P2P Techniques
for Decentralized Applications

Synthesis Lectures on Data Management

Editor
M. Tamer Özsu, ***University of Waterloo***

Synthesis Lectures on Data Management is edited by Tamer Özsu of the University of Waterloo. The series will publish 50- to 125 page publications on topics pertaining to data management. The scope will largely follow the purview of premier information and computer science conferences, such as ACM SIGMOD, VLDB-ICDE, PODS, ICDT, and ACM KDD. Potential topics include, but not are limited to: query languages, database system architectures, transaction management, data warehousing, XML and databases, data stream systems, wide scale data distribution, multimedia data management, data mining, and related subjects.

P2P Techniques for Decentralized Applications
Esther Pacitti, Reza Akbarinia, and Manal El-Dick
2012

Query Answer Authentication
HweeHwa Pang and Kian-Lee Tan
2012

Declarative Networking
Boon Thau Loo and Wenchao Zhou
2012

Full-Text (Substring) Indexes in External Memory
Marina Barsky, Ulrike Stege, and Alex Thomo
2011

Spatial Data Management
Nikos Mamoulis
2011

Database Repairing and Consistent Query Answering
Leopoldo Bertossi
2011

Managing Event Information: Modeling, Retrieval, and Applications
Amarnath Gupta and Ramesh Jain
2011

Fundamentals of Physical Design and Query Compilation
David Toman and Grant Weddell
2011

Methods for Mining and Summarizing Text Conversations
Giuseppe Carenini, Gabriel Murray, and Raymond Ng
2011

Probabilistic Databases
Dan Suciu, Dan Olteanu, Christopher Ré, and Christoph Koch
2011

Peer-to-Peer Data Management
Karl Aberer
2011

Probabilistic Ranking Techniques in Relational Databases
Ihab F. Ilyas and Mohamed A. Soliman
2011

Uncertain Schema Matching
Avigdor Gal
2011

Fundamentals of Object Databases: Object-Oriented and Object-Relational Design
Suzanne W. Dietrich and Susan D. Urban
2010

Advanced Metasearch Engine Technology
Weiyi Meng and Clement T. Yu
2010

Web Page Recommendation Models: Theory and Algorithms
Sule Gündüz-Ögüdücü
2010

Multidimensional Databases and Data Warehousing
Christian S. Jensen, Torben Bach Pedersen, and Christian Thomsen
2010

Database Replication
Bettina Kemme, Ricardo Jimenez Peris, and Marta Patino-Martinez
2010

Relational and XML Data Exchange
Marcelo Arenas, Pablo Barcelo, Leonid Libkin, and Filip Murlak
2010

User-Centered Data Management
Tiziana Catarci, Alan Dix, Stephen Kimani, and Giuseppe Santucci
2010

Data Stream Management
Lukasz Golab and M. Tamer Özsu
2010

Access Control in Data Management Systems
Elena Ferrari
2010

An Introduction to Duplicate Detection
Felix Naumann and Melanie Herschel
2010

Privacy-Preserving Data Publishing: An Overview
Raymond Chi-Wing Wong and Ada Wai-Chee Fu
2010

Keyword Search in Databases
Jeffrey Xu Yu, Lu Qin, and Lijun Chang
2009

P2P Techniques for Decentralized Applications

Esther Pacitti, Reza Akbarinia, and Manal El-Dick

ISBN: 978-3-031-00760-6 paperback
ISBN: 978-3-031-01888-6 ebook

DOI 10.1007/978-3-031-01888-6

A Publication in the Springer series
SYNTHESIS LECTURES ON DATA MANAGEMENT

Lecture #25
Series Editor: M. Tamer Özsu, *University of Waterloo*
Series ISSN
Synthesis Lectures on Data Management
Print 2153-5418 Electronic 2153-5426

P2P Techniques for Decentralized Applications

Esther Pacitti
INRIA and Lirmm, University of Montpellier 2, France

Reza Akbarinia
INRIA and Lirmm, Montpellier

Manal El-Dick
Lebanese University

SYNTHESIS LECTURES ON DATA MANAGEMENT #25

ABSTRACT

As an alternative to traditional client-server systems, Peer-to-Peer (P2P) systems provide major advantages in terms of scalability, autonomy and dynamic behavior of peers, and decentralization of control. Thus, they are well suited for large-scale data sharing in distributed environments. Most of the existing P2P approaches for data sharing rely on either structured networks (e.g., DHTs) for efficient indexing, or unstructured networks for ease of deployment, or some combination. However, these approaches have some limitations, such as lack of freedom for data placement in DHTs, and high latency and high network traffic in unstructured networks. To address these limitations, gossip protocols which are easy to deploy and scale well, can be exploited. In this book, we will give a overview of these different P2P techniques and architectures, discuss their trade-offs and illustrate their use for decentralizing several large-scale data sharing applications.

KEYWORDS

large scale data sharing, peer-to-peer systems, DHT, unstructuted overlays, gossip protocols, top-k queries, recommendation, content sharing, caching, CDN, on-line communities, social-networks, information retrieval

Contents

Preface . . . xi

Acknowledgments . . . xiii

1 P2P Overlays, Query Routing, and Gossiping . . . 1

1.1 P2P Overlays . . . 1
1.1.1 Unstructured . . . 2
1.1.2 Structured . . . 2
1.1.3 Super-peer . . . 3
1.1.4 Comparing P2P Overlays . . . 3
1.2 Query Routing . . . 4
1.2.1 Query Routing in Unstructured Overlays . . . 5
1.2.2 Query routing in DHTs . . . 8
1.2.3 Query Routing in Super-Peers . . . 12
1.3 Gossip Protocols . . . 14
1.4 Replication . . . 17
1.5 Advanced Features on P2P Overlays . . . 18
1.5.1 Locality-Aware Overlays . . . 18
1.5.2 Interest-Based Overlays . . . 20
1.5.3 P2P Overlay Combination . . . 21
1.6 Conclusion . . . 22

