

Design and Analysis of FPGA Architecture for High-Speed Network Feature Extraction

Mrs.C.Nirmala¹, K.Deekshitha²

*1 Assistant Professor, Department of ECE, Malla Reddy College of Engineering for Women.,
Maisammaguda., Medchal., TS, India
2, B.Tech ECE (19RG1A04E6),
Malla Reddy College of Engineering for Women., Maisammaguda., Medchal., TS, India*

Abstract

System include extraction includes the capacity and characterization of system parcel action. Albeit fundamentally utilized in organize interruption discovery frameworks, include extraction is likewise used to decide different angles of a system's conduct, for example, all out trafç and normal association size. Current programming techniques utilized for extraction of system highlights neglect to meet the exhibition necessities of cutting edge fast systems. Right now paper, we propose a FPGA-based reconfigurable engineering for highlight extraction of enormous fast systems. Our structure utilizes equal lines of hash capacities and sketch tables so as to process arrange bundles at a high throughput. We present a definite portrayal of our engineering and its execution on a Xilinx Virtex-II Pro FPGA board, and give cycle-precise planning results to highlight extraction of information organizing benchmark information. Our outcomes exhibit certifiable throughputs of as high as 3.32 Gbps, with speedups arriving at 18 x when contrasted with a comparable programming execution.

1 Introduction

The goal of a Network Intrusion Detection System (NIDS) is to detect attacks on any machine within the localnetwork by monitoring the network activity. In general, there are two different approaches taken when protecting networks using intrusion detection-based systems. The first approach, known as signature detection, searches for predetermined attack patterns in the network activity. The second approach, known as anomaly detection, looks for any sort of abnormal activity in the network flow, and then determines if the abnormal behavior is an attack. Signature detection-based methods are unable to detect new kinds of attacks, as well as those attacks that vary significantly from . This observation has

motivated a deeper study of anomaly-based NIDS mechanisms.

Anomaly detection typically involves two separate stages. In the first step, network features stored in packet headers are extracted and stored over an interval of time. In the second step, a change detection and classification algorithm is applied to this stored information in order to detect attacks. In large-scale high-speed networks, this first step in anomaly detection is the most crucial. An efficient NIDS must be able to store and classify network features without compromising on speed or loss of information. Considering the increasing size and speed of modem networks, general-purpose processors do not meet the requirements of the next generation of NIDSs. This has motivated researchers to explore the possibility of using dedicated hardware for anomaly detection systems in general [2, 3, 15] and feature extraction/classification in particular [10, 16]. Besides its use in anomaly detection, feature extraction is key to several other applications such as data mining [1], speech recognition [12], and image processing [11], among others. However, due to the clear needs of performance, we concentrate on using feature extraction for anomaly detection only. Particularly, we propose a reconfigurable architecture for feature extraction of high-speed networks, and implement this design using FPGAs. By making use of the inherent parallelism of FPGA hardware, we are able to speed up our application by a considerable amount as compared to an equivalent software implementation. Our architecture is pipelined to achieve a high throughput, making it suitable for application in multi-gigabit networks. We also make use of feature sketches to store network activity, thus minimizing the required memory resources. Our results show that the architecture is several times faster than

the equivalent software implementation and offers a practical solution for feature extraction of high-speed networks.

The remainder of this paper is organized as follows. Section 2 provides a brief overview of some of the main concepts behind network intrusion detection. Section 3 motivates the need for hardware implementation of the feature extraction process. In Section 4, we describe our Feature Extraction Module (FEM) architecture and its various building blocks. Section 5 presents an example NIDS implementation, demonstrating how the FEM module can be used to detect certain types of attacks. Section 6 provides details of our implementation and presents our area and performance results. Related work is discussed in Section 7, followed by the conclusion in Section 8.

