Measurement Study of Wuala, a Distributed Social Storage Service

Thomas Mager - Master Thesis

TECHNISCHE UNIVERSITÄT DARMSTADT

Advisors:
Prof. Ernst Biersack
Prof. Thorsten Strufe
Prof. Pietro Michiardi

EURECOM

Agenda

TECHNISCHE UNIVERSITÄT DARMSTADT

1. Introduction
2. Motivation
3. Basics
4. General Experimental Setup
5. Experiments
6. Conclusion

EURECOM Sophia Antipolis

Introduction

TECHNISCHE UNIVERSITÄT DARMSTADT

What is Wuala?

- On-line storage solution
- Leverages resources of peers to store files
- Development started at ETH Zürich 2004
- Company founded 2007
- Launched in 2008
- Acquired by LaCie 2009
- Pioneer in this area - first widely-used system!

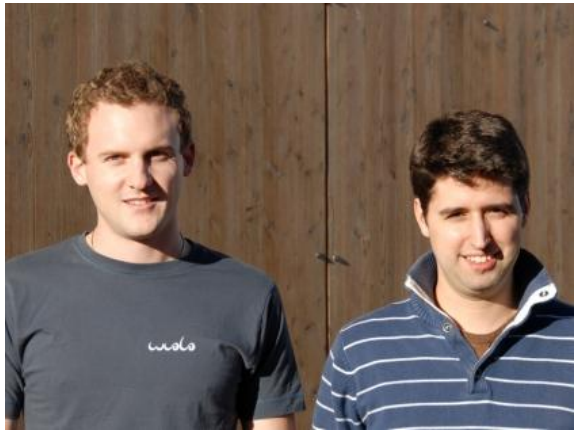

Dominik Grolimund, Luzius Meisser

ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

wuala
by LACIE

EURECOM

Features

TECHNISCHE UNIVERSITÄT DARMSTADT

- 1 GiB storage for free, else trade local storage or pay
- Automatic backup
- Share files with friends
- All files encrypted locally
- Mount on-line storage into file system hierarchy
- File versioning

EURECOM

Wuala Application

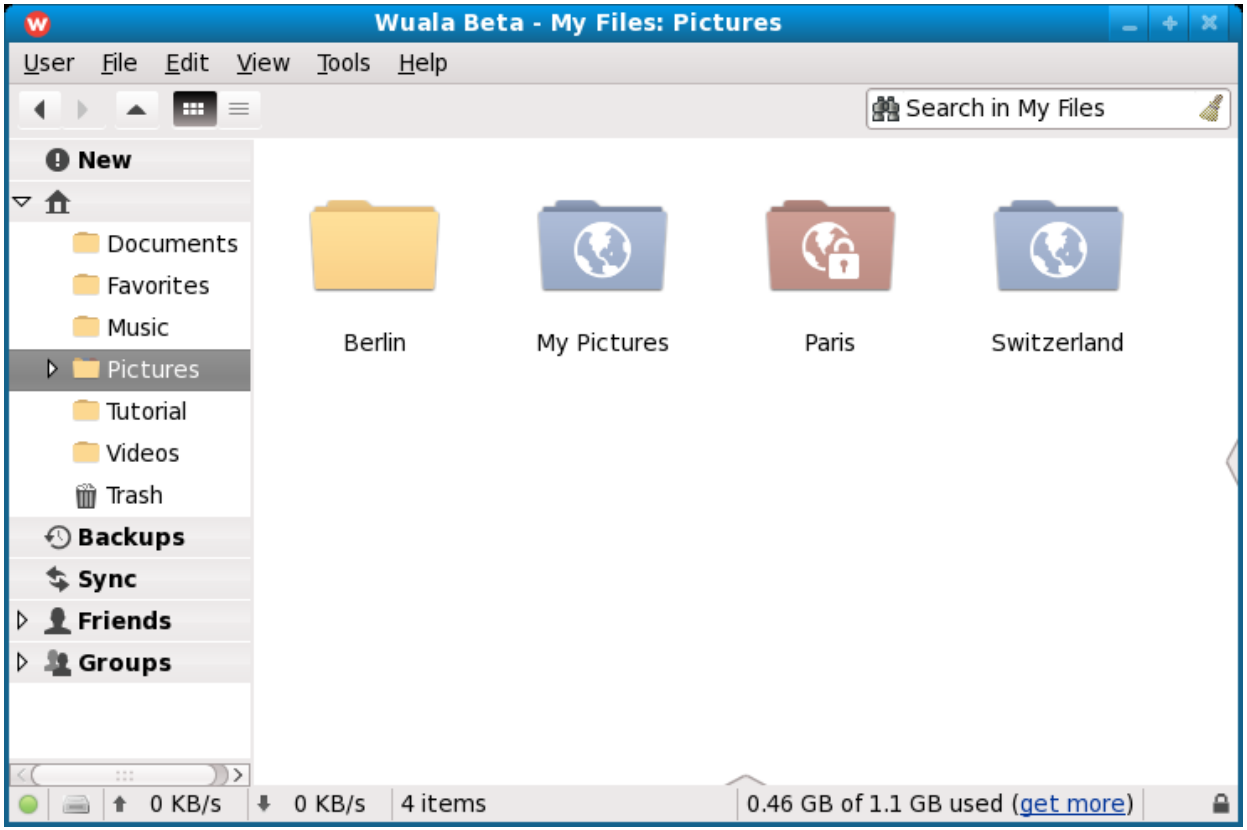

Wuala Application

General Architecture

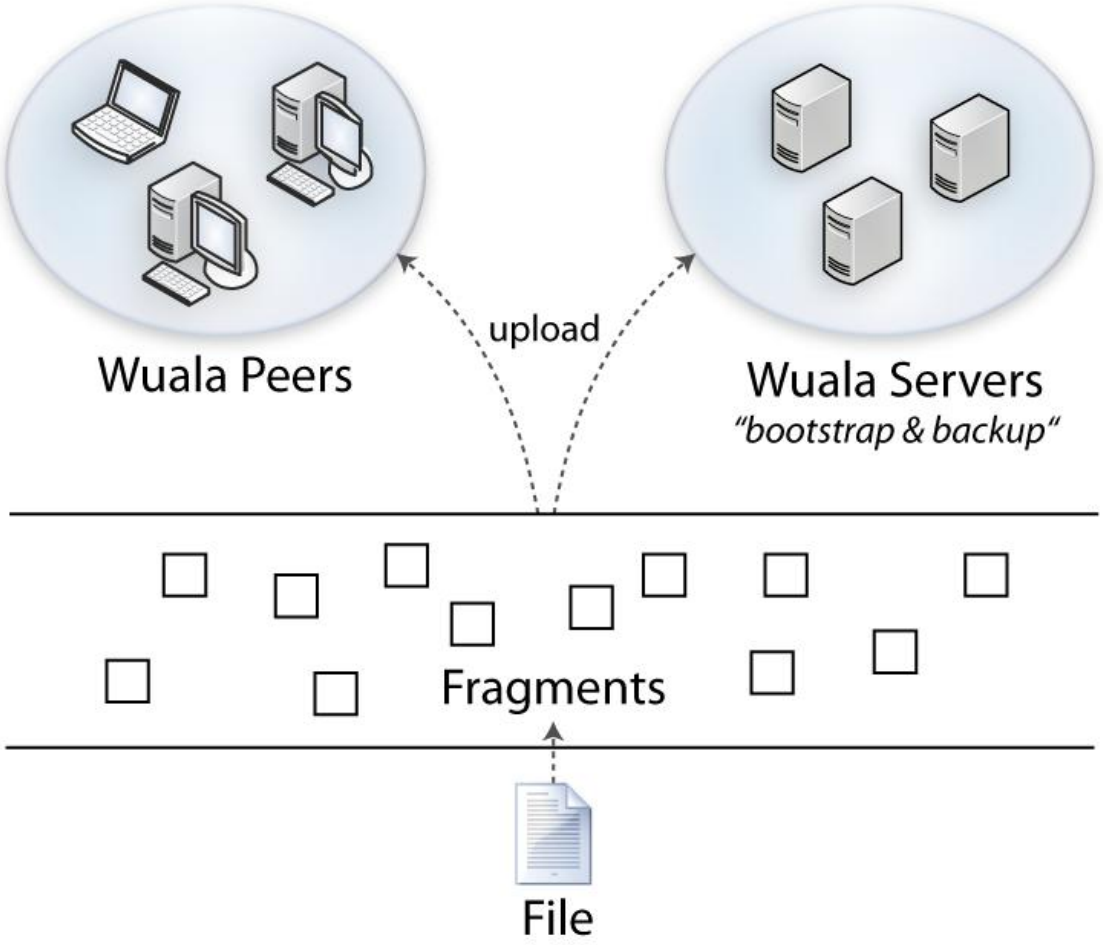

Motivation

TECHNISCHE UNIVERSITÄT DARMSTADT

How to make sure that files are always accessible?

→ Crucial because of high churn rate

Requirements:

- Availability
- Reliability
- Consistency

How do peers / servers contribute to achieve this?

EURECOM

Redundancy Required

TECHNISCHE UNIVERSITÄT DARMSTADT

Two main approaches:

- Replication
- Erasure Coding

EURECOM

Replication

$$A_{replication} = 1 - (1 - p)^r$$

p = node availability

r = redundancy factor

Replication

TECHNISCHE UNIVERSITÄT DARMSTADT

$$A_{replication} = 1 - (1 - p)^r$$

p = node availability

r = redundancy factor

Availability target of 99.9%

presuming a mean node availability of 50%:

