A Study of Prefix Hijacking and Interception in the Internet

Hitesh Ballani, Paul Francis and Xinyang Zhang
Cornell University

ACM SIGCOMM 2007

Prefix Hijacking

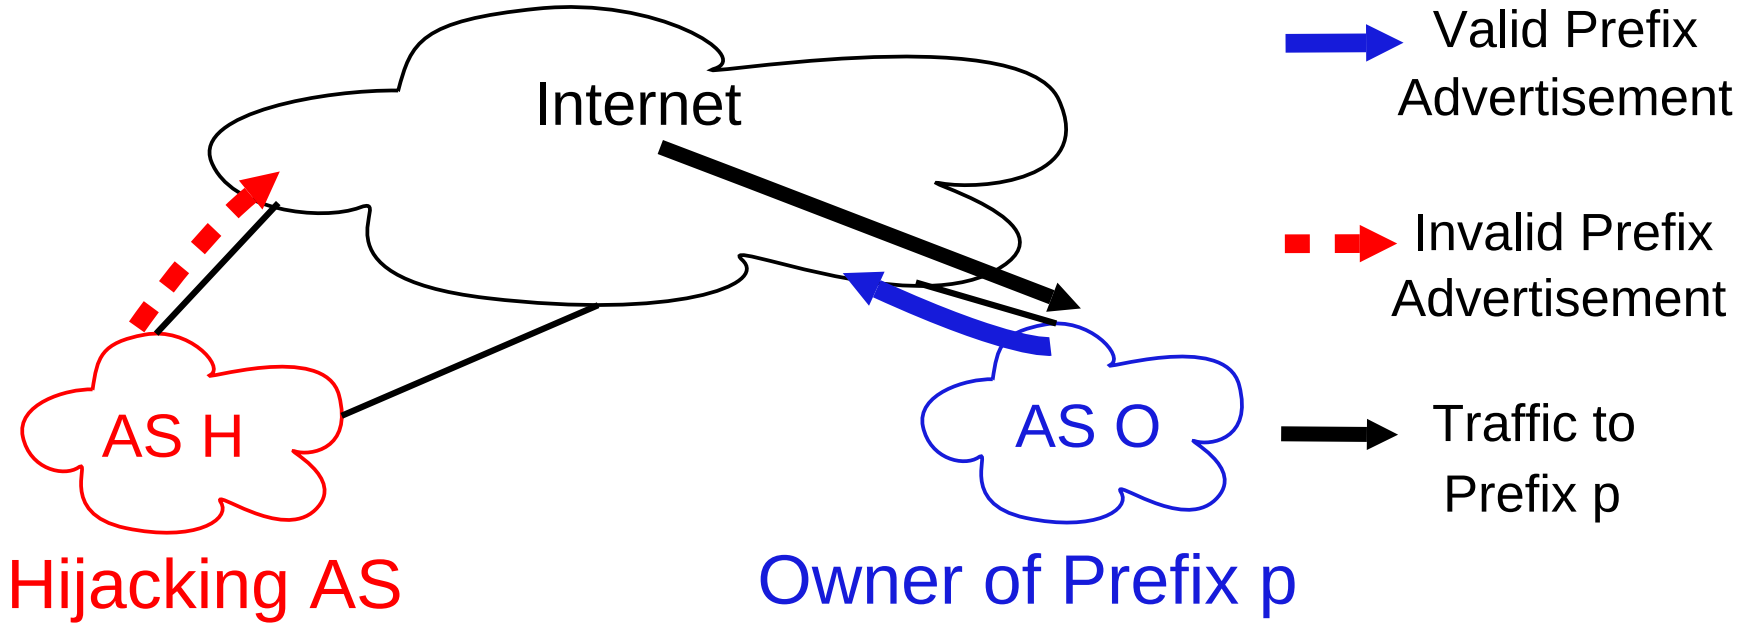

AS H advertizes a prefix owned by AS O

Prefix Hijacking

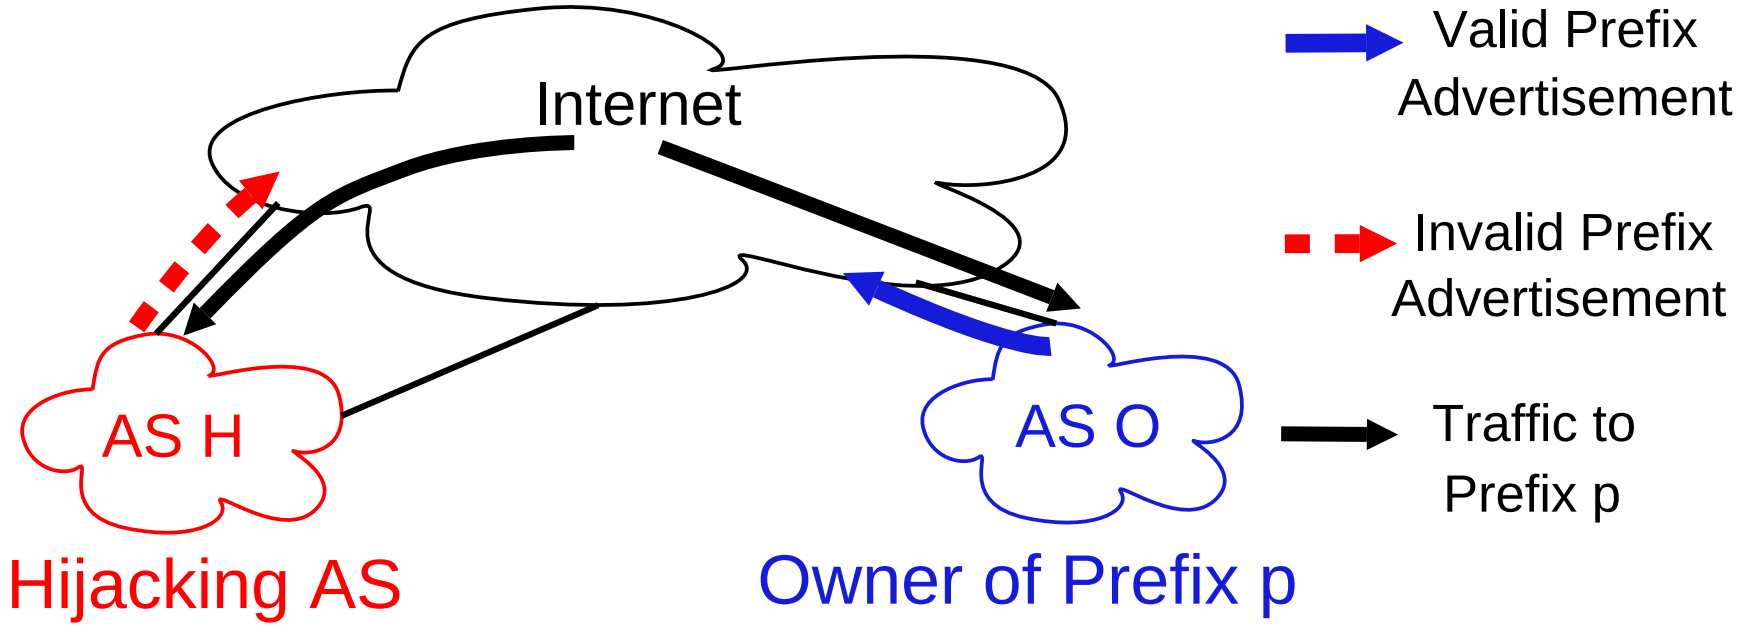

AS H advertizes a prefix owned by AS O

Fraction of traffic destined to the prefix is "hijacked"

