

Advances In Peer-To-Peer Content Search

Madjid Merabti · Zhu Liu · Heather Yu ·
Deepa Kundur

Received: 10 February 2008 / Revised: 12 December 2008 / Accepted: 19 January 2009 / Published online: 5 February 2009
© 2009 Springer Science + Business Media, LLC. Manufactured in The United States

Abstract Peer-to-peer (P2P) computer networks have recently received tremendous attention due to their inherent scalability and flexibility, which facilitates a broad spectrum of innovative multimedia applications. Such networks rely on the power of participant nodes of the network (called peers) for communications and computation. Traditional applications of P2P multimedia include decentralized file sharing and content distribution. Yet, the value of the virtually unlimited amount of data distributed in the P2P network will be sacrificed if effective and efficient ways to locate the content are missing. This challenge has stimulated extensive research in recent years, and many new P2P content search methods have been proposed. This paper provides a timely review of influential work in the area of peer-to-peer (P2P) content search. We begin with a survey of text-based P2P search mechanisms and continue with an exposition of content-based and semantic-based approaches followed by a discussion of future directions.

Keywords Peer-to-peer · Content search

1 Introduction

The power to access relevant and salient information in a timely and cost-effective manner is critical in today's information age. Models for information sharing between ad hoc groups of users are currently being investigated for collaborative and cooperative multimedia sharing applications. The most common model to date is based on the peer-to-peer (P2P) communication network model. Ideally, a P2P network consists of equal *peer* nodes that can take on the roles of both client and server to other peer network nodes. Such network operation makes use of the data acquisition capabilities, computing power and bandwidth of the network nodes rather than exploiting an existing network infrastructure for operations such as routing and information retrieval.

Since Napster made its debut in 1999, peer-to-peer (P2P) content access and distribution has been gaining increased popularity. Many other P2P systems, which span relatively simple file sharing networks, more complex multimedia content search entities and sophisticated media streaming systems, have been subsequently deployed. P2P computer networks are popular due to their inherent scalability and flexibility, which facilitates a broad spectrum of innovative multimedia applications. Such networks rely on the power of participant nodes of the network for communications and computation. This is in contrast to the traditional client-server models for multimedia communications that employ dedicated networking infrastructure. In the last few years, many P2P applications became mature and started to benefit web users enormously. Representative examples

M. Merabti
School of Computing & Mathematical Sciences,
Liverpool John Moores University,
Liverpool, UK

Z. Liu (✉)
AT&T Labs—Research,
Middletown, NJ, USA
e-mail: zliu@research.att.com

H. Yu
Huawei Technologies (USA),
Bridgewater, NJ, USA

D. Kundur
Electrical & Computer Engineering Department,
Texas A&M University,
College Station, TX, USA

include BitTorrent, eMule, Kazaa, Kontiki, PPLive, Skype, Hulu, etc.

P2P networks are capable of serving virtually unlimited amount of content, but also make content search extremely challenging. The distributed nature of content in P2P networks and the desire of the independency on search engines that may be biased or censored demand native P2P content search engines. The challenge is how to achieve the high efficiency and the high retrieval quality delivered by centralized search engines on P2P networks. Recently, tremendous research effort has been devoted to this topic.

Naturally, processes such as content search must take into account this inherent network model as well as the characteristics of the information being acquired, stored and communicated. For multimedia data, such P2P search and retrieval processes are especially challenging to design; inherent compromises among search time, retrieval accuracy, and query message bandwidth, must be accounted for. This paper addresses recent advances in P2P content search by providing an overview of influential research in the area. We begin by giving an overview of text based P2P file search and then content-based search and semantic-based search mechanisms are reviewed followed by a discussion of future directions.

2 Text Based P2P Search

One of the earliest P2P implementations that brought P2P computing into the mainstream and which sparked a large amount of media attention was Napster [1]. Napster was created purely for the distribution of MP3 audio files, and as such it was swamped with negative press because people were downloading digital content illegally and consequently ignoring content copyright. Each Napster node downloads and installs the client software used to connect the peer to the centralized Napster server. Once connected, peers share MP3 files stored locally on their hard drives, with text-based information about them being indexed and stored by the Napster server. Clients submit text-based queries to the Napster servers for a particular audio file. This results in a list of files that match, along with the connection information, username, IP and port address the querying client must use to connect to the peer hosting the file. Once the querying peer has this information it attempts to connect to the peer and transfer the target content in a P2P fashion. At this point the Napster server is no longer required [2].

Although Napster proved successful and is said to be the grandfather of modern P2P computing models it suffered from a number of limitations. The major limitation was the fact that it could only share MP3 content. In addition, its hybrid model was reliant on client–server technology—if the server became unavailable the discovery mechanism used to find content was lost.

Another hybrid protocol, similar to Napster called iMesh [3], uses a centralized server, to which clients connect in order to search for content. The iMesh model differs somewhat to Napster in two respects. First, it allows any content to be shared including MP3 audio files. Second—the reason why iMesh has not been subjected to the same legal problems as Napster—it has a mechanism to remove copyrighted files from the network.

Computational expense and scalability issues associated with the above mentioned models are well documented, which has resulted in new P2P networks devoid of any centralization. The most popular being the Gnutella protocol [3]. Like iMesh it provides a generic file sharing mechanism that allows any digital media content to be shared. However it differs from iMesh and Napster because the Gnutella protocol uses a purely decentralized model, which is not reliant on any centralized authority.

