


Understanding Internet Path Failures: Location, Characterization, Correlation

Nick Feamster, David Andersen, Hari Balakrishnan
M.I.T. Laboratory for Computer Science
{feamster,dga,hari}@lcs.mit.edu

Big Picture



What? Where? How Long? Warnings?

- Explain path failures on the Internet
 - ▶ What causes failures, and where are they happening, anyway?
 - ◆ engineering preventative measures
 - ▶ What types of links experience failures, and why?
 - ◆ understanding of why RON can route around failures
 - ◆ more informed choices about connectivity

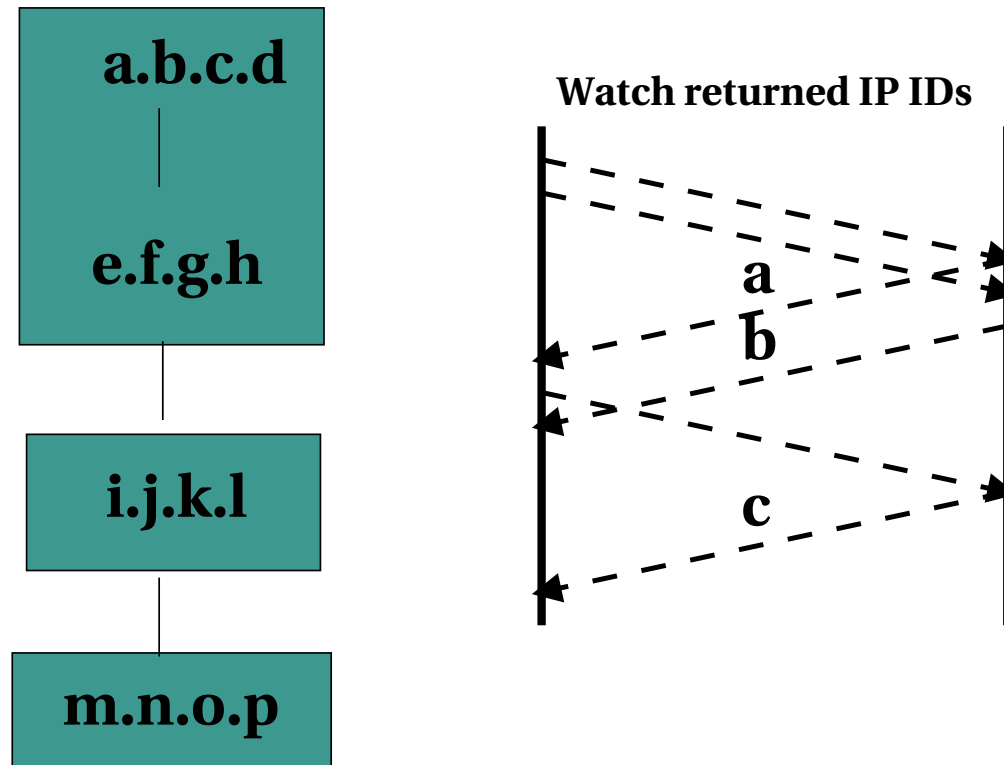
Questions

- **Locate:** Given a path failure, where did it happen?
 - ▶ With respect to end hosts?
 - ▶ What types of links? (intra-AS, etc.)
 - ▶ How can we observe this from the edge?
 - ▶ *(Traceroute is not good enough!)*
- **Characterize:** Do failures elicit specific qualities?
 - ▶ Does failure occurrence depend on location?
 - ▶ Does failure duration depend on location?
- **Correlate:** Do failures correlate with routing instability?
 - ▶ If so, under what circumstances?

