

Dissemination of Spatial-Temporal Information in Mobile Networks with Hotspots¹

Ouri Wolfson, Bo Xu, Huabei Yin

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

Abstract. In this paper we examine the dissemination of reports about resources in mobile networks with hotspots, where hotspots, vehicles and sensors communicate with each other via short-range wireless transmission. Each disseminated report represents information about a spatial-temporal event, such as the availability of a parking slot at a particular time or the detection of an injured in an earthquake damaged building. We propose an opportunistic dissemination paradigm, in which a moving object transmits the reports it carries to encountered peers and obtains new reports in exchange. We address two issues in such an environment. First, we develop an architecture that allows a moving object to receive resource reports opportunistically. Second, we study how the received reports are used by a consumer to reduce resource discovery time. The proposed system has the potential to create a completely new information marketplace.

1 Introduction

Consider an urban area with hundreds of thousands of vehicles. Drivers and passengers in these vehicles are interested in information relevant to their trip. For example, a driver would like his/her vehicle to continuously display on a map, at any time, the available parking spaces around the current location of the vehicle. Or, the driver may be interested in the traffic conditions (e.g. average speed) one mile ahead. Such information is important for drivers to optimize their travel, to alleviate traffic congestion, or to avoid wasteful driving. The challenge is processing queries in this highly mobile environment, with an acceptable delay, overhead and accuracy. One approach to solving this problem is maintaining a distributed database stored at fixed sites that is updated and queried by the moving vehicles via the infrastructure wireless networks. Potential drawbacks of this approach are (i) the responses to queries may be outdated, (ii) the response time may not meet the real-time requirements, and (iii) access to infrastructure communication service is costly, (iv) currently there is no business model to provide a return-on-investment for setting up and operating the fixed sites, and (v) the solution is not robust; particularly, it is vulnerable to failures of the fixed servers. In this paper we explore a new paradigm that is based on peer-to-peer communications.

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

We assume that each moving object (e.g. vehicle) has processing power (see [14]), and the capability of communicating with its neighbors. This communication can be enabled by a local area wireless protocol such as IEEE 802.11 [8], Ultra Wide Band (UWB) [22], or CALM [6]. These protocols provide broadband (typically tens of Mbps) but short-range (typically 50-100 meters) peer-to-peer communication. These communication capabilities exist already in experimental projects [23, 2] and are being planned for deployment on a large scale environment [9]. Similar communication capabilities are being planned between sensors in the infrastructure and moving vehicles [9]. With such communication mechanisms, a moving object receives the desired information from its neighbors, or from remote objects by multi-hop transmission relayed by intermediate moving objects. Thus, resource dissemination is performed in a *mobile peer-to-peer network*.

Compared to the traditional fixed-site based information query, this paradigm (i of the vehicles, by the vehicles, for the vehicles) has the following advantages. First, it provides better information authenticity, accuracy, and reliability, especially for real-time information. Consider for example parking space availability. Information collected from a vehicle that is leaving a parking slot tends to be more reliable than that from the fixed site. Second, it is free of charge, assuming that vehicles are willing to relay messages for free (in exchange for their messages being relayed). A back of the envelope calculation reveals that the cost (in terms of fuel) of communicating with encountered vehicles is less than a cent per day, even if the communication is continuous throughout the day.

The mobile peer-to-peer approach 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, the approach can be used in singles matchmaking; when two singles whose profiles match are in close geographic proximity, then one can call the other's cell phone and suggest a short face-to-face meeting.

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 injured. For example, scientists are developing cockroach-sized robots that are able to search victims in exploded or earthquake-damaged buildings [18]. These robots are equipped with radio transmitters. When a robot discovers a victim, it can use the mobile peer-to-peer approach to disseminate the discovery, and home in on the target far more quickly than searchers using more conventional means. Thus we use the term moving objects to refer to all, vehicles, pedestrians, and robots.

We would like to comment at this moment that in our model a peer does not have to be a moving object. 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 approach. In such an environment the mobile peer-to-peer network serves as a supplement/extension to the fixed-site based solution.

