Characterization of Failures in the Sprint IP backbone

Athina Markopoulou
Gianluca Iannaccone
Supratik Bhattacharyya
Chen-Nee Chuah
Christophe Diot

Why Study Backbone Failures?

- Backbone networks provide excellent traditional QoS
- E.g. SLAs in the Sprint's IP network
  - 0.3% packet loss
  - 55 msec delay in continental USA
  - 99.9% port availability

- Failures are poorly understood…
  - although they happen every day

Link Failures

- Definitions
  - IP link: adjacency between two IS-IS routers
  - Link Failure: loss of this adjacency
- Possible reasons
  - Fiber cuts, optical equipment failures, router problem, human error or mis-configuration, maintenance …
- Impact on the IP layer
  - Topology changes & routing re-configuration

Dealing with failures

- Potential impact on availability
  - Forwarding disrupted during route re-convergence
  - Overload/Congestion on backup paths
- Network design becomes hard
  - Protection mechanisms
  - Topology design
  - Capacity provisioning
  - Timers tuning
- **A failure model is needed !**

Outline

- Motivation
- Contributions
  - Measurement collection
  - Failures classification
  - Modeling of each class
- Conclusion

Measurements of Link Failures at ISIS level

- ISIS listeners collect flooded LSPs

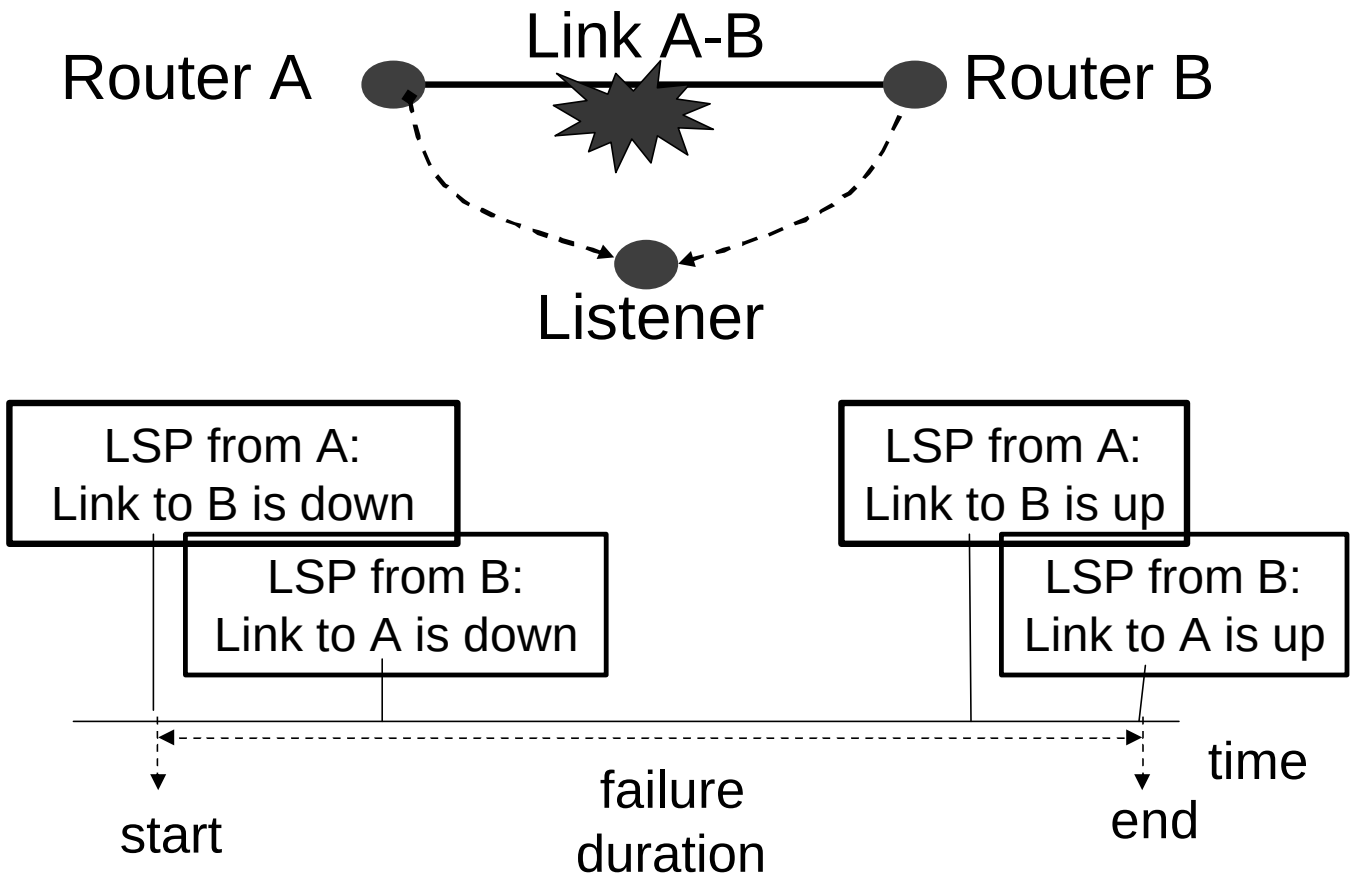

- Record: (**link A-B**, router A, router B, **start time, end time**)

