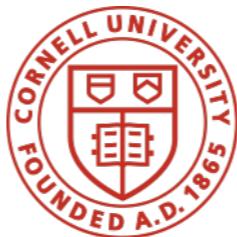


# Probabilistic Network Programming

Nate Foster  
Cornell



# Outline

## I: Introduction

- Semantics Primer
- Software-Defined Networking

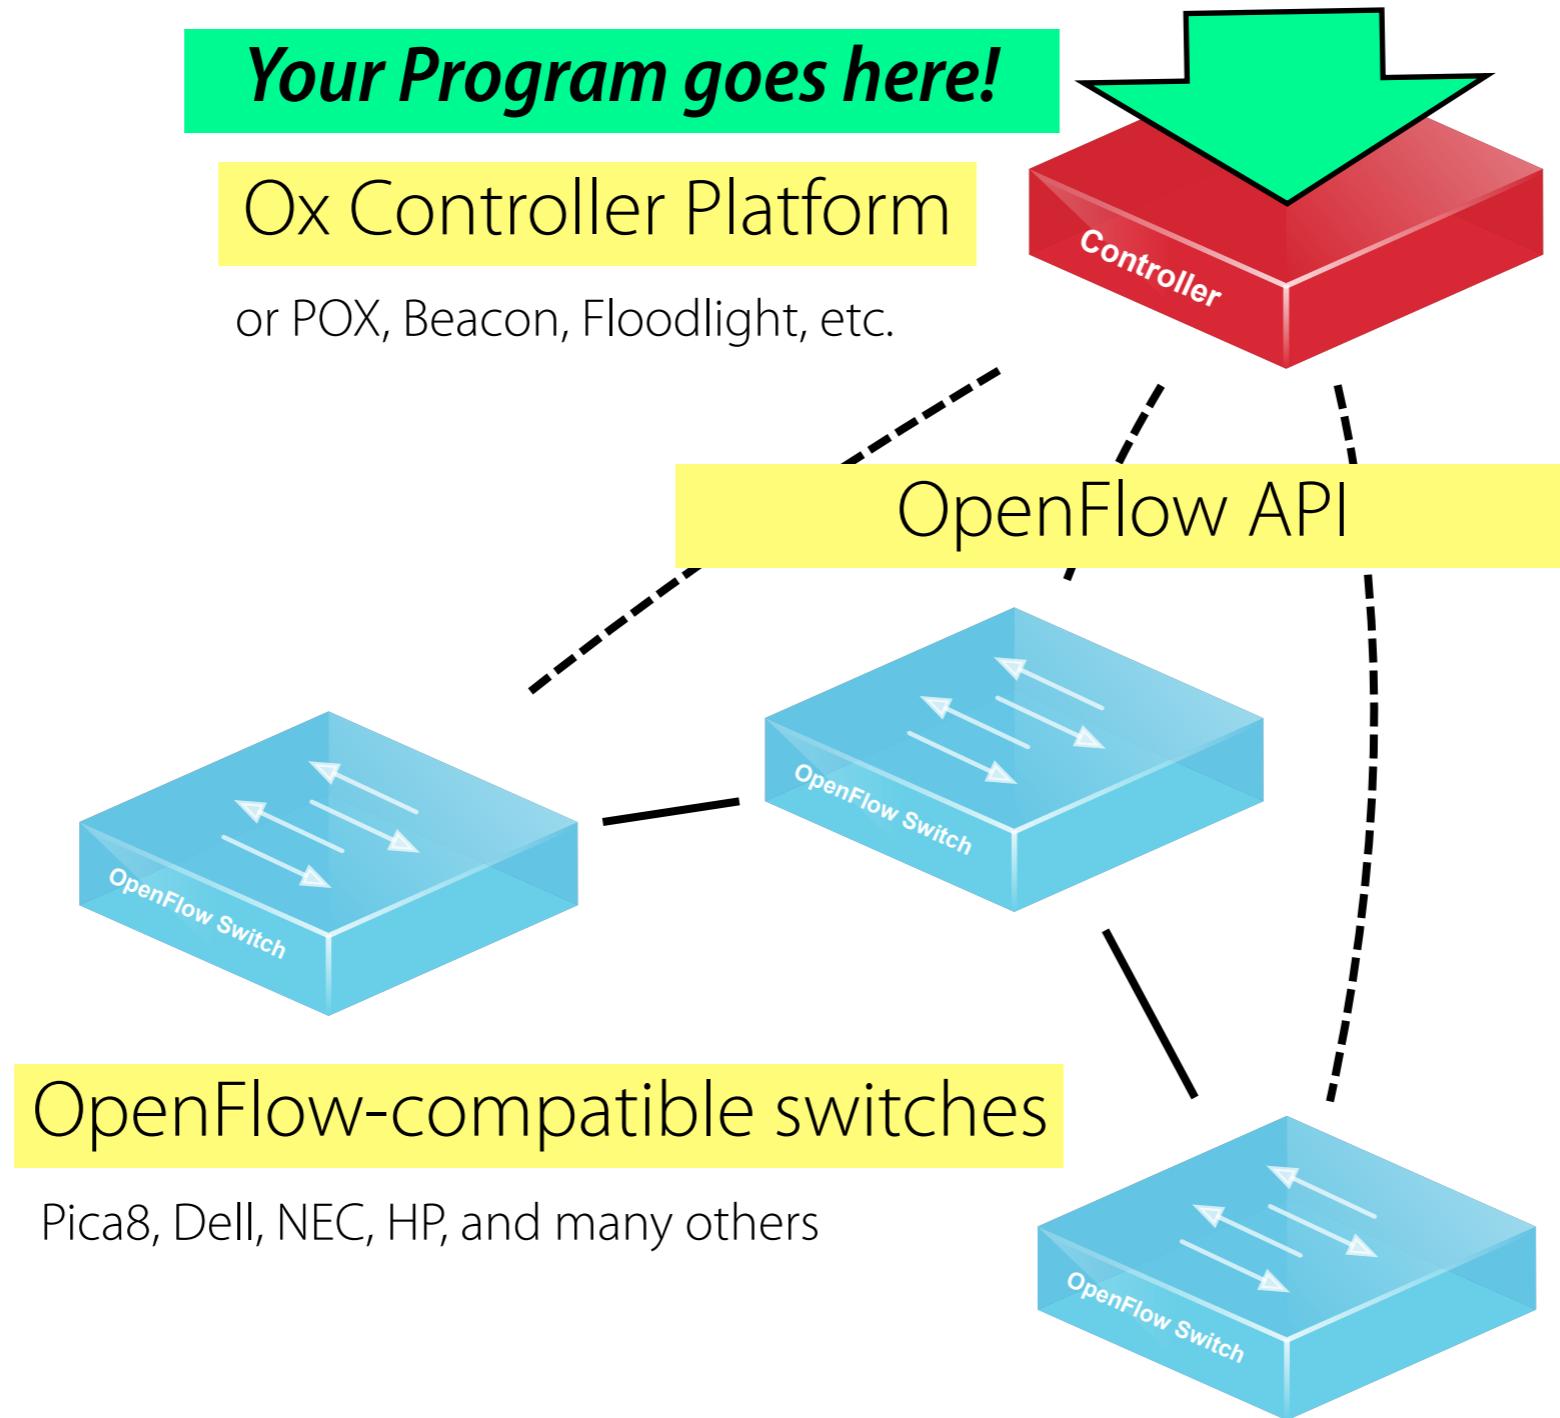
## II: NetKAT

- Language Design
- Formal Semantics

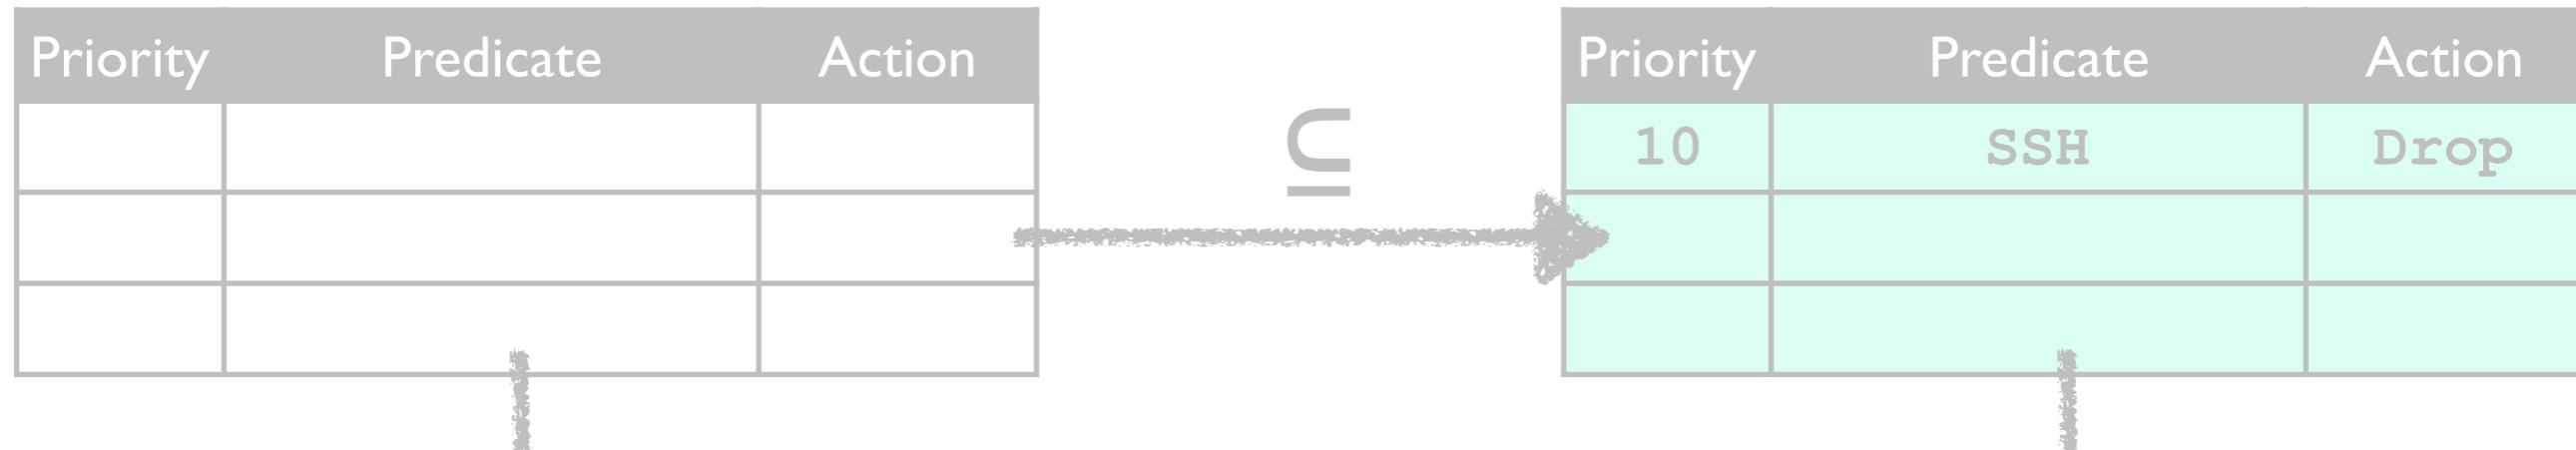
## III: Probabilistic NetKAT

- Formal Semantics
- Applications

# SDN Architecture



# Non-Atomic Updates



Must use a **barrier** to synchronize with the switch after installing the first rule...

5	dst_ip = H1	Fwd 1

*update re-ordering*

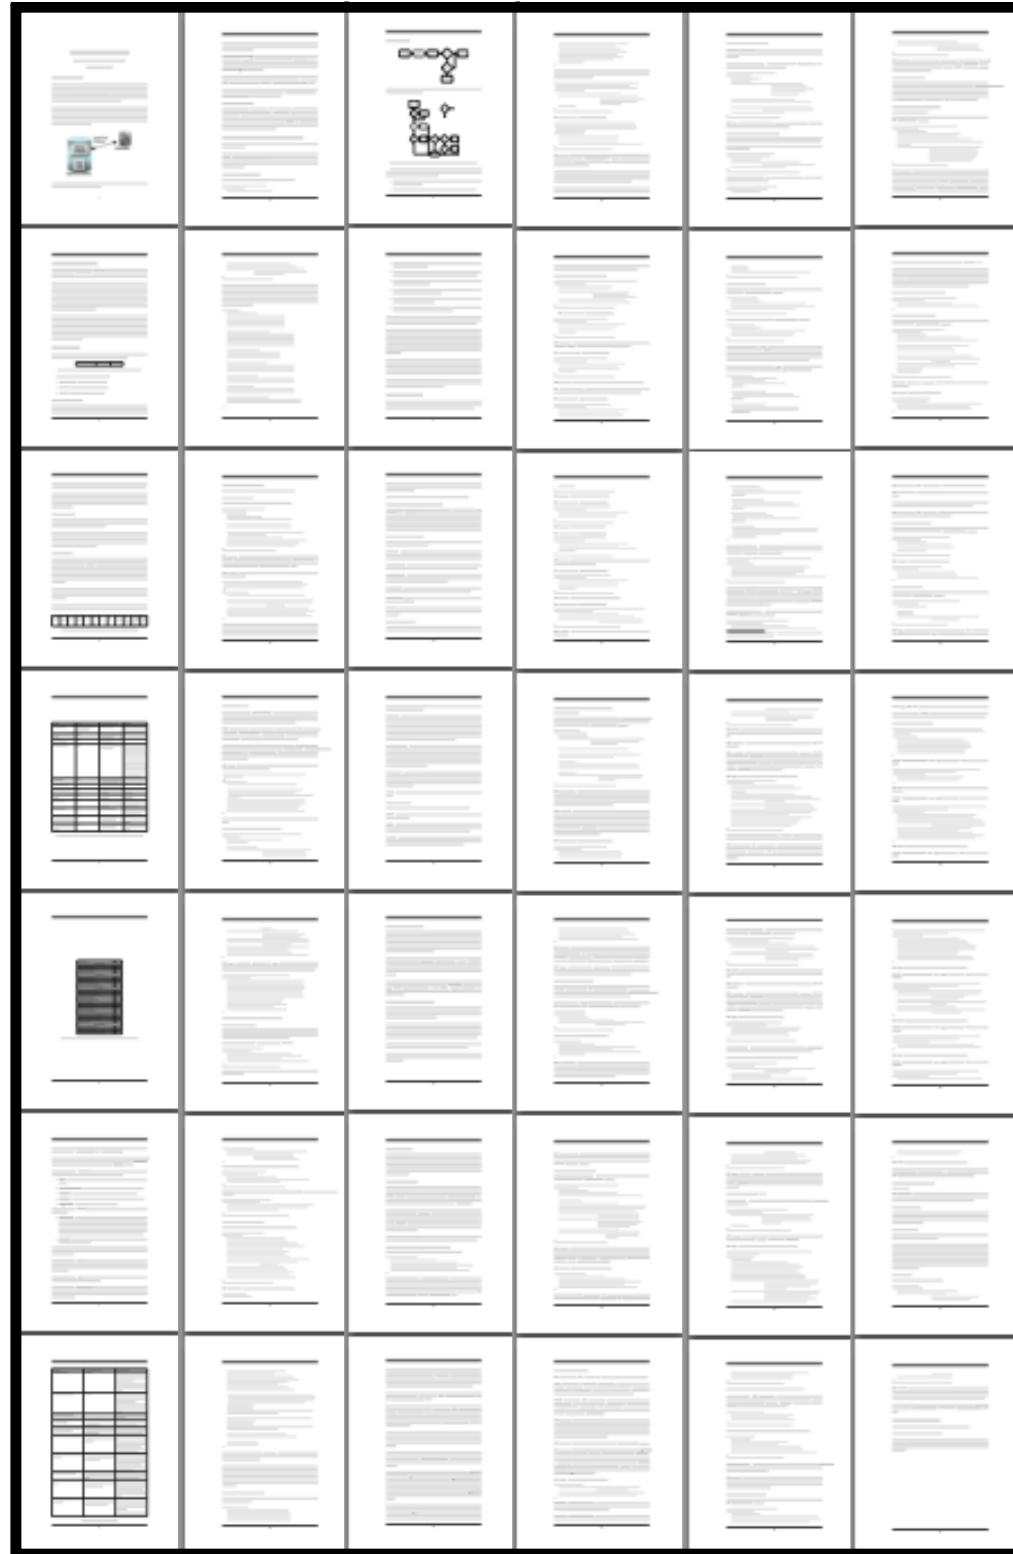
5	dst_ip = H1	Fwd 1



Priority	Predicate	Action
5	dst_ip = H1	Fwd 1
5	dst_ip = H2	Fwd 2

Priority	Predicate	Action
10	SSH	Drop
5	dst_ip = H1	Fwd 1
5	dst_ip = H2	Fwd 2

# OpenFlow Specification



42 pages...

...of informal prose

...diagrams and flow charts

...and C struct definitions

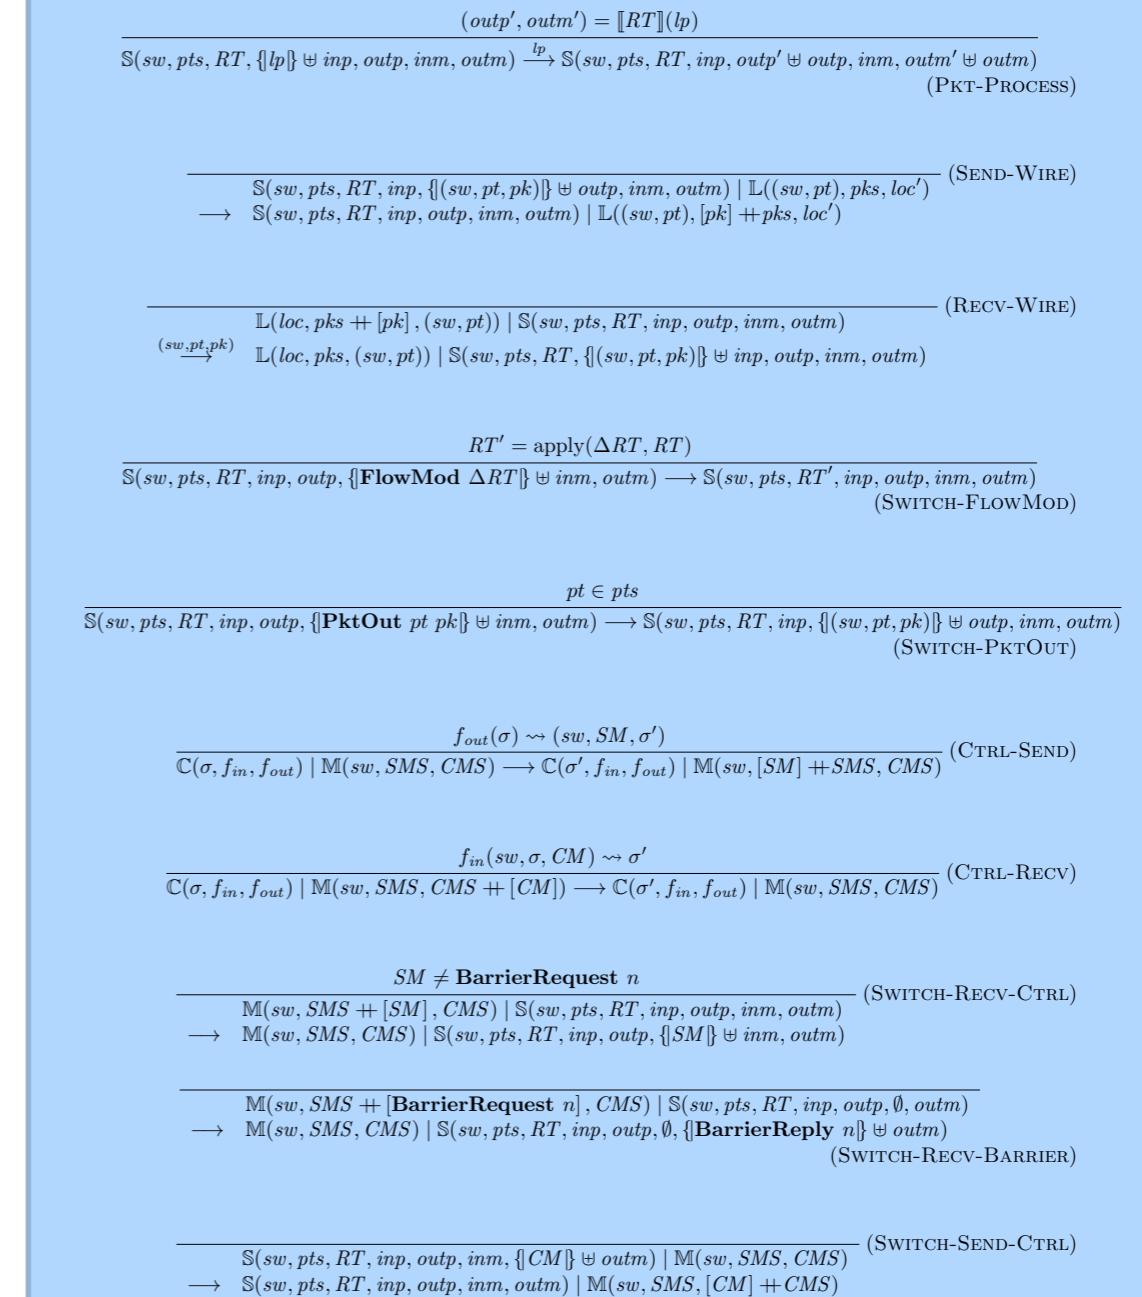
# Featherweight OpenFlow

## Syntax

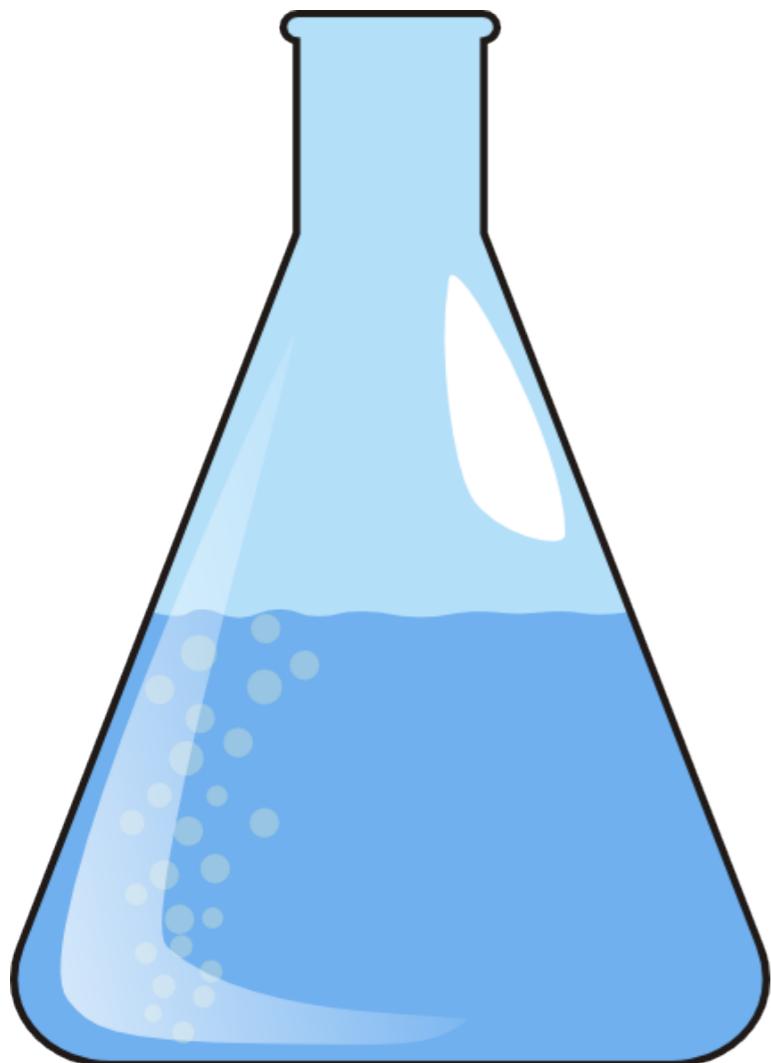
<b>Devices</b>	Switch	$S ::= \mathbb{S}(sw, pts, RT, inp.outp, inm, out)$
	Controller	$C ::= \mathbb{C}(\sigma, f_{in}, f_{out})$
	Link	$L ::= \mathbb{L}(loc_{src}, pks, loc_{dst})$
	OpenFlow Link to Controller	$M ::= \mathbb{M}(sw, SMS, CMS)$
<b>Packets and Locations</b>	Packet	$pk ::= abstract$
	Switch ID	$sw \in \mathbb{N}$
	Port ID	$pt \in \mathbb{N}$
	Location	$loc \in sw \times pt$
	Located Packet	$lp \in loc \times pk$
<b>Controller Components</b>	Controller state	$\sigma ::= abstract$
	Controller input relation	$f_{in} \in sw \times CM \times \sigma \rightsquigarrow \sigma$
	Controller output relation	$f_{out} \in \sigma \rightsquigarrow sw \times SM \times \sigma$
<b>Switch Components</b>	Rule table	$RT ::= abstract$
	Rule table Interpretation	$\llbracket RT \rrbracket \in lp \rightarrow \{lp_1 \dots lp_n\} \times \{CM_1 \dots C$
	Rule table modifier	$\Delta RT ::= abstract$
	Rule table modifier interpretation	$apply \in \Delta RT \rightarrow RT \rightarrow \Delta RT$
	Ports on switch	$pts \in \{pt_1 \dots pt_n\}$
	Consumed packets	$inp \in \{lp_1 \dots lp_n\}$
	Produced packets	$outp \in \{lp_1 \dots lp_n\}$
	Messages from controller	$inm \in \{SM_1 \dots SM_n\}$
	Messages to controller	$outm \in \{CM_1 \dots CM_n\}$
<b>Link Components</b>	Endpoints	$loc_{src}, loc_{dst} \in loc \text{ where } loc_{src} \neq loc_{dst}$
	Packets from $loc_{src}$ to $loc_{dst}$	$pks \in [pk_1 \dots pk_n]$
<b>Controller Link</b>	Message queue from controller	$SMS \in [SM_1 \dots SM_n]$
	Message queue to controller	$CMS \in [CM_1 \dots CM_n]$
<b>Abstract OpenFlow Protocol</b>	Message from controller	$SM ::= \text{FlowMod } \Delta RT \mid \text{PktOut } pt \_ i$
	Message to controller	$CM ::= \text{PktIn } pt \_ pk \mid \text{BarrierReply } n$

Models all features related  
to packet forwarding, and  
*all* essential concurrency

## Semantics

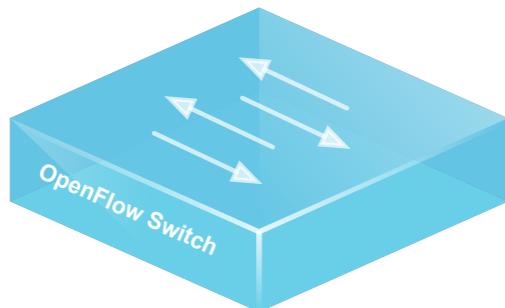


# Chemical Abstract Machine



- A general technique for modeling concurrent systems using operational semantics
- Proposed by Berry and Boudol [POPL '90]
- Intuition: computational components are “molecules” in a “solution” that can react with each other non-deterministically

# OpenFlow Molecules



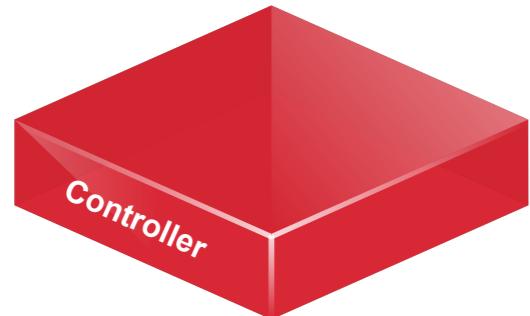
## Switches

- $S(sw, T, \text{in}, \text{out}, \text{msgs})$
- sw: unique identifier
- T: flow table
- in/out: buffered packets
- msgs: control messages

## Data Links

- $L(sw, pt, pkts, sw', pt')$
- sw/pt: source
- sw'/pt': target
- pkts: queued packets

# OpenFlow Molecules



# Controller

- $\mathbb{C}(\sigma, f_{in}, f_{out})$
  - $\sigma$ : state
  - $f_{in}$ : ingress function
  - $f_{out}$ : output function

