

AlvisP2P: Scalable Peer-to-Peer Text Retrieval in a Structured P2P Network*

Toan Luu[†], Gleb Skobeltsyn[†], Fabius Klemm[†], Maroje Puh[‡],
Ivana Podnar Žarko[‡], Martin Rajman[†], Karl Aberer[†]

[†] Ecole Polytechnique Fédérale de Lausanne (EPFL)
Lausanne, Switzerland
{firstname.lastname}@epfl.ch

[‡] University of Zagreb
Zagreb, Croatia
{firstname.lastname}@fer.hr

ABSTRACT

In this paper we present the *AlvisP2P* IR engine, which enables efficient retrieval with multi-keyword queries from a global document collection available in a P2P network. In such a network, each peer publishes its local index and invests a part of its local computing resources (storage, CPU, bandwidth) to maintain a fraction of a global P2P index. This investment is rewarded by the network-wide accessibility of the local documents via the global search facility.

The AlvisP2P engine uses an optimized overlay network and relies on novel indexing/retrieval mechanisms that ensure low bandwidth consumption, thus enabling unlimited network growth.

Our demonstration shows how an easy-to-install AlvisP2P client can be used to join an existing P2P network, index local (text or even multimedia) documents with collection-specific indexing mechanisms, and control access rights to them.

1. INTRODUCTION

In our vision, large-scale P2P text retrieval starts to represent an interesting alternative to existing centralized Web search engines. Many research results are now available in the literature, and operational systems are being deployed (e.g., Faroo[1] or YaCy[10]). In our approach to P2P retrieval we focus on distributing the indexing/retrieval load among a large number of interconnected nodes and support interesting novel usage scenarios. In such scenarios, the peers decide themselves which documents they want to make globally searchable and, more importantly, how these documents should be indexed and accessed. Thus, the effort of handling heterogenous data is distributed over the

*The work presented in this paper was (partly) carried out in the framework of the EPFL Center for Global Computing and supported by the Swiss National Funding Agency OFES as part of the European projects BRICKS (507457) and ALVIS (002068).

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the VLDB copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Very Large Data Base Endowment. To copy otherwise, or to republish, to post on servers or to redistribute to lists, requires a fee and/or special permission from the publisher, ACM.

VLDB '08, August 24-30, 2008, Auckland, New Zealand
Copyright 2008 VLDB Endowment, ACM 000-0-00000-000-0/00/00.

network and can be made more manageable. For instance, a specialized digital library might use sophisticated means for processing their local documents and use the P2P IR infrastructure to make their content searchable within the whole P2P network, possibly with specific access rights.

However, using a structured P2P network for distributing the load raises serious challenges for the design of the distributed indexing/retrieval mechanisms, especially when scalability with respect to bandwidth consumption, storage requirements, and load balancing is required. It has been shown, for example, that distributed algorithms using traditional single-term indexes in structured P2P networks generate unscalable network traffic during retrieval [11], mainly because of the bandwidth consumption resulting from the large posting list intersections required to process queries containing several frequent terms.

Our indexing strategy [6] relies on a novel approach based on two important properties:

- the generated distributed index stores posting lists for *carefully chosen indexing term combinations* (hereafter called *keys*), and
- the posting lists containing too many document references are truncated to a *bounded number of their top-ranked elements*.

We showed, both theoretically and experimentally, that these two properties guarantee acceptable storage and bandwidth requirements, essentially because the number of indexing term combinations remains scalable and the transmitted posting lists never exceed a constant size. In addition, our experimental results indicate that the retrieval quality remains comparable to state-of-the-art centralized search engines.

We have investigated two key generation techniques that are implemented in our prototype:

- indexing with *Highly Discriminative Keys* (HDKs) [7], which relies on global document frequencies to populate the index, and
- *Query-Driven Indexing* (QDI) [8, 9], which uses query popularity statistics to index only frequently queried term combinations.

We describe both techniques in more detail in Section 2.

2. DISTRIBUTED INDEXING/RETRIEVAL

In an Alvis P2P network, each peer is responsible for:

- the generation of index entries to be stored in the global distributed index for its local documents, and

- the storage and maintenance of the fraction of the global index associated with the keys that have been assigned to the peer by the hashing mechanism used in the underlying Distributed Hash Table (DHT).

Initially, the peers collaboratively build a distributed *single-term* index associating all single-term keys from the global collection with their top- k global document references. As the resulting retrieval quality might not be sufficient due to posting list truncation, the index is then additionally populated with carefully selected *multi-term* keys (indexing term combinations).

