

(19) World Intellectual Property Organization  
International Bureau



(43) International Publication Date  
13 September 2007 (13.09.2007)

PCT

(10) International Publication Number  
**WO 2007/103864 A2**

(51) International Patent Classification:  
*G06F 15/16* (2006.01)

(21) International Application Number:  
PCT/US2007/063296

(22) International Filing Date: 5 March 2007 (05.03.2007)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:  
60/778,662 3 March 2006 (03.03.2006) US

(71) Applicant (for all designated States except US): **NEW JERSEY INSTITUTE OF TECHNOLOGY** [US/US];  
323 Martin Luther King Blvd., University Heights,  
Newark, NJ 07102 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **GAO, Zhiqiang** [CN/US]; 1000 Central Avenue, Apt. 46A, Westfield, NJ 07090 (US). **ANSARI, Nirwan** [US/US]; 9 Barge Court, Montville, NJ 07045 (US).

(74) Agents: **WYCHE, Myron, Keith** et al.; Connolly Bove Lodge & Hutz LLP, 1990 M Street, N.W., Suite 800, Washington, DC 20036 (US).


(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, SV, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, MT, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

**Published:**  
— without international search report and to be republished upon receipt of that report

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: BEHAVIOR-BASED TRAFFIC DIFFERENTIATION (BTD) TO DEFEND AGAINST DISTRIBUTED DENIAL OF SERVICE(DDoS) ATTACKS



(57) Abstract: Embodiments are directed toward a method for Behavior-based Traffic Differentiation (BTD) that initially receives incoming packets and performs traffic classification to determine the protocol of the incoming packets. In addition, BTD performs bandwidth division/allocation to further support traffic classification amongst non-TCP traffic types such as UDP and ICMP. For TCP traffic, the method for BTD determines whether a TCP connection has been established and performs at least one of rate limiting, waiting time reduction for half-open connections, and incrementing backlog queue size, when the TCP connection has not been established. If the TCP connection has been established successfully, the method for BTD further includes proactive tests for traffic differentiation which identify normal traffic, which is admitted, and attack traffic, which is dropped.

WO 2007/103864 A2

**BEHAVIOR-BASED TRAFFIC DIFFERENTIATION (BTD) FOR  
DEFENDING AGAINST DISTRIBUTED DENIAL OF SERVICE (DDoS)  
ATTACKS**

5 **DESCRIPTION OF BACKGROUND ART**

The proliferation of e-business applications and mission-critical military communications has raised serious concerns on network security on the Internet. The advent of the lethal Denial of Service (DoS) attack and its advanced variant,  
10 the Distributed DoS (DDoS) attack are troublesome intruders on our usage and dependence on the Internet. The detrimental impact of DoS/DDoS attacks has been demonstrated again and again, even on such high-profile sites as Yahoo, CNN, Ebay and Amazon.

15 In a DDoS attack, an attacker sends a large volume of malicious traffic to a victim. For example, a DDoS attacker may infiltrate one or a plurality of computers at various data centers via a computer system connected to the Internet. Often the attacker will access the Internet through an Internet Service Provider (ISP). The attacker can then place the plurality of computers at the data centers  
20 under its control by use of a malicious software program. When the attacker issues a command, these computers can simultaneously send out large volumes of data at various times to the victim preventing the victim from responding to legitimate Internet traffic and messages.

25 Currently, DDoS attacks are probably the most ferocious threats to the Internet community. DDoS attacks exploit and consume the limited available resources of a system by sending superficially normal but really useless packets to degrade/corrupt the victim system, thus severely hampering the victim system from serving its normal clients. Typical resources that get drained in DoS/DDoS  
30 attacks are network bandwidth, CPU cycles of the target host, and specific TCP/IP protocol stack structure such as fragmentation buffer and TCP SYN buffer. In addition, readily accessible attack scripts prevalent in the Internet significantly decrease the technical hurdle to bombard DDoS attacks. Lack of accounting in the

Internet further allures crackers to take their chances to attack others without worrying about any subsequent penalties.

It is well known that it is rather easy to launch, but difficult to defend, a  
5 DDoS attack. The underlying reasons include: (1) IP spoofing (i.e., attack packets routinely carry forged source IP addresses, which effectively conceal the identity of attack sources and deter the efforts of detection, defense, and tracing); (2) the distributed nature of the DDoS attack (i.e., a huge number of sources generate  
10 attack traffic simultaneously that impose an overwhelming burden on any countermeasure by invoking issues of scalability to handle increasingly powerful attacks); (3) no simple mechanism for the victim to distinguish the normal packets from the lethal attack traffic.

As a result of the above, improving the sustainability of networks and  
15 hosts (especially servers) by defending against DDoS attacks are an important goal in this technology area. A critical issue in DDoS defense is how to isolate the attack traffic from the normal traffic. This issue is referred to as “traffic differentiation” and is of great importance because the goal of DDoS attack is to severely degrade the performance of target hosts and networks or even completely  
20 deprive the victim of the capability of serving its normal clients. With the knowledge of which is "attack" and which is "normal" traffic in hand, the victim is ready to defeat a DDoS attack by reacting differently according to the type of traffic (i.e., attack or normal).

25 In addition, given the diverse DDoS attack patterns, another issue of importance is how to contain as many DDoS attack patterns as possible. Exemplary DDoS attack patterns are given in: C. Douligeris and A. Mitrokotsa, “DDoS attacks and defense mechanisms: classification and state-of-the-art,” *Computer Networks*, vol.44, pp. 643-666, 2004; and J. Mirkovic and P. Reiher, “A  
30 taxonomy of DDoS attack and DDoS defense mechanisms,” *ACM SIGCOMM Computer Communications Review*, vol. 34, no. 2, pp. 39-53, Apr. 2004.

