MOBI-DIC: MOBIle DIscovery of loCal Resources in Peer-to-Peer Wireless Network [*]

Hu Cao, Ouri Wolfson, Bo Xu and Huabei Yin

Department of Computer Science
University of Illinois at Chicago
{*hcao2, wolfson, boxu, hyin*}*@cs.uic.edu*

Abstract

In this paper we examine management of databases distributed among moving objects. The objects are interconnected by a Mobile Ad Hoc Network. Several inherent characteristics of this environment, including the dynamic and unpredictable network topology, the limited peer-to-peer communication bandwidth, and the need for incentive for peer-to-peer cooperation, impose challenges to data management. In this paper we discuss these challenges in the context of a database that represents resource information. The information is disseminated and queried by the moving objects in search of resources. We are currently building such a resource discovery engine called MOBI-DIC: MOBIle DIscovery of LoCal Resources.

MOBI-DIC will enable quick building of matchmaking or resource discovery services in many application domains, including social networks, transportation, mobile electronic commerce, emergency response, and homeland security. For example, in a large professional, political, or social gathering, the technology is useful to automatically facilitate a face-to-face meeting based on matching profiles. In transportation, MOBI-DIC incorporated in navigational devices can be used to disseminate to other similarly-equipped vehicles information about relevant resources such as free parking slots, traffic jams and slowdowns, available taxicabs, and ride sharing. In mobile electronic commerce, MOBI-DIC is useful to match buyers and sellers in a mall, or to disseminate information about a marketed product. In emergency response, MOBI-DIC can be used by first responders to support rescue efforts (locate victims, and match responder capability with needs) even when the fixed infrastructure is inoperative. In homeland security, sensors mounted on neighboring containers can communicate and transitively relay alerts to remote check-points.

1 Introduction

Mobile local search is a procedure in which a mobile user searches for local resources, i.e. resources that are in geographic proximity to the user. In many situations, the local resources that are of interest to mobile users

---

Bulletin of the IEEE Computer Society Technical Committee on Data Engineering

---

[*] This research is supported by NSF Grants 0326284, 0330342, ITR- 0086144, 0513736, and 0209190.

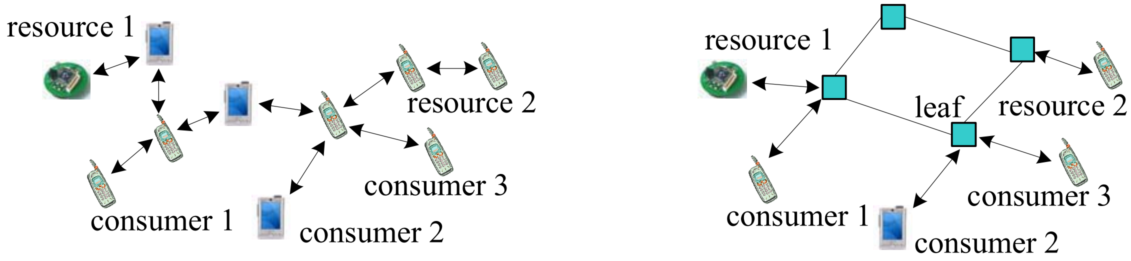

(a) Mobile ad hoc solution to local search. The nodes are sensors, PDA's, laptops, cell phones.

(b) Infrastructure based solution to mobile local search.

Figure 1: Two mobile local search solutions.

are only available during a limited period of time and these resources themselves may be mobile. For example, a cab driver wants to find a customer near by. The customer may be moving and she is available only until she hires a cab. Similarly, the current traffic speed on a road segment, the available parking slots around a driver, the available workstations in a large convention hall, are temporarily valid or available resources. We call these *spatio-temporal resources*.

Mobile local search for spatio-temporal resource is a special case of resource discovery (see e.g. [6, 9]) and publish/subscribe [14] applications, in the sense that the resources or events are spatio-temporal and limited in geographic scope. However, these characteristics make them less amenable to a centralized or hierarchical solution using the infrastructure. The reason is that installing a server in many local areas, and accessing these servers through the cellular infrastructure may be expensive. Furthermore, the temporary nature of the resources makes update and query response time critical, which again necessitates a large number of servers.

