An Active Approach to Measuring Routing Dynamics Induced by Autonomous Systems

Rocky K. C. Chang, The HK Polytechnic U.
Samantha Lo, Georgia Tech
Lorenzo Colitti, Google Inc. (with RIPE before)

APRICOT 2008, Taipei

Outline

- Problems
  - Routing dynamics induced by multihomed AS
  - BGP inbound traffic engineering
- Motivations
  - AS path prepending
- Active Measurement Methodology
  - RIPE NCC RIS
- Results and Analysis
- Conclusions and Future Works

A Doubly Homed AS

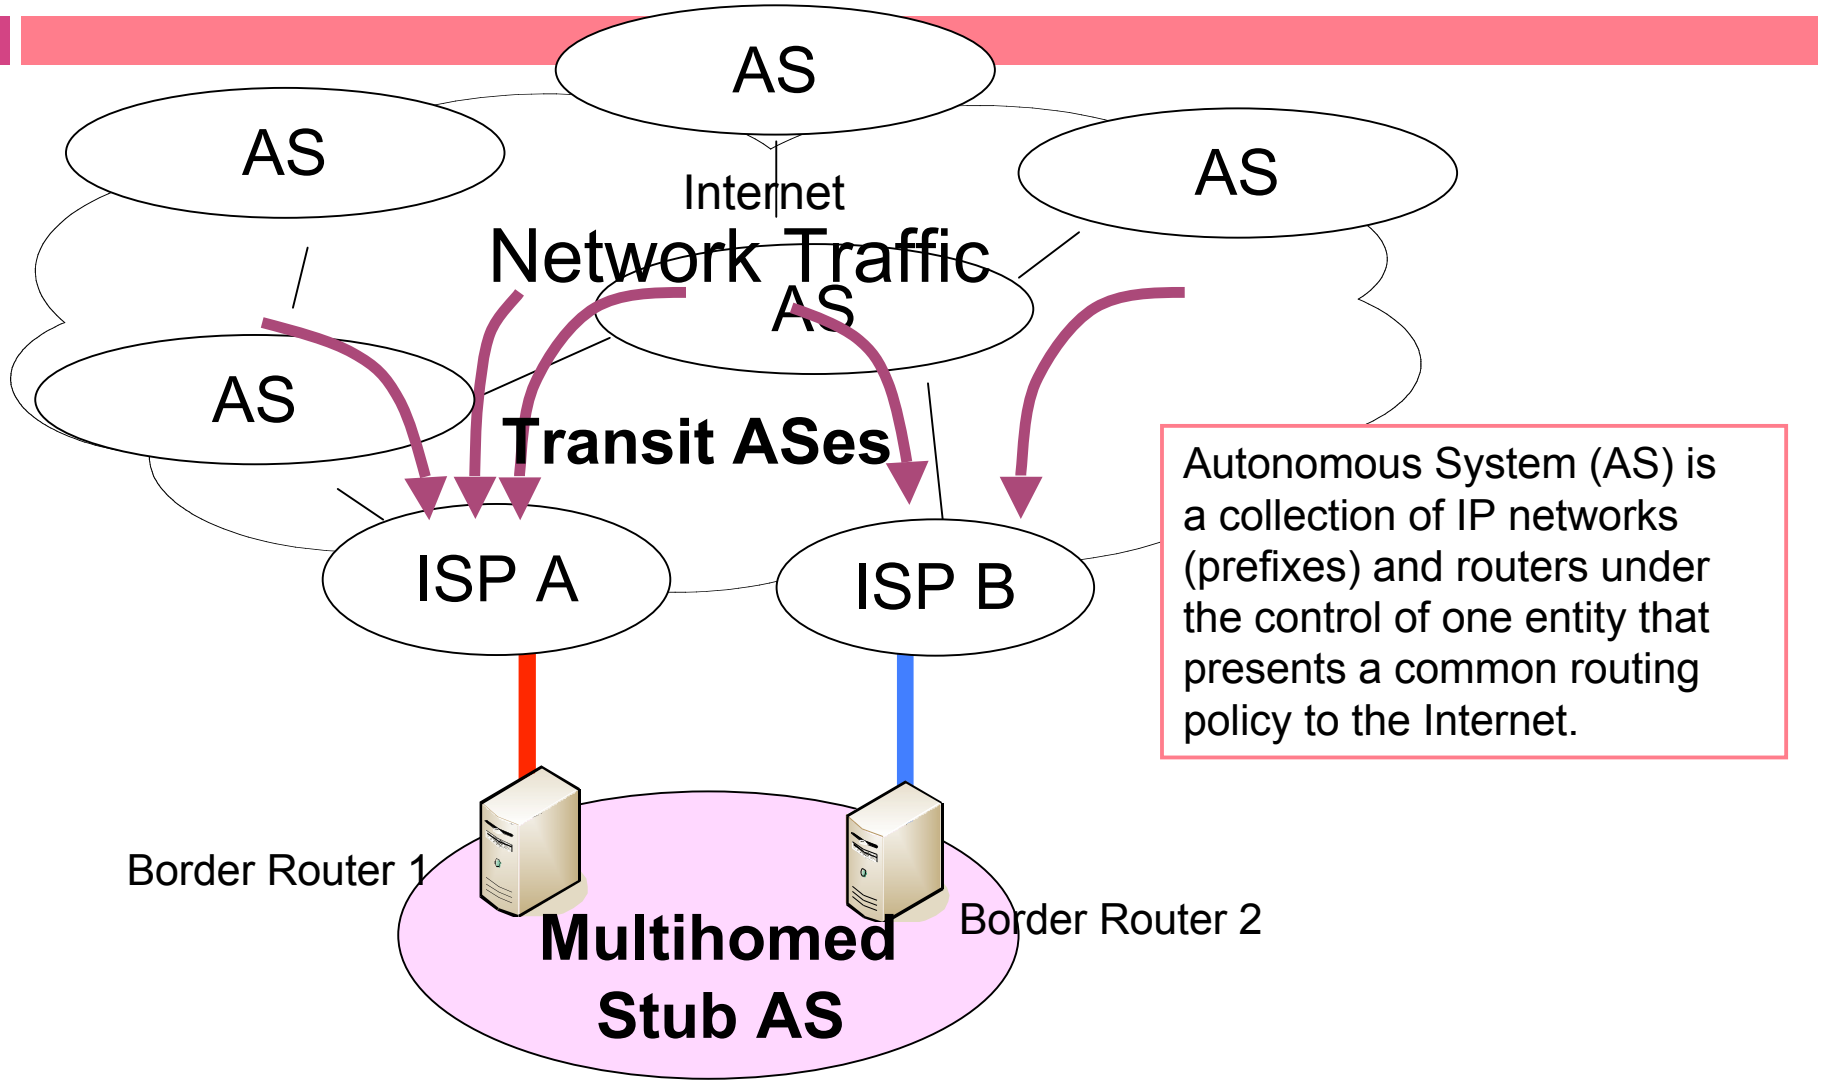

Autonomous System (AS) is a collection of IP networks (prefixes) and routers under the control of one entity that presents a common routing policy to the Internet.

