

Stabilization of Branching Queueing Networks

Tomáš Brázdil¹ Stefan Kiefer²

¹Masaryk University, Brno, Czech Republic

²University of Oxford, UK


STACS, Paris
02 March 2012

Queueing networks are **simple** and **general**.


Attractive for **modeling parallelism**:

- **Hardware**, especially large-scale multi-core systems:
 - Full-system performance simulators do not scale.
 - Queueing theory gives a more abstract analysis.
- **Software**, especially message passing:
 - asynchronous programs on multi-core computers
 - distributed programs on a network

Example: Router




Example: Router



Textual representation of the same thing:

$$\begin{array}{l} 0 \xrightarrow{1} A \\ A \xrightarrow{0.5} D \\ A \xrightarrow{0.25} S_1 \\ A \xrightarrow{0.25} S_2 \\ D \xrightarrow{0.7} M \\ D \xrightarrow{0.3} C \\ S_1 \xrightarrow{1} D \\ S_2 \xrightarrow{1} D \\ C \xrightarrow{1} M \\ M \xrightarrow{1} \varepsilon \end{array}$$

Example: Router



Textual representation of the same thing:


$$\begin{array}{l} 0 \xrightarrow{1} A \\ A \xrightarrow{0.5} D \\ A \xrightarrow{0.25} S_1 \\ A \xrightarrow{0.25} S_2 \\ D \xrightarrow{0.7} M \\ D \xrightarrow{0.3} C \\ S_1 \xrightarrow{1} D \\ S_2 \xrightarrow{1} D \\ C \xrightarrow{1} M \\ M \xrightarrow{1} \varepsilon \end{array}$$

Also given: **arrival rate** α_A

queue rates μ_A, μ_{S_1}, \dots of busy queues

→ **Continuous-Time Markov Chain** – **Infinite-State!**

Jackson Networks



Such networks are called **Jackson networks** (classical model).
Nice properties:

- **Stability** is easy to determine.
- **product form** solutions in “**steady state**”:


$$\Pr(S_1 = 3 \text{ and } D = 2) = \Pr(S_1 = 3) \cdot \Pr(D = 2)$$

But **shortcomings in modeling**:

- just one new task per one old task
- no (dynamic) control of the network

Branching Queueing Networks

New model: **Branching Queueing Networks (BQNs)**




Rates as before: $\alpha_A, \mu_A, \mu_B > 0$

Similar models:


- a bit like pushdown systems, but parallel
- a bit like stochastic Petri nets, but of a special form

Branching Queueing Networks

New model: **Branching Queueing Networks (BQNs)**




Rates as before: $\alpha_A, \mu_A, \mu_B > 0$




Branching Queueing Networks

New model: **Branching Queueing Networks (BQNs)**




Rates as before: $\alpha_A, \mu_A, \mu_B > 0$




Branching Queueing Networks

New model: **Branching Queueing Networks (BQNs)**



Rates as before: $\alpha_A, \mu_A, \mu_B > 0$



Topic of this talk: Is a given BQN **stable**??

stable $\stackrel{\text{def}}{=}$ expected return time to “completely empty” is finite

No Product Form

Consider the BQN



No product form: Take $\alpha_A = 1$, $\mu_A = \mu_B = \mu_C = 3$.

Then in steady-state:

$$\Pr(C \geq 1) = 1/3$$

$$\Pr(C \geq 1 \mid B \geq 1) \geq 3/7$$

No Product Form

Consider the BQN



No product form: Take $\alpha_A = 1$, $\mu_A = \mu_B = \mu_C = 3$.

Then in steady-state:

$$\Pr(C \geq 1) = 1/3$$

$$\Pr(C \geq 1 \mid B \geq 1) \geq 3/7$$

→ Analysis of BQNs **harder** than for Jackson networks

Balance Equations

$$0 \xrightarrow{1} A$$

$$\alpha_A = 0.2$$

$$A \xrightarrow{1/5} B, B$$

$$A \xrightarrow{4/5} \varepsilon$$

$$\mu_A = 0.4$$

$$B \xrightarrow{1/6} A, B$$

$$B \xrightarrow{5/6} \varepsilon$$

$$\mu_B = 0.3$$

(Pre-)Suppose a form of balance:

tasks processed at X per sec = # tasks arriving at X per sec

$$\text{Here: } \lambda_A = 0.2 + \lambda_B \cdot (1/6)$$

$$\lambda_B = \lambda_A \cdot (1/5) \cdot 2 + \lambda_B \cdot (1/6)$$

We call the **solution λ** the “**throughput**”.

Note: the speeds of the queues μ_A, μ_B do not occur.

$$\text{Here: } \lambda_A = 0.22$$

$$\lambda_B = 0.10$$

Proposition

Given a BQN.

