ΕΛΛΗΝΙΚΗ ΔΗΜΟΚΡΑΤΙΑ
ΠΑΝΕΠΙΣΤΗΜΙΟ ΚΡΗΤΗΣ

Δίκτυα Καθοριζόμενα από Λογισμικό

Ενότητα 1.3: OpenFlow

Ξενοφώντας Δημητρόπουλος
Τμήμα Επιστήμης Υπολογιστών

HY436: OpenFlow

Xenofontas Dimitropoulos
6/10/2014

Credits: Nick McKeown's group @ Stanford

Agenda

- SDN vs distributed routing protocols
- OpenFlow:
  - Flow table and messages
  - How it started?
    “Innovate in your campus network” use case
  - Advanced features (v1.4)

Networking industry (2007)

Routing, management, mobility management, access control, VPNs, …

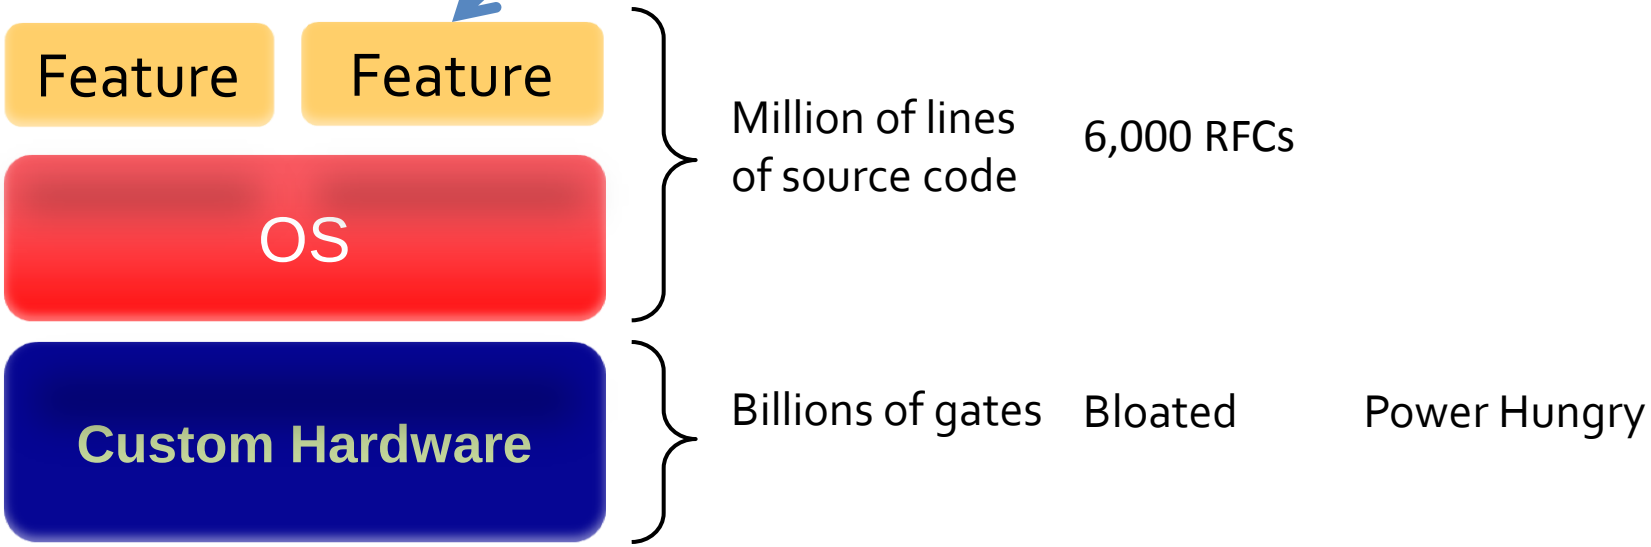

- Vertically integrated, complex, closed, proprietary
- Networking industry with “mainframe” mind-set

Present distributed routing

- Distributed protocols executed among the routers build the routing tables

- Two classes of routing protocols
  - Distance vector routing (RIP, BGP)
  - Link state routing (OSPF, IS-IS)

Distance Vector Routing

- Each router reports a list of *reachable destinations* to its neighbors
- Each router updates its internal tables according to the information received
- Key challenge: convergence can take long

RIP: example

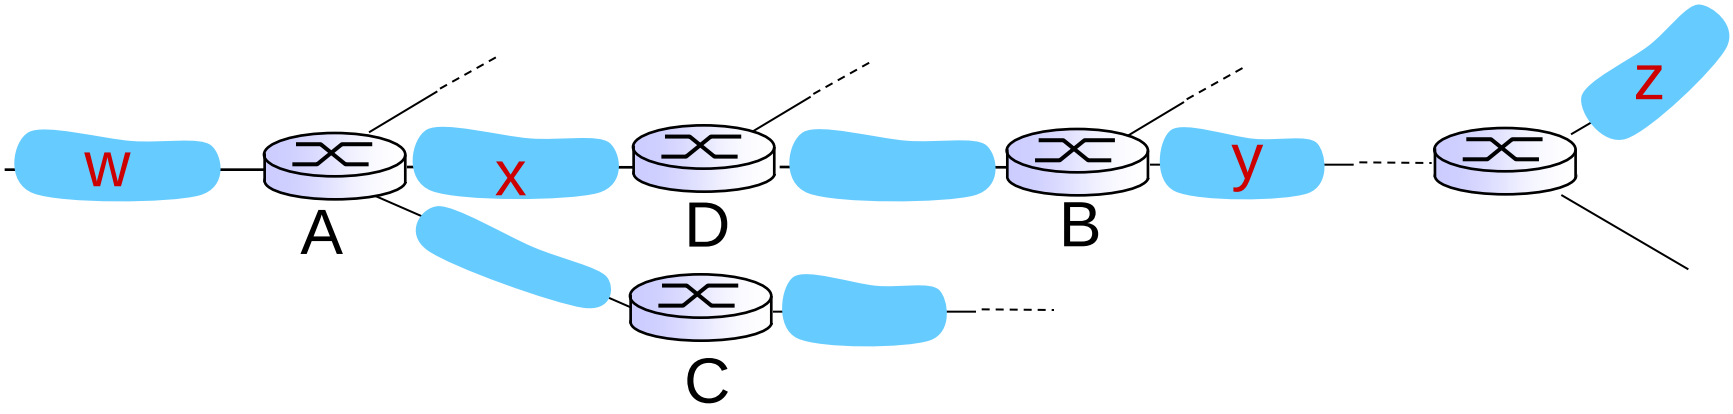

routing table in router D

destination subnet	next router	# hops to dest
w	A	2
y	B	2
z	B	7
x	--	1
….	….	….

RIP: example

A-to-D advertisement

dest	next	hops
w	-	1
x	-	1
z	C	4
....	...	...

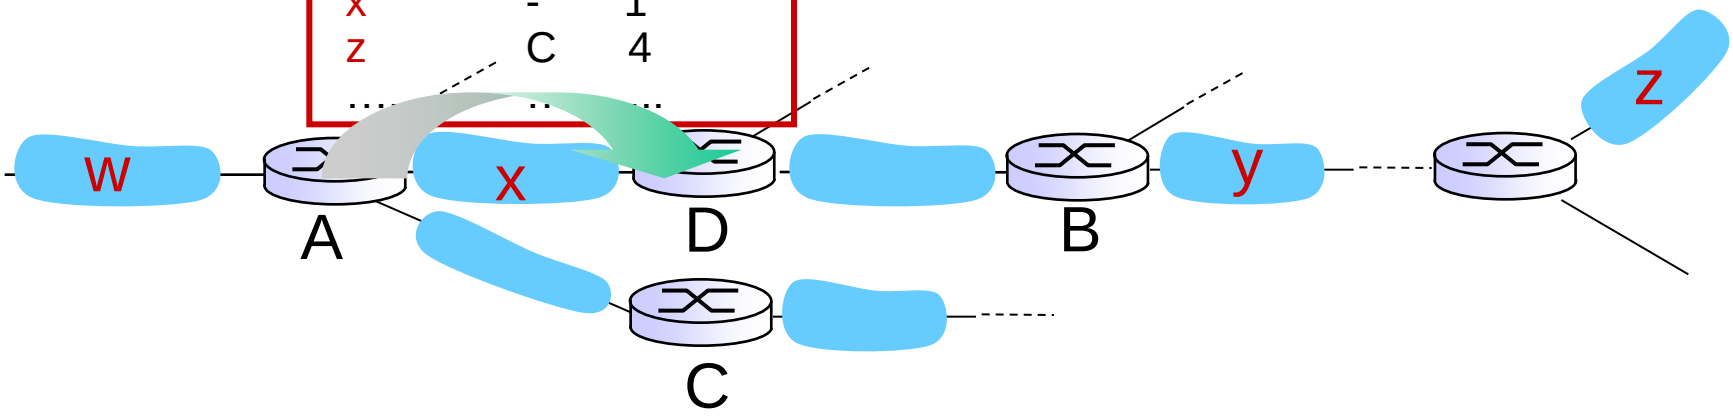

routing table in router D

destination subnet	next router	# hops to dest
w	A	2
y	B	2
z	B A	7 5
x	--	1
....	....	....

Link State Routing

- Each router:
  - compiles a list of directly connected neighbors
  - *floods* its list
  - Learns the full topology from the received lists
- Routers compute the best routes using Dijkstra
- Key challenge: controlled *flooding*

Flooding

- flooding: when node receives broadcast packet, sends copy to all neighbors
  - problems: cycles & broadcast storm

- *controlled flooding:* node only broadcasts pkt if it hasn't broadcast same packet before
  - node keeps track of packet ids already broadacsted

