

Schema-Assisted Peer Selection for XML Querying in Unstructured P2P Systems

Christos Doulkeridis¹, Kjetil Nørvåg², Michalis Vazirgiannis¹

¹Department of Informatics, Athens University of Economics and Business, Greece

²Department of Computer Science, NTNU, Trondheim, Norway
{cdoulk,mvazirg}@aueb.gr, Kjetil.Norvag@idi.ntnu.no

ABSTRACT

XML is emerging as the de-facto standard for semistructured contents and metadata. Searching this content in mobile environments is challenging, since centralized approaches are not appropriate in a very dynamic environment with limited resources available for keeping a centralized index up-to-date. A more appropriate solution is to organize the mobile devices in an unstructured peer-to-peer (P2P) network. The main challenge in the context of unstructured P2P is to determine the peers that might store documents matching a query, i.e., *peer selection*. In this paper, we propose a summary caching method for increasing the efficiency and recall of peer selection during XML querying. Our approach is based on caching parts of XML schemas along the query path, to enable subsequent *jumps* to remote peers storing content relevant to the query. We evaluate the performance improvements of our search strategy in terms of completeness of the search and reduced latency. The results show that our approach can significantly enhance a naive query mechanism such as flooding, and consistently outperform a baseline path caching technique similar to techniques used in related work.

1. INTRODUCTION

XML is emerging as the de-facto standard for semistructured contents and metadata. Searching this content in mobile environments is challenging, since centralized approaches are not appropriate in a very dynamic environment where devices have limited resources available for keeping a centralized index up-to-date. A more appropriate solution is to organize the mobile devices in a peer-to-peer (P2P) network. In this way, devices will cooperate in the search, thus avoiding the need for the central index, as well as providing an up-to-date view of contents in the network.

Thus, assuming that XML data is available on the devices and that these devices are organized in a loosely connected P2P topology, the challenge is to enable facilities for searching and querying XML data in a P2P context. In a system with high churn rate, constantly changing topology (due to the movement of the mobile devices), and high latency, maintaining a structured P2P overlay deployed as a distributed hash table is very costly, so instead we

base our approach on an unstructured (Gnutella-like) P2P network. The basic query routing mechanism in unstructured P2P networks is *flooding*. However, this technique is not scalable, and in practice, given a specified search budget in terms of number of messages, it limits the search to only a neighborhood of the querying peer. Several techniques have been proposed for reducing the search cost, and one particular class of techniques is those based on summary caching. A common approach that is widely adopted in previous work when querying XML data over a P2P network is caching XML paths on peers (henceforth referred to as *PCache*).

In this paper we present a novel *summary caching* approach of XML schemas for efficient peer selection and querying of XML data. Our new approach to summary caching is particularly useful in the context of XML data, aiming to improve *peer selection* at query time. Our approach is based on *caching (parts of) schemas of XML documents stored at remote nodes*. Schema information is cached at all peers on the query return path, in order to enable subsequent *jumps* to remote peers storing content relevant to a query. Assuming a query based on a path expression (which is the case for both XPath and XQuery), the schemas can be used for determining peers with high probability of matching the query. Thus the query is more directed than flooding, as it can be forwarded to peers that are more likely to contain relevant data. The result is that given a certain search cost, recall¹ is significantly improved in comparison to the basic routing mechanism.

Compared to previous work, our innovation lies in exploiting hierarchical schema information to cache representative summaries of peers' contents based on the query workload. This is an important difference compared to recent work on P2P keyword-based caching, since we devise caching techniques that inherently exploit the hierarchical structure of data. Moreover, contrary to previous work [2, 6, 8], we cache the structure of data, not data itself. Furthermore, peer selection is also based on the structure of data. This has two important advantages: 1) the cache is more resilient to data updates, and 2) the peer selection constitutes the first step towards schema mediation, as jumps are enabled to peers with relevant schemas. Then query routing efficiency increases, by directing the query to peers that can return results with higher probability.

The main contributions of our work are:

- A summary caching technique that uses (parts of) cached XML schemas to improve peer selection and query routing in unstructured P2P networks. This contrasts to previous approaches caching values, and has important benefits including robustness to dynamic data (schemas change less frequently than the actual data).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright 200X ACM X-XXXXX-XX-X/XX/XX ...\$5.00.

¹Recall is the fraction of all relevant material that is returned by a search.

- An algorithm that encompasses different routing variants, enabling *jumps* to topologically remote peers, based on the local cache contents.
- An extensive evaluation of the approach, through large-scale simulations. The results show that our approach: 1) reduces query processing cost and 2) increases the probability of finding the relevant XML data in a P2P network where we can not afford to search all peers. More importantly, our approach is consistently better than a baseline path caching technique, widely used in related work.

The organization of the rest of this paper is as follows. In Section 2, we overview related work. In Section 3, we define the system architecture, the data and query model. In Section 4, we describe schema-assisted P2P querying of XML data. In Section 5, we provide an evaluation of our approach based on a simulator prototype of a P2P system. Finally, in Section 6, we conclude and outline issues for further work.

2. RELATED WORK

An extensive review on P2P management of XML data can be found in [10]. One of the research challenges identified in this context is indexing XML data and subsequent efficient query routing. This is the issue addressed by our work. While several approaches in P2P XML data management assume schema-less XML documents [9], this is expected to change. So our approach capitalizes on the use of schemas. The lack of schema might have been good enough for previous document-centric XML documents, but will probably not be for data-centric documents. A schema is a very valuable resource in query processing and should always be used when available. Further, the existence of automatic schema generation techniques allows management of schema-less XML data by our approach too, without requirements of explicit schema definition by a peer.

