according to [Pawlikowski, 1990]
  - Initial transient = samples cross mean *k* times
- *k* is a (critical) user-specified parameter
  - *k=25* for *M/M/1* (Garfarian et al., 1978)
    - Hundreds of thousands samples discarded
  - *k=7* for *M/M/1/15* (Wilson & Prisker, 1978)
    - On many networks early detection during transient ramp
  - JMT sets this parameter to a conservative *k=19*

Transient Filtering (III)

- MSER [White, 1996]
  - Best truncation point in a data sequence
  - Detects the point that minimizes the width of the marginal confidence interval about the est. mean
- MSER-5 [Spratt, 1998]
  - Batches composed by 5 samples
  - We implemented an online version of the algorithm
  - Cyclically run on 5000 batches until detection
    - Increasing the number of batches has limited effect

JSIM

- Architecture and Design Choices
- Statistical Analysis of Simulation Results
- Non-Product-Form Modeling Features

JSIM NPF Models

- Main NPF modeling features
  - General Arrival and Service Processes
  - Fork-Join Centers
  - Finite Capacity Models
  - Priority Classes
  - Advanced state-dependent routing, e.g.:
    - Route to least utilized queue
    - Route to shortest queue

Arrival and Service Process

- Exponential insufficient for many models
  - Pareto, Hyperexponential, Erlang, Gamma, …
  - Custom distribution (external text file, future JWAT)
- Random number generation
  - Mersenne Twister
- Load-dependent service process
  - Server speed variable with the current queue-length
  - Building block for Hierarchical Modeling

Hierarchical Modeling

- Compact representation of large models

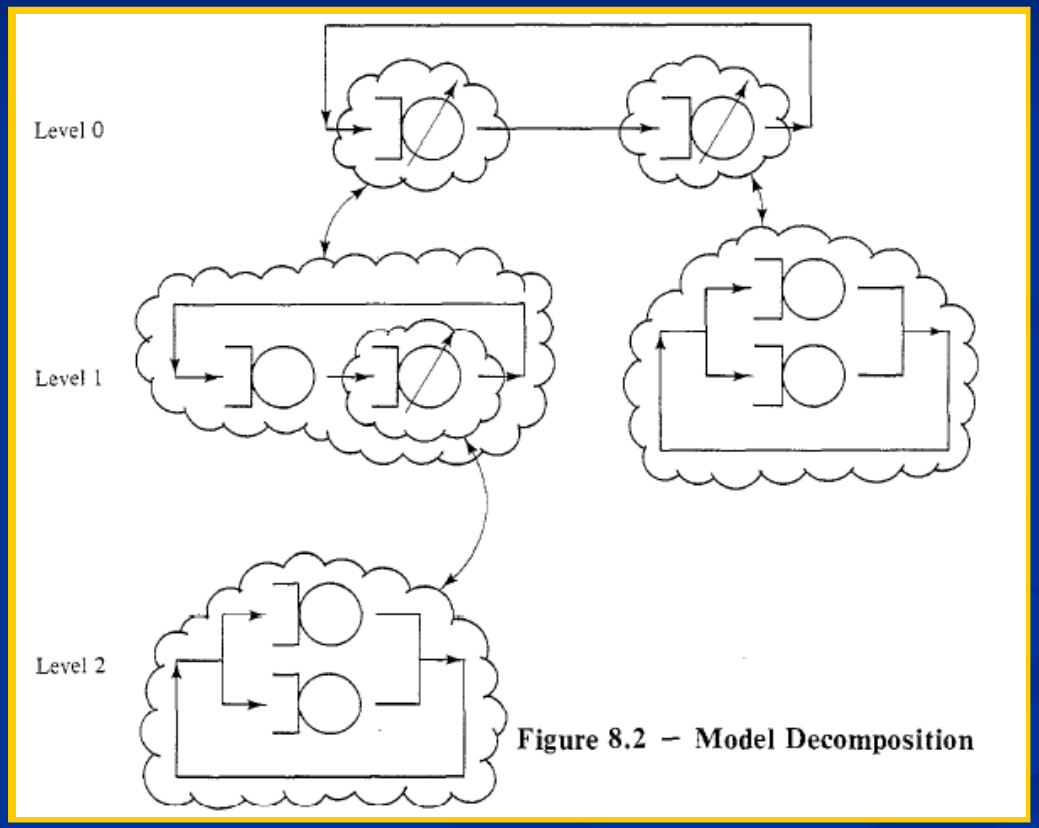

Image taken from: Lazowska et al. Quantitative System Performance. Prentice-Hall, 1984.

