Adaptive Response System for Distributed Denial-of-Service Attacks

Vrizlynn L. L. Thing, Morris Sloman and Naranker Dulay
 vriz@i2r.a-star.edu.sg, mss@doc.ic.ac.uk and nd@doc.ic.ac.uk
 Institute for Infocomm Research and Imperial College London

Abstract—This dissertation presents a **Distributed denial-of-service Adaptive Response** (DARE) system, capable of executing appropriate detection and mitigation responses automatically and adaptively according to the attacks. It supports easy integration of distributed modules for both signature-based and anomaly-based detection. Additionally, the innovative design of DARE's individual components takes into consideration the strengths and weaknesses of existing defence mechanisms, and the characteristics and possible future mutations of DDoS attacks. The distributed components work together interactively to adapt detection and response according to the attack types. Experiments on DARE show that the attack detection and mitigation were successfully completed within seconds, with about 60% to 86% of the attack traffic being dropped, while availability for legitimate and new legitimate requests was maintained. DARE is able to detect and trigger appropriate responses in accordance to the attacks being launched with high accuracy, effectiveness and efficiency. The dissertation is available at <http://pubs.doc.ic.ac.uk/VrizlynnThing-PhD-Thesis-2008/VrizlynnThing-PhD-Thesis-2008.pdf>.

Index Terms—Distributed Denial of Service, Adaptive Response System.

I. INTRODUCTION

IN recent years, high-profile DDoS attacks over the Internet have increasingly been exploited by extortionists, business rivals and rebels on websites of banking and financial companies, online gambling firms, web retailers, government sectors, and even firms providing Internet security services. The following recent studies and surveys conducted, show the prevalence of DDoS attacks in the Internet.

- 68,700 attacks against over 34,700 distinct Internet host targets over a period of 3 years from 2001 to 2004 were observed in the backscatter study.
- Symantec witnessed an average of 5,213 DDoS attacks per day between 1st July and 31st December 2006.
- 70 worldwide ISPs ranked the millions of zombie computers across the Internet as the single largest threat facing network services availability and operational security, from July 2006 to June 2007. An increase in the types of activities for which the botnets were employed for, such as launching DDoS attacks, was observed.
- Cisco Security IntelliShield Alert Manager Service issued the most alerts for DoS, buffer overflow and arbitrary code execution threats from January to October 2007.

We studied and analysed existing DDoS attack tools used by attackers to gain an in-depth understanding of how DDoS attacks are mounted and their characteristics. An evaluation

of existing DDoS attack prevention, detection and response mechanisms was also performed to determine how available techniques handle DDoS attacks and how well they are able to cope with current DDoS attacks as well as their potential mutations.

We built a flexible Adaptive Response System capable of executing appropriate detection and mitigation responses automatically to adapt to the changing conditions of the victim and its network during an attack. The system overcomes the weaknesses of existing mechanisms for handling current and future variants of DDoS attacks (for example, tweaking attack characteristics to make attack profiling difficult for defence mechanisms, spoofing of source addresses, overloading of servers with seemingly legitimate requests). The Adaptive Response System should:

- incorporate flexibility for easy integration of future components in order to evolve to cater for future changes in attack objectives and characteristics
- be distributed and scalable to handle the distributed nature of DDoS attacks
- require no change to the Internet infrastructure which raise conformance issues
- constitute simple and fast modules with minimal processing overhead for detecting and responding to attacks efficiently
- be implemented as independent management components which can be deployed in protected computers to reduce the threat from attack on the Adaptive Response System itself.

The dissertation also presents another DDoS mitigation system, Traffic Redirection Attack Protection System (TRAPS) [1], designed for the IPv6 networks. We show how TRAPS is able to provide server protection against DDoS attacks through virtual relocation using the existing Mobile IPv6 protocol. Therefore, no change is required to the end-hosts which are potentially all the Internet nodes.

DARE was implemented and deployed on Emulab to conduct large-scale experiments to test the detection and responses, in terms of correctness in executions, results accuracy, mitigation speed and efficiency, and protection of legitimate traffic (ongoing and new).

