

PERFECT SIMULATION OF M/G/c QUEUES

STEPHEN B. CONNOR,* *University of York*

WILFRID S. KENDALL,** *University of Warwick*

Abstract

In this paper we describe a perfect simulation algorithm for the stable M/G/c queue. Sigman (2011) showed how to build a dominated coupling-from-the-past algorithm for perfect simulation of the super-stable M/G/c queue operating under first-come–first-served discipline. Sigman’s method used a dominating process provided by the corresponding M/G/1 queue (using Wolff’s sample path monotonicity, which applies when service durations are coupled in order of initiation of service). The method exploited the fact that the workload process for the M/G/1 queue remains the same under different queueing disciplines, in particular under the processor sharing discipline, for which a dynamic reversibility property holds. We generalise Sigman’s construction to the stable case by comparing the M/G/c queue to a copy run under random assignment. This allows us to produce a naïve perfect simulation algorithm based on running the dominating process back to the time it first empties. We also construct a more efficient algorithm that uses sandwiching by lower and upper processes constructed as coupled M/G/c queues started respectively from the empty state and the state of the M/G/c queue under random assignment. A careful analysis shows that appropriate ordering relationships can still be maintained, so long as service durations continue to be coupled in order of initiation of service. We summarise statistical checks of simulation output, and demonstrate that the mean run-time is finite so long as the second moment of the service duration distribution is finite.

Keywords: Coalescence; dominated coupling-from-the-past; dynamic reversibility; first-come–first-served discipline; first-in–first-out discipline; Kiefer–Wolfowitz workload vector; pathwise domination; perfect simulation; processor sharing discipline; M/G/c queue; random assignment discipline; sandwiching; stable queue; stochastic ordering; super-stable queue

2010 Mathematics Subject Classification: Primary 65C05

Secondary 60K25; 60J05; 68U20

1. Introduction

Coupling methods for queues have a celebrated history, stretching back to the seminal work of Loynes [20], who discussed stability results for very general queues using what would today be described as coupling comparisons, together with recursive formulations of queueing dynamics using queues whose commencements originate further and further back in the past. More recently Wolff [32] (correcting [31]) showed how to establish comparisons between different queueing service disciplines for a single queue with multiple servers. Here the coupling argument involves assigning the n th service duration to the n th initiation of customer service (as opposed to the n th customer arrival time).

Received 4 March 2014; revision received 11 December 2014.

* Postal address: Department of Mathematics, University of York, York YO10 5DD, UK.

Email address: stephen.connor@york.ac.uk

** Postal address: Department of Statistics, University of Warwick, Coventry CV4 7AL, UK.

The work of [20], together with subsequent work on stochastically recursive systems and renovating events (for example [5]) is now viewed as one of the precursors of the celebrated *coupling-from-the-past* (CFTP) algorithm of Propp and Wilson [25], the first of a series of practical algorithms for performing exact or perfect simulation. It was thus a natural step from the first CFTP algorithms to consider how to apply CFTP to the problem of simulating from the equilibria of queueing systems: a variety of CFTP techniques applied to various queues of finite capacity were discussed in [24]. The finite capacity requirement appeared to be an inevitable constraint, imposed by the nature of the original algorithm [25], which uses maximal and minimal elements of the state-space to generate upper and lower bounding processes. Clearly, a queue without finite capacity will not in general possess a state-space with maximal element. However, the CFTP idea need not be limited to bounded processes. Working in the context of stochastic geometry, it has been shown how to replace the maximal element by a suitable dominating random process, so as to produce a *dominated* CFTP algorithm [14], [18]. Indeed, Kendall [15] has shown how in principle dominated CFTP can be applied to *any* regular geometrically ergodic Markov chain, using Foster–Lyapunov criteria, small set regeneration, and a dominating process which is in fact a D/M/1 queue. (Extensions to a class of nongeometrically ergodic chains are given in [8].) The work in [15] strongly indicates that one should search for *practical* dominated CFTP algorithms which solve the problem of simulating from the equilibria of geometrically ergodic queues. Very recently, Sigman [28] has shown how to use dominated CFTP in the so-called ‘super-stable’ case of the M/G/c queue with first-come–first-served (FCFS) discipline. Here ‘super-stable’ means that arrival and service rates are such that the queue remains stable even if $c - 1$ of the servers are removed.

Sigman’s work uses the fact that the M/G/c queue is dominated by an M/G/1 queue under FCFS discipline (equivalent in this single-server case to first-in–first-out (FIFO) discipline), which itself is stable by virtue of super-stability of the original queue; the trajectory of an equilibrium instance of the dominating queue can be generated from times in the arbitrarily distant past using the observation that the queue workload does not depend on service discipline, and therefore these trajectories can be reconstructed from a dynamically reversible relative which uses a processor sharing (PS) discipline. Thus, the strategy is as follows.

1. Given the same inputs, the total workload of a super-stable M/G/c queue is samplewise dominated by the total workload of an M/G/1 queue.
2. Using the PS discipline, an M/G/1 [PS] queue is dynamically reversible in time (and has the same workload process as the corresponding M/G/1 [FIFO] queue).
3. Thus, the M/G/1 [PS] queue can be used as a dominating process.
4. Coalescence occurs when the M/G/1 [PS] queue empties.
5. The workload process of the M/G/1 [PS] queue can be decoded to generate the arrival times and service durations of the underlying arrival process.

The second moment of the service duration must be finite if the algorithm is to have finite mean run-time.

Sigman’s approach [28] is limited to the super-stable case by the need to dominate the target queue using a simple M/G/1 queue. A different approach [29] uses regenerative techniques to extend to the merely stable case, but this different approach unfortunately results in run-times of infinite mean [30]. In this paper we show how to generalise the approach of [28] to deal with stable M/G/1 queues. The essence of the idea is to replace the dominating M/G/1

queue by an M/G/c queue run with random assignment (RA) (this may also be viewed as an independent collection of c different M/G/1 queues; hence, we write it indifferently as $M/G/c [RA]=[M/G/1]^c$); and then to use Wolff's observation [31] that samplewise monotonicity between M/G/c [FCFS] and $[M/G/1]^c$ can be arranged if service durations are assigned in order of initiation of service. The idea is simple enough; however, considerable care needs to be exercised in order to ensure that the dominating process really does dominate the target chain in an appropriate sense. (This idea, which forms the basis of our Algorithm 1, was proposed independently in [4], although that paper does not contain a proof of the algorithm's correctness.) Moreover, in order to achieve smaller run-times by using a refined algorithm, it is necessary to show that appropriate ordering (or *sandwiching*) relationships are maintained between upper and lower processes started at different initial times $-T$. The attraction of this extension to Sigman's work [28] is that it allows simulation methods to be applied precisely in the case when M/G/c queues will be most relevant; namely, when using a single server ($c = 1$) would result in loss of stability.

It is appropriate here to mention some further related papers on perfect simulation and queueing. Fernandez *et al.* [9] studied exclusion models via ensembles of Peierls contours in a spatial problem as a kind of spatially distributed loss network; however, the methods are specific to loss networks with Poisson inputs and exponential lifetimes. Blanchet and Dong [3] applied dominated CFTP to a GI/GI/c loss process, using a GI/GI/ ∞ queue as the dominating process. Rather than waiting until the dominating process empties (a time which generally grows exponentially in the arrival rate), they looked for a time interval $[a, b]$ for which all customers present at time a have departed by time b , and over the entirety of which the infinite server system has less than c customers. The two processes will have coalesced by time b , and so coalescence is determined by watching the dominating process alone. The authors overcame a significant technical difficulty in [3] by showing how to simulate a renewal process input in reverse time; however, their method for coupling target and dominating processes involved truncation of an infinite server system, and this cannot be applied in our M/G/c context. Mousavi and Glynn [22] discussed perfect simulation for reflected Brownian motion in a wedge (consequently gaining information on heavy-traffic approximation for queues). Attention was focused on stochastic differential equation problems, and links were made with the approach of [2] to exact simulation for solutions to stochastic differential equations. Blanchet and Chen [6] showed how to perform perfect simulation for the workload vector of a network of d queueing stations with Poisson inputs, in which each arrival brings a vector of service times describing the additional work to be carried out at each station. This requires the existence of a finite moment generating function for the vector of service times, although this is subsequently relaxed in [4].

We conclude this introductory section by setting out the plan of the paper. In Section 2 we review notation and fundamental facts for our target queue M/G/c and the intended dominating process $[M/G/1]^c$. In Section 3 we describe extensions (Theorems 3.2 and 3.3) of a classical domination result from queueing theory (Theorem 3.1), which prove the queue comparisons necessary to establish the required domination relationships. In Section 4 we provide a proof of a simple dominated CFTP algorithm (Algorithm 1) based on the regeneration which happens when the dominating process $[M/G/1]^c$ empties. However, the run-time for this algorithm will be large in cases when the target process M/G/c [FCFS] rarely empties, which is precisely the set of circumstances for which multi-server queues have practical utility! In Section 5 we describe and prove the validity of a refined algorithm (Algorithm 2), based on the sandwiching of the target process between pairs of upper and lower processes themselves generated from the dominating process: the regeneration used in the simple algorithm is replaced by consideration

of when these upper and lower processes agree at time 0. At the price of increased complexity (proving domination relationships hold not just between lower, target, upper, and dominating processes, but also between different pairs of upper and lower processes), the algorithm run-time can be substantially decreased. Empirical demonstrations of the savings which can be obtained, as well as the correctness of the algorithm in the computable $M/M/c$ case, are demonstrated by representative simulations in Section 6. Finally, in Section 7 we discuss further research possibilities.

2. Dynamic reversibility

In this section we work with the $M_\lambda/G/c$ queue with arrival rate λ under two different allocation rules: first-come–first-served ($M_\lambda/G/c$ [FCFS]) and random assignment ($M_\lambda/G/c$ [RA]); in this second case each arrival is assigned randomly and independently to one of the c servers without regard to load on each server. The $M_\lambda/G/c$ [RA] queue will serve as a dominating process for the dominated CFTP algorithm to be described in later sections. The queue $M_\lambda/G/c$ [RA] may be viewed as a system of c independent $M_{\lambda/c}/G/1$ [FCFS] queues, each with arrival rate λ/c . To emphasise this, we sometimes write $M_\lambda/G/c$ [RA] as $[M_{\lambda/c}/G/1]^c$.

We follow the notation of [1] and [28]. Firstly, we consider a general $\cdot/\cdot/c$ [FCFS] queue and review the Kiefer–Wolfowitz construction of a workload vector [19]. Let $V(t) = (V(1, t), V(2, t), \dots, V(c, t))$ denote the workload vector at time $t \geq 0$. To be explicit, the $V(1, t) \leq V(2, t) \leq \dots$ represent the ordered amounts of *residual* work in the system for the c servers at time t , bearing in mind the FCFS queueing discipline. Customer n arrives at time t_n (for $0 \leq t_1 \leq t_2 \leq \dots$). Interarrival times are $T_n = t_{n+1} - t_n$ (where we set $t_0 = 0$). Observing V just before arrival of the n th customer (but definitely after the arrival of the $(n - 1)$ th customer) generates the process W_n : in the $t_{n-1} < t_n$ case, we have $W_n = V(t_{n-})$. This satisfies the Kiefer–Wolfowitz recursion

$$W_{n+1} = R(W_n + S_n e - T_n f)^+ \quad \text{for } n \geq 0, \tag{2.1}$$

where

- $W_n + S_n e$ adds S_n to the first coordinate only of the vector W_n ;
- $W_n + S_n e - T_n f$ subtracts T_n from each of the coordinates of $W_n + S_n e$;
- $R(W_n + S_n e - T_n f)$ reorders the coordinates of the vector $W_n + S_n e - T_n f$ in increasing order;
- $R(W_n + S_n e - T_n f)^+$ replaces negative coordinates of $R(W_n + S_n e - T_n f)$ by 0s.

