

## UNDERSTANDING BGP

Locating Internet Routing Instabilities  
Feldmann et. al., SIGCOMM'04

BGP Routing Stability of Popular Destinations  
Rexford et. al., IMC'02

## MOTIVATION

- Interdomain routing suffers from many problems
  - Instability (Labovitz et. al.)
  - Slow convergence after changes (Labovitz et. al.)
  - Misconfigurations (Mahajan et. al.)
  - etc.....
- Poor visibility into dynamics of BGP
  - What are the primary causes of instability?
  - How does BGP respond to a routing change?
- Incomplete model leads to incomplete solutions
  - Route flap dampening

## Root-cause analysis

- The process of inferring the cause (and location) of BGP instabilities using the updates they generate
- Possible causes:
  - Session establishment/teardown/reset
  - BGP attribute or filter manipulation
  - Misconfiguration
  - Traffic engineering, eg. IGP cost change
  - etc...
- A perplexing problem ....
- NOTE: this work focuses on inferring the location of the instability

## Why identify locations of instabilities?


- A first step towards understanding BGP!
- Instabilities can lead to
  - Unreachability / poor performance
  - Route oscillation
  - BGP churn
  - Black holes
  - etc...
- Identifying the location enables corrective action
- But is passive analysis up to the task!!!

## High level approach

- An AS-path change must have occurred along some peering in the previous or the new best path!
- Similar insight for attribute changes
- Observation at multiple vantage points may help solve the puzzle
- Such an approach offers 3 dimensions of information
  1. Time
  2. Prefix
  3. View (vantage point)
- Same as the dimensions used in previous work by Chang et. al., Caesar et. al., Lad et. al.
  - differ in the order of correlation across the dimensions
  - differ in the heuristics for dealing with practical constraints


## Time

### same prefix – same observation point

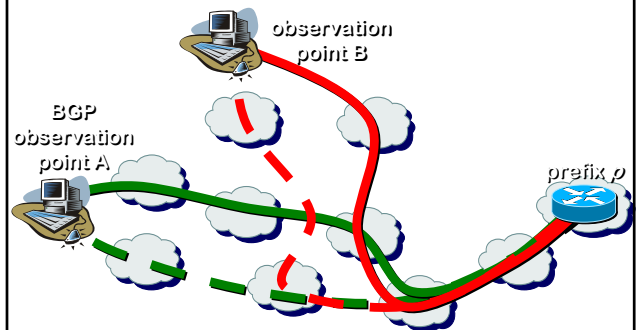


## “Update burst”

- Correlation across time leads to the concept of update burst
- Update bursts or “echoes” are multiple BGP updates for
  - same triggering event
  - at one vantage point
  - for one prefix
- Routing instability: change from “previous” to “new” path
  - “previous best” path no longer available – or – “new best” path becomes available
- UNION of AS edges as candidate set




## Views



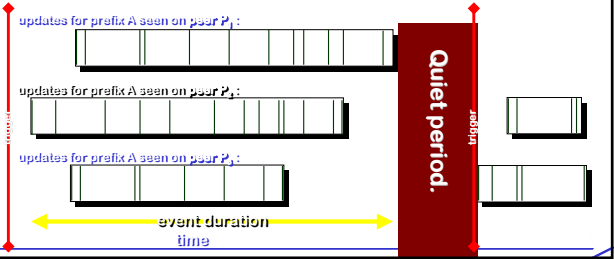
### same prefix – across observation points

## Correlation across views

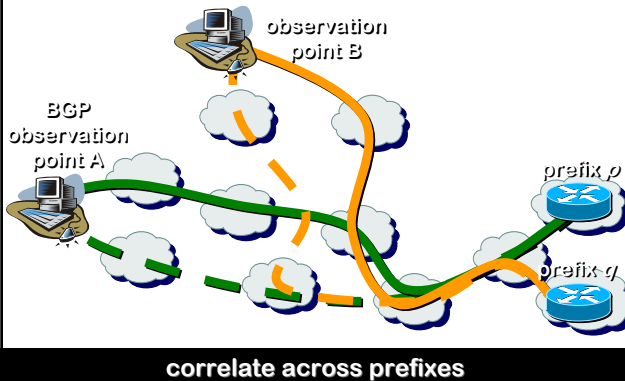


## “Event”

- > Captures update propagation
- > Clusters “updates bursts” across observation points
- > Different timeout heuristics: relative, static, adaptive



