Understanding Global Internet Routing Stability Using Link Weight

Mohit Lad, Jonathan Park, Tiziana Refice, Lixia Zhang

Feburary 2008

Goal

- Study Internet routing stability by examining data collected 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***

Our approach

- 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 we separate out which is which

Link-weight

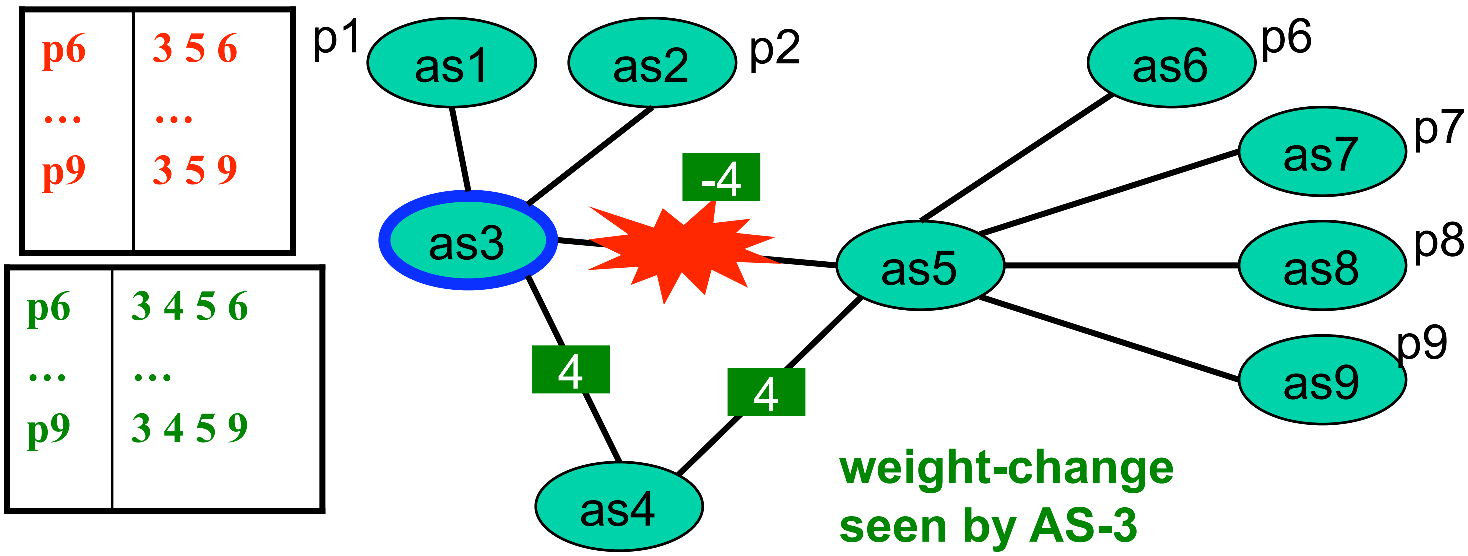

By looking at link weight changes one can

- capture **common behavior across multiple prefixes**

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

- capture **changes seen by 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 link 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 <u>each time bin</u>, <u>each AS link</u>, seen by <u>each monitor</u>
  - If a route change results in a LW change, record the prefix
  - $\delta(a,b)$ = # prefixes that caused 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 the matrix into PCA

An intuitive explanation of PCA

- link (A-B) weight changes seen by 3 monitors, showing as a point in a 3-dimension space

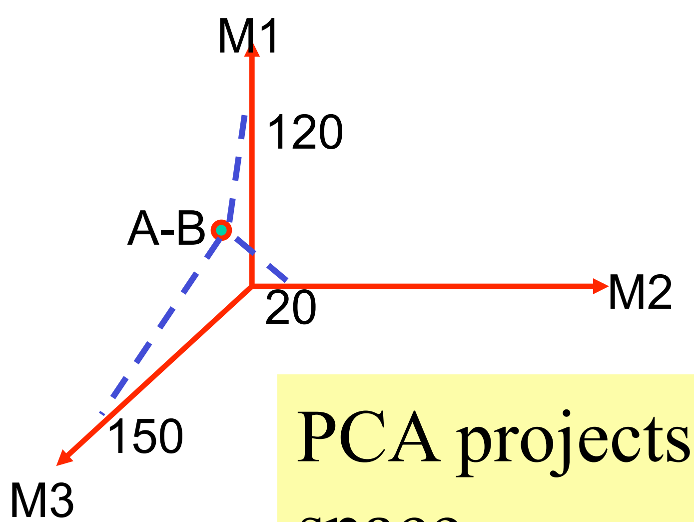

PCA projects data into a new multi-dimensional space
each axis: a linear combination of all monitors views (called "principle component", PC)
PCs are ordered by the values

