Measurements and Other Issues in Software Defined Networks

Myungjin Lee
University of Edinburgh

Contents

- My Research
- Introduction to Software Defined Network
- Research issues in SDN

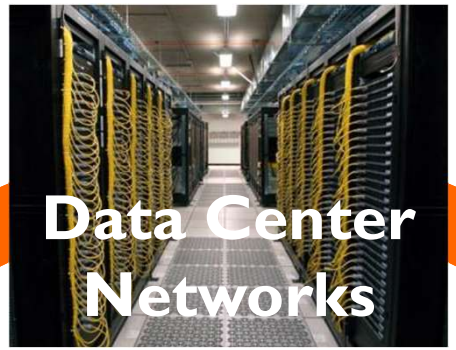

High-fidelity
latency measurements

My Research

Continuous monitoring of
networked applications

Resource provisioning for
cloud applications

Latency! Who Cares?

Applications with Partition/Aggregate workloads

Google facebook bing amazon YAHOO!

Real-time online bidding/trading systems

AdSense

AdWords

Partition/Aggregate Application Structure

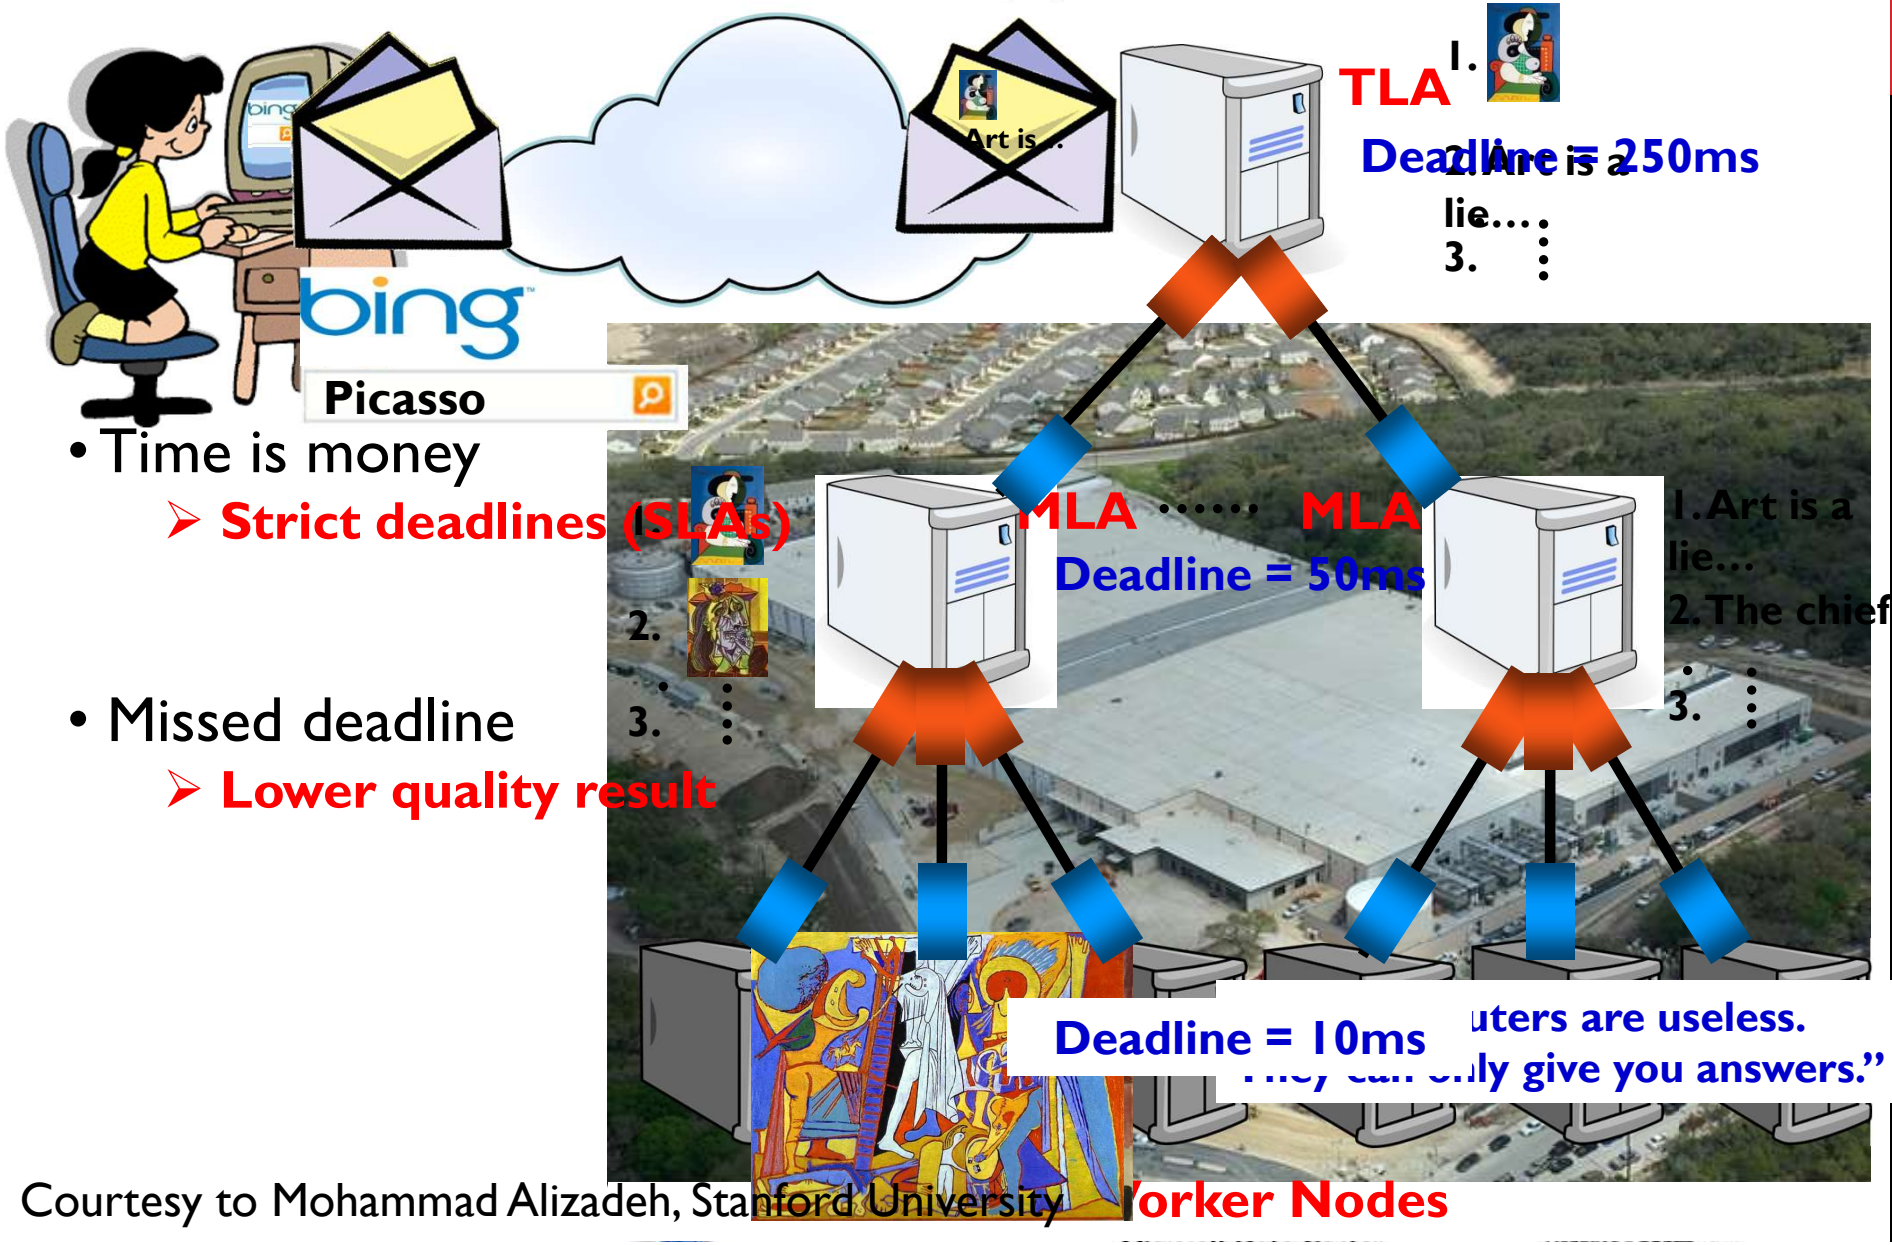

- Time is money
  - Strict deadlines (SLAs)
- Missed deadline
  - Lower quality result

Courtesy to Mohammad Alizadeh, Stanford University

How to manage low end-to-end network latencies?

