Opportunities and Research Challenges of Hybrid Software Defined Networks

Stefano Vissicchio – Université catholique de Louvain
Laurent Vanbever – Princeton University
Olivier Bonaventure – Université catholique de Louvain

SDN promises improvements…

- **Ease management**
  - remove today's challenges
- Unleash additional **flexibility**
  - more easily implement complex policies
  - support Network Function Virtualization
- **Be real**
  - commodity hardware
  - growing industrial interest

…by relying on architectural changes

- **homogeneous** devices
  - with API for programmability

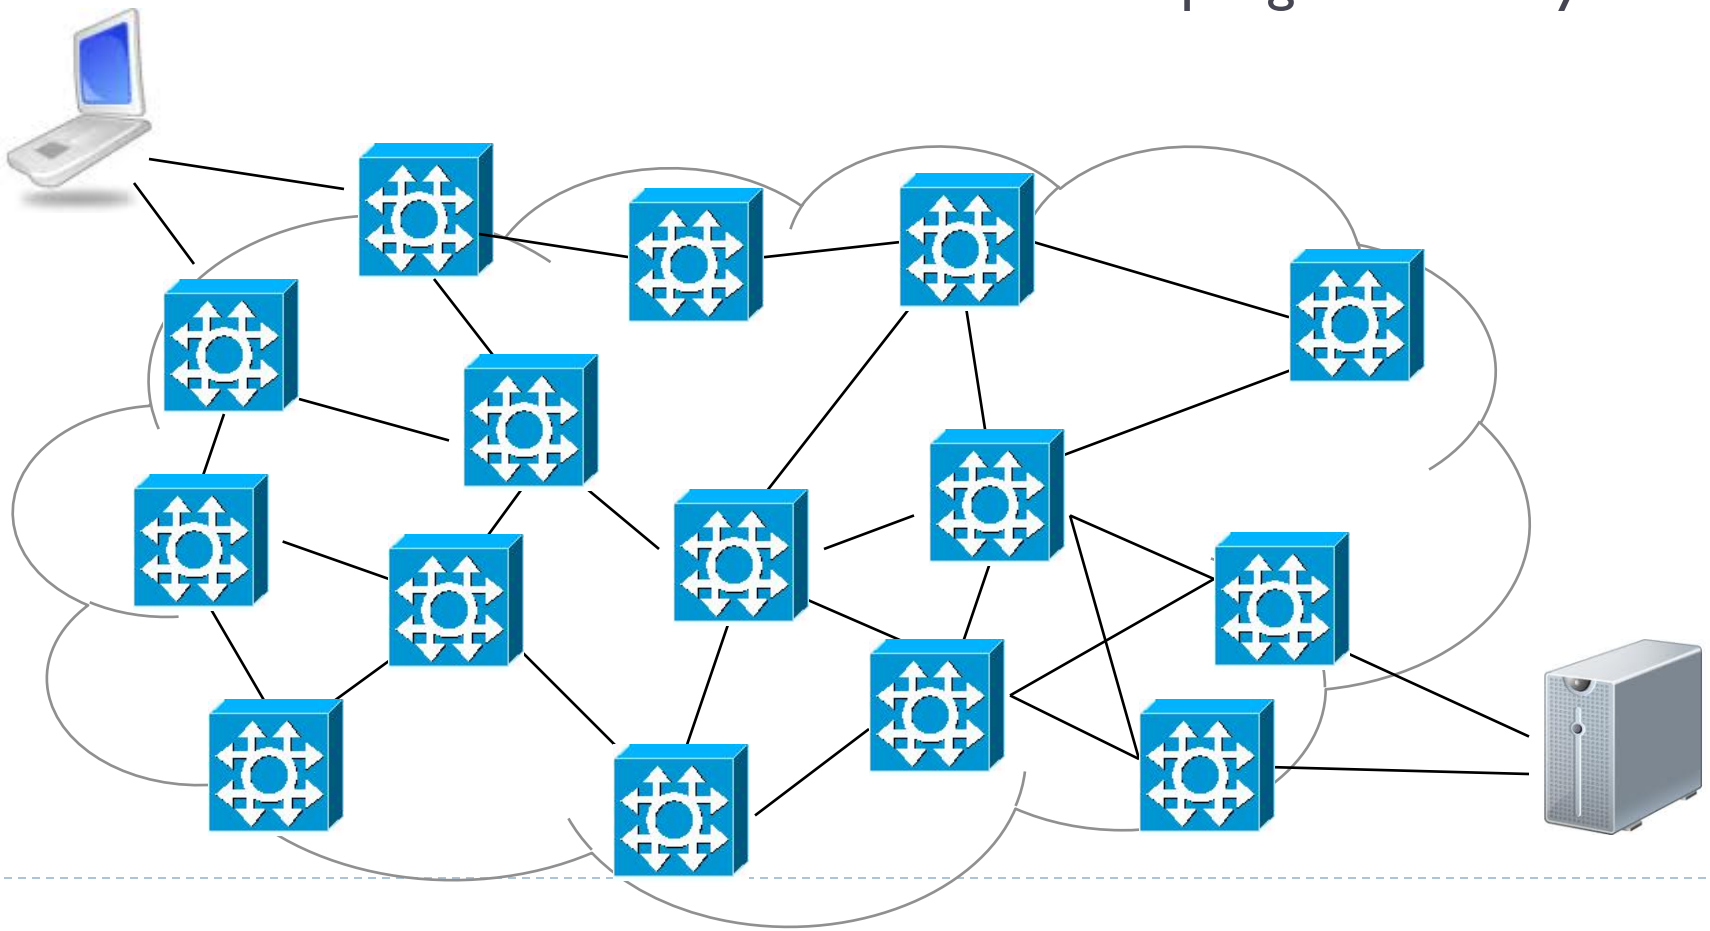

…by relying on architectural changes

- logic **centralization**
  - custom software and algorithms

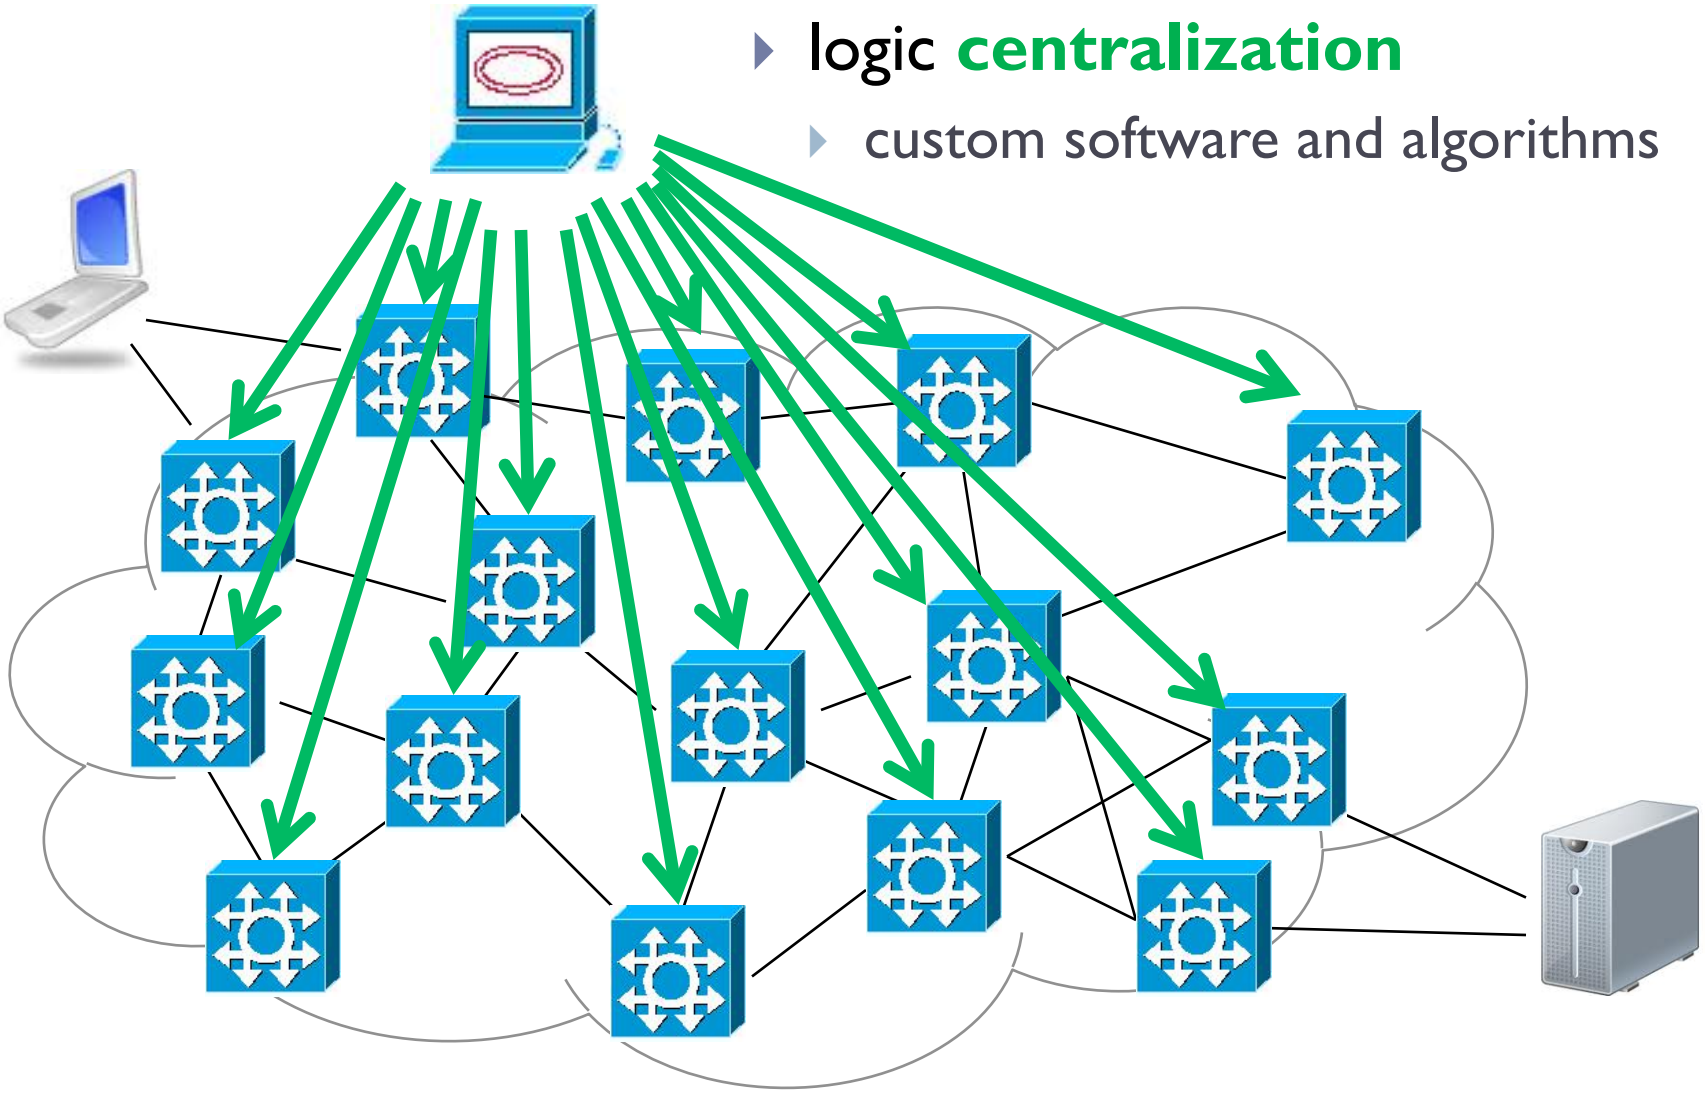

The SDN proposal

- New paradigm
  - centralized computation
  - programmatic interface to forwarding
- New abstractions
  - declarative interface to operators
- New protocols
  - new API to devices
- New hardware
  - with specific capabilities, e.g., match any packet field

New solutions → new tradeoffs

- Reaction to failures
  - querying the controller [Sharma13] → complex controller, communication overhead, expensive out-of-band network
  - pre-installing backup flow entries [Reitblatt13] → more complex protocols and hardware, not arbitrarily scalable
- Scalability
  - distributed controllers [Koponen10] → instance synchronization, state consistency
- Communication with the SDN controller
  - state of the art out-of-band network → expensive, doubles the network problems

Something to save from good old days?

- Distributed protocols provide robustness
  - per-device control-plane is a form a robustness
  - local reactions, seamless convergence techniques
    - e.g., [Filsfils12,Aceves93,Clad13,rfc4090]
- Scalability is widely studied
  - e.g., routing hierarchy + route summarization [rfc2328], route redistribution [Le07]
- No need to communicate with an external system

What about a hybrid SDN approach?

