



Queensland University of Technology
Brisbane Australia

This may be the author's version of a work that was submitted/accepted for publication in the following source:

[Paik, Hye-Young](#), Benatallah, Boualem, & Toumani, Farouk
(2005)

Toward Self-Organising Service Communities.

IEEE Transactions on Systems, Man, and Cybernetics Part A: Systems and Humans, 35(3), pp. 408-419.

This file was downloaded from: <https://eprints.qut.edu.au/23250/>

© Copyright 2005 IEEE

This work is covered by copyright. Unless the document is being made available under a Creative Commons Licence, you must assume that re-use is limited to personal use and that permission from the copyright owner must be obtained for all other uses. If the document is available under a Creative Commons License (or other specified license) then refer to the Licence for details of permitted re-use. It is a condition of access that users recognise and abide by the legal requirements associated with these rights. If you believe that this work infringes copyright please provide details by email to qut.copyright@qut.edu.au

Notice: *Please note that this document may not be the Version of Record (i.e. published version) of the work. Author manuscript versions (as Submitted for peer review or as Accepted for publication after peer review) can be identified by an absence of publisher branding and/or typeset appearance. If there is any doubt, please refer to the published source.*

<https://doi.org/10.1109/TSMCA.2004.846402>

Toward Self-Organizing Service Communities

Hye-Young Paik, Boualem Benatallah, *Member, IEEE*, and Farouk Toumani

Abstract—This paper discusses a framework in which catalog service communities are built, linked for interaction, and constantly monitored and adapted over time. A catalog service community (represented as a peer node in a peer-to-peer network) in our system can be viewed as domain specific data integration mediators representing the domain knowledge and the registry information. The query routing among communities is performed to identify a set of data sources that are relevant to answering a given query. The system monitors the interactions between the communities to discover patterns that may lead to restructuring of the network (e.g., irrelevant peers removed, new relationships created, etc.).

Index Terms—Catalog portals, e-catalogs, peer-to-peer (P2P), self-adaptivity.

I. INTRODUCTION

CATALOG¹ portals, such as amazon.com or expedia.com, are becoming a prominent feature of the World Wide Web. However, the technology to organize, search, integrate, and evolve these portals has not kept pace with the rapid growth of the available information space. Currently, users need to access a number of information portals separately, using search engines, in order to collect relevant information. For example, in a computer market scenario, users need access to several Web portals (ones that are specialized in computer retail) separately, manually filter and organize the search results to obtain the complete picture of the products on offer. Clearly, to effectively realize the potential of Web-based information access, there is a need for facilitating the *integrated* access to *relevant* Web information portals.

Another aspect to be noted in current portals is that their centralized and static categorization often results in a rigid structure that is difficult to change or maintain. The categorization is designed to reflect the system designer's point of view, who has *a priori* expectations for how catalogs will be explored by customers. It is necessary to take into consideration how the catalogs are actually utilized by the customers as well as the portal itself. This may continuously minimize the gap between expectations of the system designer and customers. For example, if observations show that a catalogs on laptops and a catalog on

handheld computers are visited together most of the time, the system designer may decide to put them in the same category rather than in separate categories.

In the project WS-CatalogNet [1], [2], we proposed a platform through which catalog portals are built and interact with one another. This enabled a potentially large number of catalogs to act as one catalog to serve customers' queries. The approach is based on a hybrid of peer-to-peer (P2P) data management paradigm and Web services architecture. It uses the notion of *catalog service communities*² where catalogs catering for similar customer needs are grouped together [3], [4] and form a single *community* (i.e., a peer node). Catalogs *register* themselves to a catalog service community (hereafter, community) as *members* and the communities themselves interact with other peers. Each peer (i.e., the community) in WS-CatalogNet can be viewed as a domain-specific data integration mediator, which holds meta-data and registry information about its members (i.e., catalogs). Queries are processed by the members of a peer, but routing of the queries is a responsibility of the peers. It should be noted that the purpose of the query routing in our system is to identify a set of members that, when put together, can satisfy *all* constraints required by a query. Hence, a routing takes place before the actual query process. Once a set of members (not necessarily from the same peer) are identified, queries are sent to each member in the set for processing. The results are combined by the original community.

WS-CatalogNet work, so far, has focused on creating, linking, and querying the community network. In this paper, we investigate the ability to evolve from the initial design of the system. The evolution may involve capabilities like "discovering" members (i.e., catalogs) or other communities (i.e., peers), splitting a community, reforming the relationships, etc. As communities are developed and relationships are formed, the interactions between communities can be monitored. We analyze the monitoring result to draw *interaction patterns* that may trigger restructuring of the network. We also offer a set of operations for restructuring community networks.

The paper proceeds as follows. Section II looks at the related work. In Section III, we present the overview of WS-CatalogNet. Section IV discusses the monitoring/adapting techniques for restructuring of the community network. Sections V and VI describe the implementation and evaluation results, followed by concluding remarks in Section VII.

II. RELATED WORK

The catalog portal framework proposed in WS-CatalogNet combines the Web Services architecture and P2P data management paradigm. Instead of developing a single global system

²The term service community comes from the fact that the community is implemented as a Web service.

Manuscript received June 8, 2004; revised October 11, 2004 and November 17, 2004. The work of B. Benatallah was supported in part by ARC Discovery under Grant-DP0452942. This paper was recommended by the Guest Editors.

H.-Y. Paik is with the School of Information Systems, Queensland University of Technology, Brisbane 4001, Australia (e-mail: h.paik@qut.edu.au).

B. Benatallah is with the School of Computer Science and Engineering, University of New South Wales, Sydney 2052, Australia (e-mail: boualem@cse.unsw.edu.au).

F. Toumani is with the LIMOS, ISIMA, University Blaise Pascal, Clermont Ferrand 63173, France (e-mail: ftoumani@isima.fr).

Digital Object Identifier 10.1109/TSMCA.2004.846402

¹We use catalogs and e-catalogs interchangeably.

or registry, we explore the possibility of using “P2P”-ed Web services as a flexible and scalable framework for sharing data across a large number of online sources.

Early P2P systems [5], [6] are based on the common objective: retrieval of a document or a file. Reference [7] offers a comprehensive survey of such systems. Our work is more in-line with current efforts that leverage database paradigm with P2P, binding a database capability with a peer [8], [9]. In this paradigm, it is assumed that each peer has data to share, which resides in a database. One of the relevant issues is to resolve the semantic interoperability between peers. This involves a technique that allows a peer to express how its data relates to the data in other peers (often known as a schema mapping). Using the mapping information, a query posed on one peer can be reformulated to a query on another peer. In many ways, the work in multidatabase systems (e.g., MOMIS [10], TSIMMIS [11], etc.), and distributed query processing [12] still apply in this paradigm. In these systems, however, a user issues a query based on a single global schema, and the system maps the query to subqueries that can be understood by the underlying data sources. Typically, the data sources are assumed to be relational tables only. Also, the global schema is created and maintained centrally by a database designer. These approaches are not directly applicable to the environment in which WS-CatalogNet operates. The data sources can be in any format, the number of participating data sources may be large, and any data source may enter or leave the system frequently. Therefore, this paradigm calls for a more flexible and scalable approach.

The PeerDB project [13] uses a relational model to describe the schema of a peer. It employs an information retrieval (IR) approach for the query routing to avoid the explicit specifications of mapping between peer schemas. The query reformulation is assisted by an agent automatically identifying the mapping based on a built-in strategy. On the other hand, Piazza [14] considers using a global schema and schema mediation techniques in P2P environments. Each peer schema is described using a relational model. It proposes a language for specifying mappings among peers and a query reformulation algorithm for translating the global schema to a local peer schema. Hyperion [15] proposes use of mapping tables which map data *values* (rather than the schema) between two peers. Such mapping tables are constructed by domain experts. These approaches assume no knowledge about the underlying schema of the peers when forwarding a query, but the efforts required to identify the correct mappings are left with the users.

