Stationary models of adversarial coupon collection with an application to moving-target defense

GEORGE KESIDIS*
EE & CSE
Penn State University
University Park, PA, USA
gik2@psu.edu

TAKIS KONSTANTOPOULOS
Mathematics
University of Liverpool
Liverpool, UK
takiskonst@gmail.com

MICHAEL ZAZANIS
Statistics
AUEB
Athens, Greece
zazanis@aueb.gr

Abstract

In cybersecurity, moving-target defense (MTD) mitigates reconnaissance attacks from botnets by periodically resetting server identities, disrupting adversarial intelligence gathering. We model this defense mechanism as a dynamic coupon collection process, where servers correspond to coupon types and botnet attacks to coupon acquisitions. The defender's intervention (removing collected coupons) represents server resets, restoring their uncompromised state. In this paper, we derive various results regarding the number of compromised servers under various stochastic attack-defense dynamics.

1 Introduction

We consider an application of moving-target defense (MTD) in a cybersecurity setting [4, 5, 14, 7, 11]. Botnets conducting reconnaissance aim to identify as many different servers as possible before launching distributed denial-of-service (DDoS) attacks, wherein they overwhelm servers, rendering them inoperable for legitimate users. To mitigate this threat, MTD periodically resets server identities (IP addresses), disrupting the botnet's intelligence-gathering process and increasing the attacker's work factor. Changing the IP address of a server may require redirection messaging to its active clients. The connection identifier of the QUIC transport protocol [8] can persist when an end-host's IP address changes. Simple bots may not be able to process redirection messages or they may not support protocols such as QUIC.

We model the foregoing interaction between attacker and defender as a dynamic coupon collection problem, where coupon types correspond to distinct servers, bots act as collectors, and acquiring a coupon represents a successful reconnaissance event. The defender periodically removes collected coupons, corresponding to identity resets that restore servers to an uncompromised state. To ensure unpredictability in the defensive strategy, reset events follow a Poisson process.

This formulation provides a probabilistic framework for analyzing the number of servers that remain uncompromised under various attack-defense strategies. We primarily study the stationary distribution of collected identities and, at times, the distribution of the number of coupons for each type. The problem connects to queueing models with "catastrophic" (or "disaster") resets [2], where accumulated service content is periodically flushed, as also observed in population dynamics [3] and some emerging applications (e.g., voiding a user cache upon teleportation in Virtual Reality). By varying assumptions on attack-defense mechanisms, we derive insights into optimal strategies for mitigating botnet reconnaissance and strengthening cyber resilience. This work establishes a mathematical model for quantitatively assessing MTD approaches, leveraging probabilistic structures to enhance adaptive defensive measures against evolving threats.

*Corresponding author.

The paper is organized as follows. We first interpret some known models in terms of coupon collecting used for MTD: Section 2 presents the case where coupon types are reset independently, while Section 3 analyzes simultaneous resets of all coupon types. An essential contribution of this paper is in Section 4, where we consider simultaneous coupon selection in fixed-size batches with synchronous resets; here, the distribution of the number of compromised servers involves Stirling numbers of the second kind. Section 5 presents numerical examples illustrating key model behaviors. Finally, we conclude with a brief summary in Section 6.

2 Fully asynchronous resetting of coupons

There are M coupon types, each available in unlimited supply. In the classical coupon collection problem, coupons are gathered until all types have been obtained. See [6] for sharp asymptotics related to this problem.

Here, coupons are continuously collected over time. Specifically, coupons of type i are gathered at the points of a renewal process Ξ_i. The times between successive points of Ξ_i are i.i.d. copies of the positive random variable X_i, with finite expectation $EX_i = 1/\lambda_i$. Coupon type i represents a server, while the process Ξ_i corresponds to an adversary launching attacks against that server. Each adversary maintains a buffer where the selected coupons are stored.

There is, in addition, a coupon reset process Φ_i. At the points of Φ_i, all coupons stored in the buffer of adversary i are immediately cleared. We assume that Φ_i is a Poisson process with times between successive points that are i.i.d. copies of an exponential random variable S_i with expectation $ES_i = 1/\mu < \infty$.

It is assumed that the $2M$ point processes Ξ_i, Φ_i, $i = 1, \ldots, M$, are independent.

