

# Modular Control Plane Verification via Temporal Invariants

Timothy Alberdingk Thijm

Princeton University

Aarti Gupta

Princeton University

Ryan Beckett

Microsoft Research


David Walker

Princeton University

PLDI 2023




# What is the Control Plane?



# What is the Control Plane?


Goal: determine routes to use to forward traffic



# What is the Control Plane?

Goal: determine routes to use to forward traffic

Send initial route announcements




# What is the Control Plane?

Goal: determine routes to use to forward traffic

Send initial route announcements

Receive announcements, process  
according to **configs**






# What is the Control Plane?

Goal: determine routes to use to forward traffic

Send initial route announcements

Receive announcements, process  
according to **configs**



```
bat examples/INTERNET2/configs/newy32aoa.cfg
5349 policy-statement CAAREN-IN {
5350   term participant {
5351     from {
5352       protocol bgp;
5353       prefix-list-filter CAAREN-PARTICIPANT orlonger;
5354     }
5355     then next policy;
5356   }
5357   term segp {
5358     from {
5359       protocol bgp;
5360       prefix-list-filter CAAREN-SEGP orlonger;
5361     }
5362     then {
5363       community add SEGP;
5364       next policy;
5365     }
5366   }
5367   term sponsored {
5368     from {
5369       protocol bgp;
5370       prefix-list-filter CAAREN-SPONSORED orlonger;
5371     }
5372     then {
5373       community add SPONSORED;
5374       next policy;
5375     }
5376   }
5377   term reject-unicast {
5378     then reject;
5379   }
5380 }
```

Distributed, written in vendor-specific, low-level language


# What is the Control Plane?

Goal: determine routes to use to forward traffic

Send initial route announcements

Receive announcements, process  
according to **configs**

Select best announcement



```
bat examples/INTERNET2/configs/newy32aoa.cfg
5349  policy-statement CAAREN-IN {
5350      term participant {
5351          from {
5352              protocol bgp;
5353              prefix-list-filter CAAREN-PARTICIPANT orlonger;
5354          }
5355          then next policy;
5356      }
5357      term segp {
5358          from {
5359              protocol bgp;
5360              prefix-list-filter CAAREN-SEGP orlonger;
5361          }
5362          then {
5363              community add SEGP;
5364              next policy;
5365          }
5366      }
5367      term sponsored {
5368          from {
5369              protocol bgp;
5370              prefix-list-filter CAAREN-SPONSORED orlonger;
5371          }
5372          then {
5373              community add SPONSORED;
5374              next policy;
5375          }
5376      }
5377      term reject-unicast {
5378          then reject;
5379      }
5380  }
```

Distributed, written in vendor-specific, low-level language

# What is the Control Plane?


Goal: determine routes to use to forward traffic

Send initial route announcements

Receive announcements, process  
according to **configs**

Select best announcement

Broadcast selected route to  
neighbors



```
bat examples/INTERNET2/configs/newy32aoa.cfg
5349 policy-statement CAAREN-IN {
5350   term participant {
5351     from {
5352       protocol bgp;
5353       prefix-list-filter CAAREN-PARTICIPANT orlonger;
5354     }
5355     then next policy;
5356   }
5357   term segp {
5358     from {
5359       protocol bgp;
5360       prefix-list-filter CAAREN-SEGP orlonger;
5361     }
5362     then {
5363       community add SEGP;
5364       next policy;
5365     }
5366   }
5367   term sponsored {
5368     from {
5369       protocol bgp;
5370       prefix-list-filter CAAREN-SPONSORED orlonger;
5371     }
5372     then {
5373       community add SPONSORED;
5374       next policy;
5375     }
5376   }
5377   term reject-unicast {
5378     then reject;
5379   }
5380 }
```

Distributed, written in vendor-  
specific, low-level language



# What is the Control Plane?

Goal: determine routes to use to forward traffic


Send initial route announcements

Receive announcements, process  
according to **configs**

Select best announcement

Broadcast selected route to  
neighbors

Converge to a **stable state**



```
bat examples/INTERNET2/configs/newy32aoa.cfg
5349  policy-statement CAAREN-IN {
5350      term participant {
5351          from {
5352              protocol bgp;
5353              prefix-list-filter CAAREN-PARTICIPANT orlonger;
5354          }
5355          then next policy;
5356      }
5357      term segp {
5358          from {
5359              protocol bgp;
5360              prefix-list-filter CAAREN-SEGP orlonger;
5361          }
5362          then {
5363              community add SEGP;
5364              next policy;
5365          }
5366      }
5367      term sponsored {
5368          from {
5369              protocol bgp;
5370              prefix-list-filter CAAREN-SPONSORED orlonger;
5371          }
5372          then {
5373              community add SPONSORED;
5374              next policy;
5375          }
5376      }
5377      term reject-unicast {
5378          then reject;
5379      }
5380  }
```

Distributed, written in vendor-  
specific, low-level language

# What is the Control Plane?

Goal: determine routes to use to forward traffic


Send initial route announcements

Receive announcements, process according to **configs**

Select best announcement

Broadcast selected route to neighbors

Converge to a **stable state**



Policies for security, traffic engineering, fault tolerance, cost, *etc.*

```
bat examples/INTERNET2/configs/newy32aoa.cfg
5349 policy-statement CAAREN-IN {
5350   term participant {
5351     from {
5352       protocol bgp;
5353       prefix-list-filter CAAREN-PARTICIPANT orlonger;
5354     }
5355     then next policy;
5356   }
5357   term segp {
5358     from {
5359       protocol bgp;
5360       prefix-list-filter CAAREN-SEGP orlonger;
5361     }
5362     then {
5363       community add SEGP;
5364       next policy;
5365     }
5366   }
5367   term sponsored {
5368     from {
5369       protocol bgp;
5370       prefix-list-filter CAAREN-SPONSORED orlonger;
5371     }
5372     then {
5373       community add SPONSORED;
5374       next policy;
5375     }
5376   }
5377   term reject-unicast {
5378     then reject;
5379   }
5380 }
```

Distributed, written in vendor-specific, low-level language

# How Do We Verify Control Planes?

```
bat examples/INTERNET2/configs/newy32aon.cfg
5349 policy-statement CAAMEN-IN {
5350   term participant {
5351     from {
5352       protocol tcp;
5353       prefix-list-filter CAAMEN-PARTICIPANT or-larger;
5354     }
5355     then next-policy;
5356   }
5357   term segg {
5358     from {
5359       protocol tcp;
5360       prefix-list-filter CAAMEN-SEGP or-larger;
5361     }
5362     then {
5363       community-add SEGP;
5364       next-policy;
5365     }
5366   }
5367   term sponsored {
5368     from {
5369       protocol tcp;
5370       prefix-list-filter CAAMEN-SPONSORED or-larger;
5371     }
5372     then {
5373       community-add SPONSORED;
5374       next-policy;
5375     }
5376   }
5377   term reject-outbound {
5378     then reject;
5379   }
5380 }
```

Network configuration files

Policies for security, traffic engineering,  
fault tolerance, cost, *etc.*

# How Do We Verify Control Planes?

```
bat examples/INTERNET2/configs/newy32aon.cfg
5349 policy-statement CAMEN-IN {
5350   term participant {
5351     from {
5352       protocol tcp;
5353       prefix-list-filter CAMEN-PARTICIPANT or-larger;
5354     }
5355     then next-policy;
5356   }
5357   term seep {
5358     from {
5359       protocol tcp;
5360       prefix-list-filter CAMEN-SEEP or-larger;
5361     }
5362     then {
5363       community add SEEP;
5364       next-policy;
5365     }
5366   }
5367   term sponsored {
5368     from {
5369       protocol tcp;
5370       prefix-list-filter CAMEN-SPONSORED or-larger;
5371     }
5372     then {
5373       community add SPONSORED;
5374       next-policy;
5375     }
5376   }
5377   term reject-outside {
5378     then reject;
5379   }
5380 }
```

Network configuration files

Policies for security, traffic engineering,  
fault tolerance, cost, *etc.*



Analyze all configurations together to  
find property violations using a  
control plane verifier  
(e.g., Batfish, ARC, Minesweeper,  
Bagpipe, Tiramisu, Plankton, Hoyan)



# How Do We Verify Control Planes?

```
bat examples/INTERNET2/configs/newy32aon.cfg
5349 policy-statement CAREN-IN {
5350   term participant {
5351     from {
5352       protocol tcp;
5353       prefix-list-filter CAREN-PARTICIPANT or-larger;
5354     }
5355     then next-policy;
5356   }
5357   term seep {
5358     from {
5359       protocol tcp;
5360       prefix-list-filter CAREN-SEEP or-larger;
5361     }
5362     then {
5363       community add SEEP;
5364       next-policy;
5365     }
5366   }
5367   term sponsored {
5368     from {
5369       protocol tcp;
5370       prefix-list-filter CAREN-SPONSORED or-larger;
5371     }
5372     then {
5373       community add SPONSORED;
5374       next-policy;
5375     }
5376   }
5377   term reject-outside {
5378     then reject;
5379   }
5380 }
```

Network configuration files

Policies for security, traffic engineering,  
fault tolerance, cost, *etc.*



Analyze all configurations together to  
find property violations using a  
control plane verifier  
(e.g., Batfish, ARC, Minesweeper,  
Bagpipe, Tiramisu, Plankton, Hoyan)

Repeat when configurations change



# Scaling Control Plane Verification



Many networks are **too big and too complex** to verify monolithically!



# Scaling Control Plane Verification



**modular verification** to the rescue!



Many networks are **too big and too complex** to verify monolithically!



# Our Contributions



demonstrate why naive stable states analysis is unsuitable for modular verification




present time-based theory for modular control plane analysis, with SMT-based verification procedure




verify properties of 2000-node data centers and complex wide-area networks in seconds!

# How (Not) to Modularly Verify



# Modular Network Verification






# Modular Network Verification

**sound modular analysis:**

captures all monolithic routing behavior




# Modular Network Verification

## sound modular analysis:

captures all monolithic routing behavior

split the network up into **node-local components** to verify independently




# Modular Network Verification

## **sound modular analysis:**


captures all monolithic routing behavior

split the network up into **node-local components** to verify independently


represent cross-component dependencies using **interfaces**




# Interfaces



User





# Interfaces



interface  $A$  over-approximates the converged states of the node  $v$  with a set of states  $A(v)$




# Bear's Modular Verification Procedure



## Verification condition:


If  $v$  receives any routes satisfying  $A(u_1), A(u_2), \dots, A(u_m)$ ,  
does its selected route satisfy  $A(v)$ ?

# An Example Network

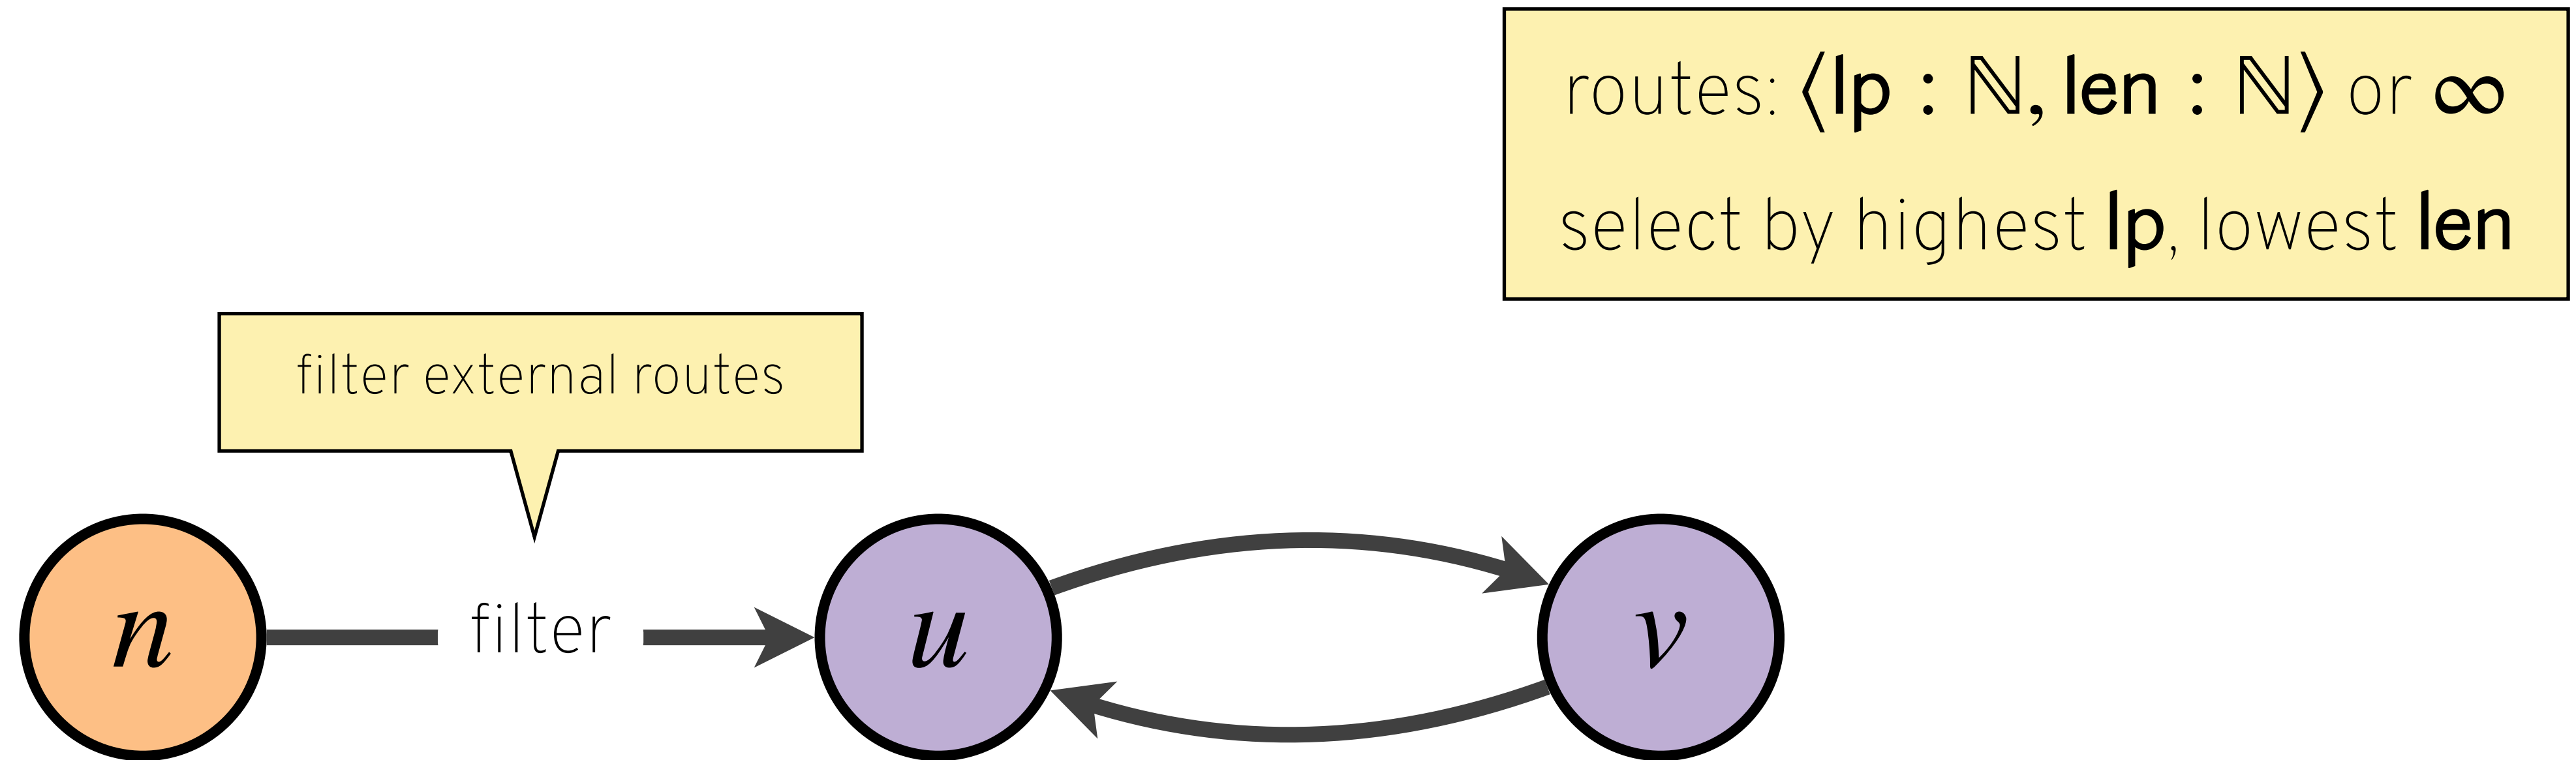


# An Example Network

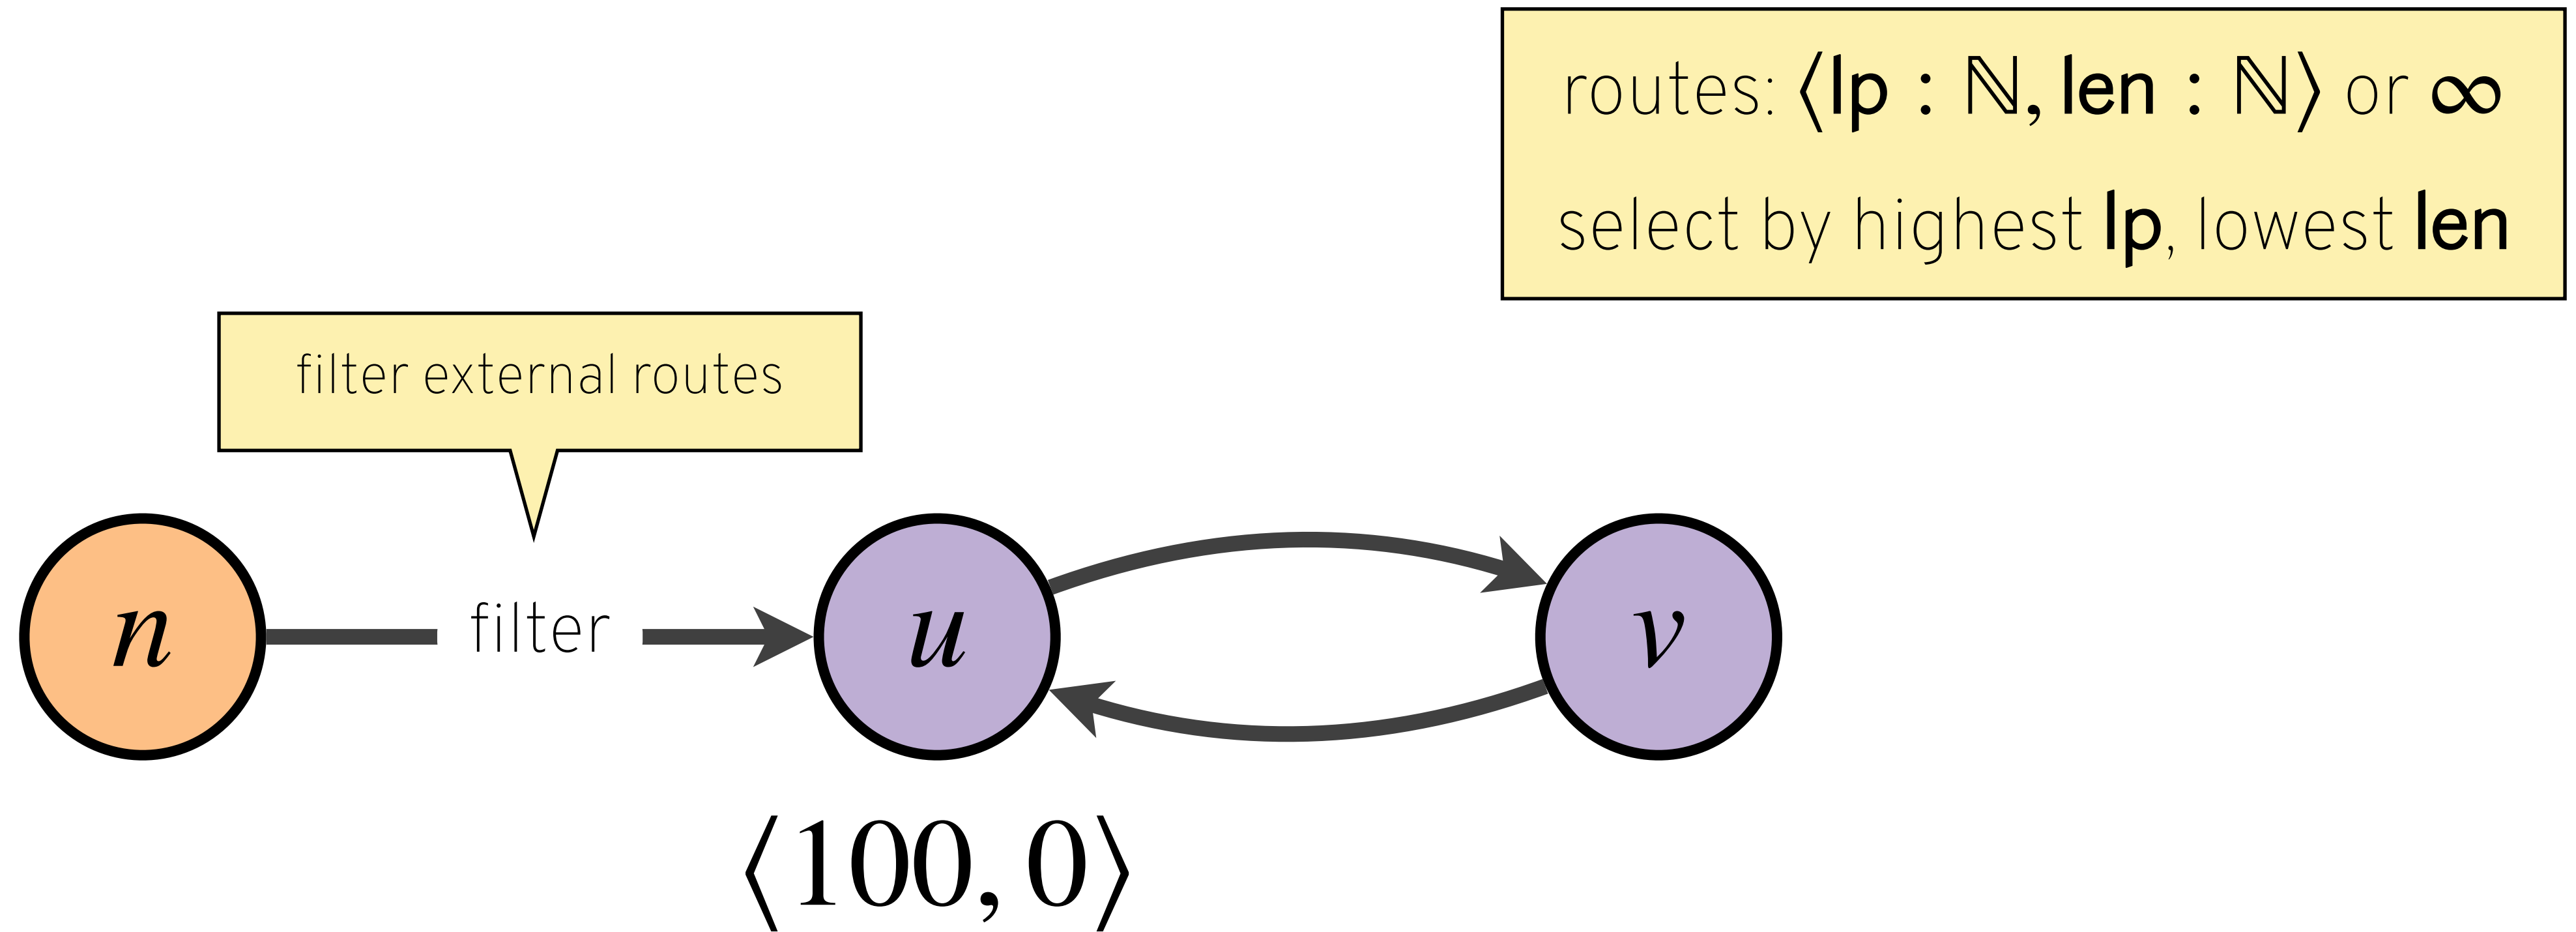
routes:  $\langle \text{lp} : \mathbb{N}, \text{len} : \mathbb{N} \rangle$  or  $\infty$   
select by highest **lp**, lowest **len**



