Towards root cause analysis of BGP routing dynamics

Matthew Caesar, Lakshmi Subramanian, Randy H. Katz

mccaesar@cs.berkeley.edu

Motivation

- Interdomain routing suffers from many problems
  - Instability
  - Slow convergence after changes
  - Misconfigurations
- Poor visibility into dynamics
  - What is the spectrum of causes of route changes?
  - What are the primary causes of instability?
  - How does BGP respond to a routing change?
- Incomplete model → incomplete solutions

How can we improve routing?

- BGP Health monitoring system
  - Collect routes from routers
  - Infer properties of network elements
  - Redistribute information
- Achieves greater visibility
- Our focus: how to do inference

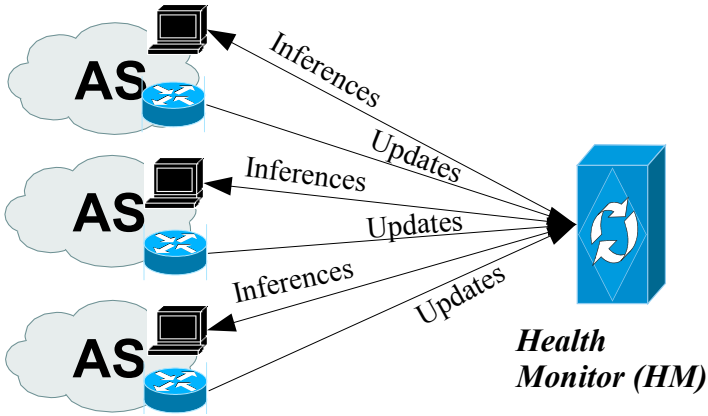

Uses of a health monitor

- Troubleshooting/debugging
- BGP policies based on historical link/node stability
- Online help to path selection
- Help customers choose upstream providers

Inference problem

- Terminology:
  - Event: an activity that generates routing updates
  - Suspect location set: set of AS's and links where event could have occurred
  - Suspect cause set: set of types of events that could have occurred
- Problem: Given route updates observed at multiple vantage points, determine the suspect set = {suspect cause set, suspect location set} of routing events that trigger each update

Correlating Observations: Main idea

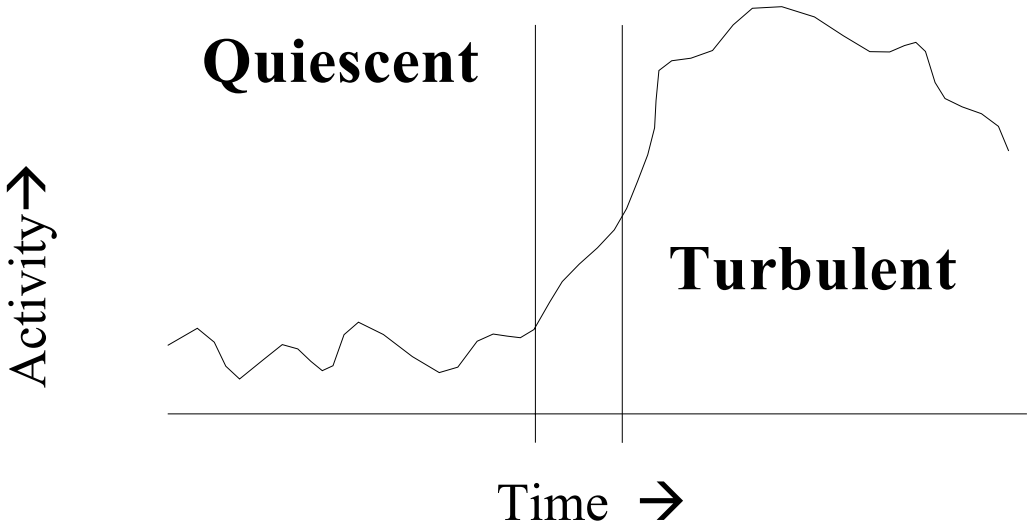

- Quiescent: Assume bursts of updates to a *single prefix* are correlated
- Turbulent: Assume updates to *many prefixes* are correlated
- Assuming correlated observations are independent worsens *precision*, but assuming independent observations are correlated worsens *accuracy*

