Understanding Global Internet Routing Stability Using Link Weight

Mohit Lad, Jonathan Park, Tiziana Refice, Lixia Zhang

Feburary 2008

Goal

- Study Internet routing stability by aggregating information from multiple vantage points
  - identify where routing changes occur
    - Locate instabilities, not explain why they occur
  - Identify repeating instabilities over time
  - Identify the scope of routing events
    - How big a splash each incident makes

Challenges: Multi-dimensional data

- Large number of destinations (> 250K)
- Multiple vantage points
  - Each sees a 2-dimensional space of the above
  - Different vantage points see very different pictures
- Changes over time

Our Approach

- Large number of destinations
  - Measure "link weight changes" to catch big routing changes
- Multiple vantage points
  - Measure link weight changes from each vantage point
- Changes over time
  - Slice time into bins and investigate each bin
- Apply Principal Component Analysis (PCA) to identify biggest routing change events
  - ***big:*** a combination of (1)the magnitude of changes; (2)number of monitors seeing the change
  - Later separate out which is which

Link-weight

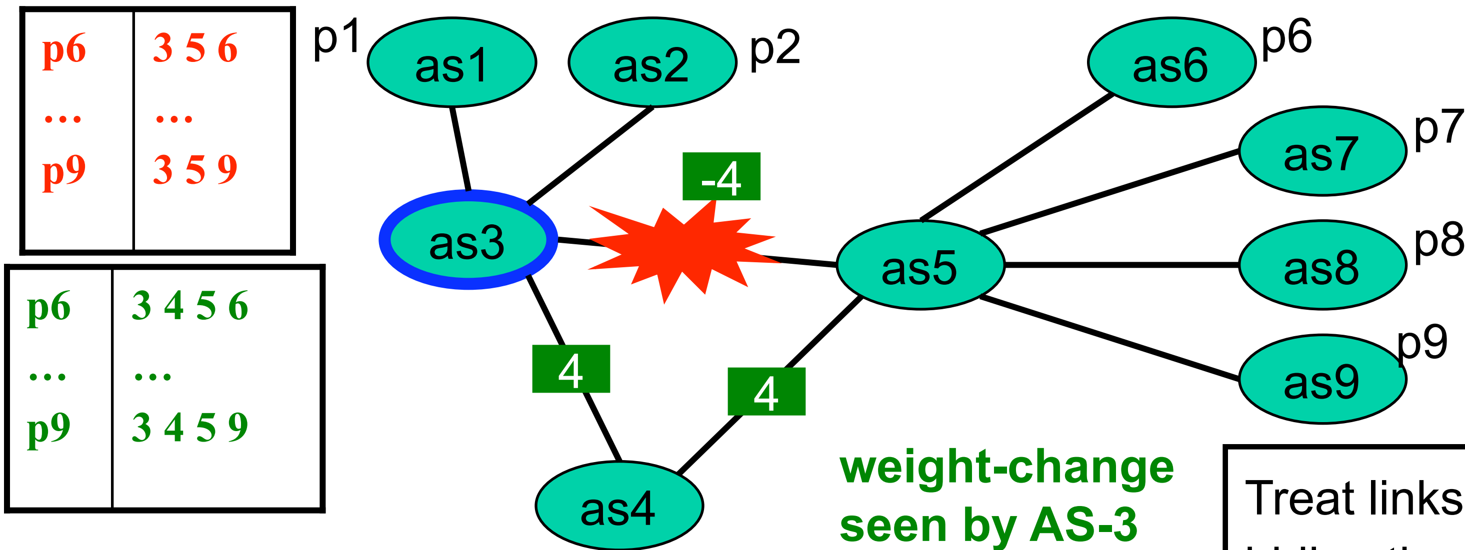

Treat links as bidirectional, A-B is the same as B-A.

By looking at link weight changes one can

- capture **common behavior across multiple prefixes**

  E.g. affected routes seen by AS 3 share a common link as3-as5.

- capture **common behavior across multiple monitors**

  Eg. AS-1 sees routes to p6,p7,p8,p9 affected, while AS-6 sees routes to p1,p2 affected, yet looking both see weight changes on link as3-as5.

