


Espresso: Comprehensively Reasoning About External Routes Using Symbolic Simulation

Dan Wang, Peng Zhang



Aaron Gember-Jacobson



Network Outages are Common and Costly

Google leaked prefixes –
and knocked Japan off
the Internet

Router Crashes Trigger Major Southwest
IT System Failure

By Chris Preimesberger - July 22, 2016

**United Airlines Grounds Flights,
Citing Computer Problems**

[United Airlines](#) grounded planes nationwide for nearly two hours Wednesday morning after a faulty computer network router disrupted its passenger reservations system.


Facebook outage: what went wrong and why did it take so long to fix after social platform went down?

Facebook issued a statement on Tuesday confirming that the cause of the outage was a configuration change to the backbone routers that coordinate traffic between the company's data centres, which had a cascading effect on a halt.

**BGP Route Leak at Angola Cables
Slows Connectivity for Many
Australians**

By Aftab Siddiqui – 25 May 2023

Control Plane Verifiers (CPVs)



Req #1: Reasoning About Failures

Google leaked prefixes –
and knocked Japan off
the Internet

Router Crashes

By Chris Preimesberger - July 22, 2016

Facebook outage: what went wrong and why did it take so long to fix after social platform went down?

Facebook issued a statement on Tuesday confirming that the cause of the outage was a configuration change to the backbone routers that coordinate traffic between the company's data centres, which had a cascading effect on the network.

United Airlines Grounds Flights, Citing Computer Problems

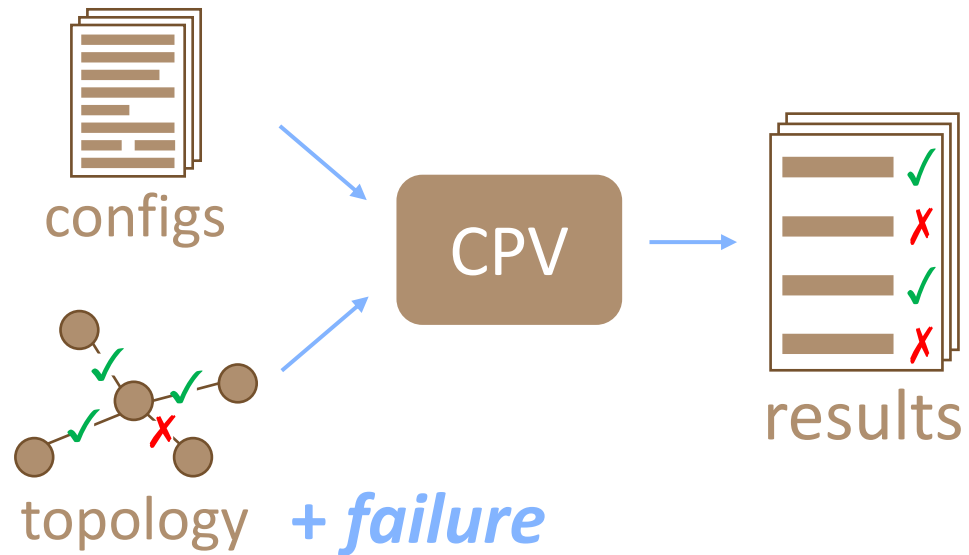
United Airlines grounded planes nationwide for nearly two hours Wednesday morning after a faulty computer network router disrupted its passenger reservations system.

BGP Route Leak at Angola Cables Slows Connectivity for Many Australians

By Aftab Siddiqui – 25 May 2023



CPV #1: Reasoning About *Failures*



Minesweeper^[3]

Hoyan^[6]

SRE^[7]

.....



Req #2: Reasoning About Config Updates

Google leaked prefixes –
and knocked Japan off
the Internet

Router Crashes Trigger
IT System Failure

By Chris Preimesberger - July 22, 2016

**United Airlines Grounds Flights,
Citing Computer Problems**

[United Airlines](#) grounded planes nationwide for nearly two hours Wednesday morning after a faulty computer network router disrupted its passenger reservations system.

Facebook outage: what went wrong and why did it take so long to fix after social platform went down?

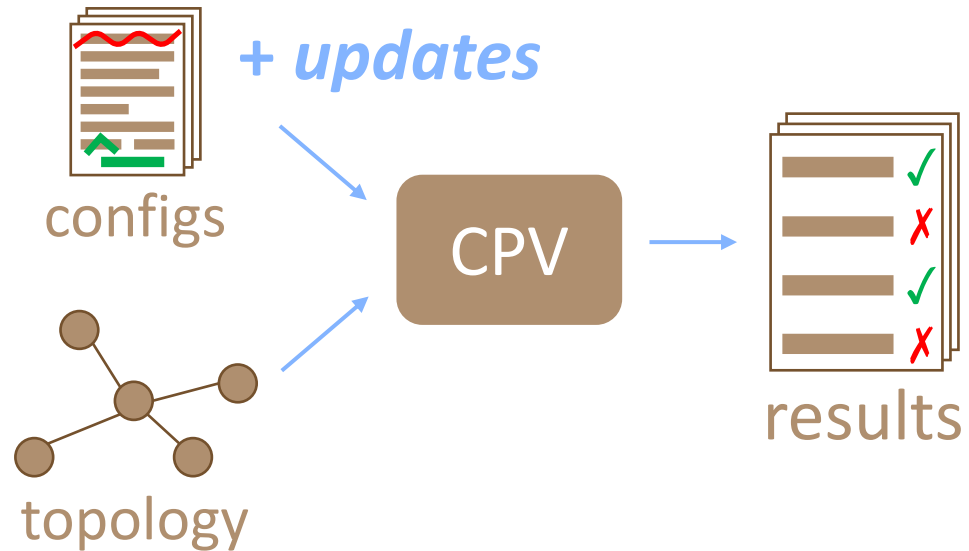
Facebook issued a statement on Tuesday confirming that the cause of the outage was a configuration change to the backbone routers that coordinate network traffic between the company's data centres, which had a cascading effect, bringing all Facebook services to a halt.

**BGP Route Leak at Angola Cables
Slows Connectivity for Many
Australians**

By Aftab Siddiqui – 25 May 2023



CPV #2: Reasoning About *Config Updates*



DNA^[8]
JinJing^[13]
.....



Req #3: Reasoning About External Routes

Google leaked prefixes –
and knocked Japan off
the Internet

Router Crashes Trigger Major Southwest
IT System Failure

By Chris Preimesberger - July 22, 2016

**United Airlines Grounds Flights,
Citing Computer Problems**

[United Airlines](#) grounded planes nationwide for nearly two hours Wednesday morning after a faulty computer network router disrupted its passenger reservations system.

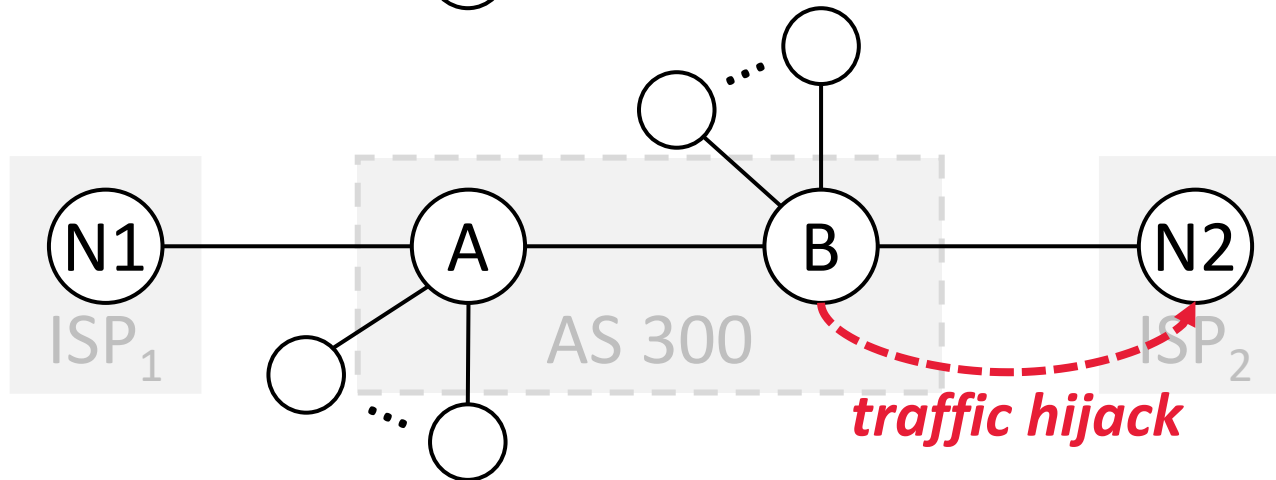

Facebook outage: what went wrong and why did it take so long to fix after social platform went down?

Facebook issued a statement on Tuesday confirming that the cause of the outage was a configuration change to the backbone routers that coordinate traffic between the company's data centres, which had a cascading effect on the network.


**BGP Route Leak at Angola Cables
Slows Connectivity for Many
Australians**

By Aftab Siddiqui – 25 May 2023

Reasoning About External Route is Important




Space of External Route is Colossal



#Prefix = $2^{33} - 1$

- 0.0.0.0/0
- 0.0.0.0/1
- 128.0.0.0/1
-
- 255.255.255.255/32


Space of External Route is Colossal



#Prefix = $2^{33} - 1$

- 0.0.0.0/0
- 0.0.0.0/1
- 128.0.0.0/1
-
- 255.255.255.255/32

Space of External Route is Colossal




#Prefix = $2^{33} - 1$

#Prefix-set = $2^{2^{33}-1}$

- 0.0.0.0/0
- 0.0.0.0/1
- 128.0.0.0/1
-
- 255.255.255.255/32

Space of External Route is Colossal



#Prefix = $2^{33} - 1$


#Prefix-set = $2^{2^{33}-1}$

- 0.0.0.0/0
- 0.0.0.0/1
- 128.0.0.0/1
-
- 255.255.255.255/32

Space size = $2^{(2^{33}-1) \times m}$

#Neighbor

Space of External Route is Colossal




How to handle such a colossal space?



$$\text{Space size} = 2^{(2^{33}-1) \times m}$$

#Neighbor


Method #1: Enumerate Concrete External Routes



SRE^[7], Hoyan^[6], Tiramisu^[9], Plankton^[10], Shapeshifter^[4], ERA^[11], Batfish^[1]

Enumeration is infeasible

Method #2: Symbolize External Routes with SMT





NV^[5], Minesweeper^[3], Bagpipe^[12]

SMT solving is Unscalable

SMT constraints

$p_1^1 = 1, \dots,$
 $len_1 \leq 16,$
 $p_2^1 = 0, \dots,$
 $16 \leq len_2 \leq 24,$
...






Express

Our Approach: Espresso



Espresso


Example Network



Example Network

3-bit IP addresses: 110, 010, ...

3-bit prefixes: 000/0, 000/1, ..., 111/3




Example Network

3-bit IP addresses: 110, 010, ...

3-bit prefixes: 000/0, 000/1, ..., 111/3

Space size = $2^{15 \times 2}$


15



Example Network

```
route-policy im1 permit node 100  
if-match prefix 100/2 110/2  
set-local-preference 200
```


```
route-policy im2 permit node 100  
if-match prefix 100/2 110/2  
set-local-preference 100
```



Example Network

```
route-policy im1 permit node 100  
if-match prefix 100/2 110/2  
set-local-preference 200
```

```
route-policy im2 permit node 100  
if-match prefix 100/2 110/2  
set-local-preference 100
```



**Preferred exit
to Internet**


Observation1: Equivalence

*Prefix
Equivalence*




```
route-policy im1 permit node 100  
if-match prefix 100/2 110/2  
set-local-preference 200
```

```
route-policy im2 permit node 100  
if-match prefix 100/2 110/2  
set-local-preference 100
```



Observation1: Equivalence


Advertiser
Equivalence



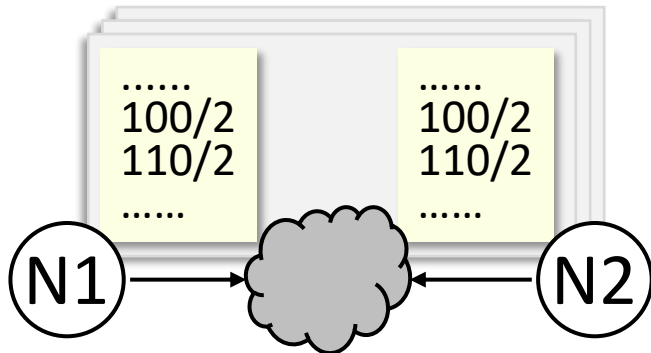
For 100/2 and 110/2

```
route-policy im1 permit node 100  
if-match prefix 100/2 110/2  
set-local-preference 200
```

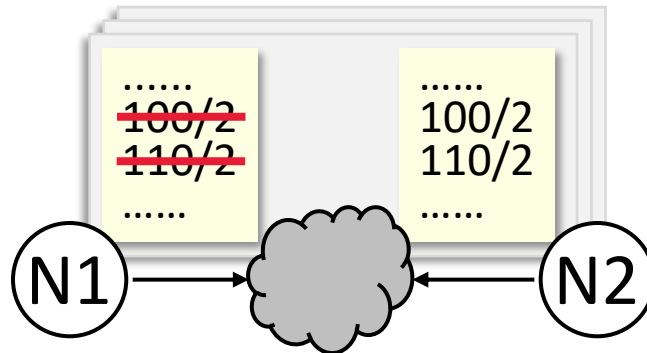
```
route-policy im2 permit node 100  
if-match prefix 100/2 110/2  
set-local-preference 100
```




How to Explore Equivalences?



