

Perfect Sampling for Infinite Server and Loss Systems

Jose Blanchet and Jing Dong

December 17, 2013

Abstract

We present the first class of perfect sampling (also known as exact simulation) algorithms for the steady-state distribution of non-Markovian loss networks. We use a variation of Dominated Coupling From The Past for which we simulate a stationary infinite server queue backwards in time and analyze the running time in heavy traffic. In particular, we are able to simulate stationary renewal marked point processes in unbounded regions. We use the infinite server queue as an upper bound process to simulate loss systems. The running time analysis of our perfect sampling algorithm for loss systems is performed in the Quality-Driven (QD) and the Quality-and-Efficiency-Driven regimes. In both cases, we show that our algorithm achieves sub-exponential complexity as both the number of servers and the arrival rate increase. Moreover, in the QD regime, our algorithm achieves a nearly optimal rate of convergence.

1 Introduction

We present the first class of exact simulation algorithms for the steady-state distribution of non-Markovian loss networks. The running time of our algorithms is analyzed in the context of many server systems in heavy-traffic; corresponding both to the so-called Quality-Driven (QD) regime, and the Quality-and-Efficiency-Driven (QED, also known as Halfin-Whitt) regime. In both cases, we show that our algorithm achieves sub-exponential complexity as the number of servers and the arrival rate increase. Moreover, in the QD regime, our algorithm achieves a nearly optimal rate of convergence. So, more broadly, our contributions are the first to provide exact simulation methodology with satisfactory running time analysis in the setting of many server queues in heavy traffic.

Exact simulation consists in sampling without any bias from the steady-state distribution of a given ergodic process. Since the inception of Coupling From The Past (CFTP), the most common exact sampling protocol, proposed in the ground breaking paper by Propp and Wilson [22], perfect sampling (also known as “exact simulation”) has become an important part of stochastic simulation. The majority of the available exact simulation algorithms for queues involve exponential distributional assumptions (on service times and/or interarrival times) and very few of such algorithms are applicable in the context of queueing networks. None of them, up to date, have been designed and analyzed in the setting of many server systems in heavy-traffic.

Foss and Tweedie [15] proved that CFTP can be applied if and only if the underlying process is uniformly geometric ergodic. Murdoch and Takahara [20] applied CFTP in the context of queueing models, but mostly with bounded state space. For instance, they consider loss queues with renewal arrivals but with bounded service times and in this case, CFTP can be easily implemented. A variation of CFTP, called Dominated CFTP (DCFTP)[17], allows one to apply CFTP-type idea to obtain unbiased samples from the steady-state distribution of ergodic processes without requiring

uniform ergodicity. A nice summary of DCFTP is given in [14]. The idea is to construct a stationary process which suitably dominates the process of interest and that can be simulated backwards in time from a stationary state at time zero. Then, a suitable lower bound process, coupled with the upper bound, must also be simulated in stationarity and backwards in time. A typical application of DCFTP involves the construction of the upper and lower bound up to a time in the past when they both meet. Then one says that the coalescence occurs. The process of interest is reconstructed forward in time from the coalescence position up to time zero, using the same input sequence that was used to simulate the upper and lower bounds. The state of the process of interest at time zero must then follow the corresponding steady-state distribution.

The paper [11] is one of the earliest to consider DCFTP in the setting of geometrically ergodic Harris recurrent Markov chains. General DCFTP algorithms have been developed more recently in [18] and [10] for Harris recurrent chains, although there are important practical limitations as outlined on p.788 in [10]. In particular, their algorithm assumes that one has analytical access to the transition kernel of the underlying Markov chain after several transitions. A recent paper by Sigman [23] provides an implementable DCFTP algorithm for multi-server queues with Poisson arrivals, but the algorithm requires rather strong conditions on stability; in [24] the conditions are relaxed (also in the setting of Poisson arrivals), using a regenerative technique but the expected termination time of the algorithm is infinite.

In connection to loss queueing systems, we have already mentioned [20], in the setting of queues with non-Poisson input, but bounded service times. In [9], the authors develop a class of CFTP algorithms which combines aggregation and multiple bounding chains, but exponential service times and interarrival times are required in their development. The papers [17], [19] and [13] are close in spirit to the main ideas of our paper as we take a point process approach to the problem. However, their approach requires the use of spatial birth and death processes (generally of poisson type) as the dominating processes and as pointed out in Section 8 of [3], the algorithms appear to significantly increase in complexity as the arrival rate increases.

We provide a practical simulation procedure that works under the assumption of renewal arrivals (having a finite moment generating function) and service time distribution with finite mean (although in our running time analysis in heavy traffic we impose additional moment conditions for service times, but we still are able to cover distributions such as log-normal, which have been observed to accurately fit service time distributions in many server applications [8]). The performance of our procedures has been successfully tested numerically in [5].

In order to implement our strategy in the setting of loss queues, we simulate a stationary infinite server queue backwards in time as our dominating process. A small variation from the standard DCFTP protocol just explained is that we use the upper bound process itself to detect coalescence, thereby bypassing the need for a lower bound process and improving the running time of the algorithm. Basically we detect coalescence over a time interval in which all customers initially present in the infinite server system leave and no loss of customers occurs.

We summarize our contributions next:

- 1) The design and analysis of the first exact sampling algorithm for the infinite server queue whose running time is shown to be basically linear in the arrival rate and thus optimal as the steady state of the infinite server queue, encoding the remaining service time of each customer, requires on average a vector which grows linearly in the arrival rate. (See Theorem 1.)
- 2) The design and analysis of the first exact sampling algorithm for loss networks under non-Markovian arrivals and a heavy-traffic environment. In the QD regime, where service utiliza-

tion is strictly less than 100%, we show that our algorithm has near optimal (linear in the arrival rate) running time. In the QED regime, when the traffic utilization converges to 100% at a square root speed as a function of the arrival rate, we show that our algorithm has a subexponential running time. (See Theorems 2 & 3).

We point out that our algorithms allow to simulate stationary renewal processes with independent and identically distributed (i.i.d.) marks in the positive line on unbounded regions (having finitely many points almost surely). This connection has been noted in [5] with a fixed region (as opposed to a moving frame going backwards in time as we include here) and without the running time analysis that we perform in this paper for high arrival rates.

The rest of the paper is organized as follows. Section 2 contains several subsections. There we introduce our notation and describe the general strategy to simulate the dominating infinite server process. Then, we describe how to detect the coalescence time using only the dominating infinite server process. Once the basic procedures and notation have been explained we proceed to give precise running time results. The whole description initially concentrates on a single station and in the end we provide the extension to the case of loss networks. In Section 3 we provide the details required to implement our general strategy outlined in Section 2 for the infinite server queue in steady state. In Section 4 we study the running time of our algorithms, some technical results in the development of this section are given in Appendix.

2 Basic strategy and main results

In this section we introduce the basic strategy to simulate the systems. We also present some results about the efficiency of our algorithms. We leave the details of the algorithms and proofs of the results to subsequent sections. We start with the strategy of a many-server loss station in steady state and then generalize our strategy to cover loss networks. The method we use is a variation of DCFTP and we use the infinite server queue as the dominating process.

To facilitate our explanation, we start with a formal description of the state of the infinite server (GI/GI/ ∞) queue.

2.1 Description of the GI/GI/ ∞ system

We first introduce some notations and assumptions next. Let $N = \{N(t) : t \in (-\infty, 0]\}$ be a one sided time stationary renewal point process. We write $\{A_n : n \geq 1\}$ for the times at which the process N jumps counting backwards in time from time zero with $A_{n+1} < A_n < 0$. Furthermore, we define $X_n = |A_{n+1} - A_n|$. Now let $\{V_n : n \geq 1\}$ be a sequence of i.i.d. random variables (r.v.'s) which are independent of the process N . Define $Z_n = (A_n, V_n)$ and consider the marked point process $\mathcal{M} = \{Z_n : n \geq 1\} \in \mathbb{R}^2$ which we call the ‘‘arriving customer stream’’. More specifically, we consider customers arriving to the system according to a renewal process with i.i.d. interarrival times X_n 's. Independent of the arrival process, their service requirements V_n 's are also i.i.d.

Figure 1 elaborates on the point process description of the infinite server queue and is important for describing our simulation strategy. In Figure 1, the point $Z_n = (A_n, V_n)$ denotes the n -th customer (counting backward in time), whose arrival time is A_n and service requirement is V_n , $n = 1, \dots, 4$. One important feature of infinite server queue is that every customer starts service immediately upon arrival (there is no queue). If we project Z_n to the horizontal axis by drawing a -45° line. The intersection of this line with the horizontal axis is the departure time of such n -th customer. We follow the technical tradition that an arrival at time t is counted in the system at

time t (closed circle) while a departure at time t is not counted (open circle). We can also draw a vertical line at any $t \in \mathbb{R}$. The height of the intersection of the -45° lines emanating from the points Z_n with $A_n \leq t$ and such vertical line, if positive, represents the corresponding remaining service time of that customer at time t .




Figure 1: Point process description of an infinite server queue

We write $G(\cdot) = P(X_n \leq \cdot)$ for the cumulative distribution function (CDF) of X_n and put $\bar{G}(\cdot) = 1 - G(\cdot)$ for its tail CDF. Similarly, we write $F(\cdot) = P(V_n \leq \cdot)$ as the CDF of V_n and $\bar{F}(\cdot) = 1 - F(\cdot)$ as its tail CDF.

The following assumption is imposed throughout our discussion:

Assumption 1. $EX_n < \infty$ and $EV_n < \infty$.

We next introduce a Markovian description of the system. Let $Q(t, y)$ denote the number of people in the system at time t with residual service time strictly greater than y . Notice that for fixed t , $Q(t, \cdot)$ is a piecewise constant step function. If we denote $\{r_1(t), \dots, r_m(t)\}$ as the ordered (positive) remaining service times of customers in the system at time t . Then $Q(t, 0) = m$ and $Q(t, y) = \sum_{i=1}^m I(r_i(t) > y)$. We also let $E(t)$ denote the time elapsed since the previous arrival at time t (i.e. $E(t) = t - \max\{A_n : A_n \leq t\}$) and $W(t) = (E(t), Q(t, \cdot)) \in \mathbb{R}^+ \times \mathcal{D}[0, \infty)$. Then $\{W(t) : t \in \mathbb{R}\}$ forms a Markov process which describes the infinite server queue.

Similarly, we denote $W^L(t) = (E^L(t), Q^L(t, \cdot)) \in \mathbb{R}^+ \times \mathcal{D}[0, \infty)$ as the state of the loss queue at time t , where $E^L(t)$ denote the time elapsed since the previous arrival, and $Q^L(t, y)$ counts the number of people in the loss system at time with residual service time strictly greater than y .

Assume both systems start empty from the infinite past, and we use the same stream of customers to update the infinite server queue and the loss queue with C servers. Then the two systems can be coupled naturally: just label the servers in the infinite server queue, assign customers to the empty server with the smallest label, and by tracking only the state of the first C servers in the infinite server queue one automatically tracks the state of the loss queue. Based on the above description, we have $E(t) = E^L(t)$ and $Q^L(t, y) \leq Q(t, y)$ for any $t \in \mathbb{R}$ and $y \geq 0$. Thus

$$W^L(t) \leq W(t)$$

We next define the coalescence time.

Definition 1. *Coalescence time is a time $T < 0$ at which the state of the loss system is completely identified from the coupled infinite server system, i.e. $W^L(t) = W(t)$.*

2.2 Coalescence time with an $GI/GI/C/C$ queue

As discussed earlier the infinite server system imposes an upper bound on the loss system. A natural way to construct the coalescence (or coupling) time would be to define the coalescence time as the first time (going backwards in time) the infinite server queue empties (assuming, say, unbounded interarrival time distribution, this will occur). However, this coalescence time generally grows exponentially with the arrival rate [16]. So, to detect the coalescence in a more efficient way, we consider the following construction. Let $R(t)$ denote the maximum remaining service time among all customers in the system at time t . And consider a random time $\tau < 0$ satisfying

- 1) $R(\tau) < |\tau|$;
- 2) $\inf_{\tau \leq t \leq \tau + R(\tau)} \{C - Q(t, 0)\} \geq 0$.

As we will show in Section 4.2, τ is well defined and our coalescence time is $T := \tau + R(\tau)$. In simple words, since the infinite server queue has less than s customers on $[\tau, \tau + R(\tau)]$, the loss queue is also operating below capacity C on that interval. Everyone who was present at time τ in the infinite server queue will have left at time $\tau + R(\tau)$. Thus the infinite server queue and the loss queue must have the same set of customers present in the system by that time. From then on we can recover the state of the loss queue at time zero using the same stream of customers as for the infinite server queue on $[\tau + R(\tau), 0]$.

2.3 Basic strategy and main results for the $GI/GI/\infty$ system

Simulating the infinite server queue in stationarity and backwards in time is not trivial, so we first need to explain how to do this task. There are two cases to be considered.

