

Slicing Match-Action Pipeline Resources for Multitenancy on Programmable Switches

Johannes Krude, Felix Frei, Pedram Ahmadiyeh,
René Glebke, Mirko Stoffers, Klaus Wehrle



RWTH Aachen University

krude@comsys.rwth-aachen.de

NetSoft, 2025-06-26

Slicing Match-Action Pipeline Resources for Multitenancy

- Programmable Switch, e.g., Tofino → guaranteed and high throughput
- But, often only a subset of resources (SRAM, TCAM, action slots,...) needed
→ Multiple tenant and programs on a shared switch



Slicing Match-Action Pipeline Resources for Multitenancy

- Programmable Switch, e.g., Tofino → guaranteed and high throughput
- But, often only a subset of resources (SRAM, TCAM, action slots,...) needed
 - Multiple tenant and programs on a shared switch



Related Work

- Modify pipeline for multitenancy [Wang et.al. 2022, MTPSA, P4VBox]
 - ▶ Not applicable to existing hardware
- Compile merged P4 program [P4Weaver, Wang et.al. 2020, P4Visor, P4Brick]
 - ▶ No guaranteed pipeline resources



Slicing Match-Action Pipeline Resources for Multitenancy

- Programmable Switch, e.g., Tofino → guaranteed and high throughput
- But, often only a subset of resources (SRAM, TCAM, action slots,...) needed
→ Multiple tenant and programs on a shared switch



Related Work

- Modify pipeline for multitenancy [Wang et.al. 2022, MTPSA, P4VBox]
 - ▶ Not applicable to existing hardware
- Compile merged P4 program [P4Weaver, Wang et.al. 2020, P4Visor, P4Brick]
 - ▶ No guaranteed pipeline resources



Proposal: Resource guarantees on existing pipeline

- Divide pipeline resources into slices
- Compile tenant program for slice

Slicing Match-Action Pipeline Resources for Multitenancy

- Programmable Switch, e.g., Tofino → guaranteed and high throughput
- But, often only a subset of resources (SRAM, TCAM, action slots,...) needed
→ Multiple tenant and programs on a shared switch



Related Work

- Modify pipeline for multitenancy [Wang et.al. 2022, MTPSA, P4VBox]
 - ▶ Not applicable to existing hardware
- Compile merged P4 program [P4Weaver, Wang et.al. 2020, P4Visor, P4Brick]
 - ▶ No guaranteed pipeline resources



Proposal: Resource guarantees on existing pipeline

- Divide pipeline resources into slices
- Compile tenant program for slice

Implemented for Tofino

- Eval on real P4 programs
- Sourcecode on GitHub

IXPs (e.g., DeCIX) need a lot of space and energy

Currently: Each customer a physical router box

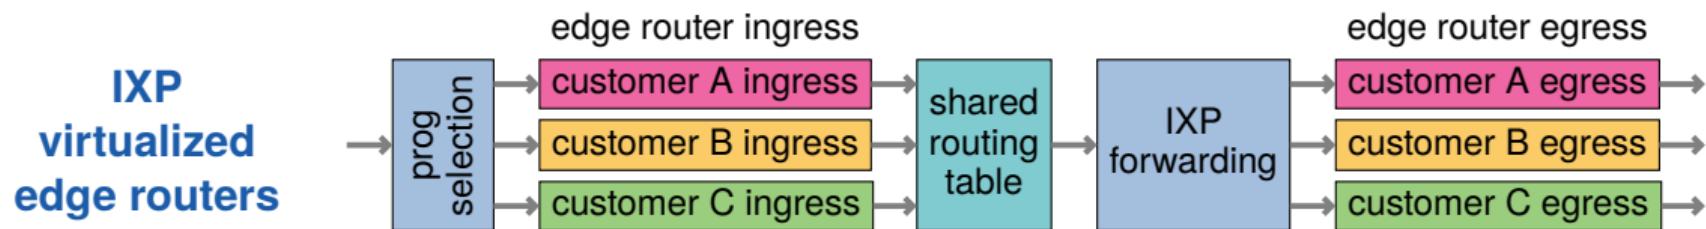
Instead: Upload P4 router to shared switch

IXP Use Case: Virtualized Edge Routers

IXPs (e.g., DeCIX) need a lot of space and energy

Currently: Each customer a physical router box

Instead: Upload P4 router to shared switch

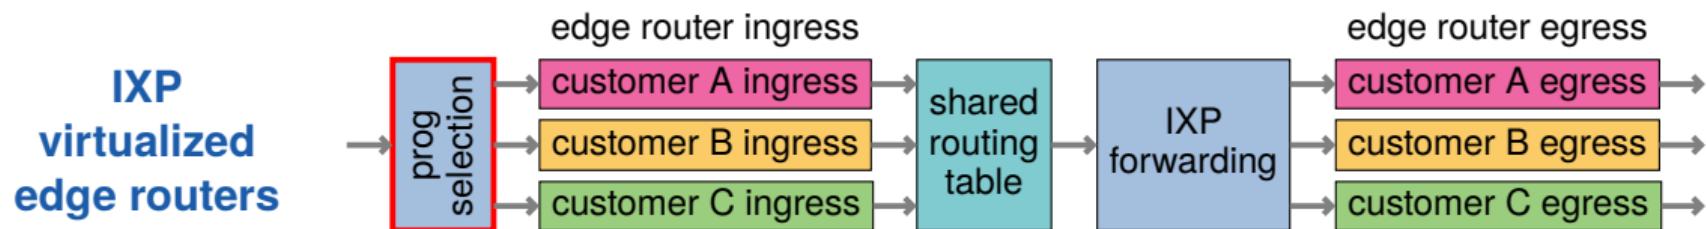


IXP Use Case: Virtualized Edge Routers

IXPs (e.g., DeCIX) need a lot of space and energy

Currently: Each customer a physical router box

Instead: Upload P4 router to shared switch

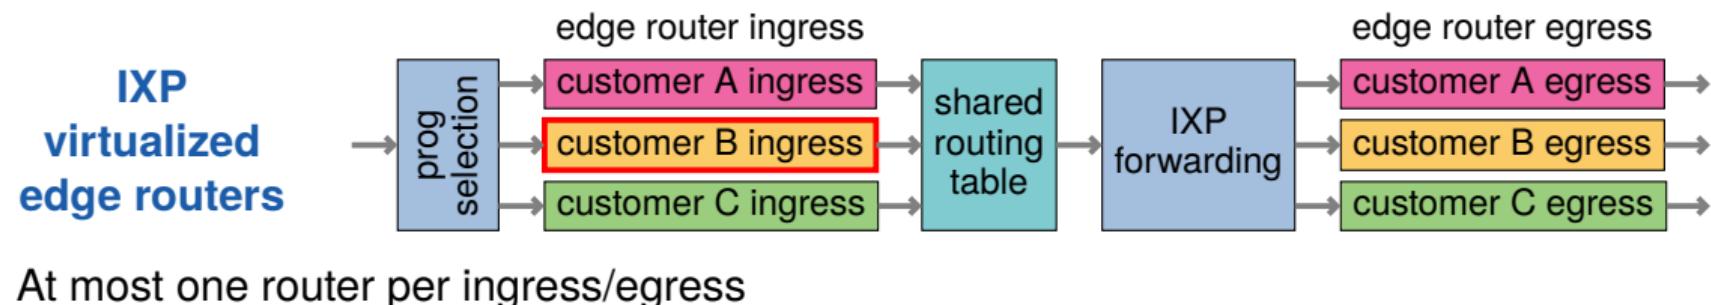


IXP Use Case: Virtualized Edge Routers

IXPs (e.g., DeCIX) need a lot of space and energy

Currently: Each customer a physical router box

Instead: Upload P4 router to shared switch

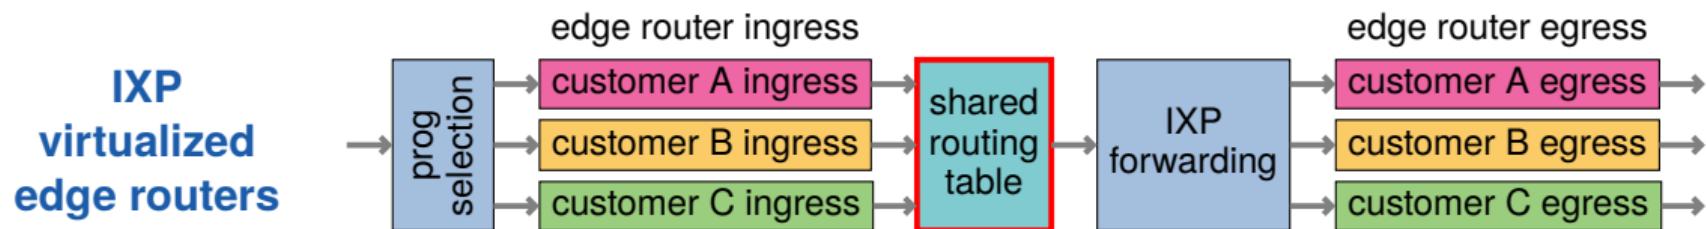


IXP Use Case: Virtualized Edge Routers

IXPs (e.g., DeCIX) need a lot of space and energy

Currently: Each customer a physical router box

Instead: Upload P4 router to shared switch



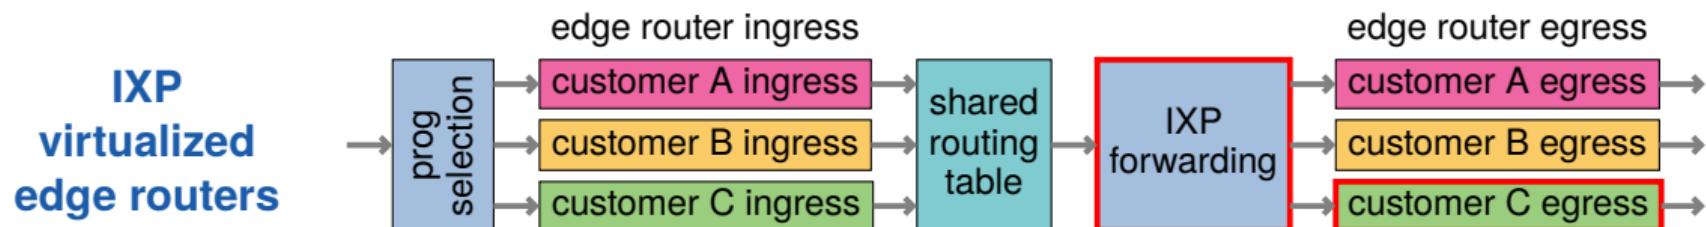
- At most one router per ingress/egress
- Sharing common routines and data

IXP Use Case: Virtualized Edge Routers

IXPs (e.g., DeCIX) need a lot of space and energy

Currently: Each customer a physical router box

Instead: Upload P4 router to shared switch



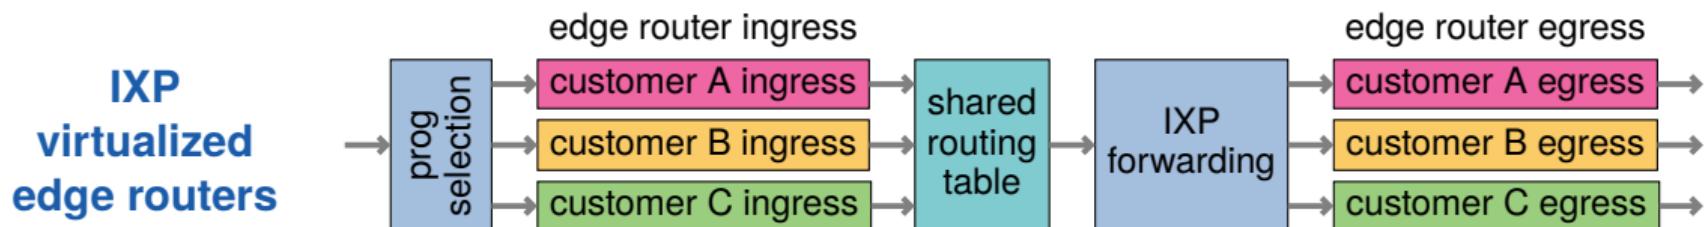
- At most one router per ingress/egress
- Sharing common routines and data