BGP's AS PATH attribute

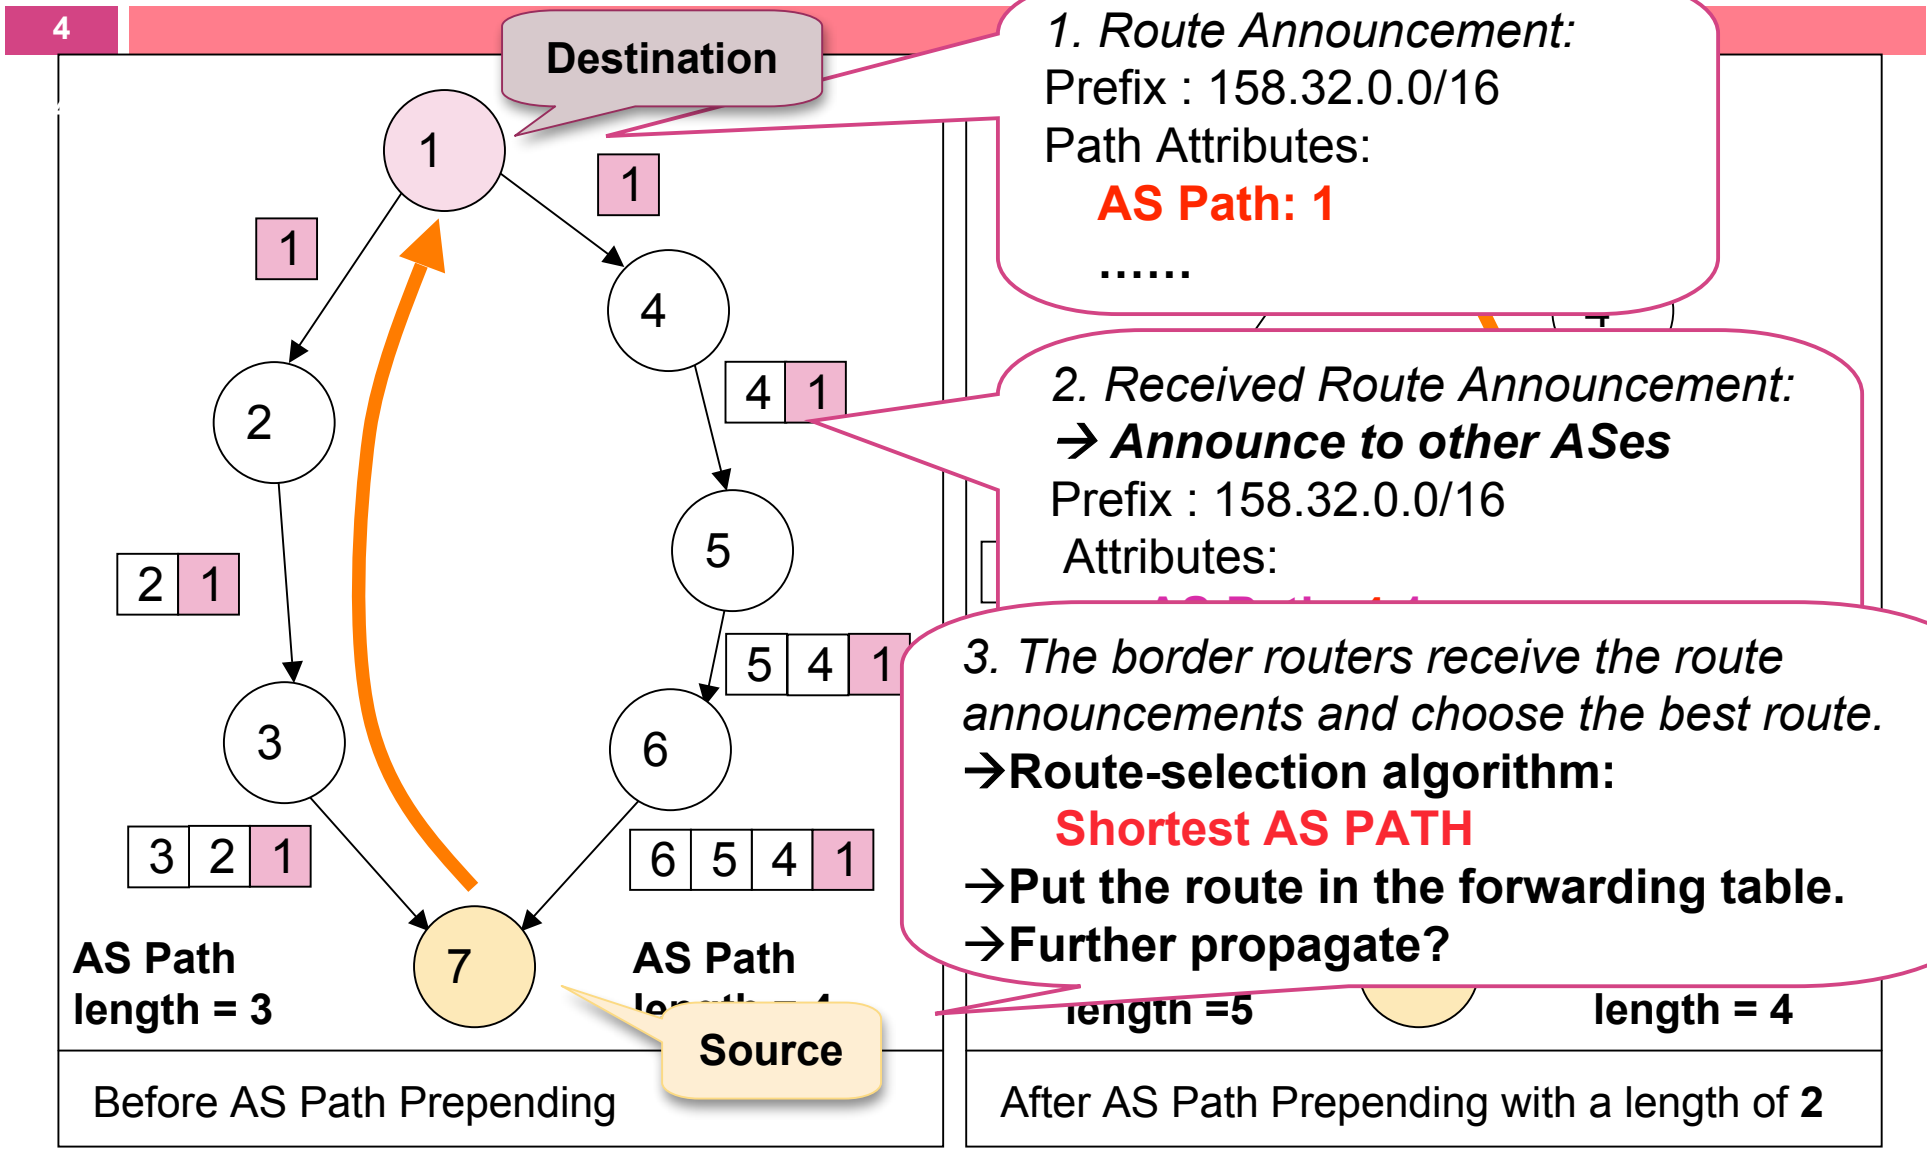

BGP routing decisions

1. Higher local preference
2. Shorter AS path
3. Lower origin type
4. Lower MED value
5. E-BGP over I-BGP routes
6. Lower IGP metric to next-hop
7. Lower BGP router ID

Outline

- Problems
  - Routing dynamics induced by multihomed AS
  - BGP inbound traffic engineering
- Motivations
  - AS path prepending
- Active Measurement Methodology
  - RIPE NCC RIS
- Results and Analysis
- Conclusions and Future Works

Motivations

- Routing dynamics affect end-to-end path performance.
  - “A measurement study on the impact of routing events on end-to-end Internet path performance”, SIGCOMM 2006.
- Routing dynamics are complex to understand and predict.
  - Routing policies
  - Topology
- Routing dynamics worsen the route convergence problems
  - “Route flap damping exacerbates Internet routing convergence”, SIGCOMM 2006.

How to study the routing dynamics induced by AS path prepending?

Contributions

- Developed an active measurement methodology to understand the prepending method.
- Deployed the methodology to RIPE.
- Insights and findings, e.g.,
  - A different classification of ASes
  - A small # ASes responsible for a large # of route changes induced by prepending
  - The convergence time prolonged

Previous Works

- Prevalence of prepending
  - Prepending is often used for backup links (Feamster et al. 2001)
  - Large amount of prepended routes (Wang et al. 2003)

- Methodology
  - Beacon prefix (Mao et al. 2003)
  - Active approach
    - AutoPrepend (Chang and Lo 2005)
    - Black box approach (Quoitin et al. 2005)
  - Optimization
    - Gao et al. 2005
    - Battista et al. 2005

Outline

- Problems
  - Routing dynamics induced by multihomed AS
  - BGP inbound traffic engineering
- Motivations
  - AS path prepending
- Active Measurement Methodology
  - RIPE NCC RIS
- Results and Analysis
- Conclusions and Future Works

Active Measurement

- Inject a BGP route with a prepending length to the Internet.
- Observe route changes induced by the route.
- Why not passive measurement?
  - Uncontrolled experiments
  - Outcomes → causes ?

