VLDB 2005

Semantic Overlay Networks

Karl Aberer and Philippe Cudré-Mauroux

School of Computer and Communication Sciences
EPFL -- Switzerland

Overview of the Tutorial

- I. P2P Systems Overview
- II. Query Evaluation in SONs
  - RDFPeers
  - PIER
  - Edutella
- III. Semantic Mediation in SONs (PDMSs)
  - PeerDB
  - Hyperion
  - Piazza
  - GridVine
- IV. Current Research Directions

What this tutorial is about

- Describing a (pertinent) selection of systems managing data in large scale, decentralized overlays networks
  - Focus on architectures and approaches to evaluate / reformulate queries

- It is *not* about
  - A comprehensive list of research projects in the area
    - But we'll give pointers for that
  - Complete description of each project
    - We focus on a few aspects
  - Performance evaluation of each approach
    - No meaningful comparison metrics at this stage

I. Peer-to-Peer Systems Overview

- Application Perspective: Resource Sharing (e.g. images)
  - no centralized infrastructure
  - global scale information systems

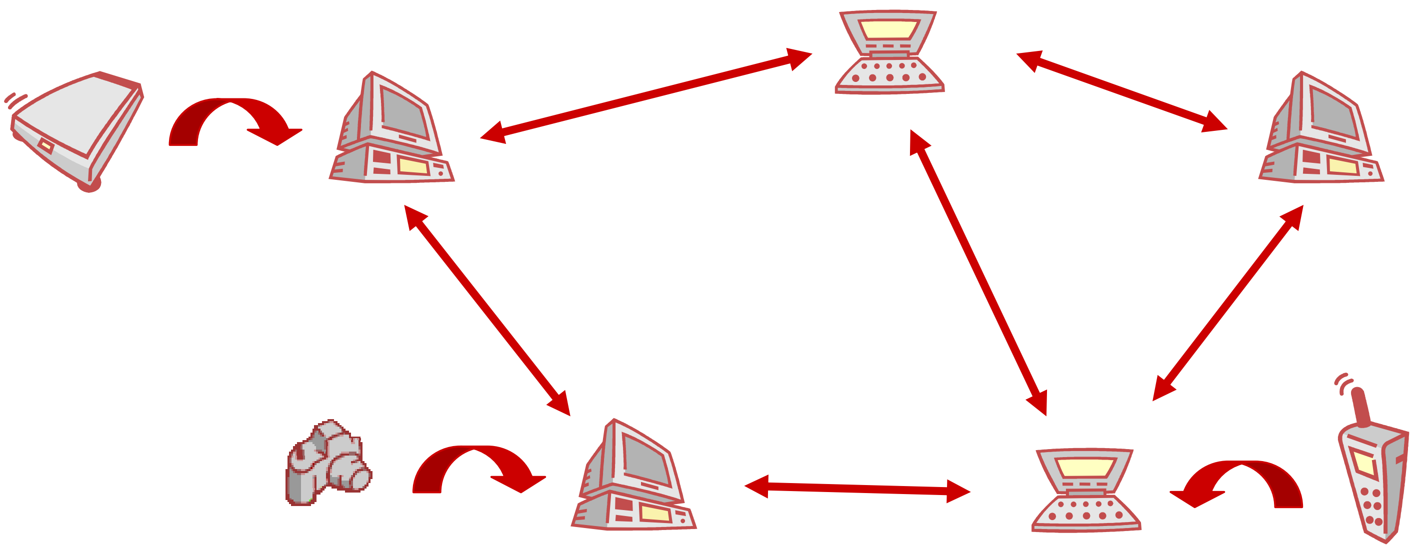

Resource Sharing

- What is shared?

knowledge

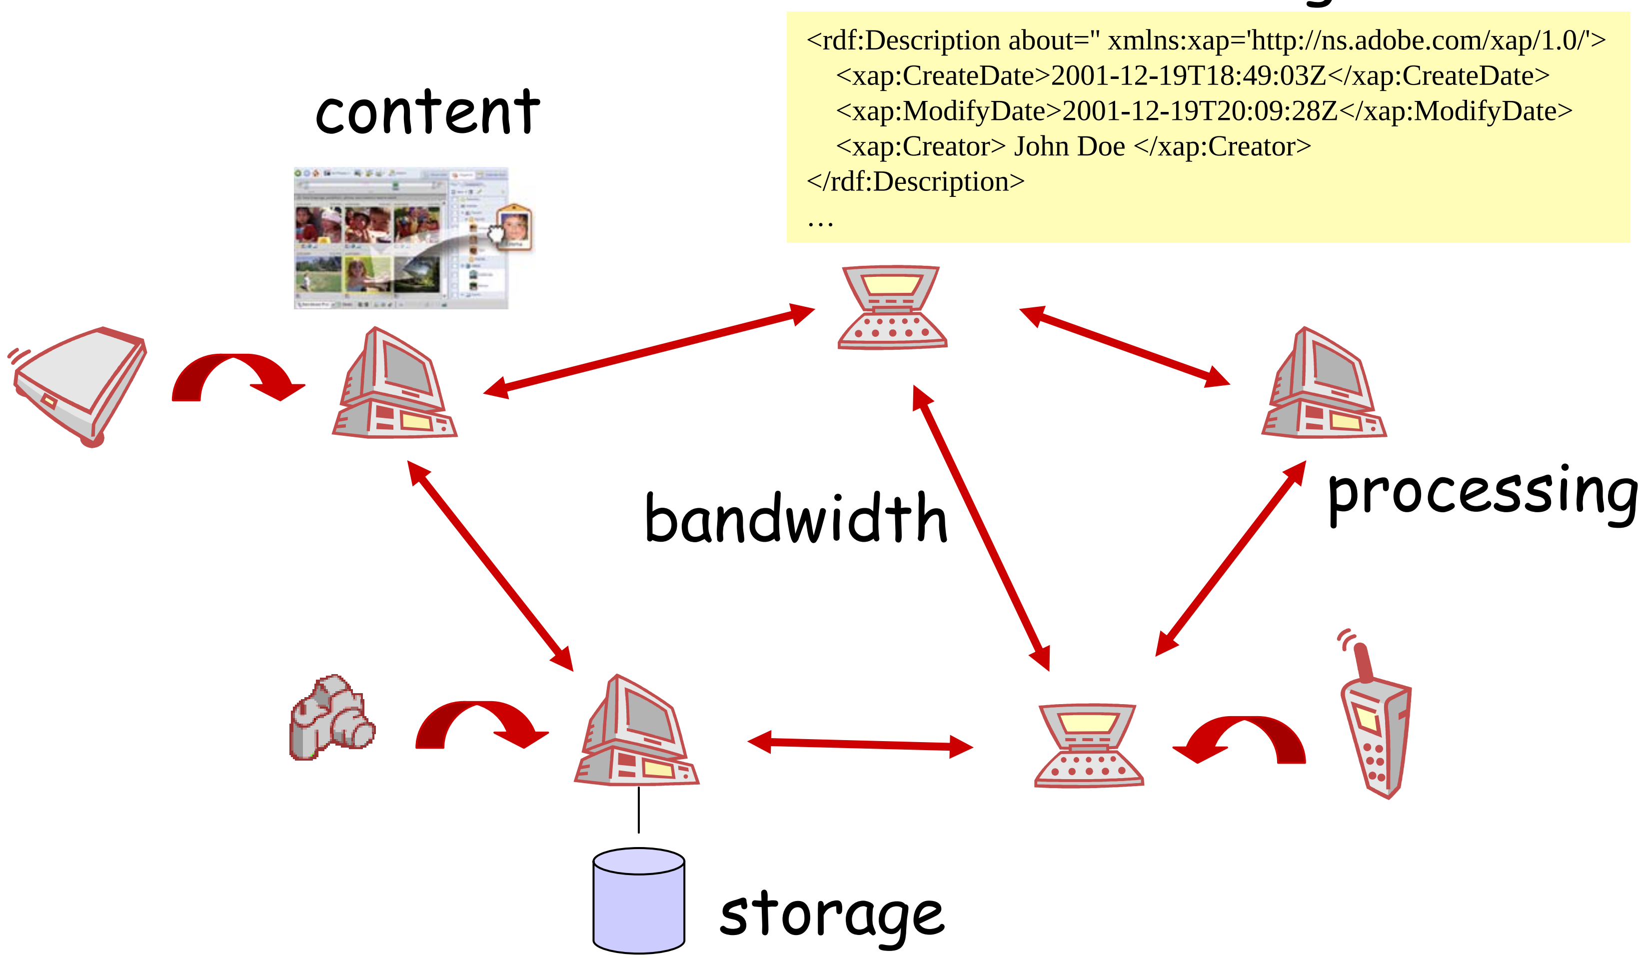

Enabling Resource Sharing

- Searching for Resources
  - Overlay Networks, Routing, Mapping
- Resource Storage
  - Archival storage, replication and coding
- Access to Resources
  - Streaming, Dissemination
- Publishing of Resources
  - Notification, Subscription
- Load Balancing
  - Bandwidth, Storage, Computation
- Trusting into Resources
  - Security and Reputation
- etc.

P2P Systems

- System Perspective: Self-Organized Systems
  - no centralized control
  - dynamic behavior

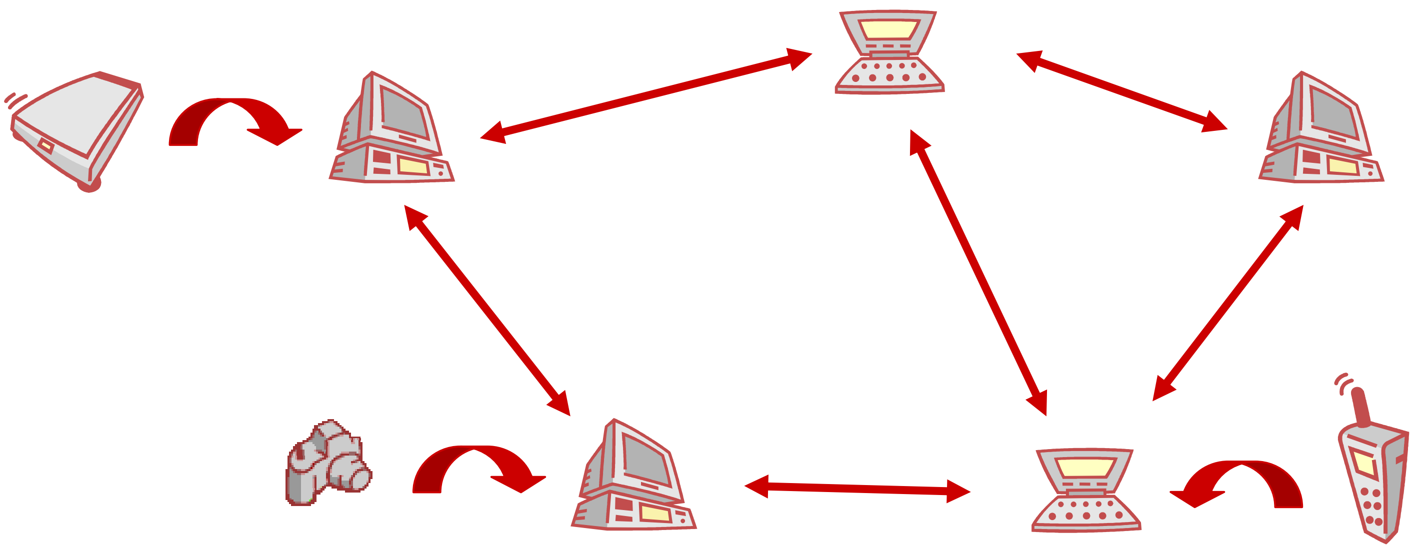

What is Self-Organization?

- Informal characterization (physics, biology,… and CS)
  - distribution of control (= decentralization)
  - local interactions, information and decisions (= autonomy)
  - emergence of global structures
  - failure resilience and scalability

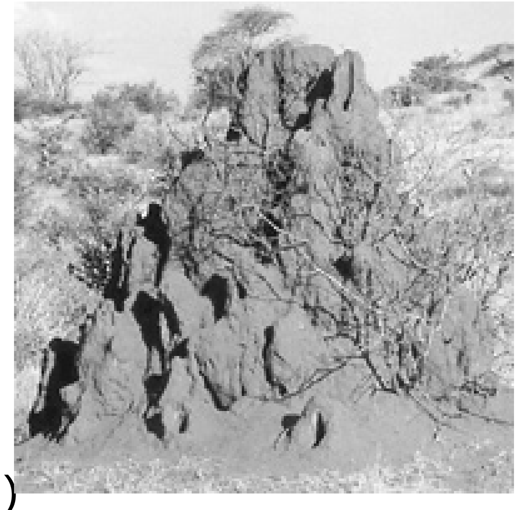

- Formal characterization
  - system evolution f_T: S $\rightarrow$ S, state space S
  - stochastic process (lack of knowledge, randomization)

$$P(s_j, t+1) = \sum_i M_{ij} P(s_i, t), \; P(s_i | s_j) = M_{ij} \in [0,1]$$

  - emergent structures correspond to equilibrium states
  - no entity knows all of S

Examples of Self-Organizing Processes

- Evolution of Network Structure
  - Powerlaw graphs: Preferential attachment + growing network [Barabasi, 1999]
  - Small-World Graphs: FreeNet Evolution
- Stability of Network
  - Analysis of maintenance strategies
  - Markovian Models, Master Equations

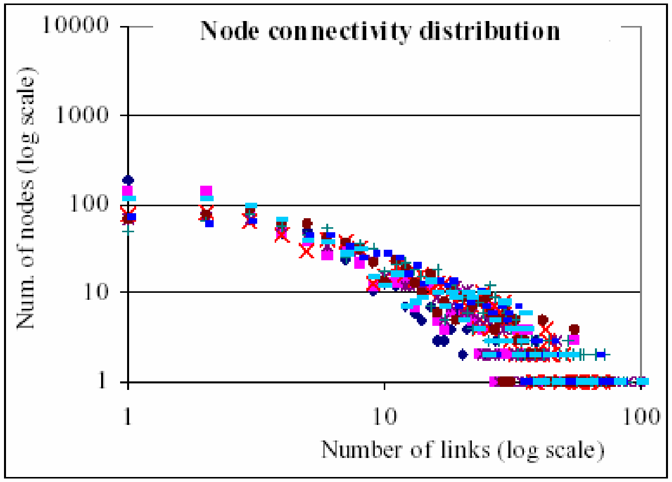

- Resource Allocation
  - game-theoretic and economic modelling
- Probabilistic Reasoning
  - Belief propagation for semantic integration (see later)

Efficiently Searching Resources (Data)

- Find images taken last week in Trondheim!

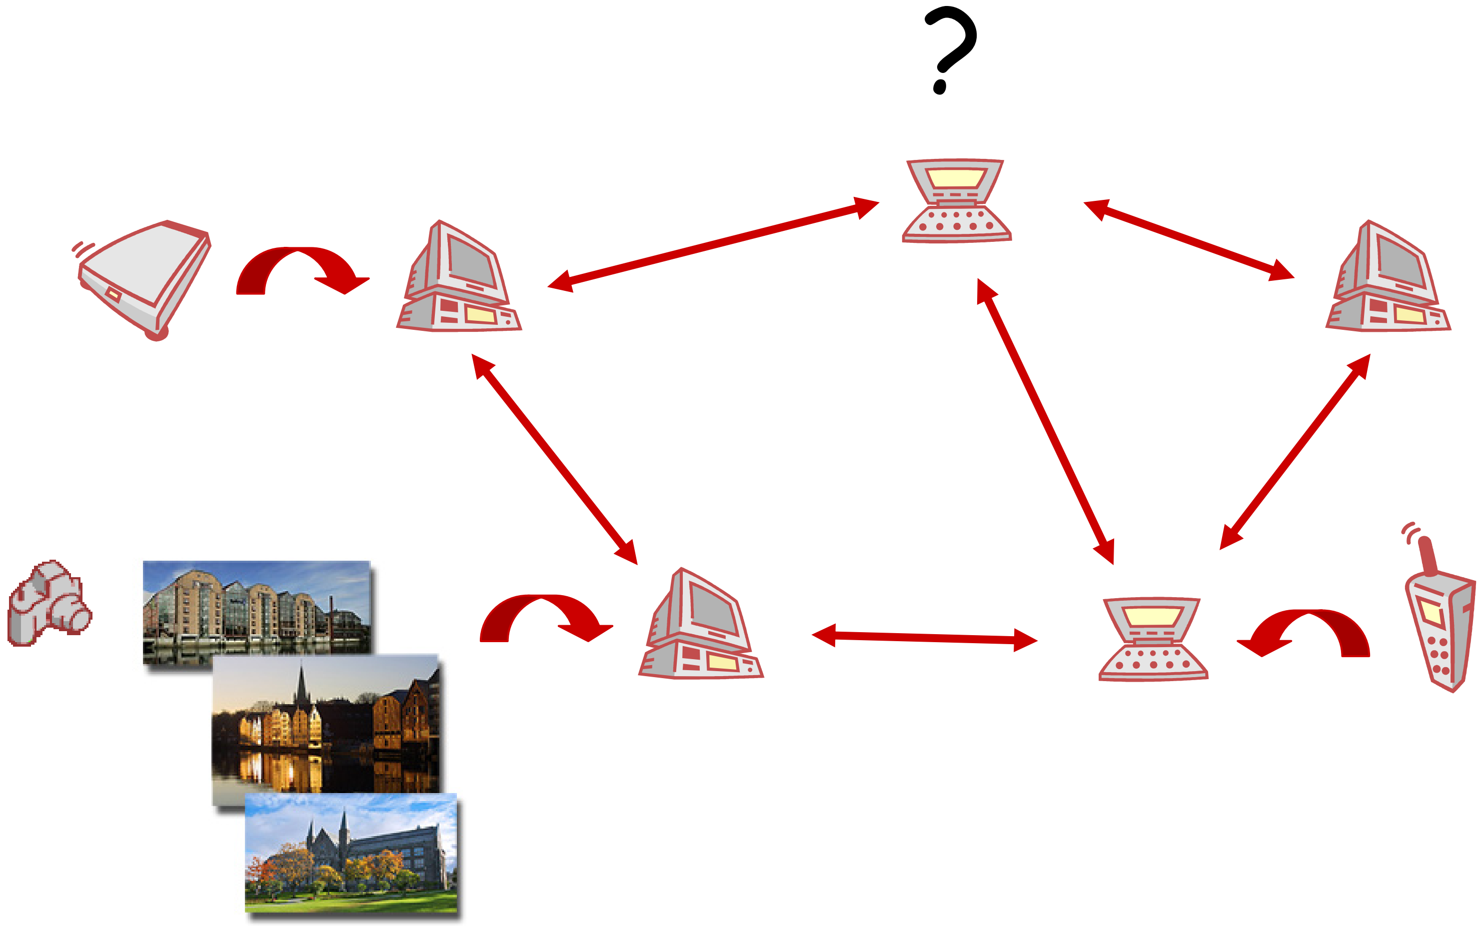

Overlay Networks

- Form a logical network in top of the physical network (e.g. TCP/IP)
  - originally designed for resource location (search)
  - today other applications (e.g. dissemination)
- Each peer connects to a few other peers
  - locality, scalability
- Different organizational principles and routing strategies
  - unstructured overlay networks
  - structured overlay networks
  - hierarchical overlay networks