"When considering how to reduce latency, the first step is to measure it."

--Joanne Kinsella
Head of Portfolio, British Telecom

Measurement Metrics

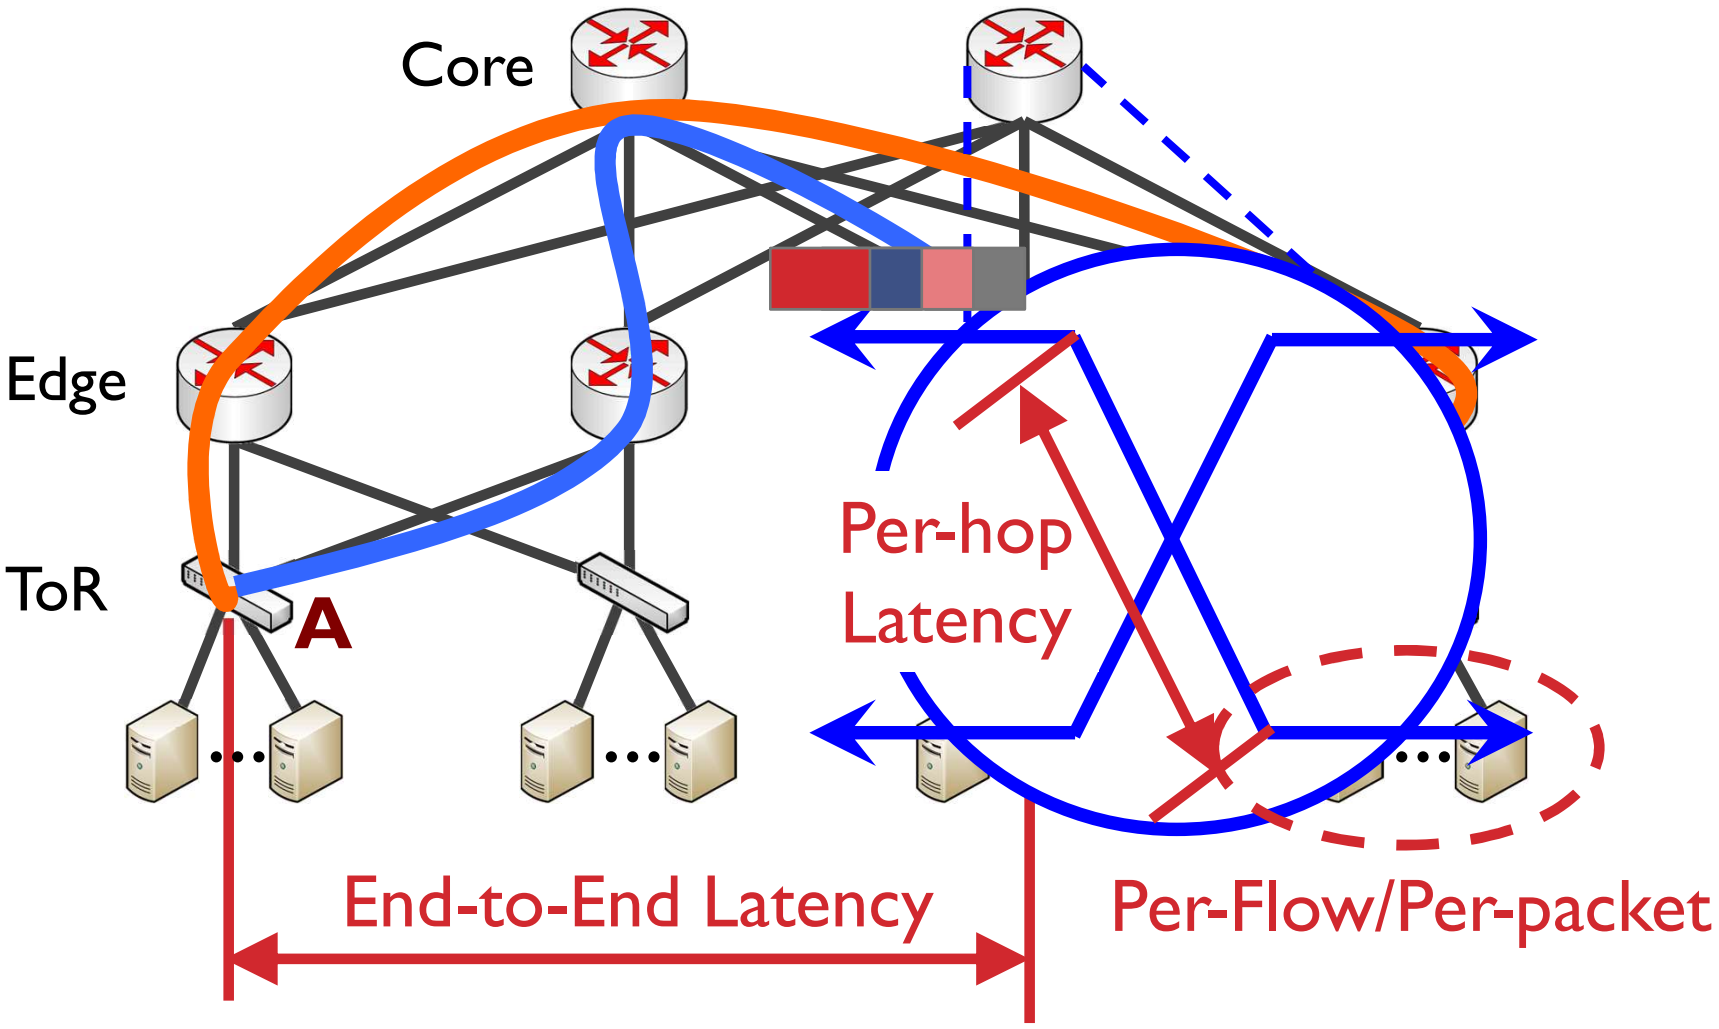

Assumption: All the nodes are time-synchronised

Key Challenge: Too Many Packets

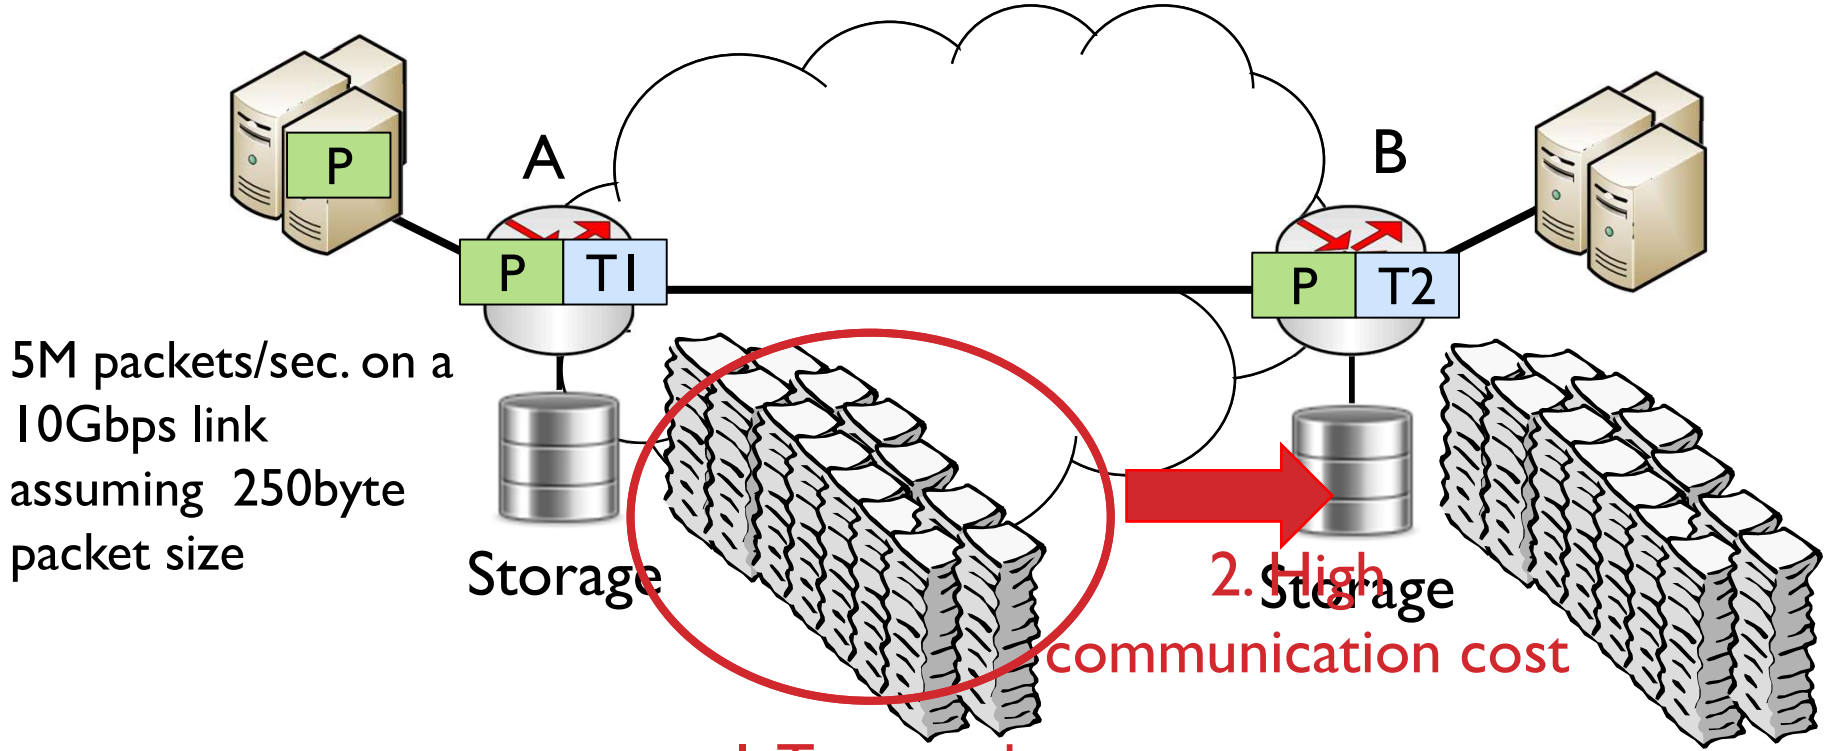

A Brief Summary of My Research

- RLI [SIGCOMM'10]
  - *Per-flow* latency measurements on a *per-hop* basis
  - Low-cost *continuous* monitoring
- FineComb [SIGMETRICS'11]
  - *Aggregate end-to-end* latency measurements
  - *Robust* to packet reordering
- MAPLE [IMC'12]
  - *Efficiently* store *per-packet* latencies with *fast* insert and lookup
  - Statistics of *any* sub-population

New Breeze in Networking: Software Defined Network

Courtesy to Nick Mckeown, Stanford University. A few slides re-use his.

Software Defined Network (SDN)

- Separate control plane from data plane
  - Centralised control of networks
  - Make networks easy to configure/manage/debug