OSPF and IS-IS

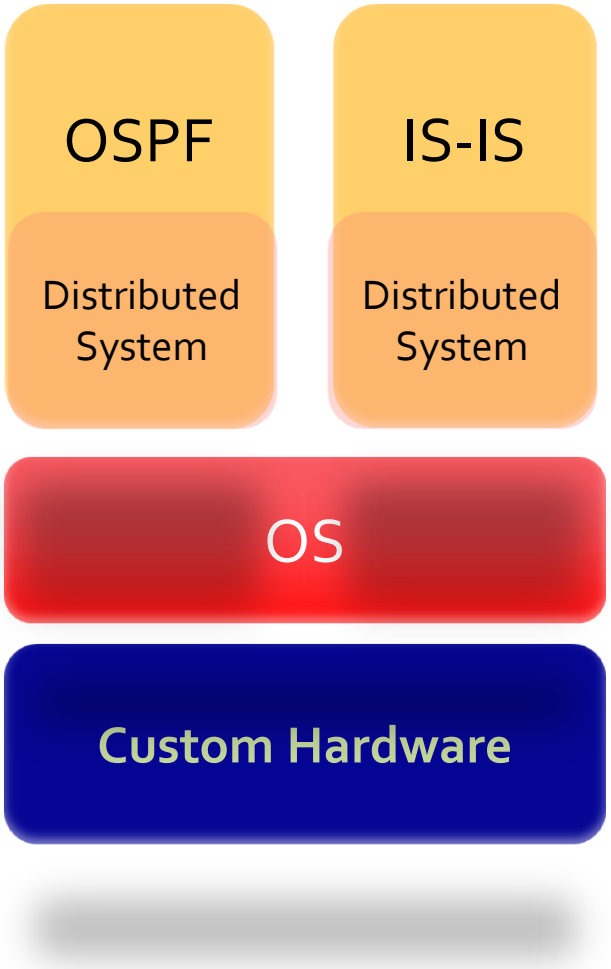

OSPF

- RFC 2328: **245 pages**

Distributed Protocol

- Builds consistent, up-to-date map of the network: **101 pages**

Dijkstra's Algorithm

- Operates on map: **4 pages**

OSPF and IS-IS in an SDN world

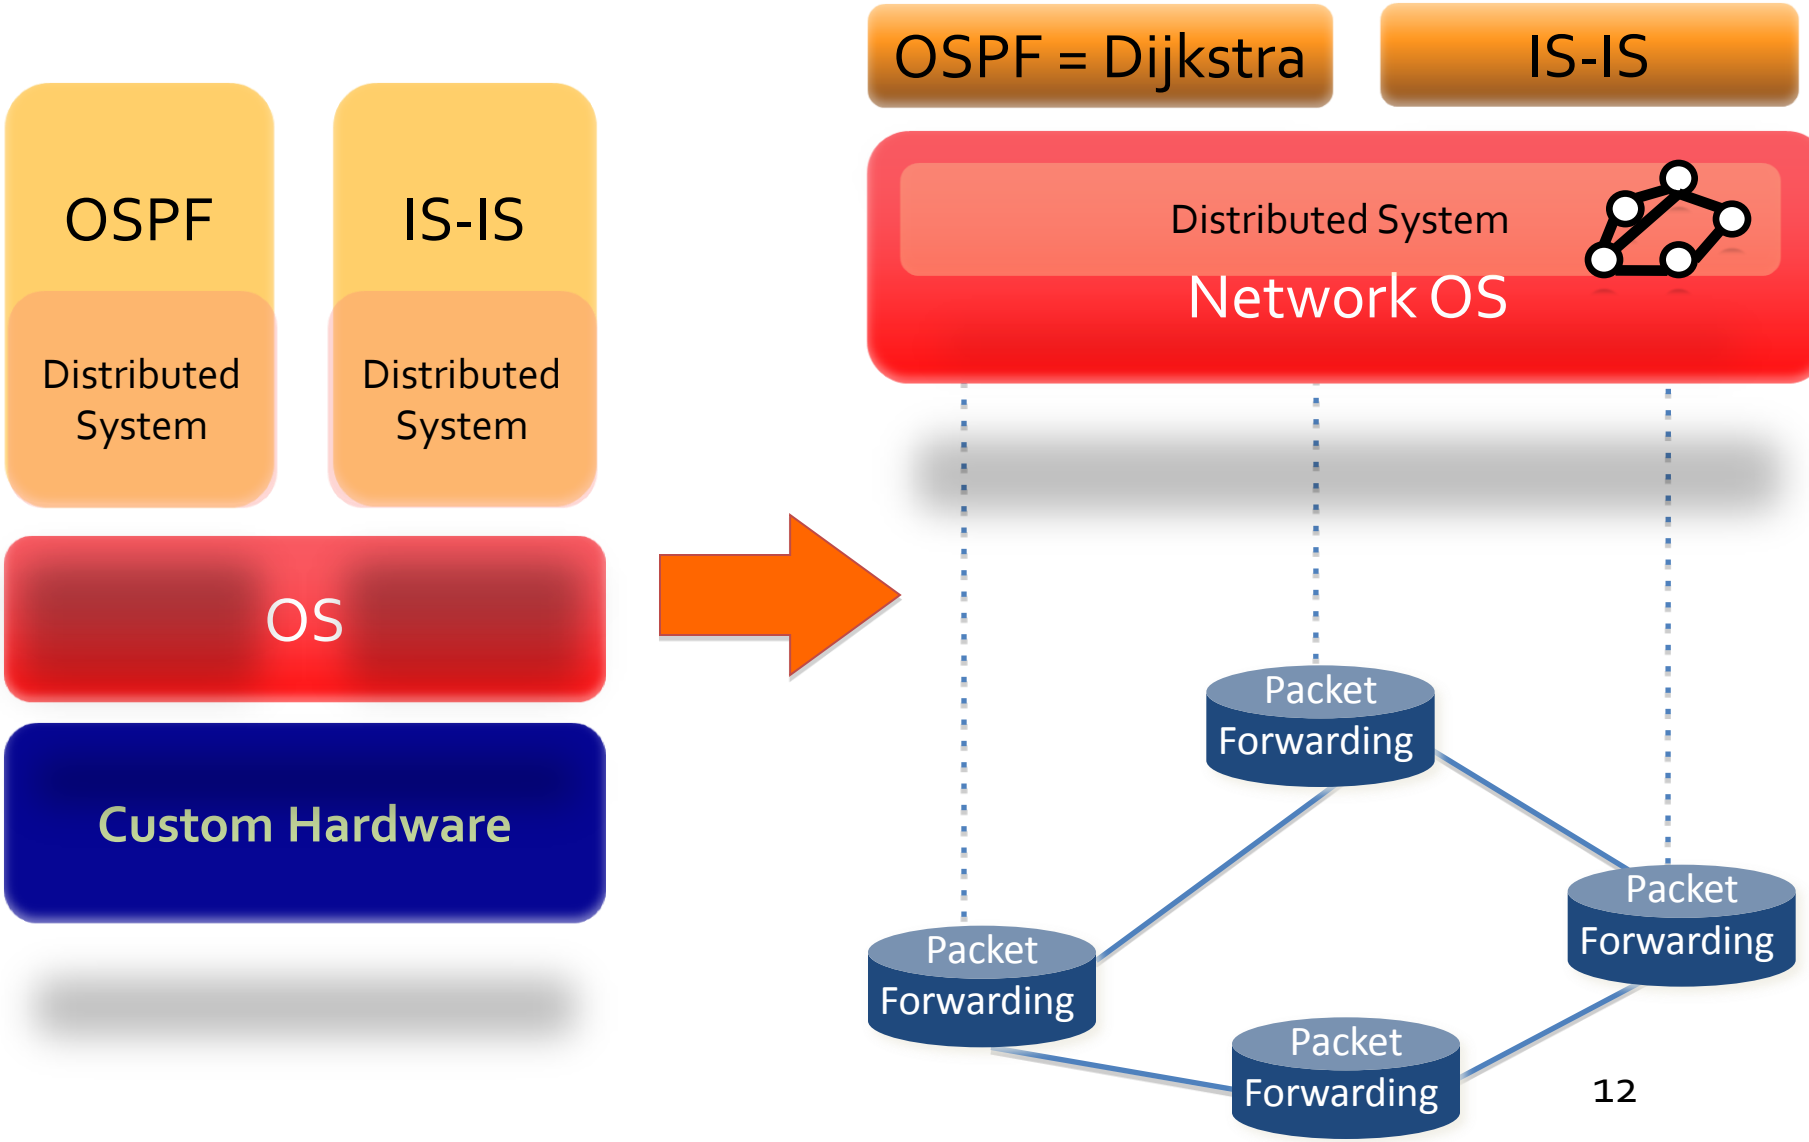

OpenFlow: A “southbound” API

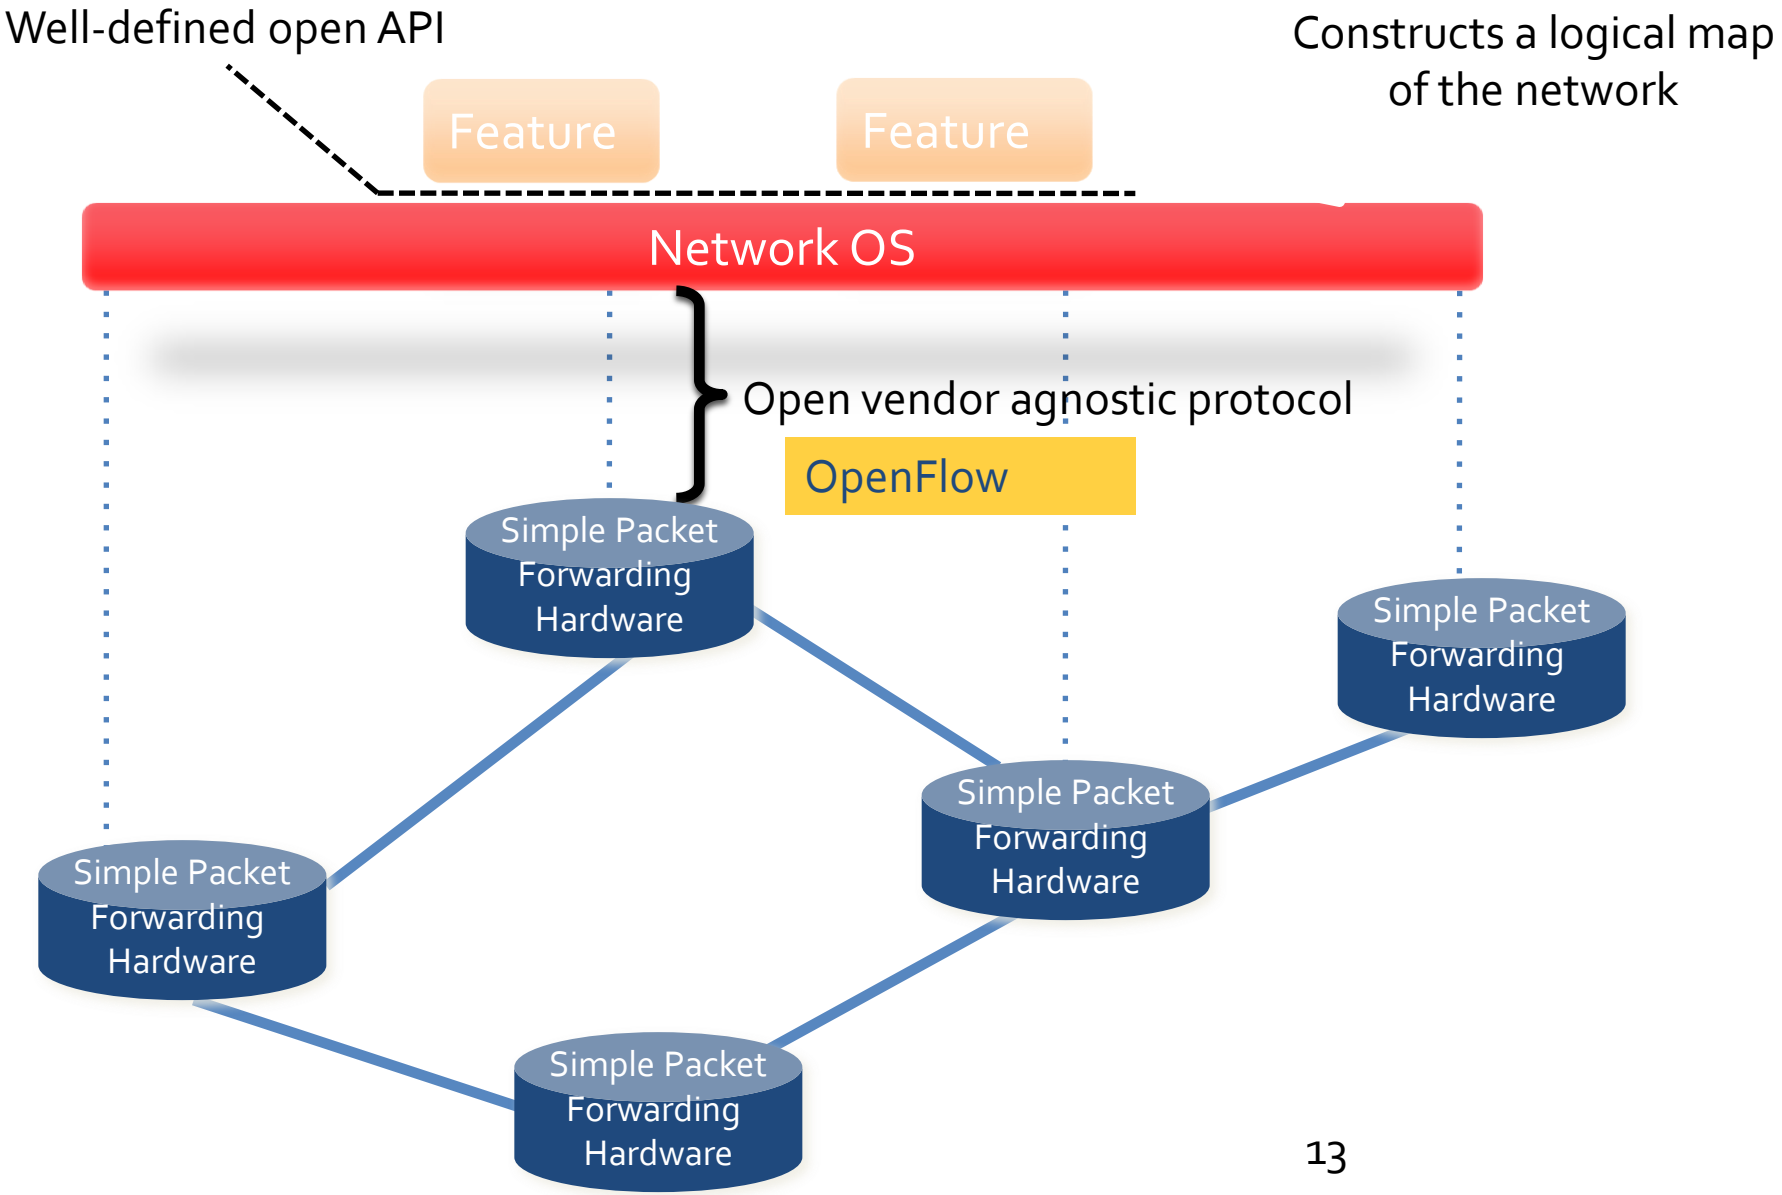

OpenFlow

- A protocol for remotely controlling the forwarding table of a switch or router