Since each of these operations is a coordinatewise monotonic function of the previous workload vector W_n and the service duration S_n , an argument from recursion shows that the coordinates of W_n depend monotonically on the initial workload vector and the sequence of service durations, once the arrival time sequence is fixed. See, for example, remarks in [23] on join shortest workload disciplines for systems of parallel FIFO queues—corresponding to $\cdot/\cdot/c$ [FCFS]. If $t_n \leq t < t_{n+1}$ then we obtain $V(t)$ from W_n by subtracting $t - t_n$ from all the workload components and then taking positive parts

$$V(t) = (W_n - (t - t_n) f)^+. \tag{2.2}$$

Arguing as before, the coordinates of $V(t)$ depend monotonically on the initial workload vector and the sequence of service durations, once the arrival time sequence is fixed.

We are specifically interested in the $M_\lambda/G/c$ [FCFS] queue with arrival rate λ and independent and identically distributed service durations S_n . Let G be the common distribution of the S_n , and set $\mathbb{E}[S] = 1/\mu$. We shall assume throughout that $\mathbb{E}[S^2] < \infty$, in order to guarantee finite mean run-time of our algorithms (as detailed in Section 6). Write $\rho = \lambda/\mu$; we consider the stable case $\rho < c$. We will compare this to the $[M_{\lambda/c}/G/1]^c$ system with total arrival rate λ and service durations as above. That is, rather than operating under FCFS, we assign incoming customers to one of c independent M/G/1 queues uniformly at random. Each of these queues sees arrivals at rate λ/c and, therefore, has subcritical traffic intensity $\lambda/(c\mu)$. As noted in [28], it is a classical fact from queueing theory that the workload of an individual $M_{\lambda/c}/G/1$ queue is invariant under changes of work-conserving discipline. We can exploit this by using the PS discipline, since under this discipline the single-server queue workload vector process can be viewed as dynamically reversible [26, Section 5.7.3]. This means that the reverse process is a system of the same type, with customers again arriving at a Poisson rate λ/c , and with workloads having the same distribution G as S_n , but with the state now representing the amount of work *already performed* on customers still in the system. Since each of the c independent copies of $M_{\lambda/c}/G/1$ is dynamically reversible under PS, it follows that the $M_\lambda/G/c$ [RA]=[$M_{\lambda/c}/G/1$]^c queue is itself dynamically reversible under PS applied to each component queue.

3. Domination of $M_\lambda/G/c$

In this section we develop results based on the observation [32] that it is possible to arrange for the $M_\lambda/G/c$ queue to be pathwise dominated by c -server queues using other queueing disciplines, if the two queues are coupled by listing initiations of service in order and assigning the same service duration to the n th initiation of service in each queue. (As noted below, this assignment in order of initiation of service is crucial.) The fundamental idea is to establish that the non-FCFS system completes less total work by any fixed time, since corresponding services initiate later (when listed in order of initiation as above).

Let Q_t denote the queue length at time t , and write $|V(t)| = V(1, t) + \dots + V(c, t)$ for the total workload (remaining work) at time t . We begin by citing a classic result proved in queueing theory monographs.

Theorem 3.1. ([1, Chapter XII].) *We consider an $M_\lambda/G/c$ queueing system under various queueing disciplines. We use \leq_{so} to refer to stochastic ordering of distribution functions, and use tildes to refer to quantities pertaining to the system when it evolves under a possibly non-FCFS allocation rule; unadorned quantities pertain to the system when it evolves under an FCFS allocation rule. For any (possibly non-FCFS) allocation rule, it holds for initially empty systems that*

$$Q_t \leq_{so} \tilde{Q}_t, \quad |V(t)| \leq_{so} |\tilde{V}(t)| \quad \text{for all } t \geq 0.$$

Similarly, $|W_n| \leq_{so} |\tilde{W}_n|$ for all n .

Note that the concept of a Kiefer–Wolfowitz workload vector is not well-defined for general non-FCFS queues: nevertheless, the total amount of work $|V(t)|$, respectively $|\tilde{W}_n|$, can be defined unambiguously as the total amount of residual work currently in the system. (Indeed, given an arrival at time t_n we can generate a service duration S_n and increase the total workload by this amount: $|V(t_n)| = |V(t_n^-)| + S_n$. Speaking algorithmically, this service duration is then stored ready for use at the time of the n th initiation of service; in non-FCFS queues this will not typically be the time at which the n th customer to arrive commences service.)

It follows immediately that the queue length and residual workload of the M/G/c queue under FCFS are stochastically dominated by those of the same queue with any alternative allocation rule. In fact this result generalises to general GI/G/c queues. However, the result does not carry over to domination in the sense of sample paths if the corresponding coupling assigns the same service duration to the same individual (where ‘same’ means same in order of arrival); see Wolff’s correction [32] of [31]. To establish such a domination, one has to take some care to link service durations between the two different systems in the right way; namely, to ensure that the same service duration is assigned to the n th initiation of service in each queue. For the purposes of our dominated CFTP argument, we need to generalise this result to cases when the allocation rule may change at some fixed time, and also to certain cases where each of an initial subsequence of service durations is reduced to 0 (this device allows us to include cases in which one of the systems is not empty at time 0).

The argument given below is a modest extension of [1, Chapter XII], but is central to the arguments of later sections of the current paper.

Theorem 3.2. *Consider an FCFS c-server queueing system viewed as a function of*

- (a) *the sequence of arrival times, (customers arriving at times $0 \leq t_1 \leq t_2 \leq t_3 \leq \dots$), and*
- (b) *the sequence of service durations S_1, S_2, S_3, \dots assigned in order of initiation of service (positive except for a possible initial subsequence of 0s).*

Then this system depends monotonically on the inputs $0 \leq t_1 \leq t_2 \leq t_3 \leq \dots$ and service durations S_1, S_2, S_3, \dots , in the sense that for each m the m th initiation of service J_m and the m th time of departure D_m are increasing functions of these inputs. Moreover, if the arrival times are fixed then for each $t \geq 0$ the Kiefer–Wolfowitz workload vector $\mathbf{V}(t)$ (considered coordinate-by-coordinate) depends monotonically on the initial workload vector (measured immediately after time 0) and the sequence of service durations corresponding to nonzero arrival times.

Proof. As noted in the discussion of (2.1) and (2.2), the coordinates of the Kiefer–Wolfowitz workload vector depend monotonically on the initial workload vector and the sequence of service durations once the sequence of arrival times is fixed. This settles the second part of the theorem. It remains to prove the first part.

Let $\tilde{t}_m \geq t_m$, $\tilde{S}_m \geq S_m$, \tilde{J}_m , and \tilde{D}_m be the time of arrival, the service duration, the time of initiation of service, and the time of departure for the m th individual of the comparison FCFS system. We need to show that $\tilde{J}_m \geq J_m$ and $\tilde{D}_m \geq D_m$ for all m . We have stipulated that the two sequences of service durations $\{S_m : m \geq 1\}$ and $\{\tilde{S}_m : m \geq 1\}$ contain positive service durations except for initial subsequences of 0s.

We will use an inductive proof, and prepare for this by establishing useful representations of D_m , \tilde{D}_m , and J_{m+c} , \tilde{J}_{m+c} . We concentrate on D_m , J_{m+c} , etc, for simplicity. First note that for any allocation policy

$$J_{m+c} \geq \max\{t_{m+c}, D_m\} \tag{3.1}$$

with equality holding in the case when the FCFS policy applies, since the $(m + c)$ th service starts either when the $(m + c)$ th customer arrives (if there is a spare server, which is to say when $D_m \leq t_{m+c}$), or if not then service initiates exactly when the relevant departure frees up a server: this happens at time D_m in the case of the FCFS policy, and otherwise can happen no earlier.

On the other hand, the m th departure time D_m is given by

$$D_m = \min^{(m)}\{J_1 + S_1, J_2 + S_2, J_3 + S_3, \dots\}, \tag{3.2}$$

where $\min^{(m)}$ denotes the operation that returns the m th order statistic. We now refine this so as to involve only finitely many times of completions of service on the right-hand side.

First note that only at most c customers can actually be in service at time J_{m+c} . Therefore,

$$D_m \leq J_{m+c}. \tag{3.3}$$

It follows that

$$D_m \leq J_{m+c} \leq J_{m+c+r} \leq J_{m+c+r} + S_{m+c+r} \quad \text{whenever } r \geq 0. \tag{3.4}$$

If $S_{m+c} > 0$ then also $S_{m+c+r} > 0$ for $r \geq 0$, so in this case $D_m \leq J_{m+c+r} < J_{m+c+r} + S_{m+c+r}$ for $r \geq 0$. Thus, in this case we can improve on (3.2) and write D_m in terms of an order statistic over a specific finite population:

$$D_m = \min^{(m)}\{J_1 + S_1, J_2 + S_2, J_3 + S_3, \dots, J_{m+c-1} + S_{m+c-1}\}. \tag{3.5}$$

On the other hand, if $S_{m+c} = 0$ then also $S_1 = S_2 = S_3 = \dots = S_{m+c} = 0$. In that case service is immediate on arrival, so $t_m = J_m = J_m + S_m = D_m$, and so (3.5) still holds (noting that there must be at least one server, so $c \geq 1$).

Consider the inductive hypothesis that $\tilde{D}_u \geq D_u$ for $u = 1, \dots, m - 1$ and $\tilde{J}_v \geq J_v$ for $v = 1, \dots, m + c - 1$. This holds for $m = 1$, since under FCFS the first c people are served at their arrival times, so $\tilde{D}_u = \tilde{t}_u + \tilde{S}_u \geq t_u + S_u$ for $u = 1, \dots, c$. Suppose that the inductive hypothesis holds for $m = n$. Then we can apply the monotonic formulae (3.1) and (3.5) and deduce that the inductive hypothesis holds for the $m = n + 1$ case too. Thus, the first part of the theorem follows by mathematical induction.

Consider two instances of M/G/c [FCFS], coupled monotonically using the construction implied in Theorem 3.2, based on the same sequence of arrival times, using sequences of service durations that agree once arrival times become positive, and such that the Kiefer–Wolfowitz workload vector of one strictly dominates that of the other at time 0+. We can remark that the queues can couple successfully (which is to say, attain the same state at the same time, called the *coupling time*) only at a time when both instances have idle servers. For the monotonicity implies that one has total workload strictly larger than the other up to the time when they first couple successfully. Since arrival times are fixed and shared by both systems, successful coupling of the two processes cannot occur at the time of an arrival (which simply increases the workload by equal amounts for each queue). On the other hand, if both queues have c or more individuals in the system then the workloads decrease at the same rate. It follows that successful coupling will occur at a time when

- (a) an arrival does not happen;
- (b) there are strictly fewer than c individuals in the smaller system (hence, the smaller system has an idle server).

