UNIVERSITÀ DEGLI STUDI ROMA TRE
Dipartimento di Informatica e Automazione

Algorithms for the Inference of the Commercial Relationships between Autonomous Systems: Results Analysis and Model Validation

Massimo Rimondini
Maurizio Pizzonia
Giuseppe Di Battista
Maurizio Patrignani

Summary

- Introduction
- Motivations
- Methodology
- Tools
- Experimental Results
- Conclusions

Introduction

- Autonomous System
  - (very) large set of network devices
    - under the control of a single administration
    - using coherent routing policies
- ASes exchange routing information using the Border Gateway Protocol
  - BGP-speaking routers establish *peering sessions*
  - the spread of BGP reachability messages can be limited using *policies*
- Internet model
  - interconnection of ASes

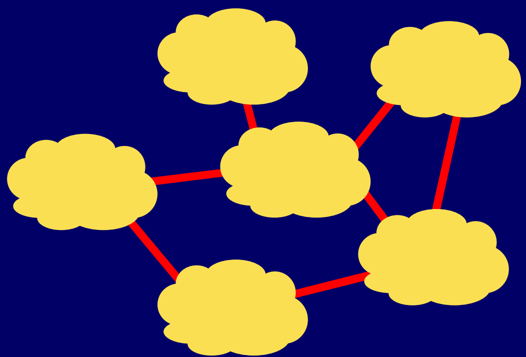

Commercial Relationships

- Contracts which are subscribed in order to obtain connectivity
- Implemented by using BGP policies [1]

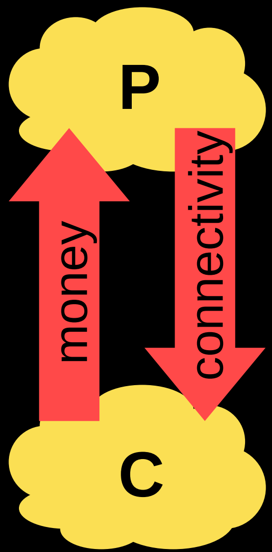

Customer-Provider

C (customer) pays P (provider) to obtain connectivity to the rest of Internet

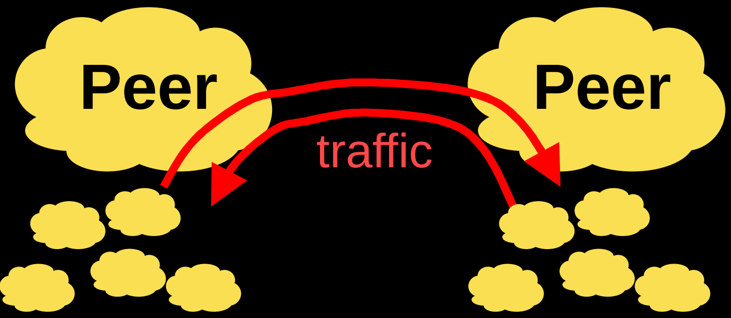

Peer-Peer

The two peers agree to exchange traffic between their customers

[1] Lixin Gao. On Inferring Autonomous System Relationships in the Internet. *IEEE/ACM Transactions on Networking*, 9(6):733–745, Dec 2001.

Commercial Relationships

- Contracts which are subscribed in order to obtain connectivity
- Implemented by using BGP policies [1]

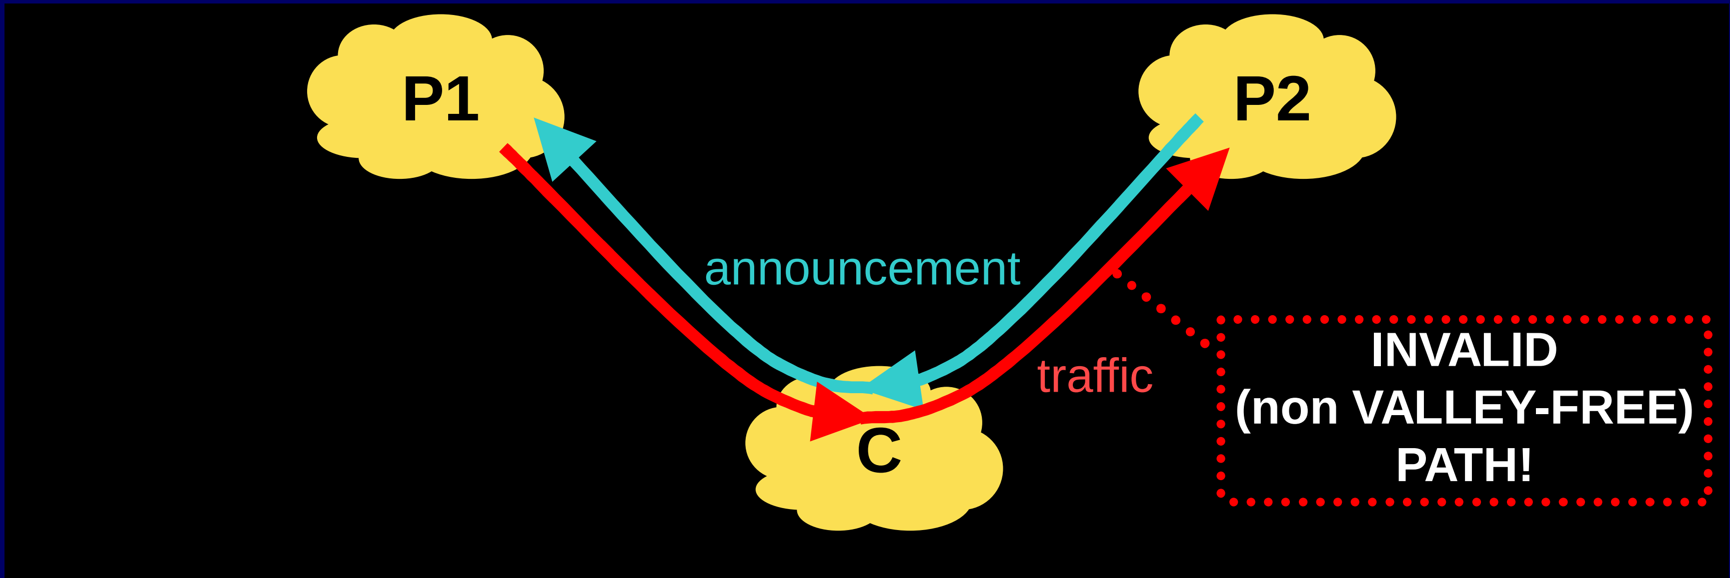

[1] Lixin Gao. On Inferring Autonomous System Relationships in the Internet. *IEEE/ACM Transactions on Networking*, 9(6):733–745, Dec 2001.

Why using Inference Algorithms

- Knowing the commercial relationships is useful, but difficult
- Asking AS authorities for the commercial relationships they establish is unfeasible
  - too numerous
  - not willing to reveal sensible information
- Routing registries may be incomplete/out of date

Inference Algorithms

- Inference algorithms take as input BGP routing data
- Produce as output a relationship assignment
- Network model: AS graph
  - vertices ⇔ ASes
  - edges ⇒ BGP peering sessions

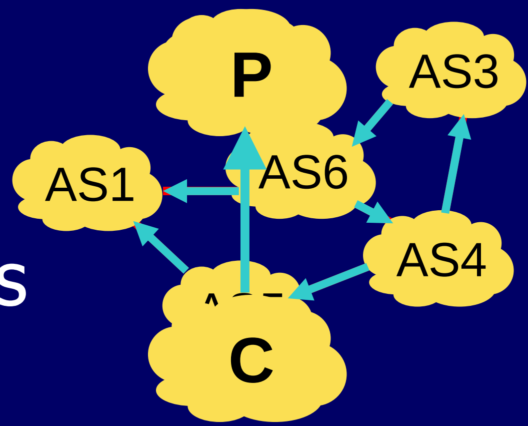

- Relationship assignment

  ↕

  graph orientation

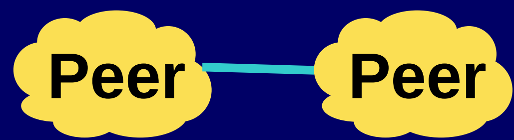

Inference Algorithms

- State of the art
  - L. Gao. On Inferring Autonomous System Relationships in the Internet. IEEE/ACM Transactions on Networking, 9(6):733–745, Dec 2001.

    degree-based

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

    ranking, multiple observation points

  - DPP — G. Di Battista, M. Patrignani, and M. Pizzonia. Computing the Types of the Relationships between Autonomous Systems. In *Proc. IEEE INFOCOM 2003.*

    satisfiability of propositional formulae, no peerings

Summary

- Introduction
- Motivations
- Methodology
- Tools
- Experimental Results
- Conclusions

Motivations

- Results Analysis
  - do inference algorithms produce realistic results?
  - stability analysis
    - evaluates the impact of routing oscillations on relationship assignments
  - independence from the algorithm
    - investigates whether different algorithms produce similar results

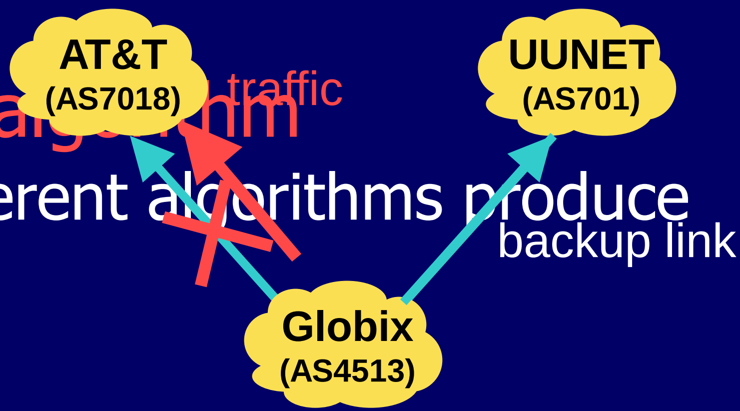

- Model Validation
  - good results $\Rightarrow$ sound models
- Towards a unified methodology

Methodology

- Possible approach:
  - considering assignments as being realistic depending on the ability to enforce the valley-free property
- Drawback:

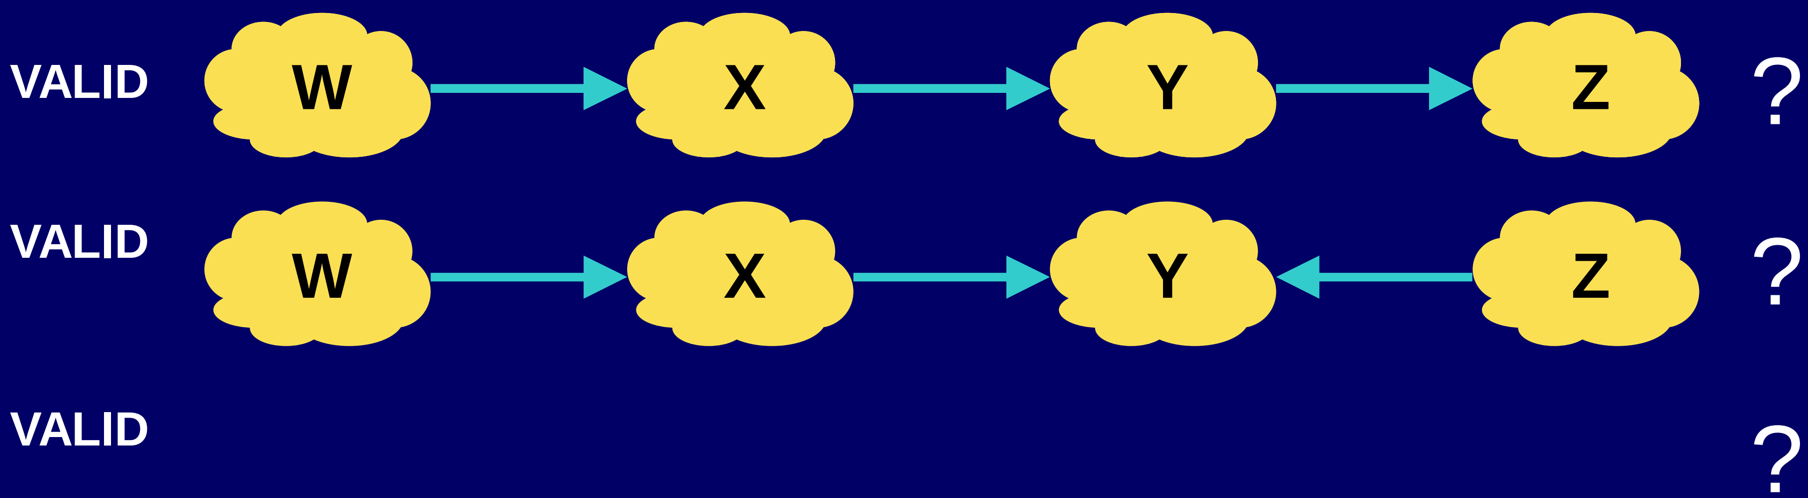

Methodology

- Another possibility:
  - validating assignments against real world data [1]
- Drawback:
  - retrieving real world information is not easy
  - not scalable

[1] Lixin Gao. On Inferring Autonomous System Relationships in the Internet. *IEEE/ACM Transactions on Networking*, 9(6):733–745, Dec 2001.