Our approach

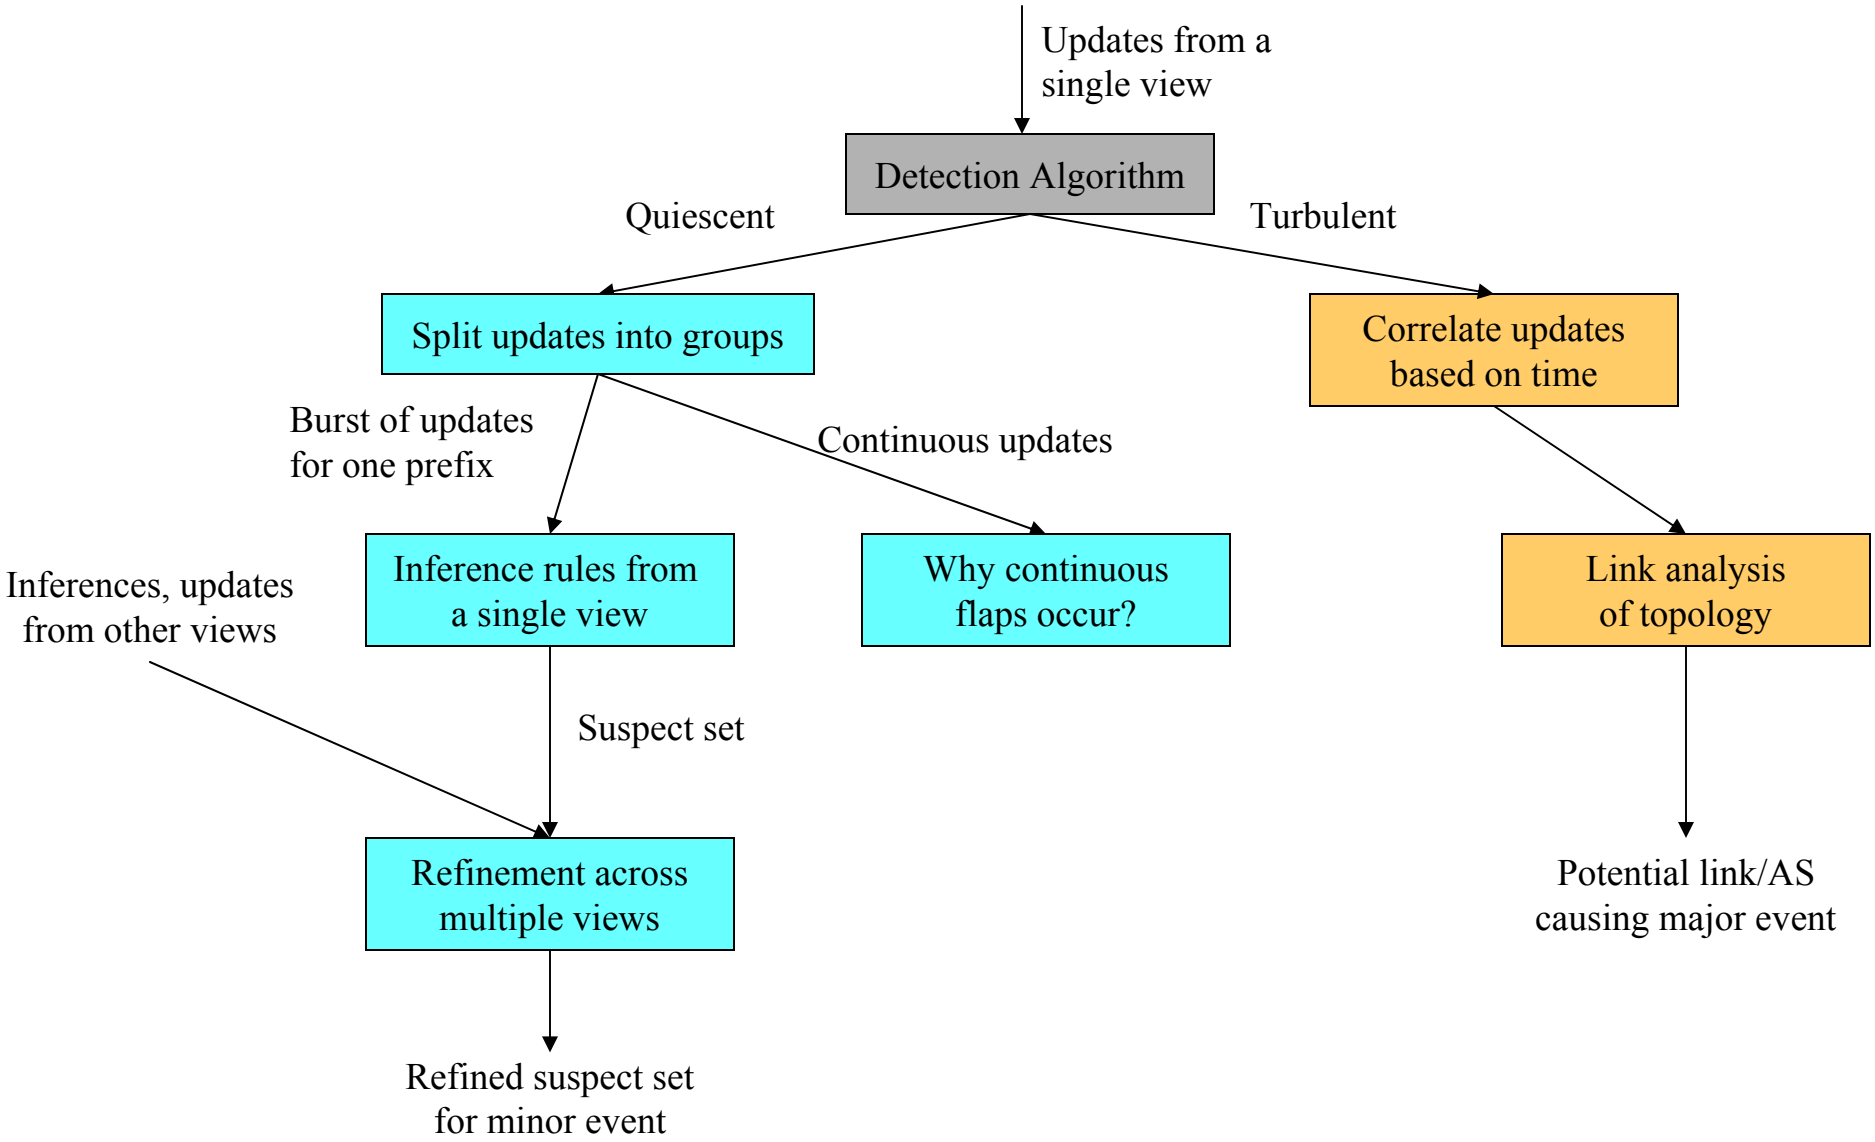

Turbulent Inference: Example

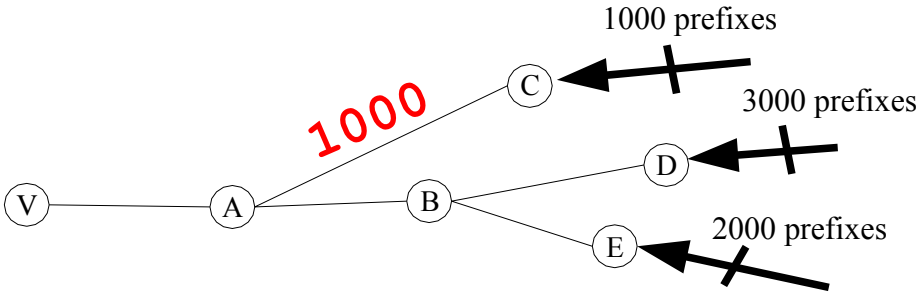

- Scenario #1:
  - ~1000 prefixes updated that used (A,C)
  - → (A,C) suspect

Turbulent Inference: Example

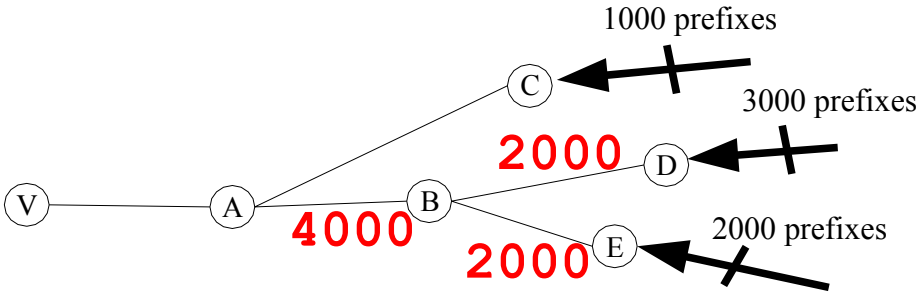

- Scenario #2:
  - ~4000 prefixes updated that used (A,B), ~2000 that used (B,E)
