

Perfect simulation of Markovian load balancing queueing networks in equilibrium

Carl Graham

► To cite this version:

Carl Graham. Perfect simulation of Markovian load balancing queueing networks in equilibrium. 2024. hal-04633838

HAL Id: hal-04633838 https://polytechnique.hal.science/hal-04633838v1

Preprint submitted on 3 Jul 2024

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 simulation of Markovian load balancing queueing networks in equilibrium

Carl Graham*

CMAP, CNRS, École polytechnique, Institut Polytechnique de Paris, 91120 Palaiseau, France, and Inria.

Abstract

We define a wide class of Markovian load balancing networks of identical singleserver infinite-buffer queues. These networks may implement classic parallel server or work stealing load balancing policies, and may be asymmetric, for instance due to topological constraints. The invariant laws are usually not known even up to normalizing constant. We provide three perfect simulation algorithms enabling Monte Carlo estimation of quantities of interest in equilibrium. The state space is infinite, and the algorithms use a dominating process provided by the network with uniform routing, in a coupling preserving a preorder which is related to the increasing convex order. It constitutes an order up to permutation of the coordinates, strictly weaker than the product order. The use of a preorder is novel in this context. The first algorithm is in direct time and uses Palm theory and acceptance-rejection. Its duration is finite, a.s., but has infinite expectation. The two other algorithms use dominated coupling from the past; one achieves coalescence by simulating the dominating process into the past until it reaches the empty state, the other, valid for exchangeable policies, is a back-off sandwiching method. Their durations have some exponential moments.

Keywords: load balancing, network topology, parallel queues, work stealing, stochastic order, perfect simulation, dominated coupling from the past, Monte Carlo estimation, performance evaluation

MSC 2020: Primary 60K30; Secondary 60E15 60G10 68M20

*e-mail: carl.graham@polytechnique.edu — July 3, 2024

1 Introduction

This paper provides perfect simulation algorithms for the invariant laws (or stationary distributions) of Markovian queueing networks managed by load balancing policies (or protocols, algorithms, etc.). It devises a general abstract definition of these networks, whereas most previous works consider specific examples only.

These invariant laws are usually not known even up to normalizing constants, which precludes the use of Gibbs samplers. The algorithms allow to draw samples for practical purposes such as the Monte Carlo estimation of quantities of interest in equilibrium, *e.g.*, quality of service (QoS) indicators for performance evaluation. They are "perfect" up to real-life constraints on computer precision, running times, pseudo-random number generators, etc.

Each network consists of $c \geq 2$ identical single-server infinite-buffer queues, and its state is given by the queue length vector. Tasks arrive according to a Poisson process of intensity $\lambda > 0$. Each task has an exponential service duration of parameter 1 (w.l.o.g., up to time-scale) after which it exits the system. Queueing discipline is work conserving, thus the total service rate is equal to the number of non-empty queues (busy servers). The initial state, arrival stream, and services are independent. The stability condition $\lambda < c$ will always be assumed to hold.

Load balancing policies observe the network state and strive to allocate tasks to server queues so as to minimize server idleness. They may have to satisfy varied constraints, such as partial observation, distributed functioning, and limited overhead. Parallel queue networks allocate each task at arrival to a server queue and leave it there until served [1, 2], but other kinds of networks may implement work stealing or jockeying at instants of service completion [3-8].

We shall determine a wide class of load balancing networks allowing quite general policies. These policies are possibly asymmetric in terms of the queues, for instance they may respect a directed graph network topology, and may thus yield non-exchangeable queue length vector processes. Most previous results are proved for specific exchangeable networks in suitable asymptotic regimes. Many interesting networks are asymmetric, non-asymptotic results are often needed, and simulation studies such as those in [9] are then quite useful. See [2, Sect. 6], [10], *e.g.*

We shall provide three perfect simulation algorithms. The first one is a directtime acceptance-rejection method based on Palm theory. Its run time is finite but has infinite mean, and it cannot be implemented in practice. The other two use backward couplings and have run times with some exponential moments. Backward couplings have long been used for invariant laws [11]. Coupling From The Past (CFTP) was devised by Propp and Wilson [12] for perfect simulation of the invariant law of a Markov chain on a finite state space with a largest and a least element, and has been extended for infinite state spaces into Dominated CFTP (DomCFTP) using suitable dominating processes simulated backward and forward in time, see [13–18].

We were inspired by the use in Connor and Kendall [13] of DomCFTP for a M/G/c [FCFS] queue, and we refer to this paper for interesting background on the topic. The service distribution has a second moment, thus the emptying time from equilibrium has a first moment. The arrival rate is λ , the mean service time is $1/\mu$, the offered traffic is λ/μ , and the queue is stable if and only if $\lambda/\mu < c$. The Kiefer-Wolfowitz workload

vectors record the evolution of the M/G/c [FCFS] queue by a Markov chain. The dominating process is given by a M/G/c queue with uniform routing, called random assignment and denoted by M/G/c [RA], which can be viewed as a system of c i.i.d. M/G/1 queues with arrival rates λ/c denoted by $[M/G/1]^c$. The invariant law of the M/G/1 queue can be perfectly simulated using the Pollaczek-Khintchine formula. The workload of the M/G/1 queue does not depend on the queueing discipline, and the dynamic reversibility of the M/G/1 [PS] queue is used to simulate the dominating process backward in time. Much care is needed to ensure that the coupling between M/G/c [FCFS] and $[M/G/1]^c$ is preserved when simulating forward in time.

The present paper is organized as follows. Section 2 provides examples of load balancing networks and then provides a general definition. Section 3 constructs a coupling in which the uniform routing (UR) network has a task backlog not less than an arbitrary load balancing network, in a preorder related to the increasing convex order which constitutes an order up to permutation of the coordinates strictly weaker than the product (coordinate-wise) order. To the best of our knowledge, the use of such a preorder is novel for perfect simulation. Section 4 explains how to manage efficiently in code the coupled processes and defines an embedded Markov chain which has same invariant law and is simpler to simulate. Section 5 exploits the domination by the UR network for an acceptance-rejection method based on Palm theory, inspired by [19]. Section 6 provides the two DomCFTP methods inspired by [13]. The dominating UR network is reversible in equilibrium and thus simulatable backward in time, and the preservation of the coupling through the backward and forward simulations is obtained by considering the law of an augmented Markov chain. The first method achieves coalescence by simulating the dominating process into the past until it reaches the empty state. The second is valid only for exchangeable networks and uses a back-off sandwiching method which is likely to be much quicker.

2 Load balancing networks

2.1 Examples of parallel server load balancing policies

Parallel server policies allocate each task at arrival to a server queue and leave it there until it ends its service and departs the network.

Uniform Routing (UR). At each arrival, the policy chooses a queue uniformly at random and allocates the task to it. This blind policy is perhaps the only one for which the invariant law is known: the *c* queues receive independent Poisson arrivals of intensities λ/c and constitute a system of i.i.d. M/M/1 queues. Under the ergodicity condition $\lambda < c$, the invariant law is $\gamma^{\otimes c}$ where γ is the geometric law given by $\gamma(n) \triangleq (1-(\lambda/c))(\lambda/c)^n$ for $n \in \mathbb{N}_0$. This allows performance evaluation in equilibrium, *e.g.*, the probability that a given queue contains at least *n* tasks and the limit proportion of such queues when *c* and λ go to infinity with λ/c fixed is given by $\gamma([n, \infty)) = (\lambda/c)^n$.

Join the Shortest Queue (JSQ). At each arrival, the task is routed to one of the queues of minimal length in the network. Here and below, ties may be broken arbitrarily, *e.g.*, uniformly at random. JSQ requires high overhead in form of intensive signaling and bookkeeping, and is practicable only in certain applications such as call centers or parallel computing of moderate sizes. JSQ is optimal among parallel queue

policies using only the past queue length information, see [20] and the coupling proof in [21] inspired by previous work of S. Foss. The invariant law is unknown, even though complex analysis provides useful transforms of these laws for c = 2, see [22, 23].

Power of d Choices (JSQ(d)). At each arrival, $d \ge 2$ queues are chosen uniformly at random, with or without replacement according to the variant chosen, and the task is routed to one of these of minimal length. The degenerate case d = 1 corresponds to UR and the case d = c without replacement to JSQ. JSQ(d) can be seen as a low overhead proxy for JSQ, and has long been intensively studied.

Idle Queue First (IQF). At each arrival, the task is routed to an empty queue (*i.e.*, idle server) in the network if any are present, else to a uniformly chosen queue. This can be seen as a low overhead proxy for JSQ in which queues only need to signal their changes of status between empty and non-empty. This policy was investigated in [24] under the name Join Idle Queue.

Idle One First (IOF). At each arrival, the task is routed to an empty queue in the network if any, else to a queue of length one if any, else to a uniformly chosen queue. This can be seen as a low overhead proxy for JSQ in which queues only need to signal their changes of status between empty, containing one task, and other. This policy was introduced and investigated in [1] under the acronym I1F, see below.

2.2 Asymptotic performance evaluation

Asymptotic performance evaluation yields guidelines and understanding. It is usually applied to an exchangeable network in which a reduced low-dimensional, ideally onedimensional, representation exists yielding a tractable limit process.

