

# Online Reprogrammable Multi Tenant Switches

Johannes Krude<sup>1</sup>, Jaco Hofmann<sup>2</sup>, Matthias Eichholz<sup>2</sup>,  
Klaus Wehrle<sup>1</sup>, Andreas Koch<sup>2</sup>, Mira Mezini<sup>2</sup>



<sup>1</sup>RWTH Aachen University, <sup>2</sup>Technische Universität Darmstadt

<https://comsys.rwth-aachen.de/>

ENCP '19, 2019-12-09



TECHNISCHE  
UNIVERSITÄT  
DARMSTADT

&




- **On switch ...**
  - ▶ ...stateful load balancer **replaces hundreds of servers** [SilkRoad 2017]
  - ▶ ...data aggregation **speeds up databases** [Lerner et.al. 2019, ...]
  - ▶ ...paxos **reduces coordination overhead** [NetChain 2018, ...]
  - ▶ ...key-value caching **improves throughput and latency** [NetCache 2017, ...]

# Programmable Switch as a Service

- **On switch ...**

- ▶ ...stateful load balancer **replaces hundreds of servers** [SilkRoad 2017]
- ▶ ...data aggregation **speeds up databases** [Lerner et.al. 2019, ...]
- ▶ ...paxos **reduces coordination overhead** [NetChain 2018, ...]
- ▶ ...key-value caching **improves throughput and latency** [NetCache 2017, ...]

## Programmable Switch as a Service



Tenant 1:  
**Load Balancer**


Tenant 2:  
**SQL group-by**

Tenant 3:  
**PAXOS node**


Tenant 4:  
**Key-Value Cache**

**Packet Forwarding**

# Current Programmable Switch Pipeline [RMT 2013, Barefoot Tofino]




# Current Programmable Switch Pipeline [RMT 2013, Barefoot Tofino]



- **Runs a single P4 program**


# Current Programmable Switch Pipeline [RMT 2013, Barefoot Tofino]



online reconfigurable


- Runs a single P4 program

# Current Programmable Switch Pipeline [RMT 2013, Barefoot Tofino]



- Runs a single P4 program
- Reprogramming causes switch and network **downtime**

# Current Programmable Switch Pipeline [RMT 2013, Barefoot Tofino]



- Runs a single P4 program
- Reprogramming causes switch and network **downtime**

## We propose to modify the programmable switch architecture

- To enable hot-pluggability of on-switch functions



## Definition

The ability to **insert**, **modify**, and **remove** on-switch functions **without affecting** other on-switch functions and packet forwarding.

Tenant 1:

**Load Balancer**

*needs high availability*

Tenant 2:

**SQL group-by**

*lifetime of seconds*

**Packet Forwarding**

Programmable Switch

## Definition

The ability to **insert**, **modify**, and **remove** on-switch functions **without affecting** other on-switch functions and packet forwarding.

## Related Work

- **Use a dedicated switch for each application** [PPS 2019]
- **Put generalized functionality permanently onto switches** [NetAccel 2019, Ports et al. 2019]
- **Emulate P4 in Match-Action Tables** [Hyper4 2016, HyperVDP 2019]
  - ▶ Excessive Resource Consumption

## Definition

The ability to **insert**, **modify**, and **remove** on-switch functions **without affecting** other on-switch functions and packet forwarding.


## Related Work


- **Use a dedicated switch for each application** [PPS 2019]
- **Put generalized functionality permanently onto switches** [NetAccel 2019, Ports et al. 2019]
- **Emulate P4 in Match-Action Tables** [Hyper4 2016, HyperVDP 2019]
  - ▶ Excessive Resource Consumption

