

# A Monitoring System for Detecting Repeated Packets with Applications to Computer Worms<sup>‡</sup>

M. Vargas Martin\*

J.-M. Robert<sup>†</sup>

P.C. van Oorschot\*

## Abstract

We present a monitoring system to detect worm propagation using Bloom filters with counters. The system is based on stateful analysis of network traffic in routers of a network. Our preliminary evaluation of the system involved real traffic from our internal lab and a well known DARPA data set. After appropriate configuration, no false alarms are obtained under these data sets. We also conduct simulations using real Internet Service Provider topologies with real link delays and simulated traffic. These simulations confirm that this approach can detect worms at early stages of propagation. We believe our approach, with minor adaptations, is of independent interest for use in a number of network applications which benefit from detecting repeated packets, beyond detecting worm propagation. These include detecting network anomalies such as dangerous traffic fluctuations, abusive use of certain services, and distributed denial-of-service attacks.

## 1 Introduction

Worms, unlike other malicious code such as viruses and trojans, are capable of self-propagating [4]. The constant threat of worms to the network infrastructure has become one of the major concerns of network equipment designers. In this paper we propose a router-based monitoring system using Bloom filters with counters (BFWC) [16] to automatically detect worms in the propagation phase. In addition, the detection mechanism provides sufficient information – e.g. identifying a port number or packet signature – to allow some possible reactions. (In contrast, e.g. detecting only that “*there is a UDP flood attack in progress*” would not suffice to allow an appropriate reaction.) The monitoring system can also detect abnormal behaviour in the network by identifying the application responsible for sending significant amounts of repeated packets. In this paper, we focus our attention largely on worm detection.

A large class of worms can be characterized by self-propagating code that selects victims by generating random IP addresses, yielding a propagation pattern giving the impression of a “fan.” For example, the Slammer worm [30] arrives at the victim network in one malicious UDP packet; once the infection is completed an infected host sends the malicious packet to as many random IP addresses as possible. This paper addresses the problem of how to efficiently detect such a propagation behaviour locally, and how to locate the monitoring mechanisms in the network infrastructure to use as few monitors as possible.

Our contribution is a monitoring system based on stateful analysis of network traffic in routers, using an extension of Bloom filters [5]. The idea is to use BFWCs [16] to count the number of times each packet has been forwarded locally. If one packet is forwarded more than a specified threshold within a given fixed period of time, the BFWC triggers an alarm. Bloom filters can indeed be implemented in hardware without considerable performance degradation in routers [13, 36]. We validate the system using real network traffic, and show that the number of false alarms may be reduced by filtering out a small number of common pre-identified non-worm repeated packets, such that they are not processed by the BFWC. We use simulations to show that deploying BFWCs in routers of an approximate minimum vertex cover may be effective at detecting worms in the early stages of propagation.

In §2 we present a brief overview of related work. In §3 we present the principles of Bloom filters, which are a fundamental part of the BFWC studied in §4. In §5 we discuss the deployment of BFWCs in routers of a vertex cover. In §6 we address some important implementation issues of the system. In §7 we evaluate the performance of the BFWC under real traffic. In §8 we report on the performance of the BFWC in a simulated network. We close with concluding remarks in §9. In the Appendix we discuss further the selection of the threshold parameter used in this application.

## 2 Related Work

In this section we present a partial summary of strategies against worm propagation, including research on intrusion

<sup>‡</sup>Version: 17th June 2004.

\*School of Computer Science, Carleton University, Canada.

<sup>†</sup>Research & Innovation Centre – Security Group, Alcatel Canada Inc.

<sup>‡</sup>Version: 17th June 2004. Author addresses: {mvargas, paulv}@scs.carleton.ca, jean-marc.robert@alcatel.com.

detection systems (IDS) based on the principle that worms (as well as other kinds of denial-of-service attacks) perform *many* similar actions within a *short* period of time, which our approach also relies on.

The CERIAS Group [9] proposes a network anomaly-based IDS based on a common propagation technique observed in worms, namely probing many random IP-addresses in a “short” period of time. It is assumed that legitimate traffic does not exhibit such probing behaviour. The idea is to have a sensor constantly monitoring the outbound traffic of a host (or a set of hosts). This sensor triggers an alert if the number of packets of a particular  $[src\_addr, dst\_port]$  pair exceeds a predefined threshold  $t$ . The counter for each  $[src\_addr, dst\_port]$  pair is set to zero every  $Q$  units of time. The authors indicate that the implementation of this approach is in progress; details are not provided.

Williamson [42] pursues an approach based on the same principle. He presents a network anomaly-based intrusion detection and response system against viruses based on the belief that a virus would open many *new connections*, i.e., connections that have not been seen recently. This behaviour is also characteristic of worms. The idea is to maintain a list of  $n$  (different) destination IP addresses – called the working set – which represents the  $n$  most recently open connections. Any connection to an IP address not found in the working set is considered a new connection. Before a connection is processed the system checks for newness; if the connection is not new, it is processed normally. Otherwise, the connection is put into a delay queue, whose elements are popped off in FIFO order every  $Q$  units of time. An implementation of this approach is presented in [40].

Following the same principle, Toth and Kruegel [39] propose a network anomaly-based IDS against worms, involving a firewall on each host, a traffic monitor per local network, and a central analyzer. Each traffic monitor listens to all packets transmitted in that network and collects tuples of the form:  $[timestamp, src\_addr, src\_port, dst\_addr, dst\_port, z\_bytes\_of\_payload]$ . These tuples are stored temporarily in a connection-history table, which is renewed every given fixed period of time. The central analyzer periodically collects lists of tuples from the traffic monitors and infers attacks based on some measured values which include the number of times a tuple is repeated, the number of connection attempts to non-existing hosts, and the number of connection attempts to non-existing services. Upon detection of a worm, the central analyzer automatically broadcasts “appropriate changes” to the firewalls’ policies to all the hosts in the network. Given the number of issues involved in the whole process (e.g., overhead of reassembling packets, storage capacity required to keep lists of tuples, time to retain the lists, etc.), the impact on network per-

formance is not clear (no performance measurements or implementation is reported).

Based on the same principle, Chen and Heidemann [10] present a worm propagation detection system based on packet matching. An enhanced router maintains two lists of counters where each counter is associated to one port; one list is for incoming connections and the other for outgoing connections. The system matches port counters in both lists and then monitors the number of packets to distinct destination addresses for each of these ports. The system relies on Williamson’s observation [42] that users (as opposed to worms propagating to random IP addresses) usually make connections to a small set of addresses. Every predetermined period of time enhanced routers check the number of connections to different addresses. Worm activity is inferred if this number is “significantly” larger than the long-term average. The system is tested with simulations using different router-level topologies. The authors test the system under different scenarios including variations on scanning strategies, network topologies, and parameters of the detection system; however, no experimentation using real traffic or real network topologies is reported. Unless routers are vulnerable to worm infections, matching incoming and outgoing port counters is not effective in routers which have no LANs attached.