From the perspective of protocol exploits, DDoS attacks may be TCP, UDP, ICMP, or other protocols based. Some attacks use the combination of different protocols as shown in **Table 1** below. From the point of view of the attack rate, most attacks are high-speed flood-based while a novel and more sophisticated attack is low-rate DDoS attack as given in: R. Chang, "Defending against flooding-based, Distributed Denial of Service attacks: a tutorial," *IEEE Communications Magazine*, vol. 40, no. 10, pp.42-51, 2002; and A. Kuzmanovic and E. knightly, "Low-Rate TCP-Targeted Denial of Service Attacks (The Shrew vs. the Mice and Elephants)," *ACM SIGCOMM 2003*, Aug. 2003, pp.75-86.

10

**Table 1. Diverse DDoS attack tools**

Protocols Used	DoS/DDoS Attack Tools/Names
TCP only	SYN flood, RST flood, mstream
UDP only	trinoo, shaft
ICMP only	Ping of death, flood pinging, Smurf
Combinations of TCP, UDP, and ICMP	TFN, TFN2k, MIX, Stacheldraht, trinity,v3

Unfortunately, in contrast to attack schemes, defense schemes do not keep pace with the evolution of DDoS attacks. Most background art DDoS defense schemes aim to address one or two types of DDoS attacks and are inefficient and ineffective to the wide spectrum of possible patterns of DDoS attacks.

20

**SUMMARY**

Embodiments are directed at overcoming the foregoing and other difficulties encountered by the background arts. In particular, embodiments of behavior-based traffic differentiation (BTD) provide a novel framework that can proactively identify and deter most of the malicious attack traffic, and tackle a variety of DDoS attack patterns as well. Two components of the framework are

25

traffic classification and traffic differentiation. Traffic classification aims to address non-TCP flood attacks while traffic differentiation is used to identify malicious TCP flows. Embodiments of BTD are implemented, for example, but not limited to ns-2. Experimental results indicate that traffic classification can improve the throughput of TCP traffic significantly (i.e., by over 70%) while traffic differentiation can quickly block malicious attack traffic. Other benefits of our embodiments of BTD include, but are not limited to: (1) minimal requirements for equipment/system modification; (2) practically deployable; (3) no issues of scalability because the deployment is at the receiver side only; (4) no issue of lack of incentives since the deployment of our scheme is served to solely protect one's own networks and hosts, not others.

### BRIEF DESCRIPTION OF THE DRAWINGS

**FIG. 1** is an exemplary illustration showing the timing of how low-rate DoS attacks are mounted by sending attack packets sporadically using an on-off model.

**FIG. 2** depicts an exemplary flowchart of an embodiment of the BTD framework.

**FIG. 3** is an exemplary illustration showing how traffic classification may be implemented at the entry point of the protected networks or systems at either a firewall or an edge router.

**FIG. 4** is an exemplary flowchart for a method for traffic differentiation in embodiments of BTD.

**FIG. 5** illustrates an exemplary simulation scenario including one FTP source and sink (using TCP), and one CBR source and sink (using UDP).

**FIG. 6(a)** is an exemplary illustration of TCP goodput and TCP throughput for the TCP flow and UDP throughput without the traffic classification policy.

**FIG. 6(b)** is an exemplary illustration of TCP goodput and TCP throughput for the TCP flow and UDP throughput with the traffic classification policy.

FIG. 7 illustrates exemplary experimental results showing the effectiveness of traffic differentiation in a simulation scenario including one FTP source and an attack source.

FIG. 8(a) are exemplary experimental results showing the throughput of the attack traffic using the FTP sink in embodiments of BTD.

FIG. 8(b) are exemplary experimental results showing the throughput of the attack traffic using TCP smart sink in embodiments of BTD.

## DETAILED DESCRIPTION

10

Embodiments of BTD attempt to isolate attack traffic and contain as many DDoS attack schemes as possible. DDoS attacks are a resource management problem, and thus embodiments adopt a QoS means to combat them. Embodiments of BTD use an integrated scheme to efficiently and effectively contain a variety of DDoS attack schemes rather than just one or two attack schemes, as in the background art. Embodiments of BTD can differentiate between malicious attack traffic and normal traffic, and can punish the malicious attack traffic. Embodiments of BTD are comprised of at least 2 components. The first component classifies different types of traffic based on the protocols used and limits their rates accordingly with bandwidth allocation. This method serves to isolate UDP, ICMP and other traffic from TCP, and is helpful in mitigating some flood-attacks based on UDP and ICMP via limiting bandwidth allocation to non-TCP traffic.

TCP is the dominant traffic on the Internet and most DDoS attacks are based on TCP. Thus, the next step is to differentiate disparate TCP traffic. All TCP traffic is categorized into two groups according to the status of the connection establishment. Among all connection-established TCP traffic, embodiments of BTD attempt to identify the properties of a flow, whether it is benign or malicious, according to its behavior. A flow is defined as "benign" or "normal" if it responds to the control signal of the other endpoint of the same connection appropriately. On the contrary, "malicious" or "attack" flows are those that do not follow the TCP congestion control principle and act aggressively.

30

It is possible that a TCP unfriendly flow is classified as a lethal flow, and is punished by embodiments of the BTM scheme. That is, a collateral side effect of embodiments of BTM is that it currently does not distinguish between a TCP unfriendly flow and an attack flow. Based on the distinction, certain penalties may be imposed to the aggressive flows of unfriendly flows.

In addition, embodiments of BTM are "victim-centered" and do not require any modification of the source endpoints or intermediate routers. This characteristic of BTM avoids the issues of scalability, cooperation between different domains, and lack of incentives. Extensive simulations have been conducted and experimental results are presented to validate the design of embodiments of BTM.

An ideal defensive scheme shall be able to distinguish between attack flows and normal ones, and isolate deleterious traffic accordingly. However, it is by no means trivial to make such a distinction. Embodiments of BTM identify malicious attack traffic from their behaviors. Besides IP spoofing, aggressiveness is the salient feature/behavior of DDoS traffic. One example of the aggressive behavior is that an attack source may not care about whether it may receive a response from the victim or not. If this is the case, an attacker can still conduct an attack whether or not a response is received by bombarding its target with a monstrous number of useless packets. Even the low-rate DoS attacks that mount attack packets sporadically using an on-off model, as shown in **FIG. 1**, demonstrate such aggressive features during the "on" stage.