Fork-Join Systems

- Popular in storage and multiprocessor models
  - Jobs are forked at fork node into P tasks
  - Synchronization at the join node before leaving
- JSIM: special ad-hoc Fork and Join components

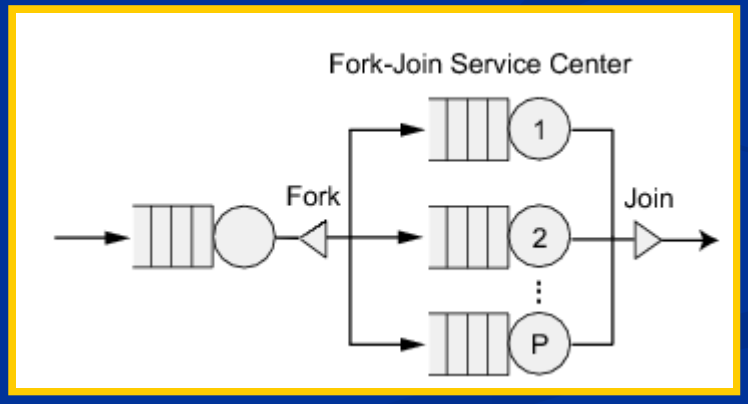

Finite Capacity Regions

- Models of admission control in networks
  - Describe well application and memory constraints
- JSIM allows to tag a group of queues as a region
  - Non-admitted jobs can be either put in a FCFS waiting buffer or dropped

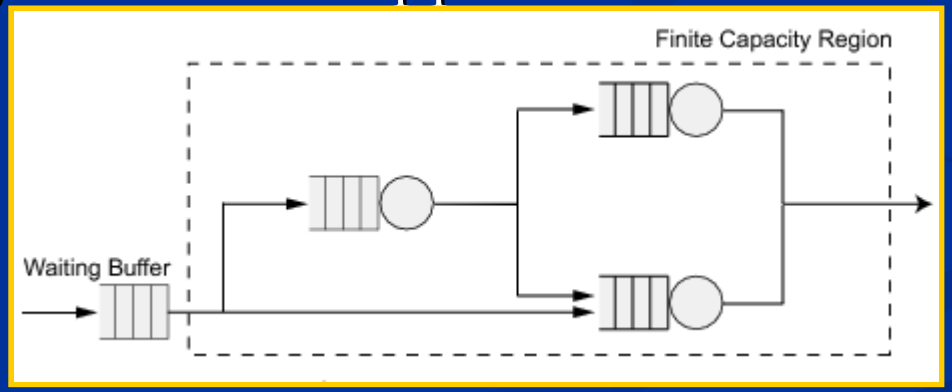

Case Study

- Multiprocessor Web server
  - Workloads: orders (class 1), backend service (class 2)
  - Constant population of requests (N_1,N_2), $N_1+N_2=N$

