Delft University of Technology

Reordering of IP Packets in Internet

Xiaoming Zhou and Piet Van Mieghem

Network Architectures and Services
TU Delft

January 29, 2004

Overview

- Introduction to Reordering
- Problem description and definitions
- Experiment results
- Conclusion

Introduction

- Existence of packet reordering (out-of-order arrival of packets at the destination)
- Main reason:
  - The Parallelism in Internet components (switches) and links

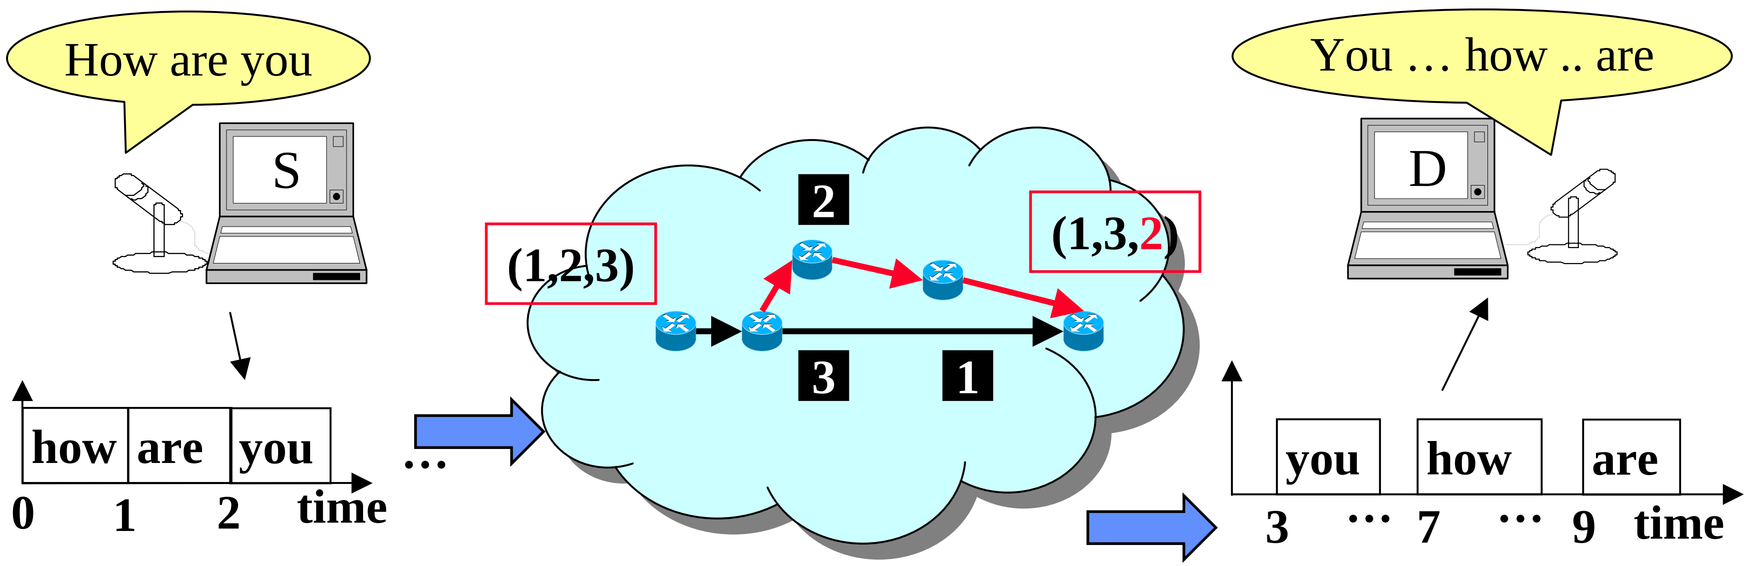

Introduction (con.)

