

Enabling Location Specific Real-time Mobile Applications

Ravi Kokku, Karthik Sundaresan, Guofei Jiang
NEC Laboratories America, Princeton, NJ, USA
{ravik, karthiks, gfj}@nec-labs.com

Abstract

As cellular networks open up and foster the deployment of richer user-centric mobile services, we envision the evolution of a new class of services that involve significant mobile user collaboration for making the services realizable and rapidly deployable. Our position is grounded on the intuition that mobile devices that users carry can act as “powerful” sensors for diverse information about a specific location, and can aid in providing real-time updates to other interested users that are not present at the location. Realizing such services is, however, challenging due to the need for several mechanisms such as (a) incentives for participation from users and network providers, (b) location-specific data collection and delivery and (c) determining trust in users providing updates. These mechanisms become significantly more complex to realize with user mobility. Unifying these mechanisms, in this paper, we propose a framework FLORA that enables rapid evolution of location-specific real-time services.

Categories and Subject Descriptors

C.2.1 [Computer-Communication Networks]: Network Architecture and Design—*Wireless Communication*

General Terms

Design, Economics, Performance

Keywords

Mobile, Applications, Open Networks, Open Services, Collaborative Services, Ubiquity, Real-time, Location, Framework, Incentives, Information Delivery

1. INTRODUCTION

Three interesting trends are emerging in the mobile network domain. (1) **Network Openness:** The recent push

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MobiArch'08, August 22, 2008, Seattle, Washington, USA.
Copyright 2008 ACM 978-1-60558-178-1/08/08 ...\$5.00.

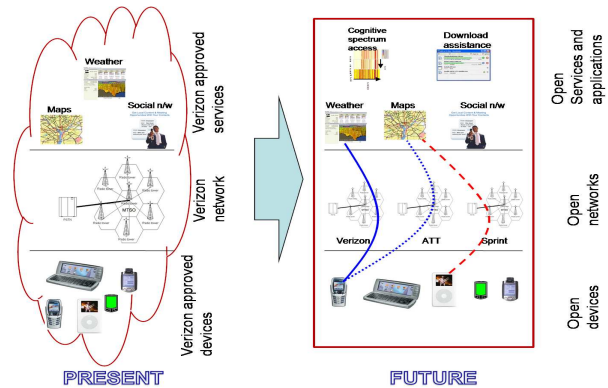


Figure 1: Towards open networks, open devices and open applications.

towards opening up wireless networks by decoupling services and applications from network providers nurtures rapid evolution of new application providers, which in turn provides greater service to cost ratio to the end customers. With such openness, we envision that various mobile applications will emerge quickly at least with the same scale and popularity as Internet services did in 1990s, mainly because the wired Internet was open. In contrast, today, customers are bound to the applications only provided by their own network providers, thereby being severely restricted in choices. Figure 1 represents the trend; in future, users will be able to choose any network dynamically to connect to services of their choice. (2) **Ubiquity:** Mobile devices (phones, headsets, gaming devices, etc.) have seen enormous growth in the recent years; just the number of mobile phone subscribers in the world has already hit 2.3 billion in 2006 and is expected to increase to 3.3 billion by 2011. With such a population size, any attractive mobile applications may lead to a big market and revenue. (3) **Resourcefulness:** While a mobile phone is used as a communication device today, we believe that it has a unique capability to serve as a location-specific sensing device as well due to its mobility with the user. With the massive number of mobile users, mobile devices can thus provide us enormous amount of spatio-temporal data in real-time that characterizes many of our physical activities.

We envision that the above trends enable a new class of applications that we call location-specific real-time ap-

lications. In these applications, a user U at a location B is interested in current information about location A . At the same time, there are users at location A that can potentially provide the necessary information to U . Examples of such location-specific real-time information include traffic situation, population density in a mall, live videos of an event such as a football game, radio spectrum availability (such as in opportunistic cognitive radio networks) and radio resource parameters (such as best base station to handoff, transmit power and bitrate) for efficient communication, etc. In effect, such real-time applications can be enabled easily by having user devices upload location-specific information as opposed to using dedicated sensor infrastructure.

