Inferring Country-Level Transit Influence of Autonomous Systems

Alexander Gamero-Garrido

Esteban Carisimo, Shuai Hao, Bradley Huffaker,
kc claffy, Alex C. Snoeren, Alberto Dainotti, and Amogh Dhamdhere

Center for Applied Internet Data Analysis

The Register®
Biting the hand that feeds IT

NTRE SOFTWARE **SECURITY** DEVOPS BUSINESS PERSONAL TECH SCIENCE

Security

DDoS attack boots Kyrgyzstan from net

Russian bears blamed

By Dan Goodin 28 Jan 2009 at 19:57

Kyrgyzstan Under DDoS Attack From Russia

TUESDAY, JANUARY 27, 2009
BY: COUNTER THREAT UNIT RESEARCH TEAM

The two primary Kyrgyzstan ISPs (www.domain.kg, www.ns.kg) have been under a massive, sustained DDoS attack . . . Few alternatives for Internet access exist in Kyrgyzstan. . . . [the attacks] essentially knocked most of the small, Central Asian republic offline.

What are the technical mechanisms at play in these disconnections?

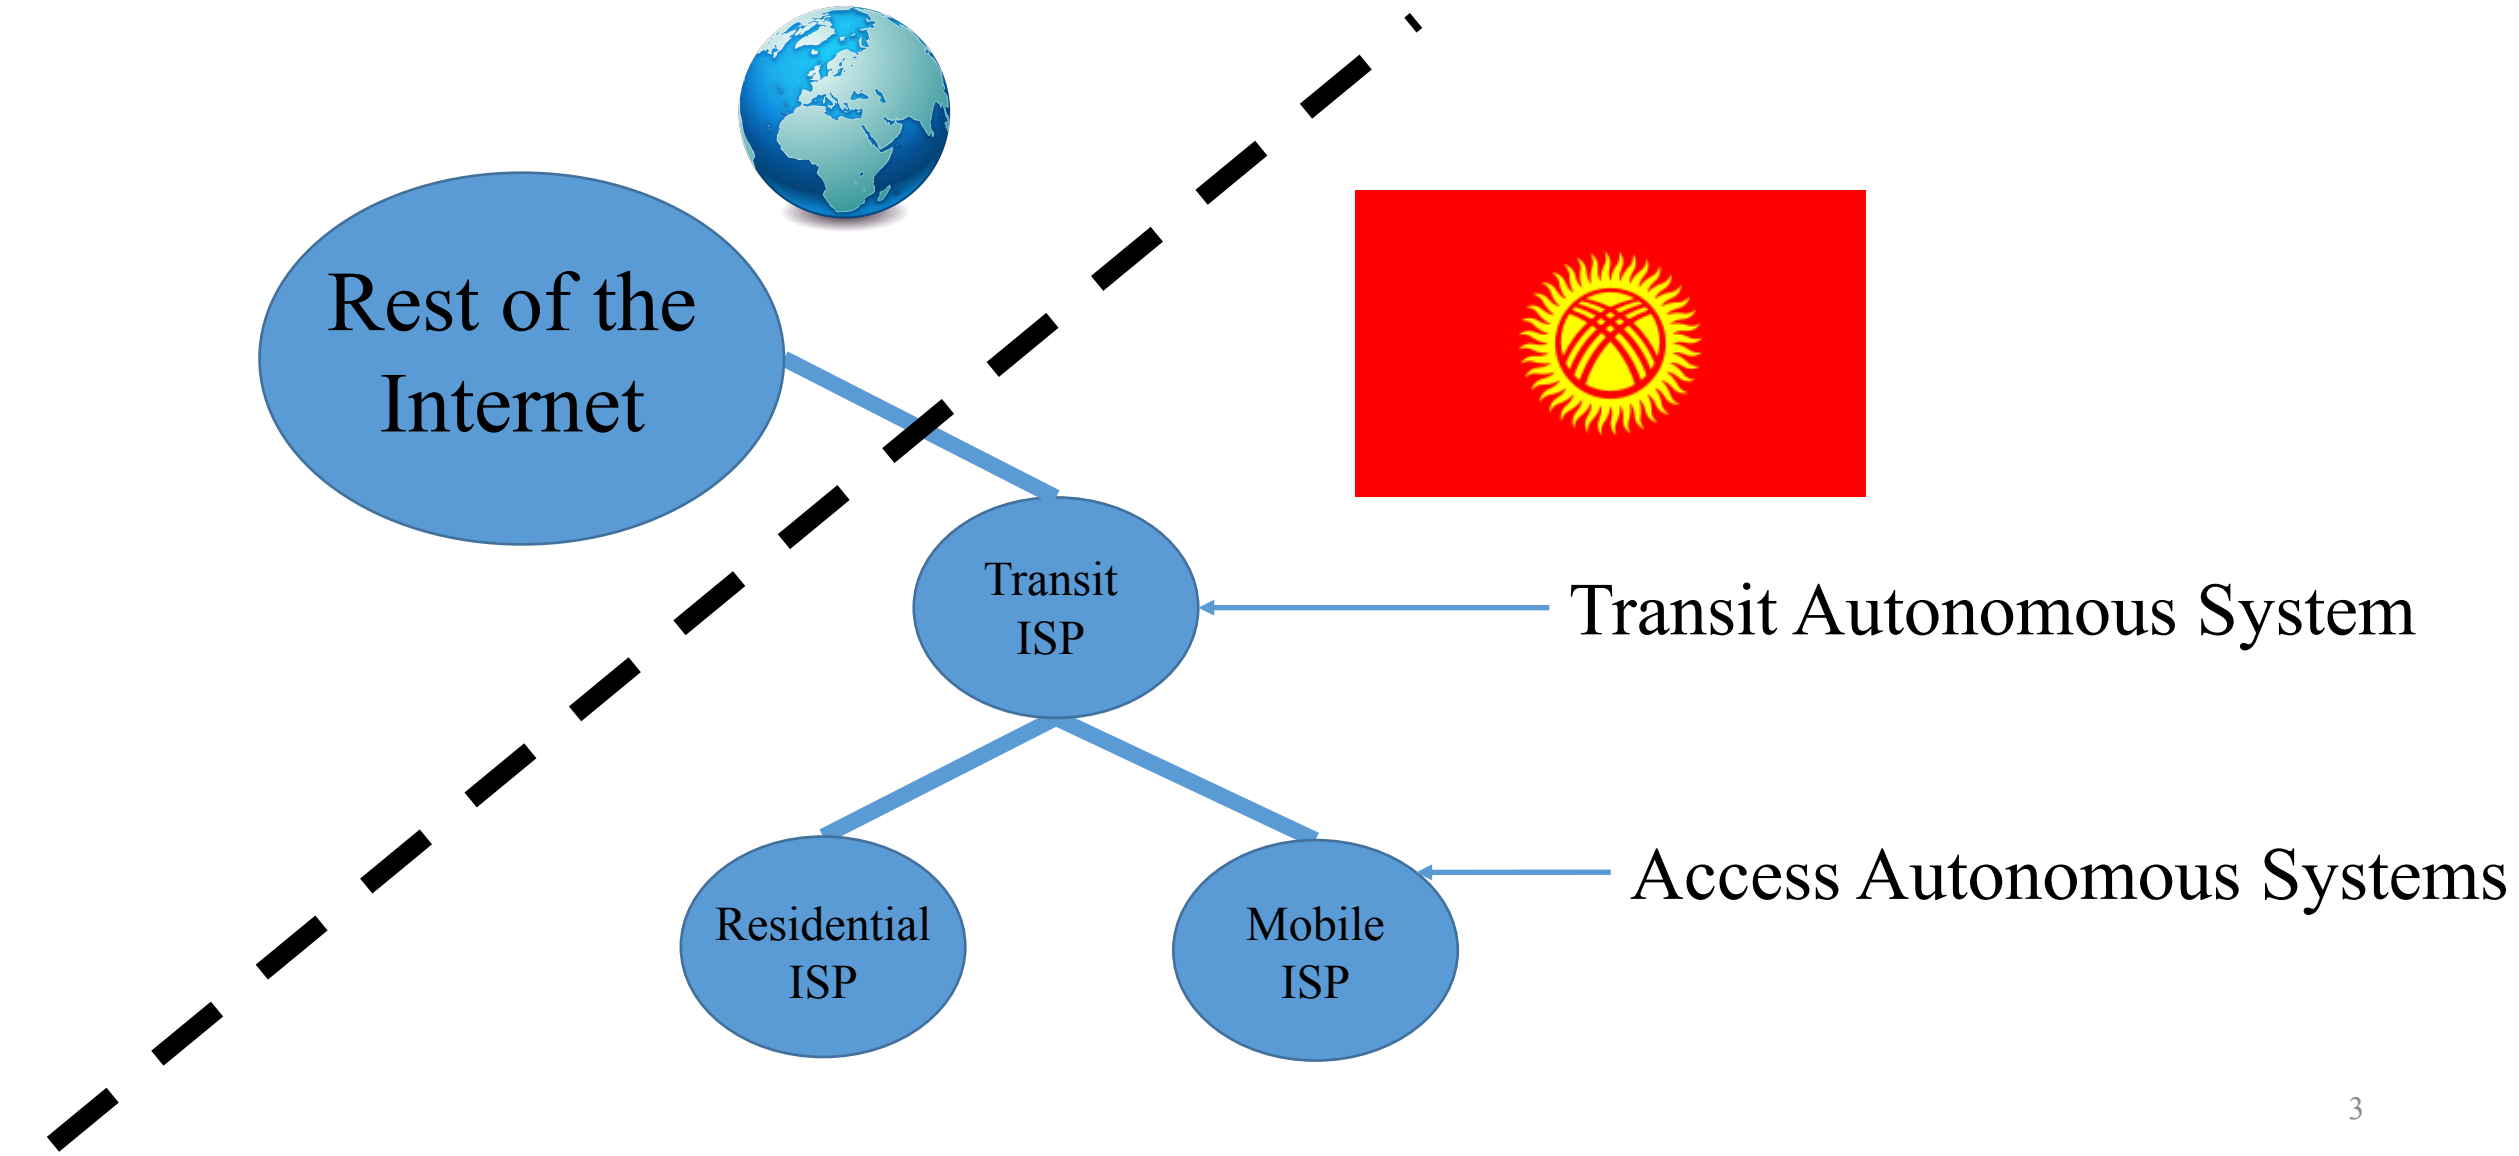

What are the technical mechanisms at play in these disconnections?