The coupling implies that at the coupling time the same will be true of the larger system (that is, it too will have an idle server).

Theorem 3.3. Consider a c -server queueing system viewed as a function of

- (a) the sequence of arrival times (customers arriving at times $0 \leq t_1 \leq t_2 \leq t_3 \leq \dots$), and
- (b) the sequence of service durations S_1, S_2, S_3, \dots assigned in order of initiation of service (positive except for a possible initial subsequence of 0s).

Consider the following cases of different allocation rules, in some cases varying over time:

Case 1: $\cdot / \cdot / c$ [RA];

Case 2: $\cdot / \cdot / c$ [RA] until a specified nonrandom time T , then switching to $\cdot / \cdot / c$ [FCFS];

Case 3: $\cdot / \cdot / c$ [RA] until a specified nonrandom time T' , $0 \leq T' \leq T$, then switching to $\cdot / \cdot / c$ [FCFS];

Case 4: $\cdot / \cdot / c$ [FCFS].

For the sake of an explicit construction, when initiations of service tie then we break the ties using order of arrival time. On change of allocation rule to FCFS, customers in the system but not yet being served are placed at the front of the queue in order of arrival-time; service initiates immediately for the appropriate number of customers if there are servers free. If all this holds then the k case dominates the $k + 1$ case, in the sense that the m th initiation of service in the $k + 1$ case occurs no later than the m th initiation of service in the k case, and the m th departure in the $k + 1$ case occurs no later than the m th departure in the k case.

Moreover, for all times $t \geq T'$ the Kiefer–Wolfowitz workload vector for case 3 dominates (coordinate-by-coordinate) that of case 4, with similar domination holding for cases 2 and 3 for all $t \geq T$.

Proof. First observe that the desired relationships between case 1 and case 2, and between case 2 and case 3, follow immediately once we have established the desired relationship between case 3 and case 4. The two compared systems evolve in exactly the same way up to time T , respectively T' , and so we may simply argue in terms of the processes started at time T , respectively T' , for example, in case 1 adjusting the sequence of arrival times by $t_n \mapsto \min\{t_n - T, 0\}$ and replacing service durations of services initiated before T by the residual service duration at T .

The argument therefore depends on the comparison between case 3 and the FCFS case 4. Letting quantities with tildes refer to case 3, and using the notation of the proof of Theorem 3.2, we have, from the arguments for (3.5),

$$\begin{aligned} \tilde{D}_m &= \min^{(m)}\{\tilde{J}_1 + S_1, \tilde{J}_2 + S_2, \tilde{J}_3 + S_3, \dots, \tilde{J}_{m+c-1} + S_{m+c-1}\}, \\ D_m &= \min^{(m)}\{J_1 + S_1, J_2 + S_2, J_3 + S_3, \dots, J_{m+c-1} + S_{m+c-1}\}, \end{aligned}$$

(noting that here the service durations agree when considered in order of initiation) and from the arguments for (3.1), and the fact that FCFS holds for case 4,

$$\tilde{J}_{m+c} \geq \max\{t_{m+c}, \tilde{D}_m\}, \quad J_{m+c} = \max\{t_{m+c}, D_m\}.$$

Consider the inductive hypothesis that $\tilde{D}_u \geq D_u$ for $u = 1, \dots, m - 1$ and $\tilde{J}_v \geq J_v$ for $v = 1, \dots, m + c - 1$. This holds for $m = 1$, since under FCFS the first c people are served at their arrival times while service cannot occur earlier under any other allocation policy, so $\tilde{J}_u \geq J_u = t_u$ and $\tilde{D}_u = \tilde{J}_u + S_u \geq t_u + S_u$ for $u = 1, \dots, c$. Suppose that the inductive

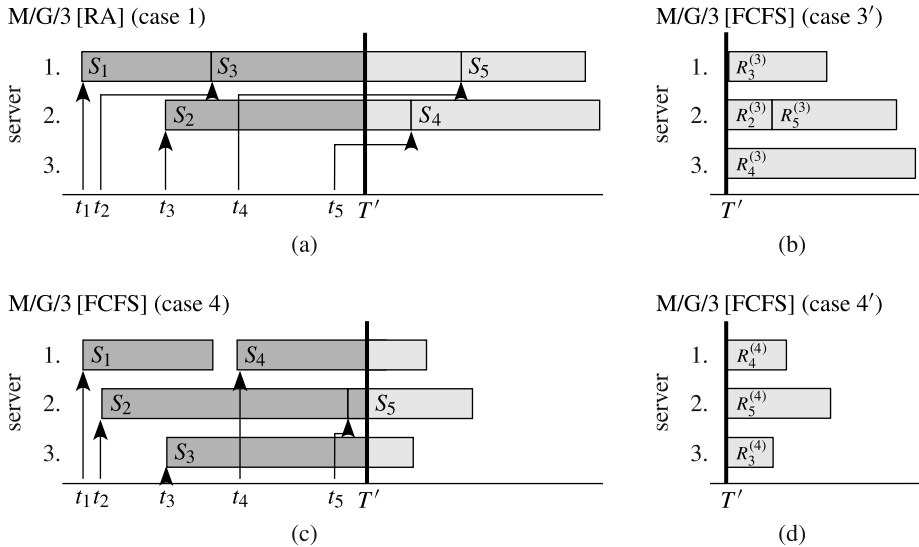


FIGURE 1: Illustration (with three servers) of how cases 3' and 4' are instantiated at time T' from the evolution over $[0, T')$ of cases 1 and 4, respectively. (a) Arrivals at times t_1, \dots, t_5 and a possible allocation to servers 1–3 in the M/G/3 [RA] process. Lengths of blocks represent size of workloads S_1, \dots, S_5 ; work completed by time T' is shaded dark grey. (b) Allocation of the residual workloads to servers at time T' in case 3' (under FCFS). (c) Illustration of how the same set of arrivals and service durations are handled by the FCFS process in case 4, and (d) how the corresponding residual workloads are used to instantiate case 4' at T' .

hypothesis holds for $m = n$. Then we can apply the above monotonic formulae and deduce that the inductive hypothesis holds for the $m = n + 1$ case too. Thus, the theorem follows by mathematical induction.

We now prove the claimed Kiefer–Wolfowitz domination between cases 3 and 4 for all times $t \geq T'$. (The proof that case 2 dominates case 3 for $t \geq T$ follows similarly.) To do this we construct two new $\cdot / \cdot / c$ [FCFS] processes (called cases 3' and 4') as follows. Both systems have arrival times $t'_i = \max\{t_i, T'\}$, and case 3' (respectively 4') has service durations given by the residual service durations in case 1 (respectively case 4) at time T' . That is, the system in case 3' has service durations $R_1^{(3)}, R_2^{(3)}, \dots$, where

$$R_i^{(3)} = \max\{J_i^{(1)} + S_i, T'\} - \max\{J_i^{(1)}, T'\},$$

case 4' has service durations $R_1^{(4)}, R_2^{(4)}, \dots$, where

$$R_i^{(4)} = \max\{J_i^{(4)} + S_i, T'\} - \max\{J_i^{(4)}, T'\},$$

and where $J_i^{(1)}$ and $J_i^{(4)}$ are the times of i th initiation of service in cases 1 and 4, respectively (see Figure 1).

Now consider the Kiefer–Wolfowitz vectors for these processes. We claim that (with subscripts corresponding in an obvious way to the cases being considered)

- (i) $V_{4'}(T') = V_4(T')$;
- (ii) $V_{3'}(T') = V_3(T')$;
- (iii) $V_{4'}(T') \leq V_{3'}(T')$, where \leq denotes coordinatewise domination.

The proof that $V_4(t) \leq V_3(t)$ for all $t \geq T'$ follows from these three claims: it is immediate that the required domination holds at time T' ; domination at all subsequent times follows by applying the final part of Theorem 3.2, since cases 3 and 4 both operate a FCFS policy over $[T', \infty)$, with the same arrival times and associated service durations over this period.

It therefore remains to prove claims (i)–(iii) above. For (i), note that any service duration S_i which has completed service by time T' under case 4 will correspond to an arrival time $t'_i = T'$ and residual service duration $R_i^{(4)} = 0$ for case 4': such service durations therefore make no contribution to workload vectors $V_{4'}(T')$ and $V_4(T')$. Next consider any pairs $(T', R_i^{(4)})$ for which $R_i^{(4)} < S_i$ (i.e. customers who arrive by time T' and who have had some, but not all, of their workload served by case 4 before this time). There can clearly be no more than c such customers, and by construction they make the same positive contribution to both workload vectors $V_{4'}(T')$ and $V_4(T')$ (subject to appropriate labelling of servers for case 4'). The only other customers who contribute to $V_4(T')$ (and $V_{4'}(T')$) are those who arrive before time T' but who are yet to start service by this time (i.e. for which $J_i^{(4)} > T'$). These customers correspond to inputs of the form (T', S_i) for case 4'. Consider the first such arrival: in case 4 this customer is placed in queue at its arrival time $t_i < T'$, being allocated to the server with the least residual workload at time t_i . But since all c servers in case 4 must be busy over the entire period $[t_i, T']$ (for if not, the customer arriving at time t_i would necessarily have commenced service by T'), this server still has the least residual workload at time T' . It follows that this customer will be allocated to the same server in case 4' at T' . Arguing inductively along these lines, it is clear that all customers arriving over $[t_i, T']$ are allocated to identical servers in cases 4 and 4', implying that $V_{4'}(T') = V_4(T')$, as required.

For (ii), note that the workload vector $V_{3'}(T')$ is instantiated using the residual workloads at T' from case 1 and then applying an FCFS policy, whereas the vector $V_3(T')$ uses the residual workloads at T' from case 3, again under FCFS (due to the change of service discipline in case 3 at time T' , as described in the statement of the theorem). But since the systems in cases 1 and 3 are identical over the period $[0, T')$, it is clear that $V_{3'}(T') = V_3(T')$.

Finally, claim (iii) follows from Theorem 3.2 applied to the two FCFS systems in cases 3' and 4', which use the same sequence of arrival times (t'_1, t'_2, \dots) and possibly different sequences of service durations. But since $J_i^{(1)} \geq J_i^{(4)}$ for all i , it follows that $R_i^{(3)} \geq R_i^{(4)}$ (see Figure 1 for an illustration), and so service durations for case 3' are at least as big as those for case 4', which provides the required monotonicity.

We close this section with a standard lemma which assures us that actual numbers of customers in the systems also obey the comparisons indicated in Theorems 3.2 and 3.3, so long as the arrival processes agree. (Note that the same is *not* true of total residual workload.)

Lemma 3.1. *Consider two queueing systems, such that arrivals happen at the same time for each system, initiations of service happen earlier in the first than in the second ($\tilde{J}_m \geq J_m$ for all m), and service durations are shorter in the first than in the second when indexed by order*

of initiation of service ($\tilde{S}_m \geq S_m$ for all m). Then numbers in the second system \tilde{X}_t exceed numbers in the first system X_t at any specific time t .