  - → (A,B) suspect

Turbulent Inference: Example

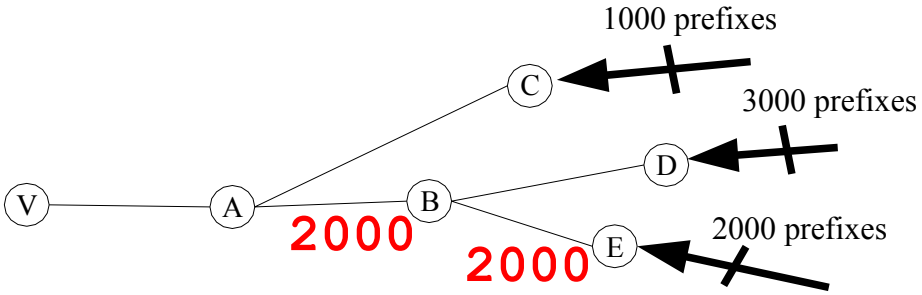

- Scenario #3:
  - ~2000 prefixes updated that used (A,B), ~2000 that used (B,E)
  - → (B,E) suspect

TurbulentInfer: Issues

- Effects of simultaneously occurring independent events are overshadowed
- Two large events simultaneously occurring
- Transition from Quiescent to Turbulent periods

Quiescent Inference: Example

- REROUTE
  - Silence, route change, silence
- Potential causes (yellow):
  - Link/router repair
  - MED/LocalPref decrease
  - Hold-down expired
- Potential causes (blue):
  - Link/router failure
  - MED/LocalPref increase
  - Hold-down triggered