- Amount of reordering is a function of:
  - Network load
  - Configuration of the hardware (i.e., multiple switches in a router) and software (i.e., class-based scheduling or priority queueing) in the routers
- Motivation
  - Reordering greatly impacts the performance of applications in the Internet
  - These reordering measurements may shed light on the underlying properties of the current topology
- Purpose:
  - Understanding the nature of reordering
- UDP 1-way measurement

Overview

- Introduction to Reordering
- Problem description and definitions
- Experiment results
- Conclusion

Problem Description and Definitions

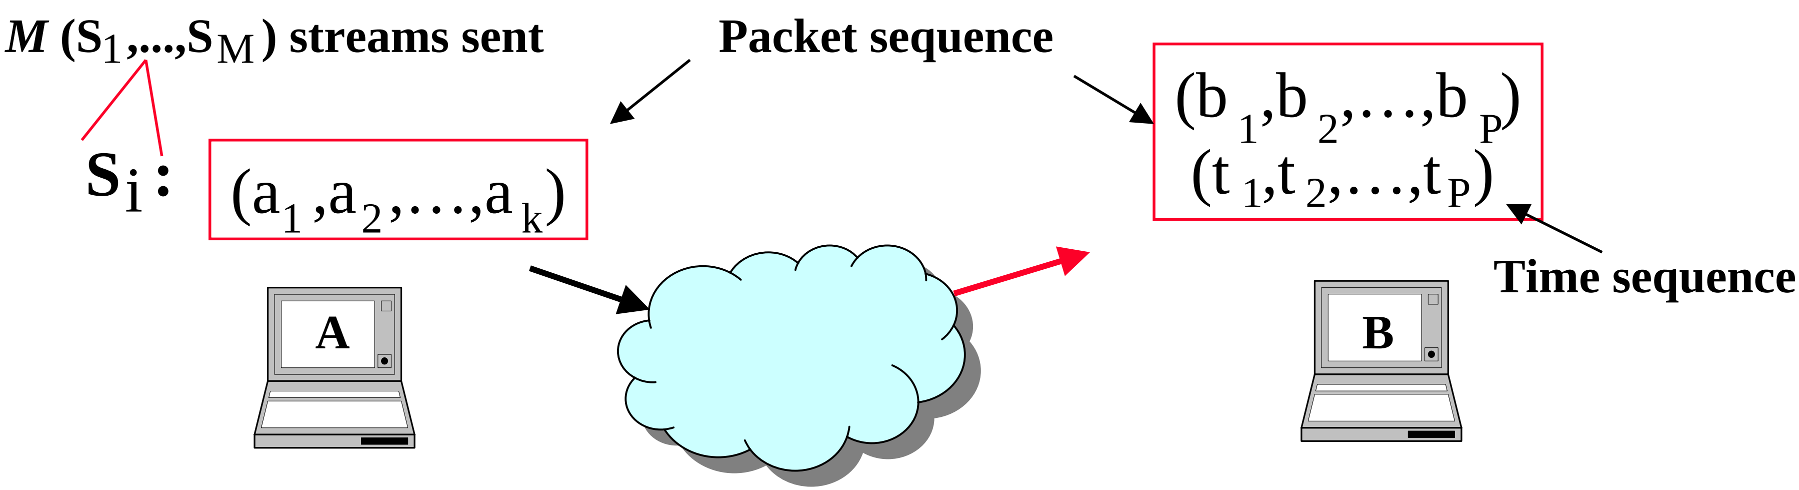

- Reordered stream and reordered packet

(1,2,3,5,4,7,6)

Reordered packets (L=2)

- Reordered packet lengths L is defined as the total number of reordered packets in an arrival stream

Definitions (con.)

