P4 On The Move!

Andy Fingerhut

intel

P4.org working groups

Applications work group

- New focus: Congestion Control
- New co-chairs

Nandita Dukkipati
Google

Rong Pan
Intel

https://p4.org/join (to join p4-apps or any of the other P4.org email lists)

API work group

- New release of version 1.4.0 of P4Runtime API specification, including
  - Simplifying the process of primary controller election
  - Added per-color counters for meter entries
  - Clarifications

- Some future-looking discussions in P4 API work group meetings
  - Table-Driven Interface (TDI) presented, and later discussed
  - Nvidia presented a proposal for optimizing CPU performance when adding many entries, e.g. millions of entries in single P4 tables in NICs

https://p4.org/specs

Architecture work group

- Focus is on defining Portable NIC Architecture
- Active participation from engineers at multiple organizations, including device vendors and users
- Plan to release an update to spec soon
  - Have been working through many details, e.g. add-on-miss tables and their use in features like connection tracking

Broadcom Intel Pensando Xilinx

Google Nvidia Vmware

Language design work group

- New release of version 1.2.3 of language specification, including
  - Improvements and fixes to the type system

- Andrew Pinski at Marvell Semiconductor
  - Actively working on a P4 compiler, implemented independently from the open source p4c front end.
  - Has been filing issues on many corner cases regarding the P4 language spec and p4c compiler (thank you!)

Open source

P4 compiler DPDK back end

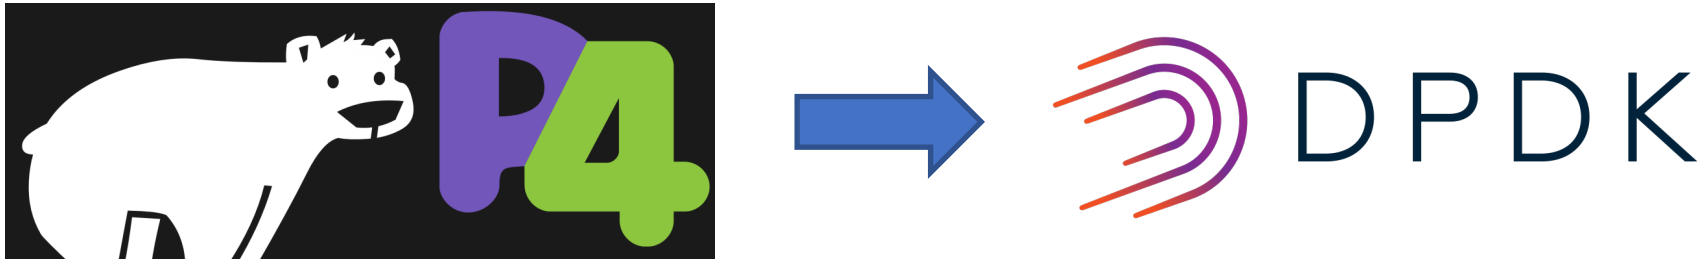

- Steady progress in the last year by a team of Intel engineers
- Implements the Portable Switch Architecture (PSA)
- Tracking the Portable NIC Architecture (PNA)

https://github.com/p4lang/p4c/blob/main/backends/dpdk/README.md
https://www.dpdk.org

P4 compiler eBPF back end

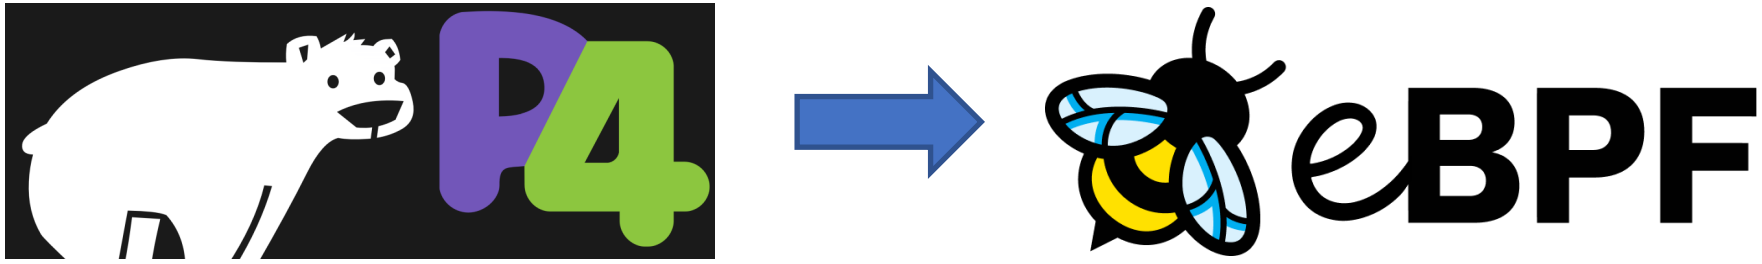

- Compiles P4 source to C source, that can be compiled to eBPF binary
  - The eBPF binary can be loaded into one of a few packet processing “hook points” in the Linux kernel
- Implements the Portable Switch Architecture (PSA)

https://github.com/p4lang/p4c/blob/main/backends/ebpf/README.md
https://ebpf.io

ONF has open sourced the entirety of its portfolio of production-ready platforms

- All of ONF's projects have now been released to open source with permissive software licenses.

