Stability Attributes of Network Availability

- Prevalence (Predictability)
  - Will you see a route again? (Probability)
- Persistence (Manageability)
  - Duration of how long a route is valid (Time)
- Prefix Availability
  - Fraction of time it is reachable
- Prefix Steadiness
  - Percentage of time continuously reachable
- Sensitivity metrics
  - routers/topologies
  - how prone is your h/w prone to traffic shifts
- Control plane + Data plane interactions
  - not completely quantified

Presentation of Results

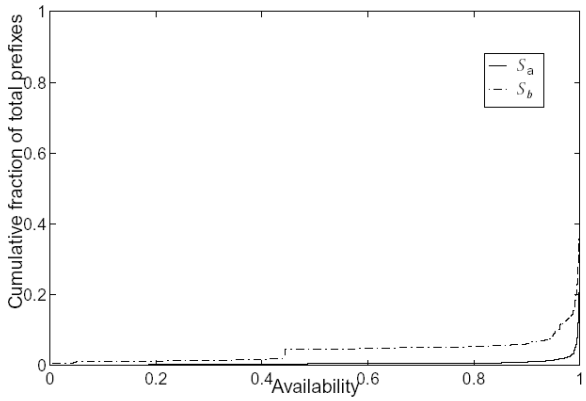

10% of prefixes available in less than 95 % of the time

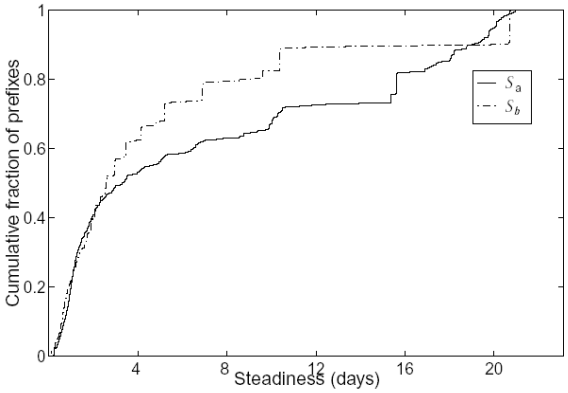

Some flaky routes (low steadiness), something in between, some very stable routes (high steadiness)

MTTF → in the order of 25 days
MTTR → in the order of 20 minutes or less

Problem Taxonomy

Pathology
- Excess work
- Inconsistency
- Vast majority of pathological changes – from small ISPs

Failures
- Routers, links
- Forwarding anomalies
    - loops
    - erroneous routing (RARE!)
    - connectivity alteration
- Hot Potato Changes
    - Shifts in net traffic → packet losses (happens infrequently: big impact)
- Labovitz's taxonomy for failures
    - Failure = path withdrawn and not replaced
    - Repair = path restoration
    - Fall-over = new path announcement
    - Experimental study of 1999:

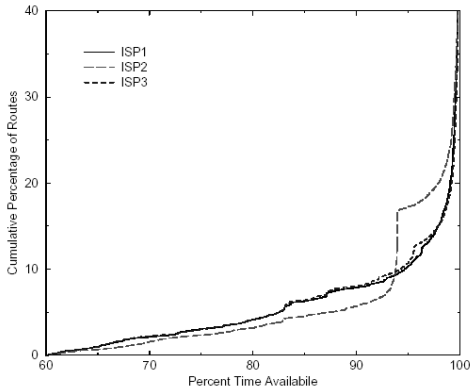

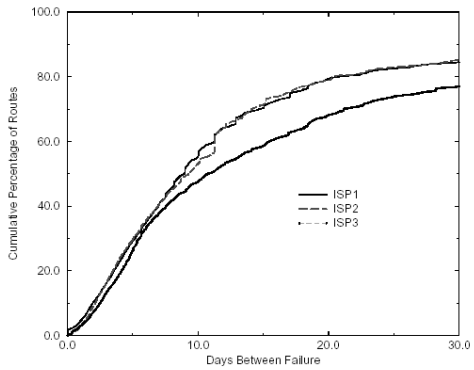

    - Top tree failure categories constitute almost 50% of failures:
    - Maintenance (16.2%)
    - Power Outage (16.0%)
    - Fiber Cut/Circuit/Carrier Problem (15.3%)

Congestion
- Will congestion ever lead to path failure?
- Artifact of protocol processing – it should actual be zero

Instability

- 1st derivative is high – routing flaps
- Policy fluctuations + normal convergence → Routing convergence
- Now we look at finer time scales
    - Who is responsible for fluctuations?
        - not dominated by small set of AS's/routes
- "Routing Stability in Congested Networks: Experimentation and Analysis"
    - Only interesting thing was methodology
    - Treated router as a black box
    - Quantifying time observations of different implementations

Anomaly Detection

- Methodology is important
- Main value: Holt-Winters approach
- Correlating multiple sources
- Combining with actual trouble ticket data
- Is there value in knowing something is higher than usual?
    - Anomaly Detection
- Previous papers concentrated on the phenomenon
- Detecting earlier can help solve the problem
- Anomalous behavior of routing announcements

Why Routing?

- Packet forwarding as an essential service
- Control plane
- Data plane → Disturbances → network availability
    - Update frequency → Route utilization (high message frequency/high updates)
    - Propagation of changes – dynamic – temporary impairments to reachability → higher than normal packet loss rate
    - Duration → stabilizes convergence
    - Intrinsic Distributed Algorithm → 30 seconds phenomenom
- Intra
- Inter domain (eBGP)

Timeline

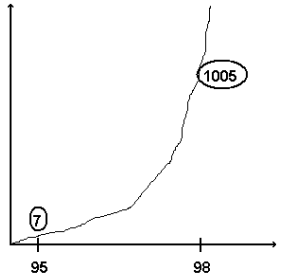

1995 Paxson Apr 1, 1995 (7 major AS) (Early day / early debugging) (200 AS)

1997 (Peering between ISPs)
1998 (1005 AS) In 1998: malicious attacks = 1-5% of failures

2001 congestion, worms
2003

2004 Managing T1 ISPs (1500 AS, 12 T1 ISPs)

End comments

- All previous papers
  - single view point
  - backbone
  - not end system to end system
- Reachability/Availability is the critical metric
  - not possible to achieve 99.9%
    - so go for redundant paths etc.