EC#1: N1 announces
100/2, 110/2



EC#2: Only N2 announces
100/2, 110/2

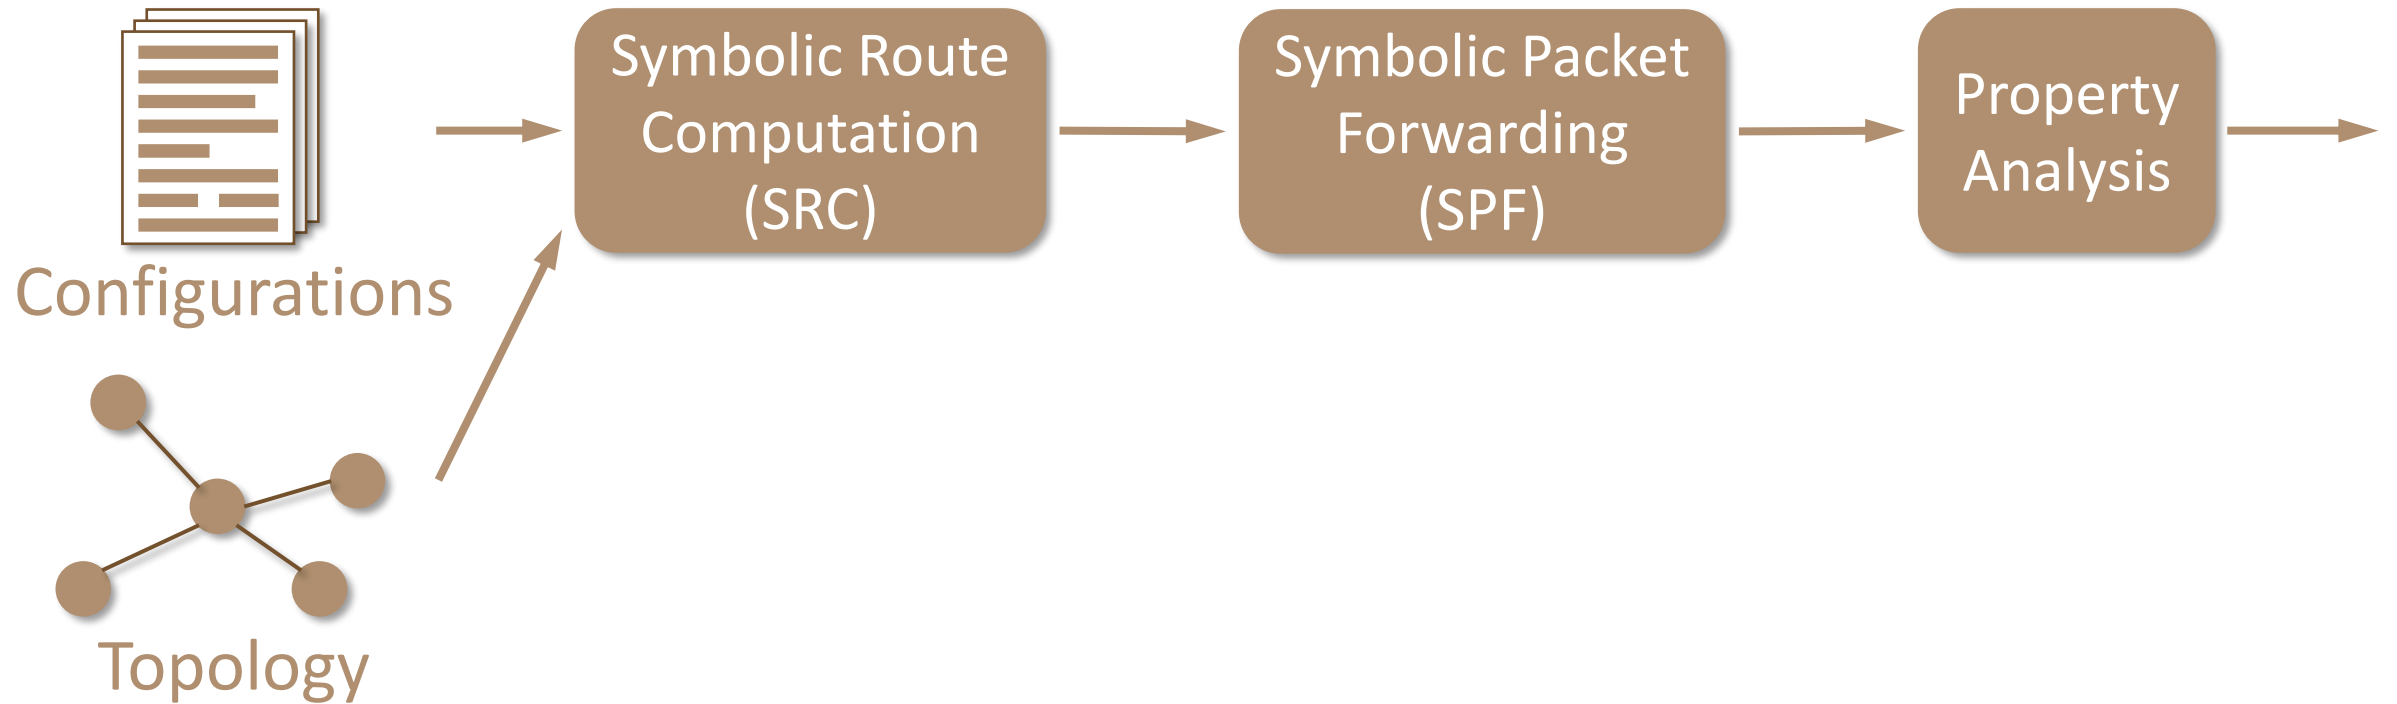


EC#3: No announcement for
100/2, 110/2

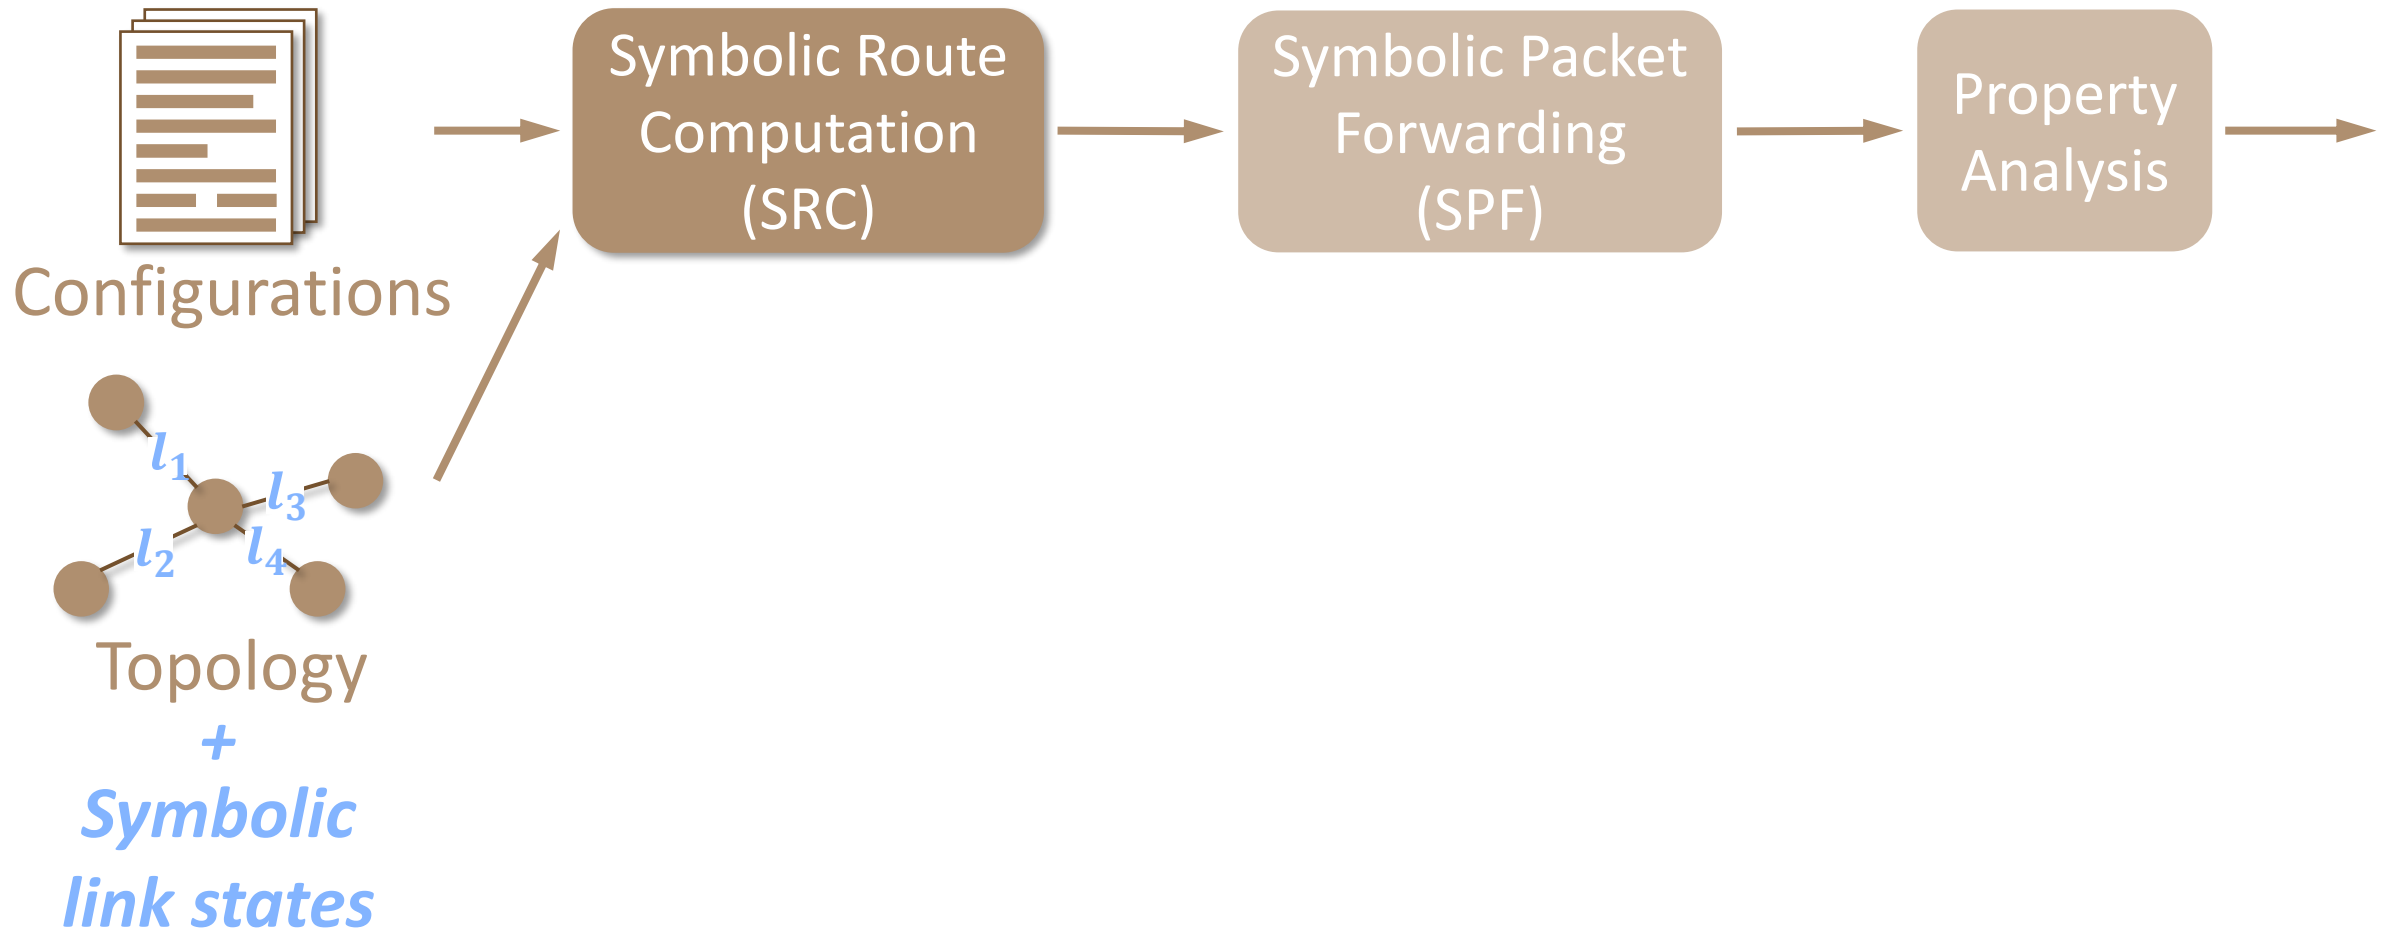


SRE^[7] SIGCOMM'22
Symbolic simulation

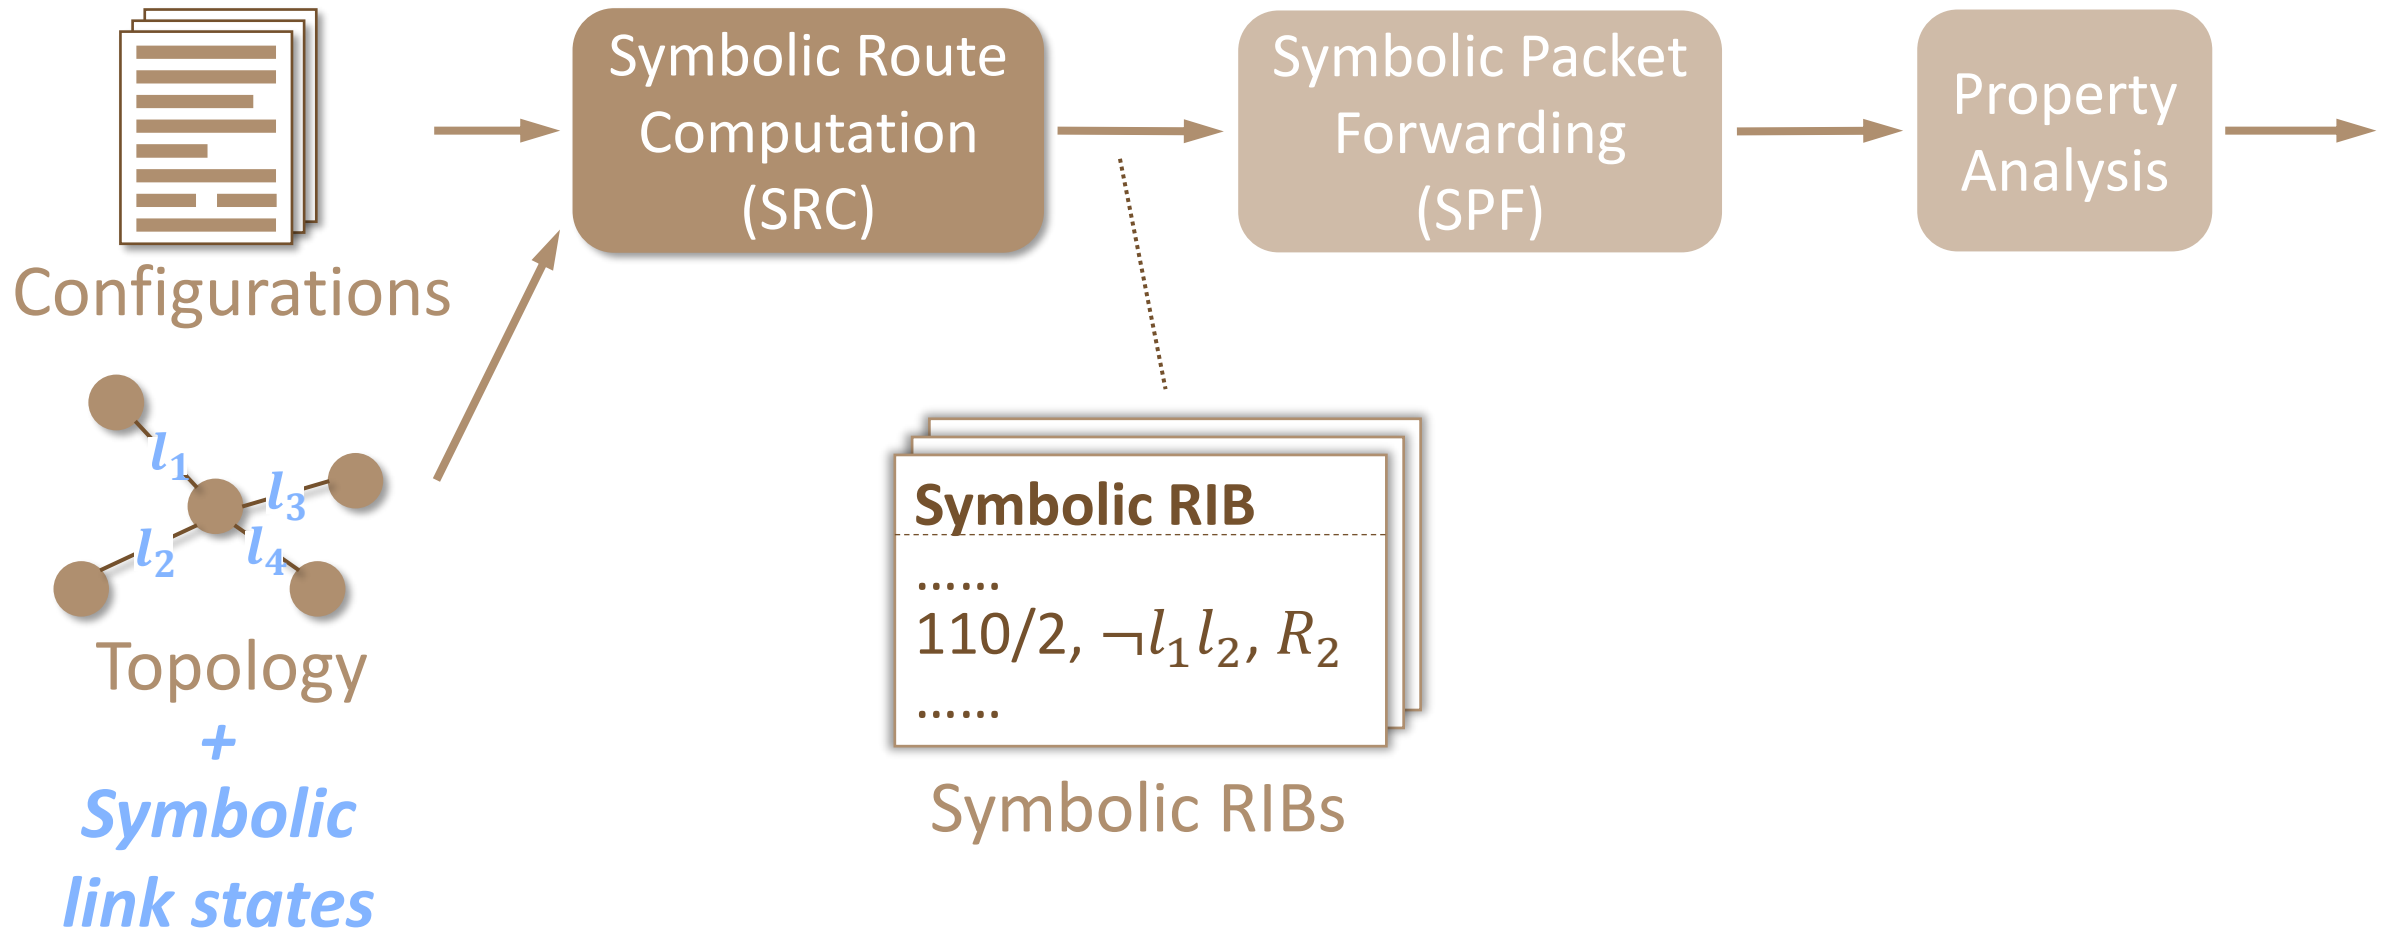
Workflow of Symbolic Simulation^[7]