The search mechanism used by Gnutella adopts a different approach to Napster in that it does not require any centralized server to manage the location of content within the network. Search packets containing text queries are used with predefined time to live (TTL) values, the default value being 7, which corresponds to the maximum number of hops the message can take before it is terminated by the network. The packet is passed to all the immediate peers' the querying peer is connected to, which in turn is passed to all the peers the peer is connected to. The *horizon*, as defined by Kan [1], given a TTL of 7 encompasses about ten thousand nodes. If a node in the horizon is found with a file name matching the query, the information is routed back to the querying peer. The file can then be downloaded directly from the target node. This is commonly referred to as blind search.

Unlike Napster, it is difficult to disrupt the network because no one single node is responsible for the search mechanism. If any given node is lost the overall search mechanism of the Gnutella network is not effected. In the worst case, one loses the content provided by that node. Consequently Gnutella provides mechanisms to counteract some of the limitations associated with Napster. As such many Gnutella clients have been developed since the protocol was first released in 2000, including Bearshare [5], Shareaza [6], and Limewire [7].

The FastTrack protocol promises better performance than Gnutella and its variants. A number of popular applications, such as Kazaa [8], Morpheus [9] and Grokster [10], use the FastTrack protocol, which divides users into two groups. The first group contains supernodes and the second contains ordinary nodes. Supernodes are defined as computers with significant computation, network and bandwidth capabilities. All supernodes are connected together to create an overlay network that acts as a hub processing all data requests received from ordinary nodes

within the network, which are inherently less capable nodes.

When a node wants to share or search for a file a request is submitted to the supernode, which in turn submits it to all other supernodes, which in turn propagate the request to the ordinary nodes they are servicing. Like Gnutella, messages are configured with a TTL value of 7, ensuring that message propagation is terminated once seven hops have been reached.

Once the content has been found it is transferred directly from the target node to the querying node using the HTTP protocol, without using the supernode. There is a subtle distinction between the FastTrack model and that of Napster in that the Napster server managed an index of audio file information thereby breaching copyright laws. The FastTrack protocol avoids this problem because it only manages a list of supernodes and not explicitly the information regarding the content itself. Supernodes are ad hoc by nature and are free to join and leave the network at any time. So information about supernodes held by the FastTrack servers continually changes. This abstraction detaches the FastTrack protocol, including the applications that use the protocol, from media content and thus some believe that FastTrack-based applications do not aid copyright infringement.

The difficulty with protocols such as Gnutella and FastTrack is that they rely on flooding or random walk trajectories for content search, with messages propagated to every peer. This results in increased costs and network traffic. Wang and Xiao [11] aim to alleviate these limitations using their proposed Differential Search (DiffSearch) algorithm. They claim DiffSearch improves search efficiency of unstructured P2P networks by giving higher querying priority to peers with high query/reply capabilities, known as ultrapeers. Ultrapeers form an overlay and serve visiting peers known as leaf nodes. The indices of leaf nodes are uploaded to ultrapeers allowing all shared content to be searched within what they call the first round. Based on tests using Gnutella, Wang argues that 1% of peers answer the main portion of queries. Consequently by routing queries to these peers it is possible to save up to 90% of query traffic. Using counters to track which files answer queries, which they call *effective files*, a matrix is created allowing ultrapeers to be self-aware by counting the number of shared files which have been visited. If the number of shared files exceeds a threshold, a peer can promote itself to ultrapeer status. This results in an overlay where members have higher priority depending on where they reside in the hierarchy. To further decrease traffic, DiffSearch hitchhikes query/response messages to perform network management task. For example, allowing ultrapeers to advertise themselves to leaf nodes and vice versa.

The Foreseer P2P system aims to address various limitations using distributed indices [12]. Cai and Wang claim their approach improves efficiency in decentralized unstructured P2P systems using two orthogonal overlays, which they term neighbor and friend overlays. To use their example, everyone has neighbors and friends which form part of an individual's social network. People tend to get to know their neighbors over time as they become more settled within their environment and make friends through social interactions. Implementing this scenario in Foreseer, friend nodes can serve future requests with a high probability (temporal locality), whilst neighbor nodes can offer quality of service, such as fast responses and low resource consumption if they are able to carry out the request (geographical locality). Locality is also discussed in Datta et al. [13], where it is considered important for scalability in data mining—the same principles apply to Foreseer.

Extending this scenario content is searched for using a collection of business cards provided by all neighbor and friend nodes. Finding content is as simple as sending/receiving a request. Using each business card, the node checks to see if a peer exists capable of service the request. If a suitable node is found the node in question is contacted. However, if the request cannot be serviced the request is passed to all its neighbors and friends. In Foreseer, business cards refer to a peer's content filter, which is derived using the Bloom filter (BF) on all the content it shares. It is therefore only able to address text-based content.

A by-product of the neighbor and friend overlays is that it provides an efficient search direction, where random walking or query flooding is unnecessary. Query requests contain one or more terms and it is these terms that are compared to the content filters for the neighbors and friends the node is aware of. In local matching a node computes the query filter by mapping the query terms and comparing it with the content filter of each node it has routing information about. If a match is found it indicates that a node may contain all the key words with high probability. If the local matching fails the query is selectively forwarded based on the results obtained from the first approach. The query in this instance is forwarded along the neighbor and/or friend links where local matching is performed.