We want: **Switch Sharing & On-Demand Instantiation & Individual Customization**




# Switch Architecture Requirements

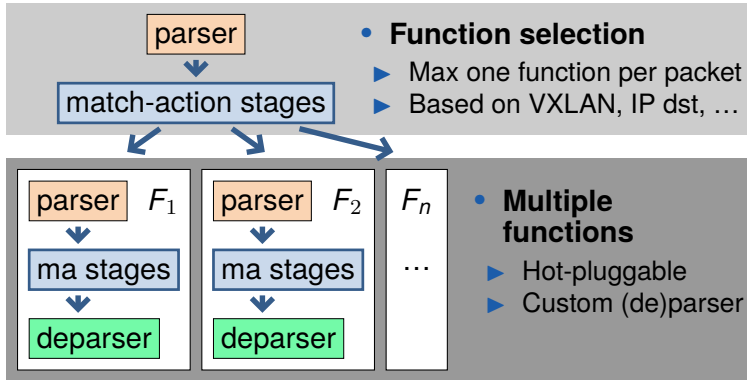





## Packet Structure



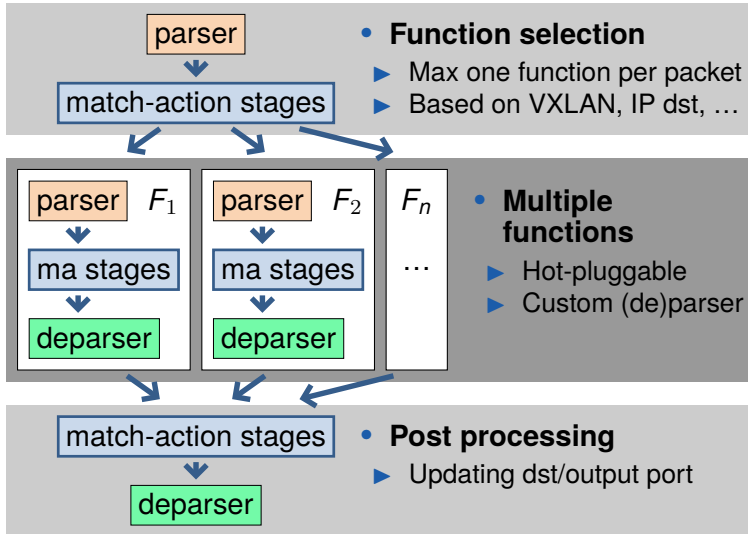
# Switch Architecture Requirements




## Packet Structure



# Switch Architecture Requirements

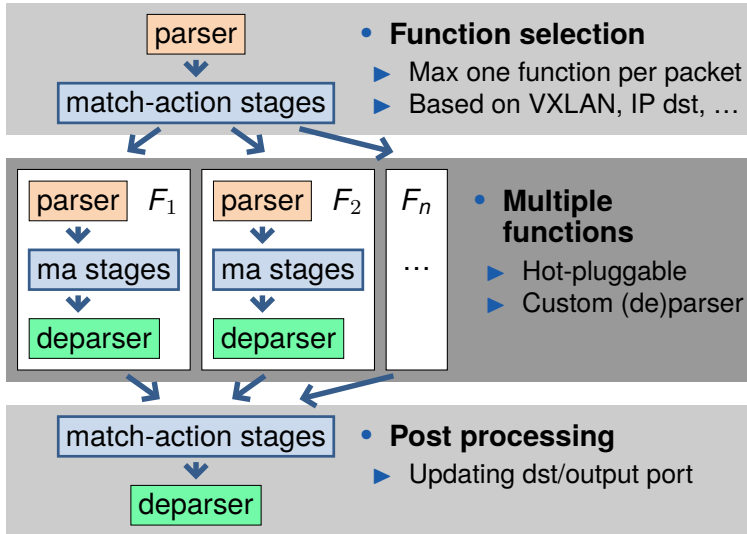


## Packet Structure






# Switch Architecture Requirements



## Packet Structure



### Program isolation

- ▶ Access to only own packets
- ▶ Limit access to outer header
- ▶ Control plane virtualization

- We present three different possible implementations
  - ▶ None of them yet implemented

## Multiple Switching ASICs

- ✓ Easily realizable
- ✗ No statefull functions


## Using FPGAs

- ✓ Realizable with FPGA knowledge
- ✗ Reduced Throughput


## An ASIC extension

- ✗ To be done by switching ASIC vendors
- ✓ High performance


- Alternate between two switching ASICs




- Alternate between two switching ASICs




- Alternate between two switching ASICs




- Alternate between two switching ASICs



- Alternate between two switching ASICs




- Alternate between two switching ASICs






- Alternate between two switching ASICs




# Multiple Programmable Switching ASICs

- Alternate between two switching ASICs
- Merge functions into single program




# Multiple Programmable Switching ASICs

- **Alternate between two switching ASICs**
- **Merge functions into single program**
- **Compiler provides isolation**
  - ▶ Restricts access to outer headers
  - ▶ Control plane mapping from table memory to function



# Multiple Programmable Switching ASICs

- **Alternate between two switching ASICs**
- **Merge functions into single program**
- **Compiler provides isolation**
  - ▶ Restricts access to outer headers
  - ▶ Control plane mapping from table memory to function



