

EXACT SAMPLING OF THE INFINITE HORIZON MAXIMUM OF A RANDOM WALK OVER A NON-LINEAR BOUNDARY

JOSE BLANCHET,* *Columbia University*

JING DONG,** *Northwestern University*

ZHIPENG LIU,*** *Columbia University*

Abstract

We present the first algorithm that samples $\max_{n \geq 0} \{S_n - n^\alpha\}$, where S_n is a mean zero random walk, and n^α with $\alpha \in (1/2, 1)$ defines a nonlinear boundary. We show that our algorithm has finite expected running time. We also apply this algorithm to construct the first exact simulation method for the steady-state departure process of a $GI/GI/\infty$ queue where the service time distribution has infinite mean.

Keywords: Exact simulation; Monte Carlo; queueing theory; random walk

2010 Mathematics Subject Classification: Primary 65C05

Secondary 60K25; 60G50

1. Introduction

Consider a random walk $S_n = \sum_{i=1}^n X_i$ for $n \geq 1$ and $S_0 = 0$, where $\{X_i : i \geq 1\}$ is a sequence of independent and identically distributed random variables with $E[X_1] = 0$. Without loss of generality, we shall also assume that $Var(X_1) = 1$. Moreover, we shall impose the following light-tail assumption on the distribution of X_i 's.

Assumption 1. *There exists $\delta > 0$, such that $E[\exp(\theta X_1)] < \infty$ for $\forall \theta \in (-\delta, \delta)$.*

In this paper, we develop the first algorithm that generates perfect samples (i.e. samples without any

* Postal address: Department of Industrial Engineering and Operations Research, 500 West 120th Street, New York, NY 20017. Email: jose.blanchet@columbia.edu

* Support from NSF grant DMS-132055 and NSF grant CMMI-1538217 is gratefully acknowledged by J. Blanchet.

** Postal address: 2145 Sheridan Road, Tech M239, Evanston, IL 60208-3109. Email: jing.dong@northwestern.edu

*** Postal address: Department of Industrial Engineering and Operations Research, 500 West 120th Street, New York, NY 20017. Email: zl2337@columbia.edu

bias) from the random variable

$$M_\alpha = \max_{n \geq 0} \{S_n - n^\alpha\},$$

where $\alpha \in (1/2, 1)$. Moreover, we will show that our algorithm has finite expected running time.

There has been substantial amount of work on exact sampling (i.e. sampling without any bias) from the distribution of the maximum of a negative drifted random walk, e.g. M_1 in our setting. Ensor and Glynn [6] propose an algorithm to sample the maximum when the increments of the random walk are light-tailed (i.e Assumption 1 holds). In [2], Blanchet et al. propose an algorithm to simulate a multidimensional version of M_1 driven by Markov random walks. In [5], Blanchet and Wallwater develop an algorithm to sample M_1 for the heavy-tailed case, which requires only that $E|X_1|^{2+\varepsilon} < \infty$ for some $\varepsilon > 0$ to guarantee finite expected termination time.

Some of this work is motivated by the fact that M_1 plays an important role in ruin theory and queueing models. For example, the steady state waiting time of $GI/GI/1$ queue has the same distribution as M_1 , where X_i corresponds to the (centered) difference between the i -th service time and the i -th interarrival time, (see [1]). Moreover, applying Coupling From The Past (CFTP), see for example [9] and [8], the techniques to sample M_1 jointly with the random walk $\{S_n : n \geq 0\}$ have been used to obtain perfect sampling algorithms for more general queueing systems, including multi-server queues [4], infinite server queues and loss networks [3], and multidimensional reflected Brownian motion with oblique reflection [2].

The fact that M_α stochastically dominates M_1 makes the development of a perfect sampler for M_α more difficult. For example, the direct use of exponential tilting techniques as in [6] is not applicable. However, similar to some of the previous work, the algorithmic development uses the idea of record-breakers (see e.g. [3]) and randomization procedures similar to the heavy-tailed context studied in [5].

The techniques that we study here can be easily extended, using the techniques studied in [2], to obtain exact samplers of a multidimensional analogue of M_α driven by Markov random walks (as done in [2] for the case $\alpha = 1$). Moreover, using the domination technique introduced in Section 5 of [4], the algorithms that we present here can be applied to the case in which the term n^α is replaced by $g(n)$ as long as there exists $n_0 > 0$ such that $g(n) \geq n^\alpha$ for all $n \geq n_0$.

We mentioned earlier that algorithms which simulate M_1 jointly with $\{S_n : n \geq 0\}$ have been used in applications of CFTP. Since the random variable M_α dominates M_1 , and we also simulate M_α jointly with $\{S_n : n \geq 0\}$, we expect our results here to be applicable to perfect sampling (using CFTP) for a wide range of processes. In this paper, we will show how to use the ability to simulate M_α jointly

with $\{S_n : n \geq 0\}$ to obtain the first algorithm which samples from the steady-state departure process of an infinite server queue in which the job requirements have infinite mean; the case of finite mean service/job requirements is treated in [3].

The rest of the paper is organized as follows. In Section 2 we discuss our sampling strategy. Then we provide a detailed running time analysis in Section 3. As a sanity check, we implement our algorithm in instances in which we can compute a theoretical benchmark, this is reported in Section 3.1. Finally, the application to exact simulation of the steady-state departure process of an infinite server queue with infinite mean service time is given in Section 4.

2. Sampling strategy and main algorithmic development

Our goal is to simulate M_α using a finite but random number of X_i 's. To achieve this goal, we introduce the idea of record-breakers.

Let $\psi(\theta) := \log E[\exp(\theta X_i)]$. As $\psi(\theta) = \frac{1}{2}\theta^2 + o(\theta^2)$, by Taylor expansion, there exists $\delta' < \delta$, such that $\psi(\theta) \leq \theta^2$, for $\theta \in (-\delta', \delta')$. Let

$$a < \min \left\{ 4\delta', \frac{1}{2} \right\}, \quad \text{and} \quad b = \frac{4}{a} \log \left(4 \left(\sum_{n=0}^{\infty} 2^n \exp(-a^2 2^{2n\alpha - n - 4}) \right) \right). \quad (1)$$

These choices of a and b will become clear in the proof of Lemma 1. We define a sequence of record-breaking times as $T_0 := 0$. For $k = 1, 2, \dots$, if $T_{k-1} < \infty$,

$$T_k := \inf \{ n > T_{k-1} : S_n > S_{T_{k-1}} + a(n - T_{k-1})^\alpha + b(n - T_{k-1})^{1-\alpha} \};$$

otherwise if $T_{k-1} = \infty$, then $T_k = \infty$. We also define

$$\kappa := \inf \{ k > 0 : T_k = \infty \}.$$

Because the random walk has independent increments, $P(T_i = \infty | T_{i-1} < \infty) = P(T_1 = \infty)$. Thus, κ is a geometric random variable with probability of success

$$P(T_1 = \infty).$$

We first show that κ is well defined.

Lemma 1. *For a and b satisfying (1),*

$$P(T_1 = \infty) \geq \frac{3}{4}.$$

Proof. We first notice that

$$\begin{aligned} P(T_1 < \infty) &= \sum_{n=0}^{\infty} P(T \in [2^n, 2^{n+1})) \\ &\leq \sum_{n=0}^{\infty} \sum_{k \in [2^n, 2^{n+1})} P(S_k > ak^\alpha + bk^{1-\alpha}). \end{aligned}$$

For any $k \in [2^n, 2^{n+1})$,

$$\begin{aligned} P(S_k > ak^\alpha + bk^{1-\alpha}) &\leq \exp(k\psi(\theta) - \theta(ak^\alpha + bk^{1-\alpha})) \\ &\leq \exp\left(2^{n+1}\psi(\theta) - \theta a 2^{\alpha n} - \theta b 2^{(1-\alpha)n}\right), \end{aligned}$$

for any $\theta \in (-\delta, \delta)$. If we set $\theta_n = a2^{(\alpha-1)n-2}$, as $a < 4\delta'$, $\theta_n < \delta'$. Then

$$\begin{aligned} P(S_k > ak^\alpha + bk^{1-\alpha}) &\leq \exp\left(2^{n+1}\theta_n^2 - \theta_n a 2^{\alpha n} - \theta_n b 2^{(1-\alpha)n}\right) \\ &= \exp\left(-a^2 2^{2n\alpha-n-3} - ab/4\right). \end{aligned}$$