Case 1 The interarrival time has finite exponential moment in a neighborhood of the origin. More specifically, define $\psi(\theta) = \log E \exp(\theta X_n)$. There exists $\theta > 0$ such that $\psi(\theta) < \infty$.

Case 2 The interarrival time does not have finite exponential moment, i.e. it has heavy-tail distribution.

As we shall explain, we can always reduce the second case to the first one by defining yet another coupled upper bound process through truncation. Specifically, define $X_n \wedge b = \min\{X_n, b\}$. We then fix a suitably large constant b and define a coupled infinite server queue with truncated interarrival times: $\{X_n \wedge b : n \geq 1\}$. This truncation essentially speeds up the arrival process. By coupling we mean we use the same stream of customers to update both the original system and the truncated one, i.e., we use (X_n, V_n) to update the original system and $(X_n \wedge b, V_n)$ to update the truncated one. We also define the event times as the arrival time and the departure time of the n th customer, $n \geq 1$ (counting backwards in time). Then the infinite server queue with truncated interarrival times imposes an upper bound, in terms of the number of customers in the system, on the original infinite server queue at the corresponding event times, i.e. $A_n = \sum_{i=1}^n X_i$ corresponds to $A_n(b) := \sum_{i=1}^n (X_i \wedge b)$ in the truncated system, and $A_n + V_n$ corresponds to $A_n^b + V_n$ in the truncated system. Notice that the actual time of the events (such as arrivals and departures) may be different for the two systems because of the truncation. But from the simulation point of view, we simulate the same amount of information to get the corresponding event times in both systems. In what follows, we shall first concentrate our discussion on Case 1 which also includes the infinite

server queue with truncated interarrival times. We then explain how to extend the result to the heavy-tailed case.

We first introduce the procedure to simulate the state of the stationary infinite server queue at time zero. We notice from Figure 2 that customers $Z_n = \{A_n, V_n\}$, with $V_n \leq |A_n|$ will have left the system by time 0. Thus if we can find a random number κ such that

$$V_n \leq |A_n| \text{ for all } n \geq \kappa,$$

then we can simulate the arrival stream backwards in time up to κ (i.e. $\{Z_n : 1 \leq n \leq \kappa\}$) to recover the state of the system at time zero. The challenge here is that κ defined above depends on future customer information, i.e. $\{Z_n : n > \kappa\}$. In what follows, we shall explain the elements behind the simulation of κ .

We write $\mu = EX_n$ and fix an $\epsilon \in (0, \mu)$. Consider any random number κ finite with probability one but large enough such that

$$A_{n+1} \geq n(\mu - \epsilon) \text{ and } V_{n+1} \leq n(\mu - \epsilon) \text{ for all } n \geq \kappa.$$

Let $\kappa(A)$ be a random time satisfying that $A_{n+1} \geq n(\mu - \epsilon)$ for $n \geq \kappa(A)$, and $\kappa(V)$ be a random time satisfying that $V_{n+1} \leq n(\mu - \epsilon)$ for $n \geq \kappa(V)$. Then we can set $\kappa = \max\{\kappa(A), \kappa(V)\}$. The following proposition states that $\kappa < \infty$ almost surely (a.s.). The proof is given in Appendix A.




Figure 2: Coupling time of the infinite server queue

Proposition 1. *Under Assumption 1, the random number κ defined above is finite with probability one.*

As $\{A_n : n \geq 1\}$ and $\{V_n : n \geq 1\}$ are independent of each other, the above construction allows us to sample $\{V_n : n \geq 1\}$ with $\kappa(V)$, and $\{A_n : n \geq 1\}$ with $\kappa(A)$ separately. We next explain the basic sampling strategies.

For the $\{V_n\}$ process, define $J(0) := 0$ and $J(l) = \inf\{n > J(l-1) : V_{n+1} > n(\mu - \epsilon)\}$ for $l = 1, 2, \dots$. It is the times at which V_{n+1} 's exceed the corresponding increasing boundary $n(\mu - \epsilon)$. Let $\gamma = \inf\{l \geq 1 : J(l) = \infty\}$. Then $\kappa(V) = J(\gamma - 1) + 1$. We first simulate $J(l)$'s for $l = 1, 2, \dots, \gamma - 1$, and then simulate the V_n 's conditional on $J(l)$'s (see Section 3.1 for details).

For the $\{A_n\}$ process, define

$$\tilde{S}_n = n(\mu - \epsilon) - (A_{n+1} - A_1) = \sum_{i=1}^n Y_i,$$

where $Y_i = (\mu - \epsilon) - X_{i+1}$. Note that Y_i 's are i.i.d. with $EY_i = -\epsilon$. Set $\tilde{S}_0 = 0$. $\{\tilde{S}_n : n \geq 0\}$ is a random walk with negative drift. $A_{n+1} = A_1 - \tilde{S}_n + n(\mu - \epsilon)$. If we can simulate some random time κ^* such that $\tilde{S}_n \leq 0$ for $n \geq \kappa^*$, then $|A_{n+1} - A_1| \leq n(\mu - \epsilon)$ for $n \geq \kappa^*$. Fix any $m > 0$. Define $\Gamma(0) = 0$ and $\Delta(l) = \inf\{n \geq \Gamma(l-1) : \tilde{S}_n \leq -m\}$, $\Gamma(l) = \inf\{n \geq \Delta(l) : \tilde{S}_n - \tilde{S}_{\Delta(l)} \geq m\}$. Let $\alpha = \inf\{l \geq 1 : \Gamma(l) = \infty\}$. We notice that \tilde{S}_n will never go above 0 from $\Delta(\alpha)$ on; which implies that we can set $\kappa(A) = \Delta(\alpha)$. As we assume the moment generating function of X_n is finite in a neighborhood of the origin, the moment generating function of Y_n is also finite in a neighborhood of the origin. We simulate \tilde{S}_n 's jointly with $\Delta(l)$'s and $\Gamma(l)$'s until α using exponential tilting and the acceptance rejection method (see Section 3.2 for details).

For the heavy-tailed case (Case 2), we can choose the truncation parameter b such that $E[X_n \wedge b] = \int_0^b \bar{G}(x)dx = \mu - 1/2\epsilon$. This is doable because we assume $EX_n = \int_0^\infty \bar{G}(x)dx < \infty$. Set $\epsilon' = 1/2\epsilon$. Then $E[X_n \wedge b] - \epsilon' = \mu - \epsilon$. Denote $A_n(b)$ as the backwards renewal times of the truncated arrival process and let $\kappa(A(b))$ be a random time satisfying that $|A_{n+1}(b)| \geq n(E[X_n \wedge b] - \epsilon')$ for $n \geq \kappa(A(b))$. Then we have $|A_{n+1}| \geq |A_{n+1}(b)| \geq n(\mu - \epsilon)$ for $n \geq \kappa(A(b))$, thus we can set $\kappa(A) = \kappa(A(b))$.

Our algorithm works only under the mild condition in Assumption 1. But we do impose stronger conditions on the service time distribution to rigorously show good algorithmic performance, especially in heavy traffic (i.e. as the arrival rate increases).

We consider a sequence of systems indexed by $s \in \mathbb{N}^+$. We shall say that s is the scale of the system. We speed up the arrival rate of the s -th system by scale s . That is, the interarrival times of the s -th system are given by $X_n^{(s)} = X_n/s$. We keep the service time distribution fixed for all systems, i.e. the service times do not scale with s . The following theorem summarizes the performance of the procedure we proposed for simulating stationary infinite server queue.

Theorem 1. *Assume $E[X_n] < \infty$, and*

(1) *if $EV_n^q < \infty$ for some $q > 2$, then*

$$E_\pi^s \kappa = O(s^{q/(q-1)});$$

(2) *if we further assume $E[\exp(\theta V_n)] < \infty$ for some $\theta > 0$, then*

$$E_\pi^s \kappa = O(s \log s).$$

We prove it by establishing two bounds for $\kappa(A)$ and $\kappa(V)$ respectively. The details is given in Section 4.1.

We next extend the procedure to simulate states of the stationary infinite server system backwards in time for time intervals of any specified length. The construction is very similar to the single time point (i.e. time zero) case explained above.

Define $\kappa_0 := 1$. We consider a sequence of random times κ_j , $j = 1, 2, \dots$, finite with probability one but large enough such that

$$|A_n - A_{\kappa_{j-1}}| \geq (n - \kappa_{j-1})(\mu - \epsilon) \text{ and } V_n \leq (n - \kappa_{j-1})(\mu - \epsilon) \text{ for all } n \geq \kappa_j. \quad (1)$$

Notice that $V_n \leq |A_n - A_{\kappa_{j-1}}|$ for $n \geq \kappa_j$. This implies that a customer who arrives before A_{κ_j} will not be in the system at time $A_{\kappa_{j-1}}$. Thus, using $\{Z_n : 1 \leq n \leq \kappa_j\}$, we can recover the system descriptor $W(t)$ for $t \in [A_{\kappa_{j-1}}, 0]$.

The κ_j 's give us some flexibility to separate the simulation of the two processes. We first simulate the service times and then conditional on the sample path of the service time we simulate the arrival process jointly with κ_j 's.

Define $J_1(0) := 1$ and let

$$\begin{aligned} J_k(l) &= \inf\{n > J_k(l-1) : V_n > (n - J_k(0))(\mu - \epsilon)\}, \\ \gamma_k &= \inf\{l \geq 0 : J_k(l) = \infty\}, \\ J_{k+1}(0) &= J_k(\gamma_k - 1) \end{aligned}$$

for $k = 1, 2, \dots$ and $l = 1, 2, \dots, \gamma_k$.

We first simulate the random time: $J_k(l)$'s for $k = 1, 2, \dots$ and $l = 1, 2, \dots, \gamma_k$, and then simulate $\{V_n : n \geq 1\}$ conditional on $J_k(l)$'s; see Algorithm I in Section 3.1 for details.

Given the sample path of $\{V_n : n \geq 1\}$ and $J_k(l)$'s, we next simulate $\{A_n : n \geq 1\}$ and κ_j 's. This is done by simulating the negative-drift random walk \tilde{S}_n jointly with its running time maximum. Define $\Delta_1(0) := 0$ and $\Gamma_1(0) := 0$. Fix $m > 0$ and let

$$\begin{aligned} \Delta_j(l) &= \inf\{n \geq \Gamma_j(l-1) : \tilde{S}_n - \tilde{S}_{\Delta_j(0)} \leq -m\}, \\ \Gamma_j(l) &= \inf\{n \geq \Delta_j(l) : \tilde{S}_n - \tilde{S}_{\Delta_j(l)} \geq m\}, \\ \alpha_j &= \inf\{l \geq 1 : \Gamma_j(l) = \infty\}, \\ \kappa_j &= \min\{J_k(0) : J_k(0) \geq \Delta_j(\alpha_j) + 1\}, \\ \Delta_{j+1}(0) &= \kappa_j - 1, \\ \Gamma_{j+1}(0) &= \Delta_{j+1}(0) \end{aligned}$$

for $j = 1, 2, \dots$ and $l = 1, 2, \dots, \alpha_j$.

Notice that the process \tilde{S}_n will never go above $\tilde{S}_{\Delta_j(0)}$ from $\Delta_j(\alpha_j)$ on. This implies that $|A_n - A_{\kappa_{j-1}}| \geq (n - \kappa_{j-1})(\mu - \epsilon)$ for $n \geq \kappa_j$. Under the light-tail assumption (Case 1), we simulate the random times $\Delta_j(l)$ and $\Gamma_j(l)$ for $j = 1, 2, \dots, l = 1, 2, \dots, \alpha_j$ and $\{\tilde{S}_n : n \geq 0\}$ by the exponential tilting and acceptance-rejection method. The details are explained in Algorithm II in Section 3.2.

For the heavy-tailed case (Case 2), we again simulate the infinite server queue with truncated interarrival times first. We carefully choose the truncation parameter b and ϵ' such $E[X_n \wedge b] - \epsilon'$ coincides with $\mu - \epsilon$. Then the $\kappa_j(b)$'s we constructed for the truncated system must automatically satisfy the conditions characterizing κ_j 's in (1) for the original system as well.

2.4 Basic strategy and main results for the $GI/GI/C/C$ system

Once we simulate the customer streams backwards in time and construct the states of the dominating stationary infinite server queue accordingly, we can check and find the coalescence time $T = \tau + R(\tau)$ where τ is defined in Section 2.2 backwards in time. Use the state of the infinite server queue at time T as the state of the many-server loss queue at the same time and go forwards in time using the same stream of customers to construct the state of the loss queue up to time 0.