In this paper we propose to examine an *opportunistic* approach to dissemination of reports regarding availability of resources (parking slot, taxi-cab customer, dermatologist, etc.). In this approach, an object propagates the reports it carries to encountered objects, and obtains new reports in exchange. For example, an object finds out about available parking spaces from other objects. These spaces may either have been vacated by these encountered objects or these objects have obtained this information from other previously encountered ones. Thus the parking space information transitively spreads out across objects. Similarly, information about an accident or a taxi cab customer is propagated transitively. In this paper we explore this information propagation paradigm, which we call *opportunistic peer-to-peer* (or OP2P).

In this paper, we will study how a resource consumer should use the received reports to discover a resource type. 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 with regard to the driver's current location. Then the question is which parking slot the driver should go to (namely, pursue). In this paper we propose an information usage strategy called *Information Guided Searching* (or IGS), that takes both the distance and the age of the resource report into considerations. With IGS, a consumer chooses resource-reports based on a spatial-temporal relevance function that represents the likelihood that the resource is available when the consumer reaches it. The consumer always pursues the resource whose relevance is the highest and is above a certain threshold. We experimentally compare IGS with the naive strategy where the information is not used, in terms of resource discovery time, namely the length of the period of time starting when the consumer starts to look for a resource type until the consumer captures a resource. The experiments show that IGS always results in reduced discovery time compared to blind search (i.e. the information is not used). In some cases IGS cuts discovery time by more than 75%.

In summary, this paper makes the following contributions. First, we introduce a data model for spatio-temporal resources in applications related to transportation, disaster recovery, and mobile electronic commerce, and we develop an architecture for opportunistic dissemination of information about these resources. Second, we propose an information usage strategy and compare it with the naive strategy for resource discovery. We show that the information usage strategy is consistently better than the naive strategy and may cut discovery time by more than 75%.

Let us emphasize that although the experiments conducted in this paper use the peer-to-peer model, the developed resource data model and the information usage strategy hold in a cellular wireless communication environment as well. These approaches are applicable to applications that need match making based on attributes, regardless whether the resource information is received from a peer or from a server through the cellular infrastructure.

The rest of the paper is organized as follows. Section 2 develops the architecture. Section 3 discusses the information usage strategy and evaluates the benefit of resource information. Section 4 discusses relevant work. Section 5 concludes the paper and discusses future work.

2 System Architecture

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. The valid duration of an accident starts when it occurs, and lasts until it is cleaned up. The valid duration of a victim-assistance-request starts when the person is injured, and lasts until a rescuer reaches him/her. We say that 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. [13]). 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 o (either hotspot or moving object) 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 o that produces validity reports, called the *report producer* for R . o is referred to as the *producer* of each report it produces. A peer may be the report producer for multiple resources. Report $a(R)$ contains the home of R and a *timestamp*. The timestamp is the time at which the report is transmitted to a peer by its producer. Report $a(R)$ is *valid* as long as R is valid. For each resource type T , a peer o has a *validity reports database*, or *reports database*.

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. When the parking slot is occupied, the report is removed from the reports database. In the cab request example, the customer who needs a cab may click an application icon on her PDA, or press a button in a station at the closest intersection (similar to road-crossing buttons in the USA). The PDA or the station produces a report that indicates the location of the customer. The report is removed from the reports database once the customer takes a cab (this can be sensed, for example, by the cab sending a notification to the station). In the car accident example, the report is produced by the sensor that deploys the air-bag.

At any point in time, the reports database for a resource type in a peer o stores two categories of validity reports. The first category are the reports produced by o for resources that are currently valid. These are referred to as the *native reports* of o . The second group are the reports received by o from other peers (when and how the reports are received is discussed in 2.3). These are referred to as the *foreign reports* of o . A native report $a(R)$ is inserted into the reports database once it is produced, and it is deleted from the reports database when R becomes invalid. For example, if the meter for a parking slot is also a hotspot (i.e. a peer), then its reports about the monitored parking slot are native, and other reports are foreign. If a hotspot monitors all the parking slots in a lot, then all the reports about these slots are native.