Section II describes the background work and gives an overview of the work in the dissertation. Section III covers our contributions. Section IV describes the DARE architecture and modules. Experiments conducted are presented in Section

V, followed by the description of TRAPS in Section VI and the conclusion is in Section VII.

II. BACKGROUND WORK AND DISSERTATION OVERVIEW

In this section, we provide an insight to the analysis we have conducted on initial background work in the area of DDoS attacks and mitigation. We also explain the overview of our research work.

A. Related Work

We classified DDoS mitigation into three categories, namely prevention, detection and responses, and studied existing techniques in each category. Although prevention is an important first line of defence against attacks, we observed that determined attackers will always try to work around the preventive measures. New attacks will also compromise preventive measures which have not been prepared to cope with them as indicated by the current prevalence of DDoS attacks. Thus it is important to be able to detect and respond to DDoS attacks. Detection includes signature and anomaly based techniques. Different types of attacks requires different detection methods to increase true positives and achieve minimal false negatives, in particular when selecting the detection parameters, thresholds and normal profiles for anomaly-based detection methods which have inherently lower reliability than signature-based ones.

After an attack has been detected, an appropriate response to handle the attack should be triggered. This could include traceback to locate the actual source of the attack traffic, especially in the case of address spoofing, so that further responses such as filtering or rate-limiting could be performed nearer to the source to reduce attack traffic propagation. Different response types also vary in the degree to which they filter off potential attack traffic corresponding to the reliability of the detection mechanism and the confidence in differentiating attack traffic from the legitimate ones.

We also studied existing automated mitigation frameworks. Among them, EMERALD and COSSACK did not provide an open architecture, were no longer maintained and supported limited detection and response modules. DiDDeM supports congestion detection but relies on signature-based response only. SNORT and Prelude allow plugins and sensors integration, respectively. However, SNORT and Prelude mainly support rule-based or signature-based detection, but not anomaly-based detection modules.

B. DDoS Attack Bots

In addition, we studied in details the source code of DDoS attack tools present in four popular bots used by attackers to launch DDoS attacks [2]. The source code was obtained from hacker web sites. Such analysis gave us an in-depth understanding of the attacks' features and characteristics.

We analysed the tools and discussed the implications of our findings on well-known DDoS mitigation techniques. For example, the partial subnet address spoofing for the attack packets, which retains the source address network prefix,

would compromise ingress filtering. The popular SYN-Dog detection mechanism would not work correctly as it is possible for attackers to send out randomised SYN and ACK packets, imitating the three-way handshake. We also emphasised the need to acquire an understanding of the attacks before being able to design more effective and efficient mitigation techniques.

From our study of the DDoS attack bots, we observed that the most common attacks present in the bots are TCP SYN and UDP flooding attacks. This observation coincides with the study by the Arbor Networks Inc. where the surveyed ISPs reported that these two attack types were the most prominent. Therefore, we concentrated on the detection and response to these attacks, as well as considered the effect of their possible mutations on mitigation techniques.

C. Analysis of Appropriate Responses

We classified DDoS responses as Traceback, Containment, Reconfiguration, Redirection, Filtering, Ratelimiting, Resource replication, Legitimacy testing, and Attackers' resource consumption. As the system performance relies on the individual components, special care was put into choosing the responses to design. We analysed and discussed when a response is appropriate and the reasons to carry out specific actions under each response type. We considered and decided on the responses necessary to effectively mitigate DDoS TCP SYN and UDP flooding attacks, as follow.

Traceback is essential in any DDoS defence system to locate the point nearest to the attack sources to perform mitigation so as to prevent attack traffic from propagating to other networks (including intermediaries). Existing mechanisms require wide spread deployment on Internet routers and the provision of an attack signature for identifying attack packets to extract traceback information. We implemented a novel traceback mechanism in DARE, Non Intrusive IP Traceback [3], [4] (in Section IV), which resolves the above-mentioned problems.

Reconfiguration techniques work within the target network to provide a resilient network in the event of an attack, by re-routing legitimate traffic and isolating attack traffic. However, existing methods such as the Secure Overlay Service only works with previously authenticated clients, while [5] only works with HTTP traffic. Therefore, we designed a new reconfiguration mechanism called the Traffic Redirection Attack Protection System (TRAPS), a standalone system, for IPv6 networks, in Section VI.