- Let λ be the *throughput*,
i.e., λ solves the balance equations $\lambda = \alpha + M\lambda$.
Suppose $\lambda < \mu$ (in all components).
Then the *BQN is stable*.
- Conversely, if the BQN is stable,
then there is λ with $\lambda = \alpha + M\lambda < \mu$.

A Stability Result

Proposition

Given a BQN.

- Let λ be the *throughput*,
i.e., λ solves the balance equations $\lambda = \alpha + M\lambda$.
Suppose $\lambda < \mu$ (in all components).
Then the *BQN is stable*.
- Conversely, if the BQN is stable,
then there is λ with $\lambda = \alpha + M\lambda < \mu$.

Theorem

Stability of a BQN can be decided in polynomial time.

Proposition


If a BQN is stable, it is “very much so”: in steady state there is an exponential moment of the total queue size, i.e., there is $\delta > 0$ such that $\sum_{\mathbf{x} \in \mathbb{N}^n} \exp(\delta \|\mathbf{x}\|) \Pr(\mathbf{x})$ exists.

Proof of the Stability Result

The stability result requires a delicate proof (no product form).

$$0 \xrightarrow{1} A \left| \begin{array}{l} A \xrightarrow{1/5} B, B \\ A \xrightarrow{4/5} \varepsilon \end{array} \right| \begin{array}{l} B \xrightarrow{1/6} A, B \\ B \xrightarrow{5/6} \varepsilon \end{array} \left| x_B \uparrow \right.$$

Consider “drift” for all points:




Proof of the Stability Result

The stability result requires a delicate proof (no product form).

$$0 \xrightarrow{1} A \left\{ \begin{array}{l} A \xrightarrow{1/5} B, B \\ A \xrightarrow{4/5} \varepsilon \end{array} \right. \left\{ \begin{array}{l} B \xrightarrow{1/6} A, B \\ B \xrightarrow{5/6} \varepsilon \end{array} \right. x_B \uparrow$$


Consider “drift” for all points:

$$\text{For } (0, 0): [0] = \begin{pmatrix} \alpha_A \\ 0 \end{pmatrix}$$



Proof of the Stability Result

The stability result requires a delicate proof (no product form).




Consider “drift” for all points:

For $(0, 0)$: $[0] = \begin{pmatrix} \alpha_A \\ 0 \end{pmatrix}$

For $(+, 0)$: $[0A] = [0] + [A]$ with

$$[A] = \mu_A \cdot \left(\begin{pmatrix} -1 \\ 0 \end{pmatrix} + \begin{pmatrix} 0 \\ (1/5) \cdot 2 \end{pmatrix} \right)$$



Proof of the Stability Result

The stability result requires a delicate proof (no product form).

$$0 \xrightarrow{1} A \left| \begin{array}{l} A \xrightarrow{1/5} B, B \\ A \xrightarrow{4/5} \varepsilon \end{array} \right| \begin{array}{l} B \xrightarrow{1/6} A, B \\ B \xrightarrow{5/6} \varepsilon \end{array} \left| \begin{array}{l} x_B \uparrow \\ 3 \\ 2 \\ 1 \end{array} \right.$$


Consider “drift” for all points:

For $(0, 0)$: $[0] = \begin{pmatrix} \alpha_A \\ 0 \end{pmatrix}$

For $(+, 0)$: $[0A] = [0] + [A]$ with

$$[A] = \mu_A \cdot \left(\begin{pmatrix} -1 \\ 0 \end{pmatrix} + \begin{pmatrix} 0 \\ (1/5) \cdot 2 \end{pmatrix} \right)$$

For $(0, +)$: $[0B] = [0] + [B]$



Proof of the Stability Result

The stability result requires a delicate proof (no product form).

$$0 \xrightarrow{1} A \left\{ \begin{array}{l} A \xrightarrow{1/5} B, B \\ A \xrightarrow{4/5} \varepsilon \end{array} \right. \left\{ \begin{array}{l} B \xrightarrow{1/6} A, B \\ B \xrightarrow{5/6} \varepsilon \end{array} \right. \begin{array}{l} x_B \uparrow \\ 3 \\ 2 \\ 1 \\ 0 \end{array}$$

Consider “drift” for all points:


For $(0, 0)$: $[0] = \begin{pmatrix} \alpha_A \\ 0 \end{pmatrix}$

For $(+, 0)$: $[0A] = [0] + [A]$ with

$$[A] = \mu_A \cdot \left(\begin{pmatrix} -1 \\ 0 \end{pmatrix} + \begin{pmatrix} 0 \\ (1/5) \cdot 2 \end{pmatrix} \right)$$

For $(0, +)$: $[0B] = [0] + [B]$

For $(+, +)$: $[0AB] = [0] + [A] + [B]$



Proof of the Stability Result

The stability result requires a delicate proof (no product form).