Like in the infinite server queue case, we again consider a sequence of systems indexed by $s \in \mathbb{N}^+$ where the arrival rate of the s -th system is scaled by s and the service rate is kept fixed. Let $\rho = E[V_n]/E[X_n]$ (the ratio of the mean service time and mean interarrival time of the base system). We analyze the system in two heavy-traffic asymptotic regimes. One is the quality driven (QD) regime where $\rho < 1$ and the number of servers in the s -th system, C_s , is s . The other is the quality and efficiency driven (QED) regime where $\rho = 1$ and the number of servers in the s -th system, C_s , is $s + b\sqrt{s}$ with $b > 0$.

Theorem 2 summarizes the performance of the coalescence time in the QD regime.

Theorem 2. *Assume $EX_n < \infty$ and X_n 's are non-lattice and strictly positive. We also assume that $EV_n^q < \infty$ for any $q > 0$ and the cumulative distribution function (CDF) of V_n is continuous. Then*

$$E_\pi^s \tau = o(s^\delta)$$

for any $\delta > 0$.

Remark 1. *The existence of all moments assumption on the service time distribution covers a range of heavy tailed distributions, such as Weibull and log-normal, which are known to fit well data in applications [8].*

Theorem 3 analyze the performance of the coalescence time in the QED regime.

Theorem 3. *Assume $EX_n^2 < \infty$. We also assume $EV_n^q < \infty$ for any $q > 0$ and the CDF of V_n is continuous. Then for b large enough, we have*

$$\log E_\pi^s \tau = o(s^\delta)$$

for any $\delta > 0$.

The main difficulty in the proof of Theorem 2 and Theorem 3 is that it involves the state of the system on an interval rather than a single point. In Section 4.2, we prove Theorem 2 by using the sample path large deviation results [4] of infinite server queue. For Theorem 3, we prove it by applying Borel-TIS inequality [1] to the diffusion limit process of infinite server queue [21]. The details is also given in Section 4.2.

2.5 Extensions and main results for the loss network

Following the definition in [16], we consider a generalized loss network with J stations, labeled $1, 2, \dots, J$ and suppose that station j comprises C_j servers. We have L possible routes, labeled $1, 2, \dots, L$ and for each route l , a J dimensional routing vector P_l . P_l is consist of 1's and 0's, where $P_l(j) = 1$ means route l requires a server at station j . A routing request l is blocked and thus lost if any station j with $P_l(j) = 1$ is full at the arrival time of the request. Customers requesting route l form a renewal process with i.i.d. interarrival times $\{X_n^{(l)} : n \geq 1\}$. The CDF of $X_n^{(l)}$ is G_l . Independent of the arrival process, the service times $\{V_n^{(l)} : n \geq 1\}$ are also i.i.d. with CDF F_l . We assume that G_l 's and F_l 's satisfy Assumption 1.

Following the same strategy as in the many-server loss queue case, we first couple the loss network with a network of infinite-server stations. Notice that no customer is blocked or lost in the infinite server system, thus it imposes an upper bound on the number of jobs in the loss system. Let $Q_j(t, y)$ denote the number of jobs in the j -th station with remaining service time strictly greater than y at time t . Note that a class l job with remaining service time greater than y in the system will be counted in all $Q_j(t, y)$'s with $P_l(j) = 1$. Let $R_j(t)$ denote the longest remaining service time among all customers in station j at time t . Let $R(t) = \max_{1 \leq j \leq J} \{R_j(t)\}$. Then similar to the many server loss queue, we define a random time τ' satisfying the following conditions:

- 1) $R(\tau') \leq |\tau'|$,
- 2) $\inf_{\tau' \leq t \leq \tau' + R(\tau')} \inf_{1 \leq j \leq J} \{C_j - Q_j(t, 0)\} \geq 0$,
i.e. all links are operating below capacity on the interval $[\tau', \tau' + R(\tau')]$.

At time $\tau' + R(\tau')$, everyone in the network of infinite-server stations will be in the loss network as well. Thus from then on (forwards in time), we can update the loss system using the inputs of the infinite-server system.

In order to simulate the network of infinite-server stations with L types of routing requests, we simulate L independent networks of infinite-server stations; each dealing with a single type of routing request. Then we do a superposition of them. The simulation of each independent network of infinite-server stations are exactly the same as what we have described in Section 2.3, as a type l routing request occupies a server from each station j with $P_l(j) = 1$ simultaneously and for the same amount of time. For the l -th system, let $Z_n^{(l)} = (A_n^{(l)}, V_n^{(l)})$ represent the arrival time and service time of the n -th routing request counting backwards in time and $\kappa^{(l)}$ be a random time satisfying that $V_n^{(l)} \leq |A_n^{(l)}|$ for all $n \geq \kappa^{(l)}$. Then following the procedure described in Section 2.3, we will be able to simulate $\kappa^{(l)}$ as the maximum of two random times associated the arrival process and service time process respectively.

We now consider a sequence of systems indexed by $s \in \mathbb{N}^+$. We speed up the the arrival rate of the s -th system by s , i.e. $X_n^{(l,s)} = X_n^{(l)}/s$, and keep the service rate fixed. The same result as in Theorem 1 will still be holding here. Specifically,

Theorem 4 (Theorem 1'). *Assume $EX_n^{(l)} < \infty$ (C).*

(1) *if $E[(V_n^{(l)})^q] < \infty$ for some $q > 2$, then*

$$E_{\pi}^s \kappa^{(l)} = O(s^{q/(q-1)});$$

(2) *if we further assume $E[\exp(\theta V_n^{(l)})] < \infty$ for some $\theta > 0$, then*

$$E_{\pi}^s \kappa^{(l)} = O(s \log s)$$

for $l = 1, 2, \dots, L$.

The proof of Theorem 4 is the same as that of Theorem 1 except for a few notational changes, thus we shall omit it here.

If we held the number of routing request types, L , fixed, as we shall explain below, similar results as in Theorem 2 and Theorem 3 for the coalescence time will be holding here as well. We again run L independent networks of infinite-server stations as described above. Network l serves routing request of type l only, for $l = 1, 2, \dots, L$. Let $Q^{(l)}(t, 0)$ denote the number of jobs in network l at time t and $R^{(l)}(t)$ denote the maximum remaining service time among all jobs in the network at time t . Then we have $R(t) = \max\{R^{(l)}(t) : 1 \leq l \leq L\}$.

We consider two asymptotic regimes. One is the QD regime where for the base system we have

$$\sum_{l=1}^L \frac{EV_n^{(l)}}{EX_n^{(l)}} P_j(l) < C_j. \quad (2)$$

For the s -th system, the number of servers in the j -th station is $C_j^s = sC_j$ for $j = 1, 2, \dots, J$.

Assign a fixed number H_l to each route l . H_l is well chosen such that $E[V_n^{(l)}]/E[X_n^{(l)}] < H_l$ and $\sum_{l=1}^L H_l P_l(j) \leq C_j$. This is doable because of (2). Let $H_l^s = sH_l$. Define a random time $\bar{\tau}'$ satisfying the following two conditions:

1) $R^{(l)}(\bar{\tau}') \leq |\bar{\tau}'|$ for $l = 1, 2, \dots, L$,

2) $\inf_{\tilde{\tau}' \leq t \leq \tilde{\tau}' + R(\tilde{\tau}')} \{H_l - Q^l(t, 0)\} \geq 0$ for $l = 1, 2, \dots, L$.

Notice that $\tilde{\tau}'$ is an upper bound on τ' . As the number of types of routing request is fixed at L (it does not scale with s), using the construction outlined in Section 4.2.1, we can show that the result in Theorem 2 holds for $\tilde{\tau}'$ as well.

Theorem 5 (Theorem 2'). *Assume $EX_n^{(l)} < \infty$ and $X_n^{(l)}$'s are non-lattice and strictly positive. We also assume $E[(V_n^{(l)})^q] < \infty$ for any $q > 0$ and F_l is continuous. Then*

$$E_{\tilde{\tau}'}^s \tau' = o(s^\delta)$$

for any $\delta > 0$.

The other asymptotic regime is the QED regime where for the base system we have

$$\sum_{l=1}^L \frac{EV_n^{(l)}}{EX_n^{(l)}} P_j(l) = C_j$$

and the number of servers in the j -th station of the s -th system is $C_j^s = sC_j + \beta_j \sqrt{s}$ for $j = 1, 2, \dots, J$

We then let $I_l = E[V_n^{(l)}]/E[X_n^{(l)}]$ and $I_l^s = sI_l + a_l \sqrt{s}$ where a_l 's are well chosen such that $\sum_{l=1}^L a_l P_j(l) \leq \beta_j$.

We define a random time $\tilde{\tau}'$ that satisfies the following two conditions:

- 1) $R^{(l)}(\tilde{\tau}') \leq |\tilde{\tau}'|$ for $l = 1, 2, \dots, L$,
- 2) $\inf_{\tilde{\tau}' \leq t \leq \tilde{\tau}' + R(\tilde{\tau}')} \{I_l - Q^{(l)}(t, 0)\} \geq 0$ for $l = 1, 2, \dots, L$.

As before, $\tilde{\tau}'$ is an upper bound on τ' . It is easy to check using the construction outlined in Section 4.2.2 that the result in Theorem 3 holds for $\tilde{\tau}'$ as well.

Theorem 6 (Theorem 3'). *Assume $E[(X_n^{(l)})^2] < \infty$. We also assume $E[(V_n^{(l)})^q] < \infty$ for any $q > 0$. Then for b_j 's large enough, we have*

$$\log E_{\tilde{\tau}'}^s \tau' = o(s^\delta)$$

for any $\delta > 0$.

We shall omit the proof of Theorem 5 and Theorem 6 as it is the same as the proof of Theorem 2 and Theorem 3 with the introduction $\tilde{\tau}'$ and $\tilde{\tau}'$ except for a few notational changes.

3 Detailed simulation algorithms

In order to provide the details of our simulation algorithms outlined in Section 2.3, we shall first work under the light-tailed case (Case 1) where we assume there exists $\theta > 0$ such that $\psi(\theta) < \infty$. The extension to the heavy-tailed case (Case 2) was introduced in Section 2 and we shall provide more details in Section 3.3.

We further impose the following assumptions on our ability to simulate the service times and interarrival times.

Assumption 2. We assume that $F(\cdot)$ is known and easily accessible either in closed form or via efficient numerical procedures. Moreover, we can simulate V_n conditional on $V_n \in (a, b]$ with $P(V_n \in (a, b]) > 0$.

Assumption 3. Suppose that $G(\cdot)$ is known and that it is possible to simulate from $G_{eq}(\cdot) := \mu^{-1} \int_0^\infty \overline{G}(t) dt$. Moreover, let $G_\theta(\cdot) = E \exp(\theta X_n - \psi(\theta)) I(X_n \leq \cdot)$ be the associated exponentially tilted distribution with parameter θ for $\psi(\theta) < \infty$. We assume that we can simulate from $G_\theta(\cdot)$.

We next introduce our algorithm to simulate $\{V_n : n \geq 1\}$. Conditional on the sample path of $\{V_n : n \geq 1\}$, we then explain how to simulate $\{A_n : n \geq 1\}$ and κ_j 's.

3.1 Simulation of $\{V_n : n \geq 1\}$ and $J_k(l)$'s for $k = 1, 2, \dots, l = 1, 2, \dots, \gamma_k$

We will first introduce the procedure to simulate $J_1(l)$ for $l = 1, 2, \dots, \gamma_1$. Recall that $J_1(0) := 1$. Let $p(n) = P(V_1 > n(\mu - \epsilon))$. Then $P(J_1(l) = \infty | J_1(l-1) = k) = \prod_{n=k+1}^\infty (1 - p(n))$. It involves the evaluation of the product of infinite terms. In Procedure A, we introduce a sandwiching approximation scheme to accomplish that.

The following lemma guarantees the termination of our procedure.

Lemma 1. *If $EV_1 < \infty$, then*

$$P(J_1(1) = \infty) = \prod_{n=1}^\infty (1 - p(n)) \geq \exp\left(-\frac{cEV_1}{\mu - \epsilon}\right) > 0 \quad (3)$$

for some constant c depending on the value of $p(1)$, and consequently

$$E\gamma_1 \leq \exp(cEV_1/(\mu - \epsilon)) < \infty.$$

Proof.

$$\begin{aligned} P(J_1(1) = \infty) &= \prod_{n=1}^\infty (1 - p(n)) \geq \prod_{n=1}^\infty \exp(-cp(n)) \\ &\geq \exp\left(-\frac{c}{\mu - \epsilon} \int_0^\infty P(V_1 > \nu) d\nu\right) = \exp\left(-\frac{cEV_1}{\mu - \epsilon}\right). \end{aligned}$$

For $l = 2, 3, \dots$, conditional on $J_1(l-1) = k$:

$$\begin{aligned} P(J_1(l) = \infty | J_1(l-1) = k) &= \prod_{n=k+1}^\infty (1 - p(n)) \\ &\geq \exp\left(-\frac{c \int_k^\infty P(V_1 > \nu) d\nu}{\mu - \epsilon}\right) \geq \exp\left(-\frac{cEV_1}{\mu - \epsilon}\right), \end{aligned}$$

thus γ_1 is stochastically dominated by a geometric random variable with parameter $p = \exp(-cEV_1/(\mu - \epsilon))$. The result then follows. \square

We next introduce our sandwiching approximation scheme. Notice that

$$\prod_{i=k+1}^h (1 - p(i)) \geq P(J_1(l) = \infty | J_1(l-1) = k) \geq \prod_{i=k+1}^h (1 - p(i)) \times \exp\left(-\frac{2 \int_h^\infty P(V_1 > \nu) d\nu}{\mu - \epsilon}\right) \quad (4)$$

for $h > k$.