r = 10

EURECOM

Erasure Coding

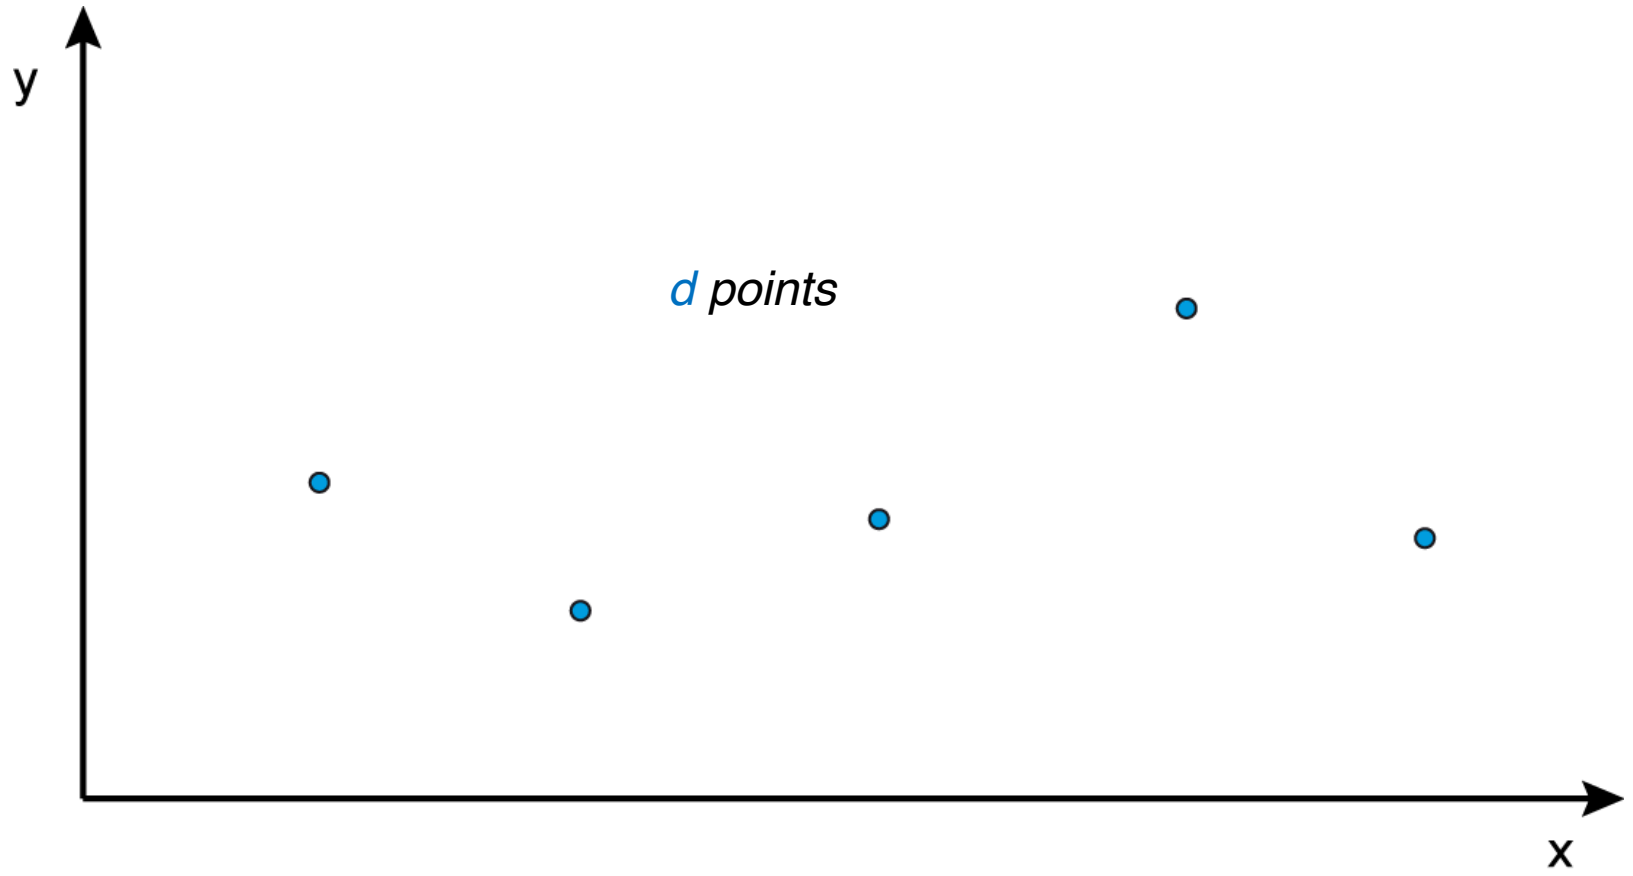

Erasure Coding

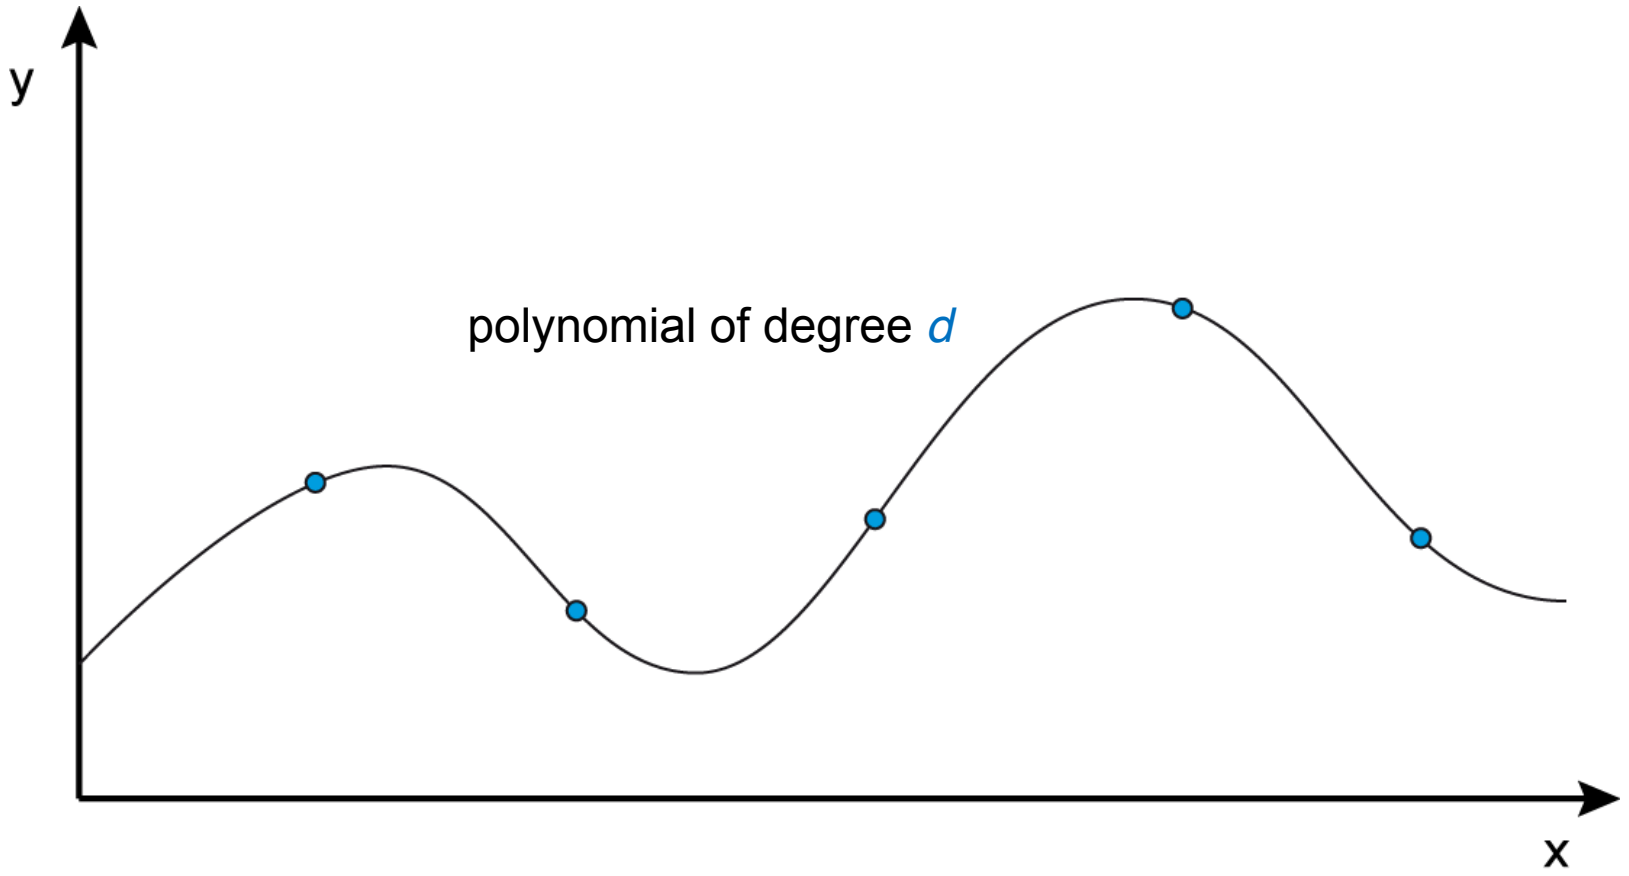

Erasure Coding

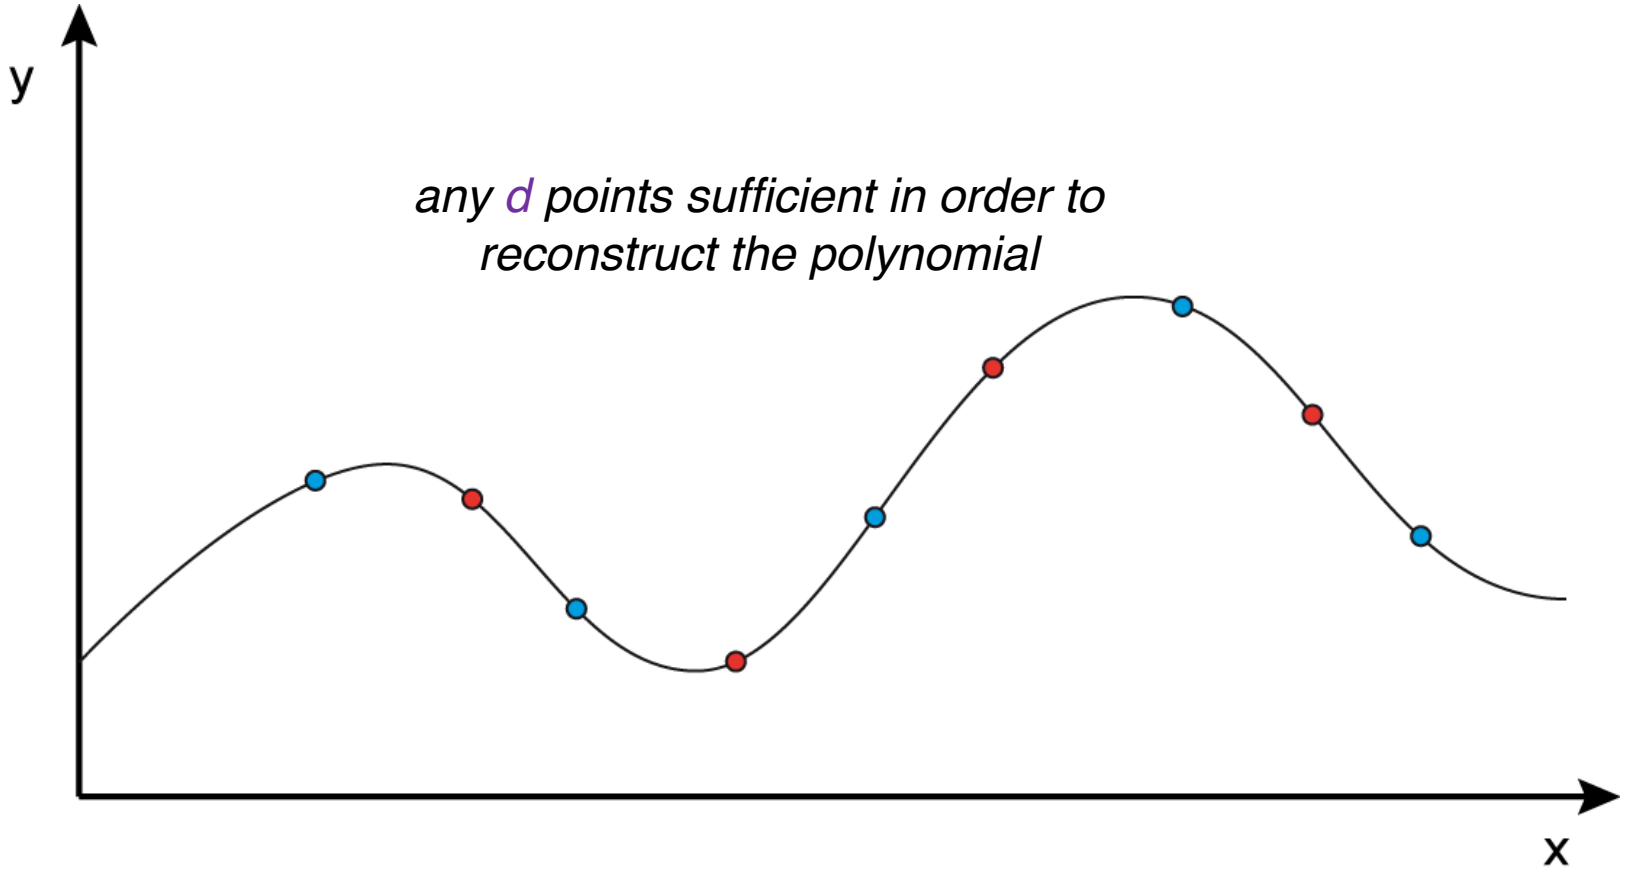

Erasure Coding

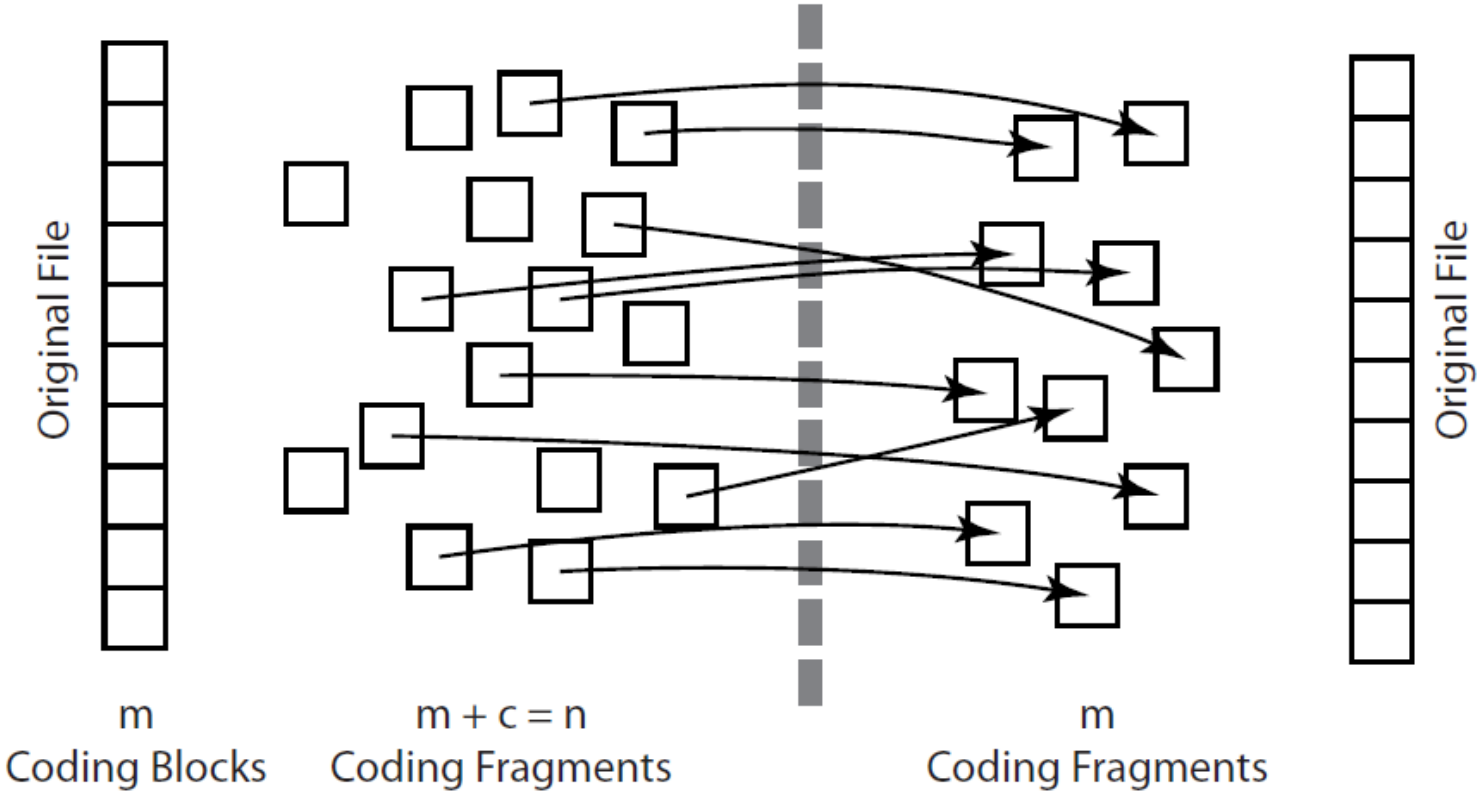

Erasure Coding

$$A_{erasurecoding} = \sum_{i=m}^{n} \binom{n}{i} p^i (1-p)^{n-i}$$

p = node availability

r = n / m = redundancy factor

Erasure Coding

TECHNISCHE UNIVERSITÄT DARMSTADT

$$A_{erasurecoding} = \sum_{i=m}^{n} \binom{n}{i} p^i (1-p)^{n-i}$$

p = node availability

r = n / m = redundancy factor

Availability target of 99.9%

presuming a mean node availability of 50%:

r = 2.49

