End-to-End Internet Video Traffic Dynamics: Statistical Study and Analysis

Dmitri Loguinov
City University of New York

Hayder Radha
Michigan State University

Overview of the Talk

- Motivation
- Experimental setup and experiment overview
- Results
  - Packet loss
  - Underflow events
  - Round-trip delay
  - Delay jitter
  - Packet reordering
  - Asymmetric paths
- Conclusion

Motivation

- Market research by Telecommunications Reports International (TRI) from August 2001

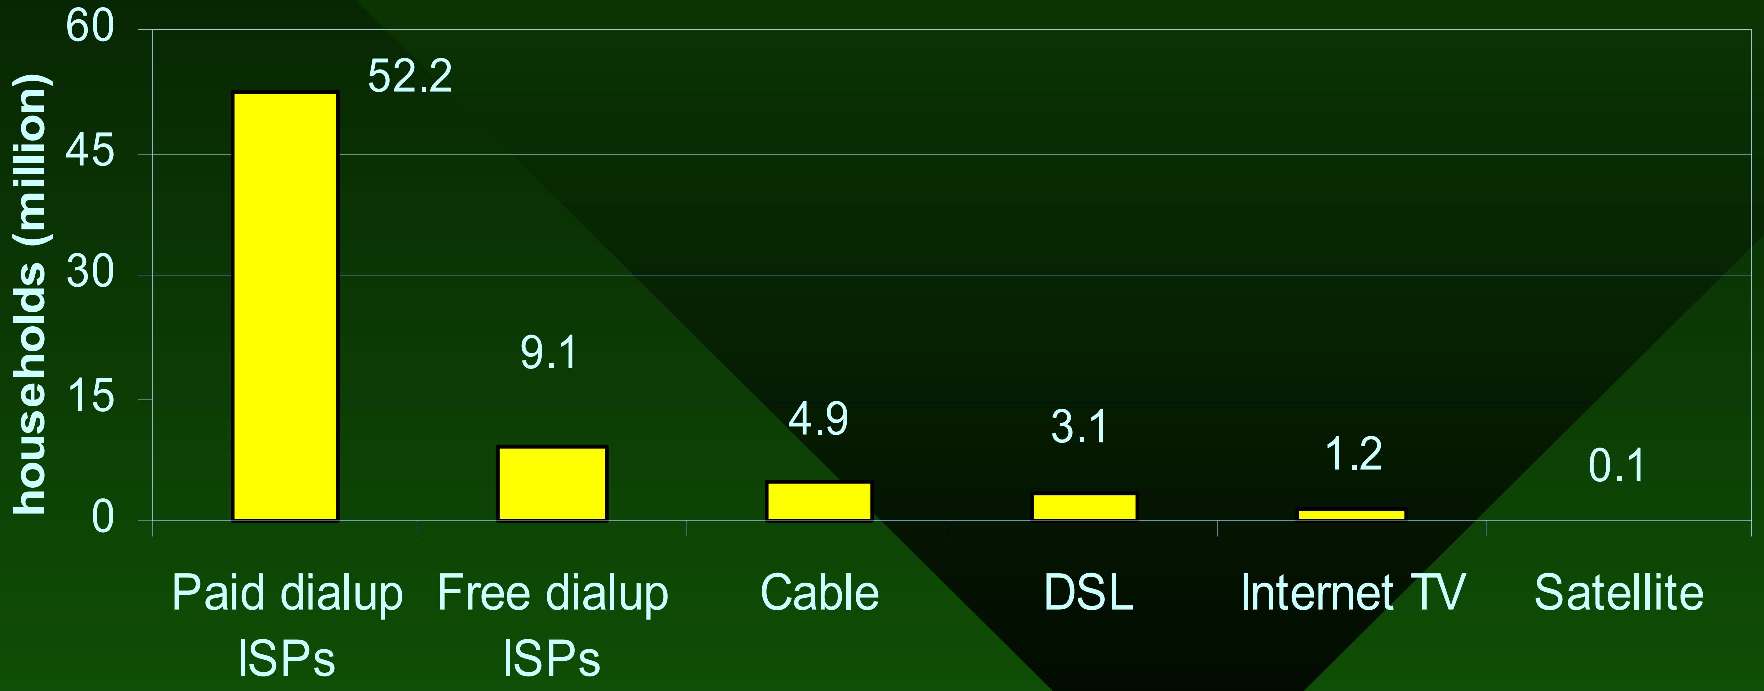

Motivation (cont'd)

- Consider broadband (cable, DSL, and satellite) Internet access vs. narrowband (dialup and webTV)
- 88% of households use modems and 12% broadband