- SDN may not be needed for everything!
- ***Hybrid SDN*** = co-existence of SDN and distributed protocols
  - we focus on IGPs as distributed protocols
- Challenge: keep only the best of each approach
  - summing resp. advantages
  - mitigating resp. limitations
  - handling the complexity of multiple paradigms

Different hybrid SDN models

- Topology-based coexistence (TB)
  - independent IGP and SDN, running on different subnetworks
- Class-based coexistence (CB)
  - independent IGP and SDN, controlling different traffic classes
- Integration (I)
  - SDN controlling IGP

TB hSDN: Basic idea

- The network is partitioned in zones
  - each device belong to only one zone
  - a zone can be managed by either SDN or IGP

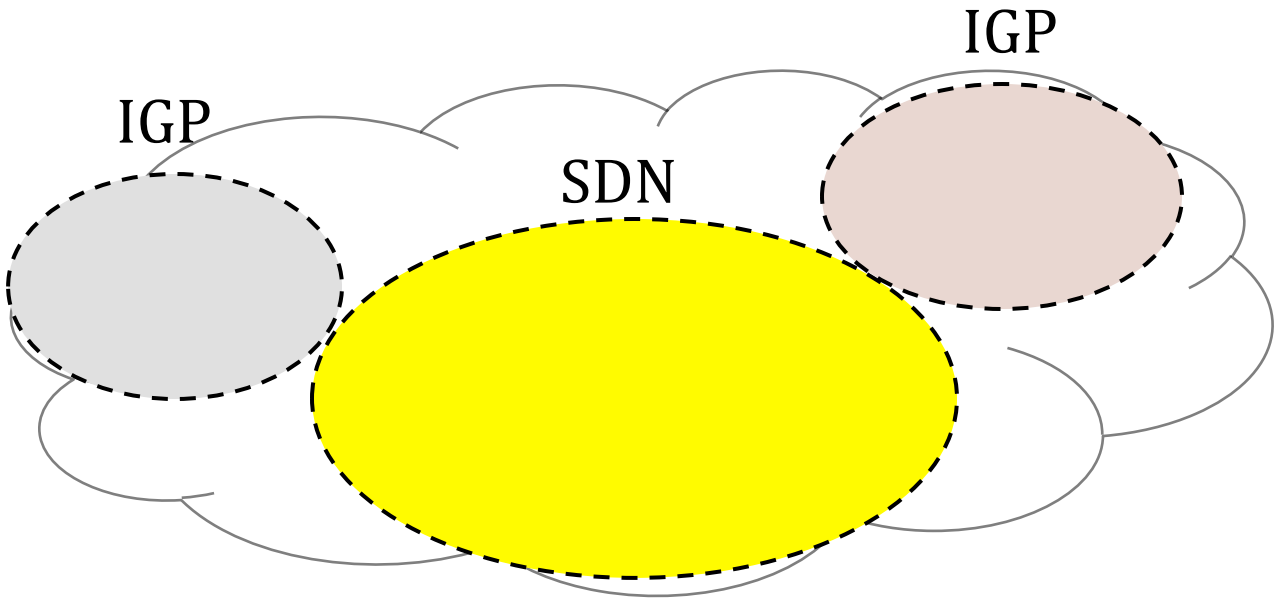

TB hSDN: Opportunities

- No upgrade **costs** where unnecessary
- Fewer requirements → **easier IGP configurations**
- SDN controller manages smaller networks
  - **less complex controllers, less scalability concerns**

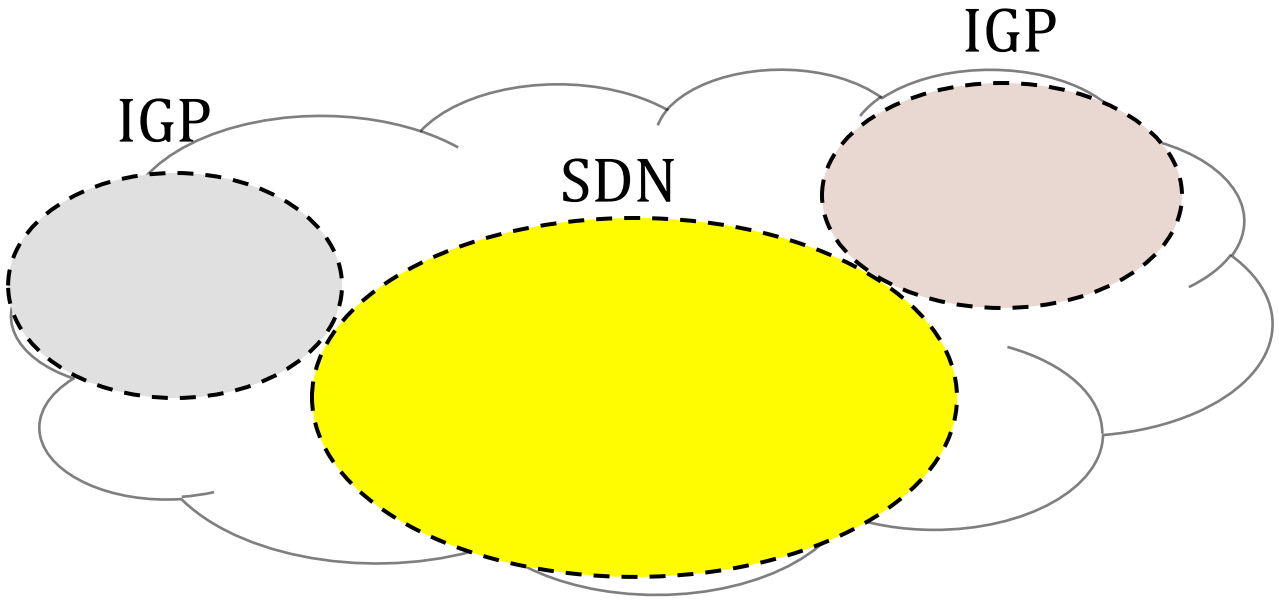

TB hSDN: Challenges

- Complex **paradigm interaction**
  - information exchanges needed between zones
- **Different system interfaces** in different zones

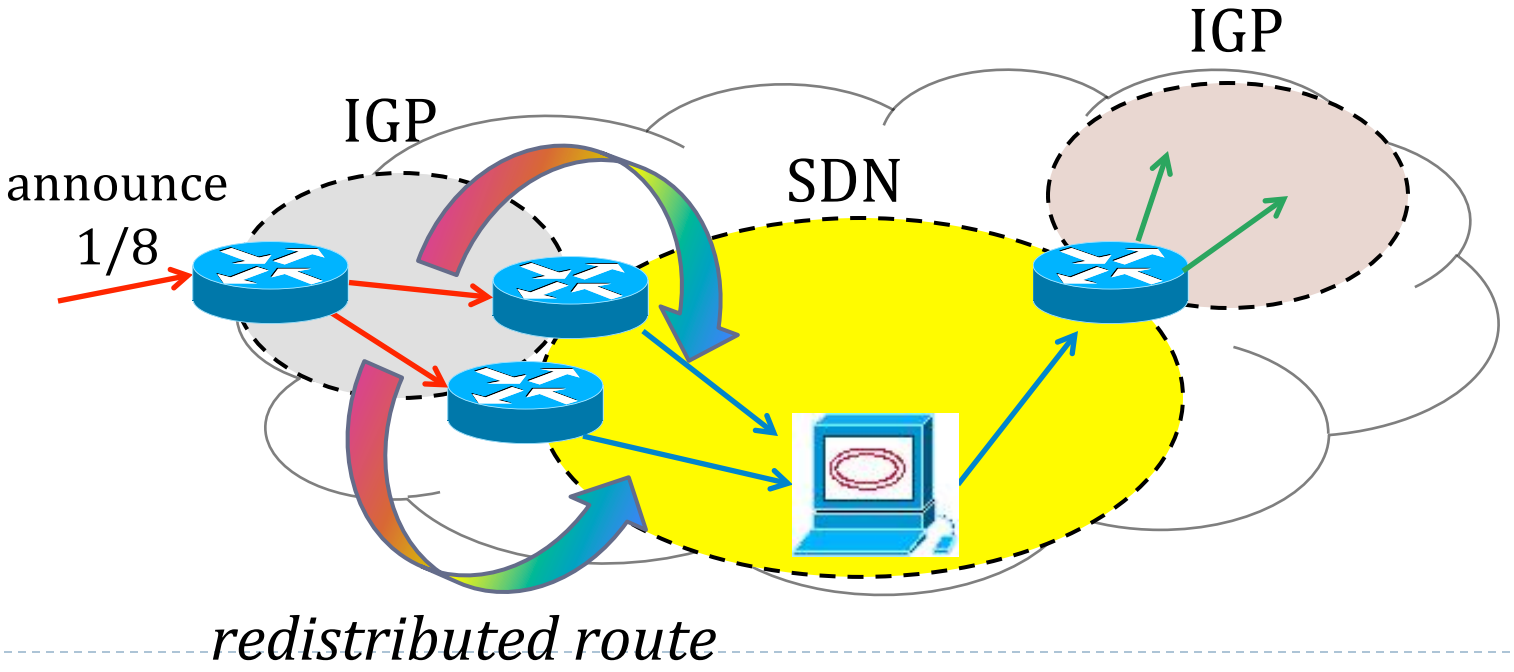

CB hSDN: Basic idea

- SDN and IGP control different traffic classes
  - on the same physical topology

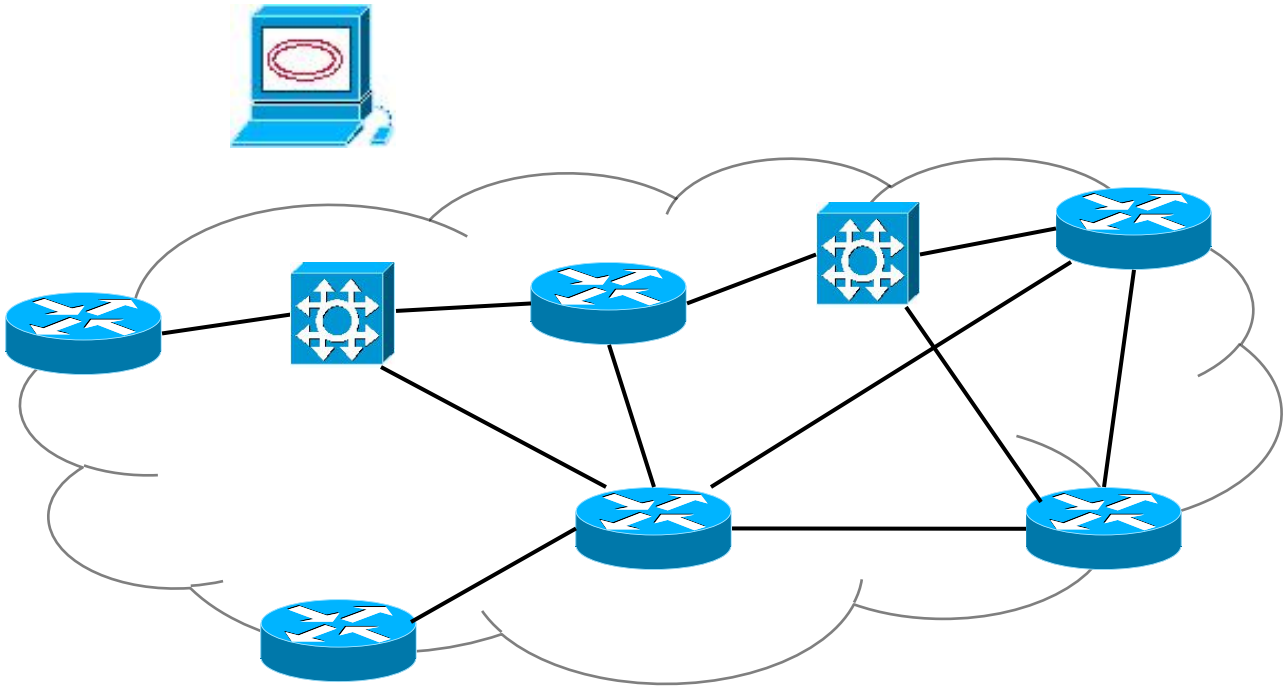

CB hSDN: Basic idea

- SDN and IGP control different traffic classes
  - on the same physical topology

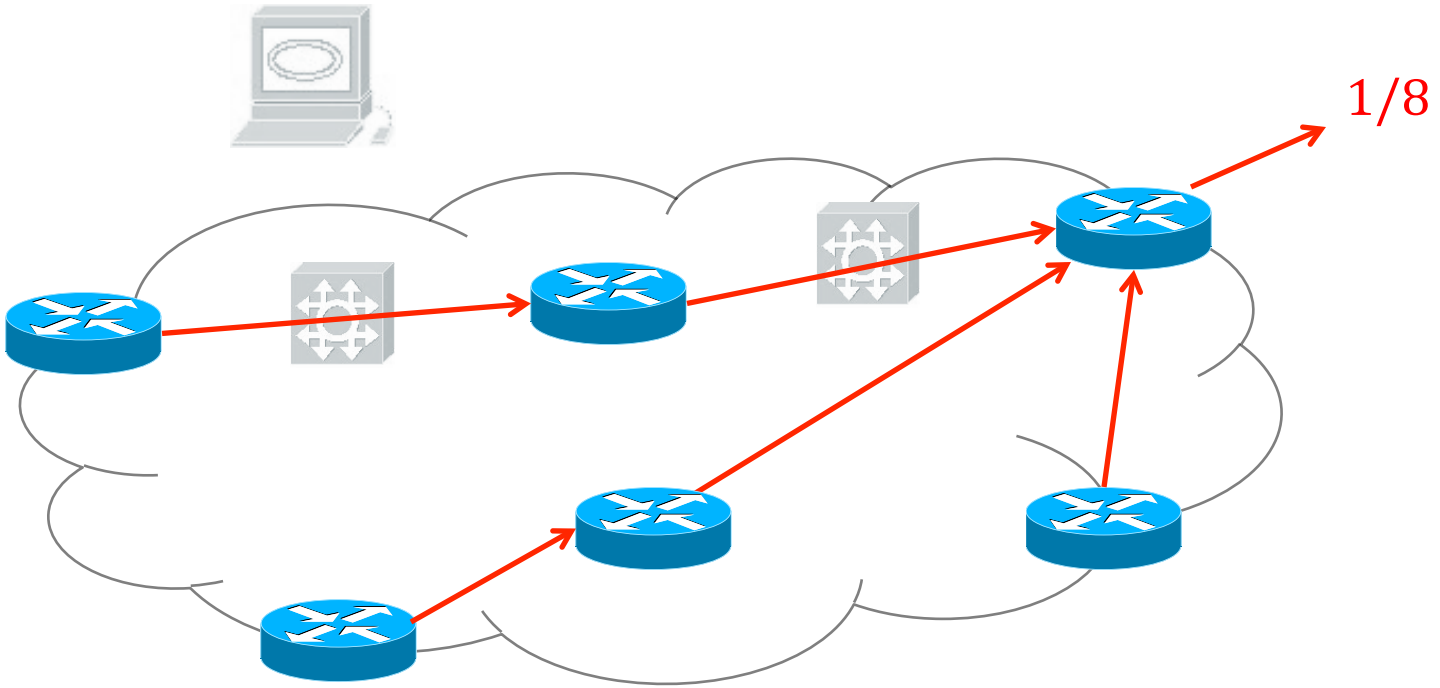

CB hSDN: Basic idea

- SDN and IGP control different traffic classes
  - on the same physical topology
- The controller installs **static routes** on IGP routers

