AS RELATIONSHIPS, CUSTOMER CONES, AND VALIDATION

Matthew Luckie, Bradley Huffaker, Amogh Dhamdhere,
Vasileios Giotsas*, k claffy
mjl@caida.org

CAIDA - University of California, San Diego
*University College London

ACM SIGCOMM Internet Measurement Conference: Barcelona, Spain, Oct 2013

CONTRIBUTIONS

- We build a new **AS relationships inference** algorithm with near-perfect accuracy

- We develop a new **customer cone inference** algorithm to address real-world complexities

- We **validate** our AS relationships to an unprecedented level

  - **99.6%** p2c, **98.7%** p2p, **34.7%** of **126,082** inferences.

- We release our code and 97% of validation data to promote **reproducibility**

`http://www.caida.org/publications/papers/2013/asrank/`

RELATED WORKS

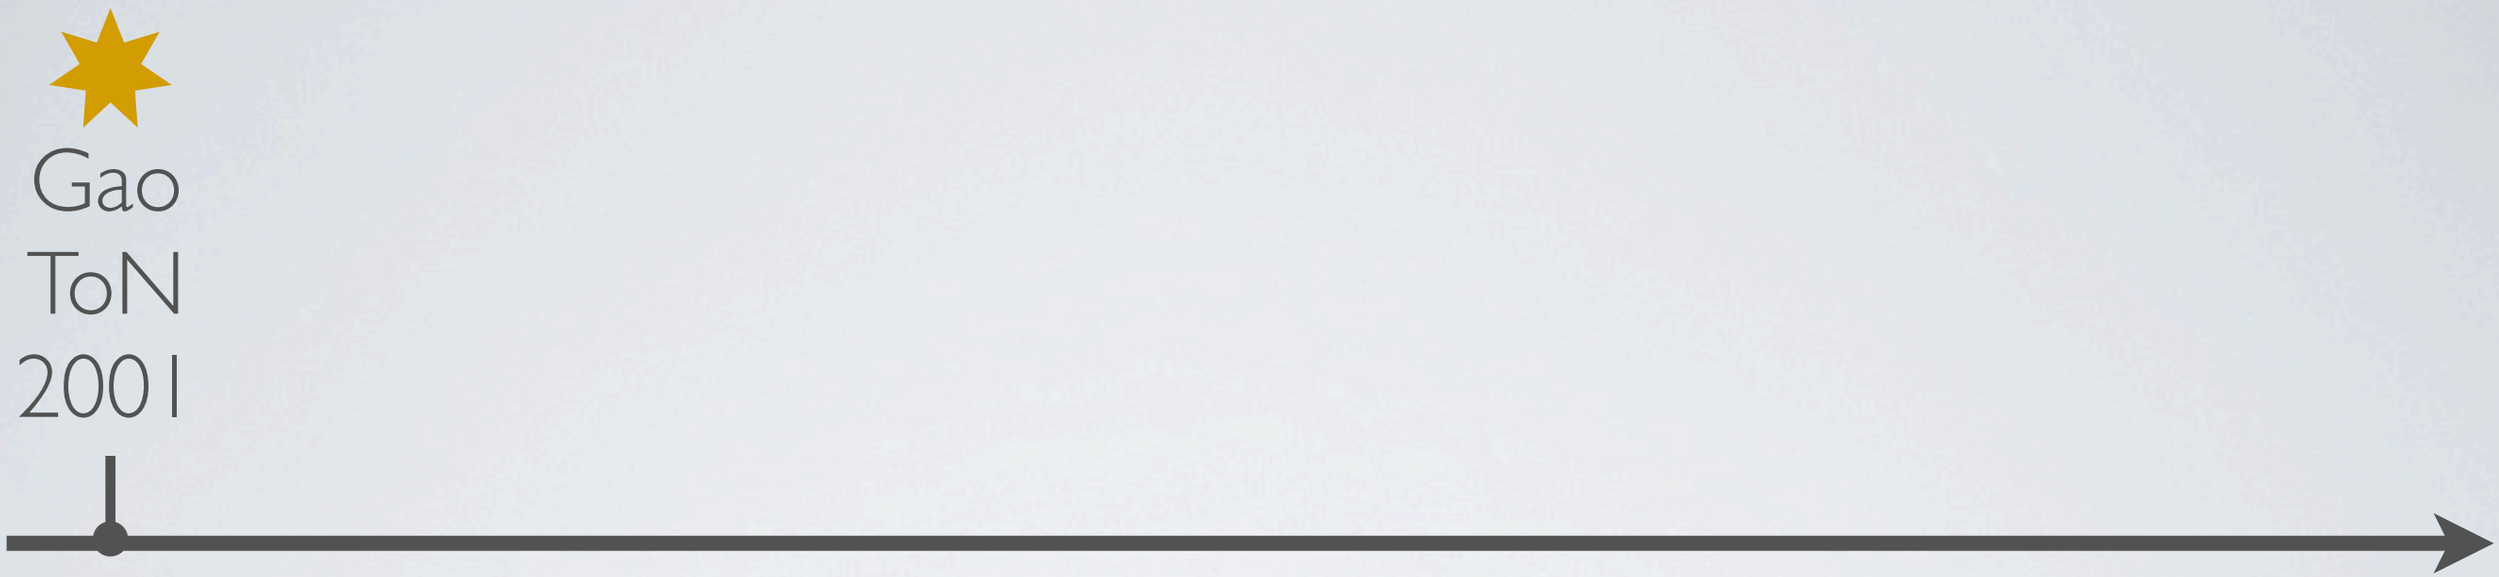

Maximise valley-free paths.
p2c, p2p, s2s
Validation: AT&T

RELATED WORKS

Formalized Gao.
Conjectured ToR NP-complete.
SARK: rank ASes by closeness to core.
p2c, p2p.
Validation: fraction of VF paths

RELATED WORKS

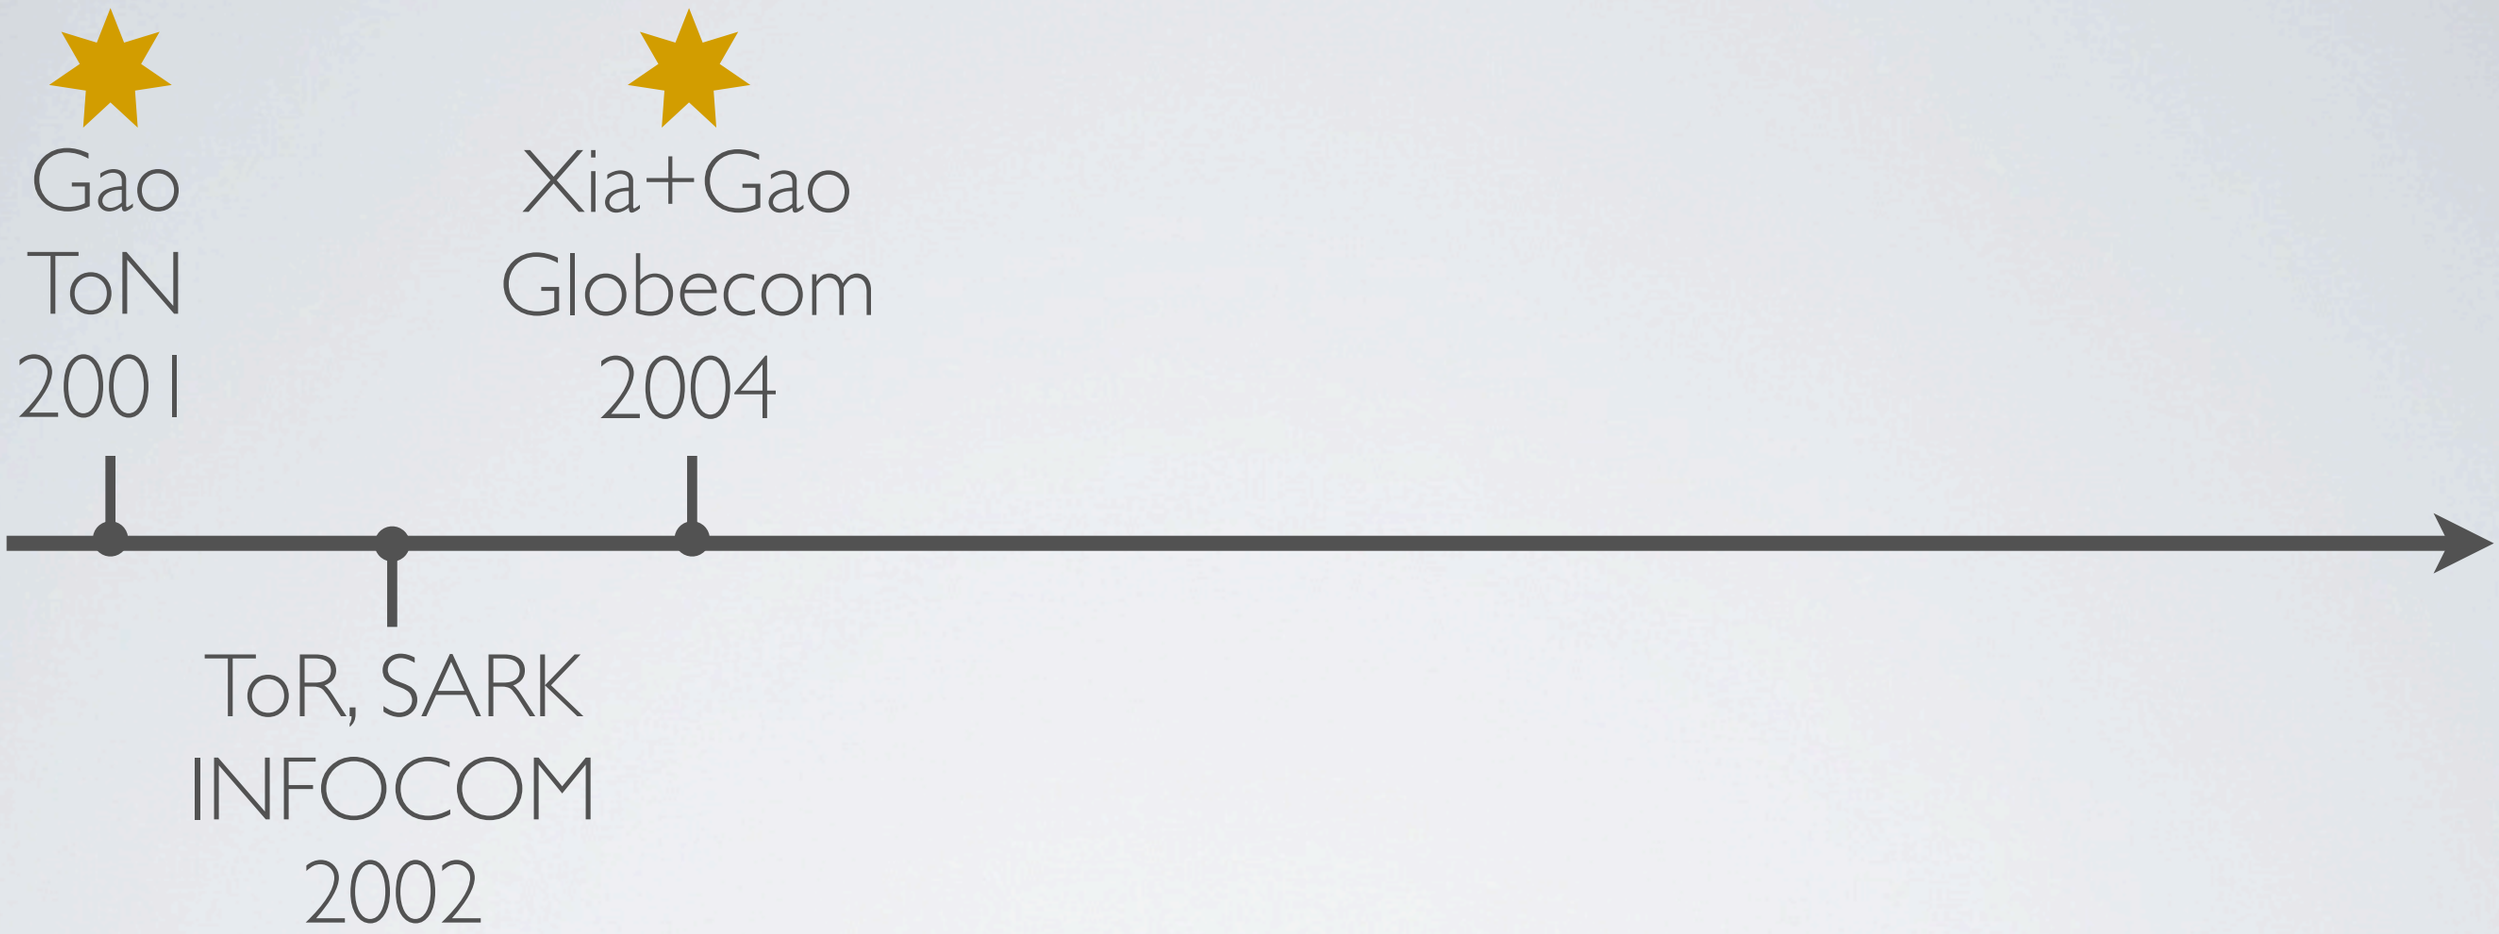

Maximise VF paths.
Seed with GT-classified links from IRR and Communities.
Validation: unused GT

RELATED WORKS

Infer clique. Links observed by clique are p2c. All others p2p. No validation.

