Locality of Interest in Peer-to-Peer Content Location

J. L. Chiang

July 1, 2004

Distributed Hash Tables

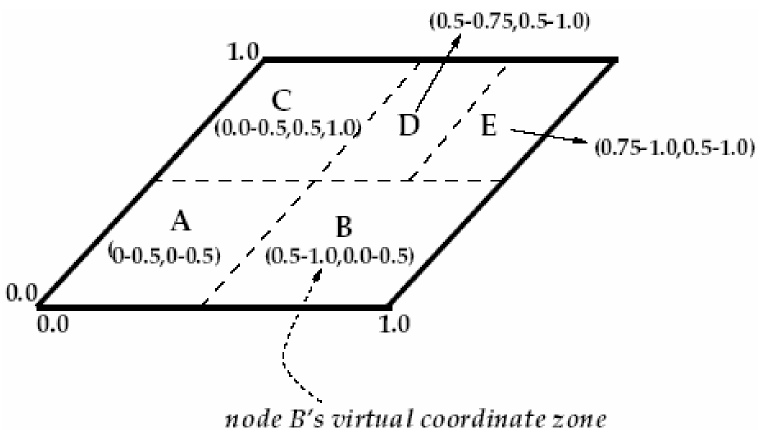

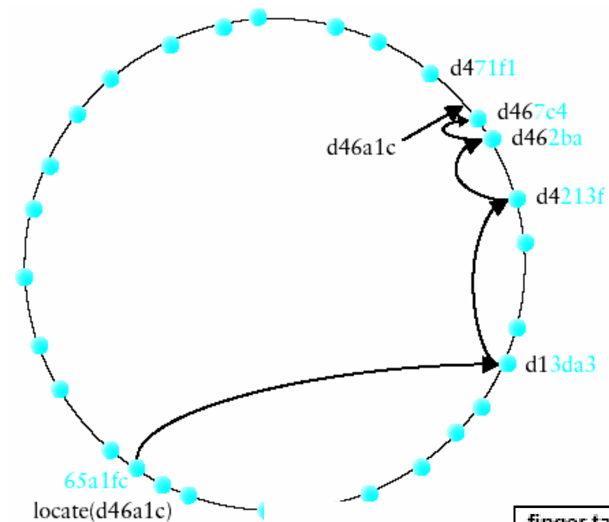

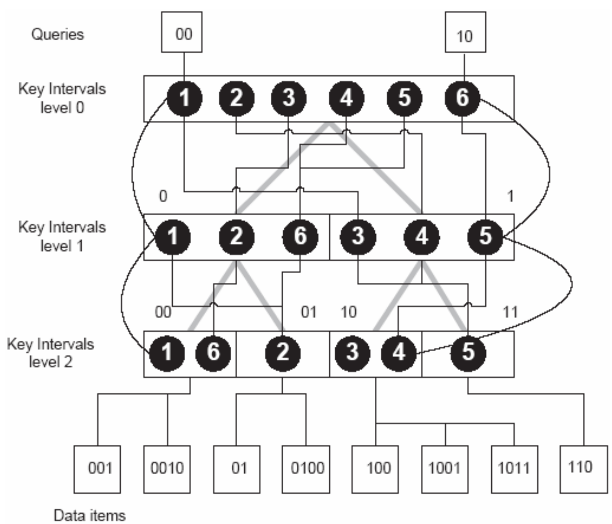

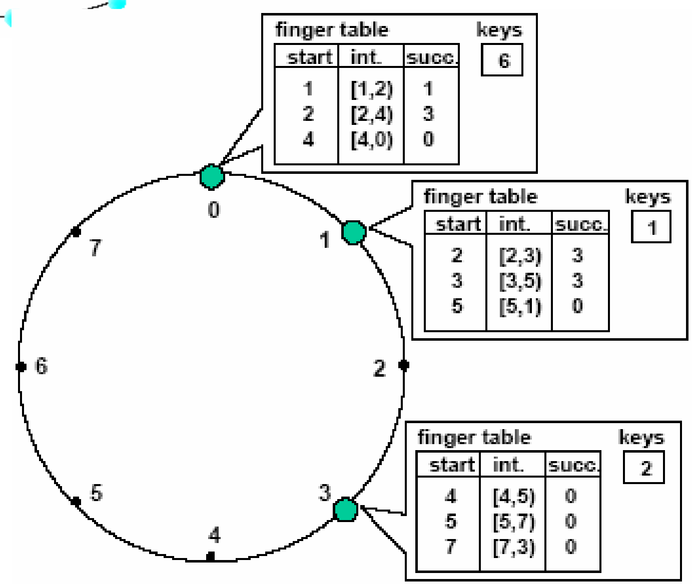

Limitations of DHTs

- DHT-based schemes only support exact-match through consistent hashing, but they can’t perform complex queries such as keyword search.
- DHT-based schemes incur larger over-head than unstructured architectures in the dynamic environment of frequent peer failure or disconnection.

Unstructured P2P

- Gnutella is a simple, robust, unstructured and fully decentralized P2P file-sharing system that easily support for complex queries and resilience to peer failures.
- The major weakness of Gnutella lies in the overhead of it's flooded messages and processing overhead.

Improving Gnutella

- Random walks[1]
- Replications[1,2]
- Shuffle[3]
- Supernode hierarchy[4,5]
- Routing indices
- Shortcuts[6]
- Interest-based locality[7-10]

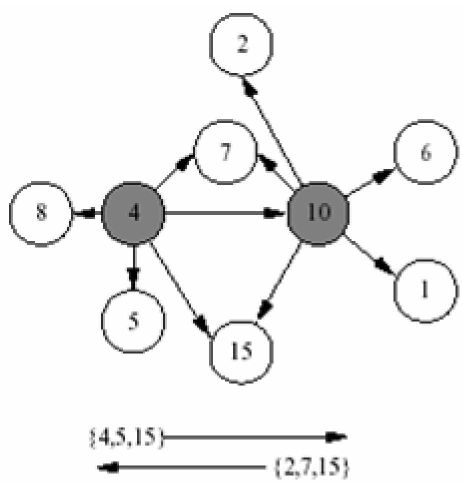

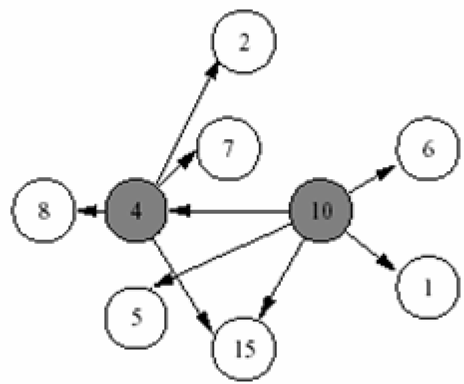

Locality of Interest

- “If a peer has a particular piece of content that one is interested in, it is very likely that it will have other items that one is interested in as well.”[6]
- “Two peers are said to have common interest if they share some of the locally stored files.”[7]
- “Nodes sharing similar interests always store similar documents.”[8]

