



HAL
open science

Perfect sampling of Jackson Queueing Networks

Ana Bušić, Bruno Gaujal, Florence Perronnin

► **To cite this version:**

Ana Bušić, Bruno Gaujal, Florence Perronnin. Perfect sampling of Jackson Queueing Networks. [Research Report] RR-8332, 2013, pp.27. hal-00851331v1

HAL Id: hal-00851331

<https://inria.hal.science/hal-00851331v1>

Submitted on 13 Aug 2013 (v1), last revised 4 Apr 2014 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Perfect Sampling of Jackson Queueing Networks

Ana Bušić, Bruno Gaujal, Florence Perronnin

**RESEARCH
REPORT**

N° 8332

July 2013

Project-Team Mescal



Perfect Sampling of Jackson Queueing Networks

Ana Bušić*, Bruno Gaujal†, Florence Perronnin‡

Project-Team Mescal

Research Report n° 8332 — July 2013 — 27 pages

Abstract: We consider open Jackson networks with losses with mixed finite and infinite queues and analyze the efficiency of sampling from their exact stationary distribution. We show that perfect sampling is possible, although the underlying Markov chain may have an infinite state space. The main idea is to use a Jackson network with infinite buffers (that has a product form stationary distribution) to bound the number of initial conditions to be considered in the coupling from the past scheme. We also provide bounds on the sampling time of this new perfect sampling algorithm for acyclic or hyperstable networks. These bounds show that the new algorithm is considerably more efficient than existing perfect samplers even in the case where all queues are finite. We illustrate this efficiency through numerical experiments. We also extend our approach to non-monotone networks such as queueing networks with negative customers.

Key-words: Perfect Simulation, Markov Chain, Jackson networks, bounding process

* Inria, Département d'Informatique de l'ENS (DI ENS) Paris, France

† Inria, Univ. Grenoble Alpes, France

‡ Université Joseph Fourier, Univ. Grenoble Alpes, France

**RESEARCH CENTRE
GRENOBLE – RHÔNE-ALPES**

Inovallée
655 avenue de l'Europe Montbonnot
38334 Saint Ismier Cedex

Simulation parfaite de réseaux de Jackson

Résumé : On considère les réseaux de Jackson avec perte comportant des files finies et infinies, et l'on s'intéresse à l'efficacité des techniques d'échantillonnage de leur distribution stationnaire exacte. Nous démontrons que la simulation parfaite est possible même si la chaîne de Markov sous-jacente a un espace d'états potentiellement infini. L'idée principale est d'utiliser un réseau de Jackson aux files infinies (qui admet une distribution de forme-produit) pour borner les conditions initiales à considérer dans l'algorithme de simulation parfaite. Nous donnons également des bornes sur le temps d'échantillonnage de ce nouvel algorithme dans le cas des réseaux acycliques, ainsi que pour des réseaux hyperstables. Ces bornes prouvent que le nouvel algorithme est considérablement plus efficace que les échantillonneurs parfaits actuels, même dans le cas où toutes les files sont finies. Nous illustrons cette efficacité par des expériences numériques. Enfin, nous généralisons notre approche au cas des réseaux non-monotones comme les réseaux aux clients négatifs.

Mots-clés : Simulation parfaite, réseaux de Jackson, chaînes de Markov, processus bornant

1 Introduction

The stationary behavior of queueing networks can only be obtained quite efficiently under specific assumptions that yield the so called product-form property; e.g., [Kelly, 1979]. This property means that the stationary probability distribution of these networks can be decomposed, up to a normalizing constant, in the product of the marginal distributions of each network node (or queue). In several cases, product-form queueing networks are restrictive because they often assume that nodes have infinite buffer sizes or that the behavior of a network node does not depend on the state of other nodes; e.g., [Jackson, 1963]. In the context of Internet networks, blocking and rejection mechanisms arise due to finite-buffer constraints and state-dependent routing. While it is possible to obtain the stationary distribution of non-product-form queueing networks through the solution of a set of linear equations, i.e., the global-balance equations [Bolch et al., 2005], the huge size of their state space makes this approach of practical interest only for small networks. In this setting, simulation is a useful approach to obtain robust measures and insights on the stationary performance.

At the cost of a slightly higher computational complexity than Monte Carlo simulation, the exact stationary distribution can be sampled in finite time using a technique called coupling from the past (see the seminal work by Propp and Wilson [Propp and Wilson, 1996]). Unlike Monte Carlo simulation, this powerful technique produces independent samples of ergodic finite Markov chains *exactly distributed* according to their stationary distribution. For this reason, this technique is also known as perfect sampling algorithm (PSA) and will be denoted by PSA in the remainder of the paper.

This technique has been used to design simulation algorithms for queueing networks with finite capacity buffers and rather general routing policies [Vincent, 2005]. The main assumption needed to make these algorithms work is that the state space is finite (or equivalently that the buffer capacities are all finite).

In this paper we present a new network simulation algorithm that can handle finite and infinite buffers at the same time, with rejection of jobs arriving at a saturated queue. We also derive a bound on its sampling time complexity and show that it does not depend on the size of the state space, when the network is acyclic (see section 4).

Related work.

As mentioned before, the original perfect sampling algorithm has been adapted for the simulation of monotone queueing networks in [Vincent, 2005]. The complexity of this algorithm has been analyzed in [Dopper et al., 2006] for acyclic networks, while cyclic networks have been studied in [Anselmi and Gaujal, 2011].

A series of papers [Gaujal et al., 2008, Pin et al., 2011, Busic et al., 2012b, Busic et al., 2008, Gaujal et al., 2011] propose new perfect sampling algorithms (introducing new ideas such as *envelopes*, *splitting* and *skipping*) for non-monotone queueing networks. However, in all cases, it is essential that the underlying Markov chain has a finite state space.

In [Kendall, 2005], the concept of bounding process is introduced. The coupling between the original process and the bounding process is different from what we propose here. Actually we could not see how to use the coupling proposed in [Kendall, 2005] in an effective algorithm.

Our technique to construct the forward and the backward chains does not use the same events for both chains but is based on the same random innovations.

This paper is an extended version of a conference paper [Busic et al., 2012a] where the backward bounding process for queueing networks was introduced. While the existence of the coupling between the forward and the backward processes was proved in our preliminary work

[Busic et al., 2012a], here, we explicitly construct the coupling between the two backward processes $Y(-t)$ and $Y^\infty(-t)$ as well as the coupling between the backward and the forward processes, $Y^\infty(-t)$ and $X^{(1)}(t)$. As a result, our main algorithm is now effective and ready for straightforward implementation. This explicit construction is also used in the proofs of Theorems 1 and 2. In this paper, we provide a detailed derivation of the complexity of the new Perfect Sampling Algorithm based on a Bounding Process (PSA-BP) in two cases: for stable acyclic networks and hyper-stable networks, given in Theorems 2 and 3 that was not present in [Busic et al., 2012a]. We also present new experiments in Section 6, showing the effect of the load in the performance of the perfect sampling algorithm as well as the tightness of the complexity bounds. Finally, the treatment of negative customers (that make the chain non-monotone), is also new.

Outline.

In section 2, we show how to use this concept for open Jackson queueing networks with mixed finite and infinite buffers and in section 3 we propose a new PSA based on this approach. We prove that this algorithm constructs samples of the exact stationary distribution of the number of jobs in all queues. In section 4 we establish a theoretical bound on the expected simulation time. Finally, in section 5 we show that in the case where all buffers are finite our algorithm clearly outperforms previous perfect samplers that all have larger simulation times. We illustrate this with numerical experiments in section 6.

2 Queueing network model

We consider an open Jackson queueing network (JQN) \mathcal{J} with M queues. The vector $\mathbf{C} = (C_1, \dots, C_M)$ denotes the buffer size of each queue. Note that for any i , $C_i \in \mathbb{N} \cup \infty$, *i.e.* the buffer capacities can either be finite or infinite. When a queue is full, all arriving jobs at this queue are lost (no blocking). For $1 \leq i \leq M$, we denote by \mathbf{e}_i the vector in \mathbb{Z}^M with all the components equal to 0, except for component i that is equal to 1.

To unify the notations, in particular with respect to the exogenous/ endogenous flows of packets, we use a modeling trick: The outside world is seen as one additional queue (numbered 0) with an infinite capacity $C_0 = \infty$, containing an infinite number of jobs initially: $X_0(0) = +\infty$.

Therefore, the input stream in the network with rate λ_0 can be seen as coming from queue 0 whose service rate is $\mu_0 \stackrel{\text{def}}{=} \lambda_0$. Figure 1 illustrates this construction.

The probability that a job joins queue i , upon arrival to the network, is p_{0i} . In queue i ($i > 0$), each job requires some processing for an exponentially distributed amount of time with mean service rate μ_i . The service discipline of each queue i is work-conserving. Upon completion of service at queue i , a job is sent to queue j with probability p_{ij} , and it is accepted if queue j has an available slot (*i.e.*, if it is non-saturated), otherwise the job is lost. The probability that a job leaves the network after service at i is p_{i0} . We make the assumption that the matrix $P = (p_{i,j})_{i,j \in \{0, \dots, M\}}$ is irreducible. With the trick about Queue 0 representing the outside world, this is a compact way to state the classical assumptions that all queues always get new jobs and that all jobs eventually leave the network.

Under the foregoing assumptions, the vector of the queue occupancies forms a continuous-time Markov chain. The state space is $\mathcal{S} \stackrel{\text{def}}{=} \{\mathbf{x} \in \mathbb{Z}^M : 0 \leq x_i \leq C_i, \forall i\}$. The notation $X(t) = (X_1(t), \dots, X_M(t))$ refers to the corresponding uniformized Markov chain in discrete time ($t \in \mathbb{N}$).