## Advantages

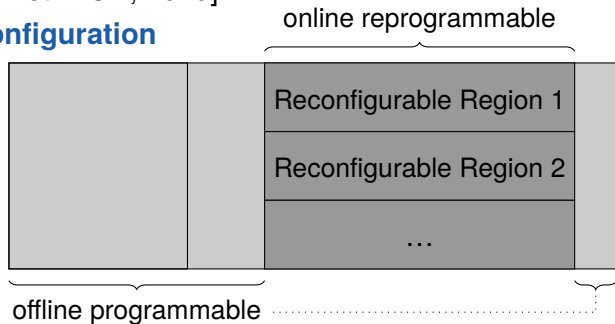
- ✓ Based on available hardware

## Limitations

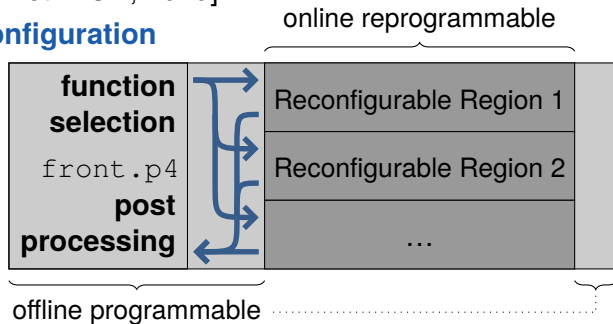
- ✗ Problematic for statefull functions

- **P4 can be executed on FPGAs** [P4→ NetFPGA, 2019]
- **FPGAs support dynamic partial reconfiguration**


- P4 can be executed on FPGAs [P4 → NetFPGA, 2019]
- FPGAs support dynamic partial reconfiguration



- **P4 can be executed on FPGAs** [P4 → NetFPGA, 2019]
- **FPGAs support dynamic partial reconfiguration**

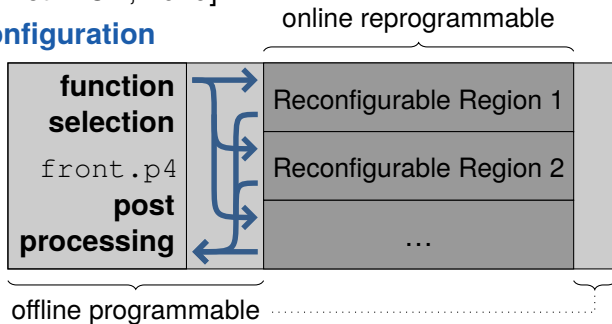


- **P4 can be executed on FPGAs** [P4 → NetFPGA, 2019]
- **FPGAs support dynamic partial reconfiguration**
- **Isolation on FPGA level**
  - ▶ Forward each packet to single region
  - ▶ Remove outer headers

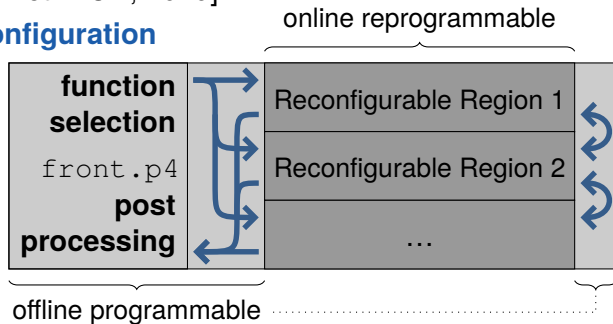




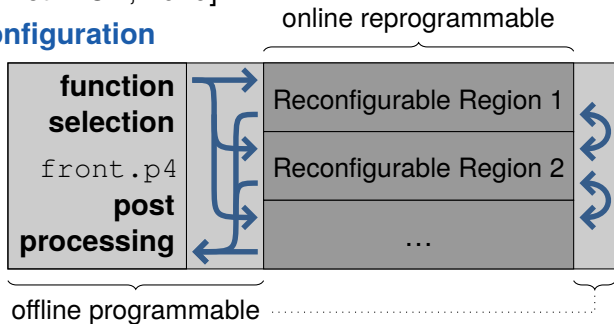
- **P4 can be executed on FPGAs** [P4→ NetFPGA, 2019]
- **FPGAs support dynamic partial reconfiguration**
- **Isolation on FPGA level**
  - ▶ Forward each packet to single region
  - ▶ Remove outer headers
- **Fixed sized reconfigurable regions**