Prefix Hijacking

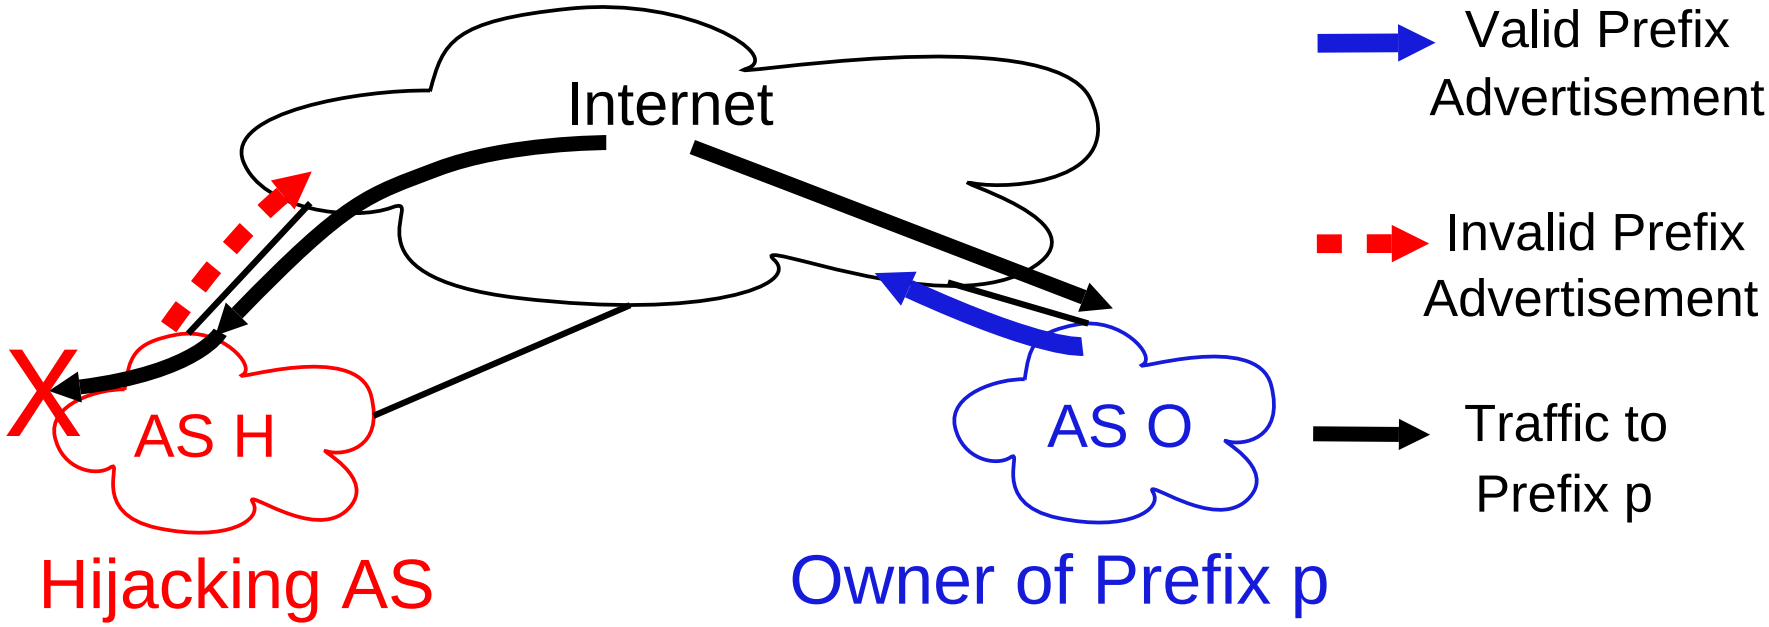

Hijacked traffic can be

- Blackholed
- Redirected
- Intercepted

Prefix Hijacking

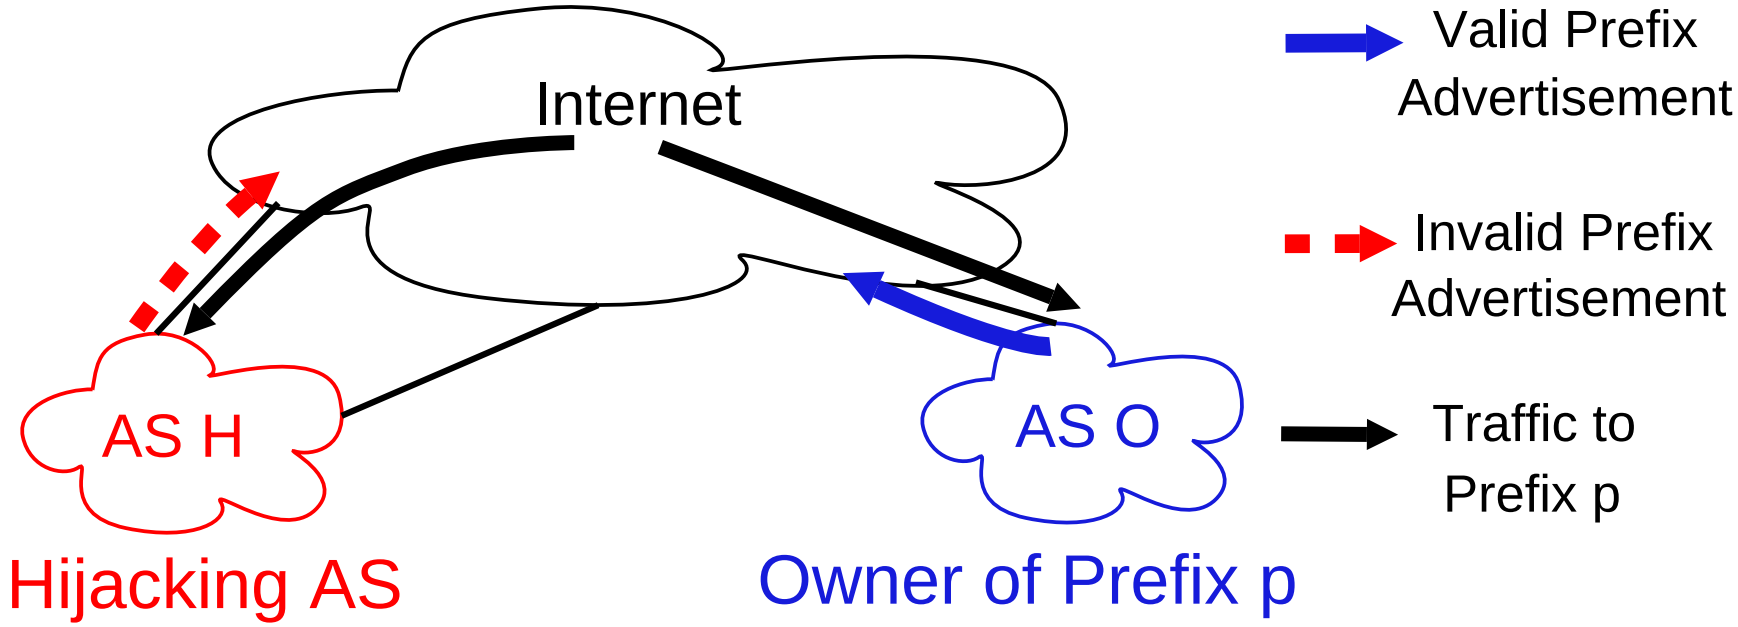

Hijacked traffic can be

- Blackholed
- Redirected
- Intercepted

Prefix Hijacking

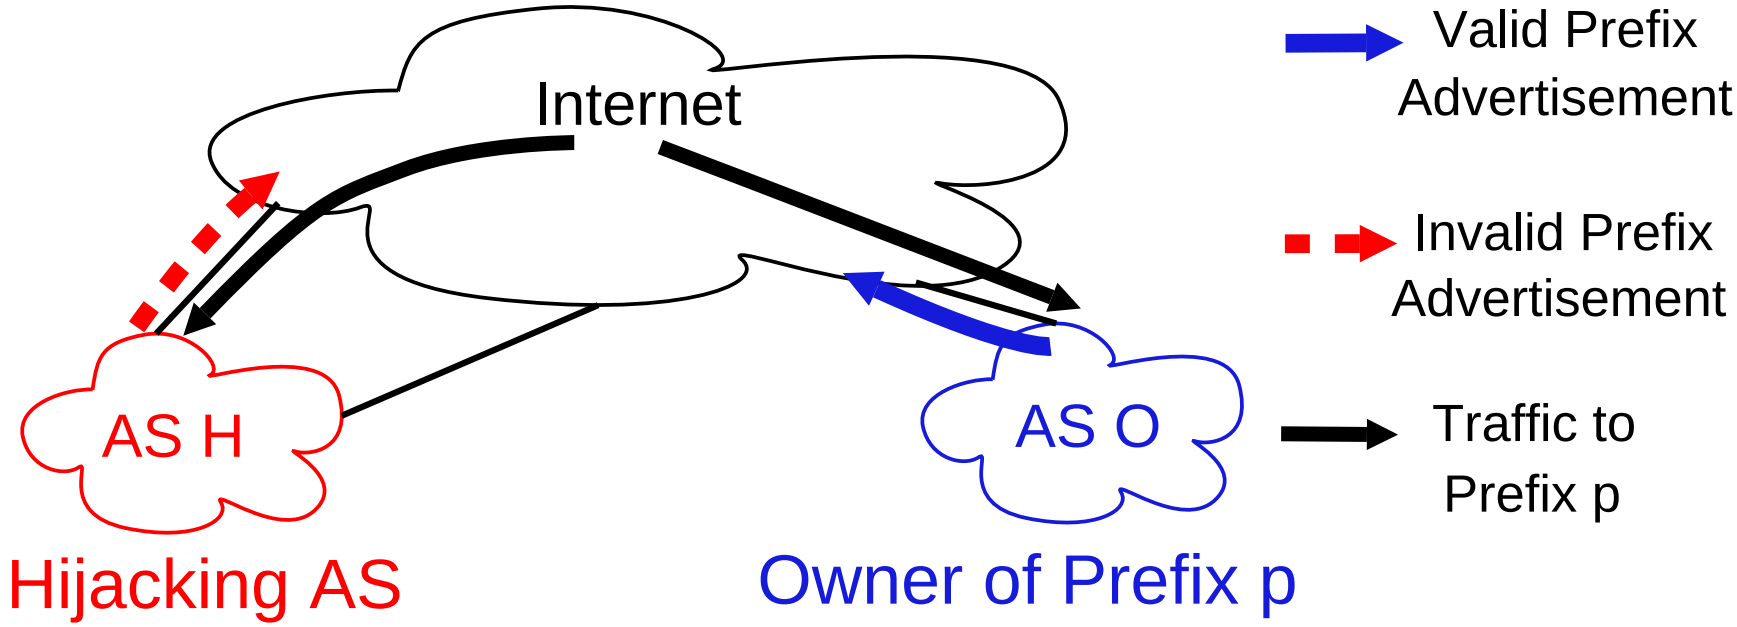

Hijacked traffic can be

- Blackholed
- Redirected

} Traffic does not reach destination

- Intercepted

Prefix Hijacking

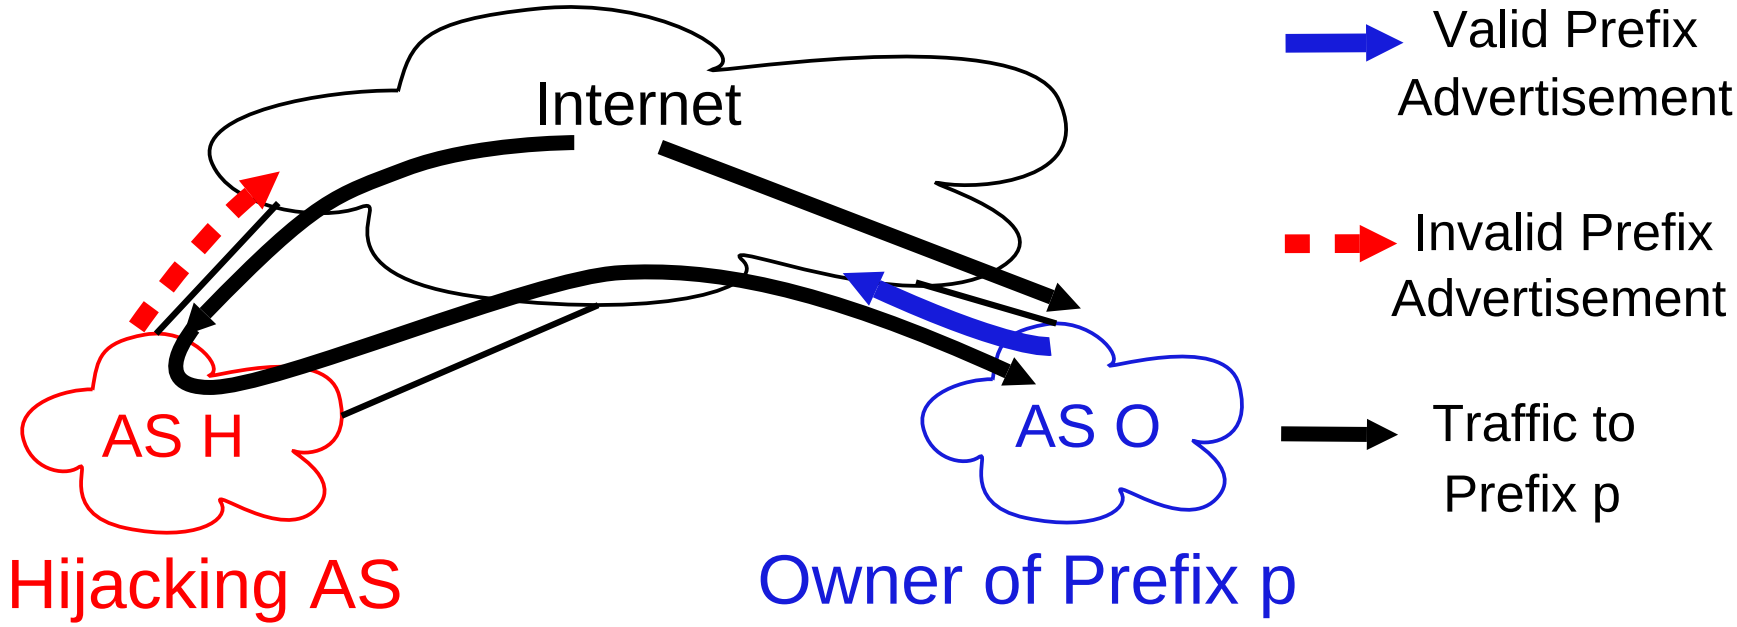

Hijacked traffic can be

- Blackholed
- Redirected

} Traffic does not reach destination

- Intercepted

} Traffic reaches its destination

A **Hypothetical** Interception Scenario

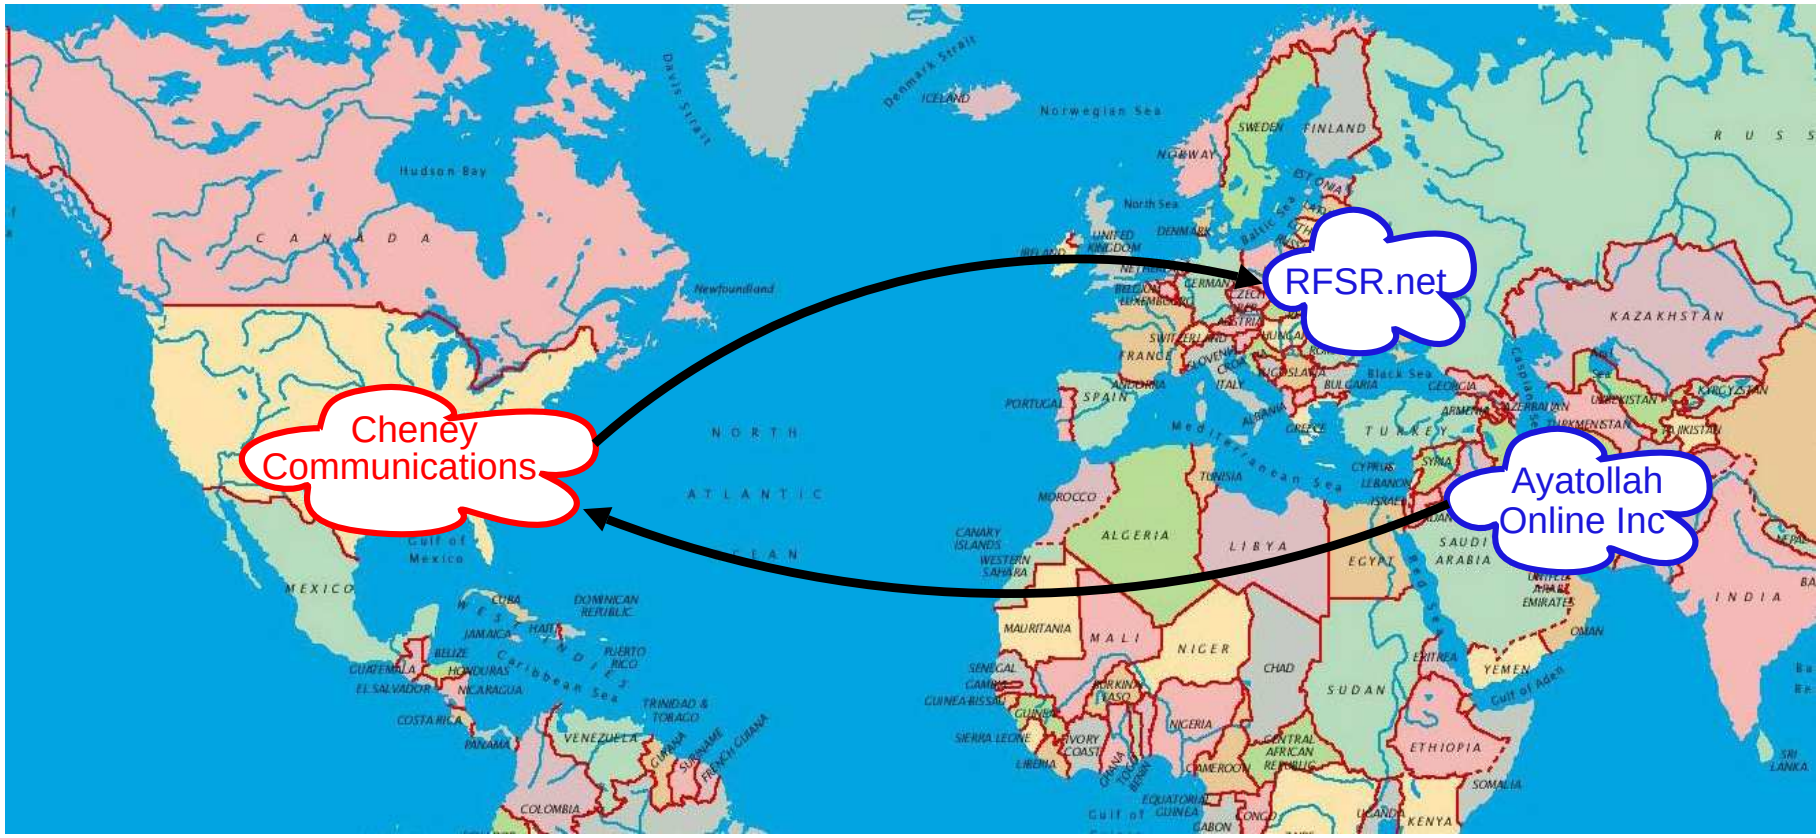

A **Hypothetical** Interception Scenario

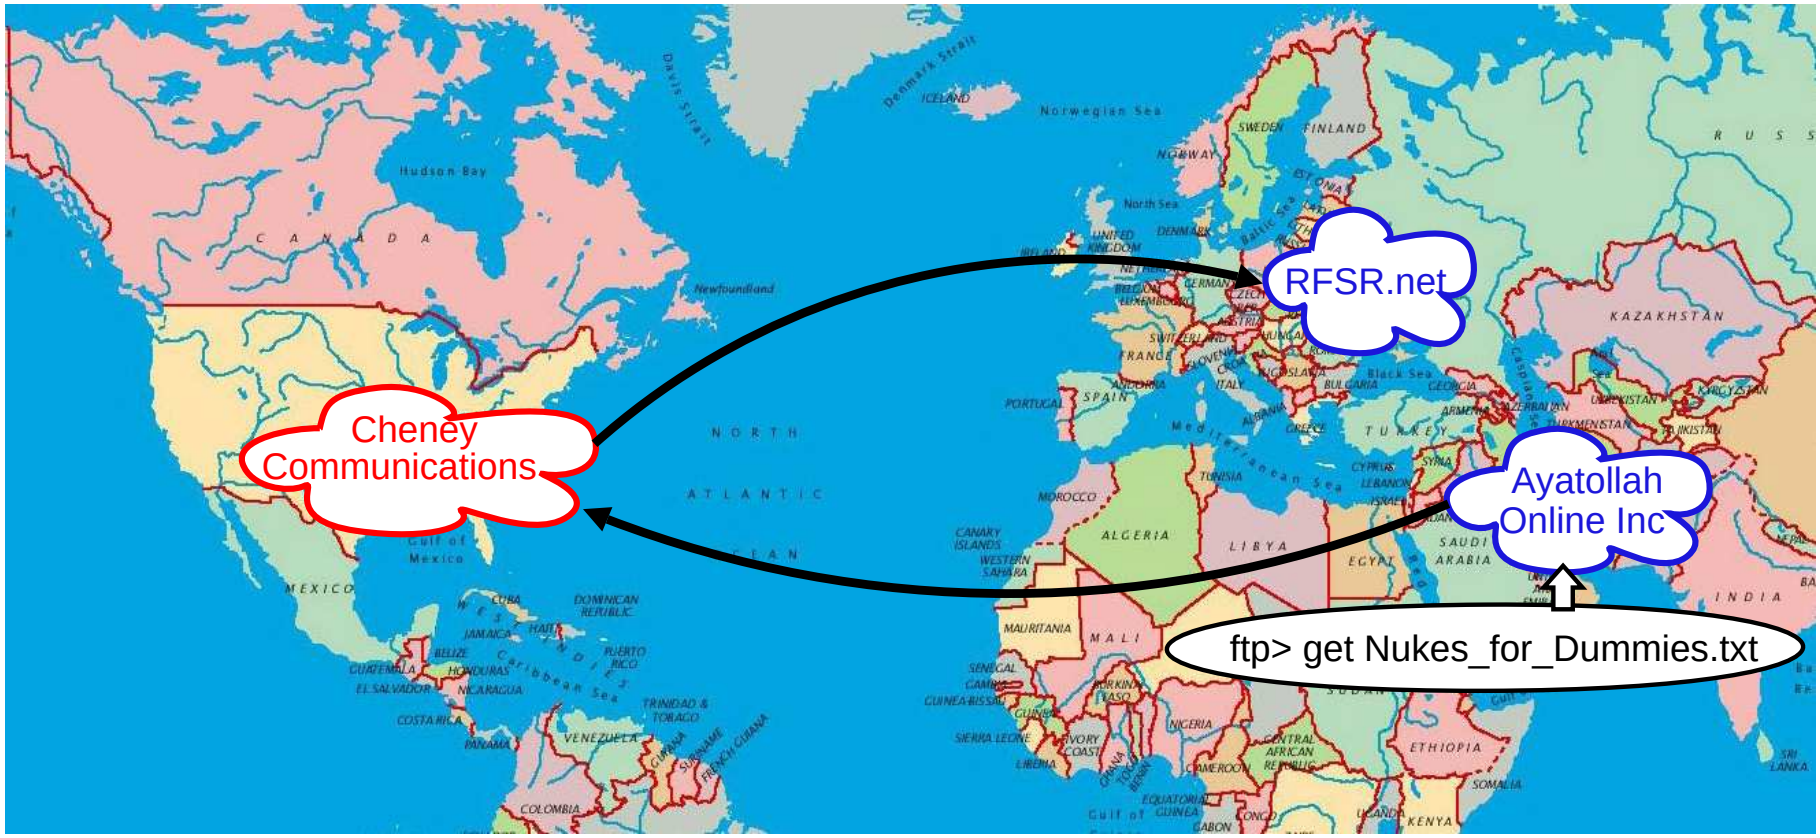

Prefix Hijack Puzzle

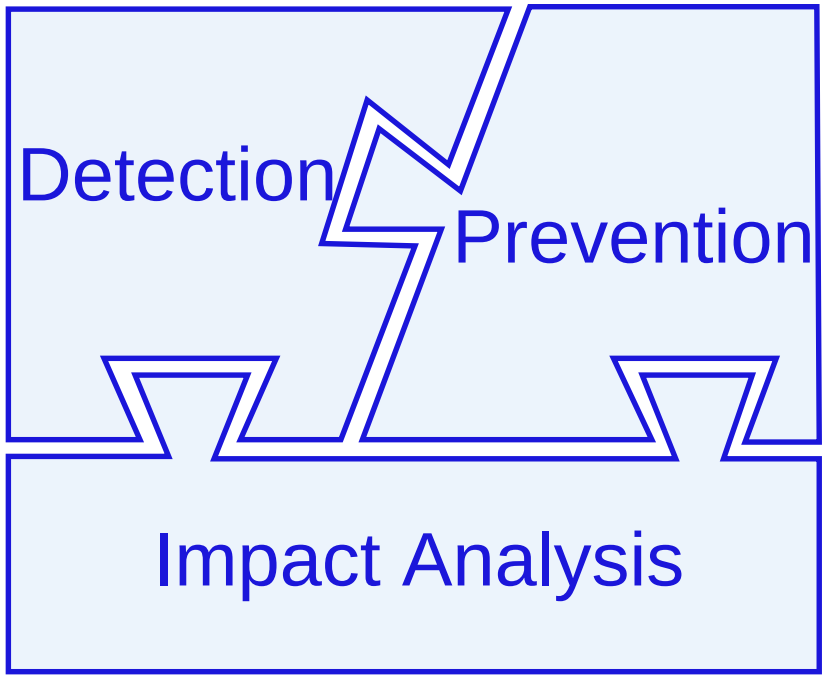

Prefix Hijack Puzzle

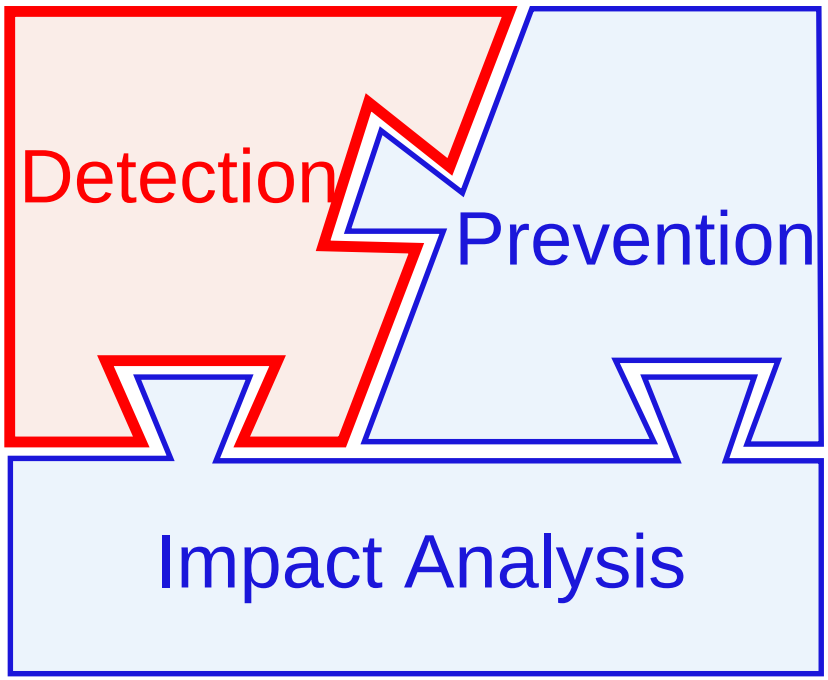

[RIPE MyASN]
[Kruegel et. al., LNCS'03]
[Teoh et. al., VDMCS'03]
[PHAS, Usenix Security'06]
[Hu et. al., IEEE Security'07]
[Zheng et. al, SIGCOMM'07]

Prefix Hijack Puzzle

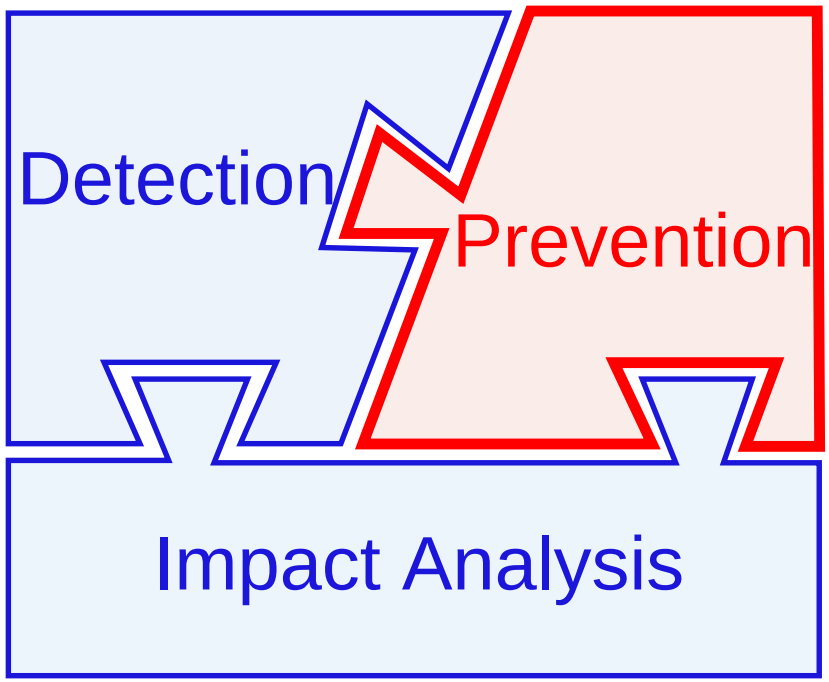

[Smith et. al., GI'96]
[S-BGP, JSAC'00]
[Zhao et. al., DSN'02]
[Wang et. al., ICDCS'03]
[Goodell et. al., NDSS'03]
[Aiello et. al., CCS'03]
[Subramanian et. al, NSDI'04]
[SPV, SIGCOMM'04]
