

Collaborative Detection of DDoS Attacks over Multiple Network Domains

Yu Chen, *Member, IEEE*, Kai Hwang, *Fellow, IEEE Computer Society*, and Wei-Shinn Ku, *Member, IEEE Computer Society*

Abstract—This paper presents a new distributed approach to detecting DDoS (*distributed denial of services*) flooding attacks at the traffic-flow level. The new defense system is suitable for efficient implementation over the core networks operated by *Internet service providers* (ISPs). At the early stage of a DDoS attack, some traffic fluctuations are detectable at Internet routers or at the gateways of edge networks. We develop a *distributed change-point detection* (DCD) architecture using *change aggregation trees* (CAT). The idea is to detect abrupt traffic changes across multiple network domains at the earliest time. Early detection of DDoS attacks minimizes the flooding damages to the victim systems serviced by the provider. The system is built over attack-transit routers, which work together cooperatively. Each ISP domain has a CAT server to aggregate the flooding alerts reported by the routers. CAT domain servers collaborate among themselves to make the final decision. To resolve policy conflicts at different ISP domains, a new *secure infrastructure protocol* (SIP) is developed to establish mutual trust or consensus. We simulated the DCD system up to 16 network domains on the Cyber Defense Technology Experimental Research (DETER) testbed, a 220-node PC cluster for Internet emulation experiments at the University of Southern California (USC) Information Science Institute. Experimental results show that four network domains are sufficient to yield a 98 percent detection accuracy with only 1 percent false-positive alarms. Based on a 2006 Internet report on *autonomous system* (AS) domain distribution, we prove that this DDoS defense system can scale well to cover 84 AS domains. This security coverage is wide enough to safeguard most ISP core networks from real-life DDoS flooding attacks.

Index Terms—Cyber defense, network security, DDoS attacks, and Internet technology.

1 INTRODUCTION

TODAY'S defense systems against *distributed denial-of-service* (DDoS) attacks are mostly built on detecting flooding consequences rather than the causes of the traffic surges [17], [19], [31], [39]. Flooding consequence is manifested by congestions on communication links [28], an overflow in a half-open SYN queue, or an imbalance between the I/O traffic on the gateways of edge networks [42]. Unfortunately, the damage has been done when the flooding consequence is observed. Thus, it is highly desirable to detect DDoS attacks at the earliest possible time, instead of waiting for the flood to become widespread [7], [8].

A comprehensive solution to DDoS attacks requires covering the global effects over a wide area of *autonomous system* (AS) domains on the Internet [3], [39]. Obviously, the global-scale defense is too costly for a real-life implementation. Even the Cyber Defense Technology Experimental Research (DETER) testbed [4] can only emulate partial

Internet activities. To implement an efficient defense system, we must leverage the network topology and use distributed traffic monitoring and detection. In reality, we build a DDoS defense system over a limited number of network domains serviced by the same *Internet service provider* (ISP). These ISP network domains cover the edge networks where the protected systems are physically connected.

In the sequel, we consider each AS a single network domain such as the core network of an ISP. According to an ISO 2006 Report [21] on AS resource allocations, there are 34,998 AS domains globally. Dimitropoulos et al. [12] identified that 67.7 percent of the AS domains belong to companies, organizations, or universities that run their own local area networks, 30.3 percent are ISP-controlled domains, the remaining 2 percent are Internet exchange points or network information centers.

Our DDoS defense is targeted for implementation in ISP core network domains [2]. The majority of ISPs do not share their AS domains with competitors. Therefore, they are unlikely to take part in collaborative DDoS defense. However, AS domains serviced by the same ISP or owned by the same company or organization can combat DDoS attacks collectively. This covers 98 percent of the Internet AS domains.

At the early stage of a DDoS attack, the traffic changes are difficult to detect because low traffic fluctuations are not observable. Monitoring Internet traffic at the individual flow level is cost prohibitive to cover all possible flows. Meanwhile, the global traffic in a wide area network is tremendously large to perform real-time detection of network anomalies effectively.

In practice, it is possible to convince a small percentage, say, 25 percent, of the ISP-controlled network domains to

• Y. Chen is with the Department of Electrical and Computer Engineering, State University of New York-Binghamton, Binghamton, NY 13902. E-mail: ychen@binghamton.edu.

• K. Hwang is with the Departments of Electrical Engineering and Computer Science, University of Southern California, 3740 McClintock, EEB 212, Los Angeles, CA 90089. E-mail: kaihwang@usc.edu.

• W.-S. Ku is with the Department of Computer Science and Software Engineering, Auburn University, Auburn, AL 36849. E-mail: weishinn@auburn.edu.

Manuscript received 13 Aug. 2006; revised 23 Dec. 2006; accepted 10 Apr. 2007; published online 9 May 2007.

Recommended for acceptance by S. Olariu.

For information on obtaining reprints of this article, please send e-mail to: tpds@computer.org, and reference IEEECS Log Number TPDS-0228-0806. Digital Object Identifier no. 10.1109/TPDS.2007.1111.

TABLE 1
Notations and Abbreviations Used in the Paper

Symbol	Definition	Abbr.	Meaning
α	Traffic inertia factor	ATR	Attack transit router
β	Router detection threshold	CAT	Change aggregation tree
θ	CAT detection threshold	SIP	Secure infrastructure protocol
R_d	DDoS detection rate	ISP	Internet service provider
R_{fp}	False-positive alarm rate	DCD	Distri. change-point detection
DR	Traffic deviation ratio in the I/O ports of a router	DETER	Cyber Defense Technology Experimental Research

join in collective effort to combat DDoS attacks. This amounts a few hundreds of domains to form a consortium in collective DDoS defense. We will prove in Section 5.1 that it would be sufficient to involve only tens of domains to work together in coping with most DDoS flooding attacks. This defense range is certainly within the coverage of a single ISP or of a few ISPs that collaborate with each other.

To be cost-effective, we propose to monitor the traffic at a *superflow* level. A superflow contains all packets destined for the same network domain from all possible source Internet Protocol (IP) addresses and applies various protocols such as Transmission Control Protocol (TCP) or User Datagram Protocol (UDP), etc. This detection level covers the aggregate from individual traffic flows. All packets of a superflow have the same prefix IP address of the same destination domain [15]. Motivated by using lightweight detection with low complexity [3], [9], [26], [42], we propose a *distributed change-point detection* (DCD) architecture using a new mechanism, called *change aggregation tree* (CAT). The concept of CAT was first presented in [11].

This CAT mechanism is designed at the router level for detecting abrupt changes in traffic flows. When a DDoS attack is launched, the routers observe changes in the spatiotemporal distribution of traffic volumes. The domain server uses the router-reported traffic surge reports to construct the CAT. Usually, these changes in traffic flows present a directionality homing toward the victim system. Random fluctuations incurred with legitimate traffic flows do not present the homing effects. For the benefit of our readers, Table 1 summarizes the basic parameters and abbreviations used in this paper.

Our DCD approach is unique and offers the very first attempt to explore DCD over collaborative network domains. We detect the start of DDoS flooding attacks by monitoring abnormal traffic flows. This monitoring and detection is performed from router to router as the CAT is dynamically constructed on the fly. On the DETER testbed, we implemented the detection scheme from 4 to 16 AS domains. We carried out intensive experiments to evaluate the DCD scheme. The performance results demonstrate high detection accuracy and low false-positive alarms.

The rest of this paper is organized as follows: Section 2 briefly reviews related works. Section 3 presents the principle of the change-point detection method and the algorithms for raising attack alerts by individual routers. Section 4 explains the CAT construction within a single network domain. The interdomain change detection algorithm is presented in

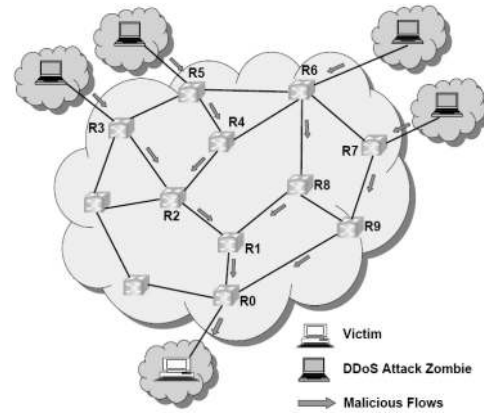


Fig. 1. A traffic superflow by DDoS flooding attacks launched from a large number of zombies toward a common victim host.

Section 5, along with a new *secure infrastructure protocol* (SIP). Section 6 reports the DETER experiments setups and performance results. Section 7 discusses scalability issues and deployment limitations. Finally, we conclude with a summary of contributions and discuss further research needed toward eventual use of the defense system in Section 8.

2 RELATED WORK AND OUR CONTRIBUTIONS

DDoS attacks often come with widespread worms [6]. The attacker often exploits the huge resource asymmetry between the Internet and the victim systems [39]. The flooding traffic is large enough to crash the victim machine by communication buffer overflow, disk exhaustion, connection link saturation, and so forth. Fig. 1 shows a flooding attack launched from four zombies. The *attack-transit routers* (ATRs) detect the abnormal surge of traffic at their I/O ports. The victim is attached with the end router R0 in Fig. 1. All the attack flows form the superflow homing toward the end router.