- $M(sw, msgs)$

- msgs: control messages

## Control Links

# OpenFlow Messages

$m ::= \text{FlowMod } f$  (\* switch-controller messages \*)

| BarrierReq n

| Barrier Rep n

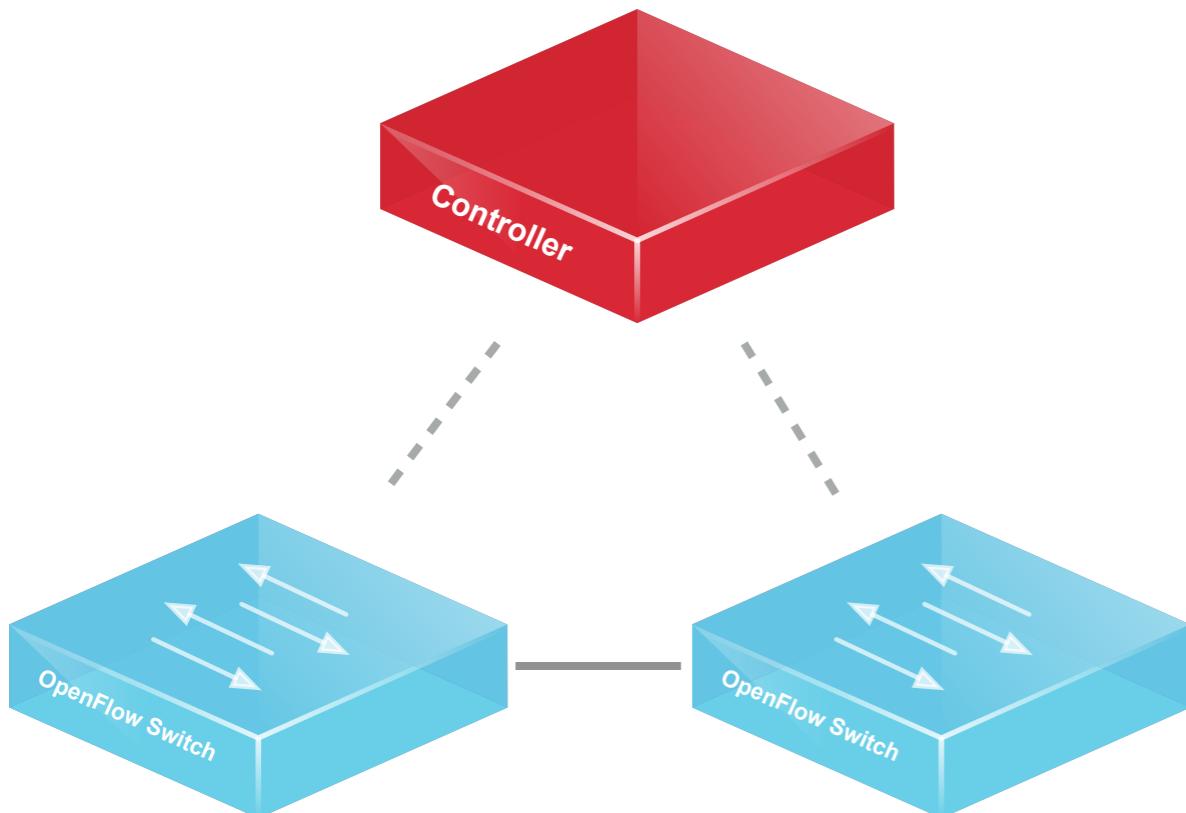
| PktIn pk pt

| PktOut pk pt

$f ::= \text{Add}(n, pat, act)$  (\* table manipulation \*)

| Del pat

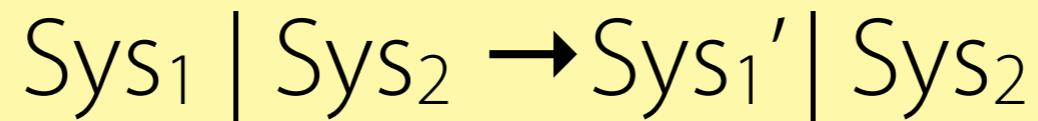
# OpenFlow Network



To model an OpenFlow network in we use:

- One  $C$  molecule
- An  $M$  molecule per switch
- An  $S$  molecule per switch
- An  $L$  molecule per link
- The overall model is a concurrent composition “|” of these molecules

# Reaction



**Intuition:** if any sub-collection of molecules in the system react with each other, then the whole system may also exhibit the same reaction

# Structural Congruence

$$\text{Sys}_1 \equiv \text{Sys}'_1 \quad \text{Sys}'_1 \rightarrow \text{Sys}'_2 \quad \text{Sys}'_2 \equiv \text{Sys}_2$$

---

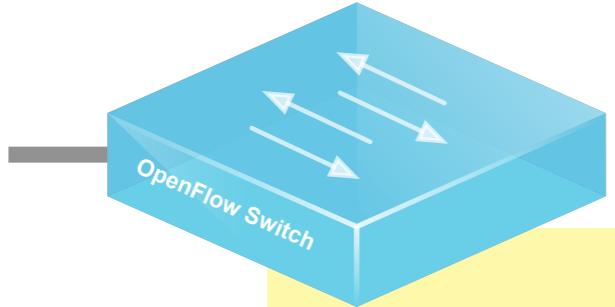
$$\text{Sys}_1 \rightarrow \text{Sys}_2$$

$$\text{Sys}_1 | \text{Sys}_2 \equiv \text{Sys}_2 | \text{Sys}_1$$

$$\text{Sys}_1 | (\text{Sys}_2 | \text{Sys}_3) \equiv (\text{Sys}_1 | \text{Sys}_2) | \text{Sys}_3$$

**Intuition:** concurrent composition ("|") is both associative and commutative

# Switch Ingress

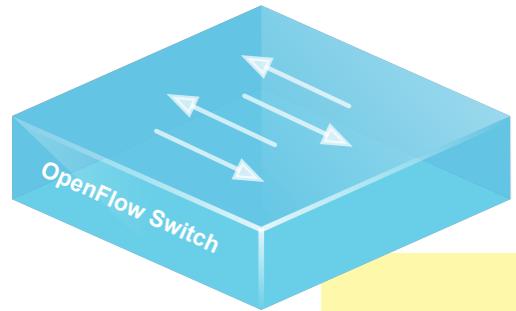

$$\mathbb{L}(\text{sw}, \text{pt}, \text{pkts} @ [\text{pkt}], \text{sw}', \text{pt}') \mid \\ \mathbb{S}(\text{sw}', \text{T}, \text{in}, \text{out}, \text{msgs})$$

→

$$\mathbb{L}(\text{sw}, \text{pt}, \text{pkts}, \text{sw}', \text{pt}') \mid \\ \mathbb{S}(\text{sw}', \text{T}, \text{in} \cup \{(\text{pkt}, \text{pt}')\}, \text{out}, \text{msgs})$$

**Intuition:** a switch can input a packet from an adjacent link into its ingress buffer

# Forwarding



$$[\![T]\!] \ (pkt, pt) = (out', msgs')$$

---

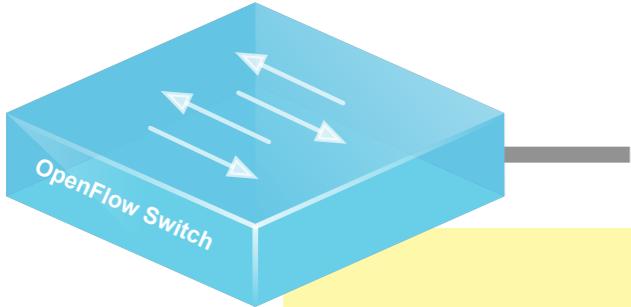
$$\mathbb{S}(sw, T, in \cup \{(pkt, pt)\}, out, msgs)$$

→

$$\mathbb{S}(sw, T, in, out \cup out', msgs \cup msgs')$$

**Intuition:** a switch processes a single packet using its flow table, which generates zero or more output packets and zero or more control messages

# Switch Egress

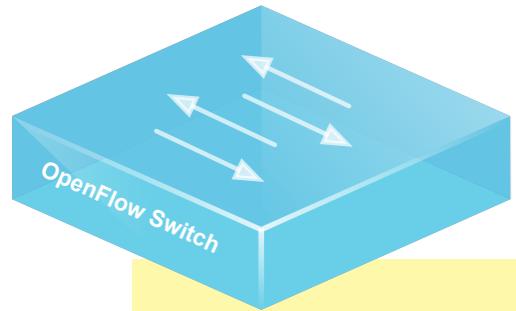

$$\mathcal{S}(\text{sw}, \text{T}, \text{in}, \text{out} \cup \{(\text{pkt}, \text{pt})\}, \text{msgs}) \mid \\ \mathcal{L}(\text{sw}, \text{pt}, \text{pkts}, \text{sw}', \text{pt}')$$

→

$$\mathcal{S}(\text{sw}, \text{T}, \text{in}, \text{out}, \text{msgs}) \mid \\ \mathcal{L}(\text{sw}, \text{pt}, \text{pkt} :: \text{pkts}, \text{sw}', \text{pt}')$$

**Intuition:** a switch can output a packet from its output buffer onto the adjacent link

# Add



$\llbracket \text{Add}(n, \text{pat}, \text{act}) \rrbracket \quad T = T'$

---

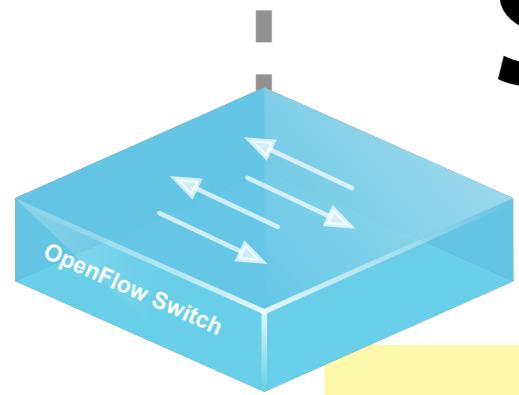
$S(sw, T, \text{in}, \text{out}, \text{msgs} \cup \{\text{FlowMod } (\text{Add } (n, \text{pat}, \text{act}))\})$

→

$S(sw, T', \text{in}, \text{out}, \text{msgs})$

**Intuition:** a switch may modify its flow table by adding a rule sent previously by the controller...

# Switch-to-Controller



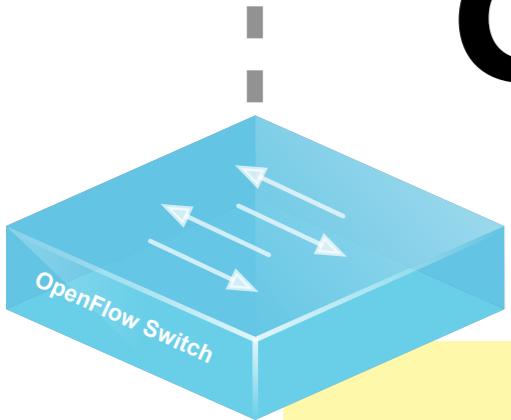
$S(sw, T, \text{in}, \text{out}, \text{msgs} \cup \{m\}) \mid$   
 $M(sw, \text{msgs})$

→

$S(sw, T, \text{in}, \text{out}, \text{msgs}) \mid$   
 $M(sw, m:: \text{msgs})$

**Intuition:** a switch can send a message to the controller by moving it from its internal buffer to the shared queue in the controller link

# Controller-to-Switch



$m \neq \text{BarrierReq } n$

---

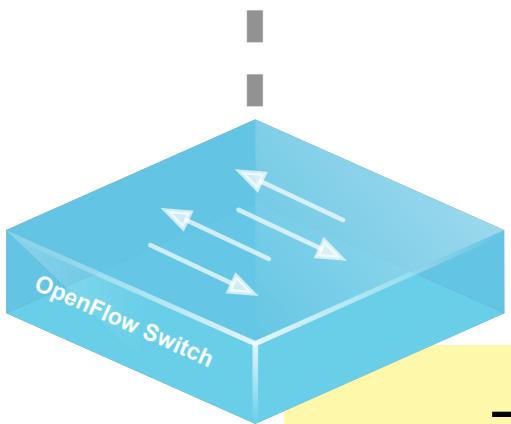
$M(\text{sw}, \text{msgs} @ [m]) \mid$   
 $S(\text{sw}, T, \text{in}, \text{out}, \text{msgs})$



$M(\text{sw}, \text{msgs}) \mid$   
 $S(\text{sw}, T, \text{in}, \text{out}, \text{msgs} \cup \{m\})$

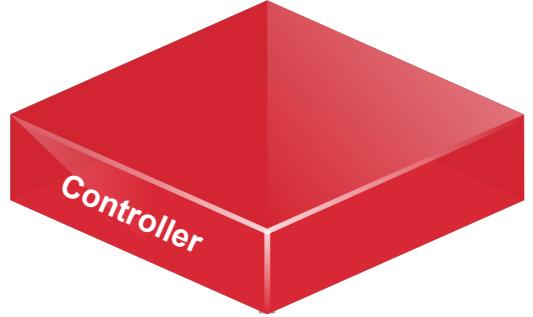
**Intuition:** a switch can receive a controller message by moving it from the controller link to its internal buffer

# Barrier


$$\mathbb{M}(\text{sw}, @[\text{BarrierReq } n]) \mid \\ \mathbb{S}(\text{sw}, T, \text{in}, \text{out}, \{\})$$

$$\mathbb{M}(\text{sw}, \text{msgs}) \mid \\ \mathbb{S}(\text{sw}, T, \text{in}, \text{out} \cup \{\text{BarrierRep } n\}, \{\})$$

**Intuition:** a switch can receive a barrier request message *only* if it has processed all outstanding locally-buffered messages



# Controller Input

:

:

$$f_{in}(\sigma, sw, m) = \sigma'$$

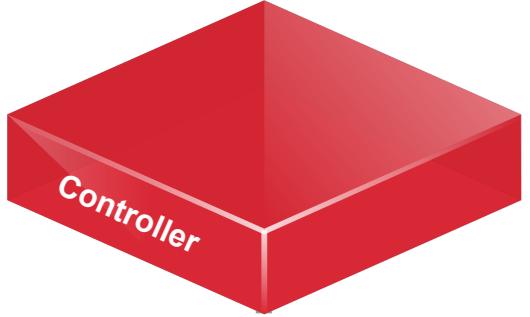
---

$$\begin{aligned} & M(sw, msgs @ [m]) \mid \\ & C(\sigma, f_{in}, f_{out}) \end{aligned}$$

→

$$\begin{aligned} & M(sw, msgs) \mid \\ & C(\sigma', f_{in}, f_{out}) \end{aligned}$$

**Intuition:** the controller processes a message received from a switch by running its input function



# Controller Output

:

:

$$f_{\text{out}}(\sigma) = (\sigma', \text{sw}, m)$$

---

$M(\text{sw}, \text{msgs}) \mid$

$C(\sigma, f_{\text{in}}, f_{\text{out}})$

→

$M(\text{sw}, m:\text{msgs}) \mid$

$C(\sigma', f_{\text{in}}, f_{\text{out}})$

**Intuition:** the controller may spontaneously execute its output function

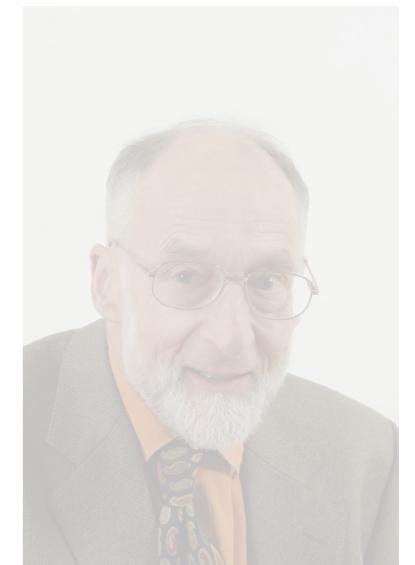
# Equivalence

$$(H_1, \text{✉}) \longrightarrow (S_1, pt_1, \text{✉}) \longrightarrow (H_2, \text{✉})$$

Given a specification of a packet-processing function, how can we establish that an OpenFlow implementation is correct?

It turns out that simple trace equivalence won't do...

# Which Equivalence?



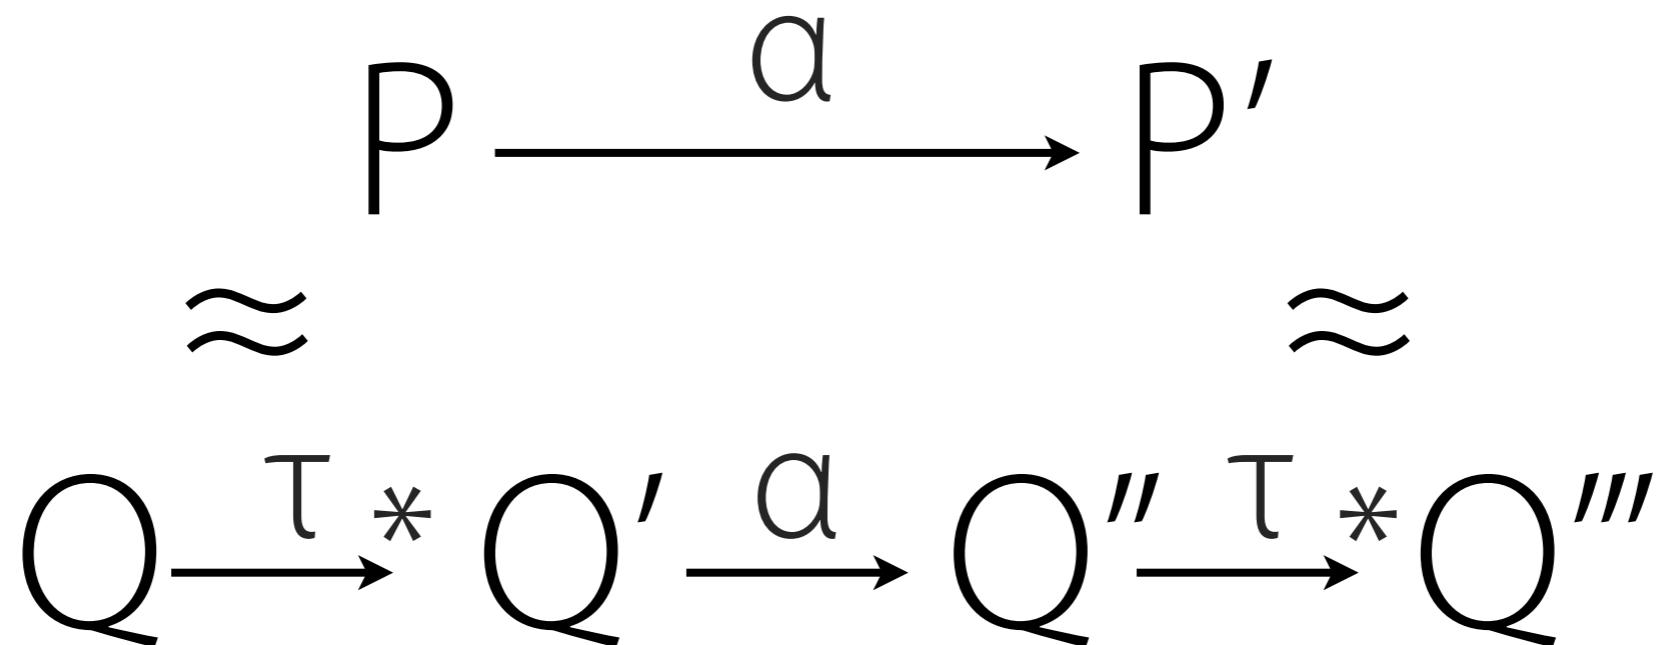
Clearly we need a finer notion of equivalence! That is, one that distinguishes between more behaviors



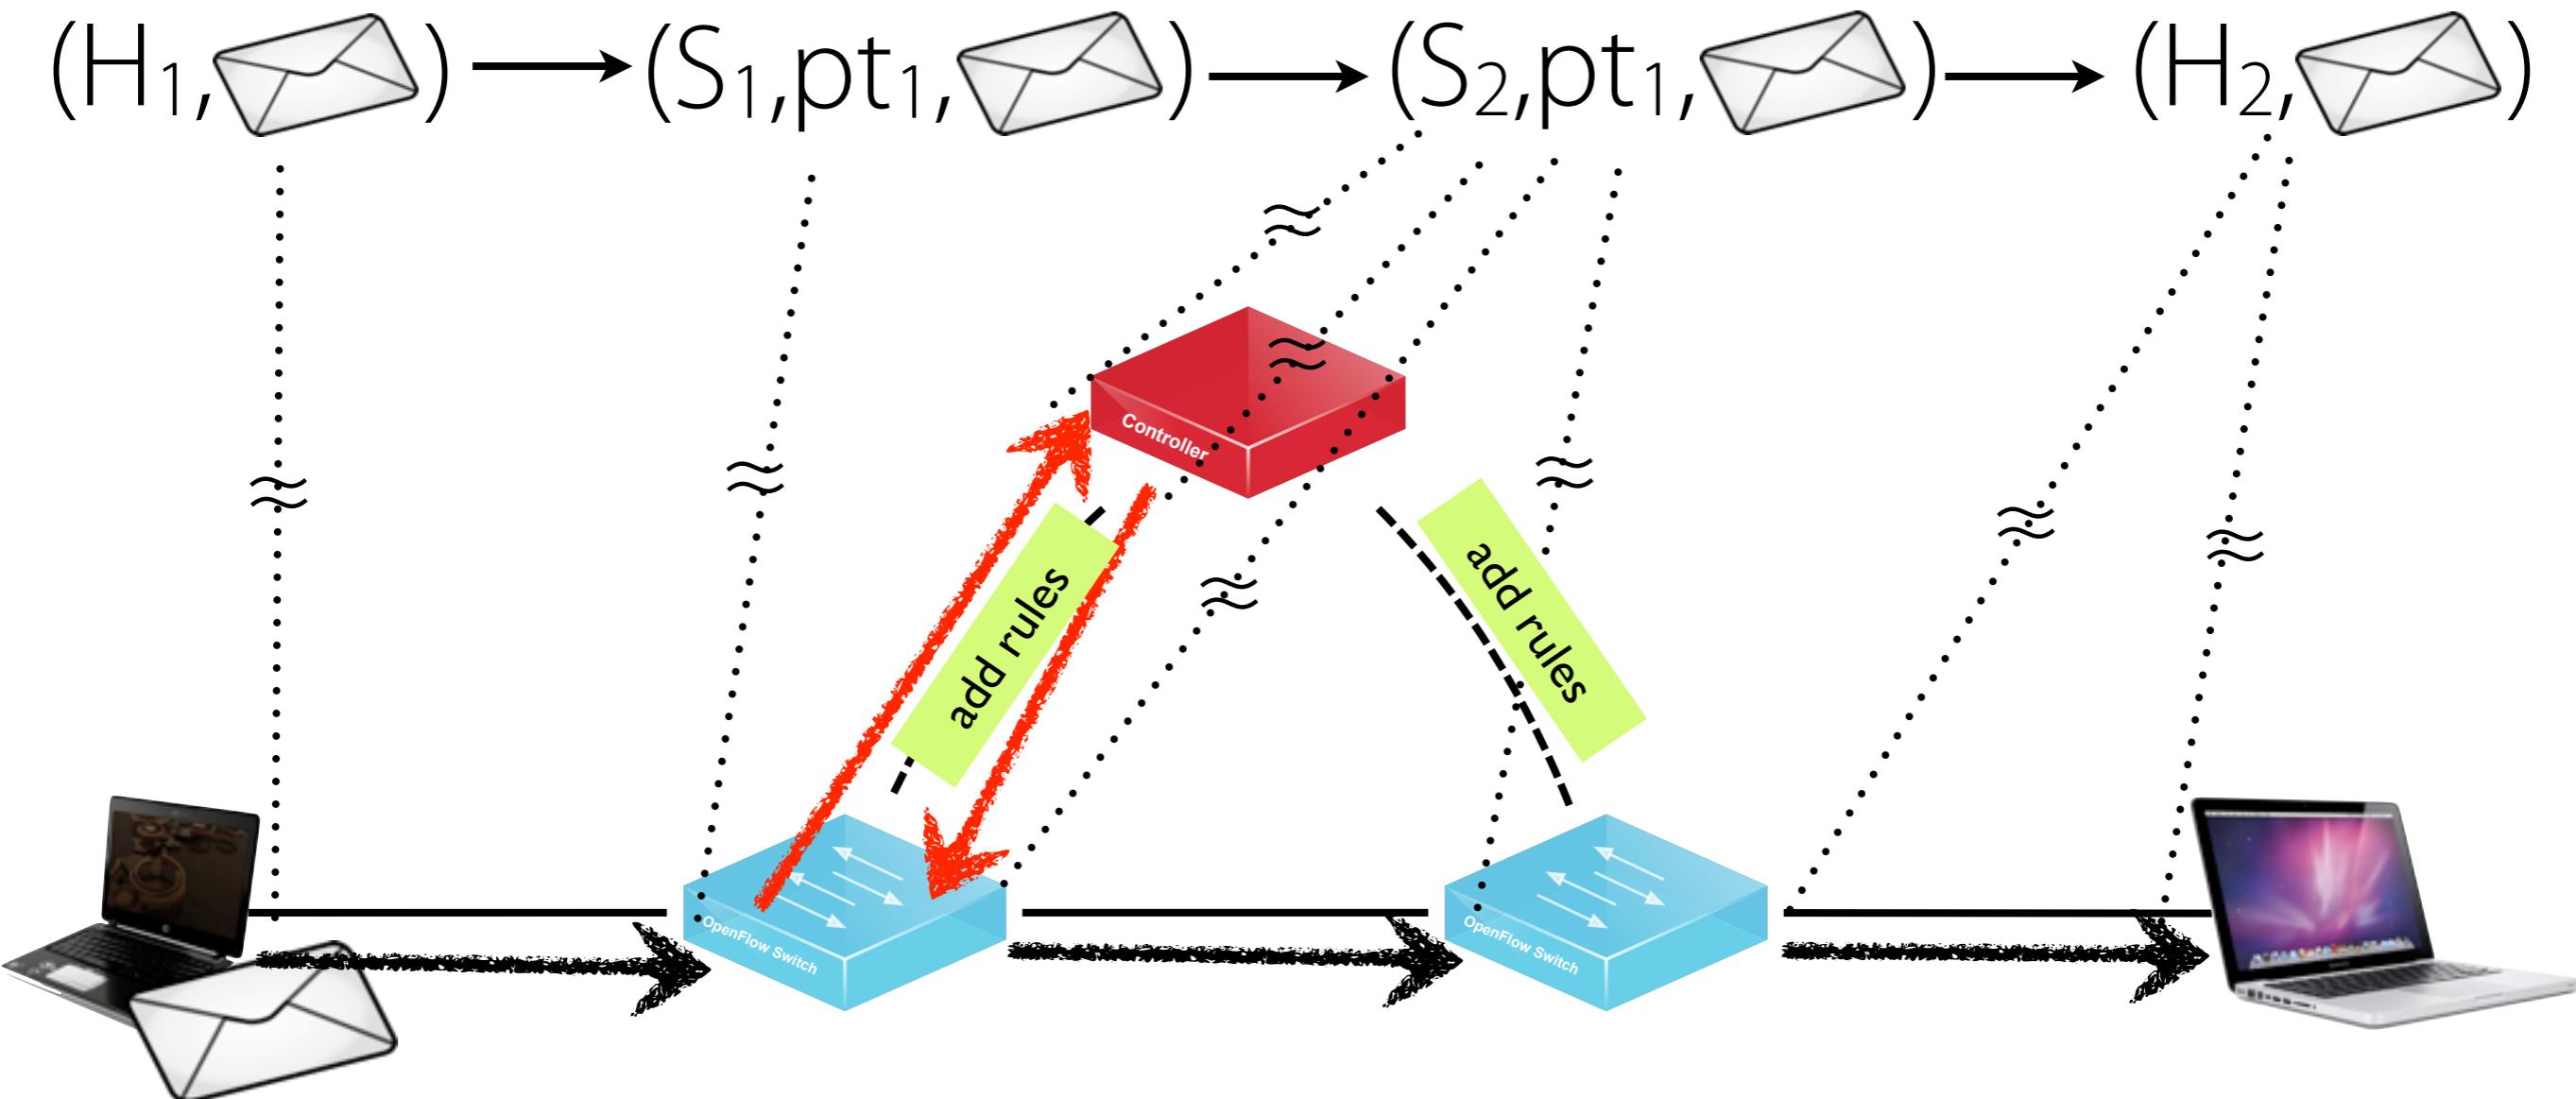
$$L(A) = L(B) \{ \$0.50 \bullet \text{Tea}, \$0.50 \bullet \text{Coffee} \}$$

# Weak Bisimulation

- **a:** observable actions such as the receipt of a packet at a given switch and port
- **τ:** “internal” actions such as a switch moving a packet into a buffer



# Weak Bisimulation



# Parameterized Weak Bisimulation

## Invariants

- *Safety*: at all times, the rules installed on switches are a *subset* of the controller function
- *Liveness*: the controller eventually processes all packets diverted to it by switches

## Theorem

```
Module RelationDefinitions :=  
  FwOF.FwOFRelationDefinitions.Make (AtomsAndController).  
  ...  
Theorem fwof_abst_weak_bisim :  
  weak_bisimulation  
  concreteStep  
  abstractStep  
  bisim_relation.
```

# Coq Implementation



```
Inductive pred : Type :=
| OnSwitch : Switch -> pred
| TDepend : Depend -> pred

Lemma inter_wildcard_other : forall x,
  Wildcard inter WildcardAll x = x
Proof
  int
  Qed.

Inductive Wildcard : Type :=
| WildcardAll
| WildcardOther

Lemma nettleServer :: ControllerRec -> IO ()
nettleServer controller = do
  nettle <- startOpenFlowServer Nothing 6633
  switchMsgs <- newChan
  forkIO (handleOFMsgs controller switchMsgs
    nettle)
  forever $ do
    (switch, switchFeatures) <- retryOnExns
    "nettle bug" (acceptSwitch nettle)
    writeChan switchMsgs (Left $ toInteger $
      handle2SwitchID switch)
    hPutStrLn stderr ("switch: " ++ (show
      (handle2SwitchID switch)))
    hFlush stderr
    return ()
  forkIO (handleSwitch switch switchMsgs)
  closeServer nettle
  
