


CAIDA's AS-rank: measuring the influence of ASes on Internet Routing

Matthew Luckie

Bradley Huffaker

Amogh Dhamdhere

k claffy



<http://as-rank.caida.org/>

Overview

1. Inferring **AS relationships** using publicly available BGP paths
 - views of ~400 ASes at Route Views and RIPE RIS
2. Inferring the **influence of ASes** based on their “customer cone”
 - Traffic in your customer cone stays on-net and is the most profitable (when it reaches you)

<http://as-rank.caida.org/>

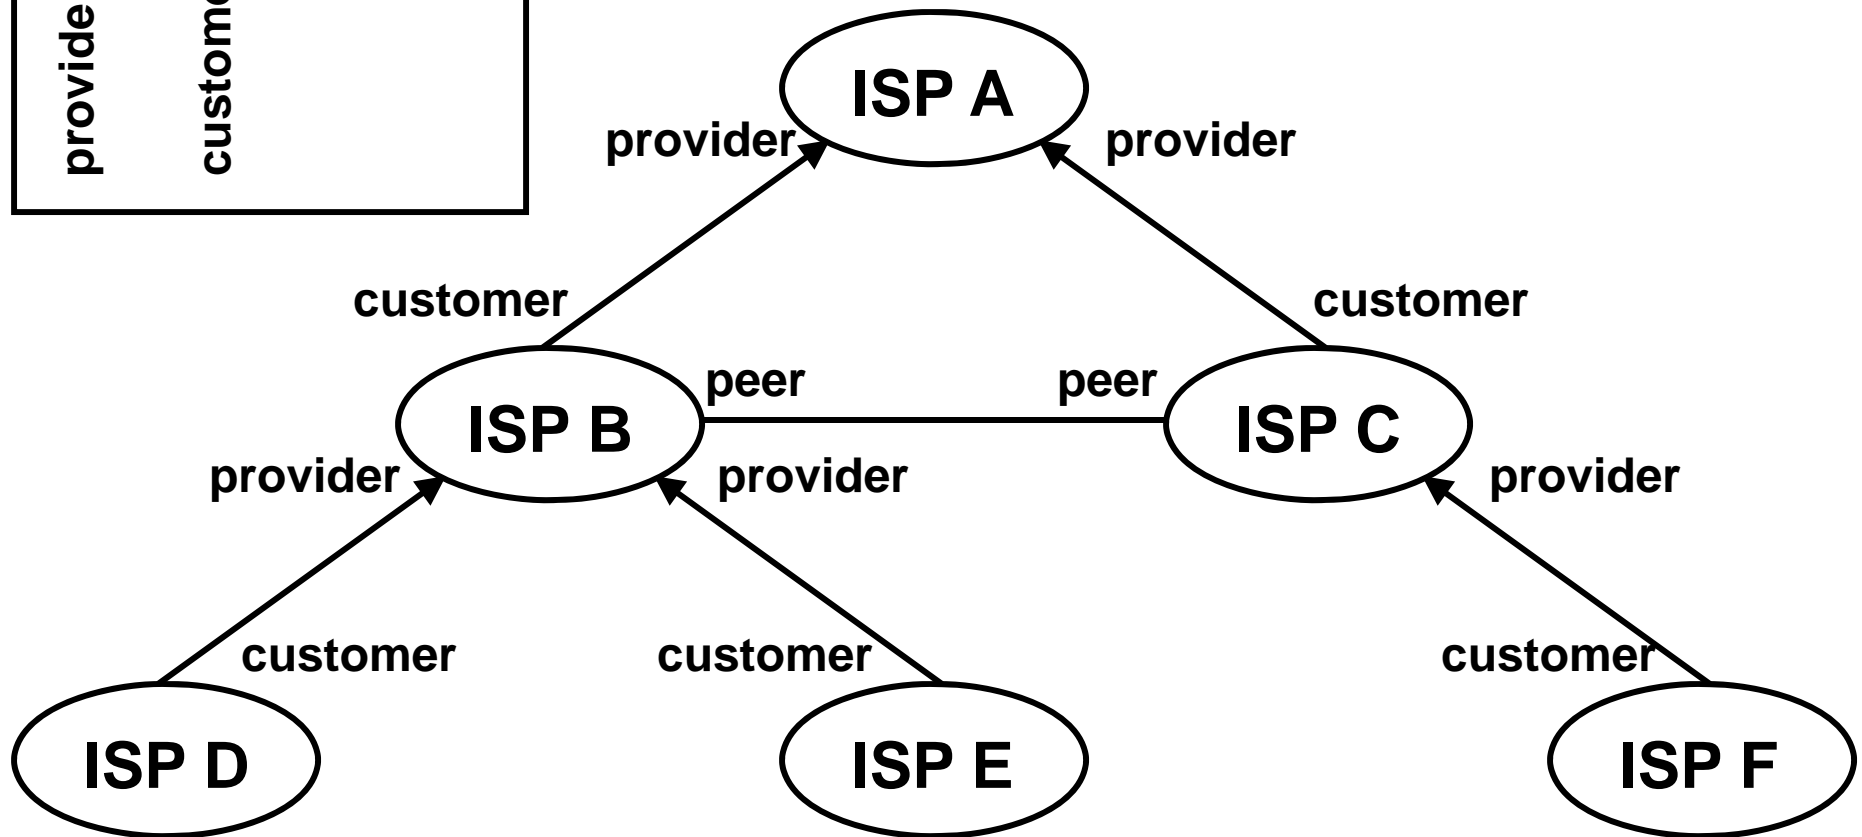
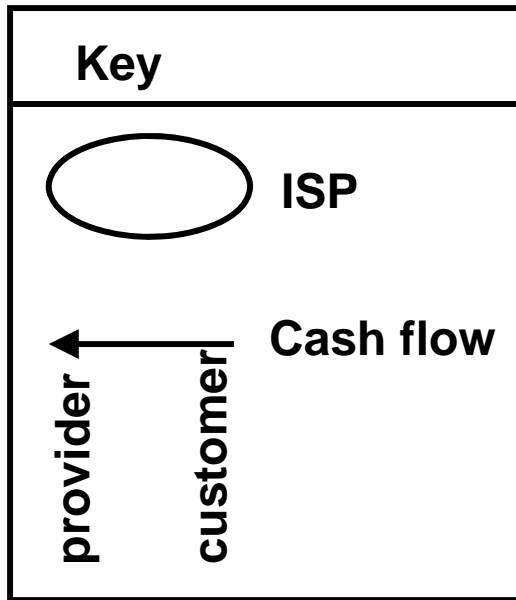
Roadmap

- Objectives
- Methodology
 - Definitions, assumptions, and caveats
 - Algorithm
 - Illustration
- Results
- Open Problems
 - How NANOG can help!

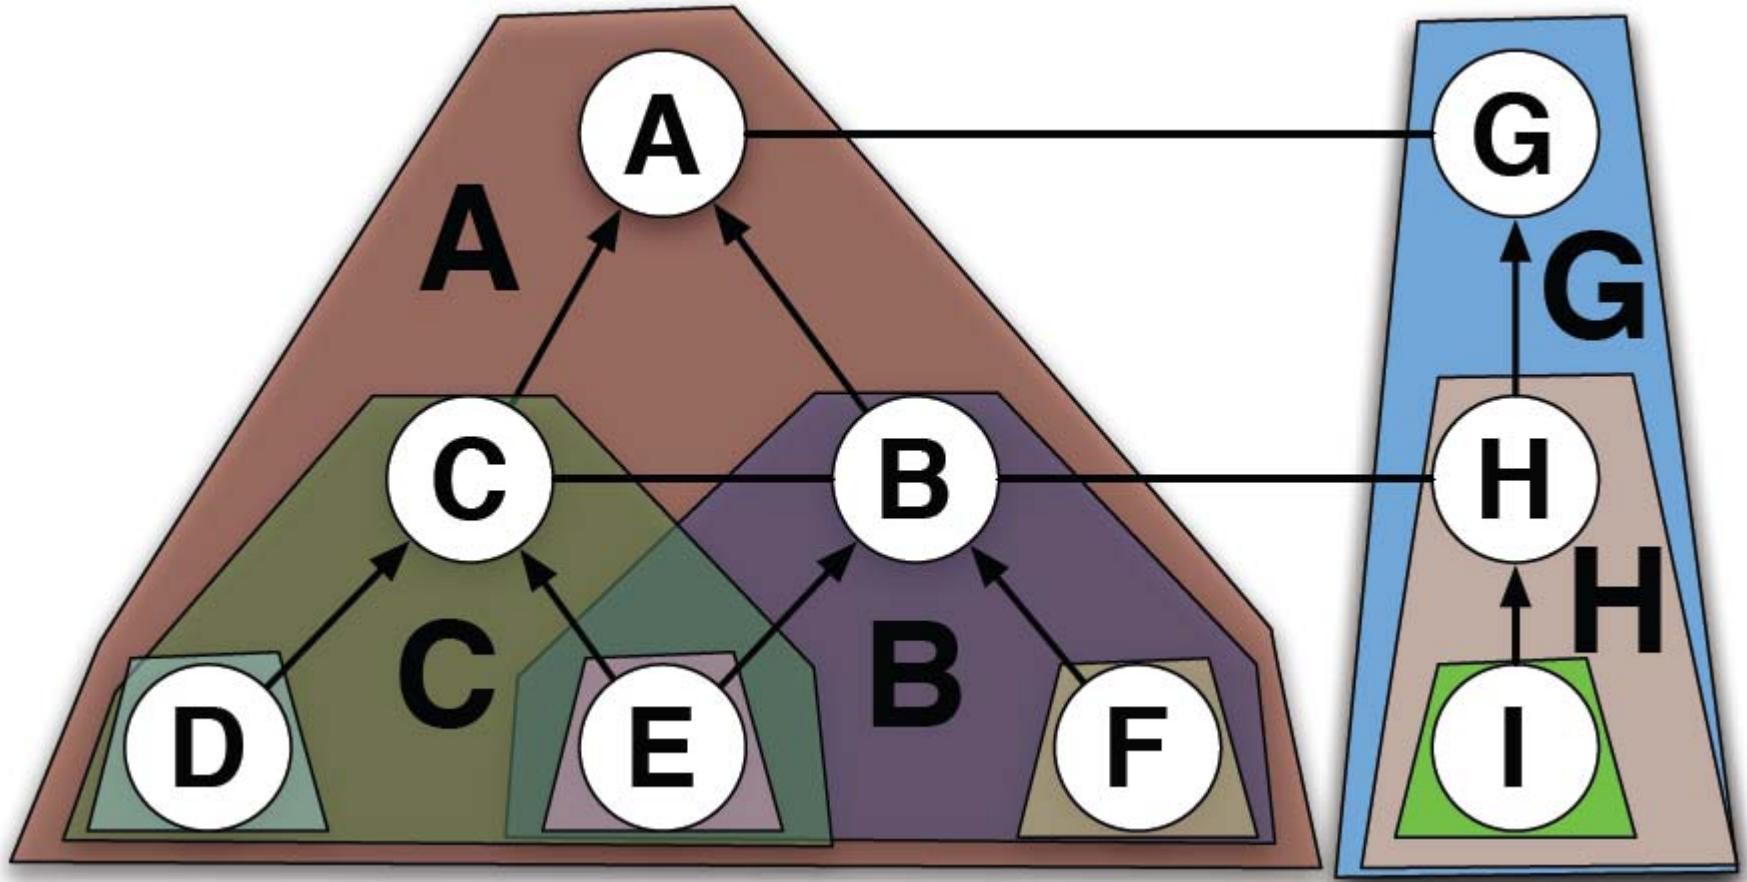
Objectives

- AS relationships determine how traffic flows through the Internet
- Having accurate AS relationships supports:
 - Modeling infrastructure **security and stability**
 - Analysis of **BGP mis-configuration**
 - Analysis of the **influence of ASes** on Internet routing using “customer cone”
 - What is the absolute number of ASes a provider could charge traffic for?

AS Relationships



Definition – Customer Cones



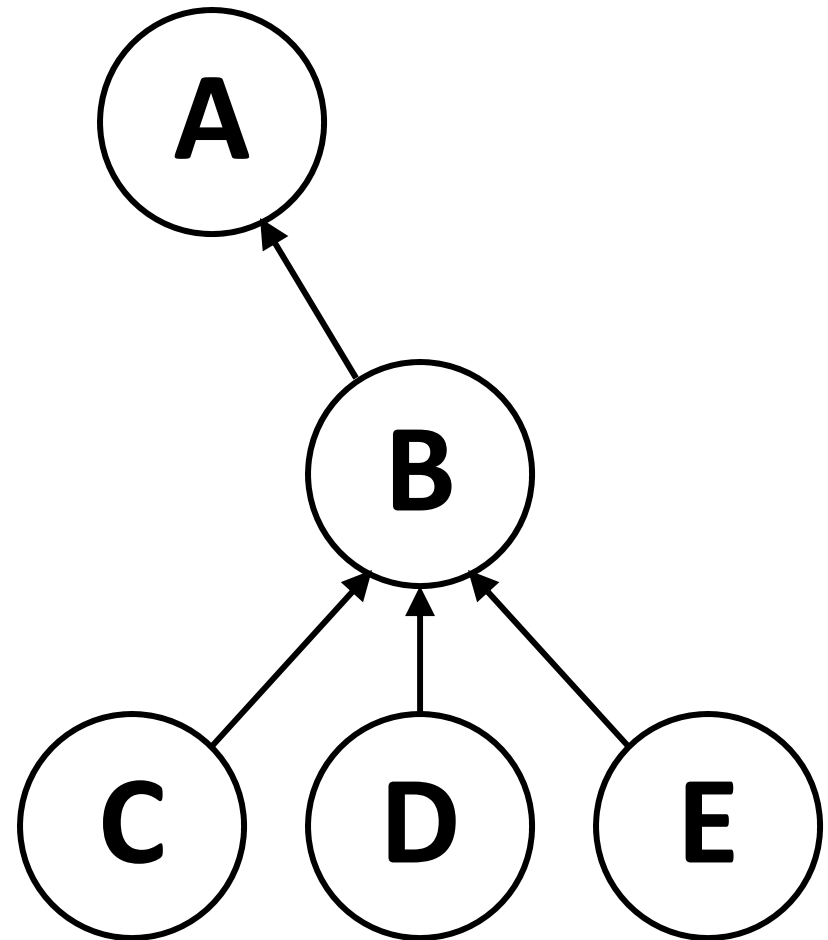
A's customer cone: A, B, C, D, E, F

B's customer cone: B, F

C's customer cone: C, D

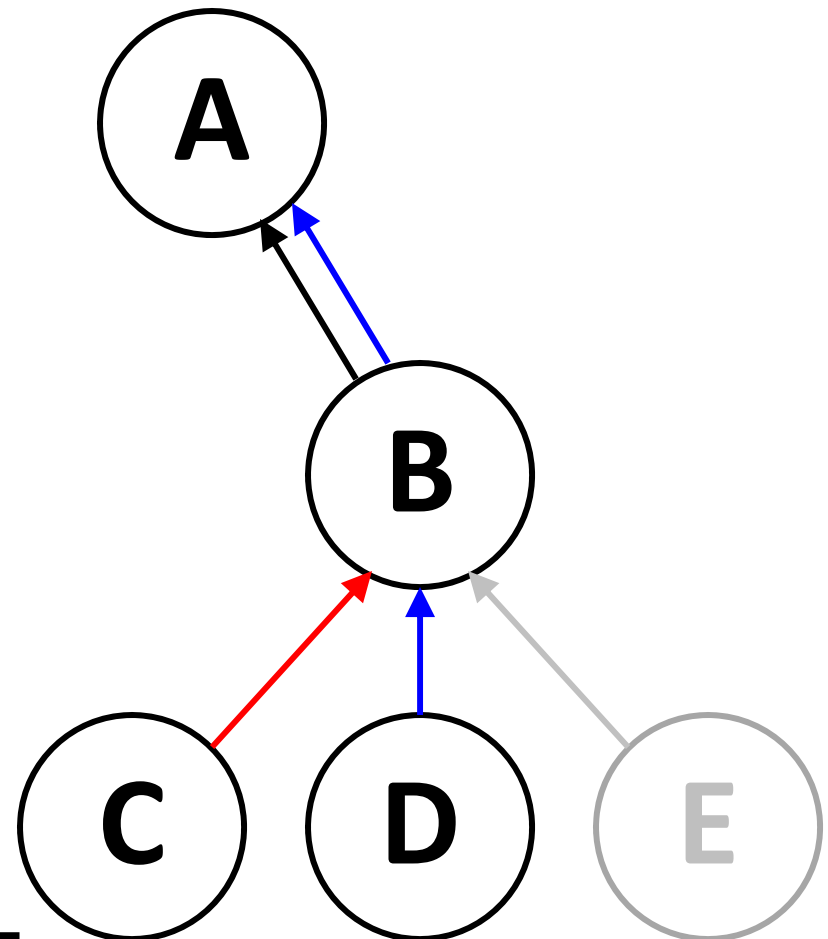
Customer Cone Computation (1)