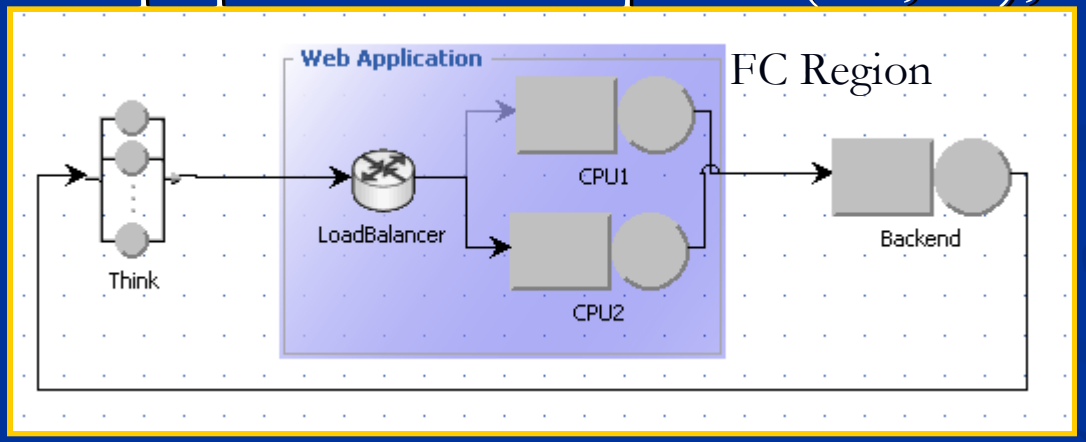

  - Finite capacity constraints
    - Limits shared by all classes, or class-dedicated

Asymptotic Analisys in PF Networks

- How does the system evolves with the mix?
  - PF=Multiple bottlenecks, still unobserved in NPF

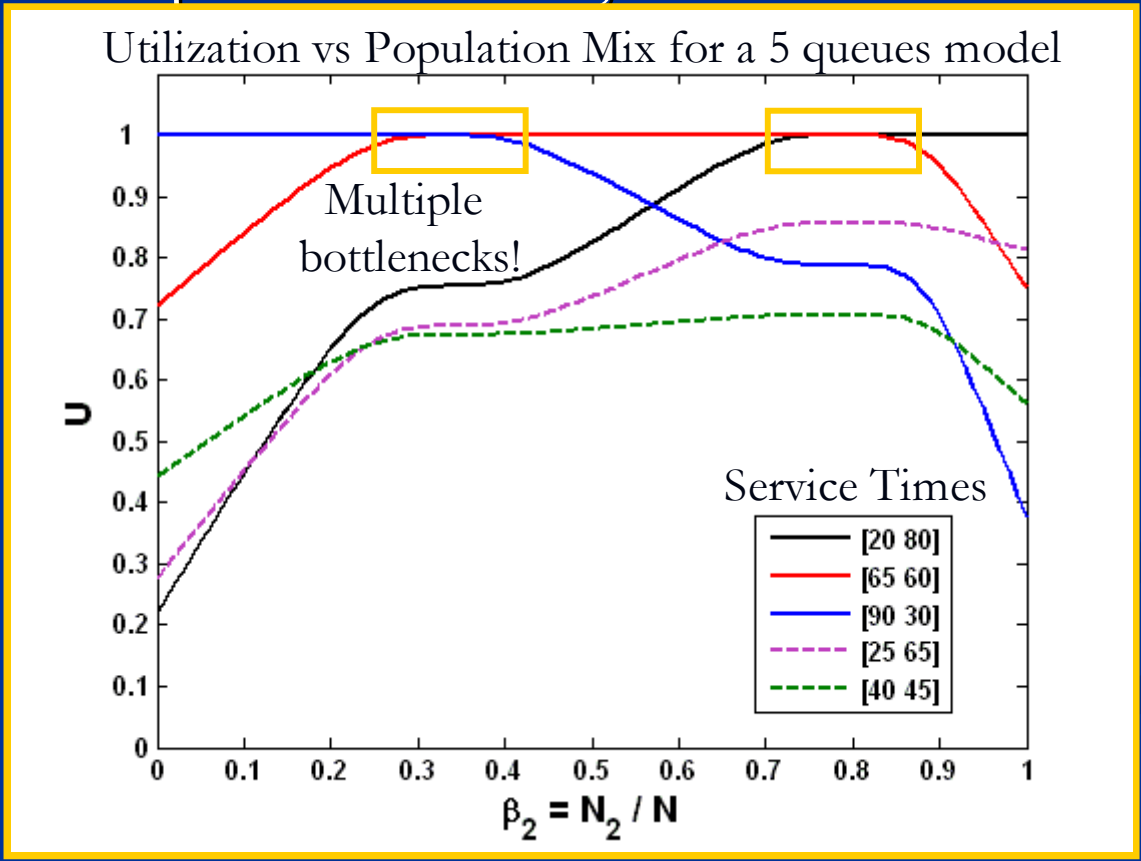

Shared constraints
(no multiple bottlenecks)