Therefore,

$$P(T_1 < \infty) \leq \left(\sum_{n=0}^{\infty} 2^n \exp\left(-a^2 2^{2n\alpha-n-3}\right)\right) \exp\left(-\frac{ab}{4}\right) \leq \frac{1}{4},$$

where the last inequality follows from our choice of b . \square

Let

$$\xi := \max_{n \geq 0} \left\{ (an^\alpha + bn^{1-\alpha}) - \frac{1}{2}n^\alpha \right\}. \quad (2)$$

Conditional on the value of κ and the values of $\{X_i : 1 \leq i \leq T_{\kappa-1}\}$, we define

$$\Gamma(\kappa) := \left\lceil (2S_{T_{\kappa-1}} + 2\xi)^{1/\alpha} \right\rceil.$$

The choice of ξ will become clear in the proof of Lemma 2. We will next establish that

$$M_\alpha = \max_{0 \leq n \leq T_{\kappa-1} + \Gamma(\kappa)} \{S_n - n^\alpha\}.$$

Lemma 2. For $n \geq T_{\kappa-1} + \Gamma(\kappa)$,

$$S_n \leq n^\alpha.$$

Proof. For ξ defined in (2), we have for any $n \geq 0$,

$$an^\alpha + bn^{1-\alpha} \leq \frac{1}{2}n^\alpha + \xi.$$


Since $T_\kappa = \infty$, for $n \geq \Gamma(\kappa)$,

$$\begin{aligned} S_{T_{\kappa-1}+n} &\leq an^\alpha + bn^{1-\alpha} + S_{T_{\kappa-1}} \\ &\leq \frac{1}{2}n^\alpha + \xi + \frac{1}{2}\Gamma(\kappa)^\alpha - \xi \\ &\leq n^\alpha \leq (T_{\kappa-1} + n)^\alpha. \end{aligned}$$

□

Figure 1 illustrates the basic idea of our algorithmic development. The solid line is n^α . The first dotted line from the left (lowest dotted curve) is the record-breaking boundary that we start with, $an^{\alpha-1} + bn^\alpha$. T_1 is the first record-breaking time. Based on the value of S_{T_1} , we construct a new record-breaking boundary, $S_{T_1} + a(n - T_1)^{\alpha-1} + b(n - T_1)^\alpha$ (the second dotted line from the left). At time T_2 , we have another record-breaker. Based on the value of S_{T_2} , we construct again a new record-breaking boundary, $S_{T_2} + a(n - T_2)^{\alpha-1} + b(n - T_2)^\alpha$ (the third dotted line from the left). If from T_2 on, we will never break the record again ($T_3 = \infty$), then we know that for n large enough (say, $n > 3000$ in the figure), S_n will never pass the solid boundary again.

FIGURE 1: Bounds for record-breakers



The actual simulation strategy goes as follows.

Algorithm 1. Sampling $\Gamma(\kappa)$ together with $(X_i : 1 \leq i \leq T_{\kappa-1} + \Gamma(\kappa))$.

- i) Initialize $T_0 = 0$, $k = 1$.
- ii) For $T_{k-1} < \infty$, sample $J \sim \text{Bernoulli}(P(T_k = \infty | T_{k-1}))$.

- iii) If $J = 0$, sample $(X_i : i = T_{k-1} + 1, \dots, T_k)$ conditional on $T_k < \infty$. Set $k = k + 1$ and go back to step ii); otherwise ($J = 1$), set $\kappa = k$ and go to step iv).
- iv) Calculate $\Gamma(\kappa)$, sample $(X_i : i = T_{k-1} + 1, \dots, T_{k-1} + \Gamma(\kappa))$ conditional on $T_k = \infty$.

Remark 1. In general, any $a < \min\{4\delta', 1/2\}$, and $b \geq \frac{4}{a} \log(4(\sum_{n=0}^{\infty} 2^n \exp(-a^2 2^{2n\alpha - n - 4})))$ would work. However, there is a trade-off. The larger the value of a and b , the smaller the value of κ , but the value of $\Gamma(\kappa)$ would be larger. We conduct a numerical study on the choice of these parameters in Section 3.1.1.

In what follows, we shall elaborate on how to carry out step ii), iii) and iv) in Algorithm 1. In particular, step ii) and iii) are outlined in Procedure A. Step iv) is outlined in Procedure B.

2.1. Step ii) and iii) in Algorithm 1

It turns out step ii) and iii) can be carried out simultaneously using exponential tilting based on the results and proof of Lemma 1.

We start by explaining how to sample the first record-breaking time T_1 . We introduce an auxiliary random variable N with probability mass function (pmf)

$$p(n) = P(N = n) = \frac{2^n \exp(-a^2 2^{2n\alpha - n - 3})}{\sum_{m=0}^{\infty} 2^m \exp(-a^2 2^{2m\alpha - m - 3})}, \quad \text{for } n \geq 1 \quad (3)$$

We can then apply exponential tilting to sample the path $(X_1, X_2, \dots, X_{T_1})$ conditional on $T_1 < \infty$. The actual sampling algorithm goes as follows.

Procedure A. Sampling J with $P(J = 1) = P(T_1 = \infty)$; if $J = 0$, output (X_1, \dots, X_{T_1}) .

- i) Sample a random time N with pmf (3).
- ii) Let $\theta_N = a2^{N(\alpha-1)-2}$. Generate $X_1, X_2, \dots, X_{2^{N+1}}$ under exponential tilting with tilting parameter θ_N , i.e.

$$\frac{dP_{\theta_N}}{dP} \mathbf{1}\{X_i \in A\} = \exp(\theta_N X_i - \psi(\theta_N)) \mathbf{1}\{X_i \in A\}.$$

Let $T_1 = \inf\{n \geq 1 : S_n > an^\alpha + bn^{1-\alpha}\} \wedge 2^{N+1}$

- iii) Sample $U \sim \text{Uniform}[0, 1]$. If

$$U \leq \frac{\exp(-\theta_N S_{T_1} + T_1 \psi(\theta_N))}{p(N)} I\{T_1 \in [2^N, 2^{N+1}]\},$$

then set $J = 0$ and output $(X_1, X_2, \dots, X_{T_1})$; else, set $J = 1$.

We next show that Procedure A works.

Theorem 1. *In Procedure A, J is a Bernoulli random variable with probability of success $P(T_1 = \infty)$. If $J = 0$, the output $(X_1, X_2, \dots, X_{T_1})$ follows the distribution of $(X_1, X_2, \dots, X_{T_1})$ conditional on $T_1 < \infty$.*

Proof. We first show that the likelihood ratio in step iii) is well-defined.

$$\begin{aligned}
& \frac{\exp(-\theta_n S_{T_1} + T_1 \psi(\theta_n))}{P(N = n)} I\{T_1 \in [2^n, 2^{n+1})\} \\
& \leq \frac{\exp(-\theta_n (a2^{\alpha n} + b2^{(1-\alpha)n}) + 2^{n+1}\theta_n^2)}{P(N = n)} \\
& = \frac{\exp(-a^2 2^{2n\alpha - n - 3} - ab/4)}{P(N = n)} \\
& = 2^{-n} \exp(-ab/4) \sum_{m=0}^{\infty} 2^m \exp(-a^2 2^{2m\alpha - m - 3}) \leq \frac{1}{4},
\end{aligned}$$

where the last inequality follows from our choice of b as in the proof of Lemma 1. We next prove that $P(J = 0) = P(T_1 < \infty)$.

$$\begin{aligned}
E[I\{J = 0\} | N = n] &= E_{\theta_n} \left[I \left\{ U \leq \frac{\exp(-\theta_n S_{T_1} + T_1 \psi(\theta_n))}{p(n)} \right\} I\{T_1 \in [2^n, 2^{n+1})\} \right] \\
&= E_{\theta_n} \left[\frac{\exp(-\theta_n S_{T_1} + T_1 \psi(\theta_n))}{p(n)} I\{T_1 \in [2^n, 2^{n+1})\} \right] \\
&= \frac{P(T_1 \in [2^n, 2^{n+1}))}{p(n)}.
\end{aligned}$$

Thus,

$$\begin{aligned}
E[I\{J = 0\}] &= \sum_{n=0}^{\infty} E[I\{J = 0\} | N = n] p(n) \\
&= \sum_{n=0}^{\infty} P(T_1 \in [2^n, 2^{n+1})) = P(T_1 < \infty).
\end{aligned}$$

Let $P^*(\cdot)$ denote the measure induced by Procedure A. We next show that $P^*(X_1 \in A_1, \dots, X_t \in A_t | J = 0) = P(X_1 \in A_1, \dots, X_t \in A_t | T < \infty)$, where t is a positive integer, and $A_i \subset \mathbb{R}$, $i = 1, 2, \dots, t$, is a sequence of Borel measurable sets satisfying $A_i \subset \{x \in \mathbb{R} : x \leq ai^\alpha + bi^{1-\alpha}\}$ for $i < t$ and