Note that "aggressiveness" is not equivalent to "high-rate". It is possible that a high-rate flow is a normal TCP flow. The receiver may identify the aggressive behavior by intentionally testing the response of a source upon certain control signals from the receiver. Any source which fails such tests is regarded as a lethal attack source and can be punished accordingly.

However, a source that passes the test above may not necessarily be benign. The receiver (usually a server) is suggested to conduct the test upon any receipt of a request/requirement for a connection. That is, a sophisticated attacker may pass an initial test by behaving well at first and perform deleterious attack operations later. To handle this situation, the receiver may increase the frequency of such tests used to detect an attack flow. An alternative solution is to introduce

some dynamics into the test and randomly determine the frequency and interval of the test for each flow, especially the high-rate one.

To accommodate high-rate legitimate traffic better, a threshold is set that defines the maximal number of successful tests for a flow. No more tests are  
5 conducted on packets from a flow once the flow successfully passes the specified number of tests. By actively testing a source, the receiver can determine with high confidence the nature of a flow (i.e., normal or attack) from that source and react accordingly.

Filtering based on behavior brings an attack source into a dilemma: (1)  
10 sending packet aggressively at the risk of being identified and punished; or (2) reducing attack rate to meet the requirements of the receiver so that the effect of an attack is diminished. In producing this dilemma for the attack source, embodiments of BTM allow the receiver to throttle the scope and impact of potential attacks.

The above design is feasible for TCP solely because TCP has the built-in  
15 congestion control and reliable transmission mechanism. Note that TCP is the dominant traffic on the Internet, and as much as 90% of DDoS traffic use TCP. Currently, TCP occupies 80% in terms of the number of flows, 90% with respect to the number of packets. Thus, it is essential for DDoS defense schemes to  
20 accommodate TCP traffic effectively and efficiently.

However, other traffic such as UDP and ICMP lack alternative mechanism  
that are exploitable by our framework. Embodiments of BTM use traffic  
classification to handle non-TCP traffic. Traffic classification is a simple and  
efficient solution to handle flood attacks based on UDP and ICMP. By limiting the  
25 resource allocated to UDP and ICMP traffic, the receiver may significantly mitigate UDP and ICMP flood attacks.

Another benefit of traffic classification is that UDP traffic is much more  
aggressive than TCP traffic. It is well established that when UDP traffic competes  
for bandwidth against TCP traffic, UDP is capable of seizing most of the available  
30 bandwidth due to its lack of any congestion control mechanism and being more aggressive in nature. Therefore, sending UDP traffic may deprive the fair share of bandwidth from the TCP traffic effectively, and degrade the performance of TCP streams. Thus, it turns out that UDP is a good vehicle for DDoS attacks. This

demonstrates the difficulty in containing different kinds of DDoS attack patterns in one scheme, and justifies the necessity to also perform traffic classification.

Other issues that need to be accounted for include, but are not limited to: (1) methods for addressing IP spoofing; and (2) determining at what level the framework of embodiments should be performing (e.g., packet level or flow level). First, one identifies attack traffic with forged sources IP addresses, and then discerns the normal traffic with genuine IP addresses by observing their behaviors. Spoofing may be addressed by means of verification while aggressive behavior may be identified by manipulating the rate of returned ACK at the recipient side for TCP traffic. This method is also capable of identifying the more subtle and sophisticated low-rate DDoS attacks.

**FIG. 2** depicts an exemplary flowchart of embodiments of the BTD framework. In **Step 2.1** of **FIG. 2**, incoming packets are received by the BTD framework. At **Step 2.2** a method for traffic classification is implemented. **Step 2.3** implements a method of bandwidth division/allocation. In **Step 2.4**, a determination of whether the TCP connection has been established is made. If the TCP connection has not been established, at least one of rate limiting; waiting time reduction for half-open connections; and incrementing backlog queue size are performed in **Step 2.5**. Alternatively, if the TCP connection has been established a method for traffic differentiation is performed in **Step 2.6**.

As the result of the method for traffic differentiation, a determination is made of whether the packet is normal traffic or attack traffic in **Step 2.7(a)** and **Step 2.7(b)**, respectively. If it is determined in **Step 2.7(a)** that normal traffic is being received, the received packet is admitted by the BTD framework in **Step 2.8(a)**. Alternatively, if it is determined in **Step 2.7(b)** that attack traffic is being received, the received packet is dropped and denied admission by the BTD framework in **Step 2.8(b)**. For a heavily loaded receiver (e.g., a server) that serves a huge amount of packets from disparate sources, embodiments of BTD shall make a fast and smart decision to admit or drop an incoming packet. The ever-increasing link-speed motivates us to contrive countermeasures at the flow level.

The methods for Traffic Classification and Bandwidth Allocation are further discussed in the following paragraphs. In particular, methods for Traffic

Classification in embodiments of BTD are discussed in the following. Since, as discussed above, UDP lacks built-in congestion control mechanisms and is aggressive in nature, embodiments of BTD first classify traffic based on protocols to isolate TCP and non-TCP traffic. Specifically, embodiments of BTD classify traffic into 4 groups: TCP, UDP, ICMP, and other. Embodiments of BTD implement this classification at least by checking the protocol field at the IP layer. Regarding bandwidth division/allocation, as a non-limiting example, a system administrator may configure its network and assign the bandwidth in the following way: TCP: 85%, UDP: 13%, ICMP 1%, and others 1%. This specific configuration is determined by the site's profile, and may be changed dynamically.