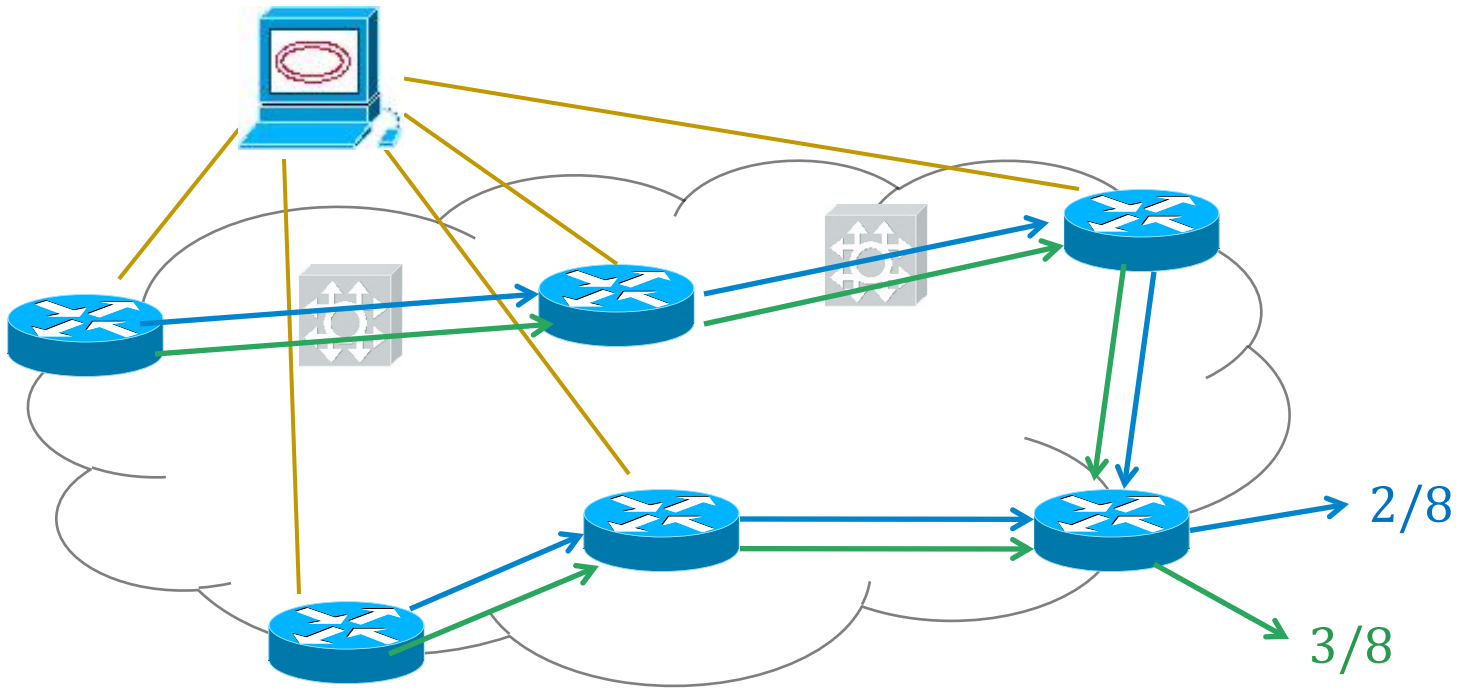

CB hSDN: Basic idea

- SDN and IGP control different traffic flows
  - on the same physical topology
- The controller programs **forwarding** on SDN devices

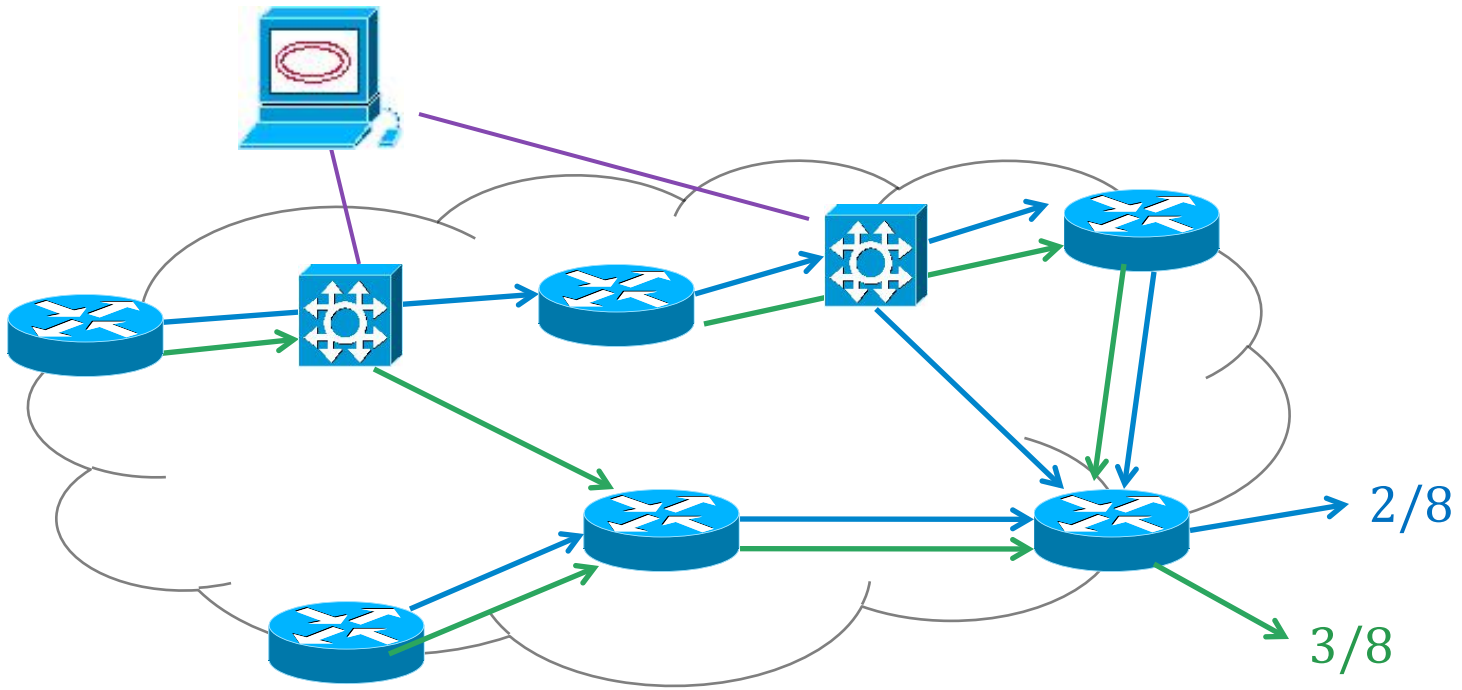

CB hSDN: Opportunities

- **Enabled SDN capabilities**
  - e.g., on critical traffic
- Some traffic **offloaded from the controller**

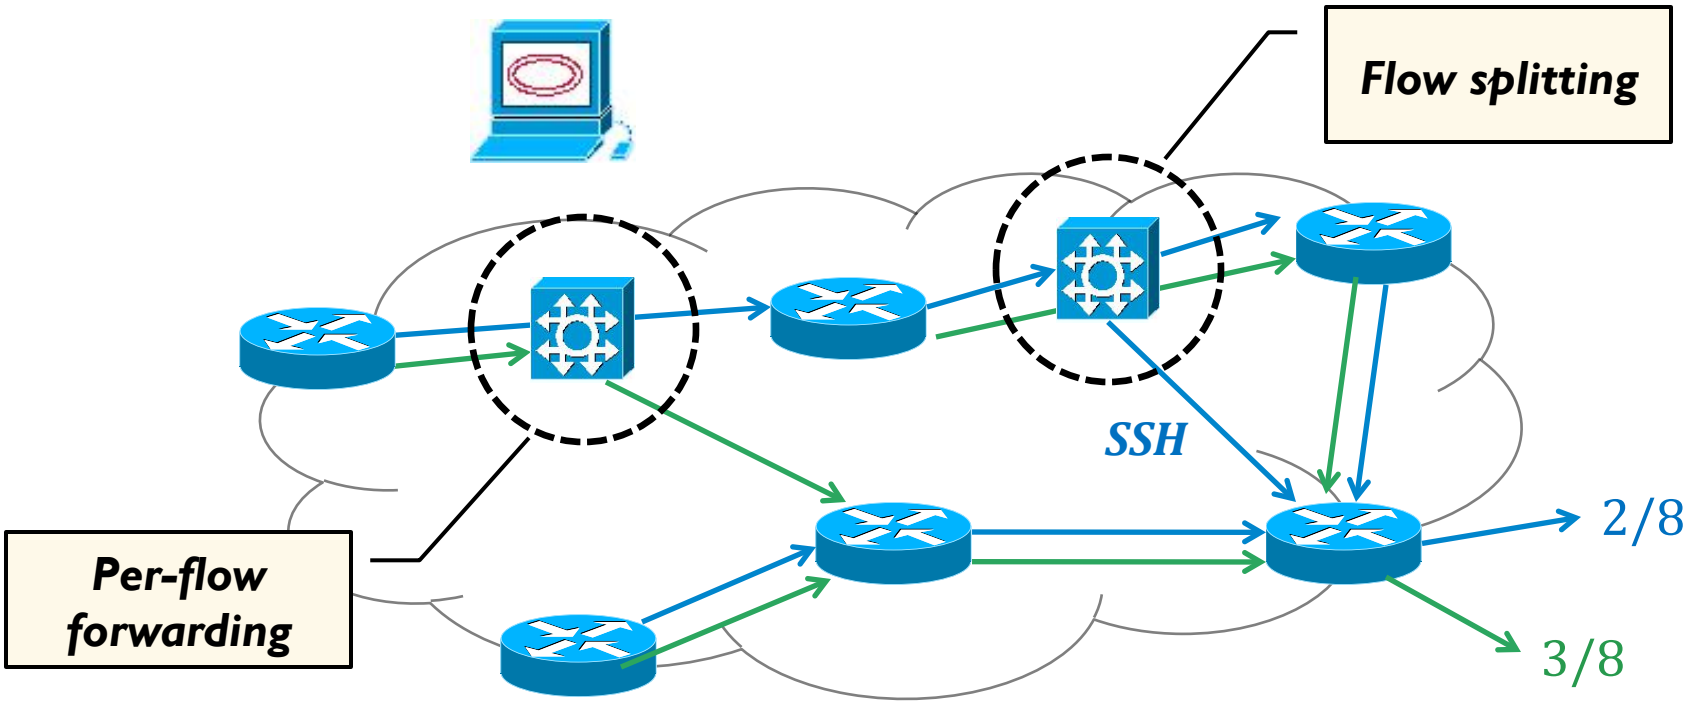

CB hSDN: Challenges

- **Control-plane coordination**
  - e.g., to transfer control of classes from SDN to IGP
- **No unified management interface**

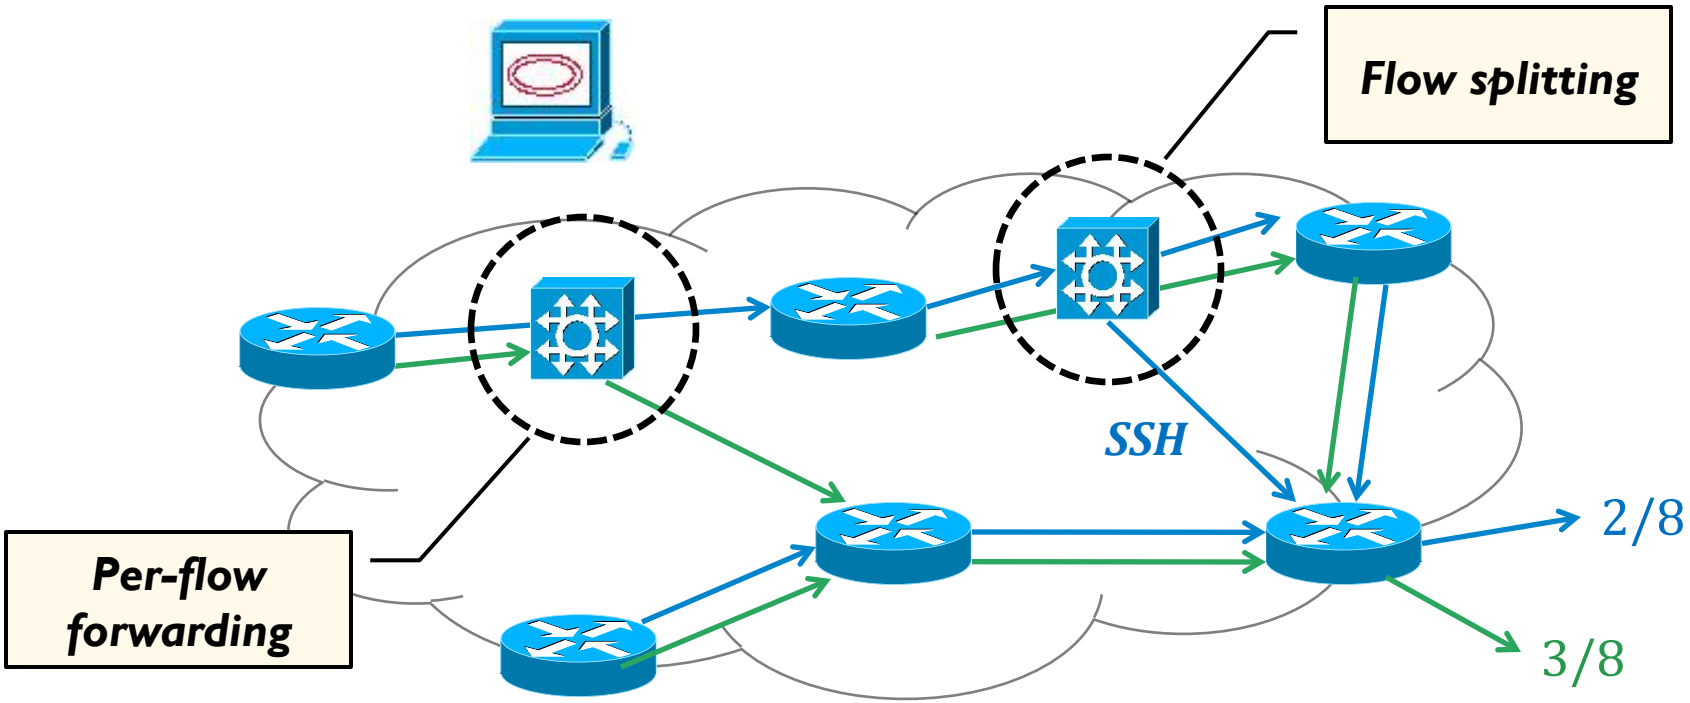

I hSDN: Basic idea

- SDN needs programmatic interface to devices
  - OpenFlow is an enabler for SDN
- IGP = distributed computation of forwarding tables
  - standardized, configurable
- What about using IGP as an API for SDN?
  - declaring forwarding requirements
  - computing paths on the SDN controller
  - implementing paths through IGP

