AIFB

Peer-to-peer (P2P)

Steffen Staab
Institut AIFB
Universität Karlsruhe (TH)

AIFB

What is P2P ?

- In a Peer- to-Peer network, end users share resources via direct exchange between computers.
- Information is distributed among the member nodes instead of concentrated at a single server.
- A pure peer to peer system is a distributed system without any centralized control, where the software running at each node is equivalent in functionality.

AIFB

Exchanging Content in P2P Networks: Vision

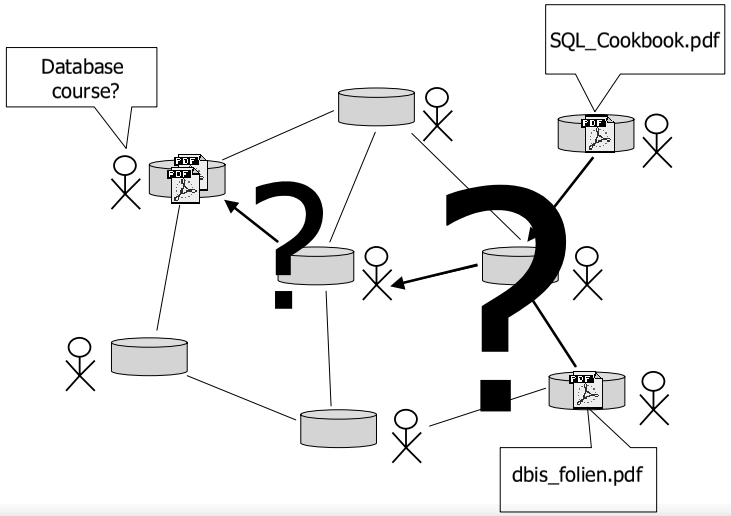

AIFB

Advantage and Disadvantage of P2P in Comparison with Client-Server

- Advantages:
  - Inherent Scalability
  - Availability of More Information
- Disadvantages:
  - No Guarantee about Quality of Service

AIFB

P2P Application Areas

- Communication
  - AOL Instant Messenger, ICQ
- Remote Collaboration
  - Shared File Editing
  - Audio-video Conferencing.
- Distributed Computing
  - SETI@home Folding@home
- File Sharing
  - Napster
  - Gnutella, Freenet
  - KazaA, Morpheus

AIFB

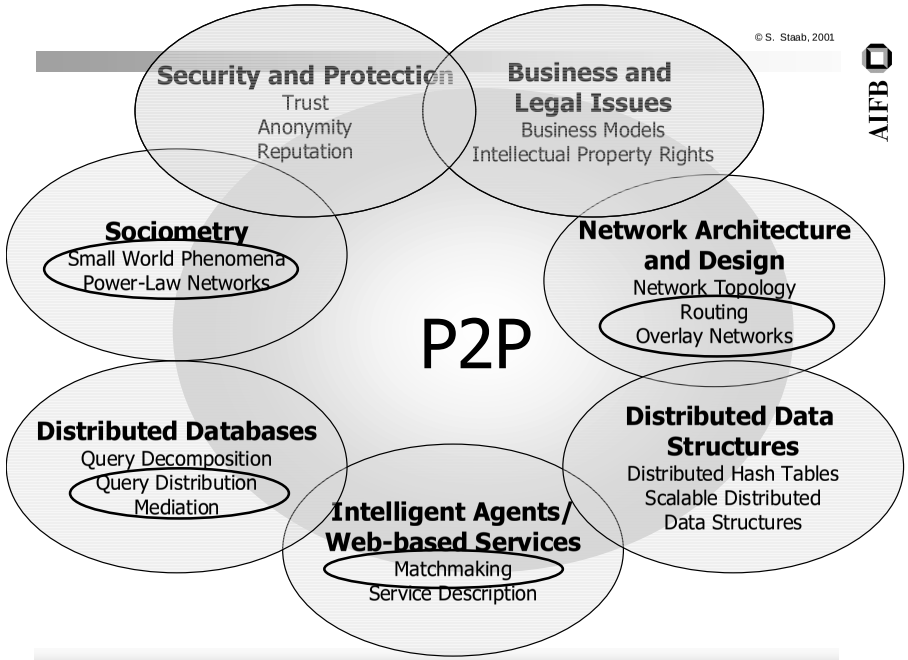

AIFB

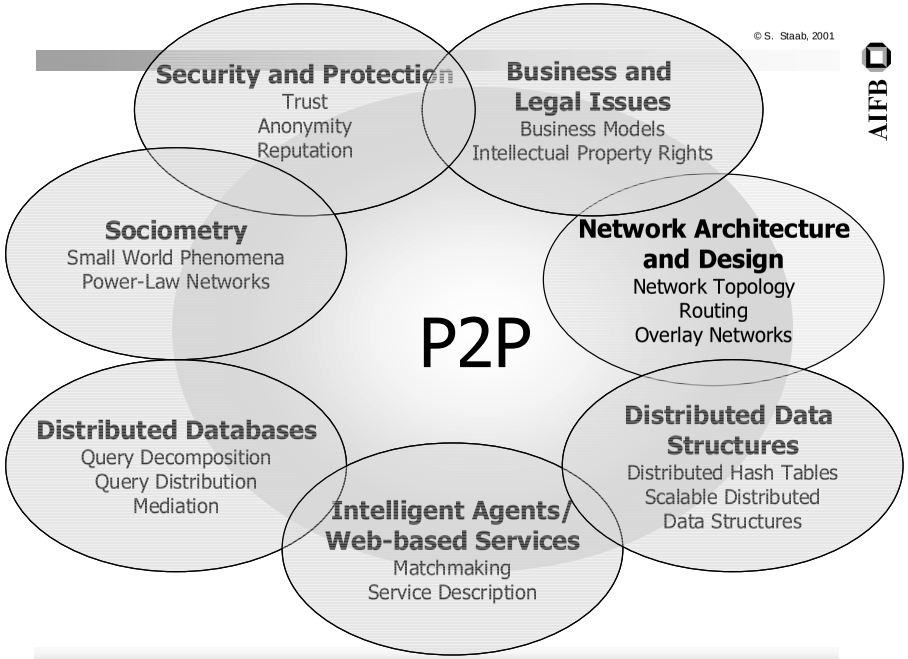

AIFB

Routing: Current State in P2P

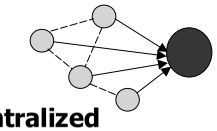

Centralized
- One or few central coordinator(s)
- e. g. Napster, Instant Messengers