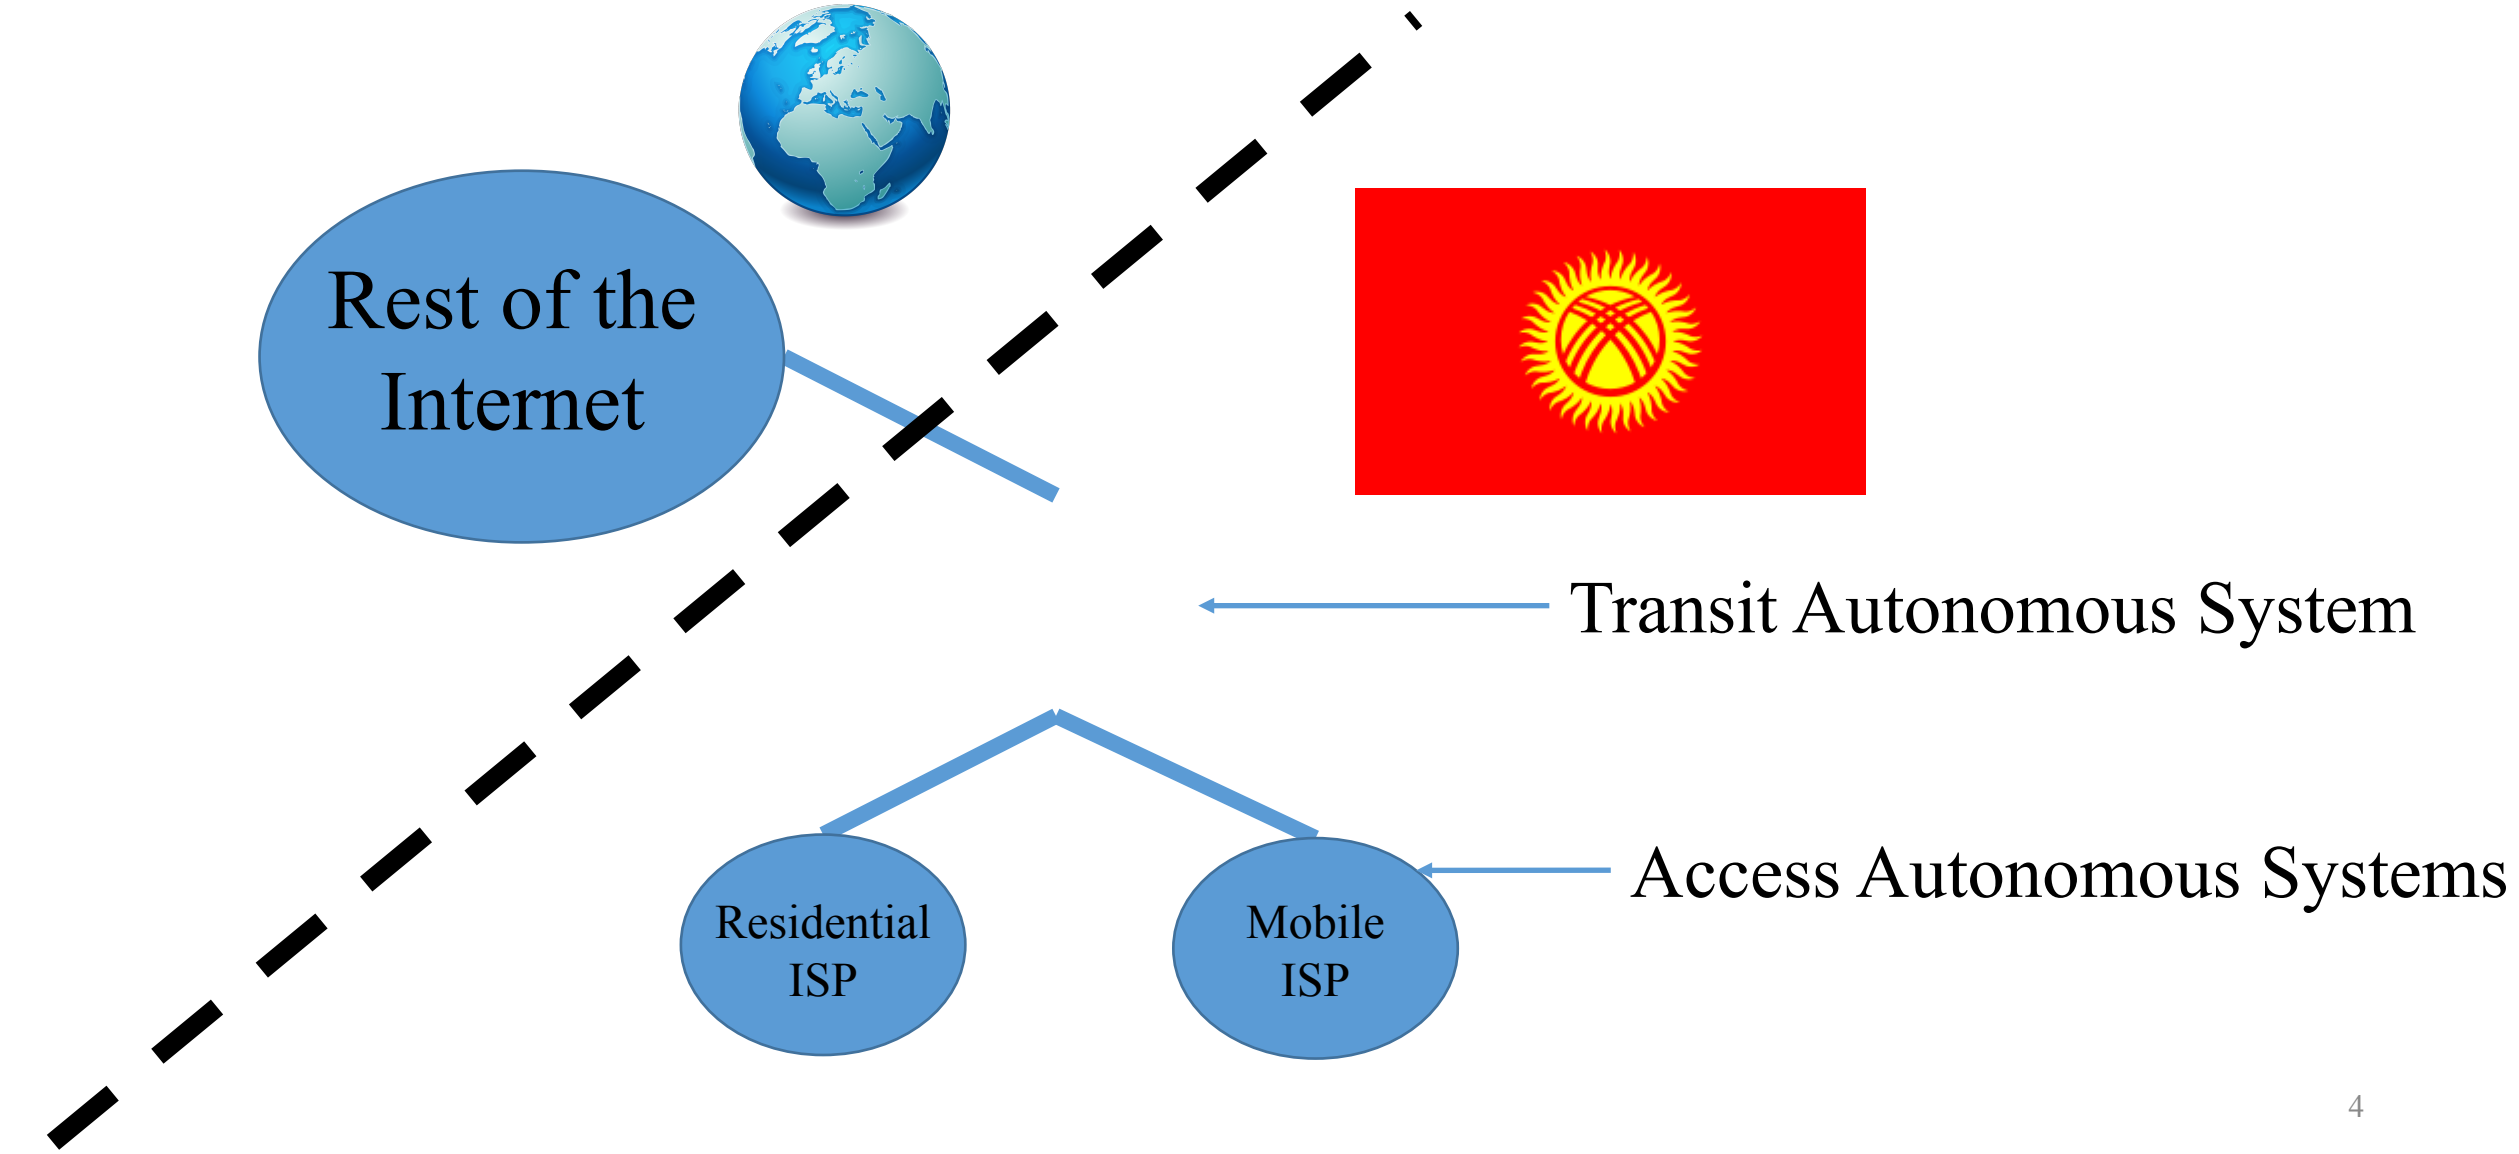

What are the technical mechanisms at play in these disconnections?