The goal of the method of traffic classification of embodiments of BTD is to determine, according to normal traffic profiles, whether it is necessary to implement a defense to mitigate the impact of possible UDP and ICMP flood attacks. Unlike TCP traffic, whose property may be identified according to its behavior, UDP is aggressive in nature. As a less "reliable" protocol, it turns out that no easy way may be taken to distinguish between normal UDP traffic and attack UDP traffic. However, by imposing a bandwidth limit, embodiments of BTD provide a simple but effective strategy to distinguish between the two and does not affect the network performance in most cases.

In addition, limiting the amount of ICMP traffic allowed into a network is a good practice because it is used solely for diagnostic purpose and ICMP-based DDoS attack are common.

**FIG. 3** shows how traffic classification may be implemented at the entry point of protected networks or systems, either at a firewall **3.1** or an edge router **3.2**. The availability of network processors and other ASICs specifically designed for classification renders it possible for embodiments of BTD to be operated at line rate.

In considering the bandwidth allocation/division function of embodiments of BTD, the quota for each protocol is configurable and it is normally determined by the routine traffic profile at one site. As discussed above and in accordance with recent measurements, the dominant traffic in the current Internet is still TCP. Thus, it is reasonable to assign most of the bandwidth to TCP traffic for a normal site because most killer applications such as web, email, telnet, and FTP all use

TCP. For example, a possible bandwidth allocation/division among different protocols is TCP: 88%, UDP: 10%, ICMP: 1%, and other: 1%. In case of a traffic pattern change in the future, what embodiments of BTD provide is a variable bandwidth allocation that is able to adjust the bandwidth allocation percentage according to the new traffic model, and thus, embodiments of BTD can be easily tailored for future traffic changes.

The method of TCP Flow Differentiation in embodiments of BTD is discussed in the following paragraphs. Connection establishment determines whether or not a connection has been established and means a lot to the receiver. A successfully established connection indicates that both ends have completed the three-way handshaking procedure.

An even more important determination in DDoS defense is whether or not IP spoofing is being used by the source. For an incomplete connection, on the other hand, the receiver shall be alert, and be conservative in its resource consumption. Possible measures to mitigate potential attacks include, but are not limited to: (1) tightening the total bandwidth allocation to all incomplete connections; (2) significantly reducing the timeout value to avoid buffers occupied by half-open connections for a long time, or providing no buffer allocation at all for half-open connections; and (3) increasing the buffer size for backlogs. These options are shown in **FIG. 2** at **Step 2.5** and discussed above

Analyzing normal or benign flows and lethal or attack flows are discussed in the following. After assigning the quota for each protocol, the next task is to isolate lethal or attack TCP flows from benign or normal TCP flows. This is done simply because TCP consumes most of the bandwidth of the protected network. As mentioned earlier, straightforwardly adopting fair packet queuing is impractical in case of flood-based DDoS attacks when a large majority of traffic is attack streams. TCP is an end-to-end solution that requires close orchestration between the sender and the receiver. To identify the nature of a TCP flow (i.e., after a successful connection), the receiver can actively test the response of the sender by delaying the ACK packets intentionally. If the sender is normal, it will take action accordingly and reduce its sending rate.

On the contrary, for a DDoS attack, two cases may occur. One is that the sender uses forged source IP addresses, and thus cannot receive the rate-reduction

message from the receiver. It has no idea of the proper sending rate. The other scenario is that the sender does receive the notification, but it neglects it and just keeps sending packets, thus violating the protocol, and it may be punished by the recipient to reduce its share or even block its traffic. This procedure is dynamic.

- 5 The protected site can decide the frequency and extent of rate-reduction so that no perpetrator can easily fool the system to believe that the traffic from the perpetrator is normal.

**FIG. 4** depicts an exemplary flowchart for a method for traffic differentiation in embodiments of BTD. Upon the arrival of a new incoming packet at **Step 4.1**, the receiver first determines where the flow of the current packet belongs by checking the tuple of (source IP address, source port number, destination IP address, destination port number) also in **Step 4.1** of **FIG. 4**. In **Step 4.2**, if the method determines the packet is the first packet of a new flow, the receiver examines whether the number of admitted flows reaches the maximal flow count in **Step 4.3** relative to a threshold set by the receiver to ensure provisioning of proper quality of service. If the maximal flow count does occur, the packet is dropped in **Step 4.5**. Otherwise, after updating the flow table maintained by the receiver, incrementing the flow count by 1, and initializing several counters, such as the number of successful tests and the number of failure tests in **Step 4.7**, the new packet is admitted in **Step 4.9**.

For the packet of an existing flow, the receiver checks the behavior history of the flow. In **Step 4.3** of **FIG. 4**, if the number of failure tests is no less than a threshold,  $f$ , and the packet will be dropped in **Step 4.5**. An integer larger than 1 is selected to prevent embodiments of BTD from falsely identifying the behavior of a flow. A low value of  $f$  may exacerbate packet dropping. In case of a false identification, subsequent packets from an innocent flow will be blocked. Selecting a too high value is unwise, either. A high  $f$  delays the packet dropping decision, and thus subsequent packets of a malicious stream may still consume system resources. Experimental results indicate that  $f=3$  provide a good balance between the proper identification rate and the acceptable performance impact.

In addition, it is worth mentioning that a receiver has a couple of options to choose to punish the source at this point. One option is to send DUPACK on purpose forcing the source into the stage of slow start. Another is to send RST to

halt the connection so that its resource is not wasted by the misbehaving sources. In **FIG. 4**, we only show the operation of packet dropping without any further punishment.

For the flow whose behavior is not so bad in the past, embodiments of  
5 BTD further examine whether the flow has passed a certain number of tests,  $h$  as shown in **Step 4.6** of **FIG. 4**. The receiver will admit directly any packet of flows having passed  $h$  tests successfully (e.g., some tradeoff has to be made to determine a proper value of  $h$ . Selecting a high  $h$  value may hurt the performance of a good but high rate flow, while the tests may be easily evaded by sophisticated attackers  
10 using a low value. In **Step 4.6**, we set  $h$  to 6 after extensive simulations).