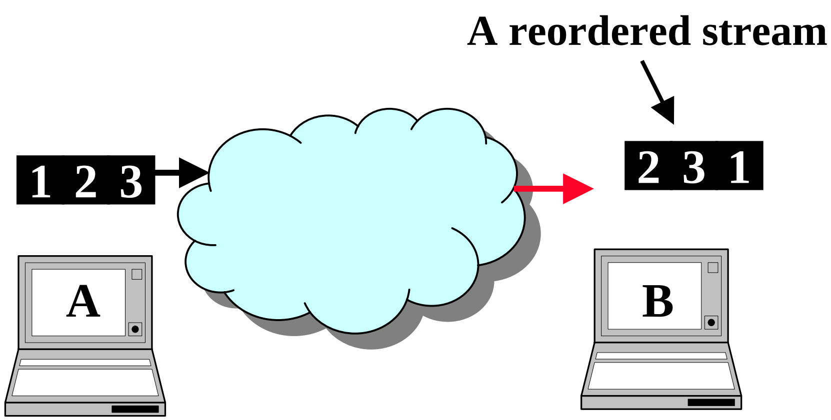

M streams sent

M_r streams

M_a reordered streams

- Reordered stream ratio

$$R_{AB} = \frac{M_r}{M_a}$$

- Asymmetric nature

$$|R_{AB} - R_{BA}|$$

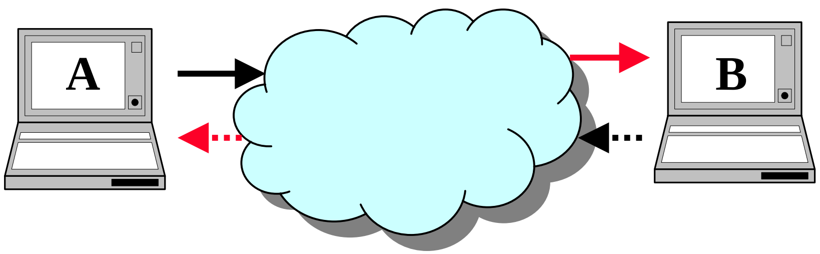

Definitions (con.)

- Question: How to predict whether a reordered packet will be useful in a receiver with a finite buffer?
  - VoIP
  - The reordered packet length does not give sufficient information

- The packet lag P_L

(2,1,3,4,5,6,7,8) ← $P_L = 1$

(2,3,4,5,6,7,8,1) ← $P_L = 7$

P_L refers to the number of packets, with sequence numbers greater than the reordered packet that are received before the reordered packet itself.

Delft University of Technology

Definition (3)

- The time lag T_L

Packet sequence

(1,2,3,4,5,6,7,8)

(2,3,4,5,6,7,8,1)

$(t_2, t_3, t_4, t_5, t_6, t_7, t_8, t_1)$

Time sequence

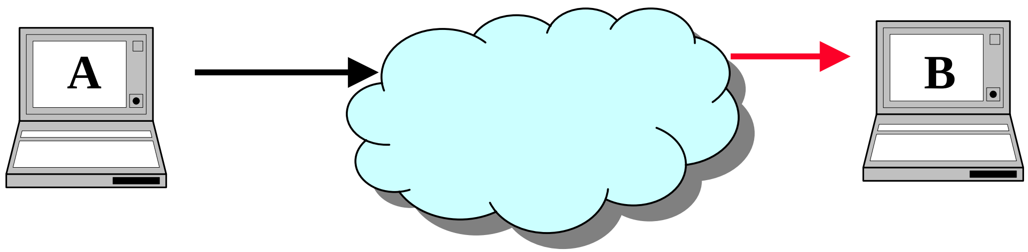

T_L **is defined as the difference between the delay t_k of the reordered packet k and its expected delay t'_k without reordering**

$$T_L = |t_k - t'_k|$$

$$= |t_k - \min(t_1, \ldots, t_P)|$$

Definition (4)

- Research question
  - How to investigate to which extent packets are reordered with respect to a sample of packets?
- Methodology (these functions are measured)
  - Reordered streams ratio
  - The asymmetric nature
  - pdf of reordered packet lengths L
  - The packet lag P_L and time lag T_L

Overview

- Introduction to Reordering
- Problem description and definitions
- Experiment results
- Conclusion

