A Self-organizing XML P2P Database System

Giovanni Conforti Giorgio Ghelli Paolo Manghi
Carlo Sartiani
Dipartimento di Informatica - Università di Pisa

Topics

- XPeer: a data management system
  - XML data
  - p2p architecture
  - self-organizing
  - zero-administration

Presentation Outline

- system architecture
- query processing & query algebra
- conclusions

System Architecture

Introduction

- an open-ended and dynamic network
- a p2p hybrid architecture
  - peer nodes
  - super-peer nodes
- self-organizing and tree-shaped overlay network
  - peers: leaves
  - super-peers: internal nodes
  - adaptation to changes in the network topology or in the workload

Main Issues

- keeping track of network topology changes
- keeping track of changes in local data
- routing queries over the network

Peer Nodes

- peers are autonomous and may carry heterogeneous data
  - free local updates
- peer content is described by a tree-shaped schema
  - tree-guide
- peer nodes also execute query plans returned by the super-peer layer
- peer nodes may replicate the content of other nodes and cache the result of previous queries

Tree-Guide

```
<market>
   <buildings>
      <building>
         <desc> Marvelous luxury house in the Hamptons </desc>
         <location> Hamptons </location>
         <price> 1600000 </price>
      </building>
      <building>
         <desc> Very nice flat in the Upper East Side </desc>
         <location> Upper East Side, Manhattan </location>
         <price> 1350000 </price>
         <type> comdo </type>
      </building>
      <building>
         <desc> Elegant luxury house in the countryside </desc>
         <location> Greensboro </location>
         <price> 1700000 </price>
      </building>
   </buildings>
</market>
```

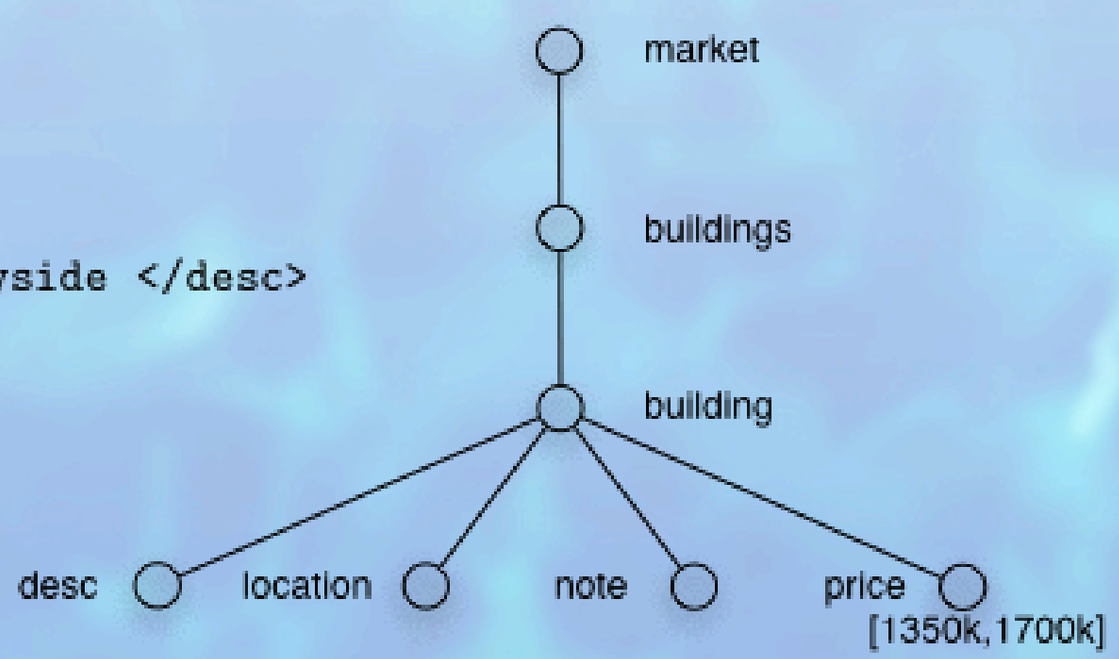

Peer Clusters

- peer nodes are organized into clusters
- peer clusters are formed on a schema-similarity basis
  - tree similarity
- each cluster is managed by a super-peer node

Super-Peer Nodes

- super-peer nodes are still peer nodes
- super-peer tasks:
  - query compilation
  - replica management
  - cache management

More on Super-Peers

- super-peers maintain schema information
  - the list of the schemas of their peers
  - the union of these schemas
- this information is used for compiling queries
  - twig matching

Super-Peer Groups

- super-peer nodes are organized into groups
- super-peers having the same father in the hierarchy form a group
- groups are intended to
  - increase the robustness of the system wrt node or network failures
  - guide the query compilation process