For other flows, embodiments of BTD further check the current state of the flow in **Step 4.8** of **FIG. 4**. If the flow is under a test, its current rate shall not exceed one half of its previous one (i.e., the receiver enforces this constraint by manipulating the reverse ACK rate). If the flow conforms to that constraint, the  
15 flow passes the current test in **Step 4.13**; the `pass_num` is incremented by 1 in **Step 4.14**; and the packet is admitted in **Step 4.16**. Otherwise, the flow fails one test in **Step 4.13**; the `fail_num` is incremented by 1 in **Step 4.15**; and the packet is dropped in **Step 4.17**.

In the case that **Step 4.8** determines the flow is not in the state of testing,  
20 its sending rate is compared with that of the fair share of each flow. The result of the comparison is used to determine the test probability for that flow in **Step 4.10**. **Step 4.12** determines whether a test is needed. Obviously, a flow with a smaller bandwidth assumption is subject to a smaller number of tests. **Step 4.1** determines the waiting time between tests. The test probability  $p$  for a high-rate  
25 flow (over the fair share) is  $1/(\text{pass\_num}+1)$ . At the very beginning, `pass_num` is 0 for all flows. Therefore, as long as a high-rate flow has not passed a test, its chance of being tested is 100%.

As the number of successful tests of a flow increases, its test probability reduces. The test probability  $p$  for less resource-consumption flow is  $1/\max(m, 2*h)$ , where  $m$  is the total number of flows. For the normal case,  $m \gg 2h$ ; thus,  
30  $p=1/m$ . We use the `max(.)` function to address the case that only a few flows exist in the system and ensure that the test probability for a low-rate flow is at most 1/2 of that of a high-rate one.

The rate of a flow is calculated according to the following formula,

$$(\text{num\_pkt} * \text{sz\_pkt}) / t,$$

(1)

5           where  $t$  is the time interval (window),  $\text{num\_pkt}$  is the number of packets received during this period,  $\text{sz\_pkt}$  the packet size. It is worth mentioning that the flow rate calculated here is not the average rate of a flow, as normally used by others, because we update the starting time of a flow once it passes a test. In so doing, we can effectively thwart low-rate DoS attack which sends a burst of attack  
10           packets to incite congestion and keeps silent for a much longer period to significantly lower its average rate in order to escape detection and filtering.

          Four scenarios may happen: (1) an attack source always behaves well, and thus the effect of an attack is greatly diminished; (2) an attack source behaves well at first and misbehaves later. In this case, the constraint that the current rate is at  
15           most 1/2 of previous rate upon the test state will not be satisfied, and the source fails the test; (3) an attack source always misbehaves, that may be easily thwarted by the fail count; and (4) an attack source misbehaves at first and behaves well later. For scenario (4), the attack source is exposed to more chances of being tested because its  $\text{pass\_num}$  is offset by the  $\text{fail\_num}$  once it fails a test. It is  
20           noted that a low-rate flow is also subject to test, though at a lower probability in our design. As time passes by, the chance that a low-rate flow has never been tested by the receiver is very low. Embodiments of BTD enforce this policy to contain the case that some low-rate streams are malicious.

          The Penalty for Bad TCP Flows is discussed in the following. A receiver  
25           has a couple of options to choose to punish the source. One option is to send DUPACK on purpose, thus forcing the source into the stage of slow start. Another is to send RST to halt the connection so that its resource is not wasted by the misbehaving sources. Embodiments of BTD send RST once the flow fails the tests 3 times. In **FIG. 4**, the operation of packet dropping is shown without any  
30           further punishment.

          Experimental results regarding traffic classification are discussed below. To test the effectiveness of traffic classification, we set up a simple simulation

scenario including 1 FTP source and sink (using TCP), and 1 CBR source and sink (using UDP), as shown in **FIG. 5**. These flows pass through the same bottleneck link. The link parameters between the FTP source and the checking point, the CBR source to the checking point, destination and CBR sink, and destination and FTP sink are all 10Mb in bandwidth, and 2 ms in delay. The bottleneck link between the checking point and the destination is set to be 1 Mb and 10ms. The CBR rate is 10Mb. The traffic classification for TCP is 90%, and UDP10% at the checking point. In the comparison of simulations, everything is the same except whether the entry point enforces traffic classification or not. The throughput of the FTP traffic is presented in **FIG. 6(a)** and **FIG. 6(b)**.

**FIG. 6(a)** depicts TCP goodput and TCP throughput for the TCP flow and UDP throughput without the traffic classification policy. In **FIG. 6(a)** and **FIG. 6(b)**, TCP traffic starts at time 0s and UDP traffic begins 1s later. During 0-1s, TCP is the only traffic in the link and its goodput and throughput reaches 12 packets per 0.1s, the maximal value (the bandwidth of the bottleneck link is 1mb/s, each packet is  $1000 \times 8 = 8000$  bits,  $1\text{mb} \times 0.1 / 8000 = 12$ ). From 1.2s on, UDP traffic starts to merge and it captures the available bandwidth so fast that TCP throughput reduces to 1 packet during 1.3-1.8s, starving since 1.9s. With traffic classification, in contrast, TCP throughput remains 3 packets per 0.1s even after 1.9s, while UDP traffic is no more than 9 packets per 0.1s. The total number of packets sent is 127 in **FIG. 6(a)**, 234 in **FIG. 6(b)** (we only depict the throughput till 1.9s in **FIG. 6(b)** for comparison). Therefore, the throughput of TCP traffic is increased by 70.8% with traffic classification. Though UDP traffic still seizes more share than TCP, as explained above, this is due to the built-in congestion control mechanism of TCP.

The following discusses TCP Flow Differentiation. To test the effectiveness of traffic flow differentiation, a simulation scenario including an FTP source and an attack source was set up, as shown in **FIG. 7**. These flows pass through the same bottleneck link. The difference is that one simulation uses a normal FTP sink to accept packets from both flows, another uses our developed TCP sink, called TCP smart sink. The simulation results are shown in **FIG. 8(a)** and **FIG. 8(b)**.