A plethora of DDoS defense and response mechanisms have been suggested in the past, including IP traceback [1], [3], [17], packet filtering [26], and flood pushback [20]. More sophisticated intrusion detection systems [19], [32] and DDoS defense schemes [10], [24], [31], [41] have been recently proposed. Researchers have attempted to combat repeated DDoS attacks [18]. Others use overlay networks [43], DDoS-resilient scheduling [36], and trust-negotiation [37] approaches to establish trust.

MULTIlevel Tree for Online Packet Statistics (MULTOPS) [16] and D-WARD [29] suggested filtering and rate limiting on suspicious flows at the source end. Security managers often focus on protecting their own networks and choose local detection approaches [7]. For instance, COSSACK [33] and DefCOM [29] deploy detectors at the victim side and send an alert to the filter or to the rate limiter located at the source side. Chen and Song [9] proposed a perimeter-based scheme for ISP to enable anti-DDoS services to their customers. Their scheme relies on edge routers to identify the sources of the flood of attack packets.

Many researchers use the change-point detection theory to detect abnormal Internet traffic caused by DDoS attacks [5], [11], [34], [42]. Lacking accurate statistics to describe the prechange and postchange traffic distributions, a nonparametric cumulative sum (CUSUM) scheme was developed

for its low computational complexity [5]. The scheme monitors the short-term behavior shifting from a long-term behavior. Once the cumulative difference reaches a certain threshold, an attack alert is raised. Wang et al. [42] suggested a centralized DDoS defense scheme to monitor the change points at the gateway level. Peng et al. [34] took a similar approach in monitoring the source IP addresses.

In this paper, we propose a new distributed aggregation scheme based on change-point detection across multiple network domains. This scheme is extended from the single-domain change-detection scheme reported in [11]. We establish cooperation among communicating network domains. This enables the building of an early warning system for DDoS defense across multiple ISP domains. Our DCD scheme is capable of tracing back automatically, once the detection is successfully carried out. The global CAT detects the network anomalies incurred on the fly.

In summary, our contributions are highlighted below in four technical aspects, the details and proofs of which are given in subsequent sections:

1. *Traffic anomaly detection at the superflow level.* Monitoring Internet traffic at routers on individual flows is identified by a 5-tuple: $\{source\ IP, destination\ IP, source\ port, destination\ port, protocol\ applied\}$. The superflow consists of those traffic flows destined for the same network domain and applied the same protocol. This level of traffic monitoring and anomaly detection is more cost-effective for DDoS defense in real-life Internet environments.
2. *Distributed change-point detection.* Considering the directionality and homing effects of a DDoS flooding attack, we propose to use collaborative routers for DCD and use the domain servers for alert correlation and aggregation.
3. *Hierarchical alerts and detection decision making.* Our system adopts a hierarchical architecture at the router and domain levels. This simplifies the alert correlation and global detection procedures and enables the DCD system implementation in ISP networks.
4. *Novelty of SIP.* We propose a new trust-negotiating SIP to secure interserver communications. The SIP has removed some of the shortcomings of the existing IP security (IPsec) and application-layer multicasting protocols [25], [43]. SIP appeals for implementation on virtual private network (VPN) tunnels or over an overlay network built on top of all domain servers.

3 DISTRIBUTED CHANGE-POINT DETECTION

The DCD scheme detects DDoS flooding attacks by monitoring the propagation patterns of abrupt traffic changes at distributed network points. Once a sufficiently large CAT is constructed to exceed a preset threshold, an attack is declared. This section presents the principles behind the DCD system. We focus on traffic pattern change detection at the router level.

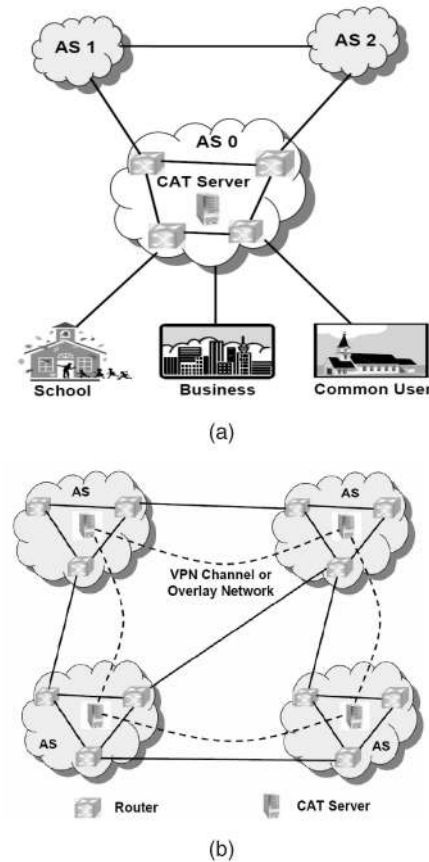


Fig. 2. Distributed change detection of DDoS attacks over multiple AS domains. (a) Multidomain DDoS defense system. (b) Interdomain communication via VPN tunnels or an overlay network atop the CAT servers in four domains.

3.1 The DCD System Architecture

Fig. 2 presents the system architecture of the DCD scheme. The system is deployed over multiple AS domains. There is a central CAT server in each domain. The system detects traffic changes, checks flow propagation patterns, aggregates suspicious alerts, and merges CAT subtrees from collaborative servers into a global CAT. The root of the global CAT is at the victim end. Each tree node corresponds to an ATR. Each tree edge corresponds to a link between the ATRs.

Our system has a hierarchical detection architecture. There are three layers in this architecture. At the lowest layer, an individual router functions as a sensor to monitor local traffic fluctuations. A change-point detection program (Algorithm 1) is executed on each router. Considering the directionality and homing effects in a DDoS flooding attack, routers check how the wavefront changes. A router raises an alert and reports an anomalous traffic pattern to the CAT server.

The second layer is at each network domain level. The CAT server constructs a CAT subtree according to the alerts collected. The subtree displays a spatiotemporal vision of the attack superflow in the domain. At the highest layer, the CAT servers at different domains form an overlay network. For security precaution, they communicate with each other through VPN channels.

All CAT servers send their locally generated CAT subtrees to the edge server in the destination domain




Fig. 3. Construction of a CAT for the flooding pattern reported by nine ATRs in Fig. 1, where the victim host is connected to router R0.

where the victim is attached. By merging CAT subtrees from cooperative domains, the destination server has a global picture of the attack. The larger is the global CAT so constructed, the higher is the threat experienced.

The CAT detection scheme does not need to specify an absolute threshold on traffic volume. The detection is done by checking the number of nodes (routers) raising the alerts from the CAT subtree. Fig. 3 illustrates how a CAT subtree rooted at the end router is constructed by merging the alert reports from nine ATRs. The upstream and downstream ATRs report to the CAT server during each monitoring cycle.

Using Algorithm 2, the server constructs a CAT rooted at the end router R0. The server recursively scans through all upstream routers to construct the tree. The CAT presents a traffic-flow tree pattern rooted at the router connected to the edge network where the victim is attached. With sufficient exchange of alert information from related domains, the system can detect the DDoS flooding attack at a very early launching stage, before the attack flows hit the victim network.

Flow-level detection can distinguish among several DDoS attacks. We monitor the traffic change based on the homing effects of the victim address. Each CAT is uniquely constructed for flooding streams toward the same destination in an edge network. When multiple DDoS attacks are launched concurrently against multiple victims, there are multiple CATs to be constructed, and they are completely distinguishable. The shape of the CAT corresponds to the attacking traffic paths.

Surely, the attacker can randomly choose zombies in an attack. In addition, the group of zombies can be changed dynamically during the attack. However, the random selection of zombies will not impact our detection results, because the CAT is constructed on the fly. Essentially, a different distribution of zombies results in a different CAT. Since the detection criterion is not the shape, but the size of the CAT, changing the zombie distribution will not weaken our detection capability.

3.2 Principles of Change-Point Detection

In change-detection problems, if prechange and postchange distributions are unknown, the CUSUM statistic has been suggested to solve the problem [5]. We adopt a nonparametric approach for simplicity. Let t_1, t_2, \dots, t_m be discrete time instants and $x(t_m, i)$ be the number of packets received by a router during time slot m at port i . The historical

estimate of the average number of packets is defined iteratively by

$$\bar{X}(t_m, i) = (1 - \alpha) \cdot \bar{X}(t_{m-1}, i) + \alpha \cdot x(t_m, i), \quad (1)$$

where $0 < \alpha < 1$ is an inertia factor showing the sensitivity of the long-term average behavior to the current traffic variation. A higher α implies more dependence on the current variation. We define below $S_{in}(t_m, i)$ as the deviation of input traffic from the average at time slot t_m :

$$S_{in}(t_m, i) = \max\{0, S_{in}(t_{m-1}, i) + x(t_m, i) - \bar{X}(t_m, i)\}. \quad (2)$$

The subscript i indicates that this is the statistics of the incoming traffic. While a DDoS flooding attack is launched, the cumulative deviation is noticeably higher than the random fluctuations. Since $S_{in}(t_m, i)$ is sensitive to the changes in the average of the monitored traffic [5], we measure the abnormal deviation from the historical average as follows: Let the *deviation from average* (DFA) be the indicator of such an attack. The incoming traffic DFA is defined below at port i at time t_m :

$$DFA_{in}(t_m, i) = S_{in}(t_m, i) / \bar{X}(t_m, i). \quad (3)$$

If the DFA exceeds a router threshold β , the measured traffic surge is considered a suspicious attack. The threshold β measures the magnitude of traffic surge over the average traffic value. This parameter is preset based on previous router use experience. In a monitoring window of 100 ms to 1 second, a normal superflow is rather smooth due to statistical multiplexing of all independent flows heading for the same destination [22]. If there is no DDoS attack, we expect a small deviation rate far below β . In general, we work in the range $2 \leq \beta \leq 5$.

