

EXACT SAMPLING OF THE INFINITE HORIZON MAXIMUM OF A RANDOM WALK OVER A NONLINEAR BOUNDARY

JOSE BLANCHET,* *Stanford University*
JING DONG* AND
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 $\alpha \in (\frac{1}{2}, 1)$ 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/ ∞ 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 $\mathbb{E}[X_1] = 0$ and $\text{var}(X_1) < \infty$. Without loss of generality, we shall also assume that $\text{var}(X_1) = 1$. Moreover, we shall impose the following light-tailed assumption on the distribution of the X_i .

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

In this paper we develop the first algorithm that generates perfect samples (i.e. samples without any bias) from the random variable

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

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

There has been a 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] proposed an algorithm to sample the maximum when the increments of the random walk are light tailed (i.e. Assumption 1 holds). Blanchet *et al.* [2]

Received 7 February 2017; revision received 6 August 2018.

* Postal address: Stanford University, 475 Via Ortega, Stanford, CA 94305, USA.

Email address: jose.blanchet@stanford.edu

Support from NSF grant DMS-132055 and NSF grant CMMI-1538217 is gratefully acknowledged.

** Postal address: Decision, Risk, and Operations Division, Columbia University, 3022 Broadway, New York, NY 10027, USA. Email address: jing.dong@gsb.columbia.edu

Support from NSF grant DMS-1720433 is gratefully acknowledged.

*** Postal address: Department of Industrial Engineering and Operations Research, Columbia University, 500 West 120th Street, New York, NY 20017, USA. Email address: z12337@columbia.edu

proposed an algorithm to simulate a multidimensional version of M_1 driven by Markov random walks. Blanchet and Wallwater [4] developed an algorithm to sample M_1 for the heavy-tailed case, which requires only that $\mathbb{E}[|X_1|^{2+\varepsilon}] < \infty$ for some $\varepsilon > 0$ to guarantee finite expected termination time.

Some of this work was 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 a GI/GI/1 queue has the same distribution as M_1 , where X_i corresponds to the (centred) 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 [5], 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 [4].

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 [5], 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. 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 . To achieve this goal, we introduce the idea of record breakers.

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

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

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, $\mathbb{P}(T_i = \infty \mid T_{i-1} < \infty) = \mathbb{P}(T_1 = \infty)$. Thus, κ is a geometric random variable with probability of success

$$\mathbb{P}(T_1 = \infty).$$

We first show that κ is well defined.

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

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

Proof. We first note that

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

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

$$\begin{aligned} \mathbb{P}(S_k > ak^\alpha + bk^{1-\alpha}) &\leq \exp(k\psi(\theta) - \theta(ak^\alpha + bk^{1-\alpha})) \\ &\leq \exp(2^{n+1}\psi(\theta) - \theta a2^{\alpha n} - \theta b2^{(1-\alpha)n}) \end{aligned}$$

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

$$\begin{aligned} kP(S_k > ak^\alpha + bk^{1-\alpha}) &\leq \exp(2^{n+1}\theta_n^2 - \theta_n a2^{\alpha n} - \theta_n b2^{(1-\alpha)n}) \\ &= \exp\left(-a^2 2^{2n\alpha-n-3} - \frac{ab}{4}\right). \end{aligned}$$

Therefore,

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

where the last inequality follows from our choice of b . □

Let

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

As $a < \frac{1}{2}$, $\xi < \infty$. Conditional on the value of κ and the values of $\{X_i: 1 \leq i \leq T_{\kappa-1}\}$, we define

$$\Gamma(\kappa) := \lceil (2S_{T_{\kappa-1}} + 2\xi)^{1/\alpha} \rceil. \tag{3}$$

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 $n \geq T_{\kappa-1} + \Gamma(\kappa)$,*

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

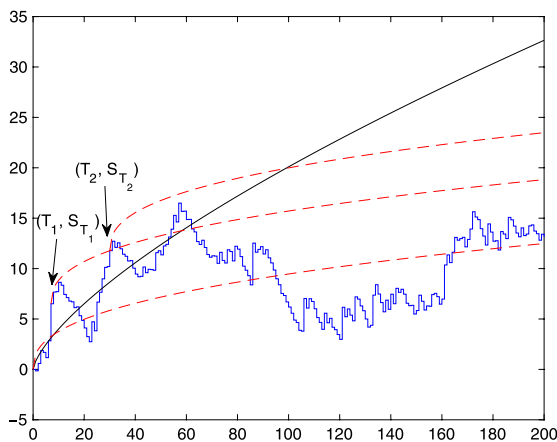


FIGURE 1: Bounds for record breakers.

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 $n \geq \Gamma(\kappa)$,

$$S_{T_{\kappa-1}+n} \leq an^\alpha + bn^{1-\alpha} + S_{T_{\kappa-1}} \leq \frac{1}{2}n^\alpha + \xi + \frac{1}{2}\Gamma(\kappa)^\alpha - \xi \leq n^\alpha \leq (T_{\kappa-1} + n)^\alpha. \quad \square$$

Figure 1 demonstrates the basic idea of our algorithmic development. (Note that the figure is rescaled for illustrative purposes. In actual simulation, the record-breaking events happen with a very small probability.) The solid line is n^α . The first dotted line from the left (lowest dashed curve) is the record-breaking boundary that we start with, $an^\alpha + bn^{1-\alpha}$. T_1 is the first record-breaking time. Based on the value of S_{T_1} , we construct a new record-breaking boundary, $S_{T_1} + a(n - T_1)^\alpha + b(n - T_1)^{1-\alpha}$ (the second dashed 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, $S_{T_2} + a(n - T_2)^\alpha + b(n - T_2)^{1-\alpha}$ (the third dashed line from the left). If, from T_2 on, we never break the record again ($T_3 = \infty$), then we know that, for large enough n (say, $n > 100$ in the figure), S_n will never pass the solid boundary again. Note that here we will need $a < 1$, which is guaranteed by (1), but a tighter constraint is imposed on a in (1) for algorithmic development and technical reasons related to Lemma 1 and 2.