Unstructured Overlay Networks

- Popular example: Gnutella
- Peers connect to few random neighbors
- Searches are flooded in the network

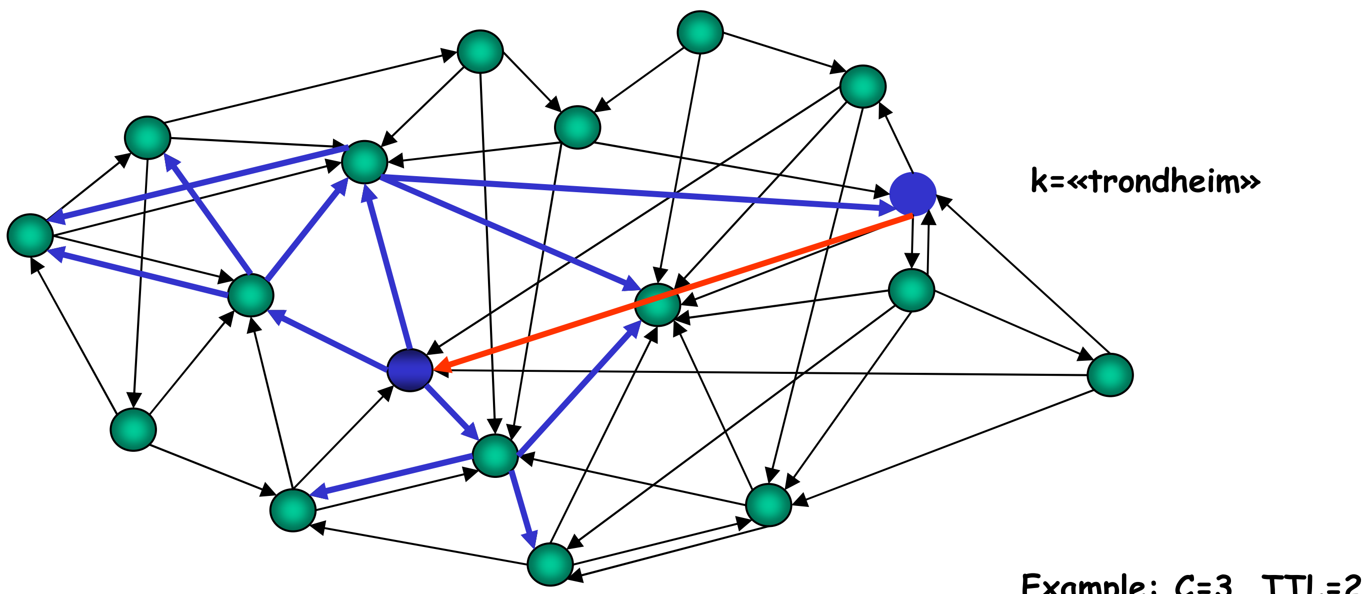

Structured Overlay Networks

- Popular examples: Chord, Pastry, P-Grid, ...
- Based on embedding a graph into an identifier space (nodes = peers)
- Peers connect to few neighbors carefully selected according to their distance
- Searches are performed by greedy routing
- Variations of Kleinberg's small world graphs:

$$P[u \rightarrow v] \sim d(u, v)^{-r}$$

r = 2

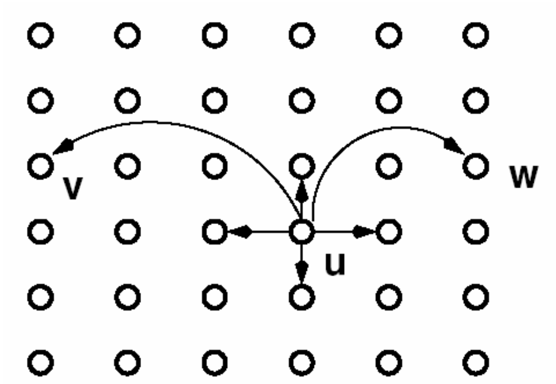

Conceptual Model for Structured Overlay Networks

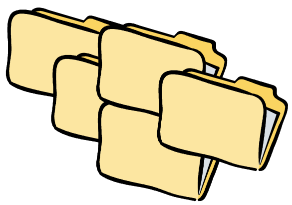

Set of resources R

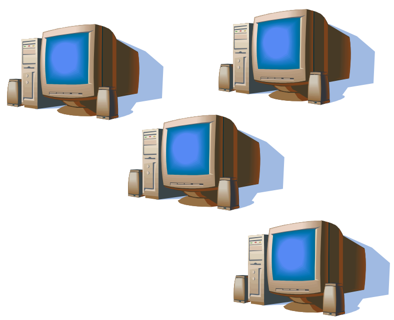

Group of peers P

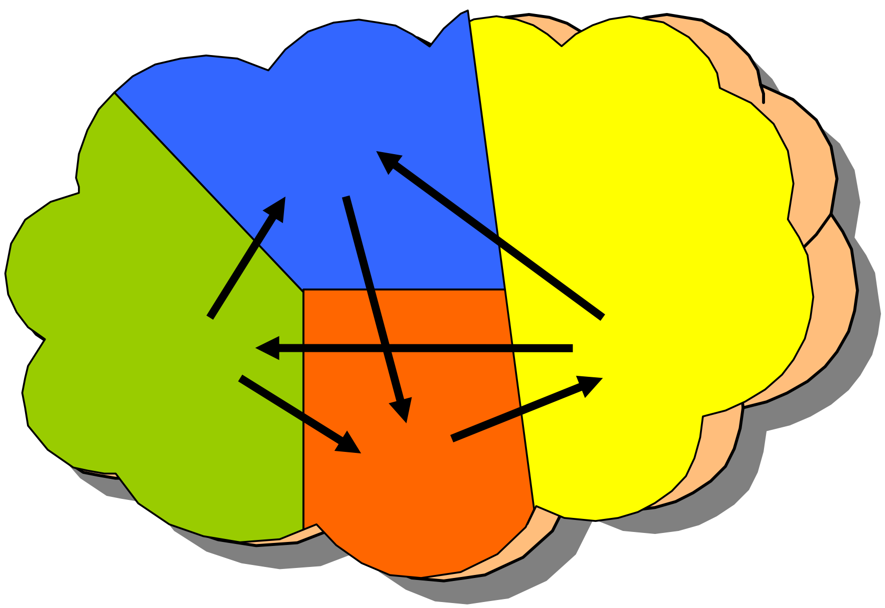

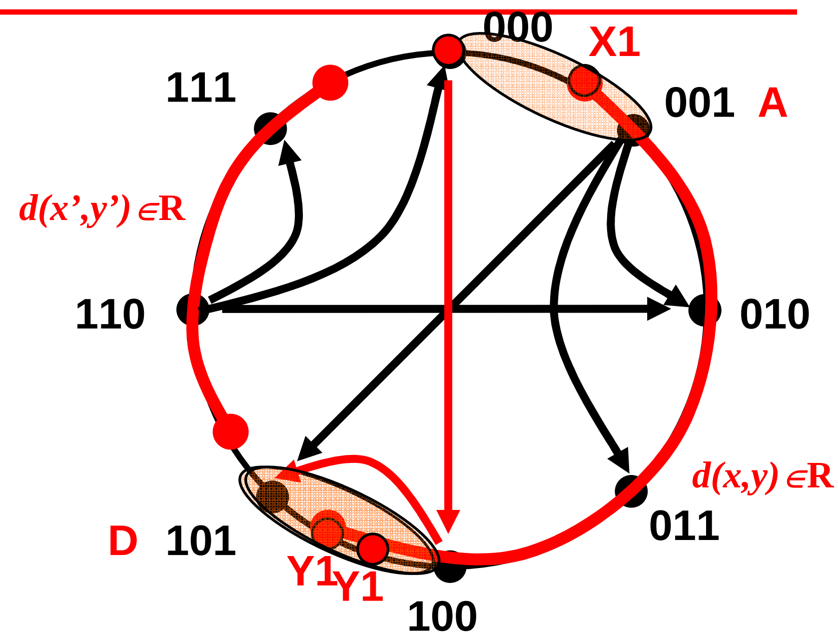

- **Six key design aspects**
  - Choice of an identifier space (I,d)
  - Mapping of peers (F_P) and resources (F_R) to the identifier space
  - Management of the identifier space by the peers (M)
  - Graph embedding (structure of the logical network) G=(P,E) (N - Neighborhood relationship)
  - Routing strategy (R)
  - Maintenance strategy

Hierarchical Overlay Networks

- Popular Example: Napster, Kaaza
- Superpeers form a structured or unstructured overlay network
- Normal peers attach as clients to superpeers

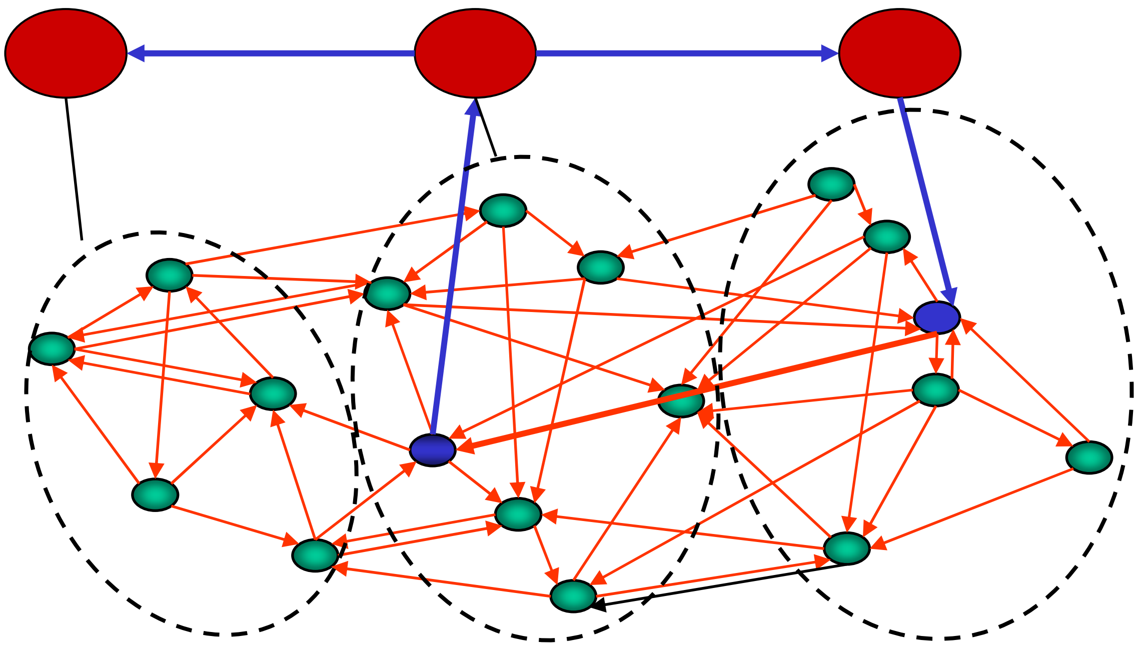

Beyond Keyword Search

⇒ searching semantically richer objects
in overlay networks

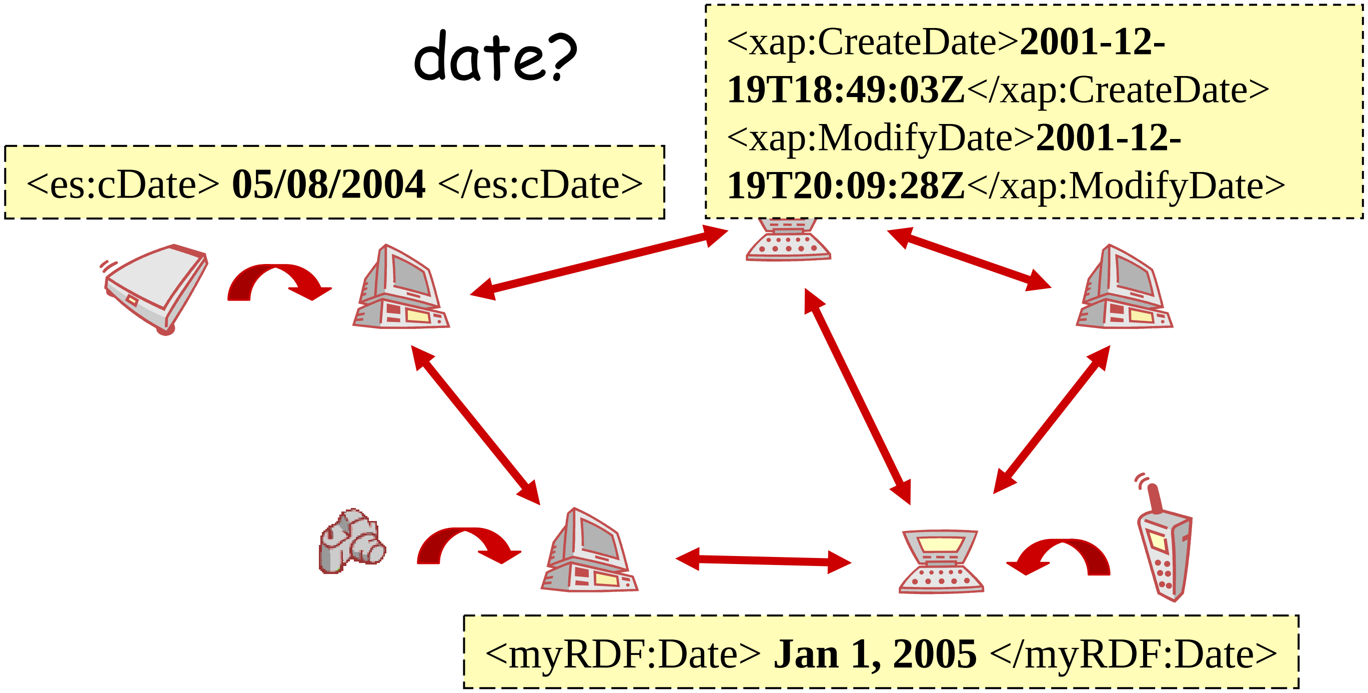

Managing Heterogeneous Data

- Support of structured data at peers: schemas
- Structured querying in peer-to-peer system
- Relate different schemas representing semantically similar information

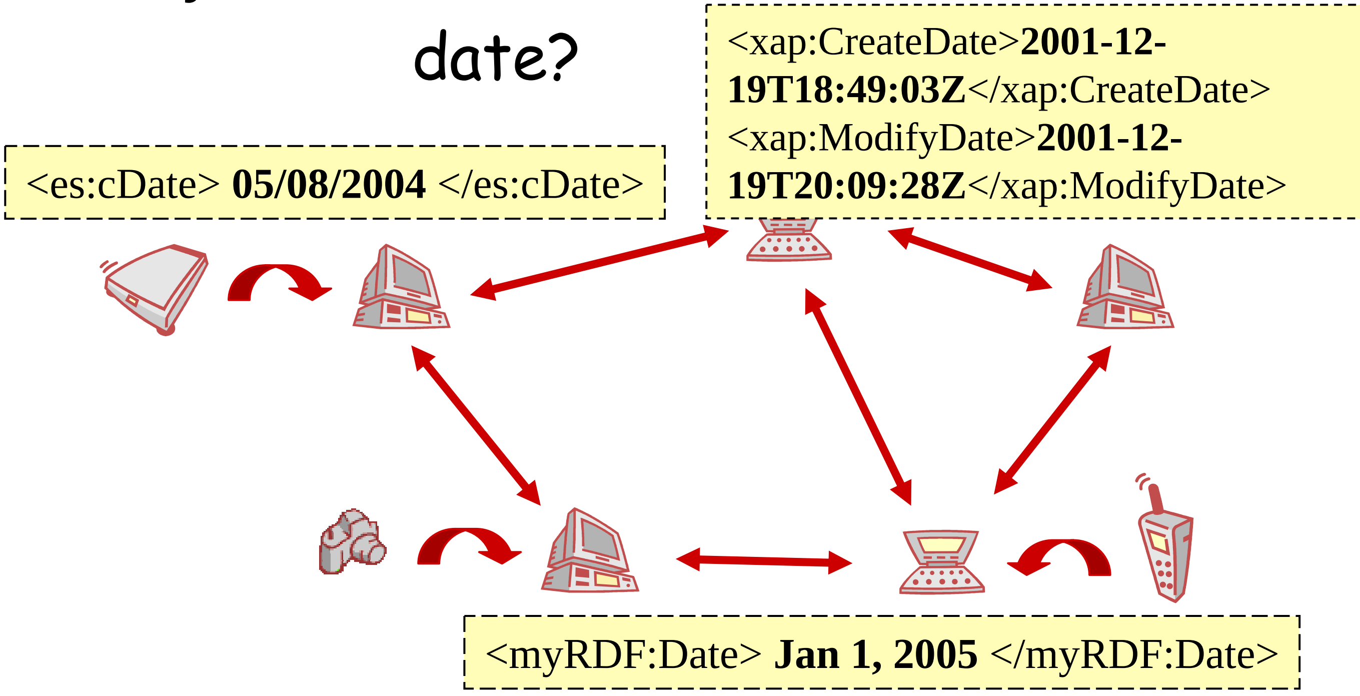

II. Query Evaluation in SONs

Beyond keyword search

$\Rightarrow$ searching complex structured data
in overlay networks

Standard RDMS over overlay networks

- Strictly speaking impossible
- CAP theorem: pick at most two of the following:
  1. Consistency
  2. Availability
  3. Tolerance to network Partitions
- Practical compromises:
  - ⇒ Relaxing ACID properties
    - ⇒ *Soft-states*: states that expire if not refreshed within a predetermined, but configurable amount of time

S. Gilbert and N. Lynch: Brewer's conjecture and the feasibility of consistent, available, partition-tolerant web services, ACM SIGACT News, 33(2), 2002.