Singh et al. [35] present the “EarlyBird System” for worm detection based on packet payload analysis. Their approach consists of detecting the most popular flows (flows are identified with packet payload hashes), counting the number of source and destination addresses for these flows, and counting the number of connection attempts of these flows to unused portions of IP addresses. The system was able to detect different malicious packets in a real network. The system requires a number of complicated tasks, which makes its applicability in high-bandwidth equipment questionable.

Park and Lee [34] show experimentally that a vertex cover of the Internet can be constructed with approximately 20% of the total number of nodes. Their results rely on the characteristics of the network underlying the Internet [15]. Park and Lee use randomly generated power-law networks to test the performance of a well-known approximation algorithm for finding a vertex cover. They deploy filtering mechanisms in the vertex cover to detect spoofed addresses. We use the vertex cover results of Park and Lee to deploy our monitoring mechanism for detecting repeated packets.

Related work on Bloom filters is discussed in §3.

### 3 Bloom Filters

In this section we review Bloom filters [5]. A Bloom filter is a hash based method for testing membership of a series of items in a large given set of items, with allowable errors.

Bloom filters have been used in several different contexts since they were introduced [14, 13, 20, 24, 8]. For example, Snoeren et al. [36] present a traceback system capable of tracing packets delivered by the network in the recent past using Bloom filters. They show that the system is effective (up to a certain number of false positives, inherent to Bloom filters), efficient (requiring storage proportional to 0.5% of the link capacity), and, perhaps most importantly, they show that it can be efficiently implemented in hardware. Using Bloom filters each router of the network keeps a record of the packets that it sends forward over a certain fixed time window. When a packet needs to be traced back, the system constructs an attack graph indicating the path(s) from which the packet was forwarded along the network. This graph is constructed by checking which routers forwarded this packet.

The idea of Bloom filters is to insert a given set of items into a bit-array, or *Bloom-table*, as follows. Each item is hashed by  $k$  independent hash functions  $h_i$ , for  $i = 1, \dots, k$ . The resulting value of each of these hash functions is interpreted as an index pointing to an entry of the Bloom-table. Then the corresponding entry of the Bloom-table is set to 1 (initially, each entry is set to 0).

To test membership of an item  $p$ , a similar procedure is followed:  $p$  is hashed and if  $h_i(p) = 1$ , for  $i = 1, \dots, k$ , then it is inferred that  $p$  is already in the Bloom-table. See Figure 1.




Figure 1: Illustration of a Bloom filter.

One way to assess the efficiency of a Bloom filter is by measuring the accuracy of its membership test. If a membership test is negative, we can be certain that it is correct, i.e., the item  $p$  has not been inserted into the Bloom-table. However, if the test is positive, there exists the possibility that the result is incorrect, i.e., arises due to several items  $p' \neq p$  collectively causing cells  $h_i(p)$  being set to 1. We define a false positive as follows.

**Definition 1 (false positive).** Insert  $n$  items  $p_1, \dots, p_n$  into a Bloom filter. If a membership test for an item  $p \neq p_i$  (for  $i = 1, \dots, n$ ) succeeds, then we call this event a *false positive*.

Note that the notion of falseness here has to do with the fact that the membership test succeeds as a result of a hashing collision rather than a repeated item.

The number of false positives depends on the number of hash functions  $k$ , the size of the Bloom-table  $m$ , and the total number of items to be inserted,  $n$ . It is desirable to design a Bloom filter which simultaneously minimizes the number of false positives, the size of the Bloom-table, and the number of hash functions. The number of false positives can be approximated as follows. When an item is inserted into the Bloom-table, the probability that any particular entry is set to 1 by one hash function is  $1/m$ , and the probability that any particular entry is unchanged is  $1 - 1/m$ . Hence, after inserting  $n$  random items using  $k$  independent hash functions, the probability that a particular entry is still 0 is  $(1 - \frac{1}{m})^{nk}$ . Thus, when we test for membership, the probability  $f$  of a false positive equals the probability that the  $k$  hash functions all point to entries with value 1, i.e.,  $f = (1 - (1 - \frac{1}{m})^{nk})^k$ . For  $m$  large,  $f$  can be approximated by  $f \approx (1 - e^{-\frac{nk}{m}})^k$  which is minimized for

$$k = (\ln 2) \frac{m}{n}. \quad (1)$$

For  $m$  large, we obtain a false positive probability of  $f \approx 0.6185 \frac{m}{n}$ . (See e.g. [29, 16].)

There are several studies regarding hardware design of Bloom filters including feasibility studies of incorporating them in network devices without considerable performance degradation (e.g., [13, 36]).

## 4 Bloom Filters with Counters

The problem of detecting the “fan” effect produced by worm propagation (see §1) consists of detecting duplicated packets leaving a network. In particular, we focus on malicious software which employs scanning methods. As reported previously [18], a common scanning strategy consists of sending out many packets with same payload and destination port, but different destination addresses within a short period of time.

Fan et al. [16] introduced BFWC applied to web cache techniques. We use BFWC in a router-based monitoring system. The idea is to enhance the routers of a network (or a subset of them — see §5) with BFWC to analyze the outbound traffic. For each outbound packet  $p$ , an *enhanced router* (i.e., a router that implements a BFWC) tracks how many times this packet has been forwarded in the recent past. If the packet has been forwarded more than  $t$  times for an appropriate *trigger threshold*  $t$  (precisely the behaviour that one would expect from a worm) an alarm is triggered.

To measure the efficiency of a BFWC we need to analyze the expected number of alarms. The value of the threshold  $t$  directly affects the expected number of alarms. To determine an appropriate value for  $t$  we need to take into account two different factors that may increase the counters

without the presence of malicious packets. First, we give some definitions.

**Definition 2 (random packet).** A random packet is a concatenation of a fixed number  $s$  of bits, where each bit may be 0 or 1 with equal probability (0.5).  $s$  is the size of the random packet.

**Definition 3 (Class I traffic).** Class I traffic is a stream of  $n$  random packets.

**Definition 4 (Class II traffic).** Class II traffic is a stream of  $n$  packets generated by a communication protocol in a network. This stream contains at least one repeated packet.

