

Scaling Open vSwitch with a Computational Cache


Alon Rashelbach

Ori Rottenstreich

Mark Silberstein

Technion, Israel


USENIX NSDI 2022



- Virtual Switch
- Open-source
- Data-centers

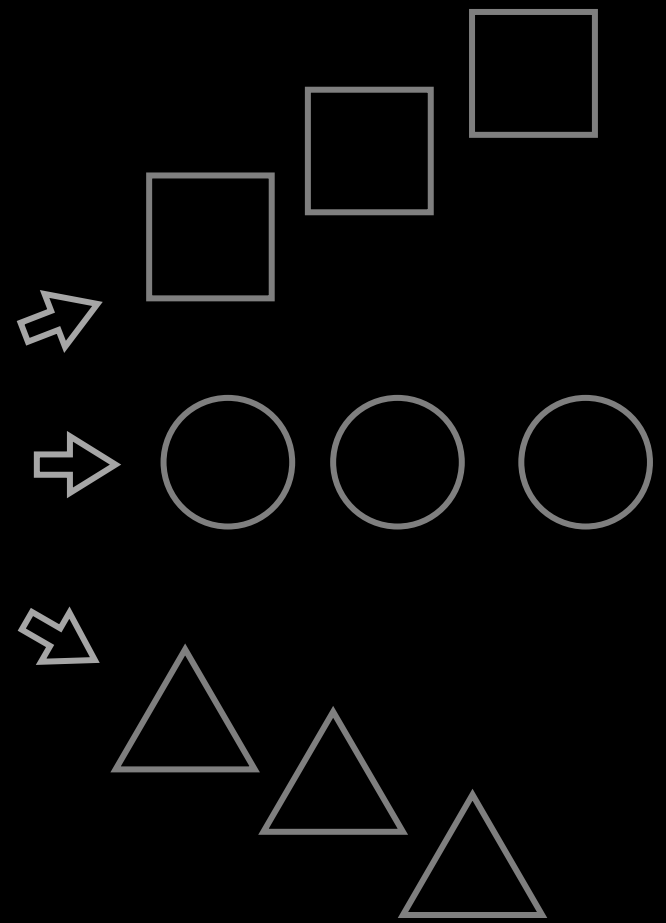
OvS


Open vSwitch



OvS

Open vSwitch






1K OpenFlow Rules



500K OpenFlow Rules




1K OpenFlow Rules



500K OpenFlow Rules

14x geomean throughput degradation

In some cases: throughput drops from a few **Mpps** to a few **Kpps**



1K OpenFlow Rules




500K OpenFlow Rules

14x geomean throughput degradation

In some cases: throughput drops from a few Mpps to a few Kpps

In this paper: a geomean speedup of 12x



1K OpenFlow Rules



500K OpenFlow Rules

14x geomean throughput degradation

In some cases: throughput drops from a few Mpps to a few Kpps

In this paper: a geomean speedup of 12x

Neural-net inference per packet in OVS data-path

Background

Challenge

This paper

Integration with OVS

Evaluation

NuevoMatch algorithm for packet classification

NuevoMatch algorithm for packet classification

OpenFlow rules

Src IP	Dst IP	Action
10.0.10.*	56.0.*.*	Port 1
10.0.*.*	56.22.7.1	Drop


⋮

NuevoMatch algorithm for packet classification

OpenFlow rules

Src IP	Dst IP	Action
10.0.10.*	56.0.*.*	Port 1
10.0.*.*	56.22.7.1	Drop

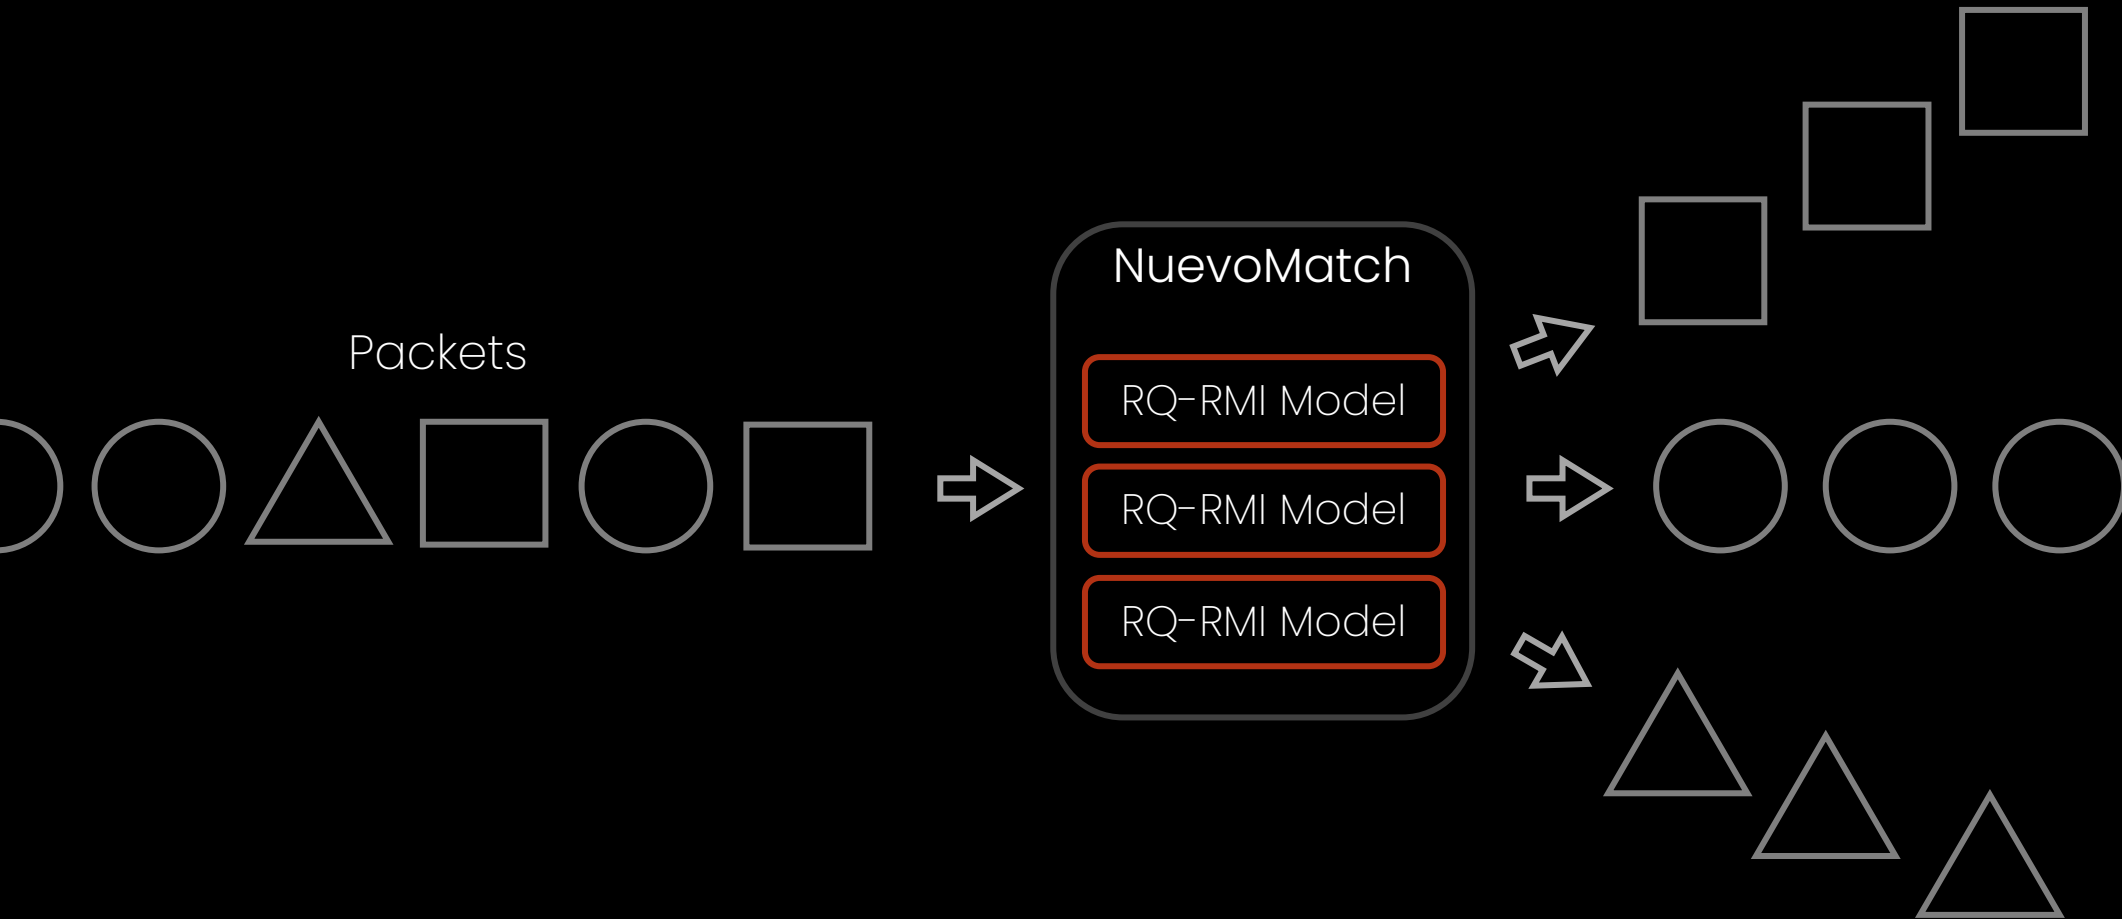
⋮



Unique model architecture
40ns NN inference on CPU




NuevoMatch algorithm for packet classification



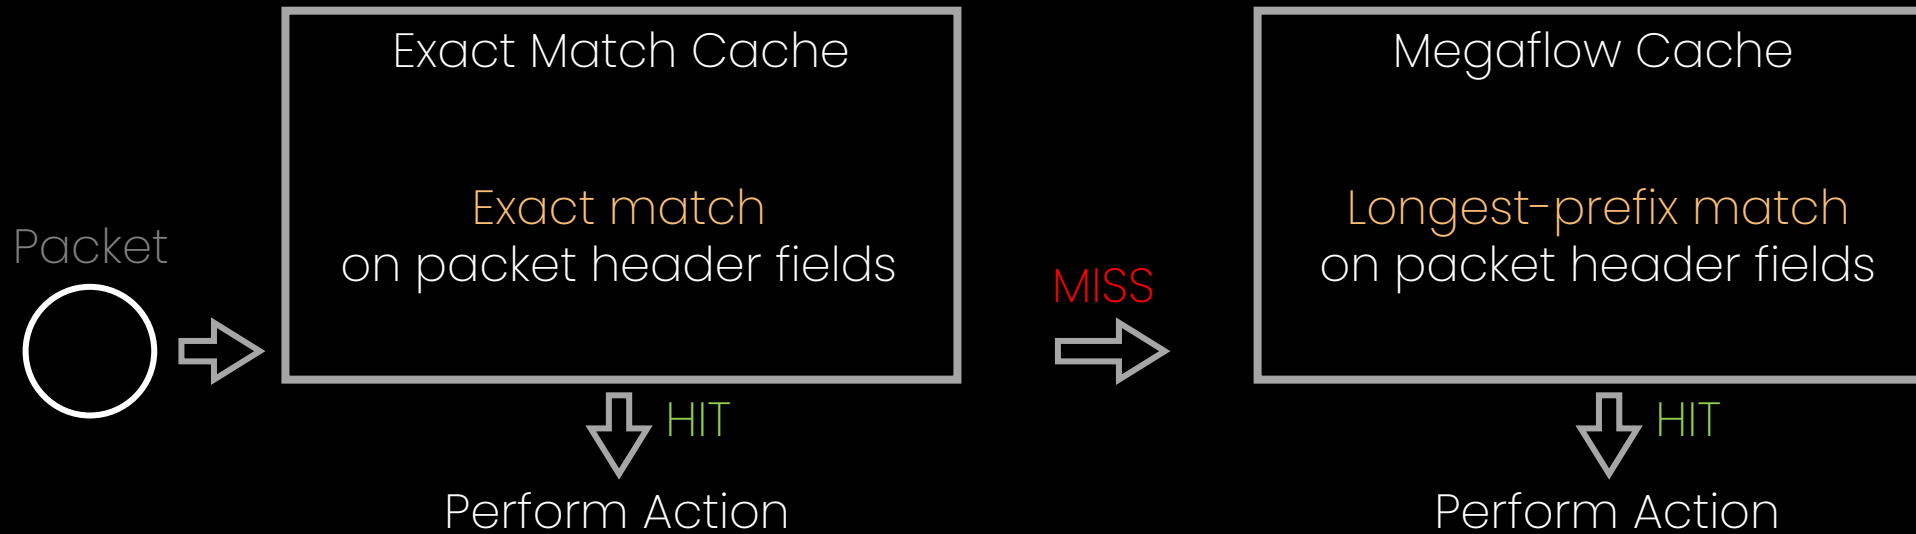
Classifies packets on the **CPU** using **Neural Network Inference**