The dynamics of the system is implicit in the above discussion. Let

$$\begin{aligned} &W_i(t) := \text{number of coupons present at time } t \text{ in the buffer of adversary } i \\ &Y_i(t) = \mathbf{1}\{W_i(t) > 0\}, \quad Y(t) := \sum_i \mathbf{1}\{W_i(t) > 0\}. \end{aligned} \tag{1}$$

It is clear, from classical applied probability arguments, that there is a stationary process $W_i(t)$, $t \in \mathbb{R}$, satisfying the system dynamics. We shall consider this process. Then $Y(t)$, $t \in \mathbb{R}$, is also stationary. Note that $Y(t)$ is the number of servers that are compromised at time t. Alternatively, $Y(t)$ is the number of coupon-types collected at time t. We refer to $Y(t)$ as the *collection-type size* at time t.

2.1 The distribution of the collection-type size

Consider the stationary process $Y(t)$, $t \in \mathbb{R}$ and let Y denote a random variable whose distribution is the distribution of $Y(t)$ for some (and hence all) t. It may be worth noting that, for each $1 \le i \le M$, the $\{0,1\}$-valued process $Y_i(t)$, $t \in \mathbb{R}$, has the distribution of the buffer contents of a stationary GI/M/1/1 queue.

Proposition 1. *For all $z \in \mathbb{R}$,*

$$Ez^Y = \prod_{i=1}^{M} \frac{E\min\{S_i, X_i\}z + E(X_i - S_i)^+}{EX_i} \tag{2}$$

Proof. Fix $1 \le i \le M$. Let $T_i(1)$ be the first positive point of Ξ_i and S_i' the first positive point of Φ_i. Let Y_i be a random variable distributed like $Y_i(t)$ for some (and hence all) t. Let P_{Ξ_i} be the Palm probability with

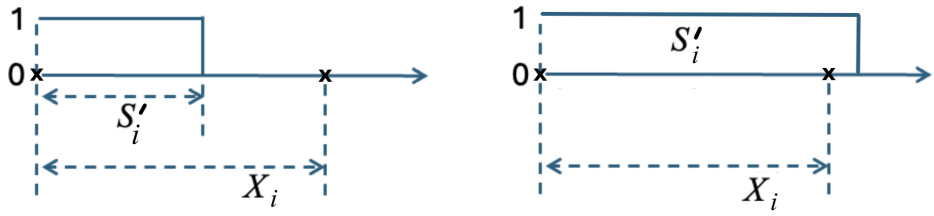

Figure 1: $Y_i(t)$ over the interval $[0, X_i]$ for cases where $S_i' < X_i$ (left) and $S_i' > X_i$ (right).

respect to Ξ_i and E_{Ξ_i} the expectation with respect to it. By the Palm inversion formula [1],

$$Ez^{Y_i} = \lambda_i E_{\Xi_i} \int_0^{T_i(1)} z^{Y_i(s)} ds = \lambda_i E_{\Xi_i} \int_0^{X_i} z^{Y_i(s)} ds,$$

where the second equality follows from the fact that $P_{\Xi_i}(T_i(1) = X_i) = 1$. We have

$$Y_i(s) = \mathbf{1}\{s \le T_i(1) \wedge S_i'\}, \quad 0 \le s \le T_i(1).$$

Combining the above and using the fact that S_i' has the law of S_i (by virtue of the memoryless property of the exponential distribution) we obtain

$$Ez^{Y_i} = \lambda_i E(X_i z \mid S_i' > X_i) P(S_i' > X_i) + \lambda_i E(S_i' z + (X_i - S_i') \mid S_i' < X_i) P(S_i' < X_i)$$

See Figure 1. Since $Y = \sum_i Y_i$ is the sum of M independent random variables, we obtain (2) after a little algebra. □

Corollary 2.

$$EY = \sum_{i=1}^{M} \frac{E \min\{S_i, X_i\}}{EX_i} \tag{3}$$

Proof. We use the simple fact that $EY = \lim_{z \to 1} \frac{d}{dz} Ez^Y$ and the expression (2). □

Corollary 3. *Let*

$$c_i := \frac{E(X_i - S_i)^+}{E \min\{S_i, X_i\}}, \quad C := \prod_{i=1}^{M} \lambda_i E \min\{S_i, X_i\}.$$

Then

$$P(Y = M - k) = C \sum_{0 \le i_1 < \cdots < i_k \le M} c_{i_1} \cdots c_{i_k}, \quad 0 \le k \le M.$$

In particular, if the Ξ_i are identically distributed and the Φ_i are identically distributed then $M - Y$ is binomially distributed. The second assertion follows immediately when we take into account that the c_i are identical under the identical distribution assumptions.