As already mentioned in Section 1, two indexing strategies can be used to populate the distributed index. The first one, hereafter called the HDK approach [7], generates new keys during the indexing phase based on observed document frequencies: each time a posting list for some key k exceeds a predefined size, new indexing keys (called expansions of k) with more terms (and thus associated with a smaller number of documents) are generated. Alternatively, in the Query-Driven approach [9], the index is populated only with frequently queried and non-redundant term combinations, and indexing is performed in parallel with retrieval. The second approach uses decentralized monitoring of query statistics to detect and index new popular keys, as well as to remove obsolete keys from the index.

In both cases, retrieval is performed in the following way: As soon as a peer receives a new query, it starts to explore the lattice of query term combinations (hereafter called the *query lattice*, see Figure 1) in decreasing combination size order, starting with the query itself. For each node in the query lattice, the querying peer requests the posting list associated with the term combination from the peer responsible for it. If the term combination is indeed present in the global index, the requested posting list is sent back to the querying peer, and if this list is not truncated, the part of the query lattice dominated by the term combination is excluded from further lattice exploration.

Additional approximations (e.g., pruning the part of the lattice dominated by a key associated with a truncated posting list) can be made to improve load balancing with an only marginal loss in retrieval precision. For example, assuming the term combination bc is indexed, while ab and ac are not, the result for the query abc is produced from the union of the truncated posting lists associated with the keys bc and a , as shown in Figure 1. Furthermore, during the exploration, each contacted peer also updates the usage statistics for the requested term combination.




Figure 1: Processing of a query $\{a,b,c\}$

Once the lattice exploration process terminates and all available posting lists relevant to the original query have been retrieved, the querying peer produces their union, ranks all the documents w.r.t the original query, and presents the top-ranked results to the user.

Finally, with the Query-Driven indexing strategy, a specific *on-demand indexing* mechanism is performed when a new popular key is detected during retrieval. The peer responsible for this key acquires a new posting list containing a bounded number of top-ranked document references [8]. The key is then considered as *indexed* and can thus be used for subsequent query processing. In general, the processing of new queries triggers the indexing of popular term combinations, which, in turn, increases the overall retrieval quality. At the same time, obsolete keys can be removed, resulting in an efficient indexing structure adaptive to the current query popularity distribution.

3. ALVISP2P ARCHITECTURE

The AlvisP2P architecture is layered in order to separate different conceptual levels, and allow the higher layers to use the functionalities provided by the lower ones. Altogether, it comprises the following layers:

- L1 A *transport layer*, which provides the means for direct communication between two peers;
- L2 A *peer-to-peer layer*, which maintains the Peer-to-Peer overlay infrastructure;
- L3 A *distributed IR layer*, which provides the basic functionalities related to document management, in particular the ones related to distributed IR;
- L4 A *ranking layer*, which implements functionalities related to distributed document ranking; and
- L5 A *local search layer*, which implements possibly sophisticated local IR models.




Figure 2: AlvisP2P architecture - layered view

Figure 2 shows all major AlvisP2P functionalities positioned in the corresponding layers. Each layer is composed of components that have been specified to fulfill the required functionalities [4]. While components in higher layers exclusively rely on the functionalities provided by lower layers, the architecture does not prevent from having different types of peers integrating in a more or less extensive way the layers from 3 above. For example, a lightweight peer could only integrate layers 1 to 4, while a peer associated with a more sophisticated local search engine could exploit all 5 layers. The discussion on the performance issues of such a system is presented in [4].

Layers 1 and 2 implement the peer-to-peer overlay infrastructure. Layer 2 (or P2P layer) consists of a Distributed Hash Table (DHT) that is able to sustain high traffic loads. Peers build routing tables of size $O(\log n)$, which results in an expected routing cost of $O(\log n)$ hops (where n is the number of peers in the network). As it uses the concept of

“hop space” for routing table construction, the DHT supports arbitrary skews in the distribution of the peers in the identifier space [3]. In addition, we integrated a congestion control mechanism into our DHT [2] to efficiently handle the large amounts of messages generated by the information retrieval application and to prevent the DHT from congestion collapses.

Layer 3 provides the features related to distributed information retrieval and implements one of the aforementioned techniques, i.e., indexing with highly discriminative keys (HDK) or query driven indexing (QDI). This layer deals with the task of *key-based indexing*, i.e., finding the set of keys and associated posting lists for a given document, and the *querying* task, i.e., given a query, finding corresponding keys in the global P2P index, retrieving the postings associated with those keys and merging the result set for ranking. Additionally, the QDI approach uses Layer 3 to collect the popularity statistics that define the keys to be indexed.

