End-to-End Routing Behavior in the Internet

Vern Paxson
University of California, Berkeley and
Lawrence Berkeley National Laboratory
vern@ee.lbl.gov

Outline

- Methodology
- Pathologies
- Stability
- Symmetry

Methodology(1/3)

- Network Probe Daemon (NPD)
  - D1: measured interval = 1~2 days
  - D2: 60% with interval = 2 hours
    40% with interval = 2.75 days
- Exponential distributed random interval
- Be Representative?
  - 37 hosts v.s. 1000 active A.S.

Methodology(2/3)

- Observation Periods
  - D1: Nov.8 – Dec.24, 1994
  - D2: Nov.3 – Dec.21, 1995
- Confidence intervals: c = 95%

$$p_l = \frac{v_2}{v_2 + v_1 Q_{F(v_1,v_2)}(1-c)}$$

$$v_1 = 2(n-k+1) \qquad v_2 = 2k$$

Methodology(3/3)

- Shortcomings
  - End-to-end observations can’t see details.
  - Centralized design, batch process is better.
  - ‘traceroute’ is not sophisticated enough.

Routing Pathologies

Routing Loops

- Types
  - Forwarding: packet travels several routers.
  - Information: Information from myself.
  - ‘traceroute’: duplicated hops.*
- Reason: routing updates inconsistency.
- Results:
  - Persistet: D1/D2 = 10/50
  - Temporary: D1/D2 =2/23
  - Location:Confined to a single A.S. (BGP protection)

Erroneous routing & Connectivity altered mid-stream

- Erroneous
  - Packet takes wrong path
  - D1/D2 = 1/0
- Altered mid-stream
  - ‘traceroute’ reports lost or altered.
  - Unable to choose path caused by outages
  - D1/D2 = 10/155

Fluttering(1/2)

- Rapidly-oscillating routing.
- Benefits: load average
- Problems:
  - Unstable network path
  - Asymmetric path
  - Difficult to calculate RTT, Avail. BW, etc.

Fluttering(2/2)

- Example
  - From St. Louis
  - To Mannheim, Germany
  - Fluttering:
    1. Washington D.C (17 hops)
    2. Anaheim (29 hops)

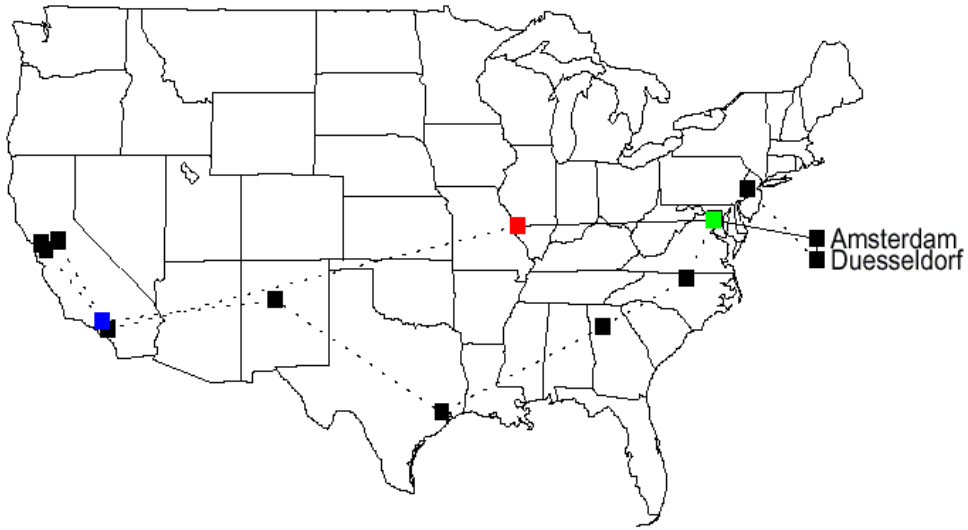

Infrastructure failures & Hop Count

- Infrastructure failures
  - Can’t reach destination hosts
  - Availability: D1/D2 = 99.7~99.9%/99.4~99.6%
  - Overestimates availability
- Unreachable due to too many hops
  - Out of the 30 hops limit of ‘traceroute’
  - Distance v.s. hops.
    - Km/Hops: 1500/3, 2000/5, 3/11

Temporary outages

- Loss network connectivity or congestion
- No looses: D1/D2 = 55%/43%
- 1~5 losses: D1/D2 = 44%/55%
- 6 or more losses: D1/D2 = 0.96%/2.2%

