


# PISCES: A P4-Enabled OVS

Muhammad Shahbaz, Cian Ferriter


Princeton, Intel

P4


# PISCES: A P4-Enabled OVS



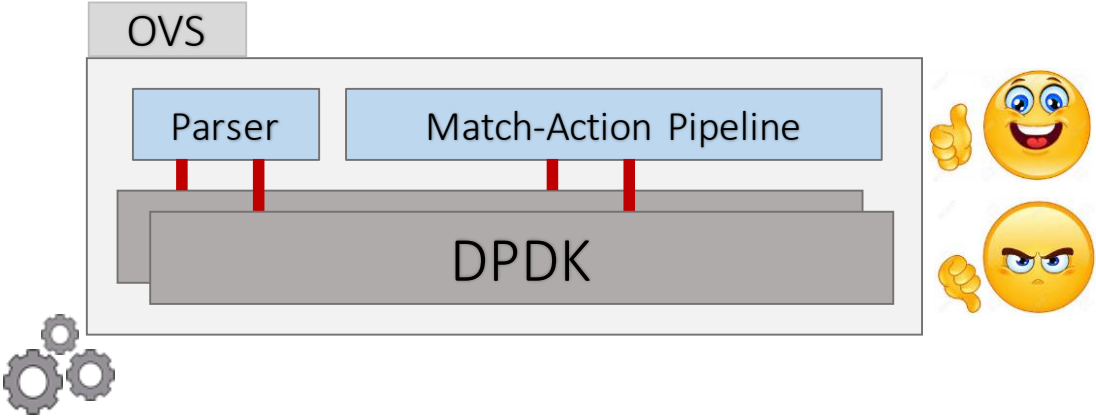
# Internal Architecture of OVS




# Internal Architecture of OVS




# Internal Architecture of OVS



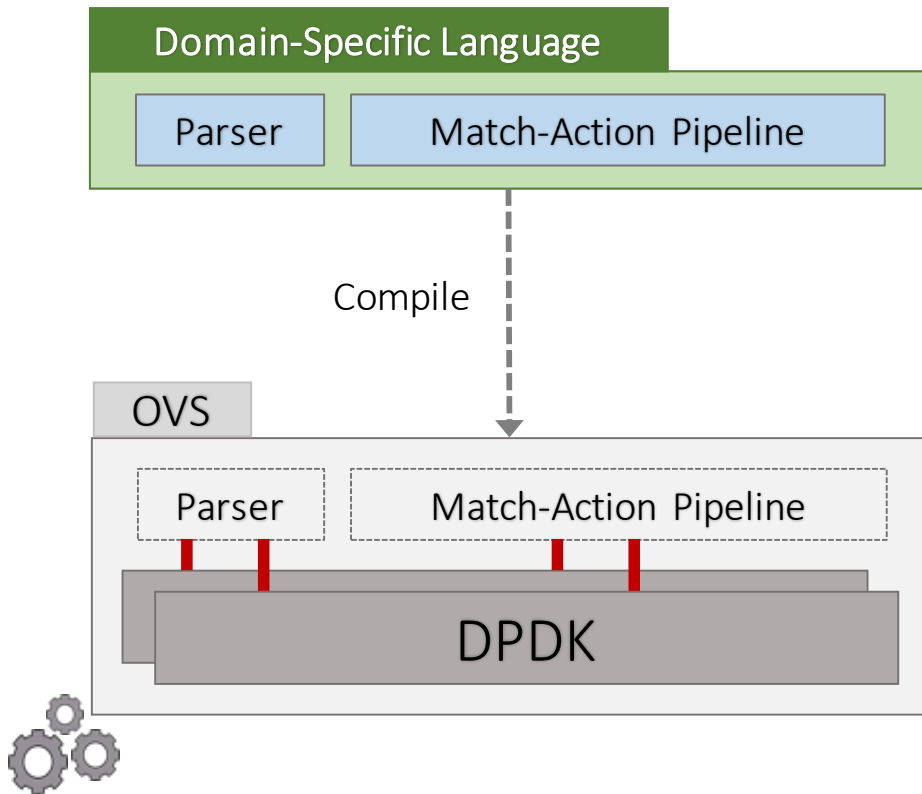
# Internal Architecture of OVS



# Internal Architecture of OVS

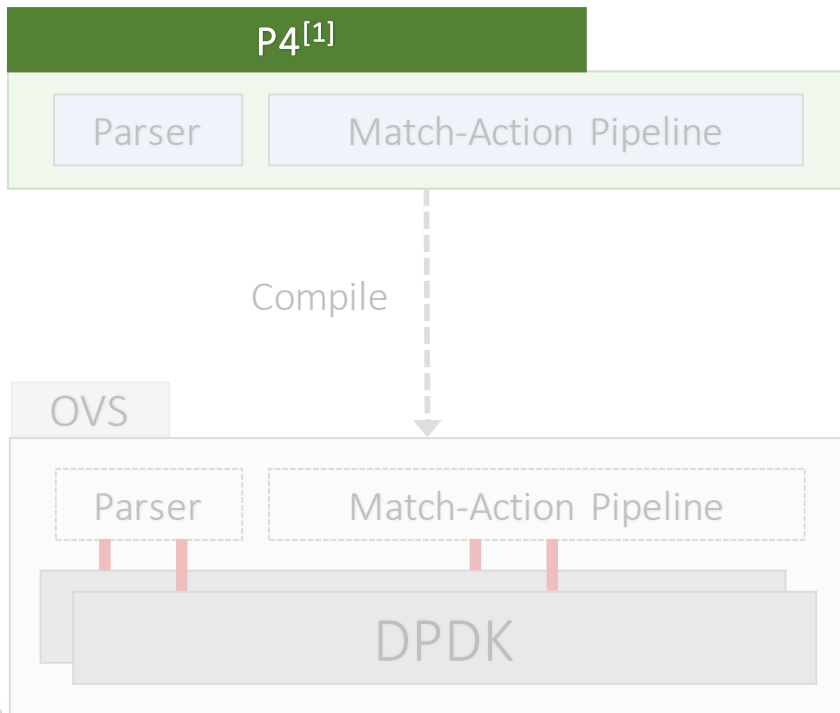


# Road to Protocol Independence





# Road to Protocol Independence



P4 is an **open-source language**.<sup>[1]</sup>


Describes different aspects of a packet processor:

- **Packet headers and fields**
- **Metadata**