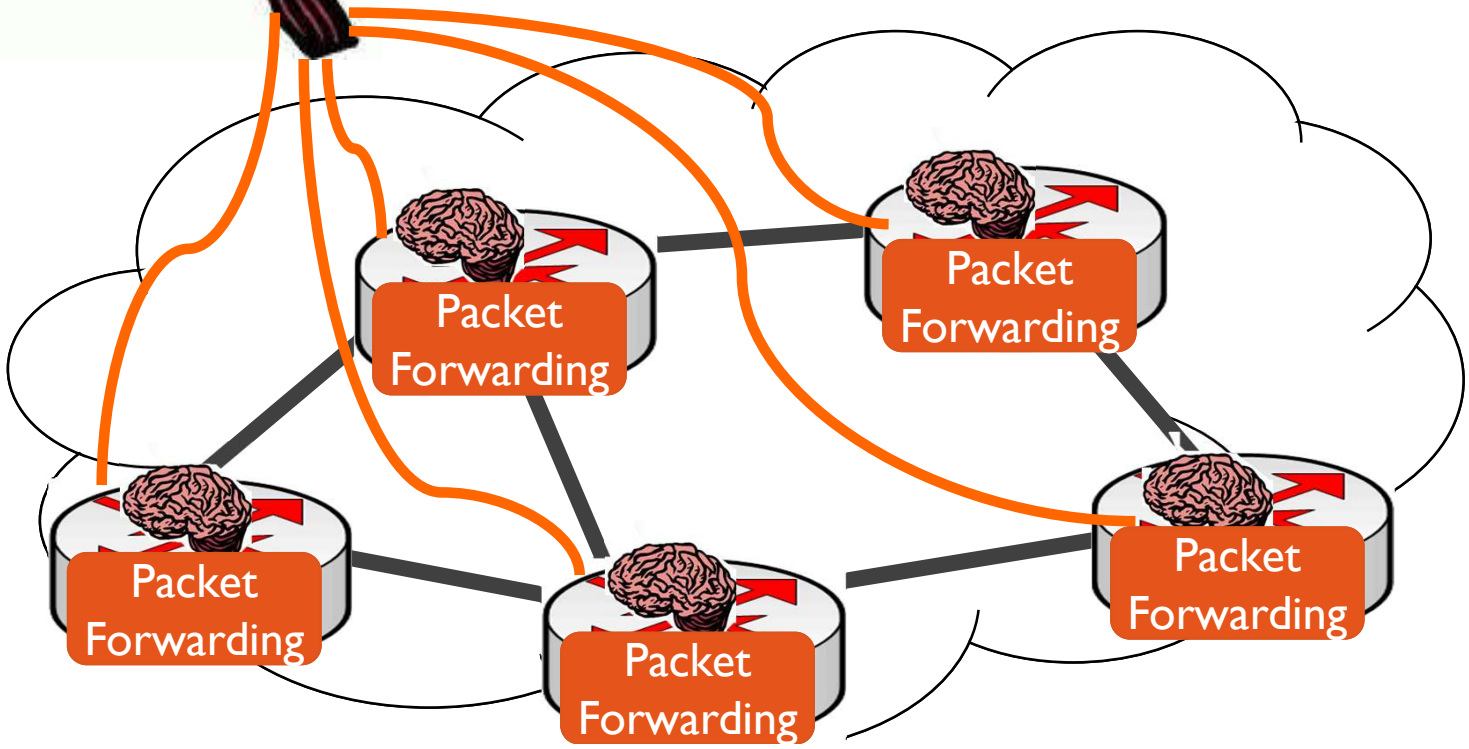

Abstract View of SDN

- Controller(s): NOX, Maestro, …
- Programmable switches: OpenFlow

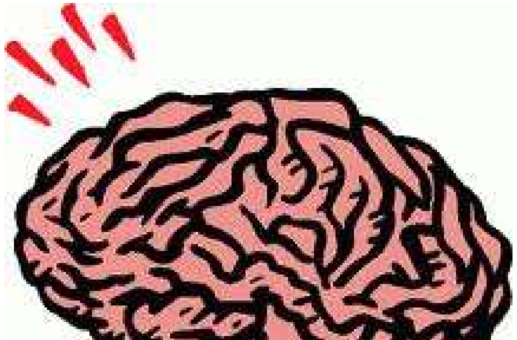

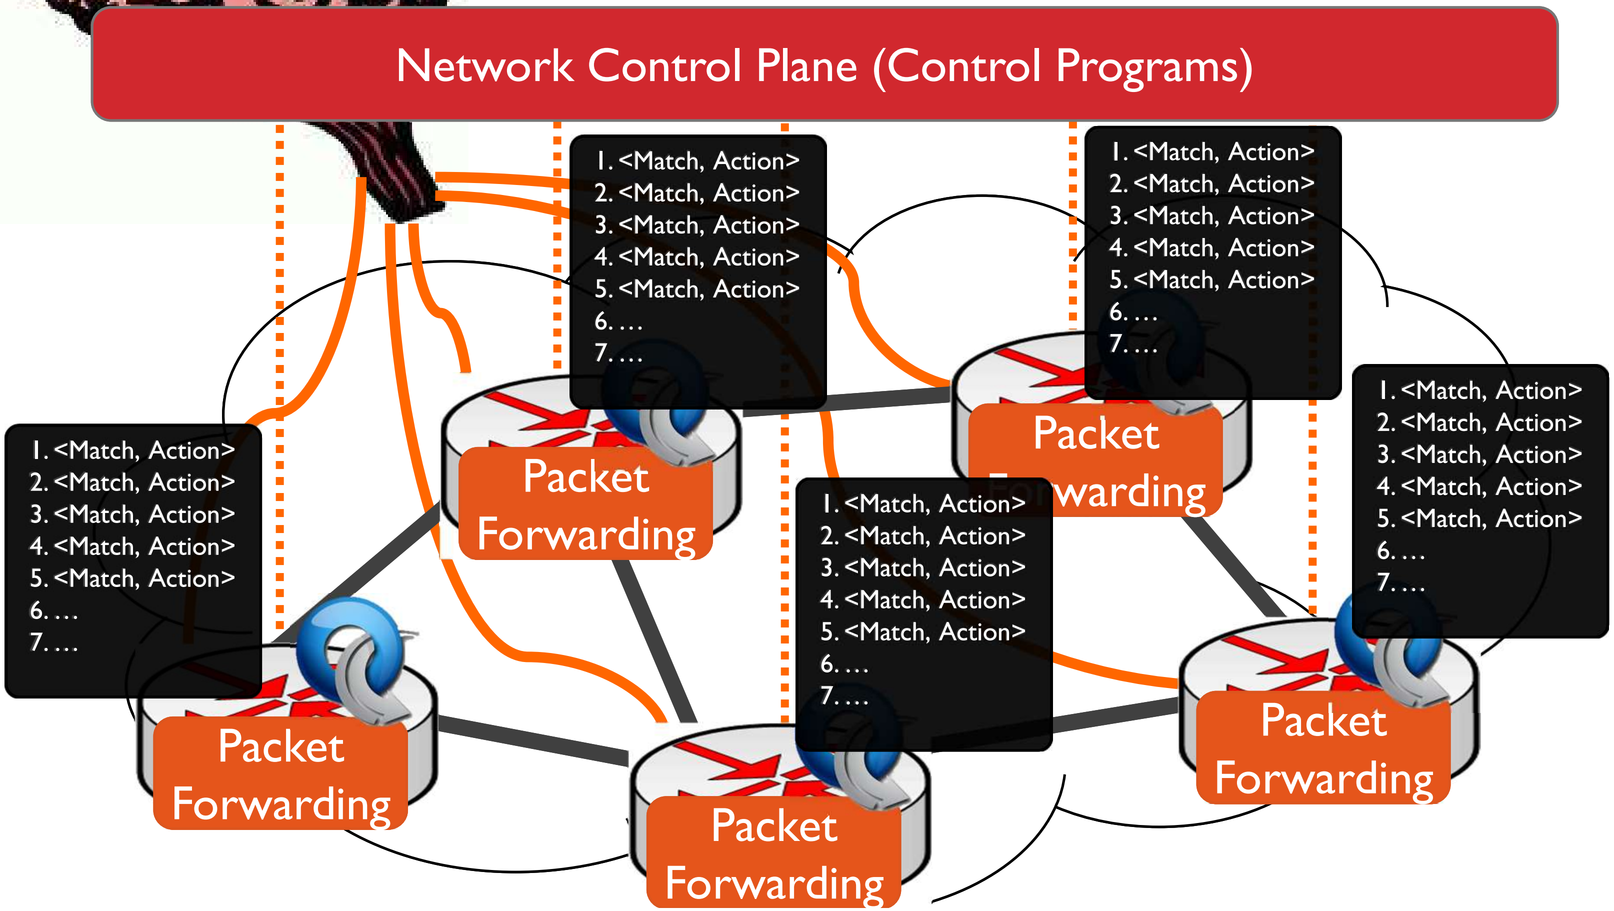

OpenFlow Basics

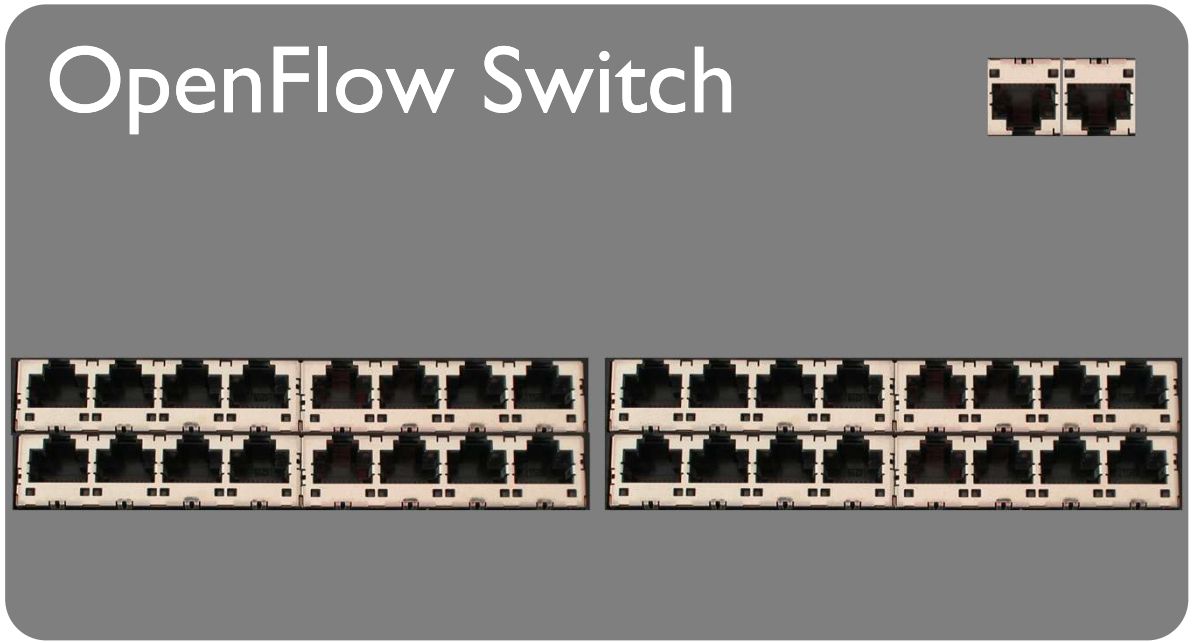

OpenFlow Basics

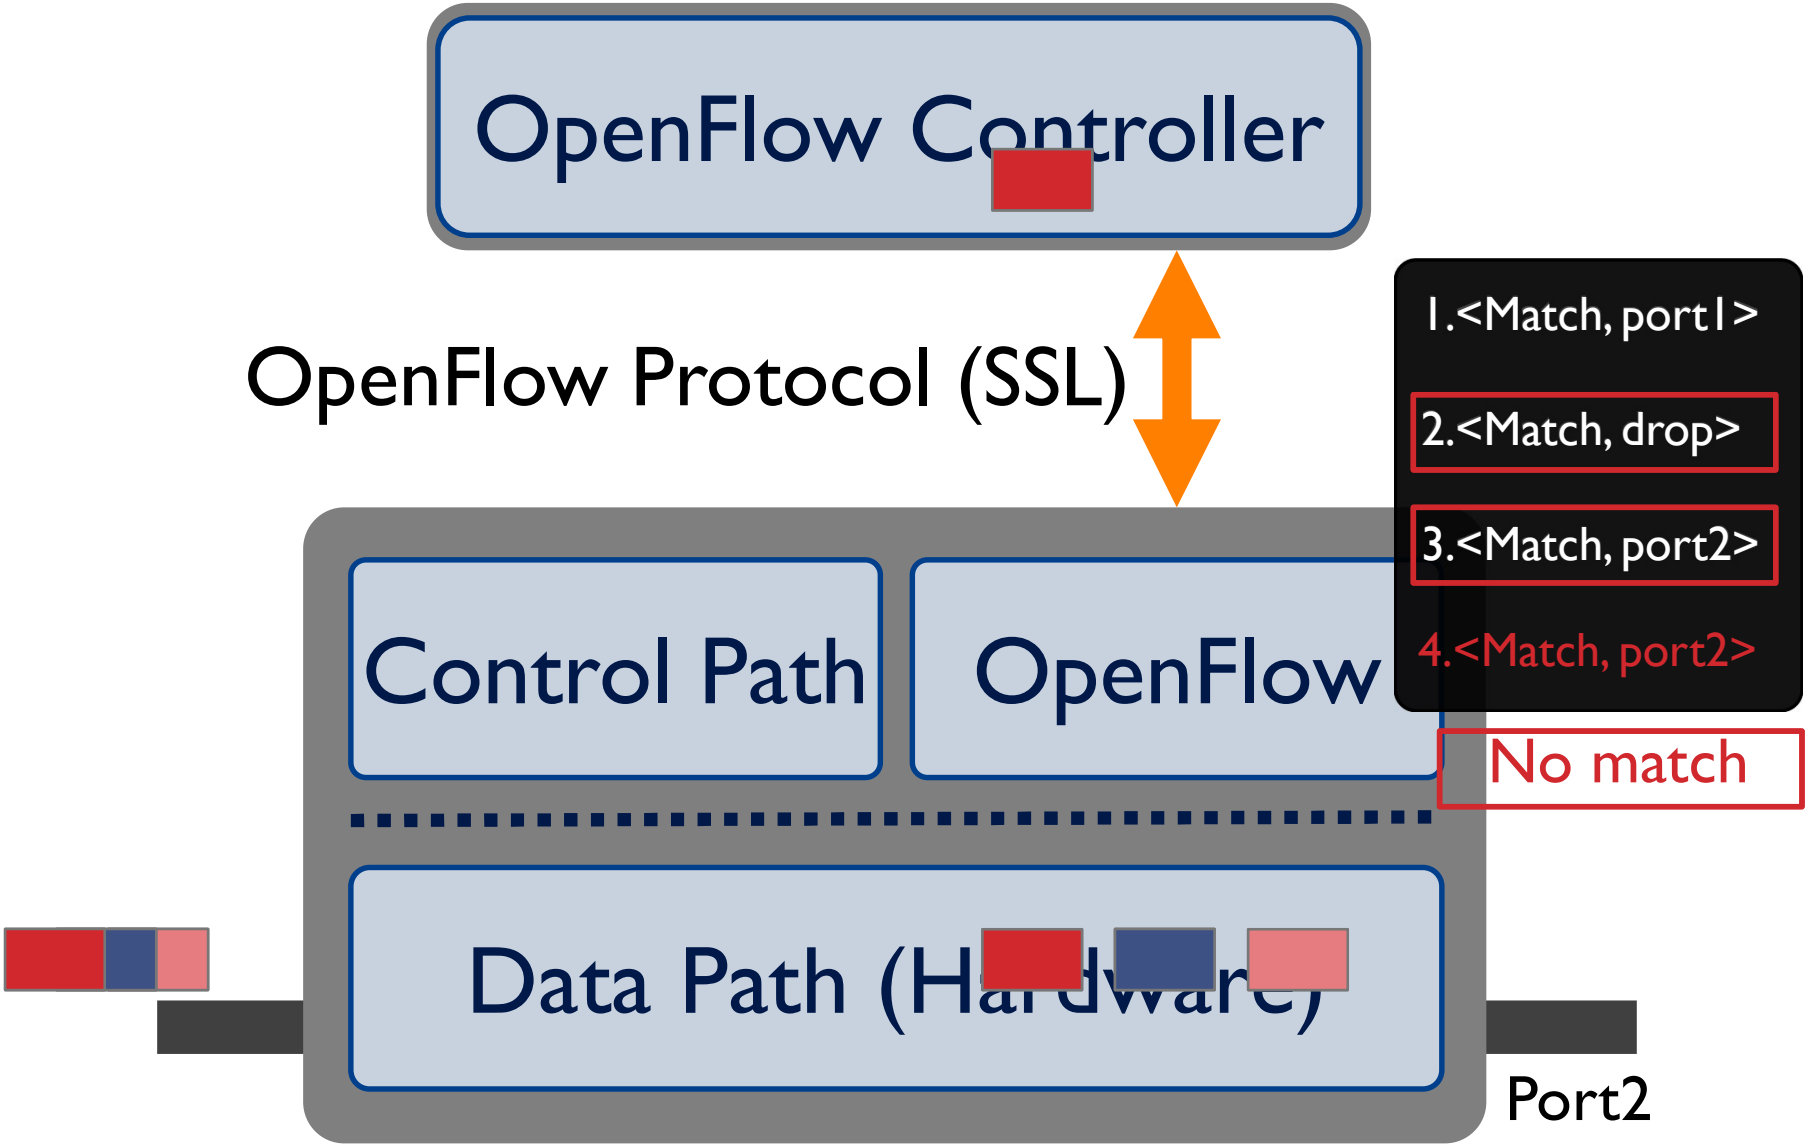

Flow Table Entry

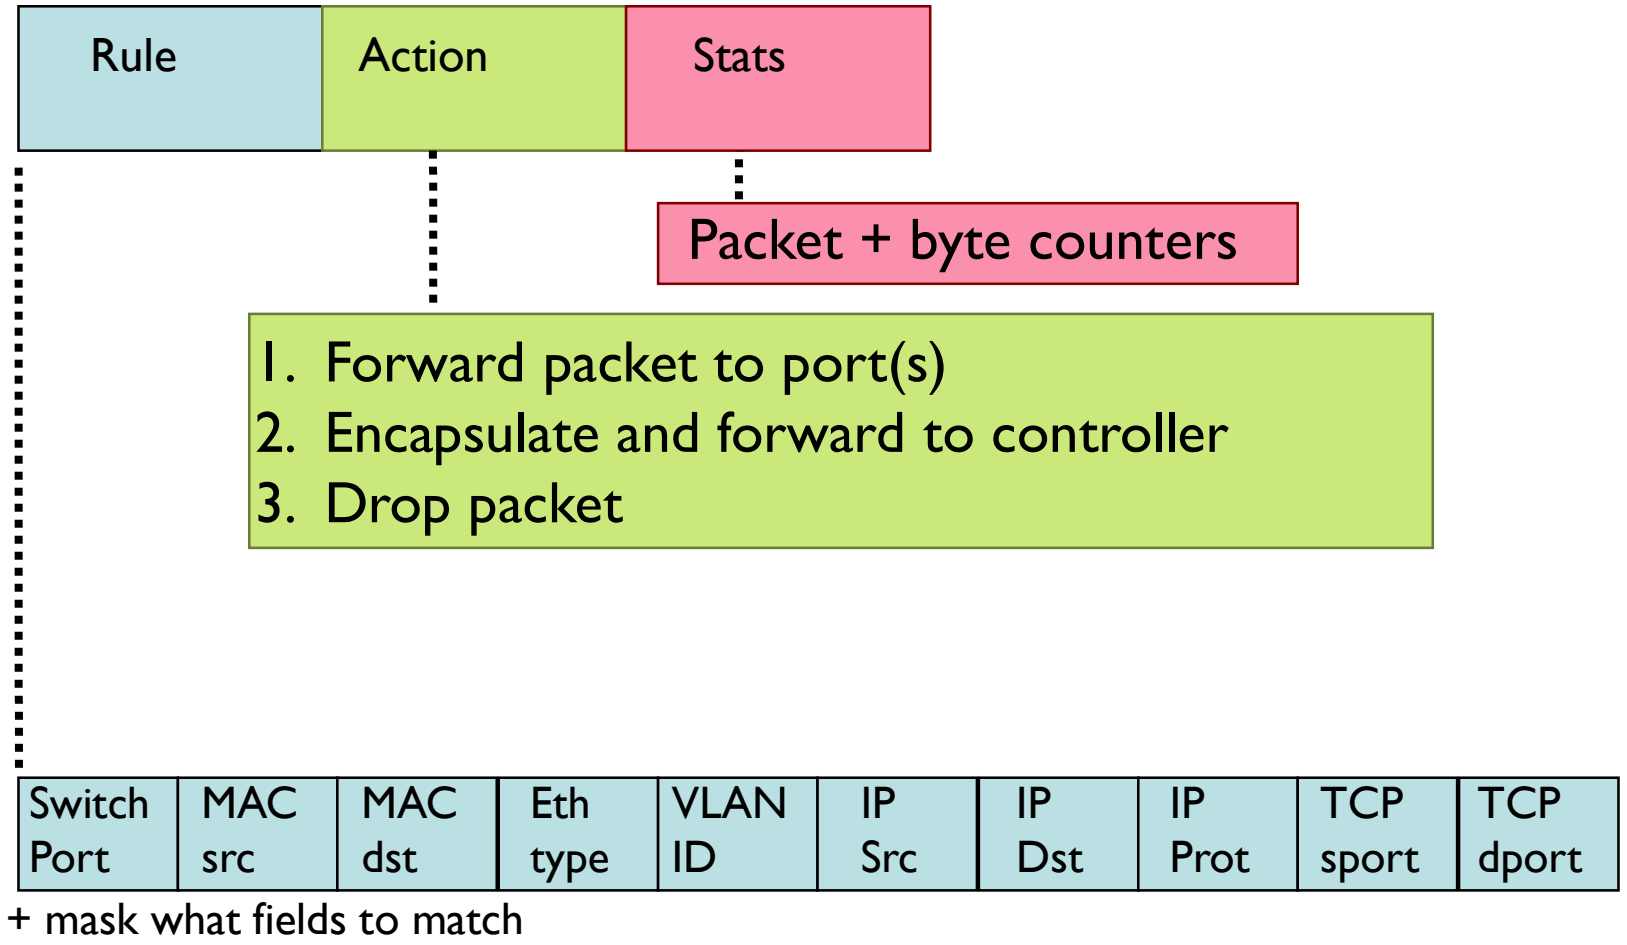

Examples

Switching

