

P4 Compiler in SDN

Federico Bruzzone,¹ PhD Student

Milan, Italy – 16 December 2024

Slides available at

federicobruzzone.github.io/activities/presentations/P4-compiler-in-SDN.pdf

¹ADAPT Lab – Università degli Studi di Milano,

Website: federicobruzzone.github.io,

Github: github.com/FedericoBruzzone,

Email: federico.bruzzone@unimi.it

Network Programmability

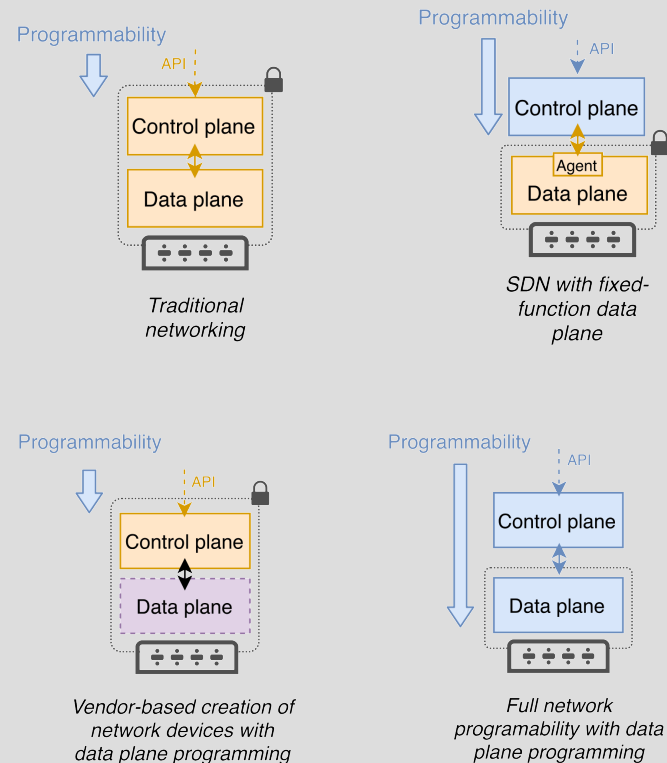
The ability of the software or the hardware to execute an externally defined processing algorithm [1]

Open Networking Foundation (ONF)

- Non-profit consortium founded in 2011
- Promotes networking through **Software Defined Networking (SDN)**
- Standardizes the **OpenFlow** protocol

Software Defined Networking (SDN)

- Born to overcome the limitations of traditional network architectures
- Decouples the control plane from the data plane
- Centralizes the control of the network



OpenFlow Protocol

- Gives access to the **forwarding plane** (data plane) of a network device
- Mainly used by switches and controllers
- Layered on top of the **Transport Control Protocol (TCP)**
- De-facto standard for SDN

OpenFlow Development

- First appeared in 2008 [2]
- In April 2012, Google deploys OpenFlow in its internal network with significant improvements (Urs Hölzle promotes it²)
- In January 2013, NEC rolls out OpenFlow for Microsoft Hyper-V
- Latest version is 1.5.1 (Apr 2015)

²[Inter-Datacenter WAN with centralized TE using SDN and OpenFlow.](#)

Fields in OpenFlow Standard

Version	Date	Header Fields
OF 1.0	Dec 2009	12 fields (Ethernet, TCP/IPv4)
OF 1.1	Feb 2011	15 fields (MPLS, inter-table metadata)
OF 1.2	Dec 2011	36 fields (APR, ICMP, IPv6, etc.)
OF 1.3	Jun 2012	40 fields
OF 1.4	Oct 2013	41 fields

More Details on the OpenFlow v1.0.0 Switch Specification³

³<https://opennetworking.org/wp-content/uploads/2013/04/openflow-spec-v1.0.0.pdf>

OpenFlow is protocol-dependent

Fixed set of fields and parser based on standard protocols

(Ethernet, IPv4/IPv6, TCP/UDP)

P4: Programming Protocol-Independent Packet Processors

Bosshart believed that future generations of OpenFlow would have allowed the controller to *tell the switch how to operate* [3]

Goals and Challenges

New Control Plane Specification: P4Runtime for controlling the data plane elements of a device defined by a P4 program

Reconfigurability: the controller should be able to redefine the packet parsing and processing in the field

Protocol Independence: the switch should process *headers* using parsing and processing using *match+action* tables

Target Independence: a compiler from *target-independent* description to *target-dependent* binary

Goals and Challenges

New Control Plane Specification: P4Runtime for controlling the data plane elements of a device defined by a P4 program