Methodology

- Our approach:
  1. identify (other) qualities of a reasonably good inference
  2. define measures which evaluate such qualities
  3. run the inference algorithms extensively on several input data sets
  4. compute the measures on the relationship assignments that have been produced
  5. draw conclusions

Methodology: Qualities

- Stability
  - several time consecutive snapshots
  - single algorithm

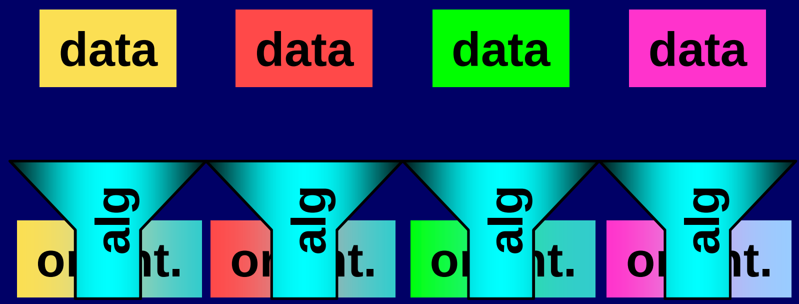

- Independence from the algorithm
  - single snapshot
  - different algorithms

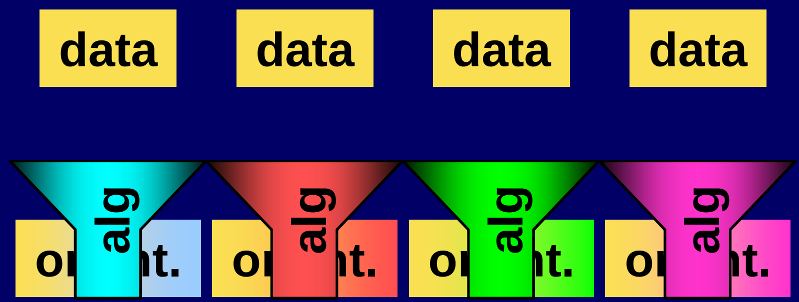

Methodology: Measures

- Undirected AS graph: $G(V, E)$
- Relationship assignment on *G*:

$$R : E \rightarrow V \cup \{peering, unknown\}$$

$$R(e) = \begin{cases} AS_2 & \text{if } AS_2 \text{ is a provider of } AS_1 \\ AS_1 & \text{if } AS_1 \text{ is a provider of } AS_2 \\ peering & \text{if } AS_1 \text{ and } AS_2 \text{ are peers} \\ unknown & \text{if no relationship is known} \end{cases}$$

Methodology: Measures

- Two relationship assignments, $R_A(\cdot)$ and $R_B(\cdot)$, defined on the same graph

$$Both = \{e \in E | R_A(e) \neq unknown \wedge R_B(e) \neq unknown\}$$
$$OnlyInA = \{e \in E | R_A(e) \neq unknown \wedge R_B(e) = unknown\}$$
$$OnlyInB = \{e \in E | R_A(e) = unknown \wedge R_B(e) \neq unknown\}$$

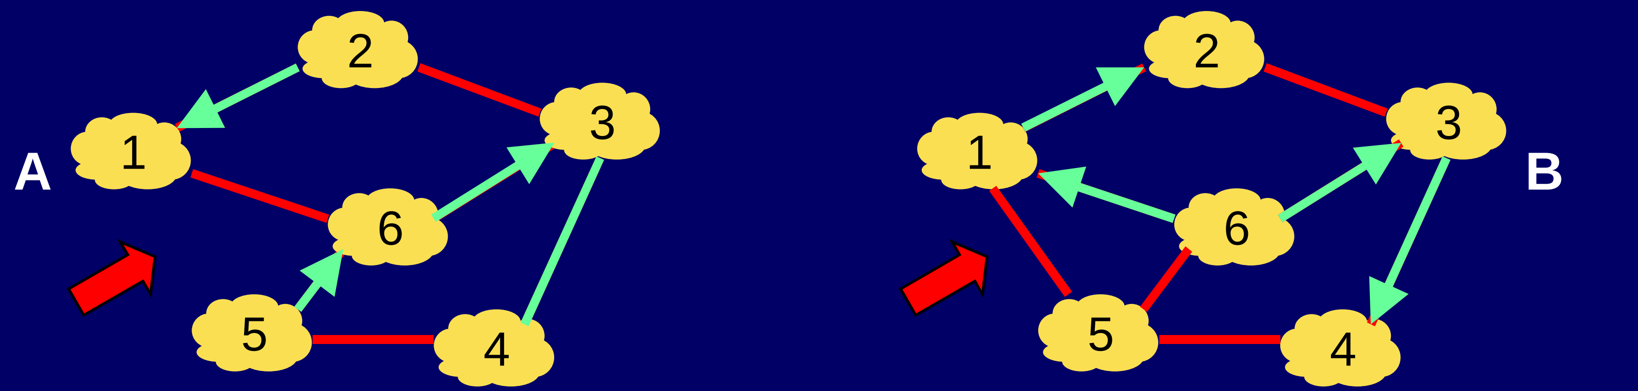

Methodology: Measures

- Two relationship assignments, $R_A(\cdot)$ and $R_B(\cdot)$, defined on the same graph

$$Both = \{ e \in E | R_A(e) \neq unknown \wedge R_B(e) \neq unknown \}$$

$$OnlyInA = \{ e \in E | R_A(e) \neq unknown \wedge R_B(e) = unknown \}$$

$$OnlyInB = \{ e \in E | R_A(e) = unknown \wedge R_B(e) \neq unknown \}$$

Methodology: Measures

$$Both = \{ e \in E \mid R_A(e) \neq unknown \wedge R_B(e) \neq unknown \}$$

$$Consistent = \{ e \in Both \mid R_A(e) = R_B(e) \}$$

$$Inconsistent = \{ e \in Both \mid R_A(e) \neq R_B(e) \}$$

Methodology: Measures

$$Inconsistent = \{ e \in Both \mid R_A(e) \neq R_B(e) \}$$

$$Opposite = \{ e \in Inconsistent \mid R_A(e) \neq peer \wedge R_B(e) \neq peer \}$$

$$PeerInA = \{ e \in Inconsistent \mid R_A(e) = peer \wedge R_B(e) \neq peer \}$$

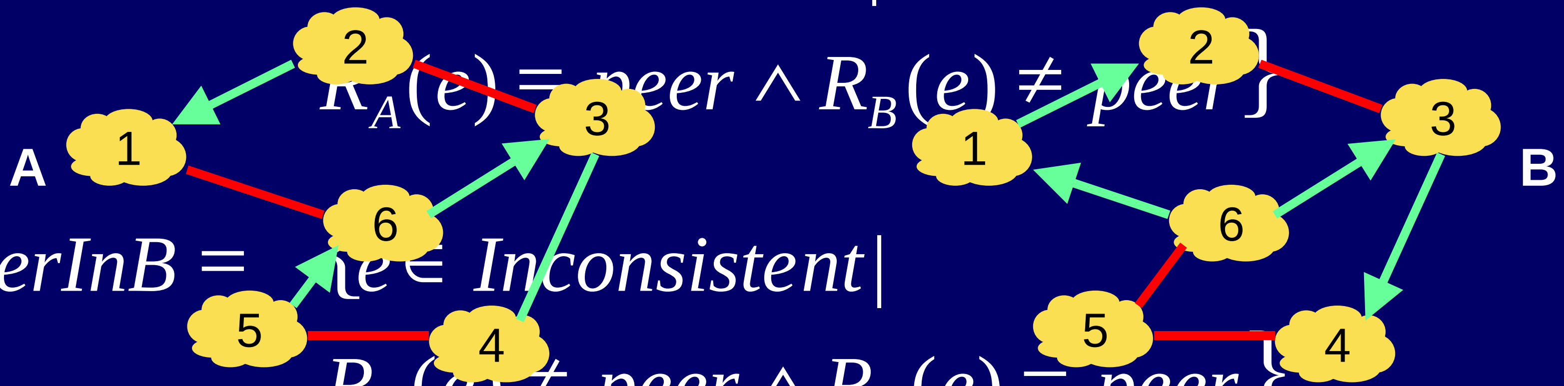

$$PeerInB = \{ e \in Inconsistent \mid R_A(e) \neq peer \wedge R_B(e) = peer \}$$