When we speak of Class I traffic, we assume that  $s$  is sufficiently large relative to  $n$  that the probability of any two packets being identical is negligible, or at least relatively small. In real networks, although most packet streams are not randomly generated, some may look like Class I traffic in the sense that no packet is repeated.

We need to consider the increments to the counters caused by “Class I traffic”, which increments the counters due to different packets resulting in the same hash value (i.e., statistical hash function collisions). Secondly, we need to consider the increments caused by “Class II traffic”, i.e. non-worm traffic that contains repeated packets due to “normal”<sup>1</sup> repeated use of particular protocols (e.g., peer-to-peer, Google visits, Yahoo connections). Figure 2 illustrates these two factors. The *safety gap* is the remaining number of increments before the threshold  $t$  is reached after considering the increments caused by “Class I” and “Class II traffic”. “Class I traffic” increments are studied in the Appendix. Increments caused by “Class II traffic” are inherent to the characteristics of the system where the BFWC is implemented and such increments are not a major focus of this paper. However, the results of testing the BFWC under real traffic (§7) make more clear how “Class II traffic” affects the counters and what needs to be done to deal with this type of traffic.




Figure 2: Types of non-worm traffic that increment a Bloom counter, and the resulting *safety gap*.

Figure 3 illustrates a Bloom filter with counters. Algorithm 1 describes the actions taken for each packet arriving

<sup>1</sup>“Normal” will vary depending on the nature of the network and its users.

at an enhanced router to be forwarded. Line 2 is a *threshold test* which triggers if a packet (worm or otherwise) is processed more than  $t$  times.




Figure 3: Bloom filter with counters (BFWC).

In the Appendix we study the probability that any  $k$  among the  $m$  counters have value greater than  $t$  each after  $n$  packets are processed.

Since memory capacity is limited, the entries of the Bloom-table are reset to 0 every fixed period of time, the *reset period*. The probability that a counter reaches its maximum value  $M$  solely based on “Class I traffic” is low because of the choice of parameters (e.g., size of Bloom-table, number of hash functions, expected number of packets). If a counter reaches its maximum value  $M$ , increments and wraps around to 0, important information is lost for the remaining time of the reset period. One way to address this is to use a design whereby the counters remain at the maximum value.

---

**Algorithm 1** BFWC( packet  $p$  )

---

1. for  $i = 1$  to  $k$   
 $counter[h_i(p)] = counter[h_i(p)] + 1$
  2. if  $counter[h_i(p)] > t$  for all  $i = 1, \dots, k$  then  
*trigger alarm*
  3. else process  $p$  as usual
- 

## 5 Enhanced Routers

Consider BFWCs (as described in §4) with “appropriate” values for  $k$  and  $m$  according to the expected traffic volume  $n$ , and a threshold  $t$  with a reset period  $r$ . Suppose that these BFWCs are deployed in each router of the network. Worms spreading at  $t$  or more worm packets per second from a host would be detected by the first router that forward these packets. In this way, enhancing all the routers of a network would facilitate detection. However,

costs and complexity of deployment might make it impossible or impractical to enhance every single router of a network. Therefore, we are interested in minimizing the number of enhanced routers while still being able to detect worms efficiently. In this section we describe how we can take advantage of Internet connectivity properties to find an appropriate small set of routers to enhance.

Consider a graph representing a network of autonomous systems where each node represents an autonomous system. Internet topology follows power-law relationships [15], including: *the number of nodes with  $d$  neighbours is proportional to  $d^{-c}$ , where  $c > 1$ .*

Several experimental studies have attempted to find an accurate value for  $c$  (e.g. [15, 7]). Since the power-law property induces hubs in the network (i.e., routers attached to “many” other routers – see [41]), we suggest that enhancing the routers of a small vertex cover including the hubs would be advantageous. A *vertex cover* of an undirected graph  $G = (V, E)$  is a set of vertices  $C \subseteq V$  such that for each link  $(u, v) \in E$ ,  $u$  or  $v$  or both are in  $C$ . Given  $G$  and a positive integer  $z \leq |V|$ , the vertex cover problem is to find a vertex cover  $C$  such that  $|C| \leq z$ , if one exists. This problem is NP-complete [19].

Park and Lee [34] show experimentally that a vertex cover of the Internet can be constructed with approximately 20% of the total number of nodes, and use random power-law networks to test the performance of a standard greedy algorithm – which we call *greedy VC* – for finding an approximated minimum vertex cover. This algorithm iteratively selects a node of highest degree and adds the node to the vertex cover, deleting the associated edges until all links are covered. We expect that enhancing the routers as selected by algorithm *greedy VC*, will be efficient in detecting the worm since the routers of the vertex cover tend to be the ones with most neighbours, and therefore the ones that will likely forward most of the traffic in the network.

## 6 Implementation Issues

In this section we discuss a few important implementation issues regarding BFWCs, four in particular: the packet-subset that is input to the hash functions, the nature of the hash functions themselves, the amount of memory required to store the Bloom-table, and the length of the reset period. While additional issues arise in an actual implementation, these are the most fundamental ones.

To be able to detect the “fan” produced by worm propagation we must look only at those portions of the packet that do not change. Our selection criteria for extracting a *packet-subset* is to retain those portions of the packet which we expect to be identical in a propagating worm (see §9 for comments regarding polymorphic worms). For ex-

ample the destination IP address would vary, therefore this field should be excluded. For most worms, the destination *port* however, is expected to be constant. At the same time we want the packet-subset to contain enough information to avoid having non-identical packets which have the same packet-subset. Thus, we propose assembling packet-subsets of the form:  $[dst\_port, payload]$ . Taking packet-subsets of this form prevents worms from evading detection by simply spoofing (varying) source IP addresses which would result in distinct hashes; i.e., worm packets not being counted by the same  $k$  counters of the Bloom-table. In addition, this packet-subset allows the detection of worm packets coming from different IP sources. However, by omitting the source IP address in the packet-subset, *flashcrowds* (i.e., sudden transmission of a significant number of packets from different IP addresses) of legal packets yielding equal packet-subsets may trigger false alarms; e.g. the same http request made at the same time by a considerable number of users in a corporation network, or any broadcasting application as on-line gaming. The destination port helps us to identify the vulnerable application.

It is important to use appropriate hash functions according to the requirements on performance and security. Several papers have discussed performance measures of hash functions implemented in both hardware and software. For example Grembowski et al. [21] show experimentally that SHA-512 implemented in hardware performs surprisingly well (i.e., with not much additional performance degradation) in terms of speed with respect to other widely-known cryptographic hash functions. Nevelsteen and Preneel [31] compare several MAC algorithms against universal hash functions implemented in software and find that universal hash functions perform better than MAC algorithms in general. However, for our present purposes, we expect that even simple linear hash functions may suffice and provide performance advantages (e.g. see [22, page 151]). For the purposes of our present research also, we are less concerned about advanced attacks such as those presented by Crosby and Wallach [11]; we would expect a worm writer would achieve far greater success by spending additional time crafting better worms than by trying to exploit details of hash functions employed in a detection mechanism deployed in some subset of a large network.