Therefore, in our project we explore a peer-to-peer (P2P) approach to local search that does not require a central server or a wireless infrastructure. However, when the infrastructure is available, the proposed P2P paradigm can augment it to make the search more efficient.

In our proposed approach (see Figure 1), a set of moving objects (denoting processors and sensors, some of which may be static) form a mobile ad-hoc network. Some moving objects are consumers and some are resources, and they communicate with each other via short-range wireless technologies such as IEEE 802.11, Bluetooth, or Ultra Wide Band (UWB). With such communication mechanisms, a moving object receives resource information from its neighbors, or from remote objects by multi-hop transmission relayed by intermediate moving objects. The cellular and mobile ad hoc approaches can be combined into an architecture in which resource-information disseminated in a mobile ad-hoc network augments the infrastructure by covering the areas that are not be covered by the hotspots, and it enhances the local search capability offered by the hotspots in areas that are covered by the infrastructure. In other words, the P2P approach can also be used to communicate among the leaves of a hierarchical cellular architecture (see Figure 1(a), further enhancing the search capability. In Figure 1(b), rectangles are leaves of a possibly fixed hierarchical infrastructure, each of which controls an area called a ”cell”.

We are currently building a software platform that supports mobile search and discovery of spatio-temporal resources in mobile ad-hoc networks, possibly in conjunction with the existing infrastructure. We call this platform MOBI-DIC (MOBIle DIscovery of loCal resources). MOBI-DIC is based on a local communication paradigm in a peer-to-peer network. Conceptually, the paradigm consists of neighbors exchanging reports, where each report is a resource description or a query. When a moving object m receives new reports, it incorporates them into its local reports database, and broadcasts the new reports database to its neighbors. Upon receiving the broadcast from m, each neighbor incorporates the received reports into its own local reports database, and broadcasts the new reports database. Thus reports transitively spread across network nodes. With this local broadcasting, the size of m's local database may continuously increase as new reports are generated. In order to limit the broadcast volume, we employ relevance functions that prioritize reports. Each moving object broadcasts only the k most relevant (top k) reports during a local broadcast. In other words, ranking is used to control flooding. We call this *paradigm rank-based broadcast (RBB)*.

It is important to note that rank-based broadcast reduces communication of unimportant and spam information. Furthermore, the user is shielded from unimportant, unwanted, and spam information by his/her query; in other words, the user is notified by his/her processor only of information that satisfies the user's query.

Compared to existing resource discovery protocols which rely heavily on constructed routing structures (see e.g. [6, 9, 14]), rank-based broadcast adapts better to high mobility and variance in network connectivity. Indeed, in a highly dynamic network that is susceptible to partitions, the constructed routing structure may easily become obsolete; whereas rank based broadcast does not rely on global structures, and a node does not need to know the global network topology. Furthermore, with the appropriate relevance function, rank-based broadcast provides spatial and temporal proximity awareness. In other words, the relevance of a report, and consequently its transmission probability, decays as it becomes older and it travels away from its point of origin.

MOBI-DIC will enable quick building of matchmaking or resource discovery services in many application domains including social networks, transportation, mobile electronic commerce, emergency response, and homeland security. For example, in a large professional, political, or social gathering, the technology is useful to automatically facilitate a face-to-face meeting based on matching profiles. In transportation, MOBI-DIC incorporated in navigational devices can be used to disseminate to other similarly-equipped vehicles information about relevant resources such as free parking slots, traffic jams and slowdowns, available taxicabs, and ride sharing. In mobile electronic commerce, MOBI-DIC is useful to match buyers and sellers in a mall, or to disseminate information about a marketed product. In emergency response, MOBI-DIC can be used by first responders to support rescue efforts even when the fixed infrastructure is inoperative; it will match specific needs with expertise (e.g. burn victim and dermatologist), and help locate victims. In homeland security, sensors mounted on neighbouring containers can communicate and transitively relay alerts to remote check-points.