2.3 Relevance Model

With OP2P, a peer constantly receives validity reports from the peers it encounters. If not controlled, the number of reports saved and communicated by a peer may continuously increase. In order to limit the data exchange volume, we employ a relevance function that prioritizes the availability reports. The relevance of a report $a(R)$ to a peer o , the *consumer*, is determined by a spatio-temporal function, which decreases as the distance of the reported resource from o and the time elapsed from report-generation increase. In this paper we use the following relevance function:

$$\text{Rel}(a(R)) = e^{-\alpha \cdot t - \beta \cdot d} \quad (\alpha, \beta \geq 0) \quad (1)$$

where t is the number of time units since $a(R)$ is transmitted by its producer, and d is the travel distance from the home-location of R to the consumer. α and β are non-negative constants that represent the decay factors of time and distance respectively. α and β may vary per resource type and per each individual peer. Observe that this function is always positive, indicating that each report always has some relevance, and it decreases as t and d increase.

Let us consider *competitive* resources, i.e. resources that require a consumer to physically reach them ahead of other consumers in order to occupy or possess them (e.g. parking slots, cab requests, or highway assistance requests). Accident reports may inform many consumers of the accident, thus they are not competitive. We now show that for a competitive resource R , under some conditions the relevance of a report $a(R)$ equals to the probability that R is valid when the consumer reaches it.

Theorem: Assume that the length of the valid duration (see subsection 2.1) of R is a random variable with an exponential distribution having mean u . Assume further that the speed of the consumer is v . If $\alpha = 1/u$ and $\beta = 1/(u \cdot v)$, then the relevance of a report $a(R)$ is the probability (at report acquisition time) that the resource R is valid when the consumer reaches R .

Proof idea: Consider a consumer that receives $a(R)$ t time units after the report has been transmitted by the producer of $a(R)$, and at that time the consumer is at distance d from R . If $\alpha = 1/u$ and $\beta = 1/(u \cdot v)$, then the relevance of the report is:

$$\text{Rel}(R) = e^{-\frac{1}{u} \cdot (t + \frac{d}{v})} \quad (2)$$

For an exponential distribution, the probability that R is still valid b time units after a validity report is transmitted is $e^{-b/u}$. Observe that the consumer will reach the resource $t+d/v$ time units after the validity report has been transmitted, thus the theorem follows.

The theorem motivates our definition of the relevance function (at least for resources with exponentially distributed valid-duration).

Observe that we implicitly made two assumptions in order to determine d and t . The first assumption is that each consumer knows its location when receiving $a(R)$, so d can be computed. The second assumption is that the clock between the report producer and the consumer is synchronized, so t can be accurately computed. Both assumptions can be satisfied if each peer is equipped with a GPS that reports both location and time.

The relevance function we use in this paper is one example in which the relevance decays exponentially per time and distance. But there are other possible types of relevance functions in which other behaviors may be exhibited. Furthermore, other factors such as the travel direction with respect to the home of a resource may be considered in the relevance function. However, in this paper we confine ourselves to time and distance alone.

2.4 Peer-to-Peer Report Exchange

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. For example, in 802.11 a node detects appearance and leaving of neighboring nodes via periodical beacon messages [8].

The user of a peer specifies to the communication module what types of validity reports she is interested in exchanging. And for each such type, the user further specifies the maximum number of reports she wishes to receive during an exchange. This number is referred to as the *interest threshold* and is denoted by M . When two peers A and B encounter each other, if both A and B have their communication module open, then A and B start a session to exchange validity reports. During each encounter, for each resource type T that B is interested in, A computes the relevance of each report $a(R)$ in its reports database. If $a(R)$ is a native report, then A updates the timestamp of $a(R)$ to be the current time before computing the relevance of $a(R)$. Finally, A chooses the top M relevant reports and transmits them to B . Upon receiving each report $a(R)$, B checks whether there is a report $a(R)'$ in its reports database that reports R . If not, B saves $a(R)$. Otherwise, B saves $a(R)$ if the timestamp of $a(R)$ is greater than that of $a(R)'$, and discards $a(R)$ otherwise.