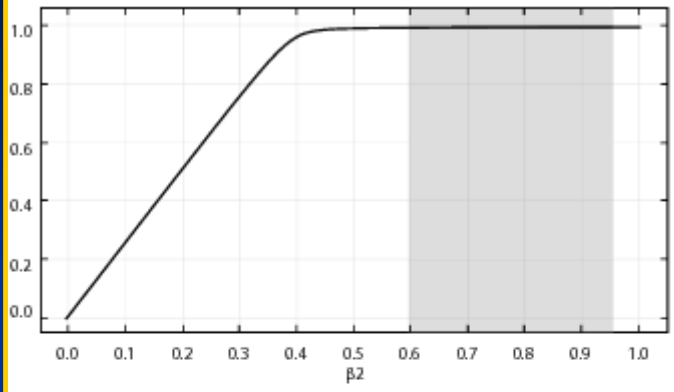

(a) Product-Form: U_{CPU}.

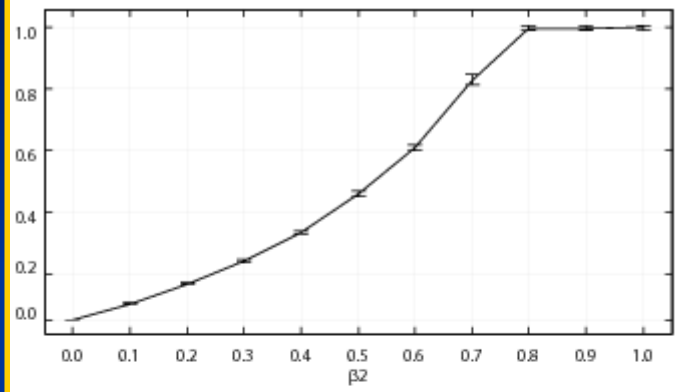

(b) Shared: U_{CPU}.

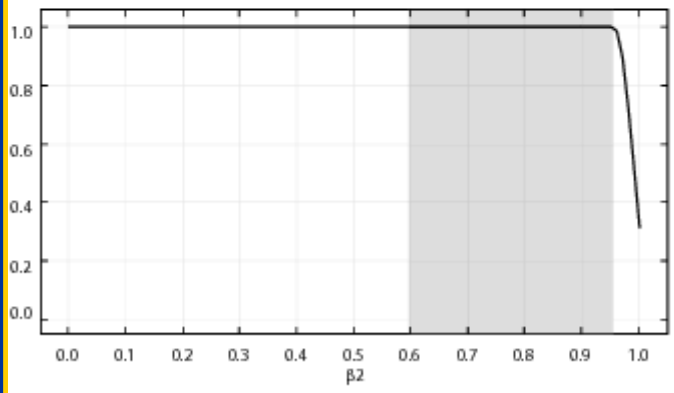

(e) Product-Form: $U_{Backend}$.

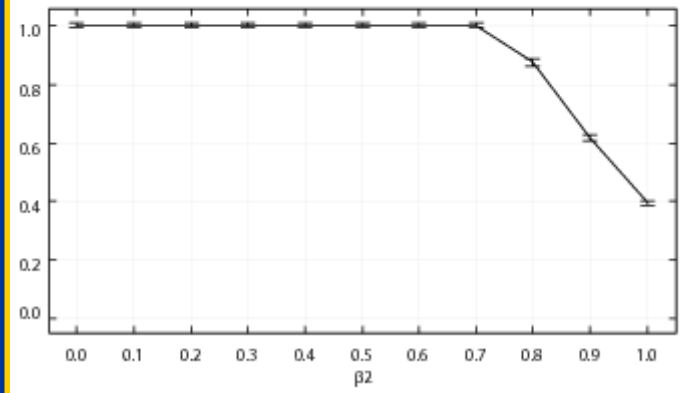

(f) Shared: $U_{Backend}$.