Example of how PCA helps

View of monitor 1273

View of monitor14608

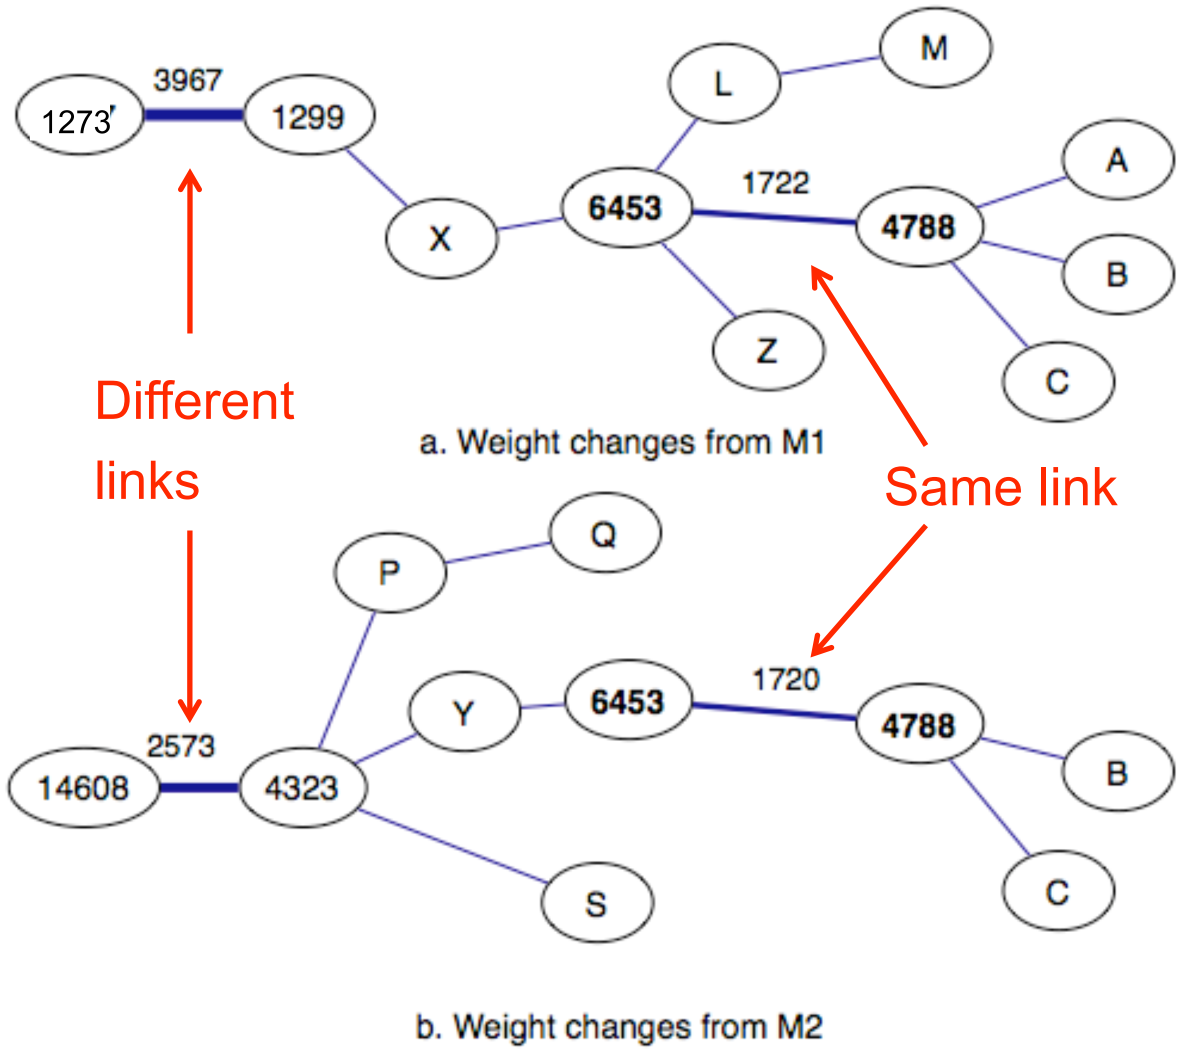

Example of how PCA helps

PCA output with 30 monitors

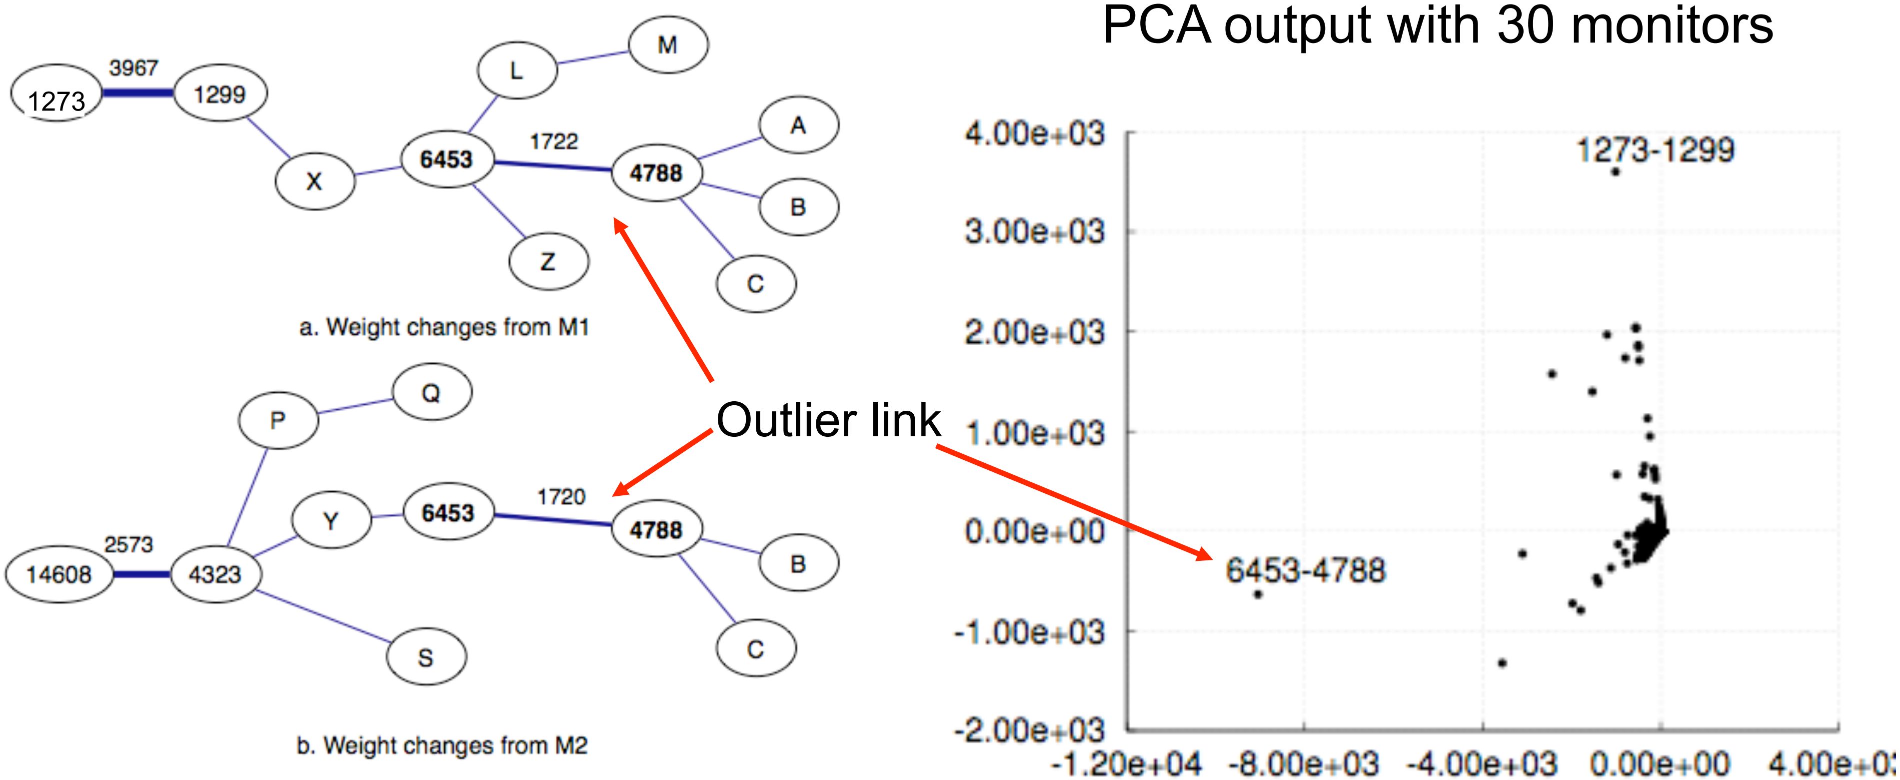

Each axis represents a combination of multiple monitors

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

30 monitors selected

1 22548-200.160.0.130-rrc15
2 4608-202.12.29.64-rrc00
3 2493-206.186.255.223-oreg
4 28895-193.232.244.138-rrc13
5 20483-193.232.244.82-rrc13
6 4777-202.12.28.190-rrc00
7 2018-196.13.250.1-oreg
8 16186-213.179.39.65-rrc00
9 39637-193.239.116.56-rrc03
10 14608-209.161.175.4-oreg
11 34347-193.203.0.124-rrc05
12 8596-193.203.0.130-rrc05
13 6881-195.47.235.100-rrc00
14 21202-194.68.123.139-rrc07
15 6695-80.81.192.158-rrc12
16 26943-195.69.145.49-rrc03
17 8419-195.66.224.121-rrc01
18 15837-80.81.192.126-rrc12
19 2857-80.81.192.8-rrc12
20 12350-192.65.185.157-rrc04
21 20932-192.65.185.142-rrc04
22 3130-147.28.7.1-oreg
23 24875-195.69.144.126-rrc03
24 3741-168.209.255.2-rrc00
25 2152-137.164.16.12-oreg
26 11686-205.241.232.55-oreg
27 293-134.55.200.31-oreg
28 6539-216.18.31.102-oreg
29 5056-167.142.3.6-oreg
30 2905-196.7.106.245-oreg

Questions from data

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

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

Magnitude analysis

Q2: are there links that appear repeatedly as outliers?