Taking an opposing view, efficiency in P2P network has been addressed using different techniques, where processes are more finely controlled so that structure can emerge. Many approaches use distributed hash tables to achieve this. Chord [14, 15] is one such protocol that adopts these principles where order emerges using their distributed hash table (DHT) routing algorithm. Its basic structure forms a ring topology, whereby each node only has to establish one connection. A consistent hashing function, such as SHA-1, is used to generate node and object identifiers known as keys. The node identifier is created using the IP address and

port. The object identifier, which can be any kind of shared content, is created using the data to be shared within the ring. Node identifiers are arranged in a circle modulo 2^m , where m is the length of the hash value. Every key k is assigned to the node whose identifier n is larger than or equal to the hash value of k . The node the key belongs to is called the successor. In Chord, node identifiers increase clockwise and keys are assigned to the first nodes that reside closest to them clockwise. Chord uses a hashing function designed to distribute keys evenly throughout the ring topology, whereby all nodes roughly receive the same number of keys.

Every node is aware of their successor and as such queries are passed from successor to successor. When a node is reached that has a hash value greater than or equal to the hash value of the key, this node can map the query to the key. In order to overcome the need to traverse every node, a node can attempt to find the predecessor of some key k using a finger table. Node n achieves this by searching its finger table for some node x that immediately precedes some key k . If it finds node x then it queries it to determine which node is closer to k . By repeating this process n moves the query closer and closer to k . In Chord this is called iterative routing.

CAN [16] is similar to Chord and also uses the DHT concept to self-organize, share content and route queries. CAN forms a P2P overlay network that stores chunks of a DHT, known as zones. The protocol is based on a virtual d -dimensional Cartesian coordinate space. This space is dynamically partitioned among all the nodes in the system, so that every node owns its own zone within the global coordinate space. This space stores key–value pairs where k_1 is mapped onto a point p in the space using a uniform hashing function. The key–value pairs are stored on the node that owns the zone in which p resides. To discover the values of some key k_1 any node can use the hash function to map k_1 onto p and retrieve the contents from p . This may be the content or a pointer to the content. If the p is not owned by the querying node or its neighbor, then the request is routed towards the node where p resides.

Pastry [17] is not too dissimilar to Chord and CAN in that it also uses a DHT-based protocol to form a self-organized overlay network. Pastry nodes are identified in the network space using a 128 bit identifier, known as the *nodeId*. The *nodeId* indicates a node's position in the circular *nodeId* space. The *nodeIds* themselves are assigned randomly when the node first connects to the Pastry network. Several mechanisms can be used to derive the *nodeId*, however typical implementations use the node's public key or IP address to create a hash. In Pastry *nodeIds* are thought of as a sequence of digits in base 2^b . Nodes within Pastry maintain their own routing table, which is organised into $128/2^b$ columns. As well as the routing table,

each node also maintains a neighborhood set M containing the *nodeIds* and IP addresses of the M nodes closest to the local node. The set is not used for routing, but rather for maintaining locality properties [18]. Nodes also maintain a leaf set L containing the set of nodes with numerically closest but larger *nodeIds* and numerically smaller *nodeIds*, relative to the present nodes *nodeId*. The leaf set is used when messages are routed. When a node receives a message it first checks to see if the key falls within the range of *nodeIds* covered by its leaf set. If it is, the message is forwarded directly to the destination node. If the key is not covered by the leaf set, the routing table is used and a message is forwarded to the node that shares a common prefix with the key by at least one more digit.

DHT-based P2P protocols are said to provide considerable benefits over previous generations, providing emergent behaviors that support order and increased performance. However, they are expensive to maintain because the network topology is continually changing. Consequently managing a consistent DHT requires considerable effort. In an attempt to leverage the benefits of DHT, but also minimize some of its inherent limitations the JXTA [18] specification have tried to create a balance by creating a hybrid system that uses a loosely consistent DHT [19]. Whilst protocols such as Chord rely on more costly mechanisms to keep the network view consistent, JXTA uses a less costly mechanism that ensures the network view is only loosely-consistent. The advantage with this approach is that it is less expensive to maintain, however the disadvantage is that it may be temporarily or permanently inconsistent.

Yang and Ho [20] enhance the DHT-based P2P networking approaches discussed above by proposing a keyword search scheme called Proof that utilizes advances in information retrieval research. The Proof protocol reduces network traffic, decreases search latency and provides high quality search results. The Proof system comprises a crawler, a database, an index generator, and a distributed P2P system. The crawler collects web pages and extracts hyperlink information for computing page rank values, whilst the index generator produces new index structures and publishes them to the P2P system. The P2P system itself assumes N peers are contained and uses a consistent hash function to assign an identifier to each peer. The system contains documents and a vocabulary, which contains all of the keywords in the documents. Each document is composed of several keywords, which are a subset of the vocabulary. Given a query containing several keywords, a subset of the vocabulary and a user-specified result threshold, the search problem is defined in terms of finding the top relevant documents that contain all the query keywords.

For a query with several keywords and a result threshold value the first processing peer, called the major query

processing peer, requests all other peers, which are responsible for the inverted lists of query words, to report the length of its inverted list. Once the lengths have been received, the first processing peer then determines a query flow from the shortest inverted list to the longest one. Yang et al. argue that Proof has a number of advantages; the query flow arrangement chooses the shortest inverted list; decreased network load and computation time; BFs are used to network traffic; and BF precision makes the Bloom membership query more accurate and further reduces traffic.