Shared constraints
(no multiple bottlenecks)

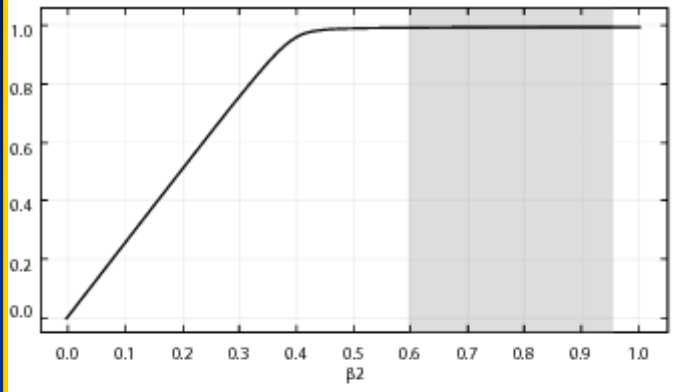

(a) Product-Form: U_{CPU}.

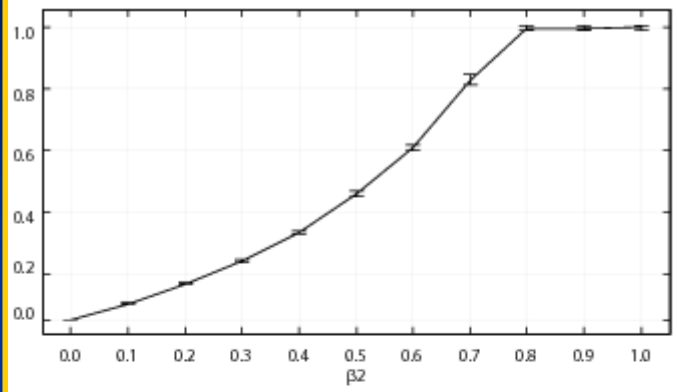

(b) Shared: U_{CPU}.

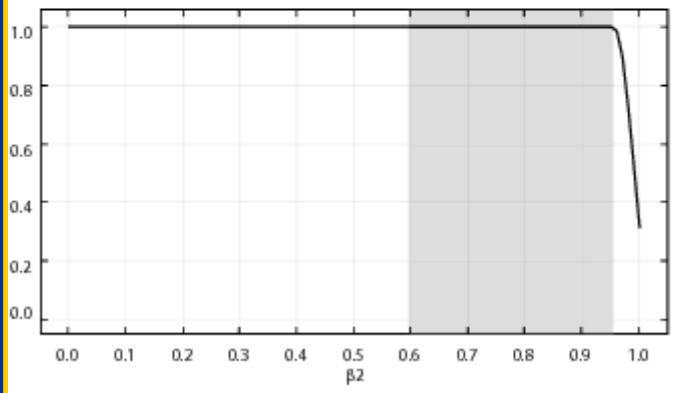

(e) Product-Form: $U_{Backend}$.

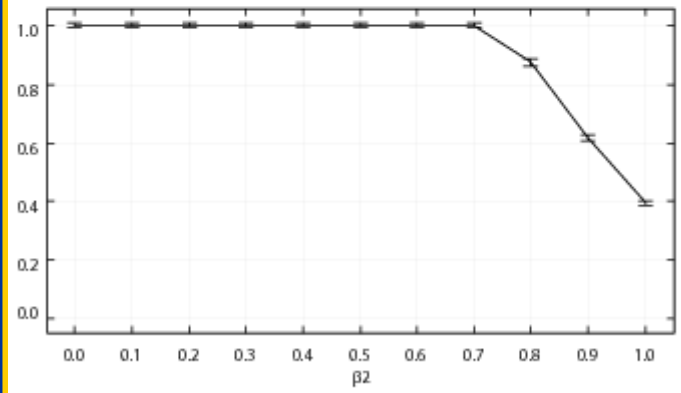

(f) Shared: $U_{Backend}$.