Another important observation is that if we let $\prod_{i=k+1}^k (1 - p(i)) = 1$,

$$\prod_{i=k+1}^{h-1} (1 - p(i)) - \prod_{i=k}^h (1 - p(i)) = p(h) \prod_{i=k}^{h-1} (1 - p(i)) = P(J_1(l) = h | J_1(l-1) = k)$$

for $h > k$.

Let

$$u(h) = \exp\left(-\frac{2 \int_h^\infty P(V_1 > \nu) d\nu}{\mu - \epsilon}\right).$$

We now propose the following procedure to simulate the value of $J_1(l)$ conditional on $J_1(l-1) = k$.

Procedure A (Simulate $J_1(l)$ given $J_1(l-1) = k$)

1. Initialize $h = k + 1$, $g = 1 - p(h)$ and $f = gu(h)$. Simulate $U \sim \text{Unif}[0, 1]$
2. While $f < U < g$,
 set $h = h + 1$, $g = g(1 - p(h))$ and $f = gu(h)$
 end while
3. If $U \leq f$, then $J_1(l) = \infty$. Otherwise, $J_1(l) = h$.

The simulation of $J_k(l)$ for $l = 1, 2, \dots, \gamma_k$ follows the same rationale. We let $p_k(n) = P(V_1 > n(\mu - \epsilon) | V_1 \leq (n + J_k(0) - J_{k-1}(0))(\mu - \epsilon))$. Then following the same argument leading to (3) and (4), we have correspondingly

$$P(J_k(1) = \infty) > 0,$$

and

$$\begin{aligned} \prod_{i=n+1}^h (1 - p_k(i)) &\geq P(J_k(l) - J_k(0) = \infty | J_k(l-1) - J_k(0) = n) \\ &\geq \prod_{i=n+1}^h (1 - p_k(i)) \times \exp\left(-\frac{2 \int_h^\infty P(V_1 > \nu | V_1 \leq \nu + (J_k(0) - J_{k-1}(0))(\mu - \epsilon)) d\nu}{\mu - \epsilon}\right) \end{aligned}$$

for $h > n$.

Let

$$u_k(h) = \exp\left(-\frac{2 \int_h^\infty P(V_1 > \nu | V_1 \leq \nu + (J_k(0) - J_{k-1}(0))(\mu - \epsilon)) d\nu}{\mu - \epsilon}\right).$$

We now propose a modification of Procedure A that allows us to simulate $J_k(l)$ conditional on $J_k(l-1) - J_k(0) = n$.

Procedure A1 (Simulate $J_k(l)$ given $J_k(l-1) - J_k(0) = n$)

1. Initialize $h = n + 1$, $g = 1 - p_k(h)$ and $f = gu_k(h)$. Simulate $U \sim \text{Unif}[0, 1]$.
2. While $f < U < g$,
 set $h = h + 1$, $g = g(1 - p_k(h))$ and $f = gu_k(h)$
 end while

3. If $U \leq f$, then $J_k(l) = \infty$. Otherwise, $J_k(l) = J_k(0) + h$.

Based on Procedure A1 and our previous analysis we have:

Algorithm I (Sample V_n 's jointly with $J_k(l)$'s)

- Step 0. Set $J_0(0) = -\infty$, $J_1(0) = 1$, $k = 1$, $l = 1$. Simulate V_1 according to its nominal distribution.
- Step 1. Simulate $J_k(l)$ conditional on the value of $J_k(l-1)$ using Procedure A1.
- Step 2. If $J_k(l) = \infty$, set $\gamma_k = l$, $J_{k+1}(0) = J_k(\gamma_k - 1)$, $k = k + 1$, $l = 1$ and go back to Step 1. Otherwise, go to Step 3.
- Step 3. Simulate V_n for $J_k(l-1) < n < J_k(l)$ by conditioning on $V_n \leq (n - J_k(0))(\mu - \epsilon)$ and simulate $V_{J_k(l)}$ by conditioning on $(J_k(l) - J_k(0))(\mu - \epsilon) < V_{J_k(l)} \leq (J_k(l) - J_{k-1}(0))(\mu - \epsilon)$. Set $l = l + 1$ and go back to Step 1.

When running the above algorithm, we specify K as the number of intervals ($[J_k(0), J_k(\gamma_k - 1)]$) we want to simulate. We then run Algorithm I from $k = 1$ till $k = K$. The program will give us $\{V_n : 1 \leq n \leq J_K(\gamma_K - 1)\}$ and $J_k(l)$'s for $k = 1, 2, \dots, K$, $l = 1, 2, \dots, \gamma_k$.

3.2 Simulation of $\{A_n : n \geq 1\}$ and $\Delta_j(l)$'s, $\Gamma_j(l)$'s for $j = 1, 2, \dots$, $l = 1, 2, \dots, \alpha_j$

Given the sample path of $\{V_n : n \geq 1\}$, we will first explain how to simulate the $\Delta_j(l)$'s and $\Gamma_j(l)$'s sequentially and jointly with the underlying random walk $\{\tilde{S}_n : n \geq 1\}$. We then simulate A_1 according to $G_{eq}(\cdot)$ and set $A_{n+1} = A_1 + n(\epsilon - \mu) - \tilde{S}_n$. The analysis and methodology in this subsection closely follows those in [12] and [7]. The same procedure can be used to simulate a negative drifted random walk, \tilde{S}_n , together with its running time maximum defined as $\max_{k \geq n} \{\tilde{S}_k\}$.

Let $\mathcal{F}_n = \sigma\{Y_1, Y_2, \dots, Y_n\}$, the σ -field generated by the Y_j 's up to time n . Let $\xi \geq 0$ and define

$$T_\xi := \inf\{n \geq 0 : \tilde{S}_n > \xi\}.$$

Then by the strong Markov property we have that for $1 \leq l \leq \alpha_j$,

$$P(\Gamma_j(l) = \infty | \mathcal{F}_{\Delta_j(l)}) = P(\Gamma_j(l) = \infty | \tilde{S}_{\Delta_j(l)}) = P(T_m = \infty) > 0,$$

where we use $P(\cdot)$ to denote the nominal probability measure.

It is important then to notice that

$$P(\alpha_j = k) = P(T_m < \infty)^{k-1} P(T_m = \infty)$$

for $k \geq 1$. In other words, α_j is geometrically distributed. The procedure that we have in mind is to simulate each stage $\Delta_j(\alpha_j)$ in time intervals, and the number of time intervals is precisely α_j .

Let $\psi_Y(\theta) = \log E \exp(\theta Y_i)$ be the log moment generating function of Y_i . As we assume $\psi_X(\theta)$ is finite in a neighborhood of zero, $\psi_Y(\cdot)$ is also finite in a neighborhood of zero. Moreover $EY_i = \psi_Y'(0) = -\epsilon$ and $\text{Var}(Y_i) = \psi_Y''(0) > 0$. Then by the convexity of $\psi_Y(\cdot)$, one can always select $\epsilon > 0$ sufficiently small so that there exists $\eta > 0$ with $\psi_Y(\eta) = 0$ and $\psi_Y'(\eta) \in (0, \infty)$. The root η allows us to define a new measure P_η based on exponential tilting so that

$$\frac{dP_\eta}{dP}(Y_i) = \exp(\eta Y_i).$$

Moreover, under P_η , \tilde{S}_n is random walk with positive drift equal to $\psi'_Y(\eta)$ [2]. Therefore $P_\eta(T_\xi < \infty) = 1$ and

$$q(\xi) := P(T_\xi < \infty) = E_\eta \exp(-\eta \tilde{S}_{T_\xi})$$

for each $\xi \geq 0$. Based on the above analysis we now introduce a convenient representation to simulate a Bernoulli random variable $J(\xi)$ with parameter $q(\xi)$, namely,

$$J(\xi) = I(U \leq \exp(-\eta \tilde{S}_{T_\xi})), \quad (5)$$

where U is a uniform random variable independent of everything else under P_η .

Identity (5) provides the basis for an implementable algorithm to simulate a Bernoulli random variable with success probability $q(\xi)$. Sampling $\{\tilde{S}_1, \dots, \tilde{S}_{T_\xi}\}$ conditional on $T_\xi < \infty$, as we shall explain now, corresponds to basically the same procedure. First, let us write

$$P^*(\cdot) = P(\cdot | T_\xi < \infty).$$

The following result provides an expression for the likelihood ratio between P^* and P_η .

Lemma 2. *We have that*

$$\frac{dP^*}{dP_\eta}(\tilde{S}_1, \dots, \tilde{S}_{T_\xi}) = \frac{\exp(-\eta \tilde{S}_{T_\xi})}{P(T_\xi < \infty)} \leq \frac{\exp(-\eta \xi)}{P(T_\xi < \infty)}.$$

Proof.

$$\begin{aligned} P(\tilde{S}_1 \in H_1, \dots, \tilde{S}_{T_\xi} \in H_{T_\xi} | T_\xi < \infty) &= \frac{P(\tilde{S}_1 \in H_1, \dots, \tilde{S}_{T_\xi} \in H_{T_\xi}, T_\xi < \infty)}{P(T_\xi < \infty)} \\ &= \frac{E_\eta[\exp(-\eta \tilde{S}_{T_\xi}) I(\tilde{S}_1 \in H_0, \dots, \tilde{S}_{T_\xi} \in H_{T_\xi})]}{P(T_\xi < \infty)}. \end{aligned}$$

□

The previous lemma provides the basis for a simple acceptance / rejection procedure to simulate $\{\tilde{S}_1, \dots, \tilde{S}_{T_\xi}\}$ conditional on $T_\xi < \infty$. More precisely, we propose $\{\tilde{S}_1, \dots, \tilde{S}_{T_\xi}\}$ from $P_\eta(\cdot)$. Then one generates a uniform random variable U independent of everything else and accept the proposal if

$$U \leq \frac{P(T_\xi < \infty)}{\exp(-\eta \xi)} \times \frac{dP^*}{dP_\eta}(\tilde{S}_1, \dots, \tilde{S}_{T_\xi}) = \exp(-\eta(\tilde{S}_{T_\xi} - \xi)).$$

This criterion coincides with $J(\xi)$ according to (5). So, the procedure above simultaneously obtains both a Bernoulli r.v. $J(\xi)$ with parameter $q(\xi)$, and the corresponding path $\{\tilde{S}_1, \dots, \tilde{S}_{T_\xi}\}$ conditional on $T_\xi < \infty$ under $P(\cdot)$ if $J(\xi) = 1$.

As $E[Y_i] = -\epsilon < 0$, by strong law of large numbers we have $\Delta_j(l) < \infty$ almost surely for $j = 1, 2, \dots$ and $l = 1, 2, \dots, \alpha_j$. We next define

$$\bar{q}(\xi) := 1 - q(\xi) = P(T_\xi = \infty)$$

and

$$P'(\cdot) = P(\cdot | T_\xi = \infty).$$

The following result provides an expression for the likelihood ratio between P' and P .

Lemma 3. *We have that*

$$\frac{dP'}{dP}(\tilde{S}_1, \dots, \tilde{S}_n) = \frac{I(T_\xi > l)\bar{q}(\xi - \tilde{S}_n)}{P(T_\xi = \infty)} \leq \frac{1}{P(T_\xi = \infty)}.$$

Proof.

$$\begin{aligned} & P(\tilde{S}_1 \in H_1, \dots, \tilde{S}_n \in H_n | T_\xi = \infty) \\ &= \frac{P(\tilde{S}_1 \in H_1, \dots, \tilde{S}_n \in H_n, T_\xi = \infty)}{P(T_\xi = \infty)} \\ &= \frac{E[I(\tilde{S}_1 \in H_1, \dots, \tilde{S}_n \in H_n)I(T_\xi > n)P(T_\xi = \infty | \tilde{S}_1, \dots, \tilde{S}_n)]}{P(T_\xi = \infty)}. \end{aligned}$$

The result then follows from the strong Markov property and homogeneity of the random walk. \square

We are in good shape now to apply acceptance / rejection to sample from P' . The previous lemma indicates that to sample $\{\tilde{S}_1, \dots, \tilde{S}_n\}$ given $T_\xi = \infty$. We can propose from the original (nominal) distribution and accept with probability $\bar{q}(\xi - \tilde{S}_j)$ as long as $\tilde{S}_j \leq \xi$ for all $0 \leq j \leq n$. And in order to perform the acceptance test we need to sample a Bernoulli with parameter $\bar{q}(\xi - \tilde{S}_n)$, but this is easily done using identity (5).