IXP Use Case: Virtualized Edge Routers

IXPs (e.g., DeCIX) need a lot of space and energy

Currently: Each customer a physical router box

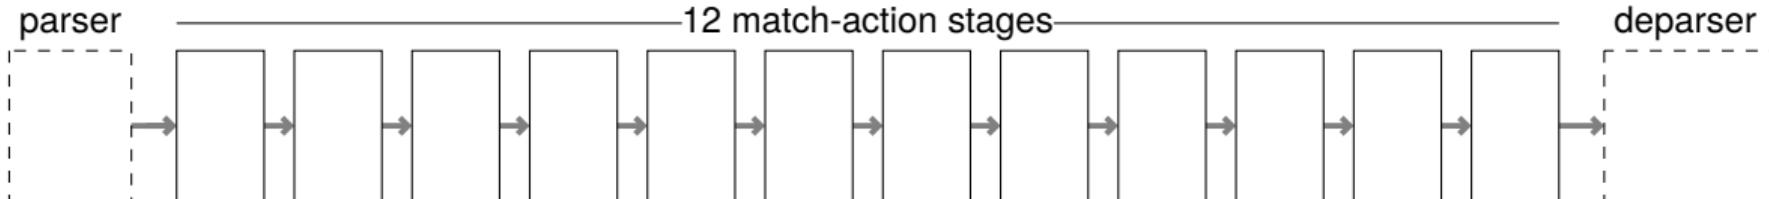
Instead: Upload P4 router to shared switch



- At most one router per ingress/egress
- Sharing common routines and data

Requirements for switch sharing

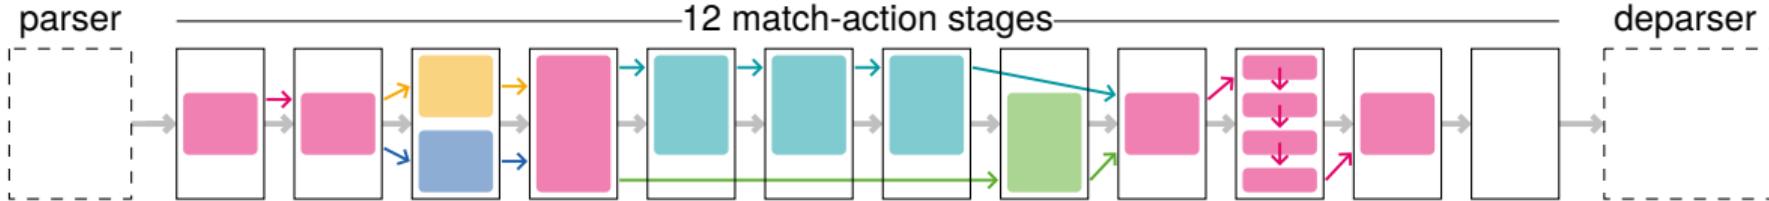
- Safely composing programs [P4Weaver, Wang et.al. 2020, P4Visor, P4Brick]
- (Our Contribution): Guaranteed switch pipeline resources for each customer



The Match-Action Pipeline

- Parser, m-a stages, deparser
- Separate resources at each stage (SRAM, TCAM, action slots, ...)
- Each stage divided into logical tables

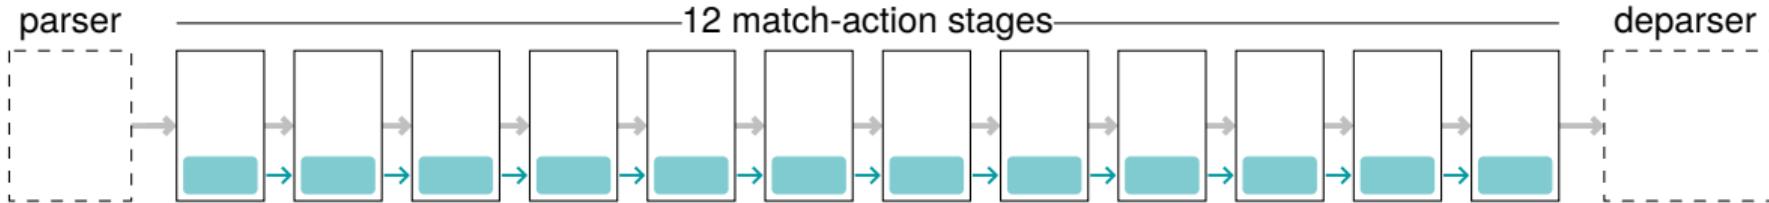
Pipeline Resources



The Match-Action Pipeline

- Parser, m-a stages, deparser
- Separate resources at each stage (SRAM, TCAM, action slots, ...)
- Each stage divided into logical tables

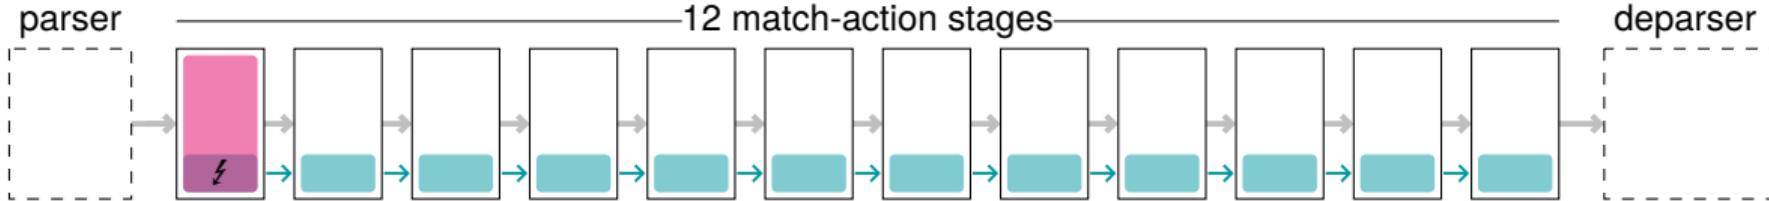
- **Compiler: fits P4 tables into stages (NP-hard)**
- **Related Work: compile merged P4 program**
 - ▶ If it does not fit, which tenant should be blamed?



The Match-Action Pipeline

- Parser, m-a stages, deparser
- Separate resources at each stage (SRAM, TCAM, action slots, ...)
- Each stage divided into logical tables

- **Compiler:** fits P4 tables into stages (NP-hard)
- **Related Work:** compile merged P4 program
 - ▶ If it does not fit, which tenant should be blamed?
- **Example:**
 - ▶ Long Program needs all stages

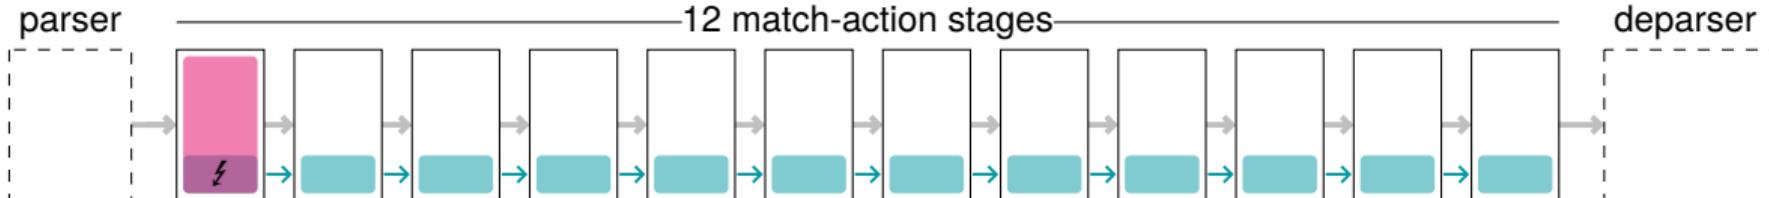


The Match-Action Pipeline

- Parser, m-a stages, deparser
- Separate resources at each stage (SRAM, TCAM, action slots, ...)
- Each stage divided into logical tables

- **Compiler:** fits P4 tables into stages (NP-hard)
- **Related Work:** compile merged P4 program
 - ▶ If it does not fit, which tenant should be blamed?
- **Example:**
 - ▶ Long Program needs all stages
 - ▶ Wide Program needs full stage

Pipeline Resources



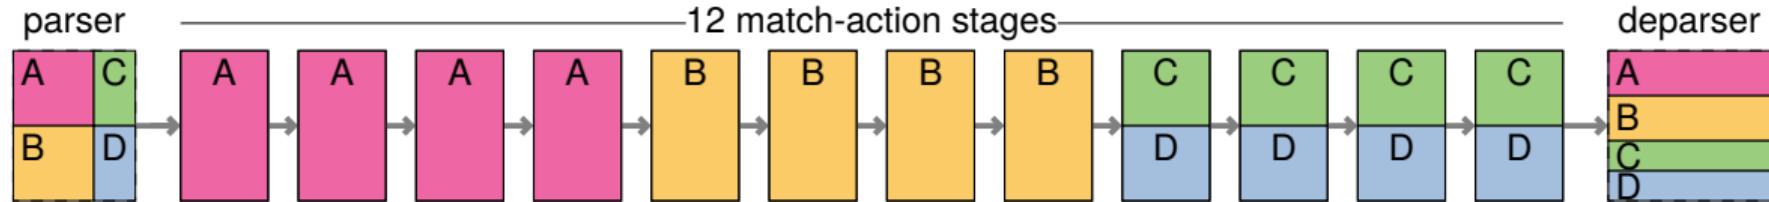
The Match-Action Pipeline

- Parser, m-a stages, deparser
- Separate resources at each stage (SRAM, TCAM, action slots, ...)
- Each stage divided into logical tables

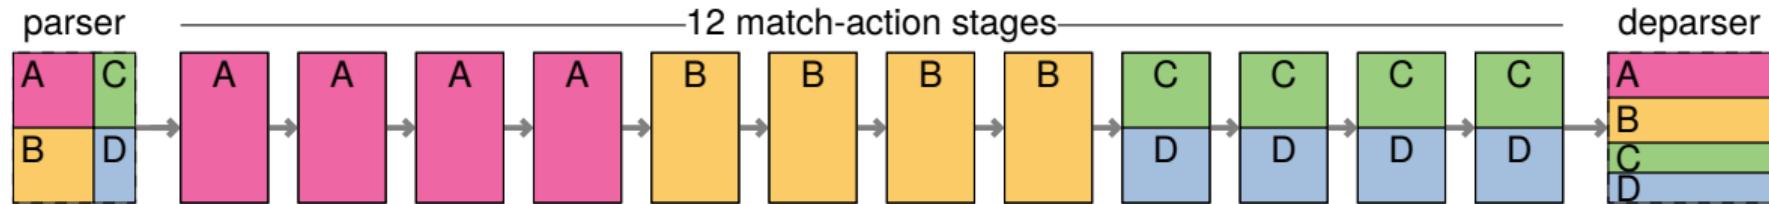
- **Compiler:** fits P4 tables into stages (NP-hard)
- **Related Work:** compile merged P4 program
 - ▶ If it does not fit, which tenant should be blamed?
- **Example:**
 - ▶ Long Program needs all stages
 - ▶ Wide Program needs full stage

