On the Co-Existence of Distributed and Centralized Routing Control-Planes

Stefano Vissicchio

UCLouvain

Infocom

28th April 2015

Joint work with

L. Cittadini (RomaTre), O. Bonaventure (UCLouvain), G. G. Xie (NPS), L. Vanbever (ETH)

Control-planes include decision-making components of network architectures

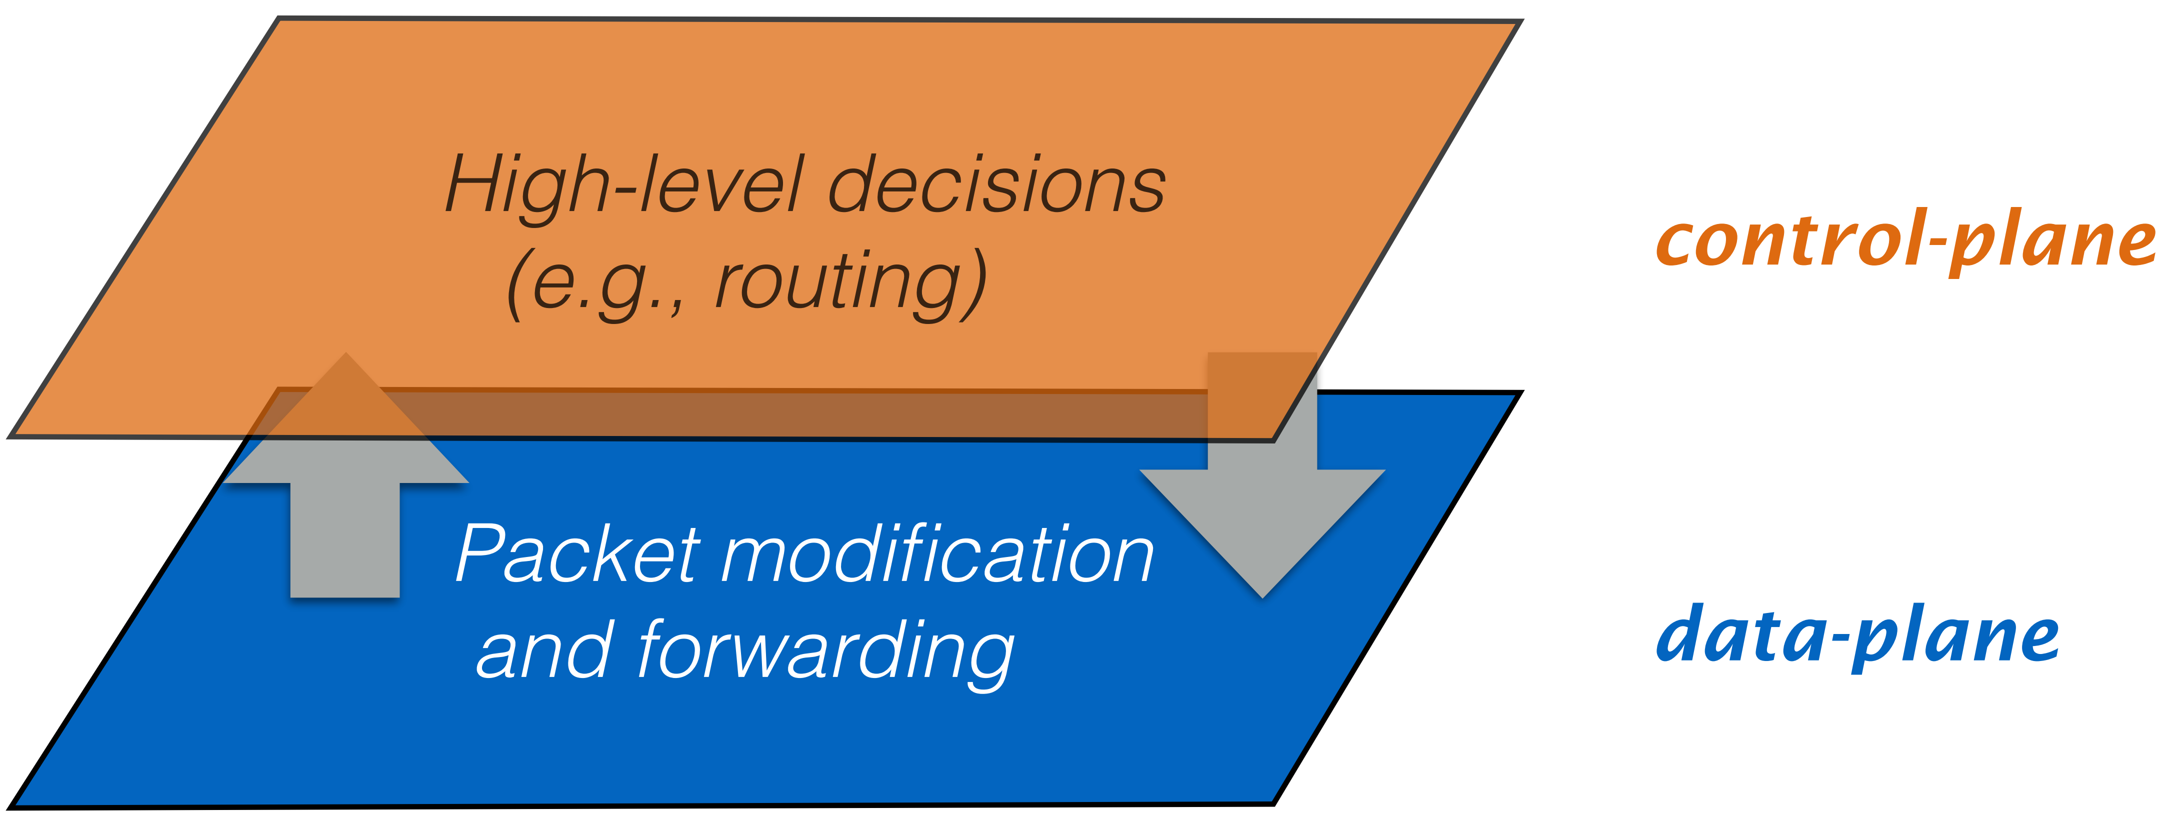

Traditional control-planes are distributed
(for example, IGP protocols like EIGRP, OSPF or IS-IS)

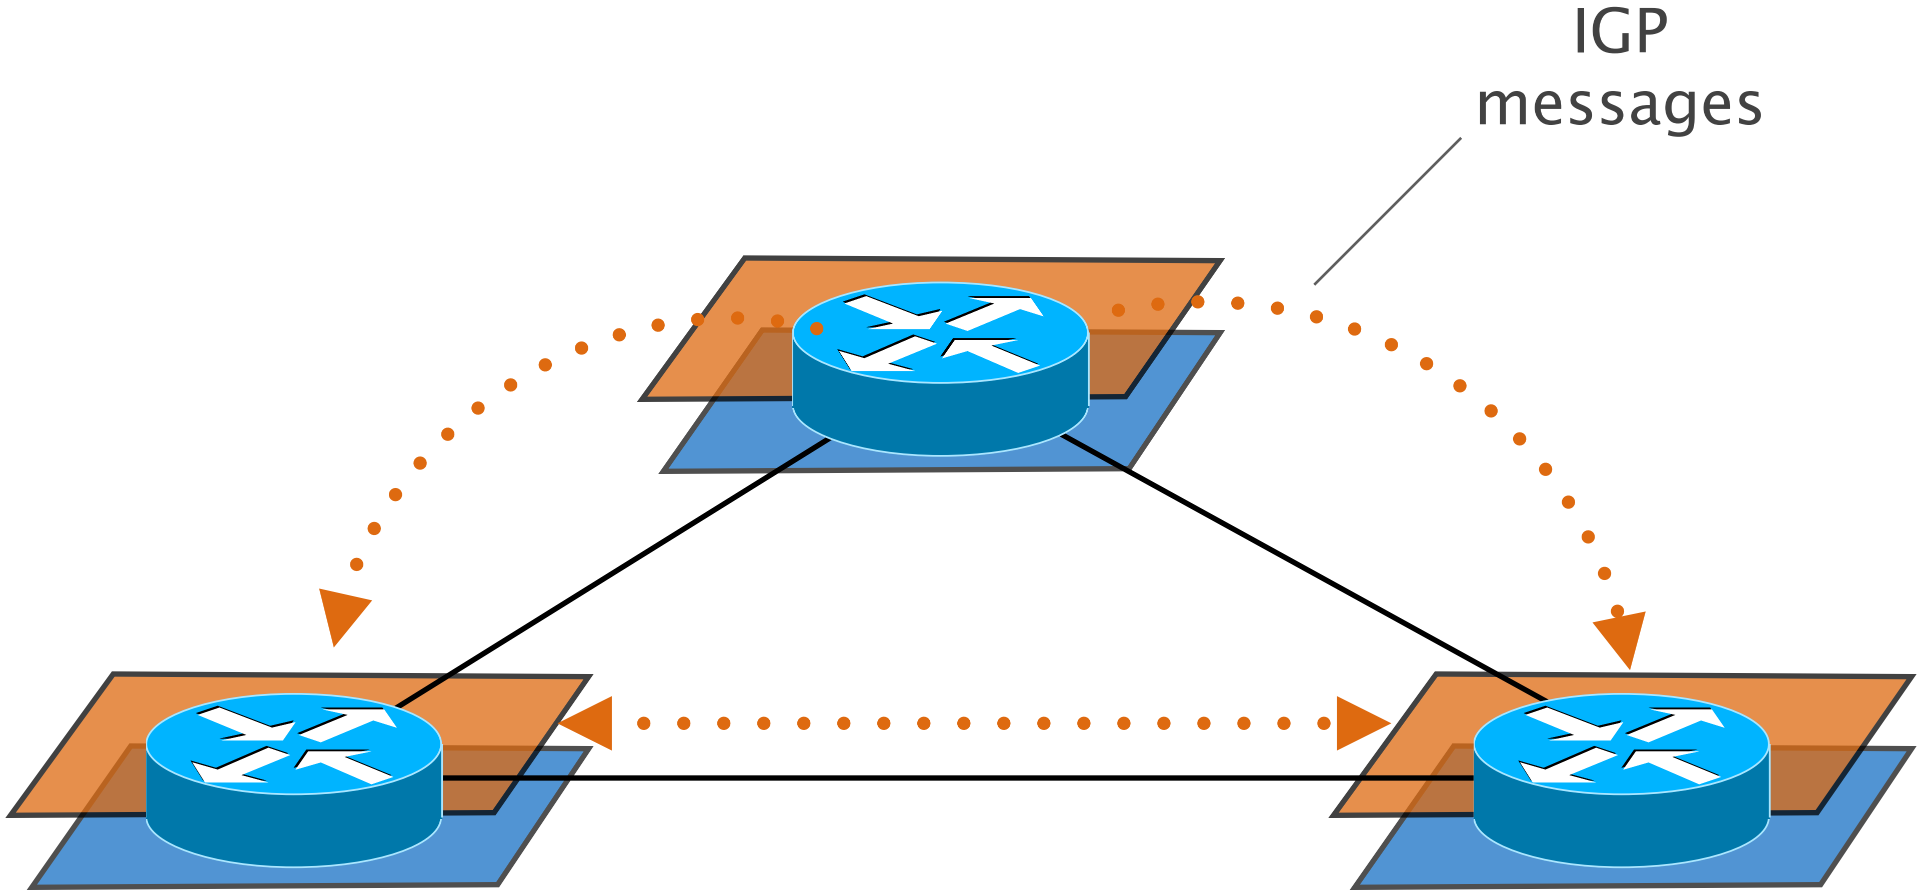

Traditional control-planes are distributed
(for example, IGP protocols like EIGRP, OSPF or IS-IS)

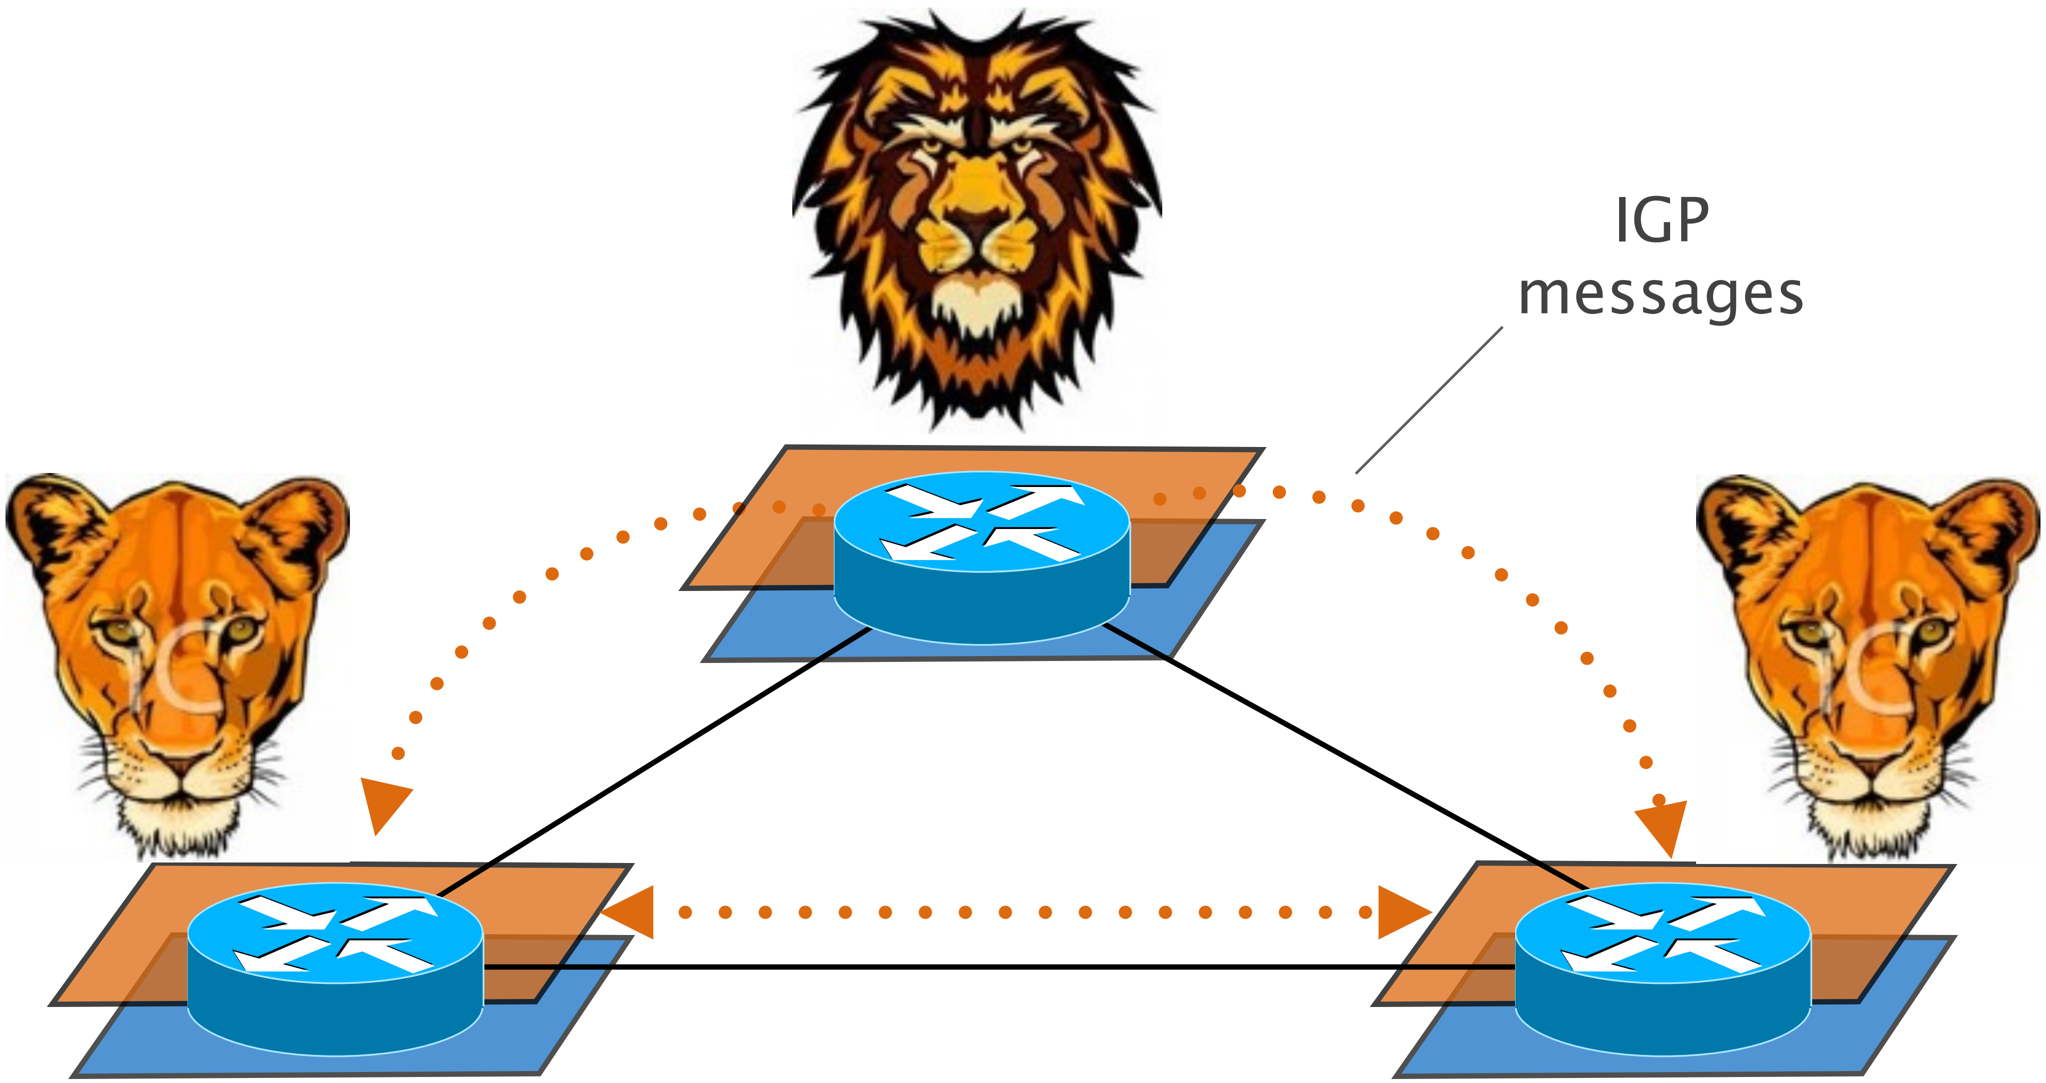

SDN is based on control-plane centralization
(as in basic OpenFlow networks)