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 trading 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.

3 Benefit of Report Dissemination

In this section we evaluate how much a consumer gains when searching for a resource using the validity reports it receives, compared to not using the resource information. Specifically, we evaluate how much time is saved when a consumer uses validity reports to capture a resource. First we describe two strategies for a consumer to discover resources, one using validity reports and the other which does not do so. Then we compare these two strategies.

3.1 Resource Discovery Strategies

The first resource discovery strategy is a naive one, called *blind search*, or BS. With this strategy, a consumer moves around the area where a resource of interest could possibly be located, and it takes possession of the first resource that is valid at the time when the consumer reaches it. For example, a driver who is looking for a parking slot simply drives around all the streets that are within walking distance from the place to visit and parks at the first parking slot that is seen available. The area within which the consumer looks for a resource is referred to as the *search space*.

The second strategy is *information guided search*, or IGS. With this strategy, a consumer starts with a blind search, and a relevance threshold H_0 . The search continues until either a resource is valid at the time the consumer reaches it, or some validity report $a(R)$ is received. In the latter case, the consumer evaluates the relevance of $a(R)$ and goes to R (i.e. attempts to capture R) if the relevance of $a(R)$ (i.e. $\text{Rel}(a(R))$) is higher than the predefined threshold H_0 . If R is invalid when the consumer reaches it, then the consumer discards $a(R)$, returns to the closest point in the search space, and continues the blind search. Clearly, if a valid resource is passed

by on the way to R , then the consumer captures it and the search ends. If another report $a(R')$ is received during the trip to R , and the relevance of $a(R')$ (i.e. $\text{Rel}(a(R'))$) is higher than $\text{Rel}(a(R))$, then the consumer goes to R' and sets the relevance threshold to $\text{Rel}(a(R'))$. Thus the relevance threshold keeps increasing while the consumer receives more relevant reports.

Our concept of the relevance threshold is used to alleviate the phenomenon of herding, which occurs when the consumers that hear about the same resource all head to that resource, leading to high contention. The relevance threshold prevents all, but the most likely to capture it, from pursuing the resource. In subsection 3.2 we will provide an approach to determining the threshold H_0 .

3.2 Comparison of Resource Discovery Strategies

In this subsection we compare IGS and BS in terms of how long it takes a consumer to discover a resource, with each one of them. First we define the performance measure, and then we describe the simulation method. Finally we present the simulation results.

3.2.1 Definition of Discovery Time

For a competitive resource, discovery means that the moving object captures the resource, i.e. it arrives to the resource while the resource is still valid. For example, discovering a parking slot means that the driver reaches the parking slot before it is occupied; discovering a cab customer resource means that the cab driver reaches the customer before the customer hires another cab; The *discovery time* is the length of the time period starting when the user starts to search the resource type, and ending when a resource of that type is captured.

3.2.2 Simulation Method




Fig. 1. The grid network, hotspots, search square, brokers and the consumer

We synthetically generated and moved objects within a 1mile×1mile grid network. The distance between two neighboring grid points is 0.1 mile (approximately the

length of one street block). Hotspots are placed on intersections (see Fig. 1). There is only one consumer in the system. The search space of the consumer is a square centered at the center of the grid network and with side length 0.6 mile. This square is referred to as the *search square*. All the other moving objects participate in opportunistic data dissemination but are not interested in capturing any resource. These objects are referred to as *brokers*.

For each broker i , we randomly chose two points on the grid network, and assigned them as the start point and the first stop of i respectively. The path of i is the shortest path between the start point and the first stop. i moves along its path from the start point to the first stop at a constant speed. When the first stop is reached, another random point is chosen as the second stop of i , and i moves from the first stop to the second stop at the same constant speed. And so on. The motion speed of a broker is randomly chosen from the interval $[v-5, v+5]$ where v is a parameter.

