Distributed Queries and Query Optimization in Schema-Based P2P-Systems

Ingo Brunkhorst, Hadhami Dhraief, Alfons Kemper, Wolfgang Nejdl and Christian Wiesner

Presented By: Mohamed Ali Soliman
m2ali@cs.uwaterloo.ca

P2P Query Processing, the Big Picture

- Earlier P2P Systems
  - Flooding Queries
  - Random Walks
- Common Data Model P2P Systems
  - Mediation Layers
  - Mutant Queries
  - Query Routing Indices
- Schema Heterogeneous P2P Systems
  - Query Reformulation
  - Data Translation Rules
- P2P Systems with More Complex Query Types
  - Range Queries
  - Multi-Attribute Queries
  - Join Queries
  - Aggregation Queries

RDF based P2P Networks [SQ03], [SP03]

- RDF is used to annotate resources providing the means to exchange and comprehend data
- The annotations about resources can be based on various schemas
- Different servers can store metadata about the same resource in different standards e.g. DC,LOM,DCQ
  - Example: (r, dc:language, “en”)
    (r, lom:context, “undergrad”)
    (r, dc:subject, “software engineering”)
- Richer descriptions of resources enable more extensive queries

Outline

- Overview
- Routing and Mediation Schema-based P2P-Systems
- Query Processing
- Query Plans
- Query Optimization
- Conclusions and Comments

Overview

- P2P-Systems are highly dynamic in nature
  - No global schema can be assumed
  - No static network topology
  - Static query plans are not suitable
- Most efforts have been in infrastructures
- Extending queries functionalities begin to receive more attention
- Database systems have also evolved from being centralized to distributed
- Marriage between database and P2P technologies is now taking place

Overview

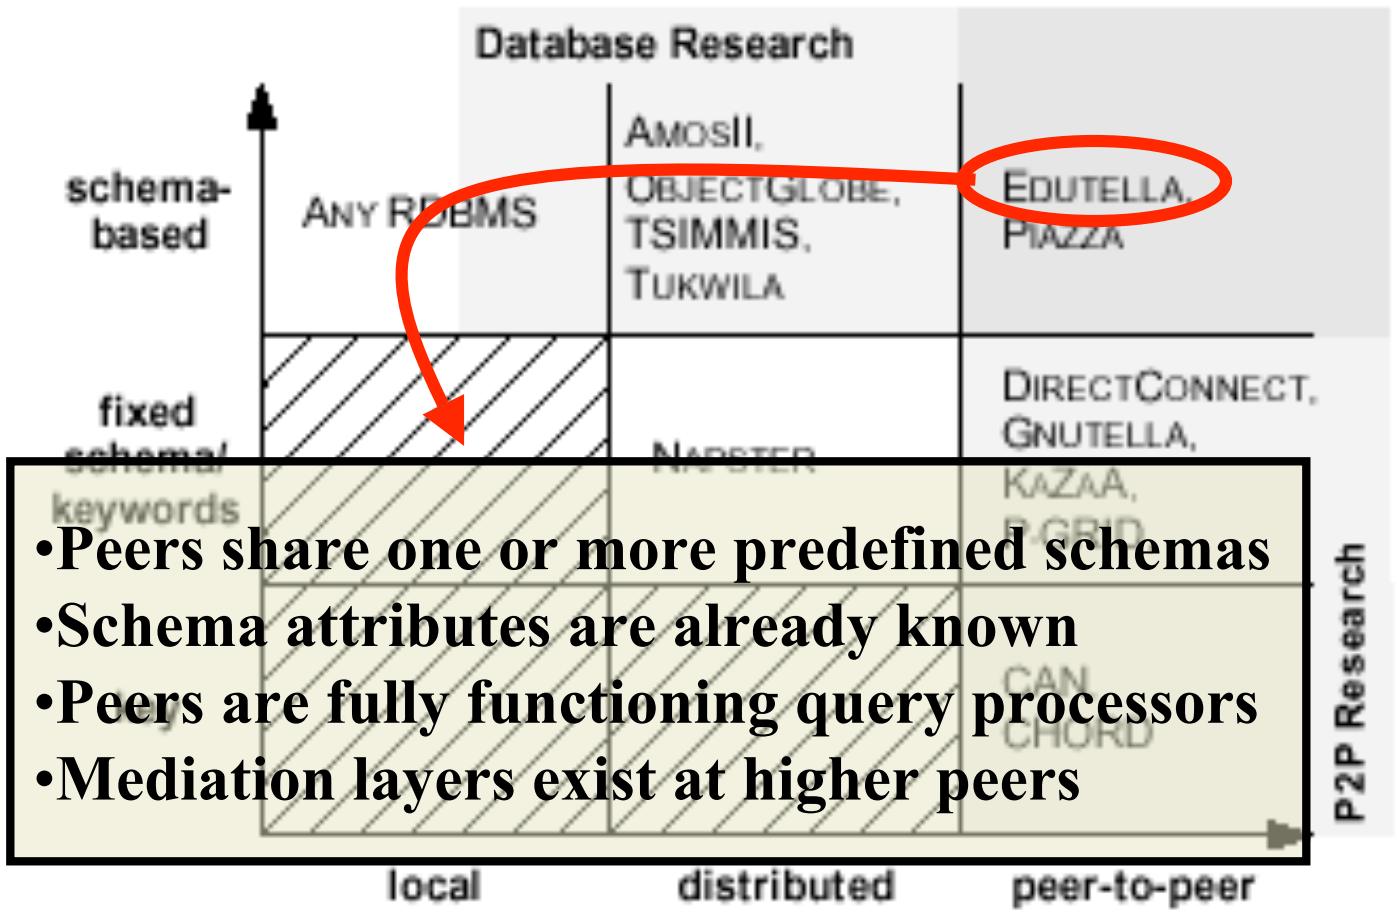

- Peers share one or more predefined schemas
- Schema attributes are already known
- Peers are fully functioning query processors
- Mediation layers exist at higher peers

