ERCIM "Alain Bensoussan" Fellowship Scientific Report

Fellow: Christos Doulkeridis
Visited Location : NTNU
Duration of Visit: March 01, 2009 – February 28, 2010

I - Scientific activity

The research conducted during the fellowship period can be categorized in three broad categories: (1) query processing in peer-to-peer (P2P) systems, (2) information retrieval in P2P networks, and (3) novel query types for business analysis.

Query processing in P2P systems. Advanced query processing in P2P systems has not been adequately addressed in the relevant research literature. We have proposed novel algorithms for processing a wide variety of queries in P2P systems, including similarity search (range queries) [J2, J3, C1] and skyline queries [J1, C5]. Additionally, we proposed novel multidimensional routing indices for P2P systems [C3], which constitute an effective routing mechanism in the case of high-dimensional data stored in a super-peer network.

Information retrieval in P2P networks. A challenging problem when searching distributed and autonomous document repositories is the normalization of document frequencies, in order to merge the retrieved lists of documents into one ranked list effectively. For this purpose, we propose a hybrid method [S5] for aggregating document frequencies: a hierarchical aggregation for frequently occurring terms [C4] coupled with an approach for gossip-based aggregation for rarely occurring terms. Our experiments show that our technique is meaningful, as it produces accurate aggregate document frequencies, thus making the quality of document retrieval comparable to the centralized case. Our article [C4] describing this work has been invited to the *Information Systems* journal (Elsevier), as one of the best papers of the WISE'09 conference.

Novel query types for business analysis. Top-k queries have been widely established as a useful query that retrieves a ranked set of results that best match the individual user's interests. However, top-k queries have been mainly studied from the aspect of the customer that wishes to retrieve the k most suitable objects. In this work, for the first time, we study top-k queries from the perspective of the product manufacturer, and we introduce *reverse top-k queries* [C2], which identify for a given object, the users that have this object in their top-k result sets. We present efficient algorithms for handling this novel query type and we demonstrate their performance gain over baseline techniques by means of extensive experiments. This work has been invited for publication in the special issue of *IEEE Transactions on Knowledge and Data Engineering* of the best papers of ICDE'10.

II- Publication(s) during your fellowship

During the tenure of the fellowship, my research work has resulted in: 3 journal publications, 5 conference publications, and 5 papers that are currently under review. These are presented in detail below.

Refereed Journal publications

J1. Akrivi Vlachou, Christos Doulkeridis, Yannis Kotidis, Michalis Vazirgiannis.
Efficient Routing of Subspace Skyline Queries over Highly Distributed Data.
In *IEEE Transactions on Knowledge and Data Engineering (TKDE).*

Data generation increases at highly dynamic rates, making its storage, processing and update costs at one central location excessive. The P2P paradigm emerges as a powerful model for organizing and searching large data repositories distributed over independent sources. Advanced query operators, such as skyline queries, are necessary in order to help users handle the huge amount of available data. A skyline query retrieves the set of non-dominated data points in a multi-dimensional dataset. Skyline query processing in P2P networks poses inherent challenges and demands non-traditional techniques, due to the distribution of content and the lack of global knowledge. Relying on a super-peer architecture, we propose a threshold-based algorithm, called SKYPEER and its variants, for efficient computation of skyline points in arbitrary subspaces, while reducing both computational time and volume of transmitted data. Furthermore, we address the problem of routing skyline queries over the super-peer network and we propose an efficient routing mechanism, namely SKYPEER+, which further improves the performance by reducing the number of contacted super-peers. Finally, we provide an extensive experimental evaluation showing that our approach performs efficiently and provides a viable solution when a large degree of distribution is required.

J2. Christos Doulkeridis, Akrivi Vlachou, Kjetil Nørvåg, Yannis Kotidis, Michalis Vazirgiannis.
Efficient Search based on Content Similarity over Self-Organizing P2P Networks.
In *Peer-to-Peer Networking and Applications, Springer*, Vol. 3, Issue 1, pages 67-79, March 2010.

The advent of the World Wide Web has made an enormous amount of information available to everyone and the widespread use of digital equipment enables end-users (*peers*) to produce their own digital content. This vast amount of information requires scalable data management systems. Peer-to-peer (P2P) systems have so far been well established in several application areas, with file-sharing being the most prominent. The next challenge that needs to be addressed is (more complex) data sharing, management and query processing, thus facilitating the delivery of a wide spectrum of novel data-centric applications to the end-user, while providing high Quality-of-Service. In this paper, we propose a self-organizing P2P system that is capable to identify peers with similar content and intentionally assign them to the same super-peer. During content retrieval, fewer super-peers need to be contacted and therefore efficient similarity search is supported, in terms of reduced network traffic and contacted peers. Our approach increases the responsiveness and reliability of a P2P system and we demonstrate the advantages of our approach using large-scale simulations.

J3. Christos Doulkeridis, Akrivi Vlachou, Yannis Kotidis, Michalis Vazirgiannis.
Efficient Range Query Processing in Metric Spaces over Highly Distributed Data.
In *Special Issue: Data Management in Peer-to-Peer Systems, Distributed and Parallel Databases*, Vol. 26, Issue 2-3, pages 155-180, December 2009.