Proof. Expanding the product of monomials on the right-hand side of (2) we obtain

$$Ez^Y = C \prod_{i=1}^{M} (z + c_i) = C \sum_{k=0}^{M} \left(\sum_{1 \le i_1 < \cdots < i_k \le M} c_{i_1} \ldots c_{i_k} \right) z^{M-k},$$

from which the probabilities can be read off directly. □

2.2 The number of coupons collected of a particular type

We now examine the number of collected coupons for each type, motivated by scenarios involving multiple independent collectors or a single collector requiring more than one coupon per type. In the context of a security application, one may interpret this as server i being rendered inoperable once the number of attacks (i.e., collected coupons) surpasses a predefined threshold.

Let $T_i(k)$ be the k^{th} point of Ξ_i and define the number of coupons of type i in the buffer of adversary i just after the k^{th} point of Ξ_i,

$$w_i(k) := W_i(T_i(k)+).$$

Hence $w_i(k) \in \mathbb{N} = \{1, 2, \ldots\}$. Clearly, $w_i(k)$, $k = 1, 2, \ldots$, is Markovian. Its (nonzero) transition probabilities are

$$\begin{aligned} P(w_i(k+1) = w_i(k) + 1 \mid w_i(k)) &= P(S_i > X_i) = G_i(\mu_i) \\ P(w_i(k+1) = 1 \mid w_i(k)) &= P(S_i < X_i) = 1 - G_i(\mu_i), \end{aligned}$$

where

$$G_i(v) := E\mathrm{e}^{-vX_i}.$$

It is then clear that the stationary distribution π_i of w_i is geometric.

Proposition 4.

$$\pi_i(n) = (G_i(\mu_i))^{n-1}(1 - G_i(\mu_i)), \quad n \in \mathbb{N}.$$

From this, we easily obtain by conditioning on whether $Y_i = 0$ (see Corollary 2):

Corollary 5. *Let W_i be a random variable whose law is the law of $W_i(t)$ for some (and hence all) t. Then*

$$\begin{aligned} P(W_i = n) &= \pi_i(n)\frac{E\min\{X_i, S_i\}}{EX_i}, \quad n \geq 1 \\ P(W_i = 0) &= \frac{E(X_i - S_i)^+}{EX_i} \end{aligned}$$

Thus, the distribution of W_i is a mixture of a geometric law and a unit mass at 0.

3 Fully synchronous resetting of coupons

We now consider a model in which all coupons are reset simultaneously at the points of a Poisson process Φ with rate μ, thereby enhancing the effectiveness of the defense system.

To ensure tractability, we assume that coupon selection takes place at the points of a single renewal process Ξ, where coupon i is selected with probability α_i. We have

$$\sum_{i=1}^{M} \alpha_i = 1.$$

We now assume that Ξ has rate $M\lambda$.

The statement at the beginning of Section 2.1 regarding the independence of different coupon types no longer holds.

3.1 The collection-type set

As before, we let $W_i(t)$ be the number of coupons of type i in the buffer of adversary i at time t. Let $Y(t)$ be the collection-type size at time t. In general, $Y(t)$, $t \in \mathbb{R}$, is not Markovian. We are thus forced to consider the *collection-type set*

$$\mathsf{Y}(t) := \{i : 1 \le i \le M,\, W_i(t) > 0\} \subset \{1, \ldots, M\},$$

so that

$$Y(t) = |\mathsf{Y}(t)|.$$

Moreover, let $T(k)$ be the k^{th} positive point of Ξ and define

$$\mathsf{y}(k) := \mathsf{Y}(T(k)+) \subset \{1, \ldots, M\} \setminus \{\varnothing\}.$$

Let X denote a random variable distributed as $T(k+1) - T(k)$ and S an independent exponential random variable with rate μ. Set

$$H(\theta) := Ee^{-\theta X}$$

for the Laplace transform of X and let

$$\alpha(\mathsf{z}) := \sum_{i \in \mathsf{z}} \alpha_i, \quad \mathsf{z} \subset \{1, \ldots, M\}.$$

We then have the following.