It should be noted that result caching is not considered in our work as an appropriate alternative solution, mainly due to the high associated storage and maintenance cost. In particular, updates of actual results can lead to excessive bandwidth consumption, depending on the size of the cached objects. Therefore we mainly consider index caching mechanisms, an approach similar to the one adopted in [21]. This protocol (DiCAS) organizes peers into disjoint groups and selectively caches index data on the query return path, to improve some limitations of uniform index caching employed by Gnutella. Compared to DiCAS, our work exploits the hierarchical structure of XML data, more precisely its schema, to cache summary information along the query return path and to subsequently support *jumps* to remote peers holding data relevant to the query. Our approach can be enhanced by this protocol, an issue that we will investigate in future work, to harness the merits of both approaches. Another approach that employs query-driven caching, in which similarly to our work the information cached is generated by the query load, is presented in [18].

Furthermore, in most P2P systems queries are keyword-based, thus supporting only exact matching, that usually does not accommodate the user's needs or requirements. It is therefore important to use more expressive and powerful languages (like XML) that go beyond keyword matching, exploiting the (hierarchical) structure of the data [10].

In the following, we first review related work in XML query processing in unstructured P2P systems, which is more tightly connected to our approach, and we then proceed with an overview of improvements to basic search in unstructured P2P networks.

2.1 XML query processing in unstructured P2P systems

Among the most closely related research papers to our work are [2, 6, 8]. In [6], it is assumed that the XML tags provided by the participating nodes have-system wide defined semantics. This is practically a text indexing approach, which does not exploit the hierarchical information. The system does not try to compensate for semantic heterogeneity, which is considered an orthogonal problem. Whenever a new node joins the system it has to propagate information to every node in the network, which clearly is not scalable. Incremental updates are used to send only changes to data and in the approach mainly stable peers are assumed.

A different family of approaches is based on the concept of routing indices [4] and their variants. In [8], the authors propose strategies for routing and query processing in unstructured P2P systems. The basic mechanism adopted is a variant of the compound routing index, following the query shipping approach, i.e., no caching is assumed. In the same context, in [2], schema-aware routing indices are used to guide queries following a super-peer based approach. Routing of XML queries in P2P databases is researched in [11]. The authors try to propose a scalable solution with respect to both query processing and data updates. They also discuss the infeasibility of flooding or global index maintenance in a dynamic P2P environment, which motivates our approach.

Koloniari and Pitoura [9] present an approach for summary indexing using multi-level Bloom filters. In their work, they consider hierarchical content-based organization of peers, to subsequently route queries efficiently based on the summaries. In [17], a super-peer approach towards a P2P XML database system is presented, where peers are organized in a hierarchy, too. Different data schemas are supported and peers are clustered together based on schema-similarity, to reduce the querying cost. Query routing is based on query propagation up to the root of the hierarchy.

In previous work [5, 14], we have studied the issue of summary caching for improved query processing in unstructured P2P systems. In [14], taxonomy caching is employed for data indexed by a taxonomy. In [5], we presented the notion of schema caching for XML data. In this paper we extend the work presented in [5, 14] with a more elaborate caching technique and algorithms for peer selection as well as a more detailed experimental evaluation.

2.2 Search in unstructured P2P systems

Several papers describe improvements [22] to the basic flooding strategy using *query jumps* (direct contact with remote peers known or believed to contain relevant information) to remote peers [13, 19]. In the approach described in [13] the object location is also stored in the return path and query jumps can be employed. However, only one answer is assumed for a given search key. In our case, we can have many results and provide mechanisms to find the most appropriate query matches. Sripanidkulchai et al. [19] also present an approach where they use *shortcuts* (direct links to remote peers) which can be used for jumps. Shortcuts are made based on successful previous queries and are not associated to the actual queries (i.e., a shortcut with a particular ranking is used for all queries). This is different from our approach, where we relate summary information to actual information contents. In [3], a dynamic network adaptation mechanism is used, which is in a sense similar to our approach. Other papers that try to improve the basic search mechanism include the approach described in [12] and APS [20]. Finally, in [7], an approach that utilizes the past behavior of peers in order to boost the search performance is presented.

3. BACKGROUND

In this section we give an overview of the system and the data- and query models.

3.1 System overview

We consider a system of peers connected in an unstructured (Gnutella-like) P2P network. A peer P that joins the P2P network first establishes connection to one or more peers, as part of the basic P2P bootstrapping protocol (the actual protocol depends on the variant of unstructured P2P network, possible techniques include use of *known peers* as well as multicasting). Thus initially, peers are only aware of their N_n immediate neighbors, where N_n is determined as a function of the basic P2P network creation and peer-join strategies.

When no performance-improving technique like caching is employed, querying in an unstructured P2P system is performed as follows: the query Q originating from the querying peer P_Q is forwarded to other peers, denoted *remote peers*. The process of deciding to which peer(s) to forward the query is called *query routing* (a basic query routing algorithm flooding). Attached to the query is a time-to-live (TTL) value which is initialized by the querying peer. The TTL is decremented each time the query is forwarded, and when it reaches zero the query is not propagated further. The peers that can be reached at a particular time from P_Q constitute the *query horizon* of P_Q .

