

Verifiable Data Plane

Thesis Mid-Defense

PhD Student

Julian Martin Del Fiore

Committee

Chadi Barakat


Olivier Bonaventure

Pascal Merindol


Cristel Pelsser




Why Verifiable Data Plane?




Why Verifiable Data Plane?



Why Verifiable Data Plane?



Why Verifiable Data Plane?



Are BGP AS Paths being followed?


- ✓ What generates discrepancies?
- ✓ Can we see it in the wild?

Are technical limitations common?

- ✓ How do they affect the forwarding?
- ✓ How can we detect them?

Why Verifiable Data Plane?

TMA 2019



Are BGP AS Paths being followed?


- ✓ What generates discrepancies?
- ✓ Can we see it in the wild?

Are technical limitations common?

- ✓ How do they affect the forwarding?
- ✓ How can we detect them?

Why Verifiable Data Plane?

TMA 2019



CoNext 2019...?

Are BGP AS Paths being followed?


- ✓ What generates discrepancies?
- ✓ Can we see it in the wild?

Are technical limitations common?

- ✓ How do they affect the forwarding?
- ✓ How can we detect them?

Why Verifiable Data Plane?

BGP
Journal...?



CoNext 2019...?

TMA 2019

Are BGP AS Paths being followed?

- ✓ What generates discrepancies?
- ✓ Can we see it in the wild?

Are technical limitations common?

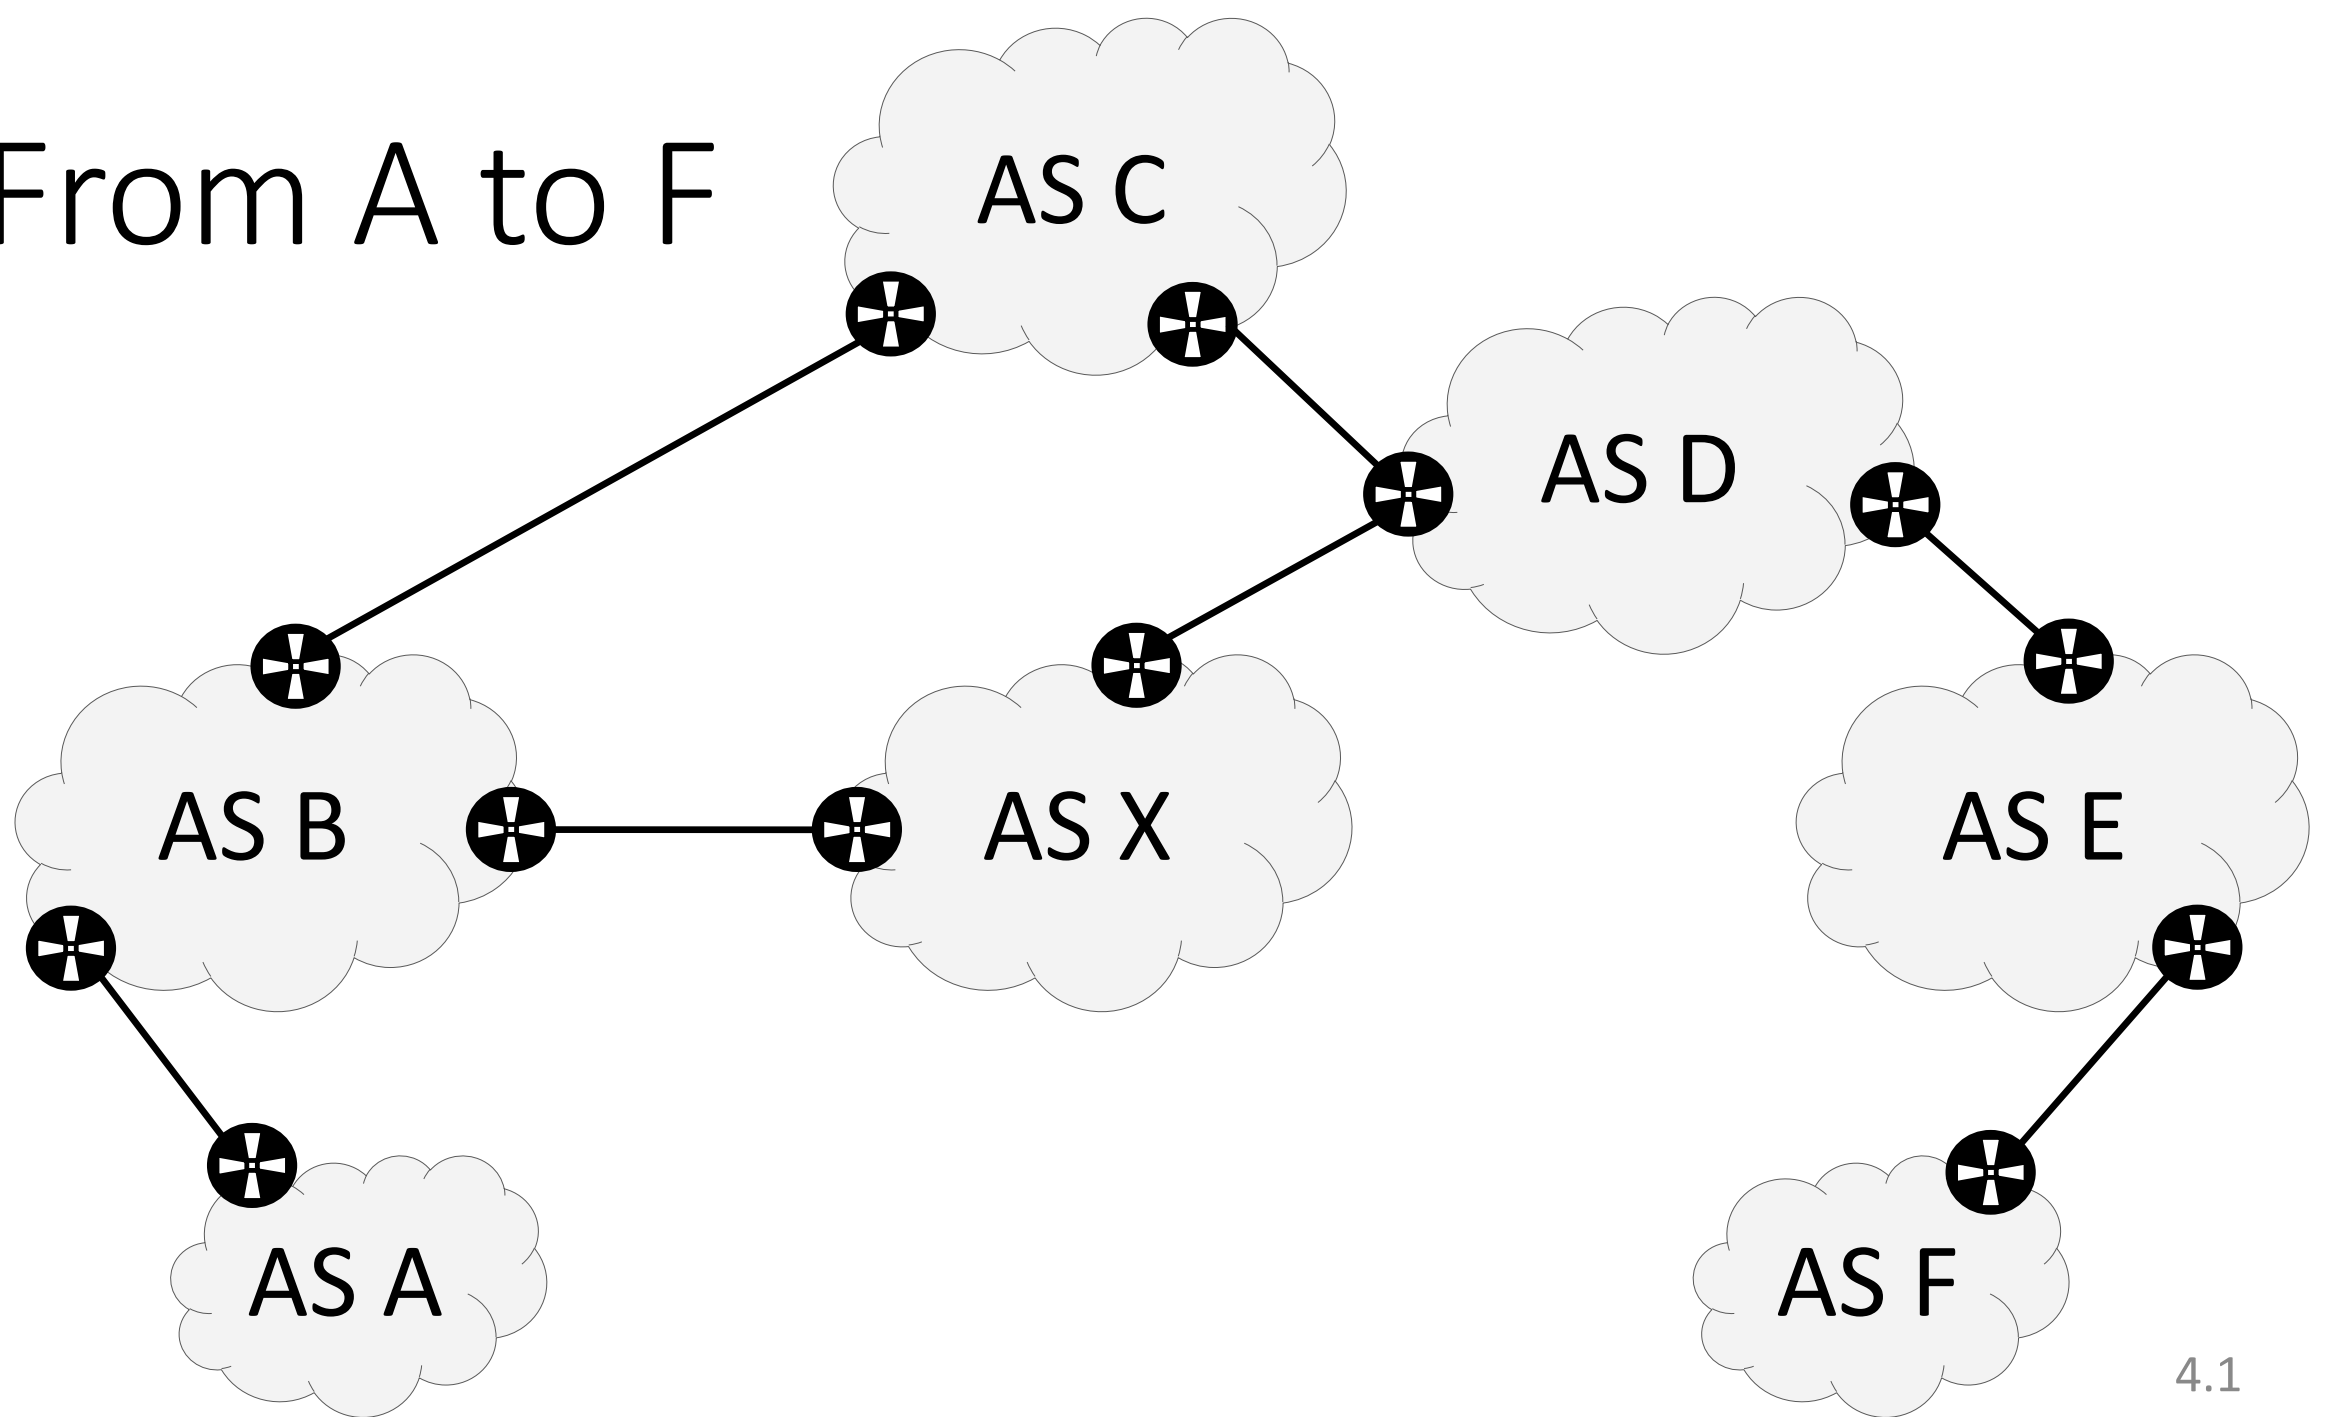
- ✓ How do they affect the forwarding?
- ✓ How can we detect them?

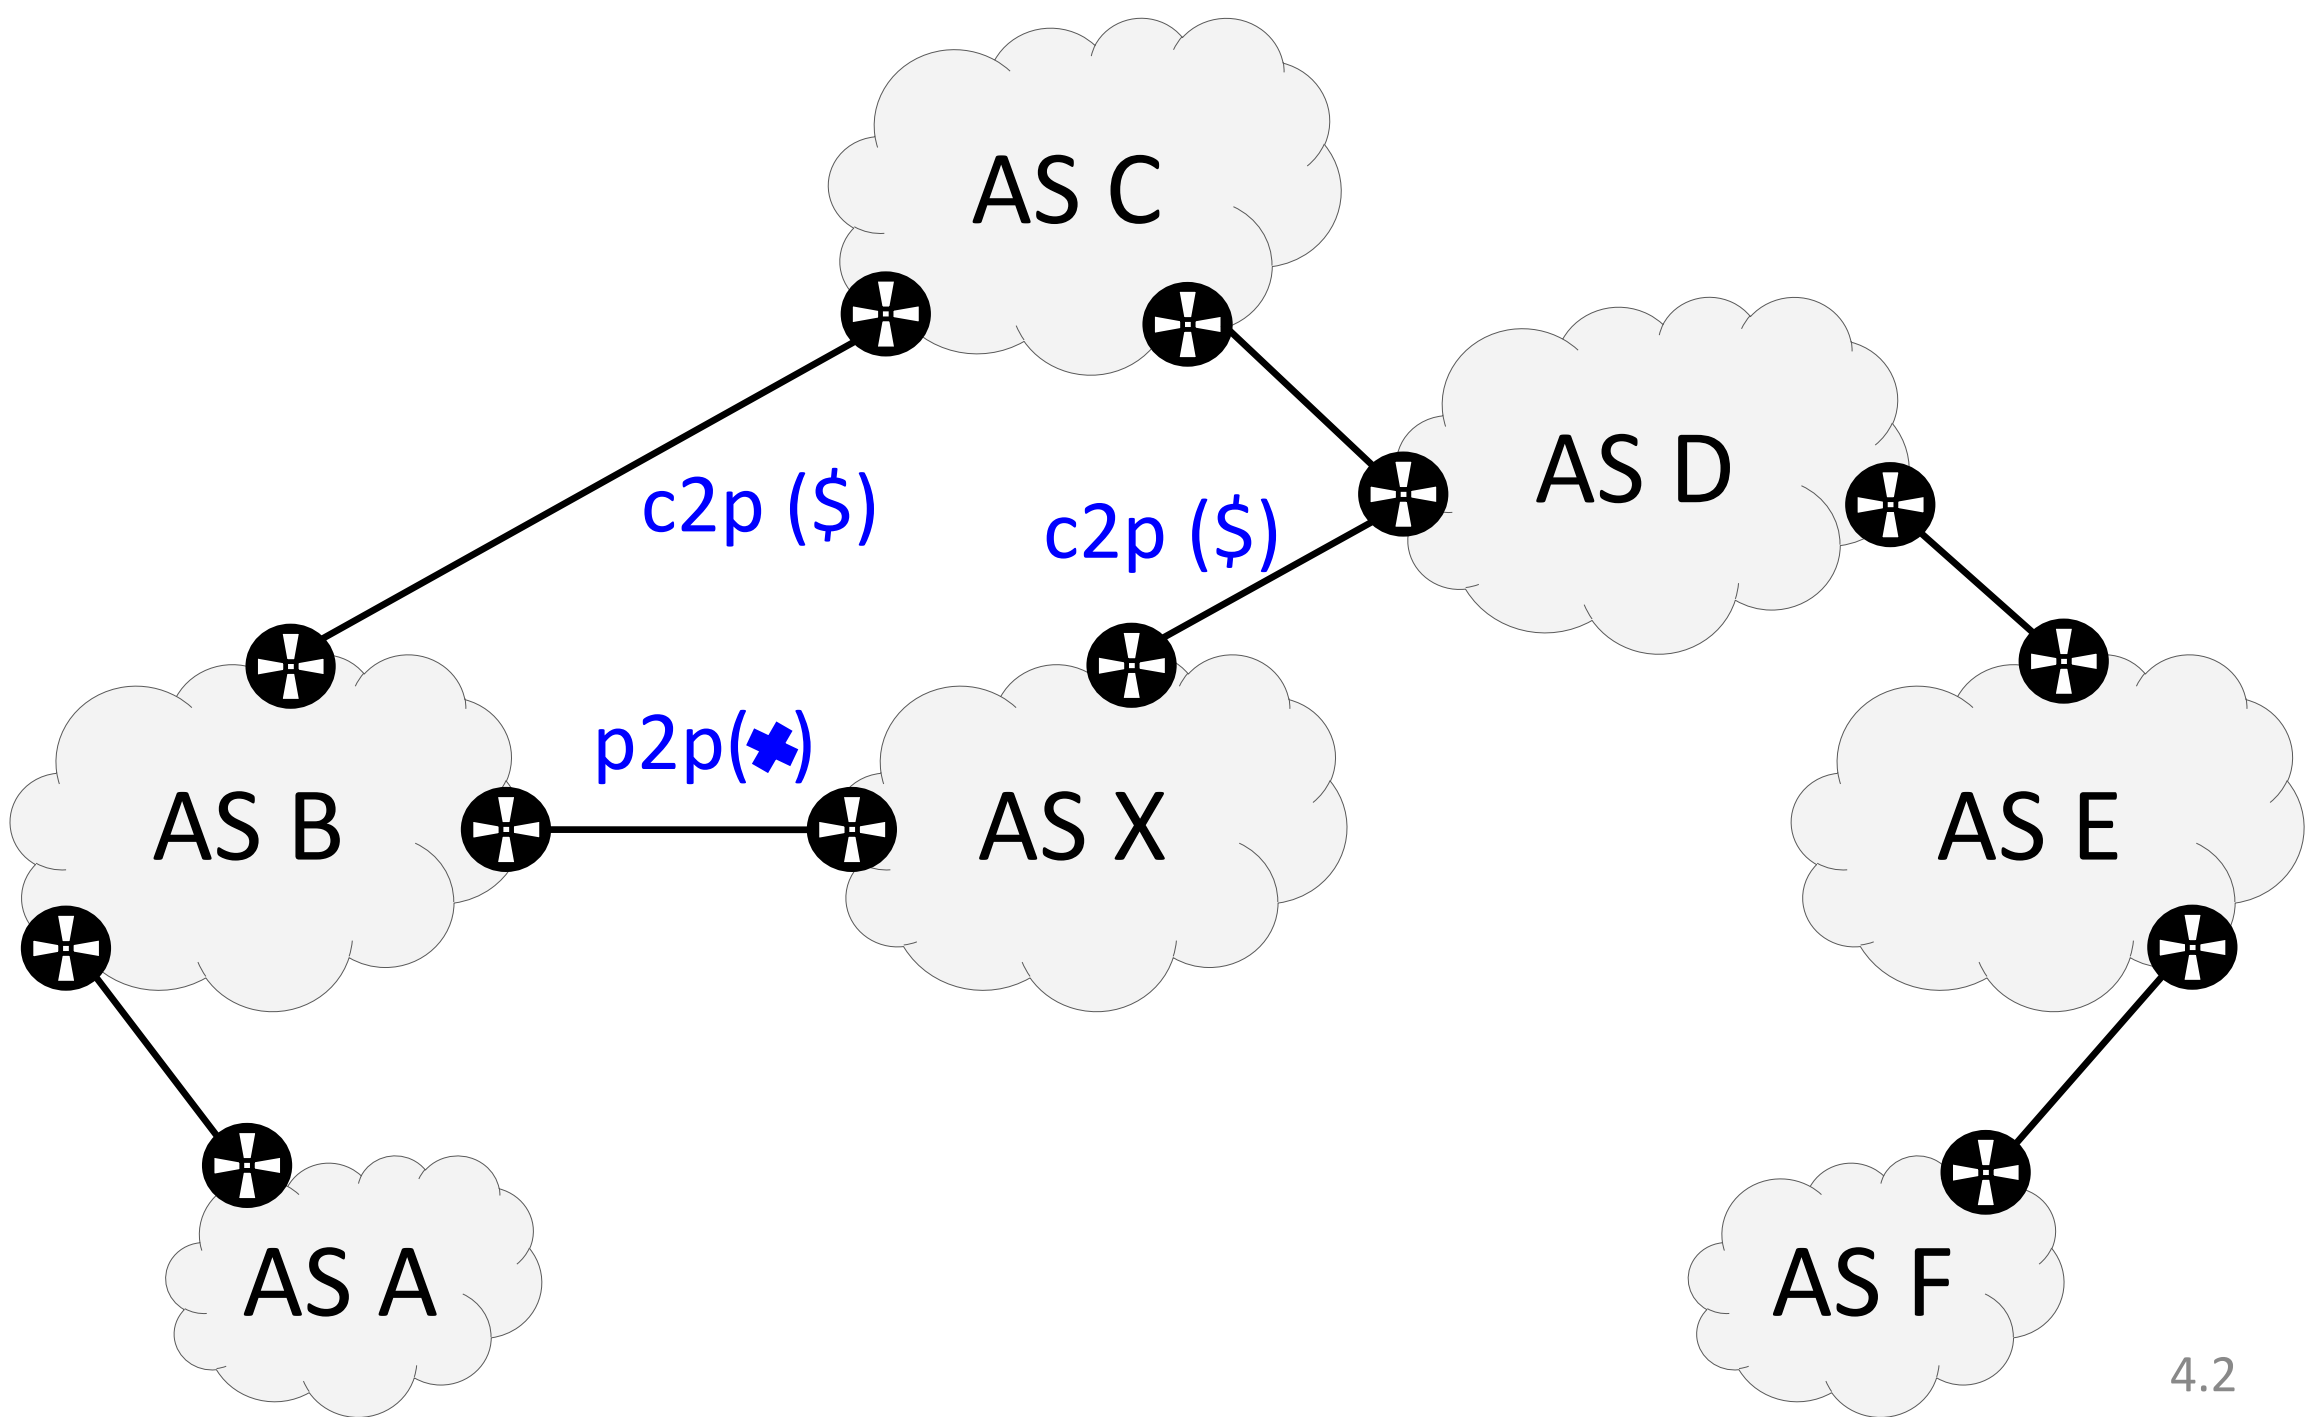
TMA 2019


“Filtering the Noise to Reveal Inter-Domain Lies”