Data Set: US Failures, Apr. - Nov. 2002

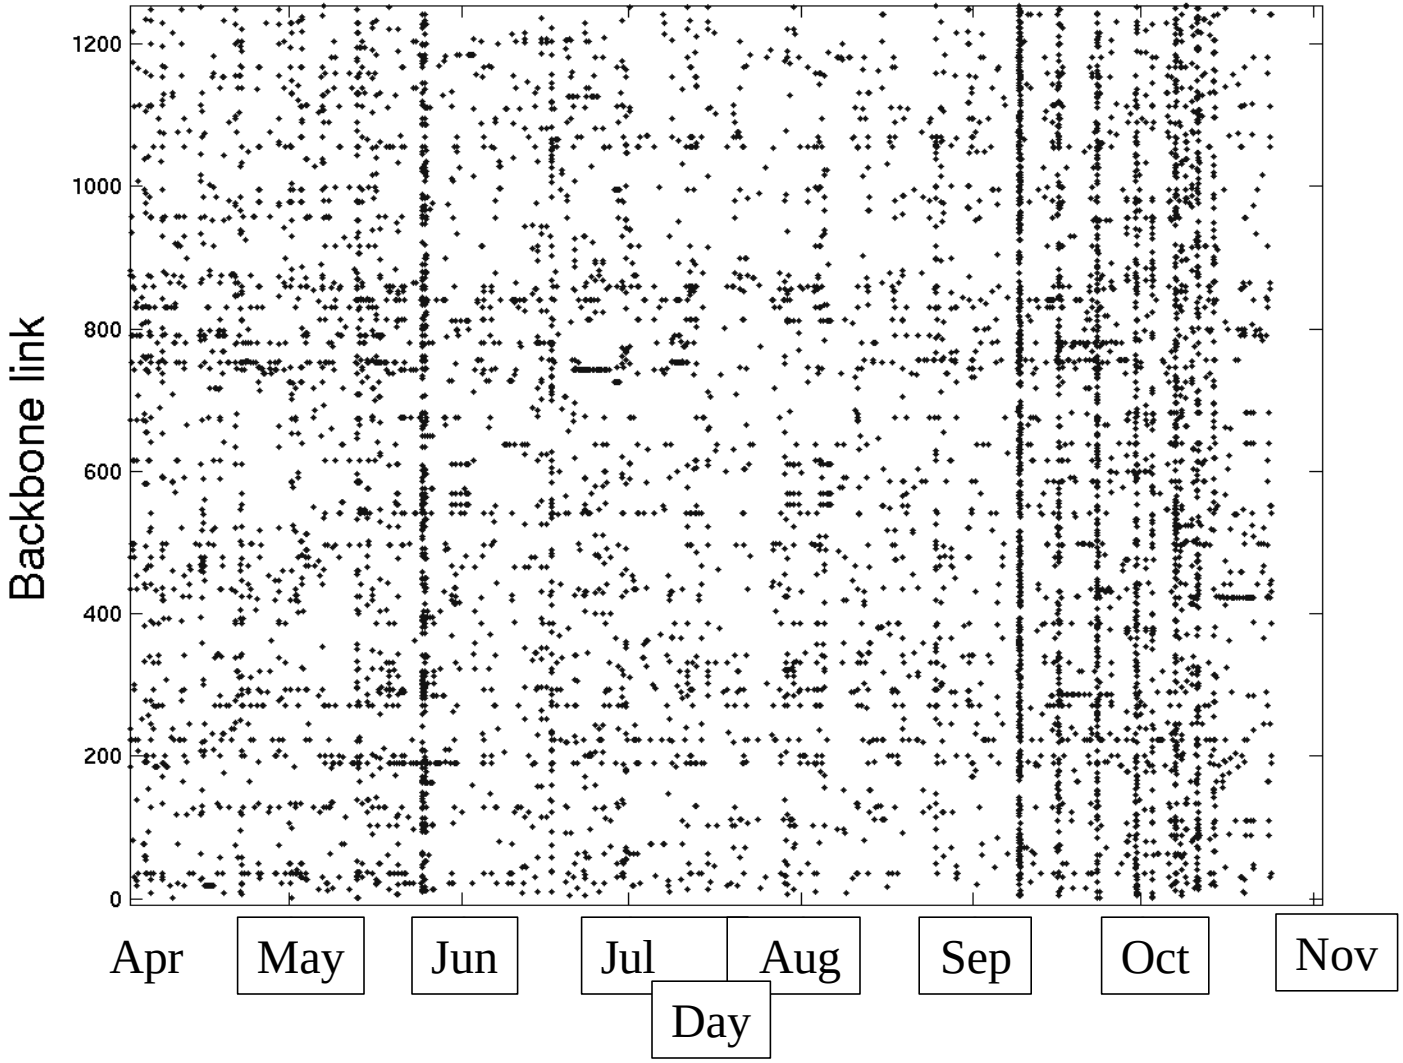

Classification Methodology

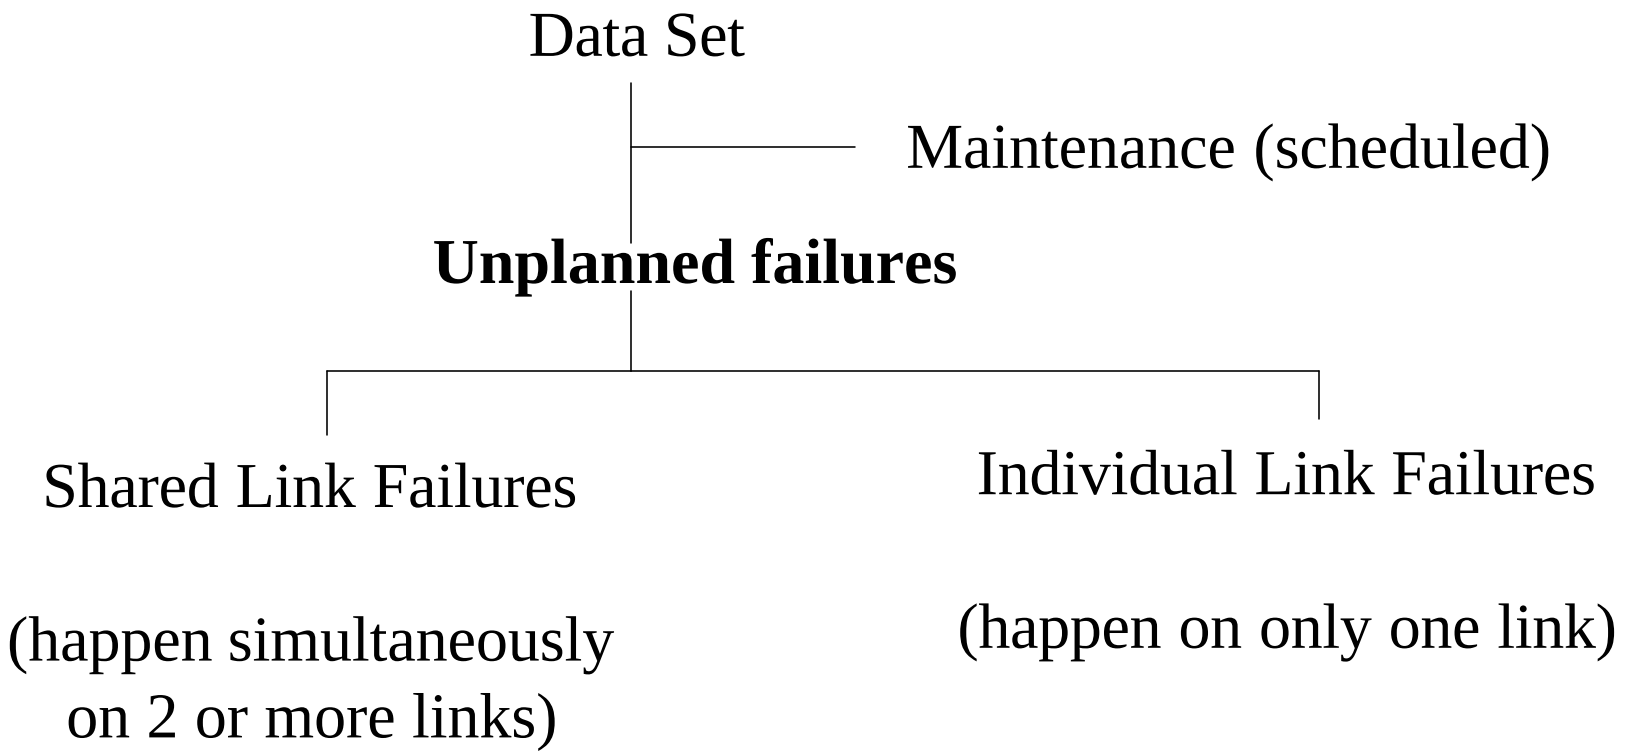

Maintenance

Weekly schedule (Mondays 5am-2pm UTC): 20% of failures

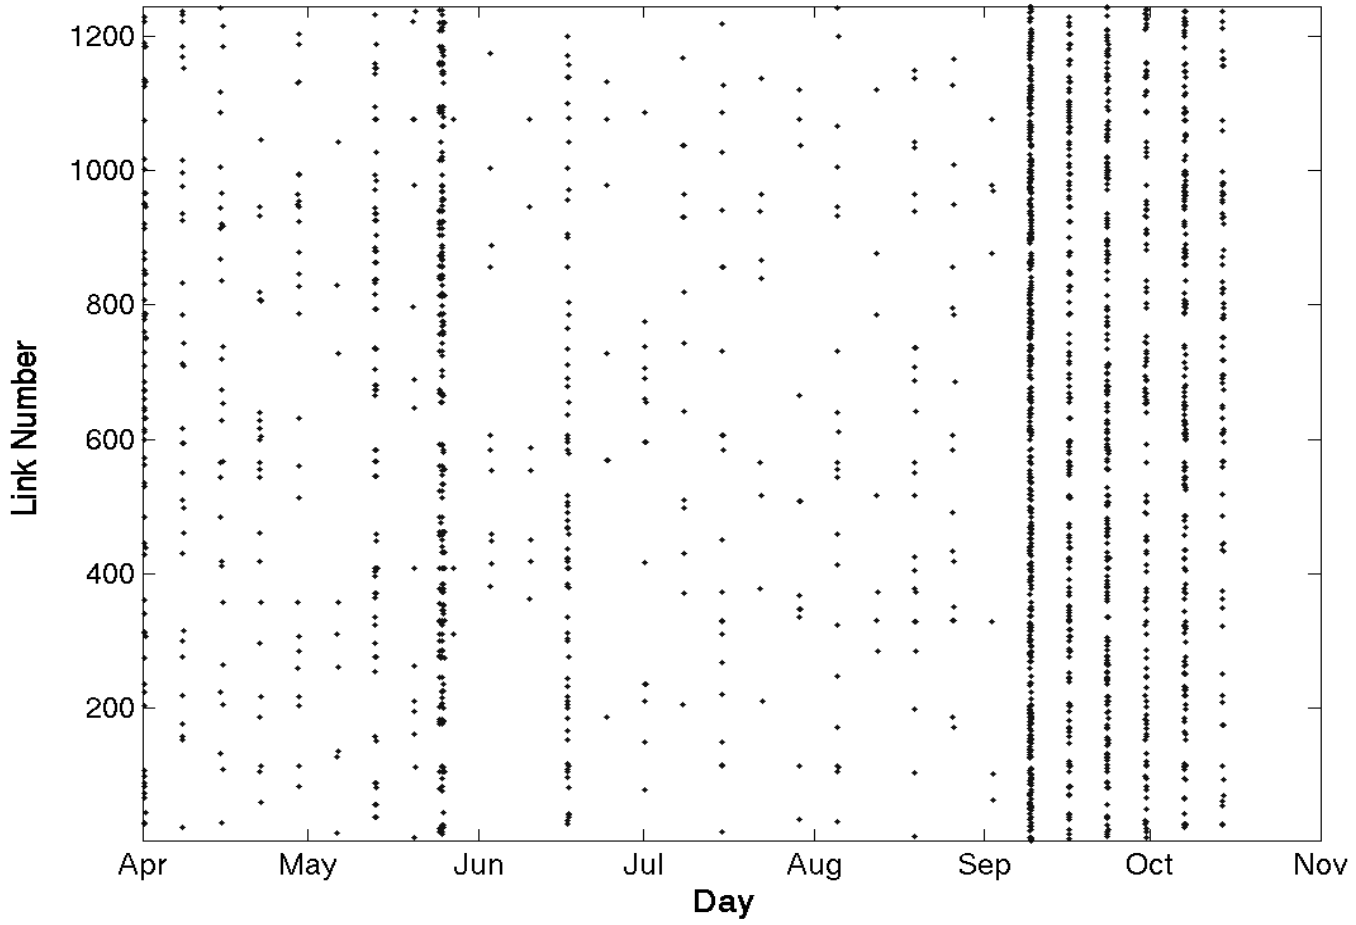

Data Set-revisited

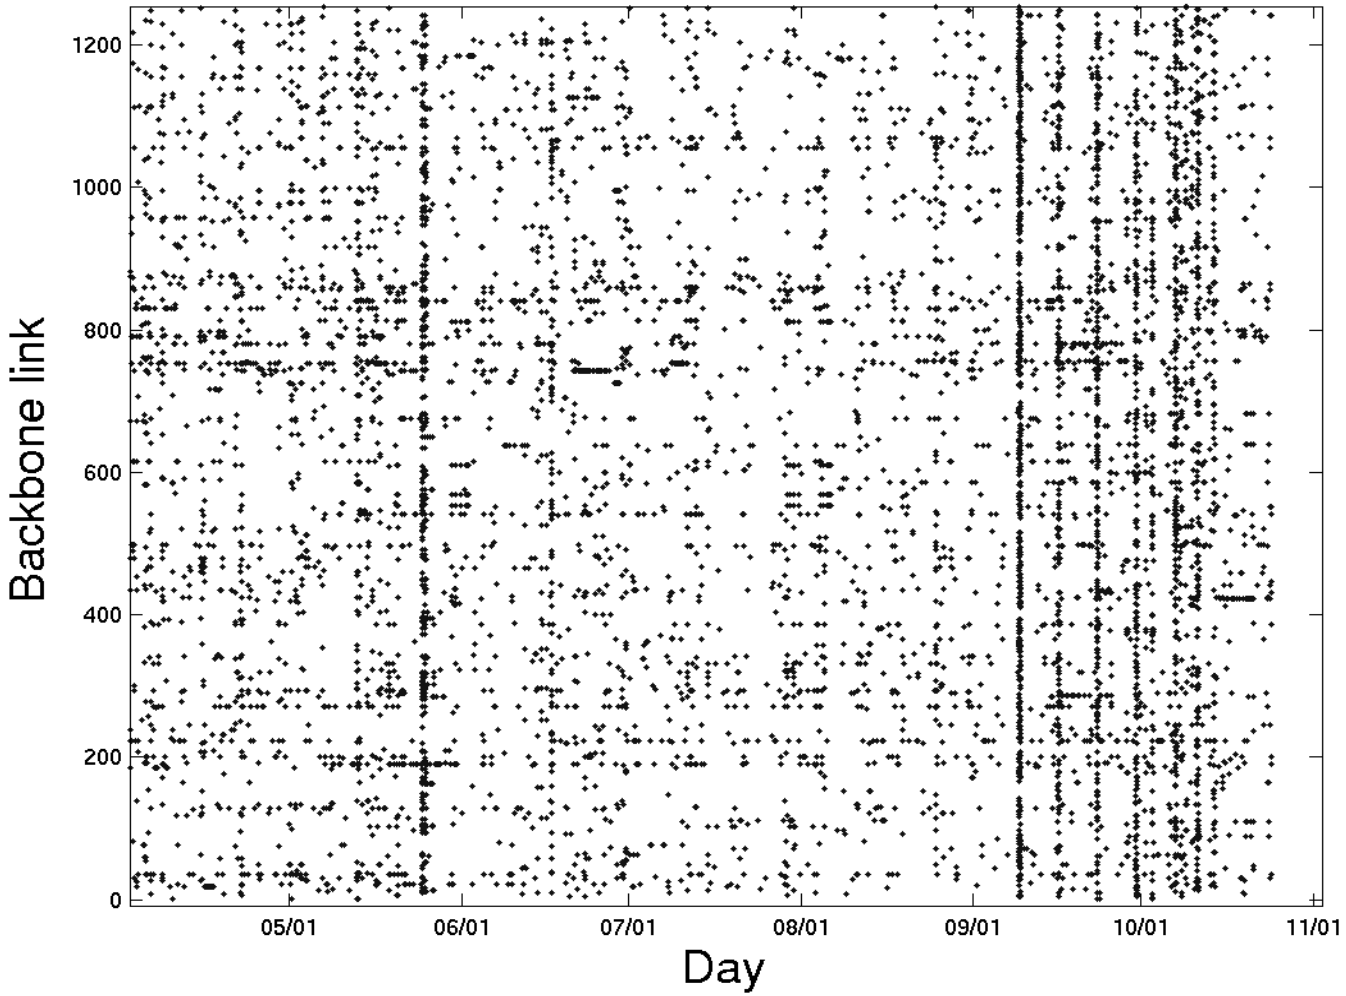

Excluding Maintenance: Unplanned failures

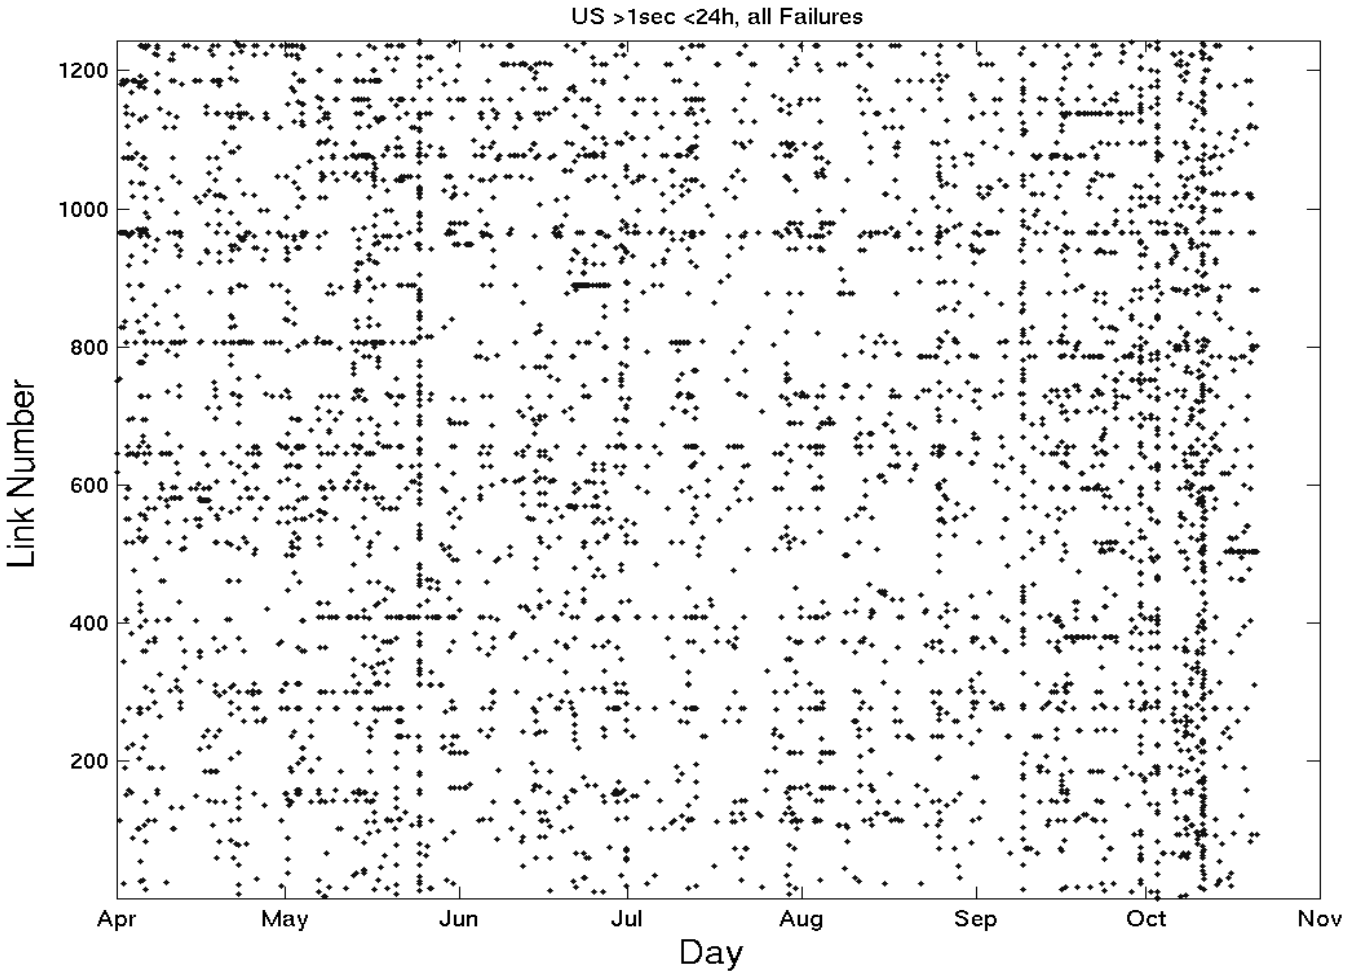

