



www.kom.tu-darmstadt.de
http://www.comp.lancs.ac.uk

Assessing Peer Reliability in Cluster based Overlay Networks

V. Darlagiannis, **A. Mauthe**, R. Steinmetz

TU Darmstadt - Darmstadt University of Technology,
Dept. of Electrical Engineering and Information Technology, Dept. of Computer Science
KOM - Multimedia Communications Lab
Merckstr. 25, D-64283 Darmstadt, Germany,
Vasilios.Darlagiannis@KOM.tu-darmstadt.de Tel.+49 6151 166155, Fax. +49 6151 166152



Lancaster University,
Computing Department,
Lancaster, LA1 4YR, UK

andreas@comp.lancs.ac.uk Tel+44/1524/5 10485, Fax. +44/1524/5 10492



Outline

Introduction

Overview of Omicron

- Architecture
- Operations

Mechanisms of Omicron

- Clusters
- Role assignment

Applying burn-in at heterogeneous peer-to-peer networks

- Uptime distribution
- Cost based modeling
- Survivability

Conclusions

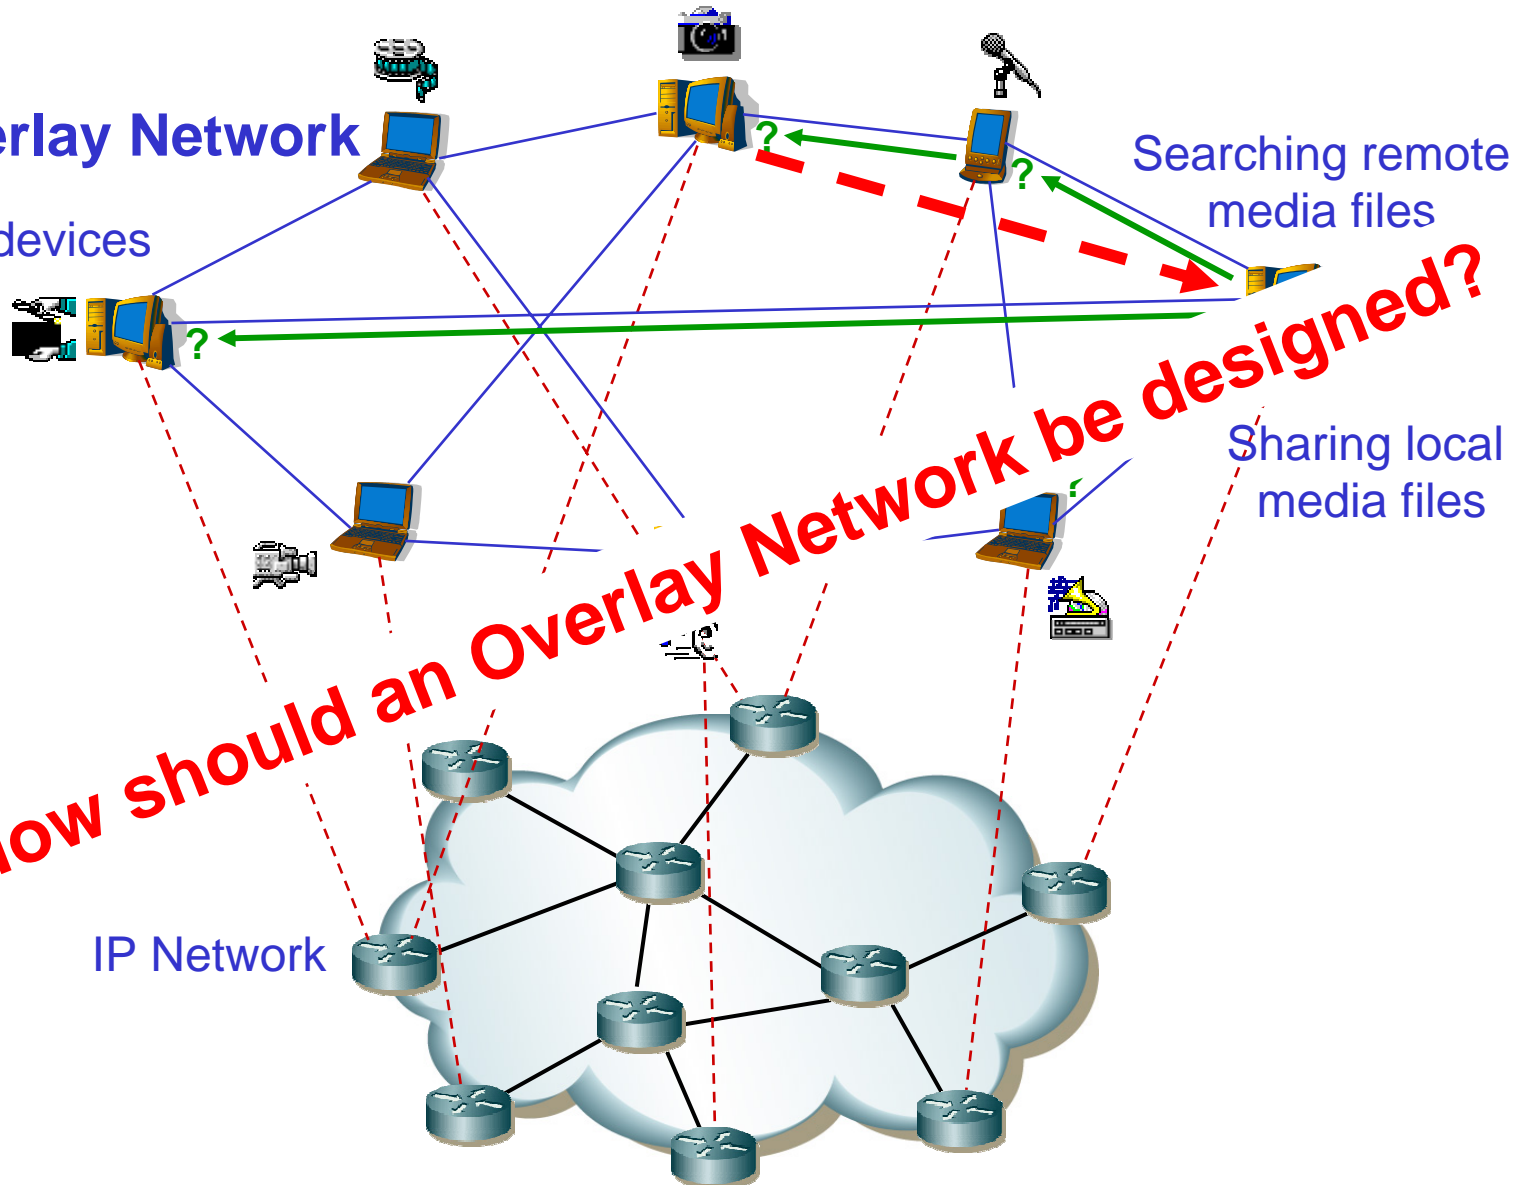


Peer-to-Peer Overlay Networks

www.kom.tu-darmstadt.de
http://www.comp.lancs.ac.uk

Overlay Network

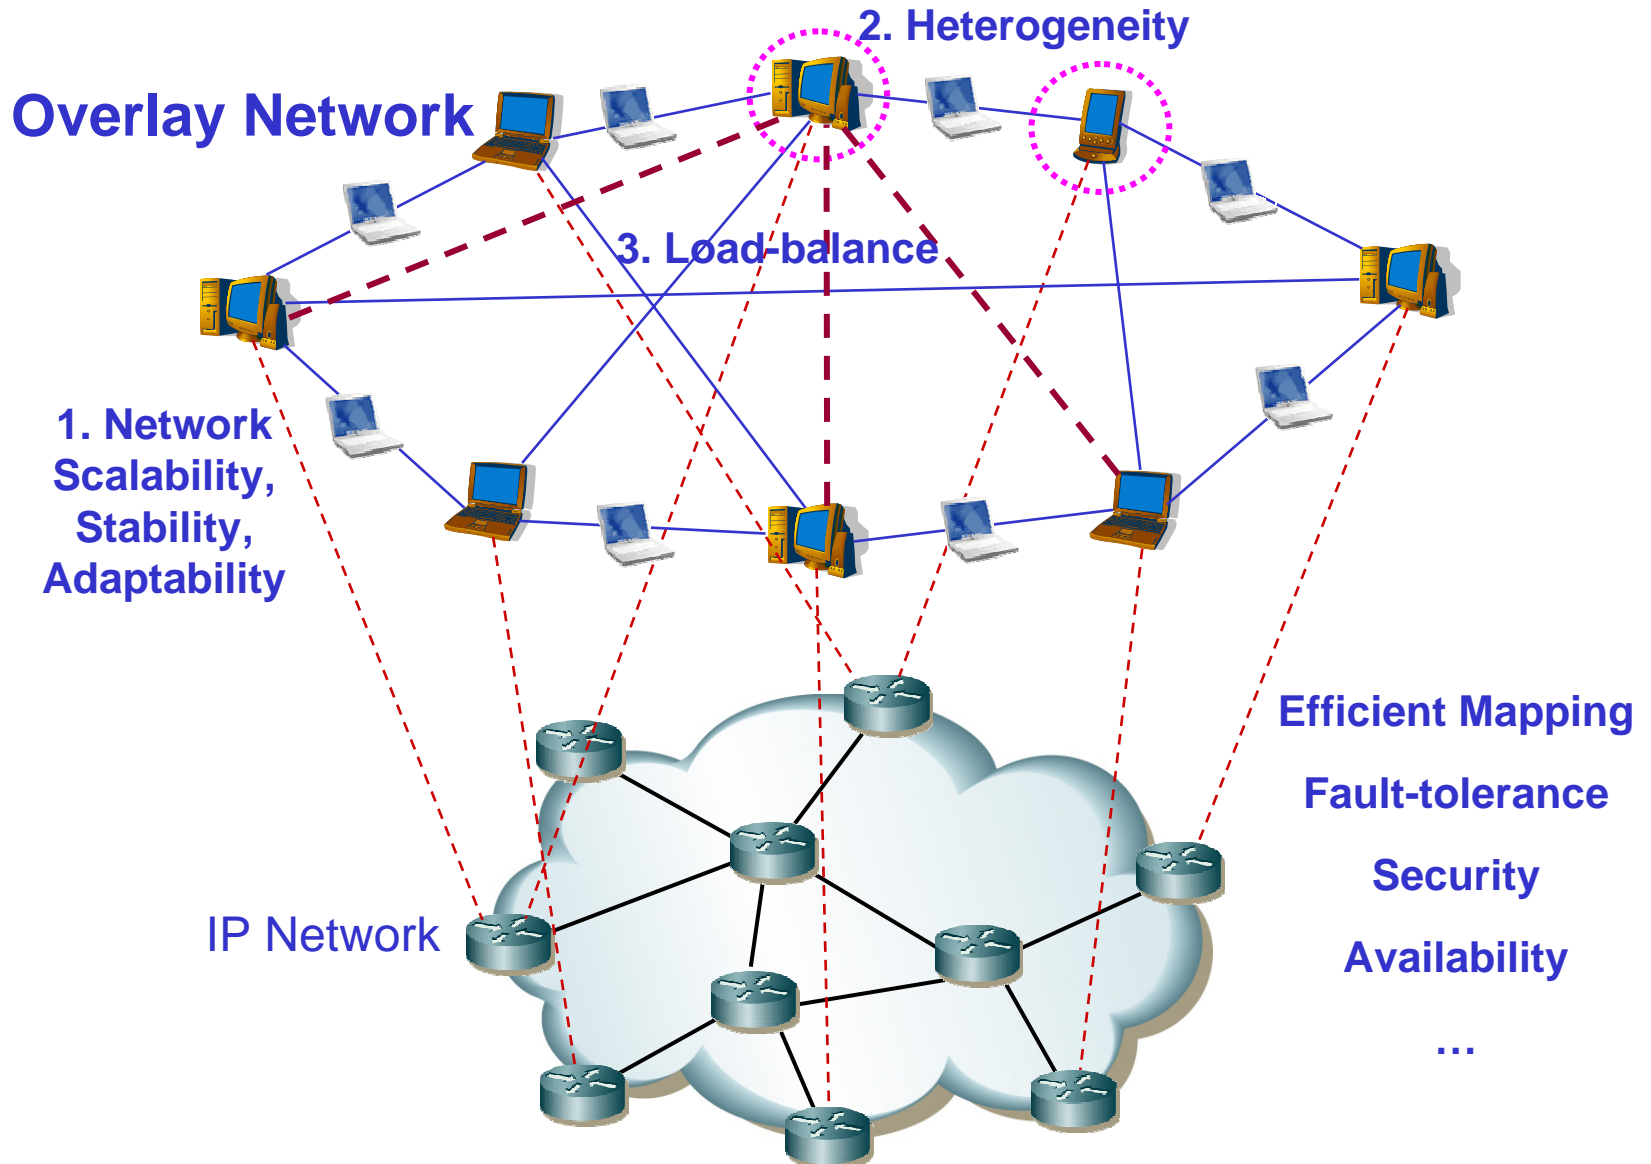
End-devices





Critical Requirements for Overlay Networks

www.kom.tu-darmstadt.de
http://www.comp.lancs.ac.uk





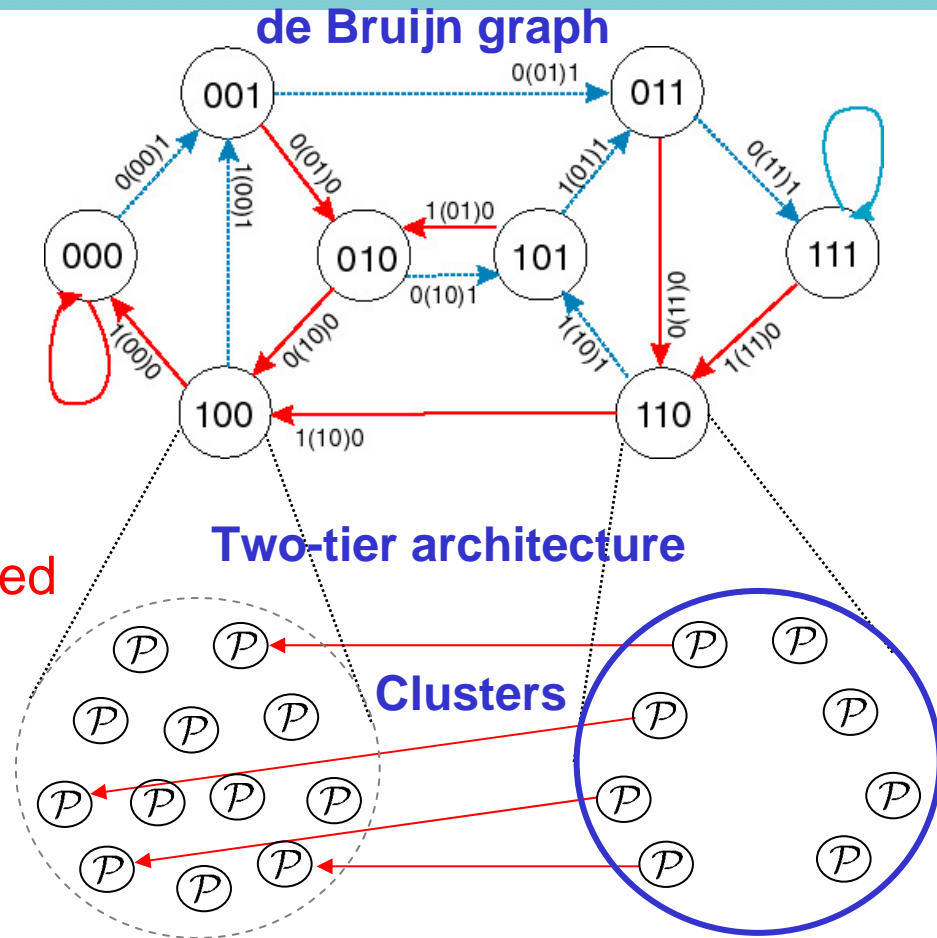
Omicron: Overlay Network Topology

Structured macro level (de Bruijn)

- Scalable
 - Asymptotically optimal
 - Diameter
 - Average node distance
 - Fixed node degree
- **Stable components required**

Clustered micro level

- Redundancy and fault-tolerance
- Locality aware
- Finer load balance
- Handling hot spots