OVS Data-path


OVS Data-path



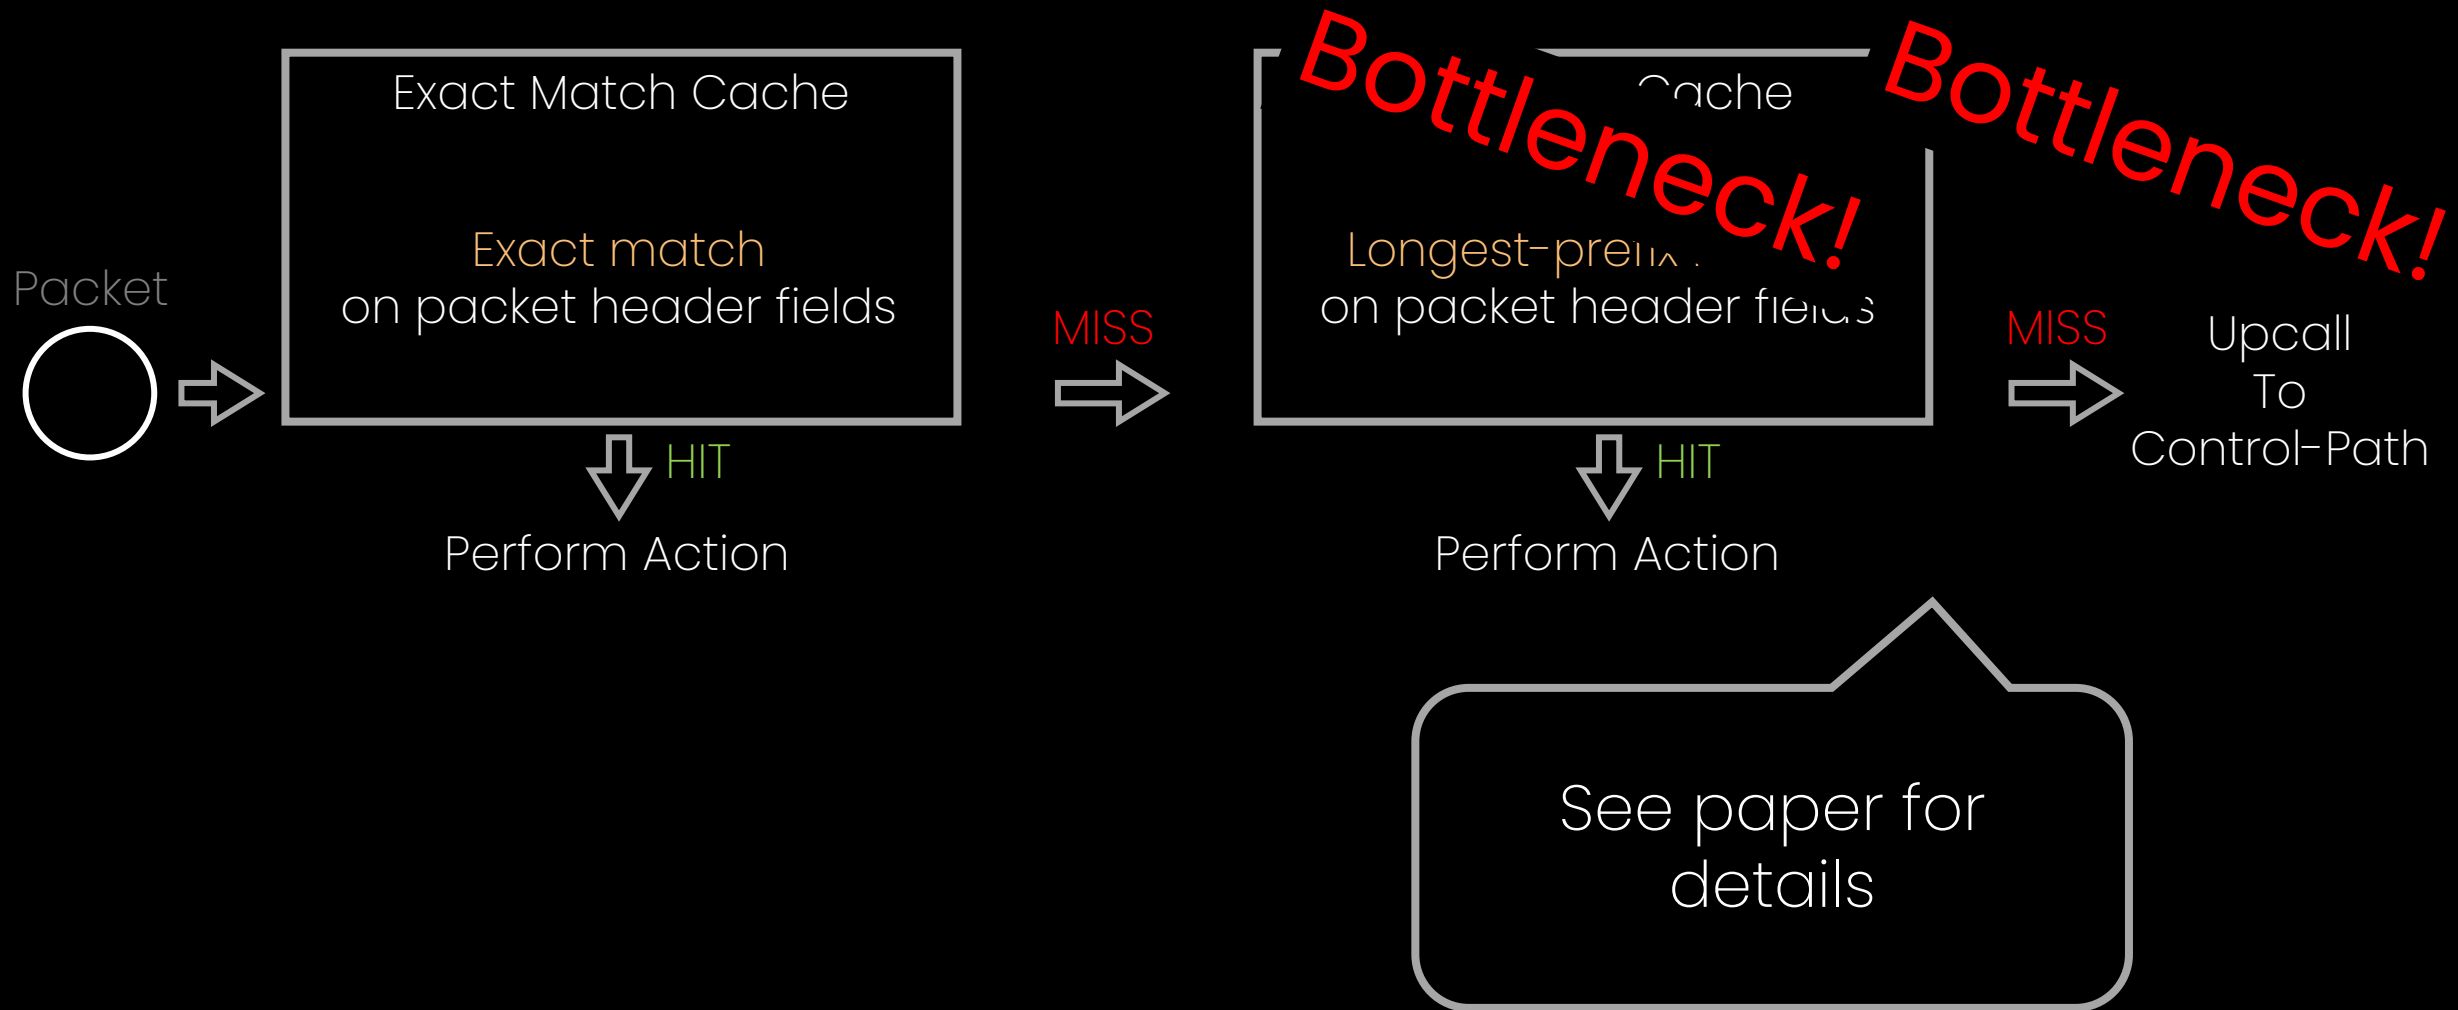
OVS Data-path



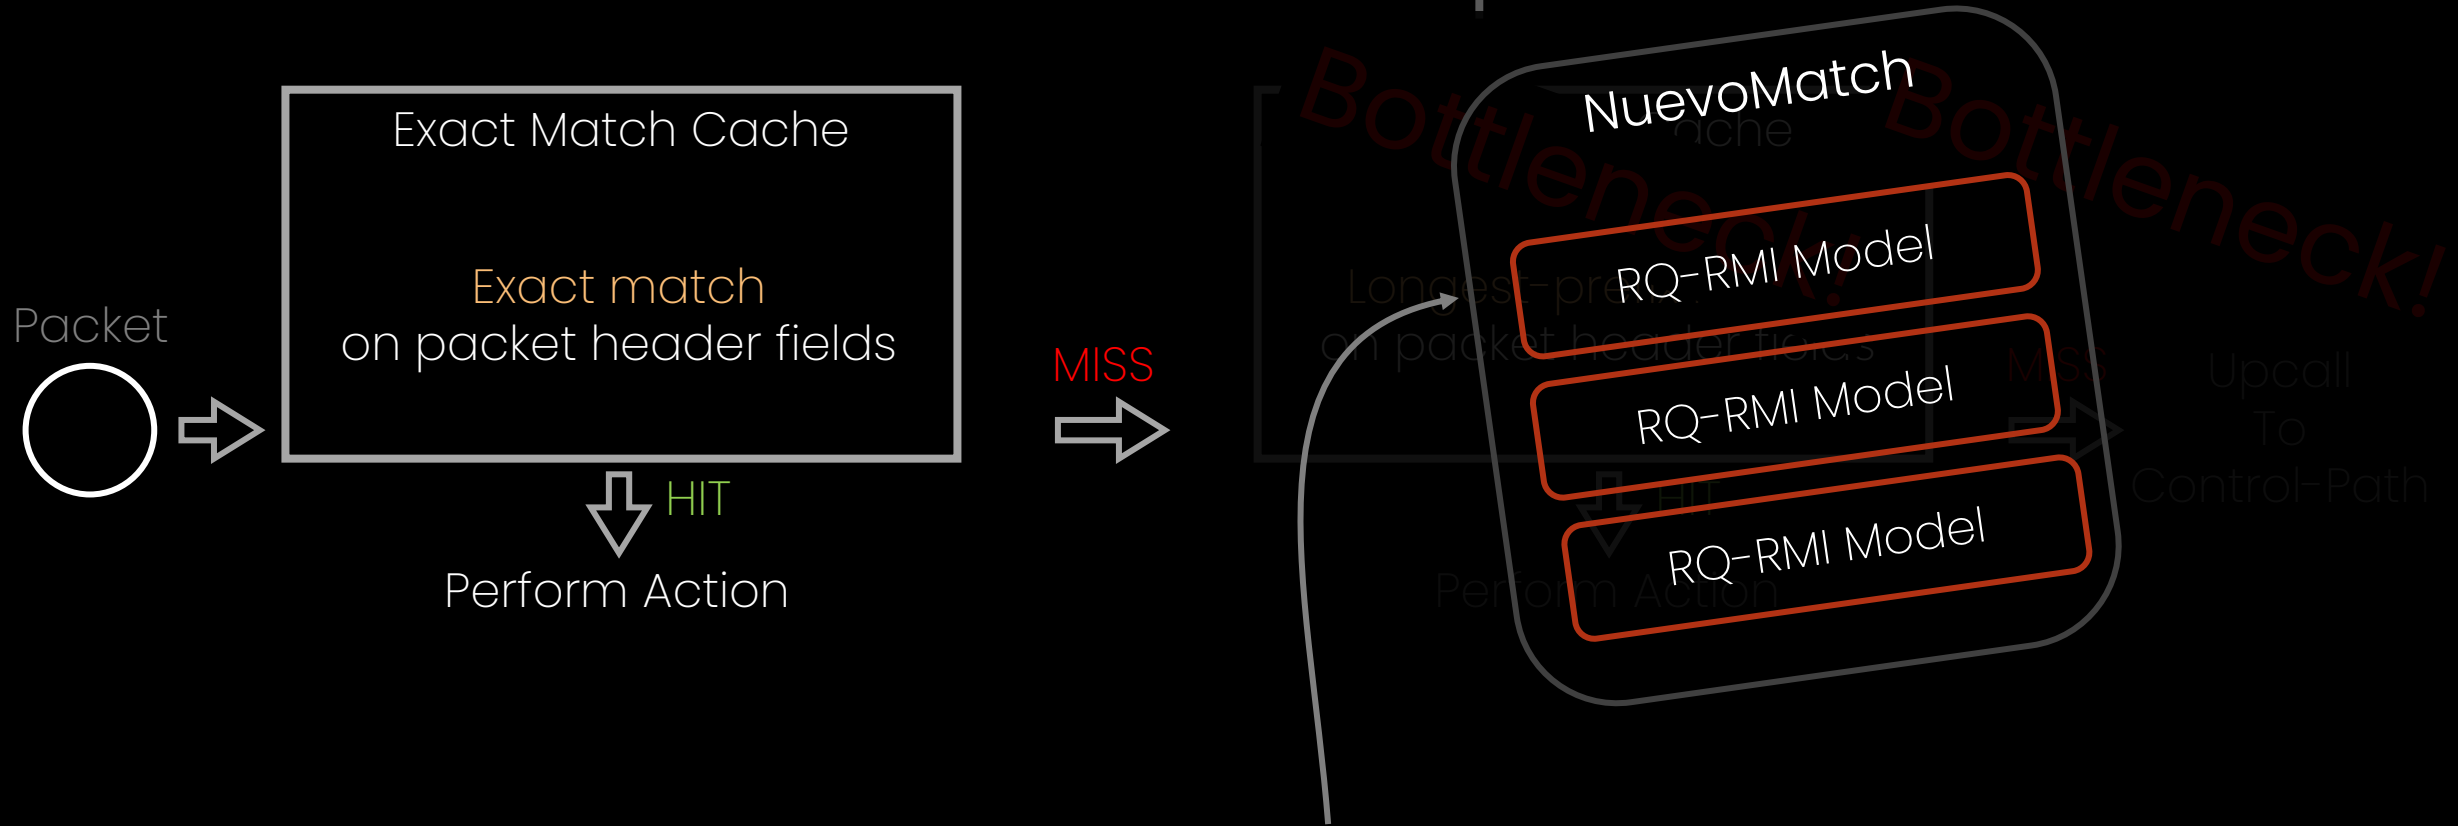
OVS Data-path



OVS Data-path



OVS Data-path



Can we use the **NuevoMatch** algorithm within the **OVS** data-path?

NO!

Can we use the **NuevoMatch** algorithm
within the **OVS** data-path?

Challenge: Slow Updates

Challenge: Slow Updates

NuevoMatch uses neural-nets
to learn rules

Neural-nets must be trained
when rules change

The training time depends
on the total number of rules

Challenge: Slow Updates

NuevoMatch uses neural-nets
to learn rules

Neural-nets must be trained
when rules change

The training time depends
on the total number of rules

Train **once**

Challenge: Slow Updates

NuevoMatch uses neural-nets
to learn rules

Neural-nets must be trained
when rules change

The training time depends
on the total number of rules

Train ~~once~~

Train fast

Must Support Frequent Updates

100K updates per **second** (on average)

Must Support Frequent Updates

100K updates per **second** (on average)

Training an RQ-RMI model with 100K+ rules
takes **5-20 minutes!**

Must Support Frequent Updates

100K updates per **second** (on average)

Training an RQ-RMI model with 100K+ rules
takes 5-20 minutes!

NuevoMatch update rate is **1000x** slower!

