CruxML
REAL-TIME COMPUTING & MACHINE LEARNING

Email: info@cruxml.com
Website: www.cruxml.com

Cybersecurity Analysis Neural Engine Whitepaper

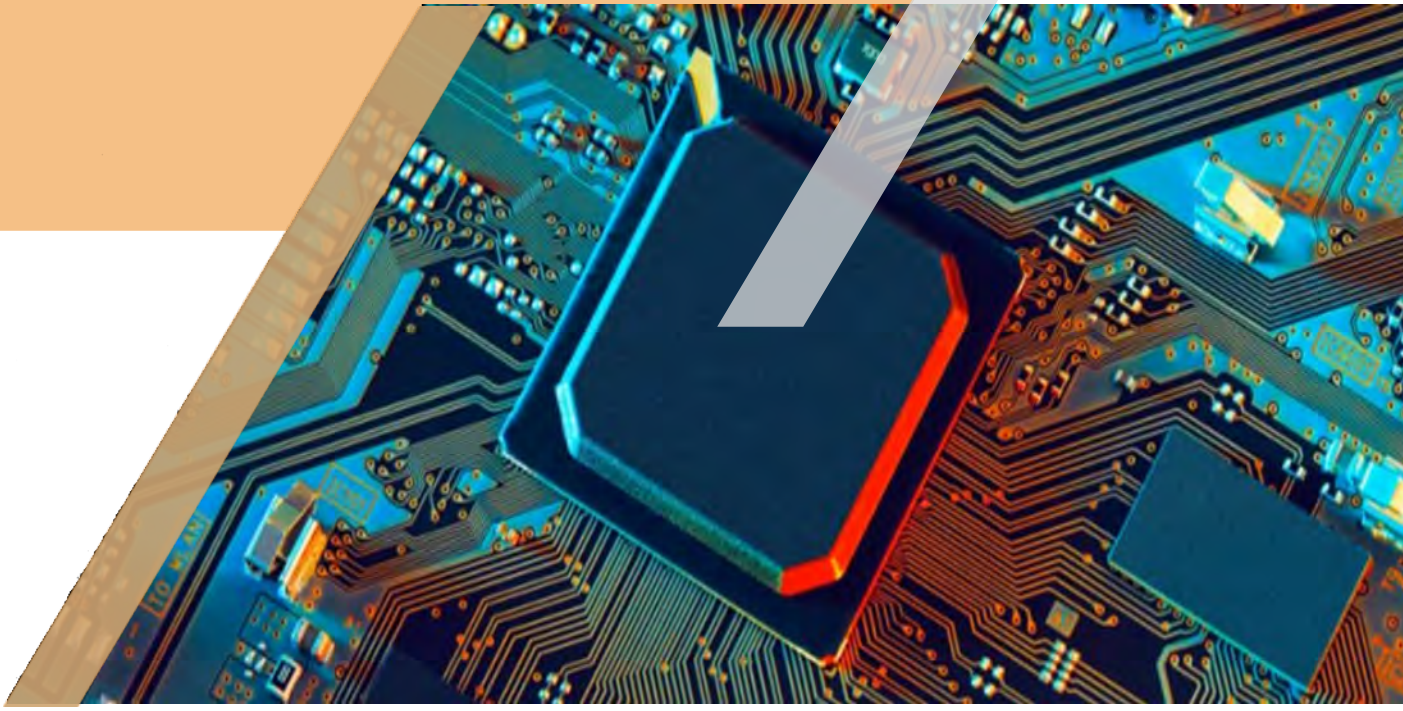

Prof. Philip Leong, CTO
Email phwl@cruxml.com
Dr. Barry Flower, CEO
Email barry@cruxml.com

This page intentionally left blank.

FPGA-based Machine Learning

Philip Leong and Barry Flower
{phwl,barry}@cruxml.com

2021/03/15

Executive Summary

A fundamental challenge with cyber security systems is the associated requirement to perform sophisticated data analysis at high speed. While machine learning (ML) is effective at addressing many Cyber problems, its computational complexity often makes its implementation infeasible at line rates. Signature-based intrusion detection systems (IDSs) identify known attacks and fall into the misuse detection class. Machine Learning approaches that learn the behaviour of the traffic flow fall into the class of anomaly detectors. We propose CruxML CANE, an FPGA-based IDS which achieves line rate speeds and combines signature (existing capability) and anomaly based (proposed capability) detectors. Compared with software implementations on processors and GPUs, our hybrid IDS (HIDS) system is more secure, accurate and performant. Moreover, it has a greatly reduced attack surface as the insertion of malicious code, injection attacks and viruses do not have an FPGA counterpart.