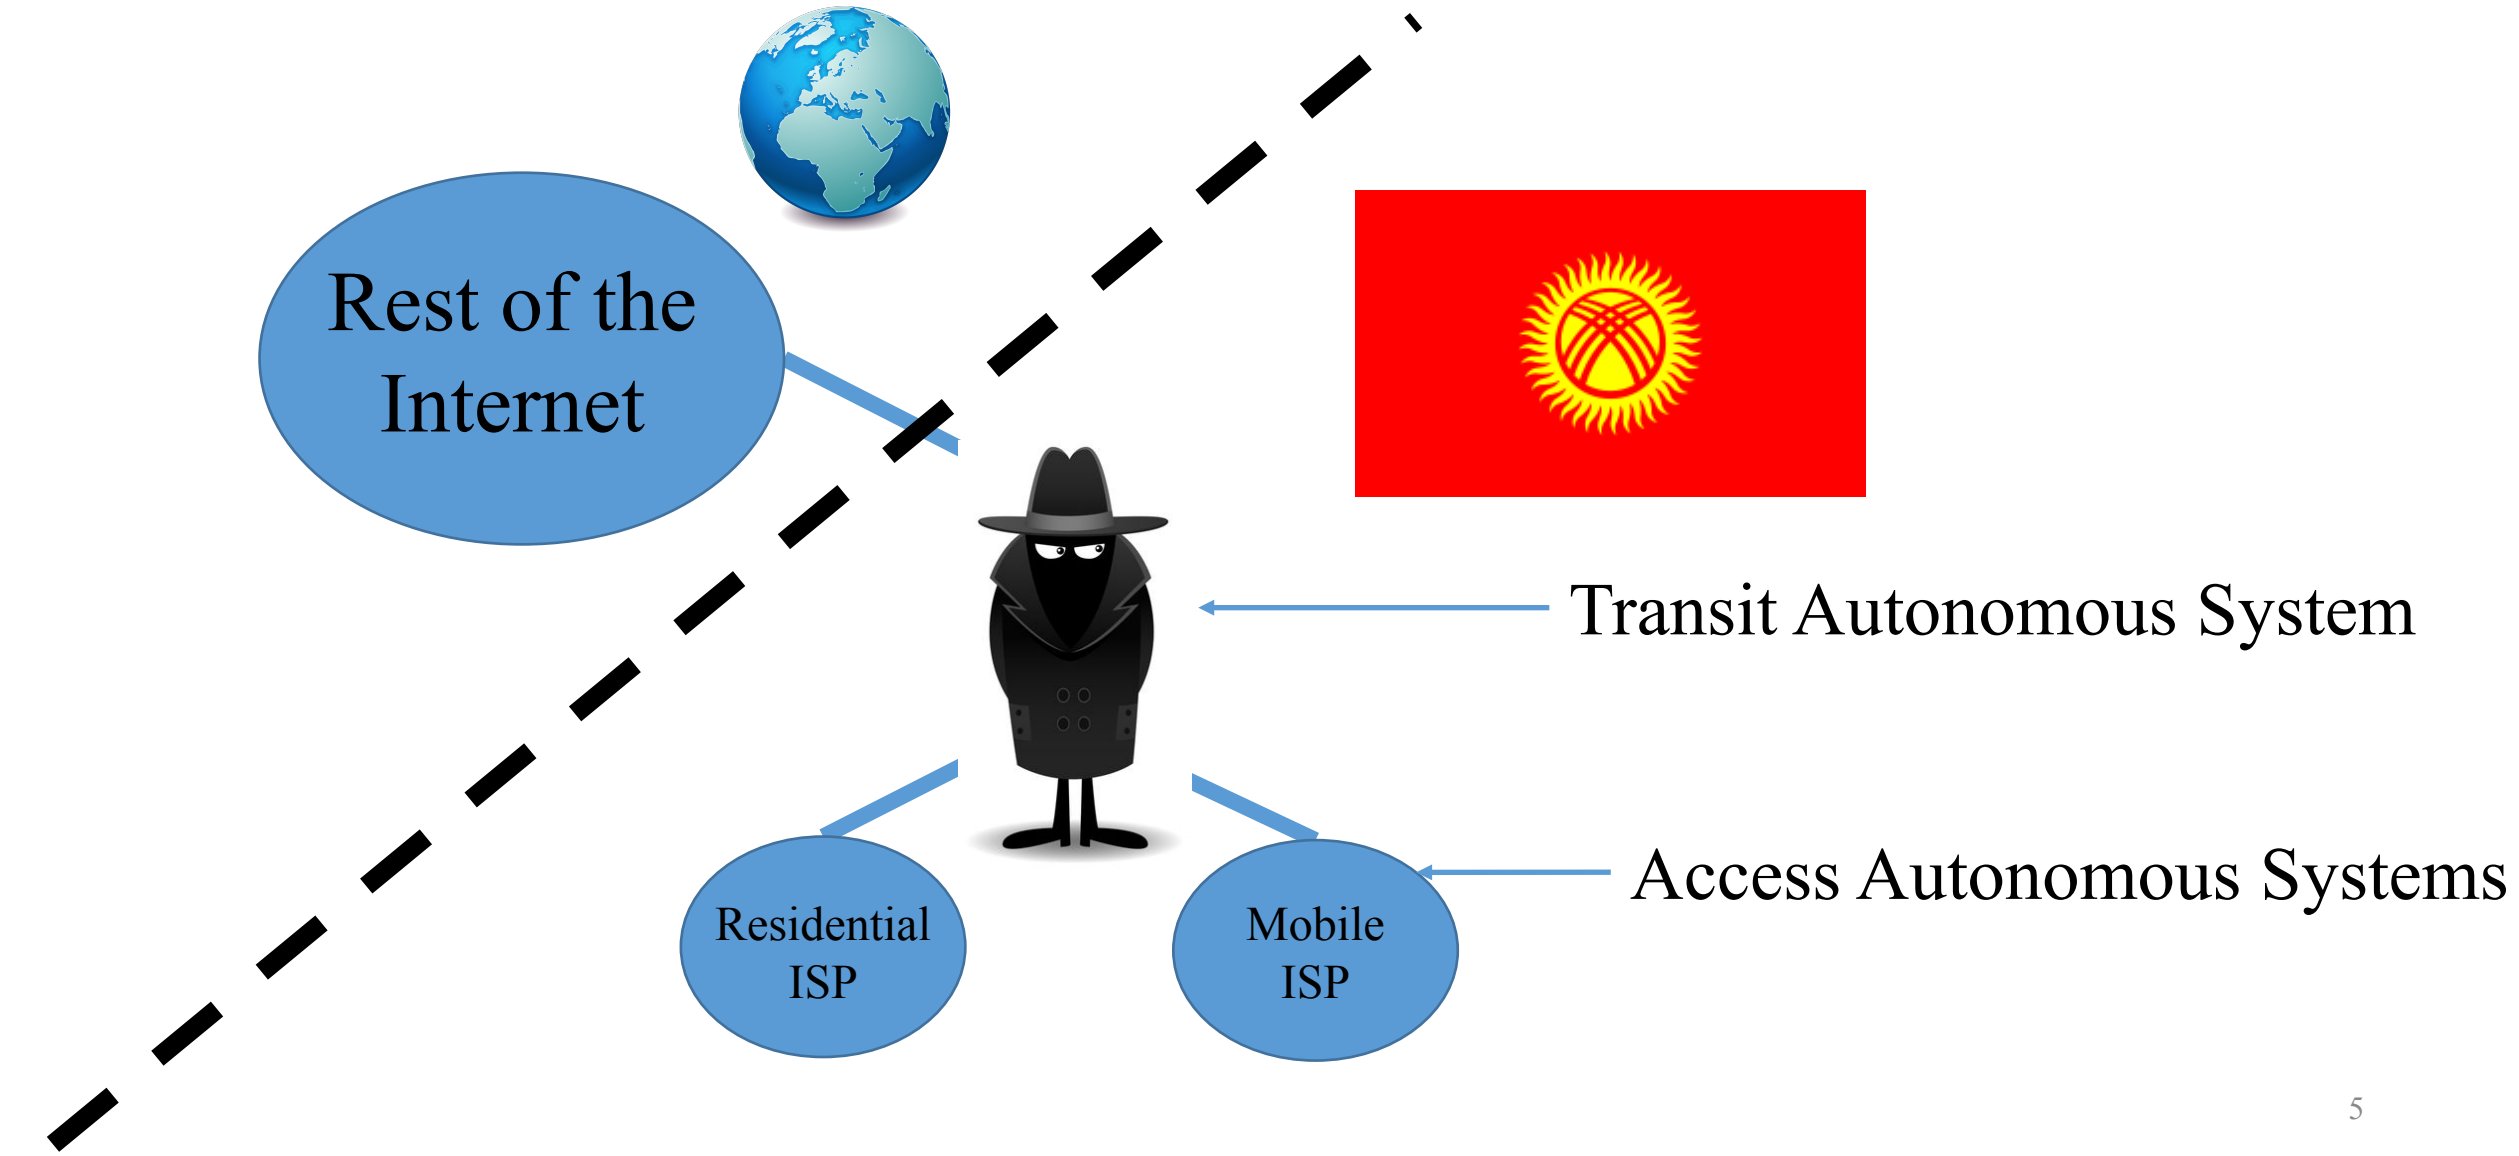

Automatic identification of the Autonomous Systems with the highest **transit influence**:

potential capability to ***observe, manipulate and disrupt*** Internet traffic flowing towards a country

Challenges in inferring transit influence

- Bias due to limited measurement infrastructure (ASes & countries)
- Measurements prone to missing both backup and preferred links
- No direct way to map prefixes to geographic location
- Massive scale of the global Internet (~10,000s of ASes and links)

Challenges in inferring transit influence

- Bias due to limited measurement infrastructure (ASes & countries)
- Measurements prone to missing both backup and preferred links
- No direct way to map prefixes to geographic location
- Massive scale of the global Internet (~10,000s of ASes and links)

Transit Influence Defined (1/3)

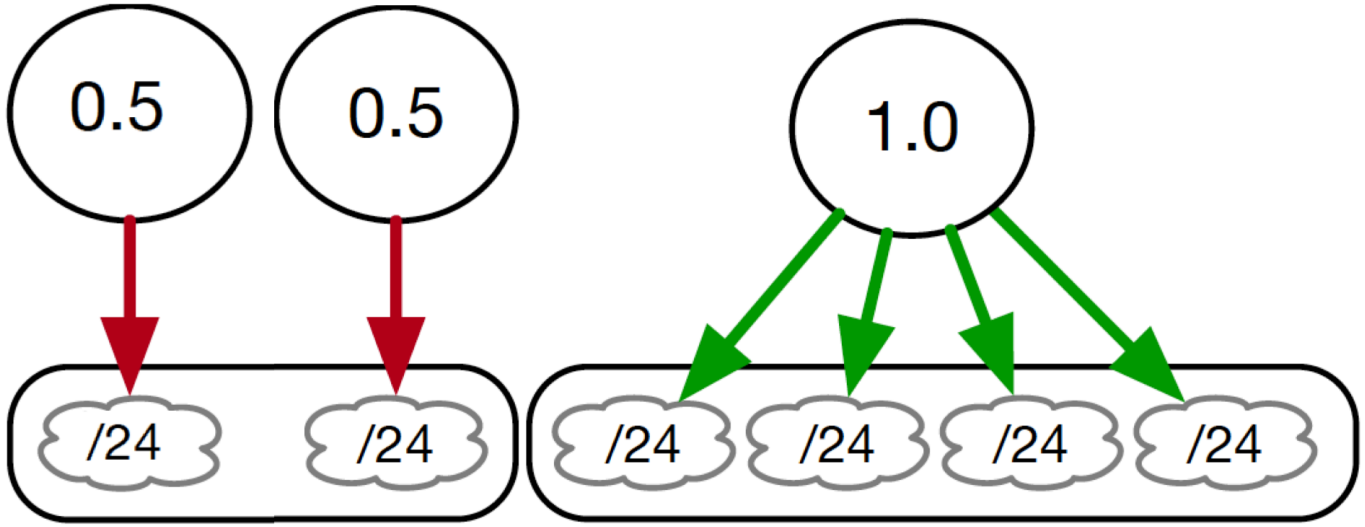

Transit influence of AS_t on AS_o :
Fraction of addresses originated by AS_o where AS_t is present as a transit provider

$$TI(AS_t, AS_o) = \sum_{m \in M} \frac{A(AS_t, AS_o)}{a(AS_o)} \cdot \frac{w(m)}{d(AS_t, AS_o)},$$

Transit Influence Defined (2/3)

$$TI(A,D) = TI(G,D) = 0.50$$

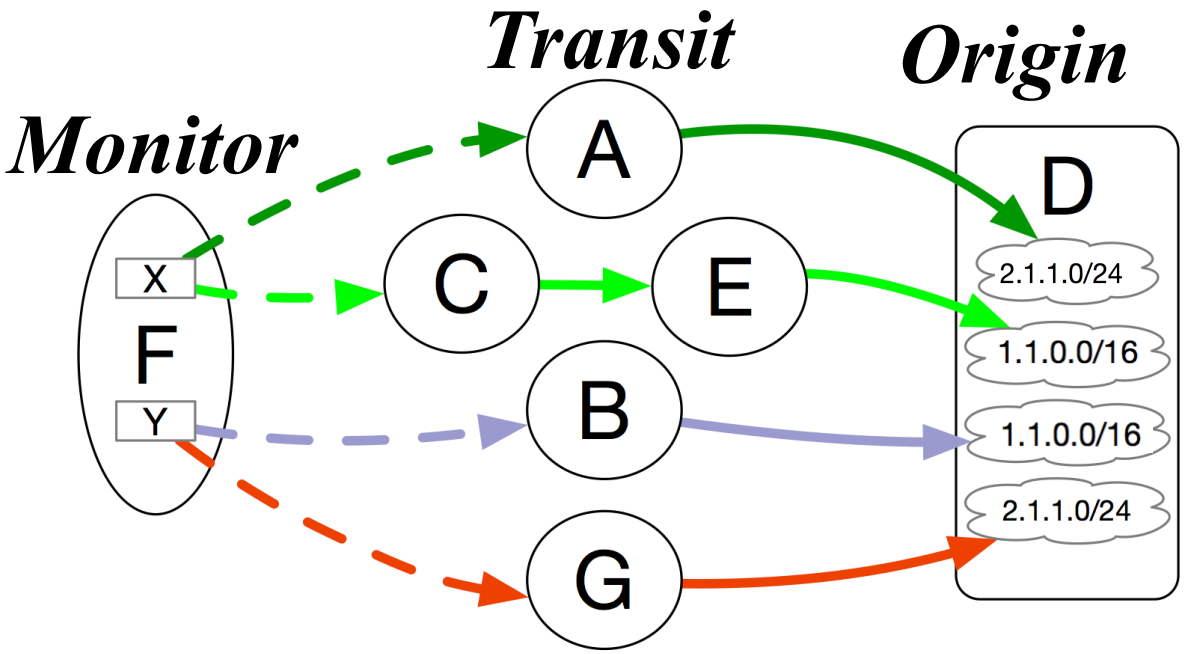

Method to filter observation bias: treat ASes with multiple monitors as a single monitor

$$TI(AS_t, AS_o) = \sum_{m \in M} \frac{A(AS_t, AS_o) \cdot w(m)}{a(AS_o) \cdot d(AS_t, AS_o)},$$

Transit Influence Defined (3/3)