Delft University of Technology

RIPE measurement configuration

- The traces and delay data measured in 12 test-boxes of RIPE TTM project

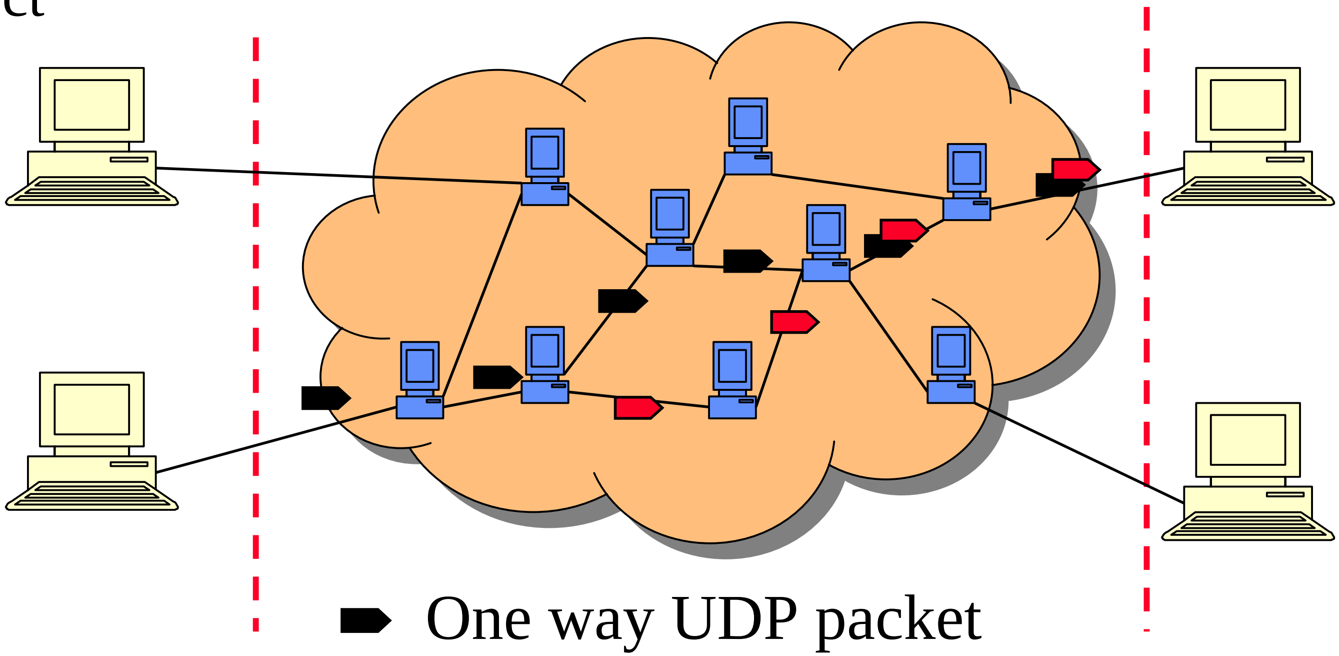

- A delay accuracy within 10 microseconds
- 12 test-boxes: 3 hosts are located in NL, 2 in GB, 1 in Sweden, Slovakia, Belgium, Australia, USA, Denmark and Greece

Two Experiments

- First send 50 100-byte UDP packets (N50)
- Tests were run from 5 to 8 PM on October 16, 2003
- Second send 100 100-byte UDP packets (N100)
- Tests were run from 5 to 8 PM on October 17, 2003
- Why 3 hours?
- The experiment consisted of 104 unidirectional paths
- Reordering does not correlate with loss

Results of reordered stream ratio

- Aim: Reordered stream ratio gives insight how often reordering happened in the probe-streams
- Principle : $R_{AB} = \dfrac{M_r}{M_a}$

Received data	N_{50}	N_{100}
UDP streams	36762	32691
Reordered stream	20445 (56%)	21649 (60%)
UDP packets	1655120	2828834
Reordered UDP packets	101018 (6%)	158413(5.6%)
Measurement duration	3 hours	3 hours

Table: Details of the packets used to measure the reordering on 104 paths

- Reordering quite often occurs in the probe-streams

Results of Reordered packet length L

Delft University of Technology

- Aim: quantify the extent of reordering
- How many reordered packets in each arrival stream?

