# On Scalable Attack Detection in the Network

Ramana Rao Kompella, Sumeet Singh, George Varghese
University of California, San Diego,
9500 Gilman Drive,
La Jolla, CA 92093
{ramana,susingh,varghese}@cs.ucsd.edu

## ABSTRACT

Current intrusion detection and prevention systems seek to detect a wide class of network intrusions (e.g., DoS attacks, worms, port scans) at network vantage points. Unfortunately, all the IDS systems we know of keep per-connection or per-flow state. Thus it is hardly surprising that IDS systems (other than signature detection mechanisms) have not scaled to multi-gigabit speeds. By contrast, note that both router lookups and fair queuing have scaled to high speeds using *aggregation* via prefix lookups or DiffServ. Thus in this paper, we initiate research into the question as to whether one can detect attacks without keeping per-flow state. We will show that such aggregation, while making fast implementations possible, immediately cause two problems. First, aggregation can cause *behavioral* aliasing where, for example, good behaviors can aggregate to look like bad behaviors. Second, aggregated schemes are susceptible to spoofing by which the intruder sends attacks that have appropriate aggregate behavior. We examine a wide variety of DoS attacks and show that several categories (bandwidth based, claim-and-hold, host scanning) can be scalably detected. By contrast, it appears that stealthy port-scanning cannot be scalably detected without keeping per-flow state.

**Categories and Subject Descriptors:** C.2.6 [**Computer Communication Networks**]: Internetworking – *Routers;*

**General Terms:** Algorithms, Design, Measurement, Performance, Security

**Keywords:** Security, Scalability, Denial of Service

## 1. INTRODUCTION

The earliest network security solutions attempted to secure Internet hosts using *anti-virus* software running at end-nodes, and *firewalls* installed at network vantage points (or, more recently, at hosts themselves). Unfortunately, end-node based approaches must be widely deployed within a network to protect against attacks. They also do very little

to mitigate against bandwidth attacks that may be blocked at the end-nodes but consume so much internal network bandwidth that the network is unusable. Similarly, most Distributed Denial of Service (DDoS) attacks and scans routinely penetrate firewalls using essential services such as HTTP or email.

For these reasons, a number of researchers and vendors have suggested perimeter defenses that sit at the entrance to networks or subnets. Besides firewalls, a traditional approach has been to do *intrusion detection* (and sometimes mitigation) at such points. Two classical approaches to intrusion detection have been *anomaly detection* and *signature detection.* Signature detection [30] is useful to detect an important class of attacks (e.g., known worms and viruses) but is not helpful in detecting other attacks (e.g., scans, DDoS attacks) which are not characterized by a signature within *a single packet*, but by unusual behavior across *a set of packets.* In this paper, we concentrate on detecting and mitigating such attacks (that can only be detected by behavior across a set of packets) at network vantage points.

While anomaly detection also targets such attacks [8], anomaly detection is often very general, and works by first automatically identifying a baseline for "normal" network behavior (using say wavelets [8] or change point detection [25]) and then flagging deviations from such behaviors as possible attacks. A difficulty with the most general approaches is the difficulty of establishing normal behaviors because most network traffic behaviors evolve in unpredictable ways.

**Network based behavioral approaches:** A simpler (but less general) approach to anomaly detection is to look for violations of specific behaviors, where a description of the bad behaviors has been *preprogrammed* into the detection device. In this paper, we will refer to such techniques as network based behavioral approaches. Such approaches have been widely deployed.

For example, in recent years a number of researchers and vendors have worked on the problem of detecting scans [44, 15], and similarly on detection of distributed denial-of-service [32, 4, 7, 45]. Detecting scans can be useful as a stepping stone for detecting various kinds of suspicious behavior including worms and attempts to check for backdoors.

Note that to be effective such a device must run at a network vantage point where it sees a lot of traffic. A minimal deployment would be at the edge of a subnet; a more useful deployment would be at the entrance to a network; a potential future deployment would be within ISPs. Notice that for some attacks such as Denial of Service (DoS) it is helpful to detect the attack as upstream in the attack path

as possible, to reduce the collateral damage caused to innocent sources that share the attack path until the detection device.[1]

Running a detection device at the edge of a network requires that the detection device operate at higher link speeds of 1 Gbps and higher. Such positions also expose the detection system to a larger number of flows making it more difficult. But most behavioral based approaches do maintain per-flow state. For example to detect port scans, Snort maintains a large vector per source to count all the ports and destinations each source talks to. Not only does this plugin take a large amount of space, but it also slows down the Snort code considerably when it is enabled. Bro [37] also maintains per-flow state in order to detect evasion and other attacks. Similar observations hold for many traditional approaches to detecting DDoS [45].

To meet the speed challenge, several vendors (e.g., NetScreen, Fortinet) and researchers are implementing detection in hardware. While signature detection is being done in hardware, it has proved difficult to speed up network based behavioral approaches.

*Scalable Attack Detection:* The main question we examine in this paper is whether it is possible to scale behavioral network detection to very high speeds. To begin to grapple with this question, it is instructive to consider other network functions that have been made to be scalable. Both in the case of IP lookups and network QoS, scalability has been achieved via *aggregation* to reduce the state used by the function so as to fit into high speed memory. For example, Internet lookups in routers use prefix aggregation to store around 150,000 prefixes for the entire Internet. Similarly, DiffServ uses class aggregation to avoid per-flow state in core routers.

Aggregation helps with forwarding performance for the following reasons. First, the number of connections/flows at network vantage points can easily scale into the millions, and this does not scale with the increases in the size of high speed memory. Second, the highest speed memories (on-chip and off-chip SRAM or cache) scale far more slowly than the number of flows.

Thus a central question we ask in this paper is: *can we use aggregation to reduce the state required for behavioral attack detection?* While several types of attacks (e.g., evasion, TCP hijacking) can be proved to be impossible to detect in scalable fashion[29], we will show in this paper that under some minimal assumptions, scan detection and DDoS detection (and perhaps several others) can indeed be detected scalably using aggregation.

The use of aggregation for scalable attack detection immediately creates two fundamental problems:

- *Behavioral Aliasing*: One form of behavioral aliasing occurs when a set of well behaved connections aggregate to look like bad behavior, creating a *false positive.* For example, when detecting a port scan, if we aggregate several sources into an aggregate, while each source may talk to only a few distinct destinations, the aggregate may look like it is talking to lots of destinations. A rigorous argument [29] can be based on this

intuition to show that attempts to scalably detect port scans using such a predicate do not work. A second form of behavioral aliasing occurs when the aggregate behavior of several badly behaved connections looks like good behavior – a *false negative.*

- *Spoofing:* Spoofing occurs when an intelligent attacker subverts the detection mechanisms by suitably spoofing the attack to appear benign. For example, the SYN-DOG approach to SYN-FLOOD detection[45] does a first level of aggregation by keeping state only on a per-destination basis (not sufficient, but a good start) for the difference between SYNs and FINs going to each destination. Such a scheme can always be spoofed by the attacker sending spurious FINs that do not serve to finish any active connection but only confuse the detection mechanism.

We believe that any scalable intrusion detection mechanism must deal with these two issues. Thus the contributions of this paper are as follows:

*1. Framework:* Our paper initiates the study of scalable attack detection schemes. We use behavioral aliasing and spoofing as a framework to analyze such techniques.

*2. Technique:* As a specific example, we focus on scalable DDoS and scan detection, and propose a specific new scalable technique called Partial Completion Filters (PCFs) and analyze behavioral aliasing and spoofing characteristics in different deployment scenarios.

*3. Evaluation:* To evaluate the efficacy of PCFs, we use a theoretical model later validated by real traces from different traces. We also show our experiences in monitoring an OC-48 traffic stream. For example, we were able to identify about 517 flows out of a total of 30.36 Million in a day.

The organization of this paper is as follows. In Section 2 we describe the kind of scanning and DoS attacks we consider in this paper. Section 3 introduces a scalable mechanism to detect Partial Completion Attacks called PCFs and describe a theoretical analysis that shows why it is resilient to behavioral aliasing and (in some deployments) to spoofing. In Section 4, we describe experimental evaluation of PCFs for scalable monitoring and detection of partial completion attacks, scanning based attacks. Section 5 contains related work followed by by conclusions in Section 6.

## 2. ATTACK CLASSIFICATION

Although, there are many different kinds of network based attacks, we restrict our focus in this paper to analyzing the following three different types of attacks with respect to scalability using the framework developed in the previous section – *Partial Completion Attacks* (e.g., TCP SYN Flood Attacks), *Attacks that use Scanning* (e.g., TCP Portscans, Worms such as CodeRed), *Bandwidth and Flooding Attacks* (e.g., Fraggle attacks, Reflector attacks). This list is by no means exhaustive; however, we use these canonical examples as a first step towards understanding the difficulties of developing scalable mechanisms for intrusion detection in the network. In this section, we present an overview of these attacks in more detail.

## 2.1 Partial Completion Attacks

Such attacks are also known as Claim-and-hold attacks. The basic theme in these attacks is to grab a precious re-

---

[1]While this might argue for moving detection close to sources, this does not work well either because often we have little control over rogue sources and their networks, and source networks may not provide a sufficient vantage point to detect distributed attacks in which each source only contributes only a small fraction.

source and not release it thereby denying service to legitimate clients. The classic example of a Partial Completion attack is Syn-flooding.

In Syn-flooding, the attacker initiates several connections to the server by sending TCP SYN packets with spoofed IP addresses and never terminates any of these connections. In a variant called Naptha[24], the attacker initiates a connection and finishes the initial three way handshake, but does not do any further activity forcing the connection to time out. In both these attacks, we can see that the precious resource namely, connection memory, is claimed but never released.