The Trouble with Traceroute

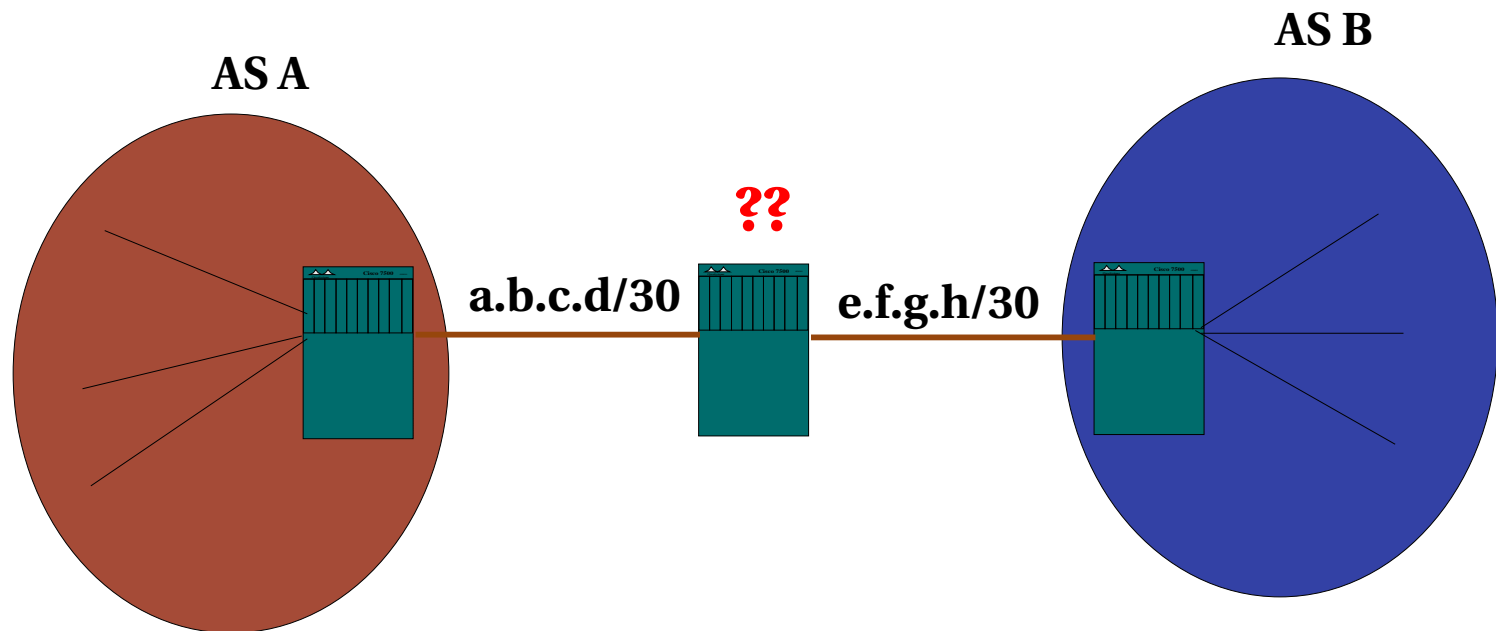
- Could reflect failure on the reverse path.
Solution: Trigger based on observed path failures.
- Tells us about interfaces, not routers.
 - ▶ Maybe even the outgoing interface of the return packet!*Solution:* Disambiguation techniques.
- Little information about AS boundaries.
Solution: Use knowledge about the topology...
- Some failures may reflect convergence issues.
Solution: BGP Hints?

Disambiguation (Alias Resolution)




- Rocketfuel's IP ID trick. Run the test 100 times per pair to gain confidence.

AS Edge Resolution with Limited Traceroutes



- Edge Resolution Algorithm:
 - ▶ Voting: IP addresses/router (1% have 3 addresses)
 - ▶ Some routers are clearly inside an AS. (~27%)
 - ▶ Voting: edges towards each AS (inductive). (~70%)
 - ▶ Last resort: traceroute to the router in question.
- From each failure: router, AS, location in AS
- Future Work: Failure Trajectory

Computing Distances



- Links directly connected to hosts have distance zero.
- New links introduce **two interfaces**.
 - ▶ At least one of these must connect to a host for which we've assigned a distance.
 - ▶ Assign minimum distance to end host.


