Data-Centric Networks and Peer-to-Peer Data Management

Wolfgang Nejdl

L3S / University of Hannover
Germany

Overview

Motivation and Background

- Networks, Databases and the Web: DHTs, PEERS, et al
- L3S P2P Background - KnowledgeWeb, Edutella, et al

Schema-Based Peer-to-Peer Networks

- Resource Description Framework (RDF) and RDF Schema
- Edutella Query Service / RDF Query Exchange Language RDF-QEL
- Subscriptions
- Efficient Routing / HyperCuP & Super-Peers
- Distributed Query Processing
- Access Control and Trust Negotiation

Summary and Conclusions

Evolution of Networks

from

- Host-centric networks (URLs & low level routing)
- Enabling and optimizing communication between network hosts

to

- Data-centric networks (Google & P2P)
- Into the Coddian world of physical data independence

(invited talk Scott Shenker, VLDB 2003)

Where Databases meet Networks: PEERS

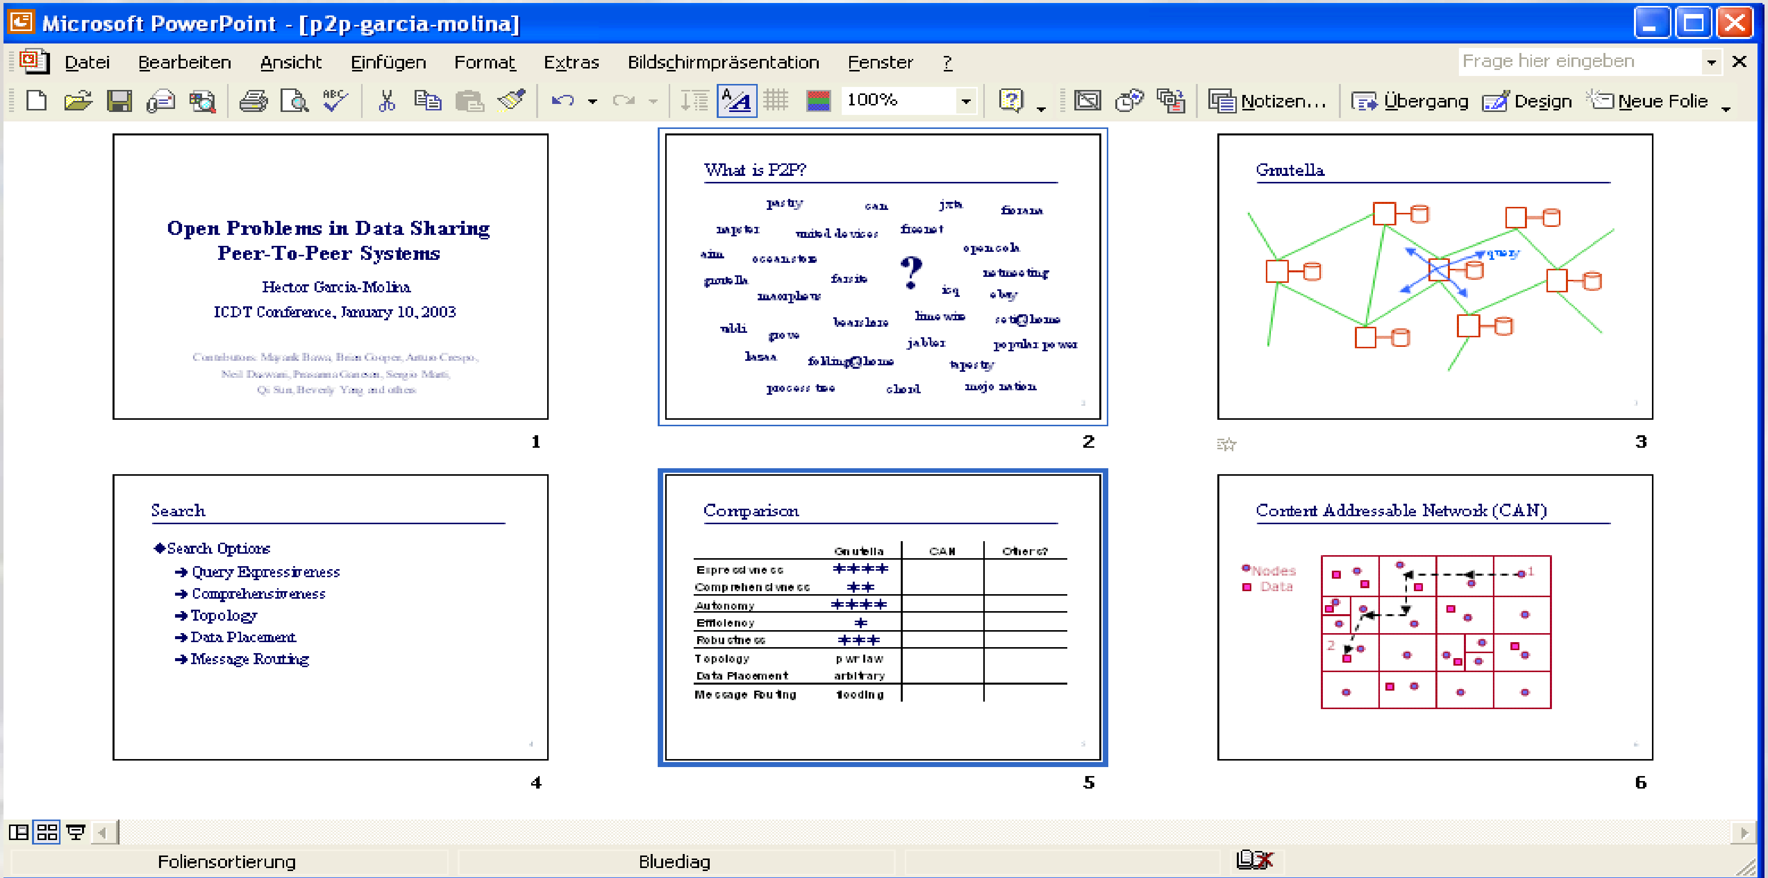

Where the Semantic Web Meets Databases: KnowledgeWeb

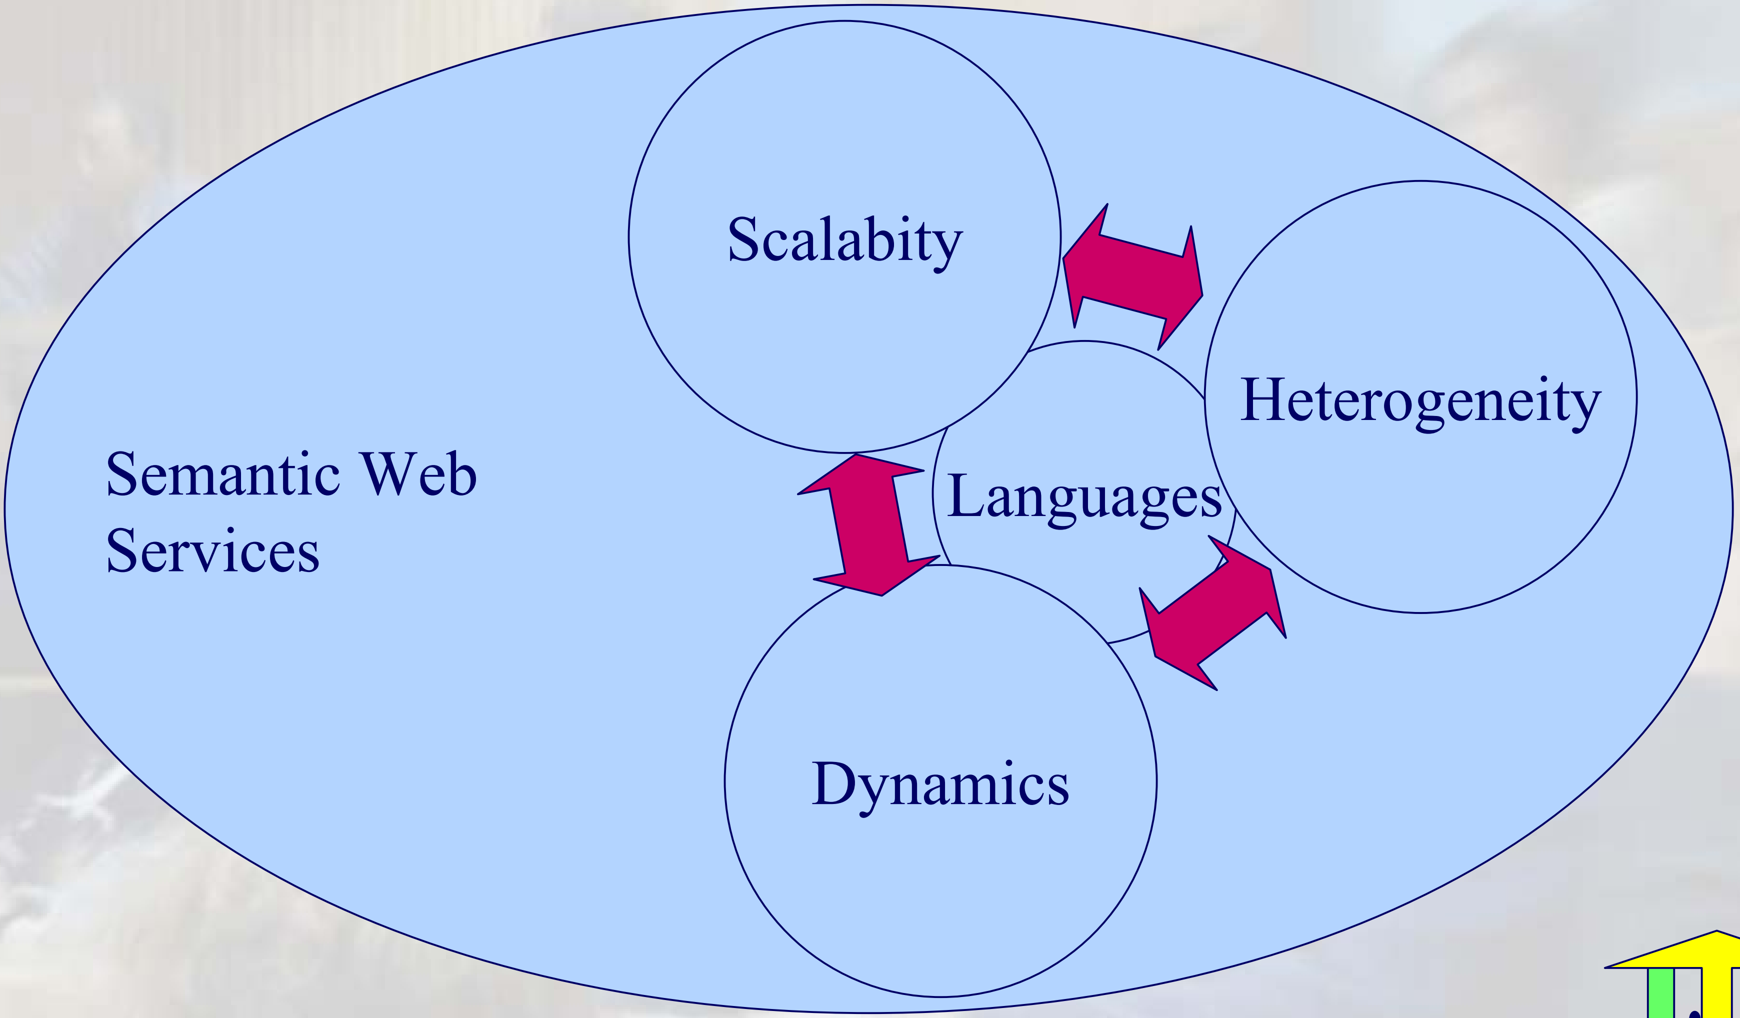

Where Databases & Rules meet the Semantic Web: REWERSE

- Reasoning on the Web with Rules and Semantics
  - How to get and retrieve data?
    - **Querying, reasoning and optimization**
  - How to protect data?
    - **Policy specification and evaluation**
  - How to integrate data?