Our main notation is summarized in Table 1.

| Glossary | | |
|------------------------------|---|---|
| JQN | Jackson queueing network | |
| PSA | Perfect sampling algorithm [Propp and Wilson, 1996] | |
| PSA-BP | Perfect sampling with bounding process (proposed algorithm) | |
| Mathematical conventions | | |
| \mathbf{e}_i | Unit vector: with zero coordinates except the i -th one that is equal to 1. | |
| $\mathbb{1}_A$ | Indicator function: 1 if proposition A is true, 0 otherwise | |
| $x \wedge y$ | componentwise minimum between vectors x and y | |
| $x \vee y$ | componentwise maximum between vectors x and y | |
| \preceq | componentwise partial order | |
| System parameters | | |
| M | Number of queues | |
| C_i | Capacity of queue i ($C_i \in \mathbb{N} \cup \{+\infty\}$) | |
| \mathbf{C} | Vector of queue capacities | |
| λ_0 | Total exogenous job arrival rate (also $\mu_0 = \lambda_0$). | |
| p_{0i} | Prob. that exogenous jobs arrive in queue i | |
| p_{ij} | Routing probability from queue i to j | |
| p_{i0} | Prob. that after service completion at queue i a job leaves the system | |
| μ_i | Service rate of queue i | |
| Algorithm and proof notation | | |
| <i>Original network</i> | <i>Reversed network</i> | |
| \mathcal{J} | \mathcal{K} | Original/reversed Jackson Network (finite and infinite queues) |
| \mathcal{J}^∞ | \mathcal{K}^∞ | Bounding infinite-queue JQN (resp. reversed) |
| p_{ij} | q_{ij} | Routing prob. from queue i to j in \mathcal{J} (resp. \mathcal{K}). $q_{ij} = p_{ji}\lambda_j/\lambda_i$ |
| λ_i | | Total arrival rate in queue i ($\lambda_i = \sum_{j=0}^M \lambda_j p_{ji}$). |
| \mathcal{R} | | Set of all possible events |
| r_{ij} | | Event “job leaving queue i and joining queue j ” |
| γ_{ij} | | Rate of event r_{ij} ($\gamma_{ij} = \mu_i p_{ij}$) |
| Γ | | Total event rate ($\Gamma \stackrel{\text{def}}{=} \sum_{ij} \gamma_{ij}$) |
| a_{ij} | b_{ij} | Prob. of event r_{ij} in the original/reversed chain: $a_{ij} = \frac{\gamma_{ij}}{\Gamma}$, $b_{ij} = q_{ij} \frac{\mu_i}{\Gamma}$. |
| A_{ij} | B_{ij} | Interval of $[0,1]$ with length $a_{i,j}$ (resp. b_{ij}) |
| v_t | w_t | Coupled sequences of events for the original / reversed chains |
| \mathcal{S} | | State space of the original and reversed CTMC |
| \mathcal{S}^∞ | | State space of the infinite-queue original and reversed CTMC |
| $X(t)$ | $Y(t)$ | Uniformized, discrete-time (reversed) Markov chain |
| $X^\infty(t)$ | $Y^\infty(t)$ | Uniformized bounding DTMC for \mathcal{J}^∞ (resp. \mathcal{K}^∞) |
| $X_t^{(1)}(\cdot)$ | | Upper trajectory starting at time $-t$ |
| $X_t^{(2)}(\cdot)$ | | Lower trajectory starting at time $-t$ |
| $\phi(\mathbf{x}, r)$ | $\beta(\mathbf{x}, r)$ | Transition function for \mathcal{J} (resp. \mathcal{K}) in state \mathbf{x} under event r |
| $\phi^\infty(\mathbf{x}, r)$ | $\beta^\infty(\mathbf{x}, r)$ | Transition function for \mathcal{J}^∞ (resp. \mathcal{K}^∞) |

Table 1: Main notation and glossary




Figure 1: An open Jackson network whose input/output flows are seen as inputs/outputs from/to an additional queue with an infinite backlog.

2.1 Discrete-event definition of JQN

The JQN \mathcal{J} with M queues described above can be seen as a discrete-event system with a single type of events, namely $\{r_{ij}, i, j \in \{0, 1, \dots, M\}\}$ corresponding to the service of one job in queue i that then joins queue j . The dummy queue 0 corresponds to the outside world: an event of type r_{0j} is an exogenous arrival in queue j and an event of type r_{i0} corresponds to the departure of a job from queue i . If queue i is empty then event r_{ij} has no effect on the system. Similarly, if queue j is full (i.e. if $X_j(t) = C_j$) then event $r_{i,j}$ only affects queue i . The set of all events is denoted by \mathcal{R} .

The rate of event r_{ij} is γ_{ij} and it is independent of M and \mathbf{C} , for any i, j . Using the previous description of a JQN, for all $i, j \in \{0, 1, \dots, M\}$, $\gamma_{ij} = \mu_i p_{ij}$. The total event rate $\Gamma \stackrel{\text{def}}{=} \sum_{i,j=0}^M \gamma_{ij}$ is finite (we set arbitrarily $\gamma_{00} = 0$).

The continuous-time Markov chain described above can be transformed into a discrete-time Markov chain $\{X(t)\}_{t \in \mathbb{N}}$ with the same stationary distribution using uniformization by constant Γ . Using the foregoing assumption that the routing matrix P is irreducible, this discrete chain is irreducible and aperiodic. The evolution of the Markov chain $X(t)$ can be written under the form $X(t+1) = \phi(X(t), r_{t+1})$ where r_{t+1} is the $t+1$ -st event, distributed according to probabilities $a_{i,j} = \gamma_{ij}/\Gamma$, $1 \leq i, j \leq M$.

The state changes according to events that correspond to the routing of one packet from queue i to queue j , as defined previously.

The forward transition function of the chain $\phi : \mathcal{S} \times \mathcal{R} \rightarrow \mathcal{S}$ is defined as follows:

$$\phi(\mathbf{x}, r_{i,j}) = \mathbf{x} + (\mathbf{e}_j \mathbb{1}_{\{x_j < C_j\}} - \mathbf{e}_i) \mathbb{1}_{x_i > 0}$$

where \mathbf{e}_i is the unit vector in direction i (\mathbf{e}_0 is the zero vector) and $\mathbb{1}_{\mathcal{P}}$ equals 1 if proposition \mathcal{P} is true and 0 otherwise.

This transition function can be extended to any finite sequence $(r_1, r_2, \dots, r_k) \in \mathcal{R}^k$, $k \in \mathbb{N}$, by defining $\phi : \mathcal{S} \times [0, 1]^k \rightarrow \mathcal{S}$ recursively:

$$\phi(\mathbf{x}, r_1, r_2, \dots, r_k) \stackrel{\text{def}}{=} \phi(\phi(\mathbf{x}, r_1), r_2, \dots, r_k).$$

2.2 Network with infinite buffers

Starting from \mathcal{J} , we construct a new network \mathcal{J}^∞ that is identical to \mathcal{J} except for the buffer sizes: in \mathcal{J}^∞ all queues have infinite capacities. The state space of this new network will be denoted by $\mathcal{S}^\infty \stackrel{\text{def}}{=} \{\mathbf{x} \in \mathbb{Z}^M : x_i \geq 0, \forall i\}$.

We further assume here that each queue of the infinite JQN is stable, i.e., the total arrival rate at queue i λ_i , determined the traffic equation $\lambda_j = \sum_{i=0}^M \lambda_i p_{ij}$ for $1 \leq i \leq M$, satisfies $\lambda_i < \mu_i$ for $1 \leq i \leq M$.

The corresponding discrete-time Markov chain is denoted $\{X^\infty(t)\}_{t \in \mathbb{N}}$. The evolution of the Markov chain $X^\infty(t)$ can be written in discrete time using the same events as for the finite case, under the form $X^\infty(t+1) = \phi^\infty(X^\infty(t), r_{t+1})$ where the function ϕ^∞ is defined by:

$$\phi^\infty(\mathbf{x}, r_{i,j}) = \mathbf{x} + (\mathbf{e}_j - \mathbf{e}_i) \mathbb{1}_{x_i > 0}$$

The functions ϕ^∞ and ϕ are related in the following way.

$$\phi(x, r) = \phi^\infty(x, r) \wedge \mathbf{C} \quad \text{for all event } r.$$

where the notation $\mathbf{x} \wedge \mathbf{y}$ (resp. $\mathbf{x} \vee \mathbf{y}$) denotes the componentwise minimum (resp. maximum) of \mathbf{x} and \mathbf{y} .

As before, $\phi^\infty : \mathcal{S}^\infty \times \mathcal{R}^k \rightarrow \mathcal{S}^\infty$ is defined recursively:

$$\phi^\infty(\mathbf{x}, r_1, r_2, \dots, r_k) \stackrel{\text{def}}{=} \phi^\infty(\phi^\infty(\mathbf{x}, r_1), r_2, \dots, r_k).$$

We consider the usual product partial order of states: for $\mathbf{x}, \mathbf{y} \in \mathcal{S}^\infty$,

$$\mathbf{x} \preceq \mathbf{y} \quad \text{iff} \quad x_i \leq y_i, \quad 1 \leq i \leq M.$$

This new chain is an upper bound for the original system:

Lemma 1 *For any $r \in \mathcal{R}$, and any $\mathbf{x} \in \mathcal{S}$, functions ϕ and ϕ^∞ satisfy:*

$$\phi(\mathbf{x}, r) \preceq \phi^\infty(\mathbf{x}, r).$$

Proof. This follows directly from the fact that for any r , and any $\mathbf{x} \in \mathcal{S}$, $\phi(\mathbf{x}, r) = \phi^\infty(\mathbf{x}, r) \wedge \mathbf{C}$. \square To establish the comparison between the two chains, we also use the fact that ϕ^∞ is monotone in \mathbf{x} . (Actually, the original system is also monotone, but we do not need it for the proof of the bound.)

Lemma 2 *For any $r \in \mathcal{R}$, and any $\mathbf{x}, \mathbf{y} \in \mathcal{S}^\infty$,*

$$\mathbf{x} \preceq \mathbf{y} \Rightarrow \phi^\infty(\mathbf{x}, r) \preceq \phi^\infty(\mathbf{y}, r).$$

The proof is straightforward.

Combining the two lemmas, we get the sample path comparison of the two systems:

Proposition 1 *For any $(r_{-t+1}, \dots, r_0) \in [0, 1]^t$, and any $\mathbf{x} \in \mathcal{S}$, $\mathbf{y} \in \mathcal{S}^\infty$,*

$$\mathbf{x} \preceq \mathbf{y} \Rightarrow \phi(\mathbf{x}, r_{-t+1}, \dots, r_0) \preceq \phi^\infty(\mathbf{y}, r_{-t+1}, \dots, r_0).$$

By induction on t . For $t = 1$, and any $\mathbf{x} \preceq \mathbf{y}$, using first Lemma 1 and then Lemma 2, we have $\phi(\mathbf{x}, r_0) \preceq \phi^\infty(\mathbf{x}, r_0) \preceq \phi^\infty(\mathbf{y}, r_0)$.