EURECOM

How can Wuala be studied?

TECHNISCHE UNIVERSITÄT DARMSTADT

Wuala is a closed source project written in Java

Java bytecode obfuscated:

- No debug information available
- Name mangling was performed

→ Decompiling not helpful

However, observing network traffic should be meaningful!

EURECOM

General Experimental Setup

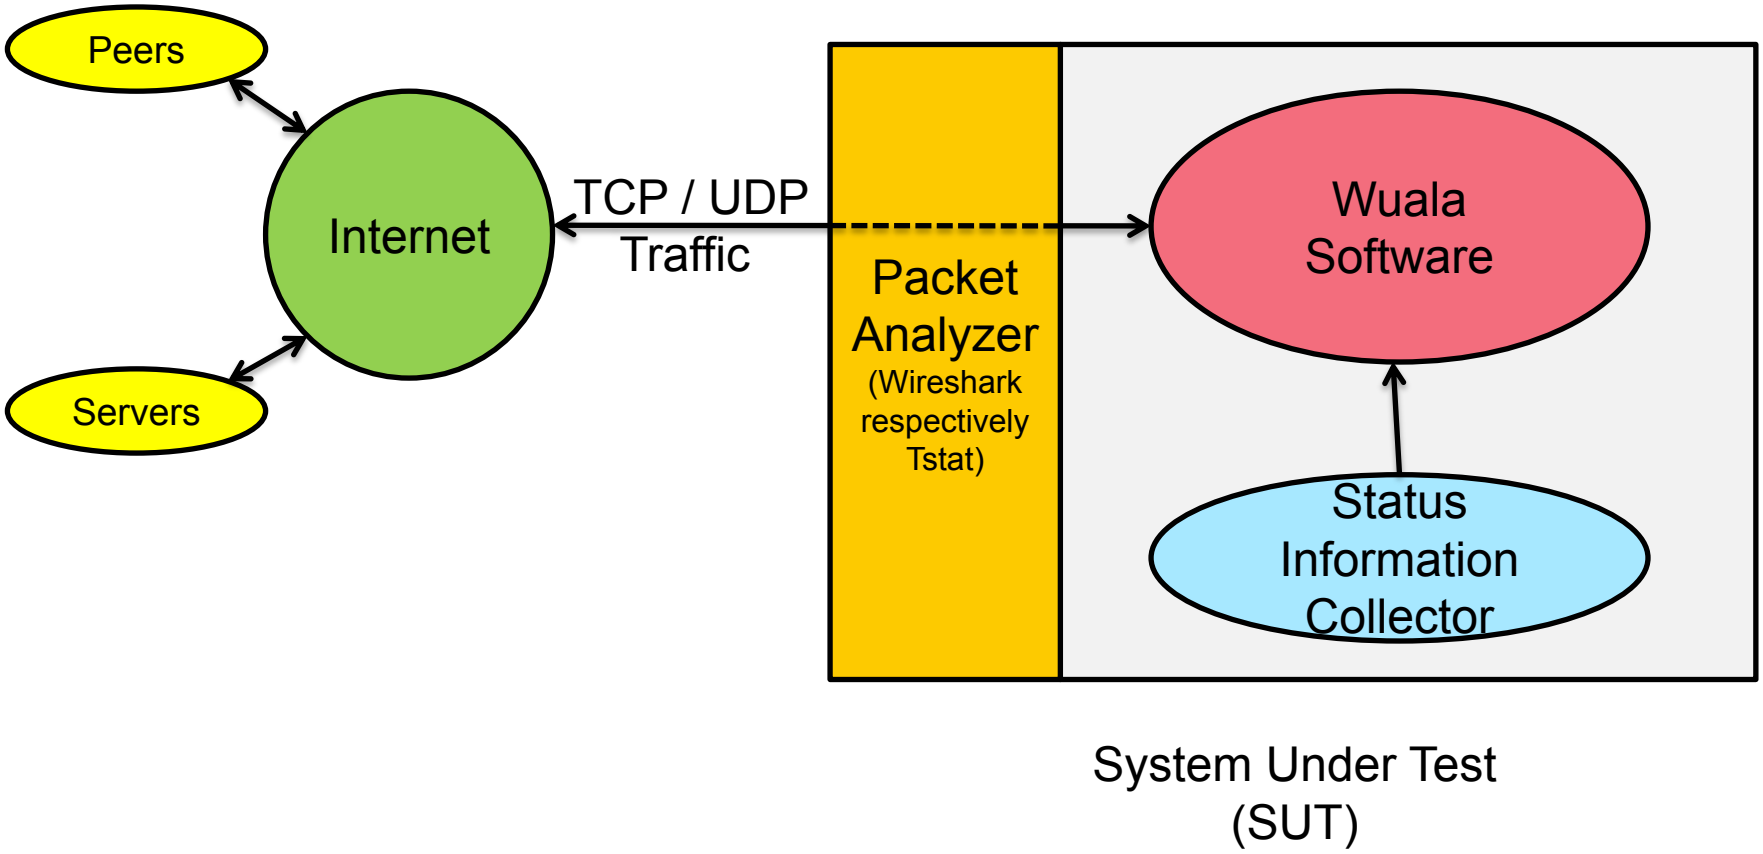

Used Packet Analyzer Depending on Purpose

TECHNISCHE UNIVERSITÄT DARMSTADT

	Wireshark	Tstat
Level of detail	complete packet dump	only summarized data
Output	binary file	text file
Content of packets	preserved	dropped
Postprocessing	possible	limited
Amount of data	large	small