As discussed in §3, the efficiency of a BFWC depends on the total number of packets  $n$  to be processed over a fixed reset period  $r$ , the number of hash functions  $k$ , the number of counters  $m$  in the Bloom-table, and the bit size of each counter. The objective is to achieve an acceptable level of false alarms with reasonable memory size.

Several factors need to be considered to design realistic and efficient Bloom filters. Suppose we have a relatively low bandwidth OC-3 router interface<sup>2</sup> and 520 KBytes of mem-

<sup>2</sup>OC stands for optical carrier. One OC-3 link has a capacity of

ory available for the Bloom-table<sup>3</sup>. First, if the threshold  $t$  can fit within 4 bits (see Appendix) then we can allocate two counters per byte, giving a Bloom-table with  $m = 1\,040\,000$  counters. Secondly, considering an average packet size of 1000 bits, this router can forward at most about 150 000 packets per second. Now, suppose we want to detect worm propagation within one-second windows, then a reset period of  $r = 1$  second yields  $n = 150\,000$  packets. Thirdly, the number of hash functions  $k$  can be obtained with (1):  $k = \ln(2) \cdot 1\,040\,000/150\,000 = 4.8$  (to reduce overhead we can choose  $k = 4$ ). Finally, an appropriate trigger threshold  $t$  under “Class I traffic” can be selected as explained in the Appendix. With these parameters and  $t = 15$  the probability of a false alarm after one second of “Class I traffic” is approximately  $1.52 \times 10^{-64}$  (see Appendix), or  $4.79 \times 10^{-57}$  in one year ( $3.15 \times 10^7$  seconds), leaving ample room to accommodate links of sufficiently higher bandwidth.

## 7 Performance of BFWCs under Real Traffic

In our evaluation we use packet-subsets of the form  $[dst\_port, payload]$ , as discussed in §6. It is clear that the BFWC would detect packets with identical packet-subsets forwarded at more than  $t$  packets per *reset period*. However, we expect that the number of non-worm packets with identical packet-subsets is not sufficiently large to trigger false alarms.

In our preliminary analysis, we use three different data sets to test a BFWC with the parameters described in Table 7: one hour of incoming traffic to our internal lab (*outside*), one day of traffic generated within our internal lab (*inside*), and one day of traffic of the well known DARPA data set [28] (the last two are known to contain no worm packets). The data sets consist of 113 463, 37 379, and 1 476 391 TCP and UDP packets, with average packet size of 10464, 7769 and 3011 bits per packet (respectively). The destination port frequency of the data sets are presented in Tables 1, 2 and 3. While our data sets fall well short of the amount of traffic in core routers, they suffice to validate the overall design principles of our approach. Similarly, despite known flaws in the DARPA data set [27], it suffices for our preliminary analysis.

To avoid false alarms due to repeated non-worm packets we need to adjust the BFWC according to the predominant (non-worm) repeated traffic in the data sets. For our lab data set, we inserted all the outside 113 463 and the inside 37 379 packets into the BFWC. For the outside traffic;

155.52 Mbps.

<sup>3</sup>An OC-3 router interface is not realistic compared to core routers which may have several OC-192 interfaces, however one OC-3 router interface can be well used as an illustrative example.

Port	Freq.	Port	Freq.	Port	Freq.
80	43 582	137	382	513	40
32777	38 815	138	346	1715	27
1985	7 771	21	291	1667	26
16080	6 455	34852	124	1699	26
32771	3 111	32781	113	1674	25
32787	3 060	53	108	1683	25
32788	2 713	22	106	1706	25
520	2 429	427	71	1689	24
631	2 324	3127	47	32772	22
135	770	445	47	1096	21

Table 1: Most frequent destination ports of the outside traffic.

Port	Freq.	Port	Freq.	Port	Freq.
53	7 830	34295	552	33298	270
631	6 118	32886	420	3932	263
9100	5 489	33175	414	68	192
32795	2 575	33788	414	67	188
138	2 088	1028	319	3179	177
123	1 544	33300	315	40216	156
38729	1 455	38833	300	3213	126
137	1 198	39135	287	1298	126
32775	615	32770	285	3159	114
32782	600	50566	272	34296	112

Table 2: Most frequent destination ports of the inside traffic.

Port	Freq.	Port	Freq.	Port	Freq.
23	398 734	123	10 511	15901	6 395
80	202 921	21262	8 983	16510	6 143
22	85 028	24638	8 580	18486	6 087
25	40 639	520	8 124	30865	5 563
53	28 354	17258	7 963	20551	5 334
161	19 992	23308	7 536	20	5 156
32770	19 413	16933	6 856	14942	5 024
15574	14 177	16586	6 757	15037	4 798
12862	10 891	29810	6 626	4702	4 356
1132	10 884	21	6 554	17782	4 349

Table 3: Most frequent destination ports of DARPA data set.

we found that packets with the following destination ports triggered false alarms: 1985 (Hot Standby Router Protocol), 520 (local routing process), and 631 (Internet Printing Protocol). We filtered out packets with these destination ports, such that they are not processed by the BFWC. After this filtering, the BFWC triggered no false alarms and moreover was able to detect Slammer packets [30] trying to infect new targets; we were not expecting these latter packets in our network, and thus were pleasantly surprised (by their detection, rather than their presence). The detection of incoming Slammer packets confirms the effectiveness of BFWCs since incoming worm packets are more difficult to detect, as they may arrive at slower rates than if they were generated from within our lab.

For the inside traffic, no packets had to be filtered out. The BFWC did not trigger false alarms under this traffic; this was as expected, in traffic with no worms.

As for the DARPA data set, we filtered out several destination ports: 520, 161 (Simple Network Management Protocol), 6667 (Internet Relay Chat), and 1270 (OPSMAN). We inserted streams of 155 550 packets each (the maximum number of packets per second that can be forwarded by a router with 3 OC-3 interfaces, considering an average packet length of 3000 bits per packet). After filtering packets with these destination ports, the BFWC triggered no false alarms, which is what we expected since the data set is known to have no malicious packets.

While obviously additional testing is required on larger data sets representing true Internet traffic, our preliminary results show the effectiveness of our system and suggest that the approach may indeed be effective in core routers.