The M/M/c queue has an optimal server utilization. It is called a "centralized system" in [1, 2] and cannot be matched by parallel server queues, but can be by a network implementing IQF at arrivals and work stealing in which any server finishing service of the last task in its queue simultaneously steals unattended work from another queue if any. It is interpreted as a "Join the Shortest Workload" policy by the Kiefer-Wolfowitz construction [13, Sect. 2]. Its invariant law is explicit up to a normalizing constant which is known for all practical purposes from the well-tabulated Erlang C formula. This allows for precise performance evaluation and provides a benchmark.

JSQ(d) has been investigated in the mean-field regime in which c and λ go to infinity with λ/c and $d \geq 2$ kept fixed. The equilibrium proportion of queues containing at least n tasks converges to $((\lambda/c)^{d^n} - 1)/(d - 1)$ with hyper-exponential decay in n, much faster than the exponential decay $(\lambda/c)^n$ for UR. See [25–28], e.g. In this regime, the asymptotic performance of IQF is perfect since the limit proportion of empty queues is non-zero and thus every task starts service at arrival.

Gupta and Walton [1] investigate parallel server policies in a much more heavily loaded regime, the nondegenerate slowdown regime [29] in which λ and c go to infinity with $c - \lambda$ going to $\alpha > 0$. They prove that the performances of JSQ and IOF are comparable and are within 15% of the performance of M/M/c, while the performance of IQF is within 100% of the performance of M/M/c. In this regime, IOF is asymptotically optimal and better than IQF, and it is pointless to implement more subtle policies, for instance with higher thresholds.

The wide survey [2] studies parallel server networks in the fluid regime as well as in the Halfin-Whitt regime in which λ and c go to infinity with $\frac{1}{\sqrt{c}}(c-\lambda)$ going to $\beta > 0$. In these regimes, IQF is already asymptotically optimal.

2.3 General definitions for load balancing networks

A filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \in \mathbb{R}_+}, \mathbb{P})$ satisfying the usual assumptions is given. It supports all random elements, all random processes are adapted, and all Poisson process are $(\mathcal{F}_t)_{t \in \mathbb{R}_+}$ -Poisson processes. Random elements are assumed to be independent from each other if not stated otherwise. Let $\mathbb{N}_0 \triangleq \{0, 1, \ldots\}$ and $(e_i)_{1 \leq i \leq c}$ denote the canonical basis of \mathbb{R}^c and \mathfrak{S}_c the permutation group of $\{1, \ldots, c\}$.

We devise a workable general definition of load balancing networks, instead of focusing on specific examples. It allows any classic load balancing policy.

Definition 1 (Load balancing network). Tasks arrive according to a $(\mathcal{F}_t)_{t \in \mathbb{R}_+}$ -Poisson process of intensity $\lambda > 0$ on $c \geq 2$ single-server infinite-buffer queues. Service duration is exponential with parameter 1. The state of the network is described by the queue length (vector) process with sample paths in $\mathbb{D}(\mathbb{R}_+, \mathbb{N}_0^c)$ given by

$$X \triangleq (X(t))_{t \in \mathbb{R}_+}, \quad X(t) \triangleq (X^i(t))_{1 < i < c},$$

where $X^{i}(t)$ is the length of (number of tasks in) queue i at time t. The load balancing networks under consideration respect the following rules.

- 1. The process $(X(t))_{t \in \mathbb{R}_+}$ is $(\mathcal{F}_t)_{t \in \mathbb{R}_+}$ -Markov.
- 2. At each arrival, the task is routed to a queue not longer than a uniformly chosen queue. Specifically, if t is a task arrival instant then the policy chooses j uniformly in $\{1, \ldots, c\}$ and determines a queue of index i such that

$$X^{i}(t-) \leq X^{j}(t-), \qquad X(t) = X(t-) + e_{i}.$$

- 3. The queueing discipline in each queue is work conserving. Thus, the total service intensity is equal to the number of non-empty queues and the subsequent service completion concerns a uniformly chosen queue among these.
- 4. At each instant of service completion, one task may be reallocated from a queue which just before had a length not less than the queue completing service to the latter. Specifically, if t is a service completion instant and queue j is completing service then the policy determines a queue of index i such that

$$X^{i}(t-) \ge X^{j}(t-) \ge 1$$
, $X(t) = X(t-) - e_{i}$.

The policy and the network are said to be exchangeable if the decision taking is invariant in law under permutations of the queue indices. In this case, if X(0) is exchangeable then X(t) is exchangeable, and it is so in equilibrium.

The uniform routing (UR) network, for which i = j in Rules 2 and 4, is an exchangeable load balancing network with queue length process denoted by

$$Y \triangleq (Y(t))_{t \in \mathbb{R}_+}, \quad Y(t) \triangleq (Y^i(t))_{1 \le i \le c}.$$

		•	
c			,

Poisson process splitting implies that $Y \triangleq (Y(t))_{t \in \mathbb{R}_+}$ constitutes a system of c i.i.d. M/M/1 queue processes with arrival rates λ/c and service rates 1. It is positive recurrent if and only if $\lambda < c$ and then it is reversible in equilibrium and has simulatable invariant law $\gamma^{\otimes c}$, the *c*-fold product of the geometric law

$$\gamma \triangleq (\gamma(n))_{n \in \mathbb{N}_0}, \quad \gamma(n) \triangleq (1 - (\lambda/c))(\lambda/c)^n.$$
 (1)

Rule 2 is satisfied by all policies in Section 2.1. Rule 4 is trivially satisfied by any parallel server policy, but allows work stealing or jockeying in which the queue of a server completing a service may be immediately joined by a task from a longer queue, such instantaneous jumps not being recorded explicitly by the queue length process. The implementation of M/M/c described in Section 2.2 satisfies Definition 1.

Remark 1. The policies in Section 2.1 are exchangeable if ties are broken uniformly at random within queues. Definition 1 allows policies which break symmetry between queues and are not exchangeable. For instance, the network may have a directed graph topology and the policies in Section 2.1 may be implemented on the neighborhood of the uniformly chosen queue. See [9], [2, Sect. 6], [10], e.g.

Most if not all asymptotic results are limited to policies which are exchangeable, or at least have an infinitely growing number of possible choices becoming exchangeable in the limit as in [2, Sect. 6] and [10]. Non-asymptotic performance evaluation may also be required. Simulation studies such as [9] usually rely on waiting long lengths of time and hoping that equilibrium has been reached without theoretical guidelines.

The present paper will provide theoretical results allowing perfect simulation for general policies. The special case of exchangeable networks will allow more efficient codes and will be investigated specifically.

3 Preorder and coupling

3.1 Preorder and order up to permutation of coordinates

Let $x \triangleq (x^i)_{1 \le i \le c}$ in \mathbb{N}_0^c be a network state and n be in \mathbb{N}_0 . Let

$$\alpha_n(x) \triangleq \sum_{i=1}^c \mathbb{1}_{\{x^i > n\}}, \qquad \beta_n(x) \triangleq \sum_{i=1}^c (x^i - n)^+ = \sum_{k \ge n} \alpha_k(x), \qquad (2)$$

denote respectively the number of queues containing more than n tasks and the number of these excess tasks. The β_n are natural quality of service (QoS) indicators, in particular β_0 counts the total number of customers and for the FCFS discipline β_n counts the number of tasks which have at least n tasks in front of them.

Theorem 1. A preorder \preccurlyeq between network states x and y in \mathbb{N}_0^c is defined by

$$x \preccurlyeq y \iff \beta_n(x) \le \beta_n(y), \ \forall n \in \mathbb{N}_0.$$

The indifference relation \sim defined by $x \sim y$ if and only if $x \preccurlyeq y$ and $y \preccurlyeq x$ satisfies

$$x \sim y \iff \exists \sigma \in \mathfrak{S}_c : (x^{\sigma(1)}, \dots, x^{\sigma(c)}) = (y^1, \dots, y^c)$$

and thus \preccurlyeq defines a partial order up to permutation of the coordinates. Moreover

$$x \preccurlyeq y \iff \sum_{i=1}^{c} f(x^{i}) \le \sum_{i=1}^{c} f(y^{i}) \text{ for any non-decreasing convex } f: \mathbb{R}_{+} \to \mathbb{R}.$$

Proof. Clearly \preccurlyeq is reflexive and transitive and is thus a preorder, and $x \sim y$ if and only if $\beta_n(x) = \beta_n(y)$ for all n in \mathbb{N}_0 . In particular

$$\max\{n \in \mathbb{N}_0 : \beta_n(x) > 0\} = \max\{n \in \mathbb{N}_0 : \beta_n(y) > 0\} \triangleq m, \quad \beta_m(x) = \beta_m(y),$$

which imply that x and y have same number of coordinates with the same maximal size m + 1. Inductively by considering $\beta_{m-1}, \ldots, \beta_0$ we deduce that x and y have same number of coordinates of every size greater than or equal to 1, and hence are equal up to a permutation of coordinates.