Similarity search in P2P systems has attracted a lot of attention recently and several important applications, like distributed image search, can profit from the proposed distributed algorithms. In this paper, we address the challenging problem of efficient processing of range queries in metric spaces, where data is horizontally distributed across a super-peer network. Our approach relies on SIMPEER, a framework that dynamically clusters peer data, in order to build distributed routing information at super-peer level. SIMPEER allows the evaluation of exact range and nearest neighbor queries in a distributed manner that reduces communication cost, network latency, bandwidth consumption and computational overhead at each individual peer. In this paper, we extend SIMPEER by focusing on efficient range query processing and providing recall-based guarantees for the quality of the result retrieved so far. This is especially useful for range queries that lead to result sets of high cardinality and incur a high processing cost, while the complete result set becomes overwhelming for the user. Our framework employs statistics for estimating an upper limit of the number of possible results for a range query and each super-peer may decide not to propagate further the query and reduce the scope of the search. We provide an extensive experimental evaluation of our framework and show that our approach performs efficiently, even in the case of high degree of distribution.

Refereed Conference publications

C1. Akrivi Vlachou, Christos Doulkeridis and Yannis Kotidis.
Peer-to-Peer Similarity Search based on M-Tree Indexing.
In Proceedings of *15th International Conference on Database Systems for Advanced Applications (DASFAA'10) (short paper)*, pp.269-275, Tsukuba, Japan, April 1-4, 2010.

Similarity search in metric spaces has several important applications both in centralized and distributed environments. In centralized applications, such as similarity-based image retrieval, usually a server indexes its data with a state-of-the-art centralized metric indexing technique, such as the M-Tree. In this paper, we propose a framework for distributed similarity search, where each participating peer stores its own data autonomously, under the assumption that data is indexed locally by peers using M-Trees. In order to support scalability and efficiency of search, we adopt a super-peer architecture, where super-peers are responsible for query routing. We propose the construction of metric routing indices suitable for distributed similarity search in metric spaces. We study the performance of the proposed framework using both synthetic and real data.

C2. Akrivi Vlachou, Christos Doulkeridis, Yannis Kotidis and Kjetil Nørvåg.
Reverse Top-k Queries.
In Proceedings of *26th IEEE International Conference on Data Engineering (ICDE'10)*, Long Beach, CA, March 1-6, 2010. **[Invited to IEEE TKDE, special issue of best papers of ICDE'10]**

Rank-aware query processing has become essential for many applications that return to the user only the top-k objects based on the individual user's preferences. Top-k queries have been mainly studied from the perspective of the user, focusing primarily on efficient query processing. In this work, for the first time, we study top-k queries from the perspective of the product manufacturer. Given a potential product, which are the user preferences for which this product is in the top-k query result set? We identify a novel query type, namely *reverse top-k query*, that is essential for manufacturers to assess the potential market and impact of their products based on the competition. We formally define reverse top-k queries and introduce two versions of the query, namely monochromatic and bichromatic. We first provide a geometric interpretation of the monochromatic reverse top-k query in the solution space that helps to

understand the reverse top-k query conceptually. Then, we study in more details the case of bichromatic reverse top-k query, which is more interesting for practical applications. Such a query, if computed in a straightforward manner, requires evaluating a top-k query for each user preference in the database, which is prohibitively expensive even for moderate datasets. In this paper, we present an efficient threshold-based algorithm that eliminates candidate user preferences, without processing the respective top-k queries. Furthermore, we introduce an indexing structure based on materialized reverse top-k views in order to speed up the computation of reverse top-k queries. Materialized reverse top-k views trade preprocessing cost for query speed up in a controllable manner. Our experimental evaluation demonstrates the efficiency of our techniques, which reduce the required number of top-k computations by 1 to 3 orders of magnitude.

C3. Christos Doulkeridis, Akrivi Vlachou, Kjetil Nørvåg, Yannis Kotidis and Michalis Vazirgiannis.
Multidimensional Routing Indices for Efficient Distributed Query Processing.
In Proceedings of *18th ACM Conference on Information and Knowledge Management (CIKM'09) (short paper)*, Hong-Kong, November 2-6, 2009.

Traditional routing indices in peer-to-peer (P2P) networks are mainly designed for document retrieval applications and maintain aggregated one-dimensional values representing the number of documents that can be obtained in a certain direction in the network. In this paper, we introduce the concept of *multidimensional routing indices* (MRIs), which are suitable for handling multidimensional data represented by minimum bounding regions (MBRs). Depending on data distribution on peers, the aggregation of the MBRs may lead to MRIs that exhibit extremely poor performance, which renders them ineffective. Thus, focusing on a hybrid unstructured P2P network, we analyze the parameters for building MRIs of high selectivity. We present techniques that boost the query routing performance by detecting similar peers and grouping and reassigning these peers to other parts of the hybrid network in a distributed and scalable way. We demonstrate the advantages of our approach using large-scale simulations.

