Locating Internet Routing Instabilities

Olaf Maennel
Technisch Universität München

Z. Morley Mao
University of Michigan

Anja Feldmann
Technische Universität München

Bruce Maggs
CMU / Akamai Technologies

Arthur Berger
MIT / Akamai Technologies

BGP

- Defacto standard inter-domain routing protocol
- Path vector protocol
- Policy routing protocol

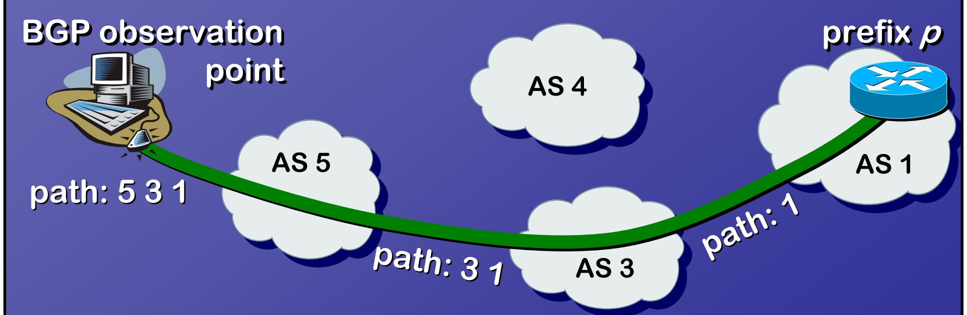

- Topology changes, etc… ⇒ BGP updates

Definitions

- **LOCATION:**
  either internal to an AS (e.g., inside AS1)
  – or –
  link between two ASes (e.g., AS1–AS2)

- **INSTABILITY:**
  any change of BGP advertisment over an e-BGP session

Why identify locations of instabilities?

- Instabilities can lead to
  - Unreachablility / poor performance
  - Route oscillation
  - BGP churn
  - Black holes
  - …

- Identifying the location enables corrective action

Causes of instabilities

- Possible causes for BGP instabilities
  - BGP session availability
    - Session establishment/teardown/reset
  - BGP session filters
    - BGP attribute or filter manipulation
    - misconfiguration
  - IGP cost change
    - IGP metric change, link or node failures
  - … *(more in the paper)*

Outline

- High-level approach
- Limitations
- Evaluation

Outline

- High-level approach
- Limitations
- Evaluation

High-level approach (I): Dimensions

Locate BGP instability by analyzing BGP updates along three dimensions:

1. Time
2. Views
3. Prefixes

We use same dimensions as:

– Caesar, Subramanian, Katz

– Chang, Govindan, Heidemann

– Lad, Nanavati, Massey, Zhang

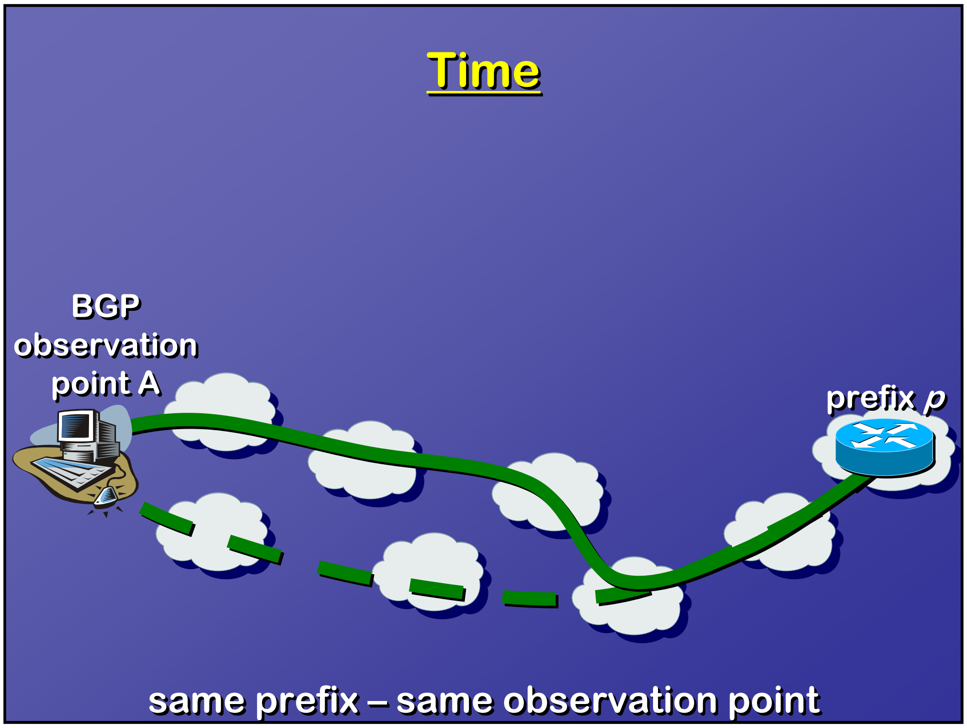
Time
BGP observation point A
prefix p
same prefix – same observation point

