Measurement Methods for Fast and Accurate Blackhole Identification with Binary Tomography

Ítalo Cunha

Thomson and UPMC Paris Universitas

Renata Teixeira

CNRS and UPMC Paris Universitas

Nick Feamster

Georgia Tech

Christophe Diot

Thomson

THOMSON TECHNICOLOR c'mon NanoDataCenters

Users can collaborate to identify failures

- Users may want to assess ISP performance
  - Rank ISPs to choose best quality of service
  - Blame them for failures

- Users have limited access to network resources, but
- Can collaborate to identify failures
  - Limited to end-to-end probes

Monitoring helps applications improve QoS

- Improve hardware and deploy applications on gateways
- Need monitoring to perform peer selection and overlay construction
  - Connectivity
  - Bandwidth and jitter for streaming
  - Latency for games
- Need end-to-end probes to maintain overlay

Operators monitor to ease troubleshooting

- **Blackholes are challenging to troubleshoot**
- **Persistent failures that raise no alerts**
  - Router software bugs or misconfiguration
  - Problems in other networks

- **Blackholes are detected by loss of end-to-end connectivity**
  - Need for end-to-end measurements

Network tomography is promising

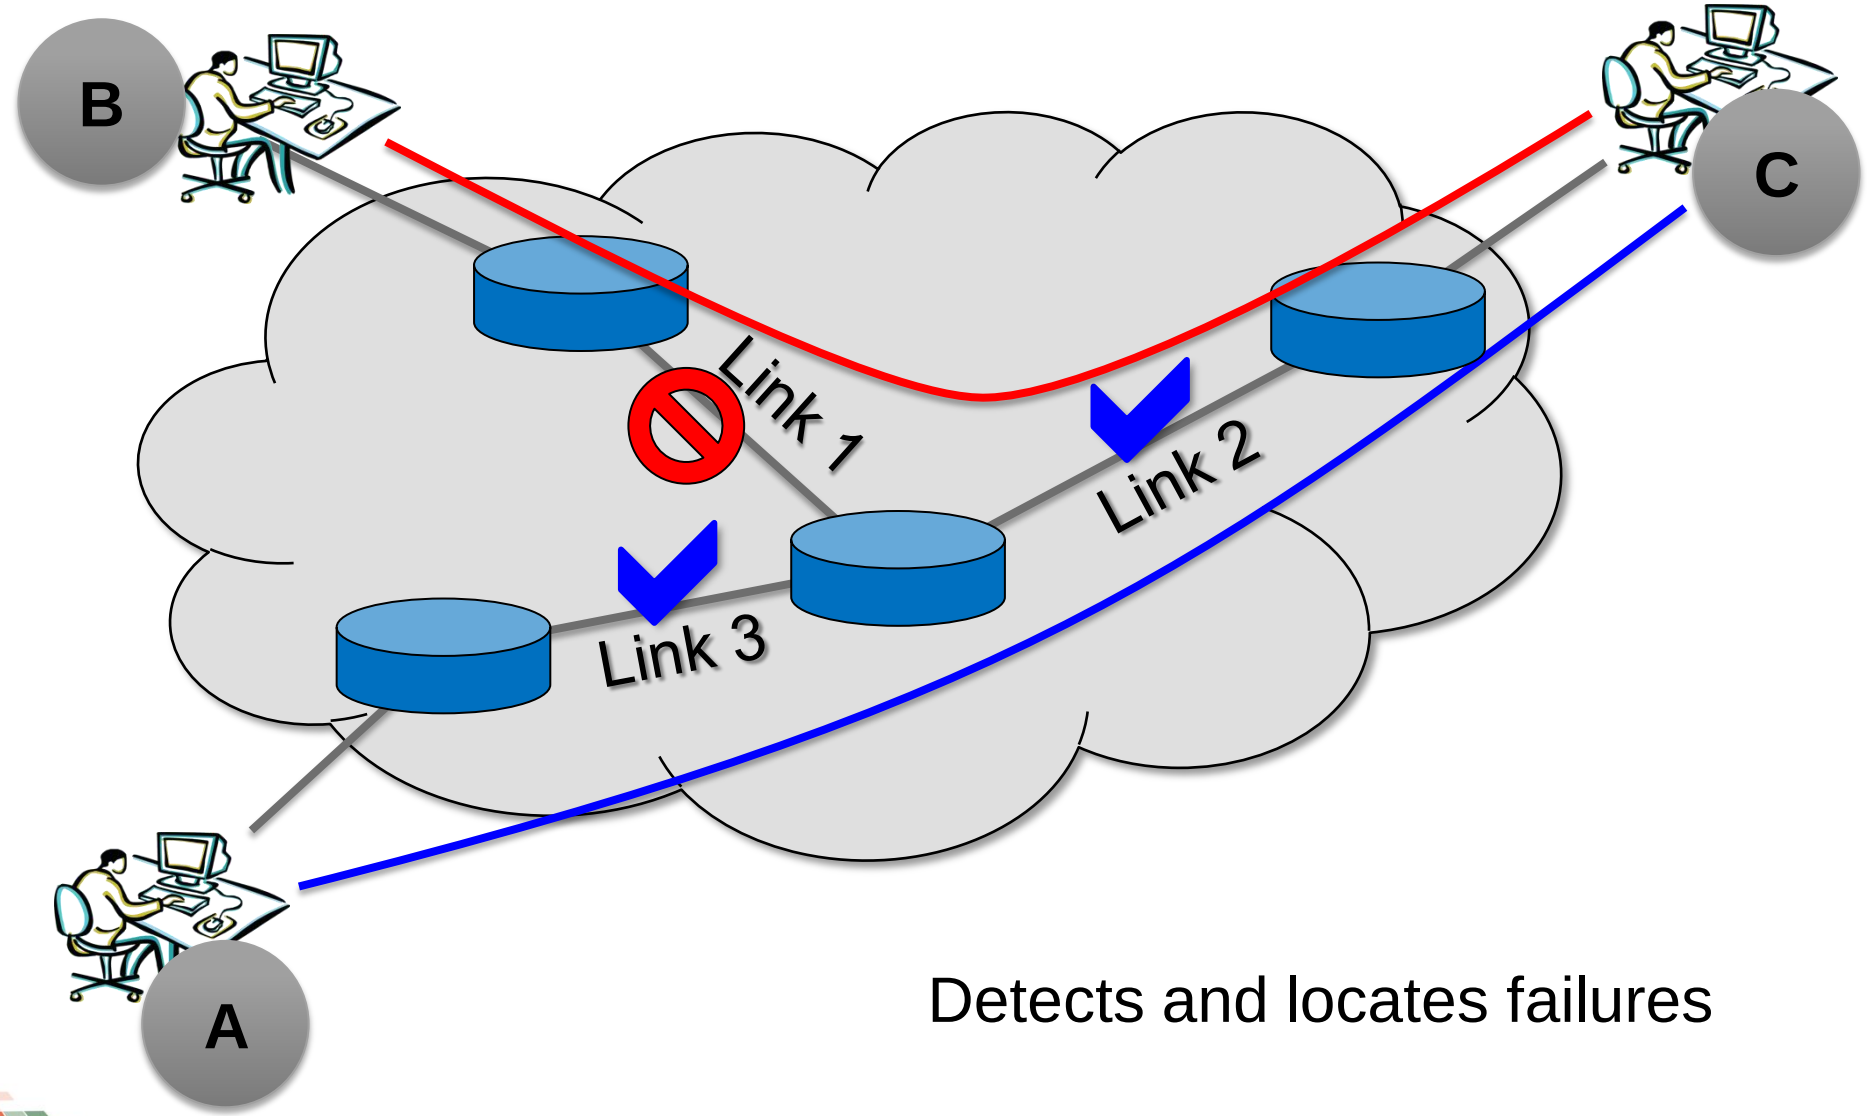

Detects and locates failures

Tomography steps

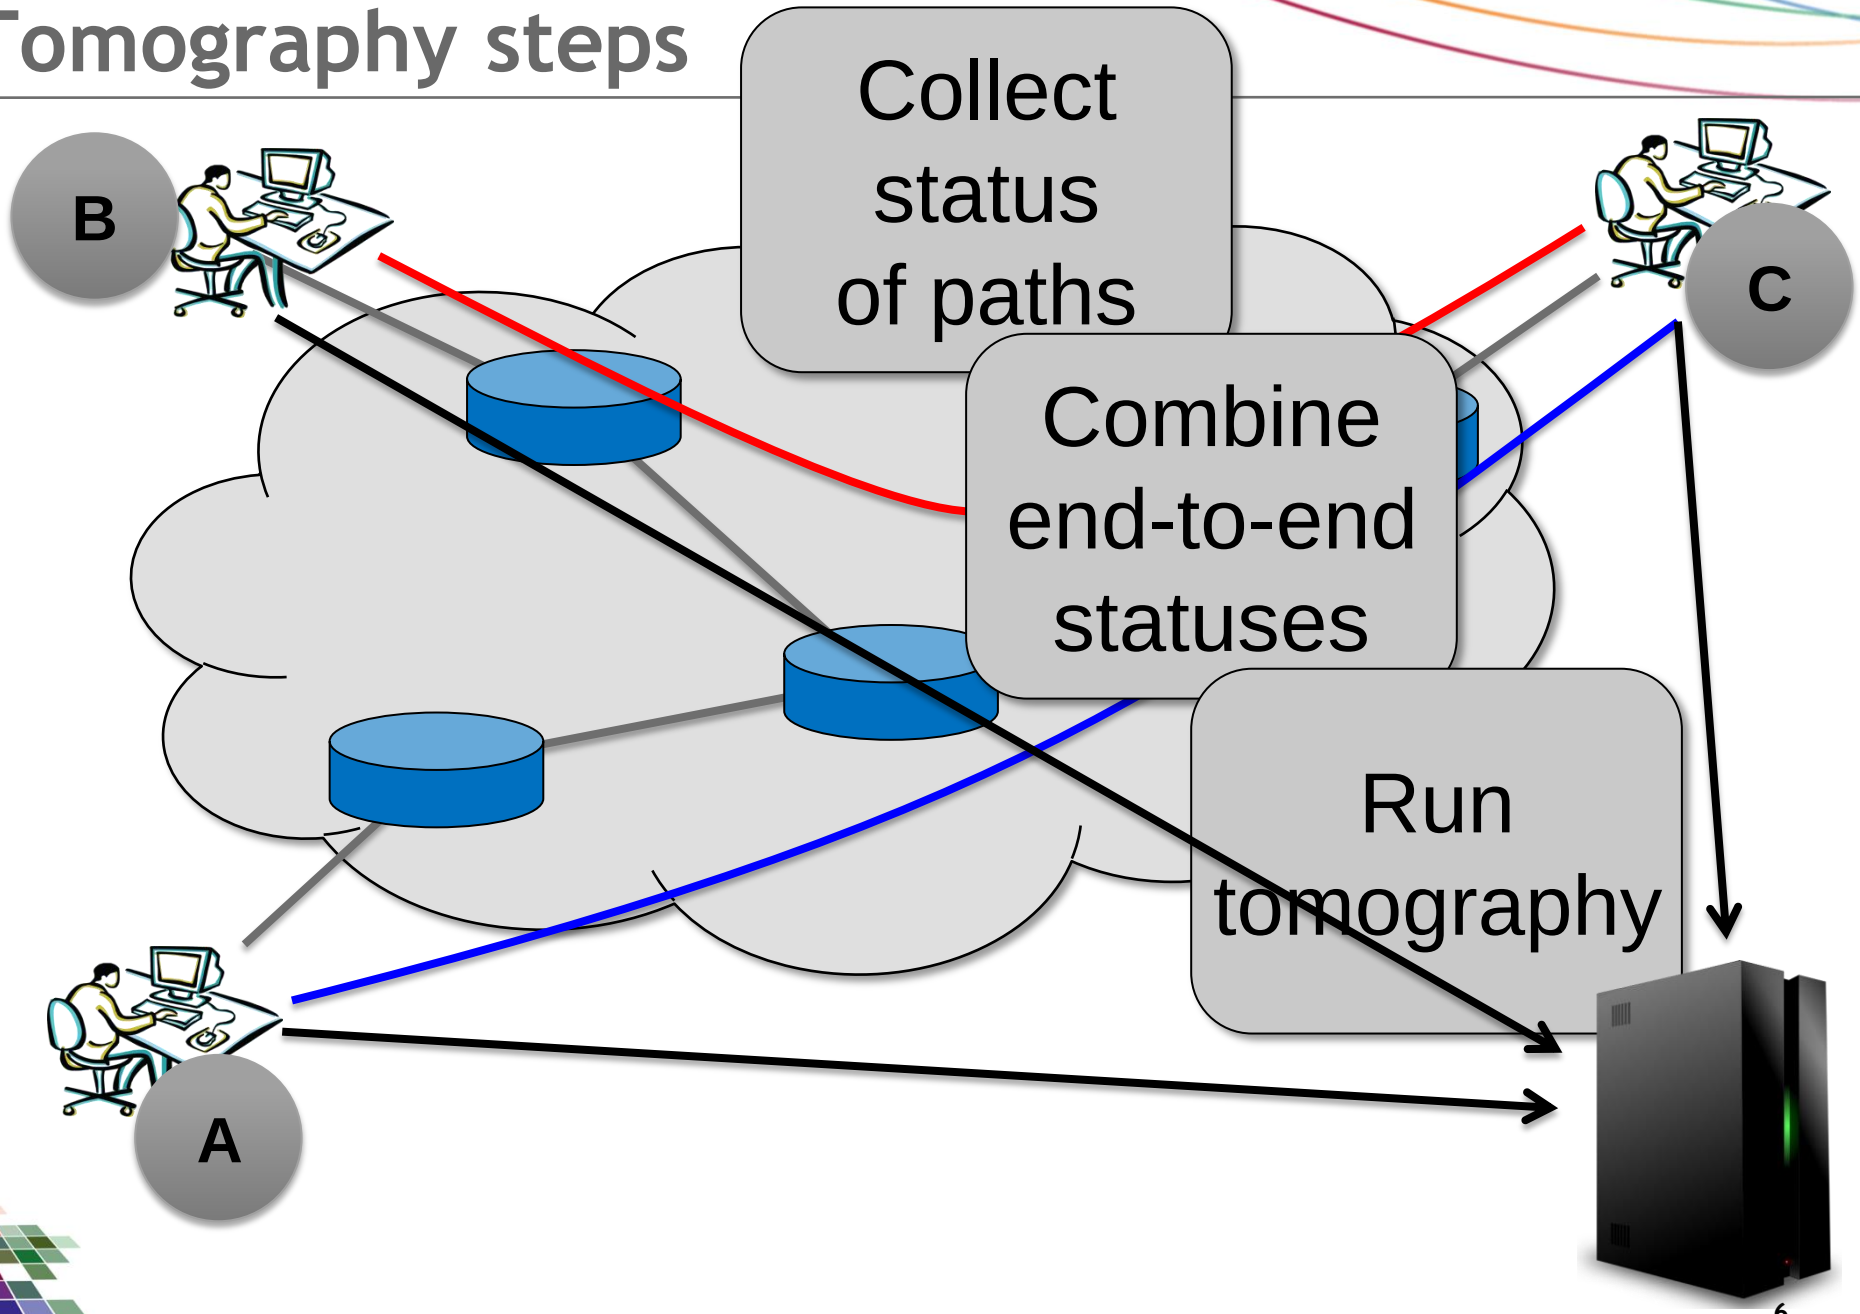

Tomography is sensitive to inaccurate inputs!

