UNIVERSITÀ DEGLI STUDI ROMA TRE
Dipartimento di Informatica e Automazione

Tracking Back the Root Cause of a Path Change in Interdomain Routing

Alessio Campisano
Luca Cittadini
Giuseppe Di Battista
Tiziana Refice
Claudio Sasso

IEEE/IFIP Network Operations & Management Symposium (NOMS 2008)
Apr 10th, 2008

Tracking Back the Root Cause of a Path Change in **Interdomain Routing**

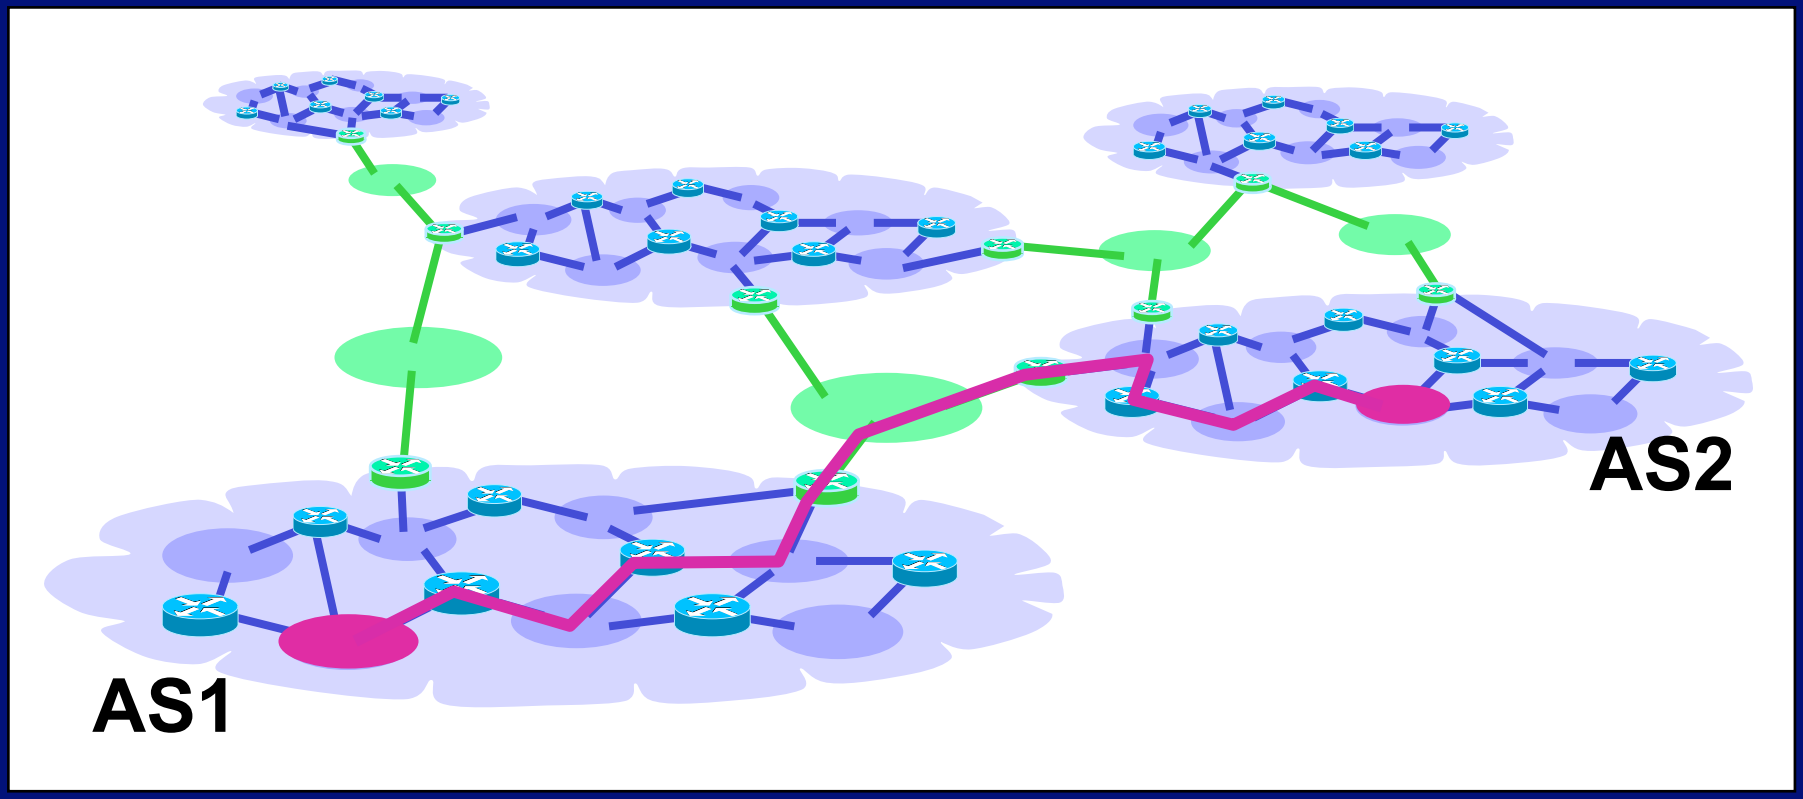

Tracking Back the Root Cause of a Path Change in Interdomain Routing