- A key SDN technology, although OpenFlow ≠ SDN

Ethernet Switch

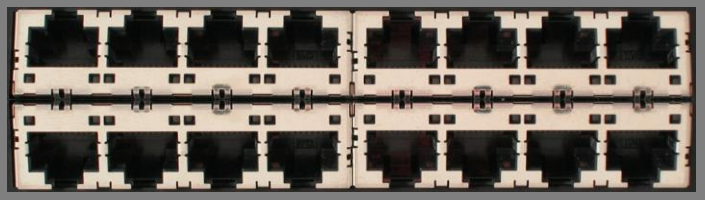

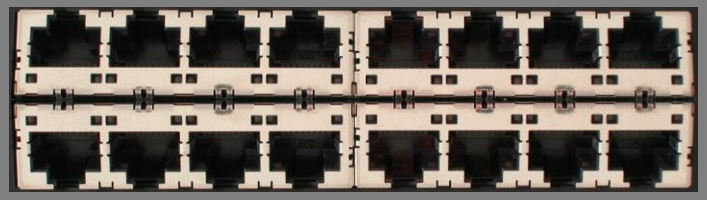


Control Path (Software)
Data Path (Hardware)

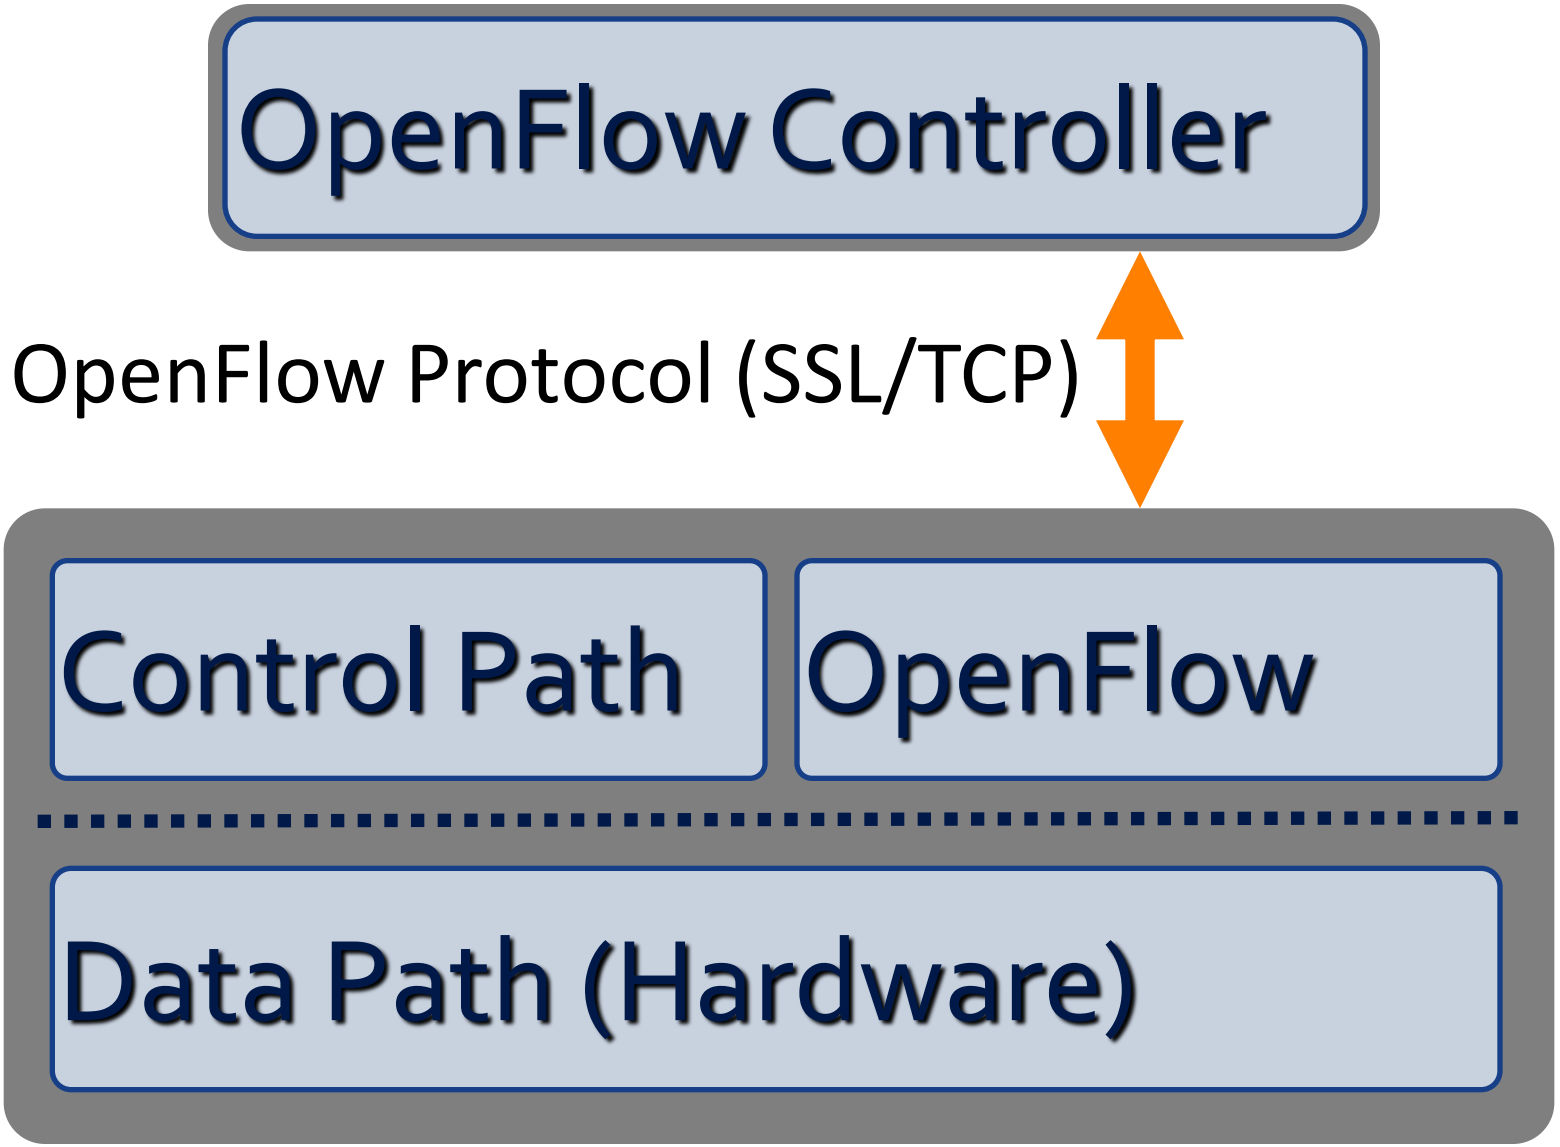
OpenFlow Controller
OpenFlow Protocol (SSL/TCP)
Control Path
OpenFlow
Data Path (Hardware)

Secure Channel

- SSL Connection, site-specific key
- Controller discovery protocol
- Encapsulate packets for controller
- Send link/port state to controller