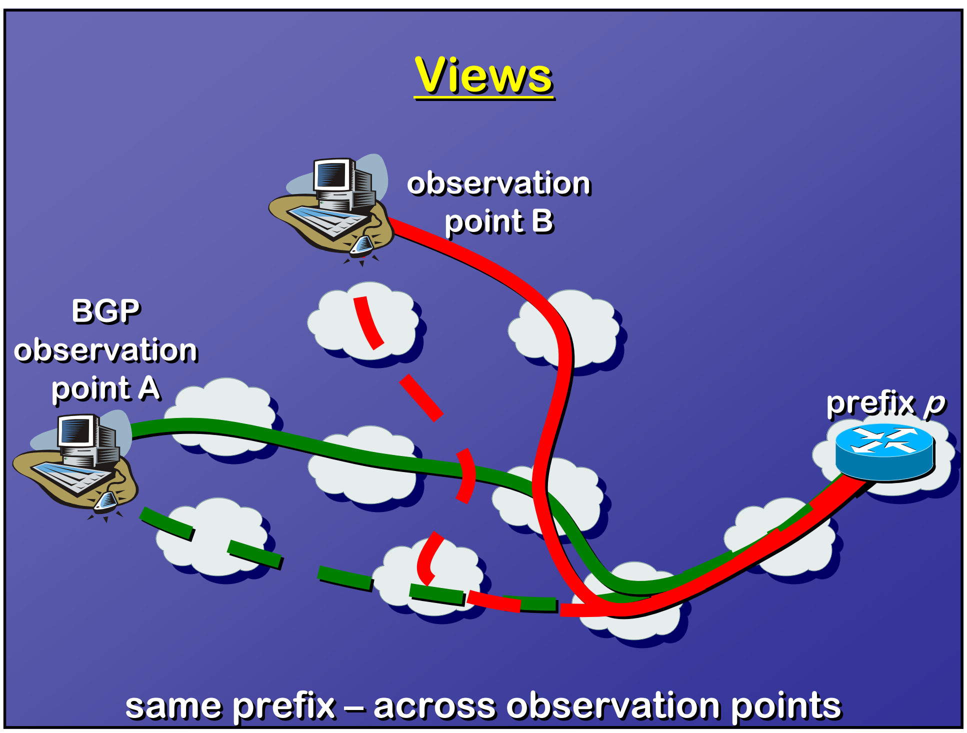
Views
observation point B
BGP observation point A
prefix p
same prefix – across observation points