Reconfigurability: the controller should be able to redefine the packet parsing and processing in the field

Protocol Independence: the switch should process *headers* using parsing and processing using *match+action* tables

Target Independence: a compiler from *target-independent* description to *target-dependent* binary

Goals and Challenges

New Control Plane Specification: P4Runtime for controlling the data plane elements of a device defined by a P4 program

Reconfigurability: the controller should be able to redefine the packet parsing and processing in the field

Protocol Independence: the switch should process *headers* using parsing and processing using *match+action* tables

Target Independence: a compiler from *target-independent* description to *target-dependent* binary

Goals and Challenges

New Control Plane Specification: P4Runtime for controlling the data plane elements of a device defined by a P4 program

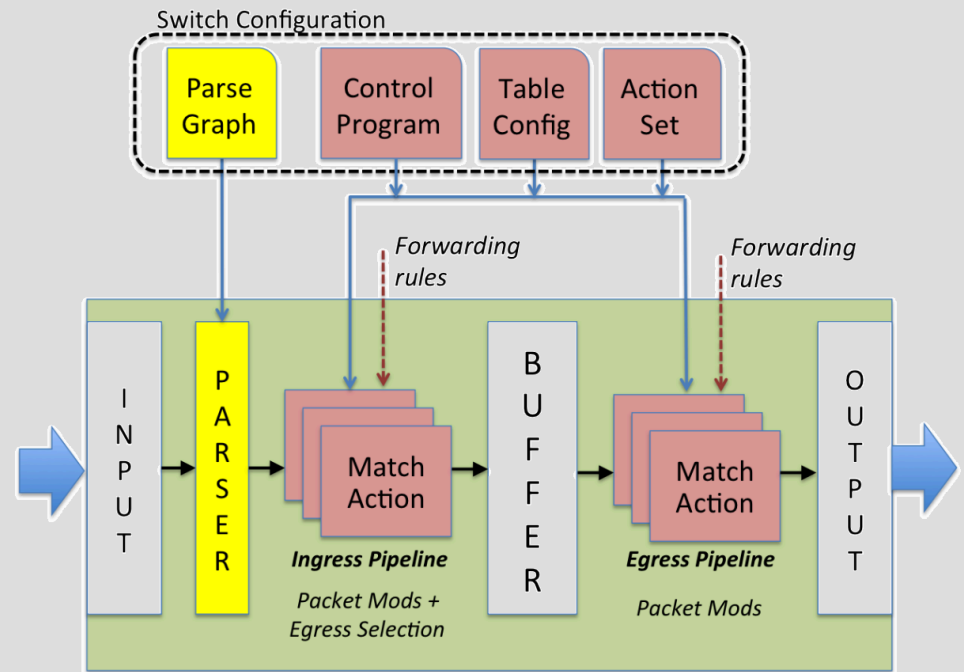
Reconfigurability: the controller should be able to redefine the packet parsing and processing in the field

Protocol Independence: the switch should process *headers* using parsing and processing using *match+action* tables

Target Independence: a compiler from *target-independent* description to *target-dependent* binary

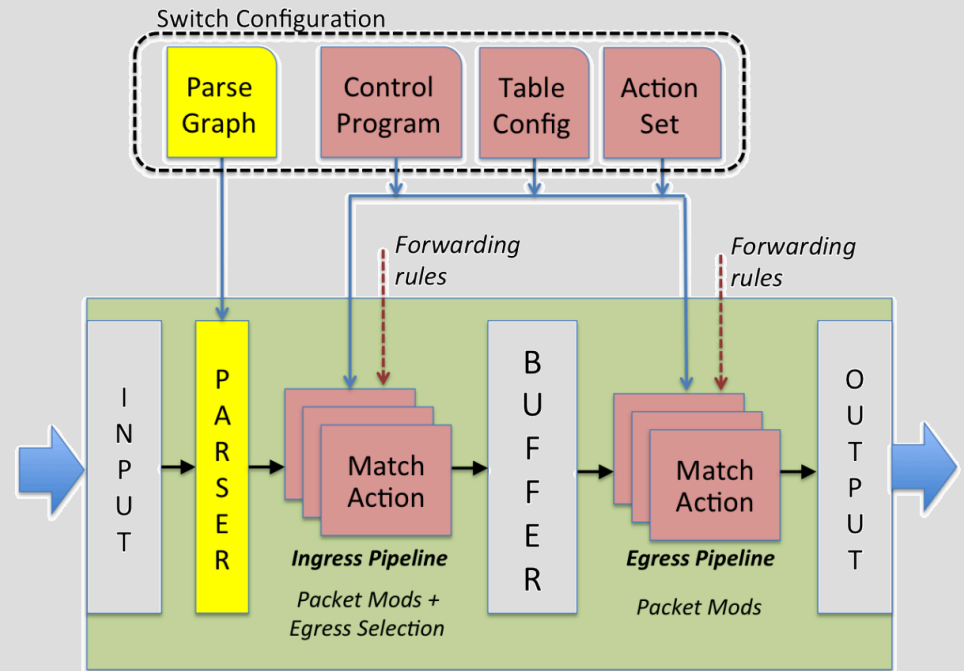
Abstract Forwarding Model (AFM)