Switch Port	MAC src	MAC dst	Eth type	VLAN ID	IP Src	IP Dst	IP Prot	TCP sport	TCP dport	Action
*	*	00:1f:..	*	*	*	*	*	*	*	port6

Flow Switching

Switch Port	MAC src	MAC dst	Eth type	VLAN ID	IP Src	IP Dst	IP Prot	TCP sport	TCP dport	Action
port3	00:2e..	00:1f..	0800	vlan1	1.2.3.4	5.6.7.8	4	17264	80	port6

Firewall

Switch Port	MAC src	MAC dst	Eth type	VLAN ID	IP Src	IP Dst	IP Prot	TCP sport	TCP dport	Action
*	*	*	*	*	*	*	*	*	22	drop

Examples

Routing

Switch Port	MAC src	MAC dst	Eth type	VLAN ID	IP Src	IP Dst	IP Prot	TCP sport	TCP dport	Action
*	*	*	*	*	*	5.6.7.8	*	*	*	port6

VLAN

Switch Port	MAC src	MAC dst	Eth type	VLAN ID	IP Src	IP Dst	IP Prot	TCP sport	TCP dport	Action
*	*	*	*	vlan 1	*	*	*	*	*	port6, port7, port9

Research Issues

- Scalability: Bottleneck at a centralised controller
- Measurements
- Scalability: Limited TCAM resource
- Control Plane Network