RELATED WORKS

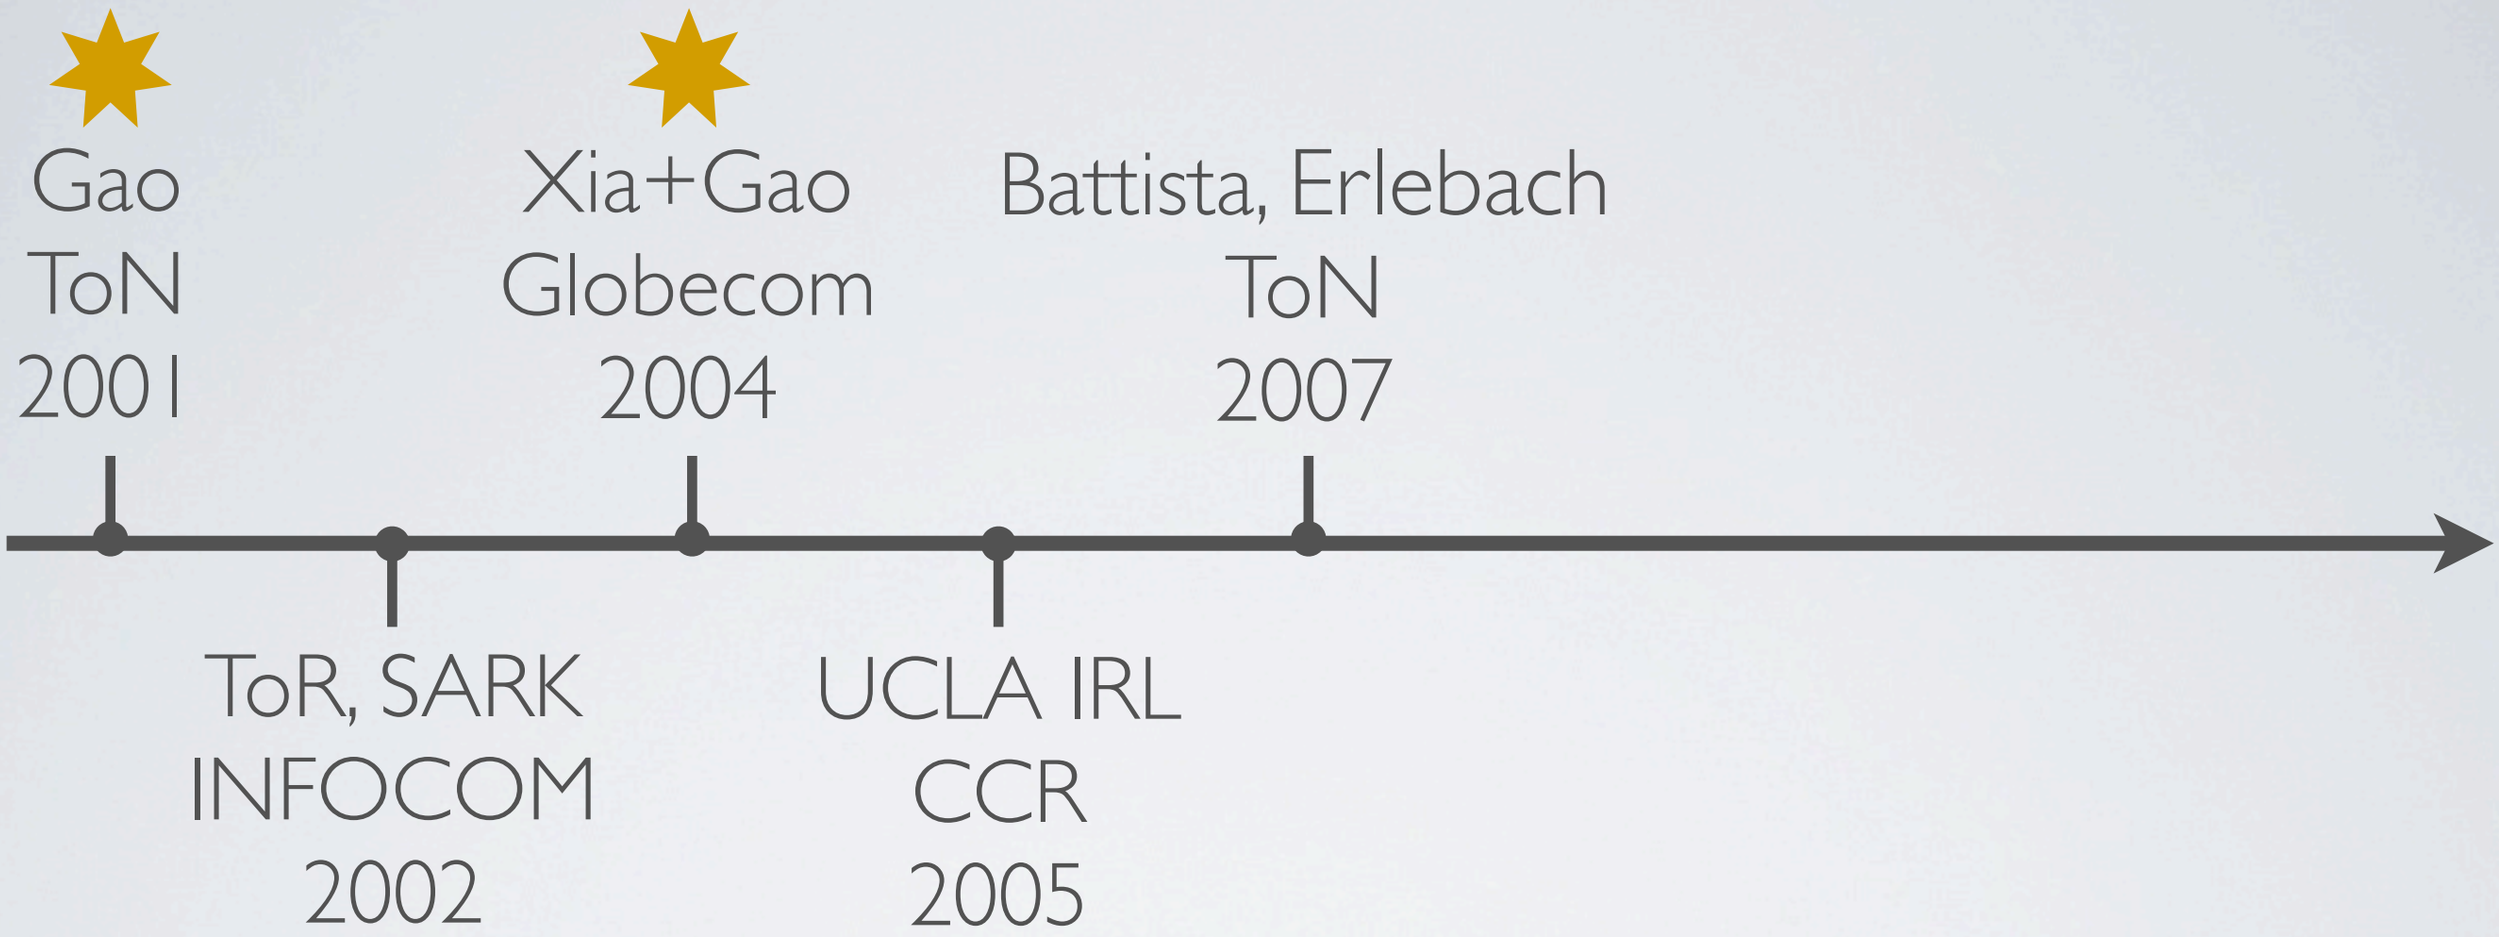

Proved ToR NP-complete.
Infer p2c, leave s2s+p2p inferences as future work

RELATED WORKS

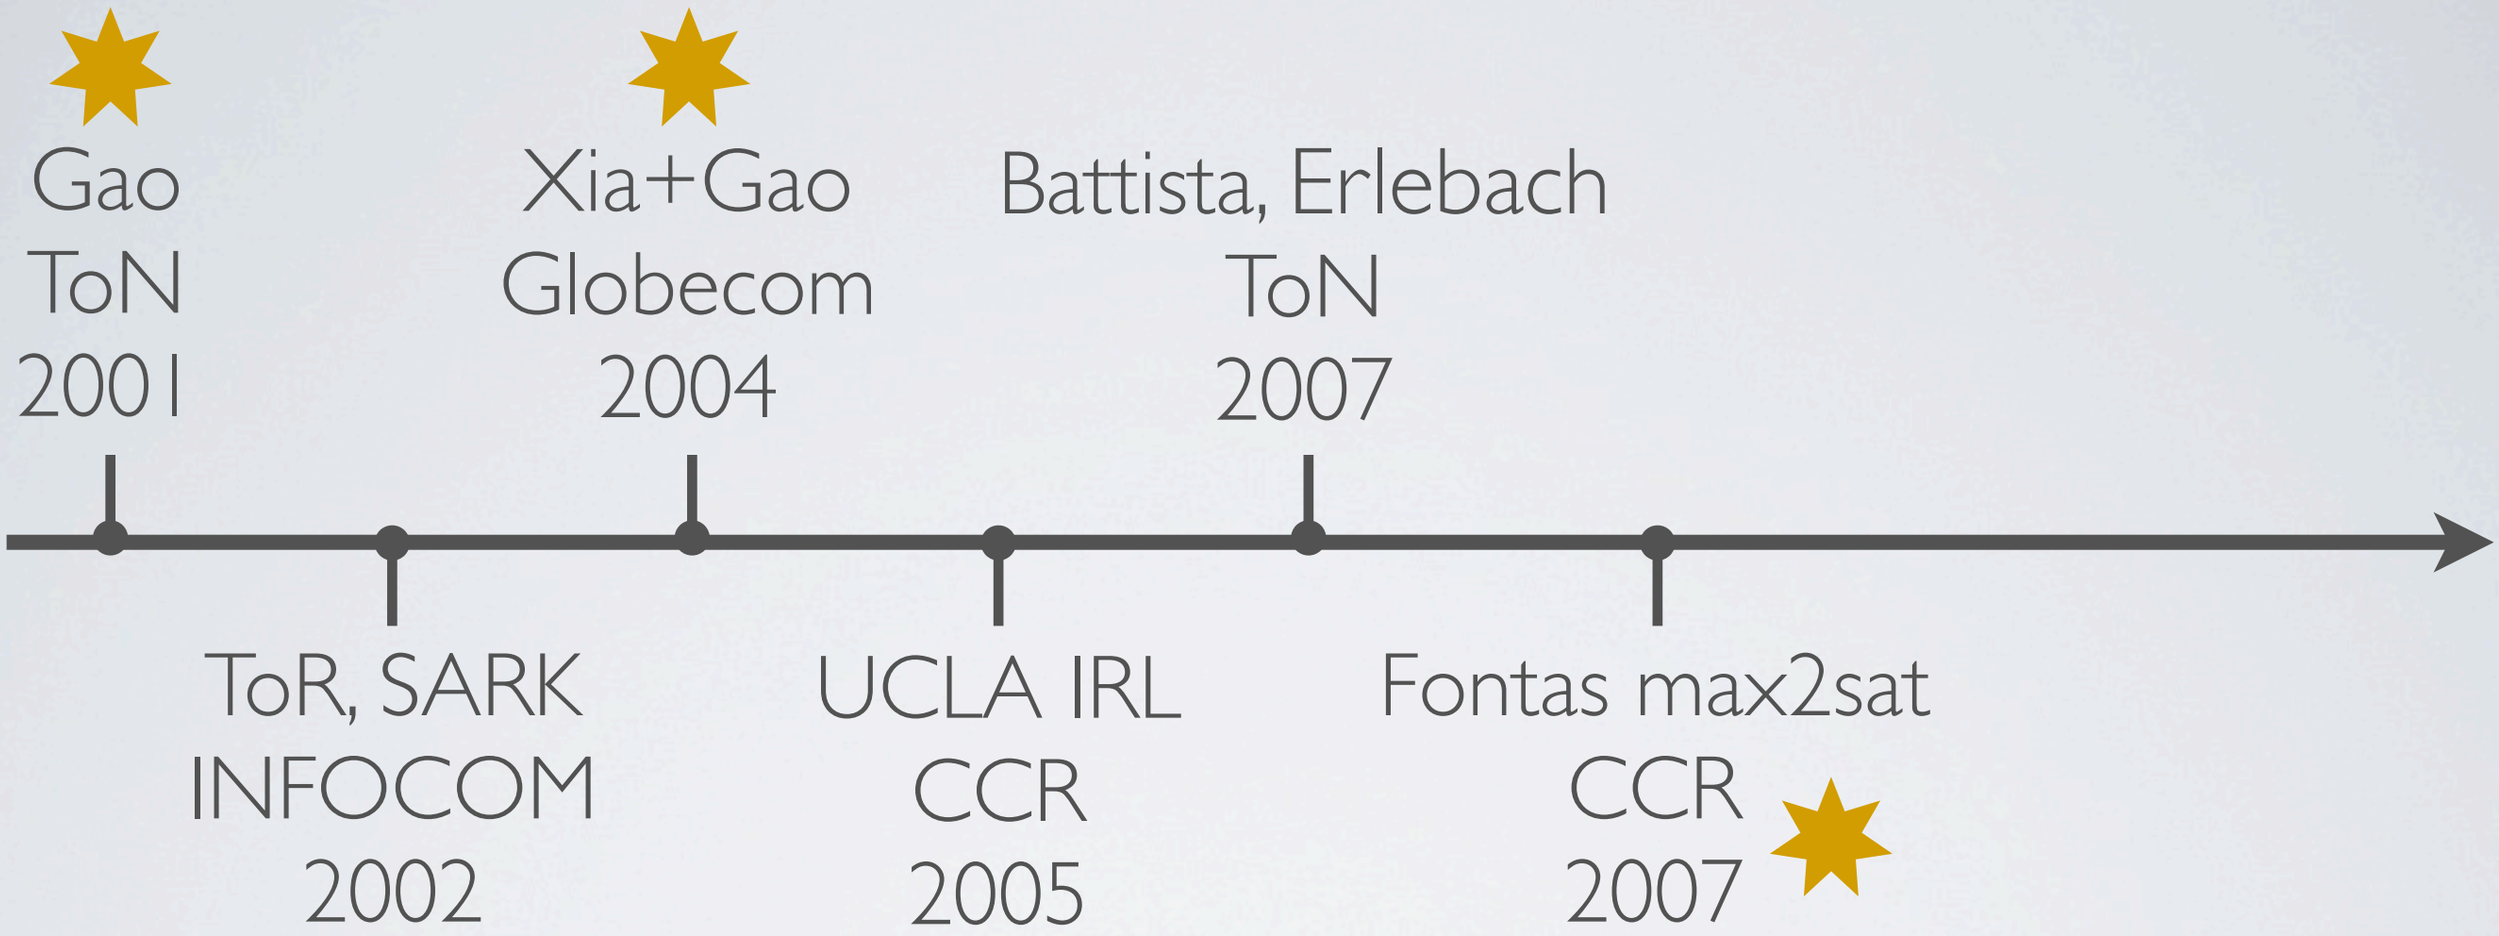

Find solution that maximises
(1) fraction of VF paths,
(2) provider deg. > customer deg.
Siblings from WHOIS.
Validation: 9.7% of inferences
MAX-2-SAT: NP-hard.

caida

RELATED WORKS

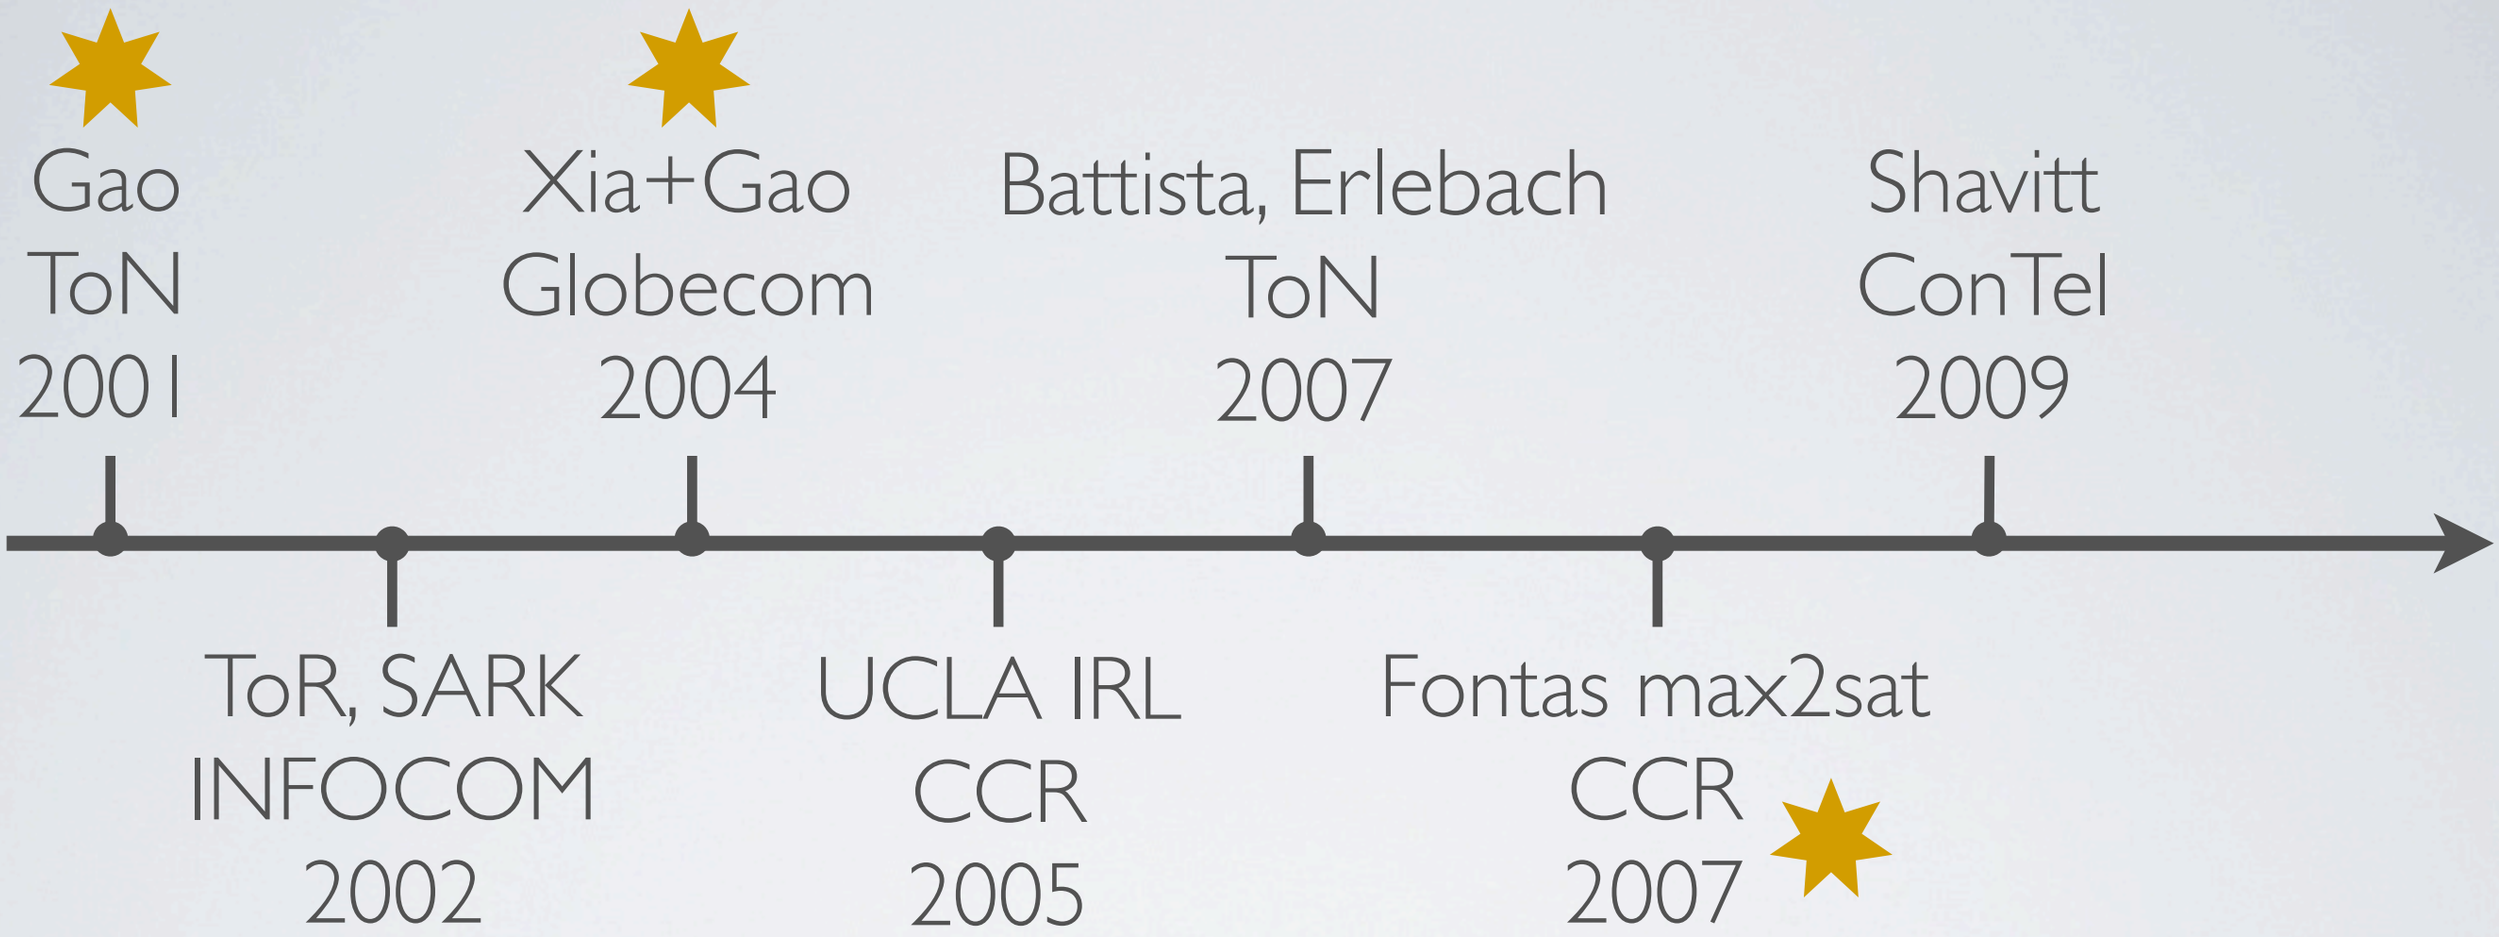

Infers core, VF paths through core, cascades VF paths. No validation.

RELATED WORKS

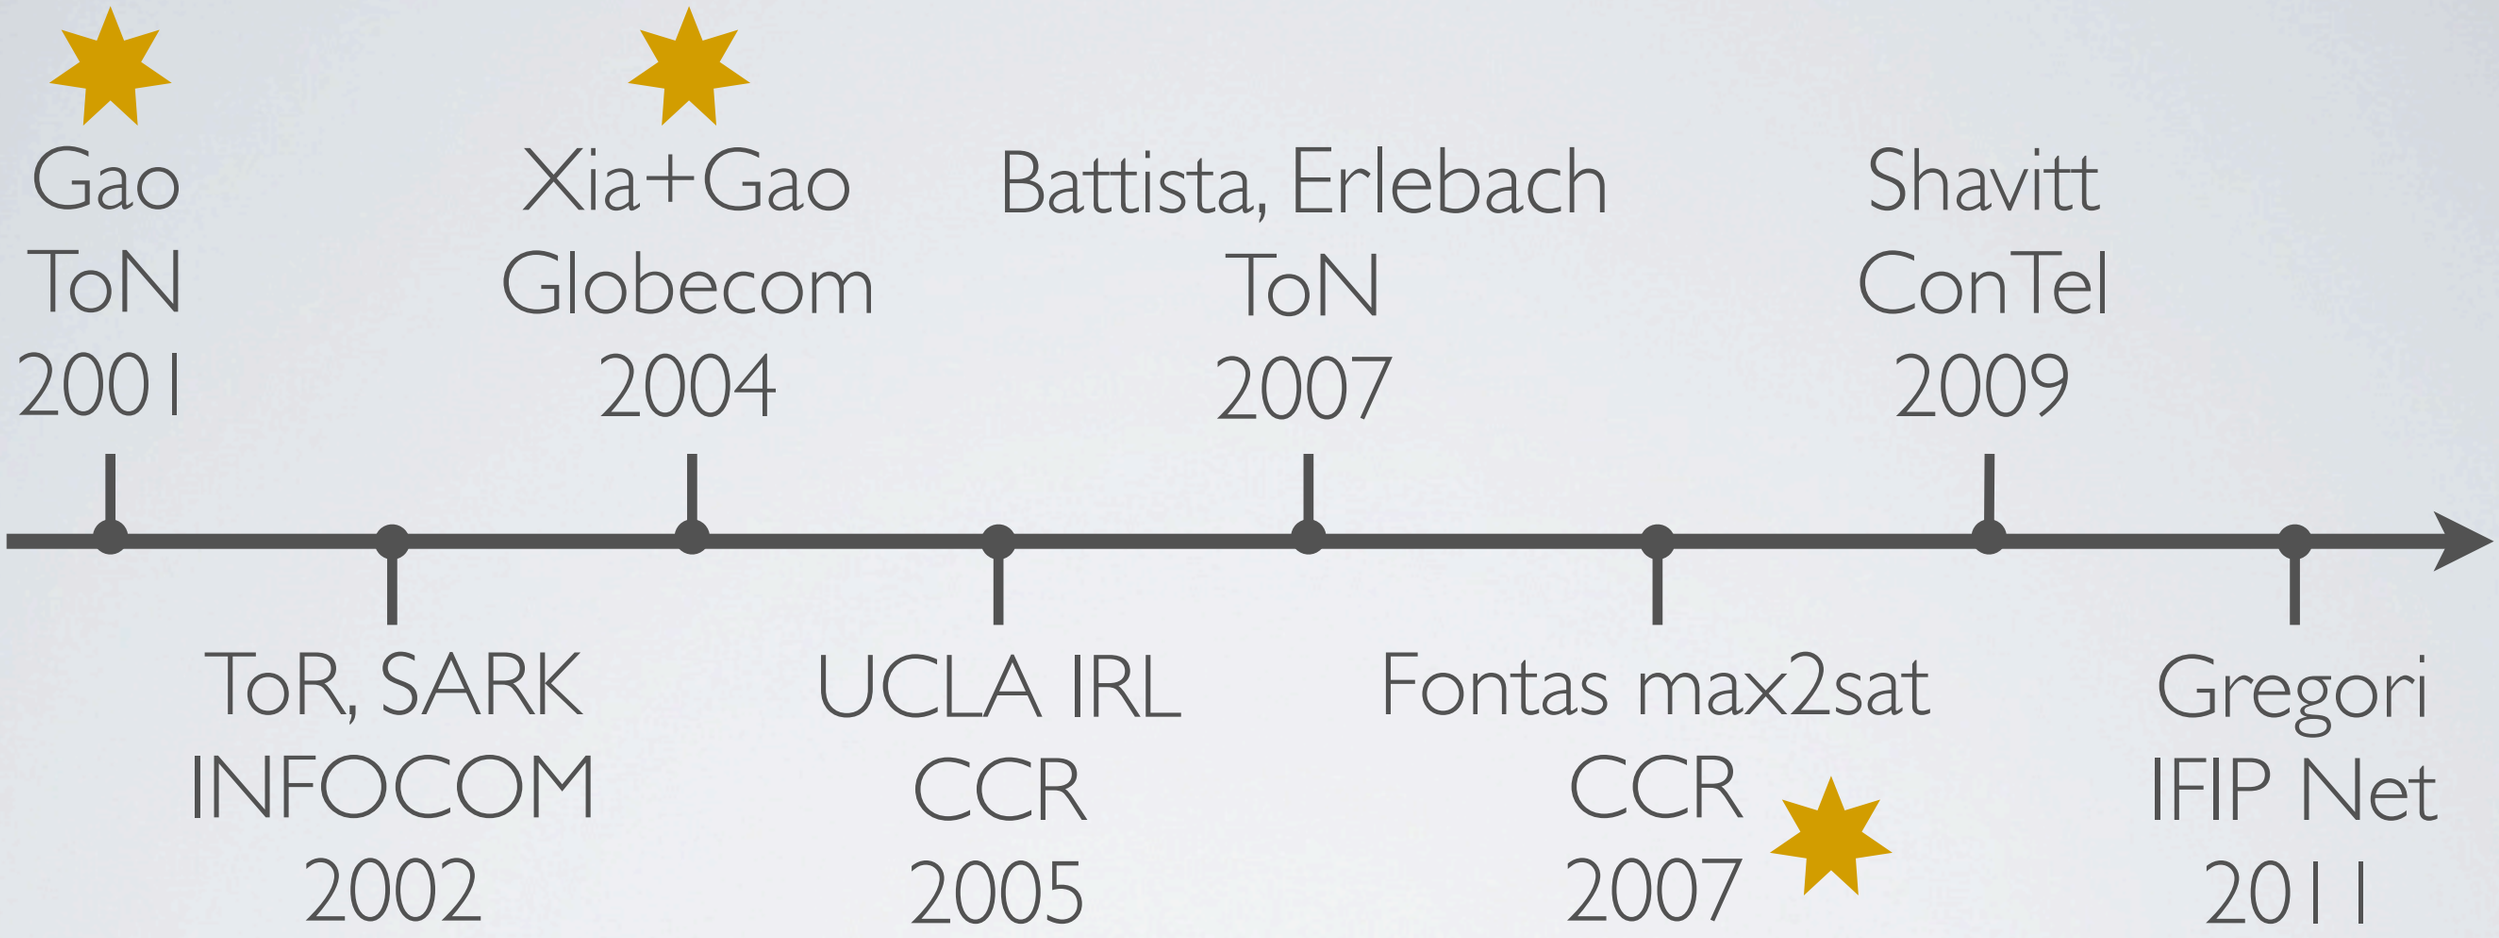

Similar to UCLA w/clique visability.
Uses lifetime of paths.
No validation.

