

LNCS 3367

Wee Siong Ng
Beng Chin Ooi
Aris Ouksel
Claudio Sartori (Eds.)

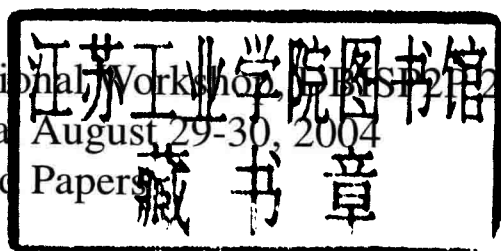
Databases, Information Systems, and Peer-to-Peer Computing

Second International Workshop, DBISP2P 2004
Toronto, Canada, August 2004
Revised Selected Papers

Wee Siong Ng Beng Chin Ooi
Aris Ouksel Claudio Sartori (Eds.)

Databases, Information Systems, and Peer-to-Peer Computing

Second International Workshop on DBISP2P 2004
Toronto, Canada August 29-30, 2004
Revised Selected Papers



Volume Editors

Wee Siong Ng
National University of Singapore
Singapore-MIT Alliance
4 Engineering Drive 3, Singapore, Malaysia
E-mail: ngws@comp.nus.edu.sg

Beng Chin Ooi
National University of Singapore
Department of Computer Science
School of Computing
Kent Ridge, Singapore 117543, Malaysia
E-mail: ooibc@comp.nus.edu.sg

Aris Ouksel
University of Illinois at Chicago
Department of Information and Decision Sciences
601 South Morgan Street, Chicago, IL 60607, USA
E-mail: aris@uic.edu

Claudio Sartori
University of Bologna
Department of Electronics, Computer Science and Systems
Viale Risorgimento, 2, 40136 Bologna, Italy
E-mail: claudio.sartori@unibo.it

Library of Congress Control Number: 2005921896

CR Subject Classification (1998): H.2, H.3, H.4, C.2, I.2.11, D.2.12, D.4.3, E.1

ISSN 0302-9743

ISBN 3-540-25233-9 Springer Berlin Heidelberg New York

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, re-use of illustrations, recitation, broadcasting, reproduction on microfilms or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable to prosecution under the German Copyright Law.

Springer is a part of Springer Science+Business Media

springeronline.com

© Springer-Verlag Berlin Heidelberg 2005
Printed in Germany

Typesetting: Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India
Printed on acid-free paper SPIN: 11404033 06/3142 5 4 3 2 1 0

Preface

Peer-to-peer (P2P) computing promises to offer exciting new possibilities in distributed information processing and database technologies. The realization of this promise lies fundamentally in the availability of enhanced services such as structured ways for classifying and registering shared information, verification and certification of information, content-distributed schemes and quality of content, security features, information discovery and accessibility, interoperation and composition of active information services, and finally market-based mechanisms to allow cooperative and non-cooperative information exchanges. The P2P paradigm lends itself to constructing large-scale complex, adaptive, autonomous and heterogeneous database and information systems, endowed with clearly specified and differential capabilities to negotiate, bargain, coordinate, and self-organize the information exchanges in large-scale networks. This vision will have a radical impact on the structure of complex organizations (business, scientific, or otherwise) and on the emergence and the formation of social communities, and on how the information is organized and processed.

The P2P information paradigm naturally encompasses static and wireless connectivity, and static and mobile architectures. Wireless connectivity combined with the increasingly small and powerful mobile devices and sensors pose new challenges to as well as opportunities for the database community. Information becomes ubiquitous, highly distributed and accessible anywhere and at any time over highly dynamic, unstable networks with very severe constraints on the information management and processing capabilities. What techniques and data models may be appropriate for this environment, and yet guarantee or approach the performance, versatility, and capability that users and developers have come to enjoy in traditional static, centralized, and distributed database environments? Is there a need to define new notions of consistency and durability, and completeness, for example?