EURECOM

Wireshark – Example Capture

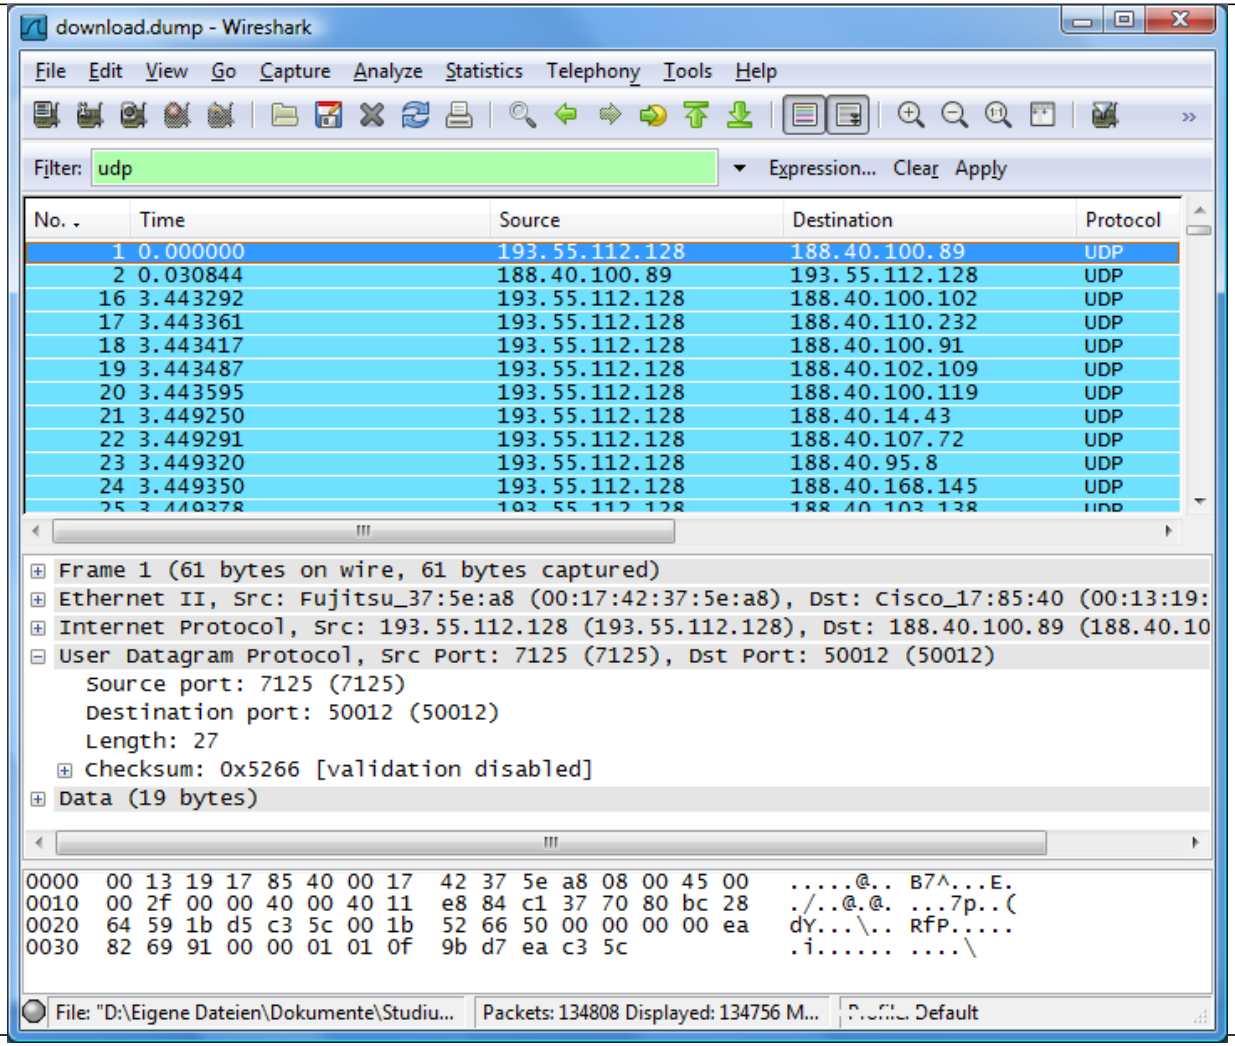

Tstat – Example Output

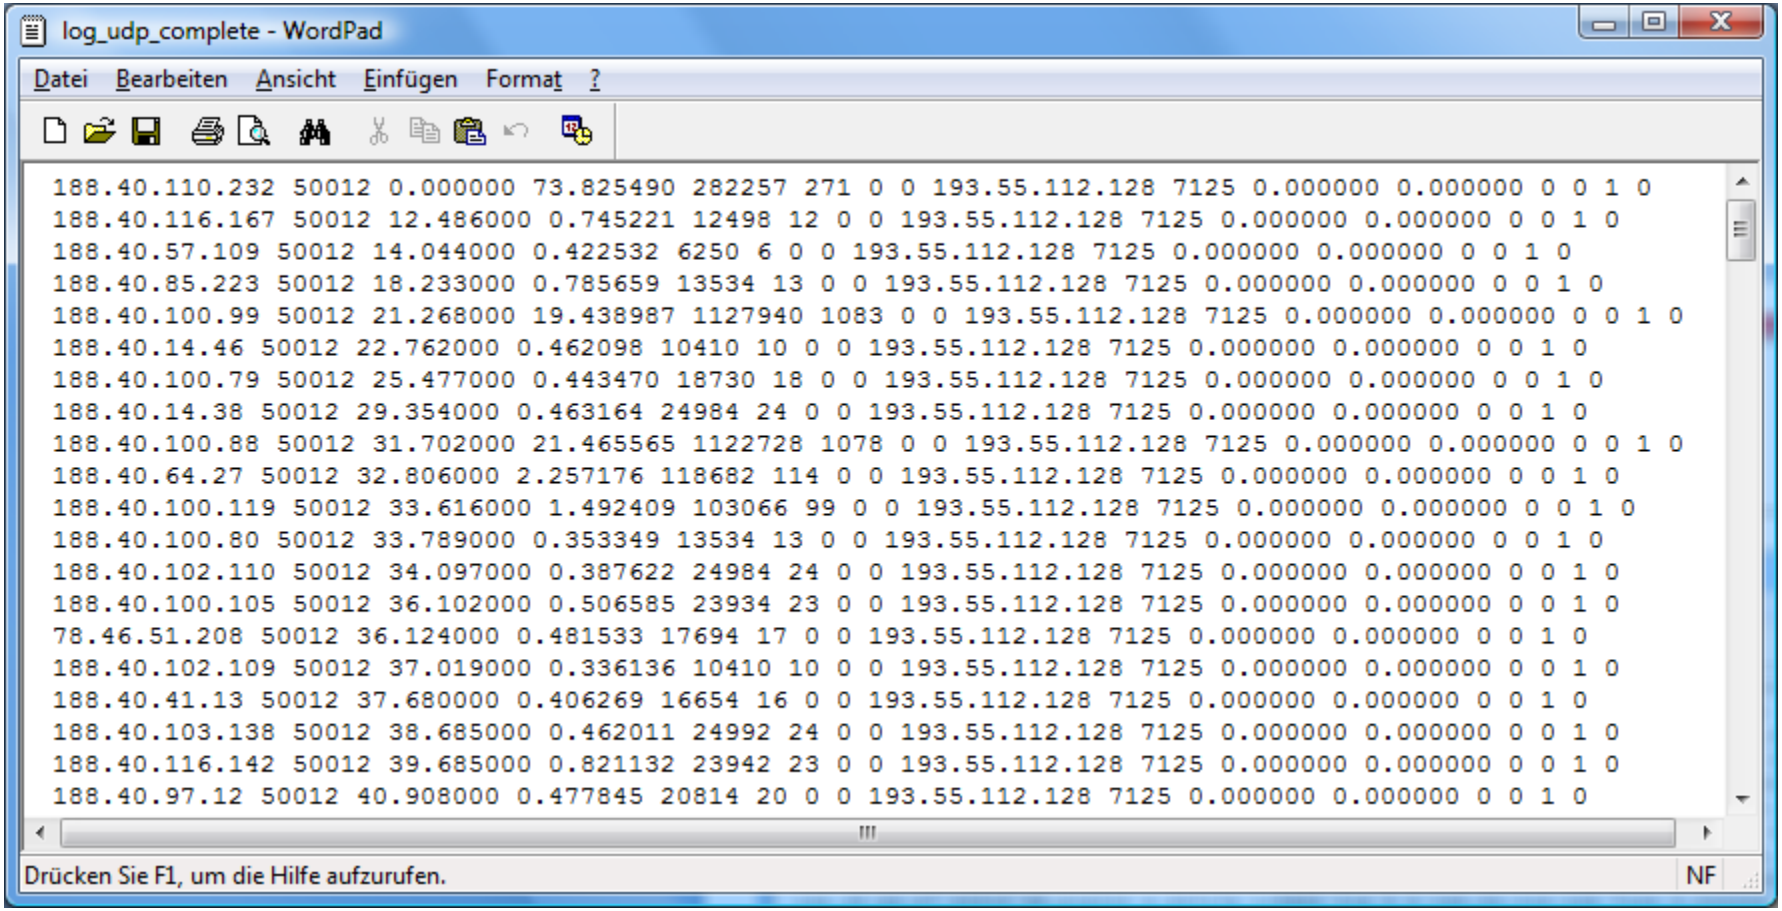

Experiment 1:
Discover Wuala Servers

TECHNISCHE UNIVERSITÄT DARMSTADT

How many servers are supporting the network?