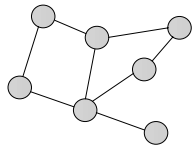

Fully Decentralized
- All peers (or none) contain routing information
- e. g. Freenet, Gutnella

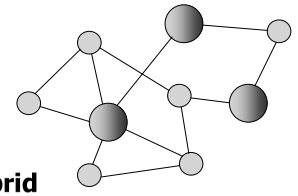

Hybrid
- Some superpeers carry indexing information
- e. g. FastTrack (Kazaa, Morpheus), Gnutella derivates

AIFB

Architecture of Existing P2P File Sharing System (1)

- Hybrid Centralized Peer-to-Peer File Sharing System: Server facilitates the interaction between peers by maintaining directories of the shared files stored on the respective PCs of registered users of the network.
  - Napster

AIFB

Architecture of Napster

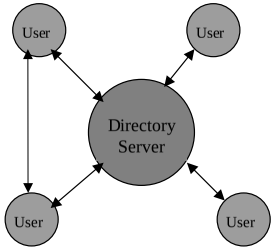

A central directory server maintain index on the metadata of all the files in the network. The metadata might include file names, creation dates, and copyright information . The server also maintain a table of user connection information including user's IP address and line speed. A file query is sent to the server first. A query consists of a list of desired words.

When the server receives a query, it searches for matches in its index. The query results including a list of users who hold the file are sent back to the user who initiated the query. The user then opens a direct connection with the peer that has the requested file for downloading

AIFB

A query in Napster

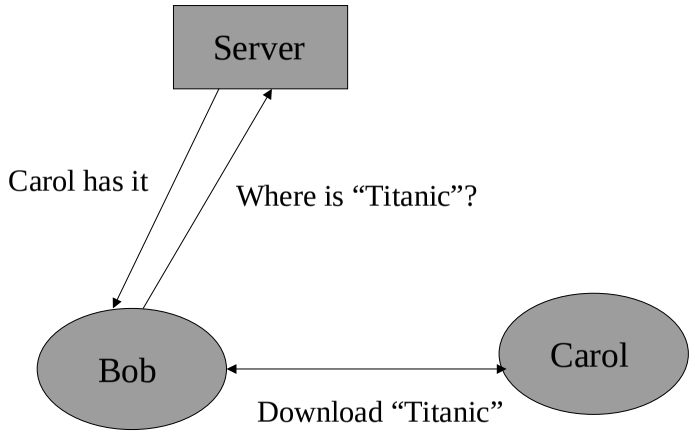

AIFB

Architecture of Existing P2P File Sharing System (2)

- Pure Decentralized Peer-to-Peer File Sharing System: Peers have same capability and responsibility. The communication between peers is symmetric. There is no central directory server Index on the metadata of shared files is stored locally among all peers.
  - Gnutella
  - Freenet
  - FreeServe
  - MojoNation

AIFB

Architecture of Gnutella and Freenet

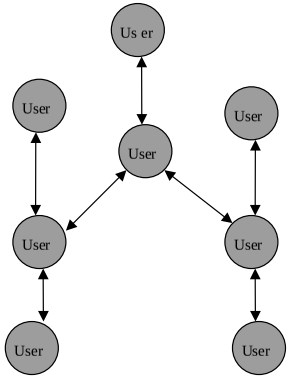

Each user acts independently.
There is no central directory server.
Peers directly change file information among each other in a query.

A query in a decentralized file sharing system

AIFB

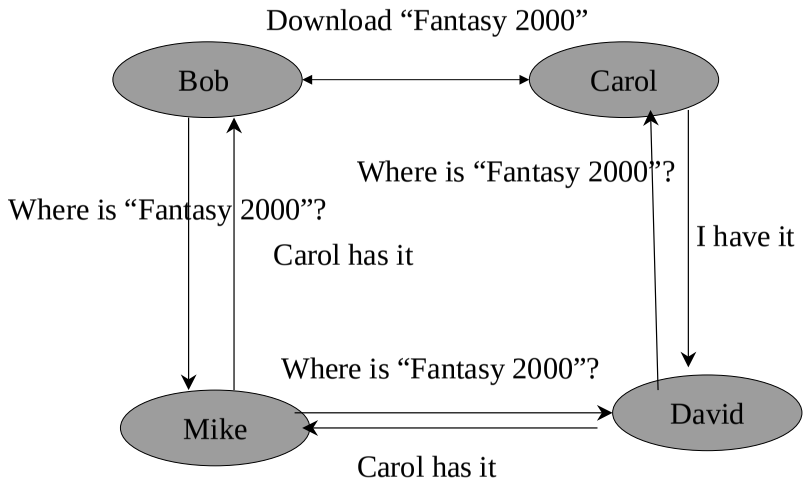

AIFB

Advantages and Disadvantages of Centralized Indexing

- Advantages:
  - Locates files quickly and efficiently
  - Searches are as comprehensive as possible
  - All users must registered to be on the network
- Disadvantages:
  - Vulnerable to censorship and technical failure
  - Slashdot effect: popular data become less accessible because of the load of the requests on a central server
  - Central index might be out of data because the central server’s database is only refreshed periodically.

AIFB

Advantages and Disadvantages of Decentralized Indexing

- Advantages:
  - Inherent scalability
  - Avoidance of “single point of litigation” problem
  - Fault Tolerance
- Disadvantages:
  - Slow information discovery
  - More query traffic on the network.

AIFB

Architecture of Existing P2P File Sharing System (3)

Partially centralized indexing system (hybrid):

A central server registers the users to the system and facilitates the peer discovery process. After a Morpheus peer is authenticated to the server, the server provides it with the IP address and port (always 1214) of one or more "SuperNodes" to which the peer then connects.

Local SuperNodes" index the files shared by local peers that connected to it and proxy search requests on behalf of these peers.