Characterization Methodology

- 20 pairwise nodes, topologically distributed
 - ▶ 6 with BGP feeds
- Periodic pairwise probing. Trigger traceroutes upon failure.
- Failure: 3 consecutive lost probes, >2 minutes
 - ▶ Results may be affected by faults (as described before).

Failure Characterization

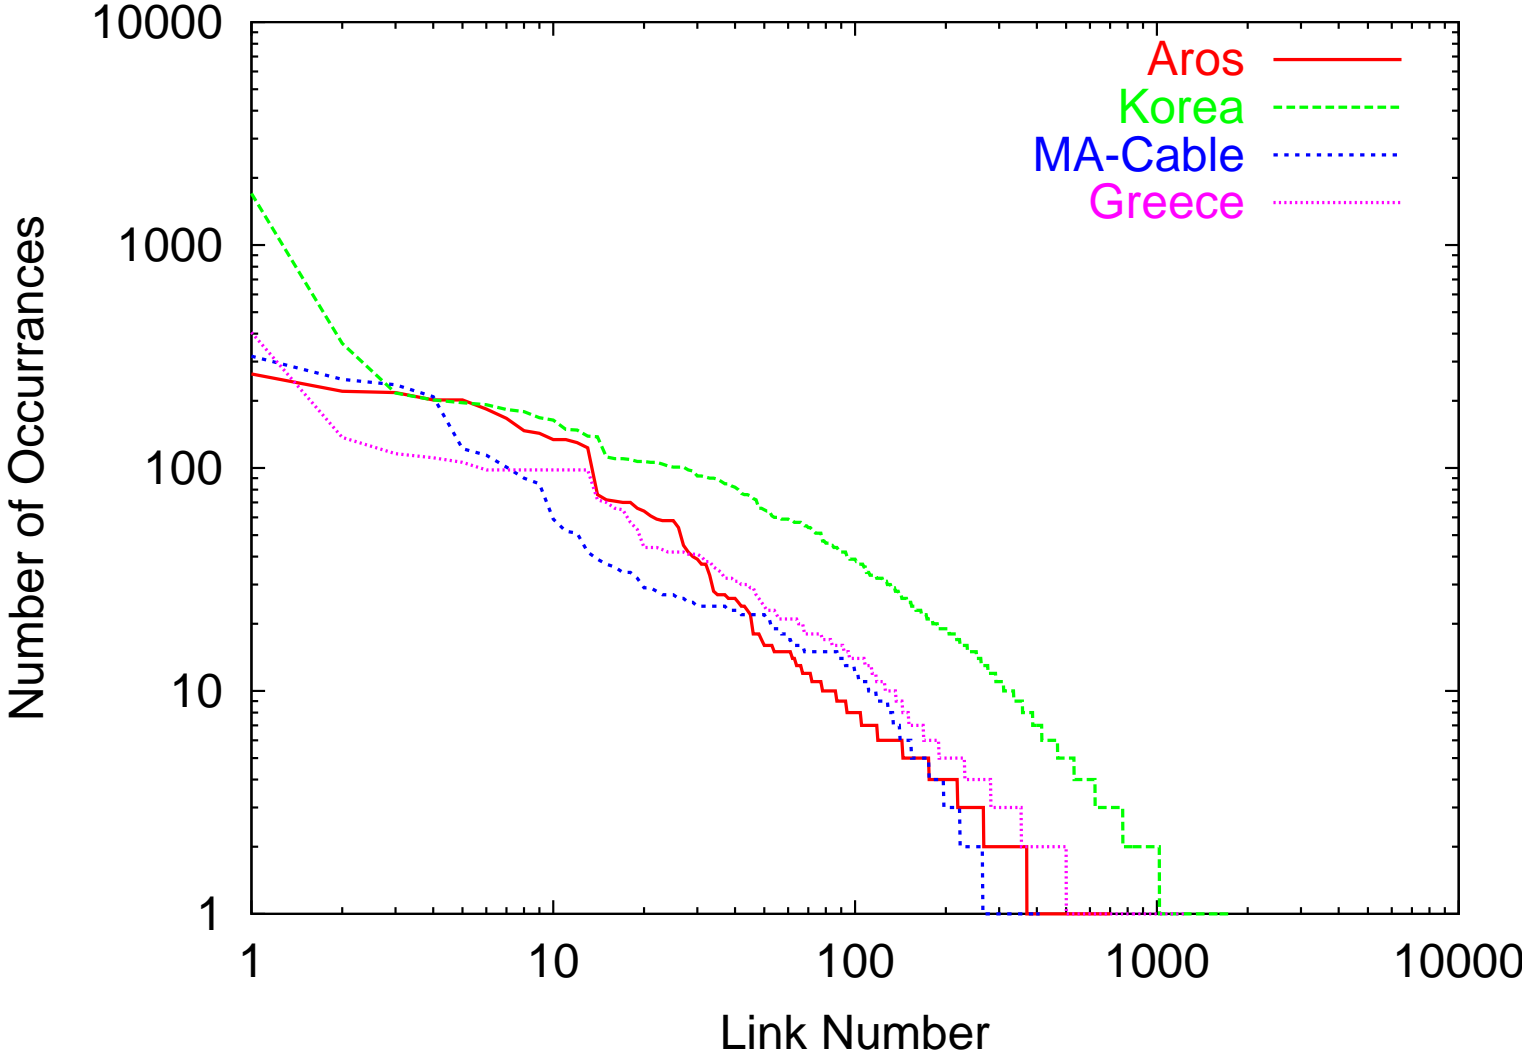
- Where are they occurring?
- How long do they last?
- How does outage duration depend on
 - ▶ link type? (edge vs. non-edge)
 - ▶ distance from last hop?