This workshop concentrated on exploring the synergies between current database research and P2P computing. It is our belief that database research has much to contribute to the P2P grand challenge through its wealth of techniques for sophisticated semantics-based data models, new indexing algorithms and efficient data placement, query processing techniques, and transaction processing. Database technologies in the new information age will form the crucial components of the first generation of complex adaptive P2P information systems, which will be characterized by their ability to continuously self-organize, adapt to new circumstances, promote emergence as an inherent property, optimize locally but not necessarily globally, and deal with approximation and incompleteness. This workshop examined the impact of complex adaptive information systems on current database technologies and their relation to emerging industrial technologies such as IBM's autonomic computing initiative.

The workshop was collocated with VLDB, the major international database and information systems conference. It offered the opportunity for experts from all over the world working on databases and P2P computing to exchange ideas on the more recent developments in the field. The goal was not only to present these new ideas, but also to explore new challenges as the technology matures. The workshop provided also a forum to interact with researchers in related disciplines. Researchers from other related areas such as distributed systems, networks, multiagent systems, and complex systems were invited.

Broadly, the workshop participants were asked to address the following general questions:

- What are the synergies as well as the dissonances between the P2P computing and current database technologies?
- What are the principles characterizing complex adaptive P2P information systems?
- What specific techniques and models can database research bring to bear on the vision of P2P information systems? How are these techniques and models constrained or enhanced by new wireless, mobile, and sensor technologies?

After undergoing a rigorous review by an international Program Committee of experts, including online discussions to clarify the comments, 14 papers were finally selected. The organizers are grateful for the excellent professional work performed by all the members of the Program Committee. The keynote address was delivered by Ouri Wolfson from the University of Illinois at Chicago. It was entitled “DRIVE: Disseminating Resource Information in Vehicular and Other Mobile Peer-to-Peer Networks.” A panel, chaired by Karl Aberer from EPFL (Ecole Polytechnique Fédérale de Lausanne) in Switzerland, addressed issues on next-generation search engines in a P2P environment. The title of the panel was “Will Google2Google Be the Next-Generation Web Search Engine?”

The organizers would particularly like to thank Wee Siong Ng from the University of Singapore for his excellent work in taking care of the review system and the website. We also thank the VLDB organization for their valuable support and the Steering Committee for their encouragement in setting up this series of workshops and for their continuing support.

September 2004

Beng Chin Ooi, Aris Ouksel, Claudio Sartori

Organization

Program Chair

Beng Chin Ooi
Aris M. Ouksel
Claudio Sartori

National University of Singapore, Singapore
University of Illinois, Chicago, USA
University of Bologna, Italy

Steering Committee

Karl Aberer
Sonia Bergamaschi
Manolis Koubarakis
Paul Marrow

EPFL, Lausanne, Switzerland
University of Modena and Reggio-Emilia, Italy
Technical University of Crete, Crete
Intelligent Systems Laboratory,
BTexact Technologies, UK
University of Bologna, Cesena, Italy
University of Illinois, Chicago, USA
North Carolina State University, USA
University of Bologna, Italy

Gianluca Moro
Aris M. Ouksel
Munindar P. Singh
Claudio Sartori

Program Committee

Divyakant Agrawal
Boualem Benattallah
Peter A. Boncz

University of California, Santa Barbara, USA
University of New South Wales, Australia
CWI, Netherlands and Alex Delis Polytechnic
University, New York, USA

Fausto Giunchiglia
Manfred Hauswirth
Vana Kalogeraki
Achilles D. Kameas
Peri Loucopoulos
Alberto Montresor
Jean-Henry Morin
Gianluca Moro
Wolfgang Nejdl
Wee Siong Ng
Thu D. Nguyen
Evaggelia Pitoura

University of Trento, Italy
EPFL, Switzerland
University of California, Riverside, USA
Computer Technology Institute, Greece
UMIST Manchester, UK
University of Bologna, Italy
University of Geneva, Switzerland
University of Bologna, Italy
Learning Lab Lower Saxony, Germany
Singapore-MIT Alliance, Singapore
Rutgers University, USA
University of Ioannina, Greece

VIII Table of Contents

Dimitris Plexousakis	Institute of Computer Science, FORTH, Greece
Krithi Ramamritham	IIT, Bombay, India and Wolf Siberski University of Hannover, Germany
Peter Triantafillou	RA Computer Technology Institute and University of Patras, Greece
Ouri Wolfson	University of Illinois, Chicago, USA
Martin Wolpers	Learning Lab Lower Saxony, Germany
Aoying Zhou	Fudan University, China