Julian M. Del Fiore, Pascal Merindol, Valerio Persico
Cristel Pelsser, Antonio Pescape

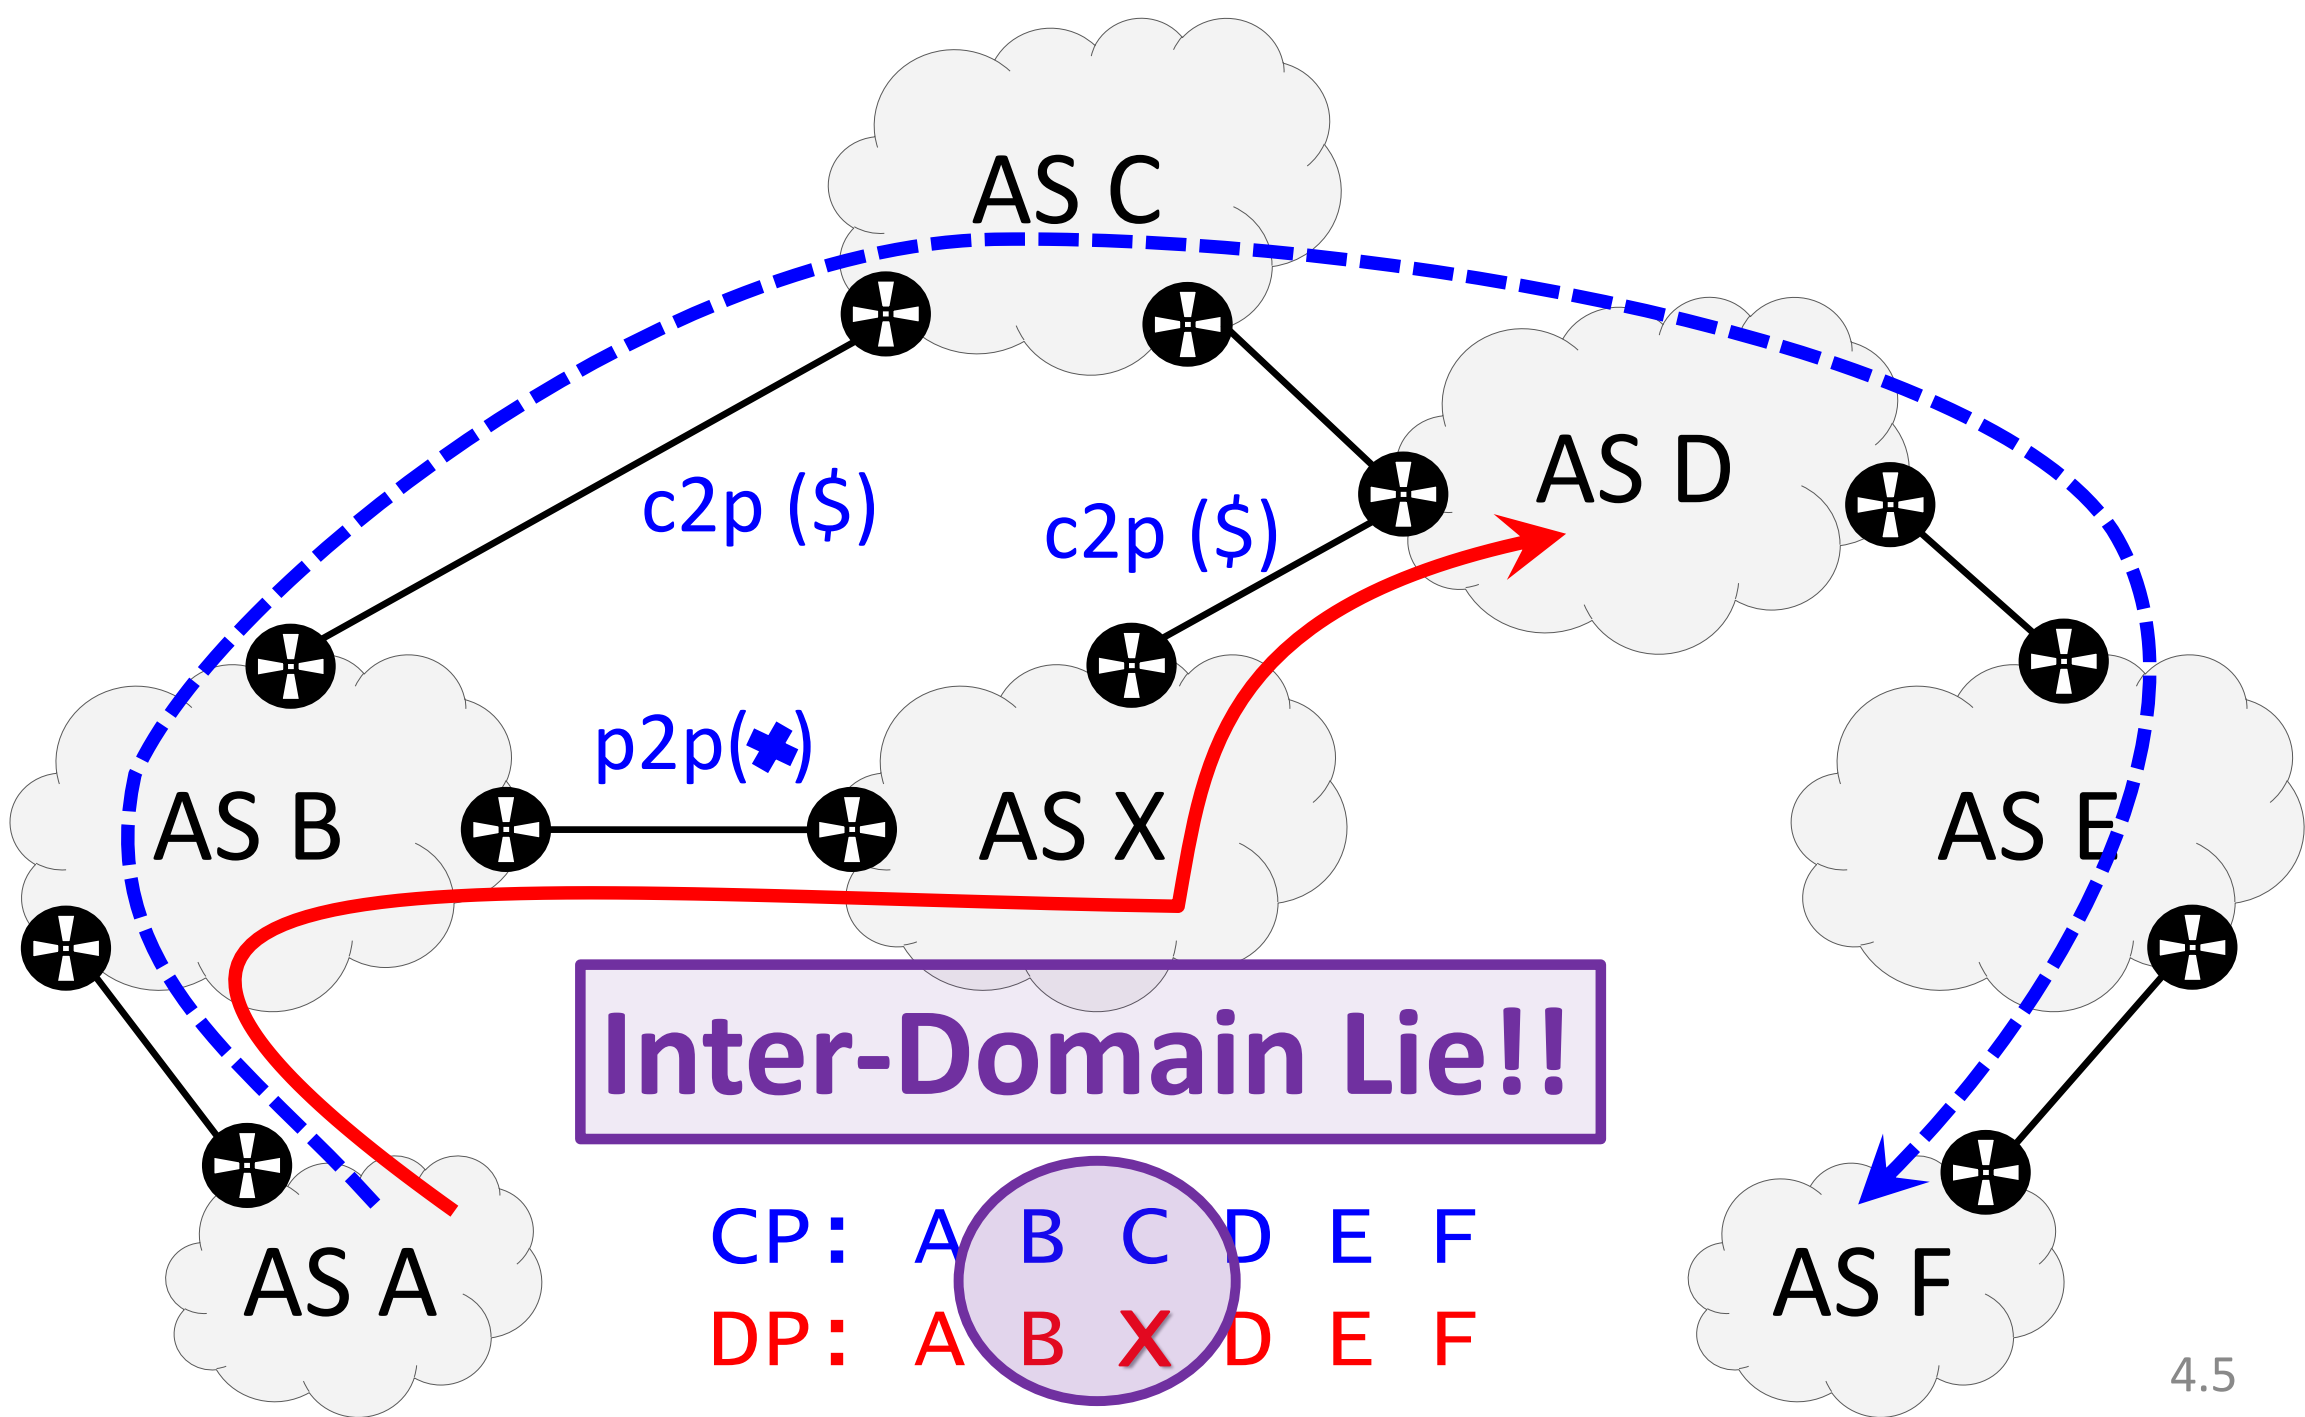
From A to F











Internet Measurements

...

BGP AS Paths



Control Paths



CPs

VS

Traceroute-AS Paths




Data Paths




DPs

Mismatches (MMs) in the Wild




Mismatches (MMs) in the Wild


Surprisingly High MM Rate!! 😞




Mismatches (MMs) in the Wild




Surprisingly High MM Rate!! 😞



But, What Causes the MMs? :/




MMs Reasons




Lies

MMs Reasons



Lies



Noise

MMs Reasons

Are we Just Capturing Noise? :S

Lies

Noise

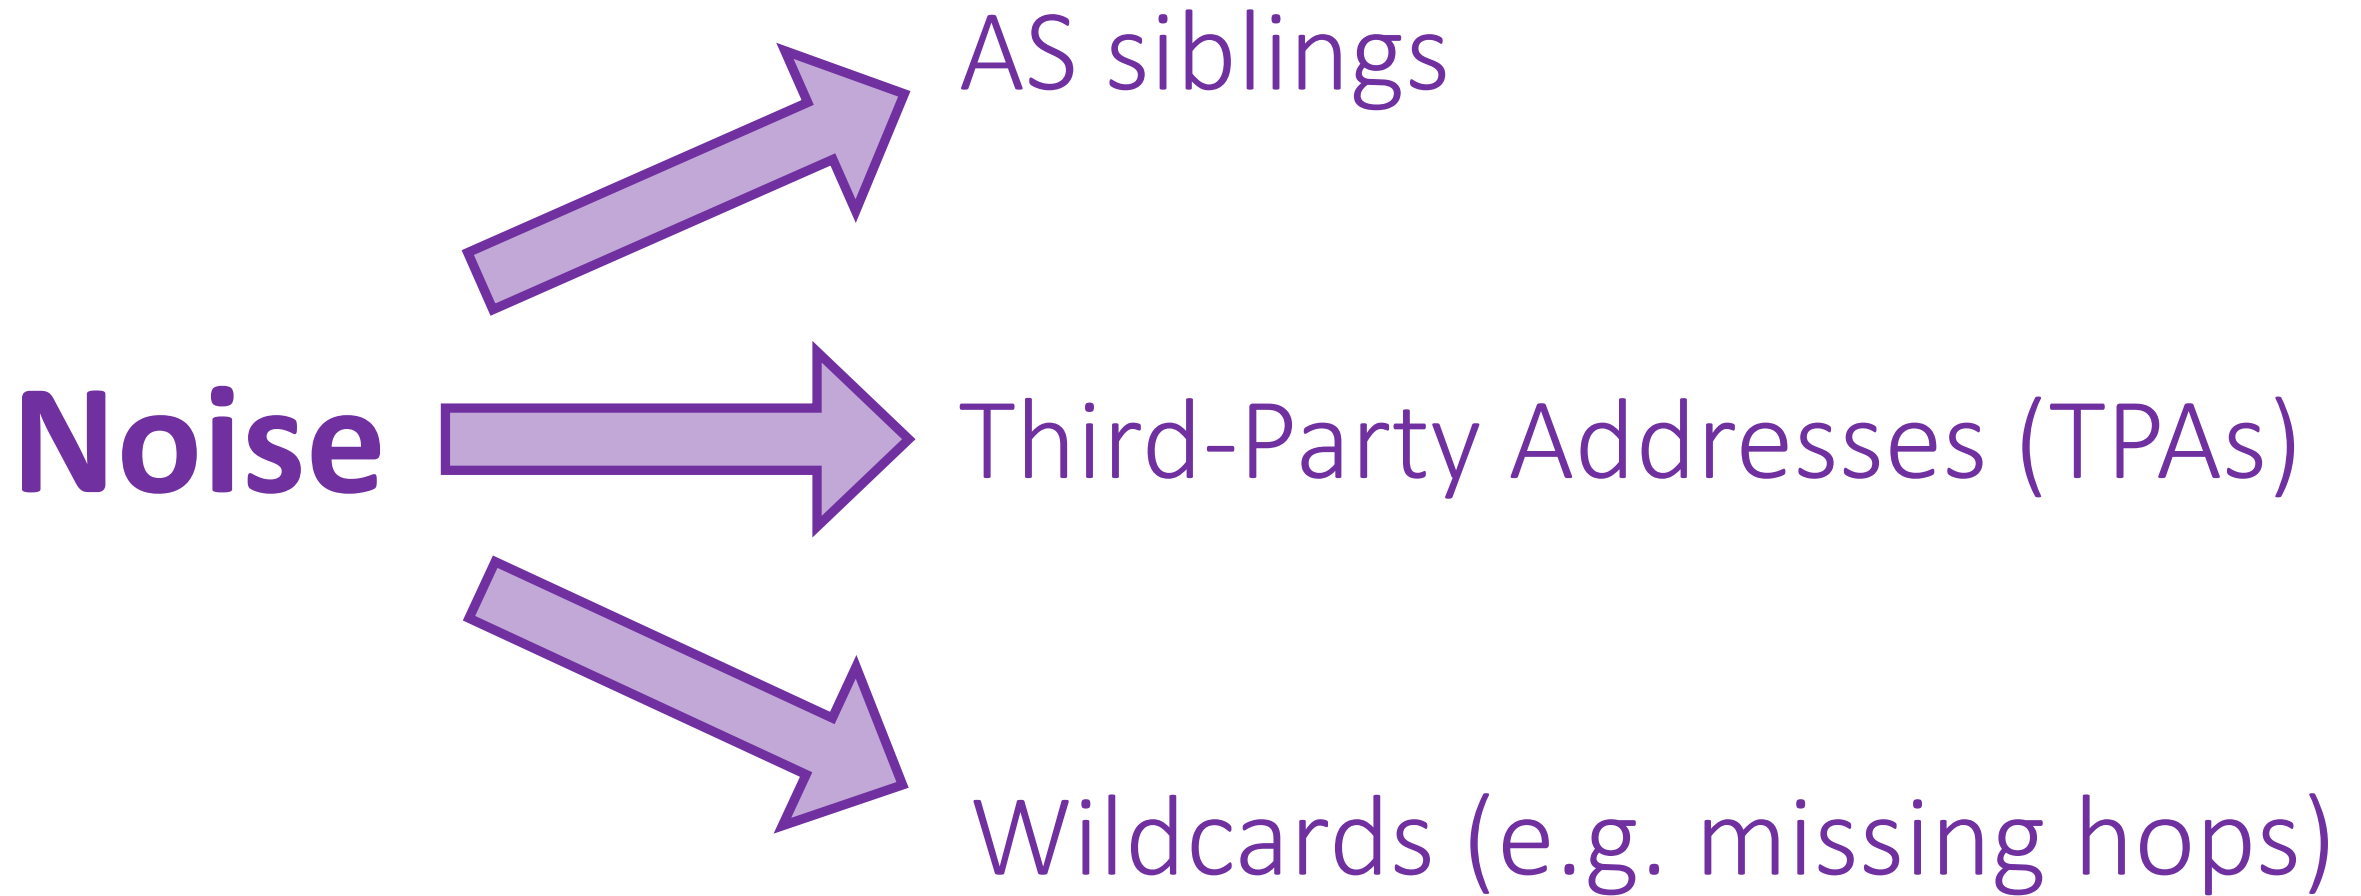
MMs Reasons

Are we Just Capturing Noise? :S


Lies

Noise

Please, Define Noise 😊




AS siblings




CP: Z A₀ A₁ B


AS siblings



TPAs




TPAs




TPAs

CP: A B C



TPAs

CP: A B C
DP: A B X C




Wildcards

Missing hops

Private IP Addresses

IP-to-AS mapping undefined




...