Distributed Hash Table Lookup

- DHT lookups designed for binary relations (key,content)
- Structured data (e.g., RDF statements) can sometimes be encoded in simple, rigid models

Index attributes to resolve queries as distributed table lookups

t = (<info:rdfpeers> <dc:creator> <info:MinCai>)

Key 1 → info:rdfpeers

Key 2 → dc:creator

Key 3 → info:MinCai

RDFPeers: A distributed RDF repository

Who?
- *U.S.C. (Information Sciences Institute)*

Overlay structure
- *DHT (MAAN [Chord])*

Data model
- *RDF*

Queries
- *RDQL*

Query evaluation
- *Distributed (iterative lookup)*

RDFPeers Architecture

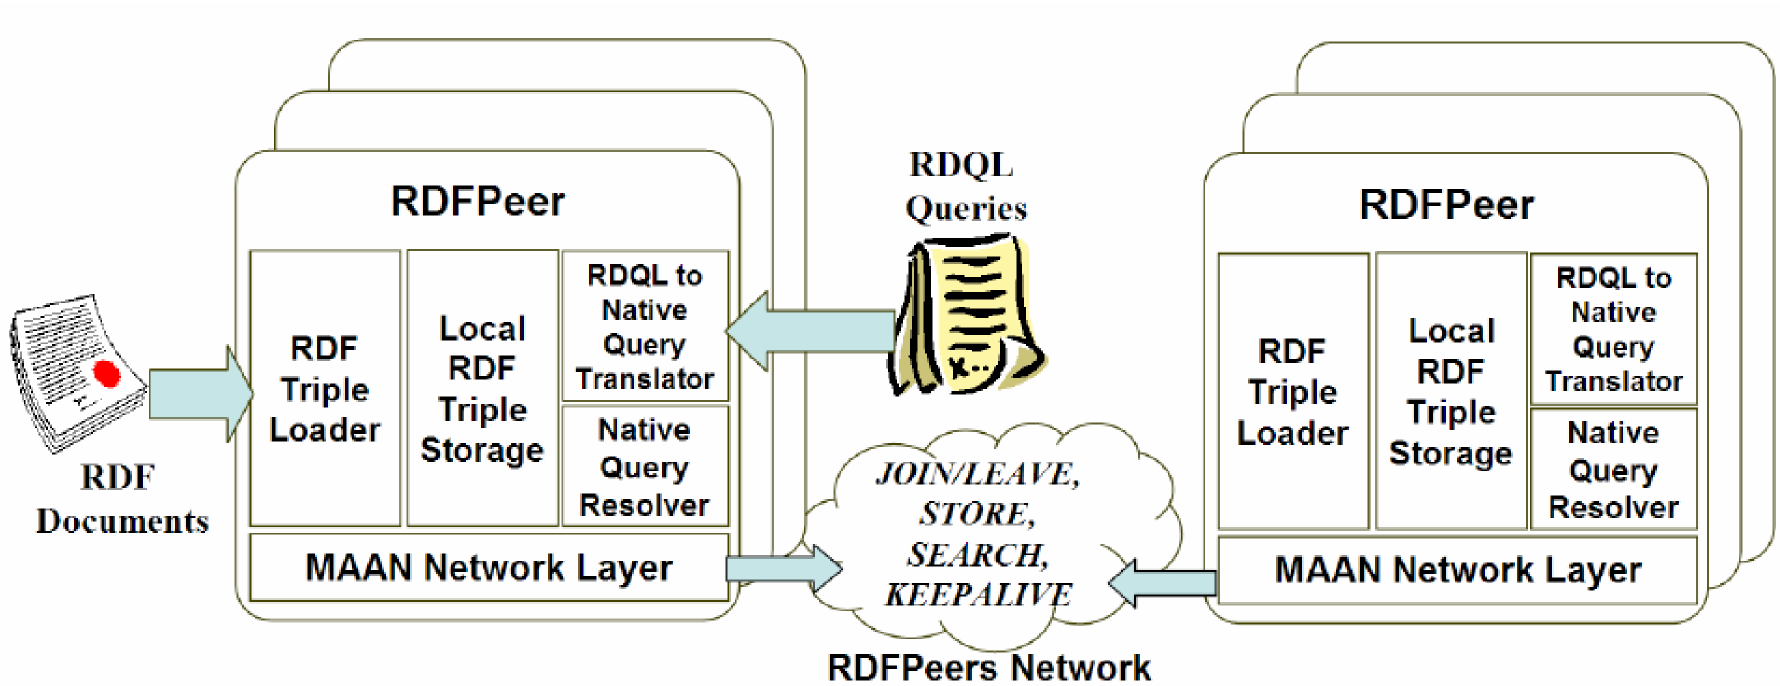

Index Creation (1)

Triple t = <info:rdfpeers> <dc:creator> <info:mincai>

Put(Hash(**info:rdfpeers**), t)

Put(Hash(**dc:creator**), t)

Put(Hash(**info:mincai**), t)

- Soft-states
  - Each triple has an expiration time
- Locality-preserving hash-function
  - Range searches

Index Creation (2)

URI / Literal	Hash Value in [0, 15]
<info:rdfpeers>	13
<info:mincai>	1
<dc:creator>	5
<foaf:name>	4
<foaf:age>	10
“Min Cai”	7
“28”	2

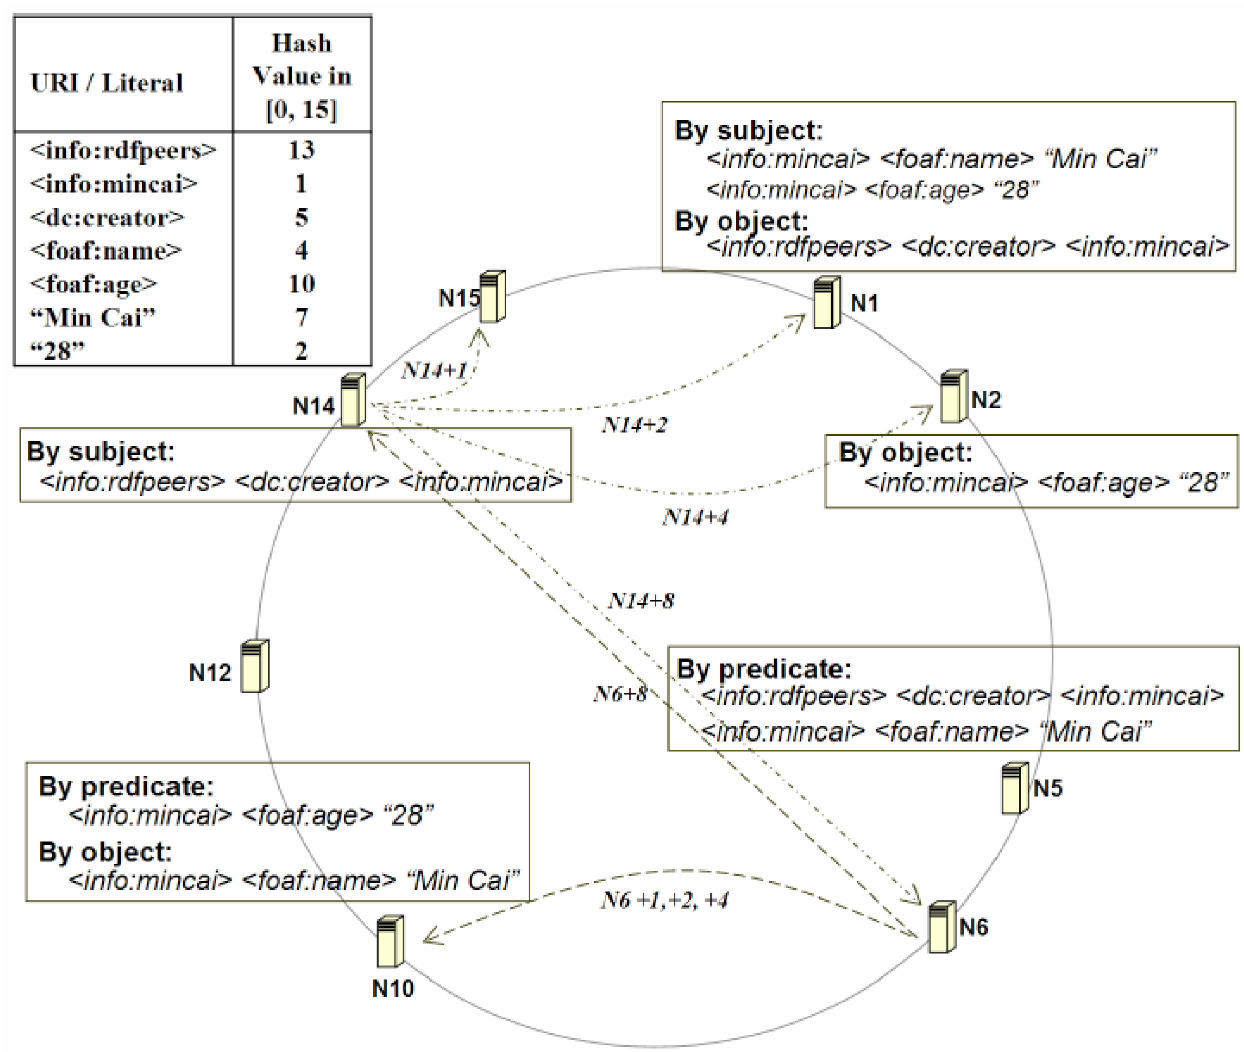

Query Evaluation

- Iterative, distributed table lookup

(?x, <rdf:type>, <foaf:Person>)
(?x, <foaf:name>, "John")

2) Results = $\pi_{subject}\sigma_{predicate=rdf.type,\ object=foaf:Person}(R)$

1) Get(foaf:Person)

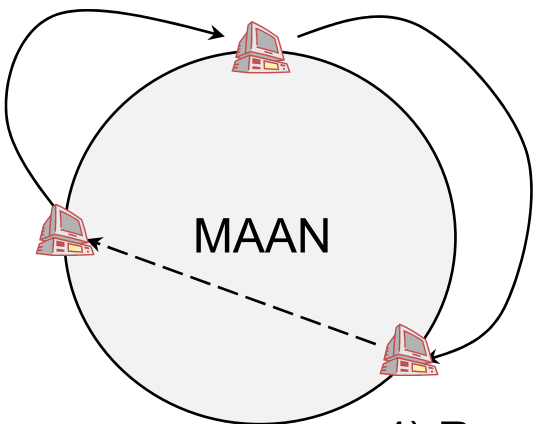

3) Get("John")

4) Results =

Results $\cap\ \pi_{subject}\sigma_{predicate=foaf:name,\ object=\text{"John"}}(R)$

Want more? Distributed RDF Notifications

- Pub/Sub system on top of RDFPeers
- Subscription = triple pattern with at least one constant term
  - Routed to the peer *P* responsible of the term
  - *P* keeps a local list of subscriptions
  - Fires notifications as soon as a triple matching the pattern gets indexed
- Extensions for disjunctive and range subscriptions

References

- M. Cai, M. Frank, J. Chen, and P. Szekely. Maan: A mulit-attribute addressable network for grid information services. Journal of Grid Computing, 2(1), 2004.

- M. Cai and M. Frank. Rdfpeers: A scalable distributed rdf repository based on a structured peer-to-peer network. In International World Wide Web Conference(WWW), 2004.

- M. Cai, M. Frank, B. Pan, and R. MacGregor. A subscribable peer-to-peer rdf repository for distributed metadata management. Journal of Web Semantics, 2(2), 2005.

DHT-Based RDMS

.

- Traditional DHTs only support keyword lookups
- Traditional RDMS do no scale gracefully with the number of nodes

Scaling-up RDMS over a DHT

– Distributing storage load

– Distributing query load

⇒ Relaxing ACID properties

The PIER Project

Who?
- *U.C. Berkeley*

Overlay structure
- *DHT (currently Bamboo and Chord)*

Data model
- *Relational*

Queries
- *Relational, with joins and aggregation*

Query evaluation
- *Distributed (based on query plans)*

PIER Architecture

- Peer-to-peer Information Exchange and Retrieval
- Relational query processing system built on top of a DHT
- Query processing and storage are *decoupled*

⇒ Sacrificing strong consistency semantics

  - Best-Effort

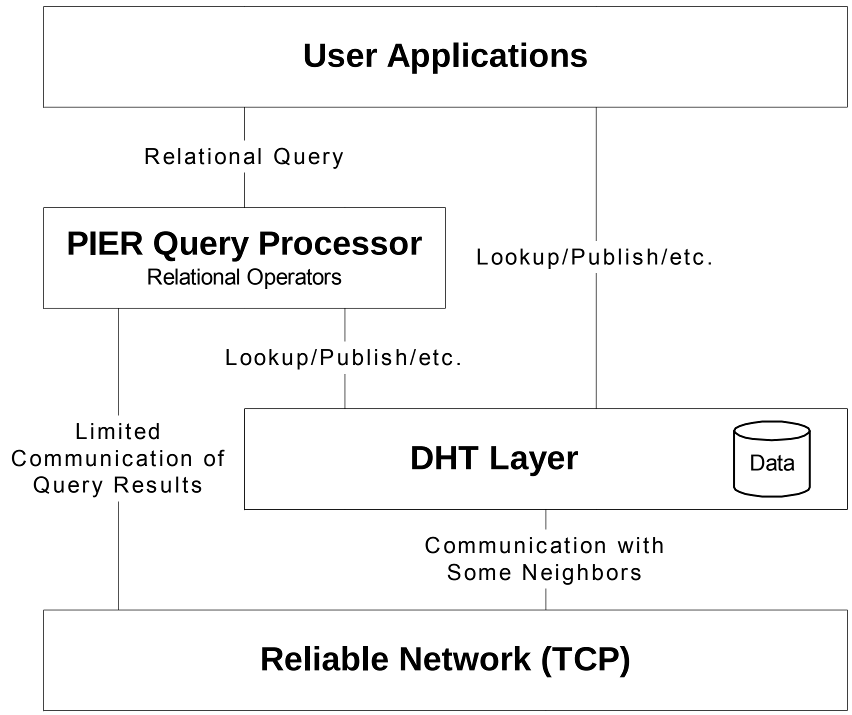

Main Index Creation: DHT Index

- Indexing tuples in the DHT (equality-predicate index)
  - Relation R1: {35, abc.mp3, classical, 1837,...}
  - Index on 3rd/4th attributes:
    - **hash key=**{R1.classical.1837,35}

resourceID namespace Partitioning key

- Soft-state storage model
  - Publishers periodically extend the lifetime of published objects
- No system metadata
  - All tuples are self-describing

Two Other Indexes

- Multicast index
  - Multicast tree created over the DHT
- Range index
  - Prefix hash tree created over the DHT

Query Evaluation

- Queries are expressed in an algebraic dataflow language
  - A query plan has to be provided
- PIER processes queries using three indexes
  - DHT index for equality predicates
  - Multicast index for query dissemination
  - Range index for predicates with ranges

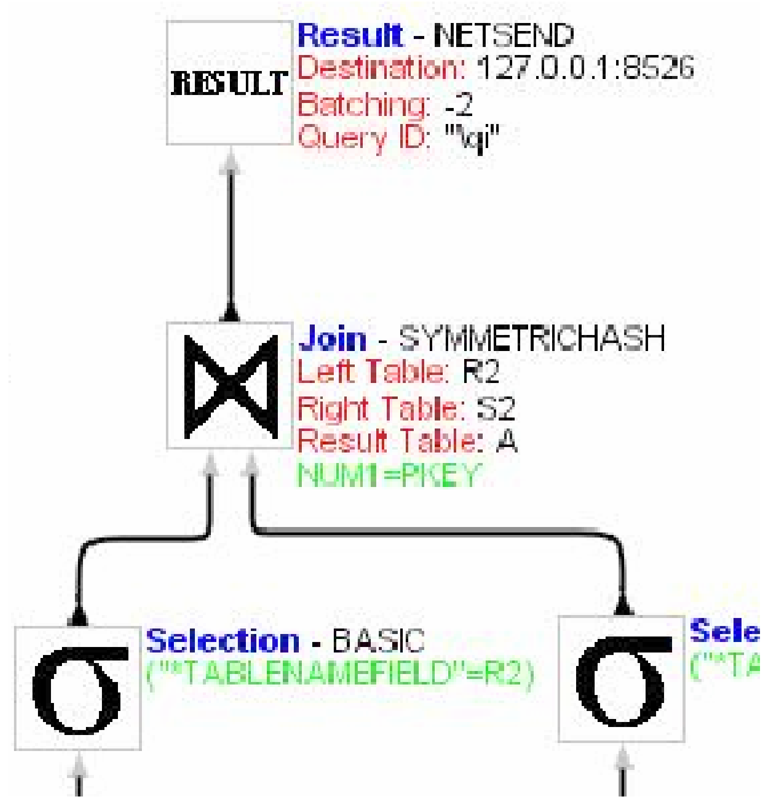

Symmetric hash join

- Equi-join on two tables R(A,B) and S(C,B)