Sponsoring Institutions

Microsoft Corporation, USA
Springer

Author Index

- Agrawal, Divyakant 61
- Böttcher, Stefan 166
- Batko, Michal 79
- Bressan, Stéphane 93
- Choe, Sungwon 215
- Cooper, Brian F. 46
- Cuenca-Acuna, Francisco Matias 200
- Dhraief, Hadhami 31
- El Abbadi, Amr 61
- Emekci, Fatih 61
- Frieder, Ophir 107
- Gennaro, Claudio 79
- Gunopulos, Dimitrios 136
- Kalogeraki, Vana 136
- Kang, Seungwoo 215
- Kemper, Alfons 31
- Koloniari, Georgia 16
- Lee, Hyonik 215
- Lee, Jinwon 215
- Liau, Chu Yee 93
- Martin, Richard P. 200
- Nejdl, Wolfgang 31
- Ng, Wee Siong 93
- Nguyen, Thu D. 200
- Ntarmos, Nikos 151
- Peery, Christopher 200
- Petrakis, Yannis 16
- Pitoura, Evaggelia 16
- Sahin, Ozgur D. 61
- Shu, Yanfeng 93, 122
- Song, Junehwa 215
- Türling, Adelhard 166
- Tan, Kian-Lee 93, 122
- Triantafillou, Peter 151
- Wiesner, Christian 31
- Wolfson, Ouri 1, 185
- Xu, Bo 1, 185
- Yee, Wai Gen 107
- Yin, Huabei 185
- Zeinalipour-Yazti, Demetrios 136
- Zezula, Pavel 79
- Zhou, Aoying 122

Table of Contents

Keynote Address

Data Management in Mobile Peer-to-Peer Networks

Bo Xu, Ouri Wolfson 1

Query Routing and Processing

On Using Histograms as Routing Indexes in Peer-to-Peer Systems

Yannis Petrakis, Georgia Koloniari, Evaggelia Pitoura 16

Processing and Optimization of Complex Queries in Schema-Based P2P-Networks

*Hadhami Dhraief, Alfons Kemper, Wolfgang Nejdl,
Christian Wiesner* 31

Using Information Retrieval Techniques to Route Queries in an InfoBeacons Network

Brian F. Cooper 46

Similarity Search in P2P Networks

Content-Based Similarity Search over Peer-to-Peer Systems

*Ozgur D. Sahin, Fatih Emekci, Divyakant Agrawal,
Amr El Abbadi* 61

A Scalable Nearest Neighbor Search in P2P Systems

Michal Batko, Claudio Gennaro, Pavel Zezula 79

Efficient Range Queries and Fast Lookup Services for Scalable P2P Networks

*Chu Yee Liau, Wee Siong Ng, Yanfeng Shu, Kian-Lee Tan,
Stéphane Bressan* 93

The Design of PIRS, a Peer-to-Peer Information Retrieval System

Wai Gen Yee, Ophir Frieder 107

Adaptive P2P Networks

Adapting the Content Native Space for Load Balanced Indexing

Yanfeng Shu, Kian-Lee Tan, Aoying Zhou 122

On Constructing Internet-Scale P2P Information Retrieval Systems <i>Demetrios Zeinalipour-Yazti, Vana Kalogeraki, Dimitrios Gunopulos</i>	136
AESOP: Altruism-Endowed Self-organizing Peers <i>Nikos Ntarmos, Peter Triantafillou</i>	151
Information Sharing and Optimization	
Search Tree Patterns for Mobile and Distributed XML Processing <i>Adelhard Türling, Stefan Böttcher</i>	166
Dissemination of Spatial-Temporal Information in Mobile Networks with Hotspots <i>Ouri Wolfson, Bo Xu, Huabei Yin</i>	185
Wayfinder: Navigating and Sharing Information in a Decentralized World <i>Christopher Peery, Francisco Matias Cuenca-Acuna, Richard P. Martin, Thu D. Nguyen</i>	200
CISS: An Efficient Object Clustering Framework for DHT-Based Peer-to-Peer Applications <i>Jinwon Lee, Hyonik Lee, Seungwoo Kang, Sungwon Choe, Junehwa Song</i>	215
Author Index	231

