

SC25 Network Research Exhibition: Demonstration Publishable Abstract

In-network payload classification on streaming scientific data

Cong Wang, RENC/UNC Chapel Hill, cwang@renci.org

Joaquin Chung, U Chicago/ANL, chungmiranda@uchicago.edu

Ganesh C Sankaran, RENC/UNC Chapel Hill, sankarang@renci.org

Abstract

This demonstration presents an in-network packet classification approach designed to operate at high data rates, identifying whether packet payloads contain new information. The method classifies payloads based on differences relative to a reference payload, motivated by distributed scientific workflows that detect rare events observable across time and space. When such workflows span multiple sites, analyzing data in transit can reduce unnecessary data movement and accelerate discovery.

A main challenge in in-network processing is that, while programmable switches operate at line rate, they have limited access to payload data. For example, Intel Tofino switches can process only the first 256 bytes of a packet, even though Ethernet frames can carry up to 1500 bytes and Jumbo frames up to 8192 bytes. To overcome this limitation, we derive a payload checksum using pre-computed IP and TCP checksum fields. By nullifying the header contributions in the TCP checksum, we isolate a signature representing only the payload contents, enabling change detection without full payload inspection.

To evaluate this approach, we generate reference payloads with known modifications and measure classification accuracy in terms of false positives, false negatives, true positives, and true negatives. The experiment is deployed on the FABRIC testbed, where two VM nodes communicate through a P4-programmable Tofino switch. Preliminary results show 100% accuracy in detecting checksum changes, demonstrating the feasibility of checksum-based payload classification for high-speed, in-network data analysis.

Goals

Measure accuracy for

1. Single reference and change
2. Multiple references and change
3. At higher data rates such as 100 Gbps

Impacts

This work advances the use of in-network processing to enable real-time data analysis and data transfer reduction for core network traffic. By classifying payload changes directly within programmable switches, we demonstrate a method to minimize unnecessary data transmission and accelerate time-to-insight for distributed experiments. This capability supports broader goals in data-driven research and contributes to the community's growing interest in edge and in-transit data processing for large-scale scientific applications.

Utilizing FABRIC research infrastructure, this demo enables a high-speed, programmable networking environment to validate our checksum-based classification method at realistic data rates, up to 100 Gbps—capabilities not available in standard laboratory settings. The next steps will focus on extending to multiple reference payloads, and integrating the approach into distributed science workflow systems, paving the way for autonomous, network-assisted scientific data pipelines.

Resources

The FABRIC testbed will be used for this demonstration. The experimental topology, shown in Figure 1, consists of sender and receiver nodes connected through a P4-programmable Intel Tofino switch over a high-throughput link. To emulate realistic network conditions, we will use tcpreplay or the TRex traffic generator to generate synthetic traffic flows from the sender to the receiver.




Figure 1. Demo topology on FABRIC

Involved Parties

- Cong Wang, RENCi/UNC Chapel Hill,
cwang@renci.org
- Joaquin Chung, University of Chicago/ANL,
chungmiranda@uchicago.edu
- Ganesh C Sankaran, RENCi/UNC Chapel Hill,
sankarang@renci.org
- Raj Kettimuthu, University of Chicago/ANL,
kettimut@mcs.anl.gov
- Paul Ruth, UNC/RENCi, pruth@renci.org
- Tom Lehman, Consultant, tom.w.lehman@gmail.com