Passive measurement results

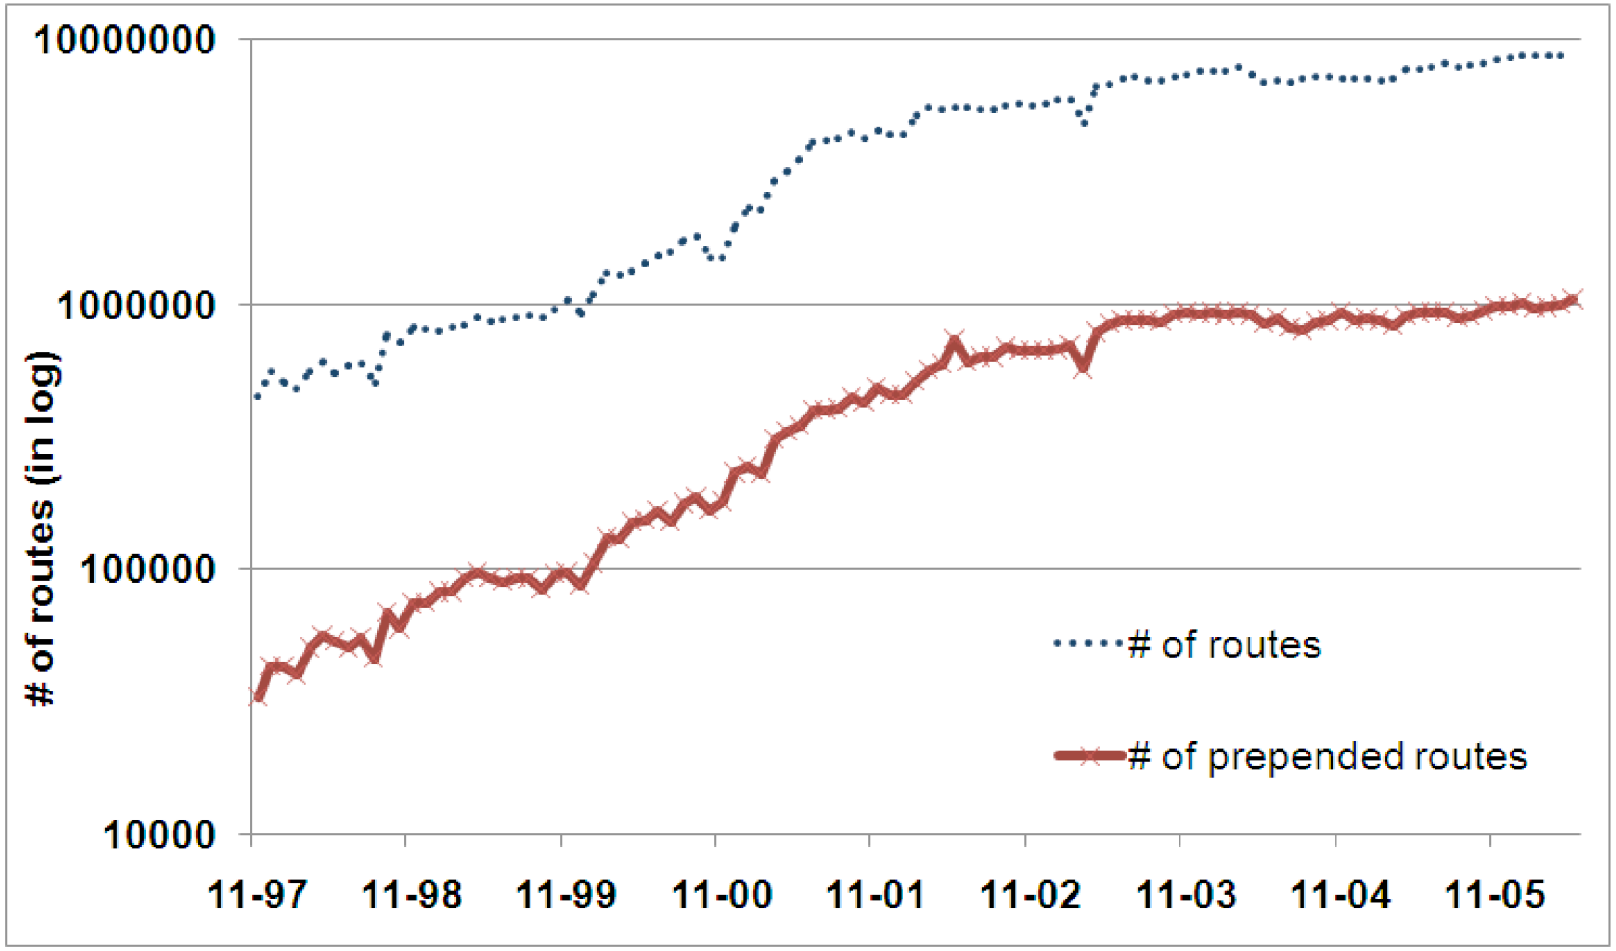

Active Measurement

1. Announce the beacon and control prefixes with 0 prepending length to the upstream ASes.

Control prefix: to observe the normal routing dynamic without prepending.

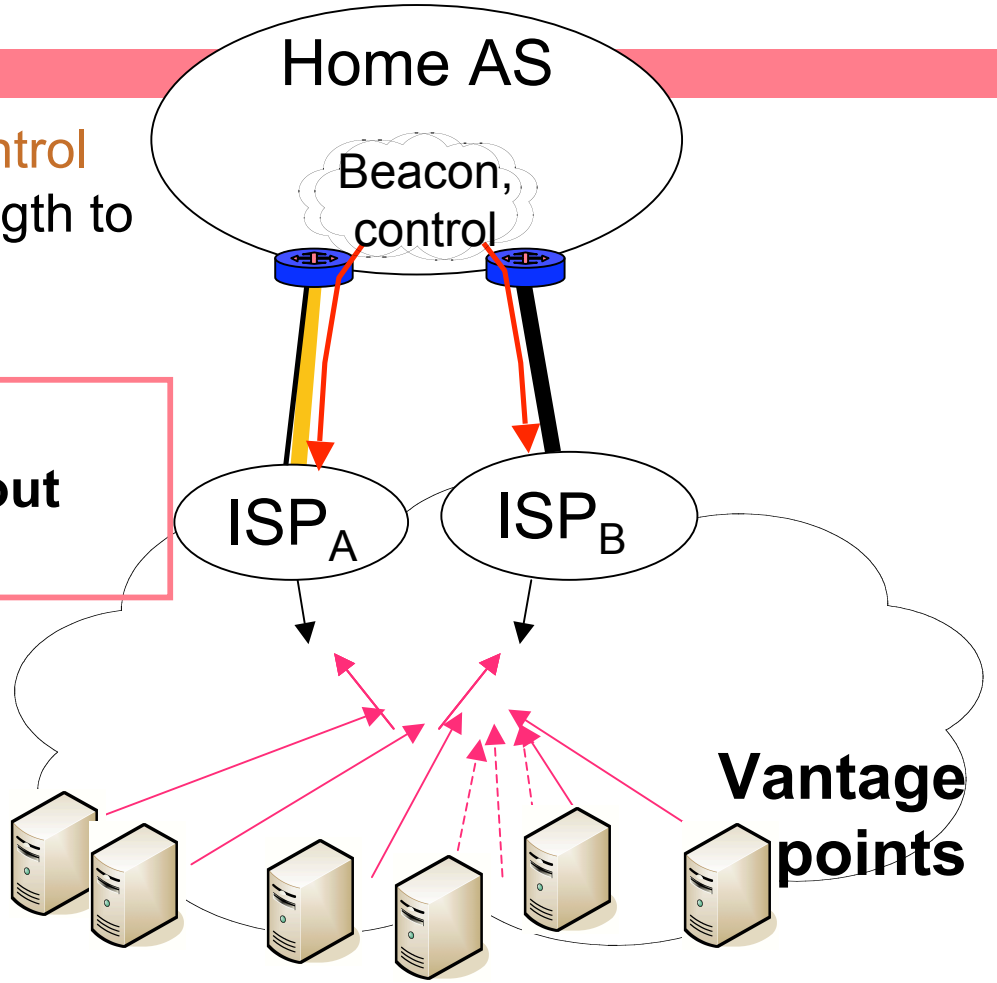

2. Query vantage points (VPs) and capture the BGP routes of the beacon and control prefixes.

3. Repeat steps 1-2 with different prepending lengths of beacon prefix on ISP$_A$ link.

Our Setup

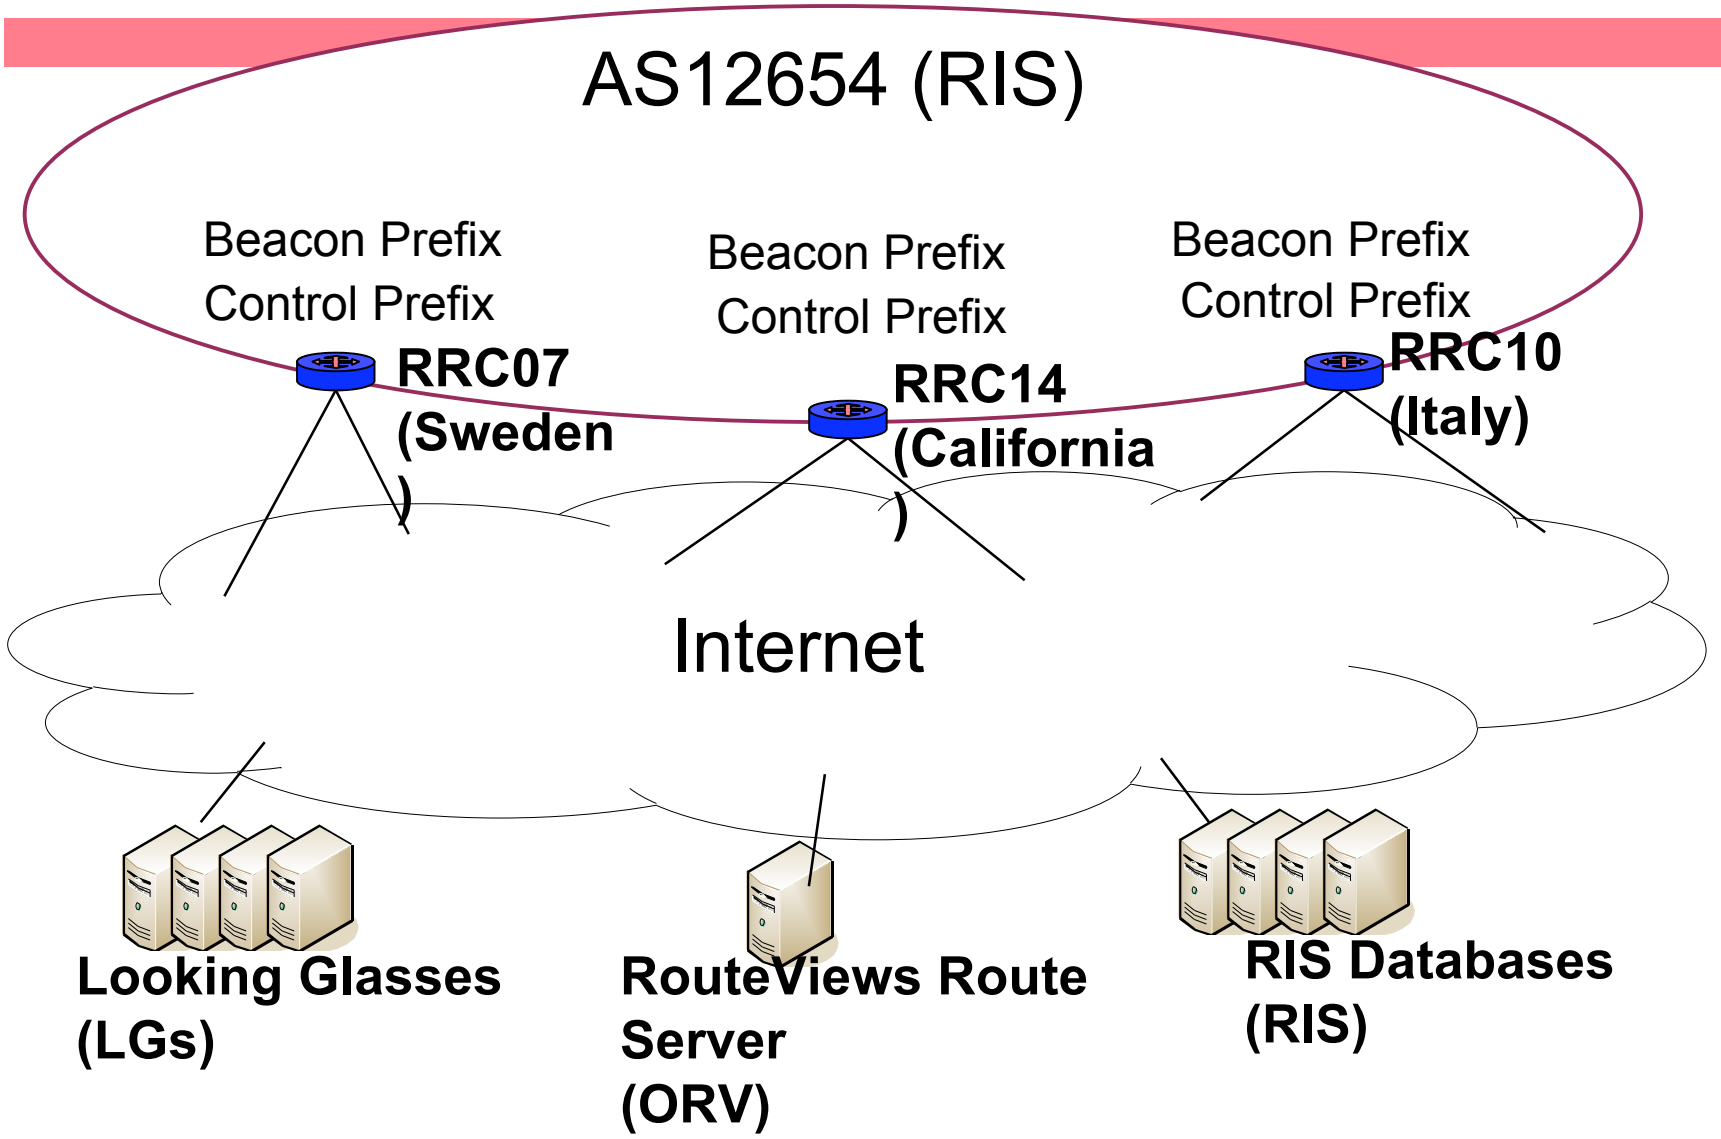
AS12654 (RIS)
Beacon Prefix
Control Prefix
RRC07
(Sweden
)
Beacon Prefix
Control Prefix
RRC14
(California
)
Beacon Prefix
Control Prefix
RRC10
(Italy)
Internet
Looking Glasses
(LGs)
RouteViews Route
Server
(ORV)
RIS Databases
(RIS)