2 Intrusion Detection Overview

There are a wide variety of attacks prevalent in modern networks. Detection of each type of attack requires the monitoring of different network features. For our implementation of a FEM-based anomaly detection architecture, we focus on two major categories of intrusions common to modern day networks: Denial-of-Service (DoS) attacks and port scanning attacks. There are many different variations of DoS attacks including the SYN flood, spoofing, smurf attack, fraggle attack, and distributed DoS attack (DDoS). Although our architecture can be configured to extract features for detection of any type of the above mentioned attacks, in our study we focus only on SYN floods and DDoS attacks in a TCP/IP network. Since the implementation of FEM depends on the type of attack, we provide a description of the various attacks and the corresponding features needed to detect them.

The understanding of these types of attacks requires some basic knowledge of the 3-way handshake by which TCP connections are established. First, the source requests a connection by sending a packet with the SYN flag set to the host computer. The host then responds with a packet having both the SYN and ACK flags set. To complete this half-open connection, the source responds with a packet having the ACK flag set. Typically the host waits for a certain duration of time for the last acknowledgement from the source, after which it closes the half-open connection and the entire process has to be repeated for a new connection. A host computer in a network can handle only a finite number of connections at a time.

The SYNflood is a very common example of a DoS attack. In the SYN flood attack, the source (attacker) sends a series of connection requests in a short duration of time to the host (victim), filling up all of its connection handling resources. When the host responds with both the SYN and ACK flags set, the source purposefully skips sending the last ACK flag to the host. Consequently the host cannot accept any requests for a new connection as all of its connection handling resources are consumed by the half-open connections. In another variation of SYN flood, the source spoofs the IP address in the packet which it sends to the host to request a connection. The host responds with a SYN/ACK packet to the modified IP address which will either be non-existent or will be ignorant of the requested connection. As a result the 3-way handshake is never completed, resulting in multiple half-open connections at the host. In either case, if the host computer happens to be a server, then the entire network is affected as it becomes unable to accept any new connection requests.

In the DDoS attack, an attacker first gains control of several computers in the network in order to attack a server using multiple parallel DoS attacks (like the SYN flood). This attack is more difficult to detect than the standard SYN flood since it originates from multiple sources, none of which is the real source of the attack. In the port scanning attack, the attacker scans for open ports on different computers on the network. During this scan the attacker sends a connection request to that particular port and determines whether the port is open from its response. Similar to the SYN flood attack, in this attack the 3-way handshake never occurs and usually ends in the second step. Since the host port's response must return to the source, the attacker cannot spoof its IP address in the initial connection request packet. Consequently, unlike in the case of a DoS attack, a port scan attacker can be traced back.

3 Application Analysis

The process of extracting features from network packets consists of two stages. The first stage involves storing the information associated with specified fields in the packet header. The fields that are stored are directly determined by the requested features. The values of these features are then computed as a function of the stored information. This is followed by a second stage in which the features are classified in order to determine relevant information about the network activity. As applied in a NIDS, this second stage uses these extracted features to detect the occurrence of an attack from

outside the network. A critical aspect of feature extraction is the speed at which it is performed. A feature extraction system that is unable to cope with the throughput of the network packet flow will result in a loss of features in some packets, thereby reducing the clarity of the overall network picture. Another interesting metric is the number of features being extracted.

Extracting more features can give a more accurate picture of the network activity. With the emergence of multi-gigabit networks, it is highly essential for the first stage in a feature extraction system to have a very high throughput. Unlike in signature detection-based approaches, anomaly detection utilizes the network features gathered over a period of time to look for the possibility of occurrence of any attack. In other words, features need not be sent to the anomaly detection.




Figure 1. Feature sketch architecture

module for each individual packet. Consequently, the second stage of the feature extraction system can afford to have a lower throughput, even if network speed is very high. General-purpose processors currently used in the first stage of an anomaly-based NIDS cannot meet the requirements of high-speed networks for several reasons. First, extracting multiple features from a single packet involves cycling the same data through the processor multiple times, reducing the effective throughput. Also, computation of complex features may involve extensive computation within the processor, again reducing the overall throughput given the limited computational parallelism available in general purpose processors.

In order to overcome these limitations, we propose a hardware design to implement the first stage of the feature extraction system. We incorporate both coarse and fine grained parallelism into the design in order to obtain a very high throughput.

