Workshop on Large-Scale Distributed Systems for Information Retrieval

Flavio P. Junqueira, Vassilis Plachouras
Yahoo! Research Barcelona
{fpj,vassilis}@yahoo-inc.com

Fabrizio Silvestri
ISTI-CNR, Pisa, Italy
fabrizio.silvestri@isti.cnr.it

Ivana Podnar
University of Zagreb, Croatia
ivana.podnar@fer.hr

Abstract

The Workshop on Large-Scale Distributed Systems for Information Retrieval was a venue for seminal ideas on the design of systems for search. The workshop focused mainly on mechanisms for P2P IR, which is currently a highly popular research area, but it also had fruitful discussions and presentations on other architectures for large-scale systems. Given the attendance and the good level of discussion, we conclude that systems for information retrieval is a growing and promising area of research.

1 Introduction

Large-scale distributed systems for information retrieval can use a number of different architectures, but recently two main architectures have the spotlight: P2P and federated data centers. P2P architectures are often based on open overlay networks that users can join and use. Their main goal is to foster the freedom of accessing data without having to search using commercial engines. In such search engines, users are free to join and contribute, and ideally all participants perform similar tasks and process similar workloads. In the second form, we have commercial search engines, such as Yahoo! and Google, that have been highly successful since the early days of the Web as they provide a simple interface to search for Web content. Such search engines make use of one or more data centers that collaborate to process user queries. Data centers often have a large number of machines, and their set up can be arbitrarily complex [1, 2].

The main goal of this venue was to attract researchers interested in such large-scale systems. As P2P architectures for search is currently a popular research topic in information retrieval, most of the discussion and the technical content presented were on this topic. More precisely, out of the six papers presented, five were on P2P, and the keynote talk also was on P2P applications. The panel at the end of the workshop, however, was centered around large-scale computations and a discussion on the quality of open search engines compared to commercial ones.

One interesting feature of the workshop was the poster presentation. Instead of having a session for posters, and making a distinction between posters and papers, we requested all authors of submitted papers to bring a poster that was hanging during the whole time of the workshop. This feature was interesting because authors could extend discussions on their work to coffee breaks and because people could have a glimpse on what the presentation would be about beforehand.

In the remainder of this report, we first present our brave research committee and the program of the workshop. We then make a few comments on the keynote, on the papers, and on the panel. We finish with some comments on lessons learned.

2 Program Committee and Program

The reviewers for this venue were:

- Karl Aberer, Ecole Polytechnique Fédérale de Lausanne
- Ricardo Baeza-Yates, Yahoo! Research, Spain and Chile,
- Fidel Cacheda, University of A Coruña, Spain
- Abdur Chowdhury, Summize, USA
- Norbert Fuhr, University of Duisburg-Essen, Germany
- Ronny Lempel, IBM Research, Israel
- Jie Lu, Carnegie Mellon University, USA
- Massimo Marchiori, University of Padova and UTILABS, Italy
- Iadh Ounis, University of Glasgow, UK
- Diego Puppin, ISTI - CNR, Italy
- Thomas Risse, L3S Research Center, Germany
- Luo Si, Purdue University, USA
- Umberto Straccia, ISTI - CNR, Italy
- Christos Tryfonopoulos, Max Planck Institute for Informatics, Germany
- Michalis Vazirgiannis, Athens University of Economics and Business, Greece
- Gerhard Weikum, Max Planck Institute for Informatics, Germany
- Pavel Zezula, Masaryk University, Czech Republic
- Justin Zobel, NICTA, Australia

Ophir Frieder gave the keynote talk, entitled: "On Peer-to-Peer Search Applications". In the continuation, we had six paper presentations in two technical sessions. The papers were:

- Wai Gen Yee, Linh Thai Nguyen, Ophir Frieder. "A View of the Data on P2P File-sharing Systems";
- Jie Yang, Jun Wang, Maarten Clements, Johan A. Pouwelse, Arjen P. de Vries, Marcel Reinders. "An Epidemic-based P2P Recommender System";

- Matthias Bender, Sebastian Michel, Peter Triantafillou, Gerhard Weikum. "Design Alternatives for Large-Scale Web Search: Alexander Great, Aeneas a Pioneer, and Anakin has the Force";
- Nuno Lopes, Carlos Baquero. "Taming Hot-Spots in DHT Inverted Indexes";
- Christian Zimmer, Christos Tryfonopoulos, Klaus Berberich, Gerhard Weikum, Manolis Koubarakis. "Node Behavior Prediction for Large-Scale Approximate Information Filtering";
- Maarten Clements, Arjen P. de Vries, Johan A. Pouwelse, Jun Wang, Marcel J.T. Reinders. "Evaluation of Neighborhood Selection Methods in Decentralized Recommendation Systems".