3.2 Data model

The peers in the system store data that is searchable by other peers in the network. This data is represented as XML documents that are stored either as files or in a database, the only requirement is that it should be possible to query these documents as will be described shortly. Although documents stored as files (and possible also XML documents stored in databases) might have a filename, in a system-wide context the file names are not of interest as our approach is data-centric². XML documents may or may not have an associated description using XML Schema. We expect that for new data-centric applications based on XML the use of schemas will be the rule (or can be automatically generated). An important aspect of schemas is that *given the schemas of the XML documents in a collection, it is possible to determine from the schemas whether there can be documents in the collection matching a particular path expression in a query*. In other words, the XML schemas are used to select which collections to query, and prune collections that certainly contain no documents matching the query.

In a P2P context all peers are autonomous. This means that there is no global schema or authority that can control or verify schemas that are used. Thus we expect that on each peer a number of schemas are used. In the general case schemas on different peers will not be related, but in many cases peers will use standardized or at least commonly accepted schemas, which is necessary in the case of communication and e-Business. Our approach specifically targets at such focused application scenarios. It should also be mentioned that although our approach is most useful when XML schemas are used, we expect that a certain fraction of documents will still be created without an associated schema. In order to improve query efficiency also for queries involving these documents, a schema can be created that is compliant with the current instance of the document.

²Data-centric XML documents are typically documents meant for computer consumption, while document-centric XML documents are typically meant for human consumption (like books, papers, etc.)

3.3 Query model

There exist a number of approaches to query XML documents. This includes standardized languages like XPath and XQuery, as well as vendor-dependent languages and SQL-extensions (which has been the typical approach in commercial relational DBMSs). Most of these approaches have in common that they employ a variant or subset of XPath, in order to filter out relevant elements from the XML documents before further processing. Our system-wide query model is also based on XPath.

The most important aspect of XPath is the notion of *location paths* (LP). Example of location paths are `/person/address/city` and `//address/city`. It is interesting to note about location paths that a path P_1 that is a prefix of a path P_2 will match a superset of the elements matched by P_2 . For instance, the number of elements matched by `/a/b` will be the same or larger than those matched by the more specific path `/a/b/c`. This is exploited by our approach, which caches information about the more specific path, in order to improve the efficiency of the cached entries.

Example: `/person/address` will match all `person` elements containing an `address` element, while `/person/address/city` will only match those `person` elements that contain an `address` element containing a `city` element. A location path can also contain predicates, which can be tests on strings as well as on numerical values.

In general, the full XPath language or XQuery will be used at a querying peer P_Q . However, given a query issued by P_Q , only the location path, possibly including predicates, is actually forwarded to other peers. When a query is issued, the following steps take place:

1. The local query processor extracts the location path LP from the query.
2. LP will be forwarded to appropriate peers (as will be described in more detail in the following sections).
3. When a LP is received at a peer, it will be applied to XML documents having a schema matching the LP . The result of applying the LP is a set of elements E (i.e., XML data). The elements E are returned to P_Q .
4. P_Q will receive a set E for each peer having matching documents. In the case of XQuery additional manipulation might be performed, in order to generate the final results.

In the description above we have assumed a *data-centric* application context where all matching data of contacted peers will be retrieved and be part of the process of generating the result. However, we note that in *document-centric* application areas, only a subset of the matching documents will be needed and retrieved. Which documents to retrieve is based on a ranking decision. For these application areas the identifiers and relevant metadata of the documents containing matching elements are returned, instead of the actual elements. In document-centric application areas also the whole documents and not only individual elements might be needed in the final result. Our proposed techniques are equally applicable for both data- and document-centric application areas, but for simplicity of presentation we will assume a data-centric context in the rest of this paper.

4. SCHEMA-ASSISTED P2P QUERYING OF XML DATA

The problem of using the basic approach for P2P querying, as described above, is the high cost associated with flooding. The

result is that in order to find as many results as possible, a very high number of peers *potentially* having relevant data have to be contacted. In order to have a more targeted query forwarding, we employ summary information retrieved during past queries to better decide on candidate peers.

4.1 Result caching vs. summary caching

One particular kind of information from previous queries is the actual results, i.e., after receiving the *query results* from a number of peers, the querying peer can cache the query results for more efficient processing of future queries as described in, e.g., [1, 16]. An expiration time (presumably with a low value) is attached to the result, to ensure that the results will be discarded after a certain amount of time. In particular for static contents and stable peers, the result caching technique can significantly reduce the search cost. However, 1) the results might soon get invalid/stale, caching the results incurs a considerable storage cost and thus the results are only kept for a certain amount of time, limiting the usefulness of the caching, 2) the stored results are only useful when the *exact* same query is issued several times, 3) result caching in itself will not improve the probability of finding contents outside the query horizon (although this can be improved on, by allowing query results to be used also for queries from remote nodes, i.e., not only for local query processing), and 4) the associated storage and maintenance cost of result caching is high. Even though secondary storage media is getting cheaper, result caching in a dynamic P2P system is not a wise design decision, since the system becomes vulnerable in terms of bandwidth consumption, depending on the rate of updates and the size of updated objects.

An alternative to result caching is to instead cache *summary information of contents located at remote peers*. This summary information can be returned together with the query results. Ideally, such summary information should be compact and robust to changes of the actual contents. In the following, we will describe our XSCache approach for summary caching.