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

Frequency analysis

Change Magnitude Distribution

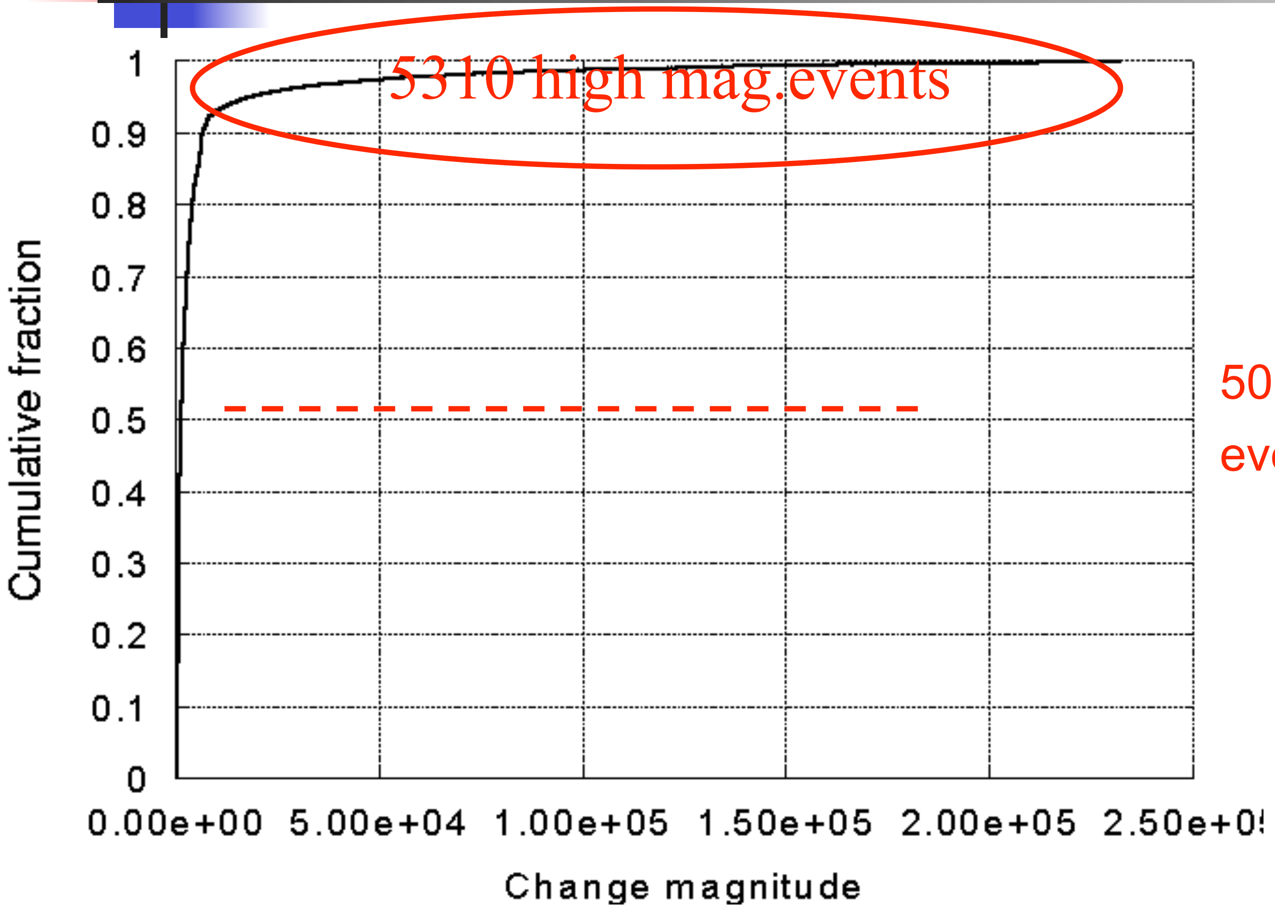

50%: do frequency analysis for events above this point

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

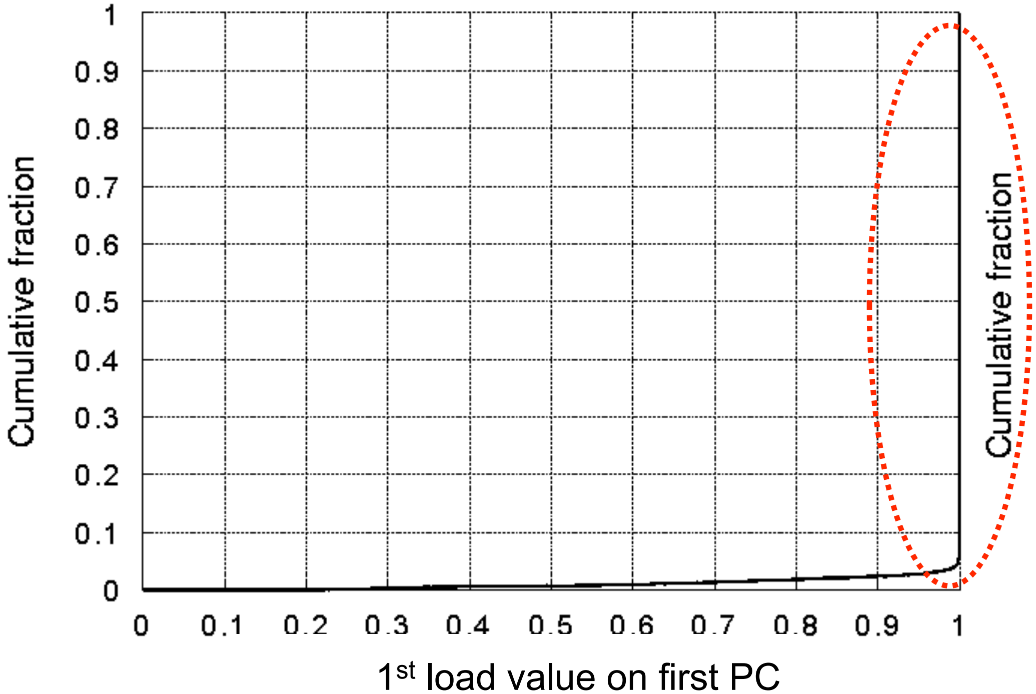

