IMPERIAL

A Family of Multiclass LCFS Networks with a Novel Product-Form Solution

Giuliano Casale

Imperial College London

12-May-2026 @ ECQT 2026, Reykjavik

Problem

- Product-form $=$ exact solution as a tractable product, e.g. for BCMP

$$\pi(\text{job distribution}) \propto \prod_i f(\text{state of queue } i)$$

- BCMP/Kelly (1975): multiclass PS, FCFS, IS, or LCFS-PR (preemptive-resume)
- Multiclass LCFS-NP (non-preemptive) lies outside the BCMP product-form

This talk:

- Two-station LCFS networks (NP+PR) admit a non-BCMP product-form solution
- Paper in QUESTA 2026 special issue on Product-Form Probability Distributions

Outline

Outline

Model

- Closed cyclic two-station model
- Single job per class
- Exponential class-specific means (results also hold if PR is Coxian)
- mean NP service time $\alpha_r = 1/\mu_r$
- mean PR service time $\beta_s = 1/\sigma_s$

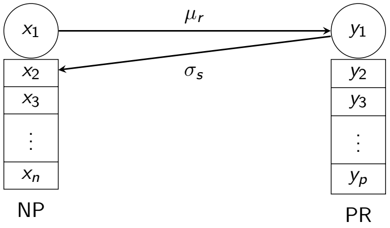

$$\boldsymbol{x} = (x_1, \ldots, x_n), \quad \boldsymbol{y} = (y_1, \ldots, y_p)$$

$$x_1 : \text{ running at NP}, \quad y_1 : \text{ running at PR}$$

$$n + p = N$$

CTMC $N = 3$

- Irreducibile CTMC
- Three partitions emerge
- Partitions fix the last job at the PR station
- Striped states bridge partitions (PR empty)

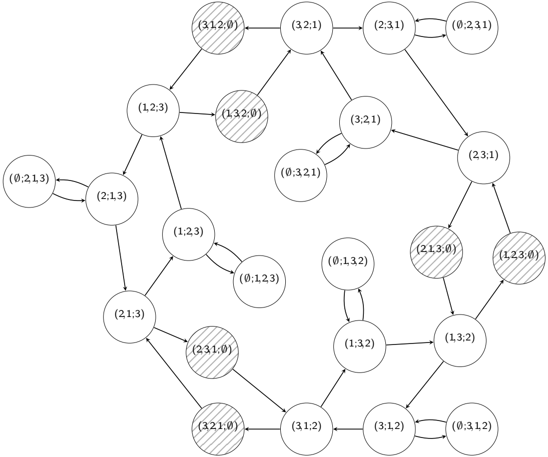

CTMC $N = 3$

- Irreducibile CTMC
- Three partitions emerge
- Partitions fix the last job at the PR station
- Striped states bridge partitions (PR empty)
- **Orange partition = class 3 last at PR**

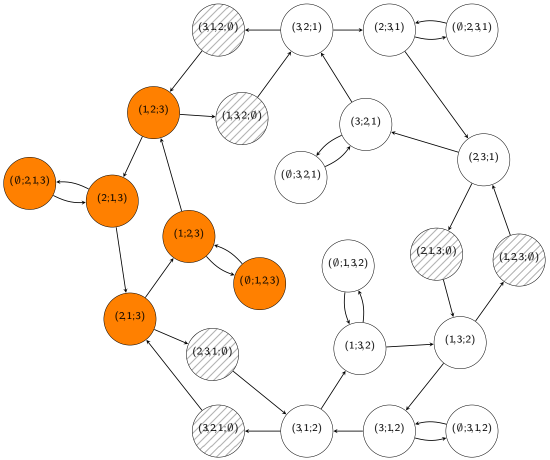

Partition Detail

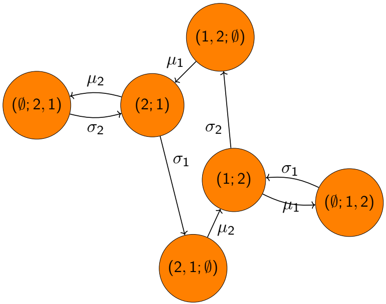

CTMC $N = 2$

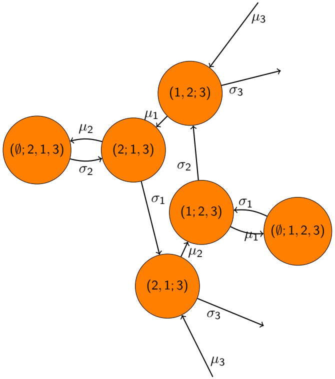

CTMC $N = 3$

Product Form

- Main result (from global balance equations)

