

Peer-to-Peer Technology in Grid Computing

—

The Igor File System and Beyond

April 2008

*Kendy Kutzner,
and Thomas Fuhrmann*




Faculty of Informatics, System Architecture Group,
Universität Karlsruhe (TH), Germany



Department of Informatics, Chair of Network Architectures, Technical
University Munich, Germany

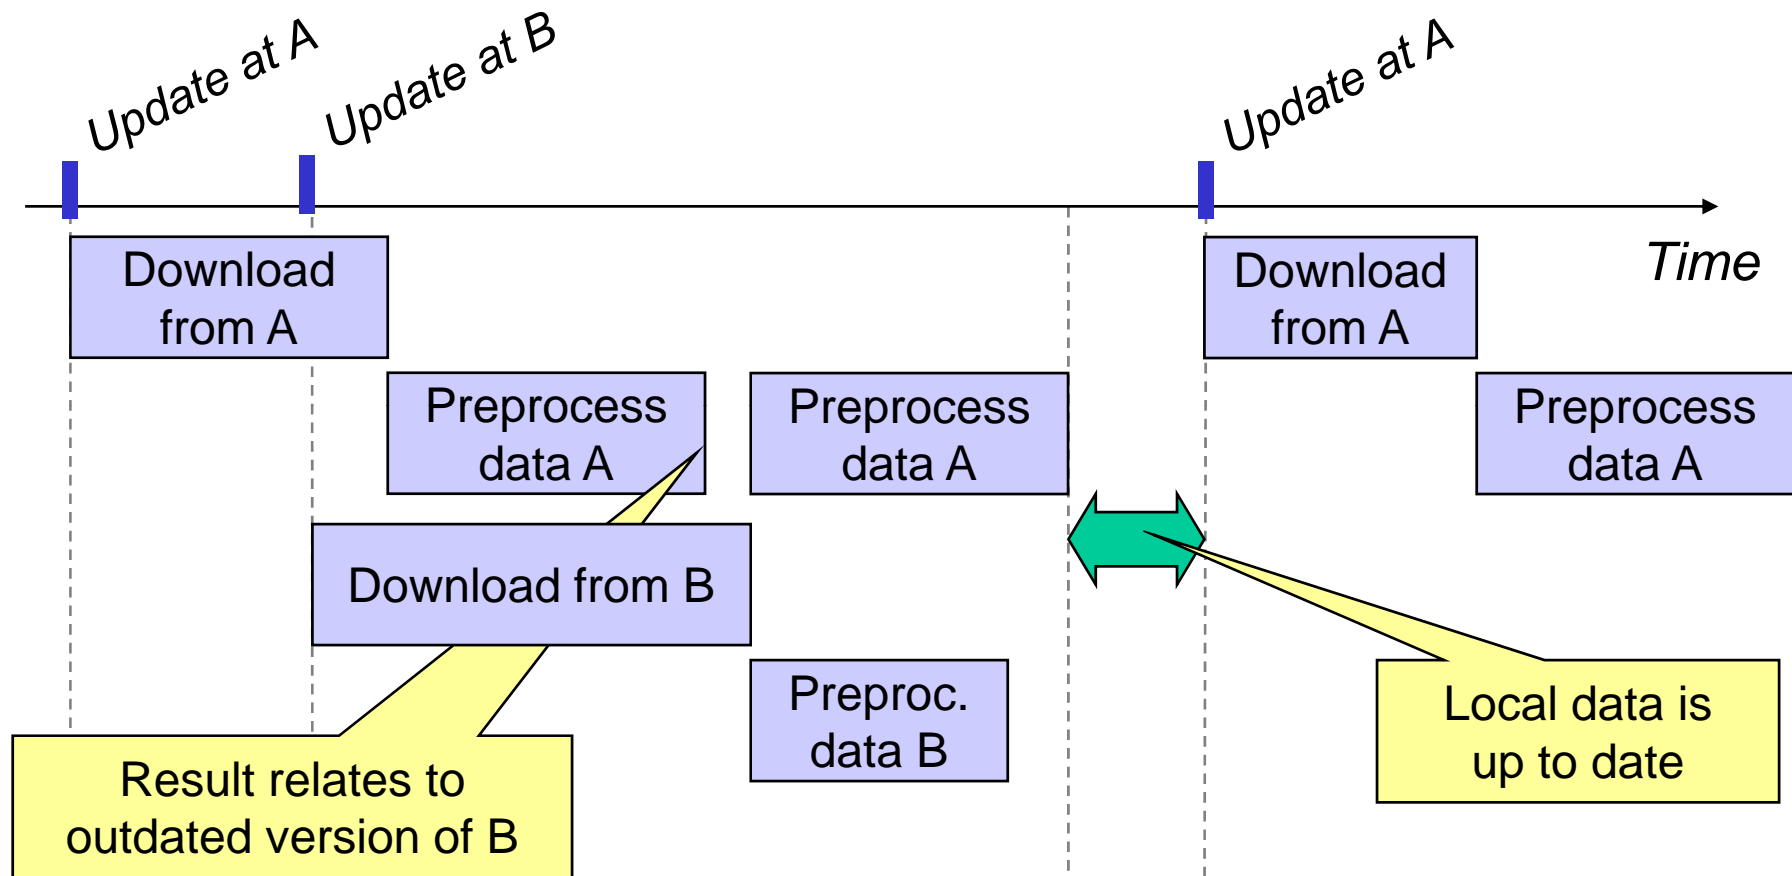
The Pharma Challenge (1)



Various sites produce public and private data that is stored and exchanged as flat files, typically by FTP.

In order to work with the data, sites need to download all the external data before they can start to work, even though they need only a small fraction of the data.


The Pharma Challenge (2)