In summary, MOBI-DIC provides mobile users a search engine for transient and highly dynamic information in a local geospatial environment. MOBI-DIC employs a unified model for both the cellular infrastructure and the mobile ad hoc environments. When the infrastructure is available, it can be augmented by the mobile ad hoc approach.

2 Concepts and Model

In MOBI-DIC there are *moving objects* and *resources* (also called events in publish-subscribe terminology). Each moving object m may generate resource descriptions for resources it learns or produces. A resource description for a resource R is denoted by $a(R)$. Each resource description $a(R)$ has a set of attributes. In addition to application specific attributes, all the resource descriptions for spatio-temporal resources have two common attributes, namely *resource-time*, and *resource-location*. Resource-time is the time when the resource description is generated. Resource-location is the location of R. If R is a mobile resource, then the location is a (possibly uncertain) trajectory that gives the current location as a function of time.

Each moving object m may also generate queries (or subscriptions in publish-subscribe terminology) which express m's interests in certain types of resources. A query may be associated with a function that computes the matching degree between the query and a resource description. We term a query and a resource description collectively as a *report*.

Each moving object m has a database that stores the reports that it has generated or has received from other moving objects. This is called the *reports database*. m is called the *producer* of the reports that it generates, the *consumer* of the reports that it is interested in receiving, and the *broker* of the reports it neither generates nor is interested in receiving. Each report has a *relevance* at location l and at time t that is determined by a *relevance*

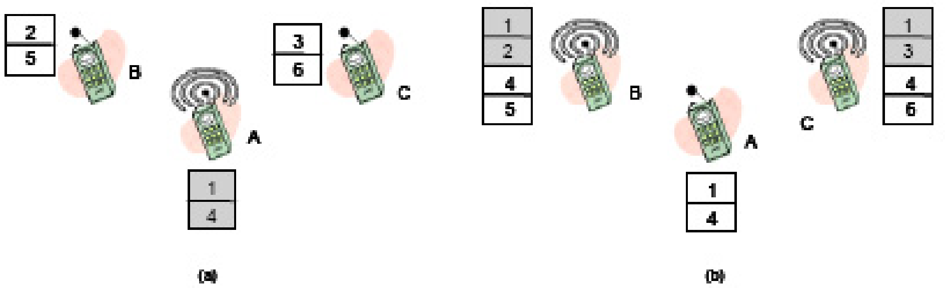

Figure 2: (a) Moving object *A* broadcasts its top two reports which are reports 1 and 4. (b) After receiving from *A*, object *B* incorporates the received reports, re-ranks, and broadcasts the top two (shadowed). The same for *C*.

function. The relevance of a report determines the priority of the report relative to other reports in the database, in allocating the available bandwidth of a particular broadcast. The relevances of all the reports are calculated for a specific location point and a specific time point, when it is determined that a broadcast has to be performed at those coordinates.

Observe that communication in mobile ad hoc networks is unreliable and disconnection may occur at any time, particularly in highly mobile environments such as vehicular networks. Furthermore, due to collisions in a dense environment, and due to energy considerations, the bandwidth available to each moving object may be limited. Therefore we postulate that ranking of reports is important, so that the most important reports are transmitted first. Therefore we propose a paradigm, *rank-based broadcast* (RBB), in which each moving object m periodically sorts the reports in its local database according to the relevance function, and broadcasts top k to its neighbors (i.e. the moving objects within the transmission range from m). The frequency of broadcasting and the size of each broadcast (i.e. k) depend on the available bandwidth and power. Figure 2 gives an example of the rank-based broadcast procedure.

Before ending this section we comment on the relationship between the RBB and flooding protocols. When energy and bandwidth are abundant, then RBB naturally reduces to flooding in the sense that each moving object caches all the reports that it receives, and relays them to each other moving object that it encounters. However, observe that as new resources are generated, it becomes increasingly unlikely that a moving object has enough power and bandwidth to relay all the reports it has ever received. Then a prioritization mechanism, such as that provided by RBB, becomes critical.

