Diagnosing Network Disruptions with Network-wide Analysis

Yiyi Huang, Nick Feamster,
Anukool Lakhina*, Jim Xu
College of Computing, Georgia Tech
** Guavus, Inc.*

Problem Overview

- Network routing disruptions are frequent
  - On Abilene from January 1,2006 to June 30, 2006
    - **379 e-mails, 282 disruptions**
- How to help network operators deal with disruptions quickly?
  - Massive amounts of data
  - Lots of noise
  - Need for fast detection

Existing Approaches

- Many existing tools and data sources
  - Tivoli Netcool, SNMP, Syslog, IGP, BGP, etc.
- Possible issues
  - Noise level
  - Time to detection
- Network-wide correlation/analysis
  - Not just reporting on manually specified traps
- **This talk:** Explore *complementary* data sources
  - First step: Mining BGP routing data

Challenges: Analyzing Routing Data

- Large volume of data
- Lack of semantics in a *single* stream of routing updates
- **Needed: *Mining*, not simple reporting**

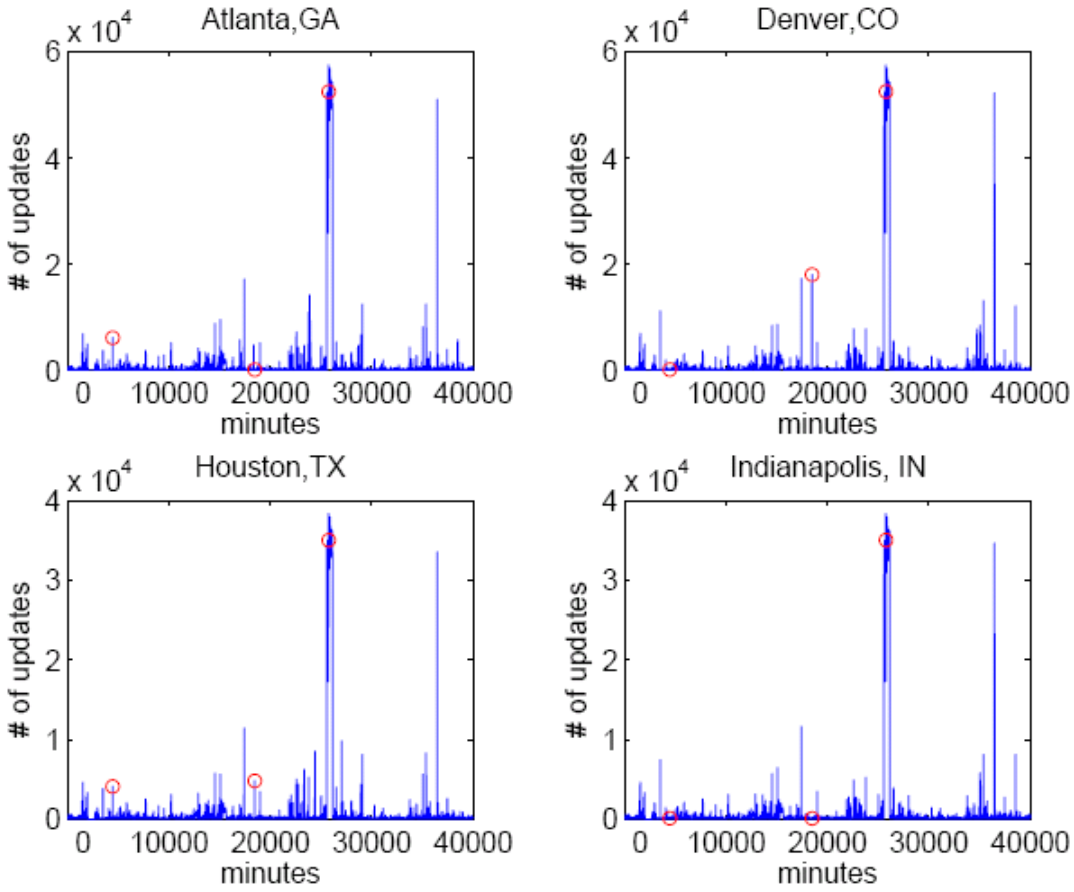

Idea: Can we improve detection by mining network-wide dependencies *across* routing streams?