Keywords: FPGA, field programmable gate array, real-time, machine learning, edge

1 Introduction

Cyber artificial intelligence (AI) is an extremely challenging problem with conflicting requirements. On one hand there is an underlying need for technology to be able to defend against attacks which are ever increasing in sophistication and this requires more and more computation. On the other hand, it must be achieved with low latency which limits the amount of computation that can be done. Moreover, the attacks change over time so the system must be field-upgradable and intrusion detection systems (IDSs) must be capable of adapting to changing attack strategies. To ensure that the cure is not worse than the disease, we need to achieve line speed and introduce negligible or minimal additional latency. Failure to achieve line speed means either sacrificing operational ability or missing threats.

Field-programmable gate arrays (FPGAs) are an enabling technology for AI in Cyber Security and Information. They can be effectively applied from the Edge (ie. Xilinx FPGA chips) to the server room (Maxeler, Intel and Xilinx appliances) and Cloud (Amazon F1, etc). Graphics processing unit (GPU) based acceleration technologies are not suited to this problem as they require data to be transferred from the network adaptor, to a processor, to the GPU and back to the processor, each step introducing unnecessary latency.

2 FPGAs

FPGAs are commercial off the shelf user-customizable integrated circuits which are ideally suited to this problem. In the Cyber AI context higher degrees of parallelism compared to GPUs and CPUs are possible leading to an ability to perform more computation in a given time and hence more advanced cyber security algorithms become feasible It is also possible to integrate all layers

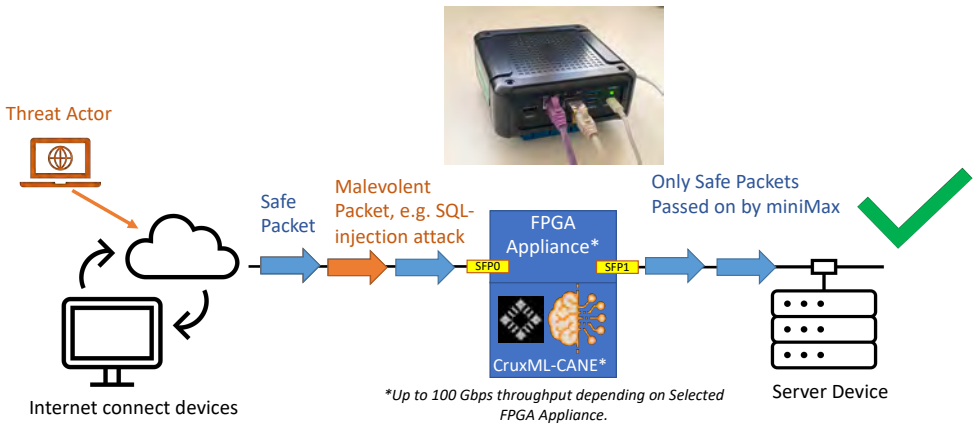

Figure 1: Snort IPS uses a series of rules that help define malicious network activity and uses those rules to find packets that match against them and generates alerts for users.

of the network stack, a pattern-based intrusion detection system like Snort, and a sophisticated DNN on the same device. Such integration improves latency and reduces power consumption. Importantly, FPGAs have a much smaller attack surface as they don't execute sequences of instructions and the skill set for attacking FPGAs is very different (and highly specialized) to that for CPU-based implementations

A good analogy is that FPGAs are like Formula 1 cars while CPUs and GPUs are respectively like passenger cars and busses (with application-specific integrated circuits being akin to a one-off rocket car).