The systems mentioned above are based on the pure P2P architecture, where individual data sources are connected directly as peers, whereas in WS-CatalogNet, each peer is a data integration mediator (similar to the concept of super-peer in [16], [17]) holding a registry/mapping information for the individual data sources. This helps in organizing the available information space into manageable, meaningful spaces. Also, the burden of query reformulation/translation is distributed among the peers, rather than the individual data sources. Another important distinction in our work is that WS-CatalogNet does not make any assumption about the nature of the local data source whereas other works we mentioned focus on either relational databases or XML/RDF resources only.

To the best of our knowledge, little work has been done in addressing the issue of restructuring catalog community networks, as tackled in this paper. There have been some related work in similar areas (tuning P2P networks and ontology restructuring). For example, [18] proposes a self-monitoring P2P network that *tunes* itself according to the workload of a peer. The tuning is based on break() and connect() operations. The network concerned is set on the earlier models of P2P systems (i.e., simple file sharing) and it does not consider ontological relationships or interaction patterns between peers. Authors of [19] offer a framework for an ontology discovery, reuse and evolution in a distributed environment. Their work could potentially be applied to ours in the continuing maintenance of the registry/mapping information in a community.

III. WS-CATALOGNET: OVERVIEW

We first set the conceptual background of WS-CatalogNet as a catalog portal framework [1]–[4]. After an overview of WS-CatalogNet, we will then focus our discussion on monitoring and restructuring. The main ideas of WS-CatalogNet are the concepts of catalog service communities and *peer relationships* among them. Further details about the framework we propose for building and peering communities can be found in [2] and [20]. Here, we only describe the concepts that are central to understanding the monitoring/restructuring communities.

A. Catalog Service Communities

Fig. 1(a) shows the metadata model implemented by WS-CatalogNet. A community is a *container* of catalogs of a specific domain (e.g., community of Airlines, CarRentals). A *community ontology* provides a description of desired products (e.g., airline tickets) without referring to an actual catalog provider (e.g., Qantas Airlines). Catalog providers register their catalogs into a community as members by *exporting* (all or part of) their descriptions. Moreover, communities themselves can be linked to facilitate interoperability across overlapping domains.

A community is described by a set of categories. A category, in turn, is described by a set of attributes. Categories within a community may be interrelated via the specialization (subsumption) relationship. For example, in Fig. 1(b), the community FlightCenter has the category Flights, which is described using such attributes as arrival, departure, and price, etc., and the category has two subcategories (i.e., specialization): domestic and international.

To provide formal semantics, which are necessary for precise characterization of queries over the communities, we use a (concept) class description language that belongs to the family of description logics [21]. In the following, we illustrate the main constructs of this language via examples³.

In WS-CatalogNet, a community ontology is described in terms of *classes* (unary predicates) and *attributes* (binary predicates). Class descriptions are denoted by expressions formed by means of the following constructors.

- 1) Class conjunction (\sqcap), e.g., the description `Travel \sqcap Accommodation` denotes the class of products which are

³A description of the syntax and semantics of the proposed language is presented in [2].

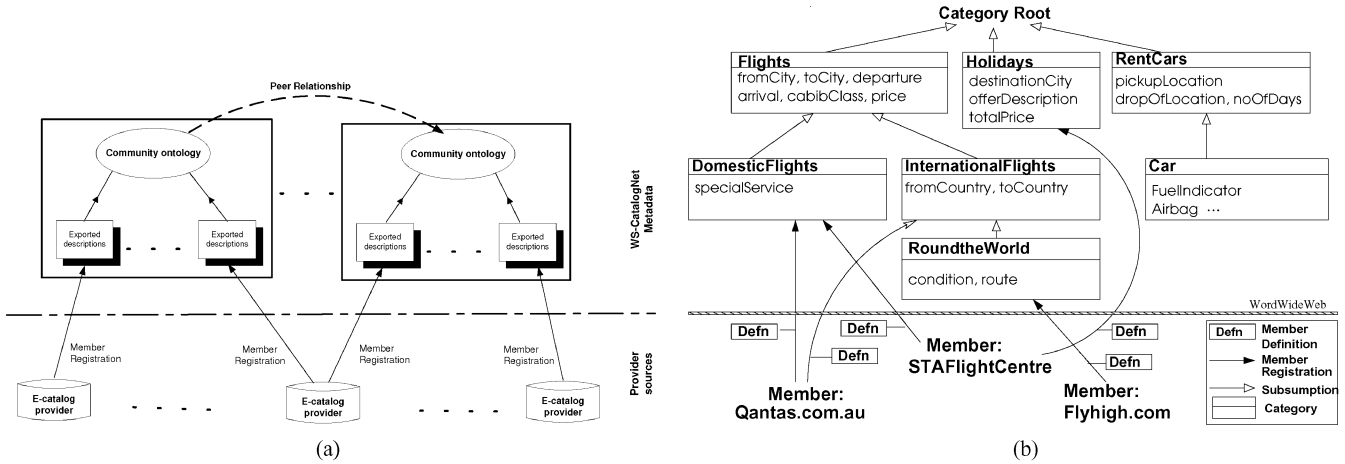


Fig. 1. Community ontology. (a) Metadata representation of the network. (b) A community definition (categories and members).

instances of the classes travel and accommodation (e.g., a hotel),

- 2) The universal attribute quantification ($\forall RC$), e.g., the description $\forall arrivalDate.Date$ denotes the class of products for which all the values of the attribute $arrivalDate$ are instances of the class $Date$ (i.e., the data type of the attribute $arrivalDate$ is $Date$),
- 3) The existential attribute quantification ($\exists R$), e.g., the description $\exists Price$ denotes the class of products having at least one value for the attribute $Price$.

Category Definition: A category definition is specified as follows. $Cname \equiv CatDescr$, where 1) $Cname$ is the name of the category and 2) $CatDescr$ is a class description that defines $Cname$.

For example, using the language, the category *Domestic Flights* in Fig. 1(b) can be described as follows:

$$\text{DomesticFlights} \equiv \text{Flights} \sqcap \\ \forall specialService.String \sqcap \exists specialService$$

This definition states that the category *Domestic Flights* inherits all the attributes of the category *Flights*, and has one additional attribute, namely *specialService*.

We would like to emphasize that the use of a formal language is transparent to community providers and users. Indeed, WS-CatalogNet provides a graphical editor that supports a community provider in creating a community and defining the community ontology. After the definition of an ontology, the editor automatically generates the class descriptions of the ontology (i.e., class descriptions of all categories).

Member Definition: When a catalog is registered to a community, the catalog provider supplies a capability description which specifies the categories and attributes of the community ontology that are supported by the catalog. This form of information is referred to as a *member definition or description*. Member definitions are also converted to class descriptions expressed using the description language and stored in the community meta-data repository.

A member definition specifies the query capabilities of a given catalog as follows: $Mname_Dname \equiv MDescr$, where $Mname_Dname$ is a member definition name made of $Mname$,

a member name (i.e., an unique identifier of a member), and $Dname$, the name of a description. $MDescr$ is a class description that specifies which products are actually provided by this member. For example, the catalog *flyhigh.com*, which offers a range of flight information, can be registered with *FlightCentre* using the following member definition:

$$\text{flyhigh.com_International} \equiv \text{InternationalFlights}.$$

This definition states that *flyhigh.com* supports all attributes in the category *InternationalFlights* as well as all attributes inherited from the category *Flights* (as *InternationalFlights* is a subcategory of *Flights*). Each member can provide several definitions. When registering with a community, the catalog providers are in charge of providing and maintaining the exported descriptions of their catalogs as well as the wrapping utilities required to translate user queries into the catalog native query models.