Key Idea: Network-Wide Analysis

Don't treat streams of data independently.
"Big" network events may cause correlated blips.

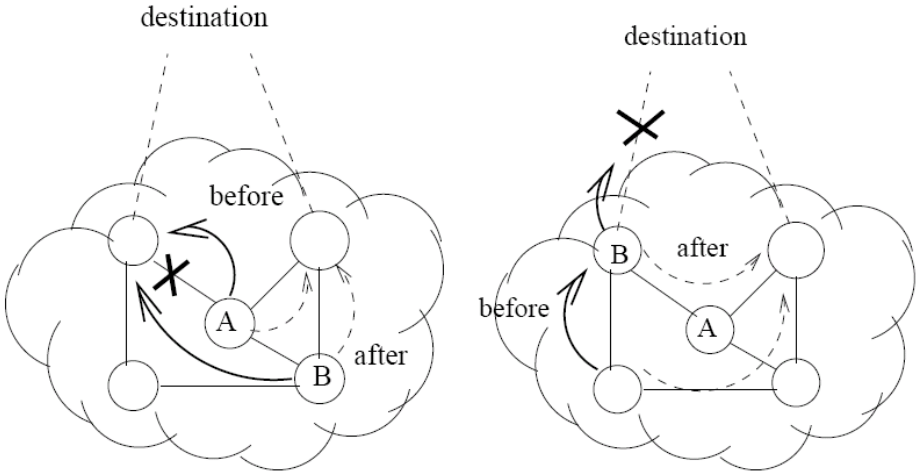

(a) Internal disruption (b) External disruption

- Structure and configuration of the network gives rise to dependencies across routers
- Analysis should be cognizant of these dependencies.