Proposition 6. *The random sequence* $\mathsf{y}(k)$, $k \in \mathbb{Z}$, *is a set-valued Markov chain. Its transition probabilities are given by*

$$\begin{aligned}
&P(\mathsf{y}(k+1) = \mathsf{y}(k) \mid \mathsf{y}(k)) = H(\mu)\, \alpha(\mathsf{y}(k)), \quad \textit{if } |\mathsf{y}(k)| > 1,\\
&P(\mathsf{y}(k+1) = \{i\} \mid \mathsf{y}(k) = \{i\}) = \alpha_i,\\
&P(\mathsf{y}(k+1) = \{i\} \mid \mathsf{y}(k)) = (1 - H(\mu))\, \alpha_i, \quad \textit{if } \mathsf{y}(k) \neq \{i\},\\
&P(\mathsf{y}(k+1) = \mathsf{y}(k) \cup \{i\} \mid \mathsf{y}(k)) = H(\mu)\alpha_i, \quad \textit{if } i \notin \mathsf{y}(k),
\end{aligned}$$

for $1 \le i \le M$. *There is a unique stationary distribution* π *for this Markov chain.*

Sketch of proof. Consider the collection-type size $\mathsf{y}(k)$ just after $T(k)$. If $\mathsf{y}(k)$ is not a singleton and if there is no coupon resetting (which occurs with probability $P(S > X) = H(\mu)$) between $T(k)$ and $T(k+1)$ then $\mathsf{y}(k+1) = \mathsf{y}(k)$ if and only of the coupon type collected at time $T(k+1)$ is an element of the set $\mathsf{y}(k)$; the probability of the latter is $\alpha(\mathsf{y}(k))$. This explains the first line. The remaining lines follow a similar argument. The Markovian property of $\mathsf{y}(k)$, $k \in \mathbb{Z}$, is obvious. The uniqueness of the stationary distribution follows from the finiteness of the state space of the chain together with the irreducibility of its transition probabilities. □

We do not attempt to obtain the stationary distribution of the above Markov chain in closed form. The above proposition allows us of course to compute it numerically. However, we easily obtain the following.

Corollary 7. *If* Y *is a random variable distributed as* $|\mathsf{Y}(t)|$ *for some (and hence all)* t *under stationarity, we have*

$$P(Y = j) = \sum_{|\mathsf{y}|=j} \pi(\mathsf{y})\, \frac{E \min\{S, X\}}{EX},\ 1 \le j \le m \tag{4}$$

$$P(Y = 0) = \frac{E(X - S)^+}{EX} \tag{5}$$

The reason that

$$y(k) = |\mathsf{y}(k)|, \quad k \in \mathbb{Z},$$

fails to be Markovian, in general, is due to asymmetry. The asymmetry vanishes when all the α_i are identical.

3.2 The collection-type size in the fully synchronous symmetric system

Proposition 8. *Assume that the system is fully symmetric, that is,*

$$\alpha_i = \frac{1}{M}, \quad 1 \le i \le M.$$

Then $y(k)$, $k \in \mathbb{Z}$, is Markovian. Its transition probabilities are given by

$$\begin{array}{lcll} P_{i,i} & = & H(\mu)\frac{i}{M} + (1 - H(\mu))\mathbf{1}\{i = 1\} & i = 1, \ldots, M \\ P_{i,i+1} & = & H(\mu)\left(1 - \frac{i}{M}\right) & i = 1, \ldots, M-1 \\ P_{i,1} & = & 1 - H(\mu), & i = 2, \ldots, M \end{array}$$

Solving the balance equations immediately yields the the stationary distribution.

Proposition 9. *Under the fully symmetric conditions, let y be a random variable distributed as the collection-type size process at a fixed (and hence) any k.*

$$P(y = i) = \frac{\left(\frac{H(\mu)}{M}\right)^{i-1} \frac{(M-1)!}{(M-i)!} \prod_{j=i+1}^{M} \left(1 - H(\mu)\frac{j}{M}\right)}{\prod_{j=1}^{M-1} \left(1 - H(\mu)\frac{j}{M}\right)}, \quad i = 1, \ldots, M$$

with $\prod_{j=M+1}^{M}(...) = 1$.