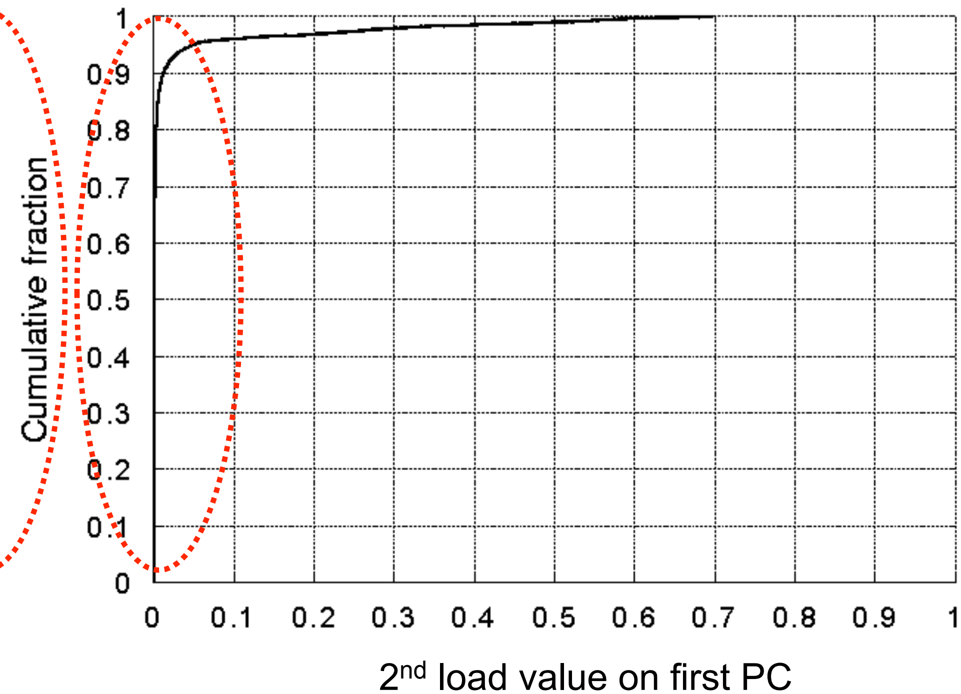

Most high magnitude change influenced by a single monitor (left)
the second most influential monitor is much farther behind (right)
Almost all high magnitude changes are local events

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.