Computing Weight changes

- Start with a full routing table: compute link weight w(a,b) for each AS link seen by each monitor
- Group BGP updates into time bins of every T seconds
- For each time bin, each AS link, seen by each monitor
  - If a route change results in a LW change, record the prefix
  - $\delta(a,b)$ = the total prefixes with weight changes on link (a,b).
- Resulting matrix: links=rows, monitors=columns
  - $X_{i,j}$: Weight change on link i seen by monitor j.
- Input into PCA

Need for PCA

- Given link weight changes seen by multiple monitors, PCA helps evaluate combined results

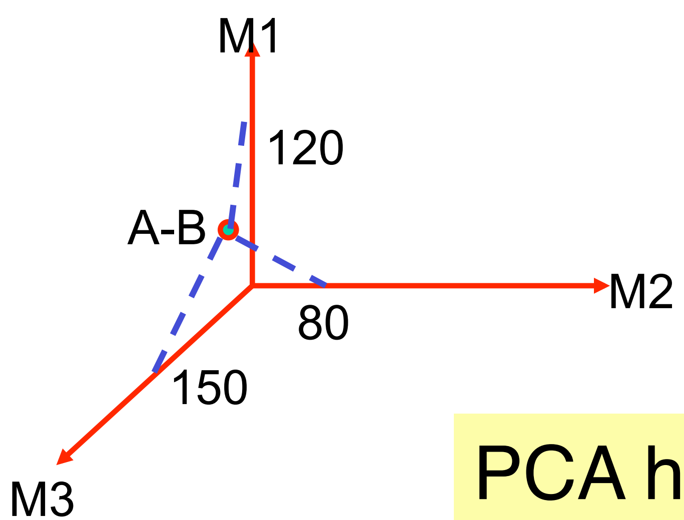

PCA helps reduce the dimensions by creating new axis representing linear combination of multiple monitors views