Background: RQ-RMI models in NuevoMatch

*Range: a
single rule field*



Row num	Range
1	0-100
2	130-700
3	800-803
4	807-890
5	970-990
6	991-991

Background: RQ-RMI models in NuevoMatch

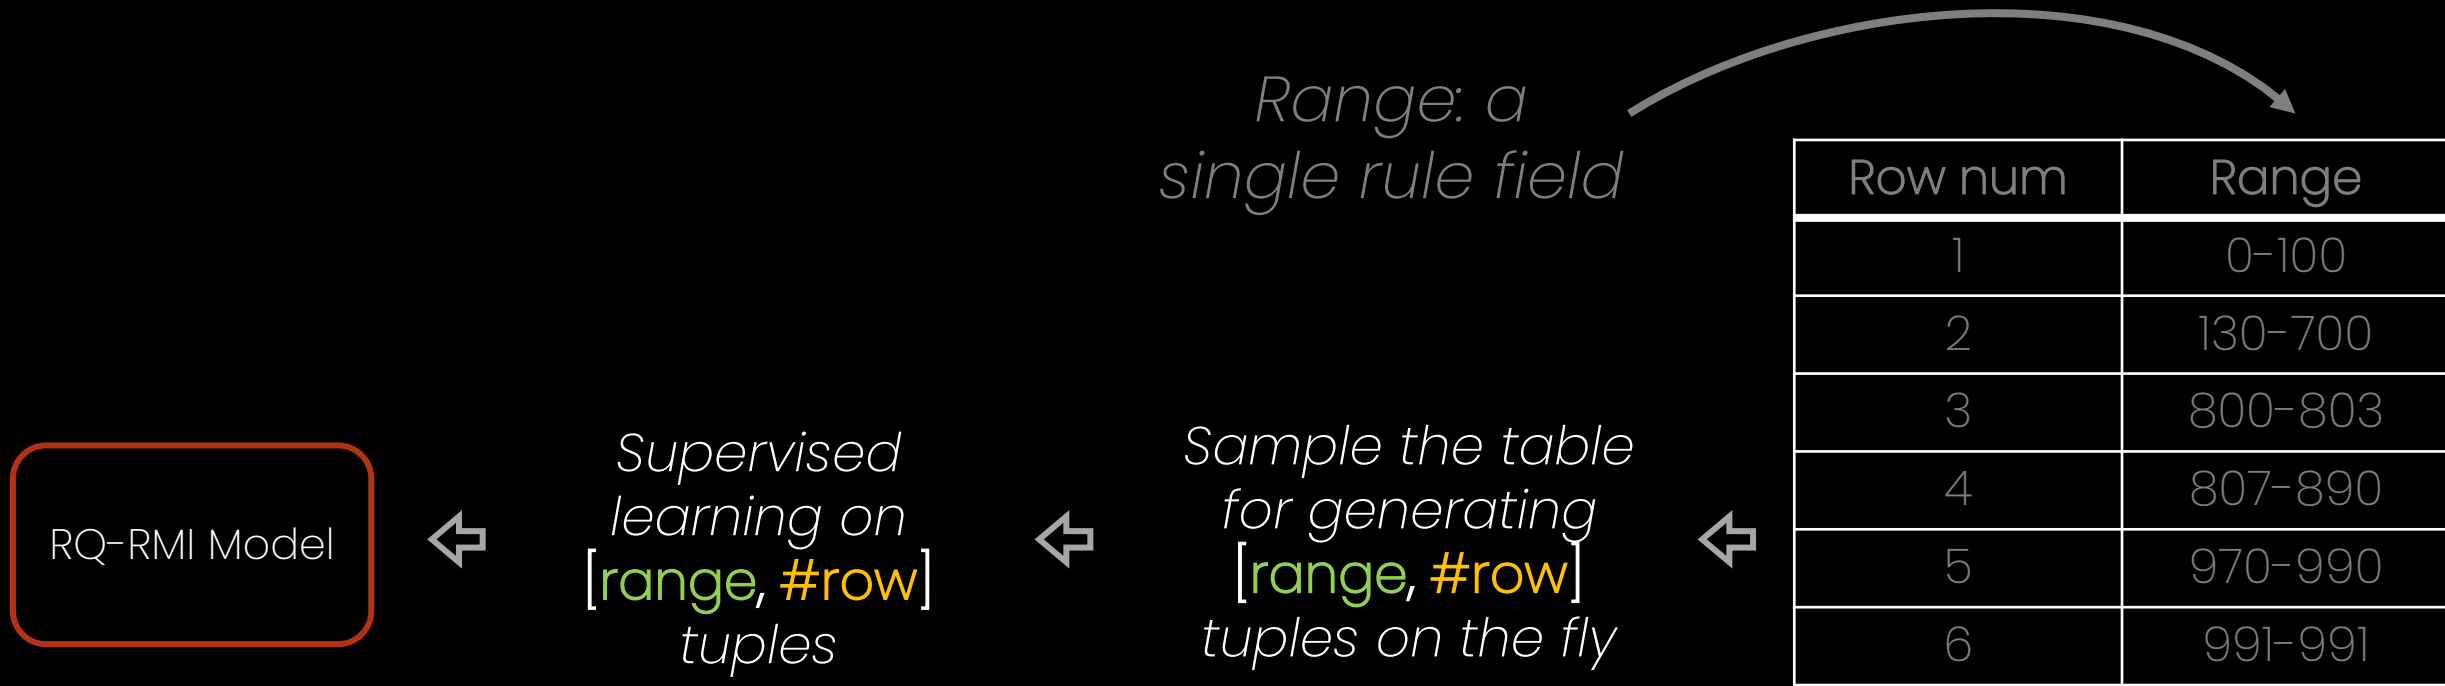
Range: a
single rule field

Sample the table
for generating
[range, #row]
tuples on the fly

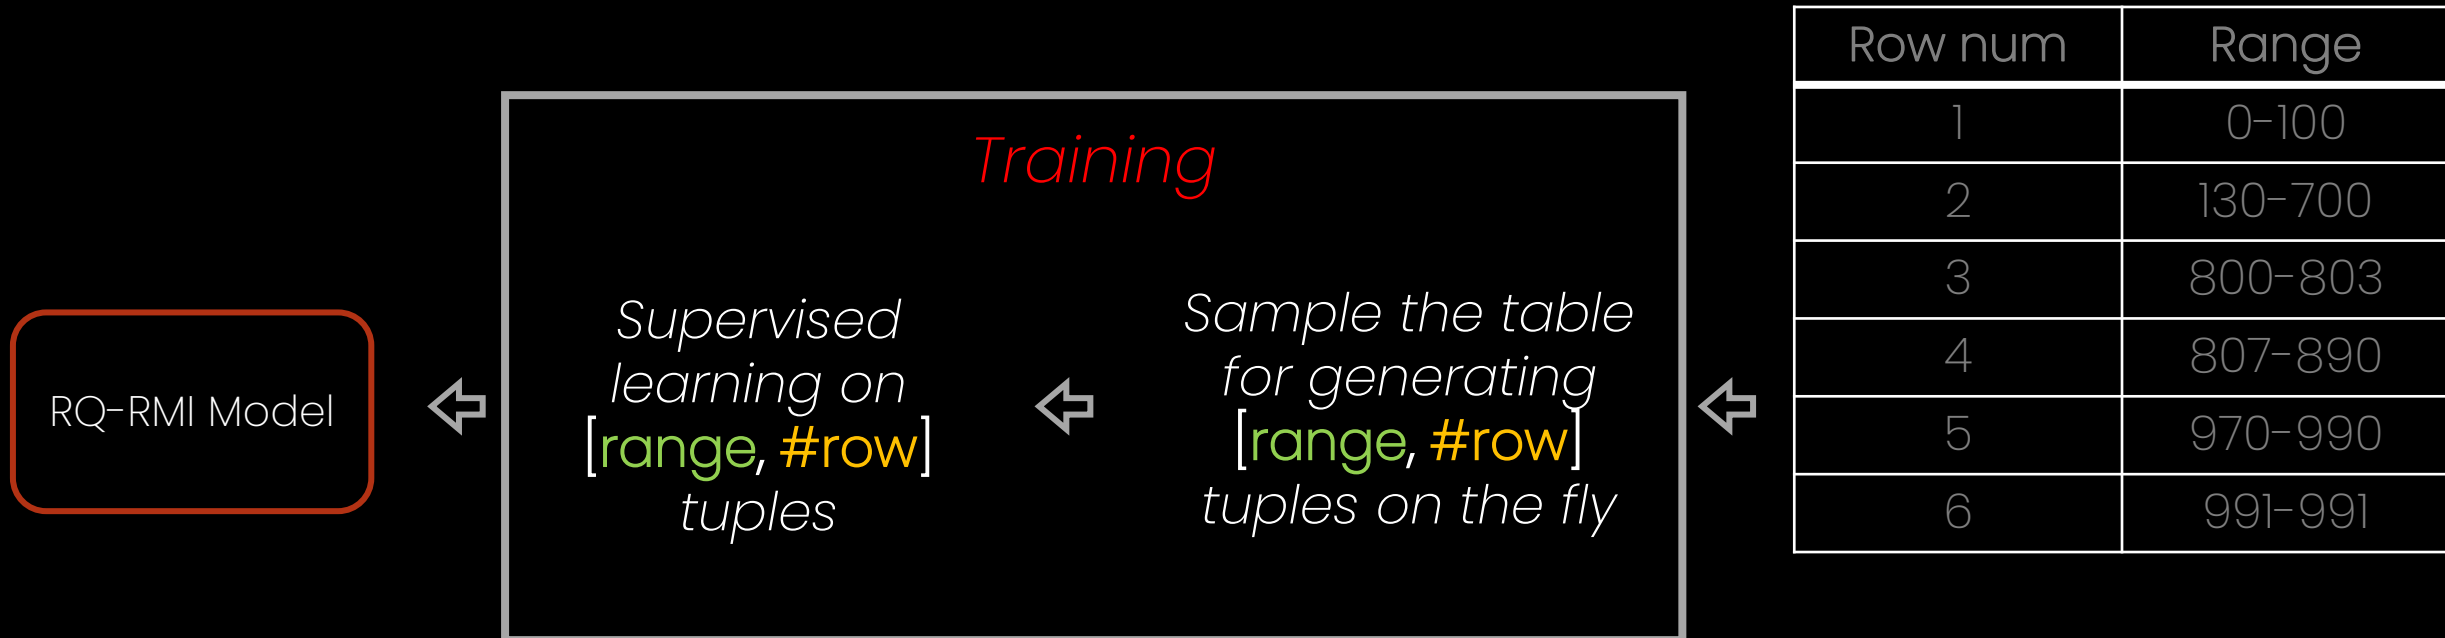


Row num	Range
1	0-100
2	130-700
3	800-803
4	807-890
5	970-990
6	991-991

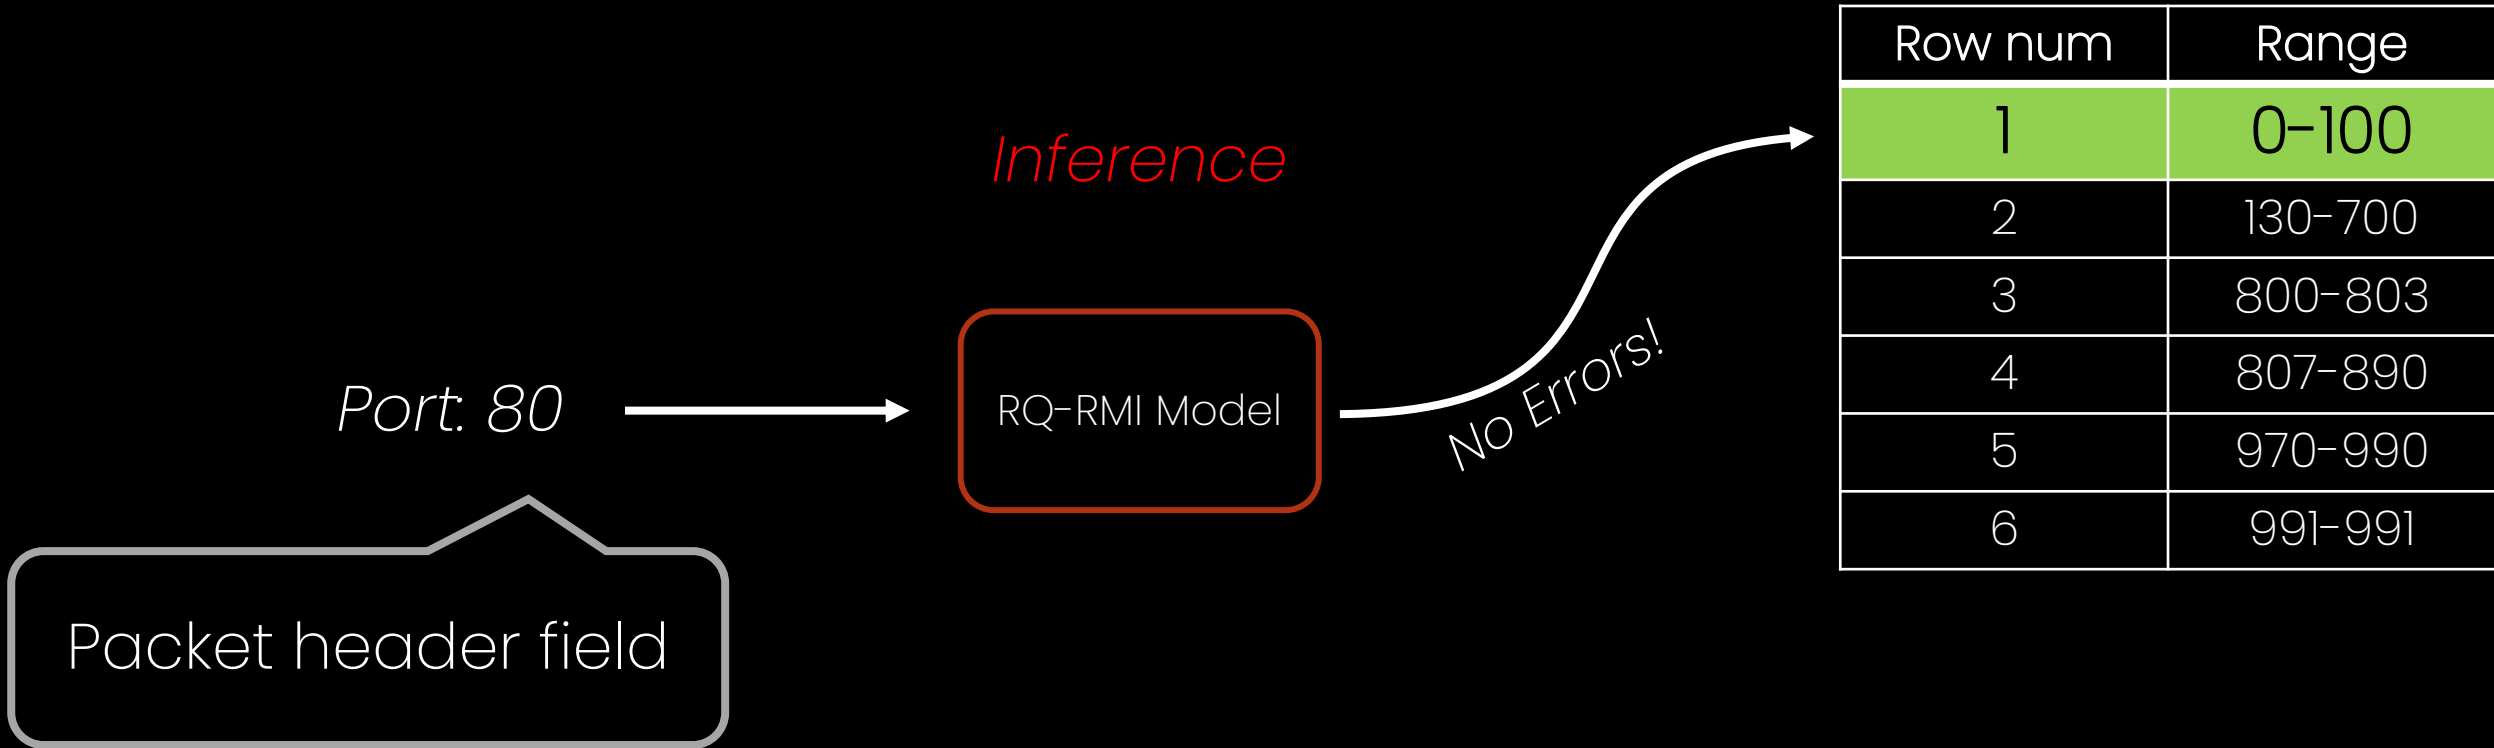
Background: RQ-RMI models in NuevoMatch



Background: RQ-RMI models in NuevoMatch



Background: RQ-RMI models in NuevoMatch



Background

Challenge

This paper

Integration with OVS

Evaluation

introducing *NuevoMatchUP*
speeding-up NuevoMatch updates

introducing *NuevoMatchUP*