- Servers should show a different behaviour
- Connection endpoints (IP addresses and ports) of particular interest
- Presumably servers located in special subnets

EURECOM

Experiment 1: Discover Wuala Servers

TECHNISCHE UNIVERSITÄT DARMSTADT

Test procedure:

- High level of participation in the network required:
  - 2 machines running as storage node for 2 months
  - 1 machine running mass download of files (20 GiB in total) for 4 days (location of data collected by crawler)
- All machines connected to the Internet via 100 Mbit/s link
- Tstat used as packet analyzer
- Database for ownership lookups of IP addresses

EURECOM

Experiment 1: Results

TECHNISCHE UNIVERSITÄT DARMSTADT

Number of IP Adresses	Country	Provider	Subnetwork
293 (!)	Germany	Hetzner Online AG	188.40.0.0/16 178.63.0.0/16 78.46.0.0/15
2	Switzerland	Atrila GmbH	188.92.144.0/21

Easy to differentiate between servers and peers:

- Servers:
  - UDP port 50012
- Peers:
  - random UDP port in range 7100-7600 (until specified in the application options)

Experiment 2: Where Are Files Stored?

TECHNISCHE UNIVERSITÄT DARMSTADT

Who is responsible to store uploaded data?

Test procedure:

- Upload files with random content and different sizes
- Sharing storage deactivated
- Use Wireshark to acquire information about endpoints, amount and content of transferred data

EURECOM

Experiment 2: Results

TECHNISCHE UNIVERSITÄT DARMSTADT

	File Size	Protocol	Stored on	Special Characteristic
Tiny File	0 – 4 KiB	TCP	Servers only	Embedded in metadata
Small File	4 KiB – 292 KiB	TCP	Servers only	Transferred to two servers
Medium File	292 KiB – 1 MiB	UDP	Servers only	Erasure coding performed
Big File	1 MiB – 14 GiB	UDP	Servers and peers	Erasure coding performed

EURECOM

Experiment 2: Results

TECHNISCHE UNIVERSITÄT DARMSTADT

- Wuala does not store data of all files on peers!
- Only data of files bigger than 1 MiB is sent to other peers
- Small files seem not to be worth the effort distributing them to peers

EURECOM

Experiment 3:
How Are Big Files Transferred?

TECHNISCHE UNIVERSITÄT DARMSTADT

How is fragmentation done?

Test procedure:

- Upload 100 MiB file with random content
- Sharing storage deactivated
- Wireshark used to generate dump file
- Tshark (delivered with Wireshark) to filter out signaling traffic
- Tstat to scan dump file
- Observe temporary folders

EURECOM

Experiment 3: General Result

TECHNISCHE UNIVERSITÄT DARMSTADT

Upload performed in following steps:

1. Encrypt file in temporary folder
2. Generation of Transmission Blocks (usage of erasure coding)
3. Upload to servers
4. Upload seems to be finished (!)
5. Subsequent upload to peers in background