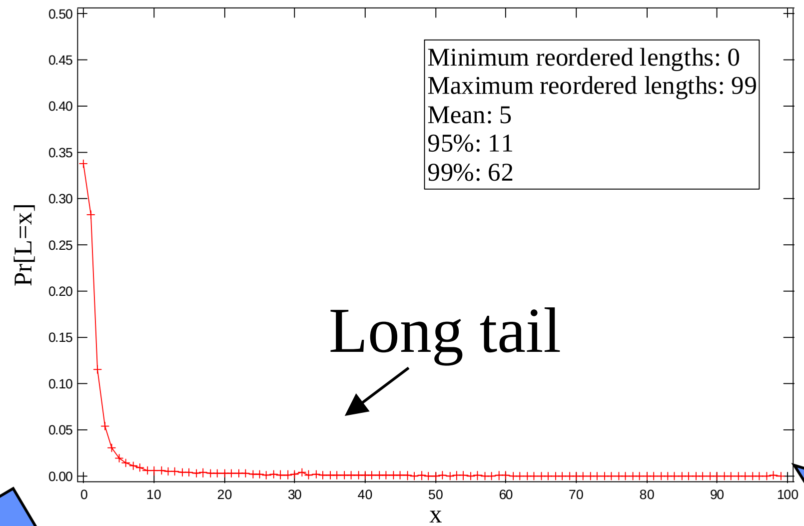

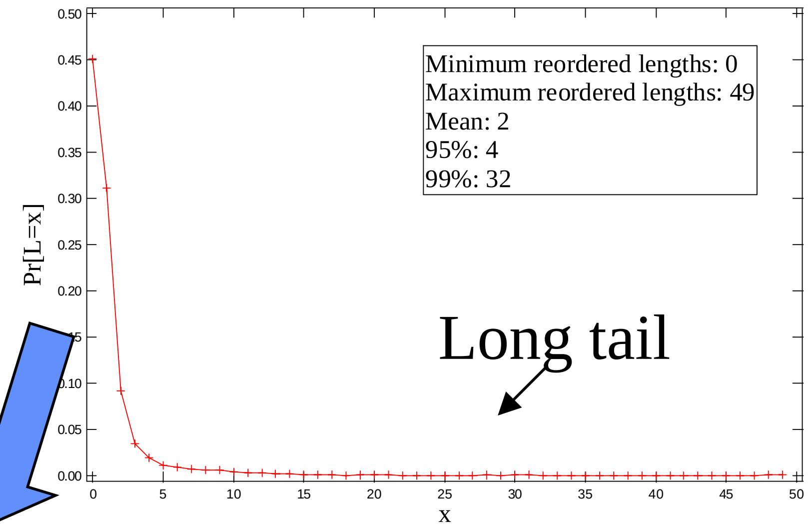

Figure: The pdf of the reordered packet lengths for the 2 data sets

	L=0	L=1	L=2	others
N_{50}	44%	32%	10%	14%
N_{100}	34%	28%	12%	26%

- Most individual streams have a relatively small number of reordering events

Results of Reordered packet length (2)

Delft University of Technology

- Fitting the pdf of *L* on a log-log scale seems to indicate power law behavior for *L*
- Power law: Pr[*L*=x]=C x $^{-b}$

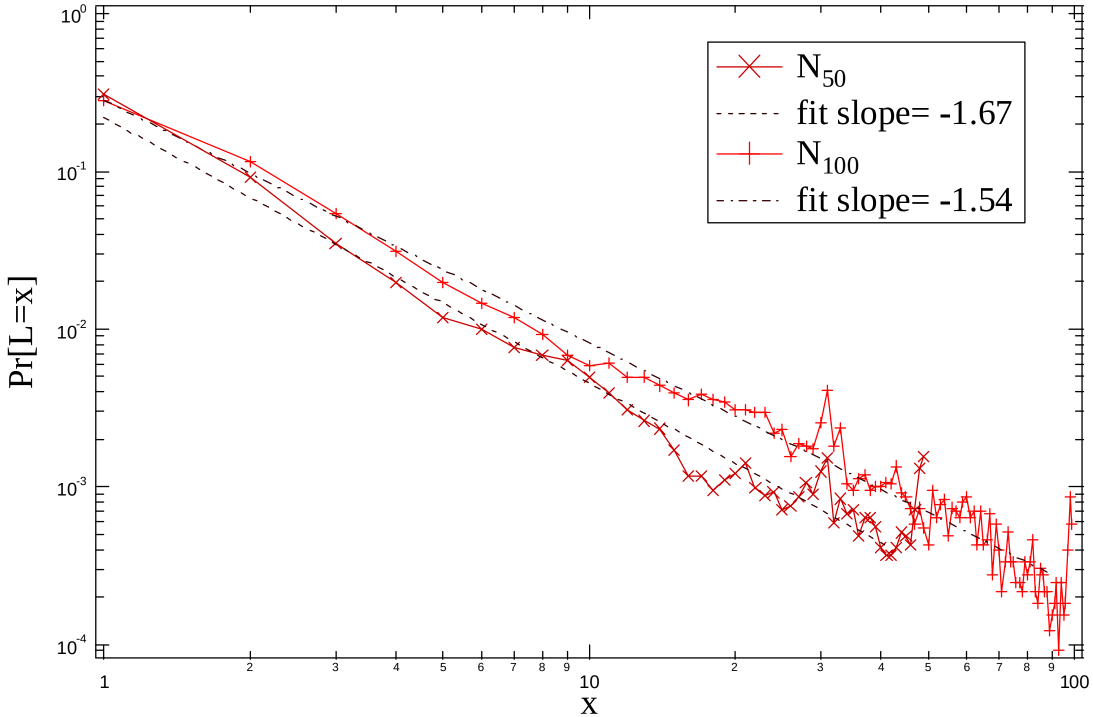

Figure: The pdf of reordered packet lengths *L* and the power law fit

Results of the packet lag P_L

- The packet lag P_L

Aim: Help to predict whether a reordered packet will be useful in a receiver buffer with finite limit