3 Results from Preliminary Work

In our prior work [29, 28, 26, 25, 27, 30, 31] we analyzed the way a report is propagated in geospace and time, the benefit of information dissemination in capturing competitive resources (i.e. resources that can be used by only one user at a time, such as a parking slot), how well the average local database reflects the status of physical resources, and how our approach compares with the central server model and with existing work on publish/subscribe in wireless ad-hoc networks.

3.1 Pattern of Report Propagation

In [29] we studied a simple situation where there is only one resource type in the system and the relevance is determined solely by a spatio-temporal function. With this function, the relevance of a report decreases as the time elapsed from its creation (i.e. age) increases, and as the distance from the resource increases. Our results show that when the memory of each moving object is limited and there are enough resources in the system, then the propagation of a report diminishes as the age and the distance increase; flooding is thus automatically limited to spatial proximity to the resource, and temporal proximity to the initial transmission-time of the report by the resource. Furthermore, the spatial and temporal boundaries automatically adapt depending on the number of resources in the system, the traffic density and speed, and other parameters that dictate the amount of storage, processing power, and bandwidth that should be allocated to each resource. For example, if the number of resources is small, each report will stay in the system longer, and spread farther.

3.2 Benefit of Resource Information

We analyzed the value of resource information in terms of how much time is saved when using the information to acquire competitive resources compared to when not using the information (see [28]). The results show that using information always reduces acquisition time compared to blind search (i.e. the information is not used). The value of information varies depending on the number of resources in the system, the traffic density and speed, and the wireless transmission range etc. In some cases by using the resource reports acquisition time is cut by more than 75%. We also developed information usage strategies to alleviate the herding effect, which occurs when a number of users all receive the same report and all decide to go to the same resource.

3.3 Accuracy of Information

We define the accuracy as the percentage of time when the report at a certain rank position is valid (i.e. the report correctly reflects the actual availability status of the reported resource). We evaluated by simulations the accuracy of information under different system environments (see [26]). The results show that in most situations, the accuracy of the first rank position is more than 85%.

We also compared by simulations the mobile ad-hoc approach and the central server approach in terms of accuracy. We found that with very reasonable object density and wireless transmission range, the accuracy of mobile ad-hoc approach reaches that of the central server approach. This indicates that the mobile ad-hoc approach could serve well as an alternative to the central server approach but with much less operational cost.

3.4 Comparison with Publish/Subscribe in Wireless Ad-hoc Networks

We compared RBB with Publish/Subscribe Tree (PSTree), a publish/subscribe algorithm designed for wireless ad-hoc networks (see [14]). We compared the two algorithms in terms of two aspects: (1) communication efficiency, namely the total relevance of resource descriptions that is delivered to consumers for each unit of bandwidth consumed; and (2) the total relevance of the resource descriptions delivered to consumers, regardless of communication cost,. The results show that in most cases RBB outperforms PSTree, in both communication efficiency and total relevance. We found that in a high mobility environment, since the network topology changes frequently, PSTree has to consume a lot of communication to update the routing structure. Furthermore, if the topology changes frequently, then relevant reports are lost anyway. However, RBB does not rely on global structures in the network, only on statistical information which is used to determine the relevance of a report.

4 Summary of the On-going Research

We are studying the following problems in this project:

1. Performance Measures. We are designing performance measures suitable for local search in a mobile ah-hoc network.

2. Report Ranking. We are investigating techniques that rank the reports in a local database. These techniques provide a total rank in terms of relevance-to-neighbors for the reports across all the resource types,

including queries, stored in a moving object's reports database. The key issue is how to quantify the relative contributions of different attributes to the utility of a report.

