A Matrix-Analytic Approximation for Closed Queueing Networks with General FCFS Nodes

Giuliano Casale[a], Gábor Horváth[b], Juan F. Pérez[c]

a: Imperial College London, Department of Computing, UK
b: Budapest University of Technology and Economics, Hungary
c: University of Melbourne, ACEMS, Australia

MAM9, Budapest, Hungary
June 29, 2016

Context and Problem

- Multiclass closed queueing networks commonly used in capacity planning and software performance engineering
- Product-form solution given by the BCMP theorem (1975)
  - Identical per-class service times at FCFS nodes
  - Exponentially-distributed service times at FCFS nodes
- Existing approximations for general FCFS nodes are brittle.
- Contribution: an accurate matrix-analytic approximation for multiclass networks with general FCFS nodes

Model

- Closed network of M FCFS queues with PH service
- Cyclic topology (method applies also to arbitrary topology)
- R job classes, each with K_r jobs
- s_{ir} : mean service time of class-r jobs at queue i
- v_{ir} : mean visits of class-r jobs at queue i
- $\theta_{ir} = v_{ir}s_{ir}$: mean demand of class-r jobs at queue i

Solving product-form models (AMVA)

- Product-form case: $s_{ir} = s_{is}$, $r \neq s$, exponential service.
- Mean value analysis (MVA) solves model using Little's law and the arrival theorem:

$$W_{ir} = \theta_{ir} + \theta_{ir} \sum_{s=1}^{R} A_{is}^{(r)}$$

  - W_{ir} : mean response time at queue i for class-r jobs
  - $A_{is}^{(r)}$: mean class-s queue seen at i by arriving class-r job.
- Approximate MVA (AMVA):
  - $A_{is}^{(r)}$ approximated with inexpensive fixed-point iteration.

Including general FCFS nodes (AMVA-FCFS)

- We now consider general FCFS nodes (arbitrary s_{ir}).
- AMVA-FCFS corrects the arrival theorem at FCFS nodes as

$$W_{ir} \approx \theta_{ir} + \sum_{s=1}^{R} \theta_{is} A_{is}^{(r)}$$

- Each queue-length is weighted differently according to class.
- Pretty good approximation, but two key problems:
  - Brittle: models with $> 15\% - 20\%$ error are not uncommon.
  - Insensitive: no moments other than means, no residual time.

Sensitivity to higher-order moments

- Decomposition method for single-class FCFS queues with PH service (Casale & Harrison, ICPE 2012).
- FCFS queues considered in isolation, fed by throughput X same as in the closed network.
  - Isolated queues treated as $MAP/PH/1$ queues
  - Approximation! True arrival rate is state-dependent.
  - Start with a guess for X, then iteratively update the guess.

Sensitivity to higher-order moments

- We approximate QBD solution with a scalar expression:

$$\tilde{p}_i(n) = \begin{cases} (1-\rho_i) & n = 0 \\ \rho_i(1-\eta_i)\eta_i^n & n > 0 \end{cases}$$

  - Similar to a diffusion approximation for closed networks.
  - ρ_i: utilization
  - η_i: tail decay rate (caudal characteristic)

Sensitivity to higher-order moments

- X iteratively updated based on the above approximation.
  - FCFS queue replaced by a load-dependent exponential node that contributes $\tilde{p}_i(n_i)$ to the product-form expression, i.e.,

$$p(n_1, \ldots, n_M) \approx \frac{\tilde{p}_1(n_1)\tilde{p}_2(n_2)\cdots\tilde{p}_M(n_M)}{G}$$

    where G is a normalising constant.
  - New value for X efficiently obtained from this approximation.

Problems towards a multiclass extension

- What is the caudal characteristic η for a multiclass queue?
  - What kind of population growth should we consider?
- Multiclass load-dependent nodes are difficult to handle, how do we iteratively update guesses on $\mathbf{X}$?

Multiclass Extension

- Network decomposed into MMAP[R]/PH[R]/1 queues.
- Arrival rates given by throughputs: $\mathbf{X} = (X_1, \ldots, X_R)$
- Arrival process obtained by scaling of input PHs and superpositions.
  - Traffic flow superpositions and aggregations needed for arbitrary topologies.

MMAP[R]/PH[R]/1 queues

- For a state $\mathbf{n} = (n_1, \ldots, n_R)$ the queue has

$$p_i(\mathbf{n}) = \begin{cases} (1 - \rho_i) & |\mathbf{n}| = 0 \\ \pi(0) \sum_{r=1, n_r>0}^{R} \mathbf{L}(\mathbf{n} - \mathbf{e}_r)\mathbf{h}_r & |\mathbf{n}| > 0 \end{cases}$$

- $\mathbf{L}(\mathbf{n})$: recursively calculated by Sylvester matrix equations.
- $\pi(0)$: initial vector of the age process.
- $\mathbf{h}_r$: 1 for states where class-r job is in service, 0 elsewhere.