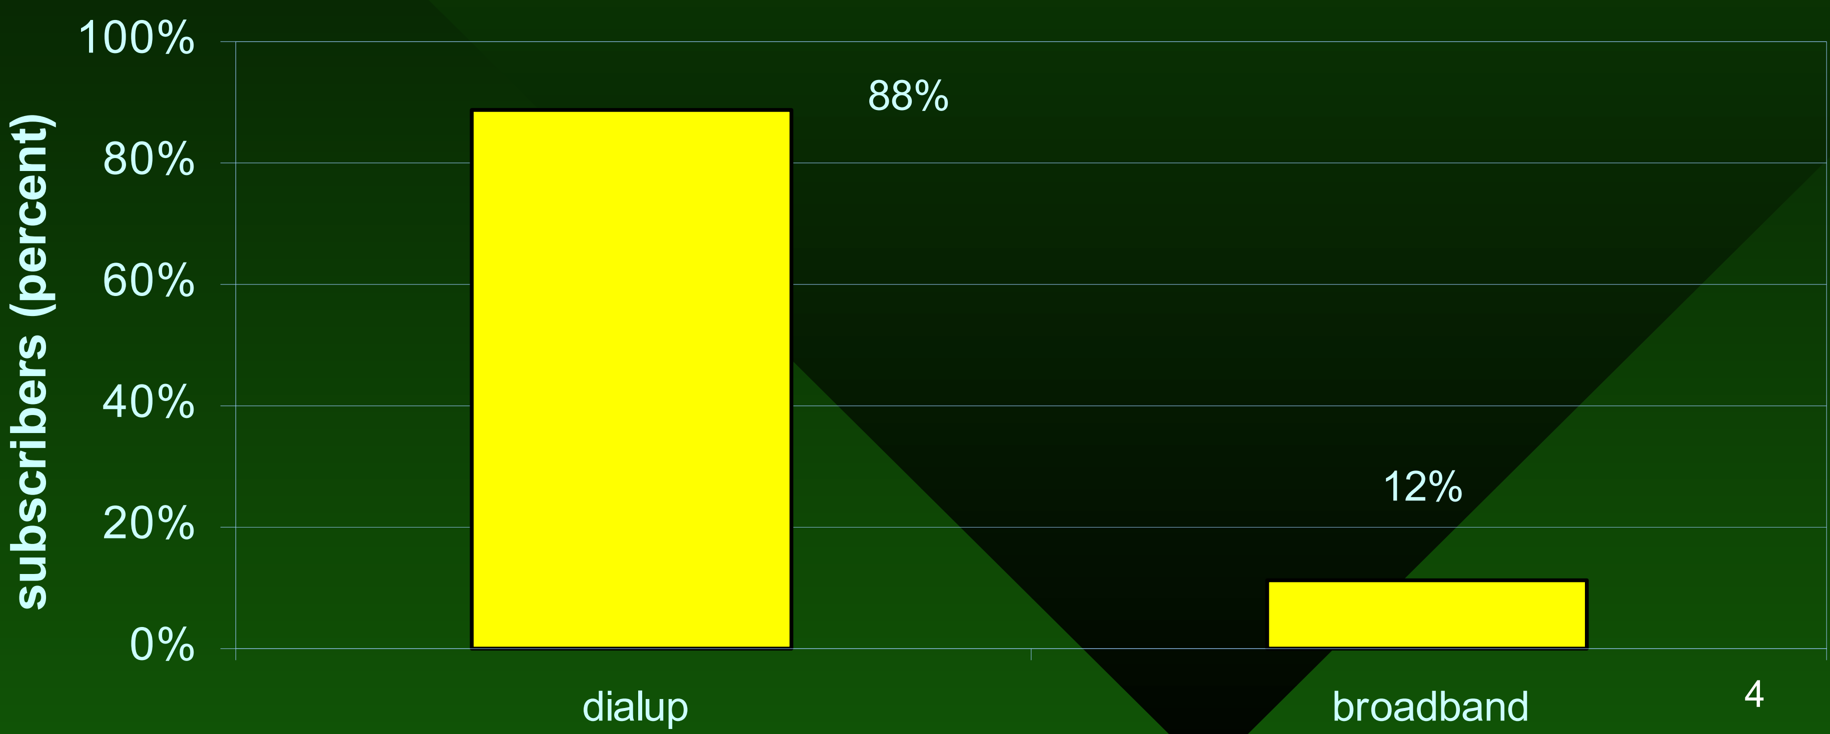

Motivation (cont'd)

- Our study was conducted in late 1999 – early 2000 when modems were more wide-spread than today
- Path properties of dialup ISPs and the view of the Internet from the angle of home users have not been documented
- Furthermore, large-scale performance of end-user video streaming in the current Internet has not been reported
- Main question – what is the main impediment to streaming? Is it delay, loss, or something else?

Experimental Setup

- MPEG-4 client-server real-time streaming architecture
- NACK-based retransmission and fixed streaming bitrate (i.e., no congestion control)
- Stream S_1 at 14 kb/s (16.0 kb/s IP rate), Nov-Dec 1999
- Stream S_2 at 25 kb/s (27.4 kb/s IP rate), Jan-May 2000

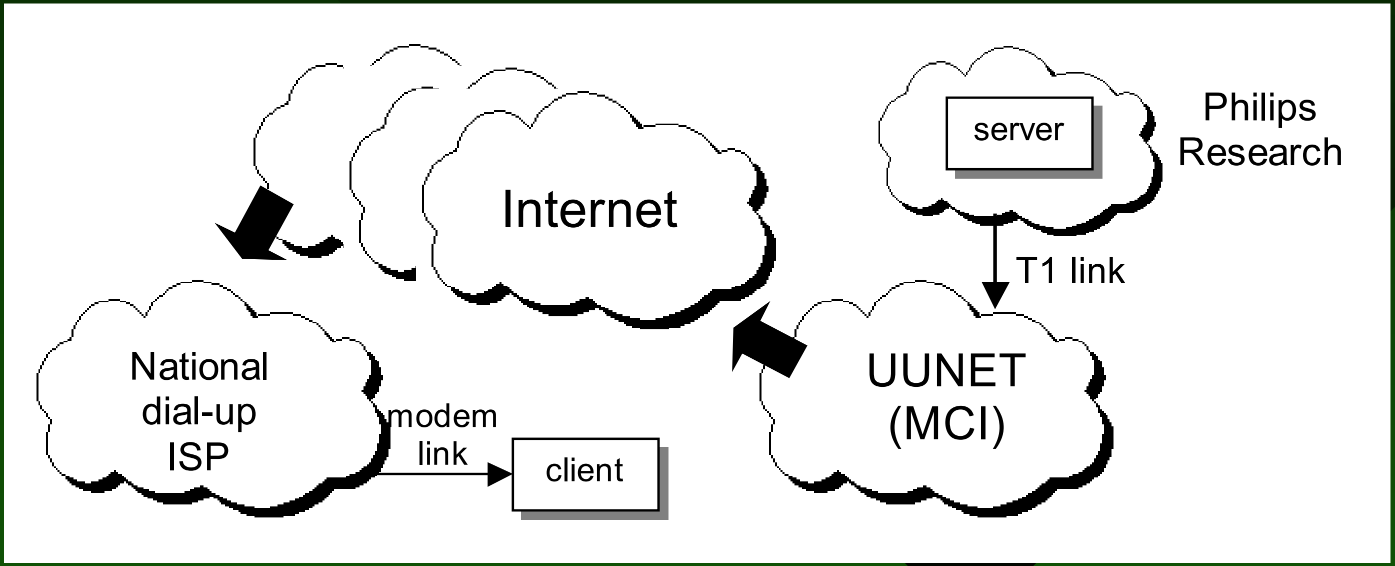

Overview

- Three ISPs (Earthlink, AT&T WorldNet, IBM Global Net) – phone database included 1,813 dialup points in 1,188 cities
- The experiment covered 1,003 points in 653 US cities
- Over 34,000 long-distance phone calls
- 85 million video packets, 27.1 GBytes of video data
- End-to-end paths with 5,266 Internet router interfaces
- 51% of routers from dialup ISPs and 45% from UUnet
- Sets D_{1p} and D_{2p} contain *successful* sessions with streams S_1 and S_2, respectively

Overview (cont'd)