[soBGP, Internet Draft'05]
[psBGP, NDSS'05]
[Karlin et. al., ICNP'06]

Prefix Hijack Puzzle

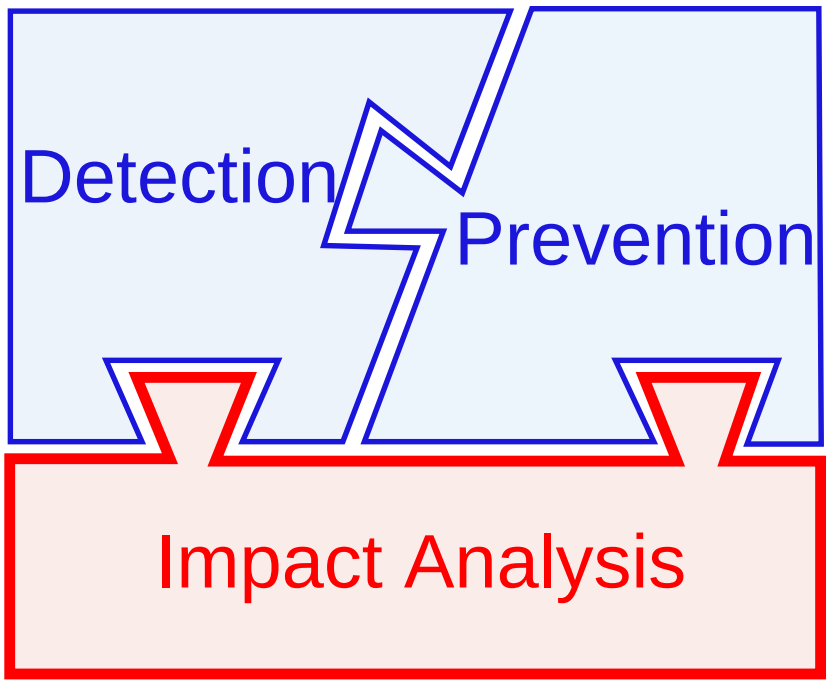

Quantification of the impact of prefix hijacks is sorely missing!

Prefix Hijacking and Interception: Unanswered Questions

What fraction of traffic can be hijacked and intercepted?

How can interception be achieved?

Is traffic on the Internet being intercepted?

Prefix Hijacking and Interception: Unanswered Questions

What fraction of traffic can be hijacked and intercepted?

- **Analyze** hijacking and interception probabilities
- **Estimate** probabilities for Route-Views ASes

How can interception be achieved?

- **Implement** interception methodology
- **Intercept** real traffic

Is traffic on the Internet being intercepted?

- **(Unsuccessful) Detection** Attempt

Talk Outline

- Introduction
- Hijacking Analysis
- Interception Analysis
- Hijacking and Interception estimates
- Hijacking and Intercepting real traffic
- Detecting Internet Interception
- Conclusions

Hijacking Analysis

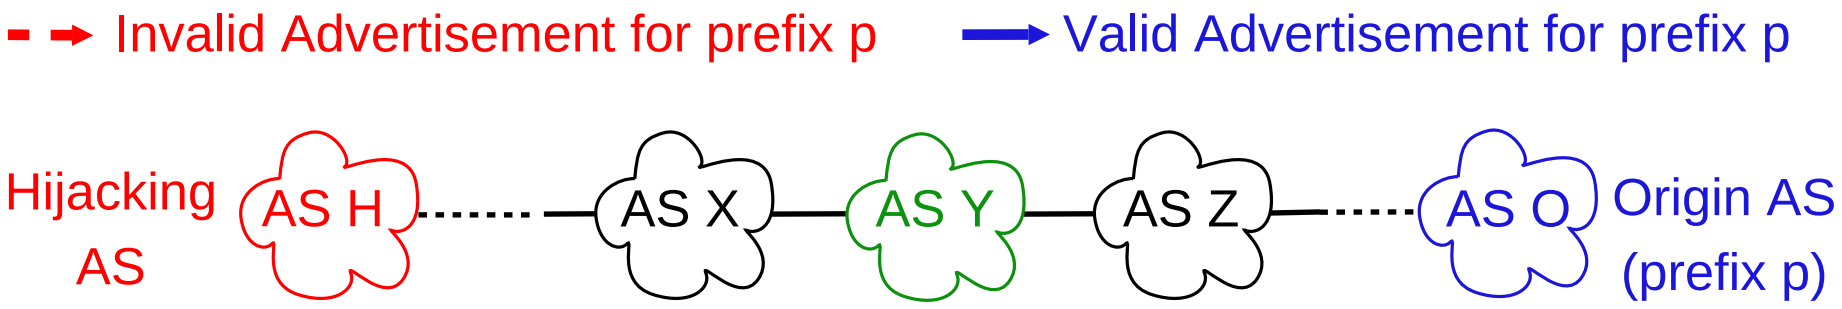

Hijacking Analysis

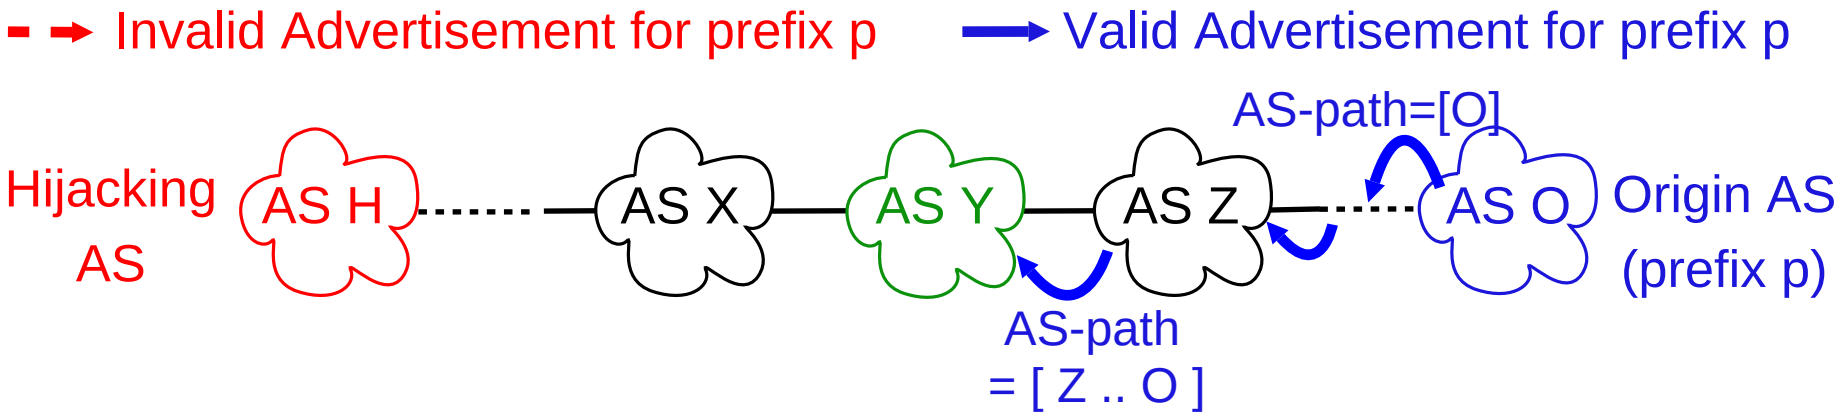

Hijacking Analysis

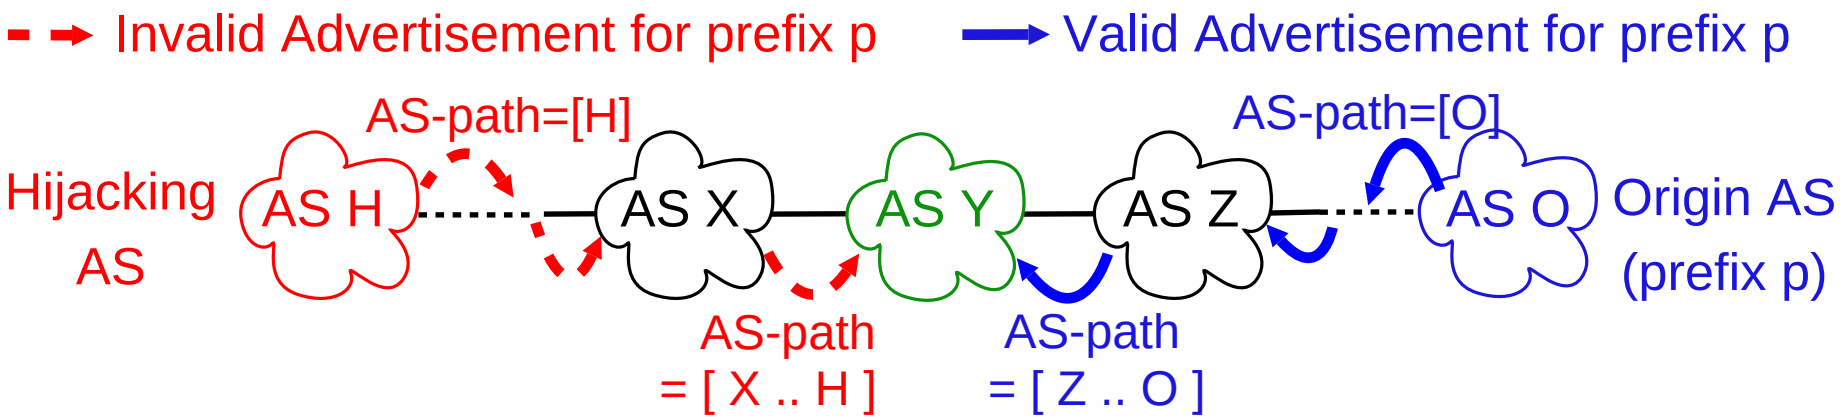

Hijacking Analysis

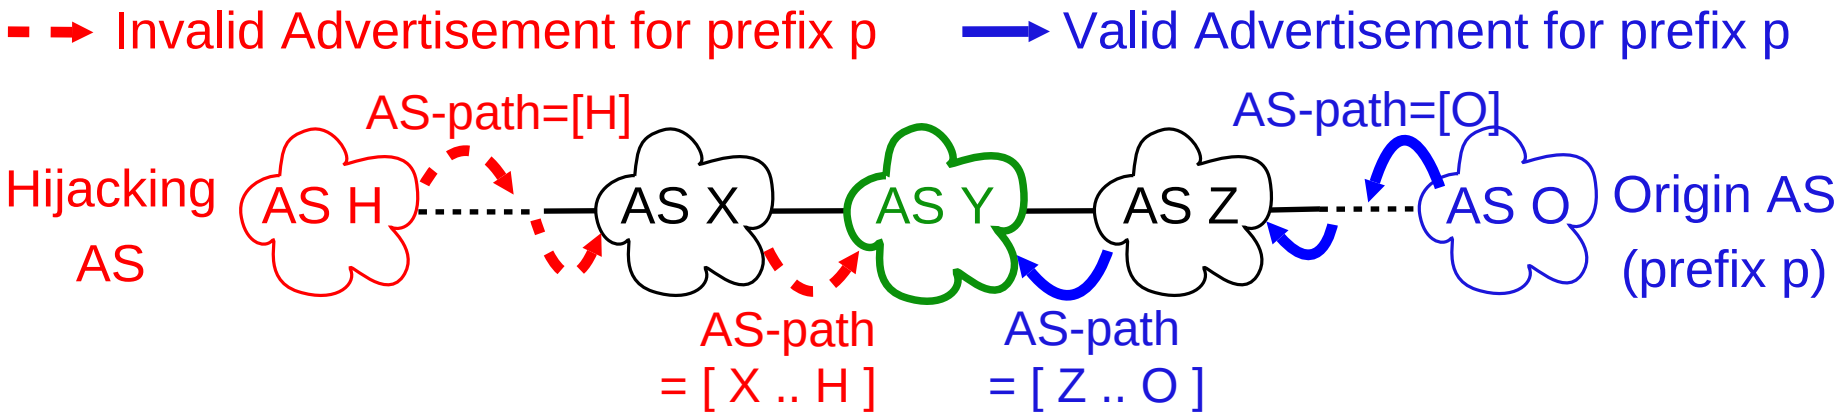

Can AS H hijack prefix *p*'s traffic from AS Y?