BGP Paths
A B C
A B D
B E



Customer Cone Computation (1)

BGP Paths
A B C
A B D
B E

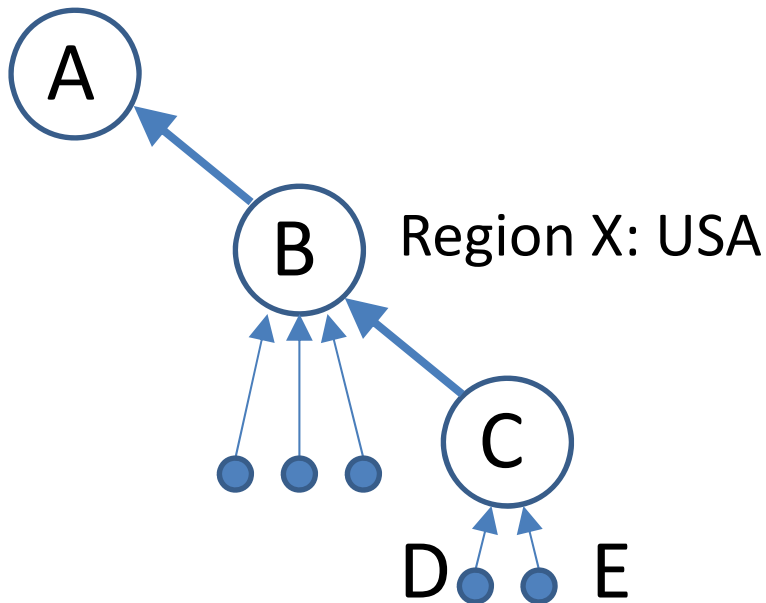


Customer cones for A:

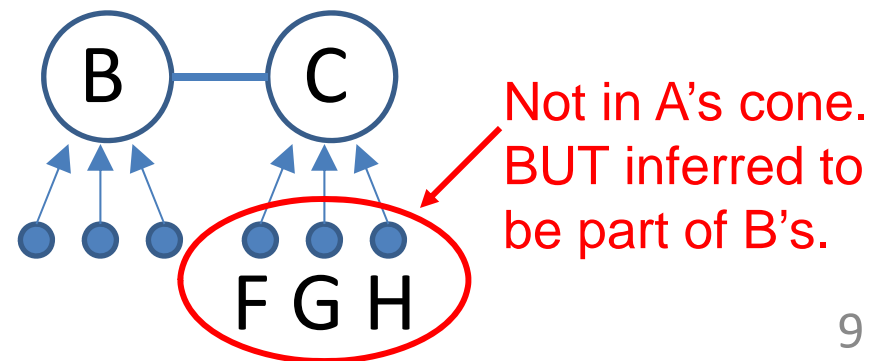
- **Recursive (5):** A, B, C, D, E
- **Observed (4):** A, B, C, D

Customer Cone Computation (2)

- AS relationships are complex: two ASes may have a c2p relationship in one location, but p2p elsewhere
- Use observed BGP paths to limit damage
 - B-C is inferred to be p2c, and B's customer cone contains all of C's customers (D, E, F, G, H)
 - A will only contain C's customers it learns from B (D+E).

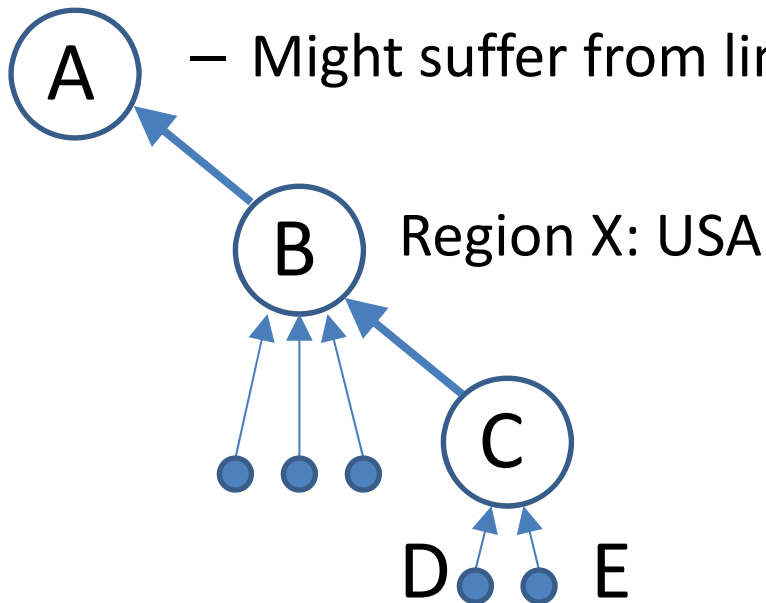


Region Y: "Europe"

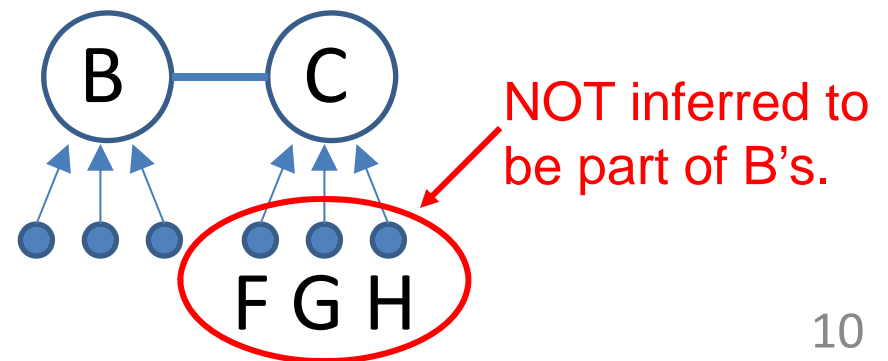


Customer Cone Computation (3)

- AS relationships are complex: two ASes may have a c2p relationship in one location, but p2p elsewhere
- Define customer cone based on provider/peer observed view of an AS
 - A sees D and E as indirect customers via B, so B's customer cone only includes D, E from C.
 - Might suffer from limited visibility