Shared failures (I): simultaneous

- 2+ links, go down/up at the exact same time
  - 16.5% of all unplanned failures

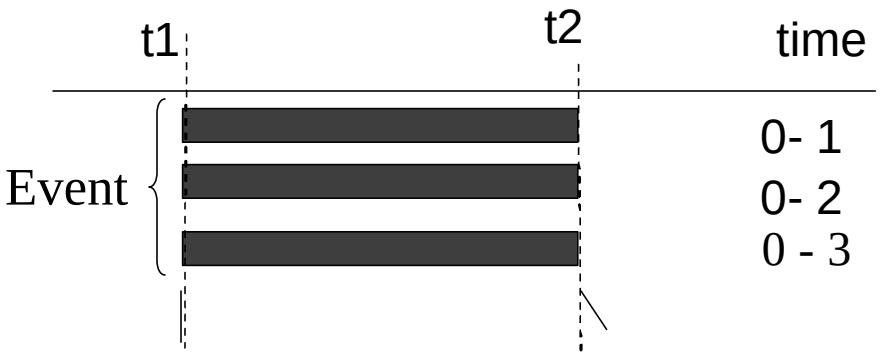

- For every such event
  - all links connected to the same router
- Indicate router-related cause

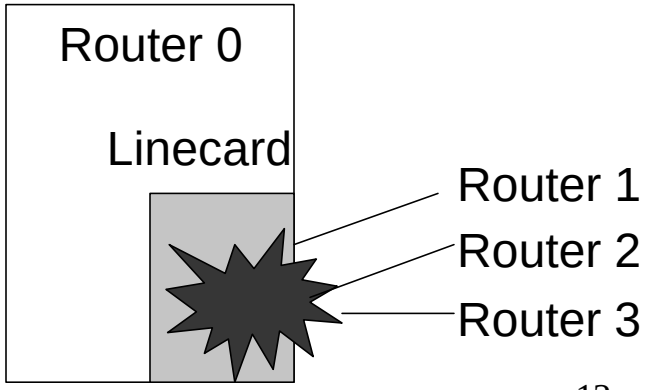

Shared failures (II): overlapping

- Some links fail almost simultaneously
  - 14.3% of all failures are grouped in overlapping events