Layer 4 is responsible for ranking. Depending on the ranking model¹, it might use global document frequencies, average document length, term frequencies and other statistical information, which are stored in the P2P network, to compute the relevant scores of documents w.r.t the query.

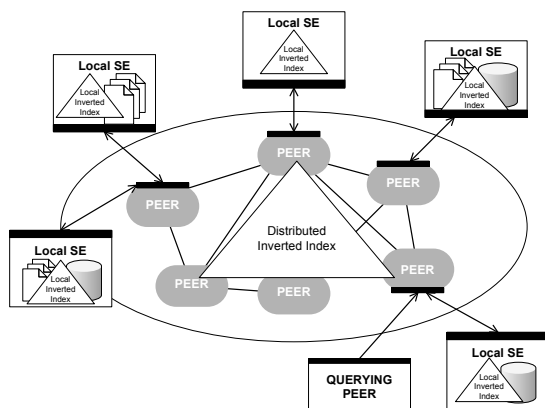


Figure 3: AlvisP2P network

Layer 5 implements a possibly sophisticated “local search engine”. For example, as shown in Figure 3, such a search engine can use specialized document processing for its local collection to build semantically rich indexes enhanced by various ranking strategies². The local search engine interacts with the associated peer through a generic API and uses a well-defined communication protocol to submit the index of its local collection to the global P2P network and to process queries.

More precisely, the answer to a given query can be:

- either produced exclusively using the information available in the distributed index and a uniform distributed ranking model; in this case the retrieval mechanism guarantees good response times, but, possibly, at the price of a lower precision;
- or refined in a second step during which the query is

¹Currently, we are using the state-of-the-art BM25 ranking function. Notice, however, that any other function could be used instead, provided that the required global statistics are available in the P2P network.

²E.g., it can support complex structured queries or/and employ a particular ranking strategy.

forwarded to the local search engines associated with the peers holding the documents found in the first step; in this case the retrieval might be slower (as it requires several interactions), but can benefit from the advanced features made available by the local engines.

4. ALVISP2P SOFTWARE

Joining an AlvisP2P network is as simple as downloading and installing the peer client software. The user only has to specify few communication parameters, such as the IP address of a contact peer and the communication port.



Figure 4: AlvisP2P peer Web-interface screenshot

The user can choose to enable a Web interface mode, which starts a Web server that can be accessed by anyone through a Web browser to query the AlvisP2P network, as shown in Figure 4. Otherwise, the default standalone client software is used, which allows only the local user to access the AlvisP2P network from this peer.

The Swing interface of the AlvisP2P client is shown in Figure 5. Once connected to the P2P network, the peer client software can be used to submit queries and browse the results. Figure 5 shows the client’s “Search” tab with a query result. For each document in the result, the following features are displayed: the URL of the hosting peer, the document title, a snippet and a relevance score. Additionally, the client software provides a shared file manager within which specific document access right can be defined.




Figure 5: Screenshot of the “Search” tab

Document access. To make documents searchable by other peers in the network, the user simply puts them in the shared directory of his/her peer and uses the AlvisP2P software to index them. A “Drag & Drop” function is also supported for this task. The corresponding “Manager of shared documents” tab is displayed in Figure 6.

Once a document is indexed, it becomes available at the URL: `http://PeerIP:Port/SharedDir/DocumentName`.

The following document types are supported: text, xml, html, doc, pdf, word or xml-based Alvis format. The index

of local shared document collection is implemented using the Terrier search engine³.

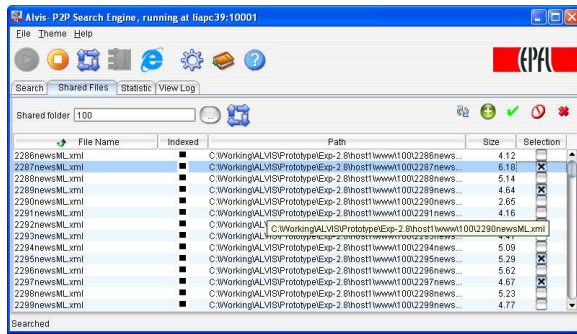


Figure 6: Screenshot of the “Shared documents” tab

External documents can be also integrated in the local document collection. To do so, the user simply needs to create an XML *file description* (following the Alvis format) containing the original URL of the document and some textual description of its content. The same procedure can also be used to publish multimedia files (video, audio).

