OpenFlow Multi-Table Lookup Architecture for Multi-Gigabit Software Defined Networking (SDN)

Document Version:
Peer reviewed version

Queen's University Belfast - Research Portal:
Link to publication record in Queen's University Belfast Research Portal

OpenFlow Multi-Table Lookup Architecture for Multi-Gigabit Software Defined Networking (SDN)

Keissy Guerra-Perez
Centre for Secure Information Technologies
Queen's University Belfast
Northern Ireland
k.guerra@qub.ac.uk

Sandra Scott-Hayward
Centre for Secure Information Technologies
Queen's University Belfast
Northern Ireland
s.scott-hayward@qub.ac.uk

ABSTRACT

The proposed multi-table lookup architecture provides SDN-based, high-performance packet classification in an Open-Flow v1.1+ SDN switch. The objective of the demonstration is to show the functionality of the architecture deployed on the NetFPGA SUME Platform.

Keywords

Packet Classification, Multi-Table Lookup, Software Defined Networking, OpenFlow

1. INTRODUCTION

The main differentiator between traditional networks and SDN is the flow-based management of the network elements e.g. switches, routers etc. In order to support the flow-based management of the switch to achieve the improved programmability, reduced latency and higher performance promised by SDN networks, improved flow classification is required. The flow classification capability is measured in terms of number of lookups, number of supported flows and the size of the table entry per flow for flow statistics and flow control. In order to improve flow processing performance, multiple table pipeline processing was introduced in the OpenFlow specification v1.1 [1]. In the pipeline process, when a packet arrives at a network device, it is passed through one or more flow tables depending on the packet processing requirements. This supports optimization of the flow processing technique by fast handling of simple flows with further processing applied to those flows which require it. The implementation of the multiple table pipeline process in hardware is challenging due to the fixed functionality of available switches and the cost of Ternary Content Addressable Memory (TCAM).

2. MULTI-TABLE ARCHITECTURE

The proposed solution is a multiple table lookup architecture for SDN-based, high-performance packet classification. The algorithmic lookup engine works by using a combination of single-field algorithm lookups in parallel. The algorithms are based on exact matching, range matching, and longest-prefix matching, as determined by the lookup field, and results are combined via the label method. The ability to perform a lookup based on a combination of single field algorithm results means that the network device can be flexible and programmed for a range of applications e.g. cyber-security/network forensics. Although the implementation uses OpenFlow, a pre-processing packet parsing function provides field entries, which means that it is reconfigurable, and not tied to OpenFlow. The potential for this lookup architecture for SDN based on simulation and synthesis has been described in [2, 3]. The hardware architecture is illustrated in Figure 1.
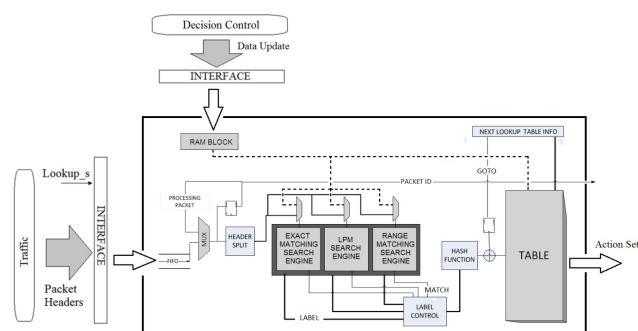

Figure 1: Hardware Architecture of Classifier based on SDN Programmability

3. DEMONSTRATION

The objective of this demonstration is to show the Multi-Table Lookup Architecture functionality on the NetFPGA SUME Platform [4]. The NetFPGA-SUME open-source platform is based on a Virtex-7 field programmable gate array (FPGA) with peripherals supporting high-performance and high-density networking designs. In this first demonstration, the NetFPGA has been configured to demonstrate a sample use-case (e.g. L2-L3-ACL) with 4 lookup tables, as illustrated in Figure 2.
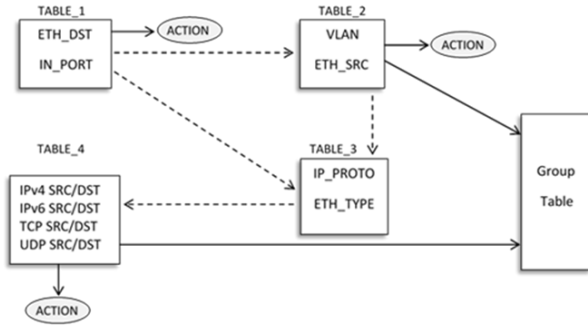
Figure 2: Example Multi-Table Use Case - L2-L3-ACL

The NetFPGA setup for the demonstration is shown in Figure 3. The Virtex-7 FPGA is programmed with the multi-table lookup function. The network traffic in the form of packet headers are stored in an embedded SRAM block.
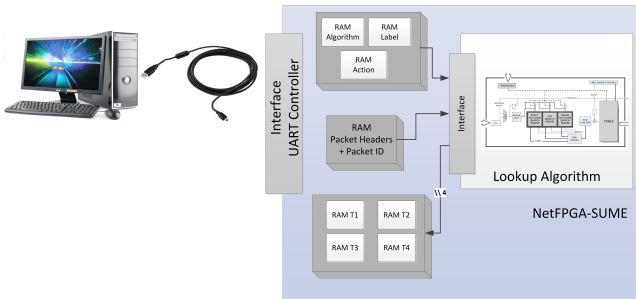

Figure 3: NetFPGA-SUME Demonstrator for Multi-Table Lookup Architecture

Performance results based on 8K flow rules are demonstrated for the following functions: (1) Lookup an existing flow entry, (2) Update an existing flow entry, (3) Insert a new flow entry, and (4) Delete an old flow entry.

3.1 Update/Insert/Delete Operation

For the update operation, the flow entries and update algorithm information are processed in the Host PC (representative of the SDN Controller). The update information is then transmitted to the NetFPGA over the UART interface and stored in an embedded SRAM block. Once the required update information is stored, the update/insert/delete process of the Multi-Table Lookup is enabled.

3.2 Lookup Operation

For the lookup operation, the packet headers are read from the relevant SRAM block and processed by the FPGA. Based on the pipeline process, multiple packets can be processed simultaneously. The lookup results are therefore output to a set of SRAM blocks and returned to the Host PC for performance measurement.

4. ACKNOWLEDGMENTS

5. REFERENCES

[1] OpenFlow Switch Specification Version 1.1, Open Networking Foundation.

[2] K. G. Perez, X. Yang, S. Scott-Hayward, and S. Sezer. A configurable packet classification architecture for software-defined networking. In *System-on-Chip Conference (SOCC), 2014 27th IEEE International*, pages 353–358. IEEE, 2014.

[3] K. G. Perez, X. Yang, S. Scott-Hayward, and S. Sezer. Optimized packet classification for software-defined networking. In *Communications (ICC), 2014 IEEE International Conference on*, pages 859–864. IEEE, 2014.

[4] N. Zilberman, Y. Audzevich, A. Covington, and A. Moore. NetFPGA SUME: Toward 100 Gbps as Research Commodity. *IEEE Micro*, 34(5):32–41, 2014.