MMAP[R]/PH[R]/1: Multinomial Approximation

- For tractability, we apply a multinomial approximation

$$p_i(\mathbf{n}_i) \approx \Big(1 - \sum_{s=1}^{R} \eta_{is}\Big) \frac{(n_{i1} + \cdots + n_{iR})!}{n_{i1}! \cdots n_{iR}!} \eta_{i1}^{n_{i1}} \cdots \eta_{iR}^{n_{iR}}.$$

- η_{ir}: tail decay of class r at queue i
- $p_i(\mathbf{n}_i)$ leads to approximate product-form equilibrium

$$p(\mathbf{n}_1, \ldots, \mathbf{n}_M) \approx \frac{\tilde{p}_1(\mathbf{n}_1)\tilde{p}_2(\mathbf{n}_2) \cdots \tilde{p}_M(\mathbf{n}_M)}{G}$$

solvable by standard algorithms, such as AMVA.

MMAP[R]/PH[R]/1: Decay rate

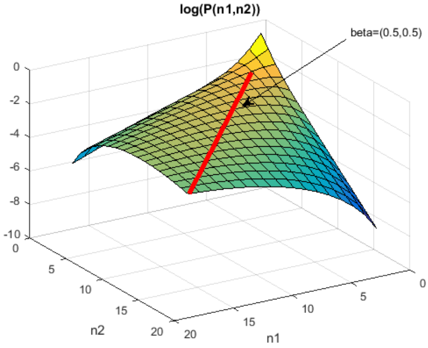

- Population growth under fixed class mix $\beta = \mathbf{n}_i/|\mathbf{n}_i|$
- η_{ir} obtain by choosing $\beta_i = E[\mathbf{n}_i]/|E[\mathbf{n}_i]|$
- Population growth limited to reachable vectors $\mathbf{n}_i : n_{ir} \leq K_r$.

MMAP[R]/PH[R]/1: Multinomial Approximation

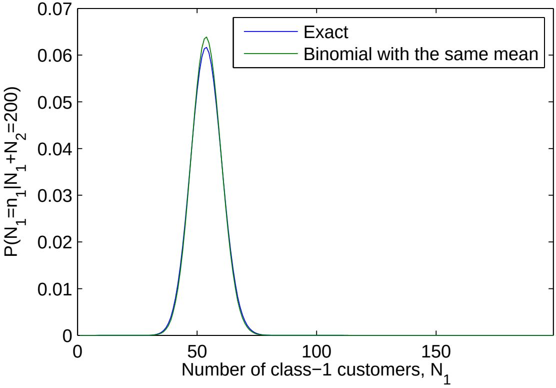

Figure: Distribution of the number of class-2 jobs given that the total number of jobs is 200

Decay Rate Approximation (DRA)

- Initial guess of throughputs $\mathbf{X}$ by AMVA-FCFS.
- Decompose into MMAP$[R]$/PH$[R]$/1 queues and find $E[\mathbf{n}_i]$
- Obtain decay rates η_{ir} under average mix
- Parameterize a product-form model with queues having demands η_{ir}/X_r
- Solve product-form model to obtain new throughputs $\mathbf{X}'$
- Repeat until minimizing weighted distance between $\mathbf{X}$ and $\mathbf{X}'$.
  - We use a non-linear constrained optimization solver.

Validation: Methodology

- 80 networks solved by simulation, AMVA, AMVA-FCFS, DRA.
- $R = 2$ job classes
- $M \in \{2, 3, 4, 8\}$ queues
- $K \in \{15, 30, 45, 60\}$ jobs, $K_2 = 2K_1$.
- One queue is hyper-exponential: $c^2 \in \{1, 2, 5\}$.
- We study errors on mean queue-lengths:

$$\text{error} = \frac{1}{2K} \sum_{i=1}^{M} \sum_{r=1}^{R} |E[n_{ir}] - E[n_{ir}^{sim}]|,$$

Validation: Results

Table: Distribution of errors for different methods across all test cases

	Method		
Error (%)	DRA	AMVA-FCFS	AMVA
0 - 5	42.5%	33.75%	20%
5 - 10	45%	30%	38.75%
10 - 15	12.5%	27.5%	26.25%
15 - 20	-	7.5%	11.25%
20 - 25	-	1.25%	3.75%

Ongoing and Future work

Ongoing work:

- Arbitrary topologies
- Reduction of computational cost of update
- Inclusion of other node types (PS, delay servers)
- Random validation on networks with arbitrary topologies

Future work:

- Generalization to priorities and fork-join queueing systems.