Tenants should know in advance the acceptable resource usage and program shape

Slicing Match-Action Pipeline Resources



Slicing Match-Action Pipeline Resources



Switch Provider

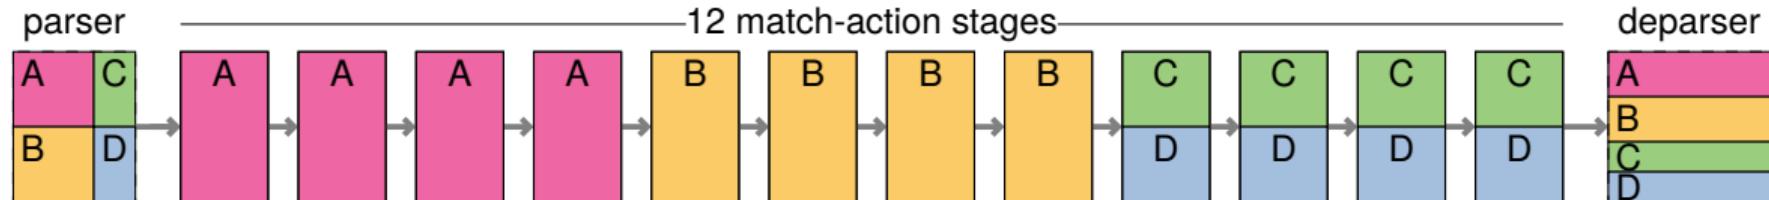
divides pipeline into slices

Tenant

creates or adapts program

negotiate pipeline slice

Slicing Match-Action Pipeline Resources



Switch Provider

divides pipeline into slices

Tenant

creates or adapts program

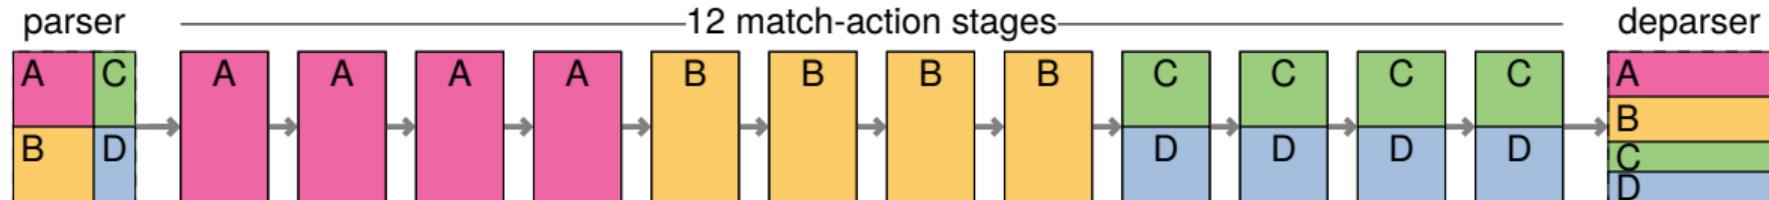
negotiate pipeline slice

submit compiled program

compiles program for slice



Slicing Match-Action Pipeline Resources



Switch Provider

divides pipeline into slices

negotiate pipeline slice

Tenant

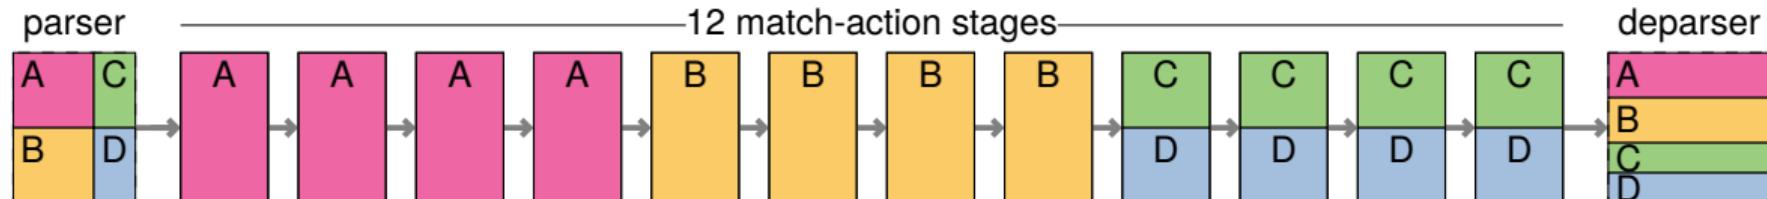
creates or adapts program

checks if program is within slice

submit compiled program

compiles program for slice

Slicing Match-Action Pipeline Resources



Switch Provider

divides pipeline into slices

negotiate pipeline slice

Tenant

creates or adapts program



checks if program is within slice

submit compiled program

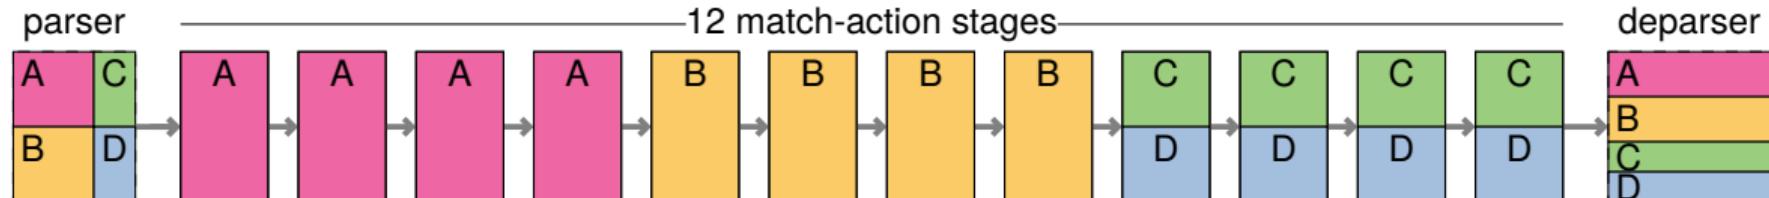


compiles program for slice



composed with other programs

Slicing Match-Action Pipeline Resources



Switch Provider

divides pipeline into slices

negotiate pipeline slice

Tenant

creates or adapts program

checks if program is within slice

submit compiled program

composed with other programs

compiles program for slice

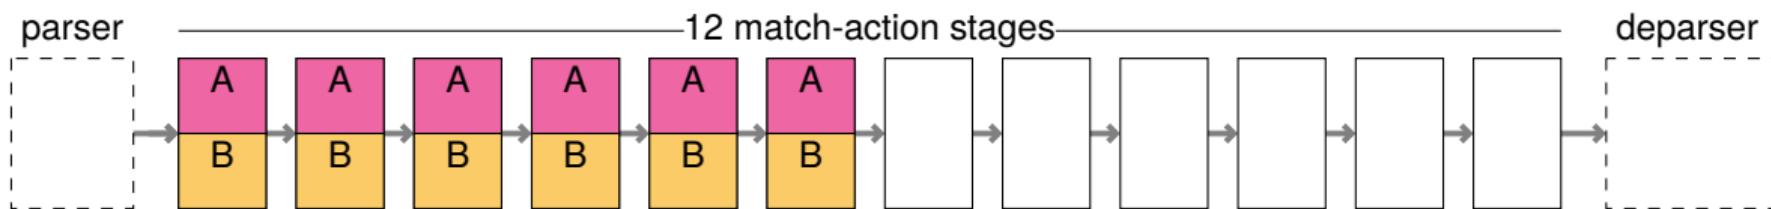
Slice checking guarantees all other tenants the availability of resources

Program and Slice Shapes

Slices are only useful if the tenant program can be shaped to fit into the slice

Same slice for each tenant → Little flexibility

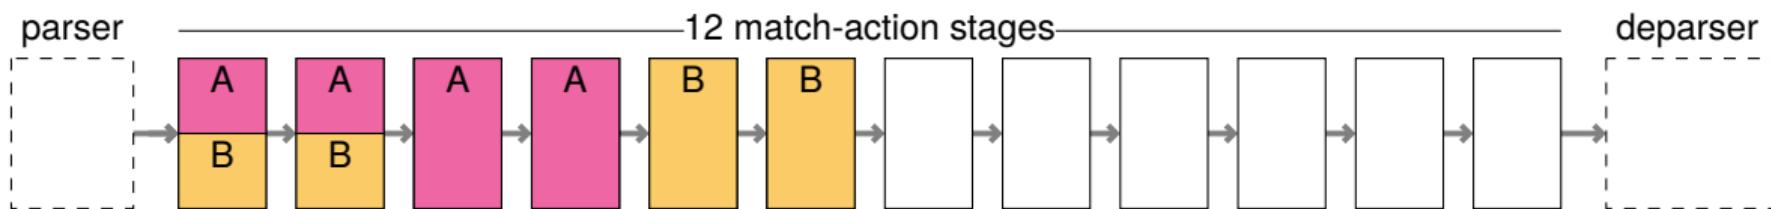
Individual slice shapes → Unallocated leftover resources



Slices are only useful if the tenant program can be shaped to fit into the slice

Same slice for each tenant → Little flexibility

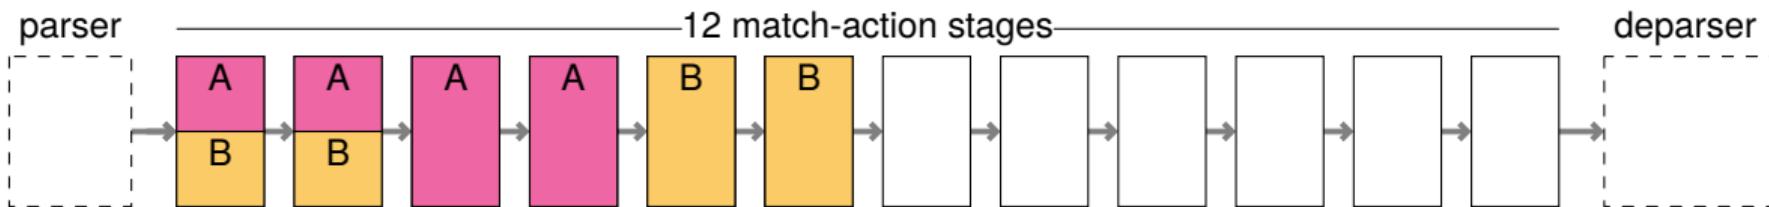
Individual slice shapes → Unallocated leftover resources



Slices are only useful if the tenant program can be shaped to fit into the slice

Same slice for each tenant → Little flexibility

Individual slice shapes → Unallocated leftover resources



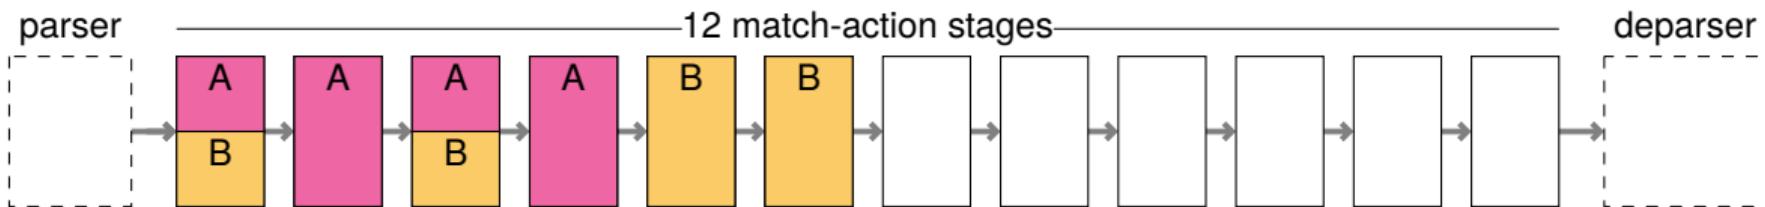
Reordering undivided stages

- A tenant may use undivided stages at any place in his program
- But, only undivided