Hijacking Analysis

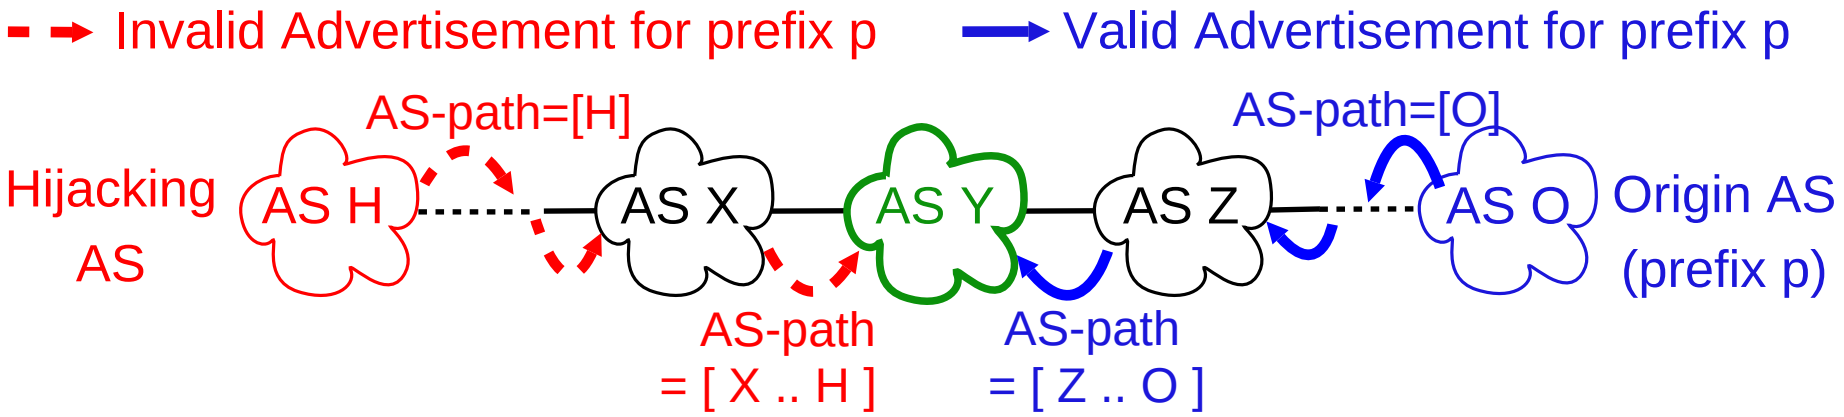

Can AS H hijack prefix p's traffic from AS Y?

AS Y needs to choose between

Invalid Route

AS-Path $= [X \ldots H]$

Length $= i$

Vs

Valid Route

AS-Path $= [Z \ldots O]$

Length $= v$

Hijacking Analysis

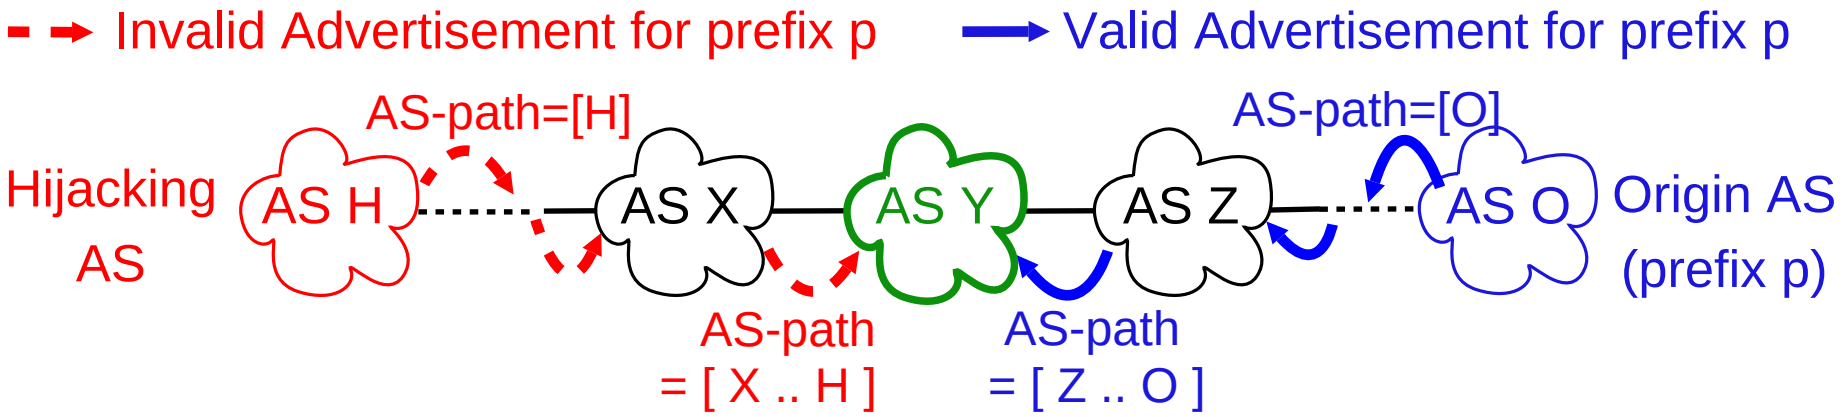

Can AS H hijack prefix p's traffic from AS Y?

AS Y needs to choose between

Invalid Route AS-Path $= [X \ldots H]$ Length $= i$

Vs

Valid Route AS-Path $= [Z \ldots O]$ Length $= v$

Assumption: AS Y has typical policies (customer $>$ peer $>$ provider)

Hijacking Analysis

Valid \ Invalid		Customer	Peer	Provider
Customer				
Peer				
Provider				

✗ : Valid route is chosen (traffic not hijacked)

✓ : Invalid route is chosen (traffic is hijacked)

Hijacking Analysis

Valid \ Invalid		Customer	Peer	Provider
Customer				
Peer				
Provider				

✗: Valid route is chosen (traffic not hijacked)

✓: Invalid route is chosen (traffic is hijacked)

Hijacking Analysis

Valid \ Invalid		Customer	Peer	Provider
Customer			✗	✗
Peer				
Provider				

✗: Valid route is chosen (traffic not hijacked)

✓: Invalid route is chosen (traffic is hijacked)

Hijacking Analysis

Invalid / Valid		Customer	Peer	Provider
Customer			✗	✗
Peer				
Provider				

✗: Valid route is chosen (traffic not hijacked)

✓: Invalid route is chosen (traffic is hijacked)

Hijacking Analysis

Valid \ Invalid		Customer	Peer	Provider
	v<i	✗		
Customer			✗	✗
Peer				
Provider				

✗: Valid route is chosen (traffic not hijacked)

✓: Invalid route is chosen (traffic is hijacked)

Hijacking Analysis

Valid \ Invalid		Customer	Peer	Provider
	v<i	✗		
Customer	v>i	✓	✗	✗
Peer				
Provider				

✗: Valid route is chosen (traffic not hijacked)

✓: Invalid route is chosen (traffic is hijacked)

Hijacking Analysis

Valid \ Invalid		Customer	Peer	Provider
	v<i	✗		
Customer	v>i	✓	✗	✗
	v=i	--		
Peer				
Provider				

✗: Valid route is chosen (traffic not hijacked)

✓: Invalid route is chosen (traffic is hijacked)

Hijacking Analysis

Valid \ Invalid		Customer	Peer	Provider
Customer	v<i	✗	✗	✗
	v>i	✓		
	v=i	--		
Peer	v<i	✗	✗	✗
	v>i	✓	✓	
	v=i		--	
Provider	v<i	✗	✗	✗
	v>i	✓	✓	✓
	v=i			--

✗ : Valid route is chosen (traffic not hijacked)

✓ : Invalid route is chosen (traffic is hijacked)

Interception Analysis

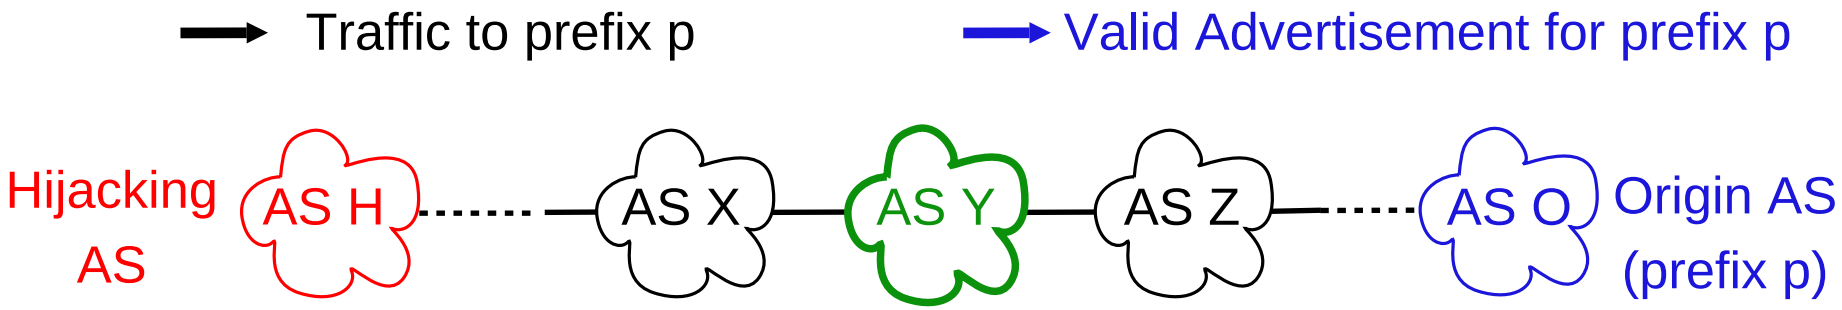

Can AS H intercept prefix *p*'s traffic from AS Y?

Interception Analysis

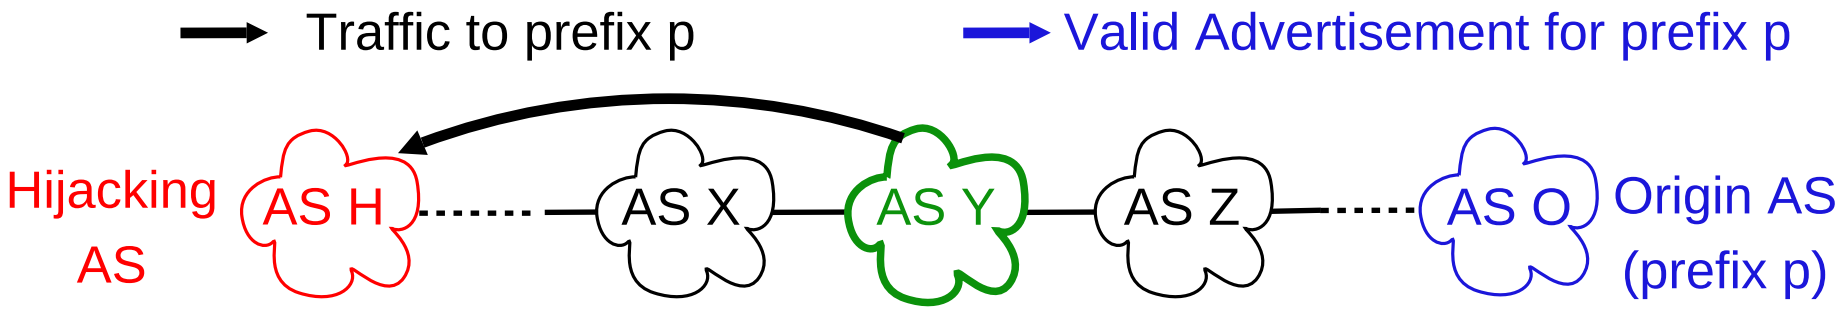

Can AS H intercept prefix *p*'s traffic from AS Y?

1. Can AS H hijack prefix *p*'s traffic from AS Y?
2. Can AS H route the hijacked traffic to back AS O?

Interception Analysis

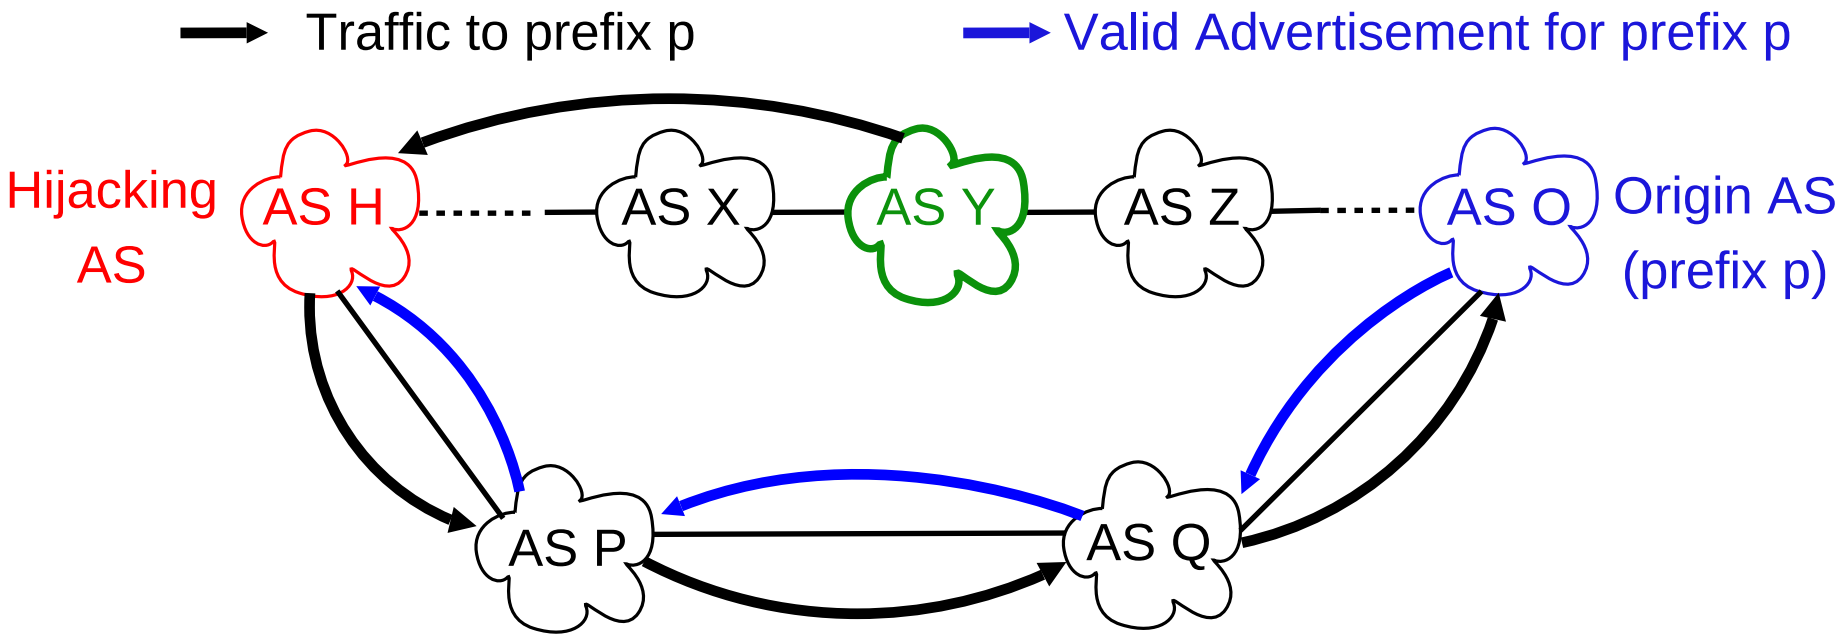

Can AS H intercept prefix p's traffic from AS Y?