Shortcuts

- Shortcut discovery
  - Piggy-backed on Gnutella
  - Shortcut list exchange
  - Content category structures

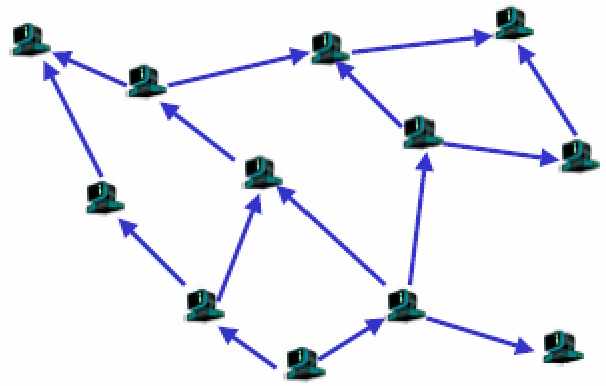

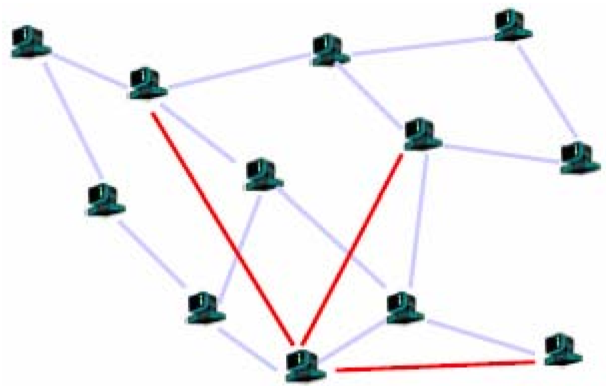

Shortcuts

- Shortcut selection (ranking)
  - Probability of providing content
  - Latency of the path to the shortcut
  - Available bandwidth of the path
  - Amount of content at the shortcut
  - Load at the shortcut
- The ratio between the number of times a shortcut was used to successfully locate content to the total number of times it was tried.

Shortcut's Shortcuts

- The success rates for discovering new shortcuts through existing shortcuts is higher than the basic algorithm.

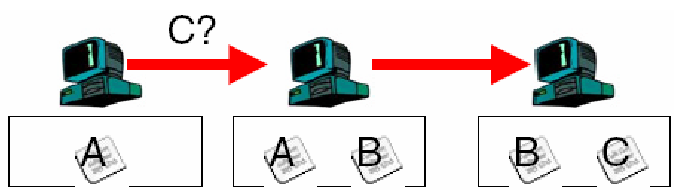

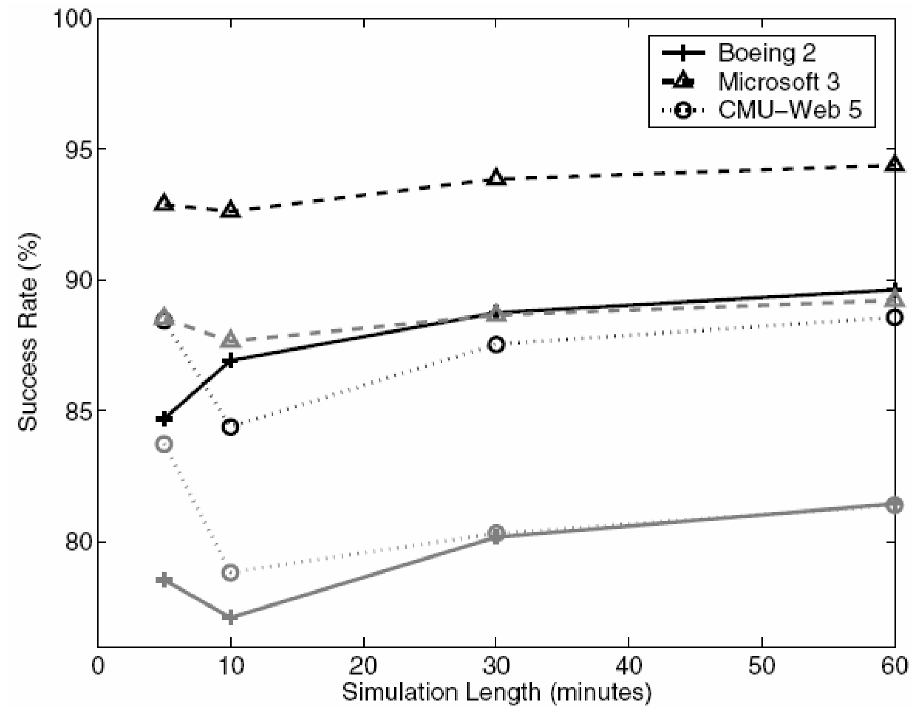

Similarity of Peers’ Interests

- Assessing the level of similarity among peers’ interests may be based on one or more of the following:
  - Query success rate (IG-1)
  - Locally-stored files (IG-2)
  - Sophisticated metadata (IG-3)
  - User Profiles (IG-4)

IG-1 [10]

- The rank of a peer can be computed as the percentage of reply messages it generates.
  - Indirect peers (s) v.s. immediate peers (q)

$$percReply_p(s) = \frac{Reply_p(s)}{\sum_i Reply_p(i)}$$

$$Imp_p(q,t) = \alpha * \frac{percReply_p(q)}{averNumHops_p(q)} + (1-\alpha) * Imp_p(q,t-1)$$

- $$averNumHops_p(q) = \sum_s \frac{Reply_{p,q}(s) * numHops_p(s)}{Reply_p(s)}$$

IG-2 [7]

- The common interest between two peers is assessed directly from the current content locally stored at both peers and is periodically re-evaluated.
- In order to limit overhead, the number of peers contacted and files compared by each peer is constrained to small values.
- The algorithm is run only after a certain fraction of the content has changed. (20%)

Creating an Interest-based Community

1: Start the community graph with the seed node at depth 0
2: For i = {0, 1} :
For each node p at depth i :
Randomly select a subset of p´s local files and a number of other peers known as participants
Send queries to each selected node with the list of selected files
For each peer q that responds with at least one shared file :
Add q to the community graph at depth i + 1
Add an edge between p e q with weight equal to the number of shared files returned by q
3: Create an artificial sink and connect each node at depth two directly to the sink with weight equal to 1.
4: Calculate the maximum flow of this graph: the seed node is the origin and the artificial node the sink.
5: Insert the nodes found as solution of step 4 in the community list of the seed node.