While attackers have undoubtedly moved on to Application level DoS attacks such as on Google's query attacks, TCP Flood Attacks remain an important class of attacks[2]. There are many ways to detect and defend against Syn flooding attacks[9, 28, 11, 34, 41]. These detection mechanisms usually rely on keeping track of individual connections on a per-flow basis making them hard to scale.

Most classical solutions are deployed close to the end-host, and thus assume the ability to observe both directions of the traffic. In the network (e.g. a core router) however, it is not often possible to look at both directions of the TCP connection due to a wide-spread prevalence of assymetric routing[36]. Hence, detection of even the simplest of known attacks in the network *with access to only one direction of traffic* is a further challenge.

## 2.2 Attacks that do Scanning

Host scanning represents an important component of several attacks including most worm epidemics[44, 47, 43][3]. Thus several recent worms such as Code Red-II[2], Nimda[3] etc., propagated by scanning other vulnerable hosts in the Internet. A second example is probing for backdoors installed on various machines either installed during worm infection or by other means such as viruses. Such activity also exhibits scanning behavior. Finally, horizontal (multiple hosts and same port) and vertical (one machine, multiple ports) port scans are often performed by attackers as preliminary reconnaissance to identify a large number of vulnerable hosts in the Internet. Henceforth, we refer to these myriad activities as just *scanning*.

In this paper, we focus on scalable detection of TCP scanning although some attacks use UDP scanning. In scanning for TCP ports, the attacker sends several SYN packets to various destinations. If that destination has a listening server on that port, the connection is established and torn down as with any other connection. However, if either the destination or the port is non existent, the connection is never established. Thus watching at a network vantage point during scanning, the detection device can observe a number of SYN packets probing a particular port with no corresponding FIN packets. This requires correlation across a set of packets.

Note that besides UDP scanning, there are several other forms of scanning that we do not consider in this paper. For example, stealthy and slow portscans are much more difficult to detect. It seems clear that if the scan rate of a source is sufficiently slow that is no more than other benign sources, no scheme can detect such a scan scalably.

## 2.3 Bandwidth Attacks

Finally, the third kind of attacks we discuss in this paper are what are commonly called *bandwidth attacks*. In such attacks, an attacker or a set of compromised slaves (zombies), continuously pound a victim with a large number of packets, crippling normal services. In other such attacks, the attacker can take advantage of other stations to amplify the magnitude of traffic directed towards a particular destination. Smurf[1], Fraggle, and Reflector attacks[38] fall into this category. The common theme in all such attacks is pure traffic volume.

# 3. DETECTION OF TCP SCANS AND PARTIAL COMPLETION ATTACKS

It is fairly immediate to see that bandwidth attacks have scalable solutions using existing techniques such as MULTOPS [18], sketches [12, 17], or multistage filters [13]. Tools such as Autofocus [14] also detect multi-dimensional heavy hitters although using off-line analysis of the traces. These existing techniques rely on hash-based aggregation(Sketches and Multistage filters) or prefix aggregation (MULTOPS): for example, in multistage filters, flows are hashed to index into a set of buckets in different stages using different hash functions. It is easy to detect such bandwidth attacks in scalable fashion by restricting carefully the traffic types fed to these filters. For example, Smurf attacks can be detected using a multistage filter that finds destinations which have a large amount of ICMP messages relative to other destinations.

On the other hand, applying techniques such as sketches, MULTOPS, or multistage filters do not work well to detect TCP Flood Attacks since the traffic volume of SYNs (especially early in the attack tree) and scans may not be large enough compared to the volume of benign traffic. Sampling [33, 13] is yet another technique that has been traditionally employed for reducing the memory and processing needs for traffic monitoring. Fundamentally, sampling also works best for bandwidth attacks because an attack with a large traffic footprint is more likely to be sampled. However, it is not at all clear how sampling can be used to detect partial completion and scan-based attacks which have much smaller traffic footprints.

Given the importance of these low traffic attacks, in this section we introduce a new data structure — *Partial Completion Filters (PCFs)* that can detect both scanning attacks and partial completion attacks even when they correspond to small traffic volumes.

## 3.1 Partial Completion Filters

A *Partial Completion Filter (PCF)* consists of parallel stages each containing hash buckets (Figure 1) that are incremented for a SYN and decremented for a FIN. The SYN and FIN packets are presented here just as an example, but typically it could be increment for any open parenthesis and decrement for the corresponding closed parenthesis. It is easy to see that the expected value of each of these counters is 0. Thus, if a destination (or source) hashes into buckets with large counters in all stages, it seems plausible that the destination is being attacked.

---

[2]Recent DoS attacks on SCO[6] for example, used SYN flooding using spoofed sources[31]

[3]Note that there exist worms that spread through other means and hence do not exhibit scanning; for example, MyDoom spread through email and a Kazaa vector[5].

Figure 1: Partial Completion Filters

Unfortunately, this argument ignores random variation. As in the well known drunkard's walk (random walk), even if the average progress after $N$ steps is zero, the standard deviation is $O(\sqrt(N))$. Thus a benign bucket may have fairly large positive counters (causing false positives) while a bucket containing an attack may be be pulled down to zero (causing a false negative). The tricky part is to show that *both* false negatives and false positives stay within control for reasonable parameter values.

One might hastily conclude that PCFs are the same as multistage filters first proposed in [13] to detect heavy-hitter flows in the network. This is not true for the following three reasons:

*1. Non-monotonicity:* In multistage filters (and in fact in all Bloom Filter[10] variants), the counters are only incremented and never allowed to be negative.

*2. False negatives:* Bloom Filters and Multistage filters have only one-sided errors; there are no false negatives. Unfortunately, since PCFs allow counters to decrease, they can cause false negatives.

*3. Different Analysis:* The analysis of PCFs using the Central Limit theorem (Section 3.2) is very different from the simple counting argument for multistage filters.

## 3.2 PCFs and Behavioral Aliasing: A Theoretical Analysis

In this section, we provide a theoretical analysis that allows us to predict the behavior of PCFs in real network settings. The first goal of such a model is to show that PCFs exhibit very little behavioral aliasing — i.e., attacks cannot aggregate to cause either false positives or false negatives. A second and no less important goal of the analysis is to determine appropriate parameter values (e.g., number of stages, memory in each stage, threshold etc.).

The analysis is in four parts. In Part 1, we will present the abstract model of PCFs using generalized up-down counters. In Part 2, we will use the Central Limit theorem and tail bounds on Gaussian distributions to bound the false negative and false positive probabilities, which in turn determines the operating range of PCFs. In Part 3, we identify how to use PCFs to detect flows that greater than a given threshold. Finally, in Part 4, we analyze the false positives and false negatives in the presence of other bad flows.

***Part 1, Generalized Counter Model:*** In the Generalized Increment-Decrement Counter model, on every trial, an event (a set of fields in a packet) is hashed using different (perfectly random) hash functions to choose one counter for each stage. An unbiased coin is tossed and if it is a head(tail), the counter is incremented(decremented) by one.



Figure 2: Part 1 of the Analysis argues that Benign but Malformed Connections contribute a FIN or SYN to an interval with equal probability.

At the end of the experiment, events that hash to buckets (that are over a specified threshold in all stages) must be output. Detection of a partial completion attack such as SYN-flooding represents an instance of this model with an increment(decrement) for SYN(FIN) and output those flows with larger number of SYNs than FINs.

In the case of SYN-FIN correlation to detect SYN-floods, an event is the combination of the Destination (or Source) IP address and Destination (or Source) Port. This event is hashed using multiple hash functions to index the counters corresponding to different stages. In case of a SYN, the counter is incremented; in case of a FIN(or RST), the counter is decremented. At the end of a sufficiently large amount of time called a *measurement interval*, flows that hash to "large" counters in all stages are passed to a further level of processing for more careful examination. Clearly, if too many flows are passed for further examination the filter is not working well and memory needs will grow commensurately. Note also that at the start of an interval all counters are set to zero.

*Benign but Malformed Connections:* Clearly, the buckets of bad flows (i.e., destinations subject to DoS attacks or sources doing scans) will have a large positive value. However, we have to worry about benign but malformed connections (Figure 2) that can add spurious "noise" to counters. There are three important cases.

First, a connection may be long-lived, in which case it contributes (see Figure 2) its SYN to one measurement interval, and its FIN to another measurement interval.

Second, a connection may retransmit its FIN. However, as a first-order approximation, we can assume that a connection is equally likely to retransmit its SYN. In practice, TCP has a built-in asymmetry that makes SYN retransmissions happen slightly more often than that of FIN retransmissions. After using our first-order model (with equal retransmission probabilities), we show how this small bias can easily be corrected for as shown later in Section 4.1.

Third, route churn may cause the SYN to be seen but not the FIN, but in that case, during another interval due to another route churn, it might see the FIN but not the SYN. On average, a set of measurement intervals should be able to smooth out this noise. In [50], the authors have experimentally verified that the routes are stable on the scale of a few minutes. So, we believe that the noise generated due to route churn is not really significant. Nevertheless, our analytical model captures this effect as well.