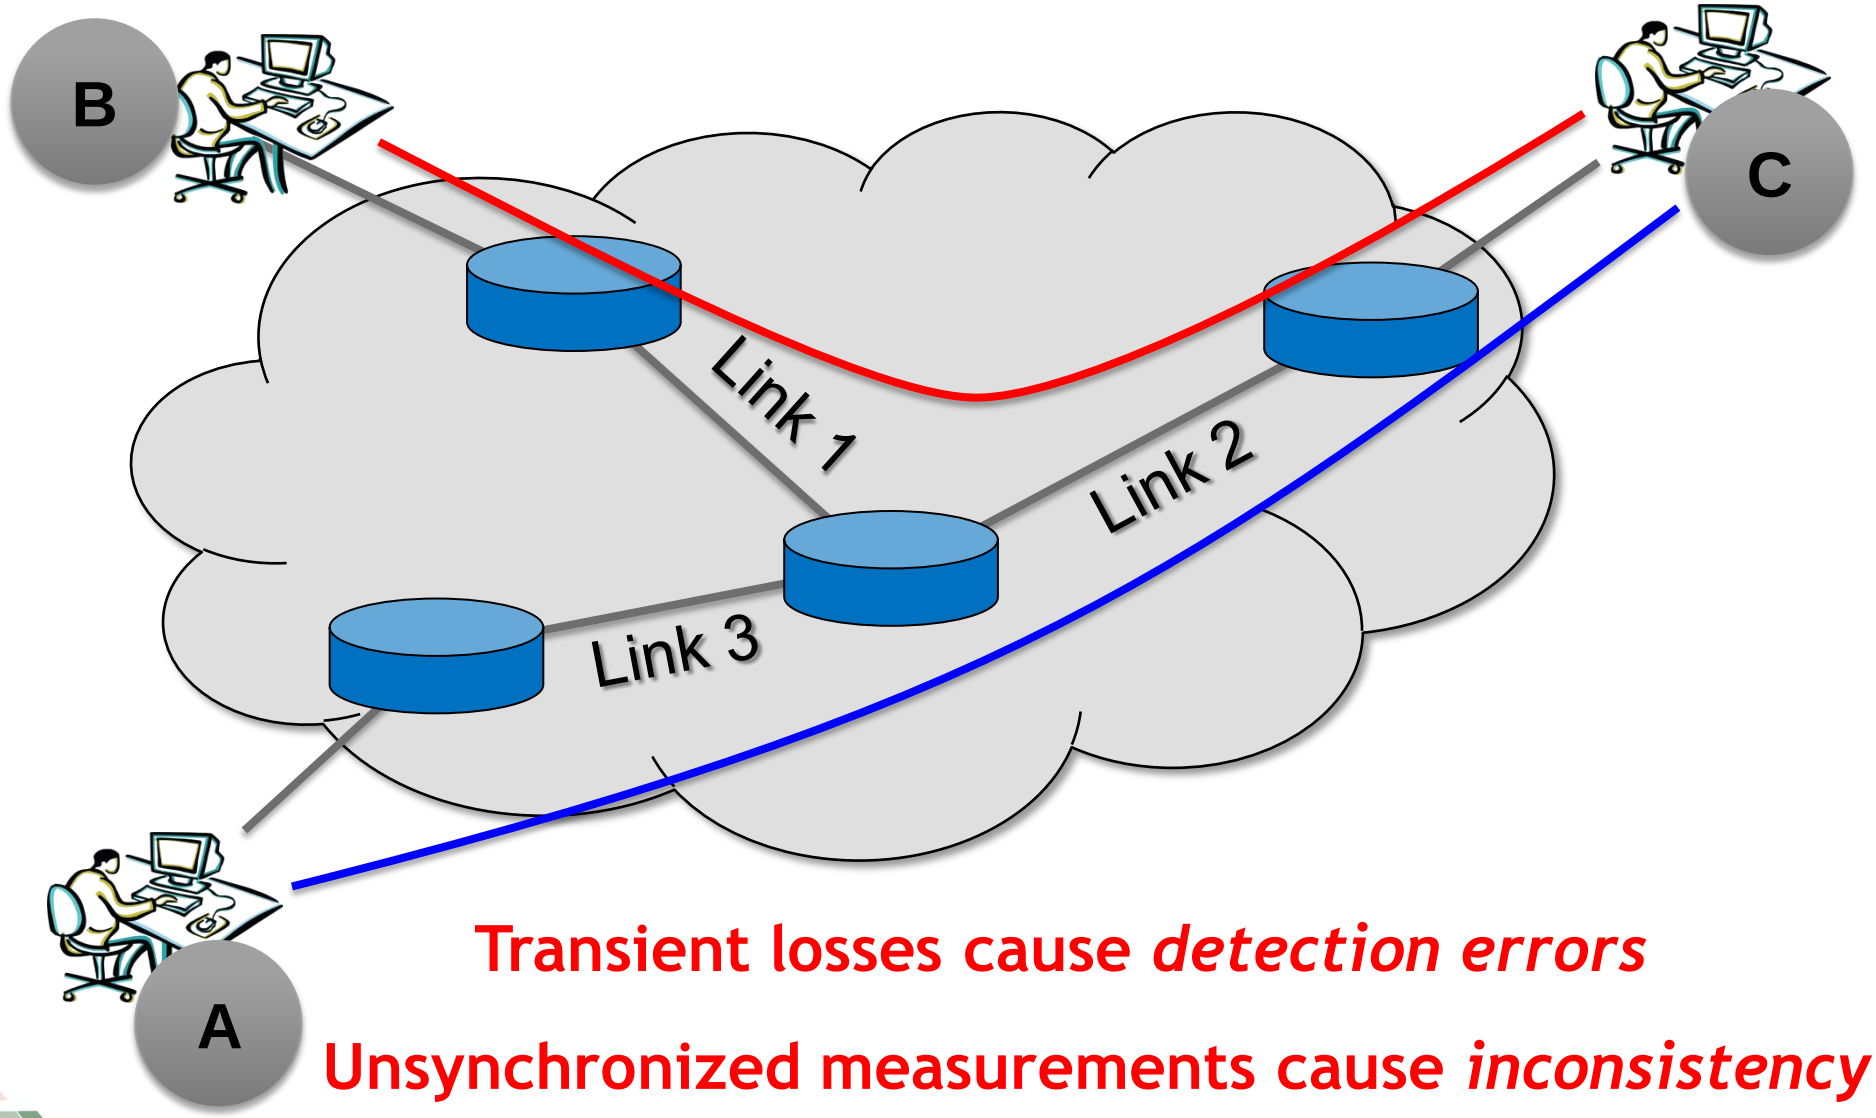

Transient losses cause *detection errors*

Unsynchronized measurements cause *inconsistency*

Remove inaccuracies to reduce false alarms

- **Failure confirmation**
  - Minimizes detection errors
  - Differentiates transient losses from persistent failures
- **Failure aggregation**
  - Aggregates unsynchronized measurement
  - Trades delay for consistency

Method

- **Analytical models**
  - How fast can we run confirmation?
  - What are the shortest failures we can identify?
- **Controlled experiments on Emulab**
  - Tomography is hard to evaluate: no ground truth
  - Measure failure identification rate and false alarms
  - Validate the analytic models
- **Deployment on PlanetLab and an enterprise VPN**
  - Assess the usefulness of the techniques in practice
  - No ground truth, but can still compare number of alarms

Failure Confirmation

THOMSON TECHNICOLOR c'mon NanoDataCenters

Dealing with probe losses

- We need low overhead and quick confirmation.

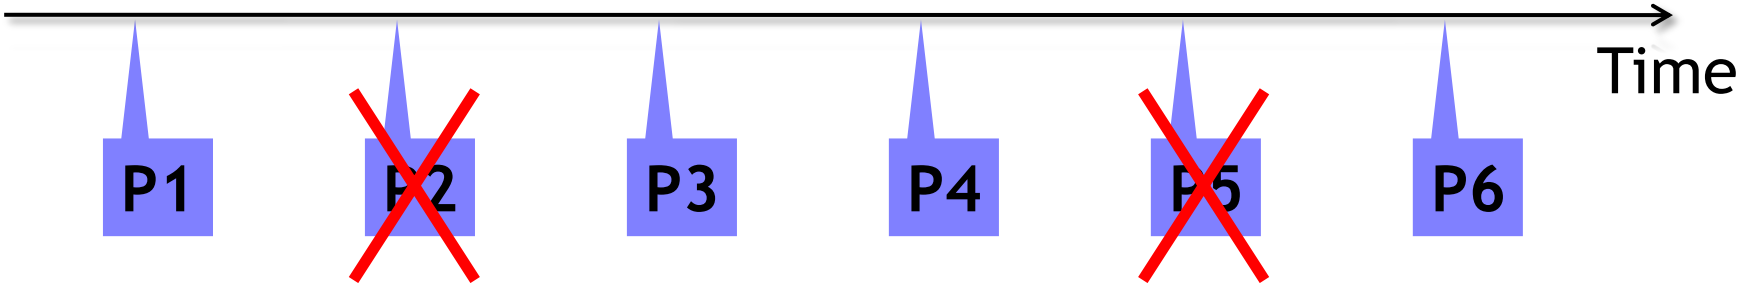

- Approach: send extra confirmation probes
  - How many?
  - When?