While many interesting solutions have been proposed in both structured and unstructured P2P networks, Ahmed and Boutaba [21] believe that flexibility and efficiency remain unsolved problems. To address these challenges Ahmed proposes their Distributed Pattern Matching system (DPMS). DPMS is based on BF based pattern matching distributed throughout the P2P network. Given a search pattern Q , DPMS tries to find peers containing some pattern P that matches Q , i.e. the 1-bits of Q for a subset of the 1-bits found in P . DPMS peers can act as either a leaf peer or an indexing peer where the former resides at the bottom level of the indexing hierarchy. This type of node advertises its indices representing the content the peer wishes to share. Whilst, indexing peers store indices received from other peers—these peers may be leaf peers or indexing peers. Peers join different levels within the hierarchy and can act both as a leaf peer and an indexing peer. Within this hierarchy indexing peers disseminate index information using repeated aggregation and replication. Replication is used for disseminating patterns from leaf peers to a large number of indexing peers. To overcome increased traffic load, DPMS combines replication with lousy-aggregation. Advertisements provided by different peers are aggregated and propagated to peers in the next level along the aggregation tree. Based on repeated lousy aggregation, information content of the aggregates is reduced as you move towards the top of the indexing hierarchy. This helps balance the system and improve fault tolerance. Furthermore, peers can route queries towards a target without having any global knowledge of the overlay topology. It also helps minimize query forwarding traffic.

The query life-cycle can be divided into three phases: ascending phase, blind search phase and descending phase. Using the ascending phase, an initial (or intermediate) peer checks its local information for the existence of a match. If a match is found the query is forwarded to the matching child, otherwise it is forwarded to any of its parents. This process continues until a query hits a peer with a match or reaches the highest level peer. Blind search is executed by a highest level peer receiving a query (from a child) that does not match any aggregate in its aggregate lists. This peer floods the query to all other peers in its group. If no peer in a group at the highest level contains a match then the search

fails. When a peer hits a peer containing the matching aggregate, it enters the descending phase. The query is forwarded to the child peer advertising the matching aggregate. This process continues until the query reaches a leaf peer.

Today, increasing numbers of P2P systems are seeking to support more powerful queries. Joost, a broadcast quality Internet TV service, for example supports phrase queries, wildcard queries, proximity queries, range queries and more. An algorithm that supports range queries over DHT is proposed in [22]. It implements range queries over DHT via a trie-based scheme in which every vertex corresponds to a distinct prefix of the data domain being indexed. A range query can be performed via longest common prefix search followed by a parallel traversal in the P2P network to retrieve all the desired items.

3 Content-based Search

Most existing P2P systems provide very limited content search capabilities, for example, search based on document title, author names, keywords, or descriptive text. To retrieve the relevant content more effectively, we need an approach that provides richer searching features. Content-based search is essential for querying textual documents, and it is also desirable for querying multimedia data when text annotations are nonexistent or incomplete. Flooding is still the main mechanism for content searching in unstructured P2P networks, e.g., the Gnutella-like networks. The performance of flooding based search can be improved by exploiting locality in peer interests, smarter message routing and data replication strategies. In structured P2P networks, inverted index of documents under single terms or term vectors are utilized with the support of DHT. In this section, we first discuss some recent work on content search for textual documents, and then extend the discussion to the domain of multimedia data.

Full-text keyword search is important for P2P web search, since nowadays users are accustomed to this service powered by Google search engine that indexes billions of web documents. Normally, an inverted index (also referred to as inverted file) is used to store the mapping from a keyword to a set of documents that contains such a keyword. In centralized systems, such inverted index tables, which are normally huge, reside on servers. In P2P systems, there are two ways for distributing the inverted index: (1) partition the document collection among peers such that each peer manages a local inverted index for a set of documents (local inverted index), and (2) partition the index terms such that each peer manages the inverted index for a set of keywords (global inverted index).

Zhang and Hu [23] proposed a search protocol called *assisted search with partial indexing* to improve content

search in unstructured P2P networks. The proposed P2P network consists of two logical overlays: an unstructured search overlay and a structured index overlay. The index only maintains information about the top interests of peers and unpopular data. The index overlay assists the search overlay to improve its search performance in three ways. First, peers communicate their interests via the index overlay, and search overlay is constructed based on peer interests. Second, due to limited node degree, peer neighbors only reflect a peer's top interests. Third, peers identify their local data that are globally unpopular and are not part of local interests. Experimental results show that such a lightweight partial indexing service can significantly improve the success rate and search speed in locating data, while reducing traffic overhead.

Chen et al. [24] proposed an efficient multi-keyword search mechanism over P2P network. The designed hybrid P2P network is a combination of an unstructured P2P network which can use a gossiping algorithm to gather global statistical information, and a BF enabled overlay based on DHT global inverted indexes. The BF is a lossy but succinct and efficient data structure to represent a set. By transmitting the encoded sets instead of raw sets for distributed intersection/union operation can effectively reduce the amount of traffic. Chen et al. showed how to optimize the setting of a BF in terms of traffic cost by the global statistical information of keywords, and designed the optimal order strategies based on the BF for queries with both “and” and “or” operators.

Multimedia content indexing and retrieval has been an active field for more than a decade. It draws tremendous research effort from the academia, the industrial, and the standard organizations. IBM Research developed a prototype multimedia analysis and retrieval system, called MARVEL [25]. It consists of two components: the multimedia analysis engine, which applies machine learning techniques to model semantic concepts in video, and the multimedia search engine, which integrates semantics-based searching with other search techniques (speech, text, metadata, audio-visual features, etc.). The Informedia II [26] digital video library at CMU is another pioneering multimedia database system that consists of more than 1,500 h of video. Informedia combines speech recognition, image understanding and natural language processing technologies to automatically transcribe, segment, index, and summarize the linear video. MIRACLE [27] is an ongoing research project at AT&T Labs aimed at creating automated content-based media processing algorithms and systems to collect, organize, index, mine, and repurpose video and multimedia information. This video search engine combines existing metadata with content-based information that is automatically extracted from the audio and video components. The MIRACLE search engine currently operates on an archive of more than 32,000 h

of video that have been collected and automatically indexed over a 10 year period.