Time-of-day Patterns

- Mean of time-of-day at SRC & DST
- Measurement times in a hour.
  - Maximum 4.5% (0:00-1:00)
  - Minimum 3.8% (13:00-14:00)
- Match daily congestion pattern
  - Maximum 8.0% (15:00-16:00)
  - Minimum 0.4% (1:00-2:00)

Pathologies Summary

Pathology	Probability(%)	Trend
Persistent loops	0.13-0.16	
Temporary loops	0.055-0.078	
Erroneous routing	0.004	
Mid-Stream change	0.16/0.44	Worse
Infrastructure failure	0.21/0.48	Worse
Outage >= 30 secs	0.96/2.2	Worse
Total pathologies	1.5/3.4	Worse

End-to-End Routing Stability

- Change frequency(probability)
- Stability
  - Prevalence
  - Persistence
- Data Source
  - D2 measurements without pathologies
  - 3-level of granularity: Host, City, A.S.

Routing Prevalence(1/2)

- Prevalence parameter
  - $\pi_r = K_r/n$
  - K_r: number of the same route ‘r’
  - n: number of observations.
- Prevalence parameter shows percentage of dominant routing path

Routing Prevalence(2/2)

- Measurement results

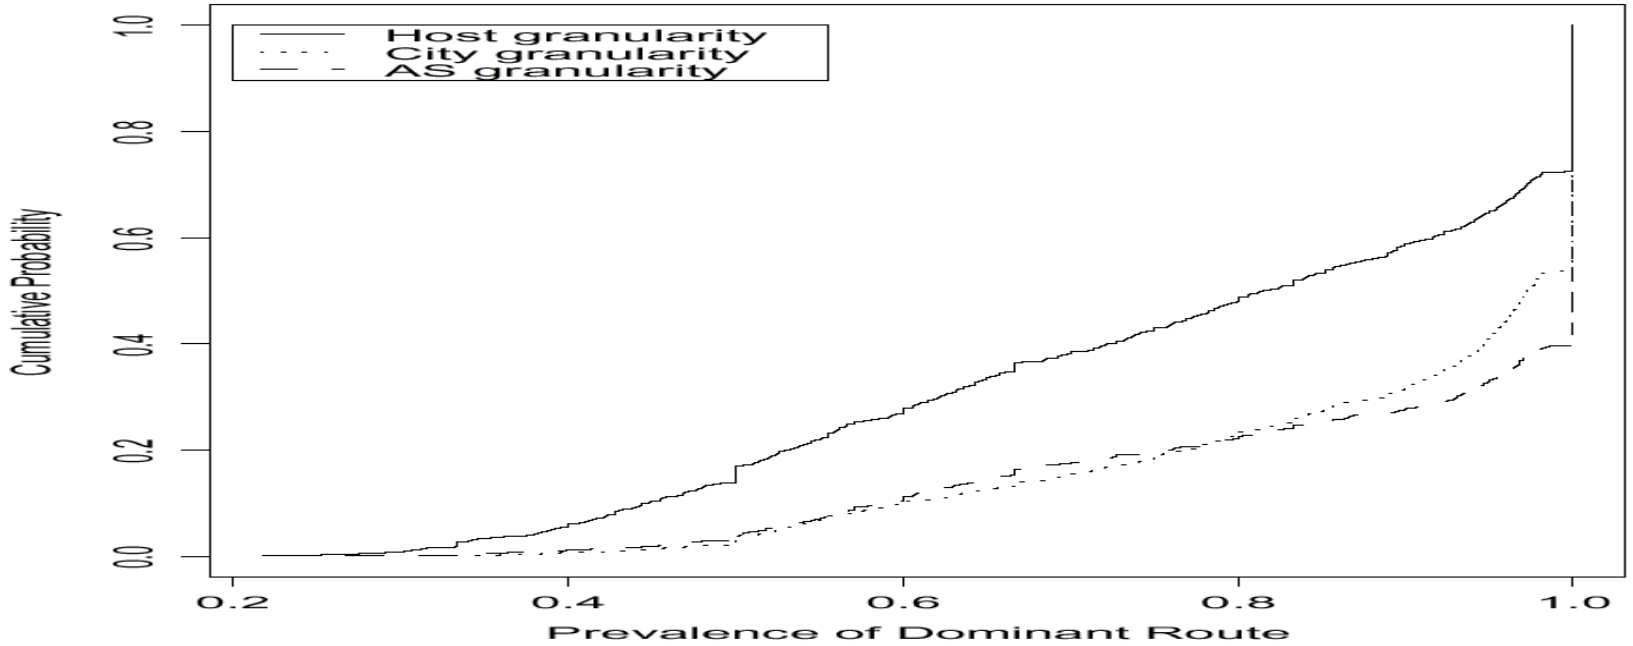

- Summary: Internet paths are strongly dominated by a single route.