- Cities per state that participated in the experiment

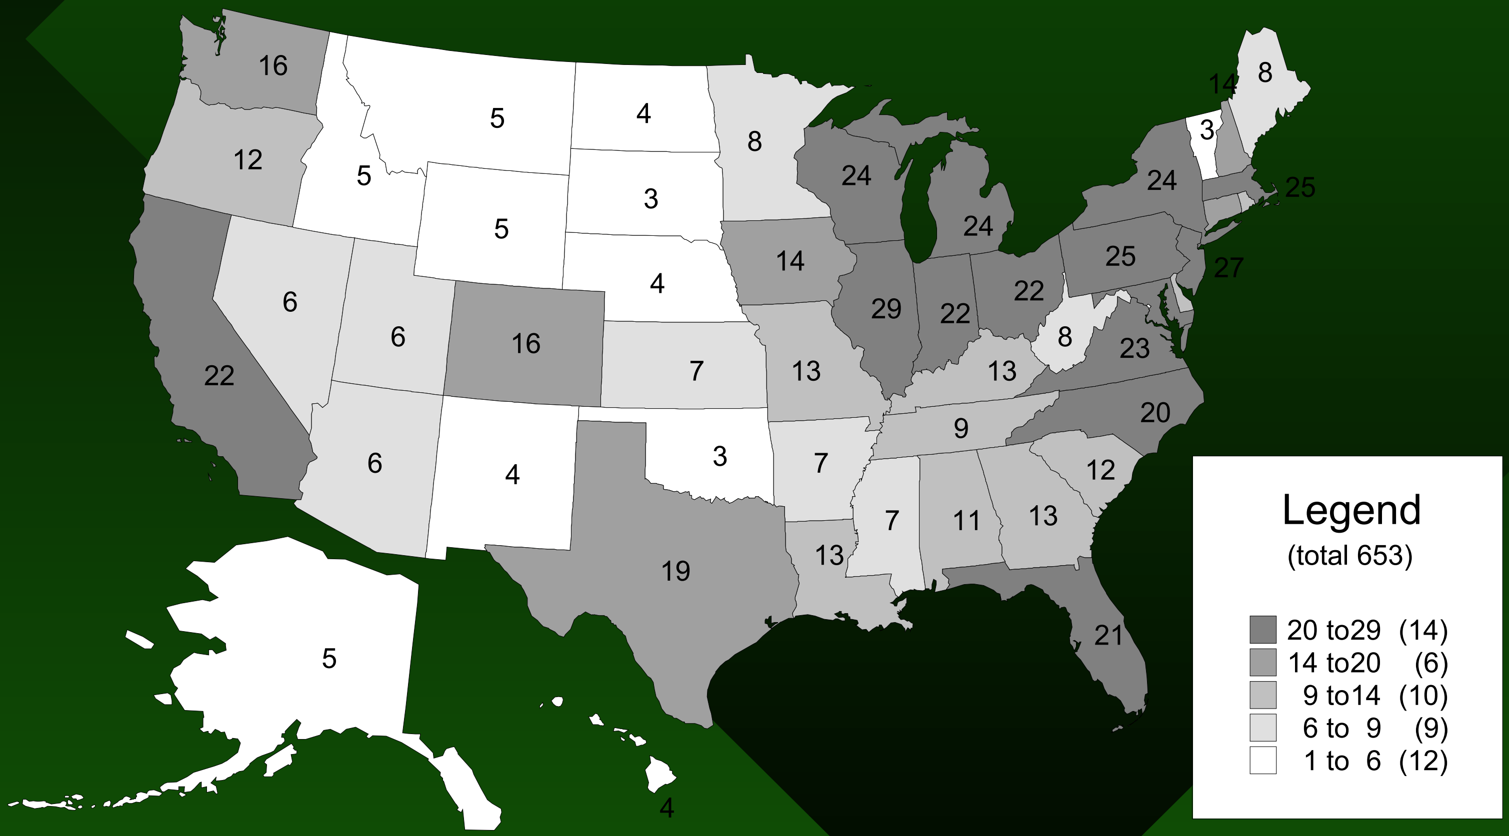

Overview (cont'd)

- Streaming success rate during the day shown below
- Varied between 80% during the night (midnight – 6 am) to 40% during the day (9 am – noon)

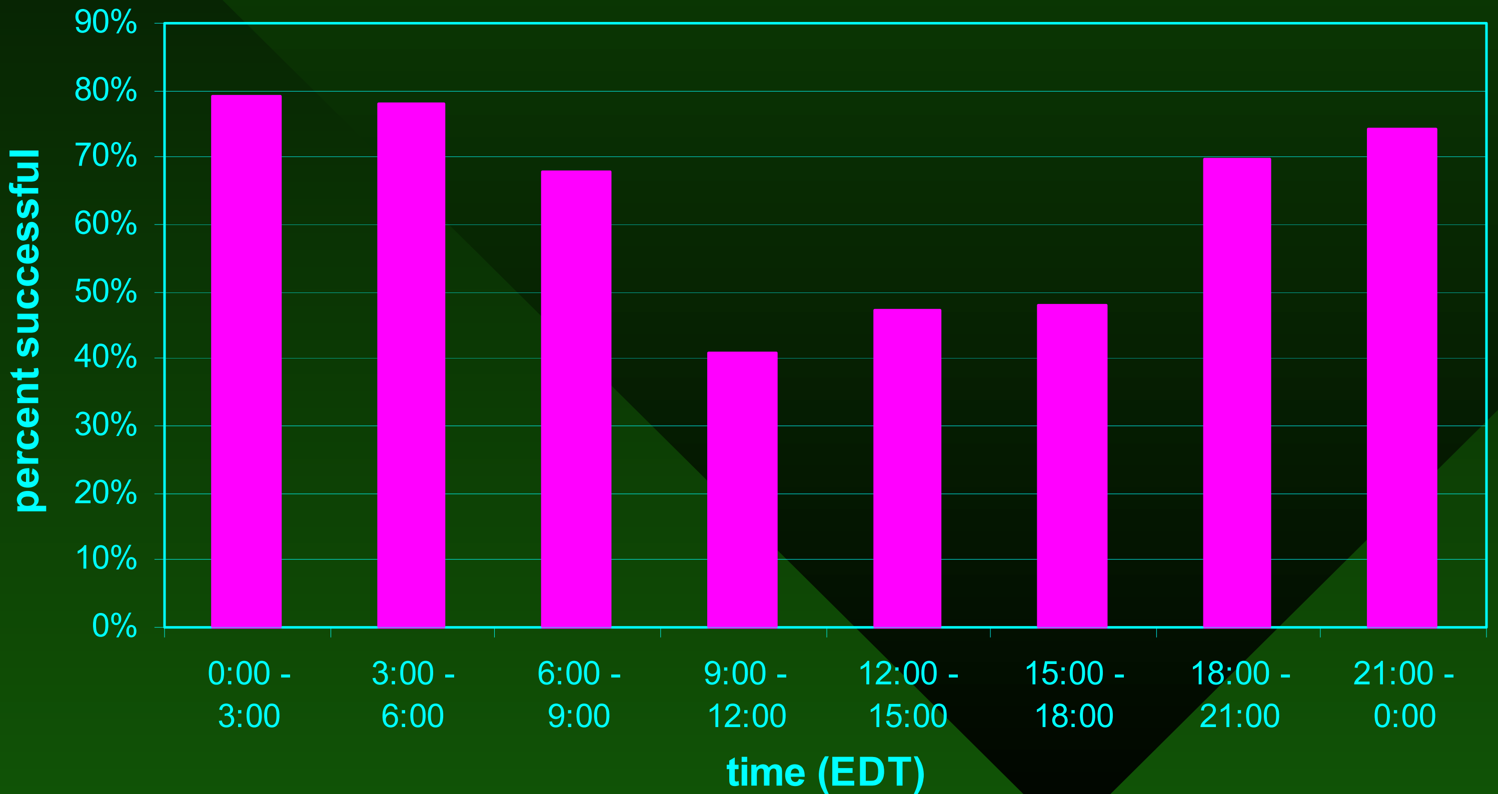

Overview (cont'd)