Proof. Let X_t, \tilde{X}_t be numbers in the system at time t . Since $\tilde{J}_m + \tilde{S}_m \geq J_m + S_m$ for any m , the representation (3.2) shows that departures happen later in the second system

$$\tilde{D}_m \geq D_m \quad \text{for any } m.$$

At a given time t we know that the same number $\#\{m : t_m \leq t\}$ of customers have entered each of the systems. However, the above inequality for departures shows that fewer have left the second system than the first. Accordingly,

$$\tilde{X}_t = \#\{m : t_m \leq t\} - \#\{m : \tilde{D}_m \leq t\} \geq X_t = \#\{m : t_m \leq t\} - \#\{m : D_m \leq t\}.$$

4. Simple dominated CFTP for M/G/c

We seek a CFTP algorithm for an M/G/c [FCFS] queueing system. The key step is to find a dominating process which is reversible. Sigman [28] showed how to do this if the system is *super-stable* (arrival rate less than service rate of single server, i.e. $\rho < 1$). Observe that the system is dominated by an M/G/1 [FCFS] queue, note that the workload process of the M/G/1 [FCFS] queue is the same as that of the same queue under processor sharing M/G/1 [PS], exploit the dynamic time reversibility of the M/G/1 [PS] workload process (exchange residual workload for work so far completed on customers in service) to simulate it backwards in time until empty, then use the observed departure times and associated service durations to evolve the Kiefer–Wolfowitz vector for the M/G/c [FCFS] queue forwards until time 0.

In this section we show how to improve on this by lifting the super-stability requirement, leaving only the minimal stability requirement (arrival rate less than total service rates of all servers, i.e. $\rho < c$). The idea is as follows: the results of the previous section show that the system is dominated by an M/G/c [RA]=[M_{λ,c}/G/1 FCFS]^c queueing system with RA allocation policy. Sigman [28] noted that naïve pathwise domination fails; however, we can and will exploit the pathwise domination which holds when service durations are assigned in order of initiation of service. Again we can extend the dominating process backwards in time using the PS representation. The simplest way to construct a dominated CFTP is then to extend backwards in time till the dominating M/G/c [RA] system becomes completely empty, because this allows us to identify service durations with initiations of service in a way which is consistent with further extensions backwards in time. In effect we are exploiting the ‘regenerative atom’ idea noted in [18]. The resulting sequences of arrival times and service durations can then be used to construct a realization of an M/G/c [FCFS] queue that is subordinate to the dominating process. Since this can be extended further back in time, using further emptying times of the dominating process, we have produced the tail end of a ‘simulation from time minus infinity’ which must therefore be in equilibrium at time 0 (for a more mathematical account of this idea, see [16]).

So the steps of the algorithm are as follows.

Algorithm 1. The algorithm description involves some random processes and associated random quantities which are run backwards in time: such quantities are crowned with a hat (for example, \hat{Y} below). We summarise the algorithm in 4 steps.

1. Consider an [M/G/1 PS]^c process \hat{Y} , run backwards in time in statistical equilibrium. Make a draw from $\hat{Y}(0)$, the state of the process at time 0.

2. Simulate the c components of the reversed-time process $(\hat{Y}(\hat{t}) : \hat{t} \geq 0)$ over the range $[0, \hat{\tau}]$, where $\hat{\tau}$ is the smallest reversed time such that all components are empty at $\hat{\tau}$.
3. Use $(\hat{Y}(\hat{t}) : \hat{t} \in [0, \hat{\tau}])$ to construct its (dynamic) time reversal and, thus, to build $(Y(t) : \tau \leq t \leq 0)$, an $M/G/c$ [RA]=[M/G/1 FCFS]^c process (here we set $\tau = -\hat{\tau}$).
4. Use Y to evolve X , an $M/G/c$ [FCFS] process, over $[\tau, 0] = [-\hat{\tau}, 0]$, started in the empty state.

Because of the comparison Theorems 3.2 and 3.3, and Lemma 3.1, we may further extend \hat{Y} forwards in reversed time, and thus Y backwards in time, and use this construction to build further variants of X started in the empty state from any time earlier than τ . Suitably extended back in time, Y dominates all these versions; moreover, agreement of any two variants $X^{(1)}$ and $X^{(2)}$ is enforced at the point when Y visits the empty state subsequent to both of their starting times. The arguments discussed in [16] then show that the common value $X(0)$ of all these variants must be a draw from the statistical equilibrium of the $M/G/c$ [FCFS] queue under consideration.

We now discuss in turn the details of each of these steps.

Step 1: generating a draw from the PS queue system in equilibrium. Thanks to the Pollaczek–Khintchine formula for an $M/G/1$ queue, we know that the equilibrium distribution for the residual workload of \hat{Y}_j at time 0 (where $j \in \{1, \dots, c\}$ denotes the j th server) is distributed as

$$\hat{Q}_j(0) = \sum_{i=1}^c \hat{S}_{j,i}^e(0).$$

Here $\hat{S}_{j,i}^e(0)$ are independent and identically distributed draws from the distribution of service durations in equilibrium, with distribution function

$$G_e(x) = \mu \int_0^x \bar{G}(y) dy \quad \text{for } x \geq 0$$

(for $\bar{G}(y) = 1 - G(y)$ the complementary distribution function of a service duration), while $\hat{Q}_j(0)$ is an independent random variable with geometric distribution given by

$$\mathbb{P}[\hat{Q}_j(0) = n] = \left(\frac{\rho}{c}\right)^n \left(1 - \frac{\rho}{c}\right) \quad \text{for } n \geq 0.$$

However, we need to know the *total* (not only residual) workload brought by each of the customers currently in service. Arguing as in [28], or using a dynamic reversibility argument, we do this by simulating from the stationary *spread* distribution for each of the $\hat{Q}_j(0)$ customers being served by server j at time 0, giving draws $H_{j,1}, H_{j,2}, \dots, H_{j,\hat{Q}_j(0)}$: these represent the total workload brought by each customer. Here the stationary spread distribution of service durations is the length-biased variant of G , with complementary distribution function given by

$$\bar{G}_s(x) = \mu x \bar{G}(x) + \bar{G}_e(x) \quad \text{for } x \geq 0.$$

(Our assumption that $\mathbb{E}[S^2] < \infty$ guarantees that the spread distribution has finite mean.) We then draw independent uniform $[0, 1]$ random variables $U_{j,i}$ and set $\hat{S}_{j,i}^e(0) = U_{j,i} H_{j,i}$ to represent the residual workloads at time 0.

Finally, since all of the servers in \hat{Y} work independently of each other, c independent draws from this distribution deliver an equilibrium draw from \hat{Y} . Set

$$\hat{Y}_j(0) = R(S_{j,1}^e(0), \dots, S_{j,\hat{Q}_j(0)}^e(0))$$

to be this draw from equilibrium, viewed as a list of workloads $S_{j,i}^e(0)$ listed in increasing order (R being the reordering operator mentioned in Section 2).

Step 2: evolving the PS queue system in reverse time till it empties. We record the [M/G/1 PS]^c queueing system as follows: at (reversed) time \hat{t} , the system is defined by

$$\begin{aligned} \hat{Q}_j(\hat{t}) &= \text{the number of customers for server } j \text{ at time } \hat{t}, \\ S_{j,i}^e(\hat{t}) &= \text{the residual workload of customer } i \text{ for server } j \text{ at time } \hat{t}, \\ \hat{Y}_j(\hat{t}) &= R(S_{j,1}^e(\hat{t}), \dots, S_{j,\hat{Q}_j(\hat{t})}^e(\hat{t})), \quad \hat{Y}(\hat{t}) = (\hat{Y}_1(\hat{t}), \hat{Y}_2(\hat{t}), \dots, \hat{Y}_c(\hat{t})). \end{aligned}$$

It is convenient to write $|\hat{Y}_j(\hat{t})| = \hat{Q}_j(\hat{t})$ and $|\hat{Y}(\hat{t})| = |\hat{Y}_1(\hat{t})| + |\hat{Y}_2(\hat{t})| + \dots + |\hat{Y}_c(\hat{t})|$.

If $|\hat{Y}(0)| = 0$ (so there is no residual workload left in the system at all) then set $\hat{\tau} = 0$ and stop simulating.

Otherwise, use event-based simulation. Calculate the next event time after \hat{t} as follows. For each $j \in \{1, 2, \dots, c\}$, all of the \hat{Q}_j customers of server j are served simultaneously by server j at rate $1/|\hat{Y}_j(\hat{t})|$ until either one of the customers of one of the servers has been completely served (and then leaves the system) or a new customer arrives (at rate λ) to be served by one or another of the servers. Reset \hat{t} accordingly.

- If the event is an *arrival*, then generate a new service duration S for the customer (using distribution G) and choose a server $j \in \{1, \dots, c\}$ to which the customer is allocated. The customer is placed in service, so increment \hat{Q}_j by $+1$ (so that the per-customer service rate of server j drops accordingly).
- If the event is a *departure*, then record the *departure time* and the *full service duration* of the departing customer; increment \hat{Q}_j by -1 (so that the per-customer service rate of server j increases accordingly). If $|\hat{Y}(\hat{t})| = 0$ then set $\hat{\tau} = \hat{t}$ and stop simulating.

If $\hat{\tau} > 0$ then record the departure times of customers as $0 \leq \hat{t}_1 \leq \hat{t}_2 \leq \dots \leq \hat{t}_k = \hat{\tau}$, and record the associated (full) service durations as S_1, \dots, S_k .

Step 3: dynamic time-reversal and construction of the M/G/c [RA]=[M/G/1 FCFS]^c dominating process. Let the M/G/c [RA]=[M/G/1 FCFS]^c system Y start from the empty state at time $\tau = -\hat{\tau}$ and run forward in time. We let $|Y(t)|$ denote the total number of customers in Y at time t . Arrivals occur at times $\tau = -\hat{t}_k \leq -\hat{t}_{k-1} \leq \dots \leq -\hat{t}_1$. The customer arriving at time $-\hat{t}_i$ has associated service duration S_i (obtained from records kept as specified in step 2 above), and is allocated to the same server that completed service S_i in \hat{Y} . Reorder the set of service durations according to the corresponding *initiation of service durations* in the forwards queueing system Y . Denote this ordered list by $\mathcal{S}' = (S'_1, \dots, S'_k)$: if J_i^Y is the time of initiation of service S'_i in Y then $\tau = J_1^Y \leq J_2^Y \leq \dots \leq J_k^Y$.

Note that it is possible for J_i^Y to be positive, in the case where Y has customers in the queue at the (terminal!) time 0 who have yet to commence service. In the event that $J_k^Y > 0$, we extend the simulation of Y further into the future by drawing extra (independent) arrival times over the period $(0, J_k^Y]$, along with associated service durations. This results in a set of additional

service durations with associated times of initiation of service: these extra services are then added to the list \mathcal{S}' in order of these times of initiation of service. (Note that this implies a potential change in index for service durations S'_i for which $J_i^Y > 0$.)