- KazaA
- Morpheus

AIFB

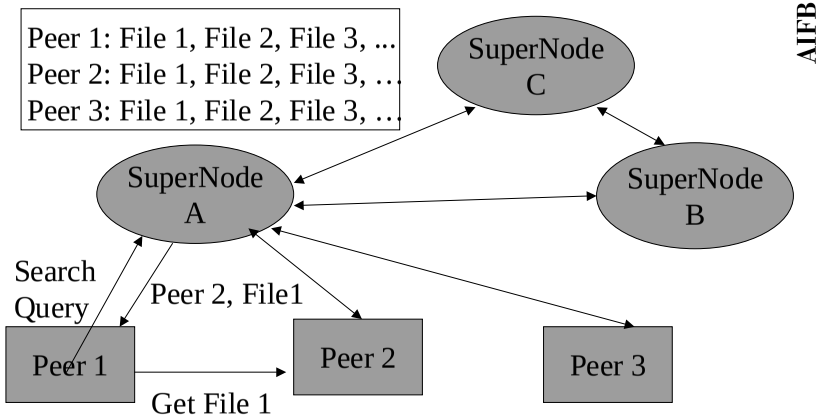

Search results in Morpheus contain the IP addresses of peers sharing the files that match the search criteria, and file downloads are purely peer-to-peer.

AIFB

Morpheus's SuperNode

- Morpheus peers are automatically elected to become SuperNodes if they have sufficient bandwidth and processing power (a configuration parameter allows users to opt out of running their peer in this mode).
- Once a Morpheus peer receives its list of SuperNodes from the central server, little communication with the server is required.
- While Morpheus is largely a decentralized system, the speed of its query engine rivals that of centralized systems like Napster because of its SuperNode.

AIFB

Advantages of Partial Centralized Indexing

- Reducing discovery time in comparison with purely decentralized indexing system such as Gnutella and Freenet
- Reducing the workload on central servers in comparison with fully centralized indexing system such as Napster.

AIFB

File Discovery Mechanisms in Freenet: Chain Mode

- A request for a file can be forwarded through many different nodes.
- If a node does not have the document that the requestor is looking for, it forwards the request to one of its neighbor that is more likely to have the document. The messages form a chain as each node forwards the request to the next node.
- Message times out after passing through a certain number of nodes, so that huge chains don't form.
- The chain ends when the message times out or when a node replies with the data.

AIFB

File Discovery Mechanisms in Freenet: Chain Mode

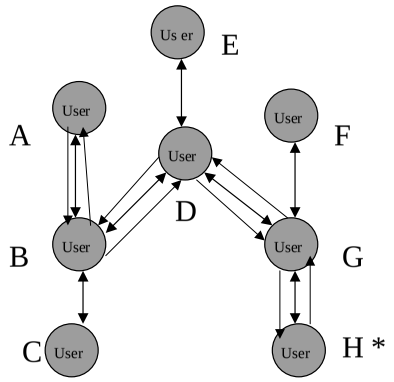

Computer A send its query to its Neighbor B, which forwards it to one of its neighbor D, which in turn forwards it to one of its neighbor G, finally, the query reach H who has the requested data. The reply is passed back through each node that forwarded the request, back to the original node that started the chain.

AIFB

Advantages and Disadvantages of Chain Mode in Searching

- Advantages:
  - Fast discovery with less network traffic in average case
  - Search stops once the requested file is found
  - More scalable
- Disadvantages:
  - Slow discovery in worst case

AIFB

Worst Case Scenario in Chain Mode Discovery

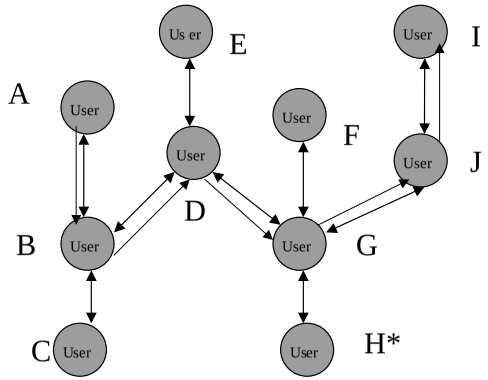

Computer H has the requested file, but the query never reaches it because of poor routing decision at computer G. instead, the query ends Up with no file found.

AIFB

File Discovery Mechanisms in Gnutella: Broadcast Mode

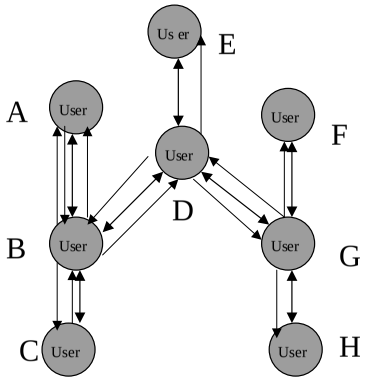

Computer A sends query request to B, which forwards the query to its neighbors C and D, which in turn forward the query to all neighbors. If a computer has a file that matches the request, it transmits the information (file name, size, etc) back through all the computers in the pathway towards A. In this example, both G and C have the requested file and both reply to A.

AIFB

Advantages and Disadvantages of Broadcast Mode in Search

- Advantages:
  - More robust
  - More comprehensive
- Disadvantages:
  - More network traffic and less scalability
  - Possible loops
  - Search continues even after the requested file is found in some nodes.

AIFB

Overlay Networks

- Layer on top of existing infrastructure
- Cope with
  - Masquerading/Dynamic IPs
    - **One address may represent many nodes**
    - **One node may appear under many addresses**
  - Firewalls
    - **Only certain types of traffic allowed**
- Employ your own
  - Adressing scheme
  - Topology
  - Routing mechanism