- **Parser**
- **Actions**
- **Match-Action Tables (MATs)**
- **Control Flow**

<sup>[1]</sup> <http://www.p4.org>



# Road to Protocol Independence



341 lines of code

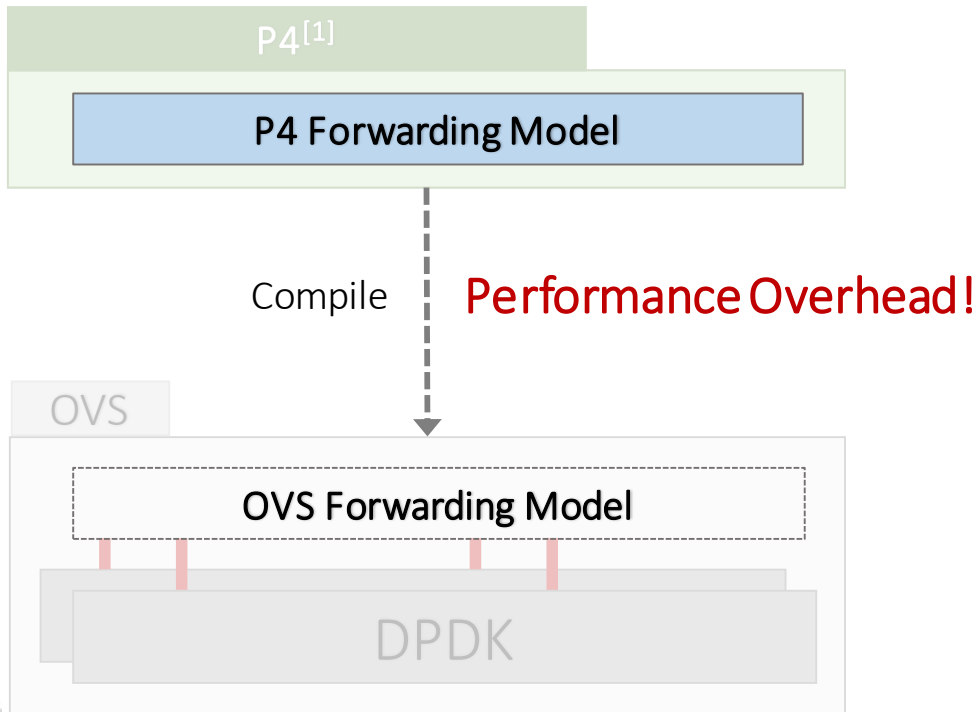
Native OVS

14,535 lines of code




<sup>[1]</sup> <http://www.p4.org>

# Road to Protocol Independence




[1] <http://www.p4.org>


# P4 Forwarding Model (Post-Pipeline Editing)




Header Fields




# OVS Forwarding Model



# OVS Forwarding Model




# OVS Forwarding Model (Inline Editing)



# PISCES Forwarding Model (Modified OVS)

- Supports both editing modes:
  - **Inline Editing**
  - **Post-pipeline Editing**