Methodology: Measures

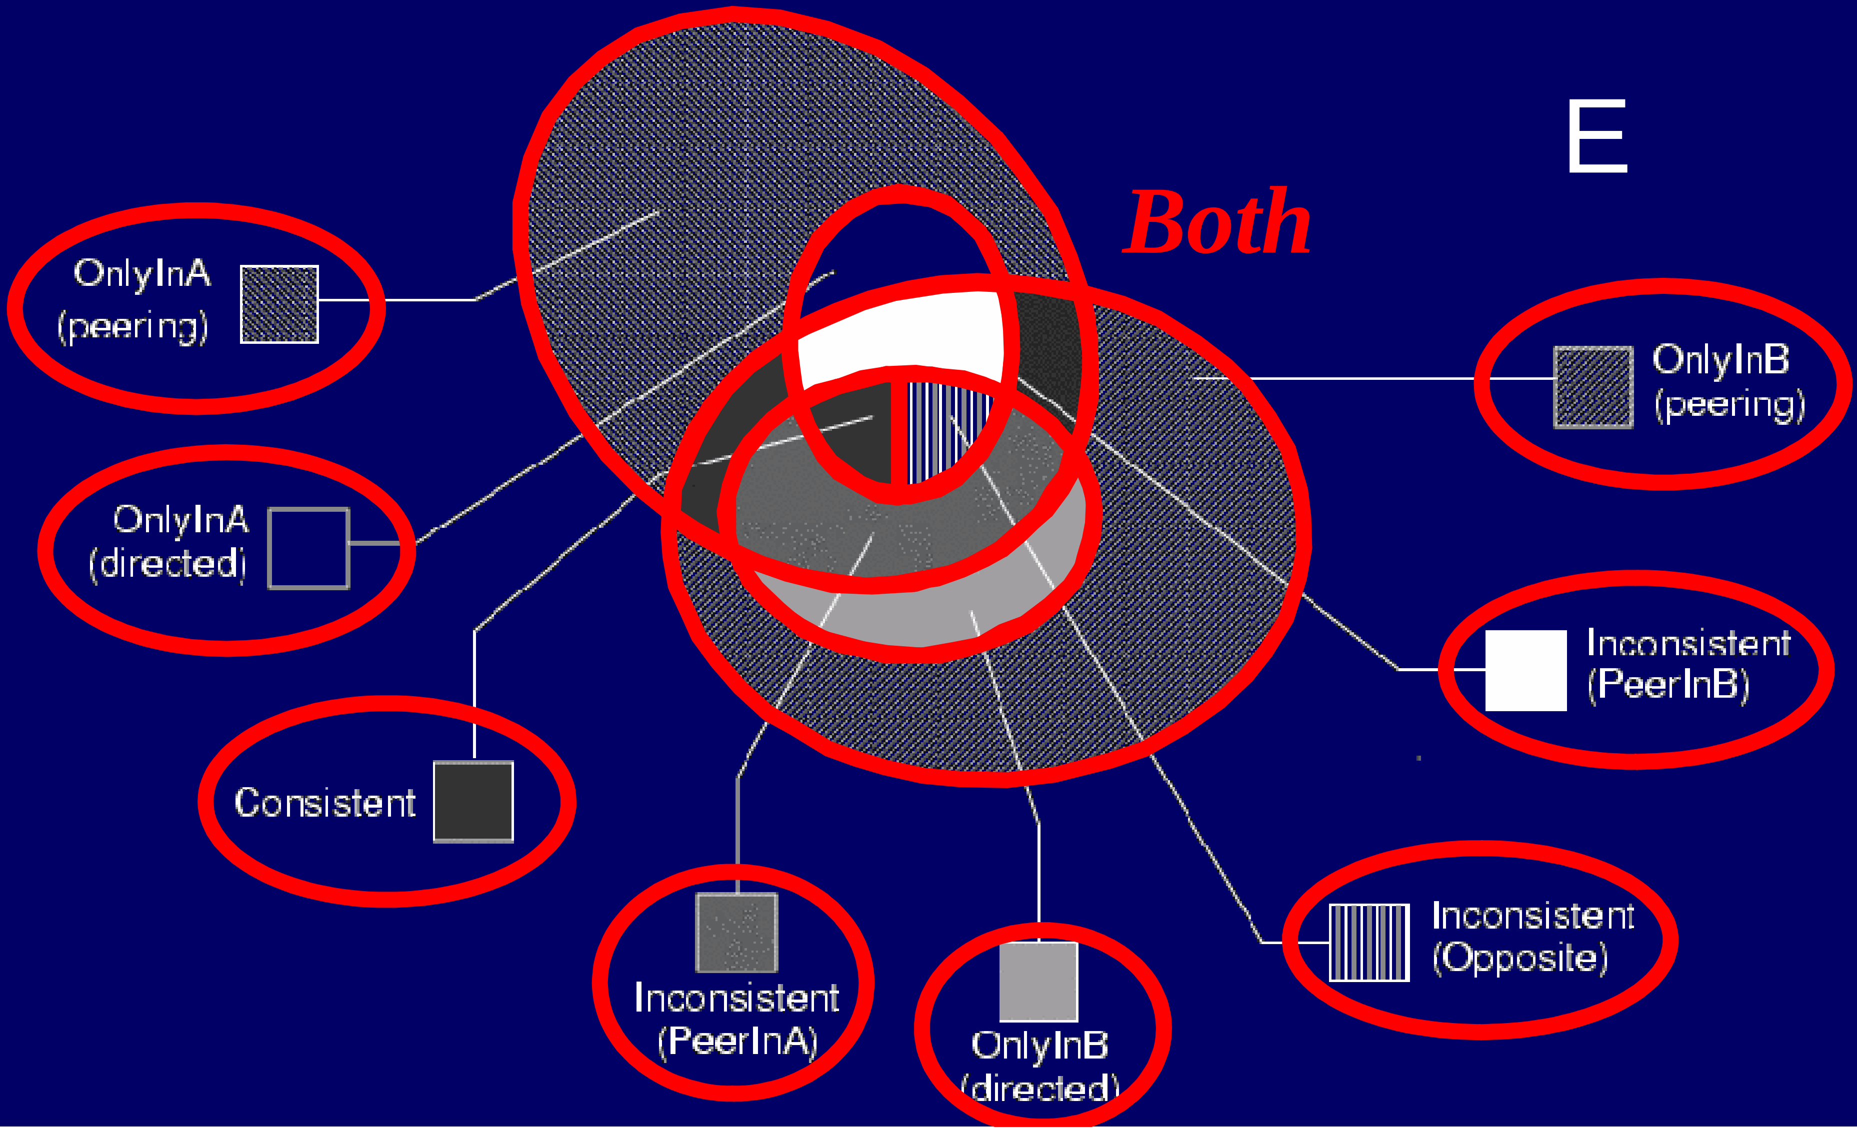

Methodology: "Historical" Measures

- Sequence of AS graphs $G_1(V_1, E_1), \ldots, G_n(V_n, E_n)$
- $\tilde{G} = (\tilde{V}, \tilde{E})$
  - $\tilde{E} = \bigcup_{i=1}^{n} E_i$
  - $\tilde{V}$ is naturally induced by $\tilde{E}$
- Sequence of relationship assignments $R_1(\cdot), \ldots, R_n(\cdot)$, defined on $\tilde{E}$
  - assumption:

$$\forall e \in \tilde{E} \setminus E_i : R_i(e) = unknown$$

Methodology: "Historical" Measures

$$J(e) = \{(j,k) \mid j,k \in \{1,\ldots,n\} \wedge j < k \wedge e \in E_j \cap E_k \wedge \forall h \in \{j+1,\ldots,k-1\}: e \notin E_h\}$$

$$Occurrences(e) = |\{i \mid R_i(e) \neq unknown\}|$$

$$Assignments(e) = |\{i \mid e \in E_i\}|$$

$$Changes(e) = |\{(i,j) \mid (i,j) \in J(e) \wedge R_i(e) \neq R_j(e)\}|$$

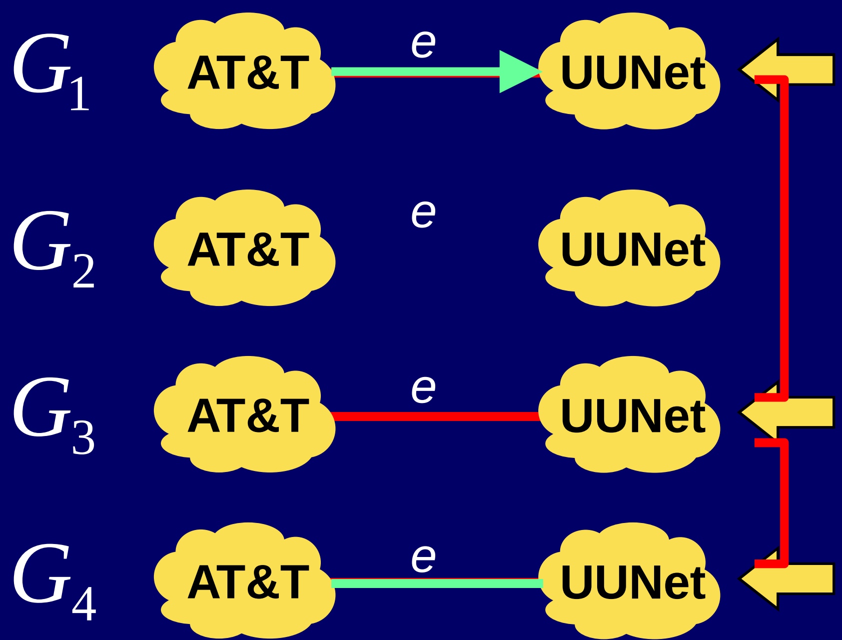

$$Occurrences(e) = 3$$

$$Assignments(e) = 2$$

$$Changes(e) = 2$$

Summary

- Introduction
- Motivations
- Methodology
- Tools
- Experimental Results
- Conclusions

Software Tools

- TORQUE (Type Of Relationship QUality Evaluation): four tools supporting the methodology [4]
  - AS path extractor: processes the data sets ('show ip bgp' dumps) to extract a list of AS paths from them (typical input of the inference algorithms)
  - graph generator: creates a labelled graph; each edge is associated with the value of $R(\cdot)$
  - graph compare tool: computes the measures we have just defined
  - graph view tool: produces reports