Data Management in Mobile Peer-to-Peer Networks¹

Bo Xu and Ouri Wolfson

Department of Computer Science, University of Illinois at Chicago
{boxu, wolfson}@cs.uic.edu

Abstract. In this paper we examine the database management of spatio-temporal resource information in mobile peer-to-peer networks, where moving objects communicate with each other via short-range wireless transmission. Several inherent characteristics of this environment, including the dynamic and unpredictable network topology, the limited peer-to-peer communication throughput, and the need for incentive for peer-to-peer cooperation, impose challenges to data management. In this paper we propose our solutions to these problems. The proposed system has the potential to create a completely new information marketplace.

1 Introduction

A mobile peer-to-peer network is a set of moving objects that communicate via short-range wireless technologies such as IEEE 802.11, Bluetooth, or Ultra Wide Band (UWB). With such communication mechanisms, a moving object receives information from its neighbors, or from remote objects by multi-hop transmission relayed by intermediate moving objects. A killer application of mobile peer-to-peer networks is resource discovery in transportation. For example, the mobile peer-to-peer network approach can be used to disseminate the information of available parking slots, which enables a vehicle to continuously display on a map to the driver, at any time, the available parking spaces around the current location of the vehicle. Or, the driver may use this approach to get the traffic conditions (e.g. average speed) one mile ahead. Similarly, a cab driver may use this approach to find a cab customer, or vice versa. Safety information (e.g. a malfunctioning brake light in a vehicle) can also be disseminated in this fashion.

A mobile peer-to-peer network can also be used in matching resource producers and consumers among pedestrians. For example, an individual wishing to sell a pair of tickets for an event (e.g. ball game, concert), may use this approach right before the event, at the event site, to propagate the resource information. For another example, a passenger who arrives at an airport may use this approach to find another passenger for cab-sharing from the airport to downtown, so as to split the cost of the cab. Furthermore, the approach can be used in social networks; when two singles whose profiles match are in close geographic proximity, one can call the other's cell phone and suggest a short face-to-face meeting.

¹ Research supported by NSF Grants 0326284, 0330342, ITR-0086144, and 0209190.

The approach can also be used for emergency response and disaster recovery, in order to match specific needs with expertise (e.g. burn victim and dermatologist) or to locate victims. For example, scientists are developing cockroach-sized robots or sensors that are carried by real cockroaches, which are able to search victims in exploded or earthquake-damaged buildings [4]. These robots or sensors are equipped with radio transmitters. When a robot discovers a victim, it can use the data dissemination among mobile sensors to propagate the information to human rescuers. Sensors can also be installed on wild animals for endangered species animal assistance. A sensor monitors its carrier's health condition, and it disseminates a report when an emergency symptom is detected. Thus we use the term moving objects to refer to all, vehicles, pedestrians, robots, and animals.

We would like to comment at this moment that in our model a peer does not have to be a moving object, and databases residing on the fixed network may be involved. In many cases there are both moving peers and fixed peers, and they collaborate in data dissemination. For example, a sensor in the parking slot (or the meter for the slot) monitors the slot, and, while unoccupied, transmits the availability information to vehicles nearby. Or all the slots in a parking lot may transmit the information to a fixed 802.11 hotspot via a wired network, and the hotspot announces the information. In either case, the vehicles that receive the information may propagate it to a wider area via the mobile peer-to-peer network approach. In such an environment the mobile peer-to-peer network serves as a supplement/extension to the fixed-site based solution.

Compared to static peer-to-peer networks and static sensor networks, mobile peer-to-peer networks have the following characteristics that present challenges to data management.

1. Dynamic, unpredictable, and partitionable network topology. In our environment the peers are physically mobile, and sometimes can be highly mobile (consider vehicles that move in opposite directions at 120 miles/hour relative speed). The traffic density can vary in a big range from rush hours to midnight. The underlying communication network is thus subject to topology changes and disconnections. Peer-to-peer or sensor network approaches that require pre-defined data access structures such as search routing tables used in Gridella [12], Chord [14] and spanning trees used in Cougar [15] and TinyDB [13, 21] are impractical in such an environment.