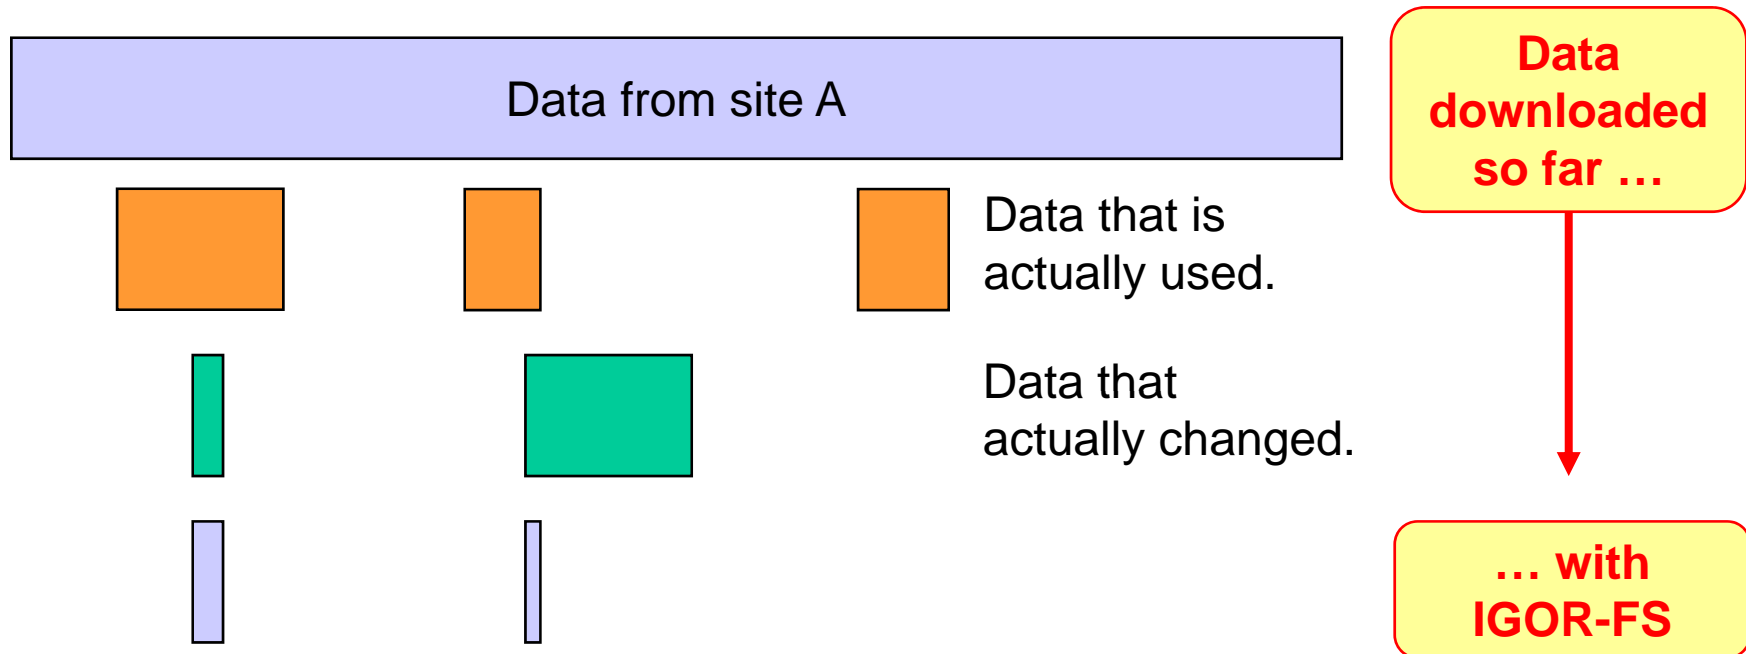
This becomes increasingly difficult: up to 1000 sources, weekly changes, preprocessing time about a day.

The IGOR-FS Solution

Today, sites use only a few sources in practice. Ideally, their system could handle cross-references between these sites.



But in fact they use only a small part of that data, and they reload the whole data even though it did not change much.



Agenda

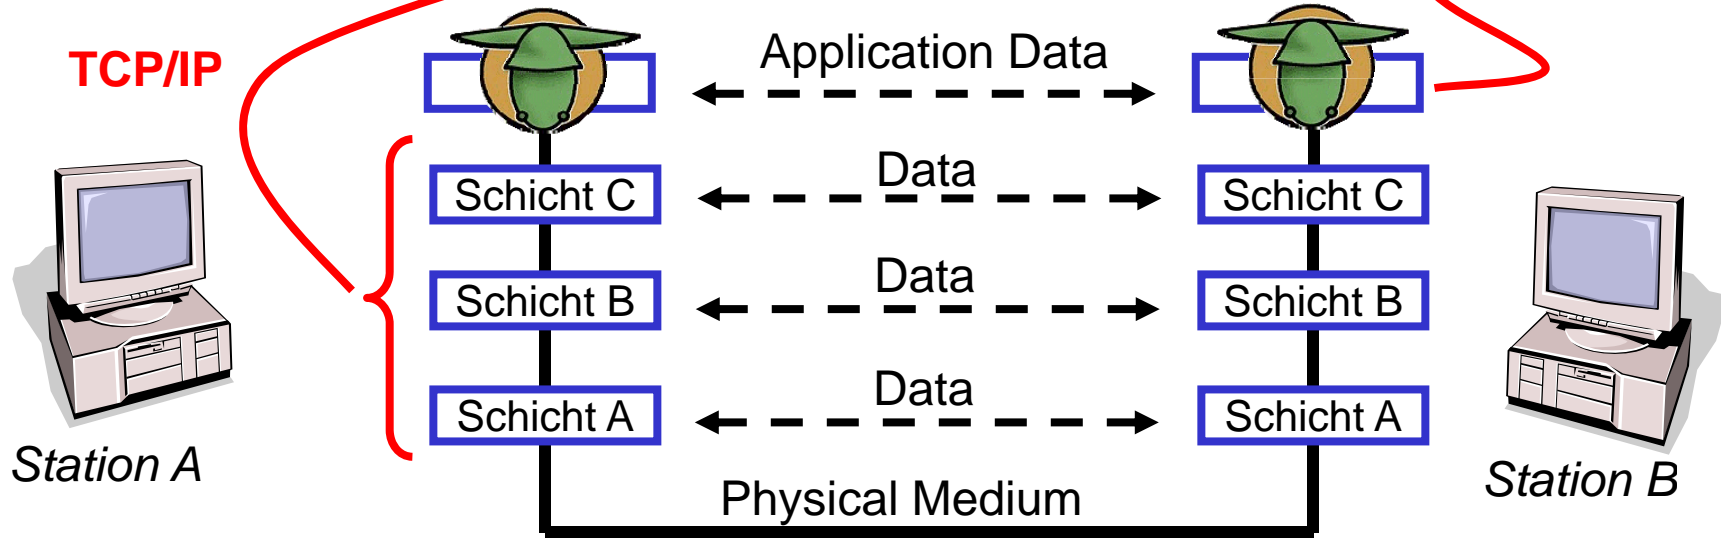
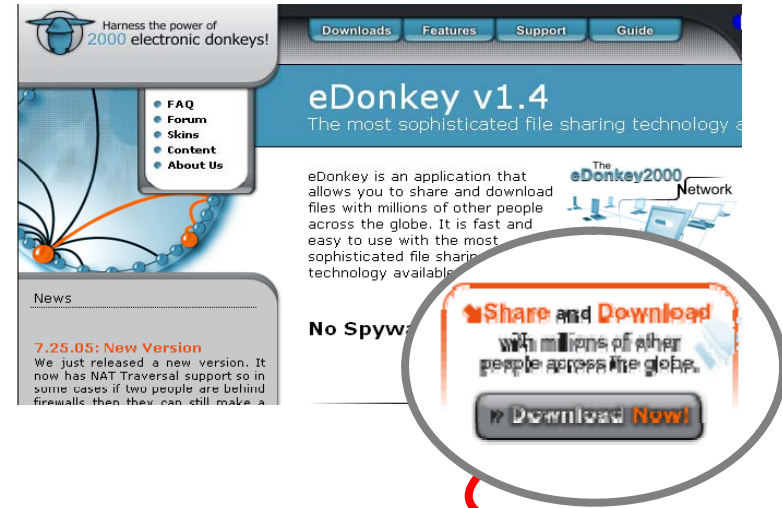

1. Motivation

2. Peer-to-Peer Overlays



3. The Igor File System

4. Outlook & Summary

What's Peer-to-Peer? (1)



What's Peer-to-Peer? (2)



Peer-to-Peer ...