As local documents always remain at the peer that holds them, the document owner can define specific access rights for them. For example, the user can choose that a document can be freely accessible or has a limited access controlled by a username and a password.

Heterogeneity support. As the software provides an interface to associate an Alvis peer with a local search engine, our framework can support heterogenous indexing models at different peers and ease their integration. The notion of *Alvis document digest* is used for this. A document digest is an explicit XML-based representation of the index of a document collection. It contains the list of the document URLs and, for each document, the list of its indexing terms, along with their positions in the document. Thus, a possibly sophisticated search engine (e.g., the one used by a digital library) can convert its local index into the Alvis digest format, and transmit it to the peer it is associated with. The peer then re-generates the local index in Terrier format and starts the distributed indexing process, which will make the document collection available for the whole P2P network.

Subsequently, if a user submits a query that retrieves one of the documents that have been published by an external search engine, he/she can click on the URL of the corresponding peer. The query is then automatically forwarded to the local search engine, which can generate an answer by the means of its local retrieval functionalities and send it back to the querying peer using the available API. This mechanism was successfully implemented in a peer-to-peer Service for Fedora [5].

5. DEMONSTRATION

Our demonstration aims at showing an operational AlvisP2P network where each participant can publish documents and search for some already available content. To do so, a large corpus of documents will be published in an AlvisP2P network running at a number of peers located in several research institutions such as EPFL and the University of Zagreb. A demonstration machine will be setup to

³<http://ir.dcs.gla.ac.uk/terrier/>

run one or several AlvisP2P clients, which will be able to join the running network through the Internet to perform network wide information retrieval and to index additional local content.

A second demonstration machine will be setup to illustrate the indexing/retrieval mechanisms implemented in our software. It will also report the current state of the network, as well as some critical statistics about bandwidth consumption, storage, etc.

The purpose of the demonstration is to let a user interact with the system to get a more detailed understanding of the distributed retrieval mechanisms. During the demonstration it will be possible to switch between the HDK and QDI approaches at any time, index some new documents, submit several queries and observe the results obtained using the distributed index.

6. CONCLUSION

In this paper, we have presented the AlvisP2P prototype for scalable full-text P2P-IR that uses carefully selected indexing term combinations associated with possibly truncated posting lists. The two presented indexing approaches (HDK and QDI) overcome the scalability problem of single-term retrieval in structured P2P networks, while preserving a retrieval quality fully comparable to state-of-the-art centralized search engine. The developed AlvisP2P prototype implements modules for distributed indexing, retrieval, and content-based ranking. This demonstration software represents a contribution to the design, development and research of realistic P2P search engine systems.

The most recent version of the prototype is available at <http://globalcomputing.epfl.ch/alvis>.

7. REFERENCES

- [1] Faroo. <http://www.faroo.com/info>, 2008.
- [2] F. Klemm, J.-Y. L. Boudec, and K. Aberer. Congestion Control for Distributed Hash Tables. In *NCA*, 2006.
- [3] F. Klemm, S. Girdzijauskas, J.-Y. Le Boudec, and K. Aberer. On Routing in Distributed Hash Tables. In *P2P*, 2007.
- [4] T. Luu, F. Klemm, I. Podnar, M. Rajman, and K. Aberer. ALVIS Peers: A Scalable Full-text Peer-to-Peer Retrieval Engine. In *P2PIR*, 2006.
- [5] G. S. Pedersen. Considerations about a Peer-to-Peer Service for Fedora, 2006.
- [6] I. Podnar, M. Rajman, T. Luu, F. Klemm, and K. Aberer. Beyond term indexing: A P2P framework for Web information retrieval. *Informatica*, 2006.
- [7] I. Podnar, M. Rajman, T. Luu, F. Klemm, and K. Aberer. Scalable Peer-to-Peer Web Retrieval with Highly Discriminative Keys. In *ICDE*, 2007.
- [8] G. Skobeltsyn, T. Luu, I. Podnar Žarko, M. Rajman, and K. Aberer. Query-Driven Indexing for Scalable Peer-to-Peer Text Retrieval. In *Infoscale*, 2007.
- [9] G. Skobeltsyn, T. Luu, I. Podnar Žarko, M. Rajman, and K. Aberer. Web Text Retrieval with a P2P Query-Driven Index. In *SIGIR*, 2007.
- [10] YaCy. <http://yacy.net>, 2008.
- [11] J. Zhang and T. Suel. Efficient Query Evaluation on Large Textual Collections in a Peer-to-Peer Environment. In *P2P*, 2005.