$A_t \subset \{x \in \mathbb{R} : x > at^\alpha + bt^{1-\alpha}\}$. Let $n_t := \lfloor \log_2 t \rfloor$.

$$\begin{aligned}
& P^*(X_1 \in A_1, \dots, X_t \in A_t | J = 0) \\
&= \frac{P^*(X_1 \in A_1, \dots, X_t \in A_t, J = 0)}{P(J = 0)} \\
&= \frac{P(N = n_t)}{P(T_1 < \infty)} E_{\theta_{n_t}} \left[I\{X_1 \in A_1, \dots, X_t \in A_t\} I\left\{U \leq \frac{\exp(-\theta_{n_t} S_t + t\psi(\theta_{n_t}))}{p(n_t)}\right\} \right] \\
&= \frac{p(n_t)}{P(T_1 < \infty)} E_{\theta_{n_t}} \left[I\{X_1 \in A_1, \dots, X_t \in A_t\} \frac{\exp(-\theta_{n_t} S_t + t\psi(\theta_{n_t}))}{p(n_t)} \right] \\
&= \frac{E[I\{X_1 \in A_1, \dots, X_t \in A_t\}]}{P(T_1 < \infty)} \\
&= P(X_1 \in A_1, \dots, X_t \in A_t | T_1 < \infty).
\end{aligned}$$

□

The extension from T_1 to T_k is straightforward: because for $T_{k-1} < \infty$, given the value of T_{k-1} and $S_{T_{k-1}}$, we essentially start the random walk afresh from each T_{k-1} on. Thus, to execute step ii) and iii) in Algorithm 1, given $T_{k-1} < \infty$, we can apply Procedure A. Based on the output, if $J = 0$, we denote $(\tilde{X}_1, \tilde{X}_2, \dots, \tilde{X}_T)$ as the output from Procedure A, and set $(X_{T_{k-1}+1}, \dots, X_{T_{k-1}+T}) = (\tilde{X}_1, \dots, \tilde{X}_T)$ and $T_k = T_{k-1} + T$, otherwise, set $\kappa = k$.

2.2. Step iv) in Algorithm 1

Sampling $(X_1, \dots, X_{T_{\kappa-1}})$ is realized by iteratively applying Procedure A until it outputs $J = 1$. Once we found κ , sampling $(X_{T_{\kappa-1}+1}, \dots, X_{T_{\kappa-1}+n})$ requires us to sample the trajectory of the random walk conditioning on that it never passes the non-linear upper bound. In particular, given $\kappa = k$, we would like to sample $(X_{T_{k-1}+1}, \dots, X_{T_{k-1}+n})$ from $P(\cdot | \mathcal{F}_{k-1}, T_k = \infty)$, where $\{\mathcal{F}_k : k \geq 0\}$ denote the filtration generated by the random walk. We can achieve this conditional sampling using the acceptance-rejection technique.

We first introduce a method to simulate a Bernoulli random variable with probability of success $P(T_1 = \infty | T_1 > t, S_t)$, which follows a similar exponential tilting idea as that used in Section 2.1.

Let

$$\tilde{T}_{t,s} := \inf \{n \geq 0 : s + S_n > a(n+t)^\alpha + b(n+t)^{1-\alpha}\}.$$

Given t , we introduce an auxiliary random variable $\tilde{N}(t)$ with pmf

$$p_t(n) = P(\tilde{N}(t) = n) = \frac{2^n \exp(-2^{-n-4} a^2 (2^n + t)^{2\alpha})}{\sum_{m=0}^{\infty} 2^m \exp(-2^{-m-4} a^2 (2^m + t)^{2\alpha})}, \quad \text{for } n \geq 1. \quad (4)$$

Given $\tilde{N}(t) = n$, we apply exponential tilting to sample $(\tilde{X}_1, \tilde{X}_2, \dots, \tilde{X}_{2^{n+1}-1})$, with tilting parameter

$$\tilde{\theta}_n(t) = 2^{-n-2} a (2^n + t)^\alpha,$$

i.e.

$$\frac{dP_{\tilde{\theta}_n(t)}}{dP} 1\{X_i \in A\} = \exp(\tilde{\theta}_n(t)X_i - \psi(\tilde{\theta}_n(t))) 1\{X_i \in A\}.$$

We also define $\tilde{S}_k := \tilde{X}_1 + \dots + \tilde{X}_k$ for $k \geq 1$, and $\tilde{T} = \inf \left\{ n \geq 0 : s + \tilde{S}_n > a(n+t)^\alpha + b(n+t)^{1-\alpha} \right\} \wedge 2^{n+1}$. In what follows, we shall suppress the dependence on t when there is no confusion. Let

$$\tilde{J} = 1 - I \left\{ U \leq \frac{\exp(-\tilde{\theta}_n \tilde{S}_{\tilde{T}} + \tilde{T} \psi(\tilde{\theta}_n))}{p_t(n)} I \left\{ \tilde{T} \in [2^n, 2^{n+1}) \right\} \right\}, \quad (5)$$

where $U \sim \text{Uniform}[0, 1]$.

Lemma 3. For \tilde{J} defined in (5), when $s < \frac{a}{4}t^\alpha$, we have

$$P(\tilde{J} = 1) = P(\tilde{T}_{t,s} = \infty).$$

Proof. We first notice that

$$\begin{aligned} & \frac{\exp(-\tilde{\theta}_n \tilde{S}_{\tilde{T}} + \tilde{T} \psi(\tilde{\theta}_n))}{p_t(n)} I \left\{ \tilde{T} \in [2^n, 2^{n+1}) \right\} \\ & \leq \frac{1}{p_t(n)} \exp \left(-\tilde{\theta}_n \left(a(2^n + t)^\alpha + b(2^n + t)^{1-\alpha} - \frac{a}{4}t^\alpha \right) + 2^{n+1} \tilde{\theta}_n^2 \right) \\ & \leq \frac{1}{p_t(n)} \exp \left(-2^{-n-3} a^2 (2^n + t)^{2\alpha} + 2^{-n-4} a^2 (2^n + t)^{2\alpha} - \frac{ab}{4} \right) \\ & = \frac{1}{p_t(n)} \exp \left(-2^{-n-4} a^2 (2^n + t)^{2\alpha} - \frac{ab}{4} \right) \\ & \leq \left(\sum_{m=0}^{\infty} 2^m \exp(-2^{-m-4} a^2 (2^m + t)^{2\alpha}) \right) \times \exp(-ab/4) \\ & \leq \left(\sum_{m=0}^{\infty} 2^m \exp(-a^2 2^{2m\alpha - m - 4}) \right) \times \exp(-ab/4) \leq \frac{1}{4}, \end{aligned}$$

where the last inequality follows from our choice of a and b . The rest of the proof follows exact the same steps as the proof of Theorem 1. We shall omit it here. \square

Let

$$L(n) = \inf \left\{ k \geq n : S_k > ak^\alpha + bk^{1-\alpha} \text{ or } S_k < \frac{a}{4}k^\alpha \right\}.$$

The sampling algorithm goes as follows.

Procedure B. Sampling (X_1, \dots, X_n) conditional on $T_1 = \infty$.

- i) Sample (X_1, \dots, X_n) under the nominal distribution $P(\cdot)$.
- ii) If $\max_{1 \leq k \leq n} \{S_k - ak^\alpha - bk^{1-\alpha}\} > 0$, go back to step i); else, go to step iii).

- iii) Sample $L(n)$ and $(X_{n+1}, X_{n+2}, \dots, X_{L(n)})$ under the nominal distribution $P(\cdot)$. If $S_{L(n)} > aL(n)^\alpha + bL(n)^{1-\alpha}$, go back to step i); else, go to step iv).
- iv) Sample \tilde{N} with probability mass function $p_{L(n)}$ defined in (4). Generate $(\tilde{X}_1, \tilde{X}_2, \dots, \tilde{X}_{2^{\tilde{N}}-1})$ under exponential tilting with tilting parameter $\tilde{\theta}_{\tilde{N}} = 2^{\tilde{N}-2}a(2^{\tilde{N}} + L(n))^\alpha$. Let $\tilde{T} = \inf\{k \geq 1 : S_{L(n)} + \tilde{S}_k > a(k + L(n))^\alpha + b(k + L(n))^{1-\alpha}\} \wedge 2^{\tilde{N}+1}$.
- v) Sample $U \sim \text{Uniform}[0, 1]$. If

$$U \leq \frac{\exp\left(-\tilde{\theta}_{\tilde{N}}S_{\tilde{T}} + \tilde{T}\psi(\tilde{\theta}_{\tilde{N}})\right)}{p_t(\tilde{N})} I\left\{\tilde{T} \in \left[2^{\tilde{N}}, 2^{\tilde{N}+1}\right]\right\},$$

set $\tilde{J} = 0$ and go back to Step i); else, set $\tilde{J} = 1$ and output (X_1, \dots, X_n) .