For the last equivalence statement, sufficiency is a special case. To prove necessity, we may and do assume that f(0) = 0. Then f is continuous on \mathbb{R} and has a right-hand derivative f'_+ which is non-negative, non-decreasing, right-continuous, and such that $f(x) = \int_0^x f'_+(y) \, dy$, see [30, Thms 24.1–2]. Integration by parts yields that

$$f(x) = xf'_{+}(x) - \int_{(0,x]} y \, \mathrm{d}f'_{+}(y) = xf'_{+}(0) + \int_{(0,\infty)} (x-y)^{+} \, \mathrm{d}f'_{+}(y) \, .$$

Thus f is in the positive cone generated by the $x \in \mathbb{R}_+ \mapsto (x-y)^+$ for $y \ge 0$.

This preorder is related to the univariate increasing convex order [31], [32, Chap. 4]. Its quotient order by \sim is strictly weaker than the product order up to permutations of the coordinates, for instance $(1,1) \leq (0,2)$.

3.2 Coupling a load balancing network with the UR network

We shall construct a coupling of the queue length processes $X \triangleq (X^1, \ldots, X^c)$ of a generic load balancing network and $Y \triangleq (Y^1, \ldots, Y^c)$ of a UR network (Definition 1) which preserves the preorder \preccurlyeq . A similar coupling is used in [9, Thm 4] and [27, Sect. 4.1] to compare networks implementing JSQ(d) for different $d \ge 1$.

The coupling will require to rearrange the coordinates of X and of Y in nondecreasing order. Ties between queues of equal length must be broken using a \mathfrak{S}_{c} valued process $(\sigma(t))_{t \in \mathbb{R}_+}$ with nice properties satisfying

$$X^{\sigma(t)(1)}(t) \le \dots \le X^{\sigma(t)(c)}(t), \tag{3}$$

and similarly for Y, with permutation process denoted by $(\tau(t))_{t \in \mathbb{R}_+}$.

For the sake of simplicity, our exposition will assume that ties are broken by ascending queue index, so that $\sigma(t)$ is a deterministic function of X(t). Section 4.3 will address some practical simulation issues. It is not efficient to sort X(t) at each update, and $\sigma(t)$ should also be stored and updated. The updates of $\sigma(t)$ when ties are broken

by ascending queue index are inefficient, and we describe there another tie-breaking rule with efficient updates to which our exposition can readily be extended.

The coupling starts by drawing X(0) and Y(0) from a suitable joint law.

The arrivals in both networks follow the same $(\mathcal{F}_t)_{t \in \mathbb{R}_+}$ -Poisson process of intensity λ marked by uniform draws in $\{1, \ldots, c\}$. Let t and k be an arrival instant and its mark. The generic policy takes as its uniformly chosen queue the k-th in the non-decreasing ordering, of index $\sigma(t-)(k)$, and determines a queue i such that

$$X^{i}(t-) \leq X^{\sigma(t-)(k)}(t-), \quad X(t) = X(t-) + e_{i}.$$

The UR network does likewise with $i = \tau(t-)(k)$.

The potential service completions in both networks follow the same $(\mathcal{F}_t)_{t \in \mathbb{R}_+}$ -Poisson process of intensity c marked by uniform draws in $\{1, \ldots, c\}$. Let t and k be a potential service completion instant and its mark. The generic network has a service completion in the k-th queue in the non-decreasing ordering, of index $\sigma(t-)(k)$, if and only if it is non-empty. In this case its policy determines a queue i such that

$$X^{i}(t-) \ge X^{\sigma(t-)(k)}(t-) \ge 1$$
, $X(t) = X(t-) - e_{i}$.

The UR network does likewise with $i = \tau(t-)(k)$.

Theorem 2. This construction yields a coupling between the queue length process $X \triangleq (X^1, \ldots, X^c)$ of a generic load balancing network and $Y \triangleq (Y^1, \ldots, Y^c)$ of the UR network. Moreover,

$$X(0) \preccurlyeq Y(0) \implies X(t) \preccurlyeq Y(t), \ \forall t \in \mathbb{R}_+$$

Proof. Classic arguments about Poisson process splitting and about uniform draws and conditioning prove that this is indeed a coupling.

Assume that $X(0) \preccurlyeq Y(0)$. Let t be either an arrival or a potential service completion instant. In a proof by contradiction, assume that

$$\forall n \in \mathbb{N}_0, \ \beta_n(X(t-)) \le \beta_n(Y(t-)), \quad \exists p \in \mathbb{N}_0, \ \beta_p(X(t)) > \beta_p(Y(t)).$$
(4)

This implies that

$$\beta_p(X(t-)) = \beta_p(Y(t-)). \tag{5}$$

First case: t is an arrival instant with mark k. The generic network allocates the task to queue i and the UR network to queue $\tau(t-)(k)$ in a way that

$$l \triangleq X^{i}(t-) \le X^{\sigma(t-)(k)}(t-), \qquad m \triangleq Y^{\tau(t-)(k)}(t-).$$
 (6)

Since $c-\alpha_n$ is the number of coordinates of size at most n, the non-decreasing orderings imply that $c-\alpha_{l-1}(X(t-)) < k$ and $k \leq c-\alpha_m(Y(t-))$ and thus that

$$\alpha_m(Y(t-)) < \alpha_{l-1}(X(t-)).$$

Moreover (2), (4), (5), and (6) imply that

$$m+1 \le p \le l$$
.

Then $p \ge 1$, (2) yields that

$$\begin{aligned} \beta_{p-1}(X(t-)) &= \alpha_{p-1}(X(t-)) + \beta_p(X(t-)) \,, \\ \beta_{p-1}(Y(t-)) &= \alpha_{p-1}(Y(t-)) + \beta_p(Y(t-)) \,, \end{aligned}$$

and (4) and (5) imply that

$$\alpha_{p-1}(X(t-)) \le \alpha_{p-1}(Y(t-)).$$

We conclude that

$$\alpha_{p-1}(X(t-)) \le \alpha_{p-1}(Y(t-)) \le \alpha_m(Y(t-)) < \alpha_{l-1}(X(t-)) \le \alpha_{p-1}(X(t-))$$

which is a contradiction. Therefore, (4) is not true.

Second case: t is a potential service completion instant with mark k. The generic network undergoes an effective service completion in the queue of index $\sigma(t-)(k)$ if and only if $X^{\sigma(t-)(k)}(t-) \geq 1$. In this case let i be the index of the queue of which the length will be reduced and else let $i = \sigma(t-)(k)$. The UR network reduces by one the length of queue $\tau(t-)(k)$ if it is not zero. In all cases

$$l \triangleq X^{i}(t-) \ge X^{\sigma(t-)(k)}(t-), \qquad m \triangleq Y^{\tau(t-)(k)}(t-).$$
 (7)

Since $c-\alpha_n$ is the number of coordinates of size at most n, the non-decreasing orderings imply that $k \leq c - \alpha_l(X(t-))$ and $c - \alpha_{m-1}(Y(t-)) < k$ and thus that

$$\alpha_l(X(t-)) < \alpha_{m-1}(Y(t-)).$$

Moreover (2), (4), (5), and (6) imply that

$$l \le p \le m - 1.$$

Then (2) yields that

$$\beta_p(X(t-)) = \alpha_p(X(t-)) + \beta_{p+1}(X(t-)), \beta_p(Y(t-)) = \alpha_p(Y(t-)) + \beta_{p+1}(Y(t-)),$$

and (4) and (5) imply that

$$\alpha_p(Y(t-)) \le \alpha_p(X(t-))$$

r	
s	

We conclude that

$$\alpha_p(Y(t-)) \le \alpha_p(X(t-)) \le \alpha_l(X(t-)) < \alpha_{m-1}(Y(t-)) \le \alpha_p(Y(t-))$$

which is a contradiction. Therefore, again (4) is not true.

Applying the falsity of (4) to the successive arrival and potential service completion instants t concludes the proof.

This yields the following important stability result.

Theorem 3. Let X and Y be the coupled queue length processes of a generic load balancing network and of a UR network. The empty state 0 is positive recurrent for X if and only if $\lambda < c$, and then X has a unique invariant law π , else there is none. The state (0,0) is positive recurrent for (X,Y) if and only if $\lambda < c$ and then (X,Y) has a unique invariant law κ with marginals π and $\gamma^{\otimes c}$, see (1), and else there is none. If $\lambda < c$ then the two-sided stationary version $(X(t), Y(t))_{t \in \mathbb{R}}$ of (X,Y) satisfies

$$X(t) \preccurlyeq Y(t), \quad \forall t \in \mathbb{R}.$$

Proof. If $\lambda < c$ then the queue length process Y of the UR network is positive recurrent since it is irreducible and has $\gamma^{\otimes c}$ as an invariant law, which is then unique. Theorem 2 yields that if (X(0), Y(0)) = (0, 0) then $X(t) \preccurlyeq Y(t)$ and in particular X(t) = 0 if Y(t) = 0. Therefore (0, 0) is positive recurrent for (X, Y) and thus 0 is positive recurrent for X. Since (0, 0) is reachable from all states for (X, Y), the complement of its recurrence class is constituted of transient states bearing mass zero for any invariant law, and by restriction to this recurrence class (necessarily irreducible) there exists a unique invariant law for (X, Y). Likewise for X. By comparison with a M/M/c queue, X is at best null recurrent if $\lambda = c$ and transient if $\lambda > c$. The result on the two-sided stationary version $(X(t), Y(t))_{t \in \mathbb{R}}$ follows from Palm theory using the hitting times of (0, 0) as regeneration times splitting the process into cycles, see [33, Sect. 8.3]. Indeed, the i.i.d. cycles start with $0 = X(0) \preccurlyeq Y(0) = 0$ and thus satisfy $X(t) \preccurlyeq Y(t)$ during their lifetimes, hence the length-biased cycle straddling time 0 must also satisfy $X(t) \preccurlyeq Y(t)$ during its lifetime. □

A load balancing parallel server queueing network is irreducible: any state leads to the empty state, and in order to attain any given state from the empty state we may load all queues beyond the maximal length of the given state and then serve the excess tasks. A network implementing work stealing may well not be irreducible.