Organic Growth

- Each machine contributes.
- Joining devices provide the resources they consume


Robust and Fault-Tolerant

- No single point of failure

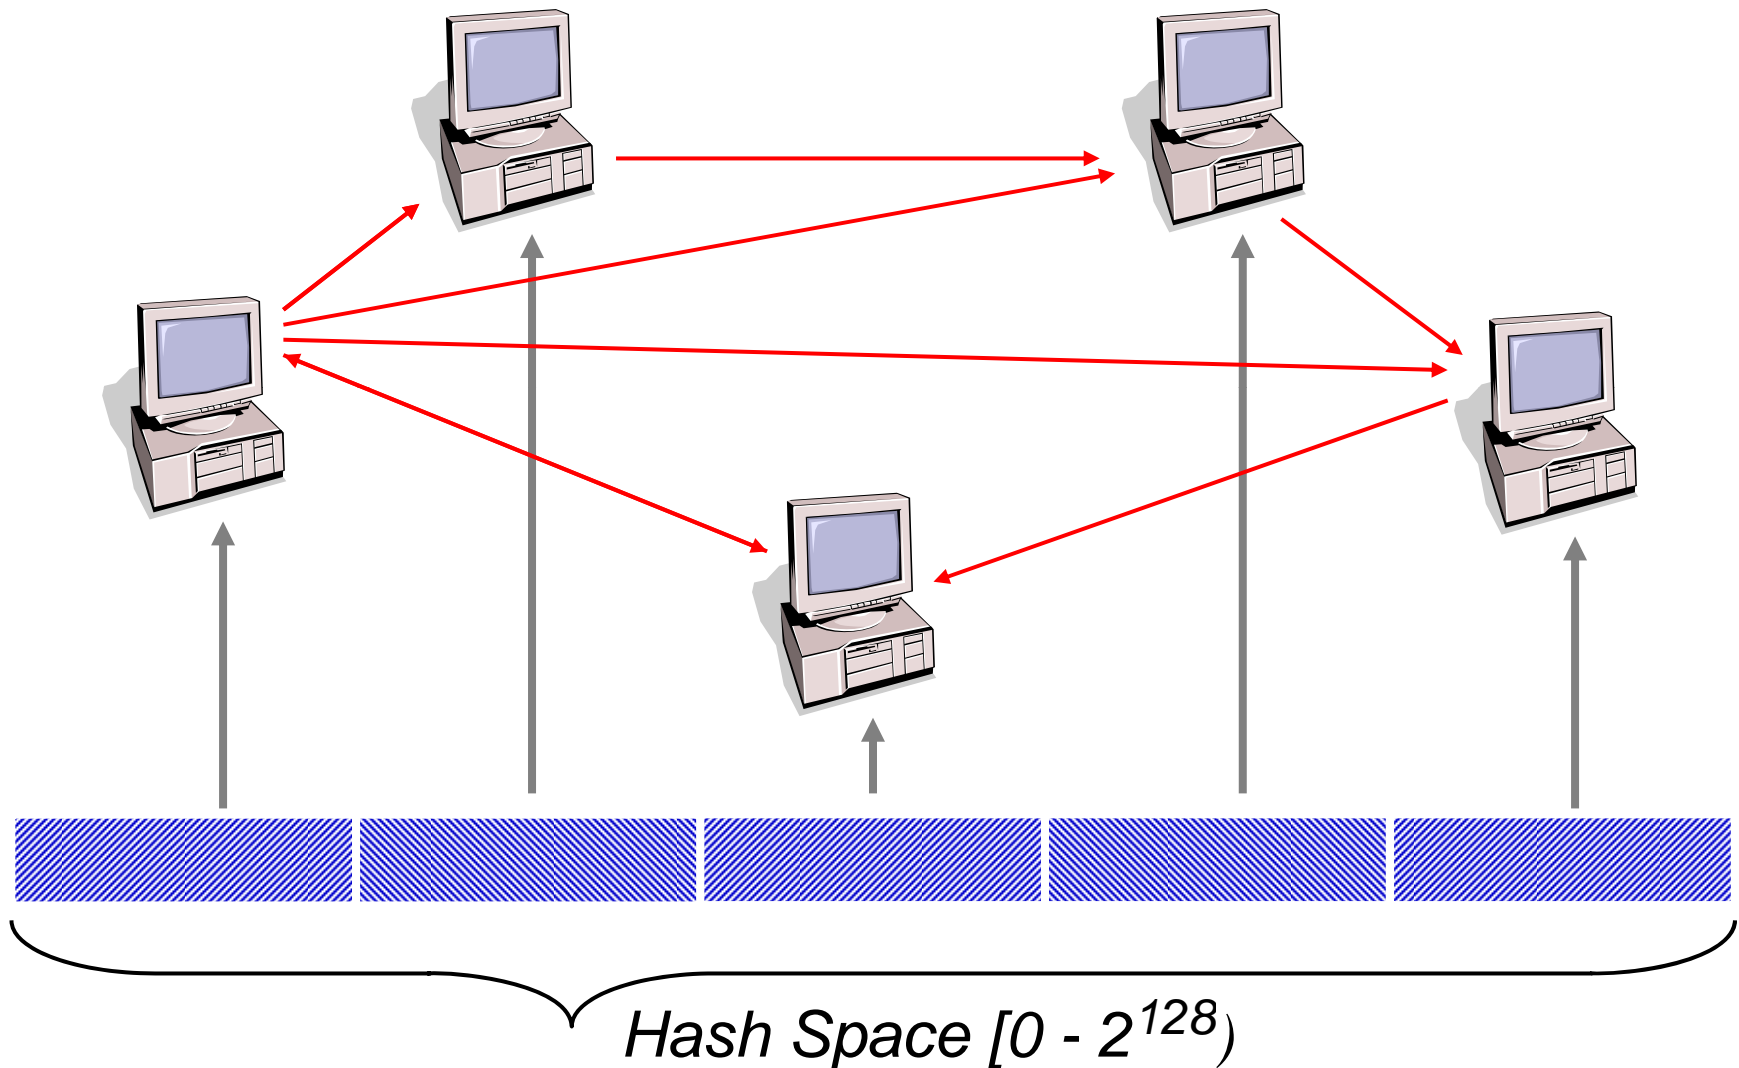
Research Question:

- Which rules make the peer-to-peer system behave as desired?