AIFB

Overlay Networks (II): JXTA

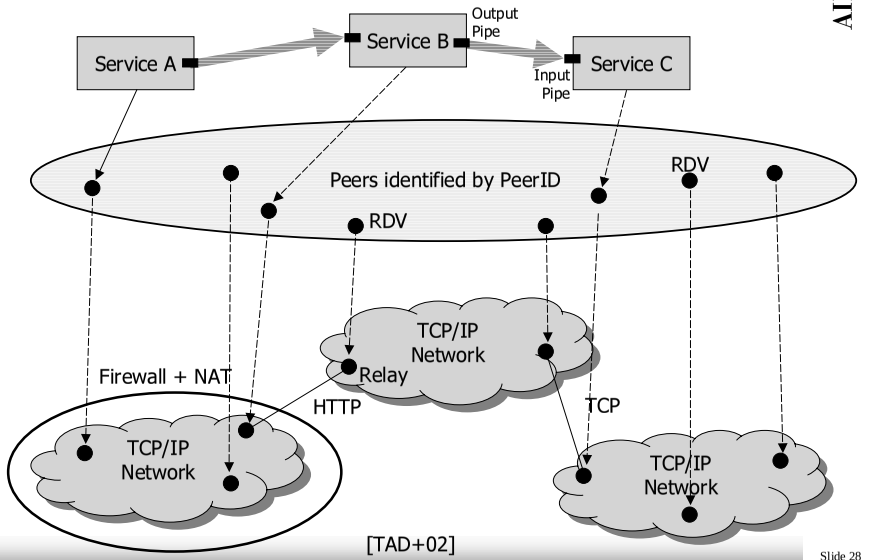

AIFB

Routing: Aggregation

- Aggregation of information about networks
  - Smaller routing tables

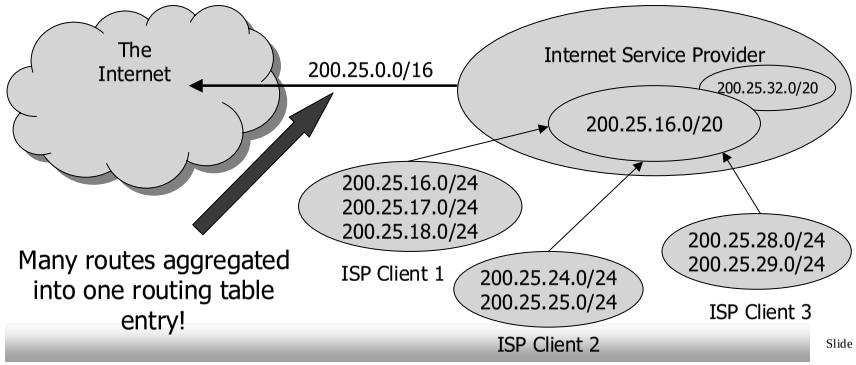

AIFB

Routing (speculation!): Towards semantic routing

- Routing Indices
  - Maintain aggregated view of term frequency count behind each link

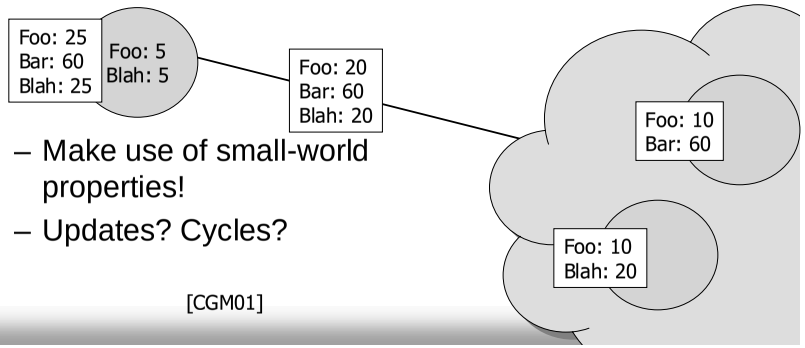

  - Make use of small-world properties!
  - Updates? Cycles?

[CGM01]

AIFB

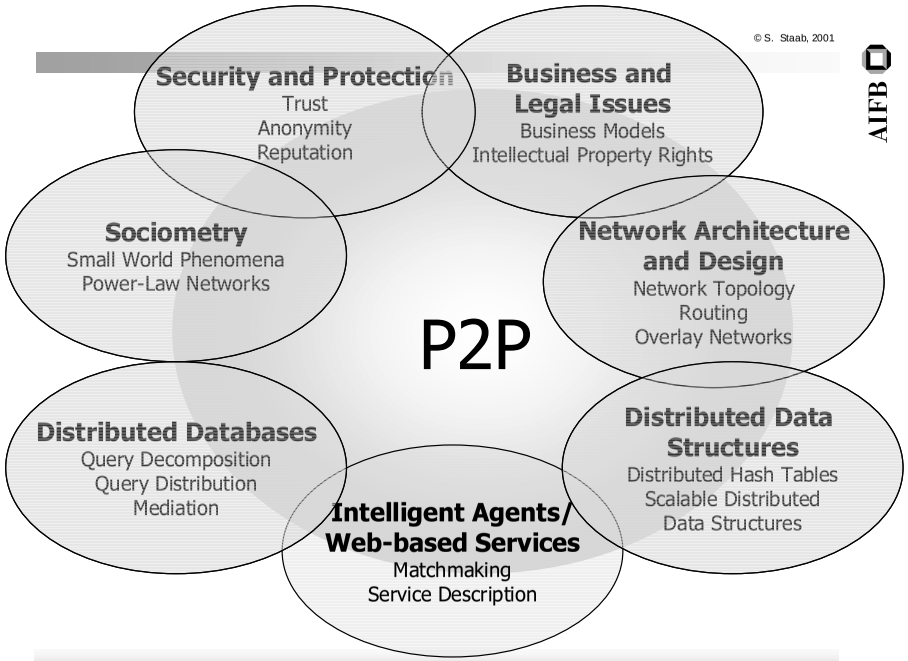

AIFB

Semantic Matchmaking

- Find agent that is capable of fulfilling a given task
- Example: LARKS
  - Service described by:
    - **Context: "Sorting"**
    - **Input/output parameters: "In i: list of real; Out o: list of real"**
    - **Input/output constraints: "len(i) <= 100, sorted(o)"**
    - **Text description: "Sorts a list of at most 100 reals"**
    - **[...]**

[SLKW99]

AIFB

Semantic Matchmaking (II)

- Match description of wanted task against available agents
- Several kinds of filters, e. g.:
  - **Similarity matching**
    - Test ontology-based similarity of words in description
  - **Signature matching**
    - Test if input and output types of descriptions are equal or subconcepts of each other
  - **Constraint matching**
    - Test if one set of constraints subsumes the other one