FIG. 8(a) shows the throughput of the attack traffic using the FTP sink while FIG. (8b) presents the throughput of the attack traffic using the TCP smart sink in embodiments of BTD. At 3.2s, the throughput of attack traffic drops drastically. After 42.3s, the attack traffic is totally blocked. In contrast, using the background art FTP sink as the receiver, the sender may keep the highest throughput during its lifetime. The result demonstrates the effectiveness of the traffic differentiation approach of embodiments of BTD.

As discussed above, embodiments of the BTD comprise two components: traffic classification and traffic differentiation. The former is used to reduce the volume of Non-TCP traffic while the latter can identify malicious TCP flows via proactive tests. The salient benefits of embodiments of BTD are listed as follows:

BTD can effectively and efficiently identify and block attack flows via proactive tests;

BTD can contain many DDoS attack patterns;

BTD requires minimal modification; and

BTD has no issues such as scalability and lack of incentives.

Simulation and the above-discussed experimental results have validated our design. In addition, the following are additional aspects of embodiments of BTD: (1) identifying the network condition automatically; (2) selecting various parameters used for proactive testing adaptively; and (3) implementing the framework in operating system environments such as, but not limited to: Linux kernel.

In summary, embodiments of BTD first try to reduce the volume of Non-TCP traffic via traffic classification, thus mitigating the popular UDP and ICMP flood attacks. For TCP flows, embodiments of BTD attempt to distinguish their nature by observing behaviors. Table 2 lists the policies adopted in the framework of embodiments of BTD to accommodate the normal traffic and to constrain the attack traffic.

**Table 2.** Measures taken by embodiments of BTD to address diverse traffic models

traffic nature	Measures
high-rate attack traffic	high test probability dynamic in terms of detection frequency and intervals after certain times of failures, all packets of the flow are dropped
high-rate good traffic	after passing the test for certain times, all packets of the flow are accepted
low-rate attack traffic	still subject to tests, though at a low probability. As time passes by, its chance of being tested increases. Once a flow passes a test, its rate is calculated from the time when it passes the last test. Thus, a flow adopting the strategy of low-rate first and high-rate later cannot escape being identified. after certain times of failures, all packets of the flow are dropped
low-rate good traffic	low test probability after passing the test for certain times, all packets of the flow are accepted

It will, of course, be understood that, although particular embodiments have just been described, the claimed subject matter is not limited in scope to a particular embodiment or implementation. For example, one embodiment may be

5 in hardware, such as implemented to operate on a device or combination of devices, for example, whereas another embodiment may be in software. Likewise, an embodiment may be implemented in firmware, or as any combination of hardware, software, and/or firmware, for example. Likewise, although claimed subject matter is not limited in scope in this respect, one embodiment may

10 comprise one or more articles, such as a storage medium or storage media. This storage media, such as, one or more CD-ROMs and/or disks, for example, may have stored thereon instructions, that when executed by a system, such as a computer system, computing platform, or other system, for example, may result in an embodiment of a method or apparatus in accordance with claimed subject

15 matter being executed, such as one of the embodiments previously described, for

example. As one potential example, a computing platform may include an apparatus or means for implementing one or more processing units or processors, one or more input/output devices, such as a display, a keyboard and/or a mouse, and/or one or more memories, such as static random access memory, dynamic  
5 random access memory, flash memory, and/or a hard drive and one or more router of firewall type equipments. For example, a display may be employed to display one or more queries, such as those that may be interrelated, and or one or more tree expressions, although, again, claimed subject matter is not limited in scope to this example. Likewise, an embodiment may be implemented as a system, or as  
10 any combination of components such as computer systems and interfaces to computer systems (for example but not limited to: routers and firewalls), mobile and/or other types of communication systems and other well known electronic systems.

In the preceding description, various aspects of claimed subject matter  
15 have been described. For purposes of explanation, specific numbers, systems and/or configurations were set forth to provide a thorough understanding of claimed subject matter. However, it should be apparent to one skilled in the art having the benefit of this disclosure that claimed subject matter may be practiced without the specific details. In other instances, well known features were omitted  
20 and/or simplified so as not to obscure the claimed subject matter. While certain features have been illustrated and/or described herein, many modifications, substitutions, changes and/or equivalents will now occur to those skilled in the art. It is, therefore, to be understood that the appended claims are intended to cover all such modifications and/or changes as fall within the true spirit of claimed subject  
25 matter.

**What is claimed is:**

1. A method for Behavior-based Traffic Differentiation (BTD), comprising:  
receiving incoming packets;  
5 performing a method for traffic classification configured to determine a protocol of the incoming packets;  
performing a method of bandwidth division/allocation;  
determining whether a TCP connection has been established, wherein when the TCP connection has not been established:  
10 performing at least one of rate limiting, waiting time reduction for half-open connections, and incrementing backlog queue size, and wherein when the TCP connection has been established:  
performing a method for traffic differentiation,  
wherein when a determination is made for normal traffic,  
15 admitting the packet; and wherein a determination is made for attack traffic, dropping the packet.
2. The method according to claim 1, wherein the method for traffic  
20 classification further comprises at least checking a protocol field at an Internet Protocol (IP) layer for the incoming packet to determine the packet classification.
3. The method according to claim 2, wherein the checked protocol field determines whether the packet is classified as one of TCP, UDP, ICMP and OTHER packets.
- 25 4. The method according to claim 3, wherein when the checked protocol field is at least one of UDP, ICMP and OTHER packets, a method of bandwidth division/allocation further comprises determining the bandwidth allocated for the incoming packet; and, classifying the incoming packet as one of UDP, ICMP and OTHER based on the determined bandwidth.