Now consider $0 \leq \xi_1 < \xi_2$, we define

$$P^o(\cdot | T_{\xi_1} < \infty, T_{\xi_2} = \infty).$$

The following result provides an expression for the likelihood ratio between P^o and P_η .

Lemma 4. *We have that*

$$\frac{dP^o}{dP_\eta}(\tilde{S}_1, \dots, \tilde{S}_{T_{\xi_1}}) = \frac{\exp(-\eta\tilde{S}_{T_{\xi_1}})\bar{q}(\xi_2 - \tilde{S}_{T_{\xi_1}})}{P(T_{\xi_1} < \infty, T_{\xi_2} = \infty)} \leq \frac{\exp(-\eta\xi_1)}{P(T_{\xi_1} < \infty, T_{\xi_2} = \infty)}.$$

Proof.

$$\begin{aligned} & P(\tilde{S}_1 \in H_1, \dots, \tilde{S}_{T_{\xi_1}} \in H_{T_{\xi_1}} | T_{\xi_1} < \infty, T_{\xi_2} = \infty) \\ &= \frac{E_\eta[I(\tilde{S}_1 \in H_1, \dots, \tilde{S}_{T_{\xi_1}} \in H_{T_{\xi_1}}) \exp(-\eta\tilde{S}_{T_{\xi_1}})P(T_{\xi_2} = \infty | \tilde{S}_1, \dots, \tilde{S}_{T_{\xi_1}})]}{P(T_{\xi_1} < \infty, T_{\xi_2} = \infty)}. \end{aligned}$$

\square

We again use acceptance/rejection to sample $\{\tilde{S}_1, \dots, \tilde{S}_{T_{\xi_1}}\}$ given $T_{\xi_1} < \infty$ and $T_{\xi_2} = \infty$. We propose $\{\tilde{S}_1, \dots, \tilde{S}_{T_{\xi_1}}\}$ from $P_\eta(\cdot)$. Then we simulate a uniform random variable U independent of all else and accept the proposal if

$$U \leq \frac{P(T_{\xi_1} < \infty, T_{\xi_2} = \infty)}{\exp(-\eta\xi_1)} \times \frac{dP^o}{dP_\eta}(\tilde{S}_1, \dots, \tilde{S}_{T_{\xi_1}}) = \exp(-\eta(\tilde{S}_{T_{\xi_1}} - \xi_1))\bar{q}(\xi_2 - \tilde{S}_{T_{\xi_1}}).$$

Based on the above analysis we propose the following algorithm.

Algorithm II (Given V_n 's and $J_k(l)$'s, sample \tilde{S}_n 's together with $\Delta_j(l)$'s, $\Gamma_j(l)$'s and κ_j 's)

- Step 0. Set $\Delta_1(0) = \Gamma_1(0) = 0$, $\tilde{S}_0 = 0$, $j = 1$, $l = 1$, $\xi = \infty$, $\gamma = -m$. Sample A_1 according to $G_{eq}(\cdot)$.
- Step 1. Simulate S_1, \dots, S_{T_γ} from the original (nominal) distribution.
- Step 2. If $S_i \leq \xi$ for all $1 \leq i \leq T_\gamma$ then sample a Bernoulli $J(\xi - S_{T_\gamma})$ with parameter $q(\xi - S_{T_\gamma})$ using (5) and continue to step 3. Otherwise (i.e. $S_i > \xi$ for some $1 \leq i \leq T_\gamma$) go back to step 1.
- Step 3. If $J(\xi - S_{T_\gamma}) = 1$, go back to step 1. Otherwise $J(\xi - S_{T_\gamma}) = 0$, let $\Delta_j(l) = \Gamma_j(l-1) + T_\gamma$ and $\tilde{S}_{\Gamma_j(l-1)+i} = \tilde{S}_{\Gamma_j(l-1)} + S_i$ for $i = 1, \dots, T_\gamma$. If $j \geq 2$, set $\xi = \tilde{S}_{\Delta_{j-1}(\alpha_{j-1})} + m - \tilde{S}_{\Delta_j(l)}$.
- Step 4. Simulate S_1, \dots, S_{T_m} from $P_\eta(\cdot)$. Sample a Bernoulli $J(\xi - S_{T_m})$ with parameter $q(\xi - S_{T_m})$ using (5) and $U \sim \text{Unif}[0, 1]$. Let $J^* = I(U \leq \exp(-\eta(S_{T_m} - m)) \times (1 - J(\xi - S_{T_m})))$.
- Step 5. If $J^* = 1$, let $\Gamma_j(l) = \Delta_j(l) + T_m$ and $\tilde{S}_{\Delta_j(l)+i} = \tilde{S}_{\Delta_j(l)} + S_i$ for $1 \leq i \leq T_m$. Set $\gamma = \min\{0, \tilde{S}_{\Delta_j(0)} - m - \tilde{S}_{\Gamma_j(l)}\}$. If $j \geq 2$, set $\xi = \tilde{S}_{\Delta_{j-1}(\alpha_{j-1})} + m - S_{\Gamma_j(l)}$. Set $l = l + 1$ and go back to step 1. Otherwise $J^* = 0$, set $\alpha_j = l$, $\kappa_j = \inf\{J_k(0) : J_k(0) \geq \Delta_j(\alpha_j) + 1\}$, $\Delta_{j+1}(0) = \kappa_j - 1$, $\xi = m$ and continue to step 6.
- Step 6. Let $h = \Delta_{j+1}(0) - \Delta_j(\alpha_j)$. Sample S_1, \dots, S_h from the original distribution.
- Step 7. If $S_i \leq \xi$ for all $1 \leq i \leq h$ then sample a Bernoulli $J(\xi - S_h)$ with parameter $q(\xi - S_h)$ using (5) and continue to step 8. Otherwise (i.e. $S_i > \xi$ for some $1 \leq i \leq h$), go back to step 6.
- Step 8. If $J(\xi - S_h) = 1$, go back to step 6. Otherwise $J(\xi - S_h) = 0$, let $\tilde{S}_{\Delta_j(\alpha_j)+i} = \tilde{S}_{\Delta_j(\alpha_j)} + S_i$ for $i = 1, \dots, h$. Set $A_{n+1} = A_1 + n(\epsilon - \mu) - \tilde{S}_n$ for $n = \Delta_j(0) + 1, \dots, \Delta_{j+1}(0)$. Set $j = j + 1$, $l = 1$, $\xi = \tilde{S}_{\Delta_{j-1}(\alpha_{j-1})} + m - \tilde{S}_{\Delta_j(0)}$, $\gamma = -m$ and go back to step 1.

When running the above algorithm, we specify K as the number of intervals $[\kappa_{j-1}, \kappa_j]$ we want to simulate and then repeat the above process from $j = 1$ till $j = K$. The program will give us $\{A_n : 1 \leq n \leq \kappa_K\}$ and $\{\kappa_j : 1 \leq j \leq K\}$.

3.3 Coupled infinite server queue with truncated interarrival times

In this subsection, we provide some additional details for simulating the coupled truncated system together with the original system.

We first explain how to simulate A_1 jointly with $A_1(b)$. The equilibrium distribution of X_n is $G_{eq}(x) = \int_0^x \bar{G}(u) du / E X_n$ and the equilibrium distribution of $X_n \wedge b$ is

$$G_{eq}^b(x) = \frac{\int_0^x \bar{G}(u) du}{E[X_n \wedge b]} I\{x \leq b\}.$$

Thus we simulate A_1 with CDF $G_{eq}(x)$, if $A_1 \leq b$, we set $A_1(b) = A_1$. Otherwise if $A_1 > b$, we keep simulating X_e with CDF $G_{eq}(x)$ until $X_e \leq b$ and set $A_1(b) = X_e$. In particular we have $A_1(b) \leq A_1$.

When simulating $X_n \wedge b$'s from the nominal distribution, we first simulate X_n with CDF $G(\cdot)$ and set $X_n \wedge b = \min\{X_n, b\}$. Denote $Y_n(b) = (E[X_n \wedge b] - \epsilon') - X_n \wedge b$ and let η_b be chosen such that $\log E \exp(\eta_b Y_n(b)) = 0$. When simulating $X_n \wedge b$'s under exponential tilting $P_{\eta_b}(\cdot)$, we first simulate $Y_n(b)$ under $P_{\eta_b}(\cdot)$ and set $X_n \wedge b = (E[X_n \wedge b] - \epsilon') - Y_n(b)$. If $X_n \wedge b < b$, set $X_n = X_n \wedge b$, otherwise ($X_n \wedge b = b$), sample X_n conditional on $X_n \geq b$ under the nominal distribution $P(\cdot)$.

4 Performance analysis

In the previous section, we provide our simulation algorithm and show that our algorithm works in the sense that the termination time is finite with probability one. In this section, we conduct some further asymptotic analysis on the performance of our algorithm. We first analyze the algorithm for the infinite server system and then conduct some analysis on the coalescence time for the many-server loss system.

4.1 Termination time for the infinite server system (Proof of Theorem 1)

Theorem 1 provides the relationship between the moment of the service times and $E_\pi^s \kappa$. We next give a proof of it. We shall omit the subscription π and s when there is no confusion for notational convenience. We first give a proof of the light tailed case. Recall that $\kappa = \max\{\kappa(V), \kappa(A)\}$, where $\kappa(V) = \inf\{k > 1 : V_{n+1} \leq n(\mu - \epsilon)/s \text{ for all } n \geq k\}$ and $\kappa(A) = \inf\{k > 1 : A_{n+1} \geq n(\mu - \epsilon)/s \text{ for all } n \geq k\}$. We prove the theorem by establishing the bounds for $\kappa(V)$ (Lemma 5) and $\kappa(A)$ (Lemma 6) respectively.

Lemma 5. *If $EV_n^q < \infty$ for some $q > 2$, then*

$$E\kappa(V) = O(s^{q/(q-1)}).$$

Proof. Let $p(n) = P(V_1 > n(\mu - \epsilon)/s)$. For k sufficiently large, we have

$$\begin{aligned} P(\kappa(V) > k) &= 1 - \prod_{n=k+1}^{\infty} (1 - p(n)) \\ &\leq 1 - \exp\left(-\frac{2s}{\mu - \epsilon} \int_{k(\mu - \epsilon)/s}^{\infty} P(V > \nu) d\nu\right). \end{aligned}$$

By Chebyshev's inequality

$$P(V_n > \nu) \leq \frac{EV_n^q}{\nu^q}.$$

Let $\delta = 1/(q - 1)$, then for s sufficiently large, we have

$$\begin{aligned} \sum_{k=s^{1+\delta}}^{\infty} P(\kappa(V) > k) &\leq \sum_{k=s^{1+\delta}}^{\infty} \frac{2s}{\mu - \epsilon} \int_{k(\mu - \epsilon)/s}^{\infty} P(V > \nu) d\nu \\ &\leq \frac{2EV_n^q s^q}{(q-1)(q-2)(\mu - \epsilon)^q} \sum_{k=s^{1+\delta}}^{\infty} \frac{1}{k^{q-1}} \\ &= O(s^{q-(1+\delta)(\delta-2)}). \end{aligned}$$

As $q - (1 + \delta)(q - 2) = 1 + \delta$,

$$\begin{aligned} E\kappa(V) &= \sum_{k=0}^{\infty} P(\kappa(V) > k) \\ &= \sum_{k=0}^{s^{1+\delta}-1} P(\kappa(V) > k) + \sum_{k=s^{1+\delta}}^{\infty} P(\kappa(V) > k) \\ &\leq s^{1+\delta} + O(s^{1+\delta}). \end{aligned}$$

□

Notice that when $E \exp(\theta V_n) < \infty$ for some $\theta > 0$,

$$P(V_n > \nu) \leq E \exp(\theta(V_n - \nu)) = E \exp(\theta V_n) \exp(-\theta \nu).$$

Similarly as above, for s sufficiently large we have

$$\sum_{k=\lceil \frac{2}{\theta(\mu-\epsilon)} s \log s \rceil}^{\infty} P(\kappa(V) > k) \leq \frac{2E \exp(\theta V_n)}{(\mu - \epsilon)^2 \theta^2}$$

and

$$E\kappa(V) = \sum_{k=0}^{s \log s - 1} P(\kappa(V) > k) + \sum_{k=s \log s}^{\infty} P(\kappa(V) > k) \leq s \log s + O(1).$$

Thus if $E \exp(\theta V) < \infty$ for some $\theta > 0$, then

$$E\kappa(V) = O(s \log s).$$

Lemma 6. *Assume there exist $\theta > 0$, such that $\psi(\theta) < \infty$, then*

$$E\kappa(A) = O(s).$$

Proof. Based on the algorithm proposed in Section 3.2, we divide the proof into two parts. We first prove that the expected number of iterations is $O(1)$. We then prove that the expected number of steps to reach $-m$ or m is $O(s)$.