EURECOM

Generation of Transmission Blocks

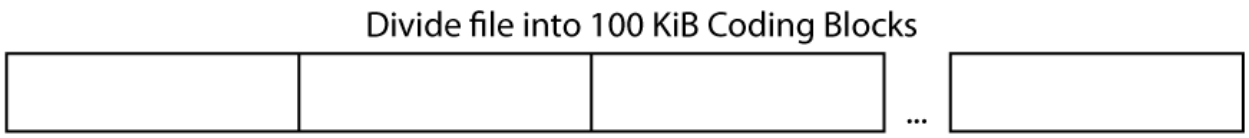

Generation of Transmission Blocks

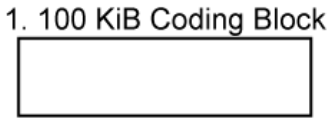

Generation of Transmission Blocks

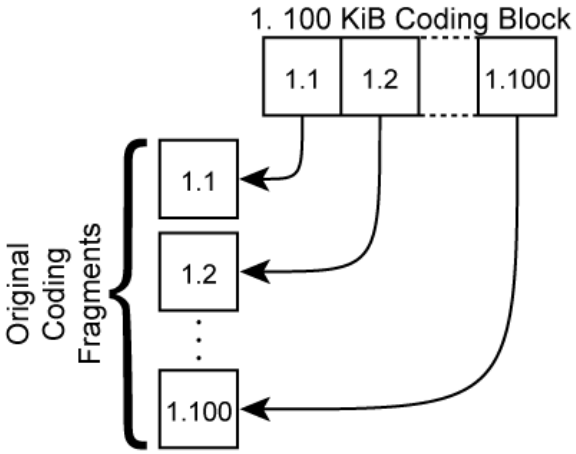

Generation of Transmission Blocks

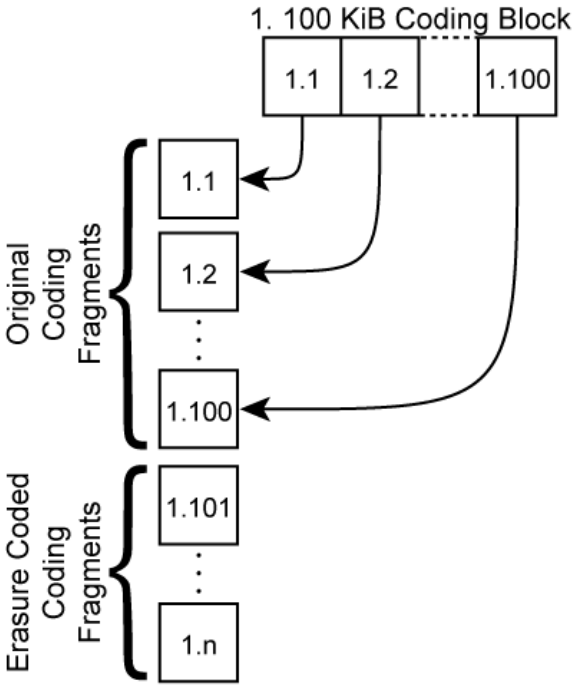

Generation of Transmission Blocks

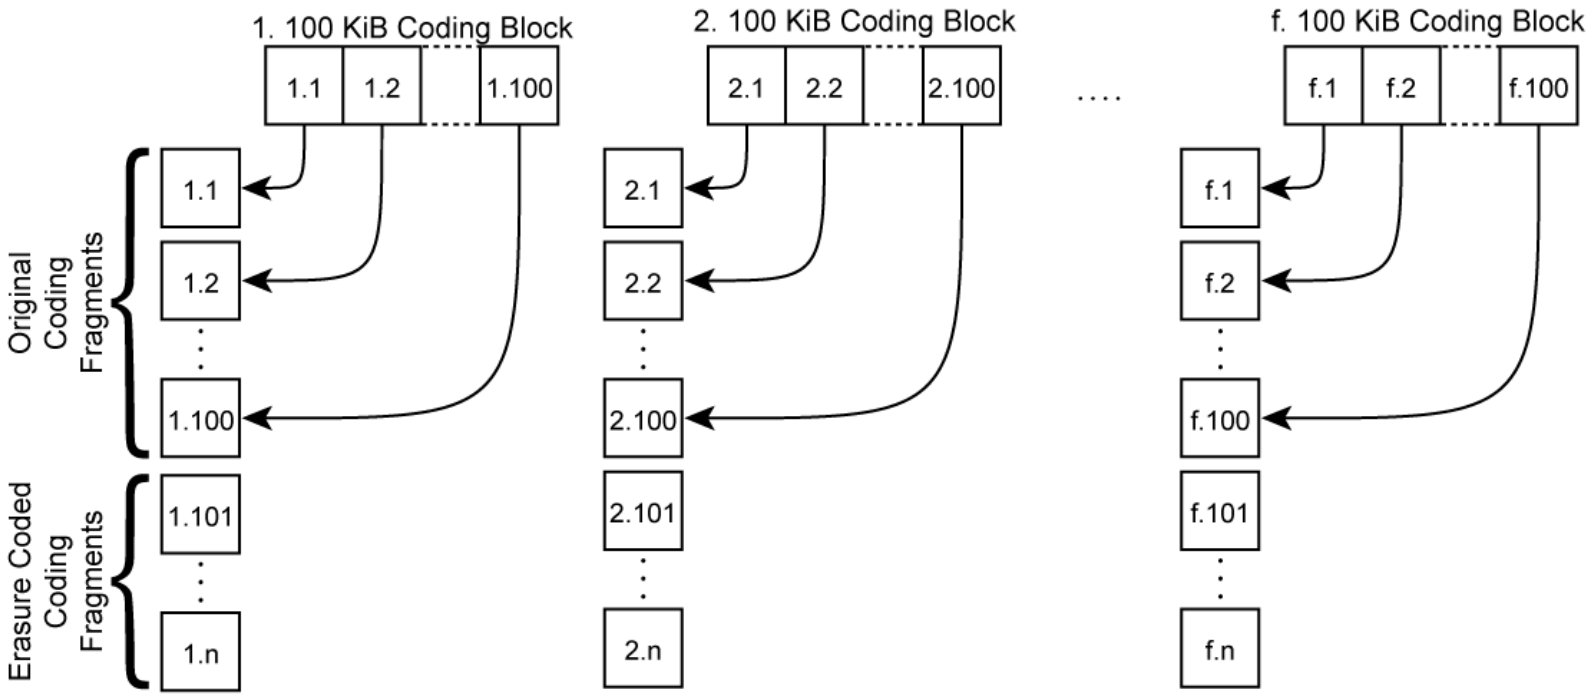

Generation of Transmission Blocks

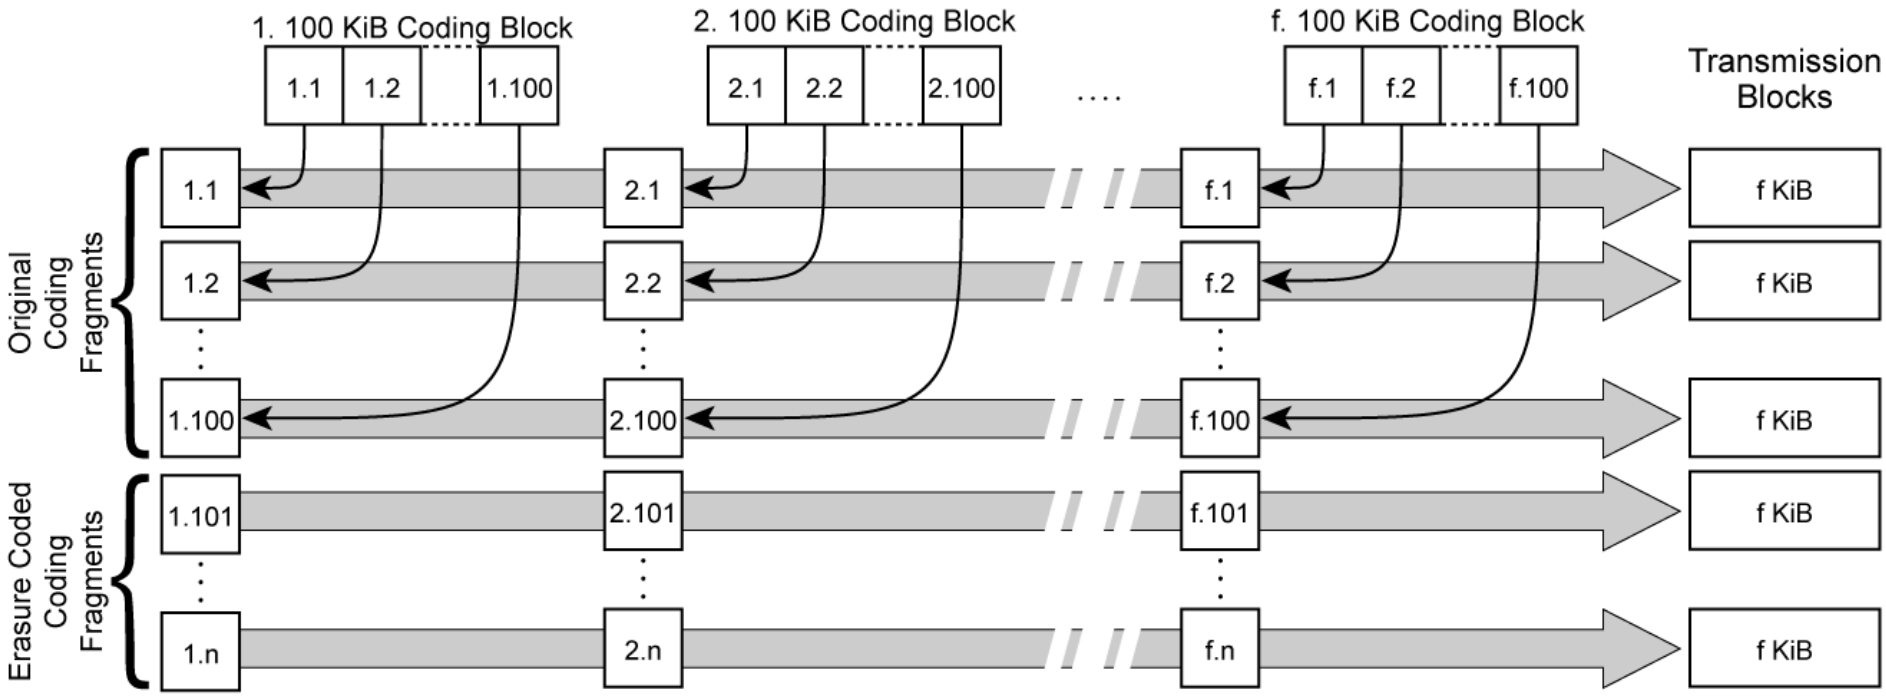

Generation of Transmission Blocks

Consequences:

- Interleaving allows streaming of files
- Download from many sources in parallel
- Switching to different sources possible
- Fast downloads

Experiment 3:
Results – Server Upload of 100 MiB file

TECHNISCHE UNIVERSITÄT DARMSTADT

First only upload to servers:

- Upload to 120 distinct servers via UDP
- Most transfer sessions ~1 MiB traffic, some servers receive this amount multiple times
- 144 MiB in total

  → redundancy on servers!
- Upload with full speed
- Concurrent transfers
- Upload completed after 50 seconds

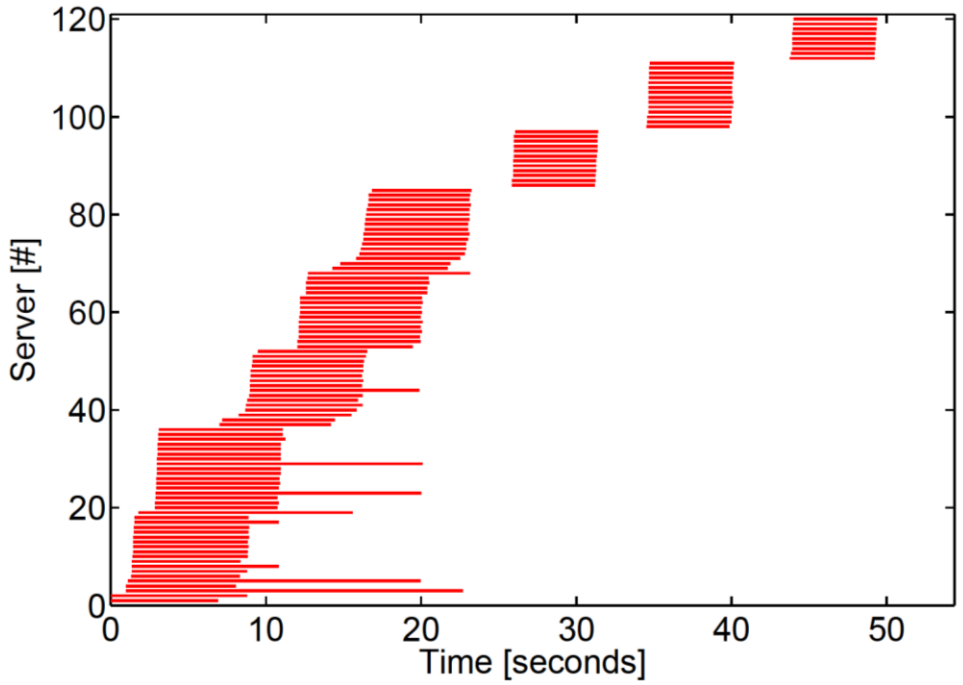

EURECOM

Experiment 3:
Results – Peer Upload of 100 MiB file

TECHNISCHE UNIVERSITÄT DARMSTADT

Subsequent upload to peers:

- Upload to 250 distinct peers via UDP
- Most transfer sessions ~1 MiB traffic
- 273 MiB in total
- Upload with limited speed
- Finished after 40 minutes

Overall ~400 MiB uploaded

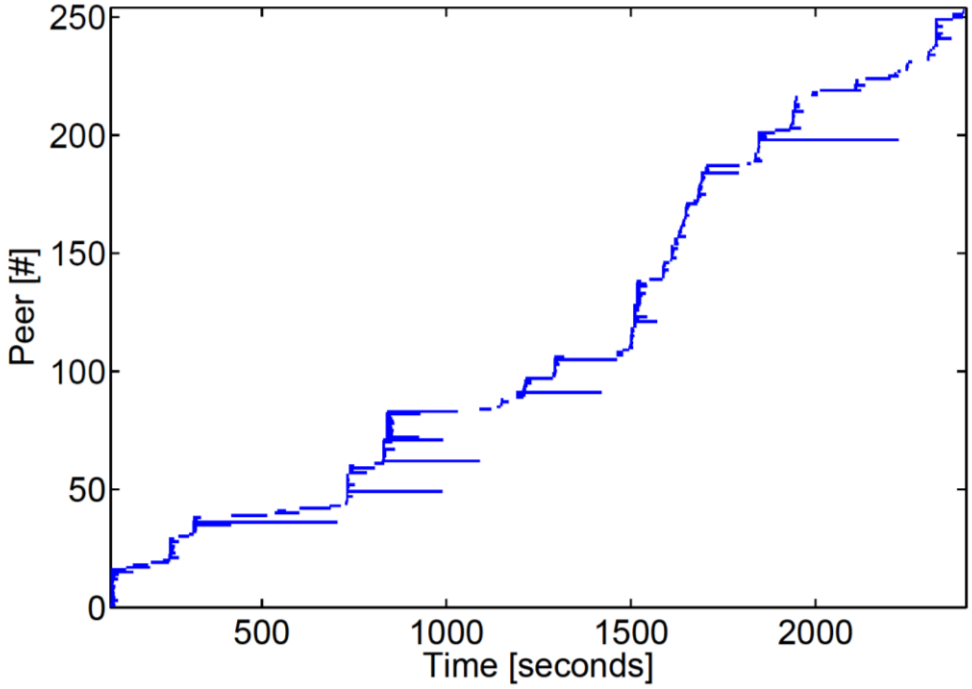

EURECOM

Experiment 4:
Who is serving data during download?

TECHNISCHE UNIVERSITÄT DARMSTADT

- Download 100 MiB random file again
  - right after upload
  - after 1 week
- Sharing storage deactivated
- Wireshark used to generate dump file
- Tshark (delivered with Wireshark) to filter out signaling traffic
- Tstat to scan dump file

EURECOM

Experiment 4:
Results – Right after Upload

TECHNISCHE UNIVERSITÄT DARMSTADT

- 21 MiB received from servers
- 97 MiB received from peers
- 118 MiB in total
- Download completed after

  72 seconds (~1,6 MiB/s)

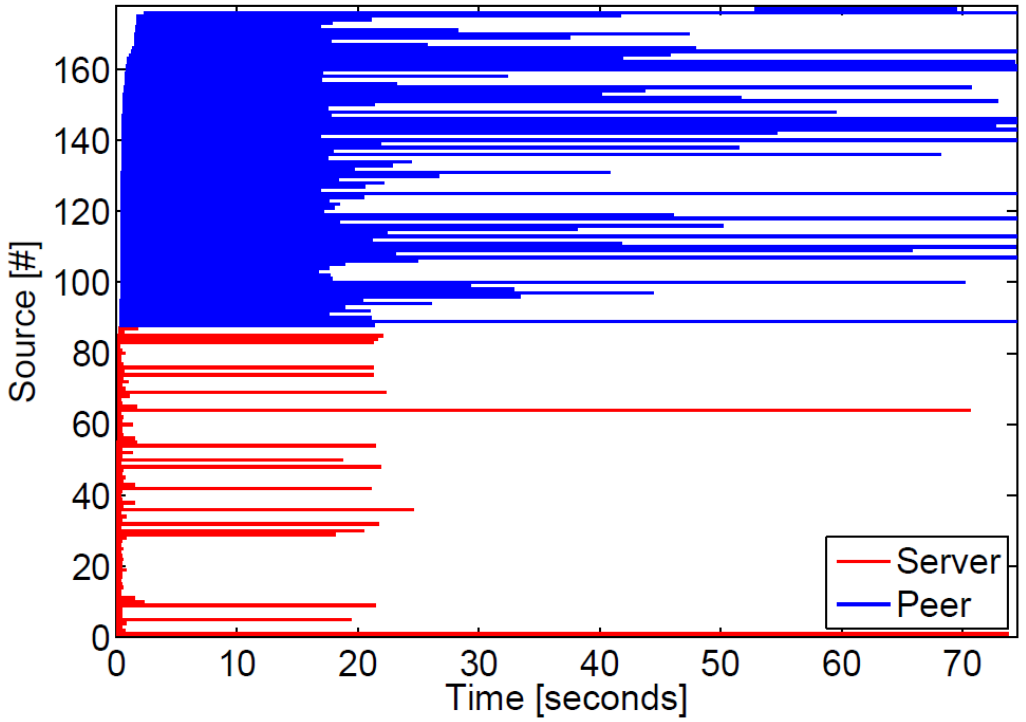

EURECOM

Experiment 4:
Results – 1 Week Later

TECHNISCHE UNIVERSITÄT DARMSTADT

- 100 MiB received from servers
- 23 MiB received from peers
- 123 MiB in total
- Download completed after

  43 seconds (~2,8 MiB/s)

→ Faster!