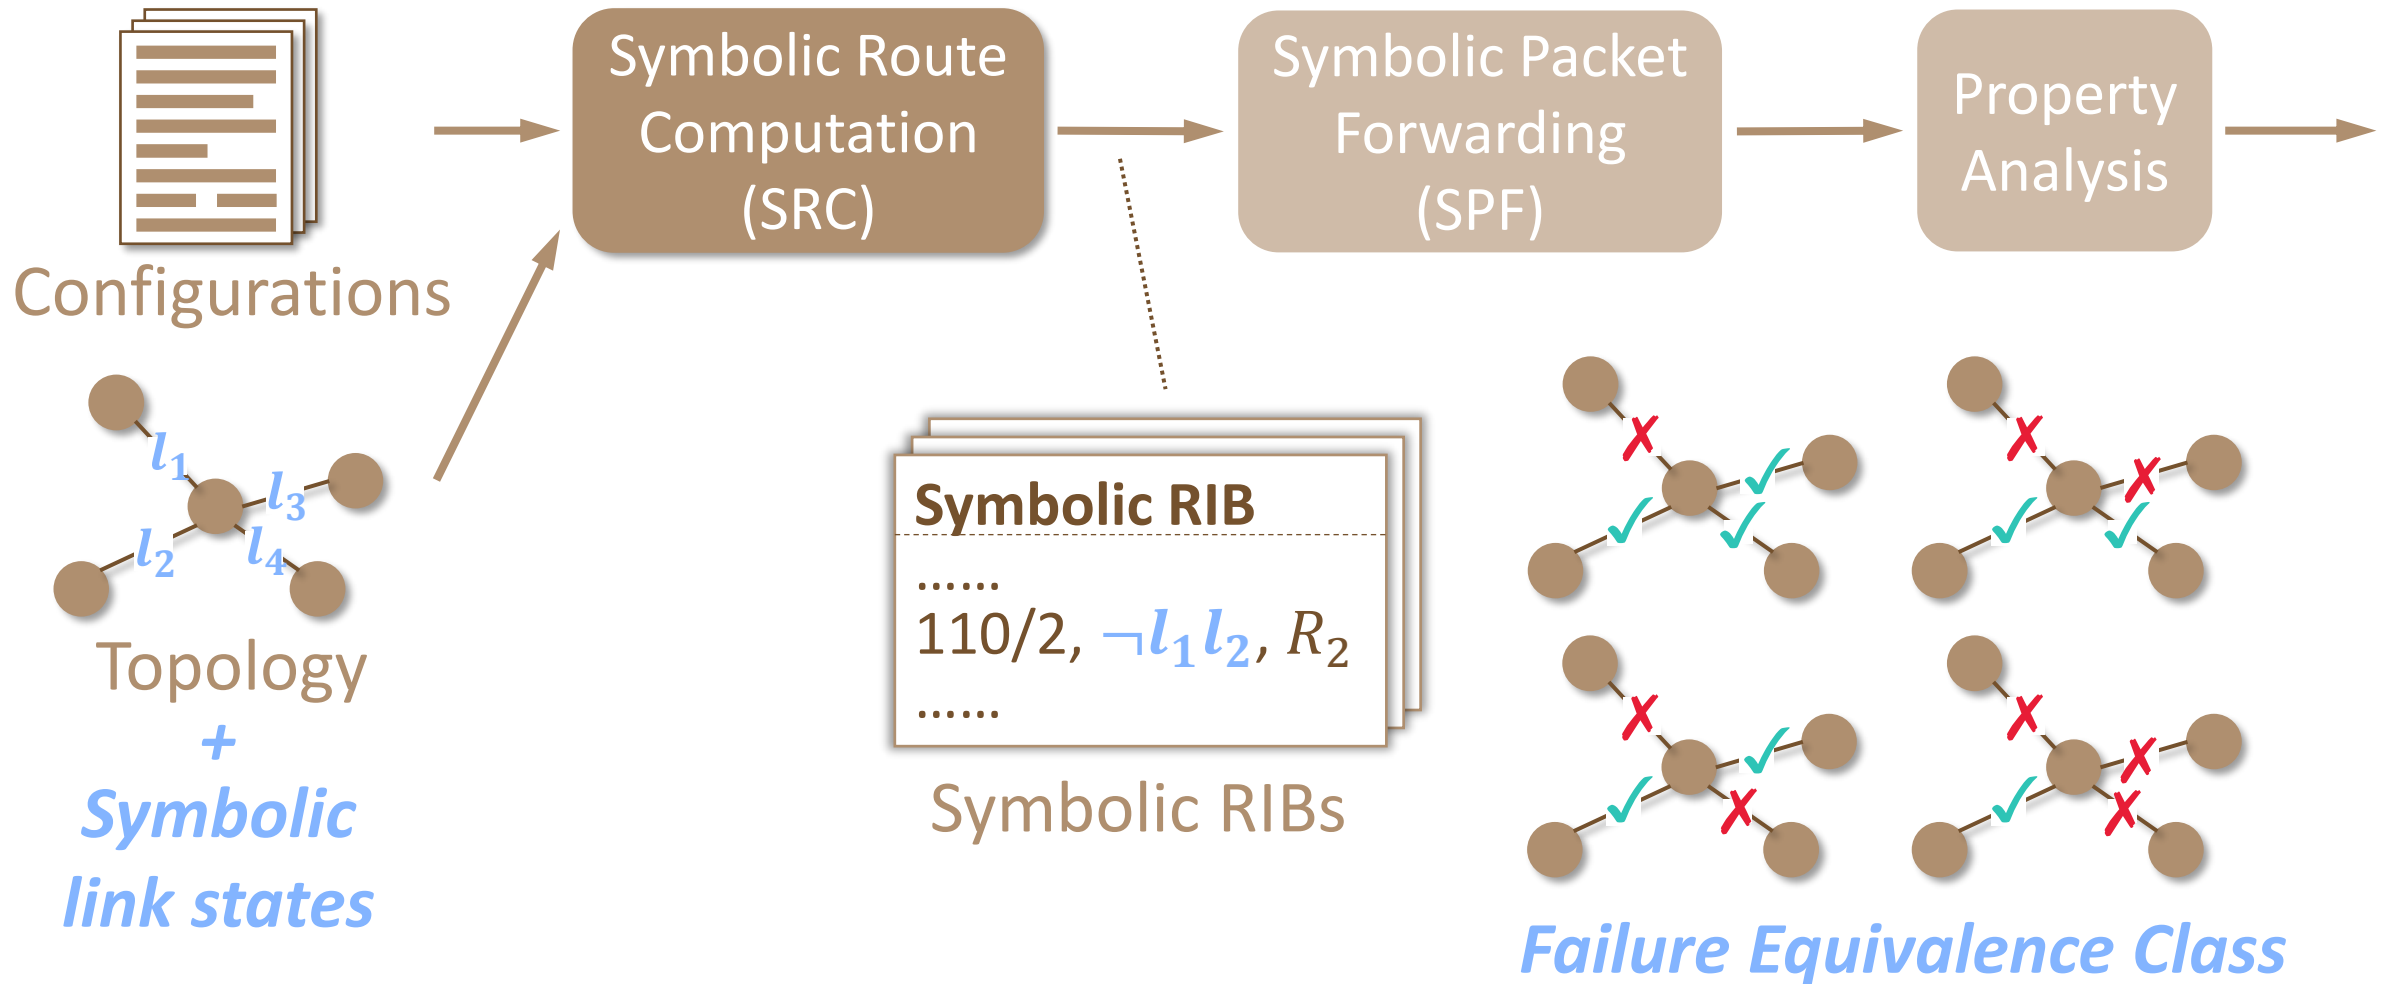
SRE Makes Link Symbolic




SRE Computes Failure Equivalence Classes



SRE Computes Failure Equivalence Classes

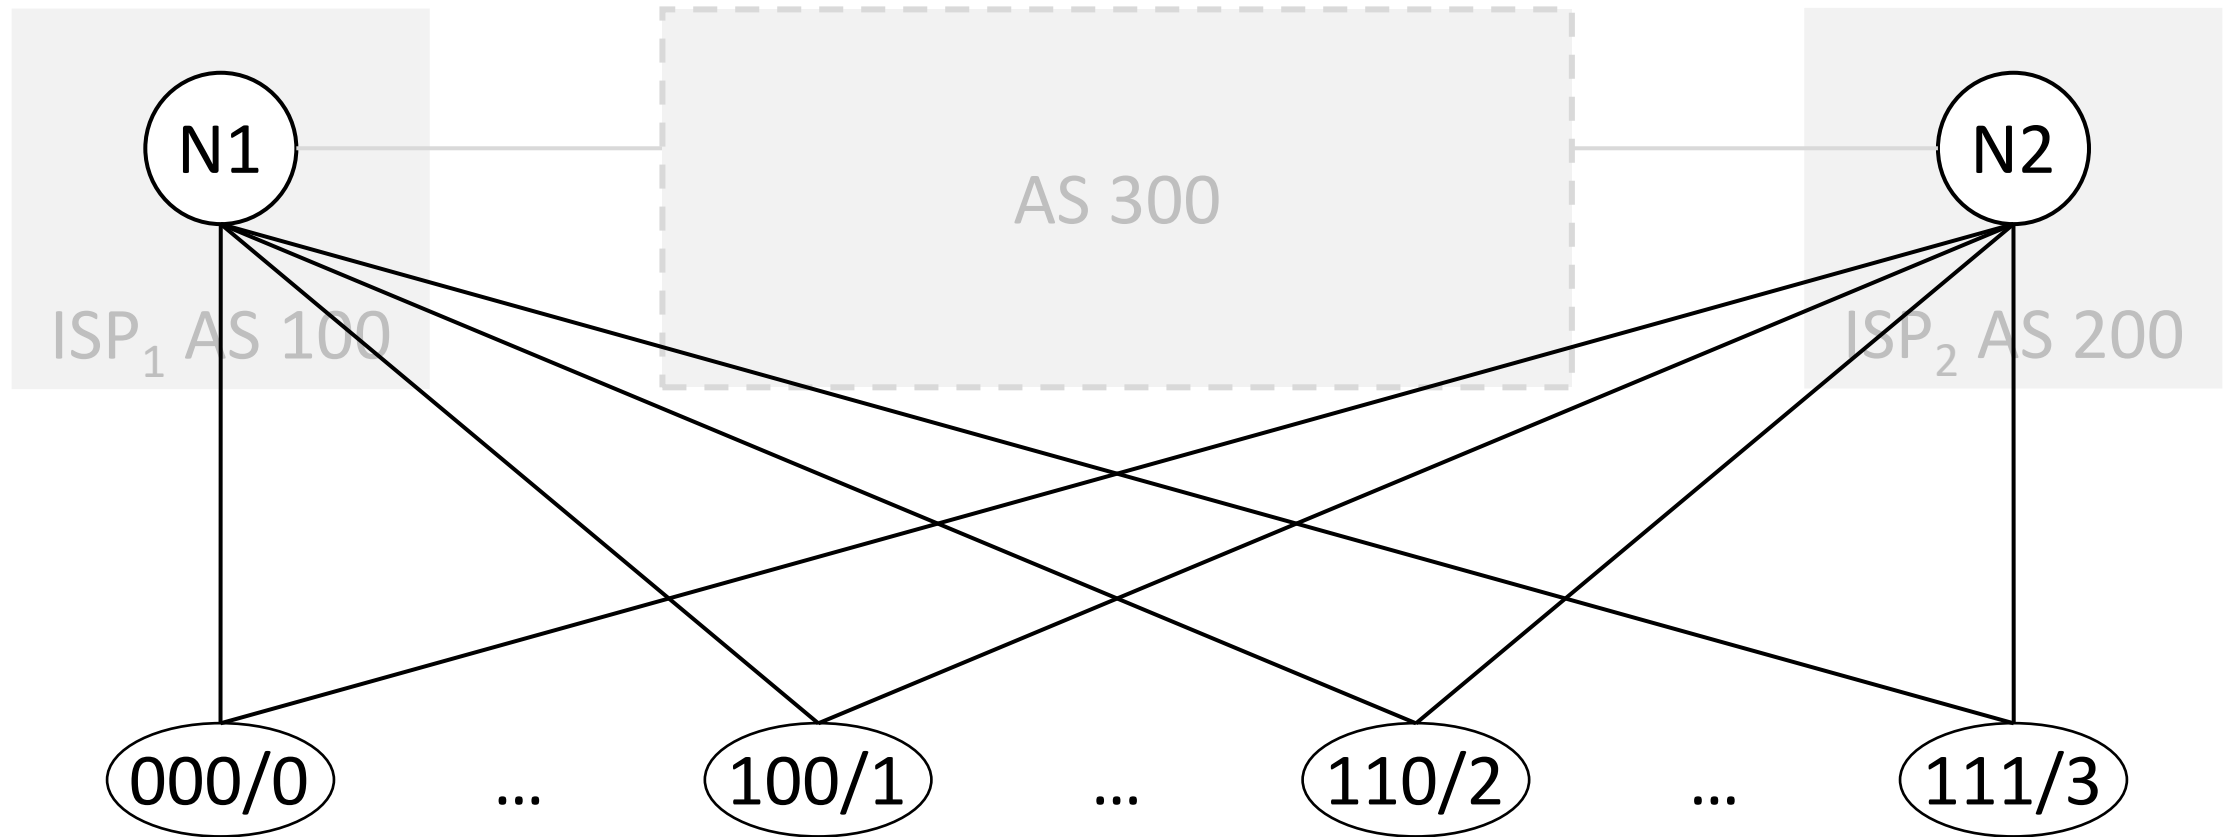


Modeling Symbolic Routes As Symbolic Links



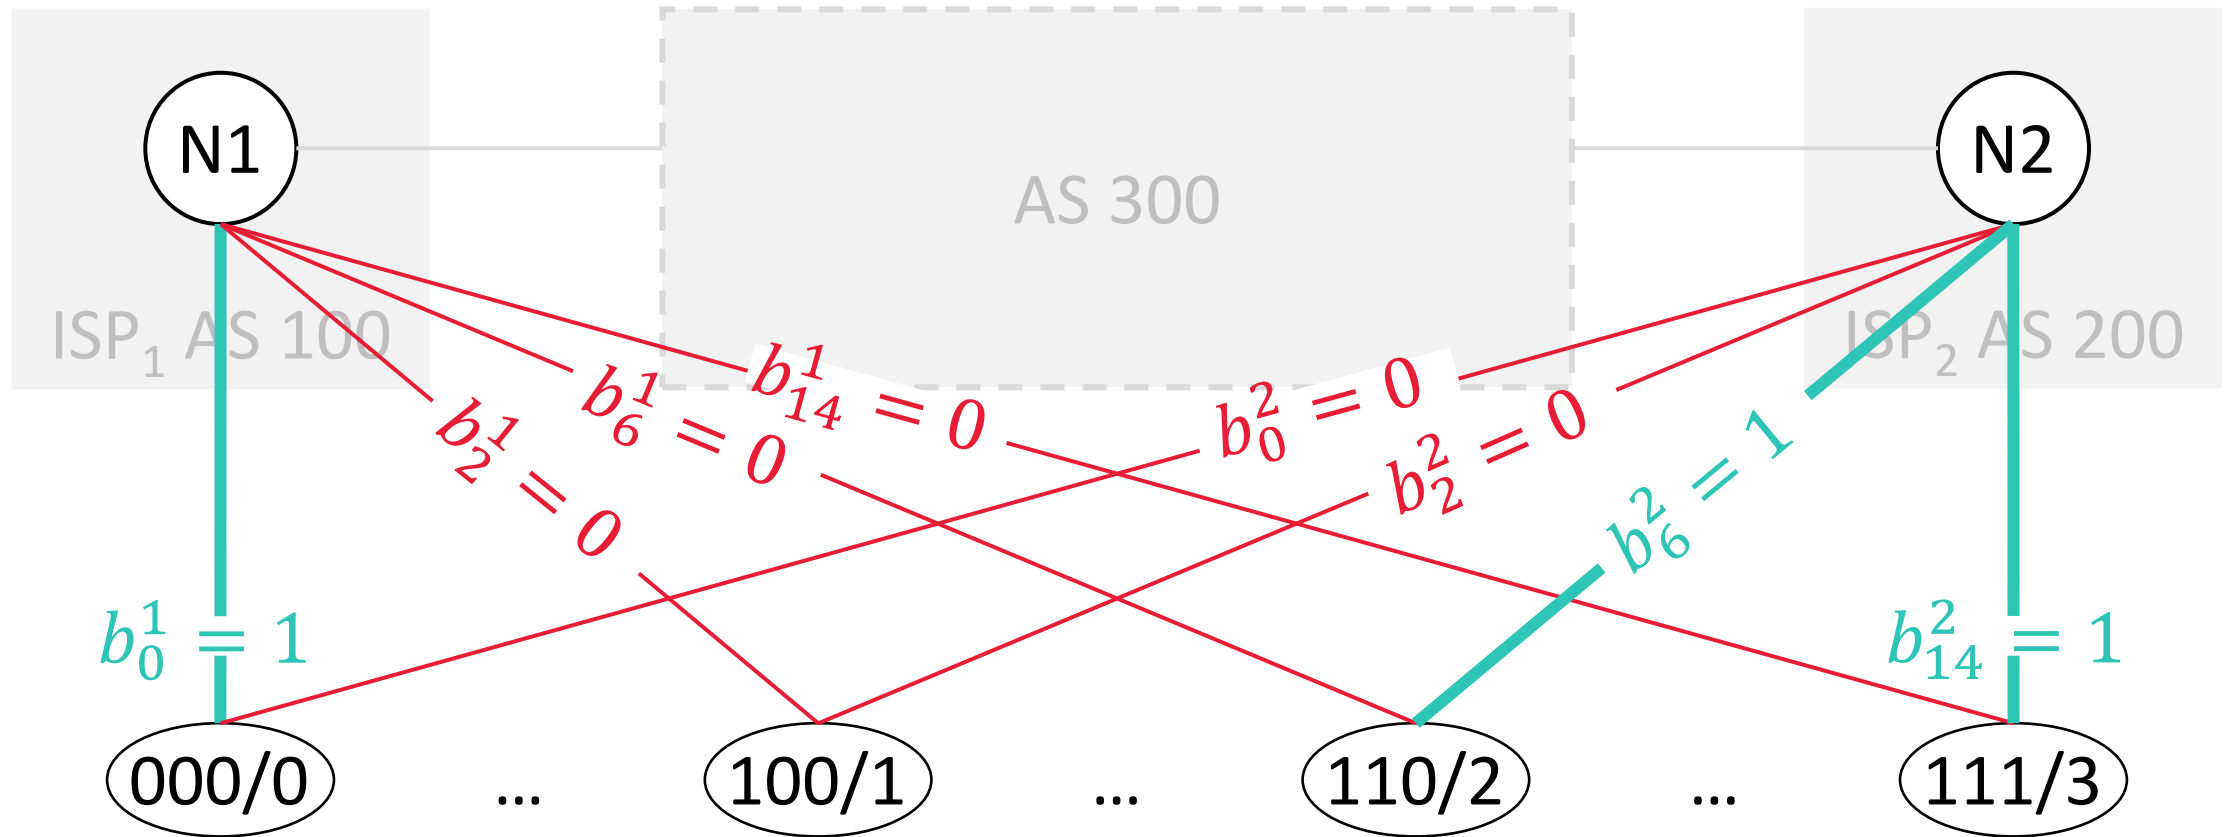
(000/0) ... (100/1) ... (110/2) ... (111/3)

Modeling Symbolic Routes As Symbolic Links

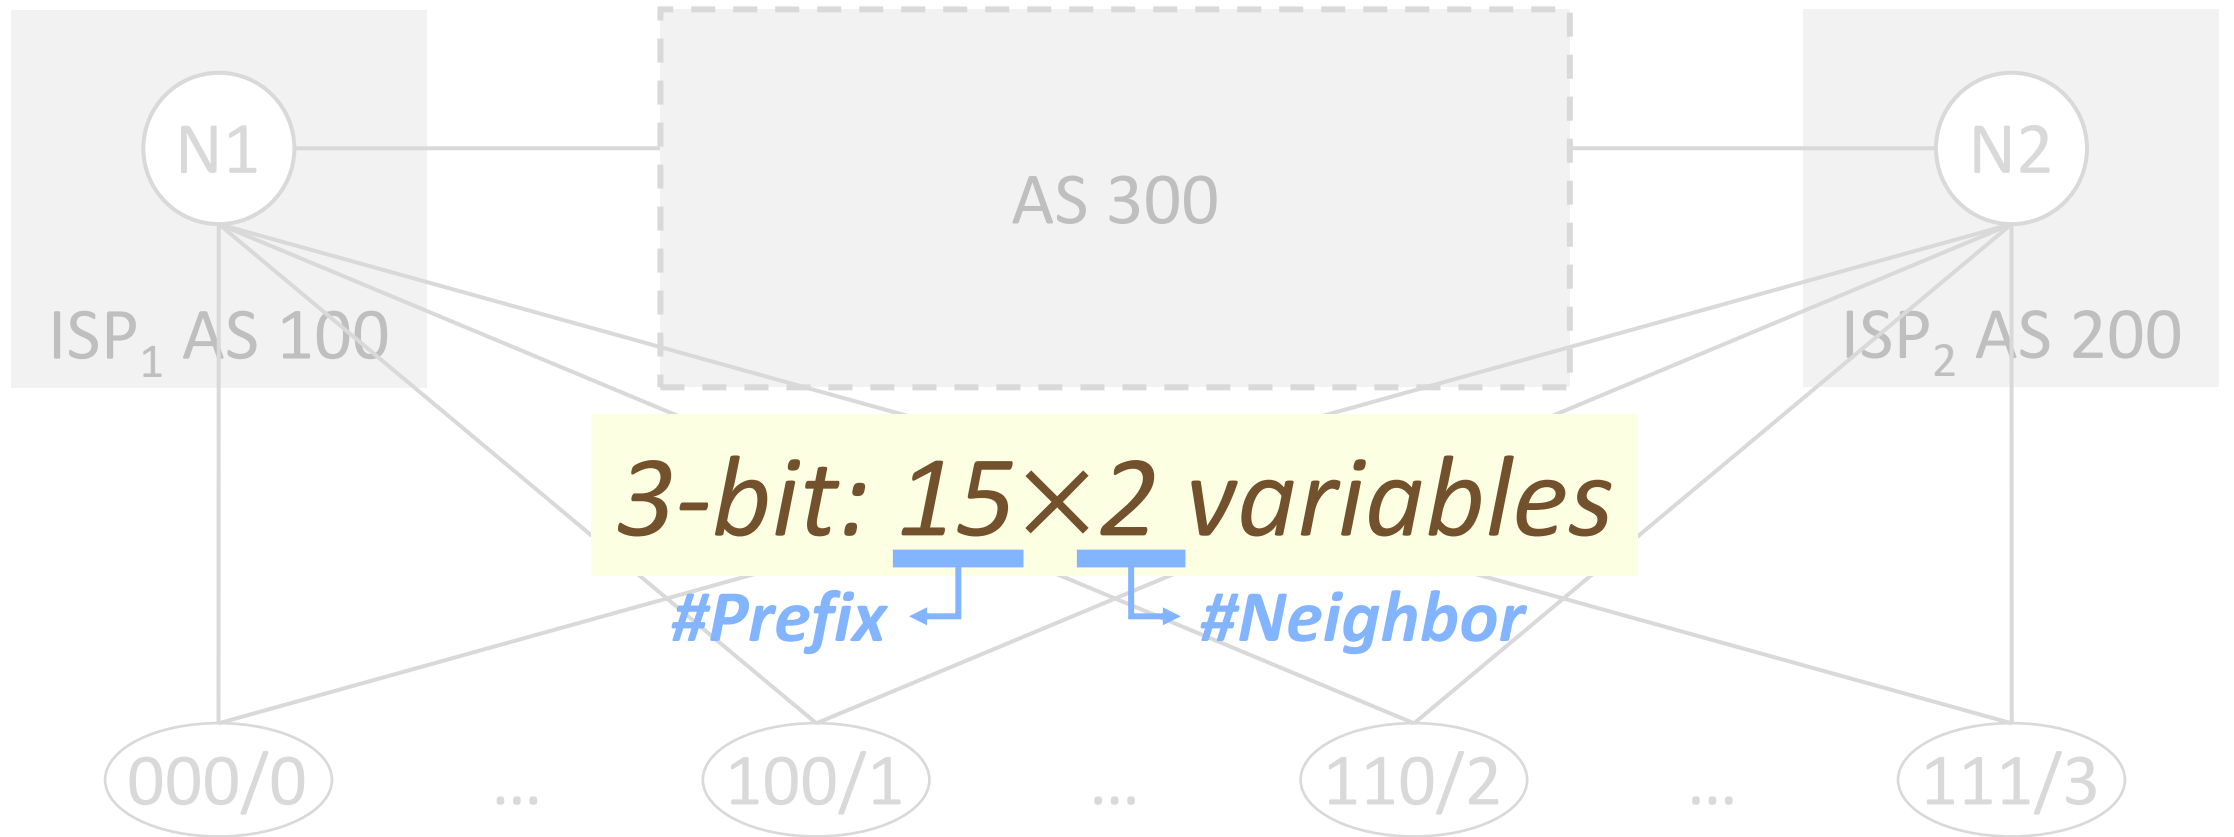


SRE Encoding: Select by Link Variables

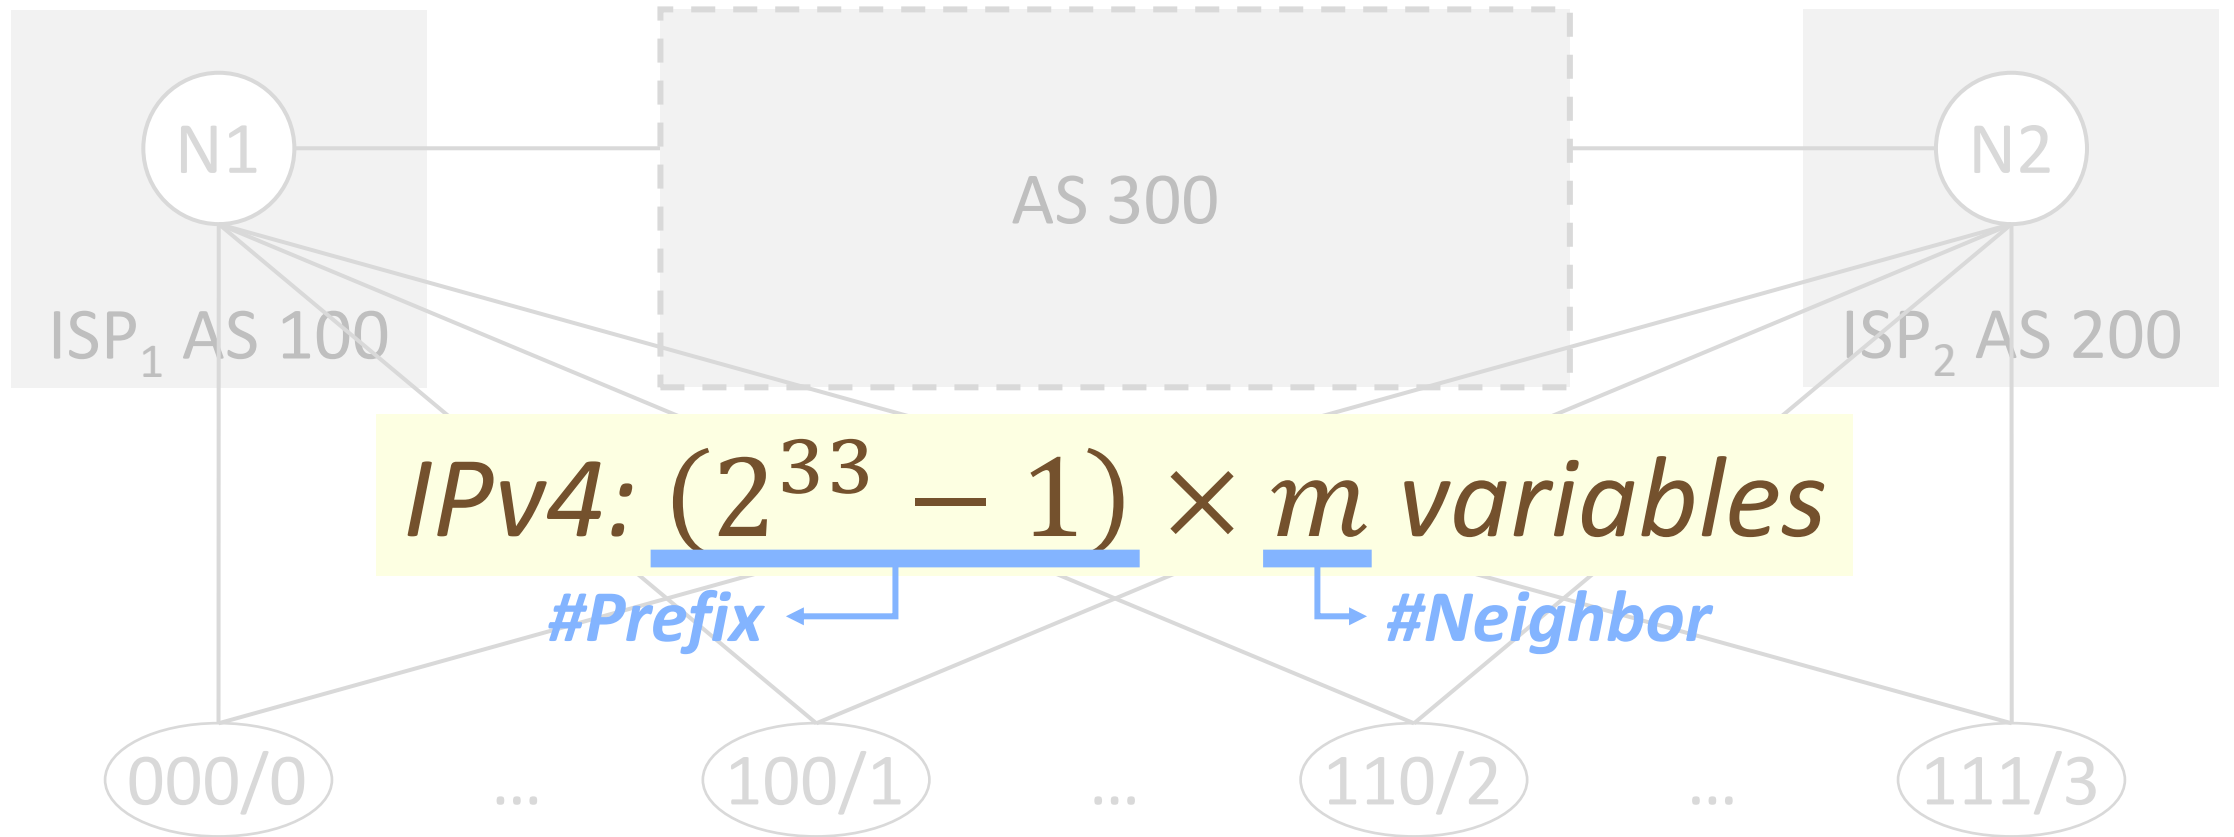
N1 announces 000/0, N2 announces 110/2 and 111/3