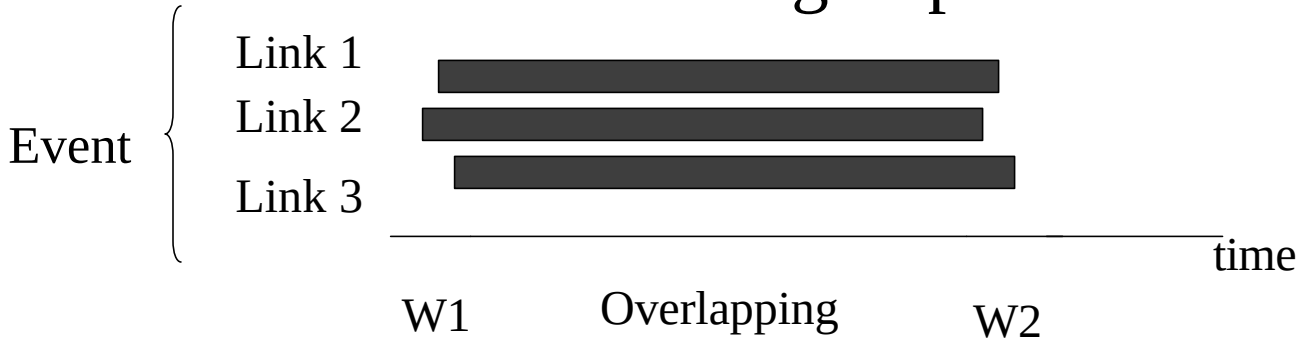

- In 80% of overlapping events
  - links have no common router
  - but all links share multiple optical components
  - Good indication for optical-related failure

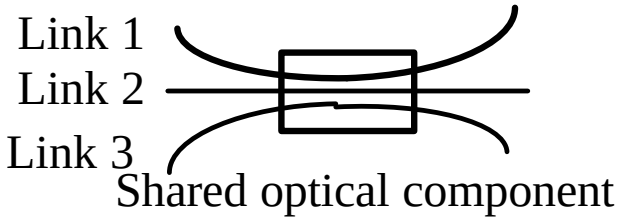

Classification

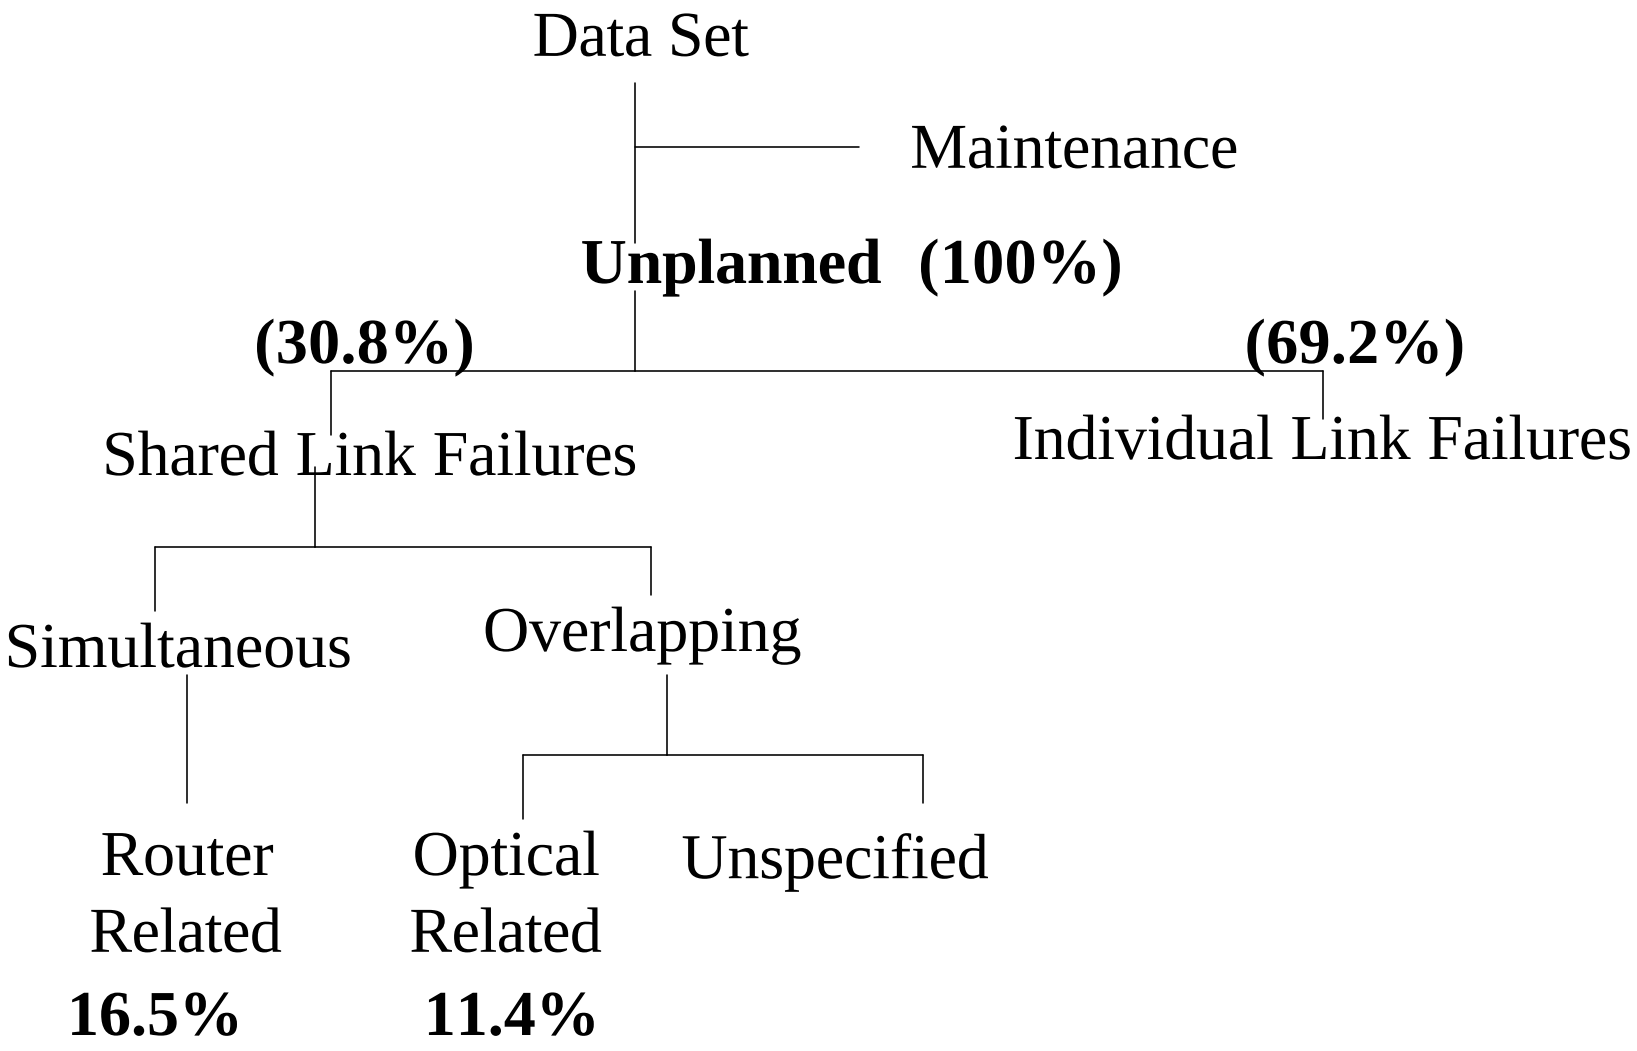

Individual Link Failures

After excluding ~~maintenance and shared (router,~~ optical)