Let $T_\xi = \inf\{n \geq 0 : \tilde{S}_n > \xi\}$. Recall that for the base system there exist $\eta > 0$ with $\psi_Y(\eta) = 0$ and $\psi'_Y(\eta) > 0$. And the number of iterations is distributed as a geometric random variable with probability of success $P(T_m = \infty) = 1 - E_\eta \exp(-\eta \tilde{S}_{T_m})$

Then for the s th system with $Y_i^s = Y_i/s$ we have $\tilde{S}_n/s > m$ is equivalent to $\tilde{S}_n > sm$. Thus the number of iterations is a Geometric random variable with probability of success

$$P(T_{sm} = \infty) = 1 - E_\eta \exp(-\eta \tilde{S}_{T_{sm}}) \geq 1 - \exp(-\eta sm).$$

Similarly, let $T'_\xi = \inf\{n \geq 0 : \tilde{S}_n < \xi\}$. Define $M_n = \tilde{S}_n + n\epsilon$, then M_n is a martingale with respect to the filtration generated by $\{Y_1, Y_2, \dots, Y_n\}$. As $EY_i = -\epsilon < 0$, $P(T_{-m} < \infty) = 1$. By the Optional Sampling Theorem, $EM_{T'_{-m}} = E\tilde{S}_{T'_{-m}} + \epsilon ET'_{-m} = 0$. Thus

$$ET'_{-m} = \frac{m}{\epsilon} - \frac{E[m - S_{T'_{-m}}]}{\epsilon}.$$

Then for the s th system we have

$$ET'_{-sm} = \frac{sm}{\epsilon} - \frac{E[sm - S_{T'_{-sm}}]}{\epsilon}.$$

$(sm - S_{T'_{-sm}})$ converges to the ladder hight Y^- distribution as $s \rightarrow \infty$ and $\sup_m E[(sm - S_{T'_{-sm}})^p] < \infty$ yields $E[(Y^-)^p] < \infty$ for $p > 1$ [2]. Therefore,

$$ET'_{-sm} = O(s).$$

□

For the heavy-tailed case, we select the truncation parameter b such that $E[X_n \wedge b] = \mu - 1/2\epsilon$. Then we set $\epsilon' = 1/2\epsilon$ and define $\kappa(A(b))$ as a random time satisfying that $|A_{n+1}| \geq n(E[X_n \wedge b] - \epsilon') = n(\mu - \epsilon)$ for $n \geq \kappa(A(b))$. As $|A_{n+1}| \geq |A_{n+1}(b)|$ under our coupling scheme, we can set $\kappa(A) = \kappa(A(b))$. By Lemma 6, we have $E\kappa(A) = E\kappa(A(b)) = O(s)$. $\kappa(V)$ is defined as before, a random time satisfying that $V_n \leq n(\mu - \epsilon)$ for $n \geq \kappa(V)$. Then $E\kappa(V) = O(s \log s)$ by Lemma 5.

As $\kappa = (\kappa(V), \kappa(A(b)))$, we have $E\kappa = O(s \log s)$. This concludes the proof of Theorem 1.

4.2 Coalescence time for the many-server loss system (Proof of Theorem 2 and Theorem 3)

As we are simulating the process backwards in time, it is natural to define the following filtration

$$\overleftarrow{\mathcal{H}}_t = \sigma\{W(-u) : 0 \leq u \leq t\},$$

for which $\overleftarrow{\mathcal{H}}_u \subset \overleftarrow{\mathcal{H}}_t$ for $0 \leq u \leq t$. τ is a stopping time with respect to $\overleftarrow{\mathcal{H}}_t$. We next try to draw connections between the backward process and some forward process. Define

$$\tau^* = \inf\{t + R(t) : \sup_{t \leq u \leq t+R(t)} \{Q(u, 0)\} < s, t \geq 0\}$$

τ^* is a stopping time with respect to \mathcal{H}_t where $\mathcal{H}_t = \sigma\{M(u) : 0 \leq u \leq t\}$. The stochastic process $\{Q(t, 0) : t \in \mathbb{R}\}$ has a piecewise constant sample path with a finite number of points of discontinuity on any finite length intervals almost surely. Thus for any fixed $T > 0$, we have

$$\begin{aligned} P_\pi(\tau > T) &= P_\pi\left(\bigcap_{-T \leq t \leq 0} (\{R(t) > -t\} \cup (\bigcup_{t \leq u \leq (t+R(t)) \wedge 0} (\{Q(u, 0) > s\})))\right) \\ &= P_\pi\left(\bigcap_{-T \leq t \leq 0} (\{R(T+t) > -t\} \cup (\bigcup_{T+t \leq u \leq (T+t+R(T+t)) \wedge T} \{Q(u, 0) > s\})))\right) \\ &= P_\pi\left(\bigcap_{0 \leq w \leq T} (\{R(w) > T-w\} \cup (\bigcup_{w \leq u \leq (w+R(w)) \wedge T} \{Q(u, 0) > s\})))\right) \\ &= P_\pi(\tau^* > T) \end{aligned}$$

The second equality holds by stationarity; this gives us $E_\pi \tau = E_\pi \tau^*$. Next, we use a special construction similar to that in Section 4 of [6] to prove the results for $E_\pi^s \tau^*$. The idea is to use a geometric trial argument. We divide the time frame into blocks that are roughly independent. And if the process is well-behaved (staying around its measure-valued fluid limit) on one block, then τ^* is reached before the end of that block.

Let $\bar{Q}(t, y)$ denote the number of customers in the infinite server queue that starts empty at time zero with remaining service time greater than y at time $t \geq 0$. For convenience, we also define $\bar{Q}_u(t, y) = \bar{Q}(u+t, y) - \bar{Q}(u, t+y)$ for $u \leq t$, as the number of customers who arrive after u with remaining service time larger than y at time $u+t$.

4.2.1 Proof of Theorem 2.

We first prove the theorem for the light-tailed case. The heavy-tail case proceeds by selecting the truncation parameter c sufficiently large.

For the QD regime, by “well-behaved”, we mean that the process does not deviate δs , for some $\delta > 0$, from its fluid limit. The following lemma states that the probability of not being well-behaved decays exponentially fast with the system scale.

Lemma 7. *Assume $\psi(\theta) < \infty$ for some $\theta > 0$ and X_n 's are non-lattice and strictly positive. We also assume the CDF of V_n is continuous. Then for any $\delta > 0$, there exist $I^*(\delta) > 0$, such that*

$$P(\bar{Q}(t, y) > (1 + \delta)\lambda s \int_y^{t+y} \bar{F}(u) du \text{ for some } t \in [0, 1], y \in [0, \infty)) = \exp(-sI^*(\delta) + o(s)).$$

The proof of Lemma 7 follows from the tow-parameter sample path large deviation result for infinite server queues in [4]. We shall omit it here.

We next introduce our construction of “blocks”. Let $l(s) = \inf\{y : (1 + \delta)s \int_y^\infty \bar{F}(u)du \leq \frac{1}{2}\}$, we define the following sequence of random times Ξ_i 's: $\Xi_0 := 0$. Given Ξ_{i-1} for $i = 1, 2, \dots$, define

$$\begin{aligned} r_i &= \inf\{k : k \geq R(\Xi_{i-1}), k = 1, 2, \dots\}, \\ z &= \inf\{k : k \geq l(s), k = 1, 2, \dots\}, \\ \Xi_i &= \Xi_{i-1} + r_i + z. \end{aligned}$$

We define a Bernoulli random variable ξ_i , with $\xi_i = 1$ if and only if

$$\bar{Q}_{\Xi_{i-1} + (k-1)t_0}(t, y) \leq (1 + \delta)\lambda s \int_y^{t+y} \bar{F}(u)du$$

for all $t \in [0, 1]$, $y \in [0, \infty)$ and every $k = 1, 2, \dots, r_i + z$.

Choose $\delta < 1/\rho - 1$. We first check that $\xi_i = 1$ implies that τ^* is reached before Ξ_i . Since $r_i \geq R(\Xi_{i-1})$, all the customers in the system at time $\Xi_{i-1} + r_i$ will be those who arrive after Ξ_i . Then $\xi_i = 1$ implies that

$$\begin{aligned} Q(\Xi_{i-1} + r_i, y) &\leq \sum_{k=1}^{r_i/t_0} \int_{(k-1)t_0+y}^{kt_0+y} \bar{F}(u)du \\ &= (1 + \delta)\lambda s \int_y^{r_i+y} \bar{F}(u)du \\ &\leq (1 + \delta)\lambda s \int_y^\infty \bar{F}(u)du, \end{aligned}$$

thus $R(\Xi_{i-1} + r_i) \leq l(s)$.

And for every $t \in (k-1, k]$, $k = 1, 2, \dots, z$

$$\begin{aligned} Q(\Xi_{i-1} + r_i + t, y) &\leq (1 + \delta)\lambda s \int_y^{r_i+t+y} \bar{F}(u)du \\ &\leq (1 + \delta)\lambda s \int_y^\infty \bar{F}(u)du, \end{aligned}$$

thus $Q(\Xi_{i-1} + r_i + t, 0) \leq (1 + \delta)\rho s \leq s$ for $t \in [0, R(\Xi_{i-1} + r_i)]$.

Now let $N = \inf\{i \geq 1 : \xi_i = 1\}$, then

$$E\tau^* \leq E \sum_{i=1}^N (r_i + z).$$

We now show a bound for $E \sum_{i=1}^N (r_i + z)$. The proof is given in the Appendix B.

Lemma 8. *Assume $\psi(\theta) < \infty$ for some $\theta > 0$ and $\psi_N(\theta)$ is continuously differentiable throughout \mathbb{R} . We also assume the CDF of V_n is continuous and $EV_n^q < \infty$ for any $q > 0$. Then*

$$E\left[\sum_{i=1}^N (r_i + z)\right] = o(s^\delta)$$

for any $\delta > 0$.

This concludes the proof of the light tailed case. We next extend the theorem to the heavy-tailed case. We prove it by drawing connection to the truncated system. Here we delicately choose the truncation parameter b so that the truncated system still operating the QD regime. More specifically, we choose b such that

$$\int_b^\infty \bar{G}(x)dx < 1/\rho - 1.$$

This can be achieved since $EX_n = \int_0^\infty \bar{G}(x)dx < \infty$. Then for fixed such b we have

$$\rho_b = \frac{E[V_n]}{E[X_n \wedge b]} = \frac{EV_n}{EX_n - \int_b^\infty \bar{G}(x)dx} < 1$$

and

$$E_\pi^s \tau(b) = o(s^\delta)$$

for any $\delta > 0$, where $\tau(b)$ denote the coalescence time in the truncated system.

We next prove by contradiction that the coalescence in the truncated system implies the coalescence in the original system with the same amount of information simulated. Recall that $\tau(b)$ is a random time satisfying that the system has less than s customers at $\tau(b)$. The maximum remaining service time among all customers in the system at time τ is denoted as $R(\tau(b))$. $R(\tau(b)) \leq |\tau(b)|$ and during $R(\tau(b))$ unites of time from $\tau(b)$ on the system always has less than s customers. We can look for $\tau(b)$ at departure times of customers. We assume the process $Q(t, y)$ is right continuous with left limit, so customers departure at time t will not counted in $Q(t, 0)$. Suppose $\tau(b)$ equals to the departure time of the n -th customer. Then every customer arriving between $\tau(b)$ and $\tau(b) + R(\tau(b))$ sees strictly less than s customers (excluding himself) when he enters the system. We set τ equal to the departure time of the n -th customer in the original system and $R(\tau)$ by definition equals to the maximum remaining service time among all customers in the system at time τ . We have $R(\tau) \leq R(\tau(b))$. We claim that every customer arriving between τ and $\tau + R(\tau)$ must see less than s customers (excluding himself) when he enters the system. Suppose this is not the case. Then there exist a customer m , $1 \leq m \leq n$ who arrives between τ and $\tau + R(\tau)$ and finds at least s customers in the system already. The customer with the same index m must have arrived between $\tau(b)$ and $\tau(b) + R(\tau(b))$ in the truncated system and $Q(A_m(b)-) \geq Q(A_m-) \geq s$. We get a contradiction. Therefore, we must have seen the coalescence in the original system as well with the same amount of information simulated.

4.2.2 Proof of Theorem 3.

For QED regime, by “well-behaved”, we mean that the process does not deviate $C\sqrt{s}$, for some $C > 0$, from its fluid limit. The following lemma states that the probability of both being well-behaved and not well-behaved are bounded away from zero.