In the sequel we assume that the stability condition $\lambda < c$ holds true.

4 Considerations for effective simulation

4.1 Sampling from the invariant law of the UR network

The perfect simulation algorithms will require to draw repeatedly from the invariant law $\gamma^{\otimes c}$ of the queue length process Y of the uniform routing (UR) network, where γ is the geometric law (1) under the stability condition $\lambda < c$.

A simple option for each draw is to simulate c independent draws from γ . A bounded time inversion method using truncated exponential random variables is given in Devroye [34, Sect. X.2.2]. It is noted in an exercise that an accuracy problem in computing $\log(1-p)$ could arise if the success probability p is small, and an expansion method is proposed, but here $p = 1 - \lambda/c$ and we can use $\log(\lambda/c) = \log \lambda - \log c$.

Another option is to first draw the total number of tasks from the negative binomial law NB($c, 1 - \lambda/c$) and then draw from a multinomial to attribute tasks to the cqueues. Draws from NB($c, 1 - \lambda/c$) can be performed by inverting the c.d.f., which is a regularized incomplete beta function, or by other uniformly fast methods as in [34, Sect. X.1.2 Example 1.5, Sect. X.4.7]. Methods for drawing from a multinomial with mean durations which are O(c) uniformly in the total number of tasks exist, such as [34, Sect. XI.1.4 Example 1.5].

4.2 Exchangeable load balancing policies

This subsection is devoted to exchangeable load balancing policies, see Definition 1. A reduced representation of $X = (X(t))_{t \in \mathbb{R}_+}$ and of $Y = (Y(t))_{t \in \mathbb{R}_+}$ "up to instantaneous permutations of the coordinates" is then adequate to study the invariant law. The exchangeability assumption is equivalent to assuming that these reduced representations are themselves Markov. An added benefit is that the preorder \preccurlyeq of Theorem 1 induces an order in this context.

For $x \triangleq (x^i)_{1 \le i \le c}$ in \mathbb{N}_0^c let $\dot{x} \triangleq (\dot{x}^i)_{1 \le i \le c}$ in \mathbb{N}_0^c be the vector of its coordinates set in non-decreasing order and

$$h_n(x) \triangleq \sum_{i=1}^c \mathbb{1}_{\{x^i \le n\}} = c - \alpha_n(x)$$

be the number of its coordinates not exceeding n in \mathbb{N}_0 . Setting $h_{-1}(x) = 0$,

$$\dot{x}^{i} = n \iff h_{n-1}(x) + 1 \le i \le h_n(x), \qquad h_n(x) = h_n(\dot{x}).$$
(8)

Note that x can be sorted into \dot{x} in $O(c \log c)$ time, and that $x \sim y \Leftrightarrow \dot{x} = \dot{y}$, see Theorem 1. Let $\dot{X}(t)$ and $\dot{Y}(t)$ be the non-decreasing reorderings of X(t) and Y(t).

The processes $X \triangleq (X(t))_{t \in \mathbb{R}_+}$ and $Y \triangleq (Y(t))_{t \in \mathbb{R}_+}$ are reduced representations of X and Y which are $(\mathcal{F}_t)_{t \in \mathbb{R}_+}$ -Markov and evolve as follows.

1. Let t > 0 be an arrival instant with mark k. The task is allocated according to policy to a queue of length $n \leq X^k(t-)$, and the single coordinate update is

$$\hat{X}^{h_n(\dot{X}(t-))}(t) = \hat{X}^{h_n(\dot{X}(t-))}(t-) + 1 = n+1.$$

2. Let t > 0 be a potential service completion instant with mark k. It is effective if and only if $\dot{X}^k(t-) \ge 1$, and then a task is removed according to policy from a queue of length $n \ge \dot{X}^k(t-) \ge 1$, and the single coordinate update is

$$\acute{X}^{h_{n-1}(\acute{X}(t-))+1}(t) = \acute{X}^{h_{n-1}(\acute{X}(t-))+1}(t-) - 1 = n-1 \, .$$

The determination of the queue of length n by the policy is Markov due to exchangeability. For instance, JSQ(d) may choose $k_1 = k, k_2, \ldots, k_d$ uniformly and take $n = X^{\min\{k_1,\ldots,k_d\}}(t-)$ in 1. The special case Y uses $n = Y^k(t-)$ in 1 and 2.

An actual code can be managed efficiently by storing $\dot{X}(t)$ and $\dot{Y}(t)$ each in an array of c cells and using (8) at each step. Indeed, $h_{n-1}(\dot{X}(t)) + 1$ and $h_n(\dot{X}(t))$ are the indices of the first and last cells storing the value n in $\dot{X}(t)$ and can be found in $O(\log c)$ time by binary search of the sorted array $\dot{X}(t)$. Usually the load balancing policy provides an index j such that $\dot{X}^j(t-) = n$, and the search for $h_{n-1}(\dot{X}(t)) + 1$ [resp. $h_n(\dot{X}(t))$] can be limited to the cells before [resp. after] the j-th.

Remark 2. We shall develop perfect simulation methods for π which use the coupling (X, Y), and actually simply (X, \acute{Y}) . The arguments are readily adapted for exchangeable policies using (\acute{X}, \acute{Y}) , providing perfect simulation methods for the invariant law $\acute{\pi}$ of \acute{X} . If f is symmetric then $\int f d\pi = \int f d\acute{\pi}$, e.g., the β_n are of this form, else a uniform random permutation of the coordinates of a sample from $\acute{\pi}$, independent of the rest, yields a sample from π .

4.3 General load balancing policies

Certain load balancing policies need consider the specific indices of the queues. For instance the network may have a directed graph structure and may implement policies such as those in Section 2.1 only on the graph neighborhood of the uniformly chosen queue. Then X is not exchangeable and \hat{X} is not Markov.

Sorting X(t) at each update would take $O(c \log c)$ time. It is usually much more efficient to store and update

$$X(t) \triangleq (X^{i}(t))_{1 \le i \le c}, \qquad \sigma(t) \triangleq (\sigma(t)(i))_{1 \le i \le c}, \tag{9}$$

satisfying (3), in two arrays of c cells, and use when needed that

$$\dot{X}^{i}(t) = X^{\sigma(t)(i)}(t), \qquad 1 \le i \le c,$$
(10)

for instance to find $h_n(X(t)) = h_n(\dot{X}(t))$ by binary search of $\dot{X}(t)$, see Section 4.2, which uses only the elements for comparison and takes $O(\log c)$ time.

For the sake of simplicity, the tie-breaking rule used for $\sigma(t)$ in the exposition was by ascending queue indices, in which case every update of $\sigma(t)$ may require up to cmemory swaps. The following tie-breaking rule has efficient updates swapping at most the contents of two memory cells. The results of the exposition can readily be extended to it by considering the $(\mathcal{F}_t)_{t \in \mathbb{R}_+}$ -Markov process $(X(t), \sigma(t))_{t \in \mathbb{R}_+}$.

1. Let t > 0 be an arrival instant with mark k. The task is allocated to a queue of index $i = \sigma(t-)(j)$ and length $n \leq X^{\sigma(t-)(k)}(t-)$ and the updates are

$$X^{i}(t) = X^{i}(t-) + 1 = n+1,$$

$$\sigma(t)(h_{n}(X(t-))) = \sigma(t-)(j) = i, \quad \sigma(t)(j) = \sigma(t-)(h_{n}(X(t-))),$$

i.e., $h_{n-1}(X(t-)) < j \le h_n(X(t-))$ and $\sigma(t)$ is obtained from $\sigma(t-)$ by swapping the index *i* in cell *j* with the index in cell $h_n(X(t-))$. This propagates (3).

2. Let t > 0 be a potential service completion instant with mark k. It is effective if and only if $X^{\sigma(t-)(k)}(t-) \ge 1$, and then a task is removed from a queue of index $i = \sigma(t-)(j)$ and length $n \ge X^{\sigma(t-)(k)}(t-) \ge 1$ and the only updates are

$$\begin{split} X^i(t) &= X^i(t-) - 1 = n-1 \,, \\ \sigma(t)(h_{n-1}(X(t-)) + 1) &= \sigma(t-)(j) = i \,, \quad \sigma(t)(j) = \sigma(t-)(h_{n-1}(X(t-)) + 1) \,. \end{split}$$

i.e., $h_{n-1}(X(t-)) < j \le h_n(X(t-))$ and $\sigma(t)$ is obtained from $\sigma(t-)$ by swapping the index *i* in cell *j* with the index in cell $h_{n-1}(X(t-)) + 1$. This propagates (3).