For outgoing traffic, we define $y(t_m, i)$ as the number of packets at time t_m leaving at port i and as the historical average of departed packets. Similarly, we have

$$\bar{Y}(t_m, i) = (1 - \alpha) \cdot \bar{Y}(t_{m-1}, i) + \alpha \cdot y(t_m, i), \quad (4)$$

$$S_{out}(t_m, i) = \max\{0, S_{out}(t_{m-1}, i) + y(t_m, i) - \bar{Y}(t_m, i)\}. \quad (5)$$

The above equations will be used to specify the change-detection algorithms in subsequent sections.

3.3 Traffic Surge Detection in Routers

Each router monitors traffic variation and counts the packet number within a monitoring window at each I/O port. We use the term *traffic pattern* to refer to the combination of traffic surges at all I/O ports of a router. In general, a router with m -input ports and n -output ports may have 2^{m+n} possible traffic patterns. The height of the black boxes in Fig. 4 signifies the magnitude of traffic volume at I/O links. The raised block height indicates the surge detected, and the lower boxes, the normal traffic.

All packets of a superflow must be homing toward the same destination network. Before entering the destination domain, the flow paths present a converging tree pattern. Only at the destination does the superflow scatter packets toward a particular edge network. There exist 16 possible traffic patterns from a 2×2 router. For simplicity, we illustrate in Fig. 4 only four basic traffic patterns at a 2×2 router with

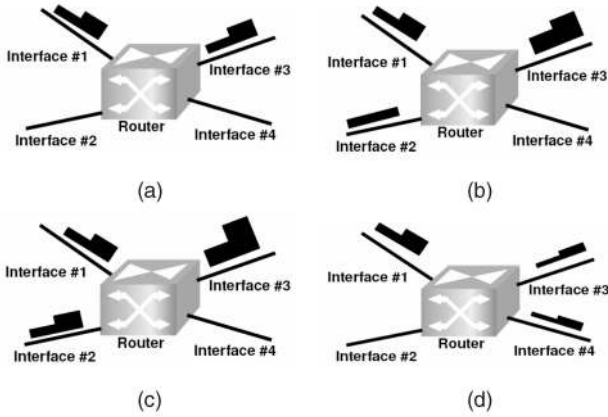


Fig. 4. Four basic patterns of traffic changes at 2×2 router I/O ports. (a) Flow through. (b) Partial aggregation. (c) Full aggregation. (d) Scatter pattern.

$m = n = 2$; the remaining 12 traffic patterns can be specified similarly:

1. *Flow-through pattern.* This traffic pattern is shown in Fig. 4a. The router forwards the traffic flow from an input port to a selected output port without subdividing or diverting the traffic to other ports.
2. *Partial aggregation pattern.* All the incoming flows are merged at one outgoing port i_{out} , not all incoming flows contain traffic surges as shown in Fig. 4b.
3. *Full aggregation pattern.* The outgoing flow merges multiple incoming flows, all containing traffic surges exceeding the threshold β . This router is considered a merge point on the attacking path (Fig. 4c).
4. *Scatter pattern.* The incoming flow scatters at this router. This is not part of a DDoS attack (Fig. 4d). This pattern is observed in the destination domain.

Another statistical parameter, *deviation ratio* (DR), is defined below to measure the ratio of incoming packets port i_{in} have propagated to output port i_{out} . DR is the ratio of traffic deviations between I/O ports:

$$DR(i_{in}, i_{out}) = S_{out}(t_m, i_{out}) / S_{in}(t_m, i_{in}). \quad (6)$$

If $DR > 1$, the router amplifies the input deviation. This corresponds to a full surge of traffic volume. $DR \approx 1$ implies that the router merely plays the role of a forwarder. This phenomenon is observed in the partial surge at one input port. The case of $DR < 1$ indicates that the incoming wave is scattered to multiple ports. It is not part of the convergence traffic of DDoS attacks. Therefore, by checking the DR value, a router determines whether the pattern is part of the traffic from a DDoS attack.

When a router detects that a DFA_{in} exceeds the deviation threshold β , it calculates the deviation rate between the outgoing and incoming ports. If DR is close to one, the traffic aggregation pattern is considered suspicious. The router generates an alert message and reports the pattern to the CAT server. Otherwise, the router sends a regular status message indicating no anomaly observed.

Presented below is a pseudocode of Algorithm 1 for local change detection at the router level.

Algorithm 1: Traffic surge detection at router level

Input: $x(t, i)$ and $y(t, i)$: Incoming and outgoing packets at time t and port i , respectively

$\bar{X}(t_{m-1}, i)$: Historical average of packet arrivals up to time $m - 1$ at port i

$\bar{Y}(t_{m-1}, i)$: Historical average of outgoing packets up to time $m - 1$ at port i

Router detection threshold β based on past experience

Output: Alert messages sent to the central CAT server.

Procedure:

01: Update historical average of I/O packets in a flow

02: Calculate DFA_{in} using (3)

03: **If** $DFA_{in} \geq \beta$ **Then** Calculate DR using (7)

04: **If** $DR \approx 1.0$ **Then** Suspicious pattern detected.
Send out an alert message to CAT server.

05: **Else** Nothing suspicious.

Send out a regular status message to CAT server.

Algorithm 1 demands lightweight computing power at ATRs. For an $m \times n$ ATR at an intermediate node of the CAT subtree, there are 2^{m+n} combinations of traffic surge patterns at the I/O ports. For each superflow, the router needs to calculate the DR at most m times. Even if a large DFA value is detected at an input port, the router uses only one output port to release the traffic surge. This will simplify the routing decision at the ATR significantly.

At the victim domain, the end-router decision is more complicated. If two or more attacks are launched concurrently toward the same destination domain, the traffic surges will pass through the end router at multiple ports. For an $m \times n$ end router, the worst case is that m input surges scatter to n output ports. The end router calculates the DR mn times. This burden is lowered by splitting the superflow to multiple destination addresses.

4 CONSTRUCTING SUBTREES AT DOMAIN SERVERS

This section describes the CAT subtree construction at each CAT server in a single network domain. Different subtrees are generated in multiple network domains. The global CAT is generated by merging all subtrees. While the flooding traffic merges at the victim end, the routers along the paths capture suspicious traffic patterns.

The router reports the identifier of a superflow causing the traffic surge. Since all routers are under the same ISP authority and work cooperatively, each router knows their immediate neighbors. Using the reported status information, the domain server detects the traffic flood based on the CAT constructed.

The alert message provides the upstream and downstream router identifiers. The alert message provides information for the CAT server to include the routers in the CAT subtree. The main purpose of sending the flow status message is to report where the suspicious flows are captured.

To indicate the location of a suspicious flow, the router identifier must send. We need to identify the superflow identifier of the n -bit prefix of the destination IP addresses. To construct the CAT, the status report provides the upstream and downstream router identifiers instead of router I/O port numbers. Using the reported status

TABLE 2
Alert Message Reported by a Router

Parameter	Brief Definition
nd_id	The router ID,
fl_id	The superflow ID
up_num	Number of upstream nodes
dn_num	Number of downstream nodes
up_id	node ID of upstream node
dn_id	node ID of downstream node
Router status	Suspicious attack or normal traffic

information, the domain server constructs the CAT gradually after receiving the alert reports from the ATRs.

Table 2 summarizes the information carried in a typical alert message from an ATR. The output of Algorithm 2 is a single-domain CAT subtree similar to the one shown in Fig. 3. The CAT is specified by a hierarchical data structure. The root node carries the superflow ID, the number of routers involved, root node ID, and the count of child nodes at the next level.

Algorithm 2: CAT Subtree Construction in a Single Domain Server

Input: Traffic alert messages received from all routers in the same AS domain

Output: A data structure describing the CAT subtree constructed in this domain

Procedure:

- 01: Read all suspicious patterns in and arrange them according to router ID
- 02: Start from the suspicious node with minimum ID R_{min}
- 03: $root \leftarrow R_{min}$
- 04: read the upstream node number up_num
- 05: read the downstream node number dn_num
- 06: $node_number \leftarrow node_number + up_num - 1$
- 07: **While** $up_num > 0$
- 08: Read in one upstream node R_{up}
- 09: Add R_{up} as a leaf node
- 10: scan through its upstream nodes
- 11: $up_num \leftarrow up_num - 1$
- 12: **End While**
- 13: **While** $dn_num = 1$
- 14: Read the downstream node R_{dn} ;
- 15: $root \leftarrow R_{dn}$
- 16: $node_number \leftarrow node_number + 1$
- 17: Scan through other upstream nodes of new root;
- 18: $dn_num \leftarrow dn_num$ of the new root
- 19: **End While**

To clarify the control flow, this construction process is specified by a flowchart in Fig. 5. The next level lists the pair of information {L1 node ID, count of children at next level L2}. This process continues until it reaches the leaf nodes of the tree. The CAT subtree is sent to the CAT server of the destination domain. In Algorithm 2, the domain server constructs the CAT subtree based on collected status reports from the routers. Routers that detected no attacks are not involved in the tree construction.