5. The method according to claim 4, wherein the method of traffic classification is implemented at a firewall or edge router.
6. The method of claim 5, wherein the method of traffic differentiation further comprises:
- 5 checking a source IP address, source port number, destination IP address, and destination port number of an incoming packet;
- determining whether the incoming packet is a first packet of a new flow or one of a plurality of admitted flows,
- determining when the number of admitted flows reaches a maximal flow
- 10 count relative to a predetermined first threshold set to ensure provisioning of proper quality of service (QoS);
- updating a maximal count flow table by incrementing the flow count by 1;
- incrementing counters for successful tests and failure tests when maximal flow count is at least one of below and equal to the predetermined first threshold
- 15 and admitting the incoming packet of the new flow;
- dropping the incoming packet when the maximal flow count exceeds the predetermined threshold;
- admitting a incoming packet from an existing flow when the successful tests exceeds a second predetermined threshold;
- 20 dropping a incoming packet from an existing flow when the failure tests is greater than a third predetermined threshold;
- checking a flow of packets to determine whether a current rate meets a constraint of exceeding by one half a previous rate;
- incrementing fail\_num counter when the current rate exceeds the
- 25 constraint; and
- incrementing a pass\_num counter when the current rate is at least one of less than and equal to the constraint;
- admitting the incoming packet when the fail\_num counter is less than a fourth predetermined threshold; and
- 30 dropping the incoming packet when the fail\_num counter is greater than a fifth predetermined threshold
7. A system for BTM, comprising:

- means for receiving incoming packets;
- means for performing a method for traffic classification configured to determine a protocol of the incoming packets;
- means for performing a method of bandwidth division/allocation;
- 5 means for determining whether a TCP connection has been established, wherein when the TCP connection has not been established:
- means for performing at least one of rate limiting, waiting time reduction for half-open connections, and incrementing backlog queue size, and wherein when the TCP connection has been established:
- 10 means for performing a method for traffic differentiation, wherein when a determination is made for normal traffic,
- means for admitting the packet; and
- wherein a determination is made for attack traffic,
- means for dropping the packet.
- 15 8. A processor-readable medium containing software code that, when executed by a processor, causes the processor to implement a method for Behavior-based Traffic Differentiation (BTD), comprising:
- receiving incoming packets;
- performing a method for traffic classification configured to determine a
- 20 protocol of the incoming packets;
- performing a method of bandwidth division/allocation;
- determining whether a TCP connection has been established, wherein when the TCP connection has not been established:
- performing at least one of rate limiting, waiting time reduction for
- 25 half-open connections, and incrementing backlog queue size, and wherein when the TCP connection has been established:
- performing a method for traffic differentiation,
- wherein when a determination is made for normal traffic,
- admitting the packet; and
- 30 wherein a determination is made for attack traffic,
- dropping the packet.

9. The processor-readable medium according to claim 8, wherein the method for traffic classification further comprises at least checking a protocol field at an Internet Protocol (IP) layer for the incoming packet to determine the packet classification.
- 5 10. The processor-readable medium according to claim 9, wherein the checked protocol field determines whether the packet is classified as one of TCP, UDP, ICMP and OTHER packets.
11. The processor-readable medium according to claim 10, wherein when the checked protocol field is at least one of UDP, ICMP and OTHER packets, a  
10 method of bandwidth division/allocation further comprises determining the bandwidth allocated for the incoming packet; and, classifying the incoming packet as one of UDP, ICMP and OTHER based on the determined bandwidth.
12. The processor-readable medium according to claim 11, wherein them method of traffic classification is implemented at a firewall or edge router.
- 15 13. The processor-readable medium of claim 12, wherein the method of traffic differentiation further comprises:
- checking a source IP address, source port number, destination IP address, and destination port number of an incoming packet;
  - determining whether the incoming packet is a first packet of a new flow or  
20 one of a plurality of admitted flows,
  - determining when the number of admitted flows reaches a maximal flow count relative to a predetermined first threshold set to ensure provisioning of proper quality of service (QoS);
  - updating a maximal count flow table by incrementing the flow count by 1;
  - 25 incrementing counters for successful tests and failure tests when maximal flow count is at least one of below and equal to the predetermined first threshold and admitting the incoming packet of the new flow;
  - dropping the incoming packet when the maximal flow count exceeds the predetermined threshold;
  - 30 admitting a incoming packet from an existing flow when the successful tests exceeds a second predetermined threshold;

- dropping a incoming packet from an existing flow when the failure tests is greater than a third predetermined threshold;
- checking a flow of packets to determine whether a current rate meets a constraint of exceeding by one half a previous rate;
- 5        incrementing fail\_num counter when the current rate exceeds the constraint; and
- incrementing a pass\_num counter when the current rate is at least one of less than and equal to the constraint;
- admitting the incoming packet when the fail\_num counter is less than a
- 10        fourth predetermined threshold; and
- dropping the incoming packet when the fail\_num counter is greater than a fifth predetermined threshold