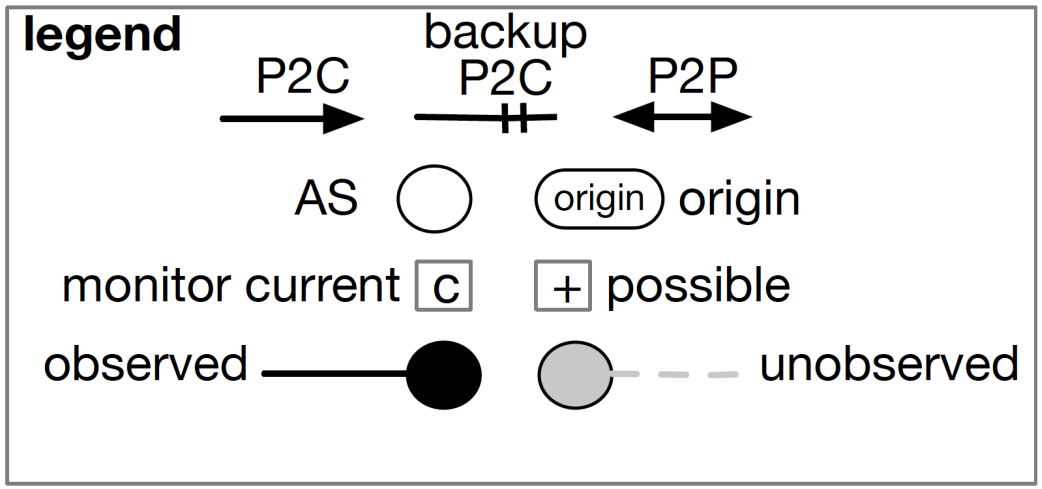

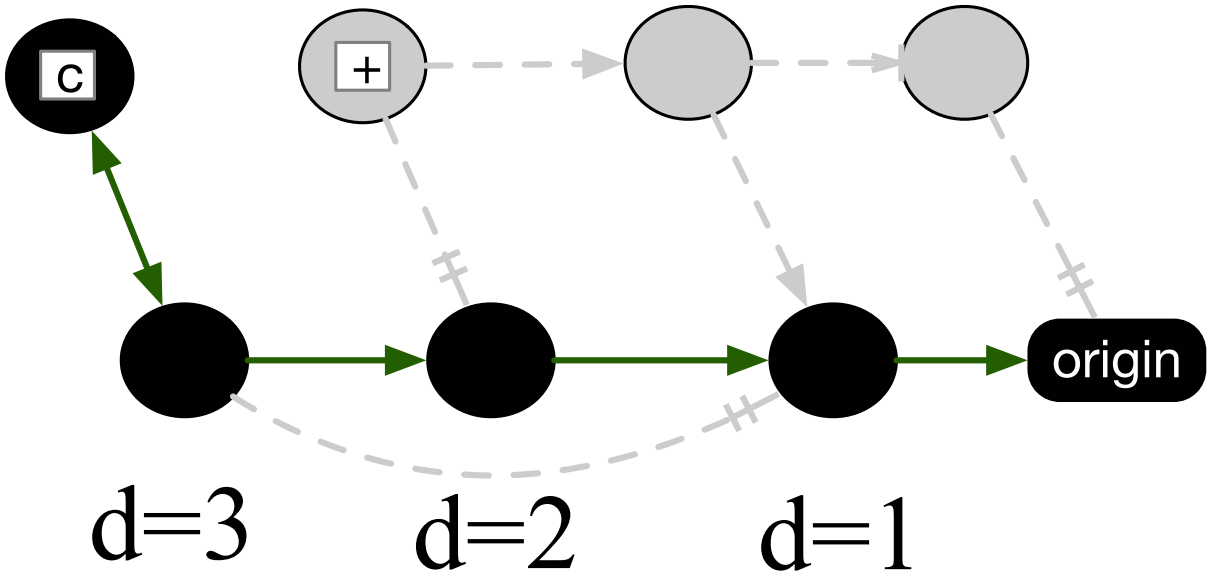

How to factor in limited visibility: discount the influence of indirect transit providers – as the likelihood of missing a link increases with the AS-level distance from the origin

$$TI(AS_t, AS_o) = \sum_{m \in M} \frac{A(AS_t, AS_o) \cdot w(m)}{a(AS_o) \cdot d(AS_t, AS_o)'}$$

Aggregate transit influence: how influential is this transit AS on the **country's** organizations?

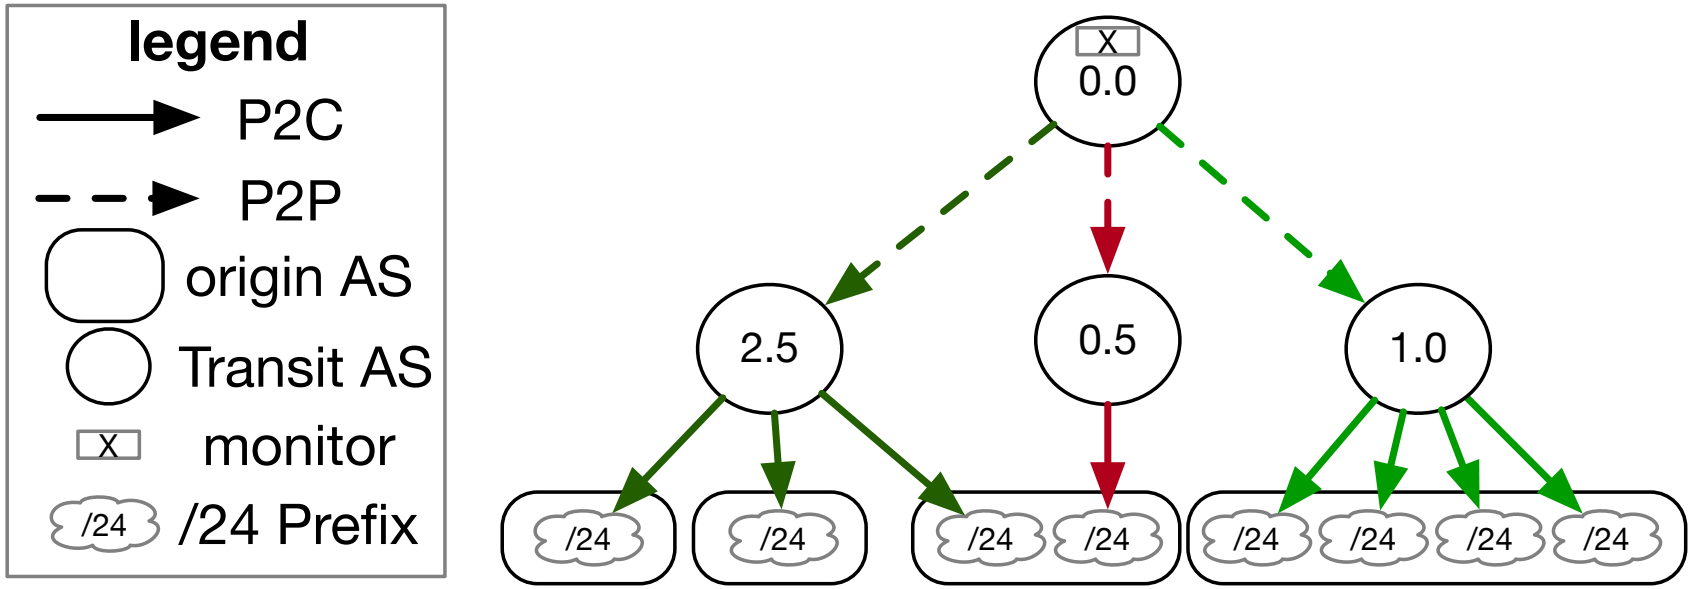

$$ATI(AS_t) = \sum_{o \in O} TI(AS_t, AS_o); AS_t \neq AS_o,$$

Add up the transit influence of AS_t on each of the country's origin ASes

How much address space does each additional compromised AS yield to the attack surface?

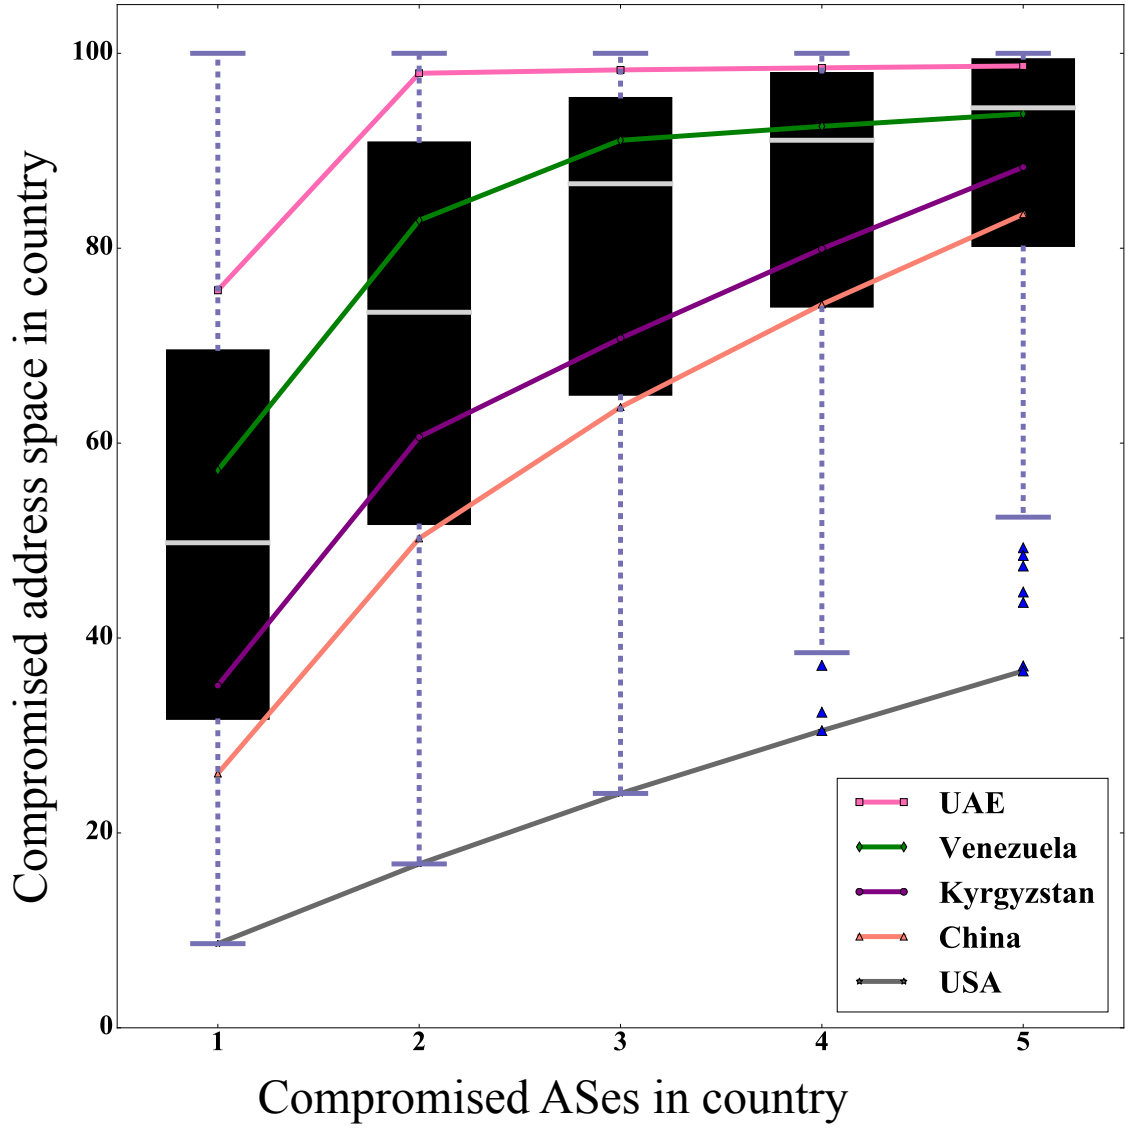

How much address space does each additional compromised AS yield to the attack surface?