In our system, we consider Redirection as Filtering. Filtering and Ratelimiting are implemented as the response modules for DARE. However, to ensure correct filtering, we devised a Bloom-based filter which relies on our Enhanced TCP SYN Attack detection [6] (in Section IV). Ratelimiting is also performed only on identified suspicious and attack flows, detected by our DDoS Attack Detector and Flow Identifier (in Section IV).

III. CONTRIBUTIONS OF DISSERTATION

The dissertation details our attacks analysis, and design and experiments of our detection and response modules in DARE

and TRAPS. The following outlines the contribution of the dissertation:

- A DDoS Bot Survey to analyse the source code of the four most popular botnet tools typically used for carrying out DDoS attacks. We identified and analysed their characteristics to cater for better understanding of current DDoS attacks. We proposed enhancements and mutations, and implemented more powerful attack tools to conduct tests on our new mitigation mechanisms. This work is to the best of our knowledge, the first published study on DDoS Bot attack tools, and will be a useful resource for the research community working on DDoS attack mitigation.
- A DDoS Adaptive ResponseE (DARE) System is designed and built to achieve the goal of establishing a tight loop of modular detection, decision making and response components. DARE composes innovative detection and response mechanisms in the form of a bandwidth monitoring module, a traceback module, a TCP SYN detection module, an adaptive system manager for coordination of various components, a filtering module and a ratelimiting module. DARE can work across multiple, distributed networking devices. The modular approach also allows for easy modification, tailoring and adaptation of the detection, decision making and response techniques. DARE stands out from existing frameworks by allowing the integration of not only signature-based but also anomaly-based detection and responses.
- A newly Enhanced TCP SYN Attack Detection and Bloom-based Filtering Mechanism to cater for the potential “mutation” of TCP SYN attacks, which would not be detected by current techniques, is designed. The mechanism allows zero false positive rate of filtering, while protecting ongoing and new legitimate traffic during attack occurrence.
- An enhanced DDoS Flood Attack Detector and Attack Flow Identification Mechanism is presented based on MULTOPS with enhancements to allow for a speedier attack detection and a more efficient memory allocation for the hierarchical data structure it uses. Legitimate traffic classified into different flows is spared from the ratelimiting response.
- A novel Non Intrusive IP Traceback technique to identify actual paths taken by attack traffic by detecting route anomalies is built and evaluated. It is used to facilitate tracing of attackers and to carry out further responses on attack traffic as close as possible to the source in order to limit attack traffic within intermediate networks. The technique is non intrusive as no change is required to the Internet routers unlike existing techniques.
- A Traffic Redirection Attack Protection System (TRAPS) aimed at next generation IPv6 networks is described. TRAPS verifies the authenticity of source addresses in order to detect address spoofing during a DDoS attack. It then performs adaptive reconfigurations to differentiate legitimate clients and attackers in order to carry out mitigation by filtering out attack traffic. It is particularly

useful in IPv6 networks where Mobile IPv6 and its route optimisation feature are built-in. Therefore, no deployment is required to the end hosts, which are potentially all the nodes in the Internet.

In the next section, we describe the DARE architecture.

IV. DARE ARCHITECTURE AND MODULES

In this section, we present the architecture overview of DARE which is based on the architecture of the EU Diadem Firewall project in which we were involved.

A. Diadem Firewall

The Diadem Firewall project [7] was an European Union funded project to develop an architecture that enables an Internet Service Provider (ISP) to protect its own networking environment as well as the connected hosts and servers of its customer against network attacks. We were one of the seven European project partners, and were responsible for the System Manager (SM) and the Non Intrusive IP Traceback module.

Fig. 1: Architecture Overview of Diadem

Figure 1 shows an architecture overview of the Diadem Firewall. It consists of distributed monitoring elements running the Versatile monitoring toolkit (VERMONT), which are network monitors configured to capture and/or aggregate traffic packets. The traffic packets or flow information is then exported to the Traffic Flow and Packet Analysis System (TOPAS), a framework for the reception and real-time analysis of the network packets and flow information. TOPAS subscribes to flow information essential for its modules to carry out analysis to detect various attacks. The existing detection modules in Diadem consist of the Web Server attack detector and the TCP SYN attack detector running in parallel within TOPAS.