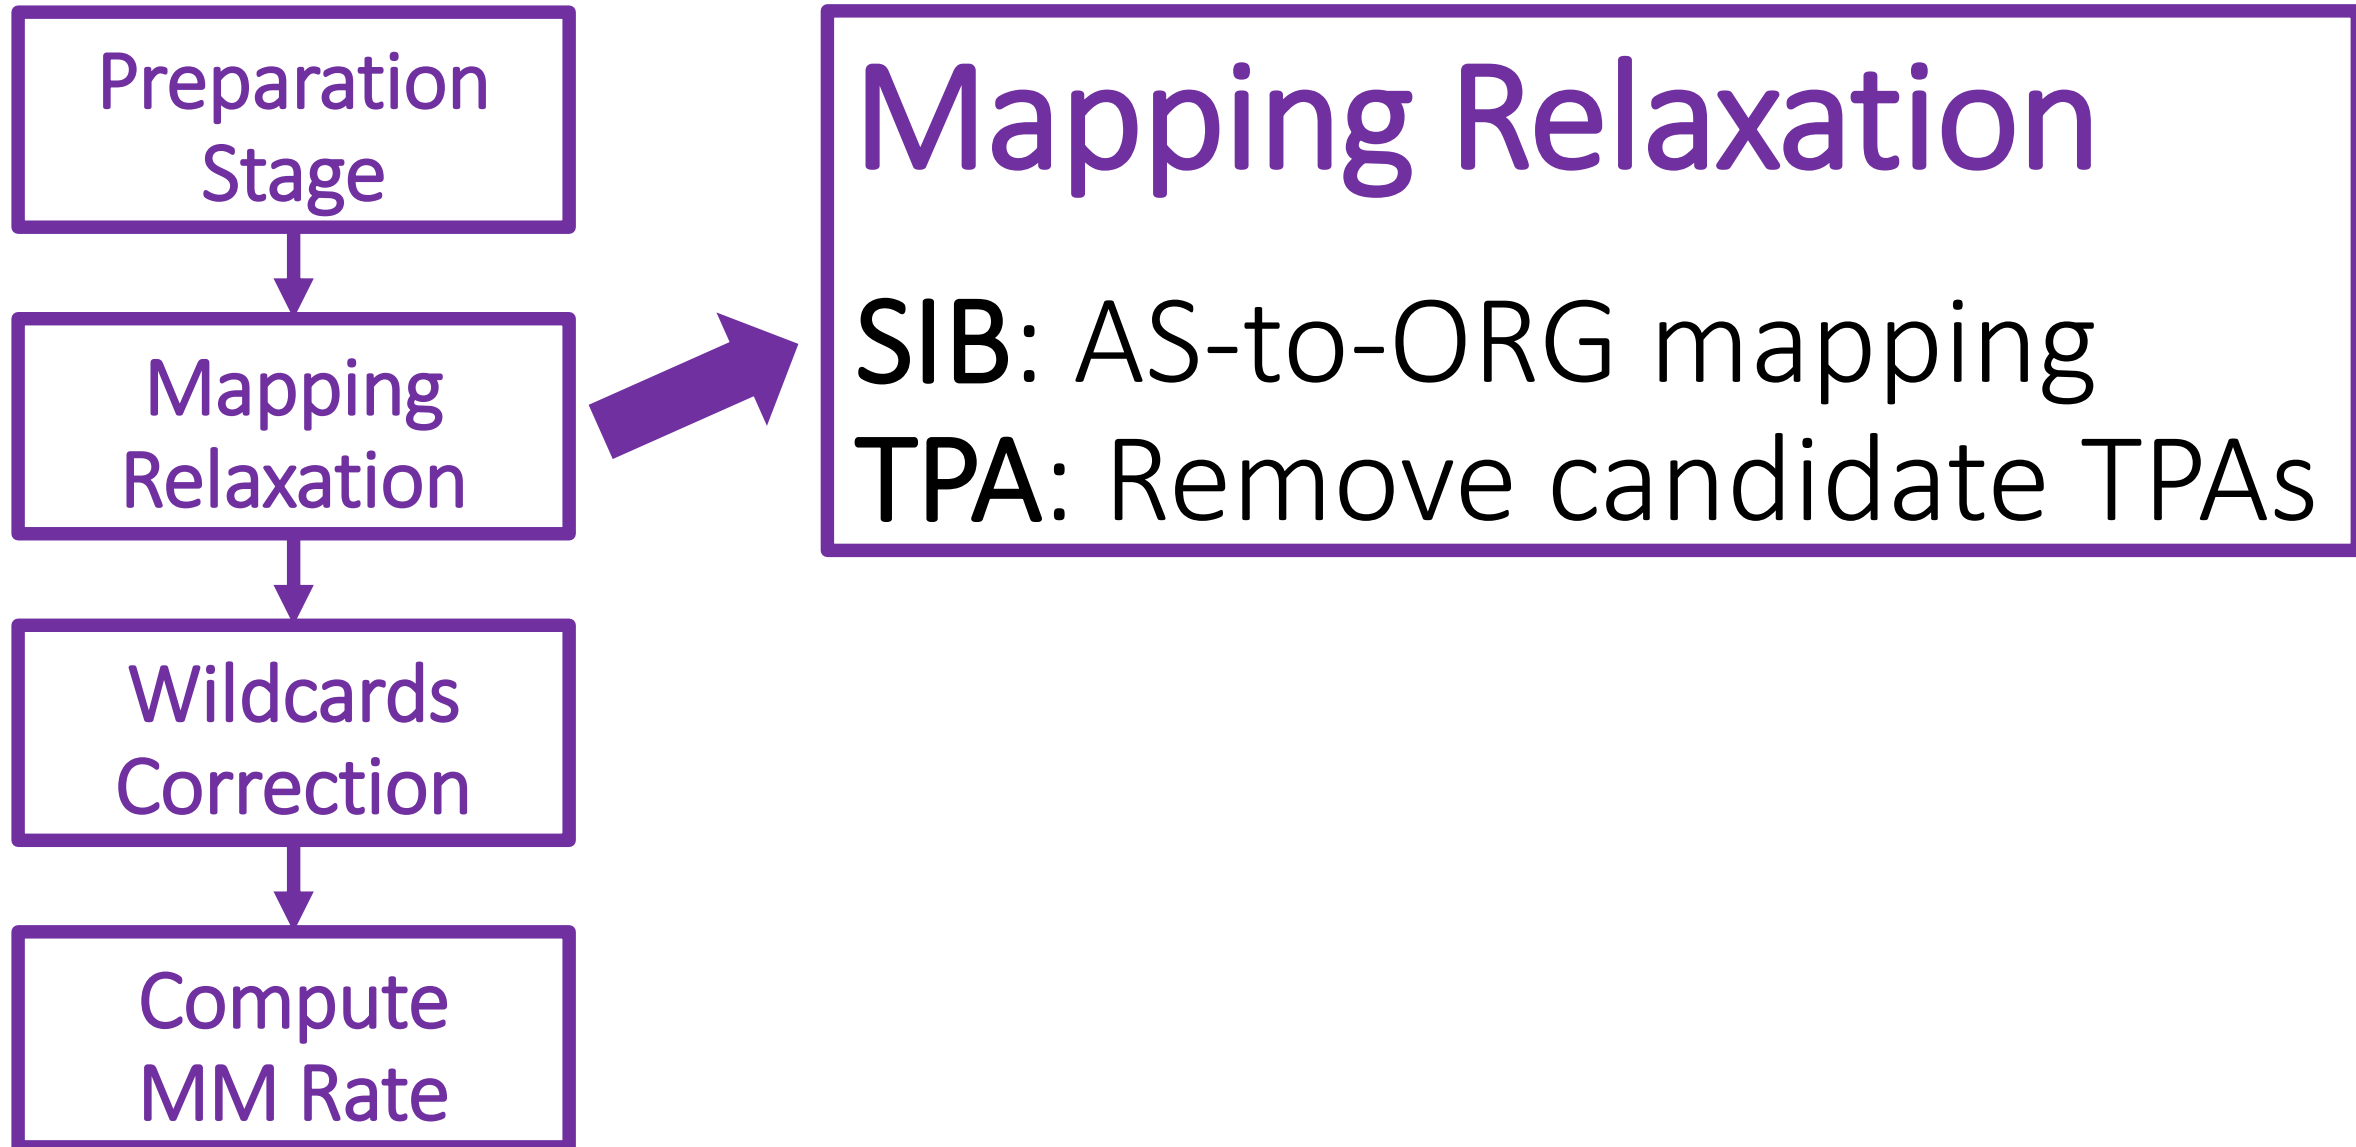
Filtering The Noise to Reveal Inter-Domain Lies

...

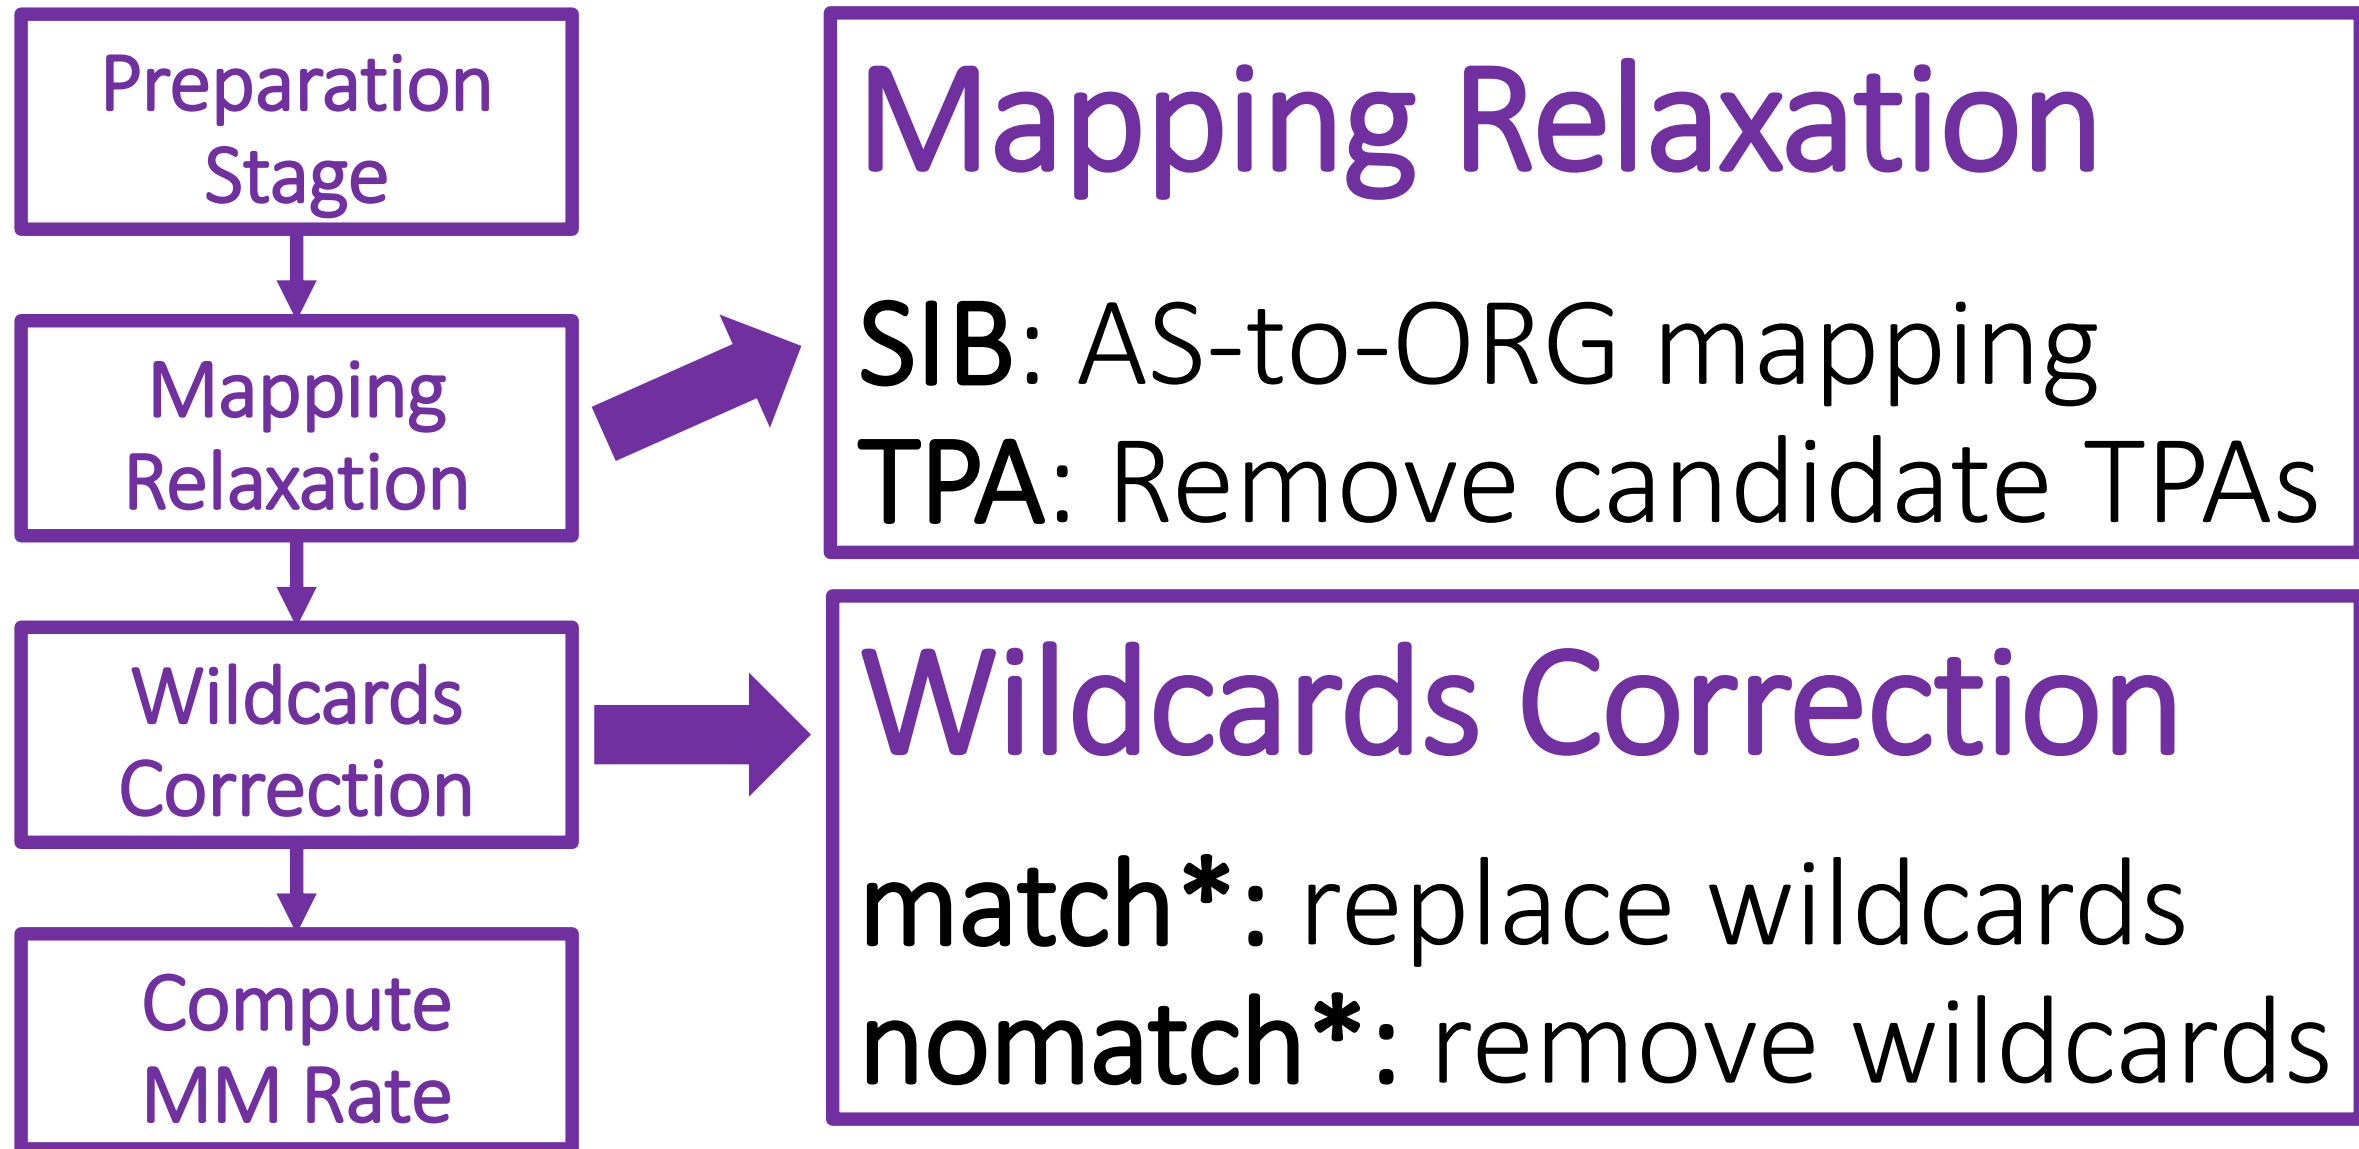
A Framework




A Framework




A Framework



MMs Bounds in the Wild




MMs Bounds in the Wild



Not zero!

MMs Bounds in the Wild



Conclusions

The Lower Bound of MMs
is not negligible


Noise usually **does not** include
AS siblings and TPAs
at the same time

Conext 2019...?

“Routing Inconsistencies at the FIB Level”

Work In Progress


What is a routing inconsistency (RI)?



What is a routing inconsistency (RI)?

BGP \longrightarrow $G(d_e)$: Gateway to reach IP d_e


IGP \longrightarrow $N(g_d)$: Next hop to reach Gateway g_d



What is a routing inconsistency (RI)?

BGP \longrightarrow $G(d_e)$: Gateway to reach IP d_e


IGP \longrightarrow $N(g_d)$: Next hop to reach Gateway g_d



What is a routing inconsistency (RI)?

BGP \longrightarrow $G(d_e)$: Gateway to reach IP d_e


IGP \longrightarrow $N(g_d)$: Next hop to reach Gateway g_d



What is a routing inconsistency (RI)?

BGP \longrightarrow $G(d_e)$: Gateway to reach IP d_e


IGP \longrightarrow $N(g_d)$: Next hop to reach Gateway g_d



What is a routing inconsistency (RI)?

BGP \longrightarrow $G(d_e)$: Gateway to reach IP d_e


IGP \longrightarrow $N(g_d)$: Next hop to reach Gateway g_d



What is a routing inconsistency (RI)?

BGP \longrightarrow $G(d_e)$: Gateway to reach IP d_e


IGP \longrightarrow $N(g_d)$: Next hop to reach Gateway g_d



What is a routing inconsistency (RI)?

BGP \longrightarrow $G(d_e)$: Gateway to reach IP d_e


IGP \longrightarrow $N(g_d)$: Next hop to reach Gateway g_d



What is a routing inconsistency (RI)?

BGP \longrightarrow $G(d_e)$: Gateway to reach IP d_e


IGP \longrightarrow $N(g_d)$: Next hop to reach Gateway g_d



What is a routing inconsistency (RI)?

BGP \longrightarrow $G(d_e)$: Gateway to reach IP d_e


IGP \longrightarrow $N(g_d)$: Next hop to reach Gateway g_d



What is a routing inconsistency (RI)?

BGP \longrightarrow $G(d_e)$: Gateway to reach IP d_e

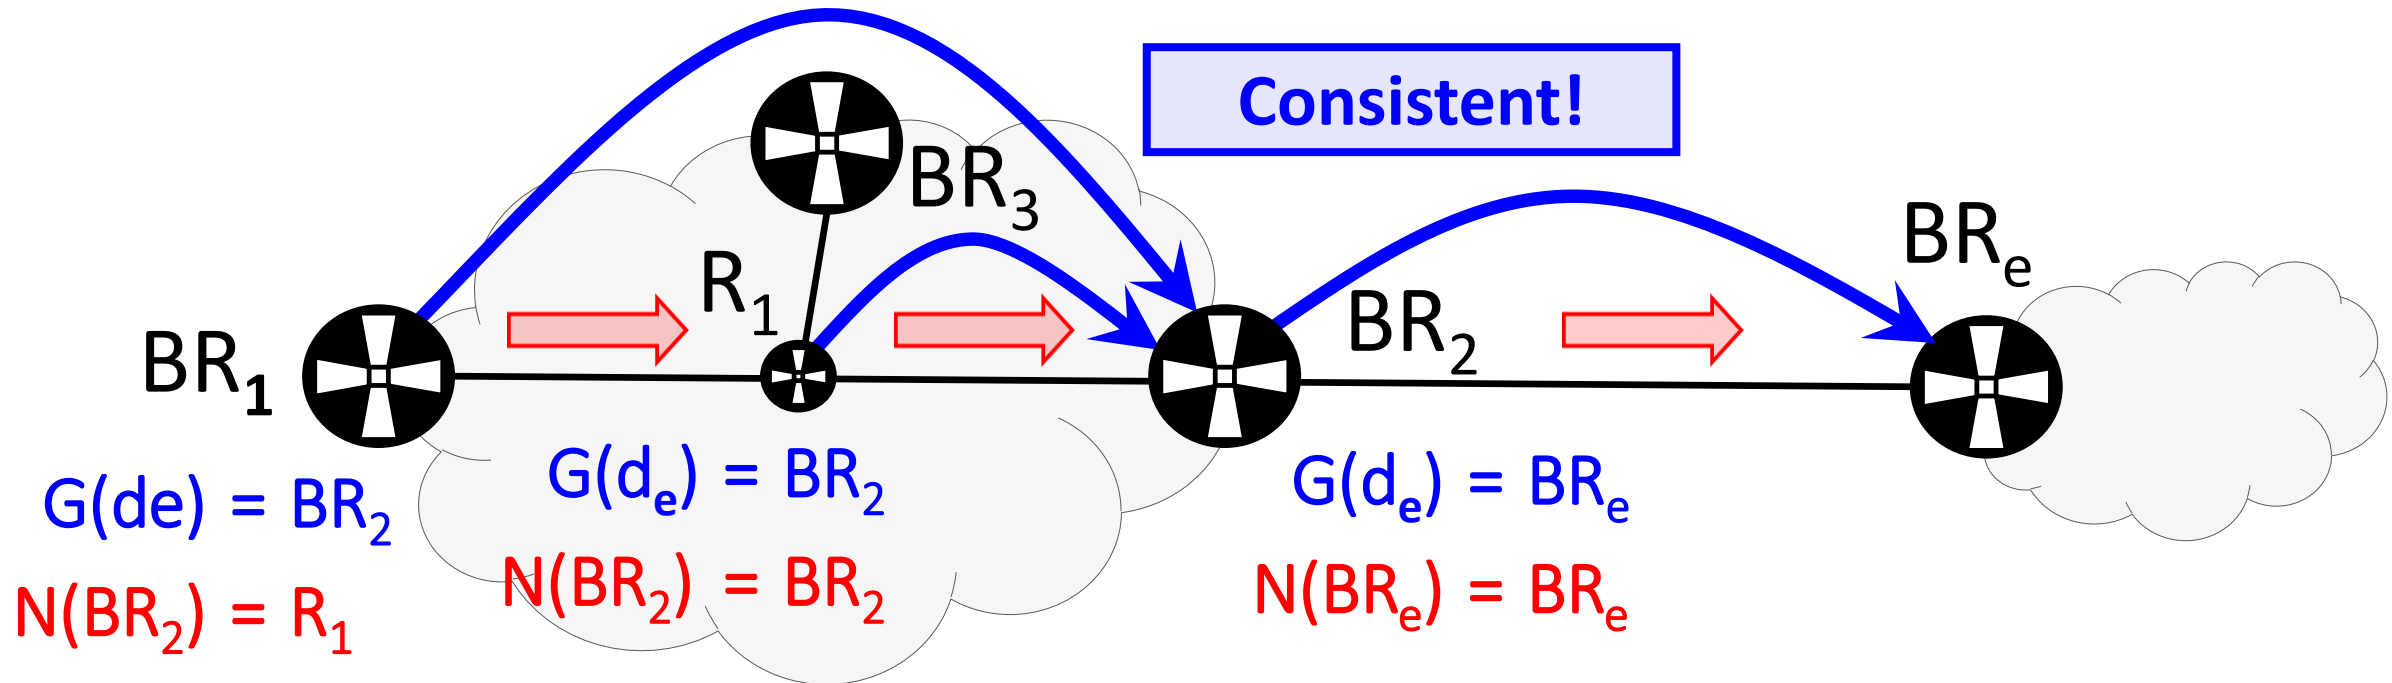
IGP \longrightarrow $N(g_d)$: Next hop to reach Gateway g_d



What is a routing inconsistency (RI)?