This gives us a method of constructing a stationary version of Y started arbitrarily far back into the past and run until the time when all customers in the system at time 0 have commenced service. For example, our simulation of \hat{Y} can be extended to the second time $\hat{\tau}' > \hat{\tau}$ of emptying, and then these additional departure times and service durations used to feed Y over the corresponding period of forward time $[\tau', \tau)$. Since the workload of the M/G/1 queue is invariant under changes of work-conserving discipline, it follows that if Y starts from empty at τ' then it will again be empty just before the arrival at time τ .

Step 4: construction of the target process M/G/c [FCFS] process. Start the M/G/c [FCFS] queue X from the empty state at time τ , and let it evolve (using (2.1)) by generating arrivals at times $\tau = -\hat{t}_k \leq -\hat{t}_{k-1} \leq \dots \leq -\hat{t}_1$ (i.e. the same arrival times as used for Y), but with service duration S'_i now allocated to the arrival $-\hat{t}_i$. Since customers are served by X in order of arrival, this means that service durations are once again allocated by time of initiation of service, i.e. $J_1^X \leq J_2^X \leq \dots$. The domination arguments of Section 3 permit us to argue that $J_i^X \leq J_i^Y$ for $i = 1, \dots, k$, and so X satisfies $|X(t)| \leq |Y(t)|$ for all $t \in [\tau, 0]$, where $|X(t)|$ denotes the total number of customers in X at time t . (Note, however, that it is certainly *not* the case that the residual workload in $X(t)$ is necessarily dominated by that in $Y(t)$.) Return $X(0)$ as a draw from equilibrium of an M/G/c [FCFS] process.

5. Sandwiching for the dominated CFTP algorithm

The algorithm described in Section 4 is inefficient, because it uses the regenerative atom which is the empty system state. For typical applications of $M_\lambda/G/c$ queueing systems, we would expect $1 \ll \rho < c$, so that the system would frequently visit states where no more than c people were in the system, but would only rarely visit the empty state.

A more efficient dominated CFTP algorithm exploits the domination results (Theorems 3.2, 3.3, and Lemma 3.1) to establish sandwiching. The idea is to stop the backward-in-time simulation of the $[M/G/1 PS]^c$ process \hat{Y} at some time \hat{T} well short of the time required to achieve empty state, but then to construct a lower envelope M/G/c [FCFS] process L (started at the empty state at time $-\hat{T}$) and an upper envelope M/G/c [FCFS] process U (started using the state of the forwards dominating M/G/c [RA]=[M/G/1 FCFS]^c process Y at time $-\hat{T}$), and to evolve these using the arrival times and service durations derived from Y in such a way that

- (a) at any given time, the number of people in L lies below the number in U which in turn lies below the number of people in Y ;
- (b) similar envelope processes begun at earlier times sandwich themselves between L and U (the so-called ‘sandwiching property’), in the sense of coordinatewise domination of Kiefer–Wolfowitz workload vectors.

It follows from the theory of dominated CFTP [16], [18] that if we then successively decrease $-\hat{T}$ till eventually $L(0) = U(0)$ then the common state of $L(0) = U(0)$ will be a draw from the equilibrium (this depends crucially on the sandwiching property mentioned above, which must not be neglected in implementation). The delicate issue in all this is exactly the requirement to maintain sandwiching. This requires us to match service durations to times of initiation of service, not just with respect to individual pairs of envelope processes, but also as between a

couple of pairs begun at different times. The trick is to extend the simulation of \hat{Y} beyond \hat{T} so that matching may be carried out in a stable way.

Algorithm 2. 1. Consider an $[M/G/1 PS]^c$ process \hat{Y} , run backwards in time in statistical equilibrium. Make a draw from $\hat{Y}(0)$, the state of the process at time 0.

2. Fix a suitable positive $\hat{T} = -T$. Evolve the queue for server j of \hat{Y} (independently of all other servers) until the first time $\hat{\tau}_j \geq \hat{T}$ that *this server* is empty for $j = 1, \dots, c$.

3. Construct Y_j , an M/G/1 [FCFS] process over the corresponding reversed time interval $[-\hat{\tau}_j, 0]$ for $j = 1, \dots, c$.

4. Produce lists of service durations and arrival times, \mathcal{L}_T^* and \mathcal{L}_T .

5. Construct an upper sandwiching process, $U_{[T,0]}$ over $[T, 0]$.

6. Construct a lower sandwiching process, $L_{[T,0]}$ over $[T, 0]$.

7. Check for coalescence.

We now discuss in turn the details of each of these steps.

Step 1: produce a sample from the stationary distribution of the $[M/G/1 PS]^c$ process \hat{Y} . This is performed exactly as in Algorithm 1.

Step 2: evolve the queue for each server of \hat{Y} independently until empty. Record departure times and associated (full) service durations for each server; simulate the queue served by server j (as in step 2 of Algorithm 1) until time

$$\hat{\tau}_j = \inf\{\hat{t} \geq \hat{T} : |\hat{Y}_j(\hat{t})| = 0\}, \quad j = 1, \dots, c.$$

Step 3: construct Y_j , an M/G/1 [FCFS] process over the corresponding reversed time interval, for $j = 1, \dots, c$. For each server $j \in \{1, 2, \dots, c\}$, we simulate Y_j starting in the empty state at time $\tau_j = -\hat{\tau}_j$, and we feed the simulation with arrival times and associated service durations corresponding to the recorded departures from \hat{Y}_j . If $t^+ > 0$ is the *positive* time by which all customers in Y at time 0 have initiated service then extend each Y_j simulation to cover the range $(\tau_j, t^+]$, as in the detail of step 3 of Algorithm 1. Note that, since the M/G/1 [FCFS] process Y_j starts from empty at time τ_j , the path of Y_j over $[0, t^+]$ will remain unchanged if we decrease the value of $T < 0$. Furthermore, since $\tau_j \leq T$ for $j = 1, \dots, c$, we have in fact established the path of Y , an $[M/G/1 FCFS]^c$ process, over the interval $[T, 0]$.

Step 4: produce lists of service and arrival times, \mathcal{L}_T^ and \mathcal{L}_T .* Form the union of all arrival times observed in each Y_j over the interval $[\tau_j, 0]$, and order them as $t_1 \leq t_2 \leq \dots \leq t_n$. Similarly, form the union of all pairs (J, S) of time J of initiation of service and associated service duration S from each Y_j over the interval $[\tau_j, t^+]$, and order these in increasing order of J . Let $\mathcal{L}_T^* = \{(t_k, J_k, S_k) : t_k \leq T\}$ and $\mathcal{L}_T = \{(t_k, J_k, S_k) : t_k > T\}$. Finally, for each $(t_k, J_k, S_k) \in \mathcal{L}_T^*$, replace the arrival time t_k by T , J_k by $J_k \vee T$, and the service duration S_k by its *residual workload* at time T . That is, replace (t_k, J_k, S_k) by $(T, J_k \vee T, R_k)$, where $R_k = (J_k + S_k) \vee T - (J_k \vee T)$.

Step 5: construct an upper sandwiching process, $U_{[T,0]}$ over $[T, 0]$. We construct an M/G/c [FCFS] process $U_{[T,0]}$ over $[T, 0]$ by starting from empty at time T and feeding it the arrival times and service durations read first from \mathcal{L}_T^* and then from \mathcal{L}_T . The intention here is that U can be seen to be a process which switches from the M/G/1 [RA] queue Y to an M/G/c [FCFS]

queue at time T : Theorem 3.3 guarantees that $|U|$ (the number of customers in the upper process) will be dominated by $|Y|$.

Step 6: construct a lower sandwiching process, $L_{[T,0]}$ over $[T, 0]$. In a similar manner we construct an M/G/c [FCFS] process $L_{[T,0]}$ over $[T, 0]$ by starting from empty at time T and feeding it the arrival times and service durations read once again from \mathcal{L}_T^* and then from \mathcal{L}_T , but now with all the service durations in \mathcal{L}_T^* set to 0. Theorems 3.2 and 3.3 ensure that $L_{[T,0]} \leq U_{[T,0]}$ (here \leq denotes coordinatewise domination of the Kiefer–Wolfowitz workload vectors).

Step 7: check for coalescence. If the residual workload vectors of $U_{[T,0]}(0)$ and $L_{[T,0]}(0)$ agree, return their common value as a draw from equilibrium of our target process X . If not, then replace \hat{T} by $\hat{T}' > \hat{T}$, and return to step 2: extend the paths of \hat{Y}_j until they have each emptied at some time $\hat{\tau}_j' \geq \hat{T}'$ etc., and continue as before.

The reader may be concerned that coalescence here occurs when the residual workload vectors first coincide, apparently without requiring equality of numbers of customers in the system. However, under FCFS, a disparity of numbers together with equality of residual workload vectors would require at least one of the two systems to have strictly more than c customers in the system. As already remarked after the proof of Theorem 3.2, coalescence for Algorithm 2 can only occur when both processes have idle servers, and in this case equality of residual workload vectors implies equality of the numbers of customers in the system.

Since $L_{[T,0]}$ is a version of our target M/G/c [FCFS] process started from empty, a standard dominated CFTP argument [16], [18] shows that the above algorithm really does return a perfect draw from the correct equilibrium distribution, as long as the upper and lower processes really do satisfy the ‘sandwiching property’. The following theorem establishes a rigorous validation of sandwiching for Algorithm 2.

Theorem 5.1. *Let $L_{[T,0]}$ and $L_{[T',0]}$ (respectively $U_{[T,0]}$ and $U_{[T',0]}$) be lower (respectively upper) sandwiching processes, defined as above, started at times $T' < T < 0$. Then, for all times $t \in [T, 0]$,*

$$L_{[T,0]}(t) \leq L_{[T',0]}(t) \leq U_{[T',0]}(t) \leq U_{[T,0]}(t),$$

where \leq once again denotes coordinatewise domination of the Kiefer–Wolfowitz workload vectors.

Proof. Let $\hat{\tau}_j$ (respectively $\hat{\tau}_j'$) be the first time after $\hat{T} = -T$ (respectively $\hat{T}' = -T'$) that \hat{Y}_j empties for $j = 1, \dots, c$. As noted at the end of step 3, when we extend the simulation of \hat{Y}_j from $[0, \hat{\tau}_j]$ to $[0, \hat{\tau}_j']$ its path over $[0, \hat{\tau}_j]$ is unchanged. It follows that the list \mathcal{L}_T (created in step 4 above) is unchanged by such an extension, i.e.

$$\mathcal{L}_T = \{(t, J, S) \in \mathcal{L}_{T'} : t > T\}.$$

Furthermore, any additional entries created in \mathcal{L}_T^* when extending from $\hat{\tau}_j$ to $\hat{\tau}_j'$ (which must satisfy $t_k \leq T$) have zero residual service durations at time T . Extending the simulation of \hat{Y}_j ($j = 1, \dots, c$) from $[0, \hat{\tau}_j]$ to $[0, \hat{\tau}_j']$ thus has no effect on the paths of $U_{[T,0]}$ and $L_{[T,0]}$. We may, therefore, assume that the lists \mathcal{L}_T^* , \mathcal{L}_T , $\mathcal{L}_{T'}^*$, and $\mathcal{L}_{T'}$ are all constructed by running each \hat{Y}_j process over the longer intervals $[0, \hat{\tau}_j']$.