CHALLENGES

- **Artifacts**: misconfigurations, poisoned paths, route leaks
- **Limited visibility**: observe a tiny fraction of peering links
- **Valley Free assumption** does not always hold
- **Complex** import and export filters cause some c2p links to be region or prefix specific
- **Siblings** may exchange routes freely between themselves

OUR APPROACH

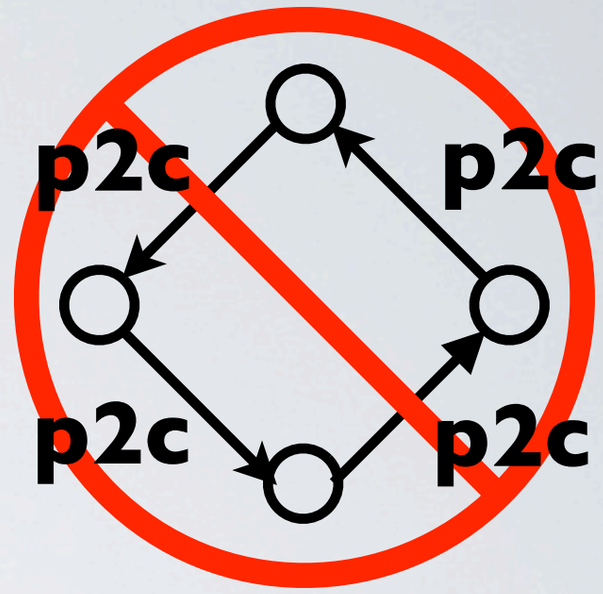

- Uses **public data**
- Three generally accepted assumptions:

  1. **clique** of large transit providers at top of hierarchy

  2. most customers purchase transit to be **reachable**

  3. AS graph **acyclic**: no cycles of p2c links

- Infers only p2c and p2p links; siblings and mutual-transit left for future work

DATA

- **BGP data**: RouteViews and RIPE RIS

Publicly available data

  - 1-5th of each month from January 1998 to August 2013

- **Validation data**, April 2012:

  - **Direct feedback** via `as-rank.caida.org`: 129 c2p, 1350 p2p
  - **RPSL** from RIPE: 6,530 c2p
  - **BGP Communities**: 23356 c2p, 16248 p2p
  - **Direct feedback** via email: 285 c2p, 689 p2p
  - **99.0% agreement** where overlap

Available as Supplementary data

TRANSIT DEGREE

(transit degree vs. node degree, also used by UCLA)

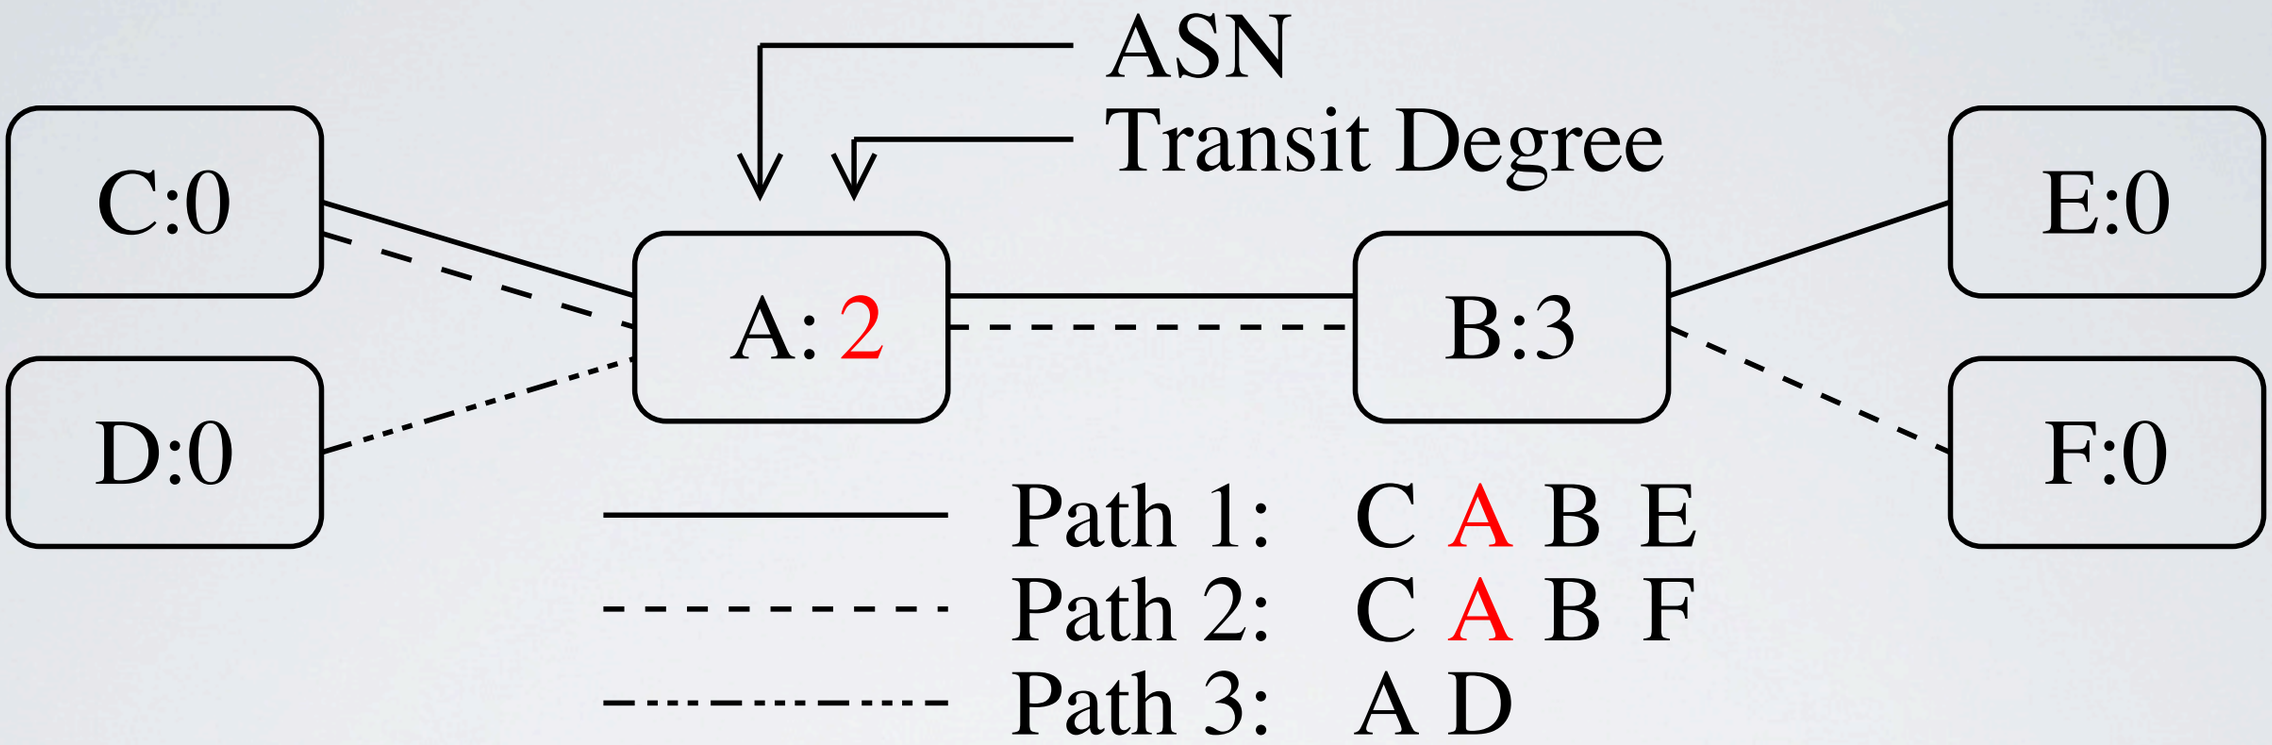

- **node degree**: the number of neighbours an AS has
- **transit degree**: the number of ASes that appear on either side of an AS in adjacent links
- Using transit degree reduces ordering errors of stub ASes with large peering visibility: i.e. stubs that provide a VP or peer with many VPs.

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
  - **56.4%** of graph @ **99.8%** PPV
  - Additional steps in algorithm for 5.9% of graph (next slides)
- All other links in graph are p2p
  - **37.7%** of graph @ **98.7%** PPV

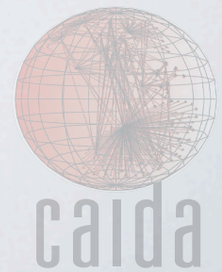

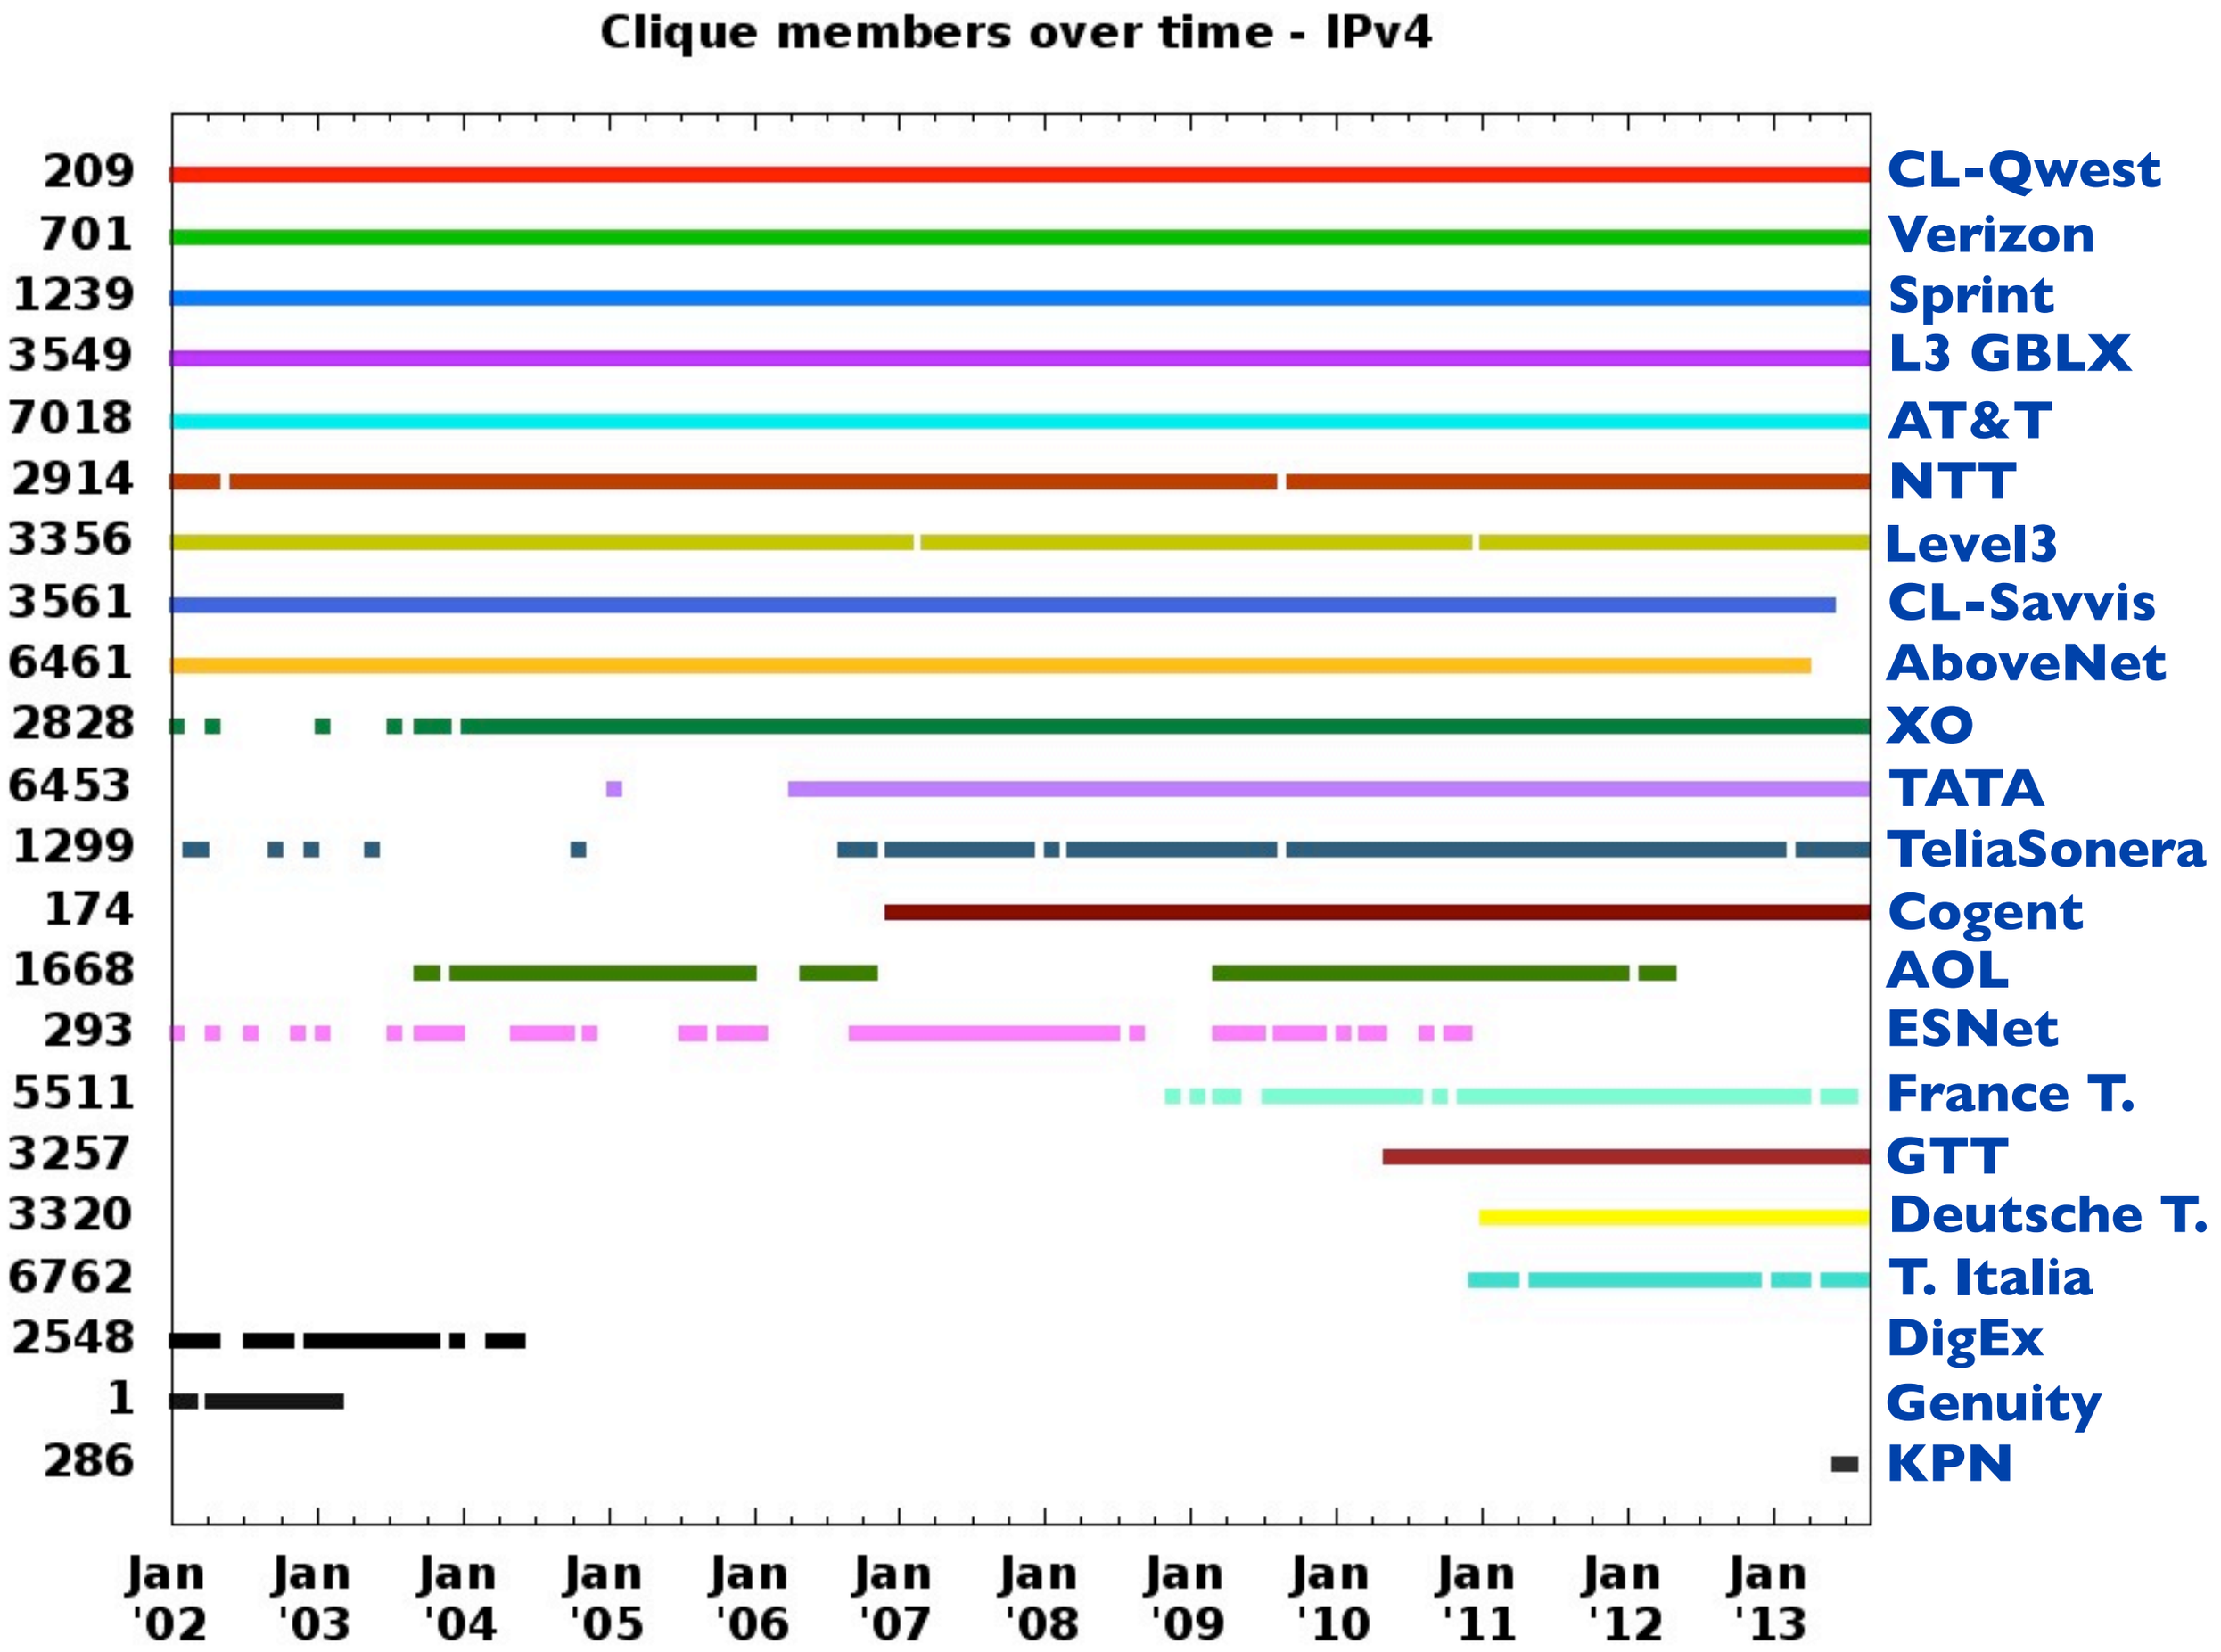
Clique members over time - IPv4
209
701
1239
3549
7018
2914
3356
3561
6461
2828
6453
1299
174
1668
293
5511
3257
3320
6762
2548
1
286
CL-Qwest
Verizon
Sprint
L3 GBLX
AT&T
NTT
Level3
CL-Savvis
AboveNet
XO
TATA
TeliaSonera
Cogent
AOL
ESNet
France T.
GTT
Deutsche T.
T. Italia
DigEx
Genuity
KPN
Jan '02
Jan '03
Jan '04
Jan '05
Jan '06
Jan '07
Jan '08
Jan '09
Jan '10
Jan '11
Jan '12
Jan '13

PROCESSING TRIPLETS

(step 5)

caida

PROCESSING TRIPLETS

(step 5)

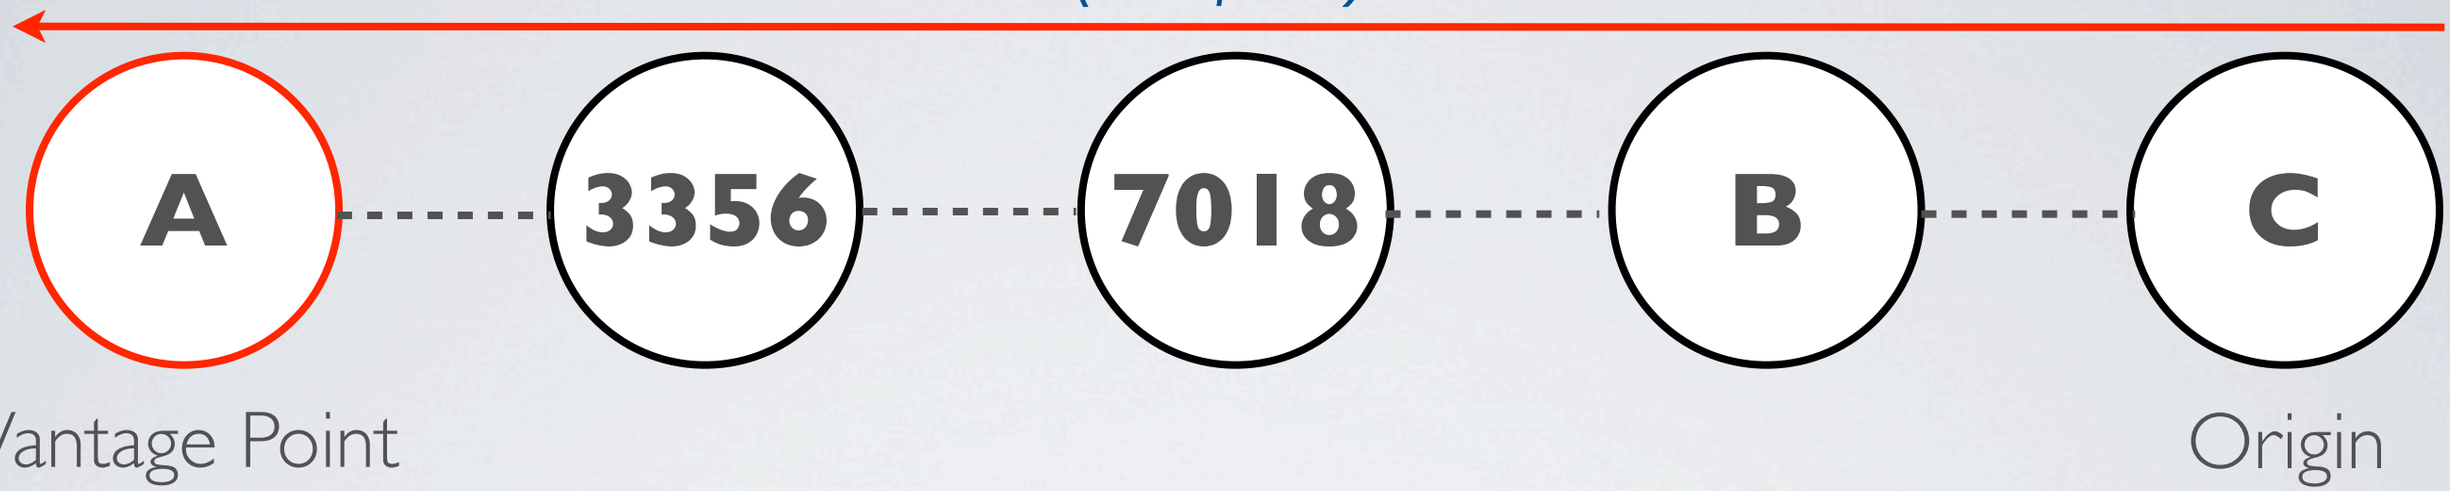

caida

PROCESSING TRIPLETS

(step 5)