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 $J \sim \text{Bernoulli}(\mathbb{P}(T_k = \infty \mid T_{k-1}))$.
 - (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 $a < \min \left\{ 4\delta', \frac{1}{2} \right\}$, and

$$b \geq \frac{4}{a} \log \left(4 \left(\sum_{n=0}^{\infty} 2^n \exp(-a^2 2^{2n\alpha - n - 4}) \right) \right)$$

would work. However, there is a trade-off. The larger the values 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.

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

2.1. Steps (ii) and (iii) in Algorithm 1

It turns out steps (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)

$$p(n) = \mathbb{P}(N = n) = \frac{2^n \exp(-a^2 2^{2n\alpha - n - 3})}{\sum_{m=0}^{\infty} 2^m \exp(-a^2 2^{2m\alpha - m - 3})} \quad \text{for } n \geq 1. \tag{4}$$

We can then apply exponential tilting to sample the path $(X_1, X_2, \dots, X_{T_1})$ conditional on $T_1 < \infty$.

When sampling the random walk, we use $\mathbb{P}(\cdot)$ to represent the measure induced by the original distribution of the random walk, which we refer to as the nominal distribution. We also denote $\mathbb{P}_\theta(\cdot)$ as the measure induced by the exponential tilting with tilting parameter θ . The actual sampling algorithm goes as follows.

Procedure A. Sampling a Bernoulli J with probability of success $\mathbb{P}(J = 1) = \mathbb{P}(T_1 = \infty)$. If $J = 0$, output (X_1, \dots, X_{T_1}) .

- (i) Sample a random time N with PMF (4).
- (ii) Let $\theta_N = a 2^{N(\alpha - 1) - 2}$. Generate $X_1, X_2, \dots, X_{2^{N+1} - 1}$ under exponential tilting with tilting parameter θ_N , i.e.

$$\frac{d\mathbb{P}_{\theta_N}}{d\mathbb{P}} \mathbf{1}\{X_i \in A\} = \exp(\theta_N X_i - \psi(\theta_N)) \mathbf{1}\{X_i \in A\}.$$

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

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

then set $J = 0$ and output $(X_1, X_2, \dots, X_{T_1})$; otherwise, 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 $\mathbb{P}(T_1 = \infty)$. 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 less than 1 almost surely. Based on this, we will then prove that $\mathbb{P}(J = 0) = \mathbb{P}(T_1 < \infty)$. That is,

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

where the last inequality follows from our choice of b as in the proof of Lemma 1.

We next prove that $\mathbb{P}(J = 0) = \mathbb{P}(T_1 < \infty)$:

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

where the second equation uses the result that the likelihood ratio is less than 1; the last equation follows from the observation that

$$\frac{d\mathbb{P}}{d\mathbb{P}_{\theta_n}}(\mathbf{1}\{T_1 \in [2^n, 2^{n+1}]\}) = \exp(-\theta_n S_{T_1} + T_1 \psi(\theta_n)) \mathbf{1}\{T_1 \in [2^n, 2^{n+1}]\}.$$

Then we have

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

Let $\mathbb{P}^*(\cdot)$ denote the measure induced by Procedure A. We next show that $\mathbb{P}^*(X_1 \in A_1, \dots, X_t \in A_t | J = 0) = \mathbb{P}(X_1 \in A_1, \dots, X_t \in A_t | T_1 < \infty)$, 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 $A_i \subset \{x \in \mathbb{R} : x \leq at^\alpha + bi^{1-\alpha}\}$ for $i < t$ and $A_t \subset \{x \in \mathbb{R} : x > at^\alpha + bt^{1-\alpha}\}$. Let $n_t := \lfloor \log_2 t \rfloor$. Then

$$\begin{aligned} &\mathbb{P}^*(X_1 \in A_1, \dots, X_t \in A_t | J = 0) \\ &= \frac{\mathbb{P}^*(X_1 \in A_1, \dots, X_t \in A_t, J = 0)}{\mathbb{P}(J = 0)} \end{aligned}$$

$$\begin{aligned}
 &= \frac{\mathbb{P}(N = n_t)}{\mathbb{P}(T_1 < \infty)} E_{\theta_{n_t}} \left[\mathbf{1}\{X_1 \in A_1, \dots, X_t \in A_t\} \mathbf{1}\left\{U \leq \frac{\exp(-\theta_{n_t} S_t + t\psi(\theta_{n_t}))}{p(n_t)}\right\} \right] \\
 &= \frac{p(n_t)}{\mathbb{P}(T_1 < \infty)} \mathbb{E}_{\theta_{n_t}} \left[\mathbf{1}\{X_1 \in A_1, \dots, X_t \in A_t\} \frac{\exp(-\theta_{n_t} S_t + t\psi(\theta_{n_t}))}{p(n_t)} \right] \\
 &= \frac{\mathbb{E}[\mathbf{1}\{X_1 \in A_1, \dots, X_t \in A_t\}]}{\mathbb{P}(T_1 < \infty)} \\
 &= \mathbb{P}(X_1 \in A_1, \dots, X_t \in A_t \mid T_1 < \infty). \quad \square
 \end{aligned}$$