Flow Table

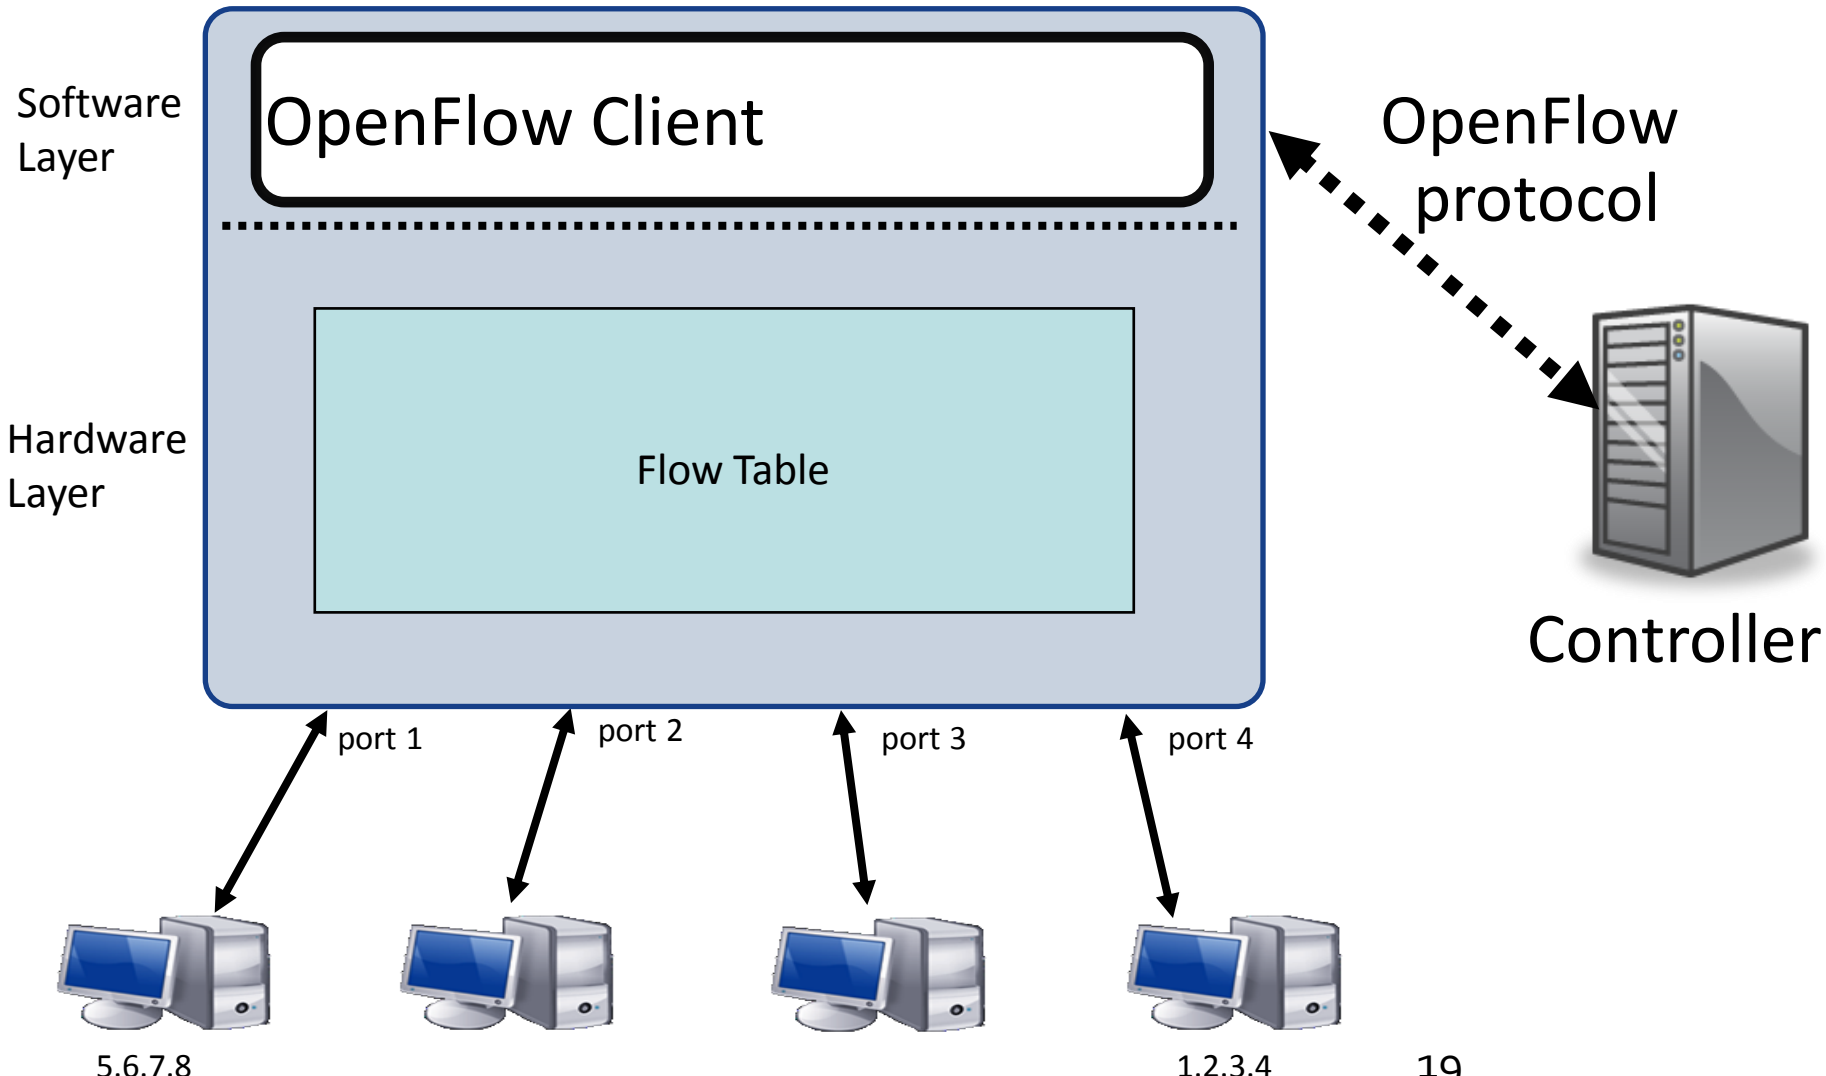
Software Layer
OpenFlow Client
OpenFlow protocol
Hardware Layer
Flow Table
Controller
port 1
port 2
port 3
port 4
5.6.7.8
1.2.3.4

Flow Table Entries

Main components of a flow entry in a flow table.

Match fields	To match against packets. These consist of the ingress port and packet headers
Priority	Matching precedence of the flow entry
Counters	e.g. packet and byte counters
Instructions	Determine action set or pipeline processing
Timeouts	Maximum amount of time or idle time before flow is expired by the switch
Cookies	Opaque data value chosen by the controller. Not used when processing packets.

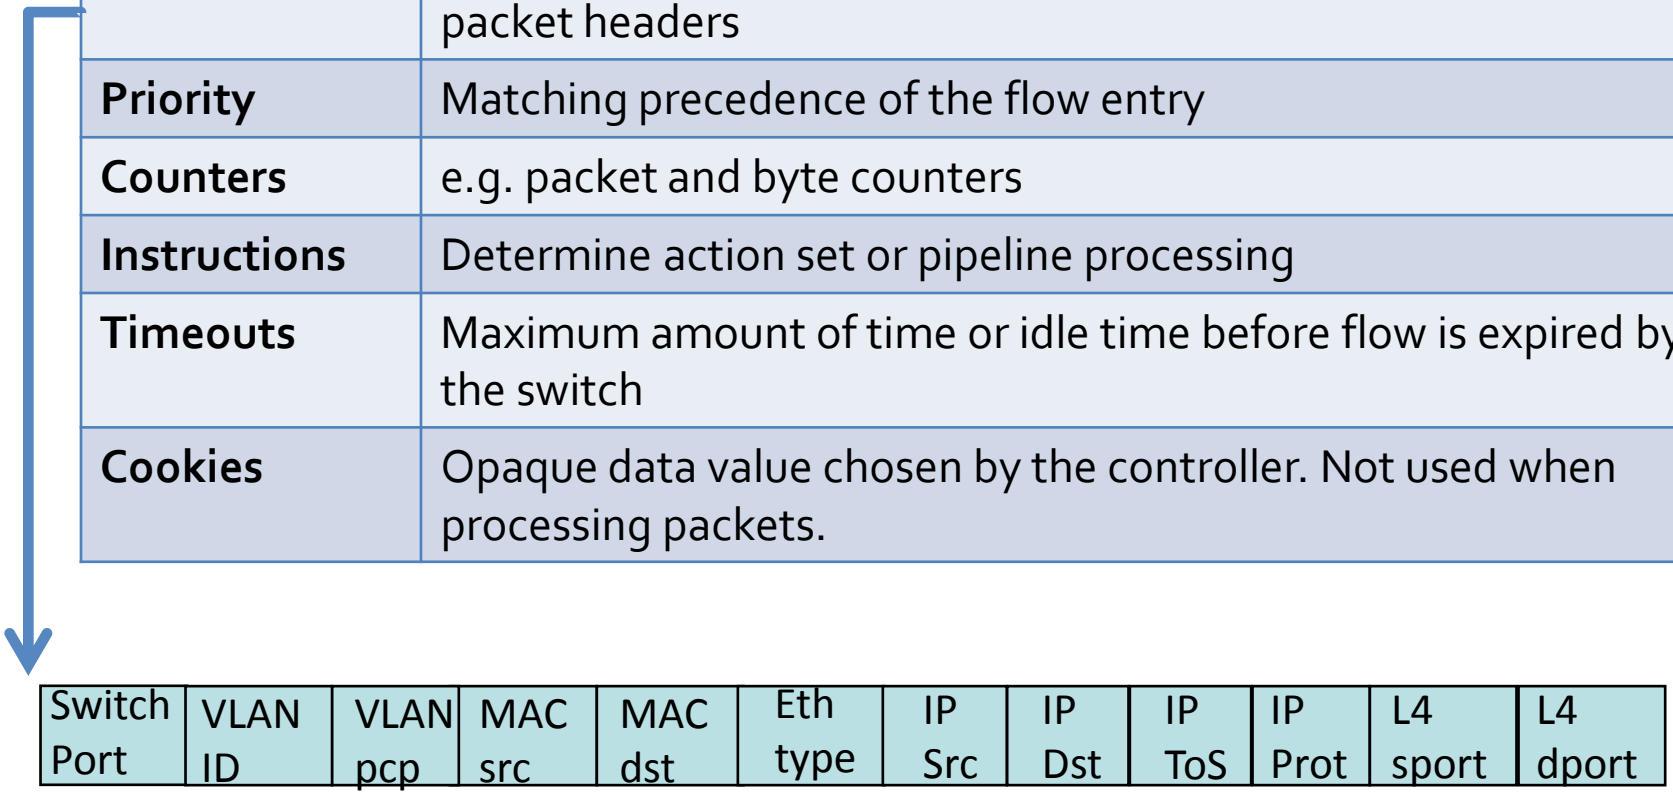

The match field contains either a specific value or a "wildcard"

Match/action 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

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

Firewall

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

Match/action examples

VLAN Switching

Switch Port	MAC src	MAC dst	Eth type	VLAN ID	IP Src	IP Dst	IP Prot	TCP sport	TCP dport	Action
*	*	00:1f..	*	vlan1	*	*	*	*	*	port6, port7, port9

Flow Switching

Switch Port	MAC src	MAC dst	Eth type	VLAN ID	IP Src	IP Dst	IP Prot	TCP sport	TCP dport	Action
port3	00:20..	00:1f..	0800	vlan1	1.2.3.4	5.6.7.8	4	17264	80	port6

Table miss

- Packets for which no flow has been defined are sent to the controller
- The controller creates one or more flow table entries
- The packet is then processed by the newly created flow entries

OpenFlow message types

- Controller-to-switch messages
  - Manage flow entries
  - Request info on switch capabilities and counters
  - Send a packet back to a switch
- Asynchronous messages
  - Send to controller a packet that does not match
  - Inform the controller that a timer has expired or that an error has occurred
- Symmetric messages
  - Hello and echo messages

OpenFlow key messages

Message	Direction	Description
Packet-In	Switch->Controller	Transfer the control of a packet to the controller. Packet-in events can be configured to buffer packets
Packet-Out	Controller->Switch	Instruct switch to send a packet out of a specified port. Send in response to Packet-in messages.
Modify-State	Controller->Switch	Add, delete and modify flow/group entries in the flow tables and to set switch port properties
Flow-Removed	Switch->Controller	Inform the controller about the removal of a flow entry from a flow table

Example on the Board

Features of Flow Table

- Backwards compatible
  - Generalization of forwarding/switch table
  - No need to change end hosts
- Easily implemented in hardware
  - e.g. TCAM flow-table in each switch
- Strong isolation of flows
  - Simple geometric construction
  - Can prove which flows can/cannot communicate

