End-to-End Routing Behavior in the Internet

Vern Paxson

Presented by Hee Dong Jung

Objective

- Understand the large-scale behavior of routing in the Internet
  - Routing behavior, not routing protocol
  - Analyze end-to-end measurements to determine:
    - Pathological conditions
    - Routing stability
    - Routing symmetry

Methodology

- Run Network Probe Daemon (NPD) on a number of Internet sites
  - Central control program: npd_control
  - Each NPD periodically measures the route to another NPD site using `traceroute`
  - How does `traceroute` work?
    - Start with a TTL (Time To Live) value of 1, get an ICMP reply from router that is 1 hop away
    - Next, use a TTL value of 2, get an ICMP message from router that is 2 hops away.
    - Continue until reach the destination

Methodology

- Two sets of measurements
  - D1: measure each virtual path between two sites with mean interval of 1-2 days
    - Each NPD `traceroute` once every two hours
    - Nov 8 to Dec 24 in 1994
  - D2: two different intervals combined
    - 60% with mean interval of 2 hours (bursts)
    - 40% with mean interval of 2.75 days
    - Paired measurements (A ⇒ B and immediately B ⇒ A)
    - Nov 3 to Dec 21 in 1995

Methodology

- Links traversed during D1 and D2

Routing Pathology

- Prevalence of routing loops
- Fluttering
- Temporary outages
- Connectivity altered mid-stream
- Infrastructure failure
- Erroneous routing
- Unreachable due to too many hops

Routing Pathology – Loops

Source	Dest.	Date	#	Location	Duration
inria	adv	Nov. 6	1	Washington	?
inria	near	Nov. 11	1	Washington	$\leq$ 3 hr
wustl	inria	Nov. 24	1	Washington	?
inria	pubnix	Nov. 12	1	Washington	?
inria	austr2	Nov. 15	1	Washington	?
sintef1	adv	Nov. 12	1	Washington	?
pubnix	sintef1	Nov. 8	1	Anaheim	?
ustutt	ucl	Nov. 11	16	Stuttgart	16–32 hr
connix	bsdi	Nov. 14	1	MAE-East	$\geq$ 10 hr
ustutt	austr	Nov. 14	1	same loop	
pubnix	sintef1	Nov. 14	1	Washington	$\leq$ 5.5 hr
austr	nrao	Nov. 15	1	College Park	?
many	oce	Nov. 23	12	Amsterdam	14–17 hr
ucol	ustutt	Nov. 24	1	San Francisco	?
ucol	inria	Nov. 27	1	Paris	$\leq$ 14 hr
mid	bsdi	Nov. 28	1	Washington	$\leq$ 3 hr
mid	austr	Dec. 6	1	Chicago	$\leq$ 3 hr
mit	wustl	Dec. 10	1	St. Louis	?
umann	nrao	Dec. 13	1	Heidelberg	?
ucl	mit	Dec. 14	1	Cambridge	$\leq$ 3 hr
near	ucla	Dec. 16	1	Los Angeles	?
sri	near	Dec. 17	1*	Palo Alto	?
near	sri	same	1*	San Francisco	?
bsdi	sintef1	Dec. 21	1	NJ, London	$\leq$ 10 hr

Table 2: Persistent routing loops in $\mathcal{D}_2$

- Persistent loops
  - Loop unsolved by end of the `traceroute`
  - 10 in D1 / 50 in D2
  - Two types of duration (≥10 hrs / ≤ 3 hrs)
  - Clustered by location / time
  - Only one span multiple cities

Routing Pathology – Loops

- Temporary loops
  - Loop resolved during the `traceroute`
  - 2 in D1 / 23 in D2
  - In the order of seconds
  - Widespread connectivity property
    - 40 sec outage ⇨ **loop in D.C. area** ⇨ **loss of connectivity all the way back to the source** ⇨ connectivity regained
    - May reflect "ripple effects"

Routing Pathology – Fluttering

- Fluttering example (large-scale):