A few such real-time applications already exist today and are developed individually [4, 6, 1, 7]. However, this class of applications have several commonalities that motivates us to explore a unified set of mechanisms that will enable us to construct more and novel applications rapidly. In particular, firstly, any such application requires a revenue model and incentive mechanisms to benefit each of the three participating entities: users, network providers and application service providers. Second, while mechanisms for data collection from user devices and application service delivery to users are *location-specific*, whereas mechanisms for incentives for user participation and trust in data uploaded by users are *user-specific*. Appropriately combining user-specific and location-specific components is necessary to make services successful, while keeping all participating entities happy. Finally, each of the above mechanisms varies slightly but remains mostly the same under different deployment models that differ in the amount of participation of each participating entity. In this paper, we develop FLORA, a Framework for Location-specific Real-time Applications that encompasses all the necessary mechanisms that are generally applicable in different deployment models. We elaborate several challenges exposed by user mobility and discuss why existing solutions in several domains with similar characteristics are not directly applicable.

The rest of the paper is organized as follows. Section 2 identifies several examples of applications that fall into the location-specific real-time category. Section 3 discusses three deployment models for realizing location-specific real-time applications. Section 4 describes our approach to build a framework that enables rapid deployment of such applications. Section 5 discusses related work and Section 6 concludes.

2. APPLICATIONS

By gathering real-time location specific information from large number of mobile phones, we can build new kind of mobile application services which fuse scattered data into useful information. In this section, we give several examples of such new mobile services to motivate the impact of the framework we propose.

1. **REAL-TIME TRAFFIC:** One can build a traffic report application that provides real-time congestion information on roads by tracking the positions of mobile users [5, 8, 14]. Several such applications are already being built. Essentially, two location coordinates within a span of say 1 second can tell us how fast the user is moving. However, we do have

to get rid of the effect of red lights and information from stopped cars etc. by taking samples from multiple users. This information can be aggregated by a server and sent to other users who are interested in the traffic situation at the current location (see Figure 2).

2. **PEOPLE DENSITY:** Using such information aggregated from many mobiles devices, one can also get an idea of how crowded a location hotspot is in real time. For example, such location hotspots could be shopping malls, movie theaters, picnic spots, gyms, swimming pools, city streets, etc. The information can also be aggregated over a longer timescale such as minutes to hours and used to determine how many people arrived at or departed from a parking lot, a movie theater or a stadium.
3. **SOCIAL NETWORKING:** Applications (a) giving indications of the availability of friends in one's social network within a chosen distance [4], (b) identifying and recommending previously unknown places or people of interest within a given area, (c) helping people to create new contacts during conferences, technical expos, etc., (d) match-making based on location and common interests, etc. Compared to existing social networking services, the location information from mobiles can enable social networking services to *dynamically* re-group people based on their current location and networking interests.
4. **SOS:** (a) Tracking lost people, even if they are unable to call, or their battery runs out. If position update is made as a fundamental requirement, a separate backup battery can be installed just for the purpose of position updates. (b) Mobiles equipped with nuclear radiation detectors can aid in national security [7]. In general, any additional sensors can be equipped on a mobile phone.
5. **TRACKING PUBLIC TRANSPORT:** Consider knowing an estimate of when a bus is going to arrive at a stop (e.g. [1]), and letting the bus driver know that a passenger is waiting close to the particular bus stop. The service above can enable people to stay at a nearby covered area such as a mall or a store, and let the passenger know the arrival time of the bus in real-time; such a service may make passengers more comfortable especially during extreme weather conditions. The passenger can also let the bus driver know that he/she intends to board the bus, so that the bus can wait "a little" longer at the stop.

3. DEPLOYMENT MODELS

Real-time location specific mobile services may be enabled under three deployment models that have varying amount of involvement from the users, network providers and service providers. These models have slightly different architectural requirements and challenges, but in general share a lot of mechanisms, thereby making a case of a unified framework like FLORA.

Service provider exclusive: In this model (DM1), a service provider collects information from the user devices




Figure 2: Traffic updates: An Example of Location Specific Real-time Services

through agents running on the devices, and provides services to the users. The users directly pay the service provider for the services. The underlying network merely serves as an access medium and hence will only influence the coverage and connectivity, for which the user pays separately. The advantage of this business model is that the separation of network providers from service providers leads to greater competition and innovation in services that leads to greater service satisfaction per unit cost for the users and also give them more flexibility to choose different service providers at fine timescales. However, there are at least two disadvantages of this approach: (1) Users have to explicitly participate in and make judgements on the implications of interacting with the service providers such as privacy issues, download/install and maintenance of service agents, etc., and (2) the customer base of network providers depends only on the coverage and connectivity and not on any additional value added services, thereby making it hard for them to differentiate from other network providers. Further, some services may be too resource intensive (e.g. P2P filesharing [12]) or be detrimental to the network provider business itself (e.g. Skype [2]). This in turn may discourage network providers from allowing the evolution of service providers.