Resources are generated only at hotspots. At each hotspot, the length of the valid duration of a resource follows an exponential distribution with mean 30 seconds, and the time length of the invalid duration follows an exponential distribution with mean $30 \times k$ seconds. k is a parameter. It represents the ratio between the length of time a resource is invalid, and the length of time it is valid. We use this ratio to model the competition for the resource generated by other consumers, and refer to this ratio as the *invalidity-to-validity ratio*. In other words, instead of simulating multiple consumers competing for resources, we simulate a single consumer; but each resource is valid for only 30 seconds approximately, and $30 \times k$ seconds it is unavailable.

Each hotspot announces exactly one resource, and the home of that resource is the location of the hotspot. The transmission range is the same for all the hotspots and all the moving objects. We use Equation (2) defined in section 2.3 as the relevance function and use the *route-distance* as the distance metric. The *route-distance* between two locations on the grid network is the length in miles of the shortest path between them on the grid network.

In each exchange a peer acquires the single most relevant report.

There are six parameters for each simulation run (see Table 1). Among these parameters, the transmission range r , the motion speed v , the broker density g , the hotspot density s , and the invalidity-to-validity ratio k are referred to as *environmental parameters* since they define the environment with respect to communication, traffic and resource generation. The relevance threshold is not an environmental parameter.

Table 1. Parameters and their values

Parameter	Symbol	Unit	Value
Transmission range	r	meter	50, 100, 150, 200
Motion speed	v	miles/hour	10, 20, 30, 40, 50, 60
Broker density	g	brokers/mile ²	0, 50, 100, 150, 200
Invalidity-to-validity ratio	k		10, 20, 30, 40, 50
Relevance threshold	H_0		0, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9
Hotspot density	s	hotspots/mile ²	17, 36, 100

If the hotspot density is 100, then there is a hotspot at each intersection. If the hotspot density is 36, then there is a hotspot for every two intersections, i.e. the distance between any two neighboring hotspots is 0.2 miles. If the hotspot density is 17, then there is a hotspot for every three intersections, i.e. the distance between any two neighboring hotspots is 0.3 miles.

Each simulation run is executed as follows. At the beginning of the simulation run, g brokers are generated and they start to move at the same time (time 0). Resources are generated and the initial state of each resource is either valid or invalid with the probability of being valid $1/(k+1)$. When two peers are within the transmission range, they exchange reports. Each exchange completes instantaneously, i.e. it takes time 0. At the 500-th simulated second, the consumer is introduced at a random point along the search square and its initial moving direction is either clockwise or counter-clockwise with equal probability. The consumer moves at the constant speed v . The consumer looks for resources using the search square as the search space. The simulation run terminates when the consumer successfully captures a resource. The length of the time period since the consumer is introduced up to the end of the simulation run is the discovery time of that run.

For each set of parameters, we ran 1000 simulation runs and averaged the discovery times of each simulation run. The average discovery time is the discovery time for that set of parameters.

3.2.3 Simulation Results

The results show that IGS consistently outperforms BS. In some cases the discovery time of IGS is less than one-fourth that of BS. This indicates that the validity reports help shortening the discovery time dramatically.

Relevance threshold. Fig. 2 shows how the discovery time of IGS changes per the relevance threshold H_0 . From the figure it can be seen that IGS has the best performance (i.e. the minimum discovery time) when $H_0=0.1$. Above 0.1, the discovery time of IGS increases as H_0 increases. This is because as the relevance threshold increases, the consumer uses fewer reports. This means that even reports with a low relevance are better than no reports at all. IGS is equivalent to BS when $H_0=1.0$, because in this case IGS degenerates to BS (both strategies capture the resource only when the consumer is at the resource). Let us point out that the optimal relevance threshold varies depending on the environmental parameters. However, our simulations provide an approach to determining the optimal H_0 given the environmental parameters.