1. Parsing the packet headers
2. The fields are passed to the match-action pipeline.
 - **Ingress:** determines the egress port/queue
 - **Egress:** per-instance header modifications
3. Metadata processing (e.g., timestamp)
4. As in OpenFlow, the queuing discipline is chosen at switch configuration time (e.g., minimum rate)



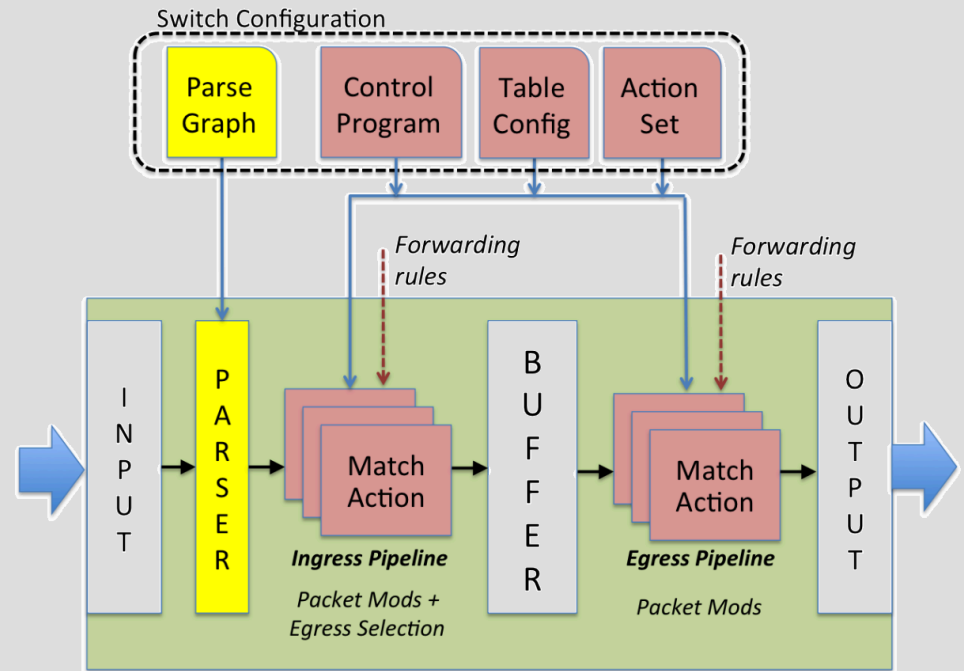
Abstract Forwarding Model (AFM)

1. Parsing the packet headers
2. The fields are passed to the match-action pipeline.
 - **Ingress:** determines the egress port/queue
 - **Egress:** per-instance header modifications
3. Metadata processing (e.g., timestamp)
4. As in OpenFlow, the queuing discipline is chosen at switch configuration time (e.g., minimum rate)