Assume now the induction statement is valid for $t - 1$. Let $\mathbf{x} \preceq \mathbf{y}$, and denote $\mathbf{x}' = \phi(\mathbf{x}, r_{-t+1})$, and $\mathbf{y}' = \phi^\infty(\mathbf{y}, r_{-t+1})$. Using Lemmas 1 and 2, $\mathbf{x}' = \phi(\mathbf{x}, r_{-t+1}) \preceq \phi^\infty(\mathbf{x}, r_{-t+1}) \preceq \phi^\infty(\mathbf{y}, r_{-t+1}) = \mathbf{y}'$. Now $\phi(\mathbf{x}, r_{-t+1}, \dots, r_0) = \phi(\mathbf{x}', r_{-t+2}, \dots, r_0)$ and $\phi^\infty(\mathbf{y}, r_{-t+1}, \dots, r_0) = \phi^\infty(\mathbf{y}', r_{-t+2}, \dots, r_0)$, so

$$\phi(\mathbf{x}, r_{-t+1}, \dots, r_0) \preceq \phi^\infty(\mathbf{y}', r_{-t+2}, \dots, r_0)$$

by induction hypothesis. □

The new Markov chain $\{X^\infty(t)\}_{t \in \mathbb{N}}$ has three interesting properties:

1. As a consequence of Proposition 1, the process $\{X^\infty(t)\}_{t \in \mathbb{N}}$ is a bounding process of the original process $\{X(t)\}_{t \in \mathbb{N}}$;
2. $\{X^\infty(t)\}_{t \in \mathbb{N}}$ has a unique stationary distribution with the product form property:

$$\pi^\infty(x_1, \dots, x_M) = \prod_{i=1}^M (1 - \rho_i) \rho_i^{x_i}, \quad (1)$$

where $\rho_i \stackrel{\text{def}}{=} \frac{\lambda_i}{\mu_i}$ is the ratio of the total arrival rate in queue i , λ_i over the service rate μ_i . The arrival rate in queue i satisfies the routing balance equations:

$$\lambda_i = \sum_{j=0}^M p_{ji} \lambda_j, \quad \text{for all } i. \quad (2)$$

3. The time-reversed chain $\{Y^\infty(t)\}_{t \in \mathbb{N}}$ of the Markov chain $\{X^\infty(t)\}_{t \in \mathbb{N}}$ can be easily constructed.

Reversed chain.

The reversed chain $\{Y^\infty(t)\}_{t \in \mathbb{N}}$ can also be modeled as a Jackson network \mathcal{K}^∞ with M queues with service rate μ_i in queue i , whose routing probabilities $q_{ij} \stackrel{\text{def}}{=} p_{ji} \cdot \lambda_j / \lambda_i$, the probability that a customer leaves the system at queue i is $q_{i0} \stackrel{\text{def}}{=} p_{0i} \cdot \lambda_0 / \lambda_i$, and the exogenous arrival rate in queue j is $p_{j0} \lambda_j$.

Therefore, the set of events that modify the state are the same as in the direct chain. It is the set of routing events from one queue to another. However these events will not have the same probability as for the direct chain. For example in the Jackson network given in Figure 1, the event $r_{2,5}$ has a positive probability for the direct chain and a probability zero for the reversed chain. The probability of event r_{ij} in the reversed chain is denoted by $b_{ij} = q_{ij} \mu_i / \Gamma$.

Similarly to the forward construction of the Markov chain, the backward transition function $\beta^\infty : \mathcal{S}^\infty \times \mathcal{R} \rightarrow \mathcal{S}^\infty$ is defined as follows:

$$\beta^\infty(\mathbf{y}, r_{ij}) = \mathbf{y} + (\mathbf{e}_j - \mathbf{e}_i) \mathbb{1}_{y_i > 0} \quad \text{for all } i, j.$$

Remark 1 β^∞ has the same mathematical expression as function ϕ^∞ , because both transition functions are defined over events rather than random innovations. The dynamics of the forward and backward Jackson networks differ through the event probabilities.

Coupling the direct and the reversed chains.

The reversed chain can be coupled with the forward chain using an event correspondence. To guarantee the fact that events will be drawn with the correct probabilities in the reversed as well as in the direct way, this coupling will depend on the trajectory of a stationary reversed chain.

Let $Y^\infty(0)$ be a sample of the infinite reversed chain, generated according to its product form stationary distribution π^∞ . Given an infinite sequence of i.i.d. uniformly distributed random variables in $[0, 1]$ U_{-1}, U_{-2}, \dots , we construct an infinite sequence of events w_{-1}, w_{-2}, \dots for the reversed chain as follows:

The probability of event r_{ij} in the reversed Jackson network \mathcal{K}^∞ is b_{ij} . The corresponding intervals B_{ij} (of size b_{ij}) are ordered to cover $[0, 1]$ using the alpha-numerical order (see Fig. 2):

$$B_{0,1}, B_{0,2}, \dots, B_{0,M}, B_{1,0}, B_{1,1}, \dots, \dots, B_{M,M}.$$

Now, if $U_{-t} \in B_{i,j}$, then $w_{-t} = r_{i,j}$.

At this point, we can construct a trajectory of the reversed chain starting from $Y^\infty(0)$ and using this sequence of events: for all $t > 0$, $Y^\infty(-t) = \beta^\infty(Y^\infty(-t+1), w_{-t})$.

As for the direct chain, the coupled event sequence $v_{-1}, \dots, v_{-t}, \dots$ is built as follows:

For the direct chain, the event r_{ij} has probability a_{ij} . We also order the intervals A_{ij} of size a_{ij} in alpha-numerical order:

$$A_{0,1}, A_{0,2}, \dots, A_{0,M}, A_{1,0}, A_{1,1}, \dots, \dots, A_{M,M}.$$

- If $Y^\infty(-t) \neq Y^\infty(-t+1)$ then v_{-t} is set to $v_{-t} = (w_{-t})^{-1}$ (the inverse of routing event $r_{i,j}$ is the routing in the other direction: $(r_{i,j})^{-1} = r_{j,i}$).
- If $Y^\infty(-t) = Y^\infty(-t+1)$ then w_{-t} must be of the form $w_{-t} = r_{i_0,j}$, where queue i_0 is empty under the state $Y^\infty(-t+1)$. This means that the random variable U_{-t} has fallen in the interval $B_{i_0,j}$.

The traffic equation (2) for the Markov chain implies that

$$\forall i, \sum_j a_{ij} = \sum_j b_{ij}. \quad (3)$$

This means that the corresponding intervals in $[0, 1]$ coincide for all i : $\cup_j A_{i,j} = \cup_j B_{i,j}$, so that U_{-t} also belongs to an interval of type $A_{i_0,k}$ for the same empty queue i_0 and some k . Then, the direct event v_{-t} is set to r_{i_0k} .




Figure 2: Event probability intervals of the reversed and direct networks for a 3-queue example network.

This coupling between the direct and the reversed chain ensures the following properties:

Property 1 *The trajectory of the reversed chain from time 0 to time $-t$, starting in $Y^\infty(0)$ and the trajectory of the direct chain from time $-t$ to time 0, starting in $Y^\infty(-t)$, coincide:*

$$\forall t' \in \{0, 1, \dots, t\} \quad X^\infty(-t') = \phi^\infty(Y^\infty(-t), v_{-t}, \dots, v_{-t'+1}) = Y^\infty(-t'). \quad (4)$$

Property 2 *For any t , the probability $\mathbb{P}(v_{-t} = r_{ij}) = a_{ij}$ and v_{-t+1} is independent of v_{-t} .*

Indeed,

$$\begin{aligned} \mathbb{P}(v_{-t} = r_{ij} | v_{-t-1}) &= \mathbb{P}(v_{-t} = r_{ij} | v_{-t-1} \wedge Y^\infty(-t) = Y^\infty(-t+1)) \\ &\quad + \mathbb{P}(v_{-t} = r_{ij} | v_{-t-1} \wedge Y^\infty(-t) \neq Y^\infty(-t+1)) \\ &= \mathbb{P}(w_{-t} = r_{ji} | v_{-t-1} \wedge Y^\infty(-t) \neq Y^\infty(-t+1)) \\ &\quad + \mathbb{P}(v_{-t} = r_{ij} | v_{-t-1} \wedge Y^\infty(-t) = Y^\infty(-t+1)) \\ &= b_{ji}\rho_j + a_{ij}(1 - \rho_i) \\ &= (\mu_j q_{ji}\rho_j + \mu_i p_{ij}(1 - \rho_i)) / \Gamma \\ &= (\mu_j p_{ij} \frac{\lambda_i}{\lambda_j} \rho_j + \mu_i p_{ij}(1 - \rho_i)) / \Gamma \\ &= \mu_i p_{ij} / \Gamma \\ &= a_{ij}. \end{aligned}$$

Actually, it is even possible to simplify the coupling between the two chains because if queue i_0 is empty, it does not matter which second queue is chosen for the routing event: both r_{i_0j} and r_{i_0k} have a null effect on the state.

Proposition 2 *Here is a simplified coupling that generates the same trajectories as the previous coupling:*

If $Y^\infty(-t) \neq Y^\infty(-t+1)$ then $v_{-t} = (w_{-t})^{-1}$

If $Y^\infty(-t) = Y^\infty(-t+1)$ then $v_{-t} = w_{-t}$.

3 Perfect Sampling Algorithm with a Bounding Process

The chain $Y^\infty(-t)$ can be used to obtain a bound for the initial condition of the original chain $X(t)$, as shown in Algorithm 1. Let us describe how this algorithm works. First, generate at time 0 a stationary sample $Y^\infty(0)$ of the infinite system according to the product form distribution π^∞ .

Starting from $Y^\infty(0)$, construct a backward trajectory up to time $-t$, using the β^∞ function described above and the sequence of events $(w_{-u})_{u \in \mathbb{N}}$: $Y^\infty(-u) = \beta^\infty(Y^\infty(-u+1), w_{-u})$.

At time $-t$, start two forward simulations of the original network \mathcal{J} using the sequence of events v_{-t} , constructed using the previous coupling given in Proposition 2, and starting from states $X^{(2)}(-t) = 0$ and $X^{(1)}(-t) = Y^\infty(-t) \wedge \mathbf{C}$. If the trajectories do not coalesce at time 0 (*i.e.* $X^{(2)}(0)$ and $X^{(1)}(0)$ are not equal), then double the length of the simulation time to $2t$. As shown in the proof below, this does not introduce a bias in the output sample.

Theorem 1 *The Algorithm PSA-BP terminates with probability 1. The output of Algorithm PSA-BP is a state whose distribution is the stationary distribution of network \mathcal{J} .*

Algorithm 1: PSA with a Bounding Process (PSA-BP)**Data:**

- (1) An infinite i.i.d. sequence $\{U_{-n}\}_{n \in \mathbb{N}}$ with U_{-n} uniformly distributed in $[0, 1]$ with the associated sequence of events $\{w_{-n}\}_{n \in \mathbb{N}} : w_{-n} = r_{ij}$ if $U_{-n} \in B_{ij}$.
- (2) A state Y^∞ generated according to the distribution π^∞ .

Result: A state sampled from the stationary distribution of \mathcal{J}

begin

```

  t := 1;
  repeat
    for k = ⌊t/2⌋ to t - 1 do
      Y∞ := β∞(Y∞, w-k);
      v-k = { (w-k)-1  if Y∞ has changed
             w-k         otherwise.
      X(1) := Y∞ ∧ C; X(2) := 0;
      for i = t - 1 downto 0 do
        X(1) := φ(X(1), v-i);
        X(2) := φ(X(2), v-i);
      t := 2t;
  until X(1) = X(2);
  return X(1);

```

Proof. We first prove that the Algorithm PSA-BP terminates with probability 1. The random variable $Y^\infty(-t)$ is distributed according to the stationary distribution π^∞ . Thus, the stability assumption for chain $\{Y^\infty(u)\}_{u \in \mathbb{N}}$ implies that we have $\mathbb{P}(Y^\infty(-t) = \mathbf{0}) > 0$ (the zero state is to be understood componentwise). In that case, since $X^{(1)}(\cdot)$ and $X^{(2)}(\cdot)$ are both bounded by $\{Y^\infty(u)\}_{u \in \mathbb{N}}$, we have that $X^{(1)}(-t) = X^{(2)}(-t)$, so we also have $X^{(1)}(0) = X^{(2)}(0)$ and the algorithm terminates. By the Borel-Cantelli lemma, this happens almost surely in finite time so Algorithm PSA-BP terminates with probability 1.