Starting from the node R_{min} with a minimum ID in Fig. 5, the CAT server takes it as the root node. The server scans

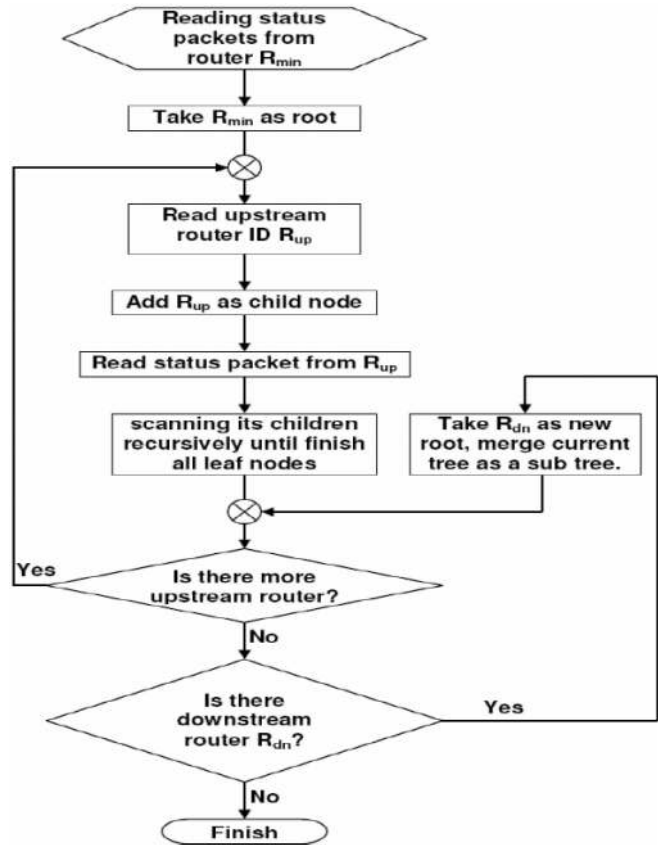


Fig. 5. Control flow in Algorithm 2 to construct a CAT subtree.

through upstream child nodes identified by up_id . This descendent search is performed iteratively until the leaf nodes are reached. If there is a downstream router R_{dn} , we take router R_{dn} as the new root and repeat the procedure. Meanwhile, the descendent search procedure is repeated for all upstream routers of root R_{dn} . Then, we check the downstream router of R_{dn} and repeat the procedure until the downstream router is out of the domain boundary.

5 MERGING TRAFFIC IN MULTIPLE DOMAINS

This section describes the extension of the single-domain detection scheme to work on multiple network domains. First, we analyze the complexity of the global CAT growth in real-life Internet domains. Then, we present the mechanisms for cross-domain attack detection. In addition, we introduce a new protocol that supports interdomain communications, trust negotiation, and collaborative detection.

5.1 CAT Growth with Domain Distribution

The complexity of the CAT growth is analyzed below based on Internet topology data available from open literature [21], [38]. Fig. 6 illustrates the process of the CAT growth out of merging subtrees from attack-transit domains. Let r be the number of hops from an AS domain to the destination domain. The server checks the received subtrees in increasing order of distance r .

The system first merges the subtrees from ASs located in 1-hop ($r = 1$) distance to form a partial global tree. Next, it merges the subtrees from domains at 2-hop distance. The merging process repeats with distances $r = 3, 4$ until all

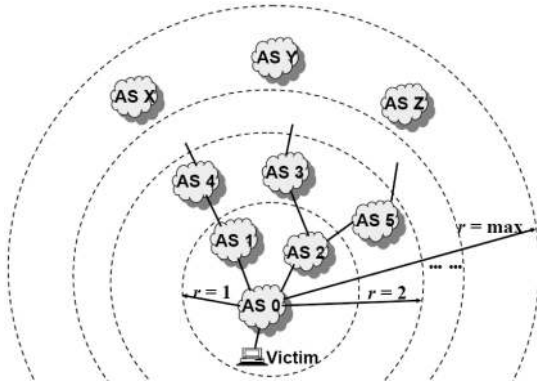


Fig. 6. Merging CAT subtrees from neighboring AS domains to outer domains to build a global CAT, where AS_0 is the victim domain, and $r_{max} = 6$ hops.

subtrees are merged into the final global CAT. We analyze below the complexity of global CAT growth at intradomain and interdomain levels. The routers monitor traffic conditions and report anomalies to their domain CAT server periodically. The local setting parameters α and β affect the size of the local CAT subtrees constructed.

Given a domain consisting of N routers, the number of alerts that the CAT server receives is proportional to N . The threshold used in CAT subtree construction (Algorithm 2) is equal to the number of alerts received by the final CAT server. Therefore, the detection time is estimated by $O(N)$ within each domain. Of course, different domain sizes (N) may require a variable subtree generation time.

At the interdomain level, the complexity of global CAT merging is highly dependent on the network topology. We model the Internet domains as an undirected graph of M nodes and E edges. The diameter of the graph is denoted by δ .

Siganos et al. [38] model the Internet neighborhood as an H -dimensional sphere with a diameter δ . The parameter H is the dimension of the network topology [14]. For example, $H = 1$ specifies a ring topology, and $H = 2$, a 2-dimensional mesh. Any two nodes are within an effective diameter, δ_{ef} hops away from each other. Faloutsos et al. estimated the magnitude of δ_{ef} by the following expression:

$$\delta_{ef} = \left(\frac{M^2}{M + 2E} \right)^{1/H}. \quad (7)$$

In 2002, the dimension of the Internet was calculated as $H = 5.7$ in an average sense. The ceiling of this diameter δ_{ef} is thus set to be 6. Let $NN(h)$ be the number of domains located at distance h from a typical domain in the Internet. Table 3 gives the domain distribution—the probability of an AS domain residing exactly h hops away from a reference domain. The numbers of domains in various distance ranges are given in the second row. It is interesting to note

that most communicating domains are within 3 or 4 hops, almost following a normal distribution centered on an average hop count of 3.5.

The number of Internet AS domains keeps increasing in time, the Faloutsos et al. reports [14], [38] indicates that this AS distribution is pretty stable over time. This implies that a packet can reach almost all domains in the Internet by traversing through 6 hops. Therefore, we set the maximum hop count $r_{max} = 6$ in Fig. 6.

Let p_h be the probability of having an AS domain located at distance h from the reference domain. Therefore, the average number of domains used to build a global CAT is upper bounded by

$$T = \sum_{h=1}^{r_{max}} NN(h) \times p_h. \quad (8)$$

Substituting the entries in Table 2 into (9), we obtain the expected domain count

$$\begin{aligned} T &= 14 \times 0.004 + 2818 \times 0.0805 + 13493 \times 0.3855 + 13342 \\ &\quad \times 0.3812 + 4445 \times 0.127 + 788 \times 0.0225 + 102 \times 0.0029 \\ &= 11,097 \end{aligned}$$

domains used in average Internet applications. This domain count posts a loose upper bound on the expected number of ISP domains involved in building a global CAT.

In reality, only a few ISP-controlled AS domains may commit to defend DDoS attacks collaboratively. On the conservative side, consider that 5 percent of ISP AS domains are committed. Thus, the above upper bound could be reduced to only 168 ISP domains, provided that they conform to the domain distribution in Table 3.

5.2 Global Tree Construction at the Victim End

In a DDoS flooding attack, the attacker often recruits many zombies distributed over the Internet. The flooding traffic may travel through multiple AS domains before reaching the edge network where the victim is physically attached. Routers at the upstream domains observe the suspicious traffic flows earlier than routers at the downstream networks.

Our DCD system was designed to have strong collaborations among all domain servers along the superflow paths. Algorithm 3 specifies the merge of CAT subtrees for detecting DDoS attacks across multiple network domains. The CAT subtrees constructed at all traversed domains must be merged to yield a global CAT at the destination domain.

Algorithm 3: Global CAT Construction and Detection Decision

Input: CAT subtree reports from participating domain servers, the server detection threshold θ .

TABLE 3
Internet Domain Distribution Reported on 28 February 2006 [21]

Hop Count, h	1	2	3	4	5	6	≥ 7
Domain Distribution, p_h	0.04%	8.05%	38.55%	38.12%	12.7%	2.25%	0.29%
Domain Count, $NN(h)$	14	2,818	13,493	13,342	4,445	788	102

Output: The global CAT over multiple AS domains. Raise the alert for an imminent DDoS attack.

