



Enabling P4 Hands-on Training using Hardware Switches in a Cloud System at the University of South Carolina

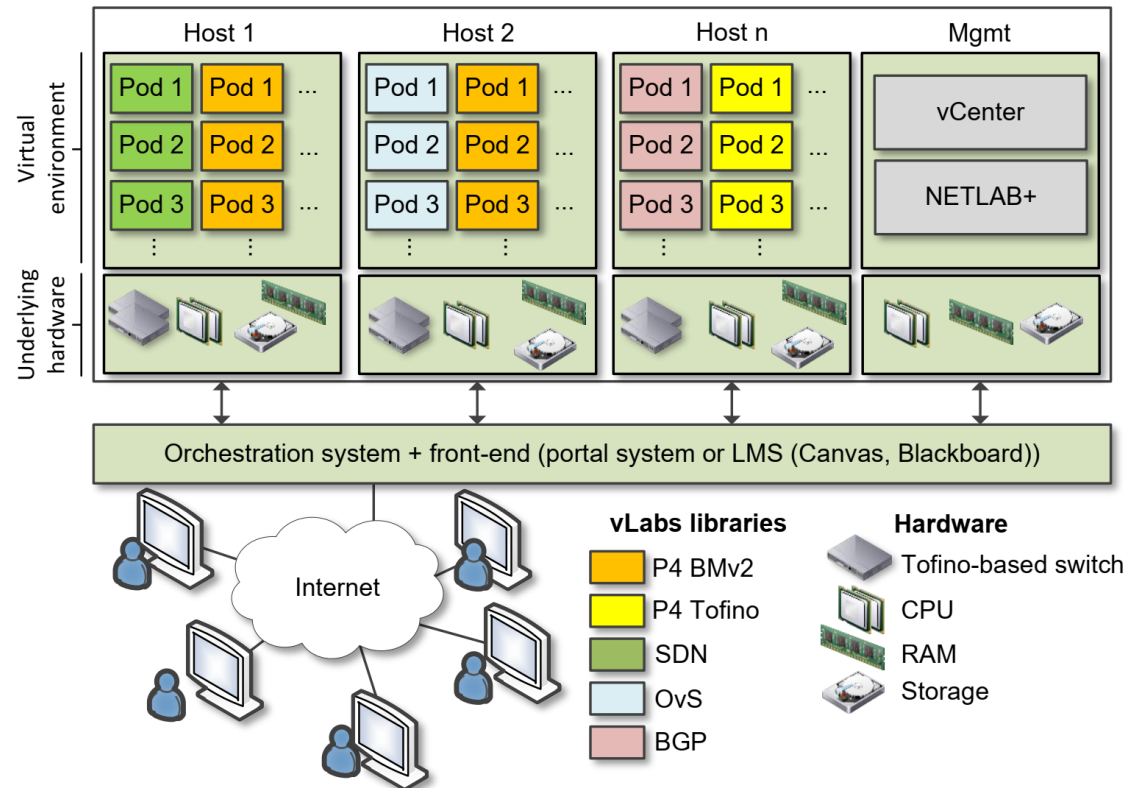
Jose Gomez, Elie Kfoury, Ali AlSabeH, Jorge Crichigno
College of Engineering and Computing, University of South Carolina

<http://ce.sc.edu/cyberinfra/>

Intel Headquarters - Santa Clara, CA
April 24-25, 2023

USC Cloud System

- The University of South Carolina (USC) deployed an Academic Cloud
- It provides remote-access capability to lab equipment via Internet
- Virtual labs on P4, routing, high-speed networks, and cybersecurity
- Hosts 1-n store equipment pods for virtual labs; Mgmt server manages the system



Open-Source Libraries

- A library consists of 10-20 activities (labs experiments + exercises)
- Each lab experiment includes a detailed, step by step manual
- Information about libraries: <http://ce.sc.edu/cyberinfra/cybertraining.html>

Library
Intro to P4 Programmable Data Plane (PDP) Switches with Tofino
Introductory to P4 PDP switches (BMv2)
Applications, Stateful Elements, and Custom Packet Processing
Security Applications of P4 PDP switches
Network Measurement Applications with P4 PDP switches
SDN Applications with ONOS
Intro to Open vSwitch (OvS)
Network Tools and Protocols