C4. Robert Neumayer, Christos Doulkeridis and Kjetil Nørvåg.
Aggregation of Document Frequencies in Unstructured P2P Networks.
In Proceedings of *10th International Conference on Web Information Systems Engineering (WISE'09)*, Poznan, Poland, October 5-7, 2009. **[Invited to Information Systems journal, special issue of best papers of WISE'09]**

Peer-to-peer (P2P) systems have been recently proposed for providing search and information retrieval facilities over distributed data sources, including web data. Terms and their document frequencies are the main building blocks of retrieval and as such need to be computed, aggregated, and distributed throughout the system. This is a tedious task, as the local view of each peer may not reflect the global document collection, due to skewed document distributions. Moreover, central assembly of the total information is not feasible, due to the prohibitive cost of storage and maintenance, and also because of issues related to digital rights management. In this paper, we propose an efficient approach for aggregating the document frequencies of carefully selected terms based on a hierarchical overlay network. To this end, we examine unsupervised feature selection techniques at the individual peer level, in order to identify only a limited set of the most important terms for aggregation. We provide a theoretical analysis to compute the cost of our approach, and we conduct experiments on two document collections, in order to measure the quality of the aggregated document frequencies.

C5. Joao Batista da Rocha Junior, Akrivi Vlachou, Christos Doulkeridis and Kjetil Nørvåg.
AGiDS: A Grid-based Strategy for Distributed Skyline Query Processing.
In Proceedings of *2nd International Conference on Data Management in Grid and P2P Systems (Globe 2009)*, Linz, Austria, August 31 - September 4, 2009.

Skyline queries help users make intelligent decisions over complex data, where different and often conflicting criteria are considered. A challenging problem is to support skyline queries in distributed environments, where data is scattered over independent sources. The query response time of skyline processing over distributed data depends on the amount of transferred data and the query processing cost at each server. In this paper, we propose AGiDS, a framework for efficient skyline processing over distributed data. Our approach reduces significantly the amount of transferred data, by using a grid-based data summary that captures the data distribution on each server. AGiDS consists of two phases to compute the result: in the first phase the querying server gathers the grid-based summary, whereas in the second phase a skyline request is sent only to the servers that may contribute to the skyline result set asking only for the points of non-dominated regions. We provide an experimental evaluation showing that our approach performs efficiently and outperforms existing techniques.

Papers under review

S1. Akrivi Vlachou, Christos Doulkeridis, Kjetil Nørvåg, Yannis Kotidis.
Identifying the Most Influential Objects with Reverse Top-k Queries.

S2. Akrivi Vlachou, Christos Doulkeridis, Neoklis Polyzotis.
Skyline Query Processing over Joins.

S3. Christos Doulkeridis, Akrivi Vlachou, Kjetil Nørvåg, Yannis Kotidis, Neoklis Polyzotis.
Efficient Rank Join Processing in Highly Distributed Systems.

S4. Joao Batista da Rocha Junior, Akrivi Vlachou, Christos Doulkeridis and Kjetil Nørvåg.
SkyPlan: Planning the Execution of Skyline Queries in Distributed Environments.

S5. Robert Neumayer, Christos Doulkeridis and Kjetil Nørvåg.
A Hybrid Approach for Estimating Document Frequencies in Unstructured P2P Networks.

III -Attended Seminars, Workshops, and Conferences

1. Participation in the 26th IEEE International Conference on Data Engineering (ICDE'10), Long Beach, CA, March 1-6, 2010.
2. Participation in the 18th ACM Conference on Information and Knowledge Management (CIKM'09), Hong-Kong, November 2-6, 2009.
3. Participation in the 8th Hellenic Data Management Symposium (HDMS'09), Athens, Greece, August 31 – September 1, 2009.
4. Participation in the 35th International Conference on Very Large Data Bases (VLDB'09), Lyon, France, August 24-28, 2009.
5. Participation in the ACM International Conference on Management of Data (SIGMOD'09), Providence, Rhode Island, USA, June 29 – July 2, 2009.

IV – Research Exchange Programme (12 month scheme)

During the fellowship, I had the opportunity to visit two research centers as part of the exchange programme. In both occasions I gave a talk about my research and I received valuable feedback that helped me improve my work. More importantly, new collaborations have emerged that will hopefully lead to joint research in the near future.

The first visit was at Ecole Polytechnique Federale Lausanne (EPFL) in the Data-Intensive Applications and Systems (DIAS) laboratory, which is headed by Professor Anastasia Ailamaki. The visit took place in the period of 13–20 May, 2009. During the stay, I was able to give a presentation about my research results in the area of top-k query processing. I was also able to learn the current research conducted in the DIAS research group at EPFL, and exchange research ideas with professors and PhD students there. Additionally, I have initiated and try to establish research cooperation with people in the group, in order to pursue research directions of common interest.

The second visit during the exchange programme was at Aalborg University in the Center for Data-Intensive Systems (DAISY), headed by Professor Christian S.Jensen. The duration of the stay was 16-25 November, 2009. I gave a talk about the computation of top-k and skyline queries in distributed and parallel environments. I managed to start cooperating with Assistant Professor Ira Assent at DAISY center and we are currently working on a research paper related to skyline computation for high-dimensional data.