Observe that in each of these malformed but benign cases, the anomaly is equally likely to add an extra SYN as it is to add an extra FIN to an interval. (Well-formed benign connections that add both a SYN and a FIN to an interval have a net contribution of zero.) Given that each interval is a random experiment in time (and this can be enforced
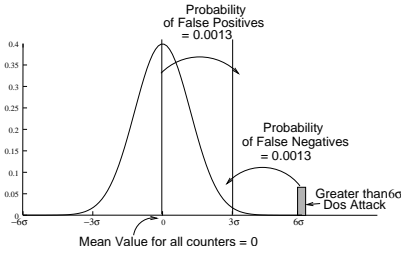
**Figure 3: Part 2 of the Analysis uses the Gaussian approximation.**

by some random dithering of interval start times), we claim that benign connections can be modeled as contributing, in the general-case, either a SYN or FIN with probability 0.5. This allows benign but malformed connections to be reduced to the generalized counter model.

*Part 2, Estimating Noise range of PCFs:* Without loss of generality, since the counters are chosen at random, let us consider one particular counter. Let $X_i$ represent the random variable that represents the value added to the counter in the $i$th trial. $X_i$ is 1 with probability 0.5 and is -1 with same probability. The expectation of $X_i$ say $\mu$ is 0, and standard deviation $\sigma$ is 1. After $n$ trials, we are interested in what the final value of $X = \sum_{i=0}^{n} X_i$ is. This represents the current counter value assuming that the counter started at zero. In the case of SYN-FIN correlation, this number represents the number of excess SYNs or FINs that have arrived for this bucket.

The exact distribution for counter values follows a binomial for which it is difficult to estimate tail probabilities in closed form. Fortunately, for sufficiently large values of $n^4$, the Central Limit theorem assures us that the binomial distribution can be approximated using a normal distribution. Thus we can use cookbook Normal tables to obtain confidence bounds on the probability that the counter value is above a particular threshold. If $X = \sum_{i=0}^{n} X_i$,

$$Pr\left[a \leq \frac{X - n \cdot \mu}{\sigma \cdot \sqrt{n}} \leq b\right] = \frac{1}{\sqrt{2\pi}} \int_a^b e^{-z^2/2} dz$$

As if we choose, $a = -3$, $b = 3$, also, $\mu = 0$, $\sigma = 1$, then

$$Pr\left[|X| \leq 3\sqrt{n}\right] = 0.9987$$

Note that these bounds can be found in statistical tables for Z-ratios (e.g., [26]). Therefore, counters of buckets containing only unattacked destinations should approximately lie within 3 times the standard deviation times $\sqrt{n}$. For example, in a measurement interval suppose there are 10 Million SYN/FIN packets. Also, let the number of buckets be 3000. The expected number of trials per bin would be 10 Million divided by 3000 which is approximately 3300 i.e, $n = 3300$.

From the above probability calculation, the probability that the counter value is less than $3 \cdot \sqrt{3300} = 172$ is rather large, and in fact 0.9987. So, even if all the connections were benign, the counters can be between -172 and 172 with very high probability. This determines the operating range of PCFs. Note that we can build more sensitive PCFs by

---
[4]And for reasonably large measurement intervals on real traffic, $n$ is indeed large enough

choosing a larger number of buckets so that the noise range becomes smaller. In fact the noise range of PCF is inversely proportional to $\sqrt{n}$, where $n$ is the number of buckets in PCF.

The probability that a benign flow maps to any counter greater than the calculated noise of $3\sigma$ is 0.0013. In order to reduce this even further, we add $s$ stages to the PCF. The probability that in $s$ stages the counter value is greater than 3500 (as in the previous example) becomes $0.0013^s$. For $s = 3$, the probability is $2.197 \times 10^{-09}$ which is extremely small. This bounds the noise range of the filter to $3\sigma$ with very high probability.

Another reason to add stages is to reduce the probability of false positives that occur because an unattacked destination hashes into a bucket containing an attacked destination. If there are $c$ concurrent destinations being attacked, this probability is $c/B$, where $B$ is the number of buckets. For example, if $c = 3$ and $B = 3000$, the probability of this happening at one stage is 1/1000. However, the probability that this happens in say 3 stages is $10^{-9}$. We will discuss this in more detail in Part 4 of our analysis.

Therefore, if the number of stages increases, the probability that a benign flow maps to all counters greater than $3\sigma$ drops *exponentially*. However, false negatives, i.e, the probability that a flow greater than $3\sigma$ goes undetected increases *linearly*. This is because the flow can map to counters all of which have negative noise that cancels out the contribution of this flow thereby appearing benign. The more stages we use, the more likely is it that at least one counter stage falls below threshold for a particular attack flow thereby going undetected.

Choosing the appropriate threshold and number of stages represents a trade-off between false positives and false negatives. For example, in Figure 3 for one stage of the filter, we assume that if the attack flow is at least 6 standard deviations from zero in its imbalance between SYNs and FINs, then the balancing strategy suggests placing the threshold at 3 standard deviations. Thus false positives would require, a shift of 3 standard deviations to the right of zero (Figure 3), while false negatives would require the same shift to the left.

Due to exponential decrease in the false positive probability and linear increase in the false negative probability, a small number of stages can drastically decrease the false positive probability without causing the false negative probability go too high. For example, we show later (in Section 4.1) that choosing 3-4 stages often reduces false positive probability while keeping false negative probability low enough.

*Dynamic setting of Threshold T:* As we have seen, the noise range of the filter is dependent on the number of packets in the given measurement interval. How then can we dynamically adapt this threshold ? There are many ways we can approximate the exact count of packets in a given interval. One way is to use the count of number of packets in the previous interval as a means to estimate the threshold. If the traffic does not exhibit drastic change in different intervals (if the intervals are suitably small), then this estimate of the threshold remains correct. One other way to choose this threshold is based on past history of this particular link. However, if we explicitly set the threshold to a reasonably large value, as we show in the next part of the analysis, there would be no need to dynamically adjust the threshold.

In the next part of the analysis, we discuss the analysis of

PCFs when they are used to detect flows which have SYN-FIN imbalance greater than a particular given threshold $T$.

**Part 3, Using PCFs to detect flows greater than a threshold $T$:** Earlier, we have obtained a theoretical bound on the noise that PCFs are susceptible to. Any flow that hashes to a counter which has higher than expected noise could potentially be malicious. In real life, we are often interested in flows that are much greater than this noise. PCFs cannot be used to identify flows that lie in the noise range without having too many false positives. Hence the threshold T, at the minimum has to lie outside the noise range.

The noise arising from other benign flows is additive with that of a malicious flow that hashes to these buckets. Hence, a flow of size $s$ can hash to buckets that lie between $s - 3\sigma$, and $s + 3\sigma$ with high probability. Hence, if we choose the PCF threshold to be $T$, then false negatives, whereby a flow of size greater than $T$ goes undetected increases. So, we should choose the PCF threshold to be at least $T - 3\sigma$. However, this increases the false positives since now, a flow of size $T - 6\sigma$ can, due to the positive noise, appear as a malicious flow with high probability.

From this analysis, we can see that if we are interested in flows greater than size T, choosing a threshold value of $T - 3\sigma$ allows us to guarantee two properties. Firstly, PCF flags all flows with greater than T with high probability. Secondly, if PCF flags a flow, the flow is at least of size $T - 6\sigma$ with high probability.

**Part 4, Estimating False Positives and False Negatives in the presence of attacks:** So far, we have analyzed those cases where we found the false positive and false negative probability in the presence of one malicious flow. In the presence of more number of attack flows, a portion of the buckets appear large increasing the chances of false positives. We now estimate false positives in the presence of a number attack flows.

For the analysis, let us assume that there are $b$ bad flows. It is easy to prove that the expected number of buckets to which these $b$ bad flows hash to is $b$ (assuming that $n$, the total number of buckets, is much larger than $b$).

Any flow that hashes to one of these $b$ buckets in all the three stages is deemed a bad flow. The probability that a flow hashes to one of these bad buckets is given by $b/n$. For $k$ stages, the probability is $\left(\frac{b}{n}\right)^k$. The expected false positives now is given by $\left(\frac{b}{n}\right)^k \cdot f$ where $f$ is the total number of benign flows, namely the bad flows subtracted from total flows.

For example, suppose the total number of flows were say 250,000 out of which 500 of them were genuinely bad flows. Let us assume that the total number of buckets is 1000 per stage and there are three stages. The expected number of buckets into which these 500 bad flows hash to is approximately 500. Therefore the expected number of false positives is $(500/1000)^3 * 249,500 = 31,187$. If there were only 100 bad flows, then the number of false positives is only about 250. As we can see from this analysis, the number of false positives increases super-linearly with the number of bad flows. In cases where the number of bad flows is only around 20, the number of false positives decreases to close to one.

A similar analysis applies to false negatives. In the presence of a reasonable number of flows with larger negative SYN-FIN differences, a portion of buckets become unreasonably small. Any flow that maps to these buckets would not be detected. However, this case is less often, since SYN packets usually dominate FIN packets for any flow.

The bottom line from the analysis is that one can tune the architecture appropriately to obtain reasonably high probabilities of catching a partial completion attack or a TCP scan, while making sure there are not too many false positives. Note that despite the small chance of missing an attack, if several routers in the attack path are using this scheme, then it becomes increasingly likely that one of them will catch the attack.

## 3.3 Applying PCFs to detect Partial Completion and Scanning Detection

**Notation:** We use PCF(A, B, C) to denote a PCF that increments(decrements) on a TCP packet with flags A(B), and uses C as the field(s) on which the packet is hashed.

*1. Partial Completion Detection:* For the detection device, the key abstract behavior that signals a SYN Flood to a destination is the presence of a *destination* that receives a large number of SYNs from various sources.[5] Thus a PCF(SYN, FIN, <DIP,DP> can be used to scalably detect a TCP SYN Flood attack by hashing based on *destination* IP address, port pairs.