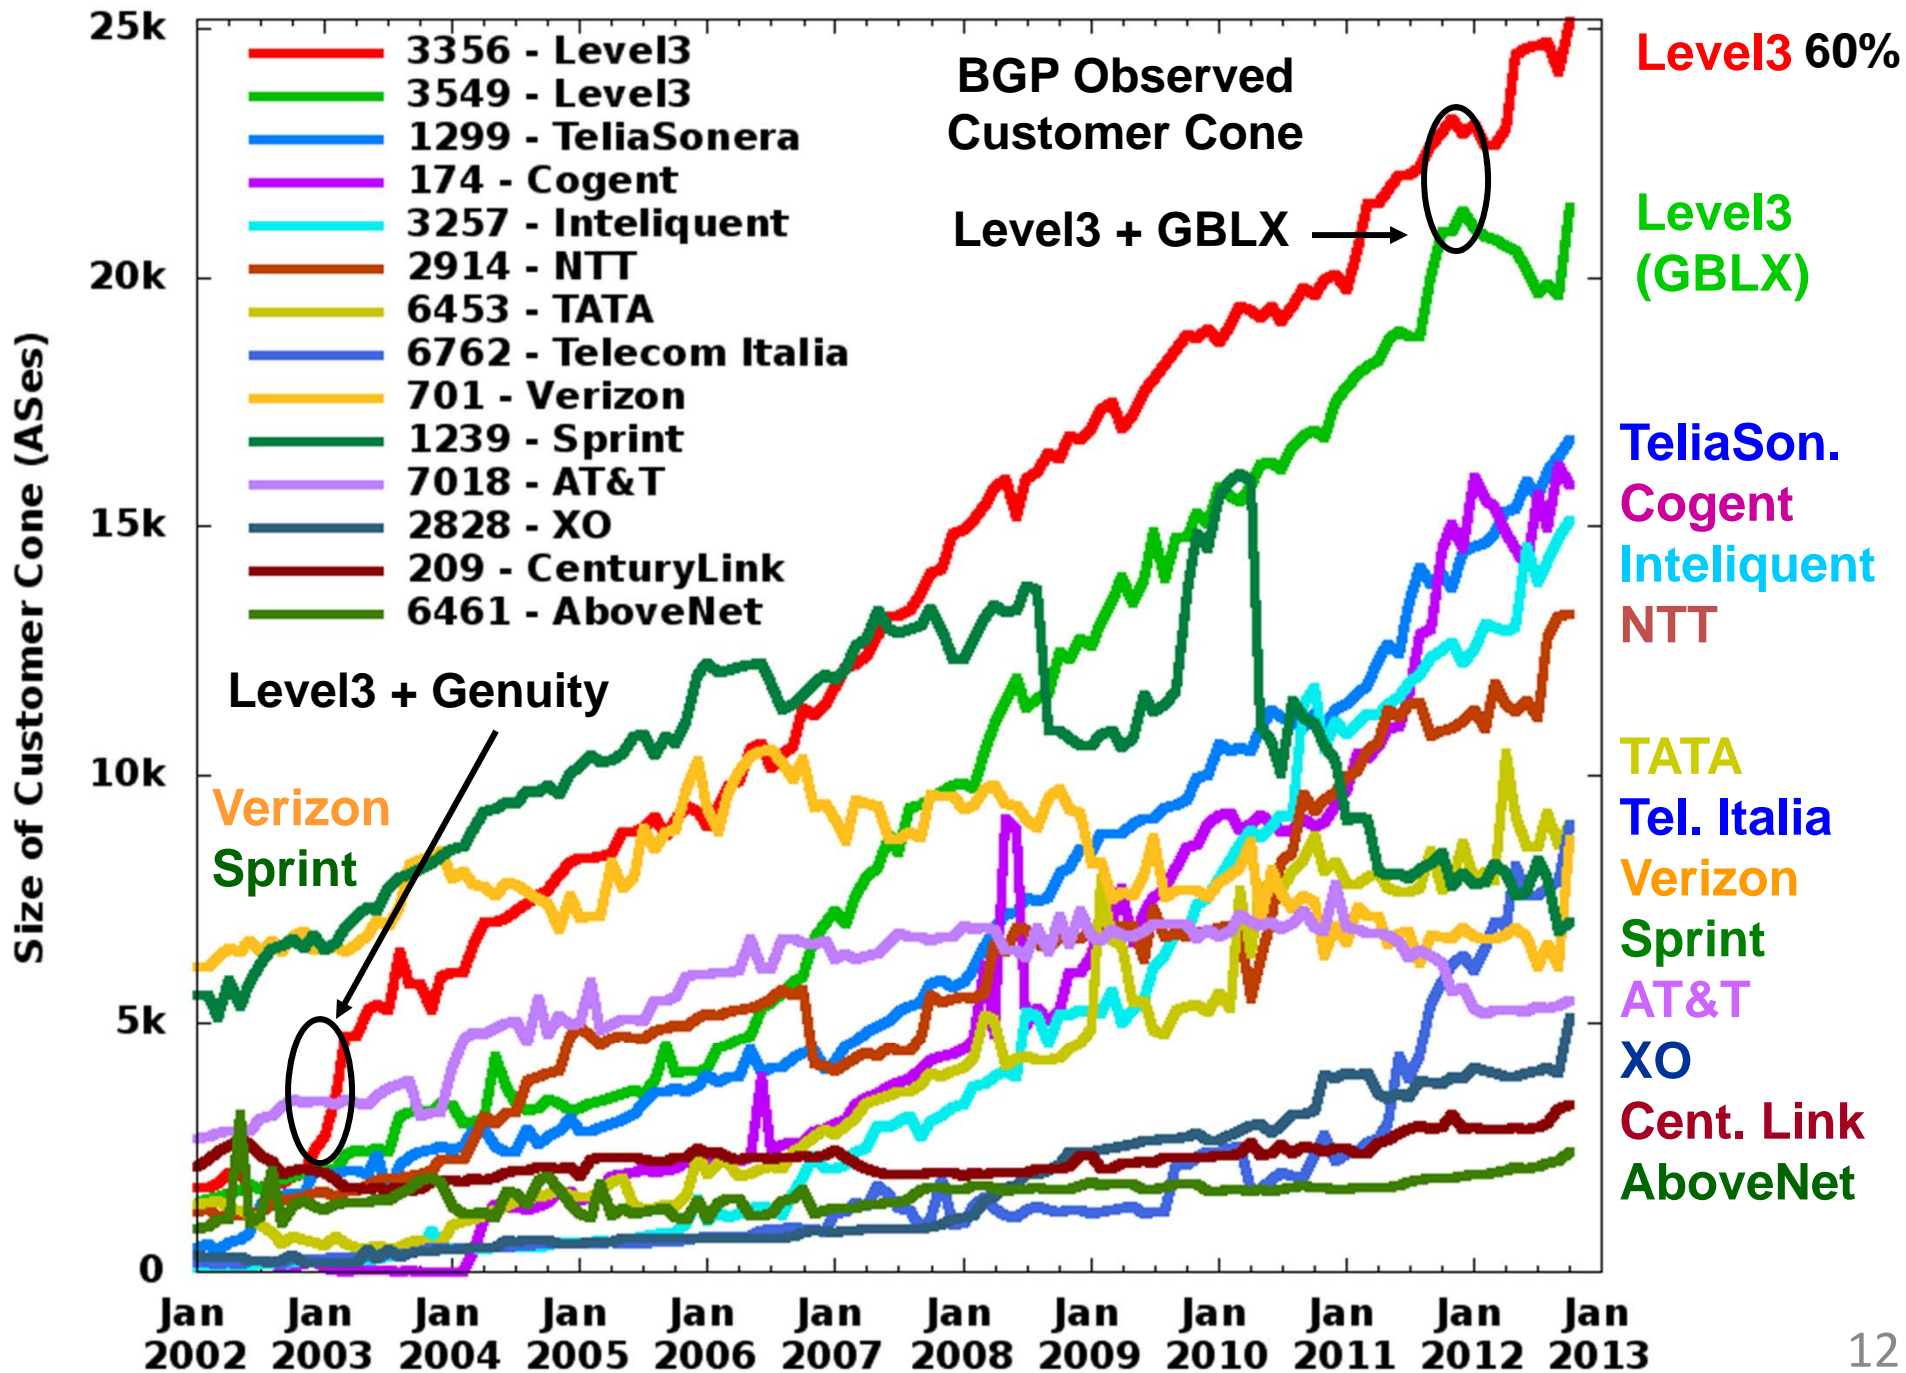


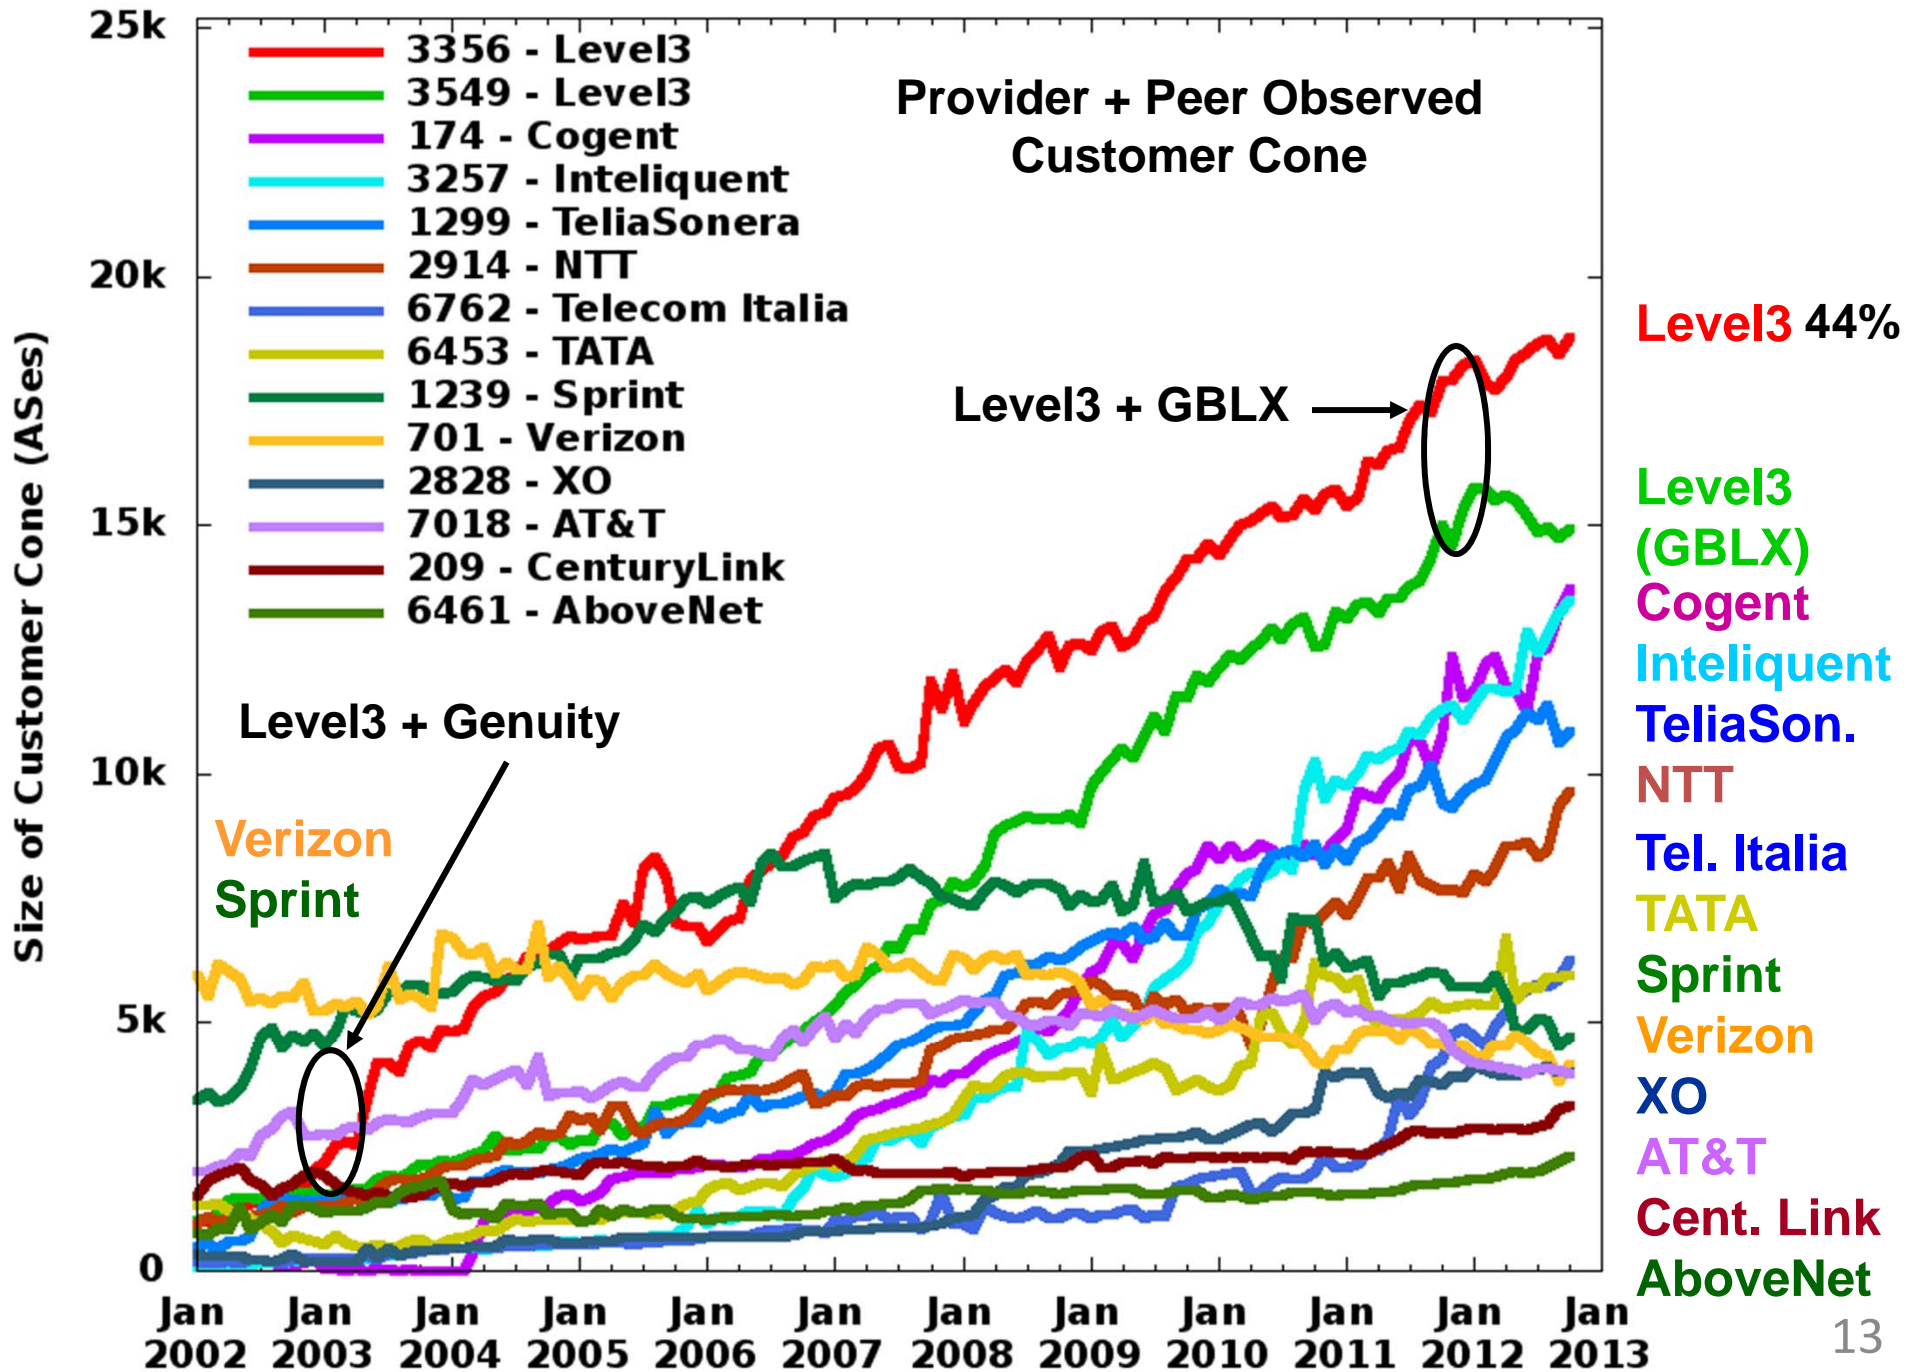
Region Y: "Europe"



Caveats

- AS Relationship ecosystem is complex
 - Different relationships in different **regions**
 - Can't differentiate between **paid-peers** and **settlement-free peers** (financial difference, not routing)
- Don't know about **traffic**
- Don't have much **visibility into peering**
- BGP paths are **messy** (poisoning, leaking)
- NOT a clear metric of market power





Existing AS-relationship algorithms
do not match ground truth very
well.

So we developed a new one.

Definition – Transit Degree

- **Transit Degree (TD)** : given set of BGP paths, the number of ASes for which an AS provides transit

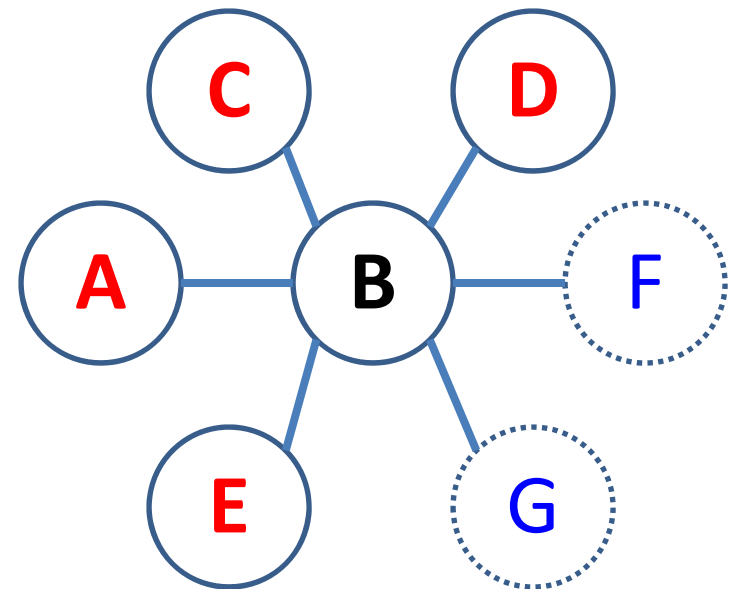
– A B C

– A B D

– A B E

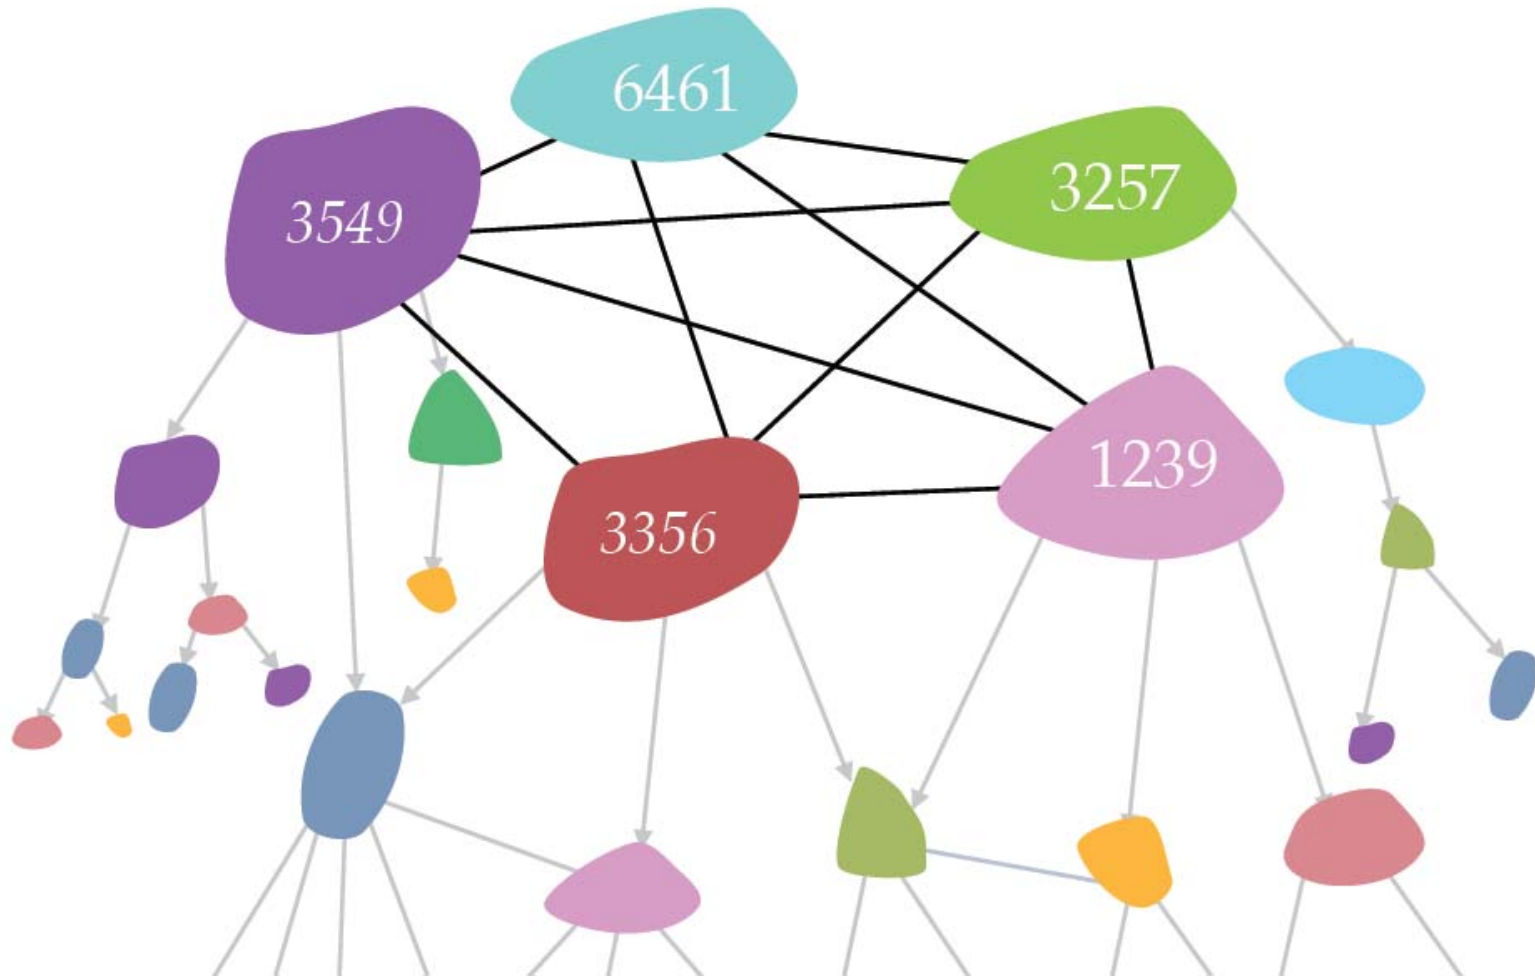
– B F

– B G



- Useful first-pass for ranking ASes:
 - B transits traffic for 4 ASes (A, C, D, E) : TD of 4.
 - B's node degree is $4 + 2 = 6$

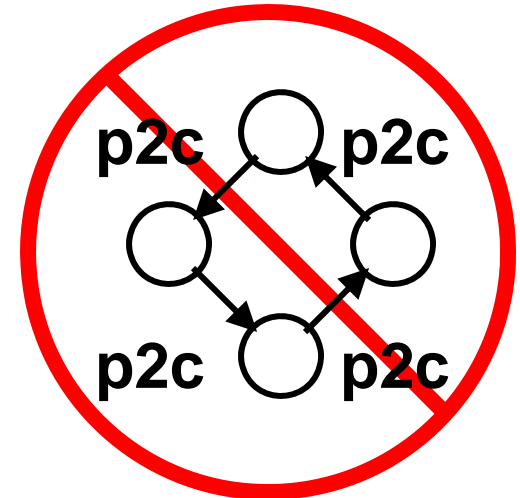
Assumption – Clique of ASes



Clique: Tier-1 ASes and their (possibly paid) peers in a full mesh. We cannot distinguish paid and settlement free peers.