Remark 10. We can easily pass on from the stationary collection-type size distribution $P(y = i)$ at the collection epochs to the stationary collection-type size distribution $P(Y = i)$ at an arbitrary point of time by Palm inversion as, e.g., in Corollary 7:

$$\begin{array}{lcll} P(Y = 0) & = & 1 - \dfrac{M\lambda}{\mu}(1 - H(\mu)) & \\ P(Y = i) & = & \dfrac{M\lambda}{\mu}(1 - H(\mu))P(y = i), & i = 1, \ldots, M \end{array}$$

Remark 11. If, in addition to symmetry, the collection-type size process Ξ is also Poisson (with rate $M\lambda$) then $H(\mu) = M\lambda/(M\lambda + \mu)$ and Y has identical law as y.

4 Batch coupon selections and fully synchronous resetting

So far, we assumed that exactly one coupon is collected at each collection time, that is, at each point of Ξ. Suppose now that N coupons are simultaneously collected, uniformly at random.

Assume that Ξ is a renewal process with rate λ and that Φ is an independent Poisson process with rate μ.

Since we are interested in the distribution of the collection-type size, we first compute the distribution of the size of different coupon types collected at a given collection time. Combinatorially, we have are faced with the problem of N unordered selections with repetition from M ordered items (the coupon types).

Proposition 12. *Let M, N be two unrelated positive integers. An urn contains M coupon types with labels $1, \dots, M$. We select N items, with repetition, and without caring about the order, uniformly at random. Let ζ be the number of different coupon types in the sample, a random integer in the set $\{1, \dots, N \wedge M\}$. Then*

$$P(\zeta = r) = \frac{S(N,r)\,(M)_r}{M^N} =: P_{M,N}(r), \quad 1 \le r \le M \wedge N, \tag{6}$$

where $(M)_r = M(M-1)\cdots(M-r+1)$ and $S(N,r)$ is the Stirling number of the second kind, that is,

$$S(N,r) = \text{ number of partitions of a set of size } N \text{ in } r \text{ parts}.$$

Moreover,

$$E\zeta = M\left[1 - (1 - \tfrac{1}{M})^N\right] =: b_{M,N} \quad (\text{with } b_{M,0} = 0). \tag{7}$$

Proof. We prove (7) first. The expectation of ζ equals M times the probability that a coupon of specific type is selected; this probability equals $1 - (1 - \frac{1}{M})^N$.

There many ways to prove (6). One way is to note that $r!S(N,r)$ is the number of surjective functions from $\{1, \dots, N\}$ onto a set of size r. This has to multiplied by $\binom{M}{r}$–the number of subsets of $\{1, \dots, M\}$ of size r. Multiplying these together and dividing by M^N–the total number of functions from $\{1, \dots, N\}$ into $\{1, \dots, M\}$–gives (6). □

If we denote explicitly the dependence of ζ on N and call it ζ_N then it is clear that ζ_N, $N = 1, 2, \dots$ is a Markov chain with $P(\zeta_{N+1} = r | \zeta_N = r) = r/M$ and $P(\zeta_{N+1} = r+1 | \zeta_N = r) = 1 - r/M$. Based on this, we can derive a recursion for $\phi_N(t) = Et^{\zeta_N}$:

$$\phi_{N+1}(t) = t\phi_N(t) + \tfrac{1}{M}t(t-1)\phi_N'(t)$$

that can, in principle, be iterated starting from $\phi_1(t) = 1$. But there is no way to get an explicit formula.

Remark 13. As a consequence of the fact that $P_{M,N}$, given by (6), is a probability distribution, we obtain that $\sum_{r \ge 1} S(N,r)\,(M)_r = M^N$. See also the classic paper [12] of Rota and [10]. This is known to hold even if M is replaced by a positive real number θ. So $P_{\theta,N}$ is a probability distribution for all $\theta > 0$ and $N \in \mathbb{N}$. It would thus be tempting to call $P_{\theta,N}$ a Stirling distribution of the second kind. However, the term is already taken: In, e.g., [9, §4.12.3], the term refers to the law of the sum of n independent Poisson random variables conditional on each being strictly positive.

As before, let $Y(t)$ be the collection-type size at $t \in \mathbb{R}$, and let $y(k) = Y(T(k)+)$, where $T(k)$ is the k^{th} point of Ξ.