How and when to send confirmation probes

- **Bursty losses modeled as a two-state Gilbert model**
  - Two parameters: path loss rate and average burst length
  - Loss bursts durations are exponentially distributed
- **Periodic confirmation probes minimize detection errors**
- **Optimization models to calculate number of confirmation probes and spacing between them**
  - Minimizing total confirmation time
  - Minimizing number of confirmation probes

PlanetLab measurement setup

- 200 PlanetLab nodes probing each other for 12 days
- Paths are probed every 60 seconds
- Run two confirmation schemes simultaneously
  - 5 back-to-back probes
  - 10 probes spaced by 200 milliseconds

Effect of failure confirmation in practice

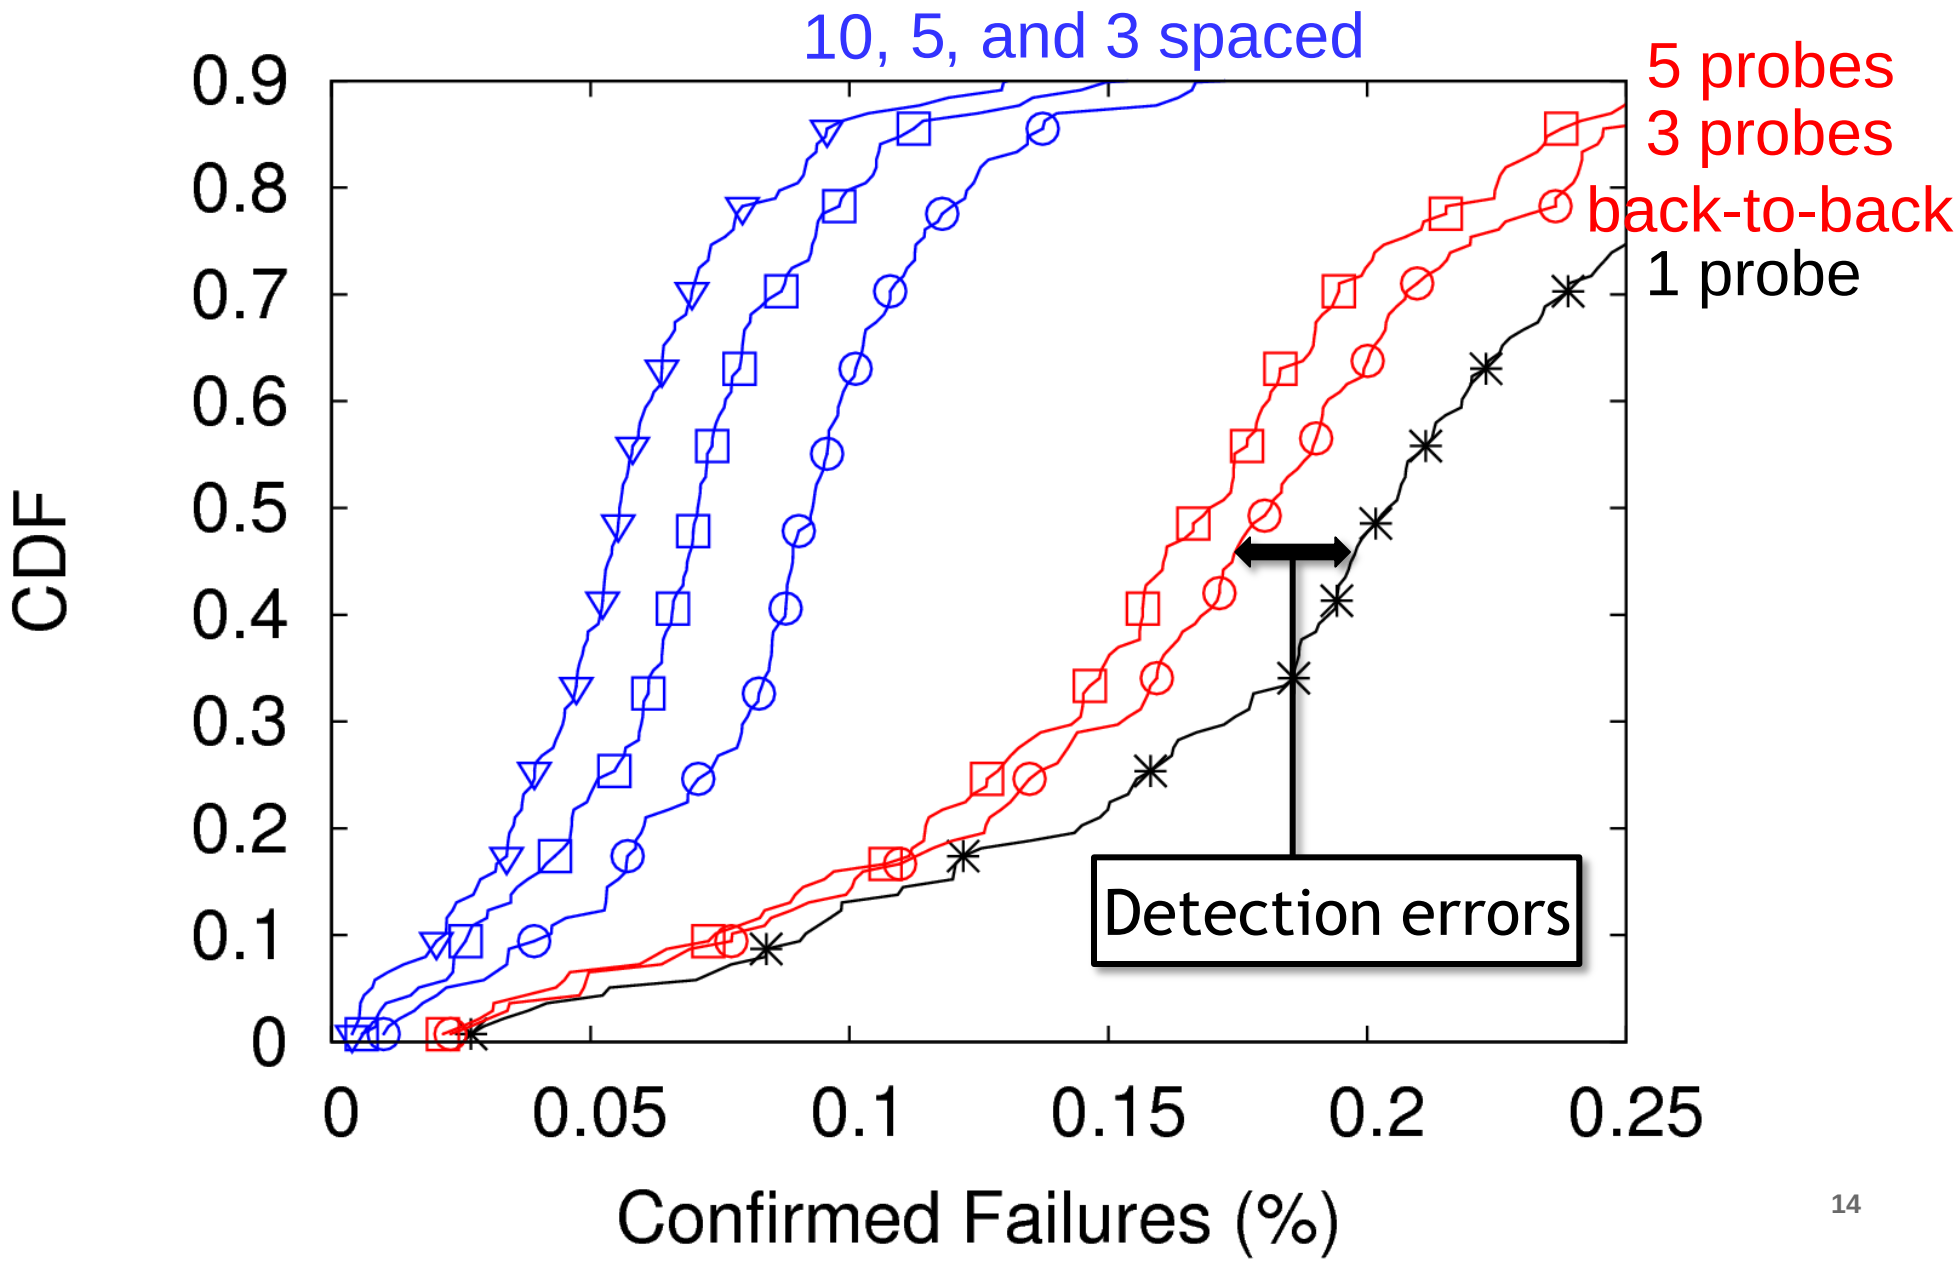

Failure Aggregation

THOMSON TECHNICOLOR c'mon NanoDataCenters

Short failures are impossible to identify consistently

- Very short failures

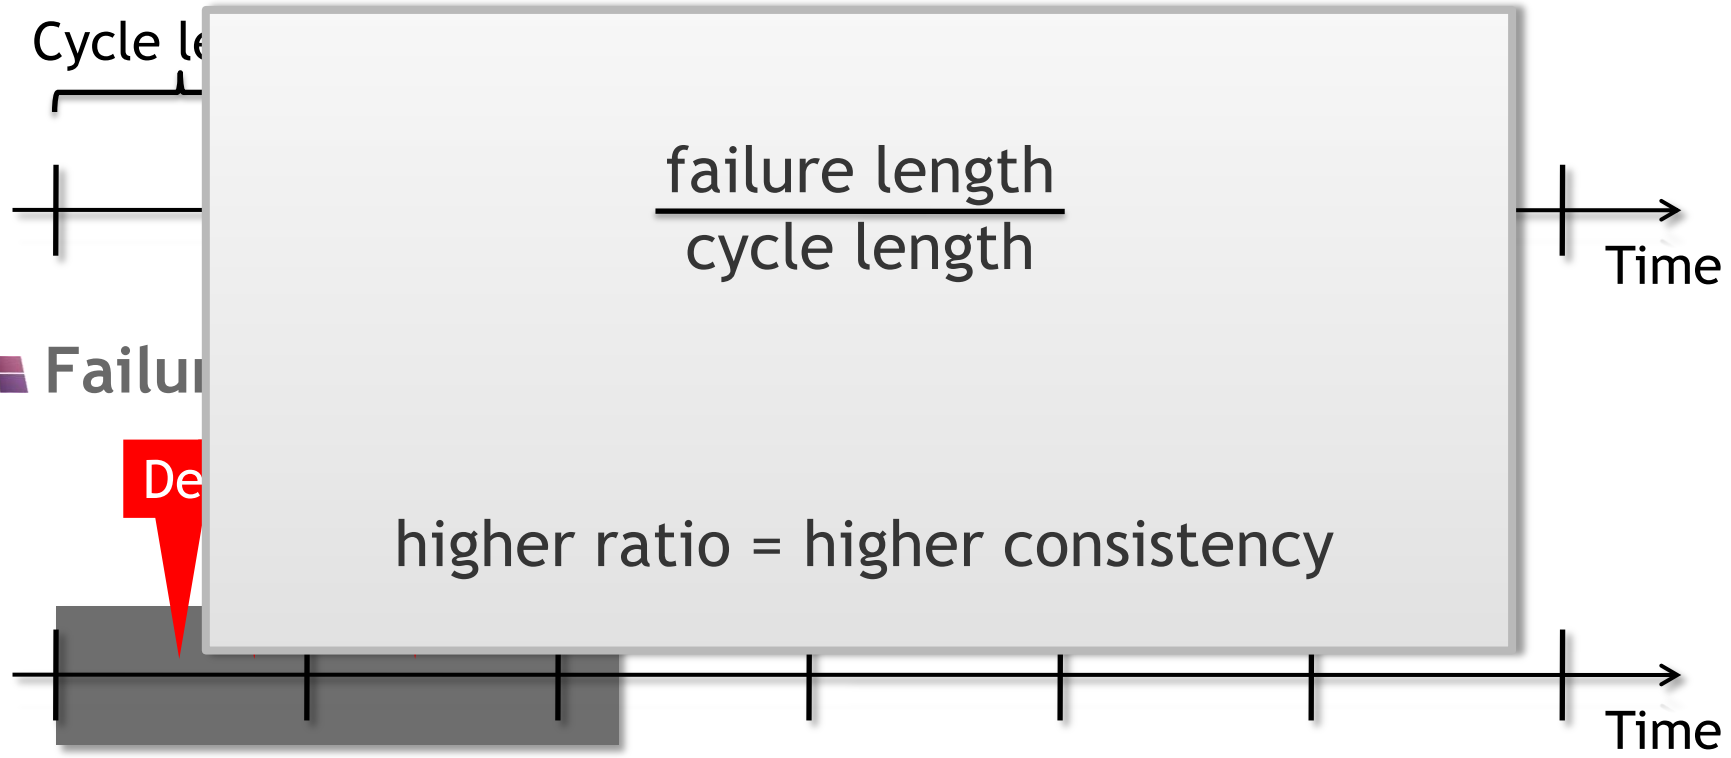

- Failu

Aggregation strategies

- **Basic aggregation**
  - After detecting a failure, measures all paths once and then runs tomography
- **Multi-Cycle aggregation (MC)**
  - Runs tomography only when n consecutive cycles have identical measurements
- **Multi-Cycle Noise-Tolerant aggregation (MC-Path)**
  - Runs tomography with paths down for n consecutive cycles