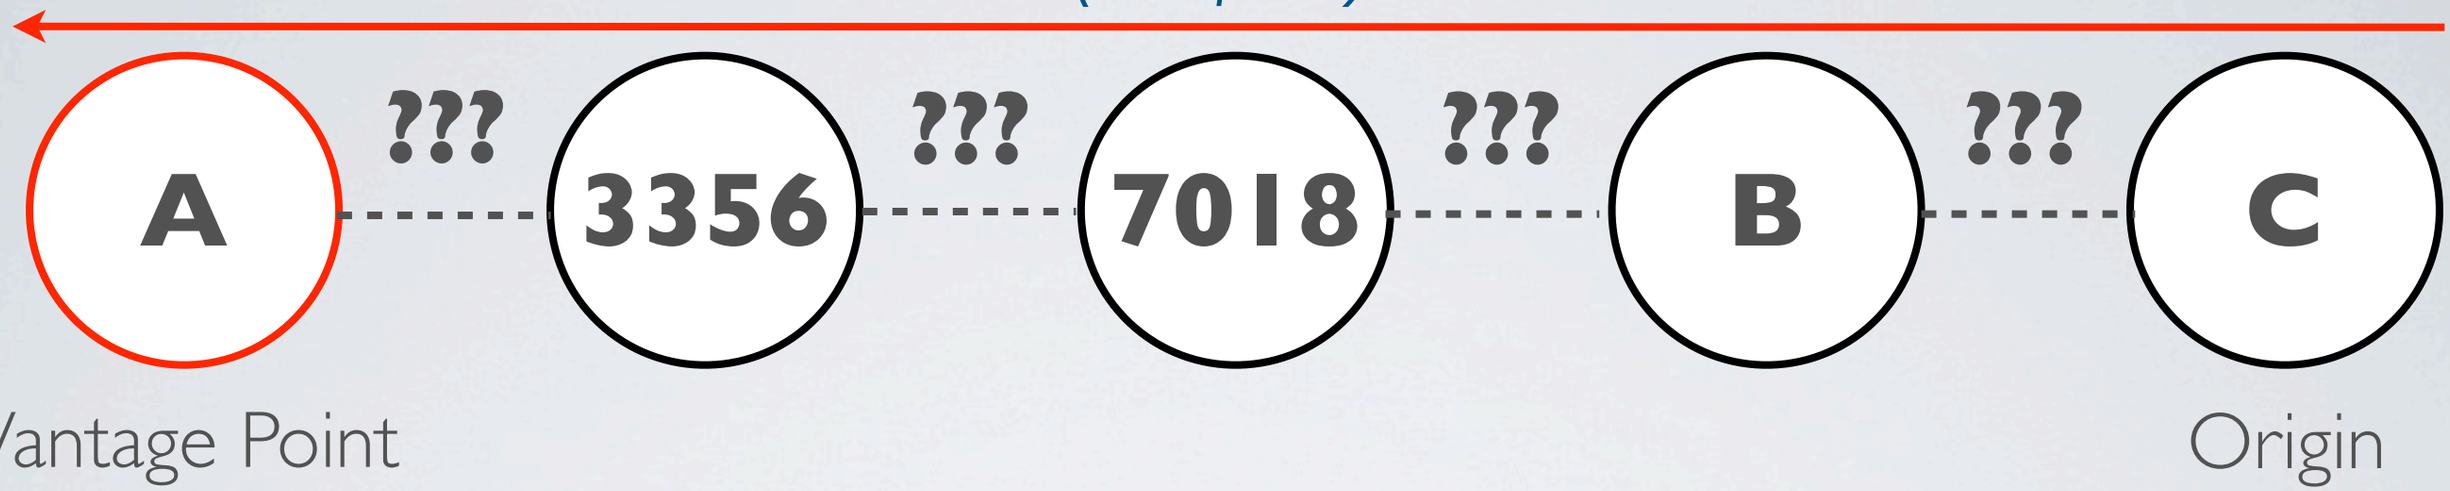

PROCESSING TRIPLETS

(step 5)

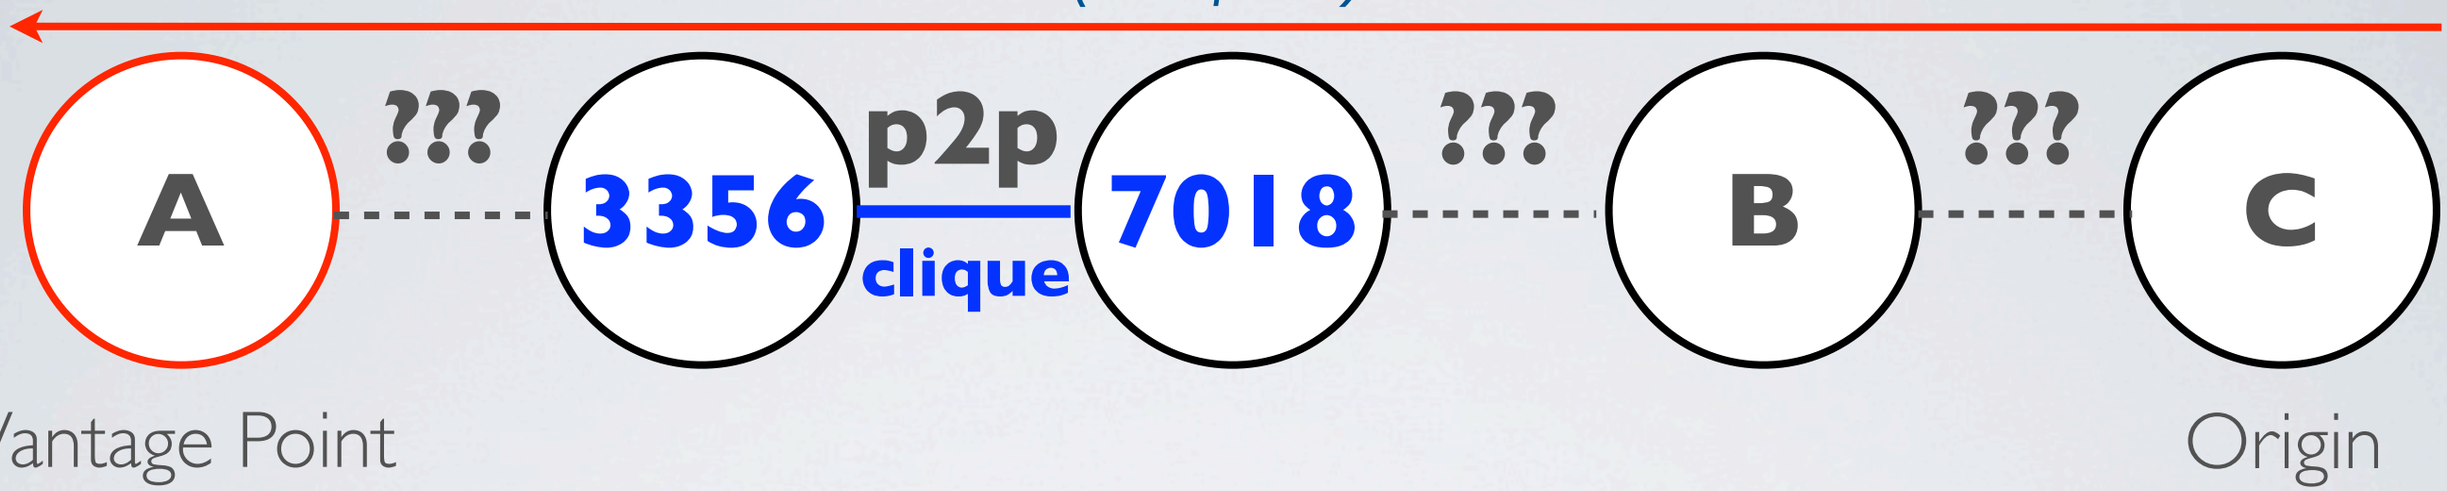

caida

PROCESSING TRIPLETS

(step 5)

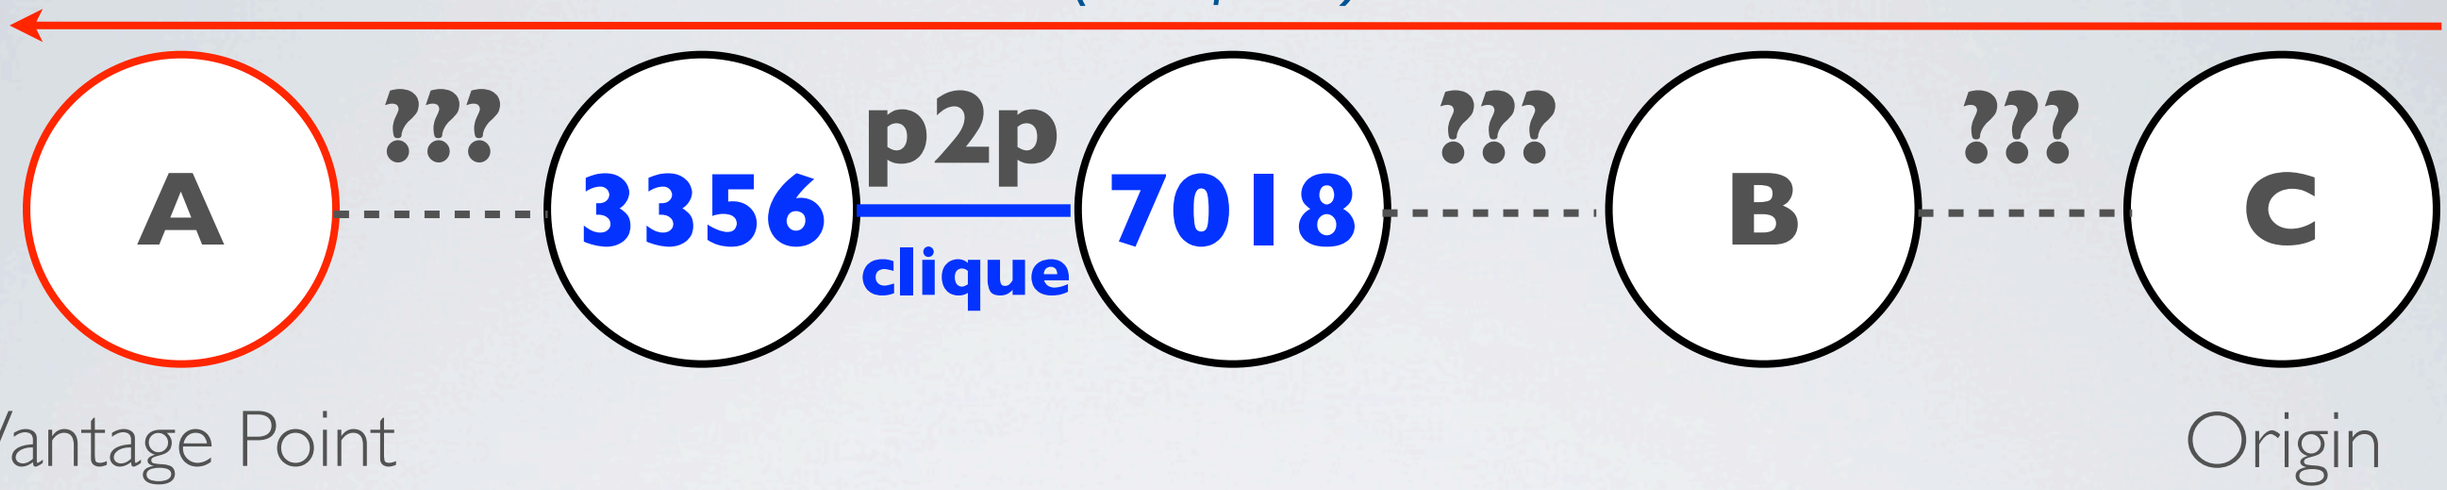

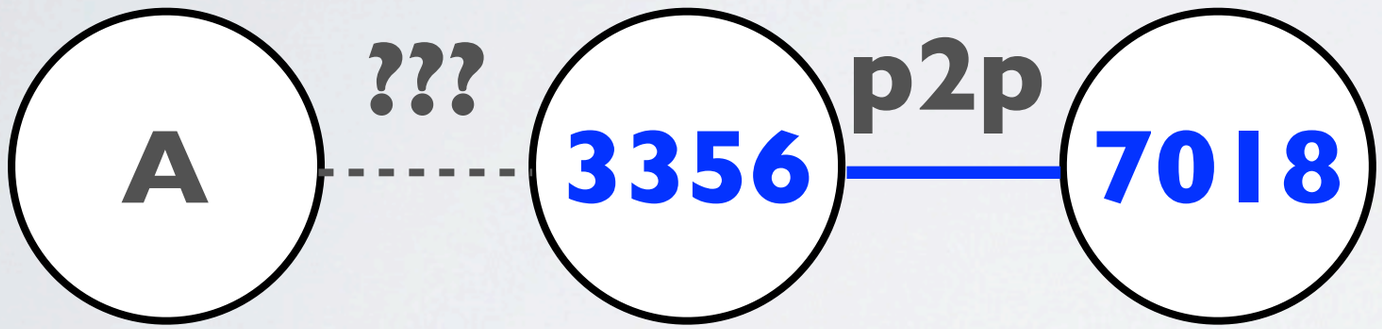

PROCESSING TRIPLETS

(step 5)

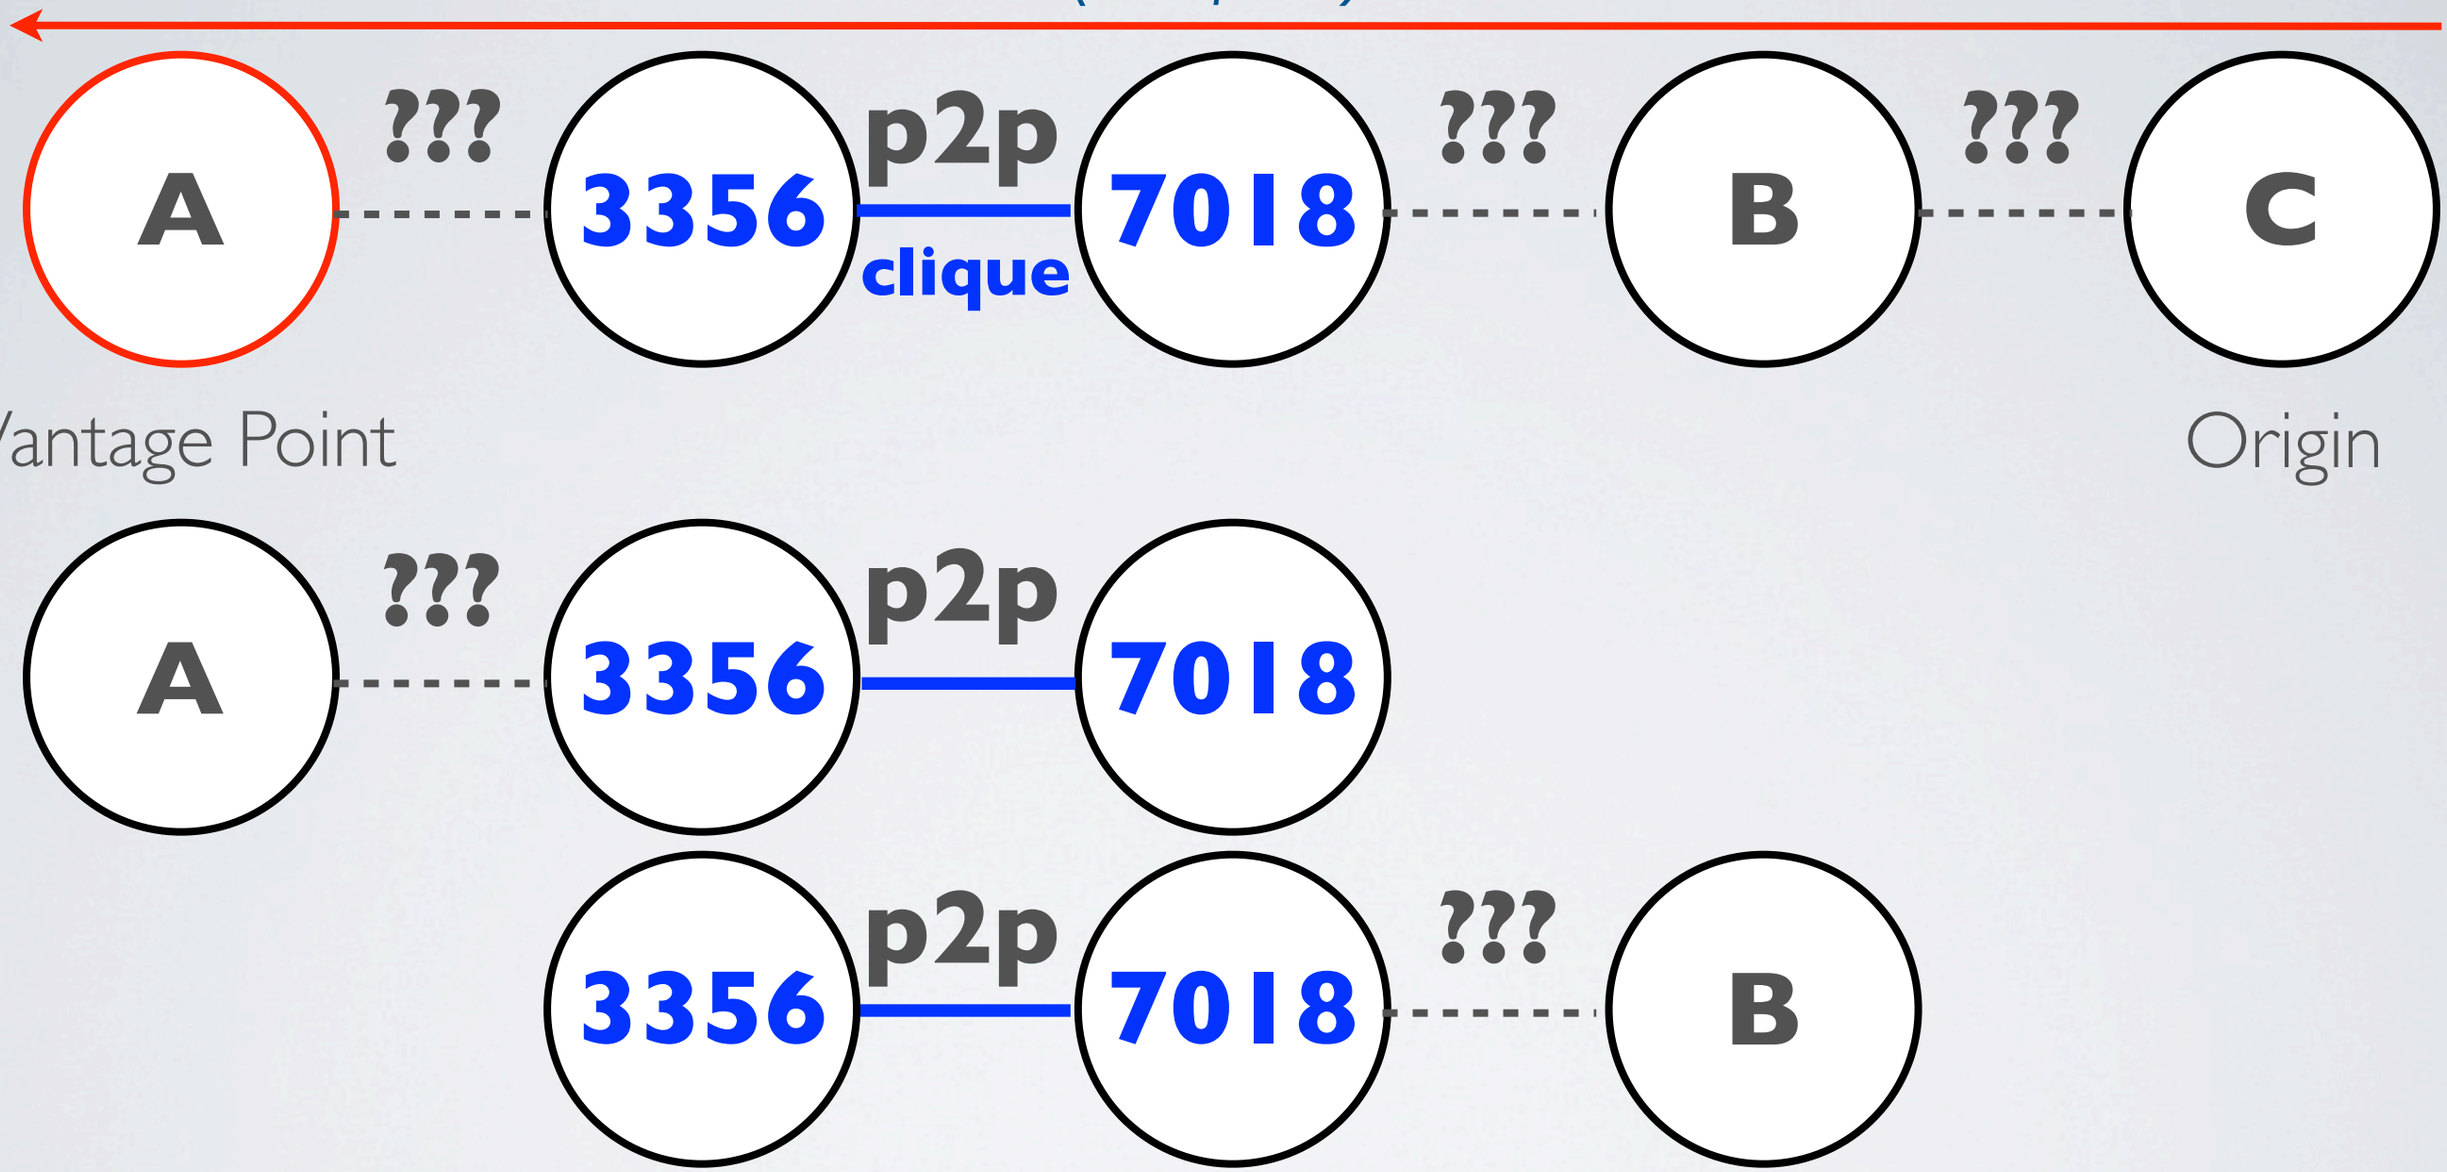

PROCESSING TRIPLETS

(step 5)

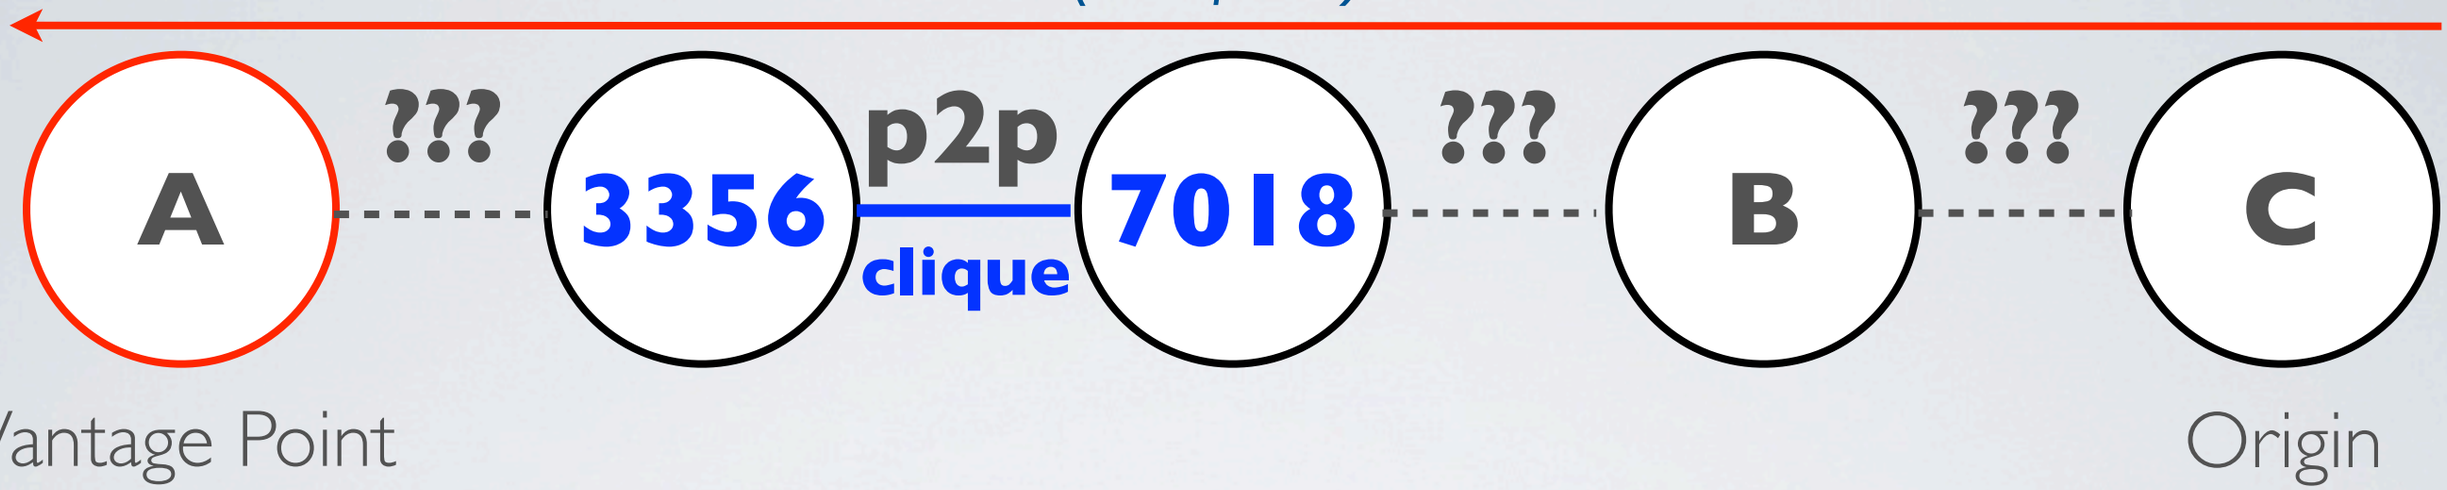

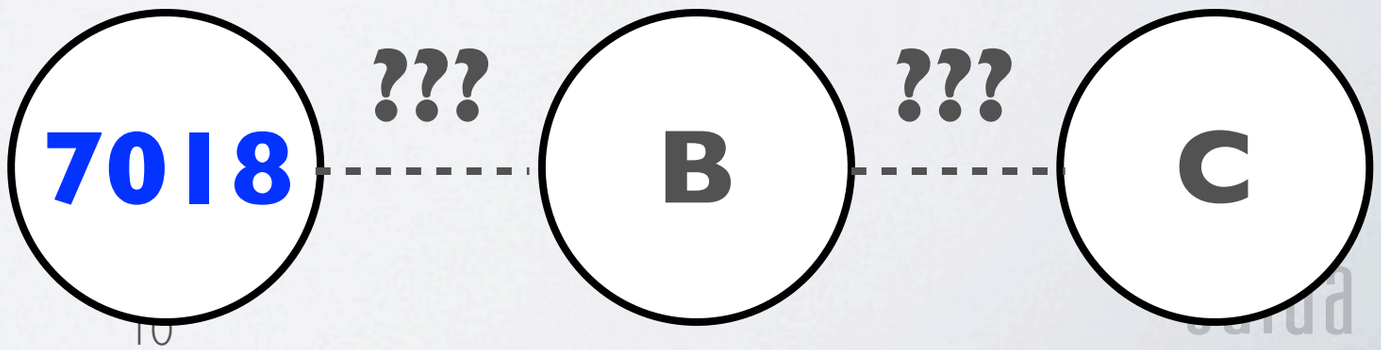

PROCESSING TRIPLETS

(step 5)