3. Rank-based Broadcast. We are studying variants of the rank-based broadcast conceptual model. These variants differ in their strategies of when to broadcast and what to broadcast.
4. Query Answering. We are studying how to deliver the match discovered at a broker to the query originator.
5. Power Management. We are incorporating power conservation mechanisms to RBB for mobile devices such as PDA and sensors.
6. Heterogeneous Environments. We are studying the RBB paradigm in environments where devices have different short range wireless capabilities, such as 802.11 and Bluetooth.
7. Augmenting the infrastructure. We are investigating methods for integrating mobile ad-hoc networks with cellular network.
8. Incentive Mechanisms. We are studying incentive mechanisms that stimulate moving objects to cooperate in report dissemination.
9. Security. We are studying solutions to various attacks from malicious or selfish parties. The solutions should minimally compromise privacy, and make minimum assumptions concerning tamper-resistant components.
10. Report Language. We are developing a language for the expression of resource description and queries.
11. Evaluation Tools. In order to evaluate our approach, we are developing three types of tools, including analytical models, a simulation testbed, and a proof-of-concept prototype.

5 Relevant Work

Resource Discovery and Publish/Subscribe Literature. The existing methods of resource discovery can be categorized into two types. Methods of one type are those relying on dedicated directory agents, e.g. SLP, Jini, Salutation, and UPnP. However, it is not clear how to make these agents available and accessible in a mobile ad hoc environment. Another type of resource discovery methods uses the peer-to-peer paradigm, however these solutions are based on fixed peers embedded in the infrastructure.

Resource discovery and publish/subscribe in mobile ad hoc networks are usually implemented by building a routing structure for resource information (see e.g. [6, 9, 14]). Most of this work uses the ”pull” mode, where a moving object searches the network for resources. As we have shown in the preliminary work, pull can be inefficient, particularly when the network topology is highly dynamic.

Work has also been done on data dissemination in mobile peer-to-peer networks [20, 21]. These methods use the gossiping/epidemic communication paradigm. However, they consider dissemination of regular data objects rather than spatio-temporal resources, and they do not rank the resources for determining what to broadcast.

Social Serendipity ([7]) is another related project. However, it focuses exclusively on social networks, whereas MOBI-DIC proposes a more general local search platform. Furthermore, the data dissemination and matchmaking in MOBI-DIC are totally decentralized, whereas Social Serendipity uses a central server approach.

Data Broadcasting. In data broadcasting [1, 2], data is periodically pushed from a server to a group of clients over a broadcast or multicast link. The major issue is how to schedule broadcast content to minimize the access time and tuning time of the client community. For example, the data disks [1, 2] scheme broadcasts hot (popular) data items with higher frequency than cool (unpopular) data items. We use the same idea in RBB. However, since we have a mobile ad-hoc environment and there are no dedicated servers, we enable the awareness of popularity by disseminating queries; In other words, we broadcast both data and queries. We divide the bandwidth capacity between resource descriptions and queries. Furthermore, in our model the importance of a resource description depends not only on its popularity, but also on many other factors such as its age and the resource location.

Mobile Ad-hoc Networks. The work in this area mainly concerns with sending a message to a specific destination given by the network address ((see [22] for a survey). In our case the network addresses of the destinations (i.e. consumers) are not known a priori. There is a body of work that deals with geographic routing (e.g. [15]) in which a message is routed from a source node to a geographic location or area. However, in most of the existing works in this area, message delivery is possible only if the source and destination are connected, namely there exists a path from the source to the destination. Li [18], Chen [5], and Vahdat [23] are among the first few works on routing in disconnected ad hoc networks. They either do not fit our context (e.g. [18] requires a moving object to actively move to reach the next object), or are too aggressive on bandwidth consumption (e.g. [23] uses flooding).

Mobile-to-mobile and Mobile-to-roadside Communication. CarTALK 2000 [4], VICS [24], and FleetNet [8] are projects aimed at designing, testing and evaluating co-operative driver assistance systems based upon vehicle-to-vehicle communication in order to improve the overall safety and convenience of the traffic participants. Infostations [10] is a mobile-to-roadside communication system that smartly allocates wireless channel depending on the positions of moving objects relative to infostations (i.e. hotspots). Our project can be considered an application that provides data management capabilities (data collection, organization, integration, modeling, dissemination, querying) on top of the underlying communication system provided by these systems.