    - **Reasoning and mediation**
- Selected applications for **proof-of-concept** purposes
  - Personalized Web systems
  - Web-based decision support
  - Bioinformatics Semantic Web

Where E-Learning meets Databases & Sem. Web: Edutella

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

60+ contributors from various institutions

Building block for the EU/IST ELENA smart learning space

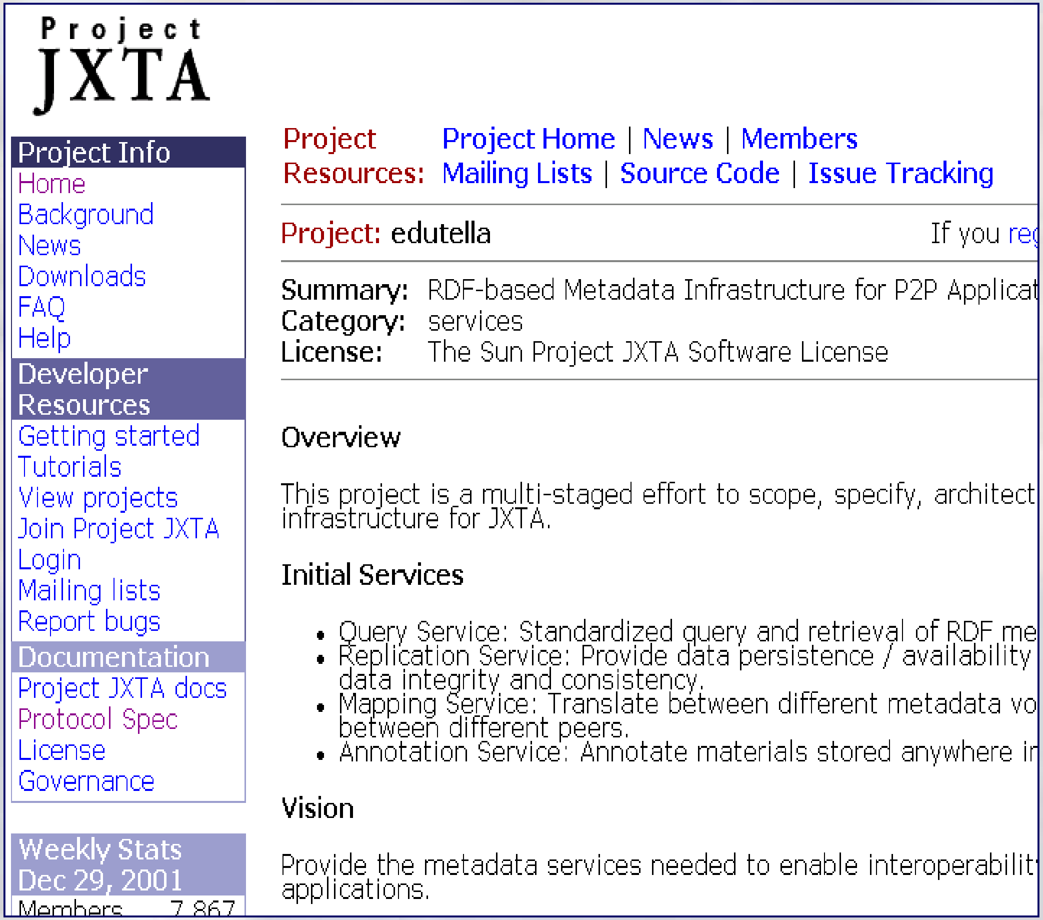

E-Learning + Infrastructures + Interoperability: PROLEARN

Working towards

- innovative elearning resources
- interoperable elearning resources and systems
- sustainable elearning infrastructures and processes for SMEs