We next show that Procedure B works.

Theorem 2. *The output of Procedure B follows the distribution of (X_1, \dots, X_n) conditional on $T_1 = \infty$.*

Proof. Let $P'(\cdot) = P(\cdot | T_1 = \infty)$. We first notice that

$$\frac{dP'}{dP}(X_1, \dots, X_n) = \frac{I\{T_1 > n\}P(T_1 = \infty | S_n, T_1 > n)}{P(T_1 = \infty)} \leq \frac{1}{P(T_1 = \infty)}.$$

Let $P''(\cdot)$ denote the measure induced by Procedure B. Then we have, for any sequence of Borel measurable sets $A_i \subset \mathbb{R}$, $i = 1, 2, \dots, n$,

$$\begin{aligned} & P''(X_1 \in A_1, \dots, X_n \in A_n) \\ &= P\left(X_1 \in A_1, \dots, X_n \in A_n | T_1 > L(n), \tilde{J} = 1\right) \\ &= P\left(X_1 \in A_1, \dots, X_n \in A_n | T_1 > L(n), \tilde{T}_{L(n), S_{L(n)}} = \infty\right) \\ &= P(X_1 \in A_1, \dots, X_n \in A_n | T_1 = \infty), \end{aligned}$$

where the second equality follows from Lemma 3, and the third equality follows from the fact that

$$P(T_1 = \infty | S_t, T_1 > t) = P\left(\tilde{T}_{t, S_t} = \infty\right).$$

□

To execute Step iv) in Algorithm 1, we apply Procedure B with $n = \Gamma(\kappa)$.

3. Running time analysis

In this section, we provide a detailed running time analysis of Algorithm 1.

Theorem 3. *Algorithm 1 has finite expected running time.*

We divide the analysis into the following steps.

1. From Lemma 1, the number of iterations between step ii) and iii) follows a geometric distribution with probability of success $P(T_1 = \infty)$.
2. In each iteration (when applying Procedure A), we will show that the length of the path needed to sample J has finite moments of all orders (Lemma 4).
3. For step iv), we will show that $\Gamma(\kappa)$ has finite moments of all orders (Lemma 5).
4. When applying Procedure B for step iv), we will first show that $L(\Gamma(\kappa))$ has finite moments of all orders (Lemma 5).
5. When applying Procedure B for step iv), we will also show that the length of the path needed to sample J has finite moments of all orders (Lemma 6).

Lemma 4. *The length of the path needed to sample the Bernoulli J in Procedure A has finite moments of every order.*

Proof. The length of the path generated in Procedure A is bounded by 2^{N+1} , where N 's distribution is defined in (3). Therefore, $\forall r > 0$,

$$\begin{aligned} E[2^{(N+1)r}] &= \frac{\sum_{m=0}^{\infty} 2^{(m+1)r} 2^m \exp(-a^2 2^{2m\alpha-m-3})}{\sum_{m=0}^{\infty} 2^m \exp(-a^2 2^{2m\alpha-m-3})} \\ &= \frac{\sum_{m=0}^{\infty} \exp(-a^2 2^{2m\alpha-m-3} + (mr + r + m) \log 2)}{\sum_{m=0}^{\infty} 2^m \exp(-a^2 2^{2m\alpha-m-3})}. \end{aligned}$$

Since for m sufficiently large,

$$\exp(-a^2 2^{2m\alpha-m-3} + (mr + r + m) \log 2) \leq \exp(-a^2 2^{2m\alpha-m-4}),$$

for fixed $r > 0$, $\exists C > 0$, such that

$$\begin{aligned} &\frac{\sum_{m=0}^{\infty} \exp(-a^2 2^{2m\alpha-m-3} + (mr + r + m) \log 2)}{\sum_{m=0}^{\infty} 2^m \exp(-a^2 2^{2m\alpha-m-3})} \\ &\leq C \frac{\sum_{m=0}^{\infty} \exp(-a^2 2^{2m\alpha-m-4})}{\sum_{m=0}^{\infty} 2^m \exp(-a^2 2^{2m\alpha-m-3})} < \infty. \end{aligned}$$

Note that this also implies that

$$E[T^r I(T < \infty)] \leq E[2^{(N+1)r} I(J = 0)] \leq E[2^{(N+1)r}] < \infty.$$

□

Lemma 5. $\Gamma(\kappa)$ and $L(\Gamma(\kappa))$ have finite moments of any order.

Proof. We start with $\Gamma(\kappa)$. Let $R_n := S_n - an^\alpha - bn^{1-\alpha}$. For $T_i < \infty$, we also denote

$$\mathcal{R}_i := S_{T_i} - S_{T_{i-1}} - a(T_i - T_{i-1})^\alpha - b(T_i - T_{i-1})^{1-\alpha}.$$

$$\begin{aligned} \Gamma(\kappa) &= \left[(2S_{T_{\kappa-1}} + 2\xi)^{1/\alpha} \right] \\ &\leq \left[\left(\sum_{i=1}^{\kappa-1} (T_i - T_{i-1})^\alpha + 2\kappa\xi + 2 \sum_{i=1}^{\kappa-1} \mathcal{R}_i \right)^{1/\alpha} \right]. \end{aligned}$$

We first prove that conditioning on $T_1 < \infty$, R_{T_1} has finite moments of every order.

$$\begin{aligned} &E[e^{\gamma R_{T_1}} I(T_1 < \infty)] \\ &= \sum_{n=0}^{\infty} E[e^{\gamma R_n} I(T_1 = n)] \\ &= \sum_{n=0}^{\infty} E[e^{\gamma(X_n + S_{n-1} - an^\alpha - bn^{1-\alpha})} I(T_1 = n)] \\ &\leq \sum_{n=0}^{\infty} E[e^{\gamma X_n} I(T_1 = n)] \\ &\leq \sum_{n=0}^{\infty} E[e^{p\gamma X_n}]^{1/p} E[I(T_1 = n)]^{1/q} \quad \text{for } p, q > 1, \frac{1}{p} + \frac{1}{q} = 1 \text{ by Hölder's inequality} \\ &\leq E[e^{p\gamma X_n}]^{1/p} \sum_{n=0}^{\infty} P(T_1 = n)^{1/q}. \end{aligned}$$

Because X_n has moment generating function within a neighborhood of 0, we can choose $p > 0$ and $\gamma > 0$ such that $E[e^{p\gamma X_n}]^{1/p} < \infty$. In the proof of Lemma 4 we showed that $\forall r > 0$, $E[T_1^r I(T_1 < \infty)] < \infty$, which implies that $P(T_1 = n) = O(\frac{1}{n^r})$. Because r can be any positive value, $\sum_{n=0}^{\infty} P(T_1 = n)^{1/q} < \infty$. By Jensen's inequality, for any $r \geq 1$,

$$\begin{aligned} &E \left[\left(\sum_{i=1}^{\kappa-1} (T_i - T_{i-1})^\alpha + 2\kappa\xi + 2 \sum_{i=1}^{\kappa-1} \mathcal{R}_i \right)^{r/\alpha} \right] \\ &\leq 3^{\frac{r}{\alpha}-1} E \left[\left(\sum_{i=1}^{\kappa-1} (T_i - T_{i-1})^\alpha \right)^{r/\alpha} + (2\kappa\xi)^{r/\alpha} + \left(2 \sum_{i=1}^{\kappa-1} \mathcal{R}_i \right)^{r/\alpha} \right]. \end{aligned} \quad (6)$$

We next analyze each of the three part on the right hand side of (6). As κ is a geometric random variable, $E[(2\kappa\xi)^{r/\alpha}] < \infty$.

$$\begin{aligned}
E \left[\left(\sum_{i=1}^{\kappa-1} (T_i - T_{i-1})^\alpha \right)^{r/\alpha} \right] &= E \left[E \left[\left(\sum_{i=1}^{\kappa-1} (T_i - T_{i-1})^\alpha \right)^{r/\alpha} \middle| \kappa \right] \right] \\
&\leq E \left[\kappa^{r/\alpha-1} E \left[\sum_{i=1}^{\kappa-1} (T_i - T_{i-1})^r \middle| \kappa \right] \right] \\
&\leq E \left[\kappa^{r/\alpha} 2^{(N+1)r+1} \right] \\
&= E \left[\kappa^{r/\alpha} \right] E \left[2^{(N+1)r+1} \right] < \infty.
\end{aligned}$$