$$\pi(\boldsymbol{x};\boldsymbol{y}) = \frac{1}{S}\left(\prod_{i=1}^{n} \alpha_{x_i}^{n-i+1}\right)\left(\prod_{j=1}^{p} \alpha_{y_j}^{N-p+j-1}\beta_{y_j}\right)$$

  where S is a normalizing constant
- Not BCMP:
  - Sensitive to the job positions within the buffers (i and j)
  - PR term also includes NP means
- As expected, identical service time means recover BCMP product-form (PR=NP)
- S is challenging to compute, state space size is $(N+1)!$

Outline

The Permanent

- Problem: how to accumulate in S the permutations of buffer positions?
- For a square matrix $\boldsymbol{A} = [a_{i,j}]$ of order n

$$\text{perm}\, \boldsymbol{A} = \sum_{\sigma \in \Sigma_n} \prod_{i=1}^{n} a_{i,\sigma(i)}$$

- Like the determinant, but all signs taken positive
- Laplace expansion by minors: $\text{perm}\, \boldsymbol{A} = \sum_{i=1}^{n} a_{i,j}\, \text{perm}\, \boldsymbol{A}(i \mid j)$

Computing the normalizing constant S

- Fix marginal population n at NP station, associated product-form factors described by

$$\boldsymbol{A}_n = \begin{bmatrix} \alpha_1 & \alpha_1^2 & \alpha_1^3 & \cdots & \alpha_1^{n-1} & \alpha_1^n & \alpha_1^n\beta_1 & \alpha_1^{n+1}\beta_1 & \cdots & \alpha_1^{N-1}\beta_1 \\ \alpha_2 & \alpha_2^2 & \alpha_2^3 & \cdots & \alpha_2^{n-1} & \alpha_2^n & \alpha_2^n\beta_2 & \alpha_2^{n+1}\beta_2 & \cdots & \alpha_2^{N-1}\beta_2 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ \alpha_{N-1} & \alpha_{N-1}^2 & \alpha_{N-1}^3 & \cdots & \alpha_{N-1}^{n-1} & \alpha_{N-1}^n & \alpha_{N-1}^n\beta_{N-1} & \alpha_{N-1}^{n+1}\beta_{N-1} & \cdots & \alpha_{N-1}^{N-1}\beta_{N-1} \\ \alpha_N & \alpha_N^2 & \alpha_N^3 & \cdots & \alpha_N^{n-1} & \alpha_N^n & \alpha_N^n\beta_N & \alpha_N^{n+1}\beta_N & \cdots & \alpha_N^{N-1}\beta_N \end{bmatrix}$$

Rows = classes, columns = buffer positions.

Laplace expansion of $\boldsymbol{A}_n$

- Fix marginal population n at NP station, associated product-form factors described by

$$\boldsymbol{A}_n = \begin{bmatrix} \alpha_1 & & & & & & & & & \\ & \alpha_2^2 & \alpha_2^3 & \cdots & \alpha_2^{n-1} & \alpha_2^n & \alpha_2^n\beta_2 & \alpha_2^{n+1}\beta_2 & \cdots & \alpha_2^{N-1}\beta_2 \\ & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ & \alpha_{N-1}^2 & \alpha_{N-1}^3 & \cdots & \alpha_{N-1}^{n-1} & \alpha_{N-1}^n & \alpha_{N-1}^n\beta_{N-1} & \alpha_{N-1}^{n+1}\beta_{N-1} & \cdots & \alpha_{N-1}^{N-1}\beta_{N-1} \\ & \alpha_N^2 & \alpha_N^3 & \cdots & \alpha_N^{n-1} & \alpha_N^n & \alpha_N^n\beta_N & \alpha_N^{n+1}\beta_N & \cdots & \alpha_N^{N-1}\beta_N \end{bmatrix}$$

Assign a job of class 1 to position 1 in NP station, then expand on the remaining submatrix.