Measurement Time Table

RRC	Upstream providers	Beacon prefix (Control prefix)	Measurement date	Max prepending length
RRC07	AS13237 (LAMBDANET) + AS16150 (PORT80)	84.205.73.0/24 (84.205.88.0/24)	8-9 May, 2006 (update every 2 hours)	6
RRC14	AS6939 (HURRICANE) + AS6762 (SEABONE-NET)	84.205.89.0/24 (84.205.95.0/24)		6
RRC10	AS12779 (ITGATE) + AS1299 (TELIANET)	84.205.88.0/24 (84.205.73.0/24)	13-15 May, 2006 (update every 3 hours)	10

“+” indicates the prepended link (PL)

Outline

- Problems
  - Routing dynamics induced by multihomed AS
  - BGP inbound traffic engineering
- Motivations
  - AS path prepending
- Active Measurement Methodology
  - RIPE NCC RIS
- Results and Analysis
- Conclusions and Future Works

Overall results

Route changes for RRC7

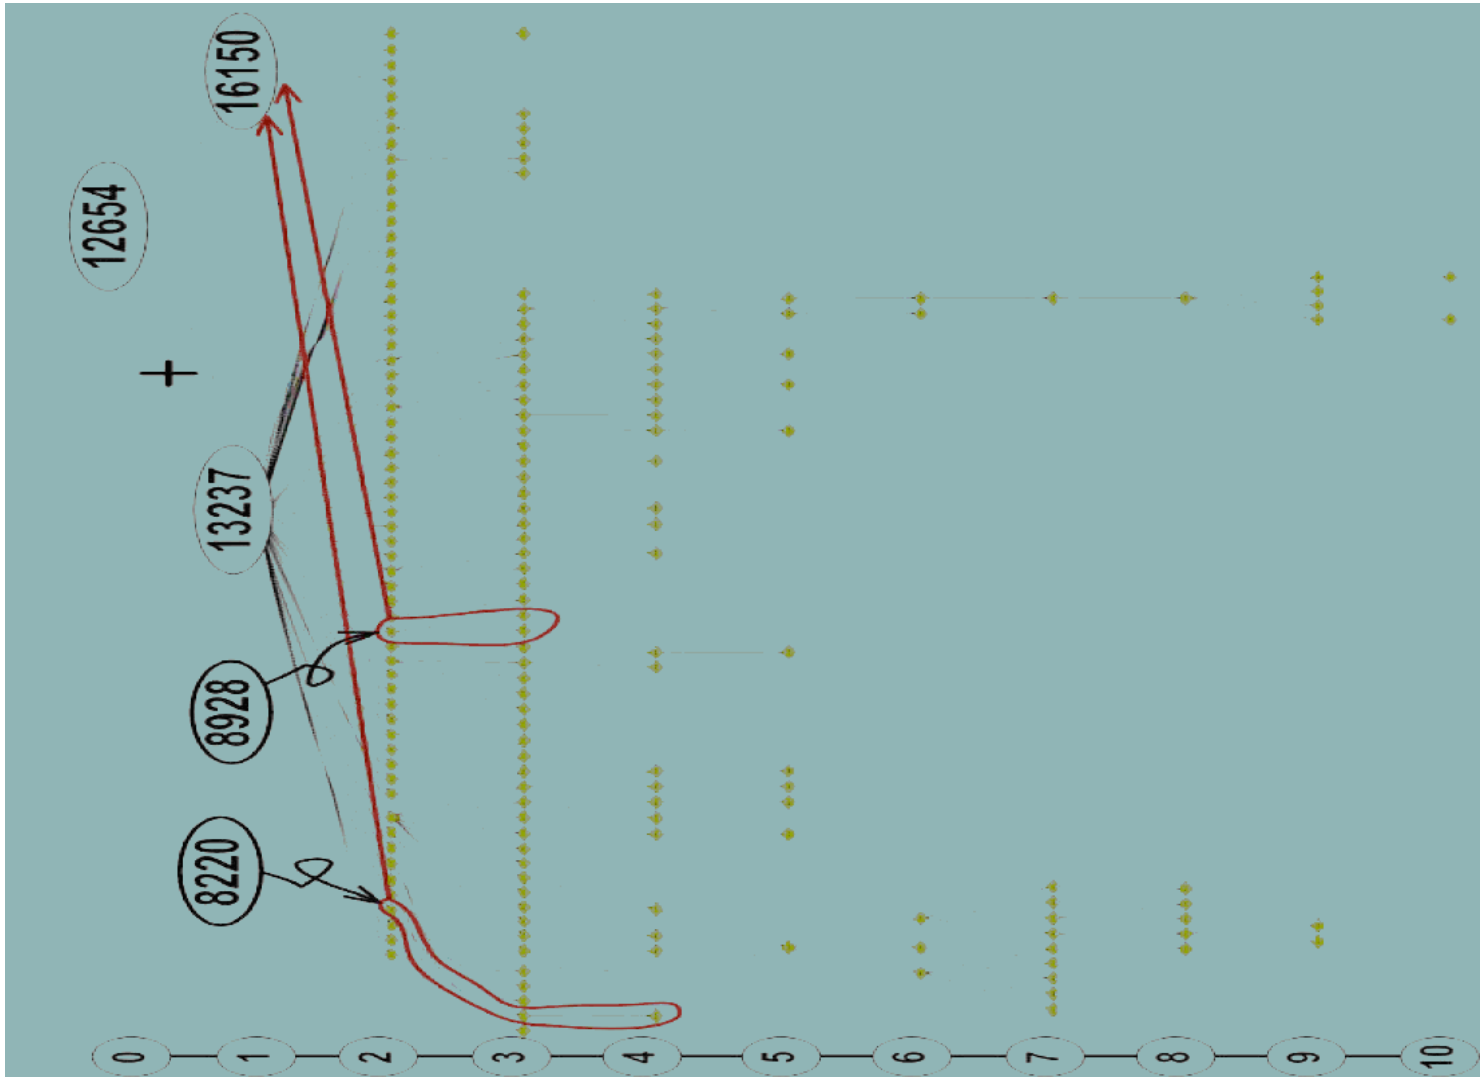

More time needed to converge?

RRC07

- Prepend on **more** VP side.
- Not a big effect
- Beacon prefix was not available in some VPs after we waited for 2 hours.