## Prefixes



## Correlation across prefixes

- Routing instability leads to changes multiple prefixes
- Intuition behind the greedy heuristic which aims to identify correlated prefixes
- Input: “events” with instability sets
- Output: “correlated events”
- The most popular AS-edge in the candidate sets for simultaneously occurring events must be responsible for the events in which this edge appeared

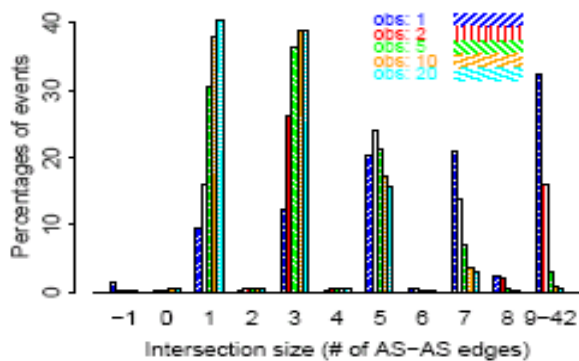
## Evaluation

- Simulation setup:
  - Inferred AS topology from BGP data
  - Incorporated AS relationships and policies
  - Randomly selected failures and vantage points
- Simulation results:
  - The methodology never excludes the simulated failure location
  - Number of observation points matter
  - Average instability sizes after intersection:
    - with only two obs.:  $\leq 7$  edges in 68%
    - with 10 obs.:  $\leq 7$  edges in 88%
  - Location of instabilities matter (in AS-hierarchy)

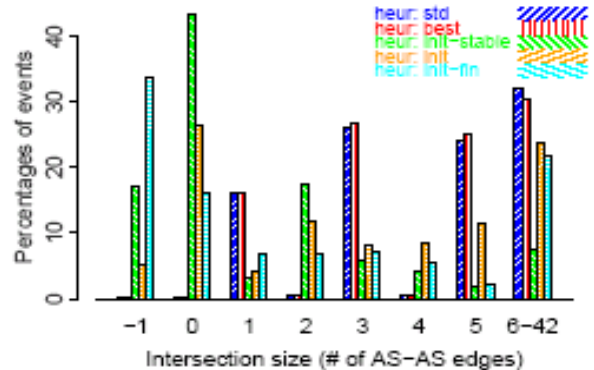
## Evaluation

- BGP data analysis:
  - BGP routing table dumps/updates from RIPE, Routeviews, and Akamai
  - Over 1,100 BGP feeds / 650 ASes (some I-BGP)
- Analysis results:
  - 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%

Instability set size for varying number of observation points




Instability set size for different heuristics



### But ....

- The analysis is too coarse for any practical use!
- The information in a BGP update decreases as it propagates through the network
  - BGP provides good instability isolation
  - Some events are not visible in BGP
  - Cannot be used for determining the cause of the instability
  - Root-cause analysis is bound to have limited utility!
  - Is it useful for the something like a 'BGP health monitor'?
- The instability originator may not lie along the union of previous and new path
  - may lead to detection of induced instability originator
  - but some view may detect the actual originator
  - correlation across views will lead to an empty candidate set

-- from "Realistic BGP Traffic for Test Labs"  
by Olaf Maennel and Anja Feldmann  
[average AS path length=3.2]



### Greedy heuristic

- Single AS edge identified for 93.4% of prefixes
- Three AS edges identified for 97.2% of prefixes
- Validation: Syslog data of tier-1 vs. Greedy results
- Crosscheck:  
Session reset on router  $\Leftrightarrow$  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%
- Successful validation on tier-1 syslog data

## BGP Routing (In)stability

- **Border Gateway Protocol (BGP)**
    - Interdomain routing protocol
    - Route updates at prefix level
    - No activity in "steady state"
  - **But, large # of BGP updates**
    - Failures, policy changes, redundant messages, ...
  - **Implications**
    - Router overhead
    - Transient delay and loss
    - Poor predictability of traffic flow
- Does instability hamper network engineering?