## 8 Simulations

BFWCs have to be tuned up to take into account environment variables and characteristics inherent to the network, such as type of traffic (e.g., predominantly peer-to-peer), network topology fluctuation (e.g., constant failures or changes will trigger routing control packets), network speed and volume of traffic, available memory in routers, etc. The parameters of the BFWC used in the simulations are determined as discussed in §6.

The goals of the simulations are (1) to determine an approximation of how efficient the BFWC is in detecting a worm when the BFWC is installed in the routers of an approximate minimum vertex cover and when all the routers of the network are enhanced, and (2) to determine an approximation of how many infections of a particular worm have been perpetrated by the time the worm is first detected by one BFWC.

Several network and worm simulators were considered [33, 38, 25, 32]. In our simulations we use NS-

2 [32]. NS-2 provides flexibility on the topology of the network, making it possible to define “large” networks with a reasonable level of detail. The simulations were performed over the backbone topologies of three real ISPs. These topologies were made available by the Rocketfuel project [37, 26, 1], an Internet measurement tool. The simulation also uses real link delays, also made available by the Rocketfuel project. The characteristics of the ISP topologies used in our simulations are summarized in Table 4.

	ISP 1	ISP 2	ISP 3
minimum link delay	1ms	1	1
maximum link delay	17ms	29	44
number of routers	108	87	79
number of links	153	161	147
routers in vertex cover	33	46	37

Table 4: Characteristics of networks used in simulation.

In our simulations, each node represents a router connected to a local area network (LAN). Internally, the simulator treats each router as a collection of hosts (e.g. clients, servers) on a LAN, as illustrated in Figure 4. To simplify the simulations, we assume that each router has a LAN associated with it.<sup>4</sup> Each router has a disjoint set of IP addresses associated, which represent the hosts of the LAN accessible through this router. To simulate IP addresses that do not reside within this ISP (i.e. foreign hosts), we assume that these addresses are reachable through arbitrary routers (assigned at the beginning of the simulation) of this ISP (these foreign hosts can also be thought of as non-vulnerable hosts residing within this ISP).




Figure 4: Example of network topology used in simulations.

We simulated a UDP worm that propagates in a single 400-byte UDP packet (not including the required headers) to a non-predetermined number of random IP addresses. The operation of this worm is described in Algorithm 2.

The propagation of a worm can be modelled with the “generic epidemic model” [12]. This model divides the hosts into two disjoint sets: *susceptible* and *infectious*. Infected hosts remain infected for an undetermined period of time (i.e., hosts do not transition from infectious to suscepti-

<sup>4</sup>For simulation purposes, routers themselves are not considered to be hosts.

---

**Algorithm 2** Simulated\_UDP\_Worm()

---

Upon reception of a worm packet, repeat forever:

1. Send copy of the worm packet to a random IP address (i.e. host) using a pseudo-random number generator.
  2. Wait for  $D$  seconds (the delay  $D$  determines the stealthiness of the worm; a larger delay makes detection more difficult).
- 

ble). Assuming that the network topology does not change (i.e., there are no new hosts or links, nor failures), and that there are no delays in the transmission lines, the number of infectious hosts at time  $T$ ,  $i(T)$  can be modelled with the following equation:  $\frac{di(T)}{dT} = \beta \cdot i(T) \cdot s(T)$ , where  $s(T)$  is the number of susceptible hosts at time  $T$  and  $\beta$  is the propagation rate of the worm (i.e., the number of victims a single infectious host attempts to infect per unit of time).

Initially, in our simulations, all the hosts of the network are susceptible, having a process listening at port 0, which is assumed to be the service being exploited by the worm. The propagation rate,  $\beta$ , depends primarily on processor power and bandwidth available. We assume that hosts from other ISPs do not try to infect hosts of the ISP in question; in other words, we assume other ISPs are effectively preventing the spread of the infection.

To simulate multiple infected hosts within a router, whenever a host becomes infected, the propagation rate of worm packets leaving this router is increased accordingly to reflect the number of infected hosts within the LAN attached to this router. If a host attempts to infect another host residing within the same LAN, the attached router will not be able to see the worm packet, and consequently, even if this router was enhanced, the BFWC would not be able to count the worm packet. Thus, we do not detect worm packets propagating within a LAN.

Constant-rate UDP packets representing non-worm traffic are transmitted between random pairs of hosts. The rate of UDP packets transmitted by a host is fixed to 200 packets per second. The payload of each packet is a string of 500 random bytes.

The traffic generated by the underlying routing mechanism should also be considered as part of non-worm traffic. The NS-2 simulator comes with several built-in routing algorithms, including distance vector algorithms [3]. Our simulation was configured to use the built-in implementation of the Distributed Bellman-Ford routing algorithm [3]. Route information packets are thus inherently added to the set of non-worm packets.

The simulation let us easily adjust the parameters of the network, the BFWC and the simulator itself. The configuration of the three ISP networks used in the simulations is shown in Table 5. The parameters of the worm and non-worm network traffic are listed in Table 6. The non-worm

traffic can be classified as “Class I traffic” (see §4). The parameters of the BFWC were set up as discussed in §6, assuming that the routers have about 520 KBytes of memory available for the Bloom-table and that their maximum throughput is around 150 000 packets per second; a network operator should be able to adjust the BFWC parameters according to the characteristics of the underlying network and its predominant traffic (see §6). The parameters of the BFWC appear in Table 7.

Parameter	Value or Configuration
address space	$\sim 3 \times 2^{16}$ (approx. 3 class B networks)
number of hosts associated with each router (see Figure 4)	100 (empirically set; Rocketfuel does not provide a value for this parameter)
which routers to “enhance”	routers as selected by an approximate minimum vertex cover

Table 5: Network parameters used in the simulations.

Parameter	Value
worm payload	fixed stream of 400 bytes within one UDP packet
infection attempts per infected host	1 per worm delay period, $D$
worm delay, $D$	0.5 seconds
non-worm payload	random stream of 500 bytes within a single UDP packet
rate of non-worm packets	200 packets/second, each 500 bytes

Table 6: Worm and non-worm traffic used in the simulations.

Parameter (see Figure 3)	Value
size of hash values, $b$	20 bits ( $b = \log_2 m$ )
entries in Bloom-table, $m$	1 048 576 ( $m = 2^b$ )
bits per counter, $c$	4
hash functions, $k$	4
threshold $t$	15
reset period $r$	1 second

Table 7: BFWC parameters used in the simulations.

The simulations were performed over three real ISP topologies with real link delays (see Table 4). The infection is started by a random host within the network. Figure 5 shows how the BFWC is able to detect the worm at its early stage of propagation. The curves show the number of infected hosts (cf. Figure 4) over time. The arrows show the first detection of the worm by an enhanced router.