```



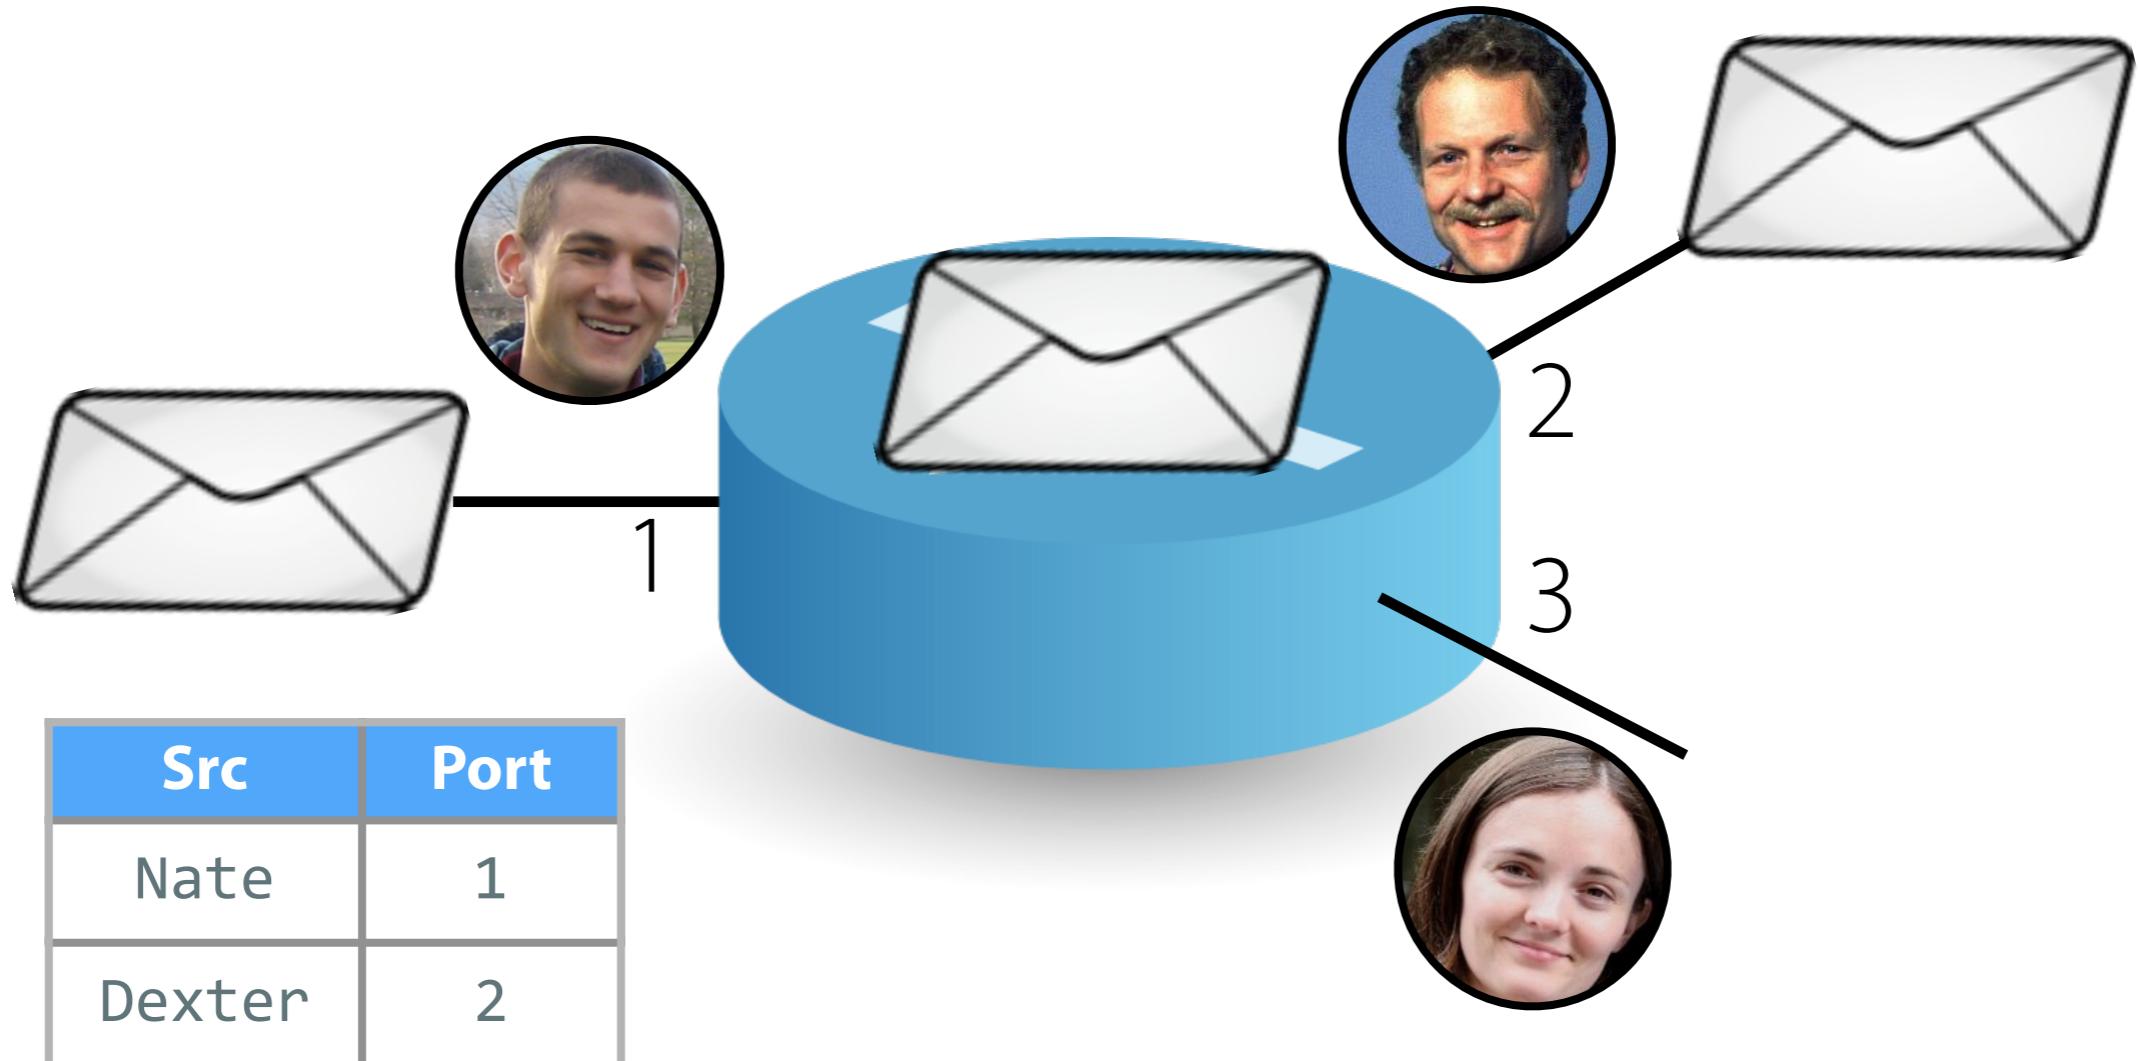
Certified  
executable

Develop, prove, extract

[github.com/frenetic-lang/featherweight-openflow](https://github.com/frenetic-lang/featherweight-openflow)

# **NetKAT Language**

# Learning Switch



- Learn locations of hosts in a table
- Flood packets going to unknown hosts; forward packets directly to known hosts

# Learning Switch

```
let known_hosts : (dlAddr, portId) Hashtbl.t = Hashtbl.create 101
let learning_packet_in (sw:switchId) (pktIn:packetIn) : unit =
Hashtbl.add known_hosts pk.dlSrc pktIn.port

let routing_packet_in (sw:switchId) (pktIn : packetIn) : unit =
try
  let out_port = Hashtbl.find known_hosts pk.dlDst in
  let sd = {match_all with dlDst=Some pk.dlDst; dlSrc=Some pk.dlSrc} in
  let ds = {match_all with dlDst=Some pk.DlSrc; dlSrc=Some pk.dlDst} in
  send_flow_mod sw 01 (add_flow 0 sd [Output out_port]);
  send_flow_mod sw 01 (add_flow 0 ds [Output pktIn.port] );
  send_packet_out sw 01
    { output_payload = pktIn.input_payload;
      port_id = None;
      apply_actions = [Output out_port] }
  with Not found ->
    send_packet_out sw 01
    { output_payload = pktIn.input_payload;
      port_id = None;
      apply_actions = [Flood] }

let packet_in (sw:switchId) _ (pk:packetIn) : unit =
  learning_packet_in sw pk; routing_packet_in sw pk
```

# PyNetKAT Learning Switch

```
table = {}

def learn(app, switch, payload, port):
    pkt = Packet.from_payload(switch, port, payload)
    mac = pkt.ethSrc
    table[switch_id][mac] = port

def mac((fwd,unk),mac):
    fwd = fwd | Filter(EthDstEq(mac)) >> SetPort(table[sw][mac])
    unk = unk & ~EthSrcEq(mac)
    return (fwd, unk)

def switch_policy(sw):
    (fwd, unk) = reduce(mac, table[sw].keys(), (drop, true))
    unk = Filter(unk) >> (SendToController("learning") | flood(sw))
    return (fwd | unk)

class LearningApp(frenetic.App):
    ...
    def packet_in(self,switch_id, port_id, payload):
        learn(self, switch_id,payload, port_id)
        self.update(Union(switch_policy(sw) for sw in table.keys()))
```

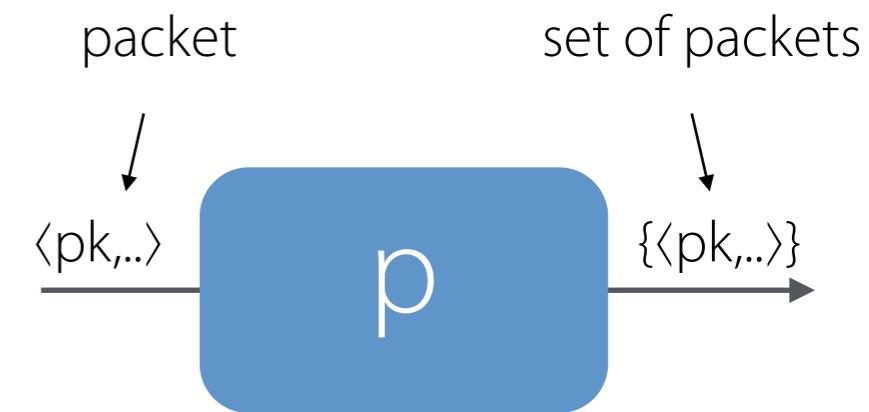
## Languages

## Machines

A *machine model* describes behavior in terms of concepts like pipelines of hardware lookup tables

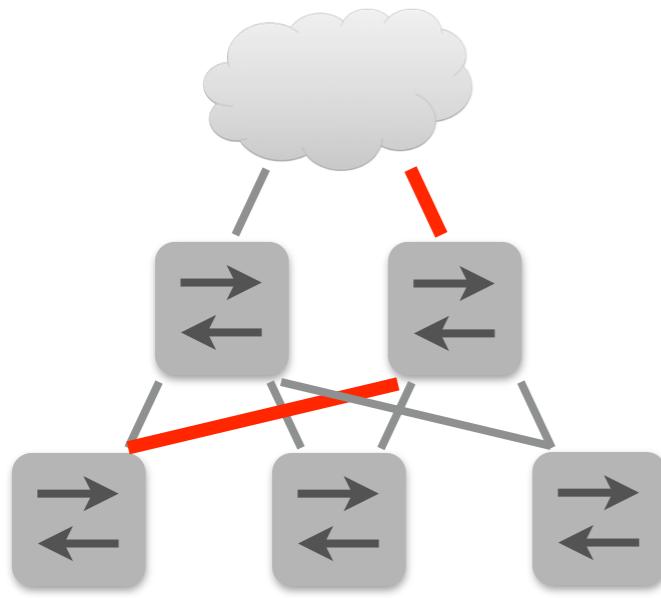
Match	Actions
ethType=0x800, ipProto=0x06, tcpDstPort=22, ethSrc=00:00:00:00:00:01	Drop
ethType=0x800, ipProto=0x06, tcpDstPort=22, ethSrc=00:00:00:00:00:02	Drop
ethType=0x800, ipProto=0x06, tcpDstPort=22,	Inport
ethType=0x800, ipProto=0x06	Inport
ethType=0x800	Inport
*	Inport

A *programming model* describes behavior in terms of concepts like mathematical functions on packets