Example of how PCA helps

View of monitor 1273

View of monitor14608

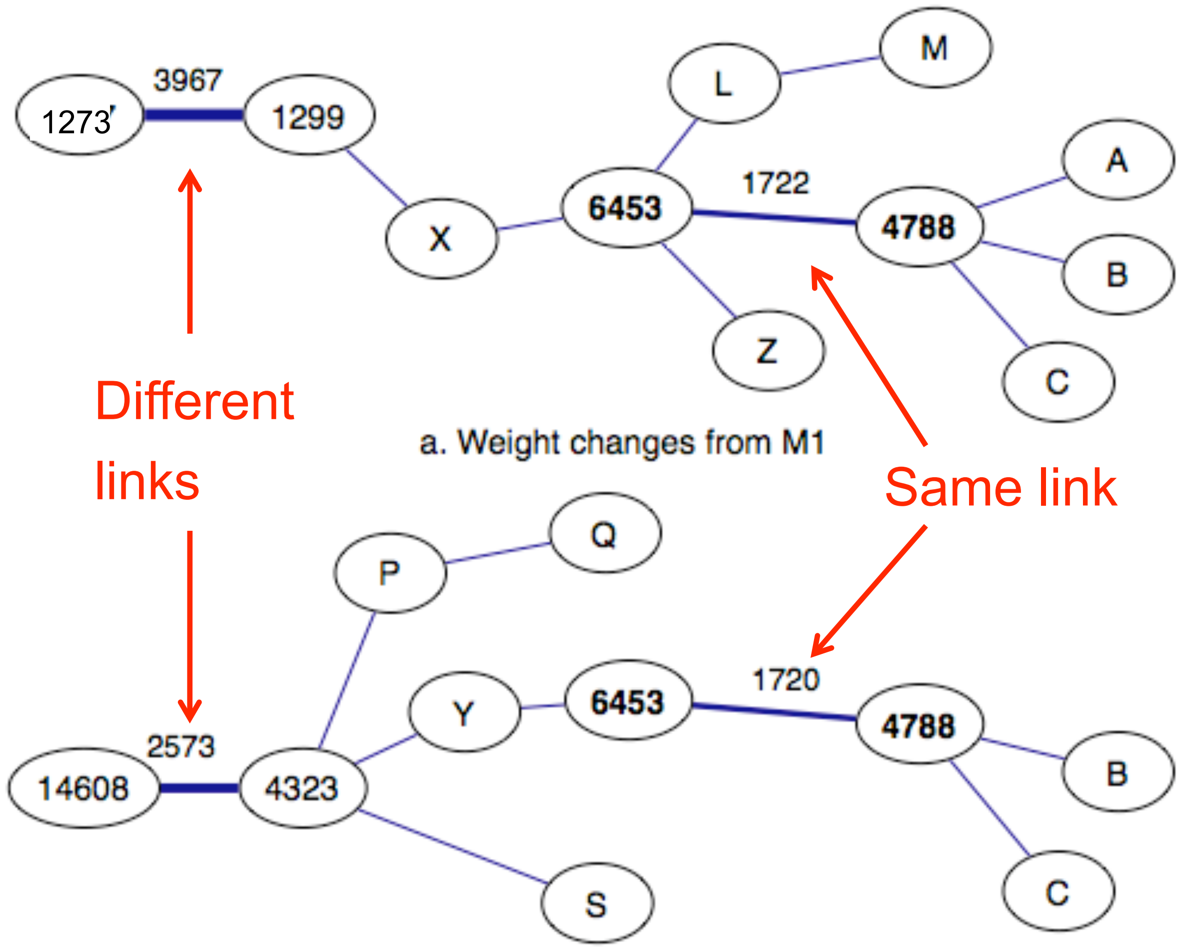

a. Weight changes from M1

b. Weight changes from M2

Example of how PCA helps

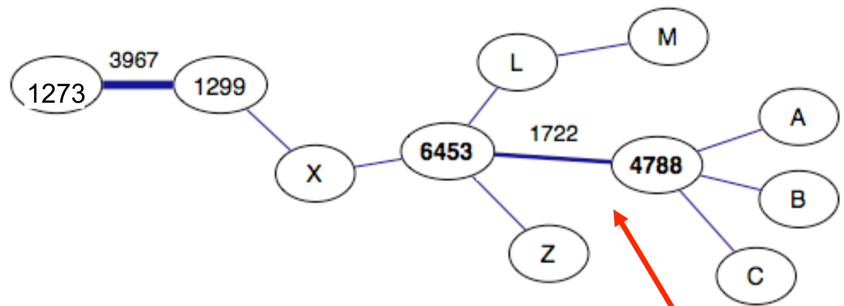

a. Weight changes from M1

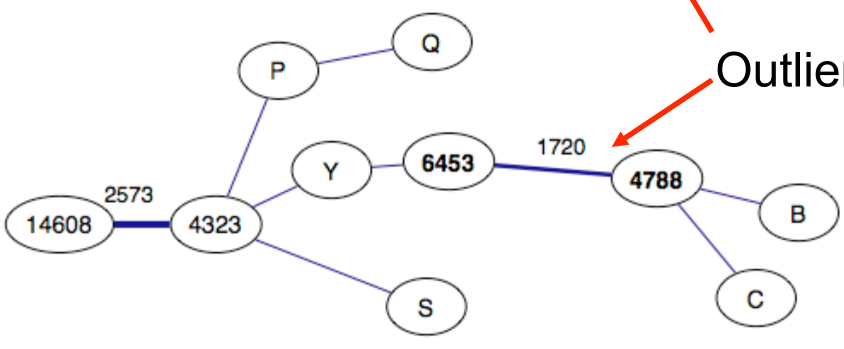

b. Weight changes from M2

Outlier link

PCA output with 30 monitors

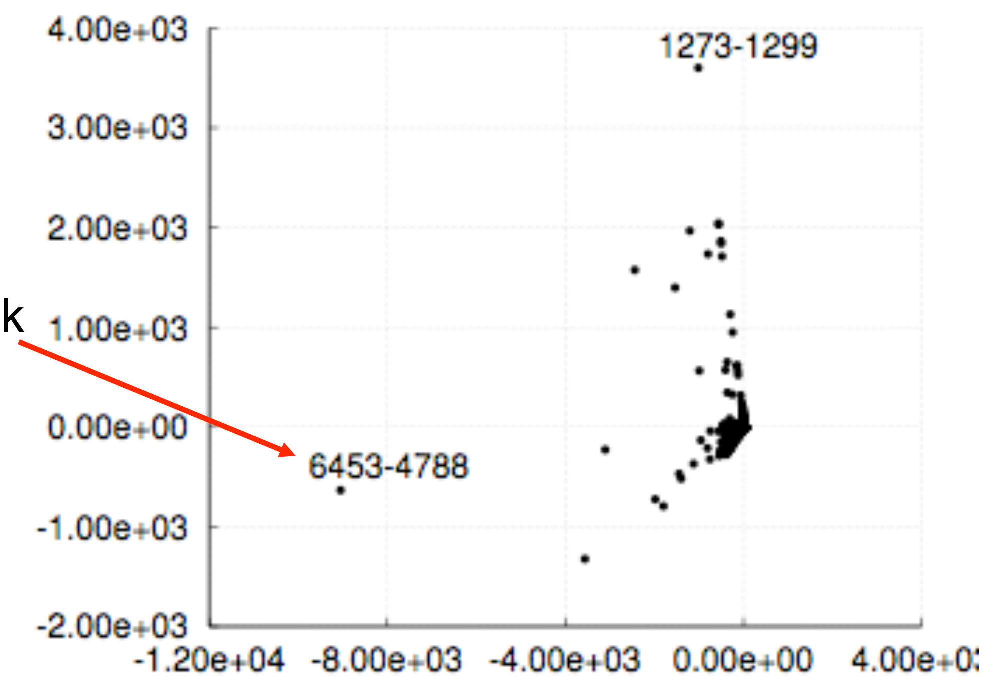

Each axis represents a combination of multiple monitors

Common outlier stands out after combining views from multiple points

Applying to long term BGP data

- Data set: RouteViews and RIPE
  - Chose a subset of 30 monitors that do not share large amounts of routes
- Duration: Jan-Dec 2007
- Group updates into 10-min bins
- Apply PCA to data in each bin to find outlier links
  - If a time bin shows noticeable magnitude changes, we call it an event

Questions from data

Q1:Are there any big events during the 1 year period?