The extension from T_1 to T_k is straightforward: 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 $S_{T_{k-1}}$ for each T_{k-1} . Thus, to execute steps (ii) and (iii) of Algorithm 1, given $T_{k-1} < \infty$, we can apply Procedure A. 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 have found κ , we achieve sampling $(X_{T_{\kappa-1}+1}, \dots, X_{T_{\kappa-1}+\Gamma(\kappa)})$ by developing a procedure that could sample $(X_{T_{\kappa-1}+1}, \dots, X_{T_{\kappa-1}+n})$ with any given $n > 0$, conditioning on the fact that the trajectory of the random walk never passes the nonlinear upper bound, $S_{T_{\kappa-1}} + a(n - T_{\kappa-1})^\alpha + b(n - T_{\kappa-1})^{1-\alpha}$. To be more precise, given $\kappa = k$, for any $n > 0$ (including $n = \Gamma(\kappa)$), we would like to sample $(X_{T_{\kappa-1}+1}, \dots, X_{T_{\kappa-1}+n})$ from $\mathbb{P}(\cdot \mid \mathcal{F}_{k-1}, T_k = \infty)$, 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 $\mathbb{P}(T_1 = \infty \mid T_1 > t, S_t)$, which follows a similar exponential tilting idea as that used in Section 2.1. In analogy to Section 2.1, we introduce a record-breaking time with a temporal–spatial shift, and an auxiliary random variable leading to the definition of the tilting parameter.

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

$$p_t(n) = \mathbb{P}(\tilde{N}(t) = n) = \frac{2^n \exp(-2^{-n-4} a^2 (2^n + t)^{2\alpha})}{\sum_{m=0}^\infty 2^m \exp(-2^{-m-4} a^2 (2^m + t)^{2\alpha})} \quad \text{for } n \geq 1. \quad (5)$$

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.

$$\frac{d\mathbb{P}_{\tilde{\theta}_n(t)}}{d\mathbb{P}} \mathbf{1}\{X_i \in A\} = \exp(\tilde{\theta}_n(t) X_i - \psi(\tilde{\theta}_n(t))) \mathbf{1}\{X_i \in A\}.$$

We also define $\tilde{S}_k := \tilde{X}_1 + \dots + \tilde{X}_k$ for $k \geq 1$, and

$$\tilde{T} = \inf\{n \geq 0 : s + \tilde{S}_n > a(n + t)^\alpha + b(n + t)^{1-\alpha}\} \wedge 2^{n+1}.$$

Let

$$\tilde{J} = 1 - \mathbf{1}\left\{U \leq \frac{\exp(-\tilde{\theta}_n \tilde{S}_{\tilde{T}} + \tilde{T}\psi(\tilde{\theta}_n))}{p_t(n)} \mathbf{1}\{\tilde{T} \in [2^n, 2^{n+1})\}\right\}, \tag{6}$$

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

Lemma 3. For \tilde{J} defined in (6), when $s < a^\alpha/4$, we have

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

Proof. We first note that

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

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

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 $\mathbb{P}(\cdot)$.
- (ii) If $\max_{1 \leq k \leq n}\{S_k - ak^\alpha - bk^{1-\alpha}\} > 0$, go back to step (i); otherwise, go to step (iii).
- (iii) Sample $L(n)$ and $(X_{n+1}, X_{n+2}, \dots, X_{L(n)})$ under the nominal distribution $\mathbb{P}(\cdot)$. If $S_{L(n)} > aL(n)^\alpha + bL(n)^{1-\alpha}$, go back to step (i); otherwise, go to step (iv).
- (iv) Sample \tilde{N} with probability mass function $p_{L(n)}$ defined in (5). Generate

$$(\tilde{X}_1, \tilde{X}_2, \dots, \tilde{X}_{2^{\tilde{N}+1}-1})$$

under exponential tilting with tilting parameter $\tilde{\theta}_{\tilde{N}} = 2^{\tilde{N}-2}a(2^{\tilde{N}} + L(n))^\alpha$. 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

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

set $\tilde{J} = 0$ and go back to step (i); otherwise, 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 $\mathbb{P}'(\cdot) = \mathbb{P}(\cdot \mid T_1 = \infty)$. We first notice that

$$\frac{d\mathbb{P}'}{d\mathbb{P}}(X_1, \dots, X_n) = \frac{\mathbf{1}\{T_1 > n\}\mathbb{P}(T_1 = \infty \mid S_n, T_1 > n)}{\mathbb{P}(T_1 = \infty)} \leq \frac{1}{\mathbb{P}(T_1 = \infty)}.$$

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

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

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

$$\mathbb{P}(T_1 = \infty \mid S_t, T_1 > t) = \mathbb{P}(\tilde{T}_{t, S_t} = \infty). \quad \square$$

To execute step (iv) of 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.

- (a) From Lemma 1, the number of iterations between steps (ii) and (iii) follows a geometric distribution with probability of success $\mathbb{P}(T_1 = \infty) \geq 3/4$.
- (b) 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).
- (c) For step (iv), we will show that $\Gamma(\kappa)$ has finite moments of all orders (Lemma 5).
- (d) When applying Procedure B for step (iv), we will show that the total length of the paths needed in Procedure B has finite moments of every order (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 the distribution of N is defined in (4). Therefore, for all $r > 0$,

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

Since, for sufficiently large m ,

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

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

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

Note that this also implies that

$$\mathbb{E}[T_1^r \mathbf{1}(T_1 < \infty)] \leq \mathbb{E}[2^{(N+1)r} \mathbf{1}(J = 0)] \leq \mathbb{E}[2^{(N+1)r}] < \infty. \quad \square$$

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 define

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

Then we have

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

where the last inequality follows from the definition of ξ in (2).