Community Definition: A community ontology consists of a tuple $CS = (C, S)$, where C is a set of category definitions, S is a set of community names that have peer relationships with the community. A catalog community consists of a tuple $CAT = (CO, M)$, where CAT is the community name, CO is the community ontology, and M is a set of member definitions. We assume that the set of class descriptions in a catalog community (i.e., category definitions and member definitions) is acyclic, i.e., there does not exist cyclic dependencies between class definitions.

B. Peering the Communities

Through a *peer relationship* between communities, they forward queries to each other. To form a peer relationship, community providers need to discover other communities whose domains are relevant/similar to their communities. The community registry is hosted by WS-CatalogNet for the purpose. Typically, a provider of a community C_i may decide to form a peer relationship with another community C_j , if C_i has categories that are considered analogous, or interchangeable to C_j 's categories (e.g., category *Accommodation* in community *Travels* and category *Budget Hotel* in community *Hotels*). When a link is formed, a weight (a value between 0 and

1) is attached to the link as a way of indicating the “relevancy” of the peer relationship.

C. Collaborative Query Processing Between Communities

This section explains our approach for the query reformulation. We use a rewriting algorithm, called BQR,⁴ that allows to reformulate a community query⁵ Q , expressed as a class description over the community ontology, into subqueries that are expressed in terms of the community member definitions. A novel feature of BQR is that besides the rewriting of the community query, it also identifies parts of the query that cannot be answered by any of the community members. This feature forms the core of our collaborative query processing.

The algorithm takes as input a catalog community definition $CAT = (CO, M)$ and a community query Q which is written in terms of the ontology CO . It then computes a set of rewritings $R(Q) = \{r_i(Q)\}$. A rewriting r_i is a couple $r_i = (Q_{local}, Q_{rest})$ where 1) $Q_{local} = \{(q_j, m_j)\}$ is a set of pairs (q_j, m_j) where q_j is the part of the query Q that can be answered by the member definition m_j and 2) Q_{rest} is the part of the query Q that cannot be answered by the members of the actual community.

Let us consider a simplified example. Assume a query “select flights from FlightCentre where fromCity = Sydney, toCity = Paris, price < 2000 and Insurance = included.” The key attributes in this query are $\{fromCity, toCity, price, Insurance\}$. Assume that the outcome of BQR is: $Q_{local} = \{(fromCity, toCity, price), Qantas\}$, $Q_{rest} = \{Insurance\}$. This means that the member Qantas of FlightCentre community has fromCity, toCity, and price information. However, there is no member in the community who can answer Insurance. Therefore, to serve the query, the community will collaborate with peers to identify external member(s) who can provide the missing information.

Note that a query Q can lead to several alternative rewritings. For example, the Q_{local} of Q is not necessarily unique. That is, there could be many members that can answer the same part of the query. In [20], we introduced an utility function that selects *only* those rewritings that maximize the user satisfaction with respect to a given quality of service criteria (called *best quality rewritings*).

In a nutshell, the query rewriting problem can be formally stated as follows. Let $\mathcal{C} = \{c_i \equiv \text{description}_i, i \in [1, n]\}$ be a set of class definitions corresponding to member definitions, and let Q be a class definition that denotes a community query. Then, can Q be reformulated as a conjunction of class names $E \equiv c_{i_1} \sqcap \dots \sqcap c_{i_m}$, with $1 \leq m \leq n$ and $c_{i_j} \in \mathcal{C}$ for $1 \leq j \leq m$, such that E contains as much as possible of *common information* with Q (E is called a rewriting of Q using \mathcal{C})?

Of course, the complexity of this algorithm is transparent to the users. The graphical query interface in WS-CatalogNet lets the user easily formulate a query (by point&clicks). The user clicks a category and then selects attributes to be queried

on, and specify values for the attributes if desired (e.g., “category: Flights, attributes: destination, departure, price, values: price \leq 2000”). The query interface automatically converts the user formulated query to a class description.

The collaborative query processing technique consists of two steps. Whenever a query is submitted to a community C , it does the following.

- Step 1) Identify the combination of catalog members whose query capabilities, when put together, satisfy all constraints expressed in the query. The catalog members can be local (i.e., belonging to C), or external (i.e., belonging to C 's peers). This step is realized by the BQR algorithm and is referred to as *Meta Query Stage*.
- Step 2) Send the subqueries to the identified members and collect the results. This step is referred to as *Actual Query Stage*.

Each community has a *forwarding policy*. This controls what should be done with Q_{rest} in the meta query stage. The forwarding policy can express: 1) which part of the query should be forwarded; 2) when the query should be forwarded (e.g., when no local members can answer, when the community is too busy etc.); 3) to which peer (e.g., all, top K , random, etc.) the query should be forwarded; and 4) how far the query should be forwarded (hop limit). From the previous example, the result of forwarding may be $\{Insurance, STA.BestTravel\}$, meaning that BestTravel (a member of community STA) can provide the Insurance information. It should be noted that, although important, the issue of assembling actual results returned by selected catalogs, is outside the scope of the discussion. As a naive solution, we assume that every product information exists in WS-CatalogNet carries a universal product identifier which is uniquely understood (e.g., in the domain of flight tickets, the flight number is unique). This identifier is used to combined (i.e., join) the results returned from different members in the actual query stage.

IV. RESTRUCTURING COMMUNITIES NET

In this section, we introduce an approach to monitor and adapt the community network. The approach we propose consists of the following three building blocks.

- 1) *Logging community events*: Any events occurring within a community or between communities are logged. Such events may include receiving a query from a user, sending a query to a peer, receiving an error from a member, etc.
- 2) *Analyzing community interaction patterns*: The event logs reflect what kind of interactions exist between communities. It is possible that over a period of time, such interactions would form certain patterns (e.g., community A sends a query to community B instead of community C. The member M of community A does not respond to queries most of the time, etc.). Among all possible patterns, we predefine *patterns of interests*. Each pattern of interest is linked with a restructuring operation. If any of the patterns of interests is observed in the log, the system will advise the administrator to perform the associated

⁴The readers are referred to [2] for details and proof of this algorithm.

⁵The term refers to the query that is submitted to a community by a user.

TABLE I
EXAMPLES OF THE COMMUNITY ACTIVITY EVENTS, NOTED \mathcal{E}

Meta Query Stage Event	Description
MQUserQryReceived(logger, when, qry, qID)	When a community receives a query from the user. <i>logger</i> is the name of the community who is logging this event, <i>when</i> records the time of the event, <i>qry</i> is a string that contains the query itself and <i>qID</i> is a unique identifier issued for the query for tracking purposes.
MQForwardedPeer (logger, to, when, qry-part, qID)	When a community forwards <i>qry-part</i> to the peer(s). Note that <i>qID</i> is a unique identifier issued for the original query.
MQAnswerReceived (logger, from, when, answer, qID)	When a community receives the corresponding result of the forwarding from one of its peers.
MQLogTimeOut (logger, from, qID)	When a community does not receive a response from a peer within a given time. <i>from</i> is the peer community that failed to send the response back.
Actual Query Stage Event	Description
AQPeerQryReceived (logger, from, when, qry, qID)	When a community receives the query that is to be locally executed by its members. <i>from</i> is the sender, <i>qry</i> is the query that is to be executed.
AQSentLocalMember (logger, to, when, qry-part, qID)	When a community sends the query to one of its members. <i>to</i> is the member of the community, <i>qry-part</i> is the query that is to be executed.
AQLocalAnswerReceived (logger, from, when, ans, qID)	When a community receives the corresponding result (<i>ans</i>) of the query processing from one of its members.
AQExtAnswerReceived (logger, from, when, ans, qID)	When a community receives the corresponding result of the query processing from an external member of a peer.

operation. The operation can also be automatically performed based on even condition action (ECA) rules defined by the administrator.

- 3) *Restructuring operations*: An adaptation process is realized by a set of community network restructuring operations such as `addPeer()`, `upgradePeer()`, `mergeCommunities()`, etc.