We identified the optimal H_0 for each set of environmental parameters. We found that in all the cases the optimal H_0 varies within a small range from 0.0 to 0.2. In many cases the optimal H_0 is 0.0. Even when the optimal H_0 is not 0.0, the discovery time obtained with the optimal H_0 is very close to that obtained with $H_0=0.0$. This is because, during the consumer's trip to the reported resource, it has the same chance of capturing a valid resource as if it uses the BS strategy. Thus going to a reported resource can only decrease but not increase the discovery time, regardless of the report relevance. Note that H_0 is the threshold for the first candidate resource that the consumer attempts to capture. However, it does not necessarily indicate the relevance (of the report) of the resource that is eventually captured.




Fig. 2. Discovery time versus initial relevance threshold H_0

From now on throughout the rest of this section, when we present the performance of IGS, we use the discovery time that is obtained with the optimal H_0 .

Impact of broker density and transmission range. The discovery time of IGS decreases as the transmission range increases, and as the broker density increases (Fig. 3 and Fig. 4). Intuitively, as the values of these two parameters increase, the interactions among peers become more frequent, and thus the newly generated reports get propagated more quickly and reach the consumer sooner. These reports have higher relevance and therefore give the consumer a higher probability of capturing a resource. The fact that the discovery time decreases as the broker density increases indicates the effect of peer-to-peer interactions on speeding up information propagation. Note that the discovery time of BS is not affected by the transmission range and the broker density.




Fig. 3. Discovery time versus broker density




Fig. 4. Discovery time versus transmission range

Impact of invalidity-to-validity ratio. Fig. 5 shows the discovery time of IGS and that of BS as functions of the invalidity-to-validity ratio. It can be seen that the discovery time of IGS and that of BS both increase as the invalidity-to-validity ratio increases, which is natural. However, the discovery time of IGS increases with a lower rate than that of BS. This suggests that IGS is particularly useful in an environment where the competition for the resources is high.

Hotspots density. Fig. 6 shows the discovery time of IGS and that of BS as functions of the hotspot density. It can be seen that the discovery time of IGS and that of BS both decrease as the hotspot density increases. This is because in our simulation resources are generated only at hotspots. So with a lower hotspot density, fewer valid resources exist in the system at a time. Further observe that the difference between IGS and BS does not change per hotspot density. This suggests that the hotspot density has little impact on the advantage of IGS over BS.

Motion speed. Fig. 7 shows the discovery time of IGS and that of BS as functions of the motion speed of the consumer. The discovery time of IGS and that of BS both decrease as the motion speed increases. This is clearly due to the fact that a higher motion speed gives the consumer a better chance to reach a valid resource before the resource is re-captured. Further observe that the ratio between the discovery time of IGS and that of BS decreases as the motion speed increases. This suggests that IGS is particularly suitable for a higher speed environment.




Fig. 5. Discovery time versus invalidity-to-validity ratio




Fig. 6. Discovery time versus hotspot density




Fig. 7. Discovery time versus motion speed

4 Relevant Work

Different resource discovery architectures (SLP [4], Jini [12], Salutation [7], and UpnP [21]) have been developed for ubiquitous computing environments over the last few years. Typically these architectures consist of a dedicated directory agent that stores information about different services or data, a set of protocols that allow resource providers to find a directory agent and to register with it, and a naming convention for resources. In highly mobile environments, due to high variability of the network topology we cannot rely on any one component being always available. Therefore, it is important to develop methods that use opportunistic exchanges rather than a dedicated directory.

Peer-to-peer networks [10, 15] and architectures [16] have all been studied in previous works. There are two major differences between our model and traditional peer-to-peer approach. First, in our environment the participating parties are physically mobile, and sometimes can be highly mobile (consider vehicles that move in opposite directions at 120 miles/hour relative speed). The object density can vary in a big range from rush hours to midnight. The underlying communication network is thus subject to topology changes and disconnections. In such an environment a moving object does not necessarily always have neighboring objects to communicate with, and even if it does, the set of the neighbors is not fixed. Furthermore, there does not always exist a communication path between a pair of peers. These characteristics defeat the applicability of typical peer-to-peer frameworks that rely on peers to forward queries (e.g. Gnutella [3], DHTs like [20]). Second, in our environment information sharing is opportunistic. Pre-defined data access structures such as search routing tables used in Gridella [1] are impractical in mobile ad hoc networks. They are replaced by opportunistic peer-to-peer interactions.