In what follows, we first prove that, conditioning on $T_1 < \infty$, R_{T_1} has finite moments of every order. That is,

$$\begin{aligned} &\mathbb{E}[e^{\gamma R_{T_1}} \mathbf{1}(T_1 < \infty)] \\ &= \sum_{n=0}^{\infty} \mathbb{E}[e^{\gamma R_n} \mathbf{1}(T_1 = n)] \\ &= \sum_{n=0}^{\infty} \mathbb{E}[e^{\gamma(X_n + S_{n-1} - an^\alpha - bn^{1-\alpha})} \mathbf{1}(T_1 = n)] \\ &\leq \sum_{n=0}^{\infty} \mathbb{E}[e^{\gamma X_n} \mathbf{1}(T_1 = n)] \\ &\leq \sum_{n=0}^{\infty} \mathbb{E}[e^{p\gamma X_n}]^{1/p} \mathbb{E}[\mathbf{1}(T_1 = n)]^{1/q} \quad \left(\text{for } p, q > 1, \frac{1}{p} + \frac{1}{q} = 1 \text{ by Hölder's inequality}\right) \\ &= \mathbb{E}[e^{p\gamma X_1}]^{1/p} \sum_{n=0}^{\infty} \mathbb{P}(T_1 = n)^{1/q}. \end{aligned}$$

Because X_1 has moment generating function within a neighbourhood of 0, we can choose $p > 0$ and $\gamma > 0$ such that $\mathbb{E}[e^{p\gamma X_1}]^{1/p} < \infty$. In the proof of Lemma 4 we showed that, for all $r > 0$, $\mathbb{E}[T_1^r \mathbf{1}(T_1 < \infty)] < \infty$, which implies that $\mathbb{P}(T_1 = n) = O(1/n^r)$. As r can be any positive value, $\sum_{n=0}^{\infty} \mathbb{P}(T_1 = n)^{1/q} < \infty$.

We next show that $\Gamma(\kappa)$ has finite moments of all orders. By Jensen's inequality, for any fixed $r \geq 1$,

$$\begin{aligned} \mathbb{E}[\Gamma(\kappa)^r] &\leq \mathbb{E}\left[\left(\sum_{i=1}^{\kappa-1} (T_i - T_{i-1})^\alpha + 2\kappa\xi + 2\sum_{i=1}^{\kappa-1} \mathcal{R}_i\right)^{r/\alpha}\right] \\ &\leq 3^{r/\alpha-1} \mathbb{E}\left[\left(\sum_{i=1}^{\kappa-1} (T_i - T_{i-1})^\alpha\right)^{r/\alpha} + (2\kappa\xi)^{r/\alpha} + \left(2\sum_{i=1}^{\kappa-1} \mathcal{R}_i\right)^{r/\alpha}\right]. \end{aligned} \tag{7}$$

We shall analyse each of the three parts on the right-hand side of (7). As κ is a geometric random variable, $\mathbb{E}[(2\kappa\xi)^{r/\alpha}] < \infty$. Hence,

$$\begin{aligned} \mathbb{E}\left[\left(\sum_{i=1}^{\kappa-1} (T_i - T_{i-1})^\alpha\right)^{r/\alpha}\right] &= \mathbb{E}\left[\mathbb{E}\left[\left(\sum_{i=1}^{\kappa-1} (T_i - T_{i-1})^\alpha\right)^{r/\alpha} \middle| \kappa\right]\right] \\ &\leq \mathbb{E}\left[\kappa^{r/\alpha-1} \mathbb{E}\left[\sum_{i=1}^{\kappa-1} (T_i - T_{i-1})^r \middle| \kappa\right]\right] \\ &\leq \mathbb{E}[\kappa^{r/\alpha} \mathbb{E}[T_1^r \mid T_1 < \infty]] \\ &= \mathbb{E}[\kappa^{r/\alpha}] \mathbb{E}[T_1^r \mid T_1 < \infty] \\ &< \infty. \end{aligned}$$

Similarly, we have

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

Therefore, we have

$$\mathbb{E}[\Gamma(\kappa)^r]<\infty.$$

As for $L(\Gamma(\kappa))$, we first note that

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

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

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

Let $w_n=a(n+n_*)^\alpha/4-n_*^\alpha/2-\xi$. If we pick $\theta=\epsilon_n|w_n/n|$, 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 note that, for large enough n ,

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

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

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

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

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

Proof. To sample the trajectory, using the notations defined in Procedure B, the length of each path generated, steps (i)–(iv), either accepted or rejected, satisfies

$$\begin{aligned} n+(L(n)-n)\mathbf{1}\{\tilde{S}_k\leq ak^\alpha+bk^{1-\alpha},1\leq k\leq n\}+2^{\tilde{N}+1}\mathbf{1}\{\tilde{S}_k\leq ak^\alpha+bk^{1-\alpha},1\leq k\leq L(n)\} \\ \leq L(n)+2^{\tilde{N}+1}, \end{aligned}$$

where \tilde{N} is sampled in step (iv) according to (5).

We start by establishing a bound for $E[2^{r\tilde{N}}]$ for any fixed $r > 0$. We have proved in Lemma 5 that, for all n , $L(n)$ has finite moments of all orders. Moreover, for any $r > 0$, $t > 0$, $\tilde{N}(t)$ generated from $p_t(\cdot)$ (defined in (5)) satisfies

$$\mathbb{E}[2^{\tilde{N}(t)r}] = \frac{\sum_{m=0}^{\infty} 2^{(r+1)m} \exp(-2^{-m-4} a^2 (2^m + t)^{2\alpha})}{\sum_{m=0}^{\infty} 2^m \exp(-2^{-m-4} a^2 (2^m + t)^{2\alpha})}. \tag{8}$$

We next prove that $\mathbb{E}[2^{\tilde{N}(t)r}] = O(t^r)$, which leads to the desired bound for $\mathbb{E}[2^{r\tilde{N}}]$. This is achieved in two steps. In the first step we show that, for large enough m , the summand in the numerator of (8) decays exponentially fast.