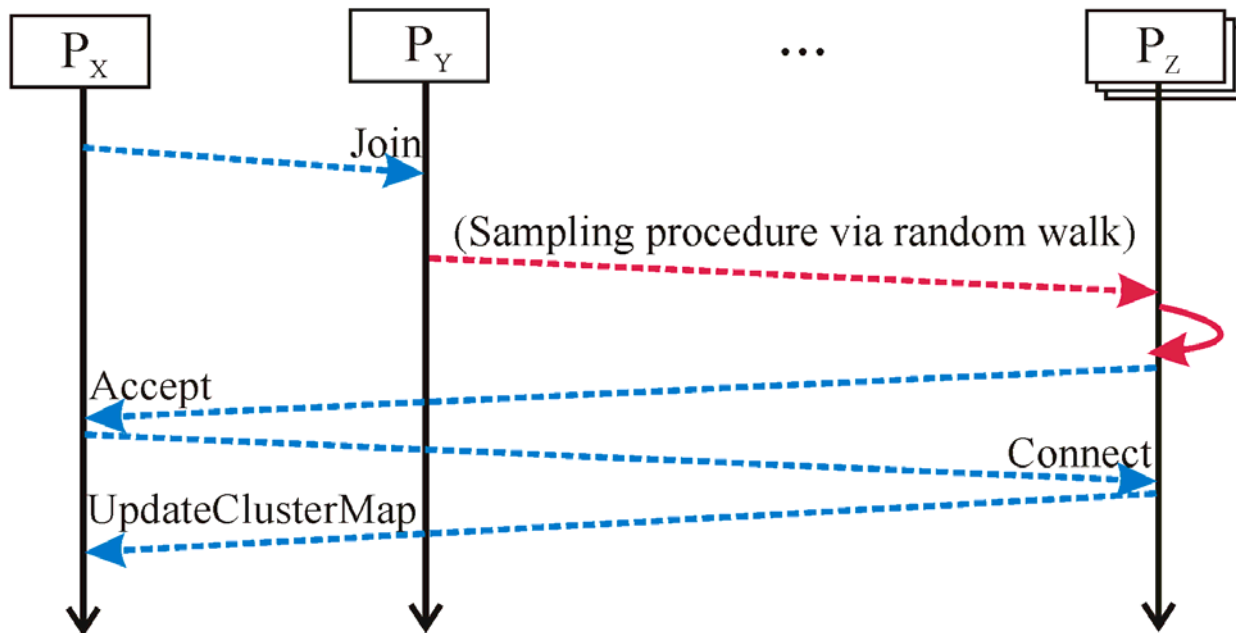
Join Process

How to select a cluster to join (or migrate)?

- Aiming at a globally stable network
- "Stable/Reliable" clusters are the target of the joining and maintenance processes

Lack of complete network stability knowledge

- Very costly to obtain
- Sampling is a promising alternative





Cluster Endurance and Network Stability

Definitions and terminology

Network stability $S_N(t)$ is the probability that the topology of the network remains *unmodified* until some time ' t '.

Peer reliability $R_P(t)$ is the probability that the peer remains *online* until some time ' t '.

$$R_P(t) = \Pr\{X > t\} = 1 - F(t).$$

Cluster endurance $E_C(t)$ is the probability that at least one peer of the cluster will remain *online* until some time ' t '.

$$E_C(t) = 1 - \prod_{i=1}^K (1 - F_i(t))$$



Omicron: Role Specialization for Peers

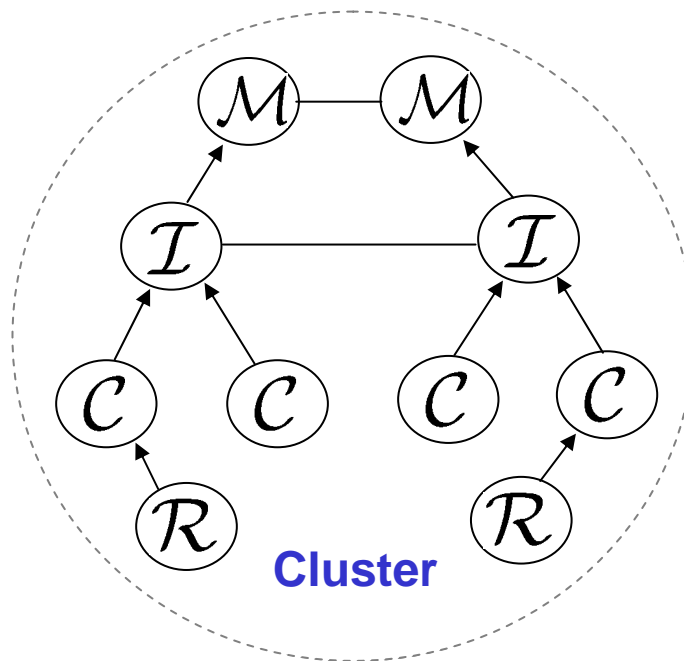
www.kom.tu-darmstadt.de
http://www.comp.lancs.ac.uk