Similarly, we have

$$E \left[\left(2 \sum_{i=1}^{\kappa-1} \mathcal{R}_i \right)^{r/\alpha} \right] \leq E \left[(2\kappa)^{r/\alpha} \right] E \left[R_{T_1}^{r/\alpha} \middle| T_1 < \infty \right] < \infty.$$

Therefore, we have

$$E[\Gamma(\kappa)] < \infty.$$

We next analyze $L(\Gamma(\kappa))$.

$$L(\Gamma(\kappa)) - \Gamma(\kappa) \leq \inf \left\{ n \geq 0 : S_{n+\Gamma(\kappa)} < \frac{a}{4}(n + \Gamma(\kappa))^\alpha \right\}.$$

Given $\Gamma(\kappa) = n_*$ and $S_{\Gamma(\kappa)} = s_*$, since $s_* < an_*^\alpha + bn_*^{1-\alpha}$,

$$\begin{aligned}
&P(L(\Gamma(\kappa)) - \Gamma(\kappa) > n \mid \Gamma(\kappa) = n_*, S_{\Gamma(\kappa)} = s_*) \\
&\leq P \left(S_n \geq \frac{a}{4}(n + n_*)^\alpha - s_* \right) \\
&\leq P \left(S_n \geq \frac{a}{4}(n + n_*)^\alpha - an_*^\alpha - bn_*^{1-\alpha} \right) \\
&\leq P \left(S_n \geq \frac{a}{4}(n + n_*)^\alpha - \frac{1}{2}n_*^\alpha - \xi \right) \\
&\leq \exp \left(n\theta^2 - \theta \left(\frac{a}{4}(n + n_*)^\alpha - \frac{1}{2}n_*^\alpha - \xi \right) \right) \text{ for } 0 < \theta < \delta'.
\end{aligned}$$

Let $w_n = \frac{a}{4}(n + n_*)^\alpha - \frac{1}{2}n_*^\alpha - \xi$. If we pick $\theta = \epsilon_n | \frac{w_n}{n} |$ where ϵ_n is chosen such that $\theta < \delta'$. Then

$$\exp(n\theta^2 - \theta w_n) \leq \exp \left(-\frac{w_n^2}{n} \epsilon_n (1 - \epsilon_n) \right) \leq \exp \left(-\frac{w_n^2}{4n} \right).$$

We notice that for n large enough,

$$w_n \leq \frac{a}{5}(n + n_*)^\alpha.$$

Thus, $\exists C > 0$, such that

$$\begin{aligned} P(L(\Gamma(\kappa)) - \Gamma(\kappa) > n | \Gamma(\kappa) = n_*, S_{\Gamma(\kappa)} = s_*) &\leq C \exp\left(-\frac{a^2}{100} \frac{(n + n_*)^{2\alpha}}{n}\right) \\ &\leq C \exp\left(-\frac{a^2}{100} n^{2\alpha-1}\right). \end{aligned}$$

This implies that, given $\Gamma(\kappa)$ and $S_{\Gamma(\kappa)}$, $L(\Gamma(\kappa)) - \Gamma(\kappa)$ has finite moments of all orders. \square

Lemma 6. *The length of the path needed to sample the Bernoulli \tilde{J} in Procedure B has finite moments of every order.*

Proof. The length of the path in Procedure B is bounded by $2^{\tilde{N}(t)+1}$, with $\tilde{N}(t)$ sampled from (4). For any $r > 0$,

$$\begin{aligned} E[2^{(\tilde{N}(t)+1)r}] &= \frac{\sum_{m=0}^{\infty} 2^{(m+1)r} 2^m \exp(-2^{-m-4} a^2 (2^m + t)^{2\alpha})}{\sum_{m=0}^{\infty} 2^m \exp(-2^{-m-4} a^2 (2^m + t)^{2\alpha})} \\ &= \frac{\sum_{m=0}^{\infty} \exp(-2^{-n-4} a^2 (2^m + t)^{2\alpha} + (mr + m + r) \log 2)}{\sum_{m=0}^{\infty} 2^m \exp(-2^{-n-4} a^2 (2^m + t)^{2\alpha})}. \end{aligned}$$

As for m sufficiently large,

$$\exp(-2^{-m-4} a^2 (2^m + t)^{2\alpha} + (mr + m + r) \log 2) < \exp(-2^{-m-5} a^2 (2^m + t)^{2\alpha}),$$

for fixed $r > 0$, $\exists C > 0$, such that

$$E[2^{(\tilde{N}(t)+1)r}] \leq C \frac{\sum_{m=0}^{\infty} \exp(-2^{-m-5} a^2 (2^m + t)^{2\alpha})}{\sum_{m=0}^{\infty} 2^m \exp(-2^{-m-4} a^2 (2^m + t)^{2\alpha})} < \infty. \quad \square$$

3.1. Numerical experiments

In this section, we conduct numerical experiments to analyze the performance of Algorithm 1 for different values of parameters. We will also conduct a sanity check of the correctness of our algorithms (empirically) by simulating the steady state departure process of an infinite server queueing model. The details of the simulation of the infinite server queue will be given in Section 4.

3.1.1. *Choice of parameter a* In Remark 1 we briefly discuss how the parameters a and b would affect the performance of Algorithm 1. We shall fix the value of b upon our choice of a as in (1), as we want to guarantee that probability of record-breaking is small enough, while keeping $\Gamma(\kappa)$ as small as possible. In this subsection, we conduct some simulation experiments to analyze the effect of different values of a on the computational cost. We first notice that the choice of a and b will affect the distribution of N ,

which is the length of trajectory generated in Procedure A. In Procedure B, the value of $\Gamma(\kappa)$, $L(\Gamma(\kappa))$ and the distribution of \tilde{N} also depend on the value of a and b .

Let $X_i \stackrel{d}{=} X - 1$, where X is a unit rate exponential random variable. Then $\psi(\theta) = -\theta - \log(1 - \theta)$, for $\theta < 1$. Let $g(\theta) := \psi(\theta) - \theta^2$. As $g'(0) = 0$, $g''(\theta) = \frac{1}{(1-\theta)^2} - 2$, we have

$$g(\theta) < 0 \quad \forall \theta \in (-1, 1 - \frac{\sqrt{2}}{2}).$$

Therefore, we can set $\delta' = 1 - \frac{\sqrt{2}}{2}$, and when $\theta \in (-\delta', \delta')$, $\psi(\theta) < \theta^2$. According to (1), $a < \min(\frac{1}{2}, 4\delta') = \frac{1}{2}$. We ran Algorithm 1 with different values of a and α . Table 1 summarizes the running time of the algorithm in different settings.

TABLE 1: Running time of Algorithm 1 (in seconds)

a	$\alpha = 0.8$	$\alpha = 0.85$	$\alpha = 0.9$	$\alpha = 0.95$
0.1	287.58	39.62	10.20	4.99
0.2	36.24	8.11	4.19	3.15
0.3	13.38	5.03	2.94	2.56
0.4	7.90	3.53	2.41	2.25
0.45	7.06	3.31	2.43	2.15

We observe that while a is away from the upper bound $\frac{1}{2}$, the running time decreases as a increases. We also observe that the decreasing rate in a is larger for smaller values of α , which in general implies greater curvature of the nonlinear boundary.