Overview

Edutella Project [Edu[1],Edu02]

- A multi-staged project to scope, specify, architect and implement an RDF-based metadata infrastructure for JXTA framework
- Peers offer a set of services
  - *Query service* based on predefined schema attributes
  - *Data Replication* to achieve load balance
  - *Mapping service* to convert queries on one schema to another
- Wrappers are used to exchange queries and results between peers and the network (Edutella Query Exchange Language)
- Applications
  - A P2P network for the exchange of educational resources between universities

[1] **http://edutella.jxta.org**

Outline

- ☑ Overview
- Routing and Mediation Schema-based P2P-Systems
- Query Processing
- Query Plans
- Query Optimization
- Conclusions and Comments

Routing in Schema-Based P2P-Systems

- Super peer architecture
- Two-phase routing architecture:
  - Route queries first in SP backbone
  - Then, distribute them to the peers
- Super-peers are arranged in the **HyperCuP** topology
- Super-peers join the HyperCuP by asking any of the already integrated super-peers
- No central maintenance is necessary for changing the HyperCuP structure

Routing in Schema-Based P2P-Systems

HyperCube Topology [Hyp02]

- A graph topology which allows for efficient search
- Nodes are organized as b nodes in each dimension
  - For b=2, it turns out to be a hypercube
- A node could cover more than one position in the cube
- Total number of nodes $N=b^{dimesnions}$, network diameter is $\log_b N$
- N-1 messages are required to span all network
- Edges are labeled to avoid sending message to the node that produced that message

Routing in Schema-Based P2P-Systems

HyperCube Construction [Hyp02]

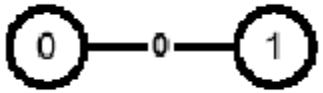

P1 contacts P0 to join

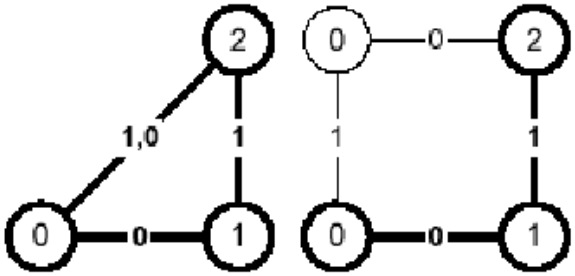

P0 covers missing dim for P2

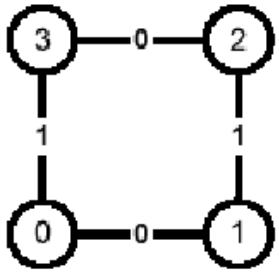

P3 contacts P0 to join

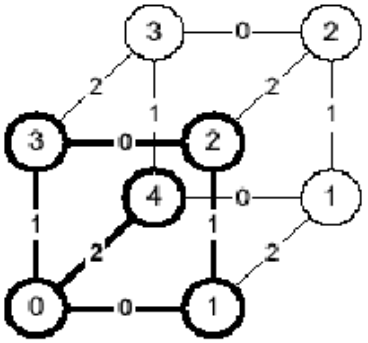

P4 contacts P0 to join

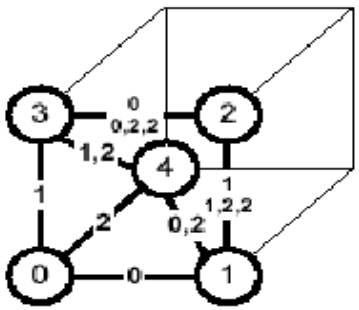

P1, P3 act as temp neighbors for P4

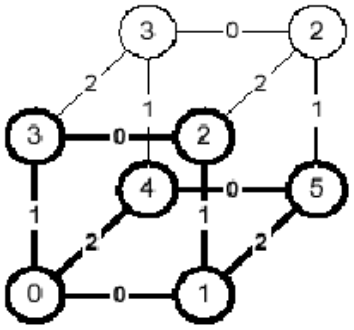

P5 contacts P1 to join

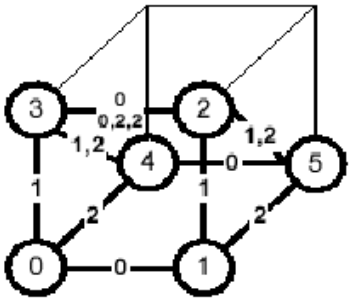

P5 takes 2 level neighbor from P1

Routing in Schema-Based P2P-Systems

- Routing indices contain metadata information about peers at different levels
- Super Peer/Peer Routing Index
  - Schema index
  - Properties names index
  - Property value range index
  - Property value index
- Different index granularities make queries more flexible
- Super Peer/Super Peer index contain index information about neighboring super peers.