AIFB

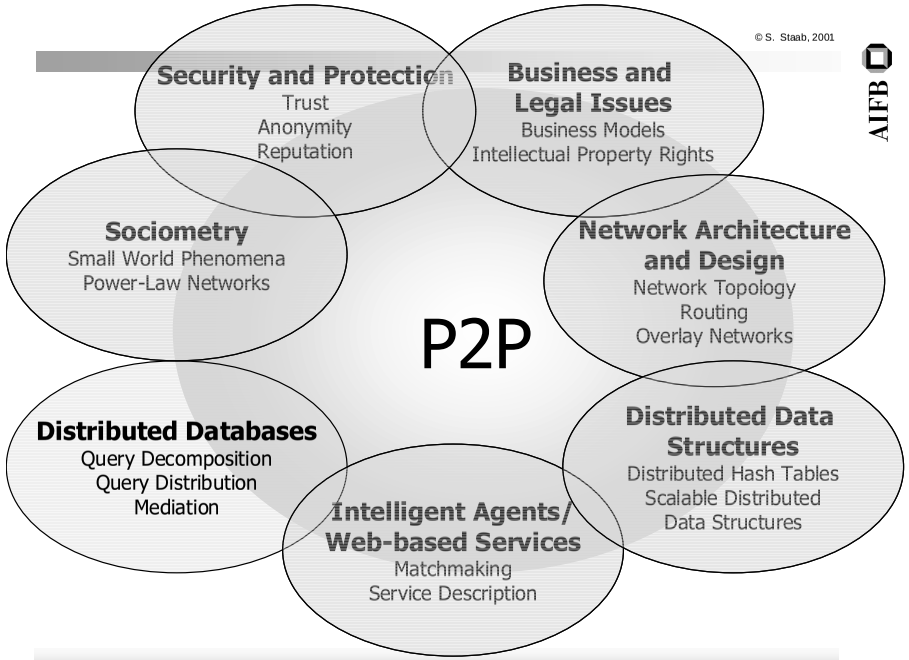

AIFB

Distributed Databases: Mediation

- AMOS II
  - Wrapper-Mediator Approach for Distributed DBMS
    - **Wrappers map data sources to common data model**
    - **Mediators distribute queries and integrate results**

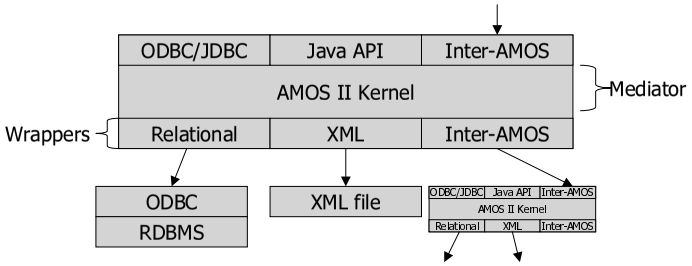

[RJ01]

AIFB

Distributed Databases (II)

- Query decomposition
  - **Predicate grouping**
    - group query predicates by node
  - **Predicate placement**
    - cost-based heuristics
    - select execution node for each predicate
  - **Cost-based scheduling**
    - compute execution tree
  - **Tree balancing and distribution**
    - optimize execution tree

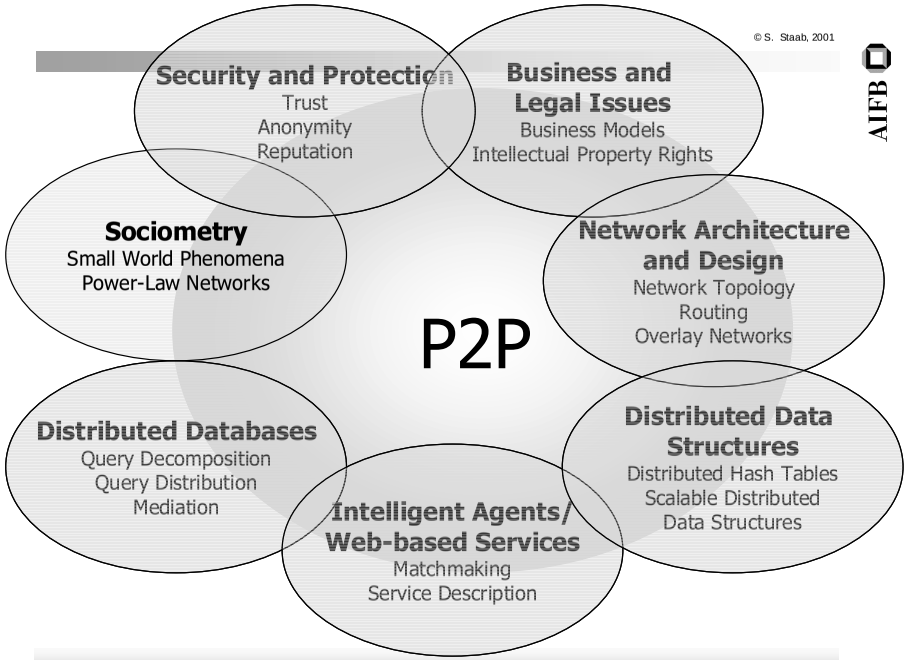

Small-World Phenomena and Power-Law Graphs

- Milgram's six degrees of separation (1967): "It's a small world"
  - Forwarding of letters from Nebraska to Massachusetts:
    - **Forward message to someone "closer" to the target**

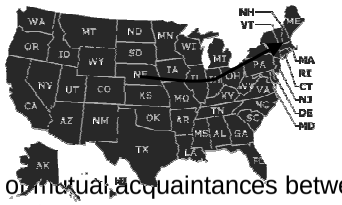

  - Average chain of mutual acquaintances between two Americans has average length 6

[WS98, FFF99]

Small-World Phenomena and Power-Law Graphs (II)

- Power-Law Graphs
  - P[node has degree k] ~ $\frac{1}{k^{\alpha}}$ for some α>0
- Found in many real-life situations
  - Neural network of worms
  - IMDb collaboration
  - Web Graph
  - AT&T Call Graph
  - Gnutella

Small-World Phenomena and Power-Law Graphs (III)

- Highly clustered, short paths
  - "short cuts": long range edges
- Milgram Experiment:
  - High-degree nodes are crucial for short paths