Putting everything together – real deployments

Number of alarms per day

	PlanetLab	Enterprise
BASIC	1500	20
MC	No alarm	1
MC-PATH	12	2

Failure identification rate - Emulab

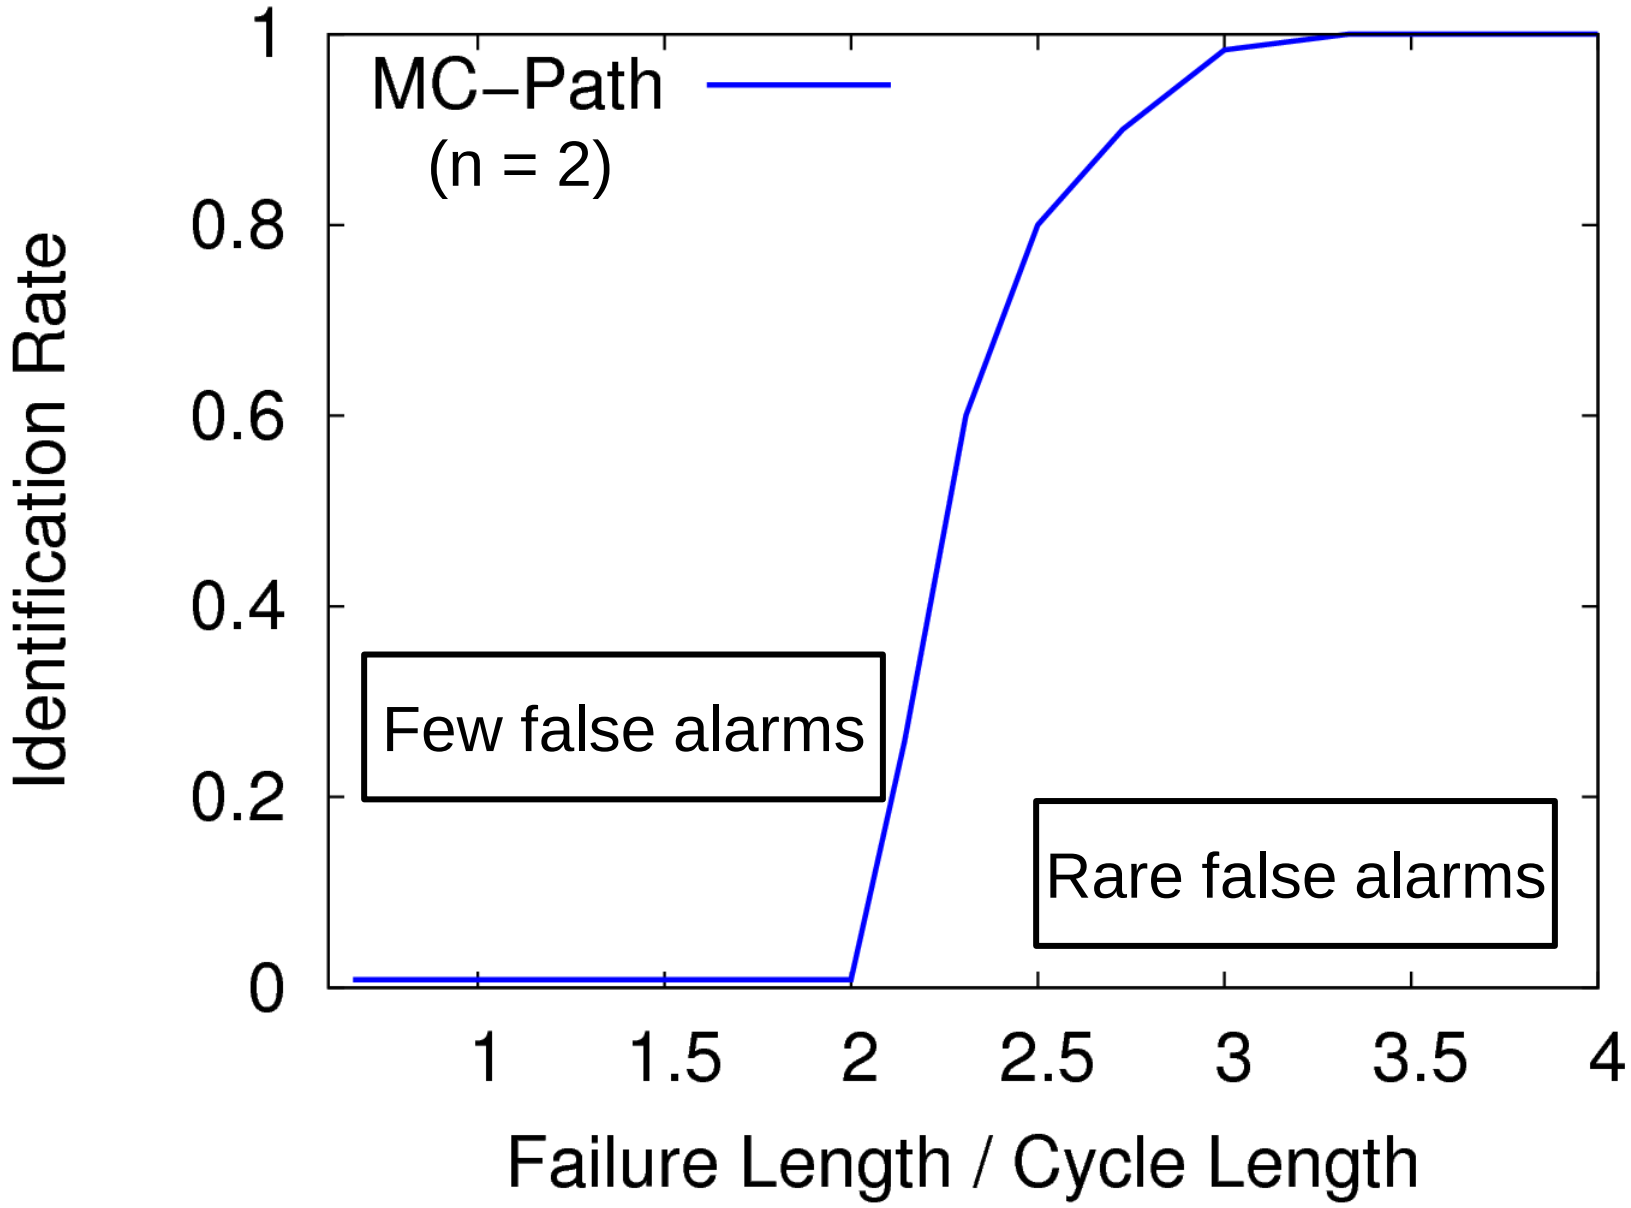

Summary

- **Tomography algorithms need accurate inputs**

- **Failure confirmation**
  - Differentiates transient losses from persistent failures
  - Minimizes the number of probes and delay

- **Failure aggregation**
  - Aggregates measurements from multiple vantage points