# An Example Network

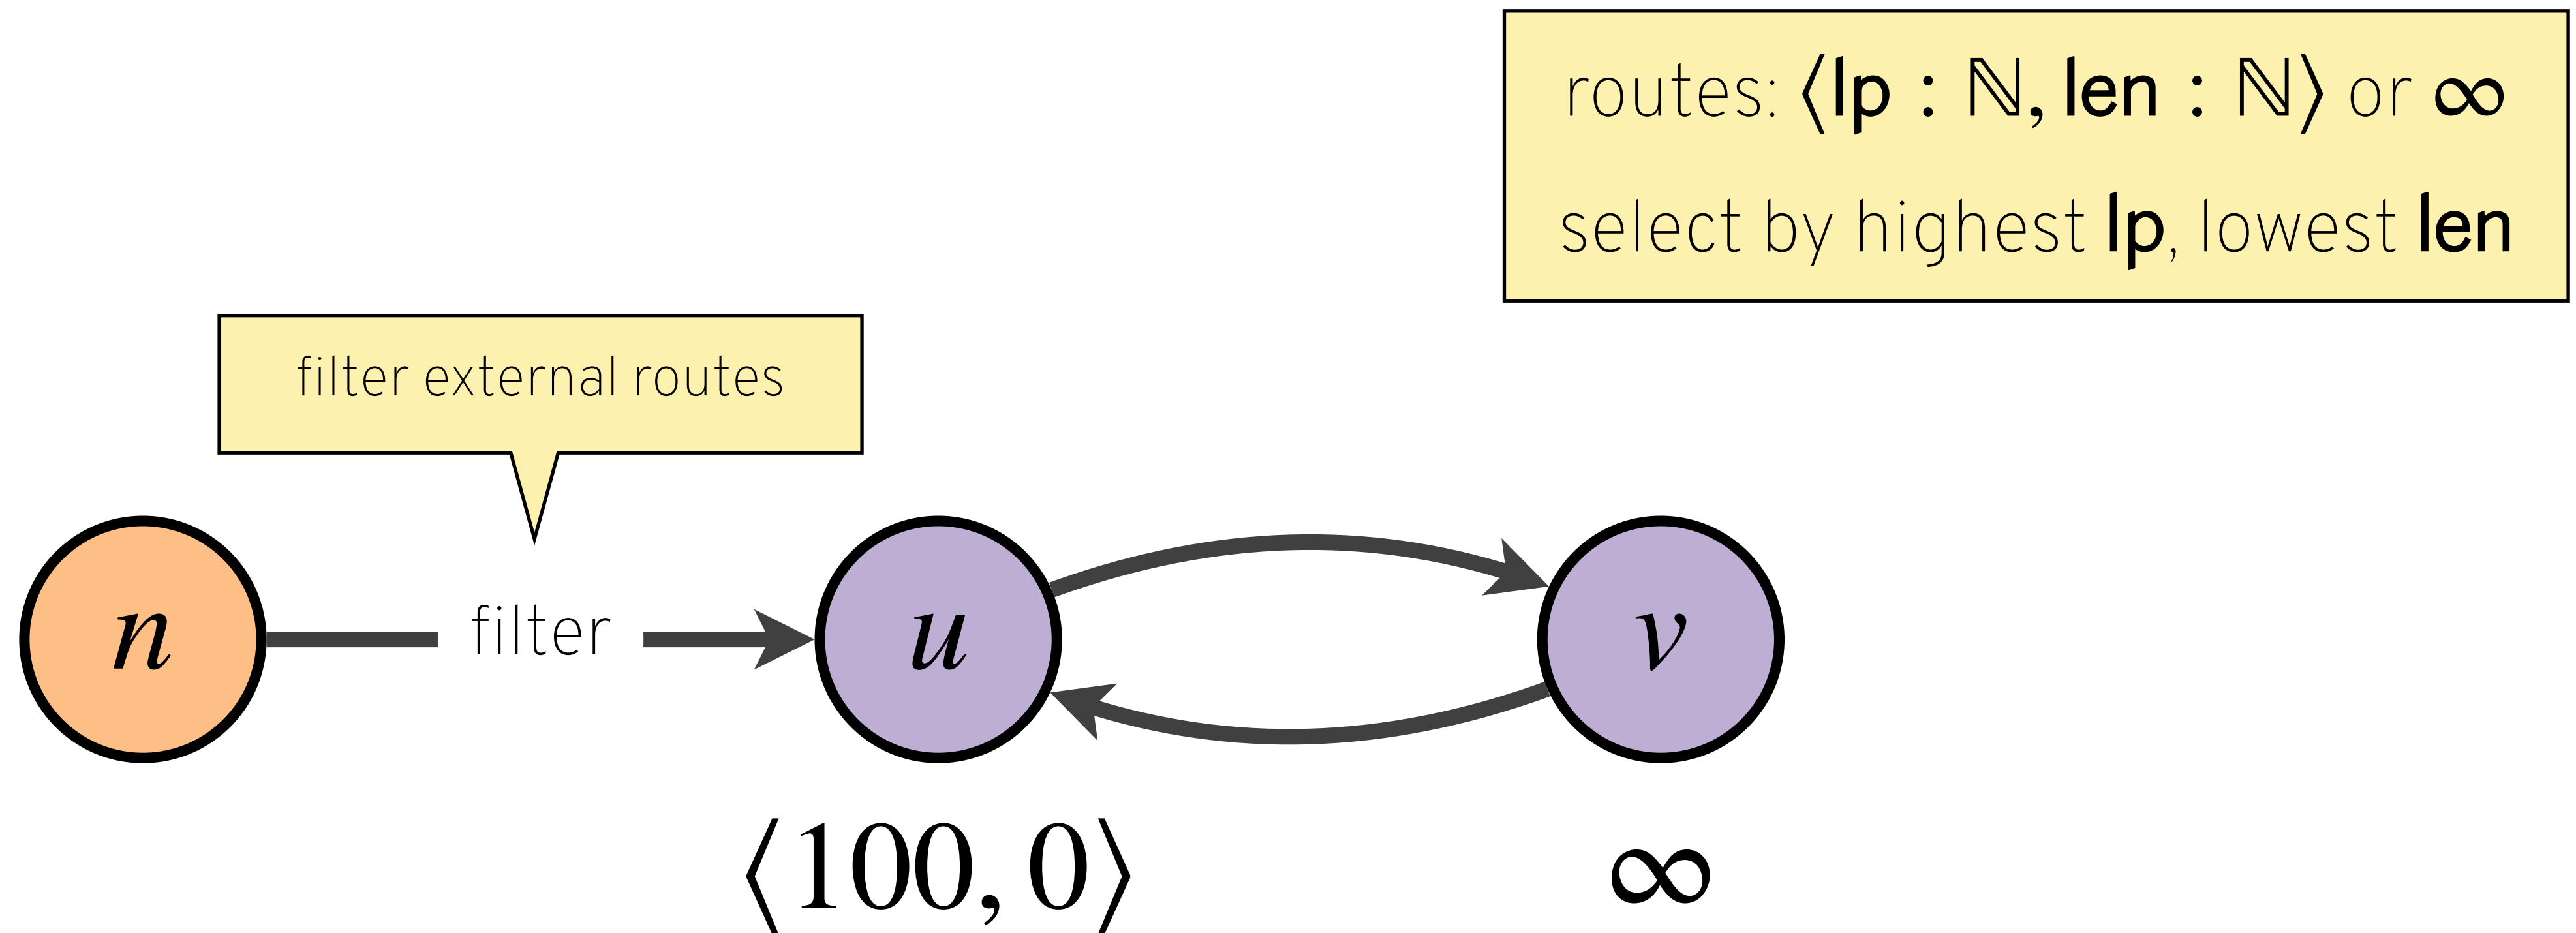


# An Example Network

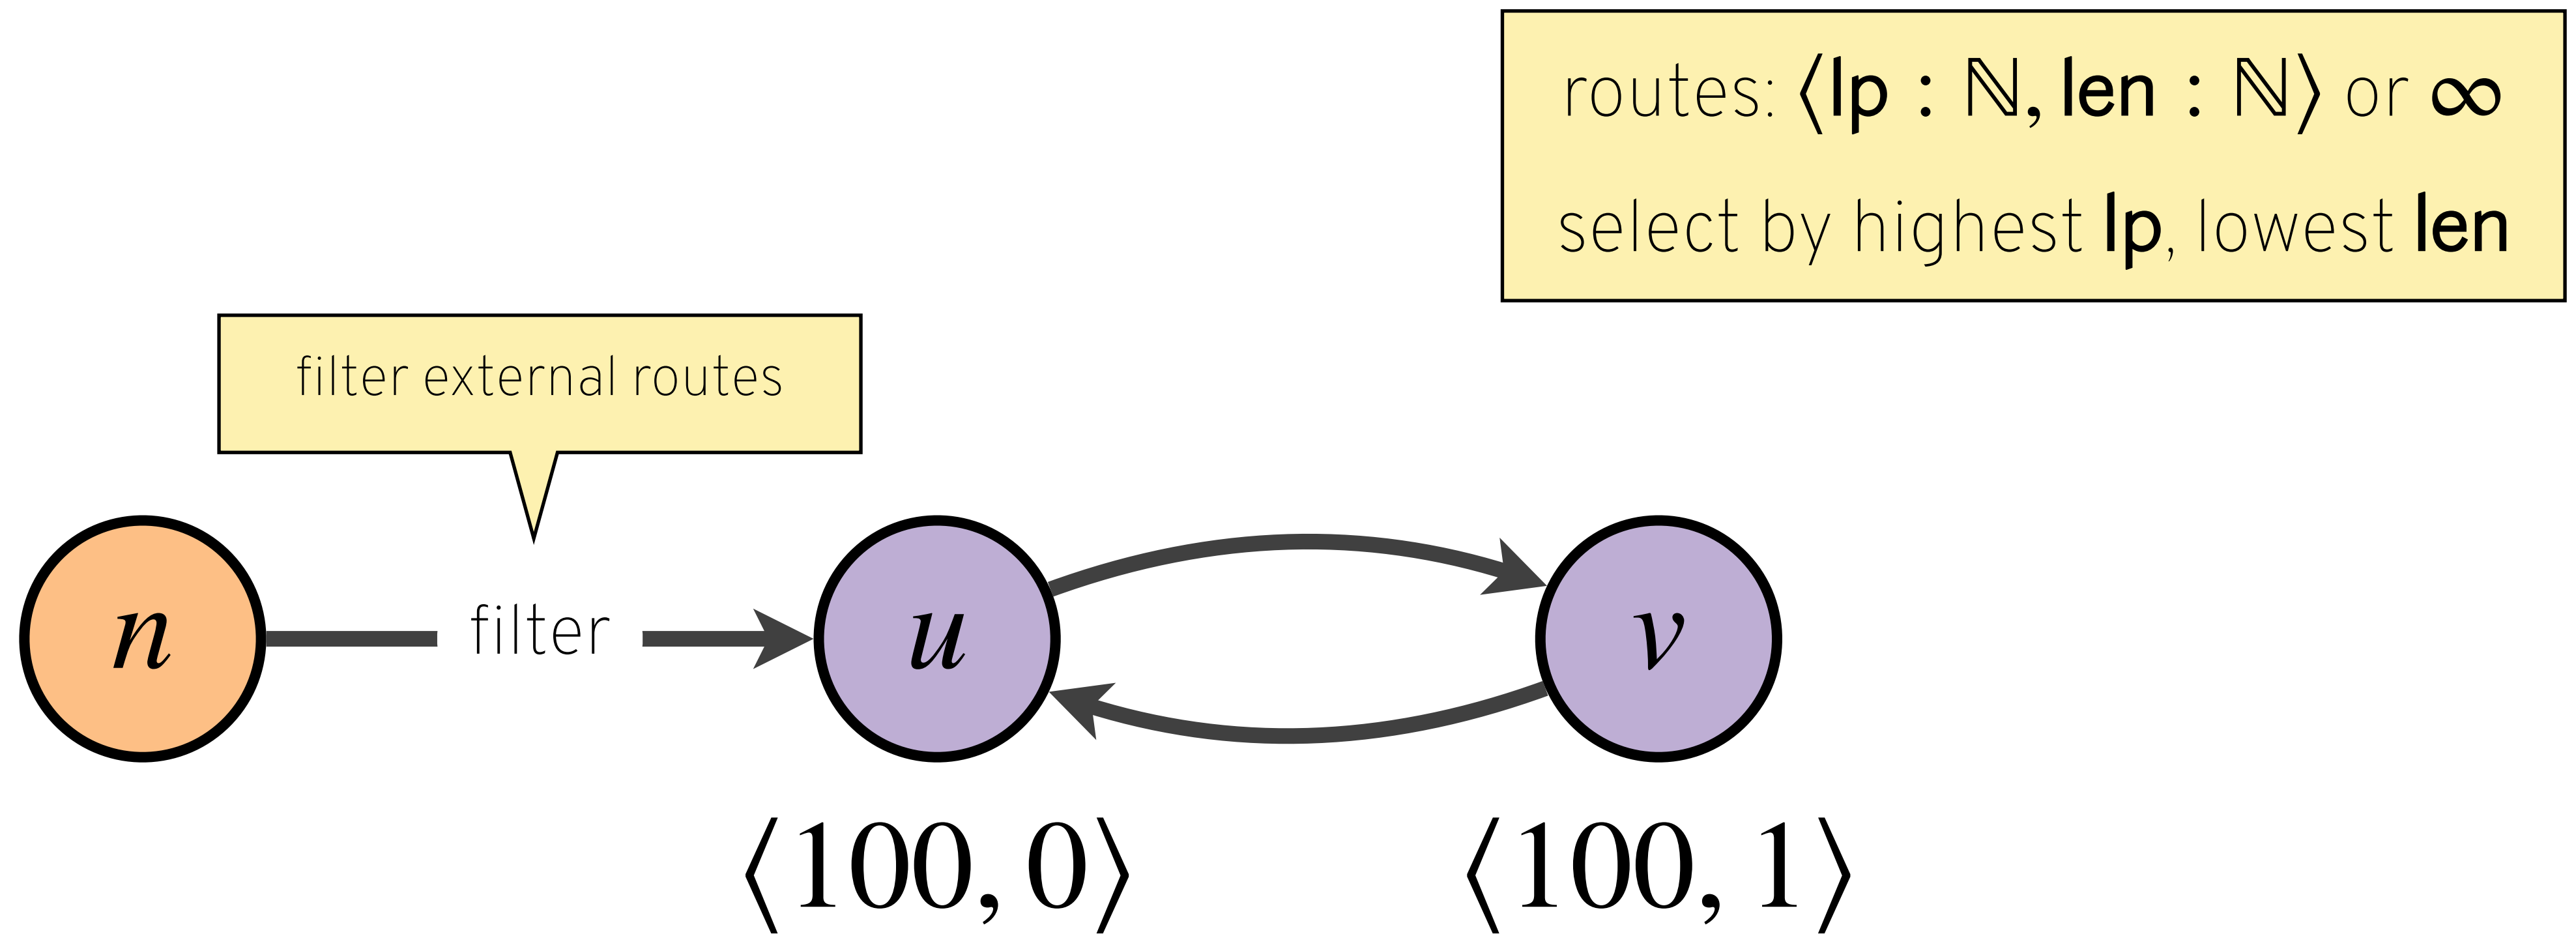





# An Example Network



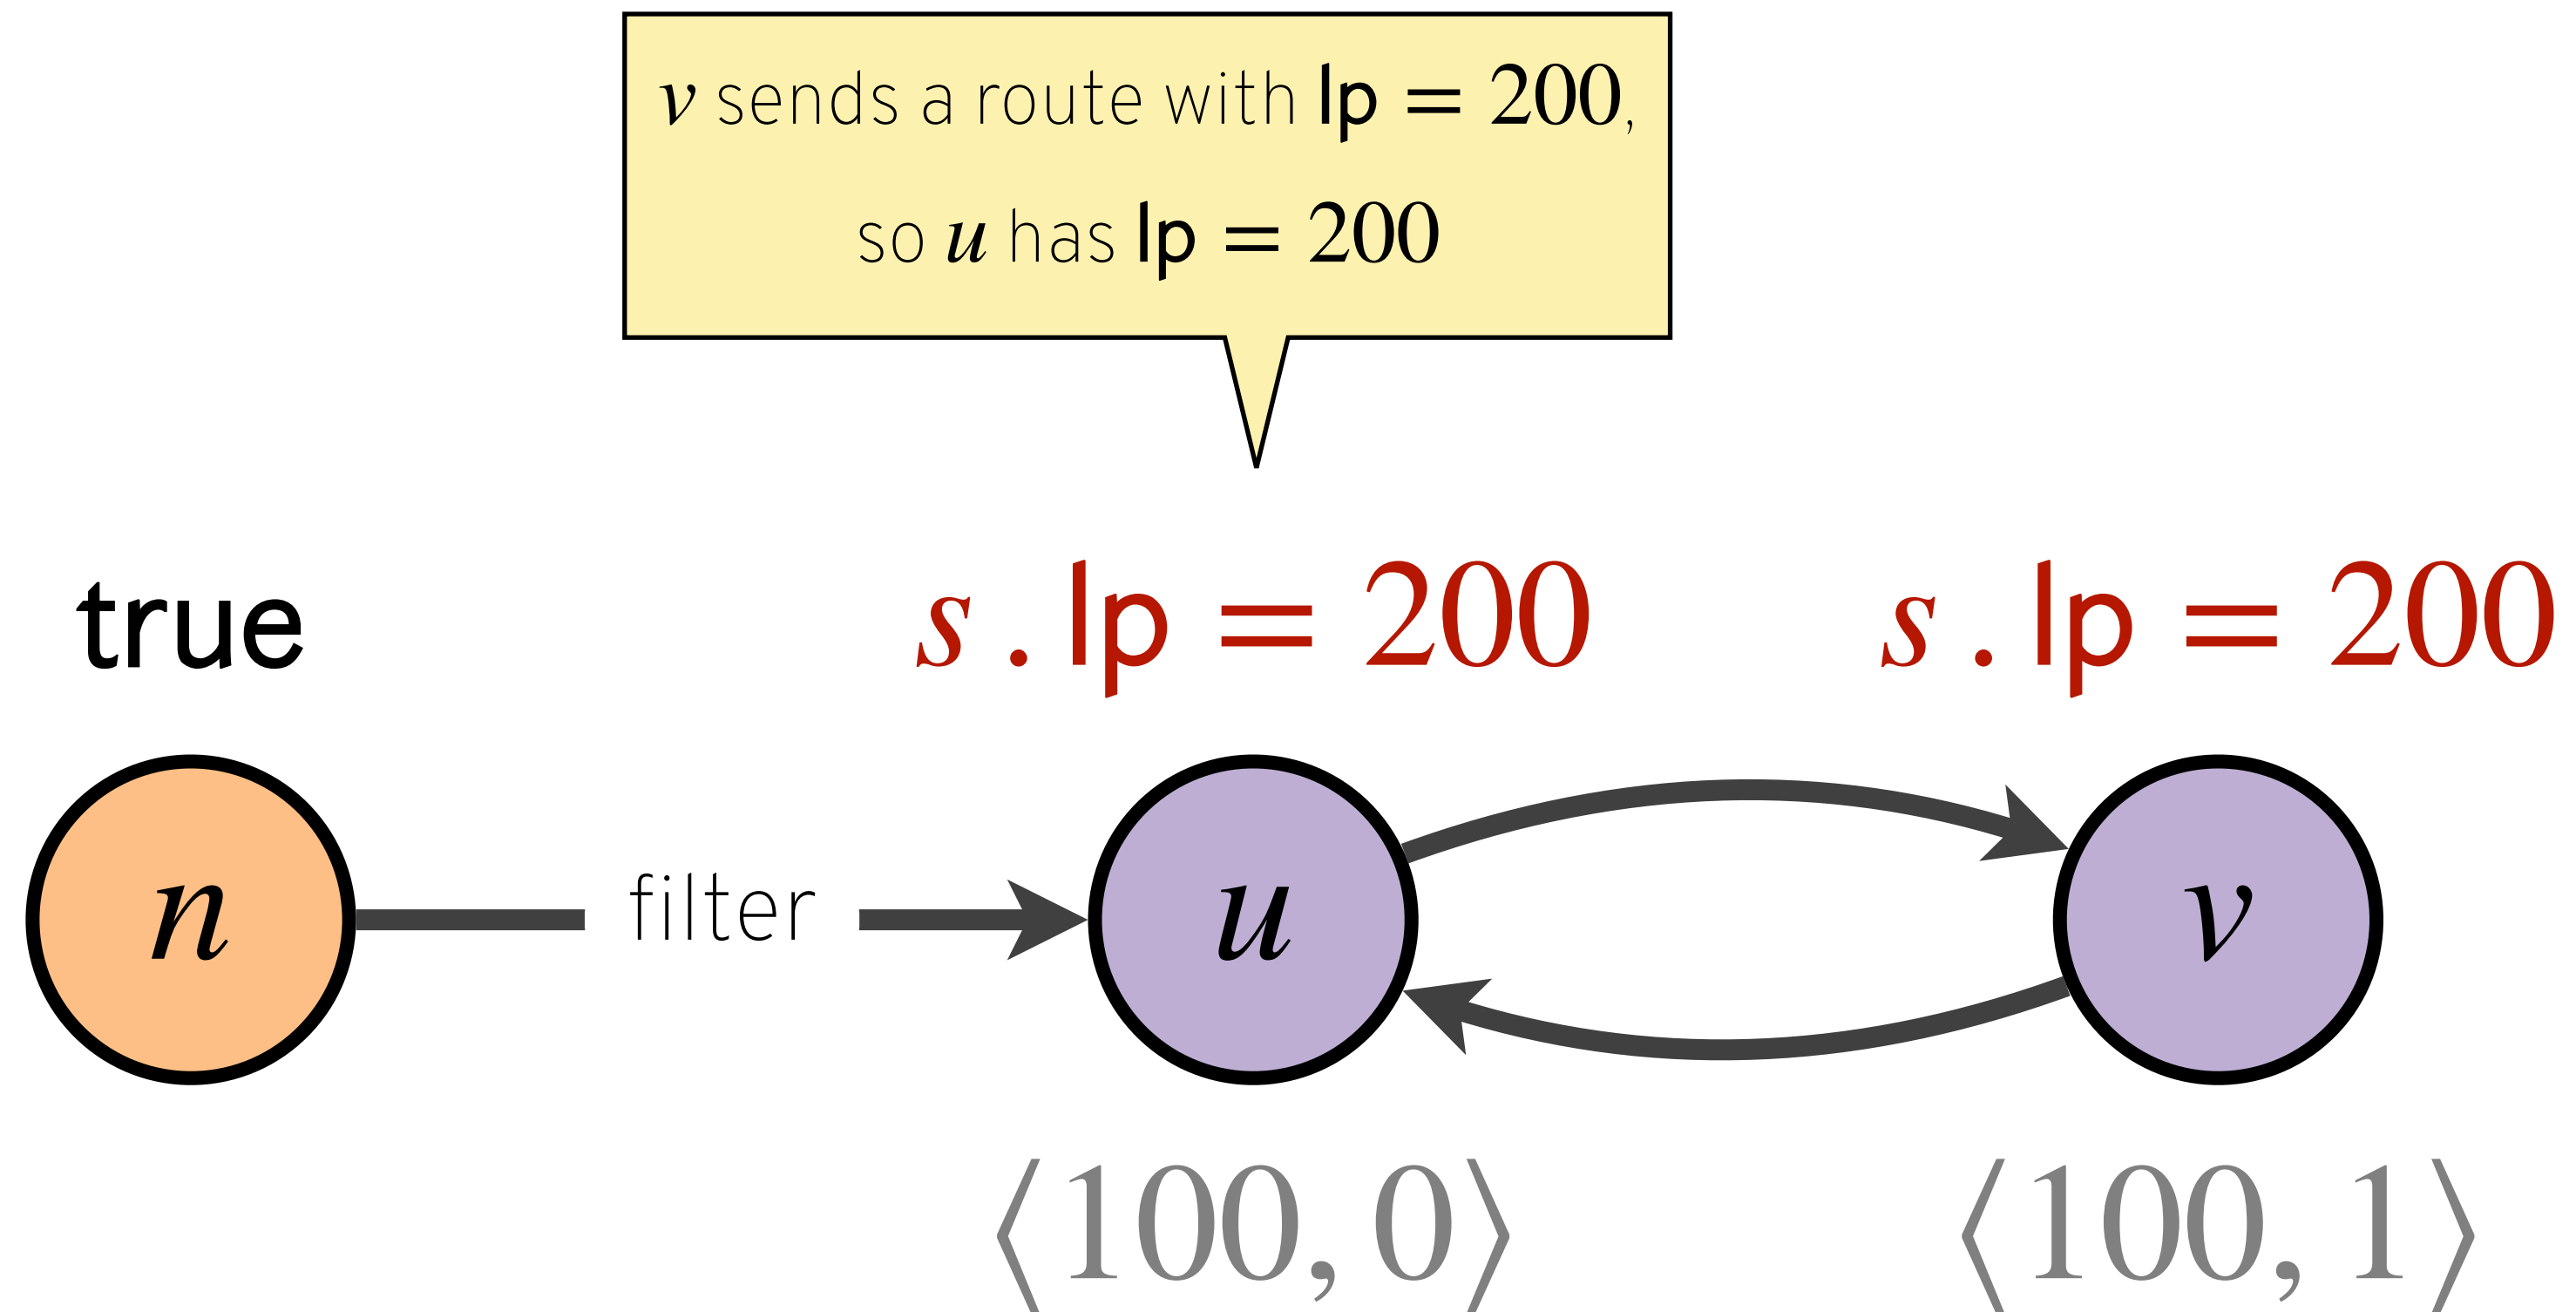
# An Example Network



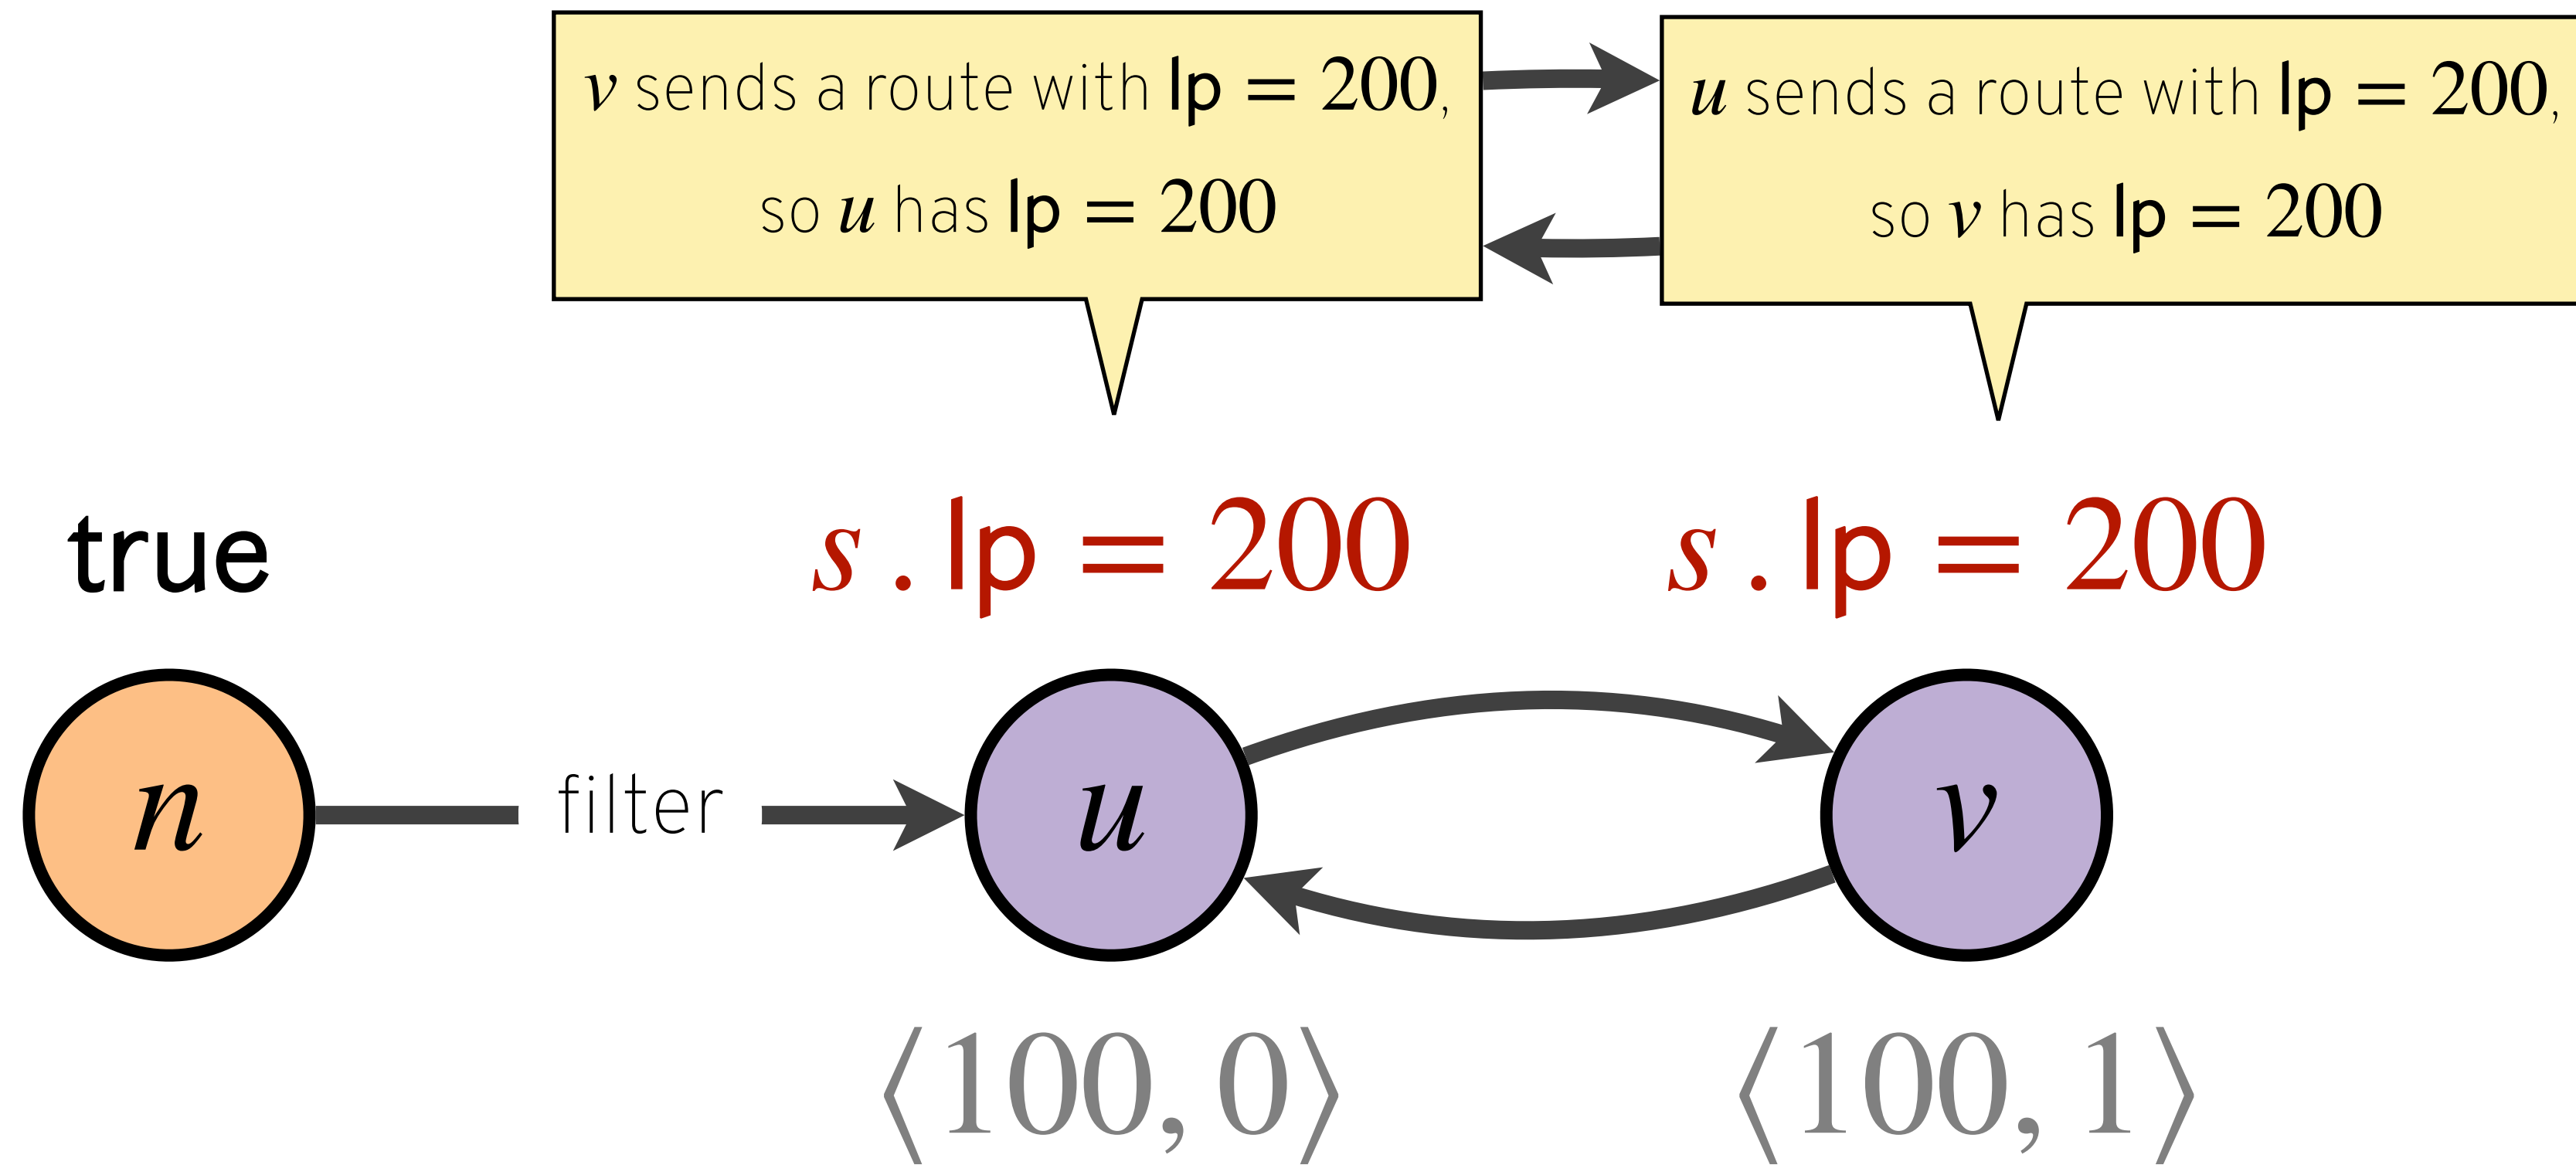
# Execution Interference



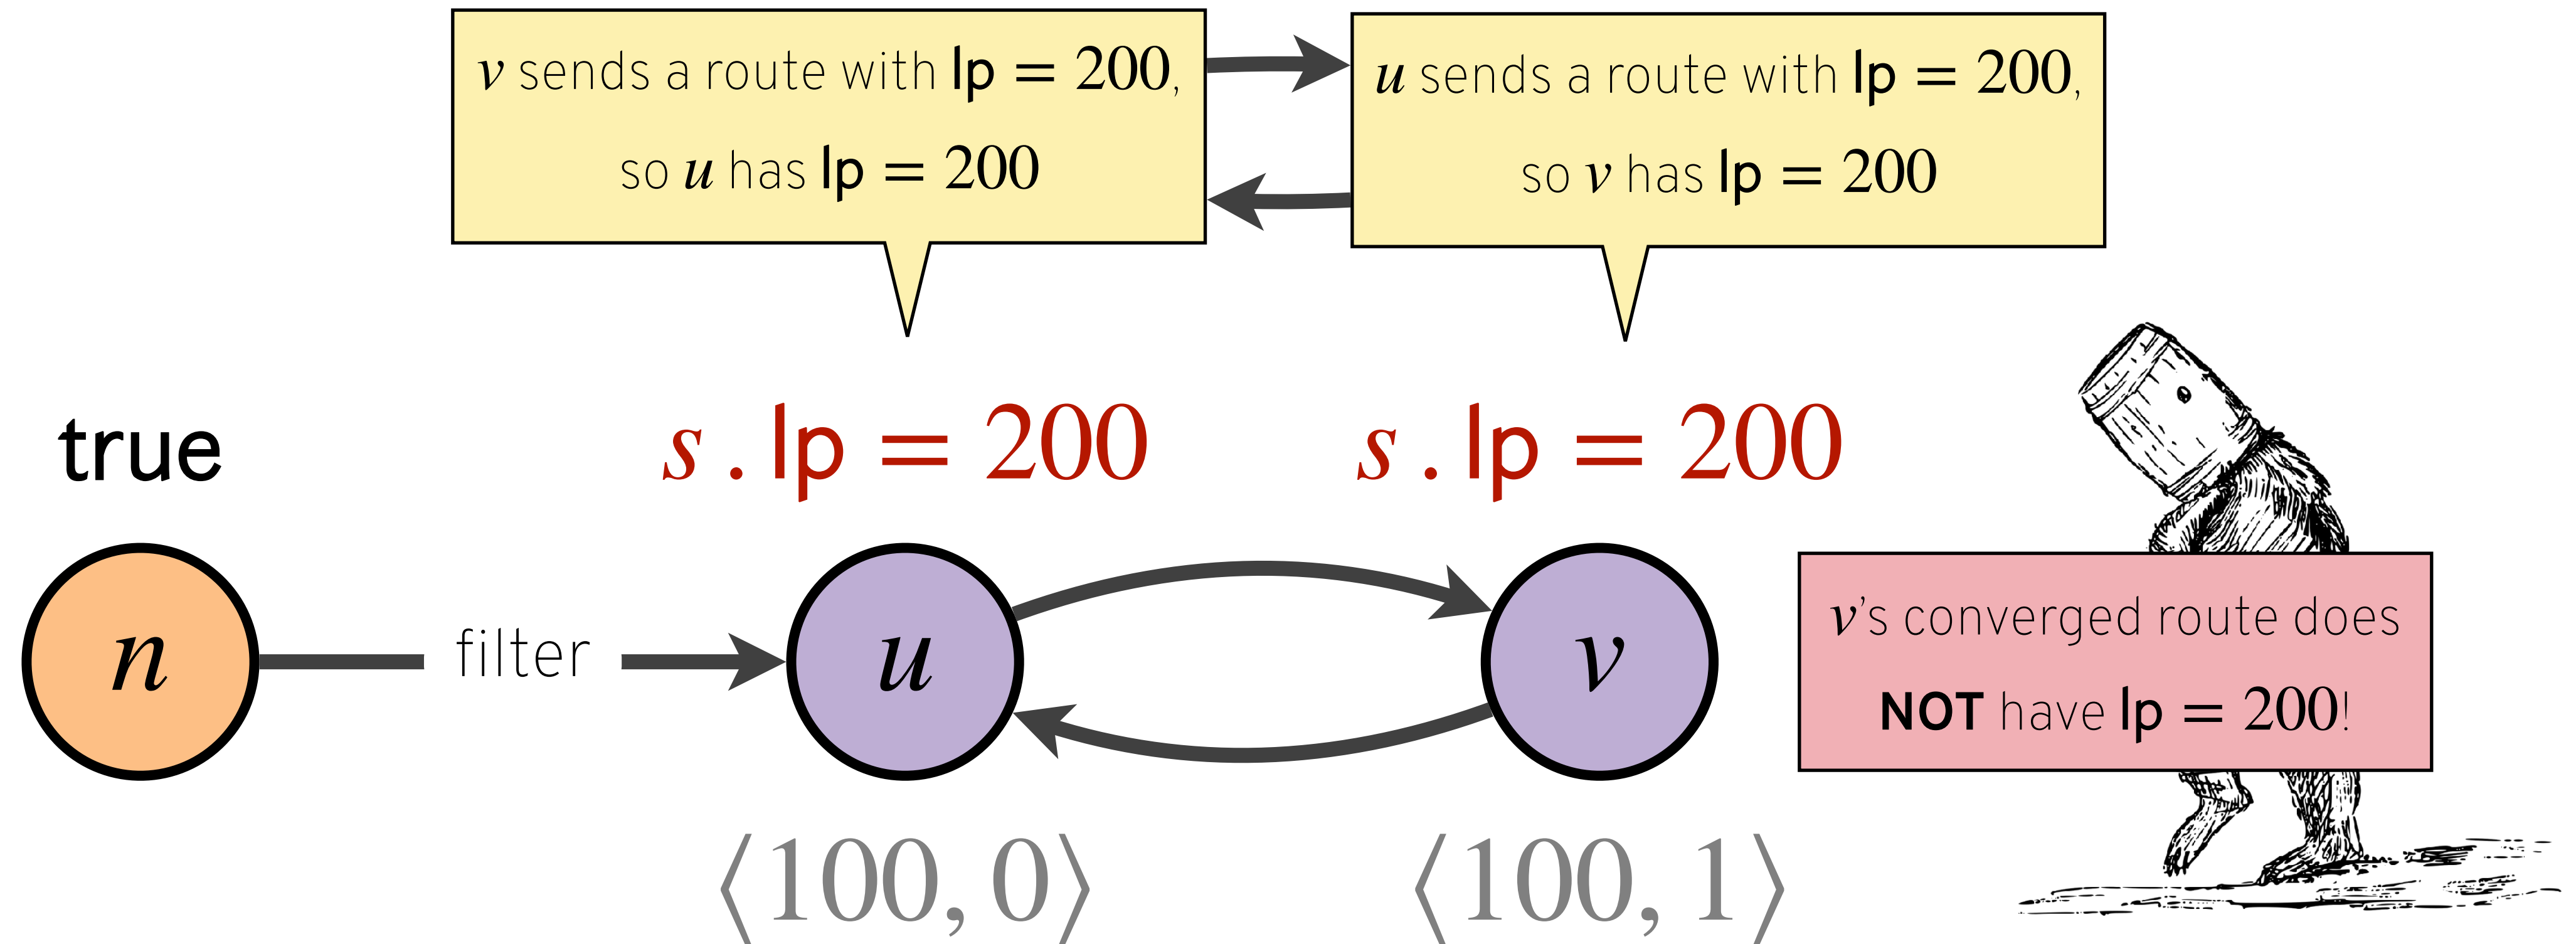
# Execution Interference



# Execution Interference

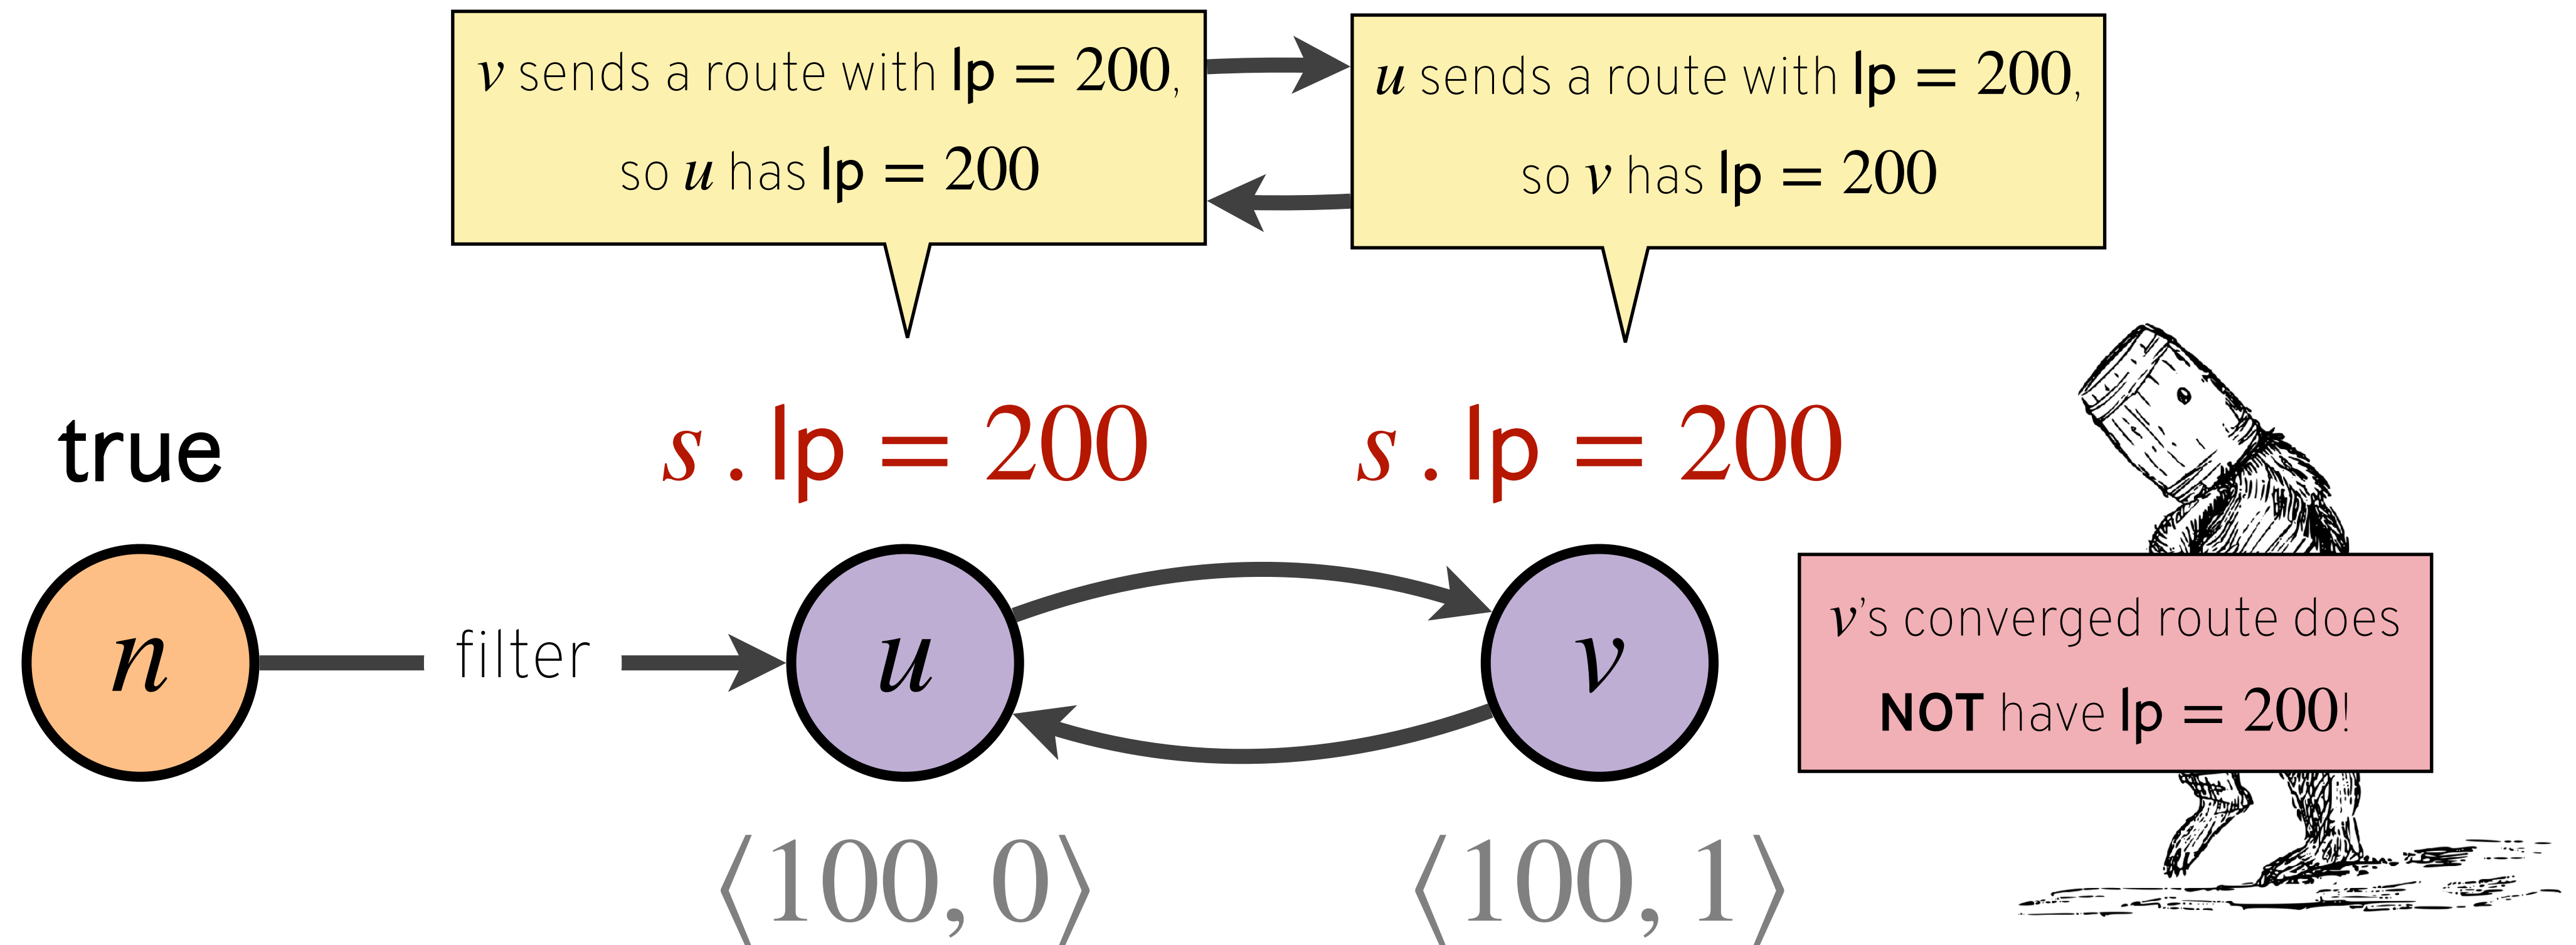


# Execution Interference






# Execution Interference




interfaces are **unsound**: exclude the legitimate converged routes, but the checks pass!

# Timepiece's Temporal Model




# Temporal Interfaces



interface  $A$  over-approximates the converged states of the node  $v$  with a set of states  $A(v)$