Now we simply observe that $U_{[T',0]}$ is (as remarked in step 5 above) a process which switches from the M/G/1 [RA] queue Y to an M/G/c [FCFS] queue at time T' , whereas $U_{[T,0]}$ switches from Y to M/G/c [FCFS] at a later time $T > T'$: Theorem 3.3 shows that $U_{[T,0]}$ must therefore

dominate $U_{[T',0]}$. Similarly, $L_{[T,0]}$ and $L_{[T',0]}$ are two M/G/c [FCFS] processes, which can be viewed as both starting from empty at time T' and with all service durations corresponding to pre- T arrival times set to 0 for $L_{[T,0]}$: Theorem 3.2 shows that $L_{[T',0]}$ dominates $L_{[T,0]}$, as required.

Bearing in mind the remark made at the end of Theorem 3.2, we see that coalescence of the two sandwiching processes can occur only when both upper and lower sandwiching processes have strictly fewer than c individuals in the system. There follows an almost obvious remark: in the $c = 1$ case Algorithm 2 offers no advantage over Algorithm 1, which itself reduces to the $c = 1$ case of [28]. However, if $c > 1$ then it is possible for Algorithm 2 to produce coalescence when started prior to the latest time (prior to time 0) at which the equilibrium queue has an idle server. For large c it follows that Algorithm 2 offers substantial practical advantages in terms of reduced run-time.

6. Assessment of algorithms for the M/M/c case

So far we have introduced, and proved the correctness of, two algorithms for perfectly sampling from the stationary distribution of the Kiefer–Wolfowitz workload vector for stable M/G/c queues. In this section we briefly investigate and compare the performance of these algorithms, mainly in the special case when service durations are exponentially distributed (i.e. for an M/M/c queue). We begin with a discussion of choice of back-off strategy for Algorithm 2, and then present some simulation results which indicate that this algorithm may be substantially faster than (the rather naïve) Algorithm 1. These observations are reinforced by theoretical bounds on the run-time of the two algorithms, which can be found in Section 6.3. We do not present here a complete analysis of our algorithms' performance, but we do elucidate the relative efficiency of Algorithm 2.

6.1. Back-off strategies

In Algorithm 2 it is necessary to specify a method for choosing the sequence of times $\{\hat{T}, \hat{T}', \dots\}$ at which to check for coalescence. We briefly discuss two options. The first is to use the well-known 'binary back-off' method (set $\hat{T} = 1$, $\hat{T}' = 2$, and continue to double in this way for as long as necessary), as is employed in many CFTP algorithms. The second is to use a sequence of stopping times determined by the dominating \hat{Y}_j processes. For $j = 1, \dots, c$ let $\hat{\tau}_j = \inf\{t > 0: |\hat{Y}_j| = 0\}$, and let $\hat{\tau}^-$ and $\hat{\tau}^+$ be the minimum and maximum of these times; suppose that server j^- is the one that empties for the first time at $\hat{\tau}^-$. The first time at which we can possibly check for coalescence is when running U and L over $[-\hat{\tau}^-, 0]$. If this does not lead to coalescence then the path of \hat{Y}_{j^-} needs to be extended until it empties once again, at which point we update the values of $\hat{\tau}^-$, $\hat{\tau}^+$, and j^- and repeat. However, since server j^- is starting from empty at time $\hat{\tau}^-$, it is quite likely to empty again after only a relatively short period of time, and it may therefore be computationally expensive to check for coalescence as soon as this server is once again empty. In what follows we make use of a binary back-off strategy whenever making use of Algorithm 2.

6.2. Example of simulation output

Both of our two algorithms produce a perfect sample from the stationary distribution of the Kiefer–Wolfowitz workload vector. In Figure 2 we show the result of using Algorithm 2 for an M/G/c queue with $\lambda = c = 25$ and service distributions following a uniform[0, 1] distribution; here we have chosen to display the last six coordinates of the workload vector (for which, recall, the coordinates are ordered monotonically by remaining workload).

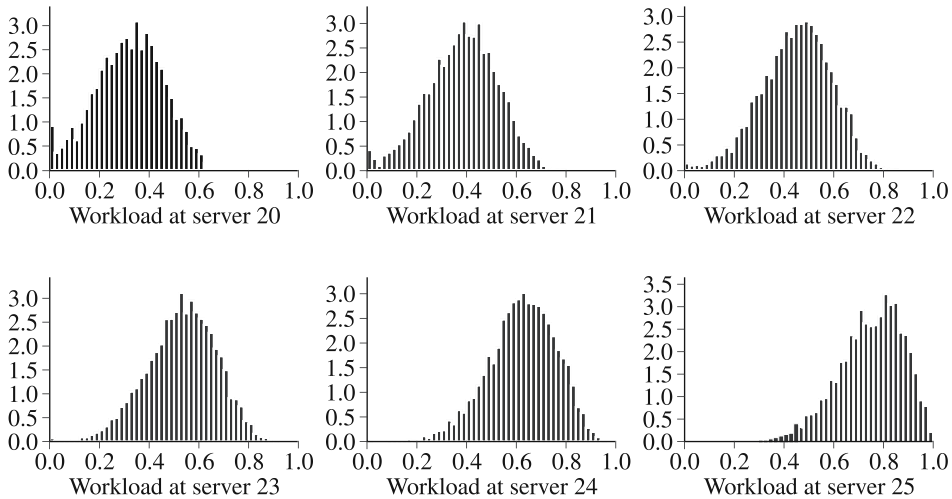


FIGURE 2: Equilibrium distribution of the final six coordinates of the Kiefer–Wolfowitz workload vector when $\lambda = c = 25$ and service durations are uniformly distributed on $[0, 1]$. (Produced from 5000 draws using Algorithm 2.)

When service durations follow an exponential distribution (i.e. X is an $M/M/c$ queue) there is a well-known closed form for the distribution of the number of customers in the system under stationarity:

$$\pi_k = \left(\frac{\rho}{c}\right)^k \frac{c^{(k \wedge c)}}{(k \wedge c)!} \pi_0 \quad \text{for } k \geq 0.$$

We have compared the theoretical distribution to the empirical distribution obtained by output from large numbers of runs of Algorithm 2 for a wide variety of different sets of parameter values and achieved good agreement: by way of illustration, the result of doing this when $\lambda = 10$, $\mu = 2$, and $c = 10$ is shown in Figure 3. Note that these parameters clearly satisfy $1 < \rho = \lambda/\mu < c$, and so this is an example of a stable, but not super-stable, queue. A chi-squared test between the theoretical and observed distributions here gave a p -value of 0.62, indicating good agreement.

It is also of interest to compare how far one has to simulate the dominating process \hat{Y} for each algorithm, and we have performed such a comparison for a variety of sets of parameter values. In Figure 4 we give an indication of how much quicker it may be to detect coalescence via Algorithm 2 rather than simply waiting for \hat{Y} to empty (as in Algorithm 1). For this example we once again set $\lambda = 10$, $\mu = 2$, and $c = 10$, and performed 5000 runs of each algorithm. For Algorithm 1 we recorded the value of $\hat{\tau}$ (the time taken for \hat{Y} to empty), while for Algorithm 2 we employed a binary back-off approach (as is common in many CFTP algorithms) and recorded the minimum value of \hat{T} needed to determine coalescence of our upper and lower sandwiching processes. Note that the binary back-off approach means that it is possible for Algorithm 2 to take longer than Algorithm 1 to detect coalescence (e.g. if \hat{Y} empties at time $0 < \hat{\tau} < 1$ then Algorithm 2 will not detect this until $\hat{T} = 1$; similar phenomena

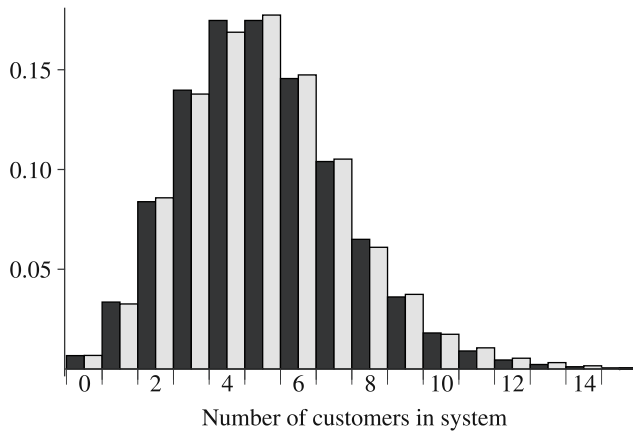


FIGURE 3: Number of customers for an M/M/c queue in equilibrium when $\lambda = 10$, $\mu = 2$, and $c = 10$. Solid bars show the theoretical number of customers in the system; shaded bars show the result of 5000 draws using Algorithm 2.

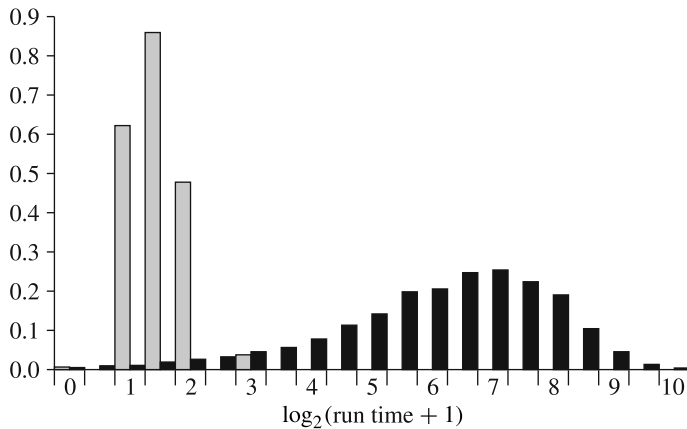


FIGURE 4: Distribution of time taken for coalescence to be detected under Algorithms 1 and 2 applied to an M/M/c queue for 5000 runs with $\lambda = 10$, $\mu = 2$, and $c = 10$. Solid bars show the distribution of $\log_2(\hat{\tau} + 1)$ for Algorithm 1, where $\hat{\tau}$ is the first time at which \hat{Y} empties; shaded bars show the distribution of $\log_2(\hat{T} + 1)$ for Algorithm 2, where \hat{T} is the smallest time needed to detect coalescence using binary back-off.

arise in several perfect simulation algorithms involving binary back-off) but that, in general, Algorithm 2 is significantly faster. In Figures 5(a) and 5(b) we show similar run-time results for Algorithm 2 using substantially larger values of λ and c (while maintaining $\rho = \lambda/2$). The coalescence time $\hat{\tau}$ here clearly does not increase significantly: in the following section we give an heuristic argument which explains why this is to be expected, at least when service times are exponentially distributed.

