Network Tomography for Fault Diagnosis

Renata Teixeira
CNRS and UPMC Paris Universitas

with
Italo Scota Cunha (LIP6, Thomson)
Nick Feamster (Georgia Tech)
Christophe Diot (Thomson)

Detection and identification of network blackholes

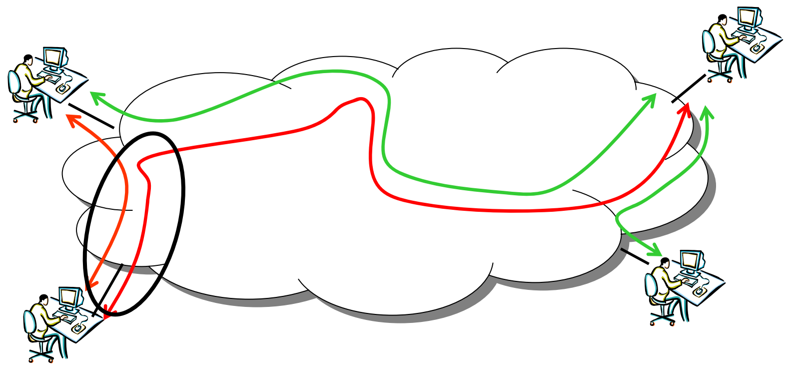

Detection: continuous path monitoring

Identification: tomography

Problem: Too many false alarms

- Applying tomography on raw measurements
  - PlanetLab: one alarm per minute
  - Thomson VPN: one alarm every two minutes
- Why?
  - Loss can be transient, topology can change
  - Different monitors see different conditions

Detection: transient losses vs. persistent failures

- Monitors ping a set of destinations
- Lost pings can have different causes
  - Congestion
  - Routing changes
  - Persistent failures
- How to know which losses are persistent?

Failure confirmation

- Upon detection of a failure, trigger extra probes
- Goal: minimize detection errors

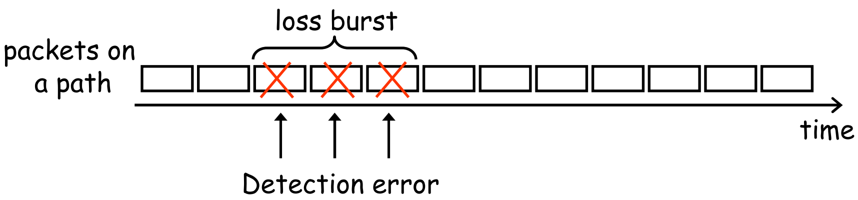

Probing strategy for failure confirmation

- Which probing process?
  - Assume link losses follow a Gilbert process
  - Periodic probing minimizes detection errors
- How many probes?
  - Confirm failures with a target detection-error rate
  - Assume independence and a given a loss rate
- How much time between probes?
  - Reduce chance that probes fall on the same loss burst
- Tradeoff: detection error and detection time

Identification through binary tomography

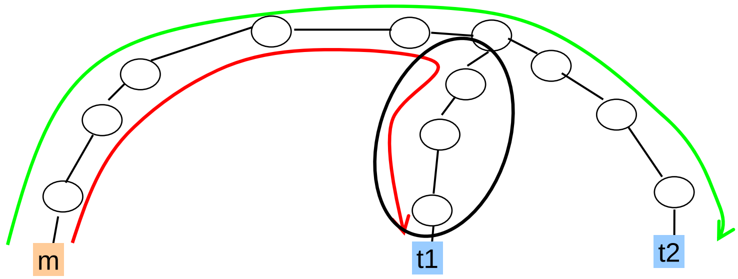

- Given: topology and end-to-end path statuses
- Find the smallest set of links that explain observations

Lack of synchronization leads to inconsistencies

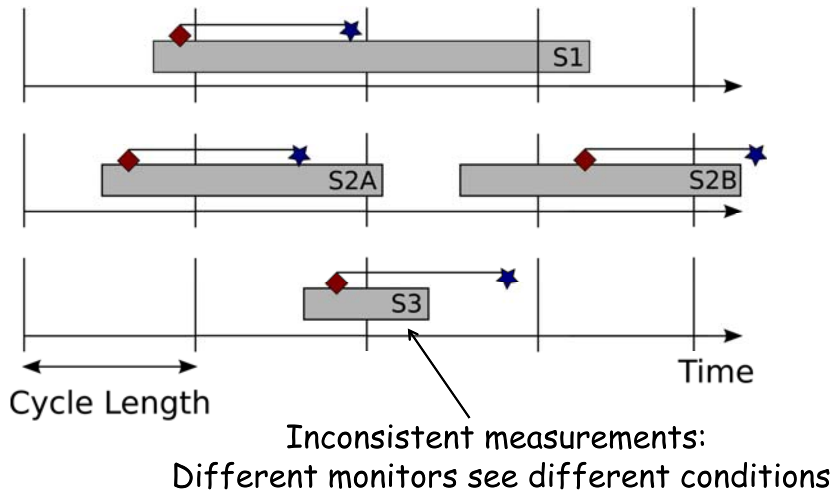

Achieving consistency: Aggregation strategies

- Basic strategy
  - Waits for one cycle
- Multi-Cycle strategy (MC)
  - Waits for n cycles with identical path statuses
- Per-Path Multi-Cycle strategy (MC-path)
  - Only considers paths that are down for n cycles

Evaluation

- Evaluation is challenging
  - Need ground truth and realistic environment
- Analytic modeling
  - Understand limits of the system
- Controlled Experiments: Emulab testbed
  - Realistic environment
  - Control over failures
- Wide-area Experiments: PlanetLab, Thomson
  - Real losses and failures, but no ground truth