I hSDN: Basic idea

- IGP forwarding depend on a logical graph
  - on a physical network

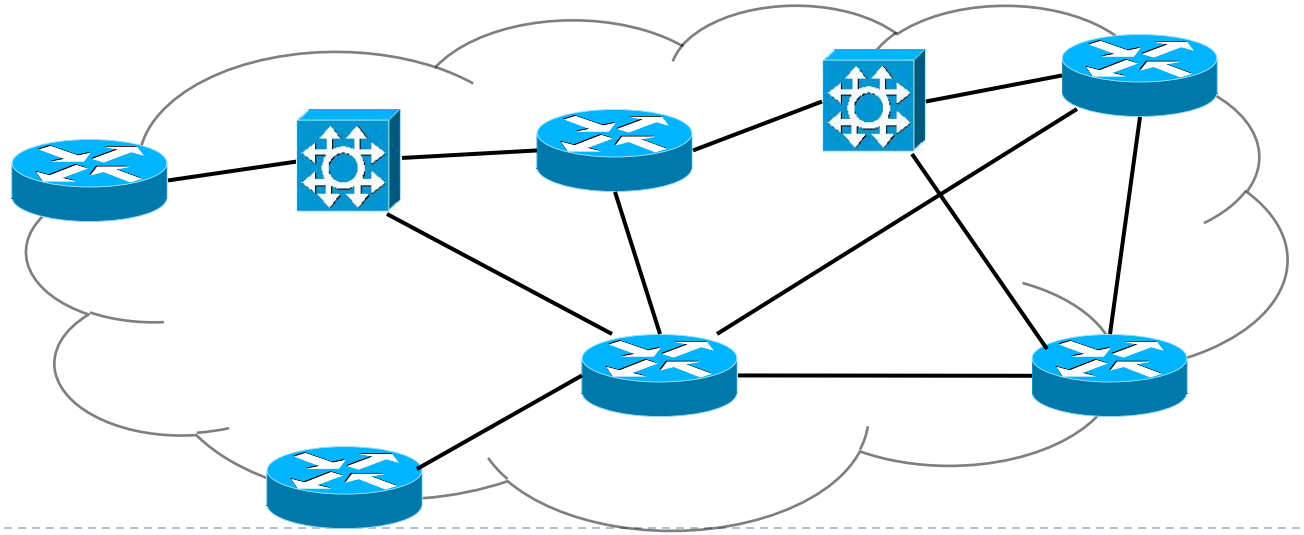

I hSDN: Basic idea

- IGP forwarding depend on a logical graph
  - on a physical network
  - the IGP builds a logical graph of adjacencies

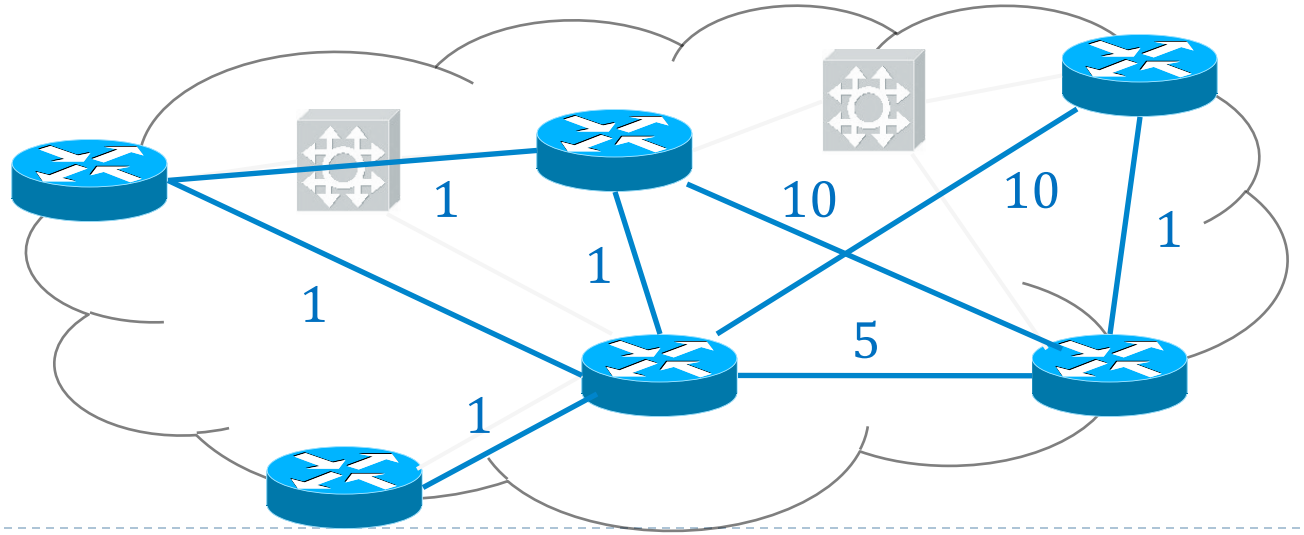

I hSDN: Basic idea

- IGP forwarding depend on a logical graph
  - on a physical network
  - the IGP builds a logical graph of adjacencies
  - IGP forwarding paths = shortest paths on the logical graph

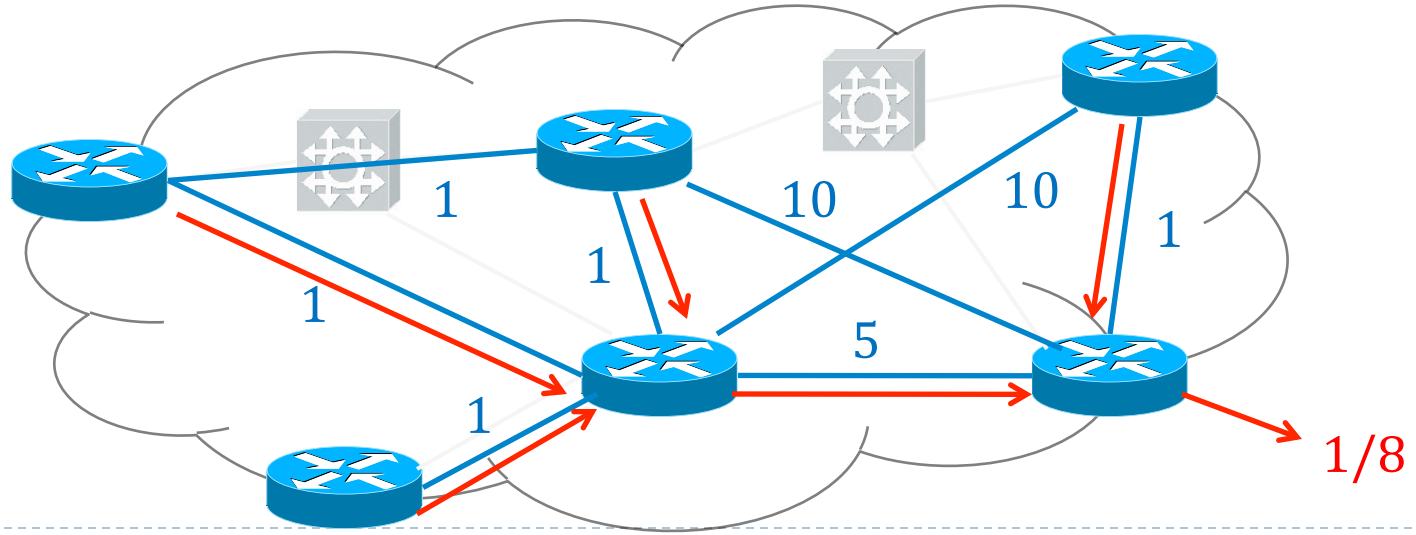

I hSDN: Basic idea

- What about faking the IGP topology?
  - adding ***fake nodes***

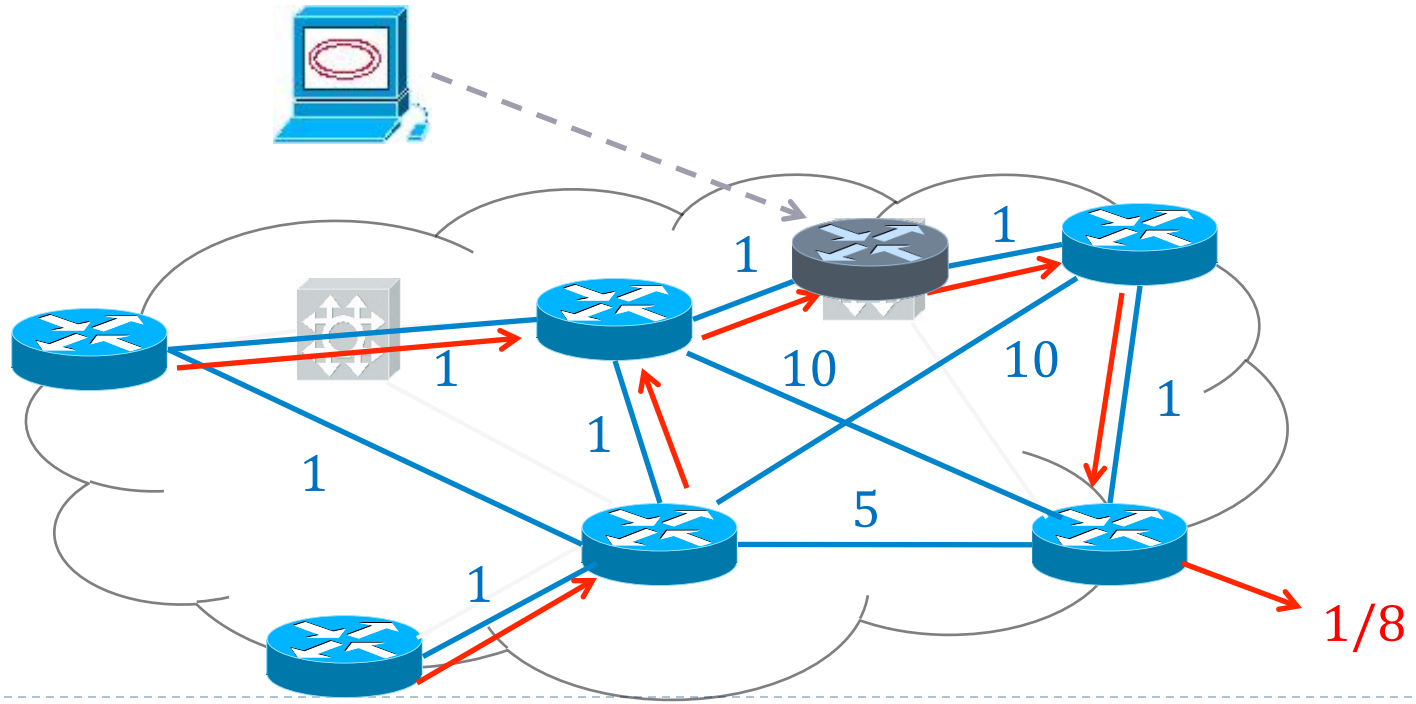

I hSDN: Basic idea

- What about faking the IGP topology?
  - adding ***fake nodes*** (e.g., to attract traffic to SDN switches)

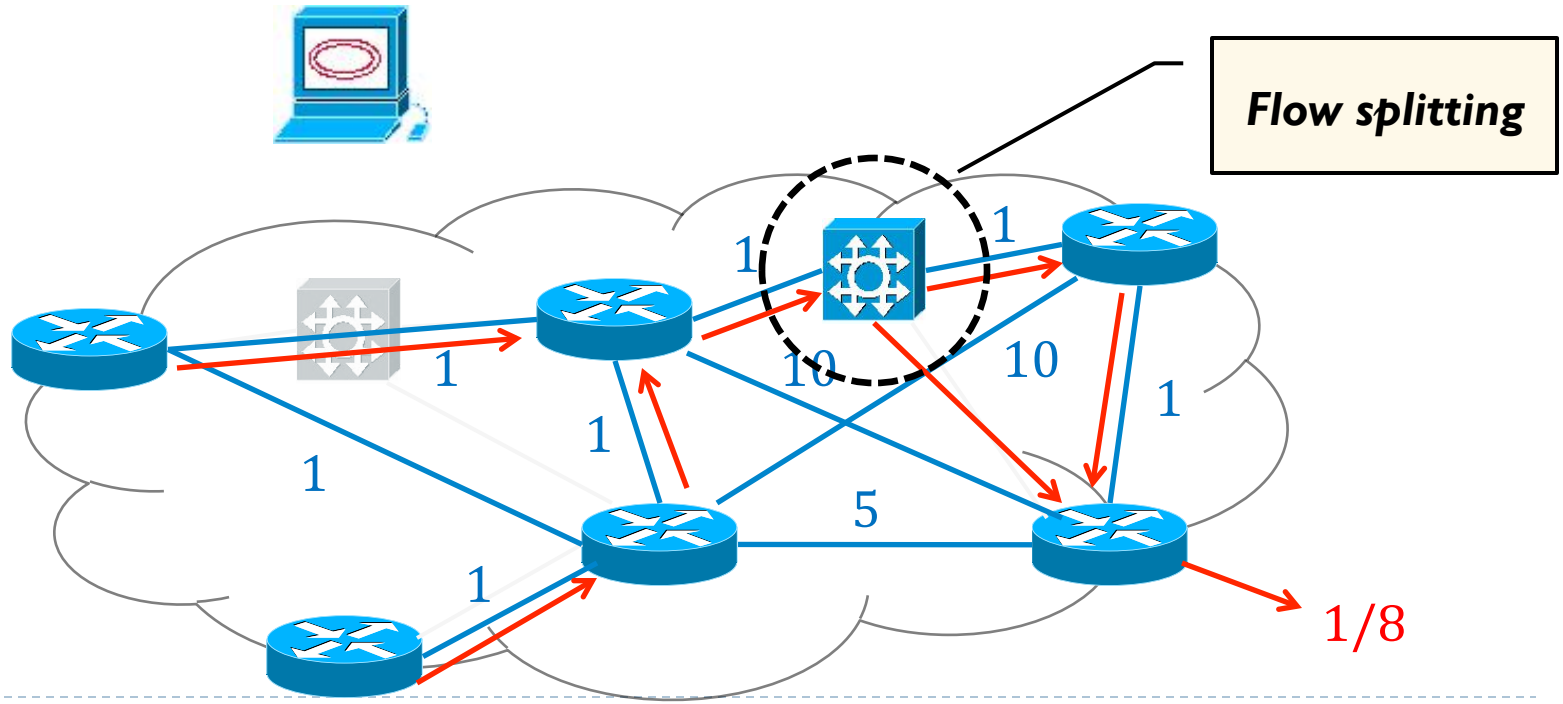

Basic idea

- What about faking the IGP topology?
  - adding fake nodes (e.g., to attract traffic to SDN switches)
  - tweaking ***fake link weights*** (e.g., to adapt to traffic matrix)