Assumptions

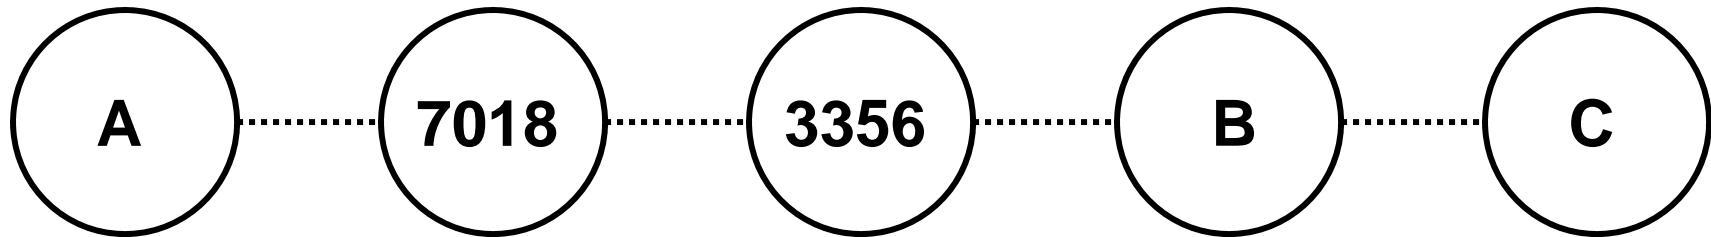
1. Full p2p mesh (clique) at top of hierarchy
 - **Otherwise Internet topology partitioned**
2. A provider announces customer routes to its providers
 - **Main point of paying a provider**
3. $TD_{\text{provider}} > TD_{\text{customer}}$ (mostly)
 - Supported by our **ground truth**
4. AS topology graph is acyclic: no p2c cycle
 - Supported by our **ground truth**



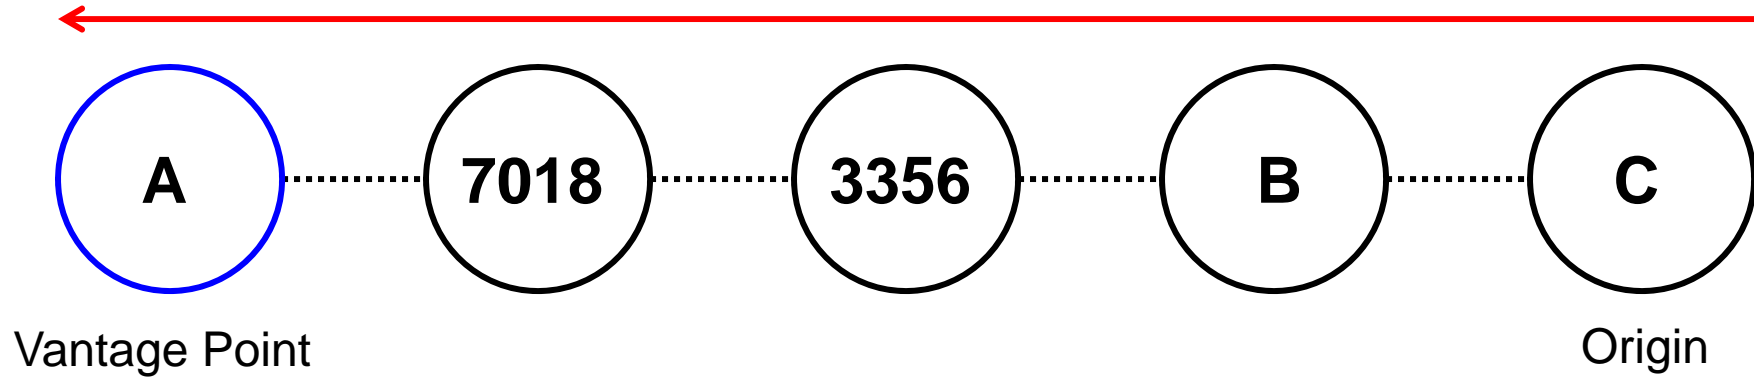
High-level Algorithm

- **Infer clique** and resulting p2p mesh
- Filter BGP paths (reserved ASes, **poisoning**)
- Break paths into **AS triplets**
- Visit ASes in order by largest **transit degree**
 - Infer c2p if
 - **neighbour passes route to a provider**, or
 - **neighbour is in clique** and passes route to another clique AS
 - 56.6% of graph / 99.8% PPV
 - Additional steps in algorithm (next slides)
- All other links in graph are p2p
 - 36.9% of graph / 98.5% PPV

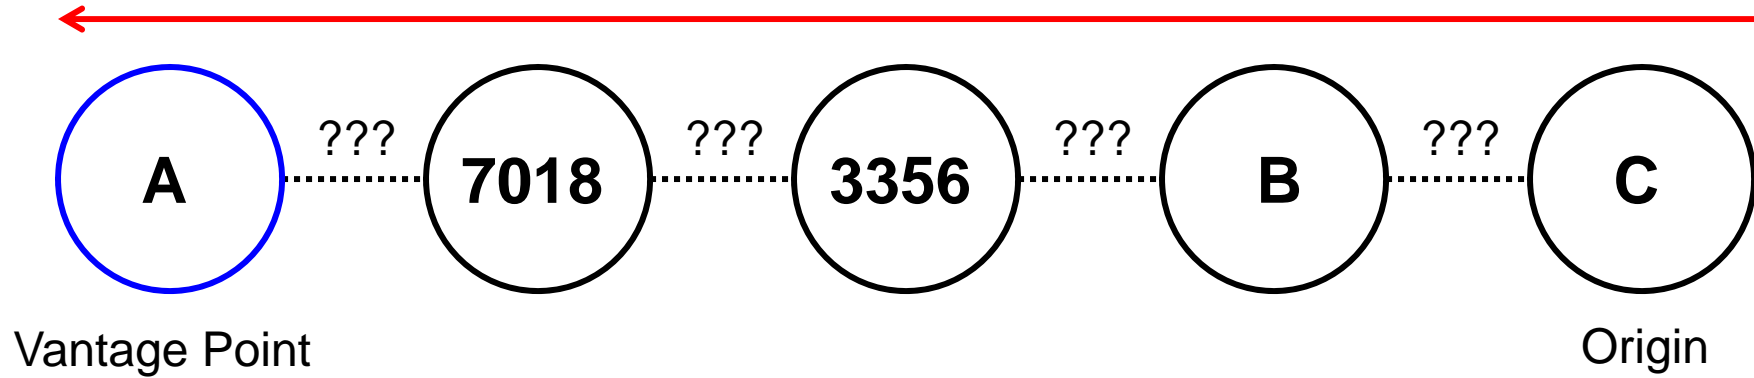
Inferring Relationships



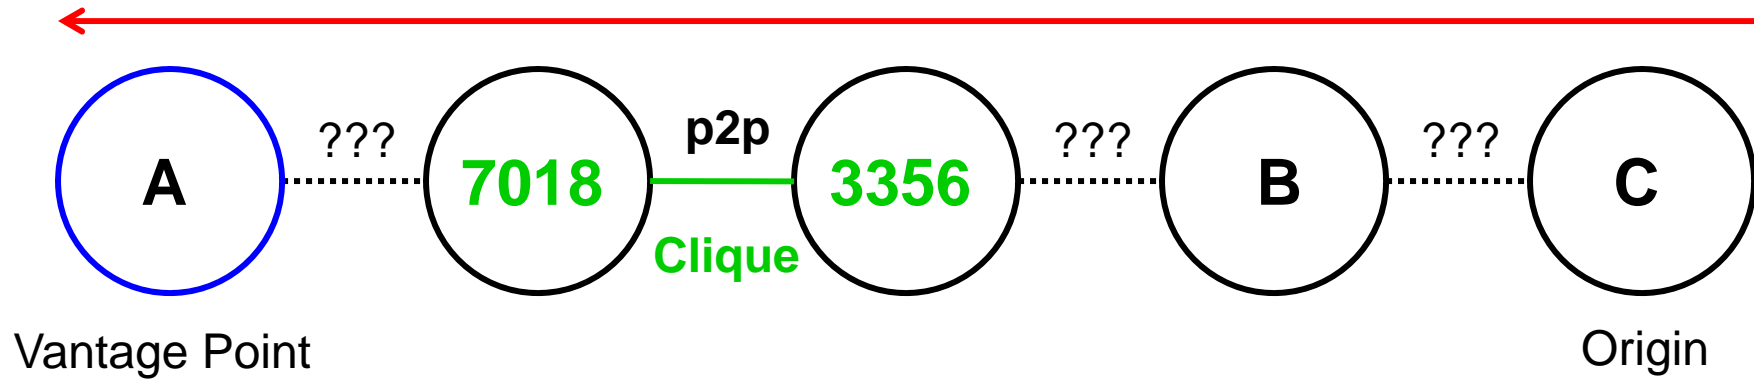
Inferring Relationships



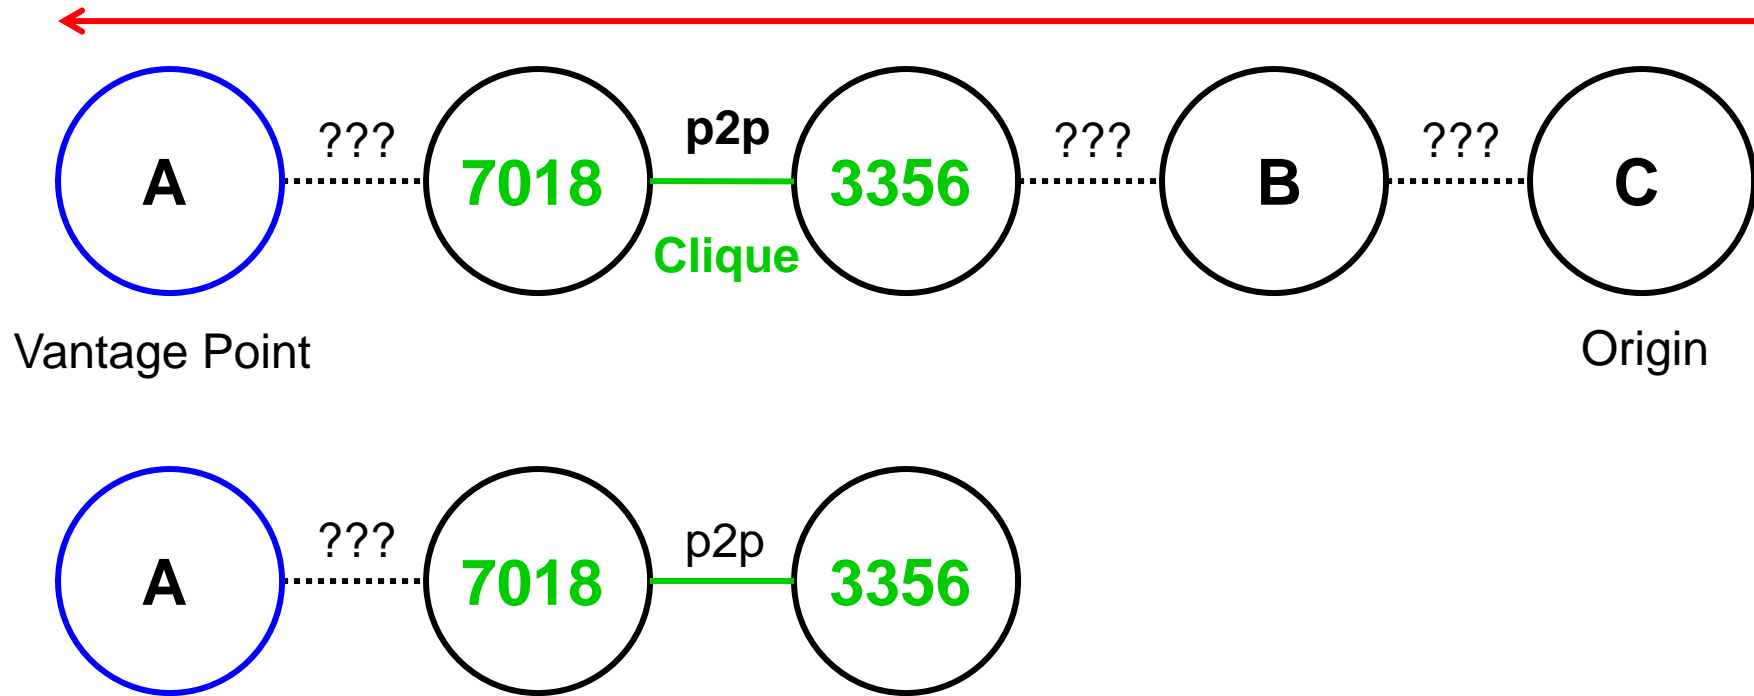
Inferring Relationships



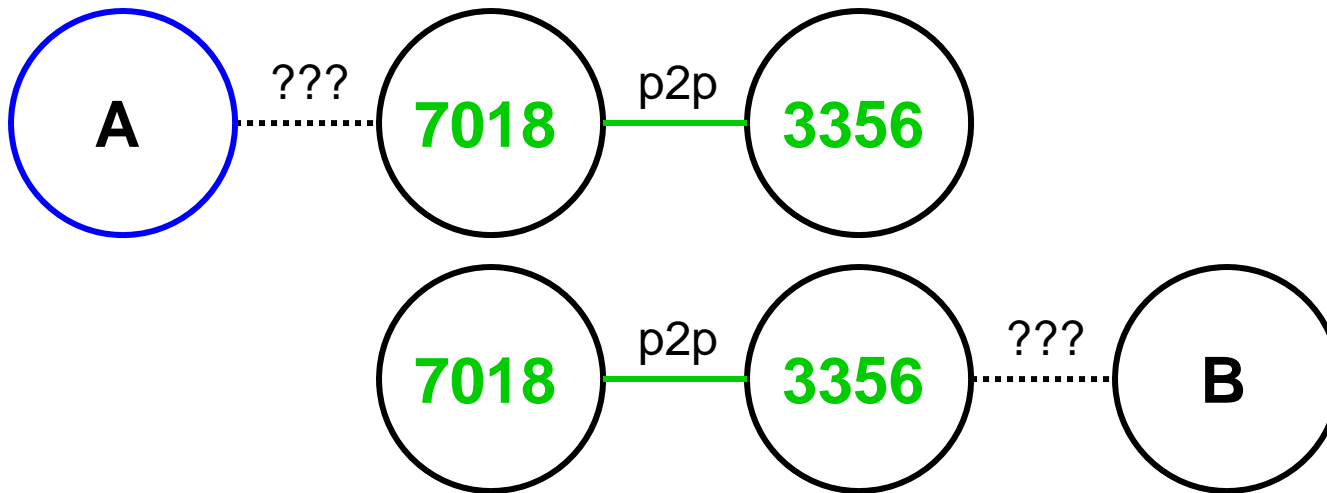
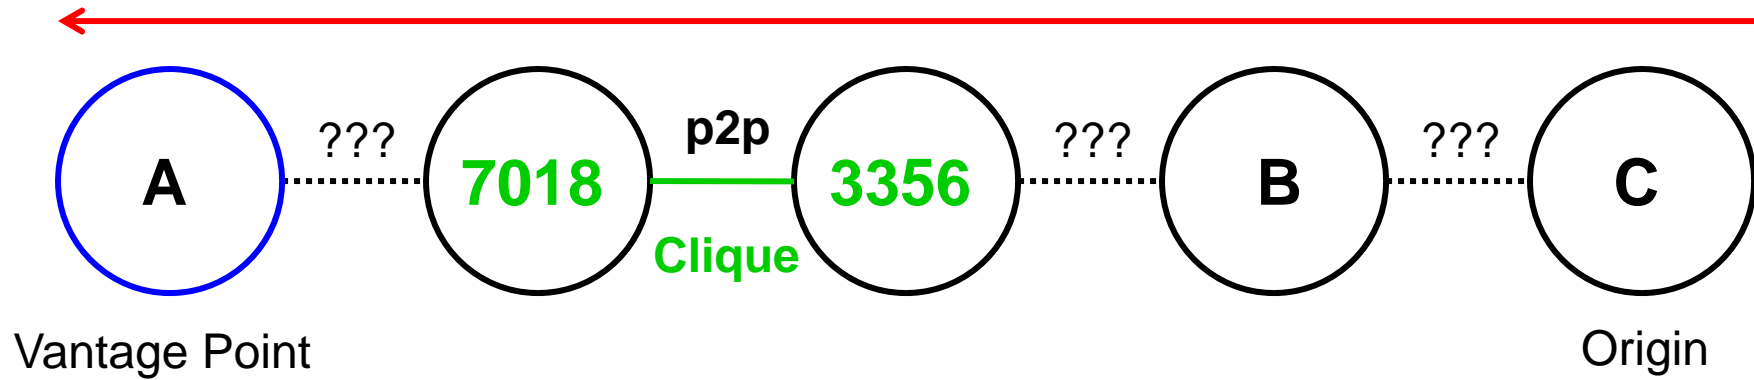
Inferring Relationships