# Temporal Interfaces



**temporal** interface  $A$  over-approximates the states of the node  $v$  **at time**  $t$  with a set of states  $A(v)(t)$

# Temporal Interfaces

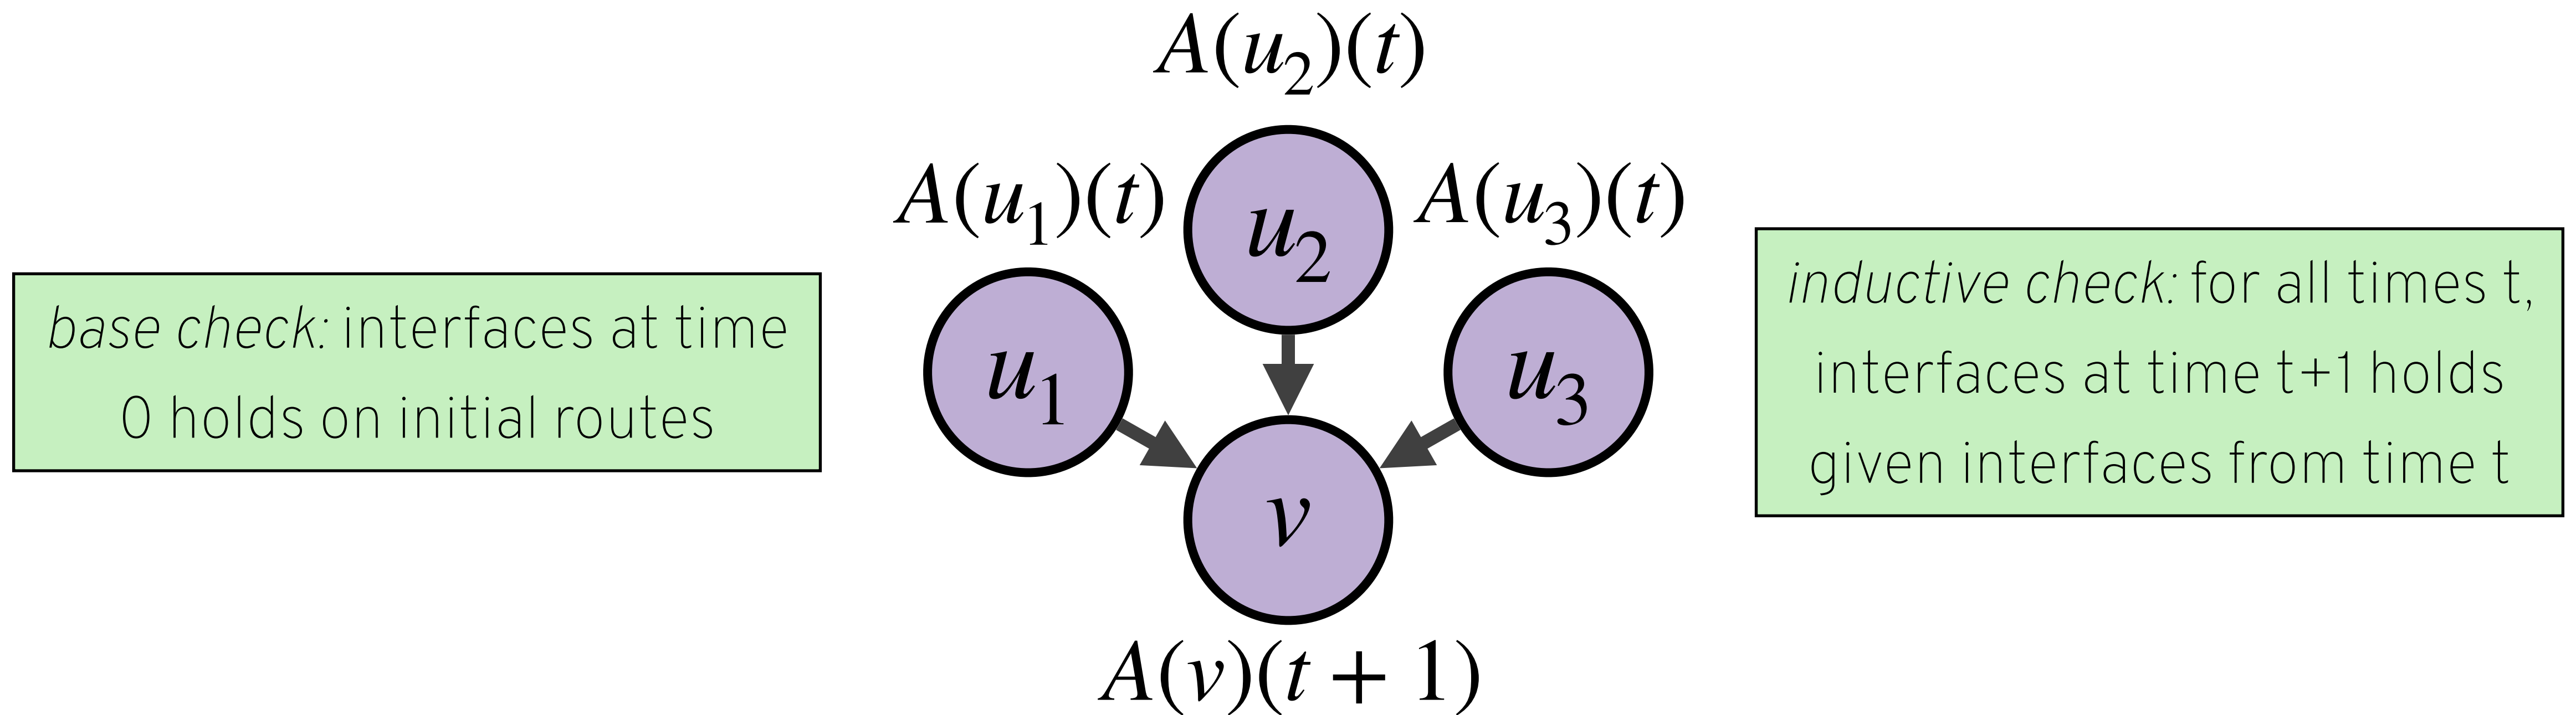
*base check: interfaces at time 0 holds on initial routes*



**temporal** interface  $A$  over-approximates the states of the node  $v$  **at time**  $t$  with a set of states  $A(v)(t)$




# Temporal Interfaces




**temporal** interface  $A$  over-approximates the states of the node  $v$  **at time**  $t$  with a set of states  $A(v)(t)$

# Preventing Interference