The simulations results suggest that enhancing the routers of an approximate minimum vertex cover with BFWCs may be effective in detecting a worm within a rea-




Figure 5: Detection of the worm enhancing the routers of an approximate minimum vertex cover.

sonably short period of time, e.g. before the worm spreads to 1% (e.g., per Figure 5, 20 infections out of  $79 \times 100$  hosts for ISP 3) of the (vulnerable) hosts in the entire network.

Also with additional experimentation we found that, with the parameters of our particular simulations, enhancing all the routers of the network is essentially no more effective than enhancing only those in an approximate minimum vertex cover. This fact becomes clear when we analyze the scenario. For example, Figure 6 depicts a portion of the (non-geographic) topology of ISP 1. According to algorithm *greedyVC*, the routers with most neighbours are most likely to be enhanced, and these routers are the ones forwarding most traffic in the network. For instance, the top right enhanced router of the figure would not be able to detect a worm propagating at less than  $t = 15$  packets per second from within its associated LAN because its Bloomtable is reset every second (reset period  $r = 1$  second). The alarms are triggered due to high volumes of worm packets collectively sent by a significant number of hosts. Therefore a worm is most likely first detected by one of the routers with most neighbours, which most probably belongs to the vertex cover.

## 9 Concluding Remarks

We have presented a router-based monitoring system using Bloom filters with counters (BFWC) for detecting worms in the propagation phase. The monitoring system is based on stateful analysis of network traffic in the routers of a network of LANs, and the principle that worms send out many similar packets within a short period of time. The approach uses a variation of Bloom filters to count the number of times that a packet is repeated. The results of our

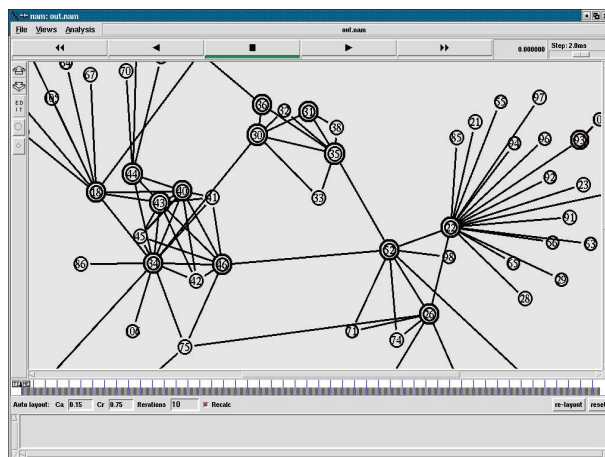


Figure 6: Topology of ISP 1. Each node represents a router. Enhanced routers are illustrated with double circle.

simulations suggest that this approach may be effective in detecting a worm at the early stages of propagation. Further experimentation or prototyping under real networks will reveal important characteristics about the efficiency of the approach.

BFWCs are not effective against polymorphic worms (i.e., worms that mutate as they are transmitted from one system to another). In this case, manual worm-specific re-configuration of the BFWC may be necessary. However, our approach in its present form is of use as a first step, and indeed is capable of detecting the “popular” recent worms, almost all of which are non-polymorphic. We believe that solving this problem is important, while acknowledging that further work is required for polymorphic worms.

Our system cannot distinguish between repeated worm packets and repeated non-worm packets. Though this situation must be appropriately addressed (see next paragraph), this “feature” allows our approach to have applications beyond worm detection such as detection of dangerous traffic fluctuations, abusive use of certain services, and distributed denial-of-service attacks. This characteristic of our system does not prevent it from being useful for worm detection since we expect the frequency of repeated worm packets to dramatically exceed that of repeated non-worm packets (see safety gap in Figure 2), except for “stealth” worms.

One solution to preventing repeated non-worm packets from triggering false alarms is to pre-identify such packets and decrement the corresponding counters of the Bloomtable every fixed period of time. We view this as part of the necessary network-specific “tuning”.

It would be interesting to test the BFWC under data sets with significant numbers of packets of particular protocols (e.g. peer-to-peer) which we expect to be susceptible to false alarms, and find out the “tuning” required under such traffic. It would also be interesting to include more fields in the packet-subset and see how this affects the rate of unde-

tected worms. It would also be of interest to run the simulations under different worm propagation rates, and with even fewer enhanced routers than those of a vertex cover, or enhance only routers with particular characteristics (e.g. capacity, potential of false alarms).

## Acknowledgements

We would like to thank Javier Govea and Tomás Pospíchal for their valuable help on the simulations and the Poisson approximation to the distribution of counters' values of the Bloom-table (respectively). The first author acknowledges the support of Alcatel Canada and NCIT. The third author is Canada Research Chair in Network and Software Security and acknowledges the support of an NSERC Discovery Grant, the Canada Research Chairs Program, and Alcatel Canada.