# PISCES: Compiling P4 to OVS

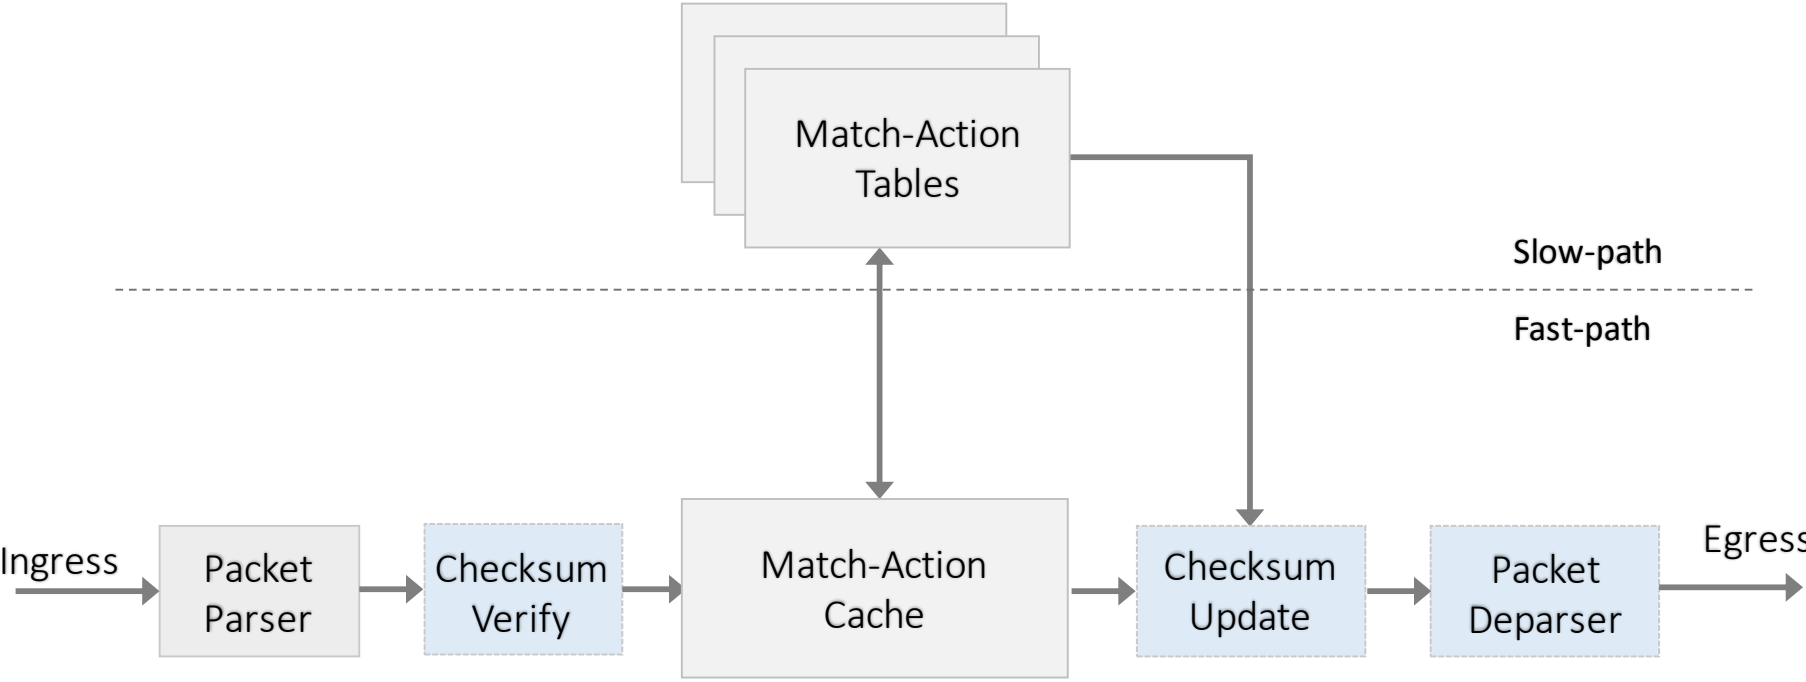
P4




modified  