Prefixes

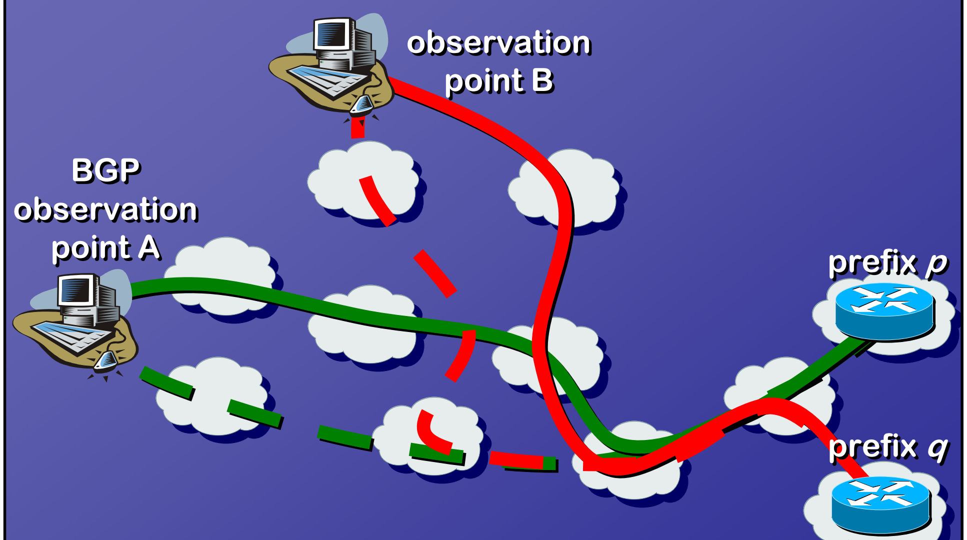

correlate across prefixes

High-level approach (II): Algorithms

- UNION heuristic (Time)
- INTERSECTION heuristic (Views)
- GREEDY heuristic (Prefixes)

UNION heuristic

- Routing instability ⇒
  change from "previous" to "new" path

  "previous best" path no longer available
  – or –
  "new best" path becomes available

  ⇒ UNION of ASes as candidates

- Input: BGP updates
- Output: "update bursts" with candidates

UNION heuristic

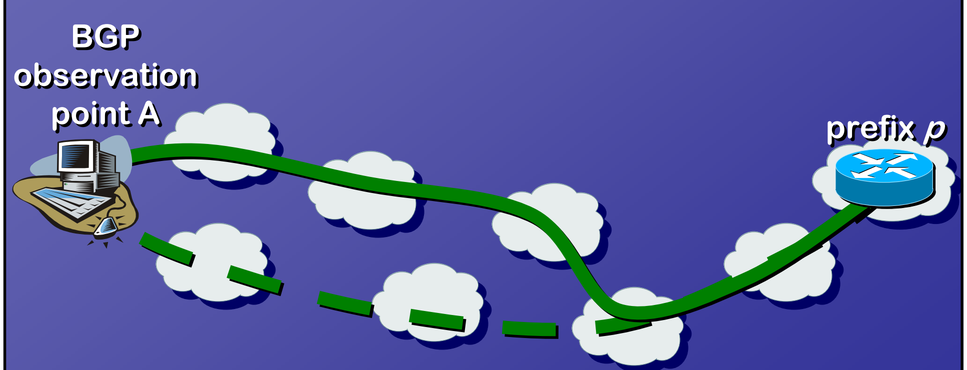

UNION of AS paths as candidates

INTERSECTION heuristic

- Routing instability ⇒
  changes at multiple observation points

  ⇒ INTERSECTION of candidate sets

- Input: "update bursts" with candidates
- Output: "events" with instability sets