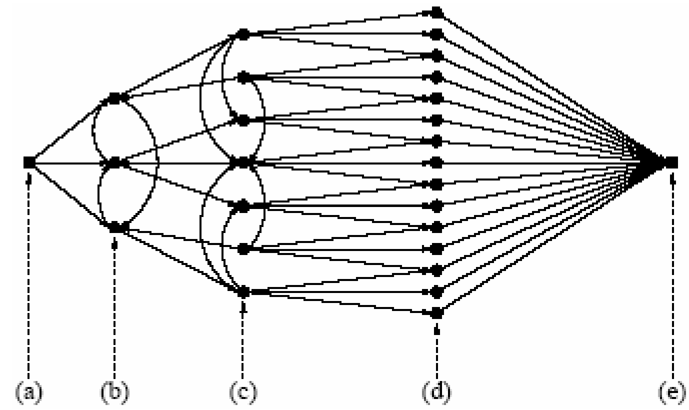

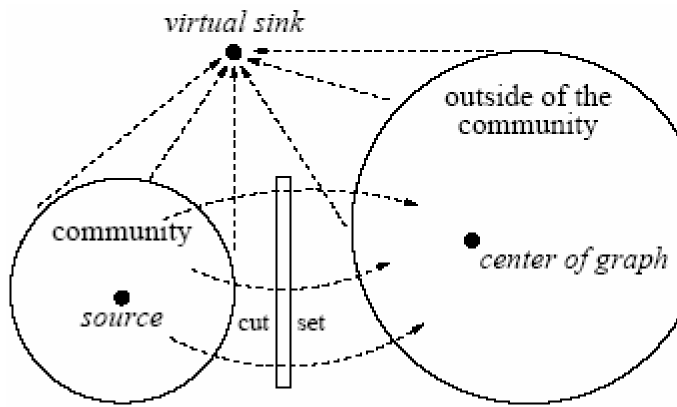

Max Flow-Min Cut

- Ford and Fulkerson proves that solving the maximum flow is identical to finding the minimum cut that separates s and t. [13]