3.Related Work

Many network applications have been implemented in hardware in order to keep up with increasing network speeds. Since the current-generation FPGAs are capable of operating at speeds up to 550Mhz, have comparable capacities to ASIC designs, and provide some flexibility in implementation, they are being actively used in developing high-speed network applications. In the most closely related work, Nguyen et al. [10] have developed a feature extraction module that utilizes multi-dimensional hashes. Although there are some similarities in the architecture, we provide a novel implementation of the underlying functions in order to obtain a higher real-world throughput while consuming significantly less area. In addition, Nguyen et al. only provide results from synthesis and do not realize their designs onto an FPGA.

FPGAs have been used in developing NIDS capable of operating up to 8 Gbps, where the network interface and intrusion detection circuitry have been integrated onto a single FPGA chip [4]. FPGAs have also been used to implement TCP/IP flow monitors [13] operating at 3 Gbps. Other network applications such as internet firewalls [9] working at 2.5 Gbps have also been implemented using FPGAs. A direct comparison of these FPGA architectures is difficult as they have different goals and are targeted towards different hardware technologies.

5. FEM Architecture

The objective of our Feature Extraction Module (FEM) is to efficiently collect all of the necessary flow information from incoming and outgoing packets in order to detect intrusions. We make use of feature sketches to store this network information. A sketch is a probabilistic approach used to summarize large amount of information given a fixed amount of memory. Researchers have shown that sketches are able to summarize large datasets (such as network activity) with a high level of accuracy [14, 15]. The FEM consists of several feature sketches connected in parallel, with each sketch customized to store a unique feature. The general architecture of our feature sketch design is shown in Figure 1. A feature sketch consists of four important blocks: the hash control, the hash function, the sketch table, and the estimate block. The only network information which the FEM needs in order to detect DoS and port scanning attacks are the source IP, destination IP, source port, destination port, and selected flags.

The multiple parallel hash functions within each feature sketch contain a combination of the input fields. The specific combination depends on the

network activity being analyzed by the FEM, and is unique for each feature sketch. The hash control block is configured as a custom input multiplexer for this purpose. For our implementation, we use Bob Jenkins' 32-bit hash function [8] because of the advantages it offers in terms of speed and fewer number of collisions.

The Jenkins hash function for hashing of three 32-bit keys (K[0], K[1], K[2]) is as shown below. The golden ratio and seeding values are random values required for the initialization of a hash function:

```
A = B = golden ratio;
C = seeding value;
A = A + K[0];
B = B + K[1];
C = C + K[2];
mix(A, B, C);
mix(A, B, C);
```

The mix function is defined as:

```
mix(a, b, c) {
a = (a - (b + c)) D (C >> 13);
b = (b - (c + a)) Q (a << 8);
c = (c - (a + b)) Q (b >> 13);
a = (a - (b + c)) Q (c >> 12);
b = (b - (c + a)) Q (a << 16);
c = (c - (a + b)) Q (b >> 5);
a = (a - (b + c)) Q (c >> 3);
b = (b - (c + a)) Q (a << 10);
c = (c - (a + b)) Q (b >> 15);
}
```

The implementation of the Jenkins hash function is done as shown in Figure 2. The initialization stage assigns appropriate values to the keys and performs a minor part of the computation of the mix function. The rest of the mix function is implemented using three mix sub-blocks as shown in Figure 2. All the mix sub-blocks are identical in structure and differ only in the number of bit shifts performed before each XOR operation. Each mix sub-block has two pipeline stages - when added to the single pipeline stage of the initialization phase, makes the hash function a thirteen stage pipelined structure. It is important to note that in this systolic style each pipelined stage has a maximum of three logical

operations, resulting in a very high clock rate and throughput.