Proposition 14. *The random sequence $y(k)$, $k \in \mathbb{Z}$, is Markovian. Its transition probabilities are given by*

$$P(y(k+1) = j \mid y(k)) = H(\mu)\, Q_{M,N}(j - y(k) | y(k))\, 1_{y(k) \le j \le M} + (1 - H(\mu))\, Q_{M,N}(j|0), \tag{8}$$

where, for $y \le M$ and $\ell \le \min\{N, M-y\}$,

$$\begin{aligned} Q_{M,N}(\ell|y) &= \frac{\sum_{i=\ell}^{N} \binom{N}{i} (1 - \frac{y}{M})^i (\frac{y}{M})^{N-i} P_{M-y,i}(\ell)}{\sum_{\ell'=0}^{N \wedge (M-y)} \sum_{i'=\ell'}^{N} \binom{N}{i'} (1 - \frac{y}{M})^{i'} (\frac{y}{M})^{N-i'} P_{M-y,i'}(\ell')}, \\ Q_{M,N}(\ell|0) &= P_{M,N}(\ell). \end{aligned} \tag{9}$$

The demonstration of the above is routine, provided that we observe that the quantity $Q_{M,N}(\ell|y)$ is simply the probability that ℓ new coupon types are collected in addition to the existing y.

Remark 15. The above Markov chain is irreducible in a finite set. So it has a unique stationary distribution which can, in principle, be derived from (8), (9). The algebra is too unwieldy, so we resort into a numerical solution of the equations, which is the subject of the last section. From the distribution of y we obtain the distribution of Y using (4) and (5). In particular, if Ξ is Poisson, the two distributions coincide.

Proposition 16. *For the model of this section with a renewal process Ξ, the number of coupon types collected in stationarity are given by*

$$\begin{aligned} P(Y(0)=0) &= 1-\frac{\lambda}{\mu}[1-H(\mu)] \\ P(Y(0)=r) &= \sum_{k=1}^{\infty}\frac{\lambda}{\mu}[1-H(\mu)]^2(H(\mu))^{k-1}\,P_{M,kN}(r),\quad r=1,2,\ldots,M. \end{aligned} \tag{10}$$

In particular, the expected number of coupon types in stationarity is

$$EY(0)=\frac{\lambda}{\mu}M[1-H(\mu)]\,\frac{1-\left(1-\frac{1}{M}\right)^N}{1-H(\mu)\left(1-\frac{1}{M}\right)^N}. \tag{11}$$

Proof. Since Ξ and Φ are two independent, stationary point processes (the first a renewal proces and the second a Poisson process with rate μ) at time 0 the number of different coupon types that have been collected is determined by considering the number of collection epochs (of Ξ) in the interval of time between the latest reset epoch (of Φ) prior to 0 and 0. If we denote the points of Ξ by $\{T(n)\}_{n\in\mathbb{Z}}$ and those of Φ by $\{U(n)\}_{n\in\mathbb{Z}}$ and using the standard convention that they are numbered so that, $P-$a.s., $T(0)\le 0<T(1)$ and $U(0)\le 0<U(1)$,

$$p_k=P(T(-k)<U(0)<T(-k+1)<0),\quad k=1,2,\ldots,\quad \text{and } p_0=P(T(0)<U(0)\le 0).$$

By time reversibility, the independence of the two point processes, and the fact that both $-U(0)$ and $U(1)$ are exponential with rate μ and independent, we have equivalently

$$\begin{aligned} p_k &= P(T(k)<U(1)<T(k+1))=E\int_0^\infty \mathbf{1}\{T(k)<t<T(k+1)\}\mu e^{-\mu t}dt \\ &= E\int_0^\infty \mathbf{1}\{T(k+1)>t\}\mu e^{-\mu t}dt - E\int_0^\infty \mathbf{1}\{T(k)>t\}\mu e^{-\mu t}dt \\ &= Ee^{-\mu T(k+1)}-Ee^{-\mu T(k)}=\frac{\lambda}{\mu}[1-H(\mu)](H(\mu))^{k-1}-\frac{\lambda}{\mu}[1-H(\mu)](H(\mu))^{k} \\ &= \frac{\lambda}{\mu}[1-H(\mu)]^2(H(\mu))^{k-1},\quad k=1,2,\ldots. \end{aligned}$$

In the above string of equalities we have also used the fact that Ξ is a *stationary* renewal process with mean time between points λ^{-1} and therefore the first point to the right of 0 has Laplace transform $Ee^{-sT(1)}=\frac{\lambda}{s}[1-H(s)]$. Similarly,

$$p_0=1-\frac{\lambda}{\mu}[1-H(\mu)].$$

Note that, on the event $\{T(-k)<U(0)<T(-k+1)<T(0)<0\}$, there are k coupon collection epochs following the last reset and hence the probability of collecting r different types for coupons is given by $P_{M,kN}(r)$ as given in (6). These considerations establish (10).