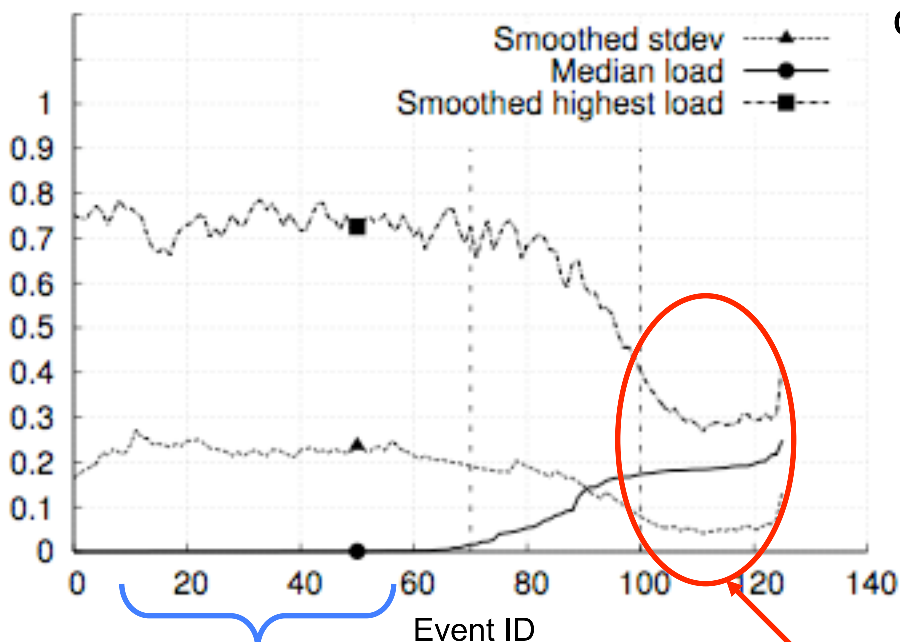

small scope events

large scope events

Typically involve new prefix announcements

Sorting out new prefixe announcements

- a set of prefixes S1 usually announced byASx
- when ASx announces a set of prefixes S2, with S2 much longer than S1 for a limited time interval
  - **Announcement of deaggregated prefixes**: if S2 covers (almost) entire prefix space as S1
  - **Announcement of new uncovered prefixes**: if there is (almost) no overlap in the address space covered by S1 and S2