Slices are only useful if the tenant program can be shaped to fit into the slice

Same slice for each tenant → Little flexibility

Individual slice shapes → Unallocated leftover resources



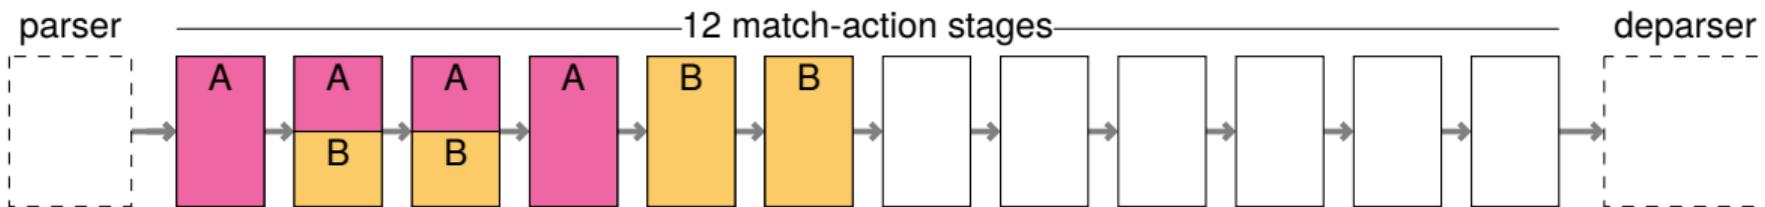
Reordering undivided stages

- A tenant may use undivided stages at any place in his program
- But, only undivided

Slices are only useful if the tenant program can be shaped to fit into the slice

Same slice for each tenant → Little flexibility

Individual slice shapes → Unallocated leftover resources



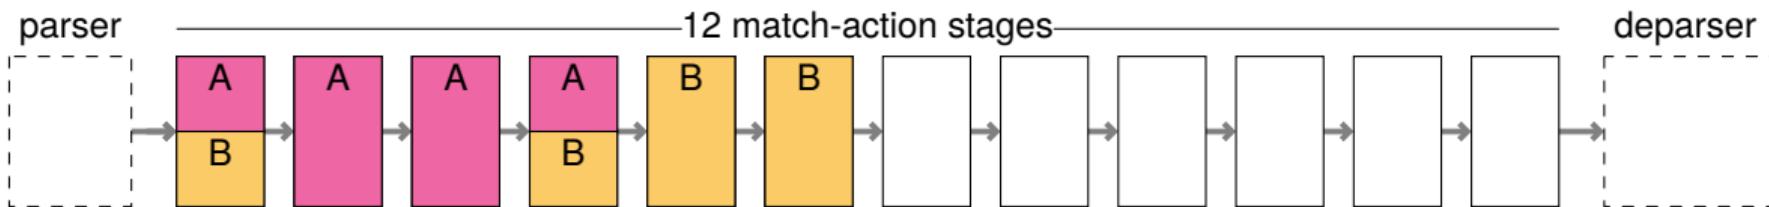
Reordering undivided stages

- A tenant may use undivided stages at any place in his program
- But, only undivided

Slices are only useful if the tenant program can be shaped to fit into the slice

Same slice for each tenant → Little flexibility

Individual slice shapes → Unallocated leftover resources



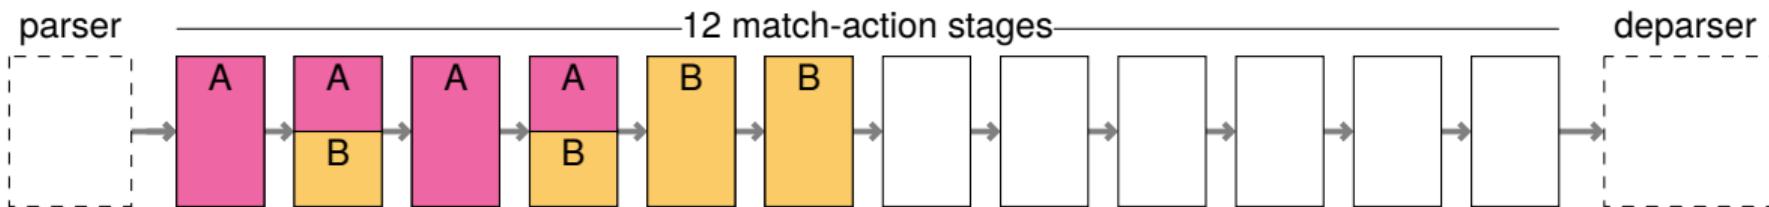
Reordering undivided stages

- A tenant may use undivided stages at any place in his program
- But, only undivided

Slices are only useful if the tenant program can be shaped to fit into the slice

Same slice for each tenant → Little flexibility

Individual slice shapes → Unallocated leftover resources



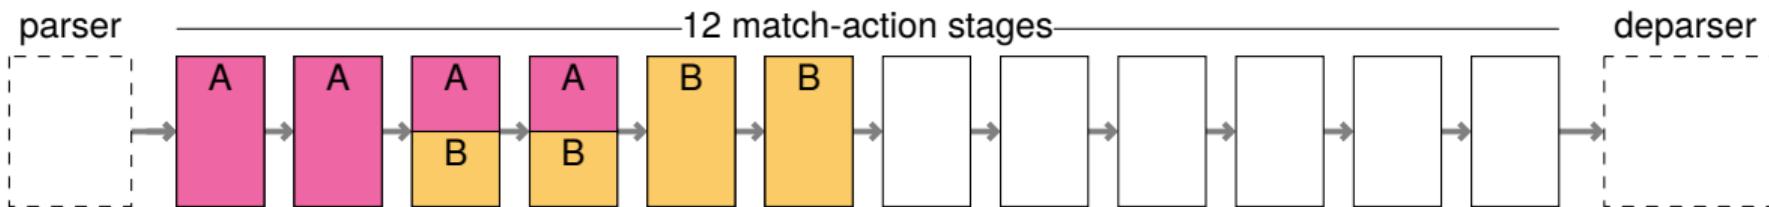
Reordering undivided stages

- A tenant may use undivided stages at any place in his program
- But, only undivided

Slices are only useful if the tenant program can be shaped to fit into the slice

Same slice for each tenant → Little flexibility

Individual slice shapes → Unallocated leftover resources



Reordering undivided stages

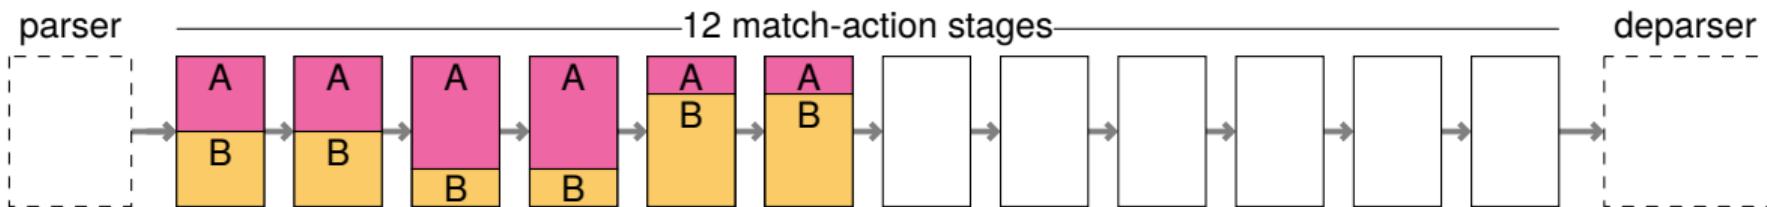
- A tenant may use undivided stages at any place in his program
- But, only undivided

Program and Slice Shapes

Slices are only useful if the tenant program can be shaped to fit into the slice

Same slice for each tenant → Little flexibility

Individual slice shapes → Unallocated leftover resources

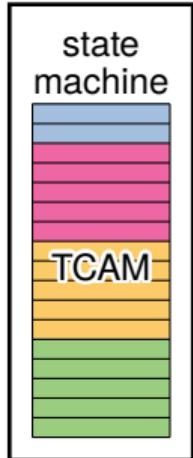


Reordering undivided stages

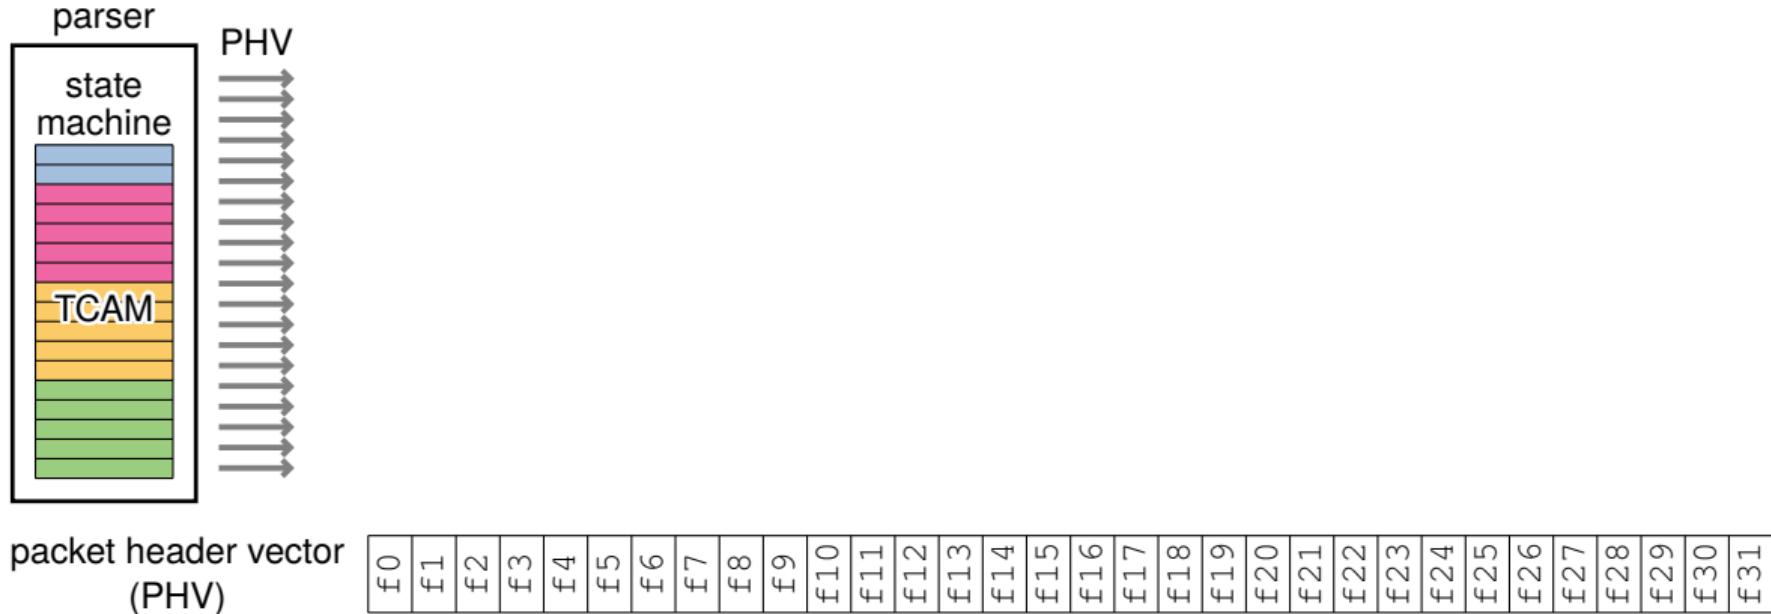
- A tenant may use undivided stages at any place in his program
- But, only undivided