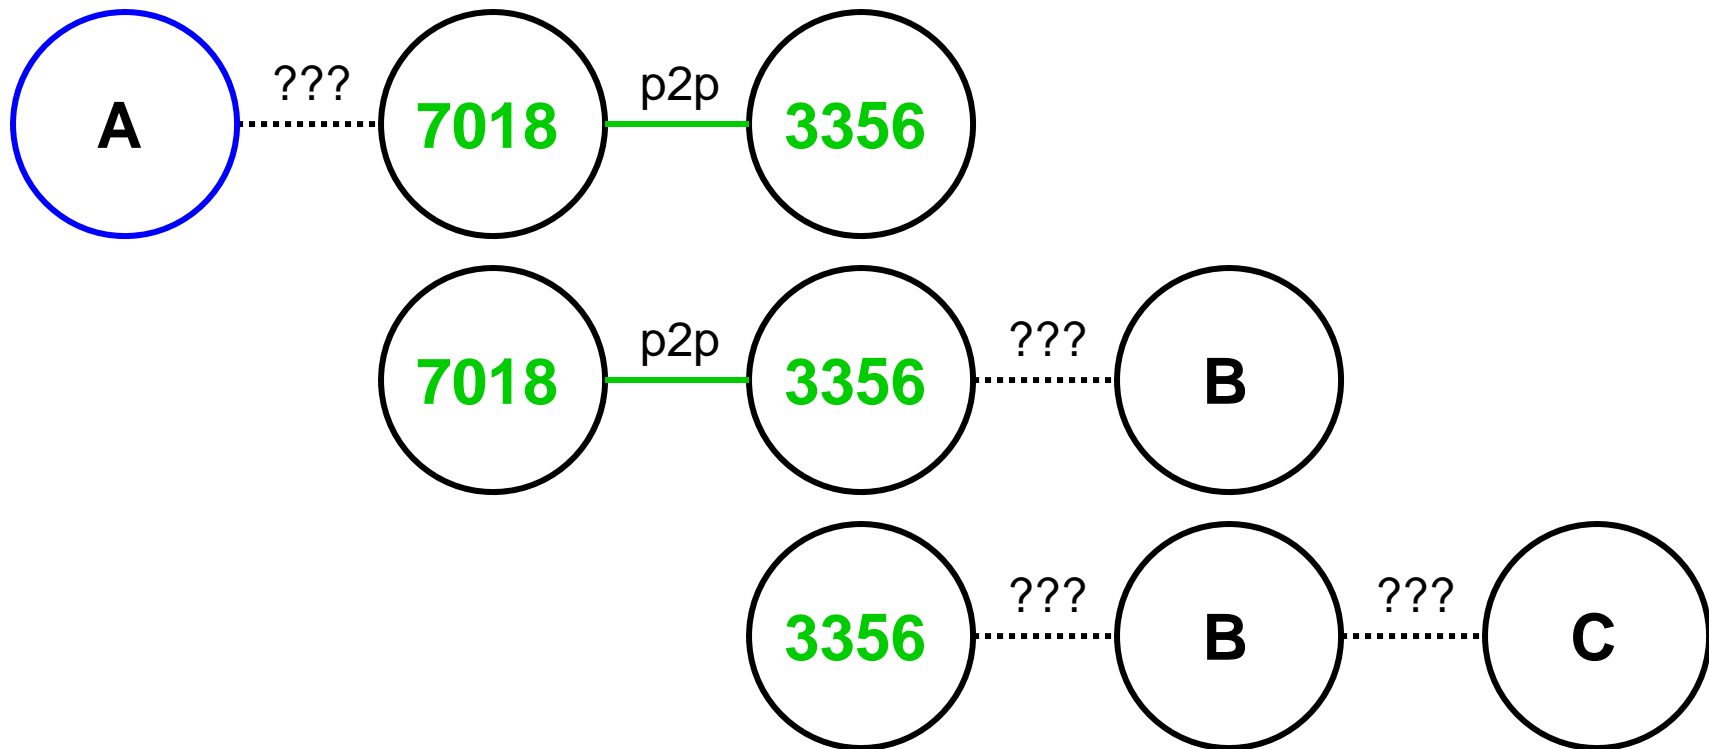
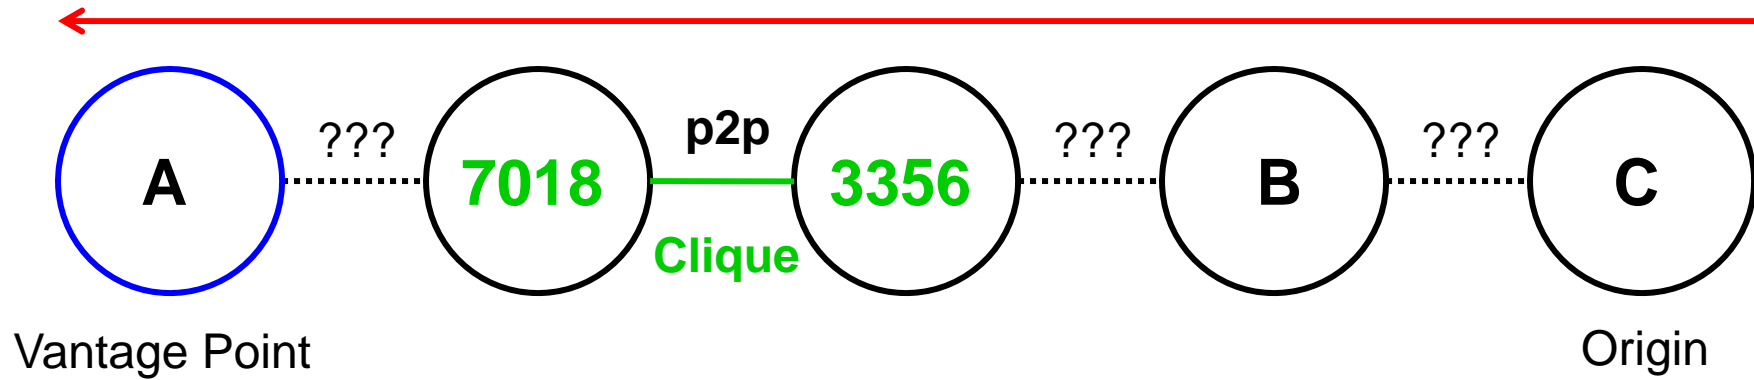
Inferring Relationships



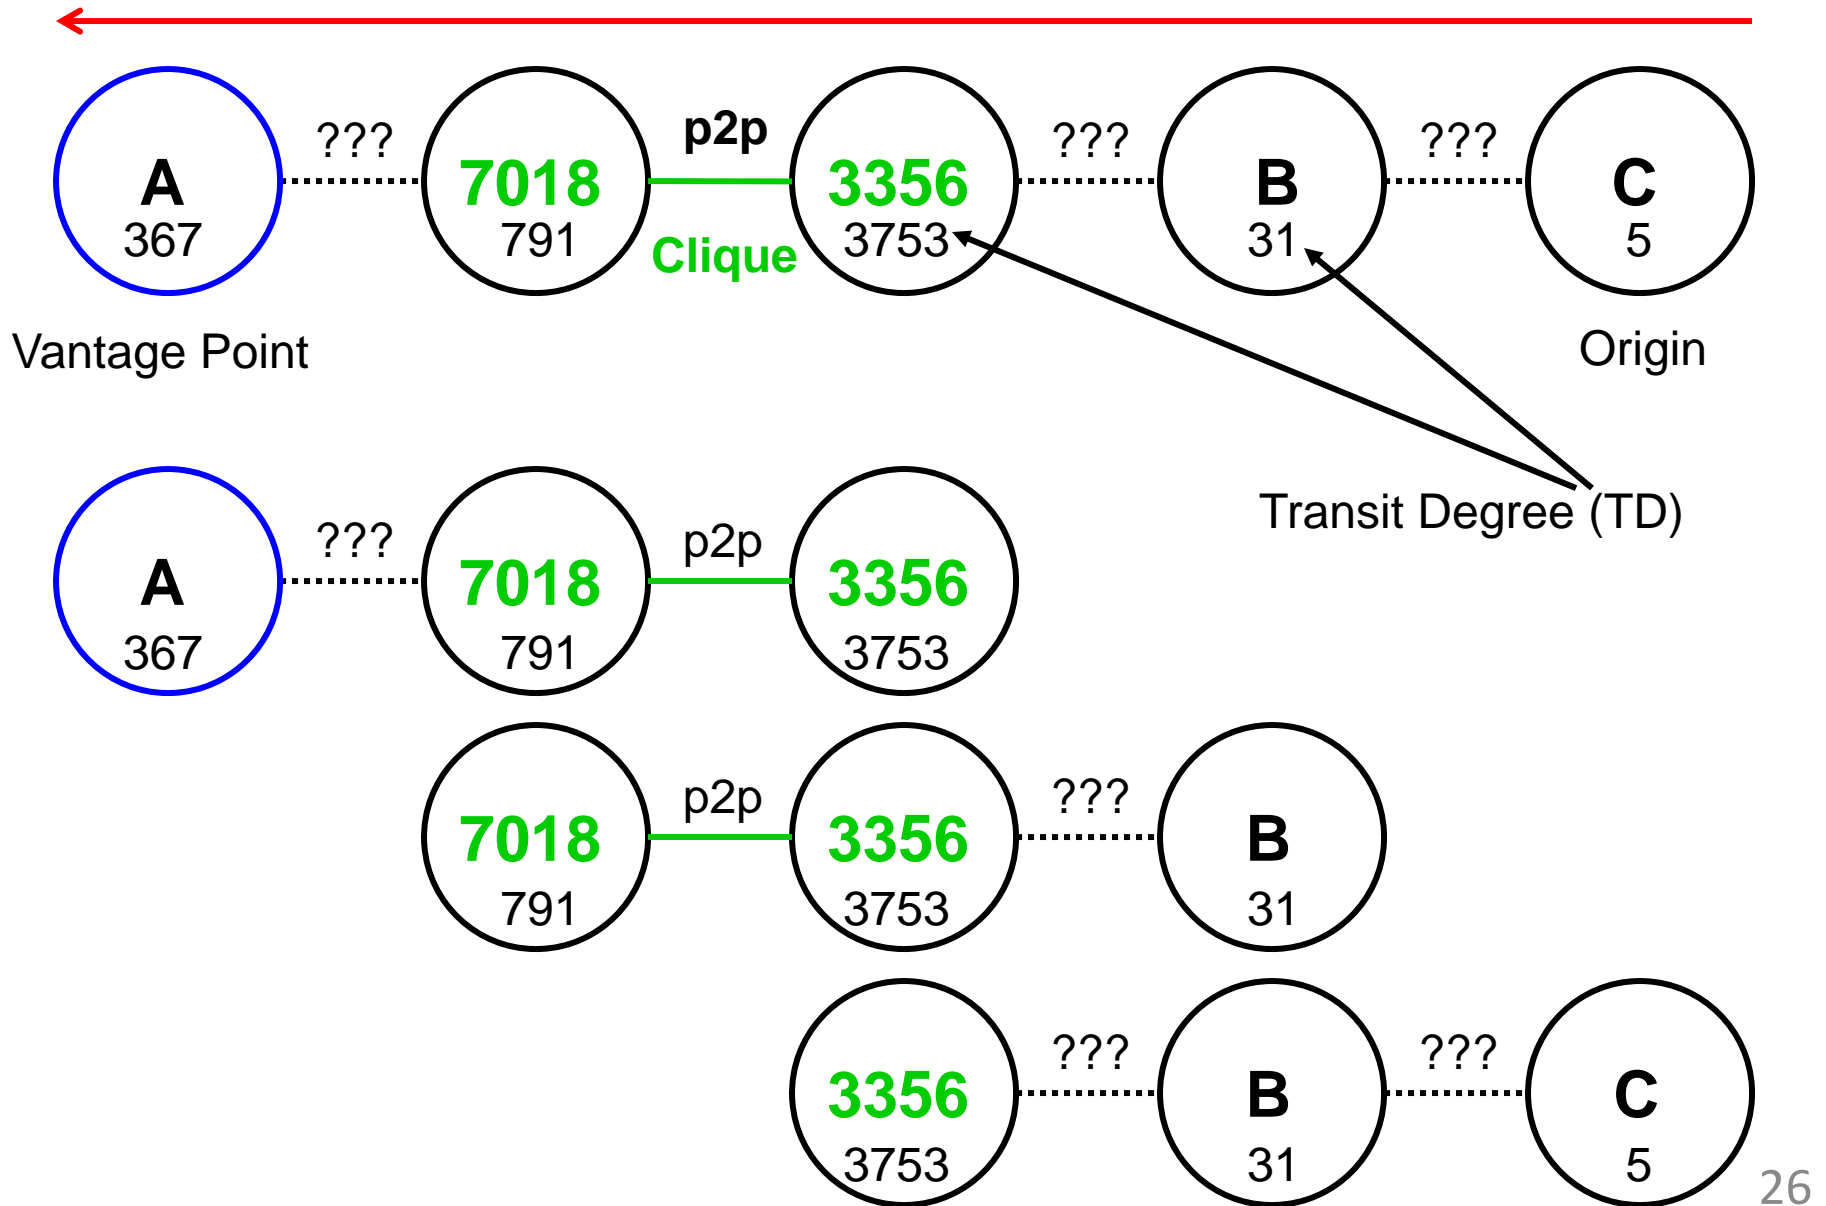
Inferring Relationships



Inferring Relationships

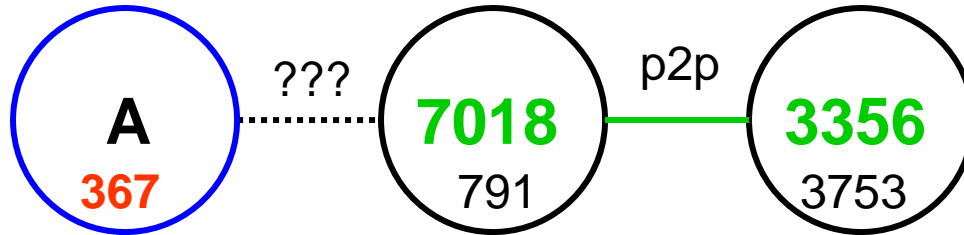


Inferring Relationships

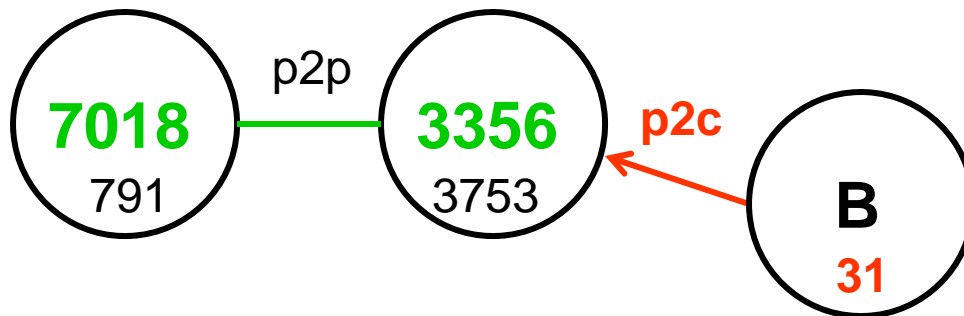


Inferring Relationships from Triplets

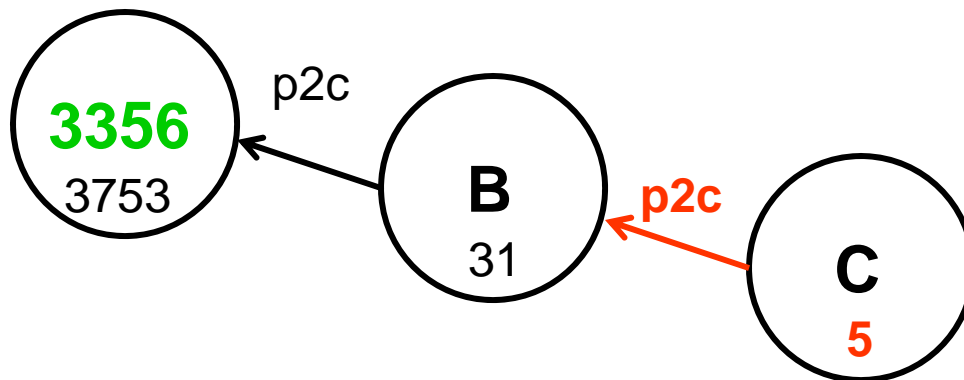
Visit ASes in order of highest transit degree



No inference made: **A** might be a peer and **7018** *might* be leaking. Need to see a path where provider is in front of its customer.