1. Can AS H hijack prefix p's traffic from AS Y?
2. Can AS H route the hijacked traffic to back AS O?

Interception Analysis

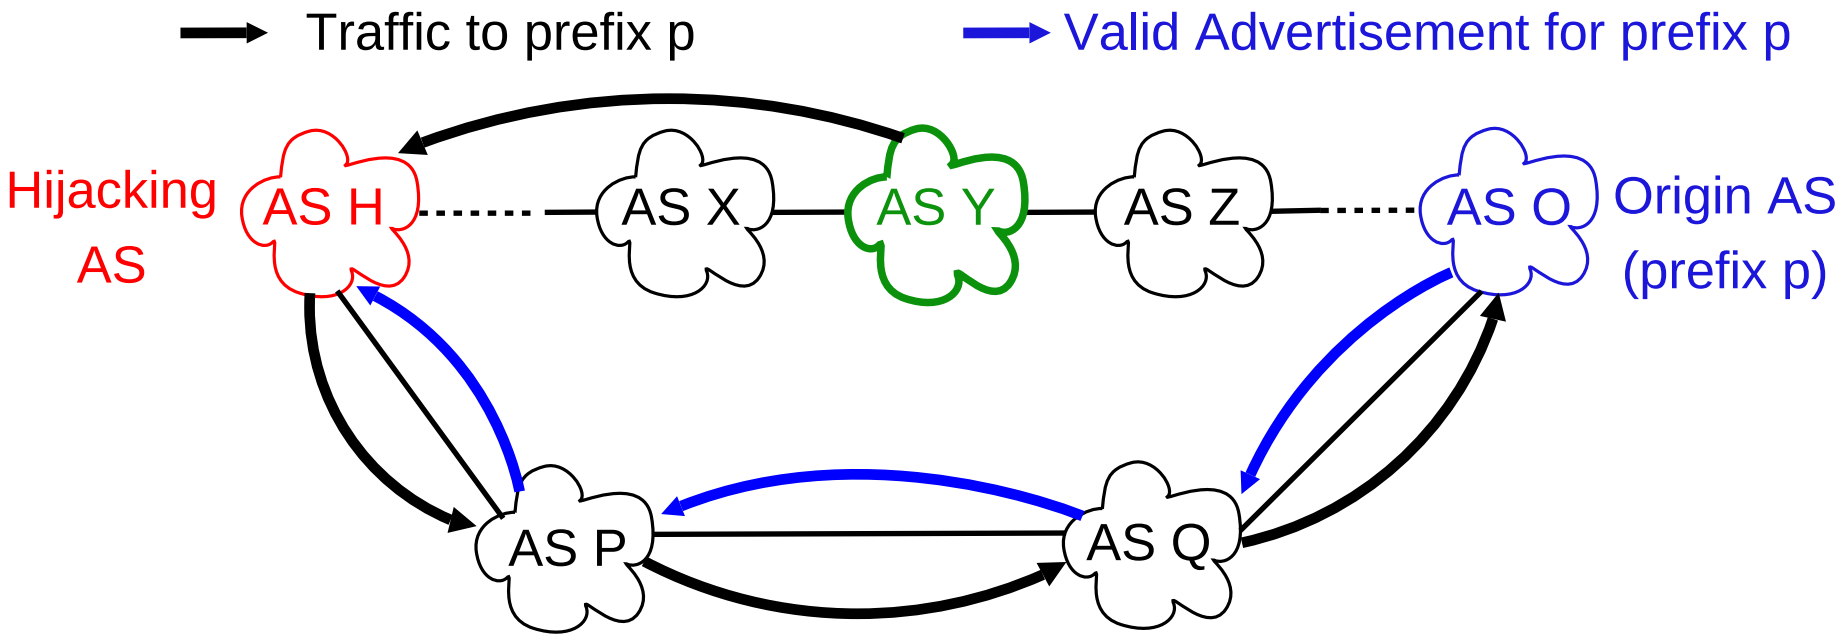

Can AS H intercept prefix p's traffic from AS Y?

1. Can AS H hijack prefix p's traffic from AS Y?
2. Can AS H route the hijacked traffic to back AS O?

Safety Condition: AS H should have a valid route for prefix p during Interception

Interception Analysis

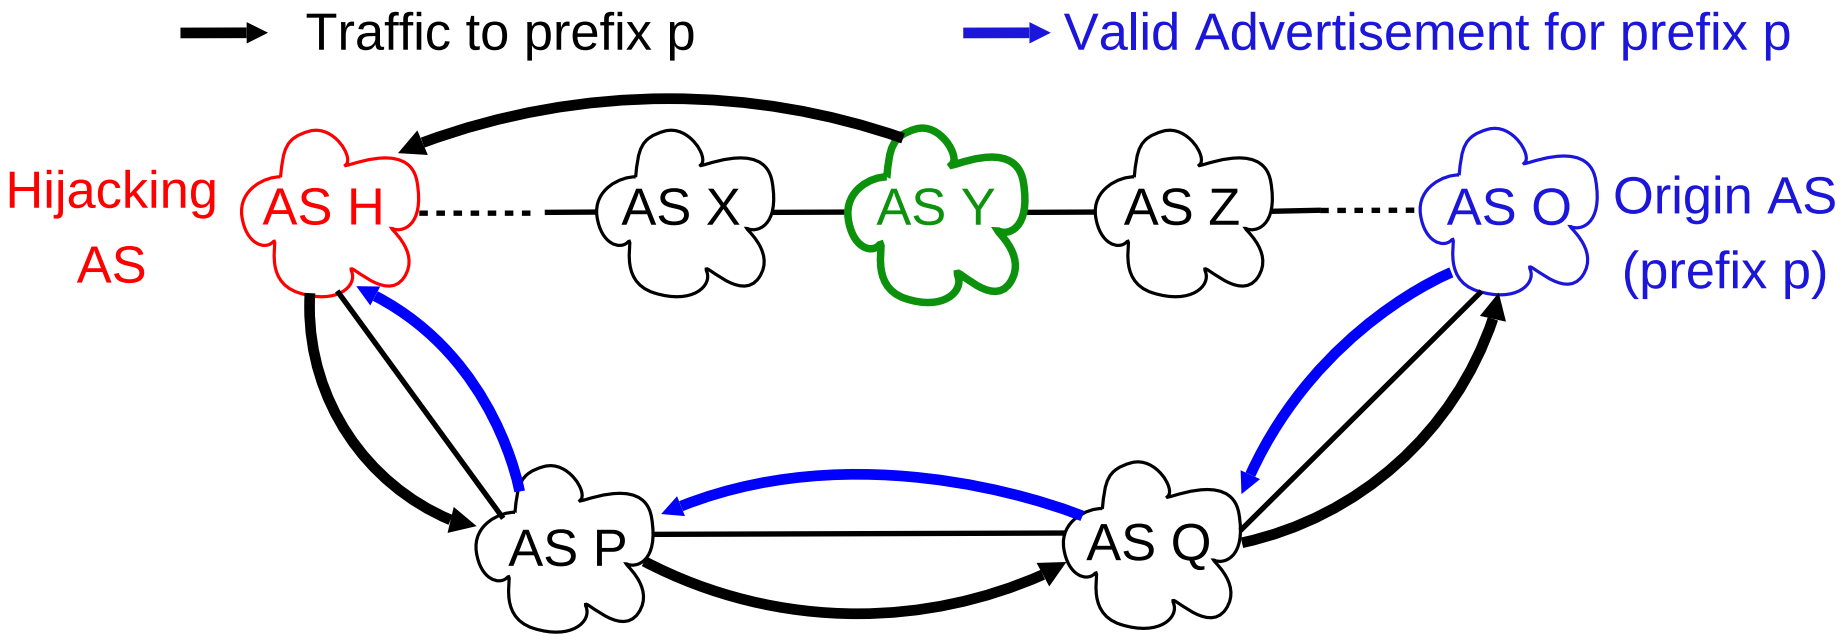

Can AS H intercept prefix p's traffic from AS Y?

1. Can AS H hijack prefix p's traffic from AS Y?
2. Can AS H route the hijacked traffic to back AS O?

Can AS H advertize the invalid route to a neighbor without impacting its valid route?

Interception Analysis

Invalid advertisement to a provider can violate the safety condition if AS H's valid route is through a provider

Talk Outline

- Introduction
- Hijacking Analysis
- Interception Analysis
- Hijacking and Interception estimates
- Hijacking and Intercepting real traffic
- Detecting Internet Interception
- Conclusions

Hijacking and Interception Estimates

Analysis results applied to Route-Views ASes

- Route-Views repository comprises of 34 ASes (RV-Set)
- 7 tier-1 ASes, 19 tier-2 ASes and 8 others
- CAIDA AS-relationship database

Hijacking and Interception Estimates

Analysis results applied to Route-Views ASes

- Route-Views repository comprises of 34 ASes (RV-Set)
- 7 tier-1 ASes, 19 tier-2 ASes and 8 others
- CAIDA AS-relationship database

Parameters of interest

1. **Probability of Hijacking**: Fraction of ASes whose traffic is hijacked by the hijacking AS, averaged across all ASes and all prefixes
   Analysis yields upper-bound (**UB**) and lower-bound (**LB**).
2. **Probability of Interception**: Defined analogously

Hijacking and Interception Estimates

Analysis results applied to Route-Views ASes

- Route-Views repository comprises of 34 ASes (RV-Set)
- 7 tier-1 ASes, 19 tier-2 ASes and 8 others
- CAIDA AS-relationship database

Parameters of interest

1. **Probability of Hijacking**: Fraction of ASes whose traffic is hijacked by the hijacking AS, averaged across all ASes and all prefixes
   Analysis yields upper-bound (**UB**) and lower-bound (**LB**).
2. **Probability of Interception**: Defined analogously