WS-CatalogNet realizes the above principles via *Community Monitor*. Every individual community is associated with a community monitor. It gives the community access to the operational knowledge, such as interaction patterns (i.e., querying for a pattern) as well as a means for analyzing event logs and invoking the restructuring operations.

Under the current implementation, an administrator (e.g., the community provider) would use the monitor to explicitly search a pattern and perform restructuring operations. However, the design of the *Community Monitor* also supports both automatic and manual execution of an *action*. The automatic actions are controlled by *monitoring policies* designed by administrators. The policies specify restructuring directives via ECA rules. For example, a monitoring policy may specify that a peer link will be disabled after a number of unsuccessful query forwarding. Therefore, the complete process of monitoring, detecting possible restructuring needs and performing the operations can be scheduled and automated via the policies.

In the remainder of the paper, a number of interaction patterns and the restructuring operations will be discussed. Before proceeding, we explain how the events are handled.

A. Event Collectors

Collecting events is facilitated by means of a logging service called *Event Collector*. The community monitor captures the events and sends them to the event collector for logging.⁶ To maintain a uniform way of tracking interactions, all event collectors in WS-CatalogNet agree on the system-wide events and their parameters to be logged. A subset of the events and their parameters are shown in Table I. Note that the events are

divided into two classes: *Meta Query Stage* events (names are prefixed with MQ) and *Actual Query Stage* events (names are prefixed with AQ). The MQ events occur during the first step of query processing where the communities collaborate with each other to identify the set of members (local and external) that are relevant to a query. The AQ events occur in the second step of query processing, where the community sends the subqueries to the selected members and collect the results.

To make use of the log data, we make necessary transformation on the log file (e.g., cleansing, sessionising). The log collected from a community is kept privately to the associated community monitor. We make an assumption that the log from all monitors can be shared by uploading them periodically to a central location for more global analysis. Hence, the log is analyzed in two scopes: *local* and *global*. In the local scope, a community's administrator uses the log from the community's monitor and focus on what kind of restructuring can be done to the immediate peers (i.e., peers that are directly linked from the community), as well as to the community. The global scope analysis is carried out by WS-CatalogNet system administrators using the log uploaded from all monitors. The restructuring in the global scope focuses on having a global view of interactions between all communities that otherwise not possible to detect from looking at the local scope only.

B. Query Interaction Patterns

We now introduce a technique to assist the administrators to make decisions as to what kind of adaptation should be done. We propose predefined interaction sequences (PIS) as a means to analyze community interaction patterns that may interest administrators. A PIS represents an interaction pattern between communities. The idea is that we associate a restructuring operation (e.g., removing a link, merging communities, adding a new member to community, etc.) with a particular PIS. If an analysis shows that a particular PIS occurs frequently in the event log, an action (i.e., the restructuring operation(s) which is associated with the PIS) is performed either automatically via a monitoring policy, or manually by the administrator. A predefined interaction sequence is formally defined as follows:

⁶Logging can be made either to a file or a database.




Fig. 2. Interaction patterns for upgrading/downgrading peer relationship.

Definition 1: Predefined interaction sequence (PIS). \mathcal{E} denotes the set of all defined activity events (see Table I). A predefined interaction sequence PIS of length n ($0 \leq n$) is a vector of ordered events $\text{PIS} = \langle e_1, e_2, \dots, e_n \rangle$ where $e_i \in \mathcal{E}$ ($i = 1, \dots, n$).

A PIS is matched against each query session⁷ in the log file to check whether the sequence exists in the session. We refer to the number of occurrence of a PIS in the log file as *Frequency*.⁸ In the following, we present some of the predefined interaction sequences.

1) *Local Scope Analysis:* For the local scope analysis, the PIS is matched against query sessions in a local log file of a single community.

a) *Upgrading/Downgrading Relationship:* Consider an interaction sequence in Fig. 2(A), this pattern is designed to capture situations where community c_i forward a query to one of its peers c_j and c_j returns an answer. The fact that this pattern occurs frequently implies that c_i collaborates with c_j frequently and c_j is responsive to c_i 's query. In this case, the administrator of c_i may choose to upgrade the weight of the relationship. We define the related PIS as follows.

Definition 2: (PIS_{up}) (PIS upgrade) The pattern represents the situation where a forwarding of a query to a peer (represented by the event MQForwardedPeer) is followed by a return from the peer (represented by the event MQAnswerReceived)

$$\text{PIS}_{\text{up}} = \langle \text{MQForwardedPeer}(c_i, c_j, \text{qid}), \text{MQAnswerReceived}(c_i, c_j, \text{qid}) \rangle$$

where c_i, c_j are the names of communities and the qid , which appears in all events, should be the same.

In the above sequence, c_i has forwarded a query to c_j and c_j returned a response.⁹ If this pattern is found to be prevalent, the weight of the relationship between c_i and c_j is increased.

On the other hand, Fig. 2(B) shows the pattern which is designed to capture situations where community c_i forward a query to c_j and c_j does not return anything (i.e., timeout). This may indicate that the given peer relationship is not responsive and does not contribute positively in query sessions. If this happens often, c_i may downgrade the peer relationship. We define the related PIS as follows.

⁷A query session is identified by the qid in the event.

⁸We refer readers to [3] for detailed description and issues related to frequency of a pattern

⁹Any number of other events may come between the two events. However, the order of events in a pattern is always respected.

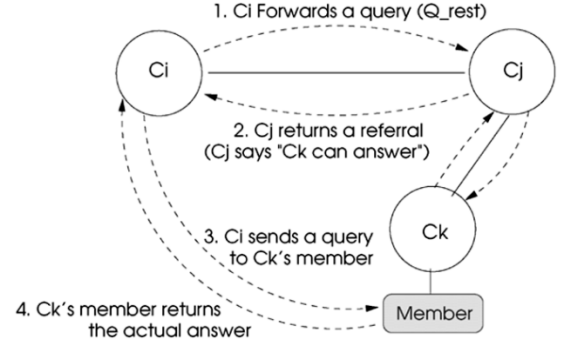


Fig. 3. Pattern to find a new member.

Definition 3: (PIS_{down}) (PIS downgrade) The pattern represents the situation where a forwarding of a query to a peer is always followed by a timeout

$$\text{PIS}_{\text{down}} = \langle \text{MQForwardedPeer}(c_i, c_j, \text{qid}), \text{Timeout}(c_i, c_j, \text{qid}) \rangle$$

where c_i, c_j are the names of communities and the qid , which appears in all events, should be the same. If this pattern is found to be prevalent, the weight of the peer relationship between c_i and c_j is decreased.

b) *Adding New Members:* Fig. 3 illustrates a situation where c_i forward Q_{rest} to peer c_j , and c_j forward it to c_k . c_k lets c_j know that a member of c_k can answer Q_{rest} . c_j relays it back to c_i . Then, c_i finally sends the query to the a member of c_k . That is, the figure summaries a situation where the part of queries that could not be answered by c_i frequently gets resolved by a member of c_k eventually. If this happens frequently, it may indicate that if the member of c_k also becomes a member of c_i , c_i would have not have to forward Q_{rest} anymore (i.e., it can be resolved locally).

Definition 4: ($\text{PIS}_{\text{newmem}}$) (PIS newMember) The pattern represents the situation where a community forward Q_{rest} to a peer who returns an identification of an external member that can answer Q_{rest} . The community sends the query to the external member and receives the results.

$$\text{PIS}_{\text{newmem}} = \langle \text{MQForwardInitiated}(c_i, c_j, \text{qid}), \text{MQAnswerReceived}(c_i, c_j, \text{qid}), \text{AQSentExtMember}(c_i, c_k.\text{member}, \text{qid}), \text{AQExtAnswerReceived}(c_i, c_k.\text{member}, \text{qid}) \rangle$$

where c_i, c_j, c_k are community names, $c_k.\text{member}$ denotes the external member belonging to c_k . qid which appears in all events should be the same.