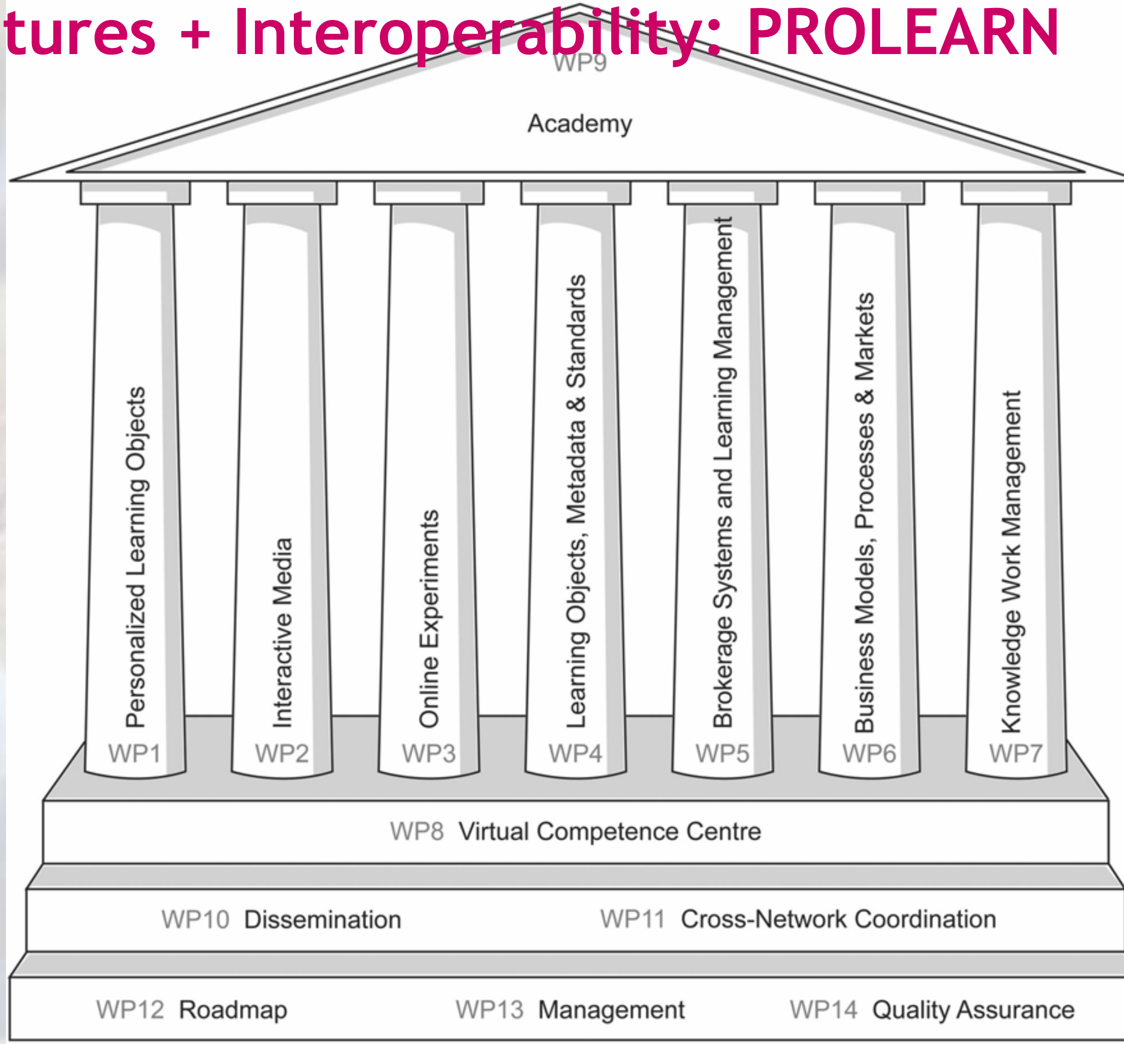

Schema-Based Peer-to-Peer Networks

User-definable schemas
Structured schemas
Query language

Decentralized control
Node autonomy
Transient peers
Self organization

Database Systems | Schema-based P2P Systems | P2P Systems

	local	distributed	peer-to-peer
schema-based	Any RDBMS ConceptBase OntoBroker	AmosII ObjectGlobe TSIMMIS Tukwila	Chatty Web Edutella Piazza
fixed schema/ keywords		Napster	DirectConnect Gnutella KaZaA P-GRID
key			CAN CHORD

(system list not complete)

RDF / RDF Schema for Describing Distributed Resources

Basic Formalisms for the Semantic Web

- URIs to identify resources
- Combine resources and annotate resources with attributes, using <Subject, Property, Value> Tuples
- Graph as basic model, easy to translate to logic facts
- RDFS allows us to define the RDF vocabulary used (classes and attributes), and thus to represent simple semantic models
- Possible extensions towards more expressive semantic descriptions, e.g. description logic (DAML+OIL / OWL)

Using RDF / RDFS in the P2P context

- Distributed annotations for distributed resources
- Flexible schema definitions, which can be uniquely identified and combined, as well as extended by additional properties

RDF-QEL: RDF Query (Exchange) Language

Datalog-based Query Exchange Language (RDF-QEL)

- RDF QEL1: conjunctive query up to
- RDF QEL5: RDF QEL4 (SQL3) + general recursion

see Nejdl et al: „EDUTELLA: A P2P Networking Infrastructure Based on RDF“, WWW 2002

- 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)