Routing in Schema-Based P2P-Systems

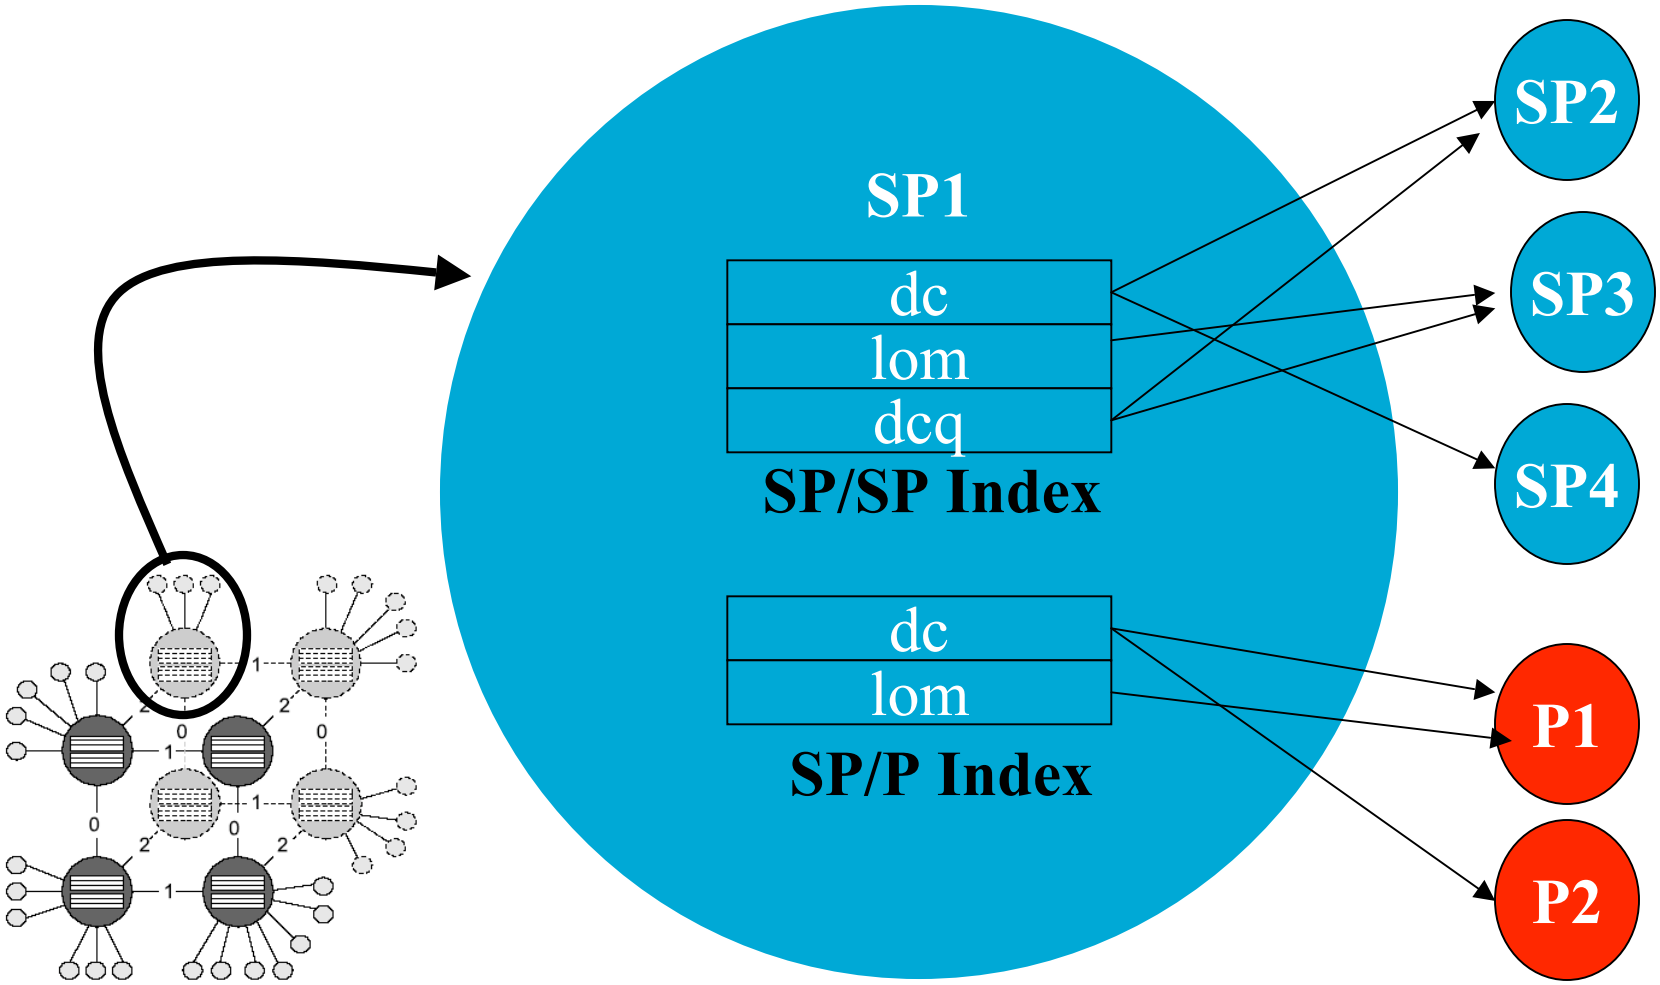

Routing indices at schema level for SP1

Routing in Schema-Based P2P-Systems

Granularity	Index of SP_2		
Schema	dc		SP_1, SP_3, SP_4
	lom		SP_1, SP_3, SP_4
	dcq		SP_3
Property	dc:subject		SP_1, SP_3, SP_4
	lom:type		SP_1, SP_3, SP_4
	dc:format		SP_3, SP_4
Property Value Range	dc:subject	ccs:dbms	SP_1, SP_2, SP_3
Property Value	lom:type	"exercise"	SP_3
	dc:language	"dc"	SP_3, SP_4

Routing indices at different levels for SP2

Routing in Schema-Based P2P-Systems

- Updating indices is made by constructing a spanning tree among super peers
- Each super peer sends update messages to the neighboring super peers
- Whenever a super peer index is not updated by the incoming message, forwarding stops