Now, let us show that the output of Algorithm PSA-BP is a state whose distribution is the stationary distribution of network \mathcal{J} .

Let us consider the Markov chain $\{Y^\infty(u)\}_{u \in \mathbb{N}}$, the reversed chain of $\{X^\infty(u)\}_{u \in \mathbb{N}}$, and $\{Y(u)\}_{u \in \mathbb{N}}$, the reversed chain of $\{X(u)\}_{u \in \mathbb{N}}$, with initial conditions $Y^\infty(0)$ and $Y(0)$ respectively, distributed according to the stationary distributions π^∞ (resp. π).

Since $X^\infty(\cdot)$ is a bounding process of $X(\cdot)$, the reversed process $Y^\infty(\cdot)$ is also a bounding process of the reversed process $Y(\cdot)$. Now Strassen's theorem and the Markovian property of the reversed processes imply that there exists:

- A joint distribution ν of $Y^\infty(0)$ and $Y(0)$, with marginals π^∞ and π , and such that $Y(0) \preceq Y^\infty(0)$ a.s.,
- and a backward transition function $\beta : \mathcal{S} \times \mathcal{R} \rightarrow \mathcal{S}$ for the chain $Y(\cdot)$ is defined by:

$$\beta(\mathbf{y}, r_{ij}) = \mathbf{y} + (\mathbf{e}_j - \mathbf{e}_i) \mathbb{1}_{y_i > 0} \wedge \mathbf{C}.$$

The backward function is such that for any $\mathbf{x} \preceq \mathbf{y}$ and any $(w_{-1}, \dots, w_{-t}) \in \mathcal{R}^t$,

$$\beta(\mathbf{x}, w_{-1}, w_{-2}, \dots, w_{-t}) \preceq \beta^\infty(\mathbf{y}, w_{-1}, w_{-2}, \dots, w_{-t}). \quad (5)$$

In the proposed PSA-BP Algorithm, the variable $Y^\infty(0)$ is generated according to the stationary distribution of $X^\infty(\cdot)$. We define $Y(0)$ such that $(Y(0), Y^\infty(0))$ are distributed according to ν . We do not need to actually construct $Y(0)$. We only need to assess its existence. By definition, $Y(0)$ has the stationary distribution of \mathcal{K} . Now, for any deterministic time t , let us define $Y(-t) = \beta(Y(0), w_{-1}, w_{-2}, \dots, w_{-t})$, using the backward transition function β of the chain $Y(t)$ (see the blue dotted line in Figure 3), so that it also has the stationary distribution of \mathcal{K} . By definition of β , from relation (5) one has $Y(-t) \preceq Y^\infty(-t)$ for all t .

Now, let us define the direct processes $\{X^{(1)}(-t; u)\}_{u \geq -t}$, $\{X(-t; u)\}_{u \geq -t}$ and $\{X^{(2)}(-t; u)\}_{u \geq -t}$. The three processes start at time $-t$ with respective initial values

$$\begin{aligned} X^{(1)}(-t; -t) &= Y^\infty(-t) \wedge \mathbf{C}, \\ X(-t; -t) &= Y(-t), \\ X^{(2)}(-t; -t) &= 0. \end{aligned}$$

For any time $-u$ such that $-t \leq -u \leq 0$,

$$\begin{aligned} X^{(1)}(-t; -u) &= \phi(X^{(1)}(-t; -t), v_{-t}, \dots, v_{-u+1}); \\ X(-t; -u) &= \phi(X(-t; -t), v_{-t}, \dots, v_{-u+1}); \\ X^{(2)}(-t; -u) &= \phi(X^{(2)}(-t; -t), v_{-t}, \dots, v_{-u+1}); \end{aligned}$$

where the sequence of events $(v_{-t}, \dots, v_{-u+1})$ is defined as in Proposition 2.

By construction, $X(-t; -u)$ has the stationary distribution of \mathcal{J} , for any t and any $u \leq t$.

Also, at the starting time in the past (*i.e.* at time $-t$), one has

$$X^{(1)}(-t; -t) \succeq X(-t; -t) \succeq X^{(2)}(-t; -t).$$

Moving forward in time, the monotony of the chain implies that at time 0,

$$X^{(1)}(-t; 0) \succeq X(-t; 0) \succeq X^{(2)}(-t; 0).$$

Therefore, if $X^{(1)}(-t; 0)$ and $X^{(2)}(-t; 0)$ coalesce, they are also equal to $X(-t; 0)$. This construction is illustrated in Figure 3.

Now, let us show that this common value does not depend on t : We will prove that if $X^{(1)}(-t; 0) = X^{(2)}(-t; 0) = X(-t; 0)$, then for any $s \geq t$, then

$$X^{(1)}(-s; 0) = X^{(2)}(-s; 0) = X(-s; 0) = X(-t; 0).$$

On one hand, $X^{(2)}(-s; -t) \succeq 0 = X^{(2)}(-t; -t)$. On the other hand, $Y^\infty(-s) \succeq X^{(1)}(-s; -s)$. By monotonicity and using the coupling events, at time $-t$, $Y^\infty(-t) \succeq X^{(1)}(-s; -t)$. This implies $X^{(1)}(-t; -t) = Y^\infty(-t) \wedge \mathbf{C} \succeq X^{(1)}(-s; -t)$. Therefore, at time $-t$, the trajectories starting further in the past at time $-s$ are sandwiched by those starting at time $-t$. This imposes coalescence for all the trajectories, starting at time $-s$ at the same point, at time 0, by monotonicity. The common value at coalescence is denoted Z in the rest of the proof.

We will show that $\mathbb{P}(Z = \mathbf{x}) = \pi_{\mathbf{x}}$, the stationary probability of state \mathbf{x} for the chain $X(t)$. The proof is similar to the classical proof of the perfect sampling algorithm (see for example [Propp and Wilson, 1996]). Let $0 < \epsilon < 1$. Since coalescence occurs in finite time almost surely, there exists T_ϵ such that $\mathbb{P}(\text{no coalescence at time } -T_\epsilon) < \epsilon$. Since coalescence implies that $X(-T_\epsilon; 0) = Z$, we get

$$\mathbb{P}(Z \neq X(-T_\epsilon; 0)) \leq \mathbb{P}(\text{no coalescence at } -T_\epsilon) < \epsilon.$$




Figure 3: Illustration of the proof. The variable $Y(-t)$ has the stationary distribution of $X(t)$ and is below $Y^\infty(-t)$. For clarity, the forward trajectory $X(-t; \cdot)$ starting from $Y(-t)$ is not displayed. This trajectory would not exactly follow back the blue dotted trajectory of $Y(-t)$, unlike the upper coupled trajectories Y^∞ and $X^{(1)}(-t; \cdot)$. The forward trajectory $X(-t; \cdot)$ will remain between the two extremas $X^{(1)}(-t; \cdot)$ and $X^{(2)}(-t; \cdot)$ and therefore end up in stationary state $X^{(1)}(-t; 0)$ instead of state $Y(0)$.

On the other hand, the construction of $X(-T_\epsilon; 0)$ makes sure that it has the stationary distribution π . Therefore,

$$\begin{aligned} \mathbb{P}(Z = \mathbf{x}) - \pi_{\mathbf{x}} &= \mathbb{P}(Z = \mathbf{x}) - \mathbb{P}(X(-T_\epsilon; 0) = \mathbf{x}) \\ &\leq \mathbb{P}(Z = \mathbf{x}, X(-T_\epsilon; 0) \neq \mathbf{x}) \\ &\leq \mathbb{P}(Z \neq X(-T_\epsilon; 0)) < \epsilon. \end{aligned}$$

Similarly,

$$\begin{aligned} \pi_{\mathbf{x}} - \mathbb{P}(Z = \mathbf{x}) &= \mathbb{P}(X(T_\epsilon; 0) = \mathbf{x}) - \mathbb{P}(Z = \mathbf{x}) \\ &\leq \mathbb{P}(Z \neq \mathbf{x}, X(-T_\epsilon; 0) = \mathbf{x}) \\ &\leq \mathbb{P}(X(-T_\epsilon; 0) \neq Z) < \epsilon. \end{aligned}$$

This is true for all ϵ so that $\mathbb{P}(Z = \mathbf{x}) = \pi_{\mathbf{x}}$. □

4 Complexity analysis

The simulation time of PSA-BP can be decomposed into three steps:

1. The generation of a sample $Y^\infty(0)$ with the stationary distribution of the bounding process. Since this distribution is product form and since for each queue the geometric law can be easily sampled in constant time, the time complexity of this first step is $\mathcal{O}(M)$.
2. To obtain random events from the innovation sequence $\{U_{-n}\}_{n \in \mathbb{N}}$, we use the Walker's alias method [Walker, 1977] (for sampling discrete random variables in $\mathcal{O}(1)$ time). The construction of alias table is linear in number of events, so this pretreatment can be done

in $\mathcal{O}(M^2)$ time. Given the events $r_{i,j}$, the values of $\phi(\mathbf{x}, r)$ and $\beta^\infty(\mathbf{x}, r)$ can be computed in $\mathcal{O}(1)$ time.

3. The construction of the backward trajectory of Y^∞ and the forward trajectories of $X^{(1)}$ and $X^{(2)}$ have the same expected duration, smaller than $10\mathbb{E}\tau(X^{(2)}, X^{(1)})$, where $\tau(X^{(2)}, X^{(1)})$ is the coalescence time of two trajectories of the Markov chain $X(t)$, starting in $X^{(1)}$ and $X^{(2)}$ respectively, under the same sequence of events. The multiplication factor 10 comes from a simple calculation: 3 trajectories, each with length at most double the coalescence time (worst case), and two of them (for $X^{(1)}$ and $X^{(2)}$) requiring a computation time equal to 2τ to compute a τ -long trajectory ($1 + 2 + 4 + \dots + \tau/2 + \tau \approx 2\tau$ steps).

Therefore, the complexity essentially depends linearly on the coalescence time $\tau(X^{(2)}, X^{(1)})$.

4.1 The acyclic case

In the following, we will show that in the acyclic case, the coalescence time can be bounded by a quadratic function of the expected size of the queues in the stationary regime. Simulations (see Section 6) suggest that this bound is tight and holds for all stable networks, including cyclic ones.