BGP \longrightarrow $G(d_e)$: Gateway to reach IP d_e

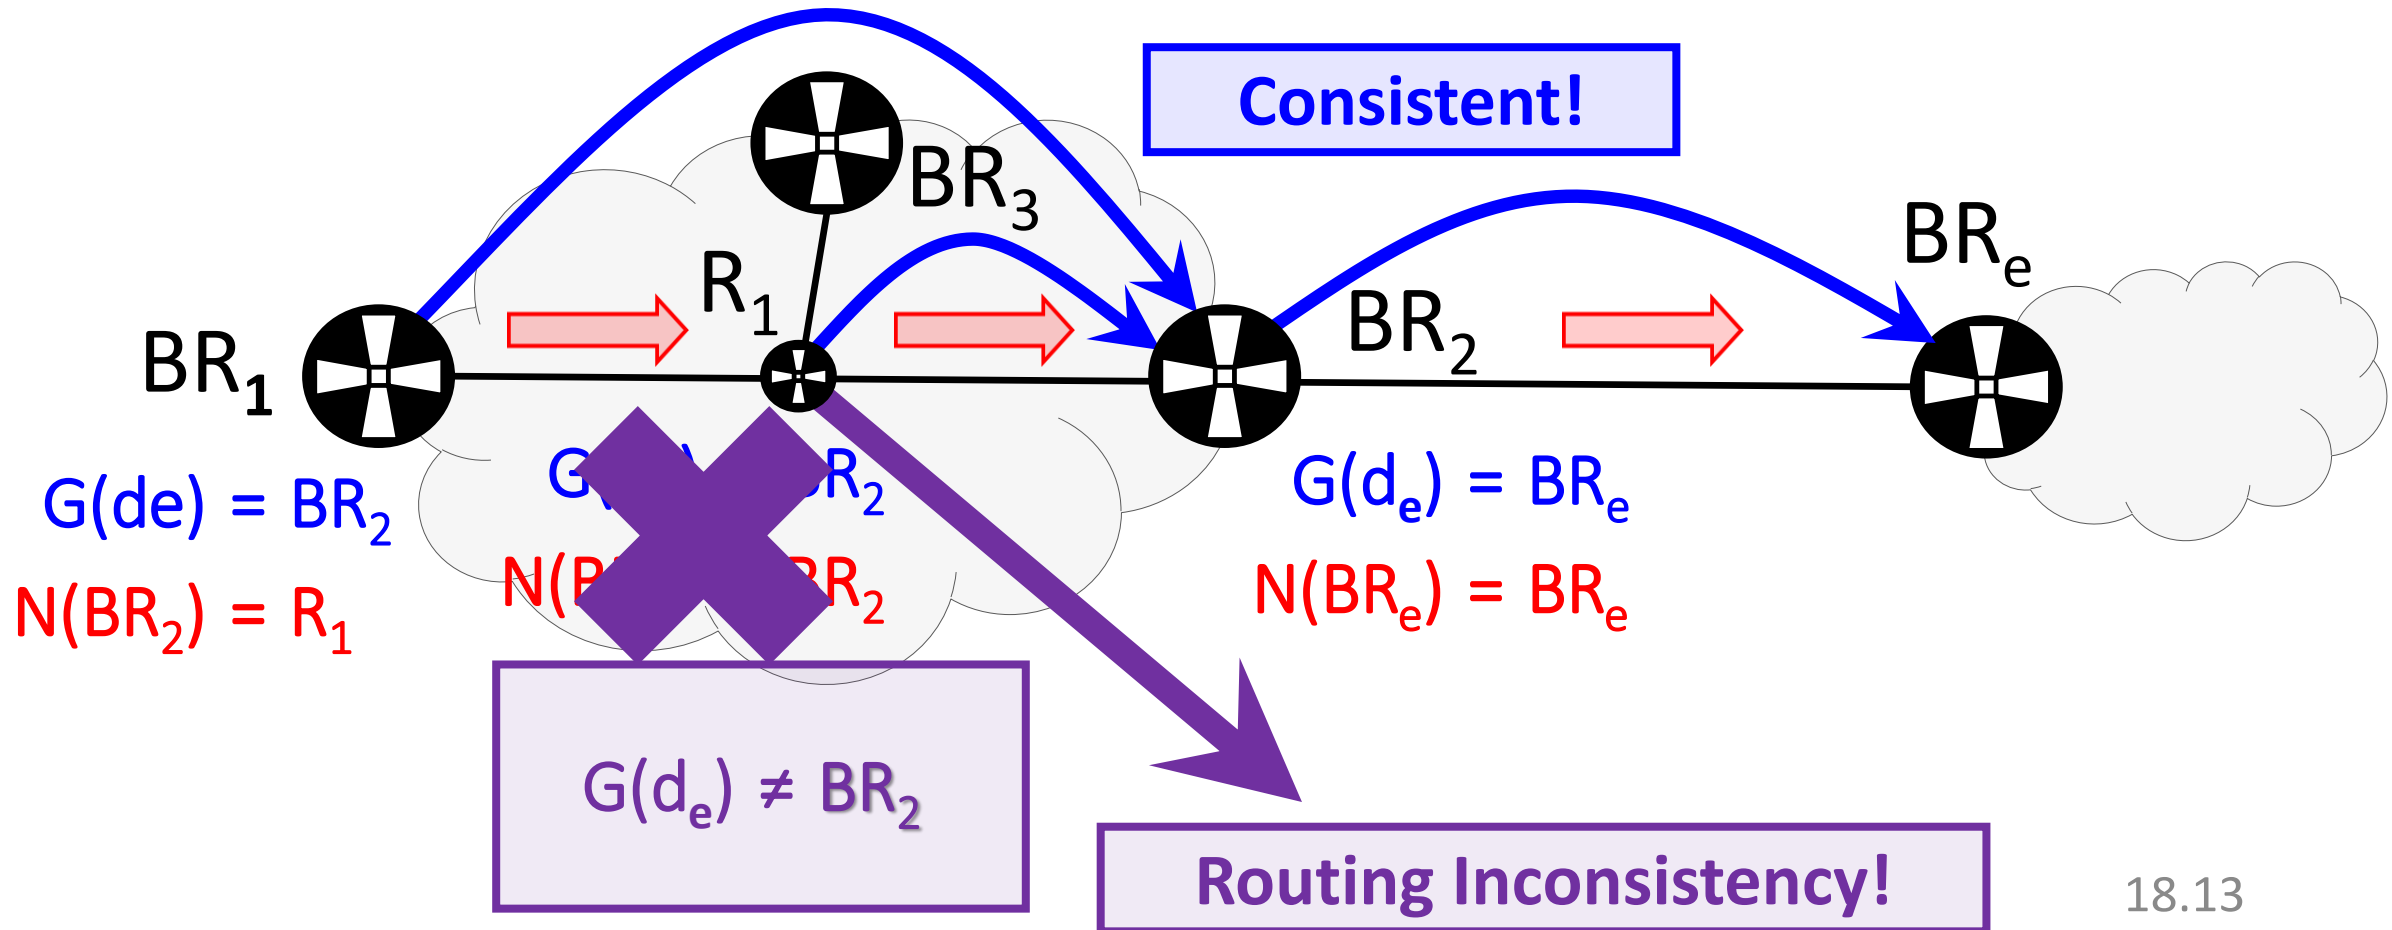
IGP \longrightarrow $N(g_d)$: Next hop to reach Gateway g_d



What is a routing inconsistency (RI)?


BGP \longrightarrow $G(d_e)$: Gateway to reach IP d_e

IGP \longrightarrow $N(g_d)$: Next hop to reach Gateway g_d




What are the effects of Rles?


Dest -> d




What are the effects of Rles?




What are the effects of Rles?



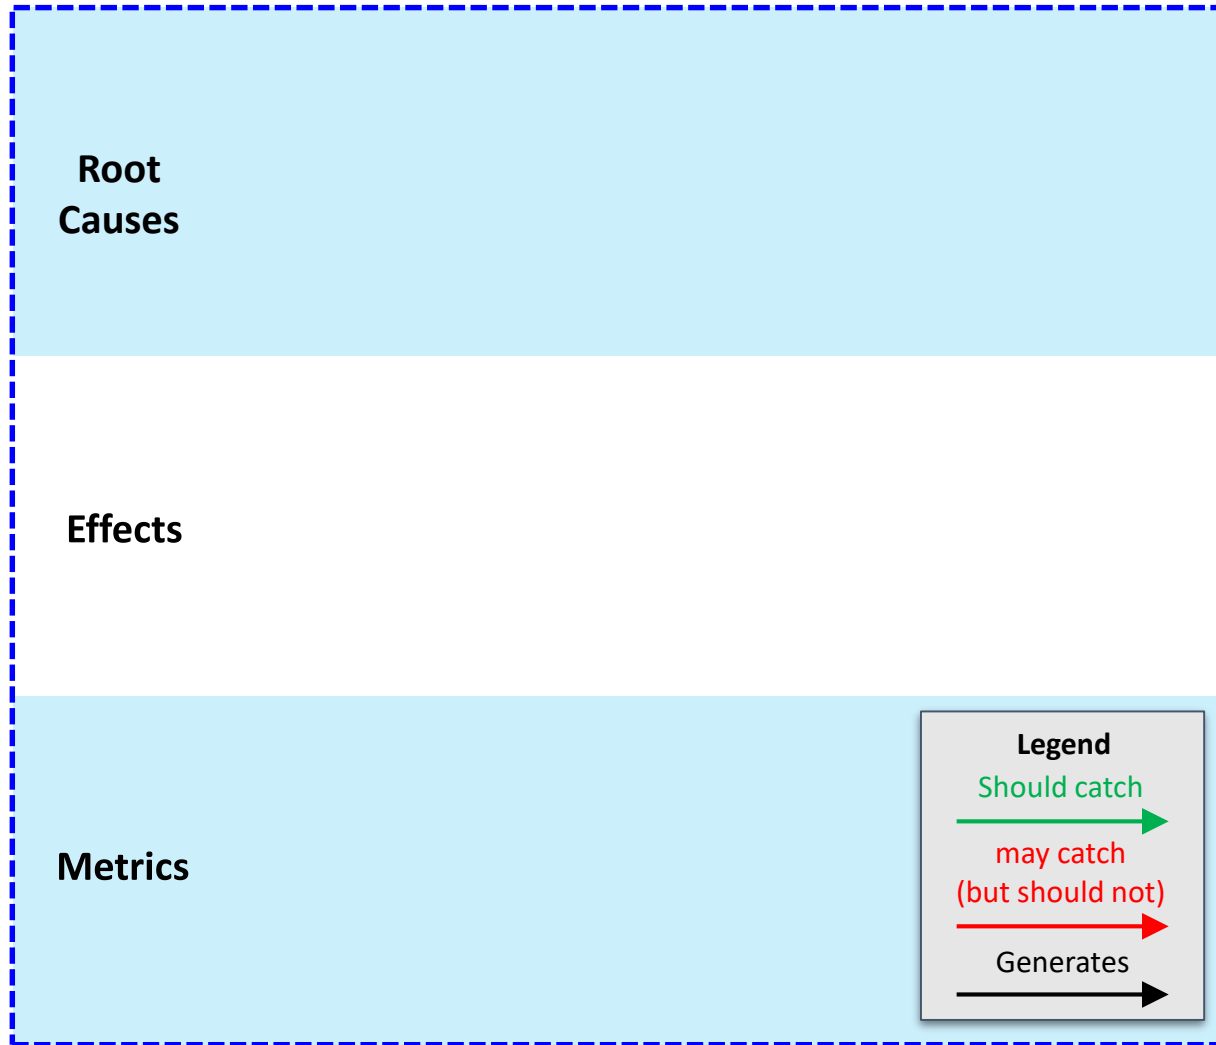
What are the effects of Rles?



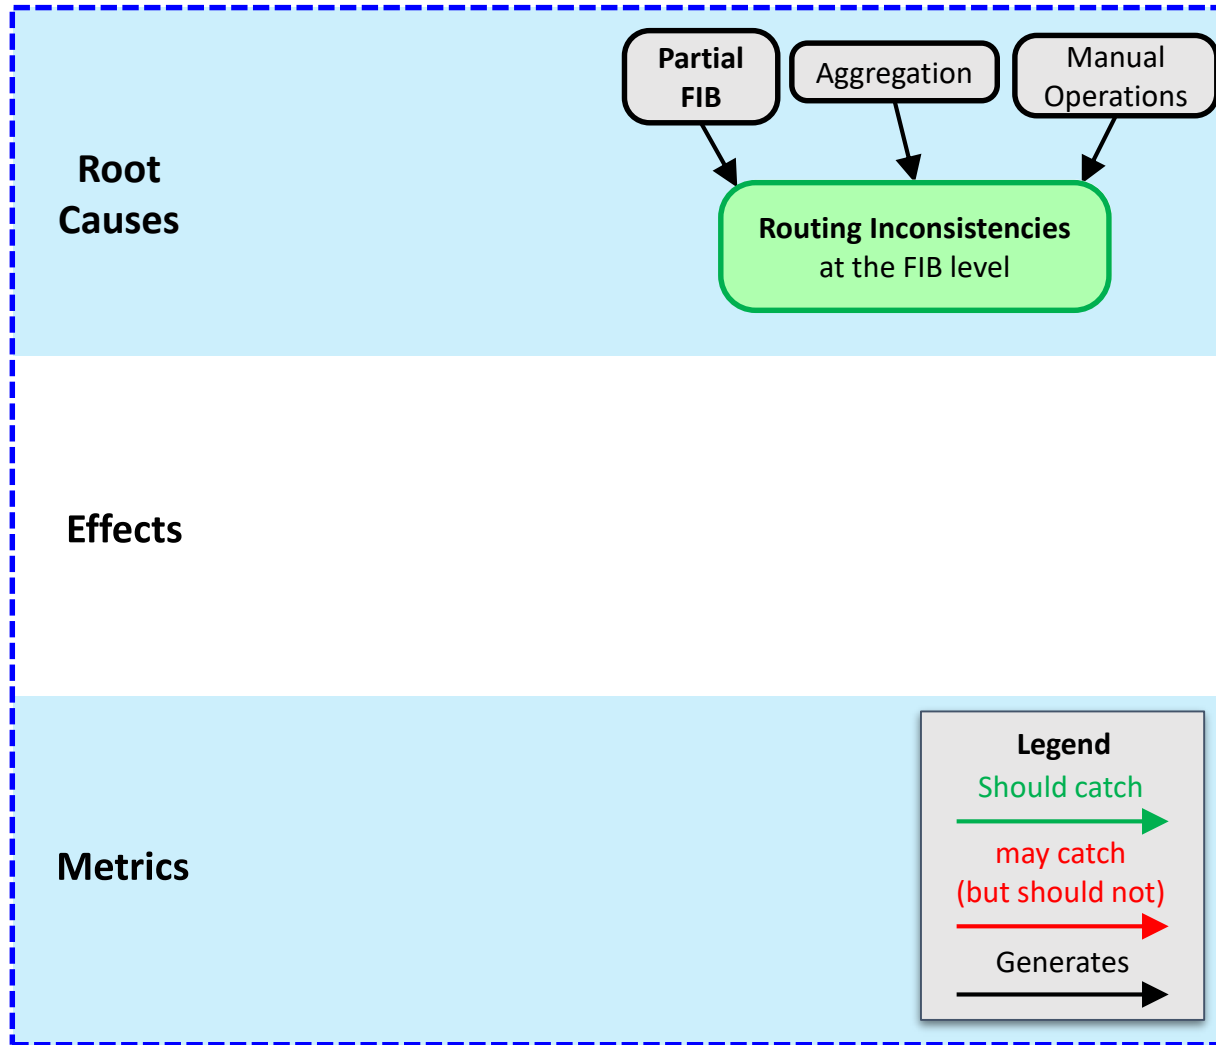
What are the effects of Rles?