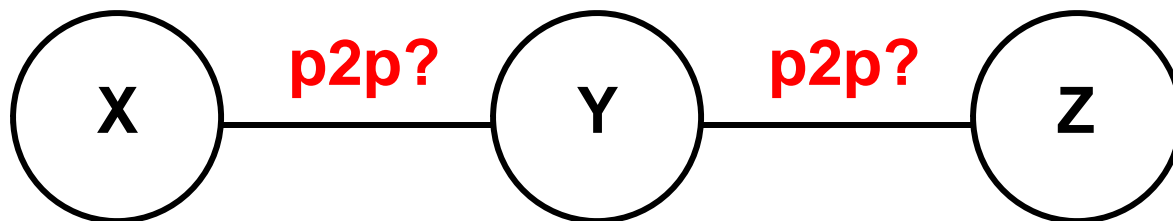
Infer **B** is a customer of **3356** because **3356** and **7018** are members of clique and **3356** advertises across clique.



Infer **C** is a customer of **B** because **B** advertises route to provider

Special Cases (6.5%)

1. VPs with no provider routes
2. Providers with smaller transit degree than customer
3. Customers for ASes with no providers (e.g. TransitRail)
4. Collapsing sequences of p2p links



Ground Truth Summary

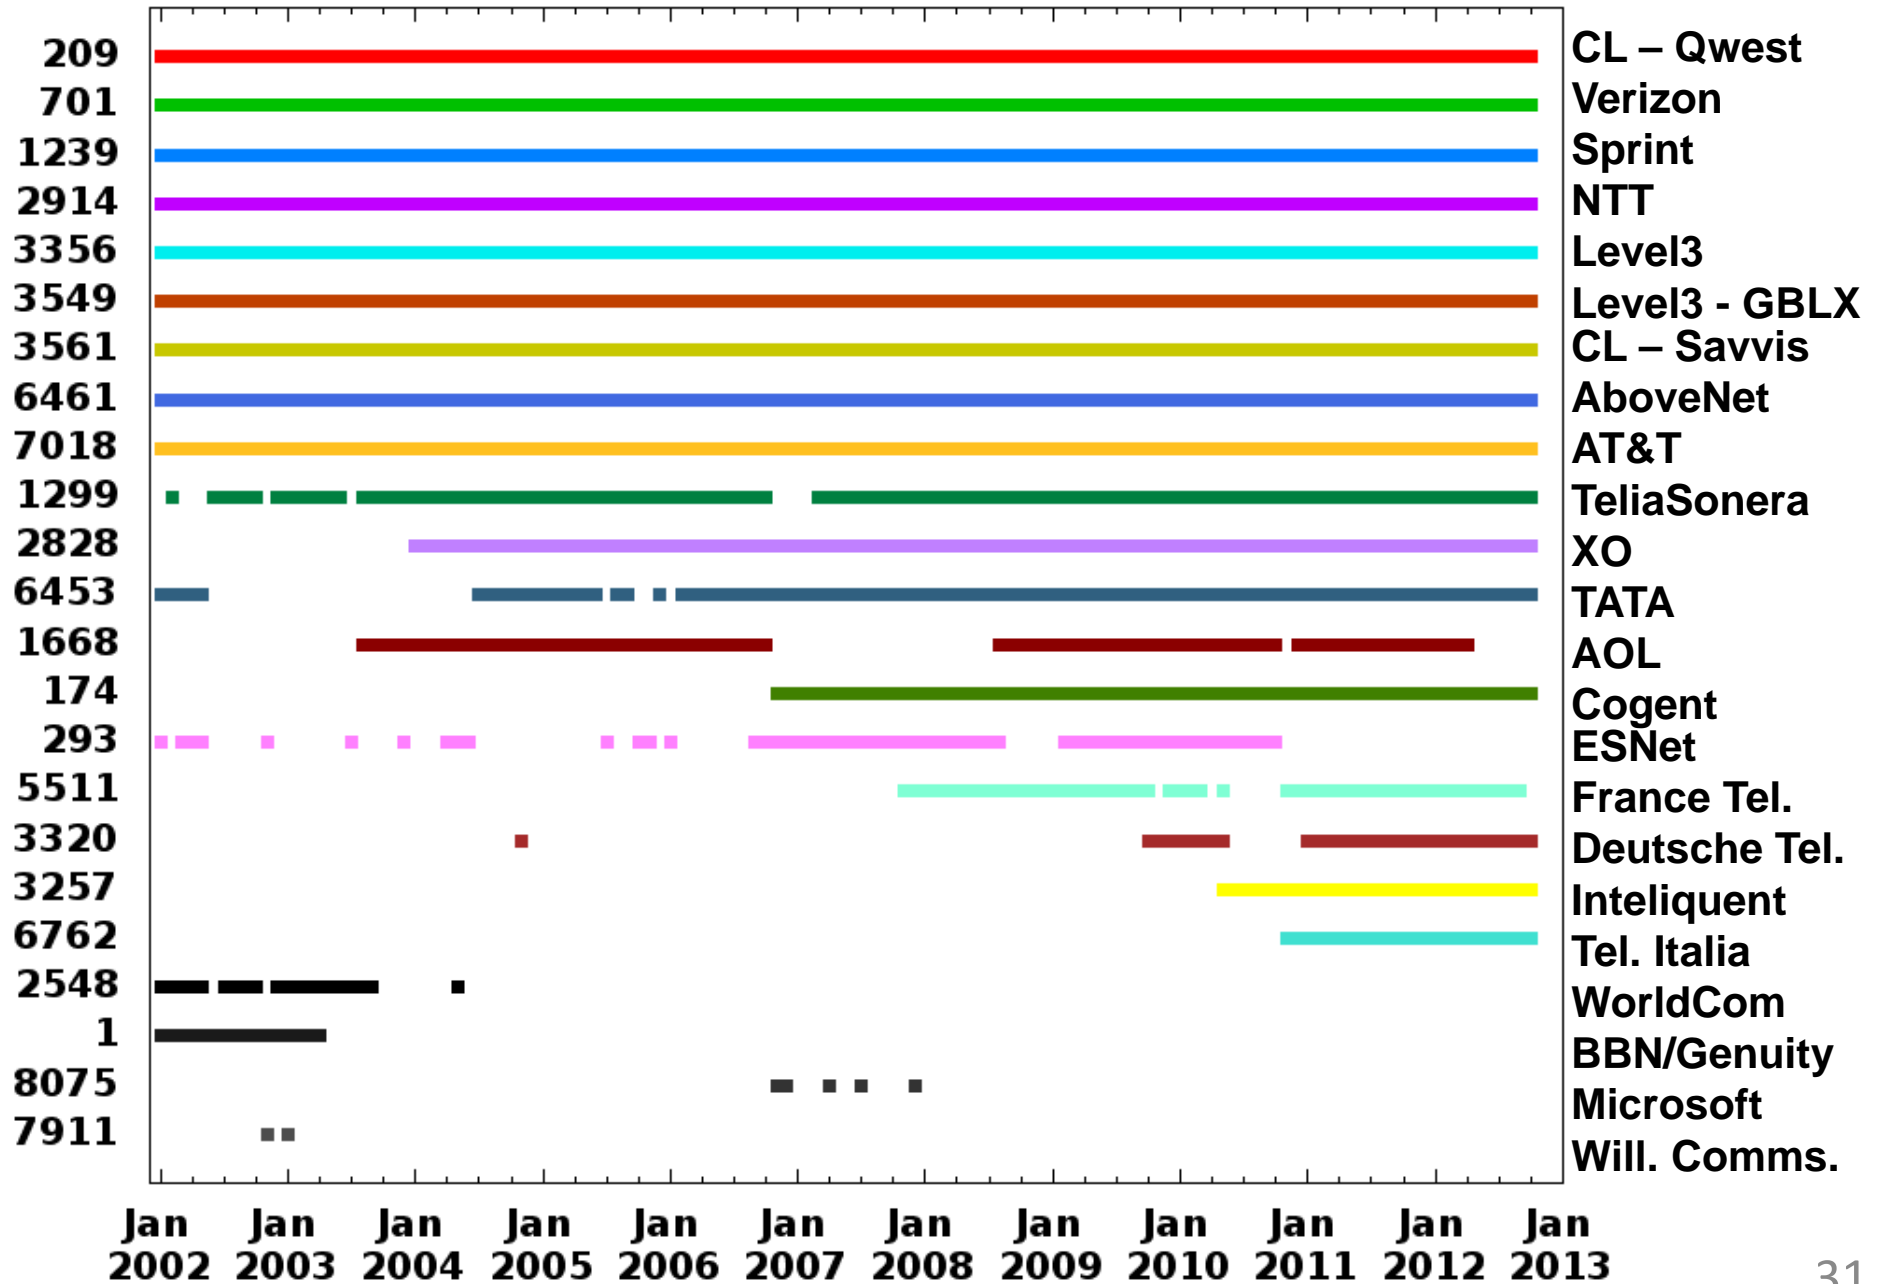
- **CAIDA:** **2,370**
 - 2010 – 2012 83% p2p
 - Most submitted via web form, some via email
- **RPSL:** **6,065**
 - April 2012 100% p2c
 - RIPE whois database, two-way handshake
- **BGP Communities:** **39,838**
 - April 2012 59% p2c
 - Dictionary of operator-published community meanings assembled by Vasileios Giotsas (UCL)
- Overall: 47,881 GT relationships, 63% p2c, 37% p2p
 - **~38% of the publicly available graph.**

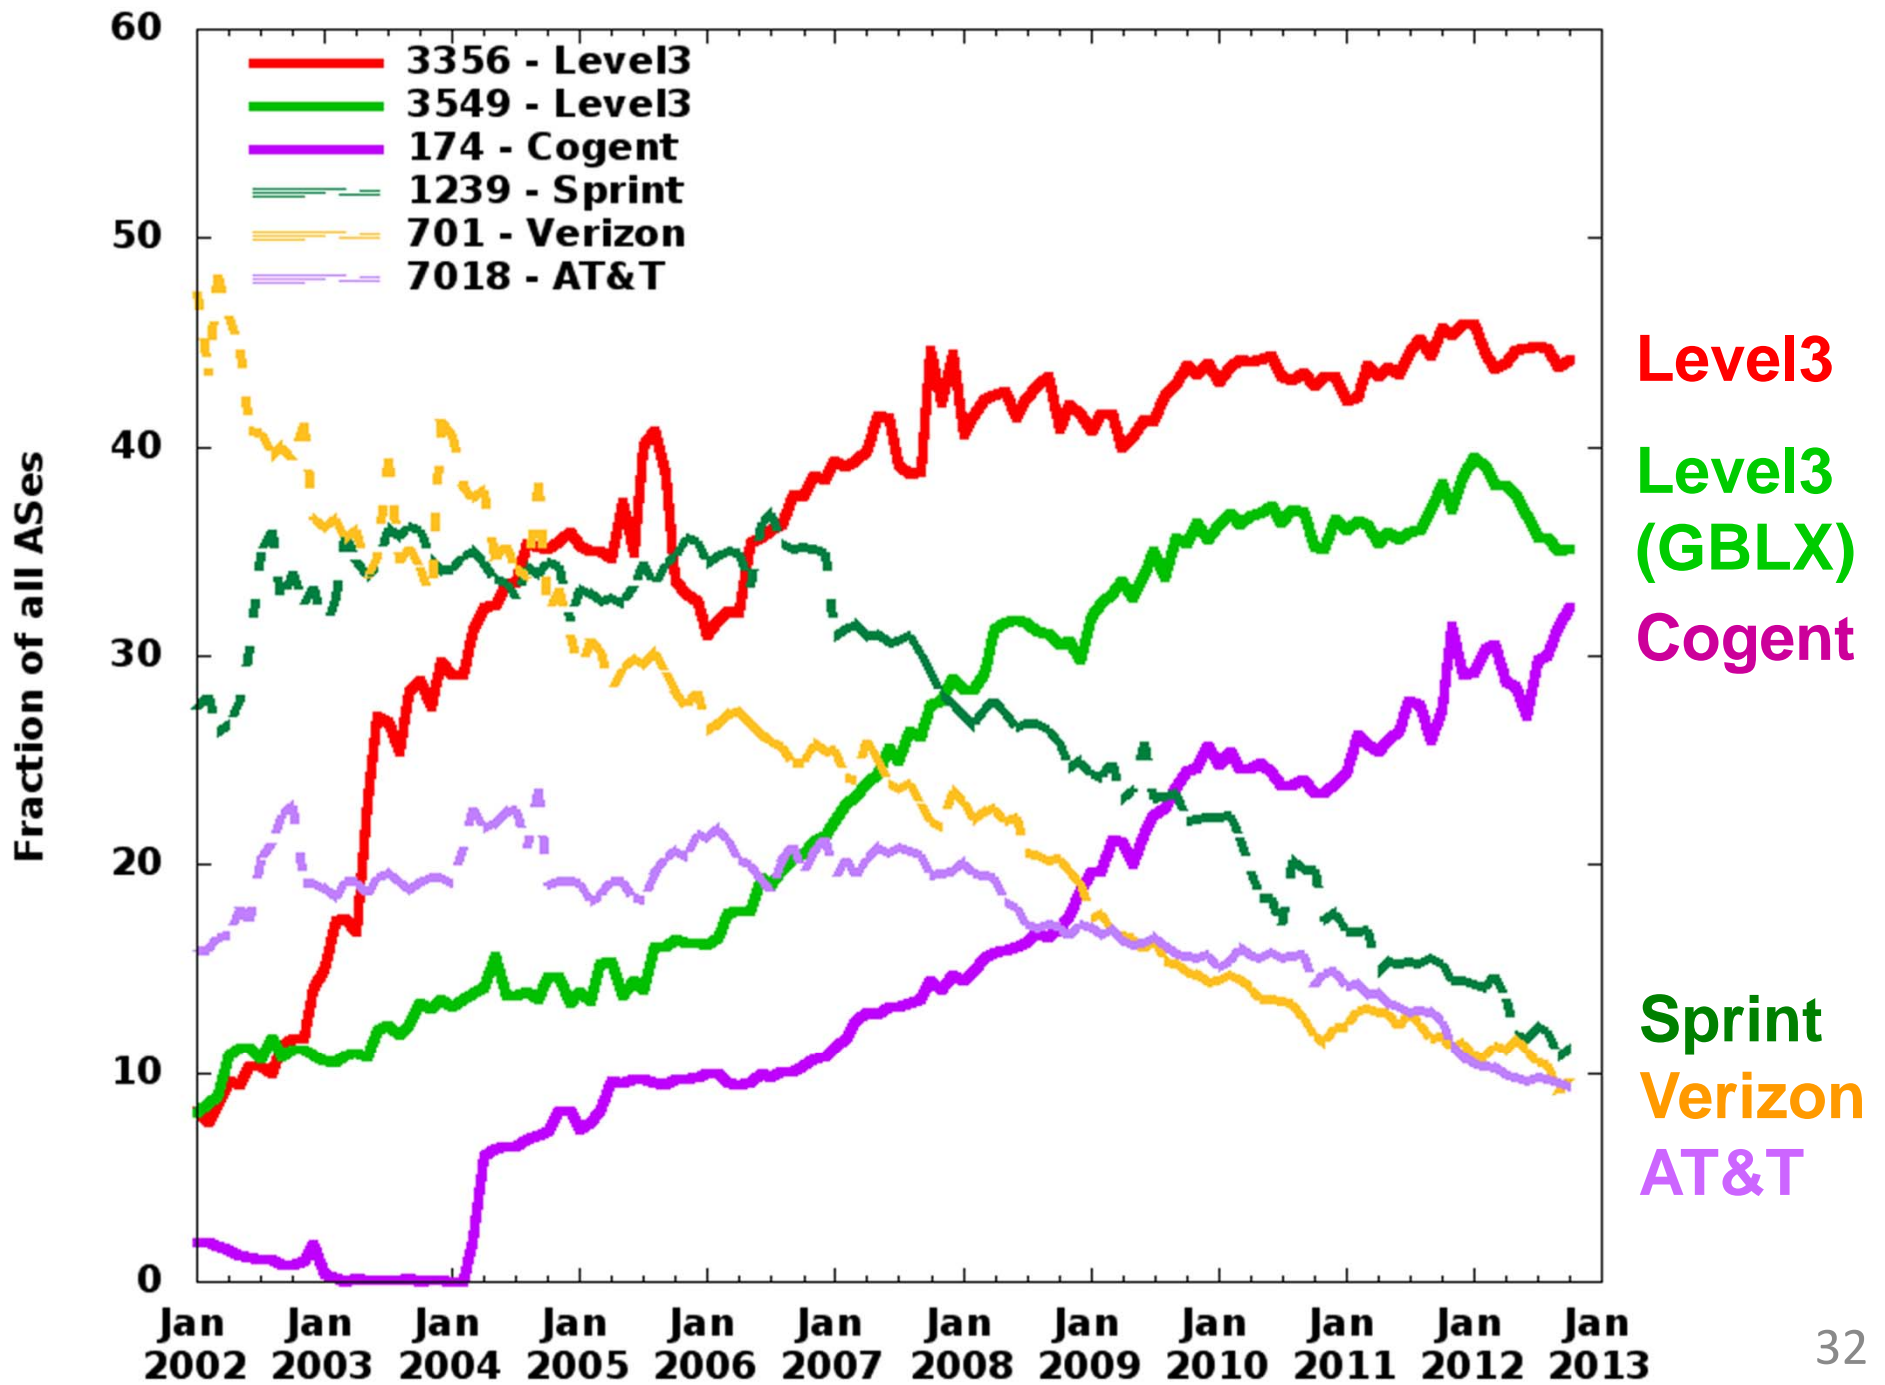
Validation (AS Relationships)