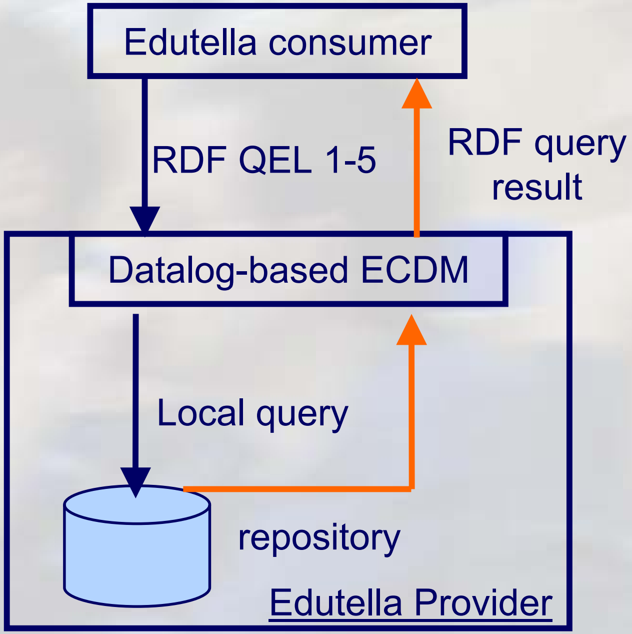

Edutella query data flow

Another Possibility: Don't query, subscribe

Subscriptions are a good idea, too (get the NYTimes each morning, get new teaching material on P2P topologies …)

Example: Selective Information Dissemination in P2P-DIET

Instead of Queries and Answers we need

- Profile forwarding
- Notification forwarding / Filtering
- Advertisement forwarding
- Dynamicity of P2P network → storing notifications / rendezvous

See e.g. Koubarakis et al: Selective Information Dissemination in P2P Networks: Problems and Solutions, SIGMOD Record, Special P2P Issue, September 2003 as well as ongoing work to integrate P2P-DIET and Edutella

See also Terpstra, Buchmann et al: A P2P-Approach to Content-Based Publish/Subscribe

P2P and Efficient Routing

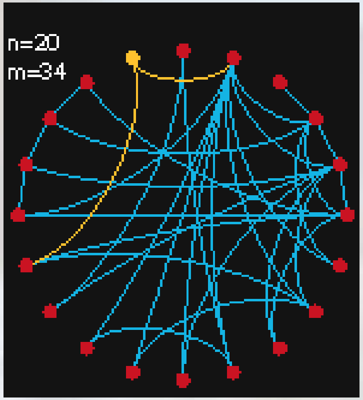

How do peer-to-peer networks scale?

Requirements:

- Symmetric topology (every node is a root)
- Low network diameter (small worlds property, should be O(log n))
- Limited node degrees (number of peer-connections from a node, should be O(log n))
- Load balancing of traffic
- Efficient broadcast (receive broadcast messages only once)
- Adaptable to dynamic number of peers

HyperCuP Peer-to-Peer Topology

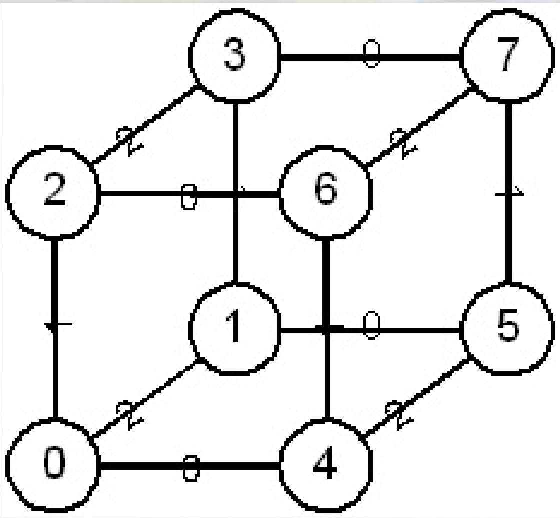

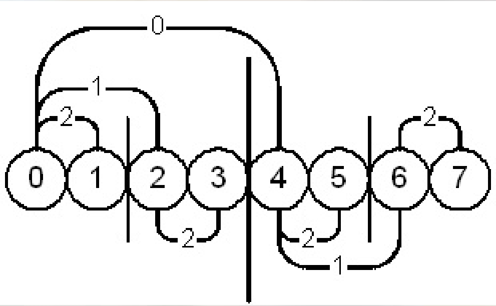

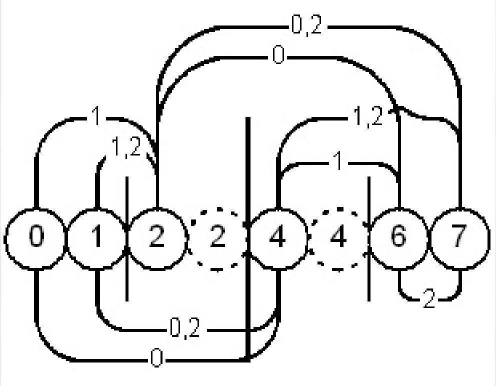

Details: see e.g. Schlosser, Sintek, Decker, Nejdl: „HyperCuP – Shaping Up Peer-to-Peer Networks", 2nd Intl. WS on Agents and P2P Computing, 2002

Hypercube Topology

Broadcast Algorithm

- Annotate messages with the “dimension” of the peer-to-peer connection, and only forward it along “higher” dimensions

Properties

- Network diameter, characteristic path length and number of nodes are $O(\log_b N)$
- Fault tolerant, vertex-symmetric