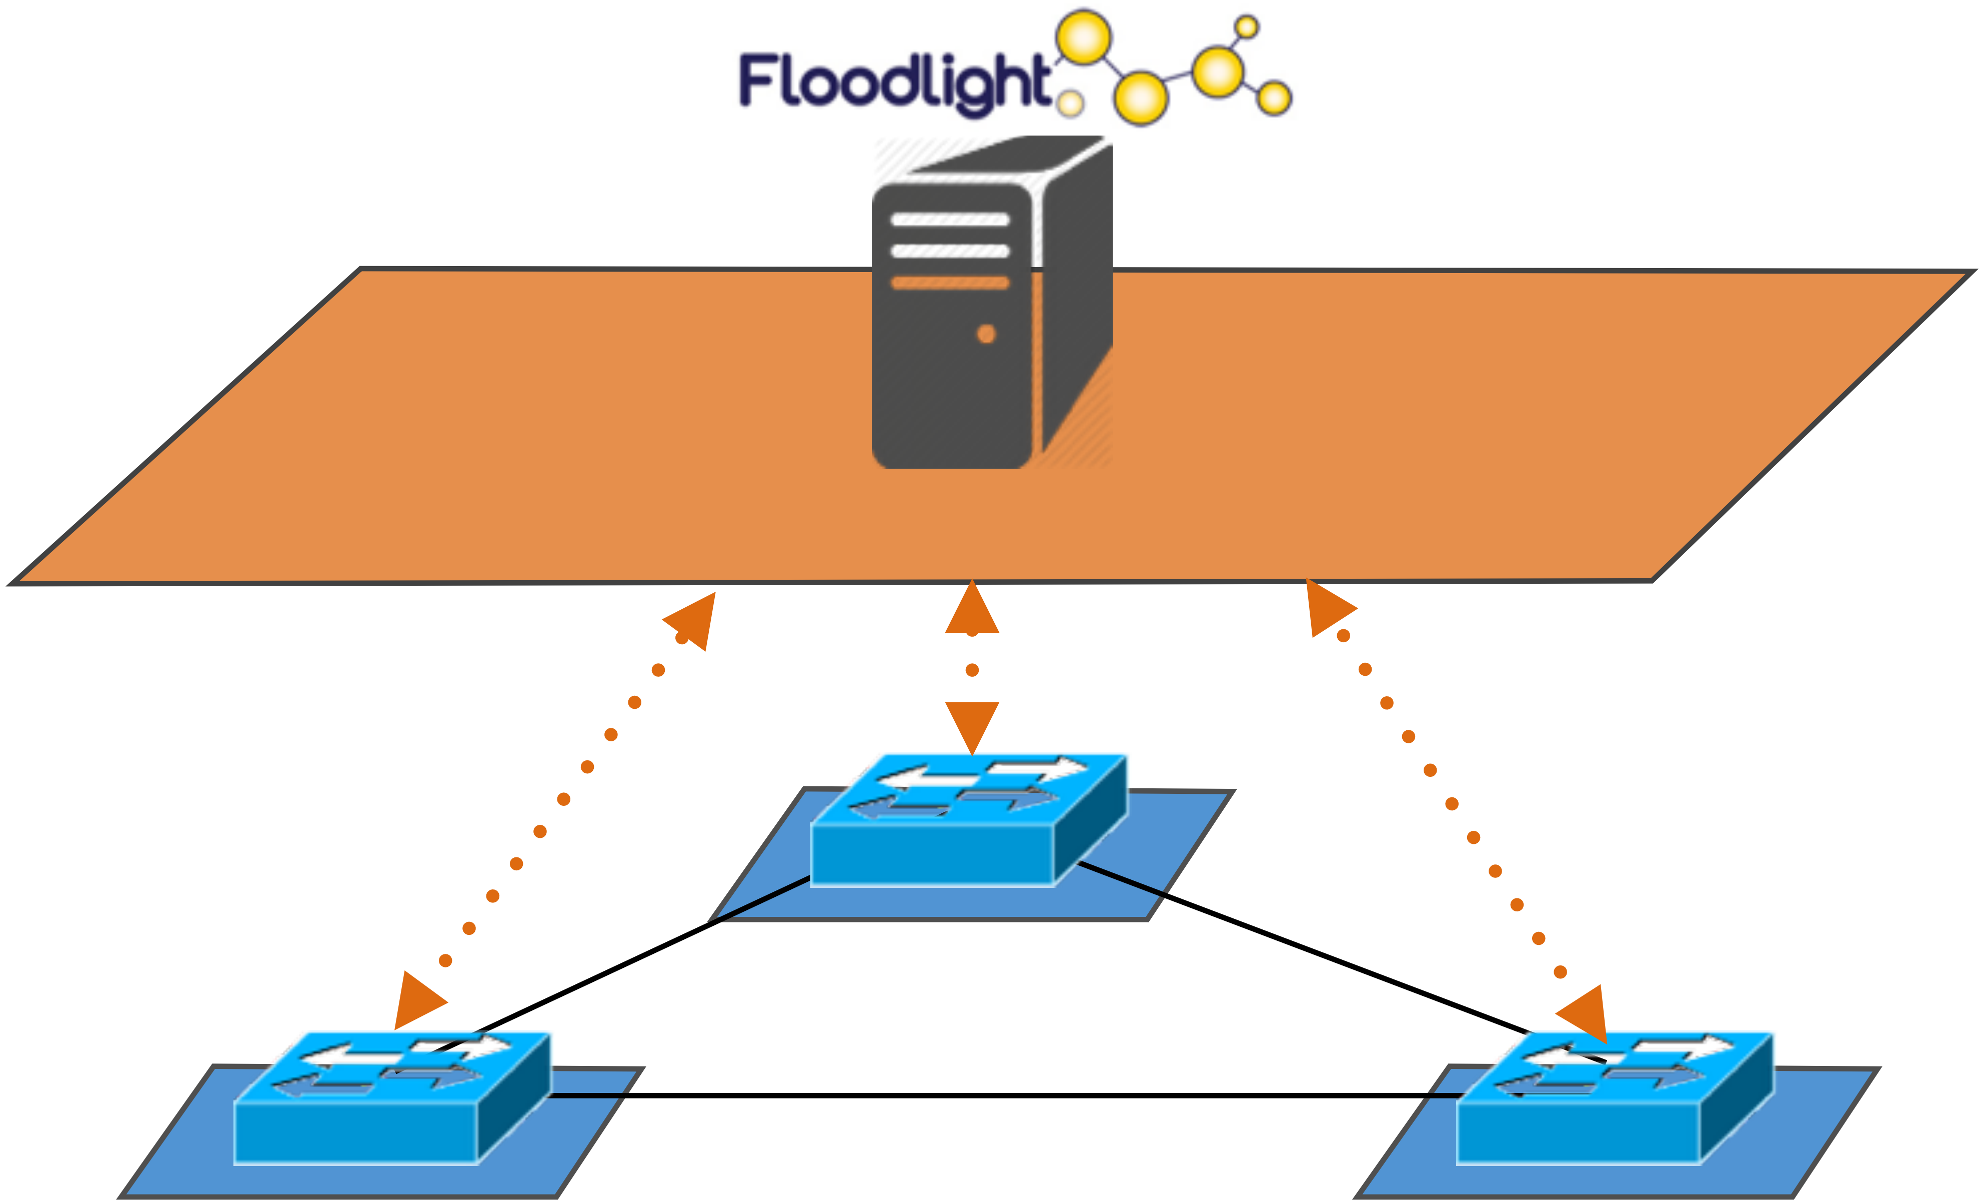

SDN is based on control-plane centralization

(as in basic OpenFlow networks)

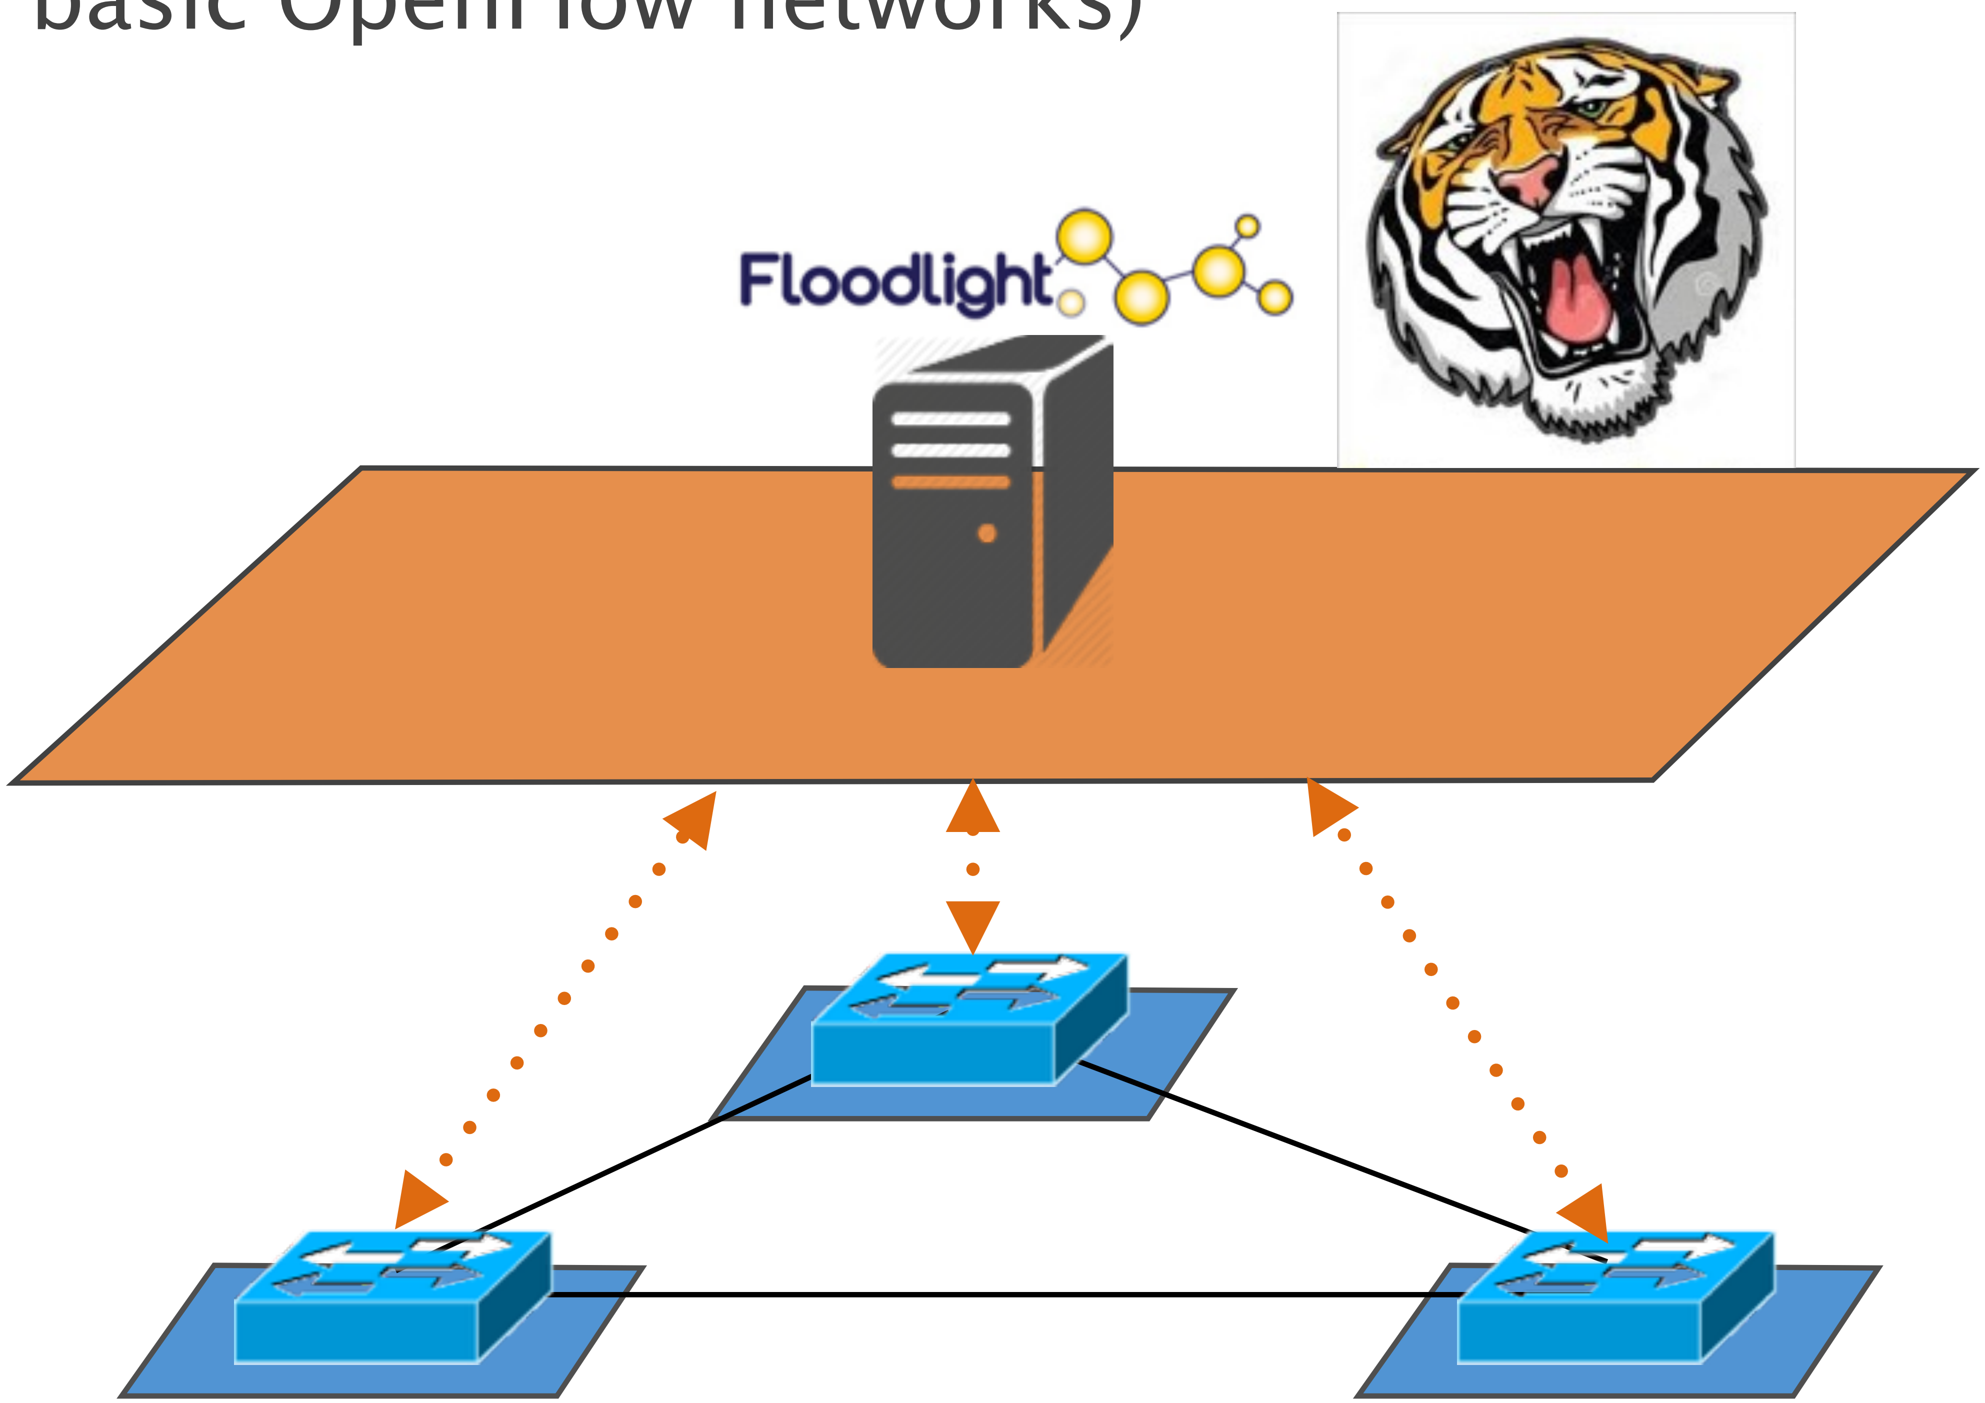

Operators can run ***coexisting control-planes***, that work independently from each other