Etisalat
(State-Owned)

CANTV
(State-owned)

Kvant Telecom
(Russian enterprise ISP)

China Telecom
(State-owned)

AT&T

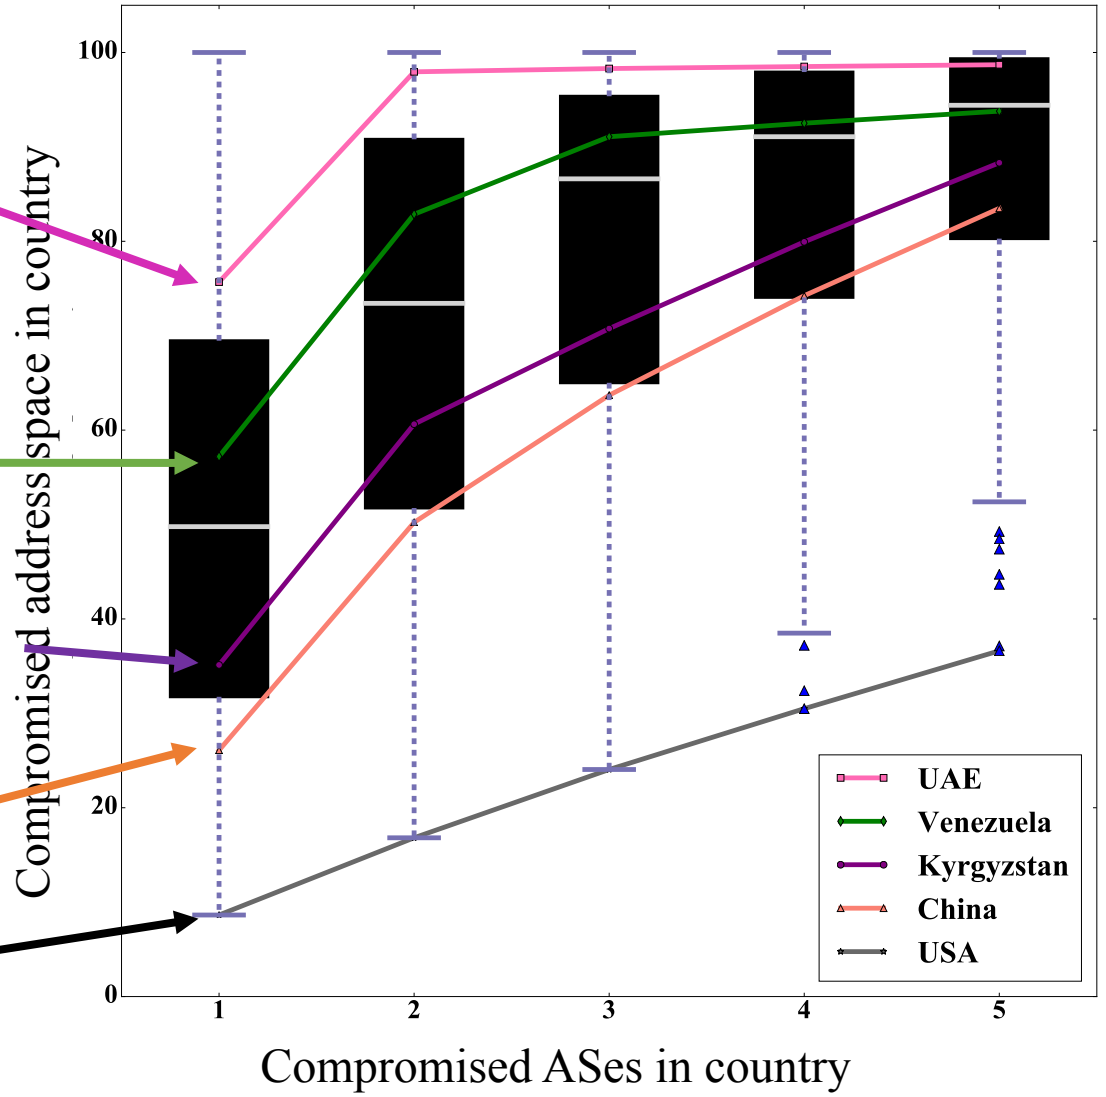

How much address space does each additional compromised AS yield to the attack surface?

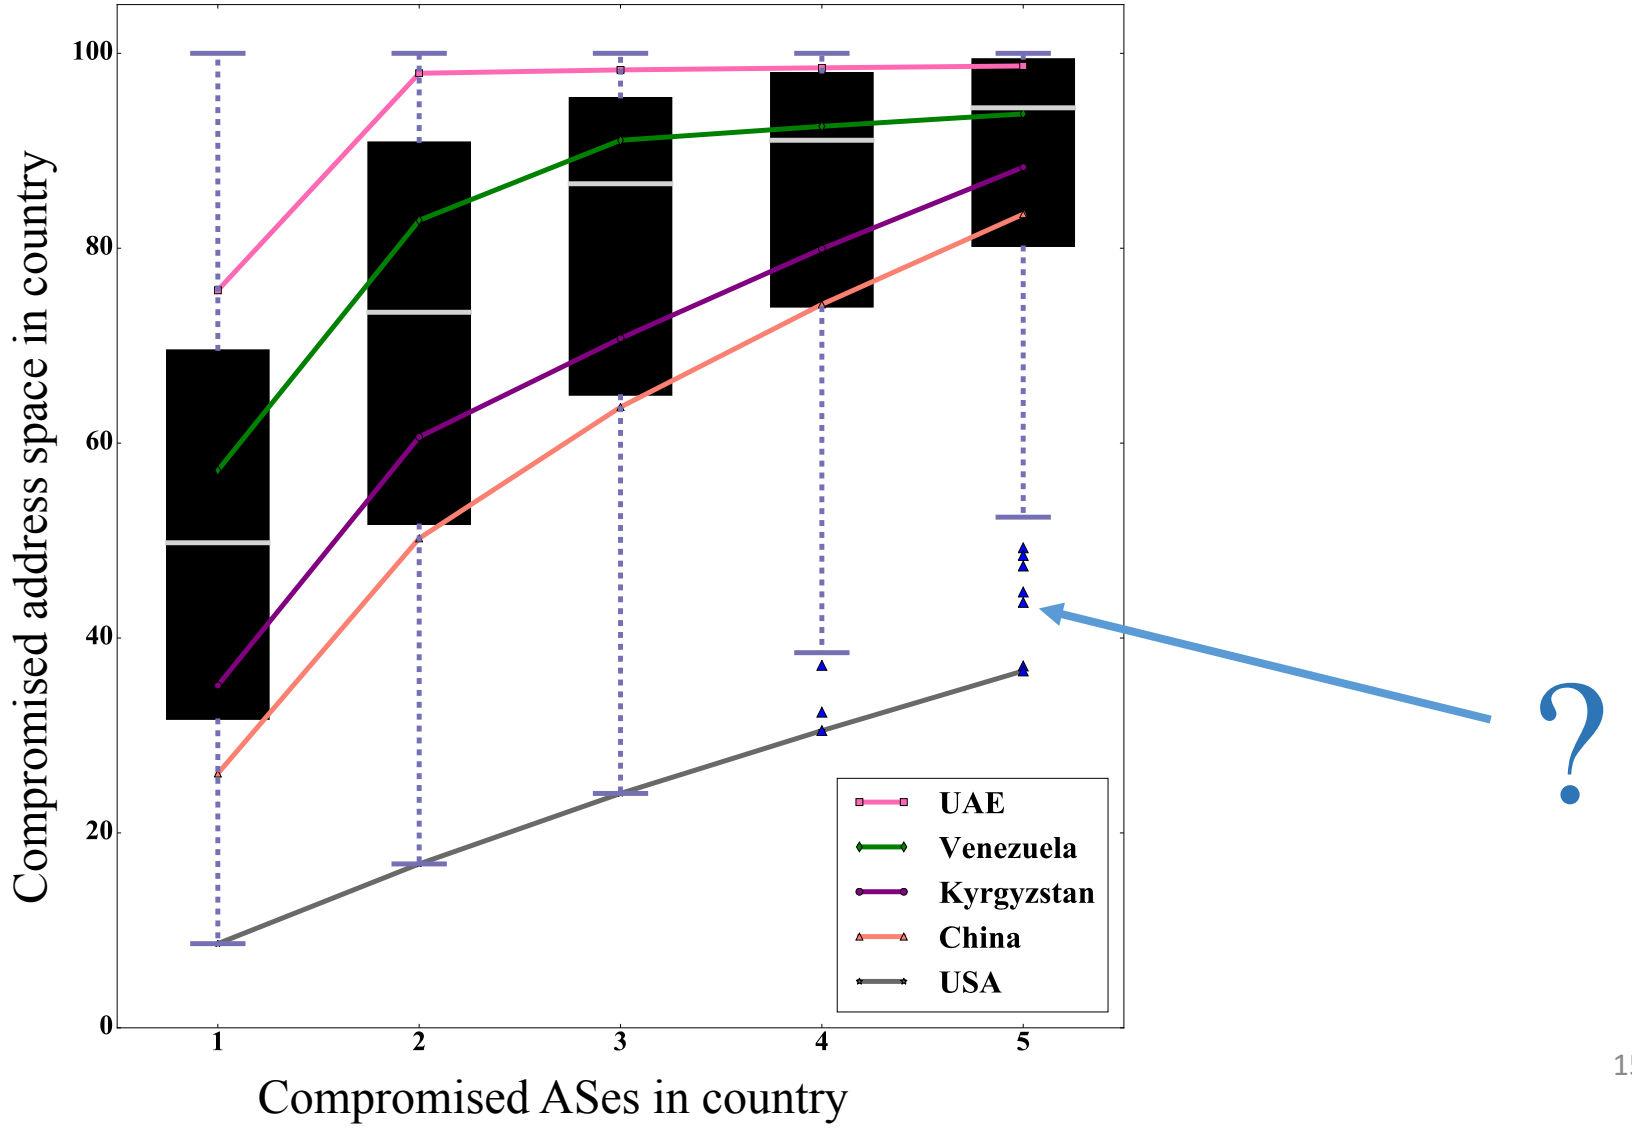

How much address space does each additional compromised AS yield to the attack surface?