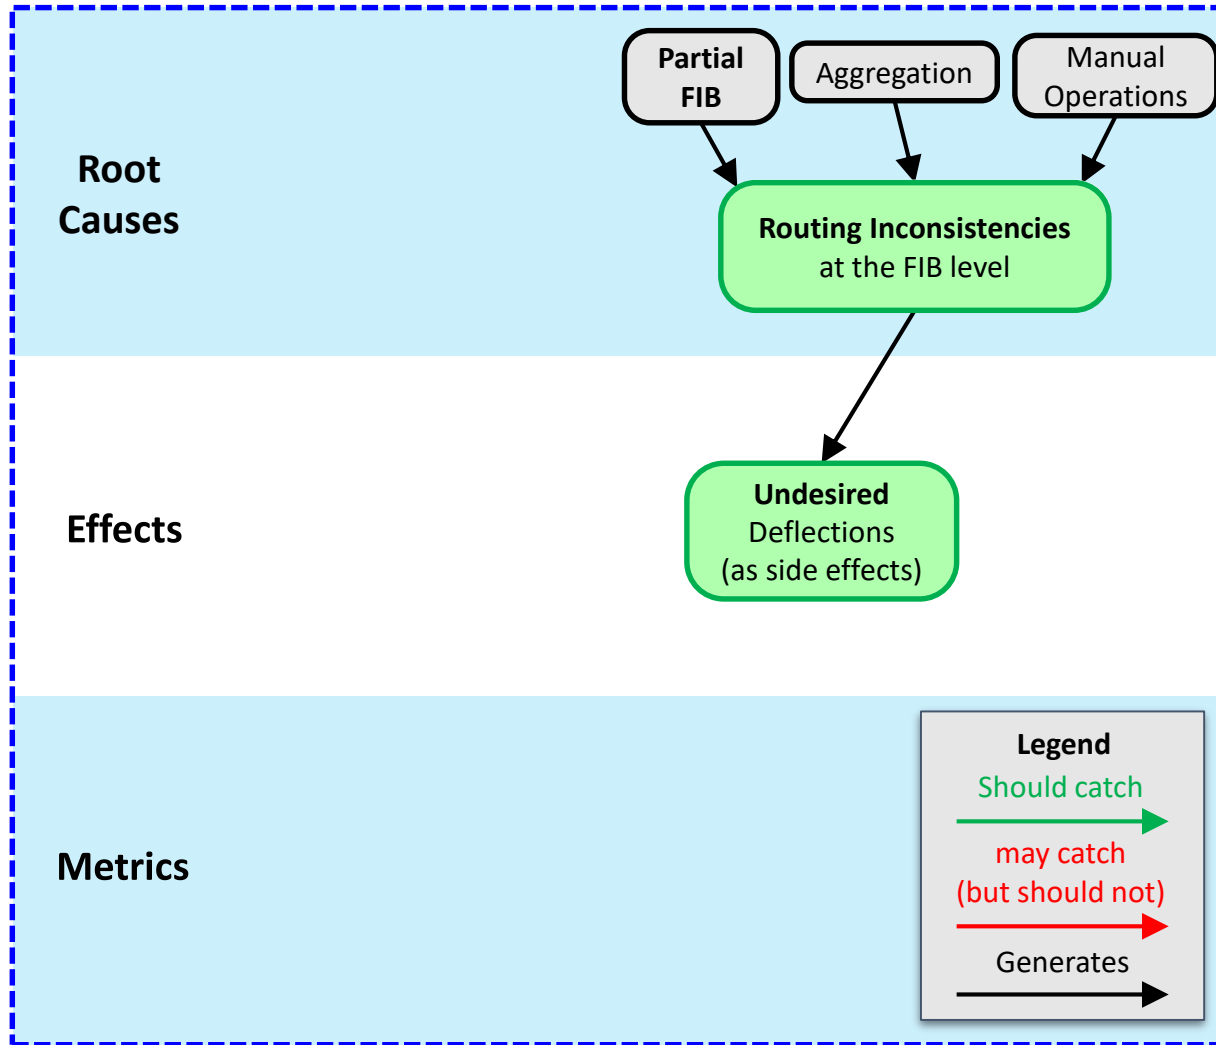
A challenging problem



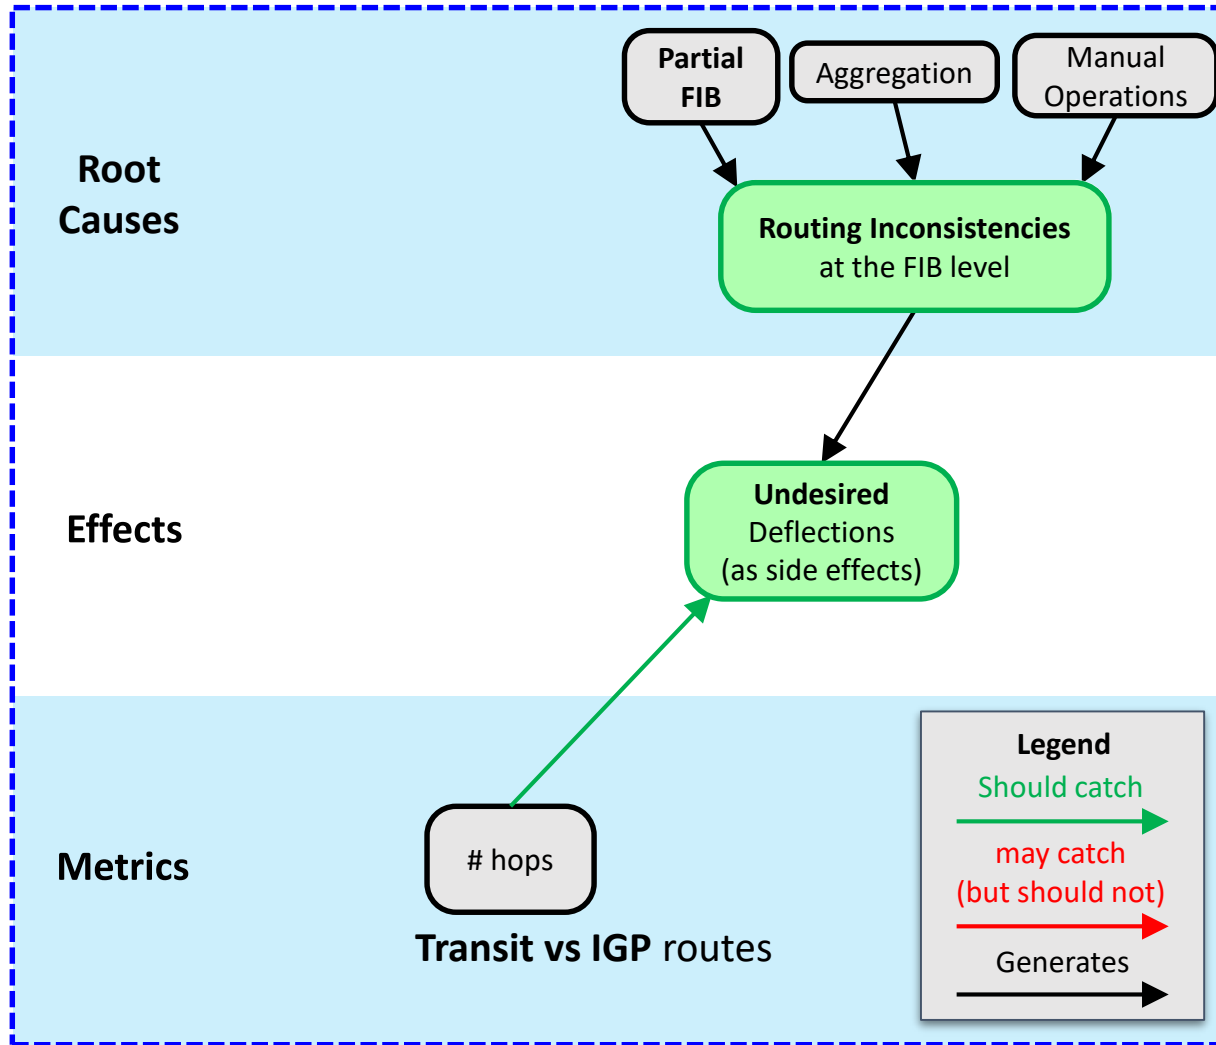
A challenging problem



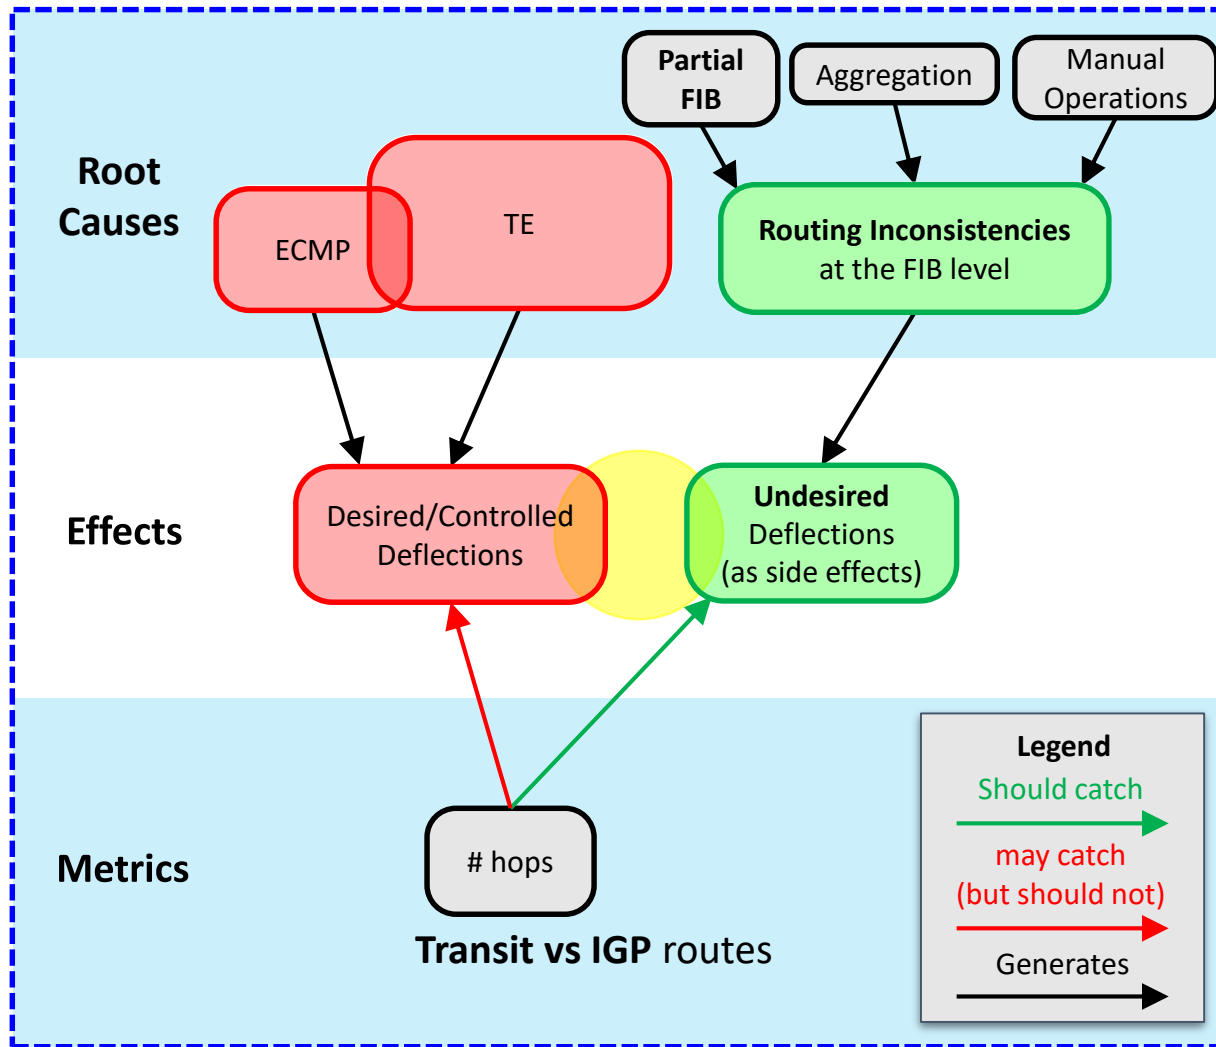
A challenging problem



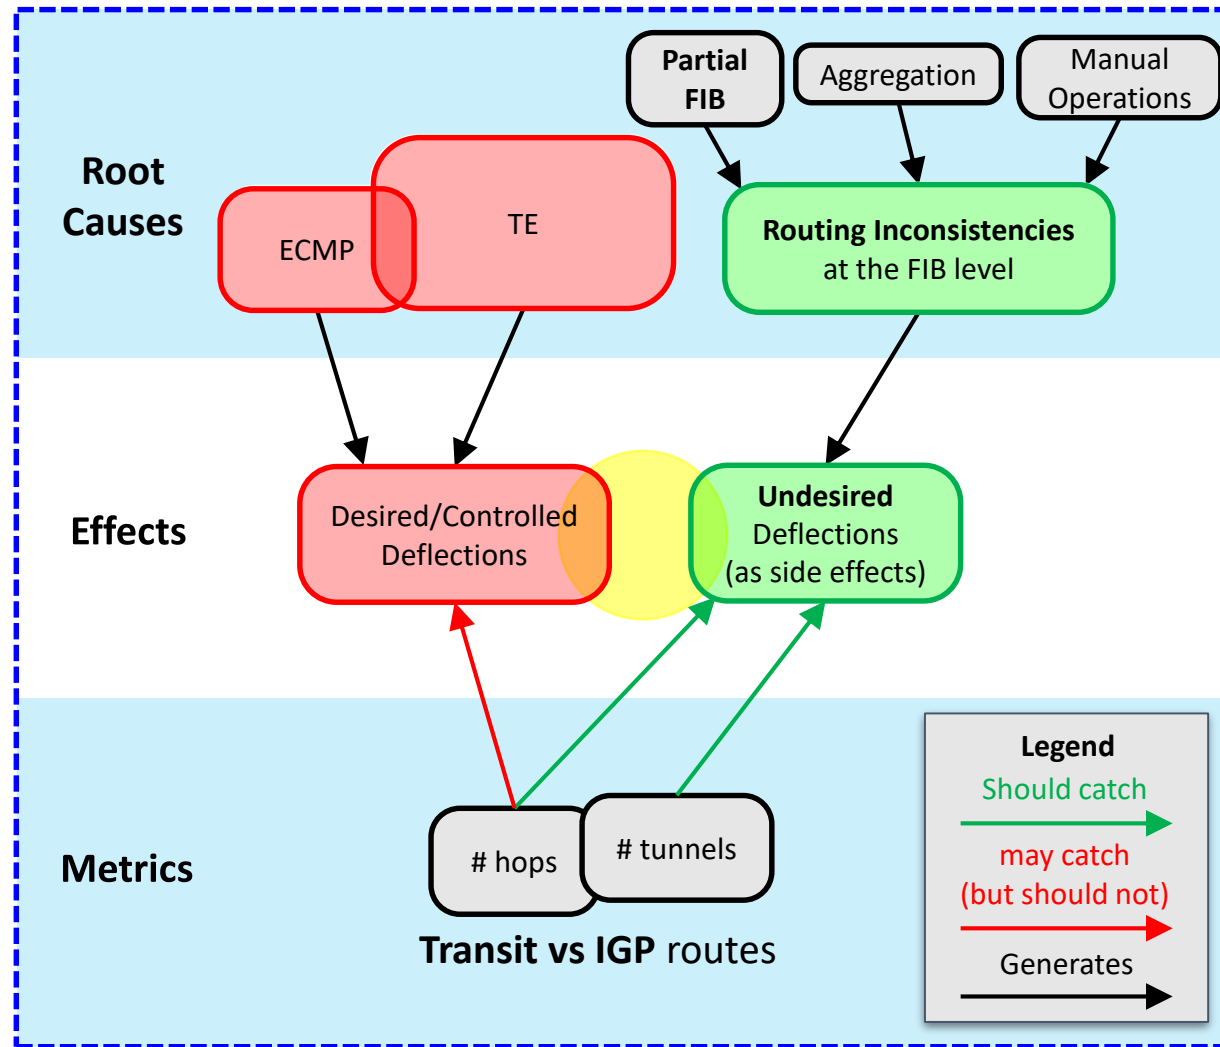
A challenging problem




A challenging problem



A challenging problem




How to Overcome ECMP?



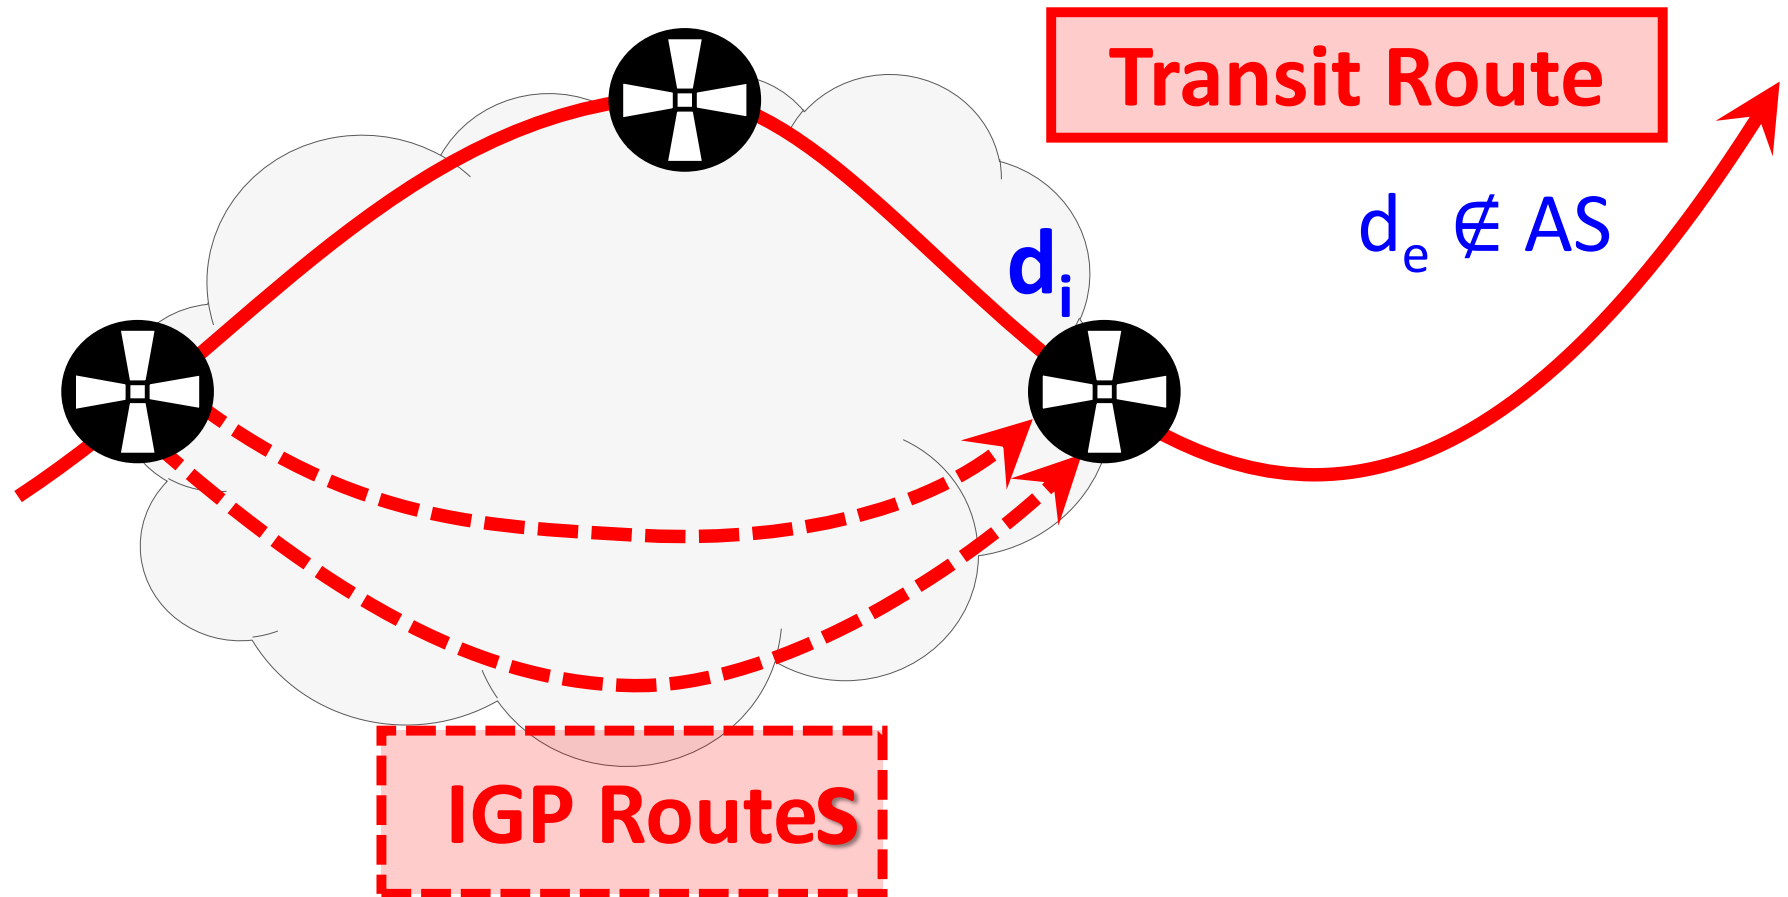
How to Overcome ECMP?