- Average end-to-end hop count 11.3 in D_{1p} and 11.9 in D_{2p}

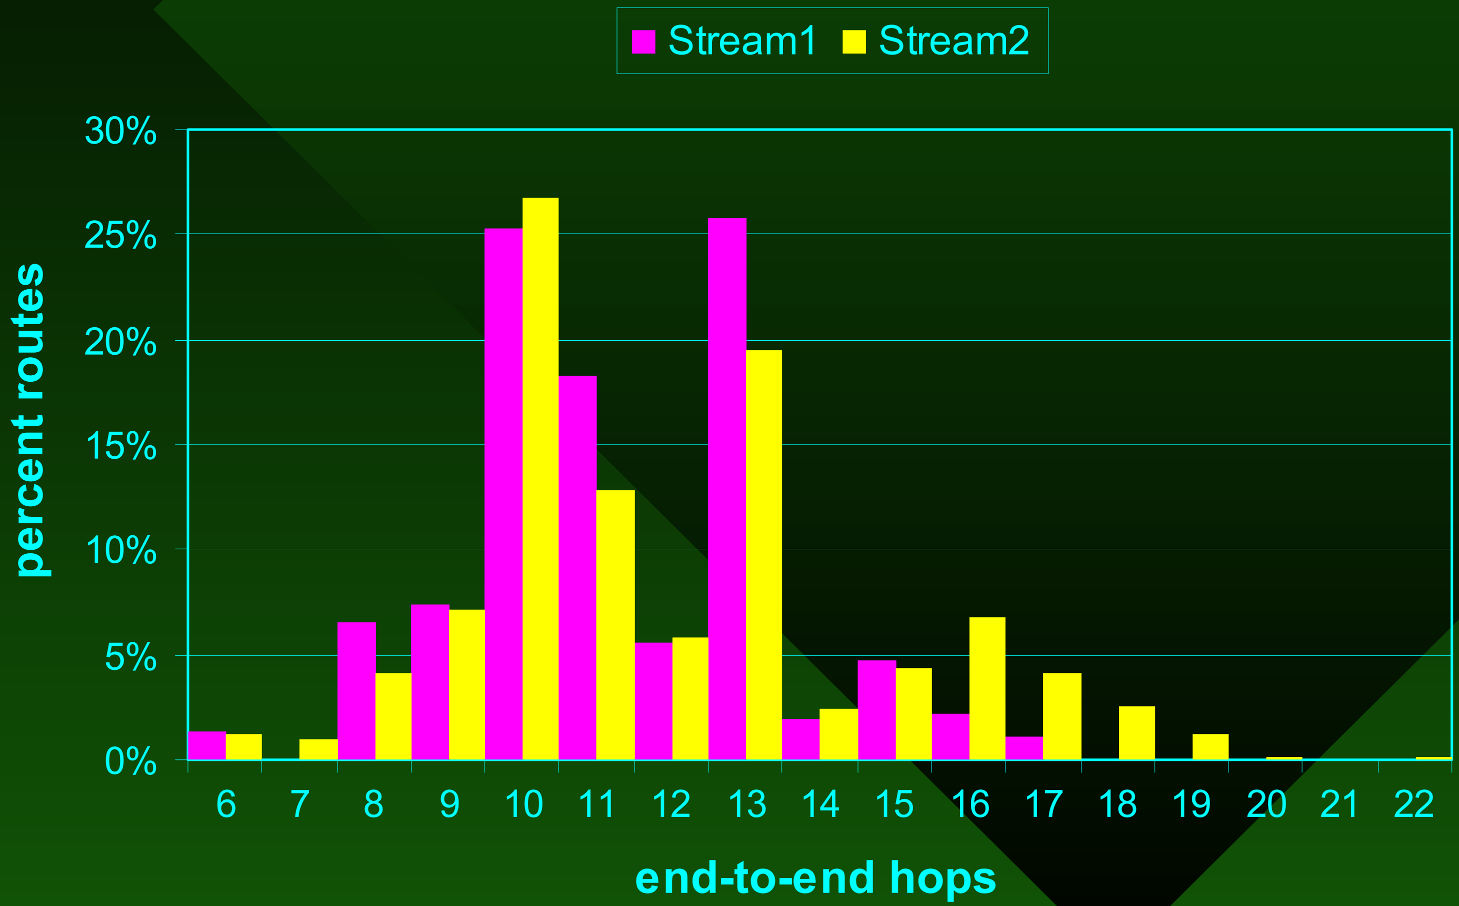

Packet Loss

- Average packet loss was 0.5% in both datasets
- 38% of sessions with no packet loss
- 75% with loss below 0.3%
- 91% with loss below 2%
- During the day, average packet loss varied between 0.2% (3 am - 6 am) and 0.8% (9am - 6pm EDT)
- Average per-state packet loss varied between 0.2% (Idaho) to 1.4% (Oklahoma), but did not depend on the average RTT or the average number of end-to-end hops in the state

Packet Loss (cont'd)