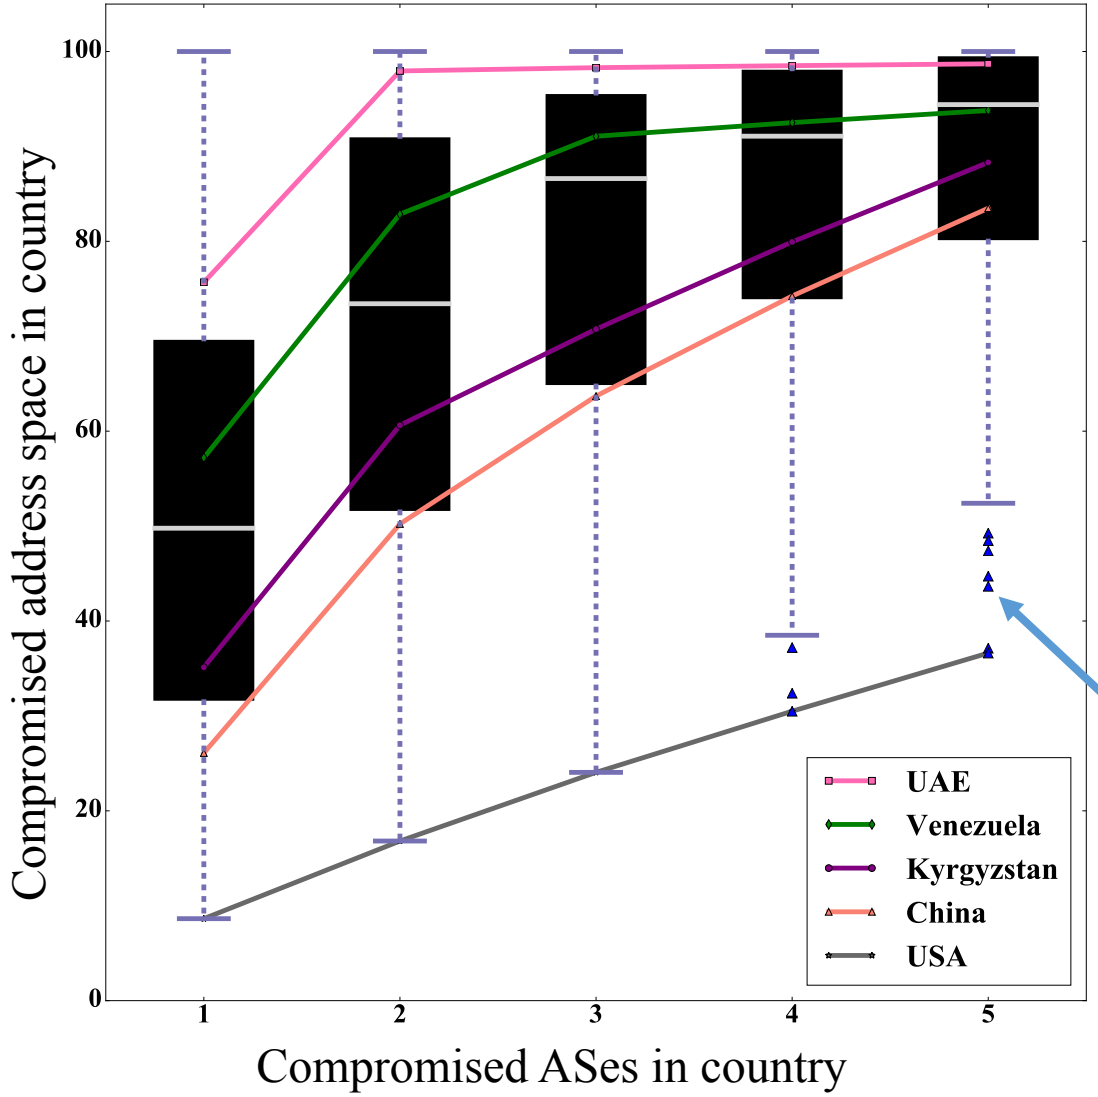

Bangladesh!

How much address space does each additional compromised AS yield to the attack surface?

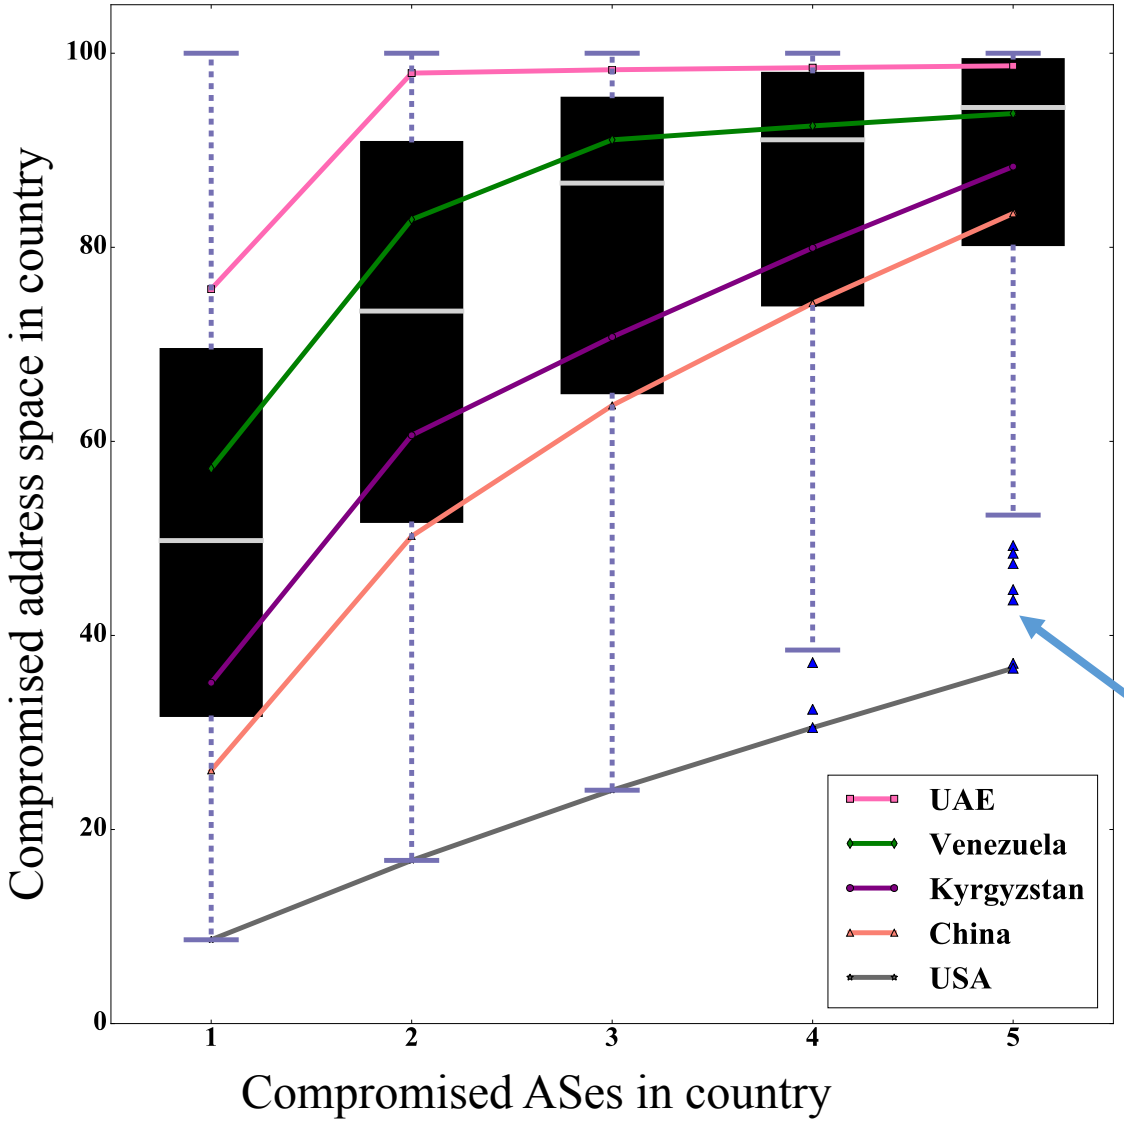

Bangladesh has a diverse ecosystem in their international connectivity that seems to have been enabled by a national government initiative to license International Internet Gateways (IIG) to:

"Ensure national security and protect national interest."

Bangladesh!

How much address space does each additional compromised AS yield to the attack surface?

- We have preliminarily found that stronger democracies tend to have less concentrated address spaces.
- Wealth is not as good a predictor of that concentration.

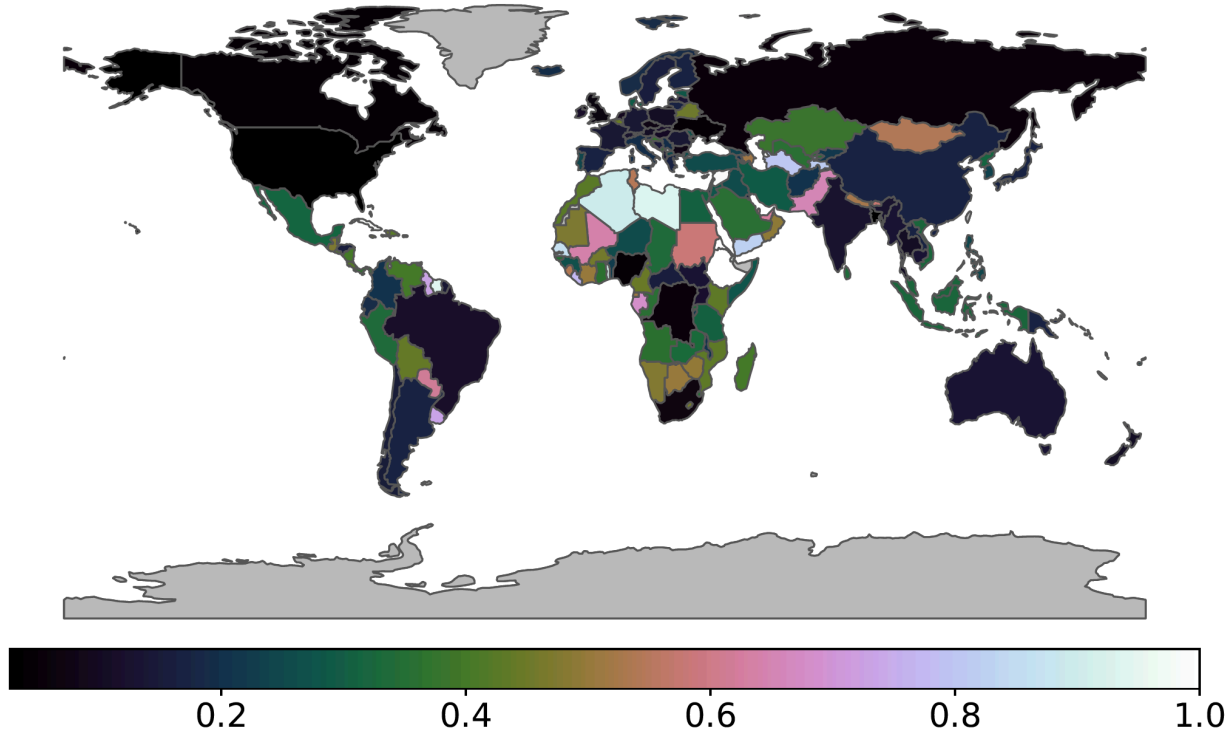

Herfindahl Index

Thank you!

alexander@caida.org

cseweb.ucsd.edu/~agamerog

Backup Slides

ATI and related metrics: Customer Cone and Hegemony

- Customer cone: who can this transit AS reach (and potentially extract revenue from) through its customers in this country?

- Hegemony: who are the most central ASes in this country's network given the bias of the measurement infrastructure?