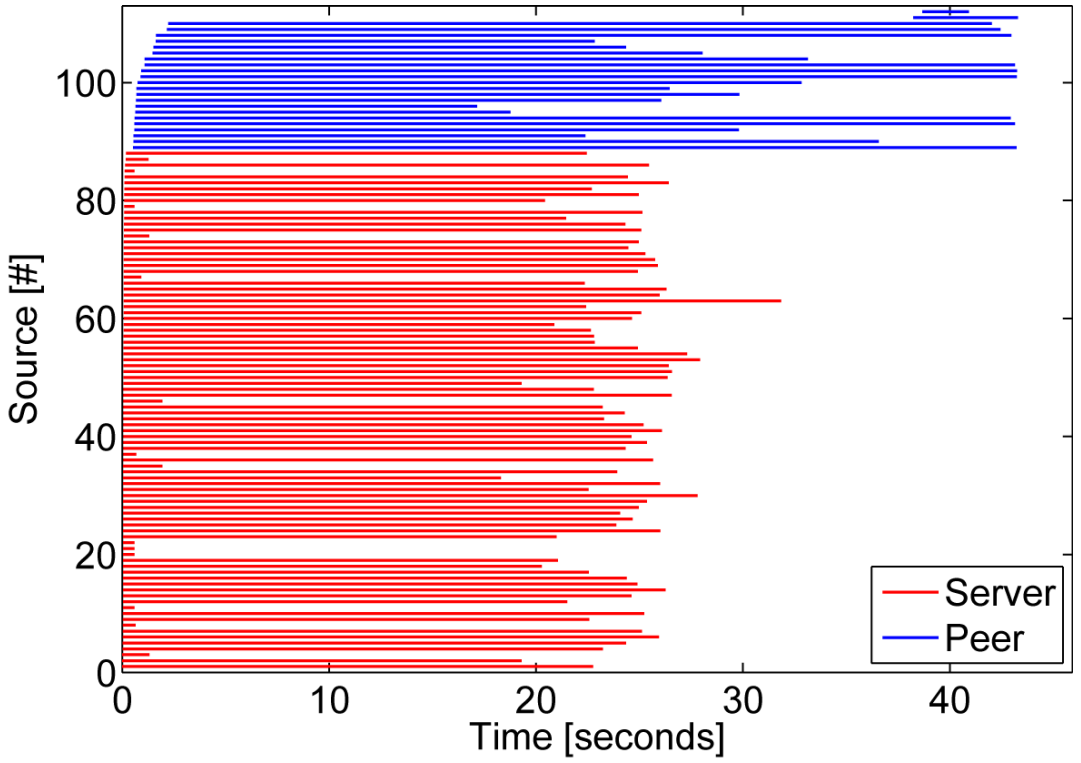

EURECOM

Experiment 5: Maintenance

TECHNISCHE UNIVERSITÄT DARMSTADT

Since peers may leave the network forever:

→ Are files maintained by a peer over time?

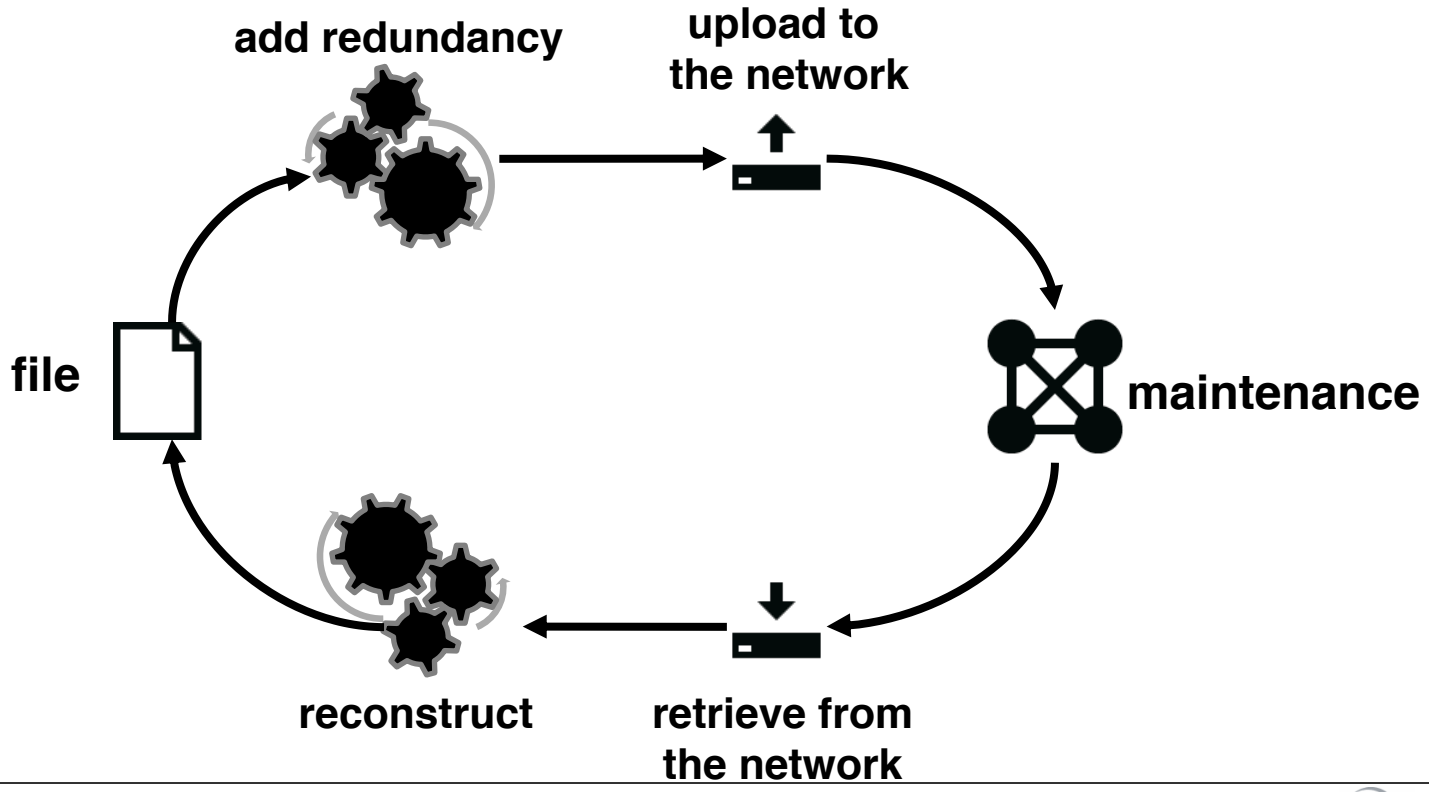

EURECOM

Experiment 5: Maintenance

TECHNISCHE UNIVERSITÄT DARMSTADT

Test procedure:

- Trading storage deactivated
- Upload 100 MiB file
- File locally deleted
- Use Tstat to observe the peer over 2 months

EURECOM

Experiment 5: Maintenance

TECHNISCHE UNIVERSITÄT DARMSTADT

Result:

No download within 2 months!

However: Subsequent Upload to Peers **after Download**

- As soon as data is not well distributed on peers
- Similar to initial peer upload
- After one week 132 MiB have been uploaded to peers again

→ Even for big files Wuala is relying on servers only!

EURECOM

Experiment 6:
How Are Sources for a Special File Located?

TECHNISCHE UNIVERSITÄT DARMSTADT

- Is there a DHT to locate sources of files?
  - Characteristic of a P2P network
- Or do servers provide this information?
  - Characteristic of a Client / Server Architecture

EURECOM

Experiment 6:
How Are Sources for a Special File Located?

TECHNISCHE UNIVERSITÄT DARMSTADT

- Long term observation revealed no communication to peers except for file transfers

  → no JOIN messages (required by common DHTs!)

- Use Wireshark to generate a dump file before initiating a download
- Put focus on first UDP packets received from servers

EURECOM

Experiment 6: Results

TECHNISCHE UNIVERSITÄT DARMSTADT

Repetition of special byte sequence obvious in datagrams

→ Presumably IP addresses of servers

First attempt:

XOR

- would result in constant cipher
- cheap usage

EURECOM

Experiment 6: Results

TECHNISCHE UNIVERSITÄT DARMSTADT

Repetition of special byte sequence obvious in datagrams

→ Presumably IP addresses of servers

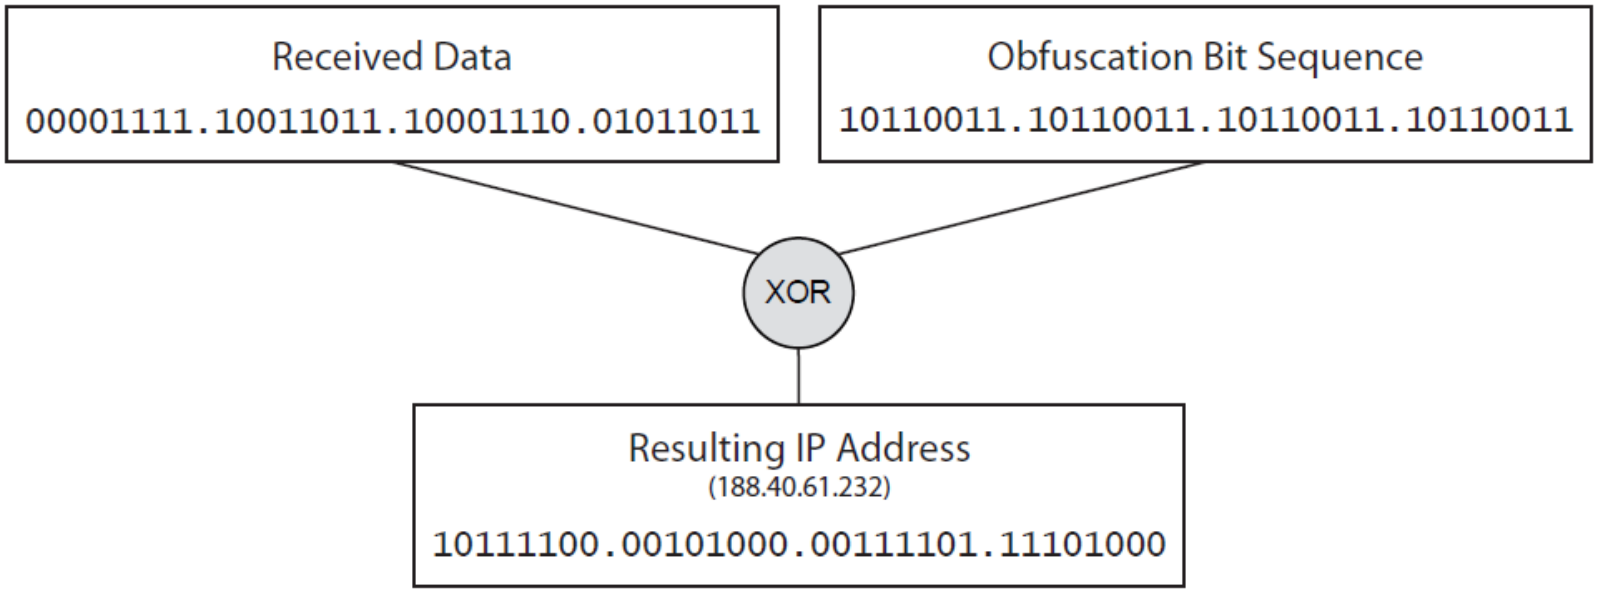

EURECOM

Experiment 7:
Earned On-line Storage vs. Occupied Local Storage

TECHNISCHE UNIVERSITÄT DARMSTADT

Is there any imbalance, pointing out that users earn more on-line storage than they actually offer?