Hijacking and Interception Estimates

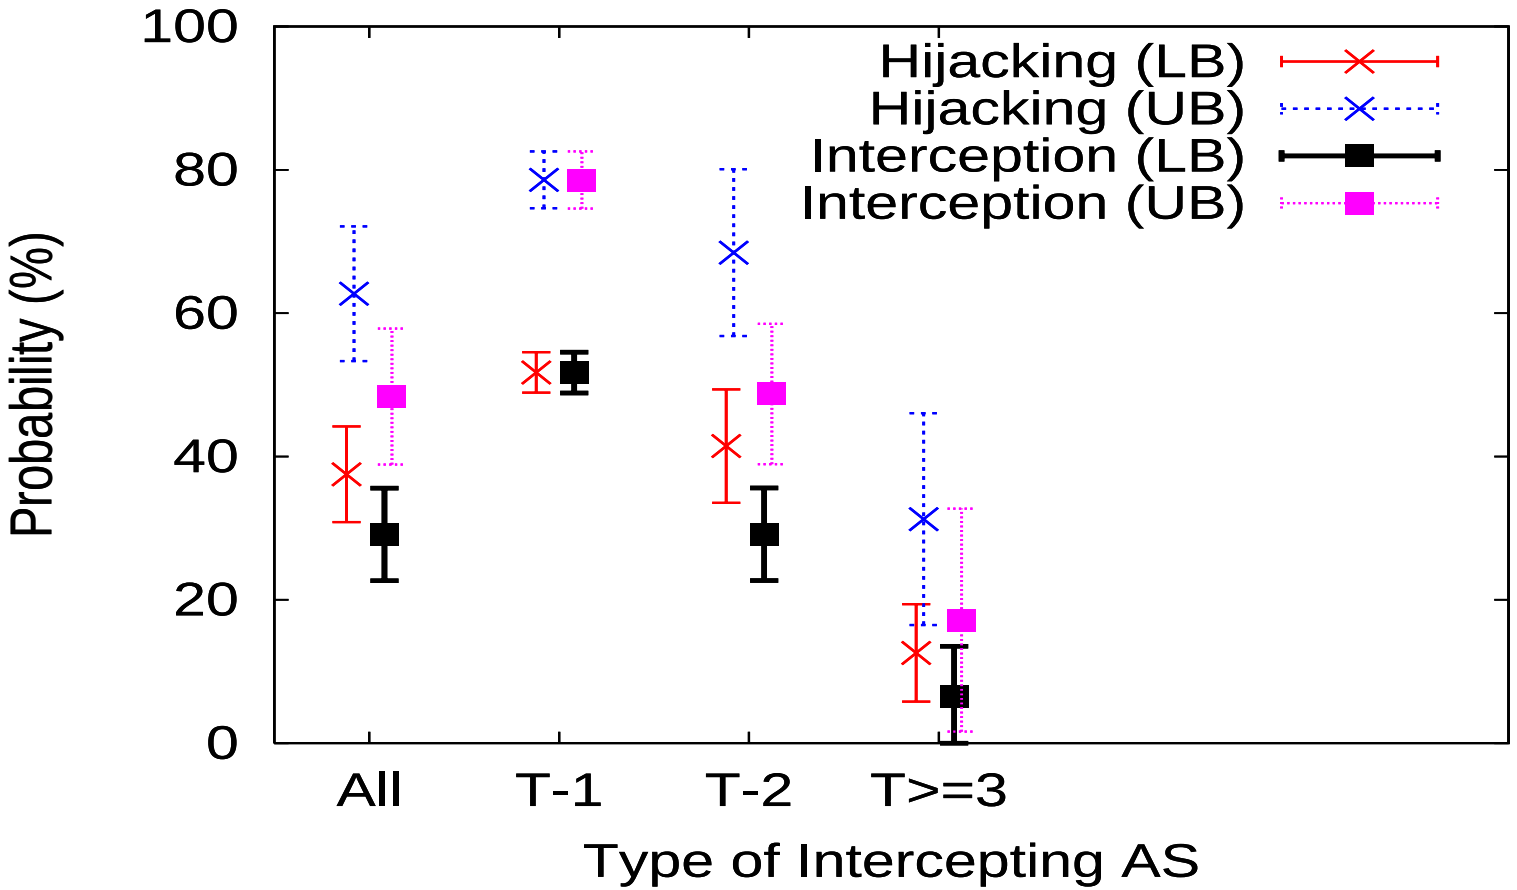

Hijacking and Interception Estimates

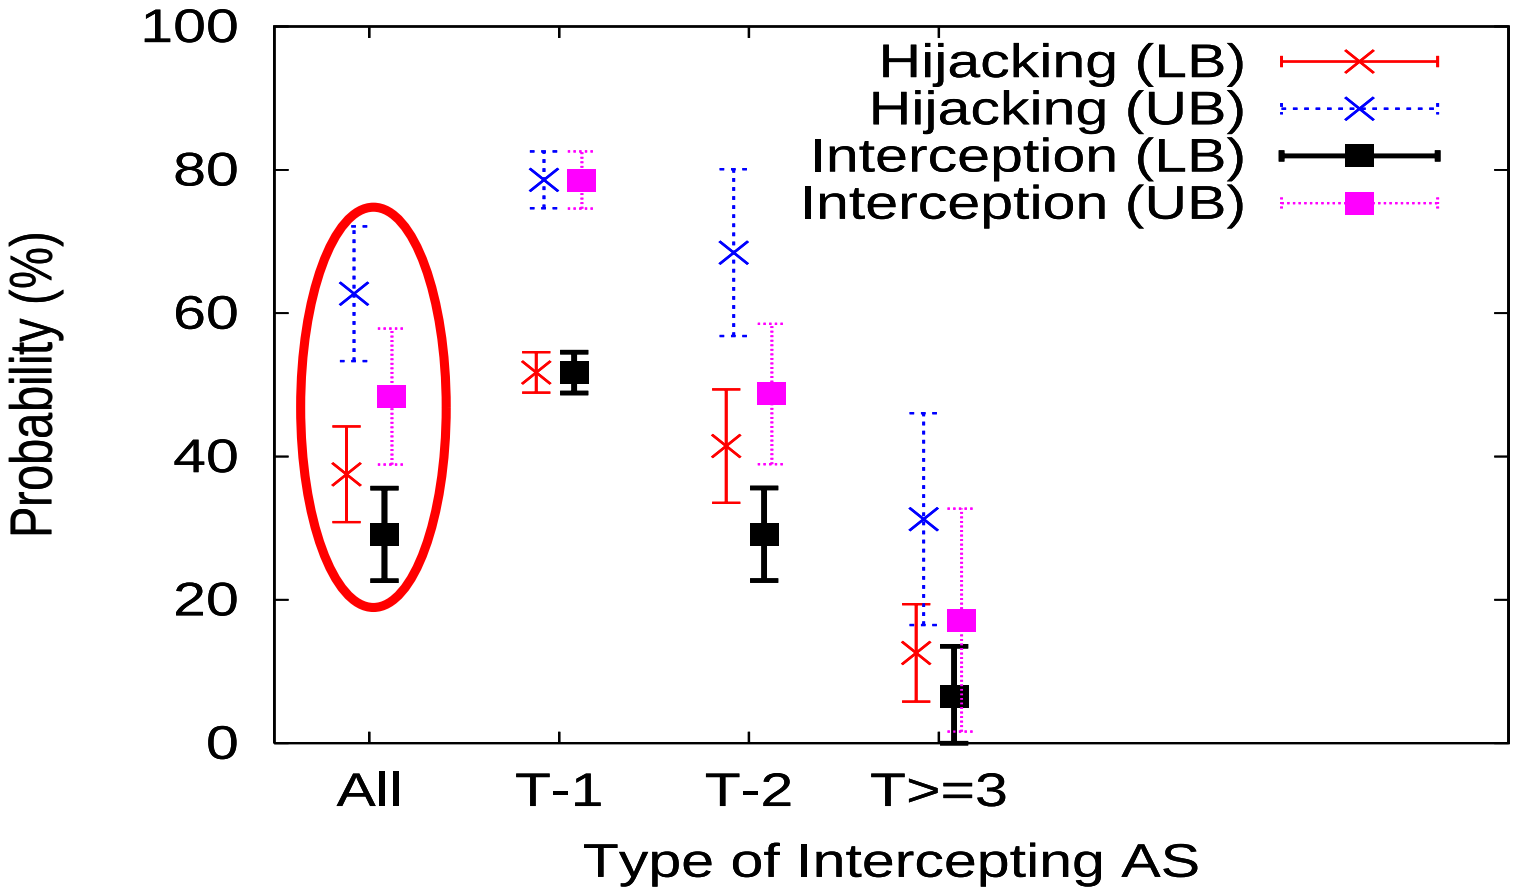

Overall probability of hijacking $\sim$ 40-60%

Overall probability of interception $\sim$ 30-50%

Hijacking and Interception Estimates

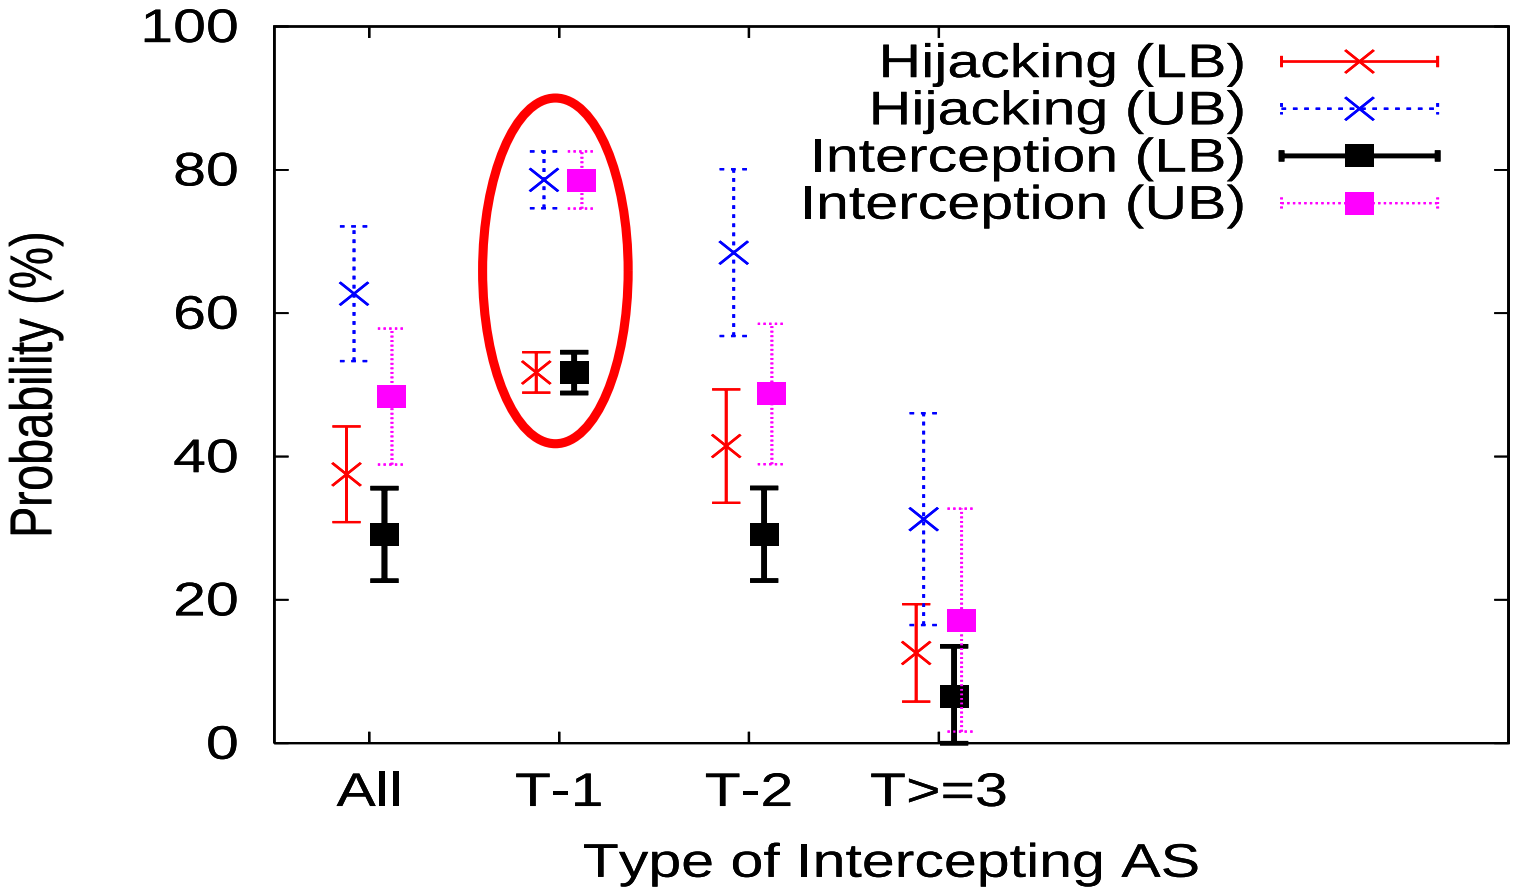

Probability of hijacking for tier-1 ISPs $\sim$ 50-80%

Probability of interception for tier-1 ISPs $\sim$ 50-80%

Hijacking and Interception Estimates

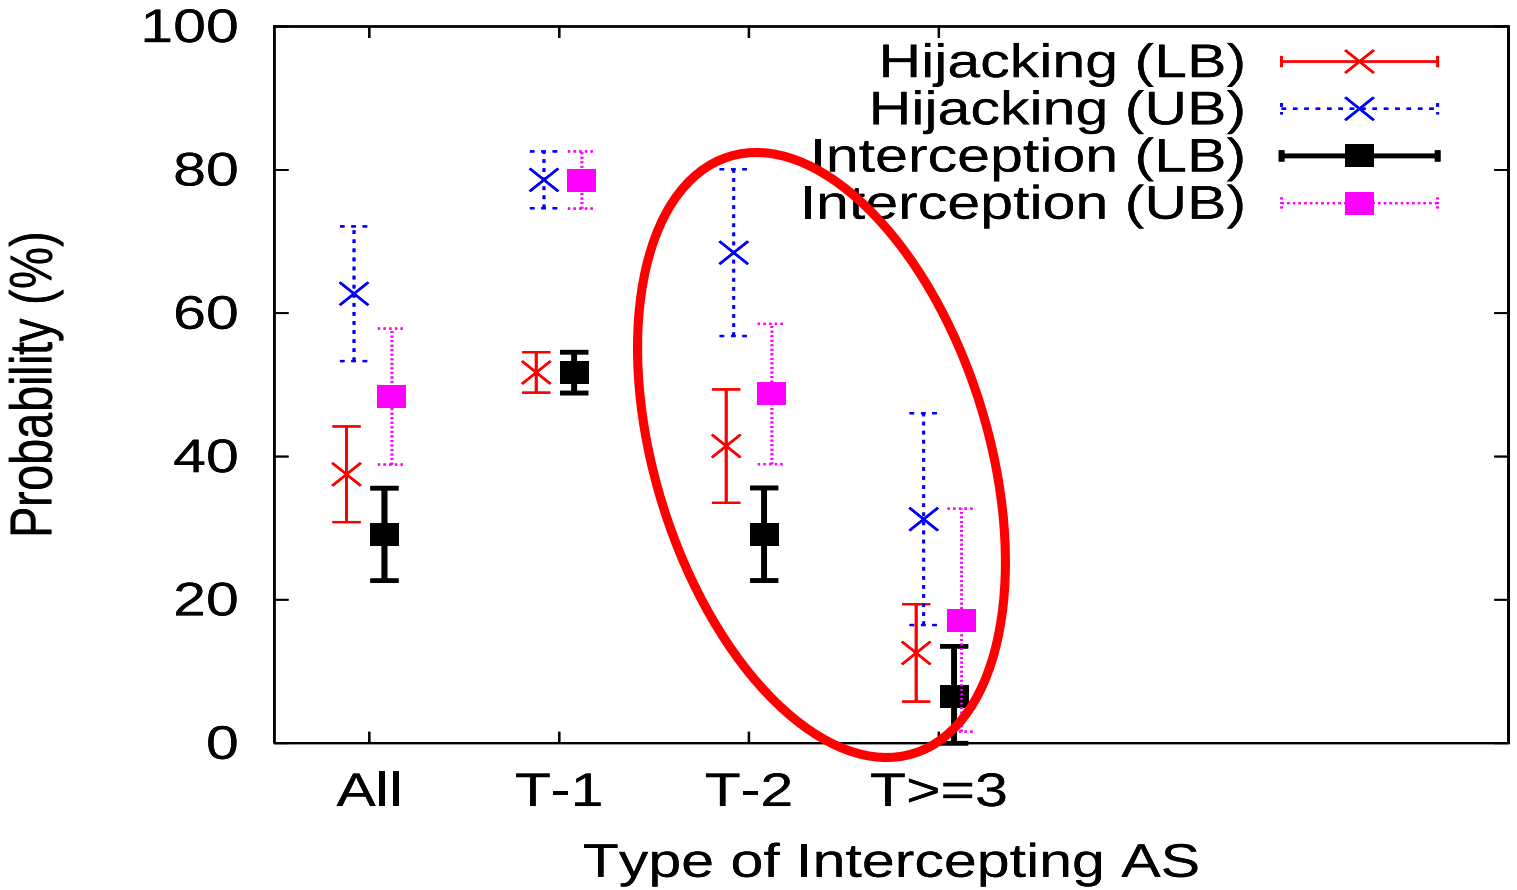

Verifying against known events

Apply analysis to known prefix hijacks

- Calculate *Actual Hijacking Percentage*
- Calculate *Estimated Hijacking Percentage (LB-UB)*

Hijack of 64.233.161.0/24 [Wan et. al., SSN'06]

- Owner AS: Google (AS 15169)
- Hijacking AS: Cogent (AS 174)
- Actual Hijacking Percentage = 45.2% (14 of 31 Route-Views ASes hijacked)
- Estimated Hijacking Percentage = 35.5-65.5%

Verifying against known events

Prefix	Owner (AS name)	Hijacker	Estimated Hijacking LB-UB %	Actual Hijack--ing (%)
64.233.161.0/24	Google	Cogent	35.5-64.5	45.2
12.173.227.0/24	MarthaStewart Living	ConEd.	36.4-84.9	42.4
63.165.71.0/24	Folksamerica	"	39.4-72.7	39.4
64.132.55.0/24	OverseasMedia	"	18.2-51.5	18.2
65.115.240.0/24	ViewTrade	"	27.2-54.5	21.2
65.209.93.0/24	LavaTrading	"	39.4-72.7	45.5
66.77.142.0/24	Folksamerica	"	90.9-90.9	90.9
66.194.137.0/24	MacKayShields	"	18.2-57.5	27.3
66.207.32.0/20	ADI	"	45.5-66.7	63.6
69.64.209.0/24	TheStreet.Com	"	72.7-81.8	84.8
160.79.45.0/24	RhodesASN	"	27.3-75.8	51.5
160.79.67.0/24	TheStreet.Com	"	60.6-75.8	69.7
192.251.16.0/24	T&TForex	"	27.3-57.6	27.3
198.15.10.0/24	TigerFund	"	0-1	60.6
204.13.72.0/24	FTENNY	"	93.9-93.9	75.8
216.223.46.0/24	SDSNY	"	51.5-78.8	18.2

Accurate prediction in 11 of the 16 cases

Talk Outline

- Introduction
- Hijacking Analysis
- Interception Analysis
- Hijacking and Interception estimates
- Hijacking and Intercepting real traffic
- Detecting Internet Interception
- Conclusions