Arguing similarly and using (7) we obtain

$$EY(0) = \sum_{k=1}^{\infty}p_k b_{M,Nk}=\sum_{k=1}^{\infty}\frac{\lambda}{\mu}[1-H(\mu)]^2(H(\mu))^{k-1}M\left[1-\left(1-\frac{1}{M}\right)^{Nk}\right].$$

Evaluating the above geometric series establishes (11). □

Remark 17. When the process Ξ is Poisson (with rate λ) the expression in (11) becomes

$$EY(0) = \frac{M\lambda\left(1-\left(1-\frac{1}{M}\right)^N\right)}{\lambda+\mu-\lambda\left(1-\frac{1}{M}\right)^N}.$$

Remark 18. Suppose we take $N \to \infty$. In this extreme case, all M types of coupons are collected at every collection point of Ξ. That is, in the proof, $b_{M,iN} = M$ for all $1 \le i \le k$ (and still $b_{M,0} = 0$). So, the Palm inversion formula gives

$$EY = \mu\left[0 \cdot E\min\{X_1, S\} + M \cdot E(S - X_1)^+\right] = \mu M \frac{\lambda}{\mu(\mu+\lambda)} = M\frac{\lambda}{\mu+\lambda},$$

with $S \sim \exp(\mu)$ and $X_1 \sim \exp(\lambda)$ independent. This is equal to the expression for EY in the statement of Prop. 16 after taking $N \to \infty$.

5 Numerical example

We conducted numerical experiments to evaluate various specific scenarios related to the cases described above. Figure 2 illustrates a representative example of a hybrid scenario in which, at coupon-type reset epochs, each coupon type is independently reset with probability q at the ticks of a Poisson process (Remark 15). Simultaneously, the collector selects fixed-size batches of coupons (N) at a constant rate (Section 4)[1].

In Figure 2, the collector selects batches of $N = 10$ coupons at a fixed rate of $\lambda = 1$ batch selections per second ($\lambda = 1/EX$). The system comprises $M = 100$ coupon types, with each type reset independently with probability $q = 0.01$ at the ticks of a Poisson process with mean rate $\mu = 1/ES$, which varies along the x-axis of the figure. Consequently, the total mean coupon selection rate is $\lambda N = 10$ coupons per second, while the mean coupon-type reset rate is $q\mu = \mu/10$ types per second.

The figure reveals a "phase transition" interval in which the coupon-type reset process must exceed a certain threshold rate to successfully "defeat" the collector. A similar phase transition is observed when $q\mu$ is held constant while varying the collection rate $N\lambda$.

Finally, consider the model presented in Section 2, *i.e.*, individual random coupon selections occurring at the ticks of a renewal process with rate λ, alongside fully asynchronous resets of the M coupon types at the ticks of independent Poisson processes with rate μ/M. If we take the special case where coupon selection follows a Poisson process, the problem reduces to an independent M/M/1/1 queue for each coupon type. In this setting, the expected number of collected coupons is $EY = M(\lambda/M)/(\lambda/M + \mu/M) = M\lambda/(\lambda+\mu)$, implying that EY is simply a convex function of μ. Thus, unlike Figure 2, no phase transition is observed in this case.

6 Summary

In summary, we analyzed various models of coupon collection and coupon-type resetting, focusing on the stationary distribution—particularly the stationary mean number of collected coupon types and the stationary mean number of coupons collected per type. The coupon-type resetting (clearance) processes were Poisson, and independent of each other and of the renewal collection process. Our models incorporated both synchronous (batch) and asynchronous mechanisms for coupon collection and coupon-type resetting. The main contribution was analysis of a case with batch coupon collection involving a Stirling distribution of the second kind. We concluded with a numerical example illustrating batch coupon collection, which involved a

[1] See [13] for an approximation of $S(N, r)$ as $N \to \infty$ that is uniformly accurate over r.