speeding-up NuevoMatch updates

Three orthogonal improvements:

Bucketizing Ranges

Efficient SIMD implementation

Optimized sampling (see paper)

introducing *NuevoMatchUP*
speeding-up NuevoMatch updates

Bucketizing Ranges

Row num	Range
1	0-10
2	20-30
3	40-50
4	60-70
5	80-90
6	100-110

introducing *NuevoMatchUP*
speeding-up NuevoMatch updates

Bucketizing Ranges

Row num	Bucket	Ranges
1	0-30	0-10 20-30
2	40-70	40-50 60-70
3	80-110	80-90 100-110

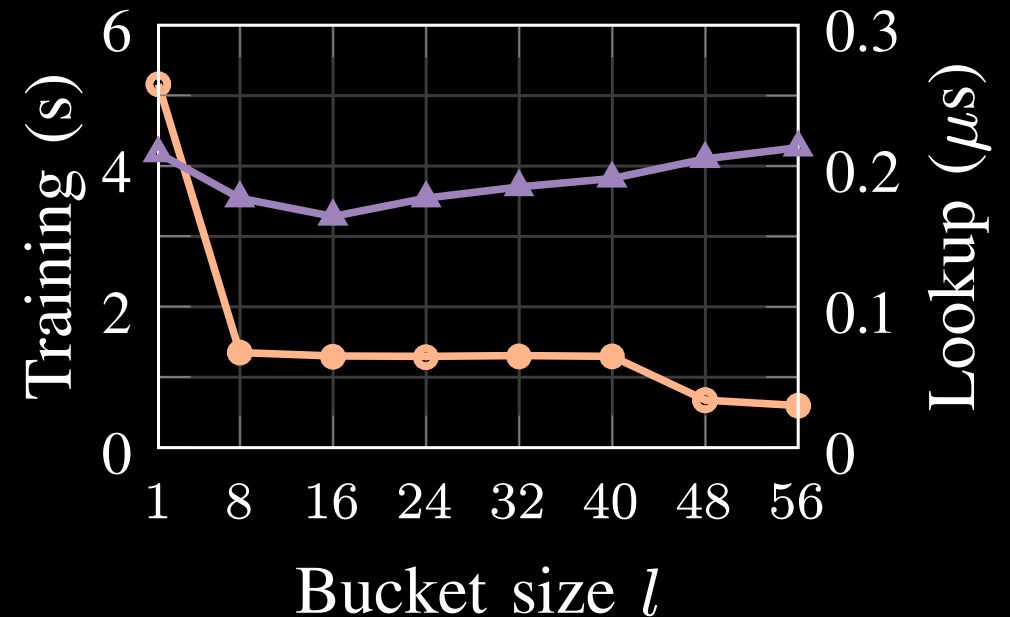


Row num	Range
1	0-10
2	20-30
3	40-50
4	60-70
5	80-90
6	100-110

introducing *NuevoMatchUP*
speeding-up NuevoMatch updates

Bucketizing Ranges

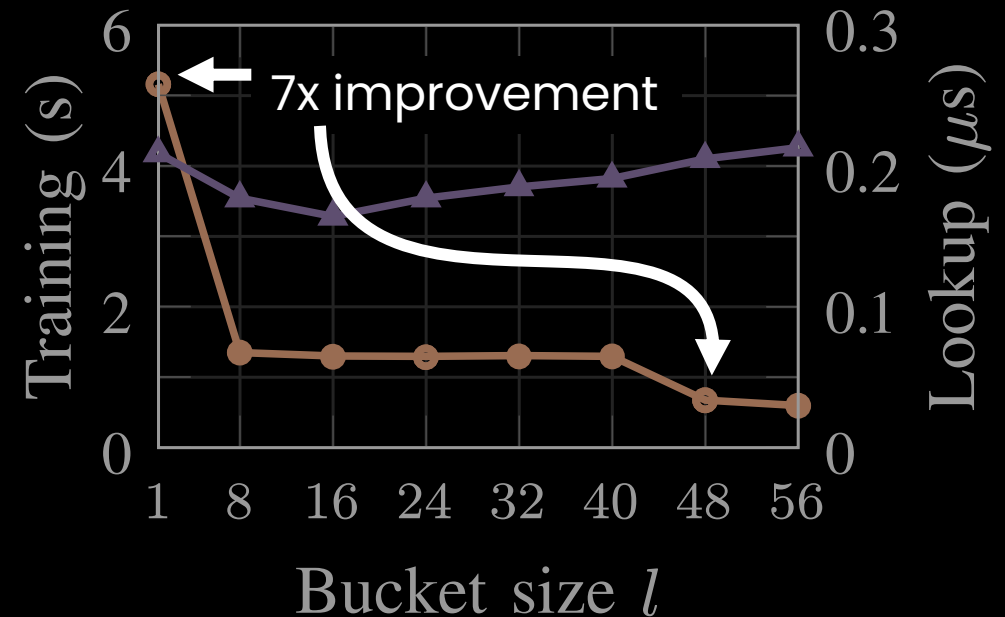

Row num	Bucket	Ranges
1	0-30	0-10 20-30
2	40-70	40-50 60-70
3	80-110	80-90 100-110