Let $\eta_1 := (2^{2\alpha} - 2)/2$. For large enough m , we have

$$\begin{aligned} 2^m &\geq \frac{(2 + \eta_1)^{1/(2\alpha)} - 1}{2 - (2 + \eta_1)^{1/(2\alpha)}} t & (9) \\ \iff 2^m (2 - (2 + \eta_1)^{1/(2\alpha)}) &\geq (2 + \eta_1)^{1/(2\alpha)} t - t \\ \iff 2^{m+1} + t &\geq (2^m + t)(2 + \eta_1)^{1/(2\alpha)} \\ \iff (2^{m+1} + t)^{2\alpha} &\geq (2 + \eta_1)(2^m + t)^{2\alpha}. \end{aligned}$$

Then we have

$$\begin{aligned} &\frac{2^{(1+r)(m+1)} \exp(-2^{-(m+1)-4} a^2 (2^{m+1} + t)^{2\alpha})}{2^{(1+r)m} \exp(-2^{-m-4} a^2 (2^m + t)^{2\alpha})} \\ &= \exp(-2^{-(m+1)-4} a^2 (2^{m+1} + t)^{2\alpha} + 2^{-m-4} a^2 (2^m + t)^{2\alpha} + (1+r) \log 2) \\ &= \exp(-2^{-m-5} a^2 ((2^{m+1} + t)^{2\alpha} - 2(2^m + t)^{2\alpha}) + (1+r) \log 2) \\ &\leq \exp(-2^{-m-5} a^2 \eta_1 (2^m + t)^{2\alpha} + (1+r) \log 2) \\ &\leq \exp(-2^{-5} a^2 \eta_1 2^{(2\alpha-1)m} + (1+r) \log 2). \end{aligned} \tag{10}$$

Note that (10) can be made arbitrarily small by having sufficiently large m . Thus, there exists $m(r)$ large enough such that, for $m \geq m(r)$,

$$2^{(1+r)(m+1)} \exp(-2^{-(m+1)-4} a^2 (2^{m+1} + t)^{2\alpha}) \leq \frac{1}{2} 2^{(1+r)m(r)} \exp(-2^{-m-4} a^2 (2^m + t)^{2\alpha}).$$

We now carry out the second step. Based on (9), let

$$\eta_2 := \frac{(2 + \eta_1)^{1/(2\alpha)} - 1}{2 - (2 + \eta_1)^{1/(2\alpha)}}.$$

Then, for large enough t , we have

$$\begin{aligned} &\sum_{m=\lceil \log(\eta_2 t) \rceil + 1}^{\infty} 2^{(1+r)m} \exp(-2^{-m-4} a^2 (2^m + t)^{2\alpha}) \\ &\leq 2^{(1+r)\lceil \log(\eta_2 t) \rceil} \exp(-2^{-\lceil \log(\eta_2 t) \rceil - 4} a^2 (2^{\lceil \log(\eta_2 t) \rceil} + t)^{2\alpha}) \sum_{k=1}^{\infty} \frac{1}{2^k} \\ &\leq 2^{(1+r)\lceil \log(\eta_2 t) \rceil} \exp(-2^{-\lceil \log(\eta_2 t) \rceil - 4} a^2 (2^{\lceil \log(\eta_2 t) \rceil} + t)^{2\alpha}). \end{aligned}$$

Thus,

$$\begin{aligned}
 & \mathbb{E}[2^{\tilde{N}(t)r}] \\
 &= \frac{\sum_{m=0}^{\lceil \log(\eta_2 t) \rceil} 2^{(1+r)m} \exp(-2^{-m-4} a^2 (2^m + t)^{2\alpha}) + \sum_{m=\lceil \log(\eta_2 t) \rceil + 1}^{\infty} 2^{(1+r)m} \exp(-2^{-m-4} a^2 (2^m + t)^{2\alpha})}{\sum_{m=0}^{\infty} 2^m \exp(-2^{-m-4} a^2 (2^m + t)^{2\alpha})} \\
 &\leq \frac{\sum_{m=0}^{\lceil \log(\eta_2 t) \rceil} 2^{(1+r)m} \exp(-2^{-m-4} a^2 (2^m + t)^{2\alpha}) + 2^{(1+r)\lceil \log(\eta_2 t) \rceil} \exp(-2^{-\lceil \log(\eta_2 t) \rceil - 4} a^2 (2^{\lceil \log(\eta_2 t) \rceil} + t)^{2\alpha})}{\sum_{m=0}^{\infty} 2^m \exp(-2^{-m-4} a^2 (2^m + t)^{2\alpha})} \\
 &\leq \frac{2^{r\lceil \log(\eta_2 t) \rceil} \sum_{m=0}^{\lceil \log(\eta_2 t) \rceil} 2^m \exp(-2^{-m-4} a^2 (2^m + t)^{2\alpha}) + 2^{(1+r)\lceil \log(\eta_2 t) \rceil} \exp(-2^{-\lceil \log(\eta_2 t) \rceil - 4} a^2 (2^{\lceil \log(\eta_2 t) \rceil} + t)^{2\alpha})}{\sum_{m=0}^{\lceil \log(\eta_2 t) \rceil} 2^m \exp(-2^{-m-4} a^2 (2^m + t)^{2\alpha})} \\
 &\leq 2^{r\lceil \log(\eta_2 t) \rceil + 1} \\
 &\leq 3\eta_2^r t^r.
 \end{aligned}$$