Library
Cybersecurity tools and applications
Intro to BGP
MPLS and Advanced BGP Topics
Interior Gateway Protocols and OSPF
Network Measurements with perfSONAR
Intrusion Detection with Zeek
Network Management Tools
Intro to High-speed Networks (including textbook)

Library on Introduction to P4


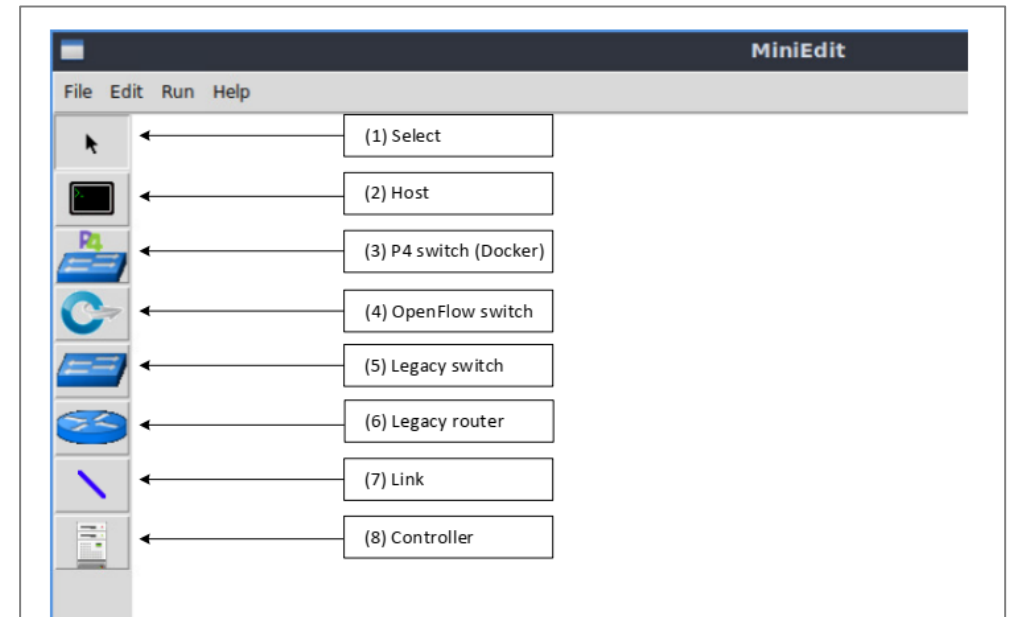
Experiments

- Lab 1: Introduction to Mininet
- Lab 2: Introduction to P4 and BMv2
- Lab 3: P4 Program Building Blocks
- Lab 4: Parser Implementation
- Lab 5: Introduction to Match-action Tables (Part 1)
- Lab 6: Introduction to Match-action Tables (Part 2)
- Lab 7: Populating / Managing Match-action Tables
- Lab 8: Checksum Recalculation and Deparsing

Exercises

- Exercise 1: Building a Basic Topology
- Exercise 2: Compiling and Testing a P4 Program
- Exercise 3: Parsing UDP and RTP
- Exercise 4: Building a Simplified NAT
- Exercise 5: Configuring Tables at Runtime
- Exercise 6: Building a Packet Reflector