A lot of work has been done on data dissemination in mobile environments (e.g. [17, 19, 11, 5]). However, all this work considers regular data items but not spatial-temporal ones. So the benefit of data dissemination is measured differently than here. In the existing work, the benefit of data dissemination is usually measured by the level of the consistency between the disseminated copies and the master copy. It does not consider how the data is used. In our work we use the spatio-temporal relevance function to express the utility of data (loosely speaking, to measure the consistency), and we analyze how the information is used, and what benefit in terms of time-saving it provides.

Goel et al [27] study the dissemination of traffic-speed information in both peer-to-peer and infrastructure environments, and evaluates the benefit of traffic-speed information in terms of travel time reduced compared to if information is not used. However, the approach considers only non-competitive resources (traffic-speeds) whereas we consider both non-competitive resources and competitive resources in our model, and we evaluate the benefit of competitive resources. Moreover, they do not discuss information prioritization of reports in terms of their relevance.

Finally, this paper differs from our prior work [24, 25, 26] in multiple aspects. The model is more general in the sense that it allows resources to be advertised by the stationary hotspots in addition to moving objects (in [24] resources are advertised by moving objects only). The hotspots not only announce reports but also exchange reports with moving objects. In other words, hotspots announce resources other than

their own (in [25] hotspots only announce resources they produce). The information usage is new, and the evaluation of the benefit of information is also new ([25] also evaluates benefit but using a much simpler simulation model and it does not discuss information usage).

5 Conclusion

In this paper we devised an architecture and a data model for dissemination of spatial and temporal resource-information in a mobile peer-to-peer communication environment, in which the database is distributed among the hotspots and moving objects. We analyzed IGS, the information guided resource discovery strategy, and determined that IGS reduces search time compared to blind search, and by how much. We also determined how the time savings varies as a function of the environmental parameters, such as the number of brokers, the competition for resources reflected in the valid/invalid time ratio, the speed of the moving objects, the number of hotspots producing resources, and the transmission range.

In general, we feel that the P2P paradigm is a tidal wave that has tremendous potential, as Napster and Gnutella have already demonstrated for entertainment resources. Mobile P2P is the next step, and it will revolutionize dissemination of spatial and temporal resources. For example, location based services have been considered a hot topic for quite some time, and it has been assumed that they have to be provided by a separate commercial entity such as the cellular service providers. The approach outlined in this paper can provide an alternative that bypasses the commercial entity. Although in this paper we used traffic examples, as mentioned in the introduction, the applications range from social networks to disaster recovery to mobile electronic commerce.

In terms of future work, much remains to be done. For example, incentive mechanisms need to be developed to stimulate peers to participate as suppliers of resource information and/or intermediaries for information propagation. Strategies that best utilize the received information need to be investigated in a realistic traffic and resource operation environment. Other forms of interactions, e.g. peer-to-peer broadcast rather than pair-wise exchange, are also worth studying. In this paper, we assumed a memory-less probabilistic distribution for the valid-time of a resource. We will examine with other distributions such as a normal distribution.

Acknowledgements: We thank Aris Ouksel for helpful discussions.

References

1. Karl Aberer, Manfred Hauswirth, Magdalena Puceva, Roman Schmidt. Improving Data Access in P2P Systems, *IEEE Internet Computing*, 6(1), January/February 2002.
2. CarTalk. <http://www.cartalk2000.net/>
3. Gnutella website. <http://gnutella.wego.com>