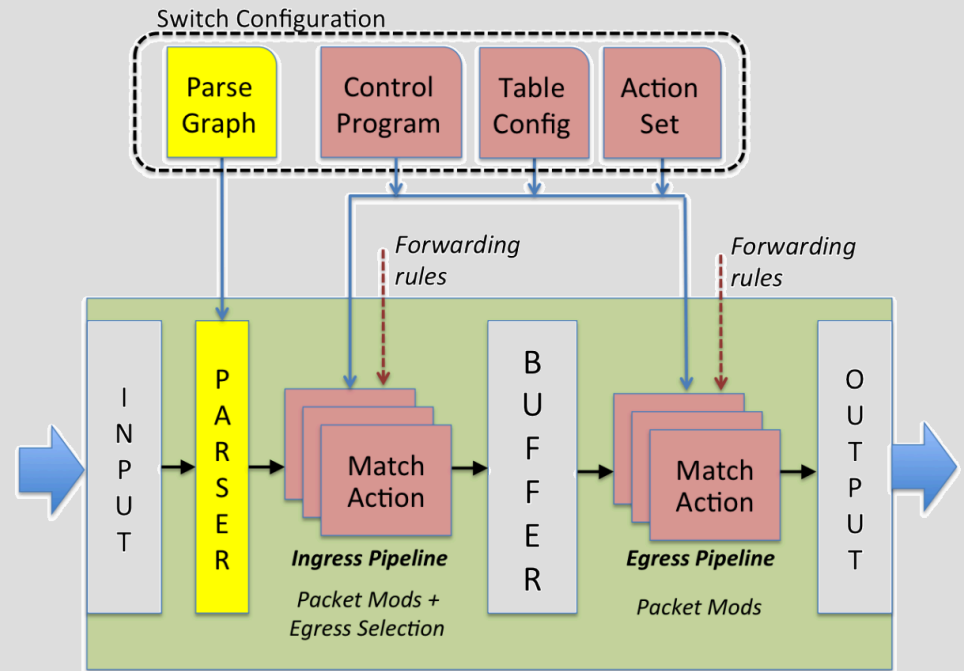
Abstract Forwarding Model (AFM)

1. Parsing the packet headers
2. The fields are passed to the match-action pipeline.
 - **Ingress:** determines the egress port/queue
 - **Egress:** per-instance header modifications
3. Metadata processing (e.g., timestamp)
4. As in OpenFlow, the queuing discipline is chosen at switch configuration time (e.g., minimum rate)



Abstract Forwarding Model (AFM)

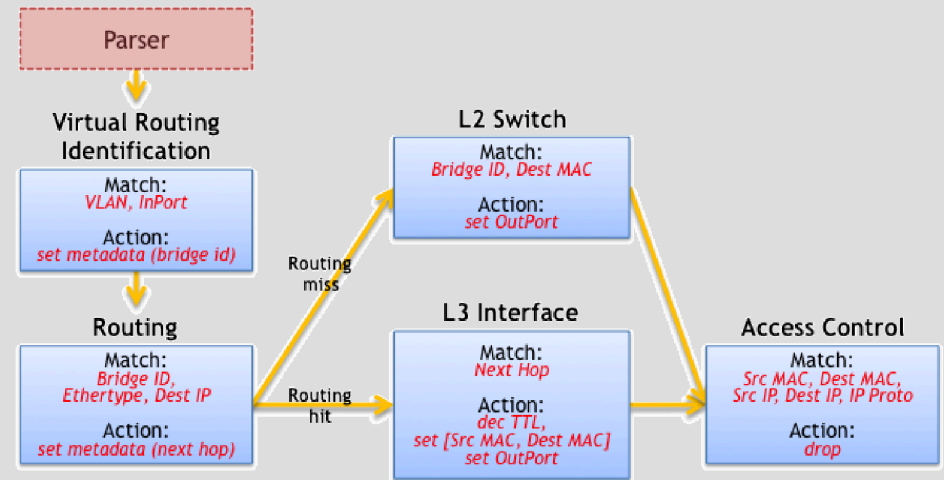
1. Parsing the packet headers
2. The fields are passed to the match-action pipeline.
 - **Ingress:** determines the egress port/queue
 - **Egress:** per-instance header modifications
3. Metadata processing (e.g., timestamp)
4. As in OpenFlow, the queuing discipline is chosen at switch configuration time (e.g., minimum rate)



Two-stage Compilation

Imperative control flow program based on AFM

1. The compiler translate the P4 program into **TDGs** (Table Dependency Graphs)
2. The **TDGs** are compiled into **target-dependent code**



Real Case Scenario

Setup: L2 Network Architecture

- *Edge (top-of-rack switches)*: connect end-hosts to the network
- *Core*: central layer that connects the edge devices

Problem: Growing End-Hosts and Overflowing Tables

- The L2 forwarding tables in the *core* are becoming too large → **overflow**
- It can cause *packet loss* and *network congestion*

Solutions: Multi-protocol Label Switching and PortLand

- *MPLS*: a technique that uses labels to make data forwarding decisions → **with multiple tags is daunting**
- *PortLand*: a scalable L2 network architecture → **rewrite MAC addresses**

P4: Language Design

Header: describes the structure of a series of fields and constraints on values

Parser: specifies how to identify headers and valid header sequences

Table: defines the fields to match on and the actions to take

Action: construction of actions from simpler protocol-independent primitives

Control Programs: determines the order of match+action tables that are applied to a packet