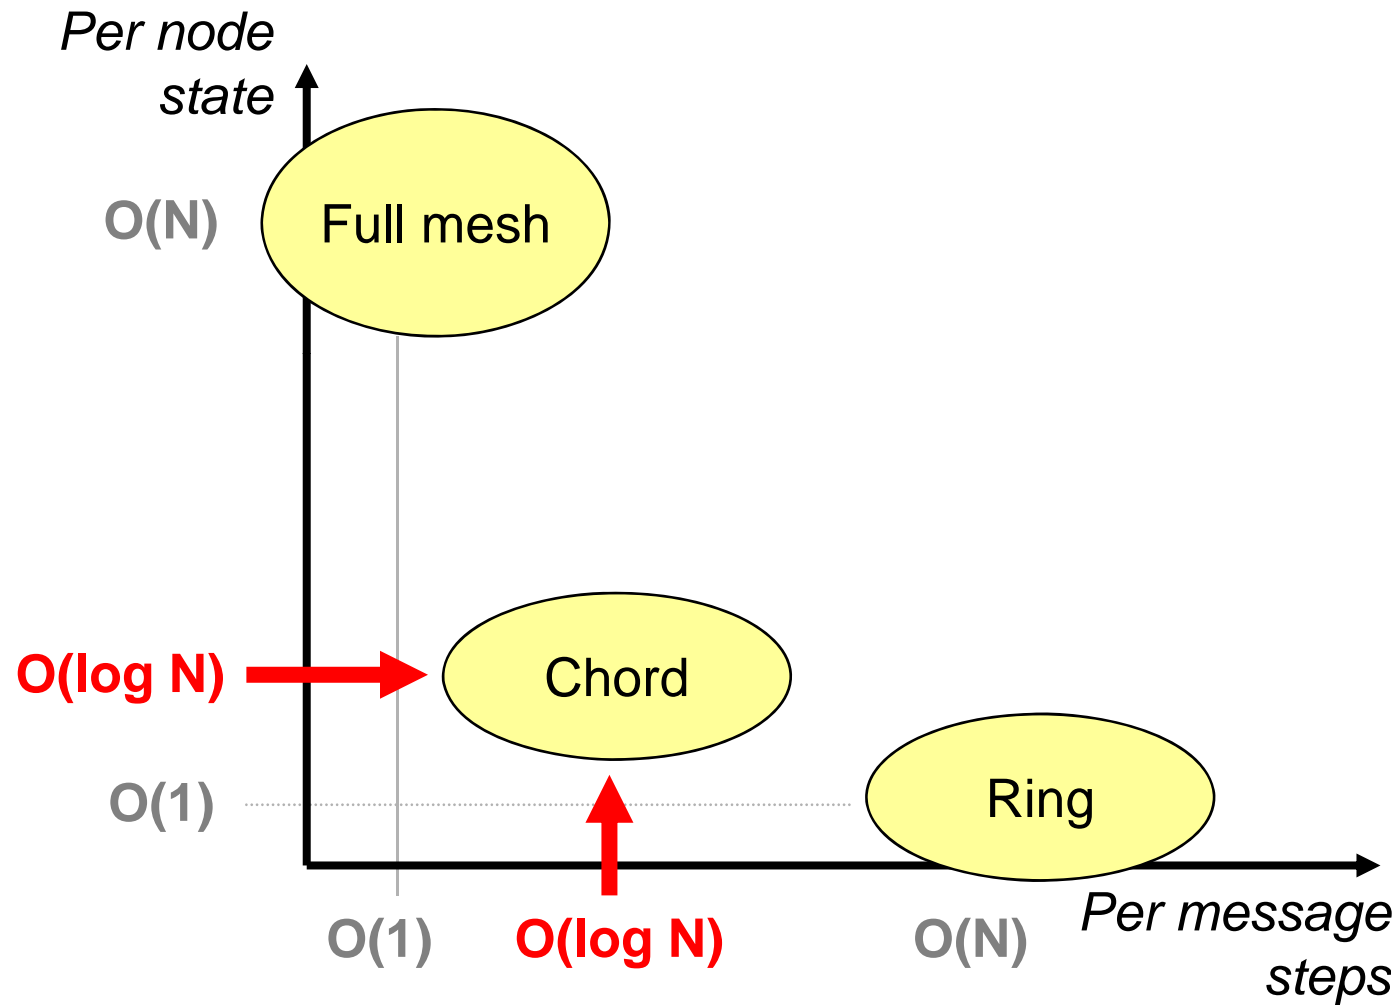
Distributed Hash Tables



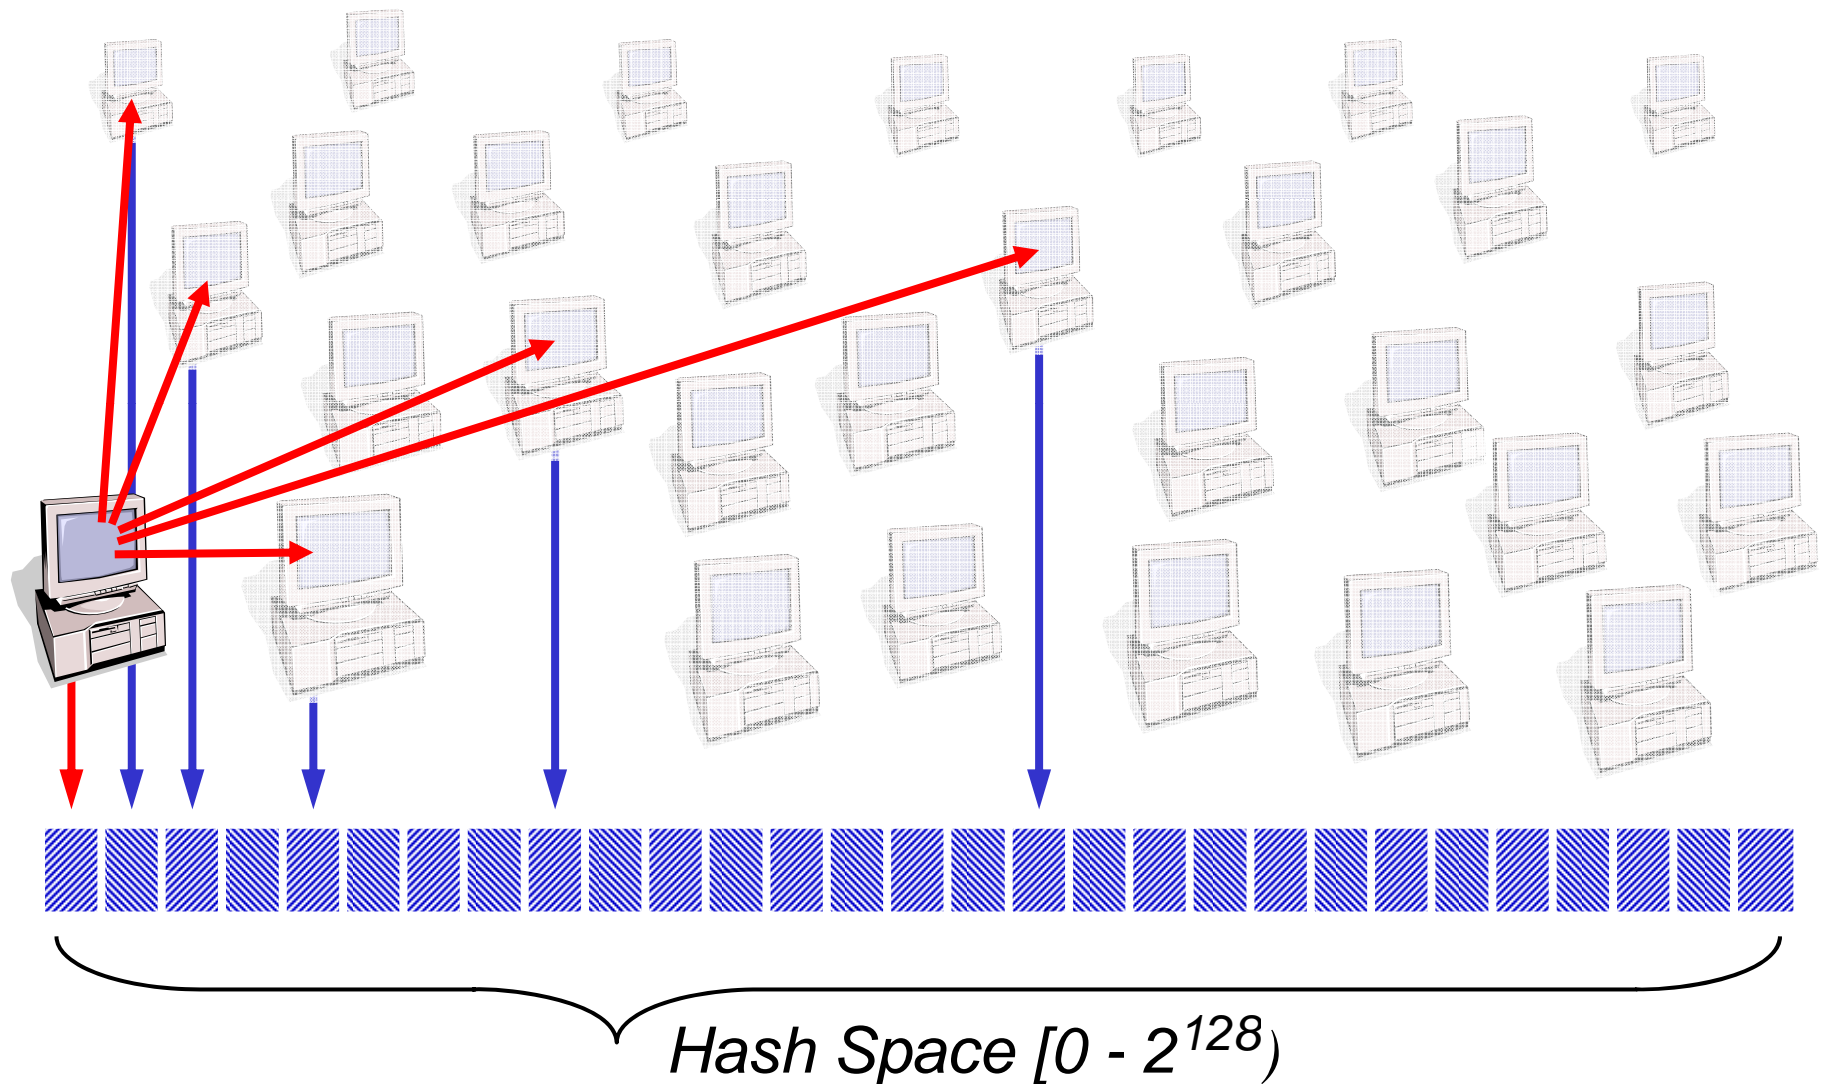
Distributed Hash Tables




Trade-Off with Structured Routing Overlays