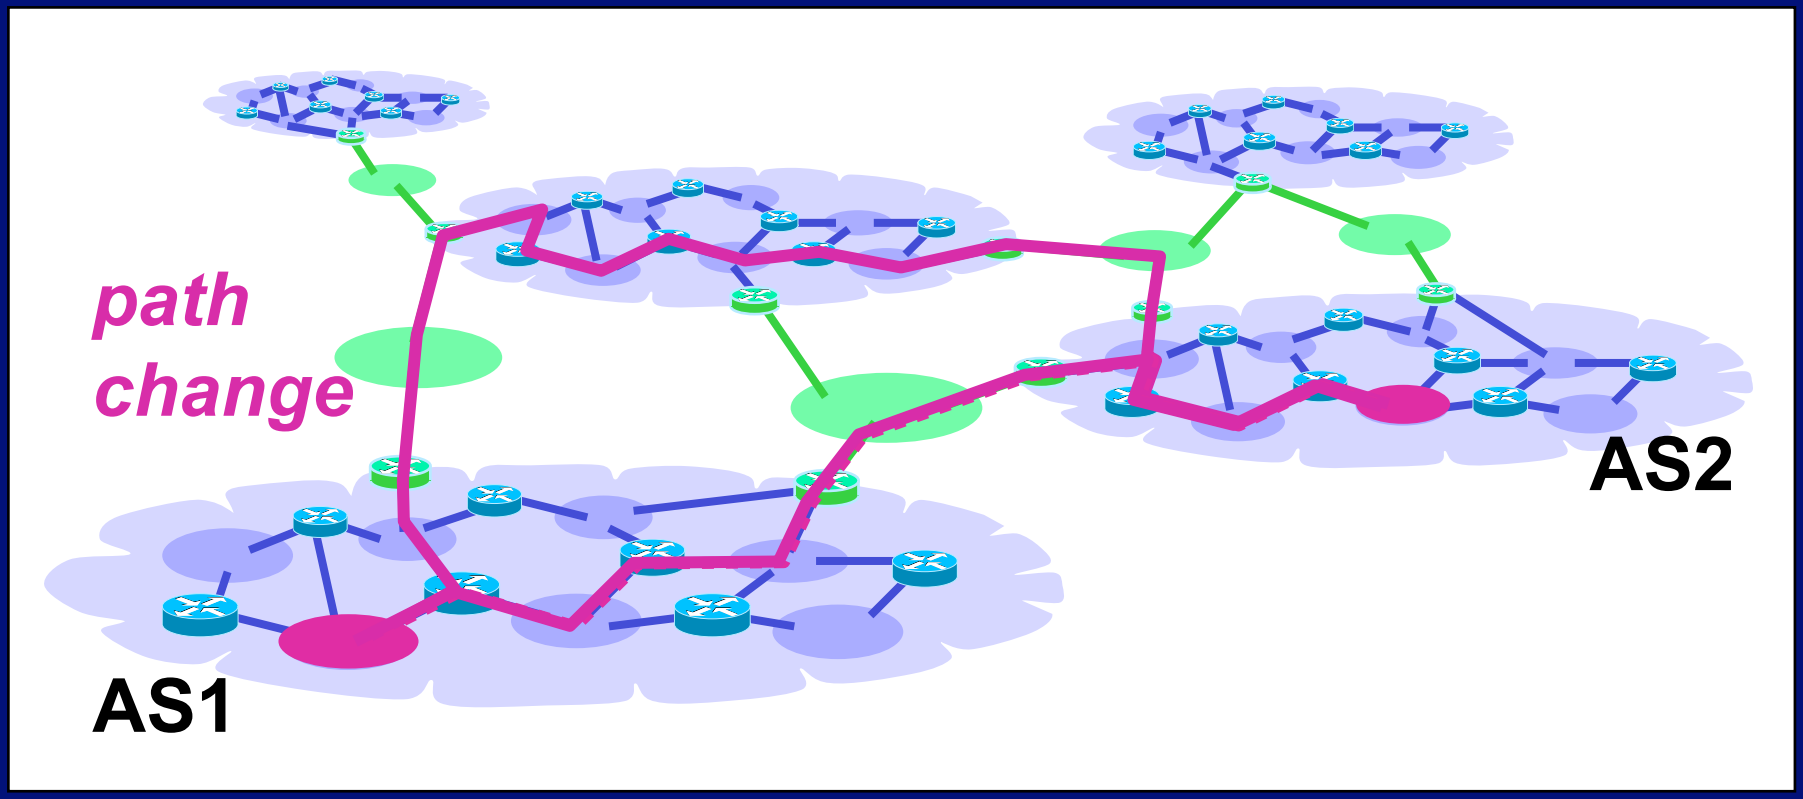

Tracking Back the Root Cause of a Path Change in Interdomain Routing

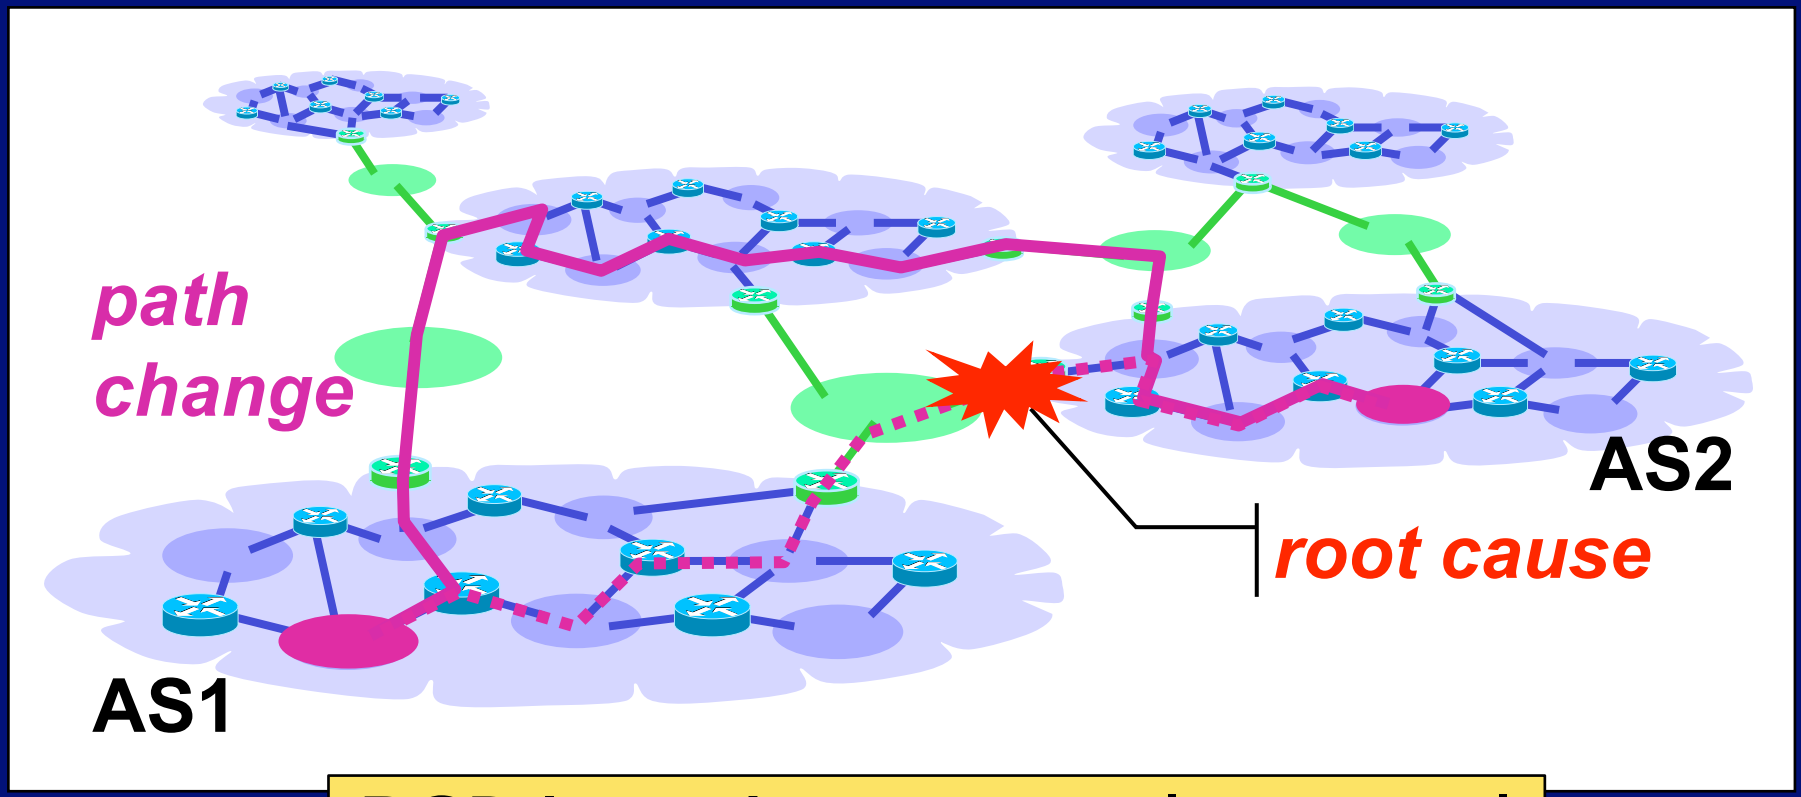

BGP is an incremental protocol

path changes *do* have an origin

Root causes - taxonomy

- Impact on the network
  - Events that affect AS-level topology
  - Events that affect routing behavior
    - policy changes
    - intra-domain events

BGP handles network dynamics for us, so … why should we even bother?