Extensions

- Dynamic hypercube
- Base=N hypercube
- Cayley graphs

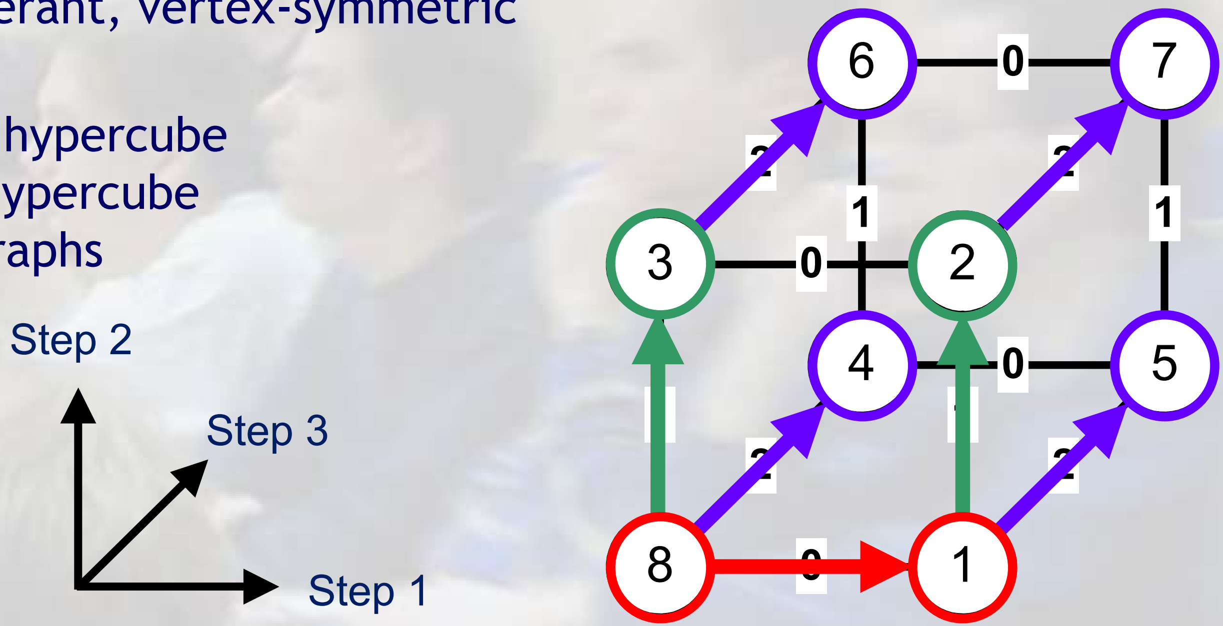

Chord Neighbor and Route selection Algorithms

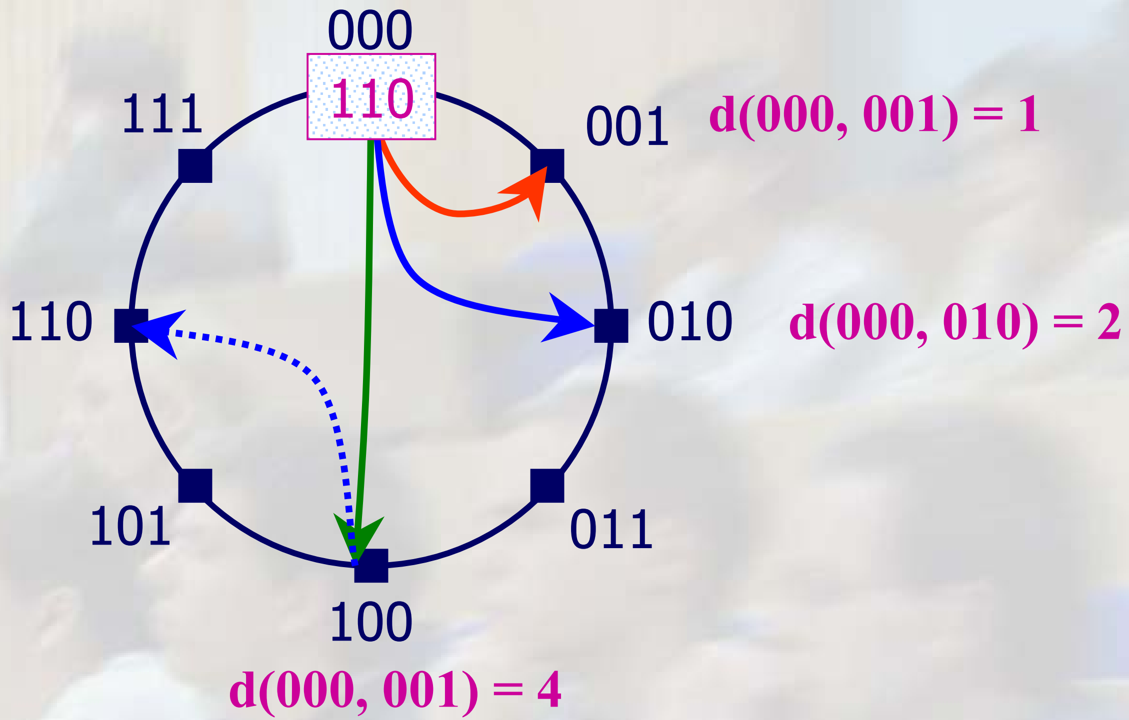

Neighbor selection: **i**th neighbor at **2**^{i} distance

Route selection: pick neighbor closest to destination

Super-Peer Networks

Observation: Peers vary significantly in availability, bandwidth, processing power, etc.