MiniEdit GUI environment with P4 switches

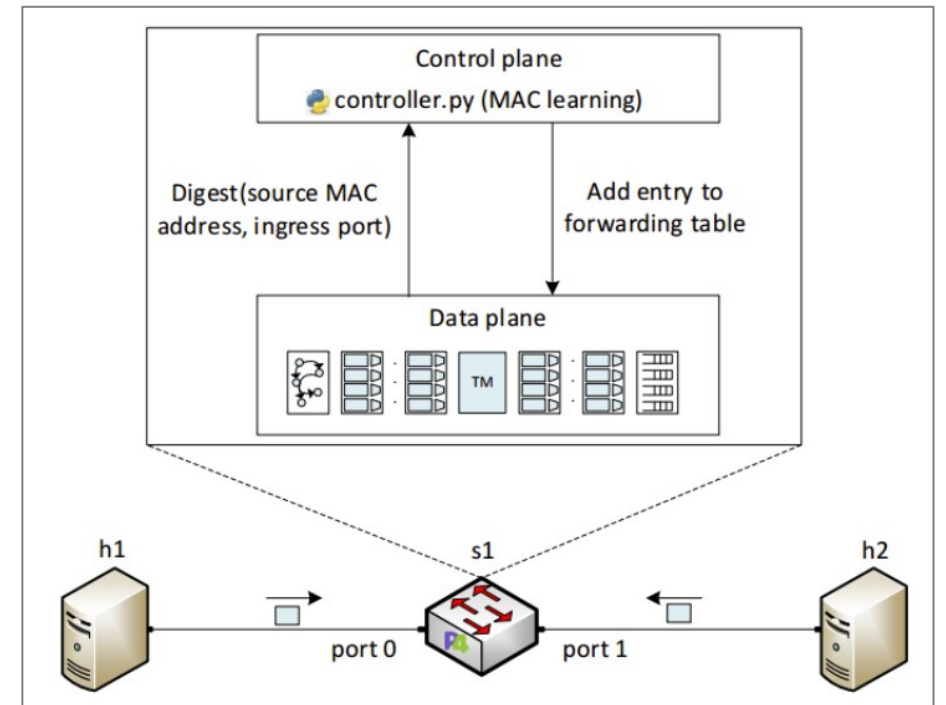


Library on Apps, Stateful Elements, and Custom Packet Processing

Experiments

- Lab 1: Introduction to Mininet
- Lab 2: Introduction to P4 and BMv2
- Lab 3: P4 Program Building Blocks
- Lab 4: Defining and Processing Custom Headers
- Lab 5: Monitoring the Switch's Queue using Standard Metadata
- Lab 6: Collecting Queueing Statistics using a Header Stack
- Lab 7: Measuring Flow Statistics using Direct and Indirect Counters
- Lab 8: Rerouting Traffic using Meters
- Lab 9: Storing Arbitrary Data using Registers
- Lab 10: Calculating Packets Interarrival Time w/ Hashes and Registers
- Lab 11: Generating Notification Messages from the Data Plane

Notification messages from the data plane




Library on Security Applications with P4

Experiments

- Lab 1: Introduction to Mininet
- Lab 2: Introduction to P4 and BMv2
- Lab 3: P4 Program Building Blocks
- Lab 4: Parser Implementation
- Lab 5: Introduction to Match-action Tables
- Lab 6: Implementing a Stateful Packet Filter for the ICMP protocol
- Lab 7: Implementing a Stateful Packet Filter for the TCP protocol
- Lab 8: Detecting and Mitigating the DNS Amplification Attack
- Lab 9: Identifying Heavy Hitters using Count-min Sketches (CMS)
- Lab 10: Limiting the Impact of SYN Flood by Probabilistically Dropping Packets
- Lab 11: Blocking Application Layer Slow DDoS Attack (Slowloris)
- Lab 12: Implementing URL Filtering through Deep Packet Inspection and String Matching