Chord – A Distributed Hash Table





Chord – A Structured Routing Overlay



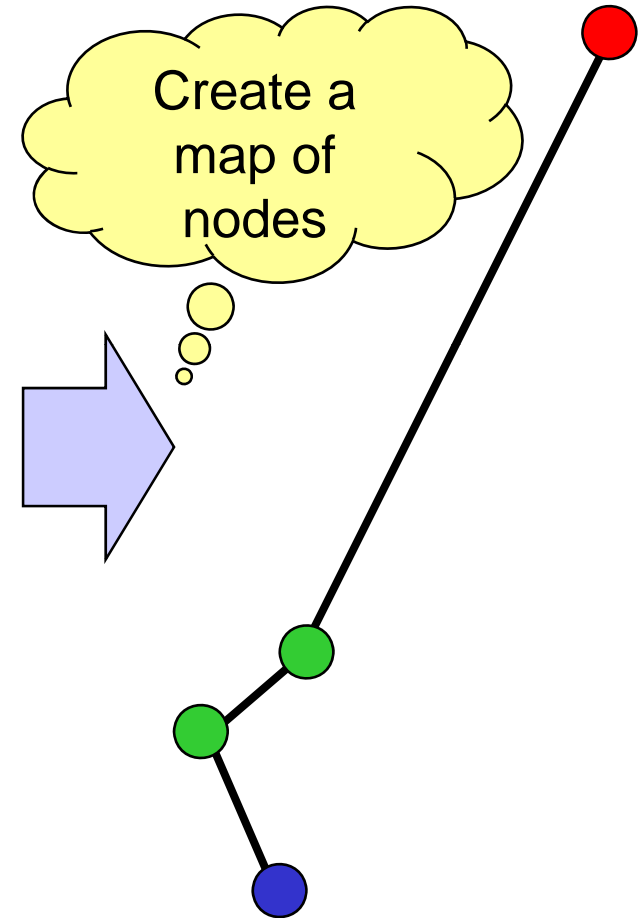
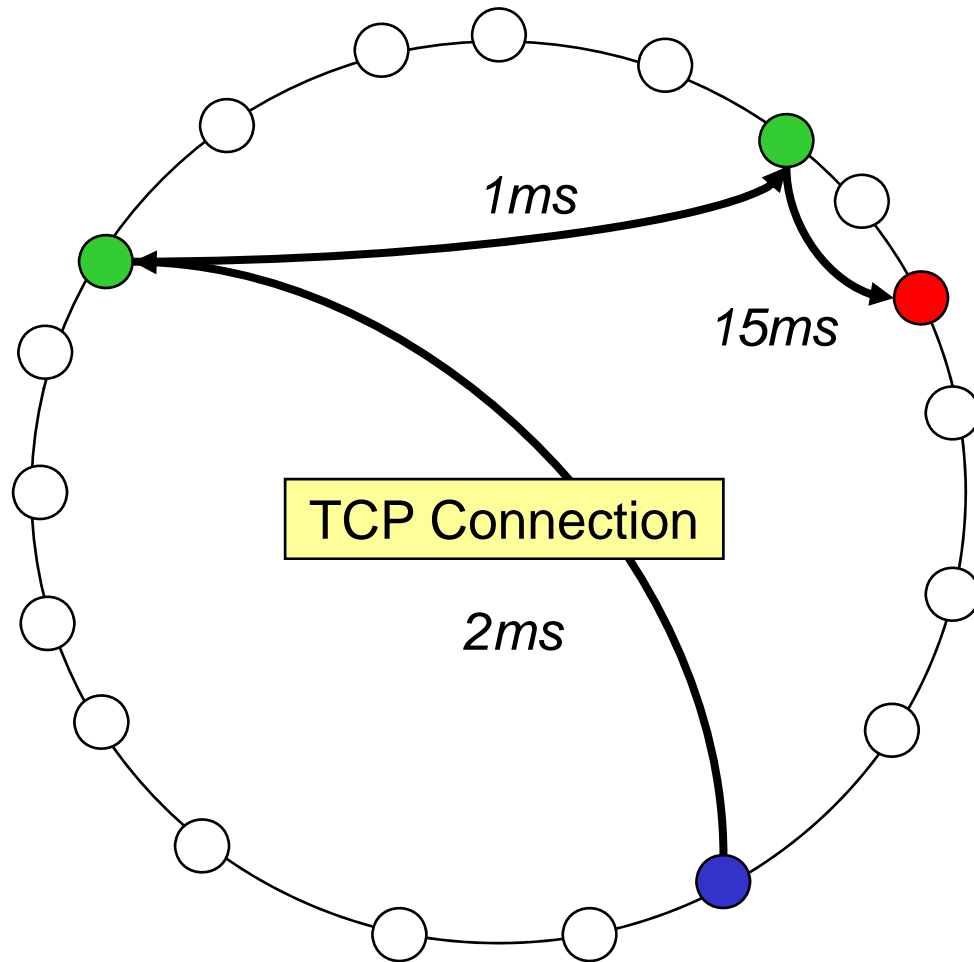
Chord – A Structured Routing Overlay (2)