- **P4 can be executed on FPGAs** [P4 → NetFPGA, 2019]
- **FPGAs support dynamic partial reconfiguration**
- **Isolation on FPGA level**
  - ▶ Forward each packet to single region
  - ▶ Remove outer headers
- **Fixed sized reconfigurable regions**
  - ▶ Forward packets between regions
  - ▶ Split function across multiple regions



- **P4 can be executed on FPGAs** [P4→ NetFPGA, 2019]
- **FPGAs support dynamic partial reconfiguration**
- **Isolation on FPGA level**
  - ▶ Forward each packet to single region
  - ▶ Remove outer headers
- **Fixed sized reconfigurable regions**
  - ▶ Forward packets between regions
  - ▶ Split function across multiple regions



## Advantages

- ✓ Readily available hardware
- ✓ Non-reconfigured regions keep state


## Limitations

- ✗ Limited throughput

- **Parser, matching, actions, and deparser are stored in SRAM and TCAM**
  - ▶ Use per entry validity bit for atomic updating [CoPTUA 2004]


# An extension to current switching ASICs

- **Parser, matching, actions, and deparser are stored in SRAM and TCAM**
  - ▶ Use per entry validity bit for atomic updating [CoPTUA 2004]
- **Insertion/Removal order to avoid inconsistent states**




# An extension to current switching ASICs

- **Parser, matching, actions, and deparser are stored in SRAM and TCAM**
  - ▶ Use per entry validity bit for atomic updating [CoPTUA 2004]
- **Insertion/Removal order to avoid inconsistent states**




# An extension to current switching ASICs

- **Parser, matching, actions, and deparser are stored in SRAM and TCAM**
  - ▶ Use per entry validity bit for atomic updating [CoPTUA 2004]
- **Insertion/Removal order to avoid inconsistent states**



# An extension to current switching ASICs


- **Parser, matching, actions, and deparser are stored in SRAM and TCAM**
  - ▶ Use per entry validity bit for atomic updating [CoPTUA 2004]
- **Insertion/Removal order to avoid inconsistent states**






# An extension to current switching ASICs

- **Parser, matching, actions, and deparser are stored in SRAM and TCAM**
  - ▶ Use per entry validity bit for atomic updating [CoPTUA 2004]
- **Insertion/Removal order to avoid inconsistent states**




# An extension to current switching ASICs

- **Parser, matching, actions, and deparser are stored in SRAM and TCAM**
  - ▶ Use per entry validity bit for atomic updating [CoPTUA 2004]
- **Insertion/Removal order to avoid inconsistent states**
- **P4 compiler tracks switch occupation**





# An extension to current switching ASICs

- **Parser, matching, actions, and deparser are stored in SRAM and TCAM**
  - ▶ Use per entry validity bit for atomic updating [CoPTUA 2004]
- **Insertion/Removal order to avoid inconsistent states**
- **P4 compiler tracks switch occupation**
- **Isolation same as with multiple switching ASICs**





# An extension to current switching ASICs

- **Parser, matching, actions, and deparser are stored in SRAM and TCAM**
  - ▶ Use per entry validity bit for atomic updating [CoPTUA 2004]
- **Insertion/Removal order to avoid inconsistent states**
- **P4 compiler tracks switch occupation**
- **Isolation same as with multiple switching ASICs**
- **Additional Pipeline for custom parser after function selection**
  - ▶ Some switches already have an additional pipeline
  - ▶ Pipelines can share resources [RMT 2013]



# An extension to current switching ASICs

- **Parser, matching, actions, and deparser are stored in SRAM and TCAM**
  - ▶ Use per entry validity bit for atomic updating [CoPTUA 2004]
- **Insertion/Removal order to avoid inconsistent states**
- **P4 compiler tracks switch occupation**
- **Isolation same as with multiple switching ASICs**
- **Additional Pipeline for custom parser after function selection**
  - ▶ Some switches already have an additional pipeline
  - ▶ Pipelines can share resources [RMT 2013]



## Limitations

- ✗ Needs to be done by ASIC vendors

## Advantages

- ✓ Same performance as current ASICs

- **Online reprogrammability is needed for "Programmable Switches as a Service"**
- **We propose an architecture for online reprogrammability**
  - ▶ No implementation yet

- **Online reprogrammability is needed for "Programmable Switches as a Service"**
- **We propose an architecture for online reprogrammability**
  - ▶ No implementation yet

## New interesting resource management questions

- Measuring & accounting resource usage
- Resource allocation
- Avoiding resource fragmentation
- ...