“Flowspace”: A way to think about packets defined by match fields

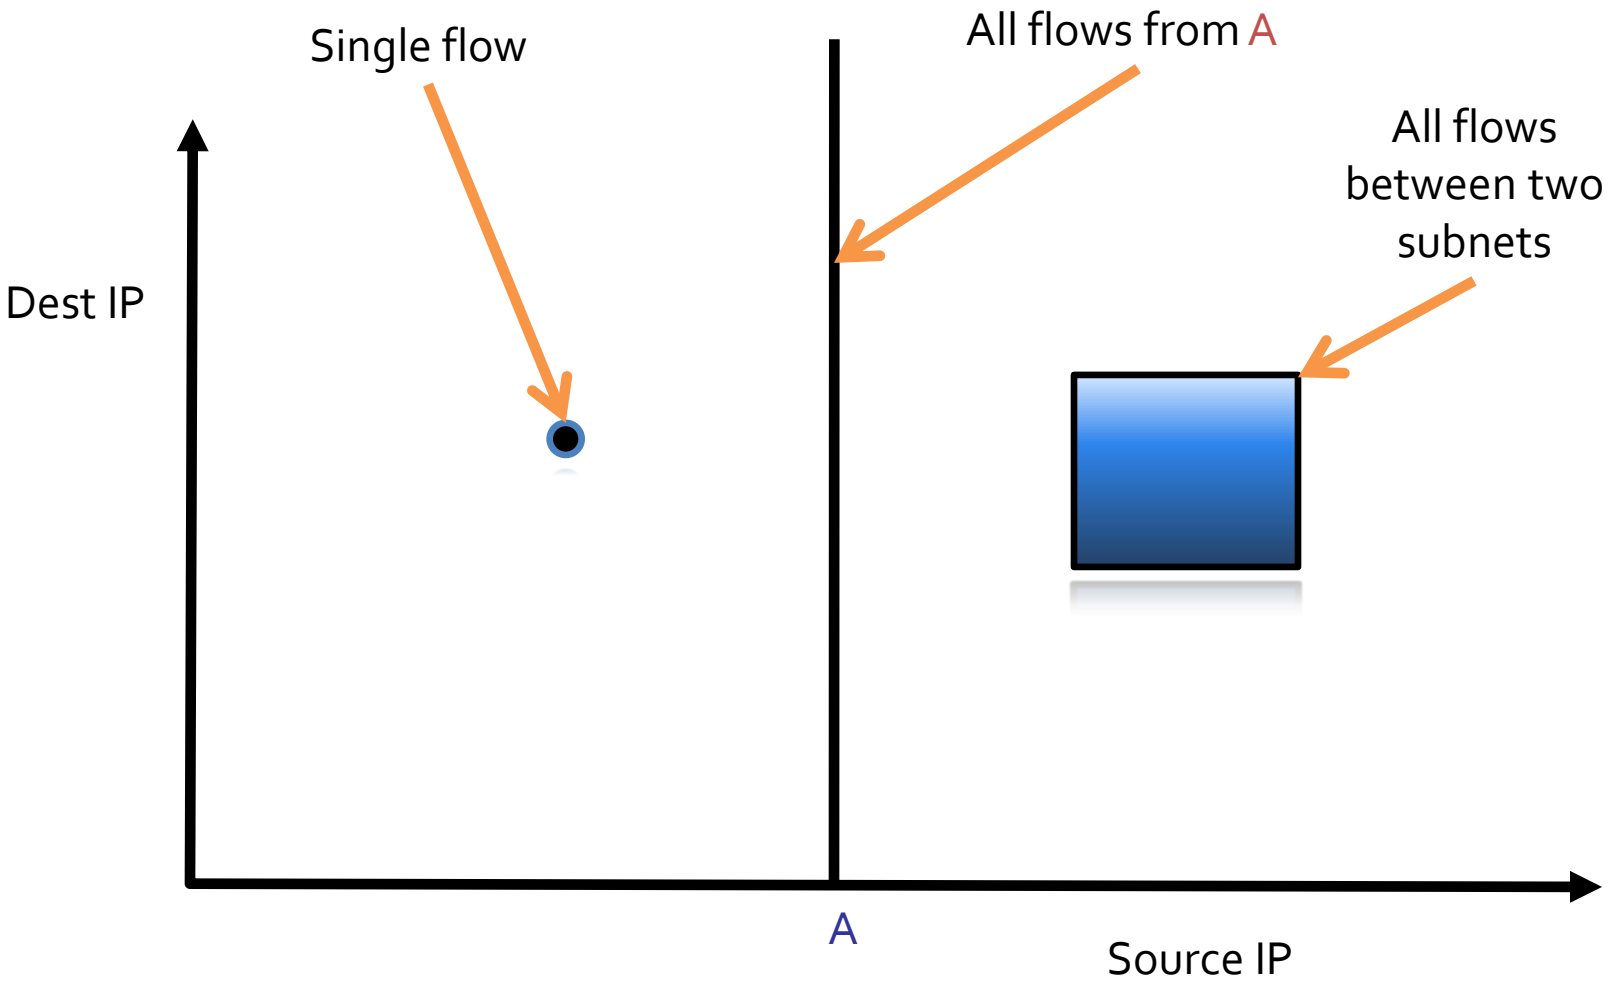

Figure thanks to Yashar Ganjali

Flowspace: Multiple Dimensions

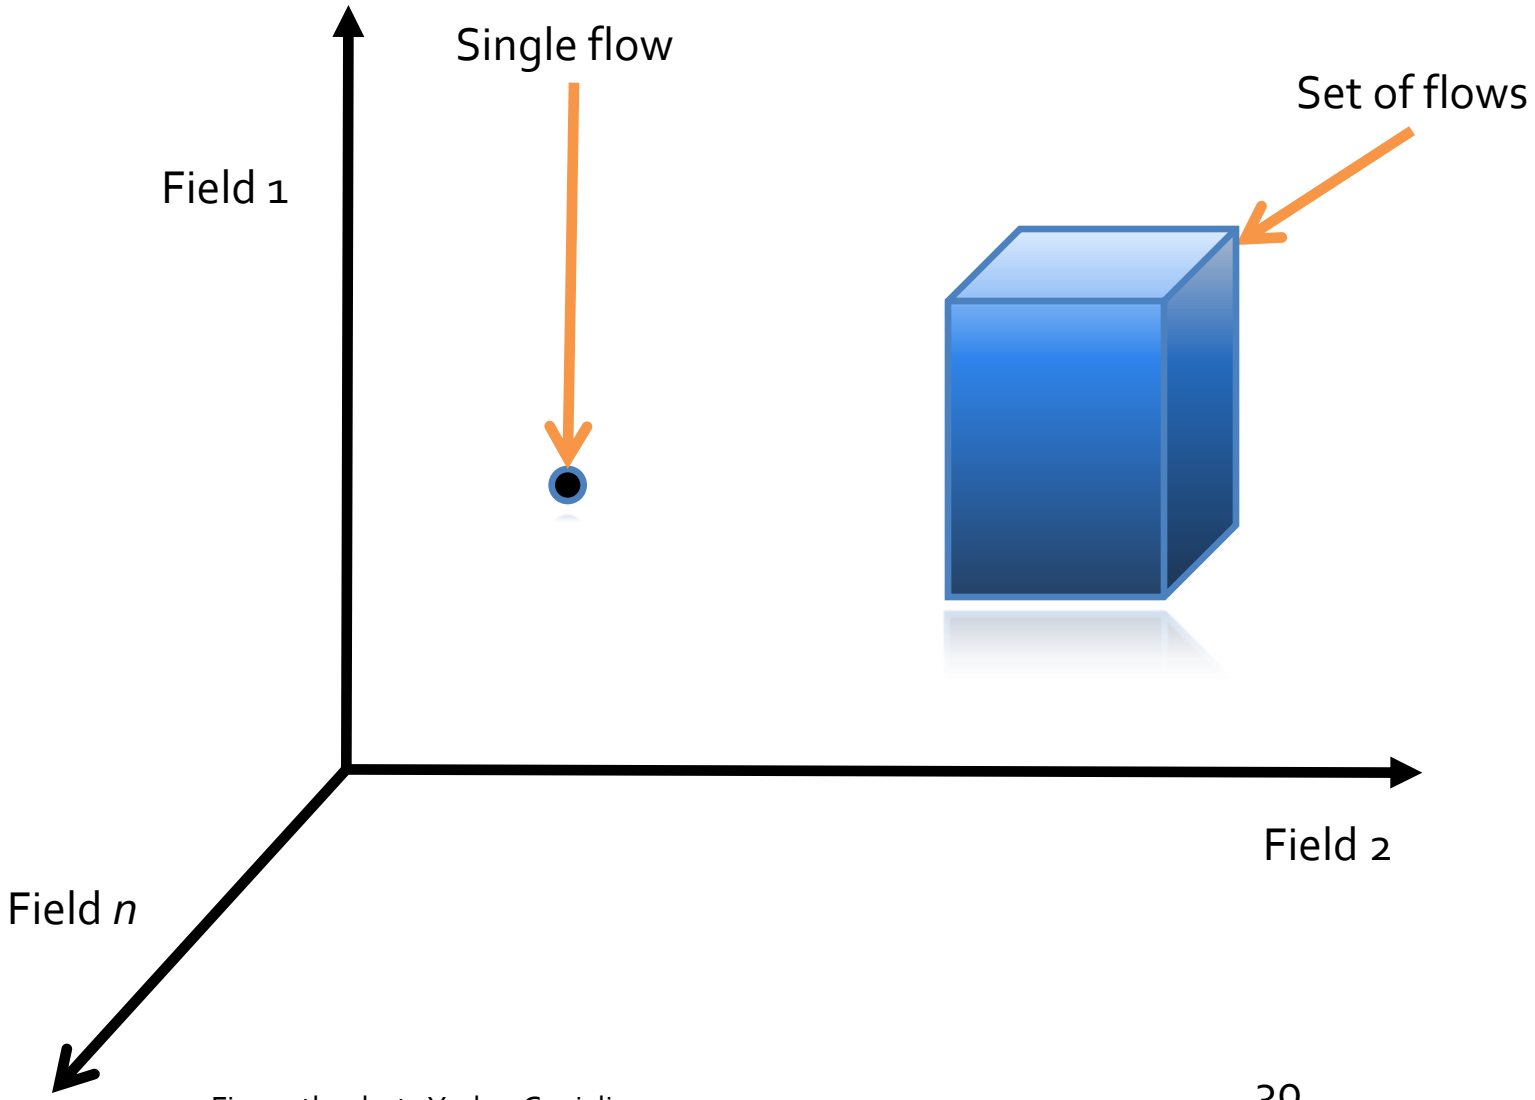

Figure thanks to Yashar Ganjali

Agenda

- SDN vs distributed routing protocols
- OpenFlow:
  - Flow table and messages
  - How it started?
    “Innovate in your campus network” use case
  - Advanced features (v1.4)

How It Started?

Based on McKeown's original 2008 slides

OpenFlow or "Why I can't innovate in my wiring closet?"