Identifying Heavy Hitters using CMS

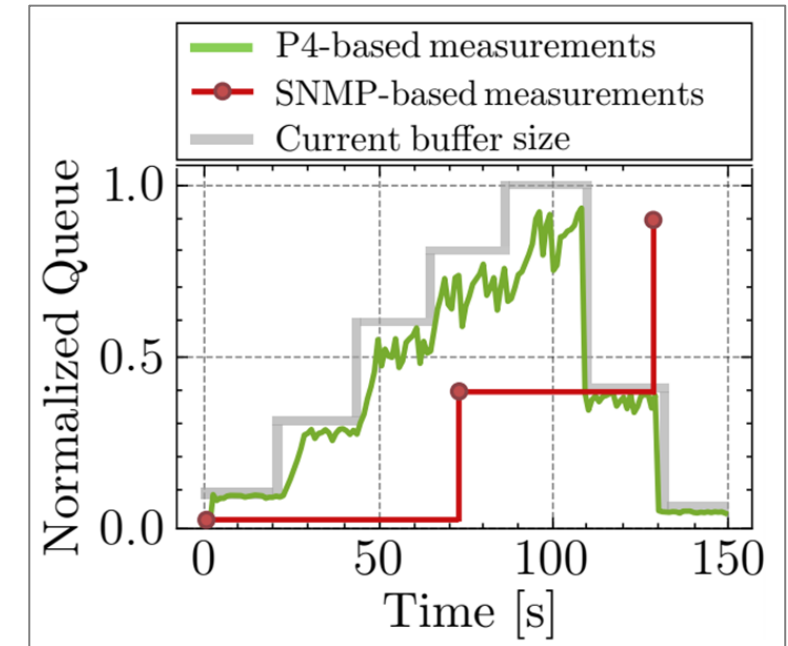


Library on Network Measurement Apps with P4

Experiments

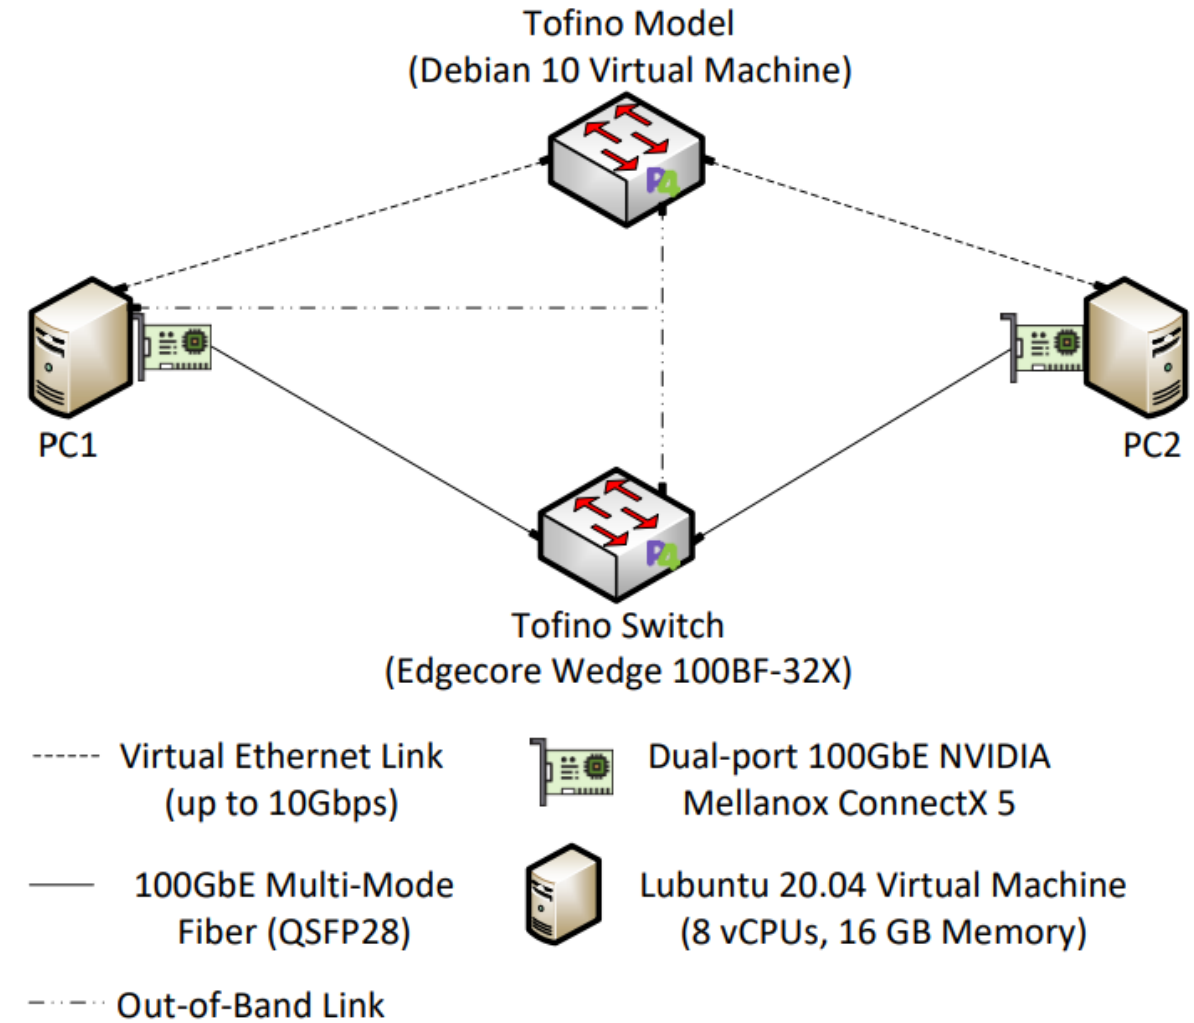
- Lab 1: Introduction to Mininet
- Lab 2: Introduction to P4 and BMv2
- Lab 3: P4 program building Blocks
- Lab 4: Defining and Processing Custom Headers
- Lab 5: Monitoring the Switch's Queue using Standard Metadata
- Lab 6: Detecting Microbursts and Identifying Flows Contributing to the Queue
- Lab 7: Measuring Flow Statistics using Direct and Indirect Counters
- Lab 8: Flow Rate Measurement and Traffic Policing using Meters
- Lab 9: Identifying Elephant Flows using Count-Min Sketches (CMS)
- Lab 10: Using In-band Network Telemetry (INT) to Track the Path Taken by Packets
- Lab 11: Measuring the round-trip time of flows
- Lab 12: Measuring the packet loss rate
- Lab 13: Measuring Packet Inter-arrival Time
- Lab 14: Identifying TCP Congestion Control Type using Inter-arrival Measurements
- Lab 15: Plotting Measurements from P4 on a Grafana Dashboard