1. Disseminate query to all nodes (multicast tree)
   - Find peers storing tuples from R or S
2. Peers storing tuples from R and S hash and insert the tuples based on the join attribute
   - Tuples inserted into the DHT with a temporary namespace
3. Nodes receiving tuples from R and S can create the join tuples
4. Output tuples are sent back to the originator of the query

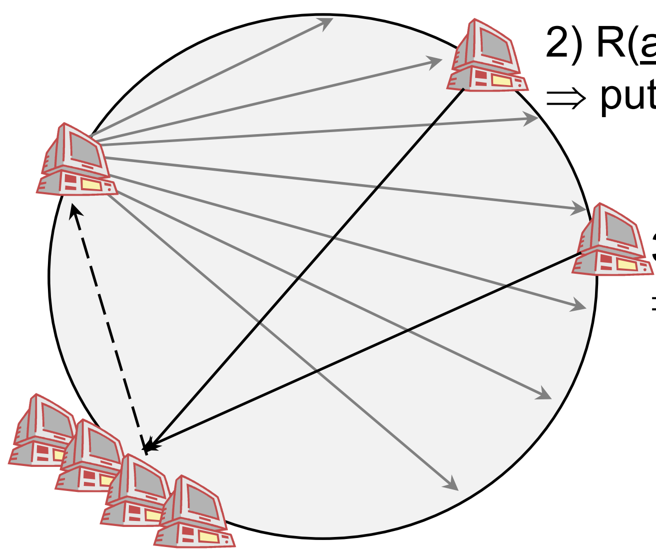

Want more? Join variants in PIER

- Skip rehashing
  - When one of the tables is already hashed on the join attribute in the equality-predicate index
- Symmetric semi-join rewrite
  - Tuples are projected on the join attribute before being rehashed
- Bloom filter rewrite
  - Each node creates a local Bloom filter and sends it to a temporary namespace
  - Local Bloom filters are OR-ed and multicast to nodes storing the other relations
  - Followed by a symmetric hash join, but only the tuples matching the filter are rehashed

References

- J. M. Hellerstein: Toward network data independence. SIGMOD Record 32(3), 2003
- R. Huebsch, J. M. Hellerstein, N. Lanham, B. Thau Loo, S. Shenker, and I. Stoica. Querying the internet with pier. In International Conference on Very Large Databases (VLDB), 2003.
- B. Thau Loo, J. M. Hellerstein, R. Huebsch, S. Shenker, and I. Stoica. Enhancing p2p file-sharing with an internet-scale query processor. In International Conference on Very Large Databases (VLDB), 2004.
- S. Ramabhadran, S. Ratnasamy, J. M. Hellerstein, and S. Shenker. Brief announcement: Prefix hash tree. In ACM PODC, 2004.
- R. Huebsch, B. Chun, J. M. Hellerstein, B. Thau Loo, P. Maniatis, T. Roscoe, S. Shenker, I. Stoica, and A. R. Yumerefendi. The architecture of pier: an internetscale query processor. In Biennial Conference on Innovative Data Systems Research (CIDR), 2005.

Routing Indices

- Flooding an overlay network with a query can be inefficient
- Disseminating a query often boils down to computing a multicast tree for a portion of the peers

Storing semantic routing information at various granularities directly at the peers

- Schema level
- Attribute level
- Value level

The Edutella Project

Who?

- *U. of Hannover (mainly)*

Overlay structure

- *Super-Peer (HyperCup)*

Data model

- *RDF/S*

Queries

- *Triple patterns (or TRIPLE)*

Query evaluation

- *Distributed (based on routing indices)*

Edutella Architecture

- An RDF-based infrastructure for P2P applications
- End-peers store resources annotated with RDF/S
- Super-peer architecture
  - HyperCup super-peer topology
  - Routing based on indices
  - Two-phase routing
    - Super-peer to super-peer
    - Super-peer to peer

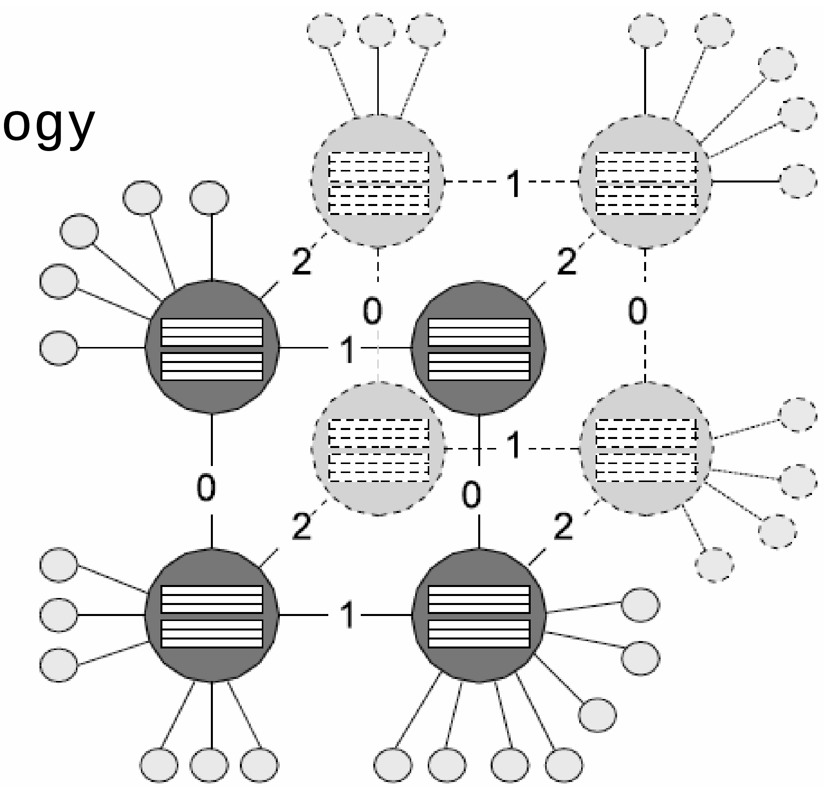

Index construction: SP/P routing indices

- Registration: end-peers send a summary of local resources to their super-peer
  - Schema names used in annotations
  - Property names used in annotations
  - Types of properties (ranges) used in annotations
  - Values of properties used in annotations
- Not all levels have have to be used
- Super-peers aggregate information received from their peers and create a local index

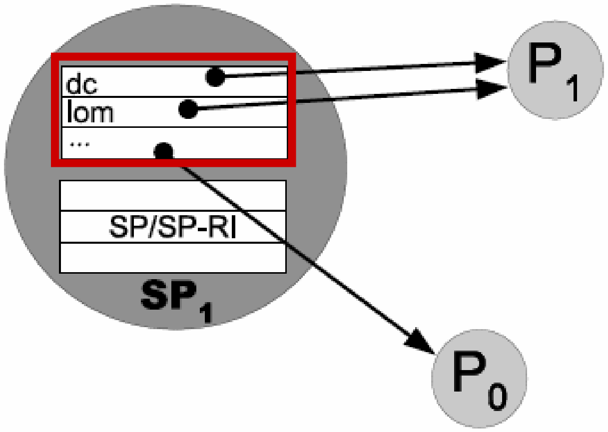

- Registration is periodic
  - Soft-states

Index Construction: SP/SP routing indices

- Super-peers propagate the SP/S indices to other super-peers with spanning trees

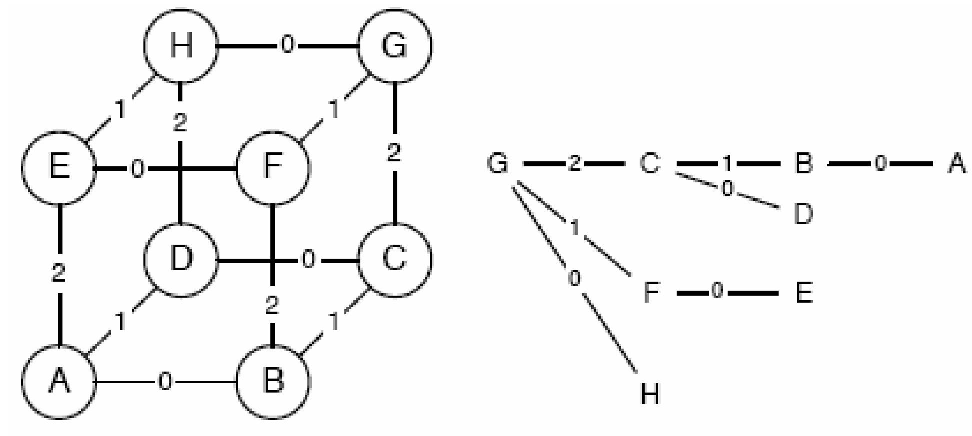

- Super-peers aggregate the information in SP/SP indices
  - Use of semantic hierarchies

Granularity	Index of A		
Schema	dc		B D E
	lom		B E
Property	dc:subject		B D E
	dc:language		B E
	lom:context		B E
Property Value Range	dc:subject	ccs:networks	D
	dc:subject	ccs:software-eng.	B E
Property Value	lom:context	"undergrad"	B E
	dc:language	"de"	B E

Query Evaluation

Q: (?, dc:language, “de”)
(?, lom:context, “undergrad”)
(?, dc:subject, ccs:softwareengineering)

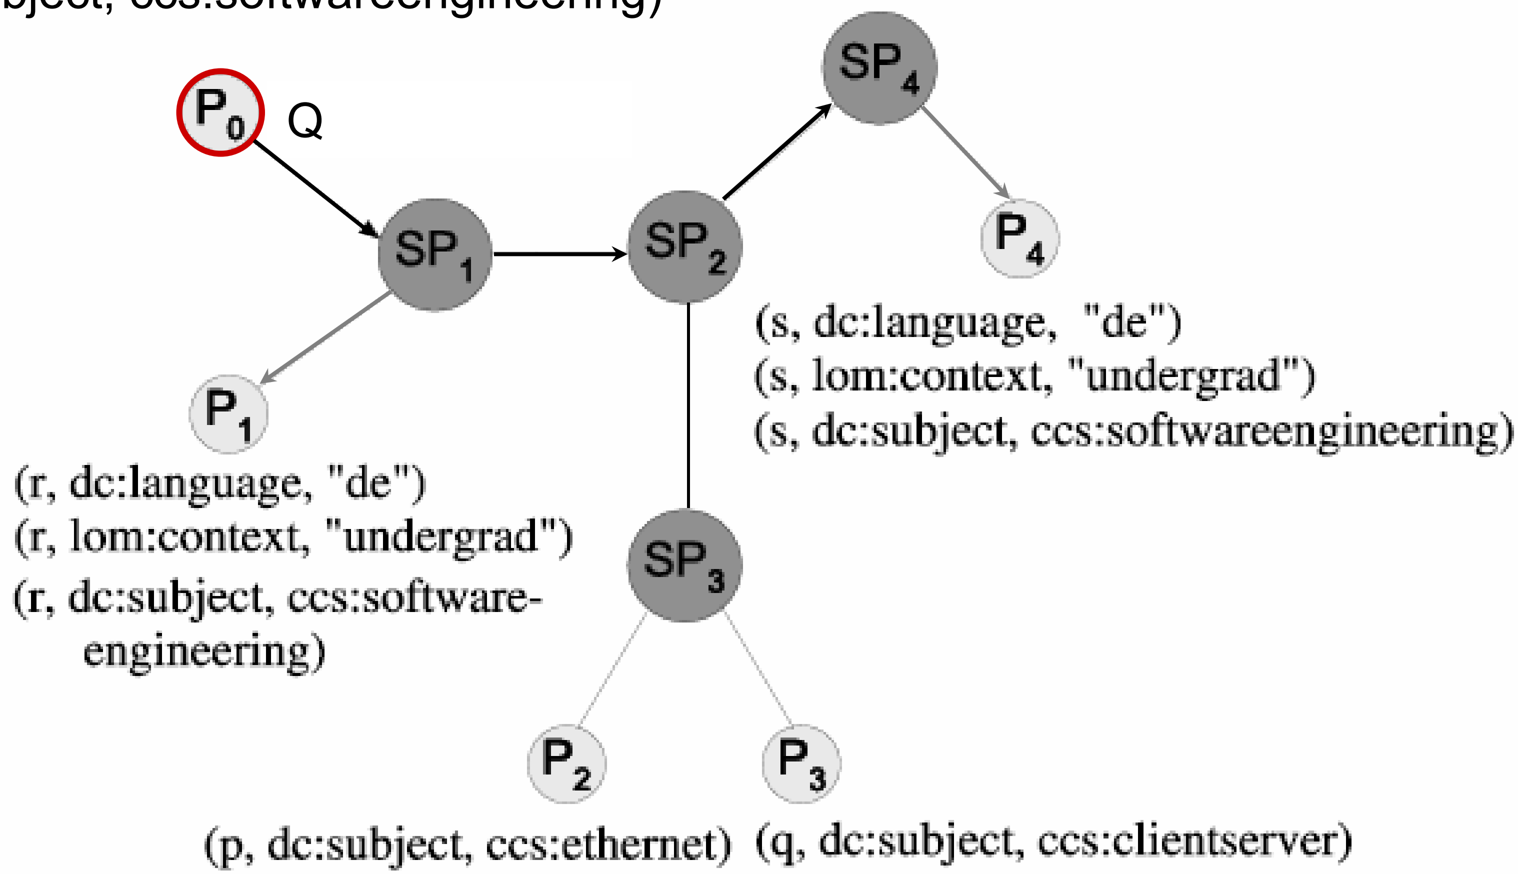

Want More? Decentralized Ranking

- Number of results returned grow with the size of the network
- Decentralized top-k ranking
  - New weight operator to specify which predicate is important
  - Aggregation of top-k in three stages
    - End-peer
    - Super-peer
    - Query originator

References

- W. Nejdl, B. Wolf, C. Qu, S. Decker, M. Sintek, A. Naeve, M. Nilsson, M. Palmer, and T. Risch. Edutella: a p2p networking infrastructure based on rdf. In International World Wide Web Conference (WWW), 2002.
- W. Nejdl, W. Siberski, and M. Sintek. Design issues and challenges for rdf- and schema-based peer-to-peer systems. SIGMOD Record, 32(3), 2003.
- W. Nejdl, M. Wolpers, W. Siberski, C. Schmitz, M. T. Schlosser, I. Brunkhorst, and A. Loser. Super-peer-based routing and clustering strategies for rdf-based peer-to-peer networks. In International World Wide Web Conference (WWW), 2003.
- W. Nejdl, M. Wolpers, W. Siberski, C. Schmitz, M. T. Schlosser, I. Brunkhorst, and A. Loser. Super-peer-based routing strategies for rdf-based peer-to-peer networks. Journal of Web Semantics, 2(2004), 1.
- W. Nejdl, W. Siberski, W. Thaden, and W. T. Balke. Top-k query evaluation for schema-based peer-to-peer networks. In International Semantic Web Conference (ISWC), 2004.
- H. Dhraief, A. Kemper, W. Nejdl, and C. Wiesner. Processing and optimization of complex queries in schema-based p2p-networks. In Workshop On Databases, Information Systems and Peer-to-Peer Computing (DBISP2P), 2004.
- M. T. Schlosser, M. Sintek, S. Decker, and W. Nejdl. Hypercup - hypercubes, ontologies, and efficient search on peer-to-peer networks. In International Workshop on Agent and P2P Computing (AP2PC), 2002.