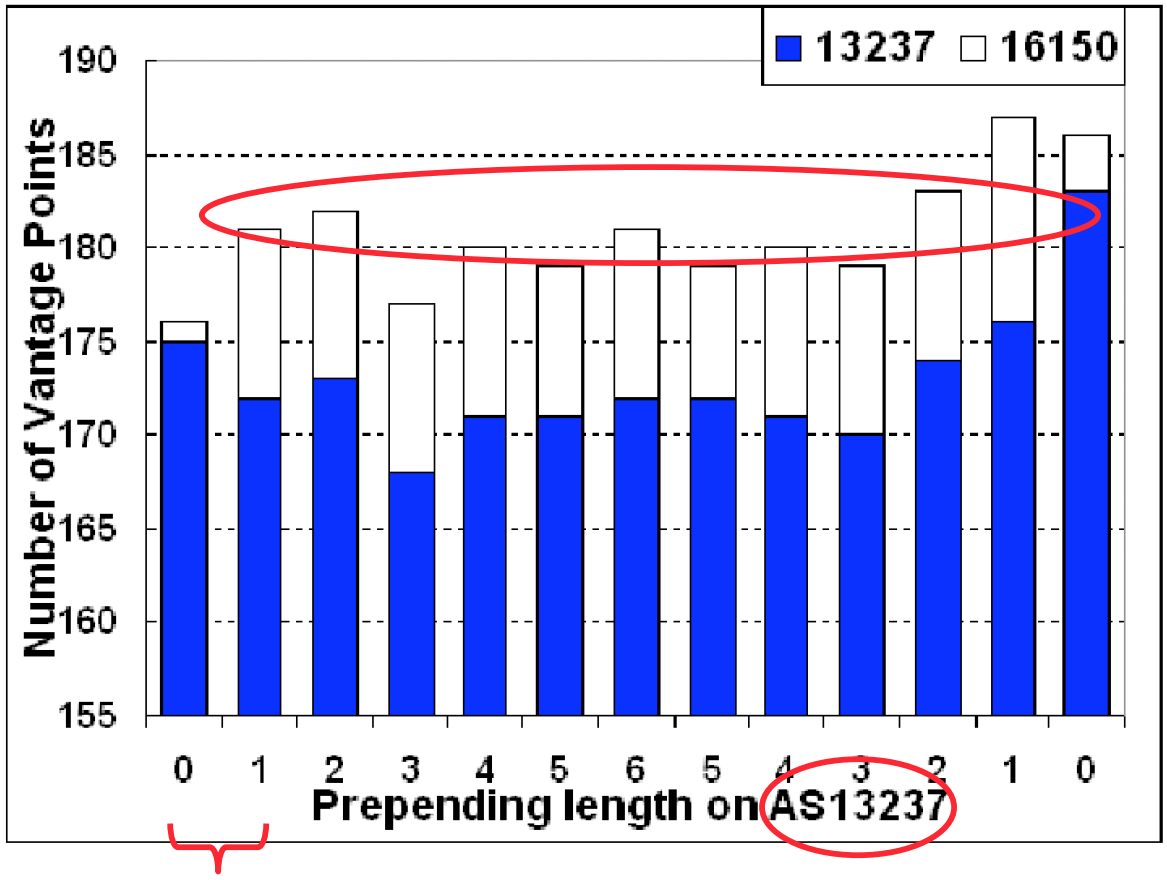

Route changes for RRC14

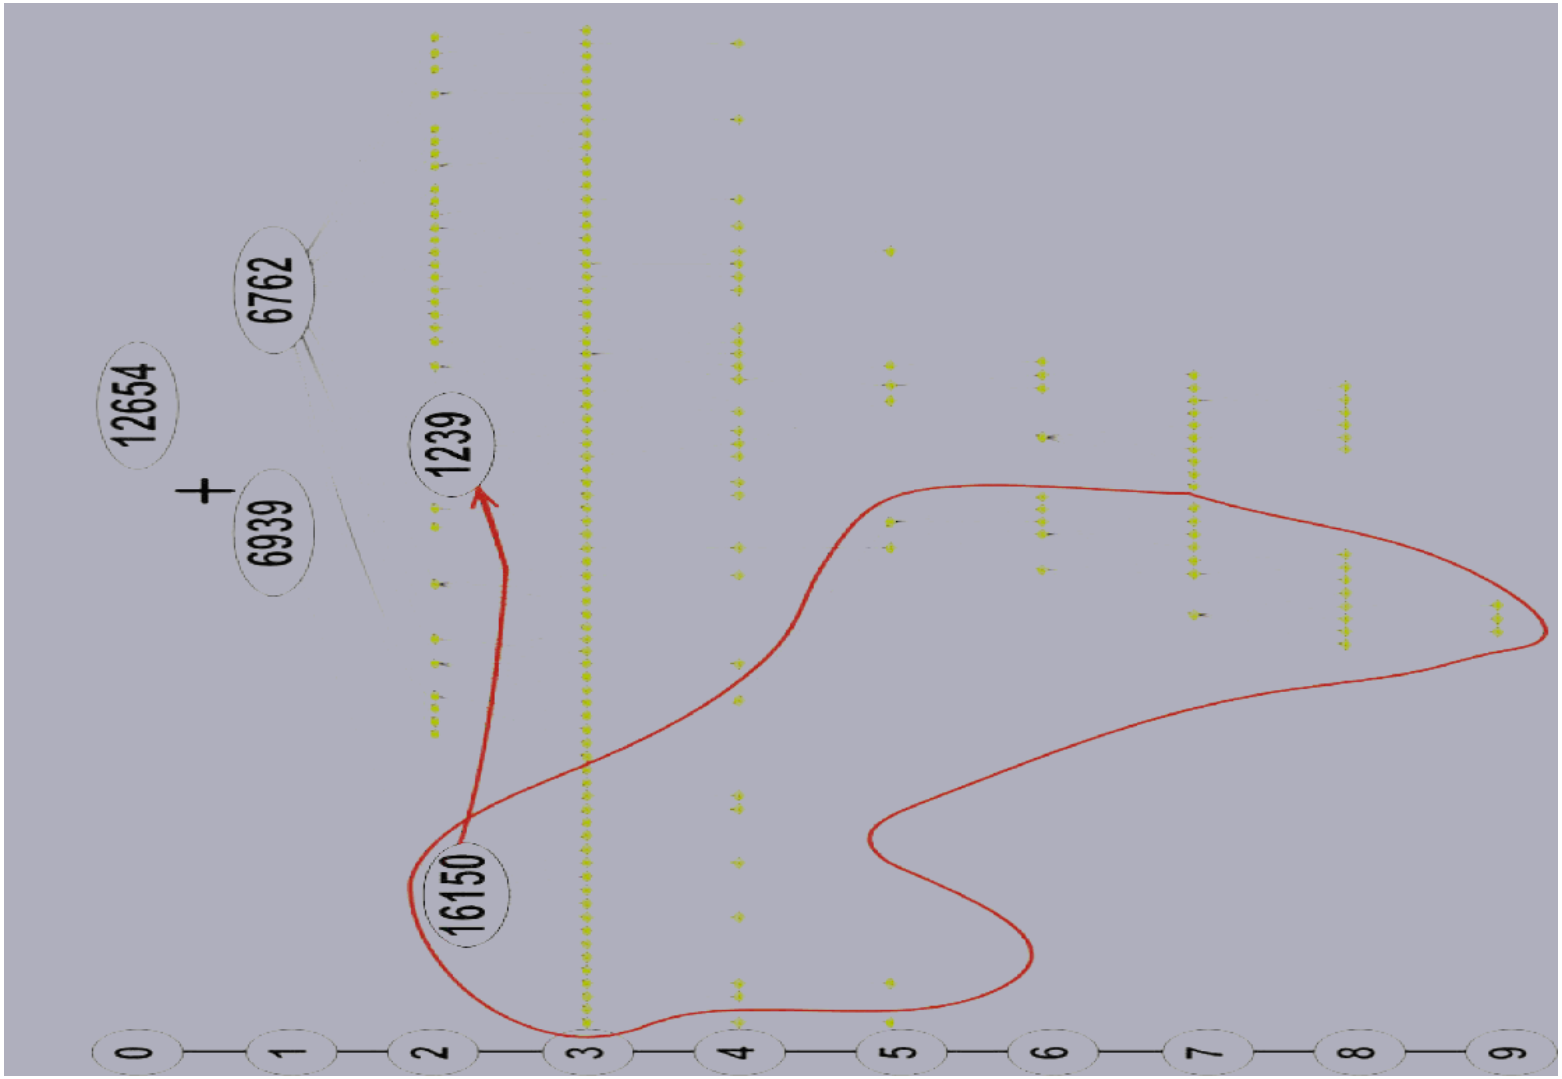

Shift all traffic to one side

RRC14

- Prepend on **less** VP side.
- All the routes changed at length 2 (Including the upstream AS6939).