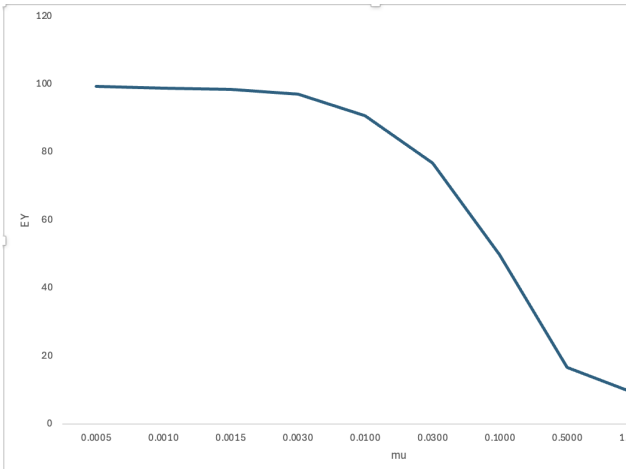

Figure 2: Typical simulation result showing the stationary mean (EY) of the number of collected coupon types (Y). Here, the collector process has fixed rate $\lambda = 1$ batches/s where each batch size is of $N = 10$. Each coupon type is randomly reset with probability $q = 0.1$ at the ticks of a Poisson process with mean rate μ (x-axis). The Stirling distribution (6) was sampled and the total simulation time was 10^6s.

Stirling distribution of the second kind. Some of these models exhibited a "phase transition" in the expected number of collected coupons as the mean collection rate to reset rate decreased.

References

[1] F. Baccelli and P. Brémaud. *Elements of Queueing Theory*. Springer-Verlag, Application of Mathematics: Stochastic Modelling and Applied Probability, No. 26, New York, NY, 1991.

[2] O.J. Boxma, D. Perry, and W. Stadje. Clearing models for M/G/1 queues. *Queueing Syst.*, 38:287–306, 2001.

[3] P.J. Brockwell, J. Gani, and S.I. Resnick. Birth, immigration and catastrophe processes. *Adv. Appl. Prob.*, 14:709–731, 1982.

[4] T.E. Carroll, M. Crouse, E.W. Fulp, and K.S. Berenhaut. Analysis of network address shuffling as a moving target defense. In *Proc. IEEE ICC*, 2014.

[5] M. Crouse, B. Prosser, and E.W. Fulp. Probabilistic performance analysis of moving target and deception reconnaissance defenses. In *Proc. ACM Workshop on Moving Target Defense (MTD)*, page 21–29, 2015.

[6] P. Erdős and A. Rényi. On a classical problem of probability theory. *Acta Math. Acad. Sci. Hung.*, 12:261—267, 1961.

[7] D. Fleck, G. Kesidis, T. Konstantopoulos, N. Nasiriani, Y. Shan, and A. Stavrou. Moving-target defense against botnet reconnaissance and an adversarial coupon-collection model. In *Proc. IEEE Conf. Dependable & Secure Computing (DSC)*, pages 1–8, 2018.

[8] Jana Iyengar and Martin Thomson. QUIC: A UDP-Based Multiplexed and Secure Transport. https://datatracker.ietf.org/doc/rfc9000/, 2022-02-19.

[9] N.L. Johnson, A.W. Kemp, and S. Kotz. *Univariate Discrete Distributons*. Wiley, 3rd edition, 2005.

[10] G. Kesidis, T. Konstantopoulos, and M.A. Zazanis. The generating functions of Stirling numbers of the second kind derived probabilistically. *Mathematical Scientist*, 43(2):82–87, 2018.

[11] G. Kesidis, Y. Shan, D. Fleck, A. Stavrou, and T. Konstantopoulos. An adversarial coupon-collector model of asynchronous moving-target defense against botnet reconnaissance. In *Proc. Intern. Conf. on Malicious and Unwanted Software (MALWARE)*, pages 61–67, 2018.

[12] G.-C. Rota. The number of partitions of a set. *Amer. Math. Monthly*, 71(5):498–504, 1964.

[13] N.M. Temme. Asymptotic estimates of Stirling numbers. *Stud. Appl. Math.*, 89(3):233–243, 1993.

[14] S. Venkatesan, M. Albanese, K. Amin, S. Jajodia, and M. Wright. A moving target defense approach to mitigate DDoS attacks against proxy-based architectures. *Proc. IEEE Conf. Comm. Netw. Security (CNS)*, pages 198–206, 2016.