$$0 \xrightarrow{1} A \left\{ \begin{array}{l} A \xrightarrow{1/5} B, B \\ A \xrightarrow{4/5} \varepsilon \end{array} \right. \left\{ \begin{array}{l} B \xrightarrow{1/6} A, B \\ B \xrightarrow{5/6} \varepsilon \end{array} \right. x_B \uparrow$$

Consider “drift” for all points:


For $(0, 0)$: $[0] = \begin{pmatrix} \alpha_A \\ 0 \end{pmatrix}$

For $(+, 0)$: $[0A] = [0] + [A]$ with
 $[A] = \mu_A \cdot \left(\begin{pmatrix} -1 \\ 0 \end{pmatrix} + \begin{pmatrix} 0 \\ (1/5) \cdot 2 \end{pmatrix} \right)$

For $(0, +)$: $[0B] = [0] + [B]$

For $(+, +)$: $[0AB] = [0] + [A] + [B]$


Find a **Lyapunov function**
 w.r.t. which the **drift** is negative



Proof of the Stability Result (cont'd)

Key steps:

- Construct **piecewise-linear Lyapunov function** w.r.t. which the **drift is negative** almost everywhere (hard).
 - use throughput λ and apply Farkas' lemma (wouldn't work for general stochastic Petri nets)
- Smooth the Lyapunov function (standard)
- Derive (strong) stability using Foster's criterion (standard)



$$0 \xrightarrow{1} A$$

$$A \xrightarrow{1/5} B, B$$

$$A \xrightarrow{4/5} \varepsilon$$

$$\sigma_1 : B \xrightarrow{1/6} A, B$$

$$B \xrightarrow{5/6} \varepsilon$$

$$B \xrightarrow{2/3} A$$

$$\sigma_2 : B \xrightarrow{1/3} \varepsilon$$

→ **Continuous-Time Markov Decision Process – Infinite-State**

Generalize balance equations:

Subdivide λ_B in $\lambda_B = \lambda_{B,1} + \lambda_{B,2}$.

Intention: $\lambda_{B,i}$ = rate of B -tasks processed according to σ_i .

$$\lambda_A = \alpha_A + \lambda_{B,1} \cdot (1/6) + \lambda_{B,2} \cdot (2/3)$$

$$\lambda_{B,1} + \lambda_{B,2} = \dots$$

Similarly to the uncontrolled case:

∃ **stabilizing** scheduler

⇔ ∃ solution λ with $\lambda_A < \mu_A$ and $\lambda_{B,1} + \lambda_{B,2} < \mu_B$

Controlled Networks

$$0 \xrightarrow{1} A$$

$$A \xrightarrow{1/5} B, B$$

$$A \xrightarrow{4/5} \varepsilon$$

$$\sigma_1 : B \xrightarrow{1/6} A, B$$

$$B \xrightarrow{5/6} \varepsilon$$

$$B \xrightarrow{2/3} A$$

$$\sigma_2 : B \xrightarrow{1/3} \varepsilon$$

→ **Continuous-Time Markov Decision Process – Infinite-State**

Generalize balance equations:

Subdivide λ_B in $\lambda_B = \lambda_{B,1} + \lambda_{B,2}$.

Intention: $\lambda_{B,i}$ = rate of B -tasks processed according to σ_i .

$$\lambda_A = \alpha_A + \lambda_{B,1} \cdot (1/6) + \lambda_{B,2} \cdot (2/3)$$

$$\lambda_{B,1} + \lambda_{B,2} = \dots$$

Similarly to the uncontrolled case:

∃ **stabilizing** scheduler

⇔ ∃ solution λ with $\lambda_A < \mu_A$ and $\lambda_{B,1} + \lambda_{B,2} < \mu_B$

LP!!

Theorem

Given a controlled BQN.

- 1 *It is decidable in polynomial time whether there exists an (arbitrary) stabilizing scheduler.*
- 2 *If it exists, one can compute in polynomial time a **static randomized** scheduler, which is stabilizing in a strong sense, i.e., in steady state there is an exponential moment of the total queue size.*

The theorem implies:

Any stabilizing scheduler can efficiently be made

- **static** and – at the same time –
- **“strongly” stabilizing.**

- Queueing networks can be used to model **parallelism**.
- Classical Jackson networks lack **branching** and **control**.
—→New model: **Branching Queueing Networks**
- **Stability** and existence of stabilizing schedulers **can be determined** in polynomial time.
- If a stabilizing scheduler does exist,
a **static** randomized scheduler suffices
and **can be computed** in polynomial time.

- Queueing networks can be used to model **parallelism**.
- Classical Jackson networks lack **branching** and **control**.
→ New model: **Branching Queueing Networks**
- **Stability** and existence of stabilizing schedulers **can be determined** in polynomial time.
- If a stabilizing scheduler does exist, a **static** randomized scheduler suffices and **can be computed** in polynomial time.

Future work:

- **Performance** beyond stability, e.g., long-term average queue size
→ Can non-static schedulers help to minimize it?

Thank you!