4.2 XSCache overview

Our approach, which is the topic of this paper, is to cache (parts of) the *schema* of remote data (instead of, for example, caching summary based on actual contents), as illustrated in Figure 1. Notice that the actual values from query results are not cached. In the case of XML, a query will involve a path as described in the previous section, and information about possible paths can be found in the schema. Thus having the schema cached locally can be used to know which peers most likely have data that matches the query. In the most extreme case (although not likely in practice), when all schemas from all peers would be available, the exact set of peers matching a query can be found before query forwarding.

The advantage of our approach is that even if contents of an XML document change frequently, schemas change less frequently. It is also the case that usually when a new document is created, it will belong to an existing schema. Although in the worst case a new document will have a new schema associated with it, this will be the case only for very few documents on a peer, as there exists no real application where each document has a separate schema.

It should also be noted that our approach can be employed in combination with result caching. This is in particular useful for very frequent queries, which then only have to be reissued at regular times. When our approach is used for query routing, the query horizon is also extended in a more robust way than when using result caching alone. It will also gradually be extended with time as schema information is further distributed.

We will now describe in more detail our new approach for query

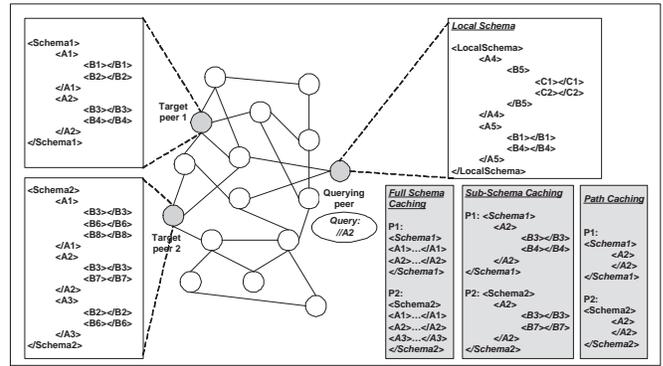


Figure 1: Schema caching variants (*full-schema, sub-schema and path caching*) as result of a query for //A2.

routing. The main contributor to increased query routing efficiency is the XML schema cache (XSCache), a structure for maintenance of remote schema information.

4.3 Schema-assisted query routing

Query routing is the process of deciding where to forward a query, performed by all involved peers, i.e., both the querying peer as well as intermediate peers. In both cases, when a lookup in the XSCache for schemas matching the XPath expression gives as result a set of matched peers, the query can be forwarded directly to one or more of these peers. This forwarding is called a *jump*. Note that unlike routing indices that only maintain information of the neighborhood and are used to choose appropriate neighbor peers for query forwarding, information about contents at very remote peers (also beyond the query horizon) can be contained in the XSCache. In the case when a peer does not find a match in the XSCache, the query is forwarded using the basic query routing algorithm. Notice that even if a query match is found at a peer and results returned to the querying peer, the query is further forwarded until the TTL reaches zero. This is because ideally in a data-centric approach, *all* relevant results to a query need to be retrieved. For an illustrative example of query routing assisted by jumps see Figure 2, where grey-colored peers are the ones contacted by the query. The figure explains how the query horizon is extended; because of the jumps and the continuation of routing at remote locations, contents of remote XSCaches can be queried and potentially result in new jumps.

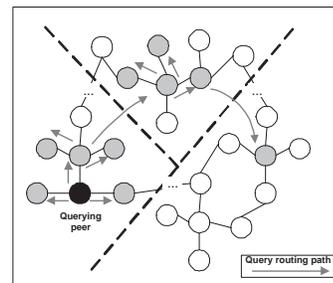


Figure 2: Query routing employing jumps (dashed lines show remote parts of the P2P network, while grey-colored peers are the ones contacted by the query).

Algorithm 1 Peer selection during query routing.

Require:

$P_N = \{P_1, \dots, P_n\}$ {Neighbors to consider, $n = N_n - 1$ }
 $XSCRV \in \{XSCacheB, XSCacheNB, XSCacheN2B, XSCacheAll\}$ {XSCache routing variant}
 $P_C = \{P_{C1}, \dots, P_{C_c}\}$ {Peers of the c matching entries in XSCache, decreasing rank}

Ensure:

P_F {Peers to which the query will be forwarded}
 $P_F = \emptyset$
if ($c < n$) **then**
 $P_R = (n - c)$ randomly selected peers from P_N
end if
if ($XSCRV = XSCacheNB$) **then**
 if ($c \geq n$) **then**
 $P_F = \{P_{C1}, \dots, P_{C_n}\}$
 else
 $P_F = P_C \cup P_R$
 end if
else if ($XSCRV = XSCacheB$) **then**
 $P_F = P_{C1} \cup P_N$
else if ($XSCRV = XSCacheAll$) **then**
 if ($c \geq n$) **then**
 $P_F = P_C$
 else
 $P_F = P_C \cup P_R$
 end if
else if ($XSCRV = XSCacheN2B$) **then**
 if ($c \geq n$) **then**
 $h = n/2$
 $P_F = \{P_{C1}, \dots, P_{C_h}\} \cup$ random selection of h other peers in P_C
 else
 $P_F = P_C \cup P_R$
 end if
end if