[4] available for download at http://www.dia.uniroma3.it/~compunet/torque

Software Tools: Usage

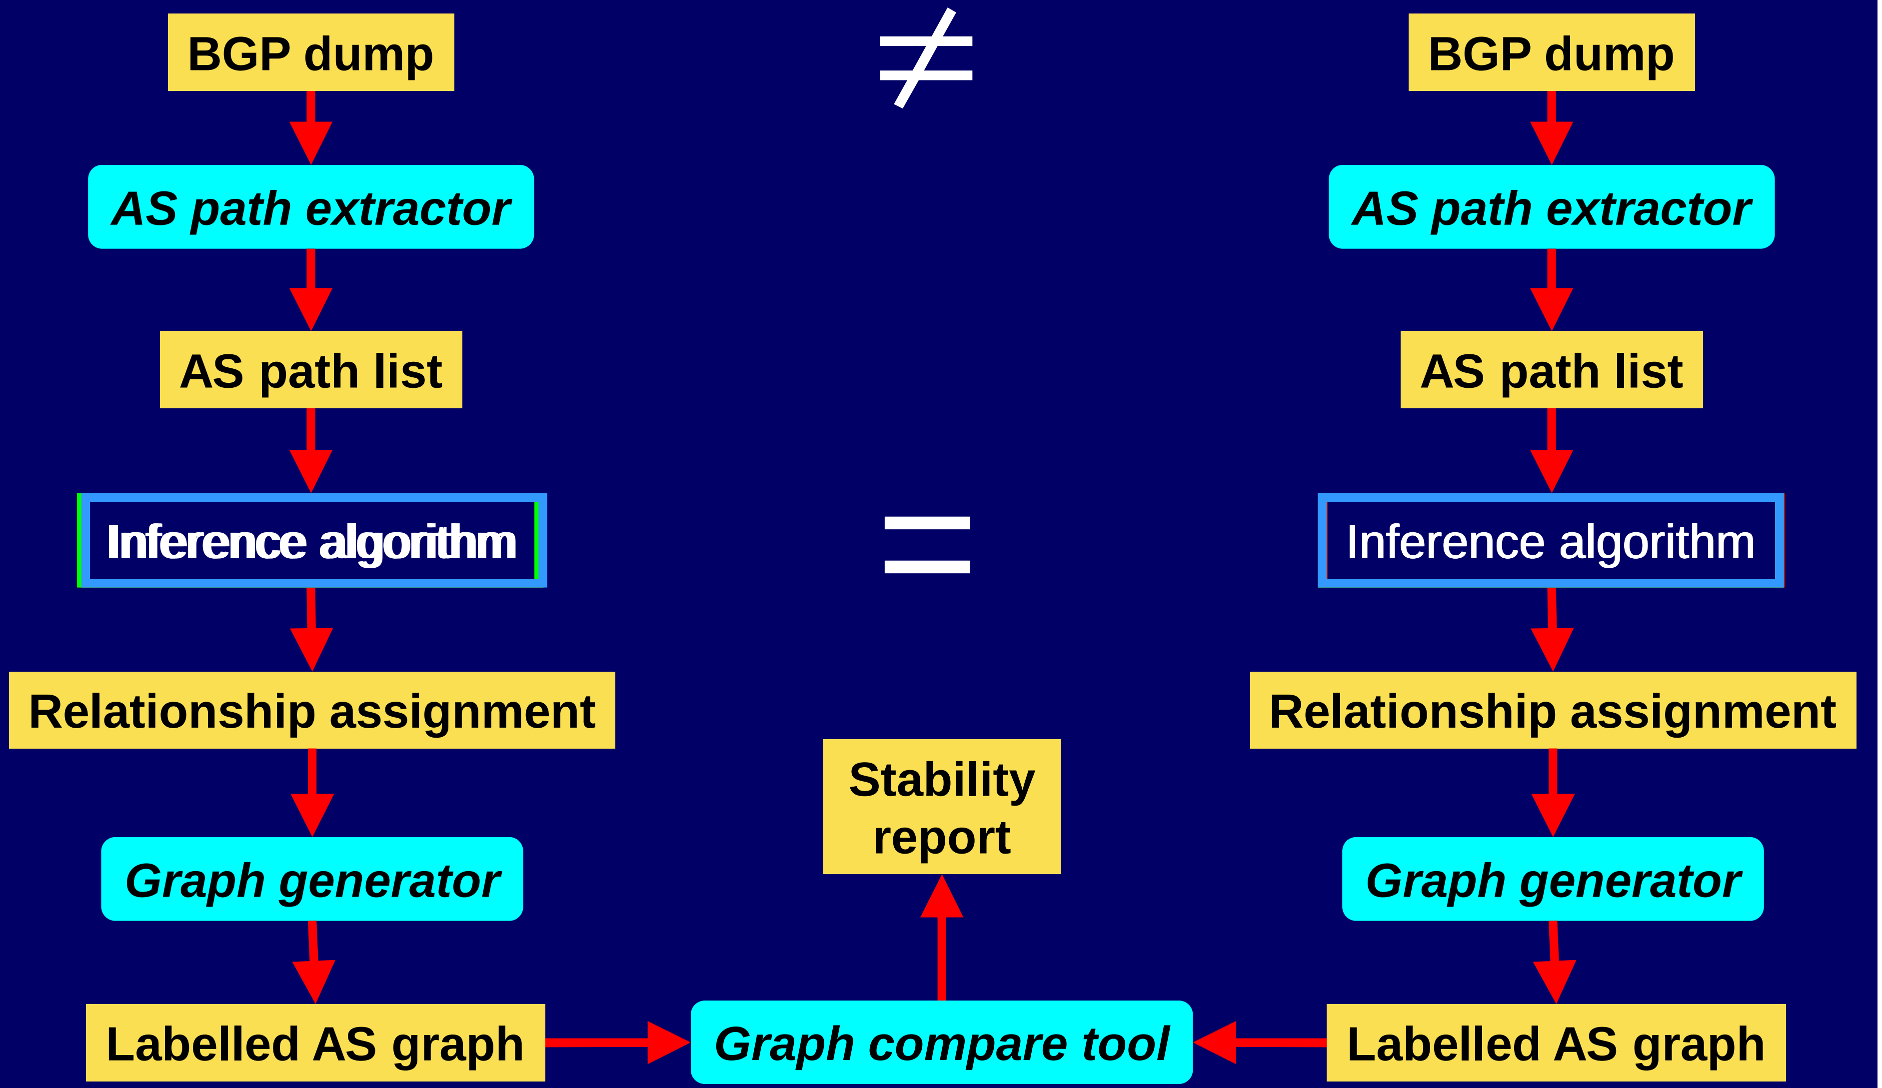

Software Tools: Usage

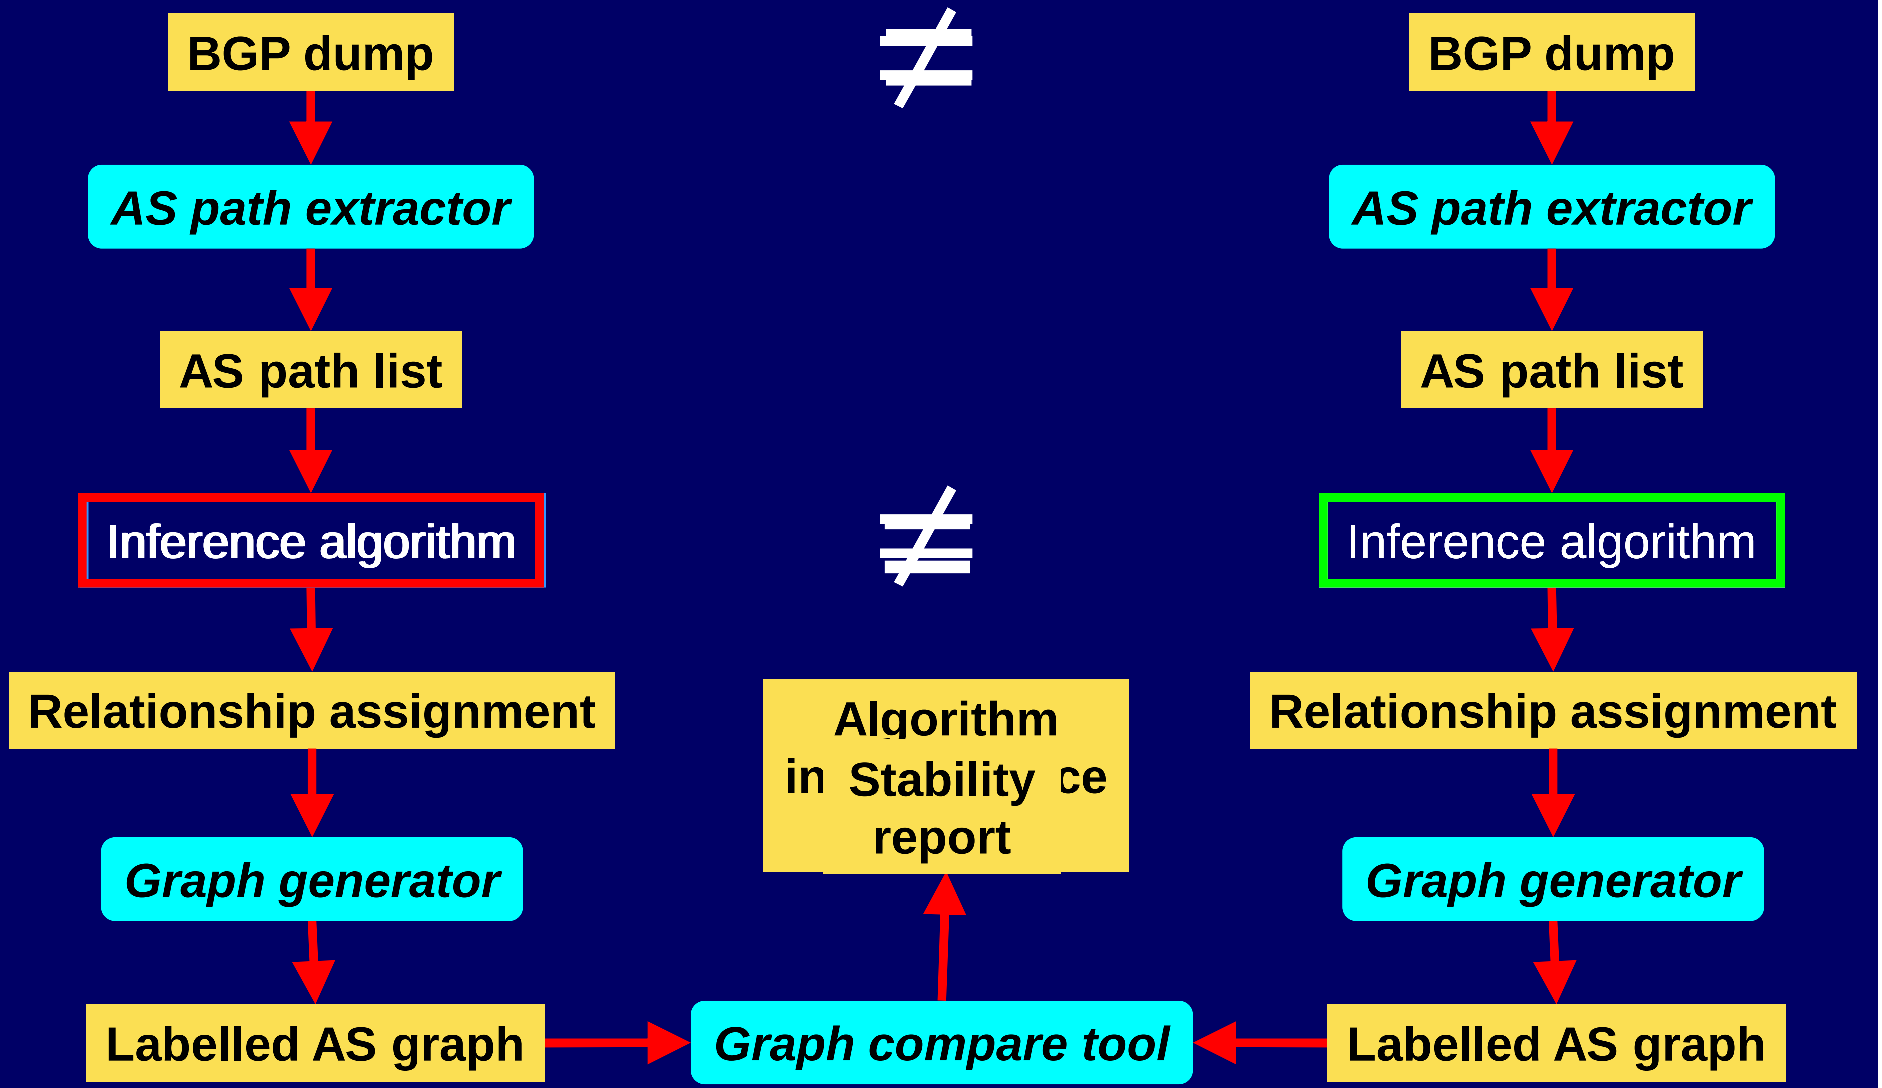

Software Tools: Usage

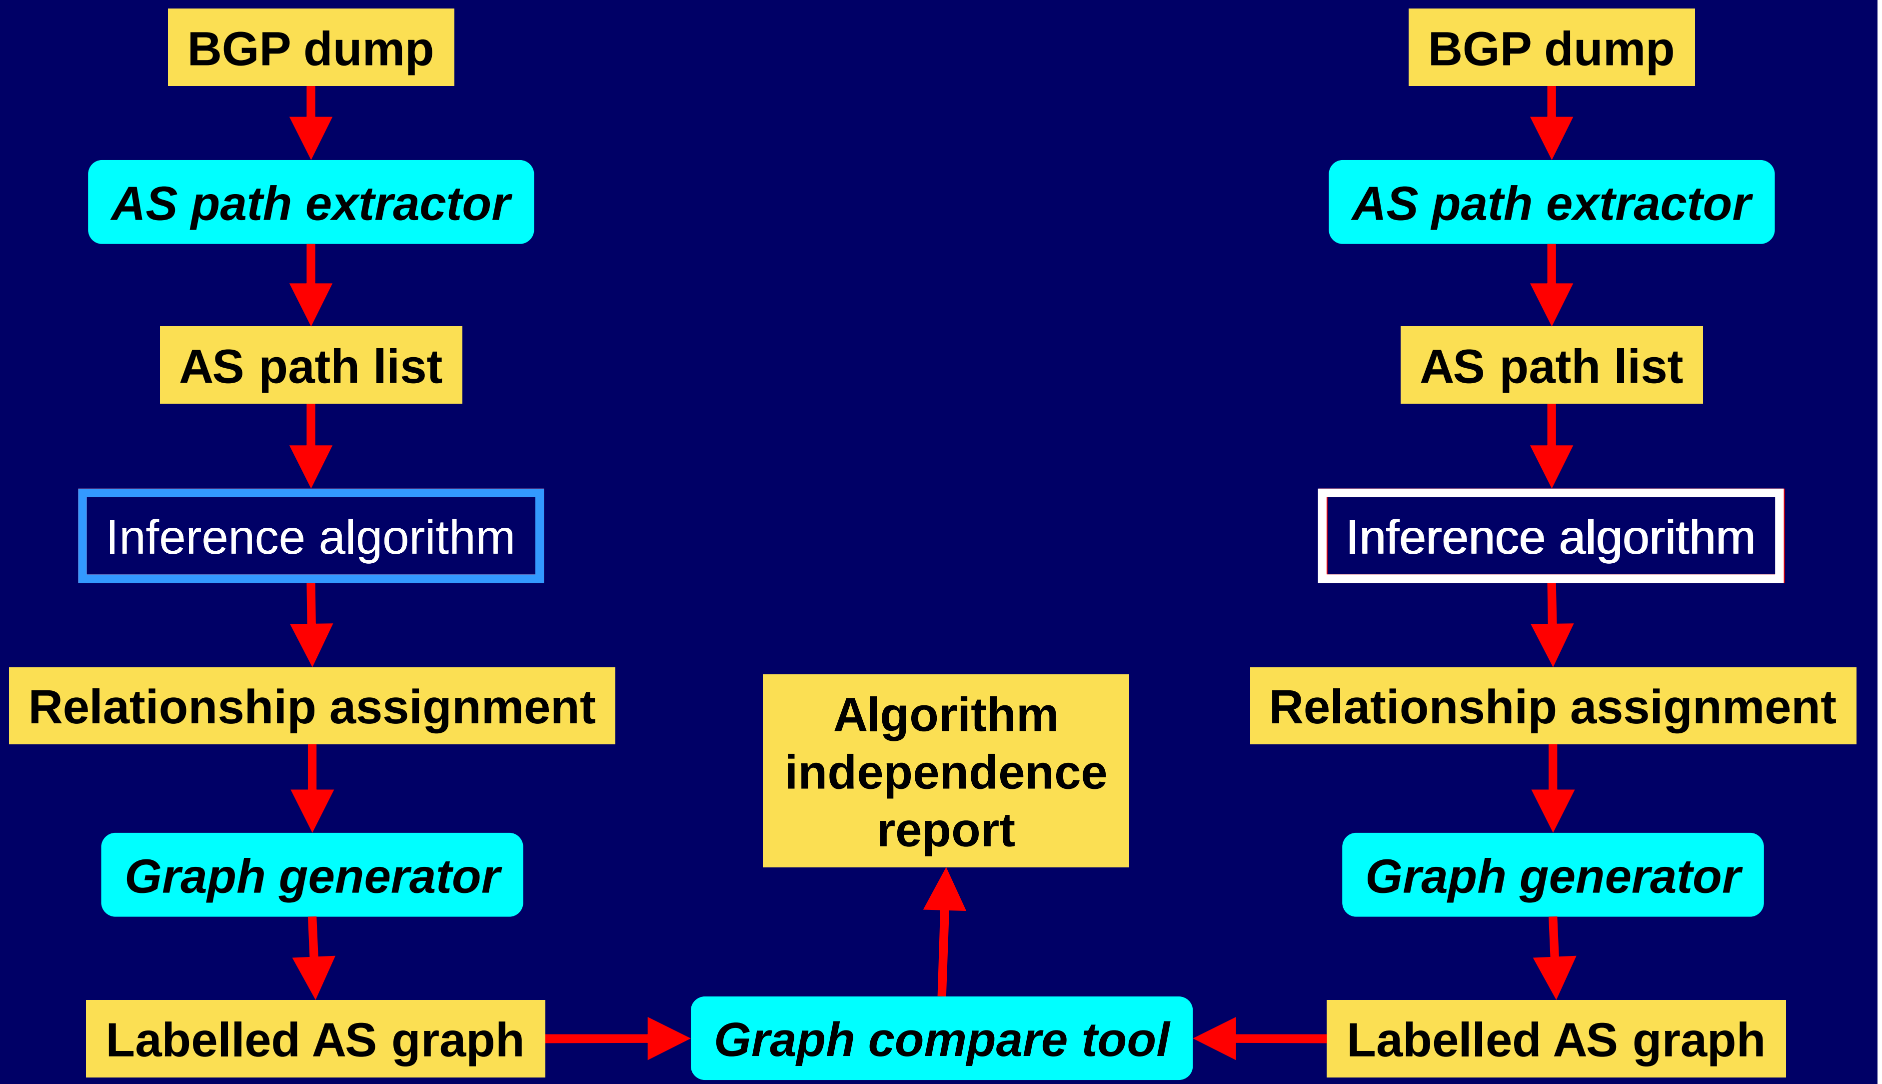

Software Tools: Usage

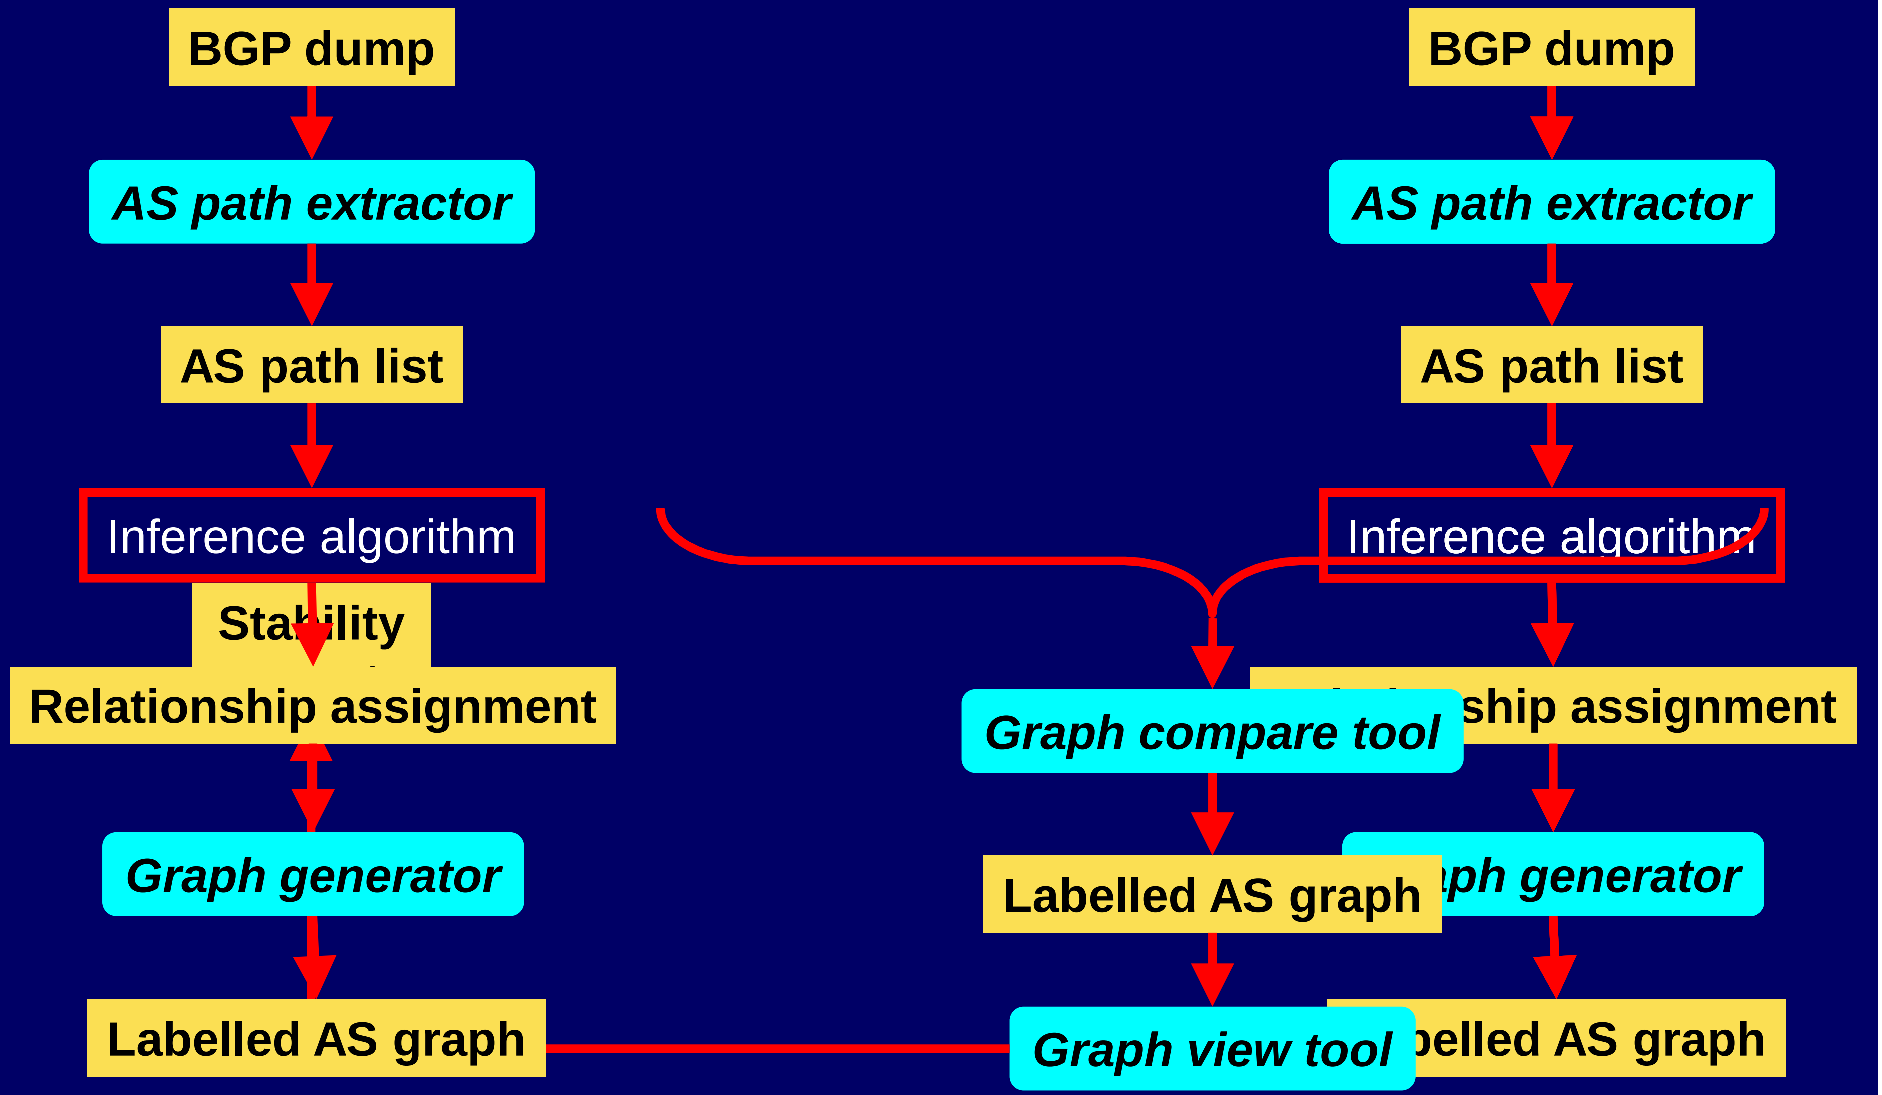

Summary

- Introduction
- Motivations
- Methodology
- Tools
- Experimental Results
- Conclusions

Experimental Results: Data sets

- http://www.cs.berkeley.edu/
  - provides snapshots collecting simultaneous BGP tables from various telnet looking glasses

MVP

		AS graph		Unique
Date	Used looking glasses	Vertices	Edges	AS paths
18 Apr 2001	1, 1740, 3549, 3582, 3967, 4197, 7018, 8220, 8709	10909	23817	511200
29 Jan, 04 Feb 2002	1, 3549, 3582, 3967, 4197, 7018, 8220, 8709	12708	27555	722481
06 Apr 2002	1, 1838, 3549, 3582, 3967, 4197, 5511, 7018, 8220, 8709, 15290	13079	28309	942382
29 Jul 2002	1, 1838, 3257, 3549, 3582, 3967, 4197, 5511, 7018, 8220, 8709, 15290	13705	29073	948720
09 Aug 2002	1, 1838, 3257, 3549, 3582, 4197, 5511, 7018, 8220, 15290	13754	29009	894396
19 Oct 2002	1, 1838, 3582, 3967, 5511, 7018, 8220, 15290	14113	29422	881836
29 Oct 2003	1, 50, 210, 553, 852, 1838, 3257, 3549, 3582, 3741, 3967, 4197, 5388, 5511, 6395, 6539, 6893, 7018, 8220, 8709, 8843, 9328, 15290	16420	37470	1143373
13 Nov 2003	1, 50, 210, 553, 852, 1838, 3257, 3549, 3582, 3741, 3967, 4197, 5388, 5511, 6395, 6539, 6893, 7018, 8220, 8709, 8843, 9328, 15290	16461	37406	1157802
28 Nov 2003	1, 50, 210, 553, 852, 1838, 3257, 3549, 3582, 3741, 3967, 4197, 5388, 5511, 6395, 6539, 6893, 7018, 8220, 8709, 8843, 9328, 15290	16316	31419	247691
12 Dec 2003	1, 50, 210, 553, 852, 1838, 3257, 3549, 3582, 3741, 3967, 4197, 5388, 5511, 6395, 6539, 6893, 7018, 8220, 8709, 8843, 9328, 15290	16585	37790	1216534
29 Dec 2003	1, 50, 210, 553, 852, 1838, 3257, 3549, 3582, 3741, 3967, 4197, 5388, 5511, 6395, 6539, 6893, 7018, 8220, 8709, 8843, 9328, 15290	16728	38162	1164370