introducing *NuevoMatchUP*
speeding-up NuevoMatch updates

Bucketizing Ranges

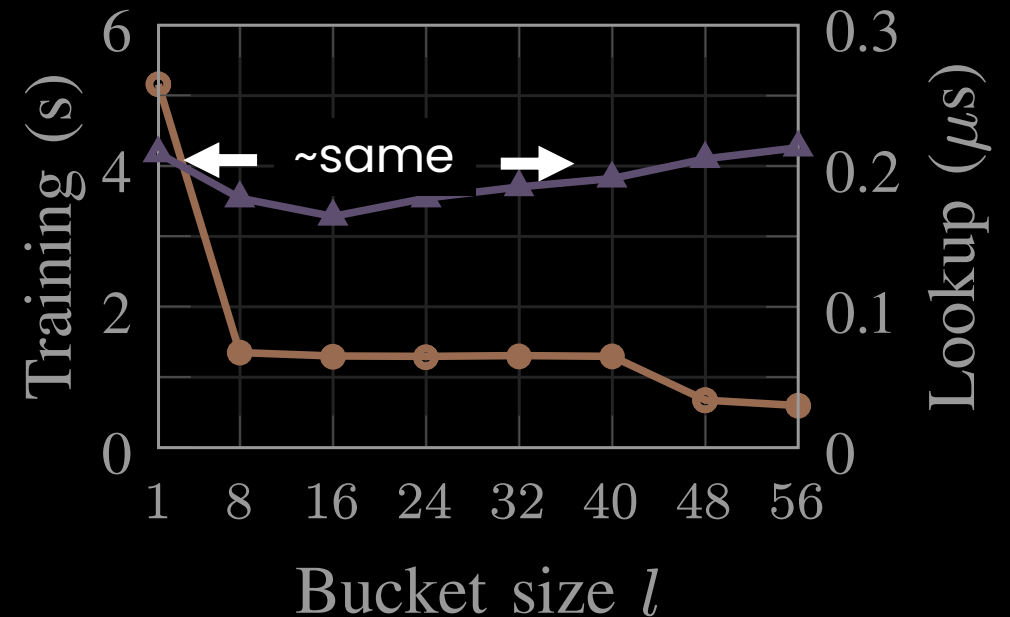

Row num	Bucket	Ranges	
1	0-30	0-10	20-30
2	40-70	40-50	60-70
3	80-110	80-90	100-110



introducing *NuevoMatchUP*
speeding-up NuevoMatch updates

Bucketizing Ranges

Row num	Bucket	Ranges
1	0-30	0-10 20-30
2	40-70	40-50 60-70
3	80-110	80-90 100-110




introducing *NuevoMatchUP*
speeding-up NuevoMatch updates

Efficient SIMD Implementation

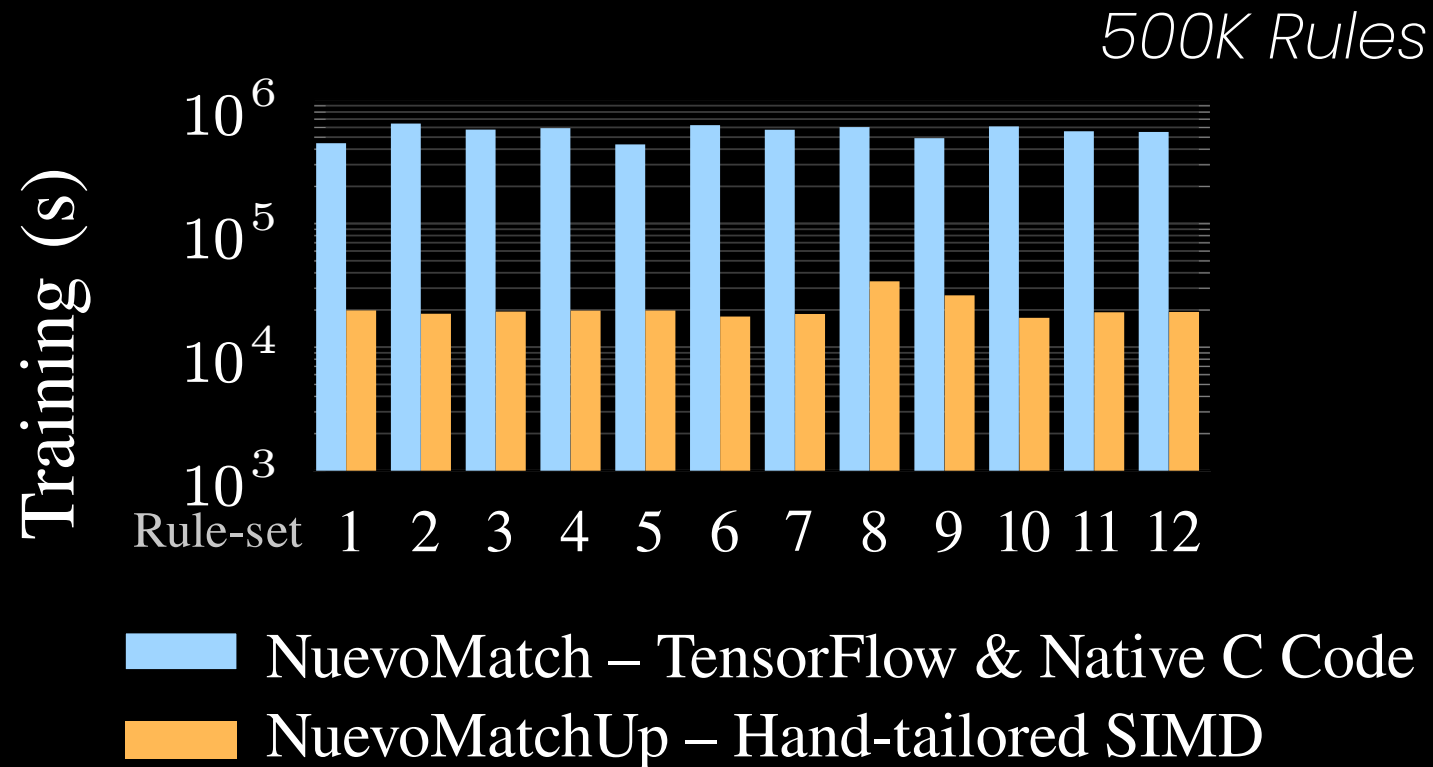
introducing *NuevoMatchUP*
speeding-up NuevoMatch updates

Efficient SIMD Implementation




introducing *NuevoMatchUP*
speeding-up NuevoMatch updates

Efficient SIMD Implementation



introducing *NuevoMatchUP*
speeding-up NuevoMatch updates

Efficient SIMD Implementation



introducing *NuevoMatchUP*
speeding-up NuevoMatch updates

Together:
Three orders of magnitude faster training
over NuevoMatch

introducing *NuevoMatchUP*
speeding-up NuevoMatch updates

Together:

Three orders of magnitude faster training
over NuevoMatch

NuevoMatchUP can be integrated into OVS data-path!