Failures Occur Near the Last Hop

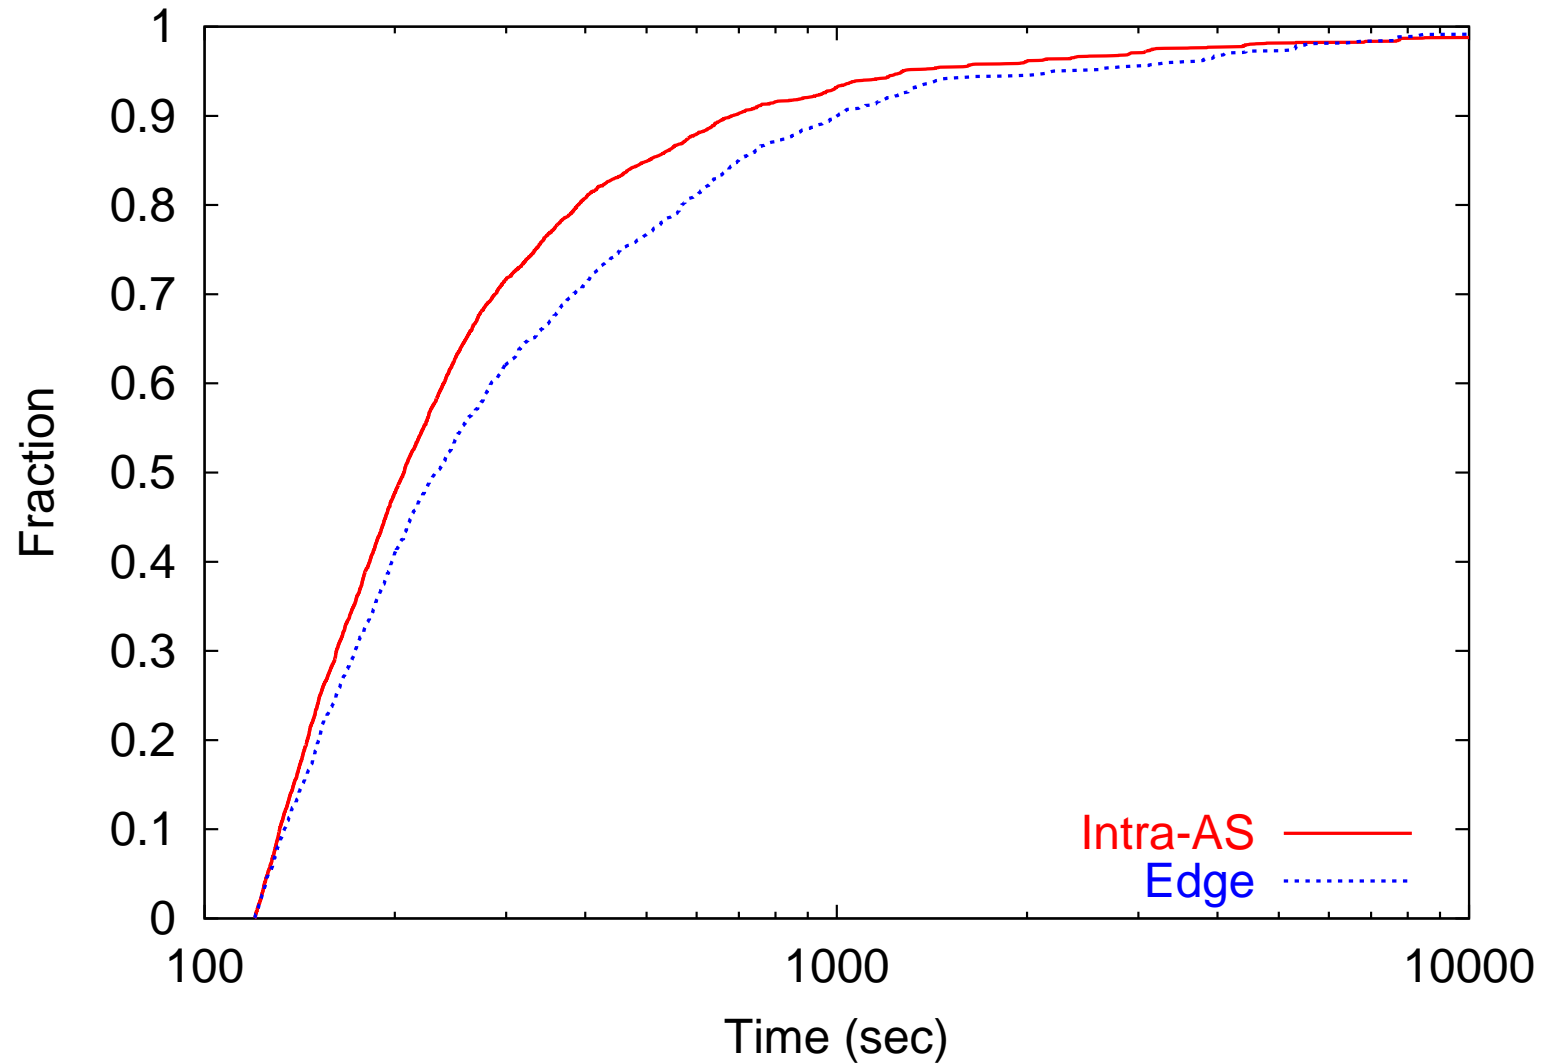


2/3 of observed failures occur intra-AS. Why?

A Few Bad Apples



Failure Duration




*Observed failures on AS boundaries last slightly longer.
Failure durations do not reflect distance from last hop.*

Correlating Failures and Routing Instability


- Do path failures correlate with routing instability? (under what circumstances)
 - ▶ Location of failure (i.e., distance from end, link type)
 - ▶ Advertisement type (e.g., degree of aggregation, etc.)
 - ▶ Path diversity
- If so, how do they correlate in time?