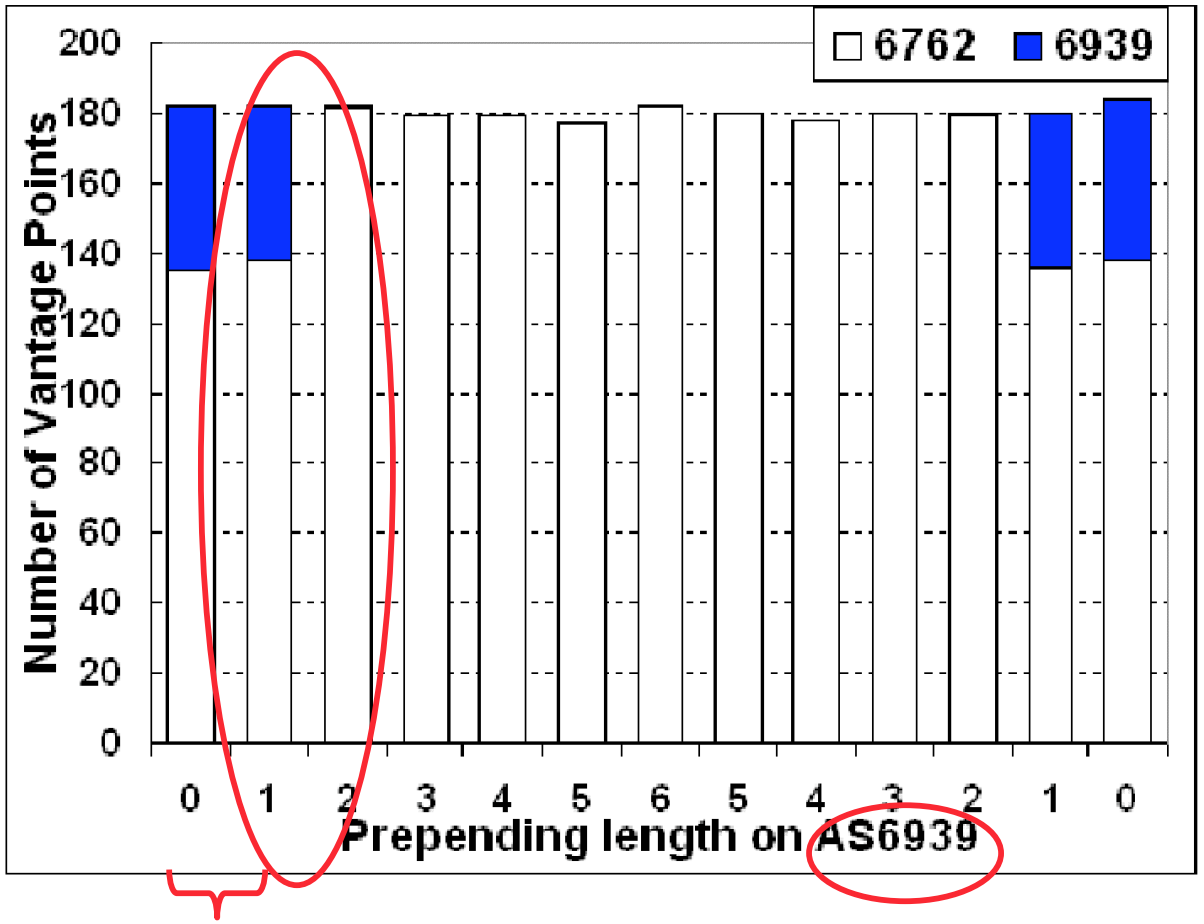

Update every 2 hr

Route changes for RRC10

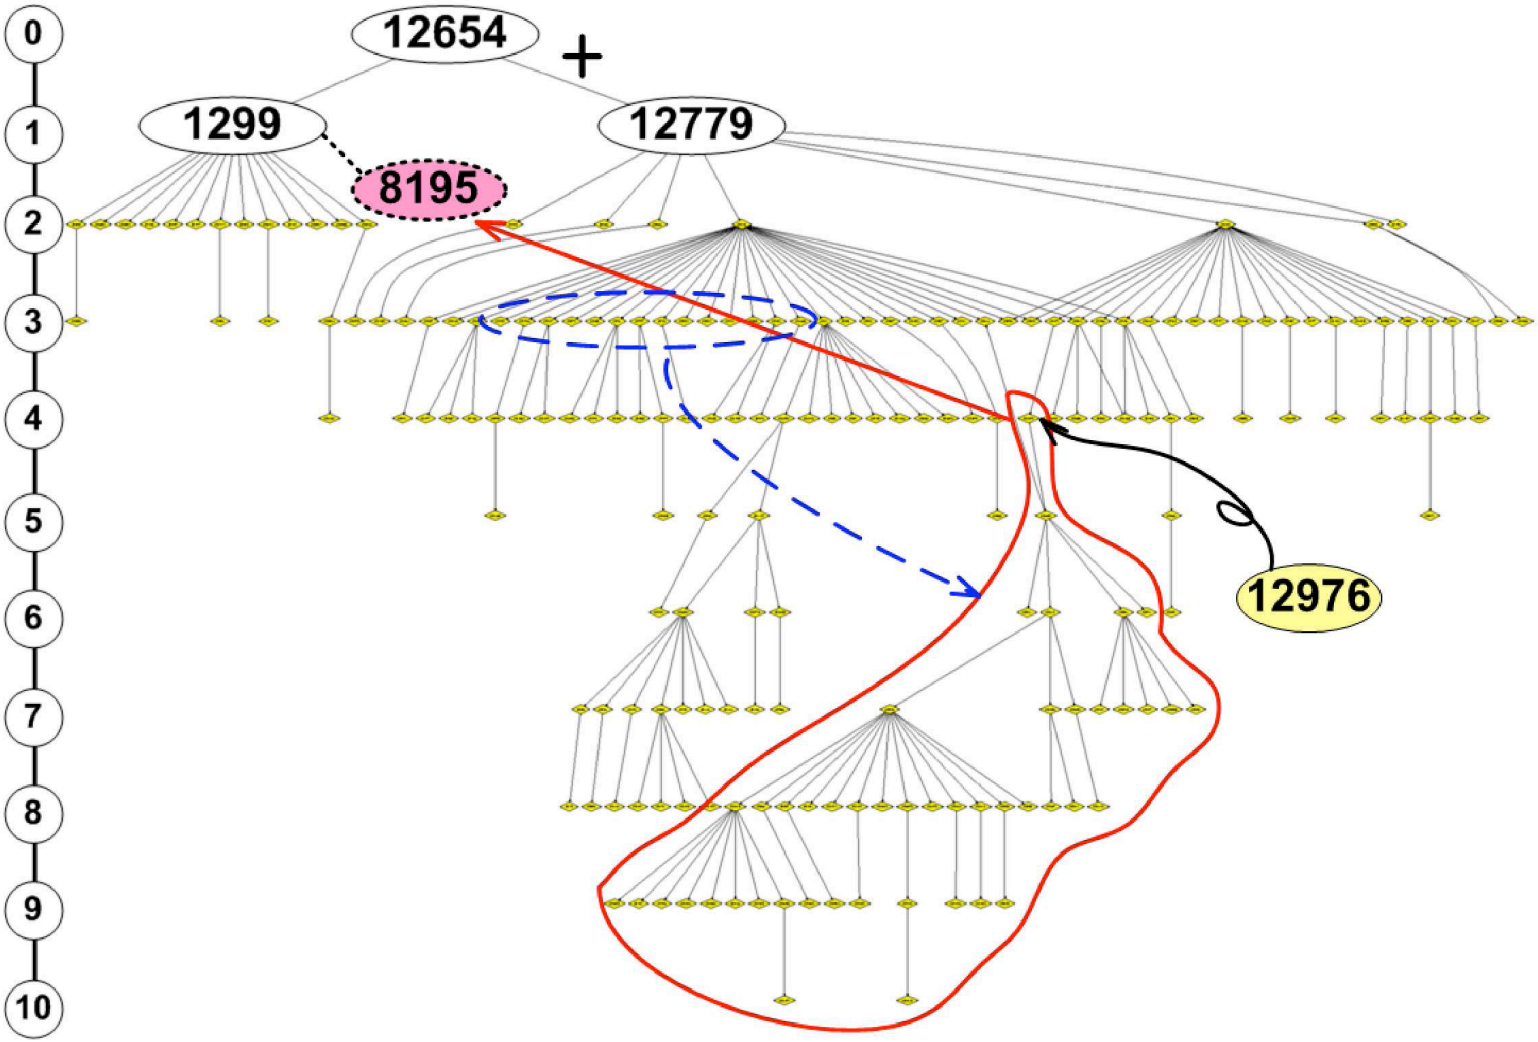

Gradual change up to length 10

RRC10

- Prepend on **more** VP side
- Greatest change on length 1
- Gradual change up to length 10
- Based on the high-impact direct responsive AS