In general, a lookup in the XSCache returns c peers having a schema matching the query, and the query is forwarded to k of these peers and in addition to a subset of the neighbors. When $k < c$ a decision has to be made to which of the peers to forward the query. Our current approach is to simply rank the candidate peers based on the number of XML document elements they contain that match the cached schema (note that this approach can easily be extended to for example consider position in the hierarchy etc.). Those that have the highest number are considered most useful for the query and are more appropriate to answer the query. The value of k is an important parameter, and different strategies for determining k results into a number of *XSCache routing variants*. Assuming N_n to be the average connectivity degree (and let $n = N_n - 1$) for the P2P network then the query is routed to the k highest ranked peers in the cache and to $n - k$ randomly selected neighbors:

1. *XSCacheB*: $k=1$, select the highest ranked peer from the cache.
2. *XSCacheNB*: $k=n$, select the n highest ranked peers from the cache.
3. *XSCacheN2B*: $k=n/2$, select the $n/2$ highest ranked peers from the cache.
4. *XSCacheAll*: $k=c$, select all peers from the cache.

The approaches are presented formally in Algorithm 1, where the set of neighbors to consider is denoted as $P_N = \{P_1, \dots, P_n\}$

(where $n = N_n - 1$) and the set of the peers of the c matching entries in XSCache, ordered by decreasing rank, as $P_C = \{P_{C1}, \dots, P_{C_c}\}$.

When forwarding the query there is also the issue on how much the TTL should be decremented. We reduce TTL by 1 during jumps, so as to enable the continuation of flooding at remote locations. This means that given a particular TTL the total flooding cost is in the same order whether jumps are performed or not. Nevertheless we emphasize that because of the jumps, more peers will be contacted, since fewer messages are wasted by reaching already contacted peers. We study the best XSCache variants (XSCacheB and XSCacheNB) in terms of performance experimentally in Section 5.

4.4 Distributing schema information

The basic mechanism for distributing schema information is by piggybacking the appropriate schema(s) with the result of the query. The query results are routed to the querying peer through the return path, possibly involving jumps. The reasons for this are: 1) to enable caching the schema(s) at peers on the return path, and 2) to make it possible for these intermediate nodes to validate or invalidate their cached routing information. In particular jumps are performed because the peer from which the jump is initiated has information about the destination peer containing data related to the query. This information is kept in the XSCache and has a validity time. By routing the information back, the peer reinitializes the time counter, when it knows the entry is valid. It is also possible to return an invalidation message indicating that the destination peer does not have any data matching the query anymore, if that is the case.

4.5 XML schema caching

The XSCache is a fairly traditional cache. When a new schema is to be inserted but the cache is full, one of the schemas is removed based on an LRU policy. Because of high cost of invalidation we use the soft-state approach which is most common in highly distributed systems, i.e., each item has an associated expiration time, so that if the validity of the schema with respect to the associated peer has not been confirmed within a certain time, the schema will be removed from the cache. The expiration time is set to a fixed value, which is higher for remote peers that are already known and where churn statistics are available, than for newly discovered peers.

When caching a schema based on query results, we have the option of caching the whole schema or just the part of the schema relevant to the current query. The first approach is called *full-schema caching* and it is considered an aggressive caching technique, while the second is called *sub-schema caching*. For completeness we also mention that an even coarser caching can be performed: *path caching*. In a path caching approach only the actual path in the query is cached, i.e., not the subpaths being part of the schema.

The three alternatives are illustrated in Figure 1. The figure shows the contents of the XSCache after a query for //A2. The grey-shadowed areas of the figure show the contents of XSCache in case of: 1) full-schema caching (the whole schema is cached), 2) sub-schema caching (only the part of the schema that contains //A2 is cached), or 3) path caching (only the path ending with //A2 is cached).

Which of the three approaches (path caching, full-schema caching, and sub-schema caching) to use depends on applications and query patterns. Path caching is the most conservative approach, however it results in savings in communication and reduces associated storage and maintenance costs. Full-schema caching has the advantage

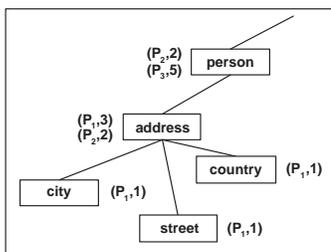


Figure 3: Tree representation of (parts of) schemas cached from different peers.

of supporting a larger number of different future queries. However, the added communication cost and memory usage can be significant, so in the general case when no assumptions can be made regarding schema sizes and access patterns, sub-schema caching is a safer approach, thus it will be the context of the following discussion.

The items in the cache, the (sub-)schemas, are represented as a tree. In order to ensure efficient access and use of memory, schemas from different peers are merged when possible. This is illustrated in Figure 3. Each node in the tree is annotated with a set of tuples (P, N) , where P is a peer identifier and N is the number of leaf elements matching the path from the root to this node in the schema tree. The value of N is used in the ranking of peers as described above. The part of the tree illustrated in Figure 3 can be the result of a number of queries, including $/person$, $/person/address$, and $//address$. However, a query $/person/address/street$ would not generate the tree depicted in Figure 3. The annotated (P, N) tuples are based on information provided together with the schemas from the remote peers.

One important aspect to note about the XSCache is that the values resulting from a query, for example the text string “Athens” in the city element returned from peers, is not cached in the XSCache. The reasons are that caching all values would result in high storage cost and more importantly high maintenance cost in the case of data updates. Also, caching the values on all peers along the return path (see above) will in general not be beneficial. However, in applications areas where caching the values *is* beneficial, result caching can be applied in addition to schema caching, most typically on the querying peer only.