Overview

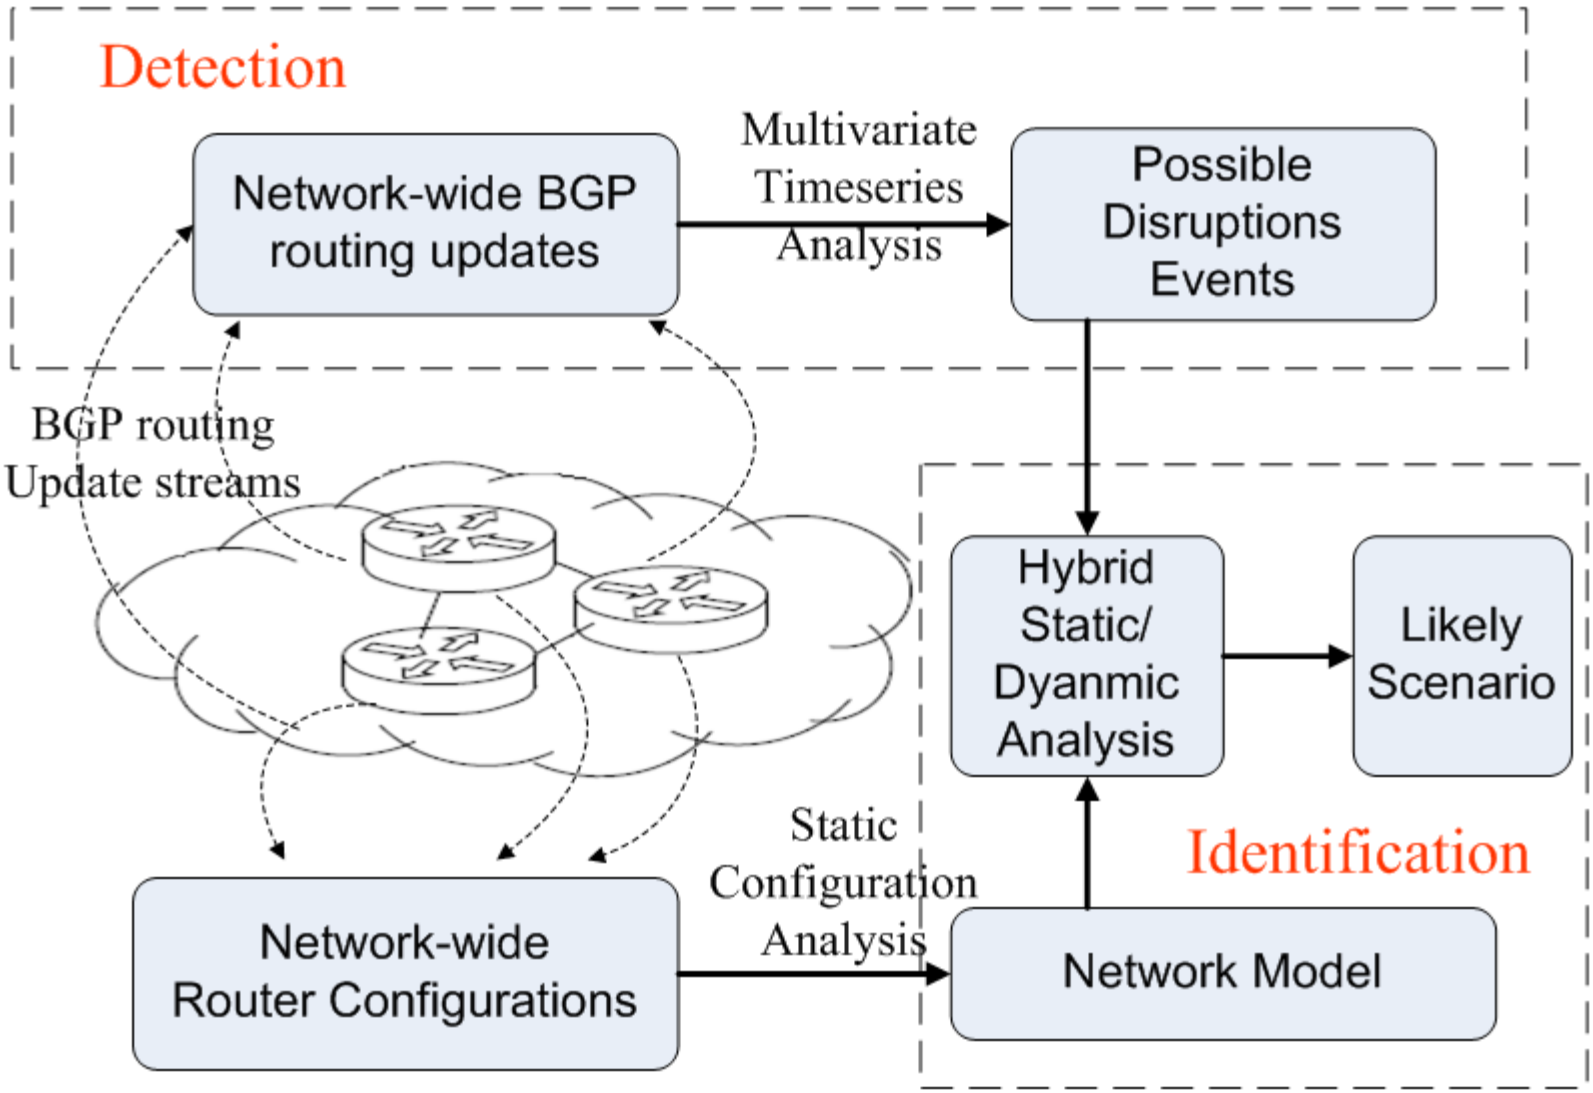
Detection
Network-wide BGP routing updates
Multivariate Timeseries Analysis
Possible Disruptions Events
BGP routing Update streams
Hybrid Static/ Dyanmic Analysis
Likely Scenario
Identification
Static Configuration Analysis
Network-wide Router Configurations
Network Model