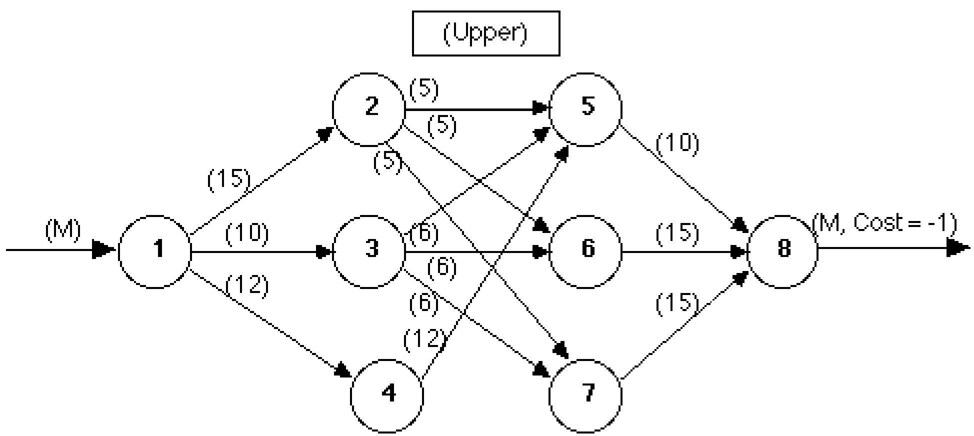

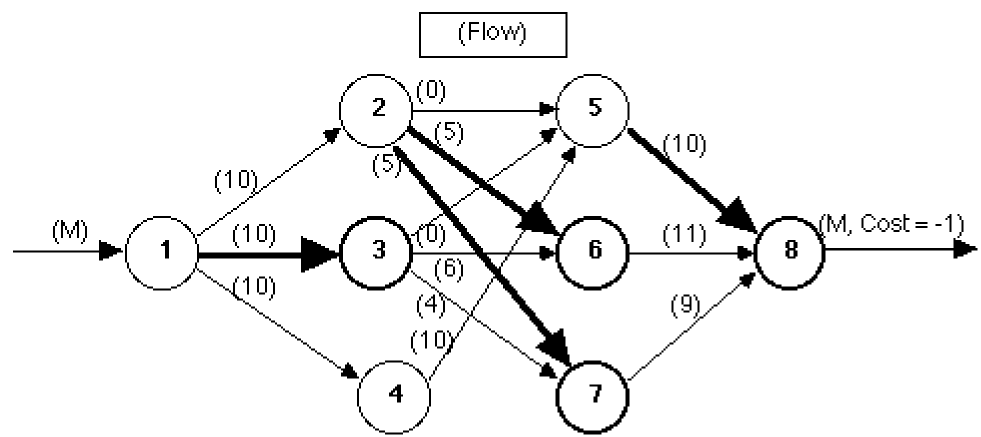

Performance

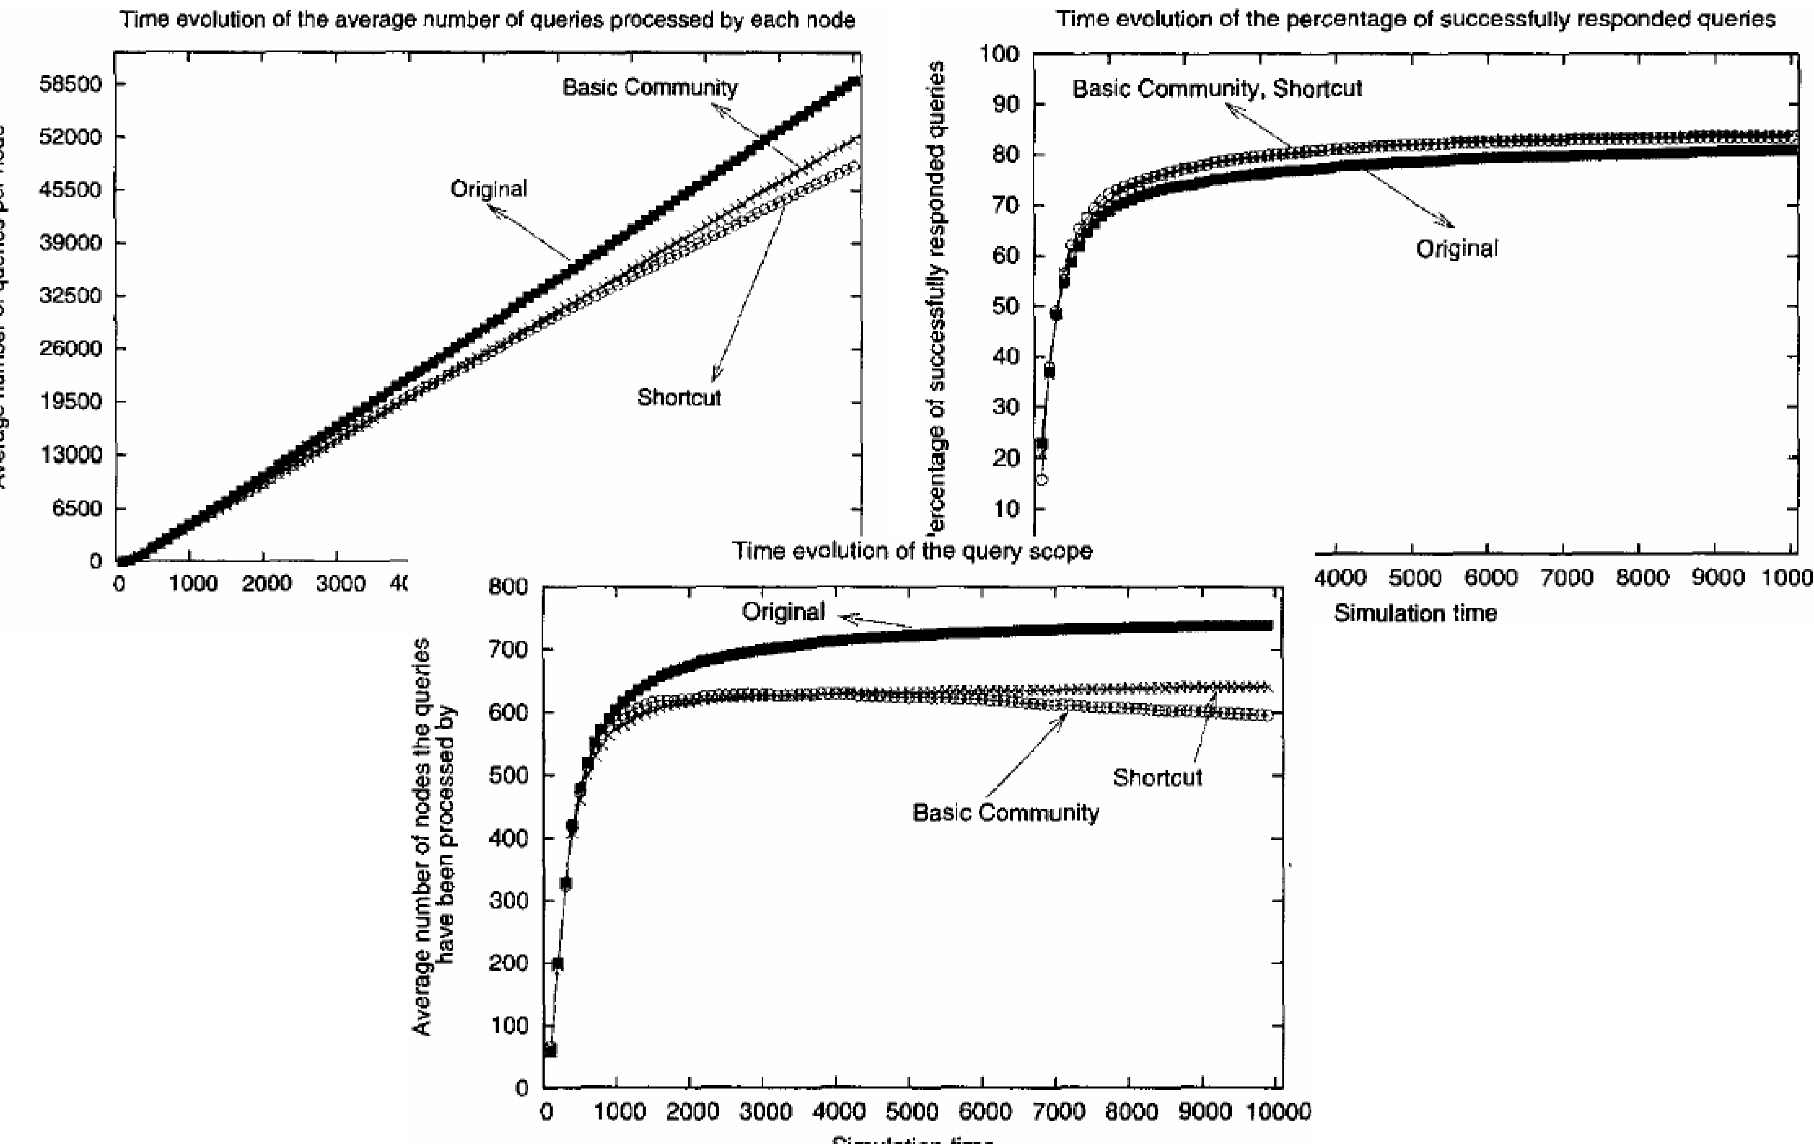

Modified

- Only n nodes that are known to share files with the seed node are added to the graph.