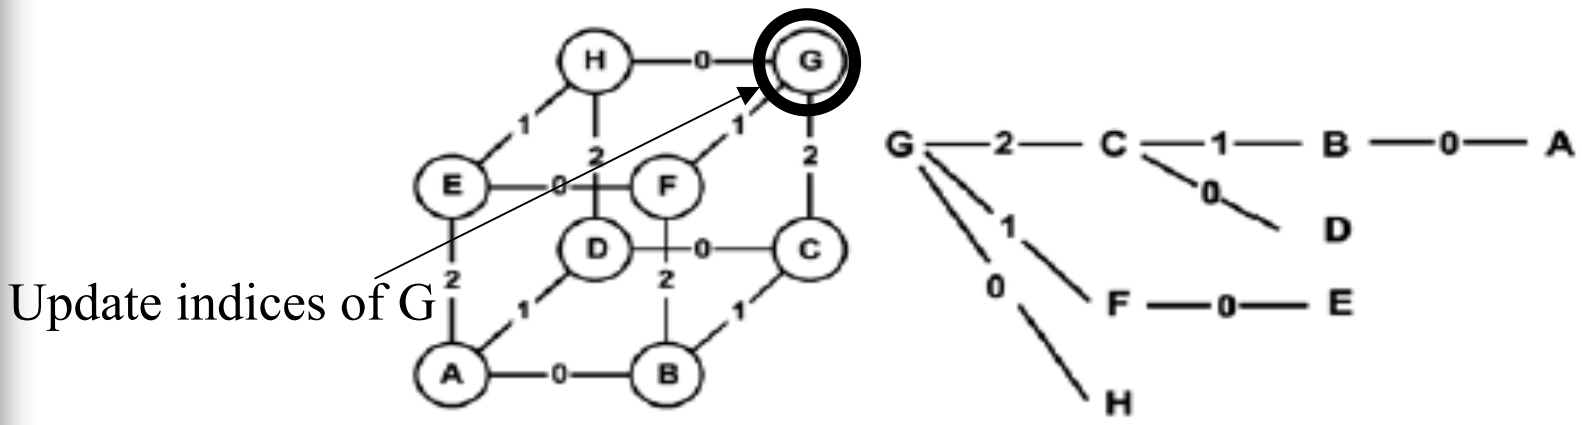

Mediation between Different Schemas

- Each peer registers with a super-peer using *advertisements* which contain the metadata schema used at the peer
- Rules are used to describe query capabilities of different peer. For example:
  - **Super peer administrator defines query schema:**
    lectures (id, lang, sub, context)
  - **Correspondences to heterogeneous peers schemas at peers P1, P2*:***
    lectures:id = dc:title
    lectures:lang = dc:lang
    lectures:sub = dc:subject

    lectures:id = lom:general.identifier
    lectures:lang= lom:general.language
    lectures:context = lom:educational.context
  - **Views can be created on peers schemas:**
    ViewDC (lectures:id, lectures:lang, lectures:sub)
    ←
    DC (dc:title, dc:lang, dc:subject)

    ViewLOM (lectures:id, lectures:lang, lectures:context)
    ←
    LOM (lom:general.identifier, lom:geneal.language, lom:educational.context)

Mediation between Different Schemas

- **Then we can construct associations between views and query schema:**

  lectures(id, lang, sub)

  ←

  ViewDC (lectures:identifier, lectures:language, lectures:subject)

  lectures (id, lang, context)

  ←

  ViewLOM (lectures:id, lectures:lang, lectures:context)

- **Super peer stores correspondences between views and peers**

  P1←ViewDC

  P2←ViewLOM

- When a super peer receives a query *lecture (id, lang, sub, context)* the super peer identifies P1 and P2 as a combination of relevant correspondences that are semantically included in the user query

Outline

- ☑ Overview
- ☑ Routing and Mediation Schema-based P2P-Systems
- Query Processing
- Query Plans
- Query Optimization
- Conclusions and Comments

Query Processing

- Flooding based P2P systems
  - Bandwidth drain
  - Uncovered network areas
  - Results are returned to client to start obtaining data
  - Results filtering and user-defined code run at client
- Schema based P2P systems
  - Distributed index is utilized
  - Selective query flooding
  - Query processing and user-defined code still run at client