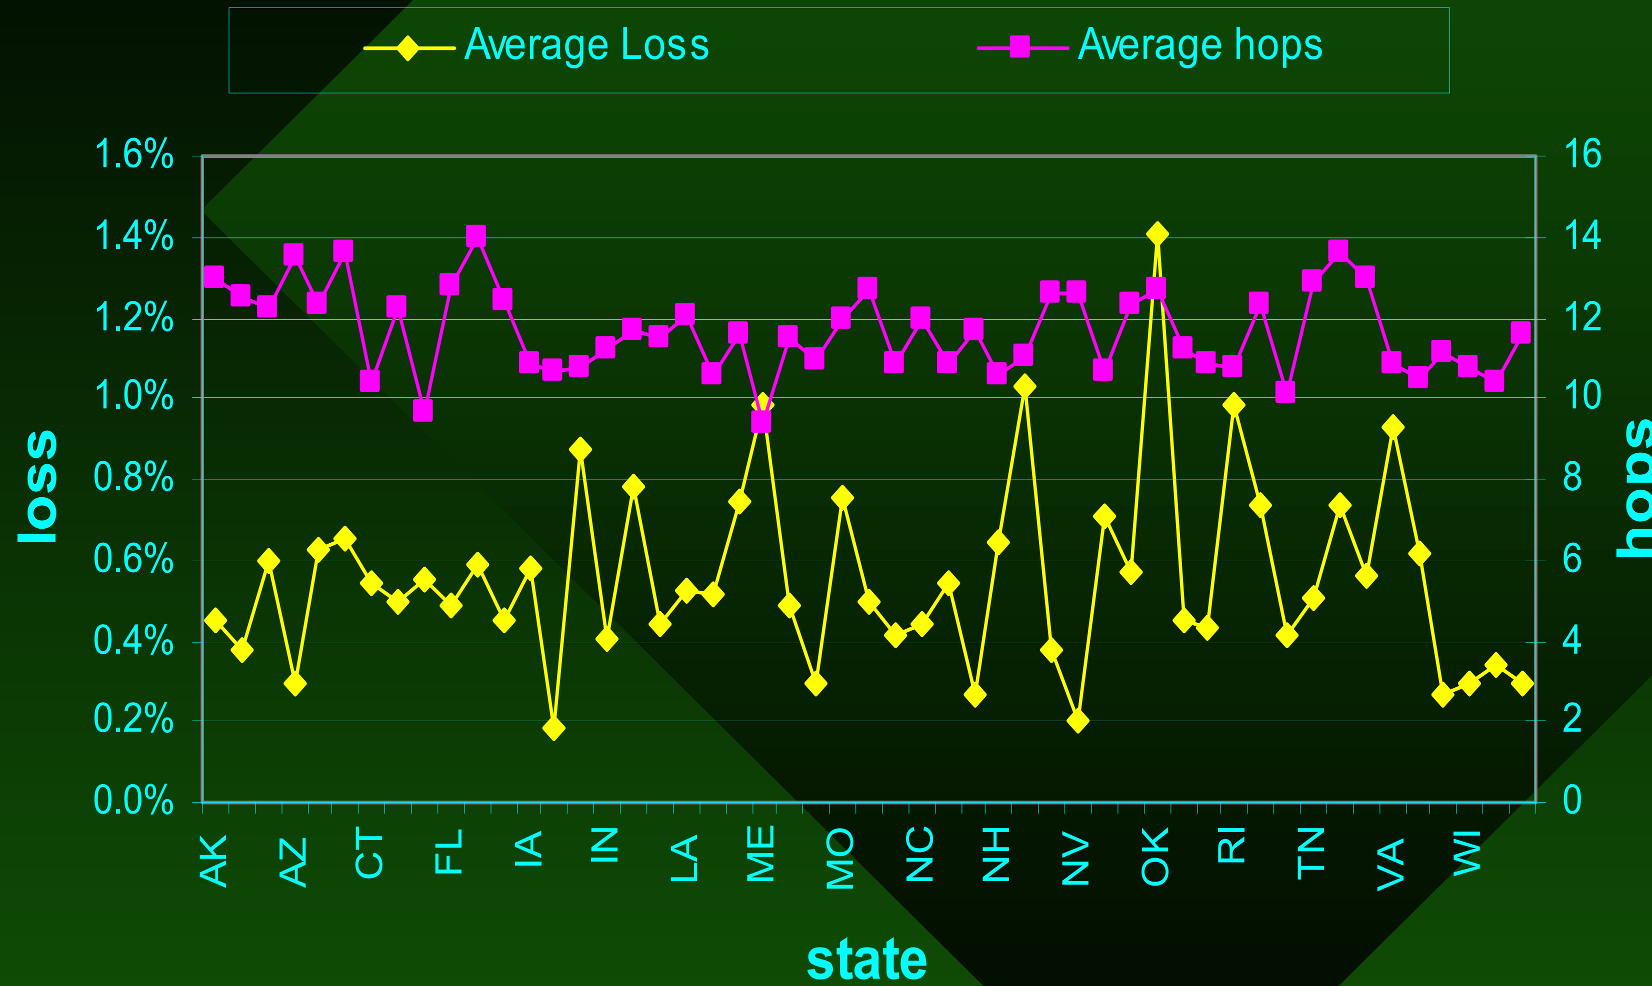

Packet Loss (cont'd)

- 207,384 loss bursts and 431,501 lost packets
- Loss burst lengths PDF:

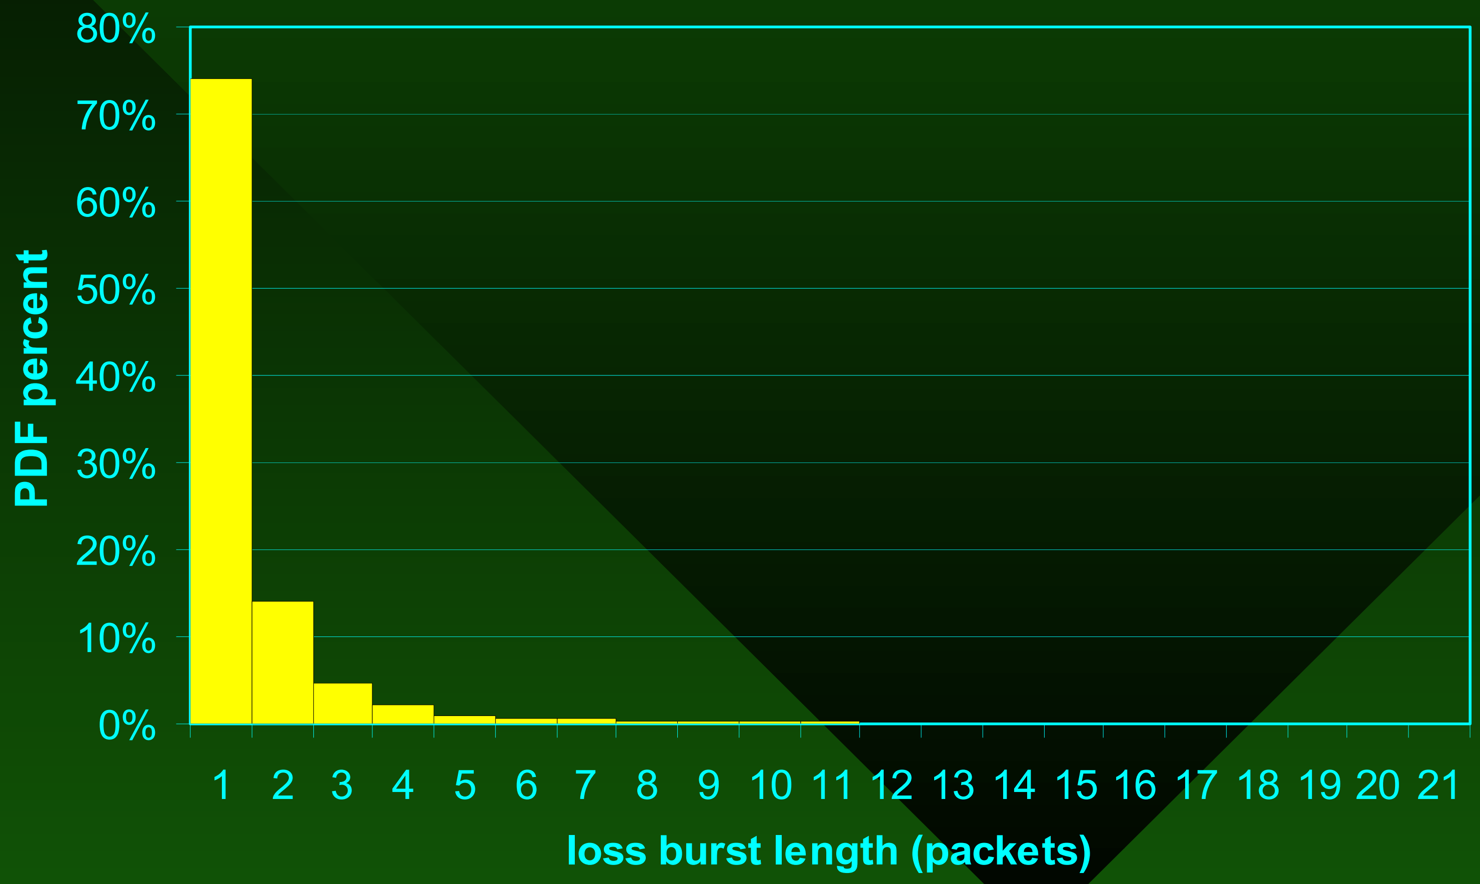

Packet Loss (cont'd)