Innovations in campus wiring closets

Experiments we'd like to do

- Mobility management
- Network-wide energy management
- New naming/addressing schemes
- Network access control

Problem with our network

- Paths are fixed (by the network)
- IP-only
- Addresses dictated by DNS, DHCP, etc
- No means to add our own processing

Experimenter's Dream
(Vendor's Nightmare)

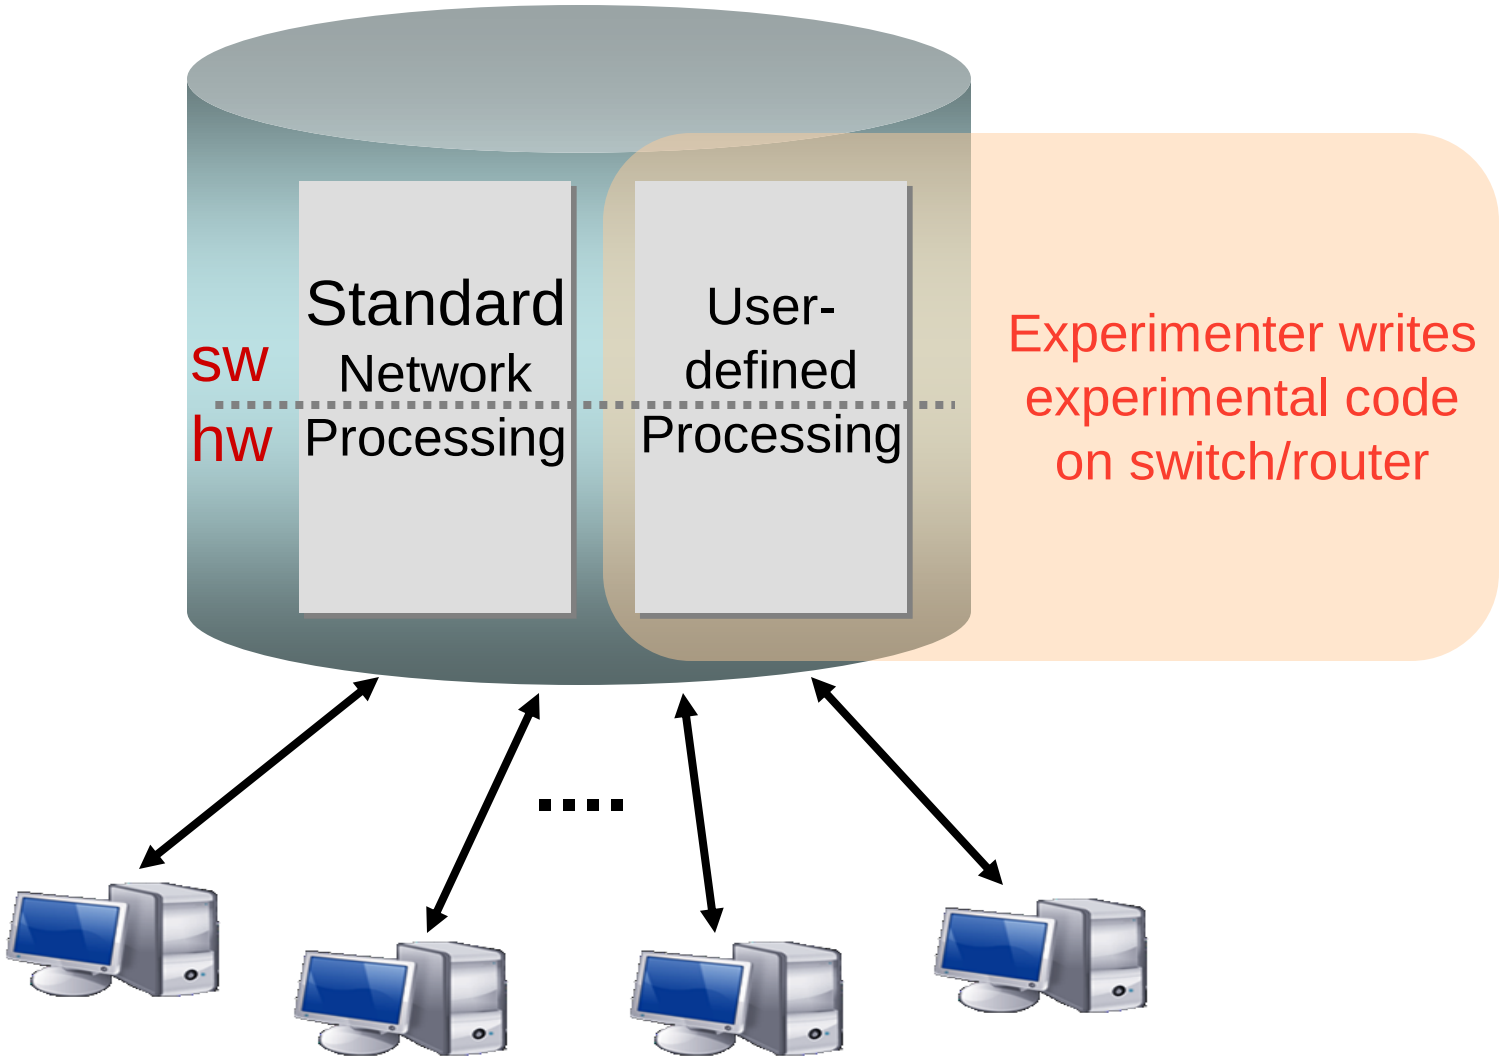

Switch Based Virtualization

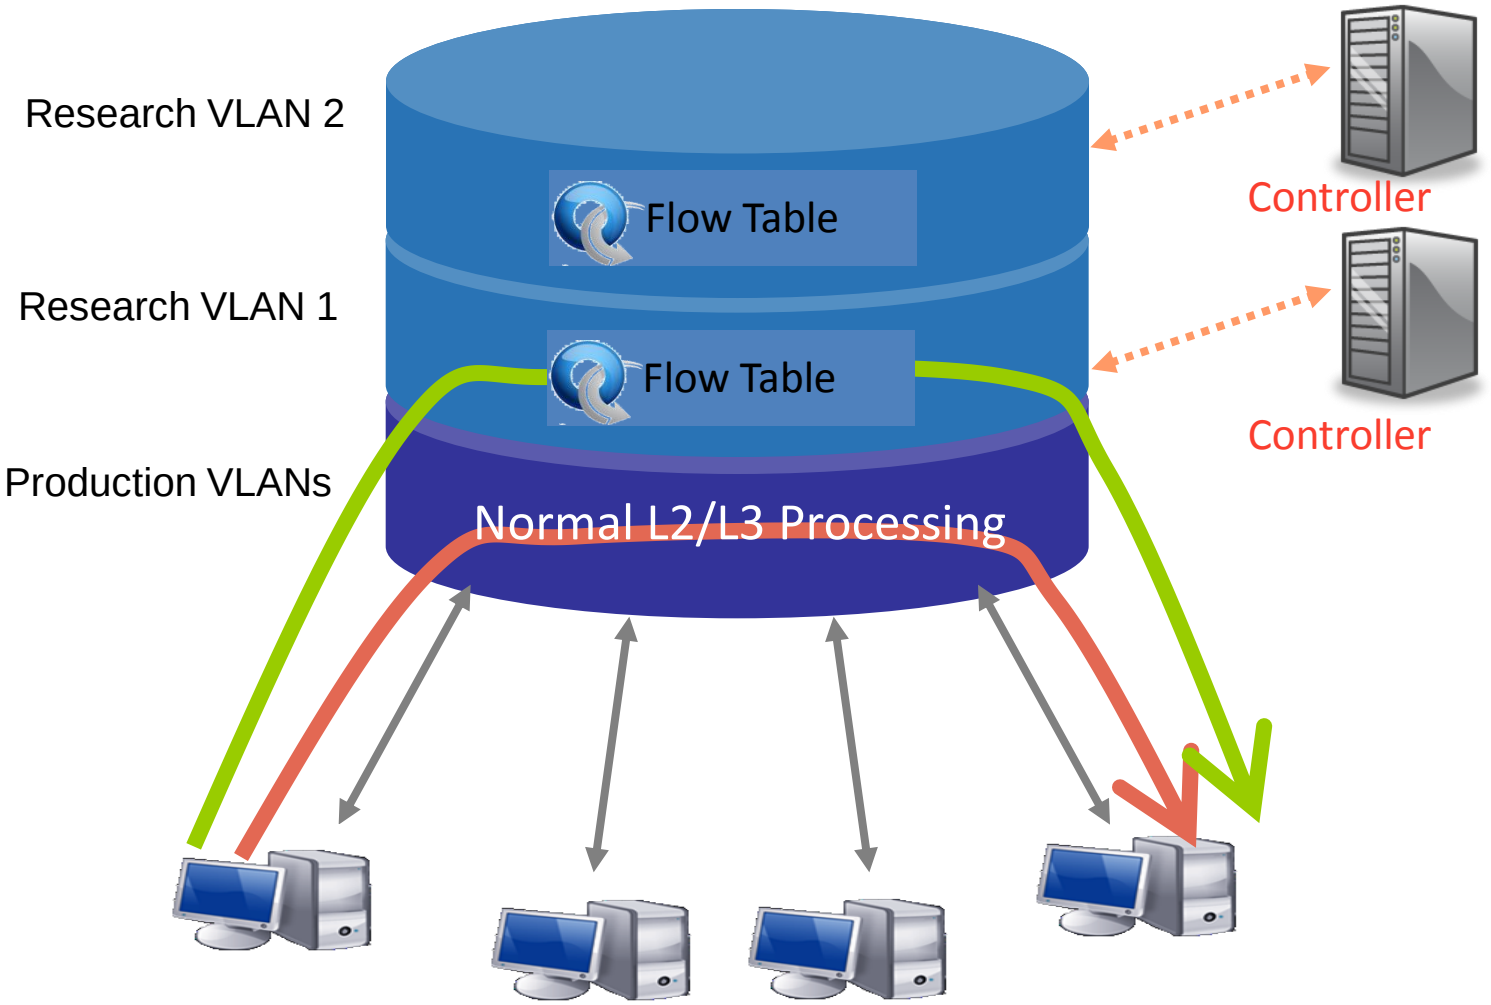

Use Case: VLAN Based Partitioning

- Basic Idea: Partition Flows based on Ports and VLAN Tags
  - Traffic entering system (e.g. from end hosts) is tagged
  - VLAN tags consistent throughout substrate

	Switch Port	MAC src	MAC dst	Eth type	VLAN ID	IP Src	IP Dst	IP Prot	TCP sport	TCP dport
Dave	*	*	*	*	1,2,3	*	*	*	*	*
Larry	*	*	*	*	4,5,6	*	*	*	*	*
Steve	*	*	*	*	7,8,9	*	*	*	*	*

Use Case: New CDN - Turbo Akamai ++

- Basic Idea: Build a CDN where you control the entire network
  - All traffic to or from Akamai IP space controlled by Experimenter
  - All other traffic controlled by default routing
  - Topology is the entire network
  - End hosts are automatically added (no opt-in)

	Switch Port	MAC src	MAC dst	Eth type	VLAN ID	IP Src	IP Dst	IP Prot	TCP sport	TCP dport
Turbo Akamai	*	*	*	*	*	84.65.*	*	*	*	*
	*	*	*	*	*	*	84.65.*	*	*	*
Default	*	*	*	*	*	*	*	*	*	*

Use Case: Your Internet Protocol

- A new layer 3 protocol
- Replaces IP
- Defined by a new Ether Type

	Switch Port	MAC src	MAC dst	Eth type	VLAN ID	IP Src	IP Dst	IP Prot	TCP sport	TCP dport
Your IP	*	*	*	YourIP	*	*	*	*	*	*
Rest	*	*	*	!YourIP	*	*	*	*	*	*