In the network, if we assume that the detection mechanism cannot see both directions of the traffic, once again it is easily spoofed. We show how to make SYN Flood detection spoof resilient using *reverse path* deployments in next section (Section 3.5), though this will require an inversion: the trick is to hash source addresses (as opposed to destinations) in the reverse path to identify victims. This is because, without collusion from inside the network, the attacker cannot force the victim to send FIN packets.

*2. TCP Scanning Detection:* At a network vantage point, during TCP scanning activity such as portscan, a detection device can observe a large number of SYN packets to a particular port but with no corresponding FIN packets that correspond to legal tearing down of the connection.

Therefore, for the detection device the key abstract behavior that signals a TCP scan is the presence of a *source* that sends a large number of SYNs to various destinations and destination ports without sending a corresponding FIN. Thus, a PCF(SYN, FIN, <SIP>) can be used to scalably detect a TCP scan by hashing based on source IP addresses and zeroing in on such sources that have a large SYN-FIN imbalance.

In the network, if we assume that the detection mechanism cannot see both directions of the traffic as we have assumed, then PCF methods are easily subject to spoofing. One approach is to ignore spoofing because most attackers employ tools such as NMAP[16] that do not spoof today; however, we will show to make detection spoof-resilient using bidirectional deployments (in scenarios where we can see both directions of traffic) in Section 3.5. Next, we apply PCFs for scalable monitoring of partial completion and scanning attacks in the network.

---

[5]Note that unlike a number of other approaches like backscatter[32] our approaches work regardless of whether the attacker employs address spoofing. Thus we can detect DDoS attacks that use a large army of zombies (increasingly common today) and attacks via reflectors, all of which often use true IP addresses.

## 3.4 Applying PCFs for Attack Monitoring

PCFs can be applied to characterize attack flows in an online fashion, in contrast to current approaches that rely on passive traces. Using PCFs, we can identify attack flows (sources and destinations using different PCFs or can be combined into one by hashing each packet twice), count the estimated size and duration of attacks in a scalable fashion. We will show in Section 4.2.2 our experiences with PCFs in scalable characterization of attack flows. Note that the attack flows triggered by PCFs can also be routed to a sinkhole[19] for further forensics, traceback and so on.

As we have seen earlier, PCF(SYN, FIN, <DIP,DP>) based on destination IP address, port pairs can detect the destinations under attack in a scalable fashion. We call this the *forward* path of the attack since we infer DoS activity based on the attack packets going towards a victim.

PCF(SYN, FIN, <SIP,SP>), based on source IP address, port pair can be effective in monitoring based on the *reverse path* of the attack. This follows the fact that a victim under attack generates several SYN-ACK packets but no corresponding FIN packets (since connection is never really established, even if it did doesn't terminate). Together, the forward and reverse path PCFs can aid in scalable monitoring and characterization of partial completion based DoS activity with an ISP network domain. The forward path PCF however is spoofable, but the reverse path PCF is *spoof-resistant* as we discuss later in Section 3.5.

Network Telescopes[32] based on "Backscatter analysis" are currently employed to infer world-wide DoS activity in a scalable fashion. Network Telescopes however cannot detect certain types of attacks such as those that do not employ random spoofed source IP Addresses. For example, reflector attacks[38] where a large number of attackers send packets to listening servers with source IP address spoofed with that of a victim thus generating a flood of responses from these listening servers directed to victim, can never be detected using a Network Telescope. PCFs on the other hand can easily detect such attacks.

A quick comparison between Backscatter Analysis and reverse path PCFs is given below, while actual trace driven comparisons between the two schemes are presented later in Section 4.2.3.

*Reliance on Source Address Spoofing:* Network Telescopes can only detect attacks that employ spoofed source addresses. PCFs, on the other hand, do not depend on this feature. Thus PCFs are capable of detecting Reflector based attacks, and other attacks using "zombies" that do not employ source address spoofing.

*Reliance on TCP:* Network Telescopes can observe source address spoofed attacks using a wide variety of protocols including TCP, UDP, ICMP as opposed to PCFs that are designed only for TCP.

*Reserved IP Space:* Network Telescopes have an inherent tradeoff between detection time and the size of the unused private address space they watch: the larger the space, the faster the detection. PCFs on the other hand, do not require any reserved IP space.

*Geographical Scope:* Network Telescopes receive the aggregate traffic sent from any source under attack that uses address spoofing towards the private IP space reserved for it. Hence, Network Telescopes have much wider geographical scope in contrast to PCFs that have only local scope.

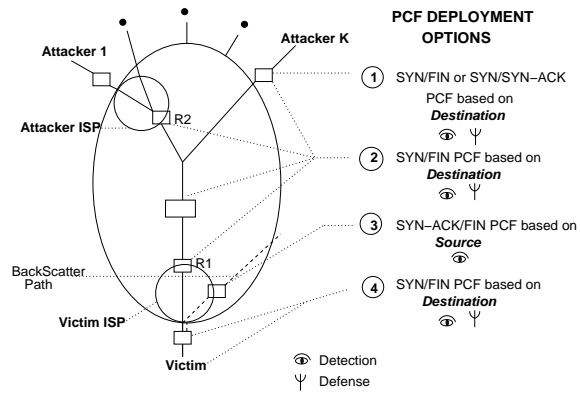As we can see from the comparison, PCFs can be a viable



**Figure 4: Deployment Options**

and scalable complementary solution for general (i.e., no assumptions about the use of address spoofing, no need for a large unused address space) TCP-imbalance detection that passes through the detection device. It can be widely deployed without any issues. On the other hand, backscatter is invaluable for the global detection of DoS attacks (based on fake source addresses) at a *few* monitoring points. A combination of Network Telescopes and PCFs can scalably detect Partial Completion Attacks in the network.

## 3.5 Deployment and Spoofability of PCFs

In this section, we apply PCFs to Partial completion attacks and scanning attacks and discuss their spoofability properties.

***Partial Completion Detection:*** Since spoofing is dependent on the particular instance of deployment, we first discuss briefly various network deployment options of interest, and discuss the spoof resistance of each option. Figure 4 shows these deployment options.

- *Near sources:* A PCF(SYN, FIN, <SIP>), where SIP is the Source IP Address, can easily identify those sources that are generating many SYN packets but no FIN packets. This deployment scenario could potentially be used to monitor sources of attacks. However this is subject to spoofing using a variety of mechanisms. Sample spoofing mechanisms include the sending of fake FIN packets, the use of carefully crafted TTLs for FIN messages that only crosses the monitor but never reaches the victim. Attempting to defend against the dazzling variety of spoofing options is a virtual impossibility; we choose instead to finesse this issue by espousing reverse path PCFs because it is harder for the attacker to control the reverse path.

- *Outgoing/Incoming Edge of an ISP or Customer Network:* A PCF(SYN,FIN, <DIP,DP>), DIP, DP being the destination IP address and port respectively, could be deployed at the edge of an an ISP network or a large customer network to detect attacks to destinations that originate or are directed towards their domain. This constitutes the *forward path* of the attack (which is controlled by the attacker). The attacker can again spoof using either FINs before SYNs, or using FINs with TTLs that expire early. Note that this particular spoofability arises from uni-directionality of

traffic. If we assume a bi-directional model, then PCF (incoming SYN, outgoing FIN, <DIP,DP>) is *not spoofable* since it is hard to forge a packet in the opposite direction.

- *Outgoing Edge of an ISP or Customer Network*: PCF (SYN, FIN, <SIP,SP>) can detect destinations in the domain under attack. This follows the fact that a victim under attack generates several SYN-ACK packets but no corresponding FIN packets (since connection is never really established, even if it did doesn't terminate). This is *spoof-resistant* [6] since the attacker cannot force the victim to generate a FIN packet without legally tearing down the connection. We call this the *reverse path* of attack. Later in Section 4.2.3 we focus on this particular deployment of PCF to scalably detect backscatter in the network

- *Near destinations:* This case is similar to that of the incoming edge of an ISP or customer network and hence spoofable in uni-directional but not spoofable in bi-directional detection model.

Fundamentally, *any approach which observes only the forward direction (Active SYN-FIN(ACK)) of TCP traffic (as is often the case in the network) to a particular destination is susceptible to spoofing.* This is because an attacker can always manufacture packets in the forward path (some of which can die before they reach the victim via low TTLs etc.) that look exactly like valid TCP connections.

**Scanning Detection:** PCF(SYN,FIN,<SIP>) can detect attacks that involve scanning hosts for vulnerabilities. However, a clever scanning approach can easily spoof the PCF (or any other full state approach) that only sees one direction of the traffic. On the other hand, by assuming a bidirectional model (where PCF sees both directions of traffic, usually true at the edge), we can scalably detect hosts that are scanning for other hosts.

In the bi-directional model, the PCF increments for the SYN packet going out and a FIN packet coming in to detect scanners inside the network. In order to detect a scanner outside the network, PCF correlates between in-bound SYN packet and out-bound FIN packet. Note that this deployment of PCF is not susceptible to spoofing since, the originator of the scan cannot force a non-existent host to generate a FIN packet in the opposite direction. Similarly, PCF(SYN, FIN, <SIP, DP>) can detect horizontal scans for specific ports in a scalable fashion too and can be made *spoof-resistant.*

# 4. MEASUREMENT ON REAL TRACES

We evaluated PCFs on a set of real traces that we obtained from two different ISPs A and B[7]. The traces are named ISP-A Dir-0, ISP-A Dir-1, ISP-B Dir-0, ISP-B Dir-1 corresponding to two different directions of traffic. All these links have OC-48 capacity.