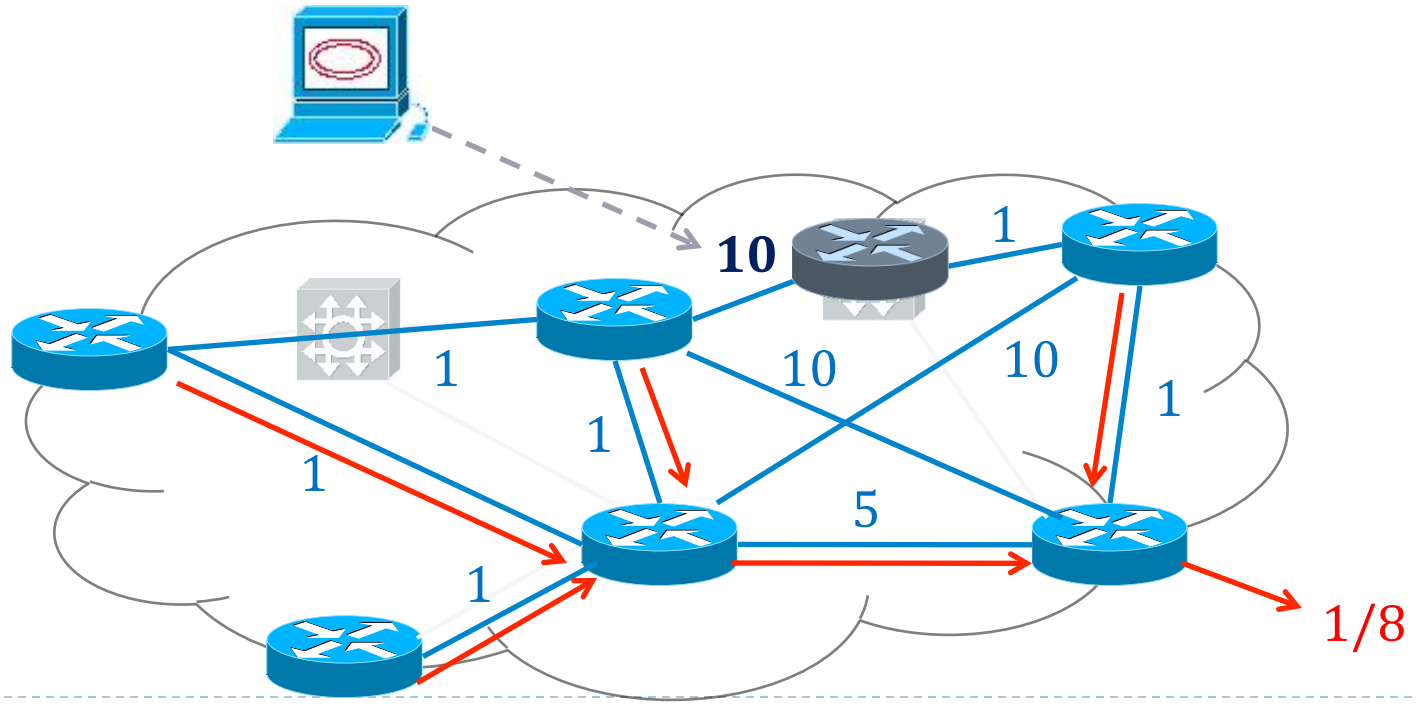

I hSDN: Basic idea

- What about faking the IGP topology?
  - adding fake nodes (e.g., to attract traffic to SDN switches)
  - tweaking fake link weights (e.g., to adapt to traffic matrix)
  - adding ***fake destinations*** (e.g., for finer-grained TE)

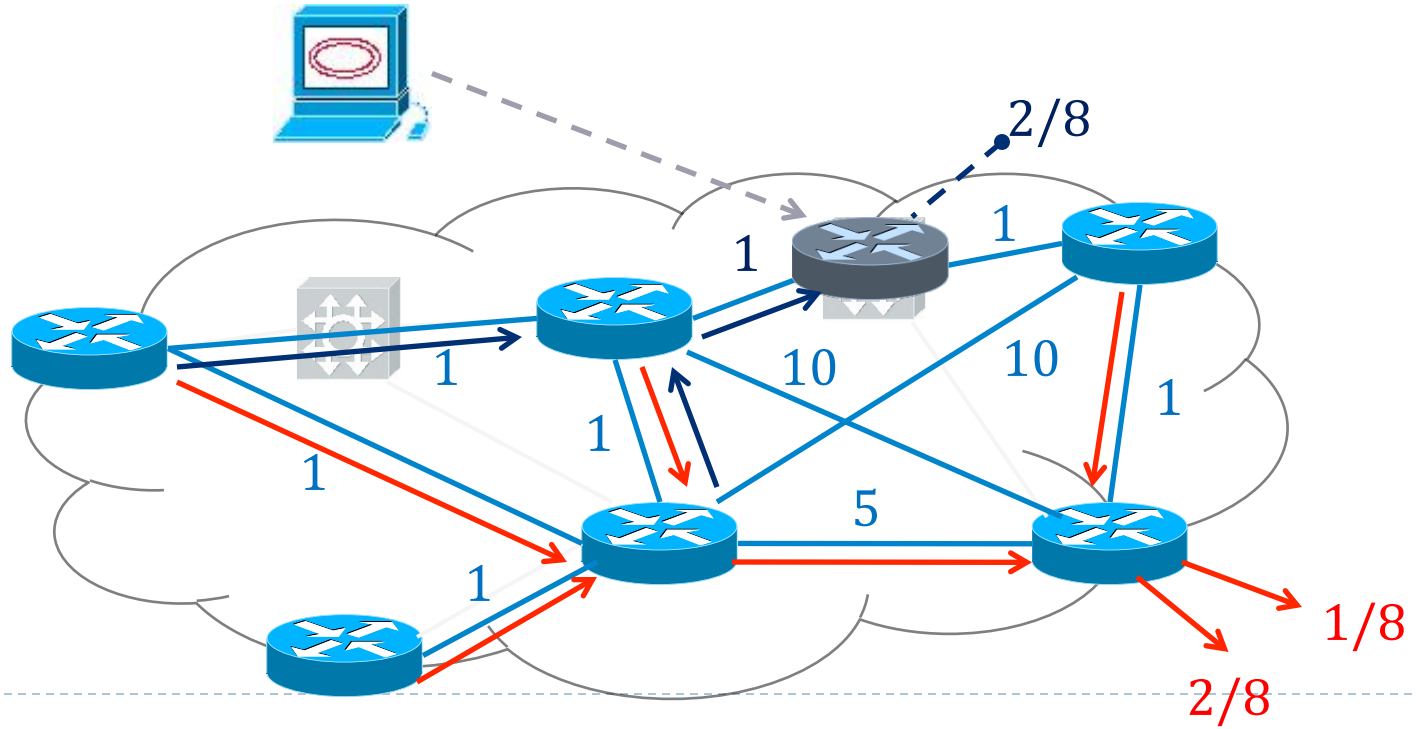

I hSDN: Opportunities

- No update **costs**
- More **powerful device API** than OpenFlow
  - some tasks can be delegated to IGP

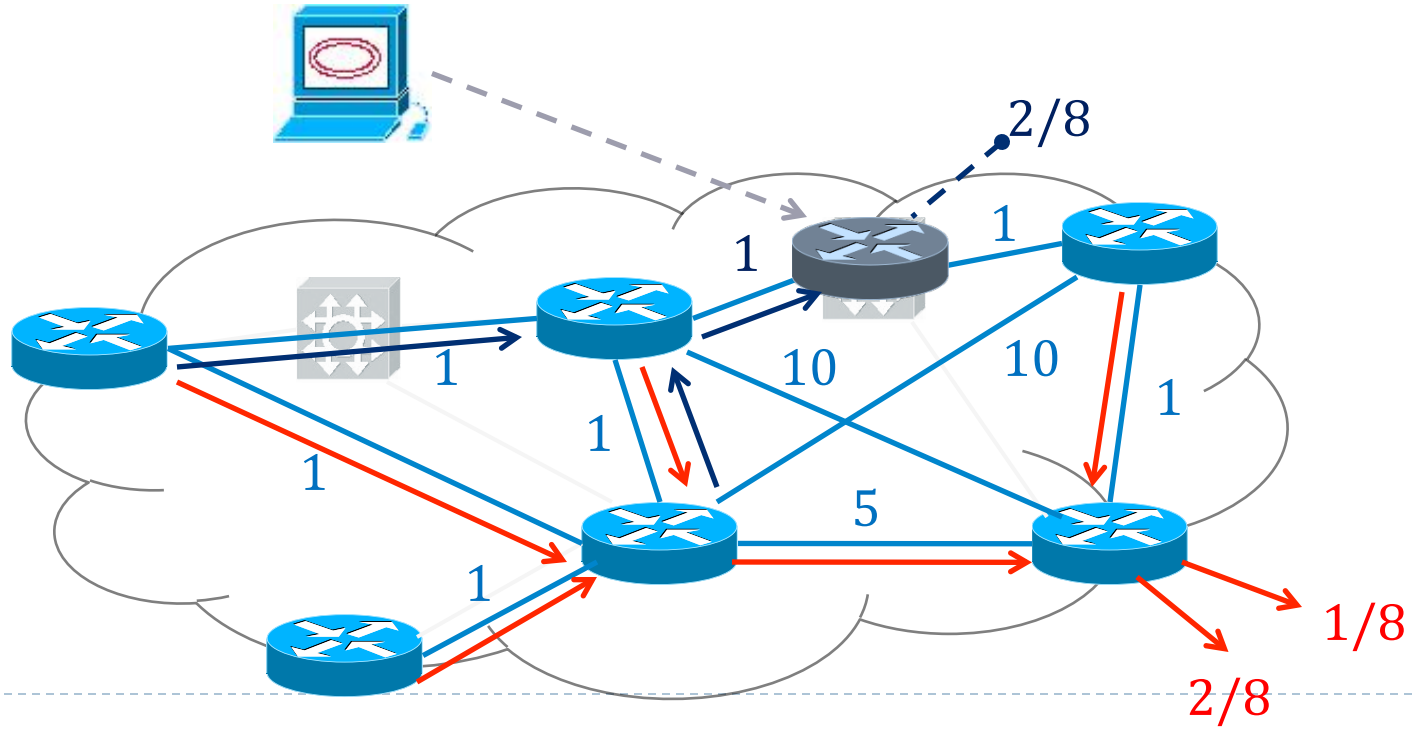

I hSDN: Challenges

- Limited **SDN capabilities**
  - e.g., SDN-specific capabilities only on some paths
- **More complex API** to devices
  - e.g., IGP specific mechanisms to be managed

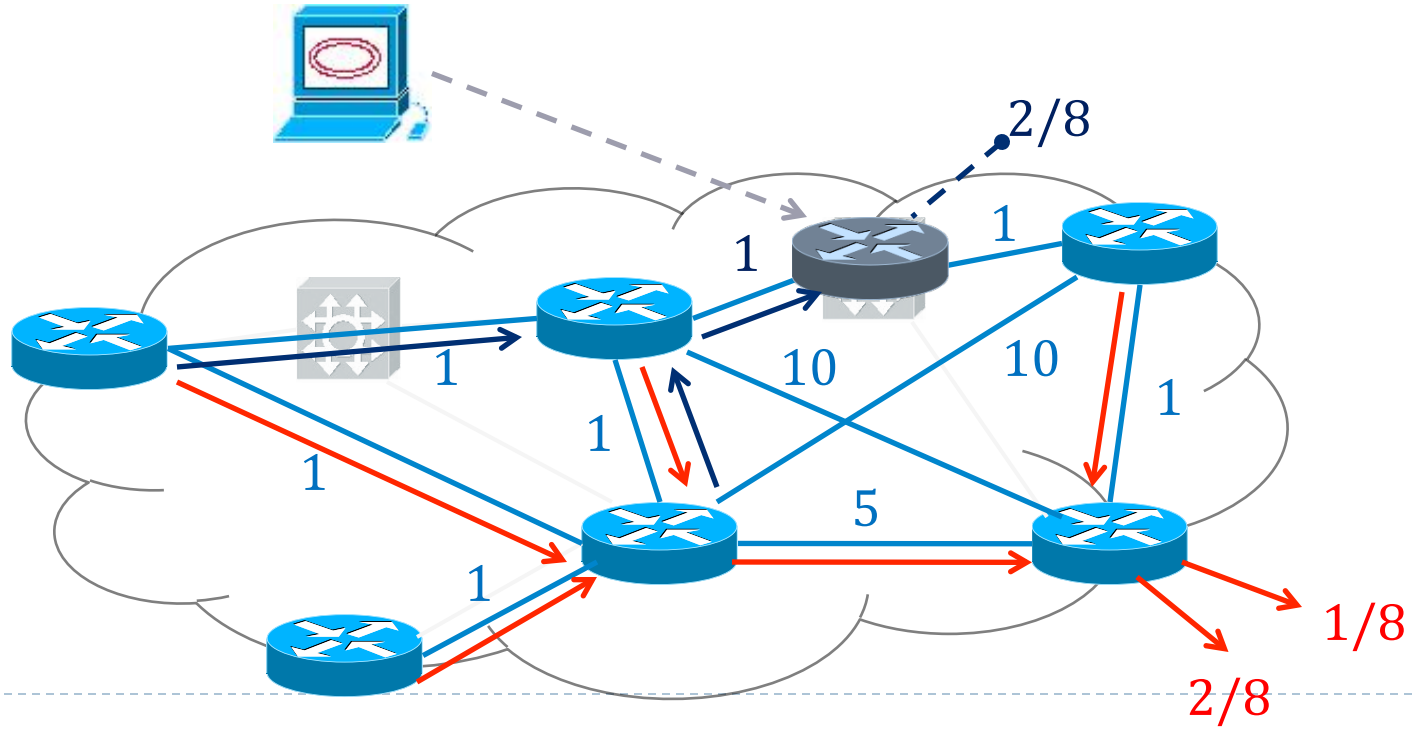

Take away

- Hybridization sacrifices SDN advantages to mitigate SDN limitations
  - different models, different tradeoffs
  - model combinations to further tradeoffs