Network provider exclusive: In this model (DM2), the network provider also acts as the service provider and provides services that are appealing to the users; the more attractive and the greater the number of services for a dollar paid by the customer, the greater will be the customer base. The advantage of this approach is that data necessary for providing services need not be exposed to a third-party service provider and hence needs no additional compromise in a user privacy. Further, the network provider can have greater control on the customer base since value-added services can be used to attract greater number of customers. The disadvantage of this approach, however, is that the network provider also has to invest in providing innovative services, and constantly increase the differentiation between services provided by others to maintain/improve the customer base. Consequently, the customer may be stuck with only a handful of services

provided by its own network provider and suffer from lack of choices and flexibility.

Collaborative: In this model (DM3), both the service and network providers share the burden of developing and deploying services and also share the revenue. A viable charging model is as follows: Network providers pay a lump-sum to the service providers for providing a particular service and also provide them with the necessary data. Since the network provider already collects large amounts of data for network resource management, and also has the necessary infrastructure for collecting extra data if necessary, the network provider may need minimum effort to enable services under this business model.

Note that one service provider can be interacting with several network providers for providing a service. Hence, a service provider can charge a network provider less if more data is provided by the network provider to make the service successful. A network provider, on the other hand, can put together services from several service providers and create appealing packages for attracting customers. The increase in customer base leads to increased revenue to the network providers. We believe that this model is a win-win situation for all the parties involved: users, network providers and service providers; consequently, this model may be the most prevalent model of service deployment in the near future.

4. FLORA

In this section, we propose FLORA, a Framework for Location-specific Real-time Applications, that encompasses the mechanisms required by the deployment models discussed in the previous section. We discuss several challenges and interesting issues for further research in realizing these mechanisms. Key to the proposal we make is the assumption that user devices have the capability to identify their location uniquely (through, GPS or assisted GPS [3] or base-station triangulation [9]). At a high level, Figure 3 shows the block-level architecture of the framework, which shows the minimum set of building blocks required: data collection, information delivery, incentives and trust, data service composition and data mining. White arrows represent data flow in models DM2 and DM3, where the network provider plays a significant role in defining the particular application. Dark arrows represent data flow in DM1 and DM3, where the service provider actively collects data using user agents. Dotted arrows represent the information flow from an application service to the mobile users. Finally, services can also be composed with information flow from other existing services.

4.1 Incentives

Arguably, the most interesting and challenging topic for exploration is the design of incentive mechanisms. Two factors influence the challenges. First, the data uploaded by users and the information delivered to the users through applications are both *location-specific*, whereas the incentives for participation of a user in uploading data is user-specific. Second, as users are mobile, the number of users and the specific users present at a particular location varies with time. Incentive mechanisms have to incorporate this dynamic behavior and create an appropriate bridging be-