Query Processing

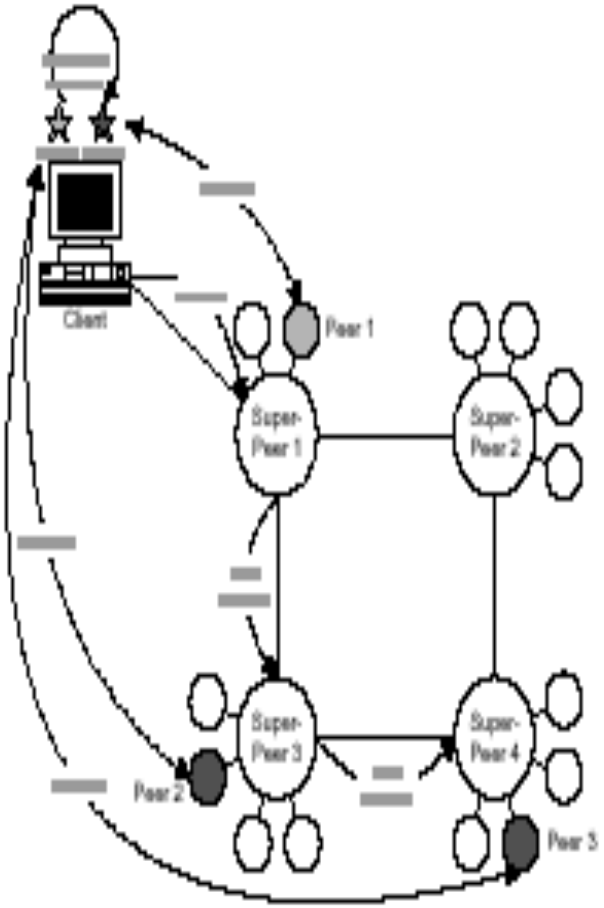

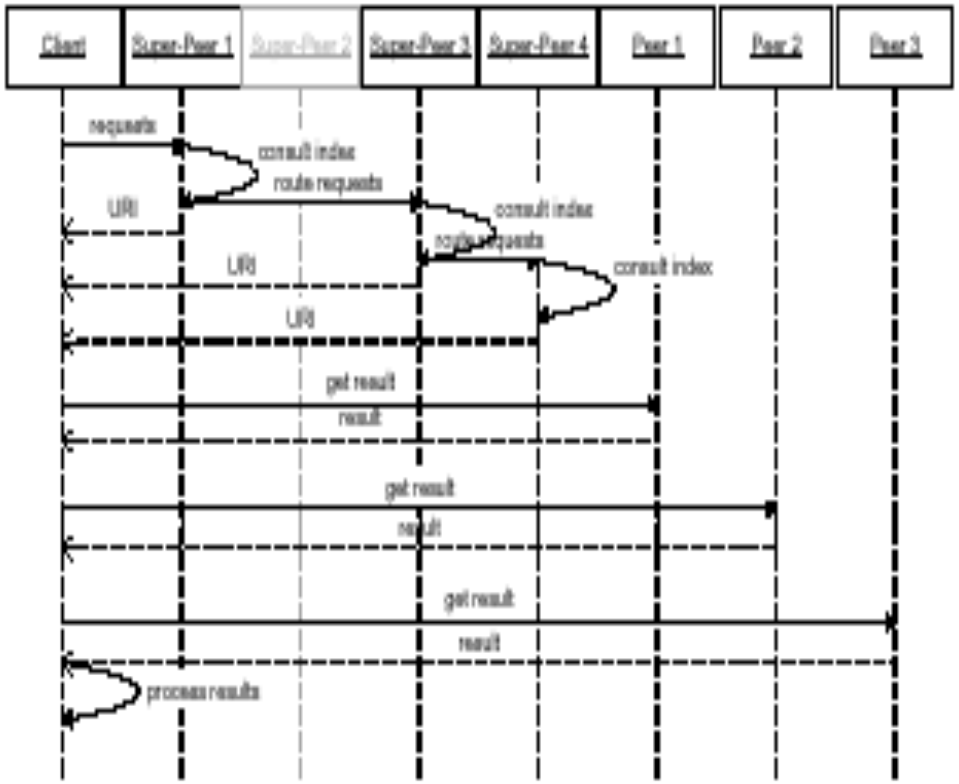

Traditional Query Processing in Schema-based P2P-Systems

Query Processing

- Query Evaluation Plans (QEPs)
  - Super Peers provide functionality for index management, query optimization and query processing
  - Peers provide query processing
  - Super peers distribute incoming queries to other peers and super peers guided by index
  - Query plans are dynamically constructed according to schema information at super peers
  - Client submits a query containing user-defined filters to super peer which decides where filters are executed