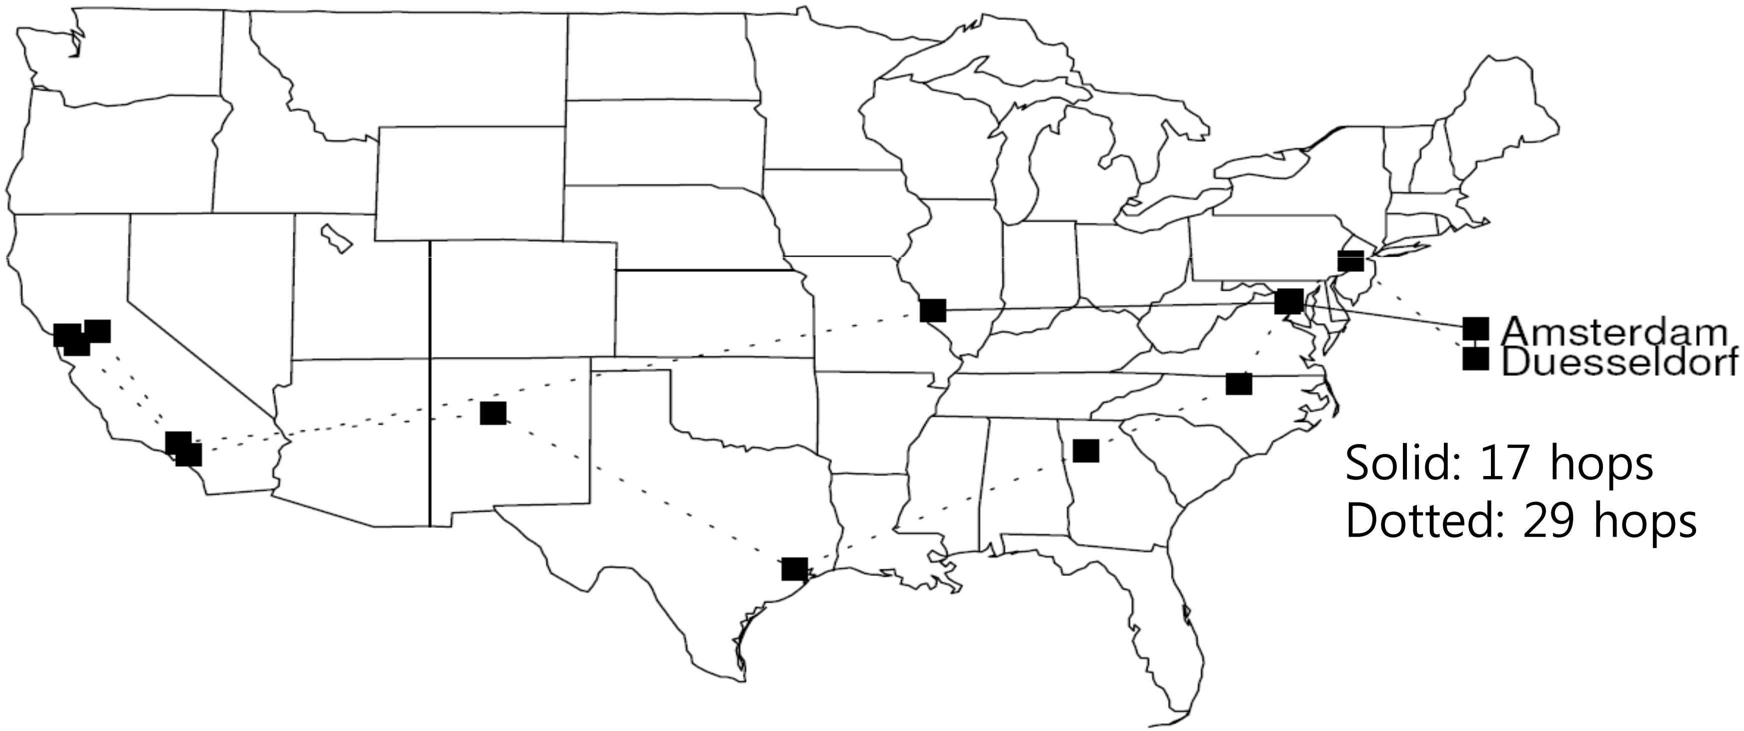

Route from St. Louis, Missouri to Mannheim, Germany

Routing Pathology – Temporary outages

- Sequence of `Traceroute` probes lost
  - Temporary loss of connectivity
  - Heavy congestion lasting more than 10 sec

- In D1, 55% had no losses, 44% had 1 to 5 losses, and 0.96% had 6 or more losses (≥ 30 sec outage)
- In D2, 43% had no losses, 55% had 1 to 5 losses and 2.2 % had 6 or more losses

- Outage more than 30 sec (6 or more losses)
  - Most prevalent pathology
  - Strong correlation with time-of-day patterns