The Web Server attack detector builds normal user behaviour models by monitoring service requests for the server objects (for example, web pages). A change-point detection algorithm checks for changes in the browsing behaviour of the users, to detect Web server overloading attacks. The TCP SYN attack detector uses the SYN-Dog (SYN-SYN&ACK) algorithm. VERMONTs at the source networks send information on outgoing SYN and incoming SYN&ACK packets to the TCP SYN attack detector for analysis. In a TCP SYN attack, the difference between the number of the outgoing SYN and incoming SYN&ACK packets will be very noticeable.

The Non-Intrusive IP Traceback module was designed and developed as a TOPAS response module. It is triggered after

an attack is detected, to identify the closest locations from the attack sources. The module depends on flow information received from VERMONTs to check for route anomalies in the source and destination IP address pairs. Details on the design are in Section V.

The external response module is a Firewall, complemented with a high-speed traffic classification engine implemented on a PCI FPGA board. When a TCP SYN attack is detected, the firewall is triggered to drop requests from sources which have not been seen recently. When a Web Server overloading attack is detected, the firewall is triggered to perform ratelimiting on suspicious traffic not matching the normal user behaviour models.

The System Manager (SM) is the system coordinator. Alert messages regarding attacks are sent to the SM, which then dispatches appropriate responses to the response modules, such as the Traceback module and/or distributed firewall elements. Results and status information from the response modules are also sent to the SM for attack status logging.

As seen in Figure 1, two distinct types of inner communication paths exist. The first path type allows the monitoring data from VERMONTs to be sent to TOPAS. The second path type allows the communications of the IDMEF (alert and response) event messages between the SM, detection modules and external response modules (residing outside TOPAS). IDMEF defines an XML message format for intrusion alerts. The event distribution between the modules is based on the XMLBlaster publish-subscribe system, where modules subscribe to receive events they are interested in (based on event topics) from a server. Modules publish events they generate to the XMLBlaster server which in turn forwards them to those modules that have subscribed to the type of events. The use of the XMLBlaster for event distribution allows Diadem to scale easily when additional detection and response modules, or even SMs, are added.

The goal of the architecture is to establish a loop of detection, decision and response actions across multiple distributed firewall and networking devices, so as to provide a comprehensive security solution. The architecture also allows development and deployment of modular network components to be integrated in an efficient manner. In this way, Diadem ensured functional flexibility, as well as scalability of the system. Additionally, it provides built-in support for traffic flow sampling, aggregating and exporting through VERMONT. With its open architecture, it also supports the design, development and integration of both signature-based and complex anomaly-based detection techniques for DDoS attacks mitigation, as well as automated response modules, into TOPAS.

Due to Diadem's architectural advantages and modular approach, as well as our knowledge of the system, it was chosen as the basis for the DARE implementation. However, several modifications and novel additions were designed and implemented within DARE to improve the detection and response mechanisms in terms of speed and mitigation efficiencies. We adopt the Diadem architectural structure, VERMONT for network traffic monitoring and TOPAS for the integration of our detection and response modules.

B. DARE

Fig. 2: Architecture Overview of DARE

We designed and implemented novel detection modules for integration into TOPAS, and external response modules which cooperate better with our new detection mechanisms. The architecture overview of DARE is shown in Figure 2 to illustrate the layout of the entire system. The sub-systems and modules that exist in Diadem are in grey, while the new ones in DARE are in black.

1) *Adaptive System Manager*: We built a new Adaptive System Manager (ASM) for coordinating the events from the detection modules and response triggering. The SM in Diadem was written in Java and it supports the alert and response event coordination between the existing modules in Diadem. The new ASM was implemented in C++ to improve the performance speed and to support the coordination between the new modules, such as triggering further responses like traceback, triggering filtering of TCP SYN attack traffic at the exact points returned by the traceback module and triggering the ratelimiters to probabilistically drop the identified suspicious/attack traffic flow.