Degree of Correlation Depends on Host




Failures inside an autonomous system are less likely to be reflected by routing instability than failures on AS boundaries.

Time-based Correlation



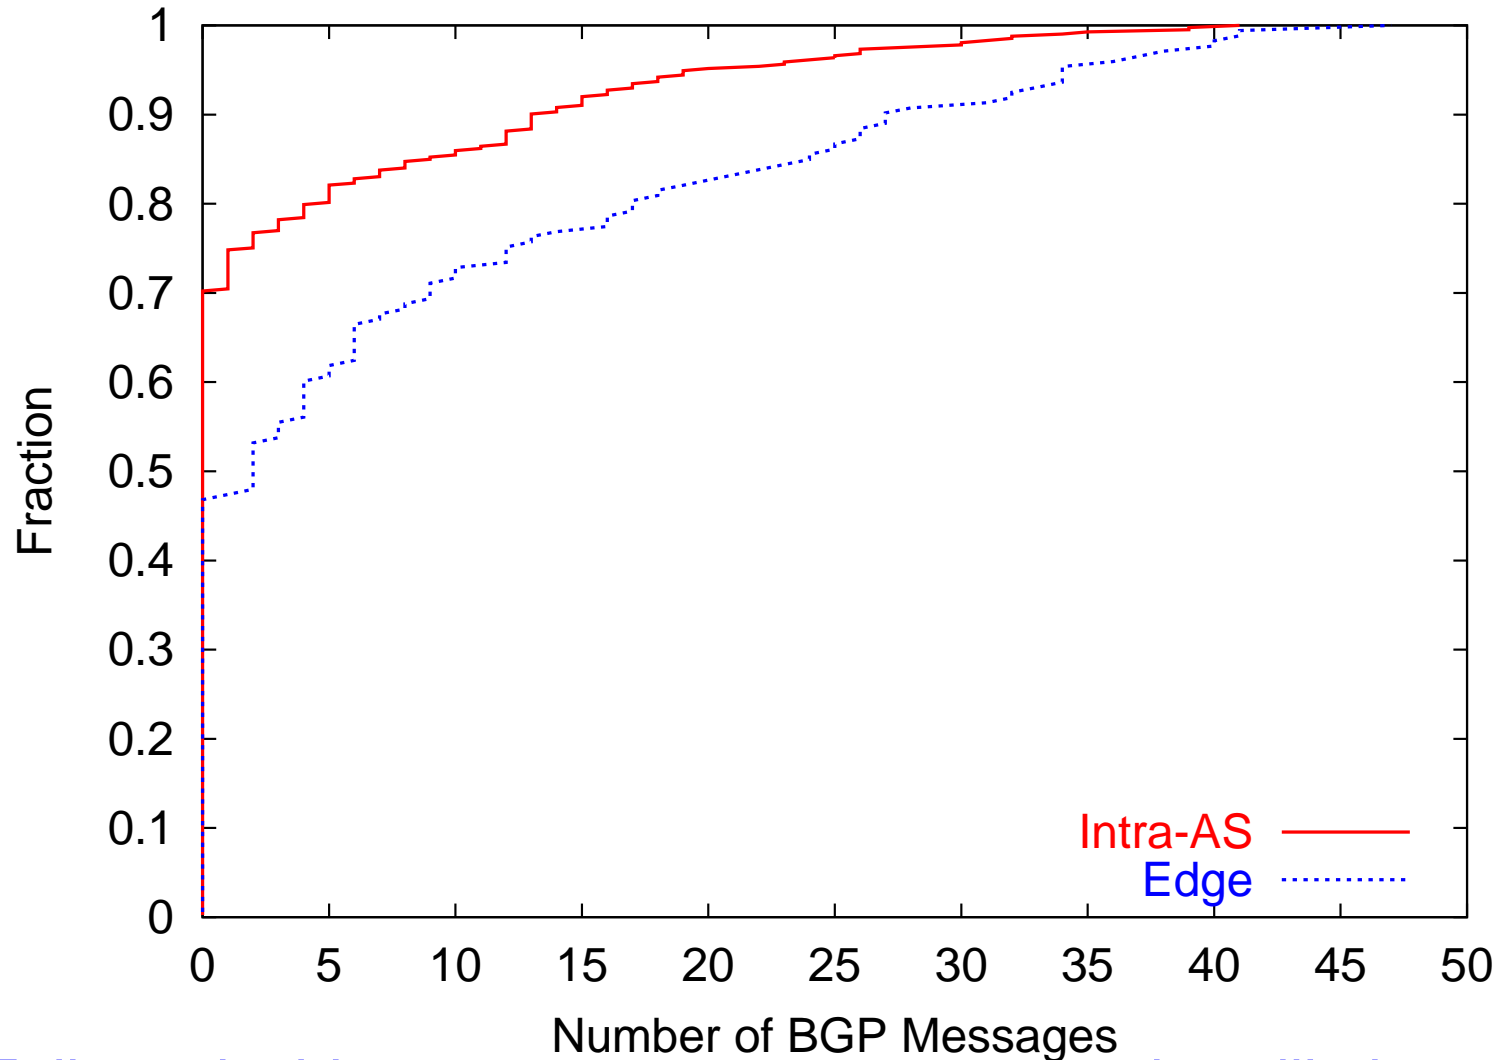
Failures occur several minutes before BGP activity.