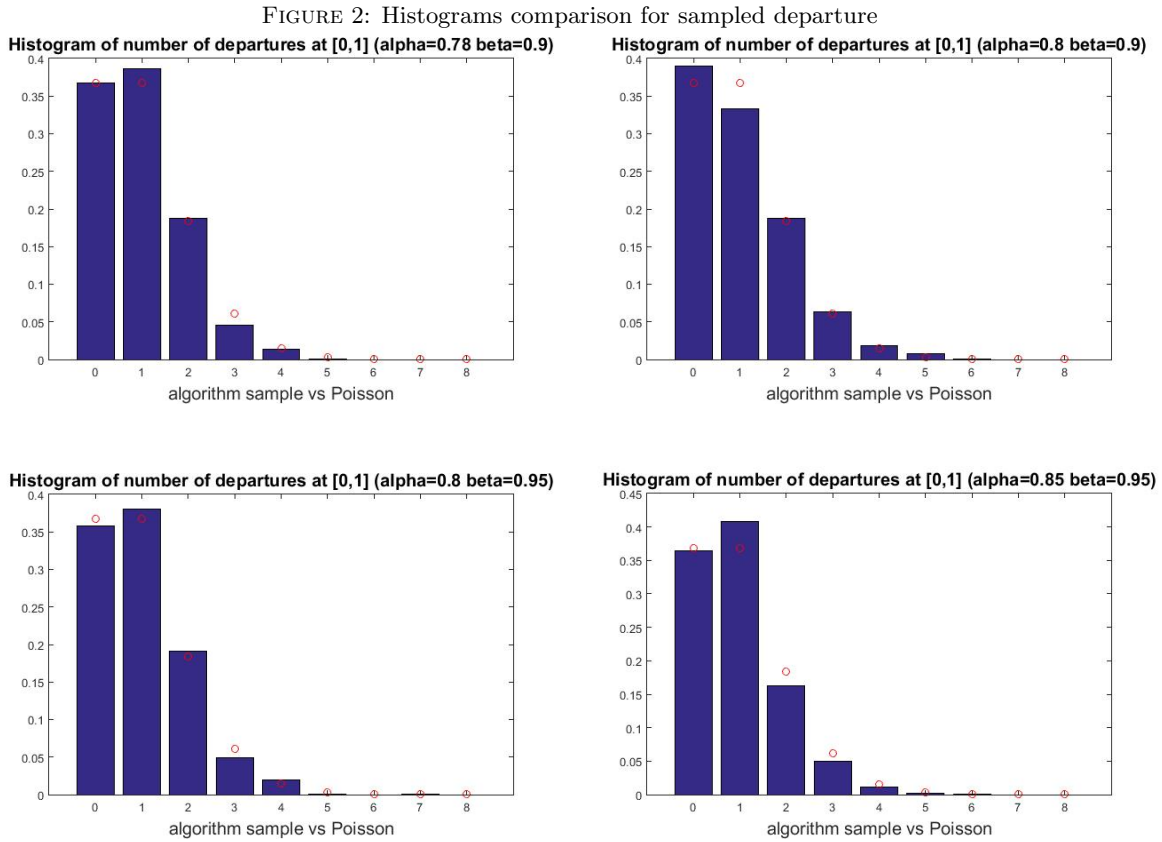
3.1.2. Departure process of an $M/G/\infty$ queue We apply a variation of Algorithm 1 to simulate the steady state departure process of an infinite server queue whose service times have infinite mean. The details of the algorithm will be explained in Section 4. In particular, we consider an infinite server queue having Poisson arrival process with rate 1, and Pareto service time distributions with probability density function (pdf)

$$f(v) = \beta v^{-(\beta+1)} I\{v \geq 1\},$$

for $\beta \in (1/2, 1)$. Notice that we already know the departure process of this $M/G/\infty$ queue should also be Poisson process with rate 1, therefore, this numerical experiment would help us verify the correctness of our algorithm.

We sample the arrival process using Procedure A' and Procedure B' described in Section 4, which are modifications of Procedure A and Procedure B to adapt for the absolute value of the random walk. We truncate the length of path at 10^6 steps. We tried different pairs of parameters α and β , and executed

1000 trials for each pair of α and β . We count the number of departures between time 0 to 1 for each run and construct the corresponding relative frequency bar plot. (Figure 2). Figure 2 suggests that the distribution of simulated departures between time 0 and 1 indeed follows a Poisson distribution with rate 1. In particular, the distribution is independent of the values of α and β , which is consistent with what we expected. We also conduct goodness of fit tests with the four groups of sampled data. The p-values for the tests are 0.2404, 0.2589, 0.4835, and 0.1137 respectively. Therefore the tests fail to reject that the generated samples are Poisson distributed.




4. Departure process of an infinite server queue

We finish the paper with an application of the algorithm developed in Section 2 to sample the steady-state departure process of an infinite server queue with general interarrival time and service time distributions. We assume the interarrival times are i.i.d.. Independent of the arrival process, the service times are also i.i.d. and may have infinite mean.

Suppose the system starts operating from the infinite past, then it would be at stationarity at time

0. We want to sample all the departures in the interval $[0, h]$. We next introduce a point process representation of infinite server queue to facilitate the explanation of simulation strategy. In particular, we mark each arriving customer as a point in a 2-dimensional space, where the x -coordinate records its arrival time and the y -coordinate records its service time (service requirement). Based on this point process representation, Figure 3 provides a graphical representation of the region that we are interested in simulation. Specifically, to sample the departure process on $[0, h]$, we want to sample all the points (customers) that fall into the shaded area. We further divide the shaded area into two part, namely H and G . Sampling the points that fall into G is easy. As G is a bounded area, we can simply sample all the arrivals between 0 and h , and decide, using their service time information, whether they fall into region G or not. The challenge lies in sampling the points in H , as it is an unbounded region.

FIGURE 3: Point process representation of infinite server queue



We mark the points sequentially (according to their arrival times) backwards in time from time 0 as $(-A_1, V_1)$, $(-A_2, V_2)$, \dots , where $-A_n$ is the arrival time of the n -th arrival counting backwards in time and V_n is his service time. Let $A_0 := 0$. We then denote $X_n := A_{n-1} - A_n$, as the interarrival time between the n -th arrival and the $(n-1)$ -th arrival.

For simplicity of notation, we write

$$\mathcal{H} := \{(-A_n, V_n) : A_n < V_n < A_n + h\}.$$

It is the collection of points that fall into region H . If we can find a random number Ξ such that

$$V_n < A_n \text{ or } V_n > A_n + h$$

for $n \geq \Xi$, then we can sample the point process up to Ξ and find \mathcal{H} .

We further introduce an idea to separate the simulation of the arrival process and the service time process. It requires us to find a sequence of $\{\epsilon_n : n \geq 1\}$, satisfying the following two properties:

1. There exists a well-defined random number Ξ_1 , such that

$$n\mu - \epsilon_n < A_n < n\mu + \epsilon_n \text{ for all } n \geq \Xi_1.$$

2. There exists a well-defined random number Ξ_2 , such that

$$V_n < n\mu - \epsilon_n \text{ or } V_n > n\mu + \epsilon_n + h \text{ for all } n \geq \Xi_2.$$

This allows us to find Ξ_1 and Ξ_2 separately and set $\Xi = \max\{\Xi_1, \Xi_2\}$. In particular, we can choose ϵ_n satisfying the following two conditions:

$$\text{C1) } \sum_{n=1}^{\infty} P(|A_n - n\mu| > \epsilon_n) < \infty,$$

$$\text{C2) } \sum_{n=1}^{\infty} P(V_n \in (n\mu - \epsilon_n, n\mu + \epsilon_n + h)) < \infty.$$

Then Borel-Cantelli Lemma guarantees that Ξ_1 and Ξ_2 are “well-defined”, i.e. finite almost surely.

We next introduce a specific choice of ϵ_n when the service times follow a Pareto distribution with shape parameter $\beta \in (1/2, 1)$. We denote the pdf of V_i as $f(\cdot)$, which takes the form

$$f(v) = \beta v^{-(\beta+1)} I\{v \geq 1\}. \quad (7)$$

We also write $\bar{F}(\cdot)$ as the tail cdf of V_i . We assume the interarrival time has finite moment generating function in a neighborhood of the origin. This is without loss of generality. Because if the interarrival time is heavy-tailed, we can simulate a coupled infinite server queue with truncated interarrival times, $X_i^C = \min\{X_i, C\}$, then it would serve as an upper bound (in terms of the number of departures) of the original infinite server queue in a path-by-path sense. Let $\mu := E[X]$ denote the mean interarrival time and $\sigma^2 := \text{Var}(X)$ denote its variance.

In this case, we can set $\epsilon_n = n^\alpha$ for $1/2 < \alpha < \beta$. We next show that our choice of ϵ_n satisfies C1) and C2) respectively. We also explain how to find Ξ_1 and Ξ_2 .