13 Jan 2004	1, 50, 210, 553, 852, 1838, 3257, 3549, 3582, 3741, 3967, 4197, 5388, 5511, 6395, 6539, 6893, 7018, 8220, 8709, 8843, 9328, 15290	16762	38205	1264677

Experimental Results: Data sets

- University of Oregon RouteViews project. http://www.routeviews.org.
  - provides dumps of the RIB of a router having about 50 active peering sessions at the time of the experiments

	Start	End	# of snapshots
RV1	25 Mar 2003 00:00	31 Mar 2003 22:00	84
RV2	25 Sep 2001 00:00	08 Oct 2001 22:00	168

  - RV1: few commercial changes
  - RV2: many commercial changes
  - Snapshots collected every 2 hours

sources: [5, 6]

[5] xSP Archives. http://www.internetnews.com
[6] ISP Business: Archives. http://isp-planet.com/

Experimental Results: Stability Analysis

- (A),(B): pair of consecutive snapshots
- R(·) is defined on the intersection of G_A and G_B
- Data set: MVP
- Algorithm: DPP
- PeerInA, PeerInB are omitted

	18 Apr 2001 (A)	29 Jul 2002 (A)	09 Aug 2002 (A)
	29 Jan, 04 Feb 2002 (B)	09 Aug 2002 (B)	19 Oct 2002 (B)
OnlyInA	7	5	3
OnlyInB	18	1	2
Both	15919	28039	25458
Consistent	15118 (95%)	27719 (99%)	24916 (98%)
Inconsistent	801 (5%)	320 (1%)	542 (2%)

	13 Nov 2003 (A)	28 Nov 2003 (A)	12 Dec 2003 (A)
	28 Nov 2003 (B)	12 Dec 2003 (B)	29 Dec 2003 (B)
OnlyInA	4	4	7
OnlyInB	2	2	2
Both	30654	31034	37038
Consistent	30120 (98%)	30509 (98%)	36653 (99%)
Inconsistent	534 (2%)	525 (2%)	385 (1%)

- Very good stability (>95%)

Experimental Results: Stability Analysis

- Data set: MVP
- Algorithm: SARK

	18 Apr 2001 (A) / 29 Jan, 04 Feb 2002 (B)	29 Jul 2002 (A) / 09 Aug 2002 (B)	09 Aug 2002 (A) / 19 Oct 2002 (B)
OnlyInA	86	80	111
OnlyInB	77	94	89
Both	15749	27789	25224
Consistent	15056 (96%)	27343 (98%)	24425 (97%)
Inconsistent	693 (5%)	446 (2%)	799 (3%)
Opposite	55 (8%)	14 (3%)	44 (6%)
PeerInA	350 (51%)	257 (57%)	367 (46%)
PeerInB	288 (41%)	175 (40%)	388 (48%)

- Very good stability (>96%)
- Most changes are due to the appearance or disappearance of peerings