# Preventing Interference

$\mathcal{G}P$  (**globally**  $P$ ): at every point in time, the predicate  $P$  holds



# Preventing Interference

$\mathcal{G}P$  (**globally**  $P$ ): at every point in time, the predicate  $P$  holds




Base checks fail: interfaces  $A(u)$  and  $A(v)$  do not hold for initial routes at time **0**




# Proving Path Length

$\mathcal{G}(\text{true})$




$\mathcal{G}(s . \text{lp} = 100 \wedge s . \text{len} = 0)$

# Proving Path Length



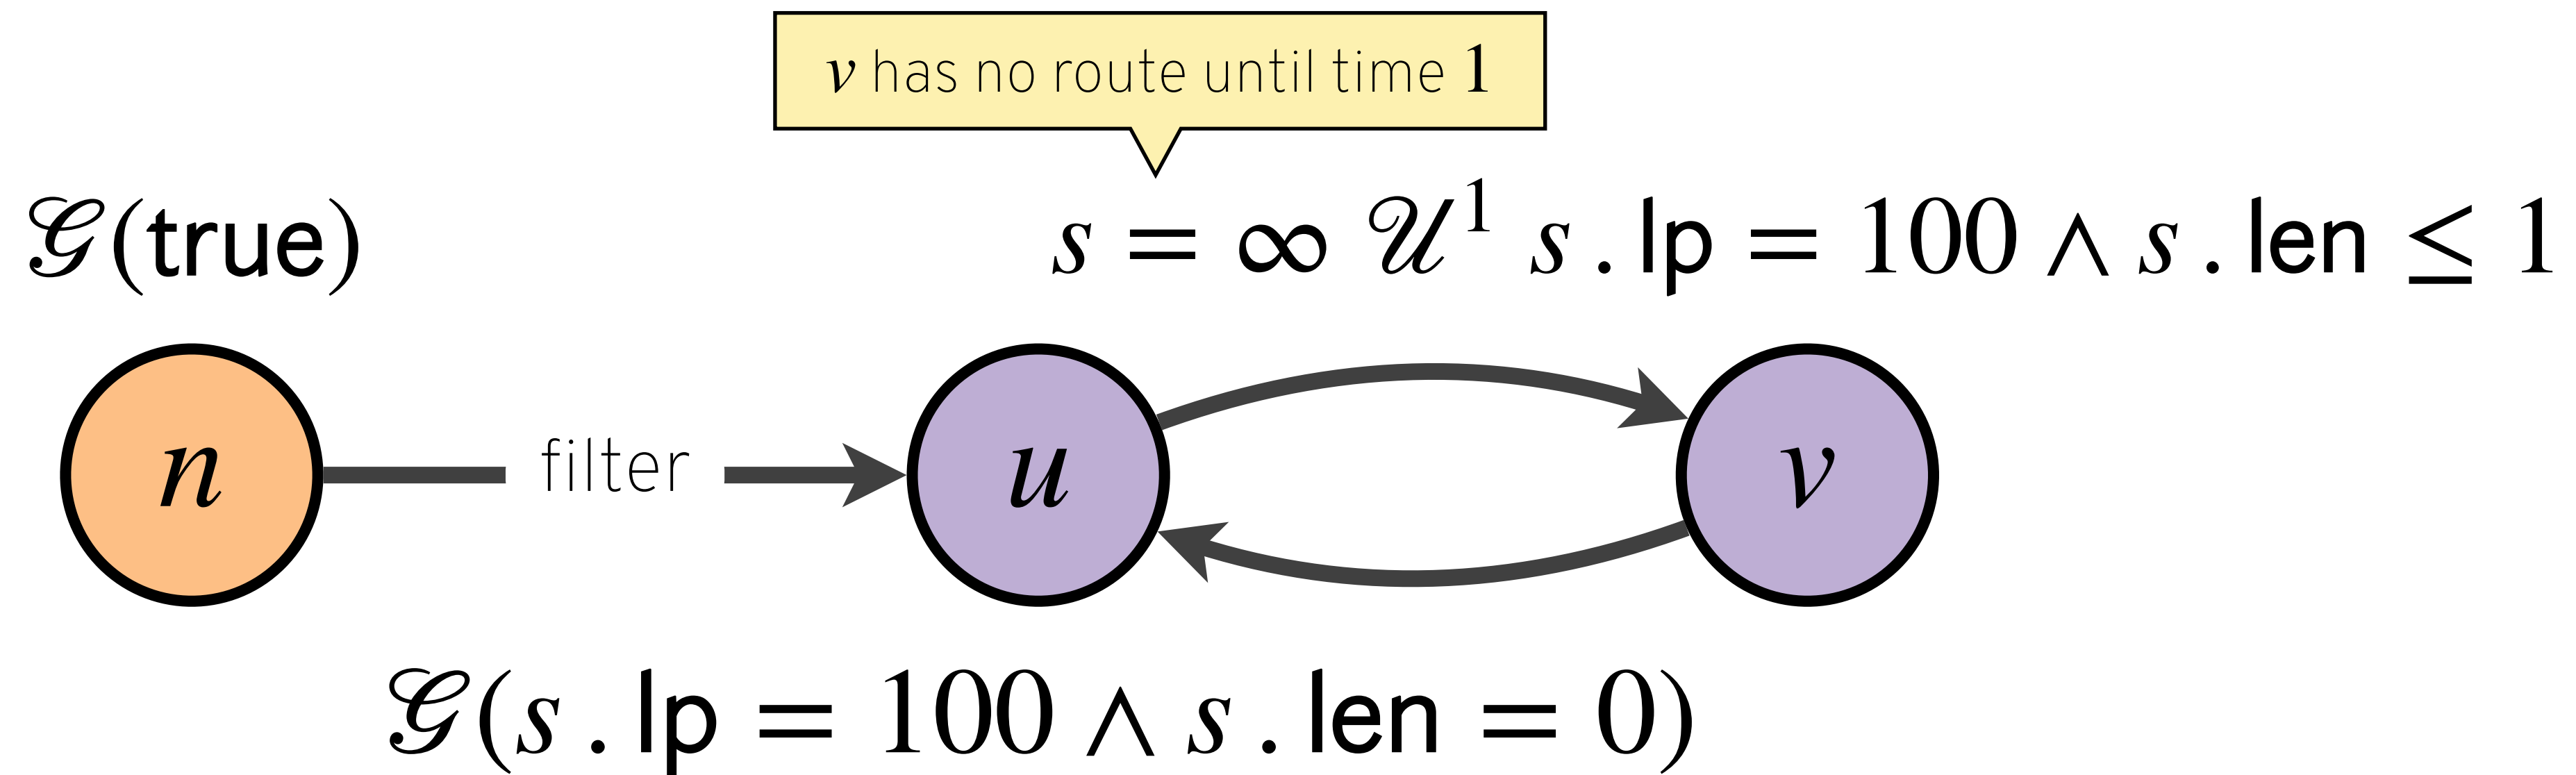
# Proving Path Length

$P \mathcal{U}^t Q$  ( $P$  **until**  $Q$  **at**  $t$ ): until time  $t$ ,  $P$  holds; at and after  $t$ ,  $Q$  holds