4.1. Sampling of the arrival process and Ξ_1

Lemma 7. *If $\epsilon_n = n^\alpha$ for $\alpha > 1/2$,*

$$\sum_{n=1}^{\infty} P(|A_n - n\mu| > \epsilon_n) < \infty.$$

Proof. We notice that $A_n = \sum_{i=1}^n X_i$ is a random walk with X_i being i.i.d. interarrival times with mean μ , *except the first one*. X_1 follows the backward recurrent time distribution of the interarrival time distribution. By moderate deviation principle [7], we have

$$\frac{1}{n^{2\alpha-1}} P(|A_n - n\mu| > n^\alpha) \rightarrow -\frac{1}{2\sigma^2}.$$

As $2\alpha - 1 > 0$, $\sum_{n=1}^{\infty} P(|A_n - n\mu| > n^\alpha) < \infty$. □

Let $S_n = T_n - n\mu$. We notice that both S_n and $-S_n$ are mean zero random walks.

$$P(|S_n| > n^\alpha) \leq P(S_n > n^\alpha) + P(-S_n > n^\alpha).$$

Thus, we can apply a modified version of Algorithm 1 to find Ξ_1 . In particular, we define a modified sequence of record-breaking times as follows. Let $T'_0 := 0$. For $k \geq 1$, if $T'_{k-1} < \infty$,

$$T'_k := \inf \left\{ n > T'_{k-1} : S_n > S_{T'_{k-1}} + a(n - T'_{k-1})^\alpha + b(n - T'_{k-1})^{1-\alpha} \right. \\ \left. \text{or } S_n < S_{T'_{k-1}} - a(n - T'_{k-1})^\alpha - b(n - T'_{k-1})^{1-\alpha} \right\};$$

else, $T'_k = \infty$. Then the modified version of Algorithm 1 goes as follows.

Algorithm 1'. Sampling Ξ together with $(X_i : 1 \leq i \leq \Xi)$.

- i) Initialize $T_0 = 0, k = 1$.
- ii) For $T_{k-1} < \infty$, sample $J \sim \text{Bernoulli}(P(T_k = \infty | T_{k-1}))$. If $J = 0$, sample $(X_i : i = T_{k-1} + 1, \dots, T_k)$ conditional on $T_k < \infty$. Set $k = k + 1$ and go back to step ii); otherwise ($J = 1$), set $\Xi_1 = T_{k-1}$ and go to step iii). (see Procedure A')
- iii) Apply Procedure C (detailed in Section 4.2) to sample Ξ_2 .
- iv) Set $\Xi = \max\{\Xi_1, \Xi_2\}$. If $\Xi > \Xi_1$, sample $(X_i : i = T_{k-1} + 1, \dots, \Xi)$ conditional on $T_k = \infty$. (see Procedure B')

We also modify Procedure A and Procedure B as follows.

Procedure A' . Sampling J' with $P(J' = 1) = P(T'_1 = \infty)$, if $J' = 0$, output $(X_1, \dots, X_{T'_1})$.

- i) Sample a random time N with pmf (3). Let $\theta_N = a2^{N(\alpha-1)-2}$. Sample $U_1 \sim \text{Uniform}[0, 1]$. If $U_1 \leq 1/2$, go to step ii a), else go to step ii b).
- ii a) Generate $X_1, \dots, X_{2^{N+1}-1}$ under exponential tilting with tilting parameter θ_N . Let

$$T'_1 = \inf\{n \geq 1 : |S_n| > an^\alpha + bn^{1-\alpha}\} \wedge 2^N.$$

- ii b) Generate $X_1, \dots, X_{2^{N+1}-1}$ under exponential tilting with tilting parameter $-\theta_N$. Let

$$T'_1 = \inf\{n \geq 1 : |S_n| > an^\alpha + bn^{1-\alpha}\} \wedge 2^N.$$

iii) Generate $U_2 \sim \text{Uniform}[0, 1]$. If

$$U_2 \leq \frac{\left(\frac{1}{2} \exp(\theta_N S_{T'_1} - \psi(\theta_N) T'_1) + \frac{1}{2} \exp(-\theta_N S_{T'_1} - \psi(-\theta_N) T'_1)\right)^{-1}}{p(N)} \times I \{T'_1 \in [2^N, 2^{N+1})\},$$

then set $J' = 0$ and output $(X_1, X_2, \dots, X_{T'_1})$; else, set $J' = 1$.

Proposition 1. *In Procedure A', J' is a Bernoulli random variable with probability of success $P(T'_1 = \infty)$. If $J = 0$, the output $(X_1, X_2, \dots, X_{T'_1})$ follows the distribution of $(X_1, X_2, \dots, X_{T'_1})$ conditional on $T'_1 < \infty$.*

The proof of Proposition 1 follows exactly the same line of analysis as the proof of Theorem 1. We shall omit it here.

Let

$$L'(n) = \inf \left\{ k > n : S_k \in \left(-\frac{a}{4} k^\alpha, \frac{a}{4} k^\alpha \right) \text{ or } S_k > a k^\alpha + b k^{1-\alpha} \text{ or } S_k < -a k^\alpha - b k^{1-\alpha} \right\}.$$

Procedure B'. Sampling (X_1, \dots, X_n) conditional on $T'_1 = \infty$.

- i) Sample (X_1, \dots, X_n) under the nominal distribution $P(\cdot)$.
- ii) If $\max_{1 \leq k \leq n} \{S_k - a k^\alpha - b k^{1-\alpha}\} > 0$ or $\min_{1 \leq k \leq n} \{S_k + a k^\alpha + b k^{1-\alpha}\} < 0$, go back to step i); else, go to step iii).
- iii) Sample $L'(n)$ and $(X_{n+1}, \dots, X_{L'(n)})$ under the nominal distribution $P(\cdot)$. If $|S_{L'(n)}| > a L'(n)^\alpha + b L'(n)^{1-\alpha}$, go back to step i); else, go to step iv).
- iv) Sample \tilde{N} with probability mass function $p_{L(n)}$ defined in (4). Set $\tilde{\theta}_{\tilde{N}} = 2^{\tilde{N}-2} a \left(2^{\tilde{N}} + L(n) \right)^\alpha$. Sample $U_1 \sim \text{Uniform}[0, 1]$. If $U_1 < 1/2$, go to step v a); else, go to step v b).
- v a) Generate $\tilde{X}_1, \tilde{X}_2, \dots, \tilde{X}_{2^{\tilde{N}+1}-1}$ under exponential tilting with tilting parameter $\tilde{\theta}_{\tilde{N}}$. Let

$$\tilde{T}' = \inf \left\{ n \geq 1 : \left| S_{L'(n)} + \tilde{S}_k \right| > a(k + L'(n))^\alpha + b(k + L'(n))^{1-\alpha} \right\} \wedge 2^{\tilde{N}+1}.$$

- v b) Generate $\tilde{X}_1, \tilde{X}_2, \dots, \tilde{X}_{2^{\tilde{N}+1}-1}$ under exponential tilting with tilting parameter $-\tilde{\theta}_{\tilde{N}}$. Let

$$\tilde{T}' = \inf \left\{ n \geq 1 : \left| S_{L'(n)} + \tilde{S}_k \right| > a(k + L'(n))^\alpha + b(k + L'(n))^{1-\alpha} \right\} \wedge 2^{\tilde{N}+1}.$$

vi) Sample $U_2 \sim \text{Uniform}[0, 1]$. If

$$U_2 \leq \frac{\left(\frac{1}{2} \exp\left(\tilde{\theta}_{\tilde{N}} \tilde{S}_{\tilde{T}'} - \tilde{\psi}(\tilde{\theta}_{\tilde{N}})\right) + \frac{1}{2} \exp\left(-\tilde{\theta}_{\tilde{N}} \tilde{S}_{\tilde{T}'} - \tilde{\psi}(-\tilde{\theta}_{\tilde{N}})\right)\right)^{-1}}{p_t(\tilde{N})} \times I\left\{\tilde{T}' \in \left[2^{\tilde{N}}, 2^{\tilde{N}+1}\right)\right\},$$

set $\tilde{J}' = 0$ and go back to Step i); else, set $\tilde{J}' = 1$ and output (X_1, \dots, X_n) .

Proposition 2. *The output of Procedure B' follows the distribution of (X_1, \dots, X_n) conditional on $T'_1 = \infty$.*

The proof of Proposition 2 follows exactly the same line of analysis as the proof of Theorem 2. We shall omit it here.