## References

- [1] T. Anderson, R. Mahajan, N. Spring, and D. Wetherall. Rocketfuel: An ISP topology mapping engine, 2003. <http://www.cs.washington.edu/research/networking/rocketfuel/> [Accessed: August 2, 2003].
- [2] A.D. Barbour, L. Holst, and S. Janson. *Poisson Approximation*. Oxford University Press, New York, USA, 1992.
- [3] D. Bertsekas and R. Gallager. *Data Networks*. Prentice Hall, Englewood Cliffs, USA, 1992.
- [4] M. Bishop. *Computer Security: Art and Science*. Addison-Wesley, Boston, USA, 2003.
- [5] B.H. Bloom. Space/time trade-offs in hash coding with allowable errors. *Communications of the ACM*, 13(17):422–426, July 1970.
- [6] M.V. Boutsikas and M.V. Koutras. On the number of overflow urns and excess balls in an allocation model with limited urn capacity. *Statistical Planning and Inference*, 104:259–286, 2002.
- [7] A. Broder, R. Kumar, F. Maghoul, P. Raghavan, S. Rajagopalan, R. Stata, A. Tomkins, and J. Wiener. Graph structure in the Web. *Computer Networks*, 33(1–6):309–320, June 2000.
- [8] A. Broder and M. Mitzenmacher. Network applications of Bloom filters: A survey, 2003. URL: <http://www.eecs.harvard.edu/~michaelm/ListByYear.html> [Accessed: May 11, 2004].
- [9] The CERIAS Intrusion Detection Research Group. Digging for worms, fishing for answers. In *Proceedings of the Annual Computer Security Application Conference (ACSAC'02)*, Las Vegas, USA, December 9–13 2002.
- [10] X. Chen and J. Heidemann. Detecting early worm propagation through packet matching. Technical Report ISI-TR-2004-585, University of Southern California, February 2004.
- [11] S.A. Crosby and D.S. Wallach. Denial of service via algorithmic complexity attacks. In *Proceedings of the 12th USENIX Security Symposium*, Washington, USA, August 4–8 2003.
- [12] D.J. Daley and J. Gani. *Epidemic Modelling: An Introduction*. Cambridge University Press, Cambridge, UK, 1999.
- [13] S. Dharmapurikar, P. Krishnamurthy, T. Sproull, and J. Lockwood. Deep packet inspection using parallel Bloom filters. In *Symposium on High Performance Interconnects (HotI)*, pages 44–51, Stanford, USA, August 2003.
- [14] S. Dharmapurikar, P. Krishnamurthy, and D. Taylor. Longest prefix matching using Bloom filters. In *Proceedings of the Special Interest Group on Data Communication (SIGCOMM'03)*, pages 201–212, Karlsruhe, Germany, August 25–29 2003.
- [15] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On power-law relationships of the Internet topology. In *Proceedings of the Special Interest Group on Data Communication (SIGCOMM'99)*, pages 251–262, Boston/Cambridge, USA, August 31 – September 1 1999.
- [16] L. Fan, P. Cao, J. Almeida, and A.Z. Broder. Summary cache: A scalable wide-area Web cache sharing protocol. *IEEE/ACM Transactions on Networking*, 8(3):281–293, June 2000.
- [17] W. Feller. *An Introduction to Probability Theory and its Applications*, volume 1. John Wiley and Sons, New York, USA, 3rd edition, 1968.
- [18] Fyodor. The art of port scanning. *Phrack Magazine*, 7(51), 1997. URL: <http://www.phrack.org> [Accessed: March 6, 2003].
- [19] M.R. Garey and D.S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman and Company, New York, USA, 1979.
- [20] E.-J. Goh. Secure indexes. Cryptology ePrint Archive, Report 2003/216, 2003. URL: <http://eprint.iacr.org/2003/216/> [Accessed: January 7, 2004].
- [21] T. Grembowski, R. Lien, K. Gaj, N. Nguyen, P. Bellows, J. Flidr, T. Lehman, and B. Schott. Comparative analysis of the hardware implementations of hash functions SHA-1 and SHA-512. In *Proceedings of Information Security Conference*, Sao Paulo, Brazil, September 30 – October 2 2002.
- [22] B. Horne, L. Matheson, C. Sheehan, and R.E. Tarjan. Dynamic self-checking techniques for improved tamper resistance. In *Proceedings of the First ACM Workshop on Digital Rights Management (DRM 2001)*, volume 2320 of *Lecture Notes in Computer Science*, pages 141–159. Springer, 2002.

- [23] K. Joag-Dev and F. Proschan. Negative association of random variables, with applications. *The Annals of Statistics*, 11(1):286–295, 1983.
- [24] A. Kumar, J. Xu, L. Li, and J. Wang. Space-code Bloom filter for efficient traffic flow measurement. In *Proceedings of IMC*, Miami Beach, USA, October 27–29 2003.
- [25] M. Liljenstam. Modeling of security and systems. A network worm modeling package for SSFNet, 2003. URL: <http://www.crhc.uiuc.edu/~mili/research/ssf/worm/> [Accessed: September 10, 2004].
- [26] R. Mahajan, N. Spring, D. Wetherall, and T. Anderson. Inferring link weight using end-to-end measurements. In *Proceedings of the Internet Measurement Workshop 2002 (IMW'02)*, Marseille, France, 2002.
- [27] J. McHugh. Testing intrusion detection systems: A critique of the 1998 and 1999 DARPA intrusion detection system evaluations as performed by Lincoln Laboratory. *ACM Transactions on Information System Security (TISSEC)*, 3(4):262–294, 2000.
- [28] MIT Lincoln Laboratory. DARPA intrusion detection evaluation: Data sets, 1999. URL: [http://www.ll.mit.edu/IST/ideval/data/data\\_index.html](http://www.ll.mit.edu/IST/ideval/data/data_index.html) [Accessed: April 1, 2004].
- [29] M. Mitzenmacher. Compressed Bloom filters. In *Proceedings of the 20th Annual ACM Symposium on Principles of Distributed Computing (PODC 2001)*, pages 144–150, Newport, USA, August 26–29 2001.
- [30] D. Moore, V. Paxson, S. Savage, C. Shannon, S. Staniford, and N. Weaver. Inside the Slammer worm. *IEEE Security & Privacy*, 1(4):33–39, 2003.
- [31] W. Nevelsteen and B. Preneel. Software performance of universal hash functions. In *Proceedings of EUROCRYPT'99*, pages 24–41, Prague, Czech Republic, May 2–6 1999.
- [32] NS-2. The network simulator - NS-2, 2003. URL: <http://www.isi.edu/nsnam/ns/> [Accessed: September 10, 2003].
- [33] OPNET Technologies Inc. Opnet modeler, 2003. URL: <http://www.opnet.com> [Accessed: September 10, 2003].
- [34] K. Park and H. Lee. On the effectiveness of route-based packet filtering for distributed DoS attack prevention in power-law internets. In *Proceedings of the Special Interest Group on Data Communication (SIGCOMM'01)*, San Diego, USA, August 27–31 2001.
- [35] S. Singh, C. Estan, G. Varghese, and S. Savage. The EarlyBird system for real-time detection of unknown worms. Technical Report CS2003-0761, University of California, San Diego, August 4 2003.
- [36] A.C. Snoeren, C. Partridge, L.A. Sanchez, C.E. Jones, F. Tchakountio, S.T. Kent, and W.T. Strayer. Hash-based IP traceback. In *Proceedings of the Special Interest Group on Data Communication (SIGCOMM'01)*, San Diego, USA, August 27–31 2001.
- [37] N. Spring, R. Mahajan, and D. Wetherall. Measuring ISP topologies with Rocketfuel. In *Proceedings of the Special Interest Group on Data Communication (SIGCOMM'02)*, Pittsburgh, USA, August 19–23 2002.
- [38] SSFNet. Scalable simulation framework network models, 2003. URL: <http://www.ssfnet.org/homePage.html> [Accessed: September 10, 2003].
- [39] T. Toth and C. Kruegel. Connection-history based anomaly detection. In *Proceedings of the 2002 IEEE Workshop on Information Assurance and Security*, New York, USA, June 17–19 2002.
- [40] J. Twycross and M.M. Williamson. Implementing and testing a virus throttle. In *Proceedings of the 12th USENIX Security Symposium*, Washington, USA, August 4–8 2003.
- [41] D.J. Watts. *Small Worlds*. Princeton Studies in Complexity, New Jersey, USA, 1999.
- [42] M.M. Williamson. Throttling viruses: Restricting propagation to defeat malicious mobile code. In *Proceedings of the Annual Computer Security Application Conference (ACSAC'02)*, Las Vegas, USA, December 9–13 2002.