2. Limited peer-to-peer communication throughput. The communication throughput between two encountered peers is constrained by the wireless bandwidth, the channel contention, and the limited connection time. For example, previous investigations into Bluetooth links have suggested 2 seconds as a typical setup time between two unknown devices [22]. This gives less than 2 seconds for data transfer when two vehicles encounter each other at 120 miles/hour relative speed (assuming that the transmission range is 100 meters). The limited throughput requires that the communication be selective such that the most important data are communicated.

3. Need for incentive for both information supplier and information propagators. Like many other peer-to-peer systems or mobile ad-hoc networks, the ultimate

success of mobile peer-to-peer networks heavily relies on cooperation among users. In P2P systems, incentive is provided for peers to participate as suppliers of data, compute cycles, knowledge/expertise, and other resources. In mobile ad-hoc networks, incentive is provided for mobile hosts to participate as intermediaries/routers. In mobile peer-to-peer networks, the incentive has to be provided for participation as both suppliers and intermediaries (namely brokers).

The objective of our Dissemination of Resource Information in Vehicular Environments (DRIVE) project is to build a software platform that addresses the above issues and can be embedded within a hardware device attached to moving objects such as vehicles, personal digital assistants (PDAs), and sensors. The DRIVE platform consists of the following components:

1. Data Model. We introduce a unified data model for spatio-temporal resources in mobile peer-to-peer applications related to transportation, disaster recovery, mobile electronic commerce, and social networks. We illustrate how the data model can be used to represent various resource types even though these resource types are utilized in quite different ways.

2. Data Dissemination. We propose an *opportunistic* approach to dissemination of reports regarding availability of resources (parking slot, taxi-cab customer, dermatologist, etc.). In this approach, a moving object propagates the reports it carries to encountered objects, i.e. objects that come within transmission range; and it obtains new reports in exchange. For example, a vehicle finds out about available parking spaces from other vehicles. These spaces may either have been vacated by these encountered vehicles or these vehicles have obtained this information from other previously encountered ones. We call this paradigm *opportunistic peer-to-peer* (or OP2P).

3. Total Ordering of Resources by Relevance. With OP2P, a moving object constantly receives reports from the objects it encounters. If not controlled, the number of availability reports saved by an object will continuously increase, which will in turn increase the communication volume in future exchanges. Thus, to deal with the throughput challenge, we investigate techniques that prioritize the reports exchanged. These techniques provide a total rank in terms of relevance for all the reports across all the resource types stored in a moving object's reports database. The key issue is how to quantify the tradeoffs between the contributions of different attributes to the utility of a report.

4. Query Language and Query Processing. With OP2P, each peer m maintains a local reports database. The collection of the local databases of all the peers forms a virtual database to the database application in m . So the query language component and the query processing component deal with how to query this virtual database and how the query is processed.

5. Economic Model. Our incentive mechanisms are based upon virtual currency [5]. Each peer carries virtual currency in the form of a coin counter that is protected from illegitimate manipulation by a trusted and tamper resistant hardware module [6]. Each coin is bought for a certain amount of real money but it cannot be cashed for real money. We analyze the requirements to the economic model and propose possible solutions.

6. Information Usage Strategy. This component deals with how a resource consumer should use the received reports to take possession of a resource. This is important when the resource can only be exclusively used by one object at one time. Consider for example a driver who is looking for a parking slot. The driver may receive reports of multiple parking slots, and these parking slots may be in different orientation and distance with respect to the driver's current location. Then the question is which parking slot the driver should go to (namely, pursue).

7. Transaction Management. This component aims to study a spectrum of solutions to transactional and consistency issues that arise in report dissemination, and minimize dependence on any centralized structure.

All the components are divided into three layers as shown in Figure 1. The bottom is the *data* layer, which implements the data model for the spatio-temporal resources. Above the data layer is the *support* layer. This layer defines how the data is disseminated and how queries are processed. It also contains transaction management. The top is the *utility* layer, which contains the modules relevant to utilization of the resource information, including relevance evaluation, query language, economic model, and usage strategies.