The main aim of these experiments is two-fold. Firstly, we wanted to validate the theoretical analysis of PCFs when applied in real settings; any deviations from analysis can be used to tune and modify PCFs. Secondly, we wanted to gather experience using PCFs to scalably detect attacks (in the categories we restricted ourselves to earlier) on real network traces.

We follow the following terminology in the experiments. Firstly, a "flow" is any unique tuple over a subset of packet contents. For example, if the aggregation is over fields SIP and SP, all packets that bear (say) <192.168.10.1, 80> as the Source IP, Port pairs (perhaps with different Destination IP addresses or ports) is termed a flow. Note that the word flow in the usual sense is often referred to as a 4-tuple; for the lack of a better word, we used this somewhat loaded term in a different way.

A flow is said to be *correctly identified* if both the full state approach and PCFs identify that the flow has a value greater than the threshold. A flow on the other hand is a *false positive* if PCF indicates that the flow has a value greater than the threshold but using the full state approach we find that the flow does not have a value greater than the threshold. Similarly, a *false negative* refers to those flows that have not been identified using PCFs but were detected using a full state approach.

## 4.1 Part I: Validation and Tuning of the model

The generalized up-down counter model with randomized events that we introduced earlier in the theoretical analysis is an *approximation* of reality. So how accurate is it really? There are a set of issues that warrant adjustment to the model. In this section, we briefly discuss these issues and discuss how to correctly adjust the behavior of PCFs for real traffic.

In the first part of our results, we validate and tune the model to possible deviations from these assumptions. These potential idiosyncrasies in real traffic determines the operating range and usefulness of Partial Completion Filters. Unless explicitly stated otherwise, we used a measurement interval of one minute, and used a total of 5000 buckets for our evaluations.

Starting with a review of these assumptions, we now address a series of questions about the theoretical analysis and parameter settings (e.g., number of stages, measurement intervals etc.).

*Q1. What is the effect of Popular Buckets?* An obvious deviation (from uniformity) is that the number of packets that hash to a particular bucket is not exactly (as we assumed above) the total number of trials divided by the number of buckets. The distribution is clearly not uniform due to the existence of popular sites like Yahoo! etc, which have a large number of packets destined to them. Also, depending on the deployment location, this distribution can often get skewed. The number of such large buckets however, is typically small. Flows that map to these buckets might be falsely classified as attack thus increasing the number of flow records that need to be maintained to monitor these flows. In other words, the number of false positives increases due to these kind of flows, however, if we were to account for it. In the traces we have used however, the distribution was more or less uniform and hence, we do not address this in more detail in this paper.

*Q2. How does the asymmetry between SYNs and FINs affect the results?* In the theoretical model, we assumed that the expected value of each counter is 0 since the packet can

---

[6]More accurately, to spoof this scheme the attacker will need to have machines under its control *in the victim domain* that can send fake FIN packets. If the attacker already has such power, it is unclear why the attacker needs to stop at simple DoS attacks of victims.

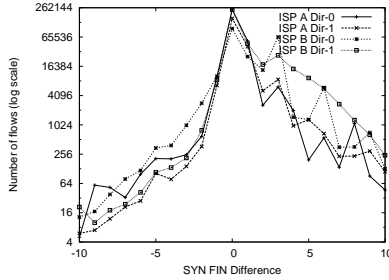[7]For anonymity, we do not present the actual names of the ISPs

**Figure 5: Flow distribution Per Source, Per Destination of the SYN-FIN, Counters. The y-axis represents the number of destinations(sources) that have the corresponding SYN-FIN difference on the x-axis**

be a SYN or a FIN with equal probability. However, in reality, this is not true. The expected value of the counters is close to zero but not quite zero. This is primarily due to the fact that there are potentially a larger number of SYN retransmissions than that of FIN retransmissions in a connection. This is plausibly because the round-trip delay has not been calculated when a SYN is sent (as opposed to a fairly accurate value when the FIN is sent); thus estimates that are too low will lead to more unnecessary SYN retransmissions.

Also, note that some connections are torn down by RSTs as opposed to FIN packets. One way to take into account this bias is to decrement for the RSTs along with FINs. We decided however to capture all these effects into a single parameter called "bias". This positive bias accounts for all irregularities that arise from SYN-FIN differences and can be appropriately set.

The problem with this positive bias is that, even a reasonable number of these flows with the positive bias can now aggregate to look like a bad bucket. Hence, in actual practice, the mean $\mu$ has to be set to a positive value. The analysis however remains the same, except that the threshold for detection has to be tuned based on the positive bias.

Figure 5 shows the number of <DIP,DP> pairs (any other kind yields similar results) for different SYN-FIN difference values ranging from -10 to 10. We chose 10 as our cut-off point since beyond this value, it is reasonable to expect that they are outliers. From Figure 5, we can observe that the area under the right half of graph (x=0 to x=10) is higher than that on the left side (x=-10 to x=0). The amount of positive bias that needs to applied to each counter per flow is equal to the weighted mean of the SYN-FIN differences.

Suppose all flows (e.g <DIP,DP> pairs) on an average have a positive bias of $x$ (average SYN-FIN difference for all flows = x). Then we have to modify the threshold to take into account this positive bias. In our counter model, suppose $f$ were the total number of flows, and $b$ be the total number of buckets. Then, the expected number of flows that map to a particular bucket is going to be $\frac{f}{b}$. But since, these flows have a genuine positive bias, the expected value of each counter is $\mu = \frac{f}{b} \cdot x$.

Therefore, instead of an expected value of 0 used in the normal approximation, if we use the new expected value (the standard deviation remains same), the probabilities would

still remain the same. In other words,

$$Pr\left[a \leq \frac{X - \frac{f}{b} \cdot x}{\sigma \cdot \sqrt{n}} \leq b\right] = \frac{1}{\sqrt{2\pi}} \int_a^b e^{-z^2/2}dz$$

Now again, for $a = 3, \ b = -3, \ \mu = 1, \sigma = 1$,

$$Pr\left[|X| \leq \frac{f}{b}x + 3\sqrt{n}\right] = 0.9987$$

Our new threshold value has to be modified to include this positive bias, resulting in a new term equal to $\frac{f}{b}x + 6\sqrt{n}$. Recall that $f$ is the number of flows, and $n$ is the total number of events per bucket (SYN-FIN packets in the case of partial completion attacks).

The theory above suggests that the model can be adjusted to account for this bias we observe in real life. In the next experiment, we evaluate the significance of this bias. Figure 6(a) shows the relationship between the bias and the false positive ratio. Note that the y-axis is on a log scale. For the traces we used in this Figure, we can see that a reasonable bias of 0.5 reduces the number of false positives dramatically, although we calculated empirically that the actual estimated bias for the ISP A and B traces was less than 1. Choosing a high value of the bias increases the threshold but the operating range of the filter becomes smaller.

*Summary:* We see that for some traces more than the others, bias effects the number of false positives significantly. Hence, this factor must be appropriately set (we assume henceforth a threshold of 0.5) to a reasonable value. For the set of traces we considered, a bias value of 0.5 worked well in all cases.

*Q3. How many stages are necessary in PCF ?* In this experiment, we evaluate empirically the number of stages required in the PCF to sustain a reasonable false positive rate. Figure 6(b) shows the variation of False Positives and Negatives with increase in the number of stages from 1 to 10. Again, the y-axis represents the average proportion of false positives to the total number of flows. From the figure as well as our theoretical analysis, 3 stages represents a good trade-off to reduce the false positives while keeping the false negatives to a minimum.

Beyond 3 stages, we can see that there is really little gain in the false positive rate. In addition, false negatives begin to increase. In fact, from the figure, 3 stages represent the "knee" of the false positives curve with diminishing returns from adding more stages. Henceforth, we operate with 3 stages for the rest of the paper for our evaluations.

*Q4. How good is the threshold we calculated ?* The threshold we calculated from the theoretical analysis only gives us a lower bound on the threshold beyond which the false positive probability becomes extremely small. This as we have pointed out before, represents the noise range of the filter. In this experiment, we wish to identify how good this lower bound is. We varied the threshold in proportion to the threshold calculated by the theoretical analysis from 0 to 2 in steps of 0.1. Figure 6(c) shows the variation in the false positives for different traces on log scale with different threshold values. Note that increasing the threshold automatically reduces the number of flows that we identify and hence correspondingly the false positives. However, when the threshold was moved from the left, exactly at the theoretical bound (that corresponds to an x-value of 1), we see
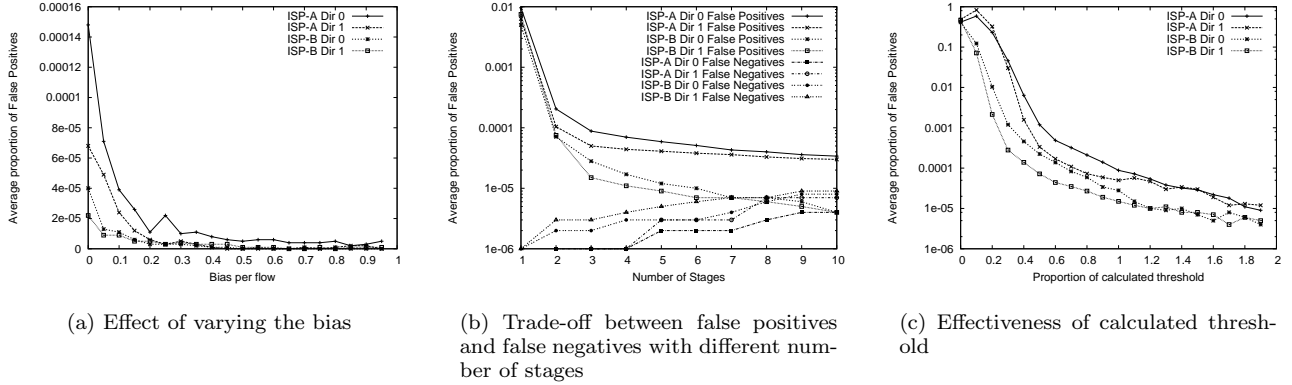
(a) Effect of varying the bias

(b) Trade-off between false positives and false negatives with different number of stages

(c) Effectiveness of calculated threshold

**Figure 6: Model validation, adjustment and tuning: We accommodate SYN-FIN assymetries using bias (a bias of 0.5 works well in practice). Three stages can reducing the false positives without giving up too much on false negatives. Finally, the calculated theoretical threshold value results in low false positives (determines the noise range of the filter).**

much smaller number of false positives as we have analyzed using the Gaussian approximation.