3 CruxML CANE Intrusion Detection System

As illustrated in Figure 1, signature-based IDSs identify known attacks. The CruxML/Maxeler IDS implements Snort Intrusion Prevention System (IPS) rules at 100 Gbps with negligible impact to throughput. It can perform arbitrary filtering operations and detect exploits such as a SQL injection attack. Snort IPS uses a series of rules that help define malicious network activity and uses those rules to find packets that match against them and generates alerts for users. An example is Snort-IPS which is computationally expensive as incoming data must be compared with all signatures in a database.

Unfortunately, fixed rules are not sufficient to ensure the detection of sophisticated attacks. Machine learning approaches, which learn the behaviour of the traffic flow, fall into the class of systems called anomaly detectors. They are challenging to implement with low latency. We propose a hybrid IDS (HIDS) system, called CruxML-CANE (Cybersecurity Analysis Neural Engine), which combines signature and machine learning based IDS at line speed [1] and is illustrated in Figure 2. This architecture leads to a more secure, accurate and performant network anomaly detector.

The previous architecture can be made capable of adapting to changing conditions using an approach we have demonstrated for radio-frequency applications [2], and illustrated in Figure 3. While CruxML-CANE continues to perform line-speed detections, any new attack exemplars are

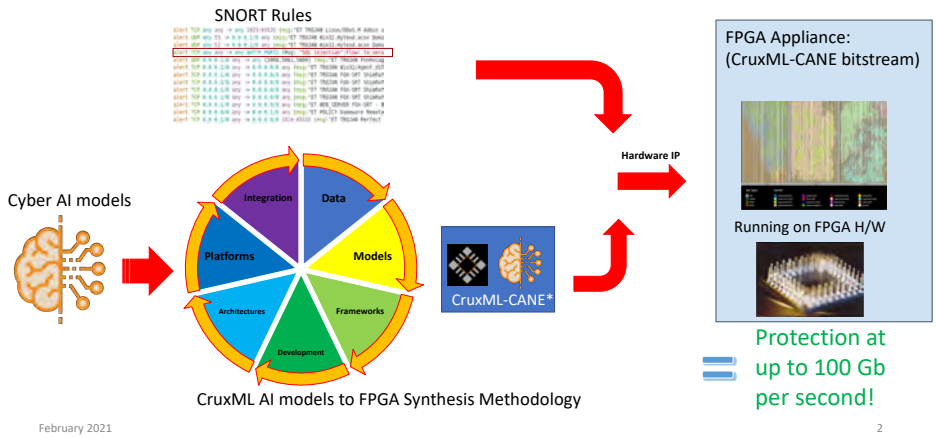

Figure 2: CruxML-CANE combines our existing signature based IDS with machine learning to detect a much broader range of attacks. .

sent to a server equipped with one or more GPUs that are used to perform off-line training on the augmented training set and does not have strict real-time requirements. After training has completed, a new FPGA design is created and uploaded to the IPS device which is now capable of detecting the new attack strategy. This approach can be fully automated or utilise semi-supervised learning, allowing the injection of expert cybersecurity operative knowledge into the system. This is an approach which can meet the Adaptation challenge of Cyber AI.

4 Conclusion

CruxML CANE will combine state of the art AI technology and a highly pipelined FPGA-based inference engine in a hybrid IDS which will protect against elaborate Cyber attacks.

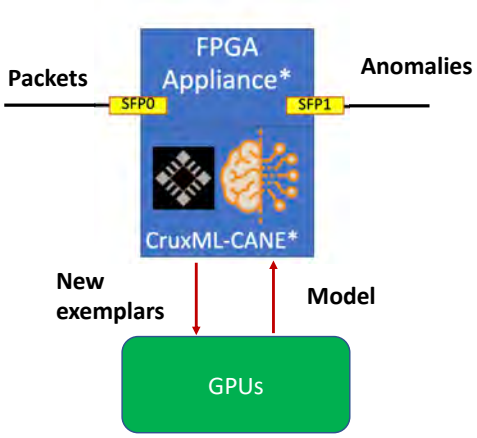

Figure 3: CruxML-CANE will perform online learning to adapt to changing conditions. .

References

[1] Y. Umuroglu et. al., “FINN: A framework for fast, scalable binarized neural network inference,” in *Proc. FPGA*, 2017.

[2] S. Tridgell et. al., “Real-time automatic modulation classification using RFSoC,” in *IEEE IPDPSW, year = 2020,*.