Of course, such a comparison does not take into account the additional computational demands of checking for coalescence (usually repeatedly) in Algorithm 2, nor the fact that some of the servers in the $[M/G/1]^c$ process may not empty until a time which is significantly greater

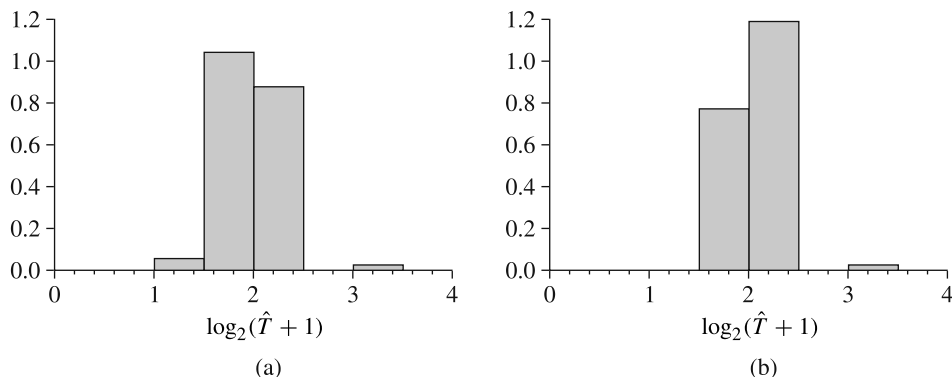


FIGURE 5: Sample of run-time distributions for a selection of M/M/c queues. (a) Coalescence time (measured as in Figure 4) for Algorithm 2 applied to an M/M/c queue for 5000 runs with $\lambda = 30$, $\mu = 2$, and $c = 30$. Chi-squared test on the distribution of the number of customers in the system gave a p -value of 0.84. (b) Coalescence time (measured as in Figure 4) for Algorithm 2 applied to an M/M/c queue for 5000 runs with $\lambda = 50$, $\mu = 2$, and $c = 50$. Chi-squared test on the distribution of the number of customers in the system gave a p -value of 0.38.

than \hat{T} (especially when ρ is close to 1), and so this is by no means a complete discussion of the relative efficiency of each algorithm. However, it emphasises just how much sooner it is possible for coalescence to be detected, without the need to wait for the dominating process to empty completely. (Note that the computational demands of Algorithms 1 and 2 may be compared as follows: running from an emptying time $\tau < 0$, Algorithm 1 requires simulation of an $[M/G/1]^c$ and an $M/G/c$ [FCFS]. Running from a time $T < 0$, Algorithm 2 requires simulation of an $[M/G/1]^c$ and *two* coupled $M/G/c$ [FCFS]. Bearing this in mind, and exploiting the remark after the proof of Theorem 3.2, the choice between Algorithms 1 and 2 should depend on an heuristic comparison of the first moments of the emptying time τ and the latest time (prior to 0) at which the equilibrium queue $M/G/c$ [FCFS] has an idle server.)

6.3. Notes on convergence rates

The run-time of Algorithm 1 is equal to the time taken for the $[M/G/1 \text{ PS}]^c$ dominating process \hat{Y} to empty. It is well known [1, Theorem 5.7] that in equilibrium the mean time for the $M/G/1$ queue to empty is finite if and only if the service duration distribution has a finite second moment. Using our standing assumption that $\mathbb{E}[S^2] < \infty$, it follows that each server in \hat{Y} will almost surely empty in finite time, and that the time taken until we see a simultaneous empty period for all c servers (that is, the time taken for the dominating process to completely empty) will itself be finite. Thus, Algorithm 1 has finite mean run-time if and only if $\mathbb{E}[S^2] < \infty$. The same observation holds for Algorithm 2, for which a finite run-time is also dependent upon each server in \hat{Y} emptying in finite time.

Stronger conditions on the moments of S allow a better bound on run-times. For example, note that the time taken for \hat{Y} to completely empty has an exponential moment (that is, \hat{Y} is geometrically ergodic) if and only if the $M_{\lambda/c}/G/1$ queue (and indeed the $M/G/c$ [FCFS] target process) is geometrically ergodic; this is equivalent to S itself having a finite exponential moment [21, Theorem 16.4.1]. More general conditions for existence of moments of the stationary waiting time for a $GI/GI/c$ queue have recently been determined in [10].

Existence of an exponential moment is rather a strong demand, but bounds on algorithm

run-time can still be produced under weaker drift conditions. In particular, if $\mathbb{E}[S^m] < \infty$ for some $m \geq 2$ then Hou and Liu [13] showed the embedded M/G/1 queue (and, hence, the embedded M/G/c queue) to be polynomially ergodic, which in turn implies that the run-time of Algorithms 1 and 2 will possess a finite m th moment. (A similar result is true in a much wider context: Connor and Kendall [8] described a generic (but impractical) perfect simulation algorithm for a class of so-called *tame* chains. This extends the work in [15] to chains which satisfy a geometric Foster–Lyapunov drift condition at a *state-dependent subsampling time*; see [7] for more details. In particular, if the M/G/c queue is tame, with $\mathbb{E}[S^m] < \infty$ for some $m \geq 2$, then [7, Proposition 4.3] can be used to show that the dominating process of [8] for the M/G/c queue is polynomially ergodic.)

The mean run-time behaviour of Algorithm 1 can be estimated using simple renewal-theoretic arguments. First, consider an M/G/1 queue with service duration distributed as the random variable S , and with arrival intensity λ/c . By Pollaczek–Khintchine theory, in statistical equilibrium the probability of this being empty is $1 - \rho/c$, where $\rho = \lambda\mathbb{E}[S]$ [1, Theorem 5.2]. Consequently, the $[M/G/1]^c$ dominating process, used in both Algorithms 1 and 2, has probability $(1 - \rho/c)^c$ of being completely empty at a given time.

Second, consider the start- and end-times of the busy periods of the whole system $[M/G/1]^c$. These form an alternating renewal process: completely empty periods have exponential(λ) durations, while busy periods are distributed as a random variable B_1 , being the time it takes for $[M/G/1]^c$ to empty completely if it starts off with just one new customer. Alternating renewal theory allows us to deduce that

$$\frac{1/\lambda}{1/\lambda + \mathbb{E}[B_1]} = \left(1 - \frac{\rho}{c}\right)^c.$$

Now let B_e be the time till the queue empties, if it is started in equilibrium.

A stochastic comparison argument shows that B_e stochastically dominates B_1 except when the equilibrium queue is empty (probability $(1 - \rho/c)^c$). Consequently, we may deduce that

$$\begin{aligned} \mathbb{E}[B_e] &\geq \left(1 - \left(1 - \frac{\rho}{c}\right)^c\right) \mathbb{E}[B_1] \\ &= \left(1 - \left(1 - \frac{\rho}{c}\right)^c\right) \frac{(1 - \rho/c)^{-c} - 1}{\lambda} \\ &\geq \frac{(1 - \rho/c)^{-c} - 2}{\lambda}. \end{aligned}$$

This carries over to a lower bound on the run-time of Algorithm 1, which is given by the complete emptying time of an $[M/G/1 \text{ PS}]^c$ system.

It is instructive to consider specific cases. In Table 1 we show how the lower bound increases quickly with the arrival rate when $\lambda = c = 2\rho$, and indicates the effect of increasing ρ when $\lambda = c$ is held constant. The lower bound when $\lambda = c = 10$ and $\rho = 5$ is comparable to the simulation results for Algorithm 1 displayed in Figure 4, for which the mean run-time was 143. Furthermore, this analysis indicates that Algorithm 1 is infeasible for large c , even when ρ/c is significantly smaller than 1.

A corresponding analysis for Algorithm 2 is more intransigent, as one has to estimate the mean coupling time of upper- and lower-processes which are coupled M/G/c, coupled by having the same arrival processes and obtaining service durations in the same sequential order from a fixed sequence. We are not yet able to give a useful analysis of this coupling, which would

TABLE 1: Lower bound on the mean run time for Algorithm 1 ($\mathbb{E}[B_e]$) for some specific queue parameters.

λ	c	ρ	Lower bound on $\mathbb{E}[B_e]$	λ	c	ρ	Lower bound on $\mathbb{E}[B_e]$
10	10	5	102	30	30	5.0	7.85
20	20	10	52 429	30	30	10.0	6392
30	30	15	3.58×10^7	30	30	20.0	6.86×10^{12}
40	40	20	2.75×10^{10}	30	30	25.0	7.37×10^{21}
50	50	25	2.25×10^{13}	30	30	29.5	7.37×10^{51}

involve consideration of the coupled Kiefer–Wolfowitz workload vectors. Instead we offer an heuristic argument, working instead with the Markov processes given by numbers of customers in the system for two coupled M/M/c queues with the same stable parameters. These queues X and Y are defined as follows: X is begun at X_0 , a draw from the stationary distribution; Y is begun at $Y_0 = 0$. Both X and Y use the same Poisson stream of arrivals. Departures from X and Y are coupled so that $X \geq Y$ always: any departure from Y always coincides with a departure from X . Thus, the continuous-time Markov chains X and Y are *immersion coupled* ([17]; this kind of coupling is also called *Markovian* or *co-adapted*). Their joint transition rates are given by

$$\begin{aligned} X &\rightarrow X + 1, & Y &\rightarrow Y + 1 & \text{at rate } \lambda, \\ X &\rightarrow X - 1, & Y &\rightarrow Y - 1 & \text{at rate } (Y \wedge c)\mu, \\ X &\rightarrow X - 1, & Y &\rightarrow Y & \text{at rate } ((X \wedge c) - (Y \wedge c))\mu. \end{aligned}$$

We wish to consider $\mathbb{E}[T_{\text{couple}}]$, where $T_{\text{couple}} = \inf\{t : X_t = Y_t\}$. Note that $\mathbb{P}[T_{\text{couple}} \leq t] = \mathbb{P}$ [upper- and lower-processes (begun at time $-t$) coalesce by time 0].

To this end, we introduce a further process $Z \geq X \geq Y$ with $Z_0 = X_0$ and coupled to Y as follows:

$$\begin{aligned} Z &\rightarrow Z + 1, & Y &\rightarrow Y + 1 & \text{at rate } \lambda, \\ Z &\rightarrow Z - 1, & Y &\rightarrow Y - 1 & \text{at rate } (Y \wedge c)\mu, \\ Z &\rightarrow Z - 1, & Y &\rightarrow Y & \text{at rate } \mu \text{ when } Y < c \text{ and } Z > Y. \end{aligned}$$

Then $\tilde{T}_{\text{couple}}$ stochastically dominates T_{couple} ; hence, $\mathbb{E}[\tilde{T}_{\text{couple}}] \geq \mathbb{E}[T_{\text{couple}}]$, where

$$\tilde{T}_{\text{couple}} = \inf\{t : Z_t = Y_t\}.$$

To estimate $\mathbb{E}[\tilde{T}_{\text{couple}}]$, it suffices to find positive constants α and β such that

$$U = \alpha(Z - Y) + \beta Y + t$$

is a nonnegative supermartingale up to the coupling time \tilde{T} . For then we can apply the methodology of the proofs of Foster–Lyapunov criteria: $\mathbb{E}[\tilde{T}_{\text{couple}}] \leq \mathbb{E}[U_{\tilde{T}_{\text{couple}}}] \leq \mathbb{E}[U_0] = \alpha \mathbb{E}[Z_0] = \alpha \mathbb{E}[X_0]$, which at last can be computed using detailed balance [1, p. 77]. Accordingly, consider the transition rates of the Markov chain (Z, Y) : using these we may deduce that, before the coupling time,

$$\mathbb{E}[U_{t+\delta t} \mid U_t] - U_t = -\alpha\mu \mathbf{1}_{\{Y_t < c\}}\delta t - \beta(\mu(Y_t \wedge c) - \lambda)\delta t + \delta t + o(\delta t).$$

TABLE 2: Heuristic overestimate of the mean run time for Algorithm 2 for some specific queue parameters.