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

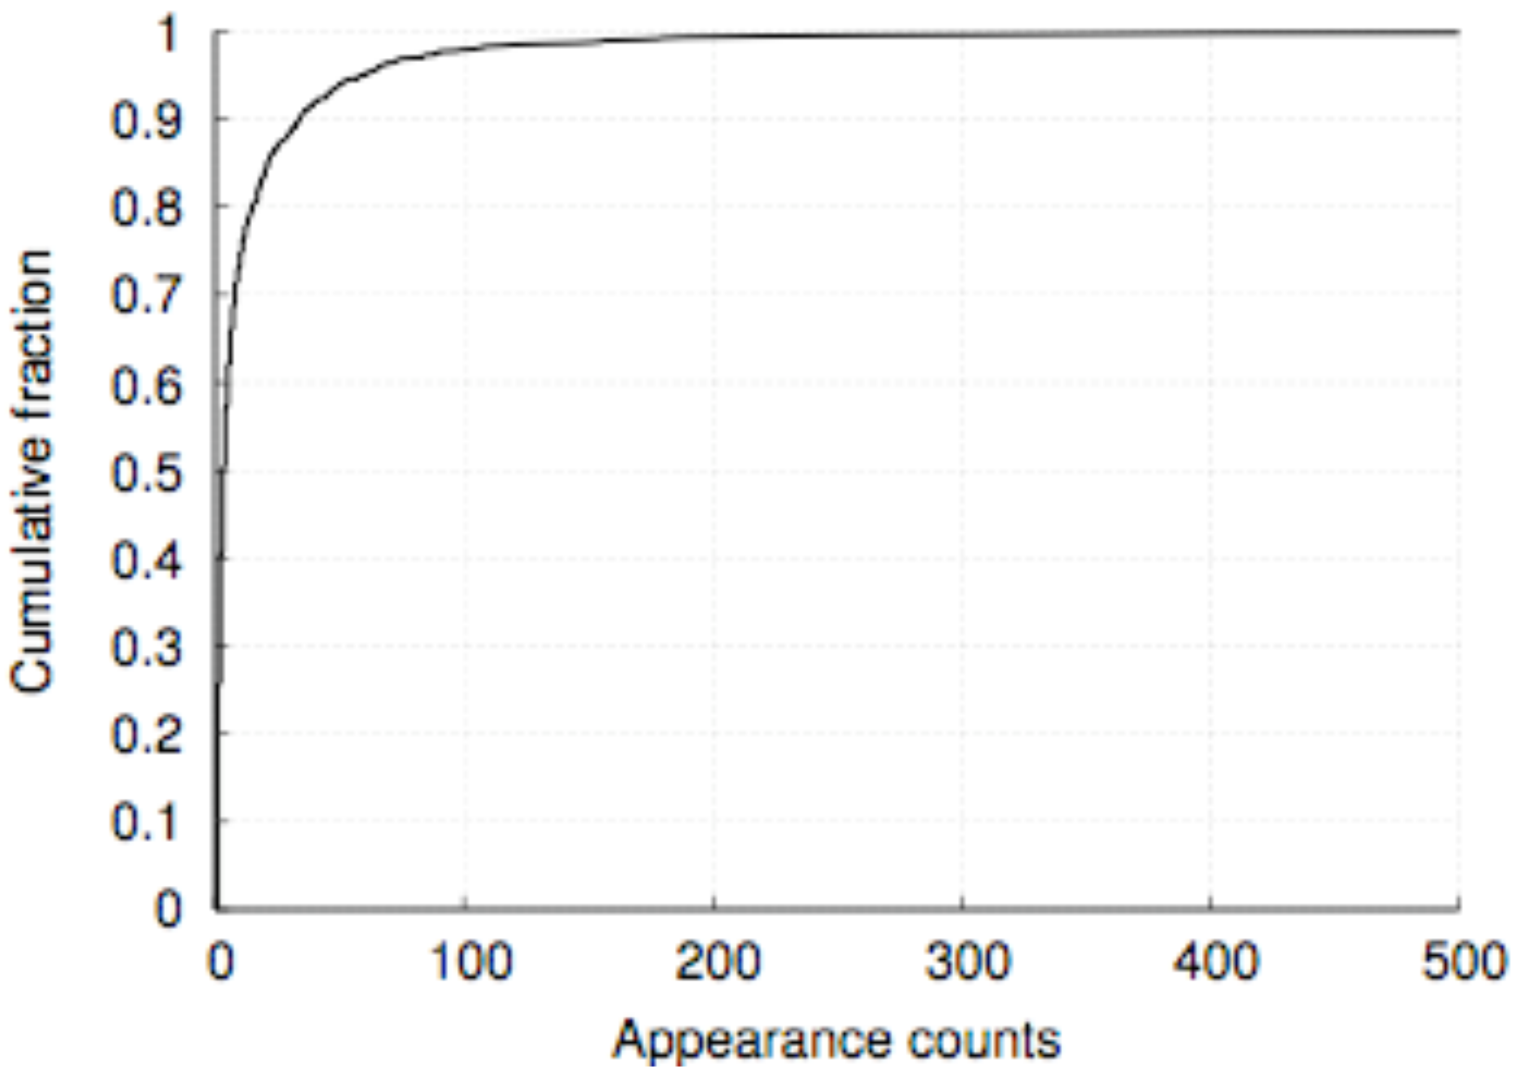

Impact scope of instable links

- 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) 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

Summary of preliminary results

- High-magnitude events occur infrequently; most of them are local in scope
- The large-scope events usually involve
  - new prefix announcements, or
  - complete loss of multiple routes (e.g. diconnection of a stub with tens of routes---rare)
- A small number of links involved in a large number of noticeable events (local impact)

What to carry away