Incentive Mechanisms. Incentive mechanisms have been studied in mobile ad-hoc networks (see e.g. [3, 33]) for stimulating intermediate nodes to forward messages to a given destination (rather than resource dissemination). We expect to be able to draw on existing ideas in this area. Incentive mechanisms have also been studied for static peer-to-peer networks (see e.g. [17]). In this case the static nature of the problem is often relied upon heavily, for example, by 'punishing' a user that is found noncooperative over time. Such a longer-time perspective is missing in our mobile ad-hoc environment, which may rarely involve the same pair of moving objects in an exchange.

Static Sensor Networks. A database approach has been applied to static sensor networks in Cougar [32], TinyDB [12], and direct diffusion [16]. Our distributed query processing relies on opportunistic interactions between mobile nodes, and is totally different.

References

[1] S. Acharya, R. Alonso, M. Franklin, and S. Zdonik. Broadcast Disks: Data Management for Asymmetric Communication Envrionments. *Proc. ACM SIGMOD Conf.*, San Jose, CA, May 1995.

[2] S. Acharya, M. Franklin, and S. Zdonik. Balancing Push and Pull for Data Broadcast. *Proc. ACM SIGMOD Conf.*, Tucson, AZ, May, 1997.

[3] L. Buttyan and J. Hubaux. Nuglets: a Virtual Currency to Stimulate Cooperation in Self-Organized Mobile Ad Hoc Networks. Technical Report DSC/2001/001, Swiss Federal Institute of Technology - Lausanne, Department of Communication Systems, 2001.

[4] CarTalk 2000, website: http://www.cartalk2000.net/

[5] X. Chen and A. L. Murphy. Enabling Disconnected Transitive Communication in Mobile Ad Hoc Networks, *ACM Principals of Mobile Computing, POMC'01*, 2001.

[6] S.M. Das, H. Pucha, and Y. Hu. Ekta: An Efficient DHT Substrate for Distributed Applications in Mobile Ad Hoc Networks. In *Proceedings of the 6th IEEE Workshop on Mobile Computing Systems and Applications (WMCSA 2004)*, English Lake District, UK, Dec. 2004.

[7] N. Eagle and A. Pentland. Social Serendipity: Mobilizing Social Software. *IEEE Pervasive Computing*, Special Issue: The Smart Phone. April-June 2005. pp 28-34.

[8] FleetNet-Internet on the Road project, website: http://www.fleetnet.de

[9] C. Frank and H. Karl. Consistency Challenges of Service Discovery in Mobile Ad Hoc Networks. In *Proceedings of the 7th ACM international symposium on Modeling, analysis and simulation of wireless and mobile systems*, pages 105-114, Venice, Italy, 2004.

[10] R.H. Frenkiel, B.R. Badrinath, J. Borras, and R. Yates. The Infostations Challenge: Balancing Cost and Ubiquity in Delivering Wireless Data. *IEEE Personal Communications Magazine*, 7(2): pp. 66-71, April 2000.

[11] I. Gruber, R. Schollmeier, and W. Kellerer. Performance Evaluation of the Mobile Peer-to-Peer Protocol. In *Proceedings of Fourth International Workshop on Global and Peer-to-Peer Computing*, Chicago, USA, Apr. 2004.

[12] J. Hellerstein, W. Hong, S. Madden and K. Stanek. Beyond Average: Toward Sophisticated Sensing with Queries. *The Second International Workshop on Information Processing in Sensor Networks*, 2003.

[13] Y. C. Hu and S. M. Das. Exploiting the Synergy between Peer-to-Peer and Mobile Ad Hoc Networks. In *HotOS-IX*, 2003.

[14] Y. Huang and H. Garcia-Molina. Publish/subscribe in a Mobile Environment. In *Proceedings of the 2nd ACM International Workshop on Data Engineering for Wireless and Mobile Access (MobiDE01)*, Santa Barbara, CA, May 2001.