INTERSECTION heuristic

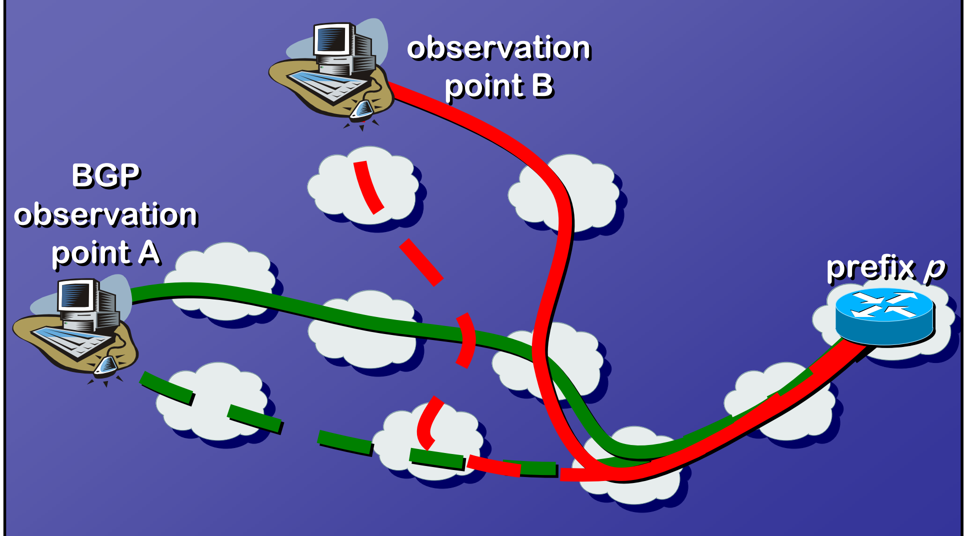

Changes observable at multiple observation points

INTERSECTION heuristic

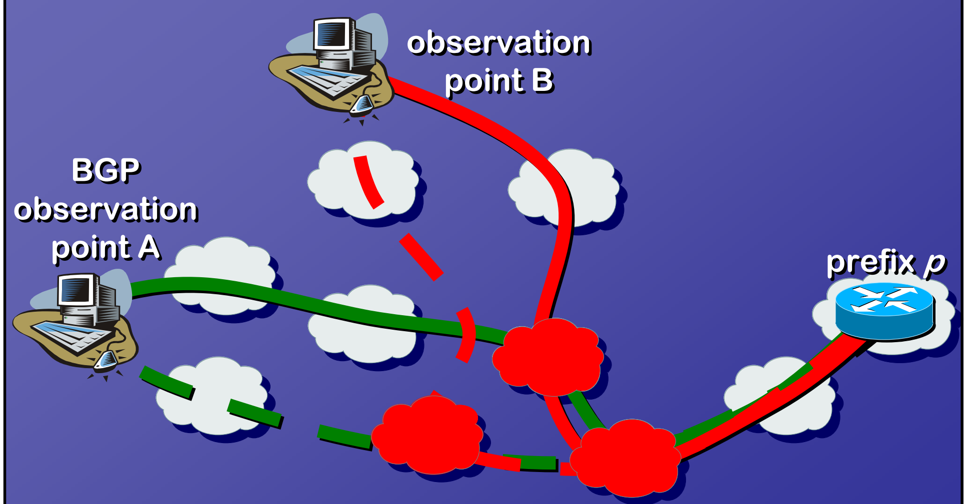

INTERSECTION of candidate sets

GREEDY heuristic

- Routing instability ⇒
  changes multiple prefixes

  ⇒ identify correlated prefixes

- Input: “events” with instability sets
- Output: “correlated events”

GREEDY heuristic

Goal: Distinguish between multiple simultaneous instabilities.

1. Determine most popular AS edge in instability sets
   - For all events and for each edge in instability set
     - Counter{edge} ++
   - Sort edges by counter values
2. Chose edge with largest counter value as candidate AS edge for associated events
3. Remove these events from the input
4. Repeat

Outline

- High-level approach
- Limitations
- Evaluation

Problems with UNION heuristic

- Location may not be in the UNION at all!
  ⇒ may lead to empty INTERSECTION
- Size of candidate set may be large

Caution: induced updates

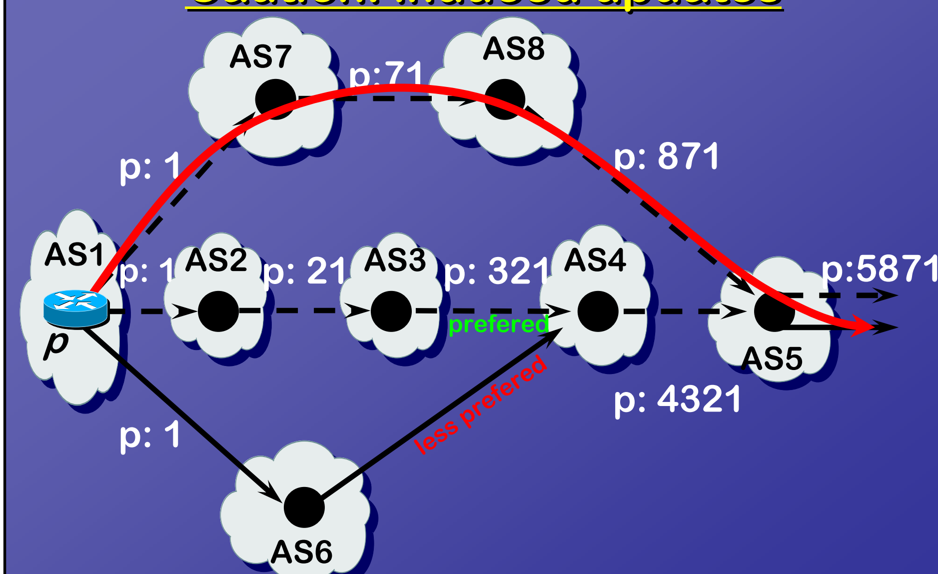

Policy: AS 4 preferes path over AS 3 instead of AS 6!