Finally, there was a panel with four participants: Ricardo Baeza-Yates (Yahoo! Research, Spain and Chile), Ophir Frieder (Illinois Institute of Technology, USA), Peter Boros (Google Inc., USA), and Vladimir Ofitserov (Yahoo! Inc., USA).

3 Peer-to-peer information retrieval

The invited talk and the six presented papers focused on two prominent topics related to P2P systems and IR, namely novel IR applications relying on P2P technology, and scalability and efficiency issues related to large-scale Web search over P2P networks. The applications discussed are search engines in P2P environments, decentralized recommender systems, and publish/subscribe systems for information filtering over P2P networks. The invited talk by Ophir Frieder motivated the discussion on P2P search: The P2P-based applications are currently increasingly popular and dominate the Internet traffic, while search accuracy in P2P file sharing applications is still rather limited. Ophir presented a number of schemes that leverage traditional IR methods to improve search performance in such environments.

The following talk, "A view of the Data on P2P File-sharing Systems", by Wai Gen Yee, complemented Ophir's presentation with a thorough analysis of the query properties and shared data characteristics within the Gnutella file-sharing system. After providing useful insights on how people share and search for files, they introduced IR-Wire, a publicly available tool for information retrieval research in P2P file sharing systems. The tool allows for collecting information on how P2P systems are used.

Two presentations discussed decentralized recommender systems. First, Jie Yang presented a novel view on how P2P can be used to build a collaborative filtering framework introducing an item ranking model inspired by the well-known Probability Ranking Principle (PRP) for information retrieval [4]. As it is pointed out in the paper, the distributed item ranking is realized by fully decomposing the computation loads of the model and preference data into the entire network. A detailed analysis of similarity functions for collaborative filtering in recommendation systems was presented by Maarten Clements in his paper entitled "Evaluation of neighborhood selection methods in decentralized recommendation systems". In this work, the authors propose a utility model to incorporate factors such as the usefulness and the confidence in the recommendations, The usefulness factor aims to improve the recommendations by favoring users with many recommendations, while the confidence factor relates to the variable number of available recommendations per item.

The last application presented during the technical sessions was on publish/subscribe style of information filtering, where users submit queries continuously and receive notifica-

tions when documents matching their interests are available. Christian Zimmer presented a ranking method which improves the peer selection process and information filtering performance within a P2P network in his talk on "Node Behavior Prediction for Large-Scale Approximate Information Filtering". The presented method improves existing methods for filtering by estimating on a per-node basis the parameters used in predicting whether a node will publish interesting information in the future.

Two of the presented papers discussed scalability and efficiency issues for large-scale Web search over P2P networks. In his talk on "Design Alternatives for Large-Scale Web Search: Alexander was Great, Aeneas a Pioneer, and Anakin has the Force", Matthias Bender introduced and analyzed three different architectures for distributed and potentially decentralized Web search engines. Load balancing problems have been discussed in this talk and especially term-based partitioning of the inverted index. The authors argued that only recent advances in distributed hash tables (DHT) and indexing may produce P2P architectures that could potentially be applied to Web search. For example, a solution for the load balancing problem in case of DHT-based inverted indexes was introduced by Nuno Lopez with "Taming Hot-Spots in DHT Inverted Indexes", who proposed a new distributed data structure based on a decentralized balanced tree.

4 Large-scale search engines

The paper by Matthias Bender *et al.* has also discussed non-P2P architectures for building a large-scale distributed IR system, as mentioned in the previous section. In particular, the solution proposed in the third alternative is based on the use of flash memory for storing the index, which is interesting because it enables random access to posting lists. Also, because of all current gadgets (cameras, cell phones, etc.) based on flash memory, it is becoming over time a very affordable option.