This empirically verifies that the threshold value we calculated using the theoretical analysis works well in practice keeping the false positives to a lower value.

*Q5. Why is the false positive rate still high?* A curious reader can immediately notice that in the earlier plots, even for the theoretical threshold, the false positive rate is much higher than that calculated theoretically in Section 3.1 for three stages $i.e. 2 \cdot 10^{-9}$. In the presence of a large number of flows that have a high imbalance in the counters, the number of false positives increases thereby translating into more memory requirements to monitor these flows. We found that the number of flows that were anomalous was close to 100. Using our analysis in Section 3.2, the estimated number of false positives in proportion to the total number of flows is about $(100/5000)^3$ which is close to $10^{-5}$. We can see that the false positive proportion is close to this value in the plots.

One way to reduce the false positives in this case is to add an extra threshold (beyond the theoretical threshold value including that calculated taking into account bias). This can potentially restrict the number of attacks of interest to a low value (only the heavier attacks will be identified – in fact maybe desirable) but it also reduces the number of flows that will monitored. We can even adaptively vary the threshold based on the amount of flow memory available. In [13], the authors suggest a heuristic algorithm that adjusts the threshold based on the amount of flow memory in the context of heavy hitter detection using multistage filters.

*Summary:* The false positive rate is higher in the presence of a large number of bad flows. A way to decrease the false positive rate to a reasonable value is to increase the threshold beyond our current theoretical value. This remains a compromise between the available memory resources, expected number of attacks and so on.

*Q6. How big should the measurement interval be ?* The next important issue is how to choose the measurement interval. A plausible first guess could be to choose a very small measurement interval. In any scheme, a small measurement interval, while facilitating fast detection, lacks a clear signature to infer an attack. A large measurement in-

terval on the other hand increases the time of detection. Usually the choice of the measurement interval is dependent on the scenario of application. We choose **one minute** as our measurement interval for all our experiments. However, we hasten to add that this can be varied depending on the situation.

## 4.2 Part II : Experience with PCFs:

In this section, we present our experience (attack identification and characterization) with PCFs in realistic settings. First, we deployed PCFs on an OC-48 link trace to identify malicious flows followed by other traces we have obtained. Thirdly, we compare our scheme with the Network Telescopes approach for scalable monitoring, after proceeding to detect scanning activity in our traces.

### 4.2.1 Experience with PCF on ISP-A OC-48 link

Through this experiment, we discuss briefly our experience with PCFs over large time periods (1 day). We set the PCF threshold to 150 (theoretical threshold is 60) , which allowed us to identify attacks of size greater than 2.5 SYNs per second. Recall once again that we use the term "flow" to represent any unique tuple over a subset of packet contents.

The total number of unique Destination IP addresses in the full trace was about 5.16 Million over the entire day. The total number of unique <DIP, DPort> pairs was over 30.36 Million. Similarly, the number of unique Sources, <SIP, SPort> pairs were found to be approximately 1.9 Million and 30.9 Million respectively.

We characterize our findings in three parts. First in Figure 7(a) , we show the per-interval summary of number of flows (unique destinations) found, correctly identified flows, false positives and false negatives In all, we have observed a total of 517 flows that were detected by PCF in the forward path. We had 6 false positives and 0 false negatives reflecting the efficacy of PCFs.

Secondly, in Figure 7(b), we summarize the cumulative number of attacks that PCFs found over a set of intervals including removing any flows that were identified in previous intervals (entire day consists of 1440 intervals, each interval with a duration of 1 minute). As we can see from the figure, initially the number of attack destinations is high after
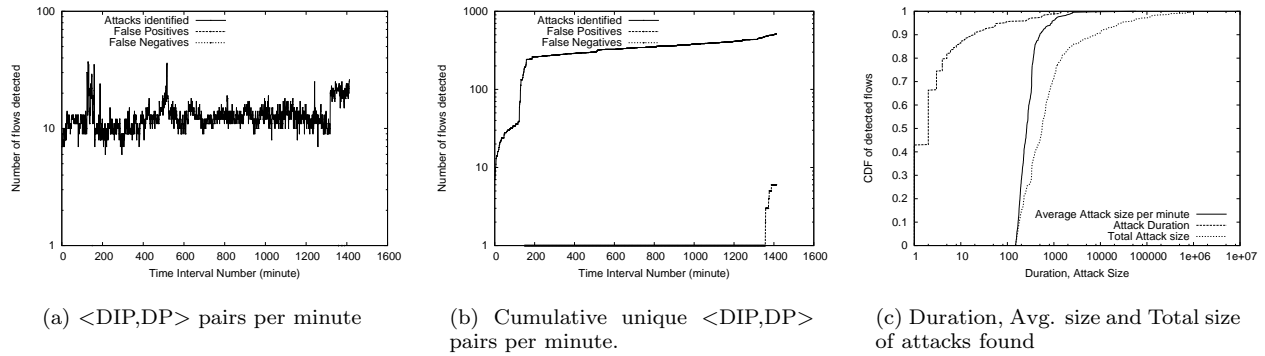
(a) <DIP,DP> pairs per minute  (b) Cumulative unique <DIP,DP> pairs per minute.  (c) Duration, Avg. size and Total size of attacks found

**Figure 7:** For a full day ISP-A trace, we show the detected flows, false positives and false negatives in every time bin (1 minute) in Sub-Figure (a). The number of false negatives in Sub-Figure (b) was 0 and false positives was very small (about 6). Further attack characterization using a CDF of durations, average size, and total size of these attacks in Sub-Figure (c)

which the number grows slowly. This is because most attacks (about 60%) last for more than 1 minute. This means on an average very small number of new attacks originate in any particular interval. Initially however, history is empty, therefore any flow detected is going to be new.

Finally, we show the CDF characterization of the duration, average size, and total size of the attacks identified in this trace. About 40% of the attacks found lasted for less than a minute and 85% less than 10 minutes. A similar observation has been made by the authors in [21].

*Summary:* These plots are illustrative of the fact that the number of "interesting" destinations and sources that exhibit anomalous behavior are like needles in a haystack. PCFs help pick these needles scalably and accurately. Such continuous monitoring to identify flows in the network also helps us understand the prevalence and dynamics of these kinds of attacks (at least on a smaller scale). Operationally, it appears that the total number of destinations or sources that need further monitoring has been brought down to an extremely small number with a negligible number of false positives.

### 4.2.2 Attack Characterization using PCFs on other traces

We now present the characterization of Partial Completion attacks that we detected using PCFs on both directions of ISP-A and B traces. Table 1 shows the number of attacks identified for each trace and the corresponding false positives and false negatives using both forward path and reverse path PCFs. As before, we have used a PCF threshold of 150 to identify these attacks.

Due to space limitations, we only show the characterization of the attack flows obtained using the reverse path PCFs. Figure 8 shows CDF plots that show the relative distribution of the attack flows PCFs identified in both directions. From the figure, we observe that in ISP-A traces, 30% of the attacks identified lasted through the complete duration of the trace (60 minutes). We have observed similar behavior in the forward path (not shown here) as well. A monitoring system based on PCFs would be able to characterize such attack statistics continuously. About 20% of attacks discovered in ISP-A trace had more than 1200 SYNs

per minute, roughly 20 SYNs per second constituting much heavier attacks.

### 4.2.3 Comparison of DoS Detection with state-of-the-art Network Telescopes

As promised earlier in Section 3.4, we conducted an experiment that provides empirical comparison between PCFs and backscatter analysis using Network Telescopes. We used a trace, BACKSCATTER obtained on Jan 22nd,2004 at 2pm from the CAIDA Network Telescope. Simultaneously, we obtained a trace from ISP A Direction 0 and Direction 1. The main intuition in comparison with the backscatter approach is that PCFs should be able to detect sources that are generating a lot of SYN-ACK packets in response to a SYN-Flood and hence should be aggregated based on sources.

On ISP-A direction 0 and 1, PCFs detected about 19 and 33 flows in the reverse path as shown in Table 1. At the same time, the telescopes detected about 776990 different flows. However, only three destinations were detected by both PCFs and Telescopes, that too in only one direction of the traffic.

Table 2 shows further description of these attacks (which have been referred to as A, B and C). The table shows the average number of SYN packets received in an interval duration of 5 minutes over 12 intervals (corresponding to one hour). Clearly, part of the issue is that the backscatter received a large number of SYN packets while the router trace showed a much smaller set of SYNs. We believe this occurred because the scope of any one router especially if it is a transit router (as in our case) is much smaller than that of the total backscatter. However, a set of these PCFs in all the peering routers would be able to capture all the backscatter originating from the domain. Unfortunately, we cannot verify this hypothesis with the traces available to us.