Motivation

- *Understand Internet dynamics*
  - Assess/debug network configurations
- *Economy*
  - Improve reliability and performance
- *Forensic analysis*
  - Identify, locate, investigate network outages

Previous approaches, in a nutshell

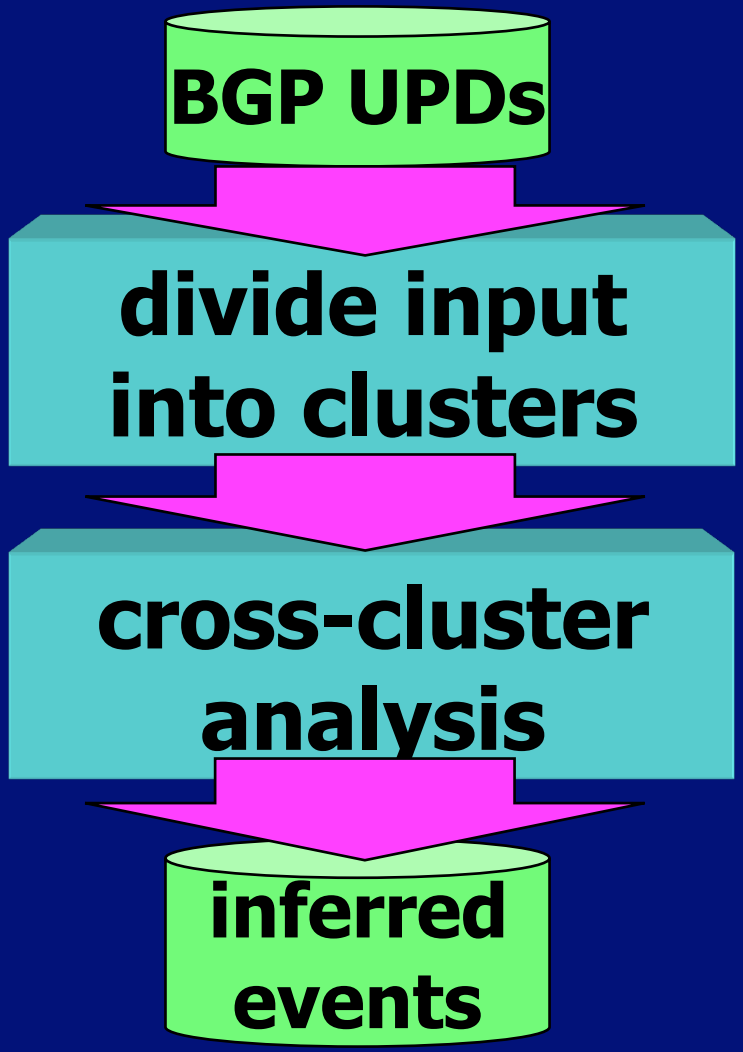

update correlation
[Feldmann2004,Caesar2003]

Principal Components Analysis
[Xu2005]
Learning-based
[J.Zhang2005]
Wavelet Transform
[K.Zhang2004]