2) *XML Subscriber and Parser*: The XML Subscriber and Parser module is created in DARE as TOPAS only comes with a built-in XML Publisher to send out IDMEF alerts. Previously in Diadem (refer to Figure 1), IDMEF alerts received by TOPAS catered only for the Non-Intrusive IP Traceback through the use of a flag in a file (read at preset interval) to trigger the traceback process. Therefore, without a subscriber and parser module, IDMEF messages could not be subscribed to by the detection and response modules within TOPAS to get loopback alert information, update messages and system configuration messages. We solved this problem in DARE by introducing our XML Subscriber and Parser utility module, which is “shared” by all the modules in TOPAS.

3) *Non Intrusive IP Traceback*: The Non-Intrusive IP Traceback is designed based on the rationale that packets belonging to a particular source-destination pair follow a relatively static path as routing tables are not updated very frequently under normal conditions. When an attacker spoofs a legitimate user's source address, the packet may pass through routers which are not on the normal source-destination routing path and this anomaly can be used to determine the attack path. Compared to existing mechanisms, our traceback scheme requires no precise attack signature from the detection modules, no modification to the Internet infrastructure and no

overwriting of data packet fields. It relieves the victim from the additional burden of the attack path construction, and is simple and fast. We modified the code of our traceback module in DARE to process the information forwarded to it through the new parser utility module.

4) DDoS Flooding Attack Detector and Flow Identifier:

The Flooding Attack detector facilitates the detection of the prominent DDoS UDP flooding attacks. The design of our detector and the Flow Identifier is based on MULTOPS but includes our new enhancements. Our module's data structure is organised as a 4-level 256-ary tree to cover the entire IPv4 address space. Each node in the tree has 2 fields: incoming packet rate and a pointer to the node in the next level of the tree. The IP addresses have a common 0-bit, 8-bit, 16-bit, or 24-bit prefix, depending on the level of the tree. The root nodes will contain the aggregate packet rates from address 0.*.*, 1.*.*, etc. If the incoming packet rate exceeds a preset threshold for the current tree level, a new subnode is created on the fly. This form of subnode creation is termed expansion. When the packet rate drops below a certain threshold, removing nodes or the entire subtree is performed and it is termed contraction. Therefore, we are able to detect flooding attack based on the unusually high incoming packet rate of particular flows. This allows the isolation of suspicious/attack traffic flows. As this detection method does not achieve zero false positive, ratelimiting is triggered as an appropriate response on the identified flows. Our design allows a faster flow identification than MULTOPS by reversing the expansion and contraction order to cater for DDoS attack detection. Due to the dynamic memory allocation for our data structure and the logging of only the incoming traffic, we are able to achieve in excess of over 50% of memory conservation, compared to MULTOPS.

5) *Enhanced TCP SYN Attack Detector and Bloom-based Filter:* While studying the DDoS attack bot tools and the existing TCP SYN detection mechanisms, we devised a new variant of TCP SYN attack, which we named the Bot Buddy attack. This attack uses a peer bot (buddy) to send a SYN&ACK packet to an attack bot, for each of the SYN packet the latter sends to the victim. It circumvents the popular SYN-Dog mechanism, which detects the abrupt rise in the difference between outgoing SYN and incoming SYN&ACK packets.

The Enhanced TCP SYN detector is a novel mechanism, which incorporates a Bloom-based hashing technique in the SYN-Dog detector to allow the detection of both conventional and Bot Buddy attacks. Instead of counting all outgoing SYN and incoming SYN&ACK packets, it tracks outgoing SYN packets and only counts valid corresponding incoming SYN&ACK packets. Storage of source-destination address pairs is not required due to the use of Bloom-based hashing on the addresses. Therefore, space and time efficiency is achieved. The number of incorrectly validated SYN&ACK packets is found to be minimal (i.e. < 8 packets after 20 minutes of attack and 72000 attack packets).

With the Bloom-based addresses hashing entries, it allows us to differentiate between legitimate and attack traffic (e.g. entries with only outgoing SYN packet are marked as '1', entries with both outgoing SYN and incoming SYN&ACK

packets are marked as '2'). Therefore, we incorporate a more effective filter based on these entries to drop attack packets while protecting legitimate ones.