Detection

- **Approach:** network-wide, multivariate analysis
  - Model network-wide dependencies directly from the data
  - Extract common trends
  - Look for deviations from those trends
- **High detection rate** (for acceptable false positives)
  - 100% of node/link disruptions, 60% of peer disruptions
- **Fast detection**
  - Current time to reporting (in minutes)

	median	minimum
node	43.35	29.17
link	57.87	14.38
peer	96.13	9.25

Identification: Approach

Goal

- Classify disruptions into four types
  - Internal node, internal link, peer, external node

Approach

Track three features

2. Global iBGP next-hops
3. Local iBGP next-hops
4. Local eBGP next-hops

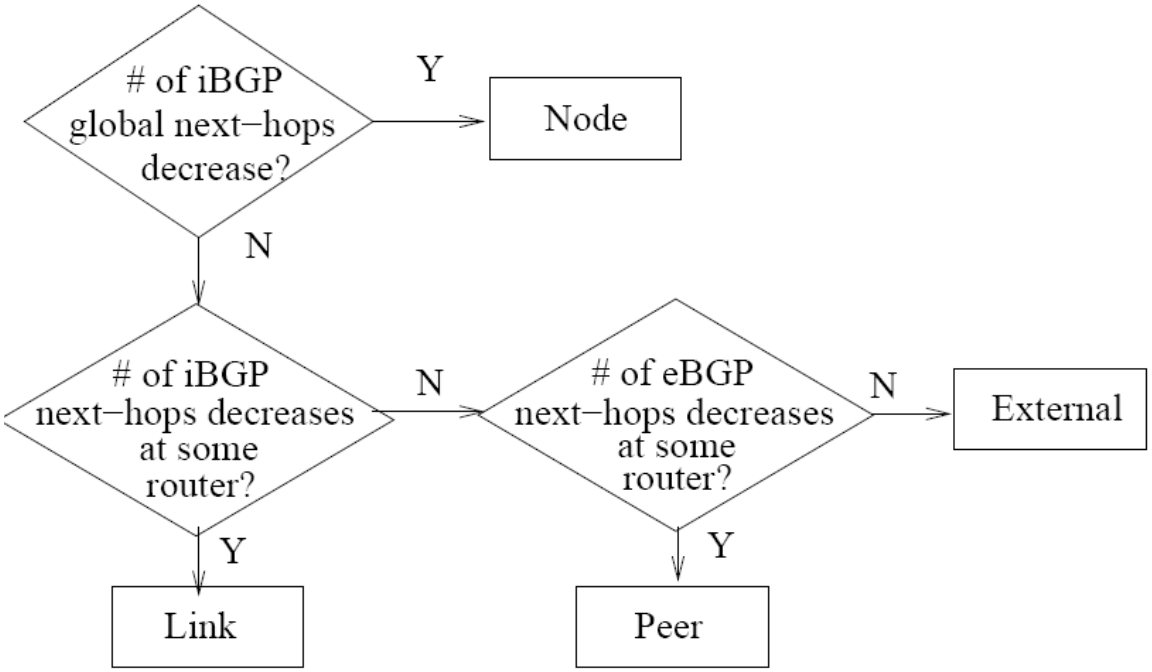

Identification: Results

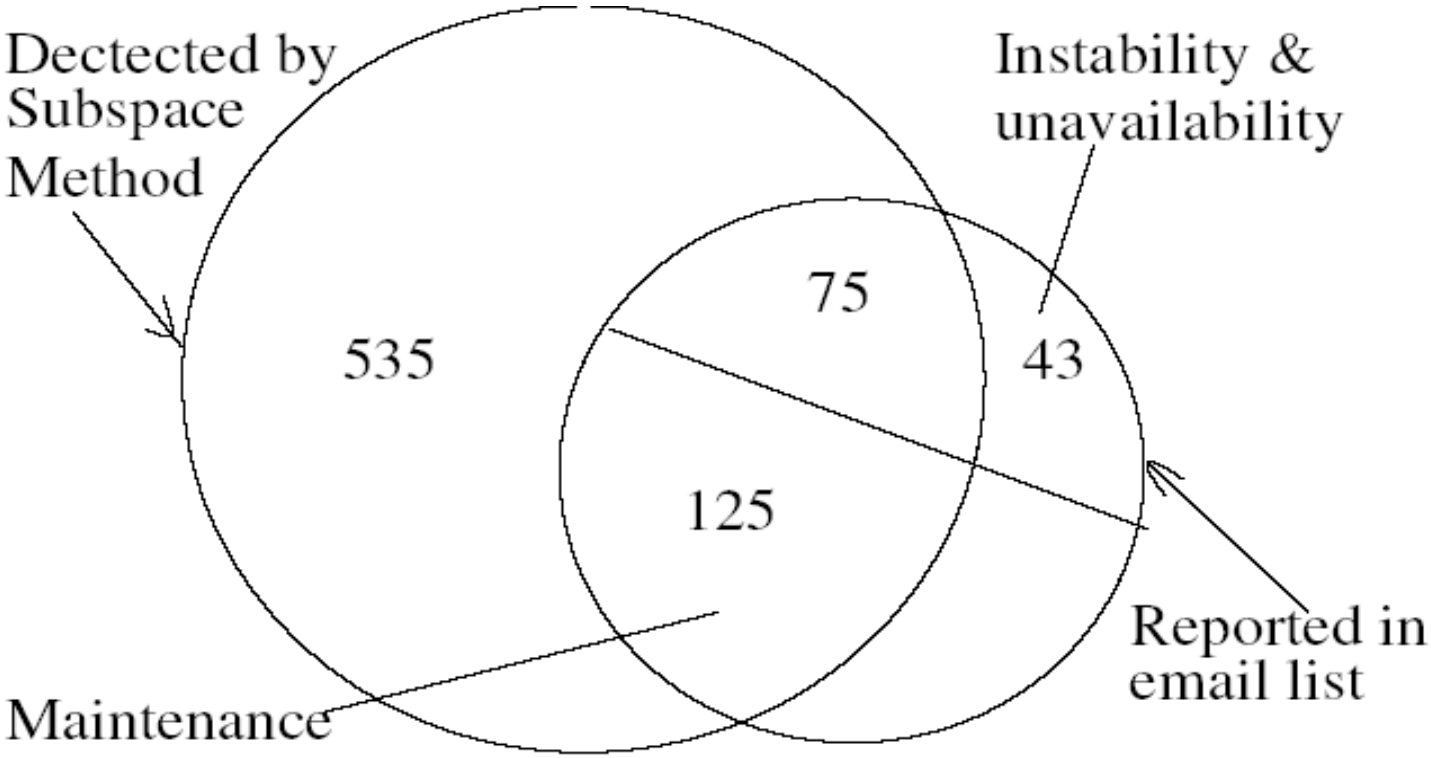

Main Results

- 90% of local disruptions are visible in BGP
  - Many disruptions are low volume
  - Disruption "size" can vary by several orders of magnitude
- About 75% involve more than 2 routers
  - Analyze data *across* streams
  - BGP routing data is but one possible input data set
- Detection:
  - 100% of node and link disruptions
  - 60% of peer disruptions
- Identification
  - 100% of node disruptions,
  - 74% of link disruptions
  - 93% of peer disruptions

Questions for Network Operators

Please send thoughts to feamster at cc.gatech.edu.

- How happy are you with existing approaches?
- What are the most common types of network faults you must diagnose?
- How effective are existing tools in terms of:
  - Reducing the noise level in reporting?
  - Fast detection?
- What about incorporating other data (*e.g.,* active probes)?
- Can you help us work with you to improve detection/identification of disruptions?

http://www.cc.gatech.edu/~feamster/papers/diagnosis07-tr.pdf