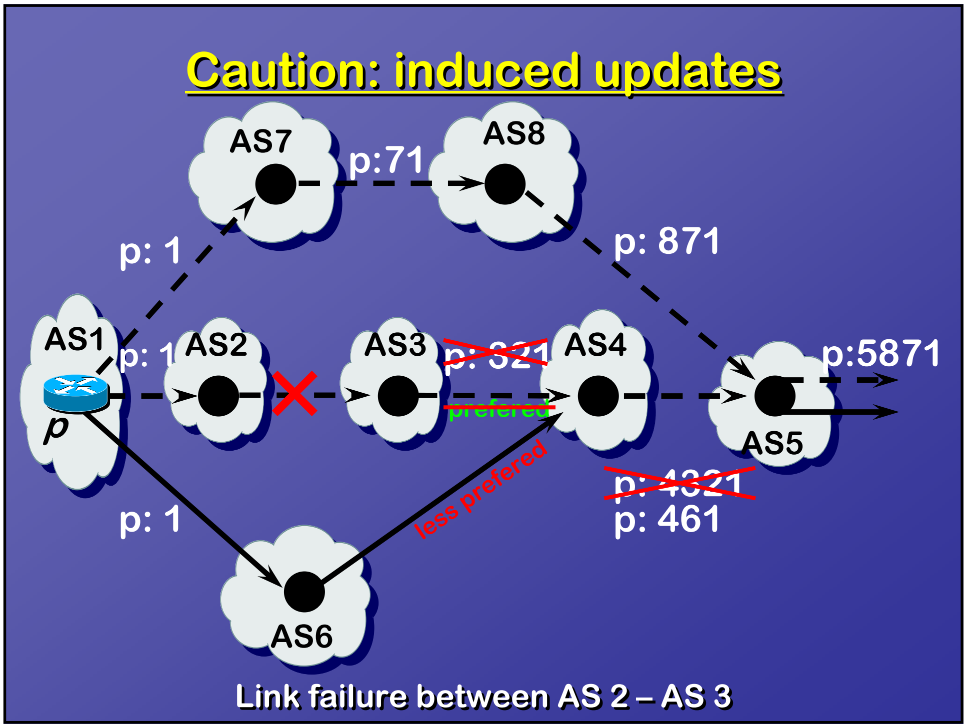
Caution: induced updates
AS7
p:71
AS8
p: 1
p: 871
AS1
p: 1
AS2
AS3
p: 321
AS4
p:5871
prefered
p
AS5
less prefered
p: 4321
p: 1
p: 461
AS6
Link failure between AS 2 – AS 3

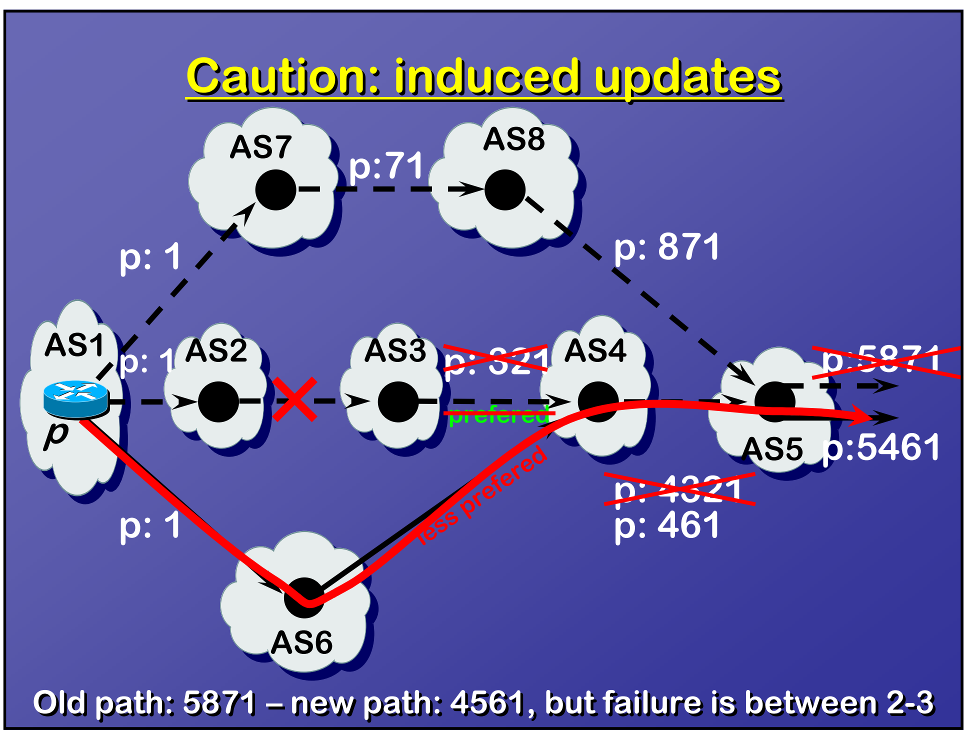
Caution: induced updates
AS7
p:71
AS8
p: 1
p: 871
AS1
p: 1
AS2
AS3
p: 321
AS4
p:5871
prefered
p
AS5
p:5461
less prefered
p: 4321
p: 1
p: 461
AS6
Old path: 5871 – new path: 4561, but failure is between 2-3