Slicing Individual Resources

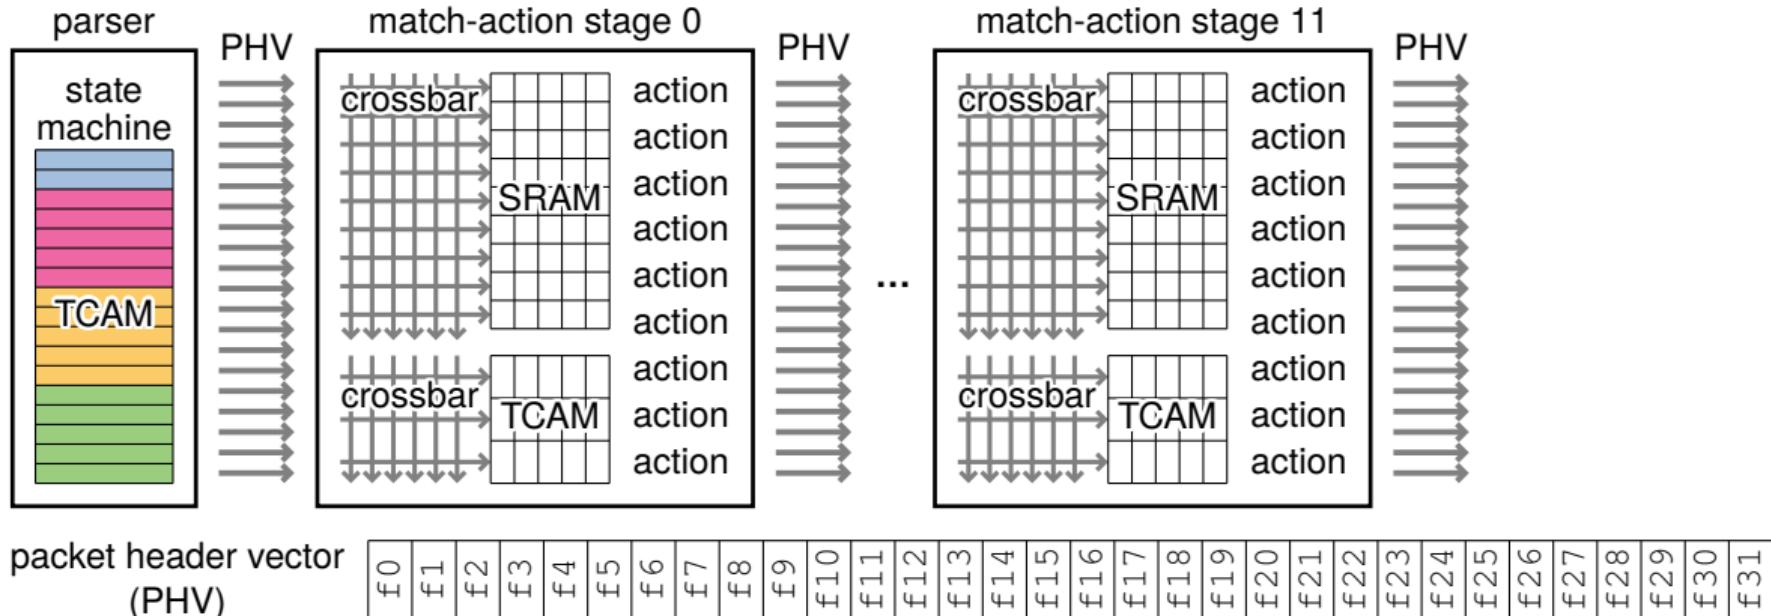
parser



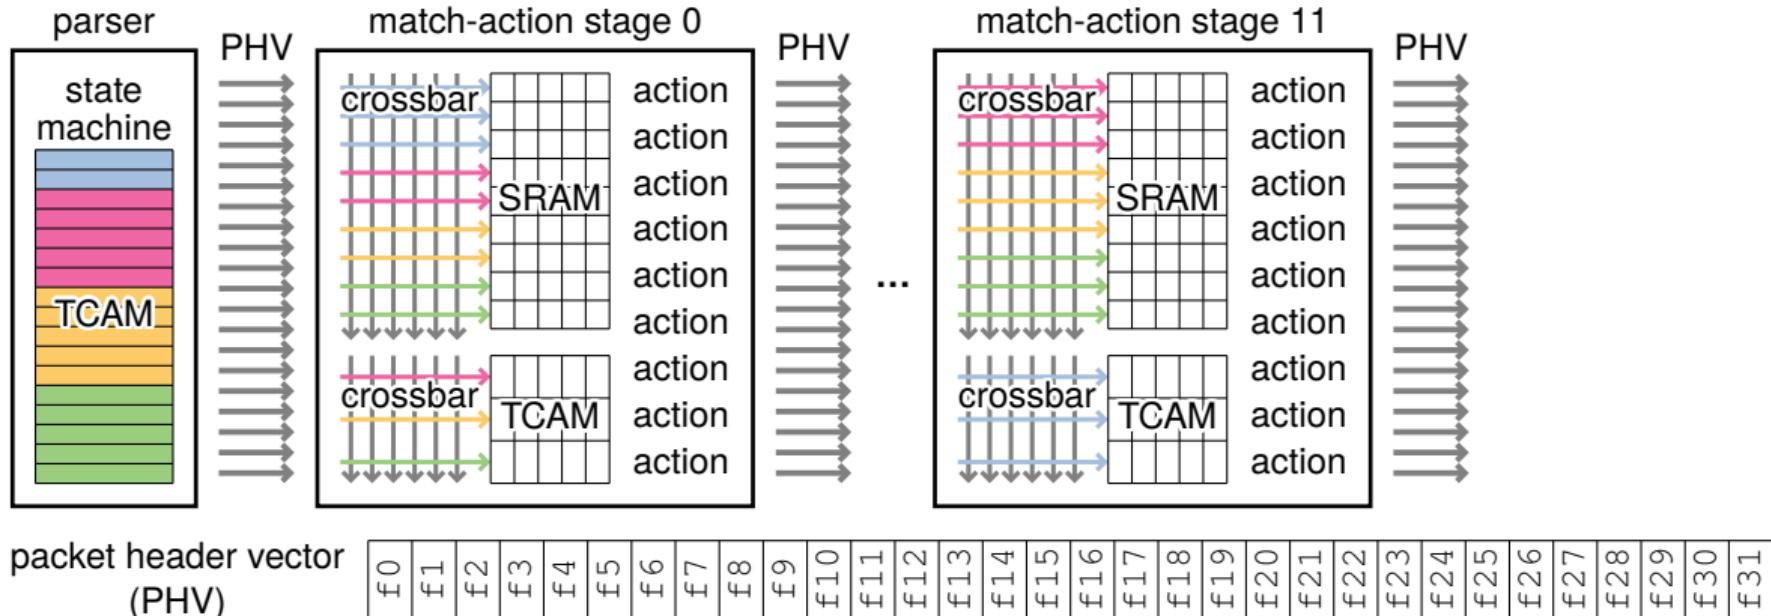
Slicing Individual Resources



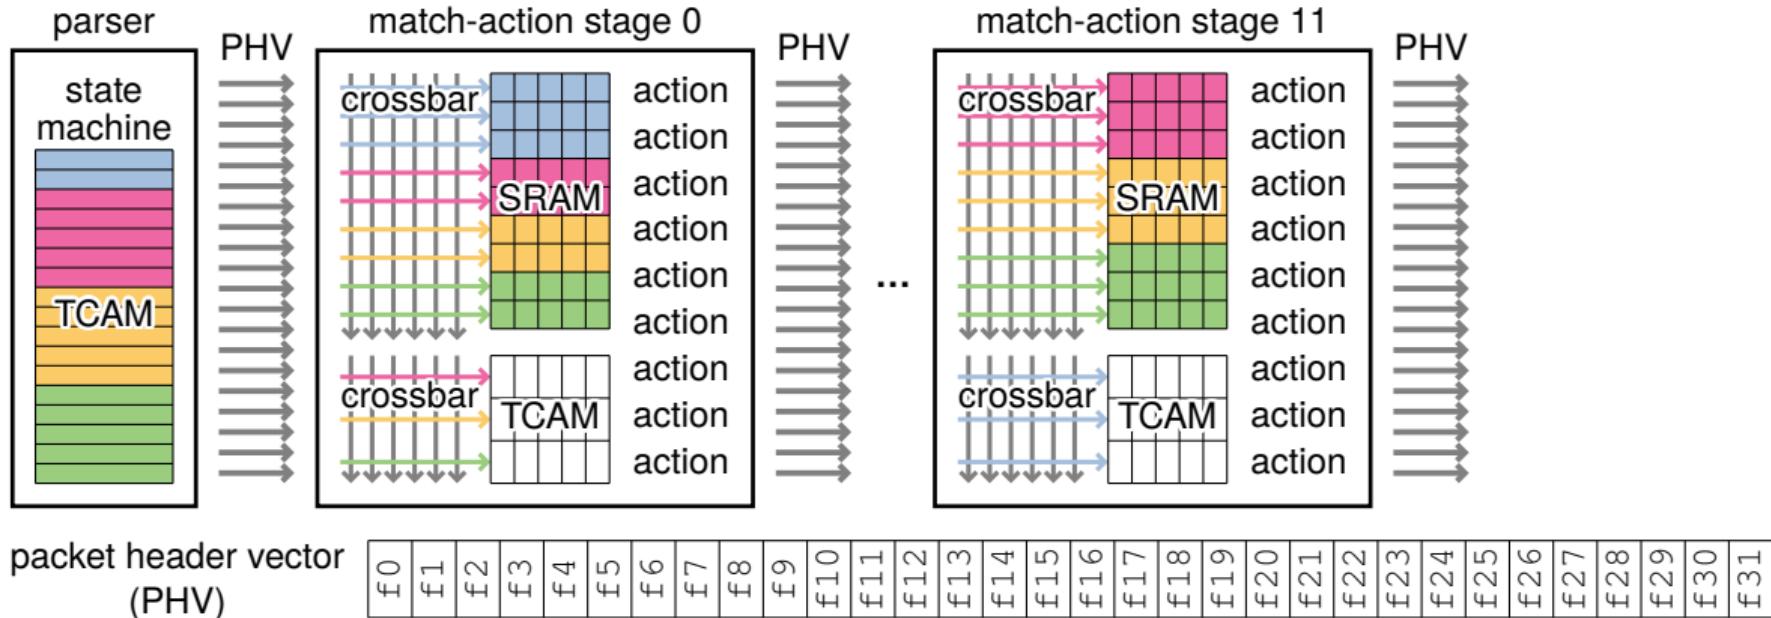
Slicing Individual Resources



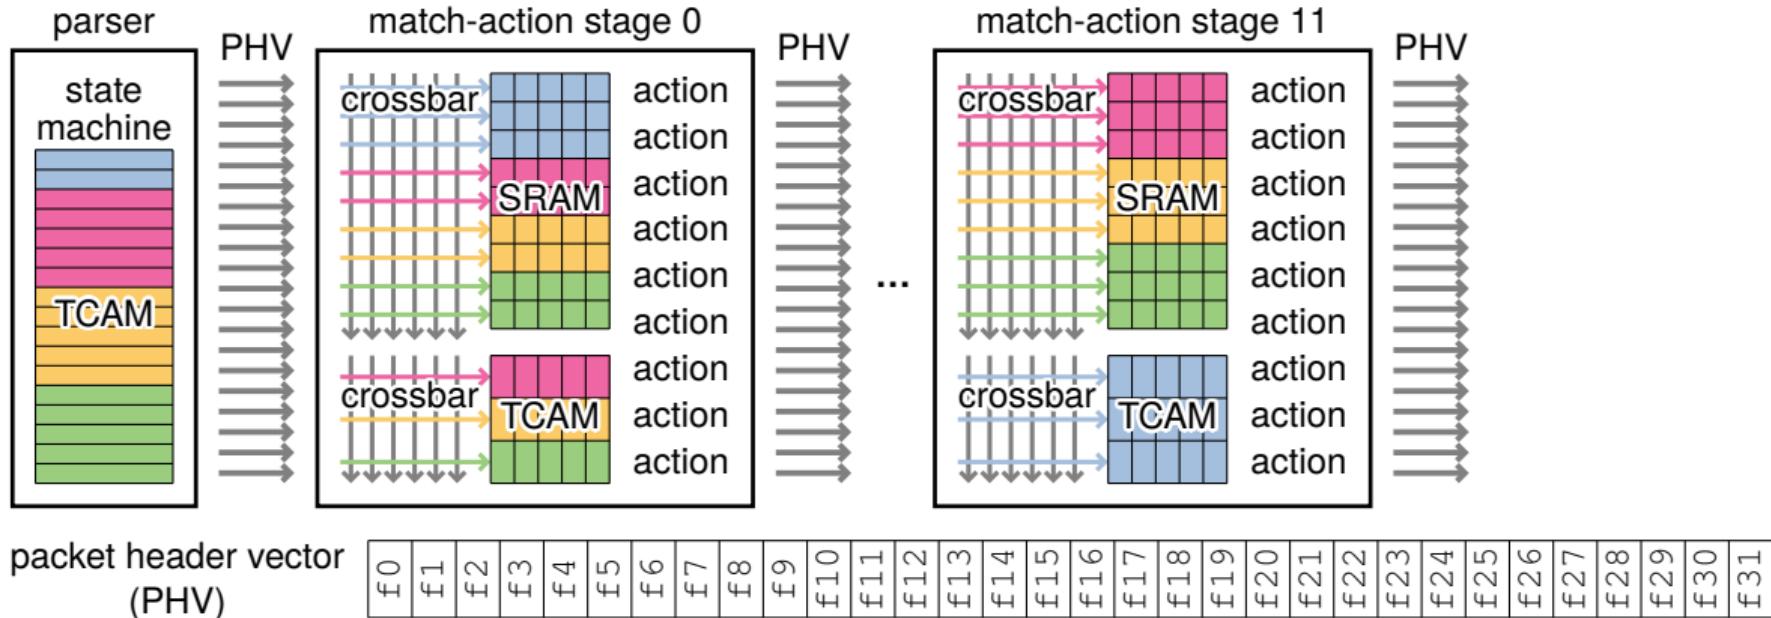
Slicing Individual Resources



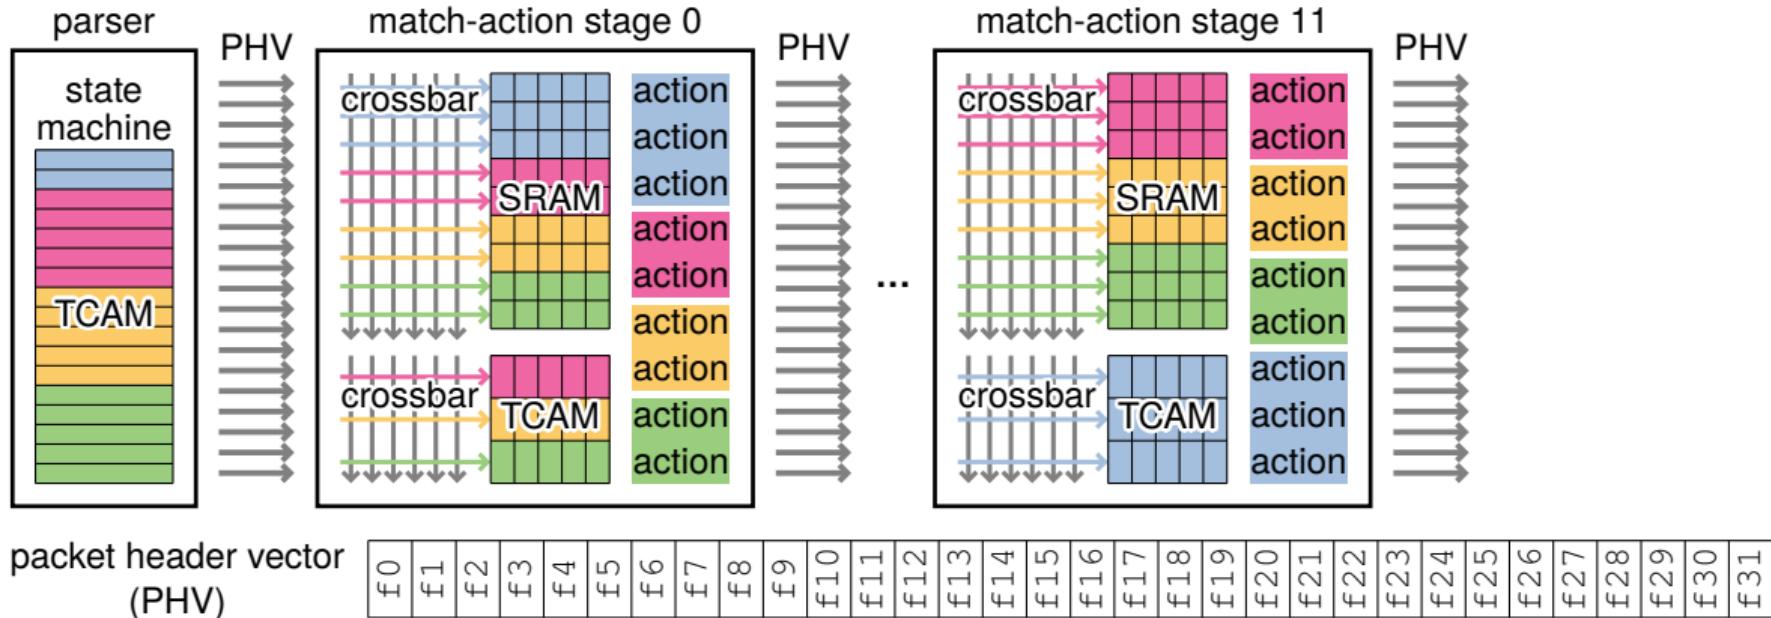
Slicing Individual Resources



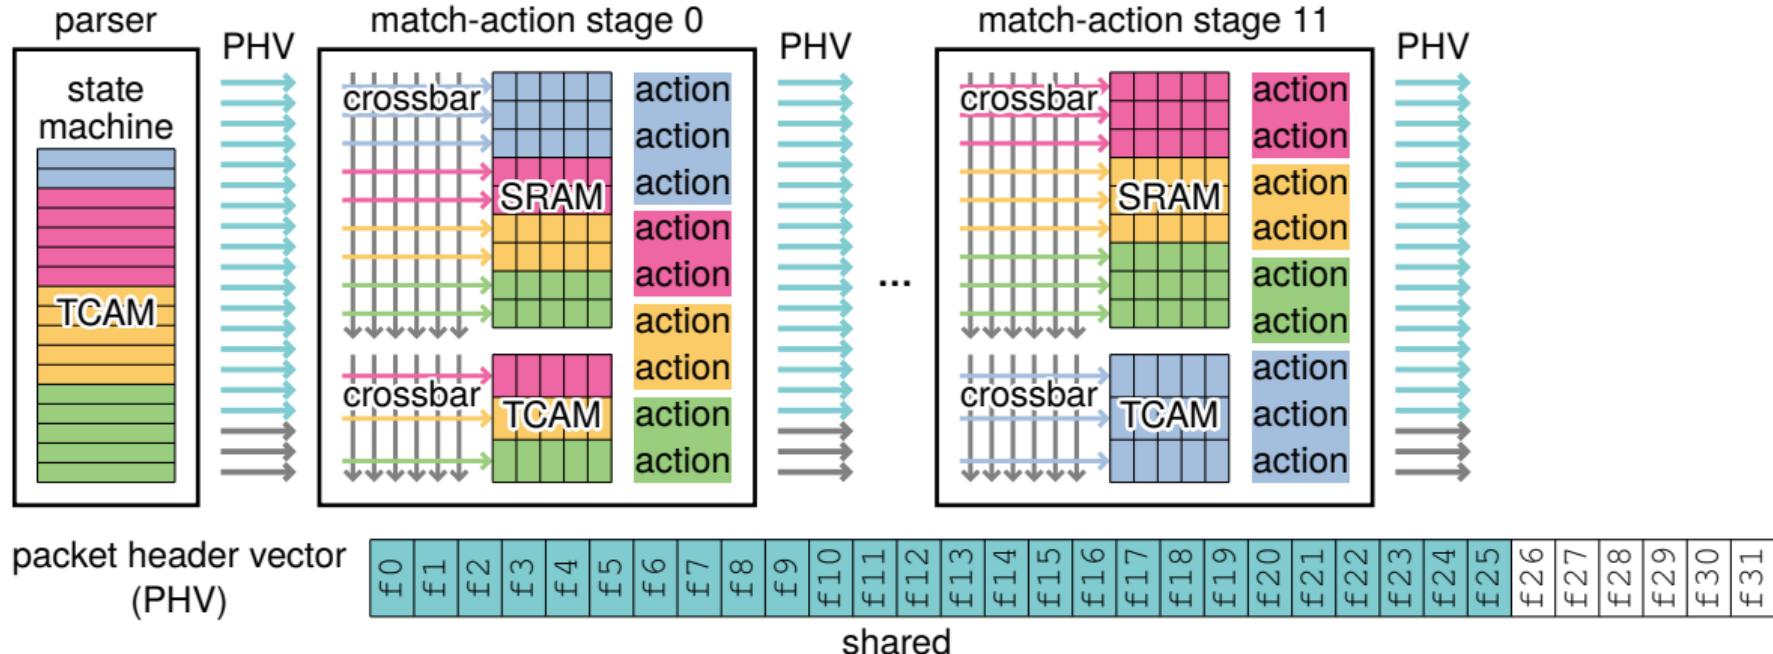
Slicing Individual Resources



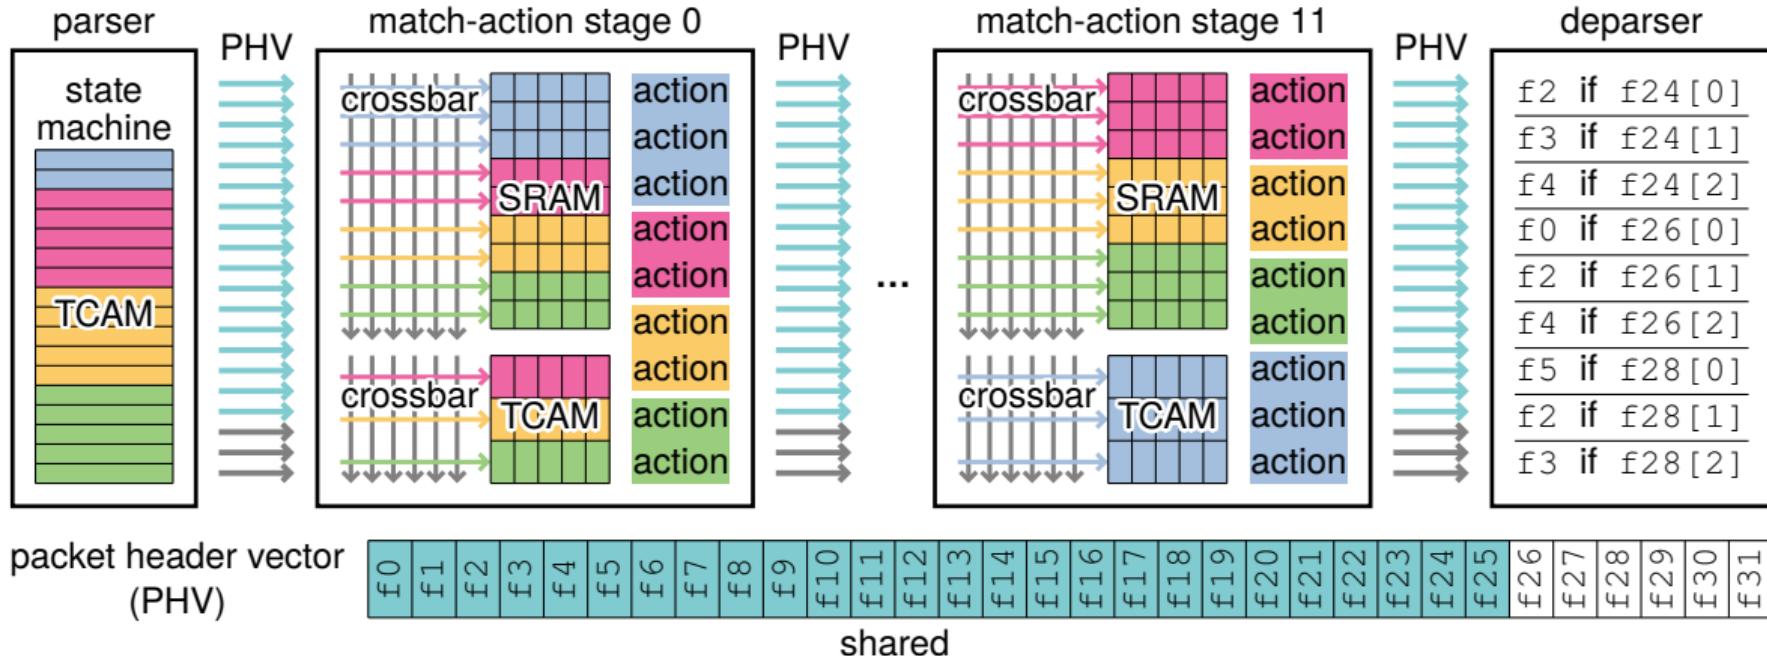
Slicing Individual Resources



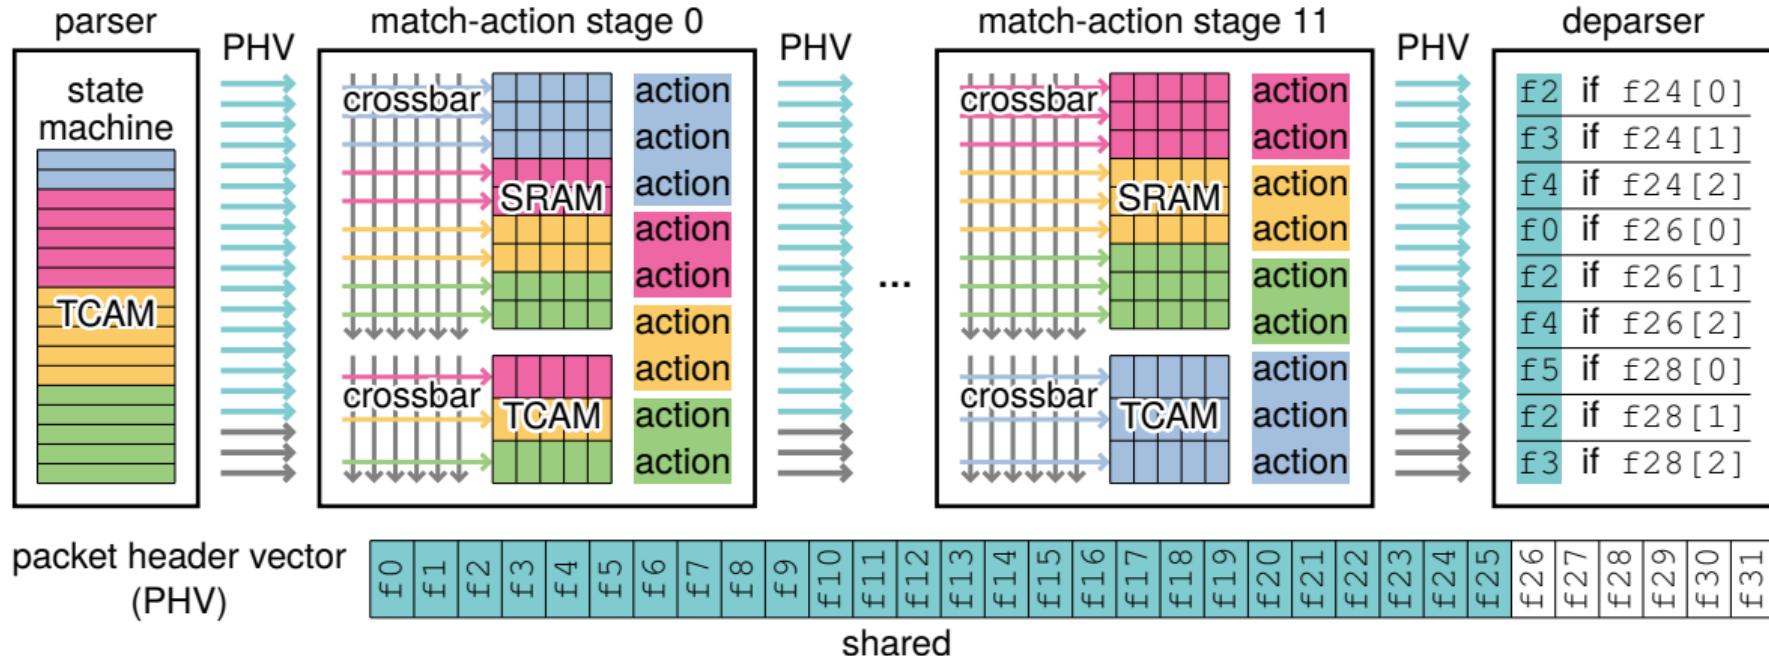
Slicing Individual Resources



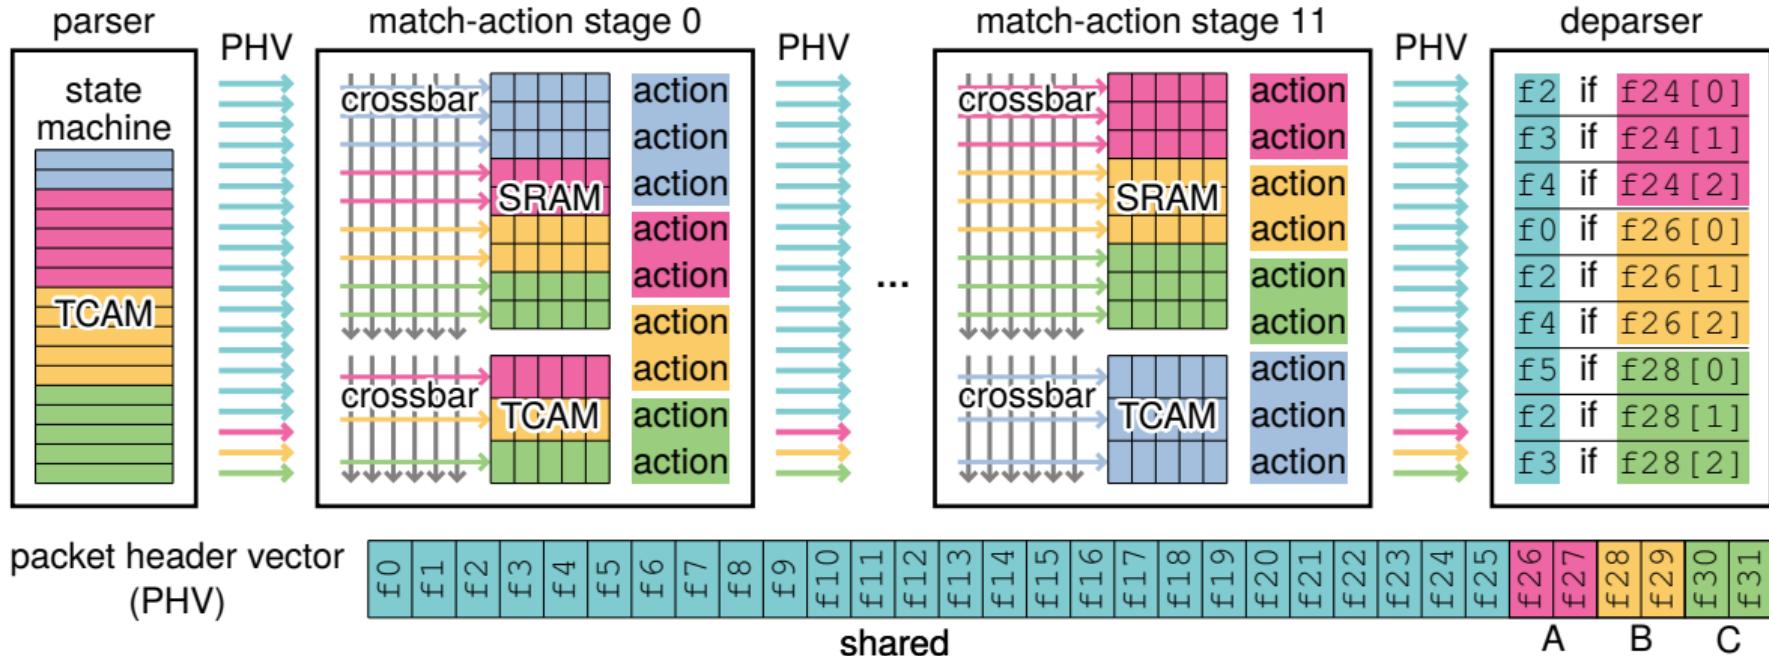
Slicing Individual Resources



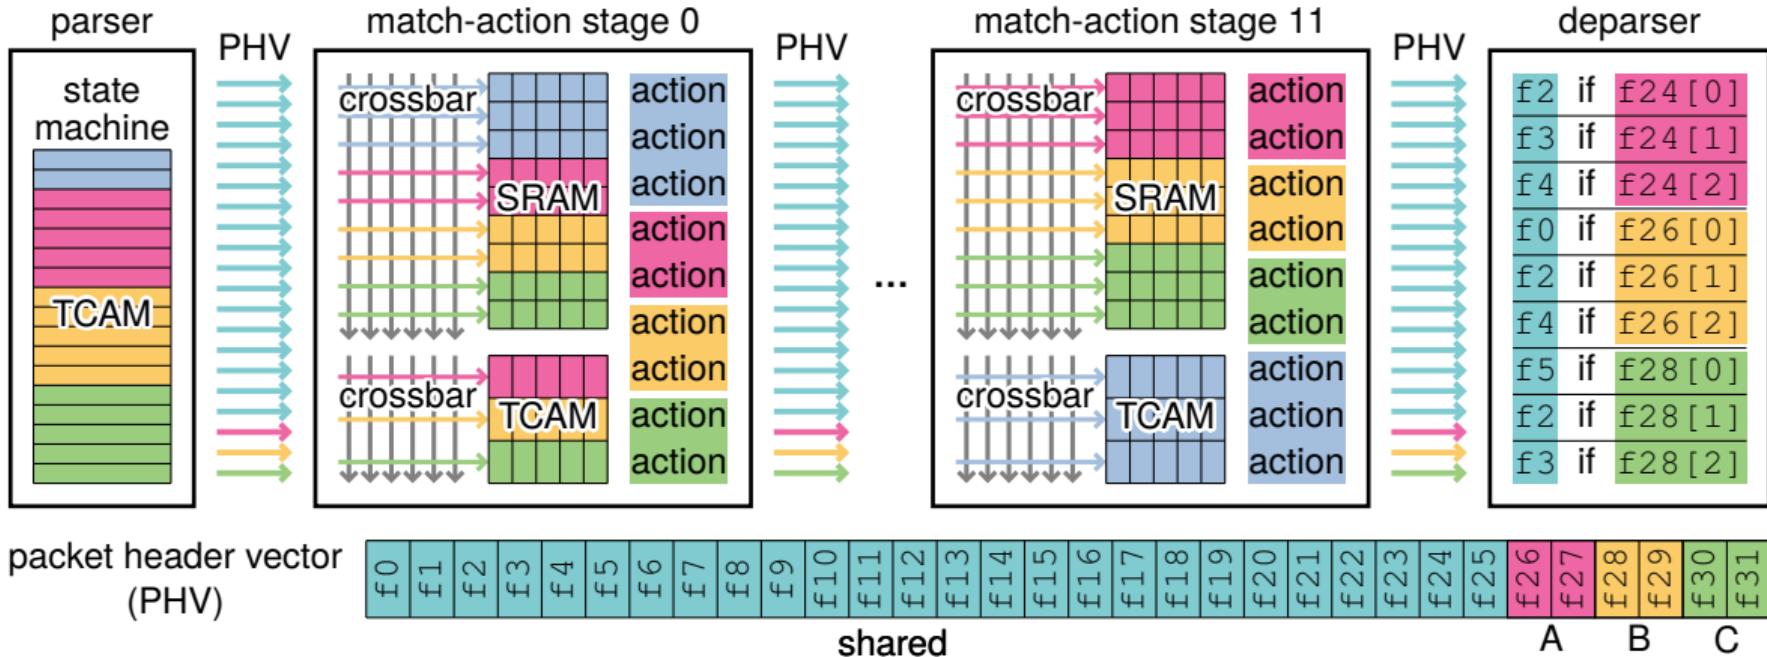
Slicing Individual Resources



Slicing Individual Resources



Slicing Individual Resources



Most of the packet header vector can be shared, except for the validity bits

- Each pipeline has 16x 100 GbE port → target up to 16 slices per pipeline

Parser

- Up to 256 slices

Deparser

- Up to 192 slices

Packet header vector

- Up to 48 slices, each with 16+16 validity bits and 128 shared fields

Slicing Tofino

- Each pipeline has 16x 100 GbE port → target up to 16 slices per pipeline

Parser

- Up to 256 slices

Deparser

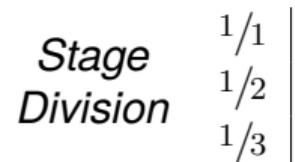
- Up to 192 slices

Packet header vector

- Up to 48 slices, each with 16+16 validity bits and 128 shared fields

Match Action Stages

- Crossbars only dividable into three identical slices



Slicing Tofino

- Each pipeline has 16x 100 GbE port → target up to 16 slices per pipeline

Parser

- Up to 256 slices

Deparser

- Up to 192 slices

Packet header vector

- Up to 48 slices, each with 16+16 validity bits and 128 shared fields

Match Action Stages

- Crossbars only dividable into three identical slices

	Stages per Slice					
	12	6	4	3	2	1
Stage Division	1/1	1/2	1/3			

Slicing Tofino

- Each pipeline has 16x 100 GbE port → target up to 16 slices per pipeline