Lemma 9. *Fix any $\eta > 0$. Let $\nu(y) = (\int_y^\infty \bar{F}(u)du)^{1/(2+\eta)}$. Assume $EX_n^2 < \infty$ and $EV_n^q < \infty$ for any $q > 0$. Then for any large enough C , there exists $\zeta_1(C) > 0$ and $\zeta_2(C) > 0$, such that*

$$P(\bar{Q}(t, y) \leq \lambda s \int_y^{t+y} \bar{F}(u)du + C\sqrt{s}\nu(y) \text{ for all } t \in [0, 1], y \in [0, \infty)) \geq \zeta_1(C) \quad (6)$$

and

$$P(\bar{Q}(t, y) > \lambda s \int_y^{t+y} \bar{F}(u)du + C\sqrt{s}\nu(y) \text{ for some } t \in [0, 1], y \in [0, \infty)) \geq \zeta_2(C). \quad (7)$$

The proof of Lemma 9 follows from the proof of Lemma 9 in [6]. Our case is actually simpler, as we are dealing with a one sided bound (upper bound) only as appose to the two sided bound in [6]. This simplification allows us to remove the light-tail assumption on interarrival time distribution required in [6]. We shall only briefly outline the procedure here.

For Inequality (6), the idea is to consider the diffusion limit of $Q(t, y)$ as a two dimensional Gaussian random field [21], and then invoke Borell-TIS inequality [1].

Assume $EX_n^2 < \infty$, $EV_n < \infty$ and the CDF of V_n is continuous. Pang and Whitt [21] has proved that for the $GI/GI/\infty$ queue with any given initial age $E(0)$,

$$\frac{\bar{Q}(t, y) - \lambda s \int_t^{t+y} \bar{F}(u) du}{\sqrt{s}} \Rightarrow R(t, y) \text{ in } D_{D[0, \infty)}[0, \infty),$$

where $R(t, y) = R_1(t, y) + R_2(t, y)$ is a Gaussian random field with $R_1(t, y) = \lambda \int_0^t \int_0^\infty I(u + x > t + y) dK(u, x)$ and $R_2(t, y) = \lambda c_a^2 \int_0^t \bar{F}(t + y - u) dB(u)$, where $K(u, x) = W(\lambda u, F(x)) - F(x)W(\lambda u, 1)$ in which $W(\cdot, \cdot)$ is a standard Brownian sheet on $[0, \infty) \times [0, 1]$ and $B(\cdot)$ is a standard Brownian motion independent of $W(\cdot, \cdot)$. The constant c_a is coefficient of variation of the interarrival times, i.e. $c_a = \sqrt{\text{Var}(X_n)}/EX_n$. We denote

$$\tilde{R}_i(t, y) = \frac{R_i(t, y)}{v(y)}$$

and define the d-metric (a pseudo-metric)

$$d_i((t, y), (t', y')) = E[(\tilde{R}_1(t, y) - \tilde{R}_2(t', y'))^2]$$

for $i = 1, 2$.

We then invoke the Borell-TIS inequality. We shall skip the verification of the conditions for such invocation here as it is tedious and detailedly proved in [6]. Let $S = [0, 1] \times [0, \infty)$. it is shown in [6] that, there exist constants $M_{i,1} > 0$ and $M_{i,2} > 0$, such that $E[\sup_S \tilde{R}_i(t, y)] \leq M_{i,1} < \infty$ and $\sup_S E[\tilde{R}_i(t, y)^2] \leq M_{i,2} < \infty$. And for $C_i \geq E[\sup_S \tilde{R}_i(t, y)]$,

$$P(\sup_S \tilde{R}_i(t, y) \geq C_i) \leq \exp\left\{-\frac{1}{2 \sup_S E[\tilde{R}_i(t, y)^2]} (C_i - E[\sup_S \tilde{R}_i(t, y)])^2\right\}$$

for $i = 1, 2$.

Let $C \geq 2 \max\{E[\sup_S \tilde{R}_1(t, y)], E[\sup_S \tilde{R}_2(t, y)]\}$. Then

$$\begin{aligned} P(R(t, y) \leq C\nu(y) \text{ for all } t \in [0, 1], y \in [0, \infty)) \\ &\geq P(\sup_S \tilde{R}_1(t, y) + \sup_S \tilde{R}_2(t, y) \leq C) \\ &\geq P(\sup_S \tilde{R}_1(t, y) \leq \frac{C}{2}) P(\sup_S \tilde{R}_2(t, y) \leq \frac{C}{2}) > 0. \end{aligned}$$

Let X_0 denote the interarrival time of the first customer and V_0 denote its service time. We also denote $\bar{Q}^0(t, y)$ as an independent infinite server process starting empty and with $E(0) = 0$. Then

for s large enough, we have

$$\begin{aligned}
P(\bar{Q}(t, y) &\leq \lambda s \int_y^{t+y} \bar{F}(u) du + C\sqrt{s}\nu(y) \text{ for all } t \in [0, 1], y \in [0, \infty)) \\
&= P(\bar{Q}^0(t - X_0, y) + 1\{V_0 > t + y\} \leq \lambda s \int_y^{t+y} \bar{F}(u) du + C\sqrt{s}\nu(y) \\
&\text{for all } t \in [X_0, 1], y \in [0, \infty)) \\
&\geq P(\bar{Q}^0(t, y) + 1\{V_0 > t + X_0 + y\} \leq \lambda s \int_y^{t+X_0+y} \bar{F}(u) du + C\sqrt{s}\nu(y) \\
&\text{for all } t \in [0, 1 - X_0], y \in [0, \infty)) \\
&\geq P(\bar{Q}^0(t, y) \leq \lambda s \int_y^{t+y} \bar{F}(u) du + C\sqrt{s}\nu(y) \text{ for all } t \in [0, 1], y \in [0, \infty)) \\
&= P\left(\frac{\bar{Q}^0(t, y) - \lambda s \int_y^{t+y} \bar{F}(u) du}{\sqrt{s}} \leq C\nu(y) \text{ for all } t \in [0, 1], y \in [0, \infty)\right).
\end{aligned}$$

It is easy to check that the set $\{f : |f(t, y)| \leq C\nu(y) \text{ for all } t \in [0, 1], y \in [0, \infty)\}$ is a continuity set, thus by the Functional Central Limit Theorem result in [21], we have

$$\begin{aligned}
P\left(\frac{\bar{Q}^0(t, y) - \lambda s \int_y^{t+y} \bar{F}(u) du}{\sqrt{s}} \leq C\nu(y) \text{ for all } t \in [0, 1], y \in [0, \infty)\right) \\
\rightarrow P(R(t, y) \leq C\nu(y) \text{ for all } t \in [0, 1], y \in [0, \infty)) > 0.
\end{aligned}$$

Inequality (7) is easy to prove as we can always isolate a point (t^*, y^*) inside S . The projection of the process on that point posses Gaussian distribution. More specifically,

$$\begin{aligned}
P(\bar{Q}(t, y) > \lambda s \int_y^{t+y} \bar{F}(u) du + C\sqrt{s}\nu(y) \text{ for some } t \in [0, 1], y \in [0, \infty)) \\
&\geq P(\bar{Q}(t^*, y^*) > \lambda s \int_{y^*}^{t^*+y^*} \bar{F}(u) du + C\sqrt{s}\nu(y^*)) \\
&= P\left(\frac{\bar{Q}(t^*, y^*) - \lambda s \int_{y^*}^{t^*+y^*} \bar{F}(u) du}{\sqrt{s}} > C\nu(y^*)\right),
\end{aligned}$$

and by Fatou's lemma

$$\liminf_{s \rightarrow \infty} P\left(\frac{\bar{Q}(t^*, y^*) - \lambda s \int_{y^*}^{t^*+y^*} \bar{F}(u) du}{\sqrt{s}} > C\nu(y^*)\right) \geq P(R(t^*, y^*) > C\nu(y^*)) > 0.$$

Let $m(s) = \inf\{y : C\sqrt{s}(v(y) + \int_y^\infty v(s) ds) \leq \frac{1}{2}\}$. Following the same construction as for the QD regime, we define the sequence of random times Ξ_i 's as follows: $\Xi_0 := 0$. Given Ξ_{i-1} for $i = 1, 2, \dots$,

$$\begin{aligned}
r_i &= \inf\{k : k \geq R(\Xi_{i-1}), k = 1, 2, \dots\}, \\
z &= \inf\{k : k \geq m(s), k = 1, 2, \dots\}, \\
\Xi_i &= \Xi_{i-1} + r_i + z.
\end{aligned}$$

We introduce a Bernoulli random variable ξ_i with $\xi_i = 1$ if and only if

$$\bar{Q}_{\Xi_{i-1}+(k-1)t_0}(t, y) \leq \lambda s \int_y^{t+y} \bar{F}(u) du + C\sqrt{s}\nu(y)$$

for all $t \in [0, 1]$, $y \in [0, \infty)$ and every $k = 1, 2, \dots, r_i + z$.

We next show that $\xi_i = 1$ implies that τ^* is reached before Ξ_i . Since $r_i \geq R(\Xi_{i-1})$, all the customers at time $\Xi_{i-1} + r_i$ will be those arrive after Ξ_i . Thus we have $\xi_i = 1$ implies that

$$\begin{aligned} Q(\Xi_{i-1} + r_i, y) &\leq \sum_{k=1}^{r_i} \left\{ \lambda s \int_{(k-1)t_0+y}^{kt_0+y} \bar{F}(u) du + C\sqrt{s}\nu((k-1) + y) \right\} \\ &\leq \lambda s \int_y^\infty \bar{F}(u) du + C\sqrt{s}(\nu(y) + \int_y^\infty \nu(u) du). \end{aligned}$$

As $\int_y^\infty \bar{F}(u) du$ decays faster than $\nu(y)$ as y grows large, for s large enough, we have

$$R(\Xi_{i-1} + r_i) < m(s).$$

Likewise for every $t \in (k-1, k]$ and $k = 1, 2, \dots, z$,

$$Q(\Xi_{i-1} + r_i + t, y) \leq \lambda s \int_y^\infty \bar{F}(u) du + C\sqrt{s}(\nu(y) + \int_y^\infty \nu(u) du).$$

Thus when $\beta > C(\nu(0) + \int_0^\infty \nu(u) du)$, we have

$$Q(\Xi_{i-1} + r_i + t, 0) \leq s + C(\nu(0) + \int_0^\infty \nu(u) du)\sqrt{s} \leq s + \beta\sqrt{s}$$

for $t \in [0, R(\Xi_{i-1} + r_i)]$.

Now let $N = \inf\{i \geq 1 : \xi_i = 1\}$. Then

$$E\tau^* \leq E\left[\sum_{i=1}^N (r_i + z)\right].$$

We now show a bound for $E\sum_{i=1}^N (r_i + z)$. The proof is given in the Appendix B.

Lemma 10. *Assume $EX_n^2 < \infty$ and $EV_n^q < \infty$ for any $q > 0$. Then*

$$\log E\left[\sum_{i=1}^N (r_i + z)\right] = o(s^\delta)$$

for any $\delta > 0$.

Notice that our proof of Theorem 3 only requires the existence of the second moment of the interarrival time distribution. We thus conclude the proof of Theorem 3.

A Proof of Proposition 1

By Chebyshev's inequality,

$$P(A_{n+1} < n(\mu - \epsilon)) \leq E[\exp(\theta(n(\mu - \epsilon) - A_{n+1}))] \leq \exp(-n(-\theta(\mu - \epsilon) - \psi(-\theta)))$$

for any $\theta \geq 0$.

Let

$$I(-\epsilon) = \max_{\theta \geq 0} \{-\theta(\mu - \epsilon) - \psi(-\theta)\}.$$

As $\psi(0) = 0$, $\psi'(0) = \mu$ and $\psi''(0) = \text{Var}(X) > 0$, $I(-\epsilon) > 0$. Then

$$P(A_{n+1} < n(\mu - \epsilon)) \leq \exp(-nI(-\epsilon))$$

and

$$\sum_{n=1}^{\infty} P(A_{n+1} < n(\mu - \epsilon)) \leq \frac{\exp(-I(-\epsilon))}{1 - \exp(-I(-\epsilon))} < \infty.$$

By Borel-Cantelli lemma, $\{A_{n+1} \geq n(\mu - \epsilon)\}$ eventually almost surely.

Similarly and independently we have

$$\begin{aligned} \sum_{n=1}^{\infty} P(|V_{n+1}| > (n(\mu - \epsilon))^\alpha) &= \sum_{n=1}^{\infty} P(|V_1|^{1/\alpha} > n(\mu - \epsilon)) \\ &\leq \frac{1}{\mu - \epsilon} \int_0^\infty P(|V_1|^{1/\alpha} > \nu) d\nu < \infty. \end{aligned}$$

Thus, again by Borel-Cantelli lemma, $\{|V_{n+1}| \leq (n(\mu - \epsilon))^\alpha\}$ eventually almost surely. Therefore, $P(\kappa < \infty) = 1$.