Fig. 1. The architecture of DRIVE

The rest of the paper is organized as follows. Section 2 introduces the data model and report ordering. Section 3 discusses OP2P data dissemination. Section 4 presents the query language and discusses query processing. Section 5 discusses the economic model. Section 6 discusses information usage strategies and transaction management. Section 7 discusses relevant work. Section 8 concludes the paper.

2 Data Model

2.1 Resource Model

In our system, resources may be spatial, temporal, or spatio-temporal. Information about the location of a gas station is a spatial resource. Information about the price of a stock on 11/12/03 at 2pm is temporal. There are various types of spatio-temporal resources, including parking slots, car accidents (reports about such resources provide traffic-jam information), taxi-cab requests, ride-sharing invitations, demands of expertise in disaster situations, and so on. Formally in our model there are N resource

types T_1, T_2, \dots, T_N . At any point in time there are M resources R_1, R_2, \dots, R_M , where each resource belongs to a resource type. Each resource pertains to a particular point location and a particular time point, e.g. a parking slot that is available at a certain time, a cab request at a street intersection, invitation of cab-sharing from airport to downtown from a passenger wishing to split the cost of the cab, or the demand of certain expertise at a certain location at a certain time. We assume that resources are located at points in two-dimensional geospace. The location of the resource is referred to as the *home* of the resource. For example, the home of an available parking space is the location of the space, and the home of a cab request or a cab-sharing invitation is the location of the customer. For each resource there is a *valid duration*. For example, the valid duration of the cab request resource is the time period since the request is issued, until the request is satisfied or canceled. The valid duration of the cab-sharing invitation starts when the invitation is announced and ends when an agreement is reached between the invitation initiator and another passenger. A resource is *valid* during its valid duration.

Let us comment further about spatial resources, such as gas stations, ATM machines, etc. In these cases the valid duration is infinite. Opportunistic dissemination of reports about such resources is an alternative paradigm to geographic web searching (see e.g. [7]). Geographic web searching has generated a lot of interest since many search-engine queries pertain to a geographic area, e.g. find the Italian restaurants in the town of Highland Park. Thus instead of putting up a web site to be searched geographically, an Italian restaurant may decide to put a short-range transmitter and advertise via opportunistic dissemination. In mobile systems, this also solves some privacy concerns that arise when a user asks for the closest restaurant or gas station. Traditionally, the user would have had to provide her location to the cellular provider; but she does not need to do so in our scheme. In our scheme, the transmission between two vehicles can be totally anonymous.

2.2 Peers and Validity Reports

The system consists of two types of peers, namely fixed hotspots and moving objects. Each peer m that senses the validity of resources produces *validity reports*. Denote by $a(R)$ a report for a resource R . For each resource R there is a single peer m that produces validity reports, called the *report producer* for R . A peer may be the report producer for multiple resources. Each report $a(R)$ contains at least the following information, namely *resource-id*, *create-time*, and *home-location*. Resource-id is the identification of R that is unique among all the resources of the same type in the system; create-time is the time when report $a(R)$ is created (it is also the time when R is sensed valid); home-location is the home of R .

In the parking slots example, a sensor in the parking slot (or the meter for the slot) monitors the slot, and, when the slot becomes free, it produces a validity report. In the car accident example, the report is produced by the sensor that deploys the air-bag.

$a(R)$ may contain other information depending on the resource type of R . For example, a parking slot report may include the time limit of the parking meter; a single-matching request may include the sender's personal information such as occupation and age; and so on. We say that $a(R)$ is a type T_i report if R is a type T_i resource.

Let $a(R)$ be a type T_i report. At any point in time, a peer m is either a *consumer* or a *broker* of $a(R)$. m is a consumer of $a(R)$, and $a(R)$ is a *consumer report* to m , if m is attempting to discover or find a type T_i resource. m is a broker of $a(R)$ and $a(R)$ is a *broker report* to m , if m is not attempting to discover/find T_i but is brokering $a(R)$, i.e. the only purpose of m storing $a(R)$ is to relay it to other peers.