OVS



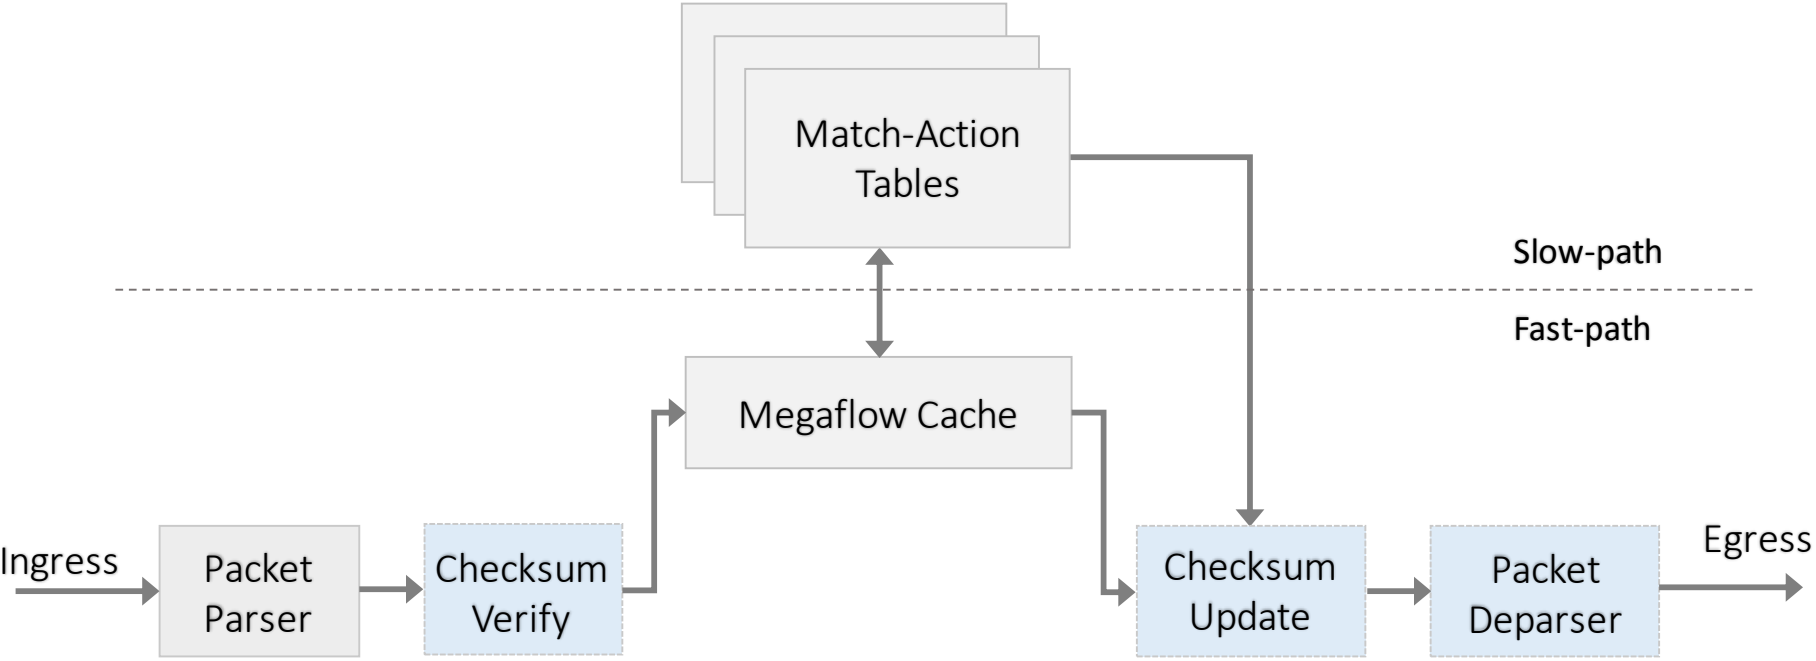
# PISCES Forwarding Model (Modified OVS)




