Mapping tc to P4

TC Workshop, November 10, 2017

NetDev 2.2, Seoul, Korea

Prem Jonnalagadda, Barefoot Networks

P4 Components

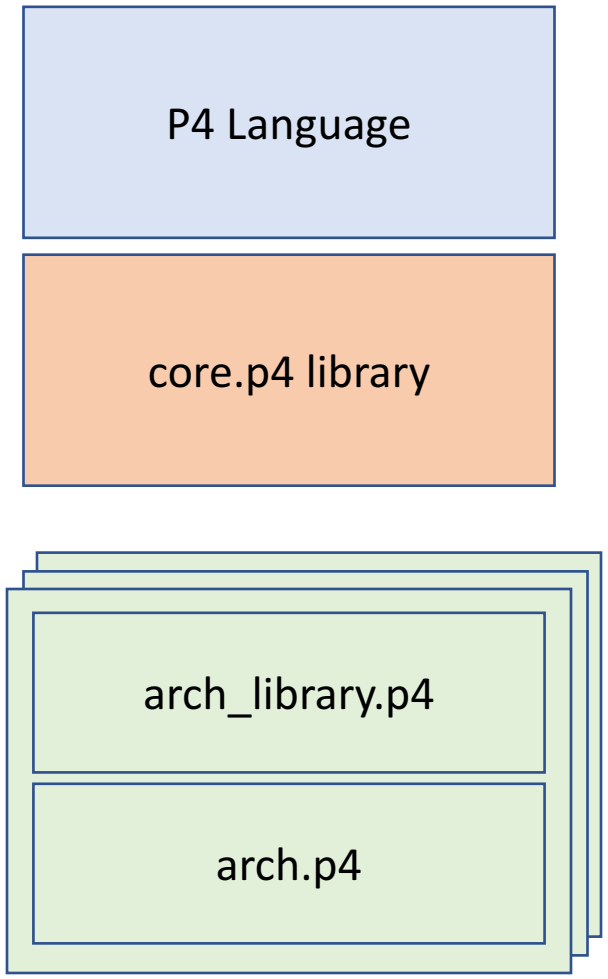