We are now ready to establish the bound for $\mathbb{E}[(L(n) + 2^{\tilde{N}+1})^r]$ for any fixed $r \geq 1$. Hence,

$$\begin{aligned}
 \mathbb{E}[(L(n) + 2^{\tilde{N}+1})^r] &\leq \mathbb{E}[2^{r-1}(L(n))^r + 2^{(\tilde{N}+1)r}] \\
 &\leq 2^{r-1}\mathbb{E}[L(n)^r] + 2^{2r-1}\mathbb{E}[2^{\tilde{N}r}] \quad (\text{by Jensen's inequality}) \\
 &\leq 2^{r-1}\mathbb{E}[L(n)^r] + 2^{2r-1}3\eta_2^r\mathbb{E}[L(n)^r] \\
 &< \infty.
 \end{aligned}$$

We have thus shown that each path has finite moments of all orders.

As for the acceptance probability in steps (ii), (iii) and (v), we note that

$$\begin{aligned}
 & \mathbb{P}(\{S_k \leq ak^\alpha + bk^{1-\alpha}, 1 \leq k \leq L(n)\} \cap \{\tilde{J} = 1\}) \\
 &= \mathbb{P}(T_1 > L(n), \tilde{T}_{L(n), S_{L(n)}} = \infty) \\
 &= \mathbb{P}(T_1 = \infty) \quad (\text{as } \mathbb{P}(T_1 = \infty \mid S_t, T_1 > t) = \mathbb{P}(\tilde{T}_{t, S_t} = \infty)) \\
 &\geq \frac{3}{4} \quad (\text{by Lemma 1}).
 \end{aligned}$$

Then the number of times a path is rejected is stochastically bounded by a geometric random variable with probability of success $\frac{3}{4}$. Therefore, the total length of paths generated in Procedure B has finite moments of all orders. □

3.1. Numerical experiments

In this section we conduct numerical experiments to analyse the performance of Algorithm 1 for different values of the parameter a . In Remark 1, we briefly discussed 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 the probability of record-breaking is small enough, while keeping $\Gamma(\kappa)$ as small as possible.

For the computational cost, we first note that the choices of a and b will affect the distribution of N , which is the length of trajectory generated in Procedure A. In Procedure B, the values of $\Gamma(\kappa)$, $L(\Gamma(\kappa))$ and the distribution of \tilde{N} also depends on the values of a and b .

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
0.49	7.25	3.06	2.19	2.11
0.499	12.81	3.79	3.49	3.12

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 $g'(0) = 0$, $g''(\theta) = 1/(1 - \theta)^2 - 2$, we have

$$g(\theta) < 0 \quad \text{for all } \theta \in \left(-1, 1 - \frac{\sqrt{2}}{2}\right).$$

Therefore, we can set $\delta' = 1 - \sqrt{2}/2$, 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.

We observe that when a is relatively far away from the upper bound $\frac{1}{2}$ (e.g. $a \leq 0.45$), the running time decreases as a increases. However, as a approaches $\frac{1}{2}$, the running time is increasing in a . This is because $\xi \rightarrow \infty$ as $a \rightarrow \frac{1}{2}$ (see (2)). We also observe that the changing rate of running time regarding a is larger for smaller values of α , which in general implies greater curvature of the nonlinear boundary.

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 independent and identically distributed (i.i.d.). Independent of the arrival process, the service times are also i.i.d. and may have infinite mean.

Suppose that 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 start by introducing a point process representation of an infinite-server queue to facilitate the explanation of the simulation strategy. We mark each arriving customer as a point in a two-dimensional space, where the x -coordinate records its arrival time and the y -coordinate records its service time (service requirement). Figure 2 provides an illustration with two points representing two arriving customers. Customer 1 arrives at $-A_1$ and has a service requirement of V_1 . Note that, as there are infinitely many servers, this customer will enter service immediately upon arrival and will leave the system at time $-A_1 + V_1$. If we draw a minus 45-degree line from $(-A_1, V_1)$, the intersection of this line with the x -axis represents customer 1's departure time. Likewise, we can also denote the departure time of customer 2 by the intersect of the minus 45-degree line starting from $(-A_2, V_2)$ with the x -axis. We observe that in this particular example, customer 1 would leave the system in the interval $[0, h]$, while customer 2 would leave the system before time 0. Based on this observation, we can draw a shaded region in Figure 2, which has the property that all the points (customers) that fall into




FIGURE 2: Point process representation of an infinite-server queue.

this region will leave the system during $[0, h]$. Therefore, to sample the departure process on $[0, h]$, we essentially would like to sample all the points (customers) that fall into the shaded area.

We further divide the shaded area into two parts, namely H and G . Points in the shaded area G are customers that arrive after time 0 and depart before time h , while points in area H are customers who arrive before time 0 and depart between times 0 and h . 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.

For the rest of this section, we explain how to sample all the points (customers) that fall into region H . 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 its service time. Let $A_0 := 0$. We then denote $X_n := A_n - A_{n-1}$ as the interarrival time between the n th arrival and the $(n - 1)$ th arrival. Let $\mu := \mathbb{E}[X]$ denote the mean interarrival time and $\sigma^2 := \text{var}(X)$ denote its variance. 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 .

The following observation builds the foundation of our simulation strategy. Suppose that we can find a random number Ξ such that

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

for $n \geq \Xi$. Then we can sample the point process up to Ξ , i.e. $\{(-A_i, V_i), 1 \leq i \leq \Xi\}$, and find \mathcal{H} . Built on this observation, 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.