Query Processing

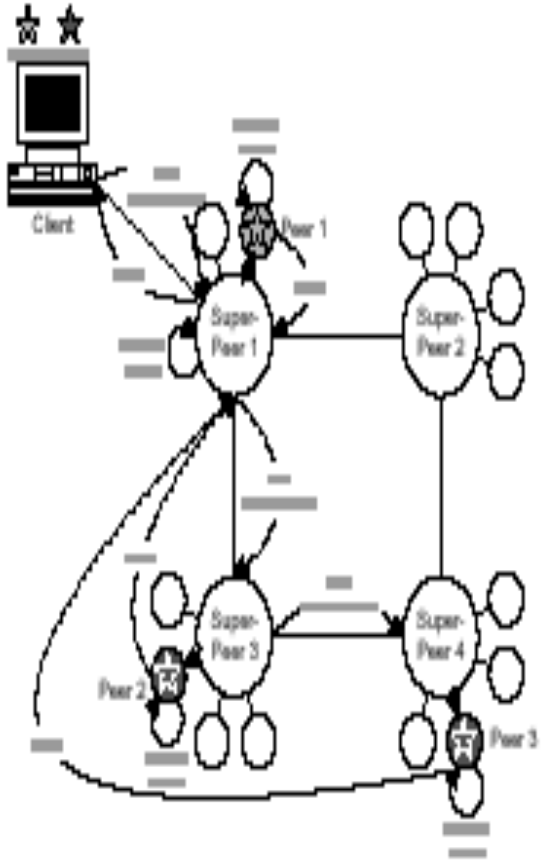

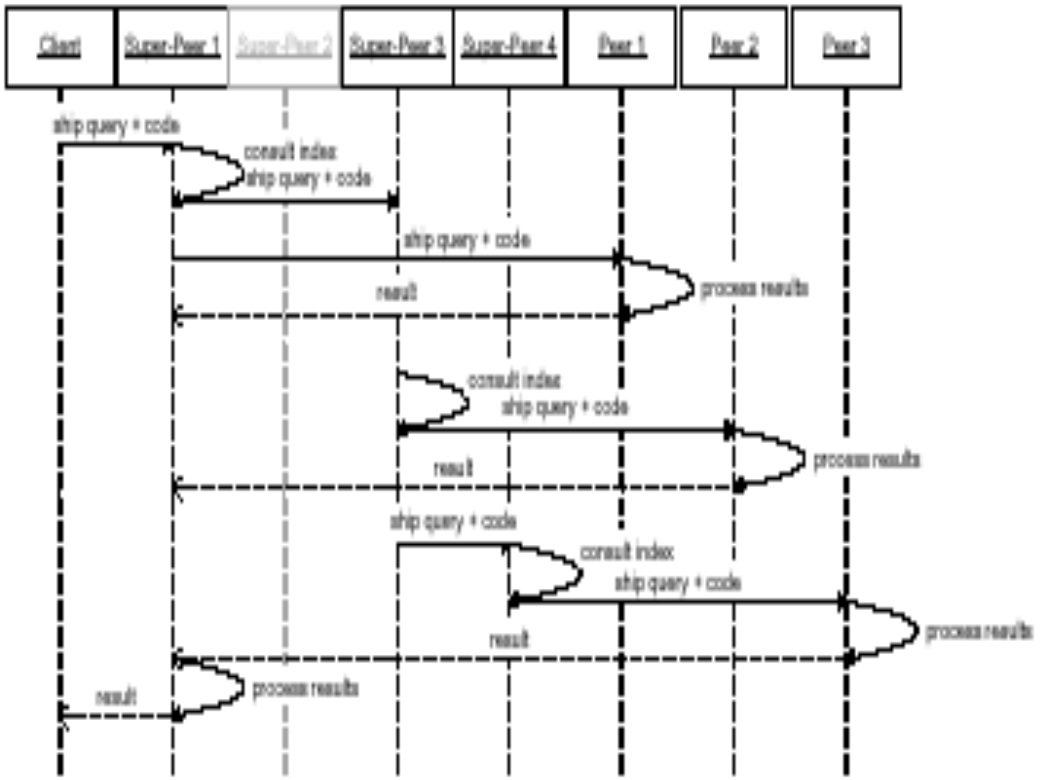

Pushing code to data sources

Outline

- ☑ Overview
- ☑ Routing and Mediation Schema-based P2P-Systems
- ☑ Query Processing
- ■ Query Plans
- ■ Query Optimization
- ■ Conclusions and Comments

Query Plans

- Query plan generation
  - Parse incoming query
  - Bind Resources based on index information
  - Generate subqueries based on bindings
  - Instantiate local plan
  - Distribute subqueries to neighboring (super)peers
- Cost optimization is based on network topology or query history
- Each query operator is annotated with the host where it is executed

Query Plans

For lom resources, return all lectures in Dutch with more than 2 occurrences of "transaction processing" and all exercises in Dutch or English ; where the author of the lecture and exercise is the same

```
select r1.data, r2.data, r3.data
from Resources r1, Resources r2, Resources r3
where r1.lom_type = "lecture" and r1.language =
"de"
and occur(r1.data, "transaction processing") > 2 and
r2.lom_type = "exercise" and r2.language = "de" and
r3.lom_type = "exercise" and r3.language = "en" and
r1.author = r2.author and r1.author = r3.author
```