SRE Encoding: Select by Link Variables

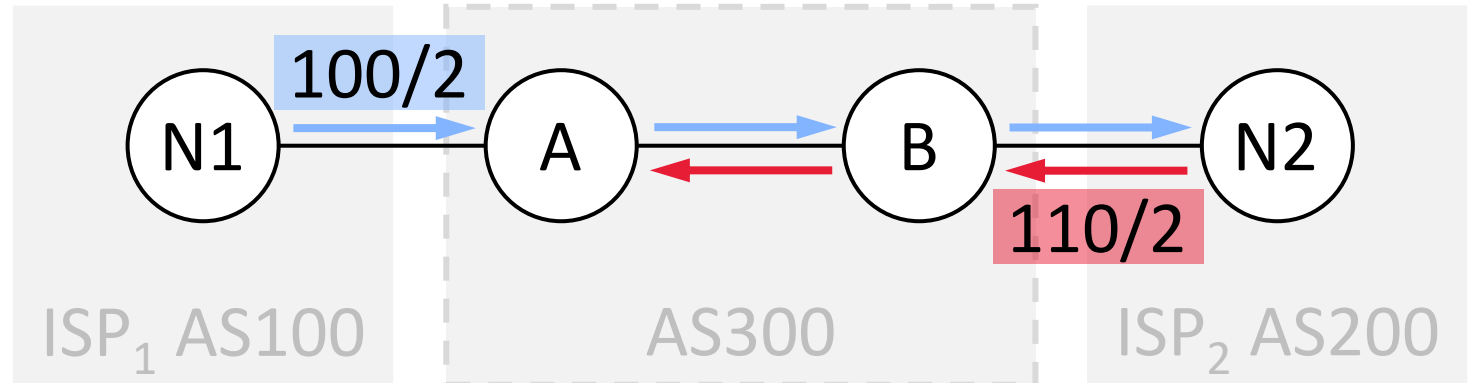


SRE Encoding: Select by Link Variables



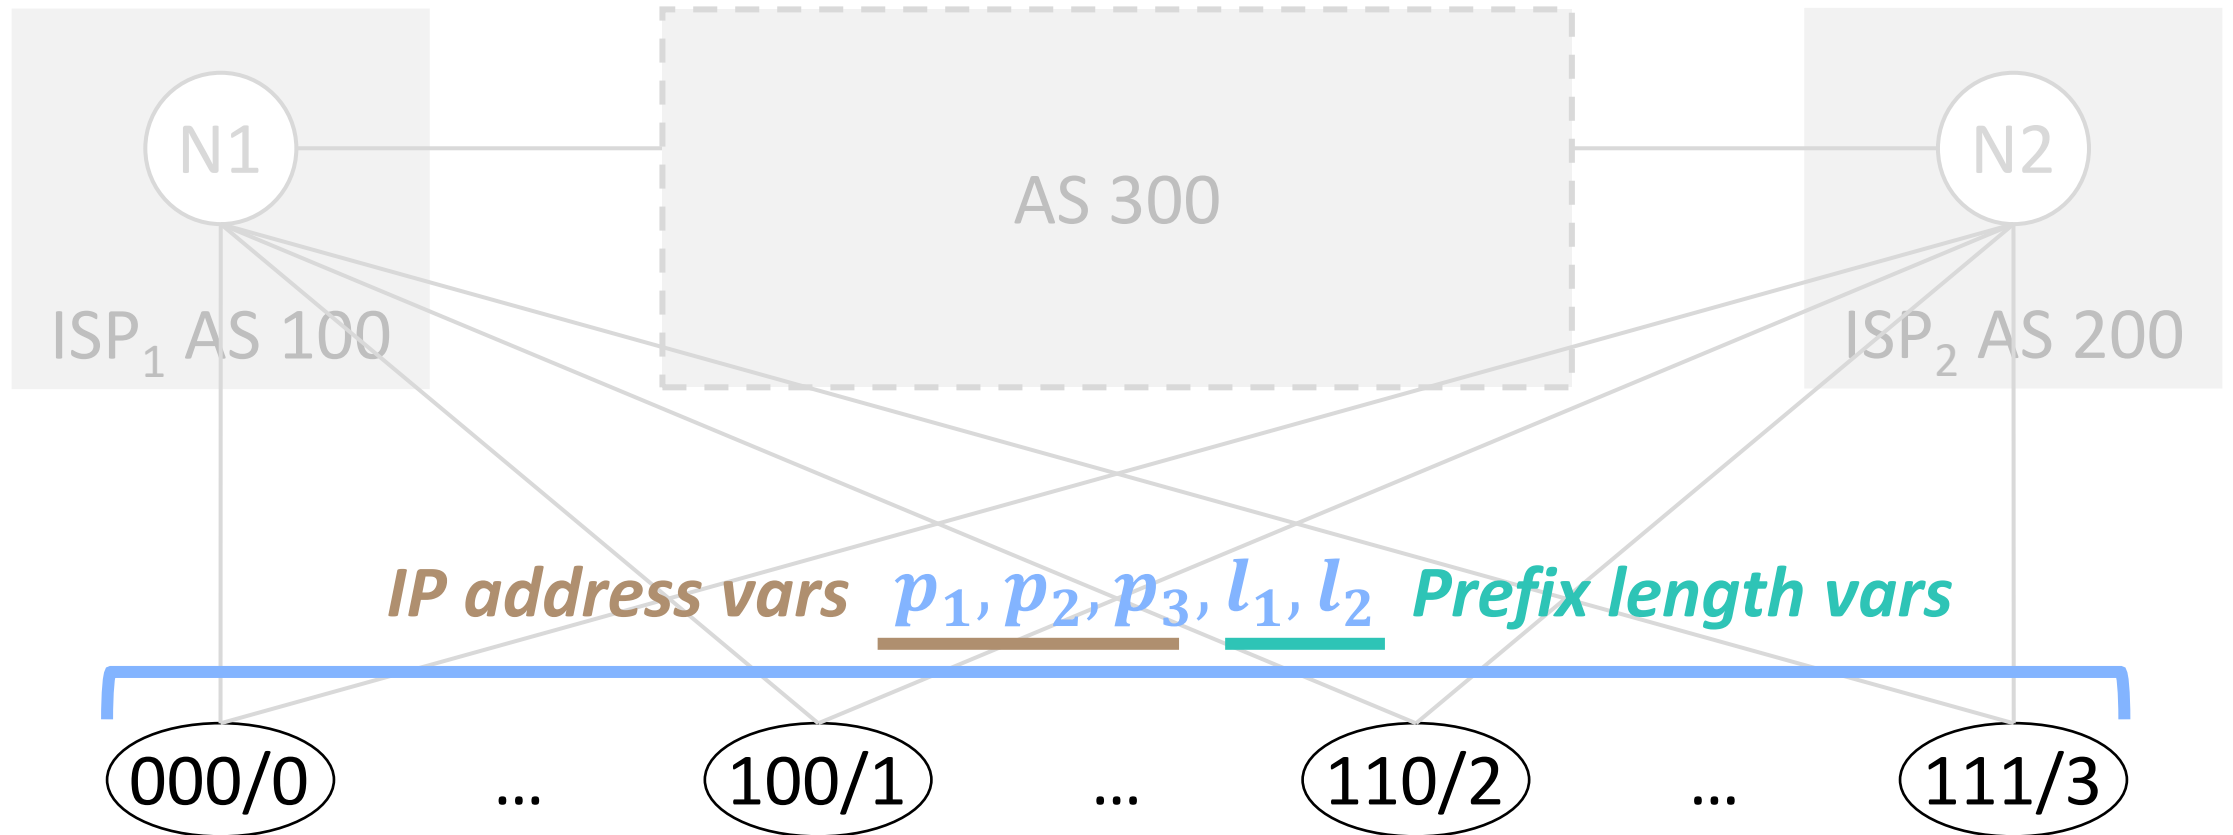
Observation2: Independency in Route Computation

*Prefixes mostly
Independent¹*




¹ Network features like route aggregation can cause dependencies, please refer to our paper for more details.