If the above pattern occurs frequently, c_i might consider adding the member of c_k to itself.

Using the community monitor the administrator would analyze which query attributes frequently appear in Q_{rest} . The attributes that most frequently appear in Q_{rest} is identified as the "missing information" in the community. Once the missing information is identified, the administrator may search, through

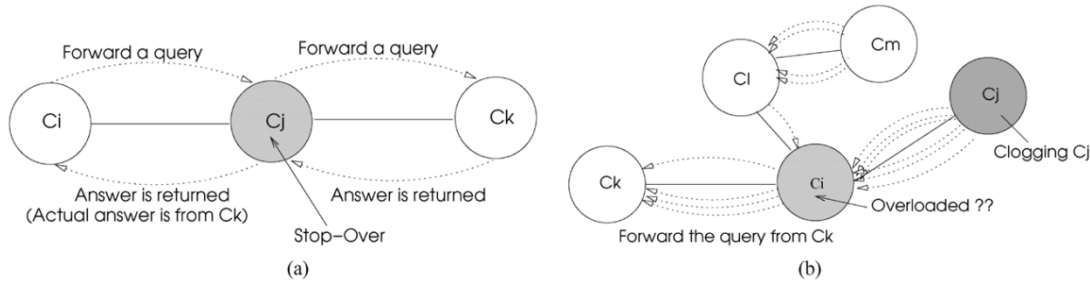


Fig. 4. Some global interaction patterns between communities. (a) Stopover (b) Clogging community.

the WS-CatalogNet's registry, for new members who are capable of supporting the attributes, so that queries can be resolved locally as much as possible.

2) *Global Scope Analysis*: The global scope analysis enables the WS-CatalogNet administrators to see the interaction between all communities involved in query sessions. We present a few examples here.

Finding Stop-Over Communities: Consider the forwarding scenario depicted in Fig. 4(a). The forwarding process that involves communities c_i , c_j , and c_k shows that the answer for c_i 's query is actually obtained from c_k via c_j . In this case, it can be said that c_j served as a stopover for c_i to reach c_k . If this pattern occurs frequently, c_i may decide to forward queries directly to c_k , bypassing c_j . To detect situations like this, we define a PIS as follows.

Definition 5: (PIS_{stopover}). (PIS Stopover) The pattern represents the situation where a forwarding process reveals a stopover peer

$$\begin{aligned} \text{PIS}_{\text{stopover}} = \langle & \text{MQForwardInitiated}(c_i, c_j, \text{qid}) \\ & \text{MQForwardedPeer}(c_j, c_k, \text{qid}) \\ & \text{MQAnswerReceived}(c_j, c_k, \text{qid}) \\ & \text{MQAnswerReceived}(c_i, c_j, \text{qid}) \rangle \end{aligned}$$

where c_i, c_j, c_k are community names and the qid , which appears in all events, should be the same.

c) *Finding Clogging Communities*: A community has a limited capability in the number of queries that it can process at a given time. In an environment like WS-CatalogNet, some communities can be overloaded with too many requests (i.e., query processing) from peers at a given time. If such overloading frequently happens to a community, the administrator of the community may consider performing some restructuring (e.g., reducing the number of incoming peers). Furthermore, it is worthwhile to investigate and identify who are the peers that are causing the community's overload (e.g., forwarding many requests to the community). Such peers are said to be *clogging* the community.

Consider a situation as illustrated in Fig. 4(b). Assume a community c_i and its peer c_j . Also, assume that c_j is clogging c_i by forwarding the most queries to c_i among other peers. Then, the event log shows that c_i subsequently forward the queries received from c_j to other peer (e.g., c_k). This means that c_i can not process the queries forwarded from c_j . Hence, by forwarding many queries to c_i , c_j is creating unnecessary workload for c_i .

If this happens often, the administrator may decide to disable the link between c_i and c_j (either temporarily, or permanently). Consider the following pattern (which is observed for a given period of time).

Definition 6: (PIS_{clog}). (PIS clog) The pattern represents the situation where a community receives a query from a peer and the query is subsequently forwarded to another peer

$$\begin{aligned} \text{PIS}_{\text{clog}} = \langle & \text{MQPeerQryReceived}(c_i, c_j, \text{time}_n, \text{qid}), \\ & \text{MQForwardedPeer}(c_i, c_k, \text{time}_m, \text{qid}) \rangle \end{aligned}$$

where $\text{time}_m, \text{time}_n$ represent the time-stamp of the event. $\text{time}_m, \text{time}_n \in [\delta, \gamma]$ where δ and γ are the period of time this pattern is considered for.

If this pattern occurs frequently and c_i has been overloaded often, the link between c_i and c_j may be disabled.

d) *Merging Two Communities*: The pattern PIS_{newmem} (Fig. 3) illustrates that c_i continuously finds that members of c_k is used to answer queries. Consider that the exactly same sequence of interactions also happens from c_k 's point of view. That is, c_k also finds continuously that members of c_i are used in resolving c_k 's Q_{rest} . If these interactions are mutual, it may be beneficial that c_i and c_k are merged and become one community. Fig. 5 depicts such scenario.

We have presented only a subset of possible patterns. There are other interaction patterns defined (e.g., patterns for monitoring member's quality, splitting a community, etc.) and the administrators may also defined and save ad hoc patterns for the future use. It is noted that the community monitor provides a user interface to the *PIS library* where patterns are managed (e.g., newly identified, created, updated, etc.)

C. Restructuring Operations

The community monitor is equipped with a set of restructuring operations which are performed on communities and their relationships. A few examples of such operations are: `addMember()`, `addPeer()`, `updatePeerRelation()`, `mergeCommunities()` etc. These operations are used, for example, to change the relationship between catalog communities (e.g., update the weight of a relationship), add a new peer to a community, or delete a peer link, etc.

The operations are divided into two classes: intracommunity and intercommunity. Intracommunity operations will be performed as a result of the local scope analysis. Intercommunity operations are performed as a result of the global scope analysis.

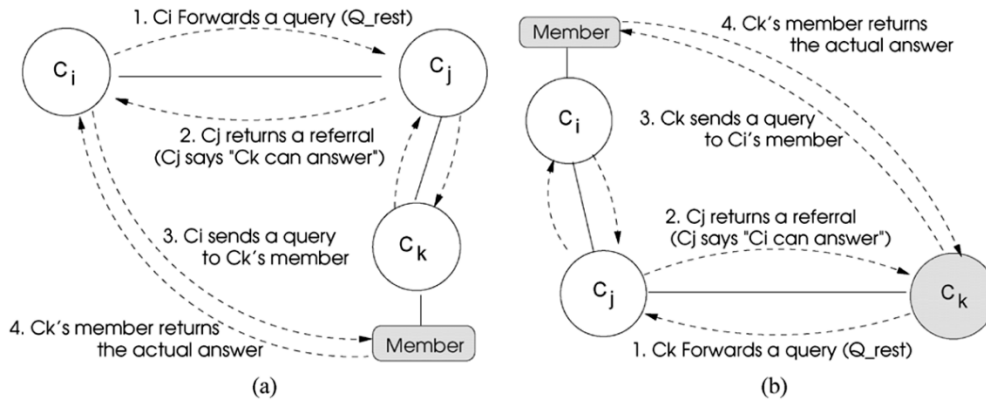


Fig. 5. A pattern for merging communities (a) Merge A (b) Merge B.

V. WS-CATALOGNET IMPLEMENTATION

In WS-CatalogNet, both catalogs and communities are presented as Web services.¹⁰ They can be described, advertised, and discovered using (XML-based) standard languages, and interact through standard Internet protocols. All services provide a SOAP-based programmatic interface for packaging requests and responses. These standards provide the building blocks for the service API descriptions and communication protocols interoperation, the two basic elements of any programmatic interaction. The universal description, discovery, and integration (UDDI) specification is used for directory functionality. We used the IBM Web Services Development Kit 5.0 (WSDK), which provides several components for developing Web services. In particular, we used the UDDI Java API (UDDI4J) to access a private UDDI registry (i.e., hosted by WS-CatalogNet), as well as the WSDL generation tool for creating the WSDL documents and SOAP service descriptors for the catalogs and communities.