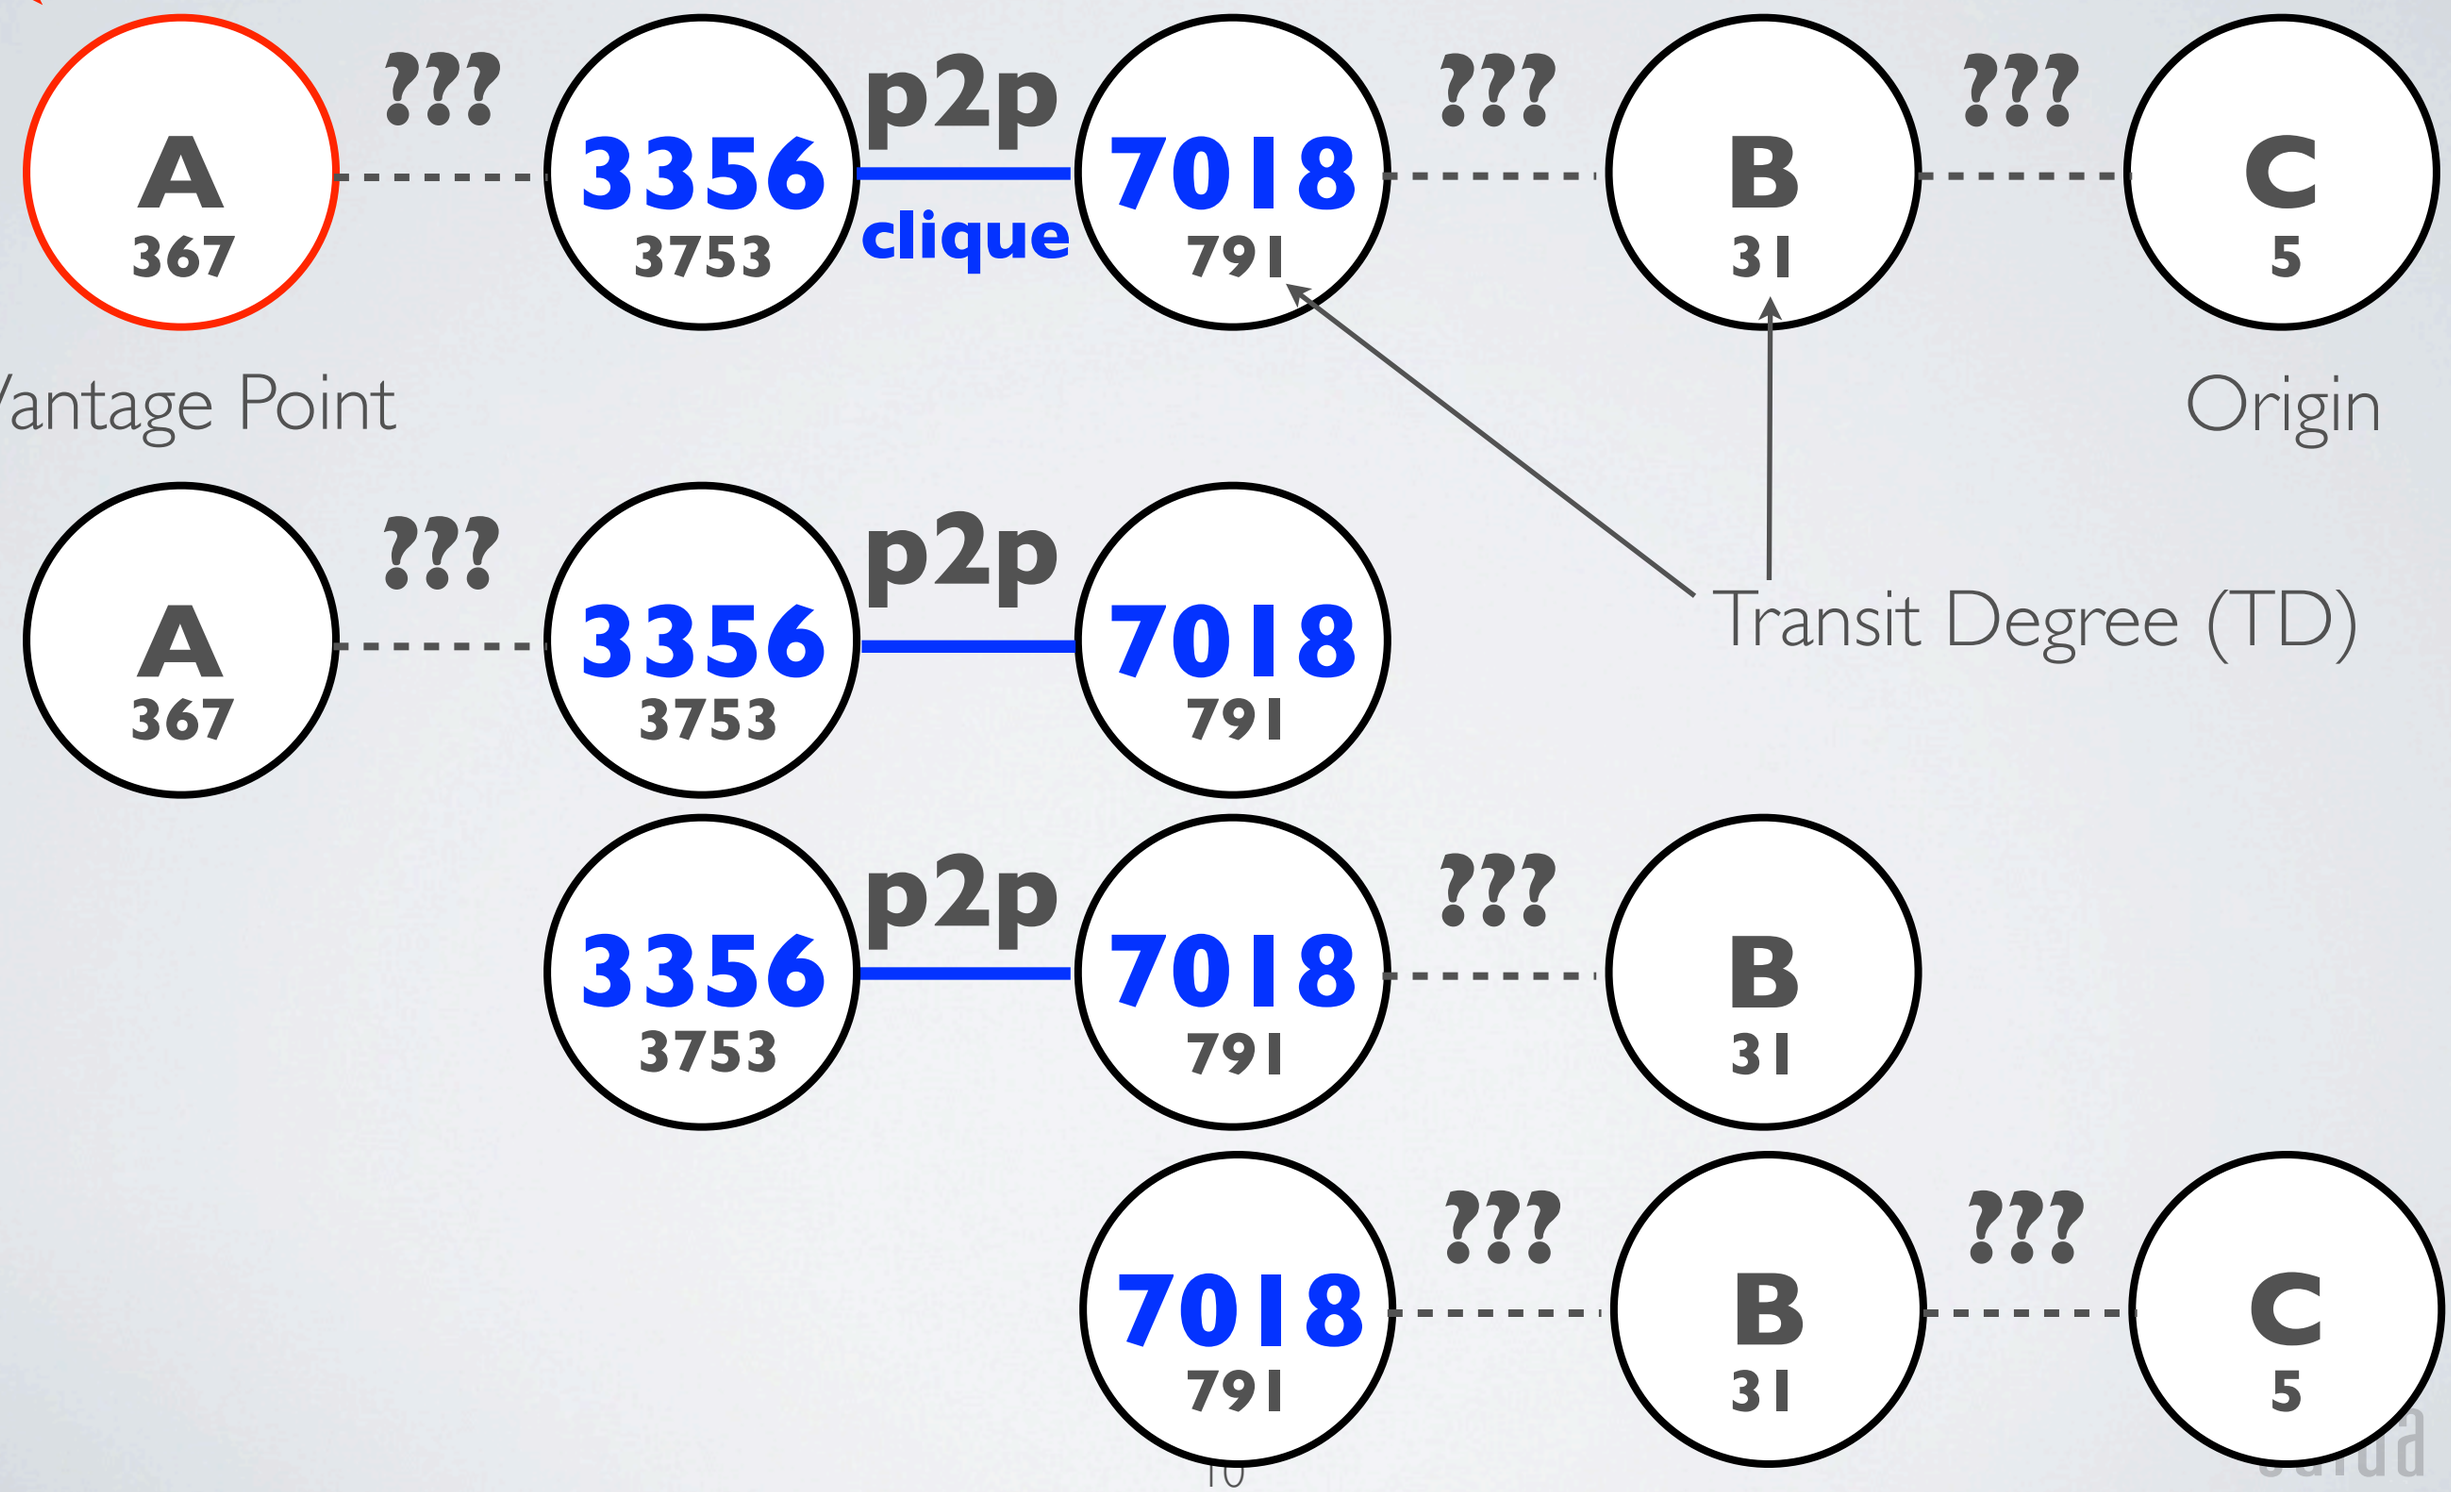

C2P INFERENCE, TOP-DOWN

(step 5, visit ASes by transit degree)

caida

C2P INFERENCE, TOP-DOWN

(step 5, visit ASes by transit degree)

No inference made: A might be a peer and 3356 *might* be leaking. Need to observe a path where provider is in front of its customer.

C2P INFERENCE, TOP-DOWN

(step 5, visit ASes by transit degree)

No inference made: A might be a peer and 3356 *might* be leaking. Need to observe a path where provider is in front of its customer.

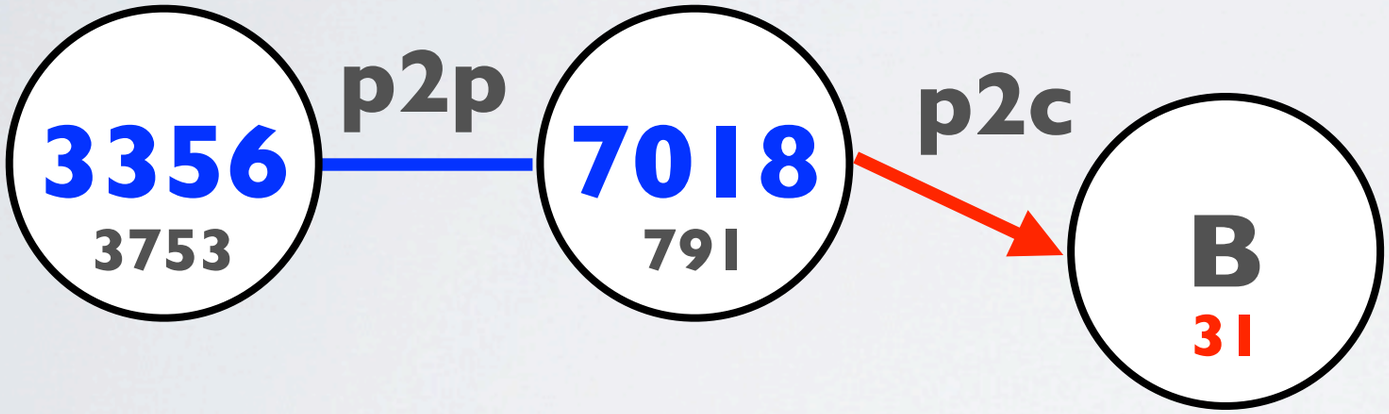

Infer B is a customer of 7018 because 7018 and 3356 are members of the clique and 7018 advertises across clique.

C2P INFERENCE, TOP-DOWN

(step 5, visit ASes by transit degree)

No inference made: A might be a peer and 3356 *might* be leaking. Need to observe a path where provider is in front of its customer.

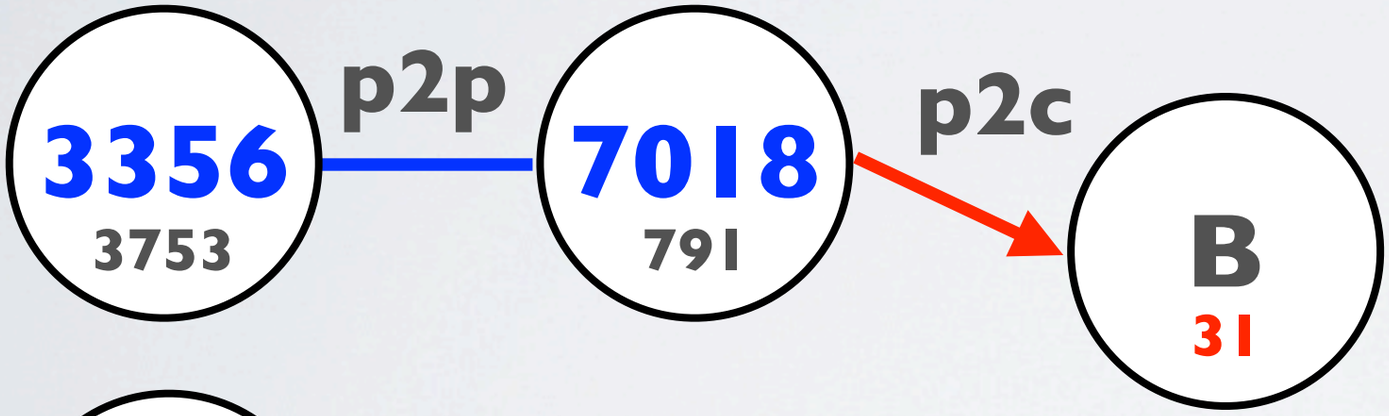

Infer B is a customer of 7018 because 7018 and 3356 are members of the clique and 7018 advertises across clique.

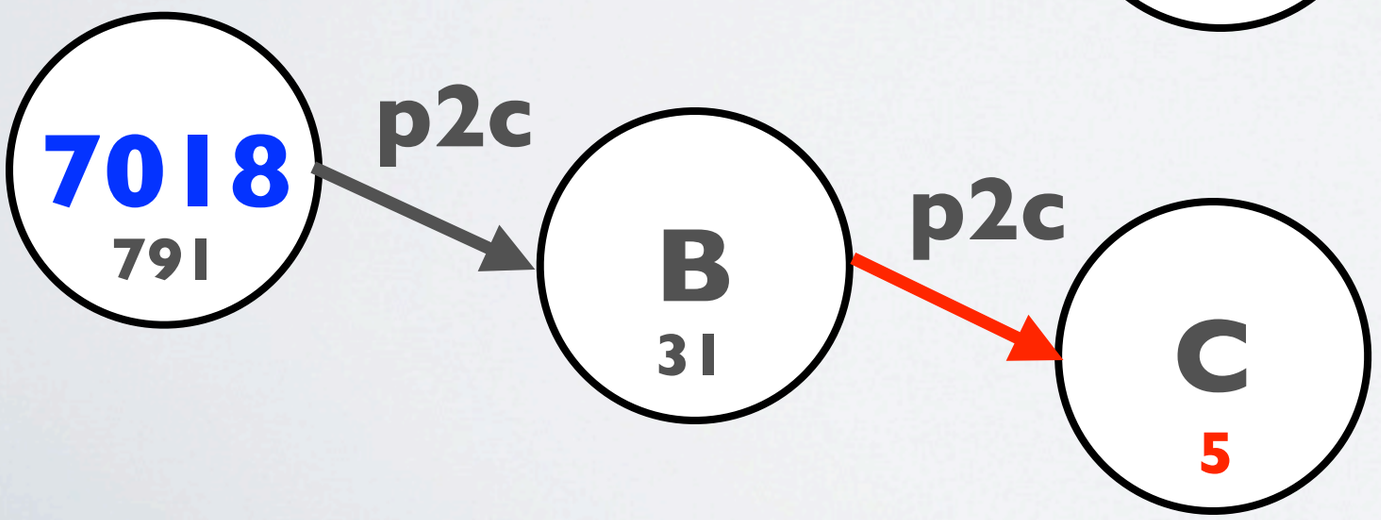

Infer C is a customer of B because B advertises route to provider (7018)

SPECIAL CASES

(see paper for full details)

step	description	PPV	fraction
6	VPs announcing no provider routes	c2p: 99.1%	0.42%
7	Smaller degree providers	c2p: 96.1%	1.92%
8	ASes with no providers	c2p: 93.3%	0.67%
		p2p: 96.7%	0.26%
9	stub-clique	c2p: 95.0%	0.52%
10	adjacent links with no relationships	c2p: 94.7%	1.96%
			5.9%

PPV: Positive Predictive Value

VALIDATION RESULTS

(April 2012)

Algorithm	c2p			p2p		
	PPV (%)	TPR (%)	Errs (1/)	PPV (%)	TPR (%)	Errs (1/)
CAIDA	**99.6**	**99.3**	**250**	**98.7**	**99.3**	**77**
UCLA	99.0	94.7	100	91.7	98.8	12
Xia+Gao	91.3	98.6	11	96.6	81.1	29
Isolario	90.3	98.0	10	96.0	82.4	25
Gao	82.9	99.8	5.8	99.5	62.5	200

TPR: True Positive Rate

caida

COMPLEX RELATIONSHIPS

(real world problems, please see paper)

- **Sibling Relationships and Mutual Transit**
  - Indistinguishable from each other, poisoning, leaking.
  - No solution currently; as2org unreliable
- **Partial Transit and Traffic Engineering**
  - Handle in "customer cone"
- **Paid Peering**
  - Unable to observe financial flows
- **Backup Transit**
  - Rare in public BGP data. Mostly inferred as p2p.

caida

CUSTOMER CONE METHODS

(recursive vs. provider-peer observed)

- Which ASes can an AS reach by following a customer link?
  - ASes generally will not peer with an AS that is not a customer but is in their customer cone
  - my customers, my customers' customers, ...
  - A recursive definition follows, but is incorrect due to real-world complexities
- We introduce the "provider-peer observed" customer cone
  - **A's customer cone contains ASes in routes announced by A's providers and peers**

CUSTOMER CONES OVER TIME

(relative, for 3 largest ASes at any time over 15 years)