MPEG-7 is a standard sponsored by the International Organization for Standardization for describing the multimedia content. It provides support to a broad range of applications, and it will make the web as searchable for multimedia content as it is searchable for text. The evolution of the World Wide Web, including the introduction of Rich Site Syndication (RSS), Web 2.0, and the semantic web, enables the web information be machine processable (rather than being only human oriented), thus permits browsers or other agents to find, share and combine information more easily. In this section we will briefly describe a few examples of content-based search in P2P systems.

Tang et al. [28] proposed PeerSearch, an efficient P2P information system that supports content and semantic search. PeerSearch extends existing information retrieval methods: the vector space model (VSM) and the latent semantic indexing (LSI) to work with the efficient routing mechanisms in a Content Addressable Network (CAN). Basically, LSI uses singular value decomposition to transform and truncate a matrix of document vectors computed from VSM to discover the semantics underlying terms and documents. The authors used the semantic vector of a document as the key to store the document index in CAN, such that the indices stored nearby in CAN are close in semantics. The same technology can be applied to audio or video data, where the semantic vectors have to be extracted from audio or video data.

Lu and Callan [29] explored content-based resource selection and document retrieval algorithms in hybrid P2P networks. In their approach, the leaf node determines the retrieval results for certain query using probabilistic information retrieval algorithm, and the directory node (supernode) builds a unified content model for all of its leaf nodes and a set of neighboring directory nodes. The content model is used for routing query messages.

Yang [30] described a content-based music retrieval system in P2P environment. Each audio document is converted into a stream of characteristic sequences, where each sequence is a vector representing a short segment of music data. All characteristic sequences are indexed using Locality-Sensitive Hashing scheme, such that similar (in terms of human perception) vectors can be hashed into the same hash value with high probability. Given a query audio, the retrieval procedure is to find a list of matches on characteristic sequences with the tolerance of tempo changes. To improve the search efficiency in a P2P environment, a two-phase search protocol was proposed. In the presearch phase, the query peer broadcasts a small subset of query vectors to all potential peers. In the actual search phase, peers with a higher chance of a hit will conduct the more rigorous search.

In [31], Lee and Guan presented a content-based image retrieval system over a P2P network. Each peer in the system maintains two look-up tables, one for generic neighbors which are typically the neighbors with the least physical hop counts, and the community neighbors which share common interests at the image content. The system uses the historical retrieval results to identify the community neighbors, such that the subsequent retrieval within the community neighbors will result better retrieval precision. The adopted visual features include color, color moments, object shape, and texture.

Ardizzone et al. [32] proposed a P2P architecture for multimedia content retrieval that exploits compact representation of multimedia resources shared by each peer and adapts the network topology dynamically to peer interests based on the query interactions among peers. In this work, video is decomposed into shots, and each shot is represented by a few keyframes. The adopted visual descriptors for each keyframe include color histogram, edge density, motion magnitude, and motion vectors direction histograms. By exploiting the results of previous interactions with its neighbors, each peer builds and refines the profiles of other peers, which concisely describe their interests and contents. Relying on these profiles, queries are forwarded to the neighbors with higher probability to satisfy the queries.

SAPIR (Search in Audio-visual content using P2P Information Retrieval) [33] is part of the CHORUS, a European Coordination Action which aims at creating the conditions of mutual information and cross fertilization between the European projects dealing with multimedia content search engines. The goal of SAPIR is to extend the power of web searches beyond centralized text and metadata searches to include distributed audio-visual content. The SAPIR consortium is built of a mix of academic and research institutes as well as industrial companies. The SAPIR media framework analyzes digital content and represents the extracted features in a common schema based on MPEG-7, which can be used to index and search content in a P2P network. A P2P architecture is being built to provide a scalable indexing structure that can be used for multi-feature search.

4 Semantic Search

Unlike text based search, semantic based search seeks to answer queries based on the meaning and its specific context and circumstances of the queries. The key challenge is how to use computer attainable means to describe the meanings and thus match the meanings of the query and that of the data in the database.

In the information retrieval arena, Vector Space Model (VSM) [34] and Latent Semantic Indexing [35] are two

popular approaches for modeling and analysis of the meaningful relationship between terms and documents. VSM is an algebraic model for representing objects, and most particularly text documents, as vectors of identifiers. With VSM, each phrase, sentence or document can be represented as a vector of multiple dimensions of terms. Similarity of phrases, sentences, and documents can be compared in the vector space based on the distances and angles. Latent Semantic Analysis (LSA) [36] takes on a different approach. It analyzes the relationship between terms and documents by producing a set of concepts related to the terms and the documents. Using LSA, one can transform the occurrence of terms into a relation between the terms and some concepts. That is, given a query of terms, LSA tools can translate it into the concept space, and find matching documents accordingly.

Recall that in a DHT based overlay network, objects are located based on the object key which is often generated via a one way mathematical function, the hash function. The hash function converts the object specific description string into a small datum, the key. Since the goal of semantic search is to locate the record based on the semantic instead of the exact key of the object, an important issue to study is how to map a pair of objects that are semantically close to a pair of keys that are mathematically close, i.e., can be identified as ‘similar’ in search. In other words, the key in any semantic search in a DHT based overlay network is to find a mapping function such that the distance between the keys of two documents is proportional to their dissimilarity in semantic.