# Proving Path Length

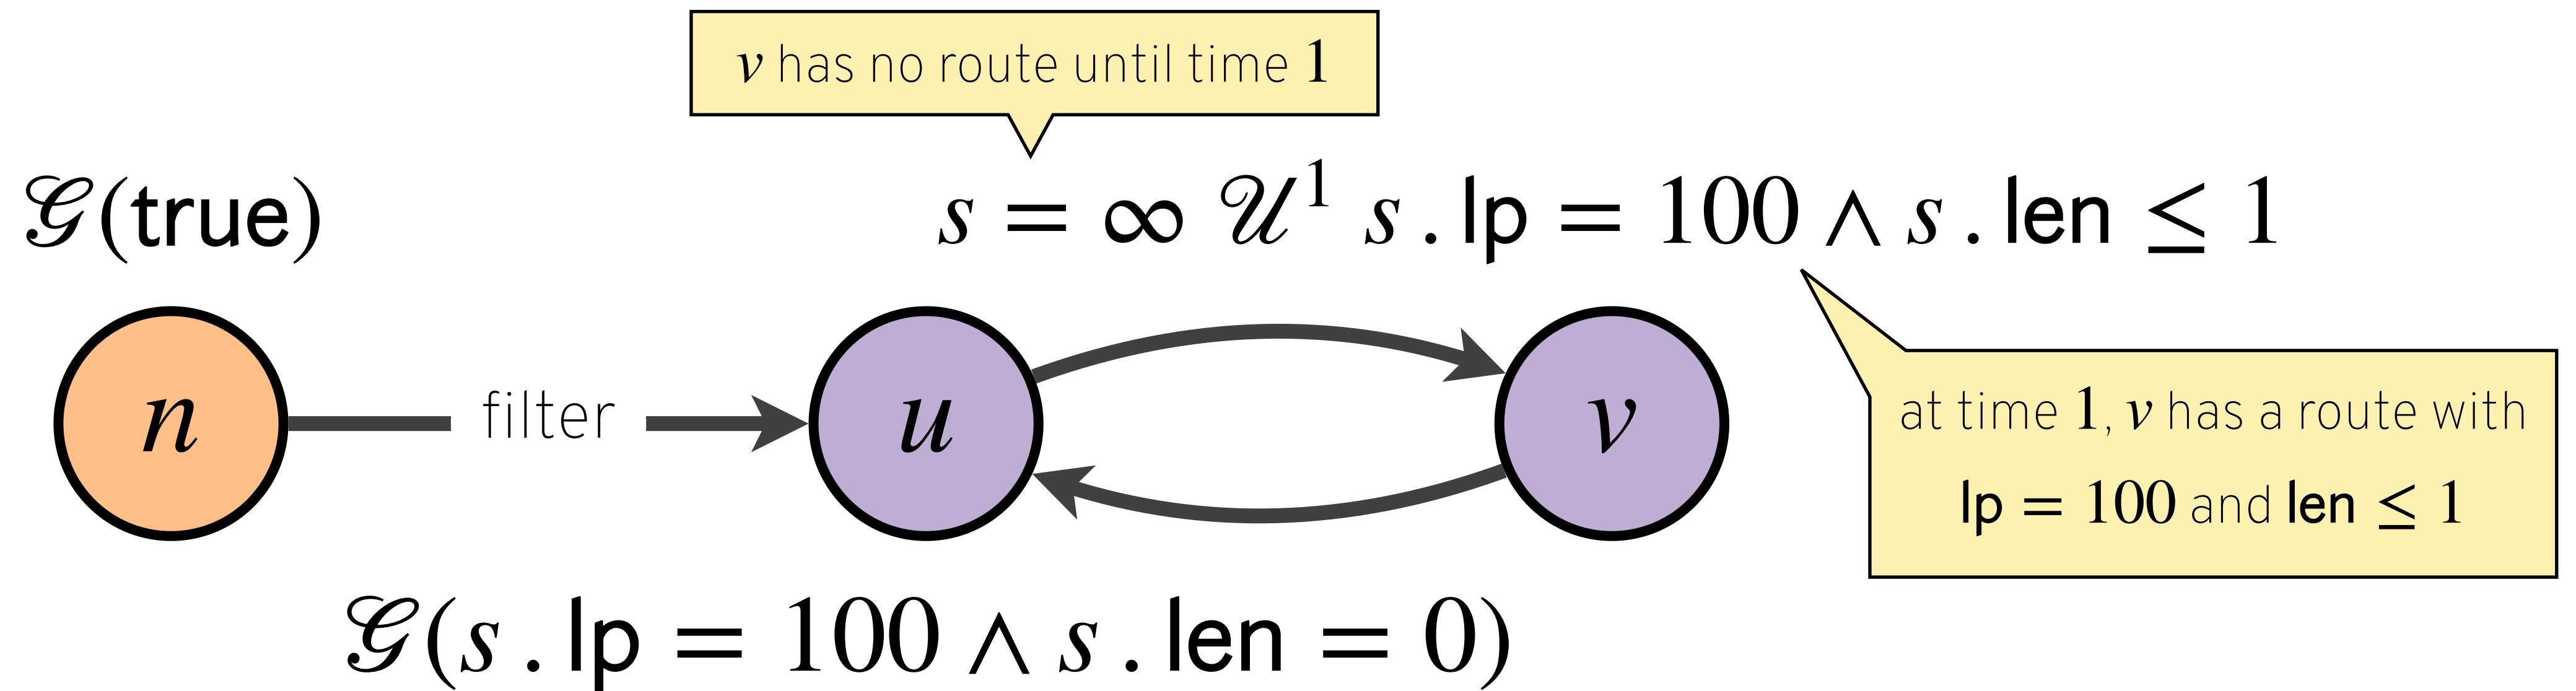
$P \mathcal{U}^t Q$  ( $P$  until  $Q$  at  $t$ ): until time  $t$ ,  $P$  holds; at and after  $t$ ,  $Q$  holds





# Proving Path Length

$P \mathcal{U}^t Q$  ( $P$  until  $Q$  at  $t$ ): until time  $t$ ,  $P$  holds; at and after  $t$ ,  $Q$  holds




# Soundness Theorem



# Soundness Theorem

If interface  $A$  satisfies the base and inductive checks for all nodes, then  $A$  includes all states computable via (monolithic) simulation.






# Soundness Theorem

If interface  $A$  satisfies the base and inductive checks for all nodes, then  $A$  includes all states computable via (monolithic) simulation.


Proof by induction on time.






# How to Use Timepiece

# How to Use Timepiece



define network semantics in C# or  
via configurations (via Batfish)

# How to Use Timepiece




define network semantics in C# or  
via configurations (via Batfish)



write interfaces using C# library of  
temporal operators ( $\mathcal{G}, \mathcal{U}^t, \mathcal{F}^t$ )



# How to Use Timepiece



define network semantics in C# or  
via configurations (via Batfish)




write interfaces using C# library of  
temporal operators ( $\mathcal{G}, \mathcal{U}^t, \mathcal{F}^t$ )



check VCs *in parallel* on every node  
using Satisfiability Modulo  
Theories (SMT) solver





# Evaluation



# Evaluation


# Evaluation

does Timepiece scale to large networks?

does Timepiece handle complex policies?

how easy is it to write invariants for different properties?

# Evaluation




Fat-tree data center networks  
C# model of eBGP routing protocol  
20–2000 nodes

does Timepiece handle complex policies?

how easy is it to write invariants for different properties?

# Evaluation




Fat-tree data center networks  
C# model of eBGP routing protocol  
20–2000 nodes



Internet2 wide-area network  
102,753 lines of Juniper configuration code  
263 nodes (10 internal, 253 external)

how easy is it to write invariants for different properties?

# Evaluation



Fat-tree data center networks  
C# model of eBGP routing protocol  
20–2000 nodes

Reachability

Path length

Valley freedom

Hijack filtering



Internet2 wide-area network  
102,753 lines of Juniper configuration code  
263 nodes (10 internal, 253 external)

No transit



# Evaluation

on Microsoft Azure D96s VM with 96 vCPUs and 384GB RAM

# Evaluation

on Microsoft Azure D96s VM with 96 vCPUs and 384GB RAM

## Benchmark

Reachability

---

Path length

---

Valley freedom

---

Hijack filtering

---

No transit

---

# Evaluation

on Microsoft Azure D96s VM with 96 vCPUs and 384GB RAM

Benchmark	Network LoC
Reachability	81
Path length	88
Valley freedom	89
Hijack filtering	146
No transit	88 (+102,753)

# Evaluation

on Microsoft Azure D96s VM with 96 vCPUs and 384GB RAM

Benchmark	Network LoC	Annotation LoC
Reachability	81	3
Path length	88	7
Valley freedom	89	12
Hijack filtering	146	21
No transit	88 (+102,753)	5

# Evaluation

on Microsoft Azure D96s VM with 96 vCPUs and 384GB RAM

Benchmark	Network LoC	Annotation LoC	Nodes
Reachability	81	3	2000
Path length	88	7	2000
Valley freedom	89	12	2000
Hijack filtering	146	21	2000
No transit	88 (+102,753)	5	263



# Evaluation

on Microsoft Azure D96s VM with 96 vCPUs and 384GB RAM

Benchmark	Network LoC	Annotation LoC	Nodes	Monolithic time
Reachability	81	3	2000	14s
Path length	88	7	2000	>2h
Valley freedom	89	12	2000	>2h
Hijack filtering	146	21	2000	>2h
No transit	88 (+102,753)	5	263	>2h

# Evaluation

on Microsoft Azure D96s VM with 96 vCPUs and 384GB RAM

Benchmark	Network LoC	Annotation LoC	Nodes	Monolithic time	Modular time
Reachability	81	3	2000	14s	28s
Path length	88	7	2000	>2h	1204s
Valley freedom	89	12	2000	>2h	398s
Hijack filtering	146	21	2000	>2h	142s
No transit	88 (+102,753)	5	263	>2h	38s

# Evaluation

on Microsoft Azure D96s VM with 96 vCPUs and 384GB RAM


Benchmark	Network LoC	Annotation LoC	Nodes	Monolithic time	Modular time
Reachability	81	3	2000	14s	28s
Path length	88	7	2000	>2h	1204s
Valley freedom	89	12	2000	>2h	398s
Hijack filtering	146	21	2000	>2h	142s
No transit	88 (+102,753)	5	263	>2h	38s

# Fat-tree Hijack Filtering

BGP misconfiguration/attack:

a “hijacker” node announces it has a path to a prefix it doesn’t own, misleading others to route through the hijacker

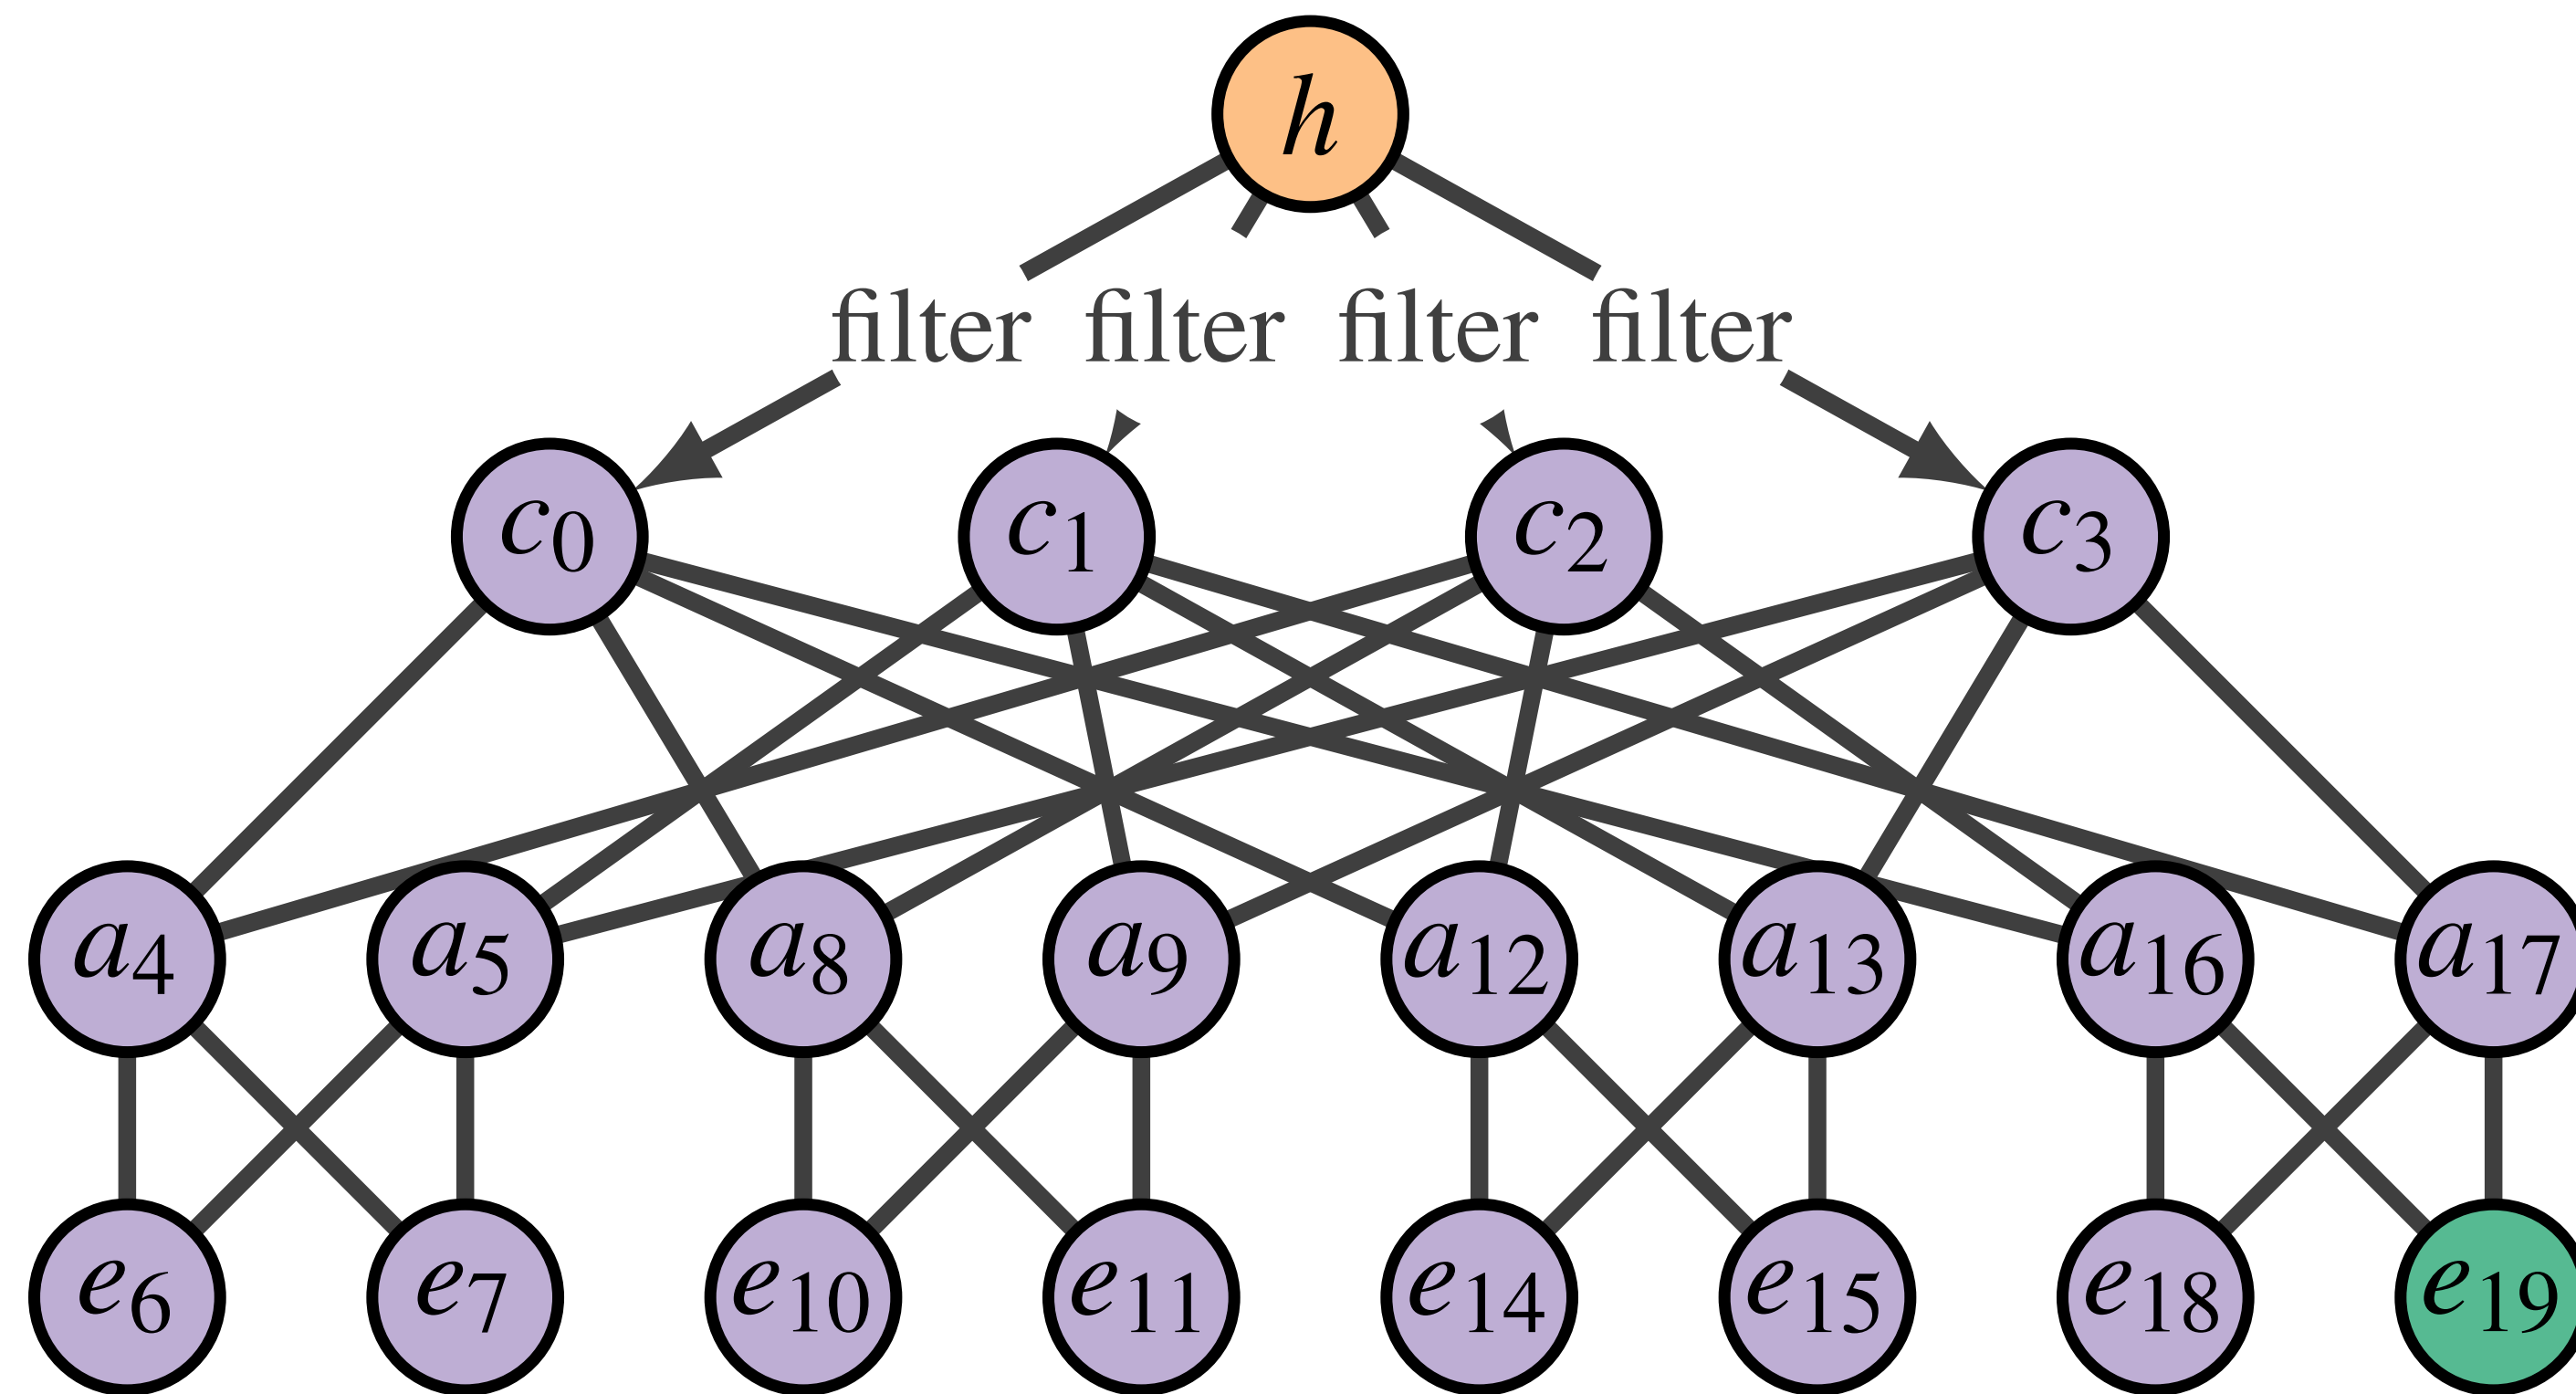
# Fat-tree Hijack Filtering



BGP misconfiguration/attack:  
a “hijacker” node announces it has a path to a prefix it doesn’t own, misleading others to route through the hijacker