Routing Persistence(1/3)

- Routing alternation
  - $N^{10}_{dst\ s}$: Number of routing path for dst=s
  - $X^{10}_{dst\ s}$: Number of routing path changes for dst=s
  - $P^{10}_{dst\ s}$: Probability of routing path changes = X/N

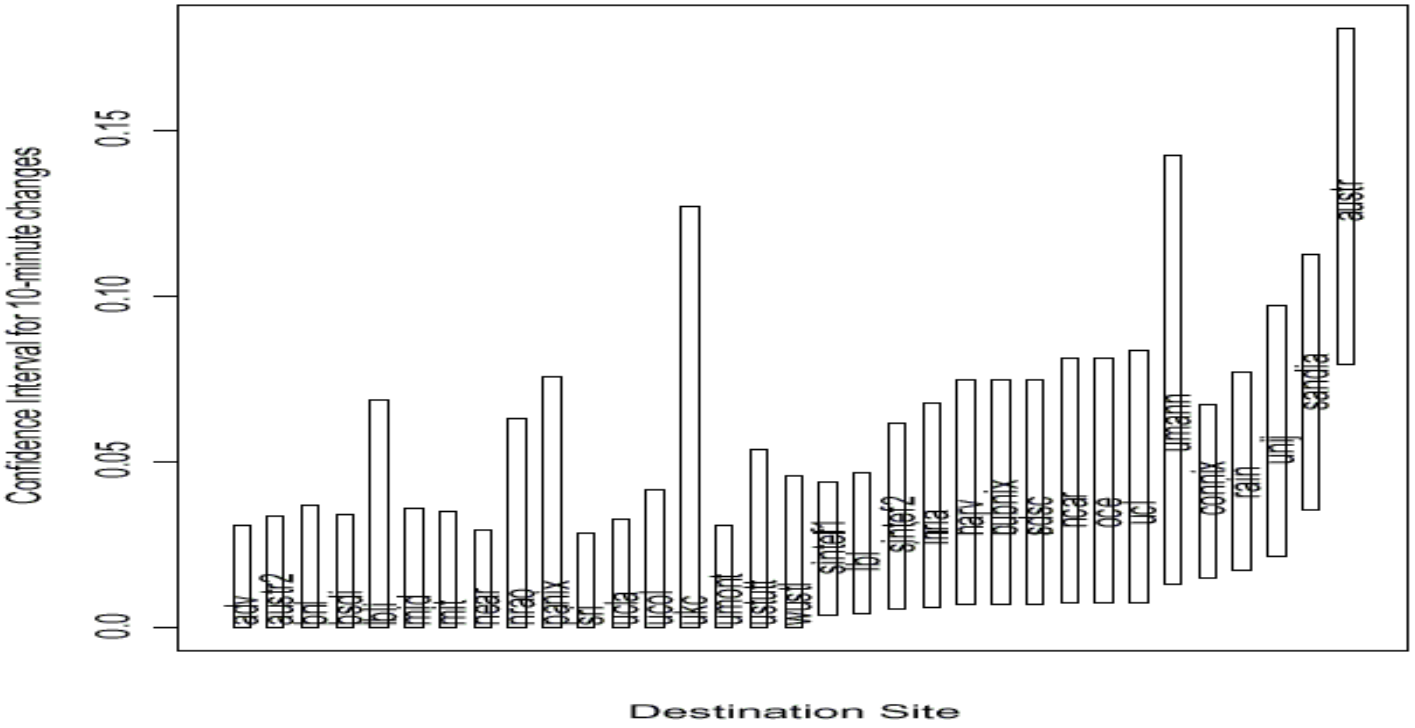

Routing Persistence(2/3)

- Duration Test Results

Time scale	Percentage
Seconds	N/A
Minutes	N/A
10‘s of minutes	9%
Hours	4%
6+ hours	19%
days	68%

Routing Persistence(3/3)

- Summary
  - Routing changes occur over a wide range of time scales.
  - Frequently path changes happens in smaller networks(Ex. Intra-AS).
  - 2/3 of internet paths persist for days or weeks.

Routing Symmetry

- Importance
  - RTT calculation.
  - Clock synchronization through network.
  - Routing asymmetry complicates network troubleshooting.
- Analysis result: D1/D2=30%/49%

Summary

- Routing pathologies increases: 1.5% to 3.4%.
- 2/3 of the Internet paths had routes persisting for either days or weeks.
- Routing asymmetry increases: 30% to 49%.
- Internet routing becomes less predictable.