Procedure:

- 01: Construct the local CAT subtree (Algorithm 2) periodically
- 02: Receiving subtrees from other CAT servers
- 03: **If** local subtree exists, **Then** Check the superflow ID,
- 04: **If** this domain is the destination domain, **Then** Set distance $r = 1$
- 05: Merge subtrees from domains at distance r to the current global tree
- 06: $r \leftarrow r + 1$
- 07: **While** (there are unchecked subtrees), generate the CAT profile
- 08: **If** CAT profile $\geq \theta$ **Then** DDoS attack is detected and raise an alert
- 09: **Else** Check the root router position
- 10: **If** root router is connected to other domain
- 11: **Then** Sent the global CAT to the destination domain server
- 12: **Else** Raise an attack alert based on the global tree merged

The final declaration of a DDoS attack is the result of threshold detection using the global CAT. Not only the victim network launches appropriate countermeasures, but also some traceback actions are to be taken by all ATRs along the superflow paths. The actions include dropping of suspicious packets or rate limiting against the flows.

The global CAT corresponds to the flooding attack flows. The leaf nodes are directly related to the zombies used. The height of the global CAT corresponds to the superflow hop count. Some highly distributed attacks may recruit hundreds of zombies; the global CAT may cover a wide area on the Internet. Therefore, we use the global CAT profile θ as a global detection threshold. The CAT profile indicates how many routers observed abnormal traffic surges. Thus, θ is an integer bounded by the number of ATRs in a domain.

The tree width and height thus reveal the scope of the DDoS attack. Through experiments on the DETER testbed, we obtain the global detection threshold value by training from some attack data sets. These threshold values have yielded the highest detection rate and lowest false-positive rate during the training period.

Upon receiving subtrees from upstream CAT servers, the CAT server in the destination domain builds the global CAT from its local subtree. Once the global CAT is formed, the server compares the CAT profile with the global detection threshold θ to decide on a DDoS attack. An alert is raised, and a necessary countermeasure is triggered accordingly. Fig. 7 shows an example network environment involving six AS domains. The victim system is located in the AS1 domain. Zombies are scattered widely in the Internet outside the illustrated domains. By detecting abnormal traffic changes in each domain, the CAT server creates a CAT subtree locally at each domain using Algorithm 2.

Fig. 7b shows the three steps taken to merge the six subtrees generated by the six CAT servers of the six AS domains. All six subtrees resulted from checking the

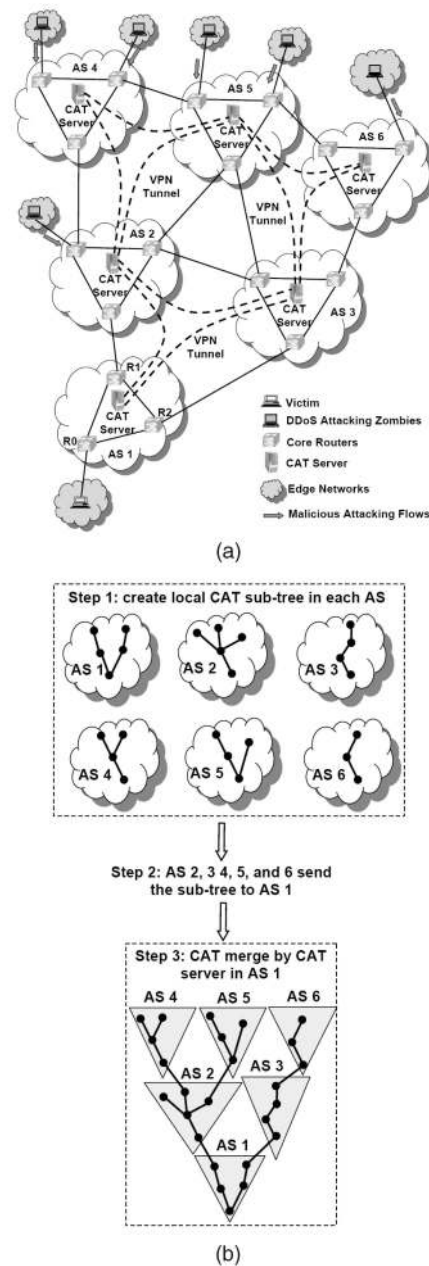


Fig. 7. An example six-domain global CAT construction environment. (a) DCD system architecture over six domains. (b) Merging six CAT subtrees to yield a global CAT.

packets belonging to the same superflow traffic destined for the same domain AS1. Five subtrees generated in AS2, AS3, AS4, AS5, and AS6 at upstream domains are sent to AS1 at Step 2. Then, the concatenated CAT subtrees are connected to the downstream subtree in AS1. Thus, the global CAT is finally rooted at the last hop router to an edge network R0 that is attached to the victim system.

5.3 Secure Infrastructure Protocol (SIP)

To support global CAT construction across multiple domains, we propose a new SIP by extending from the Internet Control Message Protocol (ICMP) [35]. The SIP is designed as an integral part of the IPv4 standard in a layer closer to the physical network. To provide a secure




Fig. 8. Three communication levels in using the SIP between two CAT servers in two AS domains.

information exchange platform, the protocol supports three levels of communication as illustrated in Fig. 8.

The lowest level enables routers in the same domain to share information in status monitoring. The second level supports communication between routers and the CAT server in each domain. Routers periodically report local traffic detection results to the domain server. The highest level supports interdomain communication among the CAT servers in trust negotiation to resolve conflicts in security policies in different domains.

Due to privacy and security concerns, ISPs are reluctant to reveal inside information to competitors, such as topology, bandwidth configuration, and capabilities. Hence, aside from managing the information exchange, servers are also in charge of trust negotiation. We adopted the trust negotiation of multilateral security suggested by Ryutov et al. [37].

Using their *Generic Authorization and Access-control* (GAA) interface, SIP can help establish trust among AS domains and enter a collaborative DDoS defense system agreement. By trust negotiating, ISPs can determine how much private information is allowed to be shared with others.

When a domain server enters a collaborative defense agreement with another domain server, the administrator sets up the trust levels required. Each server needs to perform a trust negotiation with a peer server only once with each new domain joining. The trust level determines how much sensitive information can be disclosed when exchanging detected anomalies information.

SIP is designed with three trust levels: *full trust* (FT), *basic trust* (BT), and *no trust* (NT). With FT, the SIP server will provide all necessary information that describes the characteristics of the detected anomaly. With BT, the SIP server only sends some statistics containing no private and sensitive data. When NT is set, no cooperation takes place between the CAT servers.

Existing security protocols like IPsec and ICMP focus on information security through strong cryptography to avoid eavesdropping, strict access control to block illegal access, the use of digital signatures, and so forth. ICMP is used to control error, when a network device or host requires reporting an error in datagram processing [35].

In contrast, SIP is designed to monitor the status of the network infrastructure such as link bandwidth utility, fluctuation of traffic, and so forth. Although both IPsec and SIP are implemented at the IP layer, SIP is totally deployed inside the intermediate network and does not need support from the end hosts.

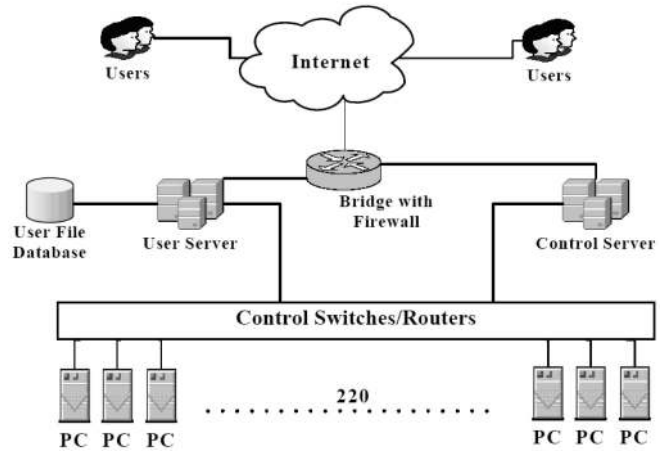


Fig. 9. Multi-PC cluster architecture of the DETER Testbed at USC/ISI.

6 RESULTS FROM DETER EXPERIMENTS

We verify the performance of the newly proposed DDoS detection scheme with network attack experiments on a 220-node DETER testbed [4] at the University of Southern California Information Sciences Institute (USC/ISI). The experimental settings and performance results are reported below.

6.1 Settings on DETER Testbed

To evaluate the performance of the CAT-based DDoS detection system, we performed experiments using variant network topology, background traffic, and attack generation. We adopt real-world ISP topologies downloaded from the Rocketfuel Project at the University of Washington [2]. The DETER testbed is essentially an Internet emulator built over a cluster of PCs to perform a broad range of network security experiments.

We report below the DETER results on 2, 4, 8, and 16 collaborative domains. Each domain typically has 7 to 12 routers. Fig. 9 shows the DETER cluster configuration built at USC/ISI. Due to limited 220 nodes in DETER, we choose the smallest ISP configuration topology from the Rocketfuel data set.

In the DETER testbed experiments, 34 routers were simulated over four AS domains. The link bandwidth among the network domains was set at 100 Mbytes/s. To generate the background traffic closer to reality, we use the OC48 trace data set from the Cooperative Association for Internet Data Analysis (CAIDA) project [30] to regenerate Internet traces using the Harpoon traffic generator [40]. To generate DDoS attacks, we use the toolkit Stacheldraht (version 4.0) [13]. Stacheldraht generates the ICMP, UDP, TCP SYN flooding, and Smurf attacks.