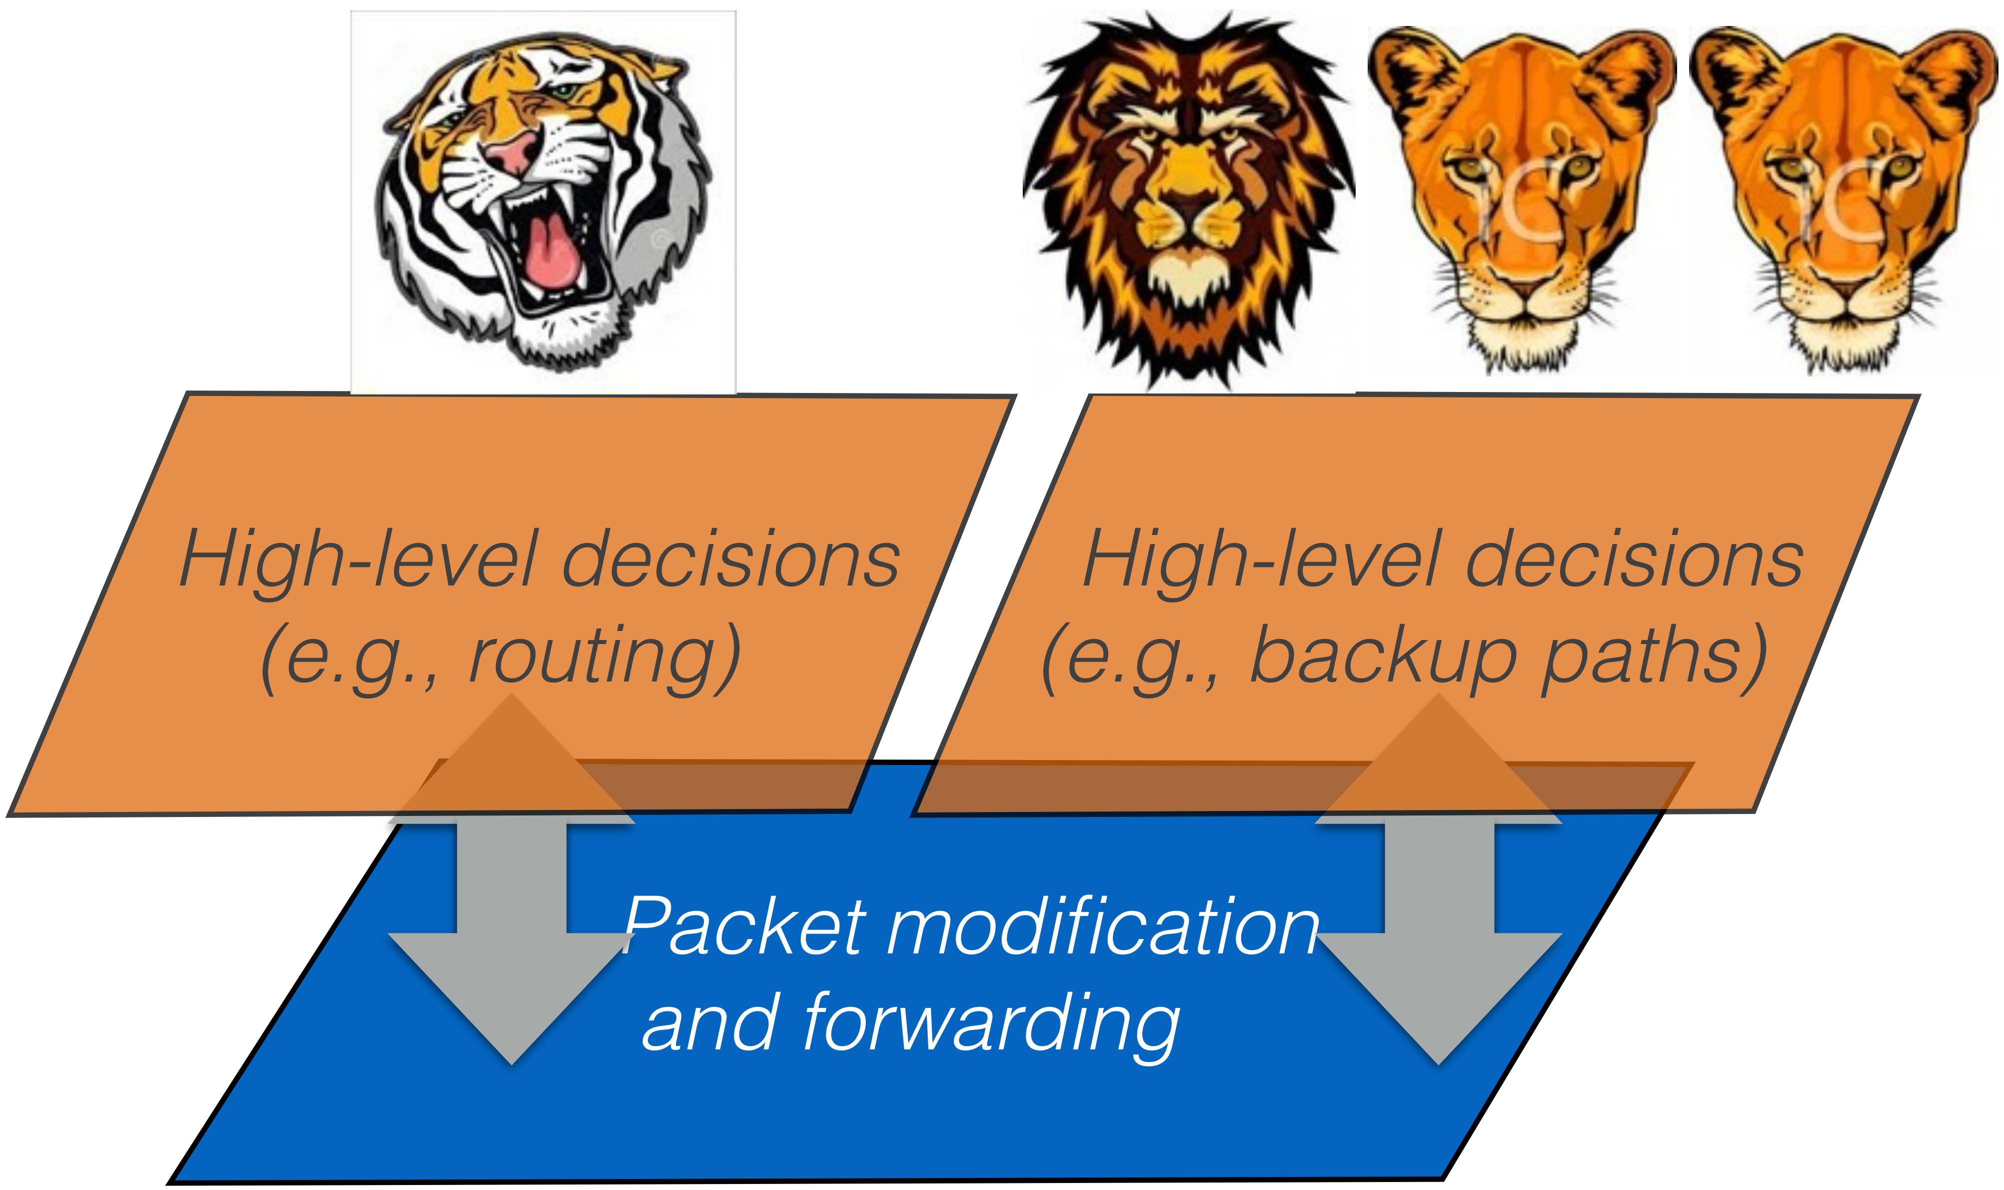

Operators can and ***do run coexisting control-planes***, that work independently from each other

- multiple IGP instances
  e.g., for resilience [Kvalbein06]

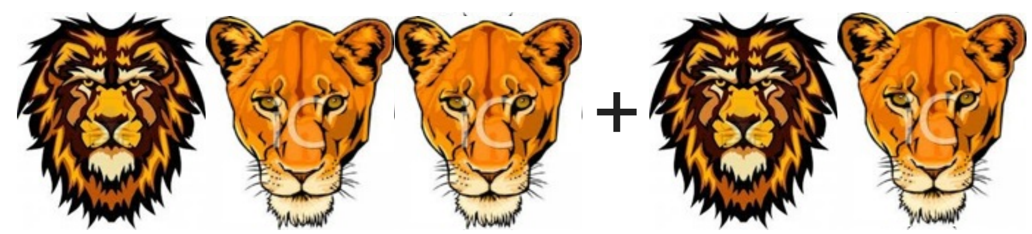

- multiple non-interacting SDN controllers
  e.g., task-specialized [Canini13]

- hybrid SDN networks
  e.g., for TE [Agarwal13]
  or robustness [Tilmans14]

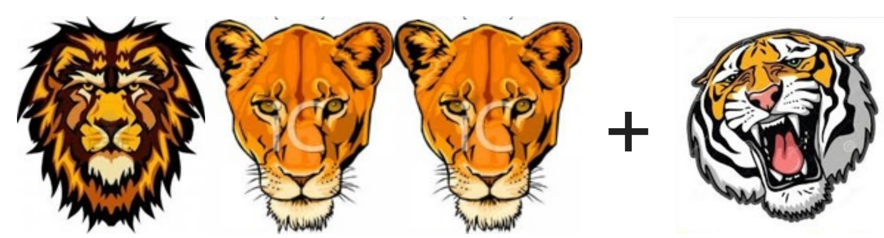

Unfortunately, control-plane coexistence can cause disruptions

Unfortunately, we don't know when and which coexisting disruptions occur and coordination is needed

- guidelines for multiple **link-state IGP** instances
  e.g., [Le08]
- theory and guidelines for **IGP control-plane interaction**
  e.g., [Le07,Le10]
- architectures to **coordinate** multiple SDN controllers
  e.g., [Canini13]

We developed a **general theory** to study disruptions due to control-planes coexistence

- **any** combination of control-planes
  existing and future
- **many** network settings
  multiple IGPs, multi-controller SDN, hybrid SDN
- both **static** and **dynamic** scenarios
  configuration guidelines and safe reconfigurations

Our contributions include modeling, formal analysis, and insight of the implications

- model for arbitrary control-planes
- characterization of coexistence anomalies
- practical applications of our theory
- analysis of the lessons learned

Our contributions include modeling, formal analysis, and insight of the implications

- model for arbitrary control-planes
- characterization of coexistence anomalies
- practical applications of our theory
- analysis of the lessons learned

The most generic router model include one FIB and multiple RIBs

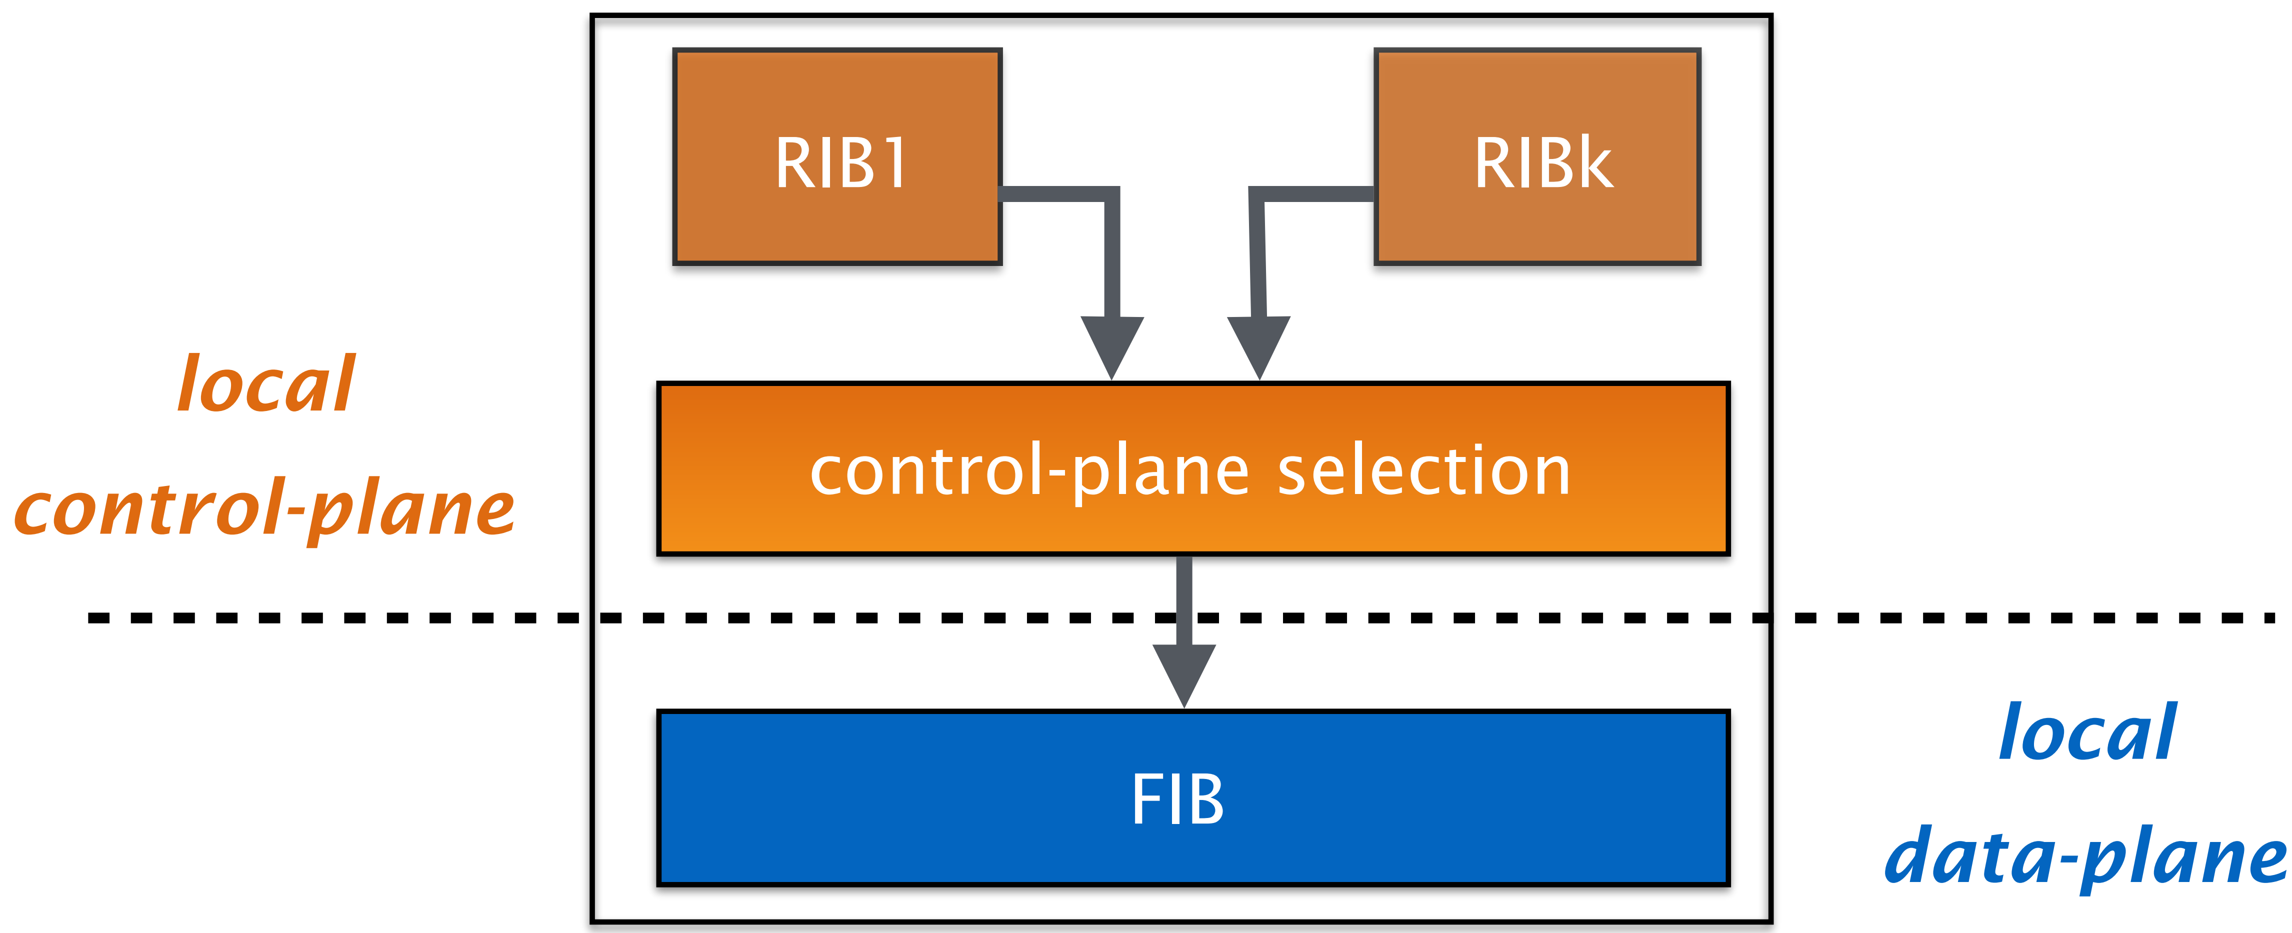

Control-planes can be classified according to their input and their output

Input (where they ***read*** from)	Output (where they ***write*** to)
▪ FIB: Fib-Aware (FA)	▪ FIB: preemptive
▪ RIB or other: Fib-Unaware (FU)	▪ RIB: non-preemptive