Dedicated – Class 1
(no multiple bottlenecks)

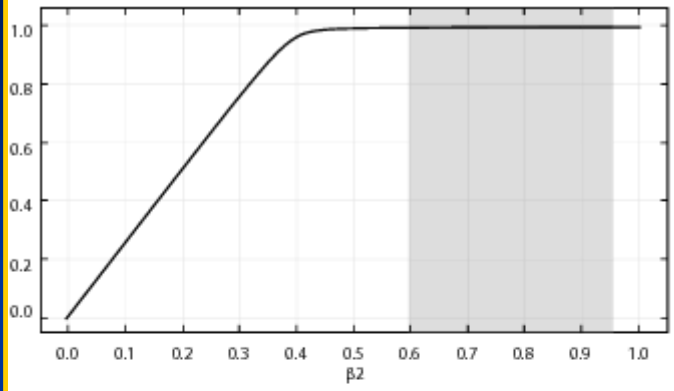

(a) Product-Form: U_{CPU}.

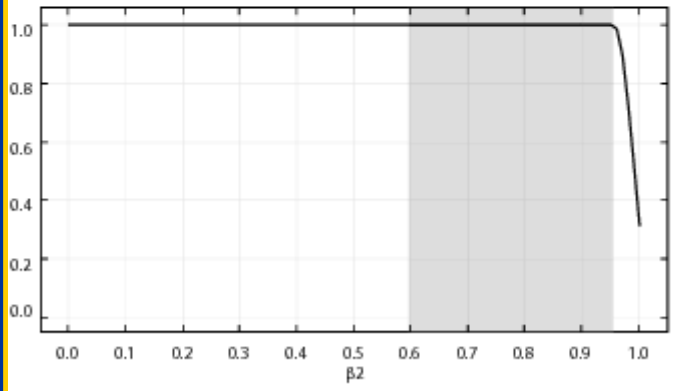

(e) Product-Form: $U_{Backend}$.

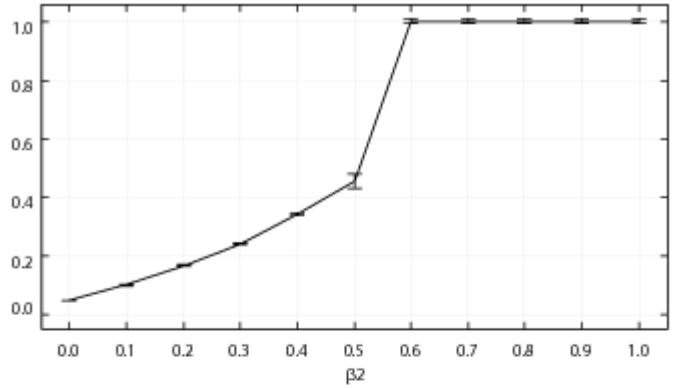

(c) Class-1 Dedicated: U_{CPU}.

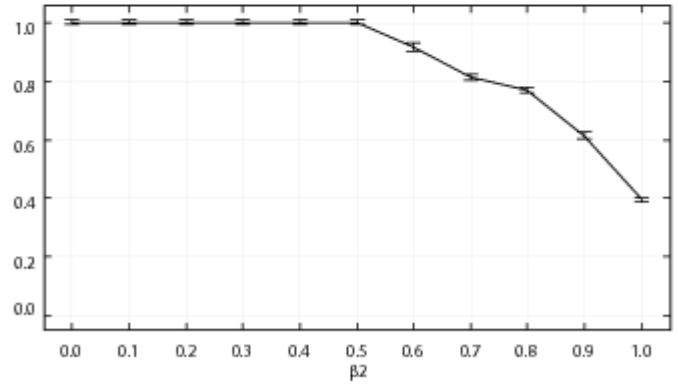

(g) Class-1 Dedicated: $U_{Backend}$.

Dedicated – Class 2
(multiple bottlenecks!)