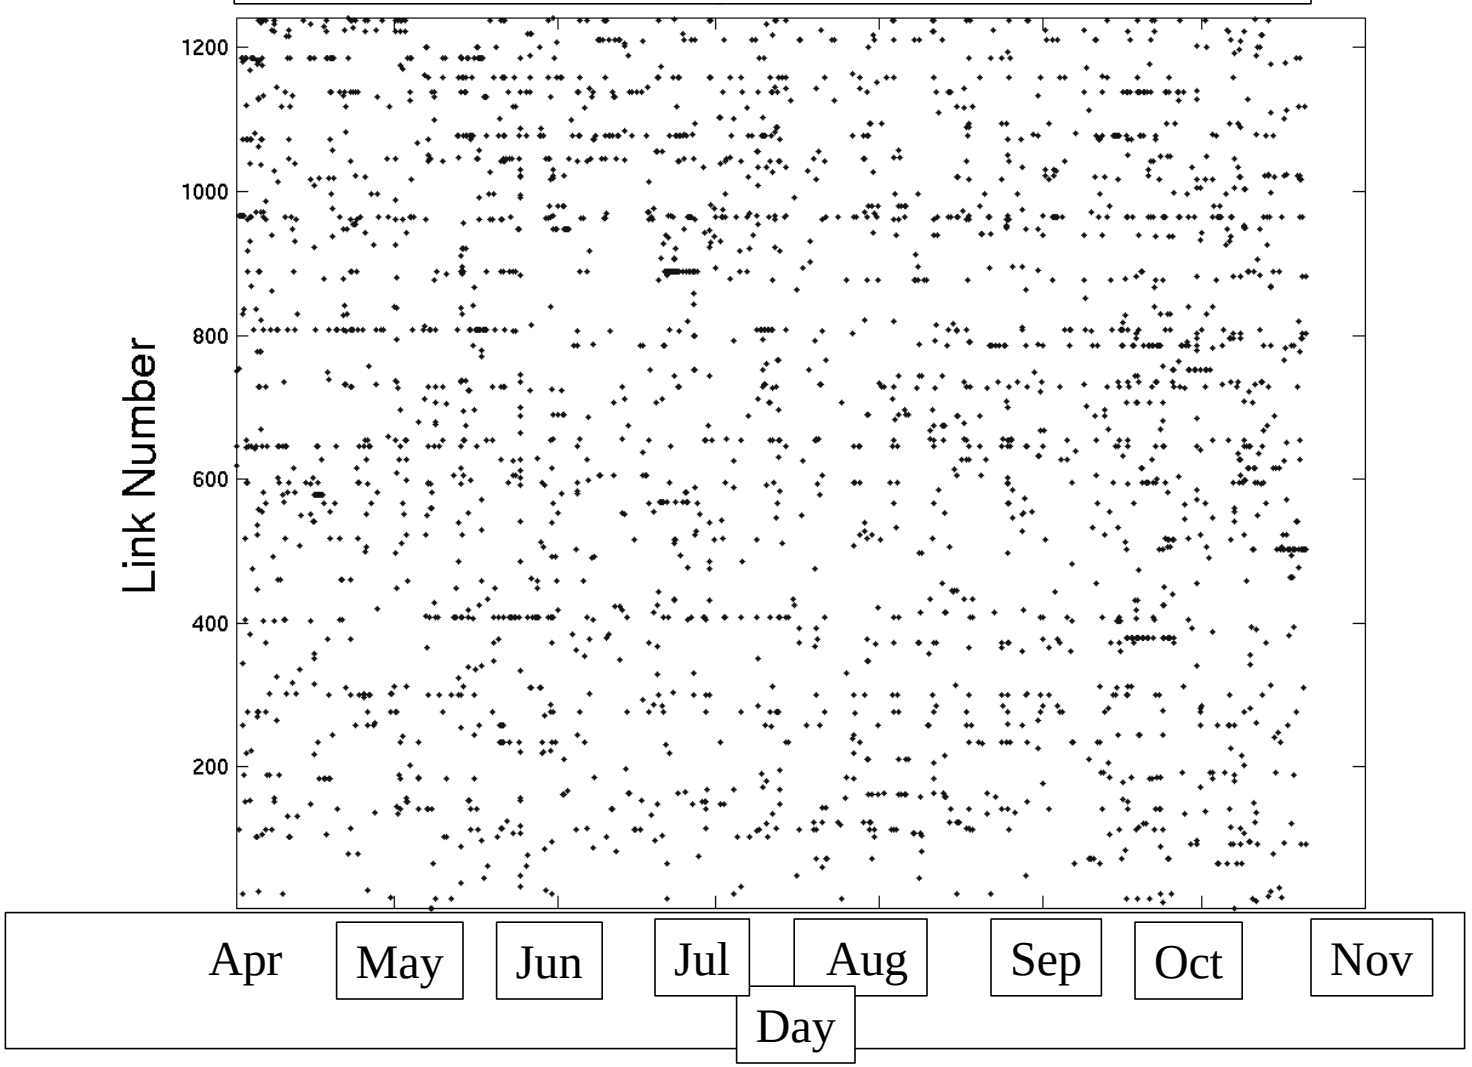

High vs. Low Failure Links

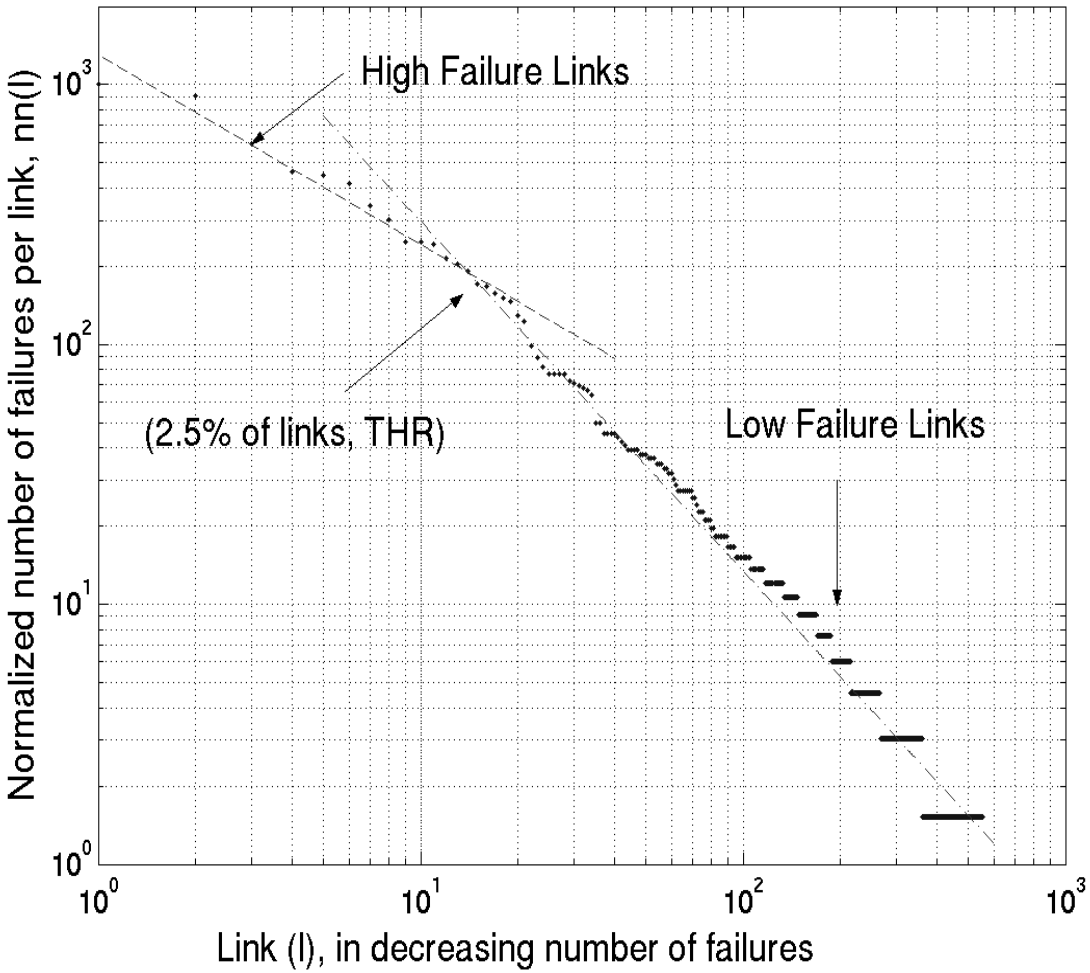

- Normalized number of failures per link
- High degree of heterogeneity
- Roughly two power-laws
- A few (2.5%) links account for half of individual failures