Header

Describes the structure of a series of fields and constraints on values

```
header ethernet {  
    fields {  
        dst_addr: 48; // bits  
        src_addr: 48;  
        ethertype: 16;  
    }  
}
```

```
header vlan {  
    fields {  
        pcp: 3;  
        cfi: 1;  
        vid: 12;  
        ethertype: 16;  
    }  
}
```

Header (Cont.)

```
header mTag {  
    fields {  
        up1: 8;  
        up2: 8;  
        down1: 8;  
        down2: 8;  
        ethertype: 16;  
    }  
}
```

- *mTag* can be added without altering the existing headers
- The core has two layers of aggregation
- Each core switch examines one of these bytes determined by its **location** and the **direction** of the packet

Parser

Specifies how to identify headers and valid header sequences

```
parser start { ethernet; }
```

```
parser ethernet {  
    switch(ethertype) {  
        case 0x8100: vlan;  
        case 0x9100: vlan;  
        case 0x800: ipv4;  
        // Other cases  
    }  
}
```

```
parser vlan {  
    switch(ethertype) {  
        case 0xaaaa: mTag;  
        case 0x800: ipv4;  
        // Other cases  
    }  
}
```

Parser (Cont.)

```
parser mTag {  
    switch(ethertype) {  
        case 0x800: ipv4;  
        // Other cases  
    }  
}
```

- Reached a state for a new header, the State Machine extracts the header and sends it to the match+action pipeline
- The parser for *mTag* is simple, it has only four states

Table

Defines the fields to match on and the actions to take

```
table mTag_table {
  reads {
    ethernet.dst_addr: exact;
    vlan.vid: exact;
  }
  actions {
    // At runtime, entries are
    // programmed with params
    // for the mTag action.
    add_mTag;
  }
  max_size: 20000;
}
```

The compiler knows what memory type use (e.g., TCAM, SRAM) and the amount of memory to allocate

- `reads`: the edge switch matches on the L2 destination address and the VLAN ID
- `actions`: selects an *mTag* to add to the header
- `max_size`: the maximum number of entries

Action

Construction of actions from simpler protocol-independent primitives

```
action add_mTag(up1, up2, down1, down2, egr_spec) {  
    add_header(mTag);  
    // Copy VLAN ethertype to mTag  
    copy_field(mTag.ethertype, vlan.ethertype);  
    // Set VLAN's ethertype to signal mTag  
    set_field(vlan.ethertype, 0xaaaa);  
    set_field(mTag.up1, up1);  
    set_field(mTag.up2, up2);  
    set_field(mTag.down1, down1);  
    set_field(mTag.down2, down2);  
    // Set the destination egress port as well  
    set_field(metadata.egress_spec, egr_spec);  
}
```

- P4 assumes parallel execution
- Parameters are passed from the match table at runtime
- The switch inserts the *mTag* after the VLAN header

Control Programs

Determines the order of match+action tables that are applied to a packet

```
control main() {  
    // Verify mTag state and port are consistent  
    table(source_check);  
    // If no error from source_check, continue  
    if (!defined(metadata.ingress_error)) {  
        // Attempt to switch to end hosts  
        table(local_switching);  
        if (!defined(metadata.egress_spec)) {  
            // Not a known local host; try mtagging  
            table(mTag_table);  
        }  
        // Check for unknown egress state or  
        // bad retagging with mTag.  
        table(egress_check);  
    }  
}
```

- *mTag* should only be seen on ports to the core
- `source_check` strips the *mTag* and records it in the metadata to avoid retagging
- If the `local_switching` table misses, the packet is not destined for a local host
- Both *local* and *core* forwarding control is handled by the `egress_check` table
- If unknown destination, the SDN controller is notified during `egress_check`

P4: Compilation Process

- The P4 compiler translates the P4 program into a *target-independent* representation (TDGs)
- The TDGs are compiled into *target-dependent* code
- The compiler can optimize the table layout to minimize the number of tables and the number of lookups
- The compiler can detect data dependencies and arrange tables in parallel or in series

Compiling Packet Parsers

- For devices with *programmable* parsers, the compiler generates the parser state machine (see PISA architecture)
- For devices with *fixed* parsers, the compiler verifies that the parser description is *consistent* with the device's fixed parser (e.g., ASICs)

Compiling Packet Parsers (Cont.)