The reason the Bloom-based filter is at the source network is due to our Enhanced TCP SYN attack detector being a source network based detection module, so TCP SYN attack packets can be isolated and dropped at the point nearest to the attacker once the Traceback module identifies the locations of the attack sources.

6) *Ratelimiter:* Our Ratelimiter module developed in C++, incorporates the Linux Traffic Control (TC) module. Information on the attack flows identified by the Flooding Detector and Flow Identifier is subscribed and parsed by our Ratelimiter. The TC rules are then formulated and executed to carry out ratelimiting. The Ratelimiter is deployed in the protected target network, where aggregated flooding traffic is detected. As flooding attack packets are usually seemingly legitimate packets, it is not feasible to ratelimit such flows at the source where traffic is minimal (traffic not converged yet). All the modules in TOPAS run in parallel and are able to detect and respond to different attacks simultaneously.

V. EXPERIMENTS

We implemented the enhanced attack tools and all the DARE modules, and evaluated the performance of DARE through experiments. We concentrated on the two most prominent DDoS attacks: TCP SYN and UDP Flooding. The system is evaluated based on the following criteria:

- Execution of the appropriate detections and responses in the situation of different attacks
- Correctness of results (that is, alert event is triggered only when there is an ongoing attack, Traceback returns correct results as to the routers forwarding attack traffic and the Enhanced TCP SYN attack detector compiles the bloom list correctly for the Bloom-based filters to forward legitimate requests and filter attack traffic)
- Speed of the attack detection and responses (and returning of the results in the case of Traceback)
- Ability to mitigate attack traffic
- Ability to protect legitimate traffic (ongoing and new)

We ported DARE to the Emulab machines to carry out the experiments. Emulab is a large-scale network testbed for testing networking and distributed systems. It allows the specifying of arbitrary network topologies, giving a controllable environment, including PC nodes with full "root" access, running an operating system of choice.

Using a comprehensive network topology and experimental scenario including a mixture of attack, legitimate and new legitimate traffic in different networks, we carried out multiple sets of experiments on TCP SYN and UDP Flooding attacks. The results from different sets of experiments were consistently close and stable. We were able to achieve correct and appropriate automated execution of the detection and responses for various attacks. The results and information returned by all the modules were verified to be accurate.

The TCP SYN attacks were detected within 6 seconds. The ASM was notified of the attacks within a second. Traceback

(and the triggering of the Bloom-based Filters at identified nodes forwarding attack traffic) was completed with the correct results returned between 2 to 6 seconds. A high percentage of the total attack traffic (86%) was filtered off. While successfully mitigating the attacks, all the legitimate and new legitimate traffic started during the attacks were fully (100%) protected.

The UDP flooding attacks were detected with the contracted attack flows identified within 7 to 27 secs. The Traceback module was also successful in returning the correct tracing results from under 1 sec to 2 secs. The Ratelimiters deployed and triggered to perform the attack traffic dropping also successfully limited the total attack traffic, resulting in approximately 60% of the attack packets being dropped.

We compared the TCP SYN attack experiment carried out in the Diadem Firewall project to DARE. Even with a much more complicated experimental scenario, DARE was able to achieve a faster reaction time in all aspects of detection, traceback and response triggering.

VI. TRAPS

Traffic Redirection Attack Protection System (TRAPS) was invented and implemented as a stand-alone DDoS attack detection and mitigation system particularly suited to IPv6 networks. In TRAPS, the victim under attack verifies the authenticity of the source by performing virtual relocations (i.e. changing of IP address). Corresponding nodes are then notified of this new “location” of the victim. Attackers spoofing their source IP addresses would not be able to receive this notification. Attack traffic can then be easily identified and isolated, with absolute confidence and be dropped, while legitimate traffic is allowed through. An advantage of TRAPS has over existing mitigation frameworks, is that the processing overhead is low due to the simplicity of the approach. It also requires minimal deployment effort and does not require modification to the Internet infrastructure. By using Mobile IPv6, we do not need special software at the correspondents (Internet end hosts). Experiments to test the feasibility of TRAPS were carried out in a testbed environment to verify that it would work with the existing Mobile IPv6 implementation. It was observed that the operations of each module were functioning correctly and TRAPS was able to successfully mitigate an attack launched with spoofed source IP addresses, while maintaining the core Mobile IPv6 functionality.