## BGP Routing and Traffic Popularity


- **A possible saving grace...**
  - Most BGP updates due to few prefixes
  - ... and, most traffic due to few prefixes
  - ... but, hopefully not the *same* prefixes
- **Popularity vs. BGP stability**
  - Do popular prefixes have stable routes?
    - Yes, for ~ 10 days at a stretch!
  - Does most traffic travel on stable routes?
    - A resounding yes!
  - Direct correlation of popularity and stability?
    - Well, no, not exactly...

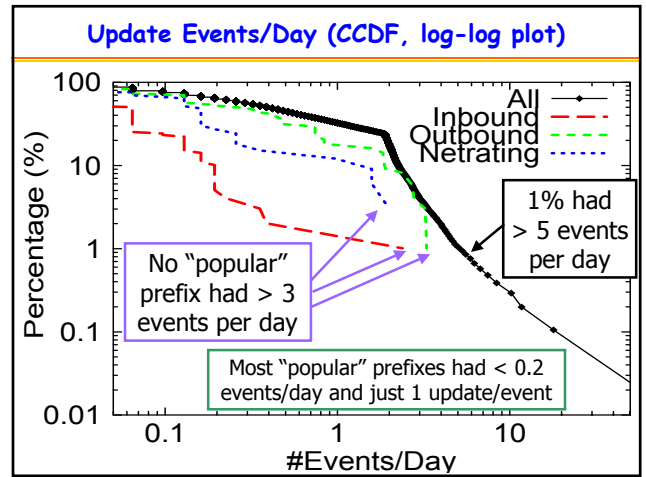
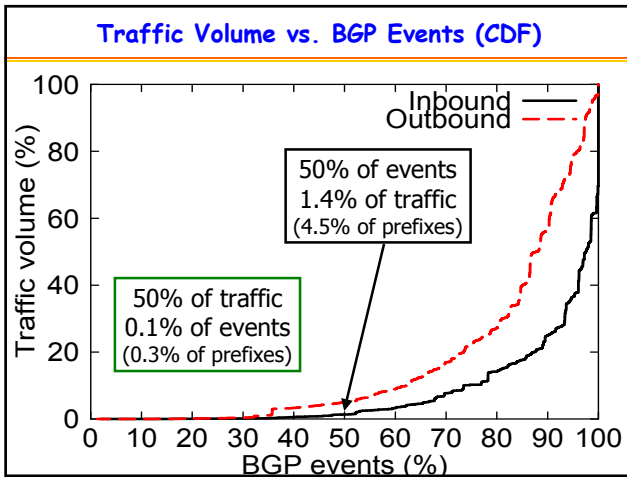
## BGP Updates

- **BGP updates for March 2002**
    - AT&T route reflector
    - RouteViews and RIPE-NCC
  - **Data preprocessing**
    - Filter duplicate BGP updates
    - Filter resets of monitor sessions
    - Removes 7-30% of updates
  - **Grouping updates into "events"**
    - Updates for the same prefix
    - Close together in time (45 sec)
    - Reduces sensitivity to timing
- Confirmed: few prefixes responsible for most events

## Two Views of Prefix Popularity

- **AT&T traffic data**
  - Netflow data on peering links
  - Aggregated to the prefix level
  - Outbound from AT&T customers
  - Inbound to AT&T customers
- **NetRatings Web sites**
  - NetRatings top-25 list
  - Convert to site names
  - DNS to get IP addresses
  - Clustered into 33 prefixes





### An Interpretation of the Results

- **Popular → stable**
  - Well-managed
  - Few failures and fast recovery
  - Single-update events to alternate routes
- **Unstable → unpopular**
  - Persistent flaps: hard to reach
  - Frequent flaps: poorly-managed sites
- **Unpopular does *not* imply unstable**
  - Most prefixes are quite stable
  - Well-managed, simple configurations
  - Managed by upstream provider

### Conclusions

- **Measurement contributions**
  - Grouping BGP updates into "events"
  - Popular prefixes from NetRatings
  - Joint analysis of popularity & stability
- **Positive result for network operators**
  - BGP instability does not affect most traffic
- **Future work**
  - Stability of the IP forwarding path
    - Does popularity imply stable forwarding path?
    - Relationship between BGP and forwarding path?
  - BGP traffic engineering
    - Tune BGP routing policies to prevailing traffic
    - Prefixes w/ stable BGP routes & high/stable volumes