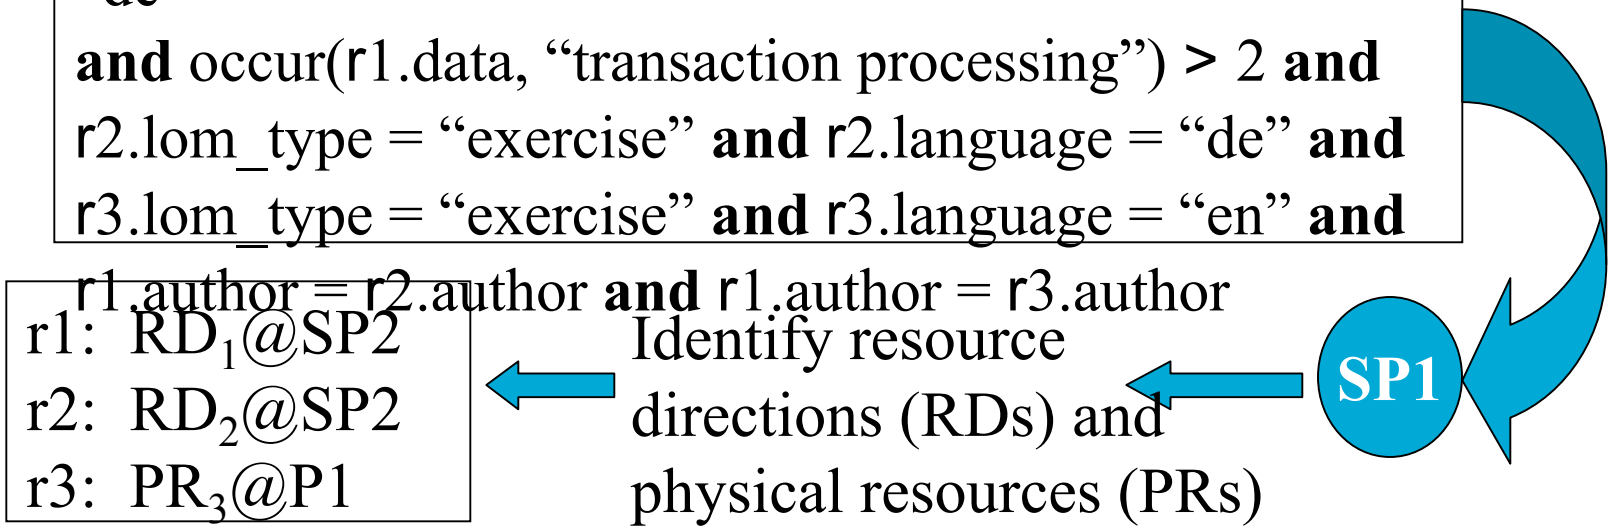

Query Plans

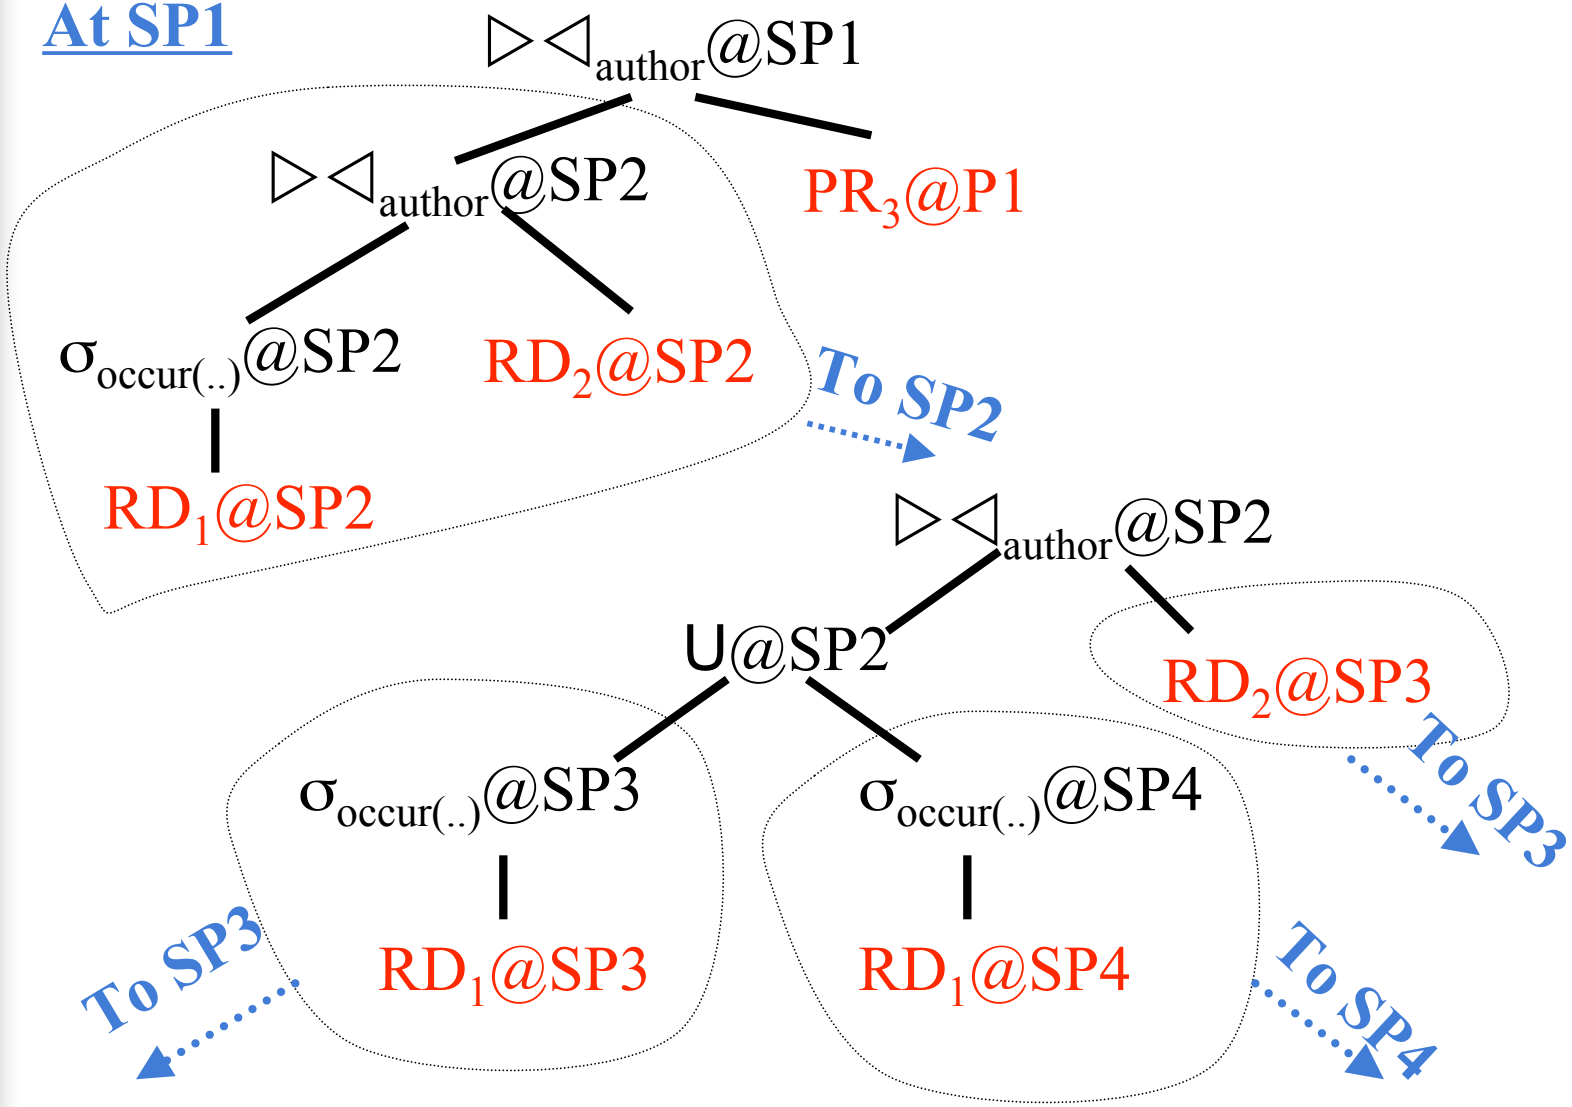
At SP1
⋈author@SP1
⋈author@SP2
PR3@P1
σoccur(..)@SP2
RD2@SP2
To SP2
RD1@SP2
⋈author@SP2
U@SP2
RD2@SP3
To SP3
σoccur(..)@SP3
σoccur(..)@SP4
To SP3
RD1@SP3
RD1@SP4
To SP4