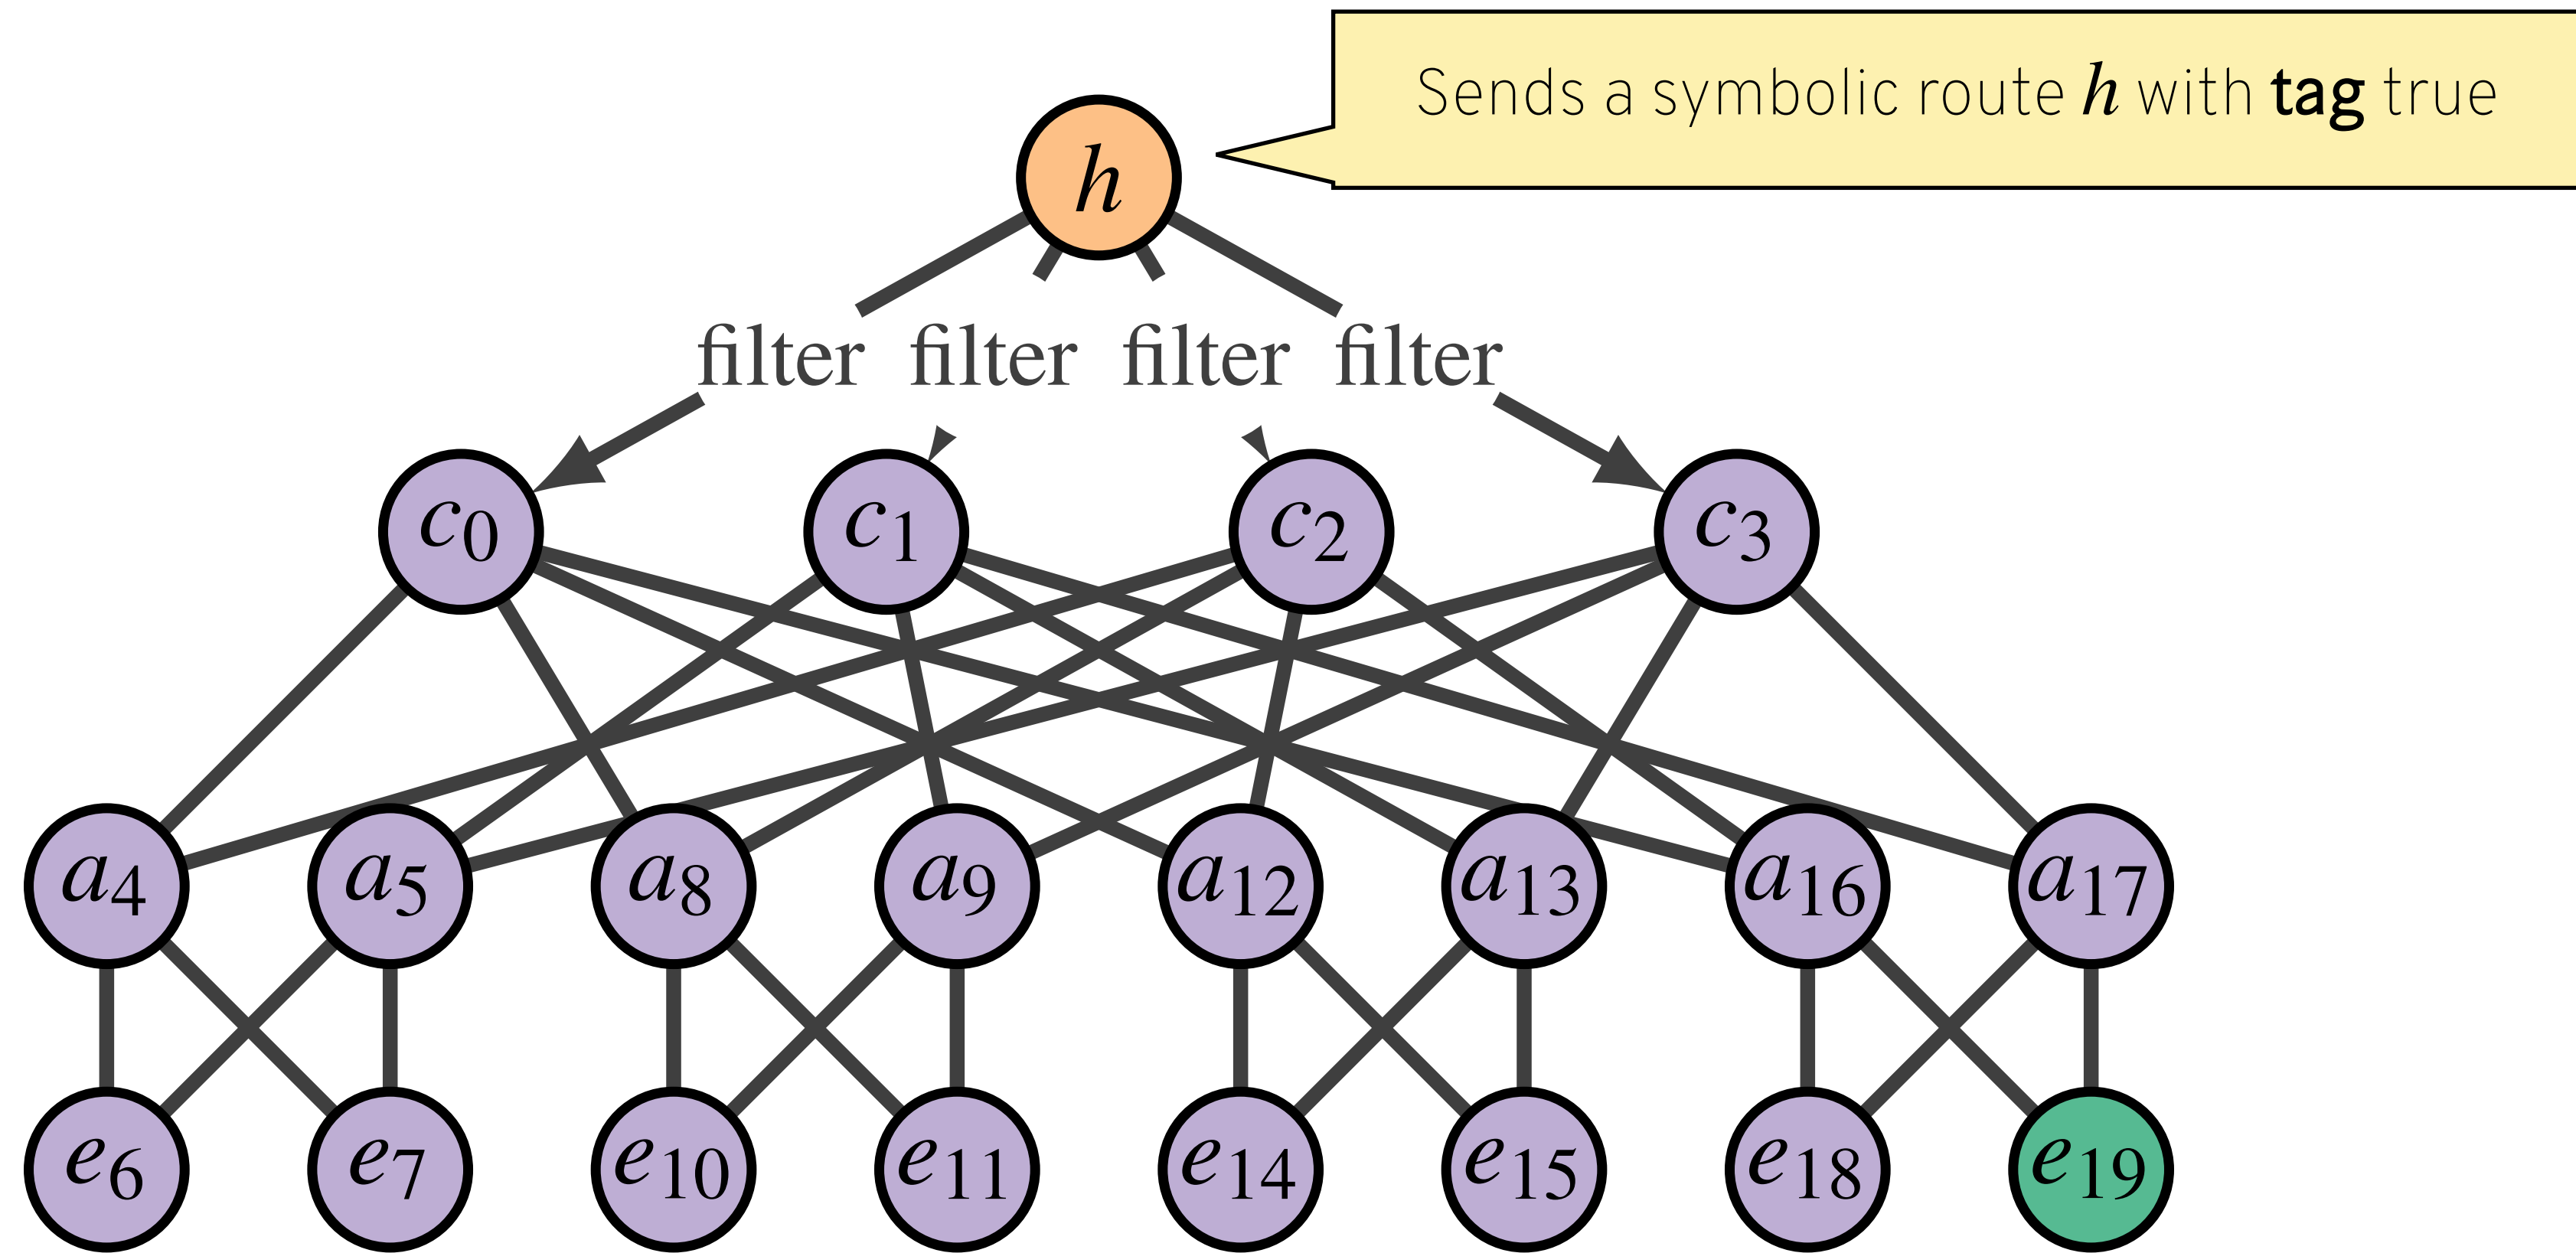
# Fat-tree Hijack Filtering



BGP misconfiguration/attack:  
a “hijacker” node announces it has a path to a prefix it doesn’t own, misleading others to route through the hijacker

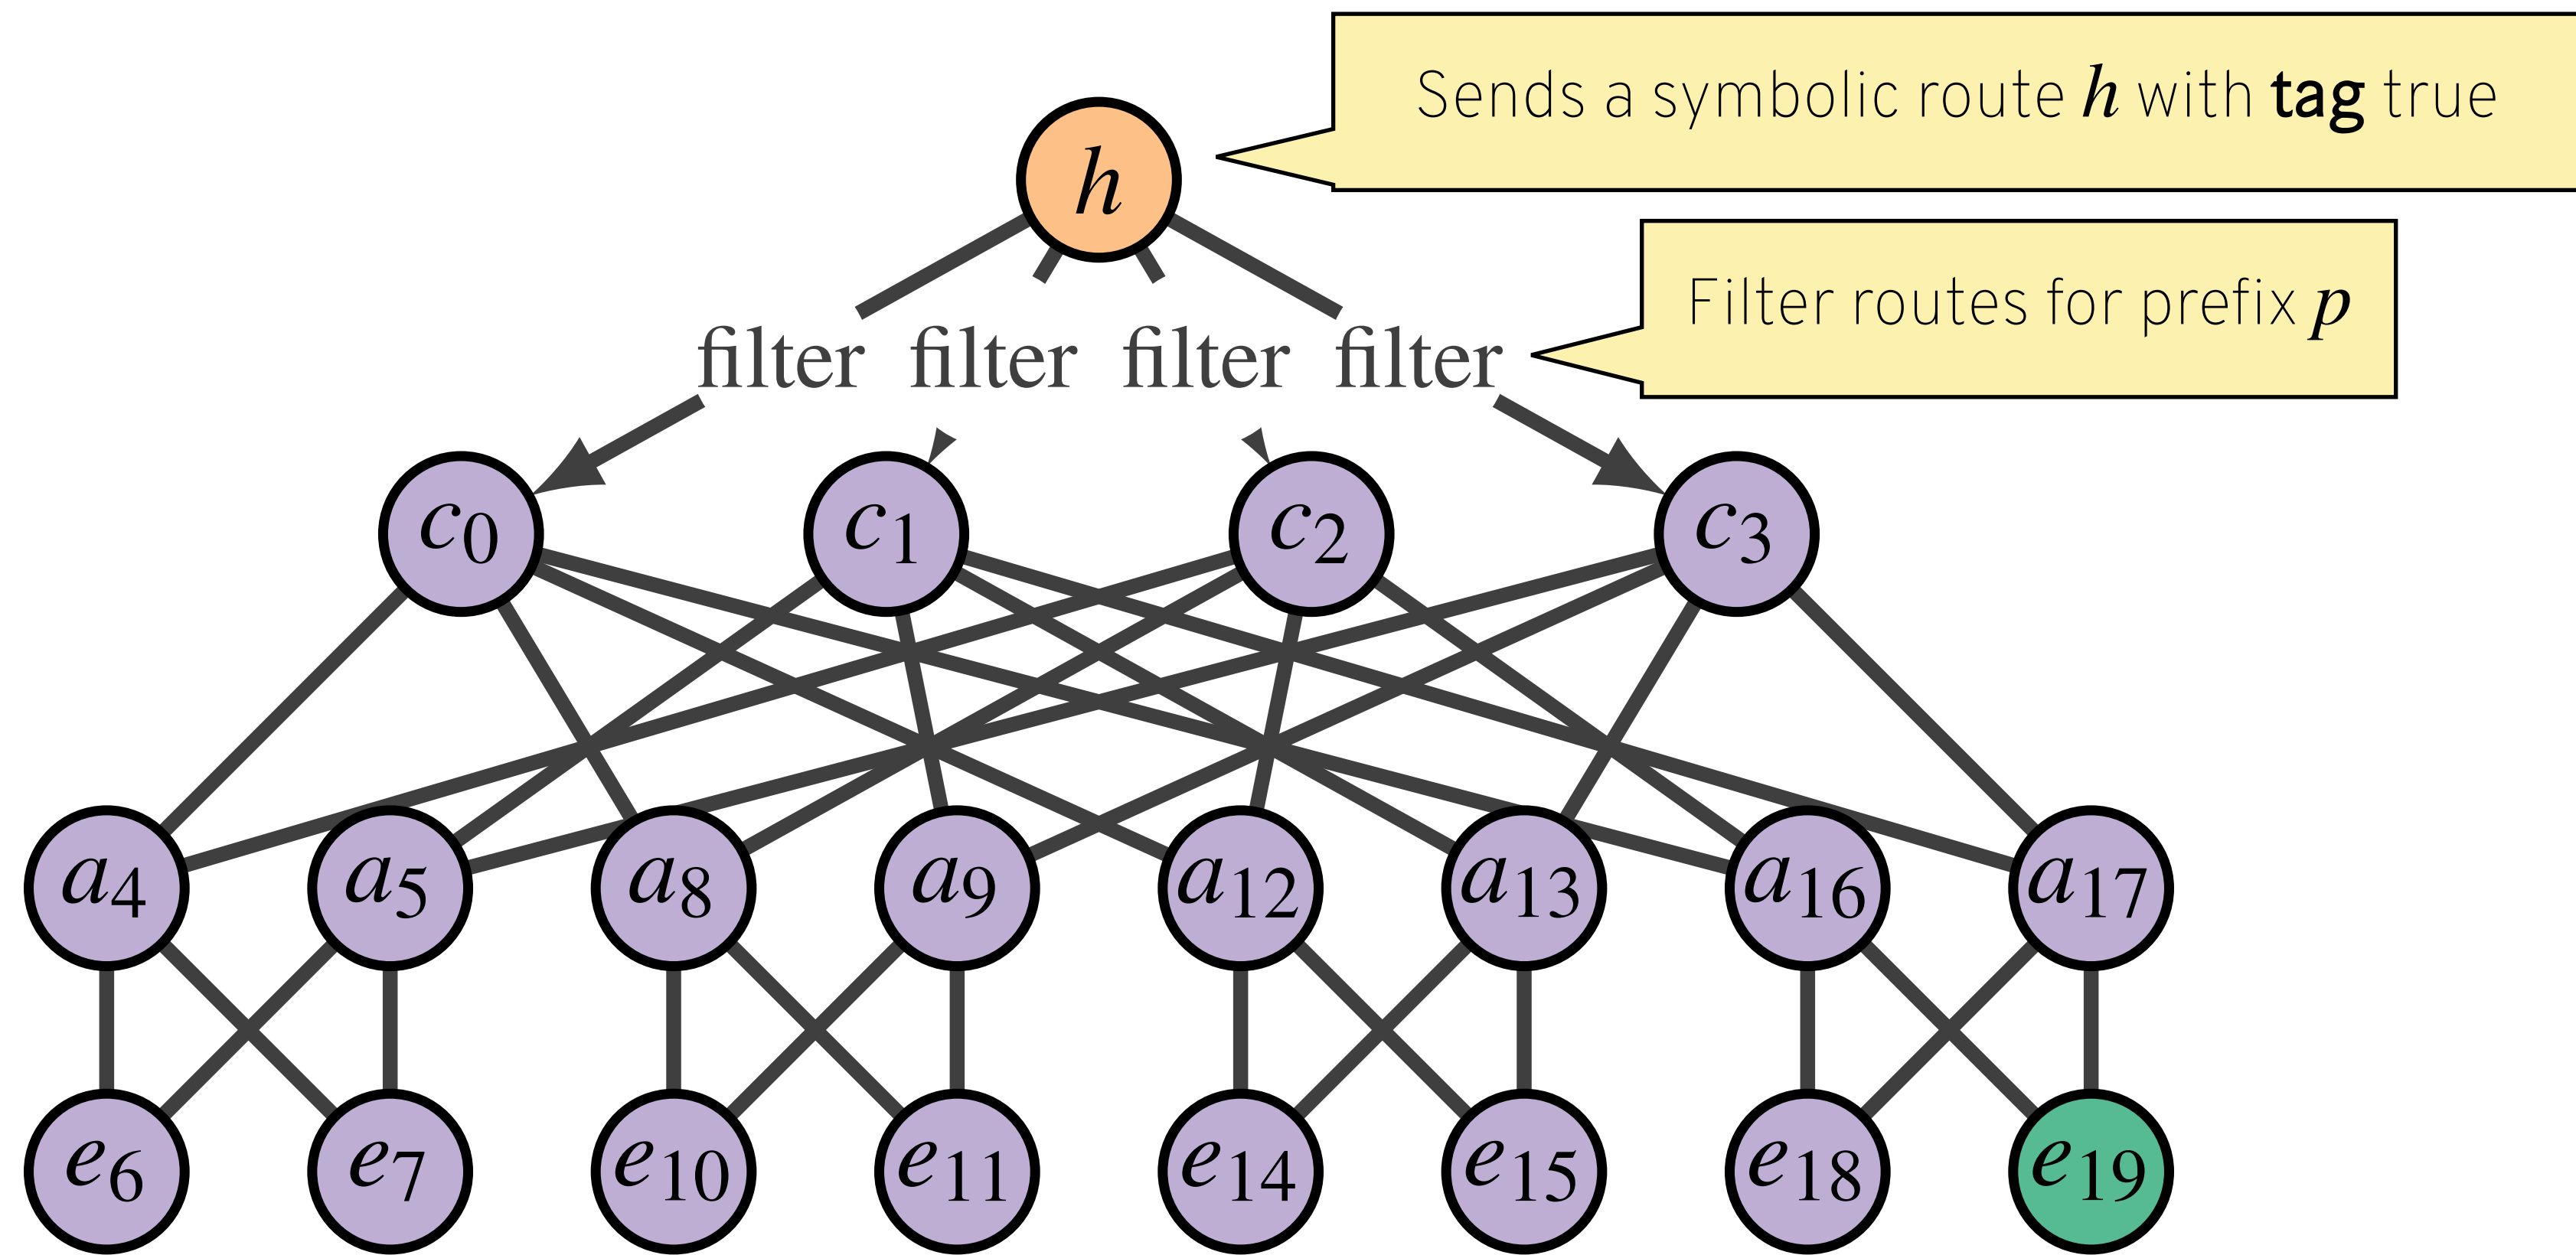
Sends a route with symbolic prefix  $p$

# Fat-tree Hijack Filtering



BGP misconfiguration/attack:  
a “hijacker” node announces it has a path to a prefix it doesn’t own, misleading others to route through the hijacker