One possible approach is to take advantage of VSM or LSI by generating a semantic vector in the concept space for each object description string and then use the semantic vector to generate a multi-dimensional DHT key of the object. In [37], an algorithm pSearch was introduced. Two basic operations are involved in pSearch: indexing and searching. Whenever a new object T is added, pSearch performs the indexing operations as follows:

1. Use LSI to derive T 's semantic vector V^T .
2. A rolling index is employed to generate K_i^T of DHT keys ($k_i^T, i=0, \dots, p-1$) from V^T .
3. Index T into the underlying DHT using these DHT keys.

Whenever a query Q is issued, pSearch performs the search operations as follows:

1. Use LSI to derive Q 's semantic vector V^Q .
2. Use a rolling index to generate K_i^Q of DHT keys ($k_i, i=0, \dots, p-1$) from V^Q .
3. Route Q to the destination nodes that are responsible for these DHT keys K_i^Q .
4. On reaching the destination, Q is either flooded to nodes within a radius r or forwarded to nodes using content-directed search.

5. All nodes that receive the query do a local search using LSI and return the matched documents to the query originator node.

Clearly, the semantic search capability of the above proposed scheme relies on the capability of the semantic vector. For instance, if one or more critical terms in the object description are not in the LSI, the recall rate of the semantic search will be greatly reduced. To improve the recall rate, ontology based analysis maybe introduced.

Compared with structured overlay, unstructured overlay has certain advantages in implementing semantic search since it may not rely on DHT for object look up and is often capable to support arbitrarily complex queries. In [38], Zhu et al. proposed a dual link approach efficient semantic search (ESS) where each node is connected both via a random link and a semantic link. The goal of their dual link based topology adaptation is to ensure that (1) semantically closed objects are organized into the same semantic groups through semantic links, and (2) high-capacity nodes have high degree and low-capacity nodes are within short reach of higher-capacity nodes. Given a query, a relevant semantic group is first located. Then the query is flooded within the semantic group to retrieve relevant documents. The search process is continued until sufficient responses are found. The intuition behind the flooding within a semantic group is that semantically associated nodes tend to be relevant to the same query. Noticeably, ESS would introduce additional complexity onto the conventional unstructured overlay searching schemes. Hence, additional strategies are needed to improve the searching efficiency in order to make such approaches more practical in real life systems.

To date, semantic based search is still a very challenging problem in the centralize database systems. The distributed nature of P2P networks add on another layer of complexity to the problem set. Nevertheless, technology advances are making intelligent search a more and more demanding functionality in many applications. How to take advantage of existing technologies and invent an efficient and effective semantic based P2P search scheme is thus an important topic to investigate.

5 What's Next

With the recent merging of content acquisition, communication networking and computation to form emerging multimedia sensor or surveillance systems, the need for advanced content mechanisms is paramount. Many of these systems are predicted to have P2P communication architectures [39] necessitating efficient P2P content-based

search for higher-level content retrieval and understanding. Imagine a world in which multimedia sensor networks (interpreted as sensor databases containing time critical information) are employed to detect emerging natural disasters or terrorist activities. The utility of such systems is, in part, determined by the ability to effectively and efficiently retrieve information for a given application given often incomplete or vague information queries.

Furthermore, for such multimedia database systems security and privacy issues are paramount [40–42]. Research into security policy development, security architectures, authenticated querying, privacy protection, and access control for information retrieval are essential during system inception in order to guarantee the most seamless, cost-effective and robust solution.

References

1. Minar, N., & Hedlund, M. (2001). A network of peers: peer-to-peer models through the history of the Internet. In A. Oram (Ed.), *Peer-to-peer: harnessing the power of disruptive technologies*. Sebastopol: O'Reilly Media.
2. Gradecki, J. D. (2002). *Mastering JXTA: building Java peer-to-peer applications*. New York: Wiley.
3. iMesh Professional 5.0. 2005, iMesh Inc., <http://www.imesh.com/>
4. The Gnutella Protocol Specification v0.4, Gnutella, http://www9.limewire.com/developer/gnutella_protocol_0.4.pdf
5. BearShare, 2005, Free Peers, <http://www.bearshare.com/>
6. Shareaza, 2005, Shareaza Development Team, <http://www.shareaza.com/>
7. LimeWire, 2005, Lime Wire LLC, <http://www.limewire.com/english/content/home.shtml>
8. KaZaa, 2003, P. Morle, A. Morris, and N. Hemming, <http://www.zeropaid.com/kazaalite/>
9. Morpheus, 2005, StreamCast Networks, <http://morpheus.com/>
10. Grokster, 2005, <http://www.grokster.com/>
11. Wang, W., & Xiao, L. (2007). An effective P2P search scheme to exploit file sharing heterogeneity. *IEEE Transactions on Parallel and Distributed Systems*, 18(2), 145–157. doi:10.1109/TPDS.2007.20.
12. Cai, H., & Wang, J. (2006). Exploiting geographical and temporal locality to boost search efficiency in peer-to-peer systems. *IEEE Transactions on Parallel and Distributed Systems*, 17(10), 1189–1203. doi:10.1109/TPDS.2006.139.
13. Datta, S., Bhaduri, K., Giannella, C., Kargupta, H., & Wolff, R. (2006). Distributed data mining in peer-to-peer networks. *IEEE Internet Computing*, 10(4), 18–26. doi:10.1109/MIC.2006.74.
14. Eberspacher, J., Schollmeier, R., Zols, S., & Kunzmann, G. (2004). "Structured P2P Networks in Mobile and Fixed Environments," *HET-NETs '04*. West Yourshire, UK.
15. Dabek, F., Brunskill, E., Kaashoek, F., & Karger, D. (2001). "Building Peer-to-Peer Systems with Chord, a Distributed Lookup Service," *HotOS-VIII* (pp. 81–86). Germany.
16. Ratnasamy, S., Fancis, P., Handley, M., & Karp, R. (2001). *A scalable content-addressable network*, *ACM SIGCOMM 2001* pp. 161–172. San Diego: ACM.
17. Rowstron, A., & Druschel, P. (2001). "Pastry: scalable, distributed object location and routing for large-scale peer-to-peer systems,"