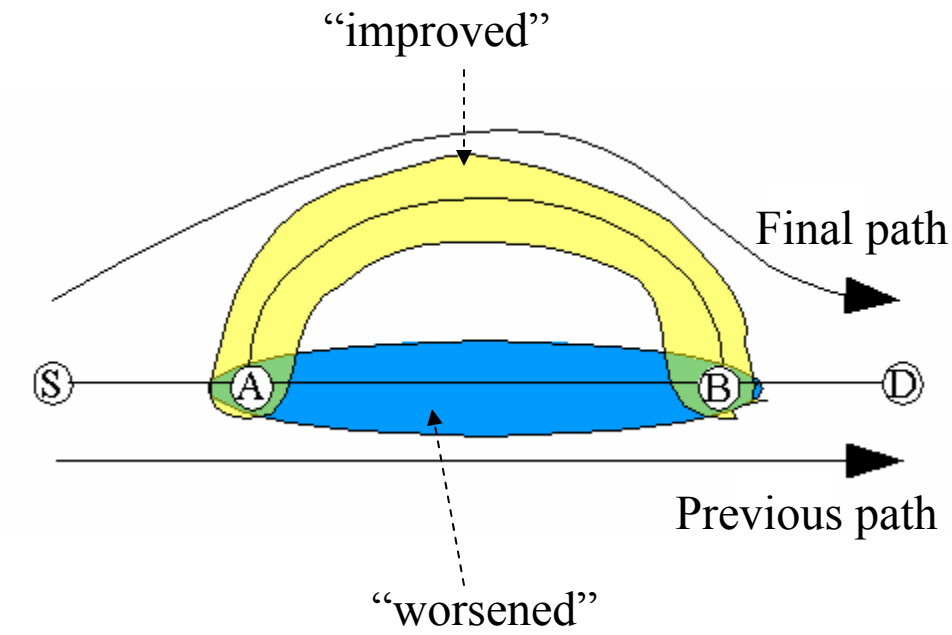

QuiescentInfer: Inference across views

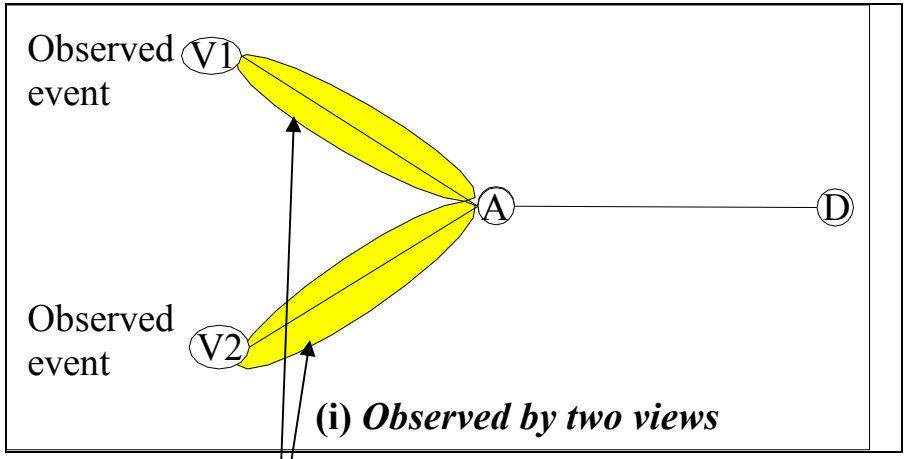

Event did *not* occur here

Event did *not* occur here

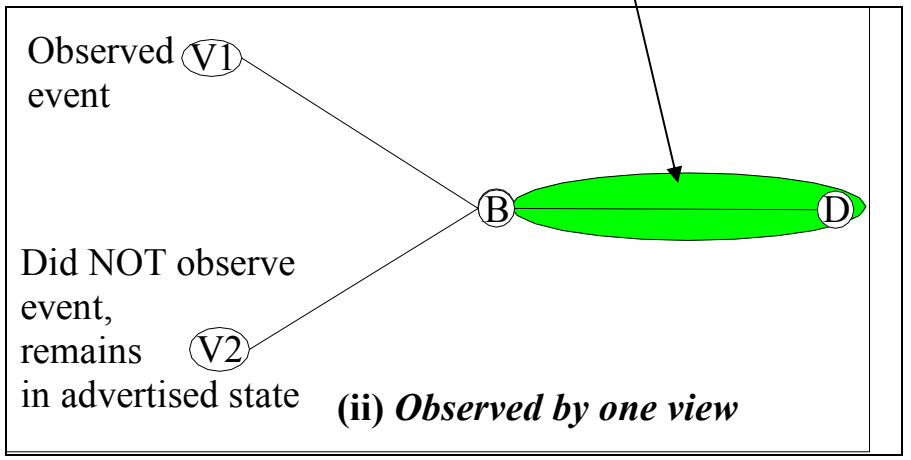

QuiescentInfer: Issues

- Simultaneous events:
  - Eg. Flap in one view, Reroute in another
  - Eg 2. Advertisement in one view, Withdrawal in another
- Community attribute changes
  - Community change can trigger reroute several hops away
- Multiple peering links

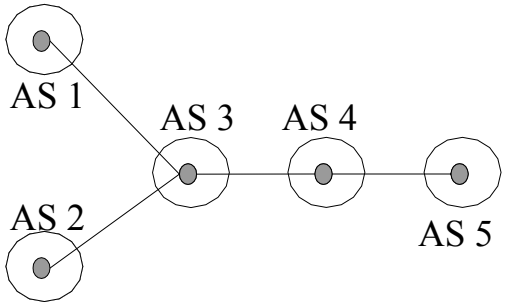

(a) Singly-peered

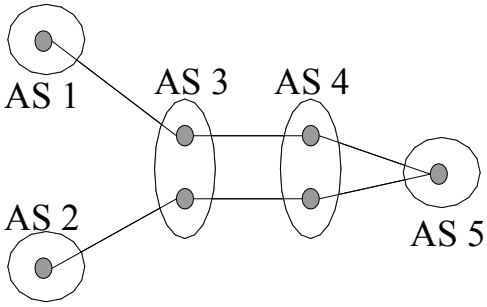

(b) Multiply-peered

Validation

- Well-known historical events
  - UUNET (10/3/02), AT&T (8/28/02) routing difficulties
  - Internet worms
- BGP beacons
- View at the origin