Theorem 2 *If \mathcal{J} is acyclic and stable, then the expected coalescence time of algorithm PSA-BP, $\tau_{\text{PSA-BP}}$ satisfies*

$$\mathbb{E}[\tau_{\text{PSA-BP}}] \leq \sum_{i=1}^M \frac{\Gamma}{\lambda_i} (\mathbb{E}[Y_i^\infty])^2, \quad (6)$$

where μ_i is the service rate in queue i , λ_i is the input rate in queue i , $\mathbb{E}[Y_i^\infty] = \frac{\lambda_i}{\mu_i - \lambda_i}$ is the expected stationary size of queue i of the infinite system \mathcal{J}^∞ , and Γ is the uniformization constant, $\Gamma = \sum_i (\mu_i + \lambda_i)$.

Remark 2 *The expected coalescence time grows as the square root of the sizes of the queues when the load increases, since the multiplicative coefficients $\frac{\Gamma}{\lambda_i}$ remain bounded when the load goes to 1.*

The proof is carried in several steps.

Lemma 3 *Let $T(s)$ be the expected time needed to empty a stable queue in discrete time with arrival (resp. service) probability $\frac{\lambda}{\lambda+\mu}$ (resp. $\frac{\mu}{\lambda+\mu}$), starting with s customers. Then $T(s) = s \frac{\lambda+\mu}{\mu-\lambda}$.*

(of the lemma). This can be considered as a classical result, the proof is given nevertheless. $T(s)$ satisfies the following one-step recurrence: $T(s) = 1 + \frac{\lambda}{\lambda+\mu}T(s+1) + \frac{\mu}{\lambda+\mu}T(s-1)$, whose solution is of the form $T(s) = C_1 \left(\frac{\mu}{\lambda}\right)^s + C_2 s$. The constants C_2 and C_1 are functions of $T(0)$ and $T(1)$. By definition, $T(0) = 0$ and $T(1)$ is such that $C_1 = 0$ (otherwise the queue cannot be stable). This implies $T(1) = \frac{\lambda+\mu}{\mu-\lambda}$. and $T(s) = s \frac{\lambda+\mu}{\mu-\lambda}$.

Another proof: The quantity $T(x)$ satisfies the recurrence: $T(x) = BP + T(x-1)$, where BP is the expected busy period of the slotted queue, known to be equal to $\frac{\lambda+\mu}{\mu-\lambda}$. The fact that $T(0) = 0$ implies that $T(x) = x \frac{\lambda+\mu}{\mu-\lambda}$. □

Let us consider the reversed network \mathcal{K}^∞ . Since \mathcal{J} is stable and acyclic, \mathcal{K}^∞ is also stable and acyclic. The queues are ordered according to the topological order induced by the network

\mathcal{K}^∞ . In each queue i of \mathcal{K}^∞ , the service rate is μ_i and the input rate is λ_i , defined by the linear system $\lambda_i = \sum_j \lambda_j q_{j,i}$ under its stationary regime.

The stability of \mathcal{K}^∞ can be translated as for all i , $\lambda_i < \mu_i$. Now, for any initial state x in \mathcal{K}^∞ , let us define the time $h(x)$ as follows: $h(x)$ is the first instant when the following sequence of events have occurred, in that order: Queue 1 hits 0 then Queue 2 hits 0 then ... then Queue M hits 0. The stability of \mathcal{K}^∞ implies that $h(x)$ is finite, with a finite expectation, for all x .

Lemma 4 *The expectation $\mathbb{E}[h(Y^\infty(0))]$, where $Y^\infty(0)$ has the stationary distribution of \mathcal{K}^∞ , is linear:*

$$\mathbb{E}[h(Y^\infty(0))] = \sum_i \frac{\Gamma}{\mu_i - \lambda_i} \mathbb{E}[Y_i^\infty(0)].$$

(of the lemma).

The proof holds by induction on the number of queues. The case $M = 1$ follows directly from Lemma 3:

$$\mathbb{E}[h(Y^\infty(0))] = \mathbb{E}[T(Y^\infty(0))] = \frac{\lambda + \mu}{\mu - \lambda} \mathbb{E}[Y^\infty(0)]$$

and $\Gamma = \lambda + \mu$.

Let us consider an acyclic network with M queues and let us cut the last queue (Queue M). The time $h(Y^\infty(0))$ is equal to the time when the first $M - 1$ queues reach 0 plus the time for the last queue to become empty from that point on. By Burke's Theorem, when the initial state S is stationary, then the state of queue M at any time is independent of the state of the other queues. Therefore, at hitting times of the other queues, the state of queue M is distributed according to the stationary distribution, and the input process in queue M is Poisson with rate λ_M .

Let us call Y_M the state of Queue M just after the first $M - 1$ st queues have emptied (in that order). Its distribution is $\mathbb{P}(Y_M = i | Y_{M-1} = 0) = \mathbb{P}(Y_M = i)$, by independence of the queues under the stationary regime, so that its distribution is stationary. From that point on, Queue M evolves as an independent queue with arrival rate λ_M and service rate μ_M . Using Lemma 3 the expected time to hit 0 for this last queue, under uniformization Γ , from that time on is:

$$\frac{\Gamma}{\mu_M - \lambda_M} \mathbb{E}[Y_M^\infty(0)].$$

As for the $M - 1$ first queues, the induction assumption says that the expected time for all queues to hit 0 is

$$\sum_{i=1}^{M-1} \frac{\Gamma}{\mu_i - \lambda_i} \mathbb{E}[Y_i^\infty(0)].$$

Adding both quantities gives $\mathbb{E}[h(Y^\infty(0))] = \sum_{i=1}^M \frac{\Gamma}{\mu_i - \lambda_i} \mathbb{E}[Y_i^\infty(0)]$. □

We are now ready for the proof of Theorem 2.

(of Theorem 2).

Let us consider the state reached by the reversed network, starting in state $Y^\infty(0)$, at time $-h(Y^\infty(0))$, namely $Y^\infty(-h(Y^\infty(0)))$ (for simplicity, we will denote this hitting time by $h^* \stackrel{\text{def}}{=} h(Y^\infty(0))$).

Now, let us consider the initial network \mathcal{J} starting at time $-h^*$ in state $Y^\infty(-h^*) \wedge \mathbf{C}$. Because of the coupling between the reversed and direct sequences of events, the trajectory $X^{(1)}(-h^*; -t)$ remains below (componentwise) the reversed trajectory, $Y^\infty(-t)$ for any t . This means that the following sequence of events must occur in the trajectory $X^{(1)}(-h^*; -t)$: queue M hits 0, then queue $M - 1$ hits 0, ..., queue 1 hits 0.

Let us now consider the direct trajectory starting in 0 at time $-h^*$, namely $X^{(2)}(-h^*; -t)$. By monotonicity, for all t , $X^{(2)}(-h^*; -t) \preceq X^{(1)}(-h^*; -t)$. Therefore, when queue M hits 0 in $X^{(1)}(-h^*; -t)$, then queue M also hits 0 in $X^{(2)}(-h^*; -t)$ at the same instant. This means that the state of queue M is the same on both trajectories. By acyclicity, from that point on, the M -th coordinate of $X^{(2)}(-h^*; -t)$ and $X^{(1)}(-h^*; -t)$ will remain equal. When queue $M-1$ hits 0, the same phenomenon will happen: The $M-1$ -st coordinate of $X^{(2)}(-h^*; -t)$ and $X^{(1)}(-h^*; -t)$ will remain equal from that point on, and so forth. This implies that the two trajectories must coalesce before time 0.

This means that the forward coalescence time $\tau(X^{(2)}, X^{(1)})$ is smaller than $h^* = h(Y^\infty(0))$. Using Lemma 4, if $Y^\infty(0)$ is stationary, then the expected coalescence time $\mathbb{E}\tau(X^{(2)}, X^{(1)})$ satisfies the required inequality:

$$\mathbb{E}[\tau_{\text{PSA-BP}}] \leq \mathbb{E}[h(Y^\infty(0))] = \sum_{i=1}^M \frac{\Gamma}{\mu_i - \lambda_i} \mathbb{E}[Y_i^\infty(0)] = \sum_{i=1}^M \frac{\Gamma}{\lambda_i} (\mathbb{E}[Y_i^\infty(0)])^2.$$

This concludes the proof. □

4.2 The hyperstable case

In this section, we consider the case where the network has an arbitrary topology (that may contain cycles) but where all queues are hyperstable (to be defined later).

In this case we also obtain a bound on the expected coupling time. This new bound is cubic in the average queue size (instead of quadratic in the acyclic case). Notice that this bound remains logarithmic with respect to the size of the state space. This shows that the algorithm PSA-BP remains very efficient even for cyclic networks.

Let $\tau_{\text{PSA-BP}}$ be the coalescence time of algorithm PSA-BP. It is a finite random variable with finite expectation when the network is stable.

On the other hand, let us define $\tau_f(\mathbf{x}, \mathbf{y})$ the meeting time of two trajectories of the Jackson network \mathcal{J} , starting in states \mathbf{x} and \mathbf{y} respectively. This is also a finite random variable with finite expectation when the network is stable.

The next lemma is a quite general result, not based on hyper stability.

Lemma 5 *Under the foregoing assumptions, if \mathcal{J} is stable, then $\tau_f(\mathbf{0}, Y^\infty(0) \wedge \mathbf{C})$ and $\tau_{\text{PSA-BP}}$ have the same distribution.*

Proof. Let T be a fixed duration. Using the notations introduced in the proof of Theorem 1,

$$\mathbb{P}(\tau_{\text{PSA-BP}} < T) = \mathbb{P}(X^{(1)}(-T; 0) = X^{(2)}(-T; 0)) \tag{7}$$

$$= \mathbb{P}(\phi(\mathbf{0}, v_{-T+1}, \dots, v_0) = \phi(Y^\infty(-T) \wedge \mathbf{C}, v_{-T+1}, \dots, v_0)) \tag{8}$$

$$= \mathbb{P}(\phi(\mathbf{0}, v_{-T+1}, \dots, v_0) = \phi(Y^\infty(0) \wedge \mathbf{C}, v_{-T+1}, \dots, v_0)) \tag{9}$$

$$= \mathbb{P}(\phi(\mathbf{0}, u_1, \dots, u_T) = \phi(Y^\infty(0) \wedge \mathbf{C}, u_1, \dots, u_T)) \tag{10}$$

$$= \mathbb{P}(\tau_f(\mathbf{0}, Y^\infty(0) \wedge \mathbf{C}) < T), \tag{11}$$

where u_1, \dots, u_T is a sequence of *iid* events such that for all t , $\mathbb{P}(u_t = r_{ij}) = a_{ij}$. Equation (9) comes from the fact that for any T , $Y^\infty(T)$ is stationary, Equation (10) is true because the sequences u_1, \dots, u_T and v_{-T+1}, \dots, v_0 have the same distribution by Property 2 and the last equation (11) follows from the definition of τ_f . □

Now, let us now consider the hyper-stable case.