Scalability Issue in SDN

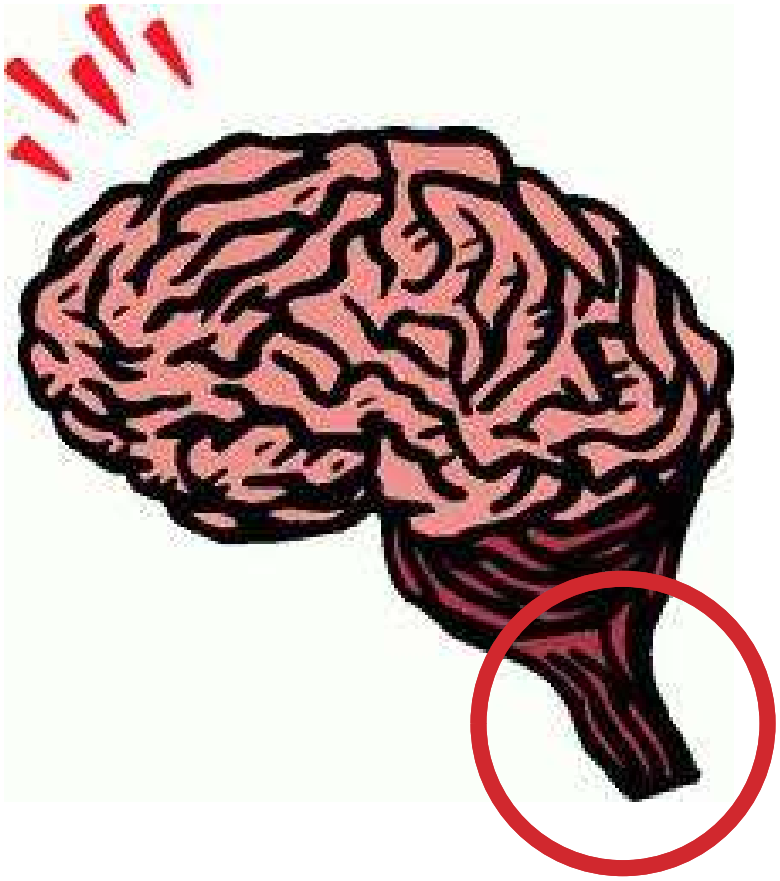

Scalability in SDN

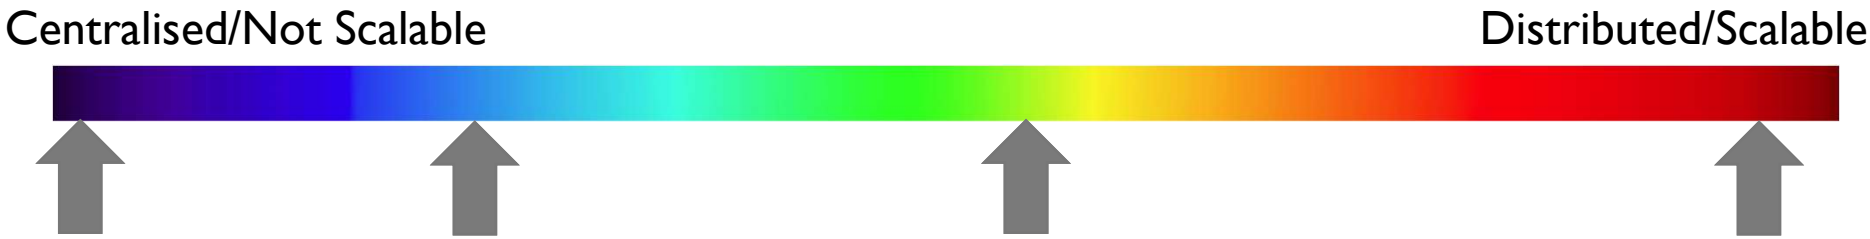

NOX

Distributed Controllers:
HyperFlow, Onyx,
Devolved Controllers

Data Plane Extensions:
DIFANE [SIGCOMM'10]
DevoFlow [SIGCOMM'11]

Hierarchical Controllers:
KANDOO [HotSDN'12]

???

Scalability vs. Visibility

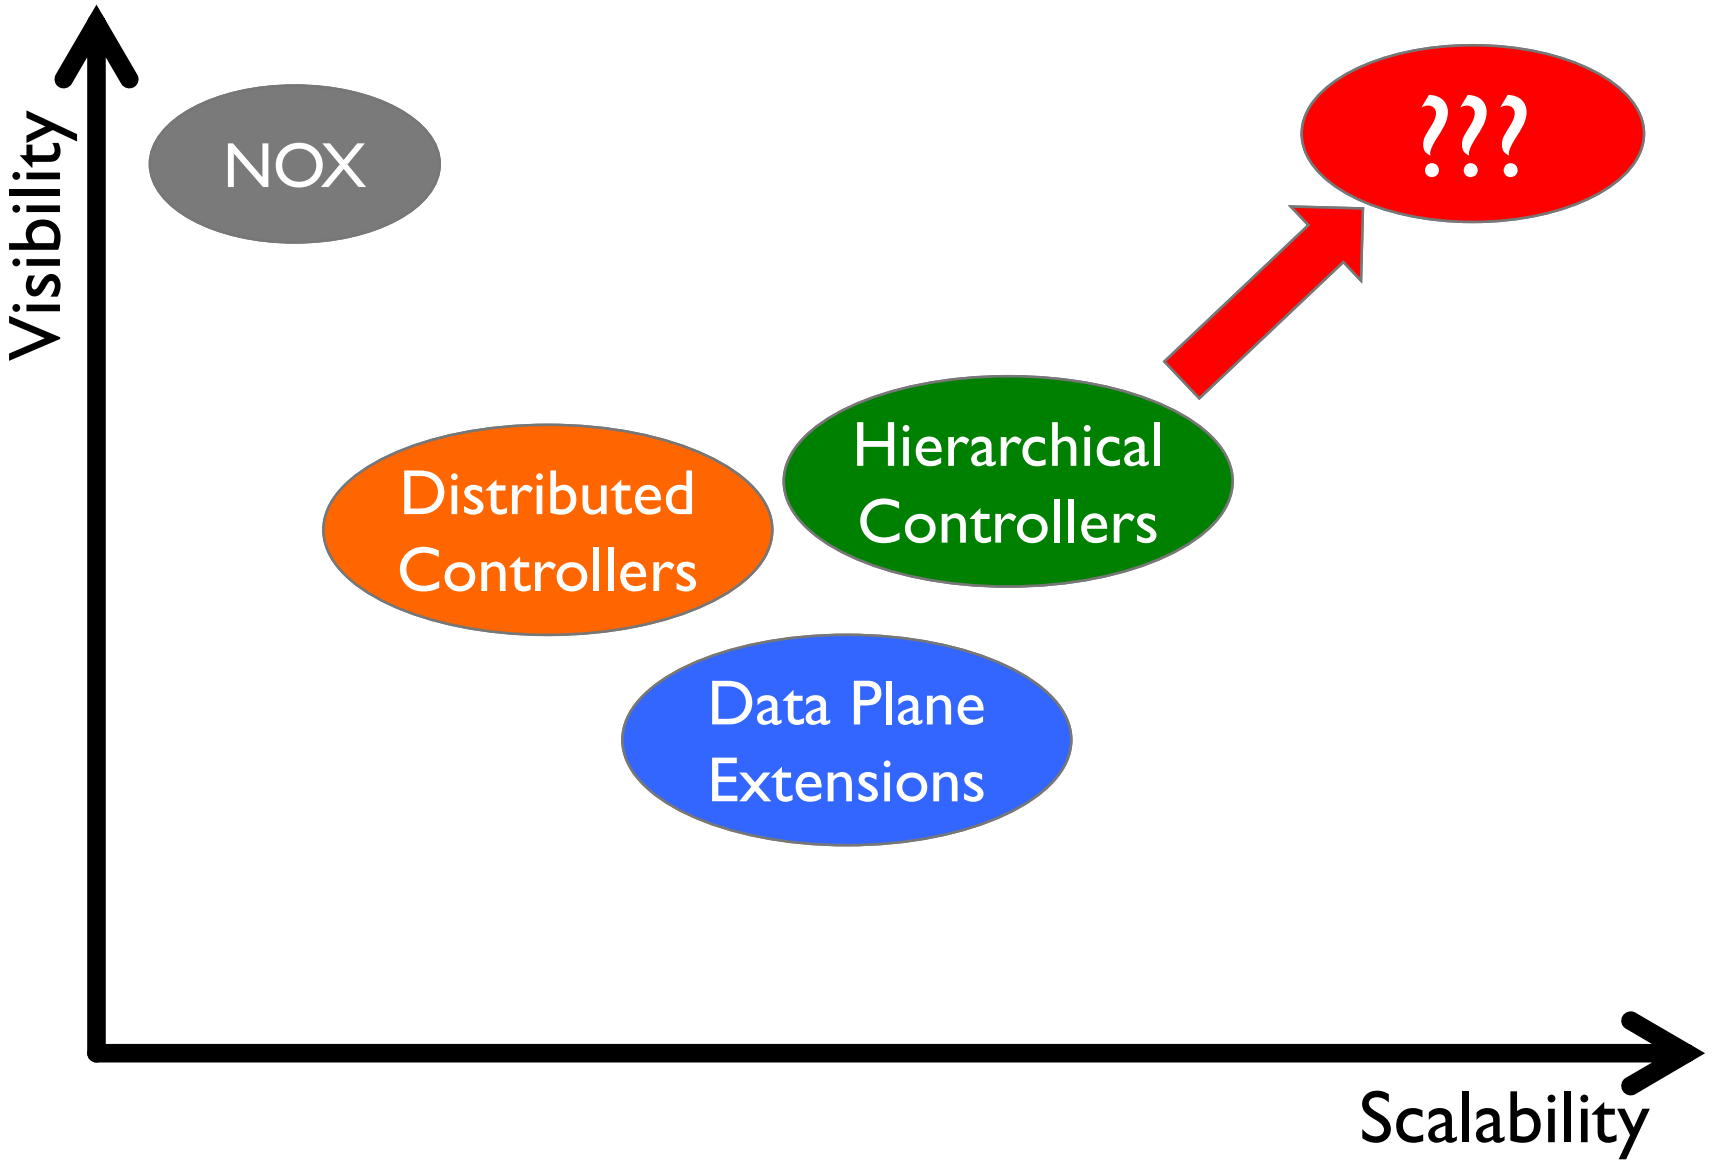

Measurements in SDN

- Basic counters
  - # of packets, bytes on per-flow, per-queue, per-table, etc.
  - # of dropped packets on per-port basis