Reducing size of candidate set

- Idea: exclude some ASes
  - e.g., initial or final shared path segment
  - Narrows the candidate set, but may exclude the location

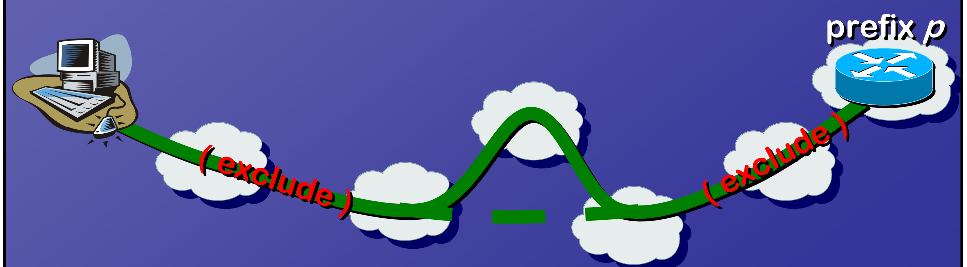

Caution on excluding candidate ASes

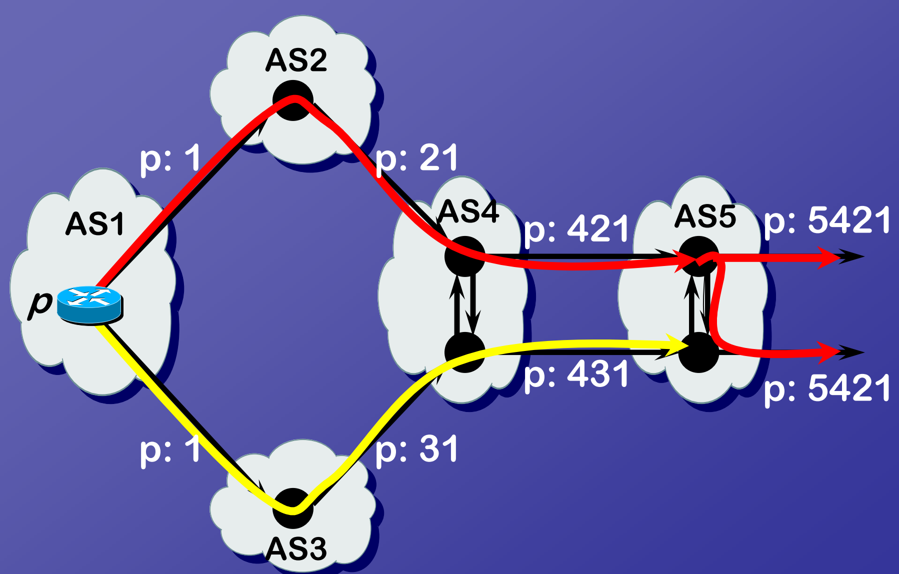

Caution on excluding candidate ASes

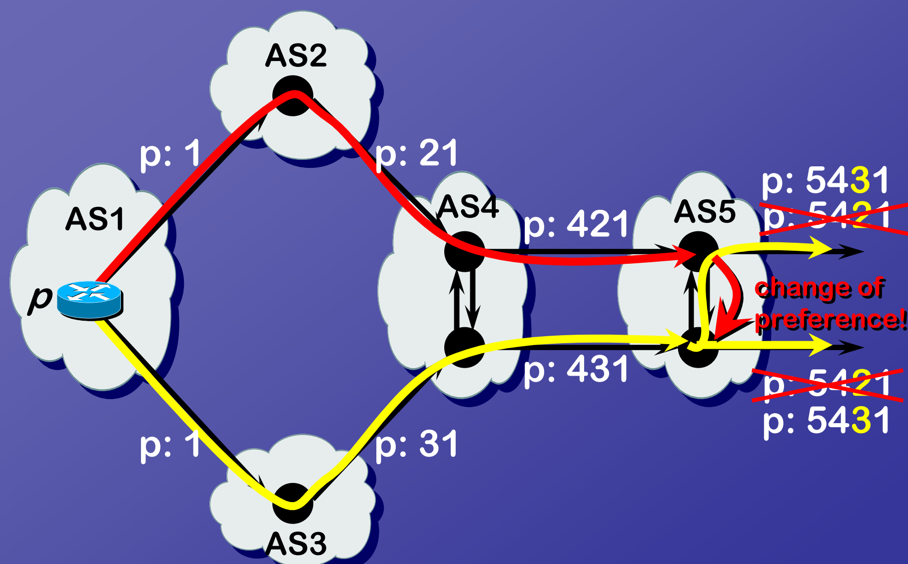

Path change: AS 2 replaces AS 3 — yet cause is AS 5

Good news

- Accurate in simulations
- Accurate when applied to real data
- Some formal justification in paper

Outline

- High-level approach
- Limitations
- Evaluation

Evaluation of methodology

- Simulation
  Setup:
  - Inferred AS topology from BGP data
  - "Routescope" Simulator
- Data analysis
  Setup:
  - BGP routing table dumps and updates from RIPE, Routeviews, and Akamai