Parser

- Up to 256 slices

Deparser

- Up to 192 slices

Packet header vector

- Up to 48 slices, each with 16+16 validity bits and 128 shared fields

Match Action Stages

- Crossbars only dividable into three identical slices

	Stages per Slice						Number of Slices
	12	6	4	3	2	1	
Stage Division	1/1	1	2	3	4	6	12
	1/2	2	4	6	8	12	24
	1/3	3	6	9	12	18	36

- Up to 36 identical slices

Slicing Tofino

- Each pipeline has 16x 100 GbE port → target up to 16 slices per pipeline

Parser

- Up to 256 slices

Deparser

- Up to 192 slices

Packet header vector

- Up to 48 slices, each with 16+16 validity bits and 128 shared fields

Match Action Stages

- Crossbars only dividable into three identical slices

	Stages per Slice						Number of Slices
	12	6	4	3	2	1	
Stage	1/1	1	2	3	4	6	12
Division	1/2	2	4	6	8	12	24
	1/3	3	6	9	12	18	36

- Up to 36 identical slices

Implemented slice checking and program composition on Barefoot Assembly (.bfa) files

Implemented 3 variants

1. Modified bf-p4c

- No PHV constraints

2. SMT table placement

- Reordering stages

3. Add additional tables

- Unmodified compiler

Compiling P4 for a Slice

Implemented 3 variants

1. Modified bf-p4c

- No PHV constraints

2. SMT table placement

- Reordering stages

3. Add additional tables

- Unmodified compiler

Compiled 19 P4 programs from various sources

Program	Smallest Slice	Program	Smallest Slice
RTT	12.1/1	IXP_routing	4.1/1
firewall	6.1/1, 12.1/2	aes	12.1/3
chord	4.1/1, 12.1/3	conquest_b.	2.1/1, 4.1/2
accelerator	4.1/2	speedtest	2.1/1, 4.1/2, 6.1/3
source_routing	4.1/3	a._tunnel	2.1/2
flowlet	2.1/2	firewall_s.	1.1/1, 2.1/2, 3.1/3
qos	2.1/3	calc	1.1/2
ecn	2.1/3	single_table	1.1/3
MPLS	1.1/2	VXLAN	3.1/2
		multicast	1.1/2

Compiling P4 for a Slice

Implemented 3 variants

1. Modified bf-p4c

- No PHV constraints

2. SMT table placement

- Reordering stages

3. Add additional tables

- Unmodified compiler