Espresso Encoding: First Select Prefixes

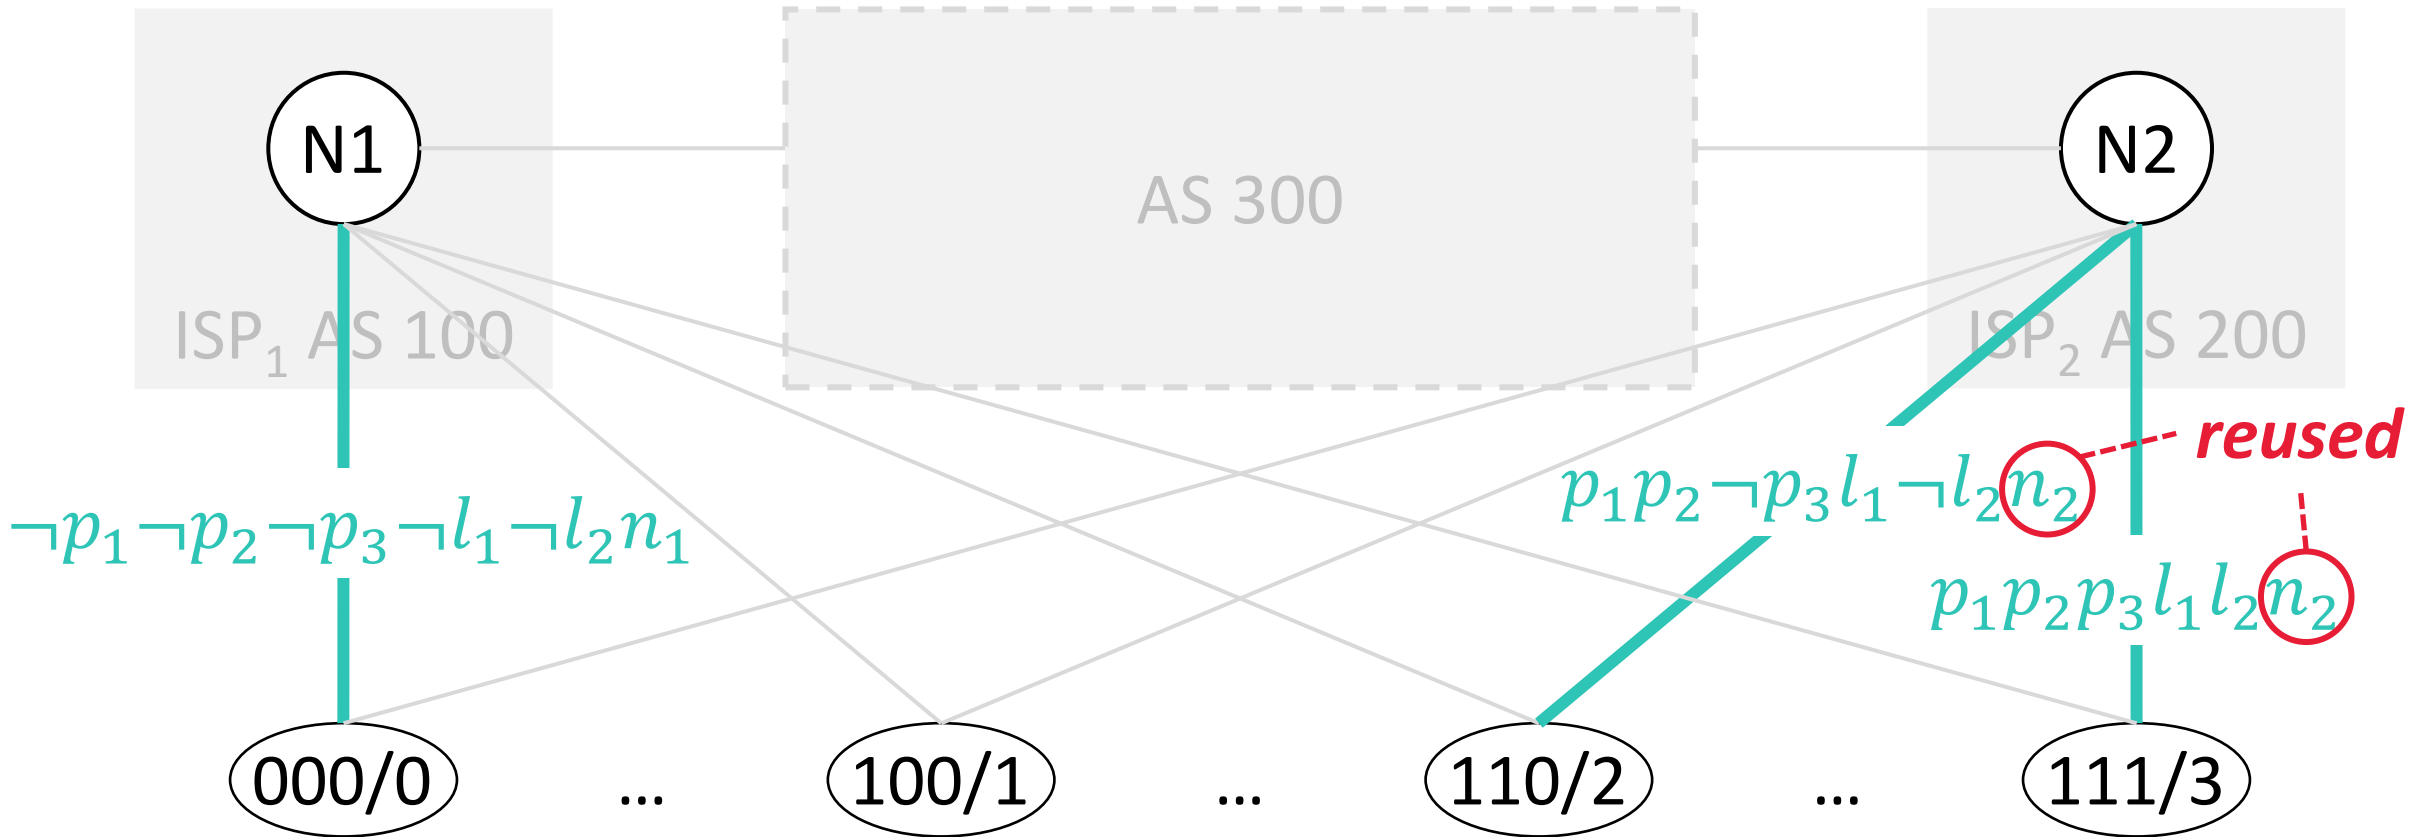


Espresso Encoding: Then Select Advertisers

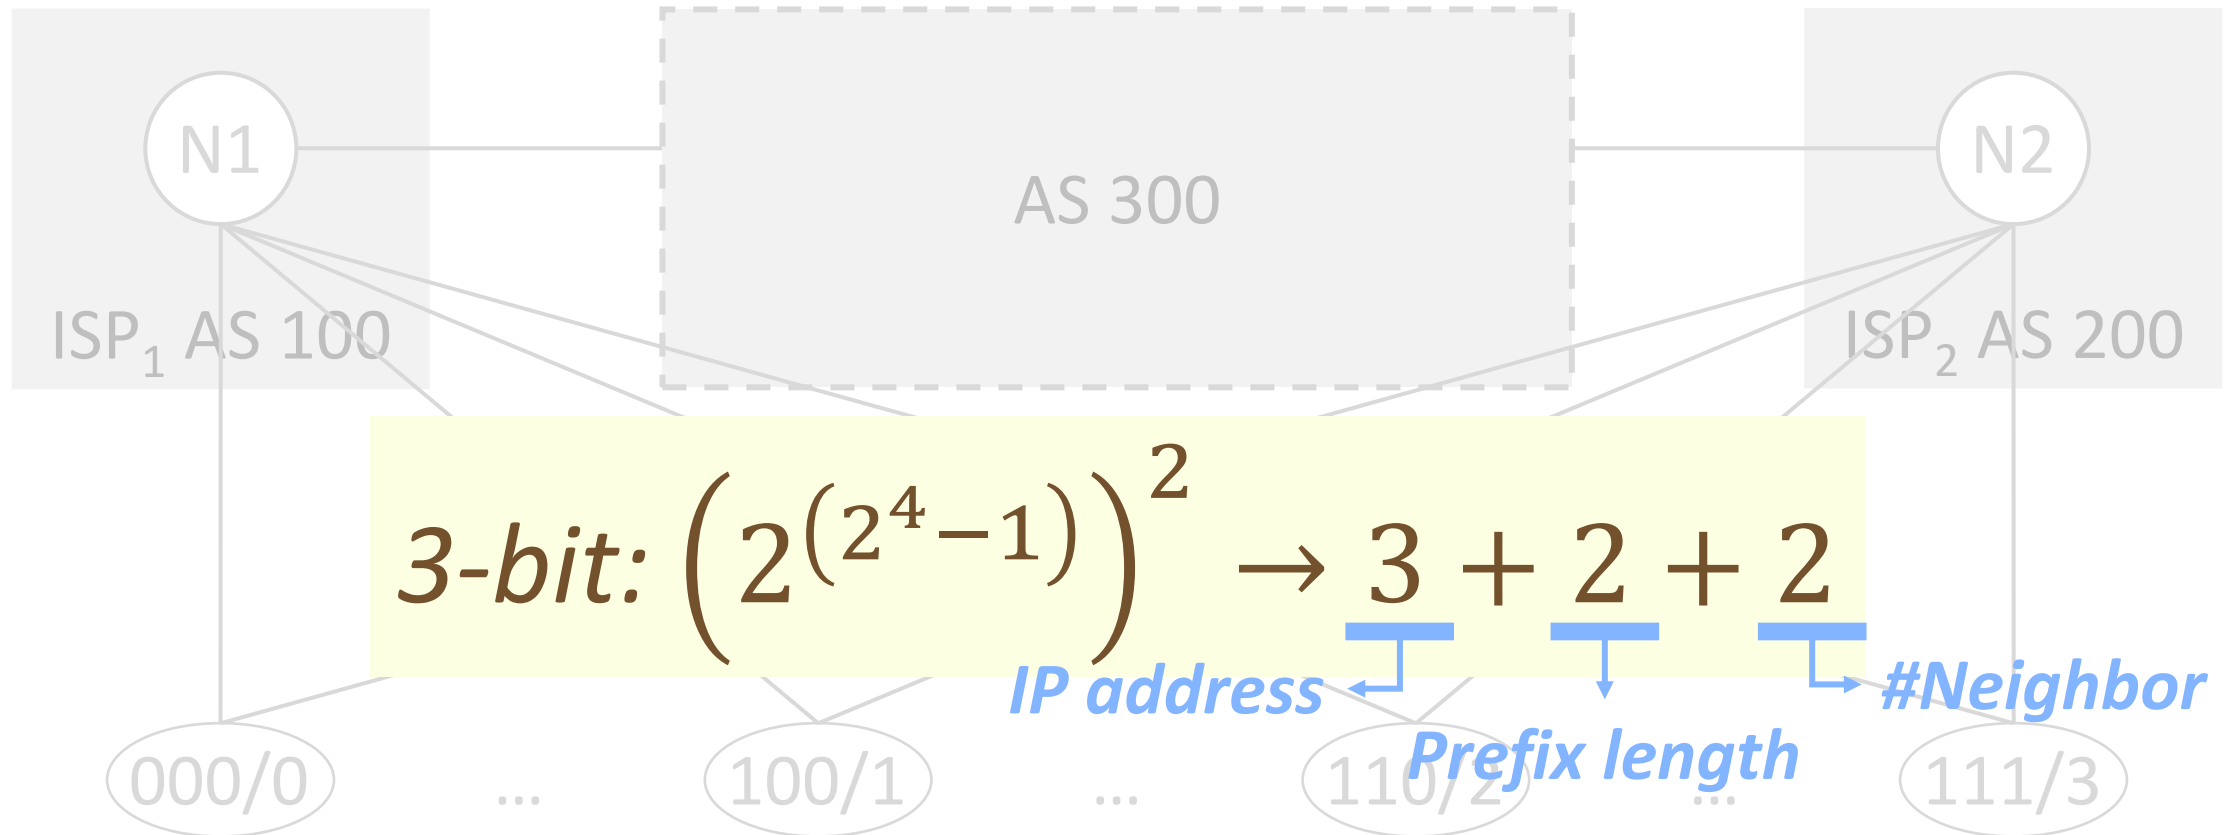


Example

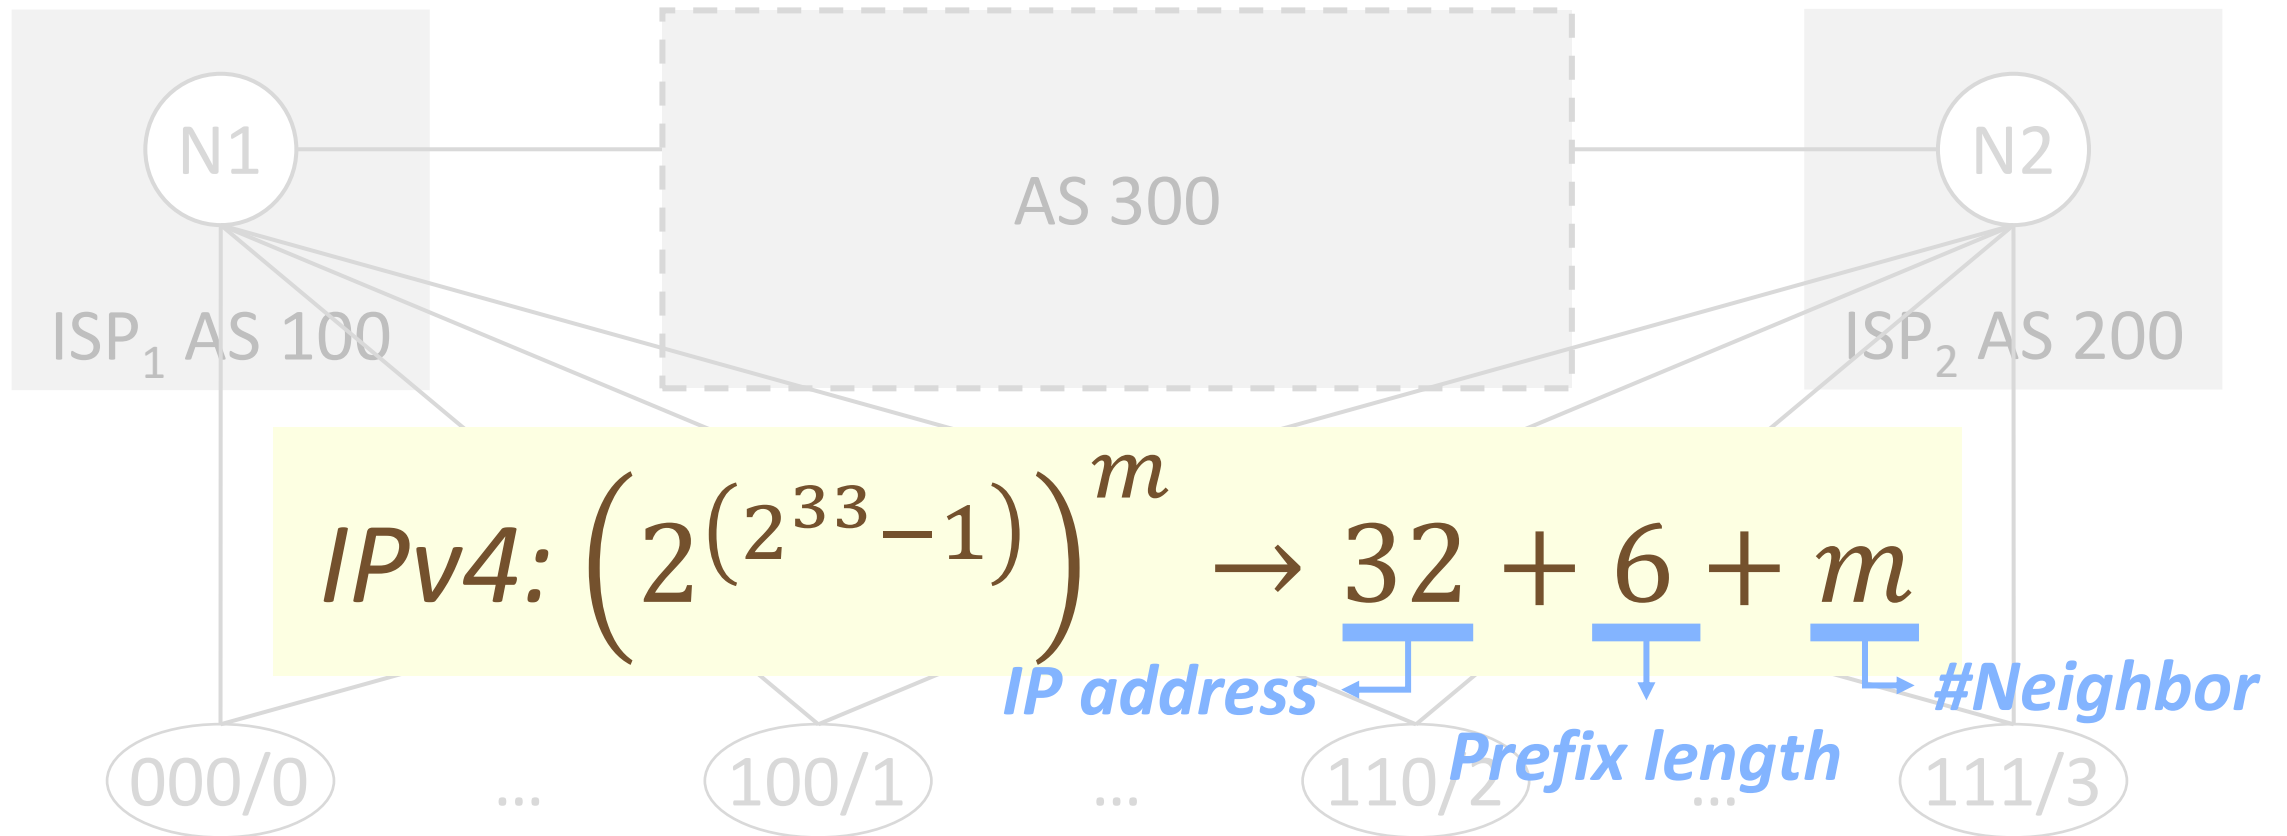
$N1$ announces $000/0$, $N2$ announces $110/2$ and $111/3$