Alternate-Route Search



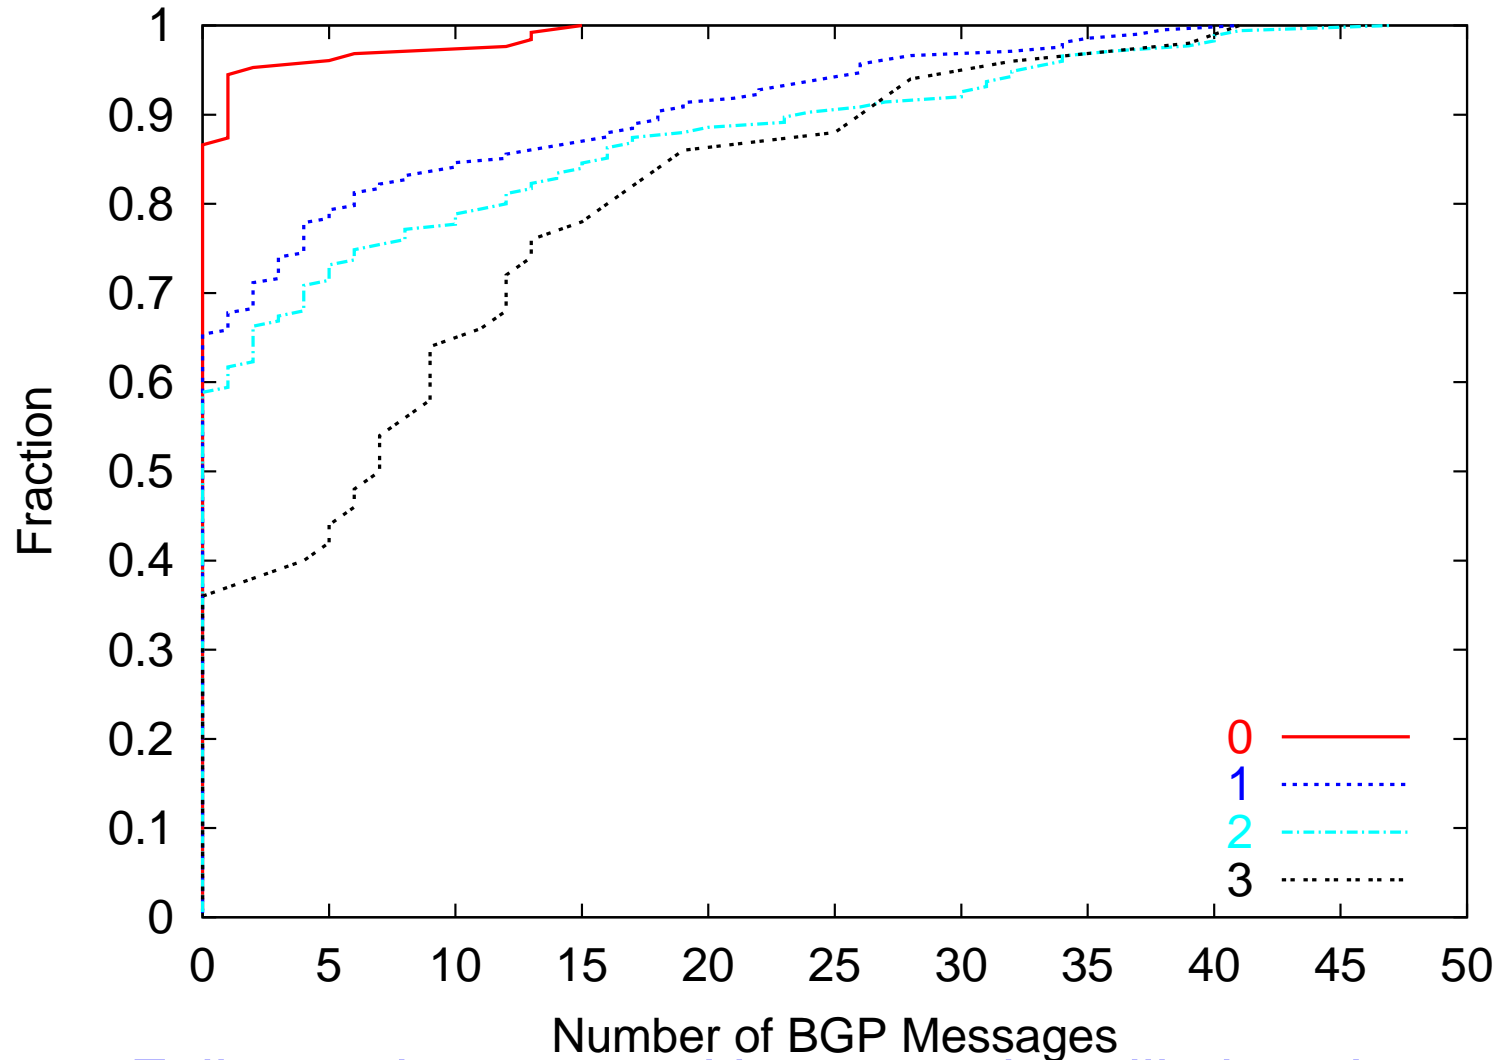
- Failures are commonly accompanied by a march through alternate routes.
- In what cases do we see this, and to what degree?

Correlation: Failures on AS Boundaries



Failures inside an autonomous system are less likely to be reflected by routing instability than failures on AS boundaries.

Correlation: Distance from Last Hop



Failures closer to end hosts are less likely to be reflected by routing instability.

Thoughts on Correlations

- Many possible explanations
 - ▶ Path Diversity
 - ▶ Level of Aggregation
 - ▶ Location of Failure
- Which of these explains well-correlated failures, in each case?
- Sometimes, continued instability is a sign of trouble to come (or continue). Predictors?

Conclusions

- Locating Failures
- Characterizing Path Failures
- Correlating with Routing Instability
- Current work
 - ▶ Also need to analyze per host to minimize bias.
 - ▶ Need to do analysis for peering/transit links.
 - ▶ How do correlation trends look across different BGP feeds?