- What is the scope of each event (how many monitors see big routing changes?)

Q2: are there links that appear repeatedly as outliers?

- What is the scope of the event (how many monitors see big routing changes?)

Magnitude analysis

Frequency analysis

Change Magnitude Distribution

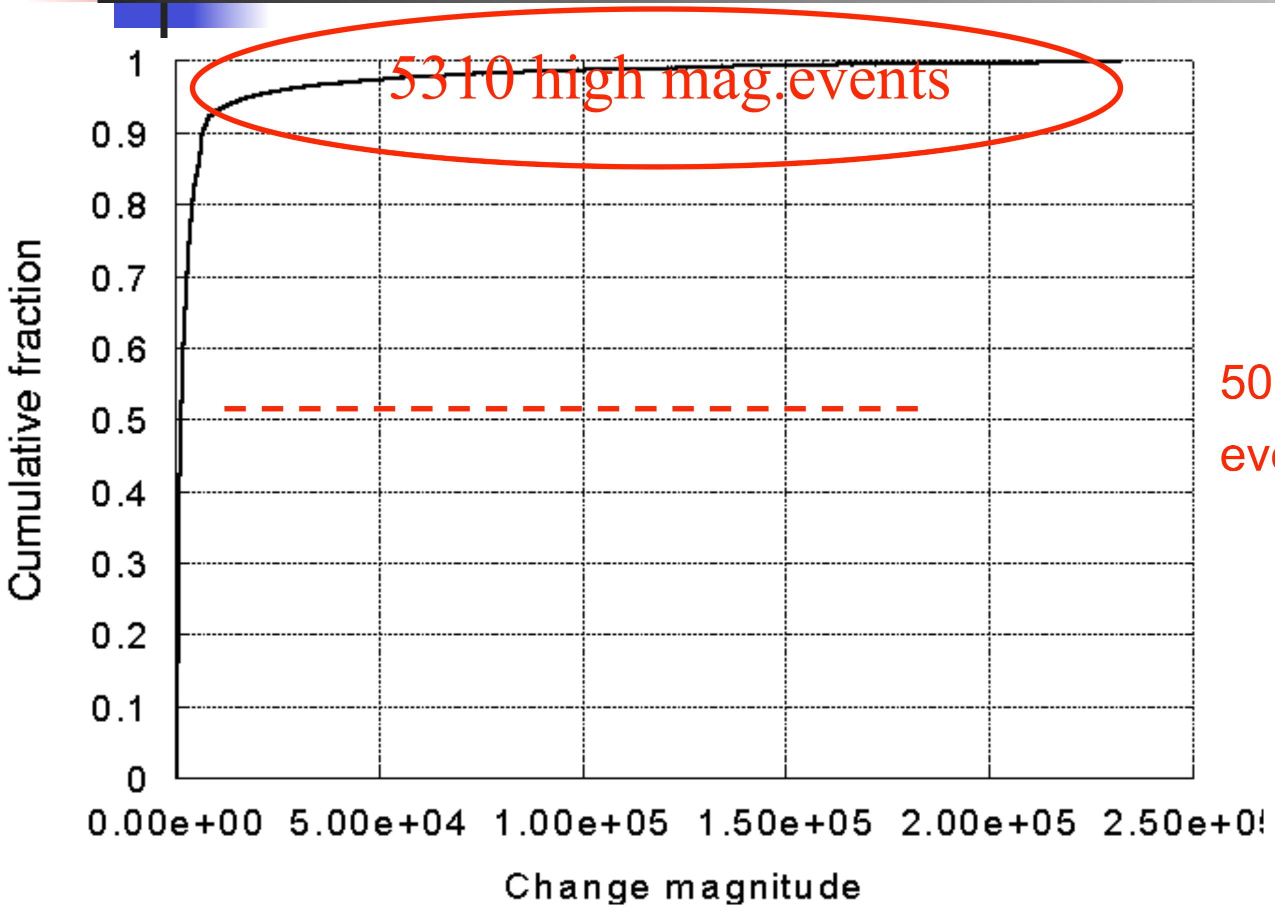

50%: do frequency analysis for events above this point

Higher x indicates bigger event

How to gauge the scope of observed changes

- For each principal component, understand how many monitors are influencing the component by looking at load values
  - If PC1=0.95 x m1 + 0.15 x m2 + 0.005 x m3, then PC1 mostly due to m1, i.e. locally observed change
  - If PC1=0.23 x m1+ 0.22 x m2 + 0.21 x m3, then non-local change, observed by multiple monitors
- Start by examining the load values of the first PC