# PISCES Forwarding Model (Modified OVS)



# PISCES Forwarding Model (Modified OVS)

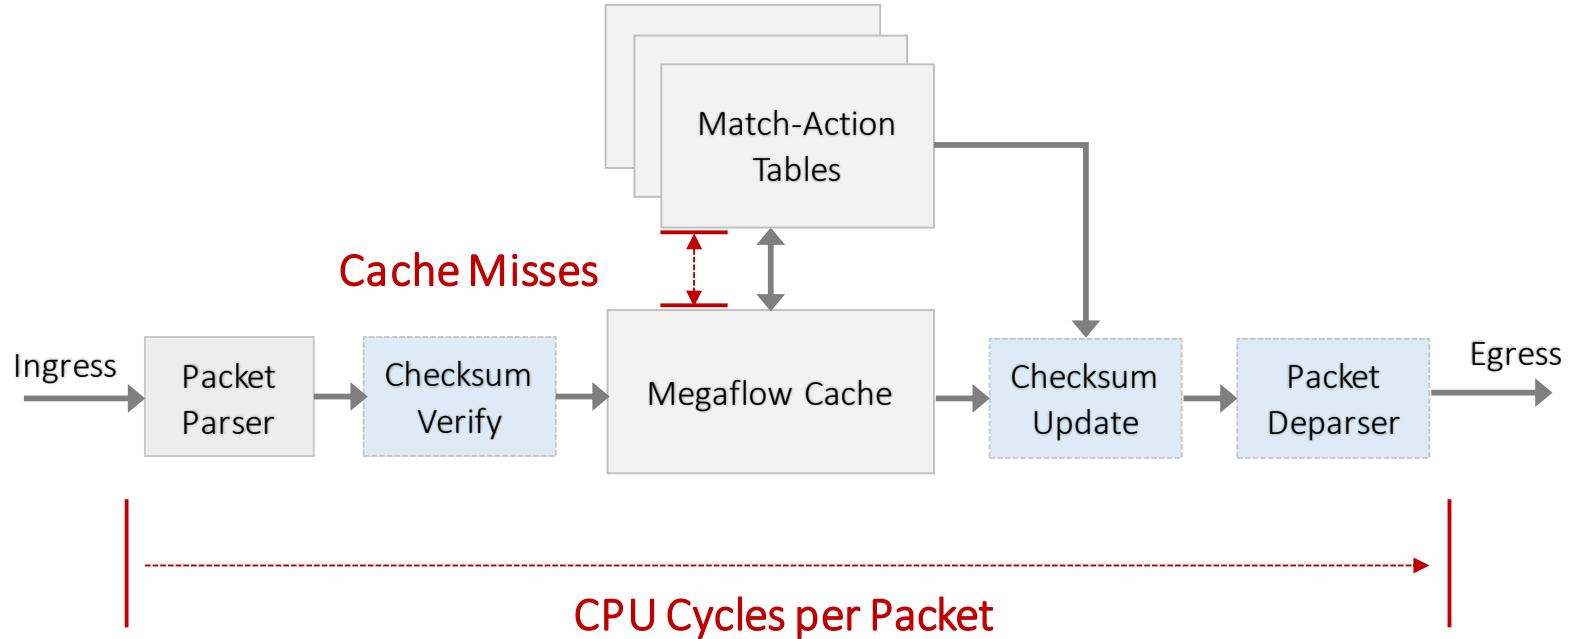


# Naïve Compilation from P4 to OVS (L2L3-ACL)



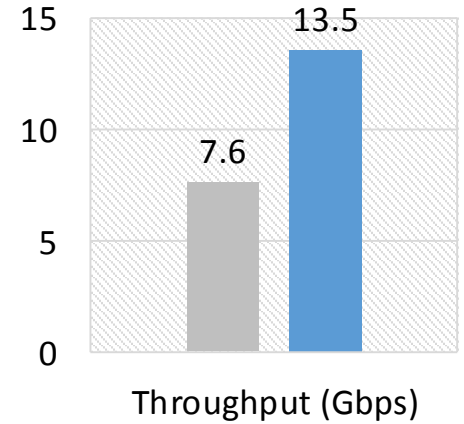
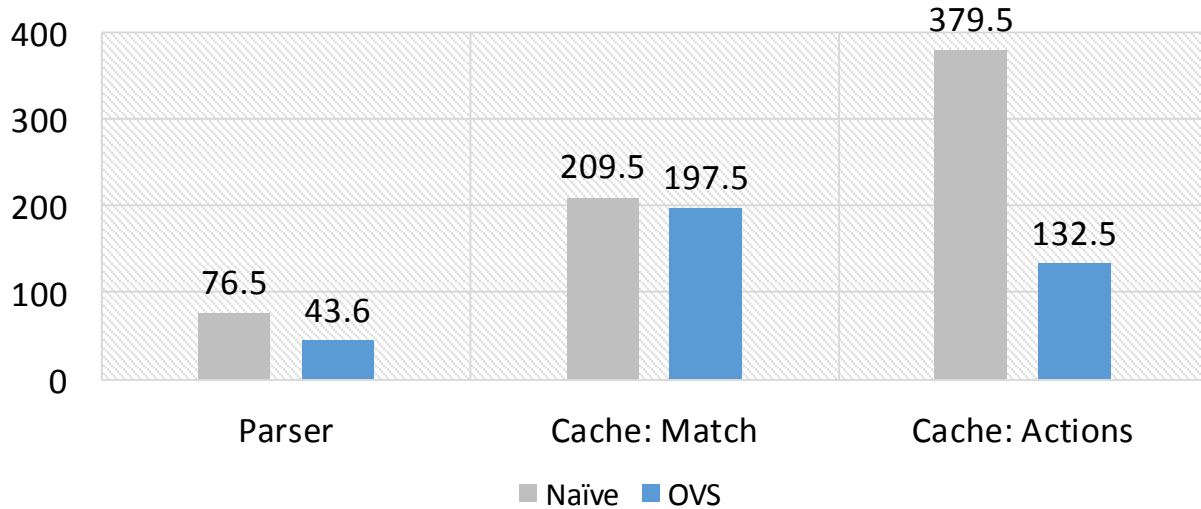
Performance overhead of  
~ 40%

# Causes of Performance Overhead



# Cause: CPU Cycles per Packet

L2L3-ACL (CPU Cycles for a 64 Byte Packet)




# Factors affecting CPU Cycles per Packet


- a. Extra copy of headers
  - b. Fully-specified Checksum
  - c. Parsing unused header fields
- and more ...