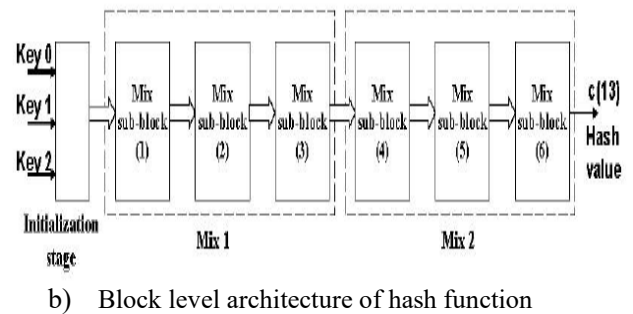
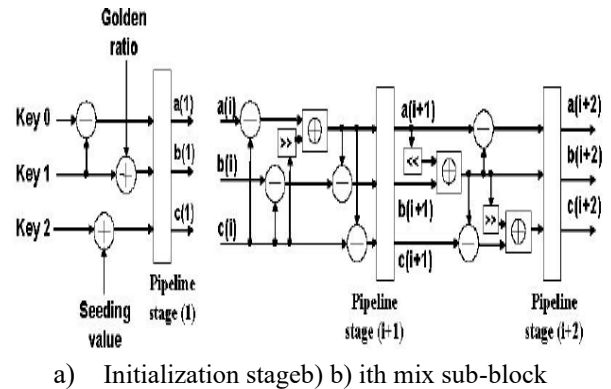


Figure 2. Hash function implementation.

In the FEM, the Jenkins' hash takes as its input the source IP, destination IP, and source/destination ports. In the configuration where any of these fields are not required by a feature sketch, they are replaced by zeroes by the hash control block. Within a single feature sketch, we use multiple hash functions with different seeding values so as to minimize the effect of hash collisions. The number of hash functions within each feature sketch can be varied depending on the accuracy requirements. The sketch table is essentially a hash lookup table which stores the network flow information contained in the flags. Its size can also be varied to reduce hash collisions, depending on the NIDS accuracy requirements.

A common characteristic of the attacks described in Section 2 is that the 3-way handshake is never fully established. This 3-way handshake involves only the SYN and ACK flags. Consequently, in a given time interval the number of incomplete connections in the network is an indication of the possibility of an attack occurring in that interval. Our FEM design applies this concept when storing the network activity. If the SYN flag is set in an incoming packet, then the value in the sketch tables corresponding to that packet is

incremented by one, and if the ACK flag is set that value is decremented by one. In other words, for every connection request the value in a given element of the feature sketch table is increased by one. Similarly, once that connection is fully established, that same value is decreased by one. As will be described in the following section, in our FPGA implementation of this FEM design the only attacks we have considered involve the monitoring of the SYN and ACK flags.

However, the FEM implementation can be easily reconfigured to analyze network activity which involves other flags. As can be seen in Figure 1, the estimate block selects the correct estimated value for each feature sketch from the sketch tables. Although there are various statistical estimation techniques that can be used for this purpose, we implement our estimate block as finding the minimum of the selected values from each sketch table of the feature sketch. We choose the minimum value as the estimate since it suffers the least amount from hash collisions. The output control block then selects the estimated value of only that feature sketch from which an estimate request was made. The FEM supports two functions for its operation:

1. update (src ip, dst ip, src port, dst port, flags)
 2. estimate (src ip, dst ip, src port, dst port, FS ID)
- The update function is a write-only operation to the

FEM where the flag information of the network packet is stored in the sketch tables of each feature sketch. During an update, the write operation is performed on all of the feature sketches of the FEM. As previously described, the value written into the sketch tables is some function of the input flags, which depends on the network feature being stored in each feature sketch. The key value in an update call (src ip, dst ip, src port, dst port) is used for determining the address of the sketch table where the value is to be written into.

6. An FEM-Based NIDS Application

In this section we discuss the practical application of the FEM architecture and how it can be used as the basis of a network intrusion detection system. From the application point of view, the FEM is placed within the network at some common node through which all of the network traffic passes. The FEM module's feature sketches are updated with the values in the network packets passing through that node. At any point in time, the FEM contains the up-to-date information of the network features, which can be obtained by estimating the feature sketches values

with an appropriate key. This information can be used to detect attacks using an anomaly-based intrusion detection algorithm. Consider the FEM having four feature sketches with keys and values as shown in Figure 3. As explained in the previous section, in our current implementation the FEM extracts only those port fields which are necessary to detect DoS or port scanning styles of attacks. Although each feature sketch stores only a specific feature, the overall nature of the network can be obtained by combining the information from all the feature sketches as illustrated below.