FIG.1

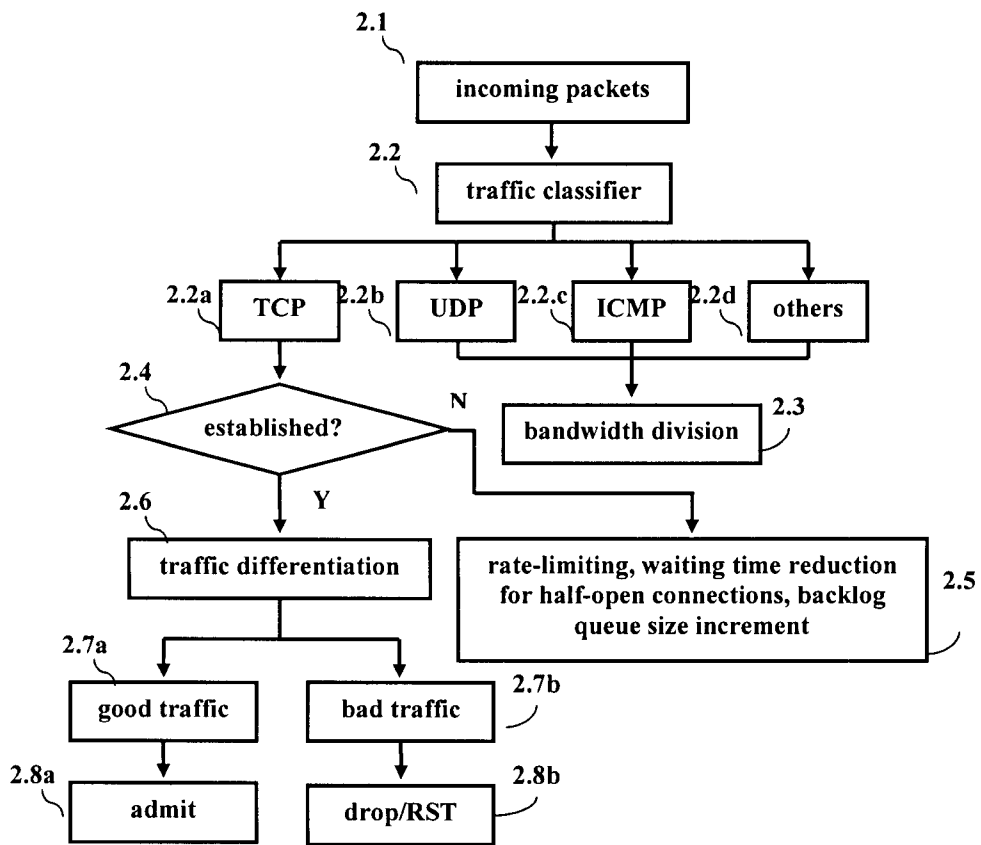


FIG. 2




FIG. 3

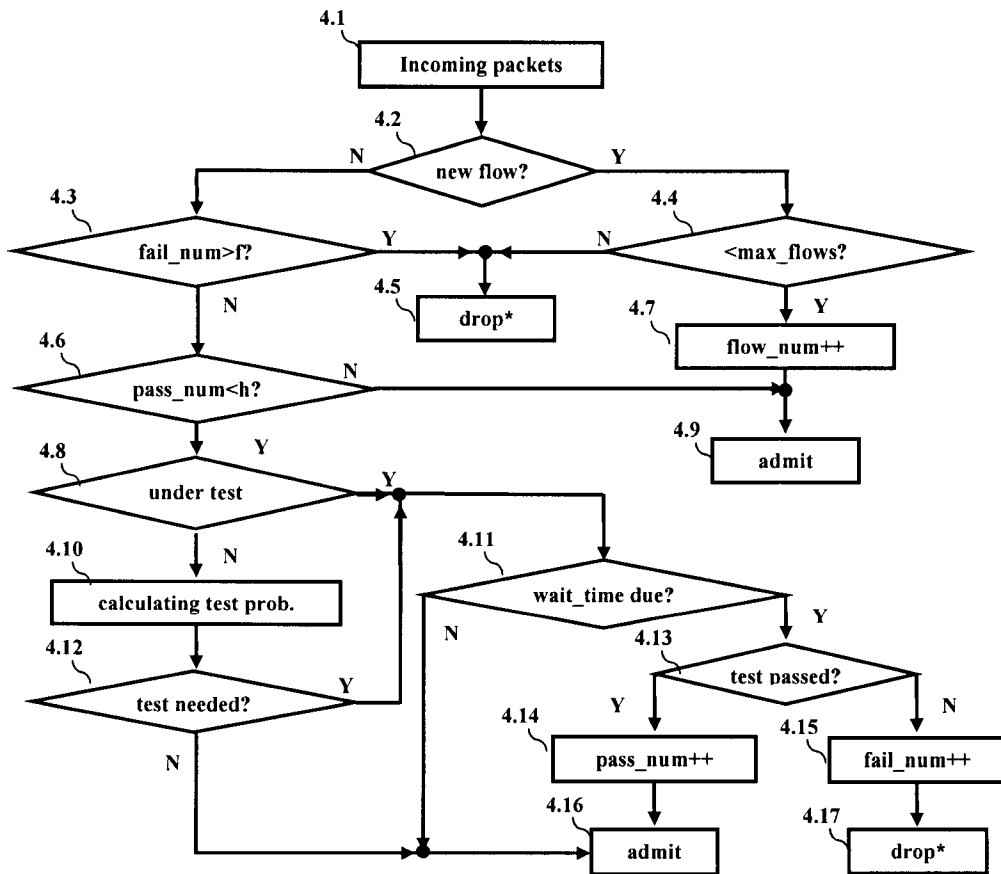


FIG. 4

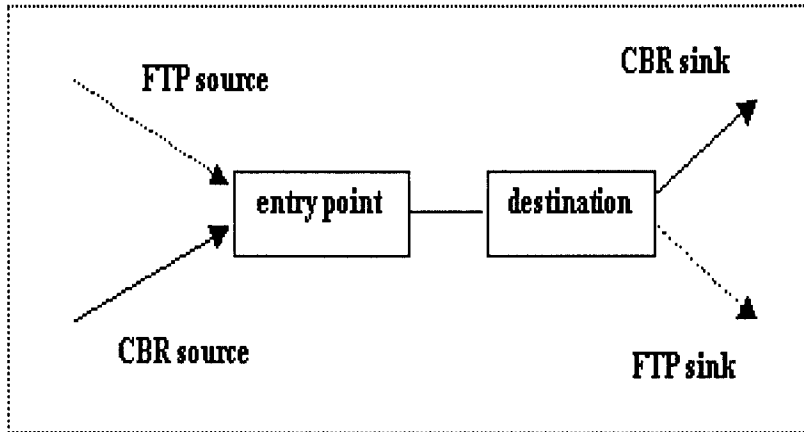


FIG. 5




FIG. 6(a)




FIG. 6(b)




FIG. 7




FIG. 8(a)




FIG. 8(b)