- IFIP/ACM International Conference on Distributed Systems Platforms (Middleware)*. Germany: Heidelberg.
18. Gong, L. (2001). JXTA: a network programming environment. *IEEE Internet Computing*, 5(3), 88–95. doi:10.1109/4236.935182.
 19. Traversat, B., Abdelaziz, M., & Pouyoul, E.: “Project JXTA: a loosely-consistent DHT rendezvous walker,” URL: <http://www.jxta.org/docs/jxta-dht.pdf>
 20. Yang, K., & Ho, J. (2006). “Proof: a DHT-based peer-to-peer search engine,” *WI 2006*, Hong Kong, pp. 702–708.
 21. Ahmed, R., & Boutaba, R. (2007). Distributed pattern matching: a key to flexible and efficient P2P search. *IEEE Journal on Selected Areas in Communications*, 25(1), 73–83. doi:10.1109/JSAC.2007.070108.
 22. Ratnasamy, S., Hellerstein, J., & Shenker, S. (2003). “Range queries over DHTs,” IRB-TR-03-009, June, 2003, available at <http://berkeley.intel-research.net/sylvia/range.pdf>
 23. Zhang, R., & Hu, Y. (2005). “Assisted peer-to-peer search with partial indexing,” IEEE INFOCOM 2005.
 24. Chen, H. et al. (2008). “Efficient multi-keyword search over P2P web,” WWW 2008.
 25. “MARVEL: Multimedia Analysis and Retrieval System,” Intelligent Information Management Dept., IBM T. J. Watson Research Center.
 26. Christel, M., & Conescu, R. (2005). “Addressing the challenge of visual information access from digital image and video libraries,” *JCDE’05*. Denver, Colorado, June 7–11.
 27. Gibbon, D., Liu, Z., & Shahraray, B. (2006). “The MIRACLE video search engine,” CCNC 2006, Las Vegas, NV, Jan. 8–10.
 28. Tang, C., Xu, Z., & Mahalingam, M. (2002). *PeerSearch: efficient information retrieval in peer-peer networks*. Palo Alto: Hewlett-Packard Labs.
 29. Lu, J., & Callan, J. (2003). Content-based retrieval in hybrid peer-to-peer networks. *Proceedings of ACM CIKM’03*, New Orleans, LA, Nov.
 30. Yang, C. (2003). “Peer-to-peer architecture for content-based music retrieval on acoustic data,” *WWW 2003*, Budapest, Hungary, May 20–24.
 31. Lee, I., & Guan, L. (2004). “Content-based image retrieval with automated relevance feedback over distributed peer-to-peer network,” *ISCAS 2004*, Vancouver, Canada, May 23–26.
 32. Ardizzone, E., Gatani, L., Cascia, M., Re, G., & Ortolani, M. (2006). A P2P architecture for multimedia content retrieval. *Lecture Notes in Computer Science*, 4351.
 33. SAPIR, 2008. <http://www.sapir.eu/index.html>
 34. Salton, G., Wong, A., & Yang, C. S. (1975). A vector space model for automatic indexing. *Communications of the ACM*, 18(11), 613–620. doi:10.1145/361219.361220.
 35. Deerwester, S., Dumais, S., Furnas, G. W., Landauer, T. K., & Harshman, R. (1990). Indexing by latent semantic analysis. *Journal of the American Society for Information Science*, 41(6), 391–407. doi:10.1002/(SICI)1097-4571(199009)41:6<391::AID-ASII>3.0.CO;2-9.
 36. Latent semantic analysis on Wikipedia, available at http://en.wikipedia.org/wiki/Latent_Semantic_Indexing
 37. Tang, C., Xu, Z., & Dwarkadas, S. (2003). ‘Peer-to-peer information retrieval using self-organizing semantic overlay networks.’ In *Proceedings of ACM SIGCOMM*, pp 175–186, Karlsruhe, Germany, August 2003.
 38. Zhu, Y., & Hu, Y. (2004). ESS: Efficient Semantic Search on Gnutella-Like P2P Systems. *Technical Report, Department of ECECS, University of Cincinnati*, March 2004.
 39. Kundur, D., Luh, W.: *Encyclopedia of Multimedia*. Springer 2006, ch. Multimedia Sensor Networks.
 40. Thuraisingham, B. (2004). Security and privacy for sensor databases. *Sensors Letters*, 2(1), 37–47. doi:10.1166/sl.2004.022.
 41. Luh, W., & Kundur, D. (2006). “Distributed Privacy for Visual Sensor Networks via Markov Shares,” *Proc. 2nd DSSNS*, Columbia, MD, April 2006.
 42. Luh, W., Kundur, D., & Zourmos, T. (2007). “A Novel Distributed Privacy Paradigm for Visual Sensor Networks Based on Sharing Dynamical Systems,” *EURASIP Journal on Applied Signal Processing Special Issue on Visual Sensor Networks*, vol. 2007.