- Stable
- Open
- Community owned

- Target specific
- Provided by Vendor or Community

Protocol Independent Switch Architecture

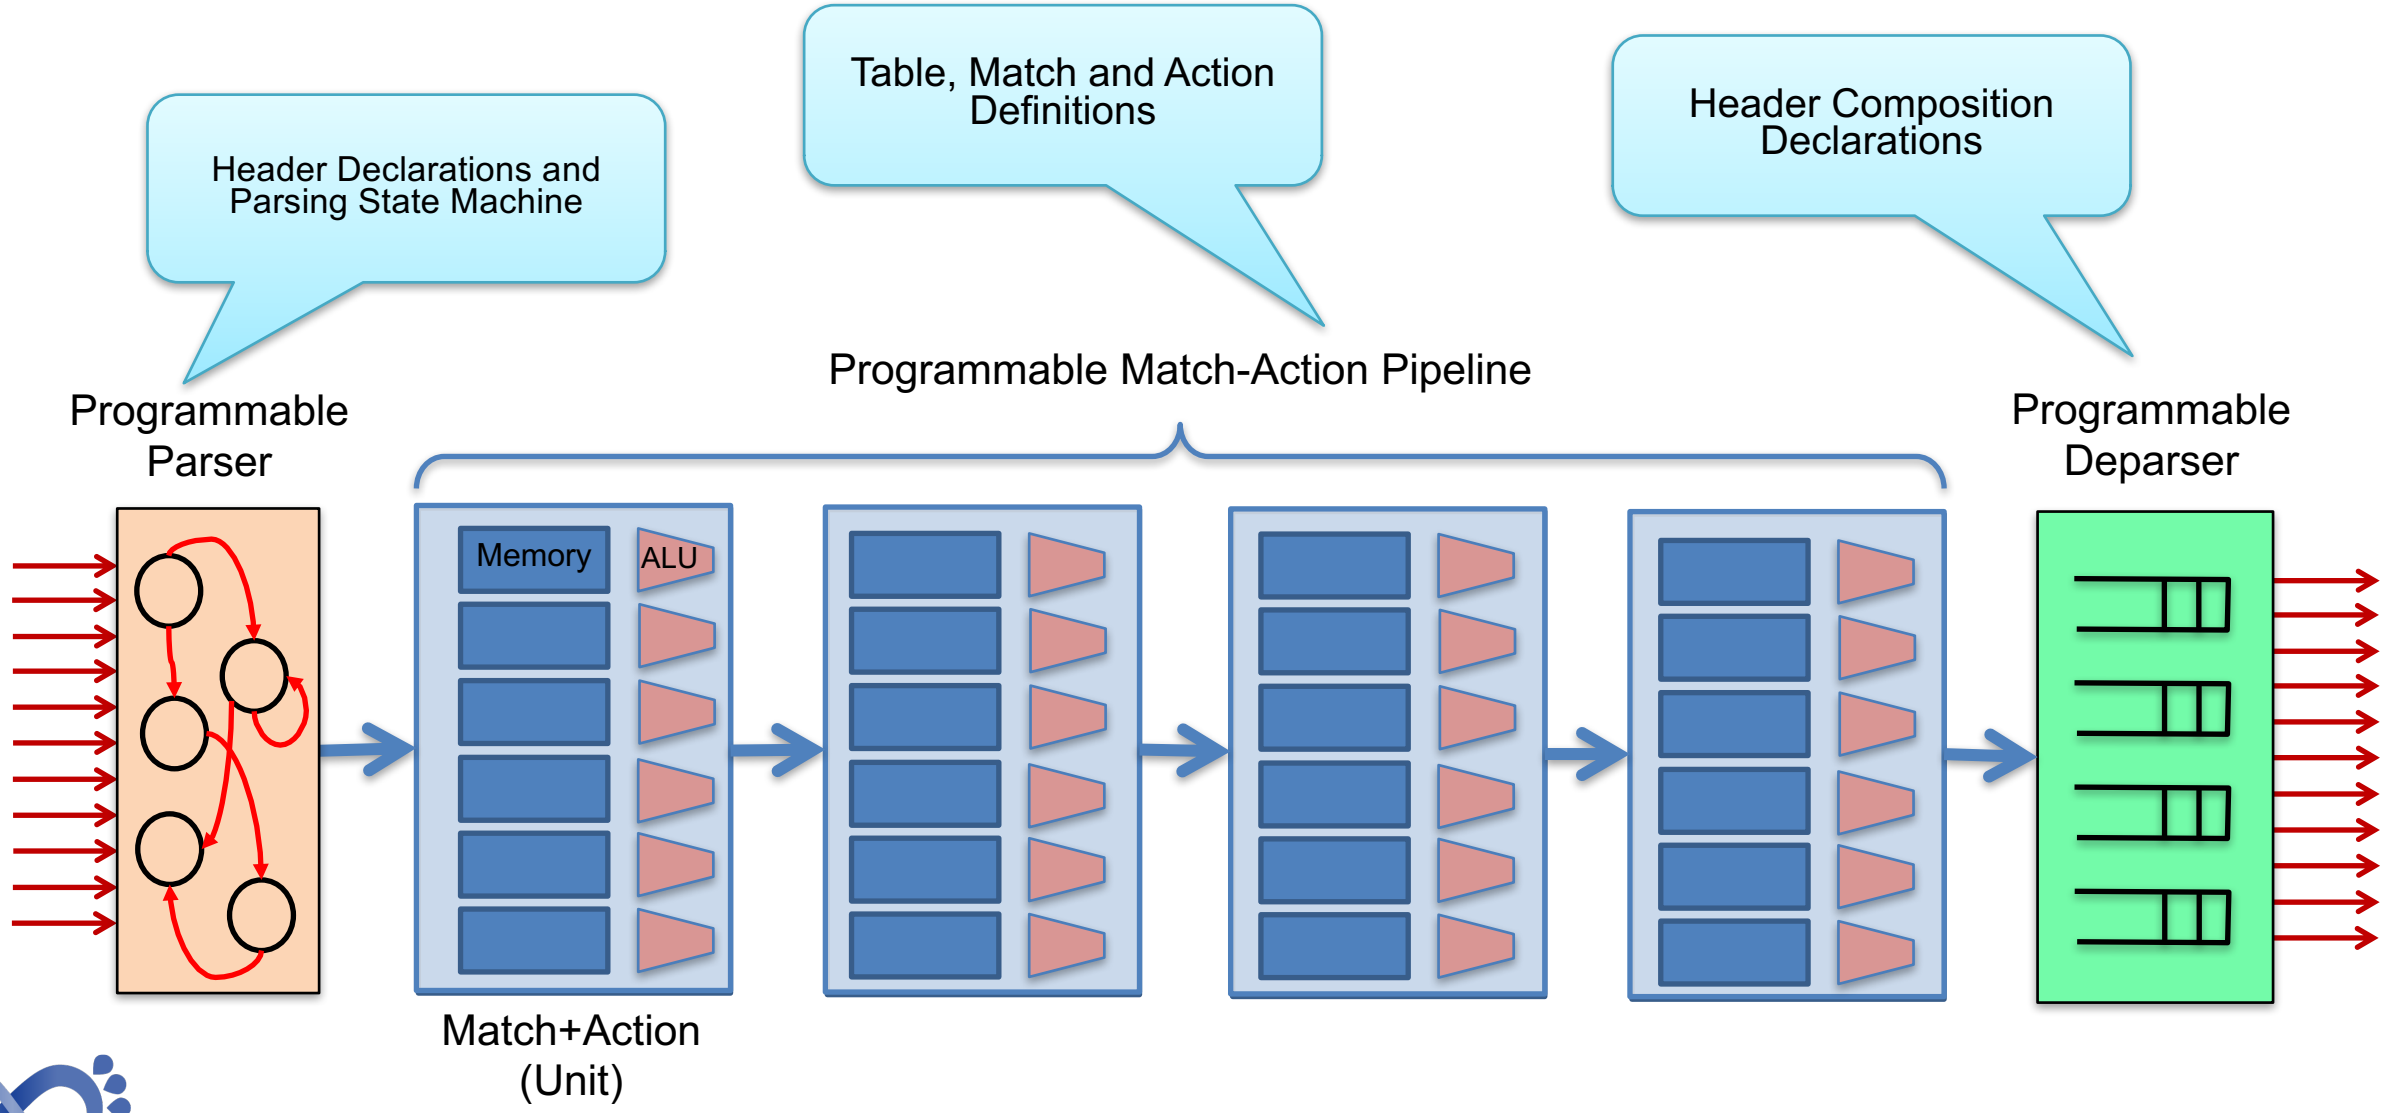

Example P4 Program

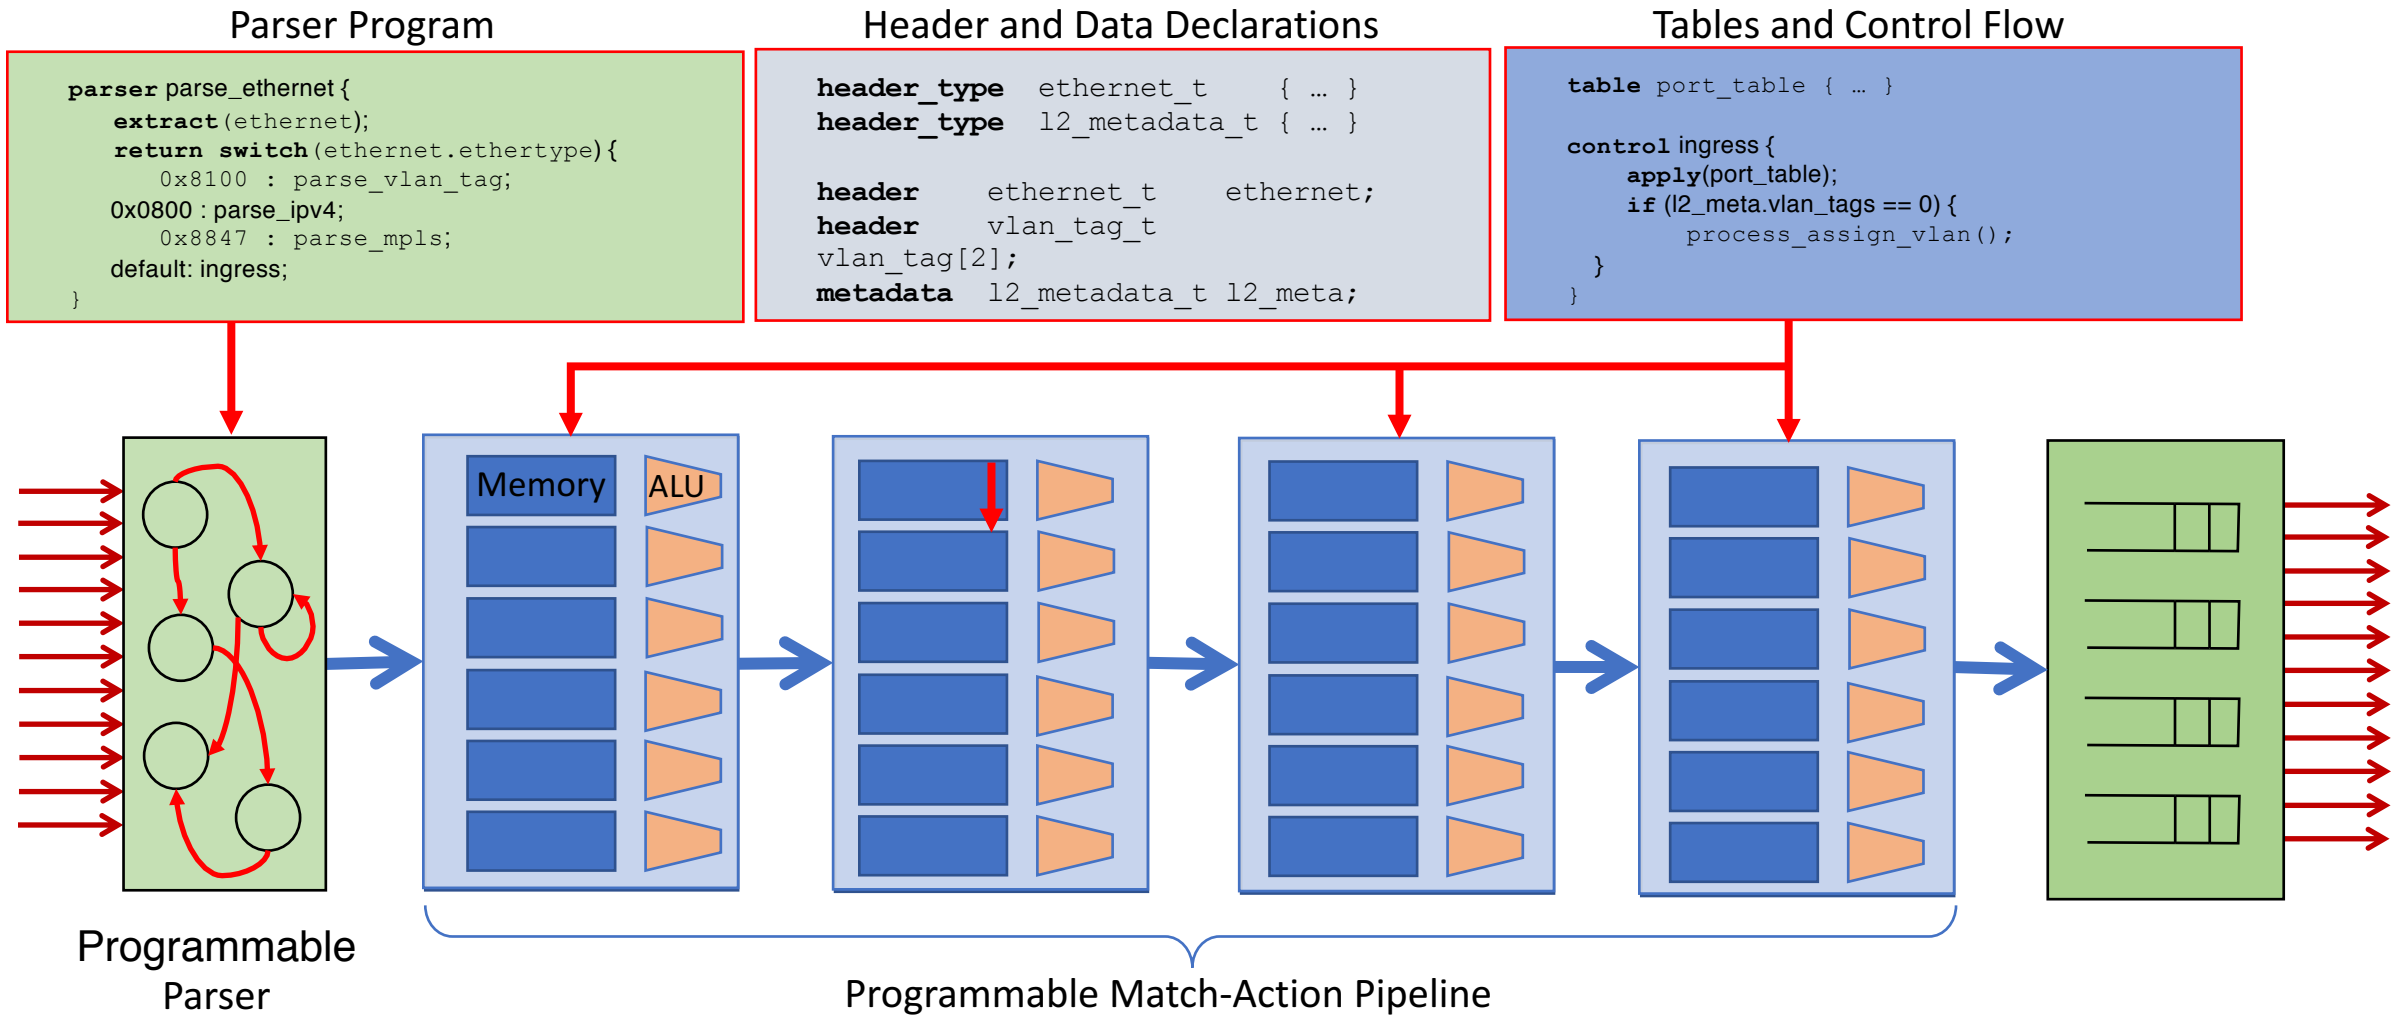

Tofino Pipeline

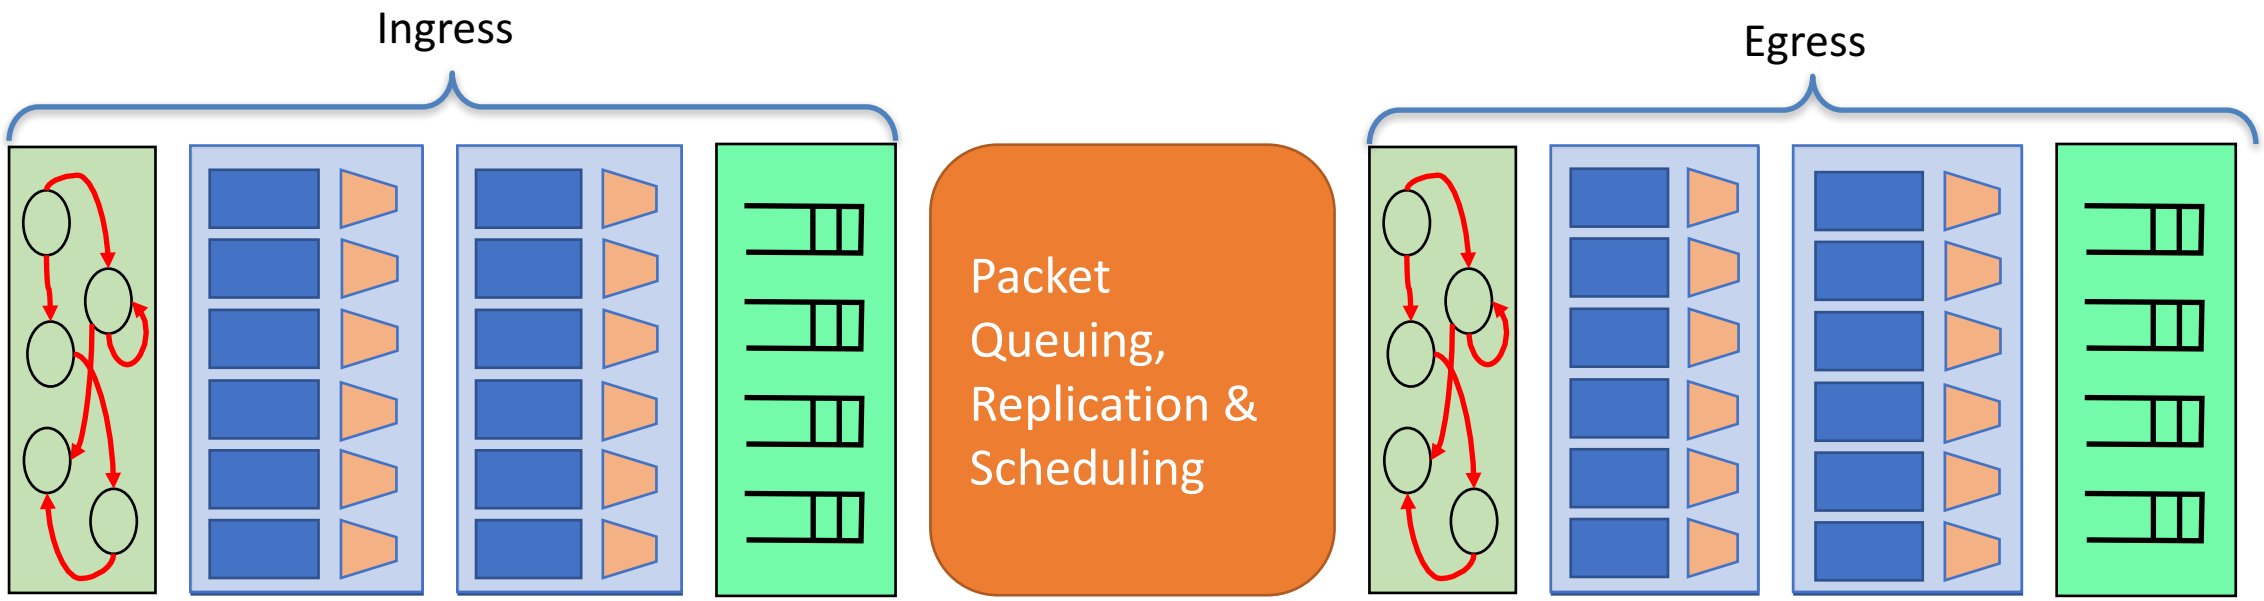