Definition 1 A queue $j \geq 1$ in a Jackson network \mathcal{J} is hyper-stable if $\sum_{i=0}^M \gamma_{ij} < \mu_j$. The network is hyper-stable if all queues $j \geq 1$ are hyper-stable. Note that hyper-stability implies stability since $\lambda_j \leq \sum_{i=0}^M \gamma_{ij}$, $\forall j \geq 1$.

Theorem 3 If \mathcal{J} is hyper-stable then the expected coalescence time of algorithm PSA-BP satisfies $\mathbb{E}\tau_{\text{PSA-BP}} \leq cM\Gamma \max_i \left(\frac{\mathbb{E}Y_i^\infty}{\lambda_i} \right) \left(\sum_{j=1}^M \mathbb{E}Y_j^\infty \right)^2$, where c is a constant (that only depends on the routing probabilities (p_{ij}) and $\mathbb{E}Y_j^\infty$ is the expected stationary size of queue j of the infinite system \mathcal{J}^∞ ($\mathbb{E}Y_i^\infty = \frac{\lambda_i}{\mu_i - \lambda_i}$).

Proof. Using Lemma 5, we will consider $\tau_f(\mathbf{0}, Y^\infty(0) \wedge \mathbf{C})$ instead of $\tau_{\text{PSA-BP}}$ since they have the same distribution. Now, in [Anselmi and Gaujal, 2011], it was proved (proof of Theorem 2 in [Anselmi and Gaujal, 2011]) that for hyperstable Jackson networks with M queues, there exists a constant c such that for all i and for any state \mathbf{y} ,

$$\mathbb{E}\tau_f(\mathbf{0}, \mathbf{y} + \mathbf{e}_i) - \mathbb{E}\tau_f(\mathbf{0}, \mathbf{y}) \leq cM\Gamma \max_i \left(\frac{\mathbb{E}Y_i^\infty}{\lambda_i} \right) \sum_{j=1}^M \frac{\lambda_j}{\mu_j - \lambda_j}.$$

By iterating this inequality and conditioning on the value of $Y^\infty(0)$, one gets

$$\mathbb{E}[\tau_f(\mathbf{0}, Y^\infty(0)) | Y^\infty(0)] \leq cM\Gamma \max_i \left(\frac{\mathbb{E}Y_i^\infty}{\lambda_i} \right) \left(\sum_{j=1}^M \frac{\lambda_j}{\mu_j - \lambda_j} \right) \sum_i Y_i^\infty(0).$$

Finally,

$$\begin{aligned} \mathbb{E}[\tau_f(\mathbf{0}, Y^\infty(0) \wedge \mathbf{C})] &\leq \mathbb{E}[\tau_f(\mathbf{0}, Y^\infty(0))] \\ &= \mathbb{E} \left[\mathbb{E}[\tau_f(\mathbf{0}, Y^\infty(0)) | Y^\infty(0)] \right] \\ &\leq M c \Gamma \max_i \left(\frac{\mathbb{E}Y_i^\infty}{\lambda_i} \right) \left(\sum_{j=1}^M \frac{\lambda_j}{\mu_j - \lambda_j} \right) \sum_i \mathbb{E}Y_i^\infty(0) \\ &= M c \Gamma \max_i \left(\frac{\mathbb{E}Y_i^\infty}{\lambda_i} \right) \left(\sum_{j=1}^M \mathbb{E}Y_j^\infty \right)^2. \end{aligned}$$

□

5 Comparison with the classical perfect sampler

As discussed earlier, our new algorithm allows for perfect sampling of Jackson networks with finite and infinite buffers. This was not possible with the classical perfect sampling algorithms that requires a finite state space. As such, our algorithm broadens the scope of perfect sampling techniques. In addition, this new approach also provides time improvements in the case where all buffers are finite, because the time complexity of the new approach is essentially independent of the capacities. In that case, it reduces the sampling time by a factor corresponding to the ratio between the maximum capacity of the buffers over the expected size of the queues under the stationary law of the bounding process, at least in the acyclic case (the comparison in the cyclic case is not as striking because of the squared term but behaves essentially in the same way).

Algorithm 2: Coupling from the past (PSA) [Propp and Wilson, 1996]

Data: An infinite i.i.d. sequence $\{U_{-n}\}_{n \in \mathbb{N}}$ with U_{-k} uniformly distributed in $[0, 1]$ (with the associated sequence of events $\{u_{-n}\}_{n \in \mathbb{N}}$).

Result: A state sampled from the stationary distribution of the JQN

begin

```

  t := 1;
  repeat
    Z(1) := C; Z(2) := 0;
    for i = t - 1 downto 0 do
      Z(1) := φ(Z(1), u-i);
      Z(2) := φ(Z(2), u-i);
    t := 2t;
  until Z(1) = Z(2);
  return Z(1);

```

Let us first recall the classical perfect sampler for monotone finite Markov chains derived from Jackson queueing networks.

It is well known that Algorithm 2 outputs samples of the stationary distribution of the network as long as all buffers are finite. Otherwise, if $X^{(1)}$ has at least one infinite coordinate, the algorithm does not converge.

Also note that, in the general case (acyclic or cyclic), the coalescence time of PSA-BP is always stochastically smaller than the coalescence time of PSA since the extreme starting point $X^{(1)}(-t) = Y^\infty(-t) \wedge \mathbf{C}$ of the former is always smaller than the starting time of the latter, $Z^{(1)}(-t) = \mathbf{C}$.

Finally, the time complexity in the case where \mathcal{J} is acyclic with finite queues has been studied in [Dopper et al., 2006]. The expected coalescence time to get one sample τ_{PSA} is shown to satisfy

$$\mathbb{E}[\tau_{\text{PSA}}] \leq \sum_{i=1}^M \frac{\Gamma}{\mu_i + \lambda_i} (C_i^2 + C_i).$$

It is also easy to find a lower bound on the expected coalescence time:

$$\mathbb{E}[\tau_{\text{PSA}}] \geq \sum_{i=1}^M \frac{\Gamma}{\mu_i + \lambda_i} C_i.$$

Indeed, one needs at least C_i departures in queues i for coalescence and the average time between two events in queue i is $\frac{\Gamma}{\mu_i + \lambda_i}$.

As shown in Theorem 2, the expected coalescence time of our new sampler is upper bounded:

$$\mathbb{E}[\tau_{\text{PSA-BP}}] \leq \sum_{i=1}^M \frac{\Gamma}{\lambda_i} \mathbb{E}[Y_i^\infty]^2$$

(where $\mathbb{E}Y_i^\infty$ is the expected stationary size of queue i in the infinite system \mathcal{J}^∞). This may result in a considerable improvement over the classical algorithm.

Indeed, the upper and lower bounds above induce the following comparison between the

coalescence time of the two algorithms:

$$\mathbb{E}[\tau_{\text{PSA}} - \tau_{\text{PSA-BP}}] \geq \max \left(0, \sum_{i=1}^M \Gamma \left(\frac{C_i}{\lambda_i + \mu_i} - \frac{\mathbb{E}[Y_i^\infty]^2}{\lambda_i} \right) \right).$$

If the arrival rate is 0.8 in every queue with service rate 1, (this is a typical situation), then $\mathbb{E}[Y_i^\infty] \leq 4$, and each term in the difference is larger than $C_i/2 - 20$ which can be very large when the buffer size of the order of 100, 1000 or more.

In the next section, we show on a numerical example that the PSA-BP algorithm 1 drastically outperforms the classical PSA Algorithm 2.

6 Numerical experiments

We have implemented algorithms 1 and 2 (PSA-BP and PSA) in C, in order to:

- demonstrate the feasibility of our approach,
- quantify the actual CPU time speedup of our approach over classic PSA,
- and estimate the tightness of the complexity bounds derived in section 4.

Experimental setup

We tested both simulation algorithms over a circular queueing network displayed in Figure 4 : the M queues form a cycle, with the same service time in each queue, $\mu_i = \mu$, the same exogenous arrival rate in each queue (λ) and the same probability of leaving the system at each queue: $p_{i0} = p$. In that case, the infinite-queue system is stable if and only if $\lambda/p < \mu$. To be consistent with the notation in Table 1, we have $\lambda_0 = M \times \lambda$.




Figure 4: Circular network \mathcal{J} used in the numerical experiments.

We have run both simulators for various queue capacities $C_i = C$ for $1 \leq i \leq M$ ranging from 1 to 600, and various traffic loads $\rho_i = \rho$ ranging from 0.6 to .99.

Our experiments were carried out on a standard laptop computer¹. All results are given with 95 % confidence intervals computed with 300 experiments per setting.

6.1 Sampling time improvement over classic PSA


The complexity analysis in Section 4 shows that the coalescence time of PSA-BP is always smaller, however we first wanted to verify that

¹Intel Core 2 Duo CPU U9400 1.40GHz, Cache 3072KB with 4.8GB RAM. Only one core is used for each sample.

- the practical sampling time (CPU time) is well predicted by the coalescence time,
- the CPU time is significantly better (the overhead due to the construction of the bounding process remains negligible).

To this end we measured the running time of both algorithms (PSA-BP and classic PSA) for the same circular network, over a wide range of queue capacities.

Figure 5 presents the CPU time ² needed per run for both PSA and PSA-BP algorithms, under average (Fig. 5(a)) and high loads (Fig. 5(b)). The capacity C of the queues ranges from 1 to 600. These capacity choices are to be compared to the expected stationary queue size of the bounding infinite-queue network, which is $\frac{\rho}{1-\rho}$, i.e. 4 jobs (resp. 49 jobs) for a 80% load (resp. 98%).



(a) Low traffic intensity ($\lambda = 0.4$) $\mathbb{E}[Y^\infty] = 4$ jobs (b) High traffic intensity ($\lambda = 0.49$) $\mathbb{E}[Y^\infty] = 49$ jobs

Figure 5: Comparative sampling time (CPU time consumption per sample) of Bounding-Process PSA and Classic PSA. Circular networks with $M=3$ queues, $\mu = 1$, $p_{i0} = 0.5$. The capacity C of each queue varies from 1 to 600. 95% confidence intervals with 300 samples per setting.

We observe that the simulation time of the classical perfect sampling algorithm grows somewhat linearly with the capacity of the buffers, as does the coalescence time (not displayed here to avoid redundancy). The succession of plateaus and steep slopes (staircase shape) comes from the doubling period trick of the perfect simulation algorithm: the coalescence time of Algorithm 2 is the smallest power of 2 larger than the actual coalescence time of the trajectory. This linear increase behavior, while the state space grows with C^M , is consistent with the findings in [Anselmi and Gaujal, 2011].

As for the proposed algorithm 1, Figure 5 perfectly illustrates the fact that the CPU sampling time remains bounded when the state space increases, as shown for the expected coalescence time by Theorem 2 in the acyclic case.