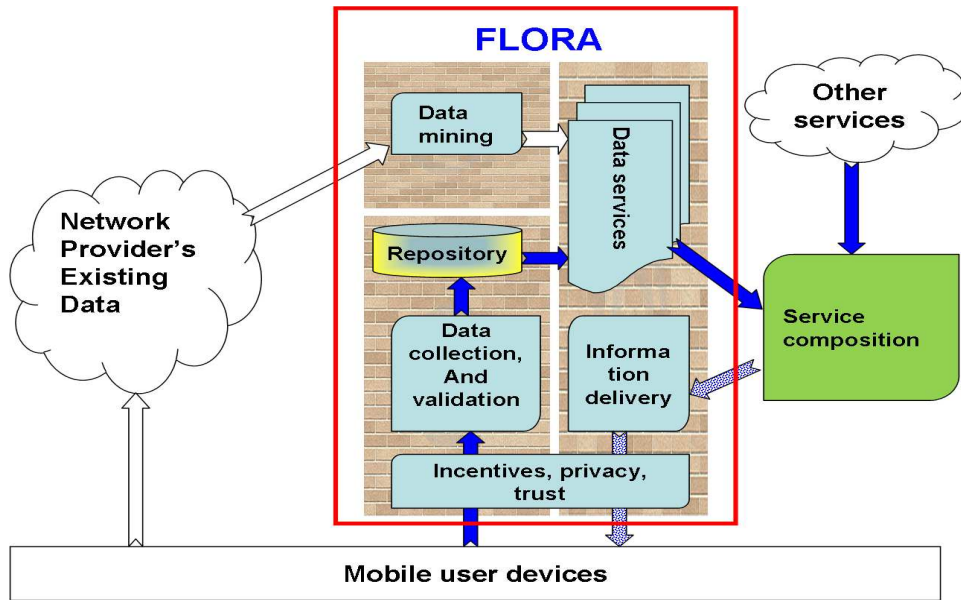


Figure 3: Block-level Architecture of the Framework for Location specific Real-time Services

tween user-specific and location-specific components of the framework. Further, while the incentive model is common to all the deployment models, its importance is more pronounced in the service-provider exclusive (DM1) and collaborative (DM3) models. In DM2, it is easier to build upon an already existing revenue model, especially since the network and service providers are the same in this case.

From the user participation perspective, while such incentive mechanisms have been explored in the past in several domains such as P2P file sharing [12], the incentive design problem in FLORA is different due to the following reasons.

- Diminishing returns for upload utility: the amount of location-specific data that needs to be uploaded per mobile device decreases with an increasing mobile population density in the location. So, the utility of an upload at a location where there are only a few other users is much higher than the utility when there are lot many users. Hence, incentives should incorporate the utility of the data provided by the users.
- Credit-based incentive policies: The dependence of utility on other user devices in a given location prevents the adoption of a pure tit-for-tat policy. Such soft credits per device may also need to be integrated with the conventional monetary charging system.
- Mobility: Finally, device mobility contributes to dynamics in the population density for a given location, requiring adaptation of the incentive model.

From the network provider perspective, especially in DM1 where the network provider only serves as an access medium, incentive design may require a monetary charging mechanism that involves either or both of users and service providers paying the network providers. E.g. Skype and/or

its users may pay Verizon an extra charge for enabling Skype calls on the phone network. The exact amount to charge again requires proper revenue model design to make it attractive to all the participating entities.

4.2 Data Collection

Since the data to be collected is location-specific, the service provider needs to configure information-upload parameters that are also location-specific, namely the definition of location such as its size, shape and granularity, identifying the specific set of users in the location, determining the data to be uploaded, rate of upload per user, etc. on a per-location basis. More dynamically, the information about users entering and leaving a location needs to be kept track of to configure the parameters on each user. However, the eventual upload of such data per user is closely tied to the incentive model. The importance of this component is more pronounced in DM1, where the wealth of existing network provider data is not readily accessible. At the same time, the underlying network resources should be used efficiently to minimize the cost paid to the network provider. These observations have several implications on the efficient collection of such location/region-specific data:

- Location-specific structure formation: service providers must form resource-efficient location-specific *trees* for collection of data.
- Localized group formation and data collection: when a group of mobile devices is to be chosen from a location/region to upload relevant information, the (i) size of the group, (ii) members of the group and (iii) amount of information (to be uploaded) from each member of the group has to be determined effectively.
- Mobility: due to device mobility and resulting variations in population density, the group size, membership and upload rate must be dynamically configured.

4.3 Information Delivery

Similar to the data collection problem, information delivery is also location-specific, and yields naturally to location-specific broadcast. However, the recipients of the information is determined based on the incentive model, and hence is user-specific. One way to realize this is to make user-specific unicast of data instead of broadcast. However, this approach may be resource inefficient, especially if the information delivery requires a lot of network resources. Another way to tie together the location-specific broadcasts and user-specific incentives is to enable information encoding and distribute user-specific keys that incorporate incentives appropriately. This approach, although being efficient in network resource usage, has several associated challenges:

- Location-specific key design and negotiation: The keys have to be associated with location-specific information, updated regularly and distributed to mobile users quickly to be able to decode the information being broadcast. The incentives (credits accrued) have to be modified as soon as we determine that a particular user used the credits for decoding certain information. Determination is not trivial because the information itself is broadcast, and the user may not respond back indicating which information is actually decoded and used.
- Resource efficiency: since information requested by mobile devices could be varying even in the same location/region, we may need to consider efficient merging of dissemination structures.
- Mobility: Finally, dissemination structures have to dynamically adapt to handle device mobility, such as the association of user-specific keys to the current location of the user.