Hijacking and Intercepting real traffic

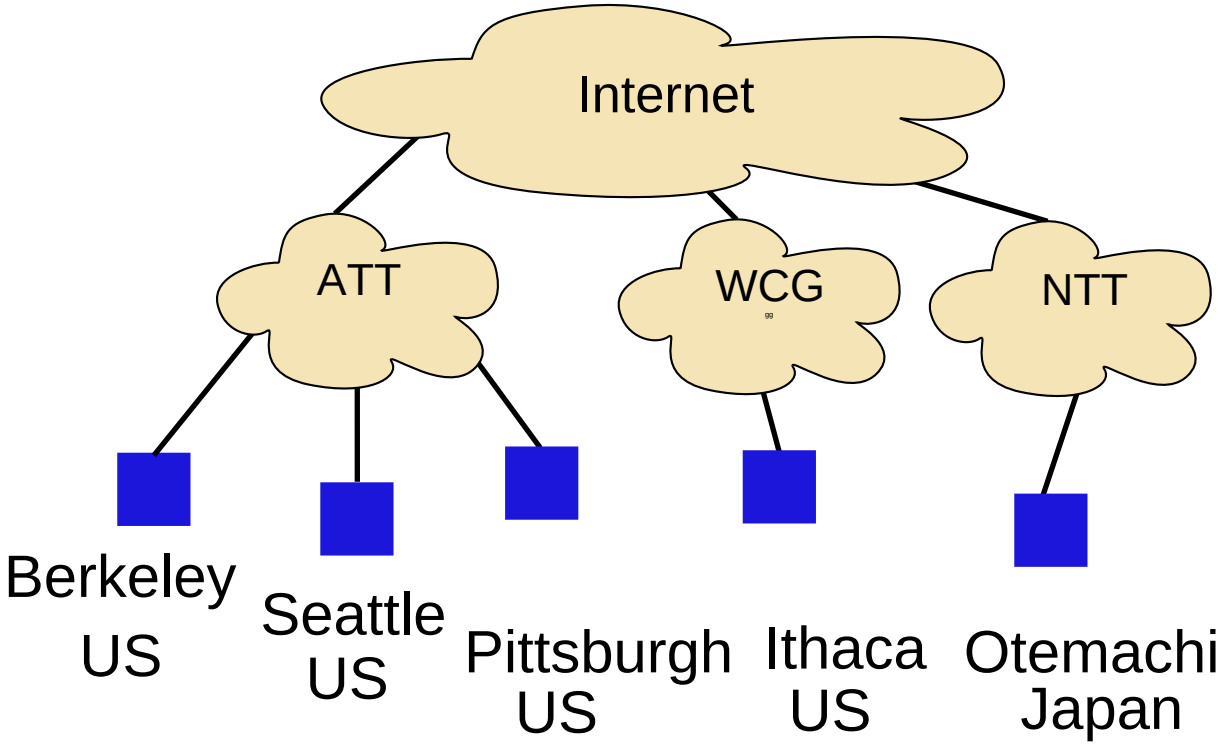

Hijacking and Intercepting real traffic

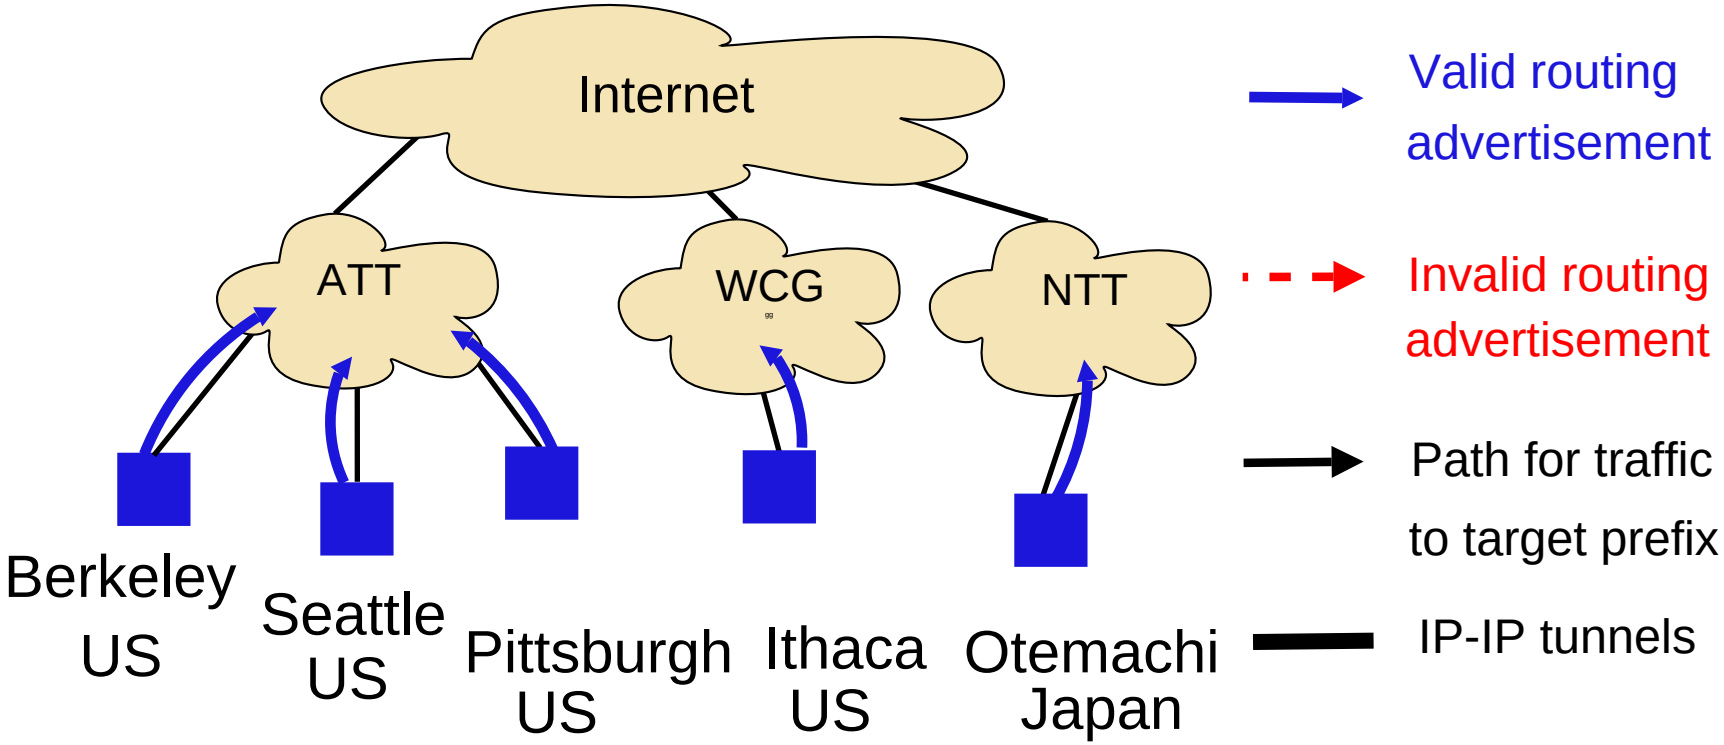

Our prefix (204.9.168.0/22) can be advertised by each of the five sites

Hijacking and Intercepting real traffic

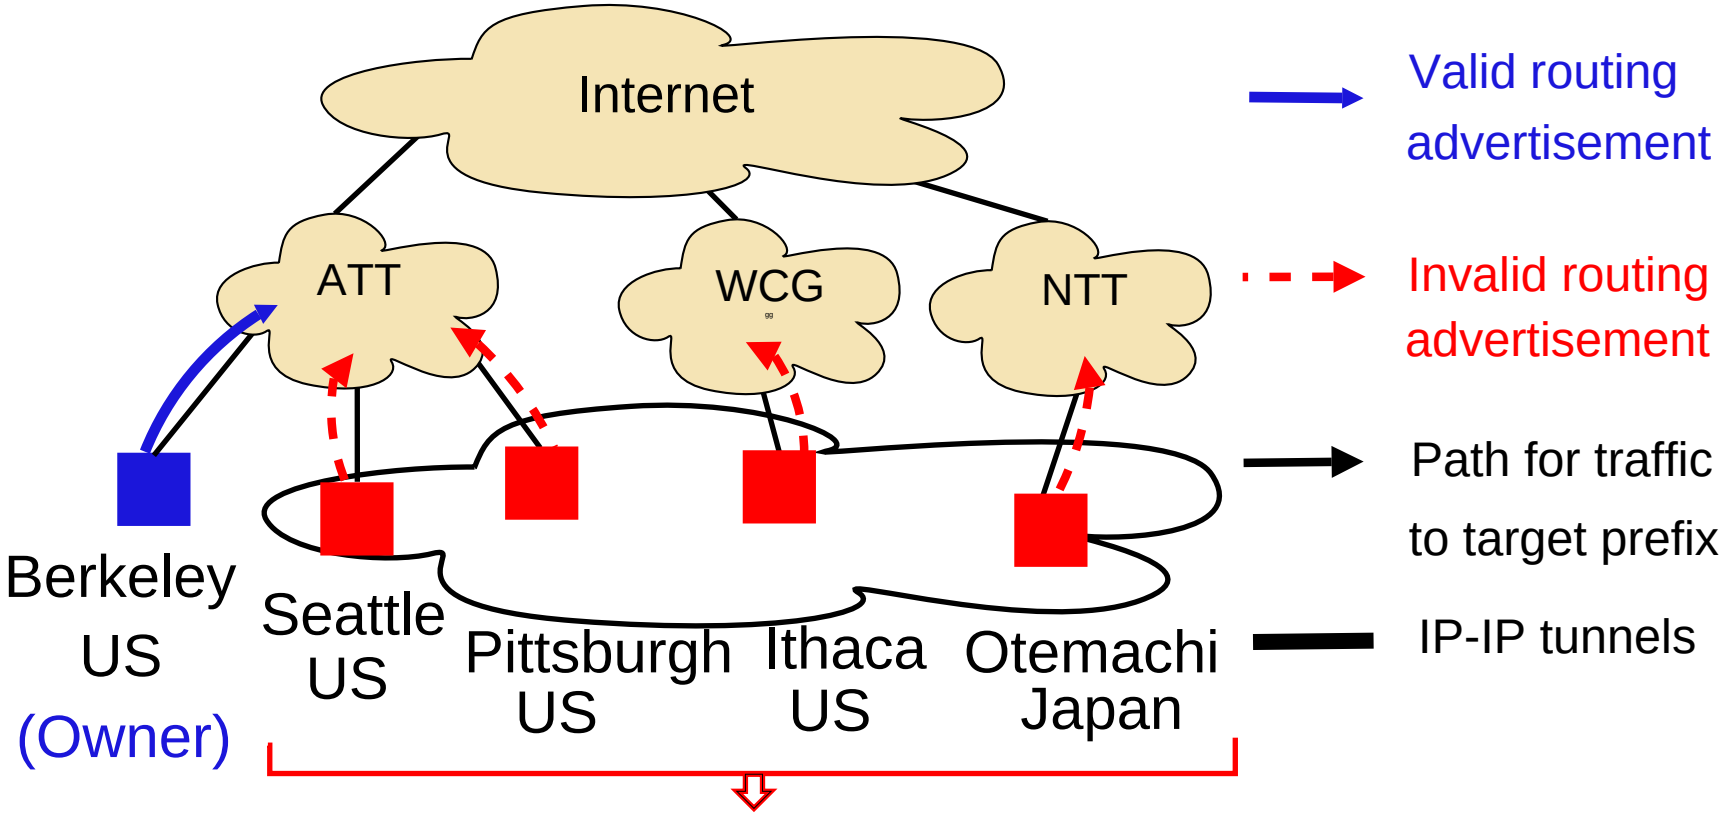

Owner AS: Berkeley site

Rest of the sites advertize prefix to hijack traffic

Hijacking and Intercepting real traffic

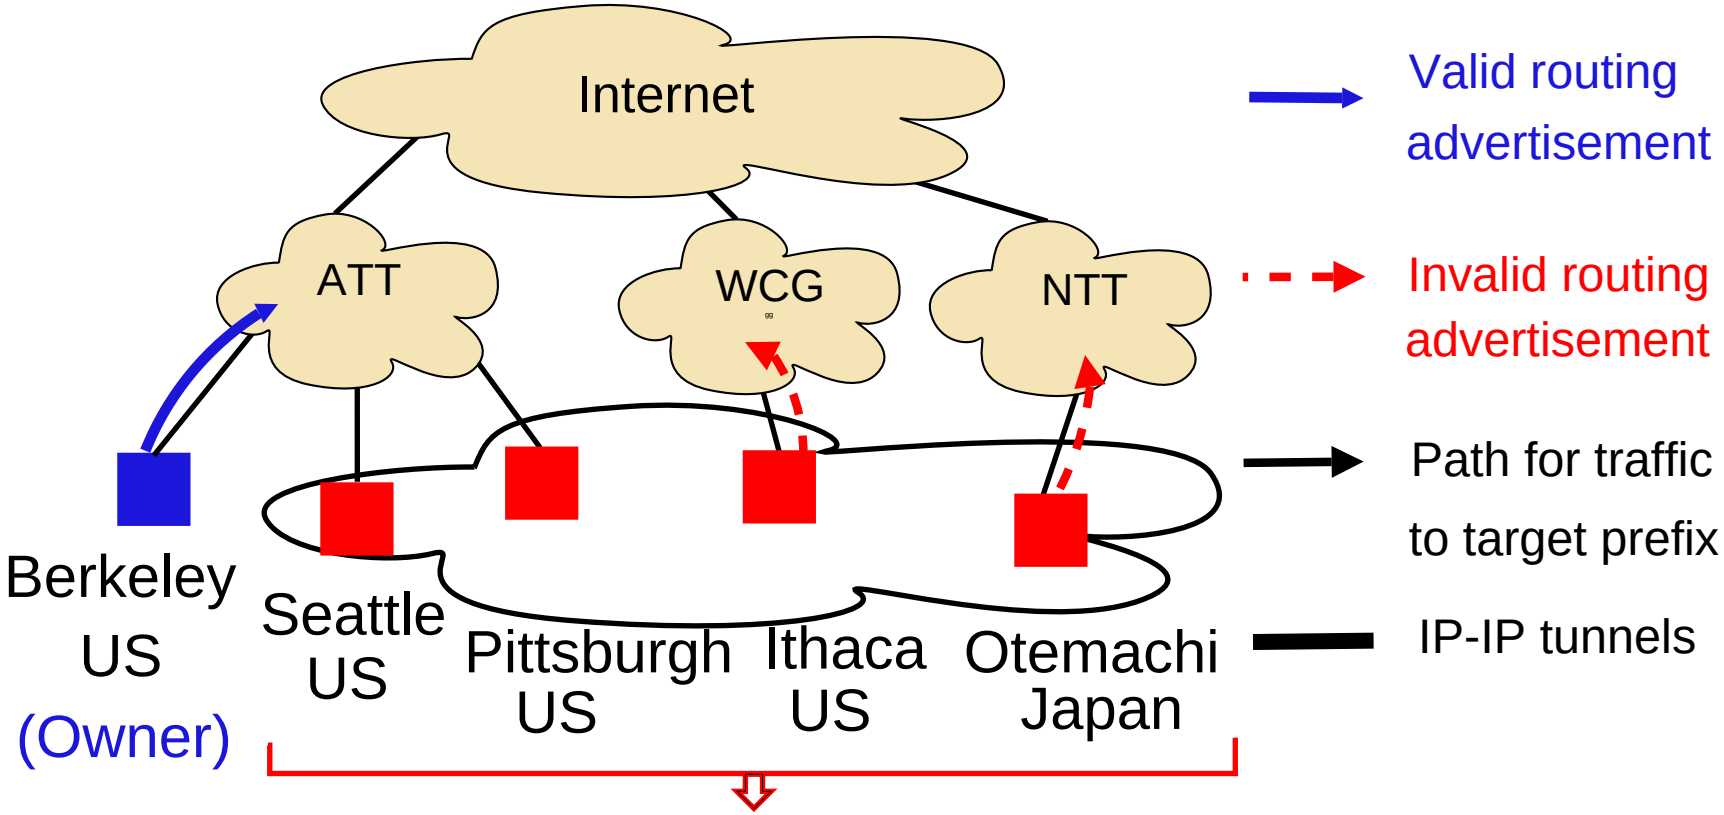

Interception of Traffic

Hijacking and Intercepting real traffic

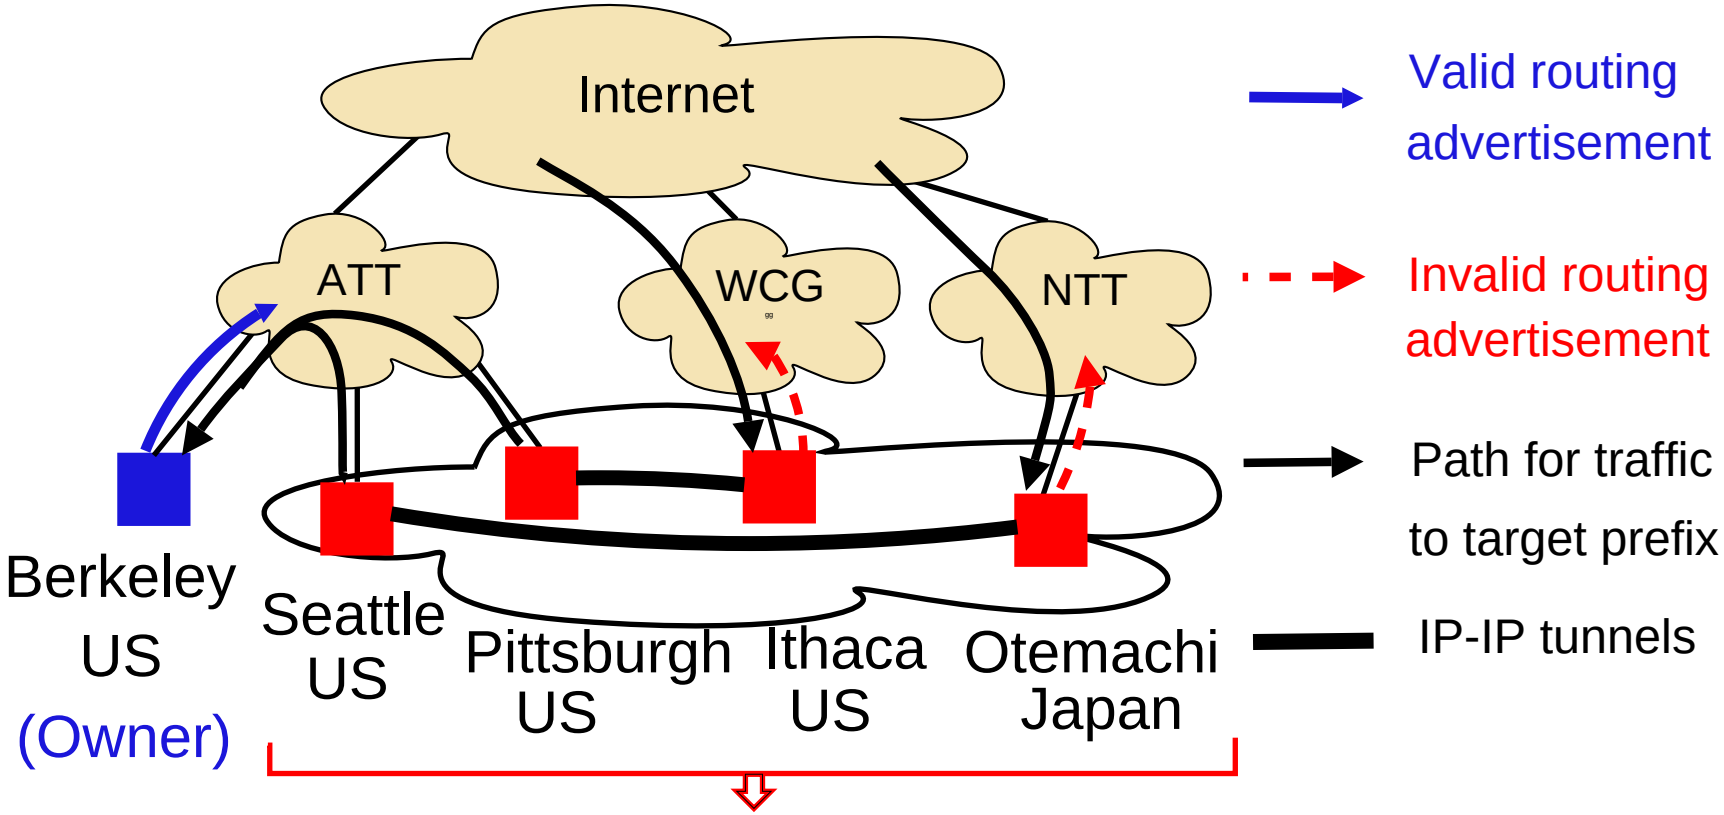

Interception of Traffic

Traffic is hijacked at Ithaca and Otemachi and routed back through Seattle and Pittsburgh

Hijacking and Intercepting real traffic

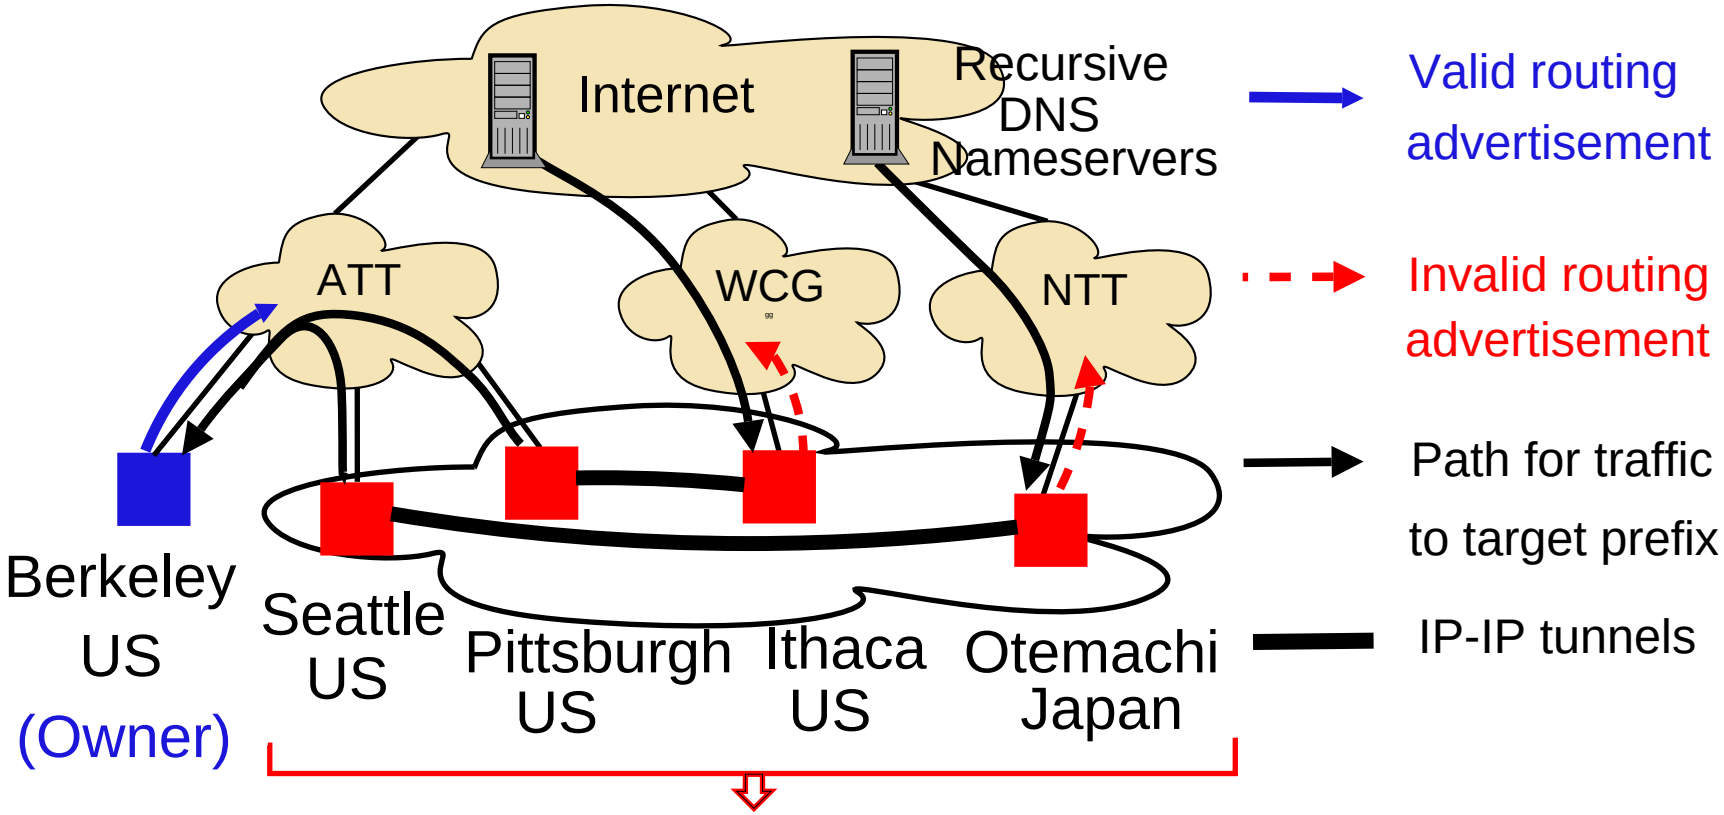

Use Recursive DNS Nameservers to generate traffic to our prefix [King, IMW'02]

Hijacking and Intercepting real traffic

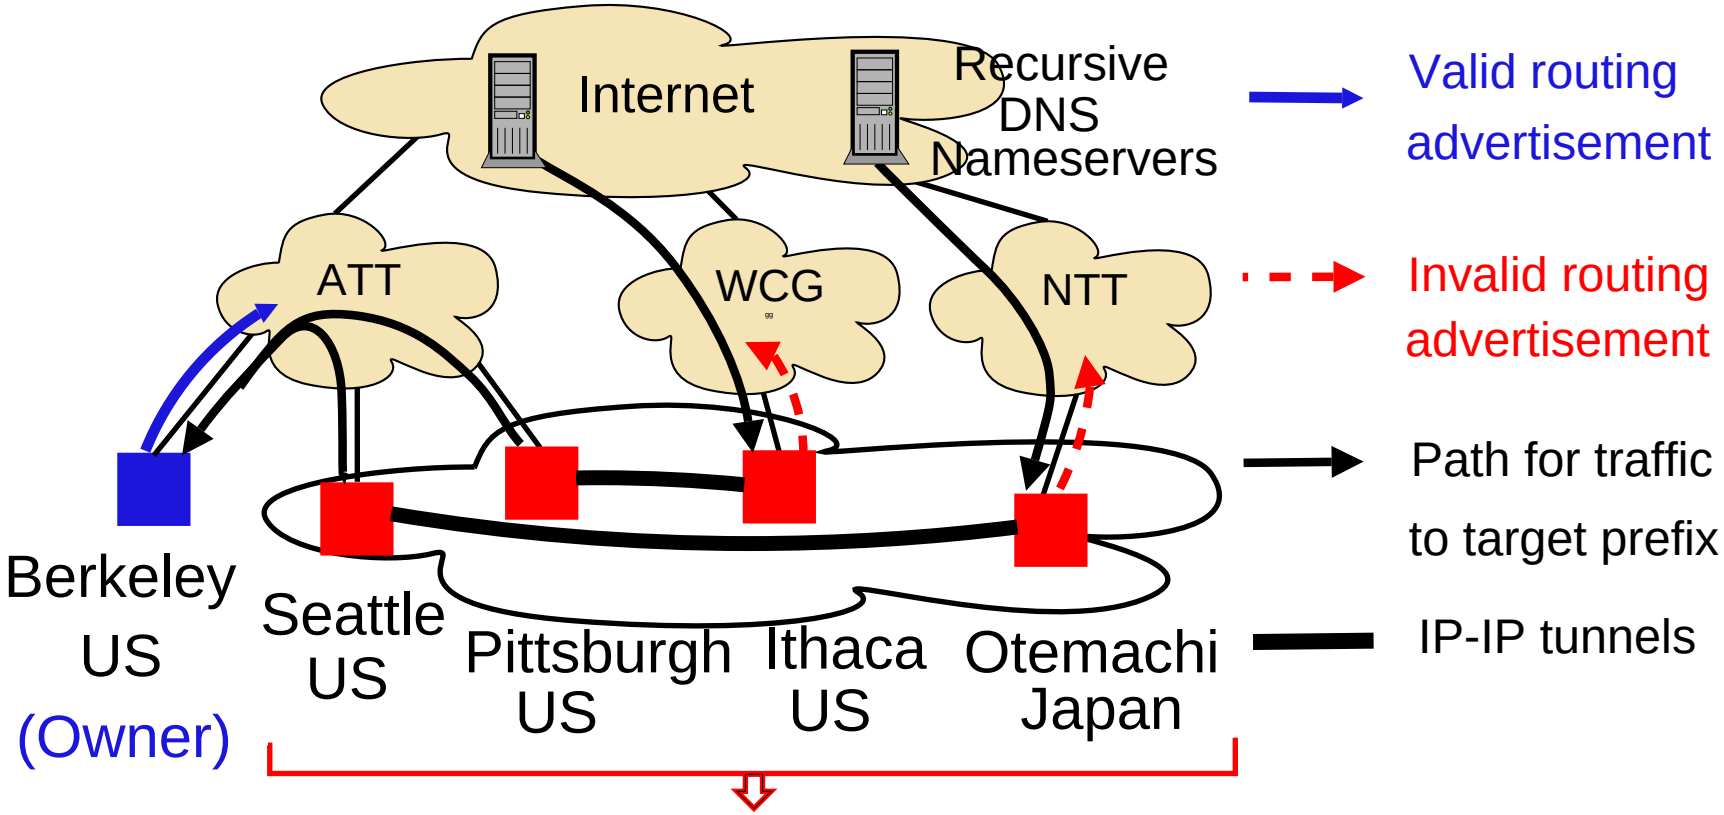

Generated traffic from 23,588 recursive nameservers

For each site as owner, hijacked and intercepted traffic using other sites

Hijacking and Intercepting real traffic

Ber	Pit	Sea	Ith	Ote	% of traffic Hijacked	% of traffic Intercepted
O	✗	✗	✓	✓	91.7	78.8
✗	O	✗	✓	✓	68.8	67.5
✗	✗	O	✓	✓	97.4	66.2
✗	✗	✗	O	✓	66.0	47.3
✓	✓	✓	✗	O	76.1	23.4

O : Site owning the prefix
✗ : Site not advertising an invalid route during interception
✓ : Site advertising an invalid route during interception

Hijacking and Intercepting real traffic

Ber	Pit	Sea	Ith	Ote	% of traffic Hijacked	% of traffic Intercepted
O	✗	✗	✓	✓	91.7	78.8
✗	O	✗	✓	✓	68.8	67.5
✗	✗	O	✓	✓	97.4	66.2
✗	✗	✗	O	✓	66.0	47.3
✓	✓	✓	✗	O	76.1	23.4

O : Site owning the prefix
✗ : Site not advertising an invalid route during interception
✓ : Site advertising an invalid route during interception

Talk Outline

- Introduction
- Hijacking Analysis
- Interception Analysis
- Hijacking and Interception estimates
- Hijacking and Intercepting real traffic
- Detecting Internet Interception
- Conclusions

Is Internet traffic being intercepted?

Use data-plane and control-plane information

- Intercepting ISP needs to route traffic back to the owner
- Data-plane AS-level path should differ *significantly* from the control-plane AS-level path

A signature for Interception of prefix *p*

- Control-Plane: Origin AS O, Next-hop ASes N_1, .., N_n (Routes for the prefix: [. . ., N_1, O], . . ., [. . ., N_n, O])