# Fat-tree Hijack Filtering



BGP misconfiguration/attack:  
a "hijacker" node announces it has a path to a prefix it doesn't own, misleading others to route through the hijacker

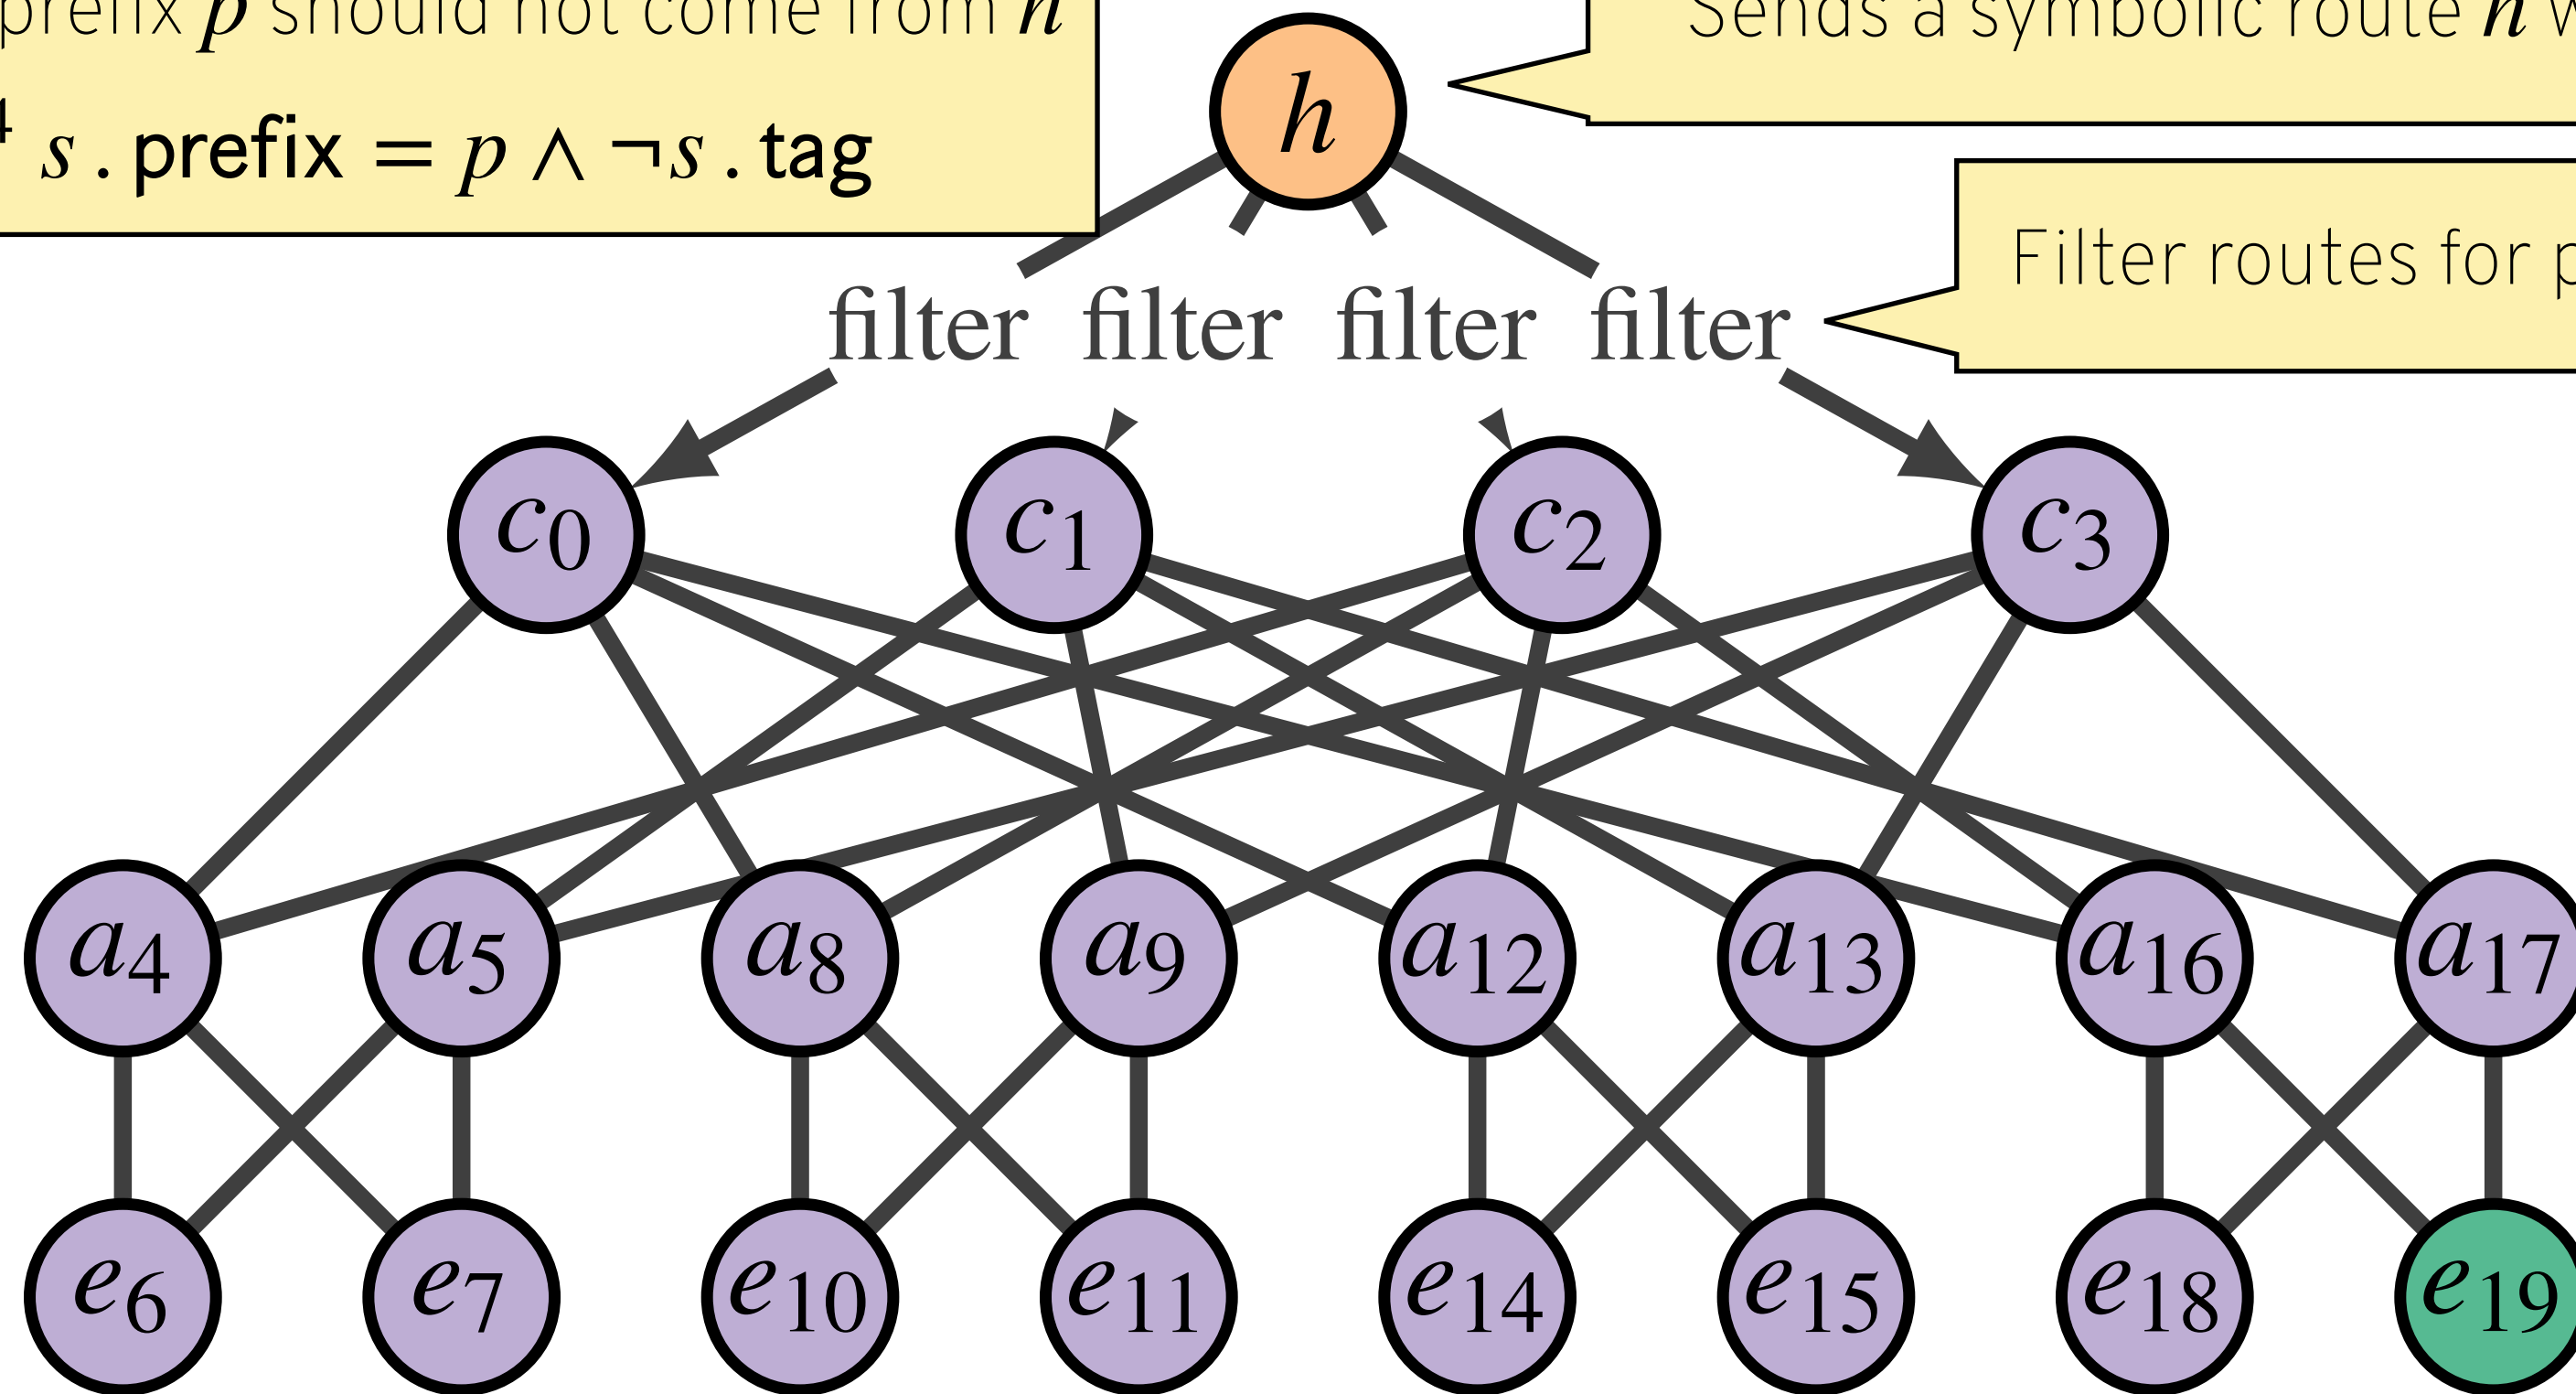
# Fat-tree Hijack Filtering

Converged routes for prefix  $p$  should not come from  $h$

$$P(v) \equiv \text{true } \mathcal{U}^4 s . \text{prefix} = p \wedge \neg s . \text{tag}$$

Sends a symbolic route  $h$  with **tag** true

Filter routes for prefix  $p$



BGP misconfiguration/attack:  
a “hijacker” node announces it has a path to a prefix it doesn’t own, misleading others to route through the hijacker

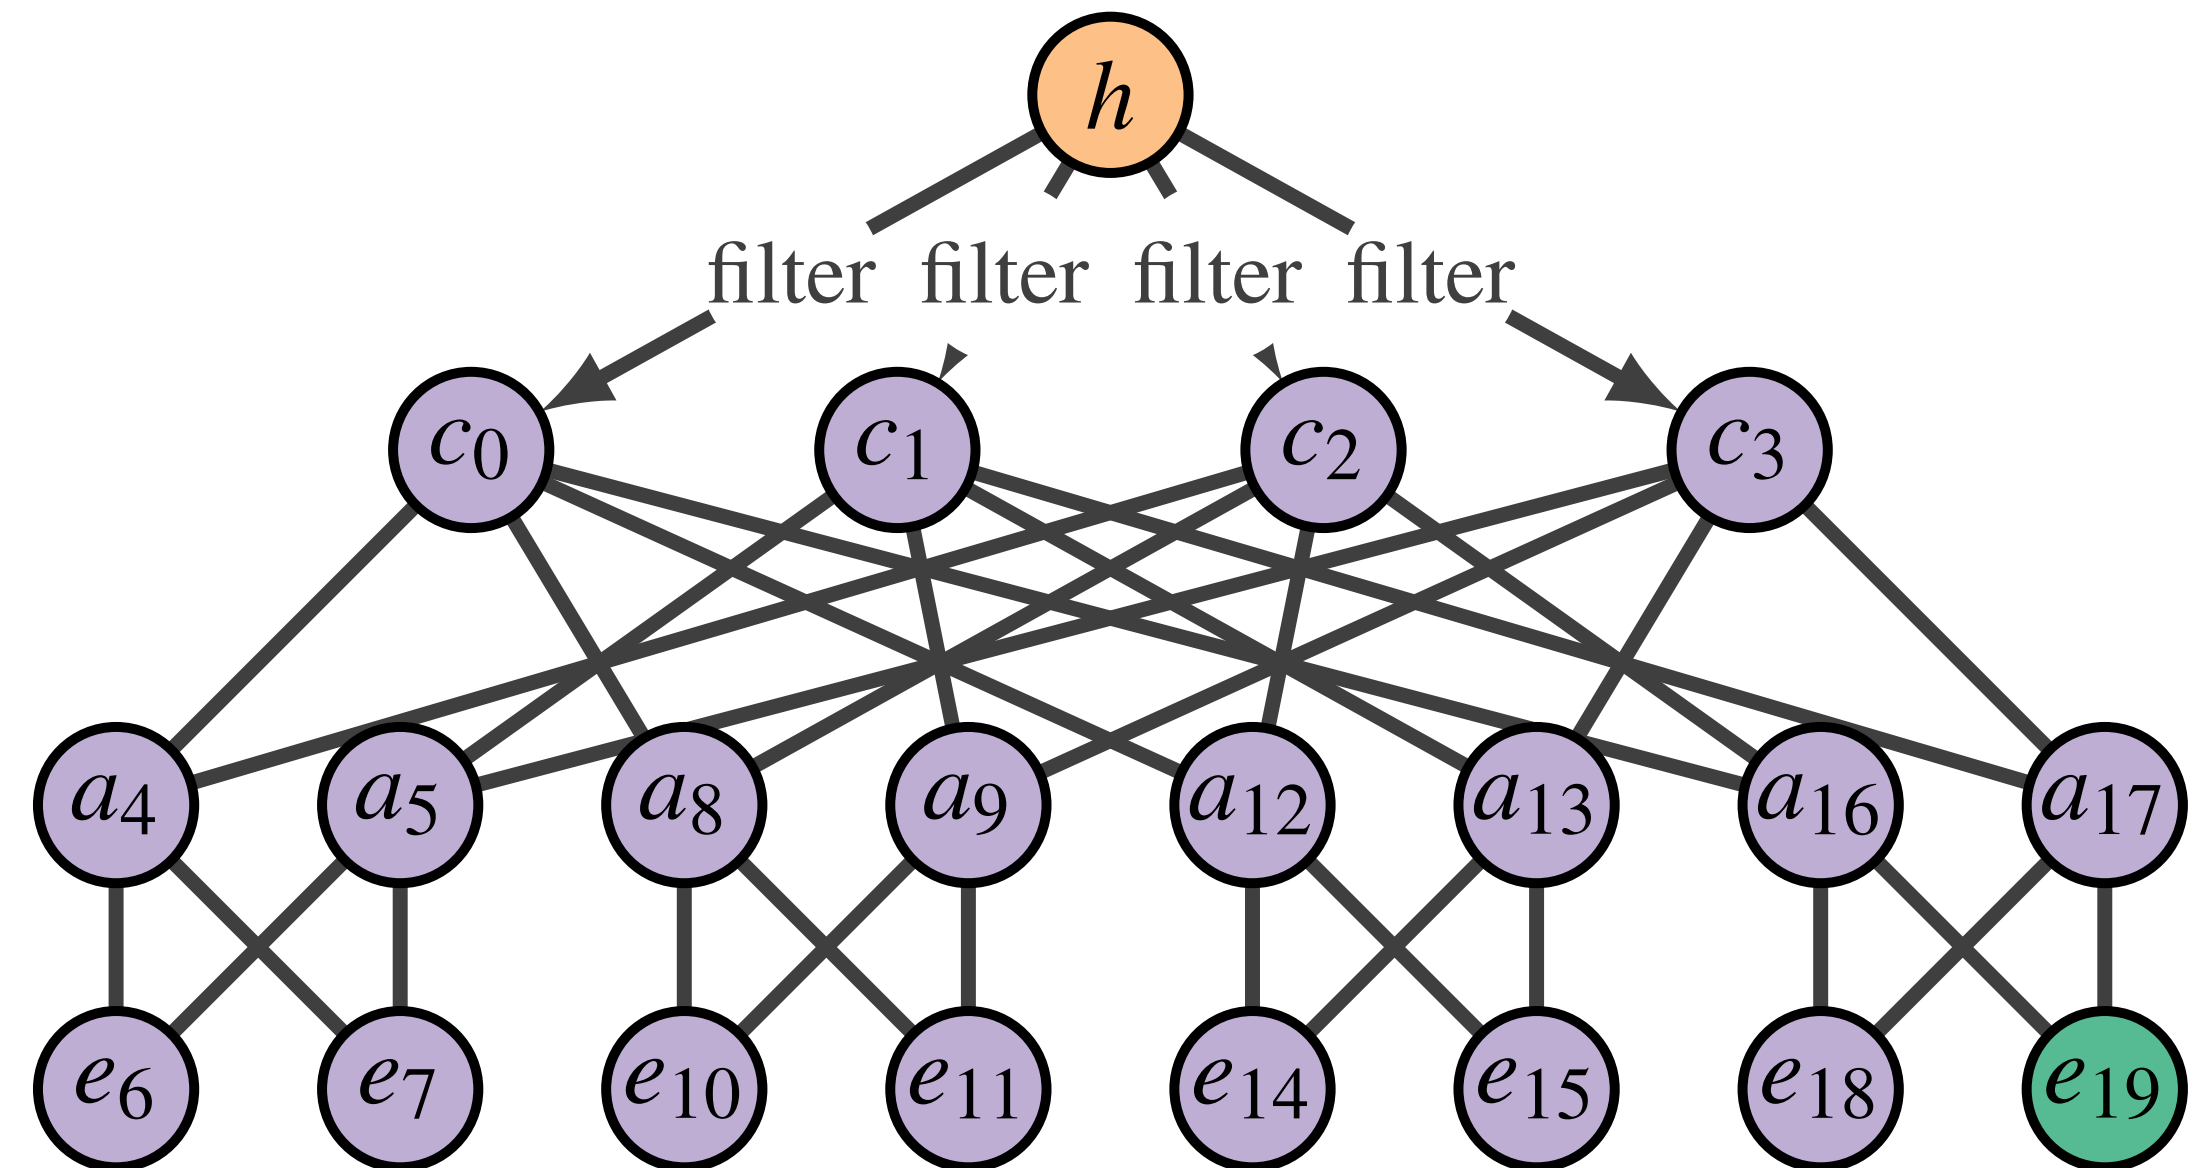
Sends a route with symbolic prefix  $p$

# Fat-tree Hijack Filtering

Converged routes for prefix  $p$  should not come from  $h$

$$P(v) \equiv \text{true } \mathcal{U}^4 s . \text{prefix} = p \wedge \neg s . \text{tag}$$