Detect all IGP routes with MDA-traceroute



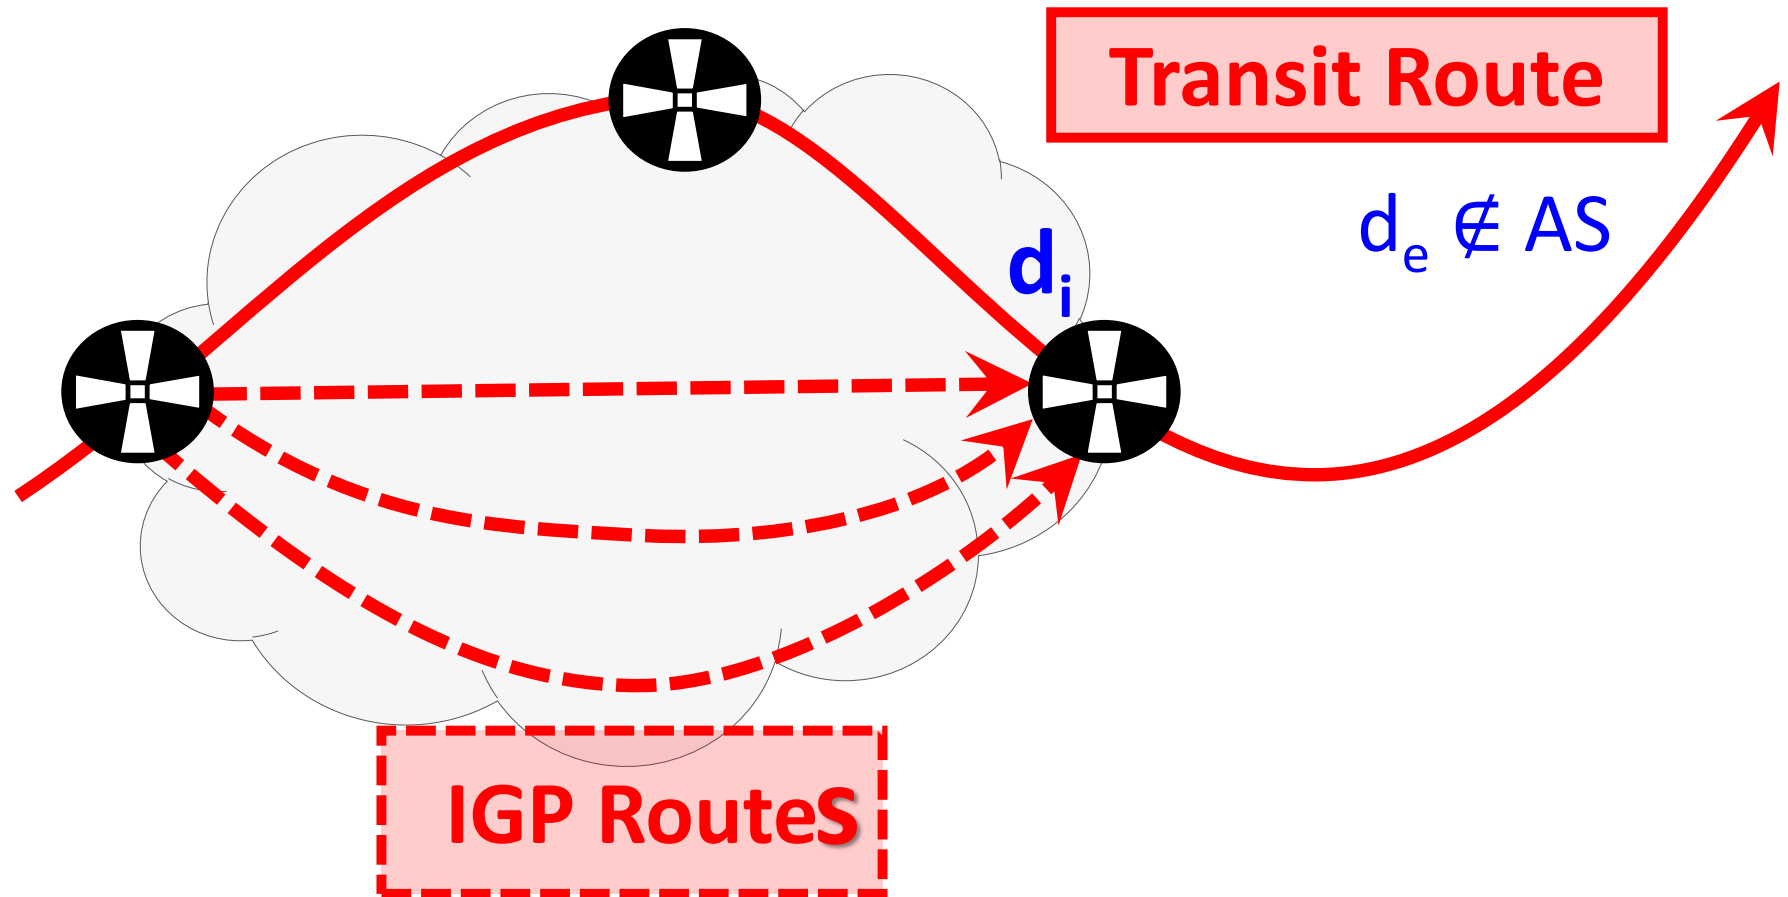
How to Overcome ECMP?

Detect all IGP routes with MDA-traceroute



How to Overcome ECMP?

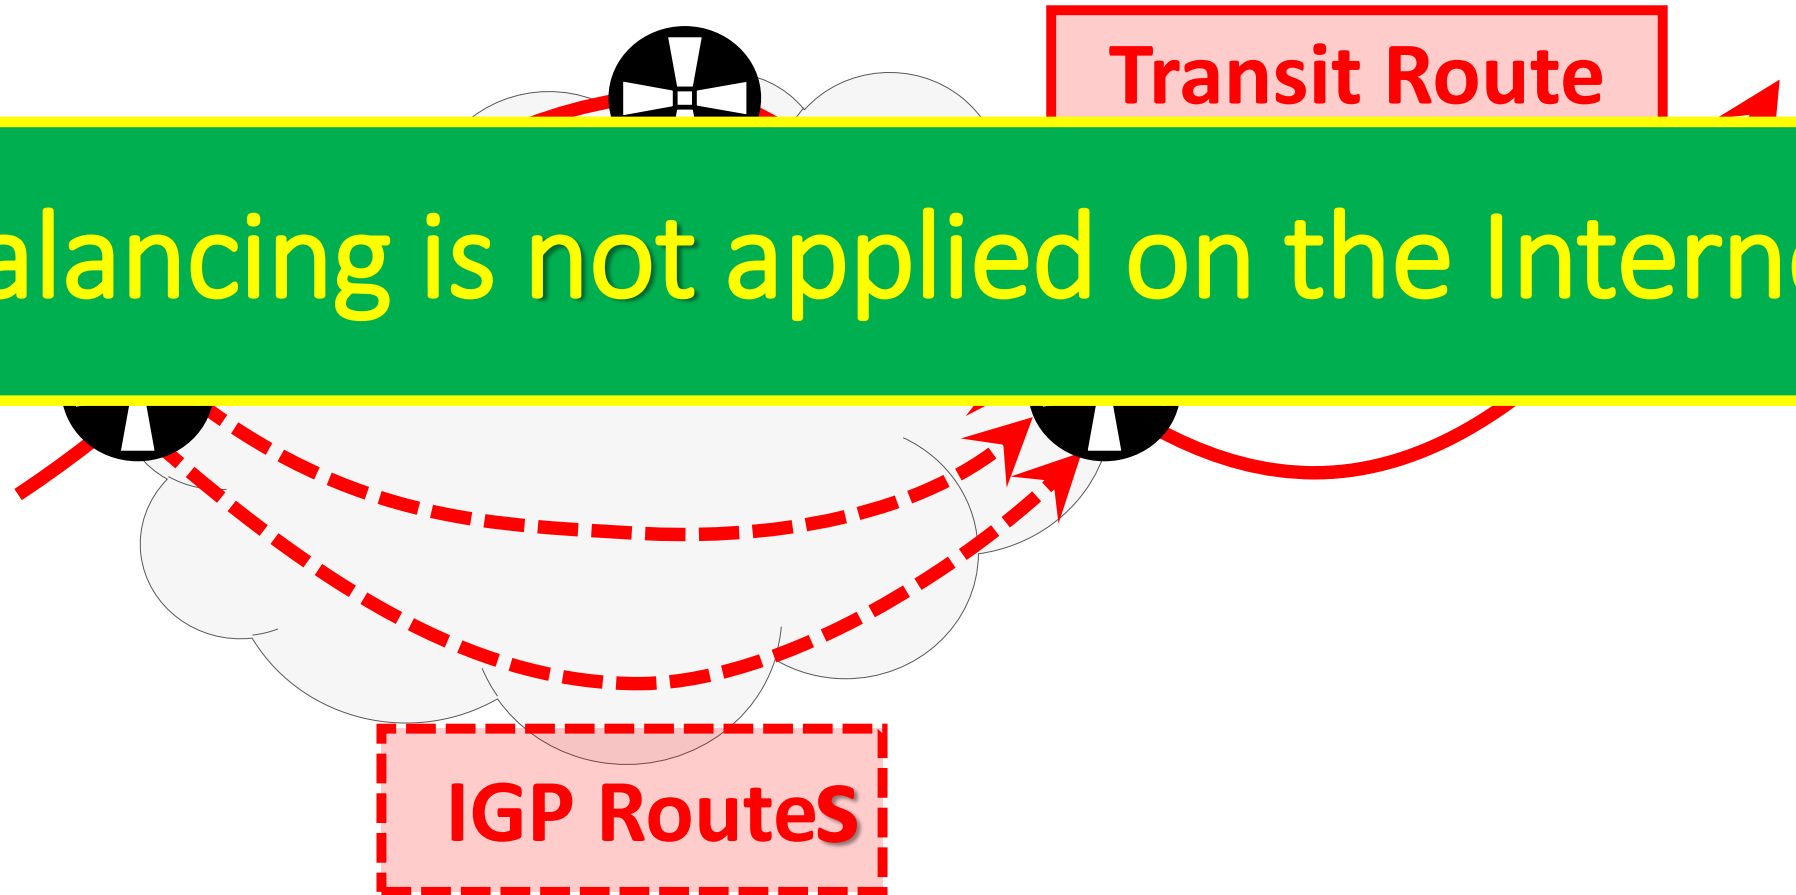
Detect all IGP routes with MDA-traceroute




How to Overcome ECMP?

Detect all IGP routes with MDA-traceroute

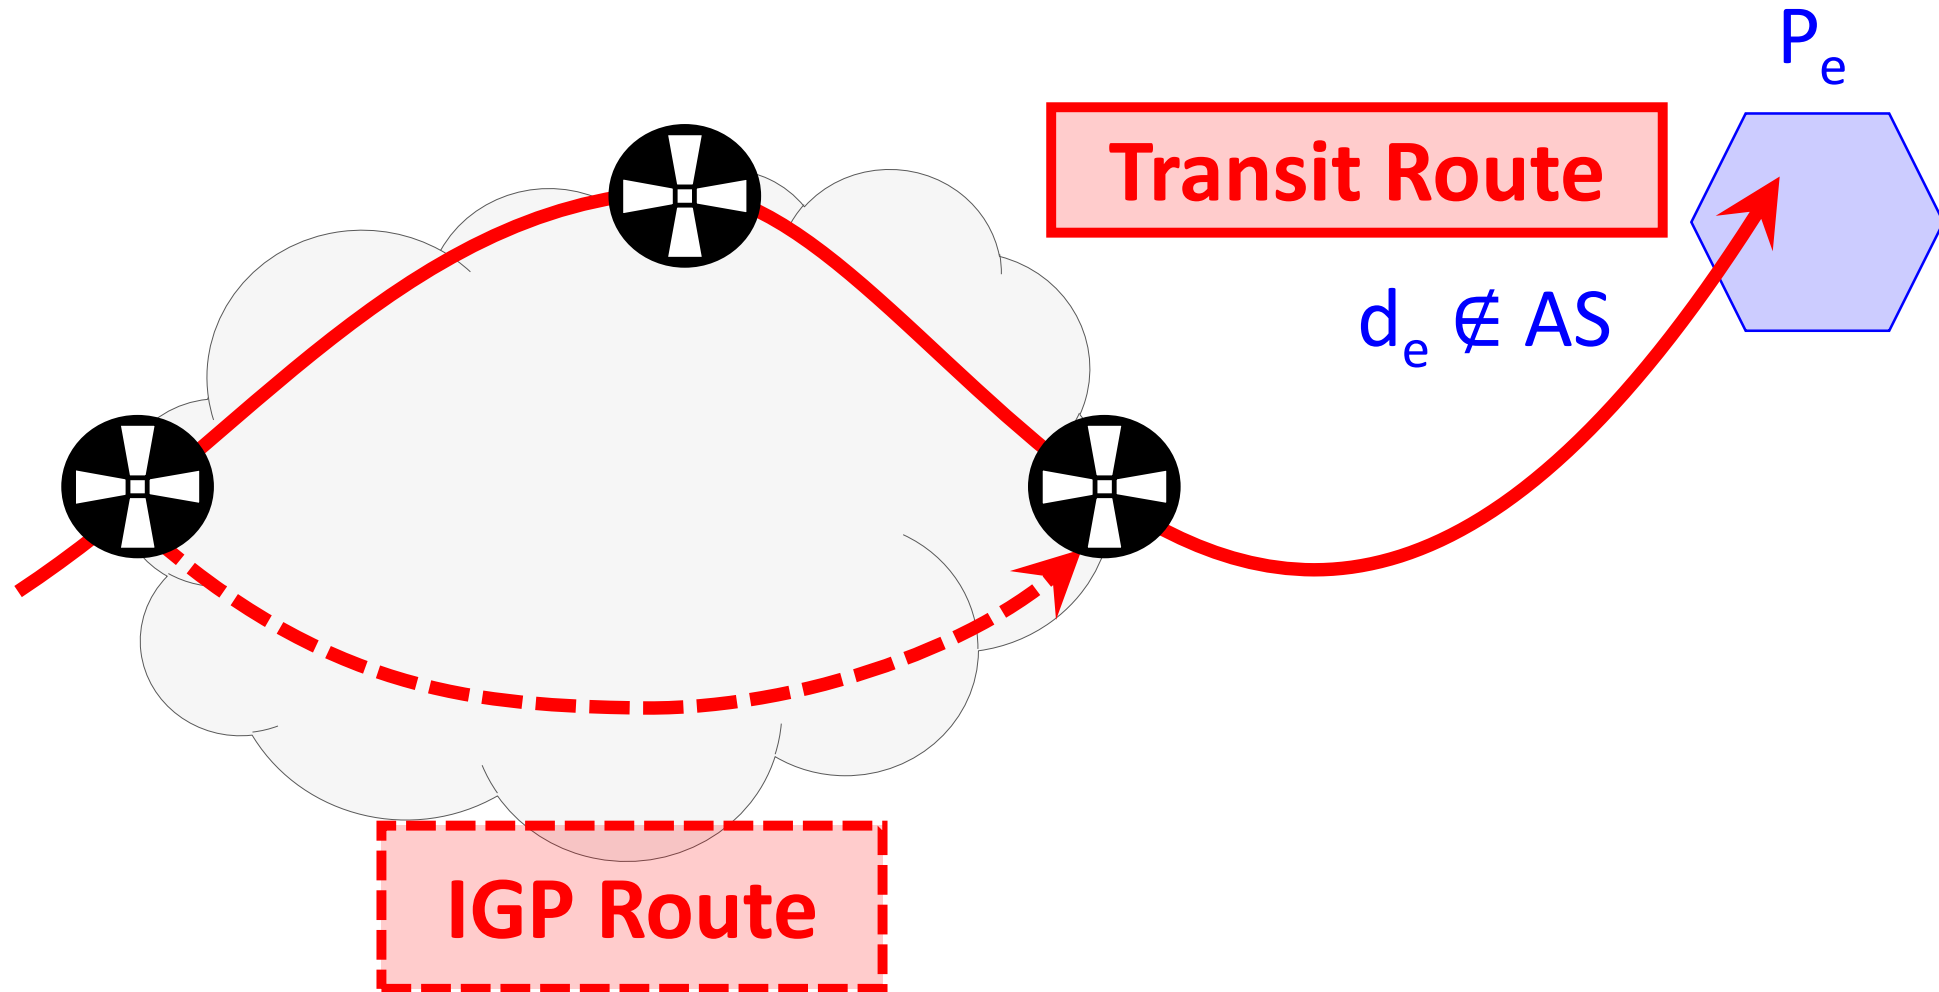
L4 load-balancing is not applied on the Internet