A second observation is that for <IP, port> pair C (as shown in Table 2), clearly, the backscatter could only see a very small portion, while the router could observe a much larger portion of the attack. This, we believe, is due to the fact that the telescope heavily depends on IP address spoofing in order to detect attacks. The higher the amount of address spoofing involved, the higher the amount of backscat-

(a) CDF of average SYN flood size per minute

(b) CDF of duration of attack

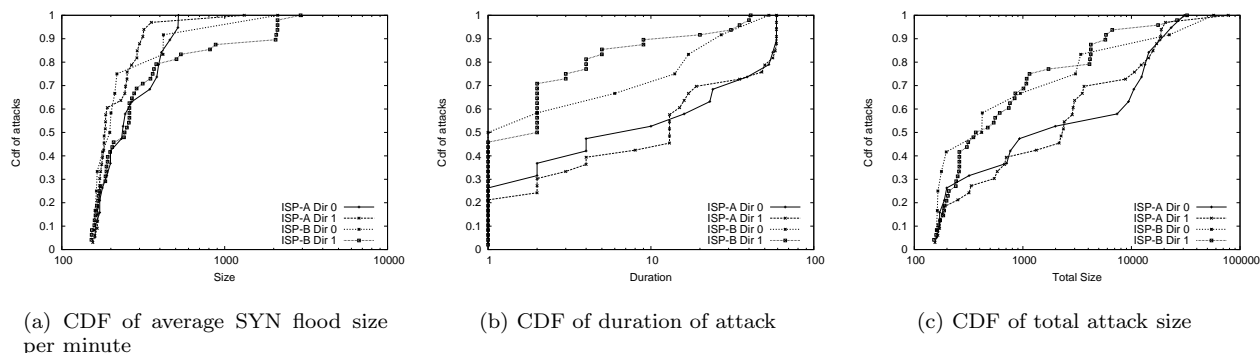(c) CDF of total attack size

**Figure 8: CDF of average attack size, total attack size, and duration of the set of attacks detected using PCFs on different traces in the *Reverse Path*.**

| Trace Code | Attacks identified | | False Positives | | False Negatives | |
|---|---|---|---|---|---|---|
| | Fw. Path | Rev. Path | Fw. Path | Rev. Path | Fw. Path | Rev. Path |
| ISP-A Dir-0 | 68 | 19 | 2 | 0 | 1 | 0 |
| ISP-A Dir-1 | 39 | 33 | 1 | 1 | 0 | 0 |
| ISP-B Dir-0 | 34 | 12 | 0 | 0 | 0 | 0 |
| ISP-B Dir-1 | 41 | 47 | 0 | 0 | 0 | 1 |

**Table 1: Summary of flows we identified using Forward and Reverse path PCFs on various traces using a threshold of 150.**

| | SYNs (ISP-A Dir 1) | SYNs (Backscatter) |
|---|---|---|
| $\langle$IP,Port$\rangle$ A | 359 | 12196 |
| $\langle$IP,Port$\rangle$ B | 309 | 14577 |
| $\langle$IP,Port$\rangle$ C | 498 | 0.09 |

**Table 2: Description of the three $\langle$IP addresses, Port$\rangle$ Pairs, that were commonly detected by both PCF as well as Backscatter.**



**Figure 9: Port scan detection using PCFs. Figure plots sources identified by PCF that $\langle$SIP$\rangle$ alone. The y-axis represents the number of scan sources identified by PCFs that have more than x destinations (for varying x)**

ter. We conjecture that this attack directed towards address port pair C did not employ enough spoofing.

We also found that a large percentage of attacks that we observed in the router trace were not observed by the telescope. This could be due to either of two reasons. One is that the attacks were mostly reflector attacks which did not employ spoofing as a method to bombard the victim. In these cases, the telescope cannot detect any backscatter. The other reason could be the presence of DoS attacks which are single or multiple source attacks that do not employ spoofing. We have manually verified the existence of both in our trace.

*Summary:* Comparison of the reverse path PCF reveals that both telescopes and PCFs can identify attacks the other can not. The telescope observes worldwide DoS activity that primarily employs packet spoofing. Any attacks such as reflector attack, or a DDoS attack with a set of zombies with no spoofing never reach the telescope. On the other hand, reverse path PCF is oblivious to address spoofing but suffers from a much smaller scope. The intersection set is rather small, making PCFs *a complementary solution* to backscatter for scalable attack monitoring. In fact, these two approaches can be combined to detect a wide variety of DoS activity.
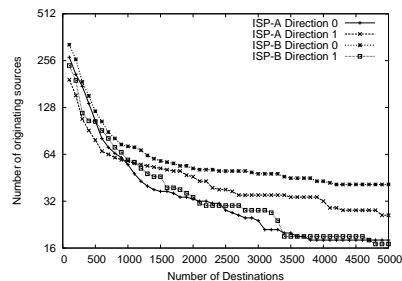
### 4.2.4 Scanning Detection

In this section, we use PCFs to scalably detect scanning in the network. We conducted two sets of experiments to validate our findings – one based on aggregation using Source IP Addresses, and the other using Source IP Address, Destination Port combinations as discussed in Section 2.2. In order to validate whether the sources we have found are indeed scanners, we counted the number of unique destinations found. Any source which generated SYNs but no FINs for more than a particular number of destinations we posit as being a true port scan for our comparison. The question at hand is whether PCFs can scalably detect such scanners. Note that PCFs based on $\langle$SIP$\rangle$, fundamentally observe sources which are generating too many failed connections generating SYNs but no FINs. This counting of destinations is only a step we can use to ascertain ourselves that these sources are indeed scanners.

This approach combines two traditional approaches of scan detection – counting events in a given time interval[20, 30], and observing failed connections[40, 37]. PCFs identifies sources that generate many SYNs but no corresponding FINs (corresponds to failed connections) in a given time interval. Note that in situations where there are a limited number of end-hosts such as in an enterprise, heavy weight schemes such as [23] might be practical. PCFs allow scalable and efficient detection of scans in situations where scalability is important (hence heavy weight approaches are impractical).

In order to establish that the sources identified by PCFs were really portscans, we plot the number of identified sources with increasing number of destinations in Figure 9. In other words, on the x-axis we vary the number of destinations in steps of 100, and we plot on the y-axis the number of sources identified by PCF that have sent SYN packets (but no corresponding FIN) to more than x destinations. From these figures, we can observe that a large number of the sources that have been identified by PCFs have failed connections to more than 500 destinations – portscans.

Due to lack of space, we only showed the results using a PCF on <SIP>. Results obtained using <SIP,DP> for horizontal portscans were similar.

## 5. RELATED WORK

The general notion of scalable attack detection has been addressed independent of our work, recently by Yaar et.al in [49]. However, their work requires routers to implement marking along with some header changes to support marking of packets. In [48], the authors also propose a Path Identification scheme for efficient identification of the sources of DDoS attacks. Their work builds on other traceback related schemes (see [48, 49] for further references on Traceback).

In [45] and [46], the authors propose simple stateless SYN-FIN and SYN-SYNACK counters in the last hop IDS solutions to detect the presence of SYN flooding presence. While their approach does a preliminary level of aggregation, it does not aggregate across destinations (or sources). They also do not consider the issue of spoofing or suggest solutions to this problem.

For SYN Flooding defense, there are several end-host solutions that have been proposed before. These include SYN-Cookies[9] and SYN-Cache[28]. These have been helpful but have not been universally deployed. "BackScatter Analysis" was proposed in [32]. This technique uses the reverse path attack properties of SYN-Floods to infer Denial of Service activity around the world. Since it is relevant to our work, we provided a comparison earlier in the paper.

Vendor based solutions such as SynKill[41], Netscreen[35] or free open source IDS tools such as Bro[37], Snort[30] can be used to detect SYN floods, port scans, but they (as far we can ascertain) employ per-flow state. A clever hop-count based method to detect spoofing in general has been proposed in [22]. Here, the main intuition is that spoofed packets typically have a wrong TTL value and hence should be identified by a previously detected TTL value for a particular IP address. This technique is fairly resilient to spoofing. MULTOPS[18] is a data-structure maintained by each network device that detects bandwidth attacks by the significant, disproportional imbalance between packet rates going to and coming from the victim or attacker.

Most scan detection techniques[20, 30] in the literature are based on detecting N events in T seconds. Another approach[40, 37] relies on failed connections as a better indicator of a scan. Leckie et.al[27] use probabilistic approaches to estimate the degree to which a given local IP address is unusual. SPICE[42] is an offline analysis algorithm to detect stealthy scans (scans which are of low rate) and cannot be performed scalably in the network. A recent paper by Jung et.al[23] apply threshold based random walks for fast portscan detection. The need to track for each remote host the different local hosts to which it has connected to makes the scheme unscalable.

## 6. CONCLUSIONS

It appears to be widely perceived that detecting intrusions scalably within the network is a bad idea. Unfortunately, that causes security devices to choose between performance (which requires low memory) and completeness (which appears to require per-flow state). This paper is a gentle first step towards suggesting that this tradeoff may not be as Draconian as is commonly thought. While the general problem is still very hard (and indeed for attacks such as evasion attacks, we believe that aggregated solutions cannot work without causing unacceptably high false positives), our paper shows some progress for bandwidth-based and partial completion DoS attacks, and scan-based attacks including worms.

This is fortunate because market researchers [39] have already begun to warn that increases in total ownership costs for end node and edge solutions require the network to play a proportionately larger share in detecting and combating network intrusions. This paper explores this possibility in the specific context of DoS attacks and scan attacks. While we have not harped on this point, doing DoS detection in the network also finesses the need for traceback and/or manual intervention, and allows enterprise networks and ISPs to automatically filter out attacks before they enter (or leave) their networks.