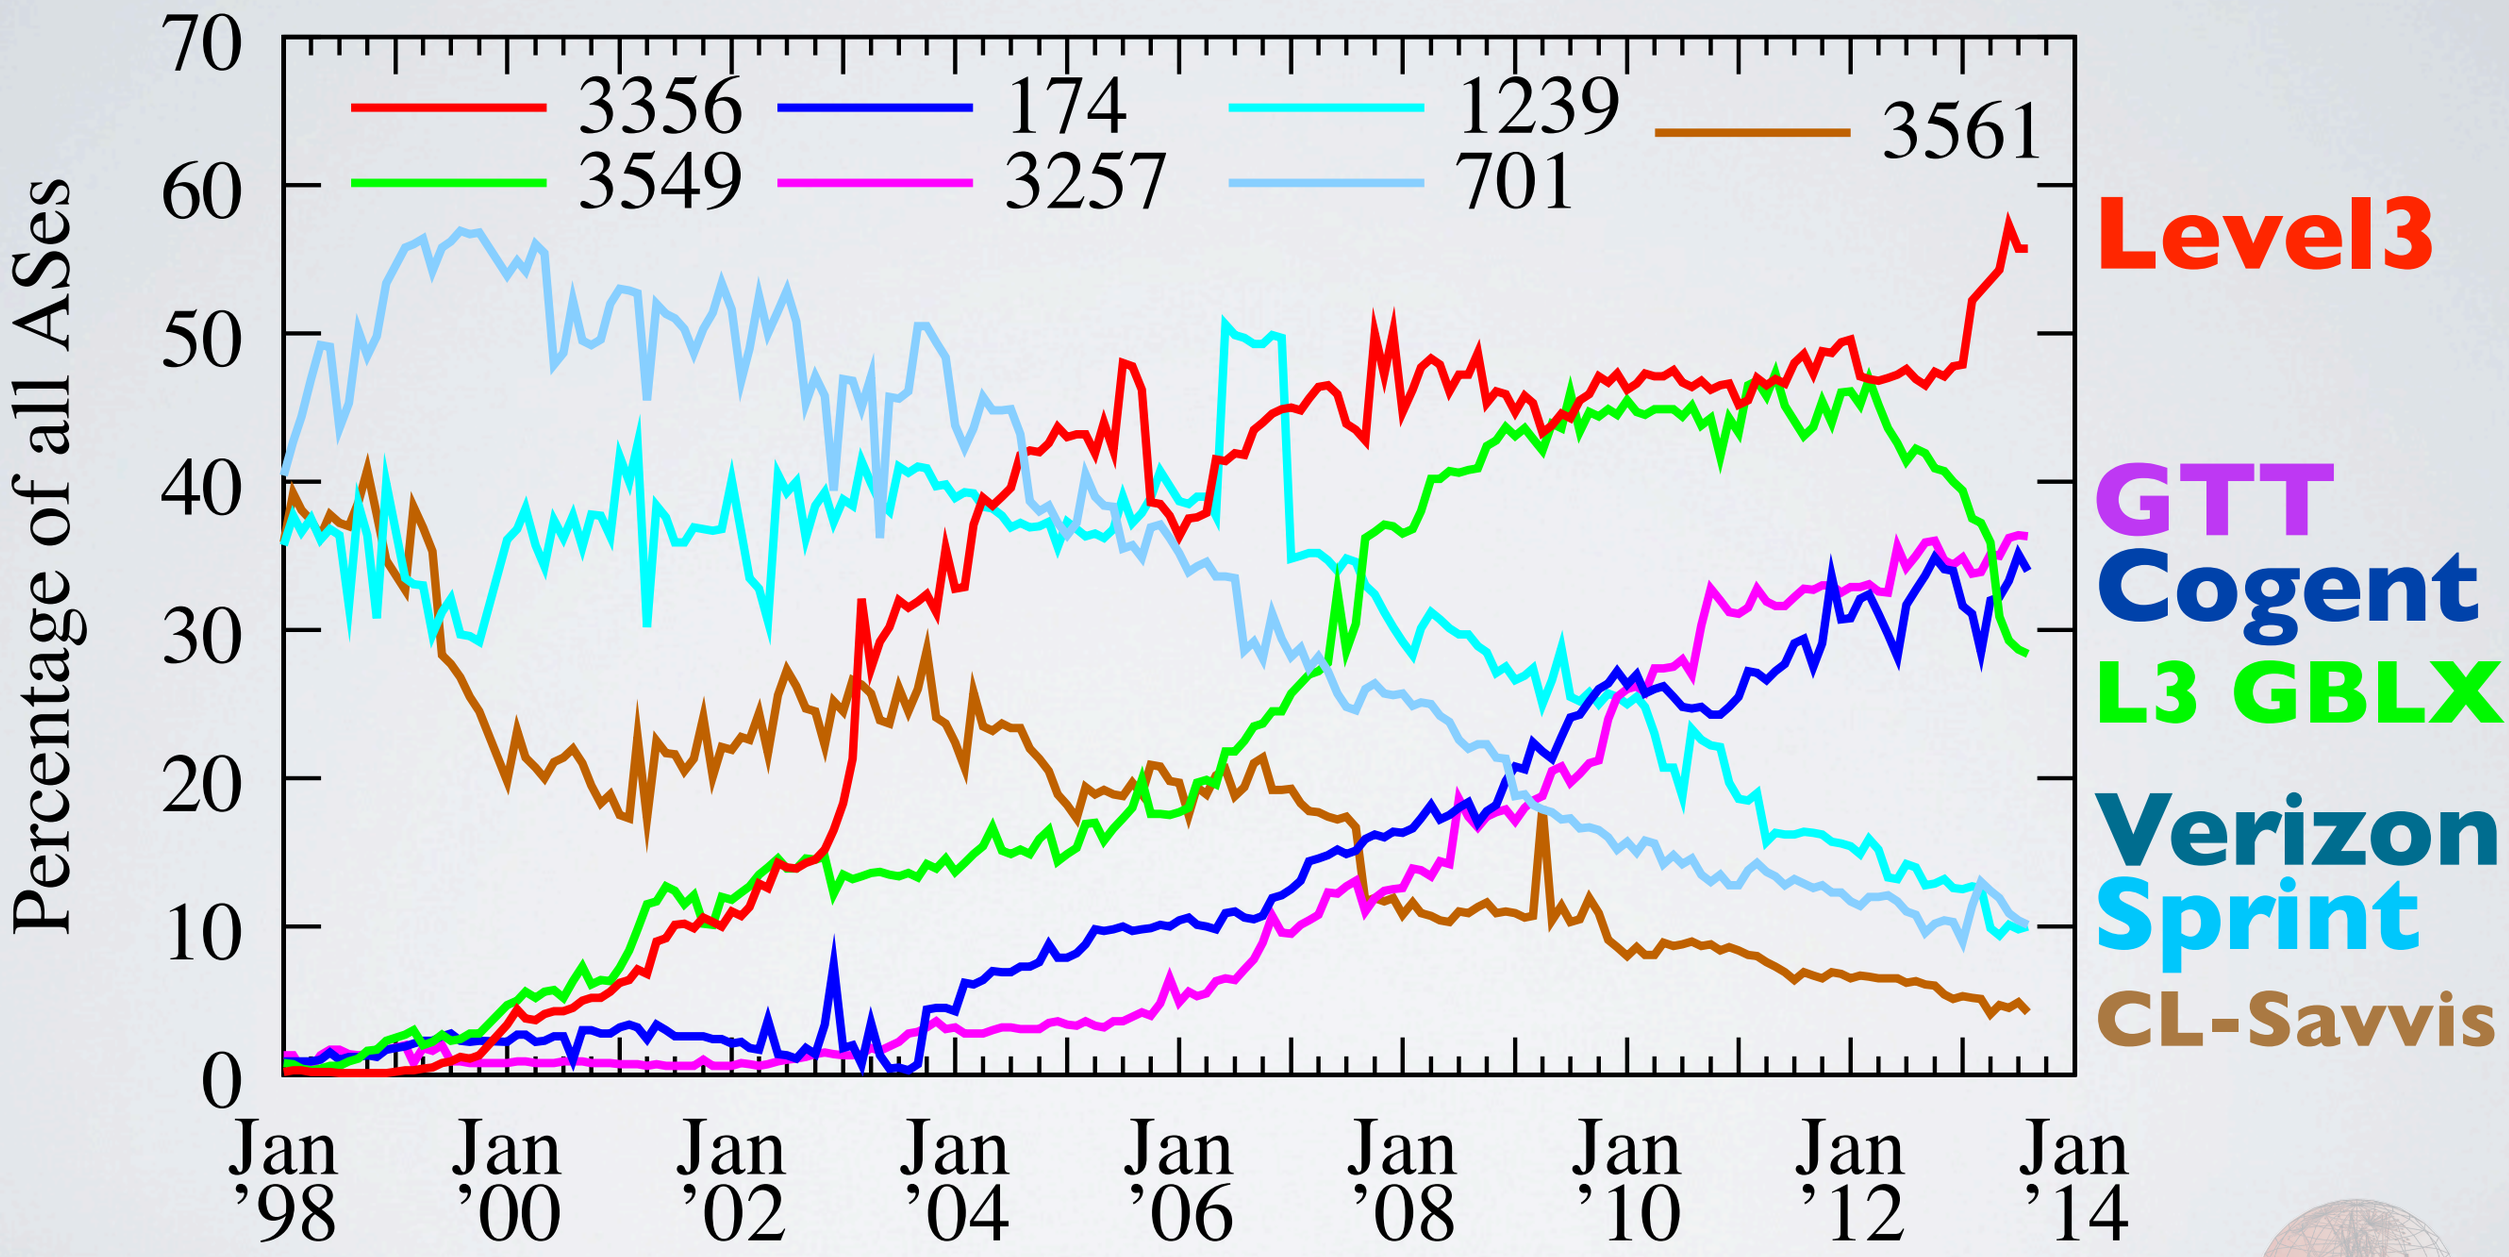

CUSTOMER CONES OVER TIME

(relative, for 3 largest ASes at any time over 15 years)

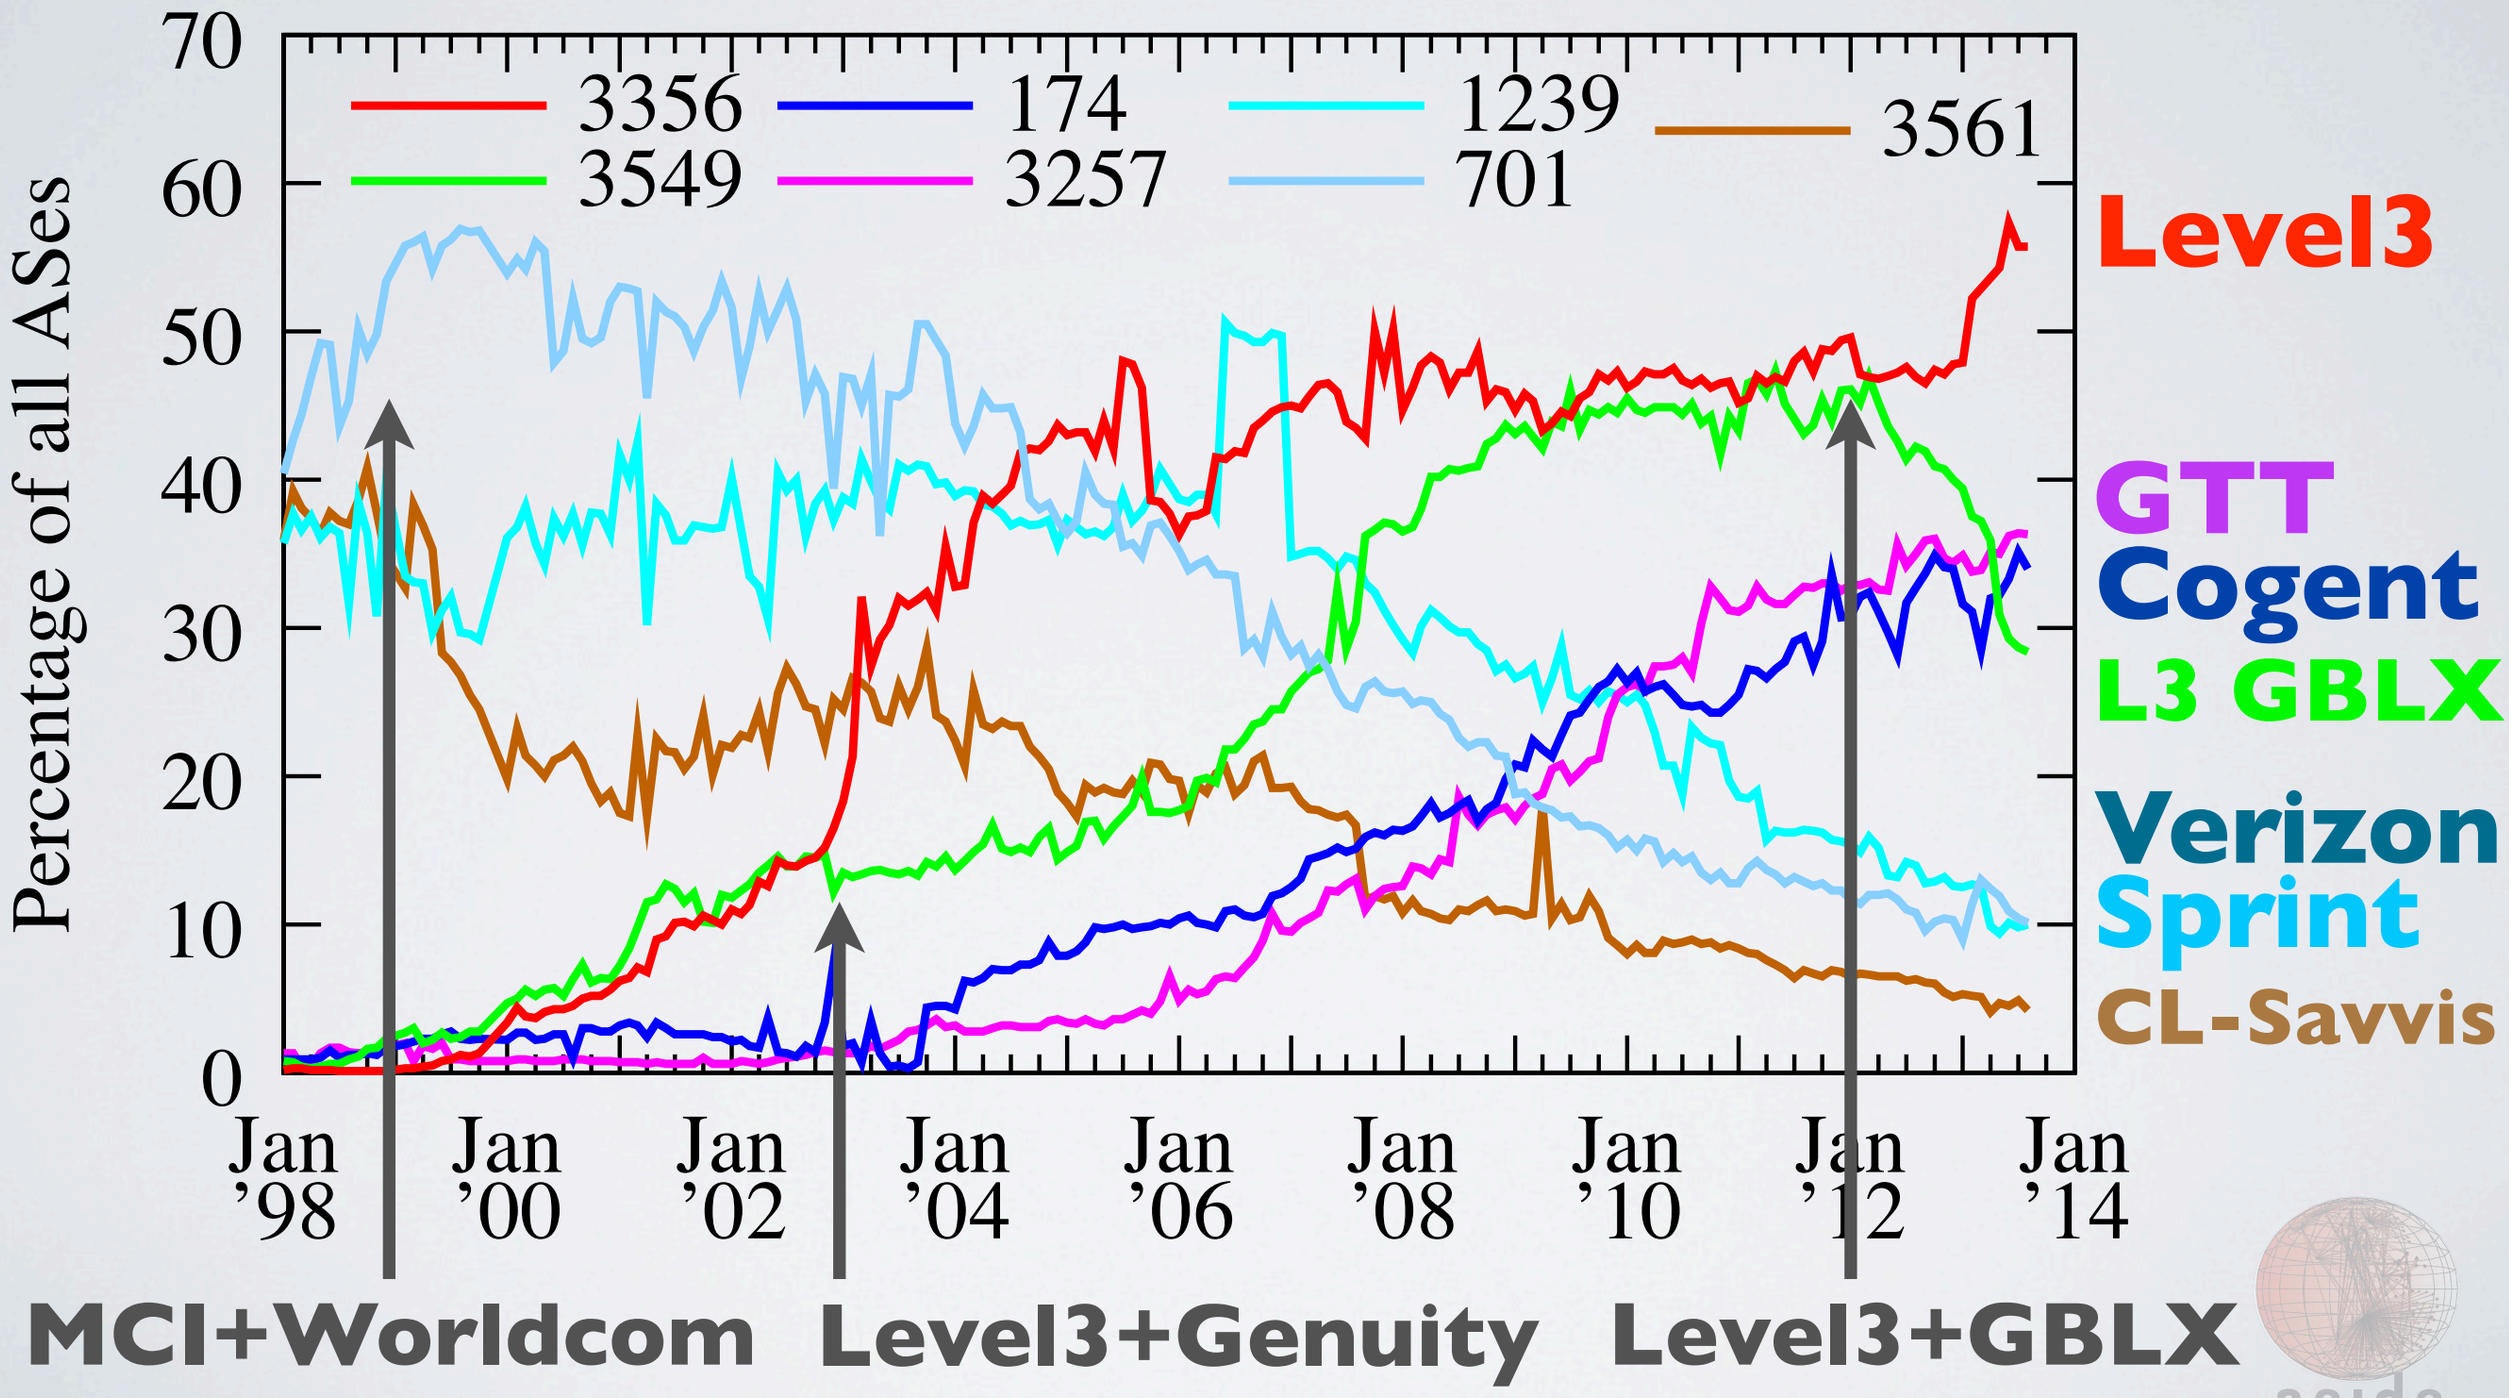

caida

CUSTOMER CONE RELEVANCE

(fraction of ASes in X's cone reached via X from an AS in X)