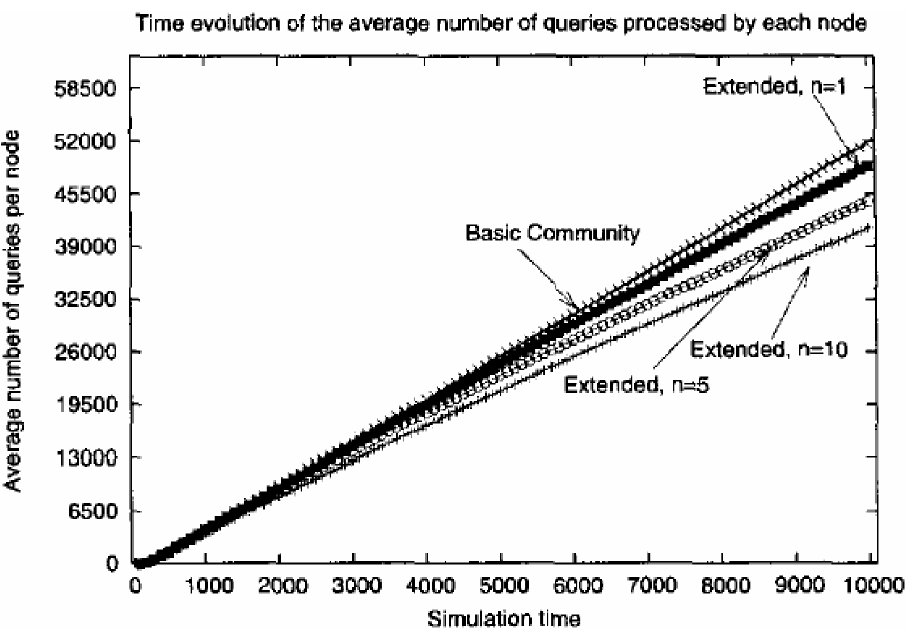

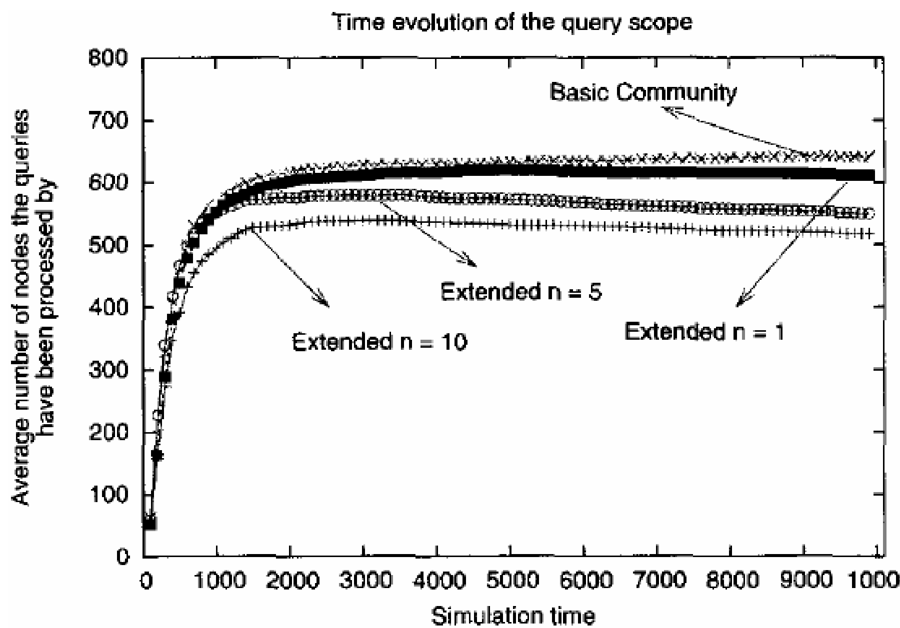

IG-3

- Metadata are used to describe and represent documents that nodes share with others.
- Metadata can be simply defined as data of data.
- The Dublin Core Metadata Element Set is a set of 15 elements (title, description, creator, data, publisher, etc.) that are useful in describing almost any web resources.

DC Metadata in the Web

```
<html>
  <head>
    <title>Distributed Metadata</title>
    <meta name="description" content="This article addresses...">
    <meta name="subject" content="metadata, rdf, peer-to-peer">
    <meta name="creator" content="Dan Brickley and Rael Dornfest">
    <meta name="publisher" content="O'Reilly & Associates">
    <meta name="date" content="2000-10-29T00:34:00+00:00">
    <meta name="type" content="article">
    <meta name="language" content="en-us">
    <meta name="rights" content="Copyright 2000, O'Reilly & Associates, Inc.">
    ...
  </head>
  ...
```

RDF

- W3C Resource Description Framework
- RDF provides a model for describing resources.
- RDF imposes formal structure on XML to support the consistent representation of semantics.
- Resource, Property-Type, Value

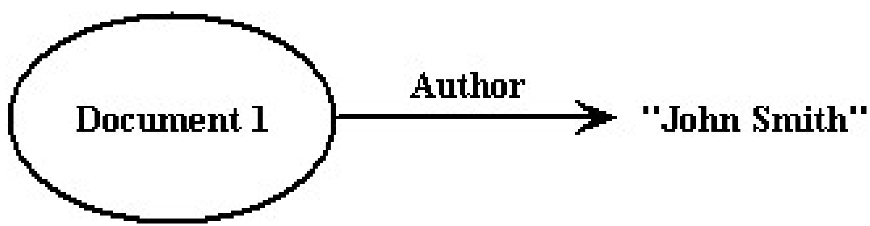

Namespace

- RDF uses the XML namespace mechanism for resource description communities in each domain to uniquely identifies their property-types.
- DC:CREATOR = “person or organization responsible for the creation of the intellectual content of the resource”

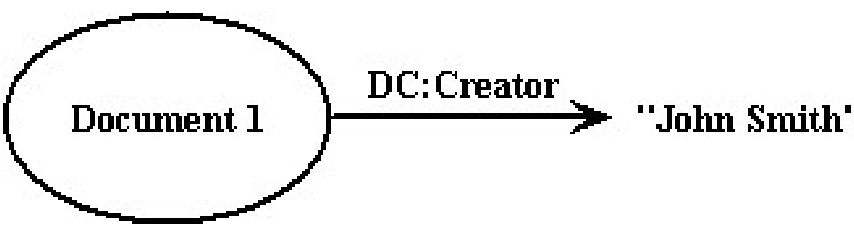

```
<?xml:namespace ns = "http://www.w3.org/RDF/RDF/" prefix = "RDF" ?>
<?xml:namespace ns = "http://purl.oclc.org/DC/" prefix = "DC" ?>
<?xml:namespace ns = "http://person.org/BusinessCard/" prefix = "CARD" ?>

<RDF:RDF>
  <RDF:Description RDF:HREF = "http://uri-of-Document-1">
    <DC:Creator RDF:HREF = "#Creator_001"/>
  </RDF:Description>

  <RDF:Description ID="Creator_001">
    <CARD:Name>John Smith</CARD:Name>
    <CARD:Email>smith@home.net</CARD:Email>
    <CARD:Affiliation>Home, Inc.</CARD:Affiliation>
  </RDF:Description>
</RDF:RDF>
```