- Projects
  - SD-RAN, SD-Core
  - Aether
  - SEBA, VOLTHA
  - SD-Fabric, P4, PINS, ONOS, Stratum

https://opennetworking.org/news-and-events/press-releases/onf-enters-a-new-era-focused-on-growing-adoption-and-community-for-its-leading-open-source-projects

P4 in academia

P4 in research

- SIGCOMM 2021
  - 1 hackathon featuring P4 on Raspberry Pi
  - 5 papers: algorithm implemented in P4
  - 2 more papers: P4 was a major part of the work
- NSDI 2021
  - 7 papers: algorithm implemented in P4
  - 1 more paper: P4 was a major part of the work
- https://p4.org/publications
  - 20 additional papers across journals, conferences, and workshops
  - 1 Ph.D. dissertation
- Intel Connectivity Research Program (ICRP - https://intel.com/icrp-papers)
  - 19 papers since May 2021

P4 in industry

“AMD completes acquisition of Xilinx”

Xilinx 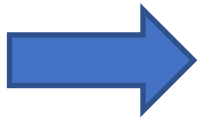 AMD

- Gordon Brebner
  - Former chair of P4 language design work group
  - Current chair of P4 architecture work group
  - Senior Fellow at AMD

https://www.amd.com/en/press-releases/2022-02-14-amd-completes-acquisition-xilinx

“AMD Expands Data Center Solutions Capabilities with Acquisition of Pensando”

Pensando 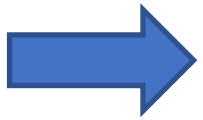 AMD

https://www.amd.com/en/press-releases/2022-04-04-amd-expands-data-center-solutions-capabilities-acquisition-pensando

DASH

- Disaggregated API for SONiC Hosts
  - “We are developing set of APIs and object models describing network services for the cloud.”
  - “The goal of DASH is to be specific enough for SMART Programmable Technologies to optimize network performance and leverage commodity **hardware** technology to achieve 10x or even 100x stateful connection performance.”

- While P4 is not a necessary part of DASH from a user’s perspective
  - Multiple vendors are interested in using P4 to define a reference model for the data plane behavior.
  - Some vendors may use P4 “under the hood” to implement compatible products.

https://github.com/Azure/DASH

Alibaba

- **Sailfish**: a cloud-scale multi-tenant multi-service gateway accelerated by programmable switches

- **Aquila**: verification system for production-scale programmable data planes
  - "Alibaba has a global network infrastructure to support its worldwide online services … which have more than one billion users. … To offer high throughput (Tbps speeds) and save CPU resources, our edge networks have widely deployed programmable switches to offload a group of network functions (e.g. load balancing, firewall and DDoS defense) from software to programmable switching ASIC hardware."

- Both at SIGCOMM 2021, in collaboration with universities

https://dl.acm.org/doi/abs/10.1145/3452296.3472889
https://dl.acm.org/doi/10.1145/3452296.3472937

Other industry & university collaborations

- **Tiara**: Stateful L4 load balancer (NSDI 2022)
  - ByteDance

- Latest progress on SRv6 open-source project to be demonstrated in this May's P4 workshop
  - China Mobile

https://www.usenix.org/conference/nsdi22/presentation/zeng

P4 on YouTube

ServeTheHome channel on YouTube

- Several recent videos mentioning P4
  - In the context of many other features and capabilities of devices like Nvidia Bluefield-2, Pensando Distributed Services Card, and others

- “Why AMD is Spending $1.9B to Buy Pensando for DPUs”
- “ZFS without a Server!?! It is DPU time!”
- “Intel Launches Next-Gen Salvo at NVIDIA & AMD with BIG New Chips”

https://www.youtube.com/watch?v=XVOvs8bsyFo
https://www.youtube.com/watch?v=NPQE0Rzh52c
https://www.youtube.com/watch?v=IhTeZj1HSCI

Jump on in!

- See calendar of meetings for anything that interests you
  - https://p4.org/working-groups
- Join mailing lists
  - https://p4.org/join

- Technical things coming up
  - P4 namespace/module system being discussed
  - Defining P4 architectures precisely in code, vs. natural language
  - Alan Lo at Nvidia proposing new syntax for defining state machines in P4

Intel notices and disclaimers

- Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.
-

Thank You