Public BGP data sources

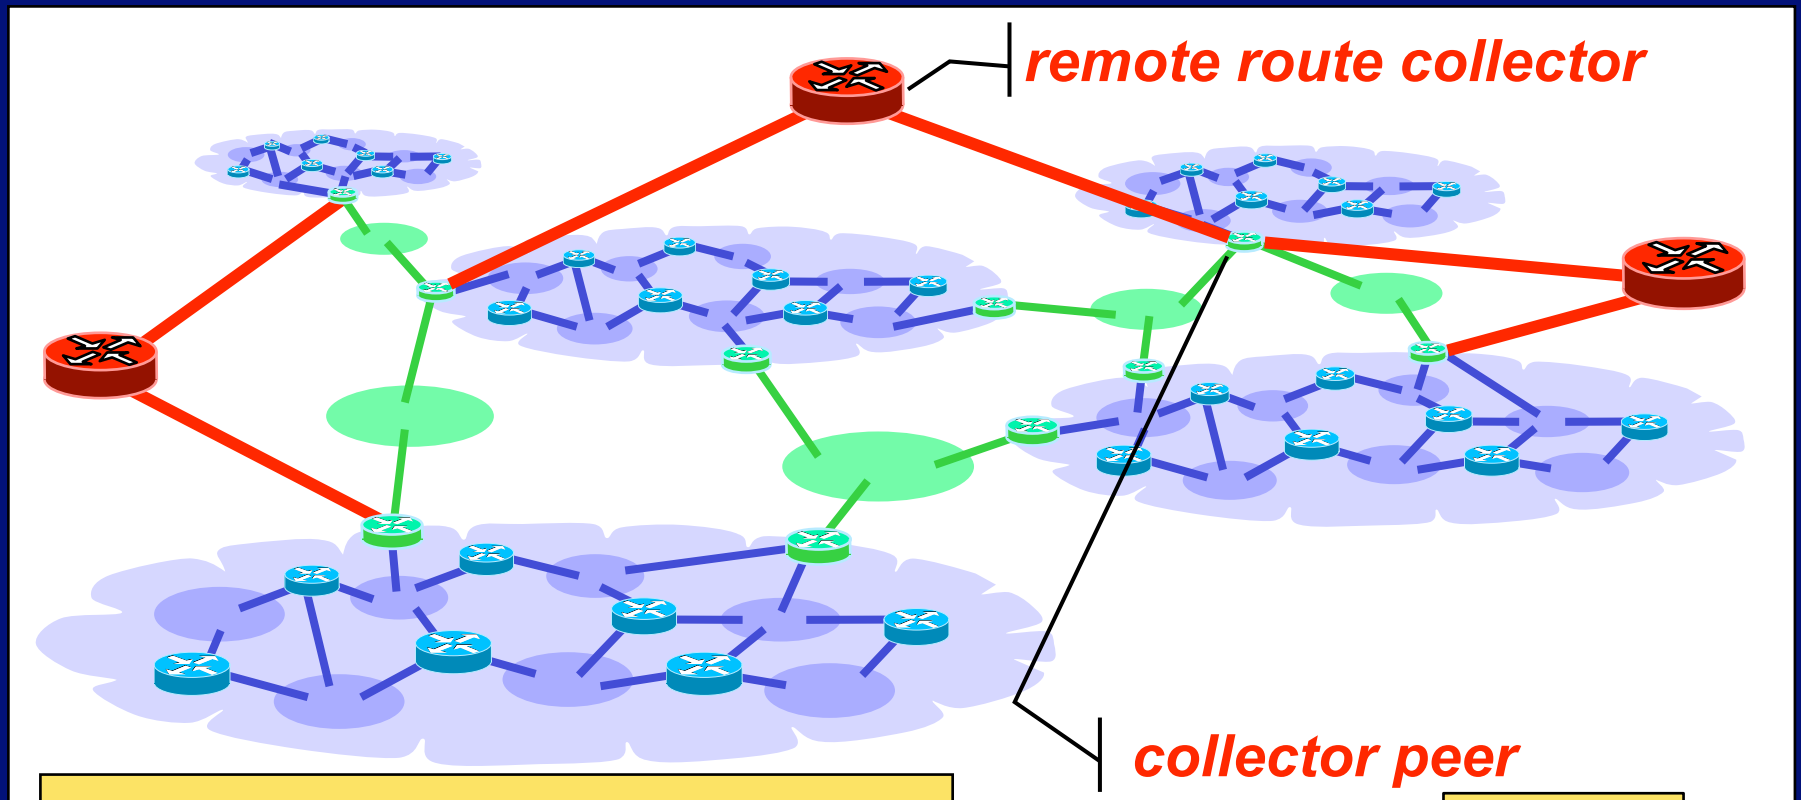

Routing Information Service
Route Views Project

~600

Why is it so hard?

- *BGP issues*
  - undisclosed policies (economical & political relationships)
  - huge network (26k ASes, 230k prefixes)
  - complex dynamics
- *Data-set issues*
  - sheer size (3GB/month)
  - unreliable collector peers
  - partial coverage of the Internet

Our contribution

- Flow-based model of a path change
  - derived from network flow theory
    - see, e.g., [Ahuja-Magnanti-Orlin93]
- Methodology to identify the root cause of a path change
- Prototype tool to support the methodology

A model for routing changes (1)

RIB snapshot => flow

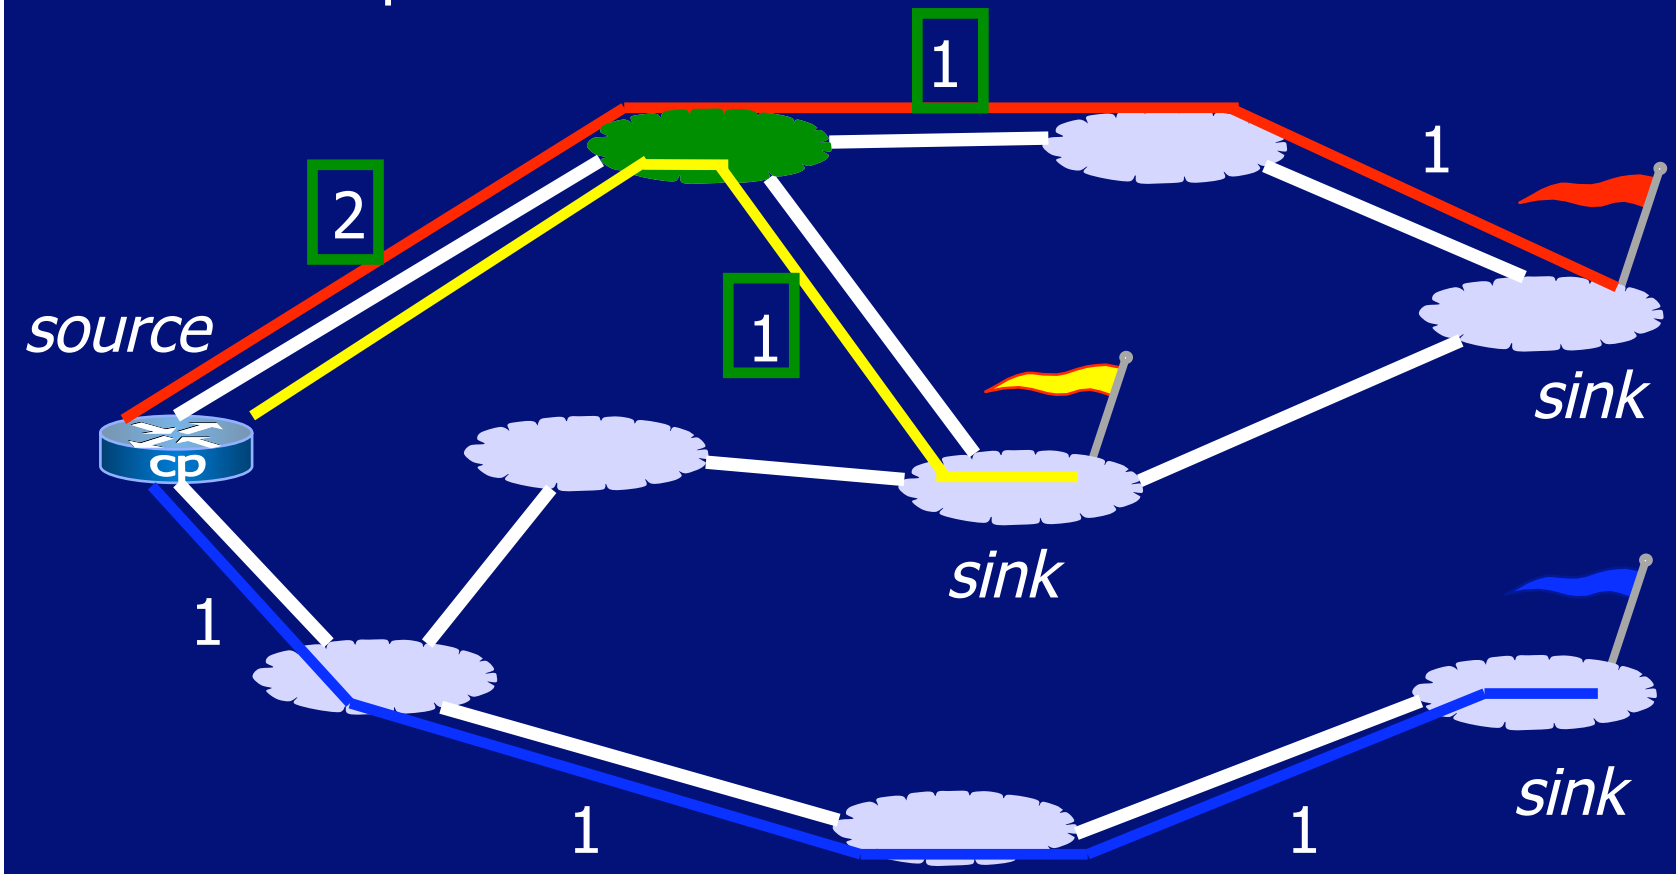

A model for routing changes (2)