Ⓜ **Maintainer**

Ⓡ **Indexer**

Ⓒ **Cacher**

Ⓡ **Router**



Organized
Maintenance,
Indexing,
Caching and
Routing in
Overlay
Networks

Common overlay network operations

- Maintaining structure (topology)
- Indexing advertised items
- Caching popular items
- Routing queries



Optimal Role Assignment

After cluster selection

- roles are assigned to the newly joined or migrated peers

Assignment of demanding roles is costly

- Such roles should be assigned to reliable peers

Goal

- Minimize the cost of assigning demanding roles to peers

Approach

- Investigate peer uptime distribution
- Exploit the distribution properties with a burn-in method



Peer Uptime Distribution

Parametric fitting of empirical observations

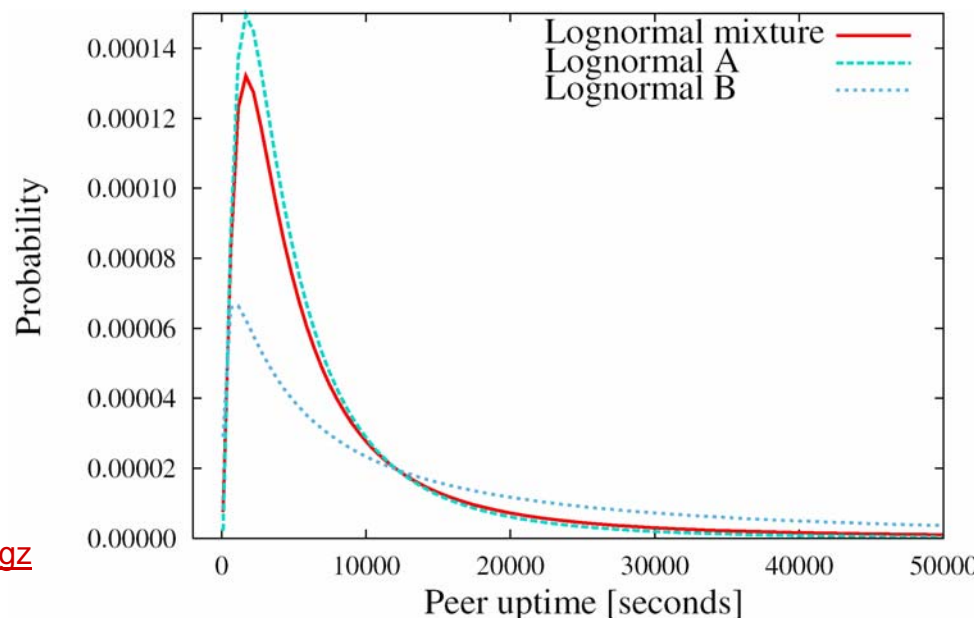
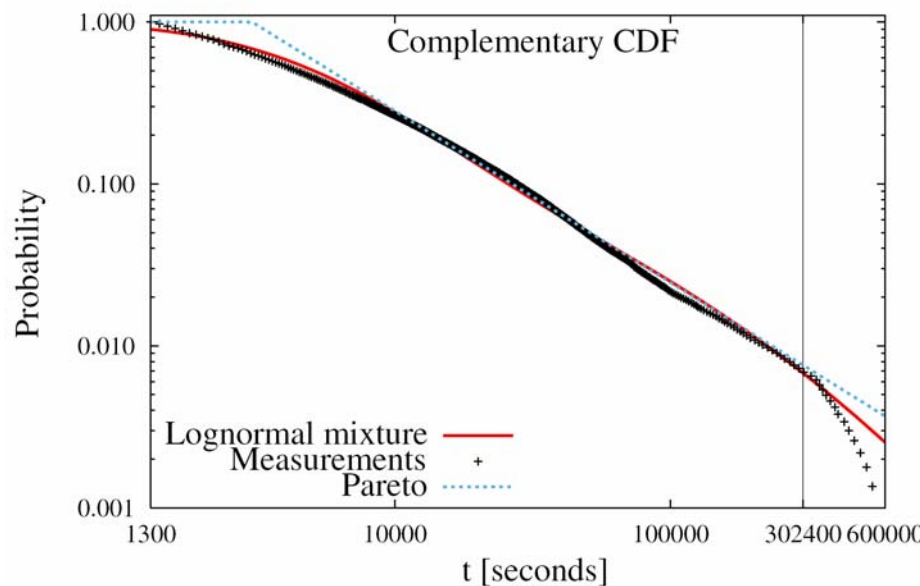
- Pareto fits well the tail
- Lognormal fits better the head