The special case Y has j = k in 1 and 2, but we shall see that we need only simulate Y as in Section 4.2 in the perfect simulation algorithms.

4.4 Subordination and embedded Markov chain

Let $(T_n)_{n\geq 1}$ denote the sequence of arrival and potential service completion instants set in increasing order and

$$\theta_n = \begin{cases} (1,k) & \text{if } T_n \text{ is an arrival instant with mark } k, \\ (-1,k) & \text{if } T_n \text{ is an potential service completion instant with mark } k. \end{cases}$$

Classically, $(T_n, \theta_n)_{n \ge 1}$ is a $(\mathcal{F}_t)_{t \in \mathbb{R}_+}$ -marked Poisson process of intensity $\lambda + c$, and the marks θ_n have common law

$$\left(\frac{c}{\lambda+c}\delta_{-1} + \frac{\lambda}{\lambda+c}\delta_{1}\right) \otimes \frac{1}{c}(\delta_{1} + \dots + \delta_{c}) \text{ on } \Theta \triangleq \{-1,1\} \times \{1,\dots,c\}.$$
(11)

The Markov process $(X(t), Y(t))_{t \in \mathbb{R}_+}$ is subordinate to $(T_n)_{n \ge 1}$, and thus has same invariant law κ as the embedded $(\mathcal{F}_{T_n})_{n \in \mathbb{N}_0}$ -Markov chain $(X_n, Y_n)_{n \in \mathbb{N}_0}$ defined by

$$(X_0, Y_0) = (X(0), Y(0)), \qquad (X_n, Y_n) = (X(T_n), Y(T_n)).$$

Simulating it avoids the simulation of exponential random variables $\mathcal{E}(\lambda + c)$.

We wish to implement the coupling $(X_n, Y_n)_{n \in \mathbb{N}_0}$ using update functions and i.i.d. random draws. For $n \geq 1$, X_n is obtained from X_{n-1} and θ_n according to the load balancing policy in a way which may involve identically distributed random variables independent of $\mathcal{F}_{T_{n-1}}$, e.g., to break ties between queues of equal length or implement JSQ(d), and Y_n is obtained from Y_{n-1} and θ_n deterministically, as in Section 3.2. We write this

$$X_n = \phi(X_{n-1}, \theta_n, U_n), \quad Y_n = \psi(Y_{n-1}, \theta_n), \quad (12)$$

using deterministic update functions $\phi : \mathbb{N}_0^c \times \Theta \times [0,1] \mapsto \mathbb{N}_0^c$ and $\psi : \mathbb{N}_0^c \times \Theta \mapsto \mathbb{N}_0^c$ and a $(\mathcal{F}_{T_n})_{n \in \mathbb{N}_0}$ -adapted sequence $(U_n)_{n \geq 1}$ of i.i.d. uniform random variables on [0,1]independent of $(X_0, Y_0, T_n, \theta_n; n \geq 1)$.

If X uses an exchangeable policy as does Y, then $(X_n)_{n \in \mathbb{N}_0}$ is a $(\mathcal{F}_{T_n})_{n \in \mathbb{N}_0}$ -Markov chain itself, see Section 4.2, and we can write

$$\dot{X}_n = \dot{\phi}(\dot{X}_{n-1}, \theta_n, U_n), \quad \dot{Y}_n = \dot{\psi}(\dot{Y}_{n-1}, \theta_n).$$
(13)

In the sequel we use $(X_n, Y_n)_{n \in \mathbb{N}_0}$ and (12) in three perfect simulation algorithms for the invariant law π of the queue length process X of a load balancing network. The invariant law κ has marginals π and $\gamma^{\otimes c}$, and $(Y_n)_{n \in \mathbb{N}_0}$ in equilibrium is reversible and will be used as a simulatable dominating Markov chain in this infinite state space. We shall see we need only simulate $(Y_n)_{n \in \mathbb{N}_0}$ using (13) for our purposes.

For the sake of simplicity, the exposition breaks ties by ascending queue index. In order to use the efficient tie-breaking rule of Section 4.3, we need only consider $(X_n, \sigma_n) = \phi(X_{n-1}, \sigma_{n-1}, \theta_n, U_n)$, etc., where $\sigma_n = \sigma(T_n)$.

Exchangeable policies require only to simulate π using $(X_n, Y_n)_{n \in \mathbb{N}_0}$ and (13).

5 Palm theory and acceptance-rejection

A classic result in Palm theory, see [33, Chap. 8, 10], is that a stationary version on \mathbb{R} of a cycle-stationary process can be obtained by length-biasing the law of the cycle-stationary process by the length of the cycle straddling the origin and then setting the origin uniformly in the length of this cycle. The trace on \mathbb{R}_+ of the cycle straddling the origin is called the delay.

In [33, Sect. 8.6.7] the simulation of the stationary delay is obtained by the following acceptance-rejection method.

- 1. Simulate the length of a stationary delay.
- 2. Simulate independent copies of the cycles, until the length of a cycle exceeds the previously simulated stationary delay length.
- 3. The part of this cycle from the stationary delay length onward constitutes a simulation of the stationary delay. In particular, the state of this cycle at the stationary delay length is a draw from the invariant law.

Moreover, the number of trials is finite, a.s., but has infinite expectation. This continuous-time result can be applied to discrete time as in [33, Sect. 10.2.6].

Sigman [19] provides an interesting implementation of this result in discrete time, and gives it a short direct proof in its Prop. 2.1. We adapt this implementation here as a warm-up for illustrative purposes, since the infinite expectation of the number of trials renders it unsuitable for actual Monte Carlo simulation. The situation is simpler than for DomCFTP since the simulations are all in direct time and there is no difficulty in preserving the coupling by using the same marks θ_n in (12).

The hitting times of (0,0) by the Markov chain $(X_n, Y_n)_{n \in \mathbb{N}_0}$ are regeneration instants. Theorems 2 and 3 yield that if $(X_0, Y_0) = (0,0)$ or follows the invariant law κ then $X_n \preccurlyeq Y_n$ for all n in \mathbb{N}_0 . Thus in both cases these regeneration instants are given by the hitting times of 0 of the Markov chain $(Y_n)_{n \in \mathbb{N}_0}$. These are the same as those of $(Y_n)_{n \in \mathbb{N}_0}$, and we need only simulate the latter.

Hence, the perfect simulation method first obtains the length L of a stationary delay by simulating an instance of $(Y_n)_{n \in \mathbb{N}_0}$ started at a draw from its invariant law, obtained by sorting in non-decreasing order a draw from $\gamma^{\otimes c}$ (Section 4.1). It then determines the first cycle of length exceeding L by simulating instances of $(Y_n)_{n \in \mathbb{N}_0}$ started at 0. When this is obtained, the method reuses the marks of this cycle in (12) to simulate the coupled X_0, \ldots, X_L , and outputs X_L as a perfect draw from π .

For exchangeable load balancing policies, the same reasoning is valid to simulate the invariant law π of \hat{X} by using (13) in place of (12); see Remark 2.

We recapitulate this in the following perfect simulation algorithm, written in mathematical fashion and not in pseudocode. It has infinite expected running time and is not to be implemented in practice.

Algorithm 1 Acceptance-rejection algorithm				
1: draw Y_0 from $\gamma^{\otimes c}$ and sort into Y_0	\triangleright see Section 4.1			
2: for $n = 1$ to $L \triangleq \inf\{n > 1 : Y_n = 0\}$				
3: draw \hat{Y}_n using \hat{Y}_{n-1}	\triangleright using (13) or otherwise			
4: for $k = 1$ to $K \triangleq \inf\{k \ge 1 : L^{(k)} = L\}$				
5: set $\dot{Y}_0^{(k)} = 0$				
6: for $n = 1$ to $L^{(k)} \triangleq \inf\{n \ge 1 : \acute{Y}_n^{(k)} = 0\} \land L$				
7: draw $\theta_n^{(k)}$ from (11)				
s: set $\dot{Y}_{n}^{(k)} = \dot{\psi}(\dot{Y}_{n-1}^{(k)}, \theta_{n}^{(k)})$				
9: set $X_0^{(K)} = 0$				
10: for $n = 1$ to L				
11: draw U_n uniformly in $[0, 1]$	\triangleright for possible random choices			
12: set $X_n^{(K)} = \phi(X_{n-1}^{(K)}, \theta_n^{(K)}, U_n)$	$\triangleright \text{ the } \theta_n^{(K)} \text{ of Loop } 4$			
13: return $X_L^{(K)}$				

Theorem 4. Let $\lambda < c$, and π be the invariant law of the queue length process X of a load balancing network of Definition 1. Algorithm 1 uses simulatable draws, terminates in an a.s. finite time with infinite expectation, and outputs a perfect draw from π , i.e., $\mathbb{P}(K < \infty) = 1$, $\mathbb{E}(K) = \infty$, and $X_L^{(K)}$ has law π . If the load balancing policy is exchangeable, Algorithm 1 with $X^{(K)}$ and (12) replaced by $\hat{X}^{(K)}$ and (13) in Loop 10 outputs a perfect draw from the invariant law $\hat{\pi}$ of \hat{X} ; see Remark 2.

Proof. The proof is given in the preceding discussion.