The performance gap with the classical algorithm becomes significant as soon as the capacity becomes larger than a few times the bounding expected stationary queue size ($C \geq 10$ for Figure

²CPU time used by the process computing one sample, measured with the `CLOCK_PROCESS_CPUTIME_ID` of the `time.h` library.

5(a), $C \geq 200$ for Figure 5(b)). Indeed, when C is small, the bounding infinite-queue process is too large and algorithm PSA-BP roughly amounts to classic PSA, except for the computation of the bounding trajectory. Even in this unfavorable case, we do not observe any significant overhead effect.

6.2 Complexity bound




Figure 6: Asymptotic behavior of (the square root of) the coalescence time (number of steps) of our PSA-BP algorithm as the expected queue size grows (increasing with load ρ). Notice that due to the doubling period artefact, the observed number of steps is at most twice the actual coalescence time of the simulated trajectories. Circular networks with $M=3$ queues, capacity $C = 500$, $\mu = 1$, $p_{i0} = 0.5$.

In Theorems 2 and 3, upper bounds on the coalescence time are given in the acyclic and hyper-stable cases, that are respectively quadratic and cubic in the expected size of the queues for the bounding network \mathcal{J}^∞ .

On the other hand, it is rather straightforward to show that for any stable network \mathcal{J} , the coalescence time is lower bounded by a sub-linear combination of the expected sizes of the queues in the network:

$$\mathbb{E}[\tau_{\text{PSA-BP}}] \geq \sum_{i=1}^M \frac{\Gamma}{\mu_i + \lambda_i} \mathbb{E}[Y_i^\infty \wedge C_i].$$

Indeed, the coalescence time in the network has the same distribution as the forward coalescence time $\tau_f(\mathbf{0}, \mathbf{Y}^\infty \wedge \mathbf{C})$ (see Lemma 5). One needs at least $Y_i^\infty \wedge C_i$ departures in queue i for coalescence and the average time between two events in queue i is $\frac{\Gamma}{\mu_i + \lambda_i}$. Therefore in expectation,

$$\mathbb{E}[\tau_f(\mathbf{0}, \mathbf{Y}^\infty \wedge \mathbf{C})] \geq \sum_{i=1}^M \frac{\Gamma}{\mu_i + \lambda_i} \mathbb{E}[Y_i^\infty \wedge C_i].$$

To test if the actual coalescence time behaves like the lower bound (linear) or the upper bound (quadratic or cubic), or does not comply with any of them in the cyclic non-hyperstable

case, we have run several experiments by letting the load grow from 50 % to 98 % and we measured the coalescence time versus the average size of the queues in the circular network of Figure 4. The results are given in Figure 6. The behavior of the square root of the coalescence time follows almost perfectly the average size of the queues, up to a multiplicative constant. This shows a quadratic behavior of the coalescence time, even though the network is not acyclic (nor hyperstable). We conjecture that the bound given in Theorem 2 is actually true for any stable network and that it is tight.

7 Extension to non-monotone networks

The perfect sampling technique presented here can be extended to non-monotone chains that admit a monotone bounding chain, by combining Algorithm 1 with the envelope technique presented in [Busic et al., 2012b].

We illustrate this extension on the case of a network with negative customers [Gelenbe, 1991]. Let \mathcal{N} be a network with finite and infinite buffer capacities and positive and negative customers: the arrival of a negative customer in a buffer decreases the buffer size by one (unless the buffer is empty), instead of increasing it by one for regular (positive) customers.

Denote by $p_{i,j}^+$ (resp. $p_{i,j}^-$) the probability that after a service in queue i the customer goes to queue j as a positive (resp. negative) customer. The discrete-event representation for this system can be obtained from the one for JQN, by dividing each event $r_{i,j}$ ($i, j \in \{0, 1, \dots, M\}$) into two new events, denoted by $r_{i,j}^+$ and $r_{i,j}^-$ of rates $\gamma_{i,j}^+ := \mu p_{i,j}^+$ and $\gamma_{i,j}^- := \mu p_{i,j}^-$. The total event rate is $\Gamma \stackrel{\text{def}}{=} \sum_{i,j=0}^M (\gamma_{i,j}^+ + \gamma_{i,j}^-)$.

We assume here that the matrix $P^+ = \{p_{i,j}^+\}_{i,j \in \{0,1,\dots,M\}}$ is irreducible. Furthermore, we assume positive stability, *i.e.*

$$\lambda_i < \mu_i, \quad 1 \leq i \leq M, \quad (12)$$

where λ_i satisfy the equation $\lambda_j = \sum_{i=0}^M \lambda_i p_{i,j}^+$.

As before, we consider a discrete-time Markov chain $\{Z(t)\}_{t \in \mathbb{N}}$, after uniformization with constant Γ . Using the assumption that the routing matrix P^+ is irreducible, this discrete chain is irreducible and aperiodic. and its evolution is given by: $Z(t+1) = \psi(Z(t), v_{t+1})$ where $(v_t)_{t \geq 0}$ are events (routing positive or negative customers, *i.e.* $r_{i,j}^+$ or $r_{i,j}^-$, $0 \leq i, j \leq M$). The forward transition function of the chain $\psi : \mathcal{S} \times \mathcal{R} \rightarrow \mathcal{S}$ is defined as follows:

- $\psi(\mathbf{x}, r_{ij}^+) = ((\mathbf{x} - \mathbf{e}_i) \vee 0 + \mathbf{e}_j \mathbb{1}_{\mathbf{x} - \mathbf{e}_i \geq 0}) \wedge \mathbf{C}$;
- $\psi(\mathbf{x}, r_{ij}^-) = (\mathbf{x} - \mathbf{e}_i - \mathbf{e}_j \mathbb{1}_{\mathbf{x} - \mathbf{e}_i \geq 0}) \vee 0$;
- $\psi(\mathbf{x}, r_{0j}^+) = (\mathbf{x} + \mathbf{e}_j) \wedge \mathbf{C}$;
- $\psi(\mathbf{x}, r_{0j}^-) = (\mathbf{x} - \mathbf{e}_j) \vee 0$;
- $\psi(\mathbf{x}, r_{i0}^+) = (\mathbf{x} - \mathbf{e}_i) \vee 0$.
- $\psi(\mathbf{x}, r_{i0}^-) = (\mathbf{x} - \mathbf{e}_i) \vee 0$.

Events $r_{i,j}^-$ are not monotone for $1 \leq i, j \leq M$. To see this, consider the states $\mathbf{x}_1 = \mathbf{e}_j$ and $\mathbf{x}_2 = \mathbf{e}_i + \mathbf{e}_j$. The inequality $\mathbf{x}_1 \prec \mathbf{x}_2$ holds. By definition, $\psi(\mathbf{x}_1, r_{ij}^-) = \mathbf{e}_j$ and $\psi(\mathbf{x}_2, r_{ij}^-) = \mathbf{0}$, so $\psi(\mathbf{x}_2, r_{ij}^-) \prec \psi(\mathbf{x}_1, r_{ij}^-)$.

It is nevertheless possible to construct a new network \mathcal{J}^∞ with infinite buffers in all queues and where negative customers are replaced by departures (events $r_{i,j}^-$ are replaced by events r_{i0}^+).

Network \mathcal{J}^∞ corresponds to a classical Jackson network with routing probabilities $p_{i,j} := p_{i,j}^+$, for $0 \leq i \leq M$ and $1 \leq j \leq M$, and $p_{i,0} := p_{i,0}^+ + \sum_{j=0}^M p_{i,j}^-$. The state of \mathcal{J}^∞ is a Markov chain whose forward transition function is denoted ϕ^∞ . Under assumption (12), network \mathcal{J}^∞ is stable. The network \mathcal{J}^∞ is also clearly monotone and with the coupling of negative customer routing with departures, it is an upper bound of \mathcal{N} , so Proposition 1 remains valid.

This allows us to construct a perfect sampler of network \mathcal{N} using the envelope technique presented in [Busic et al., 2012b] for the forward part, and the product form network \mathcal{J}^∞ , with stationary measure π^∞ , for the backward part. This perfect sampler is given in Algorithm 3, and the events for both the reversed network and the forward one are coupled using the following construction, similar to the coupling described in Proposition 2.

The reversed chain of network \mathcal{J}^∞ has the same routing events $r_{i,j}^+$, but different probabilities, whose corresponding intervals are denoted $B_{i,j}$, as in Section 2.2.

Let $Y^\infty(0)$ by a stationary sample of the infinite reversed system (with no negative customers), generated according to the product form distribution π^∞ . Given an infinite sequence of i.i.d. uniformly distributed random variables in $[0, 1]$, U_{-1}, U_{-2}, \dots , we construct an infinite sequence of events w_{-1}, w_{-2}, \dots of events for the reversed chain as follows. The intervals B_{ij} are ordered to cover $[0, 1]$ using the alpha-numerical order:

$$B_{0,1}, B_{0,2}, \dots, B_{0,M}, B_{1,0}, B_{1,1}, \dots, \dots, B_{M,M}.$$

Now, if $U_{-t} \in B_{i,j}$, then $w_{-t} = r_{i,j}^+$.

Then, we construct a trajectory of the reversed chain starting from $Y^\infty(0)$ and using the sequence of events: for all $t > 0$, $Y^\infty(-t) = \phi^\infty(Y^\infty(-t+1), w_{-t})$.

Now, the coupled event sequence for the direct chain $v_{-1}, \dots, v_{-t}, \dots$ is built as described by the following coupling procedure (C_2):

1. If $Y^\infty(-t) \neq Y^\infty(-t+1)$ then

- if $w_{-t} = r_{i,0}^+$ then $v_{-t} = r_{i,0}^+$ with probability $\frac{p_{i,0}^+}{p_{i,0}^+ + \sum_{j=0}^M p_{i,j}^-}$ and $v_{-t} = r_{i,j}^-$ with probability $\frac{p_{i,j}^-}{p_{i,0}^+ + \sum_{j=0}^M p_{i,j}^-}$.
- if $w_{-t} = r_{i,j}^+$ ($i \neq 0$) then $v_{-t} = r_{j,i}^+$

2. If $Y^\infty(-t) = Y^\infty(-t+1)$ then event w_{-t} must be of the form $w_{-t} = r_{i,j}^+$ with $i \neq 0$.

- if $w_{-t} = r_{i,0}^+$ then $v_{-t} = r_{i,0}^+$ with probability $\frac{p_{i,0}^+}{p_{i,0}^+ + \sum_{j=0}^M p_{i,j}^-}$ and $v_{-t} = r_{i,j}^-$ with probability $\frac{p_{i,j}^-}{p_{i,0}^+ + \sum_{j=0}^M p_{i,j}^-}$.
- if $w_{-t} = r_{i,j}^+$ ($j \neq 0$) then $v_{-t} = r_{i,j}^+$

Using this coupling, the following properties hold

Property 3

$$\forall t > t' > 0, \quad \psi^\infty(Y^\infty(-t), v_{-t}, \dots, v_{-t'+1}) \preceq Y^\infty(-t'). \quad (13)$$

In other words, the trajectory of the direct chain for network \mathcal{N} from time $-t$ to time 0, starting in $Y^\infty(-t)$ remains below the trajectory of the reversed chain from time 0 to time $-t$, starting in $Y^\infty(0)$.