Routing Pathology Summary

Pathology	Probability	Trend	Notes
Persistent loops	0.13–0.16%		Some lasted hours.
Erroneous routing	0.004–0.004%		No instances in $\mathcal{D}_2$.
Mid-stream change	0.16% ⊥ 0.44%	worse	Suggests rapidly varying routes.
Infrastructure failure	0.21% ⊥ 0.48%	worse	No dominant link.
Outage $\geq$ 30 secs	0.96% ⊥ 2.2%	worse	Duration exponent. distributed.
Total pathologies	1.5% ⊥ 3.3%	worse	

In 1995, the likelihood of encountering serious end-to-end routing problem (pathology) more than doubled, and was 1 in 30

Routing Stability

- Definitions
  - Prevalence: overall likelihood to observe a particular route
  - Persistence: how long a route remains unchanged
- Three levels of granularity
  - Host, City, AS

Routing Stability – Prevalence

- π_r : Steady-state probability that a virtual path at an arbitrary point in time uses a particular route r

- Unbiased estimator of π_r can be computed as
  $\hat{\pi}_r = {k_r}/{n}$

- Prevalence of dominant route p $\hat{\pi}_{domp} = {k_p}/{n_p}$

Routing Stability – Prevalence

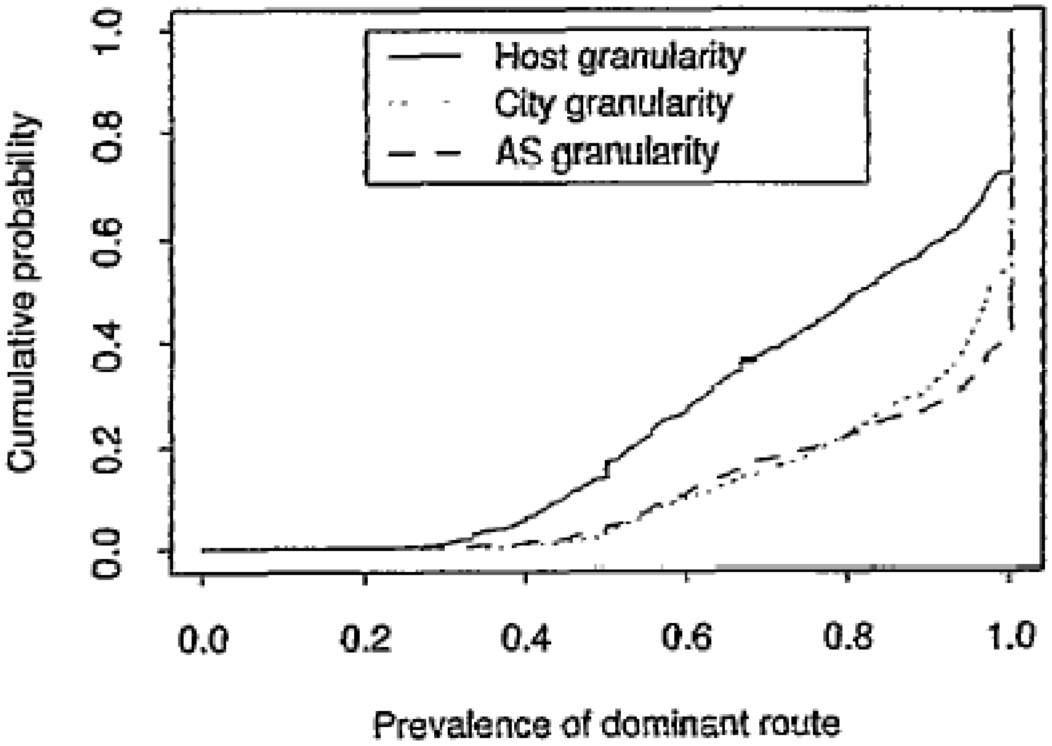

Median value : 82% (host), 97%(city), 100%(AS)

In general, Internet paths are strongly dominated by a single route

Routing Stability – Persistence

- Persistence at different time scales

Time scale	%	Notes
seconds	N/A	"Flutter" for purposes of load balancing. Treated separately, as a pathology, and not included in the analysis of persistence.
minutes	N/A	"Tightly-coupled routers." We identified five instances, which we merged into single routers for the remainder of the analysis.
10's of minutes	9%	Frequent route changes inside the network. In some cases involved routing through different cities or AS's.
hours	4%	Usually intra-network changes.
6+ hours	19%	Also intra-network changes.
days	68%	Bimodal. 50% of routes persist for under 7 days. The remaining 50% account for 90% of the total route lifetimes.