Qualifying high magnitude changes

The load value plots of the 1st PC for the 5310 high mag events

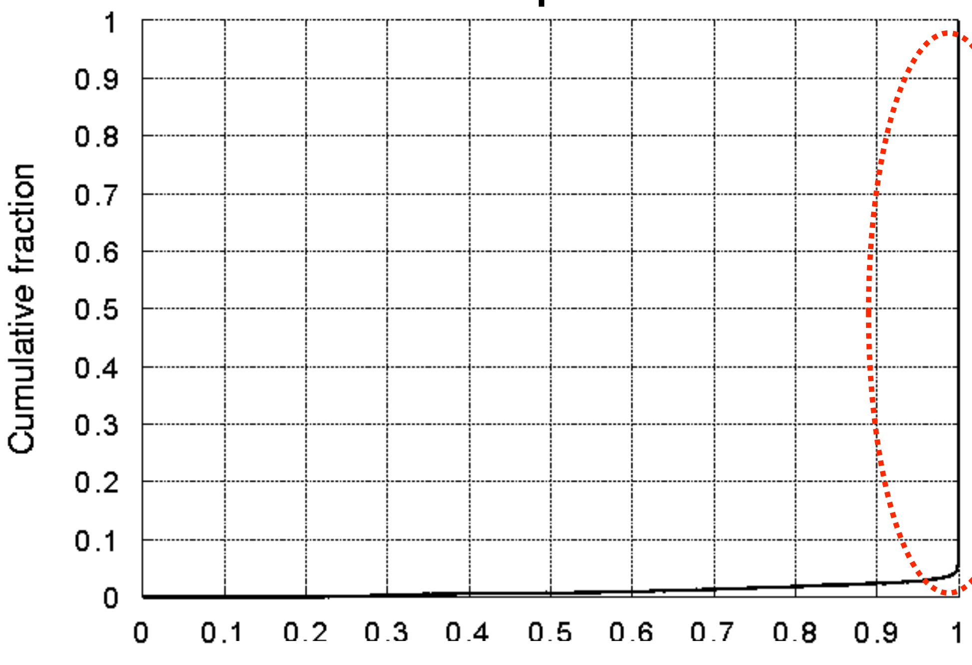

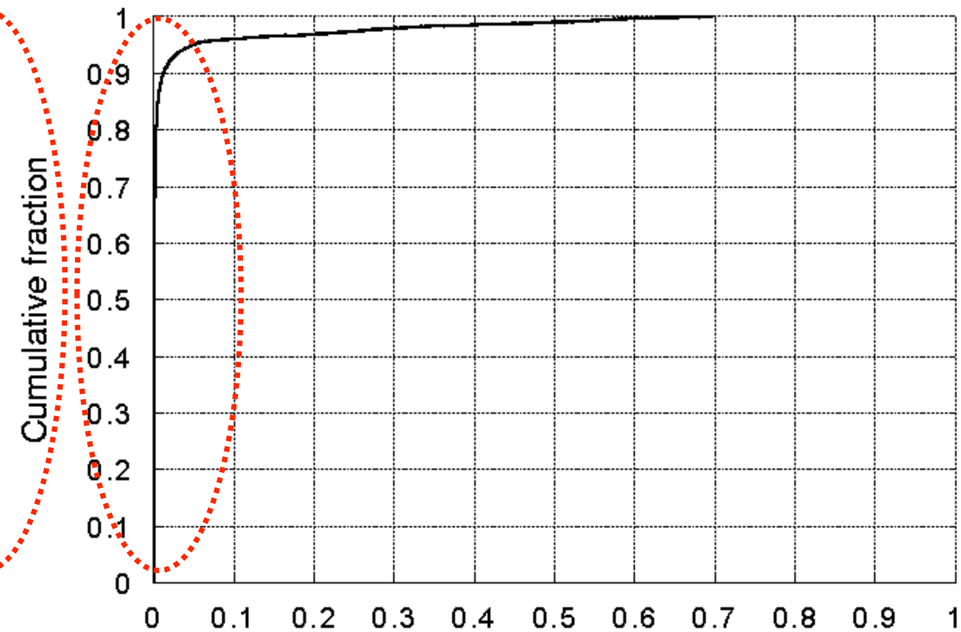

Most high magnitude change influenced by a single monitor (left)
the second most influential monitor is much farther behind (right)
Almost all high magnitude changes are local events

Non-local events

1. Load value of a monitor indicates how much it influences the component.
2. Plot median load and standard deviation of load values of monitors
3. Low standard deviation indicates monitors observe similar things.