[15] T. Imielinski, J. Navas. GPS-Based Geographic Addressing, Routing, and Resource Discovery. *Communications of the. ACM*, Vol. 42, No. 4, pp 86-92, 1999.

[16] C. Intanagonwiwat, R. Govindan and D. Estrin. Directed Diffusion: A Scalable and Robust Communication Paradigm for Sensor Networks. In *Proceedings of MobiCOM*, Boston, MA, 2000.

[17] K. Lai, M. Feldman, I. Stoica and J. Chuang. Incentives for Cooperation in Peer-to-Peer Networks. *1st Workshop on the Economics of P2P Systems*, 2003.

[18] Q. Li and D. Rus. Sending Messages to Mobile Users in Disconnected Ad-hoc Wireless Networks. *MOBICOM*, 2000.

[19] M. Mamei and F. Zambonelli. Location-based and Content-based Information Access in Mobile Peer-to-Peer Computing: the TOTA Approach. In *Second International Workshop on Agents and Peer-to-Peer Computing (AP2PC 2003)*, Melbourne, Australia, 2003.

[20] M. Papadopouli and H. Schulzrinne. Effects of Power Conservation, Wireless Coverage and Cooperation on Data Dissemination among Mobile Devices. In *MobiHoc*, 2001.

[21] F. Perich, A. Joshi, T. Finin, and Y. Yesha. On Data Management in Pervasive Computing Environments. *IEEE Transactions on Knowledge and Data Engineering*, 16(5), May 2004, pp. 621-634.

[22] E. Royer and C. Toh, A Review of Current Routing Protocols for Ad Hoc Mobile Wireless Networks, *IEEE Personal Communications*, pages 46-55, Apr. 1999.

[23] A. Vahdat and D. Becker. Epidemic Routing for Partially Connected Ad Hoc Networks, Technical Report CS-200006, Duke University, April 2000.

[24] Vehicle Information and Communication System Center (VICS Center), website: http://www.vics.or.jp/eng

[25] O. Wolfson, A. Ouksel, and B. Xu. Resource Discovery in Disconnected Mobile Ad-Hoc Networks. In *International Workshop on Next Generation Geospatial Information*, Cambridge, MA, 2003.

[26] O. Wolfson, B. Xu. Opportunistic Dissemination of Spatio-temporal Resource Information in Mobile Peer-to-Peer Networks. In *1st International Workshop on P2P Data Management, Security and Trust (PDMST'04)*, 2004.

[27] O. Wolfson and B. Xu. Data-on-the-road in Intelligent Transportation Systems. In *IEEE International Conference on Networking, Sensing, and Control (ICNSC 2004)*, Taipei, Taiwan, 2004.

[28] O. Wolfson, B. Xu, and H. Yin. Dissemination of spatial-temporal information in mobile networks with hotspots. In *Proceedings of the Second International Workshop On Databases Information Systems and Peer-to-Peer Computing*, 2004.

[29] B. Xu, A. Ouksel, O. Wolfson, Opportunistic Resource Exchange in Inter-vehicle Ad Hoc Networks. In *Proc. of 2004 IEEE International Conference on Mobile Data Management*, Berkeley, California, Jan. 2004.

[30] B. Xu and O. Wolfson. Privacy Solutions by Mobile Peer-to-Peer Communication. In *Center for Intellectual Property Law and Info. Tech., (CIPLIT) Symposium*, Chicago, Illinois, 2004.

[31] B. Xu, O. Wolfson. Data Management in Mobile Peer-to-Peer Networks. Springer Verlag Lecture Notes in Computer Science. In *Proc. of the 2nd International Workshop on Databases, Information Systems, and Peer-to-Peer Computing (DBISP2P'04)*, Toronto, Canada, Aug. 2004.

[32] Y. Yao and J. Gehrke. Query Processing in Sensor Networks. *First Biennial Conference on Innovative Data Systems Research*, 2003.

[33] S. Zhong, J. Chen and Y. Yang. Sprite: A Simple, Cheat-Proof, Credit-Based System for Mobile Ad-Hoc Networks. *In Proceedings of IEEE INFOCOM* 2003.