# Different Optimizations for L2L3-ACL




# Optimized Compilation from P4 to OVS (L2L3-ACL)



Performance overhead of  
**< 2%**

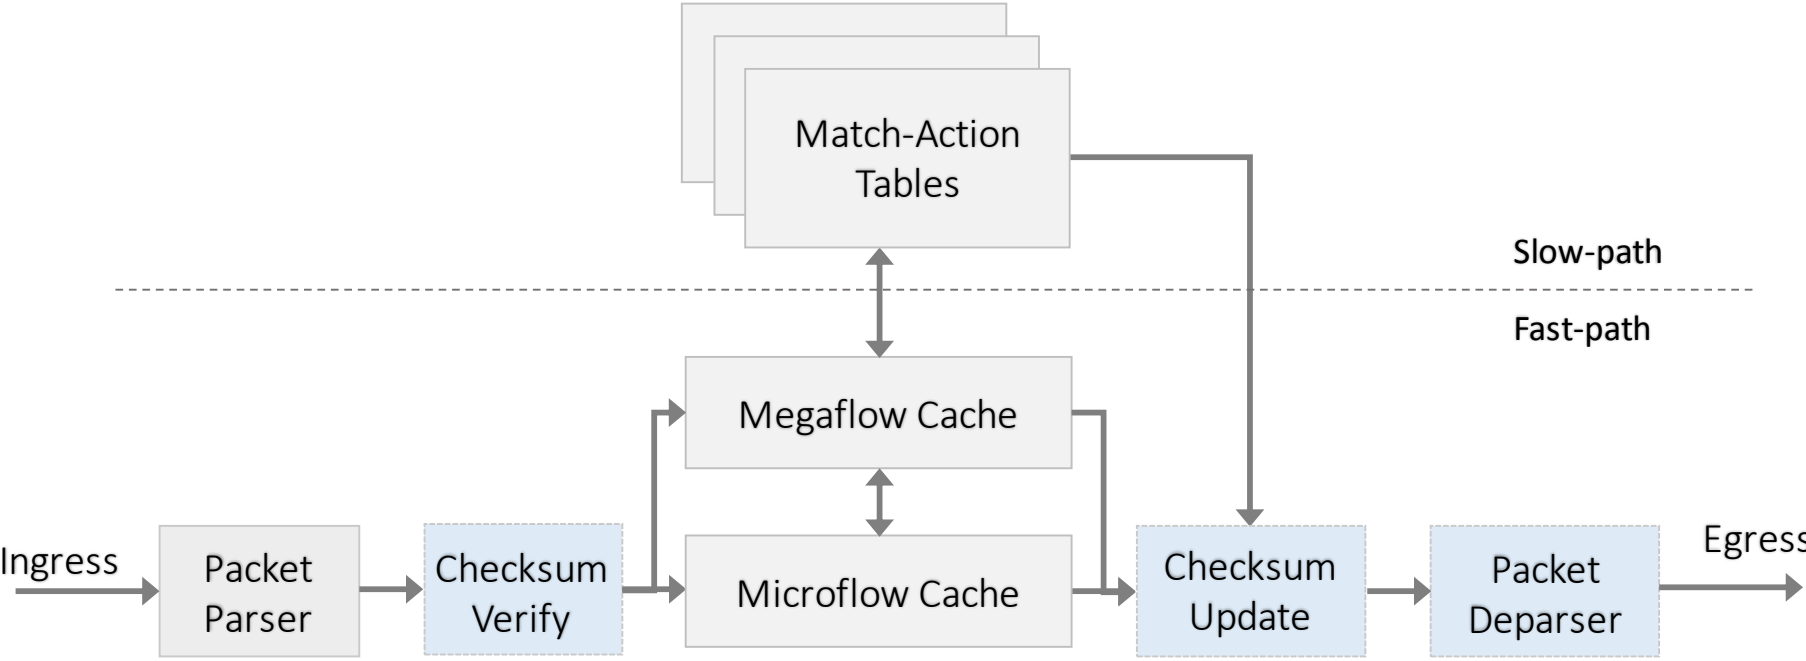
# Cause: Cache Misses



# Factors