High Failure Links

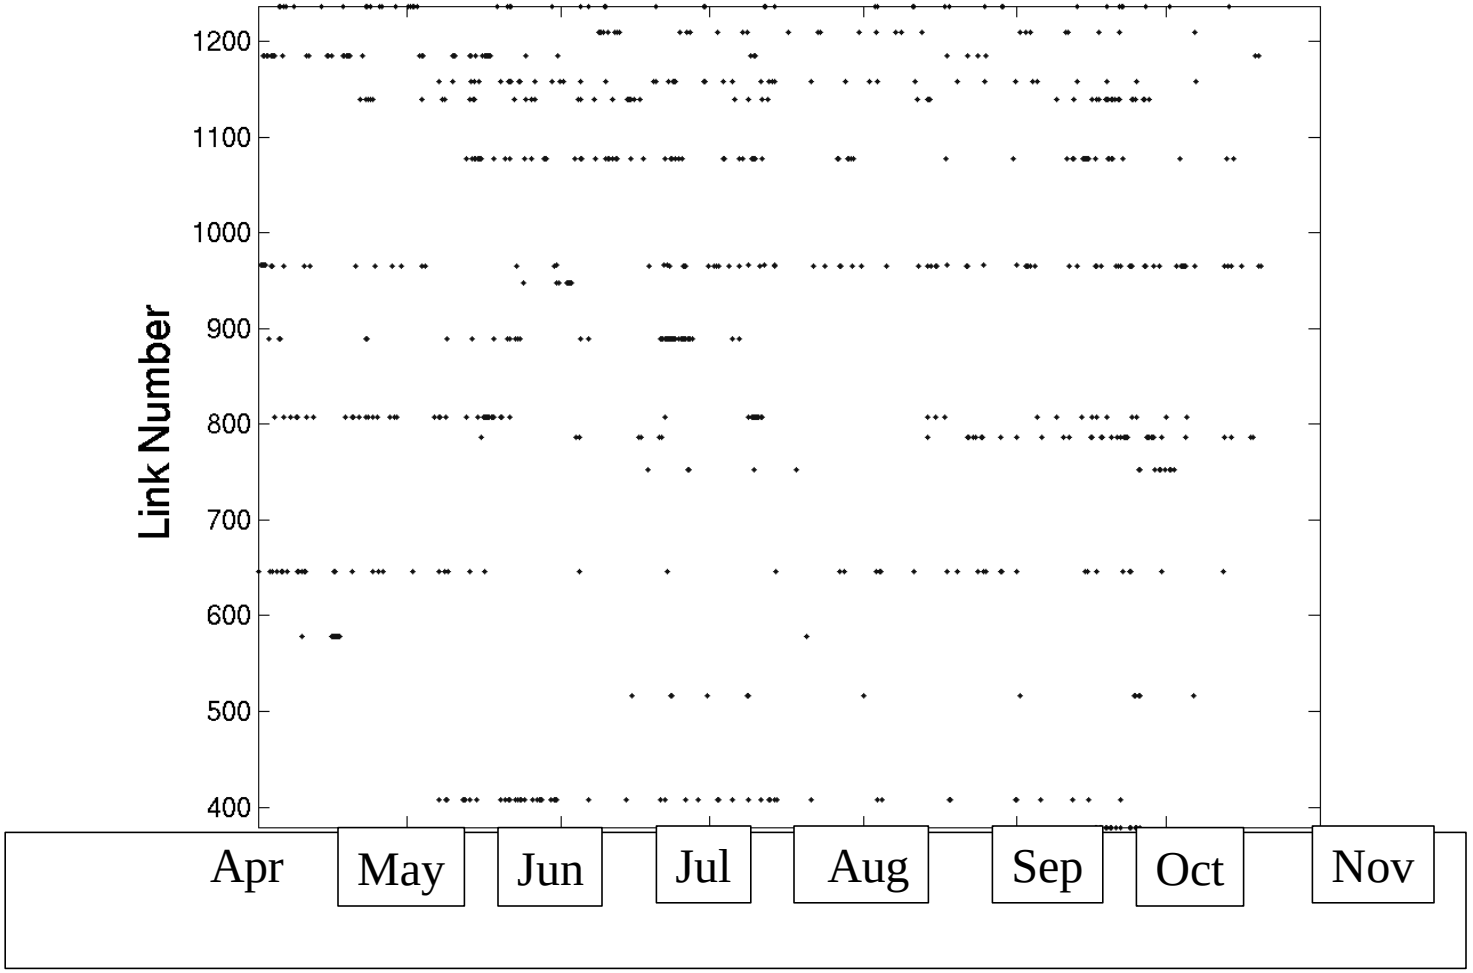

- Components may be old or undergoing upgrades

Low Failure Links

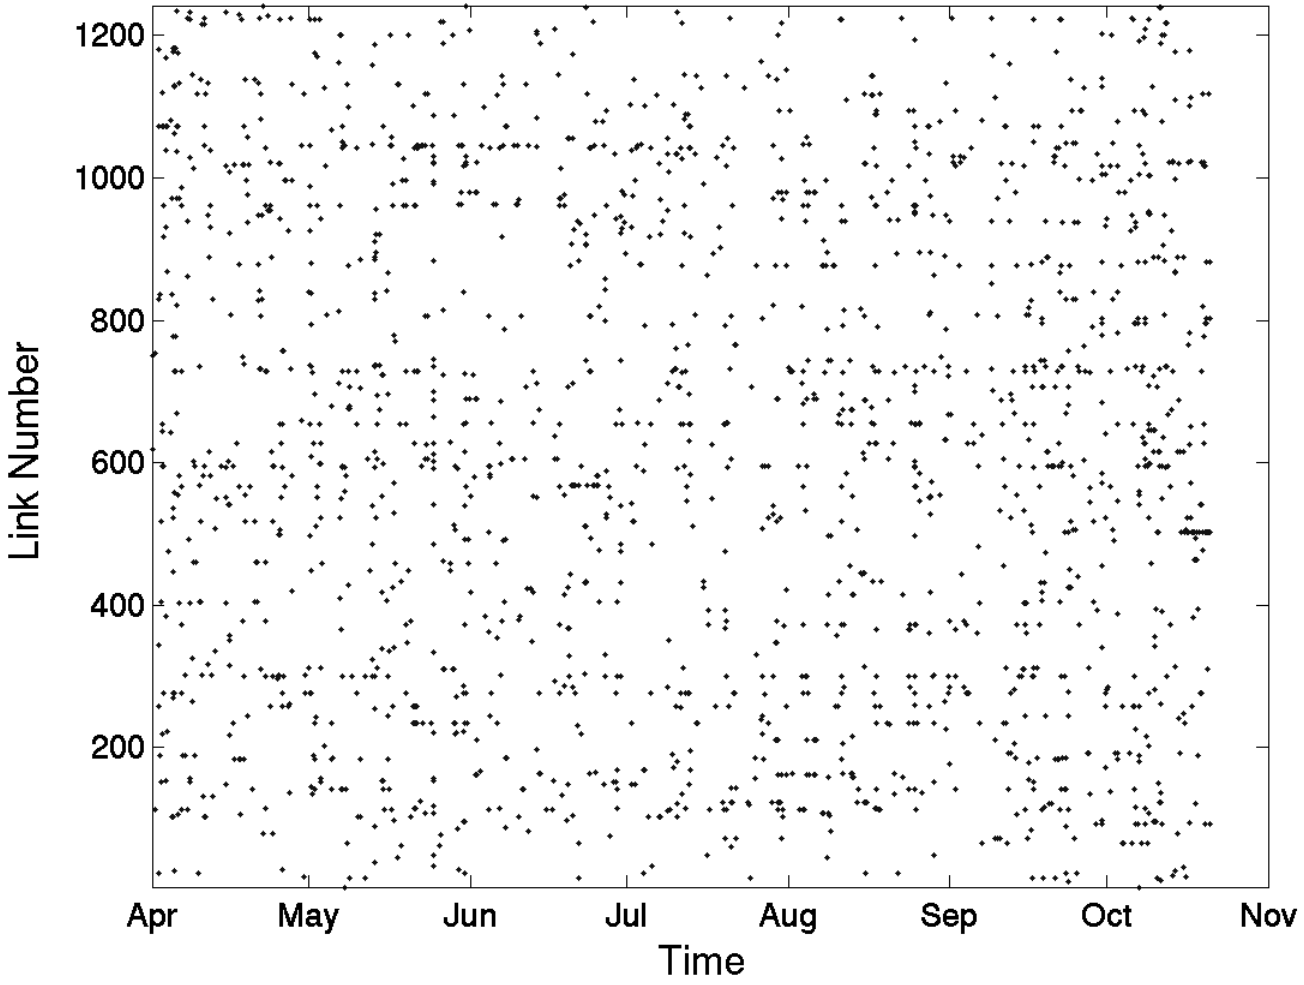

Classification Summary

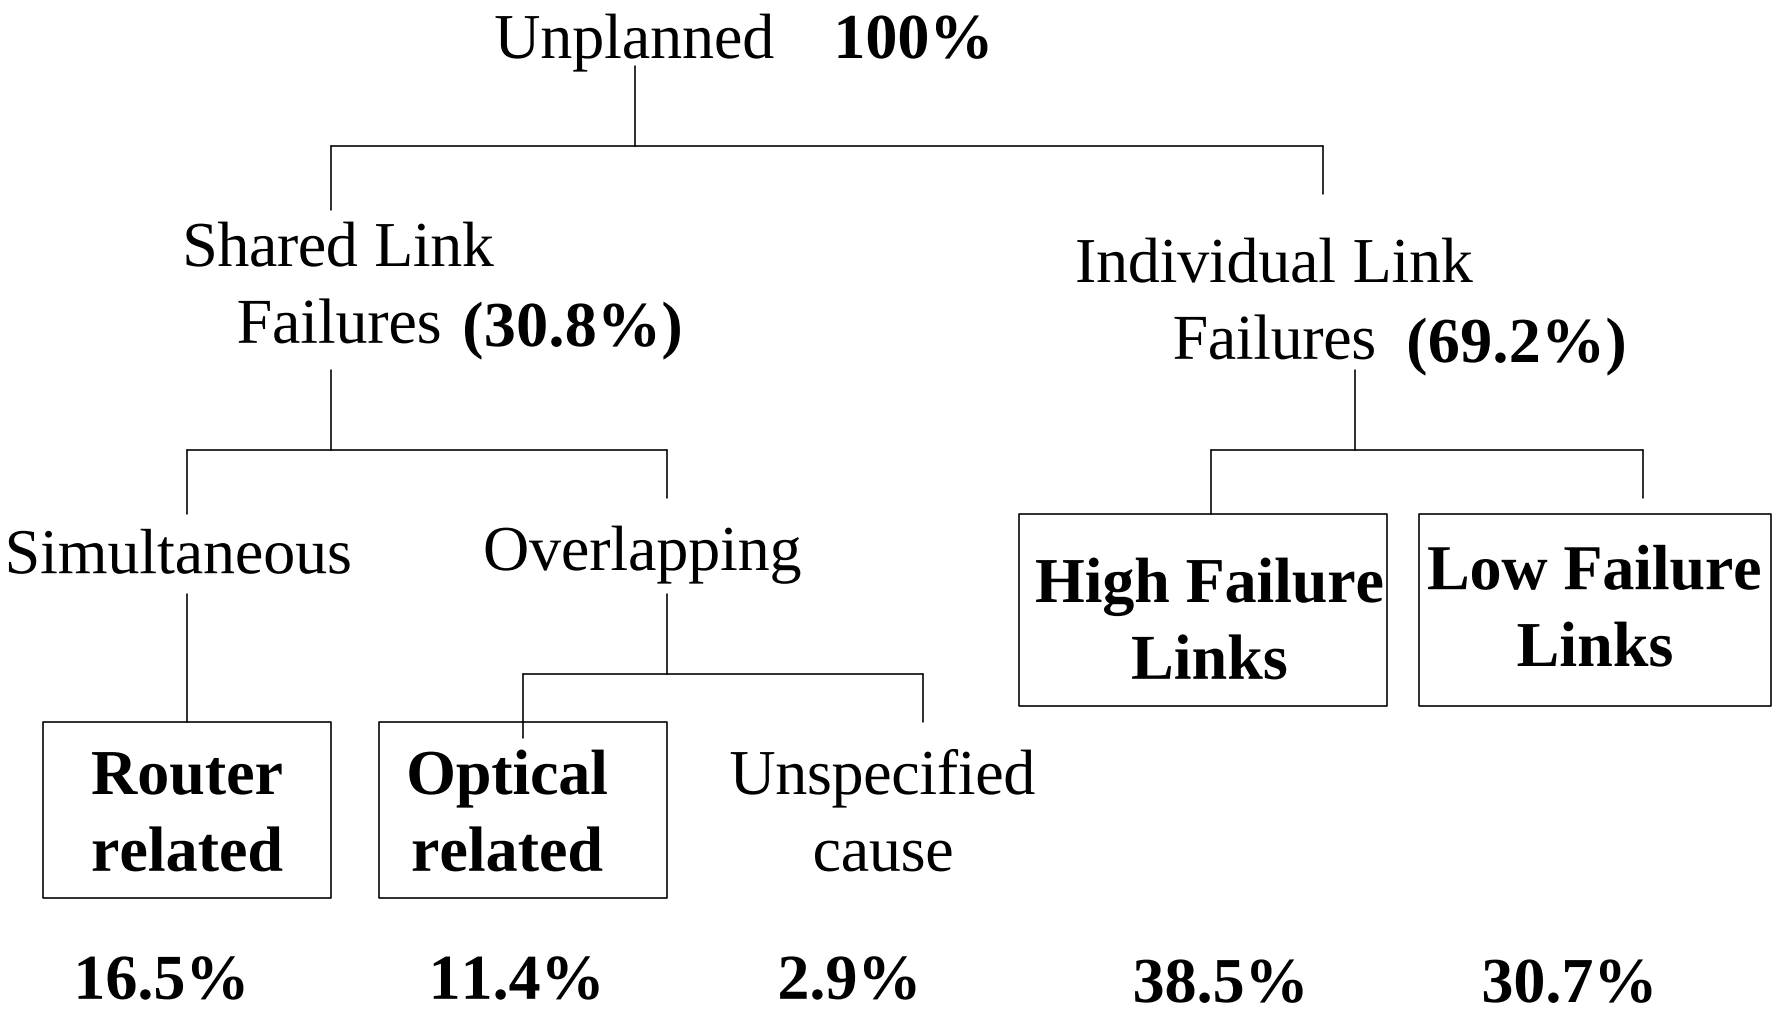

Characterize each class

E.g. low failure links - revisited

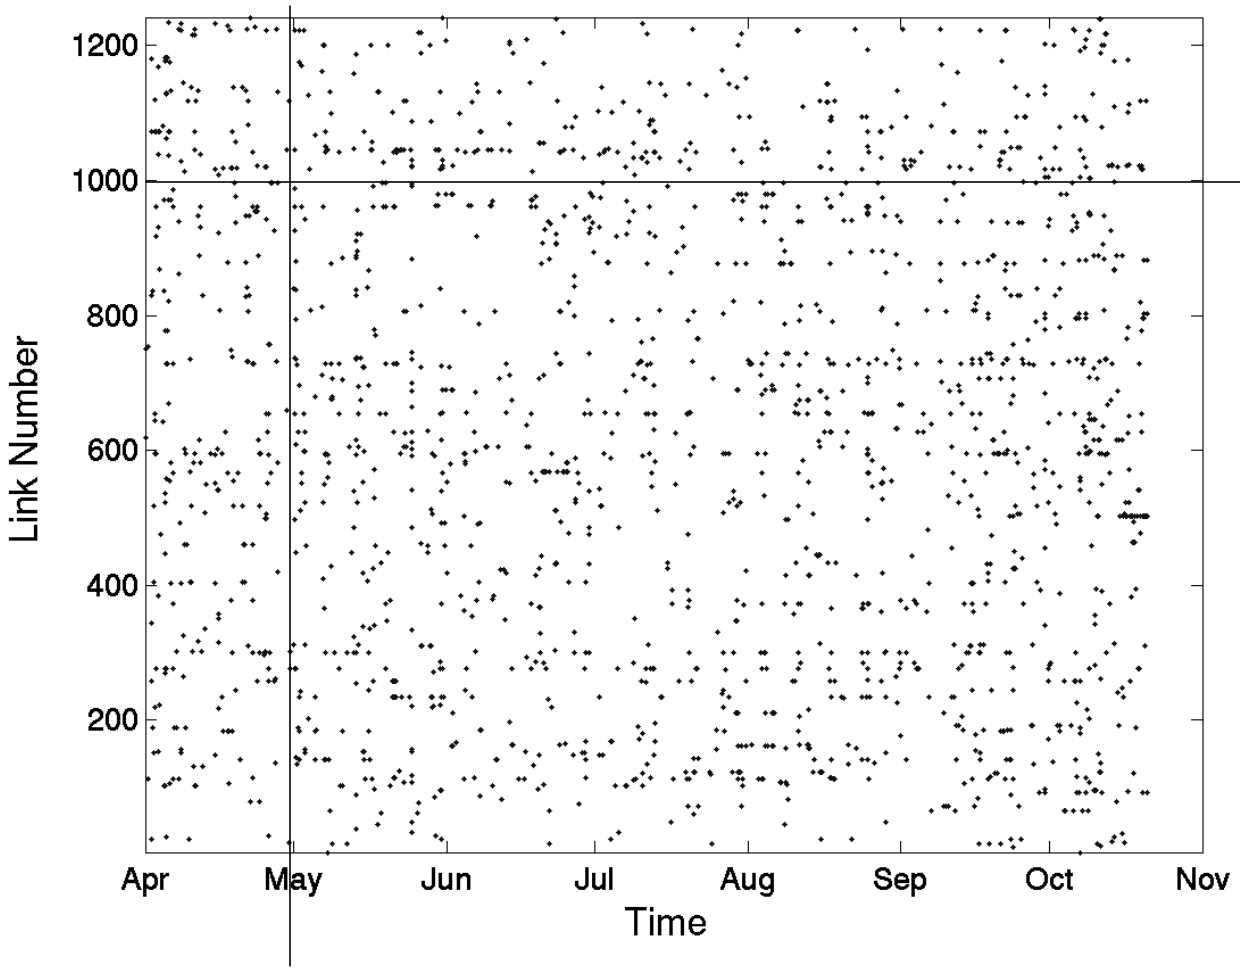

1. How often?

2. How are they spread across links?