6.2 Performance Metrics Used

The performance of our DCD detection scheme is evaluated with three metrics: *detection rate*, *false-positive alarms*, and *system overhead*. All the metrics are measured under different DDoS attacks using TCP, UDP, and ICMP. The *detection rate* R_d of DDoS attacks is defined by the following ratio:

$$R_d = a/n, \quad (9)$$

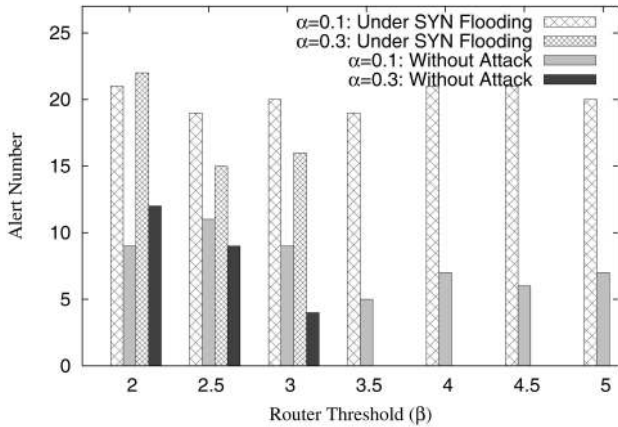


Fig. 10. Total alerts plotted against the router threshold in DETER experiments with 34 routers in four AS domains using a 500 ms monitoring window.

where a is the number of DDoS attacks detected in the simulation experiments, and n is the total number of attacks generated by the Stacheldraht toolkit during the experiments.

In addition, we are interested in the performance of our DCD scheme under normal traffic without DDoS attacks. A traffic superflow is called a *false-positive alarm* if an attack is detected out of the normal traffic without attacks. Let p be the number of false-positive alarms raised by the CAT server and m be the total number of normal traffic flow events checked by the simulator. Therefore, the ratio p/m defines the *false-positive alarm rate*:

$$R_{fp} = p/m. \quad (10)$$

The *receiver operating characteristic* (ROC) curve shows the trade-off between the detection rate and the false-positive rate. Section 6.4 reports the detection accuracy measured under different detection thresholds. Another critical issue is the time overhead to detect the launch of DDoS attacks. The average detection time measures from the start of a DDoS attack to the time of raising an alarm. The monitoring window should be chosen to be greater than this detection time.

6.3 DETER Experimental Results

To evaluate the effectiveness of the DCD detection scheme, we report the alerts raised at the routers and analyze the CAT subtree properties. The DETER experiments choose the inertia factor α in the range (0.1, 0.5) and the router detection threshold β in the range (2, 5).

6.3.1 Alert Magnitude and Router Threshold

We used 34 routers in four AS domains in the DETER testbed. When the traffic flow exceeds the router threshold β , the router raises alerts. No alert will be raised if no anomaly is detected at the routers. Fig. 10 illustrates the total numbers of alerts raised by the routers with SYN flooding compared with the case of no attack. More alerts are raised with attacks compared with alerts from no attack.

The two left bars of each group correspond to $\alpha = 0.1$, the case of heavy dependence on the past average traffic to raise an alert. The leftmost bar stays at around 20 alerts,




Fig. 11. Variations of the global CAT size over four AS domains in a 500 ms monitoring window.

which is insensitive to increasing threshold β . The second (gray) bar reduces to five alerts as β increases. This implies that $\beta = 3.5$ is an ideal threshold to distinguish attacks from no attack. The two right bars (fabric and black) demonstrate a higher inertia value setting $\alpha = 0.3$. The fabric bars with attacks are much higher than the black bars without attacks. For $\alpha \geq 0.3$ and $\beta \geq 3.5$, the alert bars reduce to zero, meaning that the attack is not detectable.

6.3.2 The Global CAT Profiles

Fig. 11 displays the global CAT profile, defined by the tree size or the number of routers that have raised the alert signals. We plot the tree profile against the router detection threshold β . For a small traffic inertia $\alpha = 0.1$, the SYN flood has an average of 20 routers being alerted to yield the CAT subtree. Without attacks, this tree profile reduces to less than five nodes.

With a higher traffic inertia $\alpha = 0.3$, the SYN attack results in a tree profile with around 10 alerts. Without attacks, the alert profile reduces to zero after the threshold β exceeds 3. Based on these results, we discover an optimal router threshold setting $\beta \geq 3.5$ with an inertia ratio $\alpha = 0.1$. When 20 out of 34 routers launched alerts, the router utilization is $20/34 = 58$ percent.

6.3.3 Effects of the Monitoring Window Size

The size of the monitoring window affects the number of alerts raised in multiple AS domains. Through experiments on the DETER testbed, we observed the optimal monitoring window size to be 100 ms. The false-positive alarm number increases steadily with the increase in the monitoring window size. However, the number of alerts with real SYN attacks remains the same for all monitoring window sizes, as shown in Fig. 12.

In Stacheldraht [13], the UDP and ICMP packet rate for each zombie is adjustable through setting different UDP and ICMP packet sizes—the longer the packet length, the lower the packet rate. The TCP SYN attacks use a fixed packet size of 64 bytes with a fixed packet rate. The maximum UDP and ICMP packet size is limited to 1,024 bytes in Stacheldraht. We observed a similar detection rate for TCP SYN and UDP/ICMP attacks with 128-byte packets.

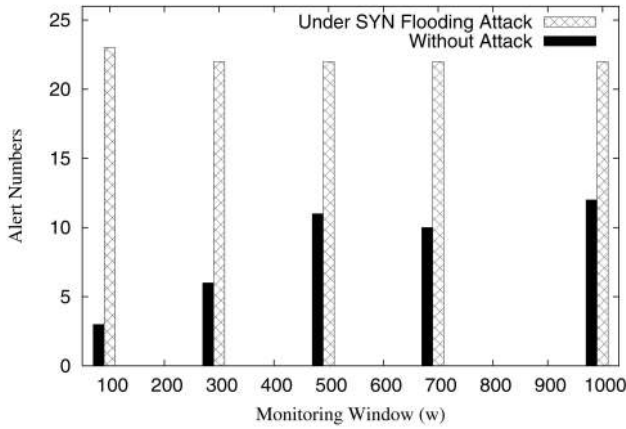


Fig. 12. The router alert number with and without DDoS attacks monitored using various window sizes.

6.4 Detection Accuracy and False Alarms

In this section, we report the detection accuracy of the DCD scheme under TCP SYN, UDP, and ICMP attacks with different packet rates. The reported results correspond to $\alpha = 0.1$, $\beta = 2.0$, and $w = 500$ ms. The detection accuracy is reflected by two factors—the detection rate and the false-positive rate. In order to achieve high detection accuracy, we have to increase the detection rate and decrease the false-positive rate. However, a trade-off exists between the two factors as discussed in the following paragraphs.

6.4.1 Detection Rate of DDoS Attacks

Fig. 13 illustrates the variances of the detection rate (7) with respect to different server detection thresholds (θ). The TCP SYN attack has the highest detection rate, which is close to 100 percent with $\theta \leq 12$. The low-rate UDP attacks have a lower detection rate than TCP attacks.

For UDP attacks of 512-byte packets, the detection rate can be kept above 80 percent with $\theta \leq 9$. When the packet size increases to 1,024 bytes, the detection rate drops to zero with $\theta \geq 7$. These results demonstrates that in order to maintain a high detection rate on TCP and UDP SYN attacks, we need to

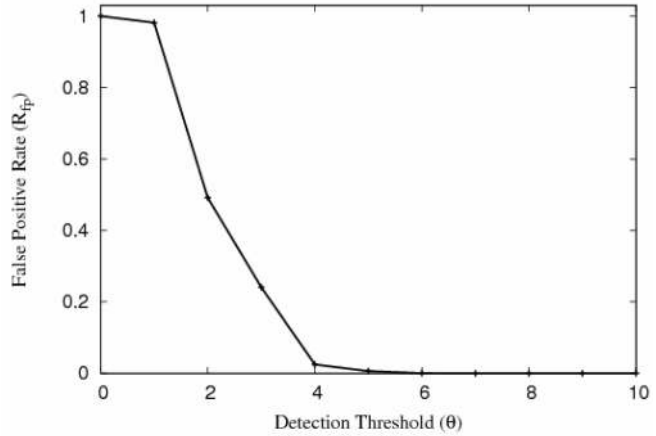


Fig. 14. Effects of the threshold on the false-positive rate in detecting TCP SYN attacks.

set θ to a small value such as $\theta = 5$ and adjust the packet size to 1,024 bytes.

6.4.2 False-Positive Alarms

Fig. 14 shows the false-positive alarm rate against the CAT server threshold θ . The number of alert generated by random fluctuation in normal traffic is small and negligible. With a server detection threshold $\theta = 4$, the false-positive rate drops to less than 1 percent. However, the real challenge lies in the fact that highly distributed attacks may use low packet rates to avoid being detected [39]. Only after sufficient attack flows are merged is the deviation detected by the routers. Hence, a small detection threshold value is required to achieve high detection accuracy with a low false-positive rate.