After quickly arguing that DoS attacks based on traffic volumes can be detected by existing techniques, we focused on the task of detecting TCP scans and partial completions attacks such as TCP Flood Attacks. The naive technique is to maintain state for each connection which is infeasible in the network. Instead, we show that using a new data structure that we call Partial Completion Filters (PCFs), that the state needed is small enough to be efficiently implemented in vanilla ASICs and FPGAs.

We believe that PCFs are of independent interest because they provide a solution to the general problem of detecting imbalanced parentheses in a streaming environment. The analysis strategy, while simple, appears novel at least in the context of streaming algorithms. We also suggest various deployment options, including bidirectional and reverse path options which appear to be the most resilient to spoofing.

More fundamental than the specific techniques discussed in this paper is the general question of scalable behavior-based detection of attacks within the network. We believe this question is interesting because many other network functions (forwarding, classification, QoS) have already received considerable attention in the research and product literature, and solutions that scale to 40 Gbps already exist. As security functions become more prevalent in the edge first and then the core, it is natural to expect the same attention to be paid to scalable security solutions.

More than just introducing the question and suggesting a

specific mechanism for some problems, our paper shows that the issues of *behavioral aliasing* and *spoofing* are key questions that must be addressed in any scalable solution, even if the only response is to simply ignore the problem. For example, it may be reasonable to ignore spoofing until the bar is raised. These two provide a simple lens to view existing and future work in attack detection, and can perhaps suggest new solutions to an even broader class of attacks.

## Acknowledgments

## 7. REFERENCES

[1] Cert. advisory ca-1998-01 smurf ip denial-of-service attacks. http://www.cert.org/advisories/CA-1998-01.html.

[2] Cert. advisory ca-2001-19 "code red" worm exploiting buffer overflow in iis indexing service dll. http://www.cert.org/advisories/CA-2001-19.html.

[3] Cert advisory ca-2001-26 nimda worm. http://www.cert.org/advisories/CA-2001-26.html.

[4] Mazu networks. http://www.mazu.com.

[5] Mydoom.b virus. http://www.us-cert.gov/cas/alerts/SA04-028A.html.

[6] Sco inc. http://www.sco.com.

[7] ARBOR NETWORKS. http://www.arbornetworks.com.

[8] BARFORD, P., KLINE, J., PLONKA, D., AND RON, A. A signal analysis of network traffic anomalies. In *Proceedings of ACM SIGCOMM Internet Measurement Workshop* (Nov. 2002).

[9] BERNSTEIN, D. J. SYN cookies. http://cr.yp.to/syncookies.html, 1997.

[10] BLOOM, B. H. Space/time tradeoffs in hash coding with allowable errors. *Communications of the ACM 13*, 7 (July 1970), 422–426.

[11] CHECK POINT SOFTWARE TECHNOLOGIES LTD. Syndefender. http://www.checkpoint.com/products/firewall-1.

[12] DATAR, M., AND MUTHUKRISHNAN, S. Estimating rarity and similarity over data stream windows. Technical report, 2001-21,, DIMACS, Nov. 2001.

[13] ESTAN, C., AND VARGHESE, G. New directions in traffic measurement and accounting. In *ACM SIGCOMM* (Aug. 2002).

[14] ESTAN, C., AND VARGHESE, G. Autofocus : A tool for automatic traffic analysis. In *Proceedings of ACM SIGCOMM* (2003).

[15] FORESCOUT TECHNOLOGIES. http://www.forescout.com.

[16] FYODOR. http://www.insecure.org/nmap.

[17] GILBERT, A., GUHA, S., INDYK, P., MUTHUKRISHNAN, S., AND STRAUSS, M. Quicksand : Quick summary and analysis of network data. Technical report, 2001-43, DIMACS, Nov. 2001.

[18] GILL, T. M., AND POLETTO, M. MULTOPS: a data-structure for bandwidth attack detection. In *USENIX Security Symposium* (2001).

[19] GREENE, B. R., AND MCPHERSON, D. Sink holes : A swiss army knife isp security tool. http://www.nanog.org/mtg-0306/pdf/sink.pdf.

[20] HEBERLEIN, L. T., DIAS, G. V., LEVITT, K. N., MUKHERJEE, B., J.WOOD, AND D.WOLBER. A network security monitor. In *Proc. IEEE Symposium on Research in Security and Privacy* (1990), pp. 296–304.

[21] HUSSAIN, A., HEIDEMANN, J., AND PAPADOPOULOS, C. A framework for classifying denial of service attacks. In *ACM SIGCOMM* (Aug. 2003).

[22] JIN, C., WANG, H., AND SHIN, K. G. Hop-count filtering: An effective defense against spoofed ddos traffic. In *ACM Conference on Computer and Communications Security (CCS)* (Oct. 2003).

[23] JUNG, J., PAXSON, V., BERGER, A., AND BALAKRISHNAN, H. Fast portscan detection using sequential hypothesis testing. In *Proceedings of IEEE Symposium on Security and Privacy* (2004).

[24] KEYES, R. The Naptha DoS Vulnerabilities. http://razor.bindview.com/publish/advisories/adv_NAPTHA.html.

[25] KRISHNAMURTHY, B., SEN, S., ZHANG, Y., AND CHEN, Y. Sketch-based change detection: methods, evaluation, and applications. In *Proceedings of the conference on Internet measurement conference* (2003), ACM Press, pp. 234–247.

[26] LARSEN, R. J., AND MARX, M. L. *An Introduction to Mathematical Statistics and Its Applications*. Prentice Hall, Upper Saddle River, NJ 07458, 2001.

[27] LECKIE, C., AND KOTAGIRI, R. A probabilistic approach to detecting network scans. In *Proceedings of the Eight IEEE Network Operations and Management Symposium* (Apr. 2002).

[28] LEMON, J. Resisting syn flooding dos attacks with a syn cache. In *Proceedings of USENIX BSDCon'2002* (Feb. 2002).

[29] LEVCHENKO, K., PATURI, R., AND VARGHESE, G. On the difficulty of scalably detecting network attacks. In *Proceedings of the ACM Conference on Computer and Communications Security, Washington, D.C.* (Oct. 2004).

[30] MARTIN ROESCH. Snort. http://www.snort.org.

[31] MOORE, D., AND SHANNON, C. Sco offline from dos attack. http://www.sco.com.

[32] MOORE, D., VOELKER, G., AND SAVAGE, S. Inferring internet denial of service activity. In *USENIX Security Symposium* (2001).

[33] NETFLOW, C. http://www.cisco.com/warp/public/732/Tech/netflow.

[34] NETSCREEN 100 FIREWALL APPLIANCE. http://www.netscreen.com.

[35] NETSCREEN TECHNOLOGIES. http://www.netscreen.com.

[36] PAXSON, V. End-to-end routing behavior in the internet. *IEEE/ACM Transactions on Networking 5*, 5 (Oct. 1997), 601–615.

[37] PAXSON, V. Bro: A system for detecting network intruders in real-time. In *Computer Networks, 31(23-24), pp. 2435-2463* (Dec. 1999).

[38] PAXSON, V. An analysis of using reflectors for distributed denial-of-service attacks. In *Computer Communication Review 31(3)* (July 2001).

[39] PESCATORE, J., EASLEY, M., AND STIENNON, R. Network security platforms will transform security markets. http://www.techrepublic.com/article.jhtml?id=r00220021223jdt01.htm&src=bc, Dec. 2002.

[40] ROBERTSON, S., SIEGEL, E., MILLER, M., AND STOLFO, S. Surveillance detection in high bandwidth environments. In *Proceedings of the 2003 DARPA DISCEX III Conference* (Apr. 2003), pp. 229–238.

[41] SCHUBA, C., KRSUL, I., KUHN, M., SPAFFORD, E., SUNDARAM, A., AND ZAMBONI, D. Analysis of a denial of service attack on tcp. In *Proceedings of IEEE Symposium on Security and Privacy* (May 1997).

[42] STANIFORD, S., HOAGLAND, J. A., AND MCALERNEY, J. M. Practical automated detection of stealthy portscans. In *In Proceedings of the 7th ACM Conference on Computer and Communications Security* (2000).

[43] STANIFORD, S., PAXSON, V., AND WEAVER, N. How to 0wn the internet in your spare time. In *Proceedings of the 11th USENIX Security Symposium* (Aug. 2002).

[44] STANIFORD, S. J. Containment of scanning worms in enterprise networks. In *Journal of Computer Security* (Nov. 2003).

[45] WANG, H., ZHANG, D., AND SHIN, K. Detecting syn flooding attacks. In *IEEE INFOCOM* (2002).

[46] WANG, H., ZHANG, D., AND SHIN, K. Syn-dog: Sniffing syn flooding sources. In *IEEE ICDCS* (Vienna, Austria, 2002).

[47] WEAVER, N., PAXSON, V., STANIFORD, S., AND CUNNINGHAM, R. A taxonomy of computer worms. In *Proceedings of the ACM Workshop of Rapid Malcode (WORM)* (2003).

[48] YAAR, A., PERRIG, A., AND SONG, D. Pi: A path identification mechansim to defend against ddos attacks. In *Proceedings of the IEEE Symposium on Security and Privacy* (2003).

[49] YAAR, A., PERRIG, A., AND SONG, D. Siff: A stateless internet flow filter to mitigate ddos flooding attacks. In *Proceedings of the IEEE Symposium on Security and Privacy* (2004).

[50] ZHANG, Y., DUFFIELD, N., PAXSON, V., AND SHENKER, S. On the constancy of internet path properties. In *Proc. ACM SIGCOMM Internet Measurement Workshop* (Nov. 2001).