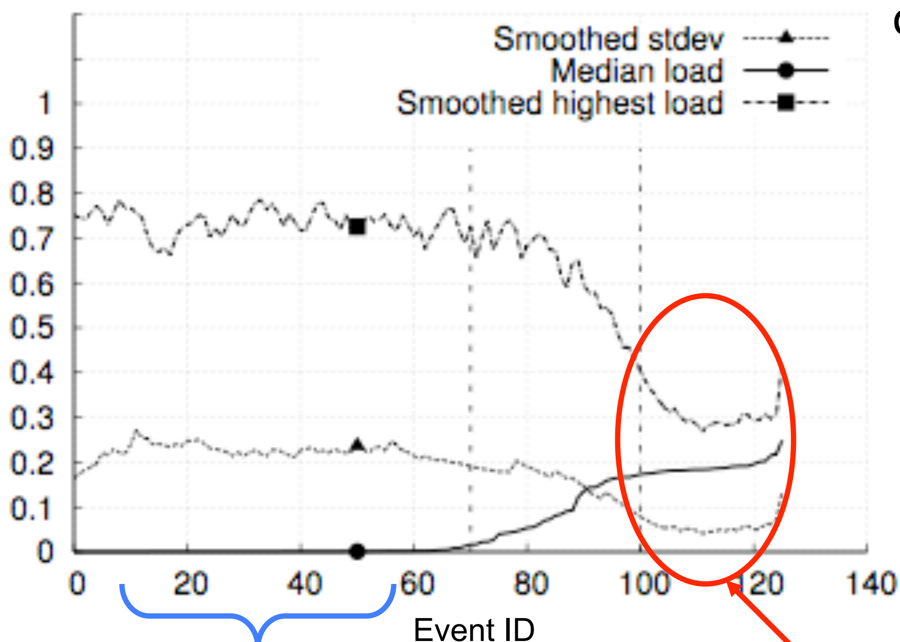

small scope events

large scope events

Typically involve new prefix announcements

- Givena set of prefixes P1 usually announced byASx, when Asx announces a set of prefixes P2, P2>>P1 (longer) for a limited time interval
  - **Announcement of deaggregated prefixes**: if P2 covers (almost) entire prefix space as P1
  - **Announcement of new uncovered prefixes**: if there is (almost) no overlap in the address space covered by P1 and P2

Here is what we caught in 2007

AS-link	Count	Origin AS	category
7018-7015	4	7015	new uncovered prefixes
2200-3356	3	3356	
3549-11456	2	11456	
1237-2200	1	2200	
28513-8151	1	8151	
6453-4788	2	4788	de-aggregation
7018-4788	1	4788	
3257-5486	2	5486	
1239-209	2	209	
17622-9394	1	9394	
7018-33650	1	33650	