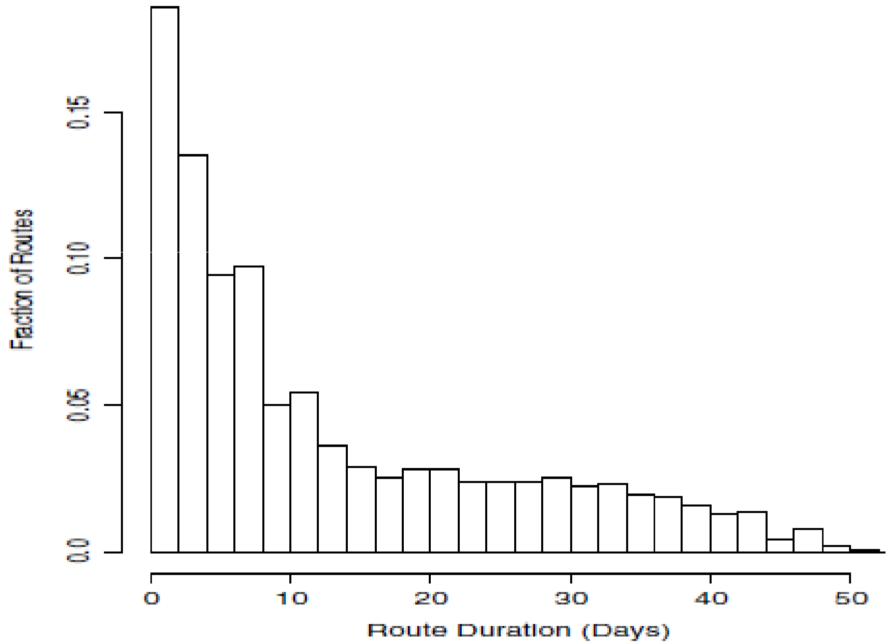

- 90% chance of observing a route with a duration of at least a week.

Routing Symmetry

- Analysis
  - Paired measurements to ensure asymmetry is actually being captured
  - Asymmetry is quite common (49% on a city granularity, 30% on AS granularity)
  - Large range of asymmetry involving different sites
- Size
  - Majority have single “hop” (one city / AS) asymmetry

Conclusion

- Likelihood of encountering routing pathology more than doubled between 1994 to 1995 (1.5% to 3.4%)
- Paths heavily dominated by single route
- Wide variation of persistence of route
- Asymmetry is common

- No typical Internet path

Discussion Points

- What are the consequences of fluttering?
  - Good or Bad?

- Implications of this paper?

- Is there a better way to learn about routing behavior?

Thank you

Questions?

Methodology (backup)

- Exponential sampling
  - Time intervals: independent, exponentially distributed
    - Additive Random Sampling: unbiased
    - PASTA (Poisson Arrivals See Time Averages) principle
- Representativeness
  - Routes include non-negligible fraction of AS's
- Devised a method to calculate and compare confidence intervals

Methodology (backup)

- Shortcomings
  - Not enough analysis provided on routing difficulties uncovered
  - Difficult to find out why and where in the path the problem occurred with end-to-end measurements
  - Centralized design issue
  - Only small subset of Internet routes
  - Only two points at a time

Routing Pathology – Route Change (backup)

- Mid Stream change
  - Route change during `traceroute` ⇒ outage
  - 10 in D1 / 155 in D2
  - Bimodal recovery times (seconds or minutes)

- Fluttering
  - Rapidly oscillating routing
  - Two cases (large-scale, localized)

Routing Pathology – Route Change (backup)

- Fluttering Problems
  - Difficulties from unstable network paths
  - Routing asymmetry problem
  - Unreliable path characteristic estimation
  - Out of order packets can lead to spurious "fast retransmissions" wasting bandwidth

- Localized fluttering is usually fine

Routing Pathology – Infra Failure (backup)

- Failure to reach destination
- Reasons other than loops and erroneous routing
- Estimated infrastructure availability
  - 99.7 ~ 99.9 % in D1
  - 99.4 ~ 99.6 % in D2
- Some correlation with time-of-day patterns
  - Peak: 1500~1600, 2nd Peak: 0600~0700, Min: 0900~1000

Routing Pathology – Too many hops (backup)

- `Traceroute` probe maximum of 30 hops
- None in D1 / 6 in D2
  - Internet has grown larger
- Hop count not necessarily correlated with distance
  - 1,500 km end-to-end route of 3 hops
  - 11 hops in 3 km distance