Headers as a protocol-agnostic collection of bits

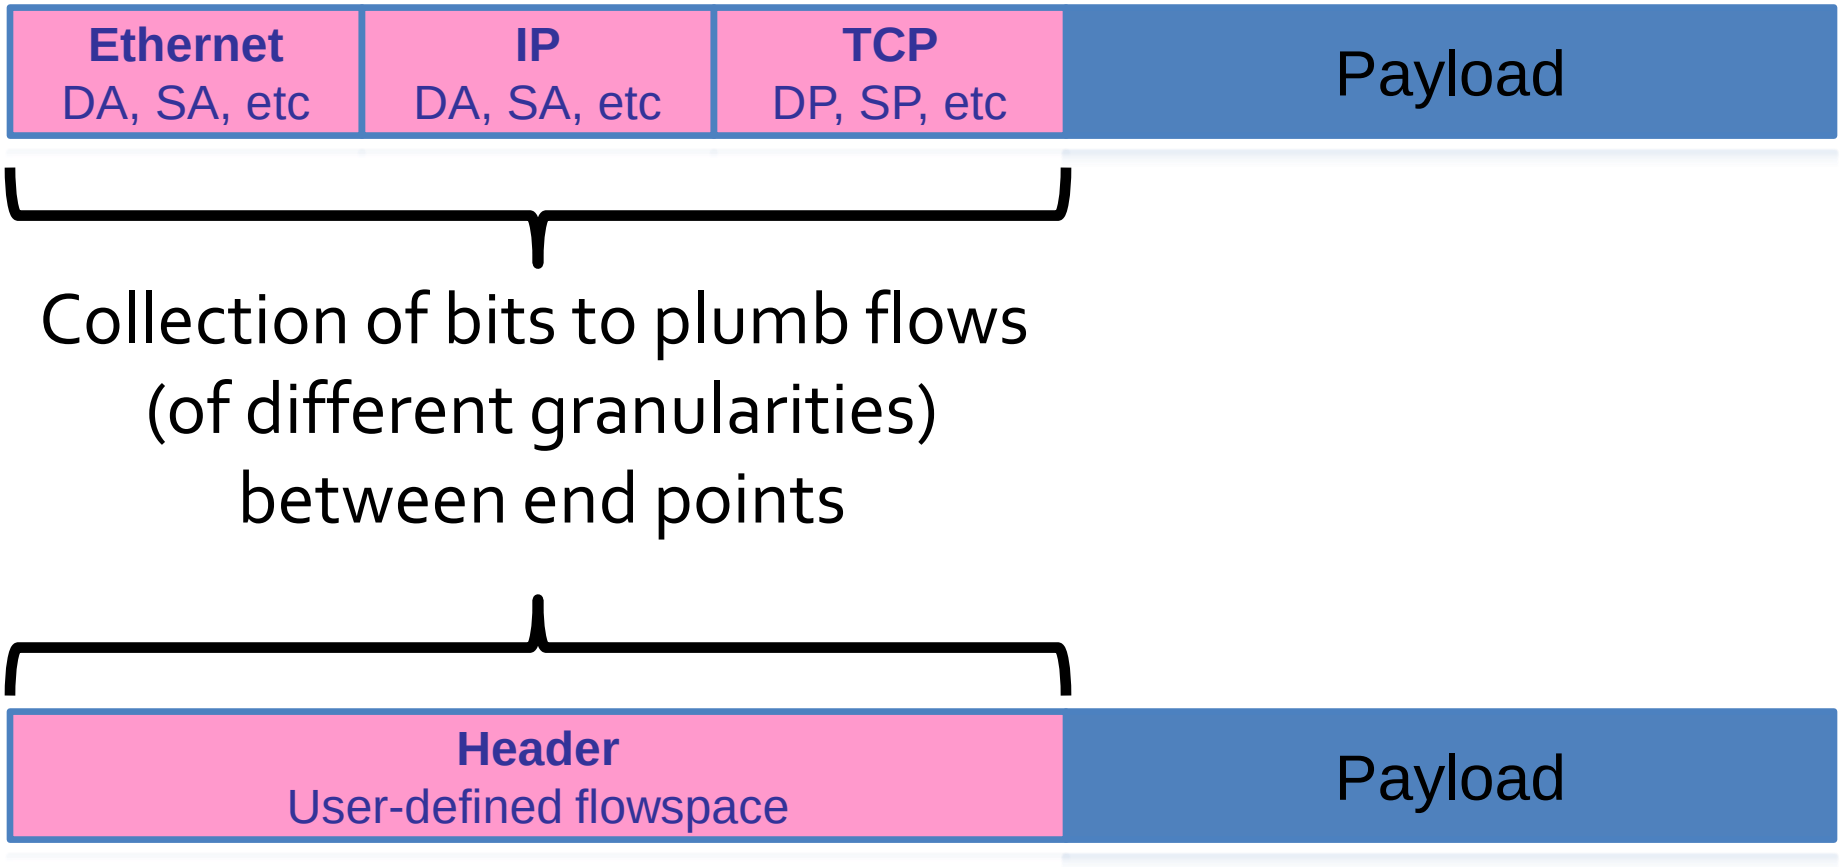

Agenda

- SDN vs distributed routing protocols
- OpenFlow:
  - Flow table and messages
  - How it started?
    “Innovate in your campus network” use case
  - Advanced features (v1.4)

Openflow's short history

- OF v1.0 (end of 2009): Single table, L2+IPv4 focused matching
- OF v1.1 (Mar 2011):
  - multiple tables, MPLS + VLAN matching, multipath forwarding: ECMP, groups
- OF v1.2 (Dec 2011): "Extensible Protocol"
  - extensible match & actions (TLV), IPv6, multiple controllers
- OF v1.3 (June 12):
  - Better expression of capabilities of a switch, meters, multiple parallel channels between switch and controller
- OF v1.4 (Aug 13):
  - Improve extensibility, better support for optical ports, many other incremental improvements