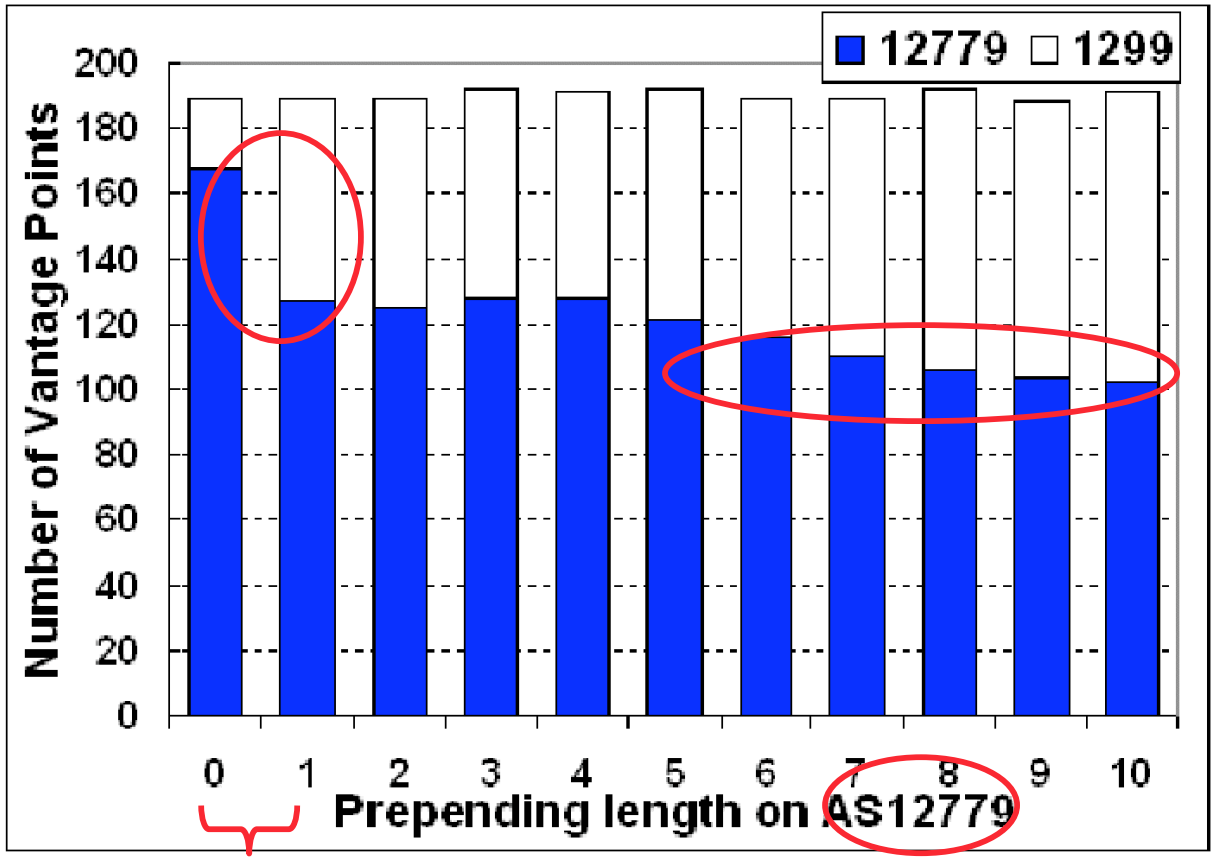

Update every 3 hr

Q1. Who were responsible for the route changes?

Three kinds of ASes

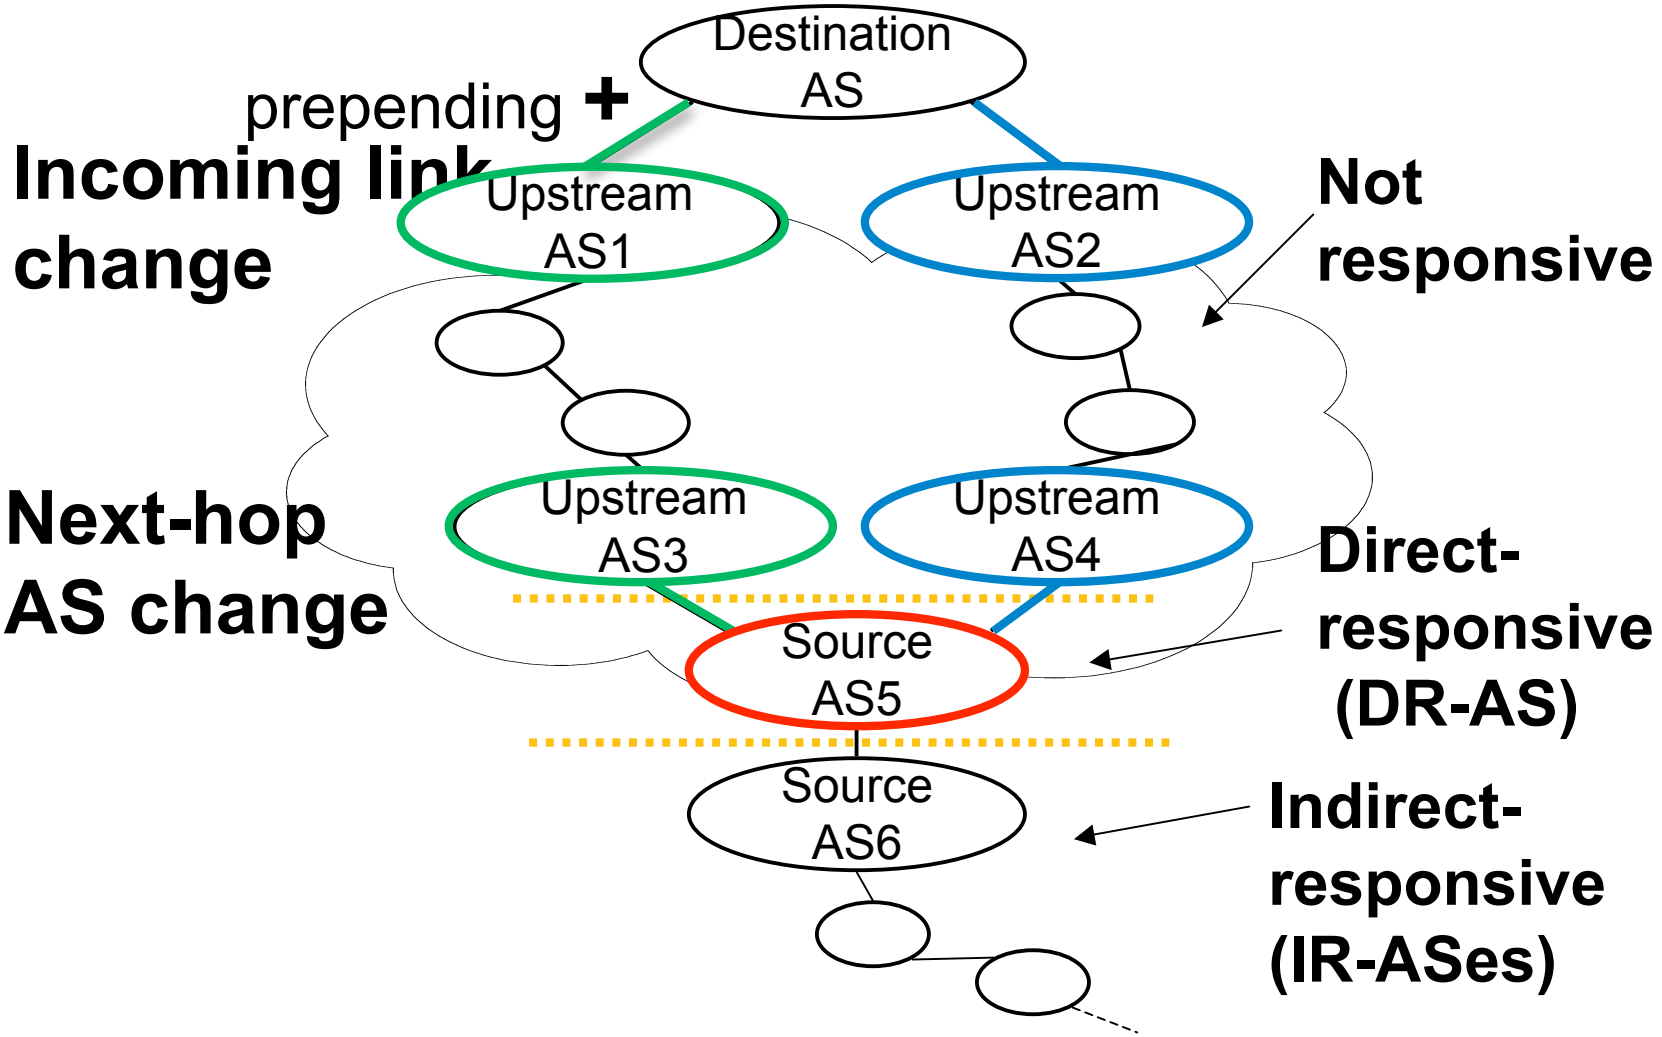

A small # of DR-ASes is responsible

	RRC07	**RRC14**	**RRC10**
# of ASes used the prepended side before	203	55	180
# of responsive ASes	12 (6%)	55 (100%)	68 (38%)
DR-ASes	7 (58.3%)	7 (12.7%)	22 (32.4%)

Q2. Is there any dominant AS that causes route change?