  - Trades delay for consistency

Thank you!

Confirmation removes false alarms – Emulab

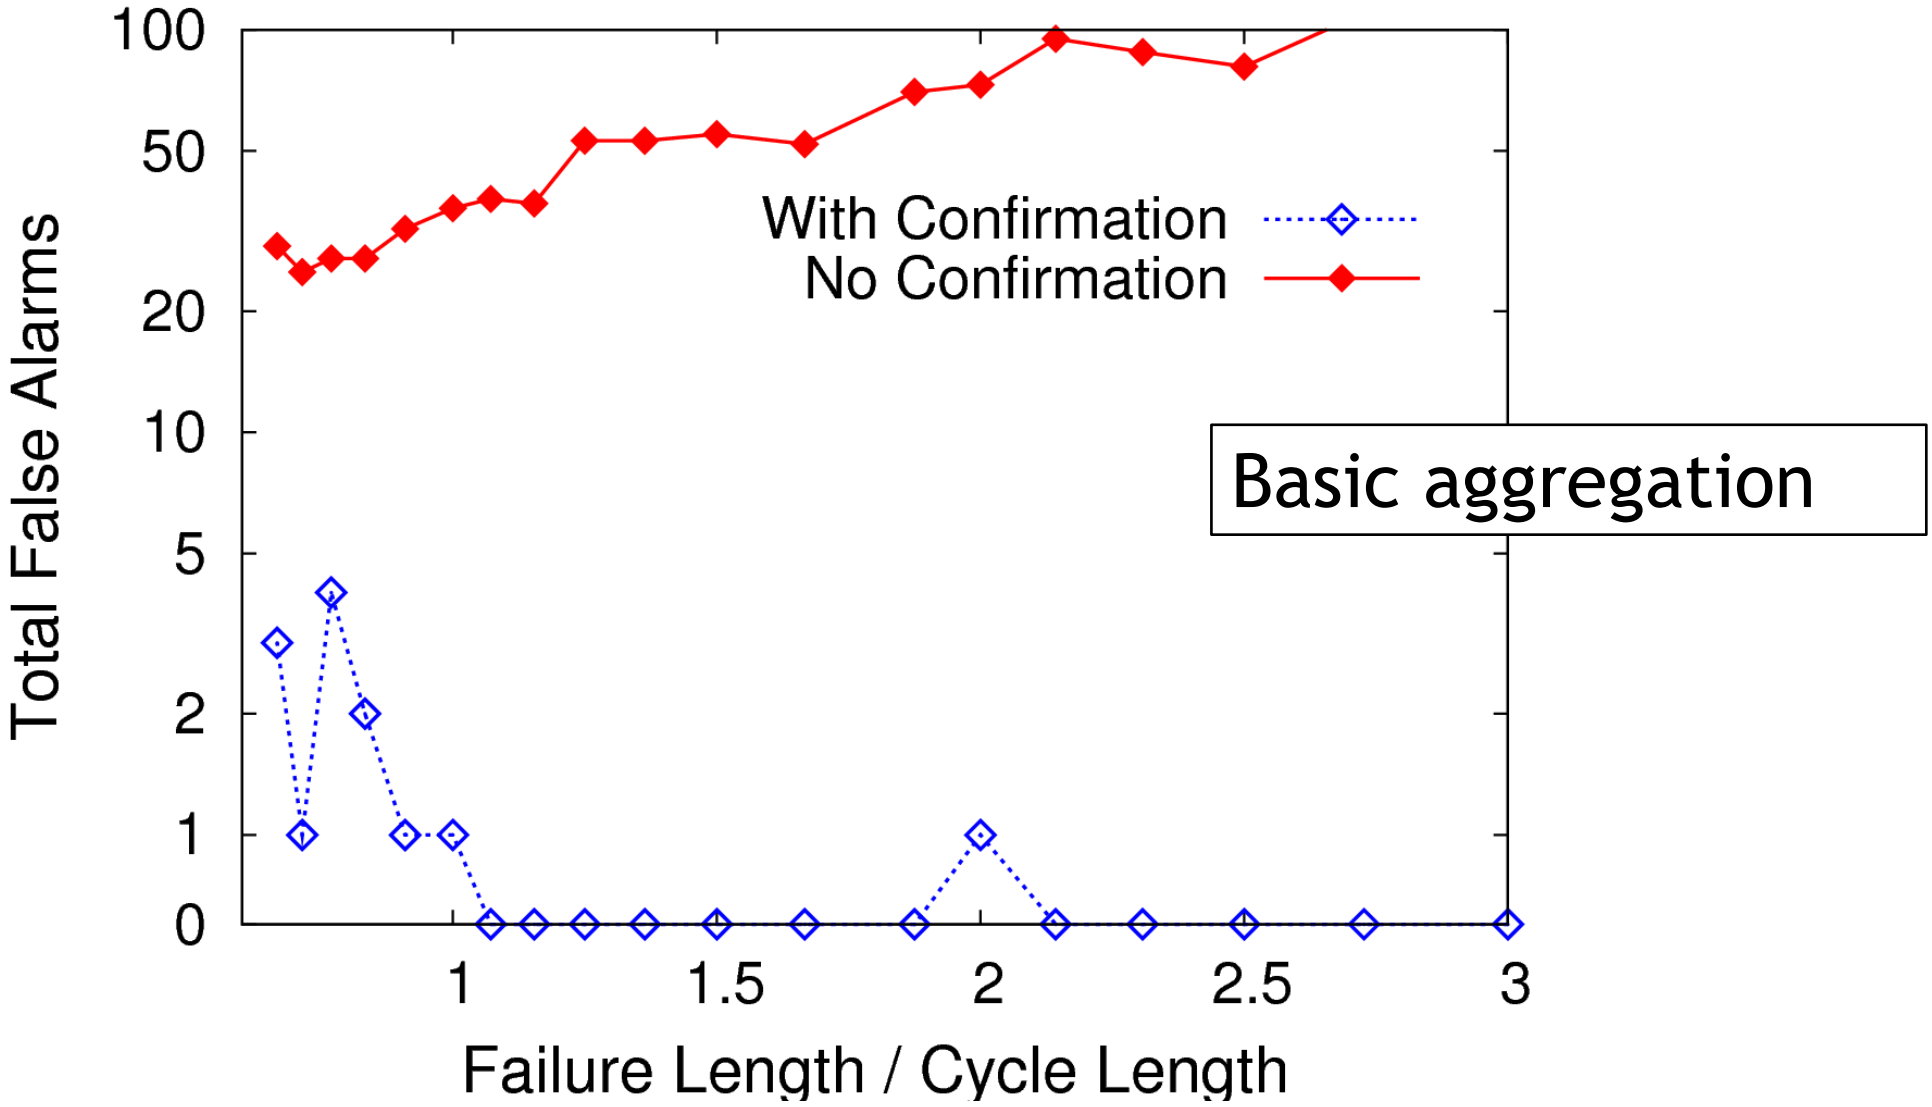

Trading delay for higher consistency - Emulab

- Abilene topology and synthetic failures
- No confirmation: many detection errors

		n				
		1	2	3	4	5
BASIC	Delay	2.51				
	Cons.	0.914				
MC	Delay		15.2	19.0	19.7	22.4
	Cons.		0.981	0.998	1	1
MC-PATH	Delay		6.92	10.0	13.0	16.0
	Cons.		0.964	0.977	0.982	0.985