Espresso Encoding




Espresso Encoding

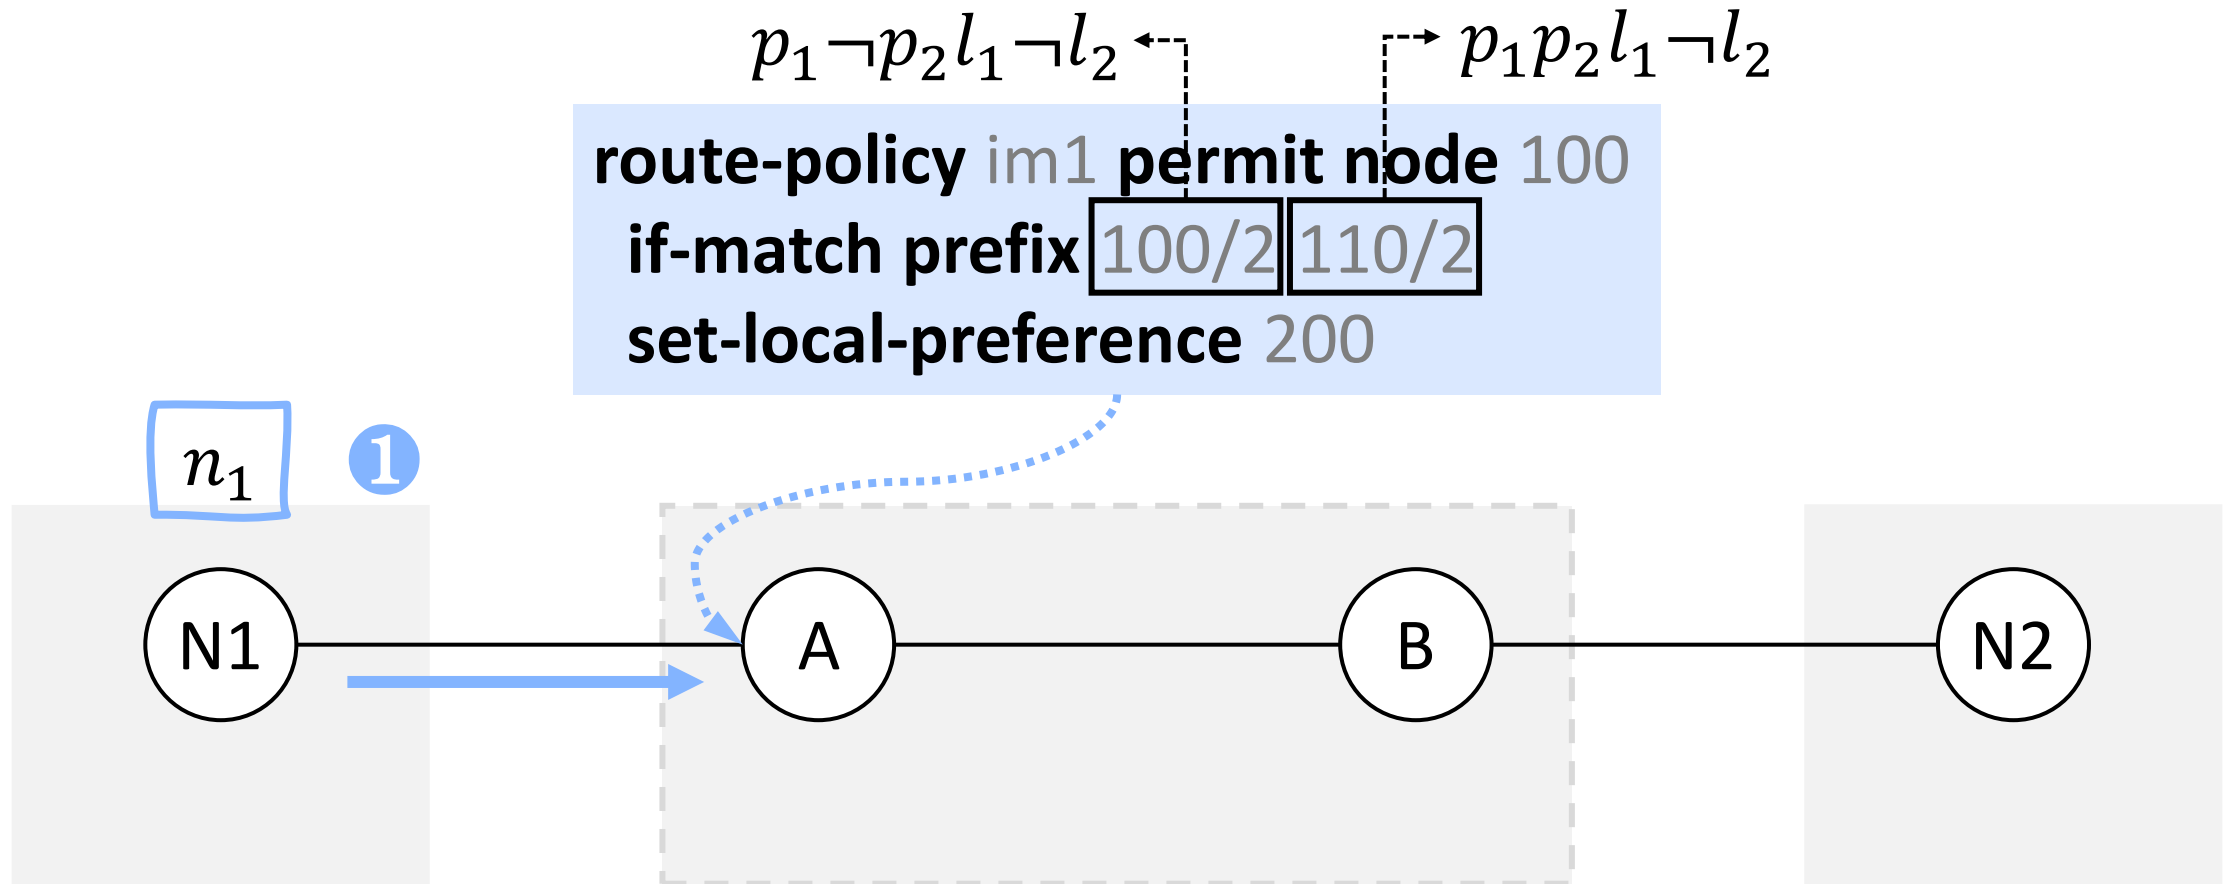


Expresso SRC: Initialize Symbolic Route

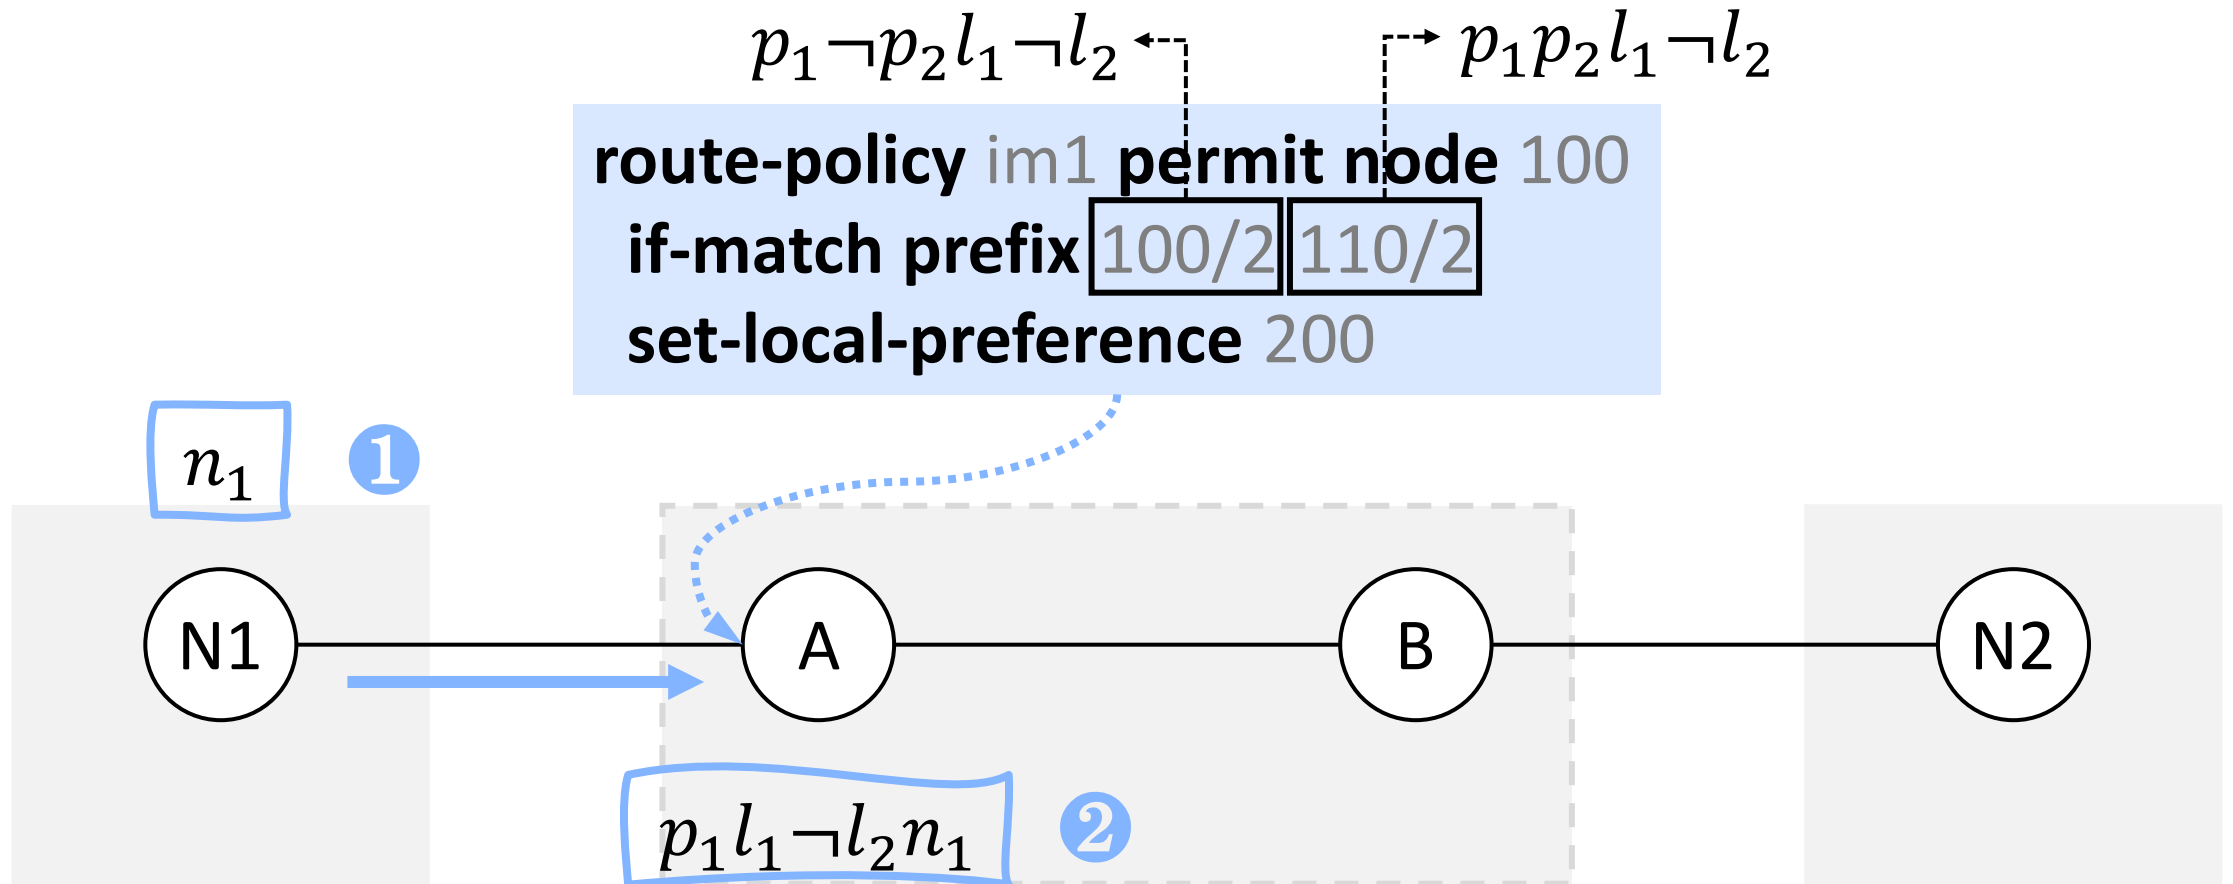
*Arbitrary prefixes
announced by N1*



Espresso SRC: Route Propagation

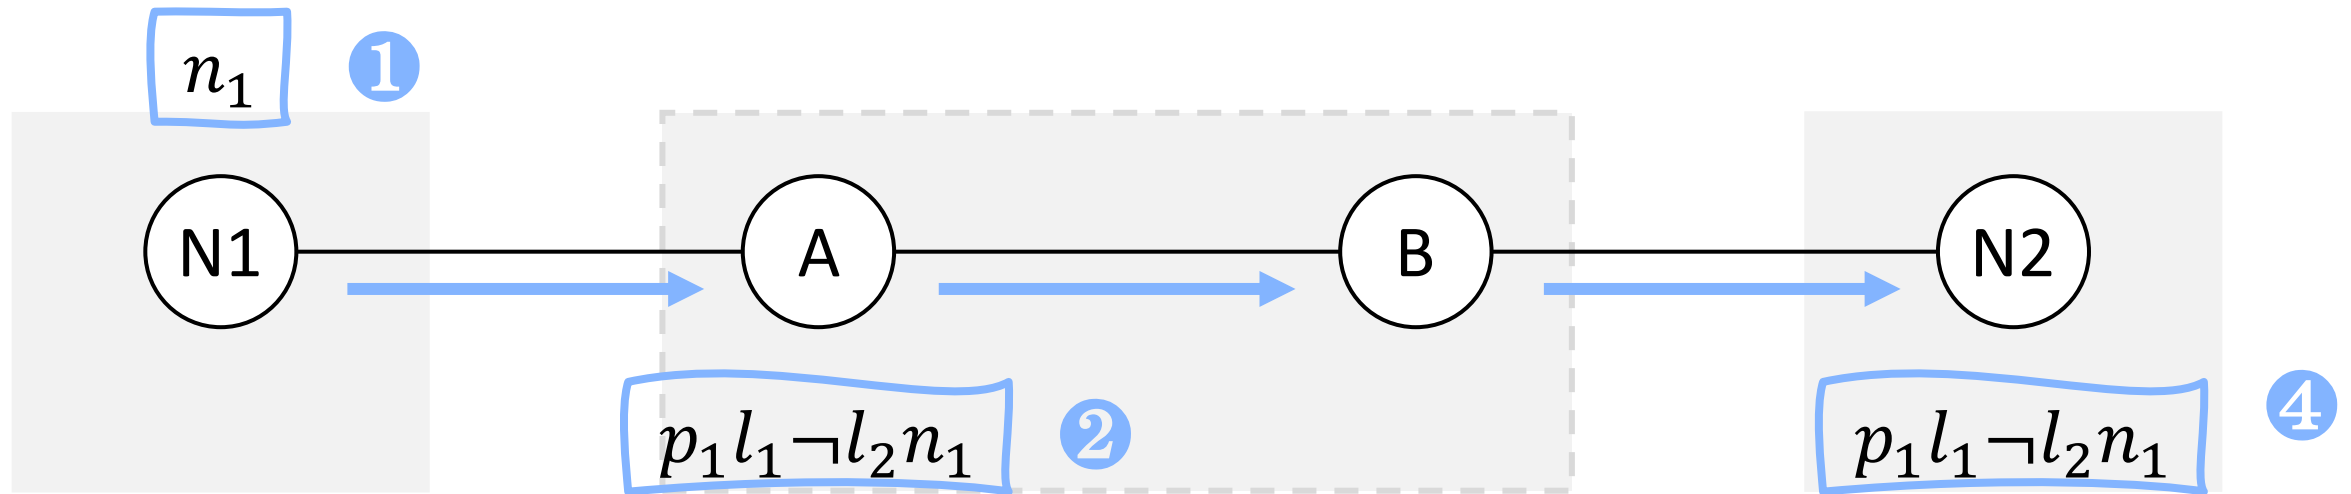


Espresso SRC: Apply Routing Policy

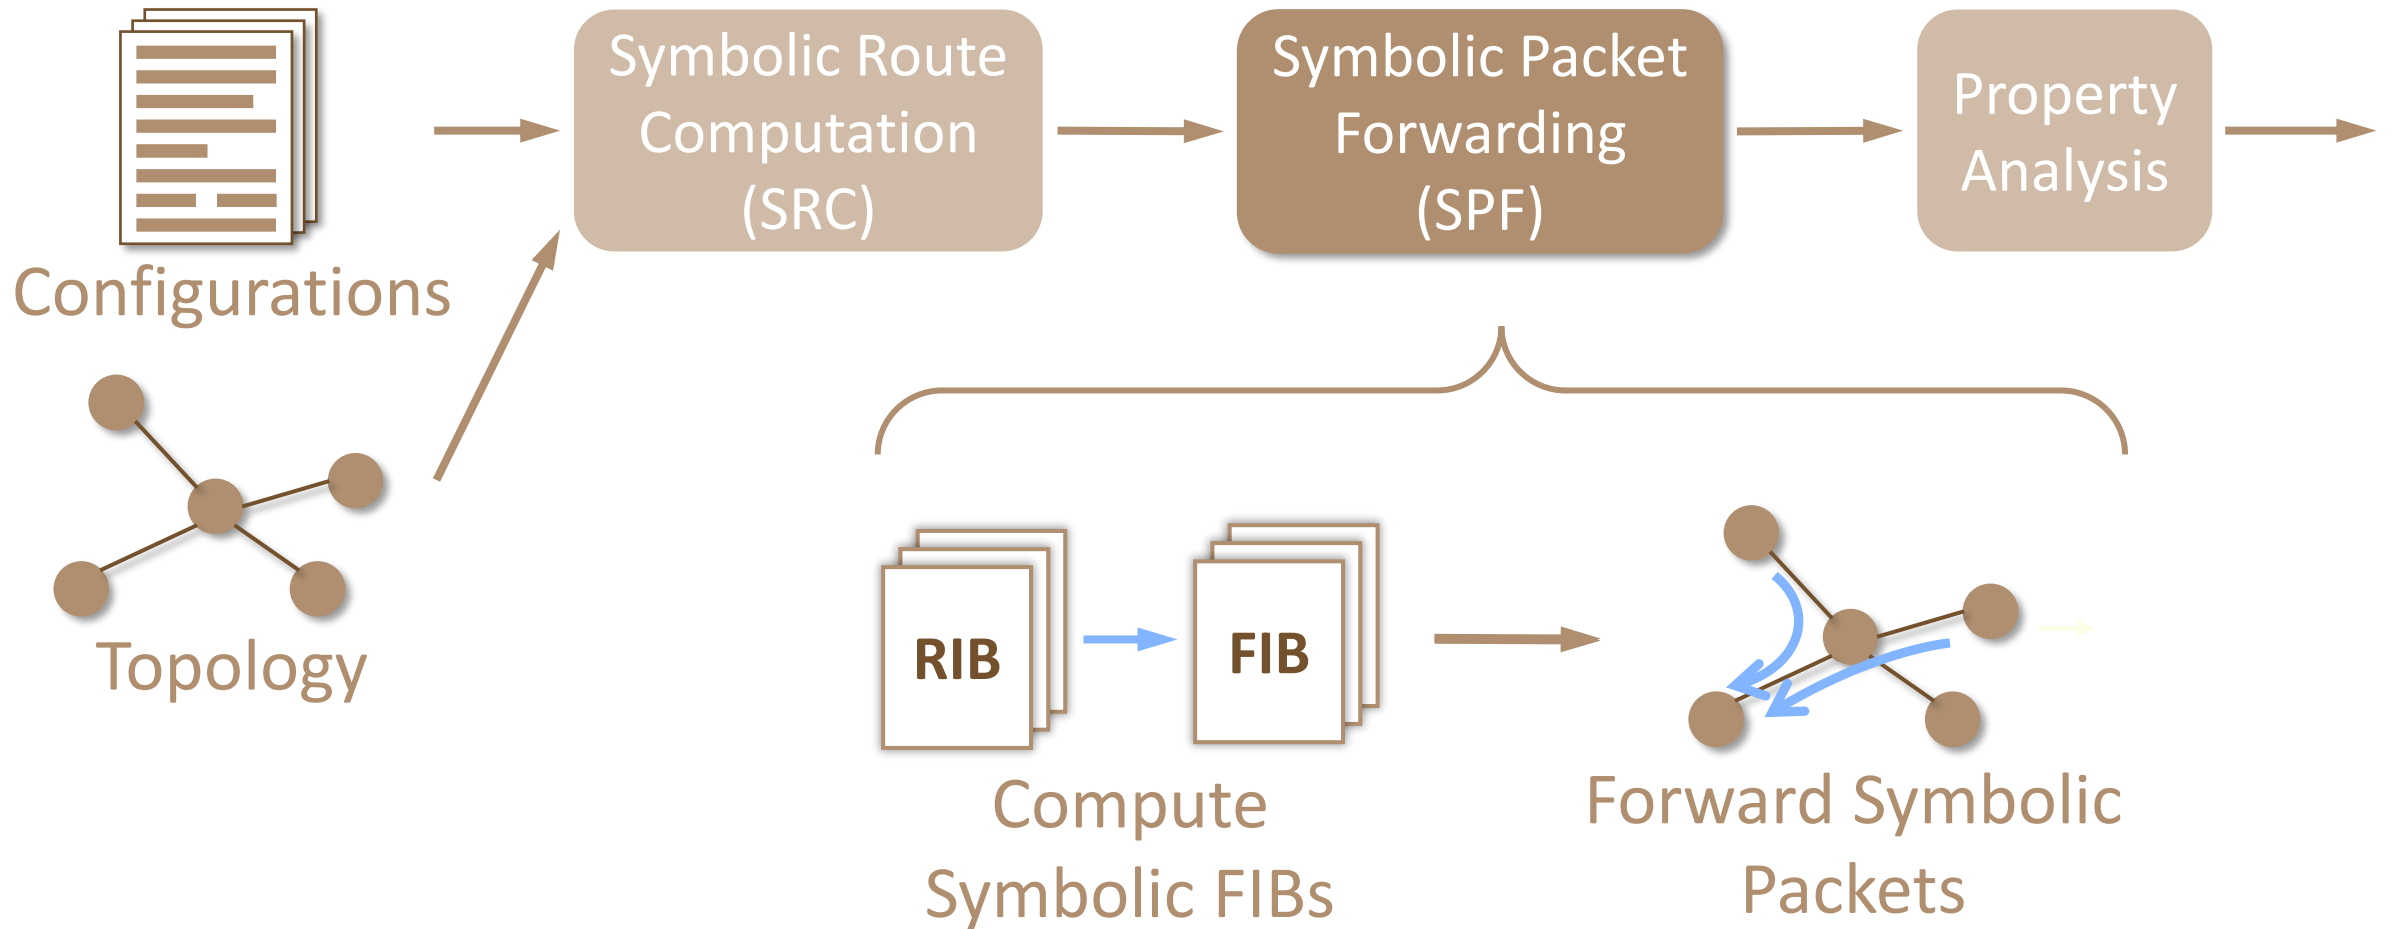


$$n_1 \wedge (p_1 \neg p_2 l_1 \neg l_2 \vee p_1 p_2 l_1 \neg l_2)$$

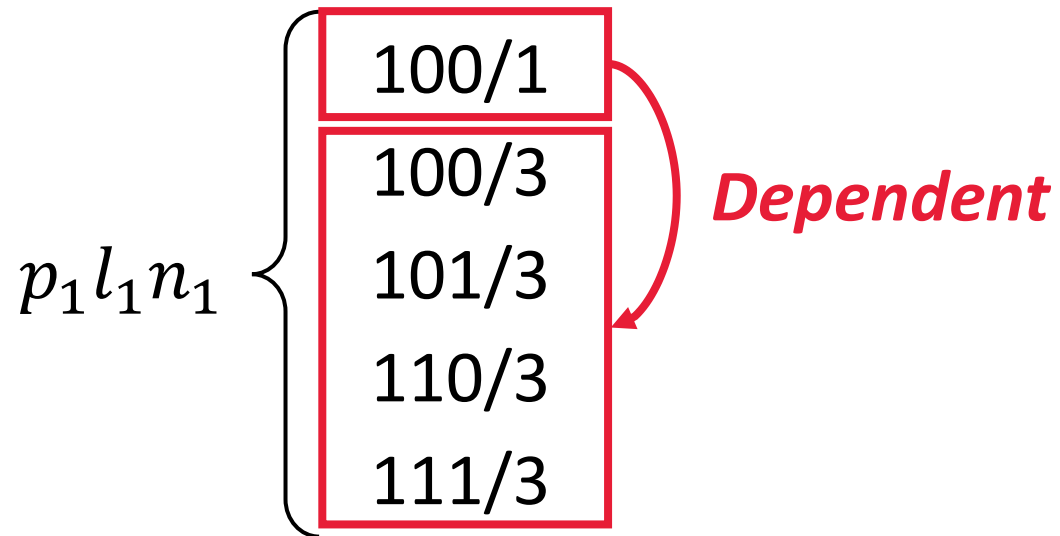
Expresso SRC: Route Propagation



Symbolic Packet Forwarding





Prefix Dependency in Packet Forwarding Due to Longest Prefix Match




1 variable per neighbor is **not enough any more.**

15 variables per neighbor now?

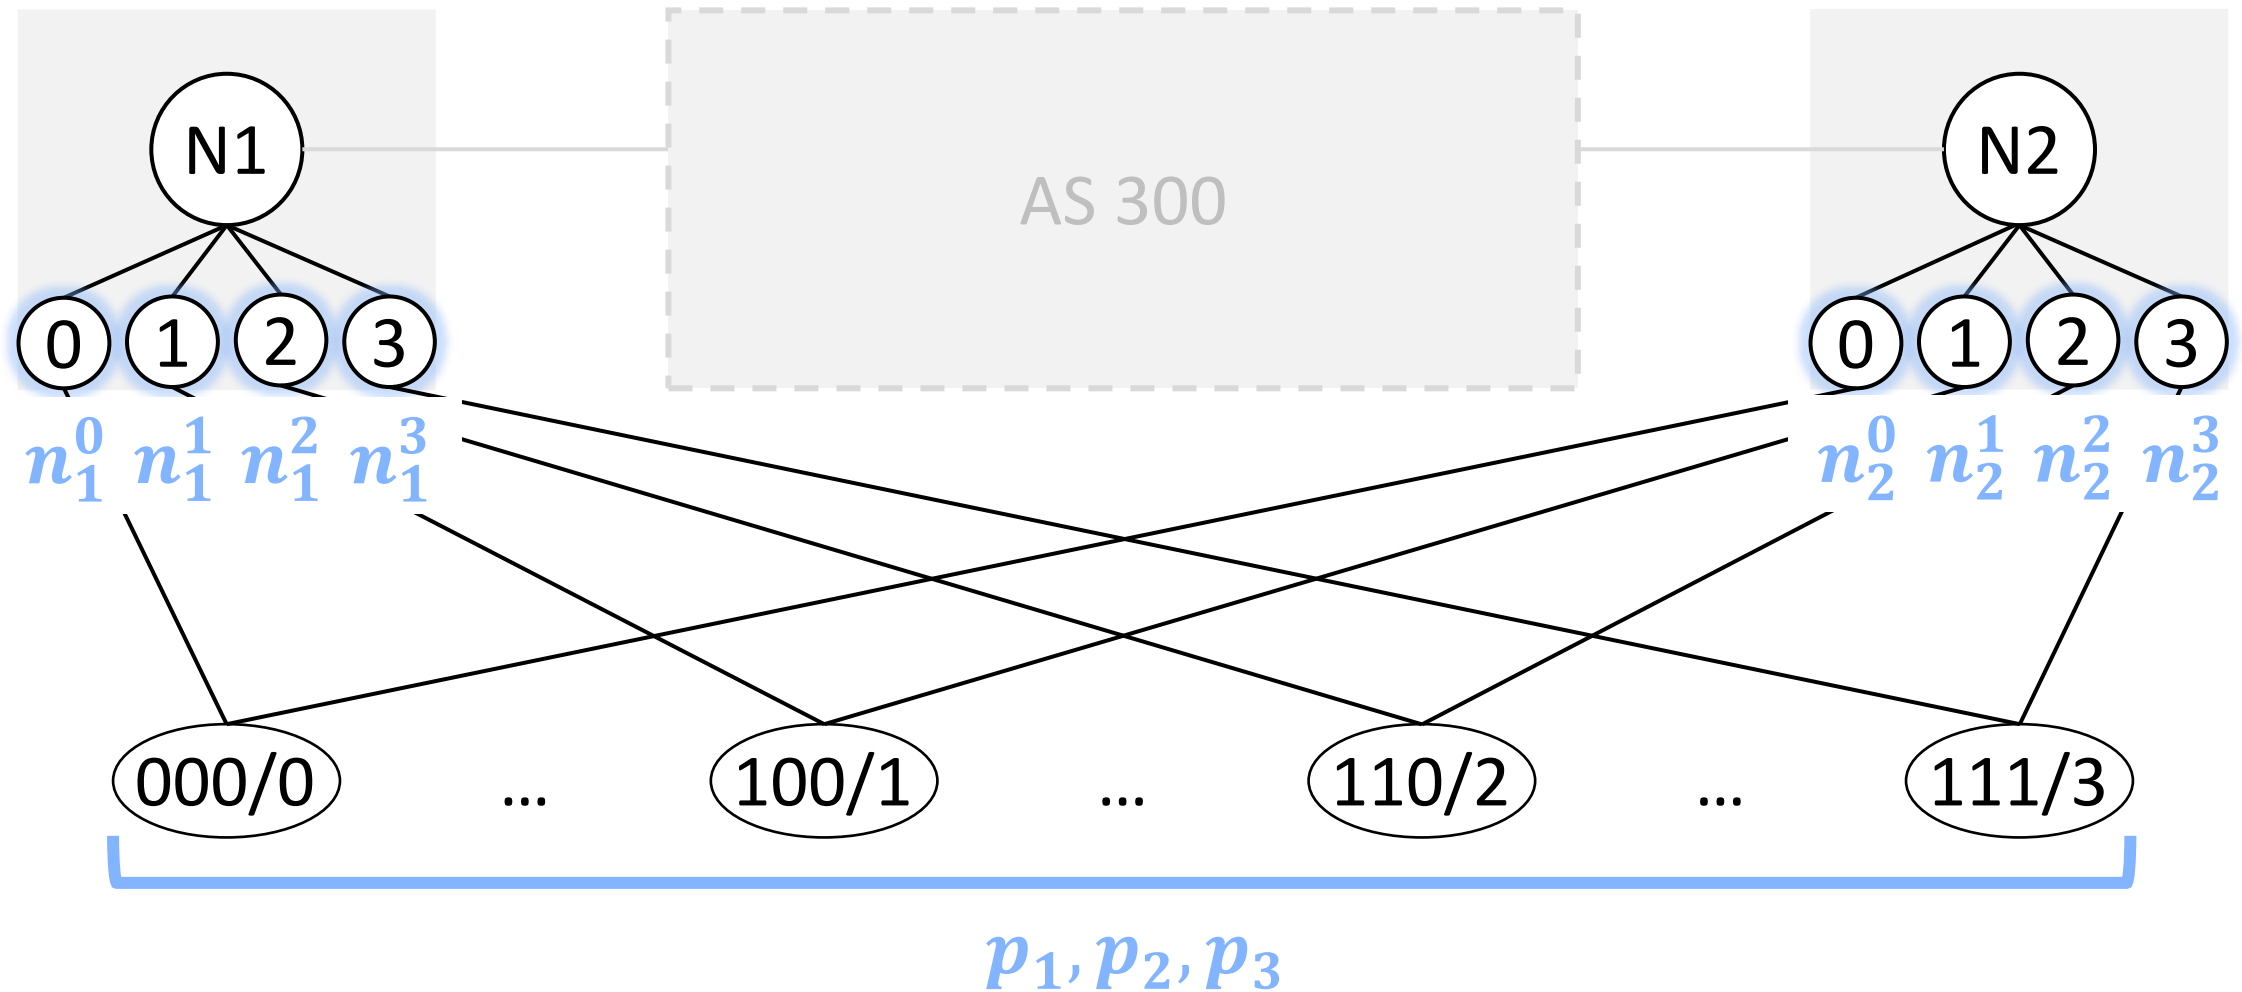
Observation 3: Dependency is Limited to Different Prefix Length