5. EXPERIMENTAL STUDY

In this section, we first present the experimental setup and we then proceed to describe and discuss the results of our experiments. An unstructured P2P network is considered with N_P peers that store XML data that conform to a variety/number of schemas. Some overlap among the schemas that belong to different peers is assumed. A fraction of peers (q_p) act as querying peers issuing randomly generated queries to the network to find relevant XML data. Several routing strategies are tested to study their comparative performance.

Our approach is a generalization of path caching, or in other words path caching is a special case of our approach. Therefore, instead of comparing against all other approaches, we study PCache in a comparative way to our approach in an unstructured P2P context.

5.1 Experimental setup

We evaluate the performance of the proposed approach through

Symbol	Description	Range of values (Default)
N_P	# of peers	2000...8000 (2000)
N_S	# of schemas	100...2000 (500)
N_D	# of data items	100000...400000 ($50 \times N_P$)
N_Q	# of queries	- (100)
q_p	Fract. of querying peers	0.2...0.5 (0.2)
a	Query distr. param.r	0.6...2 (1.4)
C_S	Cache size	- (50)

Table 1: Symbols, descriptions and default values used in the experiments.

simulations, in order to test its scalability and feasibility. We use a simulator, created by our group and written in Java, for studying the performance of different routing strategies in unstructured P2P networks.

In the experiments, two different P2P network topologies are used: 1) a grid-like topology of constant connectivity degree which is more dense, i.e., each pair of neighboring peers share 3-5 common neighbors, and 2) a random graph topology created with the GT-ITM topology generator³. Due to space constraints we only report results from the former, results from the latter can be found in the extended version of the paper⁴. We use different network sizes, up to $N_P=8000$ peers. Unless mentioned explicitly, we used the topology with a connectivity degree of 8. We also tried other connectivity degrees in our experiments leading to similar conclusions.

Peer-residing data is described by XML schemas that may have common parts. Also some peers may share a common schema, to resemble the real world case, for instance peers within the same organization. It should be mentioned that the decision to use synthetic generated schemas is due to the difficulty of finding a real XML dataset that entails a large number of schemas, as required for large-scale P2P experiments. The total number of available data values in the network is $N_D = 50 \times N_P$, where N_P denotes the number of peers. The allocation of data to peers follows the Pareto principle (also known as the 80-20 rule), i.e., 20% of the peers hold 80% of the total data. We also used uniform allocation to peers. For each data value, a path expression of the peer’s local XML schema is used to describe it.

At each simulation experiment a set of $q_p \times N_P$ ($q_p \leq 1$) peers is selected to act as querying peers. Each peer initiates N_Q queries, which are XPath expressions picked from the union of available XML schemas. In the experiments shown in this paper, we only use a simple type of XML queries, namely queries for location paths, i.e. no predicates are used. The queries are randomly generated using the zipfian distribution (with parameter a), in order to simulate the users’ interests. This is based on previous reports (e.g. see [15]), which state that file popularity on the Web follows a zipfian distribution.

Symbols of parameters, range of values, and default values are summarized in Table 1.

In our simulation study, we assess the performance of different routing strategies using as quality measures: 1) recall, which shows the completeness of the search (in contrast to file-sharing, database applications usually require finding all relevant peers, and retrieving all answers to a given query), 2) latency for: i) the first result, ii) the first 10 results, and iii) all results, in terms of number of hops

³Available at: <http://www-static.cc.gatech.edu/projects/gtitm/>

⁴Available at http://www.idi.ntnu.no/grupper/db/research/technical_report/2007/231_art8.pdf

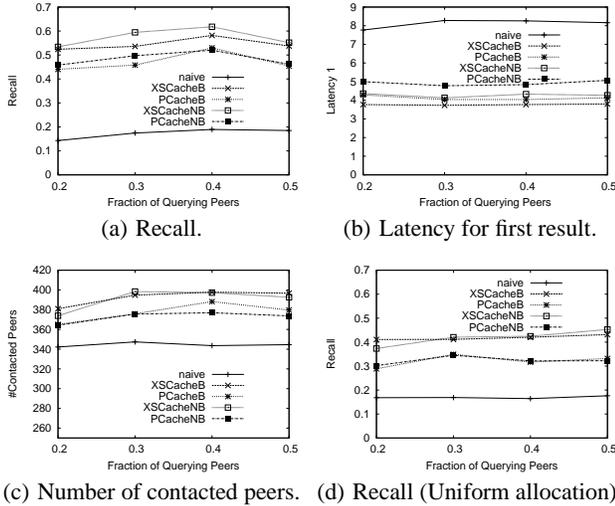


Figure 4: Measurements from using different number of querying peers, given as the fraction of peers in the network.

required to return the matching XML data to the querying peer, and 3) the number of contacted peers due to the routing strategy.

All measurements are taken after half of the total queries have been issued, in order to eliminate the warm-up effect of the cache. We study the effect of the following parameters on the performance of these strategies: 1) TTL values, 2) network sizes, 3) different topologies, 4) different skew in query distributions, 5) percentage of querying peers, and 6) total number of available schemas in the network.

In the experiments presented in this section, we present two variants of sub-schema caching (denoted *XSCacheB* and *XSCacheNB*), according to the way the cached schemas are utilized during query routing. We choose to show only two variants, in order to make the charts readable. Moreover, the results of the other XSCache variants do not offer any substantial additional insight. It should be reminded here that with each cached XML schema, also the number of associated data values is kept. Query routing variants are based on how to pick peers to forward the query. A peer is ranked according to the number of data it possesses.