High-impact AS – AS12976 in RRC10

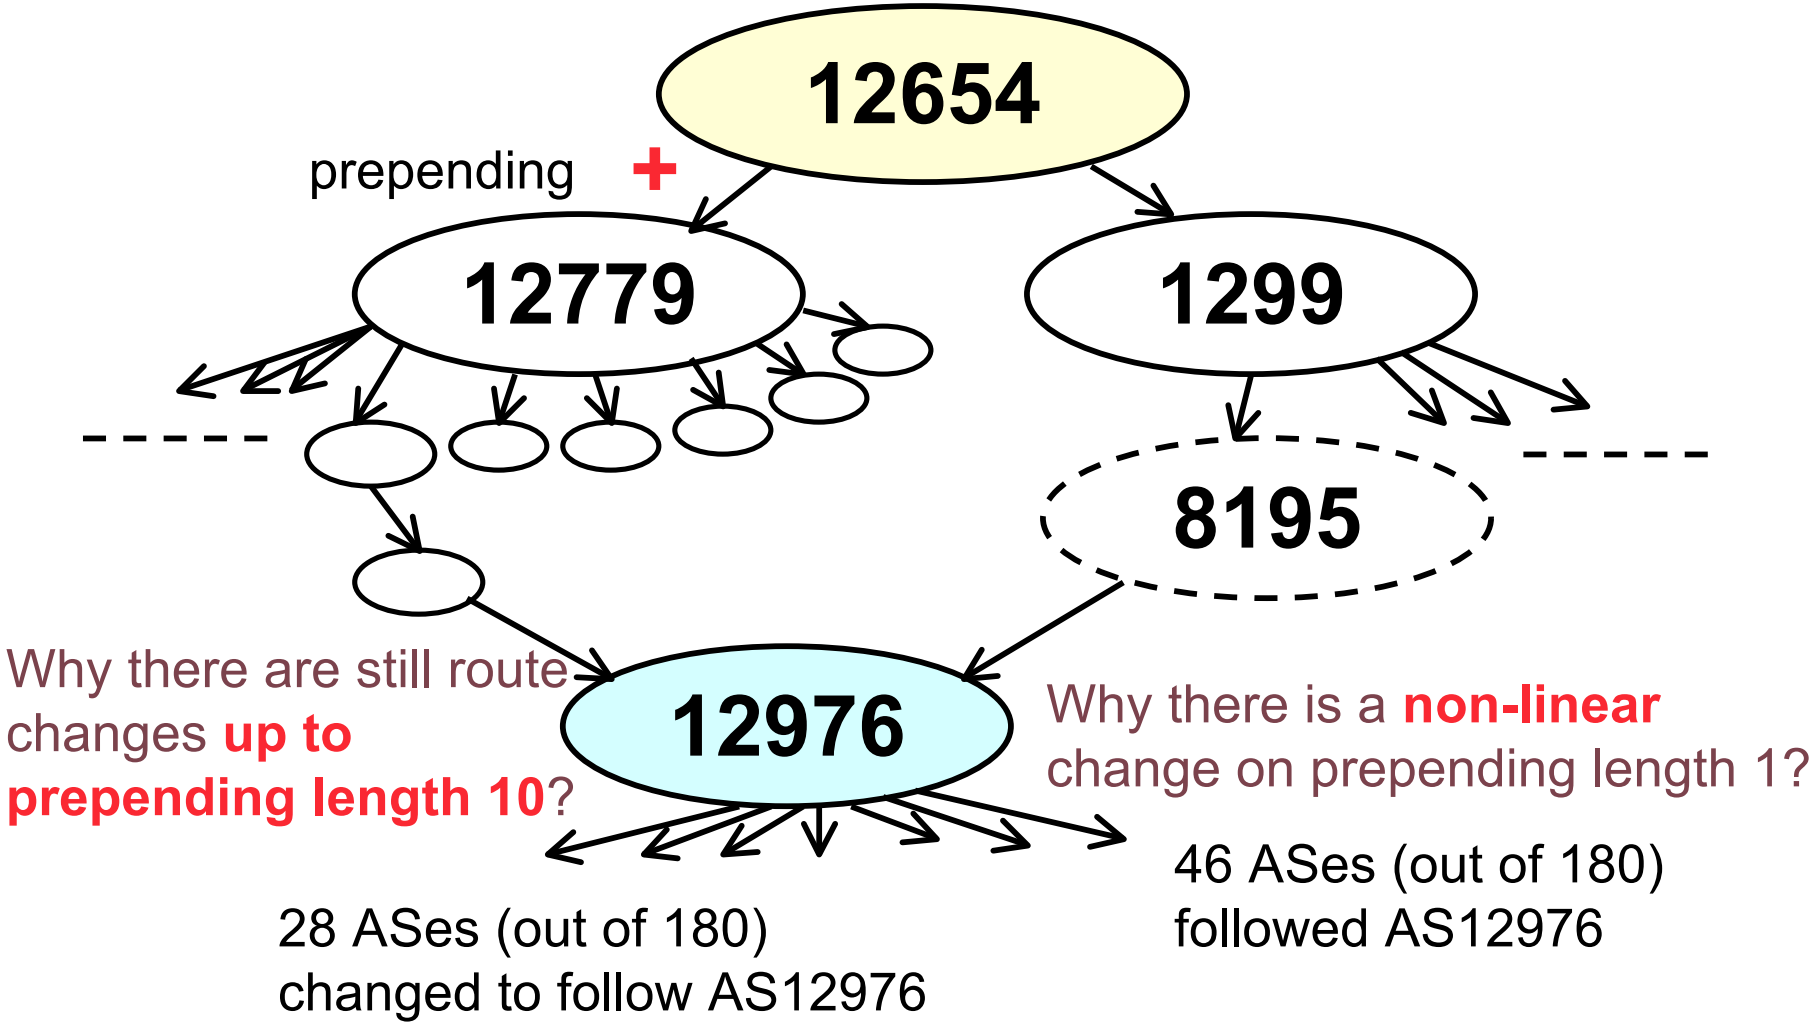

Q.3 Does prepending reveal hidden nodes and links?

Reveal hidden policies in RRC10

- Before prepending on AS12799
  - **16150→8432→2118→20483→12976→1273→1239→12779→12654**
- After prepending once on AS12799
  - **16150→8432→2118→20483→12976→8195→1299→1299→1299→12654**
- A new prepending undid an old prepending!

Undo AS1299's preference

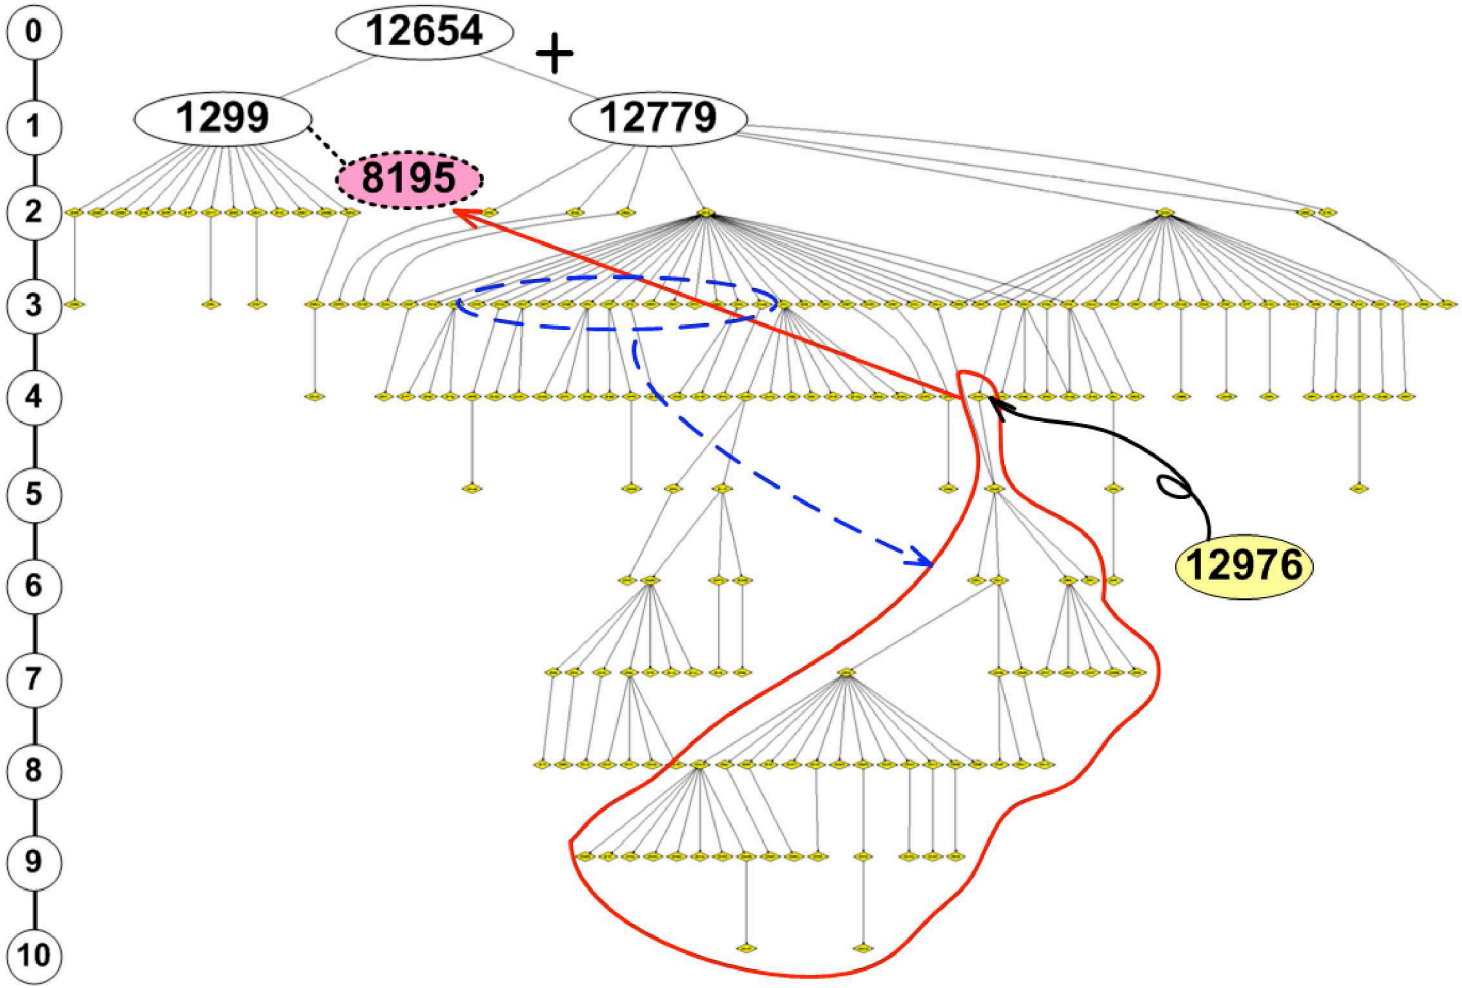

Conclusions and Future Works

- Proposed an active measurement methodology to study routing dynamics.
- Performed measurement experiments at RIPE.
  - DatCat (http://imdc.datcat.org/)
- Reported a number of unreported findings.
- Current works:
  - Measuring the interactions between data plane and control plane
  - Measuring for more than two ISP connections
  - Applying the methodology to traffic engineering

Thanks!