(P1) There exists a well-defined random number Ξ_1 such that

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

(P2) There exists a well-defined random number Ξ_2 such that

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

Now, set $\Xi = \max\{\Xi_1, \Xi_2\}$. Then we have $V_n < A_n$ or $V_n > A_n + h$ for $n \geq \Xi$. Note that, based on the introduction of the ϵ_n , we can find Ξ_1 and Ξ_2 separately.

To guarantee that Ξ_1 and Ξ_2 are well defined, i.e. finite, we need to choose ϵ_n that satisfy the following two conditions:

- (C1) $\sum_{n=1}^{\infty} \mathbb{P}(|A_n - n\mu| > \epsilon_n) < \infty$,
- (C2) $\sum_{n=1}^{\infty} \mathbb{P}(V_n \in (n\mu - \epsilon_n, n\mu + \epsilon_n + h)) < \infty$.

Under (C1) and (C2), the Borel–Cantelli lemma guarantees that Ξ_1 and Ξ_2 are finite almost surely.

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

$$f(v) = \beta v^{-(\beta+1)} \mathbf{1}\{v \geq 1\}. \tag{11}$$

We also write $\bar{F}(\cdot)$ as the tail distribution of V_i . We assume the interarrival time has a finite moment generating function in a neighbourhood of the origin. This is without loss of generality. If the interarrival time is heavy tailed, we can simulate a coupled infinite-server queue with truncated interarrival times, $X_i^C = \min\{X_i, C\}$. This coupled infinite-server queue 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.

Set $\epsilon_n = n^\alpha$ for $\frac{1}{2} < \alpha < \beta$. In what follows, we shall show that our choice of ϵ_n satisfies (C1) and (C2), respectively. We shall also explain how to find (simulate) Ξ_1 and Ξ_2 .

4.1. Sampling of the arrival process and Ξ_1

The following lemma verifies (C1).

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

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

Proof. We note 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 the moderate deviation principle [7], we have

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

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

Let $S_n = A_n - n\mu$. We note that both S_n and $-S_n$ are mean zero random walks:

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

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 $T'_{k-1} < \infty$,

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

otherwise, $T'_k = \infty$. 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 $J' \sim \text{Bernoulli}(\mathbb{P}(T'_k = \infty \mid T'_{k-1}))$.
 - (iii) If $J' = 0$, sample $(X_i: i = T'_{k-1} + 1, \dots, T'_k)$ conditional on $T'_k < \infty$ (see Procedure A' below). Set $k = k + 1$ and go back to step (ii); otherwise ($J' = 1$), set $\Xi_1 = T'_{k-1}$ and go to step (iv).
 - (iv) Apply Procedure C (detailed in Section 4.2) iteratively to sample Ξ_2 .
 - (v) 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' below).
-

We also modify Procedure A and Procedure B as follows.

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

- (i) Sample a random time N with PMF (4). Let $\theta_N = a2^{N(\alpha-1)-2}$. Sample $U_1 \sim \text{Uniform}[0, 1]$. If $U_1 \leq \frac{1}{2}$, go to step (ii(a)); otherwise, 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

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

- (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

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

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

Proposition 1. In Procedure A', J' is a Bernoulli random variable with probability of success $\mathbb{P}(T'_1 = \infty)$. 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

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

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

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

- (iii) Sample $L'(n)$ and $(X_{n+1}, \dots, X_{L'(n)})$ under the nominal distribution $\mathbb{P}(\cdot)$. If $|S_{L'(n)}| > aL'(n)^\alpha + bL'(n)^{1-\alpha}$, go back to step (i); otherwise, go to step (iv).
- (iv) Sample \tilde{N} with PMF $p_{L(n)}$ defined in (5). Set $\tilde{\theta}_{\tilde{N}} = 2^{\tilde{N}-2}a(2^{\tilde{N}} + L(n))^\alpha$. Sample $U_1 \sim \text{Uniform}[0, 1]$. If $U_1 < \frac{1}{2}$, go to step (v(a)); otherwise, 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

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

- (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

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

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

$$U_2 \leq \frac{(\exp(\tilde{\theta}_{\tilde{N}}\tilde{S}_{\tilde{T}'} - \tilde{\psi}(\tilde{\theta}_{\tilde{N}}))/2 + \exp(-\tilde{\theta}_{\tilde{N}}\tilde{S}_{\tilde{T}'} - \tilde{\psi}(-\tilde{\theta}_{\tilde{N}}))/2)^{-1}}{p_r(\tilde{N})} \times \mathbf{1}\{\tilde{T}' \in [2^{\tilde{N}}, 2^{\tilde{N}+1})\},$$

set $\tilde{J}' = 0$ and go back to step (i); otherwise, 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

We start by verifying (C2).

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

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

Proof. We have

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

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

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

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

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

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

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

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 the K_i one by one. We achieve this by finding a proper sequence of upper and lower bounds for $\mathbb{P}(K_i = \infty)$. We start with K_1 . Note that

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

Let

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

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

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

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

Then for large enough k^* such that $k^* > (2/\mu)^{1/(\beta-\alpha)}$ and

$$\frac{2(2+h)\beta}{\mu} \frac{1}{k^{*,\beta+1-\alpha}} < 1,$$

we have, for $k > k^*$,

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

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 $l(k) \leq l(k + 1) < \mathbb{P}(K_1 = \infty)$ and $\lim_{k \rightarrow \infty} l(k) = \mathbb{P}(K_1 = \infty)$.

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

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

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 $K_i = \infty$; otherwise, output $K_i = k$.

Once we find the values of the K_i , sampling the V_n conditional on the information of the K_i is straightforward. We summarize the simulation of the service time process together with Ξ in Algorithm 2.