Compiled 19 P4 programs from various sources

Program	Prog. Instances	Program	Prog. Instances
	Slicing		Slicing
RTT	1	IXP_routing	3
firewall	2	aes	3
chord	3	conquest_b.	6
accelerator	6	speedtest	6
source_routing	9	a._tunnel	12
flowlet	12	firewall_s.	12
qos	18	calc	24
ecn	18	single_table	36
MPLS	24	VXLAN	8
		multicast	24

Compiling P4 for a Slice

Implemented 3 variants

1. Modified bf-p4c

- No PHV constraints

2. SMT table placement

- Reordering stages

3. Add additional tables

- Unmodified compiler

X. Merge P4 programs

- Similar to related work
- Compile merged P4

Compiled 19 P4 programs from various sources

Program	Prog. Instances		Program	Prog. Instances	
	Slicing	Merge		Slicing	Merge
RTT	1		IXP_routing	3	
firewall	2		aes	3	
chord	3		conquest_b.	6	
accelerator	6		speedtest	6	
source_routing	9		a._tunnel	12	
flowlet	12		firewall_s.	12	
qos	18		calc	24	
ecn	18		single_table	36	
MPLS	24		VXLAN	8	
			multicast	24	

Compiling P4 for a Slice

Implemented 3 variants

1. Modified bf-p4c

- No PHV constraints

2. SMT table placement

- Reordering stages

3. Add additional tables

- Unmodified compiler

X. Merge P4 programs

- Similar to related work
- Compile merged P4

Compiled 19 P4 programs from various sources

Program	Prog. Instances		Program	Prog. Instances	
	Slicing	Merge		Slicing	Merge
RTT	1	<	2	IXP_routing	3
firewall	2	<	4	aes	3
chord	3	<	7	conquest_b.	6
accelerator	6	<	9	speedtest	6
source_routing	9		a._tunnel	12	
flowlet	12		firewall_s.	12	
qos	18		calc	24	
ecn	18		single_table	36	
MPLS	24		VXLAN	8	
			multicast	24	

Compiling P4 for a Slice

Implemented 3 variants

1. Modified bf-p4c

- No PHV constraints

2. SMT table placement

- Reordering stages

3. Add additional tables

- Unmodified compiler

X. Merge P4 programs

- Similar to related work
- Compile merged P4

Compiled 19 P4 programs from various sources

Program	Prog. Instances		Program	Prog. Instances	
	Slicing	Merge		Slicing	Merge