Repeated Outliers over time

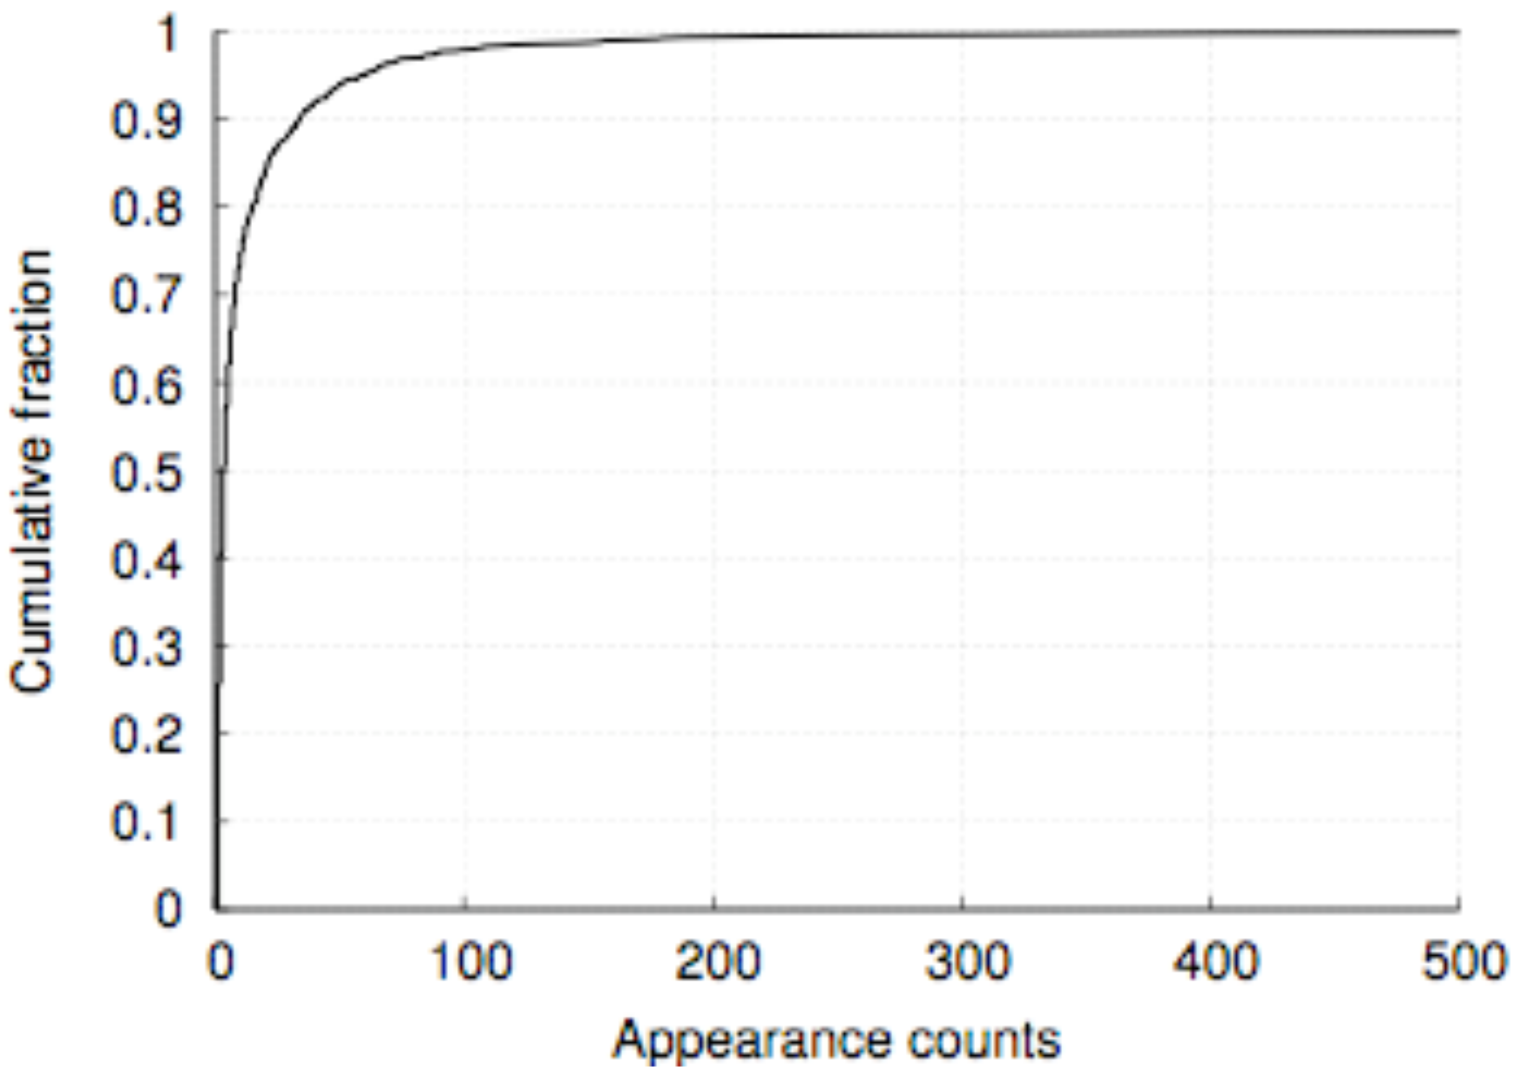

Impact scope of instable links

- Almost all the top 20 links made local impact
  - i.e. only one monitor saw big link weight change
- Handful of cases of repeated problems that are non-local in scope (seen by more than one monitor)
  - Link between AS 6453 (Teleglobe Inc) and AS 30890 (Evolva Telecom) appears 83 times
    - 2nd highest scope in the repeatedly appearing outlier link set
    - 500 routes to AS 30890 or using AS 30890 as an intermediate node in AS-PATH switched to the alternate longer route 6453-5588-5606-30890