4.2. Sampling of the service time process and Ξ_2

Lemma 8. *If $\epsilon_n = n^\alpha$ for $1/2 < \alpha < \beta$,*

$$\sum_{n=1}^{\infty} P(V_n \in (n\mu - \epsilon_n, n\mu + \epsilon_n + h)) < \infty.$$

Proof.

$$\begin{aligned} P(V_n \in (n\mu - \epsilon_n, n\mu + \epsilon_n + h)) &= \bar{F}(n\mu - \epsilon_n) - \bar{F}(n\mu + \epsilon_n + h) \\ &\leq \frac{\beta}{(n\mu - n^\alpha)^{(\beta+1)}} (2n^\alpha + h) \\ &= \frac{\beta(2 + hn^{-\alpha})}{n^{\beta+1-\alpha}(\mu - n^{-(\beta-\alpha)})^{\beta+1}}. \end{aligned}$$

As $\beta + 1 - \alpha > 1$,

$$\sum_{n=1}^{\infty} \frac{\beta(2 + hn^{-\alpha})}{n^{\beta+1-\alpha}(\mu - n^{-(\beta-\alpha)})^{\beta+1}} < \infty.$$

□

To find Ξ_2 , we use a similar record-breaker idea. In particular, we say V_n is a record-breaker if

$$V_n \in (n\mu - \epsilon, n\mu + \epsilon_n + h).$$

The idea is to find the record-breakers sequentially until there are no more record-breakers. Specifically, let $K_0 := 0$ and if $K_{i-1} < \infty$,

$$K_i = \inf\{n > K_{i-1} : V_n \in (n\mu - \epsilon, n\mu + \epsilon_n + h)\}.$$

if $K_{i-1} = \infty$, $K_i = \infty$. Let $\tau = \min\{i > 0 : K_i = \infty\}$. Then we can set $\Xi_2 = K_{\tau-1}$.

The task now is to find K_i 's one by one. We start with K_1 .

$$P(K_1 = \infty) = \prod_{n=1}^{\infty} (1 - P(V_n \in (n\mu - \epsilon_n, n\mu + \epsilon_n + h))).$$

Let

$$u(k) = \prod_{n=1}^k (1 - P(V_n \in (n\mu - \epsilon_n, n\mu + \epsilon_n + h))).$$

Then we have $P(K_1 = \infty) < u(k+1) < u(k)$ for any $k \geq 1$ and $\lim_{k \rightarrow \infty} u(k) = P(K_1 = \infty)$. We also notice that $u(k) - u(k-1) = P(K_1 = k)$.

From the proof of Lemma 8, we have for $n > (2/\mu)^{1/(\beta-\alpha)}$,

$$P(V_n \in (n\mu - \epsilon_n, n\mu + \epsilon_n + h)) < \frac{2(2+h)\beta}{\mu} \frac{1}{n^{\beta+1-\alpha}}.$$

Then for k^* large enough such that $k^* > (2/\mu)^{1/(\beta-\alpha)}$, and $\frac{2(2+h)\beta}{\mu} \frac{1}{n^{\beta+1-\alpha}} < 1$, we have for $k > k^*$.

$$\begin{aligned} & \prod_{n=k+1}^{\infty} (1 - P(V_n \in (n\mu - \epsilon_n, n\mu + \epsilon_n + h))) \\ & \geq \prod_{n=k+1}^{\infty} \left(1 - \frac{2(2+h)\beta}{\mu} \frac{1}{n^{\beta+1-\alpha}} \right) \\ & \geq \exp \left(-\frac{(2+h)\beta}{\mu} \sum_{n=k+1}^{\infty} \frac{1}{n^{\beta+1-\alpha}} \right) \\ & \geq \exp \left(-\frac{(2+h)\beta}{\mu} (k+1)^{-(\beta-\alpha)} \right). \end{aligned}$$

Let $l(k) = 0$ for $k < k^*$ and

$$l(k) = u(k) \exp \left(-\frac{2(2+h)\beta}{\mu} (k+1)^{-(\beta-\alpha)} \right)$$

for $k > k^*$. Then we have $l(k) \leq l(k+1) < P(K_1 = \infty)$ and $\lim_{k \rightarrow \infty} l(k) = P(K_1 = \infty)$.

Similarly, given $K_{i-1} = m < \infty$, we can construct the sequences of upper and lower bounds for $P(K_i = \infty | K_{i-1} = m)$ as

$$u_m(k) = \prod_{n=m+1}^k (1 - P(V_n \in (n\mu - \epsilon_n, n\mu + \epsilon_n + h)))$$

for $k > m$, and

$$l_m(k) = u_m(k) \exp \left(-\frac{(2+h)\beta}{\mu} (k+1)^{-(\beta-\alpha)} \right).$$

Based on the sequence of lower and upper bounds, given $K_{i-1} = m$, we can sample K_i using the following iterative procedure.

Procedure C. Sample K_i conditional on $K_{i-1} = m$.

- i) Generate $U \sim \text{Uniform}[0, 1]$. Set $k = m + 1$. Calculate $u_m(k)$ and $l_m(k)$.
- ii) While $l_m(k) < U < u_m(k)$
 Set $k = k + 1$. Update $u_m(k)$ and $l_m(k)$.
 end While.
- iii) If $U < l_m(k)$, output $K_i = \infty$; else, output $K_i = k$.

Once we find the values of K_i 's, sampling V_n 's conditional on the information of K_i 's is straightforward.

We next provide some comments about the running time of Procedure C. Let Φ_i denote the number of iterations in Procedure C to generate K_i . We shall show that while $P(\Phi_i < \infty) = 1$, while $E[\Phi_i] = \infty$. Take Φ_1 as an example:

$$\begin{aligned} P(\Phi_1 > n) &= P(K_1 > n) \\ &= P(l_1(n) < U < u_1(n)) \\ &\geq u_1(n) \left(1 - \exp\left(-\frac{2(2+h)\beta}{\mu}(n+1)^{-(\beta-\alpha)}\right) \right), \end{aligned}$$

with

$$1 - \exp\left(-\frac{2(2+h)\beta}{\mu}(n+1)^{-(\beta-\alpha)}\right) = O(n^{-(\beta-\alpha)}),$$

and $u_1(n) \geq P(K_1 = \infty)$ for any $n \geq 1$. As $1 < \beta - \alpha < 1$, we have $P(K_1 < \infty) = 1$, but $\sum_{n=1}^{\infty} P(K_1 > n) = \infty$. Thus, $P(\Phi_1 < \infty) = 1$, but $E[\Phi_1] = \infty$.

The fact that the Procedure C has infinite expected termination time may be unavoidable in the following sense. In the absence of additional assumptions on the traffic feeding into the infinite server queue, any algorithm which simulates stationary departures during, say, time interval $[0, 1]$, must be able to directly simulate the earliest arrival, from the infinite past, which departs in $[0, 1]$. If such arrival is simulated sequentially backwards in time, we now argue that the expected time to detect such arrival must be infinite. Assuming, for simplicity, deterministic inter-arrival times equal to 1, and letting $-T < 0$ be the time at which such earliest arrival occurs, then we have that

$$\begin{aligned} P(T > n) &\geq P(\cup_{k=n+1}^{\infty} \{V_k \in [k, k+1]\}) \\ &\geq (1 - P(V > n)) \sum_{k=n+1}^{\infty} P(V_k \in [k, k+1]) \\ &= (1 - P(V > n))P(V > n+1). \end{aligned}$$

As $\sum_{n=0}^{\infty} P(V > n) = \infty$, we must have that $E[T] = \infty$.

Remark 2. Based on our analysis above, in general, there is a trade-off between Ξ_1 and Ξ_2 in terms of ϵ_n . The smaller ϵ_n is, the larger the value of Ξ_1 and the smaller the value of Ξ_2 .

References

- [1] S. ASMUSSEN. (2003). *Applied Probability and Queues*, 2 edition. Springer, New York.
- [2] J. BLANCHET, X. CHEN, AND J. DONG. (2012). Steady-state simulation of reflected brownian motion and related stochastic networks. *The Annals of Applied Probability*, 25(6):32093250, 2015.
- [3] J. BLANCHET AND J. DONG. (2014). Perfect sampling for infinite server and loss systems. *Advances in Applied Probability*, **47(3)**, 761–786.
- [4] J. BLANCHET, J. DONG, AND Y. PEI. (2015). Perfect sampling of gi/gi/c queues. Working paper, available at <http://arxiv.org/pdf/1508.02262v1.pdf>.
- [5] J. BLANCHET AND A. WALLWATER. (2015). Exact sampling for the steady-state waiting times of a heavy-tailed single server queue. *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, **25(4)**.
- [6] K. ENSOR AND P. GLYNN. (2000). Simulating the maximum of a random walk. *Journal of Statistical Planning and Inference*, **85**, 127–135.
- [7] A. GANESH, N. O’CONNELL, AND D. WISCHIK. (2004). *Big Queues*. Lecture notes in Mathematics. Springer, Berlin.
- [8] W.S. KENDALL. (1998). Perfect simulation for the area-interaction point process. *Probability towards 2000*, pages 218–234. Springer, New York.
- [9] J. PROPP AND D. WILSON. (1996). Exact sampling with coupled Markov chains and applications to statistical mechanics. *Random Structures and Algorithms*, **9**, 223–252.