Principle: How many packets with greater sequence numbers have been received before the reordered packet

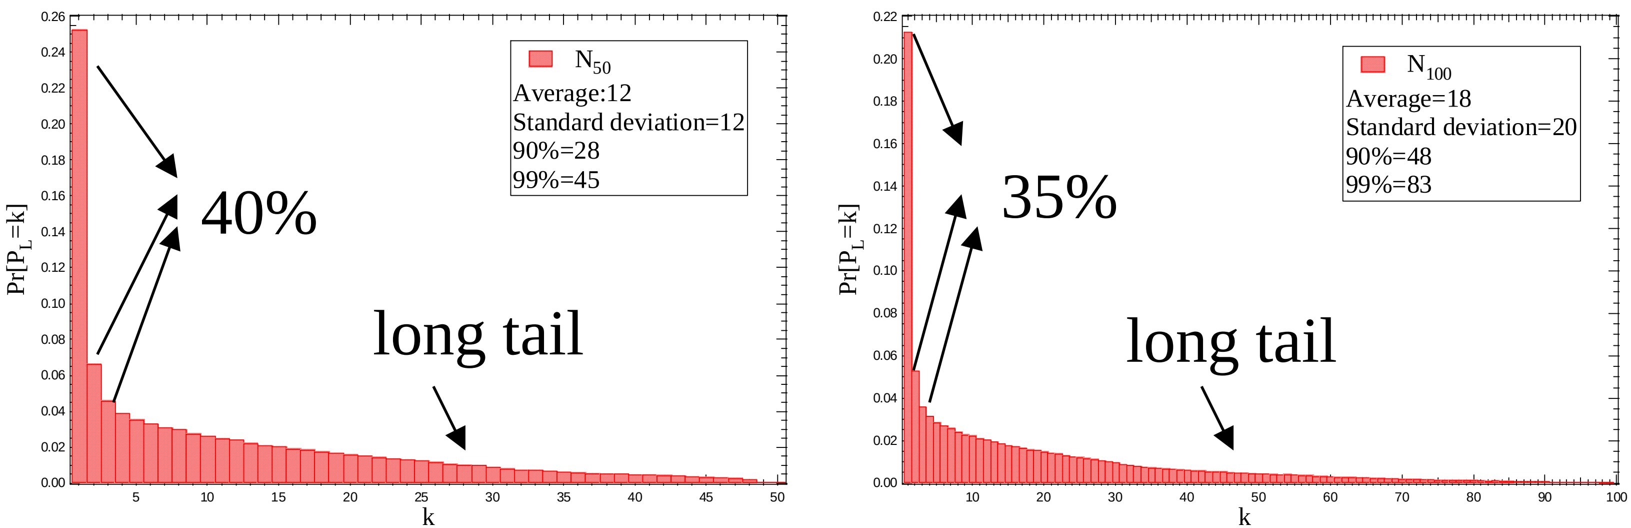

Figure: The pdf of packet lag for 2 data sets

Packet reordering has a significant impact on UDP performance since reordering adds a high cost for recovering from the reordering on the end host

Results of the packet lag P_L (2)

- Fitting the pdf of P_L on a log-lin scale seems to indicate exponential distribution for P_L
- Exponential distribution: $\Pr[P_L = k] = a\, e^{-ak}$