In UDDI registry, every Web service is assigned to a tModel. A tModel provides a semantic classification of a service's functionality and a canonical description of its interface. We have designed specific tModels for the catalogs as well as for the communities. However, since the query mechanism supported by WS-CatalogNet is more sophisticated than what UDDI provides, UDDI is only used to advertise and locate the communities. WS-CatalogNet provides a generic community service, called *CatalogNetService*. This service has a number of prebuilt operations which provide the functionality necessary for supporting query processing and the registration. In particular, the following operations are provided to support the proposed query mechanisms: *QueryProcessor()*, *QueryRouter()*, and *ResultAssembler()*. The operation *QueryProcessor()* implements the BQR algorithm. The class *QueryRouter()* operation is used for routing queries based on the forwarding policies. The operation *ResultAssembler* is used for combining and selecting relevant catalogs.

The runtime operation of a community service is defined through the combination of *CatalogNetService* and

community metadata including: community ontology, member descriptions, P2P relationships, and forwarding policies. The community ontology is encoded as XML documents. The advantage of this architecture is that the developers who want to create a new community simply need to provide the necessary meta-data information, such as a community ontology and query forwarding policies. Tasks such as implementing the query rewriting and forwarding mechanisms are delegated to the generic service, thereby simplifying the development.

The Community Monitor implements the observation and restructuring technique. It integrates modules such as the event collector, PIS library, and the interface to the restructuring operations.

VI. EXPERIMENTS AND RESULTS

We conducted several experiments and the results are reported in the following.

A. Study on Restructuring the Network

Here, we looked at the effects of restructuring the community network. Out of the patterns presented in this paper, we chose two patterns: *PIS_{stopover}* and *PIS_{overloaded}* (a variation of *PIS_{cllog}*) for the study. Through the experiments, we would like to demonstrate that, after restructuring:

- 1) the average response time per query session *decreases*. Decrease in the response time would mean that restructuring helps reduce the time that takes for a query to be resolved.
- 2) the proportion of incomplete sessions *decreases*. The incomplete sessions are the ones in which answers were not returned. The reasons for an incomplete session are time out, having no peer to which forward a query, and exceeding the hop limit. Decrease in the proportion of incomplete sessions would mean that restructuring helps improve the chance of a query being resolved.

We have conducted our initial experiments in a simulated environment, in which we created communities as Java objects. Communities have their independent states and query capabilities. They autonomously invoke operations to process query and forward queries.

¹⁰The VLDB Journal: Special Issue on E-Services, 10(1), Springer-Verlag Berlin Heidelberg, 2001

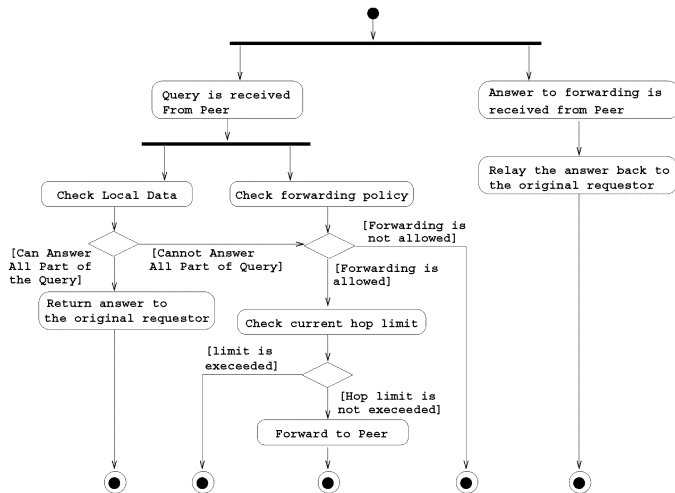


Fig. 6. Individual community's behavior.

1) *Setup*: Each community is a Java object which runs autonomously on an Intel Pentium IV 1.2-GHz machine with 512 MB of memory. A community is associated with an XML file which includes: a local data set, forwarding policy, and a list of its immediate peers. The local data set is a collection of metadata that describes community category definitions and member definitions. This is to represent the community's query capability. For the forwarding policy, each community specifies when to forward a query (e.g., only when query cannot be answered locally, or always) and how many "hops" are allowed. Queries used in the simulations are a set of keywords (randomly chosen from a pool of possible keywords). Each community has been assigned to a monitor to capture events such as receiving queries, forwarding queries etc. The monitor creates the entries to the log. The behavior of a community when processing a query is presented as a UML activity diagram in Fig. 6. Note the session will also end when the community has no peers (i.e., leaf node in the WS-CatalogNet graph). The basic steps of the experiments are as follows:

- 1) generate an instance of WS-CatalogNet.
- 2) run 10 000 query sessions.
- 3) collect the log and the measurement: *response time* and *number of incomplete sessions*.
- 4) from the log, see if the chosen pattern exists.
- 5) perform a restructuring operation according to what the pattern suggests.
- 6) with the "restructured" WS-CatalogNet, run the same 10 000 query sessions.
- 7) collect the measurements again and do a "before and after" comparison.

The program that launches each simulation takes three parameters: 1) *RLoc*, the location of the XML file; 2) *Qfile*, the file that contains keywords for querying; and 3) *MaxRun*, how many query sessions should be created. When the simulation starts each test run, it does the following.

- 1) It randomly picks, from *Qfile*, a query *Q* to be issued for the session.




Fig. 7. WS-CatalogNet topologies considered, (a) Hierarchy, (b) Centralized + Decentralized (c) Decentralized.

- 2) It randomly picks, from the XML file, the first community to which *Q* is submitted.
- 3) It repeats until *MaxRun* has reached.

Once a community receives a query from the simulation launcher, it autonomously starts processing the query according to the behavior defined as in Fig. 6.

We designed several different scenarios with the purpose of demonstrating the objectives we set out. In the first scenario, we tried the experiments on different P2P topologies as shown in Fig. 7. These topologies are the ones that we believe suit the WS-CatalogNet framework, since they are not centralized.

In the second scenario, we tried the experiments on different sizes of the network; 50, 100, 200, 400, and 800 communities. In the third scenario, with the decentralized topology, we studied the clogging pattern.

2) *First Scenario*: In this scenario, we focused on finding the stopover communities. That is, from the event log obtained from the simulation, we searched for the sequence of events specified in the $PIS_{stopover}$ pattern. For example, we looked for sequences of events like the following (here, a query was forwarded, from *c1* to *c2*, then from *c2* to *c6*):

```

⟨MQForwardedPeer(c1, c2), MQPeerQryReceived(c2, c1)
  MQForwardedPeer(c2, c6), MQPeerQryReceived(c6, c2)
  MQAnswerReceived(c2, c6), MQAnswerReceived(c1, c2)⟩.

```

To collect the "before-restructuring" data, we ran the simulation and obtained the global event log. Then, we searched the log for the $PIS_{stopover}$ pattern in various lengths (length here indicates the number of communities involved in the pattern). We wrote a Perl script for searching and counting of these patterns. The script also collected the average response time for each session, and the number of incomplete sessions. We identified the five most frequent $PIS_{stopover}$ patterns, which identified five new peer relationships between communities. We performed `createPeer()` operation, and ran the simulation again to collect the "after-restructuring" data. The result is presented in Table II.

The result showed that the restructuring performed improved both the *IncompS* and *RTime* measures in all topologies considered. However, the improvement was most significant in the decentralized model. This is an encouraging observation since the WS-CatalogNet framework assumes a decentralized model.