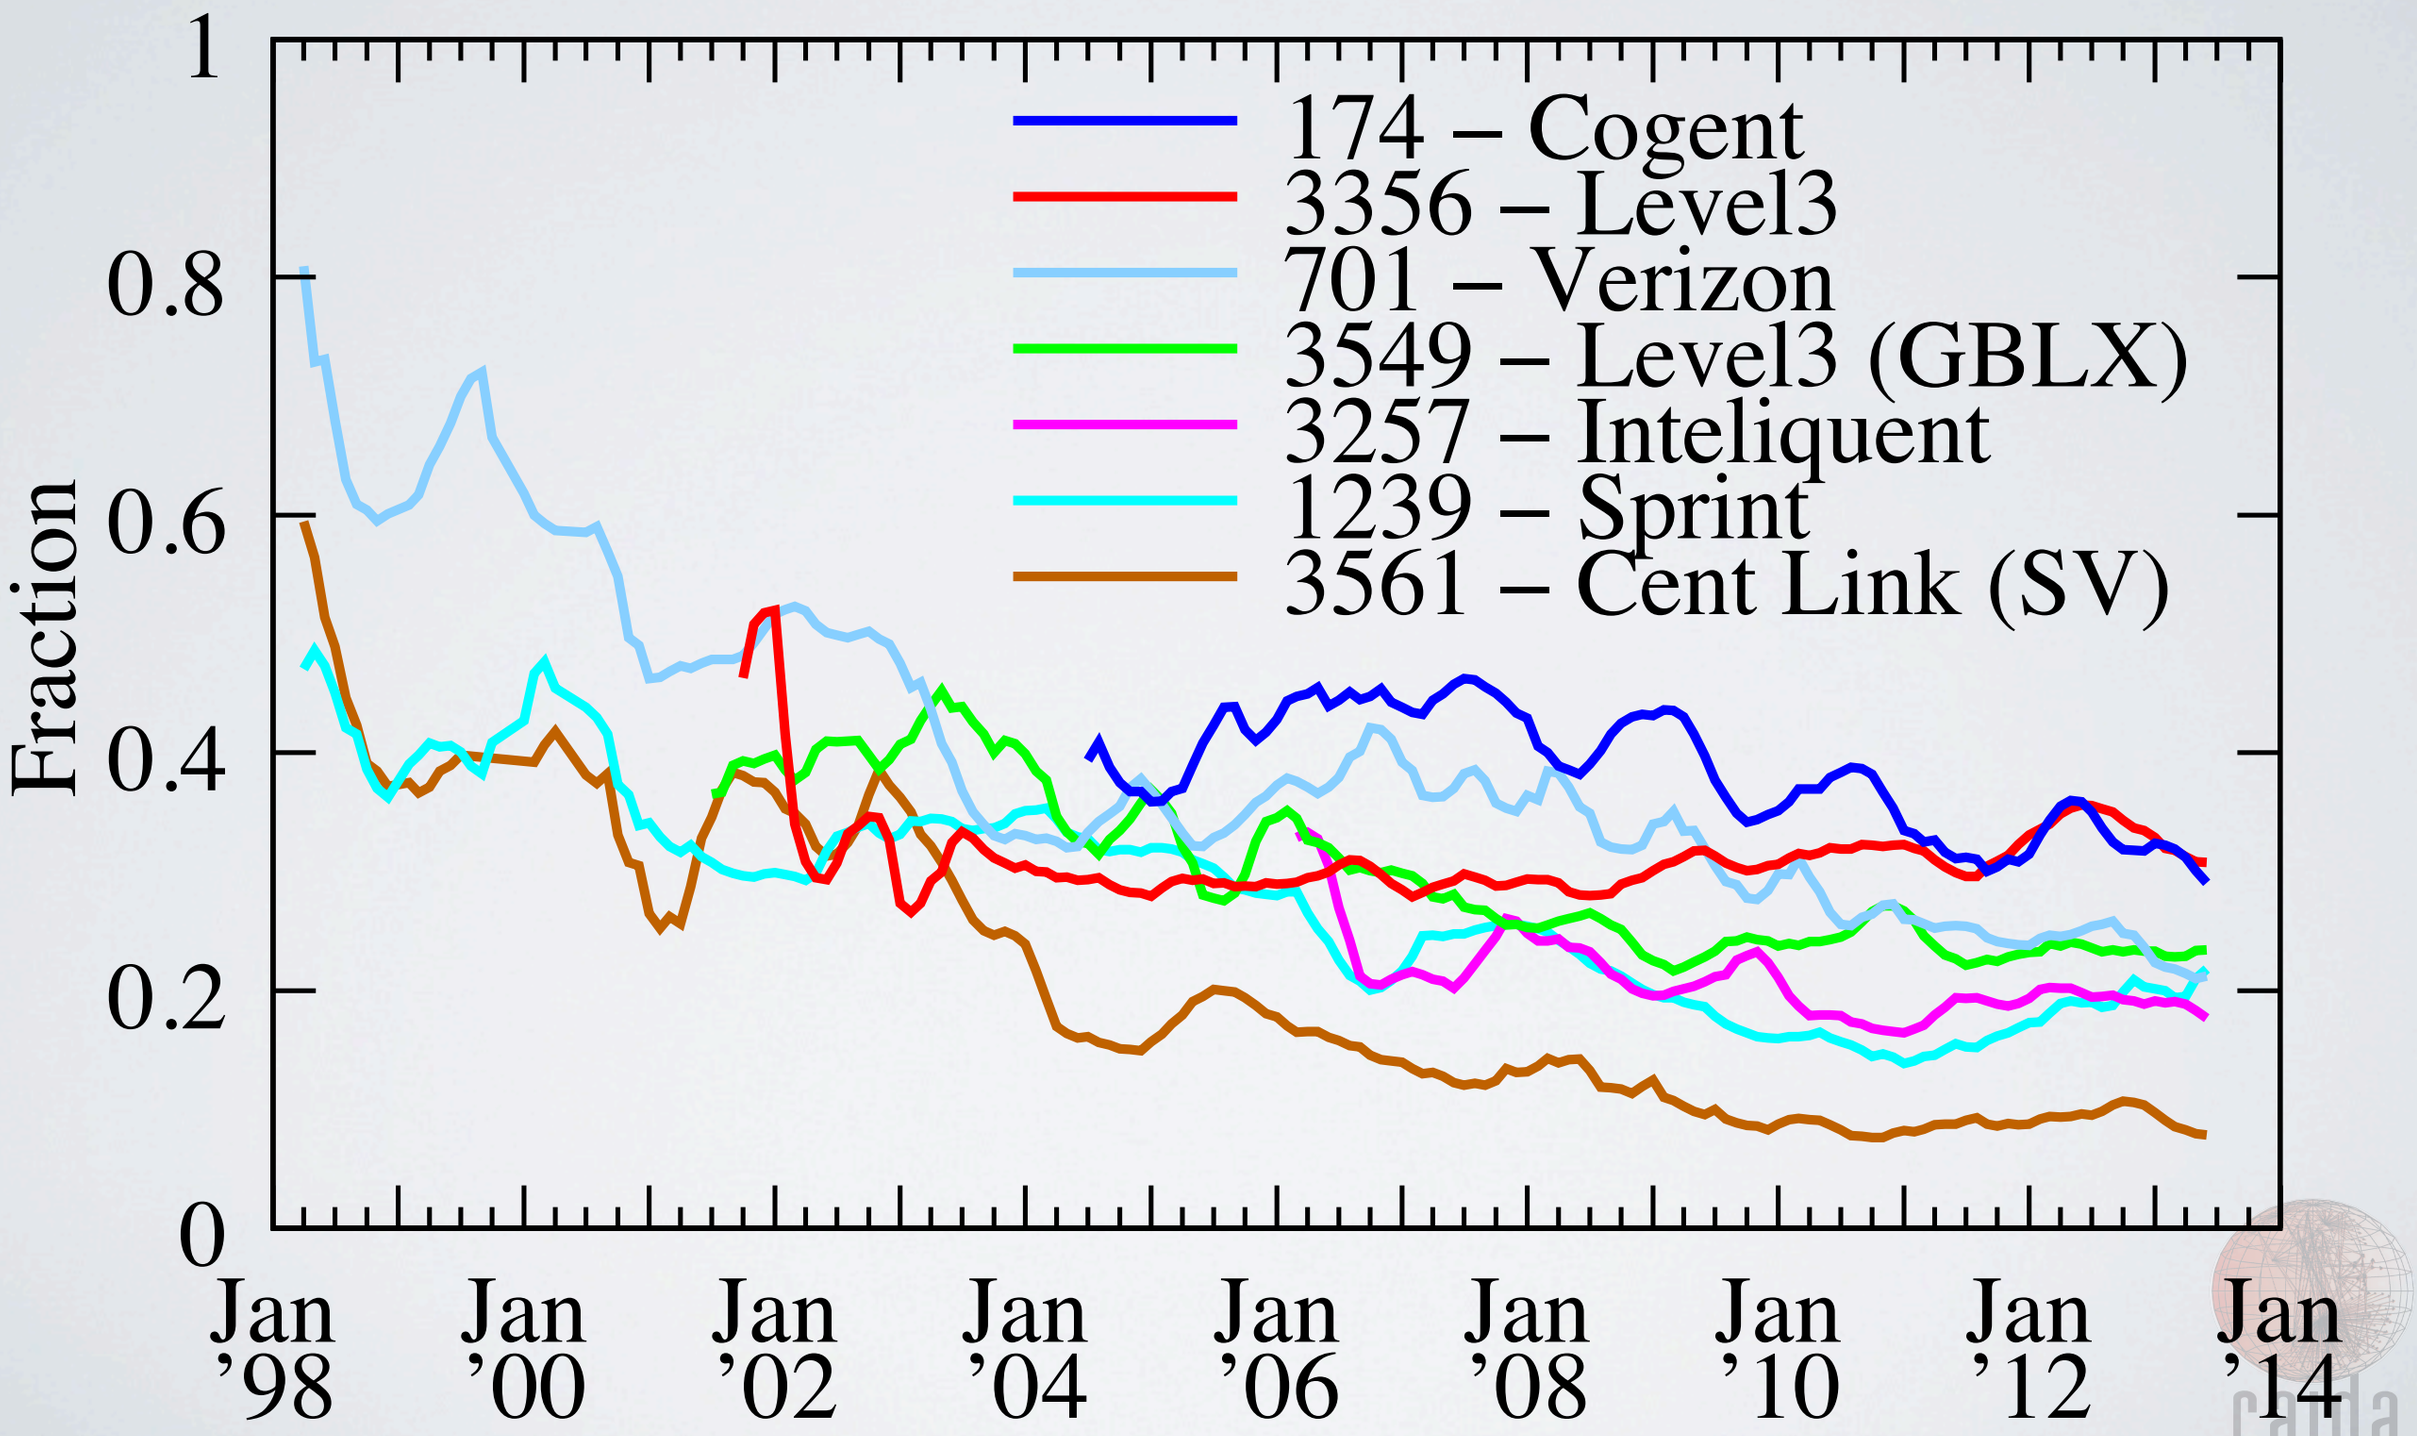

FLATTENING

(characteristics of nodes or links at top of paths)

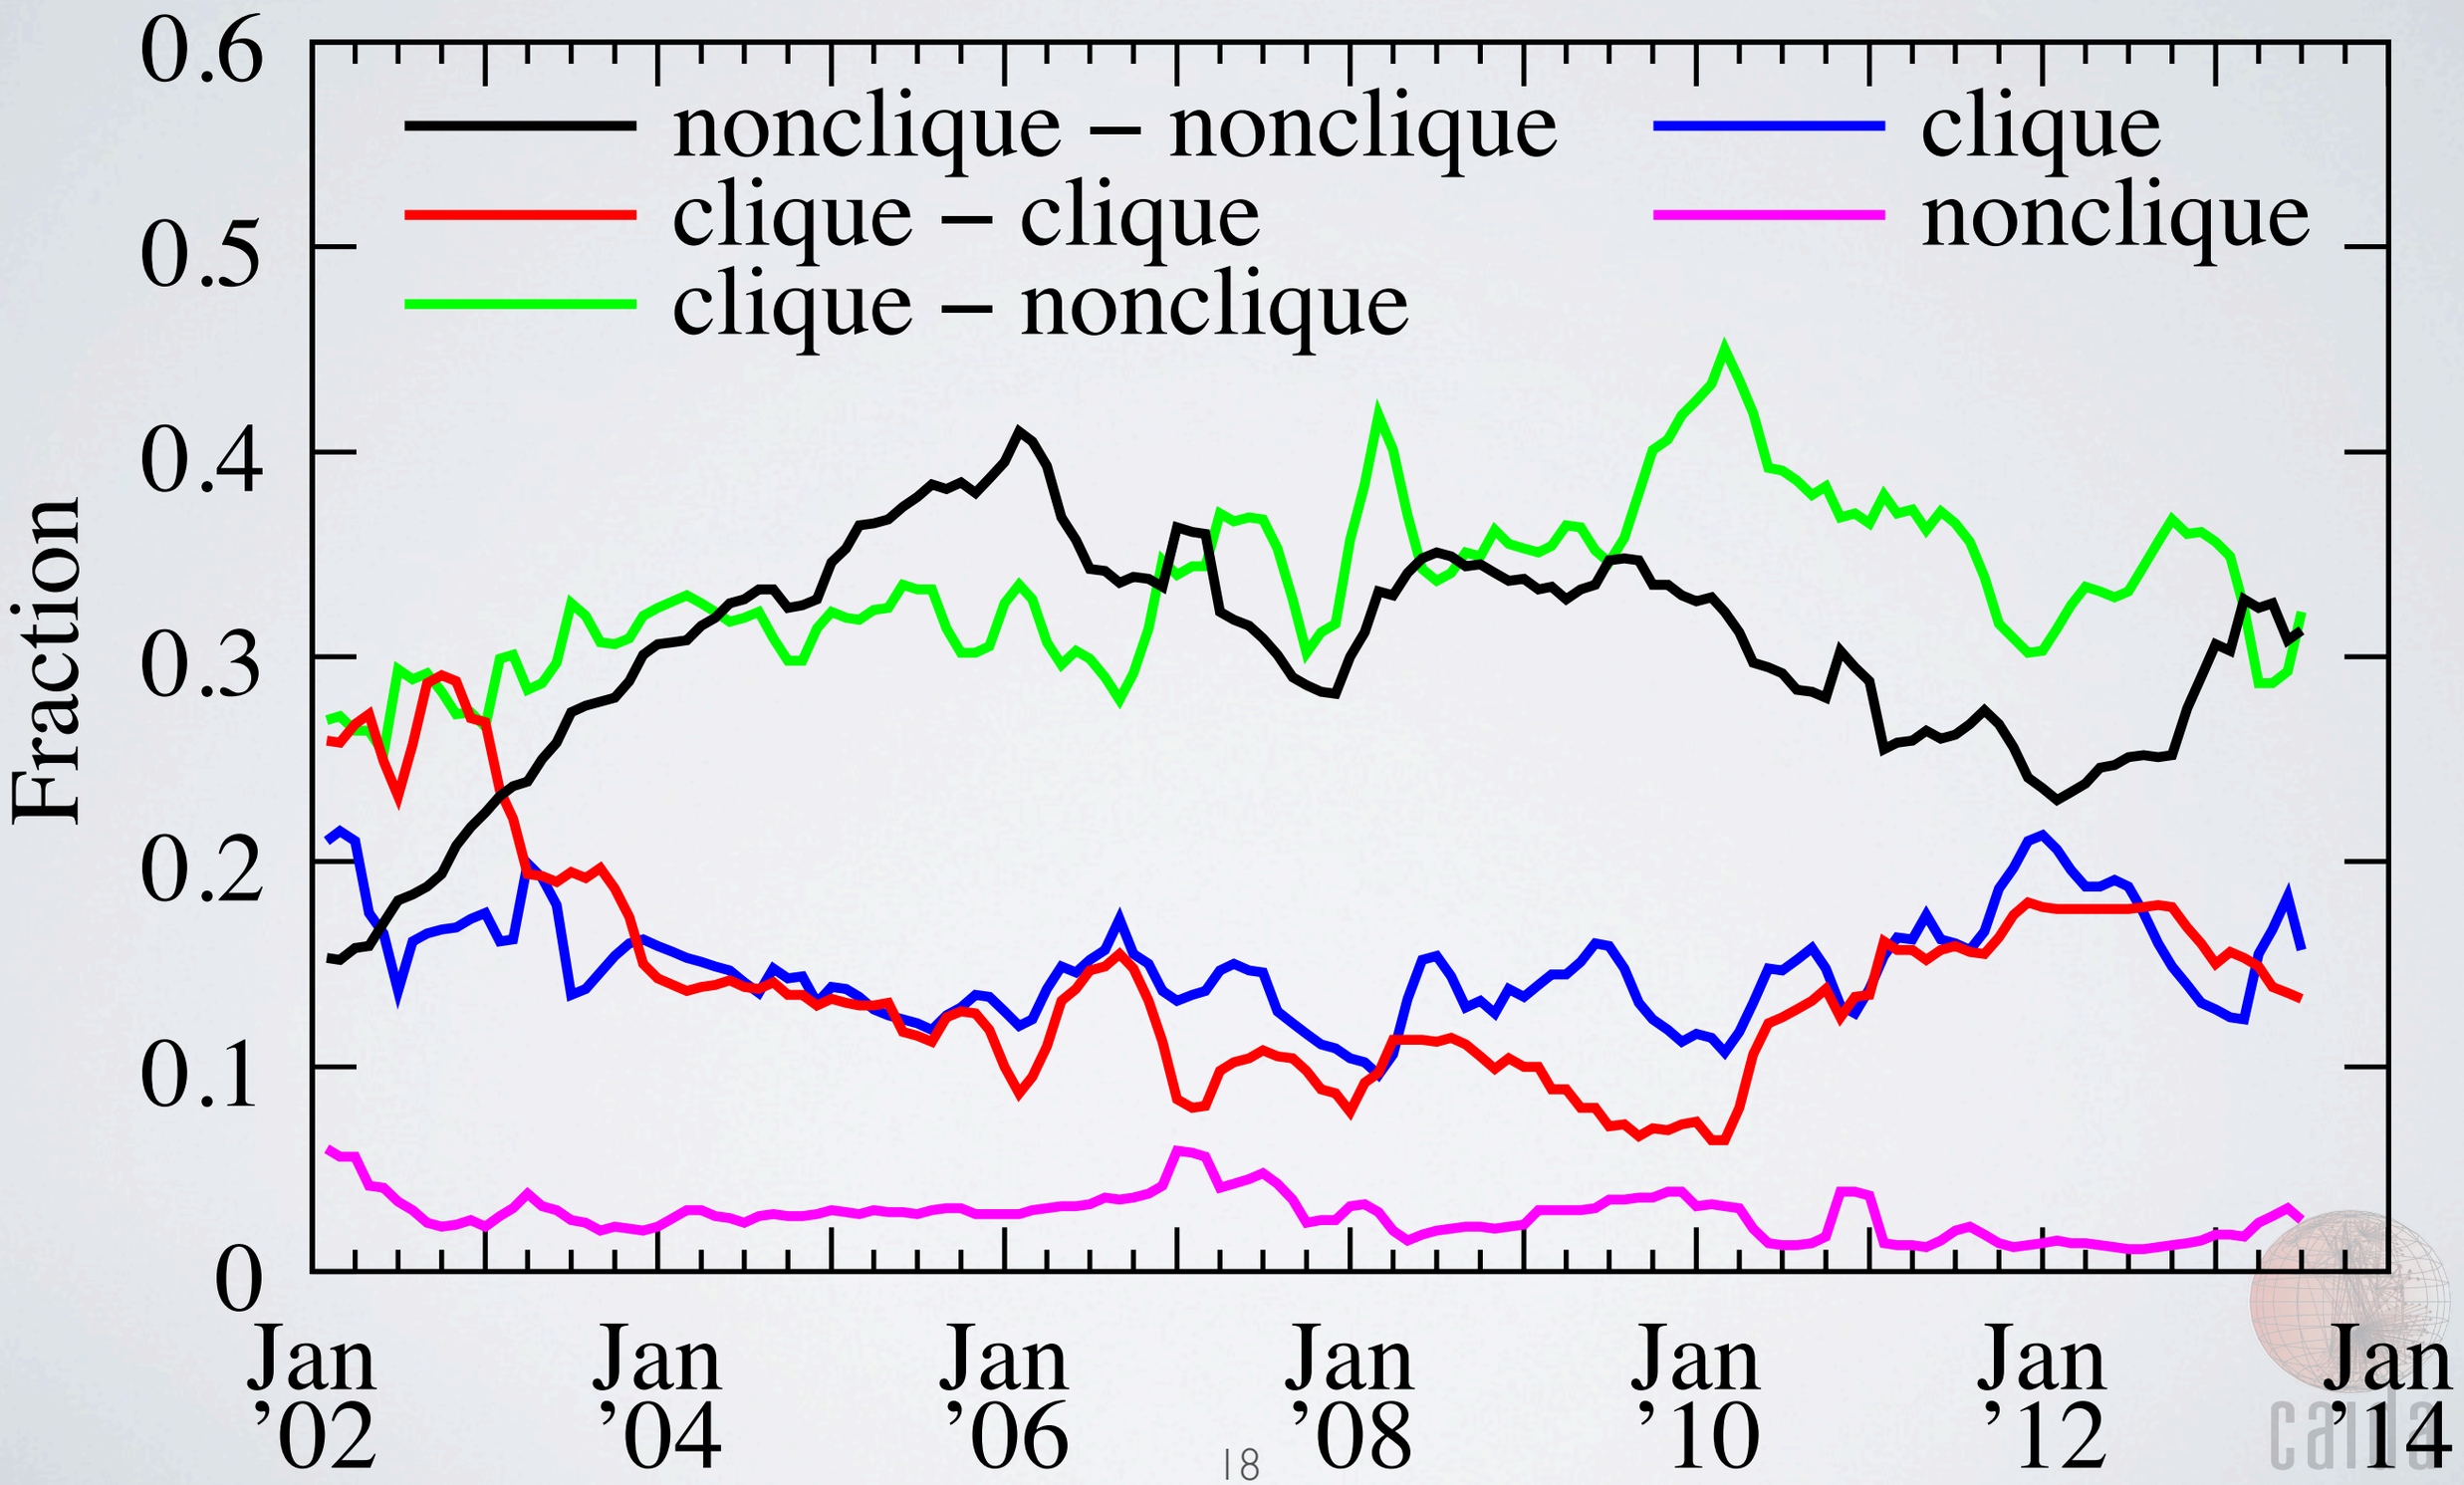

CONTRIBUTIONS

- We build a new **AS relationships inference** algorithm with near-perfect accuracy

- We develop a new **customer cone inference** algorithm to address real-world complexities

- We **validate** our AS relationships to an unprecedented level

  - **99.6%** p2c, **98.7%** p2p, **34.7%** of **126,082** inferences.