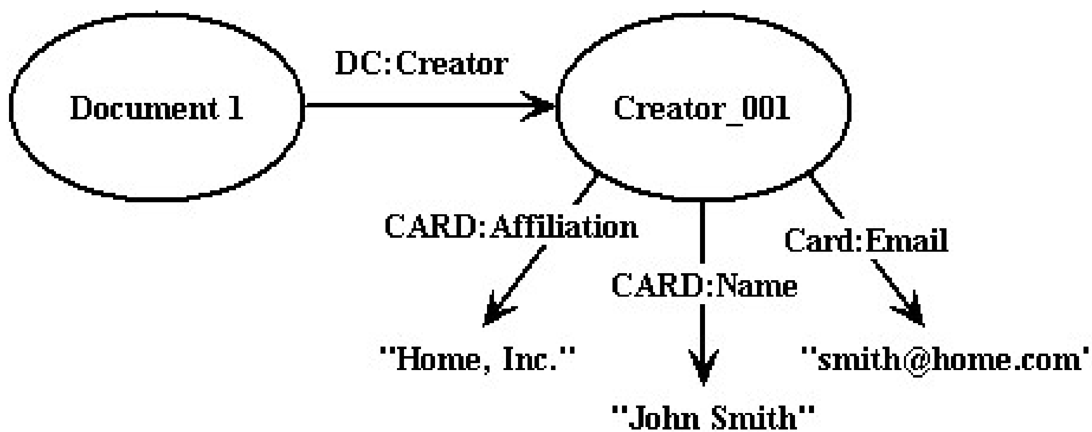

Metadata as Basis [8]

- Nodes in the same interest-group share similar interests, which indicate they have similar meta-data.
- Metadata using RDF may have very affluent descriptions about document, which makes query in interest-group can be complex and semantically rich.
- The size of metadata is usually very small, so metadata can be replicated among the interest-group (one node can store much more metadata than documents).

Notation

- ni: node; mr: metadata;
- Li maintains the (nk, mt) pairs learned by node ni;
- Mi maintains all metadata mr learned by node ni;
- Nir is the subset of nodes having metadata mr that node ni knows;
  - $N_{ir} = \{n \mid (n, m_r) \in L_i\}$

MSW

- Metadata Slide Window is proposed to select the metadata which is more effective than others for querying nodes.

8	7	6	5	4	3	2	1
m_t	m_r	m_s	m_r	m_r	m_s	m_t	m_t

- mt=8+2+1=11; mr=7+5+4=16; ms=6+3=9;

Search Procedure

1). Node n_i first matches the query against its local metadata repository M_i. If metadata m_r can be found locally, n_i completes the query successfully and goes to step 8); if not, n_i goes to step 2).

2). Node n_i checks its MSW. If MSW is empty, it chooses one metadata from its metadata repository M_i randomly; if MSW is not empty, the metadata in MSW with the heaviest or next heaviest sum of weights is selected. We denote by m_{sel} the metadata selected.

3). Using m_{sel}, n_i calculates the set $N_{i,sel}$ through the formula $N_{i,sel} = \{n \mid (n, m_{sel}) \in L_i\}$. $N_{i,sel}$ consists of nodes that store the metadata m_{sel}.

4). If $N_{i,sel}$ is not empty, then n_i randomly selects one node from $N_{i,sel}$; if $N_{i,sel}$ is empty, n_i repeats step 2) and 3) until $N_{i,sel}$ is not empty or finally after some unsuccessful retries n_i randomly selects one node from its neighbors in the P2P network. We denote by n_{next} the node selected.

5). Node n_i forwards the query to n_{next}. The search procedure is repeated from node n_{next} on until the query succeeds or some stop criteria are met and the query fails.

6). If the query succeeds, n_i goes to step 7). If the query fails, n_i goes back to step 4) and chooses another node in $N_{i,sel}$; if all nodes in $N_{i,sel}$ are selected and the query still fails, then n_i goes back to step 2) and chooses the metadata in MSW with the next-heaviest sum of weights, etc..

7). Suppose node n_{meta} is the destination node which stores metadata m_r. Node n_i contacts node n_{meta} and retrieves m_r.

8). From document location information stored in m_r, node n_i finds out the node n_{doc} which stores the desired document. Then node n_i inserts m_{sel} into the leftmost slot in its MSW as we described in section 4.2; retrieves the desired document from node n_{doc}; inserts (n_{doc}, m_r) pair into L_i; retrieves metadata repository M_{doc} and (node, metadata) pairs set L_{doc} from node n_{doc}. And M_{doc} is merged with M_i and L_{doc} is merged with L_i.

IG-4: P2People

- The objective of the P2People project is to research & develop a P2P collaborative framework and a prototype application to allow people to form "common interest" groups and provide those groups members with new ways to communicate, collaborate and make business together.
- While most innovative efforts focus on computer resources groups, P2People focuses on people.
- The P2People project provide services that are “user interest” tailored by describing the user itself through profiles.