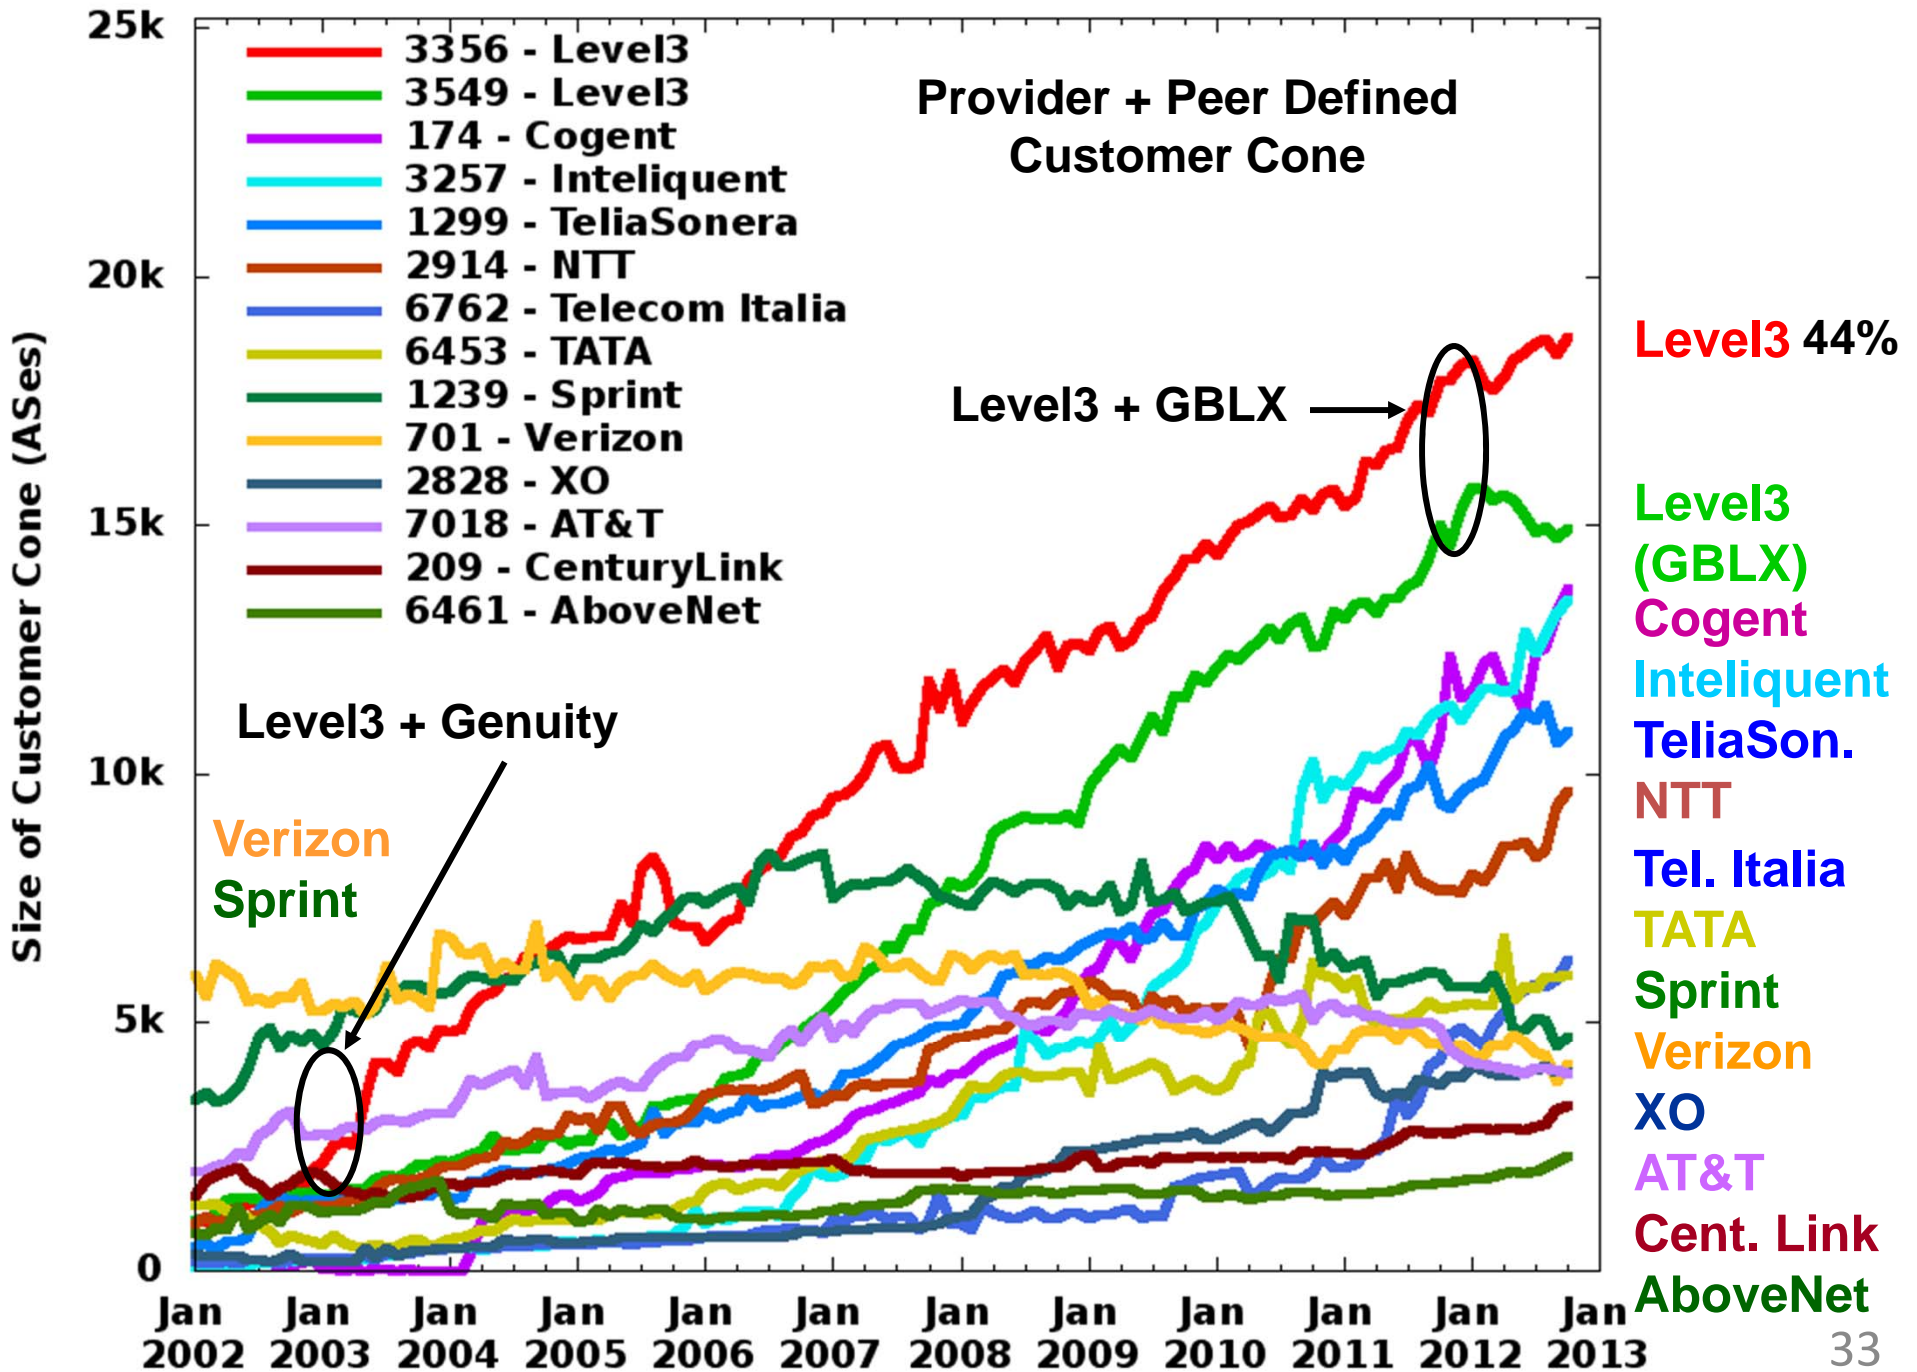
	p2c	p2p
CAIDA	99.6%	98.4%
UCLA	99.0%	90.9%
Gao	84.7%	99.5%

Percentages are Positive Predictive Value (PPV).
Take home: difficult to be accurate at
inferring both types of relationships

Clique members over time - IPv4







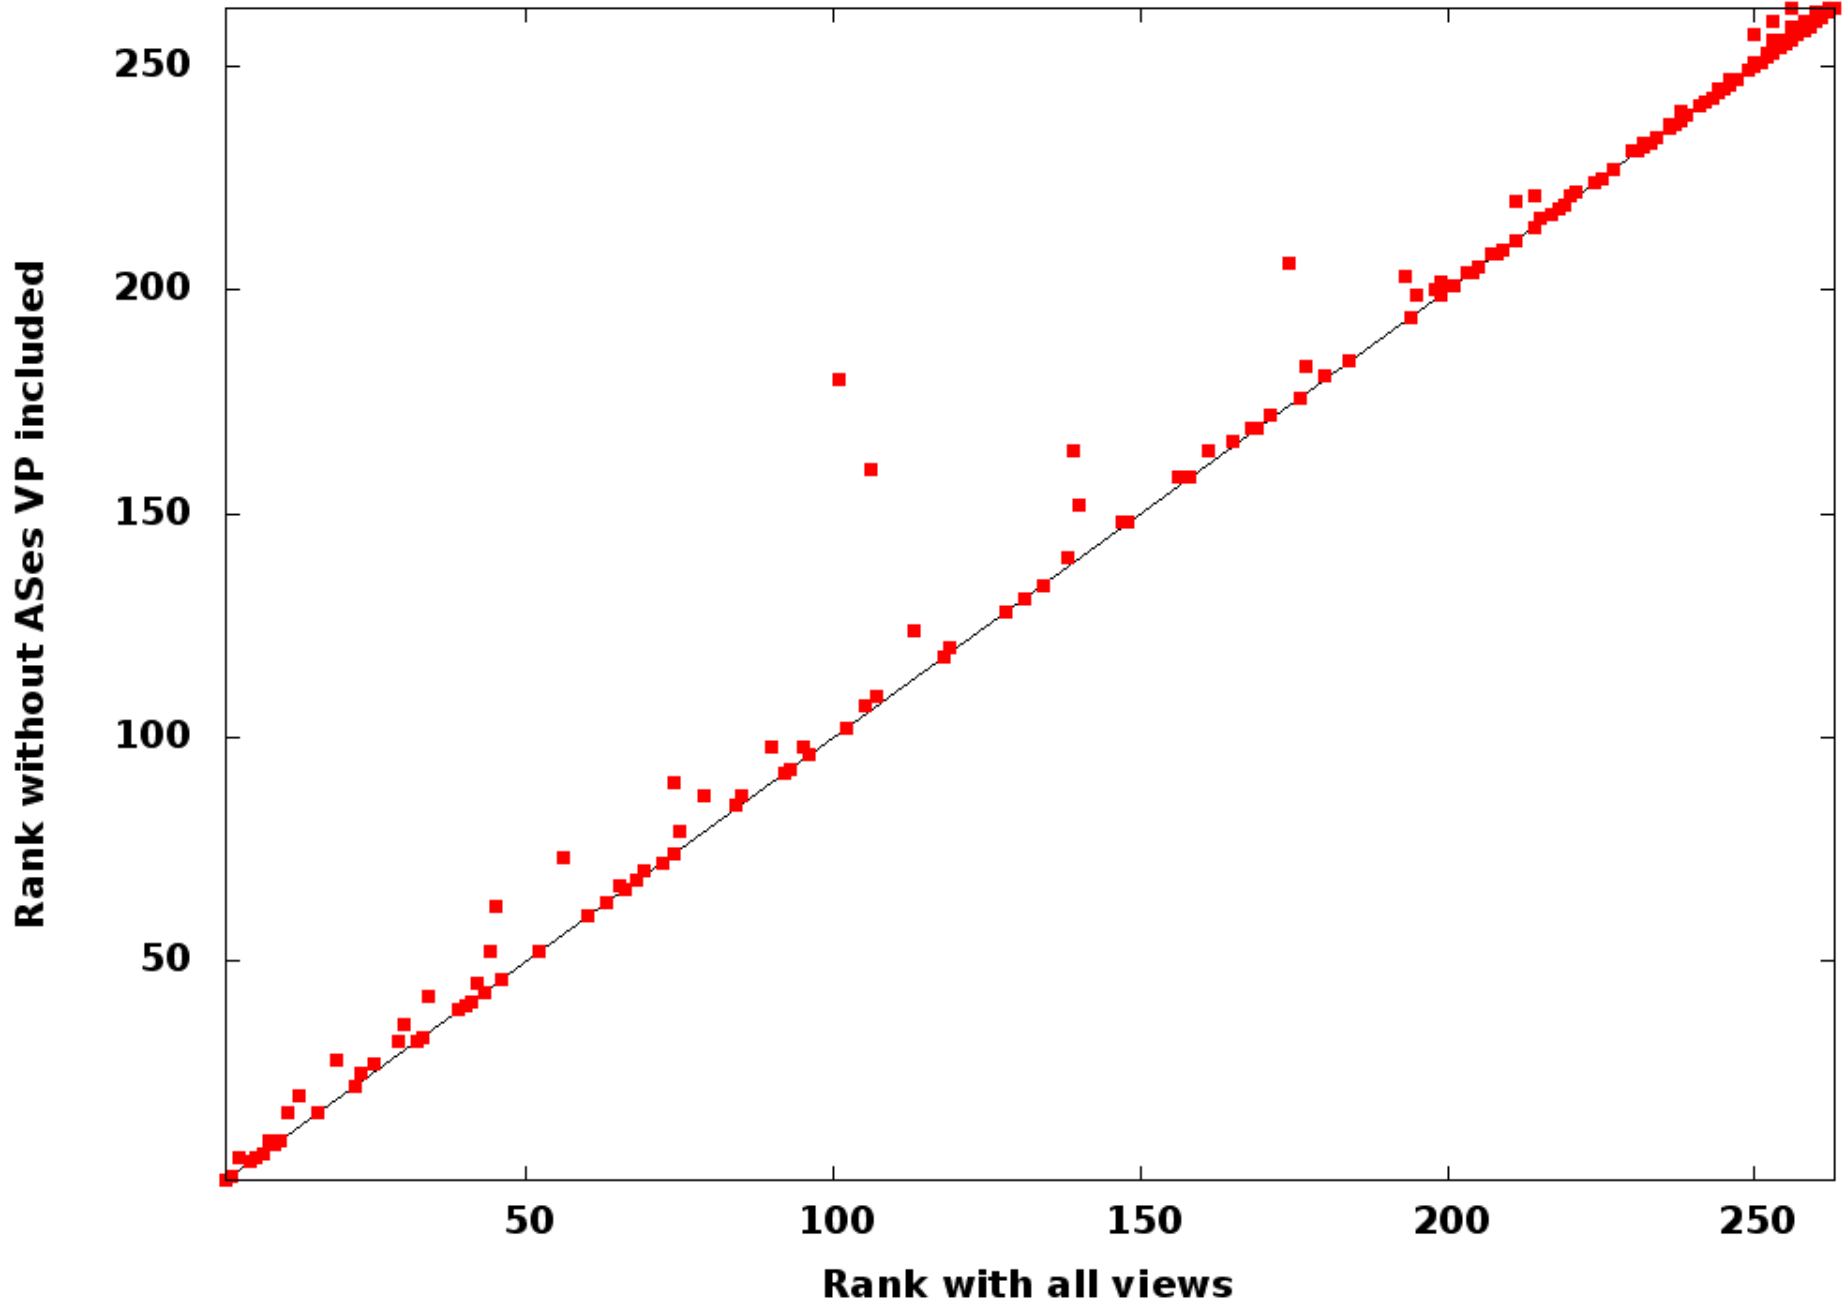
Help Wanted from NANOG

- What have we not thought of?
- How prevalent is **paid-peering**?
 - Is there any routing differences between a paid peer and a settlement-free peer?
 - Related: in what situations might a customer not be announced to its provider's providers?
- How prevalent are **complex relationships**?
 - What **granularity** is used for routing policies?
 - Region?
 - Prefix?
- More vantage points at Route Views and RIPE RIS
- Additional ground truth -- ideally submitted through <http://as-rank.caida.org/>