3. How long do they last?

1. How often?

Time between two successive failures, on any link

- Weibull: $F(x) = 1 - \exp(-(x/scale)^{shape})$
- Low autocorrelation

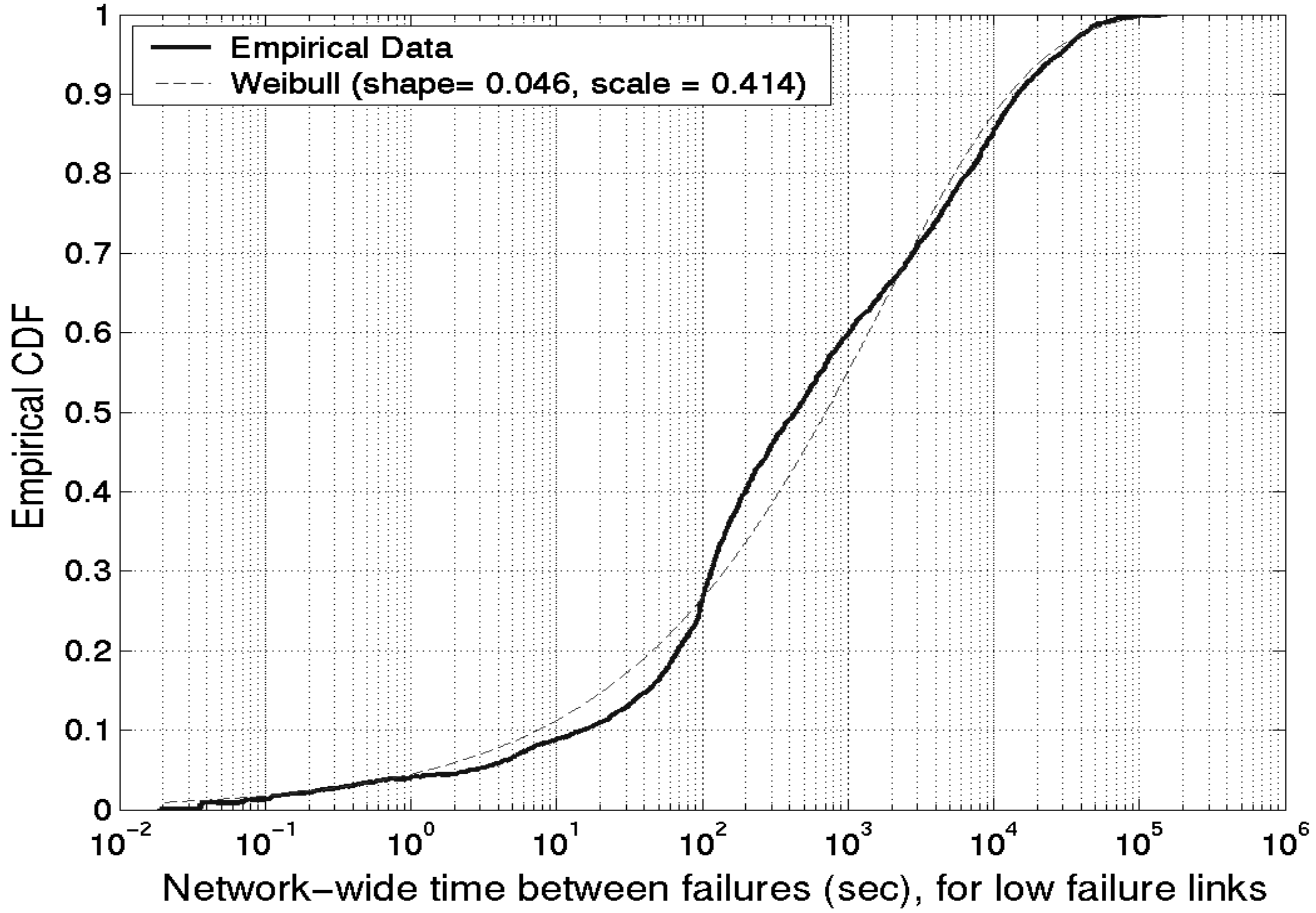

2. How are they spread across links?

- Order links in decreasing number of failures
- Link with rank l has n_l failures, s.t. $n_l \propto l^{-1.35}$