B Proof of Lemma 8 and Lemma 10

Lemma 11. *If $EV_n^q < \infty$ for any $q > 0$, then for any fixed $p > 0$,*

$$E[(\max_{k=1,2,\dots,n} V_k)^p] = o(n^\delta)$$

for any $\delta > 0$.

Proof. For any fixed $\delta > 0$ we can find $\delta' \in (0, \delta)$. Let $q = 1/\delta' + p$. By Chebyshev's inequality we have

$$\bar{F}(u) \leq \frac{EV^q}{u^q}.$$

Let $\bar{F}_n(u) = P(\max_{k=1,2,\dots,n} V_k > u)$ then

$$\begin{aligned} E[(\max_{k=1,2,\dots,n} V_k)^p] &= p \int_0^\infty u^{p-1} \bar{F}_n(u) du \\ &\leq n^{1/(q-p)} + np \int_{n^{1/(q-p)}}^\infty u^{p-1} \bar{F}(u) du \\ &\leq n^{1/(q-p)} + np \int_{n^{1/(q-p)}}^\infty \frac{EV^q}{u^{q-p+1}} du \\ &= n^{\delta'} + \frac{p}{q-p} EV^q. \end{aligned}$$

□

$$\begin{aligned}
E\left[\sum_{i=1}^N (r_i + z)\right] &= E\left[\sum_{i=1}^{\infty} (r_i + z)I\{N \geq i\}\right] \\
&\leq \sum_{i=1}^{\infty} E[(r_i + z)^2]^{1/2} P(N \geq i)^{1/2} \text{ by Holder's inequality.}
\end{aligned}$$

Lemma 12. *If $EX_n < \infty$ and $EV_n^q < \infty$ for any $q > 0$, then for any $p \geq 1$ we have*

$$E[(r_i + z)^p]^{1/p} = o(s^\delta)$$

for any $\delta > 0$.

Proof. By Minkowski inequality

$$E[(r_i + z)^p]^{1/p} \leq E[r_i^p]^{1/p} + z.$$

Using similar argument as in the proof of Lemma 11, we can show that $l(s) = o(s^\delta)$ for any $\delta > 0$, thus $z = o(s^\delta)$ for any $\delta > 0$.

For fixed $\delta > 0$, we can find $\delta' \in (0, p\delta/(1 + p\delta))$, such that

$$\begin{aligned}
E[r_i^p] &\leq E\left[E\left[\left(\max_{k=1, \dots, N_s(\Xi_{i-1}) - N_s(\Xi_{i-2})} V_k\right)^p \mid N_s(\Xi_{i-1}) - N_s(\Xi_{i-2})\right]\right] \\
&\leq CE[(N_s(\Xi_{i-1}) - N_s(\Xi_{i-2}))^{\delta'}] \text{ Lemma 11} \\
&\leq C(E[N_s(\Xi_{i-1}) - N_s(\Xi_{i-2})])^{\delta'} \text{ Jensen's inequality for concave function} \\
&\leq C\tilde{\lambda}^{\delta'} s^{\delta'} E[r_{i-1} + z]^{\delta'} \text{ Key Renewal Theorem.}
\end{aligned}$$

Let $w_i = r_i + z$ for $i = 1, 2, \dots$. As z it is a constant that only depends on s and $z = o(s^{\delta'})$, then

$$Ew_i \geq z \geq 1$$

and

$$Ew_i = Er_i + z \leq C\tilde{\lambda}^{\delta'} s^{\delta'} (Ew_{i-1})^{\delta'} + z \leq \tilde{C}s^{\delta'} (Ew_{i-1})^{\delta'}$$

where $\tilde{C} = C\tilde{\lambda}^{\delta'} + 1$.

As $E[r_1^p] = E_\pi[R(0)^p] = o(s^{\delta'})$. By iteration we have

$$Ew_i \leq \tilde{C}^{1/(1-\delta')} s^{\delta'/(1-\delta')}$$

for $i = 1, 2, \dots$.

Thus $Er_i^p = o(s^{p\delta})$ and $E[(r_i + z)^p]^{1/p} = o(s^\delta)$. □

Proof of Lemma 8. We first notice that $P(\xi_i = 0) \leq E[w_1] \exp(-sI^*(\delta) + o(s))$ by Lemma 7.

$$P(N \geq 1) = 1.$$

$$P(N \geq 2) = P(\xi_1 = 0) \leq E[w_1] \exp(-sI^*(\delta) + o(s))$$

Recall that $w_i = r_i + z$ for $i = 1, 2, \dots$.

$$\begin{aligned}
P(N \geq 3) &= P(N \geq 1)P(N \geq 3|N \geq 2) \\
&= P(\xi_1 = 0)P(\xi_2 = 0|\xi_1 = 0) \\
&\leq P(\xi_1 = 0)E[w_2|\xi_1 = 0] \exp(-sI^*(\delta) + o(s)) \\
&\leq E[w_1]E[w_2|\xi_1 = 0] \exp(-2sI^*(\delta) + o(s)).
\end{aligned}$$

We next prove that $E[w_2|\xi_1 = 0] = \exp(o(s))$. Notice that $P(\xi_i = 0) \geq \exp(-sI^*(\delta) + o(s))$ by Lemma 7. Then for any $p > 0$, $q > 0$ and $1/p + 1/q = 1$,

$$\begin{aligned}
E[w_2|\xi_1 = 0] &= \frac{E[w_2 I\{\xi_1 = 0\}]}{P(\xi_1 = 0)} \\
&\leq \frac{E[w_2^p]^{1/p} P(\xi_1 = 0)^{1/q}}{P(\xi_1 = 0)} \text{ Holder's inequality} \\
&\leq E[w_2^p]^{1/p} E[w_1]^{1/q} \exp\left(\frac{1}{p}sI^*(\delta) + o(s)\right),
\end{aligned}$$

thus

$$\frac{1}{s} \log E[w_2|\xi_1 = 0] \leq \frac{1}{s} \left(\frac{1}{p} \log E[w_2^p] + \frac{1}{q} \log E[w_1] + o(s) \right) + \frac{1}{p} I^*(\delta).$$

By sending p to infinity, we have $E[w_2|\xi_1 = 0] = \exp(o(s))$.

Similarly by iteration,

$$P(N \geq k) = \exp(-ksI^*(\delta) + o(s))$$

for $k = 4, 5, \dots$.

Then $\sum_{i=1}^{\infty} P(N \geq i)^{1/2} = O(1)$. As $E[\sum_{i=1}^N (r_i + z)] \leq \sum_{i=1}^{\infty} E[(r_i + z)^2]^{1/2} P(N \geq i)^{1/2}$ and $E[(r_i + z)^2]^{1/2} = o(s^\delta)$ for any $\delta > 0$, we have

$$E\left[\sum_{i=1}^N (r_i + z)\right] = o(s^\delta)$$

for any $\delta > 0$. □

Proof of Lemma 10.

$$P(N \geq 1) = 1.$$

$$\begin{aligned}
P(N \geq 2) &= P(\xi_1 = 0) \\
&\leq 1 - E[\zeta_1(C)^{w_1}] \text{ Lemma 9} \\
&\leq 1 - \zeta_1(C)^{E[w_1]} \text{ Jensen's inequality} \\
&= 1 - b \exp(-o(s^\delta)).
\end{aligned}$$

Moreover

$$\begin{aligned}
P(N \geq 3) &= P(N > 2|N > 1)P(N > 1) \\
&= P(\xi_2 = 0|\xi_1 = 0)P(\xi_1 = 0) \\
&\leq E[1 - \zeta_1(C)^{w_2}|\xi_1 = 0]P(\xi_1 = 0) \\
&\leq (1 - \zeta_1(C)^{E[w_2|\xi_1=0]})P(\xi_1 = 0).
\end{aligned}$$

We next show that $E[w_2|\xi_1 = 0] = o(s^\delta)$ for any $\delta > 0$. Notice that $P(\xi_i = 0) \geq \zeta_2(C)$ by Lemma 9, then

$$E[w_2|\xi_1 = 0] = \frac{E[w_2 I\{\xi_1 = 0\}]}{P(\xi_1 = 0)} \leq \frac{Ew_2}{\zeta_2(C)}$$

Similarly by iteration we have

$$P(N \geq k) \leq (1 - b \exp(-o(s^\delta)))^k$$

for any $\delta > 0$ and $k = 4, 5, \dots$.

Then

$$\log \sum_{i=1}^{\infty} P(N \geq i)^{1/2} = o(s^\delta)$$

for any $\delta > 0$.

As $E[\sum_{i=1}^N (r_i + z)] \leq \sum_{i=1}^{\infty} E[(r_i + z)^2]^{1/2} P(N \geq i)^{1/2}$ and $E[(r_i + z)^2]^{1/2} = o(s^\delta)$ for any $\delta > 0$, we have

$$\log E[\sum_{i=1}^N (r_i + z)] = o(s^\delta)$$

for any $\delta > 0$.

□

Acknowledgement: NSF support from grants CMMI-0846816 and 1069064 is gratefully acknowledged.

References

- [1] ADLER, R. J. (1990). An introduction to continuity, extrema, and related topics for general Gaussian processes. *IMS Lecture Notes: Monograph Series* **12**,.
- [2] ASMUSSEN, S. (2003). *Applied Probability and Queues* 2 ed. Springer, New York.
- [3] BERTHELSEN, K. AND MØLLER, J. (2002). A primer on perfect simulation for spatial point process. *Bull Braz Math Soc* **33(3)**, 351–367.
- [4] BLANCHET, J., CHEN, X. AND LAM, H. (2012). Two-parameter sample path large deviation for infinite server queues. *working paper*.
- [5] BLANCHET, J. AND DONG, J. (2012). Sampling point processes on stable unbounded regions and exact simulation of queues. *Proc. of Winter Simulation Conference*.
- [6] BLANCHET, J. AND LAM, H. (2012). Rare-event simulation for many-server queues. *working paper*.
- [7] BLANCHET, J. AND SIGMAN, K. (2011). On exact sampling of stochastic perpetuities. *J. Appl. Probab.* **48A**, 165–182.
- [8] BROWN, L., GANS, N., MANDELBAUM, A., SAKOV, A., SHEN, H., ZELTYN, S. AND ZHAO, L. (2002). Statistical analysis of a telephone call center: a queueing-science perspective. *Preprint*.

- [9] BUSIC, A., GAUJAL, B. AND PERRONNIN, F. (2012). Perfect sampling of networks with finite and infinite capacity queues. In *ASMTA*. vol. 7314. Springer pp. 136–149.
- [10] CONNOR, S. AND KENDALL, W. (2007). Perfect simulation for a class of positive recurrent Markov chains. *Ann. Appl. Probab.* **3**, 781–908.
- [11] CORCORAN, J. AND TWEEDIE, R. (2001). Perfect sampling of ergodic Harris chains. *Ann. of Appl. Probab.* **11**, 438–451.
- [12] ENSOR, K. AND GLYNN, P. (2000). Simulating the maximum of a random walk. *Journal of Statistical Planning and Inference* **85**, 127–135.
- [13] FERNANDEZ, R., FERRARI, P. AND GARCIA, N. (2002). Perfect simulation for interacting point processes, loss networks and ising models. *Stoch. Process. Appl.* **102(1)**, 63–88.
- [14] FILL, J. AND HUBER, M. (2010). Perfect simulation of Vervaat perpetuities. *Electron. Comm. Probab.* **15**, 96–109.
- [15] FOSS, S. AND TWEEDIE, R. (1998). Perfect simulation and backward coupling. *Stochastic Models* **14**, 187–203.
- [16] KELLY, F. (1991). Loss networks. *Annals of Applied Probability* 319–378.
- [17] KENDALL, W. (1998). Perfect simulation for area-interaction point processes. In *Probability Towards 2000*. ed. L. Accardi and C. Heyde. Spinger, New York pp. 218–234.
- [18] KENDALL, W. (2004). Geometric ergodicity and perfect simulation. *Electron. Comm. Probab.* **9**, 140–151.
- [19] KENDALL, W. 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.
- [20] MURDOCH, D. AND TAKAHARA, G. (2006). Perfect sampling for queues and network models. *ACM Transactions of Modeling and Computer Simulation* **16**, 76–92.
- [21] PANG, G. AND WHITT, W. (2010). Two-parameter heavy-traffic limits for infinite-server queues. *Queueing Systems: Theory and Applications* 325–264.
- [22] PROPP, J. AND WILSON, D. (1996). Exact sampling with coupled Markov chains and applications to statistical mechanics. *Random Structures and Algorithms* **9**, 223–252.
- [23] SIGMAN, K. (2011). Exact simulation of the stationary distribution of the FIFO M/G/c queue. *Journal of Applied Probability* **48A**, 209–216.
- [24] SIGMAN, K. (2012). Exact simulation of the stationary distribution of the FIFO M/G/c queue: The general case of $\rho < c$. *Queueing Systems: Theory and Applications* **70**,.