Main components of OpenFlow v1.4 switch

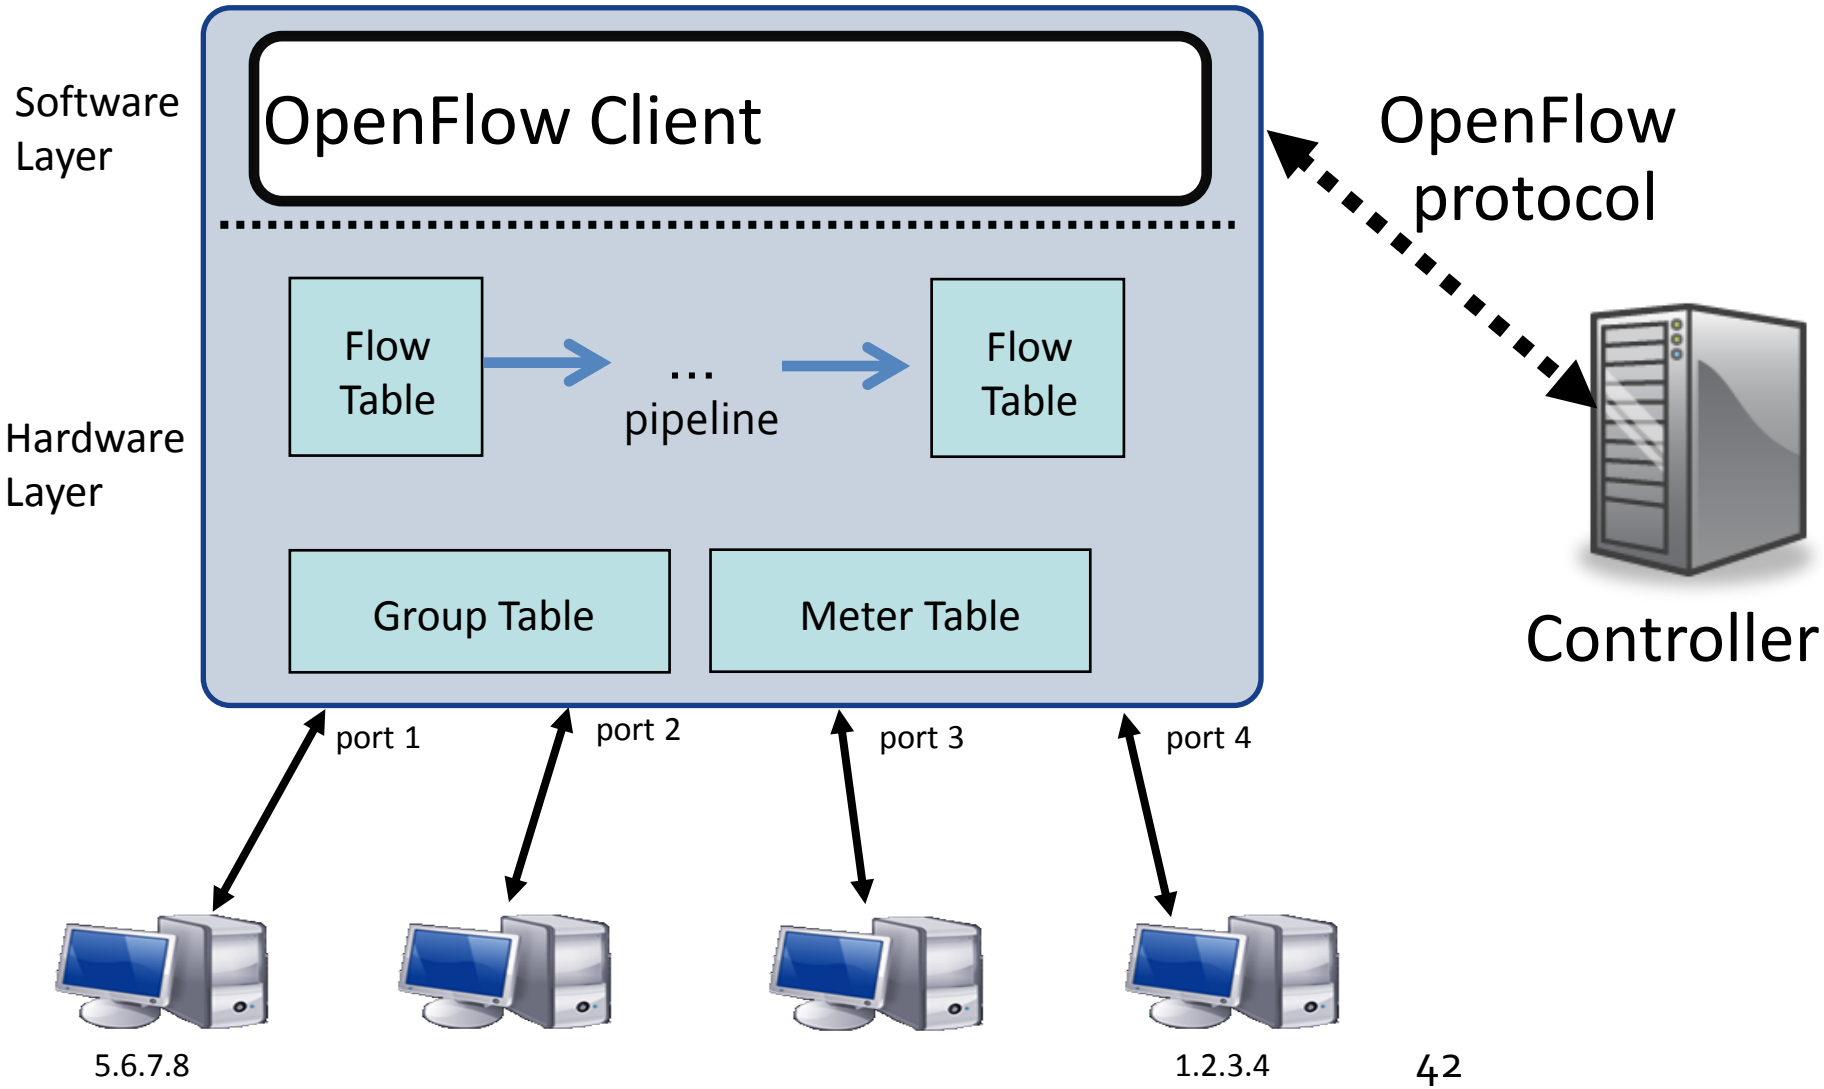

OpenFlow Pipeline

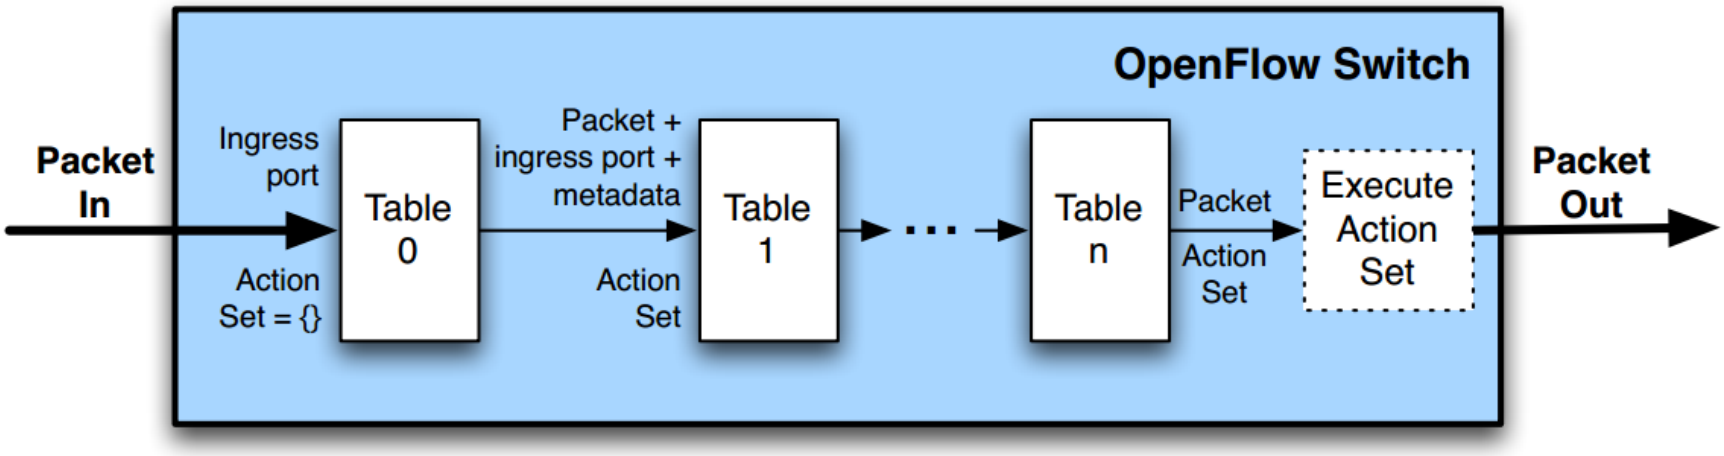

Packets are matched against multiple tables in the pipeline

OpenFlow Switch Specification Version > 1.1.0

Packet Flow

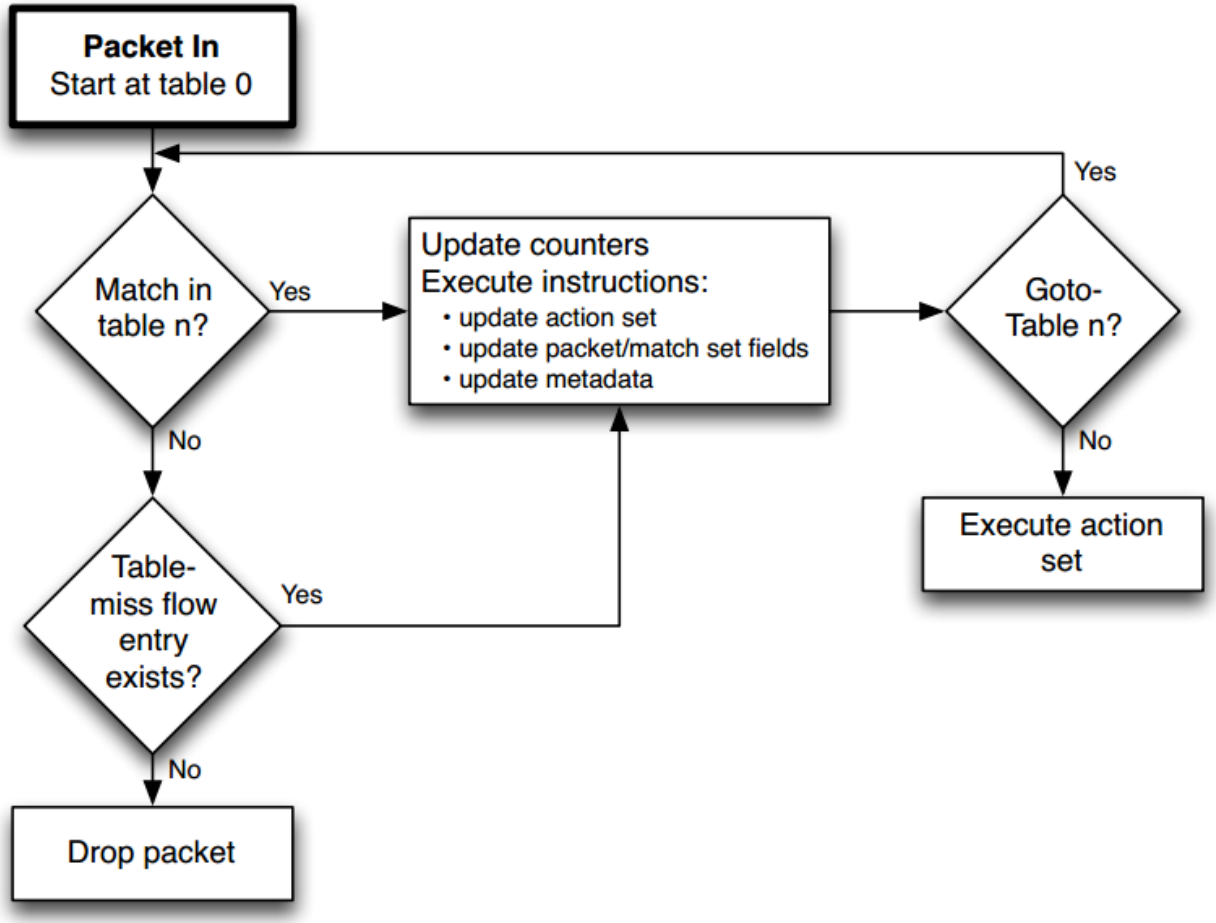

Packet Flow through OpenFlow Switch

Why multiple flow tables?

- Separate logical functions
- Example
  - Table 1: Input firewall rules
  - Table 2: Network address translation
  - Table 3: Routing
- See: “OpenFlow with multiple Flow Tables” http://www.youtube.com/watch?v=TD5wmoD7XOE

Terminology: Actions vs Instructions

- Action: decrement TTL, forward, etc.

- Instructions: Move pkt through pipeline
  - Direct to another flow table
  - Add or apply actions

Group and meter tables

- Groups represent sets of actions for more complex forwarding, e.g. flooding, multipath

- A meter table consists of per-flow meters, e.g. rate limit packets to controller

OpenFlow Support

- Open Networking Foundation was founded in 2011 to develop and standardize OpenFlow. Members include Cisco, Facebook, Google, HP, IBM and Juniper Networks.
- Juniper and start-ups Nicira and Big Switch are warm supporters of OpenFlow
- Vendors, such as the Cisco, IBM, NEC and HP, have implemented OpenFlow in existing products
- Cisco’s SDN initiative is called Open Network Environment (ONE)

Summary of key SDN/Openflow features

- Separate data from control
- Open control API
- Define a generalized flow table
  - Flexible and generalized flow abstraction
  - Unified view of layers1-7
- Backward compatible
  - Though allows completely new header
- Virtualization of the data and control plane

Further reading

- OpenFlow Switch Specification https://www.opennetworking.org/images/stories/downloads/sdn-resources/onf-specifications/openflow/openflow-spec-v1.4.0.pdf
- OVSDB: Control vs Management http://keepingitclassless.net/2014/08/sdn-protocols-3-ovsdb/
- OPFLEX: http://keepingitclassless.net/2014/09/sdn-protocols-4-opflex-declarative-networking/

Τέλος Ενότητας

Ευρωπαϊκή Ένωση
Ευρωπαϊκό Κοινωνικό Ταμείο
ΕΠΙΧΕΙΡΗΣΙΑΚΟ ΠΡΟΓΡΑΜΜΑ
ΕΚΠΑΙΔΕΥΣΗ ΚΑΙ ΔΙΑ ΒΙΟΥ ΜΑΘΗΣΗ
επένδυση στην κοινωνία της γνώσης
ΥΠΟΥΡΓΕΙΟ ΠΑΙΔΕΙΑΣ & ΘΡΗΣΚΕΥΜΑΤΩΝ, ΠΟΛΙΤΙΣΜΟΥ & ΑΘΛΗΤΙΣΜΟΥ
ΕΙΔΙΚΗ ΥΠΗΡΕΣΙΑ ΔΙΑΧΕΙΡΙΣΗΣ
ΕΣΠΑ 2007-2013
ΕΥΡΩΠΑΪΚΟ ΚΟΙΝΩΝΙΚΟ ΤΑΜΕΙΟ
Με τη συγχρηματοδότηση της Ελλάδας και της Ευρωπαϊκής Ένωσης

Χρηματοδότηση

- Το παρόν εκπαιδευτικό υλικό έχει αναπτυχθεί στα πλαίσια του εκπαιδευτικού έργου του διδάσκοντα.
- Το έργο «**Ανοικτά Ακαδημαϊκά Μαθήματα στο Πανεπιστήμιο Κρήτης**» έχει χρηματοδοτήσει μόνο τη αναδιαμόρφωση του εκπαιδευτικού υλικού.
- Το έργο υλοποιείται στο πλαίσιο του Επιχειρησιακού Προγράμματος «Εκπαίδευση και Δια Βίου Μάθηση» και συγχρηματοδοτείται από την Ευρωπαϊκή Ένωση (Ευρωπαϊκό Κοινωνικό Ταμείο) και από εθνικούς πόρους.

Ευρωπαϊκή Ένωση
Ευρωπαϊκό Κοινωνικό Ταμείο
ΕΠΙΧΕΙΡΗΣΙΑΚΟ ΠΡΟΓΡΑΜΜΑ
ΕΚΠΑΙΔΕΥΣΗ ΚΑΙ ΔΙΑ ΒΙΟΥ ΜΑΘΗΣΗ
επένδυση στην κοινωνία της γνώσης
ΥΠΟΥΡΓΕΙΟ ΠΑΙΔΕΙΑΣ ΚΑΙ ΘΡΗΣΚΕΥΜΑΤΩΝ
ΕΙΔΙΚΗ ΥΠΗΡΕΣΙΑ ΔΙΑΧΕΙΡΙΣΗΣ
ΕΣΠΑ 2007-2013
ΕΥΡΩΠΑΪΚΟ ΚΟΙΝΩΝΙΚΟ ΤΑΜΕΙΟ
Με τη συγχρηματοδότηση της Ελλάδας και της Ευρωπαϊκής Ένωσης

Σημειώματα

Σημείωμα αδειοδότησης

- Το παρόν υλικό διατίθεται με τους όρους της άδειας χρήσης Creative Commons Αναφορά, Μη Εμπορική Χρήση, Όχι Παράγωγο Έργο 4.0 [1] ή μεταγενέστερη, Διεθνής Έκδοση. Εξαιρούνται τα αυτοτελή έργα τρίτων π.χ. φωτογραφίες, διαγράμματα κ.λ.π., τα οποία εμπεριέχονται σε αυτό και τα οποία αναφέρονται μαζί με τους όρους χρήσης τους στο «Σημείωμα Χρήσης Έργων Τρίτων».

[1] http://creativecommons.org/licenses/by-nc-nd/4.0/

- Ως **Μη Εμπορική** ορίζεται η χρήση:
  - που δεν περιλαμβάνει άμεσο ή έμμεσο οικονομικό όφελος από την χρήση του έργου, για το διανομέα του έργου και αδειοδόχο
  - που δεν περιλαμβάνει οικονομική συναλλαγή ως προϋπόθεση για τη χρήση ή πρόσβαση στο έργο
  - που δεν προσπορίζει στο διανομέα του έργου και αδειοδόχο έμμεσο οικονομικό όφελος (π.χ. διαφημίσεις) από την προβολή του έργου σε διαδικτυακό τόπο

- Ο δικαιούχος μπορεί να παρέχει στον αδειοδόχο ξεχωριστή άδεια να χρησιμοποιεί το έργο για εμπορική χρήση, εφόσον αυτό του ζητηθεί.

Σημείωμα Αναφοράς