The data delivery component is vital and is common to all the deployment models.

4.4 Trust Mechanisms

The above components highlight the challenges involved in the successful composition of services, under the assumption that the location-specific data collected is indeed representative of the request. However, this may not hold true owing to various reasons: (i) ambiguity/error in the information collection, (ii) deliberate user-tampering of data to serve one's own interests. Hence, it becomes necessary to validate the correctness of the data procured with respect to the service-specific objective using data mining [22], and multiple hypotheses tracking and testing [10]. The importance of this component is more pronounced in business models where the service provider is different from the network provider and hence new trust ties/mechanisms have to be established between the service provider and the users.

4.5 Data Mining

Most network providers already collect large volumes of data during their operation. Identifying and creating meaningful services out of that data can (1) create incentives for the network provider (mainly in DM3) to actively participate in delivering richer set of applications to users,

(2) reduce the burden of collecting data from the mobiles directly, and (3) bootstrap network operations when the services are first deployed in the collaborative model.

4.6 Data services

The data services component involves composing meaningful services from the collected data, and is also equally relevant to all the deployment models. It includes the following challenges:

- Parameter(s) recognition: we have to determine how to incorporate service-specific knowledge in determining the right set of data parameters that must be (i) probed from the mobile devices, and (ii) mined from the existing network data repository, and at the same time keep the services at a high level to make the data useful for other applications.
- Data longevity: an interesting area of research is how to process the data obtained and store in a more meaningful form in repositories such that only the optimum amount of information is stored for making even long term inferences over a period of time. Examples of such inferences could be data averages, user behavior patterns over a period of time, etc.

5. RELATED WORK

Several issues in FLORA have similarities to peer-to-peer (P2P) and pub-sub systems. Similar to file exchange in P2P systems [12], users upload location-specific information (to some service-provider infrastructure) useful for others, and in return receive required information from the infrastructure. Similar to pub-sub systems [15], and in particular location-based services, users subscribe to a certain service and receive location-specific updates from the system. However, FLORA differs in the following ways

- Vs. P2P: The table below shows a plausible mapping of concepts in P2P and FLORA. In FLORA, the utility of samples (information required to represent a location) reduces and flattens with increasing number of samples, whereas in P2P, the utility of number of copies of files flattens only when all nodes have a copy of the file. Due to the flattening utility of samples, for providing sufficient information about a location, the amount of information a user uploads is inversely proportional to the number of users present at the location. As a result, the design of incentives in FLORA becomes more complex due to user mobility; the amount of information uploaded is often asymmetric with the amount of downloaded information.
- Vs. publish-subscribe: In traditional publish subscribe systems, publishers are separated from subscribers. Whereas, in FLORA, subscribers also act as publishers. As a result, the design of incentives is important in FLORA and not in traditional pub-sub systems. Further, the presence of user mobility requires the design of quick association and disassociation mechanisms (to appropriate infrastructure access points) for efficiently delivering location specific information to mobile users.

	P2P	FLORA
1	Node	User device
2	File	Location
3	Copy of a file	A sample of location specific data
4	Popularity of a file	number of user devices around a location

Table 1: Comparison of concepts with P2P.

MoB (A Mobile Bazaar for Wide-area Wireless Services) [11] deals with an infrastructure for collaborative resource sharing among wireless devices at a given location. The architecture however does not consider mobility and its impact on sharing resources, which is the main focus of our proposal. At the same time, MoB also builds on the concept of "separating network providers from service providers" to enhance the deployment of services.

Several other efforts are in progress to enable richer set of applications for mobile users [13, 20, 24, 16, 21, 18, 23, 19]. Ott [20, 17] discusses the requirements of today's application protocols that assume a predominantly fixed and connected Internet, and identifies requirements to make them work in challenged mobile environments. Roberts et al. [21] discuss the difficulty for mobility due to the traditional Internet architecture, and proposes a new architecture to support mobility in the global Internet. Mao et al [17] present a data-centric mobility infrastructure that is flexible and extensible to meet the requirements of evolving mobile applications and services. Matos et al. [18] discuss embedding identity information into all layers of a protocol stack in mobile environments. All these works discuss complementary mechanisms that are also required to make FLORA a reality.