	13 Nov 2003 (A) / 28 Nov 2003 (B)	28 Nov 2003 (A) / 12 Dec 2003 (B)	12 Dec 2003 (A) / 29 Dec 2003 (B)
OnlyInA	613	64	157
OnlyInB	58	599	168
Both	29099	29482	35739
Consistent	28748 (99%)	29146 (99%)	35100 (98%)
Inconsistent	351 (1%)	336 (1%)	639 (2%)
Opposite	8 (2%)	2 (1%)	10 (2%)
PeerInA	237 (68%)	103 (31%)	346 (54%)
PeerInB	106 (30%)	231 (68%)	283 (44%)

Experimental Results: Stability Analysis

RV1
DPP

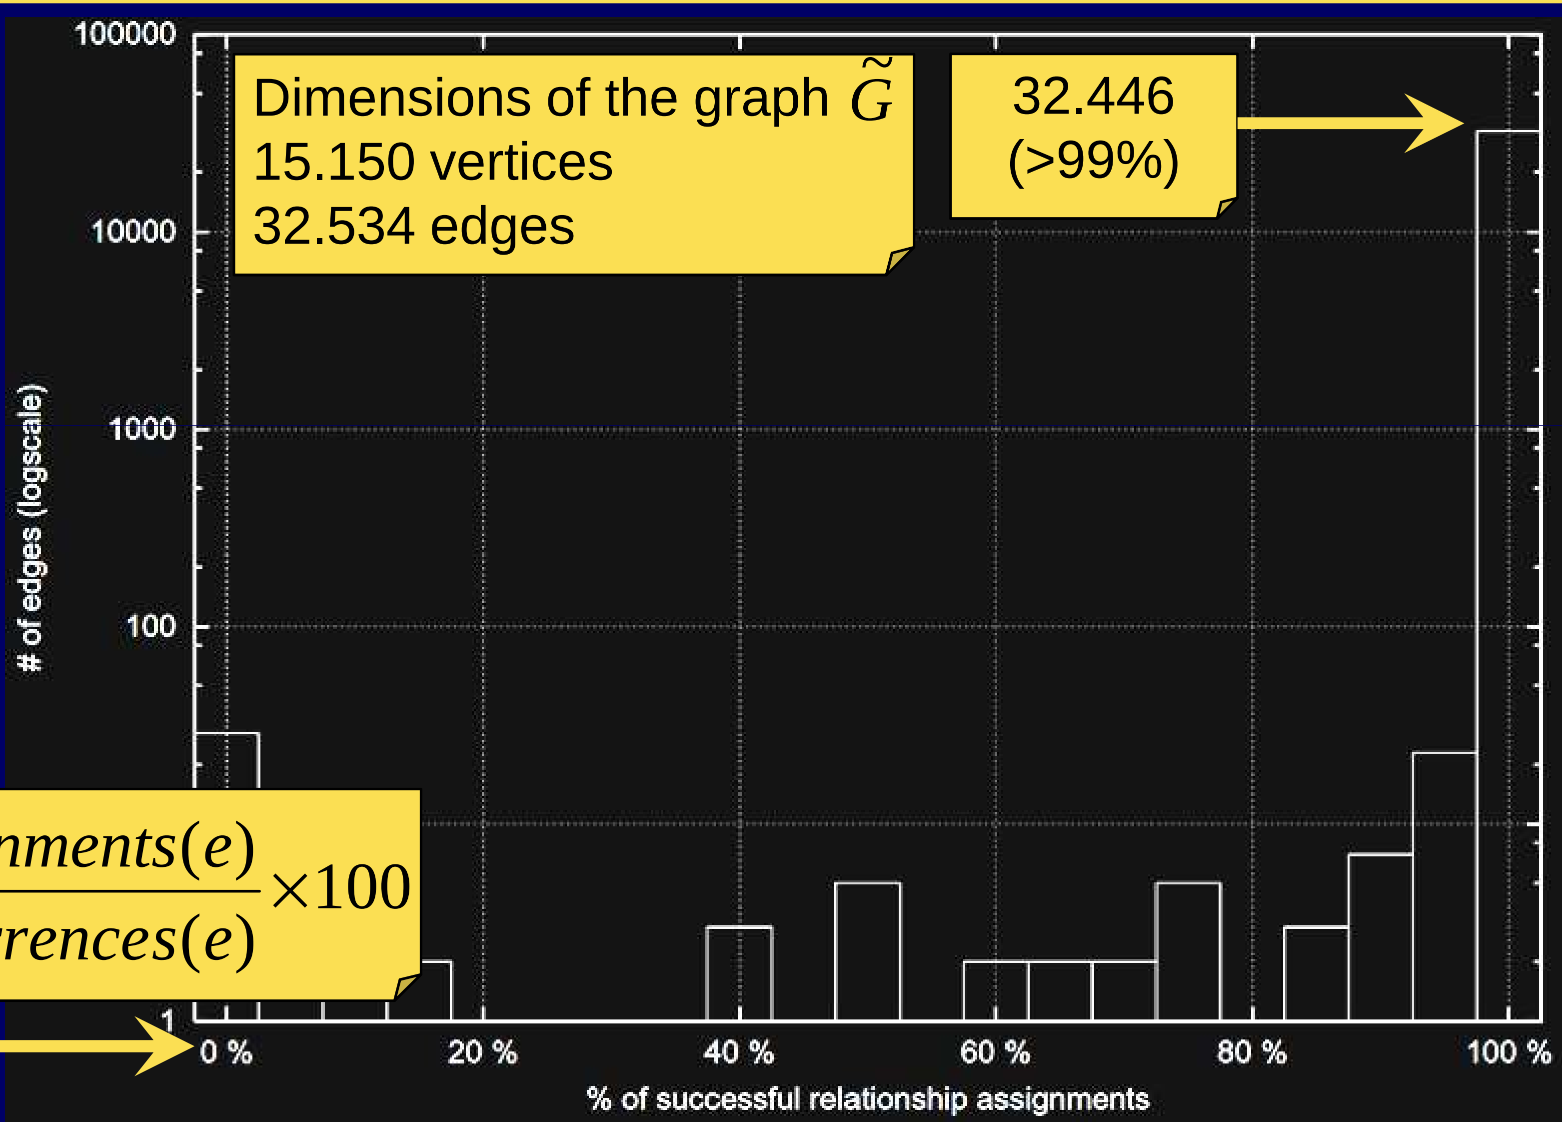

$$\frac{Assignments(e)}{Occurrences(e)} \times 100$$

Experimental Results: Stability Analysis

RV2
DPP

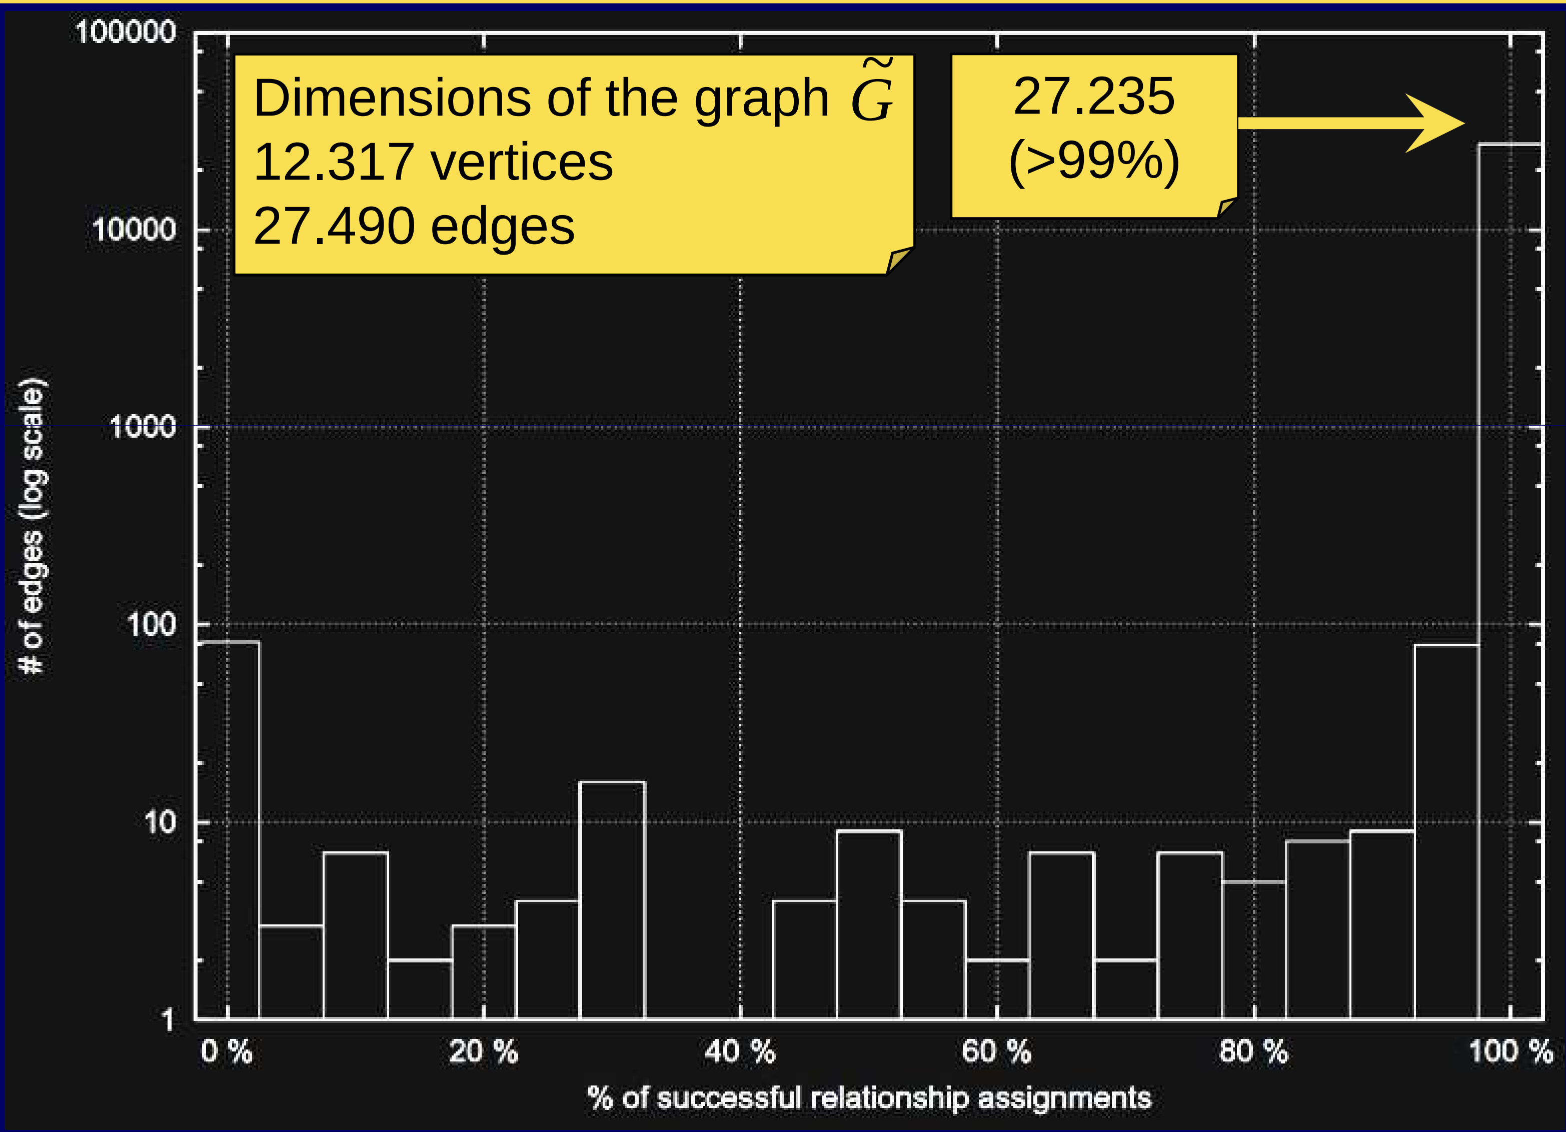

Experimental Results: Stability Analysis

RV1

DPP

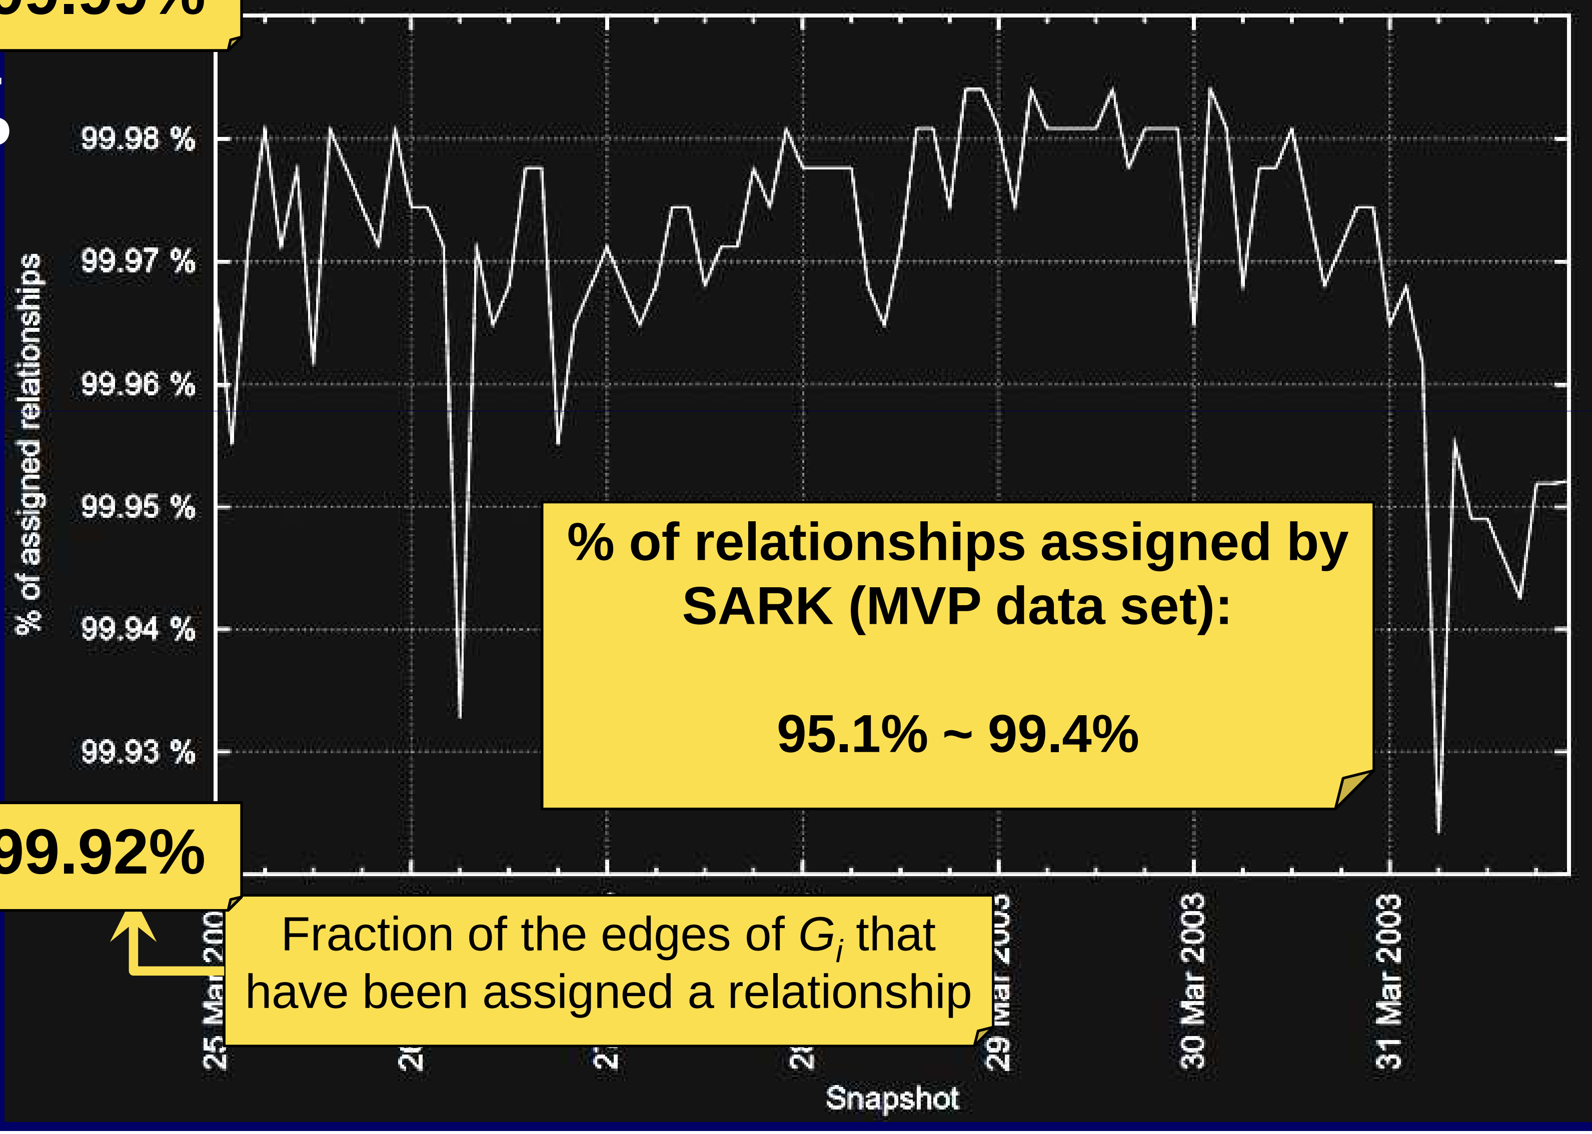

99.99%

99.92%

% of relationships assigned by SARK (MVP data set):