affecting Cache Misses

- a. Entropy of packet header fields
- b. Stateful operations in the match-action cache (or fast path).

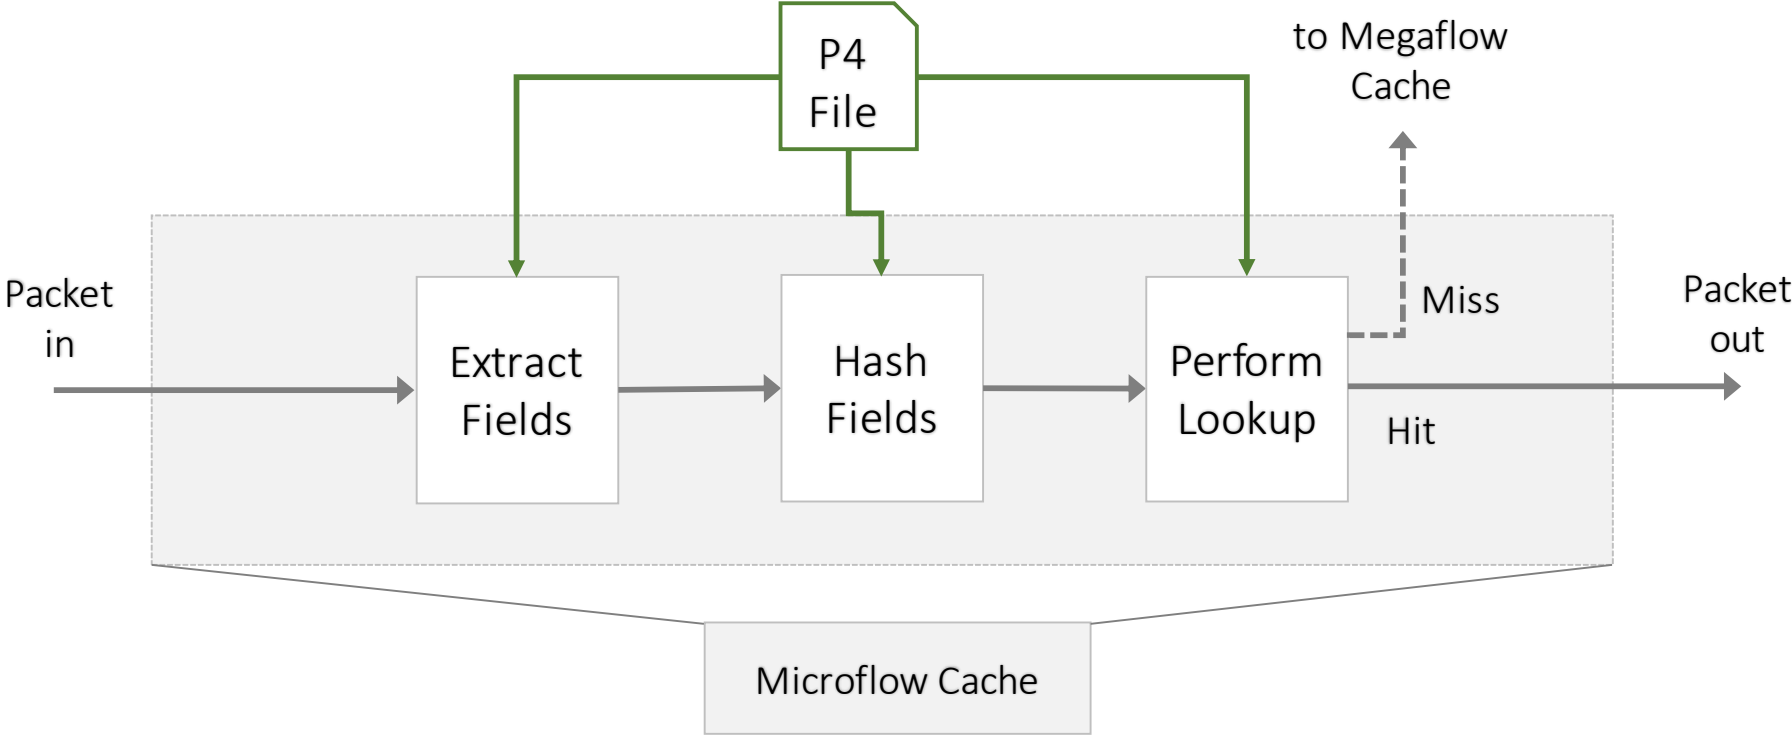
# PISCES Forwarding Model (Modified OVS)