RIB variation => flow

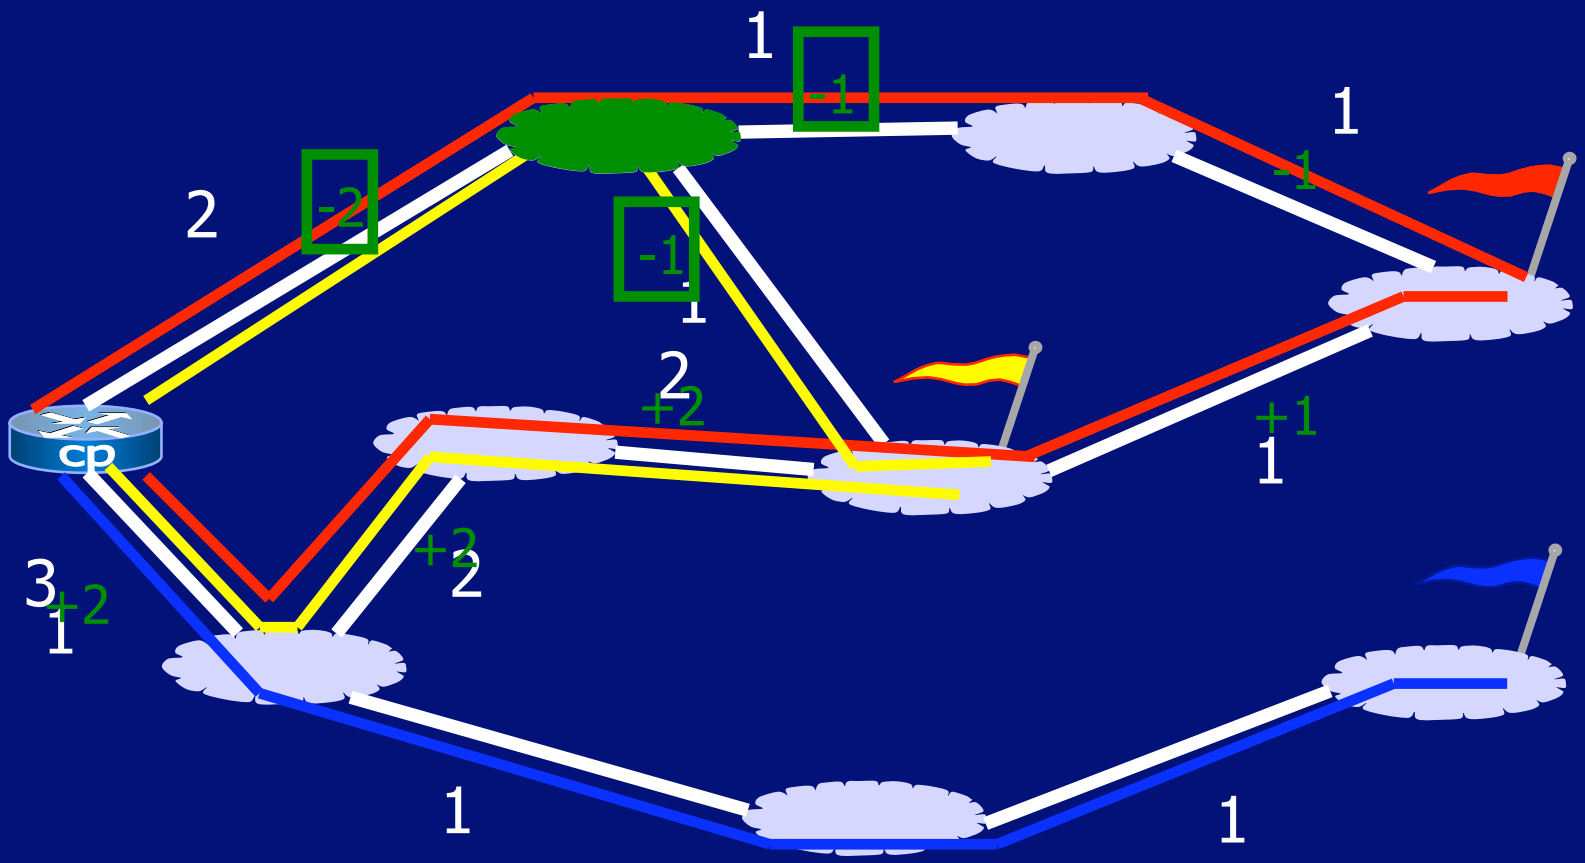

Local rank

- Local rank of a cp [Lad-Massey-Zhang2004]
  - # of prefixes on each edge
    - reflects the perspective of the cp
    - behaves like a flow system
    - depends on the specific cp
    - noise-sensitive
- How to account for multiple vantage points at the same time?

Global rank

- Idea:
  - combine different vantage points
    - merge different perspectives
  - consider distinct prefixes
- Global rank
  - # of *distinct* prefixes on an edge
- Pros and cons
  - aggregates multiple cps 
  - less noise-sensitive 
  - no locality 
    - can miss localized variations

Our approach - overview

- Input: a single BGP path change
- Methodology:
  1. check cp status
  2. identify macro-events by inspecting global/ local rank over time
  3. identify smaller events by looking for unusual flow variations
- Output: a set of links (*candidate set*)

Step 1: collector peer check

- Ignore data from unreliable cps
  - long-term unreliability: inconsistencies
  - short-term unreliability: session resets
- Locate inconsistencies
  - RIB dumps do not match update streams
- Locate session resets
  - log file, if any
  - identify BGP table transfers

Step 2: macro-events detection

- macro-events are events affecting the AS-level topology
- macro-events map to rank evolution patterns
  - link fault
  - link restoration
- Inspect the evolution of Global and Local ranks over time
  - compare with avg values to identify patterns