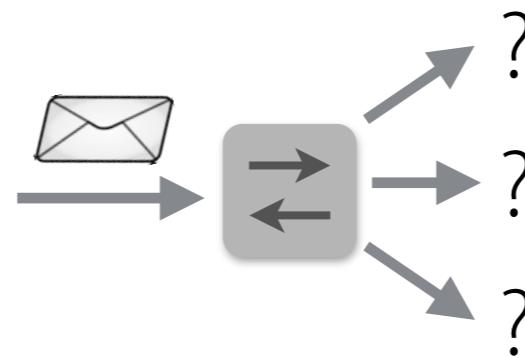


# **Language Design**

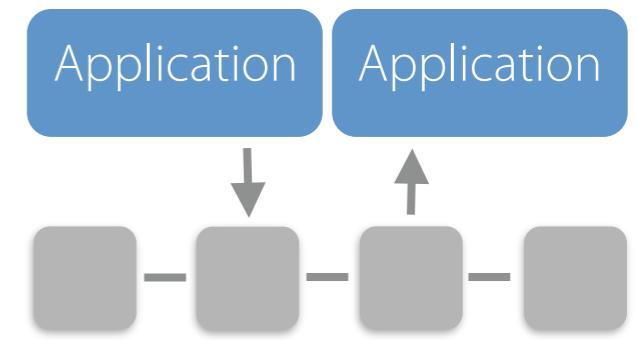
# Desiderata



High-level,  
Network-wide  
Abstractions

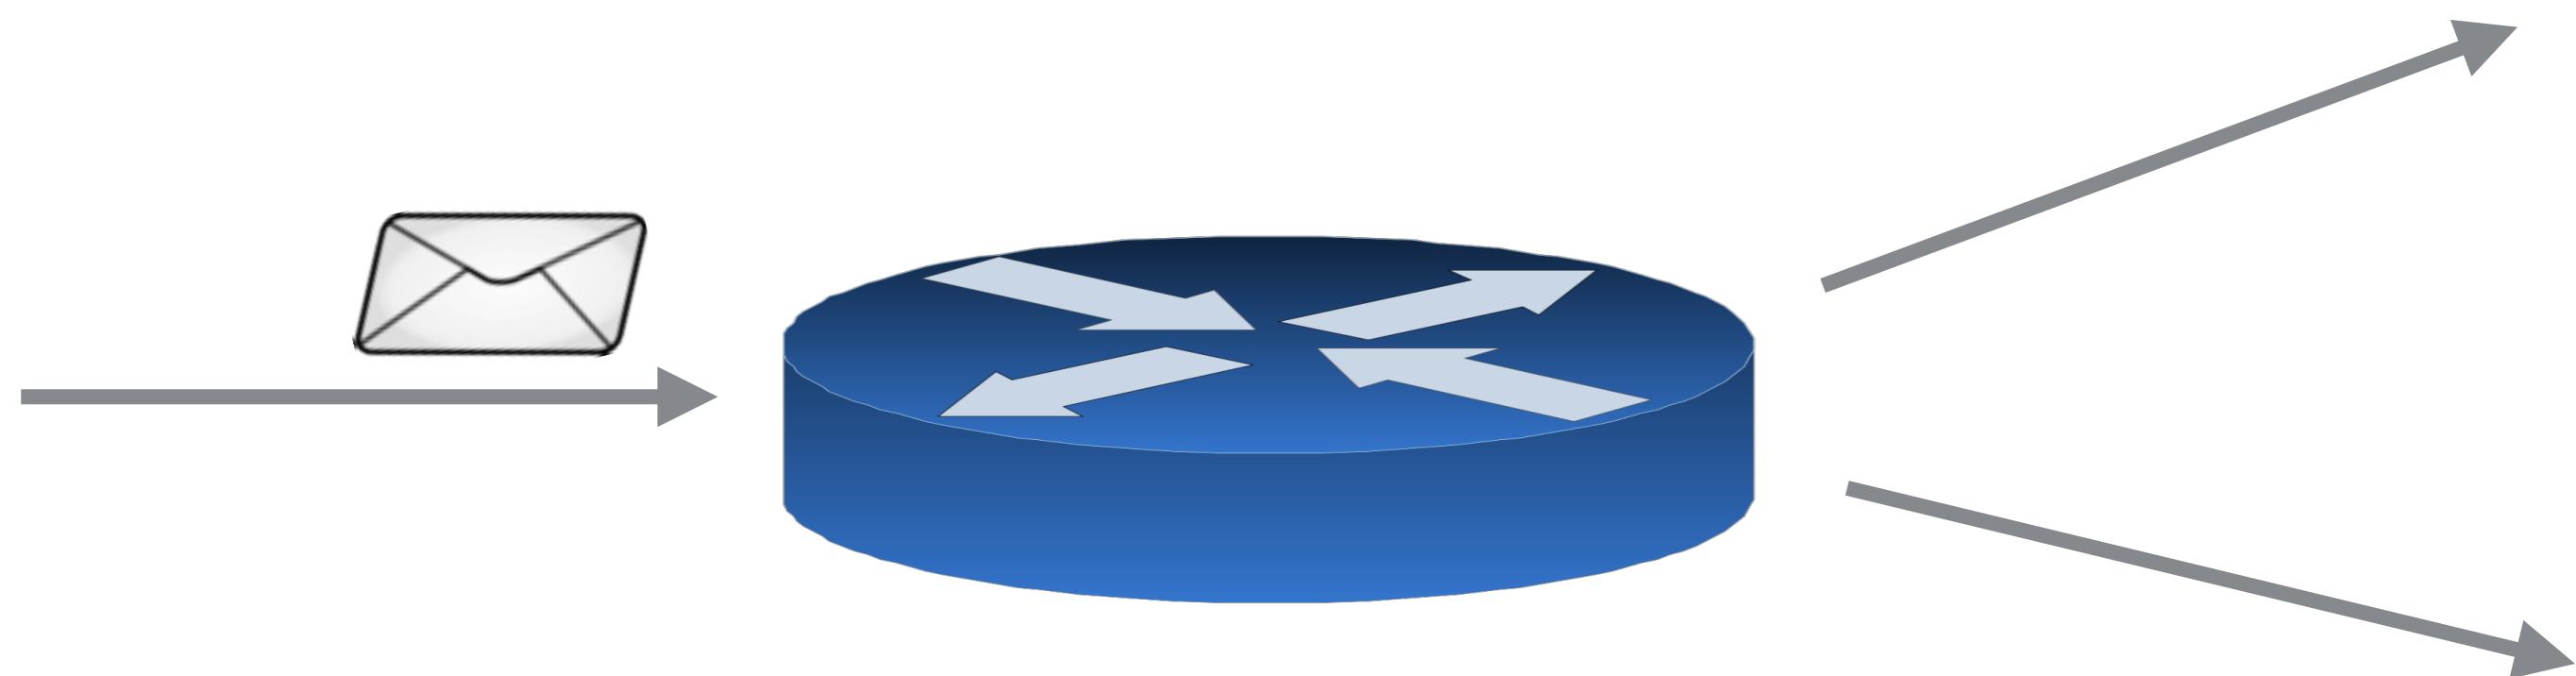


Rich Packet  
Classification



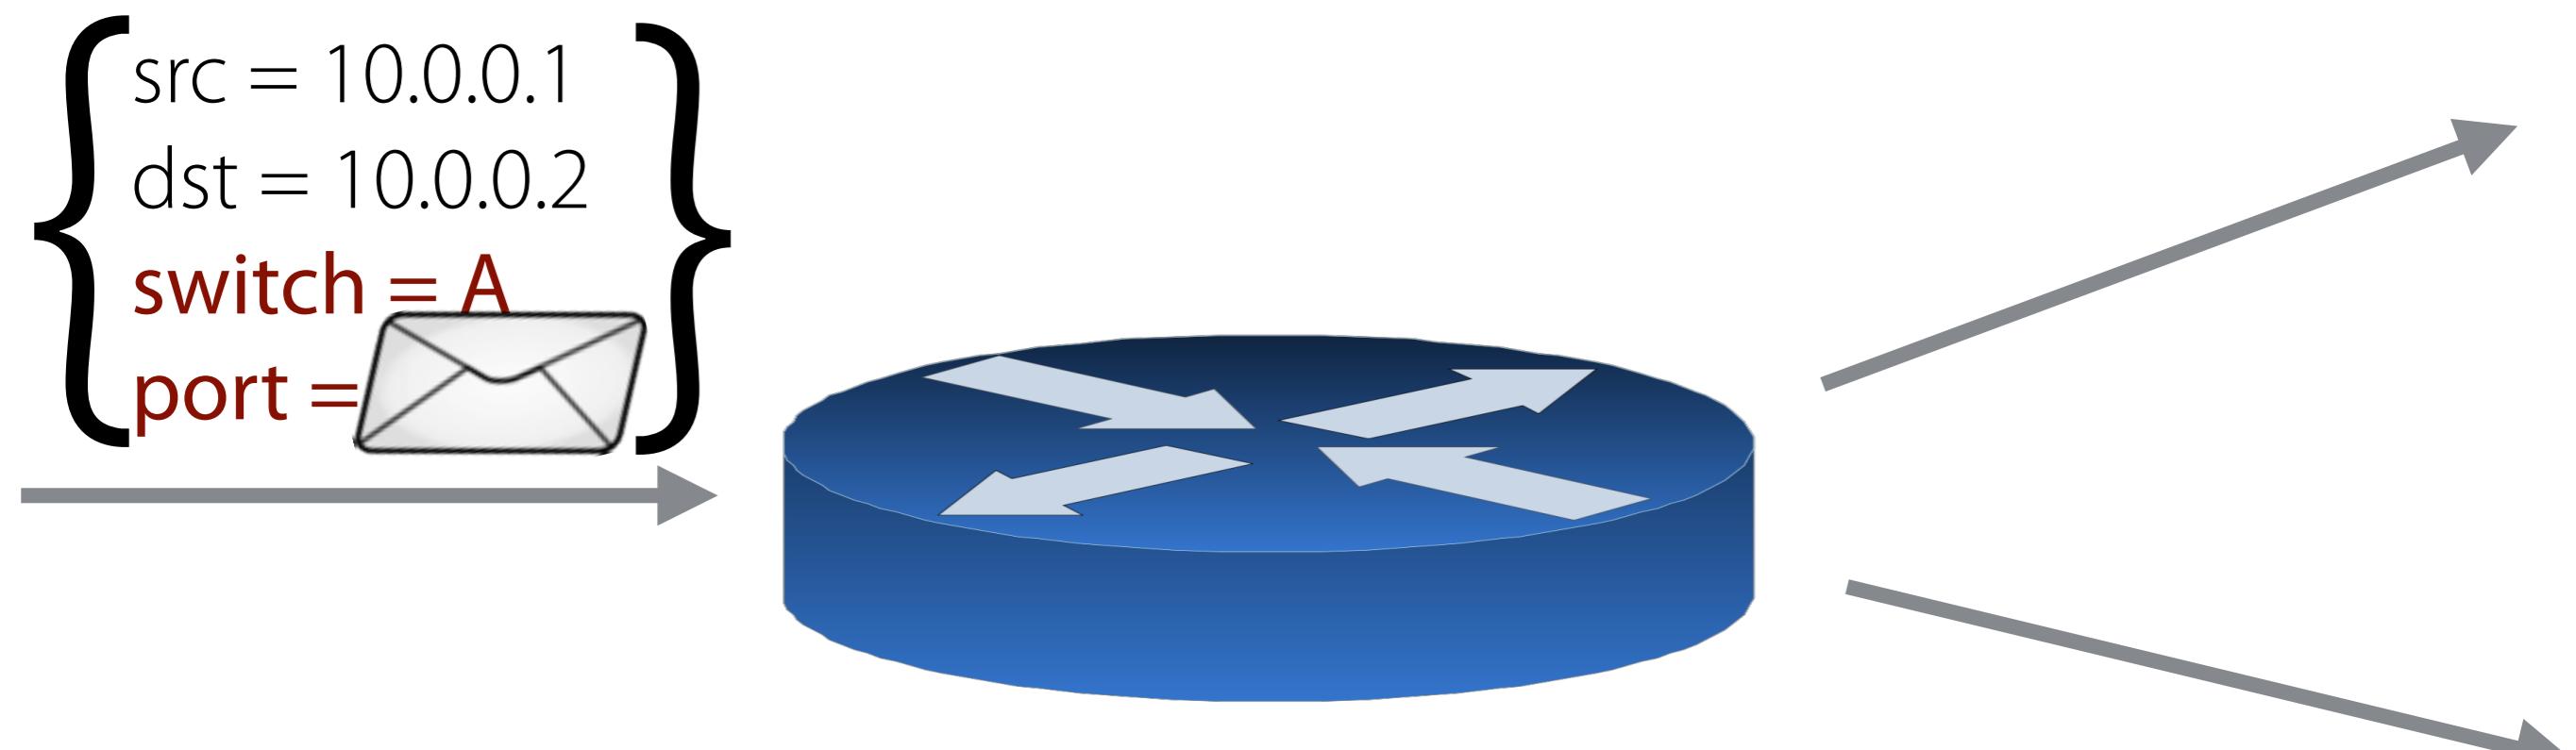
Modular  
Composition

# Model



Packets → Packets

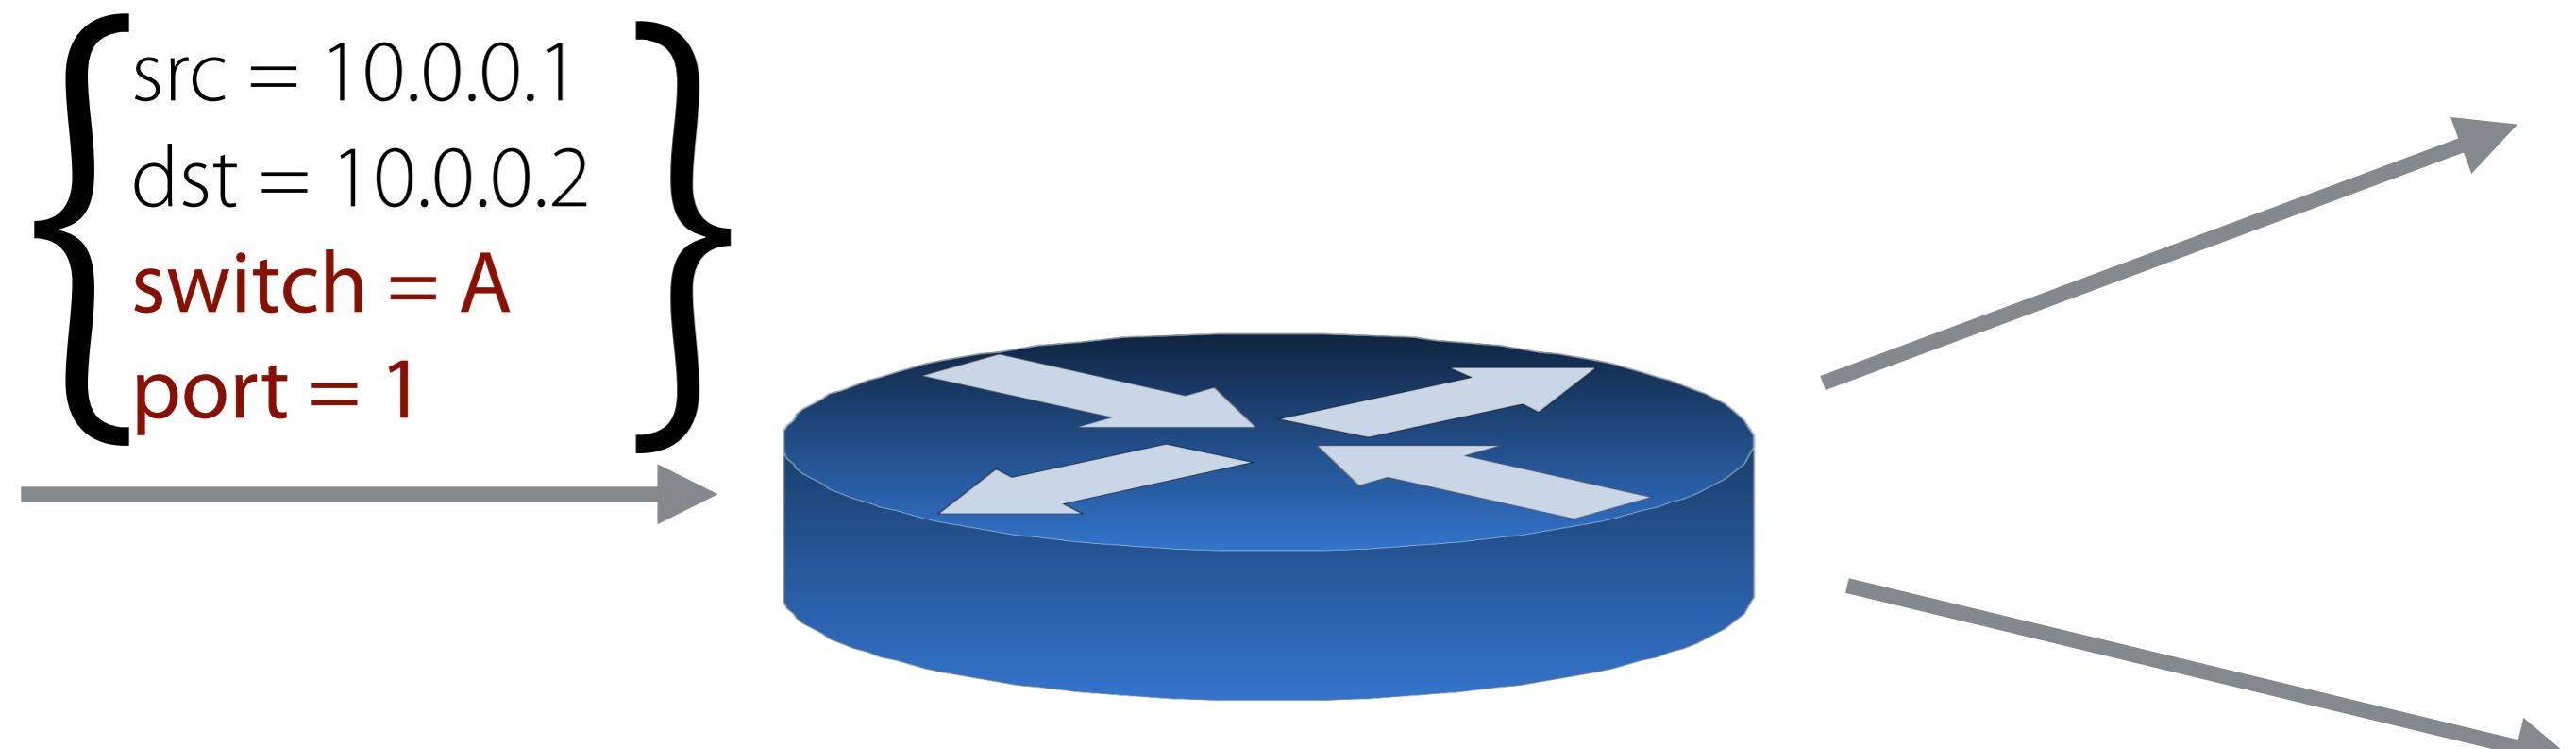
# Model



Packets → Packets

# Packet Manipulation

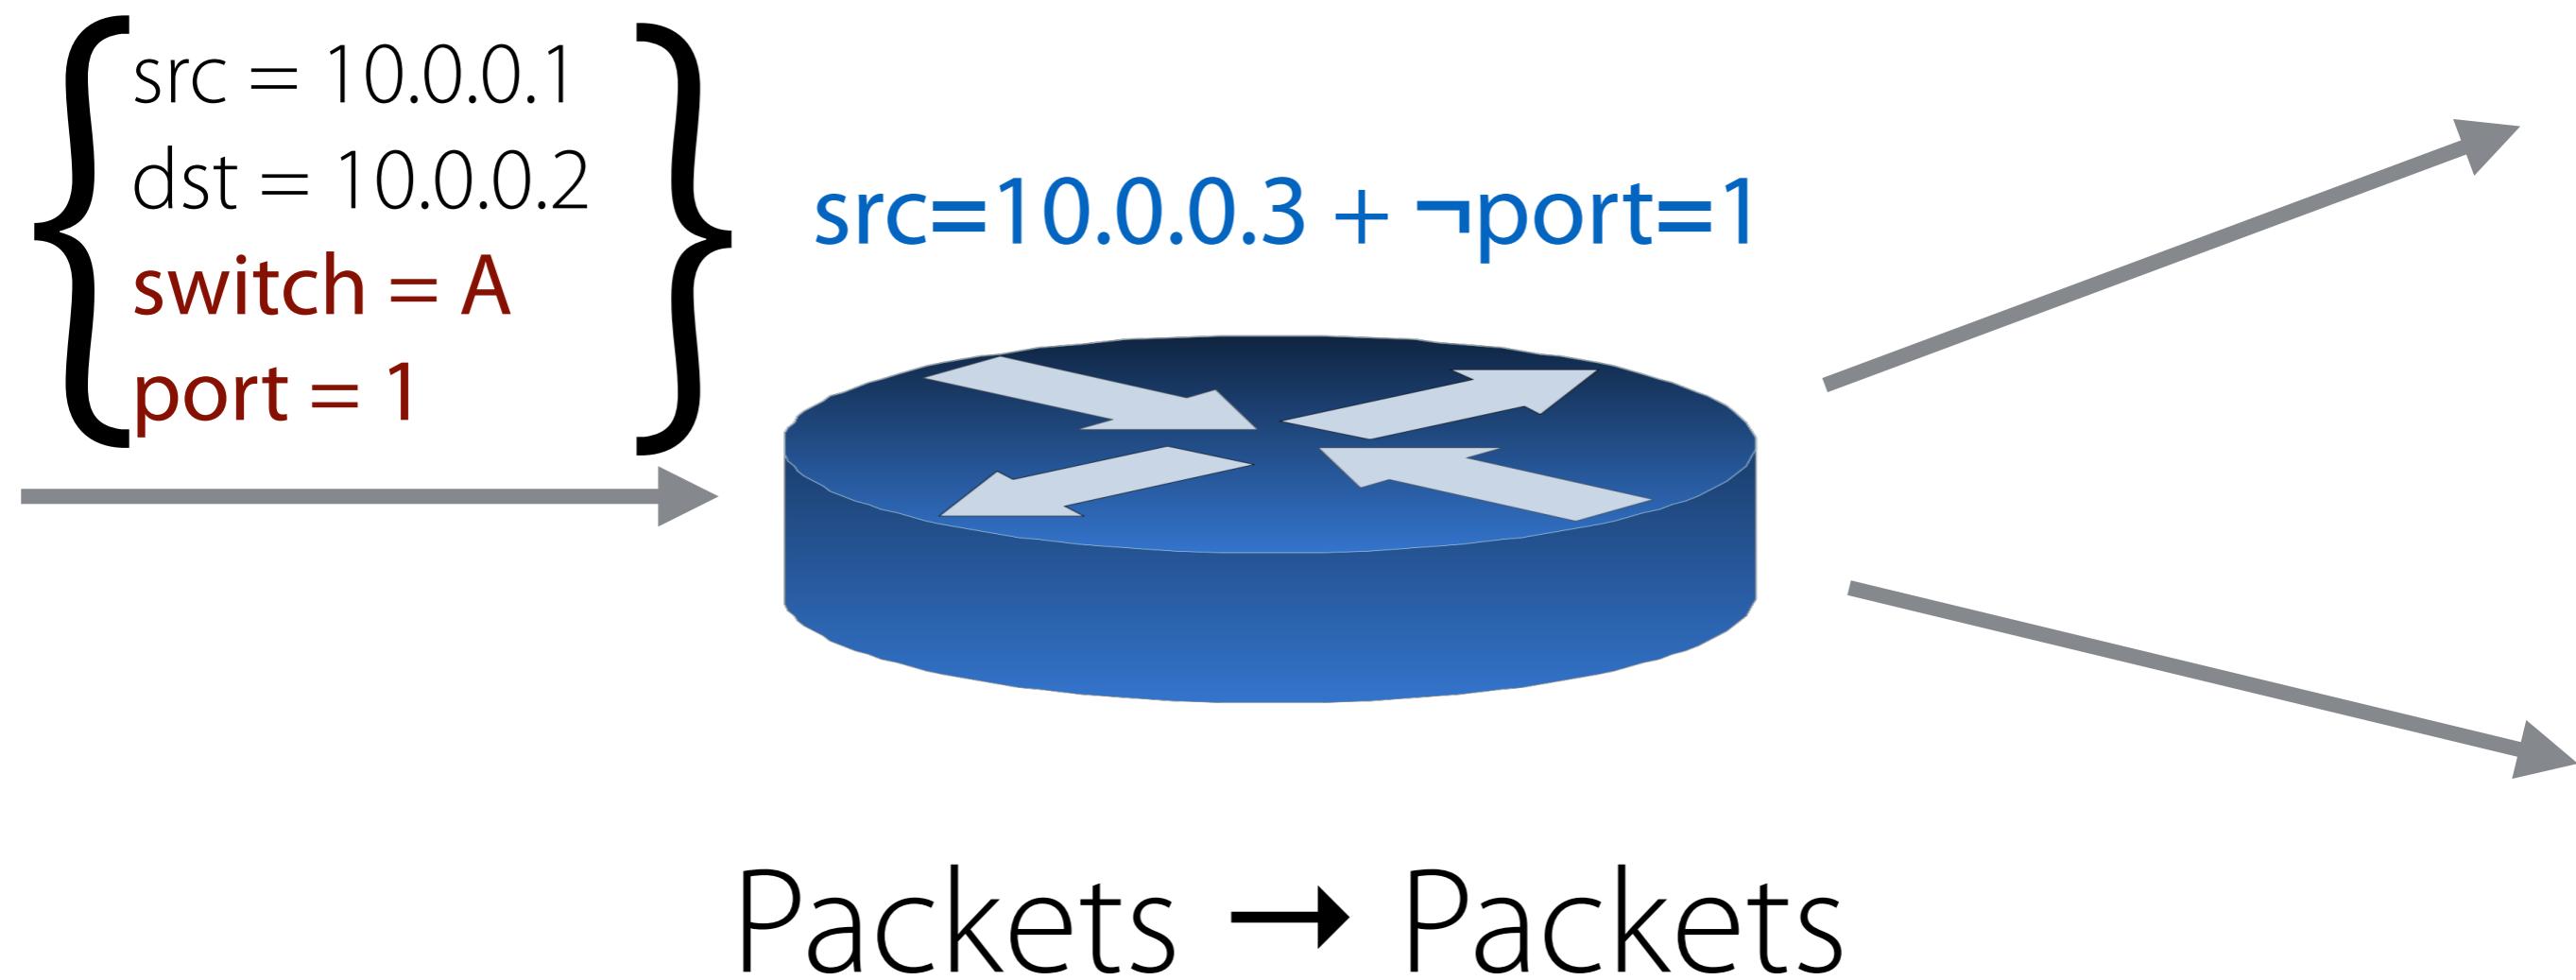
- packet filtering
- packet modification



Packets → Packets

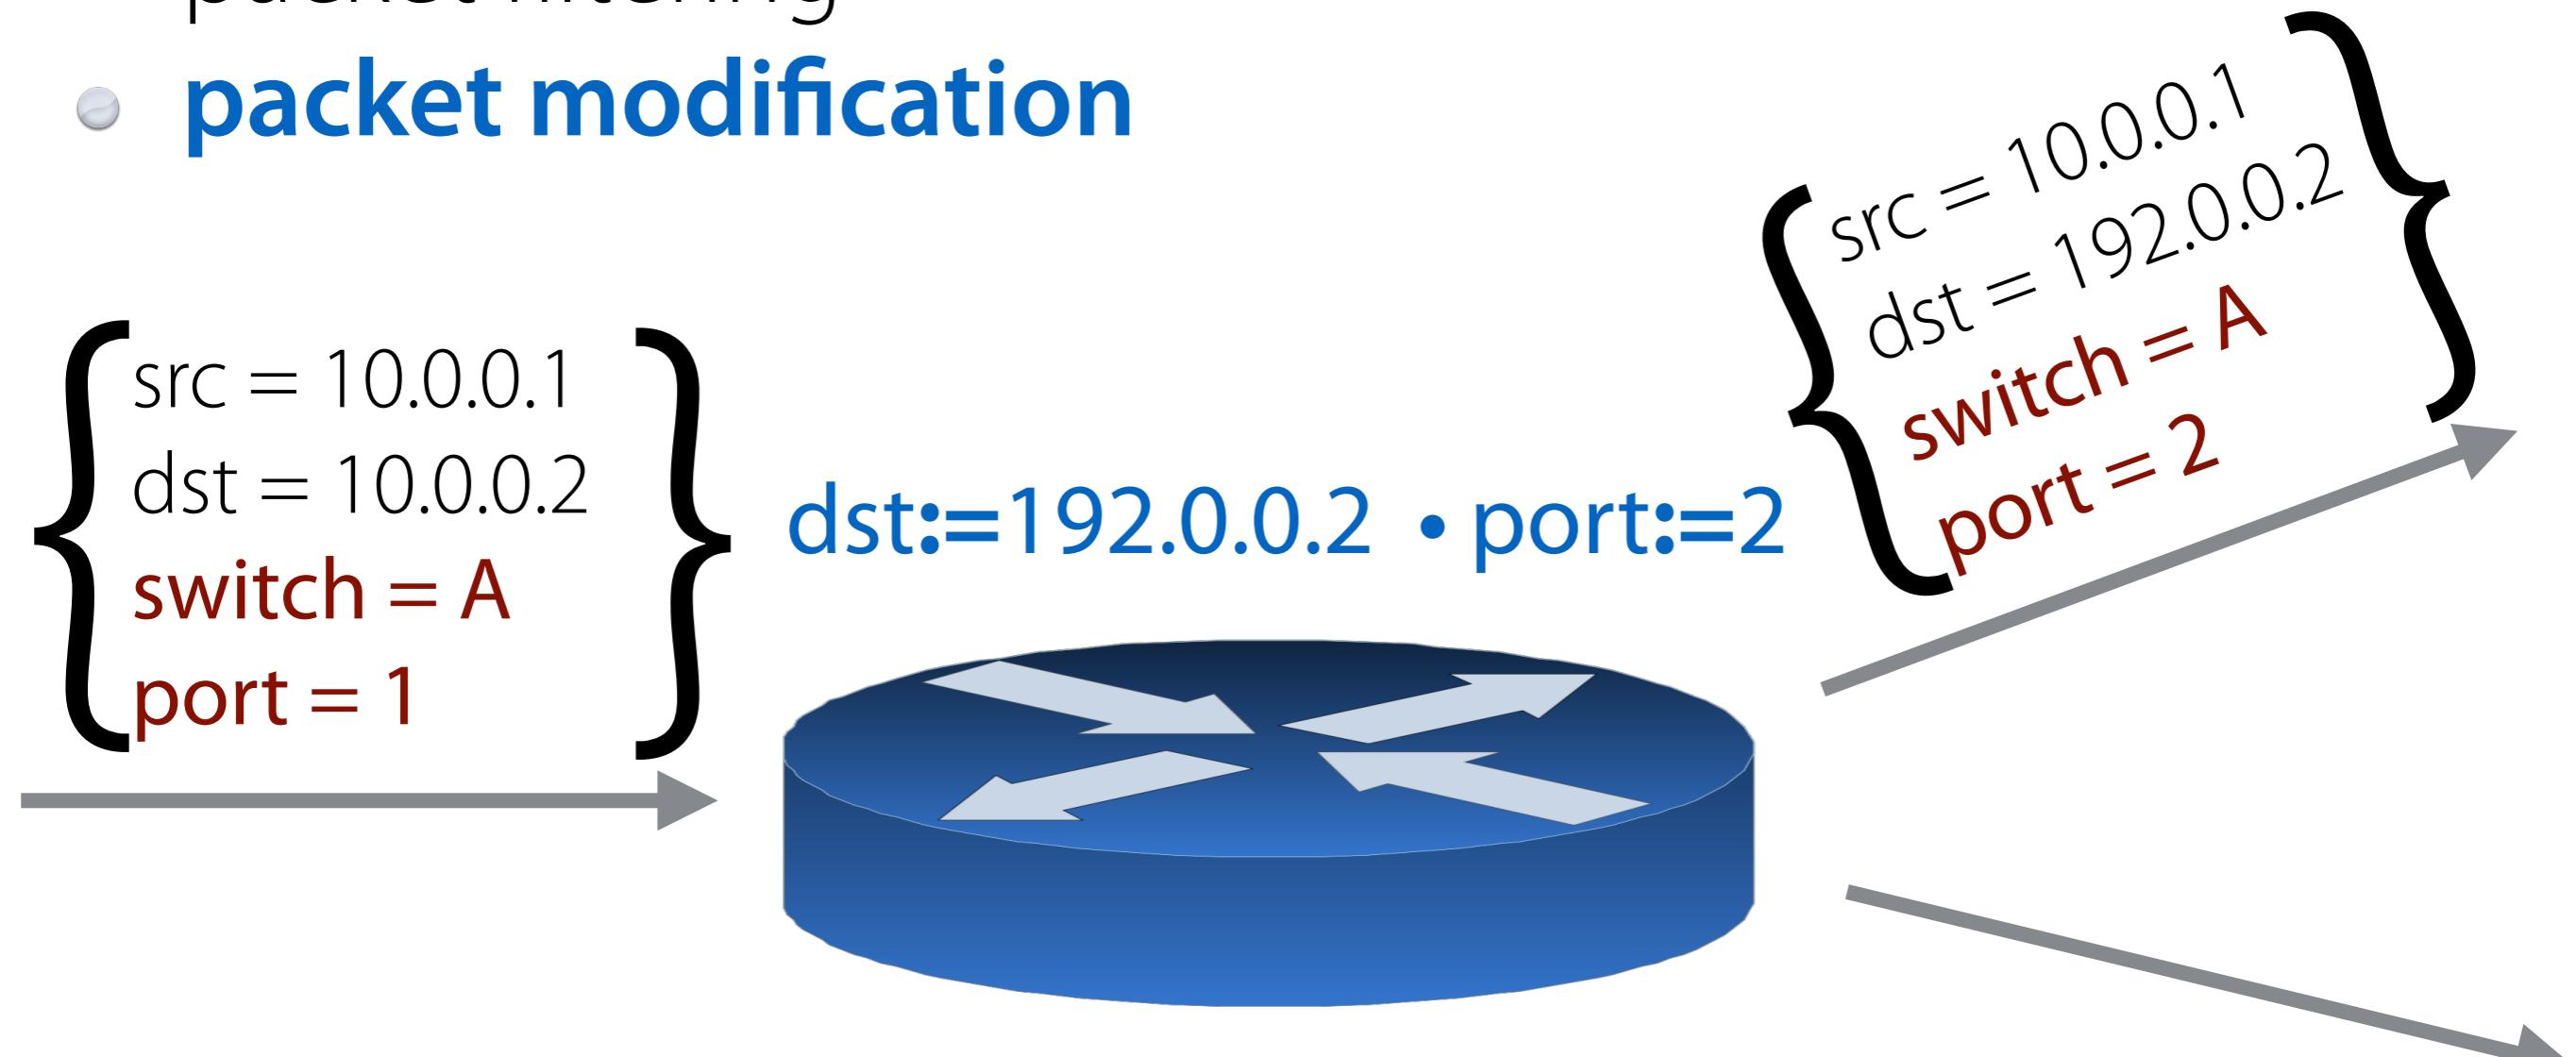
# Packet Manipulation

- **packet filtering**
- packet modification



# Packet Manipulation

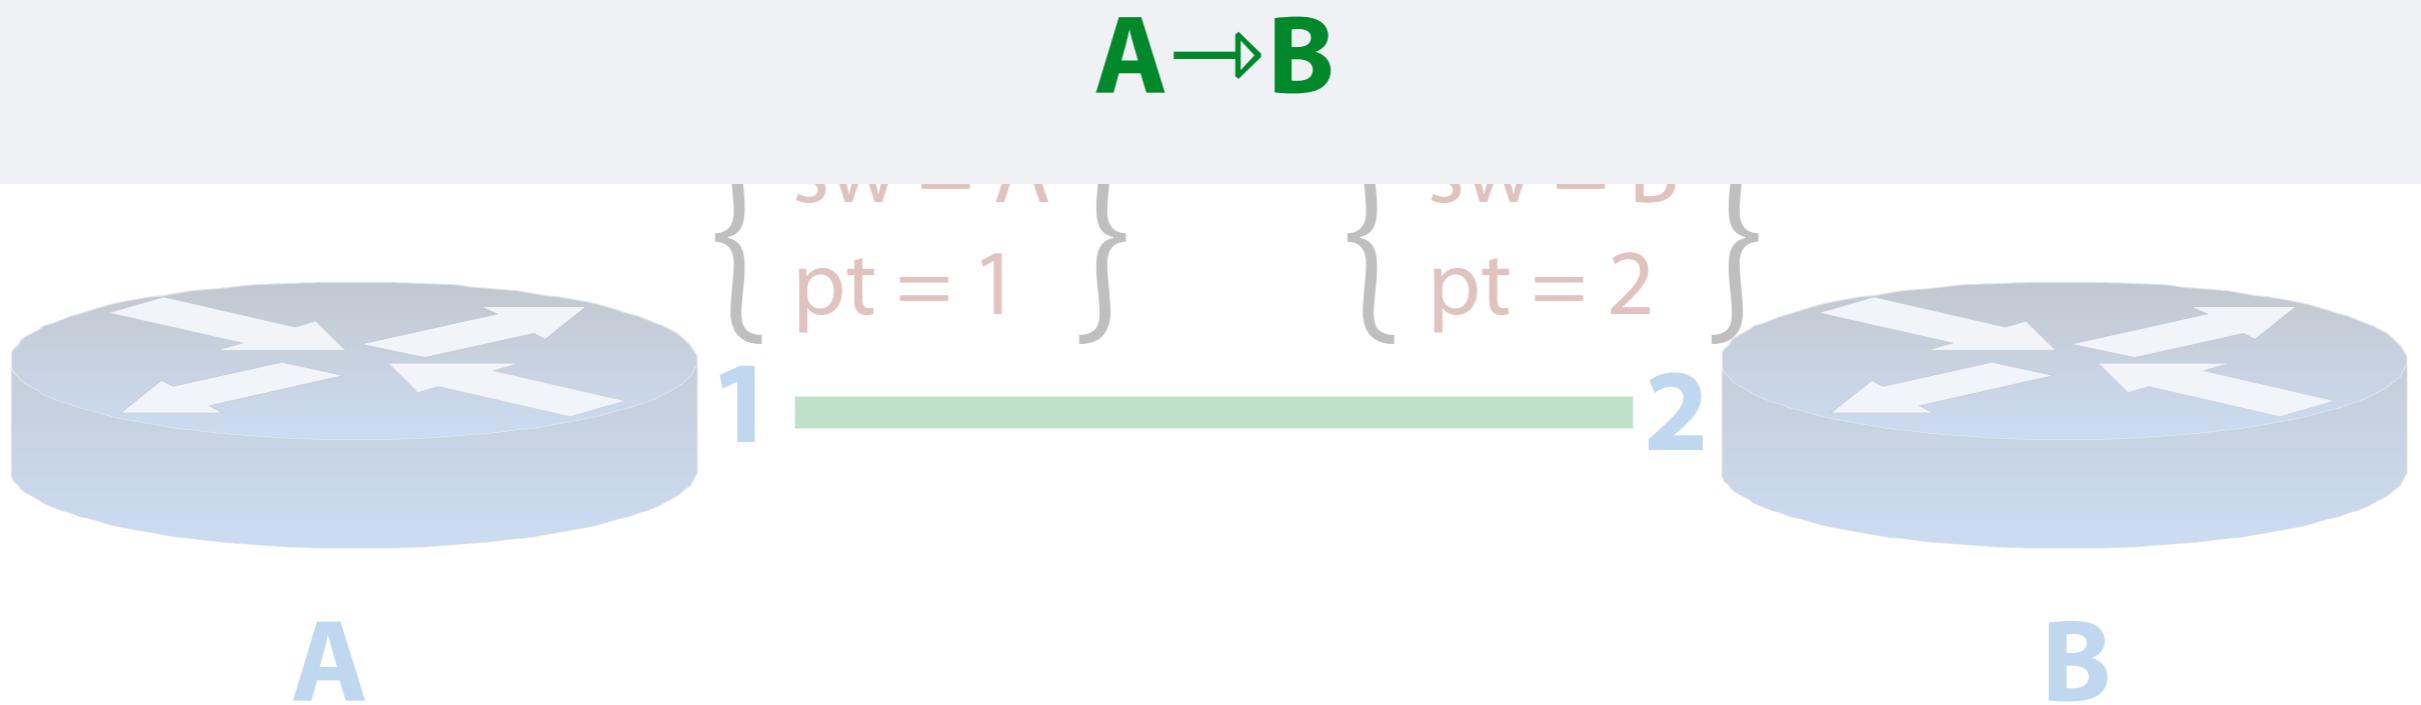
- packet filtering
- **packet modification**