Laplace expansion of $\boldsymbol{A}_n$

- Fix marginal population n at NP station, associated product-form factors described by

$$\boldsymbol{A}_n = \begin{bmatrix} \alpha_1 & & & & & & & & \\ & & \alpha_2^3 & \cdots & \alpha_2^{n-1} & \alpha_2^n & \alpha_2^n\beta_2 & \alpha_2^{n+1}\beta_2 & \cdots & \alpha_2^{N-1}\beta_2 \\ & & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ & & \alpha_{N-1}^3 & \cdots & \alpha_{N-1}^{n-1} & \alpha_{N-1}^n & \alpha_{N-1}^n\beta_{N-1} & \alpha_{N-1}^{n+1}\beta_{N-1} & \cdots & \alpha_{N-1}^{N-1}\beta_{N-1} \\ & \alpha_N^2 & & & & & & & & \end{bmatrix}$$

Assign a job of class N to position 2 in NP station, then expand on the remaining submatrix.

$\longrightarrow \cdots \longrightarrow$ product-form term for state $(\boldsymbol{x}; \boldsymbol{y})$

Computing the normalizing constant S (cont'd)

Since the permanents sums all the sequences, we have

$$\pi_{\mathrm{NP}}(n \mid \boldsymbol{N}) = \frac{\operatorname{perm} \boldsymbol{A}_n}{S} = \pi_{\mathrm{PR}}(N - n \mid \boldsymbol{N})$$

and requiring that marginal probabilities sum to 1, we get

$$S = \sum_{n=0}^{N} \operatorname{perm} \boldsymbol{A}_n,$$

Similar formulas for throughput $T_r(\boldsymbol{N})$ and mean queue-length $Q_{\mathrm{NP},r}(\boldsymbol{N})$ follow

Recurrence relation for S (Convolution algorithm)

- Expand the permanent on its last column to remove one class recursively

$$S = V + \sum_{r=1}^{N} \alpha_r^{N-1} \beta_r S_r$$

where S_r refers to a model without class-r and with

$$V = \left(\prod_{j=1}^{N} \alpha_j \right) \sum_{r=1}^{N} V_r,$$

- With multiple jobs per class, complexity becomes $O(N^R)$

Outline

Arrival Theorems

- MVA recursion directly on mean metrics,
- Follows by taking derivatives of the convolution algorithm expression
- BCMP [ReiL80]: class-k arrival sees model with one fewer class-k job

$$Q_{i,k}(\boldsymbol{N}) = U_{i,k}(\boldsymbol{N}) + U_{i,k}(\boldsymbol{N})\, Q_i(\boldsymbol{N} - 1_k)$$

- LCFS: condition on the job at the back of the PR buffer

$$Q_{\mathrm{PR},k}(\boldsymbol{N}) = B_{\mathrm{PR},k}(\boldsymbol{N}) + \sum_{r \neq k} B_{\mathrm{PR},r}(\boldsymbol{N})\, Q_{\mathrm{PR},k}(\boldsymbol{N} - 1_r)$$

- Derivatives of S do not yield a direct formula for NP metrics!

Reciprocal Model

- Replace mean service times by reciprocals: $\bar{\alpha}_r = \alpha_r^{-1}$, $\bar{\beta}_r = \beta_r^{-1}$
- Same stationary distribution, with NP and PR swapped:

$$\pi(\boldsymbol{x}\,;\,\boldsymbol{y}) = \bar{\pi}(\boldsymbol{y}\,;\,\boldsymbol{x})$$