The paper contains interesting points, some of which initiated further discussion during the time for questions after the talk. In particular, there was an interesting discussion between Ricardo Baeza-Yates and Matthias Bender on the construction of Web-scale term partitioned inverted indexes. In one of the described architectures, the authors of the paper made the assumption that the construction of a term partitioned inverted index can be performed in linear time, and did not consider the potential complications from updating such an index.

During the panel, the discussion drifted from P2P, mainly due to the presence of a Yahoo! Senior Engineer, Vladimir Ofitserov. Vladimir has driven the discussion into large-scale distributed computation as it is very important for search engines to able to process large amounts of data efficiently and in a timely fashion. Vladimir mentioned in particular the efforts of Yahoo! on the development of Hadoop, a free implementation of the map-reduce model of parallel computation [3].

5 Open versus Commercial Search Engines

One of the interesting issues that emerged during the workshop and the panel was the comparison and relation between open and commercial search engines. The term "open search engines" refers to P2P architectures for searching, in which users are free to join and leave the system. In such systems, each user shares an amount of resources for running the system, for example, for keeping part of the index and searching it. On the other

hand, commercial search engines correspond to large-scale systems with enough capacity and resources to crawl and index a large part of the Web, as well as to provide a search service to millions of users across the world. The presentations and discussion during the workshop and the panel touched upon the potential for improvements in the performance of open search architectures, the incorporation of social context in the search process, as well as using open P2P architectures for providing Web-scale search services.

Ophir Frieder in his invited speech underlined that open search systems can benefit from leveraging the research output of the IR community. For example, he pointed out that the performance of a P2P file sharing system can be improved by merely enriching the representation of the shared files. This potential for improving performance is related to the fact that many open search systems based on P2P architectures have been developed in the past independently and outside of the scientific community, employing *ad-hoc* approaches to solve the problem of search over P2P architectures. This situation is analogous to the early days of Web search engines, when engineering of large-scale systems was seen as a greater challenge than their evaluation.

One of the differences mentioned between open and commercial search engines is the integration of social aspects in the search mechanism. In the case of centralized search engines, users would have to upload their profile to the search engine. There are, however, attempts to incorporate social aspects such as friendship relationships in distributed search architectures. The aim of such attempts is to employ the user's social context to improve search effectiveness. Jie Jang presented ongoing work on the Tribler system, a P2P system in which the users' social network is employed in a collaborative filtering setting to improve the recommendation of items.

A point that was made by the panelists was that commercial Web search engines are not likely to be replaced by open search engines based on P2P technologies, because of the scale and efficiency constraints imposed by the high expectations of users, who expect answers in a fraction of a second.

6 Final remarks

The workshop has been a successful venue for the discussion and presentation of fresh ideas on large scale distributed systems for information retrieval. Most of the contributions were related to P2P architectures and their applications for searching. One of the conclusions of the workshop regarding the issue of open search architectures and commercial search engines, is that they are likely to co-exist, focusing on applications with different requirements and goals.

We conclude that the study of large-scale systems for information retrieval reveals new problems related to both the low-level implementation of search engines, as well as to their scalability and efficiency in handling massive workloads from users spread all across the globe. Equally important is the study of problems related to processing massive amounts of data in decentralized systems as several mechanisms in current search engines rely upon data collected over time and across different data centers. Such problems have not been studied sufficiently so far and provide the opportunity for further research in the area.

References

[1] L. Barroso, J. Dean, U. Holzle. Web Search for a Planet: The Google Cluster Architecture. IEEE Micro, vol. 23, no. 2, pp. 22-28, Mar/Apr, 2003

[2] E. Brewer. Lessons from Giant-Scale Services. IEEE Internet Computing, vol. 05, no. 4, pp. 46-55, Jul/Aug, 2001

[3] Hadoop. The hadoop software platform for large-scale computing. `http://lucene.apache.org/hadoop/`, 2007.

[4] S. Robertson. The probability ranking principle in IR. *Journal of Documentation*, 33(4):294–304, 1977.