6 Dominated Coupling From The Past (DomCFTP)

6.1 Background

Coupling From The Past (CFTP) was introduced by Propp and Wilson [12]. It concerns a Markov chain, on a finite state space with an order and a least and a largest state, which can be constructed using an update function preserving this order and an i.i.d. sequence of random variables. This seminal paper provided impressive applications to the Ising model close to criticality and much general insight.

There is first a *back-off* strategy. Instances of the Markov chain are repeatedly started from the least and the largest states further and further back in the past before time 0. All these instances are coupled by always using the same i.i.d. random variables driving the updates from times n - 1 to n. This is repeated until *coalescence*

happens: the two instances have coalesced by time 0, *i.e.*, they have taken the same value and necessarily remained equal ever since. Due to the preservation of the order, all instances of the Markov chain started arbitrarily at the corresponding time back in the past must have coalesced, hence the common value obtained at time 0 must coincide with the value of the Markov chain in equilibrium at time 0 and constitute a perfect draw from the invariant law.

Dominated CFTP (DomCFTP) extends this to infinite state spaces such as \mathbb{N}_0^c , see [13–18]. DomCFTP simulates backward in time a dominating Markov chain in order to obtain lower and upper bounds for the Markov chain in equilibrium, coupled in the subsequent forward simulation together with the dominating chain.

6.2 DomCFTP at the empty state

A novelty here is the use of the preorder of Section 3.1 in the truly multi-dimensional setting of queueing networks. The order "up to coordinate permutation" on the quotient set for the indifference relation \sim is strictly weaker than the classic product order, and the latter is not in general preserved by the coupling of Section 3.2.

Perfect simulation methods for non-exchangeable polices are of great interest, see Remark 1, and require the use of the preorder with its lack of anti-symmetry. If the policy is exchangeable then we can use the reduced representation of Section 4.2 on which the preorder induces an order; we shall see a use of this in Section 6.3.

A way to implement DomCFTP is to wait for the dominating chain simulated backward in time to hit 0. At this random time, the chain in equilibrium must be at state 0 (coalescence is achieved without anti-symmetry) and can be simulated forward in time up to time 0, at which time its state constitutes a perfect draw from π .

A delicate point is that the Markov chains must be coupled together through the backward and forward simulations by using the same driving i.i.d. random variables. We deal with this by considering the two-sided stationary Markov chain

$$(X_n, \dot{Y}_n, \theta_{n+1})_{n \in \mathbb{Z}}.$$

Theorem 5. The two-sided stationary Markov chain $(\acute{Y}_n, \theta_{n+1})_{n \in \mathbb{Z}}$ in equilibrium can be perfectly simulated backward in time from time 0. Using reversibility, $(\acute{Y}_n)_{n \in \mathbb{Z}_-}$ can be simulated by drawing Y_0 of law $\gamma^{\otimes c}$, sorting it into $\acute{Y}_0 \triangleq \acute{Y}'_0$, simulating a copy $(\acute{Y}'_n)_{n \in \mathbb{N}_0}$ of $(\acute{Y}_n)_{n \in \mathbb{N}_0}$, and setting $\acute{Y}_n = \acute{Y}'_{-n}$ for $n \leq 0$. Moreover,

• if $Y_{n+1} = y + e_i$ and $Y_n = y$ and $y^i = m$ for y in \mathbb{N}_0^c and i in $\{1, \ldots, c\}$, then

 $\theta_{n+1} = (1, K)$ with K uniform in $\{h_{m-1}(y) + 1, \dots, h_m(y) = i\} \neq \emptyset$,

• if $\dot{Y}_{n+1} = y - e_i$ and $\dot{Y}_n = y$ and $y^i = m \ge 1$ for y in \mathbb{N}_0^c and i in $\{1, \ldots, c\}$, then

 $\theta_{n+1} = (-1, K)$ with K uniform in $\{h_{m-1}(y) + 1 = i, \dots, h_m(y)\} \neq \emptyset$,

• if $\acute{Y}_{n+1} = y$ and $\acute{Y}_n = y$ for y in \mathbb{N}_0^c , then

 $\theta_{n+1} = (-1, K)$ with K uniform in $\{1, ..., h_0(y)\} \neq \emptyset$.

Proof. The statement about reversibility is standard. The statements about uniform laws follow from the fact that the θ_n have law (11) and (8) and conditioning. For completeness, we provide a direct proof of the last case, the others being similar. If k is in $\{1, \ldots, h_0(y)\} \neq \emptyset$ then $y^k = 0$ and

$$\begin{split} \mathbb{P}(\theta_{n+1} = (-1,k) \,|\, \dot{Y}_{n+1} = y, \dot{Y}_n = y) &\triangleq \frac{\mathbb{P}(\dot{Y}_n = y, \theta_{n+1} = (-1,k), \dot{Y}_{n+1} = y)}{\mathbb{P}(\dot{Y}_n = y, \dot{Y}_{n+1} = y)} \\ &= \frac{\mathbb{P}(\dot{Y}_n = y, \theta_{n+1} = (-1,k))}{\mathbb{P}(\dot{Y}_n = y, \dot{Y}_{n+1} = y)} \\ &= \frac{\mathbb{P}(\dot{Y}_n = y) \,\mathbb{P}(\theta_{n+1} = (-1,k))}{\mathbb{P}(\dot{Y}_n = y) \,\mathbb{P}(\dot{Y}_n = y) \,\mathbb{P}(\dot{Y}_n = y)} \end{split}$$

and $\mathbb{P}(\theta_{n+1} = (-1,k)) = \frac{1}{c(\lambda+c)}$ and $\mathbb{P}(Y_{n+1} = y \mid Y_n = y) = \frac{h_0(y)}{c(\lambda+c)}$ yield that

$$\mathbb{P}(\theta_{n+1} = (-1,k) \,|\, \acute{Y}_{n+1} = y, \acute{Y}_n = y) = \frac{1}{h_0(y)} \,. \qquad \Box$$

We should consider \dot{Y}_n instead of Y_n for practical simulation issues even though the τ_n are not needed, for instance in order to allow to determine the h_m by binary search. A possibility for simulating $(\dot{Y}_n)_{n \in \mathbb{N}_0}$ is to use (13) as $\dot{Y}'_n = \dot{\psi}(\dot{Y}'_{n-1}, \theta'_n)$ for $n \ge 1$ with θ'_n i.i.d. of law (11), which results in

$$Y_0 \sim \gamma^{\otimes c}, \qquad \acute{Y}_n = \psi(\acute{Y}_{n+1}, \theta'_{-n}), \quad n \le -1.$$
 (14)

These θ'_n are only related indirectly to the θ_n in Theorem 5.

The DomCFTP method uses Theorem 5. It simulates backwards in time $(Y_n)_{n \in \mathbb{Z}_-}$ until the backward stopping time

$$N \triangleq \sup\{n \le 0 : \acute{Y}_n = 0\} = \sup\{n \le 0 : Y_n = 0\} \le 0$$
(15)

which will be proved to be finite, a.s., Theorem 3 yields that $0 \preccurlyeq X_N \preccurlyeq Y_N = 0$ and hence that $X_N = 0$ (even though \preccurlyeq is only a preorder). Moreover, θ_{n+1} can be simulated once Y_n has been simulated from Y_{n+1} for $n \leq -1$. Then $\theta_{N+1}, \ldots, \theta_0$ and (12) allow to simulate X_{N+1}, \ldots, X_0 , and X_0 is a perfect draw from π .

This results in the following perfect simulation algorithm, written in mathematical fashion and not in pseudocode. An actual implementation should follow Section 4.

Algorithm 2 DomCFTP at the empty state

1:	draw Y_0 from $\gamma^{\otimes c}$ and sort into Y_0	\triangleright see Section 4.1
2:	for $n = -1$ down to $N \triangleq \sup\{m \le 0 : \acute{Y}_m = 0\}$	\triangleright not executed if $Y_0 = 0$
3:	draw Y_n using Y_{n+1}	\triangleright see Theorem 5 and (14)
4:	draw θ_{n+1} using Y_{n+1} and Y_n	\triangleright see Theorem 5
5:	set $X_N = 0$	
6:	for $n = N + 1$ up to 0	\triangleright not executed if $Y_0 = 0$
7:	draw U_n uniformly in $[0,1]$	\triangleright for possible random choices
8:	set $X_n = \phi(X_{n-1}, \theta_n, U_n)$	\triangleright the θ_n obtained in Loop 2
9:	return X_0	

Theorem 6. Let $\lambda < c$, and π be the invariant law of the queue length process X of a load balancing network of Definition 1. Algorithm 2 uses simulatable draws, terminates after a duration which has some exponential moments, and outputs a perfect draw from π . Specifically, there exists s > 1 such that $\mathbb{E}(s^{|N|}) < \infty$, hence $\mathbb{E}(|N|) < \infty$, and X_0 has law π . If the load balancing policy is exchangeable, Algorithm 2 with X and (12) replaced by \dot{X} and (13) in Loop 6 outputs a perfect draw from the invariant law π of \dot{X} ; see Remark 2.