<i>Heuristic upper bound</i>				<i>Heuristic upper bound</i>			
λ	c	ρ	of mean run-time	λ	c	ρ	of mean run-time
10	10	5	5.04	30	30	5.0	1.00
20	20	10	10.00	30	30	10.0	5.00
30	30	15	15.00	30	30	20.0	40.10
40	40	20	20.00	30	30	25.0	131.25
50	50	25	25.00	30	30	29.5	4853.97

For U to be a supermartingale before the coupling time, it is necessary and sufficient that this expression be nonpositive for small δt . Nonpositivity follows if

$$\beta(c\mu - \lambda) \geq 1 \quad (Y_t \geq c \text{ case}),$$

$$\alpha\mu + \beta(\mu(Y_t \wedge c) - \lambda) \geq \alpha\mu - \beta\lambda \geq 1 \quad (Y_t < c \text{ case}).$$

Using $\rho = \lambda/\mu$, we set

$$\beta = \frac{1}{\lambda} \frac{\rho}{c - \rho}, \quad \alpha = \frac{1}{\lambda} \frac{c\rho}{c - \rho}$$

(note that stability of the M/G/c queues requires $\rho \leq c$), and deduce that

$$\mathbb{E}[T_{\text{couple}}] \leq \frac{1}{\lambda} \frac{c\rho}{c - \rho} \mathbb{E}[X_0]. \tag{6.1}$$

Again, it is instructive to consider specific cases. In Table 2 we tabulate corresponding heuristic over-estimates for Algorithm 2 for the same ranges of parameter values as found in Table 1. Note that $\rho = \lambda\mathbb{E}[S]$. Note too that the large growth in mean run-time at the foot of the second column of Table 2 is an inevitable consequence of heavy-traffic in the dominating $[M/G/1]^c$ queue.

Compare the results on \log_2 run-times displayed in Figure 4 and Figure 5 (for which the mean values of \hat{T} for the results in Figure 4 ($\lambda = c = 10, \rho = 5$), Figure 5(a) ($\lambda = c = 30, \rho = 15$), and Figure 5(b) ($\lambda = c = 50, \rho = 25$) are 2.27, 2.99, and 3.32, respectively). Bearing in mind the demands of binary back-off, this suggests that this heuristic over-estimate is a reasonable indication of the feasibility of Algorithm 2 for substantial values of c .

7. Conclusion

In this paper we have described the construction of two dominated CFTP Algorithms 1 and 2 for a general stable M/G/c [FCFS] queue, and have shown that the algorithms have finite mean run-time when the typical service duration has finite second moment.

The second of these, Algorithm 2, is more complex and requires more delicacy and care in description and in implementation; however, despite this increased complexity, it demonstrates the potential for considerably reduced actual run-times compared with the first, more naïve, algorithm. In particular, Algorithm 2 will be preferable in cases when the M/G/c [FCFS] queue is stable rather than super-stable. This is because Algorithm 1 has a run-time comparable to the time at which the queue first empties; this time may be expected to be large when super-stability fails.

There has been significant contemporary interest in perfect simulation for queueing problems: we have noted the work of [3] for GI/GI/ c/c loss processes, and of [22] on a Brownian model for heavy-traffic situations. A particular motivation for this is provided in [12], where it was established that a crucial measure of queue efficiency (mean waiting time) can be substantially affected by more subtle features than simply the first two moments of service time and arrival rate. In such cases it is of clear value to have access to perfect simulation methods such as these two dominated CFTP algorithms.

We close by mentioning four further questions raised by this work.

1. It is natural to ask whether dominated CFTP methods can be extended to the case of renewal process input; the work of [3] on the same problem in the context of loss processes may be very helpful here. We believe this should be feasible, particularly because the ‘impractical’ dominated CFTP algorithms for polynomially ergodic Markov chains address similar problems [8]. However, we have not considered the details of such an extension.
2. It would be interesting to know whether anything theoretical can be said about the relative merits of the two back-off strategies for Algorithm 2 outlined in Section 6.1.
3. It is natural to ask whether it is possible to implement dominated CFTP for a suitably wide class of queueing networks. Sigman [30] describes several applications to networks of the method in [28] for the super-stable case. It is not clear to us how one would contrive to construct a dominating process for such problems in the stable case, so this question seems to us to be challenging.
4. Finally, one might ask whether in the context of M/G/ c [FCFS] it is possible to carry out dominated CFTP simultaneously for a suitable range of c (the number of servers); what might be described as ‘omnithermal dominated CFTP’, to borrow a term used to describe Grimmett’s coupling of random-cluster processes for all values of a specific parameter [11], and applied to CFTP in [25]. In his doctoral thesis [27], Shah showed how to implement omnithermal dominated CFTP for area-interaction point processes; the issue for stable M/G/ c [FCFS] queues is one of detecting at what stage there is coalescence for all c in the range. It is straightforward to carry out omnithermal dominated CFTP in the case of Algorithm 1 in a relatively efficient manner: once an emptying time τ has been established for the instance with lowest c in range, then this will serve for all higher values of c , using a simple workload domination argument; see, e.g. [23]. (Indeed, the value of τ can be updated to the most recent emptying time for the current value of c after the queue has been simulated for that value.) However, it remains an open question whether one can establish a comparably efficient omnithermal dominated CFTP based on Algorithm 2.

Acknowledgements

SBC was supported by EPSRC (grant no. EP/J009180). WSK was supported in part by EPSRC (grant no. EP/K013939).

References

- [1] ASMUSSEN, S. (2003). *Applied Probability and Queues*, 2nd edn. Springer, New York.
- [2] BESKOS, A. AND ROBERTS, G. O. (2005). Exact simulation of diffusions. *Ann. Appl. Prob.* **15**, 2422–2444.

- [3] BLANCHET, J. AND DONG, J. (2015). Perfect sampling for infinite server and loss systems. *Adv. Appl. Prob.* **47**, 761–786.
- [4] BLANCHET, J. AND WALLWATER, A. (2014). Exact sampling of stationary and time-reversed queues. Preprint. Available at <http://arxiv.org/abs/1403.8117>.
- [5] BOROVKOV, A. A. AND FOSS, S. G. (1992). Stochastically recursive sequences and their generalizations. *Siberian Adv. Math.* **2**, 16–81.
- [6] CHEN, X. AND BLANCHET, J. (2015). Steady-state simulation of reflected Brownian motion and related stochastic networks. To appear in *Ann. Appl. Prob.* **25**, 3209–3250.
- [7] CONNOR, S. B. AND FORT, G. (2009). State-dependent Foster–Lyapunov criteria for subgeometric convergence of Markov chains. *Stoch. Proces. Appl.* **119**, 4176–4193.
- [8] CONNOR, S. B. AND KENDALL, W. S. (2007). Perfect simulation for a class of positive recurrent Markov chains. *Ann. Appl. Prob.* **17**, 781–808, 1808–1810.
- [9] FERRARI, P. A., FERNÁNDEZ, R. AND GARCIA, N. L. (2002). Perfect simulation for interacting point processes, loss networks and Ising models. *Stoch. Process. Appl.* **102**, 63–88.
- [10] FOSS, S. AND KORSHUNOV, D. (2012). On large delays in multi-server queues with heavy tails. *Math. Operat. Res.* **37**, 201–218.
- [11] GRIMMETT, G. (1995). The stochastic random-cluster process and the uniqueness of random-cluster measures. *Ann. Prob.* **23**, 1461–1510.
- [12] GUPTA, V., HARCHOL-BALTER, M., DAI, J. G. AND ZWART, B. (2010). On the inapproximability of M/G/K: why two moments of job size distribution are not enough. *Queueing Systems* **64**, 5–48.
- [13] HOU, Z. AND LIU, Y. (2004). Explicit criteria for several types of ergodicity of the embedded M/G/1 and GI/M/n queues. *J. Appl. Prob.* **41**, 778–790.
- [14] KENDALL, W. S. (1998). Perfect simulation for the area-interaction point process. In *Probability Towards 2000*, Springer, New York, pp. 218–234.
- [15] KENDALL, W. S. (2004). Geometric ergodicity and perfect simulation. *Electron. Commun. Probab.* **9**, 140–151.
- [16] KENDALL, W. (2005). Notes on perfect simulation. In *Markov Chain Monte Carlo*, World Scientific, Hackensack, NJ, pp. 93–146.
- [17] KENDALL, W. S. (2015). Coupling, local times, immersions. *Bernoulli* **21**, 1014–1046.
- [18] KENDALL, W. S. AND MØLLER, J. (2000). Perfect simulation using dominating processes on ordered spaces, with application to locally stable point processes. *Adv. Appl. Prob.* **32**, 844–865.
- [19] KIEFER, J. AND WOLFOWITZ, J. (1955). On the theory of queues with many servers. *Trans. Amer. Math. Soc.* **78**, 1–18.
- [20] LOYNES, R. M. (1962). The stability of a queue with non-independent interarrival and service times. *Proc. Camb. Phil. Soc.* **58**, 497–520.
- [21] MEYN, S. AND TWEEDIE, R. L. (2009). *Markov Chains and Stochastic Stability*, 2nd edn. Cambridge University Press.
- [22] MOUSAVI, M. AND GLYNN, P. W. (2013). Exact simulation of non-stationary reflected Brownian motion. Preprint. Available at <http://arxiv.org/abs/1312.6456>.
- [23] MOYAL, P. (2013). A pathwise comparison result for parallel queues. Preprint. Available at <http://arxiv.org/abs/1301.6364>.
- [24] MURDOCH, D. J. AND TAKAHARA, G. (2006). Perfect sampling for queues and network models. *ACM Trans. Model. Comput. Simul.* **16**, 76–92.
- [25] PROPP, J. G. AND WILSON, D. B. (1996). Exact sampling with coupled Markov chains and applications to statistical mechanics. *Random Structures Algorithms* **9**, 223–252.
- [26] ROSS, S. M. (1996). *Stochastic Processes*, 2nd edn. John Wiley, New York.
- [27] SHAH, S. R. (2005). Perfection in conditioning and coverage. Doctoral thesis, Department of Statistics, University of Warwick.
- [28] SIGMAN, K. (2011). Exact simulation of the stationary distribution of the FIFO M/G/c queue. In *New Frontiers in Applied Probability* (J. Appl. Prob. Spec. Vol. **48A**), Applied Probability Trust, Sheffield, pp. 209–213.
- [29] SIGMAN, K. (2012). Exact simulation of the stationary distribution of the FIFO M/G/c queue: the general case for $\rho < c$. *Queueing Systems* **70**, 37–43.
- [30] SIGMAN, K. (2013). Using the M/G/1 queue under processor sharing for exact simulation of queues. *Ann. Operat. Res.* Available at <http://dx.doi.org/10.1007/s10479-013-1408-2>.
- [31] WOLFF, R. W. (1977). An upper bound for multi-channel queues. *J. Appl. Prob.* **14**, 884–888.
- [32] WOLFF, R. W. (1987). Upper bounds on work in system for multichannel queues. *J. Appl. Prob.* **24**, 547–551.