2.3 Reports Relations

There are two relations in the reports database of a peer m . One is the *consumer relation*, which stores all the reports that m knows about and for which m is a consumer. Another is the *broker relation*, which stores all the reports that m knows about and for which m is a broker. The two relations have a common object-relational schema. The schema contains three columns: (i) resource-type which indicates the type of the reported resource; (ii) resource-id; (iii) report-description, which is an abstract data type that encapsulates all the attributes of a report. All the report description data types inherit from a single data type called *AbstractReport*. *AbstractReport* contains two attributes, namely create-time and home-location. Thus every report description data type has these two attributes.

2.4 Report Relevance

Given the memory and communication-throughput constraints, it is desirable that the most important or useful reports are communicated during an encounter. One possible approach that appears to achieve this goal is that the receiver explicitly expresses the criteria for the reports it is interested in receiving. For example, "Give me all the reports $a(R)$ such that the distance between R and me is smaller than 1 mile and the age of $a(R)$ (i.e. the length of the time-period since the creation of $a(R)$) is less than 1 minute." However, this does not guarantee a total order of the reports; on the other hand such a total order is necessary to ensure that most relevant reports are exchanged first (such that if disconnection occurs before the exchange completes, the loss is minimal), and that the less relevant reports are purged from memory before more relevant ones.

Our approach is to rank all the reports in a peer's reports database in terms of their relevance or expected utility, and then the reports are communicated and saved in the order of their relevance. Or, the reports requested and communicated are the ones with a relevance above a certain threshold. The notion of relevance quantifies the importance or the expected utility of a report to a peer at a particular time and a particular location.

Let $a(R)$ be a type T_i report. The relevance of $a(R)$ to a consumer at time q and location p represents the importance or the expected utility of $a(R)$ to the consumer at q and p . The relevance of $a(R)$ to a broker at time q and location p represents the importance or the expected utility of $a(R)$ to future consumers of the report that the broker estimates it will encounter. The question is how to evaluate the relevance such as to provide a total order of all the reports across all the reports relations within a peer.

We consider reports ranking a multiple attribute decision making (MADM) problem [11]. We adopt a hierarchical weighting structure. At the first level of the

weighting hierarchy, each resource type T_i is assigned a weight (priority) that represents the importance of T_i relative to other resource types. At the second level, each attribute of T_i is assigned a weight that represents the importance of that attribute relative to other attributes of T_i . When ordering reports, each report is assigned a score that is a weighted aggregation of the normalized values of each attribute. Then the reports are sorted based on their scores.

3 Data Dissemination

We assume that each peer is capable of communicating with the neighboring peers within a maximum of a few hundred meters. One example is an 802.11 hotspot or a PDA with Bluetooth support. The underlying communication module provides a mechanism to resolve interference and conflicts. Each peer is also capable of discovering peers that enter into or leave out of its transmission range. Finally, each peer is equipped with a GPS system so that (i) the peer knows its location at any point in time and (ii) the clock is synchronized among all the peers.

When two peers m_1 and m_2 encounter each other, m_1 first requests consumer reports from m_2 with relevance (according to the relevance metadata of m_1) above the lowest relevance in m_1 's consumer relation. Then m_1 requests broker reports from m_2 with relevance above the lowest relevance in m_1 's broker relation.

We would like to emphasize that in our model, the interactions among peers are completely self-organized. The association between a pair of peers is established when they encounter each other and is ended when they finish the exchange or when they are out of the transmission range of each other. Other than this there is no other procedure for a peer to join or leave the network.

4 The Economic Model

In this section we introduce an economic model that stimulates peers to participate in report dissemination even if they are not interested in using a resource. The economic model needs to satisfy the following requirements:

It should handle two categories of reports, depending on whether the producer or the consumer pays for the reports. Reports that the owner is interested in advertising are *producer-paid*. Reports that the consumer is interested in knowing are *consumer-paid*. A resource may have both producer-paid and consumer-paid reports, if both the producer and the consumer are willing to pay for the reports. For example, reports that include the location of a gas station may be producer-paid because the gas station wishes to advertise them to neighboring vehicles. They may also be consumer-paid because a consumer may be willing to pay for a gas station report if he really needs one. Similarly for taxi-cab requests and reports of available parking slots.

1. It should consider peers that may be producers, consumers, and brokers. For consumer paid reports, both producers and brokers should be incentivized. For producer paid reports, brokers should be incentivized.