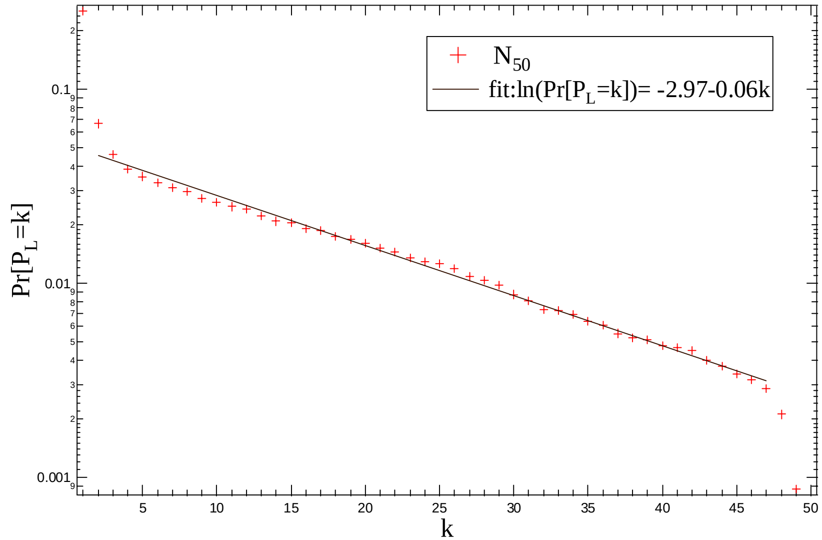

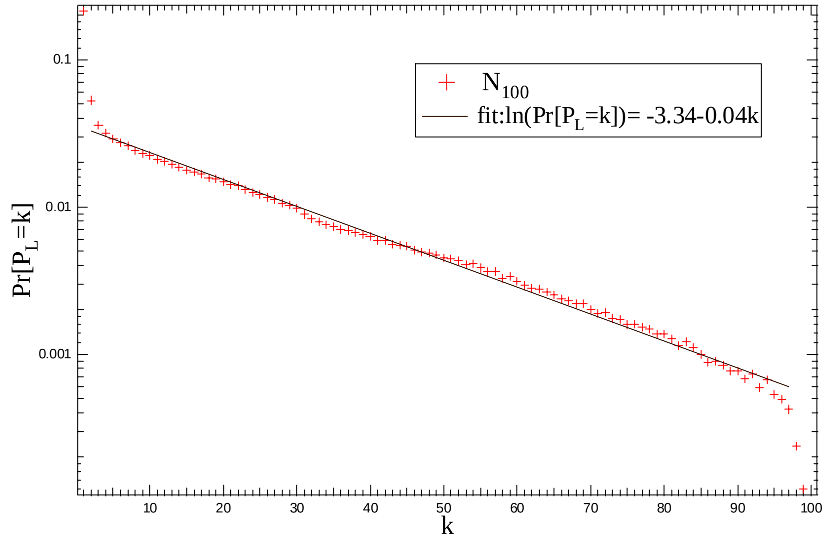

Figure: The pdf of packet lag and the exponential fit for 2 data sets

TU Delft
Delft University of Technology

Results of reordered time lag

- Aim: Time lag is a delay-based metric to more precisely evaluate the impact of reordered packets on the end hosts
- Principle: $$T = \frac{|t_k - \min(t_1, \ldots, t_P)|}{\min(t_1, \ldots, t_P)}$$