Packets → Packets

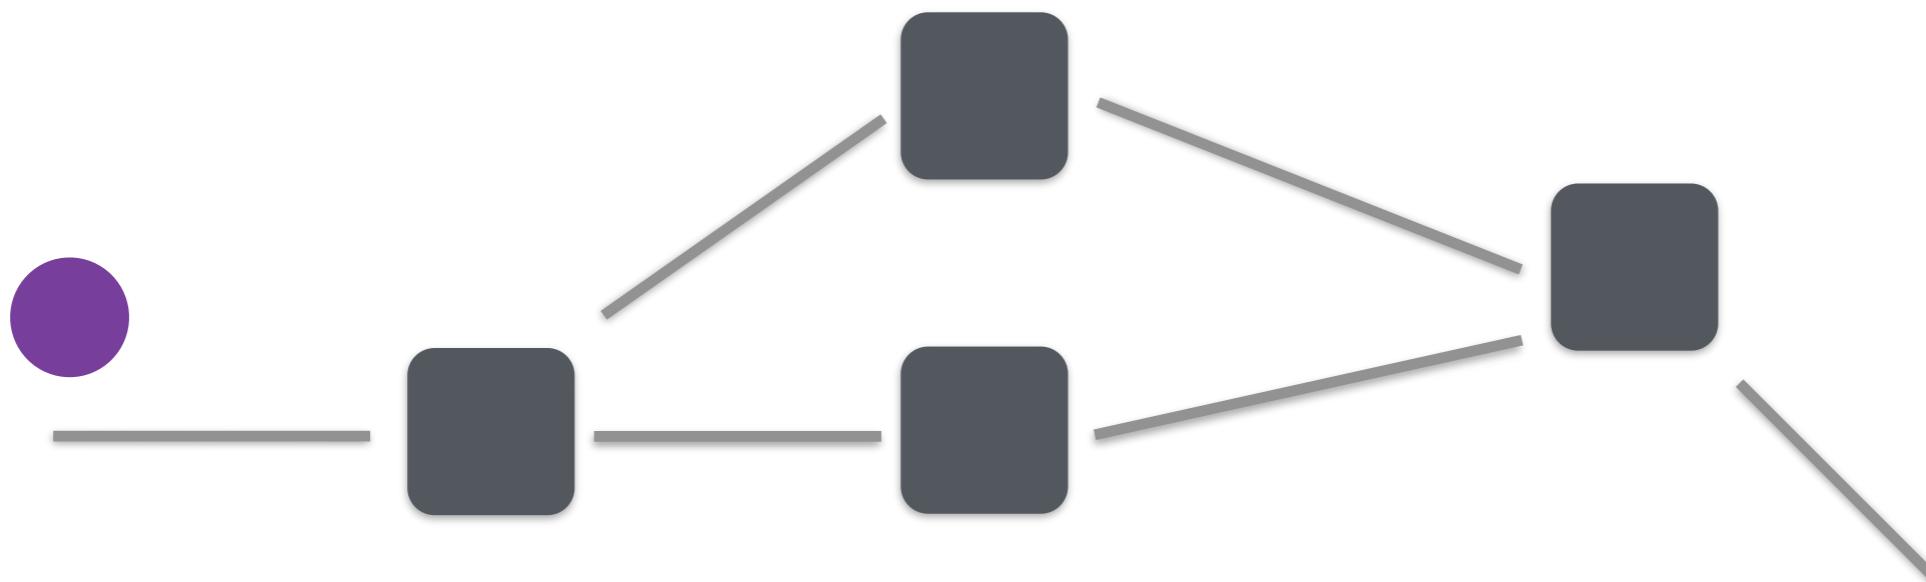
# Topology

`switch=A • pt=1 • switch:=B • pt:=2`



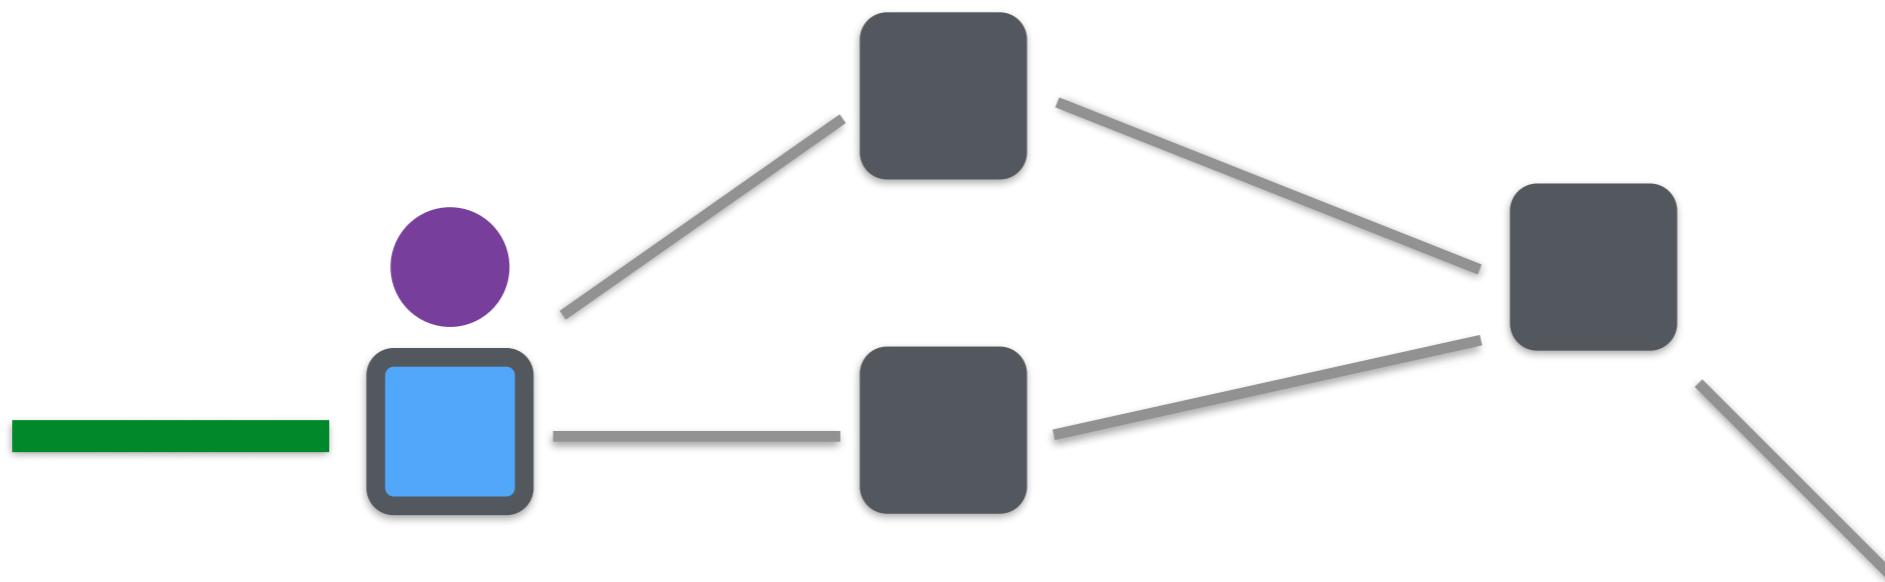
# Iteration

(topology • switch)\*



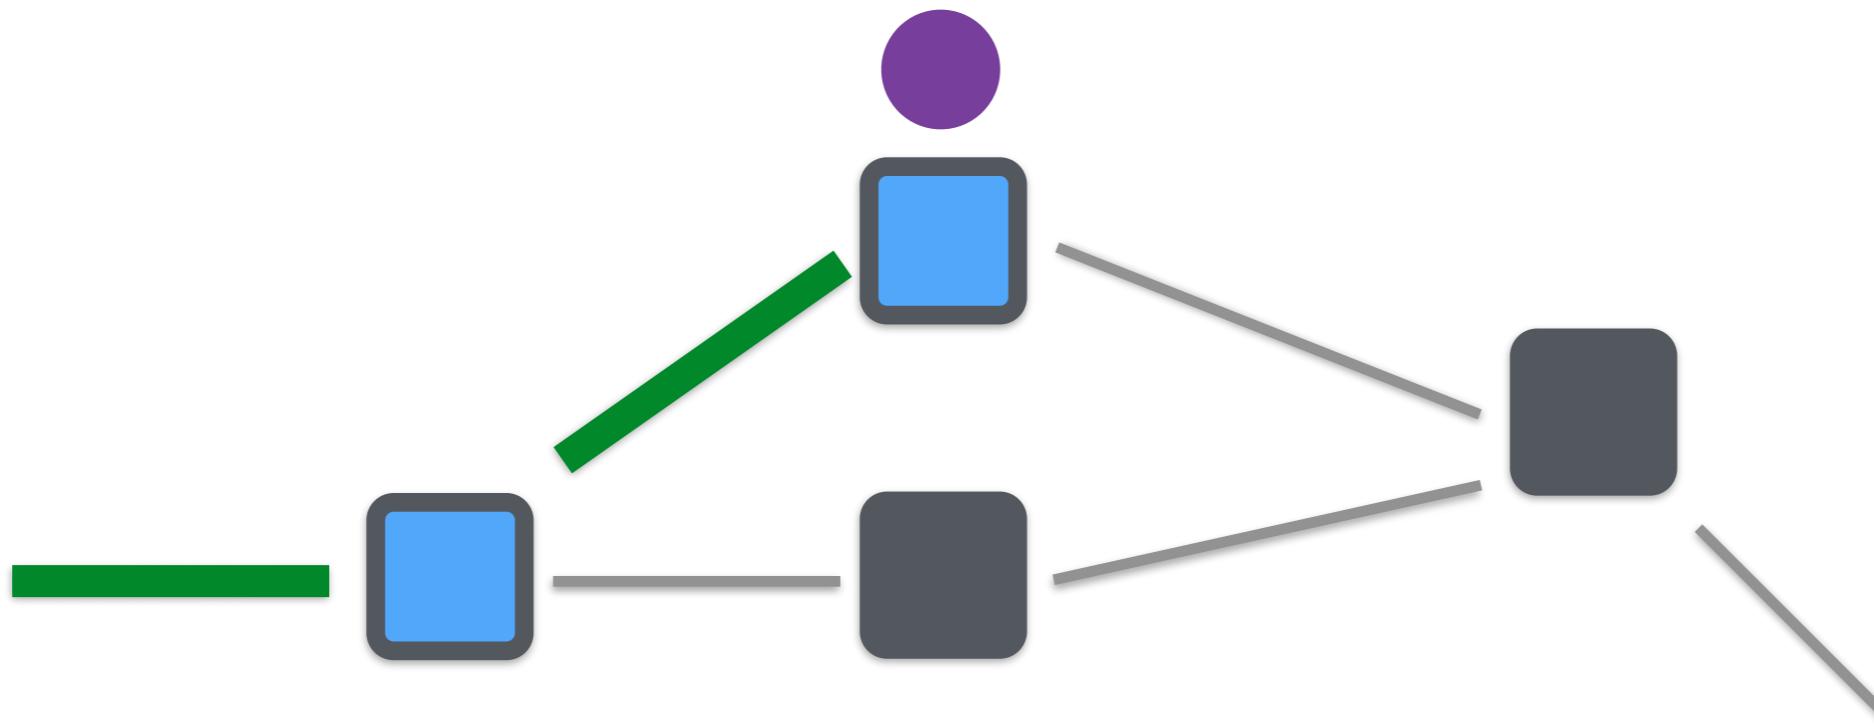
# Iteration

(topology • switch)\*



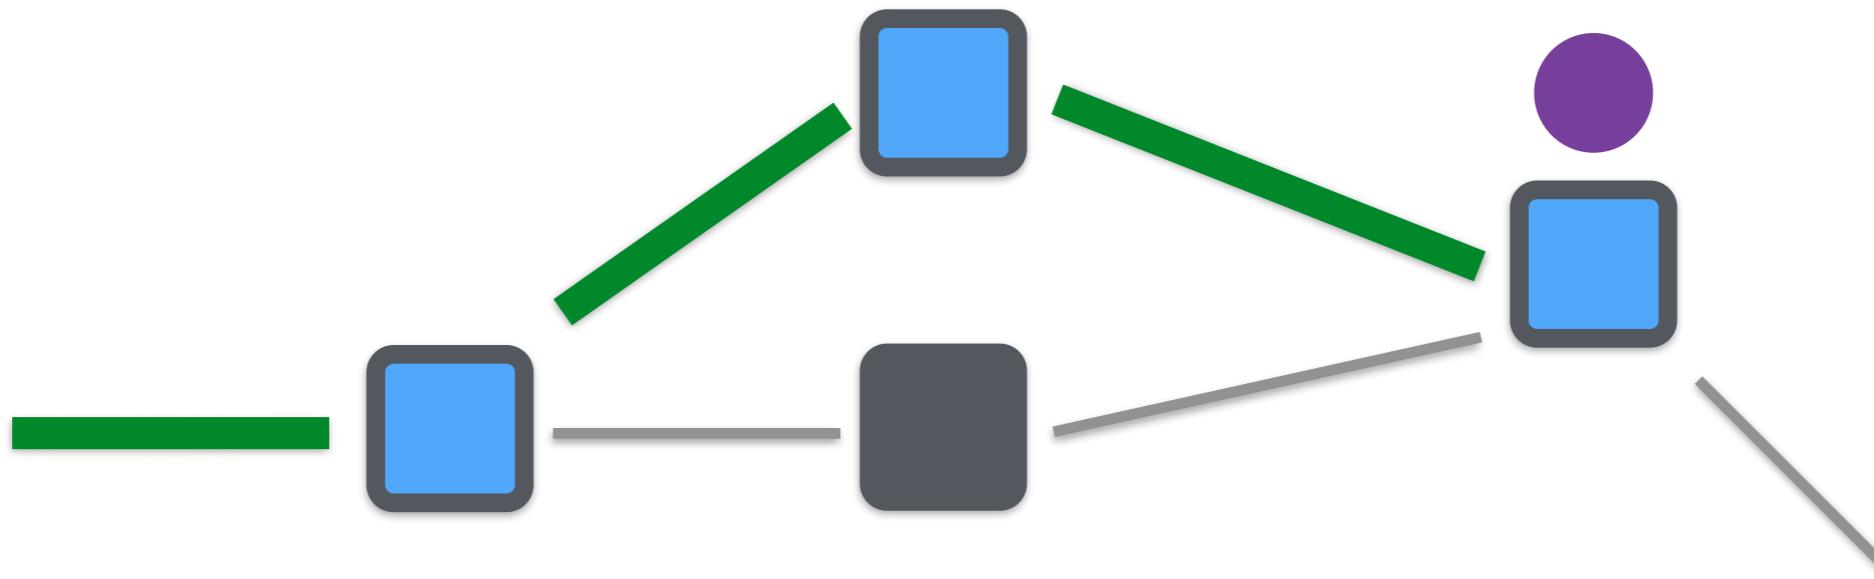
# Iteration

(topology • switch)\*



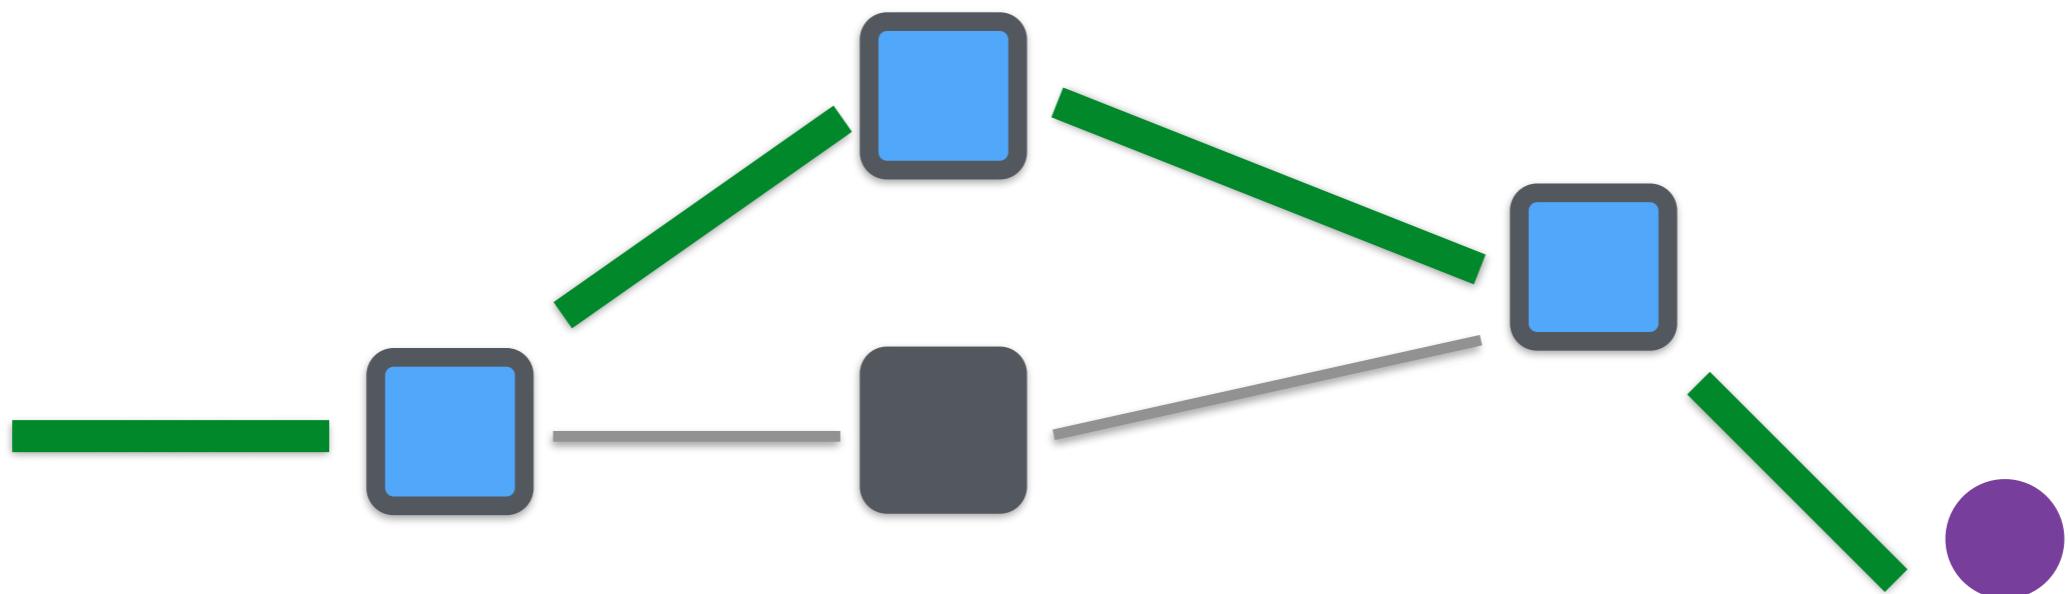
# Iteration

(topology • switch)\*



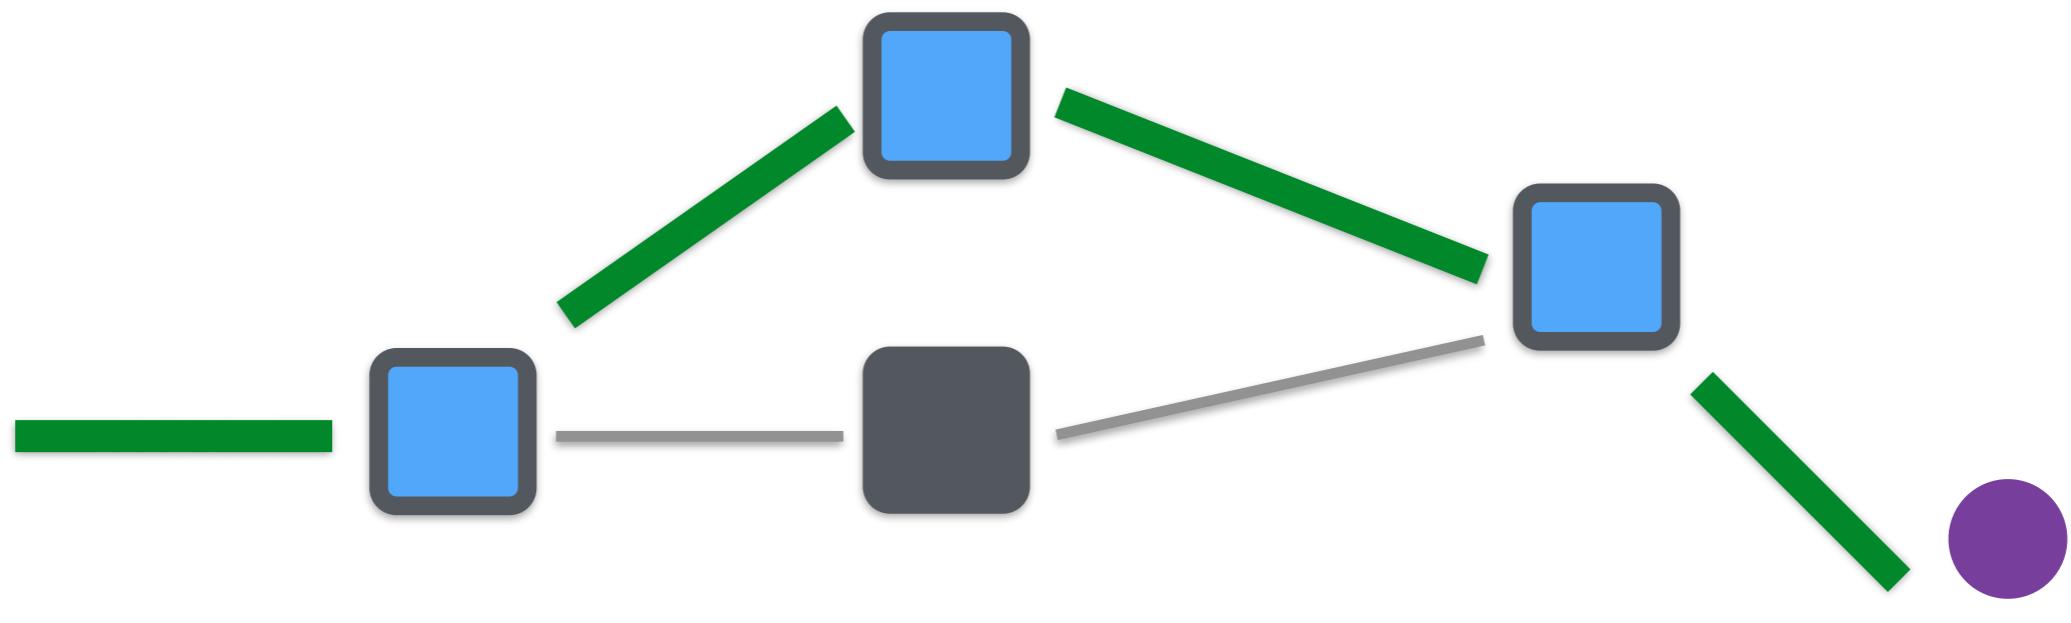
# Iteration

(topology • switch)\*



# Iteration

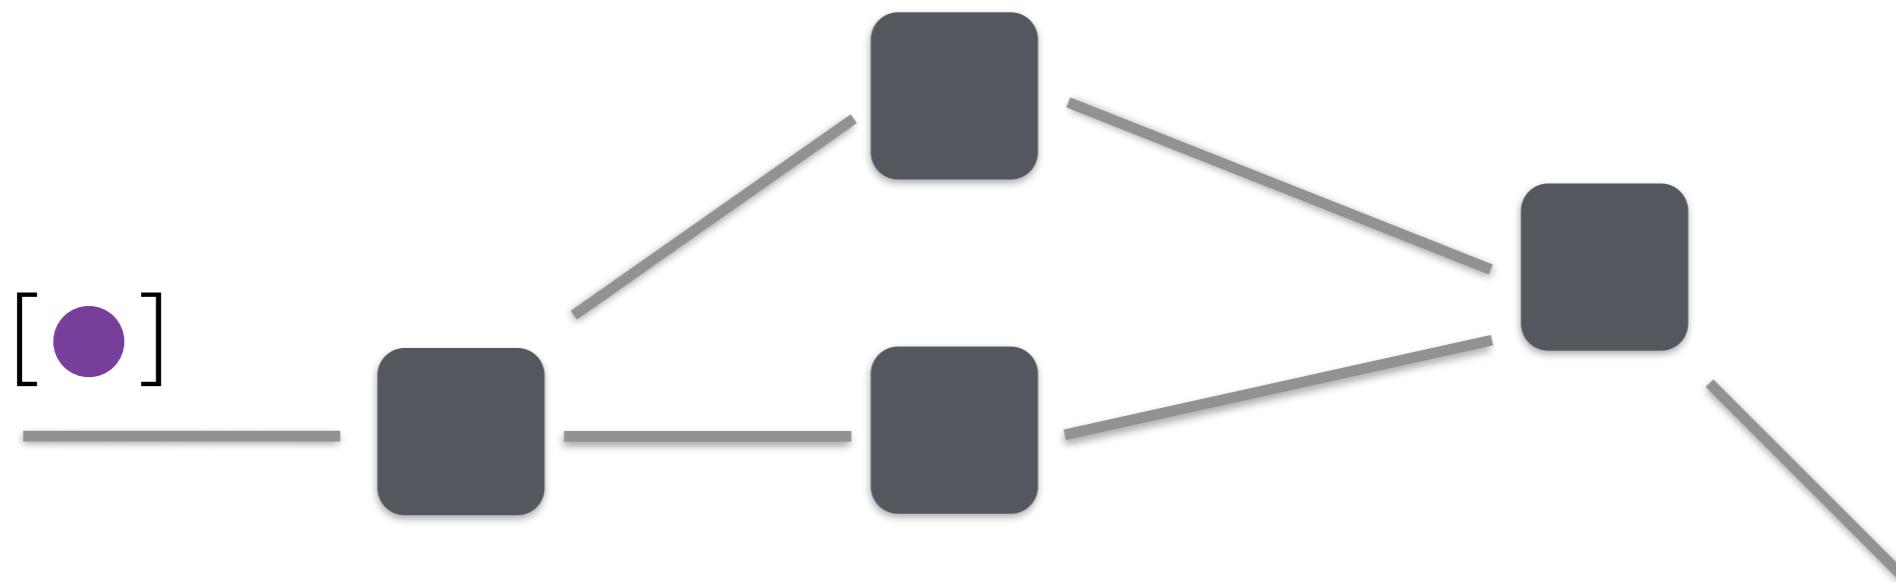
(topology • switch)\*



~~Packets → Packets~~

# Iteration

(topology • switch)\*

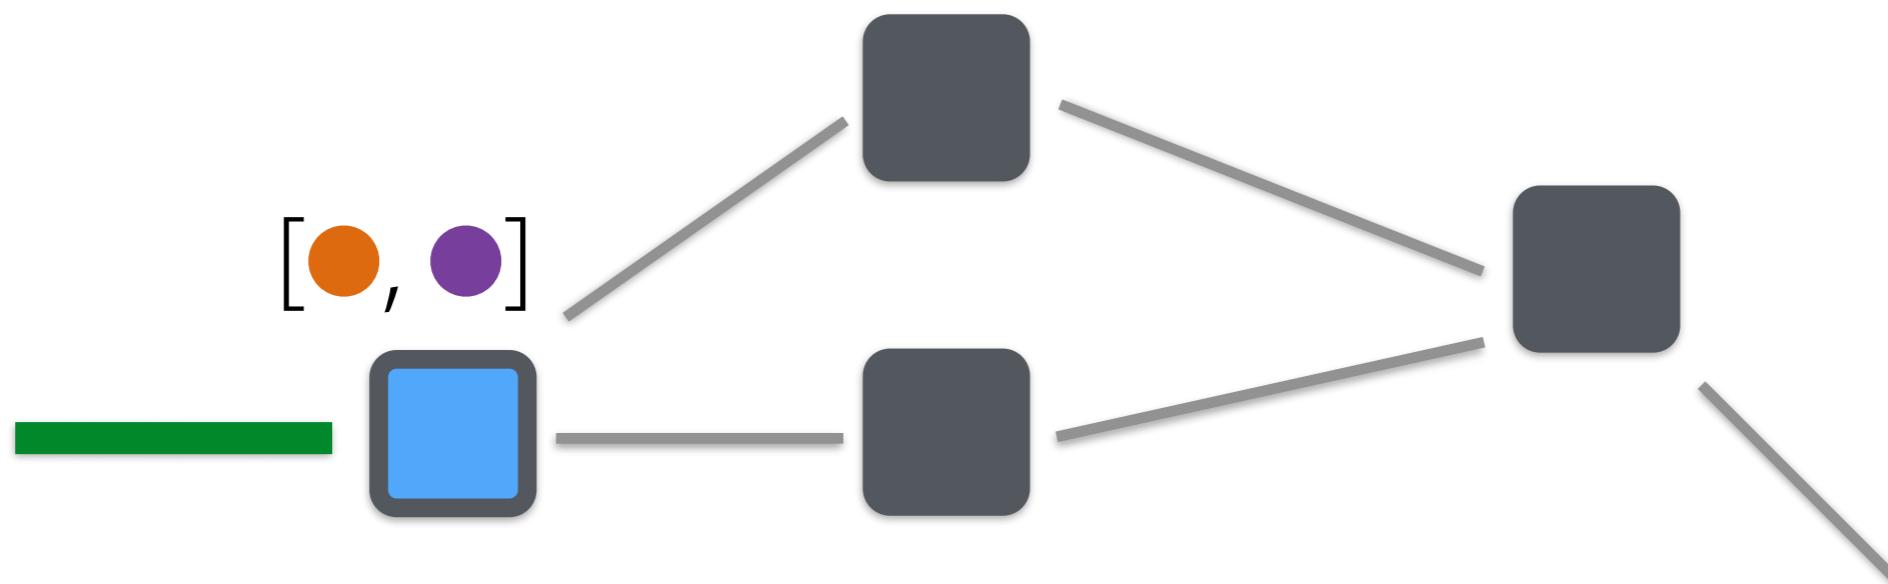


~~Packets → Packets~~

Histories → Histories

# Iteration

$(\text{topology} \bullet \text{switch})^*$

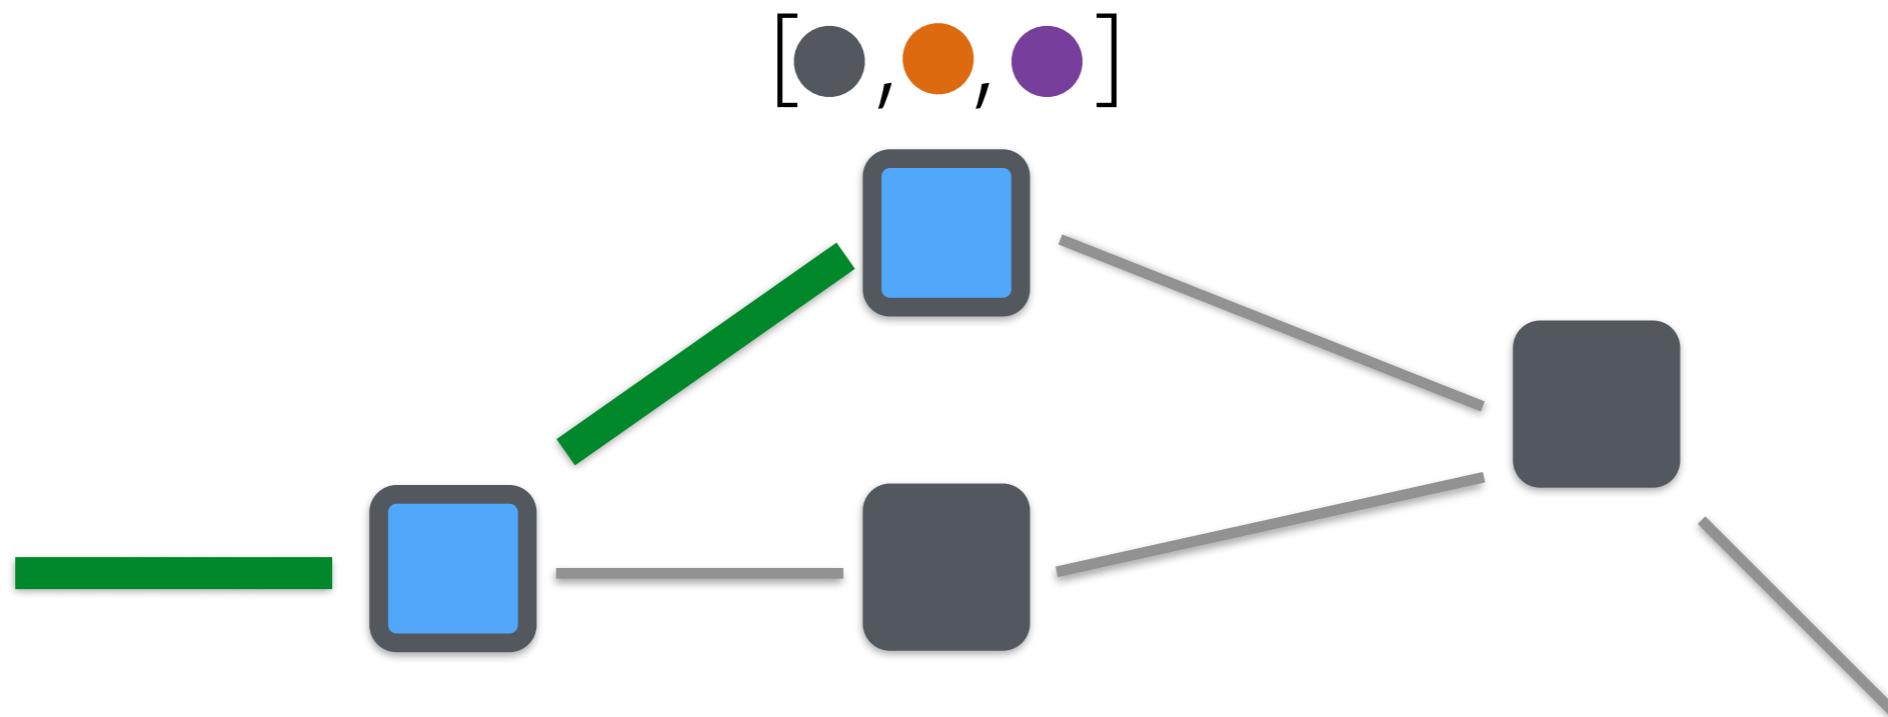


~~Packets → Packets~~

Histories → Histories

# Iteration

(topology • switch)\*