A **control-plane taxonomy** can be built upon their input / output properties

	Control-plane	Properties
SDN	OpenFlow*, ForCES	preemptive, FU
	static routes, RCP, I2RS	non-preemptive FU
IGP	OSPF, IS-IS	non-preemptive, FU
	RIP, EIGRP	non-preemptive, FA
future	BGP as IGP	non-preemptive, FU
	…	…

Our taxonomy is **general**

(covers distributed and centralized control-planes)

	Control-plane	Properties
SDN	OpenFlow*, ForCES	preemptive, FU
	static routes, RCP, I2RS	non-preemptive FU
IGP	OSPF, IS-IS	non-preemptive, FU
	RIP, EIGRP	non-preemptive, FA
future	BGP as IGP	non-preemptive, FU
	…	…

Our taxonomy is **novel**
(orthogonal to traditional classifications)

	Control-plane	Properties
SDN	OpenFlow*, ForCES	preemptive, FU
	static routes, RCP, I2RS	non-preemptive FU
IGP	OSPF, IS-IS	non-preemptive, FU
	RIP, EIGRP	non-preemptive, FA
future	BGP as IGP	non-preemptive, FU
	…	…

Our taxonomy is **exhaustive**
(enabling modeling of future control-planes)

	Control-plane	Properties
SDN	OpenFlow*, ForCES	preemptive, FU
	static routes, RCP, I2RS	non-preemptive FU
IGP	OSPF, IS-IS	non-preemptive, FU
	RIP, EIGRP	non-preemptive, FA
future	BGP as IGP	non-preemptive, FU
	…	…

Our contributions include modeling, formal analysis, and insight of the implications

- model for arbitrary control-planes
- characterization of coexistence anomalies
- practical applications of our theory
- analysis of the lessons learned