Overlay Network Example

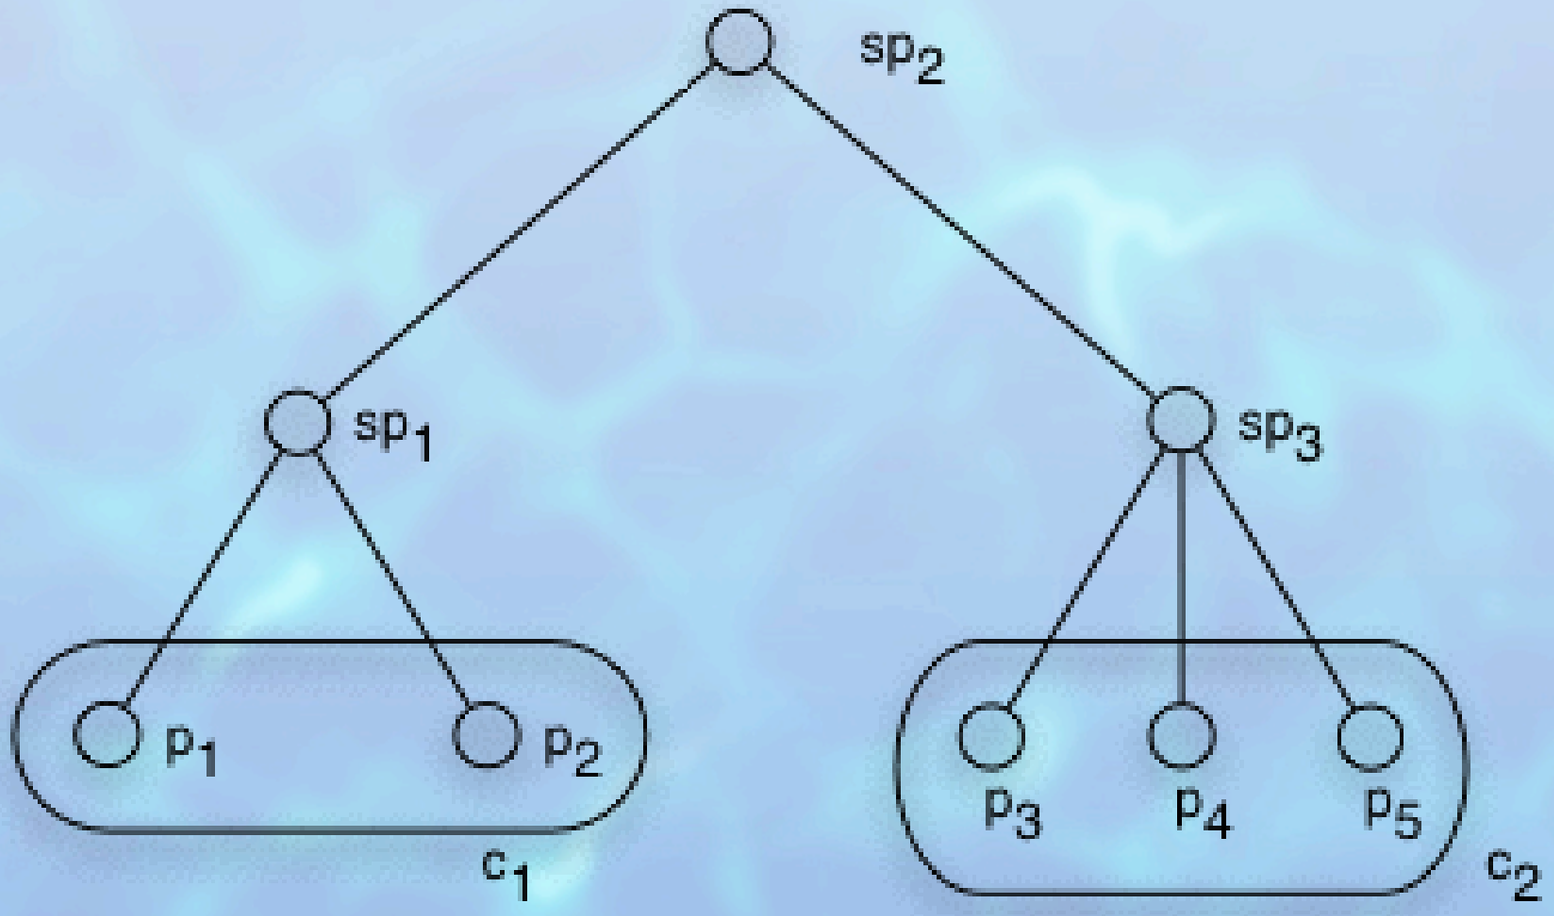

Network Evolution

- cluster splitting
- group splitting
- cluster/group merge
- vertical extension/contraction
- network re-joining

Query Processing & Query Algebra

Query Processing

- three phases
  - algebraic translation: performed locally by the submitting peer
  - location assignment compilation: performed by the super-peer network
  - query execution: coordinated by the submitting peer

Query Algebra

- extension of an existing algebra for queries over XML data
- new features
  - locations, modeling peer contents and replicas
  - freshness parameters
  - absolute time τ
  - replication constraints

Query Algebra Example

- **The XQuery binding:**

```
for $b in input()//building
let $d in $b/desc
```

- **Becomes:**

return{entry[$d,$p]}(
path{(//,$b,in)building[(/,$d,in)desc[0],
(/,$p,in)price[0]]}(
(loc1 • loc3 $|_{db1}$ loc2)(db1)))

Conclusions

- the system is being implemented
- future issues
  - correctness and completeness of queries
  - formal representation of the distributed state in p2p systems