Measurements P4 vs SNMP



Tofino Pod

- There are 8 pods that can be reserved
- End-to-end connection via 100Gbps fiber links, Tofino switch
- Users access equipment (e.g., Tofino switch) via an out-of-band connection
- Sample P4 codes for each lab (data plane)
- Lab experiments with step-by-step directions



Accessing to the Cloud System

- Users at USC can access the Tofino-based library using a web browser
- The user selects a lab activity and schedules a reservation
- A topology is deployed at the reserved time and the user can run the lab

Accessing the cloud system

The figure shows four stages of the process:

- (a)** A login page for netlab.cec.sc.edu with fields for Username and Password, and a Login button. Below the form is the text "Cyberinfrastructure Lab @ UofSC".
- (b)** A "Reservations" table with columns: Date/Time, Description, and Pod. A row shows a reservation for 2022-10-14 17:27 to 18:00 (22 mins) for the class "PDP with P4 - ASU Fall 2022" in pod "Tofino_H2_pod4". An "Enter Lab" button is visible.
- (c)** A network topology diagram titled "Lab 2: Introduction to P4 Tofino Software Development Environment (SDE)". It shows a central "Tofino Switch" (10.0.0.0/24) connected to two hosts, H1 and H2. H1 has interfaces eth0 (172.168.1.0/24), ma1, and eth1. H2 has interfaces eth0, ma1, and eth1. The switch has ports 0 and 1. A "Management Network" is also indicated.
- (d)** A terminal window showing the command prompt "root@tofino-switch:~/P4_labs/lab1/p4src#" and the output of "Status" for the Tofino Switch, PC 1, PC 2, and Tofino Model, all showing "CONNECTED".

(a) A user enters the cloud and (b) reserves a pod. (c) The scenario is realized at the reservation time. (d) The user clicks on a device (e.g., the Tofino switch), opening a console.

Selecting a lab experiment


The screenshot shows the "MyNETLAB" interface with the breadcrumb "Schedule (Self) > Select Class > Select Lab". A search bar contains the text "This pod provides P4 training using the Intel Tofino switch". Below is a table of lab experiments:

Lab Name	Action
Lab 1: Introduction to P4 and Tofino	▼
Lab 2: Introduction to P4 Tofino Software Development Environment (SDE)	▼
Lab 3: Parser Implementation	▼
Lab 4: Introduction to Match-action Tables	▼
Lab 5: Populating and Managing Match-action Tables at Runtime	▼
Lab 6: Checksum Recalculation and Packet Deparsing	▼
Lab 7: Inspecting the Resource Usage in the Tofino Switch	▼

At the bottom, there is a pagination control: "Show 25 entries Showing 1 to 7 of 7 items" and a page number "1" in a blue box.

Usage of the Cloud at USC

- Use of the virtual platform (accounting reservations for all lab libraries) at USC
- Workshops co-located / co-organized with Internet2, FABRIC, ESnet, EPOC, Research and Education Networks (TX's LEARN, CA's WASTC, GA's SOX, etc.) (BMv2 only)





UNIVERSITY OF
South Carolina



This work is supported by NSF award number 2118311

For additional information, please refer to
<http://ce.sc.edu/cyberinfra/>

Email: jcrichigno@cec.sc.edu, gomezgaj@email.sc.edu