3) *Second Scenario*: In the second scenario, we tried the same experiments as the first scenario, but on different sizes of

TABLE II
RESULT OF THE FIRST SCENARIO

Hierarchy						
	IncompS		R.Time		Summary	
	Before	After	Before	After	IncompS	R.Time
Run #1	6016	4382	9	7	Down	Down
Run #2	5593	4491	10	8	by	by
Run #3	6106	3801	9	7	28%	21%
Hybrid (Centralized+Decentralized)						
	IncompS		R.Time		Summary	
	Before	After	Before	After	IncompS	R.Time
Run #1	5629	3492	9	8	Down	Down
Run #2	6162	3580	11	9	by	by
Run #3	7780	4842	11	9	39%	16%
Decentralized						
	IncompS		R.Time		Summary	
	Before	After	Before	After	IncompS	R.Time
Run #1	5121	1183	10	8	Down	Down
Run #2	5753	1630	11	9	by	by
Run #3	5017	1390	11	9	73%	18%

R.Time: Average Response Time (in millisecond)

IncompS: No of Incomplete Sessions (out of 10,000 sessions)




Fig. 8. Second scenario: different sizes of the network.

the network (i.e., different number of nodes in WS-CatalogNet). This time, we fixed the topology to a decentralized model. The result is shown in Fig. 8. The numbers shown are the average values of three test runs in each setting (i.e., three runs for the size-50 network, both before and after restructuring). The graph shows that regardless of the size, a significant improvement was made after the restructuring.

4) *Third Scenario:* In this scenario,¹¹ we focused on finding the most *overloaded* community and identifying the peers that are overloading the community. We searched for the sequence of events specified in the PIS_{Overload} pattern. For example, we looked for sequences of events like the following (here, *c1* has forwarded a query to *c2*, giving *c2* the work):

$\langle \text{MQForwardedPeer}(c1, c2), \text{MQPeerQryReceived}(c2, c1) \rangle$.

Once such communities, that is a) overloaded ones and b) the top *K* number of peers that made them overloaded, are identified, we tried two different restructuring operations. 1) Disable,

¹¹This scenario was tested on an Intel Pentium III 863-MHz machine, with 256 MB of memory.

TABLE III
RESULT OF THE THIRD SCENARIO

Disable Link					
IncompS		R.Time		A.Wrkload	
Before	After	Before	After	Before	After
712	437	550	300	5.4	2.7
IncompS 46%↓		R.Time 39%↓		A.Wrkload 50%↓	
Change Policy					
IncompS		R.Time		A.Wrkload	
Before	After	Before	After	Before	After
712	397	550	310	5.4	2.9
IncompS 44%↓		R.Time 45%↓		A.Wrkload 47%↓	

R.Time: Average Response Time (in millisecond)

IncompS: No of Incomplete Session (out of 5,000)

A.Wrkload: Average Workload ($\frac{\text{No of requests}}{\text{Max capacity of community}}$)

temporarily, the peer links between the overloaded community and the communities in (b). 2) Change the forwarding policy of the communities in (b) so that they forward less queries to others.

To measure the effectiveness of this restructuring, we collect extra measurement, *workload* beside the usual measurements. We defined the workload as the ratio of the number of requests a community receives over its maximum capacity. Each community's maximum capacity was allocated based on the average number of requests a community received in the past (e.g., from a few runs of simulations). The expectation was to see the average workload decreasing after restructuring. The result is shown in Table III. Again, the numbers presented are the average values of three test runs. The results indicate that both "disabling links" and "changing the policy" are effective ways of reducing the workload. The disabling link is slightly better in terms of reducing the number of incomplete sessions and the average workload.

B. Study on Various Network Parameters

In this study, we focused on investigating the effects of various network parameters on the query performance. Again, a decentralized model is used for the P2P topology. The query performance is measured by comparing, for each community, the total number of queries issued, the number of queries that are resolved (*Resolved*), the number of queries that are not resolved (*Unresolved*) and the average response time in *resolved* sessions (*Rtime*), the average response time in *unresolved* sessions (*Ftime*), in difference test cases.

All communities are implemented as Web services (a Java servlet running on Apache Tomcat 4.1.18/Apache AXIS 1.1). We used 30 machines,¹² each machine hosts a single community which is running as a Web service. The network topology information is dynamically generated and loaded to a central controller node¹³ before running each batch of experiments. Each community provides a number of SOAP services such as processing queries, forwarding queries, managing a forwarding policy, logging, etc.

¹²All Intel Pentium III, 850 Mhz, 256 RAM, operating on Debian Linux.

¹³The controller node is only responsible for initializing the P2P network of communities and is not involved in routing.

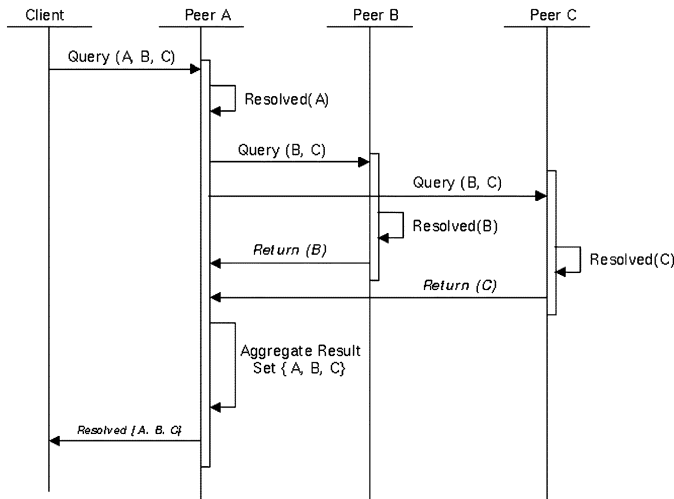


Fig. 9. Successful query process.

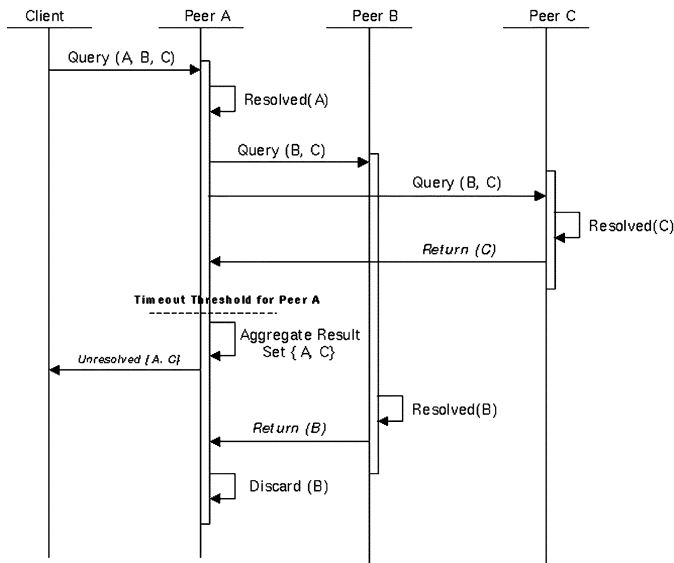


Fig. 10. Unsuccessful query process.

The forwarding policies of a community have two settings: MultiCast and UniCast. In the multicast setting, the community forward a Q_{rest} to all of its known peers, whereas in the unicast setting, a single peer is chosen at random for forwarding. The exact behavior of a community is depicted in Figs. 9 and 10. We used four sets of test cases, each containing 1000 randomized queries. To ensure that the randomly generated queries are not biased, the each test case is executed five times with a new pseudorandom generation sequence. At the end of every run, statistics are collected for each community to show the total number of queries issued, the number of queries that resolves, the number of unresolved queries, and response time. The final statistics gathered are the result of the average values from five runs. Out of all the test cases considered, we report the results on varying the hop count and the size of the network.

1) *Case One: Varying Hop Counts:* In this test case, we varied the hop count setting for each community, ranging from one to ten. The forward policy was set to UniCast and the test case included 1000 query runs. Note that the response




Fig. 11. Case one results.