Property 4 For all $\mathbf{x}, \mathbf{t}, \mathbf{s}$, the trajectories $\psi(\mathbf{x}, \mathbf{v}_{-\mathbf{t}}, \dots, \mathbf{v}_{-\mathbf{s}+1})$ are “correct” in the sense that if u_t are i.i.d. events with $\mathbb{P}(u_t = r_{ij}^{+/-}) = a_{ij}^{+/-}$, then $\psi(\mathbf{x}, \mathbf{v}_{-\mathbf{t}}, \dots, \mathbf{v}_{-\mathbf{t}'+1})$ and $\psi(\mathbf{x}, \mathbf{u}_1, \dots, \mathbf{u}_s)$ have the same distribution.

The proofs of the two properties are similar to the case without negative customers.

Algorithm 3: Envelope PSA with a Bounding Process (EPSA-BP)

Data:

- (1) An infinite i.i.d. sequence $\{U_{-n}\}_{n \in \mathbb{N}}$ with U_{-k} uniformly distributed in $[0, 1]$ (together with the sequence of backward events $\{w_{-n}\}_{n \in \mathbb{N}}$).
- (2) A state Y^∞ generated according to the distribution π^∞ .

Result: A state sampled from the stationary distribution of \mathcal{N}

begin

```

  t := 1;
  repeat
    for k = ⌊t/2⌋ to t - 1 do
      Y∞ := φ∞(Y∞, w-k);
      Compute the event v-k using the coupling procedure (C2);
      Z(1) := Y∞ ∧ C; Z(2) := 0;
      for i = t - 1 downto 0 do
        Z(1) := supx ∈ [Z(1), Z(2)] ψ(x, v-i);
        Z(2) := infx ∈ [Z(1), Z(2)] ψ(x, v-i);
      t := 2t;
    until Z(1) = Z(2);
  return Z(1);

```

Theorem 4 Algorithm EPSA-BP terminates with probability 1, and the output of EPSA-BP is a state with the stationary distribution of network \mathcal{N} .

The following theorem can be proved in a similar way as Theorem 1. The only difference of the proof is in the forward part that uses the envelope technique in [Busic et al., 2012b] and the fact that the bounding chain also bounds the upper envelope of the original chain: For any t , any sequence $(u_{-t}, \dots, u_{-1}) \in [0, 1]^t$, and any $\mathbf{x} \in \mathcal{S}$, $\mathbf{y} \in \mathcal{S}^\infty$, the backward bounding trajectory remains above the upper forward envelope.

$$\begin{aligned}
 \mathbf{x} \preceq \mathbf{y} &\Rightarrow \psi(\mathbf{x}, v_{-t}, \dots, v_{-1}) \preceq \phi^\infty(\mathbf{y}, w_{-1}, \dots, w_{-t}) \\
 &\Rightarrow \sup_{\{\mathbf{z} | \mathbf{z} \preceq \mathbf{x}\}} (\psi(\mathbf{z}, v_{-t}, \dots, v_{-1})) \preceq \phi^\infty(\mathbf{y}, w_{-1}, \dots, w_{-t}).
 \end{aligned}$$

As for termination, in general, the envelopes do not necessarily meet ($Z^{(1)} = Z^{(2)}$) so the algorithm may not terminate with probability 1. In our case however, envelopes will meet with probability one for network \mathcal{N} as soon as it is stable. This follows easily from the stability assumption (12). □

We illustrated the extension to non-monotone networks on a concrete example of negative customers. In general, this extension will be valid for any Markov chain $Z(t+1) = \psi(Z(t), R_{t+1})$ under the following assumptions:

- It is possible to find a *bounding process* $Y(t+1) = \zeta(Y(t), R_{t+1})$ such that for any t , any sequence $(r_{-t+1}, \dots, r_0) \in [0, 1]^t$, and any $\mathbf{x} \in \mathcal{S}$, $\mathbf{y} \in \mathcal{S}'$,

$$\mathbf{x} \preceq \mathbf{y} \Rightarrow \psi(\mathbf{x}, r_{-t+1}, \dots, r_0) \preceq \zeta(\mathbf{y}, r_{-t+1}, \dots, r_0),$$

where \mathcal{S} and \mathcal{S}' denote the state spaces of the original and bounding process respectively.

- It is easy to compute the backward transition function (β^∞ in the algorithm) for the bounding process and the initial sample distributed according to its stationary distribution. This second condition can be relaxed: we only need to be able to sample from an upper bound of this distribution.
- The *envelopes* for the original process couple and are easy to compute.

For negative customer events $r_{i,j}^-$, $1 \leq i, j \leq M$, we have:

$$\sup_{\mathbf{x} \in [\mathbf{y}, \mathbf{z}]} \psi(\mathbf{x}, r_{i,j}^-) = \begin{cases} \mathbf{z} - \mathbf{e}_i - \mathbf{e}_j, & y_i > 0, z_j > 0, \\ \mathbf{z} - \mathbf{e}_i, & y_i = 0, z_i > 0, \\ \mathbf{z}, & z_i = 0, \end{cases}$$

and

$$\inf_{\mathbf{x} \in [\mathbf{y}, \mathbf{z}]} \psi(\mathbf{x}, r_{i,j}^-) = \begin{cases} \mathbf{y} - \mathbf{e}_i - \mathbf{e}_j, & y_i > 0, y_j > 0, \\ \mathbf{y} - \mathbf{e}_i, & y_i > 0, y_j = 0, \\ \mathbf{y} - \mathbf{e}_j, & y_i = 0, z_i > 0, y_j > 0, \\ \mathbf{y}, & z_i = 0. \end{cases}$$

Again, the coupling time is bounded by Equation (6) and does not depend on the capacities of the queues, unlike in classical perfect samplers.

8 Conclusion

In this paper we have presented a new perfect sampling algorithm for Jackson queueing networks with finite and infinite capacities in queues. A complexity analysis of the algorithm shows that its expected sampling time does not depend on the capacities. This is a remarkable improvement over classical perfect samplers whose sampling time increases at least linearly in the capacities, or fails when some capacities are infinite. Actually our approach is quite general and should be usable for any Markov chain for which a bounding process with a computable stationary distribution can be constructed and coupled with the original chain. In particular we have shown how it applies in the case of queueing networks with negative clients.

Acknowledgements

This work was partially supported by the ANR Numerical Models Program (MARMOTE project).

References

- Anselmi, J. and Gaujal, B. (2011). On the efficiency of perfect simulation in monotone queueing networks. In *IFIP Performance: 29th International Symposium on Computer Performance, Modeling, Measurements and Evaluation*, Amsterdam. ACM Performance Evaluation Review.

- Bolch, G., Greiner, S., de Meer, H., and Trivedi, K. (2005). *Queueing Networks and Markov Chains*. Wiley-Interscience.
- Busic, A., Gaujal, B., and Perronnin, F. (2012a). Perfect Sampling of Networks with Finite and Infinite Capacity Queues. In Al-Begain, K., Fiems, D., and Vincent, J.-M., editors, *19th International Conference on Analytical and Stochastic Modelling Techniques and Applications (ASMTA) 2012*, volume 7314 of *Lecture Notes in Computer Science*, pages 136–149, Grenoble, France. Springer.
- Busic, A., Gaujal, B., and Pin, F. (2012b). Perfect sampling of Markov chains with piecewise homogeneous events. *Performance Evaluation*, 69(6):247 – 266.
- Busic, A., Gaujal, B., and Vincent, J.-M. (2008). Perfect simulation and non-monotone Markovian systems. In *3rd International Conference Valuetools'08*, Athens, Greece. ICST.
- Dopper, J., Gaujal, B., and Vincent, J.-M. (2006). Bounds for the coupling time in queueing networks perfect simulation. In *Celebration of the 100th anniversary of Markov*, pages 117–136.
- Gaujal, B., Gorgo, G., and Vincent, J.-M. (2011). Perfect sampling of phase-type servers using bounding envelopes. In Springer-Verlag, editor, *18th International Conference on Analytical and Stochastic Modelling Techniques and Applications (ASMTA'11)*, LNCS, Venise.
- Gaujal, B., Perronnin, F., and Bertin, R. (2008). Perfect simulation of a class of stochastic hybrid systems with an application to peer to peer systems. *Journal of Discrete Event Dynamic Systems*, 18(2):211–240. Special Issue on Hybrid Systems.
- Gelenbe, E. (1991). Product-form queueing networks with negative and positive customers. *Journal of Applied Probability*, 28(3):pp. 656–663.
- Jackson, J. R. (1963). Job shop-like queueing systems. *Management Sci.*, 10,131.
- Kelly, F. (1979). *Reversibility and Stochastic Networks*. John Wiley & Sons Ltd.
- Kendall, W. S. (2005). Notes on perfect simulation. Dept. of statistics, University of Warwick.
- Pin, F., Busic, A., and Gaujal, B. (2011). Acceleration of perfect sampling for skipping events. In *Valuetools*, Paris.
- Propp, J. G. and Wilson, D. B. (1996). Exact sampling with coupled Markov chains and applications to statistical mechanics. *Rand. Struct. Alg.*, 9(1-2):223–252.
- Vincent, J.-M. (2005). Perfect simulation of monotone systems for rare event probability estimation. In *WSC '05: Proceedings of the 37th conference on Winter simulation*, pages 528–537. Winter Simulation Conference.
- Walker, A. J. (1977). An efficient method for generating discrete random variables with general distributions. *ACM Trans. Math. Softw.*, 3:253–256.

Contents

| | | |
|----------|---|-----------|
| 1 | Introduction | 3 |
| 2 | Queueing network model | 4 |
| 2.1 | Discrete-event definition of JQN | 6 |
| 2.2 | Network with infinite buffers | 7 |
| 3 | Perfect Sampling Algorithm with a Bounding Process | 10 |
| 4 | Complexity analysis | 13 |
| 4.1 | The acyclic case | 14 |
| 4.2 | The hyperstable case | 16 |
| 5 | Comparison with the classical perfect sampler | 17 |
| 6 | Numerical experiments | 19 |
| 6.1 | Sampling time improvement over classic PSA | 19 |
| 6.2 | Complexity bound | 21 |
| 7 | Extension to non-monotone networks | 22 |
| 8 | Conclusion | 25 |



**RESEARCH CENTRE
GRENOBLE – RHÔNE-ALPES**

Inovallée
655 avenue de l'Europe Montbonnot
38334 Saint Ismier Cedex

Publisher
Inria
Domaine de Voluceau - Rocquencourt
BP 105 - 78153 Le Chesnay Cedex
inria.fr

ISSN 0249-6399