Macro-events detection: an example

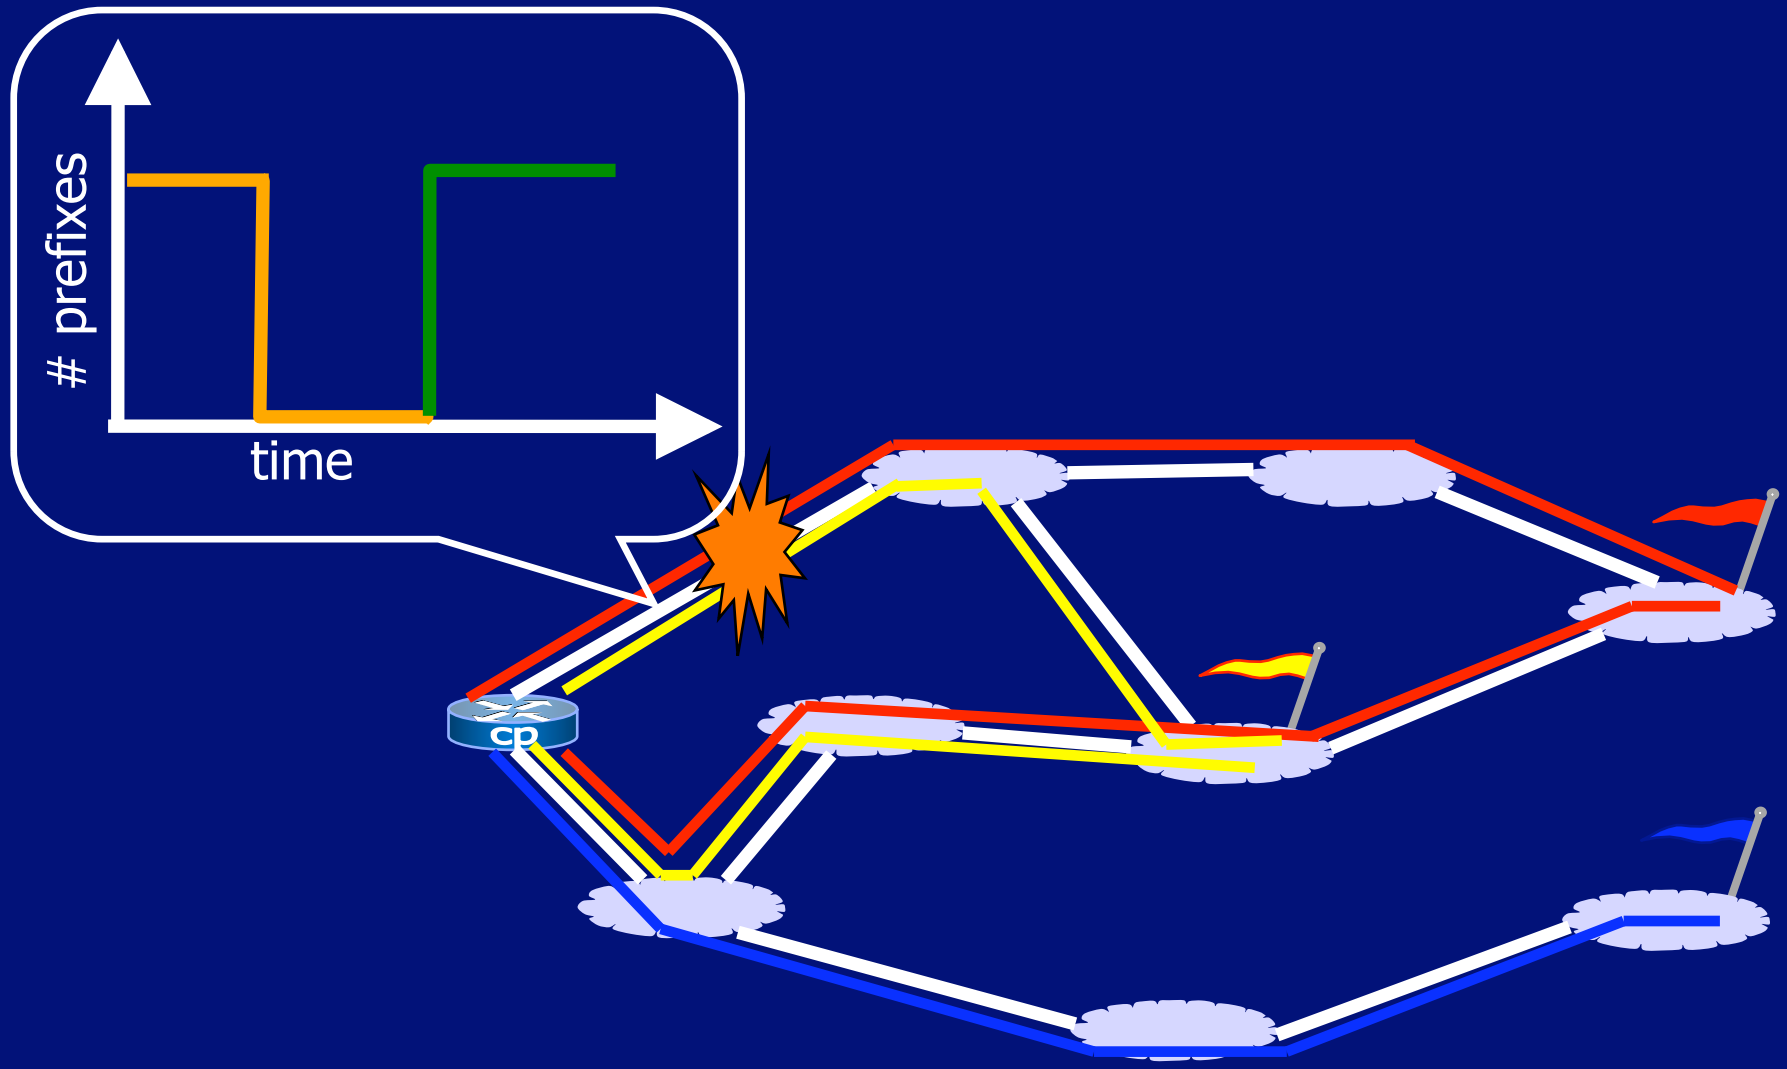

Step 3: fine-grained analysis

- Compute flow variations
  - restrict to links in the old (new) path
  - locate nodes with max inflow/outflow
  - output links in the subpath(s) containing them
- Intuition
  - nodes that move most flow are likely to be involved in the cause
  - distinct events probably do not affect the same portion of the network at the same time