Parser state table entries for the vlan and mTag sections of the parser

Current Version	Lookup Value	Next State
vlan	0xaaaa	mTag
vlan	0x800	ipv4
vlan	*	stop
mTag	0x800	ipv4
mTag	*	stop

The * is a wildcard that matches any value

The stop state indicates that the parser has finished processing the packet

Compiling Control Programs

The imperative control-flow representation does not call out dependencies between tables or opportunities for concurrency

1. The compiler analyzes the control program to determine dependencies between tables and opportunities for concurrency
2. The compiler generates the target configuration for the switch

Is this not familiar?

Two-stage compilation

Compiling Control Programs

The imperative control-flow representation does not call out dependencies between tables or opportunities for concurrency

1. The compiler analyzes the control program to determine dependencies between tables and opportunities for concurrency
2. The compiler generates the target configuration for the switch

Is this not familiar?

Two-stage compilation

Compiling Control Programs

The imperative control-flow representation does not call out dependencies between tables or opportunities for concurrency

1. The compiler analyzes the control program to determine dependencies between tables and opportunities for concurrency
2. The compiler generates the target configuration for the switch

Is this not familiar?
Two-stage compilation

Compiling Control Programs

The imperative control-flow representation does not call out dependencies between tables or opportunities for concurrency

1. The compiler analyzes the control program to determine dependencies between tables and opportunities for concurrency
2. The compiler generates the target configuration for the switch

Is this not familiar?

Two-stage compilation

1. Software Switches

- **Software Switches** provide complete flexibility:
 1. Table Count
 2. Table Configuration
 3. Parsing under SW control
- The compiler:
 1. Maps the mTag table graph to switch tables
 2. Uses table type to constrain width, height, and matching criterion
 3. Can optimize ternary matches with SW data structures

2. Hardware Switches with RAM and TCAM

In **edge** switches, the compiler configure hashing to perform efficient exact-matching using RAM

In **core** switches, which match on a subset of fields, the compiler maps the table to TCAM

3. Switches supporting parallel tables

The compiler can **detect** data dependencies and arrange tables in parallel or in series

In the mTag example, the mTag_table and local_switching tables can be executed in parallel up to the add_mTag action

4. Switches that apply actions at the end of the pipeline

The compiler can **tell** to the intermediate stages to generate metadata for the final action

In the mTag example, whether the mTag is added or not could be represented in metadata

5. Switches with a few tables

The compiler can **optimize** the table layout to minimize the number of tables and the number of lookups

When a controller installs a rule (at runtime), the compiler can generate P4 tables to generate the rules for the single physical table

In the mTag example, the `local_switching` table could be merged with the `mTag_table`

Thanks for your attention!

Slides available at

federicobruzzo.github.io/activities/presentations/P4-compiler-in-SDN.pdf

Website	federicobruzzo.github.io
Github	github.com/FedericoBruzzone
X	@fedebruzzone7
LinkedIn	in/federico-bruzzone
Telegram	@federicobruzzo
Email 1	federico.bruzzo@unimi.it
Email 2	federico.bruzzo.i@gmail.com

Bibliography

- [1] F. Hauser *et al.*, “A survey on data plane programming with p4: Fundamentals, advances, and applied research,” *Journal of Network and Computer Applications*, vol. 212, p. 103561, 2023.
- [2] N. McKeown *et al.*, “OpenFlow: enabling innovation in campus networks,” *SIGCOMM Comput. Commun. Rev.*, vol. 38, no. 2, pp. 69–74, Mar. 2008, doi: [10.1145/1355734.1355746](https://doi.org/10.1145/1355734.1355746).
- [3] P. Bosshart *et al.*, “P4: programming protocol-independent packet processors,” *SIGCOMM Comput. Commun. Rev.*, vol. 44, no. 3, pp. 87–95, Jul. 2014, doi: [10.1145/2656877.2656890](https://doi.org/10.1145/2656877.2656890).

Table (Addition)

```
table source_check {
    // Verify mtag only on ports to the core
    reads {
        mtag : valid; // Was mtag parsed?
        metadata.ingress_port : exact;
    }
    actions { // Each table entry specifies *one* action
        // If inappropriate mTag, send to CPU
        fault_to_cpu;
        // If mtag found, strip and record in metadata
        strip_mtag;
        // Otherwise, allow the packet to continue
        pass;
    }
    max_size: 64; // One rule per port
}
```

Table (Addition)

```
table local_switching {  
    // Reads destination and checks if local  
    // If miss occurs, goto mtag table.  
}  
table egress_check {  
    // Verify egress is resolved  
    // Do not retag packets received with tag  
    // Reads egress and whether packet was mTagged  
}
```