## Appendix: Determining the Threshold of Repeated Packets in Bloom Filters with Counters

In §4 we describe Bloom filters with counters (BFWC). Here, we describe a method for generating a table to allow the appropriate selection of a threshold  $t$  of repeated packets – at which an alarm is triggered – under “Class I traffic”. We first review some basic statistical concepts.

**Definition 5 (Bernoulli trials).** *Bernoulli trials are repeated independent trials with only two possible outcomes for each trial and their probabilities remain the same over time: outcome success has probability  $p$ , and outcome failure has probability  $q = 1 - p$ .*

Often one is interested in the probability of obtaining  $s$  successes out of  $n$  Bernoulli trials regardless of the order of occurrence. This probability has a *Binomial distribution*.

**Theorem 1.** [17] *Let  $b(s; n, p)$  be the probability that  $n$  Bernoulli trials with probabilities  $p$  for success and  $q = 1 - p$  for failure result in  $s$  successes and  $n - s$  failures ( $0 \leq s \leq n$ ). Then*

$$b(s; n, p) = \binom{n}{s} p^s q^{n-s}.$$

For large  $n$  and small  $p$ , a direct evaluation of  $b(s; n, p)$  may be impractical. In this case, we can use the Poisson approximation to  $b(s; n, p)$ :

$$b(s; n, p) \approx p(s; \lambda) = e^{-\lambda} \left( \frac{\lambda^s}{s!} \right) \quad (2)$$

where  $\lambda = np$ .

A generalization of the Binomial distribution for  $m$  possible outcomes for each trial is the *multinomial distribution* [17]. The possible outcomes for each trial are denoted  $E_i$  (for  $i = 1, \dots, m$ ) and the probability of each outcome is denoted  $p_i$  (for  $i = 1, \dots, m$ ,  $\sum_{i=1}^m p_i = 1$ , and  $p_i \geq 0$ ). In a multinomial distribution, the probability that in  $n$  trials  $E_1$  occurs  $s_1$  times,  $E_2$  occurs  $s_2$  times, etc., is:

$$\frac{n!}{s_1! s_2! \dots s_m!} p_1^{s_1} p_2^{s_2} \dots p_m^{s_m} \quad (3)$$

where  $s_i$  are arbitrary non-negative integers and  $s_1 + s_2 + \dots + s_m = n$ .

Observe that the number of successes for each possible outcome  $E_i$  after  $n$  trials can be approximated using (2).<sup>5</sup> Now we return to BFWCs by stating the following.

**Observation 1.** *Inserting  $N$  packets of “Class I traffic” into a BFWC with  $m$  counters and  $k$  independent hash functions that distribute the packets equiprobably into the  $m$  counters,<sup>6</sup> is equivalent to equiprobably throwing  $n = Nk$  packets into  $m$  bins.*

Below we use the terms counter and bin interchangeably. Consider the following experiment.

**Experiment 1.** *Throw equiprobably  $n$  random packets into  $m$  bins.*

We are interested in finding a threshold  $t$  such that, after performing Experiment 1, the probability that  $k$  random bins have more than  $t$  packets is “sufficiently small” to be accepted according to our context (i.e. a tolerable number of false alarms under “Class I traffic”). Applying (3) may be impractical since  $m$  and  $n$  may be large. Instead we use (2) to approximate the number of packets in each bin.

As an illustrative example, we use the parameters from earlier in the paper to generate a table allowing appropriate selection of a value for  $t$ , using the Poisson approximation given by (2).

**Example 1.** *Consider a software implementation of a BFWC with  $m = 1\,048\,576$  counters and  $k = 4$  hash functions. After inserting  $N = 150\,000$  random packets into the*

*BFWC, the mean number of packets per bin is 0.572205. We want to set the value of  $t$  such that the probability that  $k = 4$  counters taken at random have more than  $t$  packets is “sufficiently small”. The probability that  $k$  random bins all have counters exceeding  $t$ , denoted  $P[k, t, m, N]$ , is approximated using (2) as:*

$$P[k, t, m, N] \approx \left( 1 - \sum_{i=0}^t e^{-n/m} \frac{(n/m)^i}{i!} \right)^k. \quad (4)$$

*Table 8 shows this approximation for several values of  $t$ . Based on the table, a network operator can choose an appropriate value for  $t$  according to the number of false alarms tolerable over some time period, under “Class I traffic”.*

$t$	$P[k, t, m, N]$	$t$	$P[k, t, m, N]$
2	$1.75 \times 10^{-7}$	9	$1.45 \times 10^{-37}$
3	$6.50 \times 10^{-11}$	10	$1.04 \times 10^{-42}$
4	$1.03 \times 10^{-14}$	11	$5.31 \times 10^{-48}$
5	$8.02 \times 10^{-19}$	12	$1.98 \times 10^{-53}$
6	$3.43 \times 10^{-23}$	13	$5.93 \times 10^{-59}$
7	$8.69 \times 10^{-28}$	14	$2.43 \times 10^{-63}$
8	$1.38 \times 10^{-32}$	15	$1.52 \times 10^{-64}$

Table 8: Poisson approximation to  $P[k, t, m, N]$  after inserting  $N = 150\,000$  packets into a BFWC with  $k = 4$  hash functions and  $m = 1\,048\,576$  counters.

The expected number of false alarms while inserting  $N$  packets is  $\sum_{i=1}^N P[k, t, m, i]$ . For large values of  $i$ ,  $P[k, t, m, i]$  can be estimated using (4). For small values of  $i$ ,  $P[k, t, m, i]$  becomes negligibly small. Observe that  $P[k, t, m, i] > P[k, t, m, i']$  for  $i > i'$ . Therefore, a very coarse bound for the expected number of false alarms is  $\sum_{i=1}^N P[k, t, m, i] < N \cdot P[k, t, m, N]$ .

<sup>5</sup>The accuracy of the Poisson approximation to the multinomial distribution is well-studied (e.g. [23, 2, 6]). A Normal distribution may also be used as an approximation to  $b(s; n, p)$  for large  $n$  [17].

<sup>6</sup>i.e., the probability of throwing a packet into counter  $i$  is  $1/m$ , for all  $1 \leq i \leq m$ .