Coexisting disruptions depend on the **class** of the running control-planes

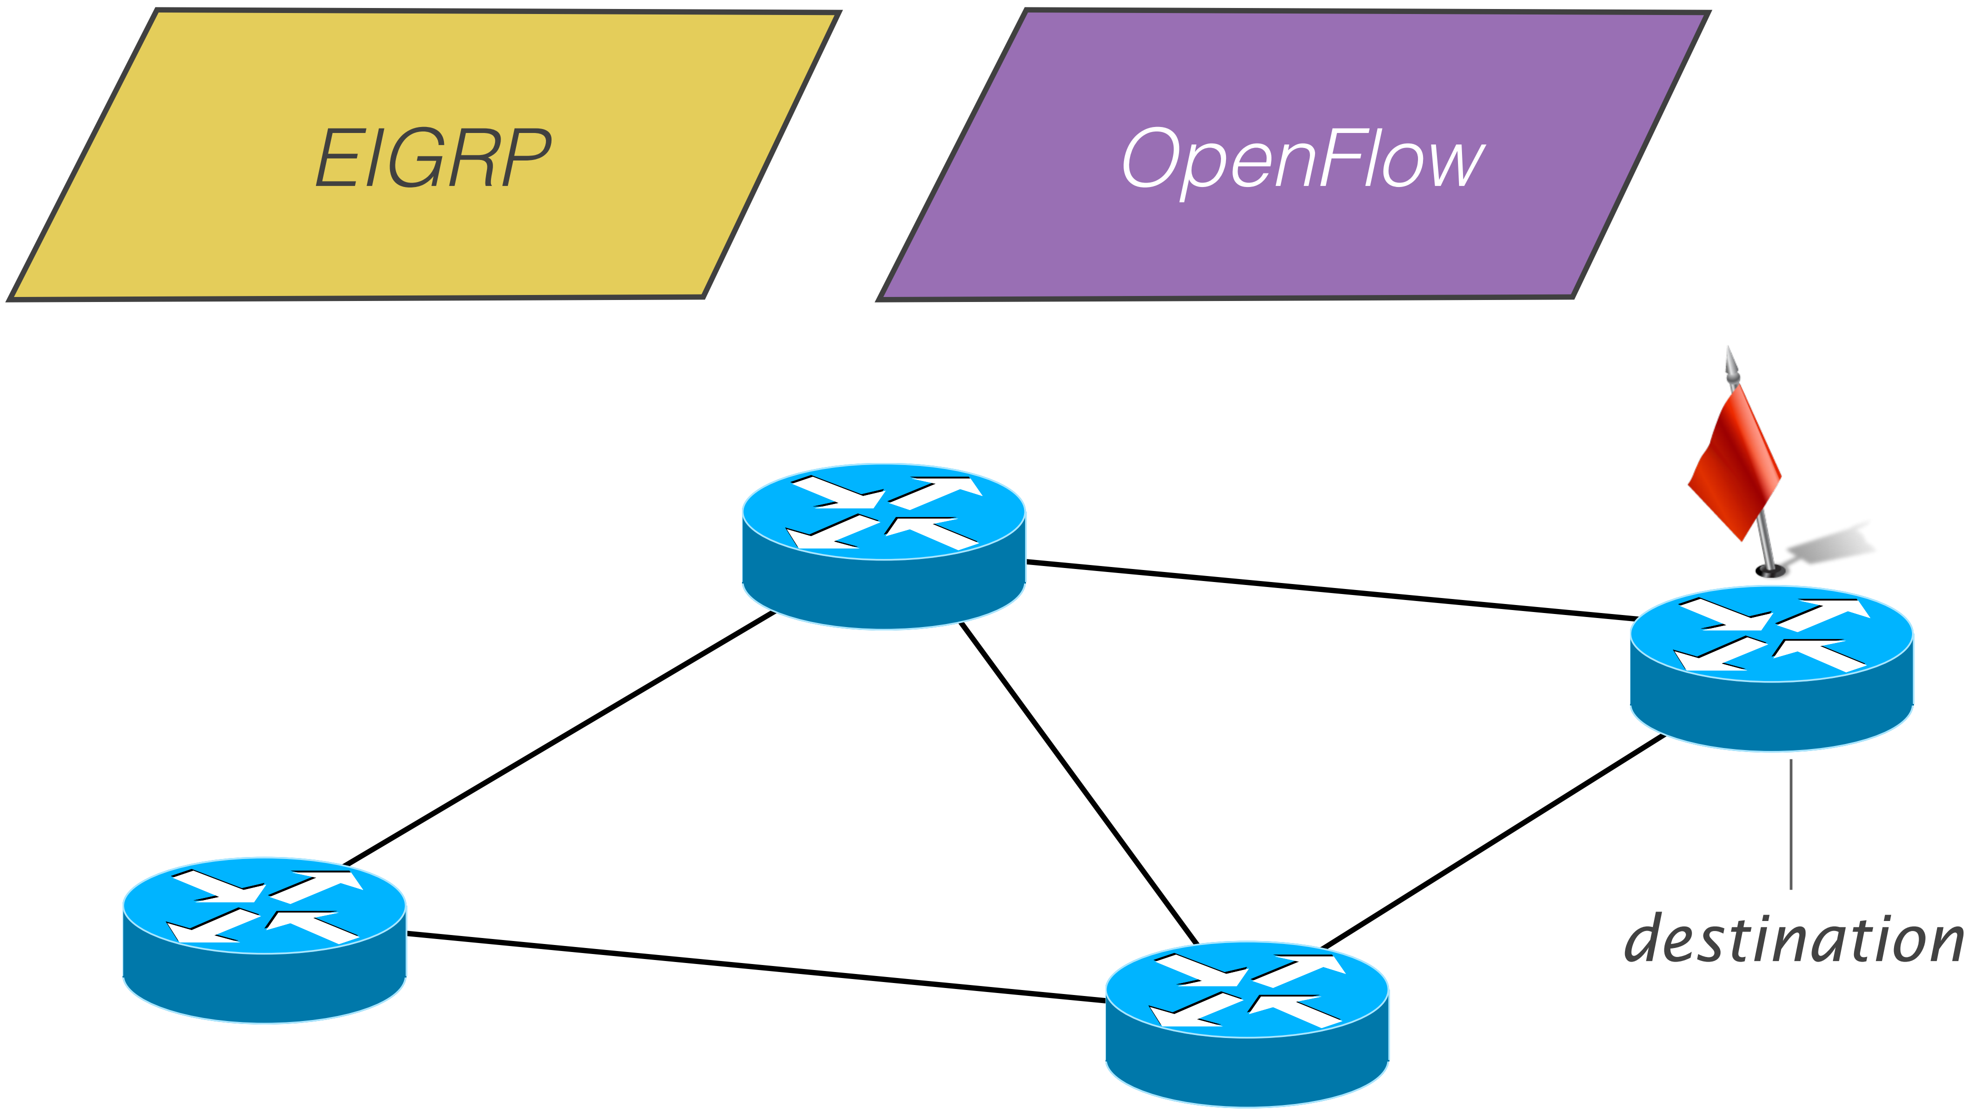

Each router has a FIB and per-destination control-plane preferences

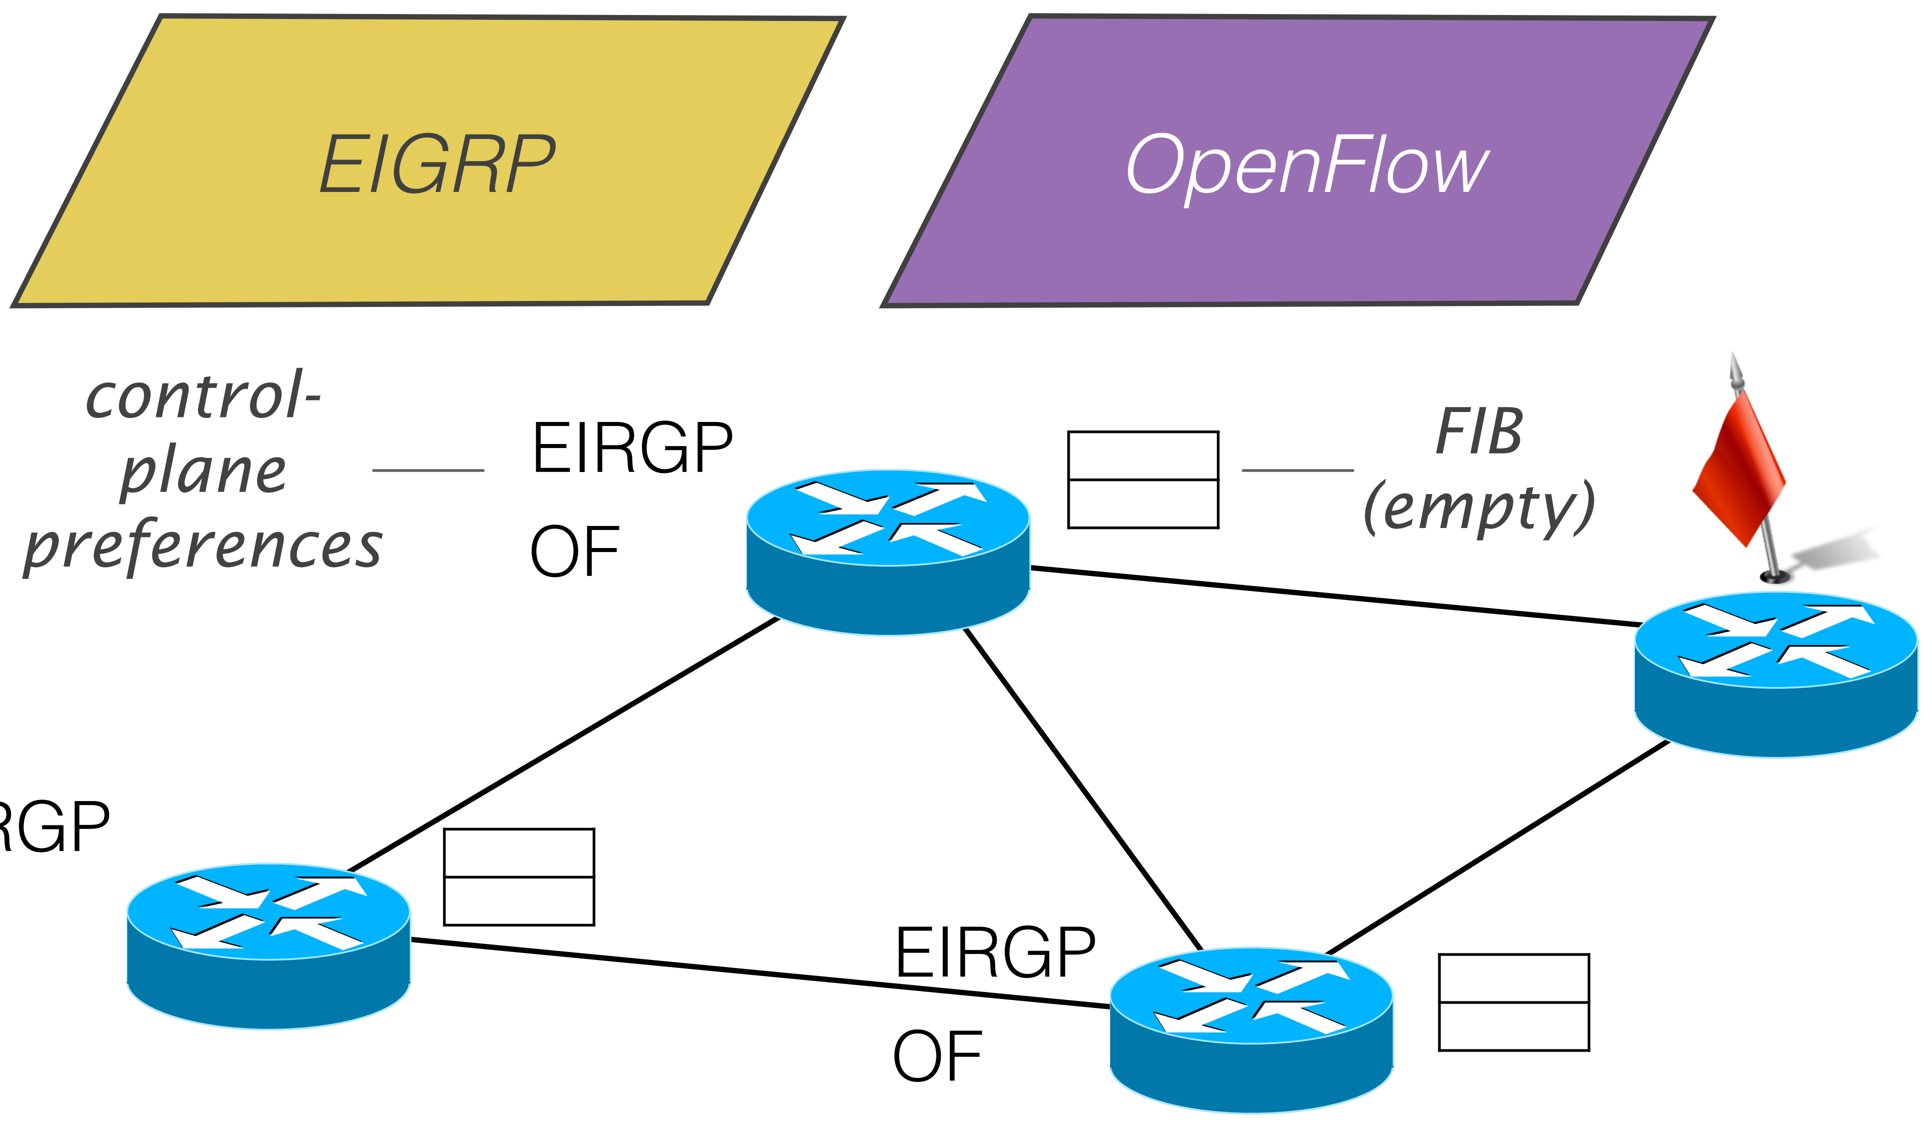

Some routers may prefer one control-plane, e.g., OpenFlow in the example

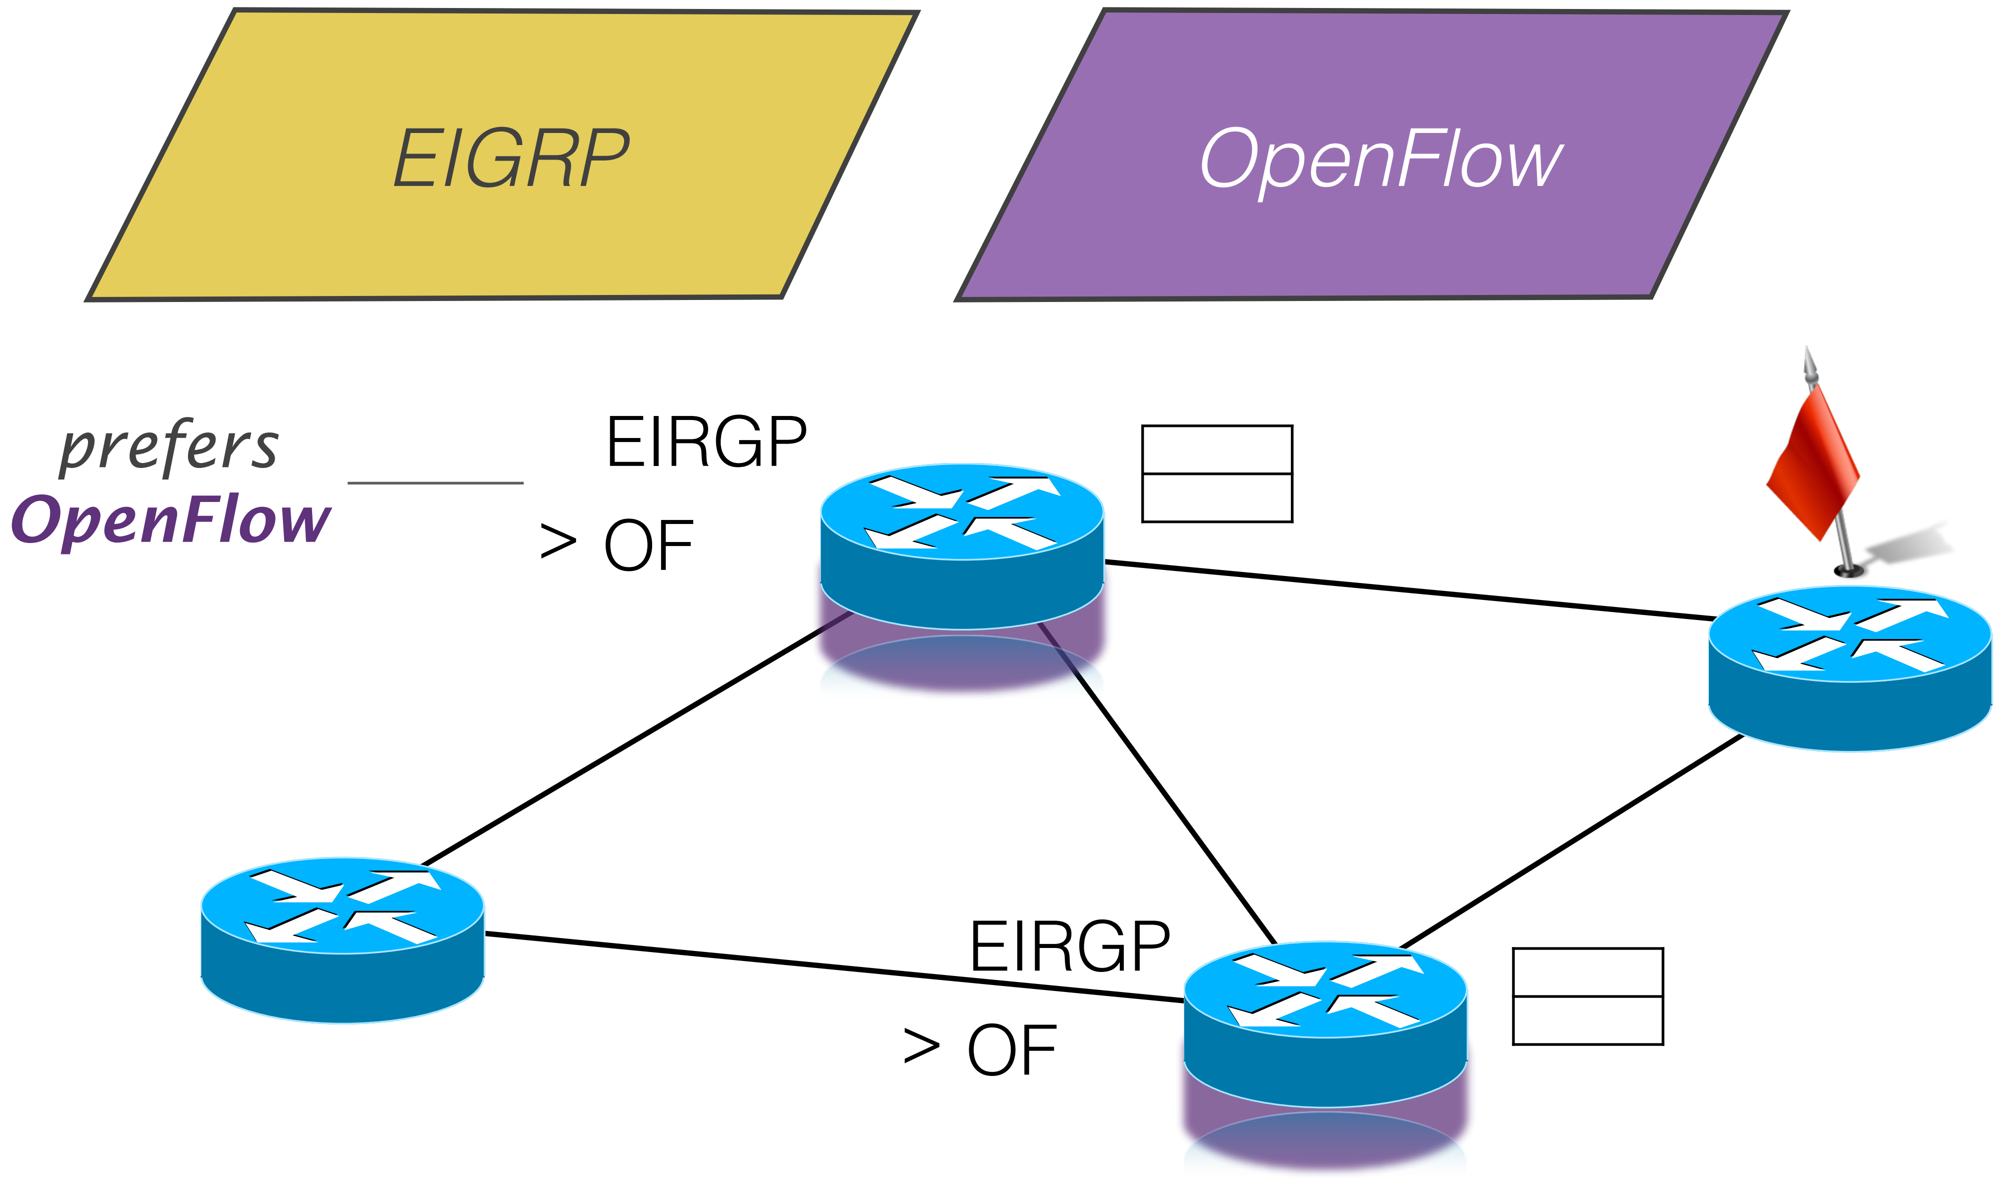