6. CONCLUSION

In this paper, we propose FLORA, a framework for location specific real-time services. We are currently exploring the design issues in this framework in detail, and also building a set of test applications to learn the challenges in realizing the mechanisms. FLORA exploits three recent trends in the mobile network domain: network openness, ubiquity of mobile devices, and expanding capabilities of mobile devices to serve as location-specific sensors. We believe that FLORA enables rapid evolution of location-specific real-time services due to (1) significant reduction in infrastructure requirements by utilizing mobile devices for obtaining location-specific information, and (2) the creation of appropriate incentives and mechanisms for efficient information collection and delivery, and for encouraging participation of users, network providers and service providers.

7. REFERENCES

- [1] <http://www.whereismybus.com/>.
- [2] <http://www.skype.com/>.
- [3] Assisted gps. www.gpsworld.com/gpsworld/article/articleDetail.jsp?id=12287.

- [4] The geopepper mobile social networking trial service. <http://www.geopepper.com>.
- [5] IntelliOne RoadSpeed Data Solutions.
- [6] Jasper wireless. <http://www.jasperwireless.com>.
- [7] Mobile phones could track bombs through radiation detectors. <http://www.mobiletor.com/2008/02/08/mobile-phones-could-track-bombs-through-radiation-detectors/>.
- [8] Nokia and uc berkeley monitors highway traffic. <http://www.mobiledia.com/news/64885.html>.
- [9] Triangulation. developers.sun.com/mobility/midp/articles/telematics/.
- [10] Y. Bar-Shalom. *Tracking and data association*. Academic Press Professional, Inc., San Diego, CA, USA, 1987.
- [11] R. Chakravorty, S. Agarwal, S. Banerjee, and I. Pratt. A mobile bazaar for wide-area wireless services. *Wirel. Netw.*, 13(6):757–777, 2007.
- [12] B. Cohen. Incentives Build Robustness in BitTorrent. In <http://www.bittorrent.com/bittorrentecon.pdf>.
- [13] D. Hadaller, S. Keshav, and T. Brecht. Mv-max: improving wireless infrastructure access for multi-vehicular communication. In *CHANTS '06: Proceedings of the 2006 SIGCOMM workshop on Challenged networks*, pages 269–276, New York, NY, USA, 2006. ACM.
- [14] B. Hellinga and P. Izadpanah. An opportunity assessment of wireless monitoring of network-wide road traffic conditions. Technical report, University of Waterloo, 2007.
- [15] Y. Huang and H. Garcia-Molina. Publish/subscribe in a mobile environment. *Wireless Networks*, 10(6), November 2004.
- [16] Y. Luo, O. Wolfson, and B. Xu. "a spatial-temporal approach for selective data dissemination in mobile peer-to-peer networks". In *Proc. of the Third International Conference on Wireless and Mobile Communications (ICWMC)*, 2007.
- [17] Y. Mao, B. T. Loo, Z. Ives, and J. Smith. The case for a unified extensible data-centric mobility infrastructure. In *ACM MOBIARCH*, 2007.
- [18] A. Matos, S. Sargento, and R. Aguiar. Embedding identity in mobile environments. In *ACM MOBIARCH*, 2007.
- [19] V. Navda, A. P. Subramanian, K. Dhanasekaran, A. Timm-Giel, and S. Das. Mobisteer: using steerable beam directional antenna for vehicular network access. In *MobiSys '07: Proceedings of the 5th international conference on Mobile systems, applications and services*, pages 192–205, New York, NY, USA, 2007. ACM.
- [20] J. Ott. Application protocol design considerations for a mobile internet. In *ACM MOBIARCH*, 2006.
- [21] P. Roberts and J. Kempf. Mobility architecture for the global internet. In *ACM MOBIARCH*, 2006.
- [22] P.-N. Tan, M. Steinbach, and V. Kumar. *Introduction to Data Mining, (First Edition)*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 2005.
- [23] O. Wolfson, B. Xu, and R. M. Tanner. "mobile peer-to-peer data dissemination with resource constraints". In *Proc. of the 8th International Conference on Mobile Data Management*, 2007.
- [24] B. Xu, O. Wolfson, and N. Rische. "benefit and pricing of spatio-temporal information in mobile peer-to-peer networks". In *Proc. of Hawaii International Conference on System Sciences (HICSS-39)*, 2006.