- The SDN controller can be offloaded from some tasks
  - delegating protocols and techniques working for years
    - e.g., for short-term decisions
  - but this requires interaction with traditional protocols

Vision

- Hybrid SDN can provide effective transitional strategies
  - deployable today
  - provides incentives for transition to SDN
  - reduces transitional costs

- Hybrid SDN can be an interesting network design point
  - combining SDN innovation and flexibility with well known and proven guarantees of distributed protocols

Network management is hard today

- Distributed computation
- Indirect control of forwarding decisions
- Low-level configuration languages
- Limited flexibility in the configuration knobs

…by relying on architectural changes

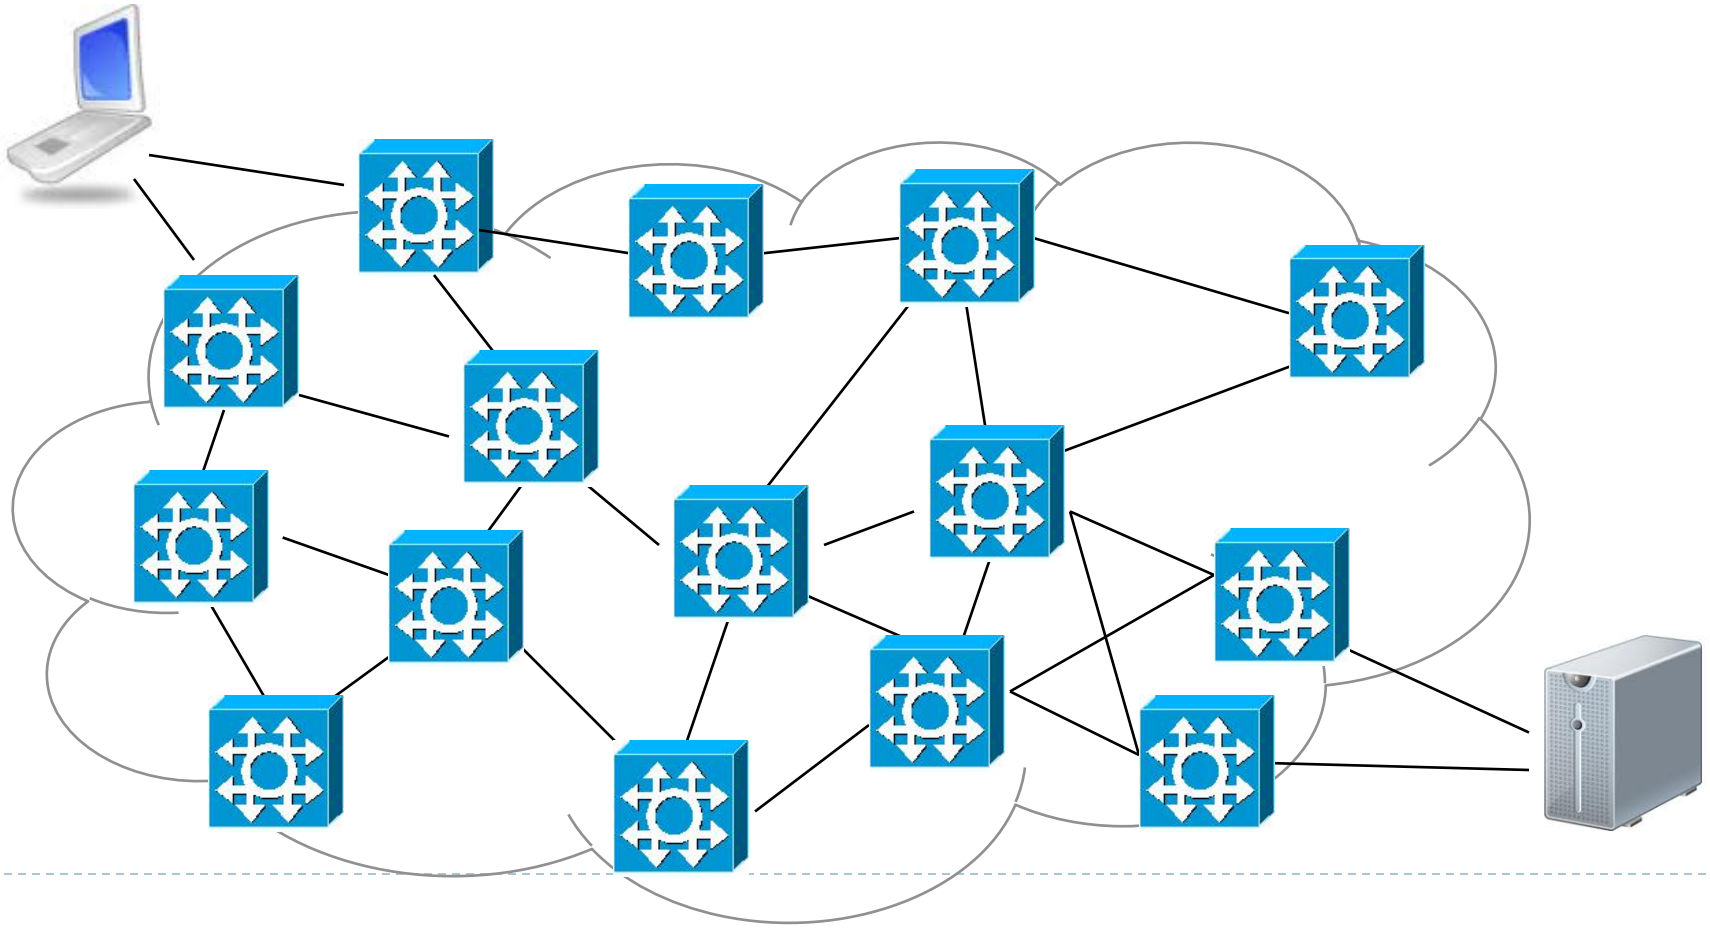

Topology-Based (TB) Hybrid SDN

Basic idea

- The network is divided in zones
  - a zone can be managed by either SDN or IGP

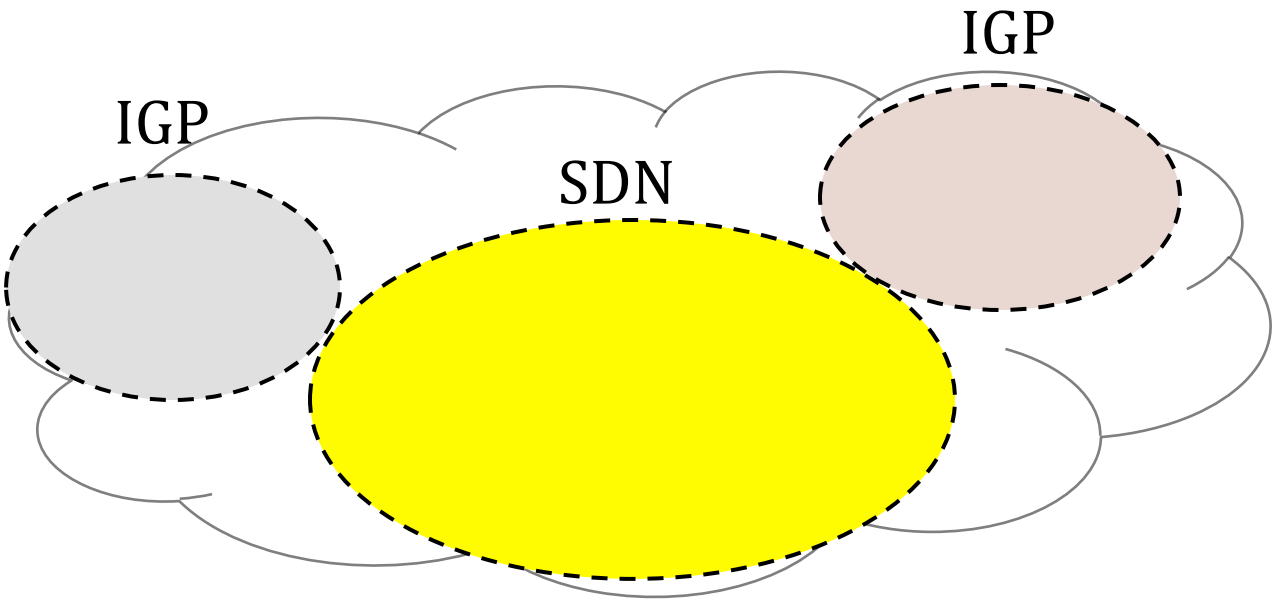

Basic idea

- The network is divided in zones
  - a zone can be managed by either SDN or IGP
- Information exchanges needed between zones
  - we assume route redistribution

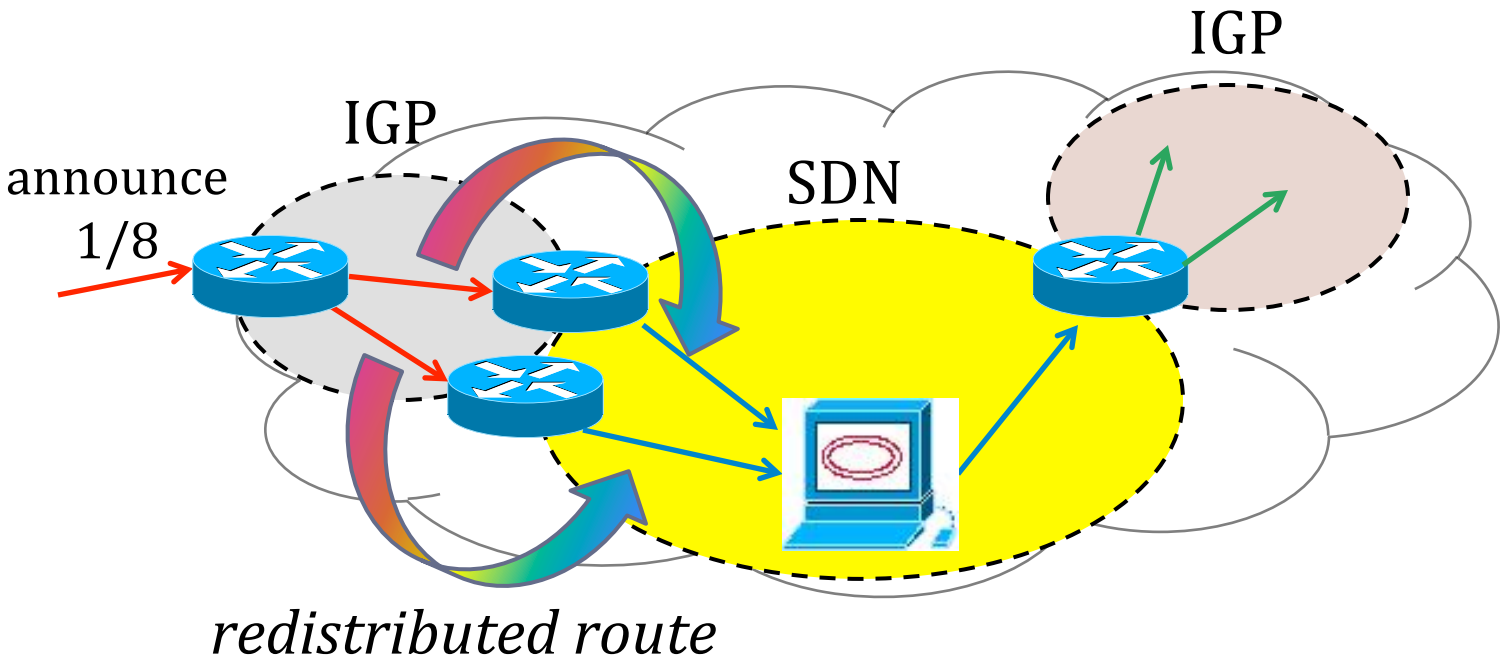

Pros and cons

- Local SDN controllers naturally match zones
- Automatic failure recovery in the IGP zones
- Improved scalability via aggregated route redistribution
- ...

- SDN capabilities are limited to SDN zones
- Additional complexity of route redistribution [Le07]
- …

Pros and cons

- Local SDN controllers naturally match zones
- Automatic failure recovery in the IGP zones
- Improved scalability via aggregated route redistribution
- ...

- SDN capabilities are limited to SDN zones
- Additional complexity of route redistribution [Le07]
  - **per-packet consistent update**
- …

Single zone updates

- Does the presence of other zones impact the update of a single zone?