Background

Challenge

This paper

Integration with OVS

Evaluation

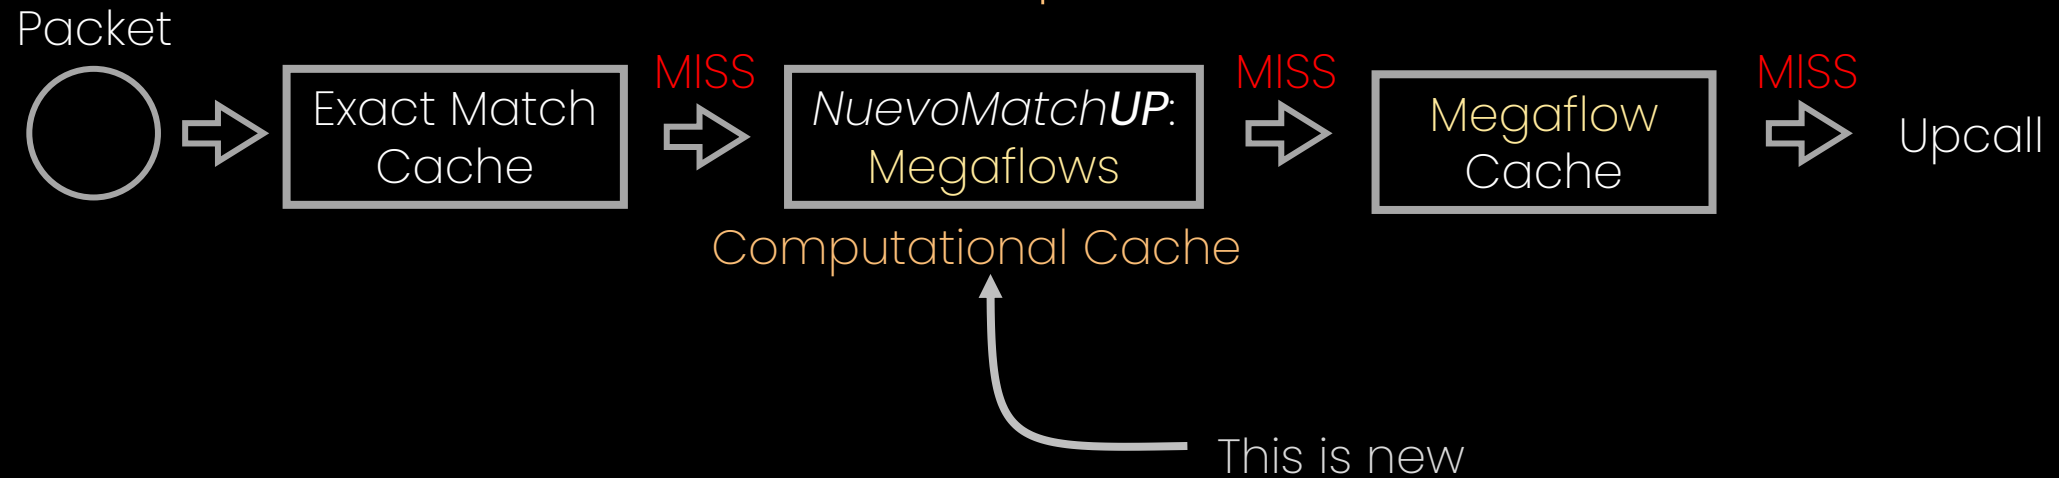
NuevoMatchUP →  **Open vSwitch**

Two Design Options
for OVS data-path

NuevoMatchUP → Open vSwitch

Two Design Options


(1) OVS with a Computational Cache



NuevoMatchUP → Open vSwitch

Two Design Options

(1) OVS with a Computational Cache



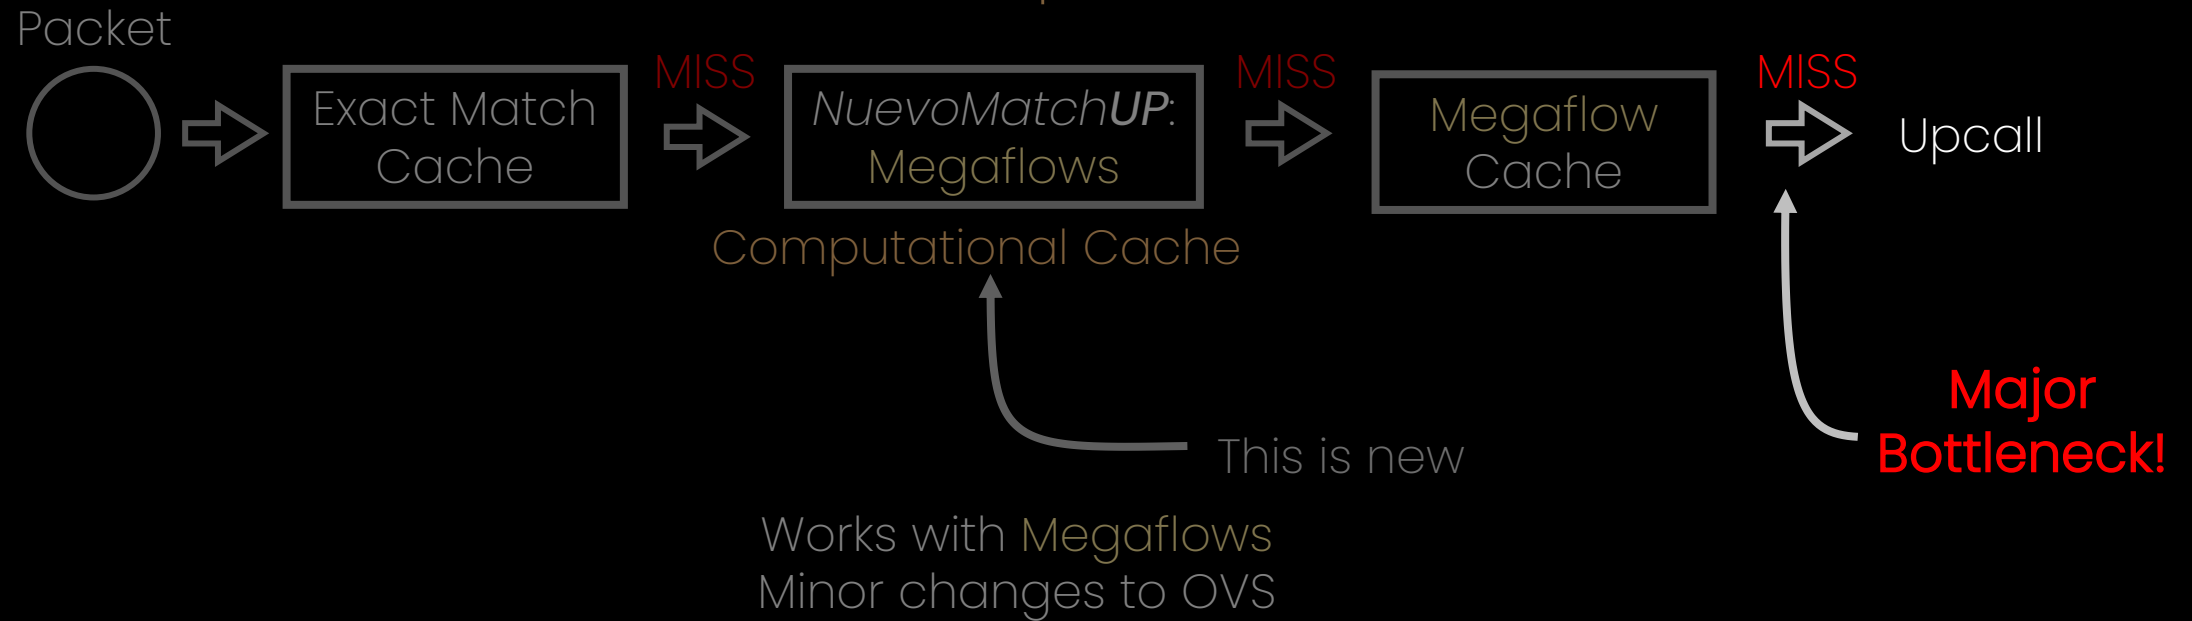
This is new

Works with Megaflows
Minor changes to OVS

NuevoMatchUP → Open vSwitch

Two Design Options


(1) OVS with a Computational Cache



NuevoMatchUP → Open vSwitch

Two Design Options

(1) OVS with a Computational Cache




(2) OVS with Computational Flows



NuevoMatchUP → Open vSwitch

Two Design Options

(1) OVS with a Computational Cache



(2) OVS with Computational Flows



A paradigm shift: data-path matching on OpenFlow rules
No Upcalls!

NuevoMatchUP → Open vSwitch

Two Design Options

(1) OVS with a Computational Cache



We design, implement, and integrate both into OVS DPDK data-path



A paradigm shift: data-path matching on OpenFlow rules
No Upcalls!

Background

Challenge

This paper

Integration with OVS

Evaluation

NuevoMatchUP →  **Open vSwitch**

End to End Results

10Gbps NICs

ClassBench

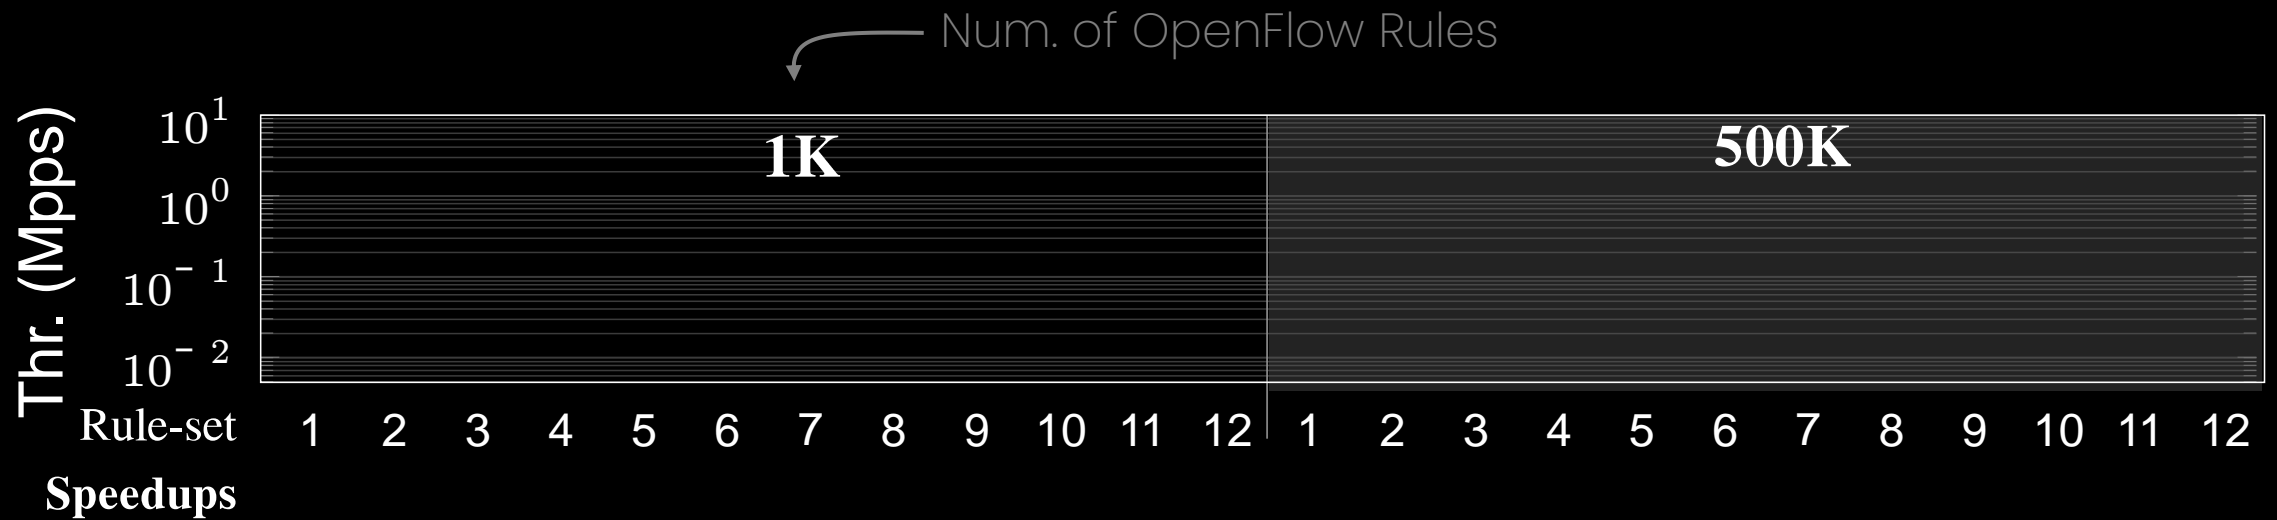
CAIDA

Max 1% Drops

NuevoMatchUP → vs


Open vSwitch

End to End Results




NuevoMatchUP → Open vSwitch

End to End Results




NuevoMatchUP → Open vSwitch

End to End Results




NuevoMatchUP → Open vSwitch

End to End Results




NuevoMatchUP → Open vSwitch

End to End Results



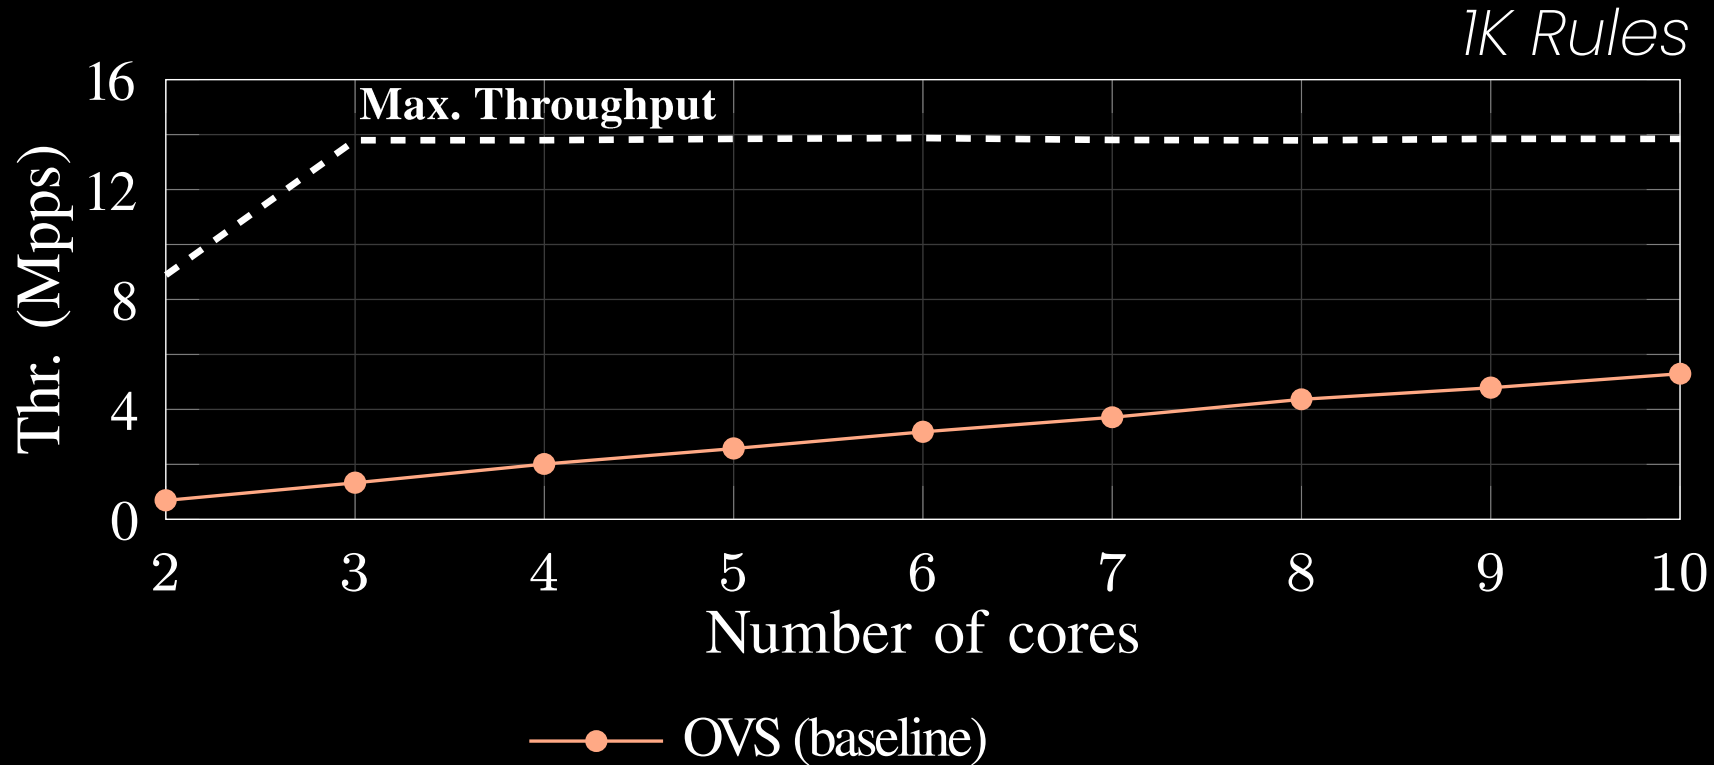
NuevoMatchUP → Open vSwitch

End to End Results



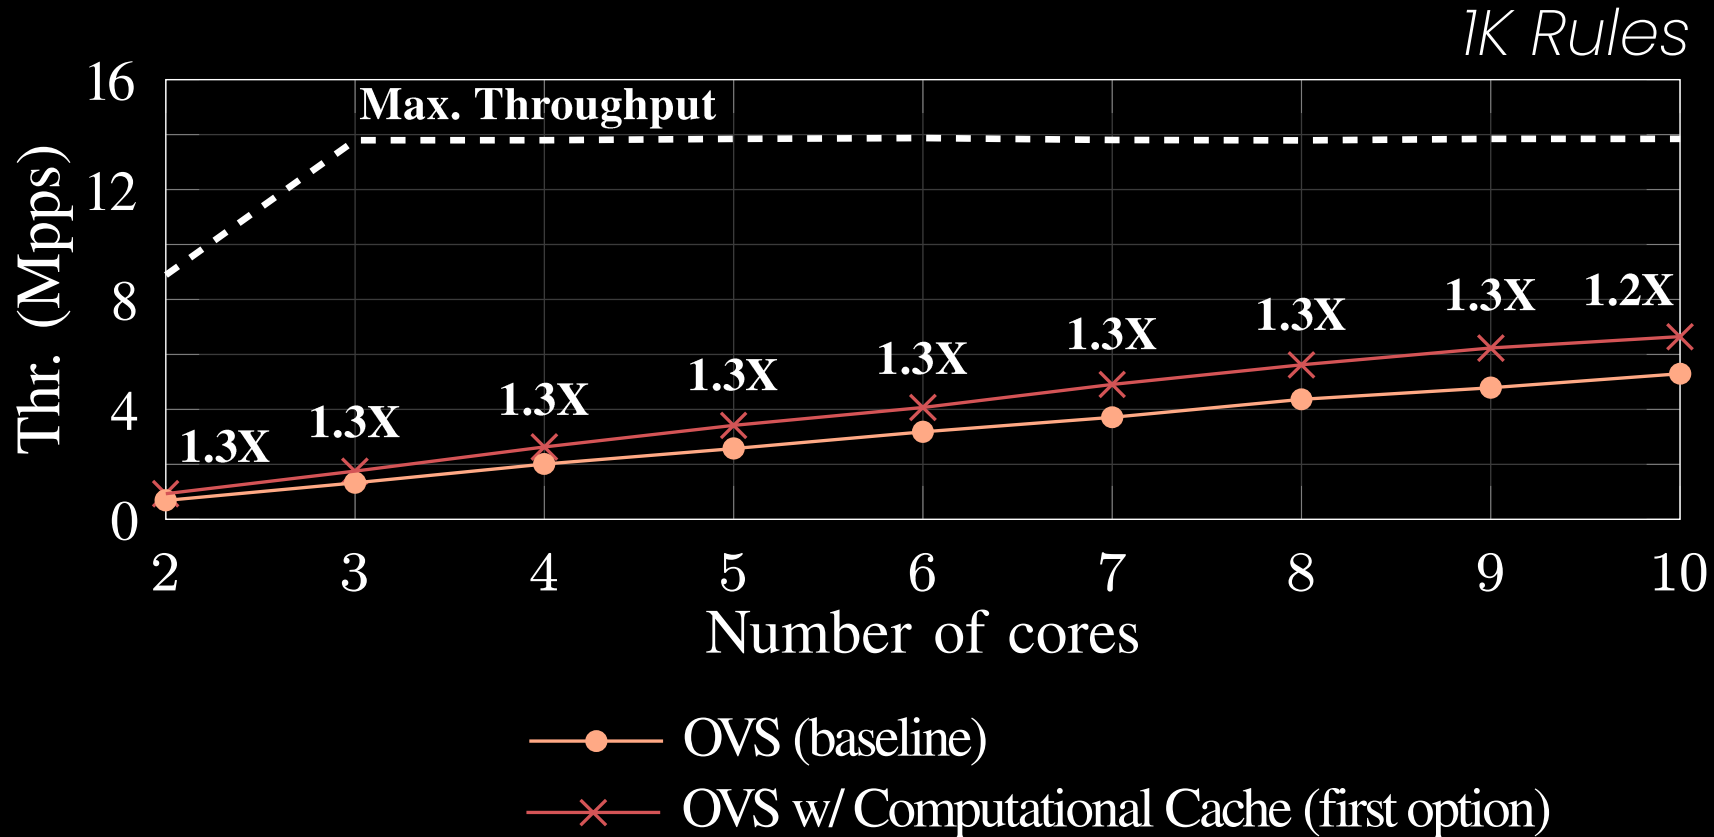
NuevoMatchUP → Open vSwitch

Scalable to #Cores




NuevoMatchUP → Open vSwitch

Scalable to #Cores



NuevoMatchUP → Open vSwitch

Scalable to #Cores



See Paper for

1. Throughput over time experiments
2. Throughput in the presence of updates
3. Microbenchmarks
4. And more...

Conclusions

1. Fast neural-net training for high update rates
2. New design of the OVS data-path
3. Readily deployable to commodity hardware

Future Work

1. NuevoMatchUP for P4
2. Reducing the need for training
3. Hardware acceleration

Future Work

1. NuevoMatchUP for P4
2. Reducing the need for training
3. Hardware acceleration

Thank You



Alon Rashelbach
alonrs@campus.technion.ac.il

Ori Rottenstreich
or@technion.ac.il

Mark Silberstein
mark@ee.technion.ac.il