Proof. The proof follows from the preceding discussion except the statements on N. Now, -N = |N| has same law as the hitting time of 0 by $(Y_n)_{n\geq 0}$ in equilibrium. Classically, when c = 1 the hitting time of 0 by $(Y_n)_{n\geq 0}$ starting at 0 has some exponential moments, and the hitting time of 0 by $(Y_n)_{n\geq 0}$ in equilibrium is a sized biased version of it multiplied by an independent uniform random variable on [0, 1] and thus also has some exponential moments. The conclusion follows from [35, Sect. II.1.4], or [33, Sect. 10.7] adapted to discrete time, and induction over $c \geq 2$.

Remark 3. This algorithm allows perfect simulation of the invariant laws of coupled networks with varied load balancing policies, for instance in order to compare their performances. The backward simulation of (Y_0, \ldots, Y_N) and of $(\theta_0, \ldots, \theta_{N+1})$ is performed once. Then the queue length processes of these networks are simulated forward in time from time N and state 0 up to time 0, possibly in parallel fashion.

6.3 DomCFTP with back-off and sandwiching

In practice we expect c to be at least moderately large and λ to be close to c. Then the time |N| that $(Y_n)_{n \in \mathbb{Z}_-}$ takes to reach the empty state 0 and the duration of the simulation procedure in Theorem 6 are likely to be quite long. It takes much shorter times for $(Y_n)_{n \in \mathbb{Z}_-}$ to reach states in which some queues but not all are empty, which are those such that a queue in $(X_n)_{n \in \mathbb{Z}_-}$ may finish serving a task while $(Y_n)_{n \in \mathbb{Z}_-}$ does not change. A corresponding back-off method may achieve a considerable speed-up.

An important issue is that the preorder \preccurlyeq can be used for these purposes only if the update function ϕ preserves the quotient order for the indifference relation \sim in Theorem 1. This is very close to requiring that the load balancing policy be exchangeable; an uninteresting generalization is for instance a policy such as JSQ or

JSQ(d) breaking ties by ascending queue index. For the sake of clarity, we henceforth focus on exchangeable load balancing networks, see Section 4.2.

We are going to provide a DomCFTP method for the perfect simulation of the invariant law π of the instantaneous non-decreasing ordering X of X using the Markov chain $(X_n)_{n \in \mathbb{Z}_-}$. It requires to choose a back-off strategy

$$-\infty < \dots < M_k < \dots < M_1 < 0 \tag{16}$$

constituted of stopping times of the dominating Markov chain $(Y_n)_{n \in \mathbb{Z}_-}$ simulated backward in time. A classic choice is binary (or doubling) back-off, in which

$$M_k = 2M_{k-1} = 2^{k-1}M_1, \quad k \ge 2.$$

It has some empirical and theoretical backing, see [12, Sect. 5.2]. Usually M_1 is deterministic but could well be random, such as the first backward time a queue empties for the dominating Markov chain.

Recall that Y_0 can be drawn from $\gamma^{\otimes c}$, see Section 4.1, and rearranged into Y_0 . If $Y_0 = 0$ then coalescence has been achieved at time $0 \triangleq M_0$ and the procedure stops, else we recursively back-off as follows. Recall Theorem 5.

For $k \geq 1$, if coalescence has not been achieved yet after simulating forward from time M_{k-1} , after which

$$\dot{Y}_0, \quad \dot{Y}_{-1}, \theta_0, \dots, \dot{Y}_{M_{k-1}}, \theta_{M_{k-1}+1}, \quad U_{M_{k-1}+1}, \dots, U_0$$

have been determined, then M_k is determined by simulating backward in time

$$Y_{M_{k-1}-1}, \theta_{M_{k-1}}, \ldots, Y_{M_k}, \theta_{M_k+1}$$

Then we construct forward in times, drawing $U_{M_k+1}, \ldots, U_{M_{k-1}}$ and reusing $U_{M_{k-1}+1},\ldots,U_0,$

$$\dot{X}_{M_k} = 0, \qquad \dot{X}_{M_k+1} = \phi(\ddot{X}_{M_k}, \theta_{M_k+1}, U_{M_k+1}), \dots, \dot{X}_0 = \phi(\ddot{X}_1, \theta_0, U_0),
\dot{X}_{M_k} = \dot{Y}_{M_k}, \qquad \dot{X}_{M_k+1} = \phi(\dot{X}_{M_k}, \theta_{M_k+1}, U_{M_k+1}), \dots, \dot{X}_0 = \phi(\dot{X}_1, \theta_0, U_0).$$

The (unknown) coupled $(X_n)_{n \in \mathbb{Z}_-}$ in equilibrium satisfies

.

$$\check{X}_n \preccurlyeq \check{X}_n \preccurlyeq \check{X}_n, \quad n = M_k, \dots, 0$$

Since the update mechanism preserves the preorder \preccurlyeq and that this constitutes an order on non-decreasing c-tuples, if $\check{X}_n = \hat{X}_n$ for some n in $\{M_k, \ldots, 0\}$ then $\check{X}_m =$ \hat{X}_m for all $n \leq m \leq 0$, which constitutes the event of coalescence, and in particular

$$\check{X}_0 = \acute{X}_0 = \hat{X}_0$$

constitutes a perfect draw from the invariant law π of X.

This back-off procedure must be repeated for $k \ge 1$ until coalescence is achieved. This happens at least as soon as the time to which the backward simulation is continued is lesser than or equal to N defined in (15), which ensures some exponential moments for the duration, see Theorem 6.

Note that in the above $U_{M_{k-1}+1}, \ldots, U_0$ do not affect the (Y_n, θ_{n+1}) , and can be redrawn anew since their reuse is not required for the coupling to be preserved.

This results in the following perfect simulation algorithm, written in mathematical fashion and not in pseudocode. An actual implementation should follow Section 4. The algorithm requires to choose a suitable back-off strategy (16).

Alg	gorithm 3 DomCFTP with back-off and	sandwiching		
1:	draw Y_0 from $\gamma^{\otimes c}$ and sort into Y_0	\triangleright see Section 4.1		
2: for $k = 1$ to $K = \inf\{k \ge 1 : \check{X}_0^{(k)} = \hat{X}_0^{(k)}\}$				
3:	determine M_k			
4:	for $n = M_{k-1} - 1$ down to M_k			
5:	draw Y_n using Y_{n+1}	\triangleright see Theorem 5 and (14)		
6:	draw θ_{n+1} using Y_{n+1} and Y_n	\triangleright see Theorem 5		
7:	set $\check{X}_{M_{k}}^{(k)} = 0$ and $\hat{X}_{M_{k}}^{(k)} = \acute{Y}_{M_{k}}$			
8:	for $n = M_k + 1$ up to $n = 0$			
9:	draw U_n uniformly in $[0,1]$	\triangleright reuse or redraw U_n for $n \ge M_{k-1} + 1$		
10:	set $\check{X}_n^{(k)} = \acute{\phi}(\check{X}_{n-1}^{(k)}, \theta_n, U_n)$	\triangleright the θ_n obtained in Loop 4		
11:	set $\hat{X}_n^{(k)} = \hat{\phi}(\hat{X}_{n-1}^{(k)}, \theta_n, U_n)$			
12:	return $\hat{X}_0^{(K)}$			

Theorem 7. Let $\lambda < c$, and π be the invariant law of the queue length process X of a load balancing network of Definition 1. Algorithm 3 uses simulatable draws, terminates after a duration which has some exponential moments, and outputs a perfect draw from π . Specifically, there exists s > 1 such that $\mathbb{E}(s^K) < \infty$, thus $\mathbb{E}(K) < \infty$, and $\hat{X}_0^{(K)}$ has law π .

Proof. The proof follows from the preceding discussion.

Remark 4. This algorithm can be used for the task described in Remark 3. The backward simulation of (Y_0, \ldots, Y_{M_k}) and of $(\theta_0, \ldots, \theta_{M_k+1})$ should be continued back in the past for $k \geq 1$ until coalescence has been achieved for all the queue length processes of the networks under consideration. In the case of parallel computing, the backward simulation of $(Y_n)_{n \in \mathbb{Z}_-}$ may be continued while forward lower and upper bound simulation at a M_k are undertaken.

6.4 Assessment of the DomCFTP algorithms

In [13, Sect. 6], the authors assess and compare the two algorithms that they have developed for M/G/c.

They do so mainly in the Markovian case M/M/c. The invariant law is well-known – its normalizing constant is obtained from the well-tabulated Erlang C formula – and used for comparisons with the simulated draws. The dominating Markov chain is given from c i.i.d. M/M/1 queues as in the present paper, and can be readily perfectly simulated without using the Pollaczek-Khintchine formula.

Simulations indicate that the more sophisticated back-off sandwiching algorithm takes considerably less time than the simpler emptying algorithm. The authors indicate that the discussion is far from complete and recall that the computational demands of the sandwiching algorithm are greater for similar durations. Some theoretic and heuristic arguments on the durations are also provided.

7 Conclusion

The main scope of this paper is to provide efficient DomCFTP methods in the truly multi-dimensional setting of a wide class of Markovian load balancing queueing networks. The policies we consider are quite general and natural, and include classic polices for parallel servers as well as policies implementing instantaneous work stealing at service completions. These policies may respect topological constraints such as a directed graph structure, which breaks the exchangeability assumption under which most asymptotic performance evaluation is achieved.