Madjid Merabti is Professor of Networked Systems and Director of the School of Computing & Mathematical Sciences, at Liverpool John Moores University, UK. He is a graduate of Lancaster University in the UK. He has over 20 years experience in conducting research and teaching in the areas of Computer Networks (fixed and wireless), Mobile Computing, and Computer Network Security. Prof. Merabti is widely published in these areas and leads the Distributed Multimedia Systems and Security Research Group, which has a number of UK Government, EU, and industry supported research projects. He is principal investigator in a number of current projects in Digital Rights Management, Games Technology, Multimedia Networking, Mobile Networks Security and Privacy Architectures and Protocols, Secure Component Composition in Ubiquitous Personal Networks, Networked Appliances, Mobile and Ad-Hoc Computing Environments, and Sensor Networks.

He is Associate Editor for IEEE Transactions on Multimedia, Area Editor for IEEE Communications Magazine series on Networked Appliances and Home Networking, Co-Editor in-Chief for the International Journal of Pervasive Computing and Communications (JPCC), Member of Editorial Boards for Springer Peer-to-Peer Networking and Applications Journal, Elsevier Journal on Computer Communications, and Wiley’s Communications and Security Journal. He was Guest Editor for the Special issue on “Research Developments in Consumer Communications and Networking”, *Multimedia Tools and Applications: An International Journal*, Kluwer, September 2005.

Madjid Merabti serves on the steering committee for the IEEE Consumer Communications and Networking (CCNC) series of conferences. He is Technical Programme Chair for: IEEE ICC Symposium on Special Areas of Communications, Germany 2009, and IASTED - EuroIMSA 2009 European Conference on Internet and Multimedia Systems and Applications. He is a member of a number of international conference programme committees on networking, security, networked appliances, digital rights management, and computer entertainment.



Zhu Liu received the B.S. and M.S. degrees in Electronic Engineering from Tsinghua University, Beijing, China, in 1994 and 1996, respectively, and the Ph.D. degree in Electrical Engineering from Polytechnic University, Brooklyn, NY, in 2001. He joined AT&T Labs - Research, Middletown, NJ, in 2000, and is currently a Principle Member of Technical Staff in the Video and Multimedia Technologies and Services Research Department. His research interests include multimedia content processing, multimedia databases, video search, pattern recognition, machine learning, and natural language understanding. He holds 8 U.S. patents and has published 1 book and more than 50 technical papers. Dr. Liu is a senior member of IEEE, and a member of ACM and Tau Beta Pi.

Dr. Liu and his colleagues won the best demonstration award in the Consumer Communication & Networking Conference 2007. He received the IP & Voice Services Research Excellence Award from AT&T in 2006, and he was selected as one of 100 outstanding young engineers to participate in the eleventh Frontiers of Engineering Symposium organized by National Academy of Engineering in 2005. He is on the editorial board of the IEEE Transaction on Multimedia and the Peer-to-peer Networking and Applications Journal. He has served as guest editor for the Multimedia Tools and Applications Journal and the International Journal of Semantic Computing. He was also on the organizing committee for the IEEE Consumer Communications & Networking Conference, the IEEE International Conference on Semantic Computing, and the IEEE International Symposium on Multimedia.



Heather Yu is a Senior Manager and the head of the Multimedia Content Networking research team at Huawei Technologies USA. She received her Ph.D. in electrical engineering from Princeton University in 1998.

Currently she is serving as Associate Editor-in-Chief of the Journal of Peer-to-Peer Networking and Applications, Chair of the new ComSoc technical subcommittee on Human Centric Communications, a voting member at the GLOBECOM/ICC Technical Content Committee, and a member of the Strategic Planning Committee at ComSoc. Her research interests include multimedia communications and multimedia content access and distribution. She has published 2 books, more than 60 technical papers and holds 23 U.S. patents.



Deepa Kundur received the B.A.Sc., M.A.Sc., and Ph.D. degrees all in Electrical and Computer Engineering in 1993, 1995, and 1999, respectively, from the University of Toronto, Canada. In January 2003, she joined the Department of Electrical Engineering at Texas A&M University, College Station, where she is a member of the Wireless Communications Laboratory and holds the position of Associate Professor. Before joining Texas A&M, she was an Assistant Professor at the Edward S. Rogers Sr. Department of Electrical and Computer Engineering at the University of Toronto where she was the Bell Canada Junior Chair-holder in Multimedia and an Associate Member of the Nortel Institute for Telecommunications.

Dr. Kundur's research interests include protection of scalar and broadband sensor networks, multimedia security, and computer forensics. She is an elected member of the IEEE Information Forensics and Security Technical Committee, vice-chair of the Security Interest Group of the IEEE Multimedia Communications Technical Committee and on the editorial boards of IEEE Communication Letters, IEEE Transactions on Multimedia, and EURASIP Journal on Information Security. She was recently a General Chair for the 2007 ACM Workshop on Multimedia and Security and a guest editor for the 2007 EURASIP Journal on Advances in Signal Processing Special Issue on Visual Sensor Networks. She has been a guest editor for the 2004 Proceedings of the IEEE Special Issue on Enabling Security Technologies for Digital Rights Management and the recipient of the 2005 Tenneco Meritorious Teaching award, the 2006 Association of Former Students College Level Teaching award, and the 2007 Outstanding Professor Award in the ECE Department at Texas A&M University.