We use flooding as the naive baseline approach. A more sophisticated baseline than *naive* is considered as well. This is achieved by adapting the basic technique used in practically the majority of related work (namely XML path caching) into our simulator, in order to compare against. We refer to this strategy as *PCache*. It is interesting to note that *PCache* is actually a special case of XSCache, namely it corresponds to the path caching variant. In full accordance with the two variants of XSCache, *PCache* is used in our experiments with two variants: *PCacheB* and *PCacheNB*. The variants of our approach are compared against these two approaches in the experiments.

It should also be stressed that we compare the performance of the aforementioned approaches *using the same number of messages* and the same setup parameters, in order to present directly comparable results. As can be seen, our approach both enhances the naive search strategy and outperforms *PCache* in all experiments.

5.2 Experimental results

In this section we study the effect of the most important parameters for a system employing the XSCache.

In Figure 4, we study the effect of increasing the number of

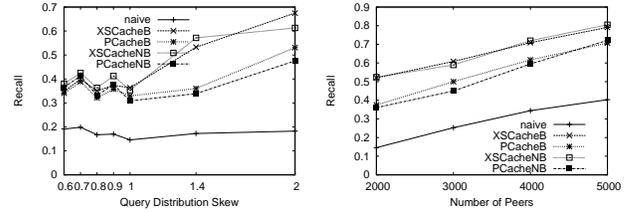


Figure 5: Recall from using different skew in the query distributions (left) and for different network sizes (right).

querying peers in the network. The number of querying peers is increased from 20% to 50% of the total peers. The results show that recall increases with the percentage of querying peers (Figure 4(a)). In a real P2P system all peers are expected to issue queries for data, so this is a strong argument in favor of the scalability of our approach with the number of querying peers. Notice also that our approach presents significantly lower latency for the first result compared to the naive approach (Figure 4(b)). XSCache also outperforms *PCache* in all other latency-related metrics (first 10 and all results). Another interesting observation comes from Figure 4(c), namely that XSCache manages to achieve 4 times the recall of flooding, by only contacting approximately 40 extra (out of 380) peers, using the same number of messages. This means that XSCache manages to find a higher number of peers that actually match the query. Furthermore, in the case of flooding, the number of contacted peers (given a certain number of messages) totally depends on the underlying topology. In contrast, XSCache is oblivious to topological parameters. In other words, XSCache overcomes some of the problems of a basic routing mechanism related to the structure of the underlying topology, such as the density of the network topology.

Finally, we tested uniform allocation of data to peers in Figure 4(d), which is exactly the same experiment as shown in Figure 4(a) except that the data allocation is uniform. The comparative results are practically the same with 80-20 allocation, only the absolute values are lower when the distribution is uniform. We use the 80-20 allocation rule in the rest of the experiments.

In Figure 5 (left), we gradually increase the skew in the query distribution by increasing the value of the query distribution parameter a , to study the behavior of our approach. As expected, when the query distribution is more skewed, recall increases, as a result of many recurring queries. This is due to the fact that the cached XML schemas become useful for more queries.

In Figure 5 (right), the scalability of our approach with respect to the number of peers is illustrated. We stress here that the increasing recall values with network size are not expected, yet this is explained by the fact that we also increase the TTL as we increase the network size. Unfortunately, the TTL does not increase analogously to the number of peers, rather it increases faster, hence recall increases with network size in the chart. The important finding of this experiment is that the higher recall achieved by XSCache relative to *PCache* (and naive) is sustained as the network size increases.

In Figure 6 (left), we study the effect of increasing TTL values on the performance of the search strategies. The effect of varying the number of available schemas in a P2P network of 2000 peers is depicted in Figure 6 (right). We use up to 2000 different schemas in this experiment, which is a reasonable value since even if data in a P2P system is some orders of magnitude larger than the number of peers, the number of schemas employed is definitely comparable

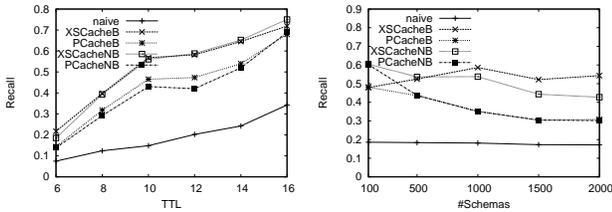


Figure 6: Recall when using different TTL values in the search (left) and with different number of available schemas in the network (right).

to the number of peers.

We see that for small number of schemas (100) our approach performs similar to PCache, but far better than naive. However as the available number of schemas increase, XSCache constantly outperforms PCache. This recall does not increase monotonically, as can be seen for 2000 schemas, where the achieved recall admittedly drops, however it still surpasses PCache. The reduced recall as the number of schemas increase is expected, as in the extreme case that each peer has a completely different schema from other peers, any schema caching mechanism will only benefit queries for this peer only. Moreover, when the number of available schemas decreases, as would be the case when schema mediation is employed, the efficiency of XSCache is not reduced significantly. All in all, with increasing number of schemas, our approach achieves 2-4 times the recall of naive, and 1.5-2.5 times the recall of PCache.