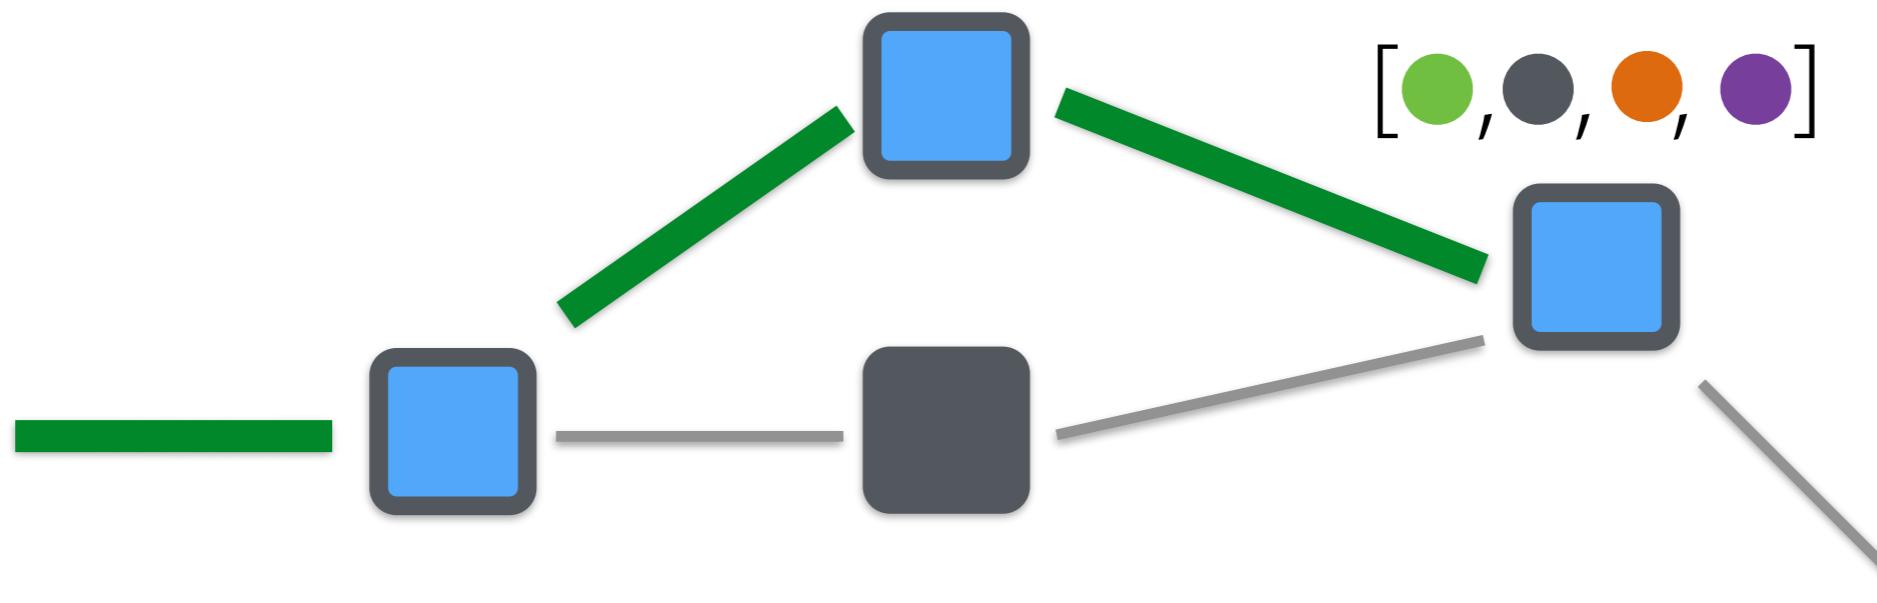


~~Packets → Packets~~

Histories → Histories

# Iteration

$(\text{topology} \bullet \text{switch})^*$

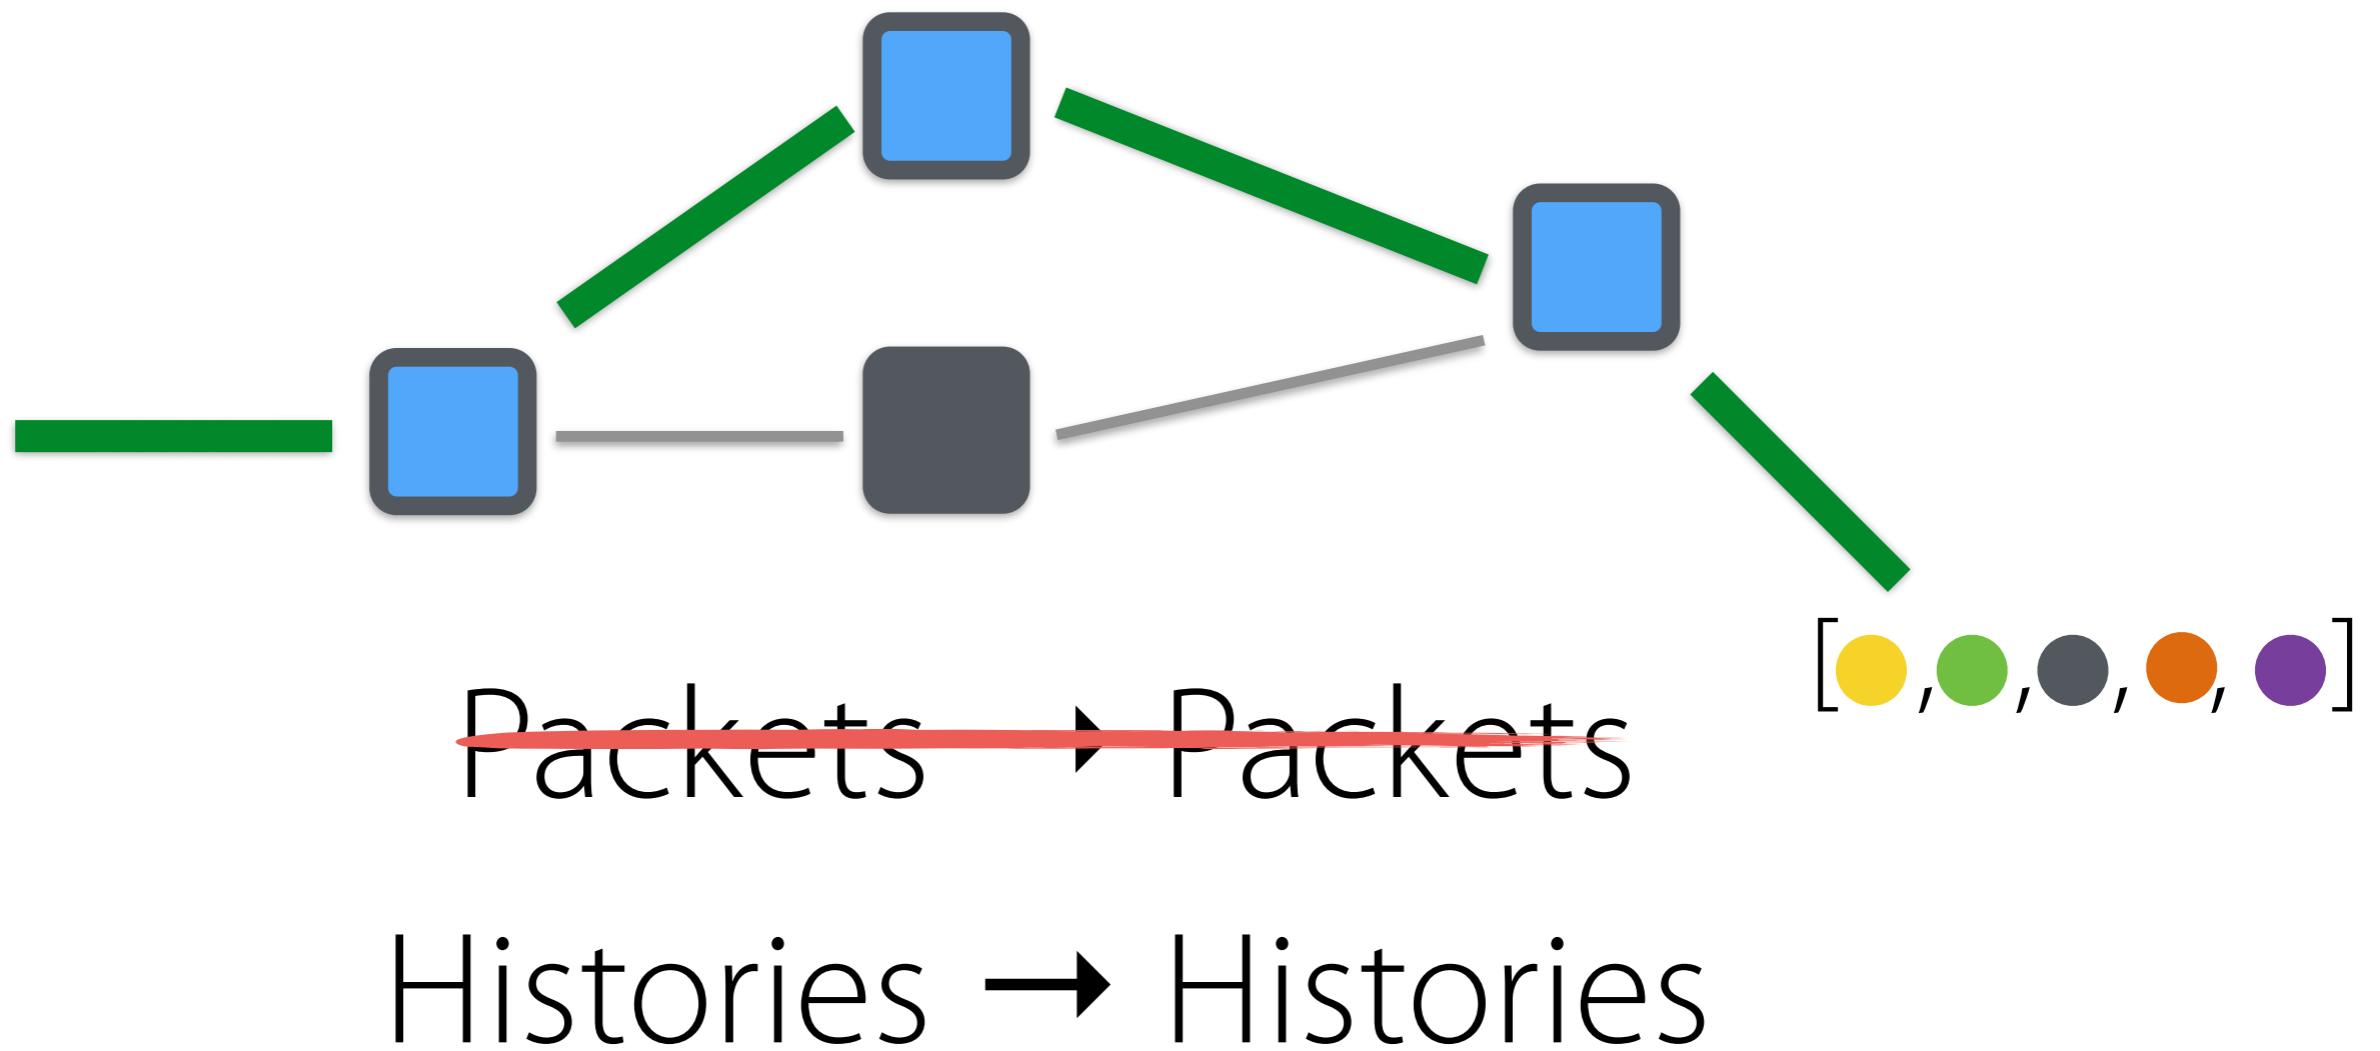


~~Packets → Packets~~

Histories → Histories

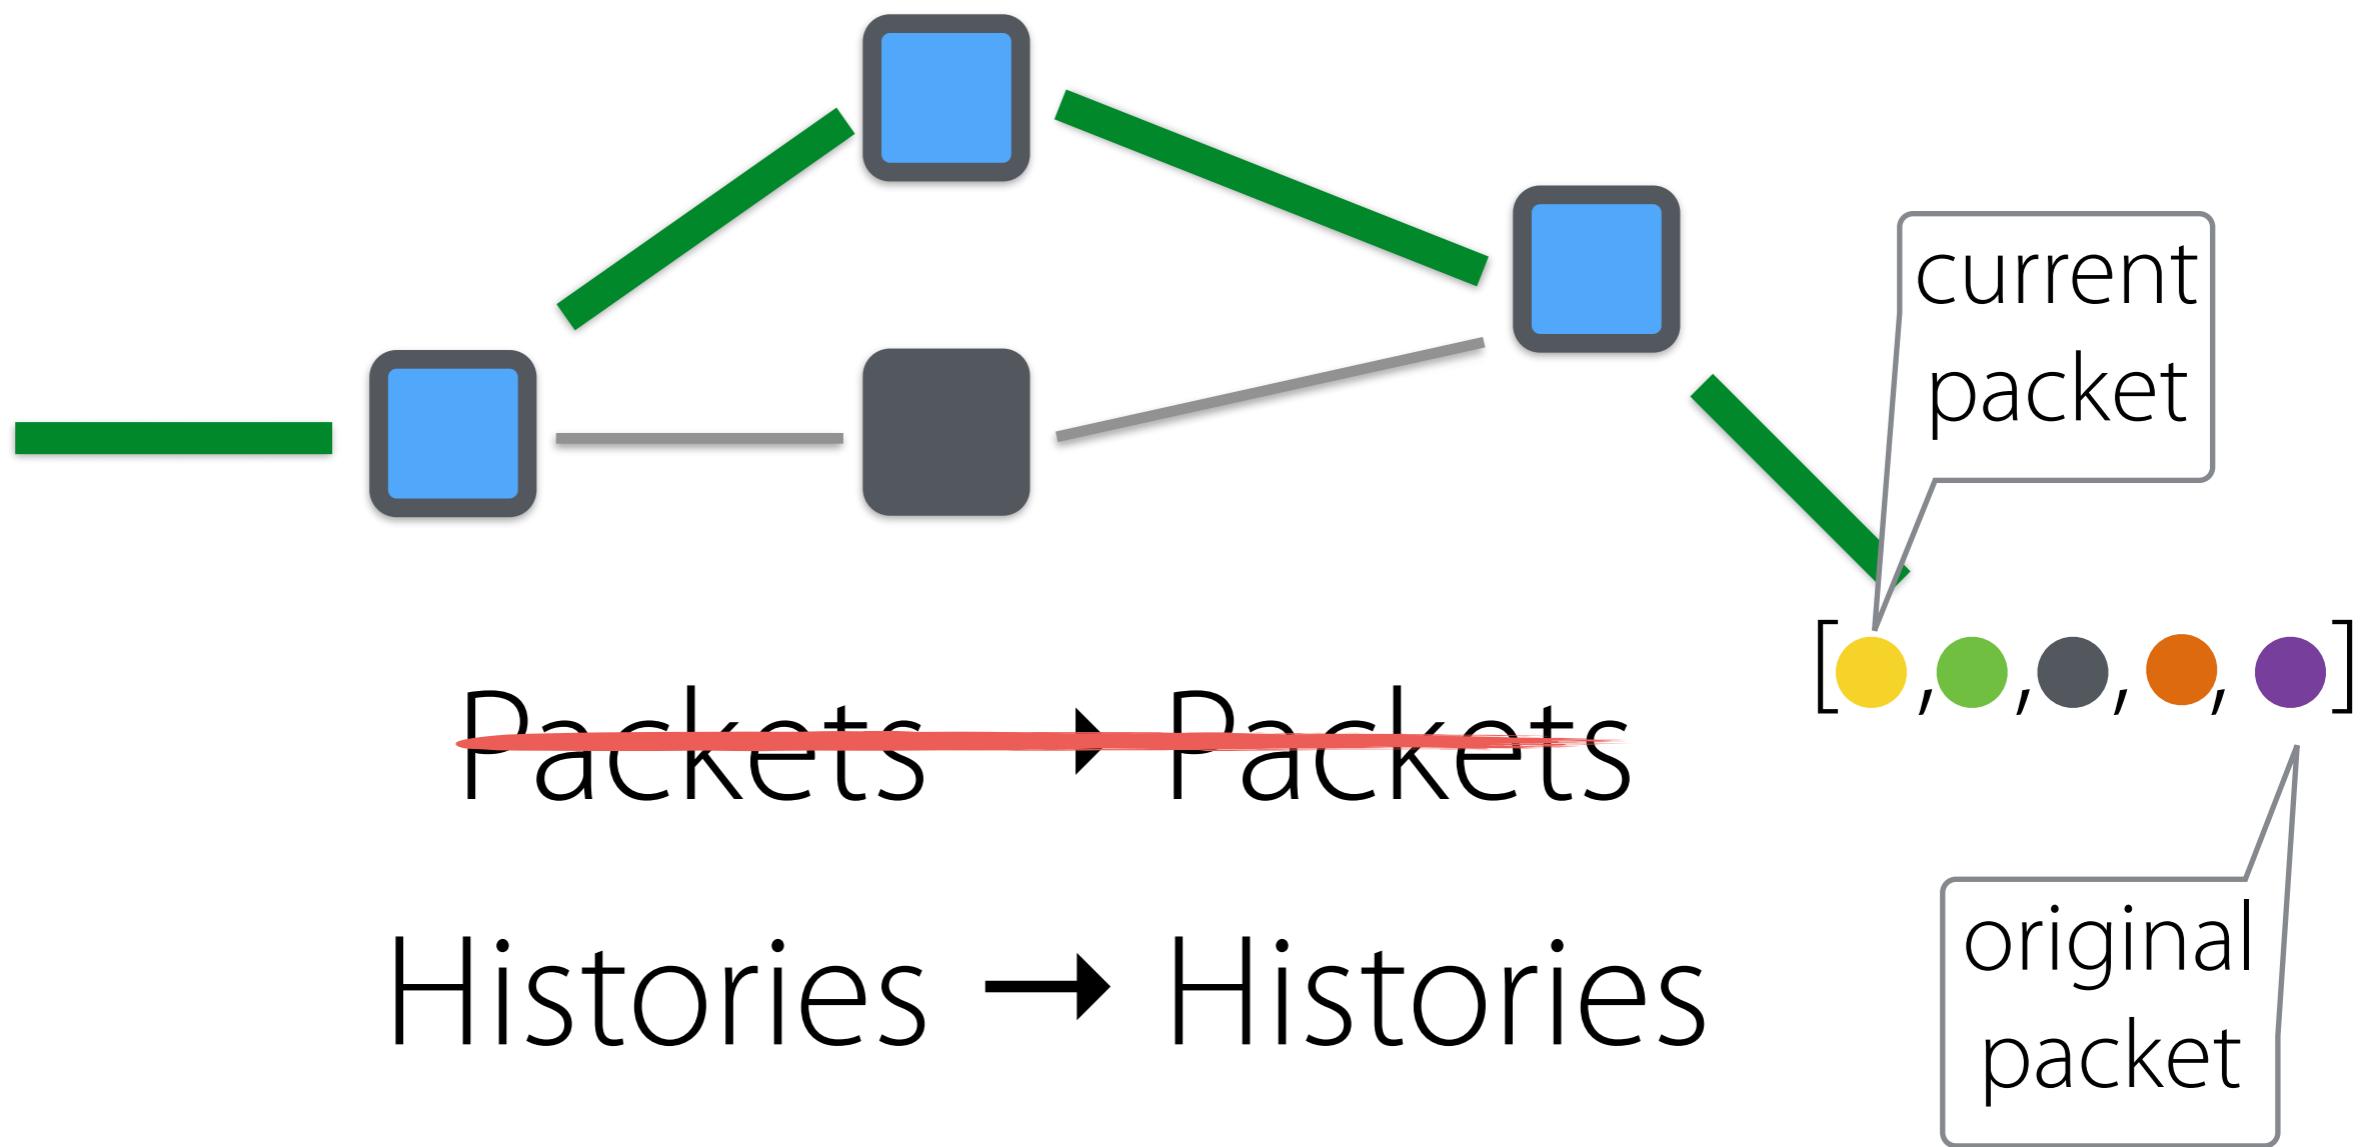
# Iteration

$(\text{topology} \bullet \text{switch})^*$



# Iteration

$(\text{topology} \bullet \text{switch})^*$

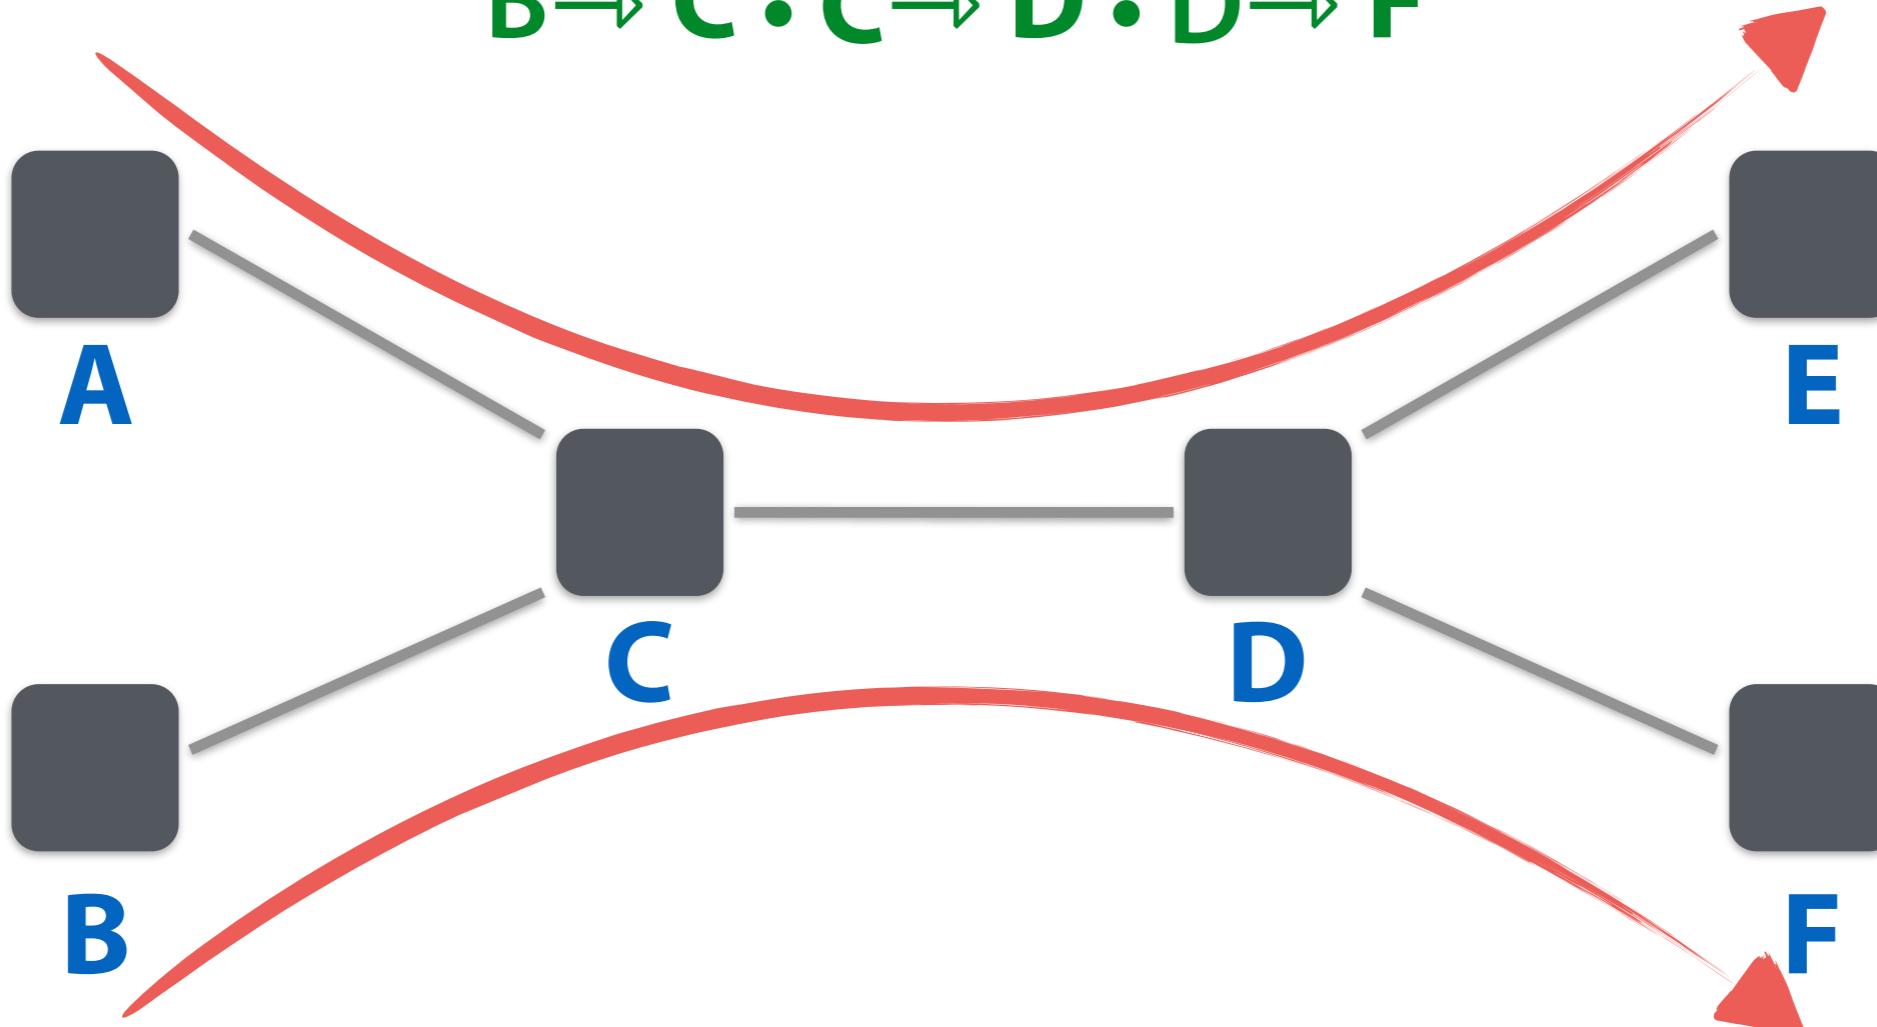


# Regular Paths

$A \rightarrow C \cdot C \rightarrow D \cdot E \rightarrow E$

+

$B \rightarrow C \cdot C \rightarrow D \cdot D \rightarrow F$



# NetKAT Syntax

$\text{pol} ::= \text{false}$   
|  $\text{true}$

Boolean  
Predicates



**if**  $b$  **then**  $p_1$  **else**  $p_2 \triangleq (b \bullet p_1) + (!b \bullet p_2)$

**while**  $b$  **do**  $p \triangleq (b \bullet p)^* \bullet !b$

$S \rightarrow S' \triangleq sw=S \bullet \mathbf{dup} \bullet sw:=S' \bullet \mathbf{dup}$