In Figure 3, FS1 contains information regarding the number of incomplete connections at any port of any computer within the network. So any (dst ip, dst port) pair in FS1 having a value above some threshold level is a likely victim of a SYN flood attack. Similarly, a (dst ip) in FS2 having a high value is also a probable candidate for a SYN flood attack or it may be being scanned for any open ports. For any victim (dst ip) in FS2, the corresponding (dst ip, dst port) values in FS 1 give the information regarding which ports are being attacked. For any possible victim (dst ip) in FS2, if there is a corresponding (src ip, dst ip) in FS4 having a high value then it can be known which computer is port scanning or performing a SYN flood attack on that particular victim. On the other hand, if there is a possible victim (dst ip) in FS2 but there is no corresponding (src ip, dst ip) in FS4 with a high value, it implies that the destination IP may be

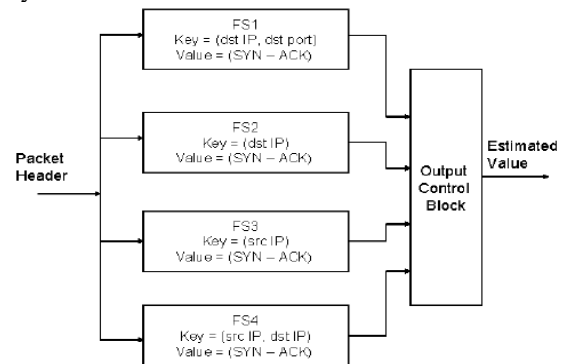


Figure 3. Sample application study

6.2 Area and Performance Results

Table 1 shows the variation in resource utilization and performance with respect to the size of a hash table (K) when FS and H are set to 2. From the trends it is observed that increasing K while keeping all other parameters constant does not result in significant variation in slice utilization.

FS	H	N_{slices}	f_{max} (MHz)	Throughput (Gbps)	C_{hw}	C_{sw}
1	1	2621	102.59	3.28	433	631
1	2	3428	103.82	3.32	433	1162
1	4	5035	101.48	3.25	433	2024
2	1	3403	101.72	3.26	433	1215
2	2	5010	102.67	3.29	433	2259
2	4	8197	101.39	3.24	433	3988
4	1	4991	103.54	3.31	433	2370
4	2	8140	102.86	3.29	433	4619
4	4	13694	99.60	3.19	433	8087

Table 2. FEM area and performance summary

This is because the hash tables are mapped onto Block RAMs which can be verified by noting the increase in Block RAM utilization when the hash table size is increased. For the different sizes of the hash table considered, the throughput is almost constant, indicating that the hash tables are not on the critical path of the FEM.

Table 2 shows the area and performance results for implementation of different configurations of the FEM on FPGA. From these results it can be observed that the number of slices utilized is directly proportional to both FS and H. This is due to the fact that an increase in FS or H directly corresponds to an additional hash function and sketch table. By considering the increments in the slices when FS and H are increased, we can see that a single hash function along with one sketch table takes up around 800 slices, which is about 6% of the total number of available slices on the Xilinx XC2VP30 FPGA. As a result, the configuration with FS=4 and H=4 occupies over 9900 of the slices available in our FPGA.

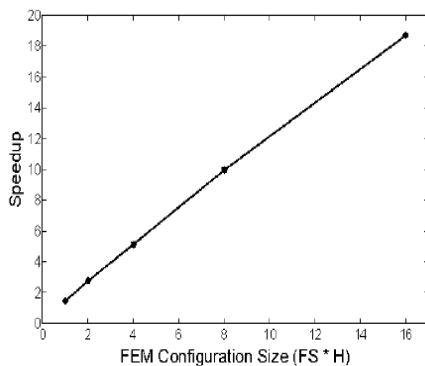


Figure 5. Hardware speedups for different FEM sizes

7. Conclusion

Feature extraction is an important component of various applications in domains such as networking,