Failure confirmation reduces false alarms

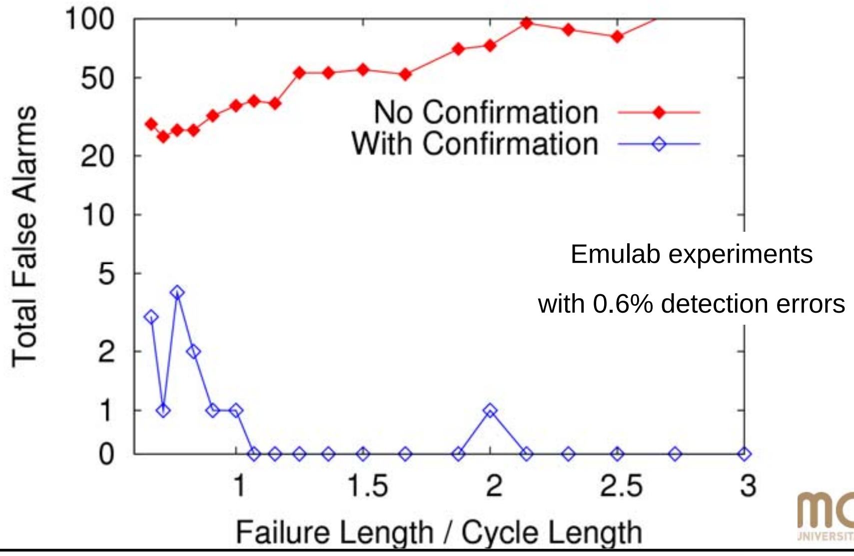

Aggregation strategies identify most long failures

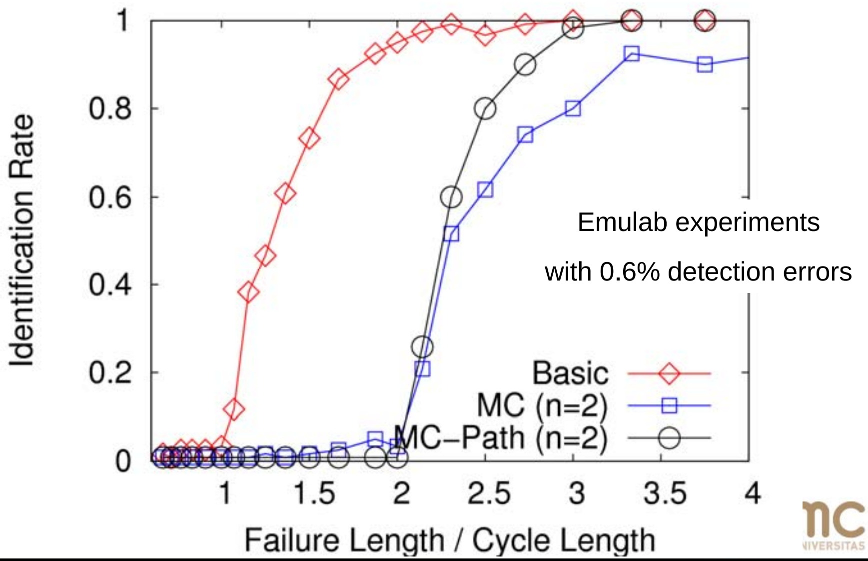

Multi-cycle aggregation reduces false alarms

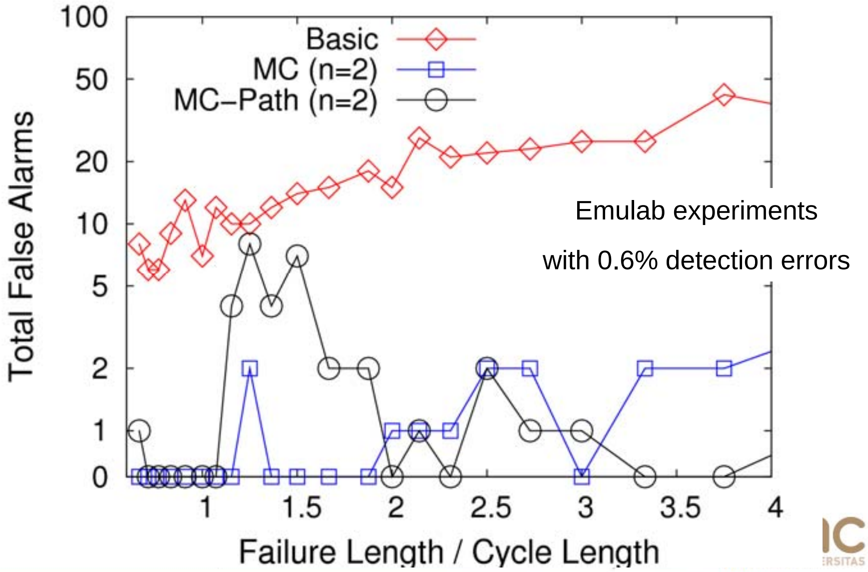

Number of alarms in wide-area experiments

	PlanetLab		Thomson	
	NO CONF.	CONF.	NO CONF.	CONF.
BASIC	16,261	16,260	6,256	225
MC	—	—	13	17
MC-PATH	251	135	27	27

- PlanetLab
  - 39,800 paths
  - Cycles: 60 seconds
- Thomson
  - 56 paths
  - Cycles: 5 seconds

Summary

- Tomography with raw data leads to false alarms
- Two techniques to reduce false alarms
  - Failure confirmation
    - Distinguishes transient losses and persistent
  - Aggregation
    - Combines measurements from different monitors

Two deployment scenarios

NanoDataCenters

grenouille.com
la météo du net depuis 2000