Fine-grained analysis: an example

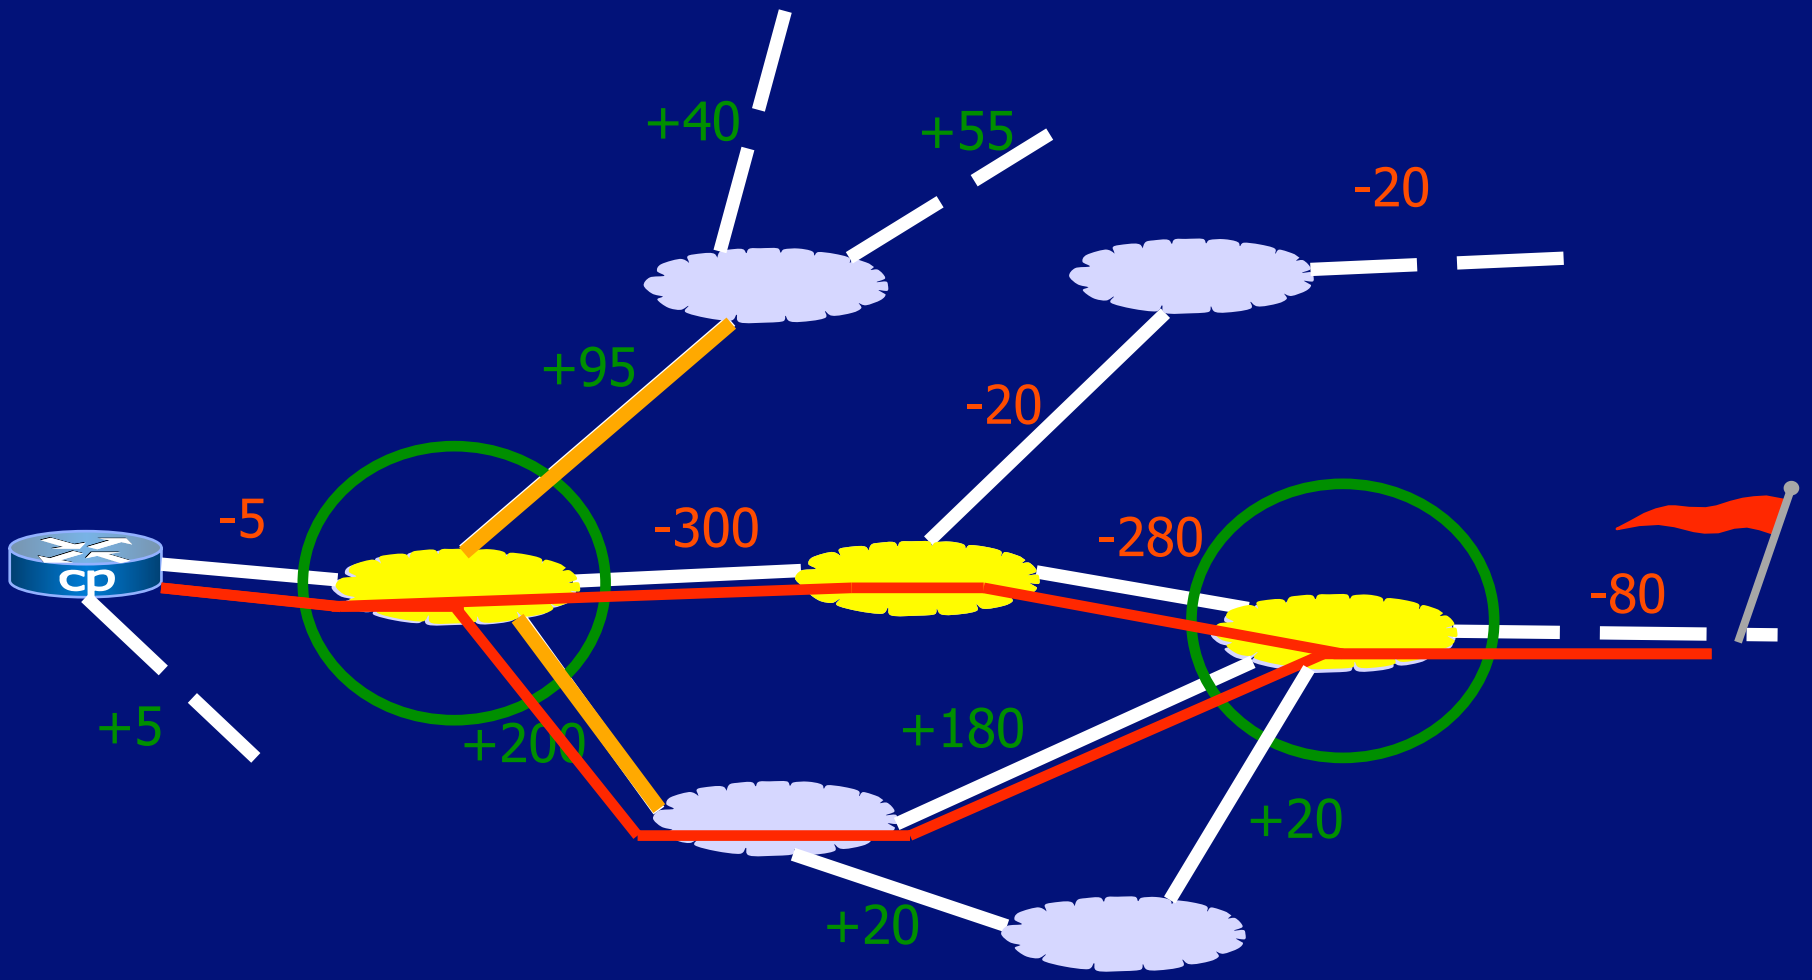

Simulation

- *Set-up:* real Internet topology using C-BGP [Quoitin-Uhlig2005]
  - 25k ASes (CAIDA)
  - 52k links (CAIDA + policies)
- *Simulated events*
  - type of event
    - link fault/restoration
    - loc-pref + hard reset
    - loc-pref + soft reset
  - location in the topology
    - Tier1-Tier1
    - Tier1-transit AS
    - transit AS-transit AS
    - transit AS-stub AS

Simulation - results (1)

- *Accuracy*

  updates whose root cause appears in the candidate set

  - link fault/restoration: **100%**
  - loc-pref + hard reset: **99%**
  - loc-pref + soft reset: **93%**

Simulation - results (2)