Proximity Awareness with Chord



Network Coordinates



Agenda

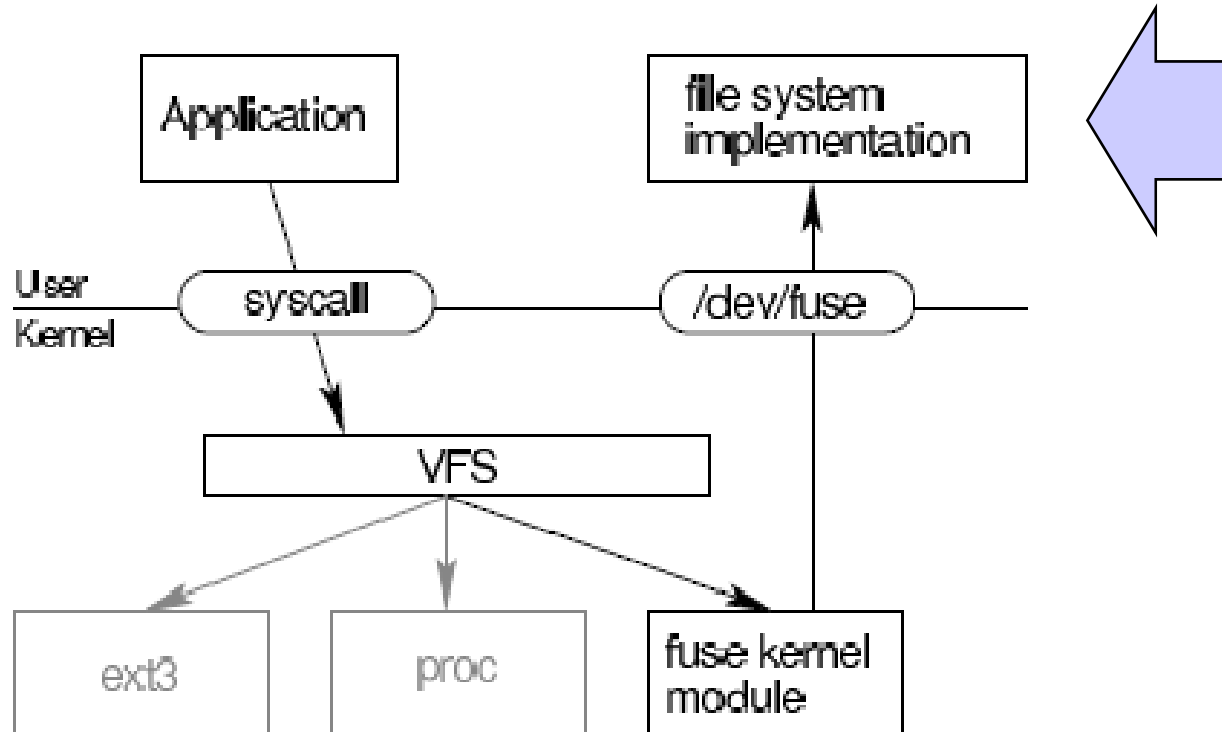
1. Motivation

2. Peer-to-Peer Overlays

3. The Igor File System

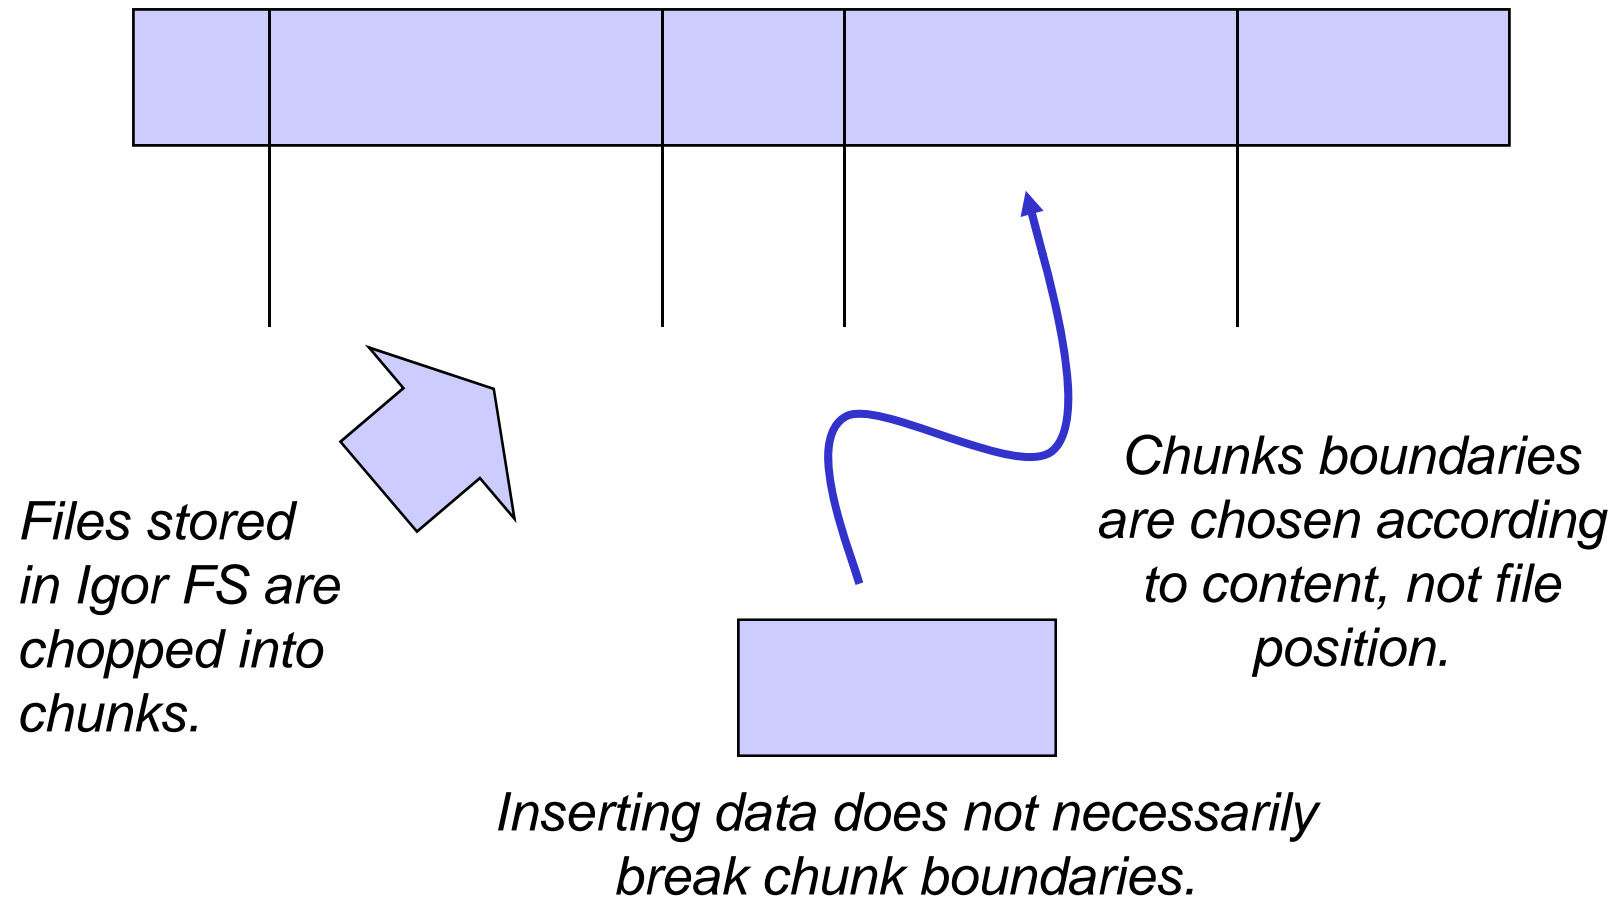
