Cybertraining on P4 Programmable Devices using an Online Scalable Platform with Physical and Virtual Switches and Real Protocol Stacks

Jorge Crichigno (PI), Neset Hikmet (Co-PI)

College of Engineering and Computing - University of South Carolina

Project website: https://research.cec.sc.edu/cyberinfra/cybertraining

Award 2118311

Motivation

- Traditionally, the data plane has been designed with fixed functions to forward data packets, using a small set of communication protocols.
- This closed-design paradigm has limited the capability of switches to costly proprietary implementations which are hard-coded by vendors.
- P4 programmable switches permit a programmer to program the data plane:
  - Define and parse new protocols.
  - Customize packet processing functions.
  - Measure events occurring in the data plane with high precision (nanosecond resolution).
  - Offload applications to the data plane (terabits per second processing rates).

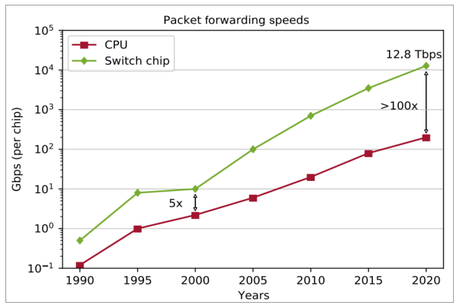

Packet forwarding speeds. Reproduced from https://tinyurl.com/bdzh9zy3

Goals and Project Overview

Goals

1. Facilitate the adoption of programmable P4 devices by CI professionals and by learners in general, by developing virtual labs.
2. Promote the integration of P4 and virtual labs into academic degrees

Overview

- Virtual lab libraries run on USC's cloud platform, referred to as the Academic Cloud, and on FABRIC.
- The Academic Cloud is used by colleges and universities, researchers, and the CI community for academic courses, workshops, and self-paced training.
- The libraries on FABRIC use JupyterHub notebooks , a web-based interface.
- Virtual lab libraries interact with FABRIC using FABRIC's APIs to deploy the learner's customized topology and embed P4 code alongside explanatory documentation, screenshots, videos, and other visual aids.
- Virtual lab libraries use state-of-the-art, production-grade devices such as software switches (e.g., Open vSwitch), hardware switches based on Intel's Tofino chips, and open-source operating systems and controllers (e.g., Open Network Linux, Open Network Operating System).

P4 programmable chip (Intel's Tofino)

Commercial switches based on Intel's Tofino chip

Background P4 Language

- P4 stands for Programming Protocol-independent Packet Processors.
- P4 is the *de facto* language for programming the data plane of switches.
- The programmable data plane can be considered as a domain-specific processor for networking.
- The processor model is referred to as Protocol Independent Switch Architecture (PISA).
- Programmable data planes are optimized to perform operations over packet headers / header fields (increment, parse headers, shift bits, etc.).
- Since they execute simple operations and use a simple Reduced Instruction Set Computer (RISC)-type instruction set, they can be implemented in the silicon at a minimal cost and size.
- P4 is distributed as open-source and is maintained by the P4 Project (https://p4.org/).

P4 code

```
state start {
    transition parse_ethernet;
}
state parse_ethernet {
    packet.extract(hdr.ethernet);
    transition select(hdr.ethernet.etherType) {
        TYPE_IPV4: parse_ipv4;
        default: reject;
    }
}
state parse_ipv4 {
    packet.extract(hdr.ipv4);
    transition accept;
}
```

Corresponding state machine

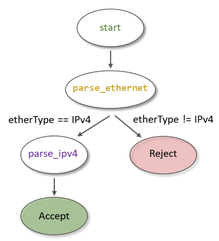

Academic Cloud

- Virtual lab libraries run on the Academic Cloud.
- The Academic Cloud is a virtual platform that manages and orchestrates computing resources to support virtual laboratories.
- Computing resources are distributed across four data centers in the United States.
- The figure below illustrates the Academic Cloud architecture at one of the four data centers. On average, each data center has approximately 10 servers.
- A pod is a collection of resources: virtual machines, virtual switches, virtual links, physical Tofino-based switches, physical fiber links, etc.
- Pods are orchestrated by the Academic Cloud platform on demand.
- NETLAB is a customized management application.

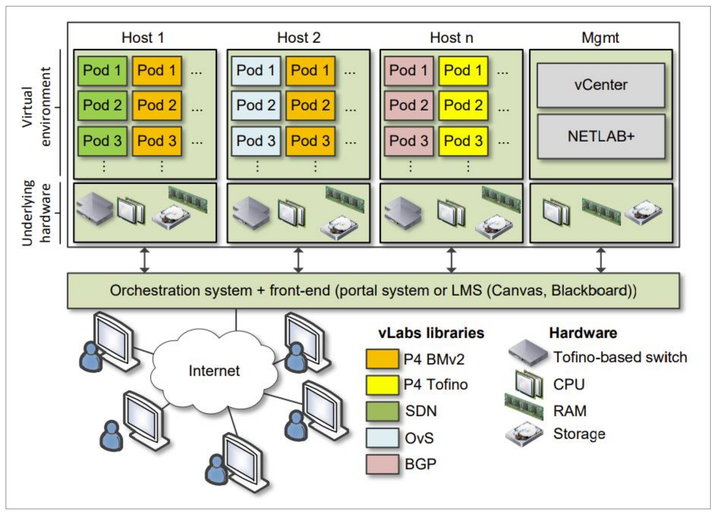

Academic Cloud

Feature	Academic Cloud	Public Cloud
Allocation of resources	Very granular allocation of physical resources	Not granular (access to the physical resources requires additional fees)
Custom pods	Easy to create custom pods	Difficult to design complex topologies
Cost	Cost-effective when used extensively	Cost-effective for individual / small VMs; costly for large VMs over time
Presentation layer for pedagogy	Topology is graphically presented to the learner using a regular browser	Not flexible; limited to providers' interface, e.g., command-line interface
Time-sharing resource feature	The owner controls who can access resources	Cloud provider controls who can access resources
Integration of physical devices	Easy; physical hardware can be integrated into pods	Difficult; no subscription plan permits integrating customized physical devices
Target	Specially built for pedagogy (education, research, and training)	General, used by a large variety of users
Typical users	From entry-level learners to PhD researchers	More experienced professionals, educators, students

Pod Example and Libraries

- The figure below shows a pod with two servers, two Tofino-based switches, and 100Gbps multi-mode fiber connections.