User Profile

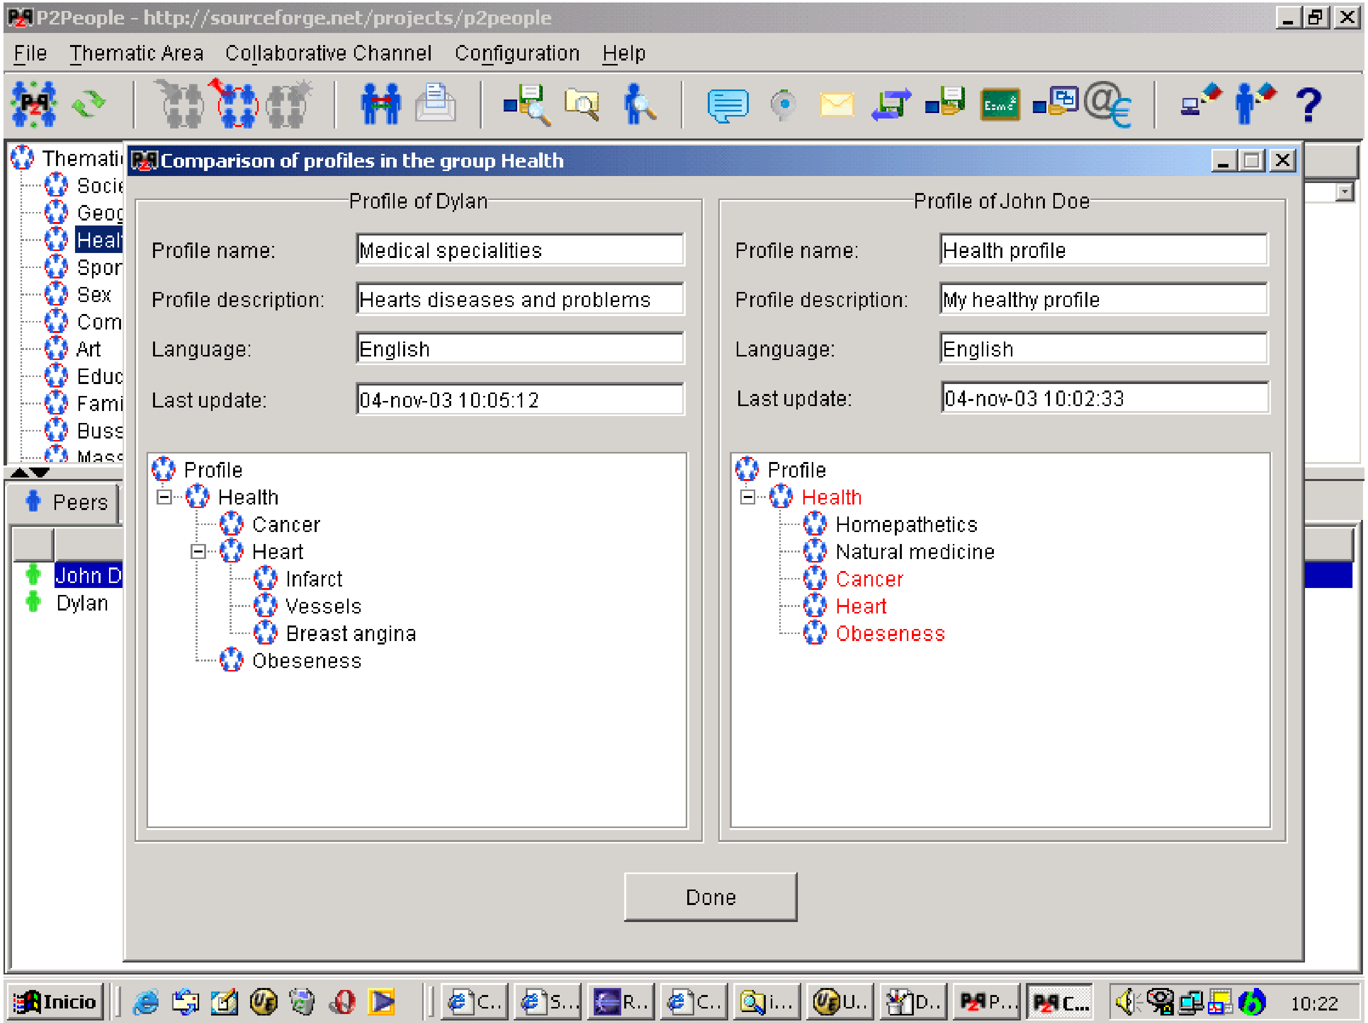

Profile Matching

- Common interest groups are formed by profile-matching among users.
  - Fixed tree: well-defined categories
  - Interest sub-tree: personalized themes

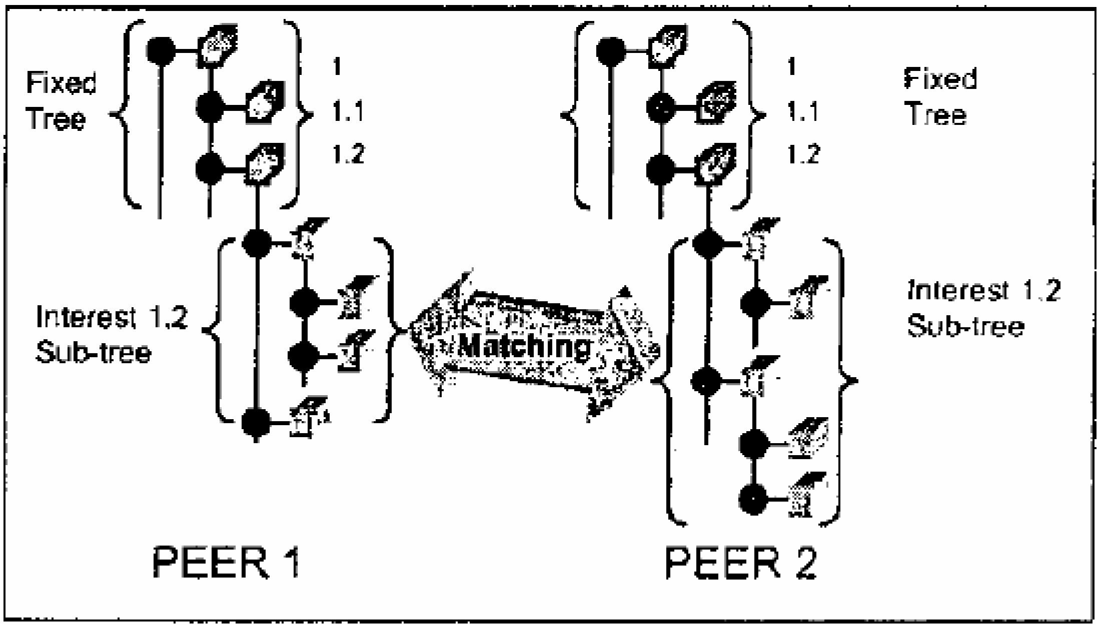

P2People Services

- A presence service
- A profile matching service
- **A reputation service**
- A messaging (mail) service
- A chat service
- A File Sharing service
- A File Browsing service
- A File Transfer service
- A screen sharing service
- A screen capturing service
- A voice&video communication service
- **A Payment service**