- Data-plane trace wherein packets traverse AS N_i after traversing AS N_j ($j \neq i$) is a **next-hop anomaly**

Detecting Internet Next-hop Anomalies

Control-plane information

- Route-Views repository
- 43 BGP sources belonging to 34 distinct ASes
- Provides control-plane AS-level path to each prefix

Data-plane information

- IPlane project: daily traceroutes to ≈100,000 route prefixes from ≈200 Planetlab hosts
- Data-set for each day of analysis ≈ 20 million IP-level traceroutes
- Mapped IP-level traceroutes to AS-level traceroutes

Detecting Internet Next-hop Anomalies

Control-plane information

- Route-Views repository
- 43 BGP sources belonging to 34 distinct ASes
- Provides control-plane AS-level path to each prefix

Data-plane information

- IPlane project: daily traceroutes to ≈100,000 route prefixes from ≈200 Planetlab hosts
- Data-set for each day of analysis ≈ 20 million IP-level traceroutes
- Mapped IP-level traceroutes to AS-level traceroutes

Looked for next-hop anomalies in Oct-Dec, 2006

Observed Next-hop Anomalies

Errors in IP-to-AS mappings

- "Towards an Accurate AS-level traceroute"
  [Mao et. al., SIGCOMM'03]
- For example, IXPs, Sibling ASes, etc.

Traffic Engineering induced anomalies

- For example, multihomed origin AS using a next-hop AS as a backup by advertizing a longer route to it

Observed Next-hop Anomalies

Unexplained anomalies

- 16 unexplained next-hop anomalies
- E-mail survey: 3 responses indicating false-positives
- No conclusive evidence of Interception

Study does not rule out ongoing Interception

- Many assumptions about Intercepting AS's behavior

Observed Next-hop Anomalies

Unexplained anomalies

- 16 unexplained next-hop anomalies
- E-mail survey: 3 responses indicating false-positives
- No conclusive evidence of Interception

Study does not rule out ongoing Interception

- Many assumptions about Intercepting AS's behavior

Study highlights some of the challenges posed by the Interception Detection problem

Conclusions

Prefix Hijacking and Interception estimates

- Tier-1 ASes can hijack and intercept significant fraction of traffic to any prefix
- Small ASes can hijack and intercept a non-negligible amount of traffic
- Verified using known prefix hijacking events

Implemented Interception methodology

- Intercepted real traffic
- ASes can intercept traffic using the existing routing set-up

Study to detect ongoing interception

- Highlights challenges posed by Interception detection


ありがとう