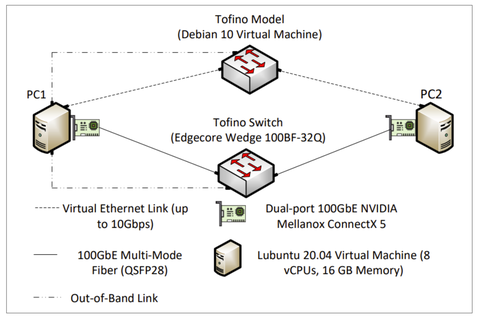

Tofino-based pod

- A virtual lab library consists of 8-15 experiments, systematically organized.
- Currently, the project has developed six virtual lab libraries:
  - Introduction to P4 Programmable Data Plane Switches.
  - P4 Programmable Data Planes: Applications, Stateful Elements, and Custom Packet Processing.
  - Cybersecurity Applications with P4 Programmable Data Planes.
  - Writing Fine-grained Measurements Apps with P4 Switches.
  - Introduction to Software-defined Networking.
  - Network Tools and Protocols.

Impact

Workshops

- 21 hands-on workshops were organized between 2021-2023.
- Workshops are face-to-face and online, using the Academic Cloud.
- Nearly 1,500 learners attended the workshops.
- Workshops were organized as stand-alone events or co-located to conferences, such as "Internet2 Technology Exchange."

Partners

- Minority-Serving Cyberinfrastructure Consortium (MS-CC), Internet2, FABRIC, ESnet / Lawrence Berkeley National Laboratory (LBNL), Oak Ridge National Laboratory (ORNL), Western Academy Support and Training Center (WASTC), over ten Research and Education Networks (RENs) from Texas, Louisiana, New York, Georgia, Missouri, Florida, Colorado, California, Indiana.

Internet2 Technology Exchange Conference, Session on P4 Switches. Denver – CO, Dec. 2022

FABRIC Community Workshop KNIT 6, Session on P4 Switches. Austin – TX, Apr. 2023

Academic Courses

- The virtual labs and companion materials (guided lab experiments, electronic book / booklets, open-ended exercises are now used in academic programs at colleges and universities.

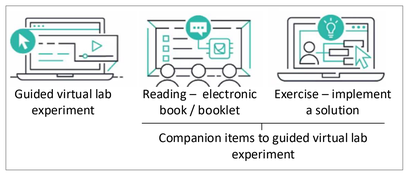

Virtual labs and companion material

Book, High-Speed Networks w/ virtual labs

Usage of the Academic Cloud for Training on P4 Technology

- Trainings with virtual labs on P4 switches started in 2022.
- At the USC site of the Academic Cloud alone, learners conducted 13,416 laboratory experiments, for a total of 66,072.42 hours.

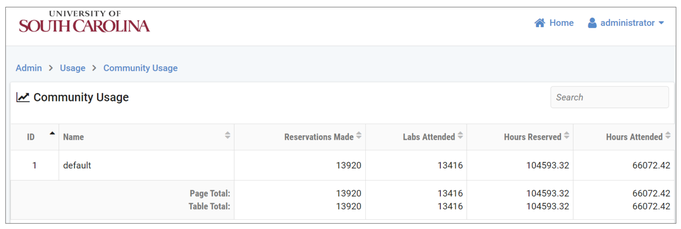

Usage, Academic Cloud

Publications

Journal Papers

- An Exhaustive Survey on P4 Programmable Data Plane Switches: Taxonomy, Applications, Challenges, and Future Trends", IEEE Access, Jun. 2021.
- "A Survey on Security Applications of P4 Programmable Switches and a STRIDE-based Vulnerability Assessment", Computer Networks, Jan. 2022.
- "A Survey on TCP Enhancements using P4-programmable Devices", Computer Networks, May 2022.
- "P4Tune: Enabling Programmability in Non-Programmable Networks", IEEE Communications Magazine, Feb. 2023.
- "A Survey on Rerouting Techniques with P4 Programmable Data Plane Switches", Computer Networks, Apr. 2023.

Conference Papers

- "Enhancing Blockage Detection and Handover on 60 GHz Networks with P4 Programmable Data Planes", IEEE BlackSeaCom, Istanbul, Turkey, Jul. 2023.
- "Understanding the Performance of TCP BBRv2 using FABRIC", IEEE BlackSeaCom, Istanbul, Turkey, Jul. 2023.
- "Effective DGA Family Classification using a Hybrid Shallow and Deep Packet Inspection Technique on P4 Programmable Switches," IEEE ICC 2023, Rome, Italy, Jun. 2023.
- "P4CCI: P4-based Online TCP Congestion Control Algorithm Identification for Traffic Separation", IEEE ICC, Rome, Italy, Jun. 2023.
- "Enabling P4 Hands-on Training in an Academic Cloud", International Workshop on Test and Evaluation of Programmable Networks, Marina del Rey, CA, May 2022.

Book

- "High-Speed Networks: A Tutorial (Practical Networking)", 1st Edition, Springer, 2022.