VII. CONCLUSION

Our analysis of the DDoS attack tools provided a useful resource for understanding how the code was structured and what design decisions were made by the Botnet authors. Based on it, we pointed out modifications which could be made to strengthen attacks and ensured that enhanced forms of the attacks were used in our experiments. There is a dilemma in suggesting and publishing suggested new variations and “improvements” of attacks. However, our intention is to inform the research community of the limitations of existing techniques and we have provided techniques for protection from these new types of attacks.

We have designed and built DARE, which is a flexible Adaptive Response System, capable of executing appropriate detection and mitigation responses automatically according to the observed changing conditions of the victim and its network during an attack. It overcomes the weaknesses of existing mechanisms to handle current and future variants of DDoS attacks, incorporating flexibility for easy integration of future components in order to evolve to cater for future changes in attack objectives and characteristics, and is distributed and scalable to handle the distributed nature of DDoS attacks. Large scale experiments were also conducted to demonstrate its high reliability, effectiveness and speed of detection and responses.

We also looked into the protection of next generation IPv6 networks from DDoS attacks. We proposed TRAPS which can be easily deployed, utilising the built-in Mobile IPv6 feature, to verify the authenticity of the source by performing virtual relocations of the server. Therefore, attack traffic can be correctly isolated and filtered off, while protecting legitimate traffic. TRAPS was implemented and experiments were conducted to prove its feasibility.

We have submitted a journal paper on the Non Intrusive IP Traceback (under review), while work is ongoing on journal papers submissions detailing analysis and experiments on the other specific modules in DARE. The Non Intrusive IP Traceback, and the Enhanced TCP SYN Attack Detector and Bloom-based Filter have also been submitted to the Tan Kah Kee Young Inventors Award Competition (under review).

ACKNOWLEDGEMENT

We gratefully acknowledge the support from the EU funded DIADEM Distributed Firewall FP6 IST-2002-002154. This research continued through participation in the International Technology Alliance sponsored by the U.S. Army Research Laboratory and the U.K. Ministry of Defence. We would also like to thank Emulab for providing the facility for the experiments.

REFERENCES

- [1] V. L. L. Thing, H. C. J. Lee, and M. Sloman, “Traffic redirection attack protection system (TRAPS),” in *IFIP International Information Security Conference (SEC)*, (Makuhari-Messe, Chiba, Japan), Springer-Kluwer, May 2005.
- [2] V. L. L. Thing, M. Sloman, and N. Dulay, “A survey of bots used for distributed denial of service attacks,” in *22nd IFIP International Information Security Conference (SEC)*, (Sandton, Gauteng, South Africa), May 2007.
- [3] V. L. L. Thing, M. Sloman, and N. Dulay, “Non-intrusive IP traceback for DDoS attacks,” in *ACM Symposium on Information, Computer and Communications Security*, (Singapore), Mar. 2007.
- [4] V. L. L. Thing, M. Sloman, and N. Dulay, “Network domain entry-point/path determination for DDoS attacks,” in *IEEE/IFIP Network Operations and Management Symposium (NOMS)*, (Salvador, Bahia, Brazil), Apr. 2008.
- [5] J. Xu and W. Lee, “Sustaining availability of web services under distributed denial of service attacks,” *IEEE Transactions on Computers*, vol. 52, no. 2, pp. 195–208, Feb. 2003.
- [6] V. L. L. Thing, M. Sloman, and N. Dulay, “Enhanced TCP SYN attack detection,” in *IEEE/IST Workshop on Monitoring, Attack Detection and Mitigation (MonAM)*, (Toulouse, France), Nov. 2007.
- [7] G. Munz, A. Fessi, G. Carle, O. Paul, D. Gabrijelcic, Y. Carlinet, S. Yusuf, M. Sloman, V. L. L. Thing, J. v. Lunteren, P. Sagmeister, and G. Dittmann, “Diadem firewall: Web server overload attack detection and response,” in *Broadband Europe (BBEurope)*, (Bordeaux, France), Dec. 2005.