III. Semantic Mediation in SONs

- What if (some) peers use different schemas to store semantically related data?
  - Need ways to relate schemas in decentralized settings

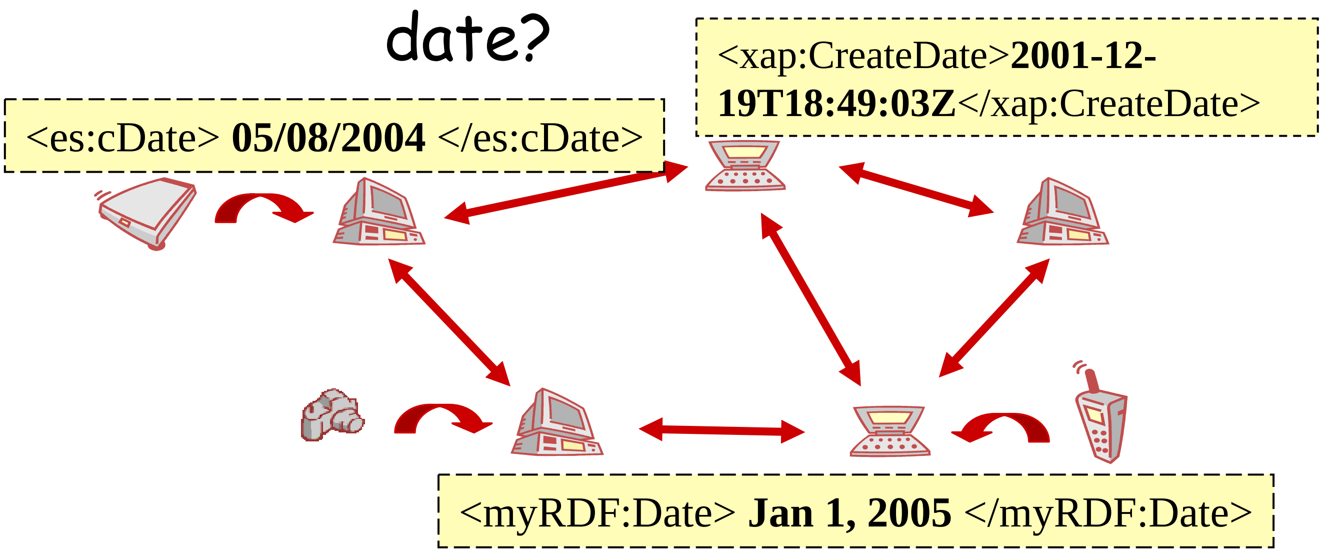

⇒ unstructured overlay network at the semantic layer

⇒ **Peer Data Management Systems (PDMS)**

Semantic Mediation Layer

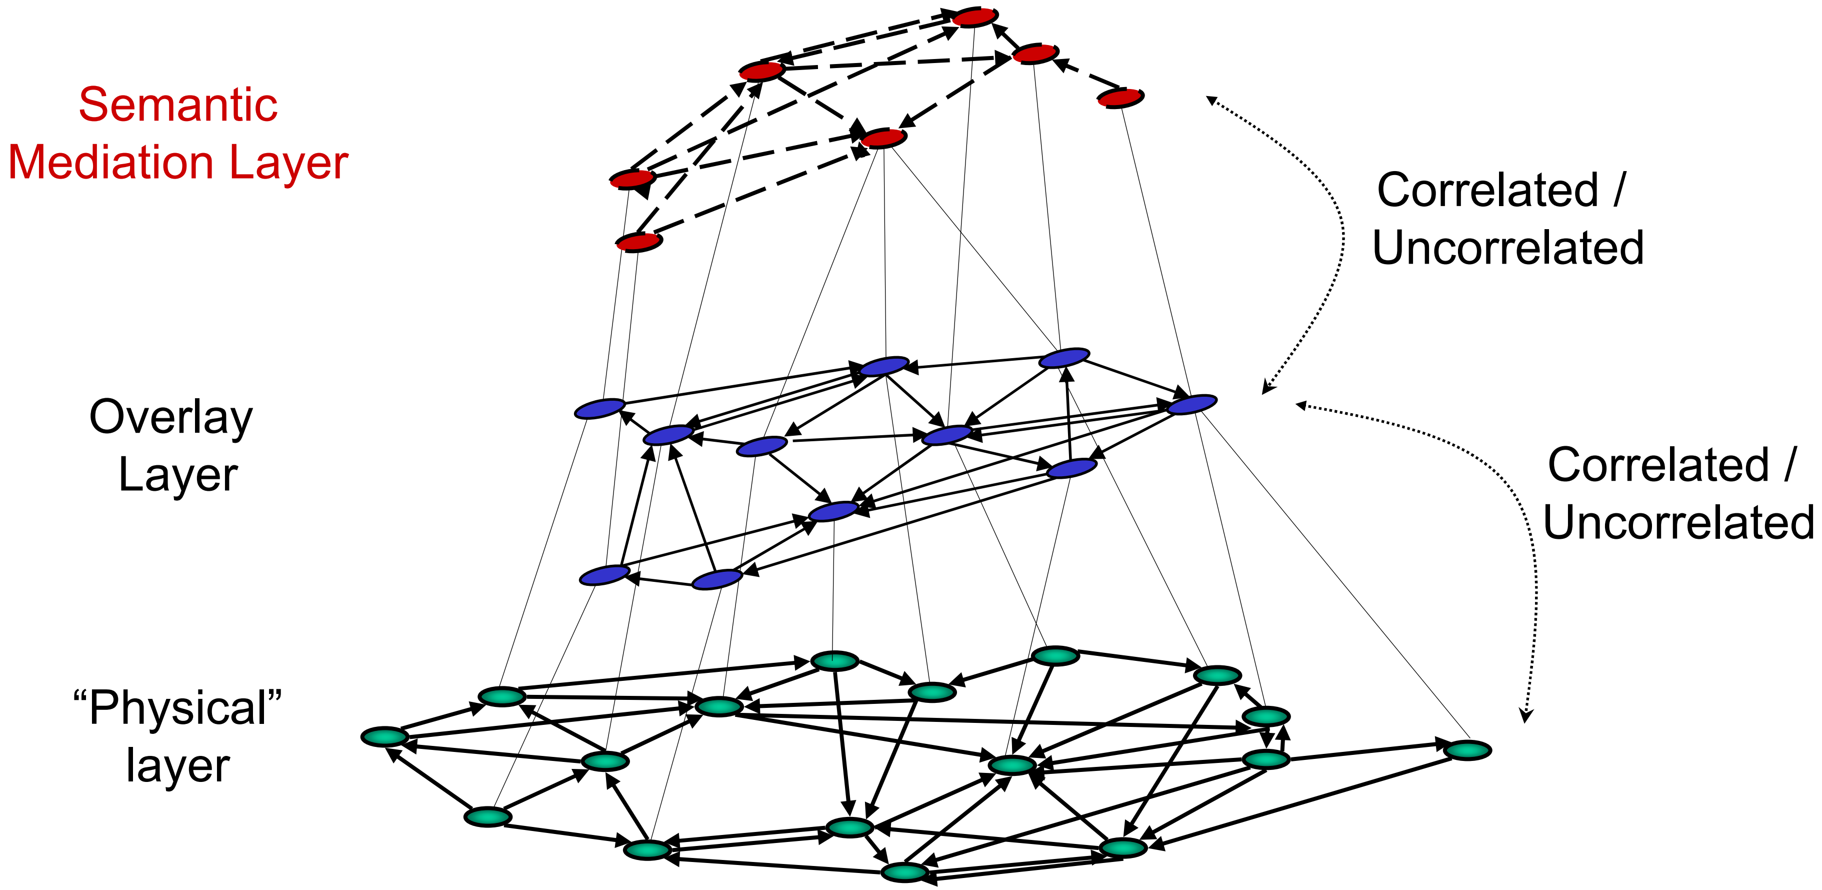

Source Descriptions

- Heterogeneous schemas can share semantically equivalent attributes
- On the web, users are willing to annotate resources or filter results manually

Let users annotate their schemas

- Search & Match similar annotations
- Use IR methods to rank matches
- Let users filter out results

PeerDB

Who?
- *National U. of Singapore*

Overlay structure
- *Unstructured (BestPeer)*

Data model
- *Relational*

Mappings
- *Keywords*

Query reformulation
- *Distributed*

Query evaluation
- *Distributed*

PeerDB architecture

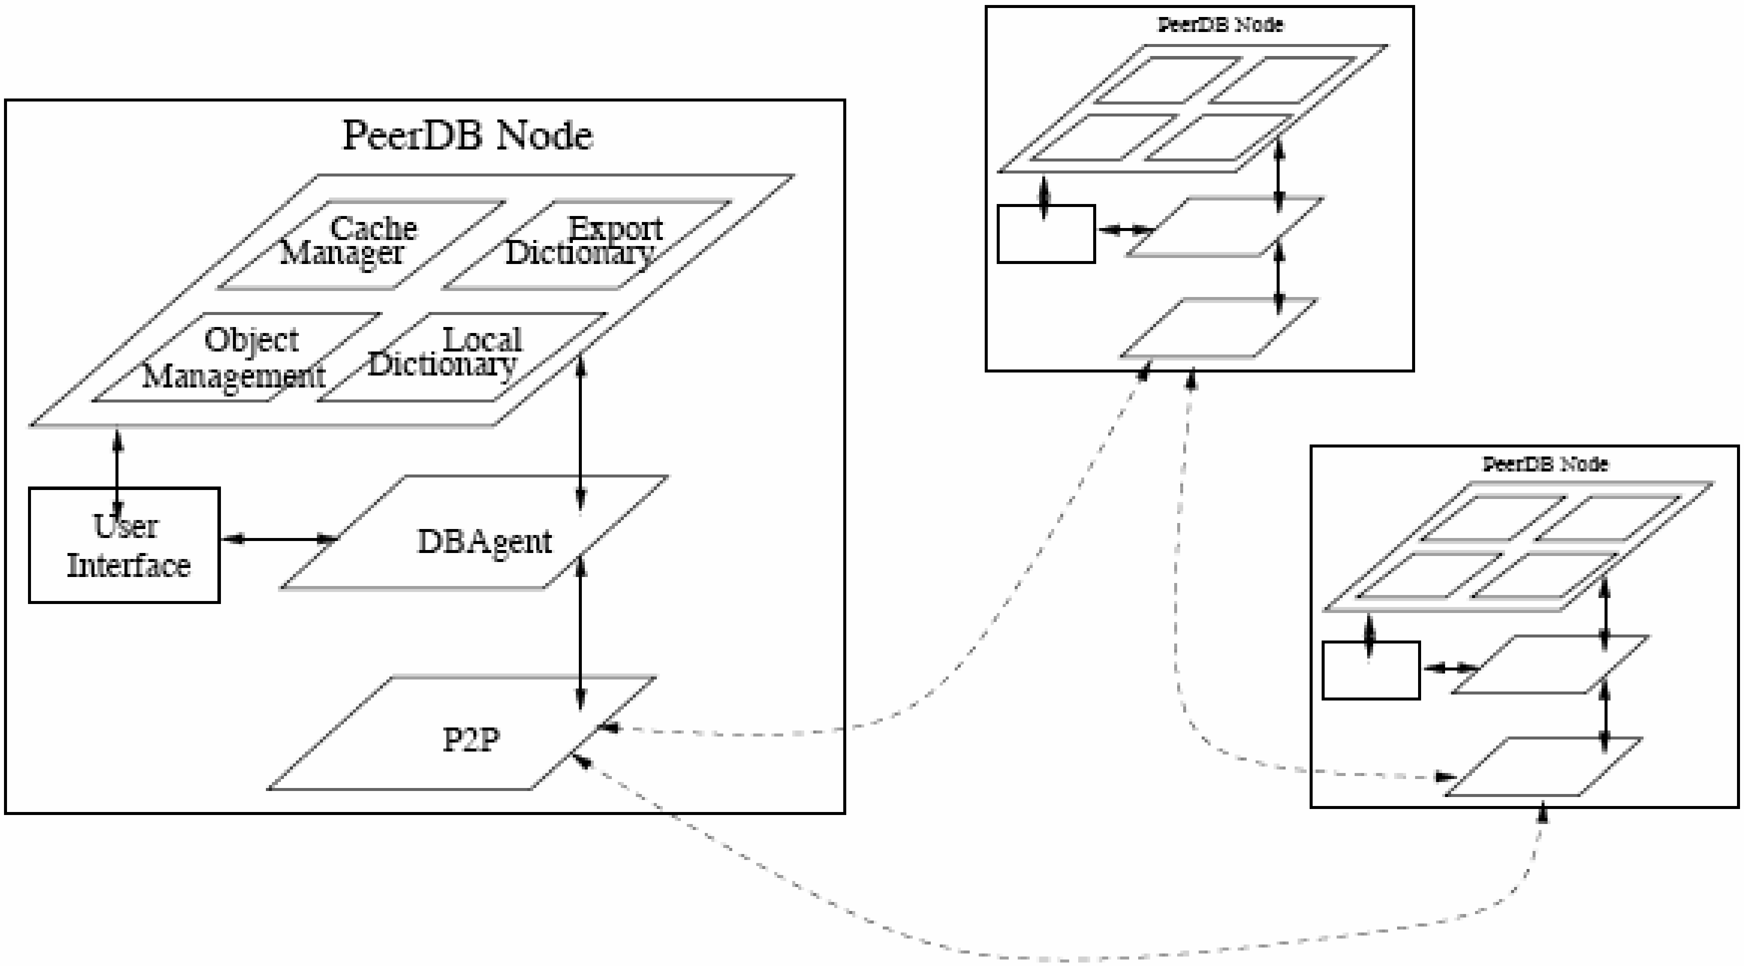

Index Construction

- Peers store keywords related to local relations / attributes

Names	Keywords
Kinases	protein, human
SeqID	key, identifier, ID
length	length
proteinSeq	sequence, protein sequence

Attribute names

Provided by experts

Query Reformulation (1)

- Local query *Q(R,A)*
  - R: set of local relations
  - A: set of local attributes

- Agents carrying the query are sent to neighbors

- Relations *D* from neighboring peers are ranked w.r.t. a matching function *Match(Q,D)*
  - Higher matching values if R's keywords can be matched to relation names / keywords of the neighbor
  - Higher matching values if A's keywords can be matched to attributes names / keywords of the neighbor

Query Reformulation (2)

- Promising relations with *Match(Q,D) > threshold* are returned to the user (query originator)
  - User filters out false positives manually at the relation level
- At the neighbor, the agent reformulates the query with local synonyms for R, A
  - Attributes might be dropped if no synonym is found
  - Results are returned to the query originator
- Query is forwarded iteratively in this manner with a certain TTL

Want More? Network Reconfiguration

- Result performance depends on the semantic clustering of the network
- PeerDB network is reconfigurable according to three strategies:
  - MaxCount
    - Choose as direct neighbors the peers which have returned the most answers (tuples / bytes)
  - MinHops
    - Choose as direct neighbors those peers which returned answers from the furthest locations
  - TempLoc
    - Choose as direct neighbors those peers that have recently provided answers.

References

- W. Siong Ng, B. Chin Ooi, K. L. Tan, and A. Ying Zhou. Bestpeer: A selfconfigurable peer-to-peer system.In International Conference on Data Engineering (ICDE), 2002.
- B. Chin Ooi, Y. Shu, and K. L. Tan. Db-enabled peers for managing distributed data. In Asian-Pacific Web Conference (APWeb), 2003
- B. Chin Ooi, Y. Shu, and K. L. Tan. Relational data sharing in peer-based data management systems. SIGMOD Record, 32(3), 2003.
- W. Siong Ng, B. Chin Ooi, K. L. Tan, and A. Ying Zhou. Peerdb: A p2p-based system for distributed data sharing. In International Conference on Data Engineering (ICDE), 2003.

Mapping Tables

- Semantically equivalent data values can often be mapped easily one onto the other

Specification of P2P mappings at the data value level

- Reformulate queries based on these mapping tables

Ids from the GDB relation at Peer P1

GDB_id	SwissProt_id
GDB:120231	Q14930
GDB:120231	Q9UMK3

Semantically equivalent Ids from SwissProt relation at peer P2

The Hyperion Project

Who?

- *U. of Toronto*
- *U. of Ottawa*
- *U. of Edinburgh*
- *U. of Trento*

Overlay structure

- *Unstructured*

Data model

- *Relational*

Queries

- *S+J algebra with projection*

Query reformulation

- *Distributed*

Query evaluation

- *Distributed*

Hyperion: Architecture

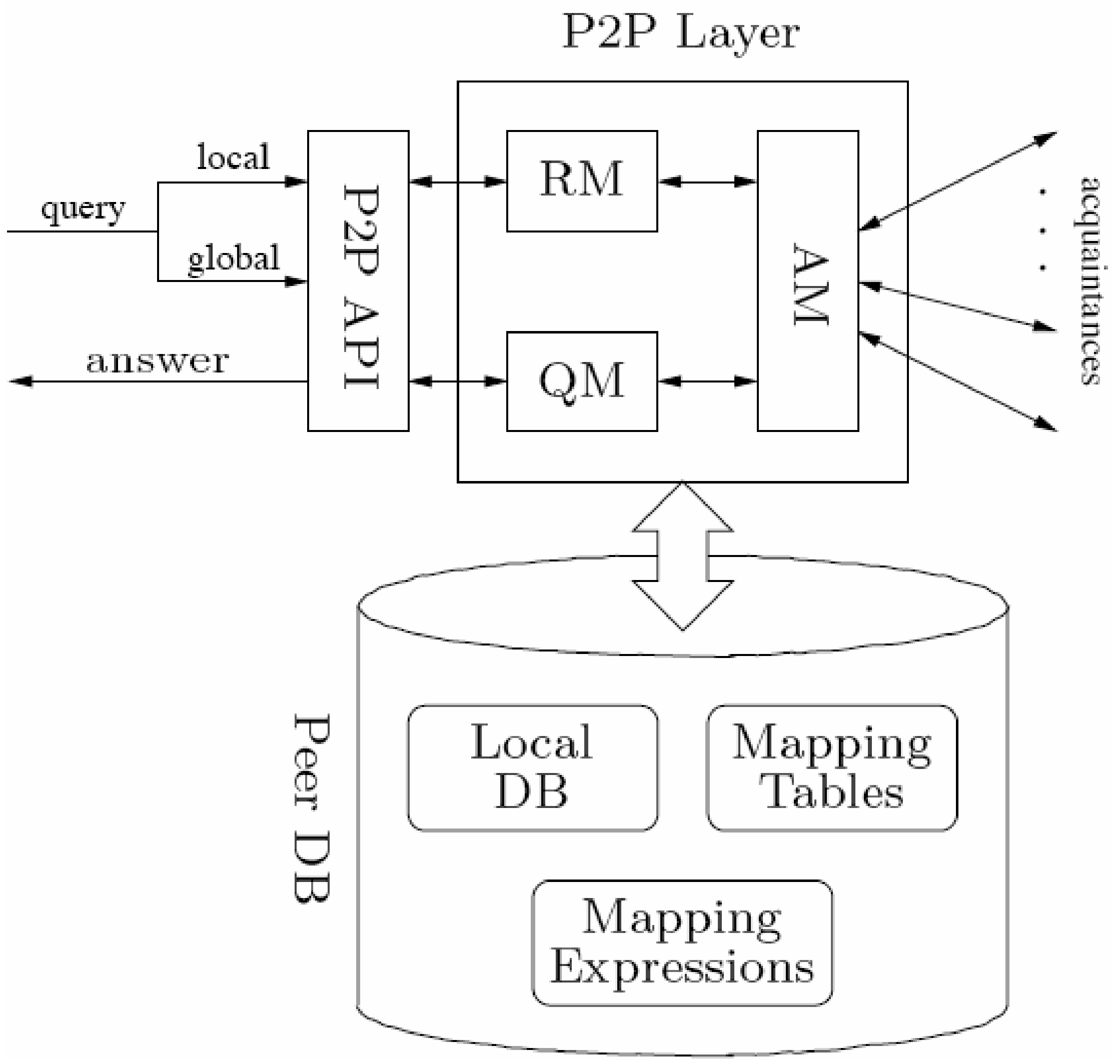

Creating Mapping Tables

- Initially created by domain experts
- Mapping tables semantics:

A	B
X_i	Y_j

	Open-world	Closed-world
present X-value	Any Y-value	indicated Y-values
missing X-value	Any Y-value	No Y-value

  - *Closed-open-world* semantics
    - Partial knowledge
  - *Closed-closed-world* semantics
    - Complete information
- Common associations, e.g., identity mappings, can be expressed with unbound variables

$year$	py
$\mathcal{X}$	$\mathcal{X}$

- Efficient algorithm to infer new mappings or check consistency of a set of mappings along a path

Query Reformulation

- Query posed over local relations only
  - S+J algebra with projection
- Iterative distributed reformulations
  - P2P propagation based on acquaintance links
- Local algorithm ensures sound and complete reformulation of query q_1 at P_1 to query q_2 at P_2
  - Soundness: only values that can be related to those retrieved at P_1 are retrieved at P_2
  - Completeness: retrieving all possible sound values

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

Query Reformulation with multiple tables

- Transform the query in its equivalent disjunctive normal form and pick the relevant tables only

keyword	*kw*
OPH	APH
OPH	AARE
NGF receptor	p75 ICD
G9 sialidase	Sialidase 1

article_id	*paper_id*
20185348	10719179

year	*py*
$\mathcal{X}$	$\mathcal{X}$

$$\begin{array}{lll} Q_9: & \textbf{\textit{select}} & * \\ & \textbf{\textit{from}} & MedLine \\ & \textbf{\textit{where}} & keyword = \text{“OPH”}\ \textbf{\textit{AND}}\ year = \text{“1999”} \end{array}$$

$$\begin{array}{lll} Q_{10}: & \textbf{\textit{select}} & * \\ & \textbf{\textit{from}} & PubMed \\ & \textbf{\textit{where}} & (kw = \text{“APH”}\ \textbf{\textit{OR}}\ kw = \text{“AARE”}) \\ & & \textbf{\textit{AND}}\ py = \text{“1999”} \end{array}$$

Want More? Distributed E.C.A. Rules

- When views between schemas are defined, Consistency can also be ensured via a distributed rule system
  - Event-Condition-Action rule language and execution engine
  - Events, conditions and actions refer to multiple peers

$$AA\_Passenger(p, n) \supseteq BA\_Passenger(p, n)$$

```
create trigger passengerInsertion
after     insert on BA_Passenger
          referencing new as NewPass
for each row
begin
          insert into AA_Passenger values NewPass
          in Alpha-Air DB;
end
```

References

- P. A. Bernstein, F. Giunchiglia, A.s Kementsietsidis, J. Mylopoulos, L. Serafini and I. Zaihrayeu. Data Management for Peer-to-Peer Computing: A Vision. In WebDB 2002.
- A. Kementsietsidis, M. Arenas, and R. J. Miller. Managing data mappings in the hyperion project. In International Conference on Data Engineering (ICDE), 2003.
- A. Kementsietsidis, M. Arenas, and R. J. Miller. Mapping data in peer-to-peer systems: Semantics and algorithmic issues. In ACM SIGMOD, 2003.
- M. Arenas, V. Kantere, A. Kementsietsidis, I. Kiringa, R. J. Miller, and J. Mylopoulos. The hyperion project: From data integration to data coordination. SIGMOD Record, 32(3), 2003.
- V. Kantere, I. Kiringa, J. Mylopoulos, A. Kementsietsidis, and M. Arenas. Coordinating peer databases using eca rules. In International Workshop on Databases, Information Systems and Peer-to-Peer Computing (DBISP2P), 2003.
- A. Kementsietsidis and M. Arenas. Data sharing through query translation in autonomous sources. In International Conference on Very Large Data Bases (VLDB), 2004.

Extending Data Integration Techniques

- Centralized data integration techniques take advantage of views to reformulate queries in efficient ways

Extending query reformulation using views to semantically decentralized settings

The Piazza Project

Who?
- *U. of Washington*

Overlay structure
- *Unstructured*

Data model
- *Relational (+XML)*

Queries
- *Relational*

Query reformulation
- *Centralized*

Query evaluation
- *Distributed*

An example of semantic topology

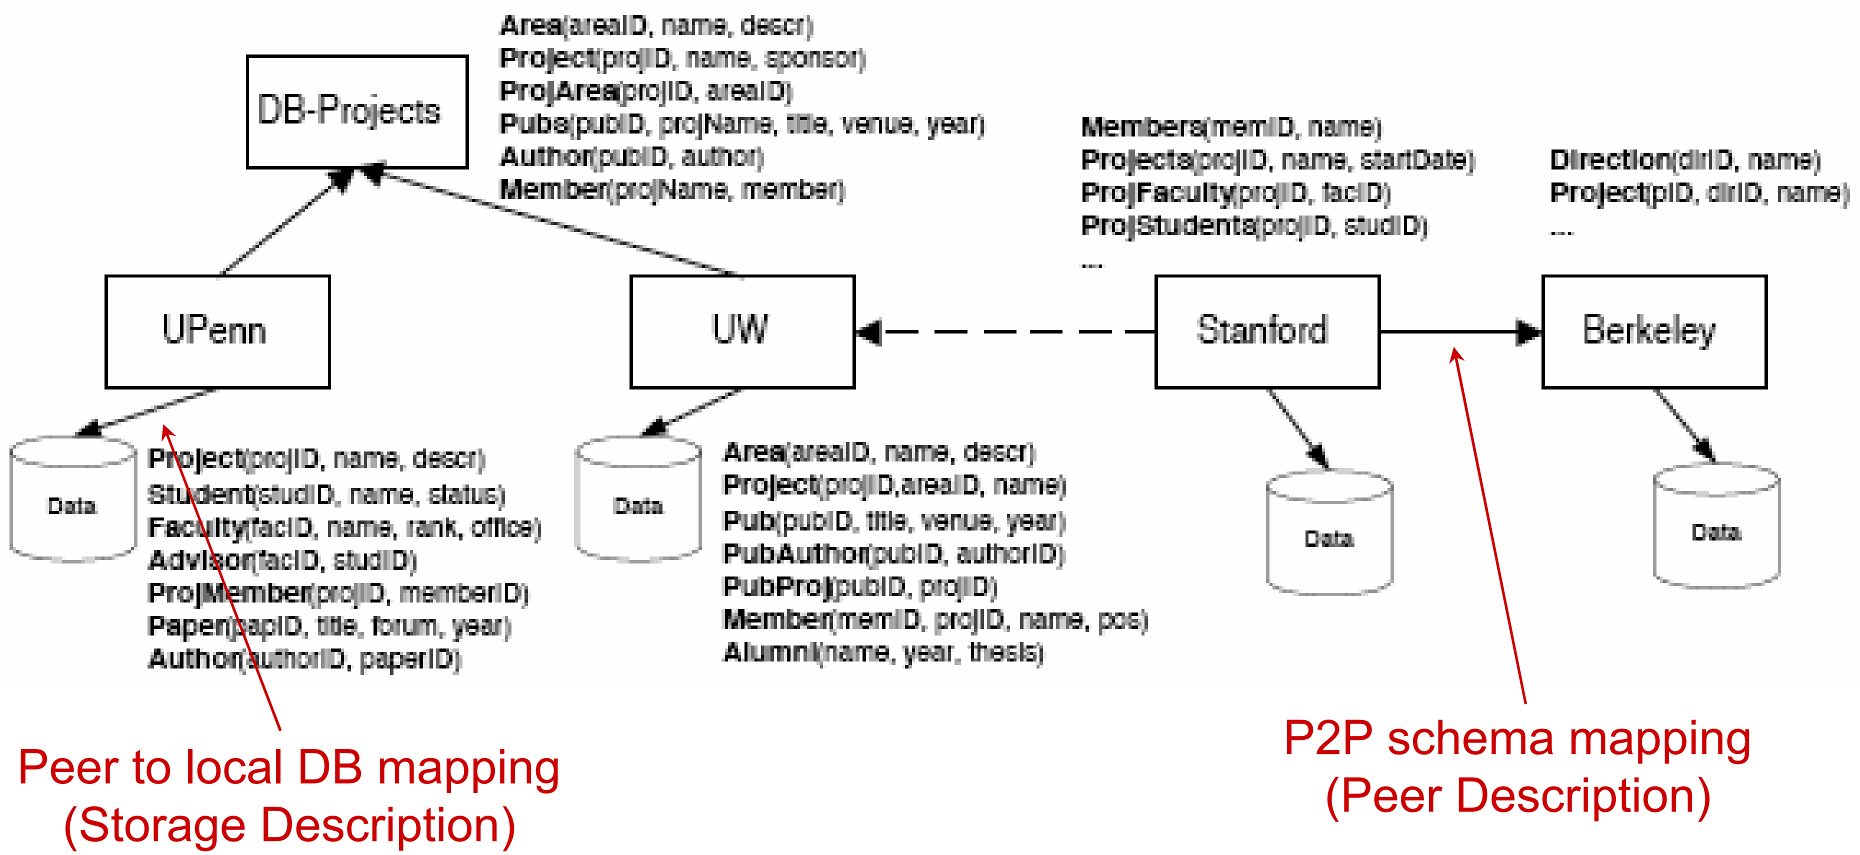

Creating Mappings in Piazza

- Mappings = views over the relations
  - Cf. classical data integration
- Supported mappings:
  - Definitions (GAV-like)

$$\begin{aligned} \text{9DC : SkilledPerson(PID, “EMT”)} \quad :- \\ \text{FS : Schedule(PID, vid)}, \\ \text{FS : 1stResponse(vid, s, l, d)}, \\ \text{FS : Skills(PID, “medical”)} \end{aligned}$$

  - Inclusions (LAV-like)

$$\begin{aligned} \text{LH : CritBed(bed, hosp, room, PID, status)} \quad \subseteq \\ \text{H : CritBed(bed, hosp, room)}, \\ \text{H : Patient(PID, bed, status)} \end{aligned}$$

Posing queries in Piazza

- Local query iteratively reformulated using the mappings

- Reformulation algorithm
  - Input: a set of mappings and a conjunctive query expression Q (evt. with comparison predicates)
  - Output: a query expression Q' that only refers to stored relations at the peers

- Reformulation is centralized

Query reformulation in Piazza

- Constructing a rule-goal tree:

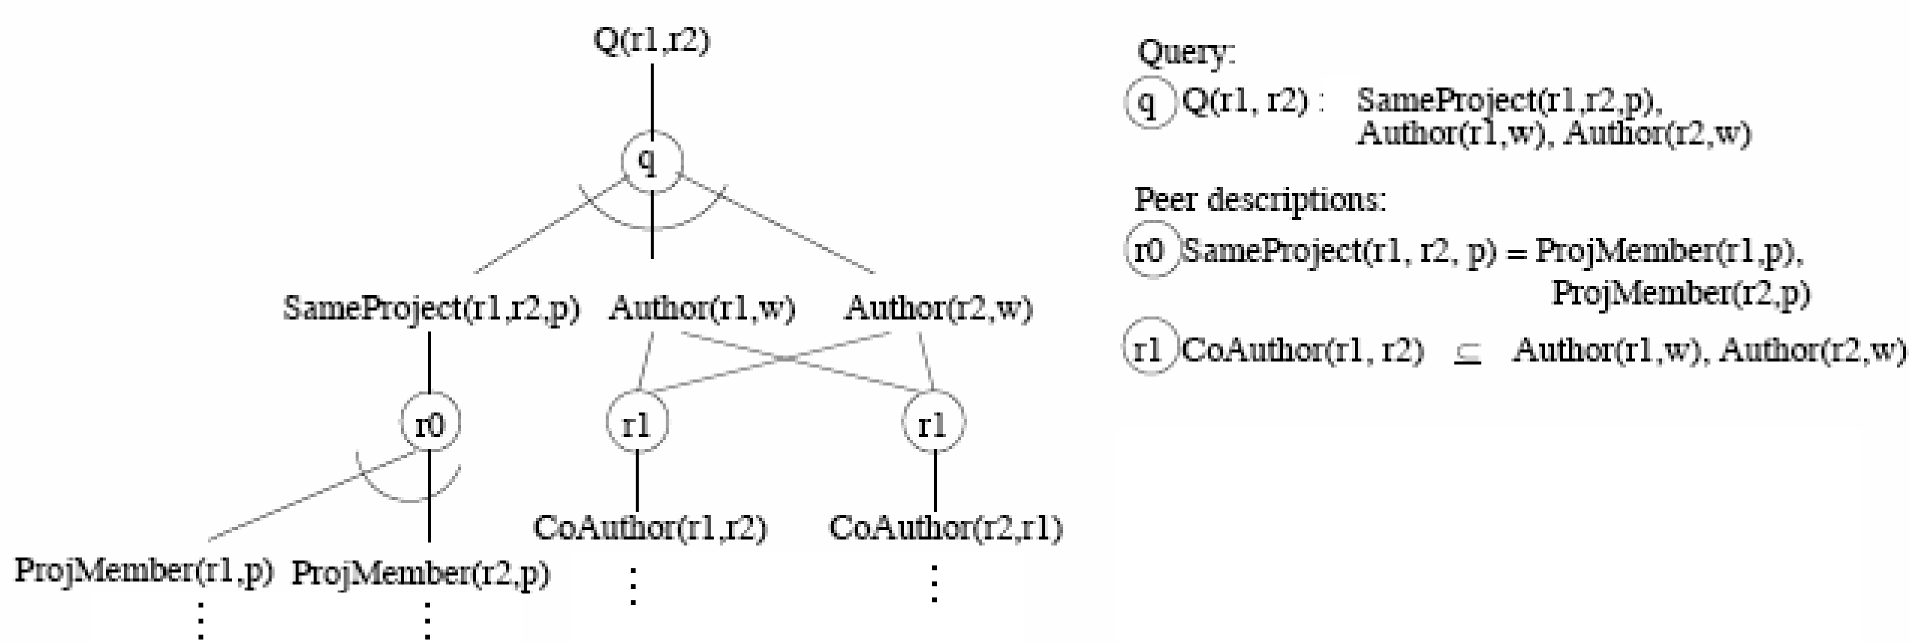

Reformulated query:

Q'(r1,r2): ProjMember(r1,p),ProjMember(r2,p),CoAuthor(r1,r2) U
ProjMember(r1,p),ProjMember(r2,p),CoAuthor(r2,r1)

More? Piazza & XML

- Piazza also considers query reformulation for semi-structured XML documents
- Mappings expressed with a subset of XQuery
  - Composition of XML mappings
- Containment of XML queries

References

- A. Y. Halevy, Z. G. Ives, P. Mork, and I. Tatarinov. Schema mediation in peer data management systems. In International Conference on Data Engineering (ICDE), 2003.
- A. Y. Halevy, Z. G. Ives, P. Mork, and I. Tatarinov. Peer data management systems: Infrastructure for the semantic web. In International World Wide Web Conference (WWW), 2003.
- I. Tatarinov, Z. Ives, J. Madhavan, A. Halevy, D. Suciu, N. Dalvi, X. Dong, Y. Kadiyska, G. Miklau, and P. Mork. The piazza peer data management project. SIGMOD Record, 32(3), 2003.
- I. Tatarinov and A. Halevy. Efficient query reformulation in peer data management systems. In ACM SIGMOD, 2004.
- X. Dong, A. Y. Halevy, and I. Tatarinov. Containment of nested xml queries. In International Conference on Very Large Databases (VLDB), 2004.