Outline

- ☑ Overview
- ☑ Routing and Mediation Schema-based P2P-Systems
- ☑ Query Processing
- ☑ Query Plans
- ■ Query Optimization
- ■ Conclusions and Comments

Query Optimization

- Naïve approach is to union all physical resources before further operations
  - Incurs extensive data transmission
- Naïve strategy is acceptable for large data distribution
  - One host collect data from other hosts
  - Collecting host may change during plan generation

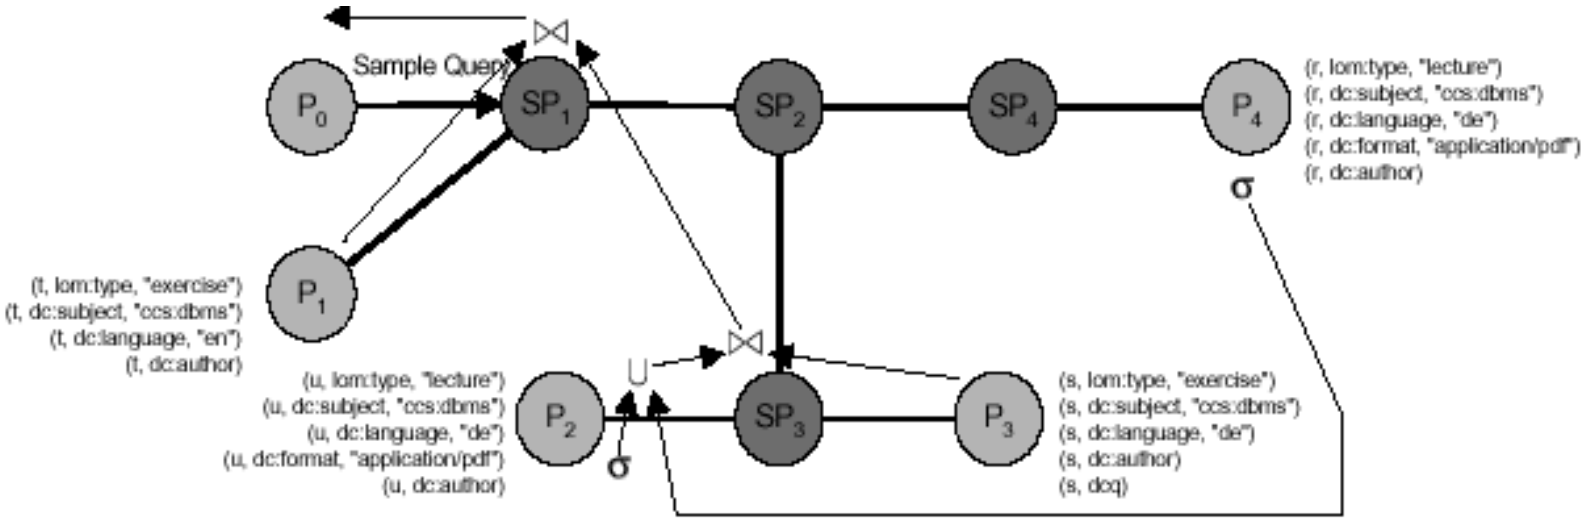

Outline

- ☑ Overview
- ☑ Routing and Mediation Schema-based P2P-Systems
- ☑ Query Processing
- ☑ Query Plans
- ☑ Query Optimization
- ■ Conclusions and Comments

Conclusions

- Distributed database query processing techniques need to adapt to P2P environments
- Distributing queries to promising peers only can be very beneficial for large P2P networks
- Pushing query operators next to data sources makes use of distributed peers processing power
- Remote results combination and filtering saves bandwidth by reducing data shipped through network

Comments

- Performance study was needed to evaluate the suggested techniques regarding:
  - Network latency
  - Bandwidth savings
  - Remote filtering
  - Response time
- Assumption of peers willingness to provide processing power for external queries
- No consideration was made to heterogeneous peers capabilities
- How can peers be evenly distributed among super peers
- Clustering content-related (super) peers

References

- **[SQ03]** Distributed Queries and Query Optimization in Schema-Based P2P-Systems. *I. Brunkhorst et al.* In Proc. Int. Workshop On Databases, Information Systems and Peer-to-Peer Computing, 2003.
- **[Edu]** The Edutella Project: http://edutella.jxta.org
- **[Edu02]** EDUTELLA: A P2P Networking Infrastructure Based on RDF, *Wolfgang Nejdl et al.*,In WWW 2002.
- **[Hyp02]** HyperCuP—Hypercubes, Ontologies and Efficient Search on P2P Networks. *M. Schlosser et al.*, In Intl. Workshop on Agents and P2P Computing, 2002.
- **[SP03]** Super-Peer-Based Routing and Clustering Strategies for RDFBased PeerToPeerNetworks, *Wolfgang Nejdl et al.*, In WWW, 2003.