Being preemptive, OpenFlow directly writes to the FIBs of the routers

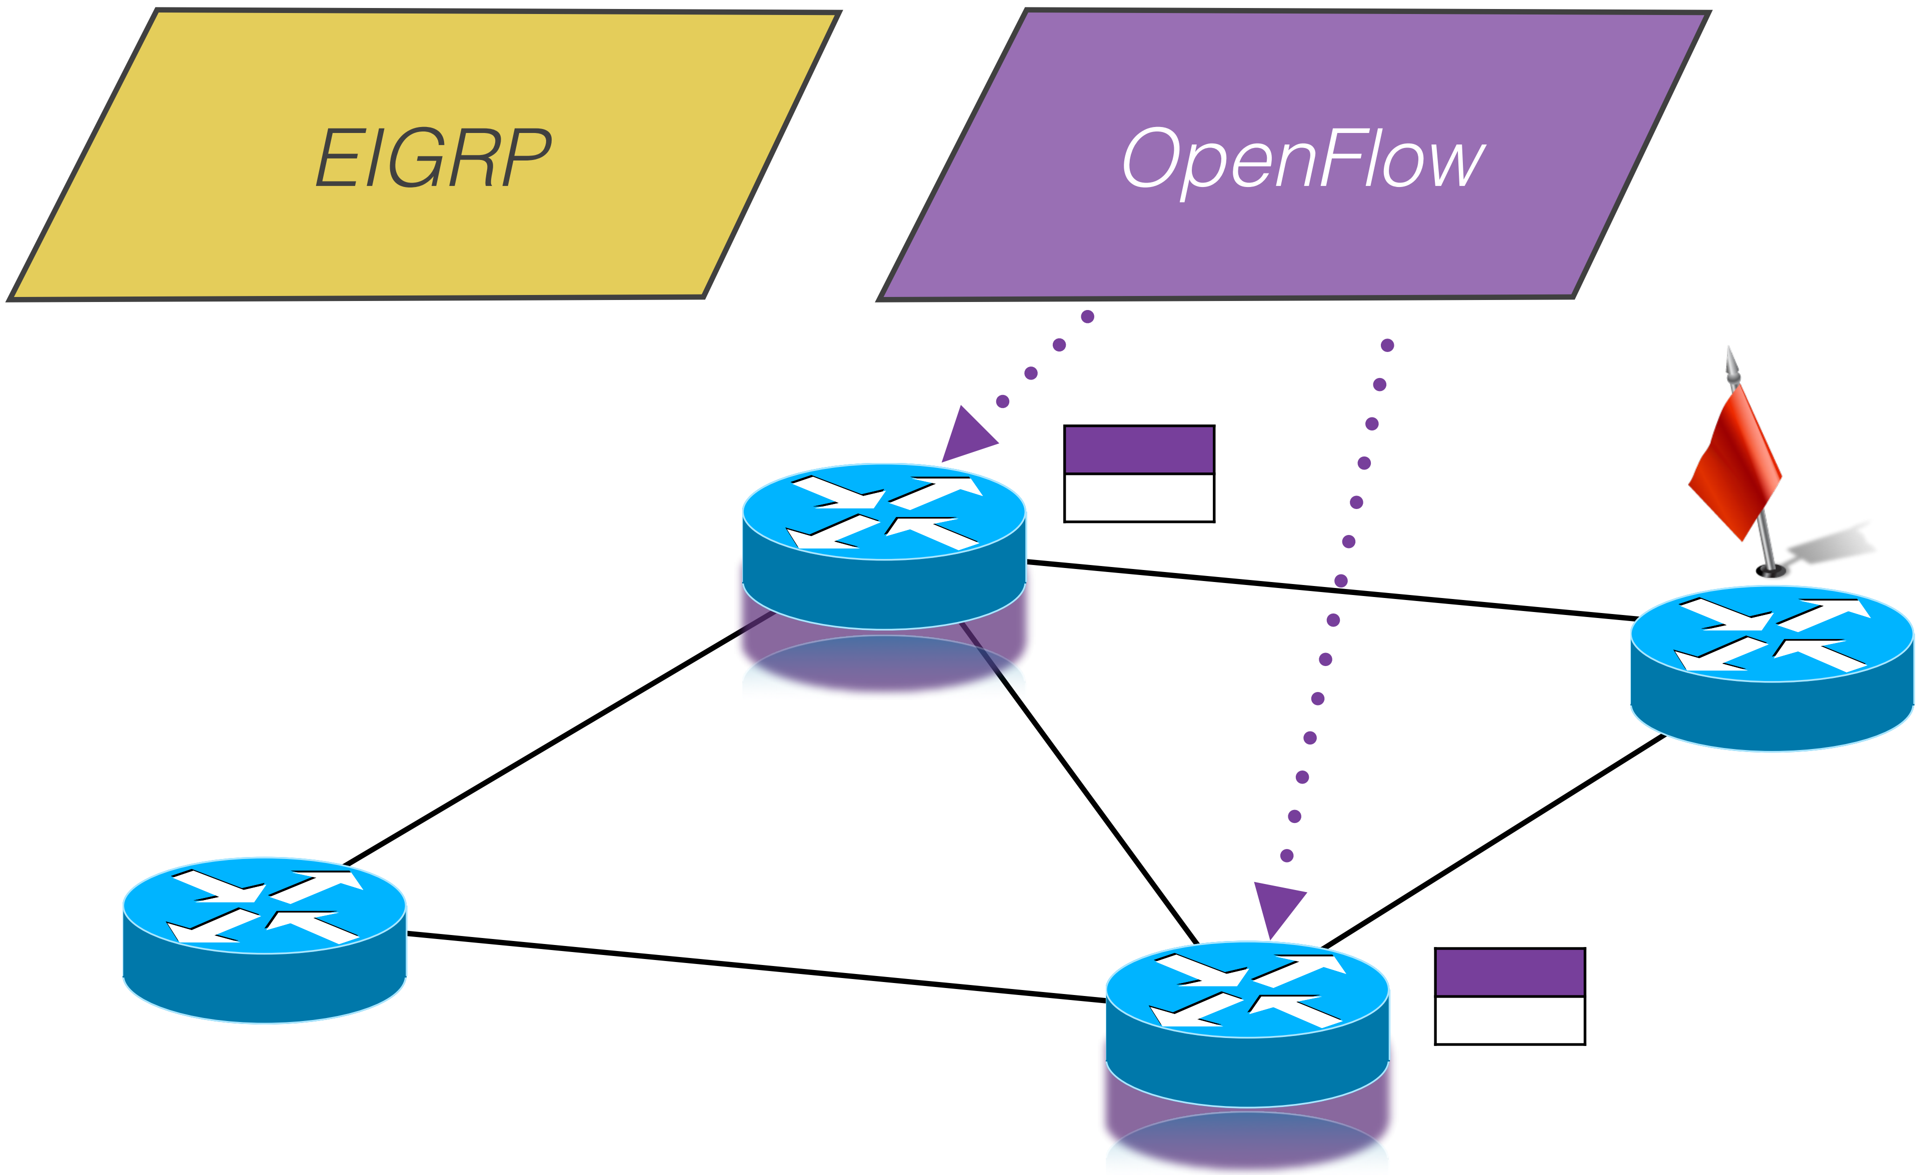

Other routers may prefer another control-plane, e.g., EIGRP in the example

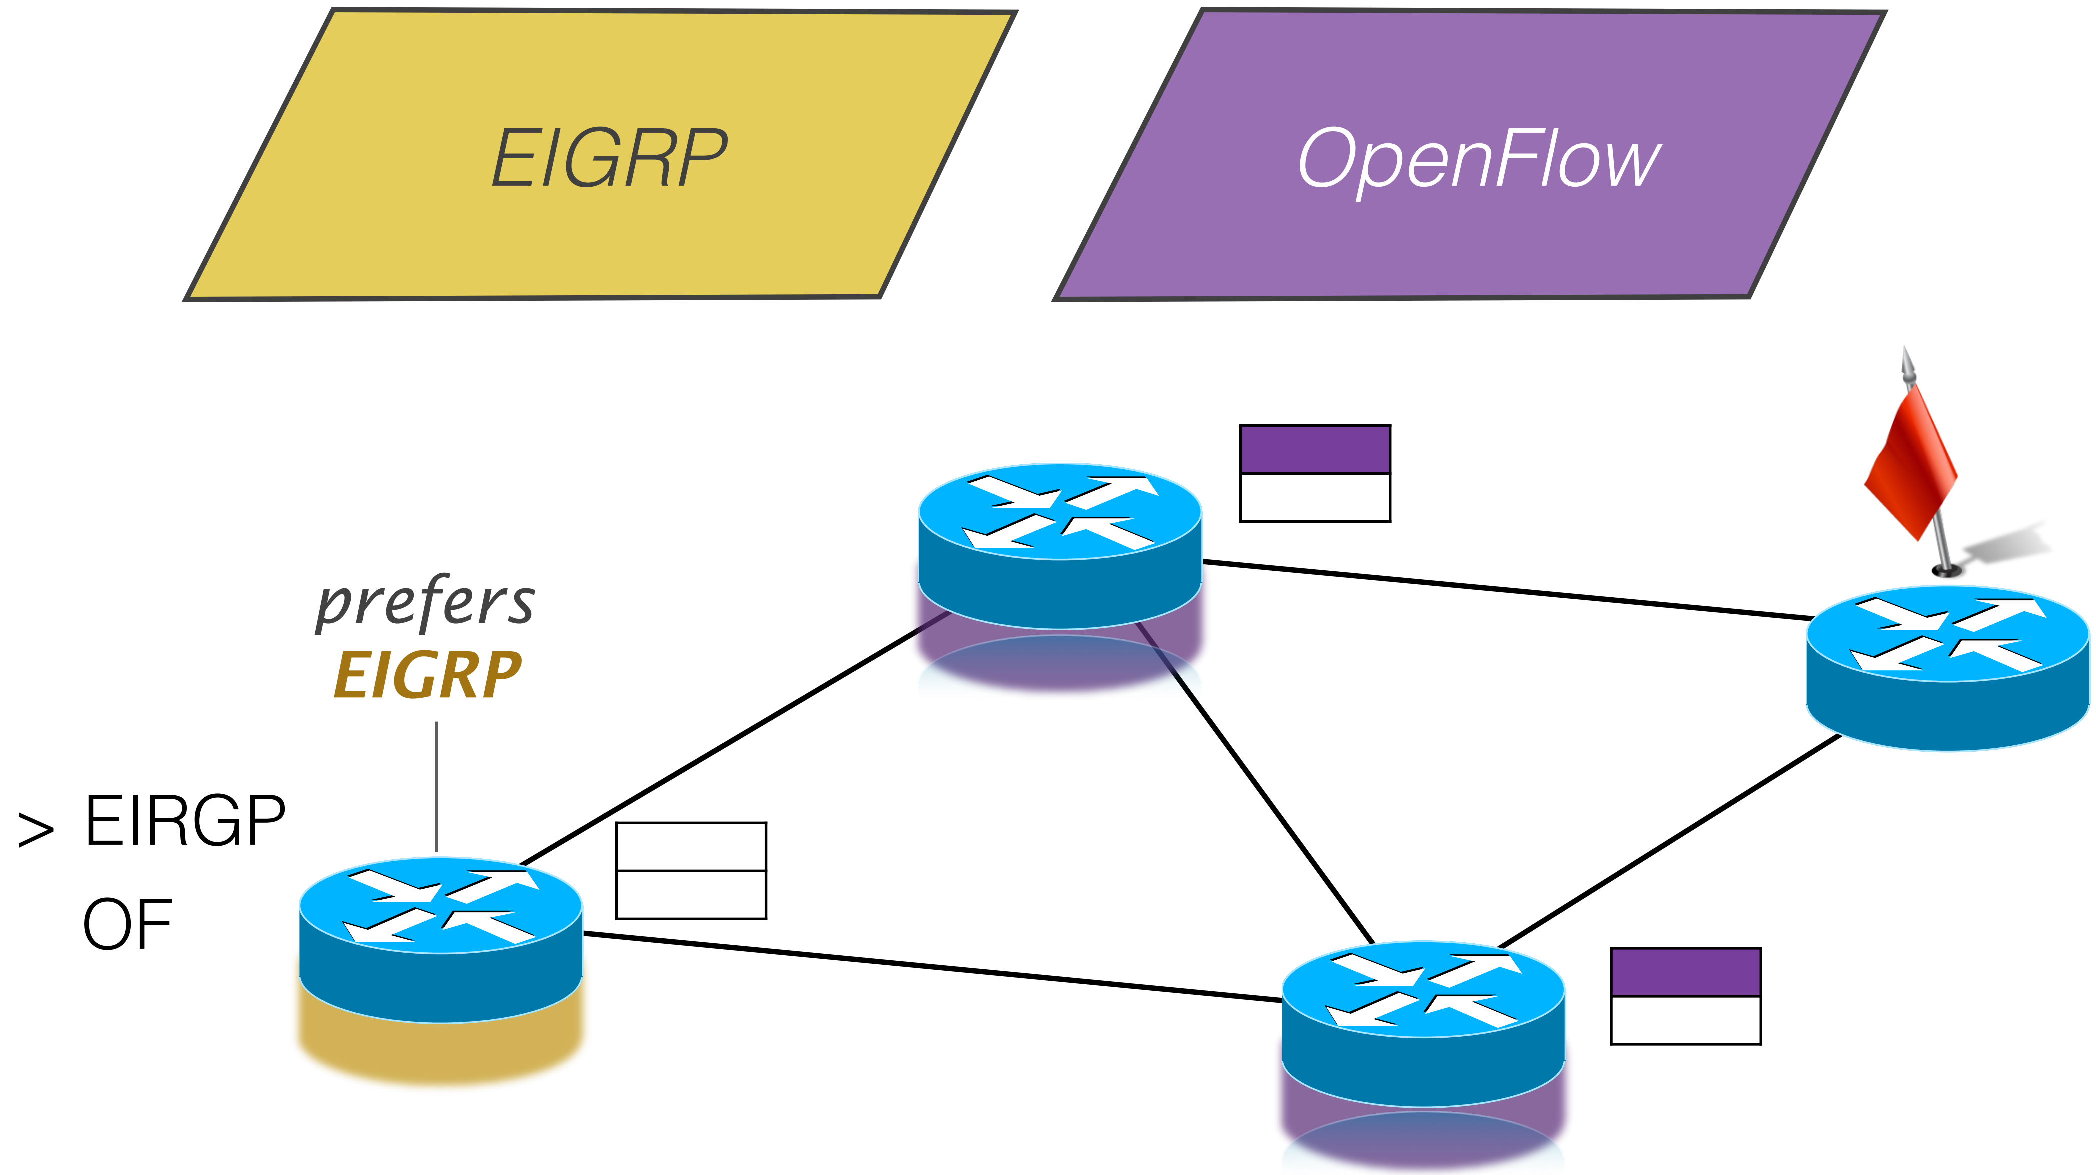

Being FA, EIGRP reads from routers' FIBs trying to build EIRGP-only paths

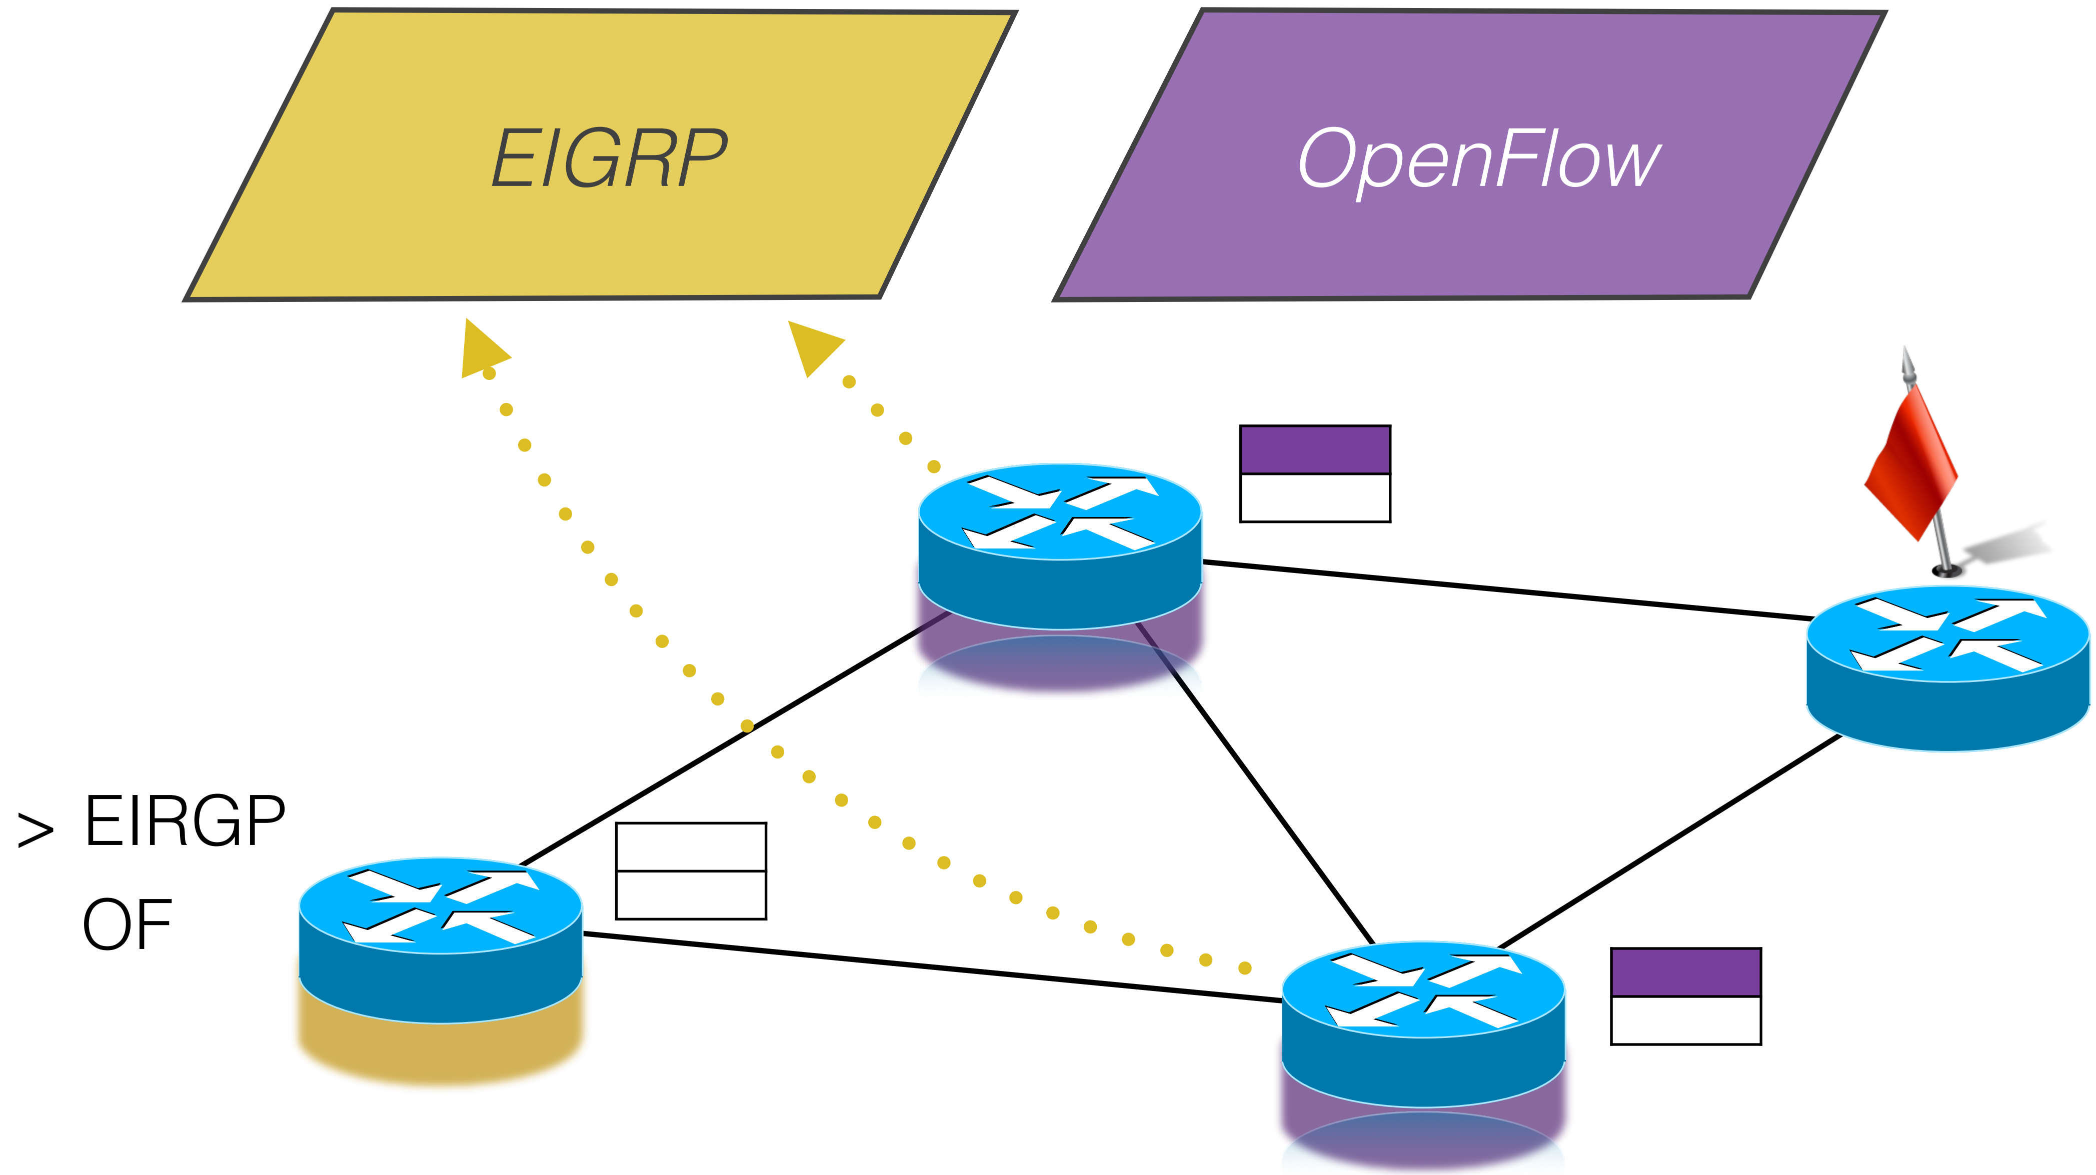