For discussion relative to tc pipeline

P4 -> tc -> Switch

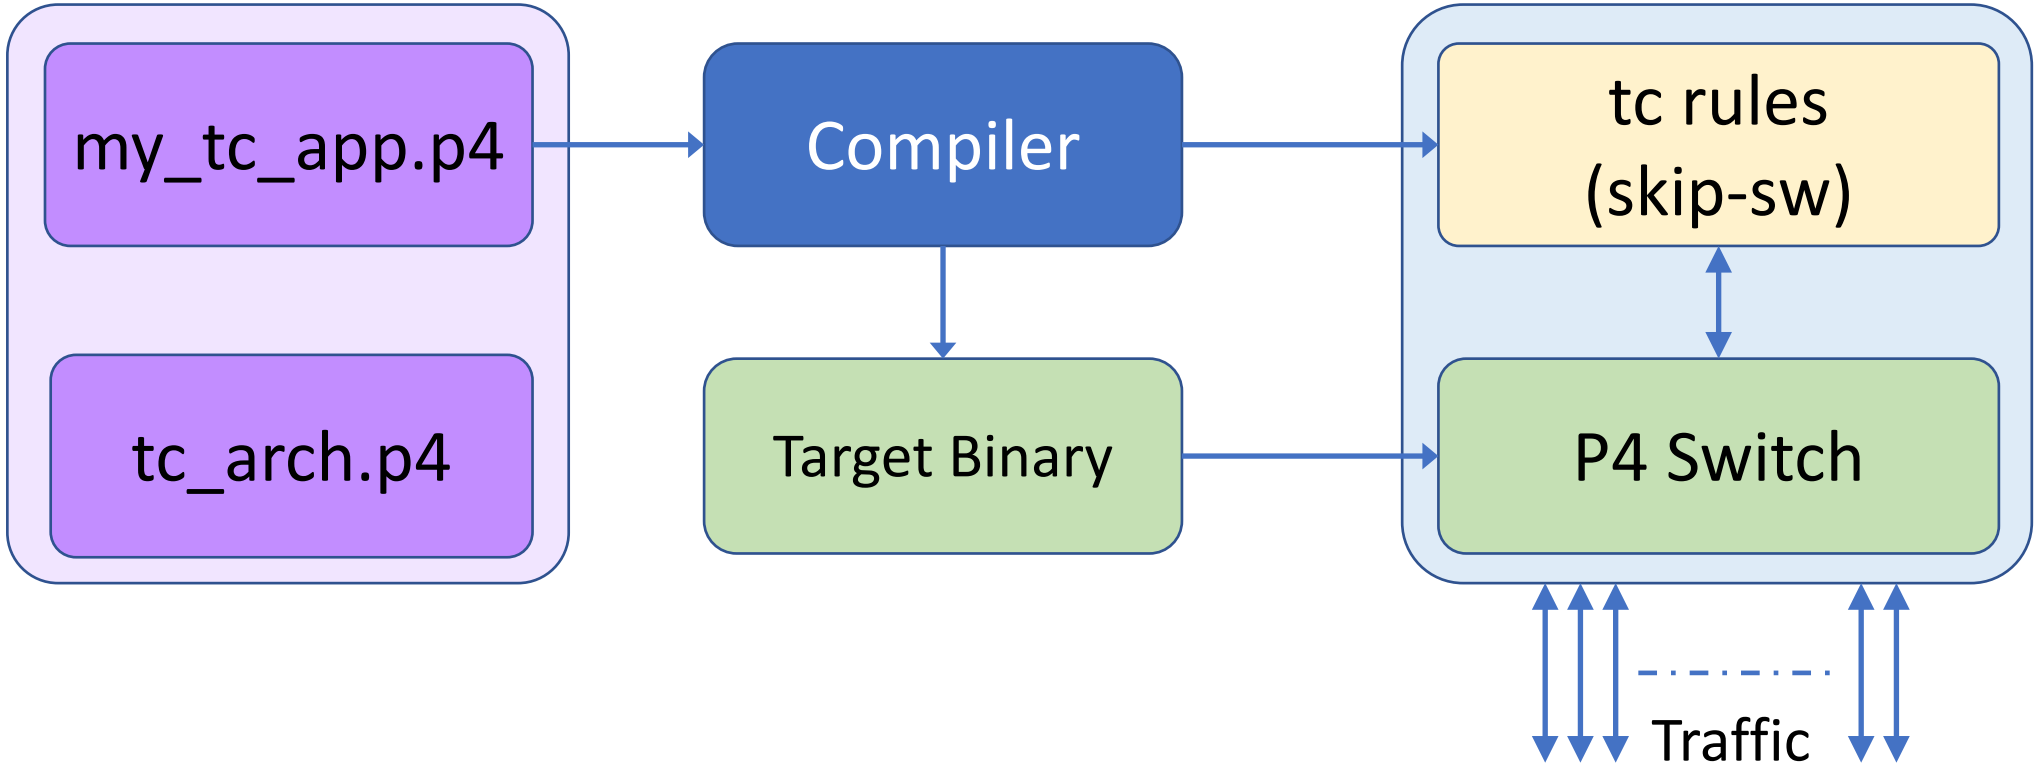

6.5Tb/s Tofino™ Summary

- State of the art design
  - Single Shared Packet Buffer
  - TSMC 16nm FinFET+
- Four Match+Action Pipelines
  - Fully programmable PISA Embodiment
  - All compiled programs run at line-rate.
- Port Configurations
  - 65 x 100GE/40GE
  - 130 x 50GE
  - 260 x 25GE/10GE
- CPU Interfaces
  - PCIe: Gen3 x4/x2/x1
  - Dedicated 100GE port