Putting it all together - Emulab

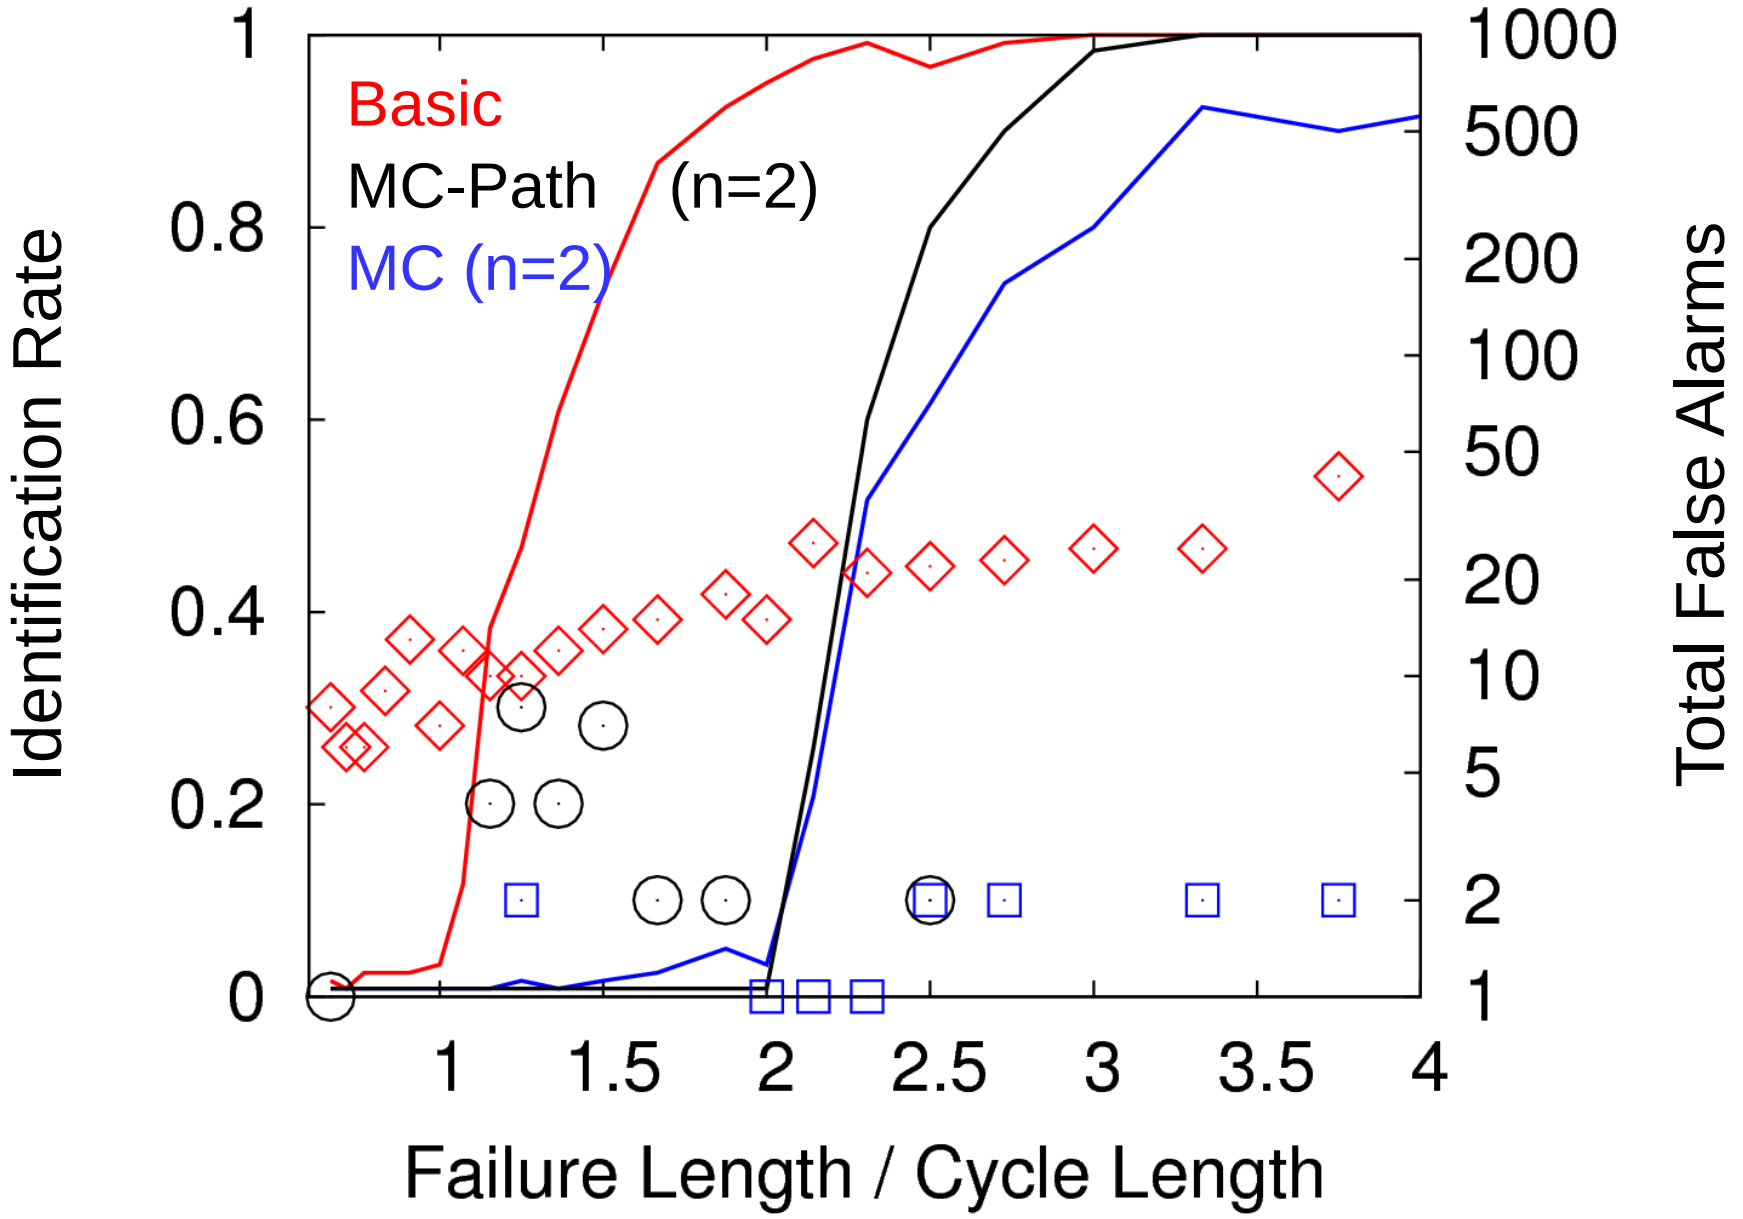

From loss model to probing scheme

- **Bursty losses modeled as a two-state Gilbert model**
  - Loss bursts durations are exponentially distributed
  - Prob(losing confirmation probe | previous probe lost)