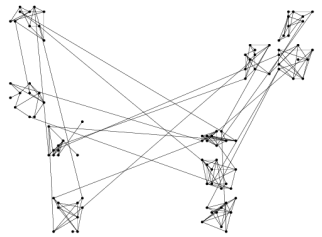

AIFB

Edutella: Project Context

- Submodule of PADLR Project (Personalized Access to Distributed Learning Repositories) (www.learninglab.de/english/projects/padlr.html)

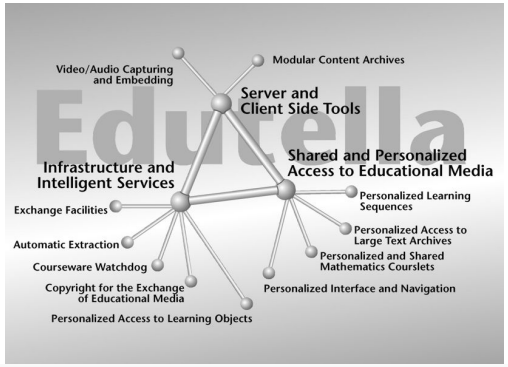

Overview of PADLR Project (Stanford, Hannover, Braunschweig, Karlsruhe, Stockholm, Uppsala)

AIFB

PADLR: Basic Assumptions

- Assumption 1/2: lots of learning resource repositories, which typically employ various back-ends, various meta-data schemas, and various architectures, etc., have already existed in many institutions. (trouble: isolate information islands, lack of interoperability between each other)

Assumption 2/2: Many institutions are reluctant to give up their control over learning resources, which is currently troubling many central-server based approaches to learning resource sharing, e.g., eLearning „portals“. (trouble: „portals“ are costly but unprofitable)

Solutions:

- P2P: enable institutions to actively participate in a global sharing network without losing the control over their learning resources
- RDF: describe heterogeneous material and collections in distributed learning resource repositories to provide basic interoperability

AIFB

Edutella: Using a Peer-to-Peer Approach

- Peer-to-peer computing is the sharing of computer resources and services by direct exchange between systems
- Edutella connects highly heterogenous peers (heterogenous in uptime, performance, storage size, functionality, number of users...)
- Goal: making distributed nature of Edutella services (e.g. repository storage) completely transparent to Edutella clients
- Means to get there: specification and implementation of a set of Edutella services

AIFB

Edutella: Enabling Technology 1/2 (JXTA)

- Project JXTA (www.jxta.org)
  - An open source programming platform to enable P2P services and applications
  - Interoperability, Platform independence, Ubiquity
  - Layered approach

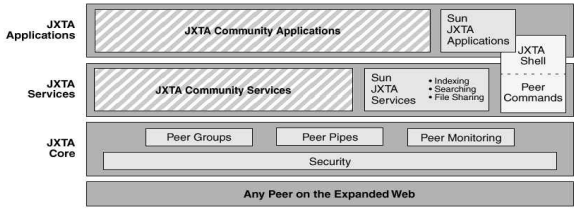

Source: Li Gong, Project JXTA: A Technology Overview

AIFB

Edutella: JXTA Layers

- JXTA-enabled Edutella Services

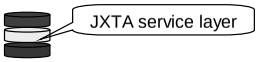

  - Edutella Query Service: query Edutella provider peers and retrieve query results
  - Edutella Annotation Service: annotate meta-data stored in repositories
  - Edutella Update and Replication Service: update and replicate (meta-) data
  - other services (mediation, mapping, etc.)

JXTA-enabled Edutella Applications

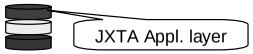

  - "Provider" Applications: Edutella Provider, Edutella Hub
  - "Client" Applications: Conzilla (Query GUI, KTH, Stockholm), Edutella Shell (Query GUI, L3S, Hannover), Ontomat (Annotation tool, AIFB, Karlsruhe), etc.

JXTA Core

  - Security, pipe-based communication, Rendezvous-based discovery mechanism

AIFB

Edutella: Enabling Technology 2/2 (RDF)

- Resource Description Framework
  - Key representation language in the Semantic Web
  - URI-based identifying mechanism to describe distributed resources, and state relationships between these resources
  - RDF graph model makes it easy to integrate a number of other formats for recording information, e.g. tables/tuples in RDB, simple assertions in formal logic, etc.,
  - RDFS offers a mechanism to define specific RDF vocabularies, enabling schema / ontology level meta-data mapping.

RDF strength in eLearning area and P2P environment

  - Availability of RDF-binding to most learning resource specifications (DC, IEEE LOM/IMS, ADL SCORM (ongoing)).
  - As „distributed XML", RDF perfectly fits into decentralized P2P environments, enabling distributed queries spanning various repositories.

AIFB

Edutella: Goal and Approach

- Specify and implement a RDF-based meta-data infrastructure for P2P networks
- Developed as part of the open source peer-to-peer project JXTA
- edutella.jxta.org
- Contributors from various institutions