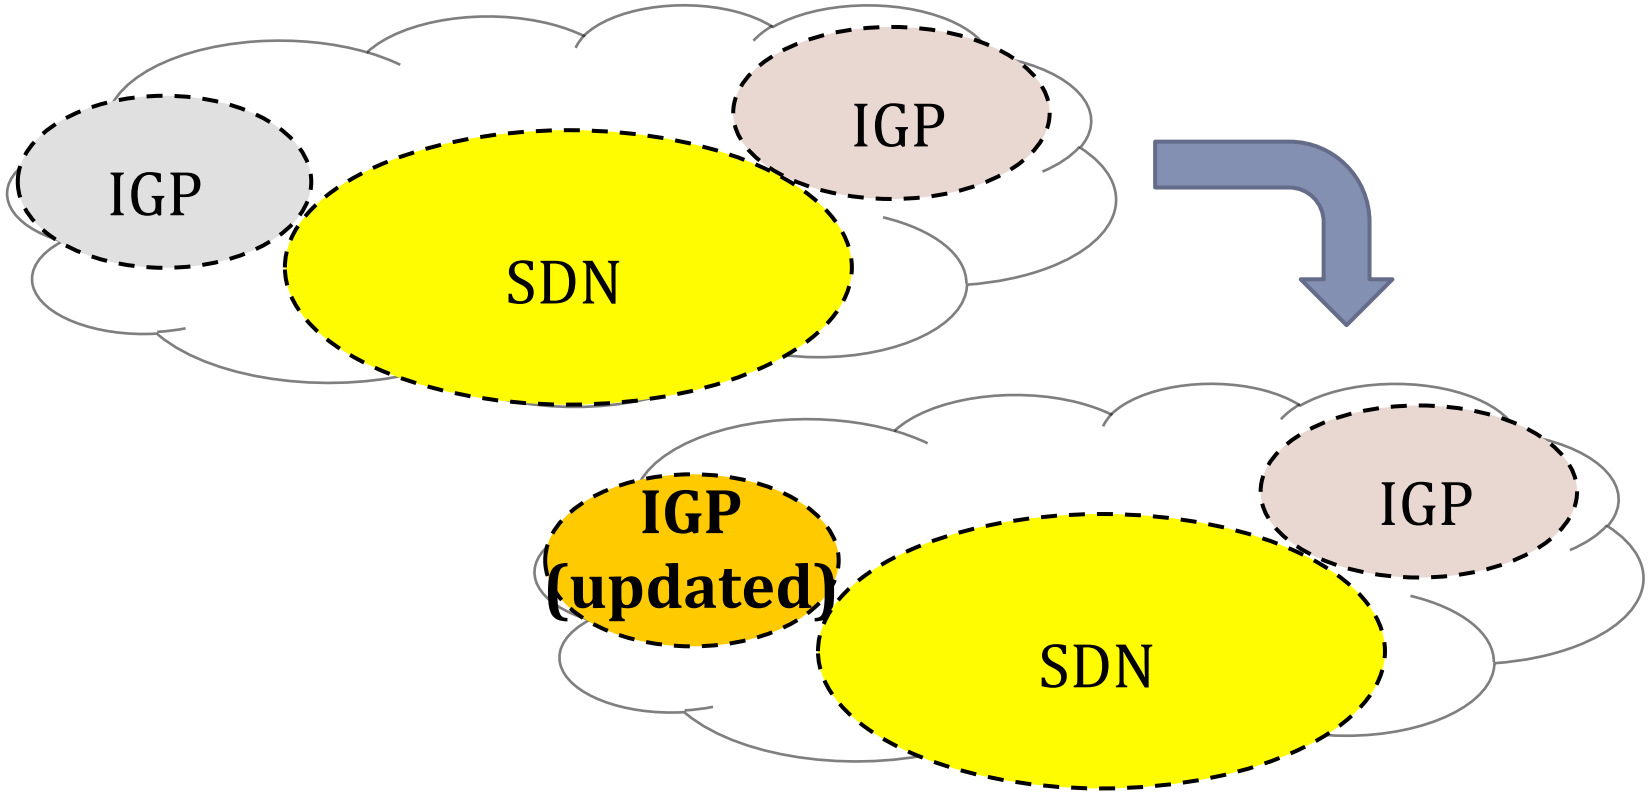

Single zone updates

- Does the presence of other zones impact the update of a single zone? ***YES, unfortunately***

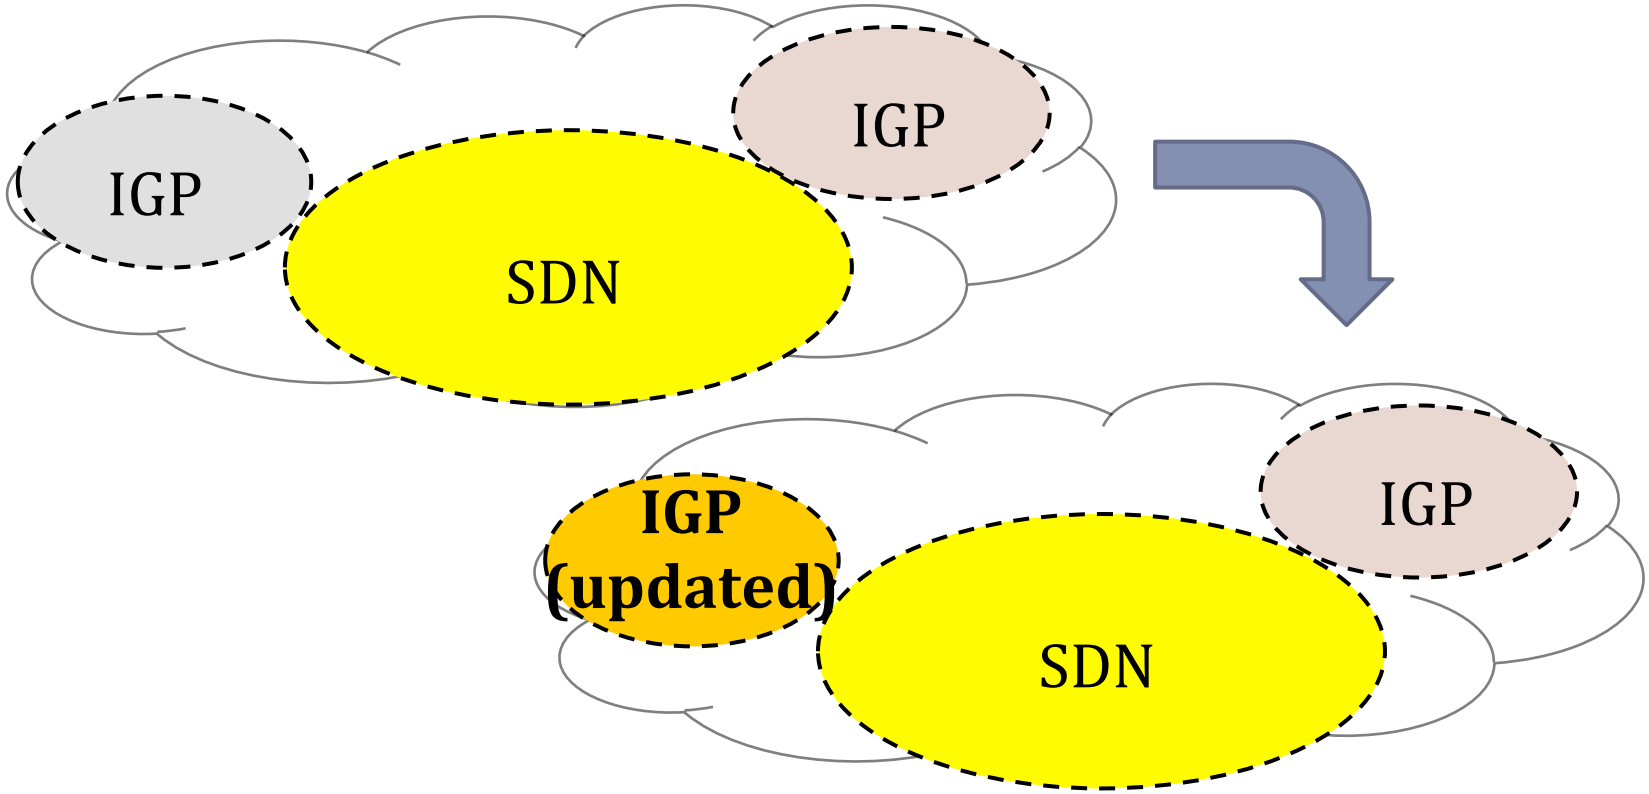

Safe techniques can be disrupted

- The update can affect forwarding in non-updated zones
  - multiple zones → conflicting preferences on the borders
  - redistributed routes → inconsistencies in non-updated zones
- Forwarding anomalies can occur in the update
  - i.e., forwarding loops → traffic losses
  - in remote zones or across different zones
  - with safe techniques that change protocol preferences (e.g., [Herrero10, Vanbever11])

Multi-zone updates

- How to perform updates spanning multiple zones?

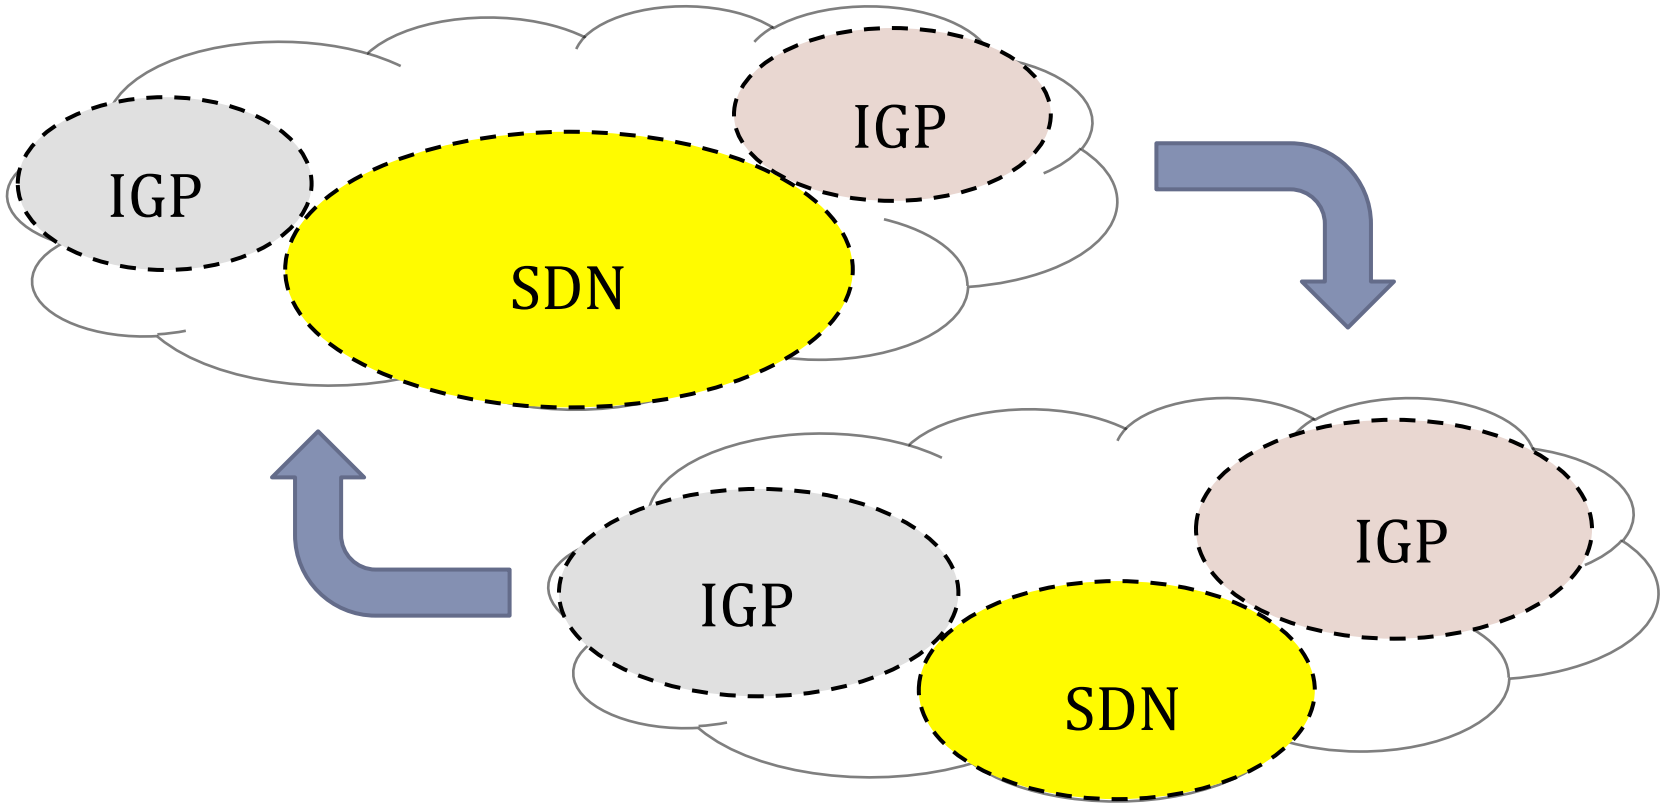

Known techniques are not easy to extend

- Additional update knobs
  - activating / de-activating route redistribution
- Additional interactions to be considered
  - route selection → route redistribution → route availability → route selection
- Intermediate forwarding paths
  - intermediate routes can leak from one zone to another

Safe updates are possible

- Generalized single-zone techniques
  - with additional constraints on update operations
- Multiple-zone procedures
  - based new sufficient conditions for route redistribution correctness
- Prototype system
  - which automates the procedures

Safe updates are possible

- Generalized single-zone techniques
  - with additional constraints on update operations
- Multiple-zone procedures
  - based new sufficient conditions for route redistribution correctness
- Prototype system
  - which automates the procedures

our techniques are provably correct and do not require to duplicate forwarding entries