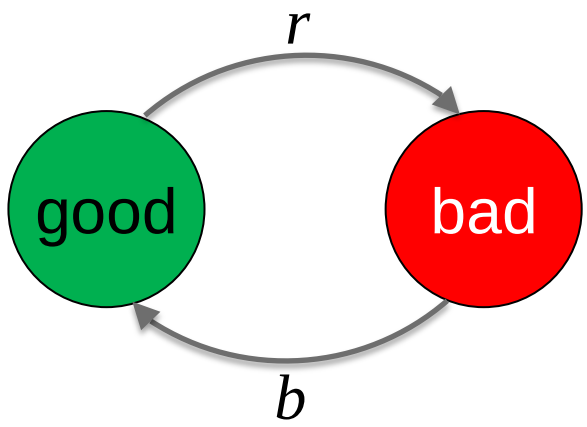

Short failures are impossible to identify consistently

- Failures shorter than one cycle

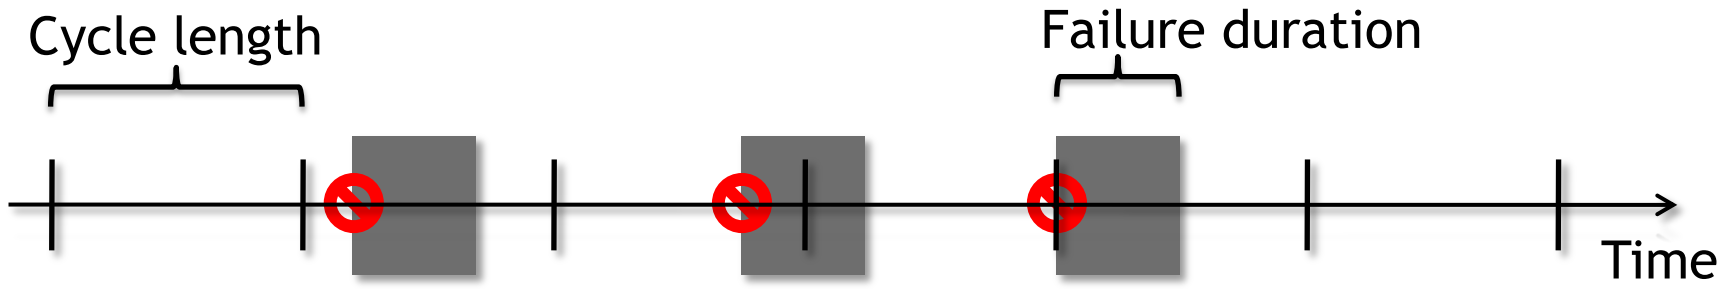

- Failures shorter than two cycles

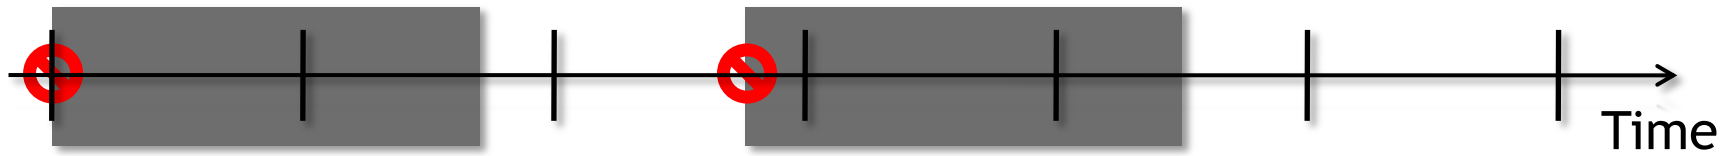

- Failures longer than two cycles

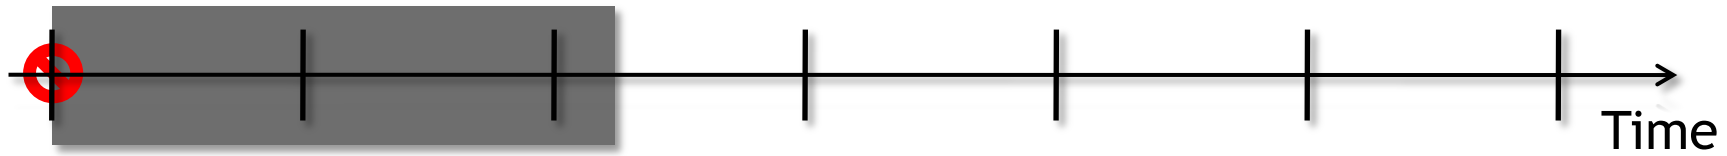

Detection errors can also create false alarms

- Detection errors reduce consistency

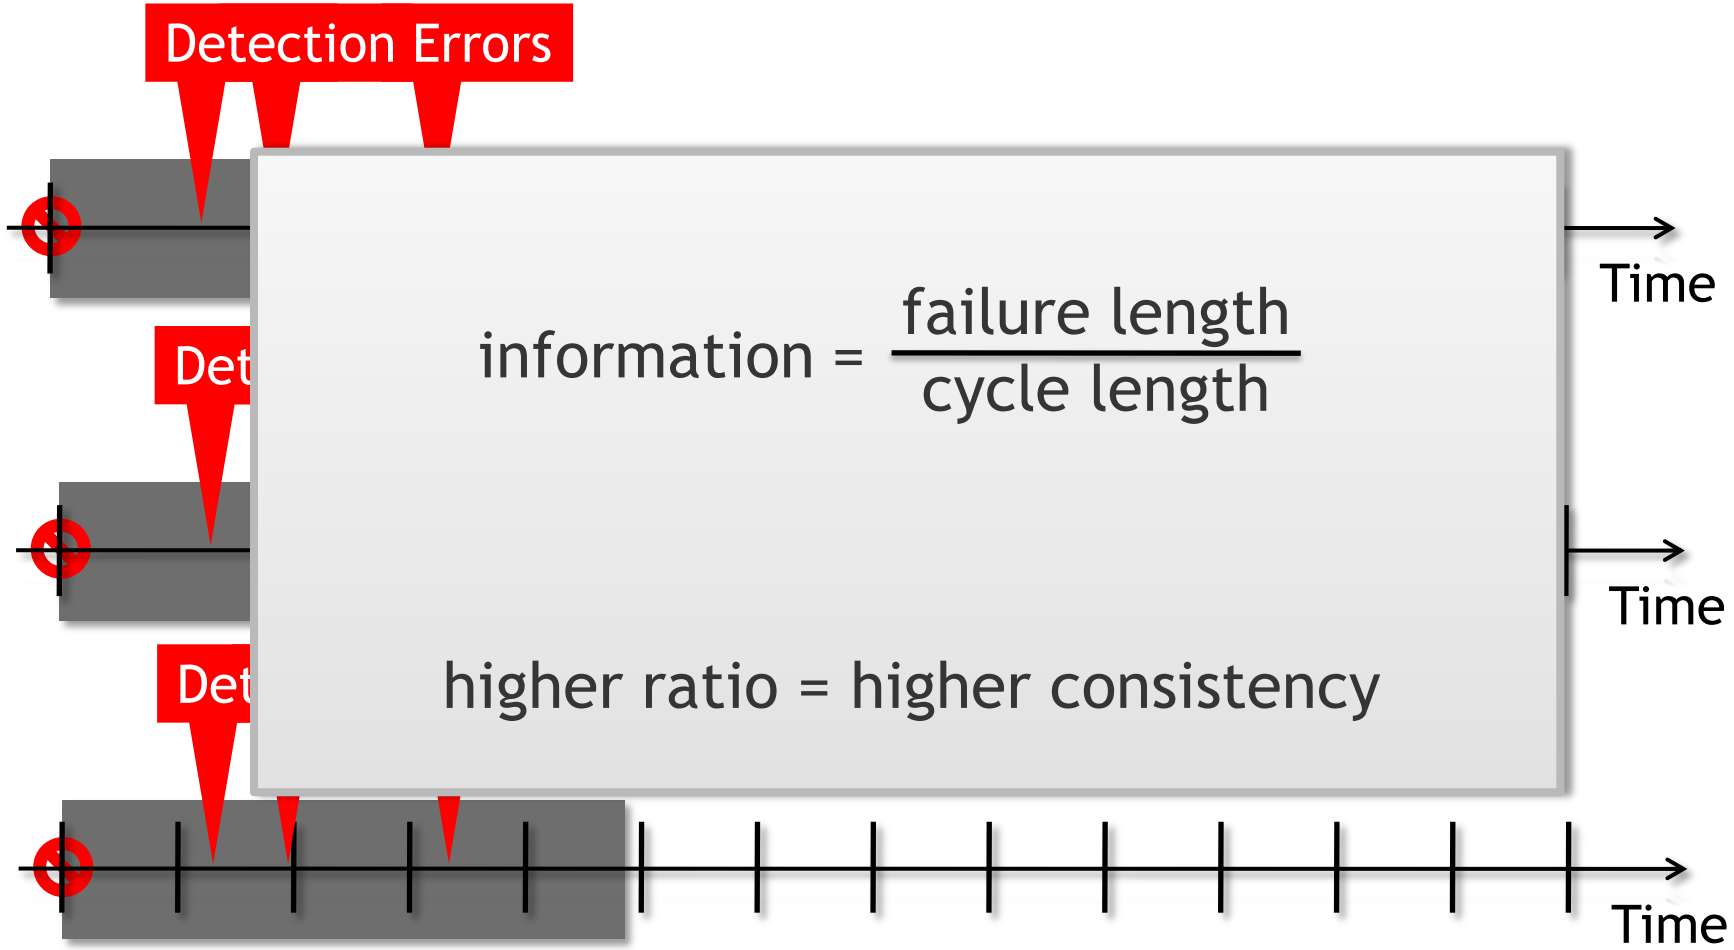

Future Work

- Tomography algorithms also need up-to-date topologies
- Develop confirmation and aggregation techniques for non-binary metrics like loss rate and bandwidth
- Deploy these techniques in conjunction with tomography algorithms to build a real-time system

Putting everything together - real deployments

	PlanetLab		Enterprise	
	No Conf.	Confirmation	No Conf.	Confirmation
BASIC	1500/day	1500/day	500/day	20/day
MC	No alarm	No alarm	1/day	1/day
MC-PATH	23/day	12/day	2/day	2/day

THOMSON TECHNICOLOR