Semantic Gossiping (Chatty Web)

- Schemas might only partially overlap
- Mappings can be faulty
  - Heterogeneity of conceptualizations
  - Inexpressive mapping language
  - (Semi-) automatic mapping creation

Self-organization principles at the semantic mediation layer

- Detect inconsistent mappings
- Per-hop semantic forwarding
  - Syntactic criteria
  - Semantic criteria

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

GridVine

Who?
- *EPFL*

Overlay structure
- *DHT (P-Grid)*

Data model
- *RDF (annotations) RDFS (schemas) OWL (mappings)*

Queries
- *RDQL*

Query reformulation
- *Distributed*

Query evaluation
- *Distributed*

GridVine Architecture

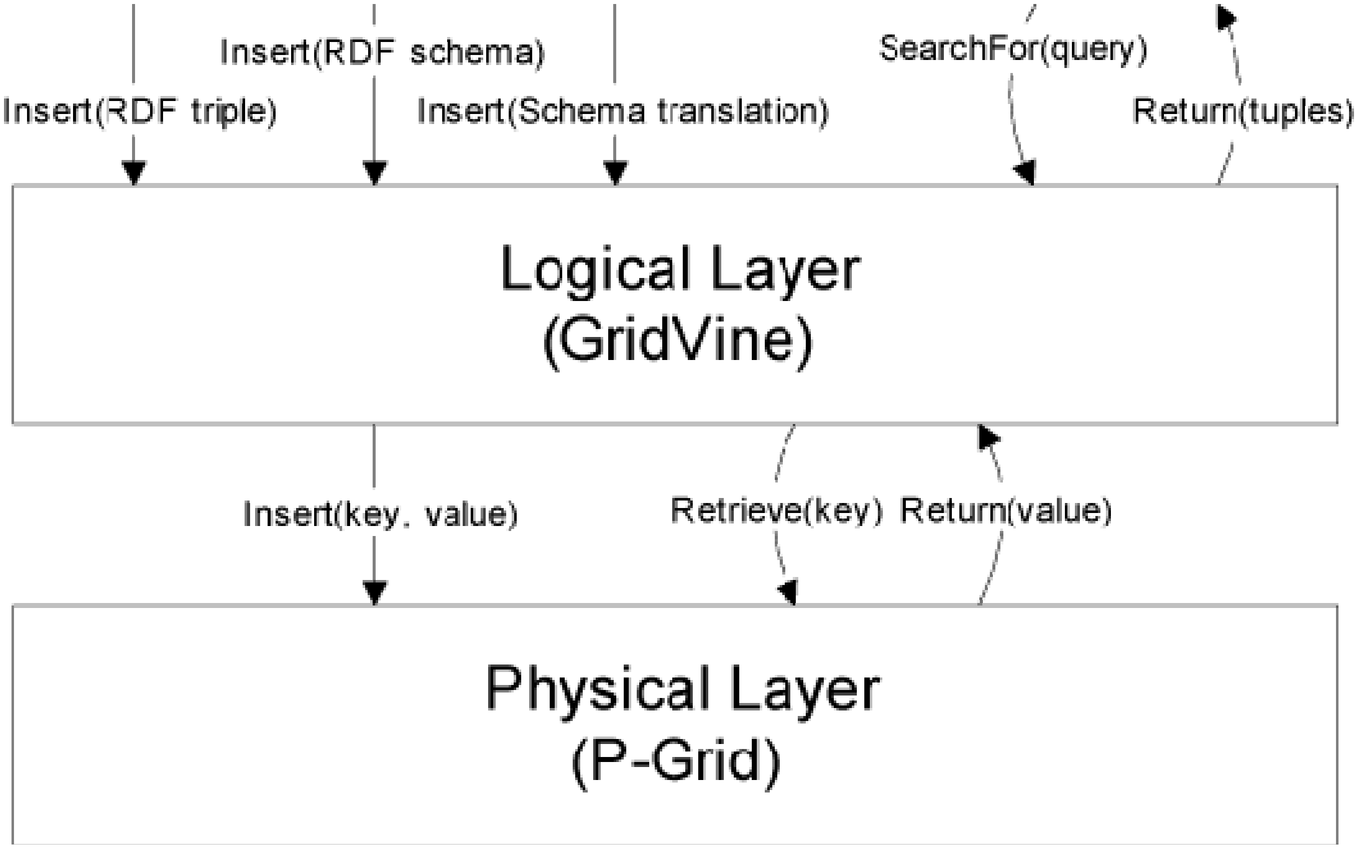

- Data / Schemas / Mappings are all indexed

⇒ *Decoupling*

Deriving Routing Indices (semantic layer)

- Automatically deriving quality measures from the mapping network to direct reformulation
  - Cycle / parallel paths / results analysis

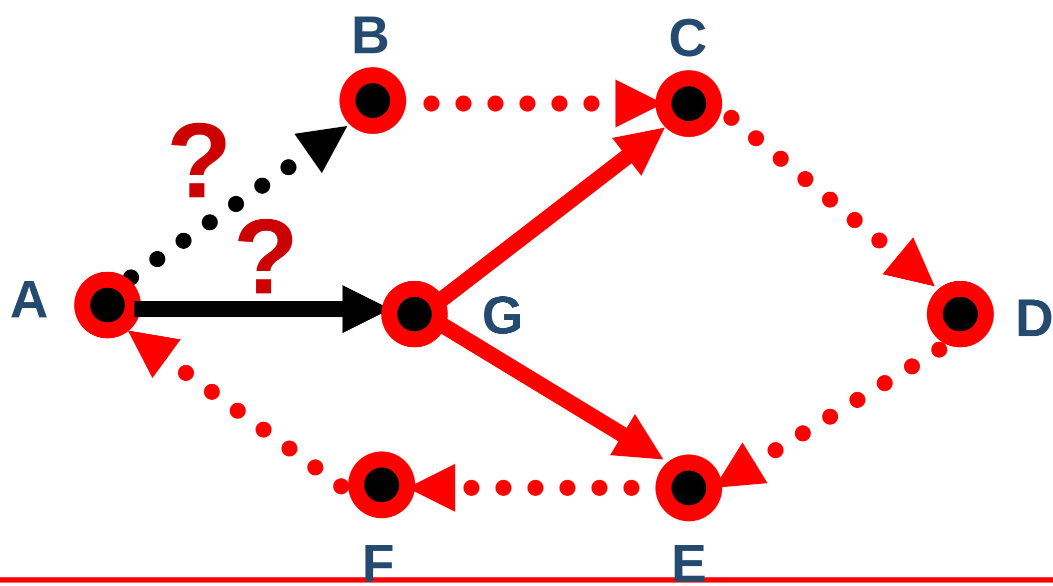

Example: Cycle Analysis

- What happened to an attribute A_i present in the original query?
  - $(T_{1\rightarrow\ldots\rightarrow n\rightarrow 1})$ (Creator) = (Creator) √
  - $(T_{1\rightarrow\ldots\rightarrow n\rightarrow 1})$ (Creator) = (Subject) X
  - $(T_{1\rightarrow\ldots\rightarrow n\rightarrow 1})$ (A_i) = ∅

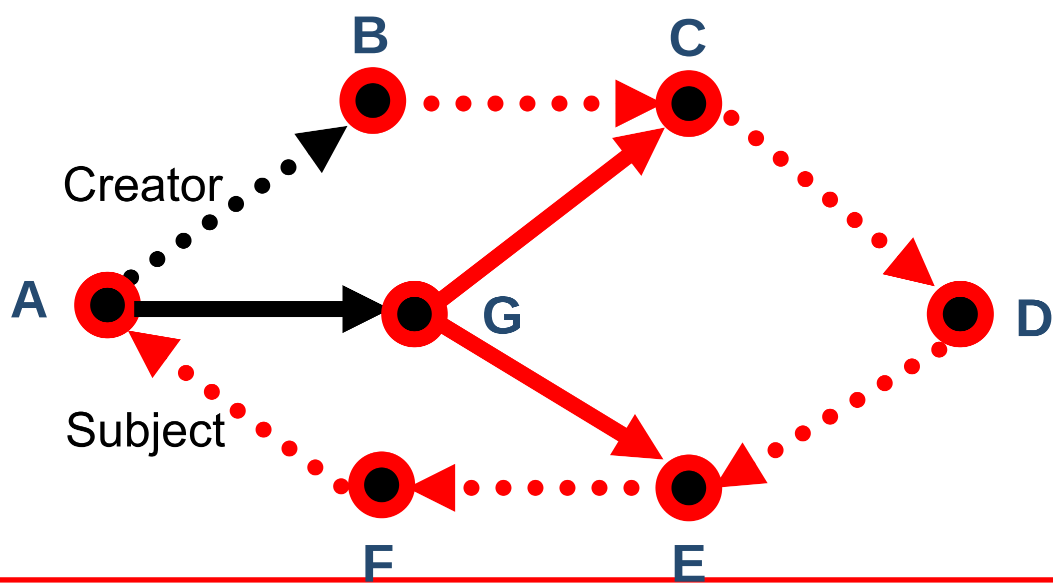

Example: Cycle Analysis

- In absence of additional knowledge:
  - "Foreign" links have probability of being wrong ε_{cyc}
  - Errors could be "accidentally" corrected with prob δ_{cyc}
    - Probability of receiving positive feedback (assuming A→B is correct) is $(1-\varepsilon_{cyc})^5 + (1-(1-\varepsilon_{cyc})^5)\,\delta_{cyc} = pro^+(5, \varepsilon_{cyc}, \delta_{cyc})$

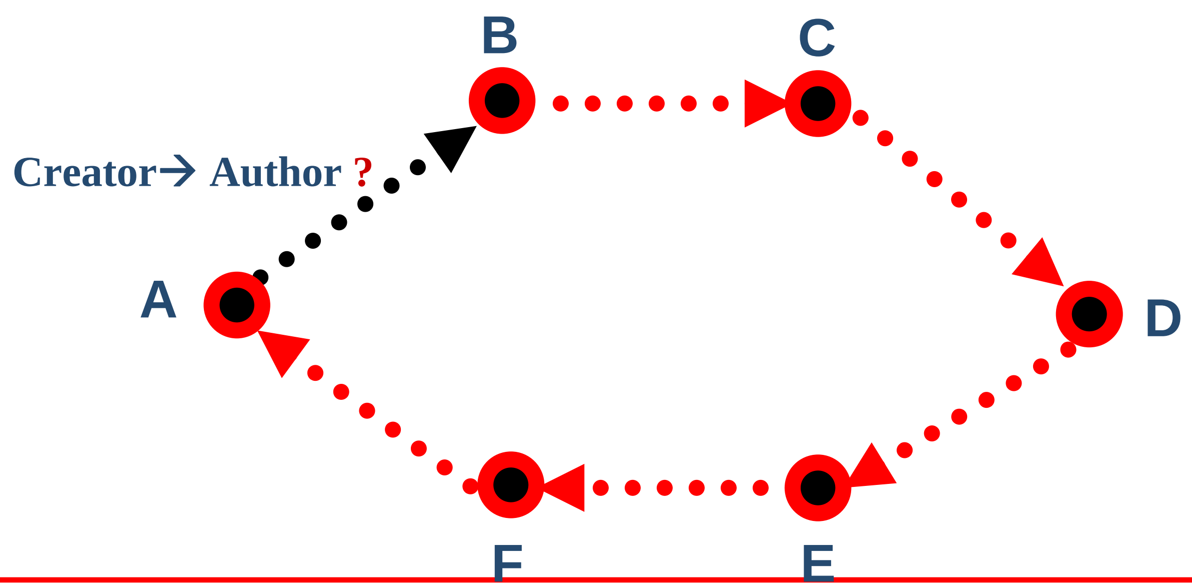

Example: Cycle Analysis

- Likelihood of receiving series positive and negative cycle feedback $c_1, \dots c_k$:

$l\,(c1,\dots, ck)$ =

$$(1-\varepsilon_s)\prod_{ci \in C+} pro^+(|c_i|, \varepsilon_{cyc},\delta_{cyc})\ )\prod_{ci \in C-} 1-pro^+(|c_i|, \varepsilon_{cyc},\delta_{cyc})$$

$$+\ \varepsilon_s\prod_{ci \in C+} pro^-(|c_i|, \varepsilon_{cyc},\delta_{cyc})\ )\prod_{ci \in C-} 1-pro^-(|c_i|, \varepsilon_{cyc},\delta_{cyc})$$

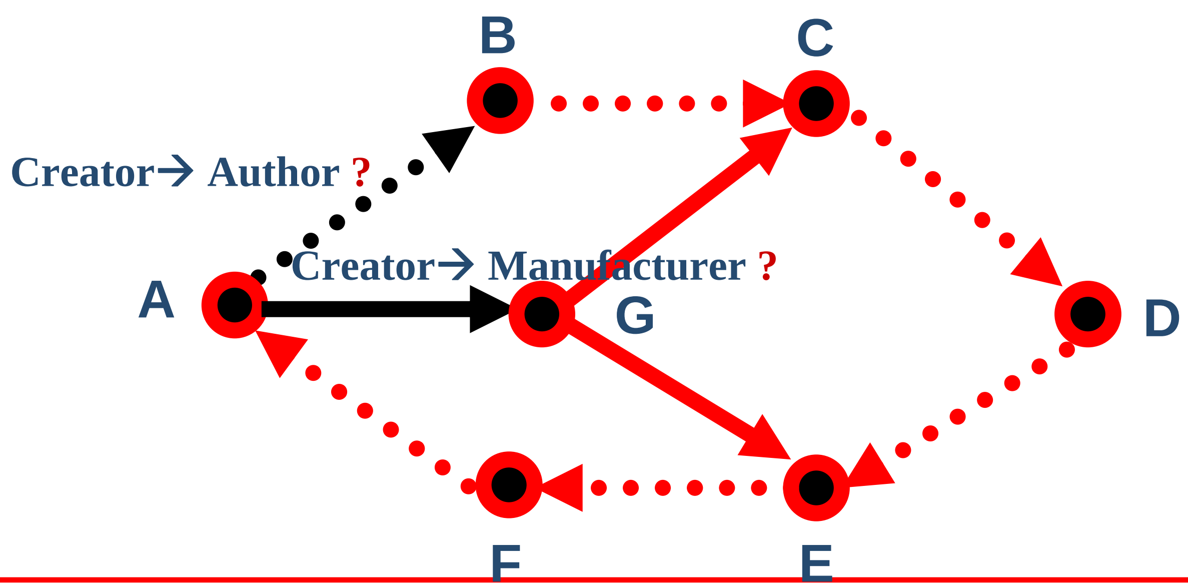

Which Link to Trust?

- Without other information on ε_{cyc} and δ_{cyc} , likelihood of our link being correct or not:

$$p^+ = \lim_{\varepsilon s \to 0} \int_{\delta cyc} \int_{\varepsilon cyc} l\,(c1,\ldots,ck)\, d\varepsilon_{cyc}\, d\delta_{cyc}$$

$$p^- = \lim_{\varepsilon s \to 1} \int_{\delta cyc} \int_{\varepsilon cyc} l\,(c1,\ldots,ck)\, d\varepsilon_{cyc}\, d\delta_{cyc}$$

$\Rightarrow \gamma = p^+ / (p^+ + p^-)$

ABCDEFA: √
AGEFA: X
AGCDEFA: X

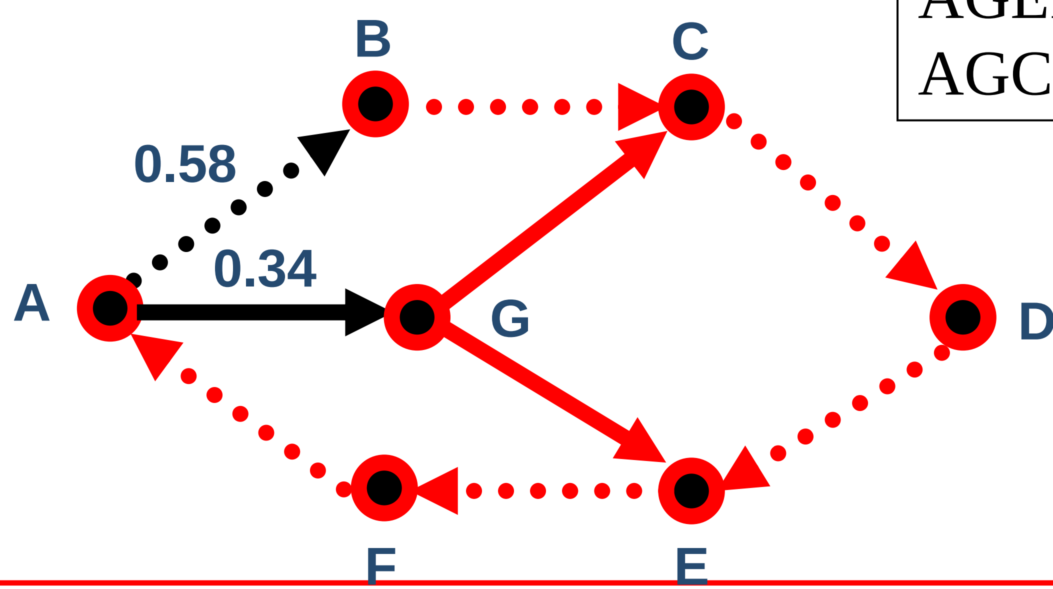

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

Reformulating query: Semantic Gossiping

- Selectively forward queries at the semantic mediation layer
  - Syntactic thresholds
    - Lost predicates
  - Semantic thresholds
    - Results analysis
    - Cycles analysis
- Drop/Repair faulty mappings
  - Self-organized semantic layer