- We release our code and 97% of validation data to promote **reproducibility**

`http://www.caida.org/publications/papers/2013/asrank/`

COMPLEX RELATIONSHIPS

(sibling relationships and mutual transit)

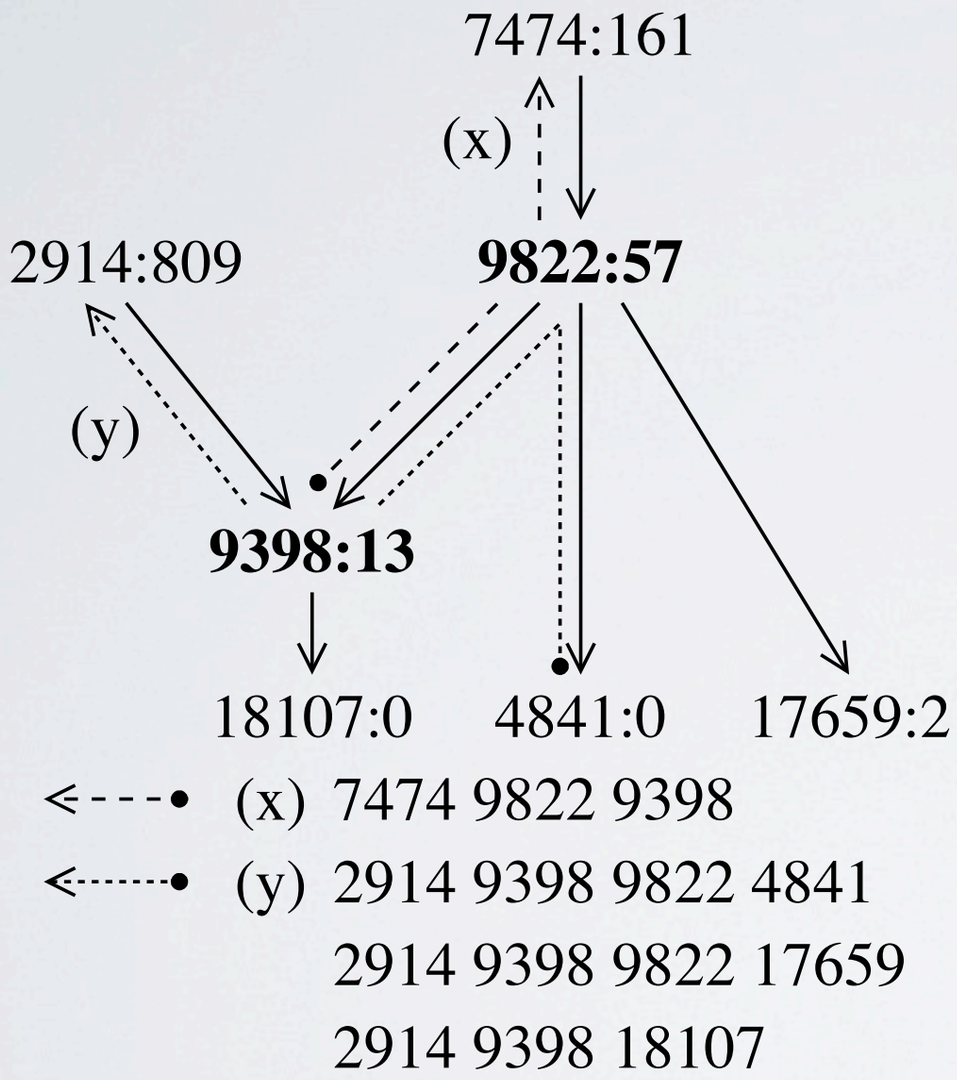

(a) siblings

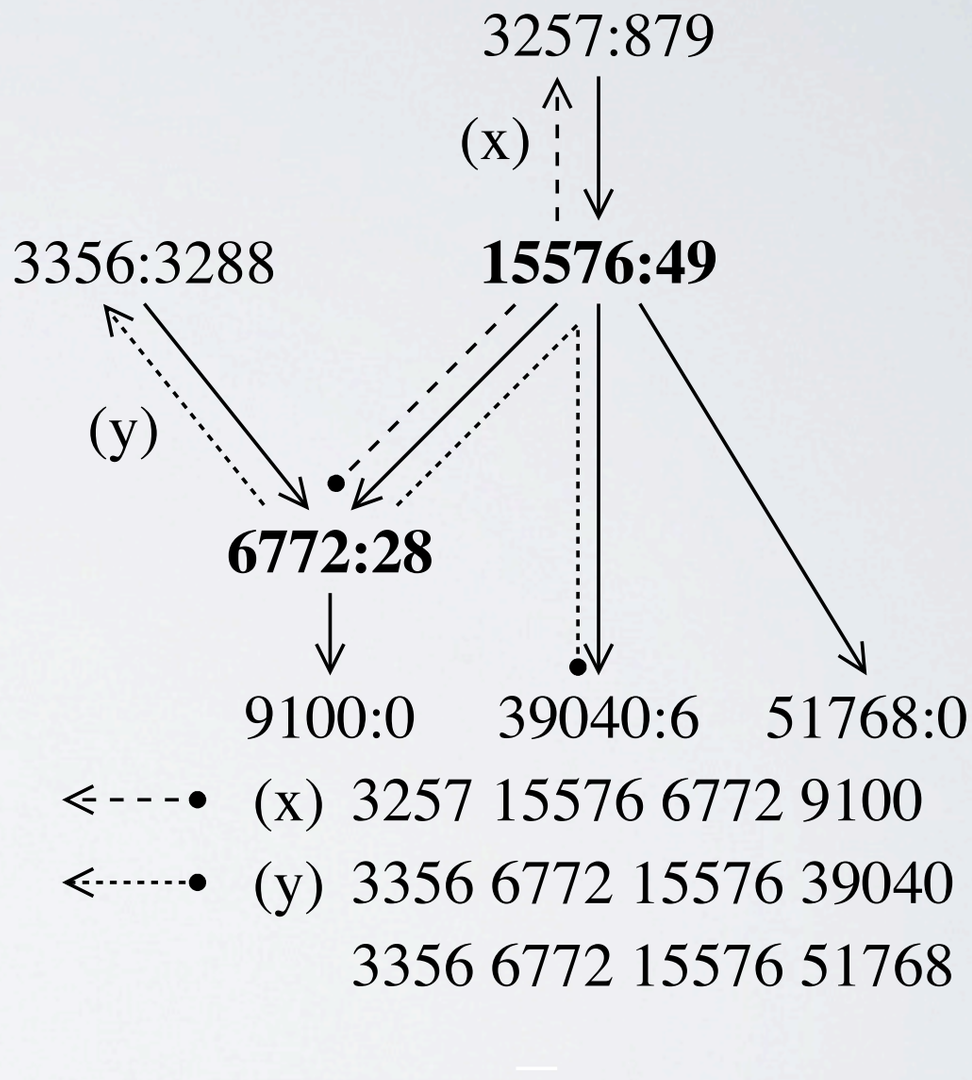

(b) mutual transit

COMPLEX RELATIONSHIPS

(partial transit, hybrid relationships, traffic engineering)

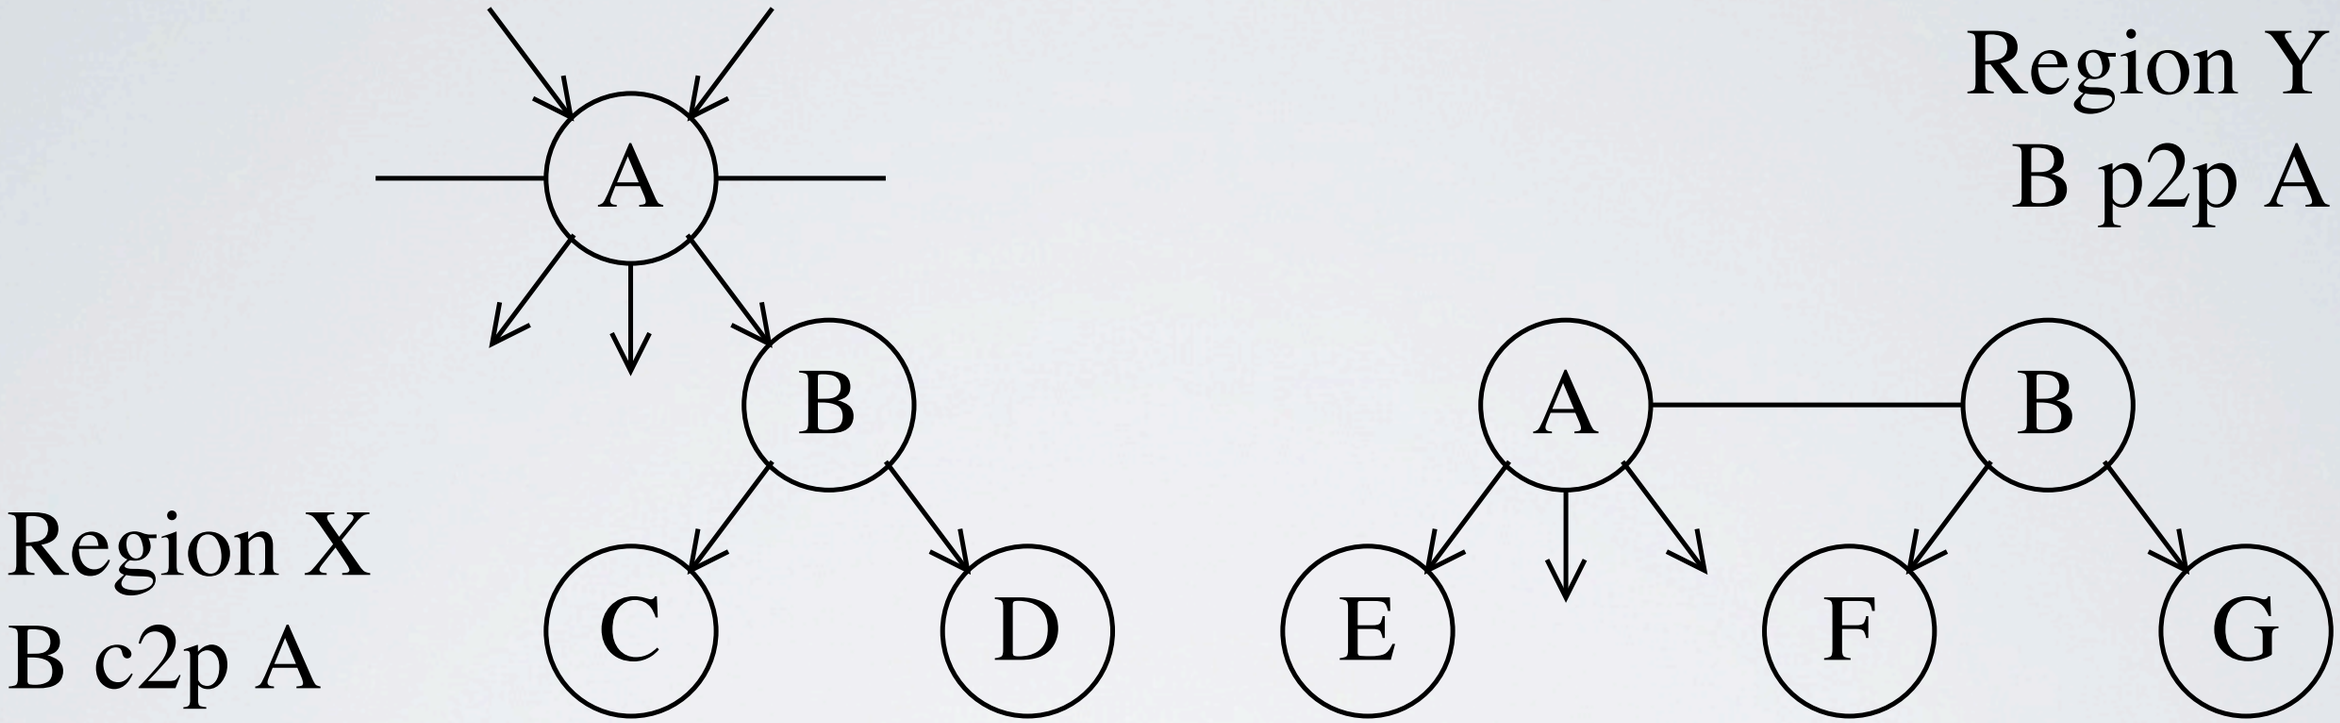

- B is a customer of A in Region X, peers elsewhere
- We assign a single relationship (p2c) because some transit is observed
- If A routes rationally, it will only advertise paths to F and G from B to customers
- We consider hybrid relationships when computing the customer cone

VALIDATION DATA BIAS

(full disclosure)

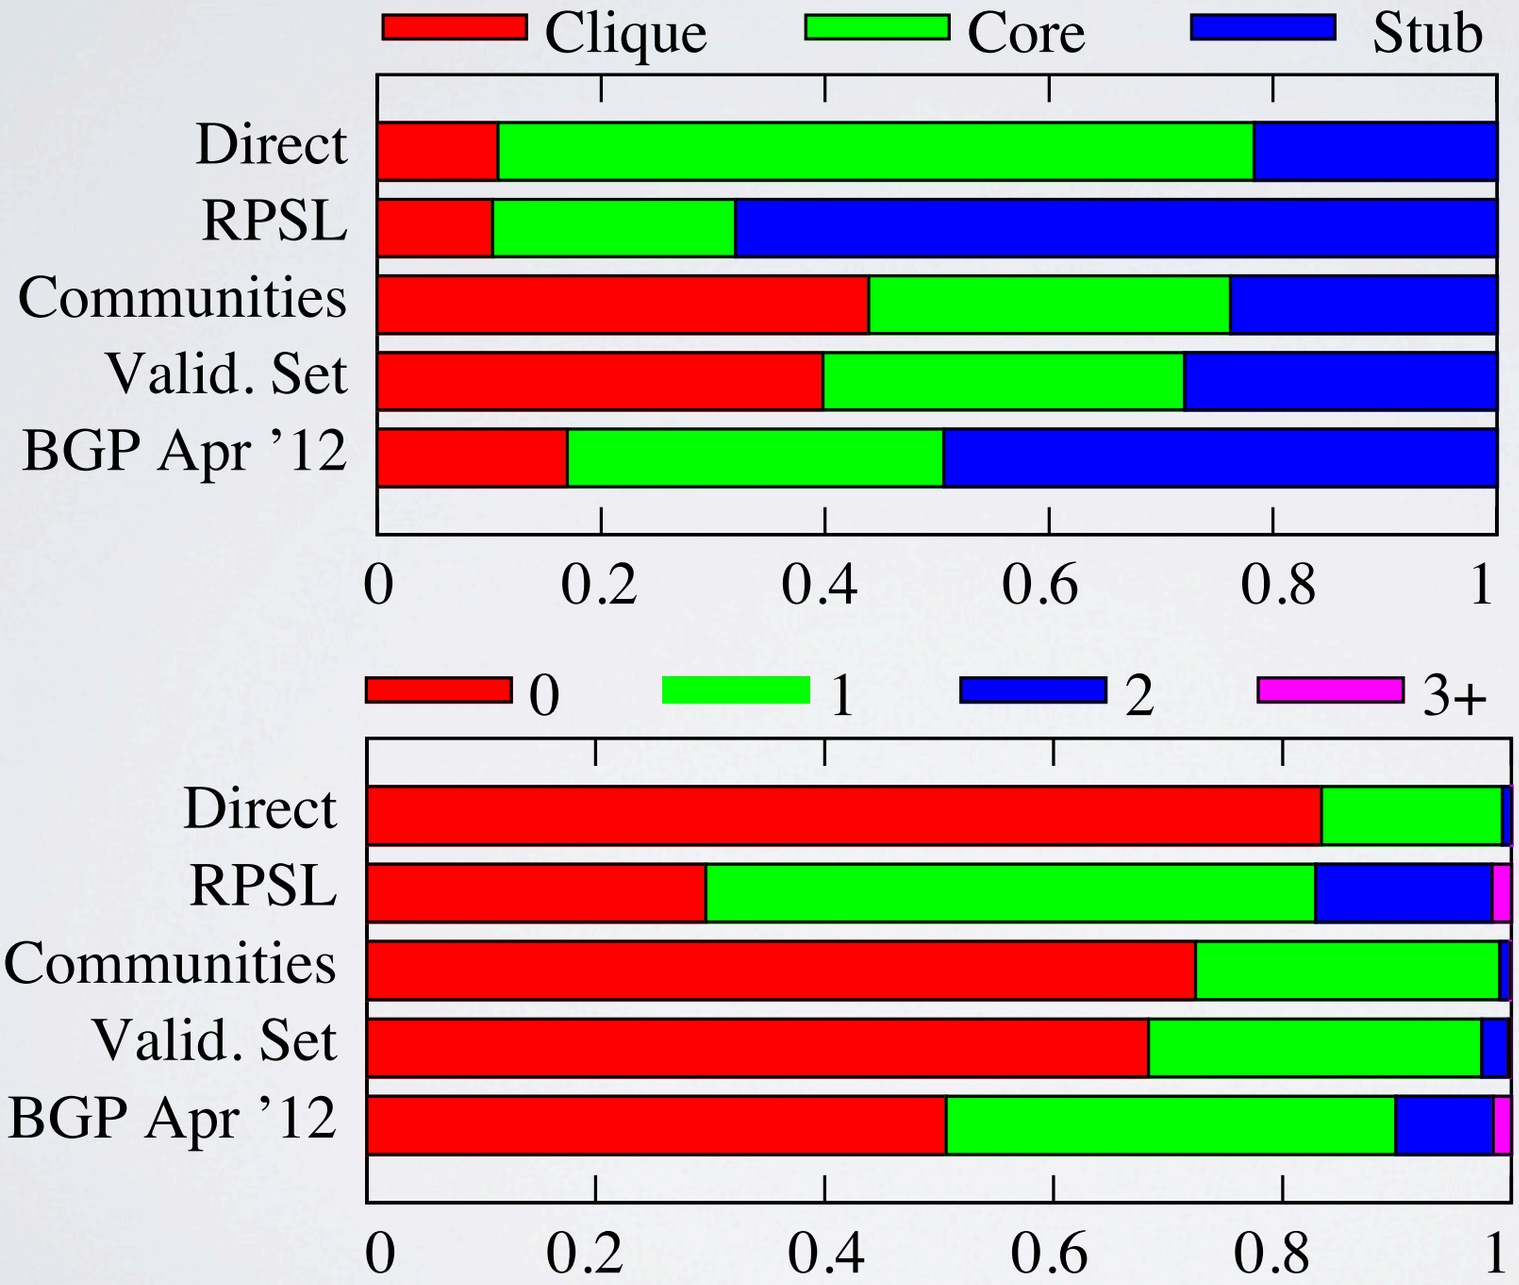

Link classification

Clique links over-represented
Stub links under-represented

Minimum link distance from VP

Directly connected links over-represented, both in directly reported and communities.

VALIDATION AGREEMENT

(99.0% agreement where overlap)

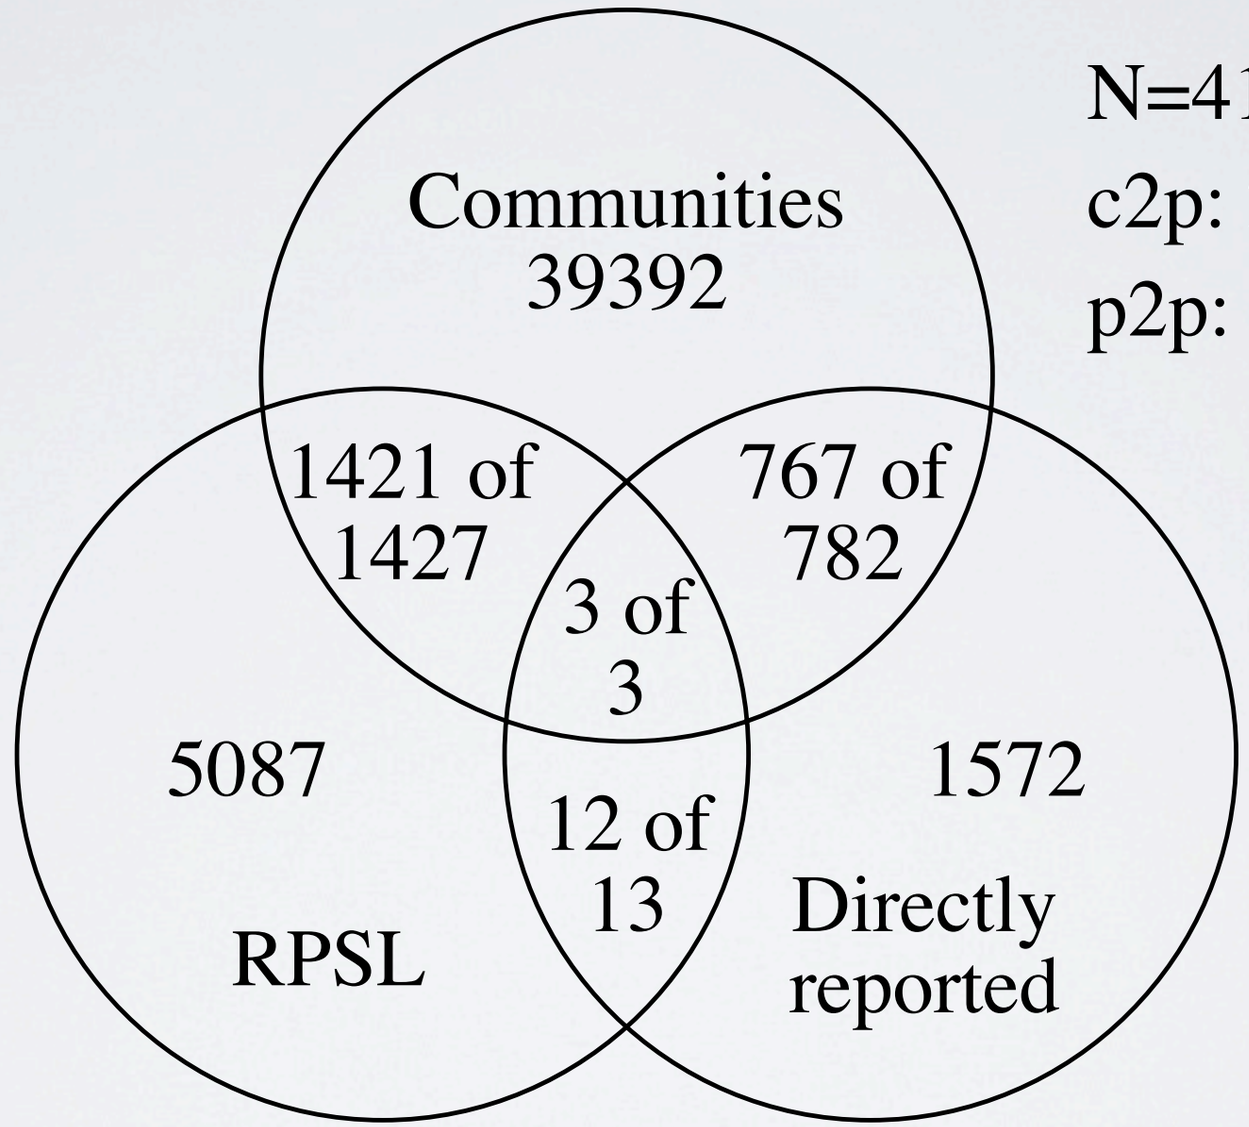

N=41604
c2p: 25356
p2p: 16248

N=6530
c2p: 6530
p2p: 0

N=2370
c2p: 396
p2p: 1974

CLIQUE INFERENCE

1. Apply Bron/Kerbosch 1973 clique detection algorithm to links involving 10 largest ASes by transit degree.
   - clique with largest transit degree sum stored in C_1
2. Test every other AS in order by transit degree to complete clique.
   - Z is added to C_1 provided it does not appear to receive transit from an existing member of C_1.
   - Z is added to C_2 if it would be admitted to C_1 except for a single missing link
3. Apply B/K to links between C_1 and C_2.
   - clique with largest transit degree sum is returned.

caida

RELATED WORKS

- **Initial heuristic techniques**

  - L. Gao. *"On Inferring Autonomous System Relationships in the Internet"*, IEEE/ACM Transactions on Networking, 2001

  - J. Xia, L. Gao. *"On the evaluation of AS relationship inferences"*, IEEE Globecom, 2004

Maximise valley-free paths.
p2c, p2p, s2s
Validation: AT&T

Maximise VF paths.
Seed with GT-classified links from IRR and Communities.
Validation: unused GT

- **ToR formulation**

  - L. Subramanian, S. Agarwal, J. Rexford, R.H. Katz. *"Characterizing the Internet Hierarchy from Multiple Vantage Points"*, IEEE INFOCOM, 2002

  - G.D. Battista, T. Erlebach, A. Hall, M. Patrignani, M. Pizzonia, T. Schank. *"Computing the Types of the Relationships Between Autonomous Systems"*, IEEE/ACM Transactions on Networking, 2007

Formalized Gao.
Conjecture ToR NP-complete.
SARK: rank ASes by closeness to core.
p2c, p2p.
Validation: fraction of VF paths

Proved ToR NP-complete.
Infer p2c, leave s2s+p2p inferences as future work

caida

RELATED WORKS

- **Recent heuristic techniques**

  - B. Zhang, R. Liu, D. Massey, L. Zhang. *"Collecting the AS-level Topology"*, ACM/SIGCOMM CCR, 2005

  - X. Dimitropoulos, D. Krioukov, M. Fomenkov, B. Huffaker, Y. Hyun, k claffy. *"AS Relationships: Inference and Validation"*, ACM/SIGCOMM CCR, 2007

  - Y. Shavitt, E. Shir, U. Weinsberg. "*Near-Deterministic Inference of AS Relationships*", ConTel, 2009

  - E. Gregori, A. Improta, L. Lenzini, L. Rossi, L. Sani. "*BGP and Inter-AS Economic Relationships*", IFIP Networking, 2011

UCLA IRL
Infer clique. Links observed by clique are p2c. All others p2p. No validation.

MAX-2-SAT: NP-hard. Find solution that maximises (1) fraction of VF paths, (2) provider deg. > customer deg. Siblings from WHOIS. Validation: 9.7% of inferences

Requires traceroute. No validation.

Similar to UCLA IRL. Uses lifetime of paths. No validation.