- But, what about latency or loss measurements?
  - Per-hop, per-flow, end-to-end, …

Limited TCAM Size

- What OpenFlow switch CAN do
  - Switching
  - Routing
  - VLAN
  - Firewall
  - Basic measurements
  - Multicast, Flooding, Multipath, Waypoint
  - …
- CAN OpenFlow switch do all of them? Maybe not
  - NEC PF5820: 750 12-tuple entries, or 80,000 layer2 entries
  - Brocade MLX: 4,000 12-tuple flow entries

Control Plane Network

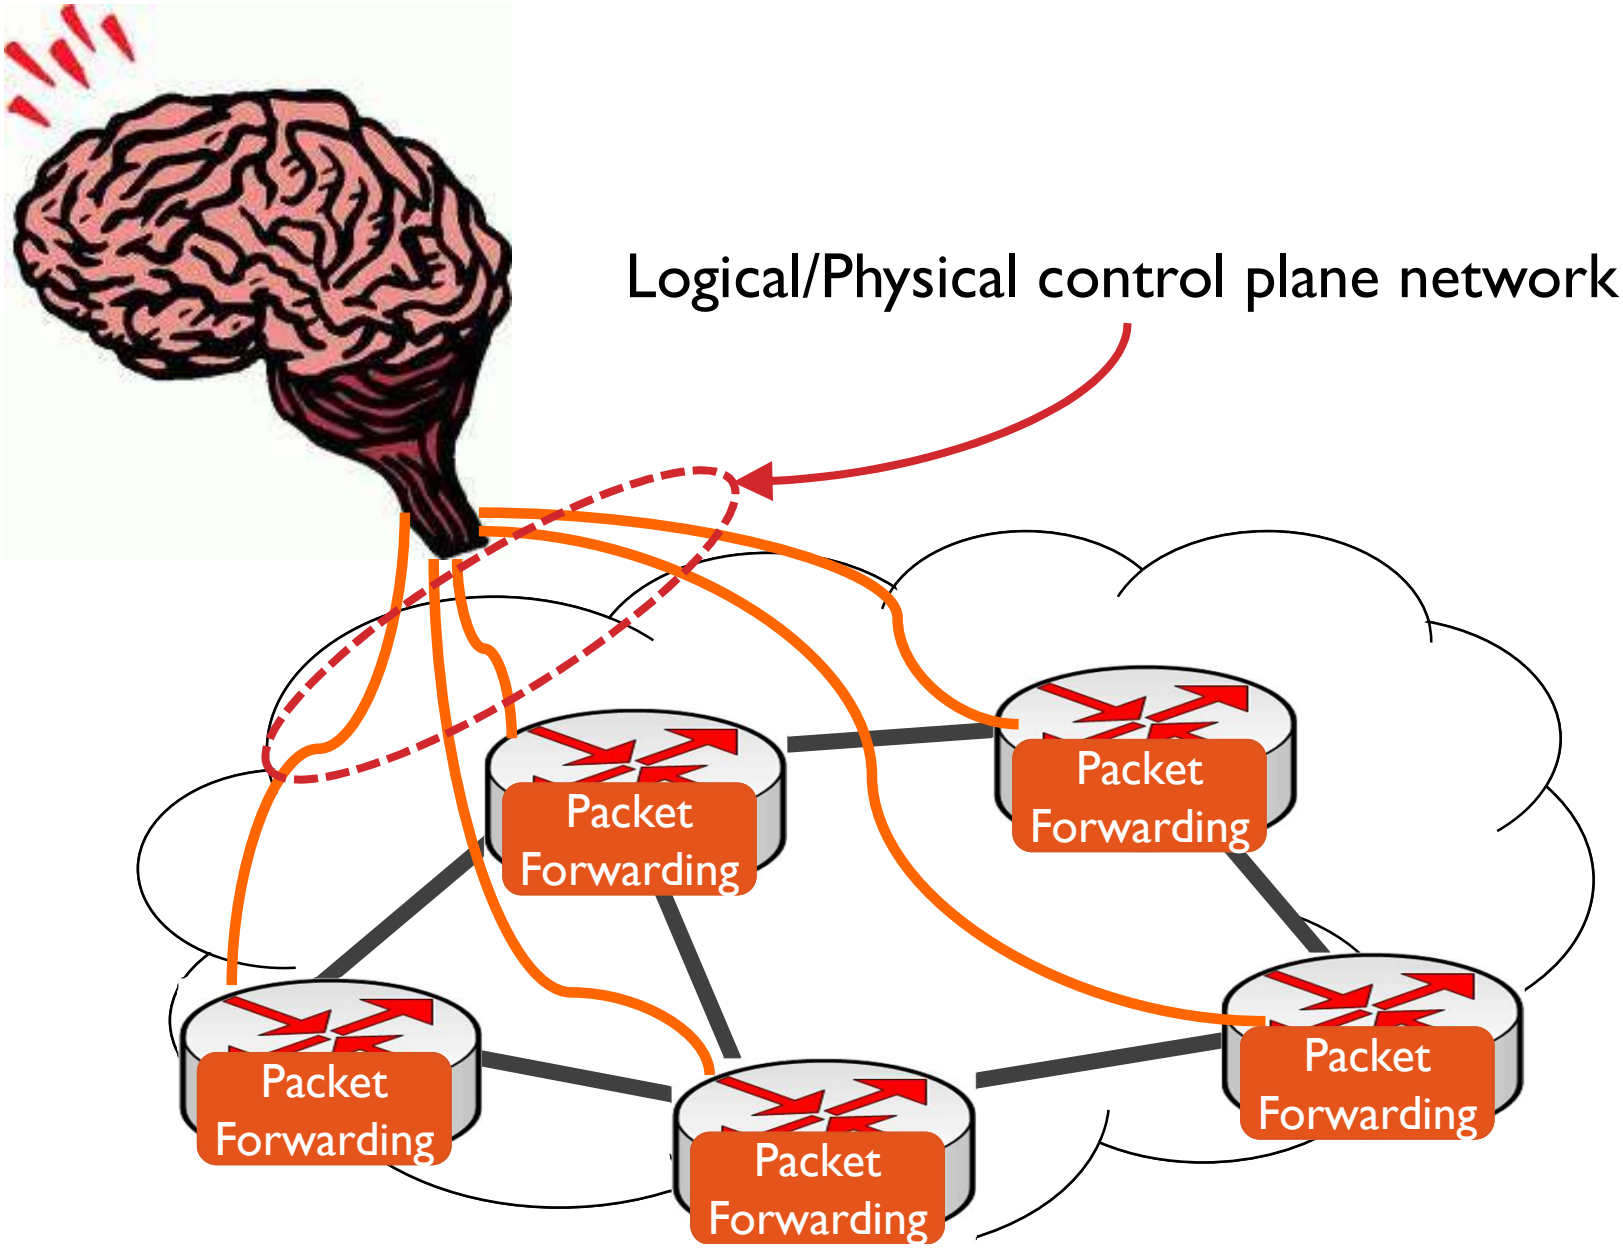

Control Plane Network

- Nobody talks about a *separate* control plane network
  - Sharing a data plane network may be sufficient
- Does this mean there is nothing to do?
- What about the following questions?
  - What is bandwidth/delay requirement for control messages?
  - Can a control plane network be useful to address scalability and visibility issues?
  - What functions do switches in control plane should support?

Summary

- Latency measurements
  - An *important* performance metric in data centre networks
- SDN is becoming a *storm*, not a breeze in networking
- Research *issues* in SDN
  - Scalability issue with a centralised controller
  - Measurements in SDN
  - Scalability issue with limited TCAM size
  - Control plane network

Thank you! Questions?

myungjin.lee@ed.ac.uk