  - Over 1,100 BGP feeds / 650 ASes (some I-BGP)

Simulation

- Topology
  - Inferred AS topology
  - Single node AS
- Policies
  - Inferred AS relationships
  - Prefer customer routes over peers over upstreams
  - Predicted routes agree with 90% of actual routes
- Link failures
  - Randomly selected
- Observation points
  - Randomly selected

Simulations: UNION + INTERSECTION

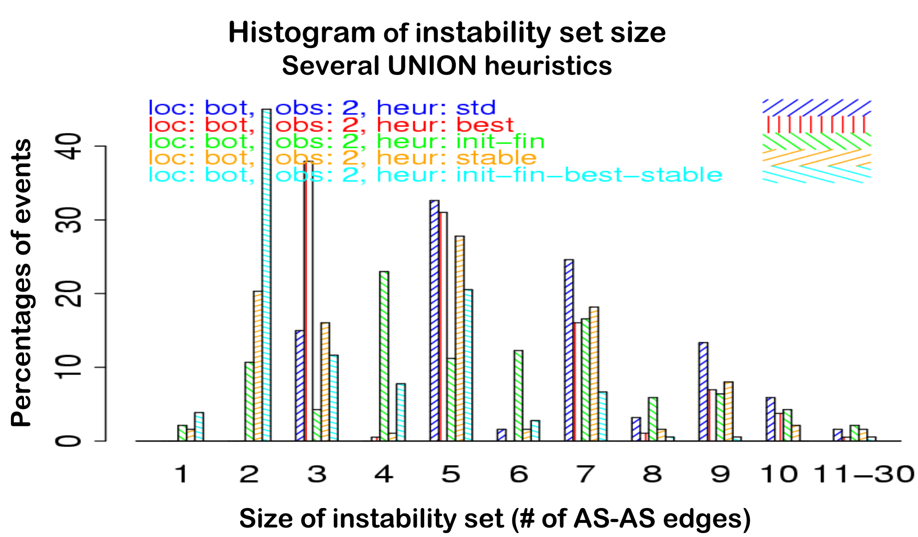

Choice of heuristic matters

Simulations: summary

- Location of observation points matter (in AS-hierachy)
- Number of observation points matter
- Average instability sizes after intersection:
  with only two obs.: ≤7 edges in 68%
  with 10 obs.: ≤7 edges in 88%
- The methodolgy never excludes the simulated failure location