2 Content Distribution in P2P Systems . . . 25

2.1 Introduction . . . 25
2.2 Insights on Traditional Content Distribution . . . 26
2.2.1 Background on Web Caching . . . 26
2.2.2 Overview of CDN . . . 27
2.2.3 Requirements and Open Issues of CDN . . . 27
2.3 P2P Content Distribution . . . 29
2.3.1 Advanced Features Used by Large-Scale P2P CDN . . . 29
2.3.2 P2P CDN Solutions . . . 31
2.4 Conclusion . . . 39

3 Recommendation Systems **41**

3.1 Overview of Recommendation 42

3.1.1 Collaborative Filtering 43

3.1.2 Content-based Filtering 44

3.1.3 Social Networks 45

3.2 P2P Content Management 46

3.2.1 Clustering Overlays 47

3.2.2 Short link overlay 48

3.3 P2P Recommendation 49

3.3.1 Basic P2P prediction 49

3.3.2 Social P2P Prediction Systems 51

3.4 Conclusion 54

4 Top-k Query Processing in P2P Systems **57**

4.1 General Model for Top-k Queries 58

4.2 Top-k Queries In Distributed Systems 61

4.3 Top-k Queries In P2P Systems 64

4.3.1 Top-k Queries in Unstructured Overlays 64

4.3.2 Top-k Queries in Super-peer Overlays 69

4.3.3 Top-k Queries in DHTs 70

4.4 Conclusion 72

Bibliography **75**

Authors' Biographies **89**

Preface

The Web 2.0 has brought a paradigm shift in how people use the Web. Before this Web evolution, users were merely passive consumers of content that is provided to them by a set of websites. In a nutshell, Web 2.0 offers an architecture of participation where individuals can participate, collaborate, share and create content. Web 2.0 applications deliver services that get better the more people use it, while providing their own content and remixing it with others content. Today, there are many emerging websites that have helped to pioneer the concept of participation in Web 2.0. Popular examples include the online encyclopedia Wikipedia that enables individuals to create and edit content (articles), social networking sites like Facebook, photo and video sharing sites like YouTube and Flickr, as well as wikis and blogs. Social networking is even allowing scientific groups to expand their knowledge base and share their theories which might otherwise become isolated and irrelevant.

With the Internet reaching a critical mass of users, Web 2.0 has encouraged the emergence of peer-to-peer (P2P) technology as a new communication model. The P2P model stands in direct contrast to the traditional client-server model, as it introduces symmetry in roles, where each peer is both a client and a server. Whereas a client-server network requires more investment to serve more clients, a P2P network pools the resources of each peer for the common good. In other terms, it exhibits the network effect as defined by economists: the value of a network to an individual user scales with the total number of participants. In theory, as the number of peers increases, the aggregate storage space and content availability grow linearly, the user-perceived response time remains constant, whereas the search throughput remains high or even grows. Therefore, it is commonly believed that P2P networks are naturally suited for handling large-scale applications, due to their inherent self-scalability. Since the late 1990s, P2P technology has gained popularity, mainly in the form of file sharing applications where peers exchange multimedia files. Chapter 1 covers the most relevant P2P concepts and overlays.

Under the Web 1.0 context, the content of web-servers is distributed to large audiences via Content Distribution Networks (CDN). The main mechanism is to replicate popular content at strategically placed and dedicated servers. As it intercepts and serves the clients queries, a CDN decreases the workload on the original web-servers, reduces bandwidth costs, and keeps the user-perceived latency low. Given that the Web is witnessing an explosive growth in the amount of web content and users, P2P networks seem to be the perfect match to build low cost infrastructures for content distribution. This is because they can offer several advantages like decentralization, self-organization, fault-tolerance and scalability. In a P2P system, users serve each other's queries by sharing their previously requested content, thus distributing the content without the need for powerful and dedicated servers. Chapter 2 presents an overview of P2P solutions for CDN decentralization over different P2P overlays.

More recently, P2P technologies have also been exploited for on-line communities, where participants are willing to post contents in order to share them. Interestingly, some on-line communities' participants prefer to keep and share their contents in their own workspace. For instance, in modern e-science, such as bio-informatics, physics and environmental science, scientists must deal with overwhelming amount of content (experimental data, documents, images, etc.) wishing to keep their contents in their own PC's instead of storing it in untrusted servers. Again, this seems a perfect match to P2P networks. P2P File-sharing systems have proven very efficient at locating content given specific queries. However, few solutions exist that are able to recommend the most relevant documents given a keyword-based query. This requires the use if recommendation methods. Chapter 3 presents some interesting P2P solutions for decentralized recommendation.

In very large-scale P2P systems, for each user's query there may be a huge number of answers most of which may be uninteresting for the user. Top-k queries have proved to be very useful to avoid overwhelming the user with large numbers of uninteresting answers. In addition, by filtering useless results they can significantly reduce the network traffic in P2P systems. By definition, a top-k query returns only the k data the most relevant to the users query. The relevance of data can be measured by a scoring function that the user specifies. In Chapter 4, we present some interesting approaches for top-k query processing in P2P networks.

A very interesting lecture on P2P Data Management can be found in Aberer [2010]. The authors focus on P2P management for data management, data integration and documents retrieval systems. Different from Aberer [2010], our goal is to show how different P2P technologies can be used generically for application decentralization focusing on Top-k, CDN and Recommendations systems.

Esther Pacitti, Reza Akbarinia, and Manal El-Dick
April 2012

Acknowledgments

We would like to acknowledge Fady Draidi for his very useful inputs for recommendation systems.