data mining, and signal processing. In this paper we propose a reconfigurable architecture for real-time feature extraction of high speed networks used for anomaly detection, and evaluate its performance using an FPGA-based hardware development platform. We show how the FEM can be used to detect various types of network attacks, and that the reconfigurability of the architecture provides the flexibility to store various network features by making minor changes to the feature sketches. Our results clearly demonstrate that this architecture is several times faster than an equivalent software implementation (up to 18 x). We have also observed that the relative performance of our hardware implementation improves as the number of features is increased, which is desirable for better accuracy. It is also seen that the throughput (as high as 3.32 Gbps) is unaffected by the number of features monitored, thereby making the architecture suitable for high performance network intrusion detection systems.

References

- [1] NIPS 2003 Workshop on Feature Extraction. Available at <http://Hclopinnet.com/isabelle/Projects/NIPS2003, 2003>.
- [2] Z. Baker and V. Prasanna. Time and area efficient pattern matching on FPGAs. In Proceedings of the International Symposium on Field Programmable Gate Arrays (FPGA), February 2004.
- [3] C. Clark, W. Lee, D. Schimmel, D. Contis, M. Kone, and A. Thomas. A hardware platform for network intrusion detection and prevention. In Proceedings of the Third Workshop on Network Processors and Applications (NP3), February 2003.
- [4] C. Clark, C. Ulmer, and D. Schimmel. An FPGA-based network intrusion detection system with on-chip network interfaces. International Journal of Electronics, 93(6), June 2006.
- [5] A. DeHon. The density advantage of configurable computing. IEEE Computer, 33(4), April 2000.
- [6] Z. Guo, W. Najjar, F. Vahid, and K. Vissers. A quantitative analysis of the speedup factors of FPGAs over processors. In Proceedings of the International Symposium on Field- Programmable Gate Arrays (FPGA), February 2004.
- [7] Institute for Visualization and Perception Research, University of Massachusetts, Lowell. Contents of network intrusion collected data. Available at <http://Hivpr.cs.uml.edu>, 2006.
- [8] B. Jenkins. Hash functions and block ciphers. Available at <http://Hburtleburtle.net/bob/hash>, 2006.
- [9] J. Moscola, J. Lockwood, R. Loui, and M. Pachos. Implementation of a content-scanning module for an internet

firewall. In Proceedings of the IEEE Symposium on Field-Programmable Custom Computing Machines (FCCM), April 2003.

[10] D. Nguyen, G. Memik, S. Memik, and A. Choudhary. Real-time feature extraction for high speed networks. In

Proceedings of the International Symposium on Field- Programmable Logic and Applications (FPL), August 2005.

[11] M. Nixon and A. Aguando. Feature Extraction and Image Processing. Elsevier, Inc., 2004.

[12] G. Saha, P. Kumar, and S. Chakroborty. A comparative study of feature extraction algorithms on ANN based speaker model for speaker recognition applications. In Proceedings of the International Conference on Neural Information Processing (ICONIP), November 2004.

[13] D. Schuehler and J. Lockwood. A modular system for FPGA-based TCP flow processing in high-speed networks. In Proceedings of the International Symposium on Field- Programmable Logic and Applications (FPL), August 2004.

[14] R. Schweller, A. Gupta, E. Parsons, and Y Chen. Reversible sketches for efficient and accurate change detection over network data streams. In Proceedings of the ACM Internet Measurement Conference (IMC), October 2004.

[15] H. Song, S. Dharmapurikar, J. Turner, and J. Lockwood. Fast hash table lookup using extended bloom filter: An aid to network processing. In Proceedings of ACM SIGCOMM, August 2005.

[16] H. Song and J. Lockwood. Efficient packet classification for network intrusion detection using FPGA. In Proceedings of the International Symposium on Field Programmable Gate Arrays (FPGA), February 2005.

[1] NIPS 2003 Workshop on Feature Extraction. Available at <http://Hclopinnet.com/isabelle/Projects/NIPS2003>, 2003.