FERMILAB-POSTER-2017-008-CD

Deep Packet/Flow Analysis using GPUs for High-Bandwidth Networks

Q. Gong, W. Wu, L. Zhang, S. Sasidharan, P. DeMar (Fermilab)

Motivation

- Deep packet inspection (DPI) is widely used in content-aware network applications, such as surveillance, statistics gathering, and traffic control.

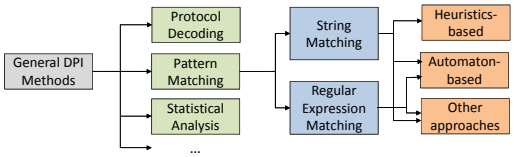

- Rising sophistication in the types of network attacks has driven a need for packet processing & inspection tools that perform at flow level.
- Packet processing devices face severe performance & scalability challenges in high-bandwidth network environments.
- GPUs work exceptionally well for packet-based network analysis applications; but a solid flow-based GPU packet inspection tool is missing.

Challenges in Deep Packet Inspection for TCP

TCP Stream Reassembly

Payload of packet affiliated to the same TCP stream need to be assembled before matching against pre-defined patterns.

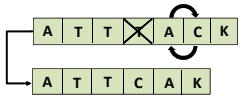

Data structure used in stream reassembly:
- Hash table
- TCB (i.e., connection states, sequence number, etc.)

Memory/Performance Burden in Buffering Out-of-order Data

Conventional packet normalizers have to buffer all packets following a missing packet, until they become in-sequence again, to prevent TCP fragmentation evasion attacks.

GPU-based Deep Packet Analysis: Framework

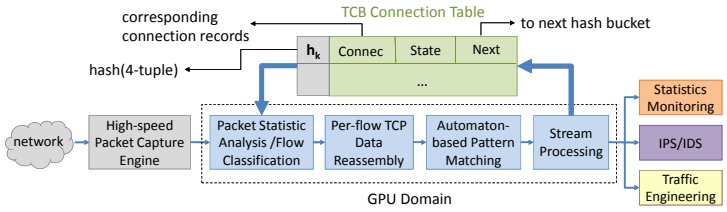

Key Functions

Flow Classification and Reordering:
- Inter-flow classification: sort TCP streams according to their *TCP 4-tuple*
- Intra-flow reorder: sort same-stream affiliated packets by their *sequence #*

Packet Normalization: drop duplicate packets and merge the overlapping payload

Pattern Matching: match the payload against fixed strings

Stream Processing: keep the internal states between consecutive batches

GPU-based Deep Packet Analysis: Key Mechanisms

GPU-based Solutions

- Parallel sorting for intra-batch flow classification and packet reordering
- Parallel pattern matching based on AC algorithm
- Lock-free parallel hash table (by allowing only one thread to access a hash entry at a time) for maintaining the inter-batch stream connections
- AC and AC-suffix automaton methods to preserve the inter-batch states for flow-based pattern matching

Key Features

- Process packets in batch on GPUs
- Keep track of the states of intra- and inter-batch packets without buffering or dropping the out-of-order ones

Flow Classification

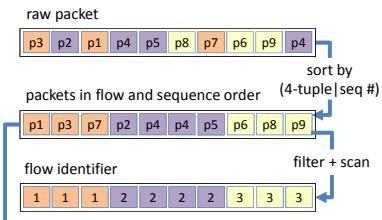

Stream Normalization

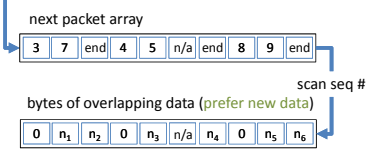

Pattern Matching

Intra-batch: AC automaton

Keywords: *X* = {he, his, she, hers}

state transition automaton

- One thread per packet
- Each thread scans extra N bytes towards its consecutive packet, where N is the maximum branch length of the AC tree
- N equals zero for UDP packets

Inter-batch: AC automaton & Suffix-AC automaton

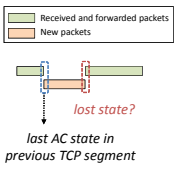

Suffix Pattern Tree (PST)

Suffix set of *X*: {e,is,s,he,ers,rs}

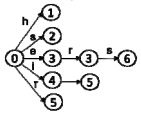

suffix string = *path*(state)

Out-of-order Packets

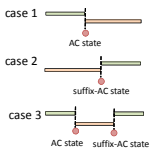

Performance

Throughput — ***Support for 40 Gbps Ethernet!***

- Traffic source: real traffics mirrored from the Fermilab gateway
- Traffic pattern (average per batch):

# of packets	1 million
# of data packets	776,207
mean packet length	1415-byte
# of connections	15,500

- Base system: Intel Xeon CPU E5-2650 @ 2.30 GHz, NVIDIA GPU K40
- Flow classification & TCP reassembly: **72.96** Mpps (×192 speedup comparing to libnids on CPU)
- State management: 286.85 Mpps
- Pattern matching: 5.83 Mpps

Resist to ordinary adversaries
- Robust stream reassembly in facing of out-of-order packets
- Immune to SYN flood and 'cold start' in doing normalization
- Exempt from attacks on available buffer memory with timeout and connection eviction mechanisms

Our Prior Works

Selective presentations:

[1] Wu, et al. "Application-Oriented Network Traffic Analysis based on GPUs", GCASR'2016
[2] Wu, et al . "Packet-based network traffic monitoring and analysis with GPUs," GTC'14

Selective publications:

[1] Wu, et al. "Network traffic monitoring and analysis with GPUs." *SC* 2013.

This manuscript has been authored by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics.