

ShadowSwitch:

enhancing hardware SDN switches performance with software

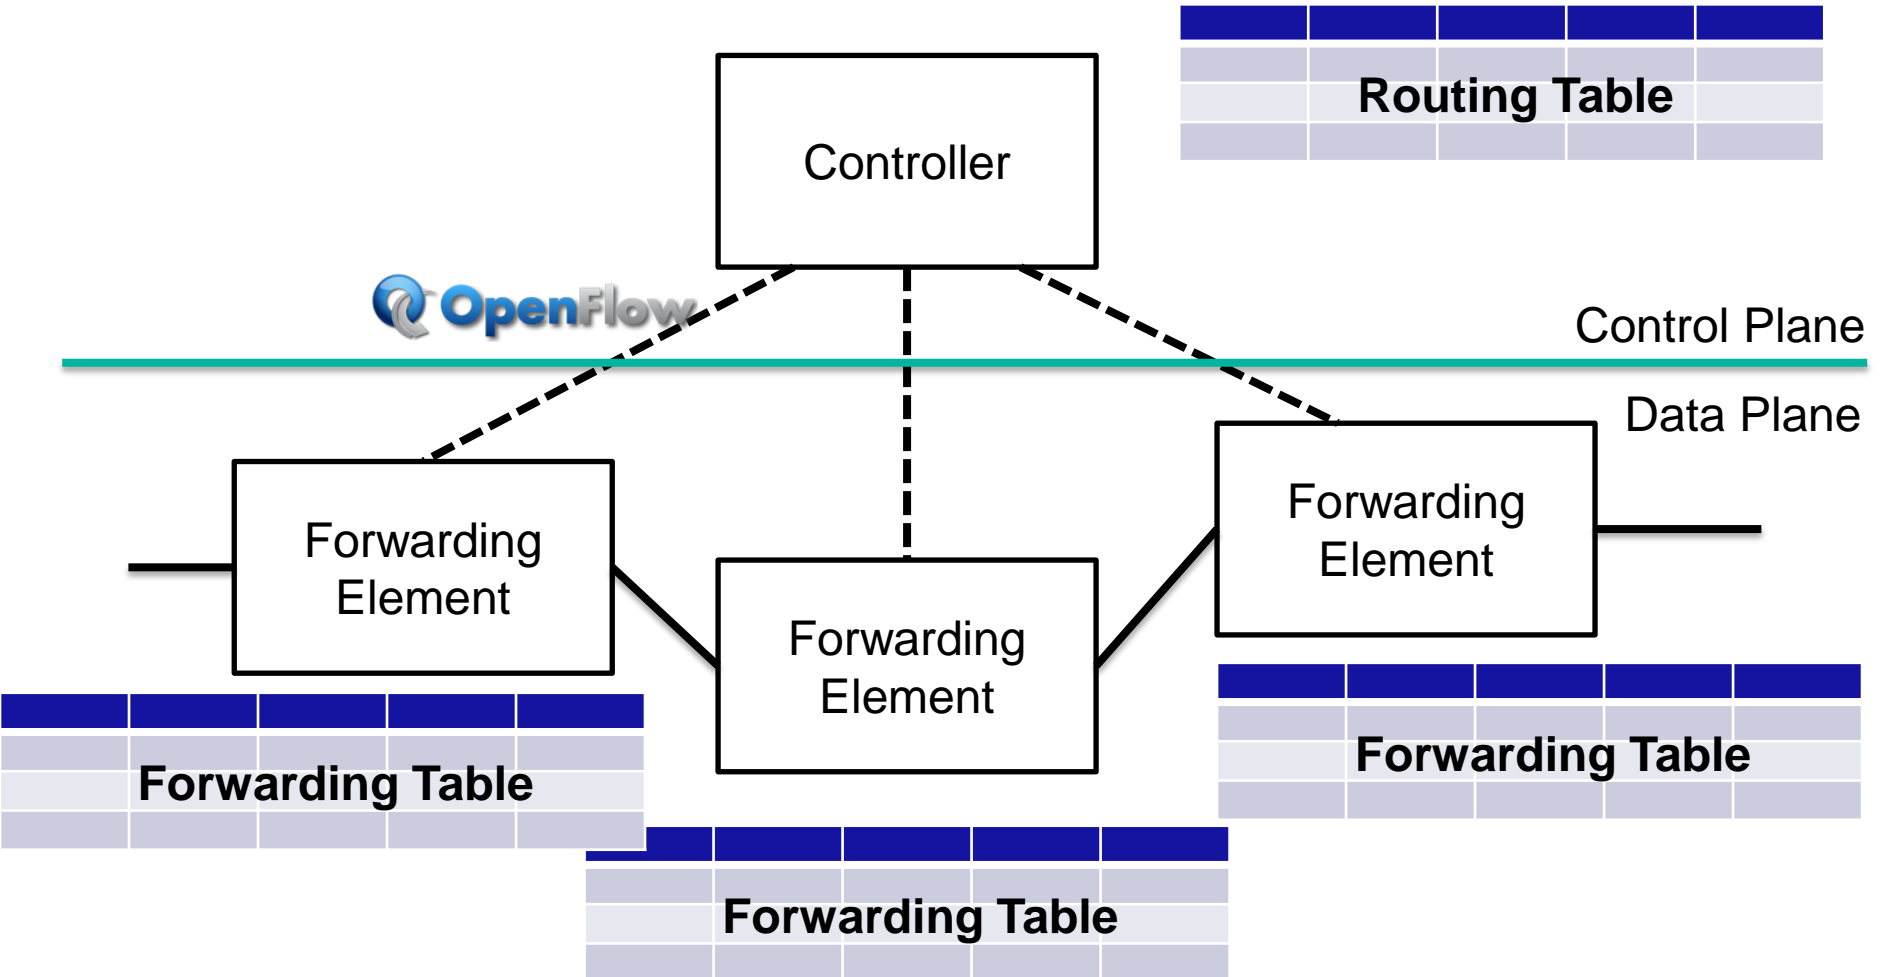
Roberto Bifulco, Anton Matsiuk

Roberto.Bifulco@neclab.eu


NEC Laboratories Europe

Background: SDN and OpenFlow

SDN: the network control plane is separated from the network data plane

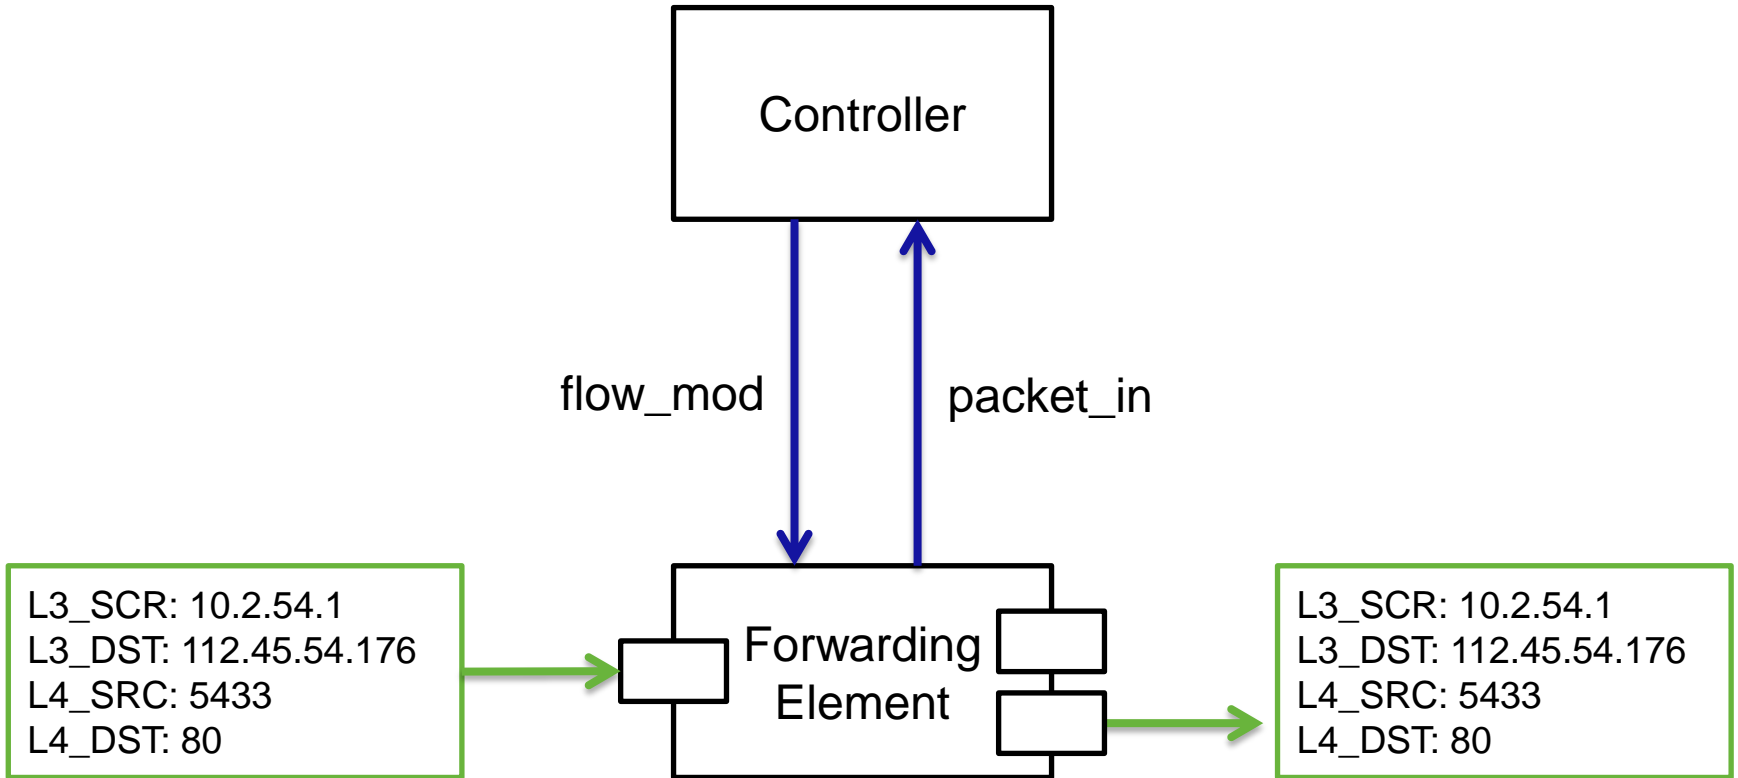


OpenFlow: proactive approach



...	L3_SRC	L3_DST	L4_SRC	L4_DST	...	Action
	Any	112/8	Any	Any		Fwd-to: 2

OpenFlow: reactive approach



...	L3_SRC	L3_DST	L4_SRC	L4_DST	...	Action
	Any	112/8	Any	Any		Fwd-to: 2

Hardware switch limitations

Number of **flow table entries**

- Current limit: **10k** Required: **100k-1M**

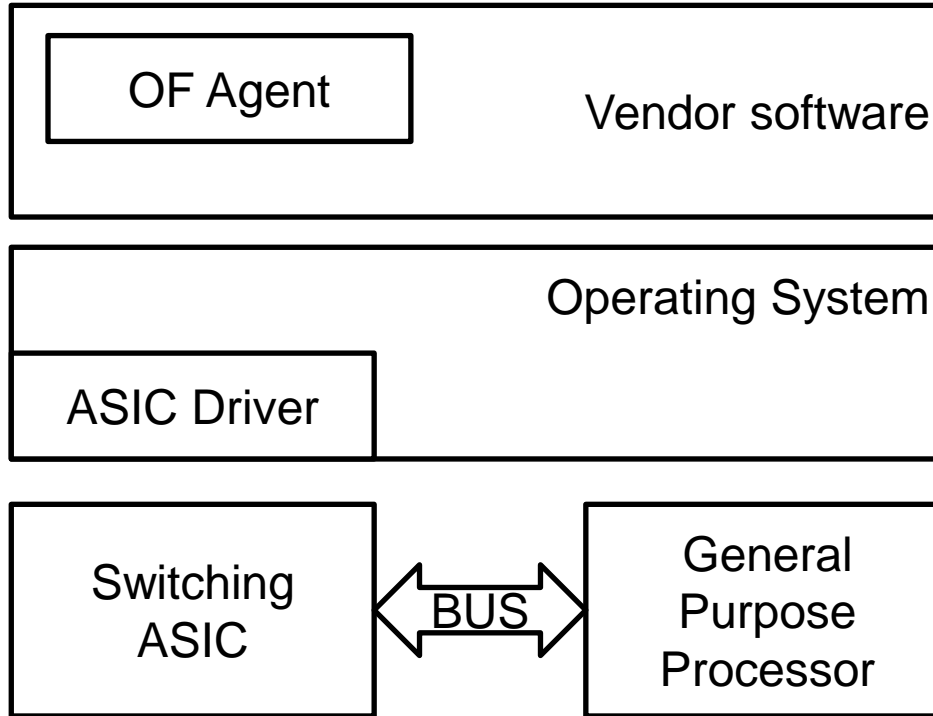
Flow table **update speed**

- Current limit: **100 entries/s** Required: **>1k entries/s**


Switch model	flow_mod/sec
HP ProCurve 5406zl	275
HP ProCurve J9451A	40
Fulcrum Monaco Ref.	42
Quanta LB4G	38
OpenVSwitch	408
NoviSwitch 1248	1000
NoviSwitch 1132	500

These numbers are not capturing the actual behavior of switches

Switch overview



Switching Chip



CAM

TCAM

CAM

CAM
SRAM

Table
space

~10k

~2-4k

~100k

~20-30k

Update
procedures


1 w/r op

Multi.
w/r ops

1 w/r op

1 w/r op

ShadowSwitch



Hardware switches have **limited** Flow Table size

- When the Flow Table is full, packets are **dropped**.

ShadowSwitch combines **hardware** and **software**

- **Hardware**: can process packets at high rates
- **Software**: can host many flows

ShadowSwitch is fast because it is **intelligent**

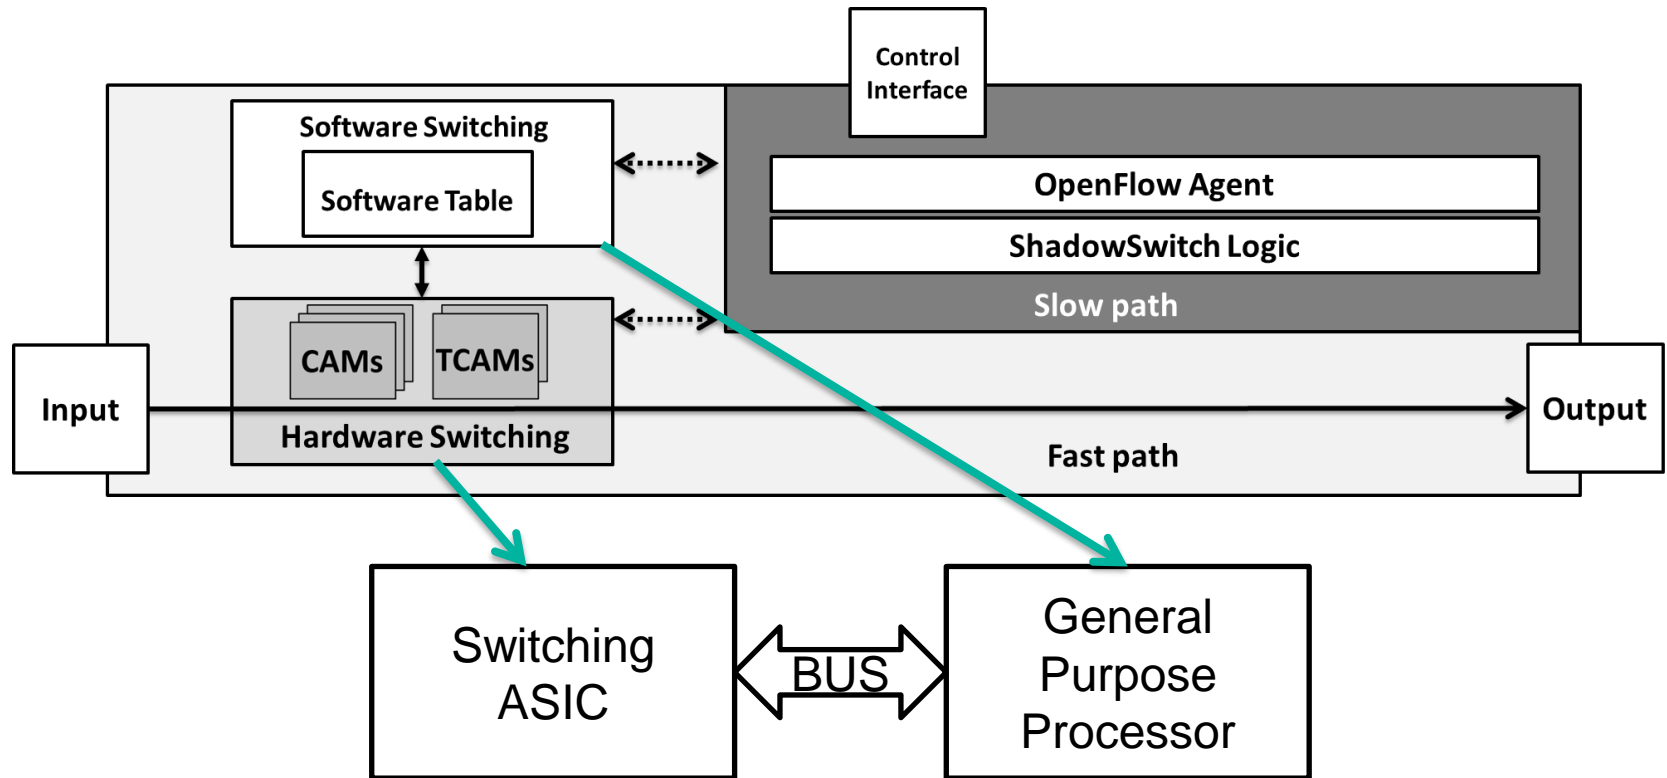
- **Small flows** are handled in **software**
- **Big flows** are handled with the full power of **hardware**
 - **90%** of the flows today are **small flows!**

Hybrid design:
hardware + software

Advanced optimization
algorithms

SDN switch design


- A **software forwarding table** is added to the switch's lookup pipeline
- The hardware and software tables are managed by a **smart logic**



Flow table update speed

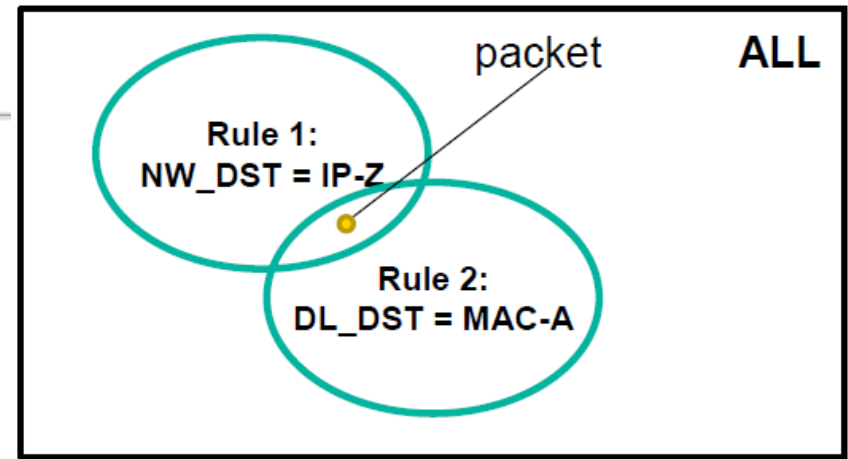
The performance of a switching chip in updating the forwarding table depends on:

- Type of entries
- Type of update
- Number of entries in the table
- Relative order of entries

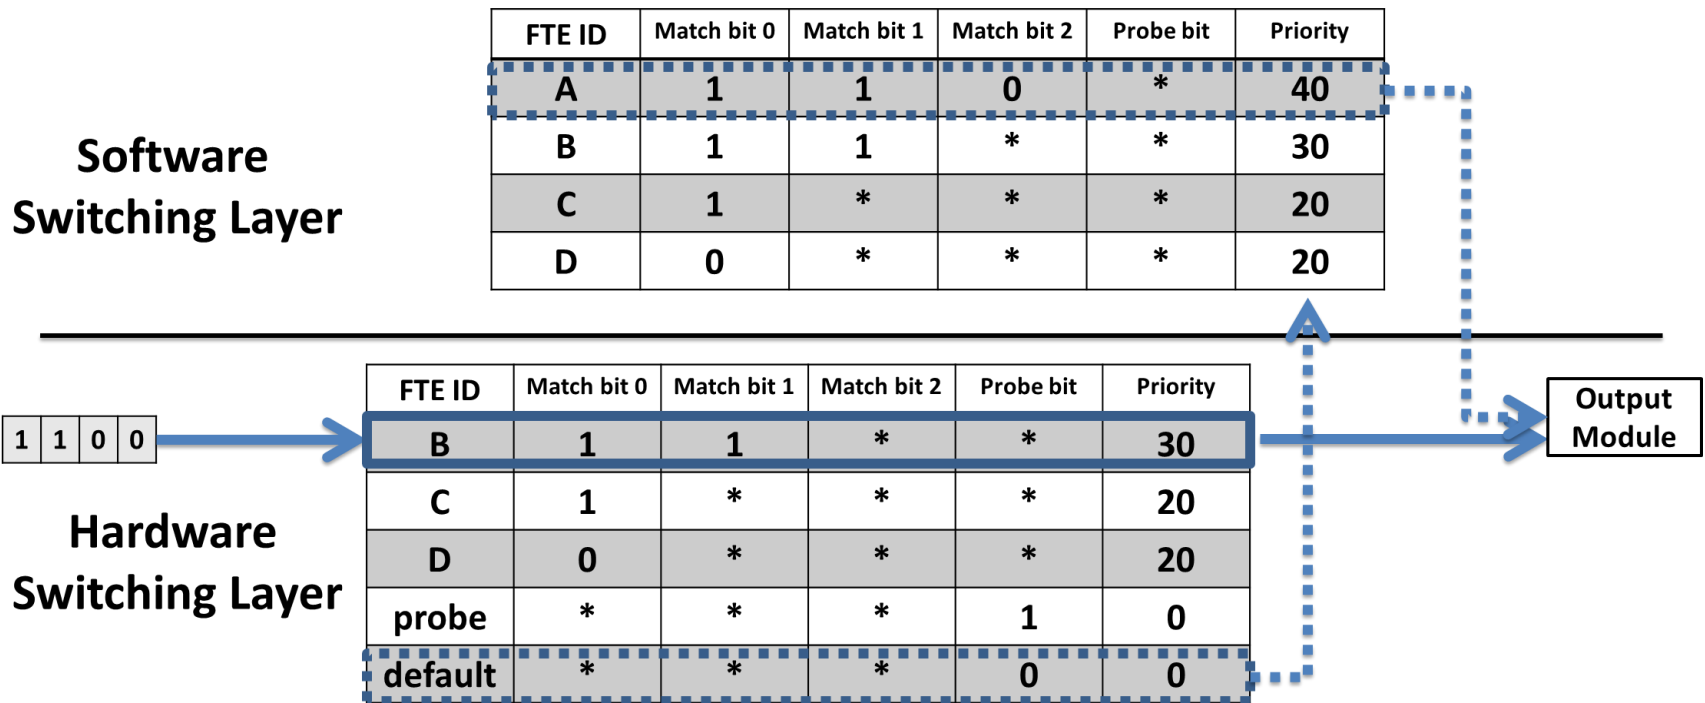


Entries dependencies issue

- Switch behavior depends on the entries relative **priorities**
- Entries in the hardware tables may **hide** the ones in the software table



2D representation of “Flowspace”
(not to scale)




Dealing with the dependencies

Observations:

- The entries in the software table are always in the **correct order**
- Deleting a entry from hardware tables is (**always**) relatively fast

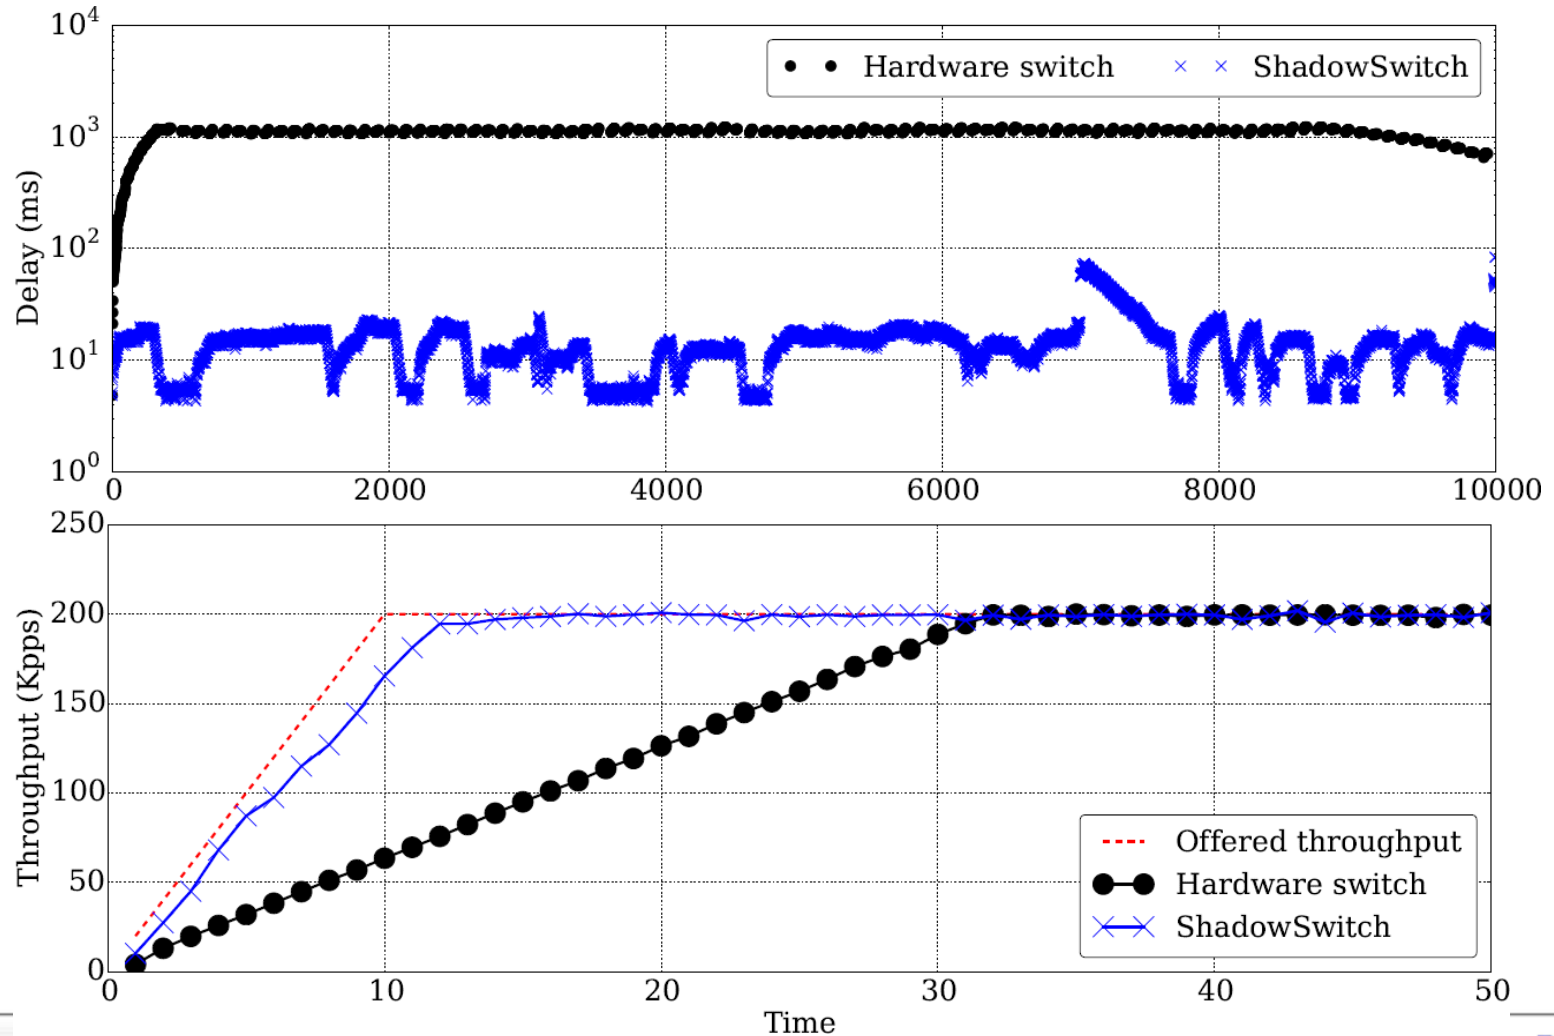
When a new entry is installed/modified:

- Installation happens in the software table
- Dependent entries, with lower priorities, are **deleted** from the hardware tables



Black-box evaluation

10k flows generated @ 1000 flows/s. Total aggr. throughput 200kpps (1 Gbit Ethernet port @ line rate)



Summary

■ New network applications use devices in different ways: new requirements

- increased forwarding table size
- increased number of updates

■ Combining different technologies may help in overcoming stringent hardware constraints

- Software could be **faster** than hardware!

■ ShadowSwitch improves on commercial hardware switches:

- Flow installation delay reduced by 100x
- Flow installation time independent from the number of already installed flows

Acknowledgment: this work has been partially supported by the **BEBA** EU project



Orchestrating a brighter world

NEC brings together and integrates technology and expertise to create the ICT-enabled society of tomorrow.

We collaborate closely with partners and customers around the world, orchestrating each project to ensure all its parts are fine-tuned to local needs.

Every day, our innovative solutions for society contribute to greater safety, security, efficiency and equality, and enable people to live brighter lives.



Empowered by Innovation

NEC