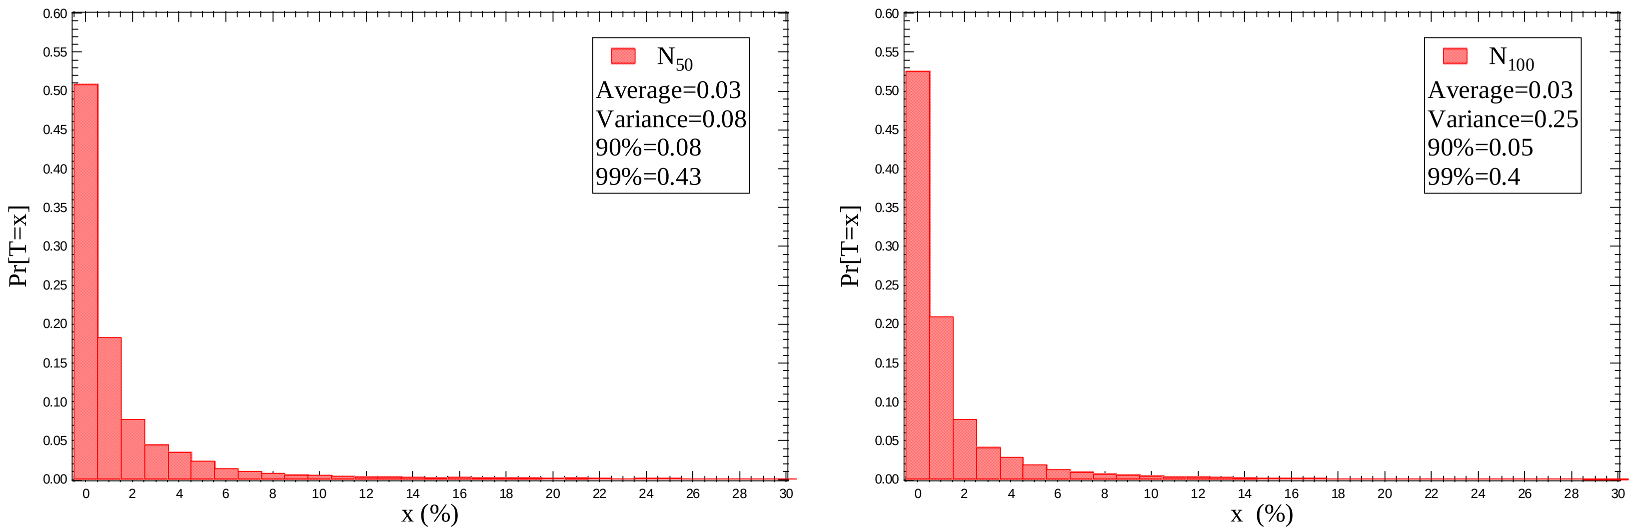

Figure: The pdf of normalized time lag for 2 data sets

Packet reordering does not have a significant impact on UDP delay since reordering does not add large delay on the end host

Results of Asymmetry nature

- Omit pairs for which the probe-streams in one of the direction were missing, leaving the data from in total 39 pairs
- Degree of asymmetry of reordering (DAR): $DAR=\frac{|R_{AB}-R_{BA}|}{\min(R_{AB},R_{BA})}$

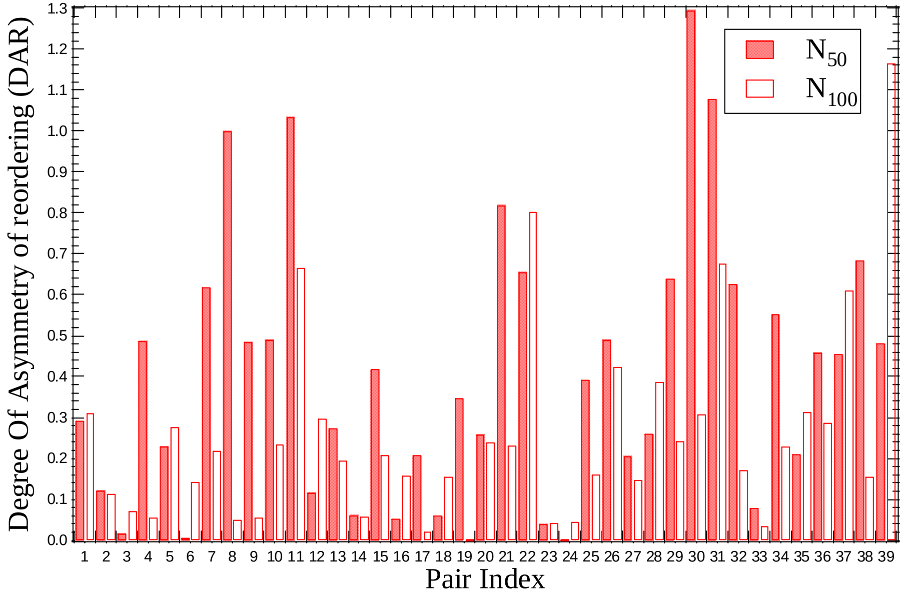

Figure: Degree of Asymmetry of Reordered streams in all 39 symmetric traces

The asymmetry of reordered streams ratios exists on all experiment pairs, but it varies greatly from testbox-to-testbox

- Routing policies

Conclusion

- Reordering is a frequent phenomenon in Internet
- Most individual streams have a relatively small number of reordering events
- Packet reordering has a significant impact on UDP performance but it does not add a large delay on the end hosts
- The asymmetry of reordered streams ratios exists on all experiment pairs, but it varies greatly from testbox-to-testbox