E-Commerce

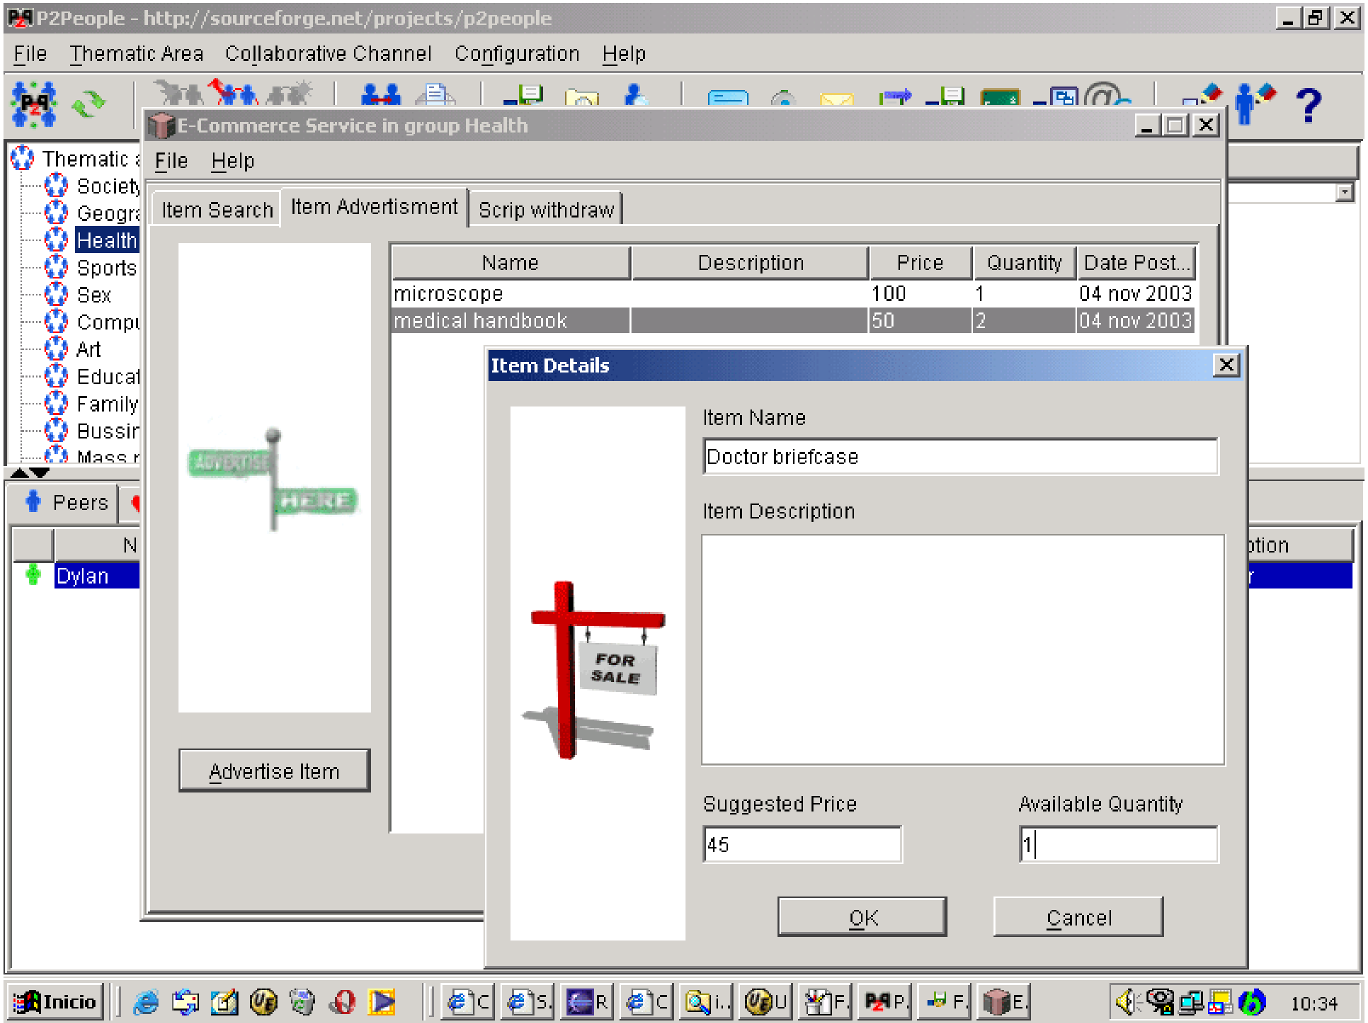

P2People - http://sourceforge.net/projects/p2people
File Thematic Area Collaborative Channel Configuration Help
E-Commerce Service in group Health
File Help
Item Search
Item Advertisment
Scrip withdraw
Name
Description
Price
Quantity
Date Post...
microscope
100
1
04 nov 2003
medical handbook
50
2
04 nov 2003
Advertise Item
Item Details
Item Name
Doctor briefcase
Item Description
Suggested Price
45
Available Quantity
1
OK
Cancel
FOR SALE
Peers
Dylan
Health
Sports
Sex
Art
Family
Inicio
10:34

Conclusions

- Well-defined categories and personalized themes should coexist with each other to reduce the complexity of the system while not to impede the uniqueness of users/events.
- Assessing the similarity of peer’s interests through locally store contents can achieve more precise result but it encounters enormous overhead.
- Metadata present the potential to efficiently locate any networked objects through resource description service provided by communities of each domain.

References

1. Q. Lv, P. Cao, E. Cohen, K. Li, and S. Shenker. Search and Replication in Unstructured Peer-to-Peer Networks. Proc. 16th ACM Int'l Conf. Supercomputing, 2002.
2. E. Cohen and S. Shenker. Replication Strategies in Unstructured Peer-to-Peer Networks. Proc. ACM SIGCOMM, 2002.
3. A. Stavrou, D Rubenstein, and S. Sahu. A Lightweight, Robust P2P System to Handle Flash Crowds. IEEE JSAC, 2004.
4. eDonkey. http://www.edonkey.com/
5. KaZaA. http://www.kazaa.com/
6. K. Sripanidkulchai, B. Maggs, and H. Zhang. Efficient Content Location Using Interest-based Locality in Peer-to-Peer Systems. Proc. IEEE INFOCOM, 2003.
7. M. W. Barbosa, M. M. Costa, J. M. Almeida, and V. A. F. Almeida. Using Locality of Reference to Improve Performance of Peer-to-Peer Applications. Proc. 4th Int'l Workshop on Software and Performance, 2004

References

8. J. Yang, Y. Zhong, and S. Zhang. An Efficient Interest-Group Based Search Mechanism in Unstructured Peer-to-Peer Networks. Proc. IEEE int'l conf. Computer Networks and Mobile Computing, 2003.
9. R. M. Ruiz, R. B. Vidal, and C. C. Palma. P2PEOPLE: P2P Common Interest Search Engine and Collaboration Platform. IEEE EUROCON, 2003.
10. V. Kalogeraki and F. Chen. Managing Distributed Objects in Peer-to-Peer Systems. IEEE Network, 2004
11. P2People project. http://www.p2people.org/index.html
12. G. W. Flakes, S. Lawrence, and C. L. Giles. Efficient Identification of Web Communities. ACM KDD, 2000.
13. L. R. Ford Jr. and D. R. Fulkerson. Maximal flow through a network. Canadian J. Math., 1956.