Interface **composes** an “eventual invariant” with a  
“safety invariant”





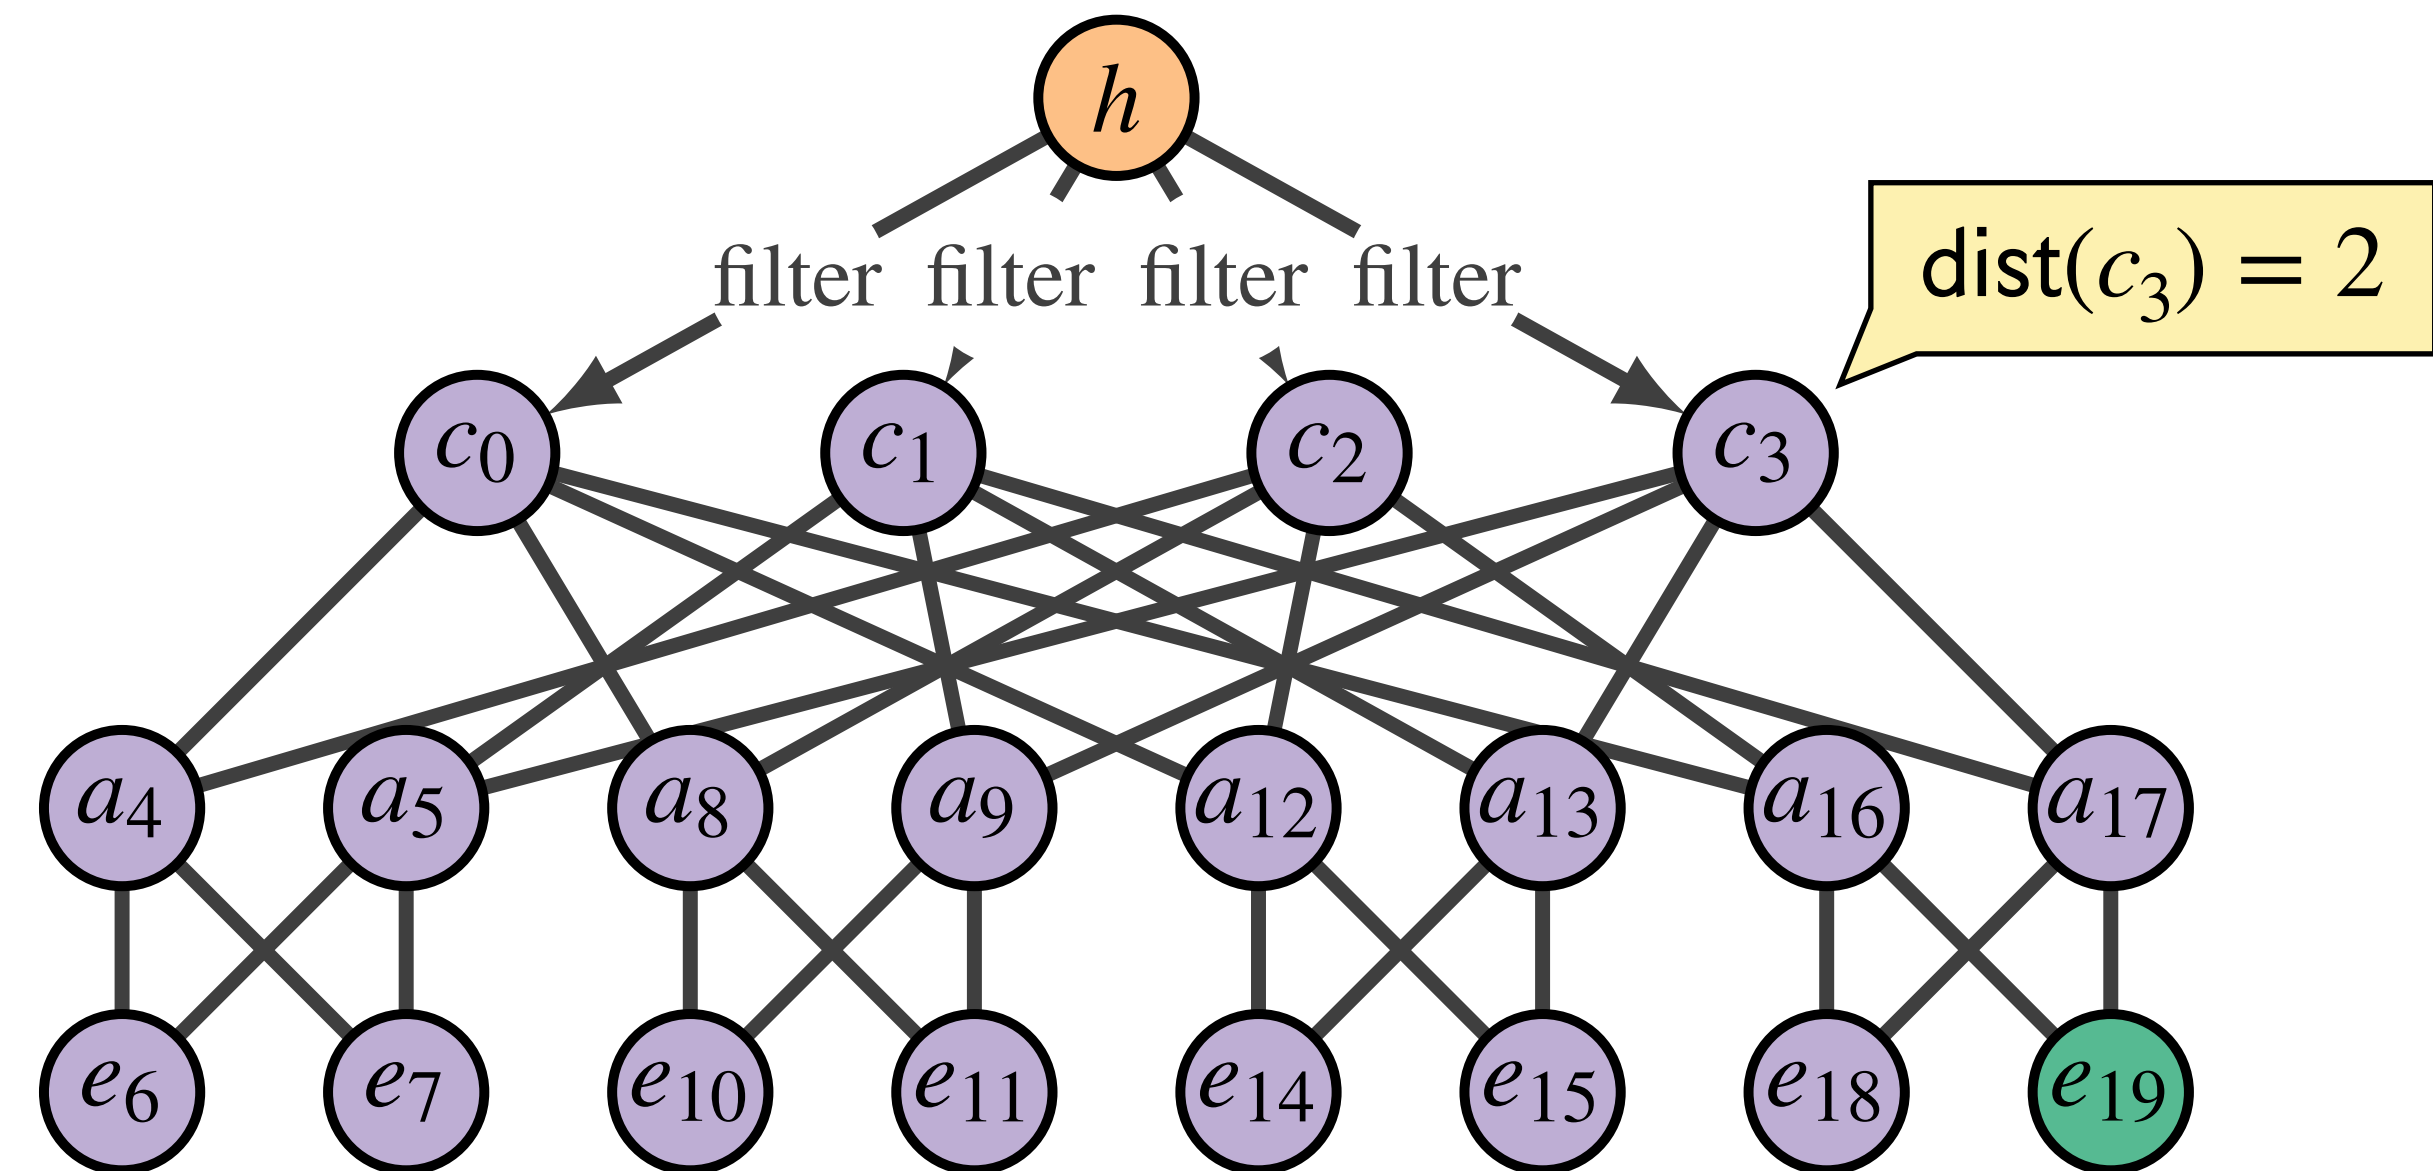
# Fat-tree Hijack Filtering

Converged routes for prefix  $p$  should not come from  $h$

$$P(v) \equiv \text{true } \mathcal{U}^4 s . \text{prefix} = p \wedge \neg s . \text{tag}$$

Interface **composes** an “eventual invariant” with a “safety invariant”

All nodes' interfaces are parameterized by their distance  $\text{dist}(v)$  from  $e_{19}$



# Fat-tree Hijack Filtering

Converged routes for prefix  $p$  should not come from  $h$


$$P(v) \equiv \text{true } \mathcal{U}^4 s . \text{prefix} = p \wedge \neg s . \text{tag}$$

Interface **composes** an “eventual invariant” with a “safety invariant”

All nodes’ interfaces are parameterized by their distance  $\text{dist}(v)$  from  $e_{19}$

Nodes are **eventually** “internally reachable”

$$\text{true } \mathcal{U}^{\text{dist}(v)} s . \text{prefix} = p \wedge \neg s . \text{tag}$$



# Fat-tree Hijack Filtering

Converged routes for prefix  $p$  should not come from  $h$

$$P(v) \equiv \text{true } \mathcal{U}^4 s . \text{prefix} = p \wedge \neg s . \text{tag}$$

Interface **composes** an “eventual invariant” with a “safety invariant”


All nodes’ interfaces are parameterized by their distance  $\text{dist}(v)$  from  $e_{19}$

Nodes are **eventually** “internally reachable”


$$\text{true } \mathcal{U}^{\text{dist}(v)} s . \text{prefix} = p \wedge \neg s . \text{tag}$$

Nodes **never** use hijacking routes

$$\mathcal{G}(s . \text{prefix} = p \rightarrow \neg s . \text{tag})$$

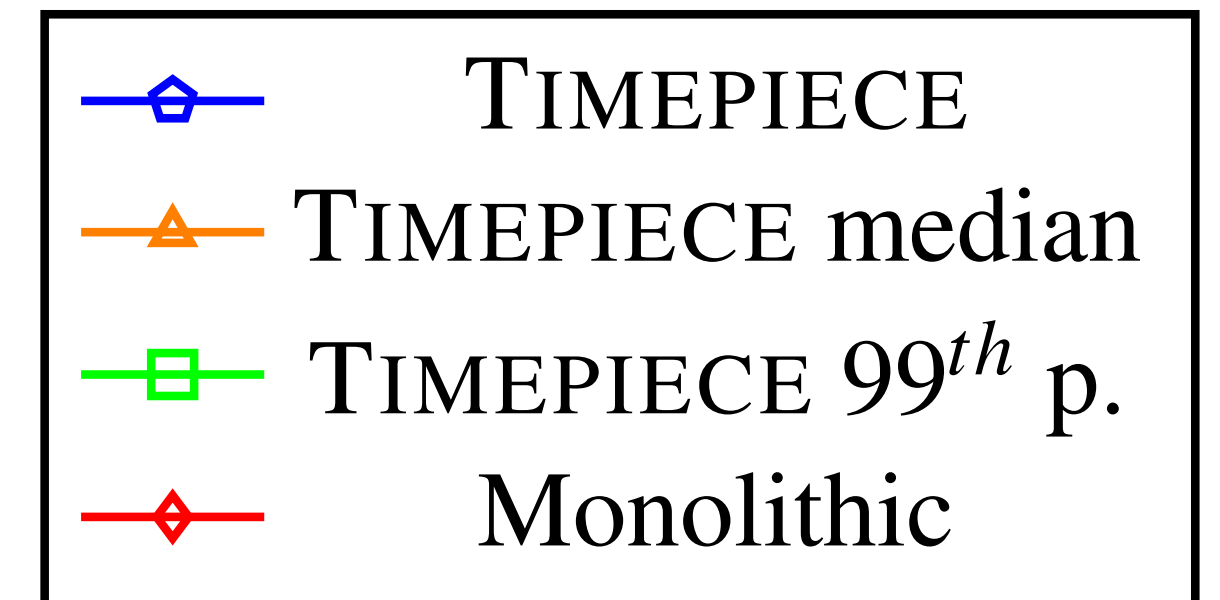



# Fat-tree Hijack Filtering




# Fat-tree Hijack Filtering

Monolithic verification times  
out (>2h) at 80 nodes







# Fat-tree Hijack Filtering



# Fat-tree Hijack Filtering



# Fat-tree Hijack Filtering




# Takeaways



# Takeaways

Big, complex control planes need modular tools






# Takeaways

Big, complex control planes need modular tools

Temporal invariants provide a **correct basis for modular verification**




# Takeaways

Big, complex control planes need modular tools

Temporal invariants provide a **correct basis for modular verification**

Scale to thousands of nodes & complex policies




# Takeaways

Big, complex control planes need modular tools

Temporal invariants provide a **correct basis for modular verification**

Scale to thousands of nodes & complex policies

Read the paper to learn more!





# Thank You!



Our paper

I'm looking for a job!  
[cs.princeton.edu/~tthijm](https://cs.princeton.edu/~tthijm)



**Tim Alberdingk Thijm**  
Princeton



**Ryan Beckett**  
Microsoft Research



**Aarti Gupta**  
Princeton



**Dave Walker**  
Princeton

# Extra slides




# Closed Completeness Theorem



# Closed Completeness Theorem

Starting from fixed initial routes, if  $\sigma(v)(t)$  is the (monolithic) state of node  $v$  at time  $t$ , then the interface  $A(v)(t) = \{\sigma(v)(t)\}$  satisfies the base and inductive checks for all nodes.






# Closed Completeness Theorem

Starting from fixed initial routes, if  $\sigma(v)(t)$  is the (monolithic) state of node  $v$  at time  $t$ , then the interface  $A(v)(t) = \{\sigma(v)(t)\}$  satisfies the base and inductive checks for all nodes.

Proof by induction on time.



# Evaluation

# Evaluation

## Benchmark

Reachability

Path length

Valley freedom

Hijack filtering

No transit



# Evaluation

Benchmark	Nodes
Reachability	20-2000
Path length	20-2000
Valley freedom	20-2000
Hijack filtering	20-2000
No transit	263

# Evaluation

Benchmark	Nodes	Network LoC
Reachability	20-2000	81
Path length	20-2000	88
Valley freedom	20-2000	89
Hijack filtering	20-2000	146
No transit	263	88 (+102,753)

# Evaluation

Benchmark	Nodes	Network LoC	Annotation LoC
Reachability	20-2000	81	3
Path length	20-2000	88	7
Valley freedom	20-2000	89	12
Hijack filtering	20-2000	146	21
No transit	263	88 (+102,753)	5

# Evaluation

Benchmark	Nodes	Network LoC	Annotation LoC	Monolithic hits 2h timeout?
Reachability	20-2000	81	3	No (fixed dest.) 80 nodes (symbolic dest.)
Path length	20-2000	88	7	80 nodes (fixed dest.) 20 nodes (symbolic dest.)
Valley freedom	20-2000	89	12	180 nodes
Hijack filtering	20-2000	146	21	80 nodes (fixed dest.) 20 nodes (symbolic dest.)
No transit	263	88 (+102,753)	5	Yes

# Evaluation

Benchmark	Nodes	Network LoC	Annotation LoC	Monolithic hits 2h timeout?	Modular verification time
Reachability	20-2000	81	3	No (fixed dest.) 80 nodes (symbolic dest.)	28s (fixed 2000 nodes) 336s (symbolic 2000 nodes)
Path length	20-2000	88	7	80 nodes (fixed dest.) 20 nodes (symbolic dest.)	1204s (fixed 2000 nodes) 3953s (symbolic 2000 nodes)
Valley freedom	20-2000	89	12	180 nodes	398s (fixed 2000 nodes) 3506s (symbolic 1280 nodes)
Hijack filtering	20-2000	146	21	80 nodes (fixed dest.) 20 nodes (symbolic dest.)	142s (2000 nodes) 2196s (symbolic 2000 nodes)
No transit	263	88 (+102,753)	5	Yes	38s



# Related Work

# Related Work

## Satisfiability Modulo Theories (SMT)-based verification

### Scalable Verification of Border Gateway Protocol Configurations with an SMT Solver

Konstantin Weitz   Doug Woos   Emina Torlak  
Michael D. Ernst   Arvind Krishnamurthy   Zachary Tatlock  
University of Washington, USA  
{weitzkon, dwoos, emina, mernst, arvind, ztatlock}@cs.washington.edu



### A General Approach to Network Configuration Verification

Ryan Beckett Princeton University	Aarti Gupta Princeton University
Ratul Mahajan Microsoft Research & Intentionet	David Walker Princeton University

# Related Work

## Satisfiability Modulo Theories (SMT)-based verification

### Scalable Verification of Border Gateway Protocol Configurations with an SMT Solver

Konstantin Weitz   Doug Woos   Emina Torlak  
Michael D. Ernst   Arvind Krishnamurthy   Zachary Tatlock  
University of Washington, USA  
{weitzkon, dwoos, emina, mernst, arvind, ztatlock}@cs.washington.edu



### A General Approach to Network Configuration Verification

Ryan Beckett   Aarti Gupta  
Princeton University   Princeton University  
Ratul Mahajan   David Walker  
Microsoft Research & Intentionet   Princeton University

## simulation-based verification

### Plankton: Scalable network configuration verification through model checking

*Santhosh Prabhu, Kuan-Yen Chou, Ali Kheradmand, P. Brighten Godfrey, Matthew Caesar*  
*University of Illinois at Urbana-Champaign*

### Tiramisu: Fast Multilayer Network Verification

Anubhavnidhi Abhashkumar\*, Aaron Gember-Jacobson<sup>†</sup>, Aditya Akella\*  
*University of Wisconsin - Madison\*, Colgate University<sup>†</sup>*

# Related Work

## Satisfiability Modulo Theories (SMT)-based verification

### Scalable Verification of Border Gateway Protocol Configurations with an SMT Solver

Konstantin Weitz   Doug Woos   Emina Torlak  
Michael D. Ernst   Arvind Krishnamurthy   Zachary Tatlock  
University of Washington, USA  
{weitzkon, dwoos, emina, mernst, ztatlock}@cs.washington.edu



### A General Approach to Network Configuration Verification

Ryan Beckett   Aarti Gupta  
Princeton University   Princeton University  
Ratul Mahajan   David Walker  
Microsoft Research & Intentionet   Princeton University

## simulation-based verification

### Plankton: Scalable network configuration verification through model checking

*Santhosh Prabhu, Kuan-Yen Chou, Ali Kheradmand, P. Brighten Godfrey, Matthew Caesar*  
*University of Illinois at Urbana-Champaign*

### Tiramisu: Fast Multilayer Network Verification

Anubhavnidhi Abhashkumar\*, Aaron Gember-Jacobson<sup>†</sup>, Aditya Akella\*  
*University of Wisconsin - Madison\*, Colgate University<sup>†</sup>*

## modular SMT-based verification


### LIGHTYEAR: Using Modularity to Scale BGP Control Plane Verification

Alan Tang   Ryan Beckett   Karthick Jayaraman  
*University of California, Los Angeles*   *Microsoft*   *Microsoft*  
Todd Millstein   George Varghese  
*UCLA / Intentionet*   *UCLA*

### Kirigami, the Verifiable Art of Network Cutting

Timothy Alberdingk Thijm   Ryan Beckett   Aarti Gupta   David Walker  
*Princeton University*   *Microsoft Research*   *Princeton University*   *Princeton University*  
Princeton, USA   Redmond, USA   Princeton, USA   Princeton, USA  
tthijm@cs.princeton.edu   ryan.beckett@microsoft.com   aartig@cs.princeton.edu   dpw@cs.princeton.edu

# Challenges



finding the correct invariants



synchronous network semantics