Scope of reoccurring events

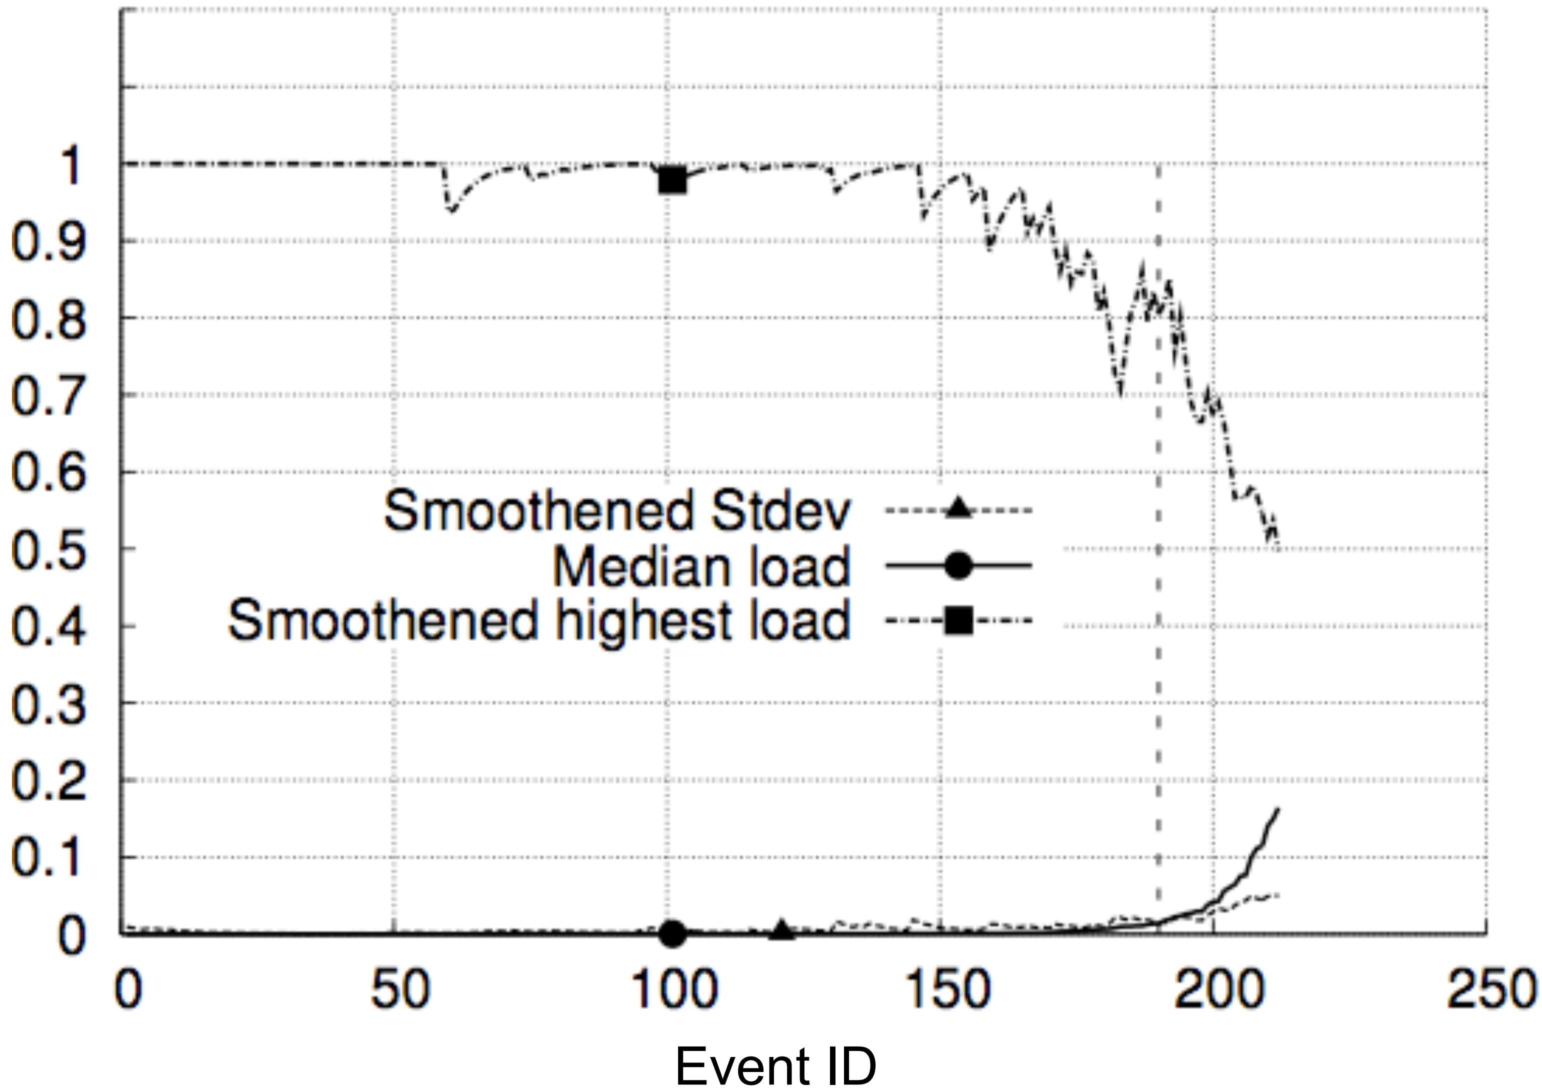

Summary of preliminary results

- High-magnitude events occur infrequently
- Most high-magnitude events are local in scope
- The large-scope events usually involve
  - new prefix announcements,
    - route leakages, prefix de-aggregation
  - or loss of multiple routes (e.g. failure of link of single homed stub carrying tens of routes.)
- There exist a small number of links involved in a large number of noticeable events (w/ local scope)

http://www.cs.ucla.edu/~mohit/submitted/bgpPCA.pdf