- Most bursts contained no more than 7 packets (however, the tail reached to over 100 packets)
- RED was disabled on the backbone; still 74% of loss bursts contained only 1 packet apparently dropped in FIFO queues
- Average burst length was 2.04 packets in D_{1p} and 2.10 packets in D_{2p}
- Conditional probability of packet loss was 51% and 53%, respectively
- Over 90% of loss burst *durations* were under 1 second (maximum 36 seconds)
- The average distance between lost packets was 21 and 27 seconds in D_{1p} and D_{2p}, respectively

Packet Loss (cont'd)

- Apparently heavy-tailed distributions of loss burst lengths
- Pareto with α = 1.34; however, note that data was non-stationary (time of day or access point non-stationarity)

Underflow events

- Missing packets (frames) at their decoding deadlines cause *buffer underflows* at the receiver
- Startup delay used in the experiment was 2.7 seconds
- 63% (271,788) of all lost packets were discovered to be missing before their deadlines
- Out of these 63% of lost packets:
  - 94% were recovered in time
  - 3.3% were recovered late
  - 2.1% were never recovered
- Retransmission appears quite effective in dealing with packet loss, even in the presence of large end-to-end delays

Underflow events (cont'd)

- 37% (159,713) of lost packets were discovered to be missing after their deadlines had passed
- This effect was caused by large one-way delay jitter
- Additionally, one-way delay jitter caused 1,167,979 *data* packets to be late for decoding
- Overall, 1,342,415 packets were late (1.7% of all sent packets), out of which 98.9% were late due to large one-way delay jitter rather than due to packet loss combined with large RTT
- All late packets caused the "freeze-frame" effect for 10.5 seconds on average in D_{1p} and 8.5 seconds in D_{2p} (recall that each session was 10 minutes long)

Underflow events (cont'd)

- 90% of late *retransmissions* missed the deadline by no more than 5 seconds, 99% by no more than 10 seconds
- 90% of late *data* packets missed the deadline by no more than 13 seconds, 99% by no more than 27 seconds

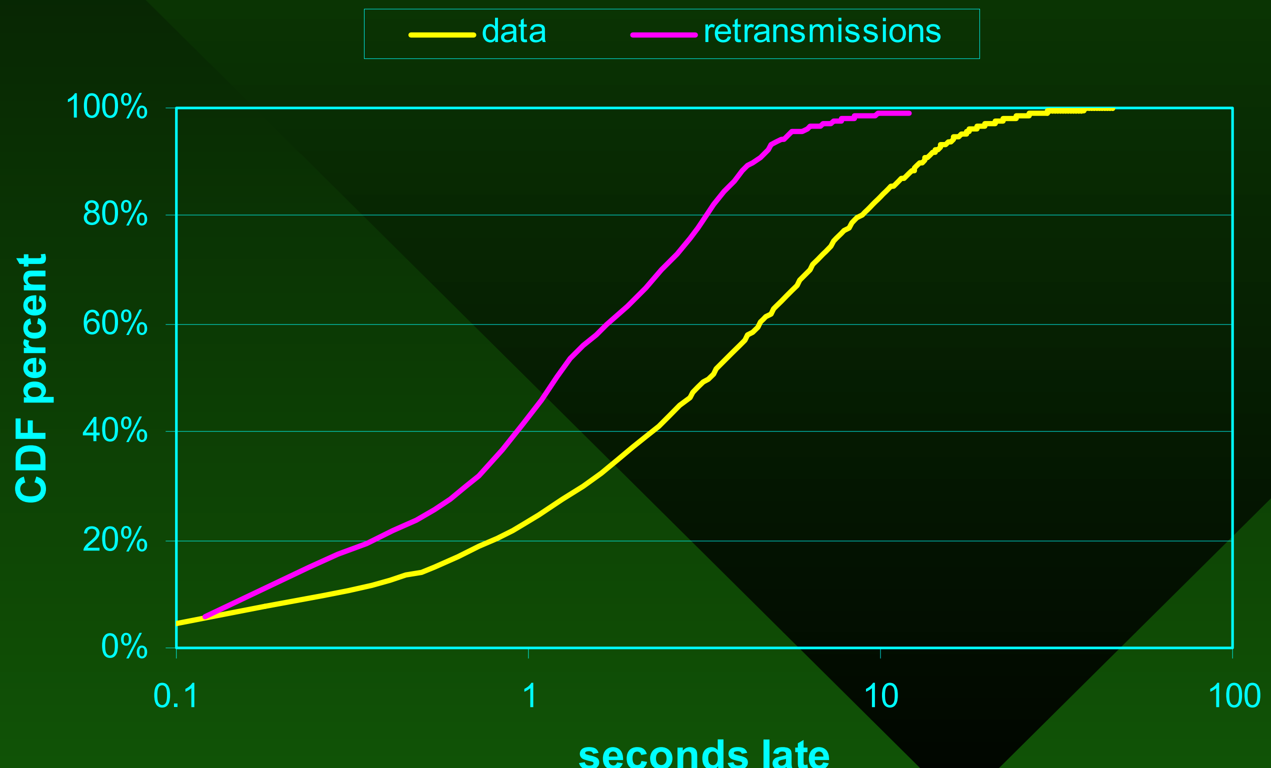

Round-trip Delay

- 660,439 RTT samples
- 75% of samples below 600 ms and 90% below 1 second
- Average RTT was 698 ms in D_{1p} and 839 ms in D_{2p}
- Maximum RTT was over 120 seconds
- Data-link retransmission combined with low-bitrate connection were responsible for pathologically high RTTs
- However, we found access points with 6-7 second IP-level buffering delays

Round-trip Delay (cont'd)

- Distributions of the RTT in both datasets (PDF) were similar and contained a very long tail