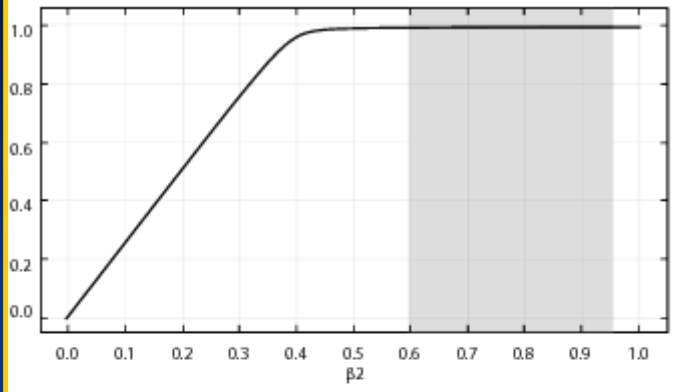

(a) Product-Form: U_{CPU}.

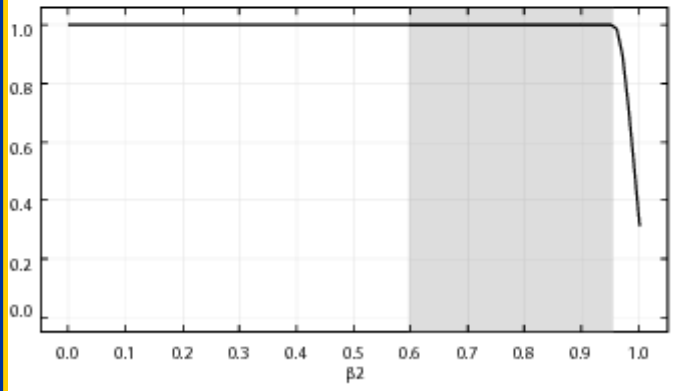

(e) Product-Form: $U_{Backend}$.

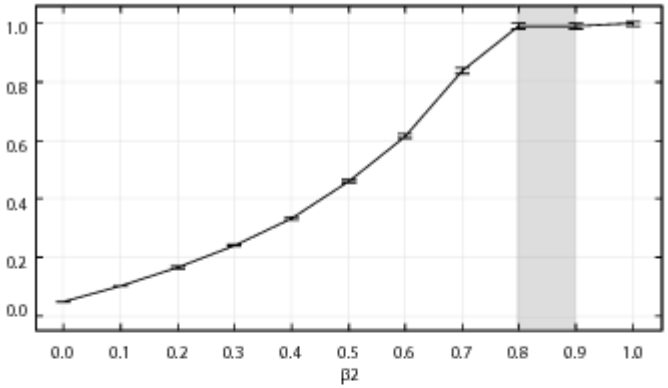

(d) Class-2 Dedicated: U_{CPU}.

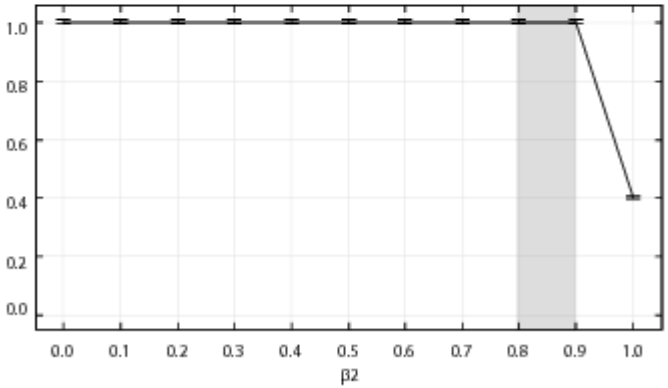

(h) Class-2 Dedicated: $U_{Backend}$.

Observations

- Theoretical interest to better understand multiclass models
- Class 1 has bottleneck outside the FC region
- Class 2 has bottleneck inside the FC region
- FC region creates a closed PF sub-model
  - Behavior of PF models may apply inside the region
  - We expect to observe the effect also in real systems

Conclusion

- JSIM: advanced queueing network simulation
- Free, open source, GNU GPL project

http://jmt.sourceforge.net

- External contributors are welcome
- Current version 0.7, new releases to come…