The ROC curve in Fig. 15 explains the trade-off between the detection rate and false-positive rate under various attacks. Our detection scheme achieves a detection rate as high as 99 percent with less than 1 percent false-positive rate for high-rate DDoS attacks. All three curves support this observation. Even for low-rate UDP attacks, our choice of a low CAT threshold (θ) accomplishes a detection rate of 91 percent at a false-positive rate of 23 percent. This result proves the effectiveness of the DCD detection mechanism




Fig. 13. Effects of the server threshold on the detection rate of three DDoS attack types.




Fig. 15. ROC curves showing the trade-off between the detection rate and false-positive rate of DDoS attacks.

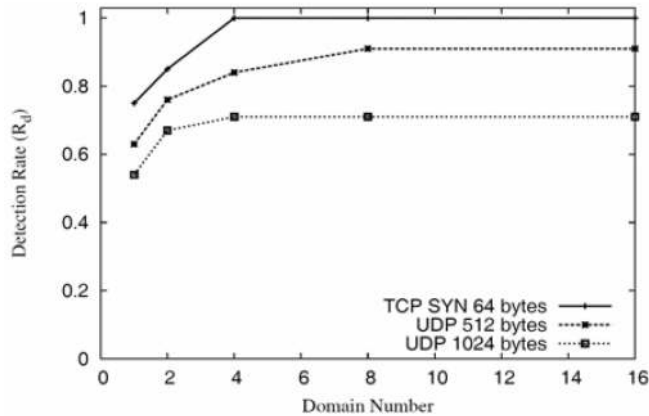


Fig. 16. Scalability in using the DCD defense system for DDoS defense, where four AS domains are sufficient to yield a 68-98 percent detection rate based on DETER simulation results.

7 SCALABILITY AND DEPLOYMENT LIMITATIONS

To deploy a distributed security scheme in ISP core networks, the scalability is related to the network size, domain number, data rate, link capacity, or router number involved. This section studies the scalability of the DCD scheme in terms of detection performance and the system overhead experienced. Then, we discuss the issues of flash crowd, security holes, and implementation limitations of the DCD system.

7.1 Domain Scalability in ISP Coverage

The study of domain scalability is driven by reducing the costs in system implementation and maintenance. One advantage of collaborative detection is its enlarged area of protection coverage. We have to use the CAT subtrees constructed by upstream domains to assess the earlier impact of a superflow generated by a DDoS attack. Even before the target network is overwhelmed, an early warning can prevent catastrophic disasters.

Fig. 16 plots the detection rates of three DDoS attack types against the number of domains used. Through experiments on the DETER testbed, we studied the effectiveness of cross-domain cooperation up to 16 AS domains. The detection rate becomes saturated after a sufficient number of AS domains are involved. The results are obtained under the system settings $\alpha = 0.1$, $\beta = 2.0$, $w = 500$ ms, and $\theta = 5$.

Recall that we simulated 8 to 10 routers per domain in scalable DETER experiments. With a small AS domain containing eight routers, $\theta = 5$ implies that more than half of the routers generated alerts as the attack flows approaches the root of the CAT. For 64-byte attacks, the optimal domain size is 4. For UDP 512-byte packets, the detection rate saturates at eight domains. For UDP 1,024-byte packets, four AS domains are sufficient to achieve the saturated detection rate.

The above analysis implies that 25 percent (4 out of 16) to 50 percent (8 out of 16) of participating network domains would be sufficient to detect a DDoS attack. Based on this level of domain participation, we find in Table 3 that the DCD system can scale well to cover $T \times 30.3$ percent \times 5 percent \times (25 to 50 percent) = 42 to 84 domains in commercial ISP core networks. This result is rather encouraging in practical applications.

This number is manageable, considering that the added monitoring burden of routers and the role of the CAT server in each AS domain. In reality, the decision process is ended much earlier by using a lower threshold θ . This leads to the conclusion that the DCD system scales well to cover a wider area of flooding DDoS attacks.

7.2 Implementation and Limitations

It is a big challenge to discriminate DDoS attacks from the fluctuation of legitimate traffic patterns, called *shrew attacks* [10] and *flash crowds* [23]. When a flash crowd happens, the CAT server creates a similar tree and could raise a false alarm. We suggest adding a few new features to separate the real DDoS attack traffic from the flash crowd.

The idea is to check newly appeared source IP addresses. For each superflow, in addition to traffic volume, we need to monitor the distribution of source IP addresses. Packet content matching offers another potential solution. However, this option is limited by the packet payload being encrypted.

Compromised insiders are often the most difficult problem to solve in security control. Infected routers may hide suspicious traffic patterns or send false alarms to the CAT server. These false alarms can weaken the use of the CAT as a means of attack detection.

We can make the DCD system more robust by introducing a topology verification procedure. Knowing the network topology, the CAT server is capable of rectifying the falsely reported traffic patterns according to the upstream and downstream relationship. A single or limited number of Byzantine defectors could be blocked this way.

The CAT server in the destination domain merges all received CAT subtrees to make a global detection. The global CAT provides useful information for traceback or pushback. In the literature, packet marking offers another option to trace the path of attack flows [1], [20]. The victim launches traceback or pushback operations only after a DDoS attack is fully detected. In contrast, our DCP system finishes the traceback task as soon as the merged CAT becomes sufficiently large. Our system pushes the defense line to upstream domains, where the traffic surges are first observed.

Internet AS resource distribution [21] suggests the use of up to 84 domains to cope with TCP SYN and UDP flooding attacks. This corresponds to a saturated detection rate (98 percent) and low false-alarm rate (below 1 percent). If we lower to a 65 percent detection rate with 15 percent false alarms, the server detection threshold (θ) can be further lowered. This implies that fewer AS domains could be used to make a decision. Based on our DETER experiments, it is sufficient to involve only 28 ISP domains in detecting most of the known DDoS flooding attacks.

8 CONCLUSIONS AND FURTHER RESEARCH

It is crucial to detect the DDoS flooding attacks at their early launching stage before widespread damage is done to legitimate applications on the victim system. We develop a DDoS detection system based on a new CAT mechanism. In conclusion, we elaborate on potential impacts and applications of the system:

1. *Detecting traffic changes at ATRs.* Based on the anomaly pattern detected in related network domains, our scheme detects a DDoS attack before the victim is overwhelmed. This approach captures the abrupt traffic changes at ATRs. The high detection rate of

DDoS attacks is expected with very low false-positive alarms.

2. *Scalable performance over multiple ISP domains.* Our DCD detection scheme is suitable for deployment at the ISP core networks. The provider-level cooperation eliminates the need for intervention by edge networks. Our DETER experimental results prove that four to eight domains are sufficient to yield a 98 percent detection rate of TCP SYN and UDP flooding attacks. Based on a recently reported Internet AS domain distribution, we expect the defense scheme to scale well to cover up to 84 ISP-controlled domains in a real-life Internet environment.
3. *SIP for interdomain trust negotiation.* To support inter-AS collaboration, SIP is proposed to resolve policy conflicts and regulate the alert message format. Our SIP is part of the USC/ISI effort in securing an Internet infrastructure against DDoS or worm attacks [19], [37] that threaten the availability, reliability, and dependability of Internet services.
4. *Valuable parameters from DETER experiments.* We have verified the effectiveness of the DCD scheme through emulation experiments on the DETER testbed. The engineering data on the traffic inertia factor α , router threshold β , global detection threshold θ , and monitoring window size w are very useful design parameters for building the real DCD prototype or production systems against DDoS attacks in the future.

Our distributed detection scheme automatically performs the traceback during the detection of suspicious traffic flows. Once a DDoS flooding attack is detected, we know the exact router or network domain where the anomaly was observed. In related projects, we have developed a spectral analysis method [10] to filter out *shrew* DDoS attacks with a low attack rate. Network worm containment helps in the defense against DDoS attacks, as reported in [6].

We suggest a hardware approach to implementing the CAT mechanism and SIP by using network processors or reconfigurable field-programmable gate array (FPGA) devices. This may demand the integration of signature-based IDs with anomaly detection systems [19]. The ultimate goal is to promote real-time detection and response against DDoS attacks with automatic signature generation.

ACKNOWLEDGMENTS

This work was supported by the US National Science Foundation (NSF) ITR Grant ACI-0325409 in the USC Internet and Grid Research Laboratory. The authors thank Cliff Neuman and Kevin Leahy at USC/ISI for assisting them to carry out the large-scale DETER experiments reported here.