Create network backbone from highly available and powerful peers to distribute load better.

See also Yang, Garcia-Molina: Improving Search in P2P Systems, Intl. Conf. on Distributed Computing Systems, Vienna, 2002, or file sharing networks like KaZaa

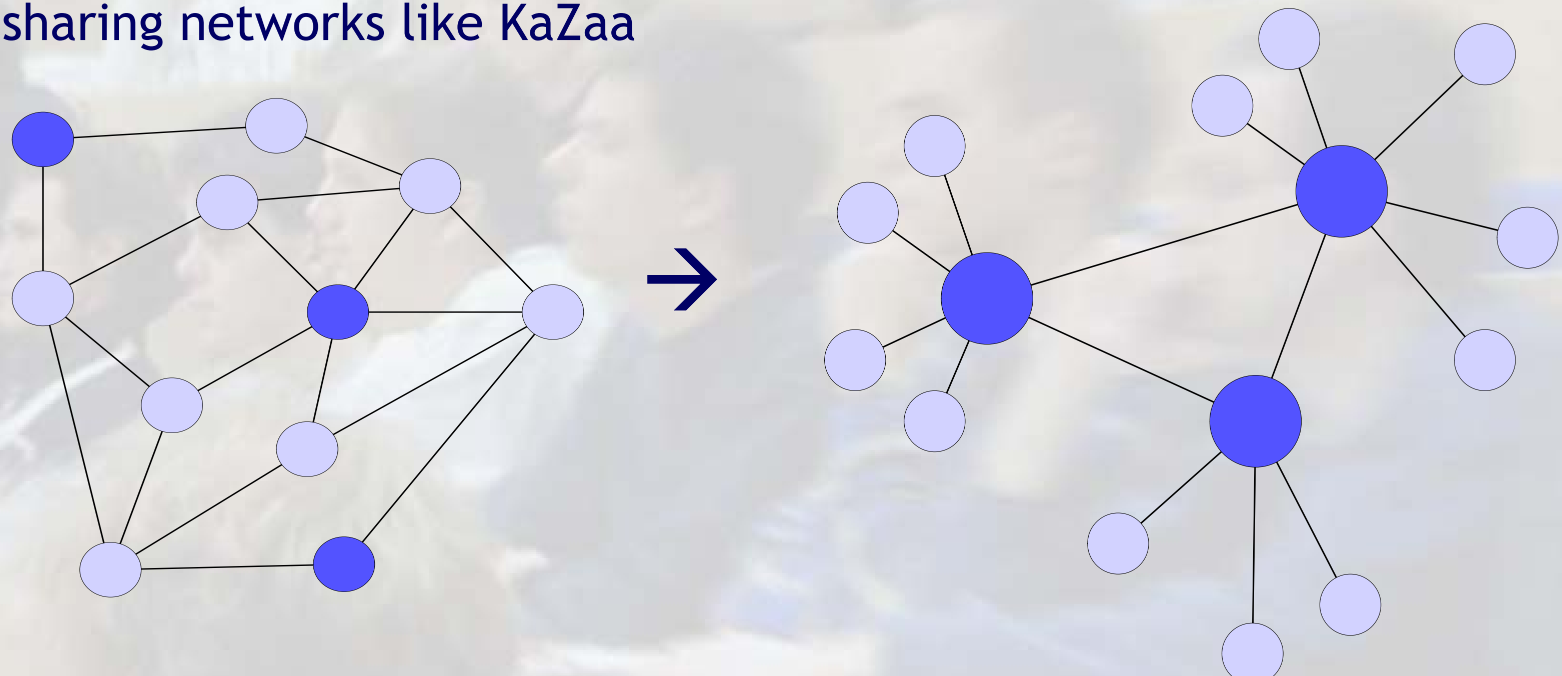

Super-Peers and Routing Indices

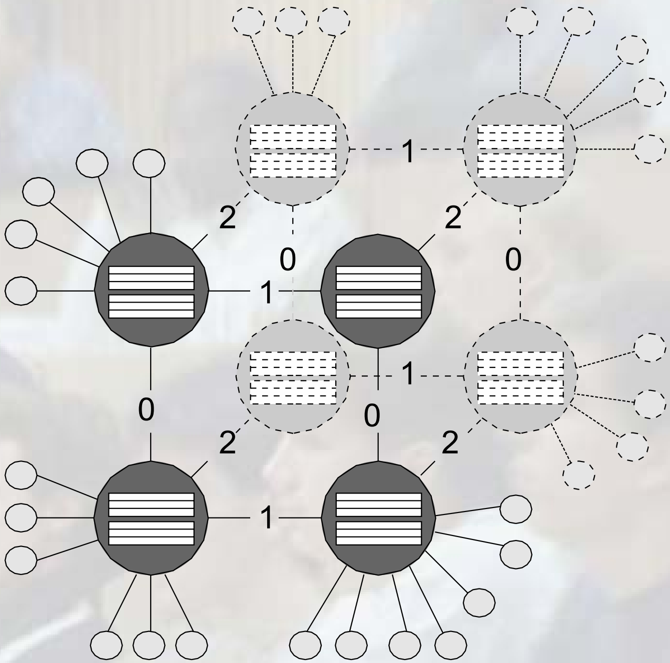

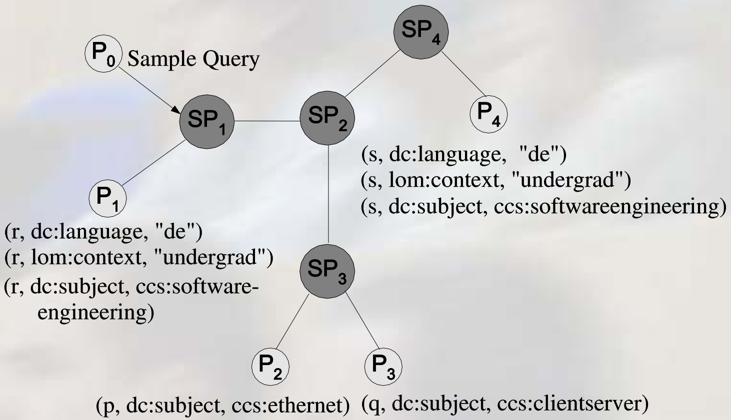

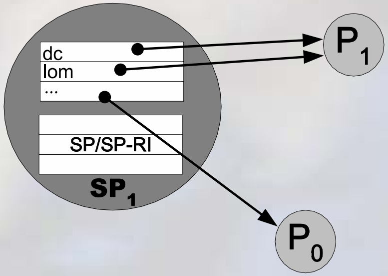

Nejdl et al. Super-Peer-Based Routing and Clustering Strategies for RDF-Based Peer-To-Peer Networks. WWW 2003

Extension to Distributed Query Processing

Interleave P2P techniques and query processing

- Push abstract query plans through the super peer network
- Super peers pick and expand those parts of the query plan that can be executed locally
- On the fly distribution and expansion of query plans
- See Brunkhorst, Dhraief, Kemper, Nejdl, Wiesner: Distributed Queries and Query Optimization in Schema-Based P2P-Systems, VLDB-P2P-Workshop

Query Optimization exploits clustering strategies

- Access-path clustering: attribute-based clustering using per-attribute hypercubes (using the hypercube as a balanced n-ary search tree) (see Dhraief, Kemper, Nejdl, Wiesner: Distributed Queries and Query Optimization in Schema-Based P2P Systems, submitted)

Access Control and Automated Trust Negotiation

- Goal → protect resources from unauthorized access
- Establish trust between strangers
  - Initial trust among nodes is not necessary
  - No need for prior registration
- Use and interchange of credentials: online analogue to the paper credentials in real life.
- Negotiation according to policies
  - Access control policies can be used in both sides (requester and provider)
- Delegation
- Automated Trust Negotiation → iterative exchange of digital credentials.
  - Iterative disclosure of policies and credentials

Credentials and Policies

- Property-based credentials
  - Describe one or more properties / attributes of the owner asserted by the issuer, signed with the private key of the issuer
  - As credentials contain sensitive information, they are not shown until the other part demonstrates that it is qualified to have such sensitive information.
- Access Control Policies
  - Protect a resource or a credential
  - Specify credentials that the other negotiation participant must provide in order to get access
  - Several policies can be involved during the negotiation.