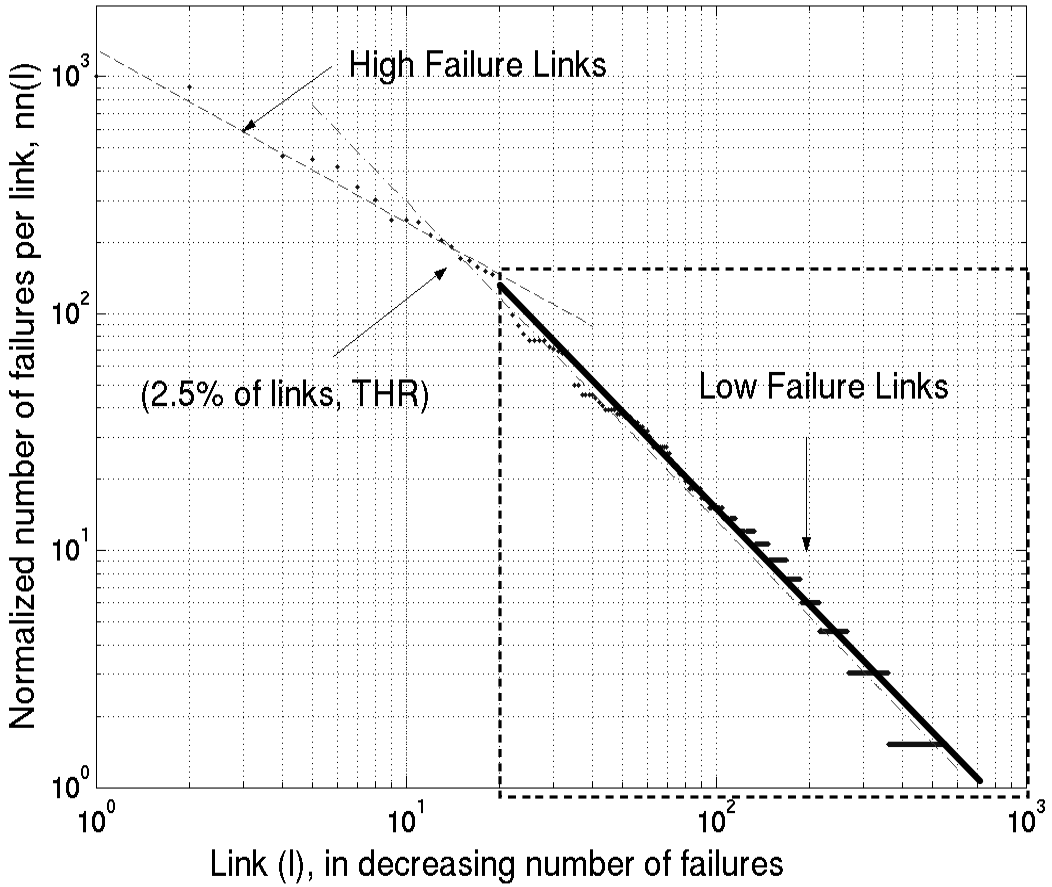

If

- long time T($p_l \sim \frac{n_l}{T}$)
- independent links

Then

- given that there is a failure, it happens on link l, with probability:

$$p_l^{cond} = \frac{n_l}{n_1+n_2+...+n_L}$$

Characterizing the Properties of each Class

	1.Time between failures	2. Num. of failures per link (or num. of events per router)	3. Duration	4. Number of links in the same event
Low failure links	Network-wide: Weibull, low autocorr.	Power-law	Empirical	N/A
High failure links	Per link: Empirical	Power-law	Empirical	N/A
Router related	Network-wide: Weibull, low autocorr.	Power-law	Empirical	Empirical
Optical related	Network-wide: Weibull, low autocorr.	-	Empirical	Empirical

Failure Durations – all groups

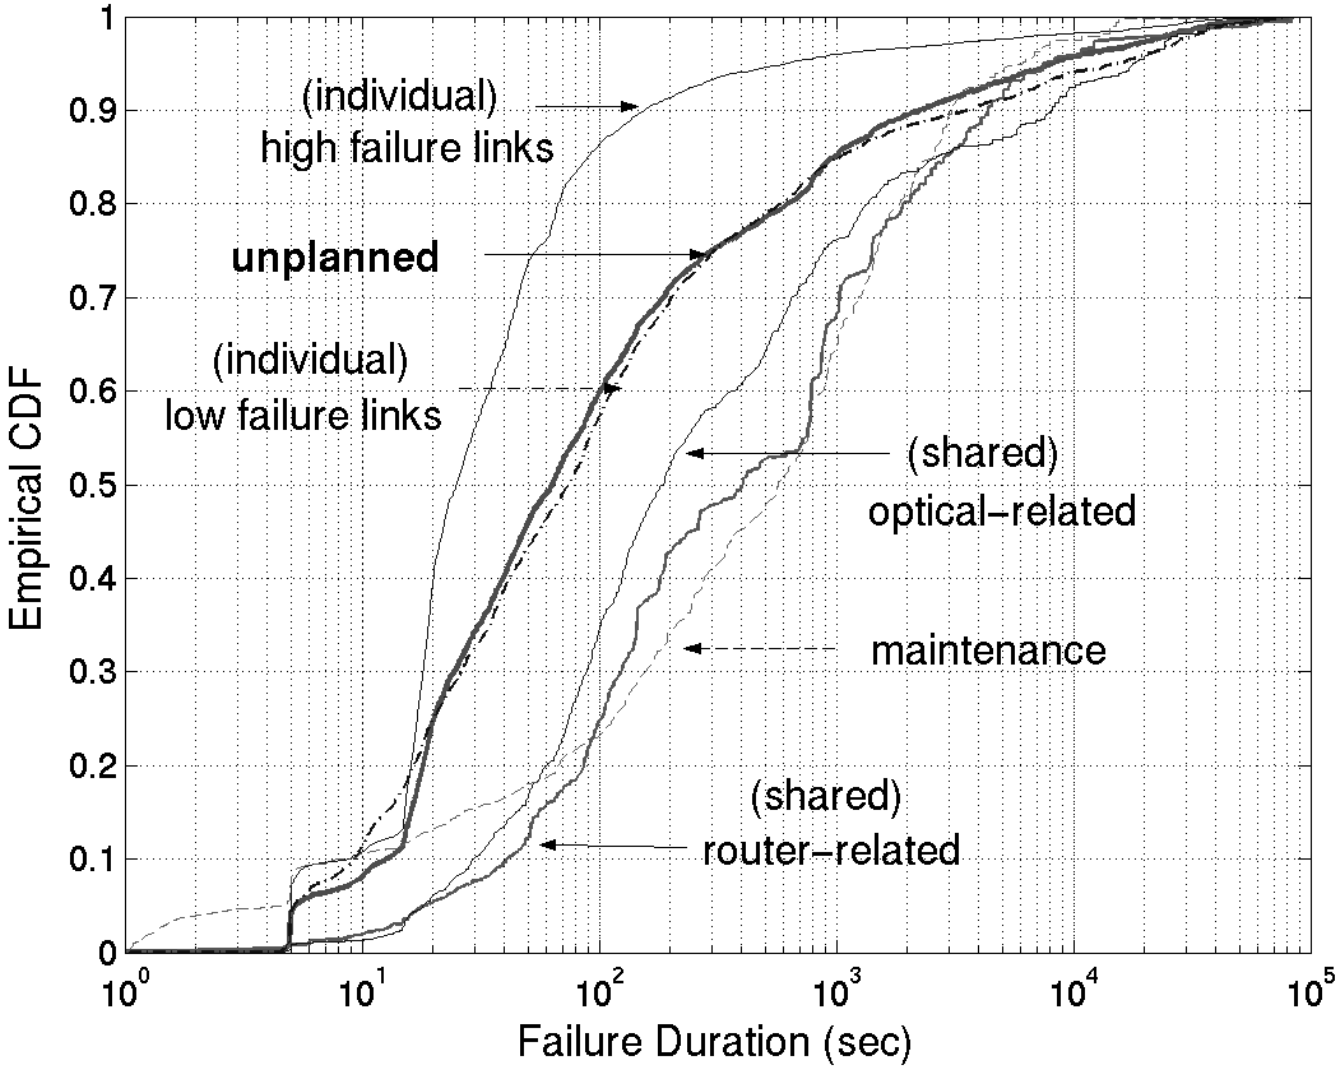

Conclusion

- Summary
  - Measurements
  - Classification
  - Characterization
- Implications
  - Contribution: A Failure Model
  - Problem Area: Network Reliability

Thank you!

amarko@stanford.edu