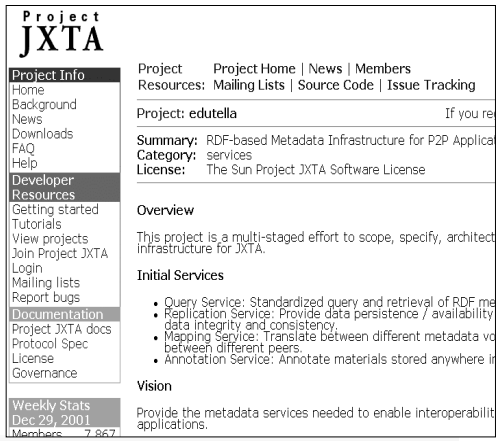

AIFB

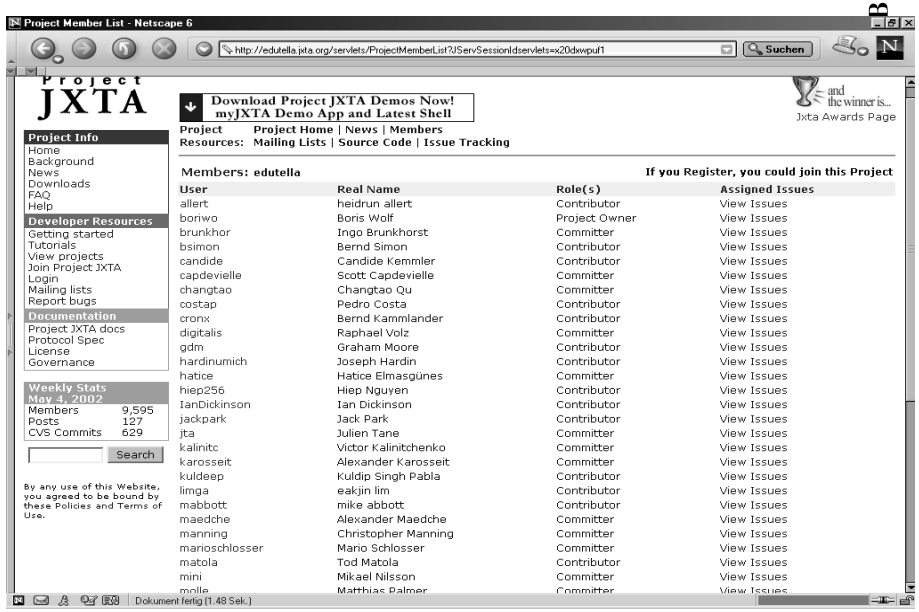

AIFB

RDF/RDFS: What is missing?

- RDF/RDFS: describe distributed resources on the Web and to describe the vocabulary (properties) and constraints (classes, domain/range) for these descriptions
- How do I query these distributed sources ???
  - How do my queries reach their destination?
  - How do I ensure they are understood at their destination?

AIFB

RDF-QEL: RDF Query Exchange Language

- RDF-based Query Exchange Language (RDF-QEL)
  - RDF QEL1: conjunctive query
  - RDF QEL2: RDF QEL1 + disjunctive query
  - RDF QEL3: RDF QEL2 + negation (SQL92)
  - RDF QEL4: RDF QEL3 + transitive closure (SQL3)
  - RDF QEL5: RDF QEL4 + general recursion
  - Datalog is used as the internal data model (ECDM: Edutella Common Data Model) and provided as a set of Java classes
  - RDF is used to represent the queries transmitted between the peers
  - Wrappers for other RDF query languages (RQL, TRIPLE, etc.) and XML query languages (like Xpath)

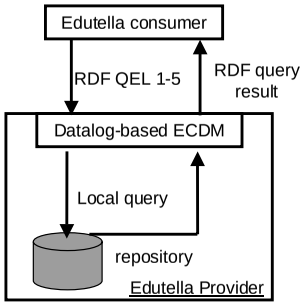

Edutella query data flow

AIFB

Edutella: Query Service Architecture

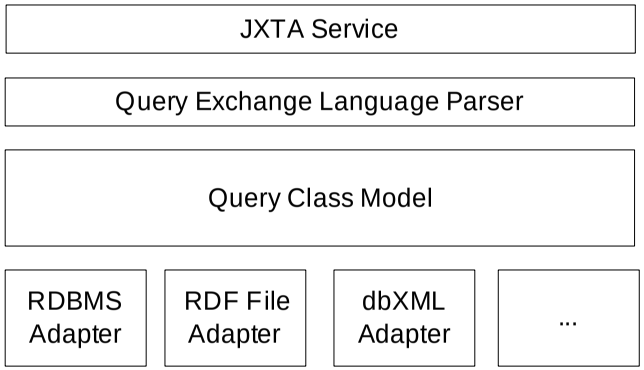

AIFB

Edutella: Query Class Model (ECDM)

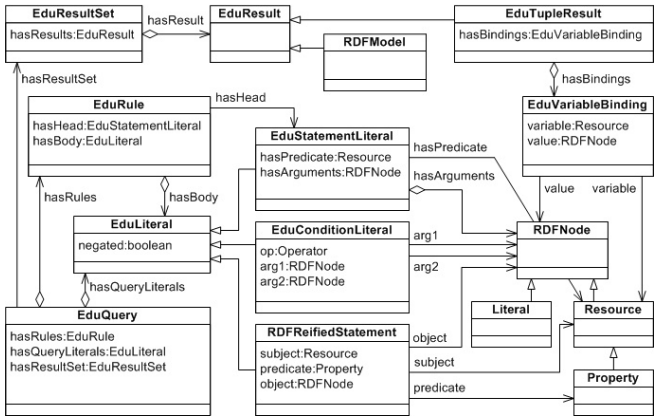

AIFB

Edutella: Working Scenario

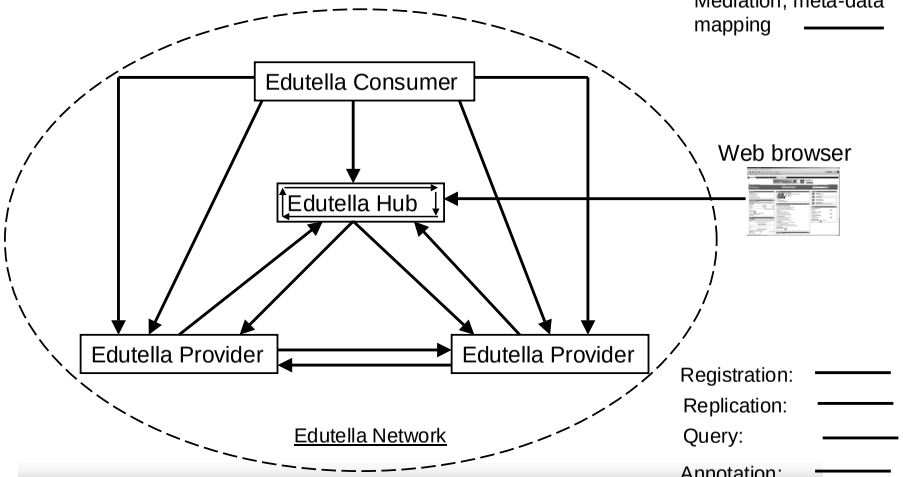

AIFB

Edutella: Current Status

- 4 Edutella provider peers have been implemented, permanent peers have been running since March 7

OLR (Open Learning Repository)