|  $\text{pol}$   
|  $\mathbf{dup}$

Primitives

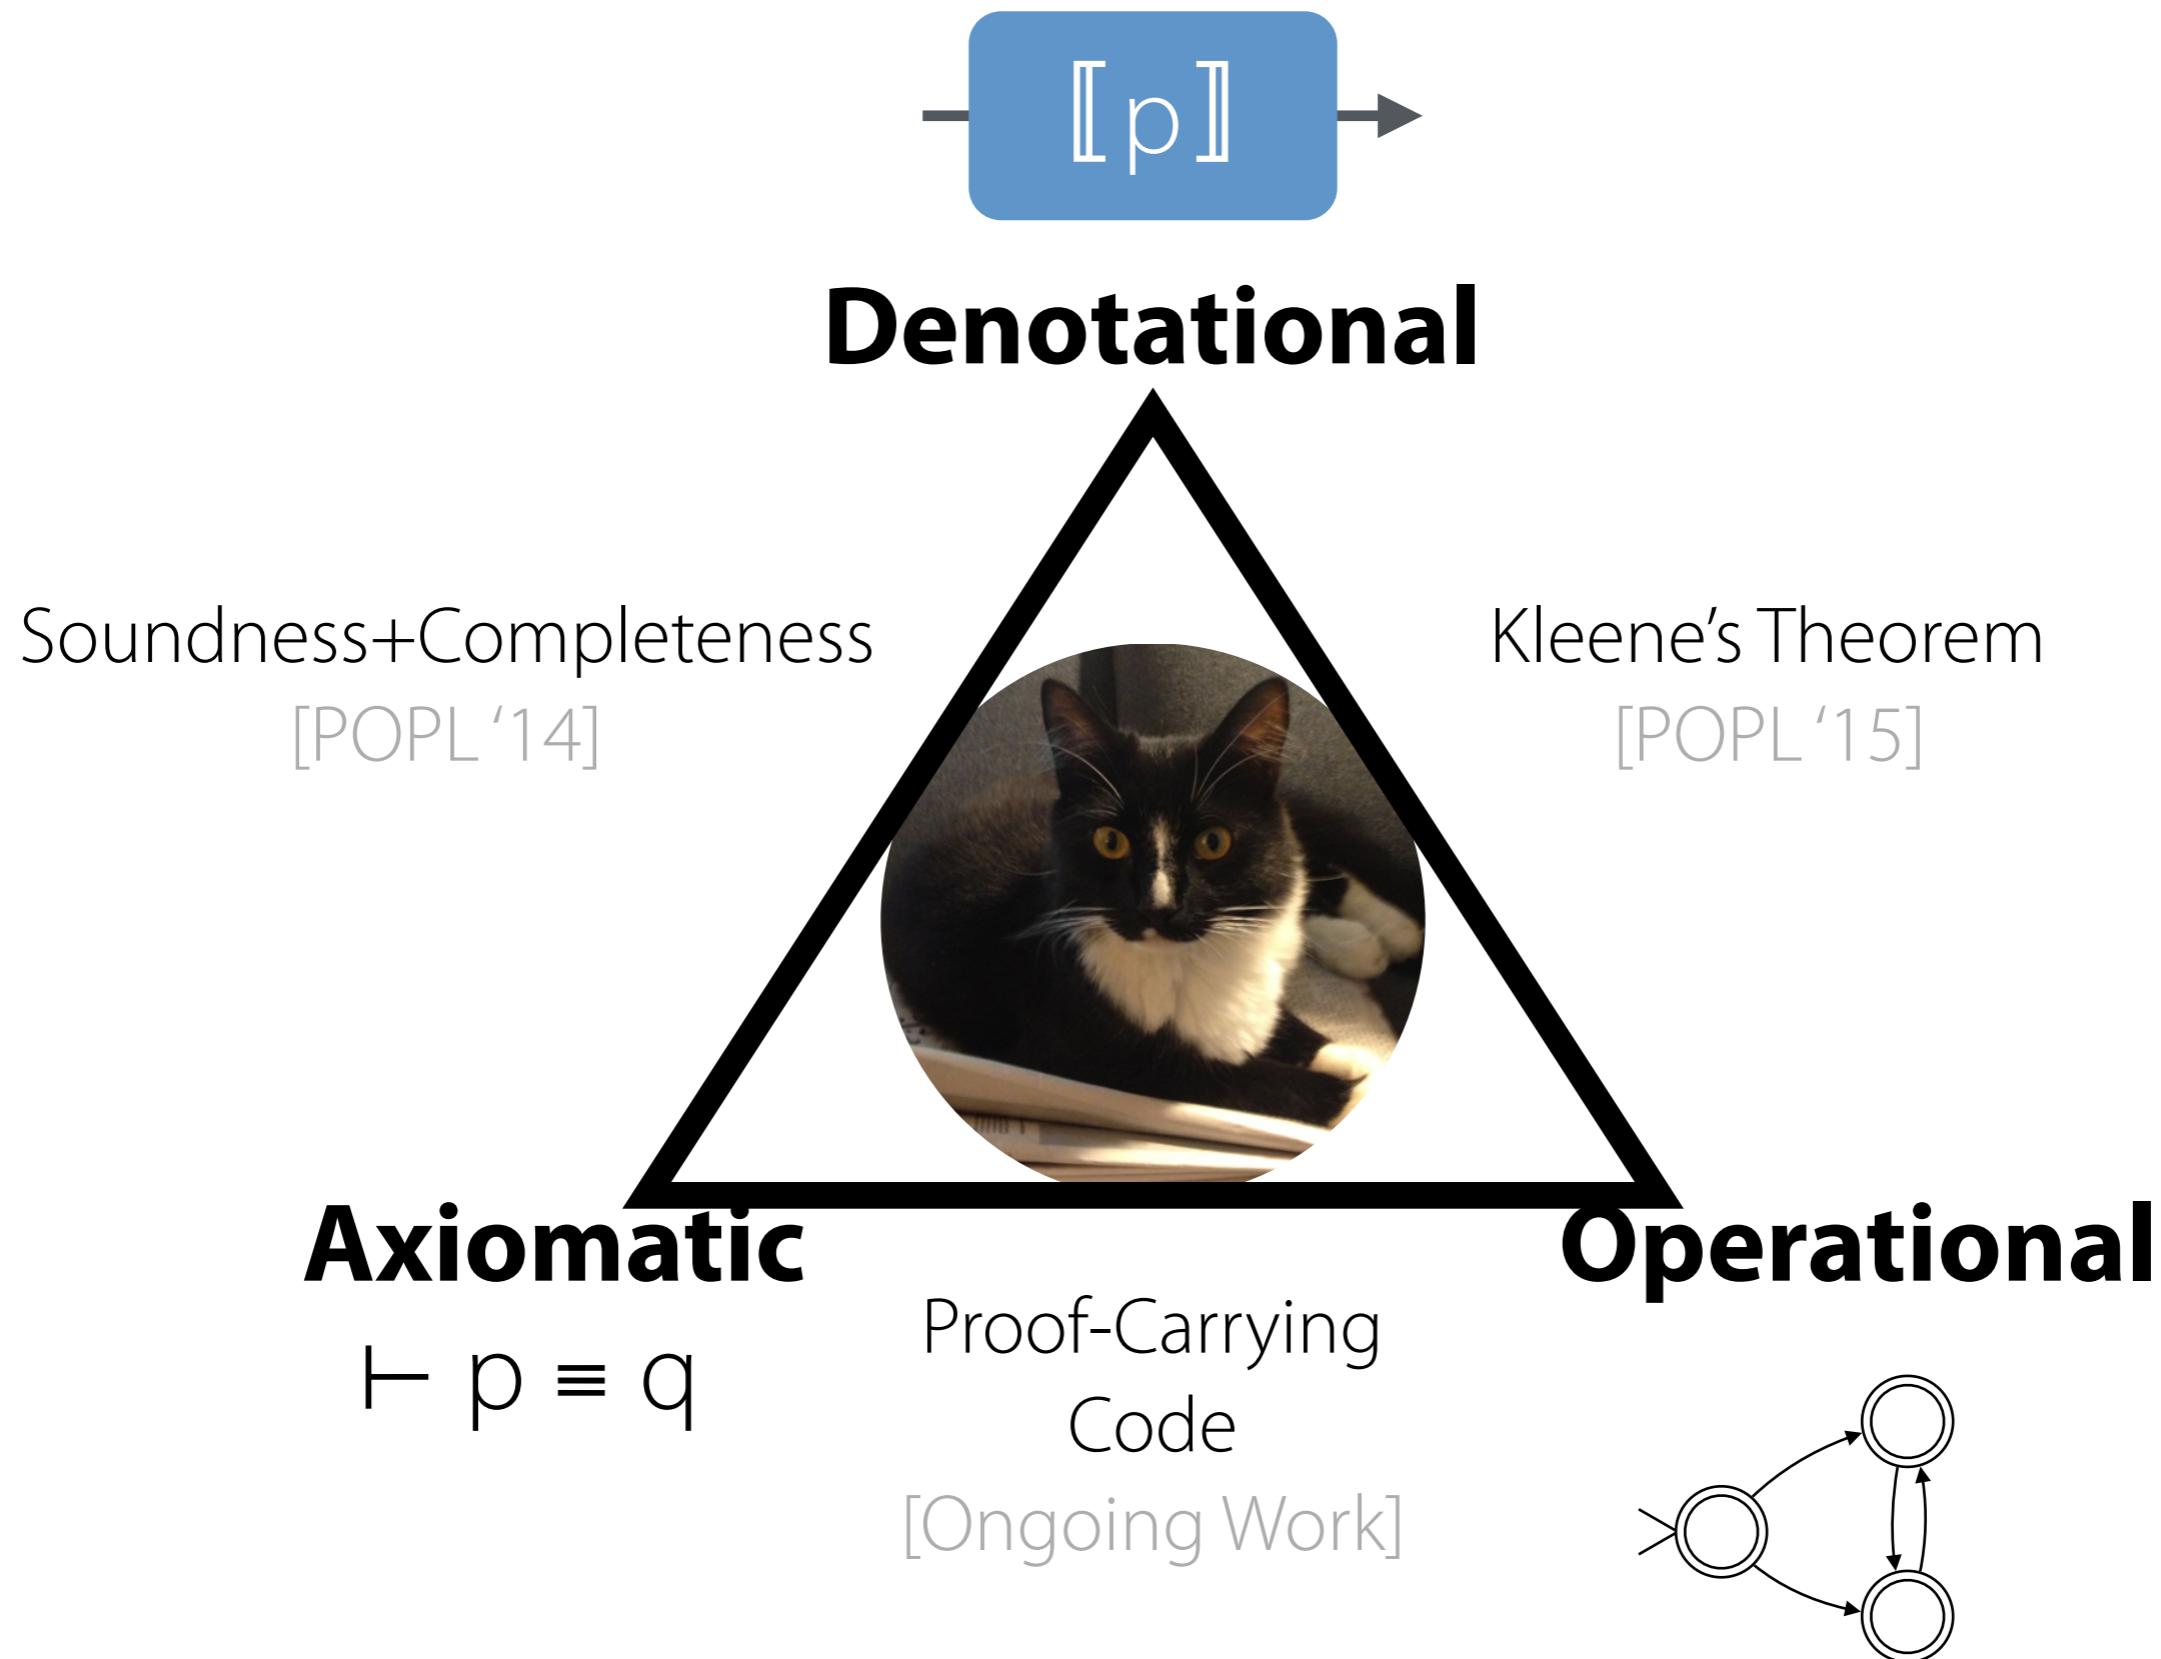


Negation may only be applied to Boolean *predicates*:  
**true**, **false**,  $f = n$ , closed under  $+$ ,  $\bullet$ , and  $!$

# **NetKAT**

# **Semantics**

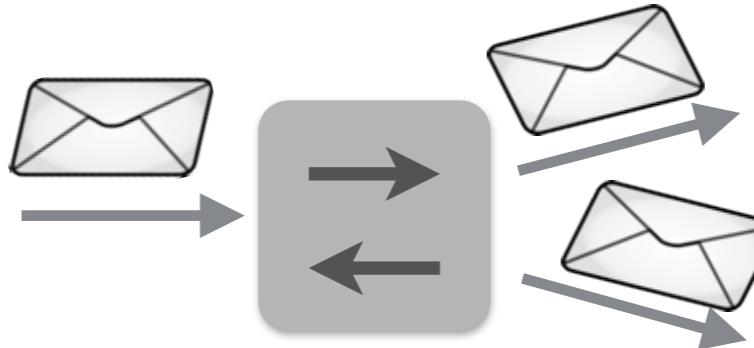
# NetKAT Semantics



# Denotational Semantics

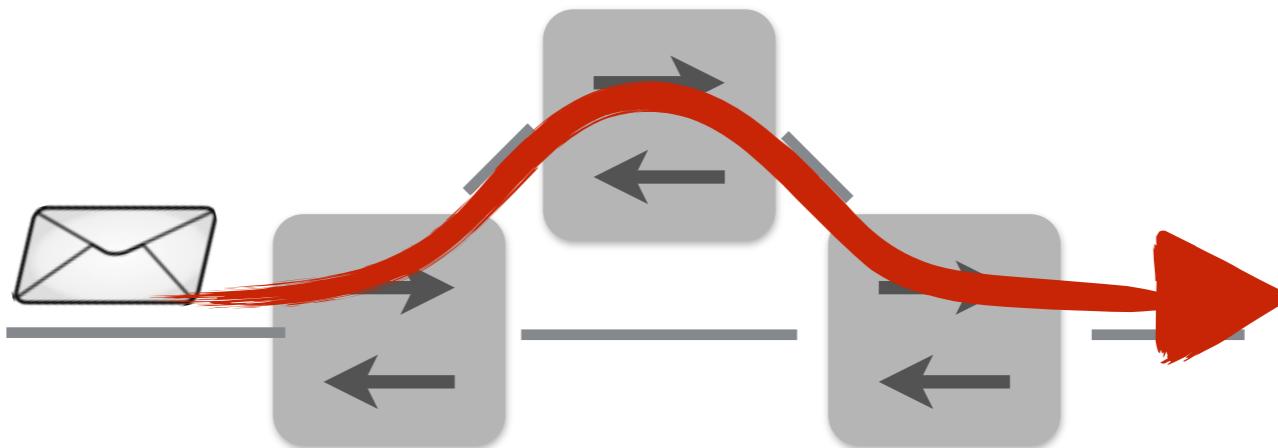
```
pol ::=  
| false  
| true  
| field = val  
| field := val  
| pol1 + pol2  
| pol1 • pol2  
| !pol  
| pol*  
| dup
```

Local: input-output behavior of switches



$\llbracket \text{pol} \rrbracket \in \text{Packets} \rightarrow \text{Packets}$

Global: network-wide paths



$\llbracket \text{pol} \rrbracket \in \text{Histories} \rightarrow \text{Histories}$

# NetKAT Semantics

$\llbracket \text{pol} \rrbracket \in \mathbf{History} \rightarrow \mathbf{History\ Set}$

$\llbracket \text{true} \rrbracket h = \{ h \}$

$\llbracket \text{false} \rrbracket h = \{ \}$

$\llbracket f = n \rrbracket pk :: h = \begin{cases} \{ pk :: h \} & \text{if } pk.f = n \\ \{ \} & \text{otherwise} \end{cases}$

$\llbracket ! \text{pol} \rrbracket h = \{ h \} \setminus \llbracket \text{pol} \rrbracket$

$\llbracket f := n \rrbracket pk :: h = \{ pk[f:=n] :: h \}$

$\llbracket \text{pol}_1 + \text{pol}_2 \rrbracket h = \llbracket \text{pol}_1 \rrbracket h \cup \llbracket \text{pol}_2 \rrbracket h$

$\llbracket \text{pol}_1 \bullet \text{pol}_2 \rrbracket h = (\llbracket \text{pol}_1 \rrbracket \bullet \llbracket \text{pol}_2 \rrbracket) h$

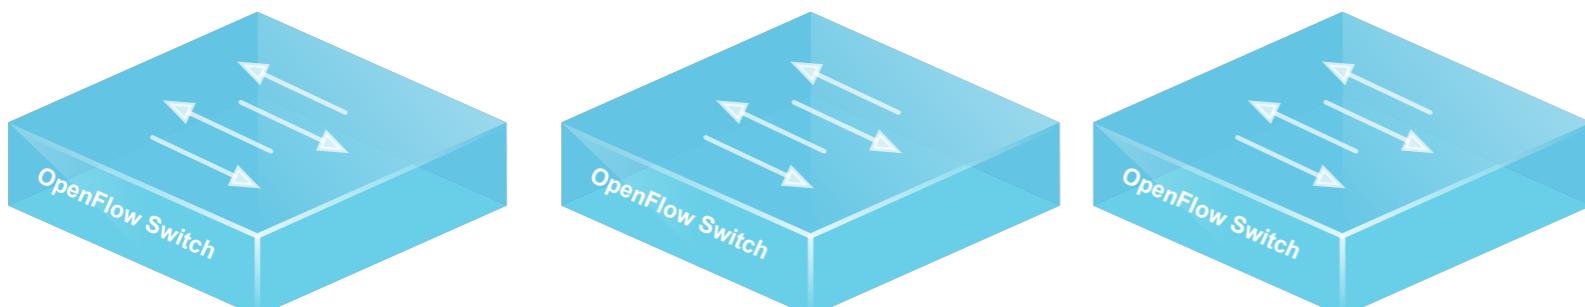
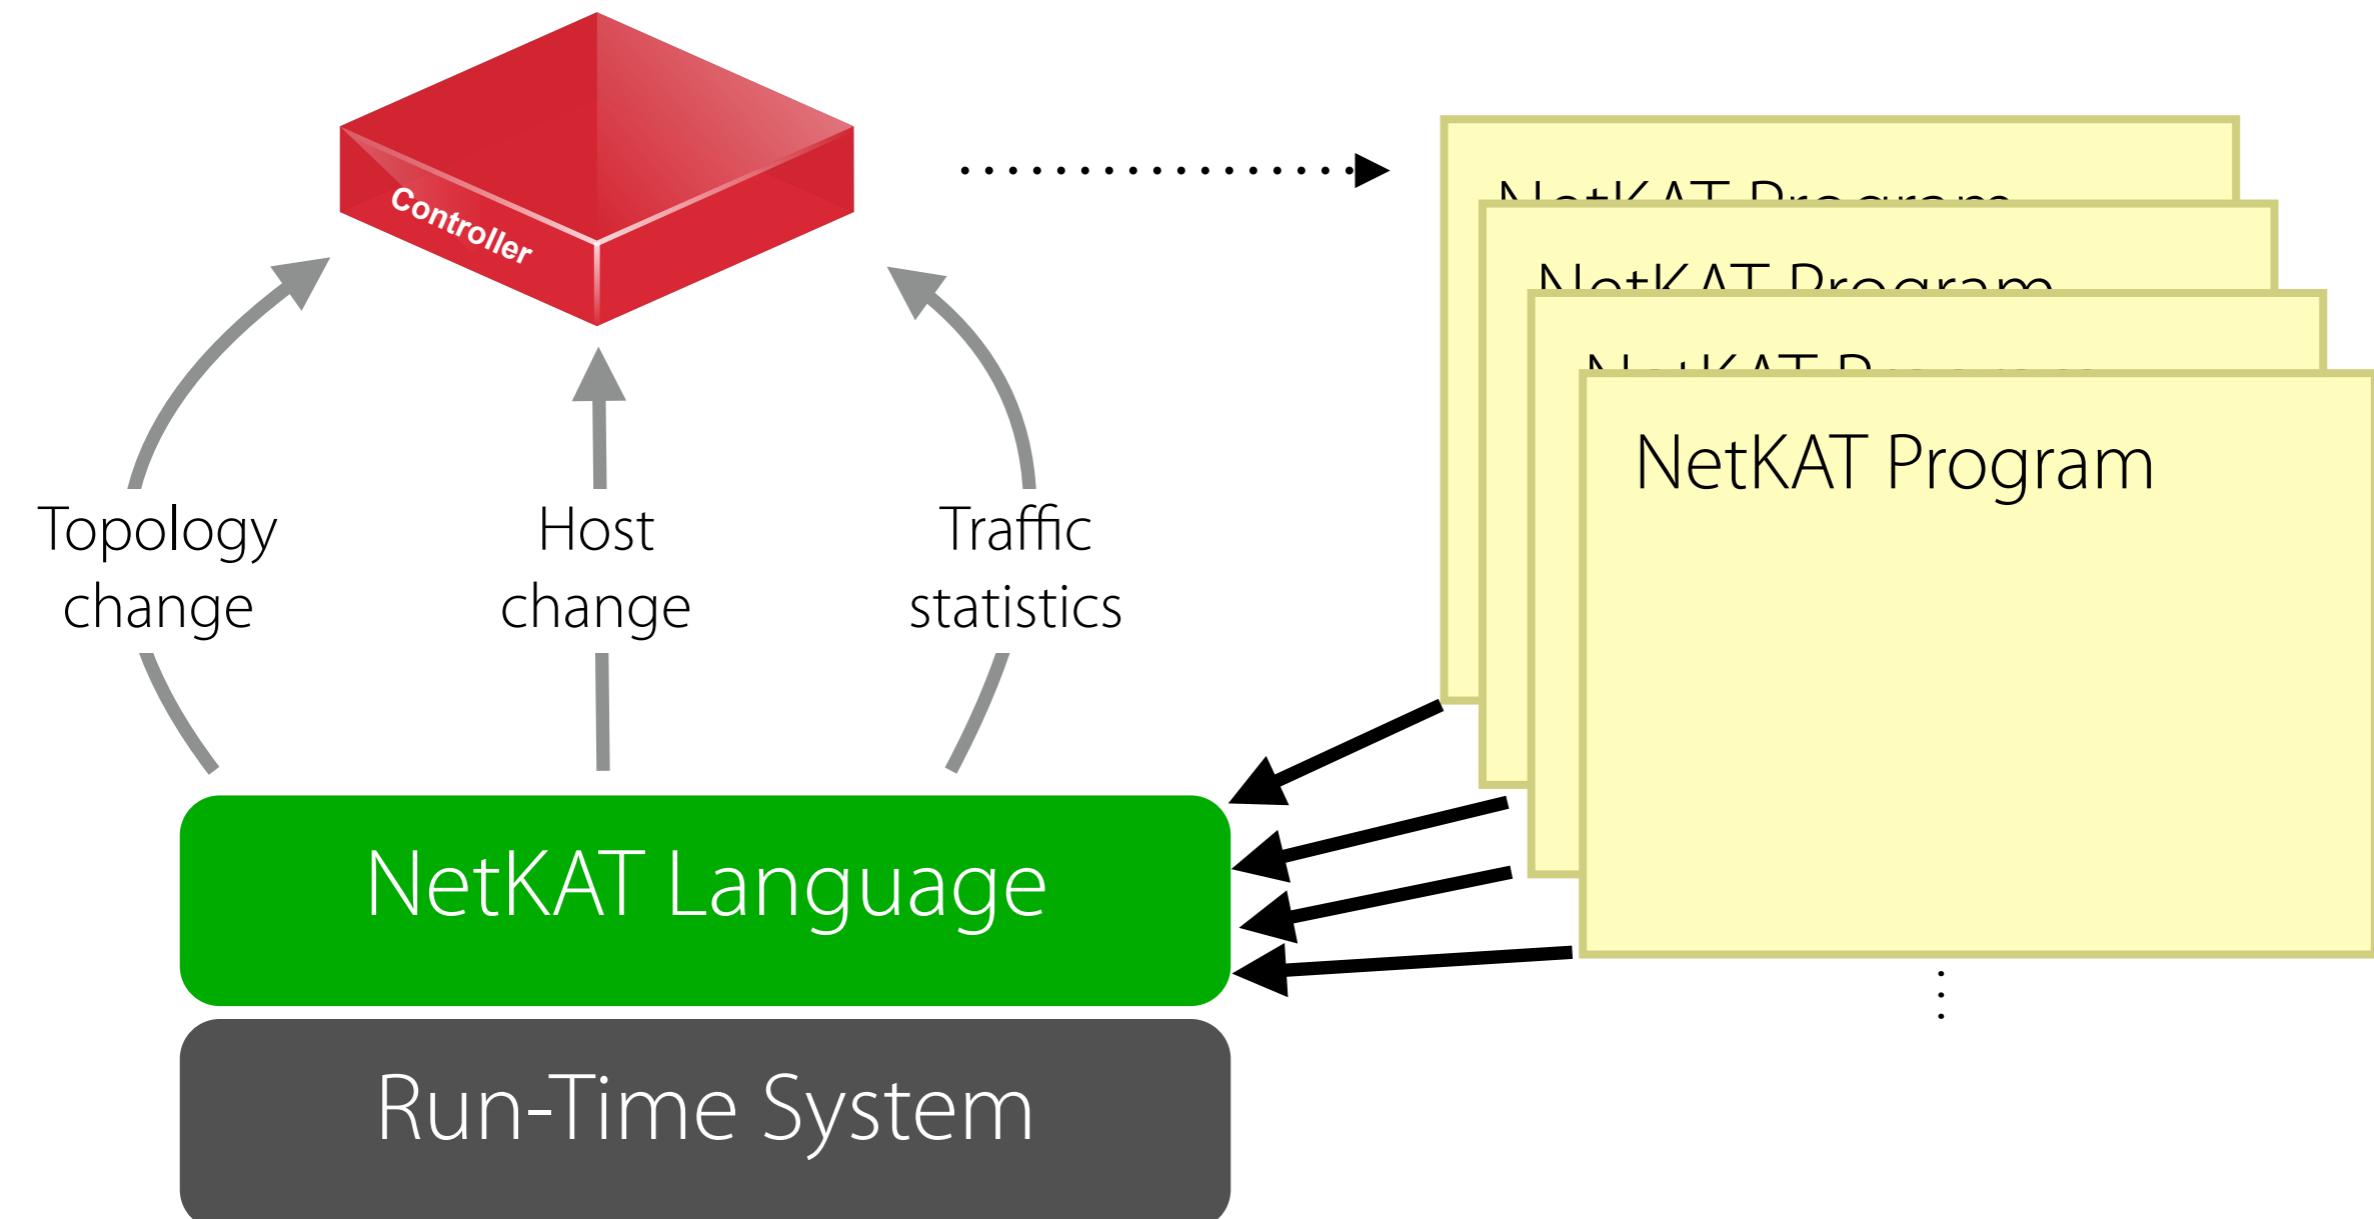
$\llbracket \text{pol}^* \rrbracket h = (\bigcup_i \llbracket \text{pol} \rrbracket^i h)$

$\llbracket S \rightarrow S' \rrbracket pk :: h = \begin{cases} \{ pk[\text{sw}:=S'] :: pk :: h \} & \text{if } pk.\text{sw} = S \\ \{ \} & \text{otherwise} \end{cases}$

$f, g \in \mathbf{History} \rightarrow \mathbf{History\ Set}$

$(f \bullet g) h = \bigcup \{ g h' \mid h' \in f h \}$

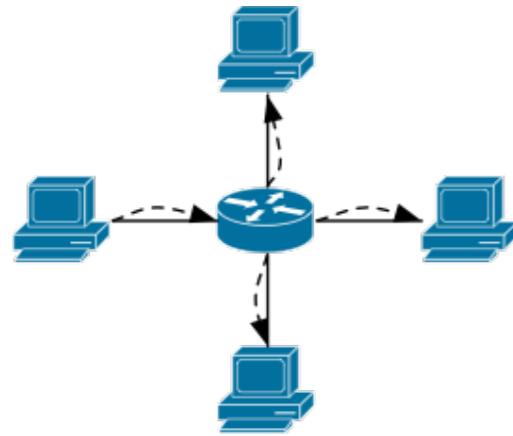
# Stateful Applications



# **Modular Composition**

# Example

## Topology



## Specification

- Forward packets to hosts 1-4
- Monitor traffic to unknown hosts
- Flood broadcast traffic to all hosts
- Disallow SSH traffic from hosts 1-2

## Flow Table

```
{pattern={ethSrc=00:00:00:00:00:01,ethTyp=0x800,ipProto=0x06, tcpDstPort=22},action=[]}  
{pattern={ethSrc=00:00:00:00:00:02,ethTyp=0x800,ipProto=0x06, tcpDstPort=22},action=[]}  
{pattern={ethDst=00:00:00:00:00:01},action=[Output(1)]}  
{pattern={ethDst=00:00:00:00:00:02},action=[Output(2)]}  
{pattern={ethDst=00:00:00:00:00:03},action=[Output(3)]}  
{pattern={ethDst=00:00:00:00:00:04},action=[Output(4)]}  
{pattern={ethDst=ff:ff:ff:ff:ff:ff,port=1},action=[Output(4), Output(3), Output(2)]}  
{pattern={ethDst=ff:ff:ff:ff:ff:ff,port=2},action=[Output(4), Output(3), Output(1)]}  
{pattern={ethDst=ff:ff:ff:ff:ff:ff,port=3},action=[Output(4), Output(2), Output(1)]}  
{pattern={ethDst=ff:ff:ff:ff:ff:ff,port=4},action=[Output(3), Output(2), Output(1)]}  
{pattern={ethDst=ff:ff:ff:ff:ff:ff},action=[]}  
{pattern={},action=[Controller]}
```

# Example: Forward

```
let forward =  
    if ethDst = 00:00:00:00:00:01 then  
        port := 1  
    else if ethDst = 00:00:00:00:00:02 then  
        port := 2  
    else if ethDst = 00:00:00:00:00:03 then  
        port := 3  
    else if ethDst = 00:00:00:00:00:04 then  
        port := 4  
    else  
        false
```

# Example: Broadcast

```
let flood =
  if port = 1 then
    port := 2 + port := 3 + port := 4
  else if port = 2 then
    port := 1 + port := 3 + port := 4
  else if port = 3 then
    port := 1 + port := 2 + port := 4
  else if port = 4 then
    port := 1 + port := 2 + port := 3
  else
    false

let broadcast =
  if ethDst = ff:ff:ff:ff:ff:ff then
    flood
  else
    false
```

# Example: Routing

```
let route = forward + broadcast
```

# Example: Monitor

```
let monitor =
  if !(ethDst = 00:00:00:00:00:01 +
      ethDst = 00:00:00:00:00:02 +
      ethDst = 00:00:00:00:00:03 +
      ethDst = 00:00:00:00:00:04 +
      ethDst = ff:ff:ff:ff:ff:ff) then
    port := unknown
  else
    false
```

# Example: Firewall

```
let firewall =
  if (ethSrc = 00:00:00:00:00:01 +
      ethSrc = 00:00:00:00:00:02) ;
    ethTyp = 0x800 ;
    ipProto = 0x06 ;
    tcpDstPort = 22 then
      false
    else
      true
```

# Example: Main Policy

```
let main = (route + monitor); firewall
```

compiles to...