4. Outlook & Summary

Igor File System

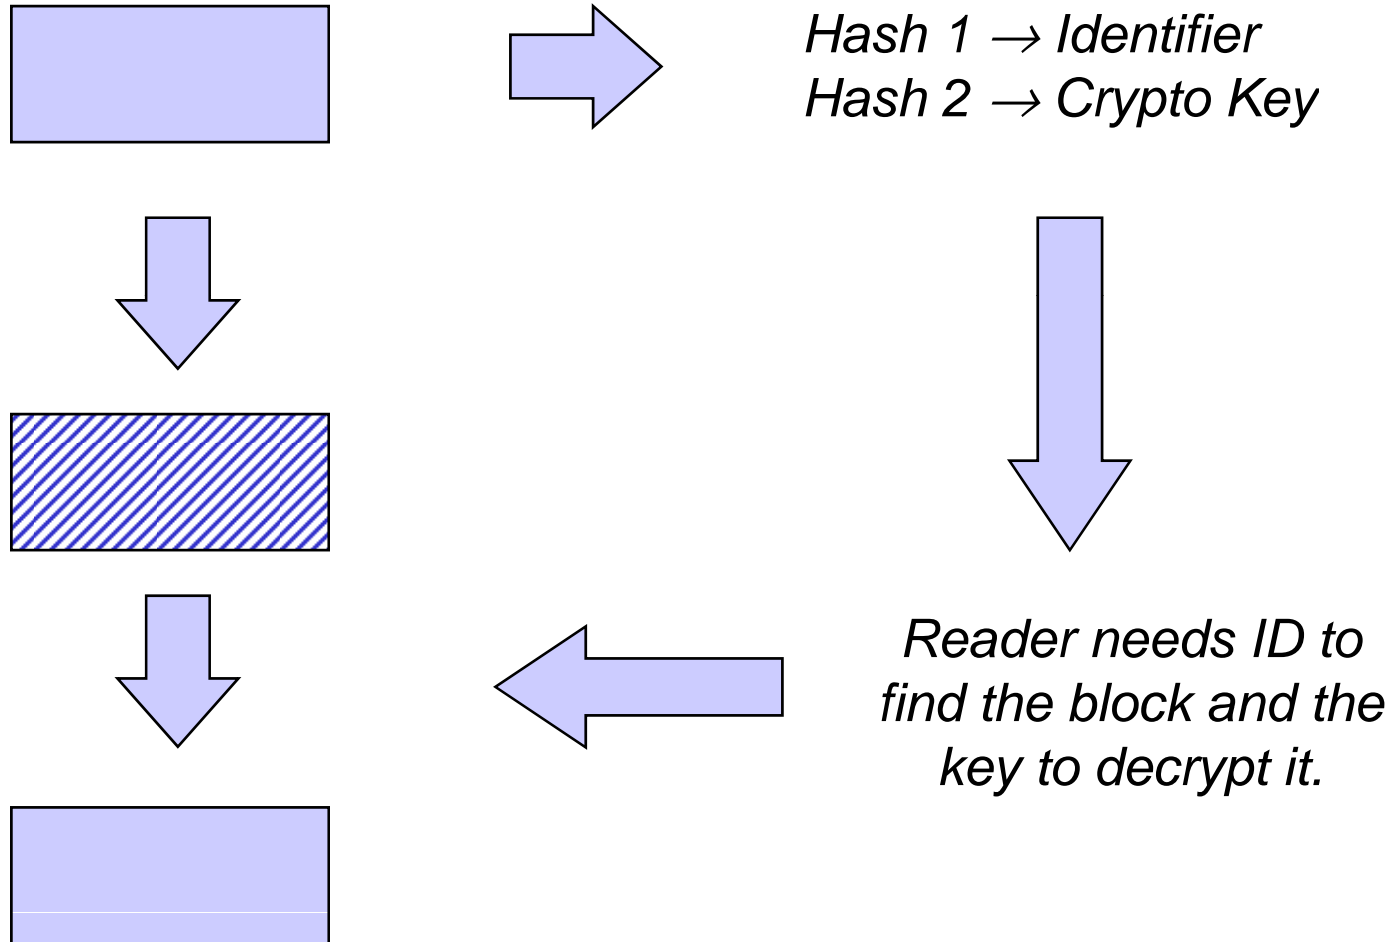


Igor FS sits on top of FUSE. Thus it appears as normal file system to the operating system.