Lognormal mixture model

- Swarm-based file sharing
 - Multiplicative process generative model
- Peer heterogeneity

Outcome

- Lognormal mixture fits best





Burn-in: Method and Conditions

Burn-in method

- Tests the reliability of manufactured components
- Identifies defected components
- Applies effectively when defectiveness appears early
- Suitable for heterogeneous populations

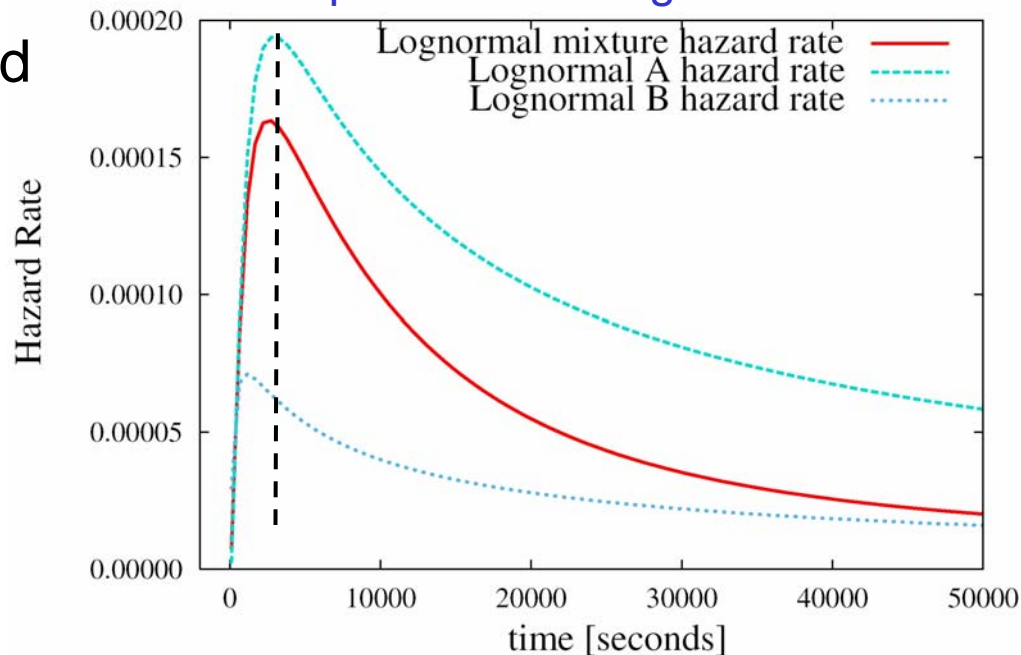
Condition

- Hazard rate (failure rate) shape

Lognormal mixture hazard rate shape

- Failure is less probable as peers survive (DFR)
- Strongest component convergence

Optimal role assignment time?





Cost-based Burn-in Concepts

Peer role utility, e.g. $u_R = \# \text{ routed messages}$

Role assignment cost: $c_I = \text{required indexing information}$

Mission time: t_m

- Minimum utilization time

Burn-in time: t_b

- Maximize the probability that a peer will survive for the additional mission time after it has survived the burn-in process