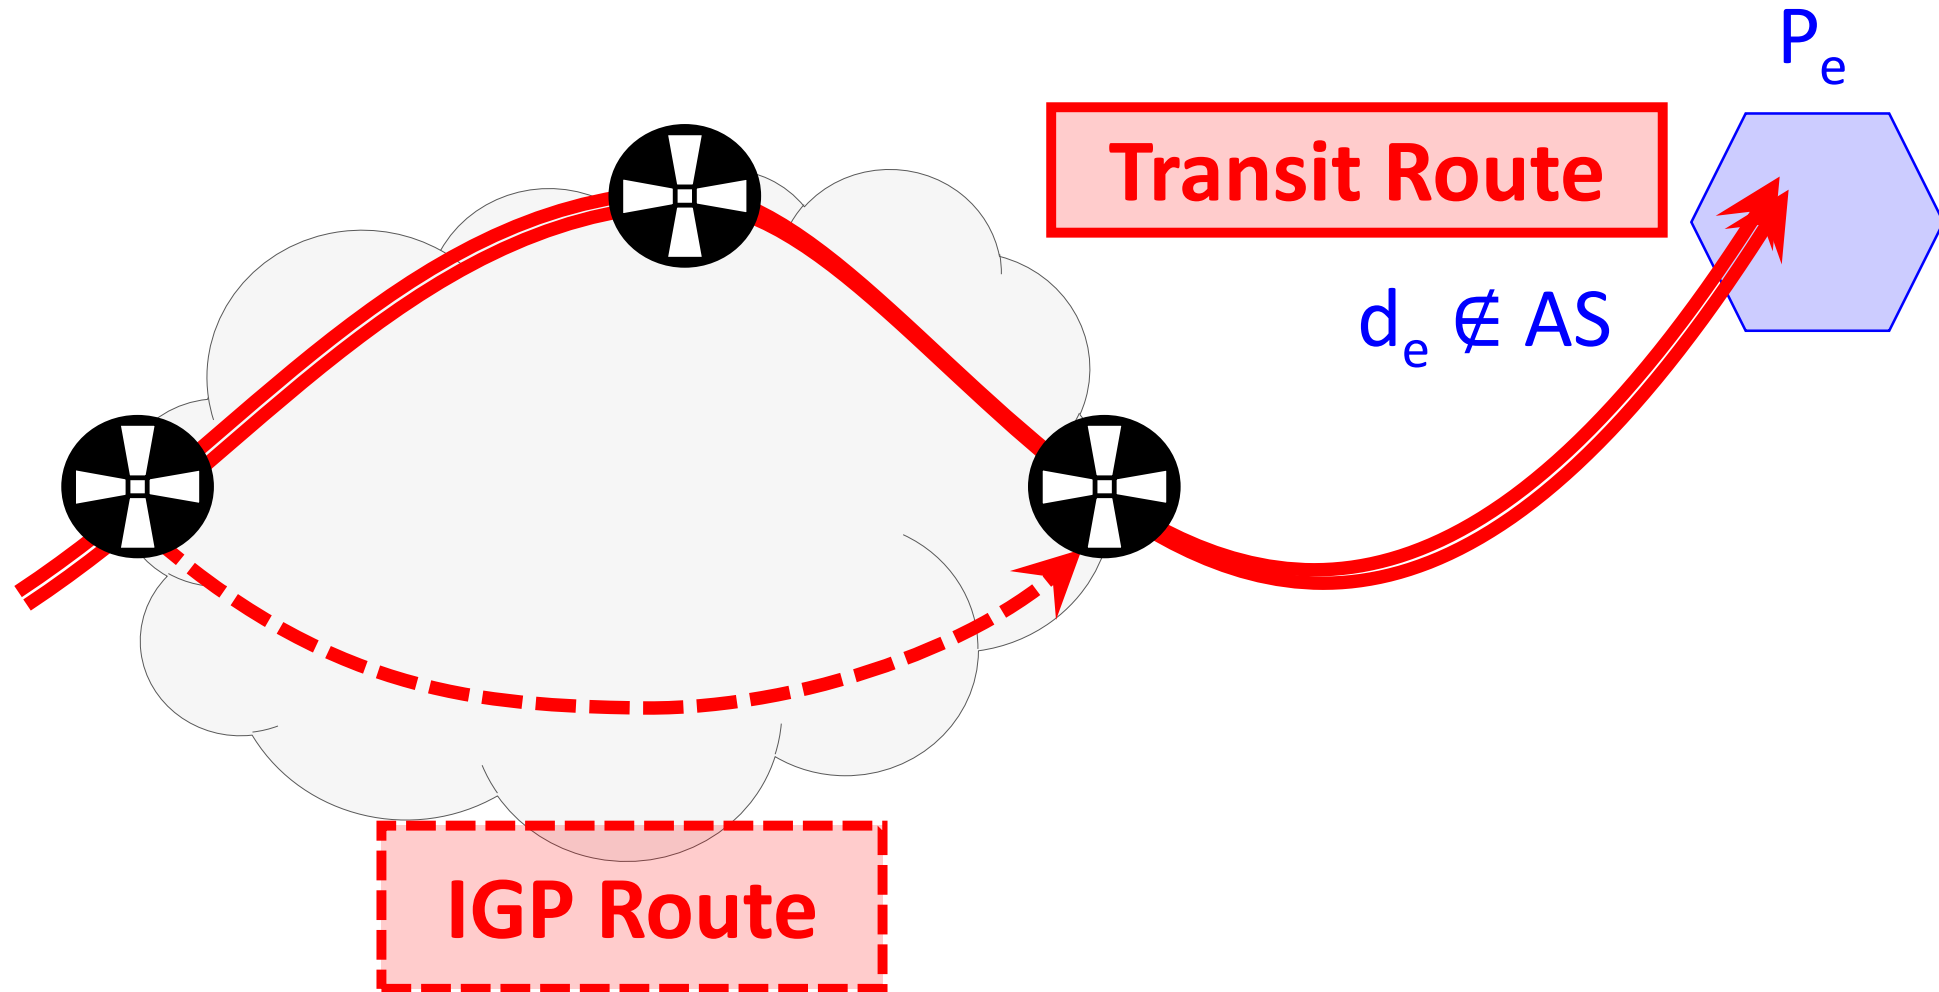
Per-Destination MDA-Traceroute




Per-Destination MDA-Traceroute



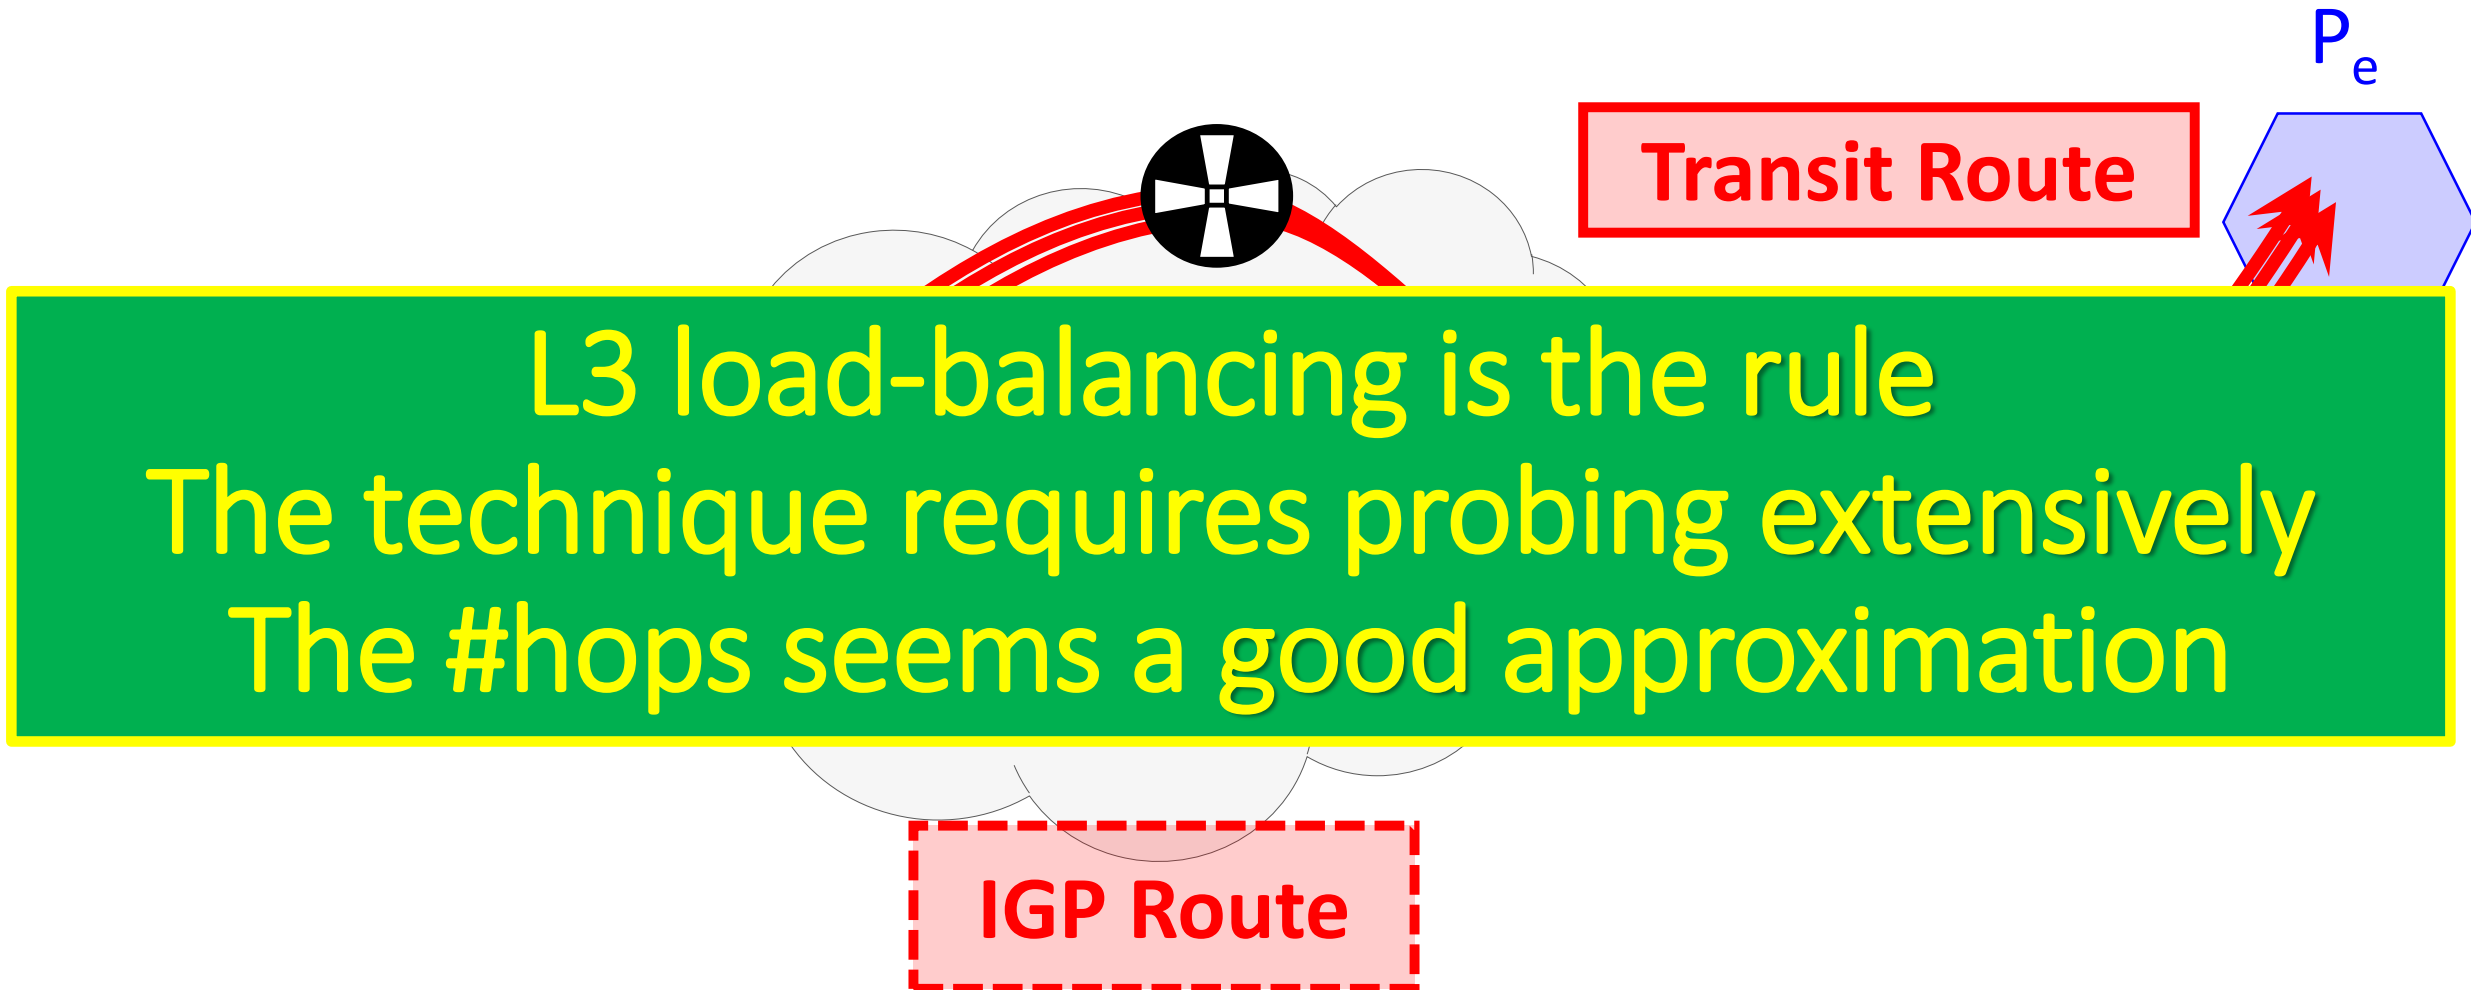
Per-Destination MDA-Traceroute



Per-Destination MDA-Traceroute



Per-Destination MDA-Traceroute



Preliminary Results

- 23-25% of 85 ASes raise a RI alarm
- #tunnels metric not so useful

N	ASN	#transit traces	#hops (%)	#tunnels (%)	OR (%)
1	20773	28	100	0.00	100
2	12965	445	61.35	35.73	66.74
3	3491	18.655	40.20	0.00	40.20
4	174	148.308	9.05	0.00	9.05
5	1299	106.421	2.82	1.81	3.00

- Many ASes deploy MPLS
- Each VP allows to measure “well” up to 3 ASes.
- It is not so clear that $\text{len(Transit)} > \text{len(IGP)}$

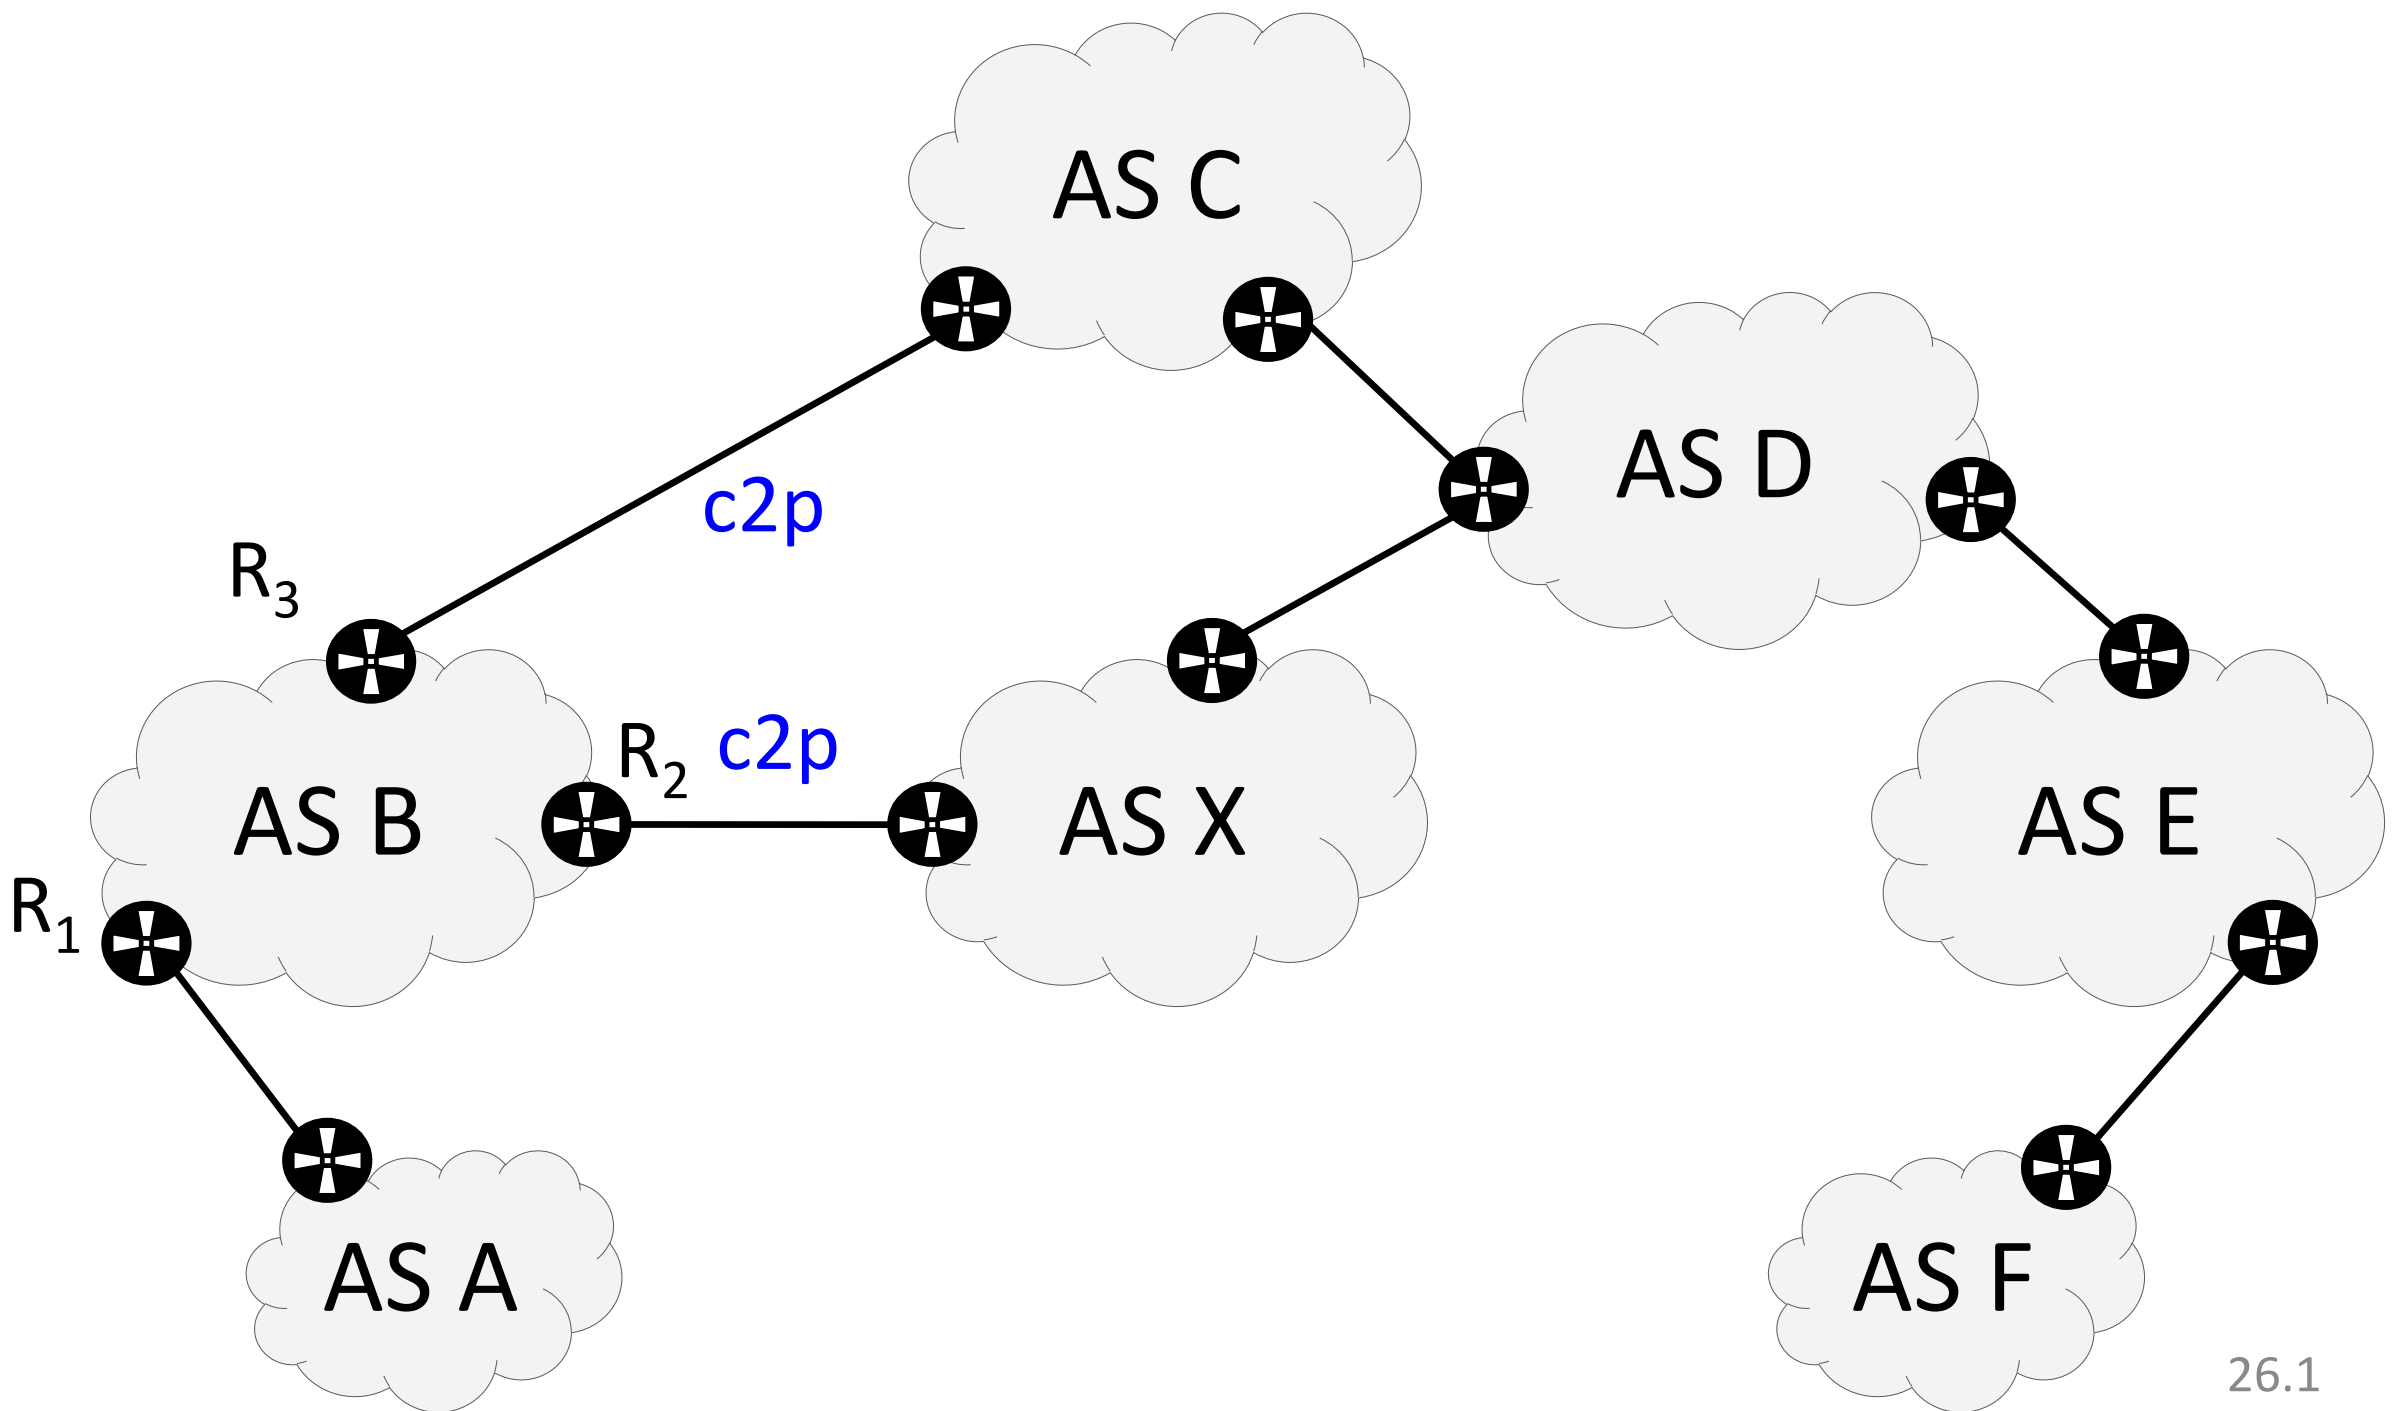
Journal...?