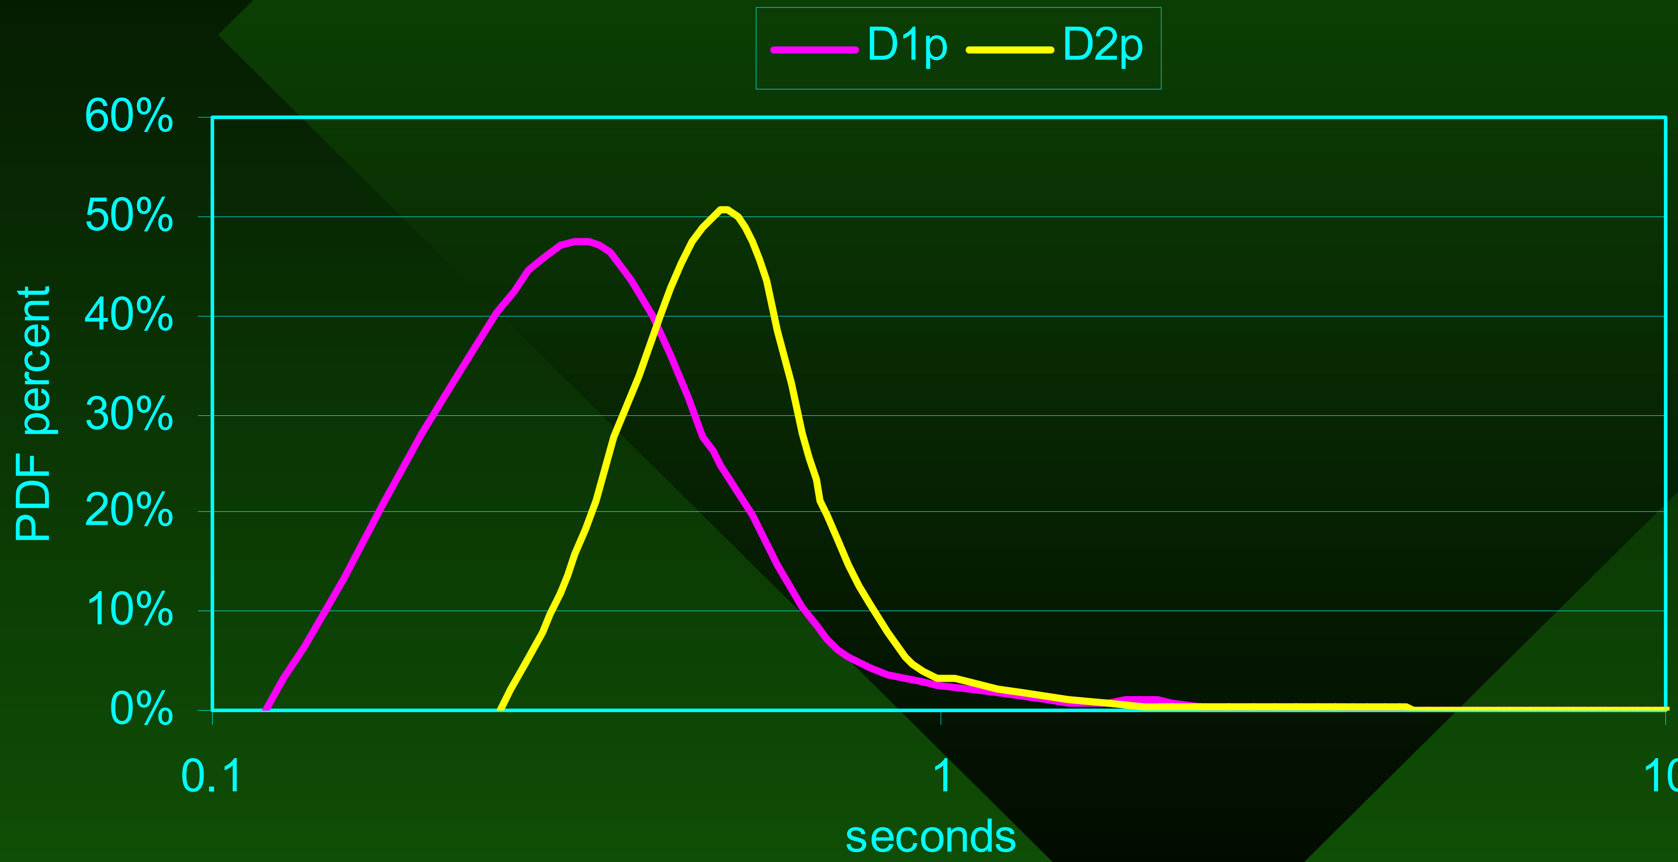

Round-trip Delay (cont'd)

- Distribution tails closely matched hyperbolic distributions (Pareto with α between 1.16 and 1.58)

Round-trip Delay (cont'd)

- The average RTT varied during the day between 574 ms (3 am – 6 am) and 847 ms (3 pm – 6 pm) in D_{1p}
- Between 723 ms and 951 ms in D_{2p}
- Relatively small increase in the RTT during the day (by only 30-45%) compared to that in packet loss (by up to 300%)
- Per-state RTT varied between 539 ms (Maine) and 1,053 ms (Alaska); Hawaii and New Mexico also had average RTTs above 1 second
- Little correlation between the RTT and geographical distance of the state from NY
- However, much stronger positive correlation between the number of hops and the average state RTT: $\rho = 0.52$

Packet Reordering

- Average reordering rates were low, but noticeable
- 6.5% of missing packets (or 0.04% of sent) were reordered
- Out of 16,852 sessions, 1,599 (9.5%) experienced at least one reordering event
- The highest reordering rate *per ISP* occurred in AT&T WorldNet, where 35% of missing packets (0.2% of sent packets) were reordered
- In the same set, almost half of the sessions (47%) experienced at least one reordering event
- Earthlink had a session where 7.5% of *sent* packets were reordered

Packet Reordering (cont'd)

- Reordering delay D_r is time between detecting a missing packet and receiving the reordered packet
- 90% of samples D_r below 150 ms, 97% below 300 ms, 99% below 500 ms, and the maximum sample was 20 seconds