Data analysis: UNION+INTERSECTION

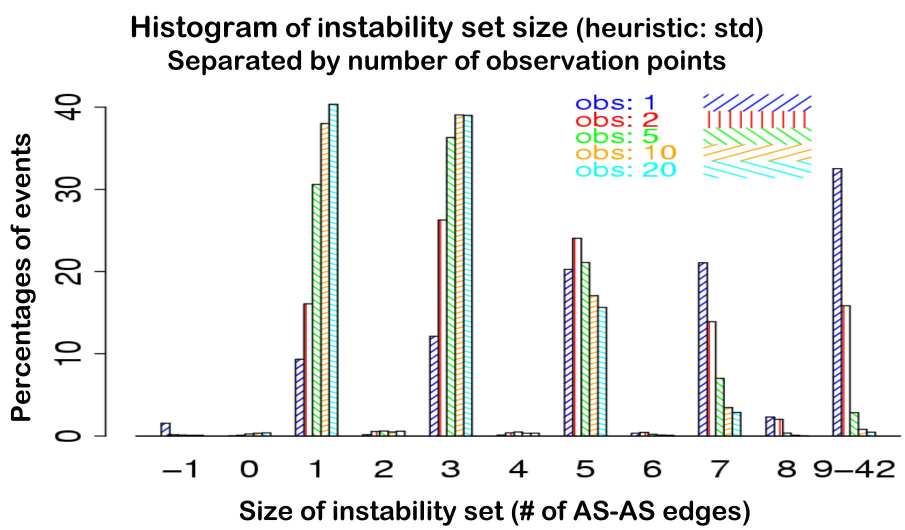

Number of observation points matter

Data analysis: UNION+INTERSECTION

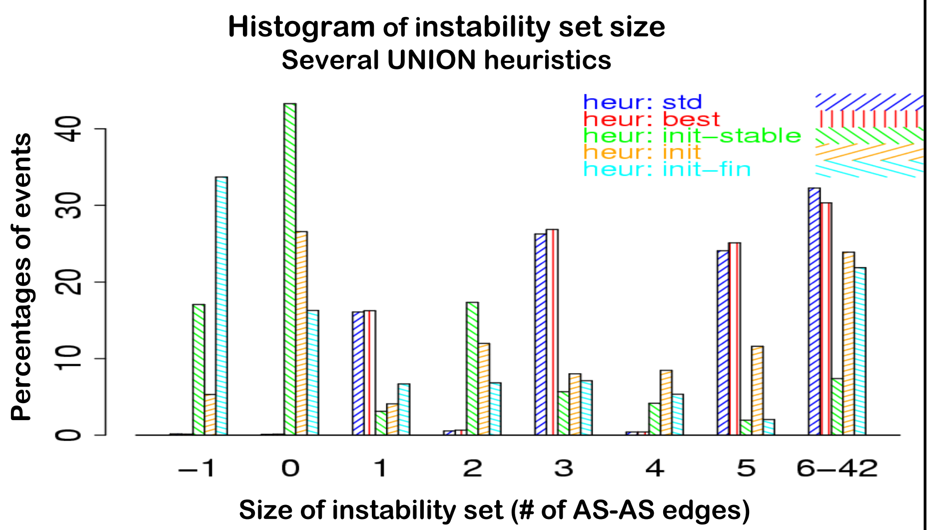

More aggressive heuristics are dangerous

UNION + INTERSECTION + GREEDY

- Zipf's law seems to apply to the distribution of correlated events across prefixes
- Single AS edge identified for 93.4% of prefixes
- Three AS edges identified for 97.2% of prefixes
- If restricted to at least 100 correlated prefixes:
  - Single AS edge identified for 96.3% of prefixes

Validation

- Syslog data of tier-1 vs. Greedy results
  - Crosscheck:
    Session reset on router $\Rightarrow$
    event within 5 minutes
  - Result:
    - Checked 35 events
    - Found 26 events $\Leftrightarrow$ 74% of the events

Summary

- Proposed methodology
  Time $\rightarrow$ Views $\rightarrow$ Prefixes
- Ideal-world study: Simulation
  - UNION / INTERSECTION heuristics
    $\leq$ 7 AS edges for 88% (10 obs.)
- Real-world study: Data analysis
  - UNION / INTERSECTION heuristics
    - Beacons: $\leq$ 3 AS edges for 76% (2 obs.)
    - All prefixes: $\leq$ 5 AS edges for 90% (5 obs.)
  - UNION / INTERSECTION / GREEDY heuristic
    - All prefixes: 1 AS edge for 93.4%
- Successful validation on tier-1 syslog data