4. E. Guttman, C. Perkins, J. Veizades, M. Day, Service Location Protocol, Version 2. *RFC2608*, June 1999. <http://www.ietf.org/rfc/rfc2608.txt>
5. H. Hayashi, T. Hara, and S. Nishio. *Cache Invalidation for Updated Data in Ad Hoc Networks*. Proc. Int'l Conf. on Cooperative Information Systems (CoopIS'03), 2003.
6. http://www.etsi.org/etsi_radar/cooking/rub11/transport_a.htm
7. <http://www.salutation.org/>
8. IEEE Computer Society. *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*. 1997.
9. Intelligent Safety Efforts in America. <http://www.its.dot.gov/speeches/madridvii2003.ppt>
10. V. Kalogeraki, A. Delis, D. Gunopulos: Peer-to-Peer Architectures for Scalable, Efficient and Reliable Media Services. IPDPS 2003.
11. G. Karumanchi, S. Muralidharan, and R. Prakash. Information Dissemination in Partitionable Mobile Ad Hoc Networks. *Proc. Symposium on Reliable Distributed Systems (SRDS'99)*, pp.4-13, 1999.
12. W. Keith Edwards, *Core JINI*, Prentice Hall, 1999.
13. A. Markowetz, et al. Exploiting the Internet As a Geospatial Database, *International Workshop on Next Generation Geospatial Information*, 2003.
14. New Microsoft Goal: A computer in every car. http://www.usatoday.com/tech/news/2003-12-01-ms-cars_x.htm
15. W.S. Ng, B. C. Ooi, K.L. Tan, A. Zhou. PeerDB: A P2P-based System for Distributed Data Sharing. International Conference on Data Engineering (ICDE'2003), Bangalore, 2003.
16. M. Papazoglou, B. Krüner, J. Yang. Leveraging Web-Services and Peer-to-Peer Networks. In *Proceedings of CAiSE 2003*, June, 2003.
17. M. Papadopouli and H. Schulzrinne. Effects of Power Conservation, Wireless Coverage and Cooperation on Data Dissemination Among Mobile Devices. *MobiHoc 2001*, October 4-5, 2001, Long Beach, California.
18. Robo-rescuers increase disaster victims' chances. http://firechief.com/ar/firefighting_roborecuers_increase_disaster/
19. K. Rothermel, C. Becker, and J. Hahner. Consistent Update Diffusion in Mobile Ad Hoc Networks. *Technical Report 2002/04*, Computer Science Department, University of Stuttgart, 2002.
20. Stoica, R. Morris, D. Karger, M. Kaashoek, and H. Balakrishnan. Chord: A Scalable Peer-to-Peer Lookup Service for Internet Applications. In *Procs. ACM SIGCOMM*, 2001.
21. Universal Plug-and-Play (UPnP) Forum, *Microsoft Corporation*. <http://www.upnp.org>.
22. Ultra-wideband (UWB). <http://www.ubisense.net/technology/uwb.html>
23. E. Welsh, P. Murphy, P. Frantz. A Mobile Testbed for GPS-Based ITS/IVC and Ad Hoc Routing Experimentation. *International Symposium on Wireless Personal Multimedia Communications (WPMC)*, Honolulu, HI, Oct. 2002.
24. B. Xu, A. Ouksel, O. Wolfson, Opportunistic Resource Exchange in Inter-vehicle Ad Hoc Networks, *Proc. of 2004 IEEE International Conference on Mobile Data Management*, Berkeley, California, Jan. 2004.
25. O. Wolfson, B. Xu, P. Sistla. An Economic Model for Resource Exchange in Mobile Peer-to-Peer Networks. *Proceedings of the 16th International Conference on Scientific and Statistical Database Management (SSDBM'04)*, Santorini Island, Greece, June 2004.
26. O. Wolfson, B. Xu. Opportunistic Dissemination of Spatio-Temporal Resource Information in Mobile Peer-to-Peer Networks. Accepted, to appear in *1st International Workshop on P2P Data Management, Security and Trust (PDMST'04)*, Zaragoza, Spain, Sept. 2004.
27. S. Goel, T. Imielinski, K. Ozbay, and B. Nath, Grassroots: A Scalable and Robust Information Architecture. *Technical Report DCS-TR-523*, Department of Computer Science, Rutgers University, June 2003. <http://paul.rutgers.edu/~gsamir>