95.1% ~ 99.4%

Fraction of the edges of G_i that have been assigned a relationship

Experimental Results: Stability Analysis

RV2
DPP

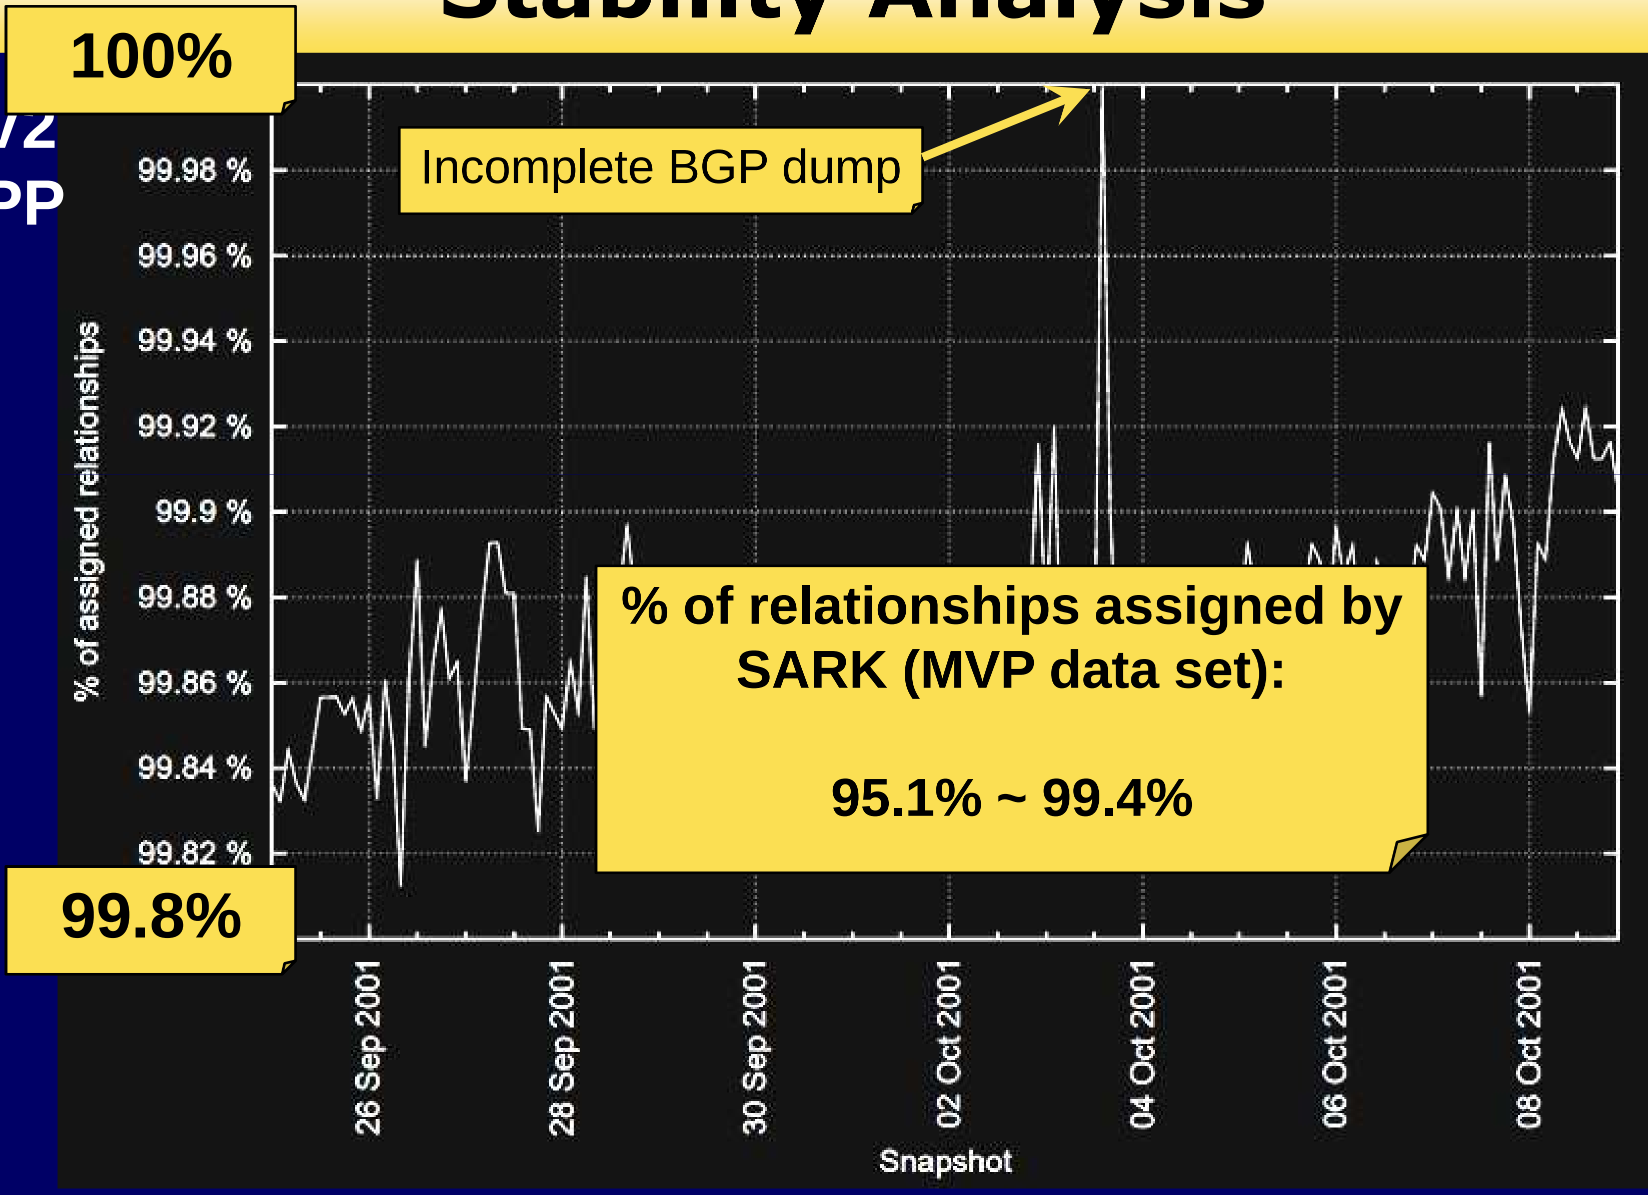

Experimental Results: Stability Analysis

RV1
DPP

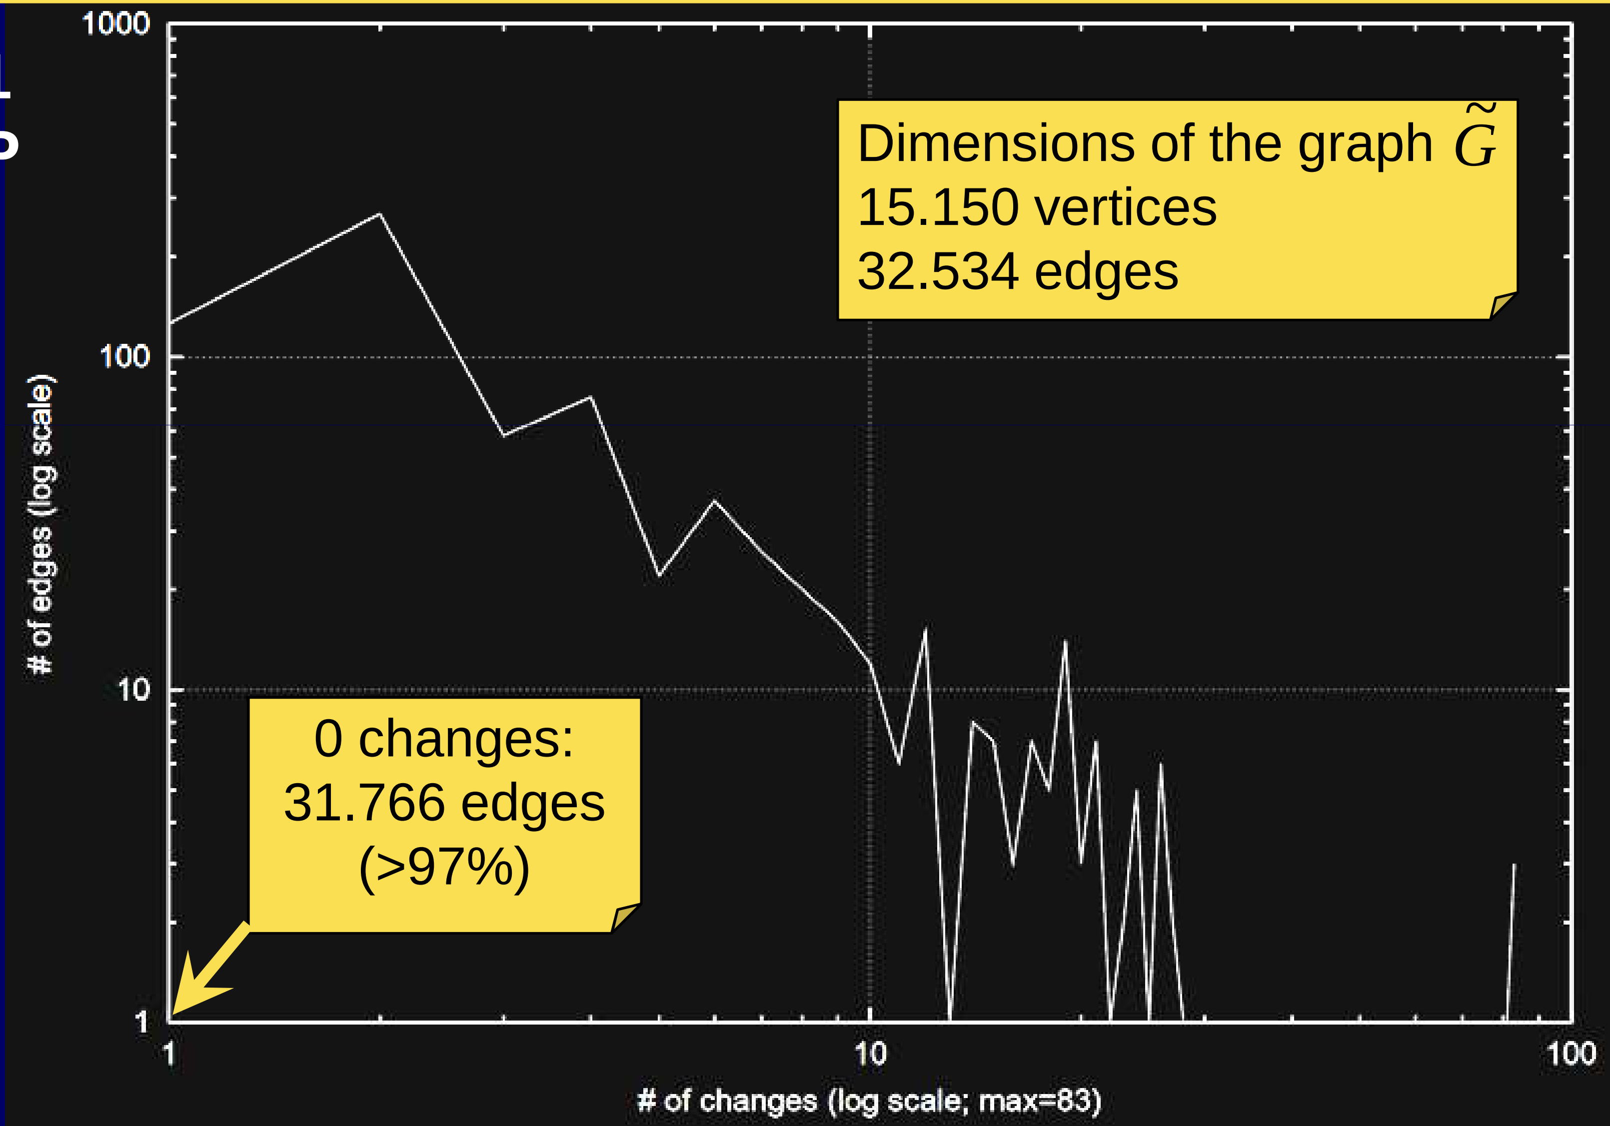

Experimental Results: Stability Analysis

RV2
DPP

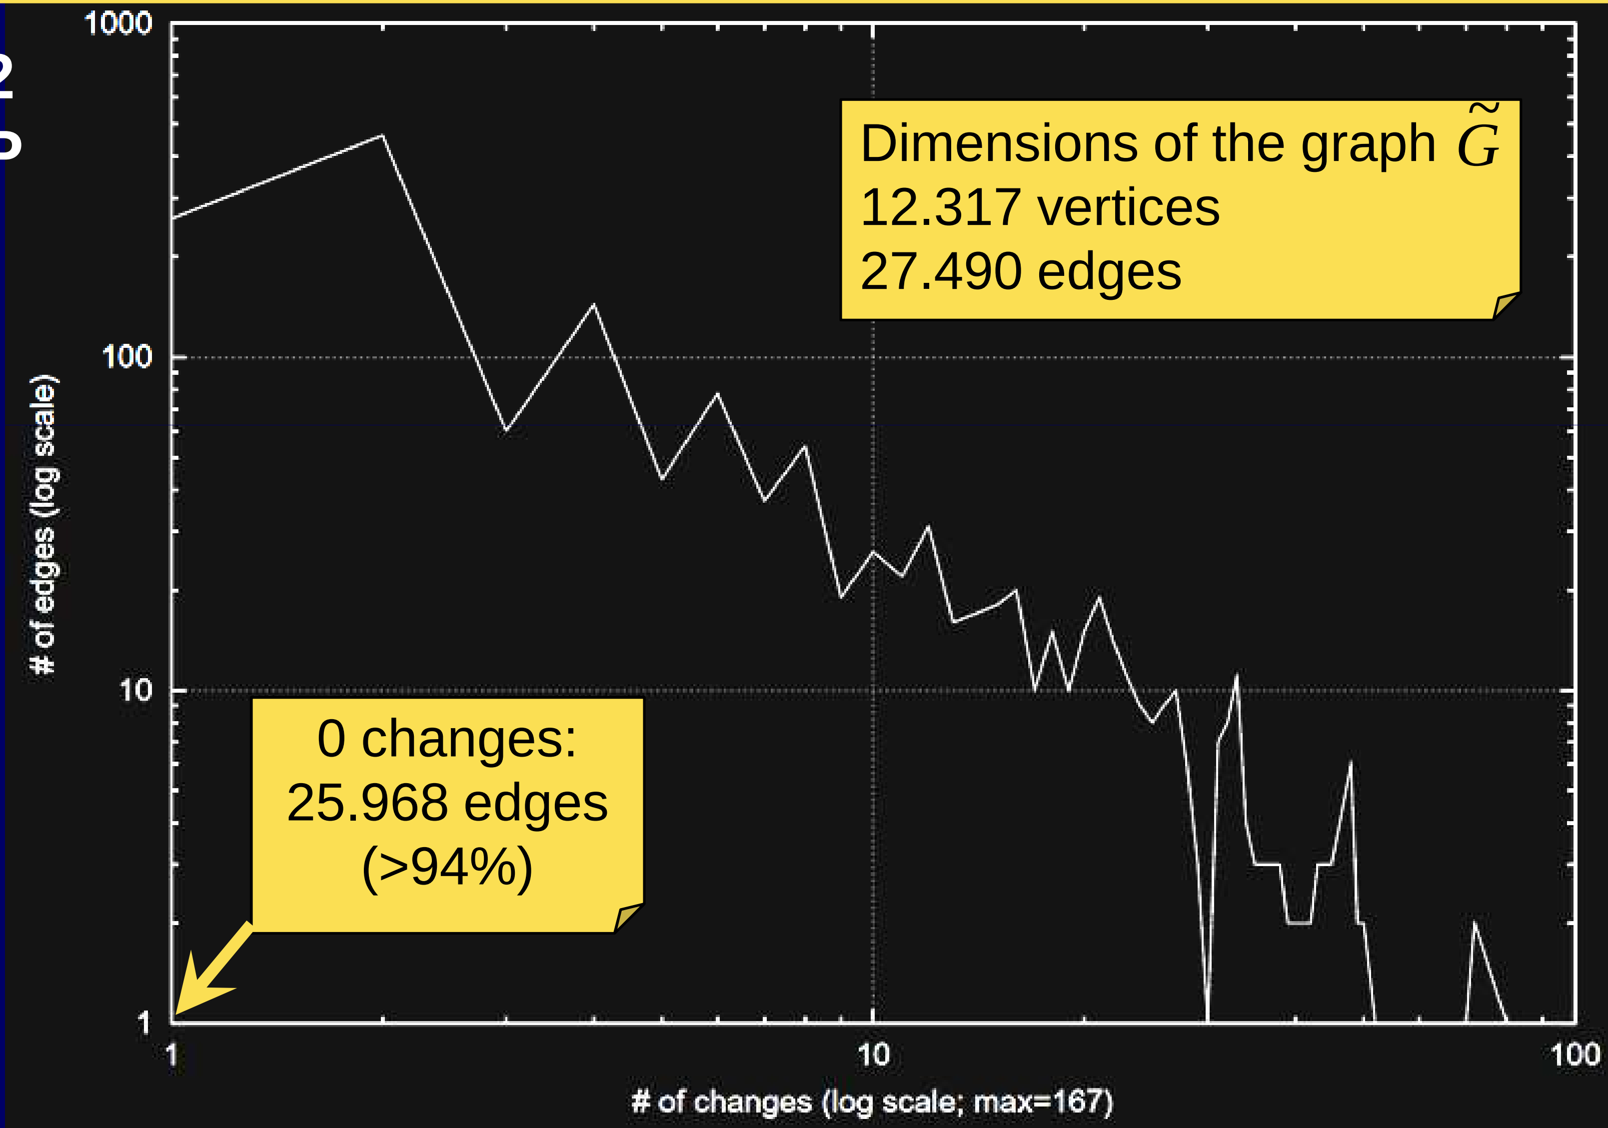

Experimental Results: Independence from the Algorithm

- (A): DPP algorithm
- (B): SARK algorithm
- Data set: MVP
- R(·) is defined on the intersection of G_A and G_B
- PeerInA is omitted

	18 Apr 2001	29 Jan, 04 Feb 2002	29 Jul 2002	19 Oct 2002
OnlyInA	192	213	200	196
OnlyInB	39	21	6	14
Both	23584	27317	28866	29212
Consistent	21487 (91%)	24991 (91%)	26261 (91%)	26659 (91%)
Inconsistent	2097 (9%)	2326 (9%)	2605 (9%)	2553 (9%)
Opposite	993 (47%)	1047 (45%)	1176 (45%)	1247 (49%)
PeerInB	1104 (53%)	1279 (55%)	1429 (55%)	1306 (51%)

- Fair independence (>90%)
- Differences are equally shared between c-p and p-p

	29 Oct 2003	13 Nov 2003	28 Nov 2003	12 Dec 2003
OnlyInA	1205	1173	1548	1367
OnlyInB	28	25	5	12
Both	36235	36208	29866	36409
Consistent	32683 (90%)	32651 (90%)	28332 (95%)	32754 (90%)
Inconsistent	3552 (10%)	3557 (10%)	1534 (5%)	3655 (10%)
Opposite	1689 (48%)	1738 (49%)	829 (54%)	1783 (49%)
PeerInB	1863 (52%)	1819 (51%)	705 (46%)	1872 (51%)

Summary

- Introduction
- Motivations
- Methodology
- Tools
- Experimental Results
- Conclusions

Conclusions

- Analysis of (the results of) state-of-the-art algorithms for the inference of commercial relationships between Autonomous Systems
  - stability analysis
  - independence from the algorithm
- A new methodology
  - a set of measures
  - a procedure (and software tools) to compute them
- Application of the methodology
  - stability analysis: DPP and SARK
  - independence from the algorithm: DPP against SARK
- Outcome of the analysis
  - very good stability (>94%)
  - fair independence from the algorithm (>90%)
- Based on this work, we think (these) inference algorithms are trustworthy

Future Work

- Analyzing the results produced by the algorithm in [1]
- Degrees of freedom
- Game theory
- …

[1] Lixin Gao. On Inferring Autonomous System Relationships in the Internet. *IEEE/ACM Transactions on Networking*, 9(6):733–745, Dec 2001.