Such perfect simulation methods are valuable for Monte Carlo estimation of quantities of interest in equilibrium, most notably for performance evaluation. This is achieved on this infinite state space using an adequate coupling providing a dominating Markov chain in an original preorder. It corresponds to an order up to permutations of the coordinates which is strictly weaker than the classic product order.

The paper by Connor and Kendall [13] has been an inspiration, and we are glad to have provided an answer to Question 3 in its Section 7. There are many remaining questions left unanswered, and we list a few.

- 1. Provide mathematical evaluation of the benefits of the back-off and sandwiching algorithm over the emptying algorithm ([13, Sect. 7 Question 2]).
- 2. Implement niceties about DomCFTP such as recycling, interruptible algorithms, read-once, see [18, Chap. 5], simultaneous implementation for a suitable range of values of c called "omnithermal DomCFTP" in [13, Sect. 7 Question 4]), etc. We have noted that we may reuse or redraw the U_n at will.
- 3. Extend results to non-Markovian queueing networks with general service distributions or general interarrival distributions or both, see [2, Sect. 9], [36].
- 4. Tell interesting things about asymmetric queueing networks, see [2, Sect. 9], beyond the fact that the first two perfect simulation methods in this paper apply to the networks of Definition 1.
- 5. Find other wide natural classes of networks, couplings, and orders for which the ideas in this work can be adapted.

Acknowledgments. The author thanks Stephen Connor for a private communication about the mean run time of the DomCFTP algorithms in [13]. This paper was an inspiration for us.

References

- Gupta, V., Walton, N.: Load balancing in the nondegenerate slowdown regime. Oper. Res. 67(1), 281–294 (2019) https://doi.org/10.1287/opre.2018.1768
- Boor, M., Borst, S.C., Leeuwaarden, J.S.H., Mukherjee, D.: Scalable load balancing in networked systems: a survey of recent advances. SIAM Rev. 64(3), 554–622 (2022) https://doi.org/10.1137/20M1323746
- [3] Elsayed, E.A., Bastani, A.: General solutions of the jockeying problem. Eur. J. Oper. Res. 22, 387–396 (1985) https://doi.org/10.1016/0377-2217(85)90258-9
- [4] Zhao, Y., Grassmann, W.K.: Queueing analysis of a jockeying model. Oper. Res. 43(3), 520–529 (1995) https://doi.org/10.1287/opre.43.3.520
- [5] Mitzenmacher, M.: Analyses of load stealing models based on families of differential equations. Theory Comput. Syst. 34(1), 77–98 (2001) https://doi.org/10.1007/s002240010010
- [6] Tchiboukdjian, M., Gast, N., Trystram, D., Roch, J.-L., Bernard, J.: A tighter analysis of work stealing. In: Algorithms and Computation. 21st International Symposium, ISAAC 2010, Jeju, Korea, December 15–17, 2010. Proceedings, Part II, pp. 291–302. Springer, Berlin (2010). https://doi.org/10.1007/978-3-642-17514-5_25 . infoscience.epfl.ch/record/183602
- [7] Suksompong, W., Leiserson, C.E., Schardl, T.B.: On the efficiency of localized work stealing. Inf. Process. Lett. 116(2), 100–106 (2016) https://doi.org/10.1016/j.ipl.2015.10.002
- [8] Kielanski, G., Van Houdt, B.: Performance analysis of work stealing strategies in large scale multi-threaded computing. In: Quantitative Evaluation of Systems. 18th International Conference, QEST 2021, Paris, France, August 23–27, 2021. Proceedings, pp. 329–348. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-85172-9_18
- [9] Turner, S.R.E.: The effect of increasing routing choice on resource pooling. Probability in the Engineering and Informational Sciences 12(1), 109–124 (1998) https://doi.org/10.1017/S0269964800005088
- [10] Budhiraja, A., Mukherjee, D., Wu, R.: Supermarket model on graphs. The Annals of Applied Probability 29(3), 1740–1777 (2019) https://doi.org/10.1214/18-AAP1437
- [11] Loynes, R.M.: The stability of a queue with nonindependent inter-arrival and service times. Proc. Camb. Philos. Soc. 58, 497–520 (1962)

- [12] Propp, J.G., Wilson, D.B.: Exact sampling with coupled Markov chains and applications to statistical mechanics. Random Structures Algorithms 9(1-2),223 - 252(1996)and https://doi.org/10.1002/(sici)1098-2418(199608/09)9:1/2(223::aid-rsa14)3.0.co;2-o
- [13] Connor, S.B., Kendall, W.S.: Perfect simulation of M/G/c queues. Adv. Appl. Probab. 47(4), 1039–1063 (2015) https://doi.org/10.1239/aap/1449859799
- [14] Kendall, W.S.: Perfect simulation for the area-interaction point process. In: Probability Towards 2000, pp. 218–234. Springer, New-York (1998)
- [15] Kendall, W.S., Møller, J.: Perfect simulation using dominating processes on ordered spaces, with application to locally stable point processes. Adv. Appl. Probab. 32(3), 844–865 (2000) https://doi.org/10.1239/aap/1013540247
- [16] Thönnes, E.: A primer in perfect simulation. In: Mecke, K.R., Stoyan, D. (eds.) Statistical Physics and Spatial Statistics, pp. 349–378. Springer, Berlin (2000)
- [17] Kendall, W.S.: Notes on perfect simulation. In: Markov Chain Monte Carlo, pp. 93–146. World Scientific, Hackensack NJ (2005)
- [18] Huber, M.L.: Perfect Simulation. Monogr. Stat. Appl. Probab., vol. 148. CRC Press, Boca Raton, FL (2016). www.crcnetbase.com/isbn/9781482232455
- [19] Sigman, K.: Exact simulation of the stationary distribution of the FIFO M/G/c queue: the general case for $\rho < c$. Queueing Syst. **70**(1), 37–43 (2012) https://doi.org/10.1007/s11134-011-9266-6
- [20] Winston, W.: Optimality of the shortest line discipline. J. Appl. Probability 14(1), 181–189 (1977) https://doi.org/10.1017/s0021900200104772
- [21] Liu, Z., Nain, P., Towsley, D.: Sample path methods in the control of queues. Queueing Systems Theory Appl. 21(3-4), 293–335 (1995) https://doi.org/10.1007/BF01149166
- [22] Fayolle, G., Iasnogorodski, R.: Two coupled processors: the reduction to a Riemann-Hilbert problem. Z. Wahrsch. Verw. Gebiete 47(3), 325–351 (1979) https://doi.org/10.1007/BF00535168
- [23] Kurkova, I.A., Suhov, Y.M.: Malyshev's theory and JS-queues. Asymptotics of stationary probabilities. Ann. Appl. Probab. 13(4), 1313–1354 (2003) https://doi.org/10.1214/aoap/1069786501
- [24] Lu, Y., Xie, Q., Kliot, G., Geller, A., Larus, J., Greenberg, A.: Join-idle-queue: A novel load balancing algorithm for dynamically scalable web services. Performance Evaluation 68(11), 1056–1071 (2011)

- [25] Vvedenskaya, N.D., Dobrushin, R.L., Karpelevich, F.I.: A queueing system with a choice of the shorter of two queues: an asymptotic approach. Problems of Information Transmission 32(1), 15–27 (1996)
- [26] Mitzenmacher, M.D.: The power of two choices in randomized load balancing. PhD thesis, University of California, Berkeley (1996)
- [27] Graham, C.: Chaoticity on path space for a queueing network with selection of the shortest queue among several. J. Appl. Probab. 37(1), 198–211 (2000) https://doi.org/10.1017/s0021900200015345
- [28] Graham, C.: Functional central limit theorems for a large network in which customers join the shortest of several queues. Probab. Theory Related Fields 131(1), 97–120 (2005) https://doi.org/10.1007/s00440-004-0372-9
- [29] Atar, R.: A diffusion regime with nondegenerate slowdown. Oper. Res. 60(2), 490–500 (2012) https://doi.org/10.1287/opre.1110.1030
- [30] Rockafellar, R.T.: Convex Analysis. Princeton Mathematical Series, No. 28, p. 451. Princeton University Press, Princeton NJ (1970)
- [31] Stoyan, D.: Comparison Methods for Queues and Other Stochastic Models. Wiley Ser. Probab. Math. Stat., p. 217. John Wiley & Sons, Chichester (1983)
- [32] Shaked, M., Shantikumar, J.G.: Stochastic Orders. Springer Series in Statistics. Springer, New York (2007). https://doi.org/10.1007/978-0-387-34675-5
- [33] Thorisson, H.: Coupling, Regeneration. Proba-Stationarity, and bility and itsApplications, p. 517. Springer, New York (2000).https://doi.org/10.1007/978-1-4612-1236-2 https://doi.org/10.1007/978-1-. 4612-1236-2
- [34] Devroye, L.: Non-uniform Random Variate Generation. Springer, New York (1986)
- [35] Lindvall, E.T.: Lectures on the Coupling Method. Wiley Ser. Probab. Math. Stat. John Wiley & Sons, New York (1992)
- [36] Blanchet, J., Pei, Y., Sigman, K.: Exact sampling for some multi-dimensional queueing models with renewal input. Adv. Appl. Probab. 51(4), 1179–1208 (2019) https://doi.org/10.1017/apr.2019.45