  - Several policies for the same resource or credential.
  - Policies can be protected like any other resource.

Example „Alice & E-Learn"

Step 1: Alice requests to access E-Learn's free Spanish course

Step 2: E-Learn replies with policy protecting this resource

- Requests police badge to prove police officer status
- Requests driver's licence to prove California residence status

Step 3: Alice views her driver's license as non-critical, but needs to protect her police officer credential

- Discloses driver's license
- Requests E-Learn membership proof from the Better Business Bureau

Step 4: E-Learn agrees

- Discloses Better Business Bureau membership card

Step 5: Alice finds her policy satisfied

- Discloses police badge

Step 6: E-Learn finds its policy satisfied

- Makes Spanish course available

Automated Trust Negotiation among Peers on the Web

Design policy language to express trust negotiation

- Delegation, policy protection, negotiation strategies
- Based on guarded distributed logic programs

Develop run-time system for automated trust negotiation

- Based on Prolog meta interpreter embedded as Java library in Applet / Server (WWW) or Peer-to-Peer (Edutella) environment

Currently two application areas

- eLearning (ELENA, EU/FP5 @ L3S)
- Emergency management (ITR @ DAIS/UIUC (M. Winslett))

See e.g. Yu, Winslett, and Seamons. Supporting Structured Credentials and Sensitive Policies through Interoperable Strategies for Automated Trust Negotiation, *ACM Transactions on Information and System Security*, February 2003.

See also: Nejdl, Olmedilla, Winslett: Automated Trust Negotiation among Peers on the Semantic Web, submitted

Summary and Conclusions

Schema-based P2P networks and P2P-based data management infrastructures build upon traditional P2P networks and distributed / heterogeneous database research, while posing new challenges as well as additional functionalities

Building blocks are flexible / extendable schema languages, expressive query and reasoning languages, efficient network topologies as well as routing and clustering algorithms, data integration and mediation functionalities, query optimization, and last but not least, dezentralized access control and trust negotiation mechanisms

See also SIGMOD Record September 2003, Special P2P Issue: Nejdl, Siberski, Sintek: „Design Issues and Challenges for RDF- and Schema-Based Peer-to-Peer Systems“