Algorithm 2. Sampling Ξ together with $(V_i : 1 \leq i \leq \Xi)$.

- (i) Initialize $K_0 = 0, i = 1$.
 - (ii) Given the value of $K_{i-1} < \infty$, sample K_i using Procedure C.
 - (iii) If $K_i < \infty$, set $i = i + 1$ and go back to step (ii); otherwise, set $\Xi_2 = K_{i-1}$ and go to step (iv).
 - (iv) Apply Algorithm 1' to find Ξ .
 - (v) Sample $(V_i : i = 1, 2, \dots, \Xi)$ conditional on $(K_1, K_2, \dots, K_{i-1})$ using the acceptance–rejection method with the nominal distribution of the service times as the proposal distribution.
-

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 $\mathbb{P}(\Phi_i < \infty) = 1$, $\mathbb{E}[\Phi_i] = \infty$. Taking Φ_1 as an example,

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

with

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

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

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 the arrivals are simulated sequentially backwards in time, we now argue that the expected time to detect such an arrival must be infinite. Assuming, for simplicity, deterministic interarrival times equal to 1, and letting $-T < 0$ be the time at which such earliest arrival occurs, then we have

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

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

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

4.3. Numerical experiment on the departure process of an M/G/∞ queue

In this section we apply Algorithms 1' and 2 to simulate the steady-state departure process of an infinite-server queue whose service times have infinite mean.

We consider an infinite-server queue having Poisson arrival process with rate 1, and Pareto service time distributions with probability density function

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

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

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 times 0 and 1 for each run and construct the corresponding relative frequency bar plot in Figure 3. We observe that the distribution of simulated departures between times 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 the χ^2 test as goodness of fit tests with the four groups of sampled data against the Poisson distribution. The corresponding p -values are 0.2404, 0.2589, 0.4835, and 0.1137, respectively. Therefore, the tests fail to reject that the generated samples are Poisson distributed.

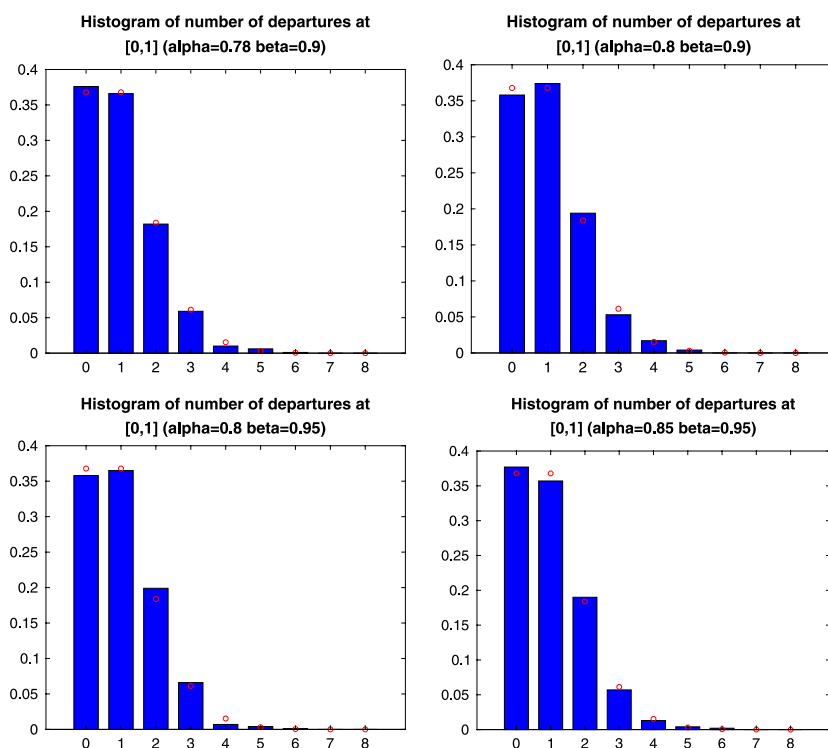


FIGURE 3: Histograms comparison for sampled departure.

References

- [1] ASMUSSEN, S. (2003). *Applied Probability and Queues*, 2nd edn, Springer, New York.
- [2] BLANCHET, J. AND CHEN, X. (2015). Steady-state simulation of reflected Brownian motion and related stochastic networks. *Ann. Appl. Probab.* **25** (6), 3209–3250.
- [3] BLANCHET, J. AND DONG, J. (2014). Perfect sampling for infinite server and loss systems. *Adv. Appl. Prob.* **47** (3), 761–786.
- [4] BLANCHET, J. AND WALLWATER, A. (2015). Exact sampling for the steady-state waiting times of a heavy-tailed single server queue. *ACM Trans. Model. Comput. Simul.* **25** (4), art. 26.
- [5] BLANCHET, J., DONG, J. AND PEI, Y. (2018). Perfect sampling of GI/GI/c queues. *Queueing Systems* **90** (1–2), 1–33.
- [6] ENSOR, K. AND GLYNN, P. (2000). Simulating the maximum of a random walk. *J. Statist. Plann. Inference* **85**, 127–135.
- [7] GANESH, A., O’CONNELL, N. AND WISCHIK, D. (2004). *Big Queues* (Lecture Notes Math. **1838**). Springer, Berlin.
- [8] KENDALL, W. S. (1998). Perfect simulation for the area-interaction point process. In *Probability Towards 2000* (Lecture Notes Statist. **128**), pp. 218–234. Springer, New York.
- [9] PROPP, J. AND WILSON, D. (1996). Exact sampling with coupled Markov chains and applications to statistical mechanics. *Random Struct. Algorithms* **9**, 223–252.