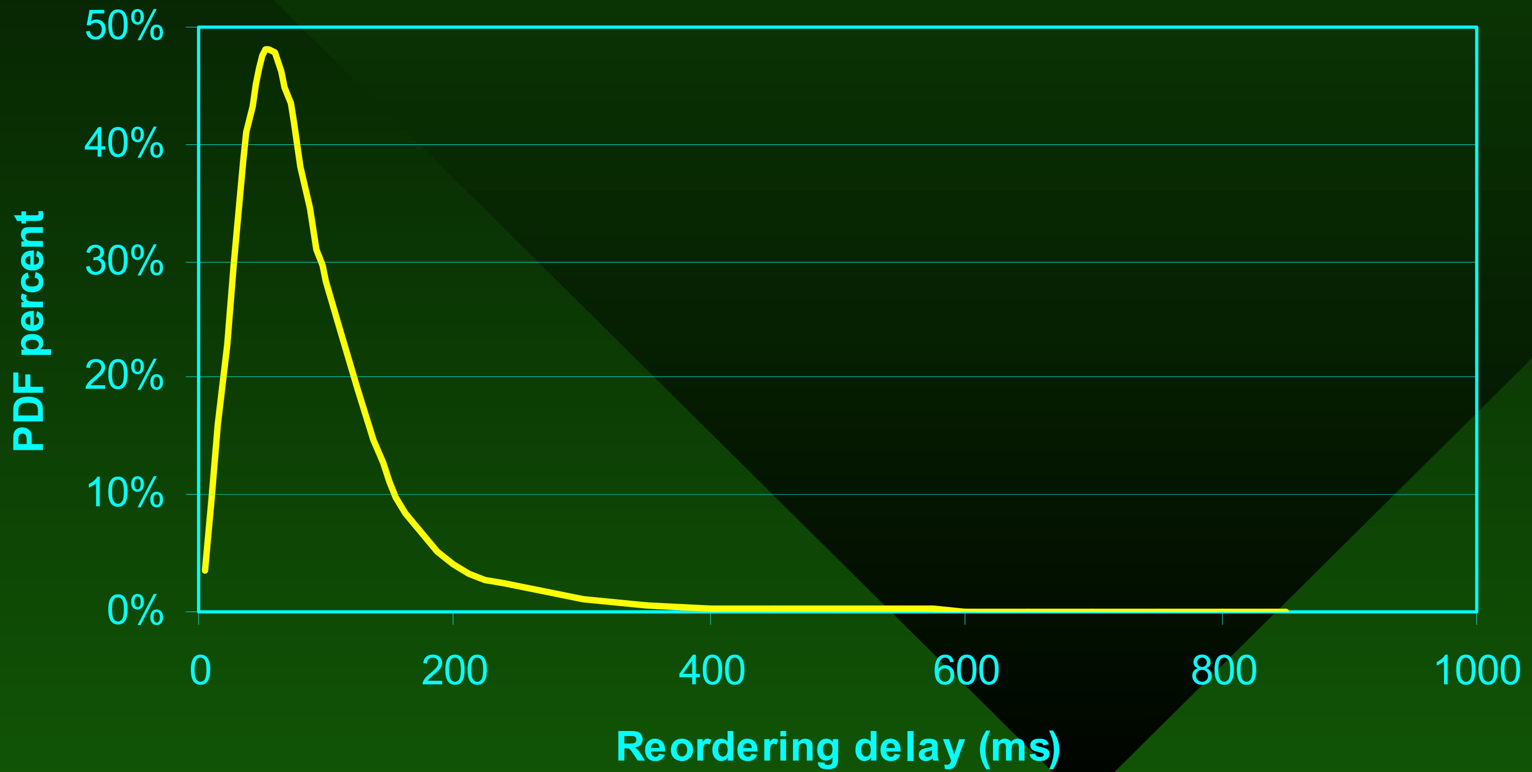

Packet Reordering (cont'd)

- Reordering distance is the number of packets received during the reordering delay (84.6% of the time a single packet, 6.5% exactly 2 packets, 4.5% exactly 3 packets)
- TCP's triple-ACK avoids 91.1% of redundant retransmits and quadruple-ACK avoids 95.7%

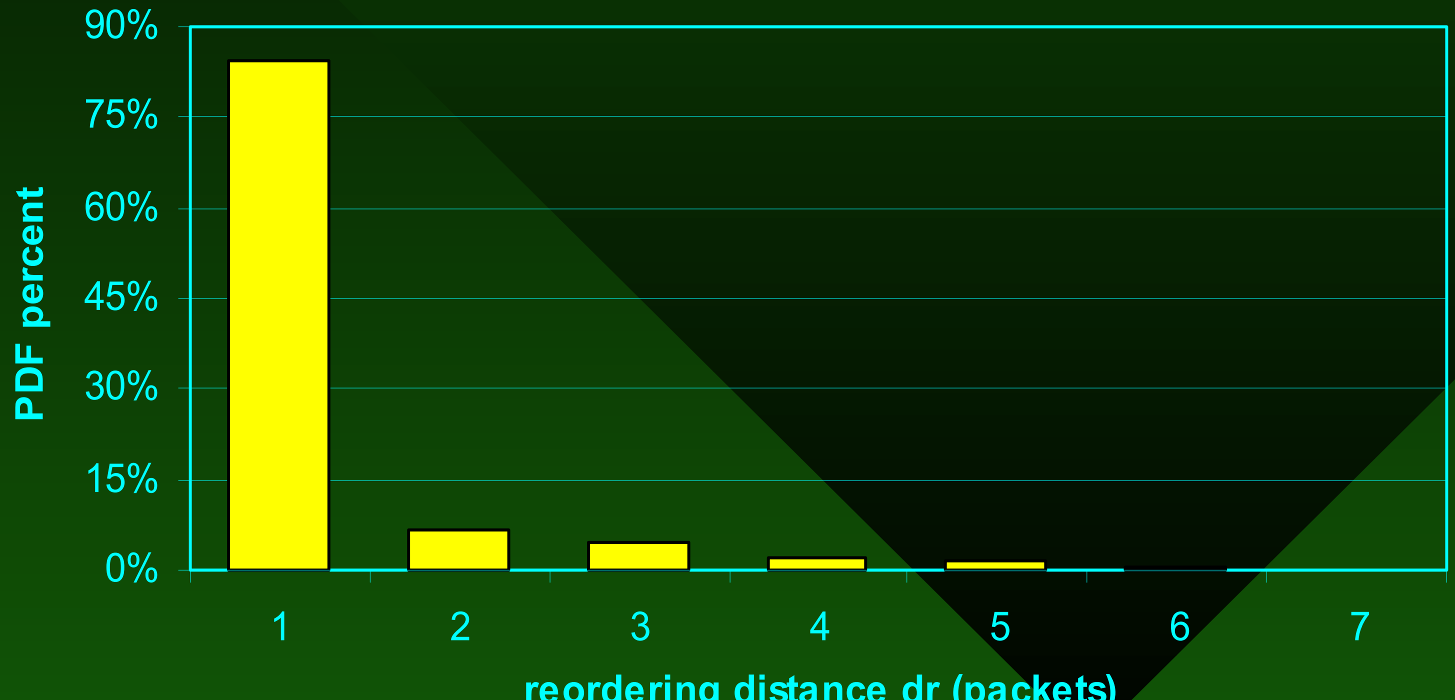

Path Asymmetry

- Asymmetry detected by analyzing the TTL of the returned packets during the initial traceroute
- Each router reset the TTL to a default value (such as 255) when sending a “TTL expired” ICMP message
- If the number of forward and reverse hops was different, the path was “definitely asymmetric”
- Otherwise, the path was “possibly (or probably) symmetric”
- No fail-proof way of establishing path *symmetry* using end-to-end measurements (even using two traceroutes in reverse directions)

Path Asymmetry (cont'd)

- 72% of sessions operated over <u>definitely</u> asymmetric paths
- Almost all paths with 14 or more end-to-end hops were asymmetric
- Even the shortest paths (with as low as 6 hops) were prone to asymmetry
- "Hot potato" routing is more likely to cause asymmetry in longer paths, because they are more likely to cross AS borders than shorter paths
- Longer paths also exhibited a higher reordering probability than shorter paths

Conclusion

- Dialing success rates were quite low during the day (as low as 40%)
- Retransmission worked very well even for delay-sensitive traffic and high-latency end-to-end paths
- Both RTT and packet-loss bursts appeared to be heavy-tailed
- Our clients experienced huge end-to-end delays both due to large IP buffers as well as persistent data-link retransmission
- Reordering was fairly frequent even given our low bitrates
- Most paths were in fact asymmetric, where longer paths were more likely to be identified as asymmetric