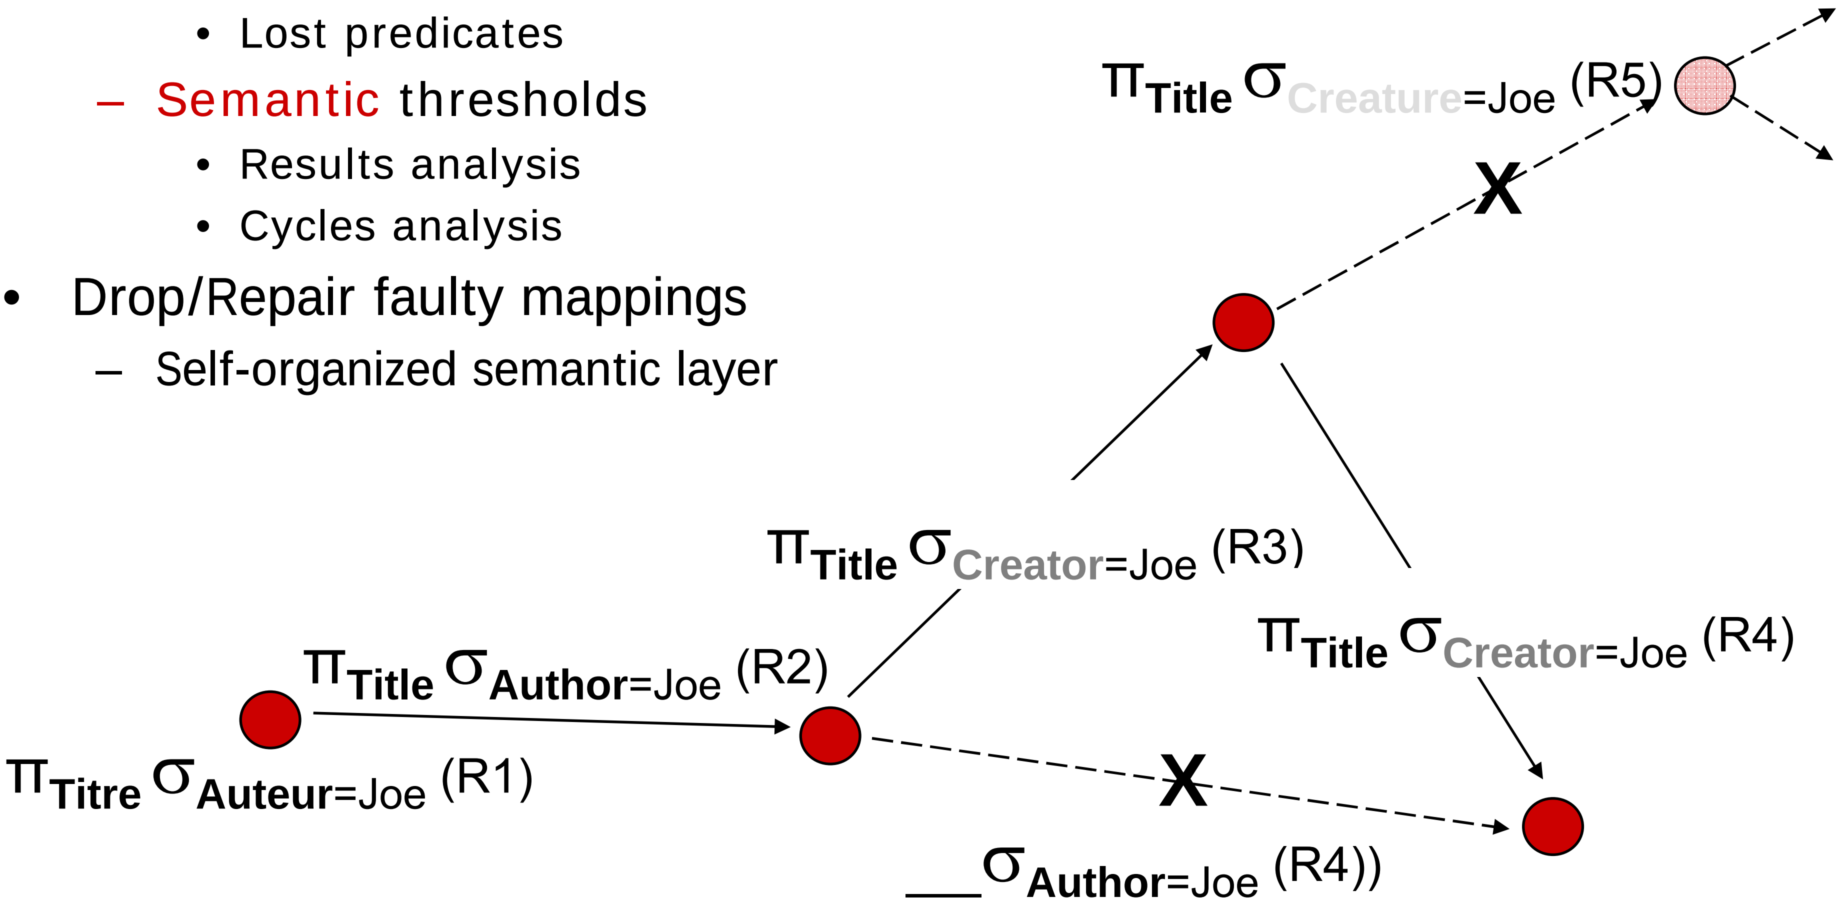

Decentralized Query Resolution: Overview

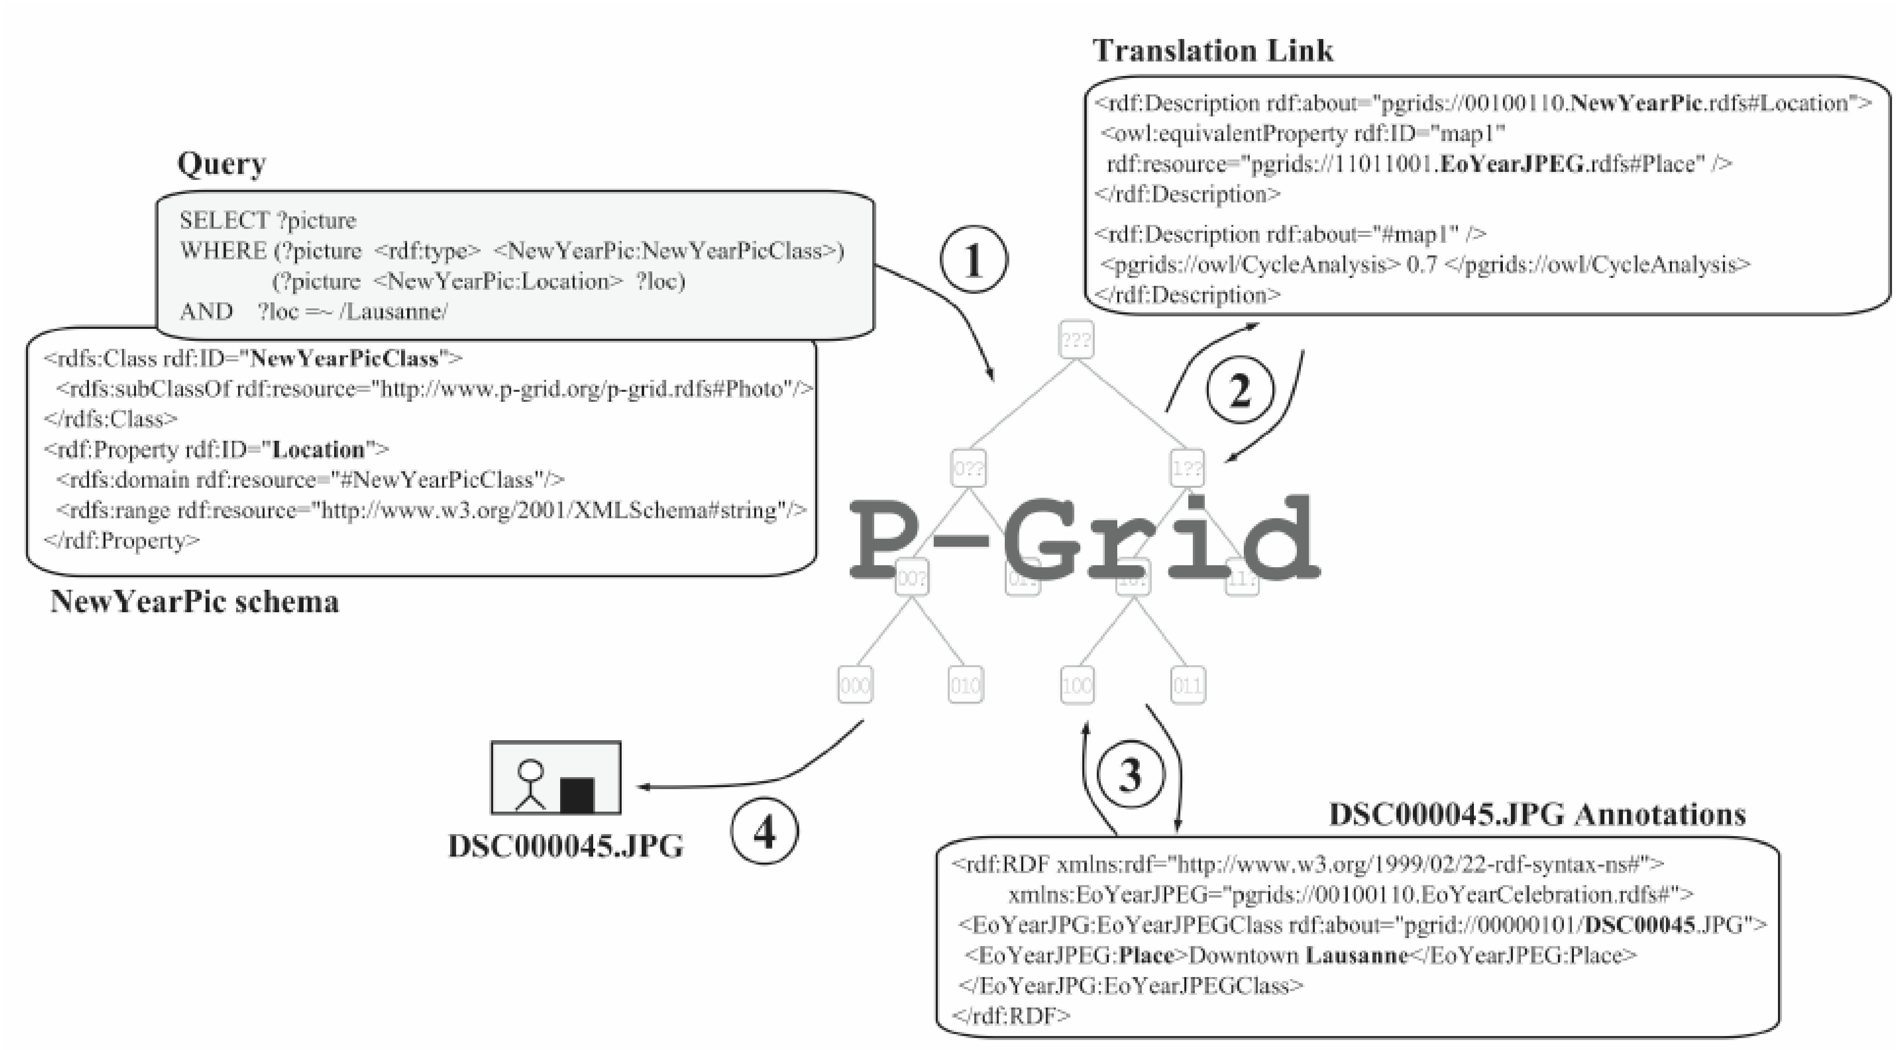

Want more? Belief Propagation in SONs

- Inferring global mapping quality values from a decentralized message-passing process

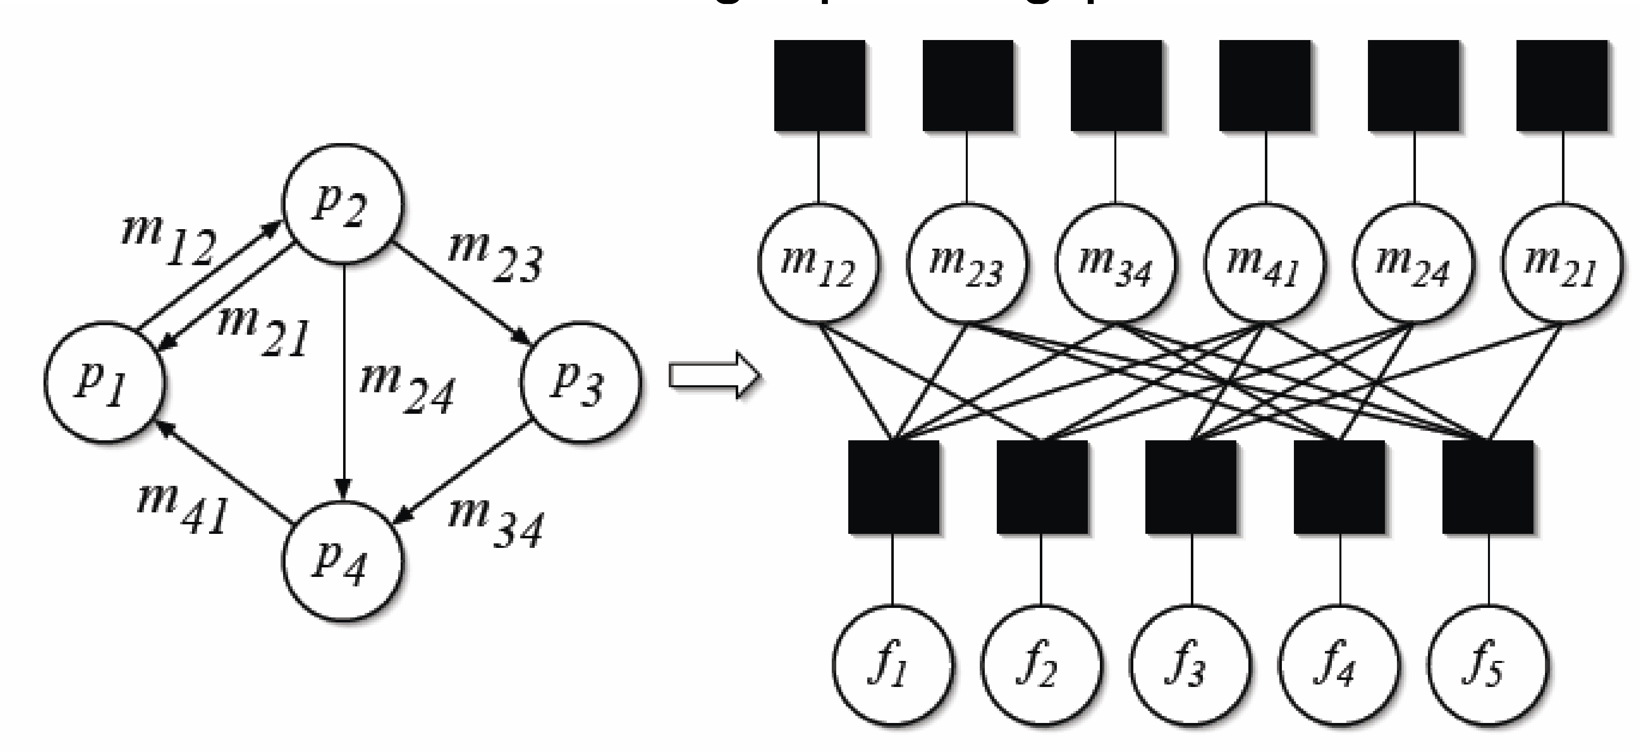

References

- K. Aberer, P. Cudre-Mauroux, and M. Hauswirth. A Framework for Semantic Gossiping. SIGOMD RECORD, 31(4), 2002.
- K. Aberer, P. Cudre-Mauroux, A. Datta, Z. Despotovic, M. Hauswirth, M. Punceva, and R. Schmidt. P-grid: A self-organizing structured p2p system. SIGMOD Record, 32(3), 2003.
- K. Aberer, P. Cudre-Mauroux, and M. Hauswirth. The Chatty Web: Emergent Semantics Through Gossiping. In International World Wide Web Conference (WWW), 2003.
- K. Aberer, P. Cudre-Mauroux, and M. Hauswirth. Start making sense: The Chatty Web approach for global semantic agreements. Journal of Web Semantics, 1(1), 2003.
- K. Aberer, P. Cudre-Mauroux, M. Hauswirth, and T. van Pelt. GridVine: Building Internet-Scale Semantic Overlay Networks. In International Semantic Web Conference (ISWC), 2004.
- P. Cudre-Mauroux, K. Aberer and A. Feher. Probabilistic Message Passing in Peer Data Management Systems. In International Conference on Data Engineering (ICDE), 2006.

IV. Current Research Directions

Emergent Semantics

- Semantic Overlay Networks can be viewed as highly dynamic systems (churn, autonomy)
- Semantic agreements can be understood as emergent phenomena in complex systems

⇒ Principles
  - mutual agreements for meaningful exchanges
  - agreements are dynamic, approximate and self-referential
  - global interoperability results from the aggregation of local agreements by self-organization

K.Aberer, T. Catarci, P. Cudré-Mauroux, T. Dillon, S. Grimm, M. Hacid, A. Illarramendi, M. Jarrar, V. Kashyap, M. Mecella, E. Mena, E. J. Neuhold, A. M. Ouksel, T. Risse, M. Scannapieco, F. Saltor, L. de Santis, S. Spaccapietra, S. Staab, R. Studer and O. De Troyer: Emergent Semantics Systems. International Conference on Semantics of a Networked World (ICSNW04).

SON Graph Analysis

- Networks resulting from self-organization processes
  - powerlaw graphs, small world graphs
- Structure important for algorithm design
  - distribution, connectivity, redundancy

$\Rightarrow$ Analysis and Modeling of SON from a graph-theoretic perspective

P. Cudré-Mauroux, K. Aberer: "A Necessary Condition for Semantic Interoperability in the Large", CoopIS/DOA/ODBASE (2) 2004: 859-872.

Information Retrieval and SONs

- Combination of structural, link-based and content-based search
- Precision of query answers drops with semantic mediation

⇒ IR techniques to optimize precision/recall in SONs

- Distributed ranking algorithms
- Content-based search with DHTs
- Peer selection using content synopsis

M. Bender, S. Michel, P. Triantafillou, G. Weikum and C. Zimmer: Improving Collection Selection with Overlap Awareness in P2P Search Engines. SIGIR2005.

J. Wu, K. Aberer: "Using a Layered Markov Model for Distributed Web Ranking Computation", ICDCS 2005.

Corpus-Based Information Management

- Very large scale, dynamic environments require on-the-fly data integration
- Automated schema alignment techniques may perform poorly
  - Lack of evidence

⇒ Using a preexisting corpus of schema and mapping to guide the process
  - Mapping reuse
  - Statistics offer clues about semantics of structures

J. Madhavan, P. A. Bernstein, A.i Doan and A. Y. Halevy: Corpus-based Schema Matching. ICDE 2005

Declarative Overlay Networks

- Overlay networks are very hard to design, build, deploy and update

$\Rightarrow$ Using declarative language not only to query, but also to express overlays

  - Logical description of overlay networks
  - Executed on a dataflow architecture to construct routing data structures and perform resource discovery

B. Thau Loo, T. Condie, J. M. Hellerstein, P. Maniatis, T. Roscoe, I. Stoica: Implementing Declarative Overlays. ACM Symposium on Operating Systems Principles (SOSP), 2005

Internet-Scale Services

- Many infrastructures tackle today data management at Internet scale
  - Semantic Web
  - Web Services
  - Grid Computing
  - Dissemination Services

⇒ SONs as a generic infrastructure for very large-scale data processing

Further References

- Length limits constrained the number of approaches we could discuss...

⇒ http://lsirwww.epfl.ch/SON

For a more complete list of research projects in the area of Semantic Overlay Networks