- ATI: whose international traffic can this transit AS potentially observe, manipulate or disrupt?
  - Draws on both customer cone and hegemony
  - Expands on prioritizing measurement AS diversity and factoring in unobserved links

- An AS with a high customer cone but low ATI may reach most of the origin ASes in this country indirectly
  - through its customers, as with AS A
  - AS A has a high AS-footprint in the country, and high potential revenue as the AS would be paid for forwarding traffic towards its customers

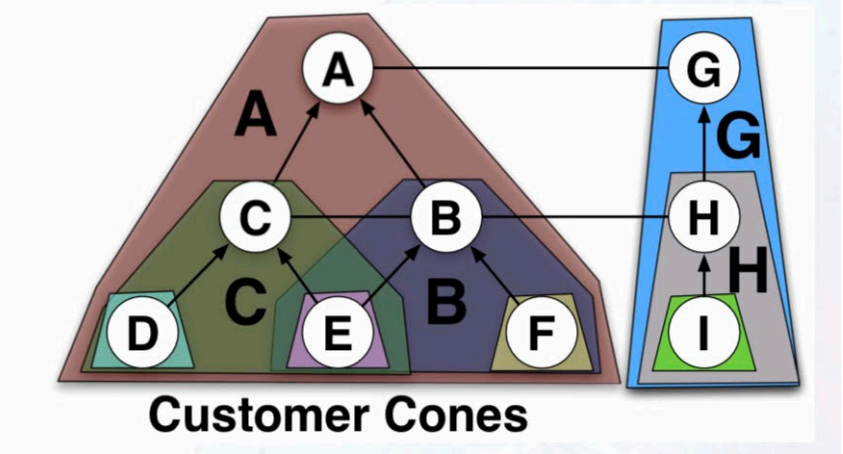

- Since AS A is an indirect provider of ASes D,E,F, it would have a lower ATI than both AS C and AS B
  - The probability that there are unseen backup or preferred links increases with the AS-level distance
  - Direct transit providers are more likely to have the capability to observe, manipulate or disrupt the origin's traffic

Diagram from Huffaker et. al 2018

- Hegemony: who are the most central ASes in this network given the bias of the measurement infrastructure?
  - Consistency within tiers
  - Filter centrality of ASes with monitors

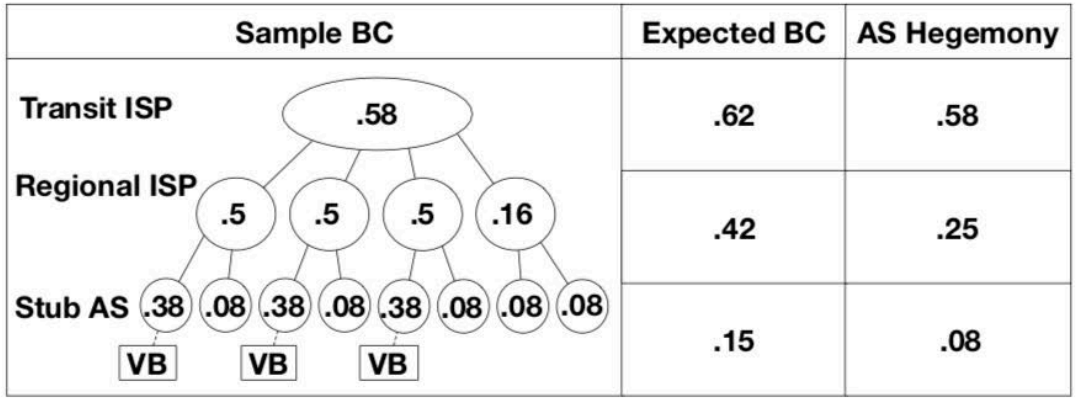

Diagram adapted from Huffaker et. al 2019

- ATI: whose international traffic can this transit AS potentially observe, manipulate or disrupt?
  - Expand on filtering of observation bias by treating multi-monitor ASes as one monitor
  - Coarsely factor in unobserved backup links with indirect transit filter, dividing the TI by the AS-level distance from the origin

- In the above diagram, ATI would assign no influence to the origin ASes (as we are concerned primarily with international transit), and would filter the TI of the "Transit ISP" at the top given that it is never directly connected to the origin AS