No packet loss with our procedures

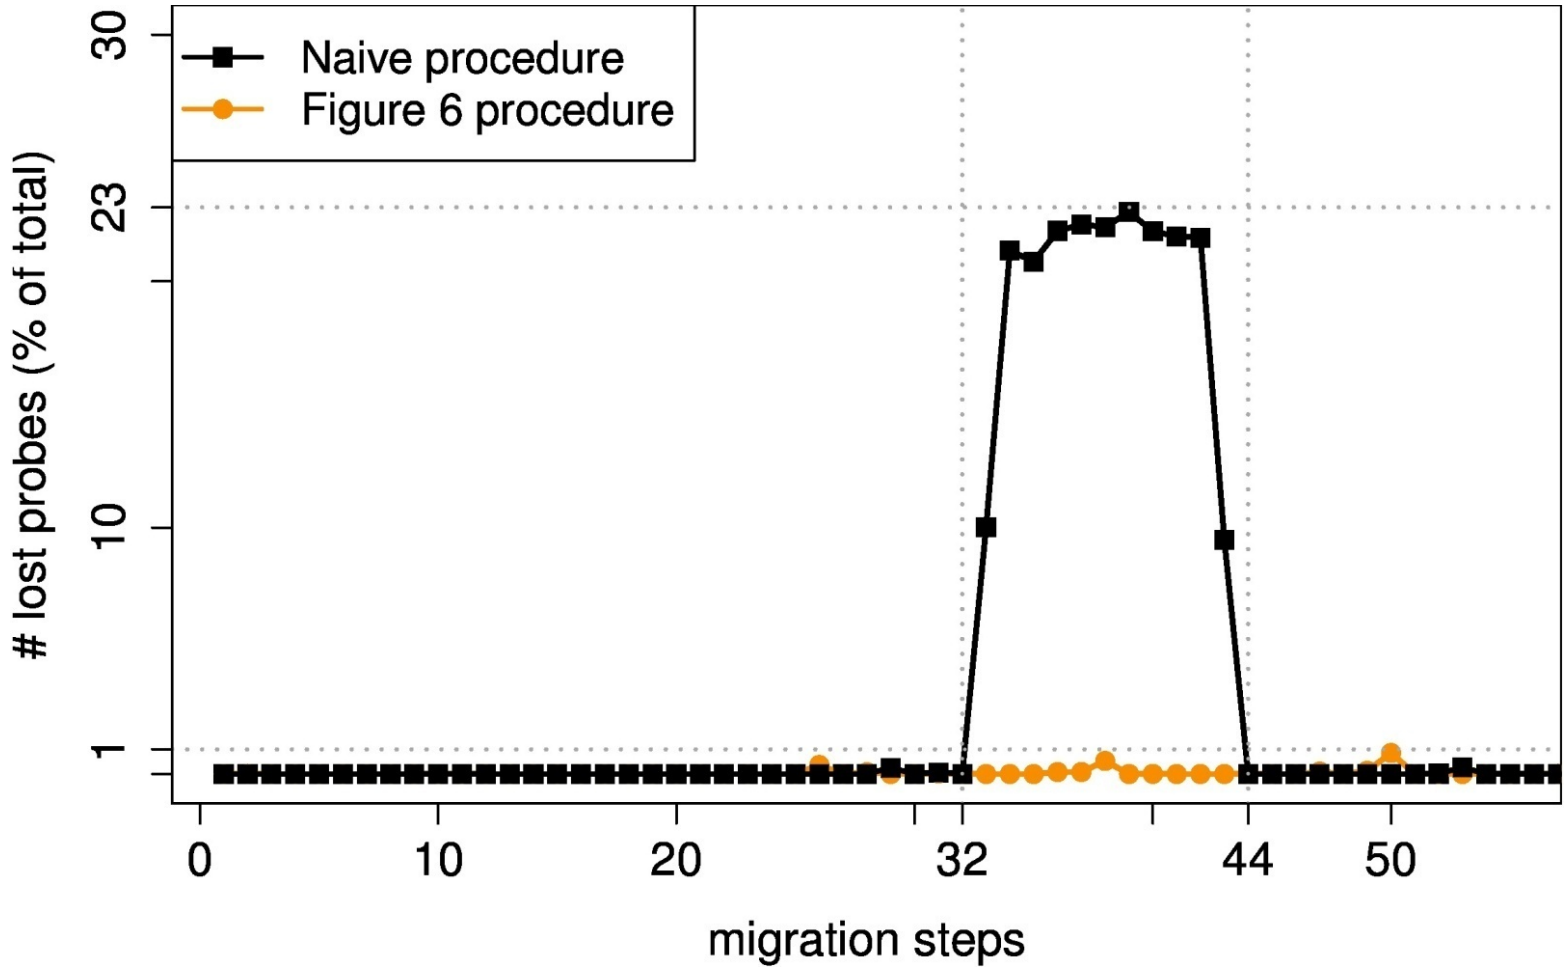

Service-Based (SB) Hybrid SDN

Pros and cons

- IGP and SDN can provide different services
- Automatic failure recovery for IGP controlled flows
- SDN controller off-loaded for non-critical flows
- …

- SDN capabilities are limited by the SDN devices
- Additional complexity of parallel paradigms
- …

Pros and cons

- IGP and SDN can provide different services
- Automatic failure recovery for IGP controlled flows
- SDN controller off-loaded for non-critical flows
- …

- SDN capabilities are limited by the SDN devices
- Additional complexity of parallel paradigms
  - **per-packet consistent update**
- …

Update of SB hybrid networks

- SDN-controlled flows ←→ IGP-controlled flows
  - e.g., because a service/dest becomes critical
  - e.g., in response to traffic matrix changes

- Inconsistencies can occur during the update
  - forwarding loops
  - intermediate forwarding paths
    - access policy violation, congestion, TCP issues, …

Safe updates are possible

- Mixed update approach
  - update sequencing algorithm + fallback solutions
  - to limit forwarding entry duplication
- Algorithm to compute maximal safe update sequences
  - typically, most routers can be updated w/o forwarding duplication
- Fallback solutions
  - applied only when a complete update sequence does not exist

Safe updates w/o forwarding doubling

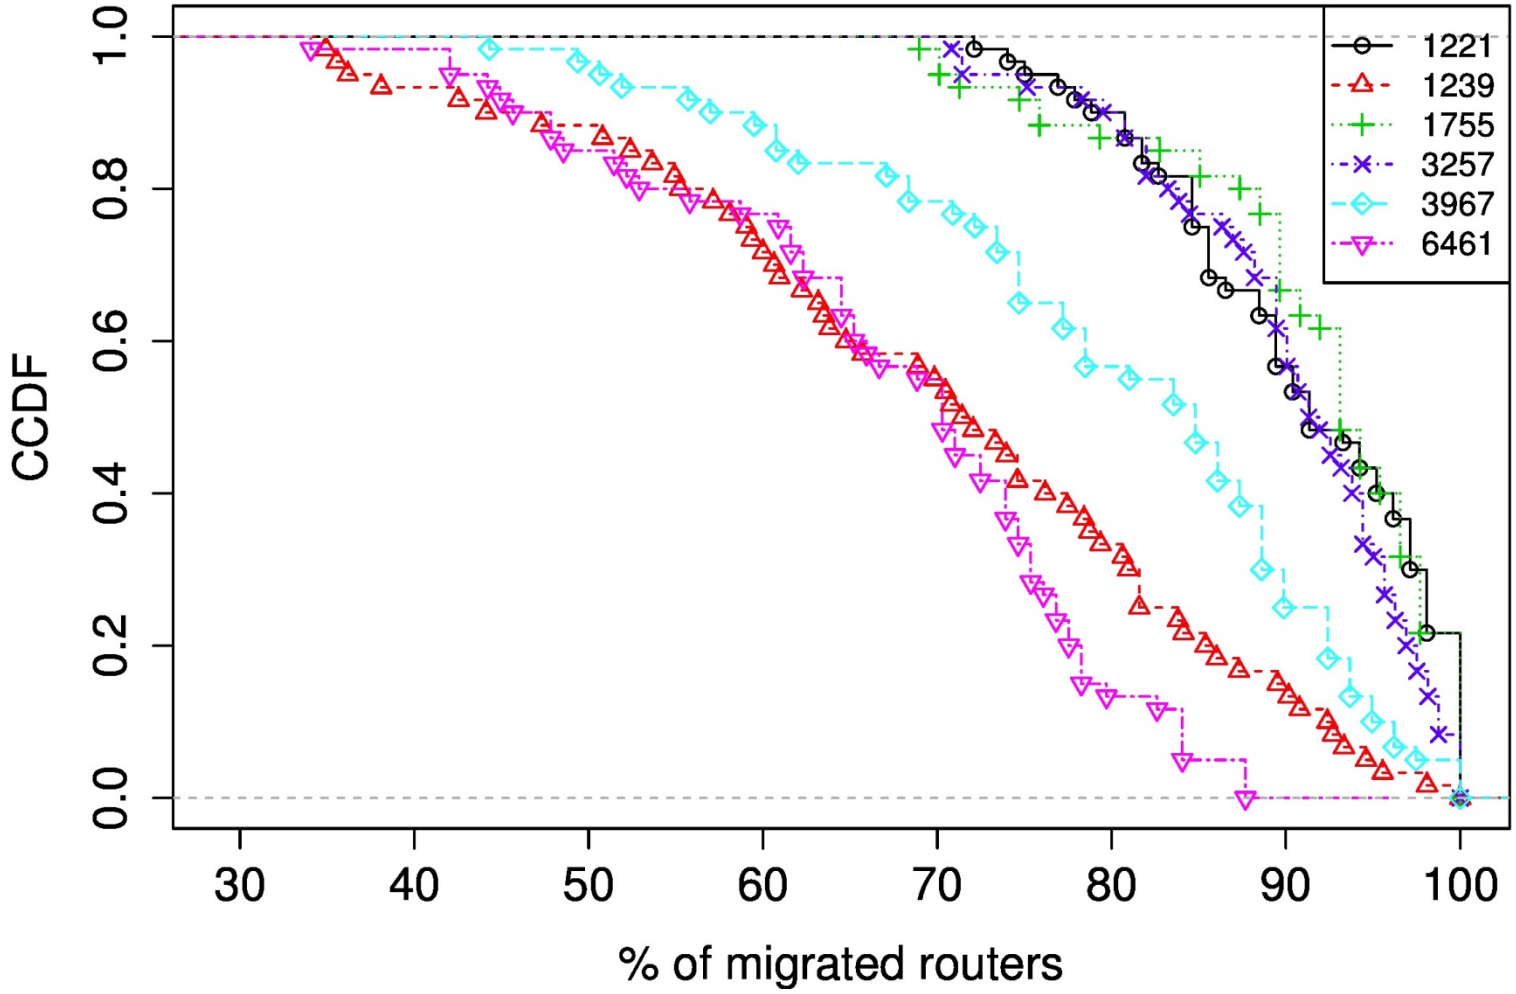

Pros and Cons

- Pros
  - Local SDN controllers can be distributed to different zones
  - Automatic failure recovery in the IGP zones
  - Improved scalability through aggregated route redistribution
  - …
- Cons
  - SDN capabilities are limited to SDN zones
  - additional complexity of route redistribution [Le07]
    - per-packet consistent updates
  - …

Basic Idea

- SDN and IGP control different traffic flows
  - on the same physical topology
- The controller installs static routes on IGP routers

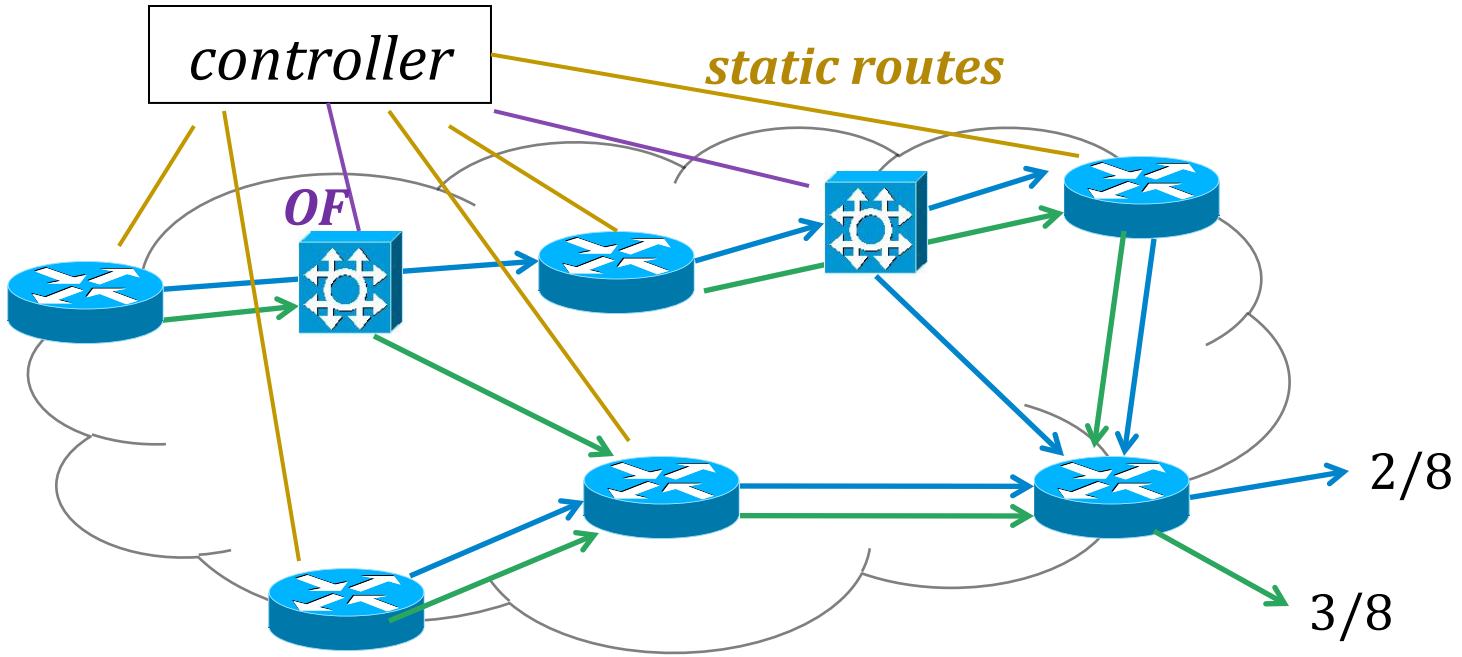

Pros and Cons

- Pros
  - IGP faking is powerful
    - we can modify the IGP graph for per-flow forwarding paths, traffic steering, load balancing on paths and on middle-boxes
    - except violations of the shortest path rules on chains of IGP routers
  - controlled reaction to failures
    - backup paths can be explicitly realized

- Cons
  - SDN capability depends on the deployed OF switches
    - traffic must flow through them
  - increased size of the new IGP graph

Integrated Hybrid SDN

Pros and cons

- IGP faking is powerful
  - per-flow traffic engineering, load balancing, middleboxing
  - controlled reaction to failures (e.g., backup paths)
- Safe update by tweaking the fake IGP graph
  - easily re-route many flows
- …

- SDN capability limited by SDN devices
- Increased size of the new IGP graph
- …

Pros and cons

- IGP faking is powerful
  - per-flow traffic engineering, load balancing, middleboxing
  - controlled reaction to failures (e.g., backup paths)
- Safe update by tweaking the fake IGP graph
  - easily re-route many flows
- …

- SDN capability limited by SDN devices
- **Increased size of the new IGP graph**
- …

IGP faking made practical

- Path requirement language
  - specifying the shape of {primary, ECMP, backup} paths
  - on a per-flow granularity
- LP algorithm to compute a minimal fake graph
- Prototype to realize path requirements
  - from the language to a minimal fake IGP graph

Take away

- Hybrid SDN can be an effective transitional mechanism
  - deployable today → reduces deployment costs
  - provides incentives for transition to SDN

- Hybrid SDN can be an interesting design point
  - hybridization sacrifices some SDN capabilities to mitigate some SDN limitations
    - it can combine the SDN innovation and flexibility with basic guarantees of distributed protocols
  - different hybridization models fit different networks

Comparison of Hybrid SDN Models

- Hybridization sacrifices some SDN capabilities to mitigate some SDN limitations
  - different models → different tradeoffs

	SDN capabilities	Network Update	Controller Scalability	Costs
TB H-SDN	local to SDN zones	carefully defined procedures	SDN spans a subset of devices	progressive with the number of SDN devices
SB H-SDN	depend on the SDN devices	carefully defined procedures	SDN controls a subset of flows	
Integrated	depend on the SDN devices	carefully defined fake topologies	proactive faking for given "events"	