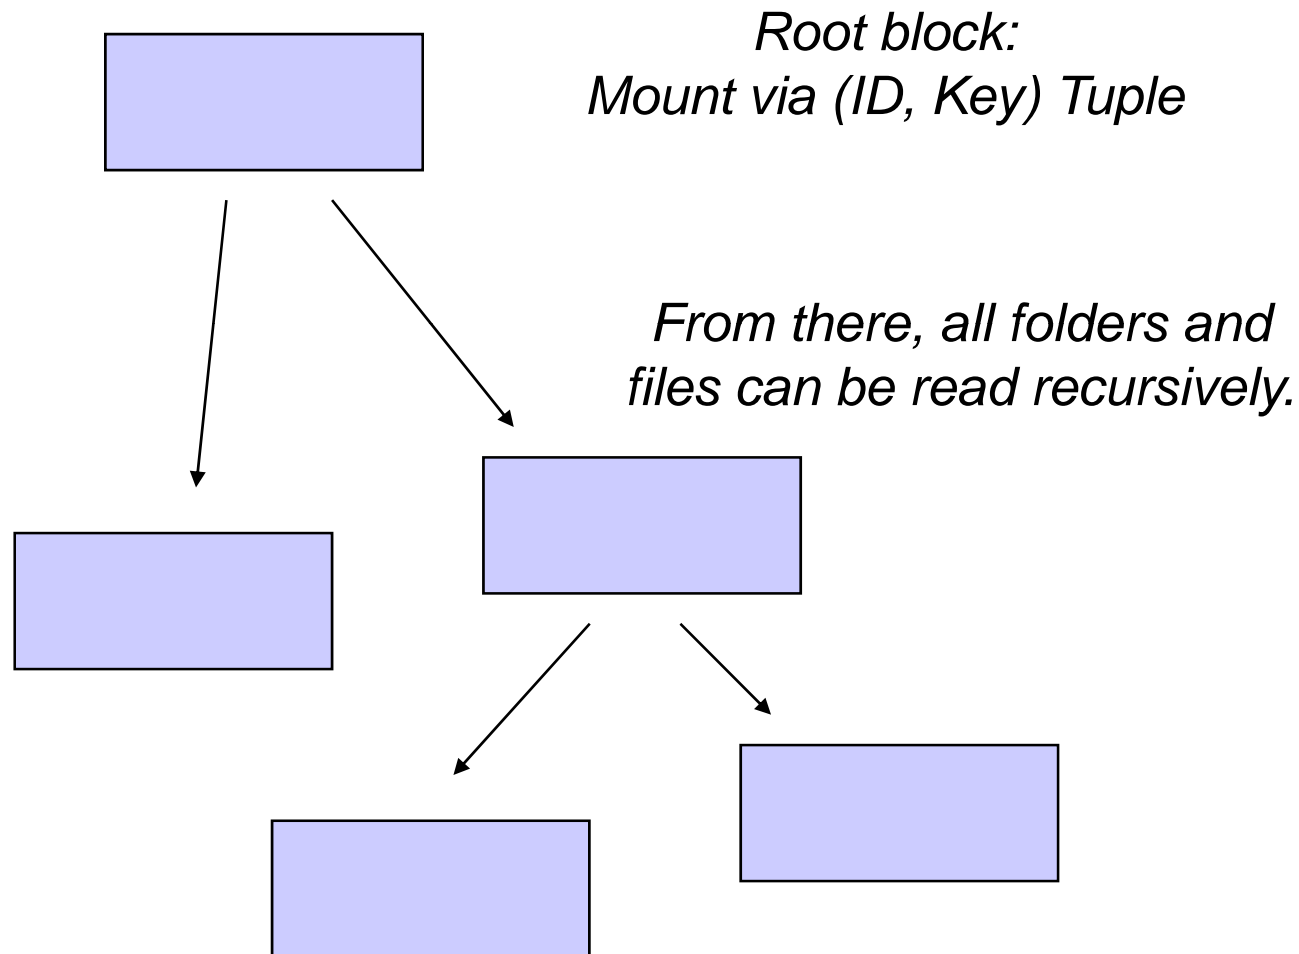
Chopping up Files



Encryption of Data



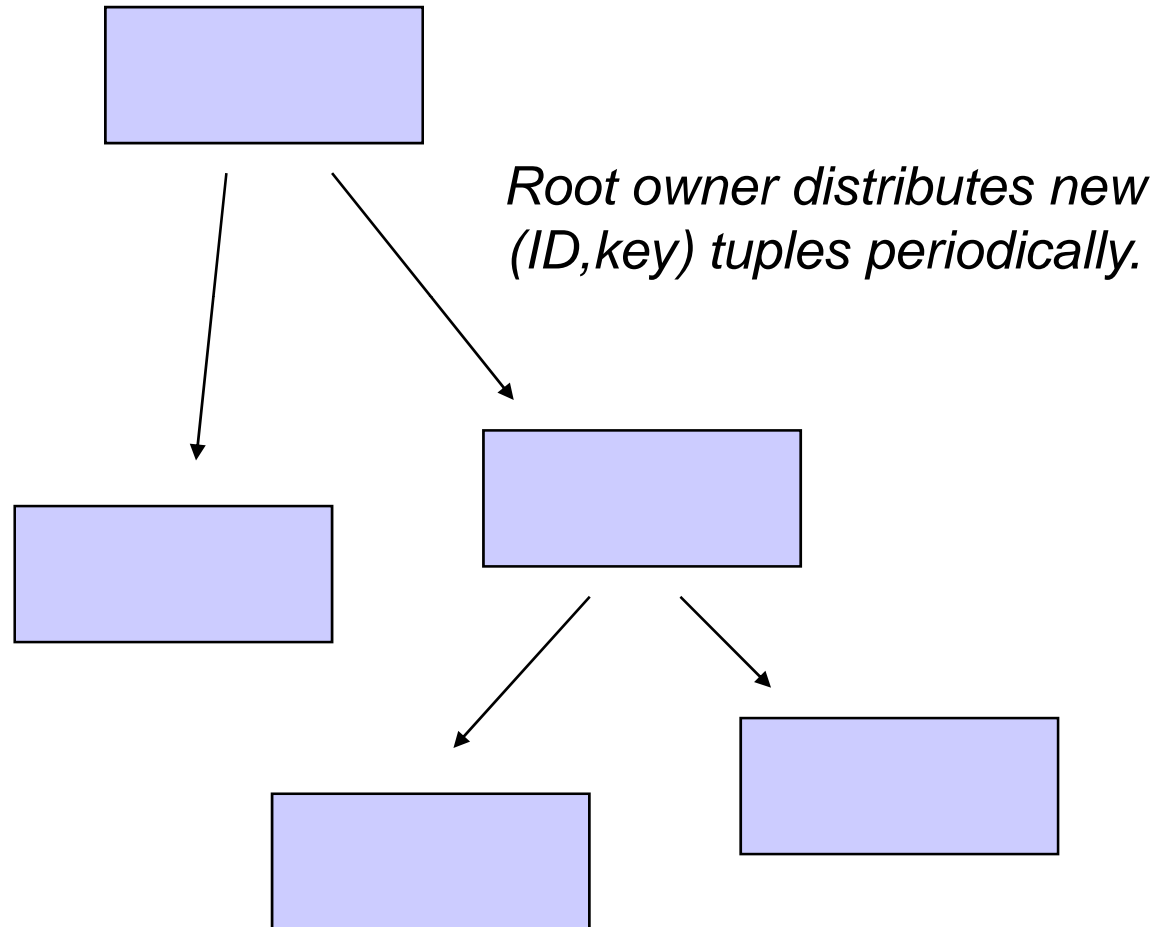
File System Structure



Backup and Versioning

- Blocks are identified by their hashed content.
- Thus, modifying the block, changes the ID.
- Thus, blocks are immutable.
- Writing to Igor FS creates a new root block.
- This means that nothing gets lost as long as you have the (ID,key) tuple.

Root Key Distribution



- Caveat:
 - So far, Igor FS is a prototype system only (PhD thesis project).
 - It is being used in an EU project by several pharma research institutions.
- Reading (block-wise) 23.5 MByte/s
- Writing (block-wise) 87.6 MByte/s
 - Note that actual persistence on hard disk and distribution in the network is asynchronous. Thus data rates drop only when the IgorFS cache memory of the machine is exhausted.
- Note that
 - NFS in the same setting (local LAN) has read/write performance of about 20 MByte/s
 - EXT3 on a local hard disk has 100MByte/s and above
- Conclusion:
 - IgorFS already levels with other distributed file systems.
 - But it can not and will not outperform local file systems.

Agenda

1. Motivation

2. Peer-to-Peer Overlays

3. The Igor File System

4. Outlook & Summary

Outlook to Future Enhancements


- So far, only one writer.
 - In future: Multiple writers! – Need to solve consistency issues.
- So far, data pulled by the reader.
 - In future: Important data pushed by the writer to ensure redundant copies exists as soon as possible.
- So far, all data kept forever.
 - In future: Only that data kept that readers will request in future.

Outlook to Future Enhancements

- Compute processes are just data:
 - Code, heap, stack, execution frames, ...
- Our Igor system move data to the place where it is needed.
- Can it move around threads, too?
- We believe, it can!

- Ongoing project to use our technology in sensor actuator networks.
- New project to use this technology with IBM cell processor.
(Grant currently under negotiation)

IGOR-FS Feature Summary




IGOR-FS works automatically and fully transparently as a distributed file system – *as if you had magically already downloaded everything you need.*



Distributed storage



Strong cryptography



Backup & versioning

Thank you!

Questions?



Thomas Fuhrmann

CS VIII - Network Architectures
Technical University Munich, Germany

IBDS System Architecture
University of Karlsruhe, Germany

fuhrmann@net.in.tum.de