Since EIGRP-only paths cannot be built, EIGRP does not write in the FIB

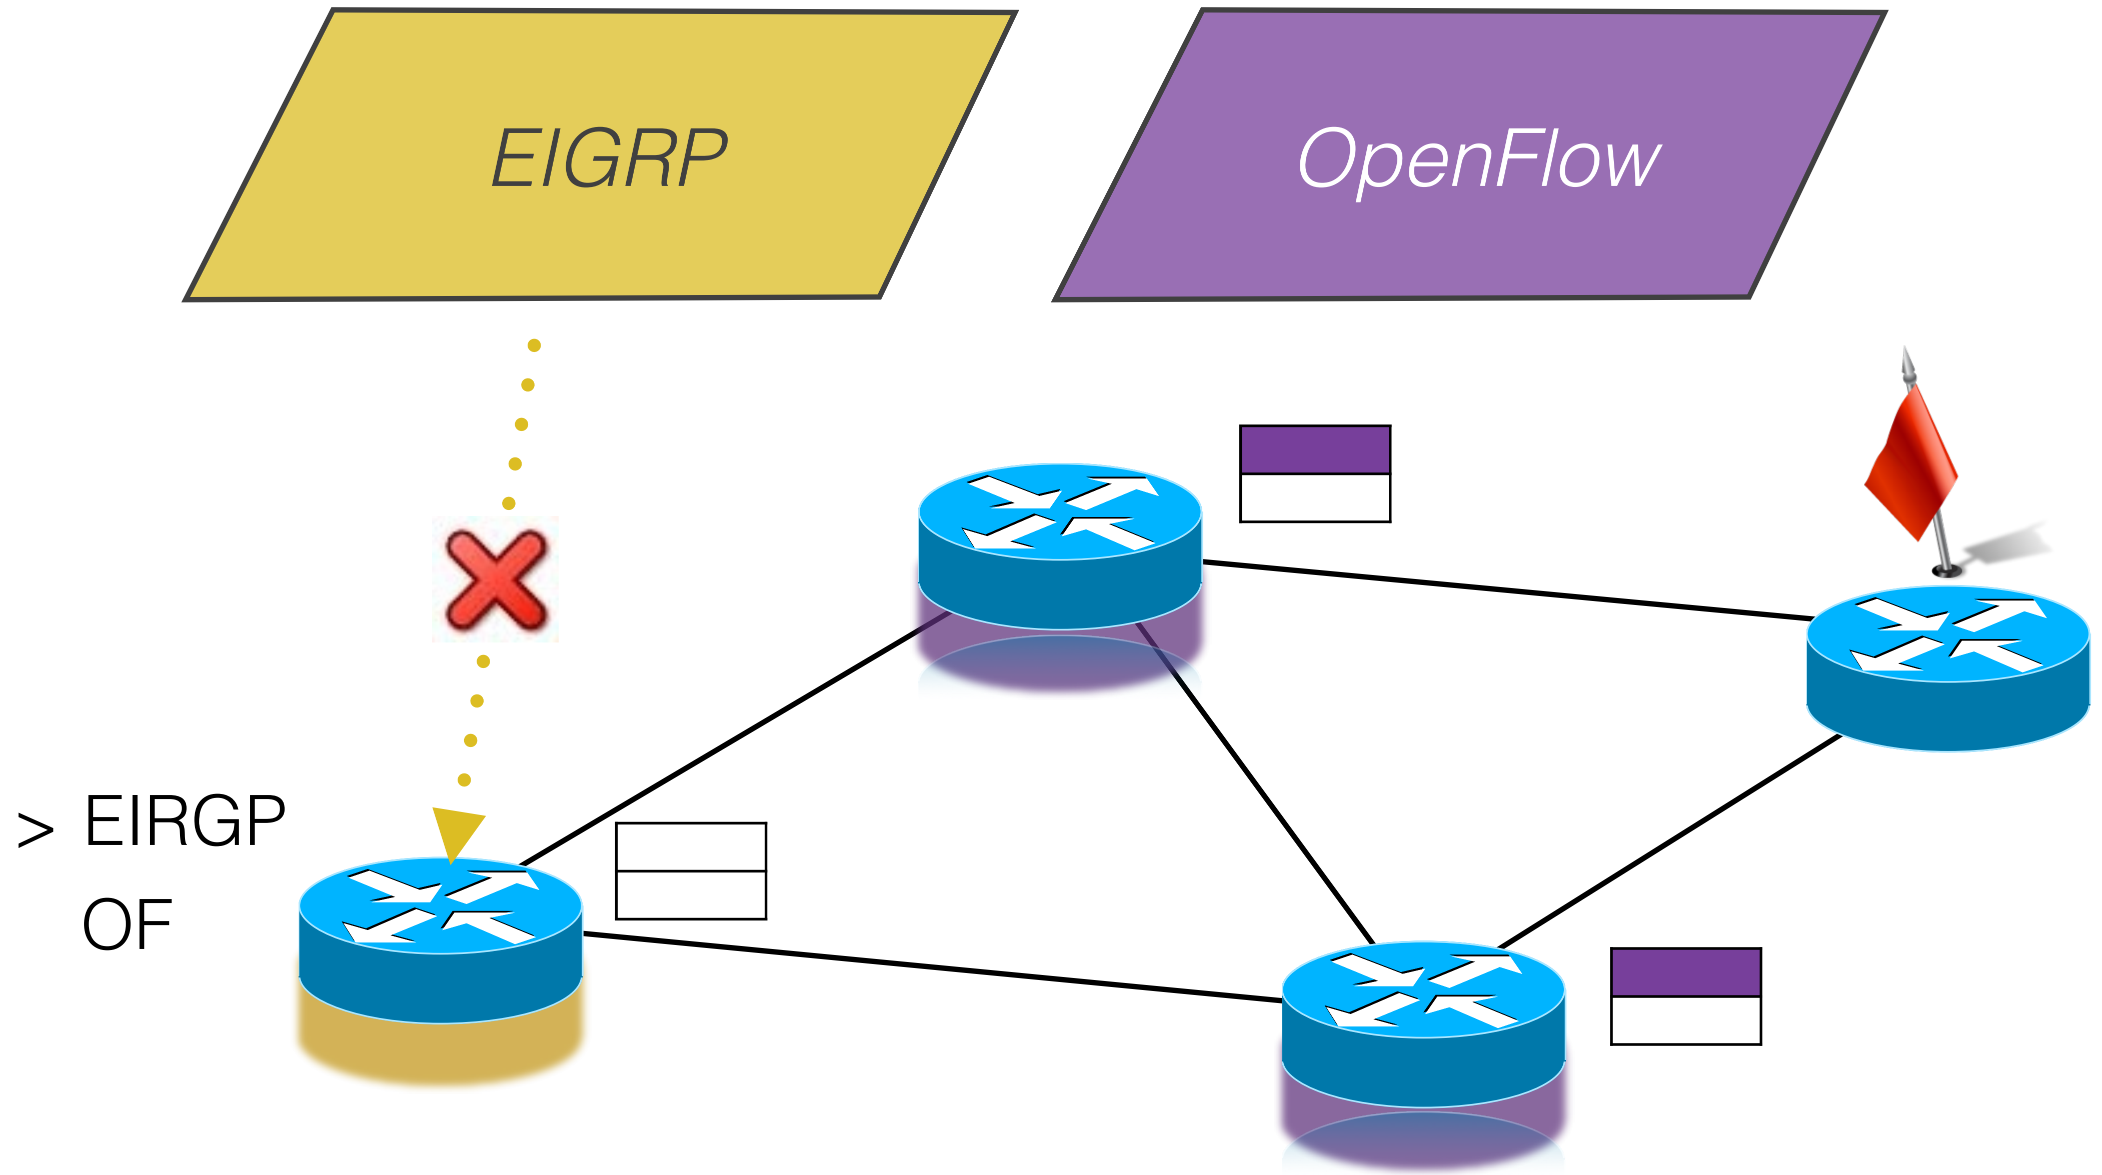

OpenFlow does not write to the FIB,
since it is not the most preferred control-plane

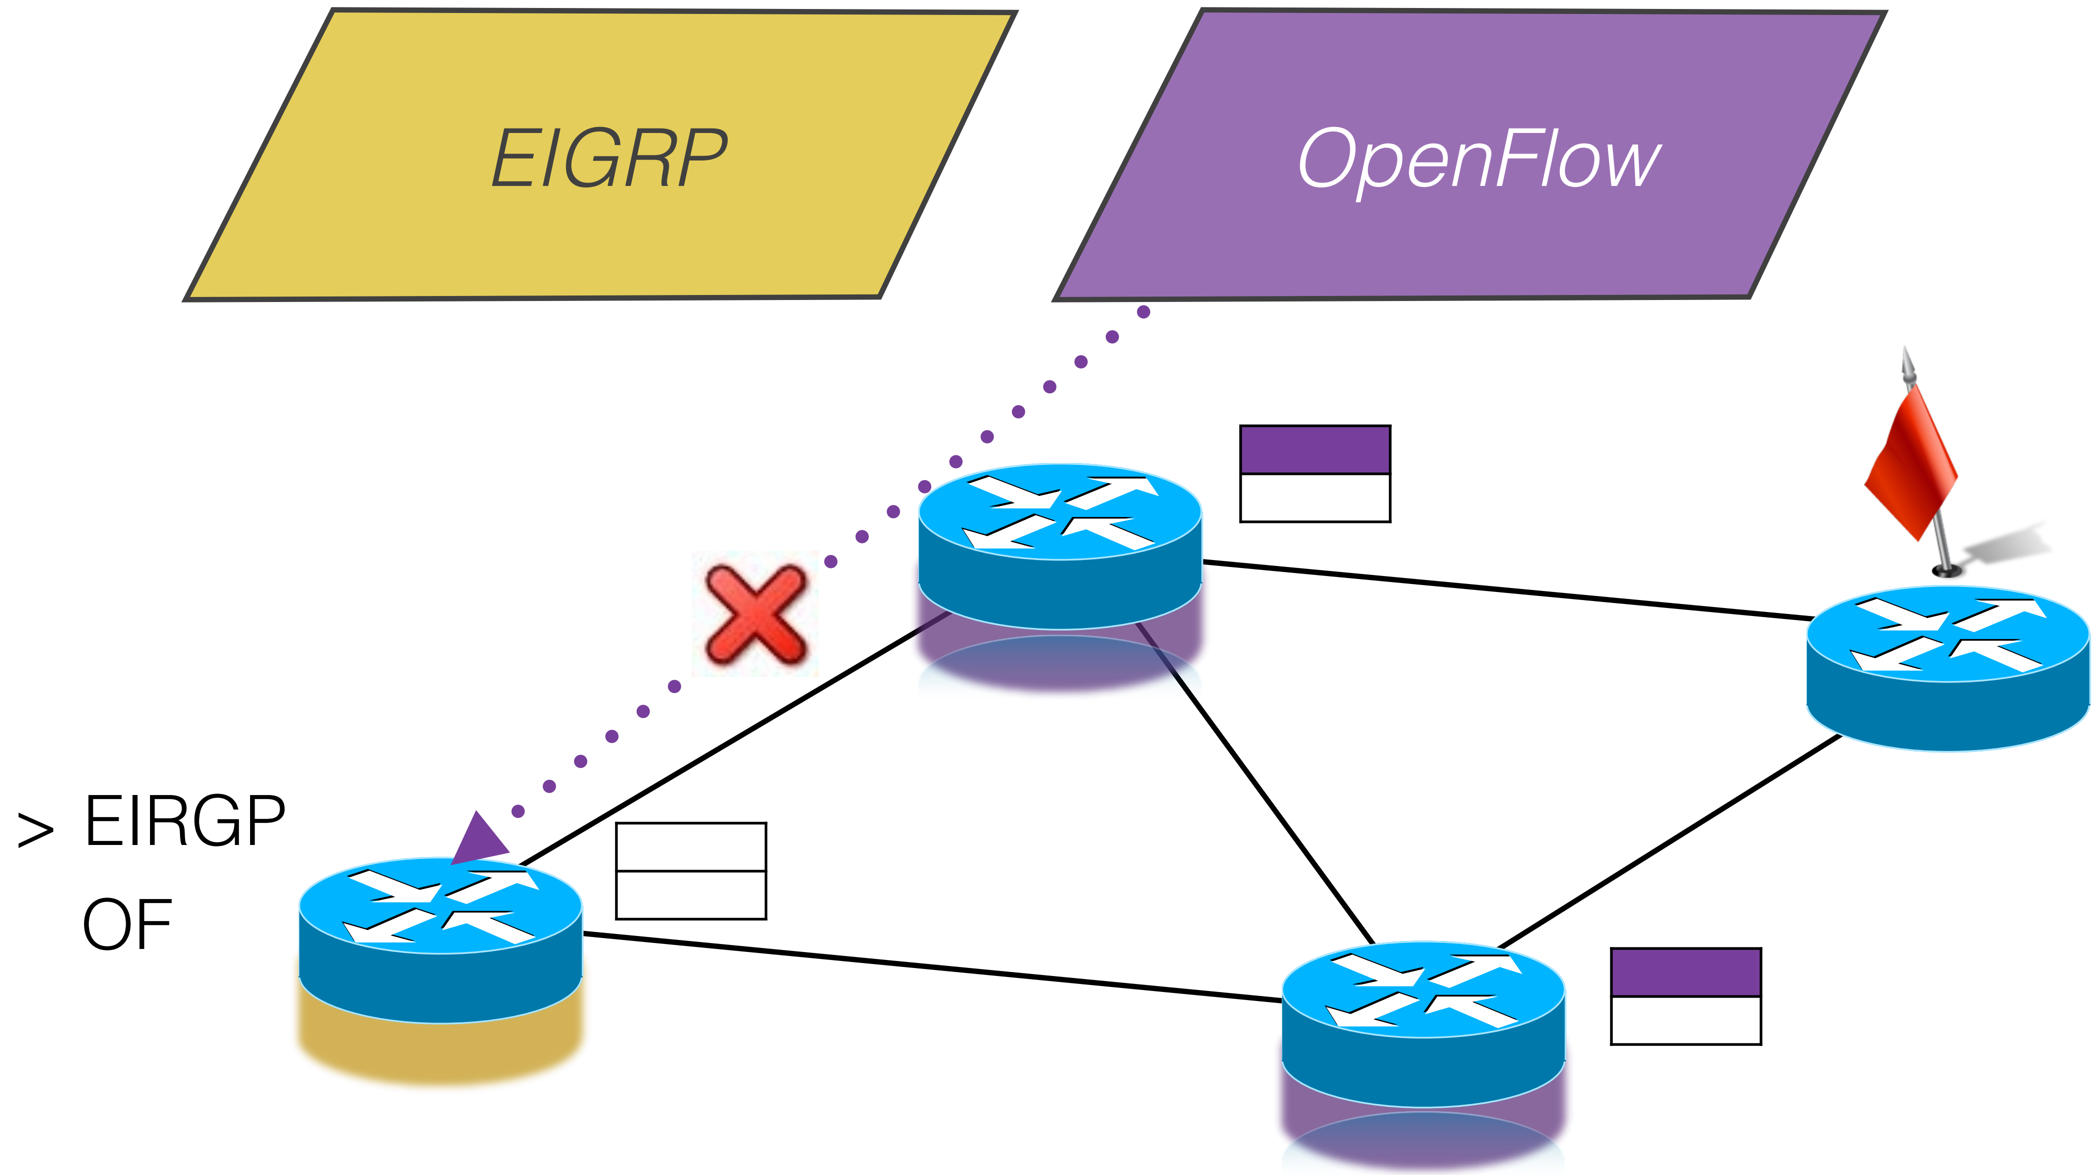

Thus, the EIGRP-preferring router has no FIB entry for the destination, which creates a blackhole!