Back-end: Oracle 8i
Meta-data: RDF
Local query : SQL
Implementer: L3S, Hannover

ConceptBase

Back-end: ConceptBase
Meta-data: RDF
Local query : O-Telos
Implementer: L3S, Hannover

RDQL

Back-end: file system
Meta-data: RDF
Local query : Jena RDQL
Implementer: L3S, Hannover

dbXML

Back-end: Apache Xindice
Meta-data: RDF
Local query : XPath
Implementer: L3S, Hannover

AIFB

Edutella: Current Status

3 Edutella provider peers are under development

AMOS II

Back-end: Multi data sources
Meta-data: RDF
Local query: AMOSQL
Implementer: Uni. of Uppsalla, Uppsalla

KAON

Back-end: RDB
Meta-data: RDF
Local query: SQL
Implementer: AIFB, Karlsruhe

Tamino

Back-end: Tamino
Meta-data: RDF
Local query: XQuery
Implementer: L3S, Hannover

AIFB

Edutella: Current Status

4 Edutella Consumers have been implemented

Conzilla (Query GUI)
Implementer: KTH, Stockholm

Ontology Query (Query GUI)
Implementer: Stanford

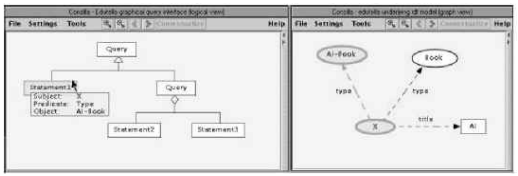

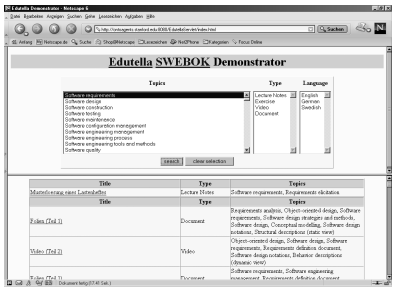

AIFB

Edutella: Current Status

4 Edutella Consumers have been implemented

Edutella shell (Query GUI)
Implementer: L3S, Hannover

Ontomat (Annotation tool)
Implementer: AIFB, Karlsruhe

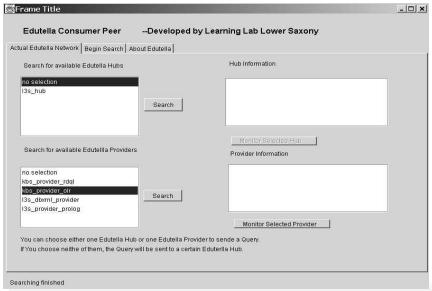

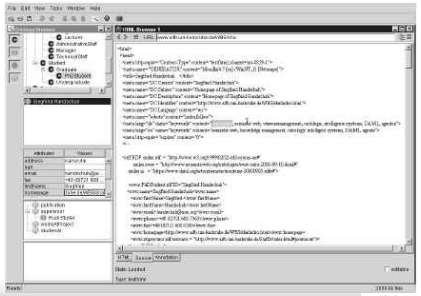

AIFB

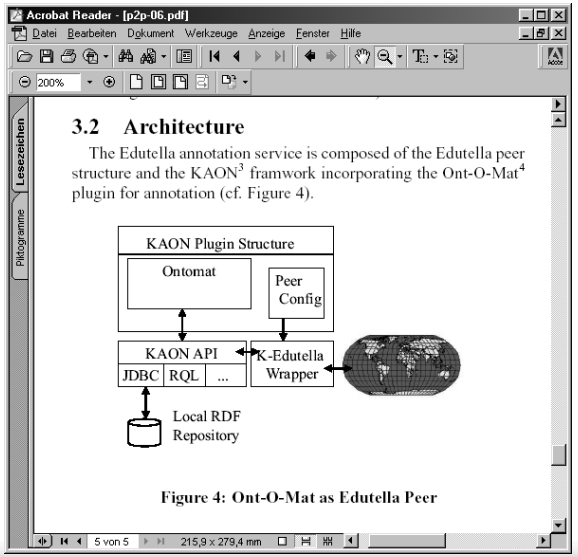

AIFB

Links

- All source code downloadable from the Edutella Homepage under edutella.jxta.org
- PADLR Project Homepage: www.learninglab.de/english/projects/padlr.html

New Edutella contributors are welcome!

Objective: Ad-hoc Information System

AIFB

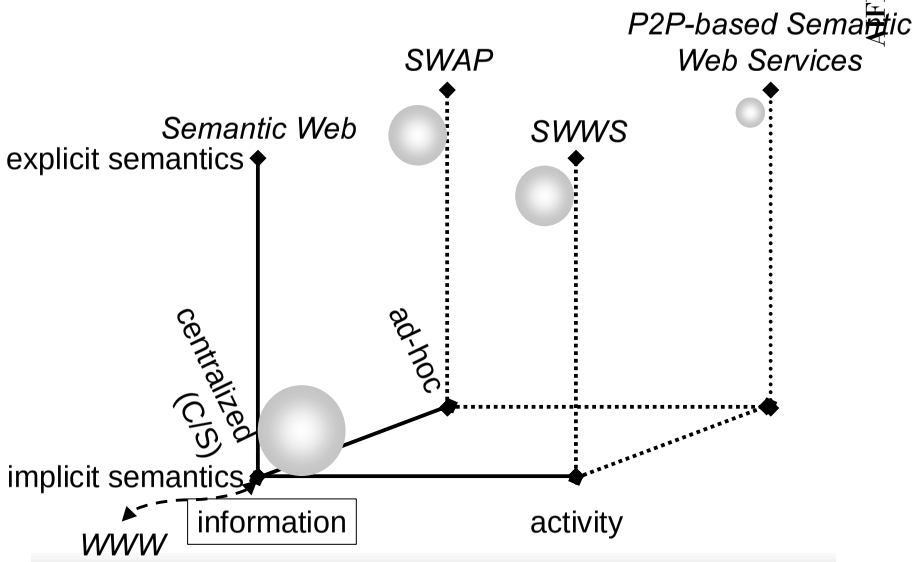

AIFB

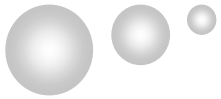

Complexity/size of manageable service(s)

AIFB

AIFB

AIFB