# PISCES Forwarding Model (Modified OVS)


Microflow Cache

# Internals of the Microflow Cache



# Performance with the Microflow Cache

Phy-Phy, L3 Router Case, 64B






# Cause of Performance Degradation

Cacheline	64 Bytes		
0	Metadata		
1	Metadata		Ethernet Header
2	IPv4 (1 <sup>st</sup> 16Bytes)	IPv4 + L4 Proto	Empty

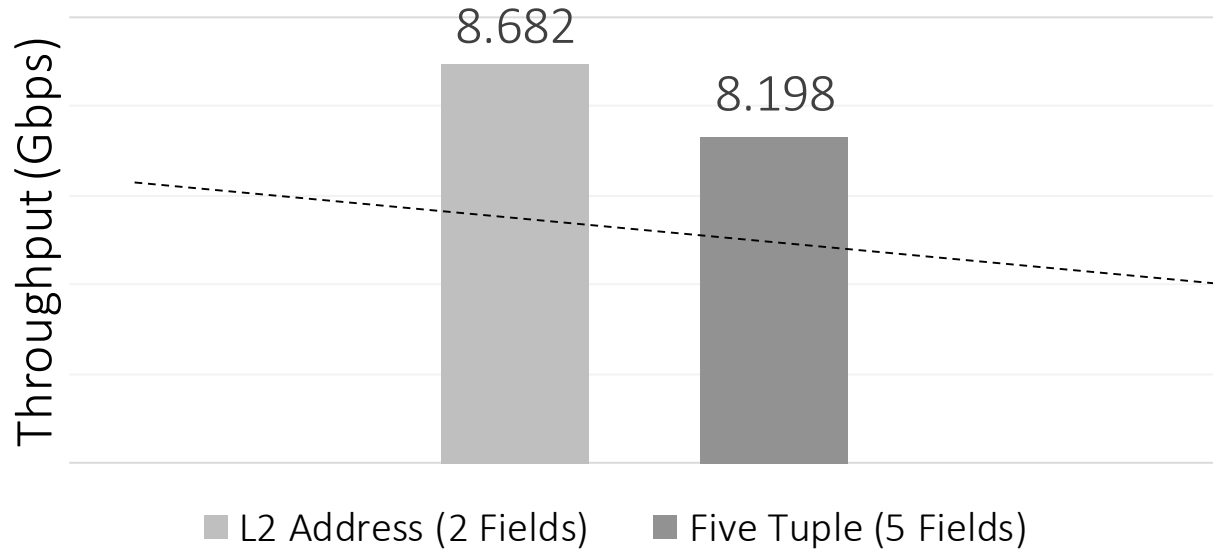
Simplified “flow” Structure

# Performance with the Microflow Cache

Phy-Phy, L3 Router Case, 64B



# Varying the Number of Hash Fields



Questions?

# Disclaimers

Intel does not control or audit third-party benchmark data or the web sites referenced in this document. You should visit the referenced web site and confirm whether referenced data are accurate.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. **No computer system can be absolutely secure.** Check with your system manufacturer or retailer or learn more at [[intel.com](https://www.intel.com)].