- Example: $N = 2$, $\alpha = (2, 3)$, $\beta = (1, 4)$, $S = 80 \Rightarrow \bar{\alpha} = \left(\frac{1}{2}, \frac{1}{3}\right)$, $\bar{\beta} = \left(1, \frac{1}{4}\right)$, $\bar{S} = \frac{5}{9}$:

$$\pi(1; 2) = \frac{\alpha_1\,\alpha_2\,\beta_2}{S} = \frac{3}{10} = \frac{\bar{\alpha}_2\,\bar{\alpha}_1\,\bar{\beta}_1}{\bar{S}} = \bar{\pi}(2; 1)$$

- Holds in two-station BCMP models too

Reciprocal Model and MVA

- In our LCFS model: PR in reciprocal = NP in original $\Rightarrow$ get NP metrics
- Implies

$$Q_{\mathrm{PR},k}(\boldsymbol{N}) = B_{\mathrm{PR},k}(\boldsymbol{N}) + \sum_{r \neq k} B_{\mathrm{PR},r}(\boldsymbol{N})\, Q_{\mathrm{PR},k}(\boldsymbol{N} - 1_r)$$

$$Q_{\mathrm{NP},k}(\boldsymbol{N}) = B_{\mathrm{NP},k}(\boldsymbol{N}) + \sum_{r \neq k} B_{\mathrm{NP},r}(\boldsymbol{N})\, Q_{\mathrm{NP},k}(\boldsymbol{N} - 1_r)$$

with

$$B_{\mathrm{NP},k}(\boldsymbol{N}) = \left(\prod_{j=1}^{N} \alpha_j\right) \frac{S_k}{S}, \qquad B_{\mathrm{PR},k}(\boldsymbol{N}) = \alpha_k^{N-1} \beta_k \frac{S_k}{S}$$

- Similar recursions for throughput

Outline

Numerical examples: Balanced ($\alpha_r = \beta_r = r$)

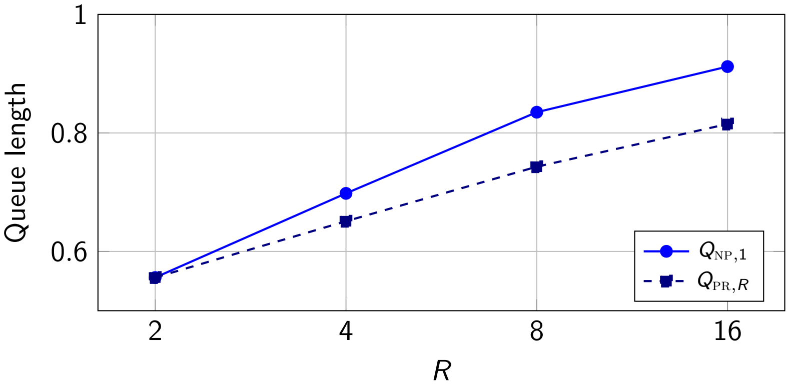

Class 1 is the *fastest* (smallest service mean), class R the *slowest*:

- Fastest class concentrates at NP; slowest drifts to PR
- Effect grows with R: at $R = 16$, $Q_{\mathrm{NP},1} = 0.91$ vs 0.50 under BCMP

Numerical examples: PR-heavy ($\alpha_r = r,\ \beta_r = r + 1$)

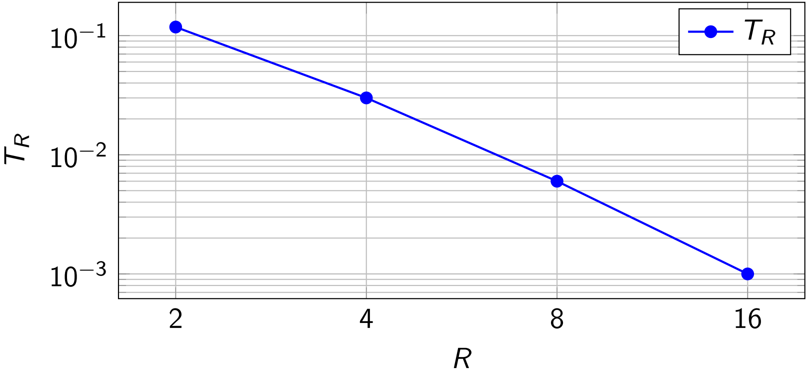

Slow classes are *starved* at the PR station:

- Throughput of the slowest class, T_R, decays by orders of magnitude with R
- Total throughput can exceed BCMP at large R

Conclusion

Extensions shown in the paper:

- PR Coxian retains product-form
- PR immediate feedback retains product-form

Future work:

- Find applications where a NP and PR combination arises naturally
- Investigate more complex networks (e.g., three stations)
- Check if specialized mechanisms retain product-form (e.g., blocking, priorities, etc)

Implemented in the open source LINE solver: `http://line-solver.sf.net`

Selected References

- G. Casale. *A family of multiclass LCFS networks with a novel product-form solution.* Queueing Systems (QUESTA), 2026.
- F. Baskett, K. M. Chandy, R. R. Muntz, F. G. Palacios. *Open, closed, and mixed networks of queues with different classes of customers.* JACM, 1975.
- F. P. Kelly. *Networks of queues with customers of different types.* J. Appl. Prob., 1975.
- M. Reiser, H. Kobayashi. *Queueing networks with multiple closed chains: theory and computational algorithms.* IBM J. Res. Dev., 1975.
- M. Reiser, S. S. Lavenberg. *Mean-value analysis of closed multichain queueing networks.* JACM, 1980.
- R. B. Bapat. *Permanents in probability and statistics.* Linear Algebra Appl., 1990.

The Permanent with Repeated Rows

- With R distinct rows, each repeated N_r times, $\sum_r N_r = N$:

$$\operatorname{perm} \boldsymbol{A} = (-1)^N \sum_{(u_1,\dots,u_R)\in\mathcal{U}} (-1)^{u_1+\cdots+u_R} \prod_{r=1}^{R} \binom{N_r}{u_r} \prod_{j=1}^{N} \Big(\sum_{i=1}^{R} u_i \, a_{i,j} \Big)$$

  where $\mathcal{U} = \{(u_1, \dots, u_R) : 0 \leq u_r \leq N_r\}$
- Cost reduces to $O\big(NR \prod_r (N_r + 1)\big)$: polynomial in N for fixed R
- Key for our setting: lets us solve models with more than one job per class

Numerical Examples: NP-heavy ($\alpha_r = r + 1,\ \beta_r = r$)

- Bottleneck at the NP station; fast classes still favor NP

Scenario	R	T	T_1	T_R	Q_{NP}	$Q_{\mathrm{NP},1}$	$Q_{\mathrm{NP},R}$	Q_{PR}	$Q_{\mathrm{PR},1}$	$Q_{\mathrm{PR},R}$
NP-heavy	2	0.345	0.207	0.138	1.345	0.724	0.621	0.655	0.276	0.379
NP-heavy	4	0.285	0.096	0.054	2.713	0.838	0.542	1.287	0.162	0.458
NP-heavy	8	0.189	0.032	0.019	5.424	0.946	0.432	2.576	0.054	0.568
NP-heavy	16	0.109	0.009	0.006	10.790	0.984	0.327	5.210	0.016	0.673

- Total queue-lengths only marginally deviate from BCMP
- No starvation: total throughput *degrades* relative to BCMP at large R

Numerical Examples: Opposite ($\alpha_r = r,\ \beta_r = R - r + 1$)

- Symmetric demands per station: $\sum_r \alpha_r = \sum_r \beta_r$

Scenario	R	T	T_1	T_R	Q_{NP}	$Q_{\mathrm{NP},1}$	$Q_{\mathrm{NP},R}$	Q_{PR}	$Q_{\mathrm{PR},1}$	$Q_{\mathrm{PR},R}$
Opposite	2	0.444	0.222	0.222	1.000	0.444	0.556	1.000	0.556	0.444
Opposite	4	0.297	0.080	0.079	2.008	0.528	0.531	1.992	0.472	0.469
Opposite	8	0.151	0.032	0.025	4.208	0.615	0.508	3.792	0.385	0.492
Opposite	16	0.070	0.018	0.009	8.565	0.567	0.529	7.435	0.433	0.471

- Asymmetric placement: $Q_{\mathrm{NP}} \neq Q_{\mathrm{PR}}$ at large R, unlike BCMP
- Throughput gap with BCMP grows with R