South Africa: the rise of **Seacom**

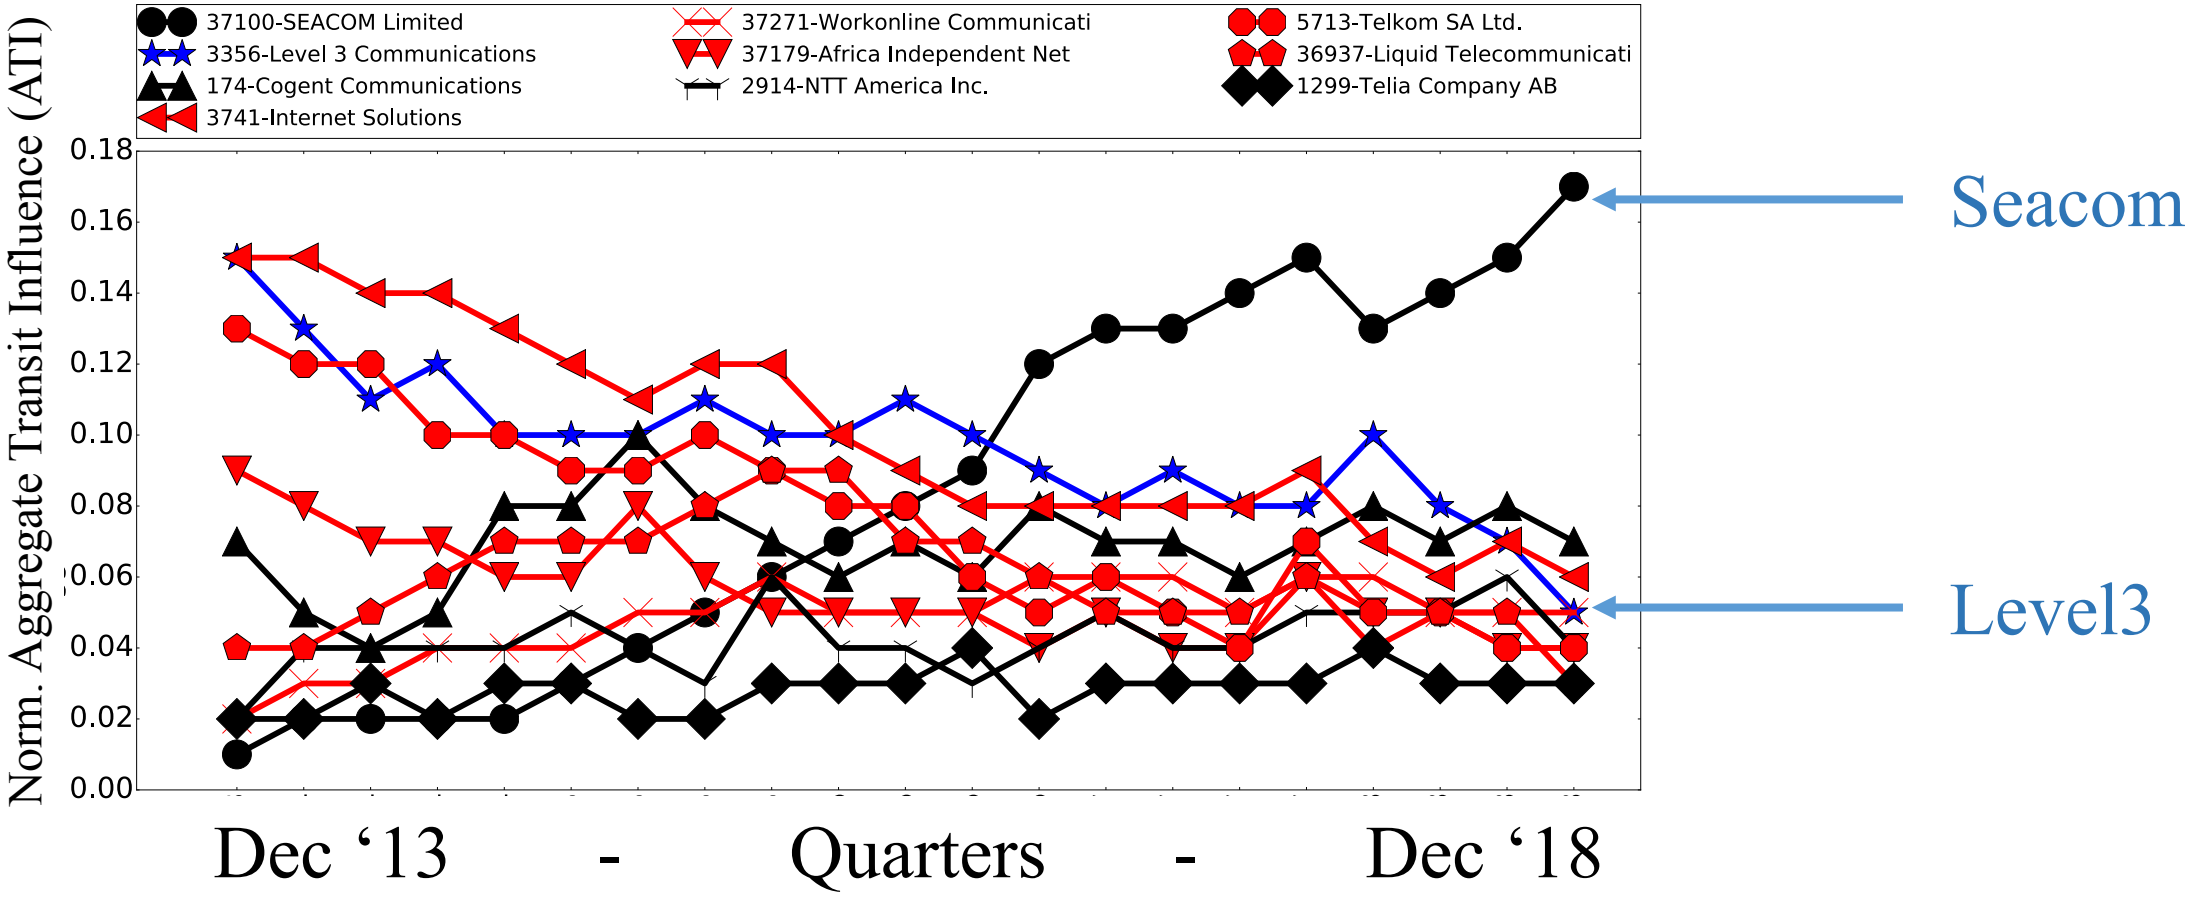

Russian ecosystem is dominated by domestic ASes

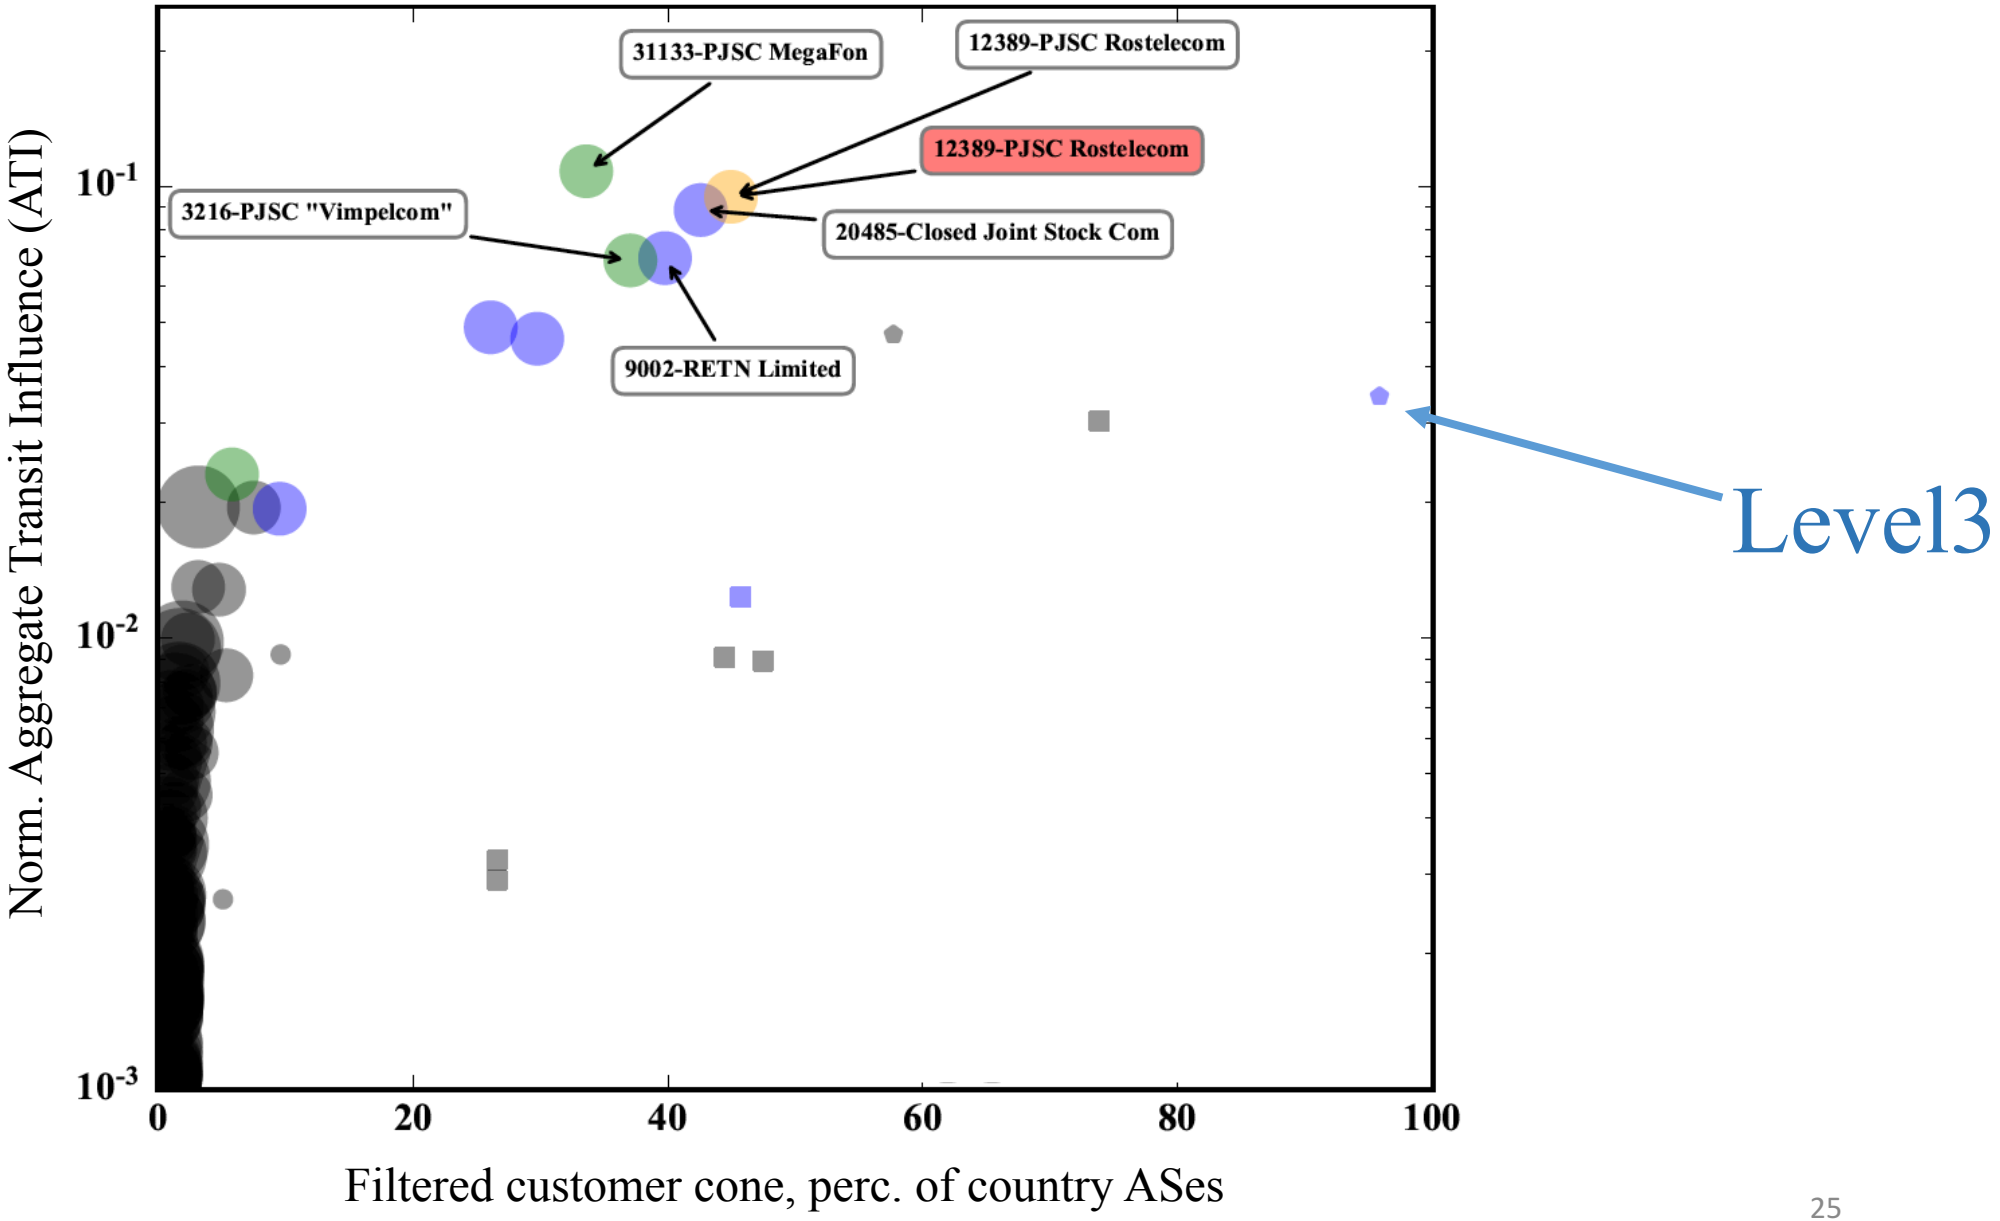

How do ASes communicate?

- ASes announce routing policies using the **Border Gateway Protocol** (BGP)
- Announcements are made to neighbors

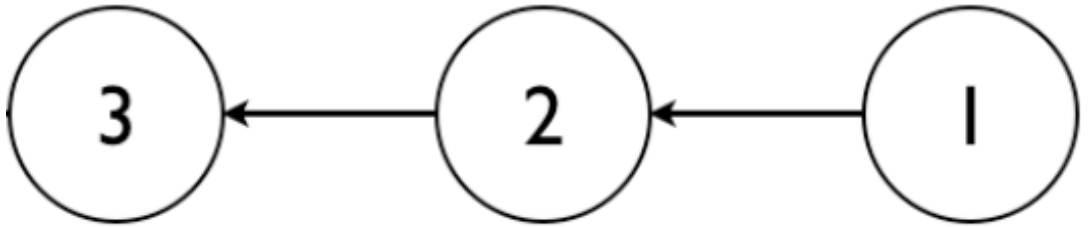

1.2.3.0/24

- In this toy example, AS1 announces that it "owns" the 256 addresses
  - AS2 has a direct link to AS1 and can just forward packets to it
  - AS3 knows it can direct traffic towards those IP addresses to AS2

Image slightly modified from Roberts '11

$$TI(AS_t, AS_o) = \sum_{m \in M} \frac{p(AS_t, AS_o) \cdot w(m)}{a(AS_o) \cdot d(AS_t, AS_o)}, \text{ where}$$

- *TI* is the **transit influence** ~ [0,1]
- AS_t is the transit AS
- AS_o is the origin AS
- $p(AS_t, AS_o)$ is the number of the origin's addresses that the transit serves
- $a(As_o)$ is the total number of addresses originated by As_o in this country
- $d(AS_t, AS_o)$ is the AS-level distance between As_o and AS_t
- *m* is the BGP monitor in the set of *M* available monitors

BGP Announcements and Transit Influence

- Goal: extract country-level transit influence of ASes from Border Gateway Protocol (BGP) announcements
- Find prefixes and ASes that are relevant to each country (*geolocate*)
- Find transit providers servicing many of those prefixes and ASes

Sample parsed BGP announcement

```
ripe/rrc00|5 1836|13030|61716 131.72.23.0/24 i 146.228.1.4
```

Collector | Host_AS | Transit AS | Origin_AS | Prefix | Monitor IP

Latin American Domestic Providers – Top 10 by ATI

Country	ATI Rank	AS Number – AS Name	Primary Purpose	Revenue USD (millions)
Brazil	3	16735-ALGAR TELECOM S/A	Telephony	500
Brazil	4	7738-Telemar Norte Leste S. ^	Telephony	7,200
Brazil	6	262589-INTERNEXA Brasil Opera	IP Transit	1,700
Brazil	**8**	**10429-Telefonica Data S.A.**	IP Transit	**60,300***
Mexico	3	8151-Uninet S.A. de C.V.	IP Transit	47,350~
Mexico	4	18734-Operbes S.A. de C.V.	IP Transit	5,300
Mexico	*6*	*11172-Alestra S. de R.L. de*	*IP Transit*	*415*
Mexico	7	32098-Transtelco Inc	IP Transit	40
Mexico	*8*	*6503-Axtel S.A.B. de C.V.*	*Telephony*	*415*
Mexico	9	17072-TOTAL PLAY TELECOMUNIC ^	Telephony	4.5
Colombia	2	23520-Columbus Networks USA	Cable	20,000
Colombia	5	18678-INTERNEXA S.A. E.S.P	IP Transit	1,700
Colombia	**10**	**3816-COLOMBIA TELECOMUNICAC ^**	Telephony	**60,300***
Peru	**9**	**6147-Telefonica del Peru S. ^**	Telephony	**60,300***
Venezuela	1	8048-CANTV Servicios Venez	Telephony	3,750
Venezuela	8	11562-Net Uno C.A.	Telephony	3,000
Chile	2	14259-Gtd Internet S.A.	Enterprise Broadband	32
Chile	**3**	**7004-CTC Transmisiones Regi**	Telephony	**60,300***
Chile	9	27978-Telmex Servicios Empre	IP Transit	47,350~
Chile	10	6471-ENTEL CHILE S.A.	Telephony	2,900

Global revenue of Telefonica* and America Movil~

Top Domestic Companies by ATI. There is diversity in their size; revenues range from USD 32 million to over 60 billion.
There are also three state-owned providers (in red). Four of these companies are owned by Telefonica (bold), two by America Movil (underlined), and two by Alfa (Italics).
All but four (marked with ^) of these companies are known to provide IP transit commercially.

State-Owned Providers

Country	ATI Rank	AS Number – AS Name	Avg. Influence	Perc. ASes in Country	Revenue USD
Brazil	6	**262589-INTERNEXA Brasil Opera**	0.22	17.53	1,700,000,000
Colombia	5	18678-INTERNEXA S.A. E.S.P	0.81	5.65	1,700,000,000
Venezuela	1	8048-CANTV Servicios Venez	0.65	22.22	3,750,000,000

- Internexa, a subsidiary of a company owned by the Colombian government, provides transit to about 18% of Brazilian Ases with an average influence of 0.22.
- CANTV is the largest telecommunications provider in Venezuela with a large AS footprint (22%) and dominant influence (Avg. Influence 0.65).
- CANTV also has 16 million direct customer accounts.

Telefonica and America Movil Subsidiaries

Country	ATI Rank	AS Number – AS Name	Commercial Name	Avg. Influence	Perc. ASes in Country	Subscribers (millions)
Brazil	8	10429-Telefonica Data S.A.	Telefonica Vivo	0.147	24.97	NA (Enterprise)
Colombia	10	3816-COLOMBIA TELECOMUNICAC	Movistar	0.811	5.65	18
Peru	9	6147-Telefonica del Peru S.	Movistar and Tuenti	0.928	3.49	21
Chile	3	7004-CTC Transmisiones Regi	Movistar	0.48	17.06	10
Mexico	3	8151-Uninet S.A. de C.V.	Uninet	0.577	16.25	NA (Enterprise)
Chile	9	27978-Telmex Servicios Empre	Telmex	0.283	13.99	NA (Enterprise)

- Four domestic companies are subsidiaries of Telefonica. Three of them are major mobile operators, with upwards of 10M subscribers.
- Telefonica's subsidiaries are also dominant providers, with avg. influences ranging from 0.48 to 0.93.
- Another two companies are subsidiaries of America Movil (a Mexican company) in Mexico and Chile.