BGP Beacon results

Suspect set size	1	2	3	4	all
Percent of bursts	61%	92%	93%	97%	100%
Inferences containing beacon AS	100%	100%	100%	100%	100%

View at the origin results

Suspect set size	1	2	3	4
Avg. suspect set size	7%	20%	36%	54%
Incorrect inferences	0%	0%	0%	0%

Results

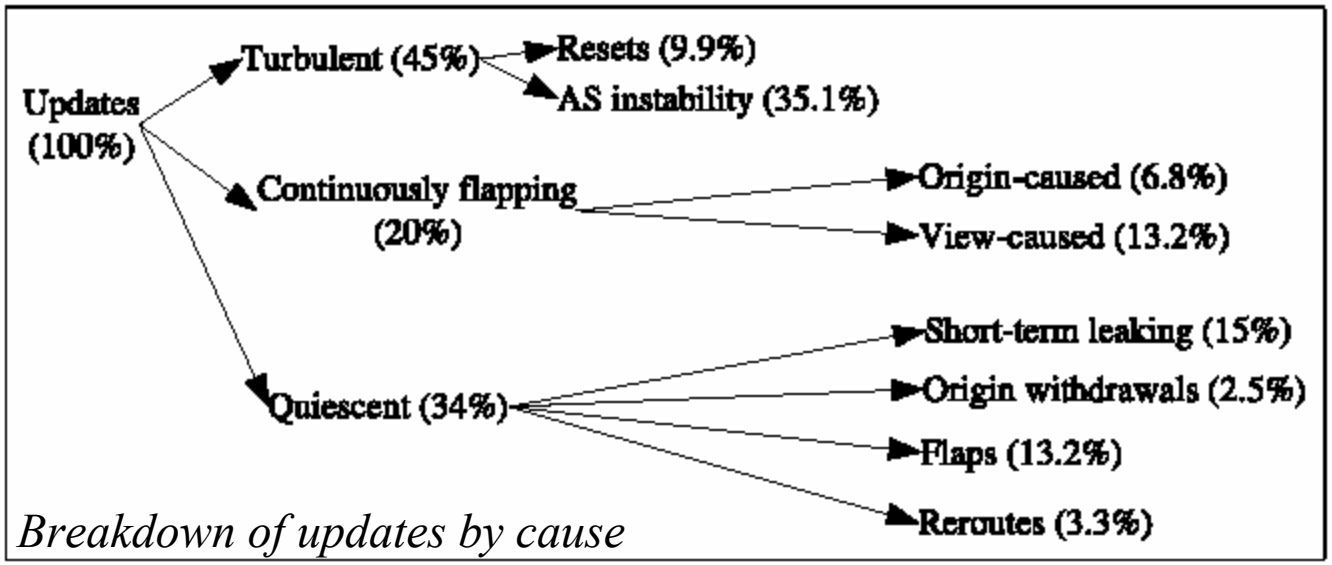

Breakdown of updates by cause

- 70% of updates can be pinpointed to a single inter-AS link (pair of AS's)
  - More precise inference for more major events

Results

- Few AS's, links causing majority of updates
- Many previously unknown major events
  - Peering instability: July 21 2003: link between two tier 1's affected reachability to several popular prefixes over 10 hour period
  - Misconfiguration: Jan 26 2003: AS erroneously advertised 500 prefixes
  - Reroute: Jan 23 2003: link between two tier 1's causes many prefixes to switch to alternate paths

Future work

- Investigate continuously flapping prefixes
- Apply statistical inference techniques
- Placement of views
- Alarms/Triggers to detect unhealthy behavior
- Real time analysis
  - http://www.cs.berkeley.edu/~mccaesar/hmon.html