- PCA: a useful tool to deal with data from large number of vantage points
- Link weight and weight changes as a simple way to measure routing dynamics
  - the scripts of doing all the computations in this talk are available online
  - LinkRank: visualizing link weight changes

http://linkrank.cs.ucla.edu/

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

View activity snapshots for all Oregon peers

144.228.241.81 [Mar-04-2008, Mar-10-2008] (7 day span)

164.128.32.11 [Mar-04-2008, Mar-10-2008] (7 day span)

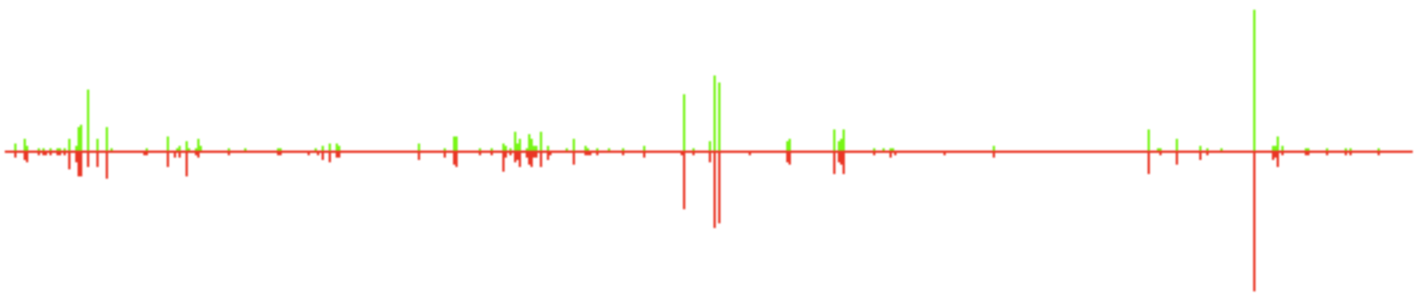

202.232.0.3 [Mar-04-2008, Mar-10-2008] (7 day span)