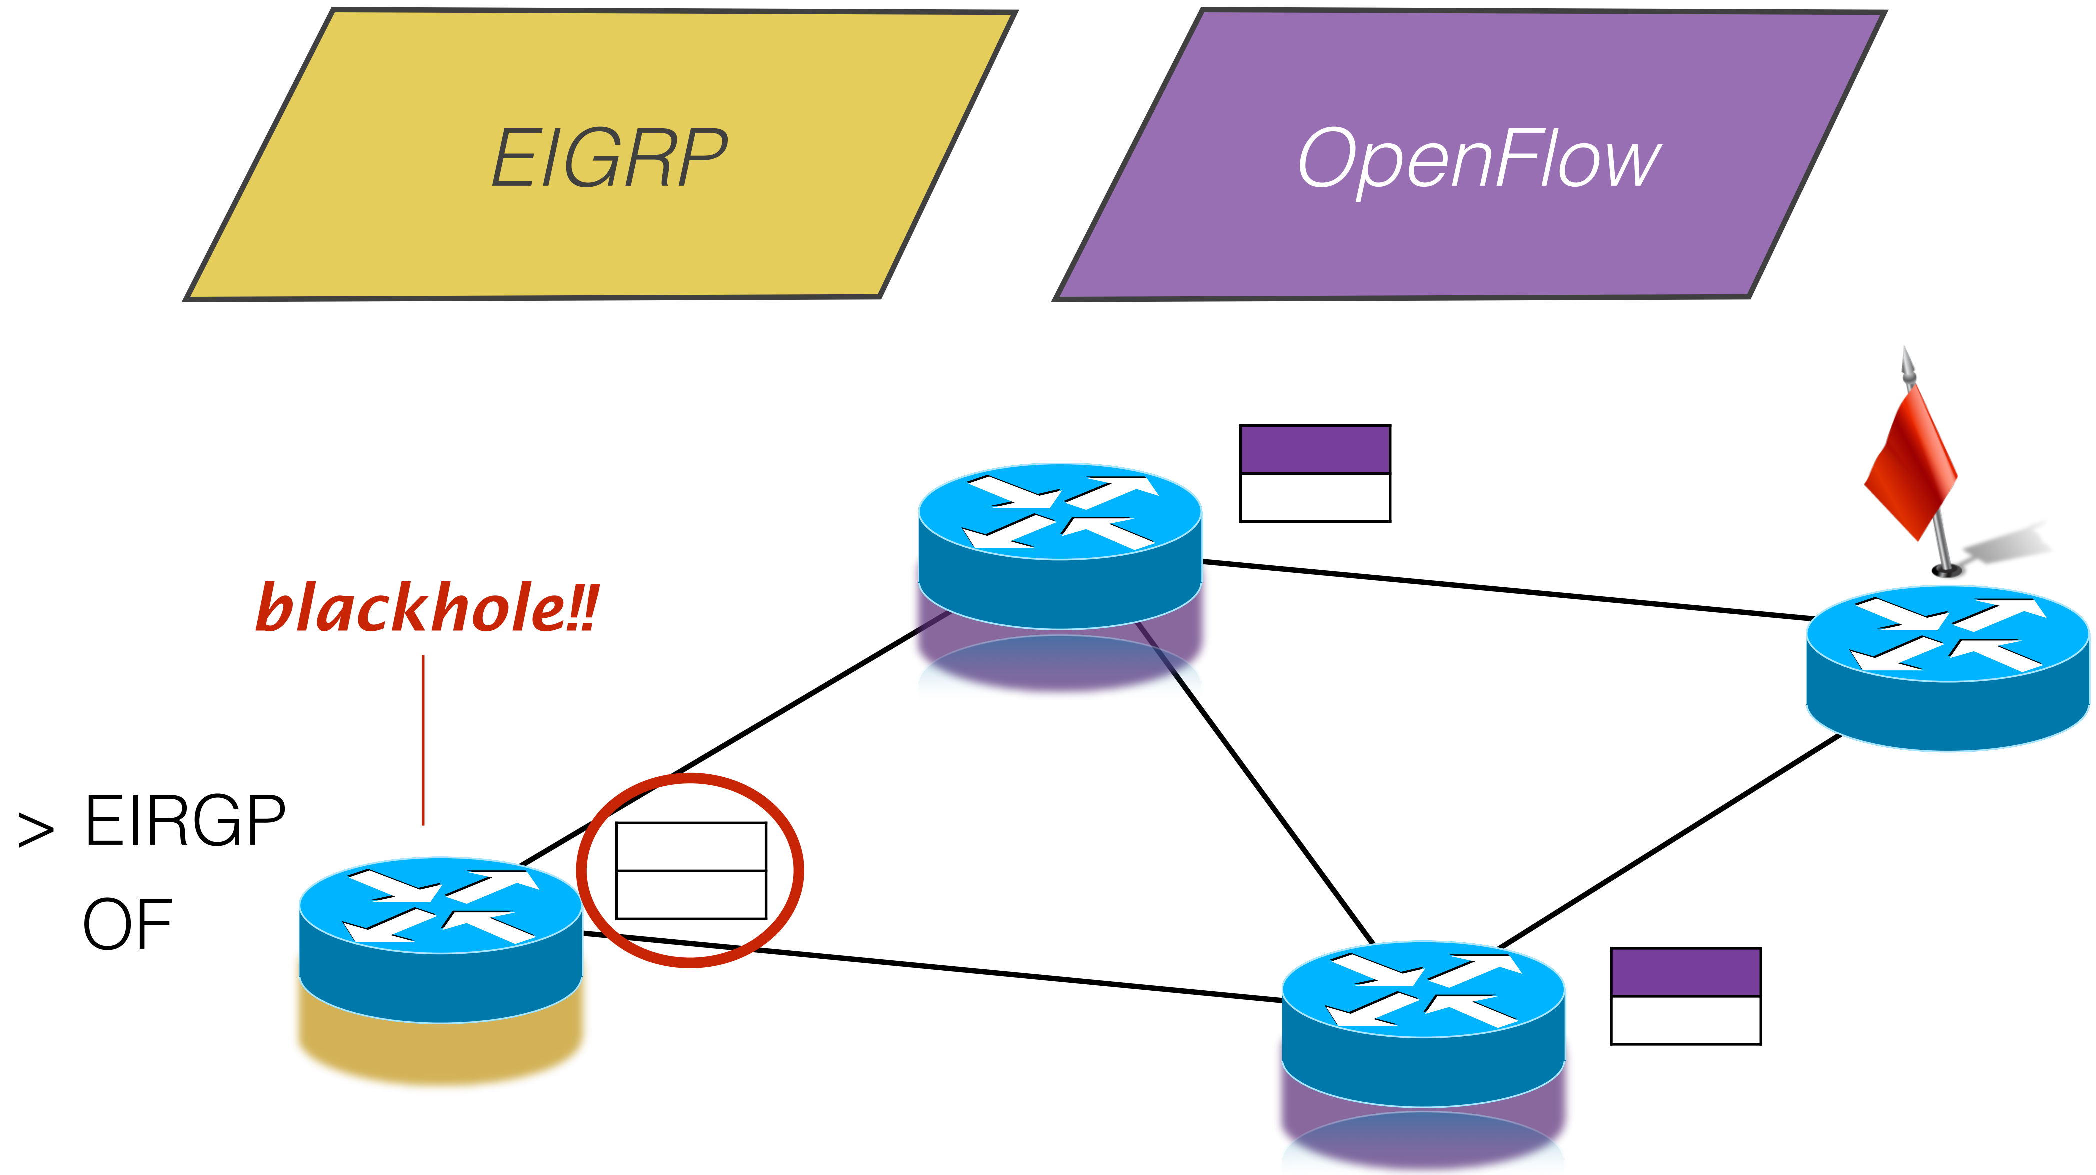

If IS-IS instead of EIGRP, no blackhole;
(nFU control-planes always provide routers with a route)

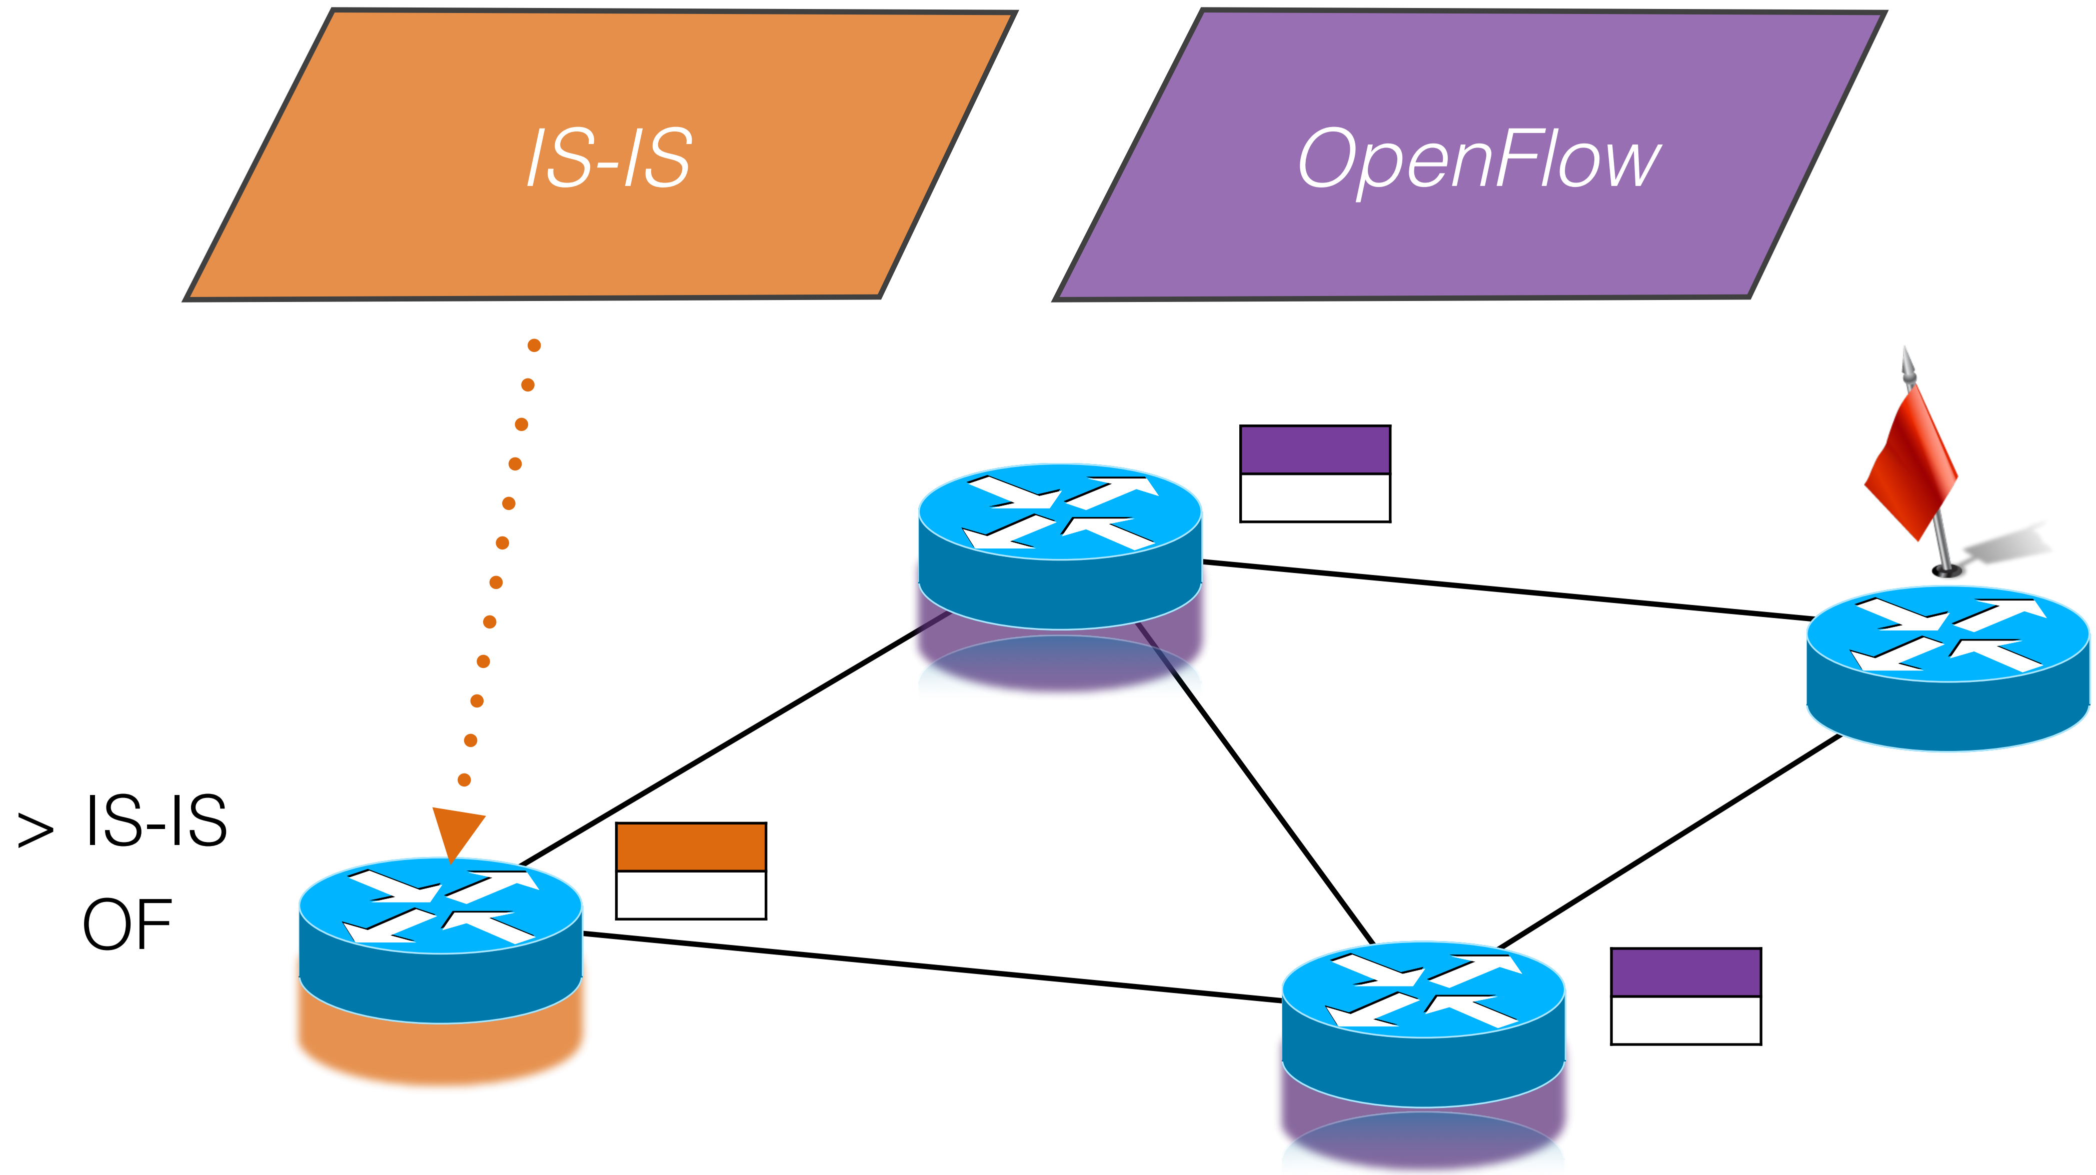

We proved that our taxonomy **characterizes** coexistence anomalies

- *Theo. 0*: No routing anomalies
  no information exchange between control-planes

- *Theo.1*: No blackholes guaranteed iff
  (i) at least one non-preemptive FU control-planes, OR
  (ii) no preemptive control-plane M1 + FA control-plane M2

- *Theo.2*: No loop guaranteed iff at most one FU control-plane

For example, our theorems can be applied to fully characterize two coexisting control-planes

	pFA (FIB-reacting SDN)	**pFU** (OpenFlow)	**nFA** (RIP, EIGRP)	**nFU** (OSPF, IS-IS)
pFA (FIB-reacting SDN)	blackholes	blackholes	blackholes	
pFU (OpenFlow)	blackholes	loops	blackholes	loops
nFA (RIP, EIGRP)	blackholes	blackholes		
nFU (OSPF, IS-IS)		loops		loops

Our findings highlights which coexisting control-plane combinations are inherently safe

	pFA (FIB-reacting SDN)	**pFU** (OpenFlow)	**nFA** (RIP, EIGRP)	**nFU** (OSPF, IS-IS)
pFA (FIB-reacting SDN)	blackholes	blackholes	blackholes	
pFU (OpenFlow)	blackholes	loops	blackholes	loops
nFA (RIP, EIGRP)	blackholes	blackholes		
nFU (OSPF, IS-IS)		loops		loops

For the remaining ones, we also provided sufficient conditions to avoid disruptions

	pFA (FIB-reacting SDN)	**pFU** (OpenFlow)	**nFA** (RIP, EIGRP)	**nFU** (OSPF, IS-IS)
pFA (FIB-reacting SDN)	blackholes	blackholes	blackholes	
pFU (OpenFlow)	blackholes	loops	blackholes	loops
nFA (RIP, EIGRP)	blackholes	blackholes		
nFU (OSPF, IS-IS)		loops		loops

Our contributions include modeling, formal analysis, and insight of the implications

- model for arbitrary control-planes
- characterization of coexistence anomalies
- practical applications of our theory
- analysis of the lessons learned

We applied our theory to both static and dynamic settings

For arbitrary combinations of control-planes, we provide

- Safe configuration guidelines
- Support for network reconfigurations

Safe coexistence guidelines to avoid blackholes can be derived from our theorems

*To avoid blackholes, apply **any** of the following*

- ***A1***: No preemptive control-planes
- ***A2***: At least one non-preemptive FU control-planes
- ***A3***: Subdivide the network in connected components, s.t.
  (i) for each component, one control-plane is preferred, AND
  (ii) each component is connected to a set of routers globally announcing all destinations

Safe coexistence guidelines to avoid loops can be derived from our theorems

*To avoid loops, apply **any** of the following*

- ***B1***: At most one FU control-plane
- ***B2***: Configure FU control-planes so that their combined routes do not contain loops for any destination

We apply our theory to reconfigurations from *any* combination of control-planes to *any* other

Leveraging our characterization, we can

- *predict* possible anomalies
  occurring during reconfigurations
- devise a generic reconfiguration *procedure*
  preserving forwarding correctness

Our contributions include modeling, formal analysis, and insight of the implications

- model for arbitrary control-planes
- characterization of coexistence anomalies
- practical applications of our theory
- analysis of the lessons learned

Lesson 1:
Design protocols with coexistence in mind

For example, straightforward deployment of current OpenFlow can jeopardize coexisting control-planes!!

- blackholes AND loops are possible in OpenFlow + IGP
  (disincentive to migrate to SDN)
- in comparison, I2RS prevents blackholes and
  FIB-reacting controllers avoid loops

Lesson 2:
Design networks with coexistence in mind

For example, operators should evaluate coexistence when choosing control-plane protocols

- e.g., safe coexistence == easy reconfigurability
- while possibly profitable, coexistence imposes a tradeoff between correctness and manageability

Lesson 3:
Define control-plane inputs/outputs unambiguously

From RFC, it is unclear if RIP's input should be the RIP RIB or the router's FIB

- RIP is FA in Cisco/Juniper routers, but FU in Quagga (hard to catch even for interoperability tests)
- a RIP network with both Cisco and Quagga routers would be unpredictably hard to update!!

On the Co-Existence of Distributed and Centralized Routing Control-Planes

- model for arbitrary control-planes
- characterization of coexistence anomalies
- practical applications of our theory
- analysis of the lessons learned

Stefano Vissicchio

stefano.vissicchio@uclouvain.be