REFERENCES

- [1] H. Aljifri, "IP Traceback: A New Denial-of-Service Deterrent," *IEEE Security and Privacy*, pp. 24-31, May/June 2003.
- [2] T. Anderson et al., "Rocketfuel: An ISP Topology Mapping Engine," <http://www.cs.washington.edu/research/networking/rocketfuel/>, 2006.
- [3] S. Bellovin, J. Schiller, and C. Kaufman, *Security Mechanism for the Internet*, IETF RFC 3631, 2003.
- [4] T. Benz et al., "Experience with DETER: A Testbed for Security Research," *Proc. Second IEEE Conf. Testbeds and Research Infrastructures for the Development of Networks and Communities (TridentCom '06)*, 2006.
- [5] R. Blazek et al., "A Novel Approach to Detection of DoS Attacks via Adaptive Sequential and Batch-Sequential Change-Point Detection Methods," *Proc. IEEE Workshop Information Assurance and Security*, June 2001.
- [6] M. Cai, K. Hwang, J. Pan, and C. Papadopoulos, "WormShield: Fast Worm Signature Generation with Distributed Fingerprint Aggregation," *IEEE Trans. Dependable and Secure Computing*, vol. 4, no. 2, Apr./June 2007.
- [7] G. Carl, G. Kesidis, R. Brooks, and S. Rai, "Denial-of-Service Attack Detection Techniques," *IEEE Internet Computing*, Jan./Feb. 2006.
- [8] A. Chakrabarti and G. Manimaran, "Internet Infrastructure Security: A Taxonomy," *IEEE Network*, Nov. 2002.
- [9] S. Chen and Q. Song, "Perimeter-Based Defense against High Bandwidth DDoS Attacks," *IEEE Trans. Parallel and Distributed Systems*, vol. 16, no. 6, June 2005.
- [10] Y. Chen and K. Hwang, "Collaborative Detection and Filtering of Shrew DDoS Attacks Using Spectral Analysis," *J. Parallel and Distributed Computing*, special issue on security in grids and distributed systems, pp. 1137-1151, Sept. 2006.
- [11] Y. Chen and K. Hwang, "Collaborative Change Detection of DDoS Attacks on Community and ISP Networks," *Proc. IEEE Int'l Symp. Collaborative Technologies and Systems (CTS '06)*, May 2006.
- [12] X. Dimitropoulos, D. Krioukov, G. Riley, and K. Claffy, "Revealing the Autonomous System Taxonomy: The Machine Learning Approach," *Proc. Passive and Active Measurement Workshop (PAM '06)*, 2006.
- [13] D. Dittrich, "The 'Stacheldraht' Distributed Denial of Service Attack Tool," <http://staff.washington.edu/dittrich/>, 2000.
- [14] M. Faloutsos, C. Faloutsos, and P. Faloutsos, "On Power-Law Relationships of the Internet Topology," *Proc. ACM SIGCOMM '99*, Aug. 1999.
- [15] V. Fuller, T. Li, J. Yu, and K. Varadhan, *Classless Inter-Domain Routing (CIDR): An Address Assignment and Aggregation Strategy*, IETF RFC 1519, 1993.
- [16] T. Gil and M. Poletto, "MULTOPS: A Data-Structure for Bandwidth Attack Detection," *Proc. 10th Usenix Security Symp.*, Aug. 2001.
- [17] K. Houle et al., "Trends in Denial of Service Attack Technology," www.cert.org/archive/pdf/, 2001.
- [18] A. Hussain, J. Heidemann, and C. Papadopoulos, "Identification of Repeated Denial of Service Attacks," *Proc. INFOCOM '06*, Apr. 2006.
- [19] K. Hwang, M. Cai, Y. Chen, and M. Qin, "Hybrid Intrusion Detection with Weighted Signature Generation over Anomalous Internet Episodes," *IEEE Trans. Dependable and Secure Computing*, vol. 4, no. 1, pp. 41-55, Jan.-Mar. 2007.
- [20] J. Ioannidis and S.M. Bellovin, "Implementing Pushback: Router-Based Defense against DDoS Attacks," *Proc. Network and Distributed System Security Symp. (NDSS '02)*, Feb. 2002.
- [21] ISO 3166 Report, "AS Resource Allocations," <http://bgp.potaroo.net/iso3166/ascc.html>, 2006.
- [22] H. Jiang and C. Dovrolis, "Why Is the Internet Traffic Bursty in Short Time Scales," *Proc. ACM SIGMETRICS '05*, June 2005.
- [23] J. Jung, B. Krishnamurthy, and M. Rabinovich, "Flash Crowds and Denial-of-Service Attacks: Characterization and Implications for CDNs and Web Sites," *Proc. 11th Int'l World Wide Web Conf. (WWW '02)*, 2002.
- [24] S. Kandula, D. Katabi, M. Jacob, and A. Berger, "Botz-4-Sale: Surviving Organized DDoS Attacks That Mimic Flash Crowds," *Proc. Second Symp. Networked Systems Design and Implementation (NSDI '05)*, May 2005.
- [25] S. Kent and R. Atkinson, *Security Architecture for the Internet Protocol*, IETF RFC 2401, 1998.
- [26] Y. Kim, W.C. Lau, M.C. Chuah, and H.J. Chao, "PacketScore: Statistics-Based Overload Control against Distributed Denial of Service Attacks," *Proc. INFOCOM '04*, 2004.
- [27] T. Law, J. Lui, and D. Yau, "You Can Run, But You Can't Hide: An Effective Statistical Methodology to Trace Back DDoS Attackers," *IEEE Trans. Parallel and Distributed Systems*, Sept. 2005.
- [28] R. Mahajan, S. Bellovin, S. Floyd, J. Ioannidis, V. Paxson, and S. Shenker, "Controlling High Bandwidth Aggregates in the Network," *Computer Comm. Rev.*, July 2002.

- [29] J. Mirkovic and P. Reiher, "D-WARD: A Source-End Defense against Flooding DoS Attacks," *IEEE Trans. Dependable and Secure Computing*, pp. 216-232, July 2005.
- [30] T. Monk and K. Claffy, "Cooperation in Internet Data Acquisition and Analysis," *Proc. Coordination and Administration of the Internet Workshop*, (CAIDA Project), <http://www.caida.org/>, Sept. 1996.
- [31] D. Moore, G. Voelker, and S. Savage, "Inferring Internet Denial-of-Service Activity," *Proc. 10th Usenix Security Symp.*, 2001.
- [32] P. Ning, S. Jajodia, and X.S. Wang, "Abstraction-Based Intrusion Detection in Distributed Environment," *ACM Trans. Information and System Security*, pp. 407-452, Nov. 2001.
- [33] C. Papadopoulos, R. Lindell, J. Mehringer, A. Hussain, and R. Govindan, "COSSACK: Coordinated Suppression of Simultaneous Attacks," *Proc. Third DARPA Information Survivability Conf. and Exposition (DISCEX-III '03)*, pp. 2-13, 2003.
- [34] T. Peng, C. Leckie, and K. Ramamohanarao, "Detecting Distributed Denial of Service Attacks by Sharing Distributed Beliefs," *Proc. Eighth Australasian Conf. Information Security and Privacy (ACISP '03)*, July 2003.
- [35] J. Postel, *Internet Control Message Protocol*, IETF RFC 792, 1981.
- [36] S. Ranjan, R. Swaminathan, M. Uysal, and E. Knightly, "DDoS-Resilient Scheduling to Counter Application Layer Attacks under Imperfect Detection," *Proc. INFOCOM '06*, Apr. 2006.
- [37] T. Ryutov, L. Zhou, C. Neuman, T. Leithead, and K.E. Seamons, "Adaptive Trust Negotiation and Access Control," *Proc. ACM Symp. Access Control Models and Technologies (SACMAT '05)*, June 2005.
- [38] G. Siganos, M. Faloutsos, P. Faloutsos, and C. Faloutsos, "Power-Laws and the AS-Level Internet Topology," *ACM/IEEE Trans. Networking*, pp. 514-524, Aug. 2003.
- [39] S.M. Specht and R.B. Lee, "Distributed Denial of Service: Taxonomies of Attacks, Tools and Countermeasures," *Proc. 17th Int'l Conf. Parallel and Distributed Computing Systems (PDCS '04)*, Sept. 2004.
- [40] J. Sommers and P. Barford, "Self-Configuring Network Traffic Generation," *Proc. ACM Internet Measurement Conf. (IMC '04)*, Oct. 2004.
- [41] M. Walfish, M. Vutukuru, H. Balakrishnan, D. Karger, and S. Shenker, "DDoS Defense by Offense," *Proc. ACM SIGCOMM '06*, Sept. 2006.
- [42] H. Wang, D. Zhang, and K. Shin, "Change-Point Monitoring for the Detection of DoS Attacks," *IEEE Trans. Dependable and Secure Computing*, vol. 1, Oct.-Dec. 2004.
- [43] X. Wang, S. Chellappan, P. Boyer, and D. Xuan, "On the Effectiveness of Secure Overlay Forwarding Systems under Intelligent Distributed DoS Attacks," *IEEE Trans. Parallel and Distributed Systems*, vol. 17, no. 7, July 2006.



Yu Chen received the PhD degree in computer engineering from the University of Southern California (USC) in 2006. He is an assistant professor of electrical and computer engineering at the State University of New York at Binghamton. His research interest includes network security, distributed denial-of-service (DDoS) defense, and Internet technology. He is a member of the IEEE and the ACM.



Kai Hwang received the PhD degree from the University of California, Berkeley, in 1972. He is a professor of electrical engineering and computer science at the University of Southern California. He specializes in computer architecture, Internet security, and distributed and peer-to-peer computing. He is a fellow of the IEEE Computer Society. For more details, visit <http://GridSec.usc.edu/Hwang.html>.



Wei-Shinn Ku received the MS and PhD degrees in computer science from the University of Southern California in 2003 and 2007, respectively. He is an assistant professor of computer science and software engineering at Auburn University. His research interests include spatial and temporal data management, geographical information systems, and network security. He is a member of the ACM and the IEEE Computer Society.

► For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.