Concluding, our approach significantly outperforms the baseline search strategies in terms of recall and associated latency (for the first, the first 10 and for all results), while utilizing the same number of messages. In other words, given a certain recall, XSCache needs a lower number of messages to achieve it. Furthermore, our approach increases the number of peers that need to be contacted, with the same search cost.

Comparing the individual variants of XSCache (XSCacheB and XSCacheNB), we conclude that using the highest ranked entry in the cache performs similarly to the approach that uses n entries. As a rule of thumb, even one cached entry (the highest ranked) is enough to achieve high recall and low latency. This conveys that XSCacheB is a suitable approach when designing a search protocol for P2P networks similar to the ones tested in our simulations.

6. CONCLUSIONS AND FUTURE WORK

In this paper, we presented a new approach aiming at reducing XML query processing costs in mobile environments. We introduced the usage of XML schema caching for query routing as well as the underlying mechanisms. We also performed an extensive analysis of the approach through large-scale simulations. The results showed that our approach: 1) reduces query processing cost and 2) increases the probability of finding the relevant XML data in a P2P network where we can not afford to search all peers.

Plans for future work include more sophisticated ranking of candidate peers for peer selection, and integrating actual XML processors in an implementation to develop a P2P XML search engine.

7. REFERENCES

- [1] B. Bhattacharjee, S. Chawathe, V. Gopalakrishnan, P. Keleher, and B. Silaghi. Efficient Peer-to-Peer Searches Using Result-caching. In *Proceeding of IPTPS'03*, 2003.
- [2] I. Brunkhorst, H. Dhraief, A. Kemper, W. Nejdil, and C. Wiesner. Distributed queries and query optimization in

schema-based P2P-systems. In *Proceedings of DBISP2P'03*, 2003.

- [3] Y. Chawathe, S. Ratnasamy, L. Breslau, N. Lanham, and S. Shenker. Making Gnutella-like P2P systems scalable. In *Proceedings of SIGCOMM'03*, 2003.
- [4] A. Crespo and H. Garcia-Molina. Routing indices for peer-to-peer systems. In *Proceedings of ICDCS'02*, 2002.
- [5] C. Doulkeridis, K. Nørnvåg, and M. Vazirgiannis. Schema caching for improved XML query processing in P2P systems. In *Proceedings of IEEE P2P'06*, 2006.
- [6] L. Galanis, Y. Wang, S. Jeffery, and D. DeWitt. Processing queries in a large peer-to-peer system. In *Proceedings of CAISE'03*, 2003.
- [7] V. Kalogeraki, D. Gunopoulos, and D. Zeinalipour-Yazti. A Local Search Mechanism for Peer-to-Peer Networks. In *Proceedings of CIKM'02*, 2002.
- [8] M. Karnstedt, K. Hose, and K.-U. Sattler. Query routing and processing in schema-based P2P systems. In *Proceedings of DEXA workshops 2004*, 2004.
- [9] G. Koloniari and E. Pitoura. Content based routing of path queries in peer-to-peer systems. In *Proceedings of EDBT'04*, 2004.
- [10] G. Koloniari and E. Pitoura. Peer-to-peer management of XML data: issues and research challenges. *SIGMOD Rec.*, 34(2):6–17, 2005.
- [11] N. Koudas, M. Rabinovich, D. Srivastava, and T. Yu. Routing XML queries. In *Proceedings of ICDE'04*, 2004.
- [12] Q. Lv, P. Cao, E. Cohen, K. Li and S. Shenker. Search and Replication in Unstructured Peer-to-Peer Networks. In *Proceedings of ICS'02*, 2002.
- [13] D. A. Menascé and L. Kanchanapalli. Probabilistic scalable P2P resource location services. *SIGMETRICS Performance Evaluation Review*, 30(2):48–58, 2002.
- [14] K. Nørnvåg, C. Doulkeridis, and M. Vazirgiannis. Taxonomy caching: A scalable low-cost mechanism for indexing remote contents in peer-to-peer systems. In *Proceedings of IEEE ICPS'06*, 2006.
- [15] V. N. Padmanabhan and L. Qiu. The content and access dynamics of a busy web site: findings and implications. In *Proceedings of SIGCOMM'00*, 2000.
- [16] S. Patro and Y. C. Hu. Transparent query caching in peer-to-peer overlay networks. In *Proceedings of IPDPS'03*, 2003.
- [17] C. Sartiani, P. Manghi, G. Ghelli, and G. Conforti. XPeer: A self-organizing XML P2P database system. In *Proceedings of EDBT Workshops 2004*, 2004.
- [18] G. Skobeltsyn and K. Aberer. Distributed Cache Table: Efficient query-driven processing of multi-term queries in P2P networks. In *Proceedings of P2PIR'06*, 2006.
- [19] K. Sripanidkulchai, B. Maggs, and H. Zhang. Efficient content location using interest-based locality in peer-to-peer systems. In *Proceedings of INFOCOM'03*, 2003.
- [20] D. Tsoumakos and N. Roussopoulos. Adaptive Probabilistic Search for Peer-to-Peer Networks. In *Proceedings of IEEE P2P'03*, 2003.
- [21] C. Wang, L. Xiao, Y. Liu, and P. Zheng. DiCAS: an efficient distributed caching mechanism for p2p systems. *IEEE Transactions on Parallel and Distributed Systems*, 17(10):1097–1109, 2006.
- [22] B. Yang and H. Garcia-Molina. Improving search in peer-to-peer networks. In *Proceedings of ICDCS'02*, 2002.