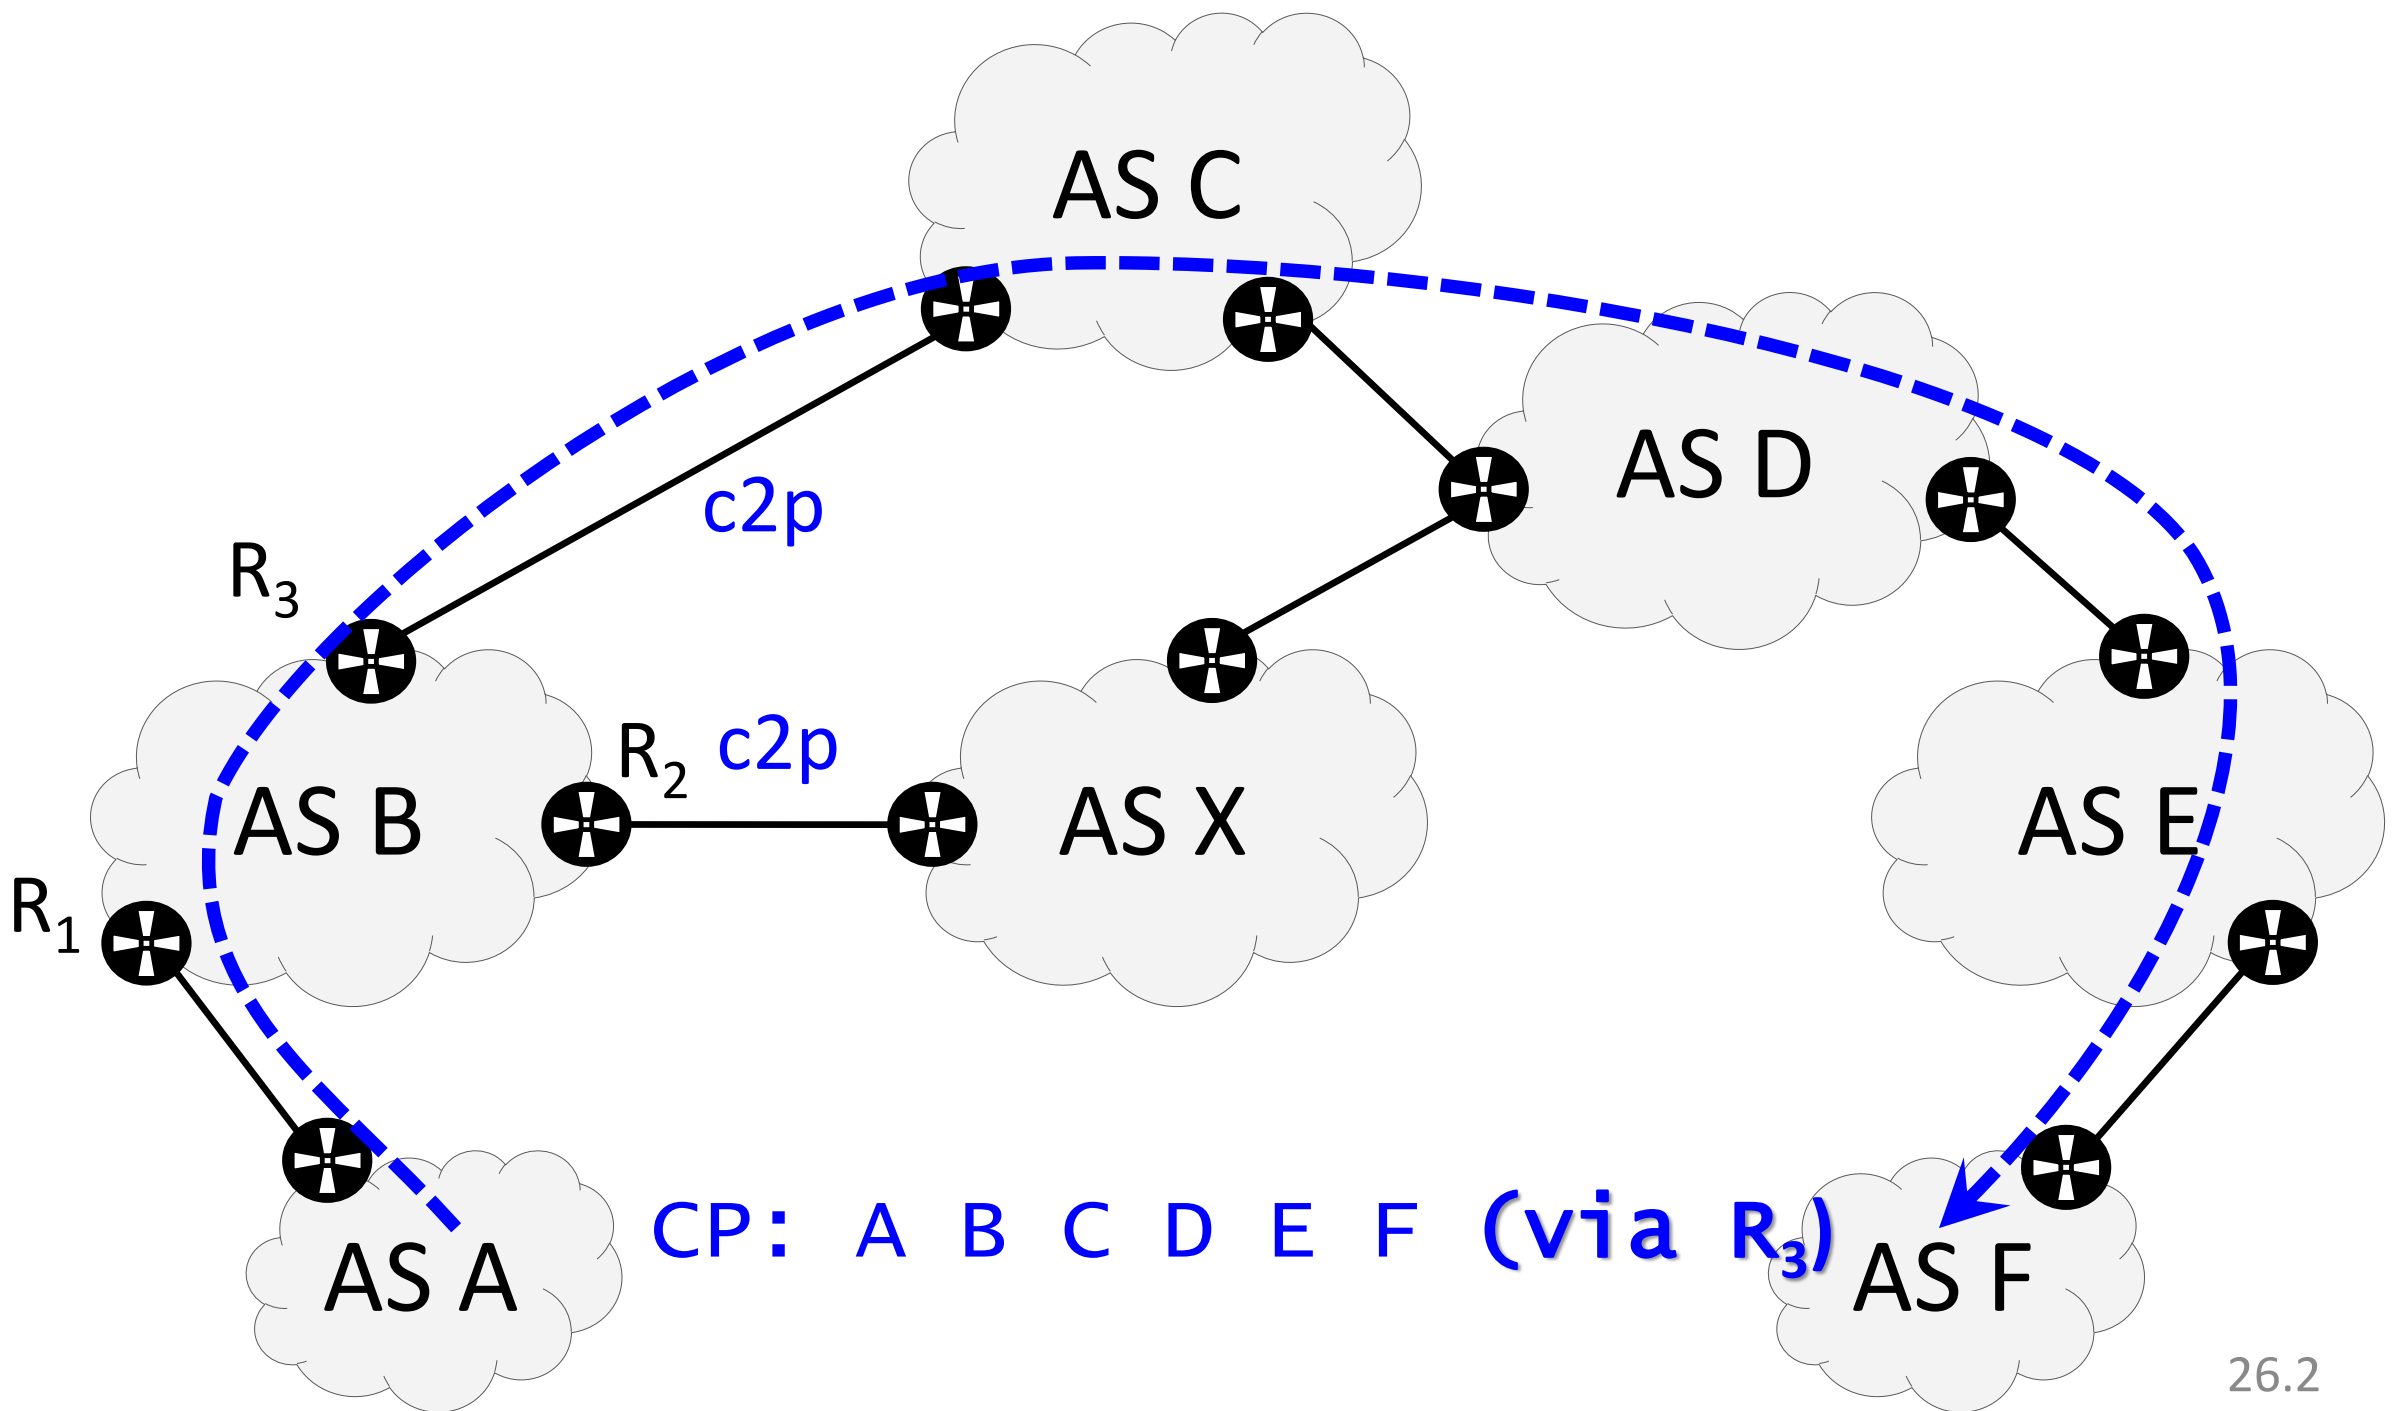
Coming soon


...


Rles might generate Inter-Domain Lies

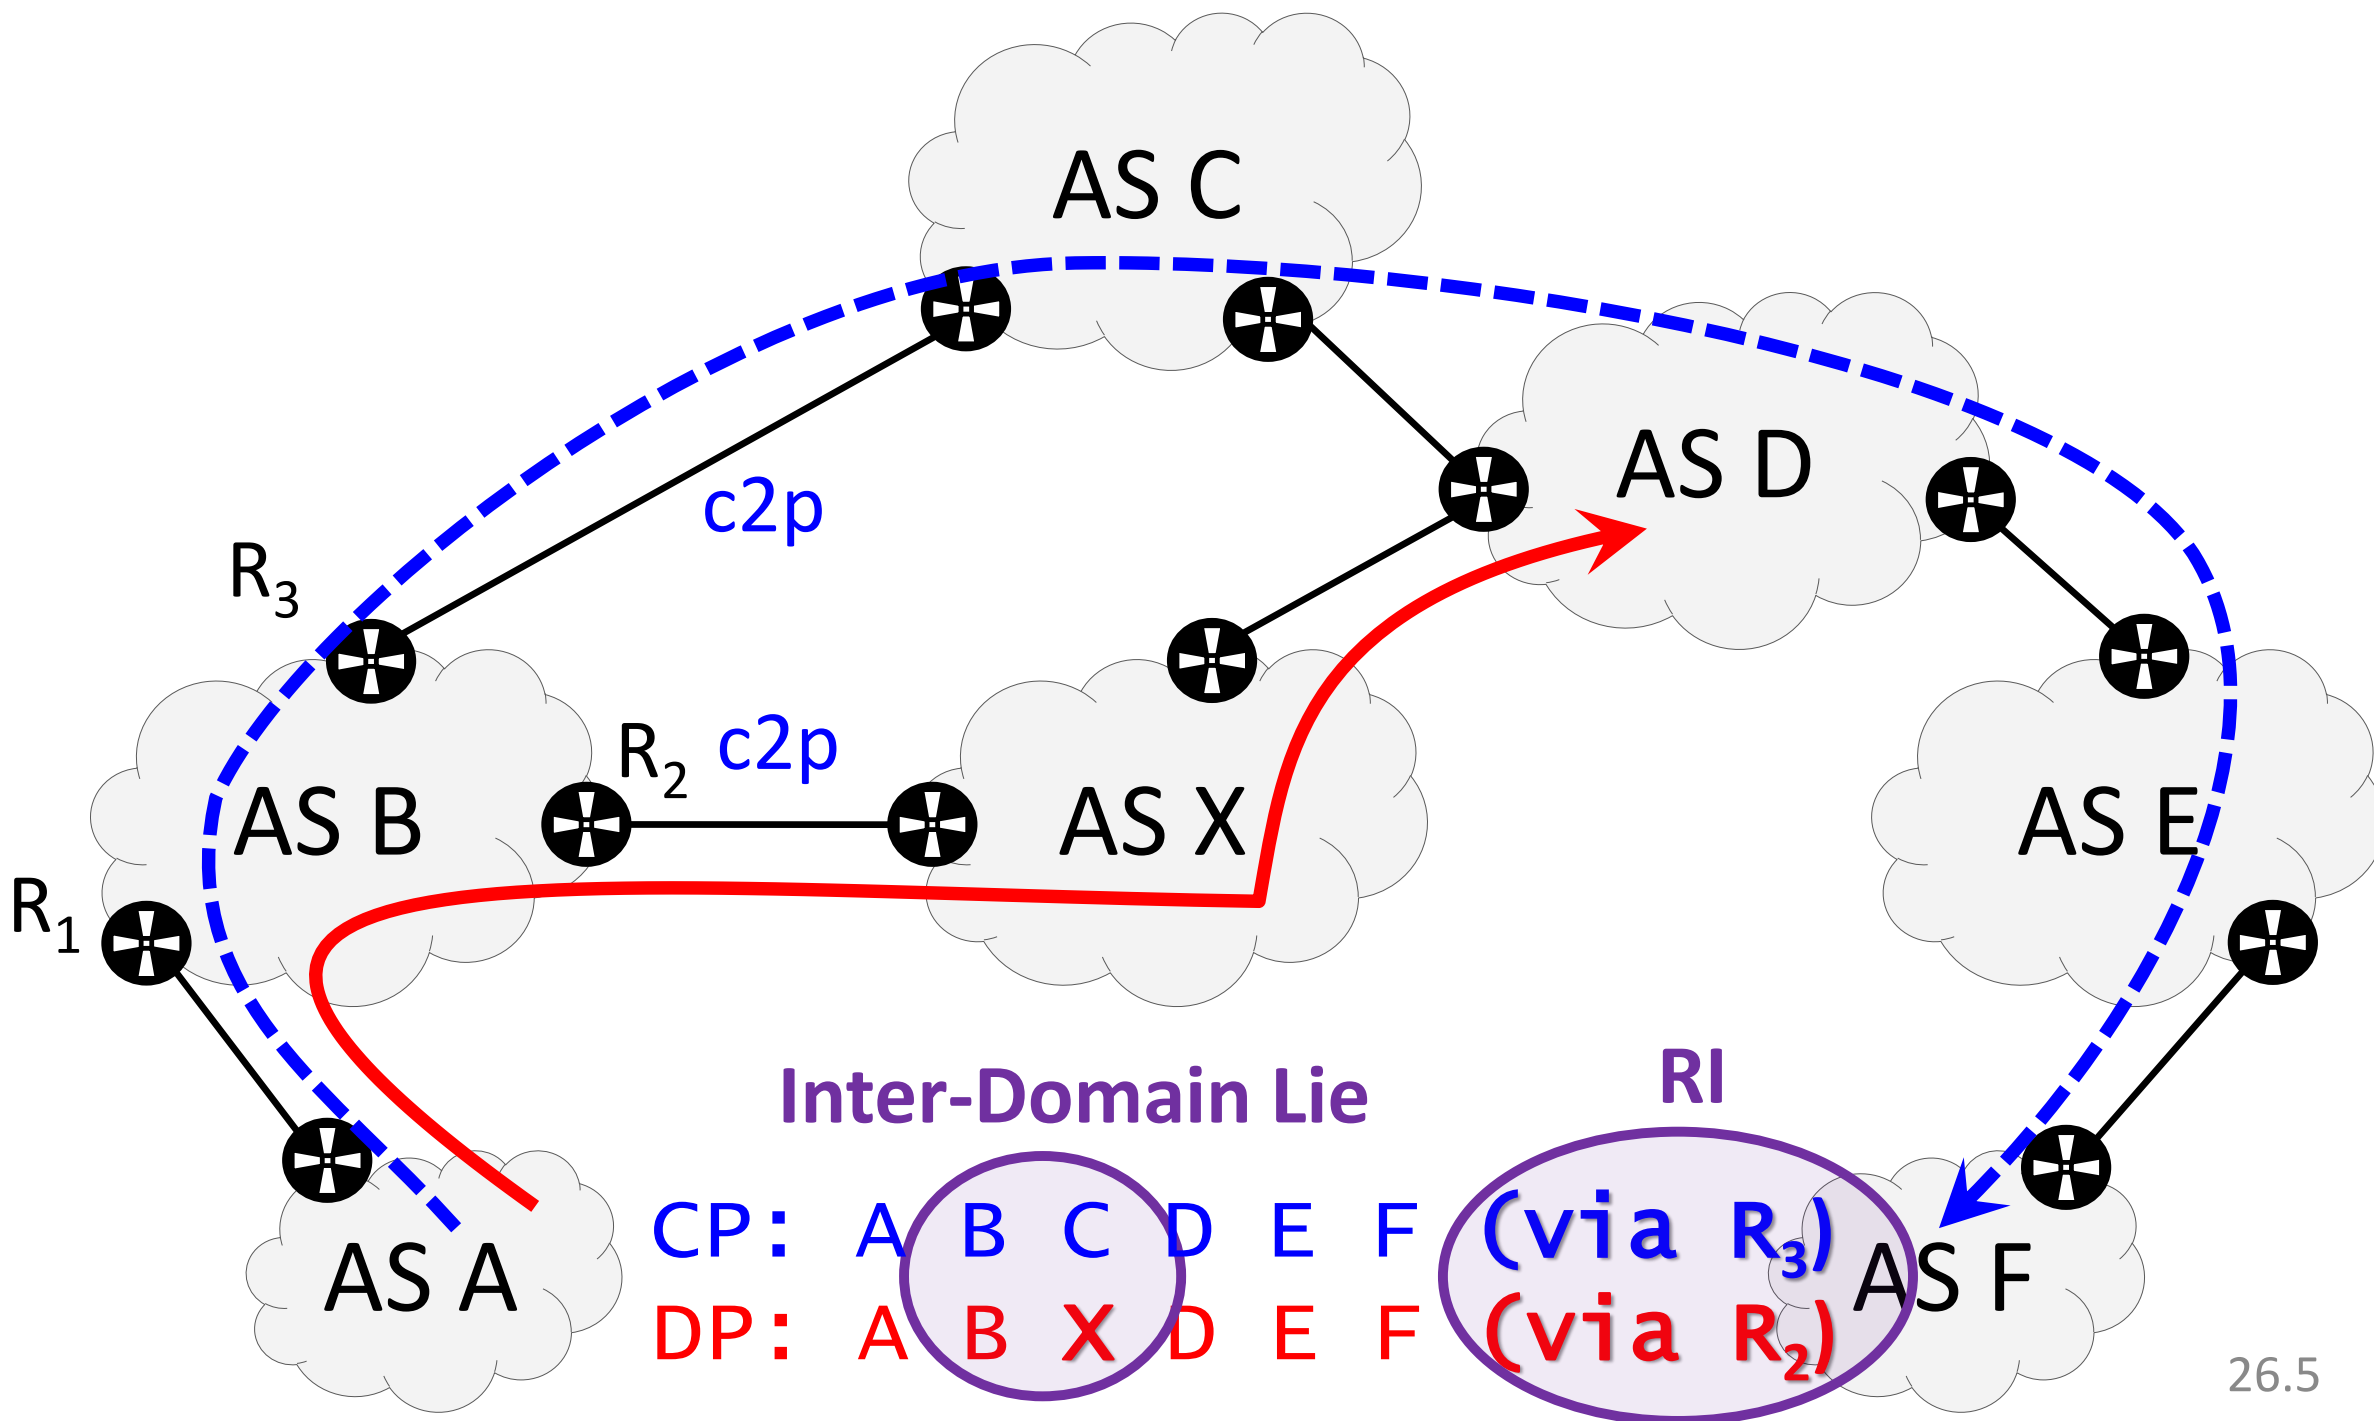
...











Work Extension

1. How to **discriminate** TE
2. **Pinpoint** the deflection point
3. Detect the **type** of RI (p-FIBs, ...)
4. Determine if lie is **deliberate** or not
5. **RTT** Analysis: High vs Low TTL

Final Goal

Security Mechanism/Protocol

- Consider **sophisticated** liars:
 - Traffic vs Traceroute DPs
 - Liars cannot “hide” in the noise

The work of the PhD

“Filtering the Noise to Reveal Inter-Domain Lies”, TMA 2019

University of Strasbourg/ICube, University of Napoli Federico II

“Routing Inconsistencies at the FIB level”, Under submission in ???

University of Strasbourg/ICube, University of Napoli Federico II

“A first Look at The Latin American IXPs”, Under submission in IMC 2019

University of Buenos Aires/CONICET, University of Strasbourg/ICube, University Diego Portales

64 hs lessons

Attended TMA PhD School 2018, 2019 (Presented Posters)

20 day internship in University of Napoli Federico II

Organized a Seminar on Cryptocurrencies

3-month Internship in Telefonica Research (coming soon)

The End Questions?