Espresso Data Plane Encoding



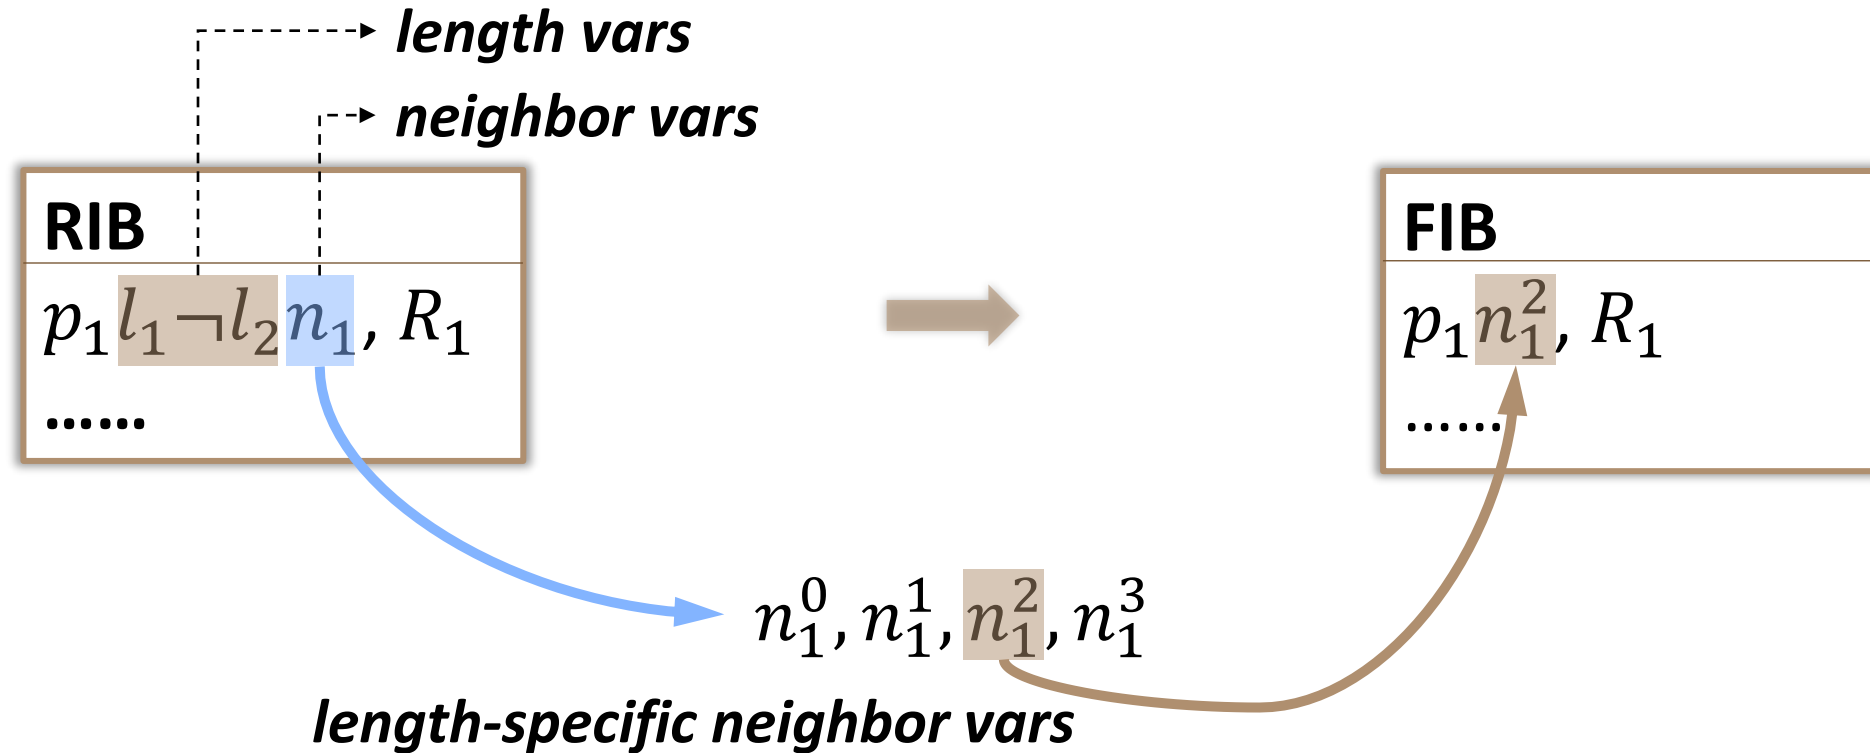
Espresso Data Plane Encoding



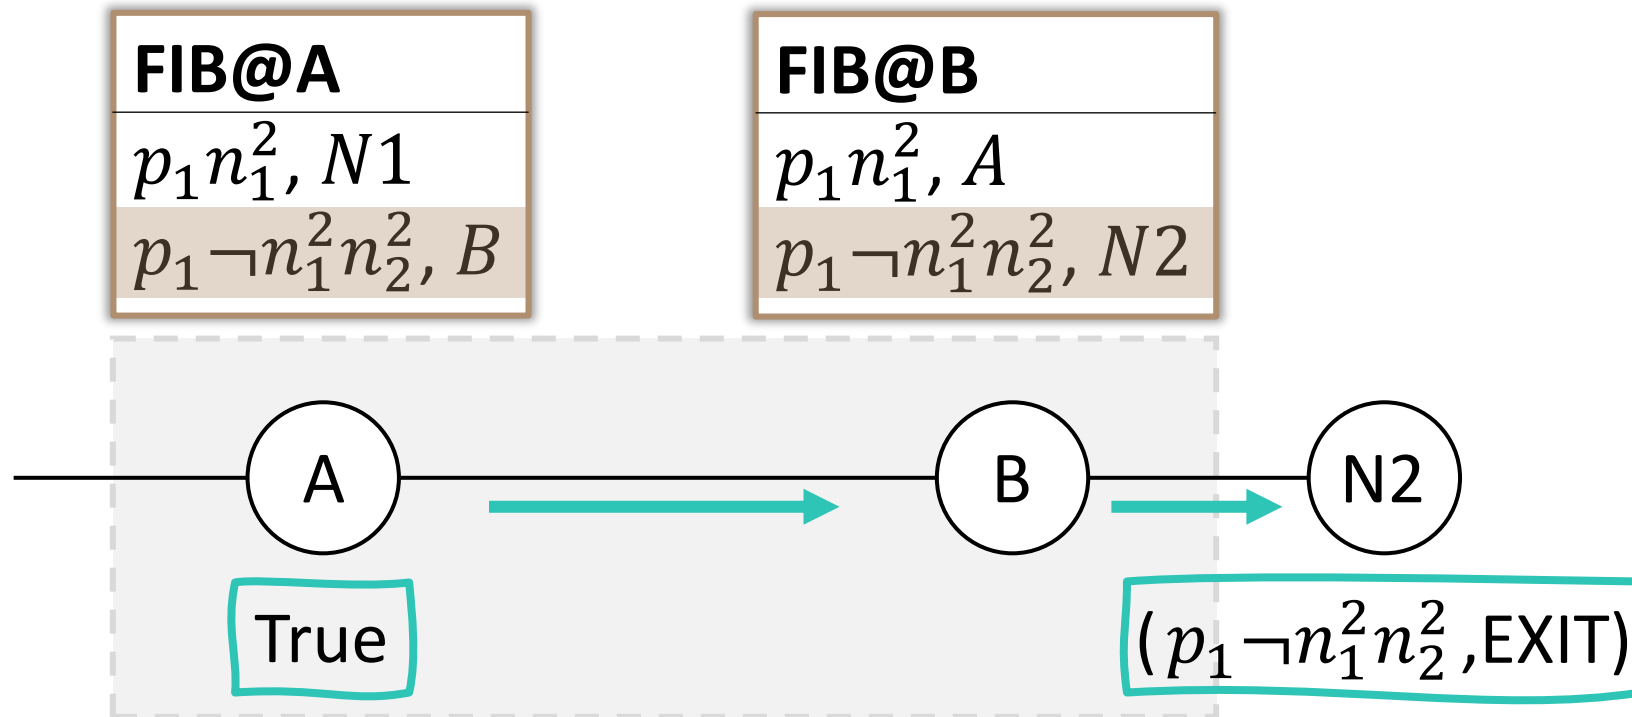
Expresso Compute Symbolic FIBs



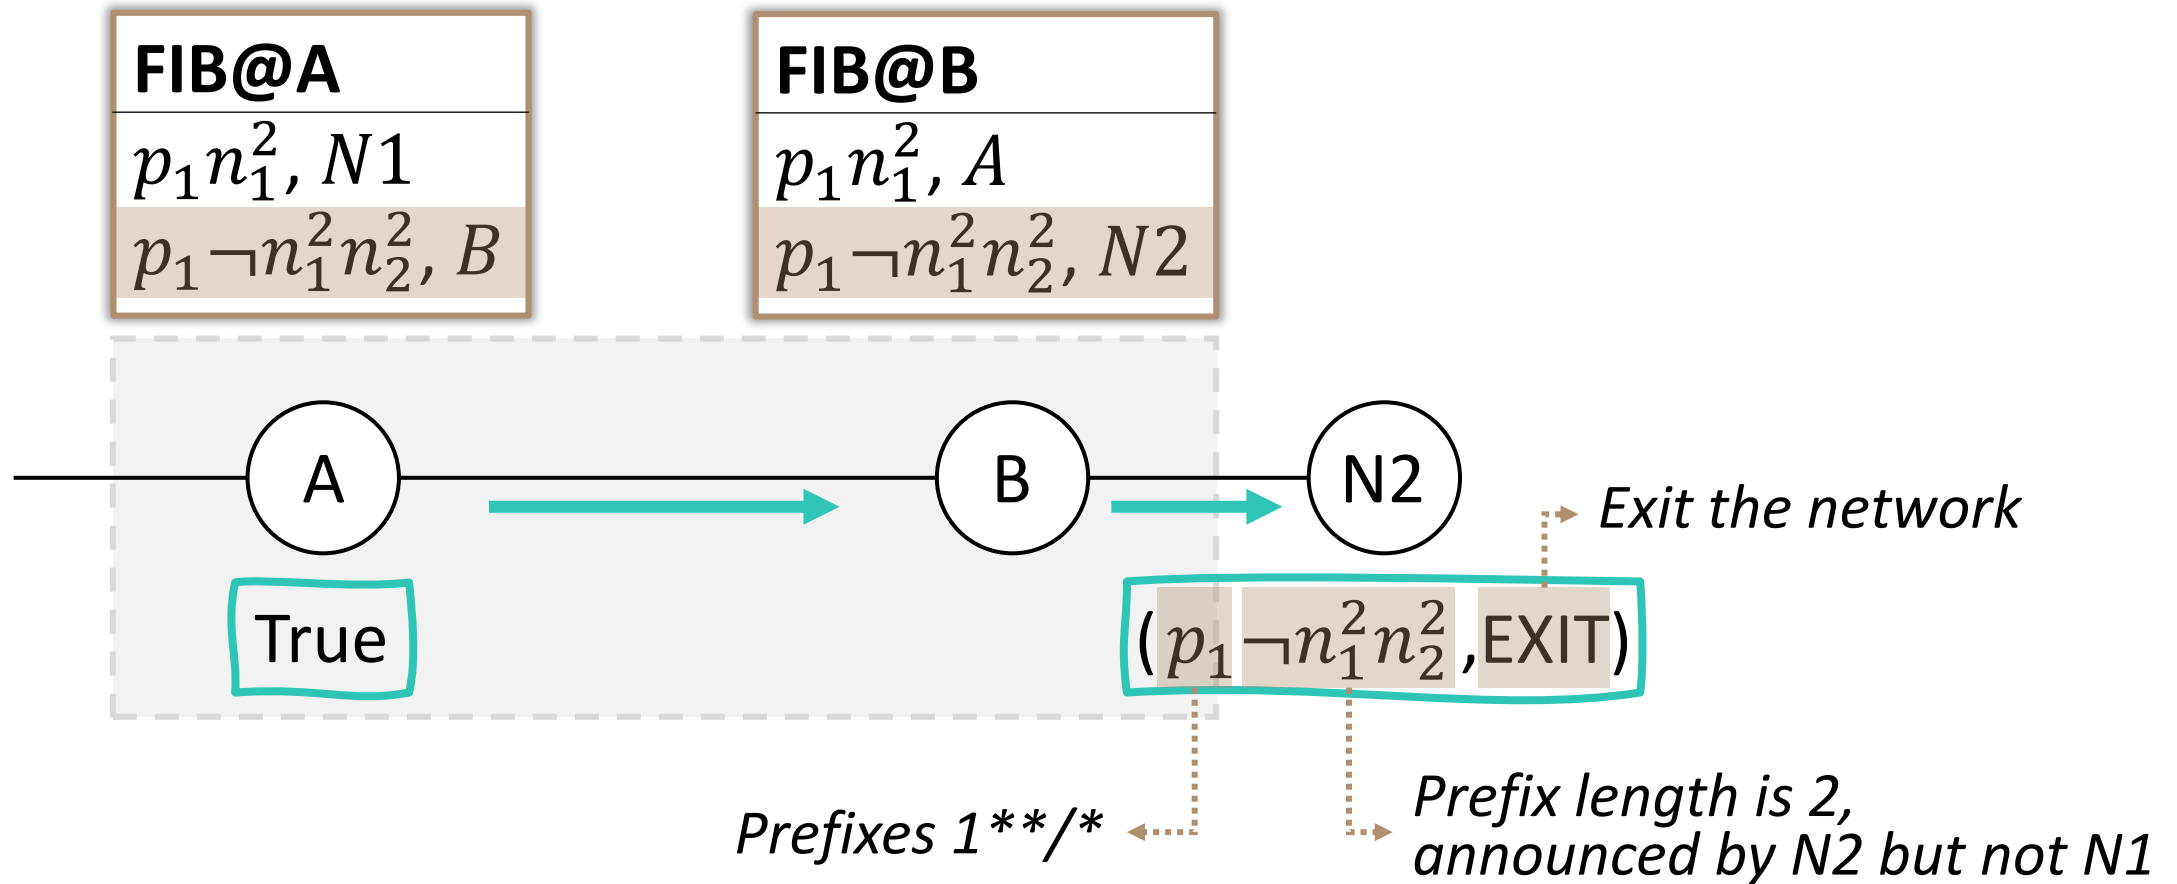
Expresso Compute Symbolic FIBs




Expresso compute PECs




Expresso compute PECs




Property Analysis




Route Leak




Route Hijack



Traffic Hijack



Egress Preference

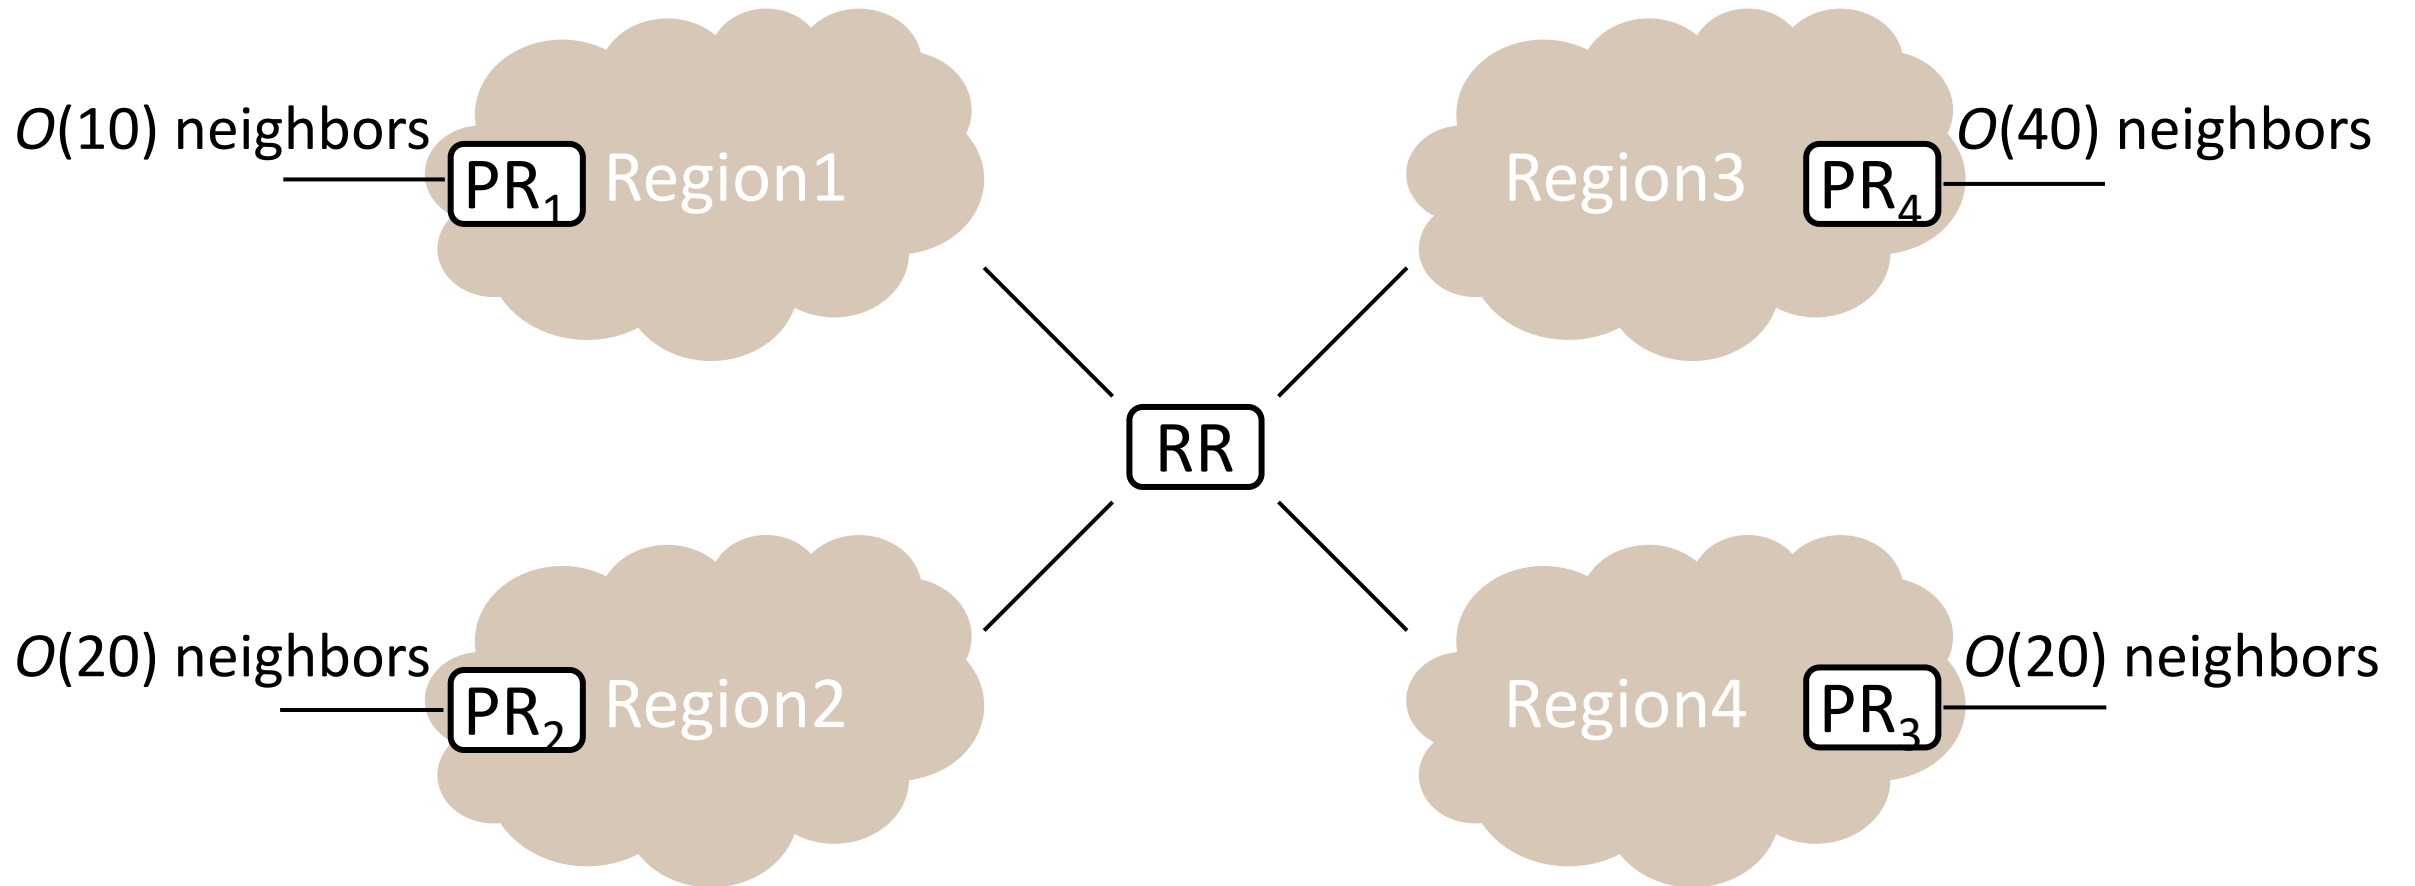


Block To External

Evaluation Results

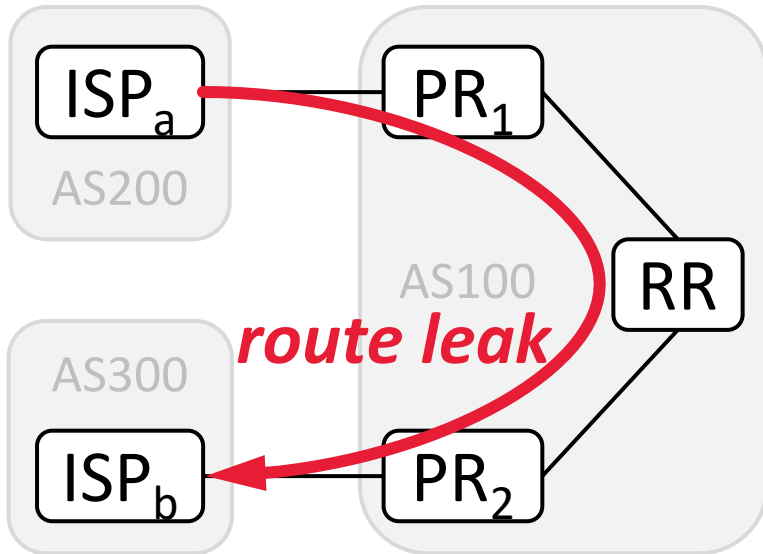
On the WAN of a Cloud Service Provider (CSP)