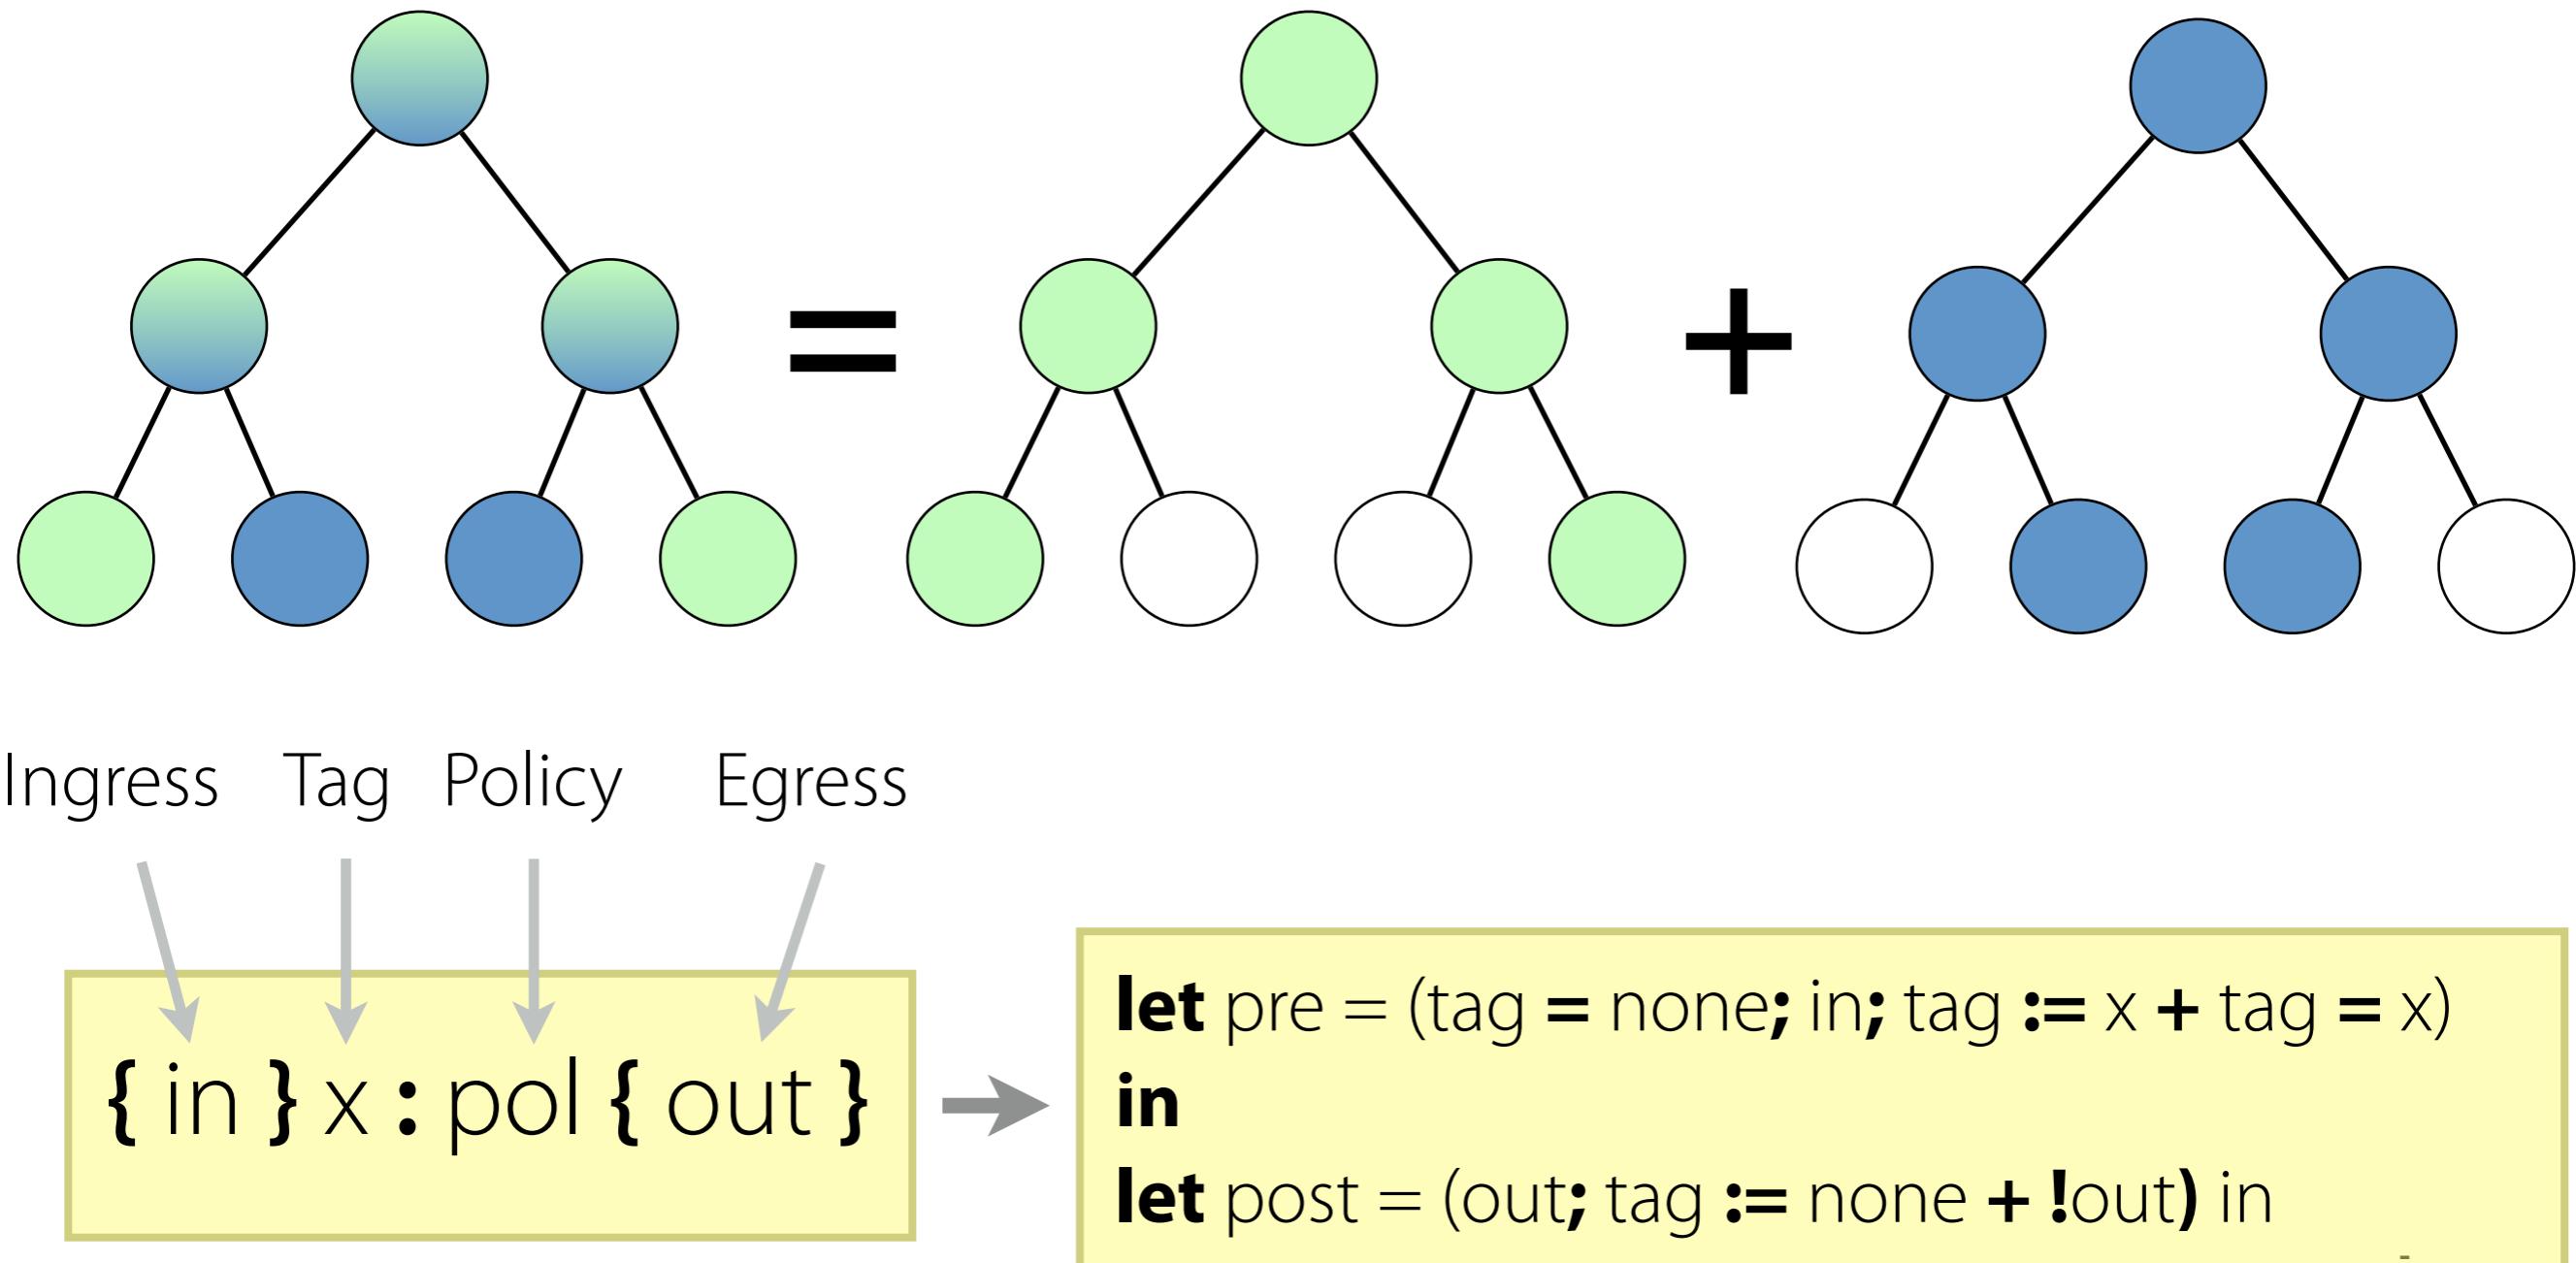


```
{pattern={ethSrc=00:00:00:00:00:01,ethTyp=0x800,ipProto=0x06, tcpDstPort=22},action=[]}
{pattern={ethSrc=00:00:00:00:00:02,ethTyp=0x800,ipProto=0x06, tcpDstPort=22},action=[]}
{pattern={ethDst=00:00:00:00:00:01},action=[Output(1)]}
{pattern={ethDst=00:00:00:00:00:02},action=[Output(2)]}
{pattern={ethDst=00:00:00:00:00:03},action=[Output(3)]}
{pattern={ethDst=00:00:00:00:00:04},action=[Output(4)]}
{pattern={ethDst=ff:ff:ff:ff:ff:ff, port=1},action=[Output(4), Output(3), Output(2)]}
{pattern={ethDst=ff:ff:ff:ff:ff:ff, port=2},action=[Output(4), Output(3), Output(1)]}
{pattern={ethDst=ff:ff:ff:ff:ff:ff, port=3},action=[Output(4), Output(2), Output(1)]}
{pattern={ethDst=ff:ff:ff:ff:ff:ff, port=4},action=[Output(3), Output(2), Output(1)]}
{pattern={},action=[Controller]}
```

# **Applications**

# Isolated Slices [POPL '14]

In many situations, multiple tenants must share the network...  
...but we don't want their traffic to interfere with each other!



# Virtualization [ICFP '15]

Often useful to programs against a simplified network topology



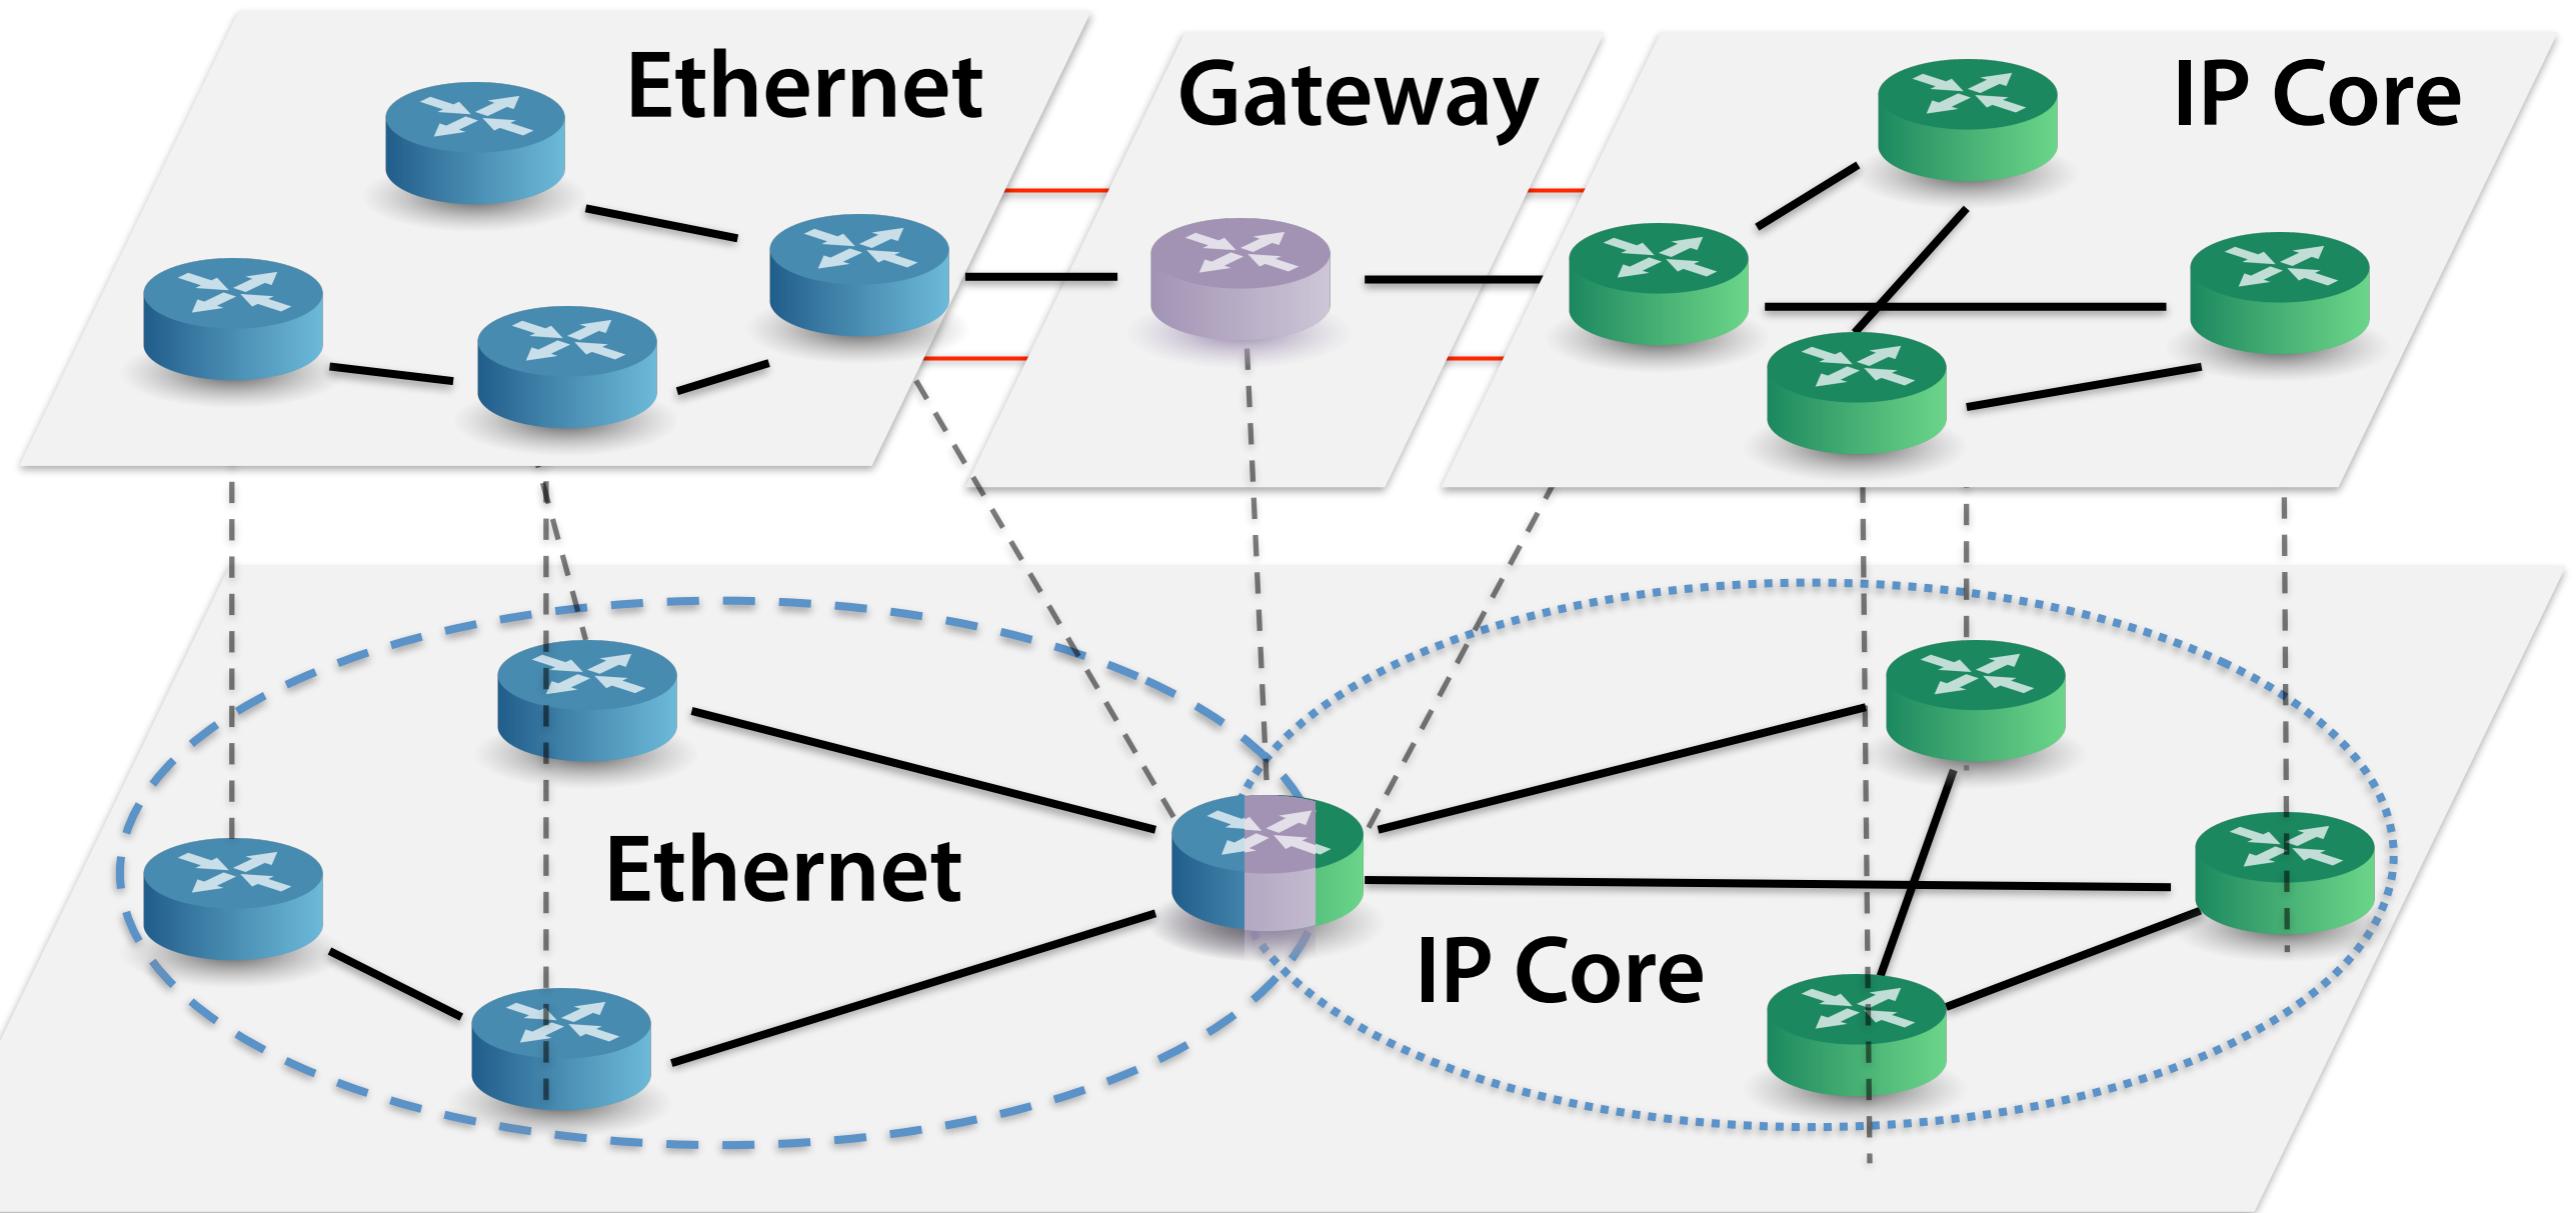
Concrete

Abstract

## Benefits:

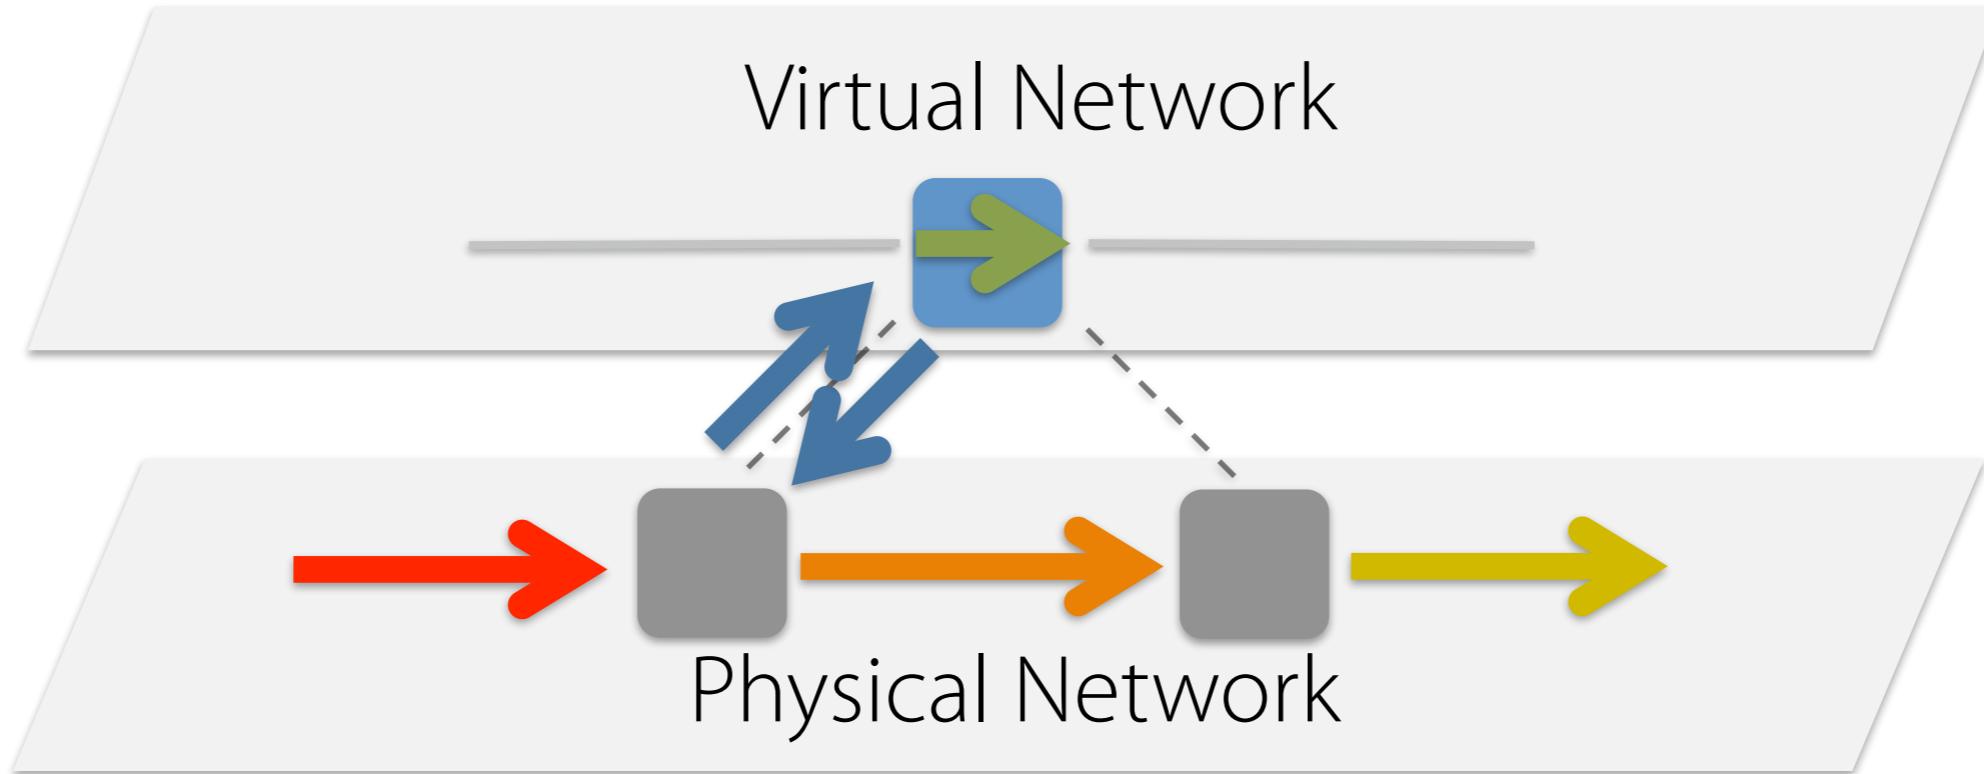
- *Information hiding*: limit what modules see
- *Protection*: limit what modules can do
- *Code reuse*: limit what dependencies modules have

# Example: Gateway



- **Left:** learning switch on MAC addresses
- **Middle:** ARP on gateway, plus simple repeater
- **Right:** shortest-path forwarding on IP prefixes

# Implementing Virtualization



This idiom can be implemented in NetKAT!

```
ingress;  
(raise; application; lower; fabric)*;  
egress
```

# Fault Tolerance [HotSDN '13]

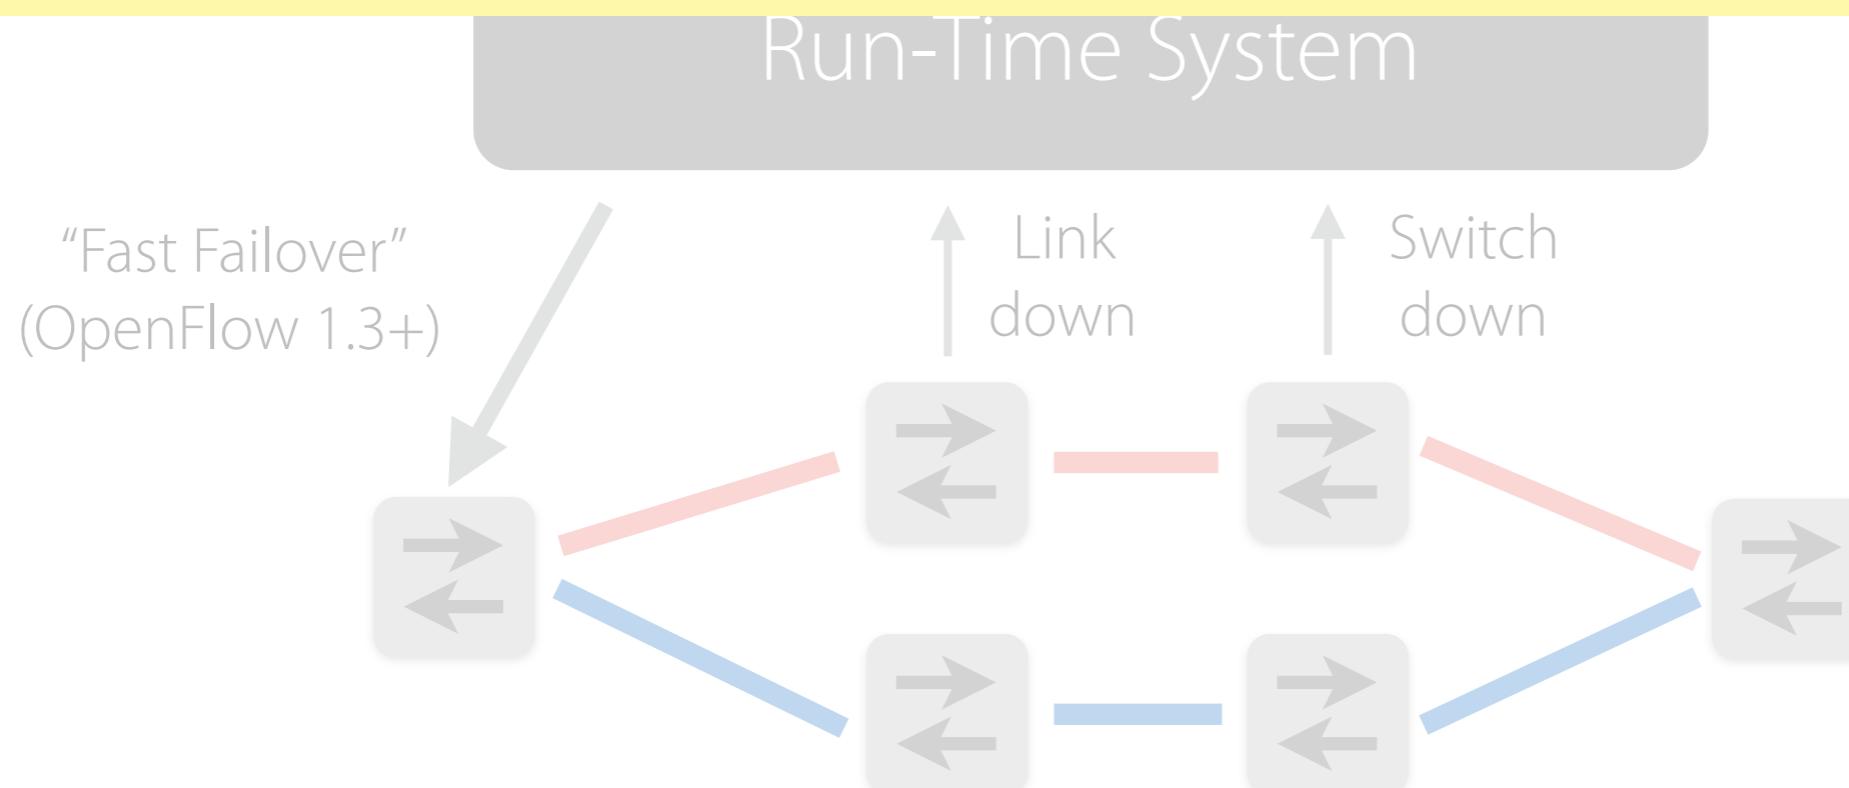
Path abstraction  
hides low-level faults

Red path

or

Blue path

NetKAT's path abstraction can also be used  
to build fault-tolerant networking applications



# **Wrapping Up...**

# Conclusion

- Programming languages have key role to play in the design and evolution of SDN programming interfaces
- NetKAT policy language provides a solid foundation for expressing and reasoning about packet-processing functions
- Many higher-level abstractions can be built on top of NetKAT
  - Isolated Slices
  - Virtual Networks
  - Fault Tolerance
  - and many others...

# Reading

- Carolyn Jane Anderson, Nate Foster, Arjun Guha, Jean-Baptiste Jeannin, Dexter Kozen, Cole Schlesinger, and David Walker.  
**NetKAT: Semantic Foundations for Networks.** In ACM SIGPLAN--SIGACT Symposium on Principles of Programming Languages (POPL), January 2014.
- Steffen Smolka, Spiridon Eliopoulos, Nate Foster, and Arjun Guha.  
**A Fast Compiler for NetKAT.** In ACM SIGPLAN International Conference on Functional Programming (ICFP), September 2015.
- Mark Reitblatt, Marco Canini, Arjun Guha, and Nate Foster. **FatTire: Declarative Fault Tolerance for Software-Defined Networks.** In ACM SIGCOMM Workshop on Hot Topics in Software Defined Networking (HotSDN), August 2013.

# Acknowledgments

- Carolyn Anderson (UMass)
- Shrutarshi Basu (Cornell)
- Marco Canini (KAUST)
- Arjun Guha (UMass)
- Jean-Baptiste Jeannin (Michigan)
- Dexter Kozen (Cornell)
- Robert Kleinberg (Cornell)
- Mark Reitblatt (Facebook)
- Cole Schlesinger (Amazon)
- Emin Gün Sirer (Cornell)
- Robert Soulé (Lugano)
- Dave Walker (Princeton)

# Later today...

## I: Introduction

- Semantics Primer
- Software-Defined Networking

## II: NetKAT

- Language Design
- Formal Semantics

## III: Probabilistic NetKAT

- Formal Semantics
- Applications