Procedure:

- Tstat observing a machine operating as storage node
- Duration of 90 days for data acquisition
- Log all status information provided by the Wuala application
- Log amount of used local storage

EURECOM

Experiment 7: Results

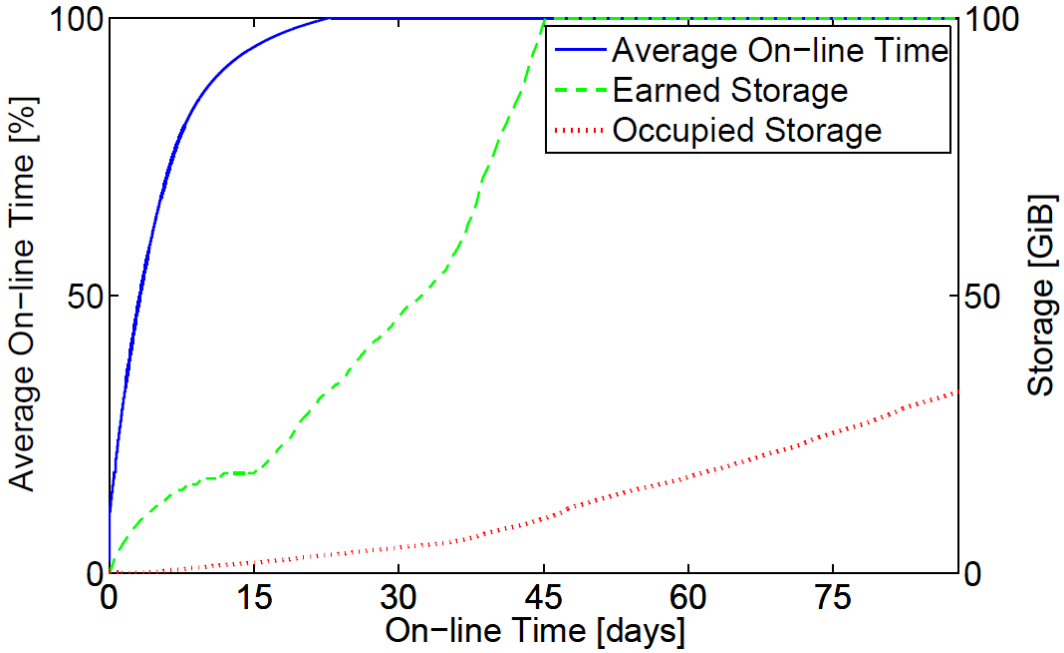

First only depending on average on-line time and maximum shared space:

$$earned\ storage = on\text{-}line\ time \times shared\ space_{maximum}$$

After 2 GiB data has been stored:

$$earned\ storage = \min\{on\text{-}line\ time \times shared\ space_{occupied} \times 10; 100\}$$

Overall a long-lasting gap between earned storage and occupied storage

Incoming Traffic over 120 Days

TECHNISCHE UNIVERSITÄT DARMSTADT

Incoming traffic increases over time - indicates rating of peers

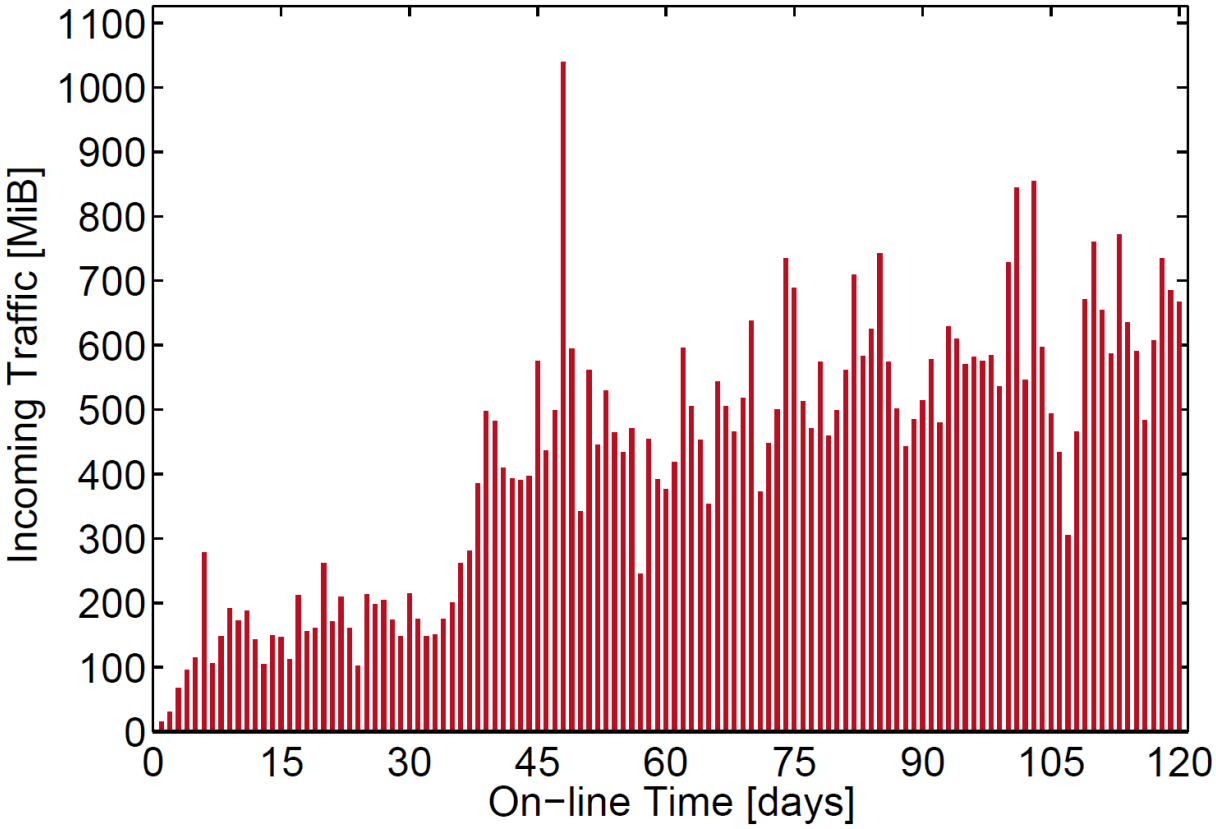

EURECOM

Outgoing Traffic over 120 Days

TECHNISCHE UNIVERSITÄT DARMSTADT

Correlation: More data stored locally → more requests for stored data
~20 KiB/s after 4 months

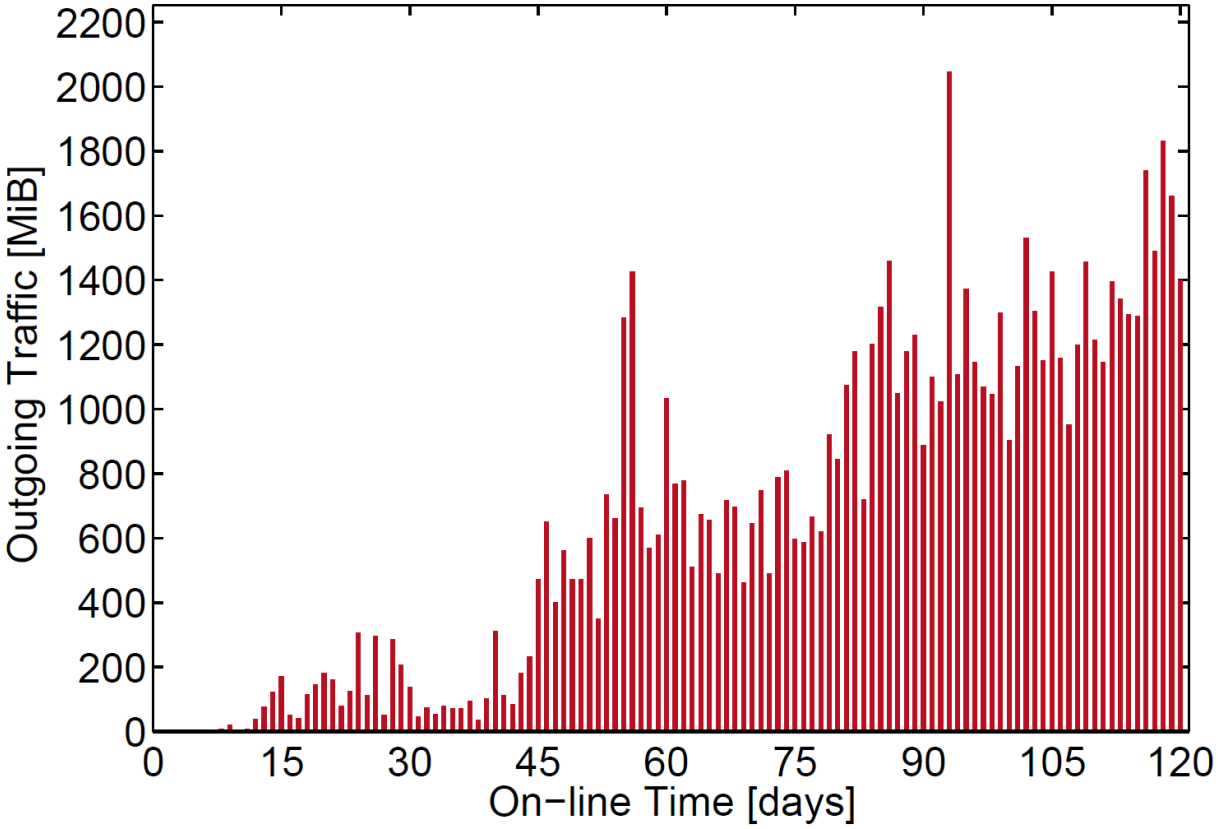

EURECOM

Conclusion

TECHNISCHE UNIVERSITÄT DARMSTADT

- Wuala only **peer assisted**, far away from pure P2P
- **No DHT used on peers**
- Lots of servers responsible to guarantee access on files in long term
- Not all data is uploaded to peers
- Imbalance between offered local storage and on-line storage
- Fast download / streaming possible
- Service dependent on company
- **Peers used to save traffic and load on servers**

EURECOM

Thank you!

TECHNISCHE UNIVERSITÄT DARMSTADT

Questions?

EURECOM