A Cloud Service Provider's WAN

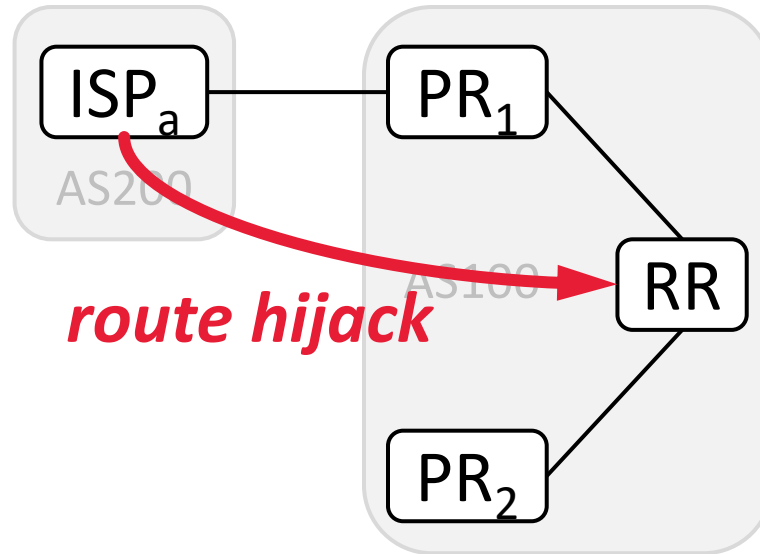


PR: Peering Router, RR: Route Reflector

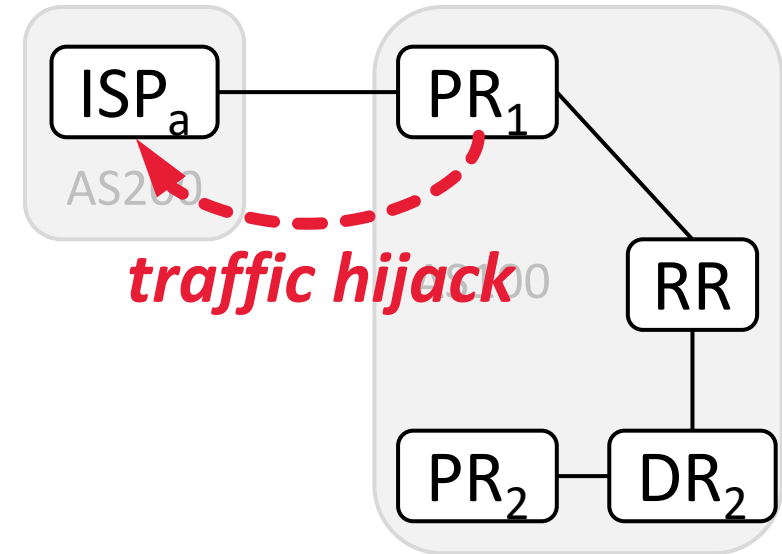
Found Violations



6 confirmed

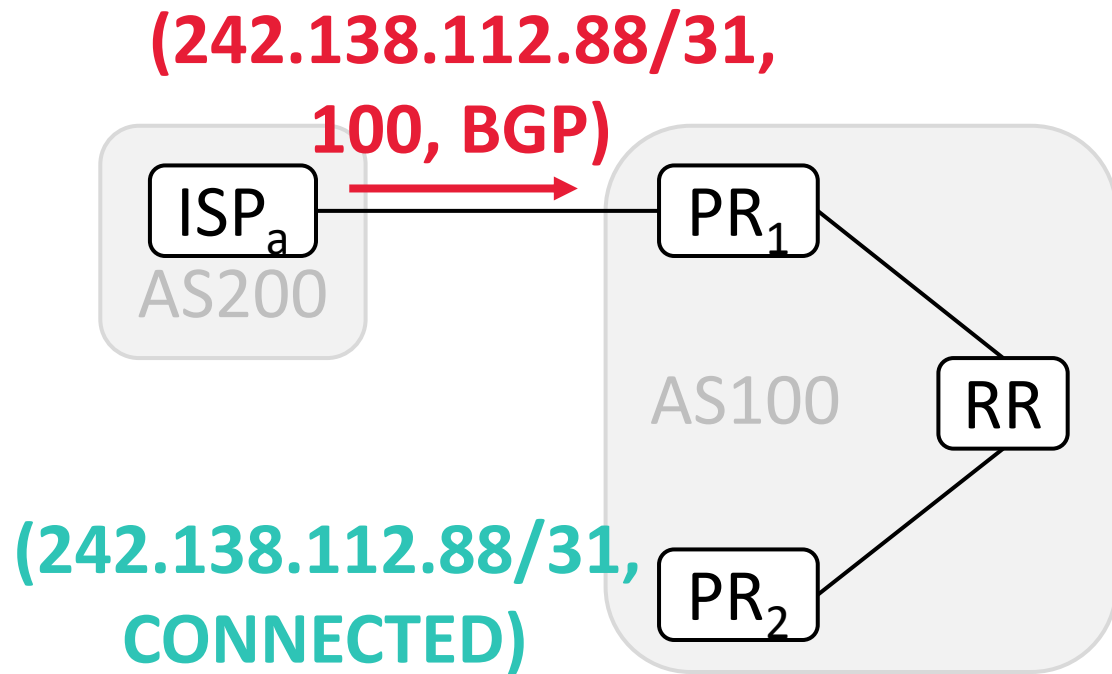


17 confirmed



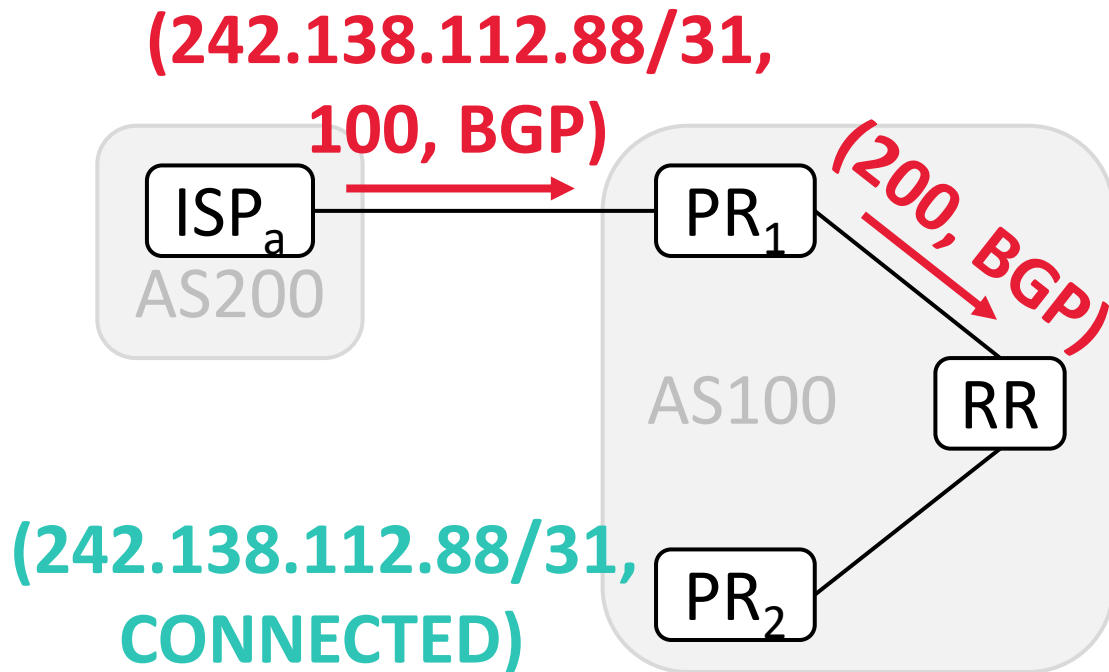
19 confirmed

Case Study (A Route Hijack)



PR: Peering Router, RR: Route Reflector

Case Study (A Route Hijack)

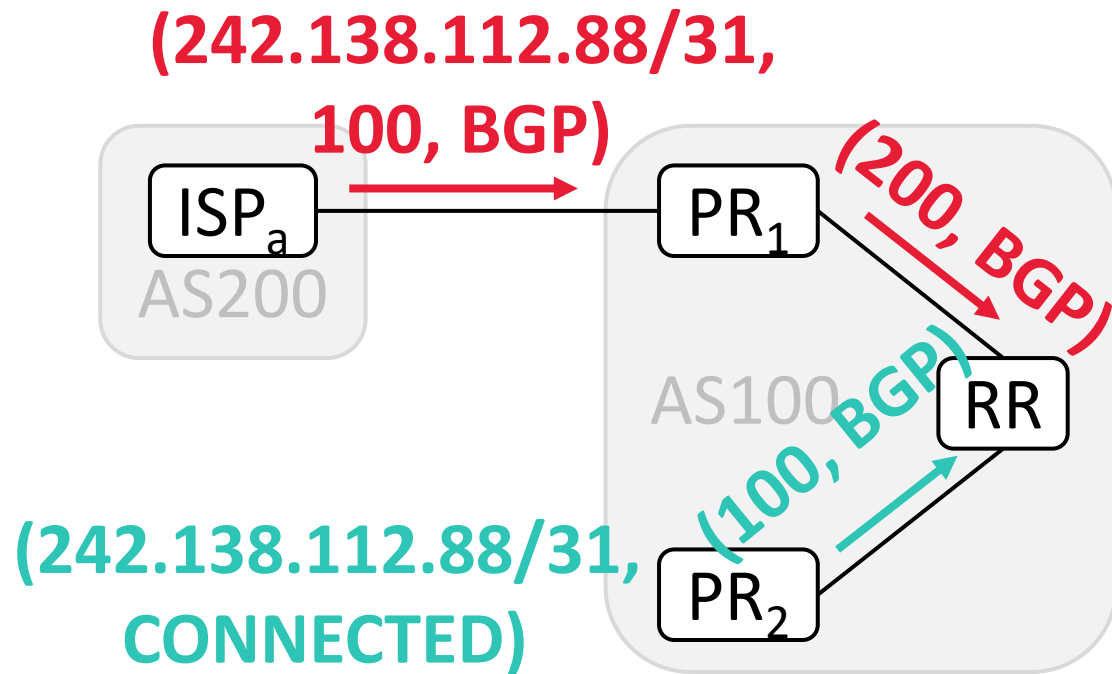


```
route-policy import node 100
  if-match prefix-list list1
  set-local-preference 200
```

```
ip prefix list1 node 100 deny 114.247.96.0 20 le 32
ip prefix list1 node 200 deny 123.29.0.0 20 le 32
ip prefix list1 node 300 deny 137.155.0.0 18 le 32
missing ip prefix list1 node 400 deny 92.230.128.0 18 le 32
ip prefix list1 node xxx deny 242.138.120.0 20 le 32
ip prefix list1 node 500 permit 0.0.0.0 0 le 32
```

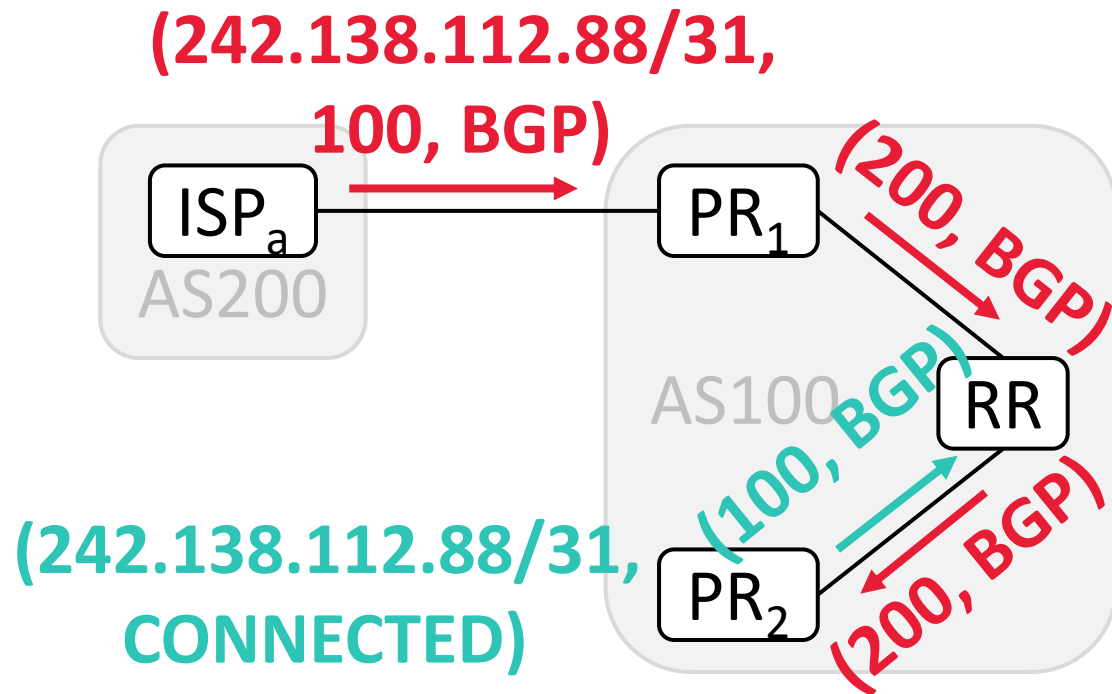
PR: Peering Router, RR: Route Reflector

Case Study (A Route Hijack)




PR: Peering Router, RR: Route Reflector

Case Study (A Route Hijack)




PR: Peering Router, RR: Route Reflector




Case Study (A Route Hijack)



PR: Peering Router, RR: Route Reflector

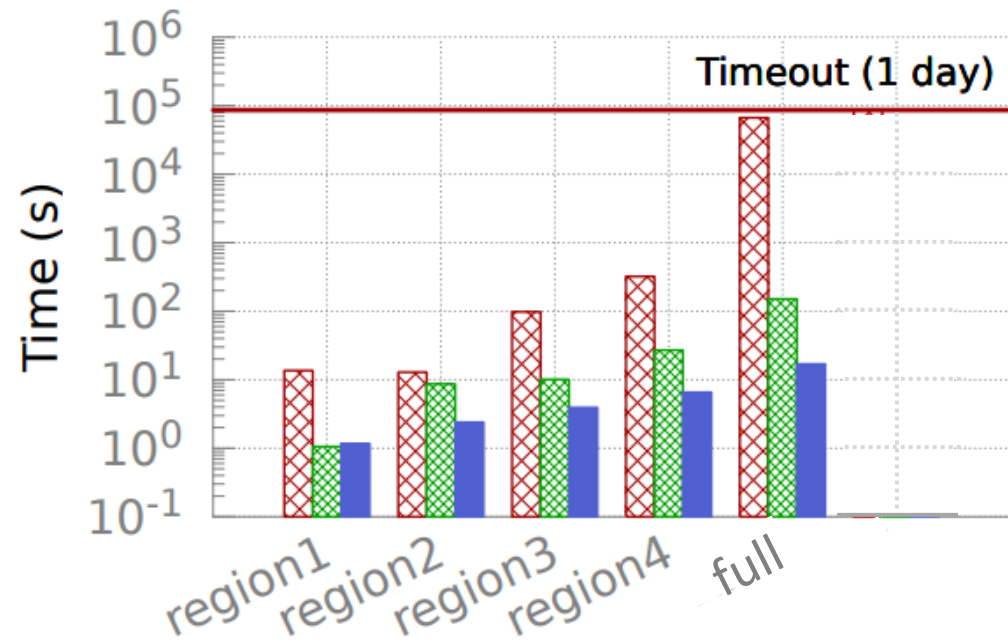
Performance



-  Minesweeper*
-  Espresso
-  Espresso-

Minesweeper ≈ 1 day
Espresso < 1 h

Performance



- Minesweeper*
- Espresso
- Espresso-

Espresso is 1-4 orders of magnitude faster

Summary

- *External route caused incidents* are common and costly.
- *External route environment space is colossal* and *existing verifiers* cannot explore it efficiently.
- Espresso uses *equivalences* and *independencies* to efficiently explore the space of external routes through *symbolic simulation*.
- Espresso finds *real violations* on a CSP's WAN.

Espresso: Comprehensively Reasoning About External Routes Using Symbolic Simulation



Dan Wang

Xi'an Jiaotong University



Peng Zhang



Aaron Gember-Jacobson

Colgate University

Reference

- [1] Fogel, Ari, et al. "A general approach to network configuration analysis." *12th USENIX Symposium on Networked Systems Design and Implementation (NSDI 15)*. 2015.
- [2] Weitz, Konstantin, et al. "Scalable verification of border gateway protocol configurations with an SMT solver." *Proceedings of the 2016 acm sigplan international conference on object-oriented programming, systems, languages, and applications*. 2016.
- [3] Beckett, Ryan, et al. "A general approach to network configuration verification." *Proceedings of the Conference of the ACM Special Interest Group on Data Communication*. 2017.
- [4] Beckett, Ryan, et al. "Abstract interpretation of distributed network control planes." *Proceedings of the ACM on Programming Languages* 4.POPL (2019): 1-27.
- [5] Giannarakis, Nick, et al. "NV: An intermediate language for verification of network control planes." *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation*. 2020.
- [6] Ye, Fangdan, et al. "Accuracy, scalability, coverage: A practical configuration verifier on a global wan." *Proceedings of the Annual conference of the ACM Special Interest Group on Data Communication on the applications, technologies, architectures, and protocols for computer communication*. 2020.

Reference

- [7] Zhang, Peng, Dan Wang, and Aaron Gember-Jacobson. "Symbolic router execution." *Proceedings of the ACM SIGCOMM 2022 Conference*. 2022.
- [8] Zhang, Peng, et al. "Differential network analysis." *19th USENIX Symposium on Networked Systems Design and Implementation (NSDI 22)*. 2022.
- [9] Abhashkumar, Anubhavnidhi, Aaron Gember-Jacobson, and Aditya Akella. "Tiramisu: Fast multilayer network verification." *17th USENIX Symposium on Networked Systems Design and Implementation (NSDI 20)*. 2020.
- [10] Prabhu, Santhosh, et al. "Plankton: Scalable network configuration verification through model checking." *17th USENIX Symposium on Networked Systems Design and Implementation (NSDI 20)*. 2020.
- [11] Fayaz, Seyed K., et al. "Efficient network reachability analysis using a succinct control plane representation." *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*. 2016.
- [12] Weitz, Konstantin, et al. "Scalable verification of border gateway protocol configurations with an SMT solver." *Proceedings of the 2016 acm sigplan international conference on object-oriented programming, systems, languages, and applications*. 2016.
- [13] Tian, Bingchuan, et al. "Safely and automatically updating in-network acl configurations with intent language." *Proceedings of the ACM Special Interest Group on Data Communication*. 2019. 214-226.