- Impact of the event location

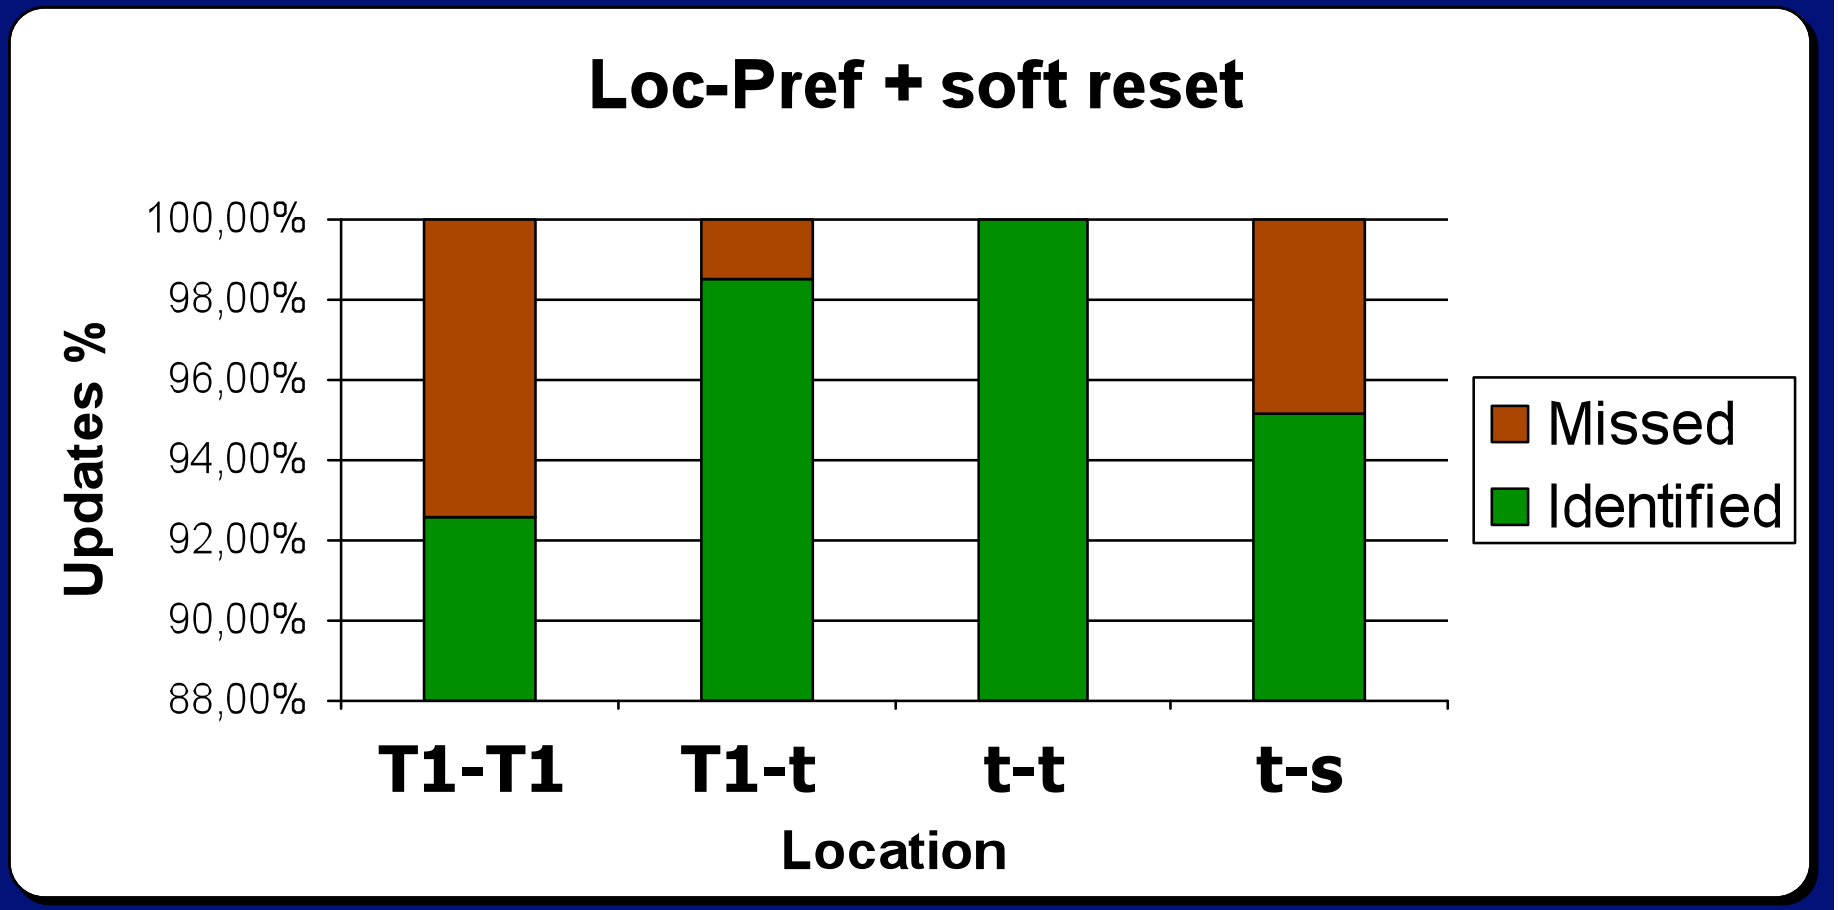

Simulation - results (3)

- *Precision*
  candidate set size (CDF)

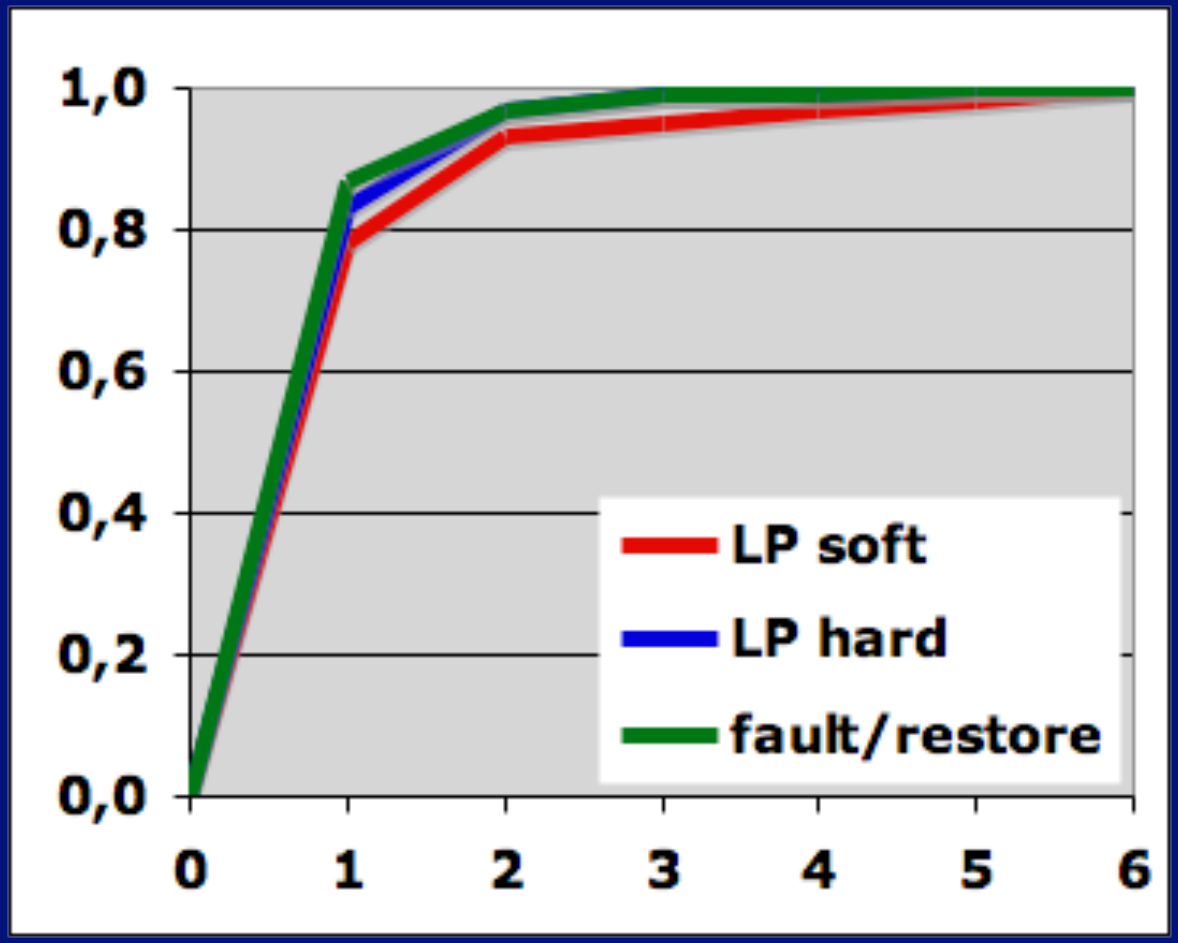

Conclusions & future work

- Summary
  - Flow-based model of a path change
  - Methodology to identify the root cause of a path change
  - Internet scale simulation
  - Prototype tool to support the methodology
- Future work
  - Extend the model and the methodology with new patterns
  - Fully automate the methodology in the tool

Thanks!

- Questions?

Real world data

- No training dataset
- Taiwan earthquakes, Dec 2006
  - bug in RIS collectors
- Mediterranean cable cut, Jan 2008
  - FLAG peerings down
  - most of macro-events are located in surrounding areas