Fig. 12. Case two results.

time is not considered here because it is clearly expected that the higher the hop counts, the longer the response time. The performance of a community is measured by Resolved (see Fig. 11). Note that the numbers in the figure are the average values of five runs in each setting.

The results have shown that with the increasing number of hops, there is a proportional increase in the number of resolved queries. However, after 5-hop and onwards the increase has slowed down and by the 8-hop setting, the number of resolved queries did not show much differences.

2) *Case Two: Varying the Size of the Network and the Number of Constraints in A Query:* In this test case, we studied the relationship between the community's performance and the size of the network and the number of constraints in the queries. The size of the network varied from 10 to 30 peers and the number of constraints in a query ranged from 10 to 50.¹⁴ From the result (see Fig. 12), it can be clearly seen that the 15-constraint setting seemed to result in the highest number of resolved queries (for all sizes of the network). The size of the network does not appear to affect the number of resolved queries. As shown in Fig. 12, there is a slight tendency

¹⁴The number of community category attributes in a query is used as the constraints.

for the performance to degrade as the number of the constraints or the number of peers increases, but this is definitely sublinear.

VII. FUTURE WORK

We have proposed a framework which supports integration and evolution of catalog portals using the P2P data management paradigm. Future work includes: 1) managing changes in community ontology (e.g., propagation of the changes to peers and its members); 2) automated discovery of potential community members and monitoring community-member interactions; and 3) further simulation studies on restructuring in a very large-scale systems over time.

ACKNOWLEDGMENT

The authors thank M. P. Domez and V. Liu for their invaluable contribution in the experiments.

REFERENCES

- [1] K. Baina, B. Benatallah, H. Paik, F. Toumani, C. Rey, A. Rutkowska, and H. Susanto, "WS-CatalogNet: An infrastructure for creating, peering, and querying e-Catalog communities," in *Proc. VLDB*, Toronto, Canada, Aug. 2004, pp. 1325–1328. demonstration paper.
- [2] B. Benatallah, M.-S. Hacid, H. Paik, C. Rey, and F. Toumani, "Toward semantic-driven, flexible and scalable framework for peering and querying e-Catalog communities," *Inform. Syst. J.*, 2005, to be published.
- [3] H. Paik, B. Benatallah, and R. Hamadi, "Dynamic restructuring of e-Catalog communities based on user interaction patterns," *WWW J.*, vol. 5, no. 4, pp. 325–366, 2002.
- [4] A. Bouguettaya, B. Benatallah, L. Hendra, M. Ouzzani, and J. Beard, "Supporting dynamic interactions among web-based information sources," *IEEE Trans. Knowledge Data Eng.*, vol. 12, no. 5, pp. 779–801, Sep./Oct. 2000.
- [5] B. Yang and H. Garcia-Molina, "Comparing hybrid peer-to-peer systems," in *Proc. VLDB*, Rome, Italy, Sep. 2001, pp. 561–570.
- [6] A. Crespo and H. Garcia-Molina, "Routing indexes for peer-to-peer systems," in *Proc. ICDCS*, Vienna, Austria, Jul. 2002, pp. 23–35.
- [7] S. Joseph and T. Hoshiai, "Decentralized meta-data strategies: Effective peer-to-peer search," *IEICE Trans. Commun.*, vol. E86-B, no. 6, pp. 1740–1753, Dec. 2003.
- [8] P. Bernstein, F. Giunchigiloa, A. Kementsietsidis, J. Mylopoulos, L. Serafini, and I. Zaihrayeu, "Data management for peer-to-peer computing: A vision," in *Proc. WebDB*, Madison, WI, Jun. 2002, pp. 89–94.
- [9] S. Bergamaschi, F. Guerra, and M. Vincini, "A peer-to-peer information system for the semantic web," in *Proc. AP2PC*, Jul. 2003, pp. 113–122.
- [10] D. Beneventano, S. Bergamaschi, F. Guerra, and M. Vincini, "The momis approach to information integration," in *Proc. ICEIS*, Setúbal, Portugal, Jul. 7–10, 2001, pp. 194–198.
- [11] S. S. Chawathe, H. Garcia-Molina, J. Hammer, K. Ireland, Y. Papakonstantinou, J. D. Ullman, and J. Widom, "The TSIMMIS project: Integration of heterogeneous information sources," in *Proc. IPSJ*, Tokyo, Japan, Oct. 1994, pp. 7–18.
- [12] D. Kossmann, "The state of the art in distributed query processing," *ACM Comput. Survey*, vol. 32, no. 4, pp. 422–469, 2000.
- [13] W. Ng, B. Ooi, K. Tan, and A. Zhou, "PeerDB: A P2P-based system for distributed data sharing," in *Proc. ICDE*, Bangalore, India, Mar. 2003, pp. 633–644.
- [14] A. Halevy, Z. Ives, D. Suciu, and I. Tatarinov, "Schema mediation in peer data management systems," in *Proc. ICDE*, Bangalore, India, Mar. 2003, pp. 505–518.
- [15] 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*, vol. 32, no. 3, pp. 53–58, 2003.
- [16] W. Nejdl, M. Wolpers, W. Siberski, C. Schmitz, M. Schlosser, I. Brunkhorst, and A. Lser, "Super-peer-based routing and clustering strategies for RDF-based peer-to-peer networks," in *Proc. WWW*, Budapest, Hungary, May 2003, pp. 536–543.
- [17] B. Yang and H. Garcia-Molina, "Designing a super-peer network," in *Proc. ICDE*, Bangalore, India, Mar. 2003, pp. 49–62.
- [18] B. Cooper and H. Garcia-Molina, "Ad Hoc, Self-Supervising Peer-to-Peer Search Networks," Database Research Group, Tech. Rep., Stanford Univ., Stanford, California, Tech. Rep., Feb. 2003.
- [19] A. Maedche, B. Motik, L. Stojanovic, R. Studer, and R. Volz, "An infrastructure for searching, reusing and evolving distributed ontologies," in *Proc. WWW*, Budapest, Hungary, May 2003, pp. 439–448.
- [20] B. Benatallah, M.-S. Hacid, H.-Y. Paik, C. Rey, and F. Toumani, "Peering and Querying e-Catalog Communities," Tech. Rep. UNSW-CSE-TR-0319, Comput. Sci. Eng., Univ. New South Wales, Sydney, Australia, ftp://ftp.cse.unsw.edu.au/pub/doc/papers/UNSW/0319.pdf, 2003.
- [21] *The Description Logic Handbook. Theory, Implementation and Applications*, Cambridge Univ. Press, Cambridge, U.K., 2003.



Hye-Young (Helen) Paik received the Ph.D. degree in computer science from the University of New South Wales, Sydney, Australia.

She is a Lecturer at the School of Information Systems, Queensland University of Technology, Brisbane, Australia. Her research interests include peer-to-peer data management systems and personalization issues in Web services.

Prof. Paik is a Member of IEEE Computer Society and the ACM.



Boualem Benatallah (M'01) received the Ph.D. degree in computer science from Grenoble University, IMAG, Grenoble, France, in 1996.

He is a Senior Lecturer at the University of New South Wales, Sydney, Australia. He has published widely in international journals and conferences including the IEEE TRANSACTIONS ON KNOWLEDGE AND DATA ENGINEERING, the *Very Large Data Base* journal, the International Conference on Data Engineering, and the ACM WWW. His research interests include Web services, semantics Web, and

autonomous data sharing.

Prof. Benatallah is a Member of the ACM.



Farouk Toumani is a Senior Lecturer at the School of Engineering in Computer Science, Modeling and Applications (ISIMA), University Blaise Pascal, Clermont Ferrand, France.

He has published widely in international journals and conferences such as the International Conference on Data Engineering, *IEEE Internet Computing*, the *Very Large Data Base*, and the *Information Systems Journal*. His research interests include Web services, semantics Web, and knowledge representation for databases.