Backup Slides Follow

Rank without BGP Vantage Point



Algorithm Details – Preparing BGP paths for input

- Filter paths with
 - Reserved ASNs (23456, 61440-65534, etc) 0.1%
 - AS Loops (<0.1%)
 - If any two Clique ASes are separated by another
 - Assume leak or poisoning (0.01%)
 - More than 5 ASes in the path (14.7% of paths)
- Remove IXP ASNs from paths
 - Manually defined list of 25 ASes

The top ASes ranked by customer cone size are displayed below.

Dataset: 2012.06.01

For information about a specific AS, enter its AS name, its AS number, or the name of the Org of which the AS is a member.

Look up an AS by number or name

Table shows of 41203 ASes, sorted by

AS rank	AS number	AS name	Org name	customer cone						AS degree
				Number of			Percentages of all			
				ASes	IPv4 Prefixes	IPv4 Addresses	ASes	IPv4 Prefixes	IPv4 Addresses	
1	3356	LEVEL3	Level 3 Communications	24,632	273,357	1,557,093,299				3331
2	3549	LVLT-3549	Level 3 Communications	20,747	238,941	1,029,028,043				1519
3	1299	TELIANET	TeliaNet Global Network	16,021	174,713	749,588,007				683
4	3257	TINET-BACK...	Tinet SpA	14,594	179,792	794,176,517				884
5	174	COGENT-174	Cogent/PSI	14,058	164,072	667,755,471				3537
6	2914	NTT-COMMUN...	NTT America, Inc.	11,690	160,779	757,127,906				814
7	6453	AS6453	TATA Communications	8,265	126,544	509,082,236				573
8	1239	SPRINTLINK	Sprint	7,791	126,000	958,565,637				972
9	6762	SEABONE-NET	TELECOM.ITALIA SPARKLE S.p.A.	7,567	105,985	379,857,433				263
10	701	UUNET	MCI Communications Services, Inc. d/b/a Verizon Business	5,468	83,096	607,375,447				1812

data sources

geolocation	database	2012.06.25	netacuity
organization	whois	0000.00.00 2012.06.29	JPNIC, KRNIC, LACNIC AFRINIC, APNIC, ARIN, LACNIC, RIPE
topology	BGP	2012.06.01, 2012.06.02, 2012.06.03, 2012.06.04, 2012.06.05	ripe rrc00, rrc01, rrc03, rrc04, rrc05, rrc06, rrc07, rrc10, rrc11, rrc12, rrc13, rrc14, rrc15 routeviews eqix, isc, linx, routeviews2, saoppaulo, sydney

The top Organizations ranked by customer cone size are displayed below.

Dataset: 2012.06.01

For information about a specific Org, enter its name:

Look up an Org by name Search

Table shows 10 of 43174 Orgs, sorted by number of ASes in customer cone update view

Org rank	Org name	Num. ASes	customer cone						AS degree	Org degree
			Number of			Percentages of all				
			ASes	IPv4 prefixes	IPv4 addresses	ASes	IPv4 Prefixes	IPv4 Addresses		
1	Level 3 Communications	18	31,141	330,943	1,770,445,516	75%	79%	69%	4,375	3,872
2	TeliaNet Global Network	5	16,073	175,312	759,252,039	39%	41%	29%	742	662
3	Tinet SpA	2	14,595	179,793	794,176,773	35%	43%	31%	884	798
4	Cogent/PSI	3	14,060	164,075	667,755,727	34%	39%	26%	3,535	3,267
5	NTT America, Inc.	7	11,697	160,793	757,129,442	28%	38%	29%	809	703
6	TATA Communications	6	8,528	131,962	518,890,651	20%	31%	20%	842	777
7	Sprint	16	7,844	129,977	965,836,541	19%	31%	37%	1,023	906
8	TELECOM ITALIA SPARKLE S.p.A.	3	7,569	105,987	379,857,945	18%	25%	14%	262	236
9	Qwest Communications Company, LLC	11	7,253	122,871	759,408,992	17%	29%	29%	1,754	1,596
10	Init7 Global Backbone	424	6,744	72,220	395,631,197	16%	17%	15%	2,564	2,305

data sources

geolocation	database	2012.06.25	netacuity
organization	whois	0000.00.00	JPNIC, KRNIC, LACNIC
		2012.06.29	AFRINIC, APNIC, ARIN, LACNIC, RIPE
topology	BGP	2012.06.01, 2012.06.02, 2012.06.03, 2012.06.04, 2012.06.05	ripe rrc00, rrc01, rrc03, rrc04, rrc05, rrc06, rrc07, rrc10, rrc11, rrc12, rrc13, rrc14, rrc15
			routeviews eqix, isc, linx, routeviews2, saoppaulo, sydney

Support for this work is provided by the U.S. Department of Homeland Security's [Science and Technology Directorate \(Project N66001-08-C-2029\)](#), the National Science Foundation [Internet Laboratory for Empirical Network Science \(Project CNS-0958547\)](#), and Cisco's [University Research Program](#).

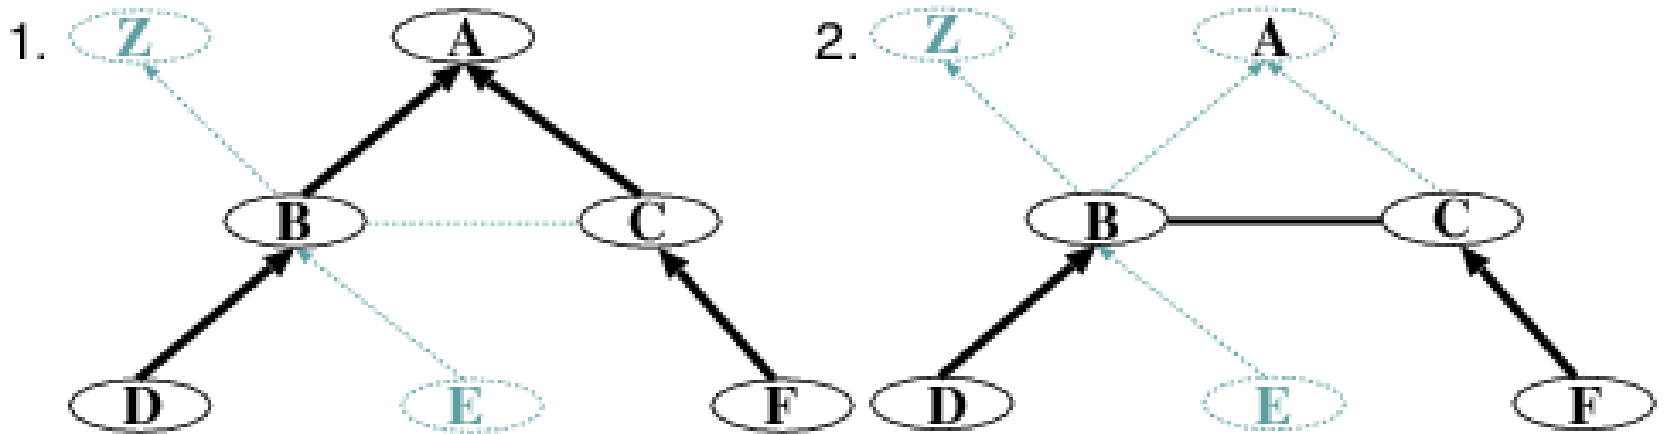
Provide corrections to AS relationships

If you know that some of our **inferred relationship type** for neighbor ASes to be inaccurate, please set its **actual relationship type** with the pull-down menu indicate what you believe to be the correct peering relationship between that neighbor and 3356 (e.g., <neighbor AS> is a <customer> of AS 3356). If you can verify that our inferred peering **relationship type** is correct, please indicate them as "(correct)".

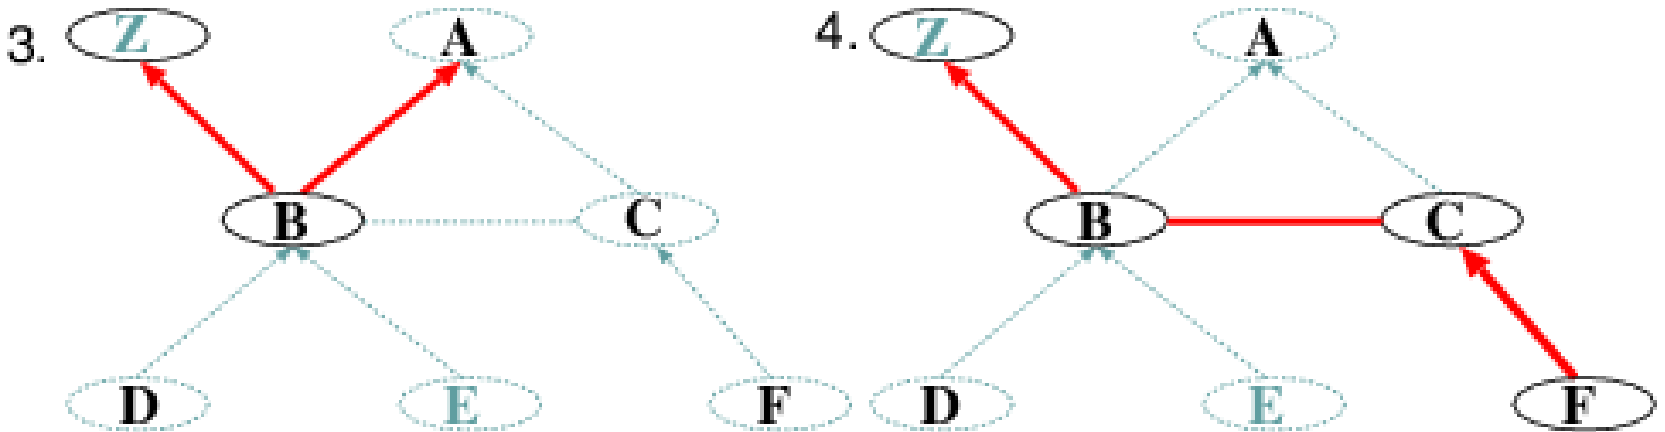
neighbor				inferred relationship type	actual relationship type
AS rank	AS	AS name	Org name		
2	3549	LVLT-3549	Level 3 Communications, Inc.	↔ sibling	<input type="text"/>
3	1299	TELIANET	TeliaNet Global Network	↔ peer	(correct)
4	174	COGENT-174	Cogent Communications	↔ peer	↓ customer
5	3257	TINET-BACKBONE	Tinet Spa	↔ peer	↑ provider
6	2914	NTT-COMMUNICATIONS-2914	NTT America, Inc.	↔ peer	↔ peer
7	701	UUNET	MCI WorldCom	↔ peer	↔ sibling
8	1239	SPRINTLINK	U.S. Sprint	↔ peer	(remove entry)
9	6762	SEABONE-NET	Info-tel Communication S.r.l.	↔ peer	<input type="text"/>
10	6453	AS6453	TATA Communications formerly VSNL is Leading ISP	↔ peer	<input type="text"/>
11	3320	DTAG	Deutsche Telekom AG	↔ peer	<input type="text"/>
12	2828	XO-AS15	XO Communications	↔ peer	<input type="text"/>
13	3561	SAVVIS	CenturyTel Internet Holdings, Inc.	↔ peer	<input type="text"/>
14	7018	ATT-INTERNET4	AT&T Services, Inc.	↔ peer	<input type="text"/>
15	209	ASN-QWEST	CenturyTel Internet Holdings, Inc.	↔ peer	<input type="text"/>
20	6461	ABOVENET	Abovenet Communications, Inc	↔ peer	<input type="text"/>
29	12389	ROSTELECOM-AS	Neology (Pty) Ltd	↔ peer	<input type="text"/>
31	5511	Opentransit	eTel Austria Gesmbh u. CO KG	↔ peer	<input type="text"/>

Background – Valid paths

valid

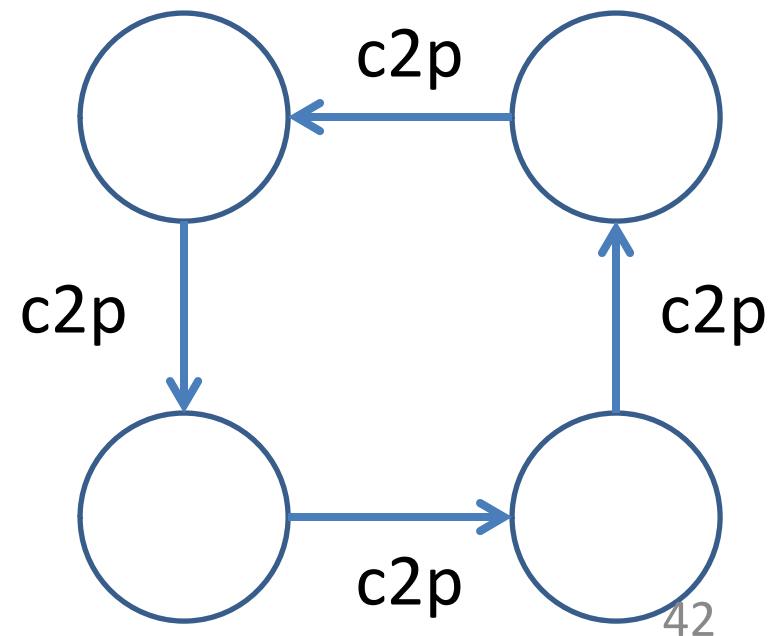


invalid



Definition – Strongly Connected Components

- Assumption: AS topology is hierarchical
 - Therefore AS graph is a DAG
- Avoid creating a “Strongly Connected Component” of c2p links between ASes



Inferring Clique

- Infer clique of ASes which includes Tier-1 ASes
 1. Apply Bron/Kerbosch 1973 algorithm to top 10 ASes by transit degree.
 2. Do not admit an AS to the clique if it appears to receive transit from other members of the clique
 - Add additional ASes to the clique provided they do not break these rules
 - Infer p2p between all members of the clique