Expected burn-in cost: Burn-in failure cost Mission time failure cost

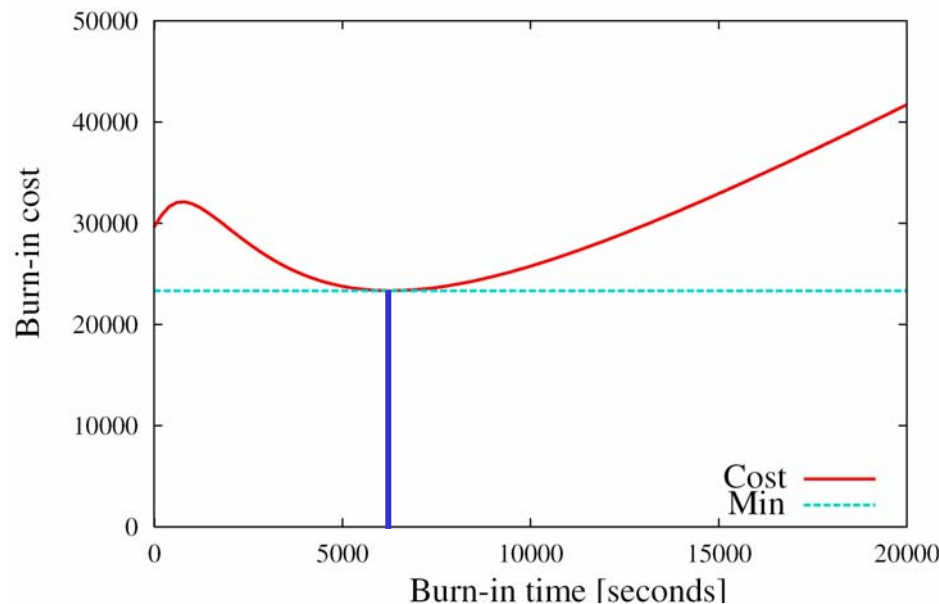
$$E[C_b(t_b)] = \underbrace{u_c t_b}_{\text{Lost utility}} + \underbrace{c_{bf} F(t_b)}_{\text{Failure during burn-in}} + \underbrace{c_{mf} [F(t_b + t_m) - F(t_b)]}_{\text{Failure during mission time}}$$



Evaluation: Peer Survivability

Burn-in cost

- Example use case:
 $c_a = 10000$ units
 $u_c = 2$ units/sec
 $c_{mf} = 60000$ units
- Results:
 $t_m = 5000$ secs
 $t_b \simeq 6300$ secs



Survivability

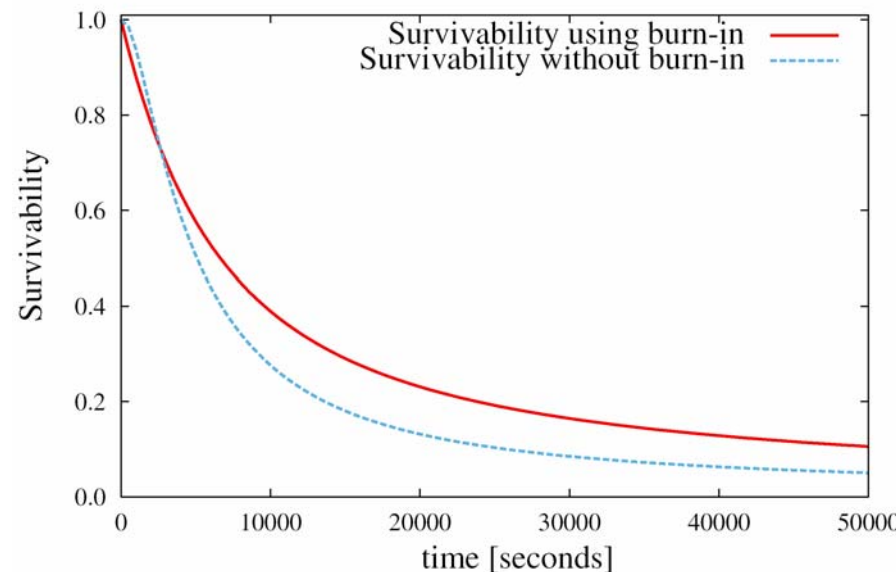
- Without burn-in:

$$R(t) = 1 - F(t)$$

- Using burn-in:

$$R_t(y) = \frac{R(t+y)}{R(t)}$$

Burn-in time





Conclusions

Summary

- Omicron
 - Hybrid overlay network
 - Designed to meet several critical requirements for P2P systems
- Sampling mechanism of Omicron
 - Employed to provide a well balanced and stable network
 - Low operation cost
- Role Specialization
 - Capitalizing on peer uptime distribution
 - Maximizing peer utilization by addressing heterogeneity
- Burn-in can be used to derive more stable and efficient systems when heterogeneous peers are present

Outlook

- Real role assignment costs
- Phase type based modeling



Thank You!

www.kom.tu-darmstadt.de
http://www.comp.lancs.ac.uk



MMCN 2007

<http://www.ifi.uio.no/mmcn2007/>