RTT	1	<	2	IXP_routing	3
firewall	2	<	4	aes	3
chord	3	<	7	conquest_b.	6
accelerator	6	<	9	speedtest	6
source_routing	9	≈	10	a._tunnel	12
flowlet	12	≈	13	firewall_s.	12
qos	18	≈	21	calc	24
ecn	18	≈	21	single_table	36
MPLS	24	≈	32	VXLAN	8
				multicast	24

Compiling P4 for a Slice

Implemented 3 variants

1. Modified bf-p4c

- No PHV constraints

2. SMT table placement

- Reordering stages

3. Add additional tables

- Unmodified compiler

X. Merge P4 programs

- Similar to related work
- Compile merged P4

Compiled 19 P4 programs from various sources

Program	Prog. Instances		Program	Prog. Instances	
	Slicing	Merge		Slicing	Merge
RTT	1 <	2	IXP_routing	3 =	3
firewall	2 <	4	aes	3 =	3
chord	3 <	7	conquest_b.	6 =	6
accelerator	6 <	9	speedtest	6 =	6
source_routing	9 ≈	10	a._tunnel	12 =	12
flowlet	12 ≈	13	firewall_s.	12 =	12
qos	18 ≈	21	calc	24 =	24
ecn	18 ≈	21	single_table	36 =	36
MPLS	24 ≈	32	VXLAN	8	
			multicast	24	

Compiling P4 for a Slice

Implemented 3 variants

1. Modified bf-p4c

- No PHV constraints

2. SMT table placement

- Reordering stages

3. Add additional tables

- Unmodified compiler

X. Merge P4 programs

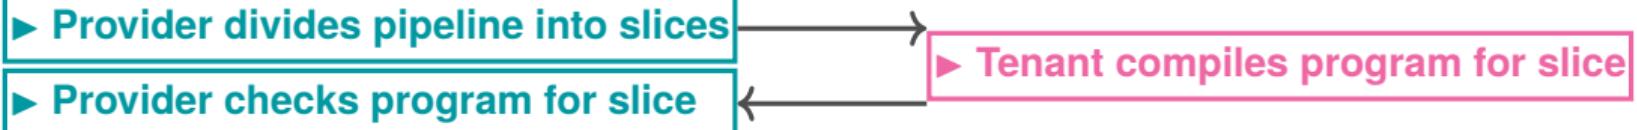
- Similar to related work
- Compile merged P4

Compiled 19 P4 programs from various sources

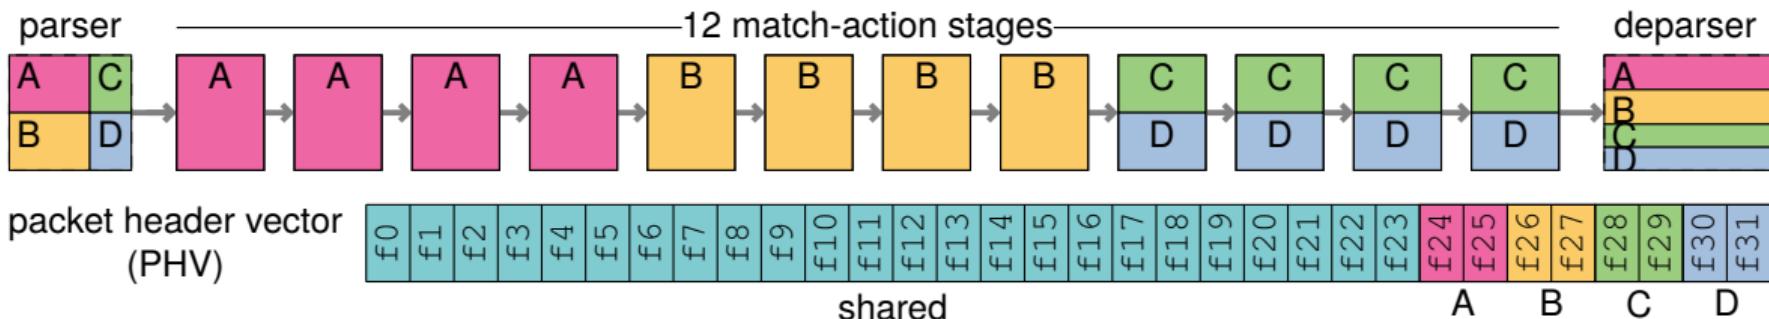
Program	Prog. Instances		Program	Prog. Instances	
	Slicing	Merge		Slicing	Merge
RTT	1 <	2	IXP_routing	3 =	3
firewall	2 <	4	aes	3 =	3
chord	3 <	7	conquest_b.	6 =	6
accelerator	6 <	9	speedtest	6 =	6
source_routing	9 ≈	10	a._tunnel	12 =	12
flowlet	12 ≈	13	firewall_s.	12 =	12
qos	18 ≈	21	calc	24 =	24
ecn	18 ≈	21	single_table	36 =	36
MPLS	24 ≈	32	VXLAN	8 >	6
			multicast	24 >	15

Conclusion

- Resource guarantees for multitenancy on match-action pipelines



- Our packet header vector sharing approach can increase switch utilization



Full implementation, all example programs, and scripted eval on GitHub