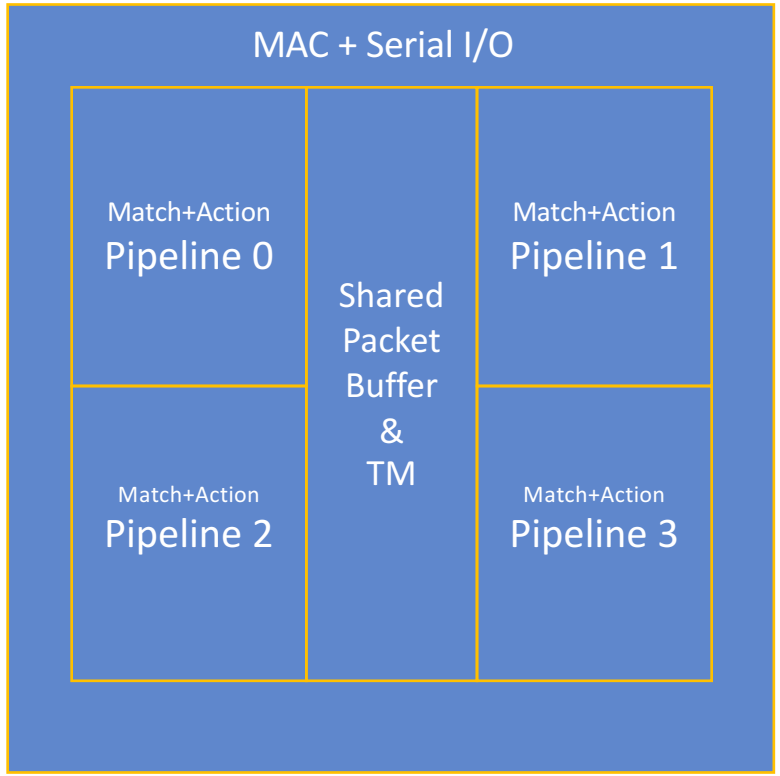

Tofino. Simplified Block Diagram

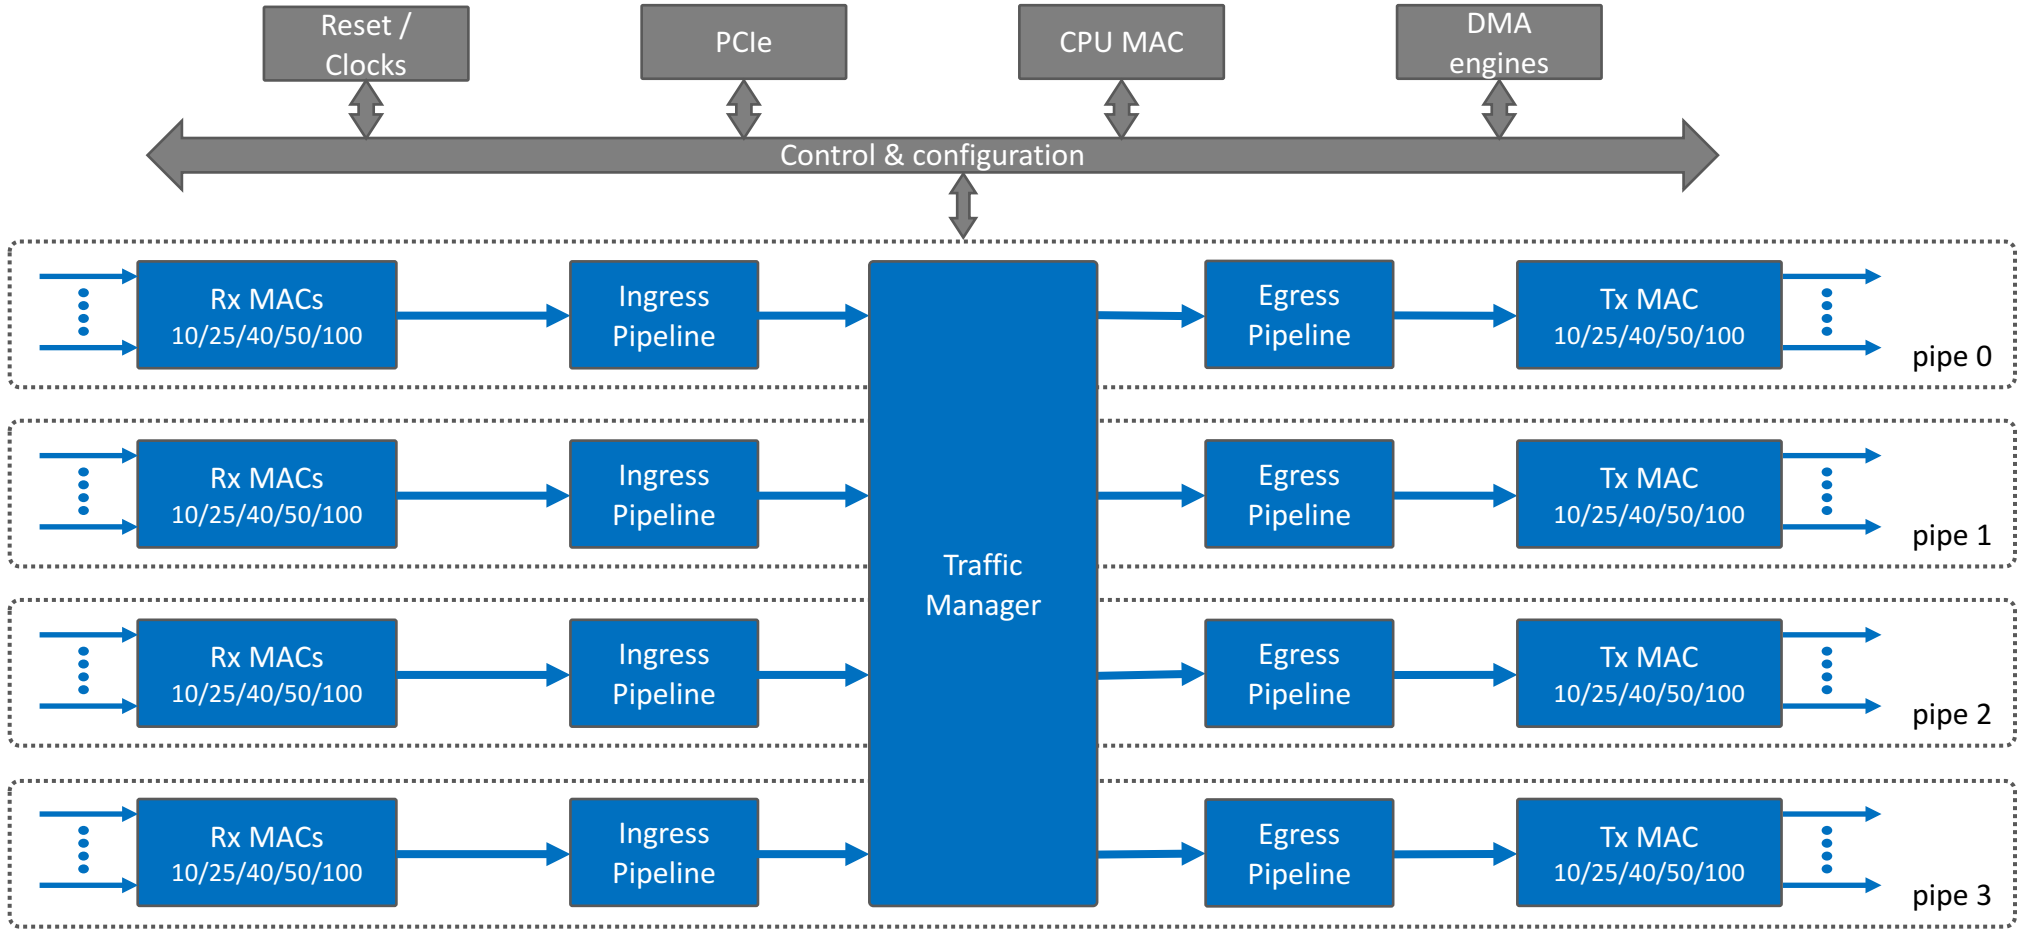

Future Work

- Based on discussions and conclusions in the workshop.

P4 Information, Links & Resources

- Programming Protocol-independent Packet Processors – P4
- A high-level networking domain-specific programming language
- Governed by P4 Language Consortium – an independent California Non-Profit organization
- Apache 2.0 CLA
- Language Spec - https://p4lang.github.io/p4-spec/docs/
- GitHub repo – https://github.com/p4lang
- Compiler - https://github.com/p4lang/p4c
- Community – http://lists.p4.org

Thanks

prem [at] barefootnetworks [dot] com