

Opportunistic Dissemination of Spatio-temporal Resource Information in Mobile Peer to Peer Networks

Ouri Wolfson and Bo Xu

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

Abstract

In this paper we examine the dissemination of availability reports about resources in mobile peer-to-peer networks, where vehicles, pedestrian or 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 objects and obtains new reports in exchange. Least relevant reports are discarded after each exchange such as to limit the communication data volume of future exchanges. We analyze by simulations the information accuracy provided by this approach. We also develop and analyze a method to determining the optimal size of the local reports database at a moving object such that a certain accuracy level is reached. 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 Potential drawbacks of this approach are (i) the responses to queries may be outdated, (ii) the response time may not meet the real-time requirements, and (iii) access to infrastructure communication service is costly, (iv) currently there is no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.

We assume that each moving object (e.g. vehicle) has processing power (see [1]), and the capability of communicating with its neighbors. This communication can be enabled by a local area wireless protocol such as IEEE 802.11 [2], Ultra Wide Band (UWB) [3], or CALM [5]. 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 [8] and are being planned for deployment on a large-scale environment [4]. Similar communication capabilities are being planned between sensors in the infrastructure and moving vehicles [4]. 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, information dissemination is performed in a *mobile peer-to-peer network*.

Compared to the traditional fixed-site based information query, this paradigm (“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 lot 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 injured in exploded or earthquake-damaged buildings [9]. These robots are equipped with radio transmitters. When a robot discovers an injured, 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.

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, a vehicle finds out about available parking spaces from other vehicles. These spaces may either have been vacated by these encountered vehicles or these vehicles have obtained this information from other previously encountered ones. 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).

With OP2P, a moving object constantly receives reports from the objects it encounter. In order to limit the data exchange volume, least relevant reports are

discarded after each exchange so that only the top M relevant availability reports are kept in the moving object. We call this method *Opportunistic Data Dissemination* (ORD).

Now consider the accuracy or reliability of the information recorded at o . At any time instance an availability report Q in o 's reports database may be obsolete, in the sense that the reported resource that shows as available, has actually become unavailable, but o does not know about it. Therefore an important question arising with ORD is what level of accuracy ORD provides, and how the accuracy is affected by various system environmental parameters such as the density of moving objects, their motion speeds, the wireless transmission range, and so on. We will study these issues in this paper. We will also propose and analyze an approach to determining the optimal size of the reports database that satisfies a certain accuracy constraint.

The rest of the paper is organized as follows. Section 2 introduces the model and the ORD algorithm. Section 3 evaluates the information accuracy. Section 4 discusses the determination of the optimal database size. Section 5 concludes the paper.

2. The Model

2.1. Resource Model

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. A *spatio-temporal resource*, or a *resource* for short, is a piece of information about an event. Spatio-temporal resources include, for example, the availability of a parking slot, a car accident, the speed of a vehicle at a particular time-point, a taxi-cab request at a street intersection, an 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. The event is specific to a certain location that 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. Each resource has a time duration for which it is valid. This duration is referred to as the *valid duration*. For example, the valid duration of the cab request resource is the time period since the request is issued, until a cab is hired or the request is 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 one or more other passengers. The valid duration of an accident starts when it occurs, and lasts until it is cleaned up. The valid duration of an injured person starts when the person is detected by a sensor robot, and lasts until a rescuer reaches him.

2.2. Availability Reports and Their Generation

Each resource R , once available, is constantly announced by the producer to all the moving objects within the transmission range of the producer. Each announcement is referred to as an *availability report*. The availability report contains the home of the resource and the time when the report is transmitted. The latter is referred to as the *transmission-time* of the report. With such information each report transmission confirms the availability of the resource at the transmission time. In the parking slots example, a sensor in the parking slot (or the meter for the slot) monitors the slot, and, while unoccupied, announces to the neighboring vehicles the location of the slot. In the cab request example, the customer who needs a cab may press a button in a station at the closest intersection (similar to road-crossing buttons in the USA). The station announces the location of the intersection.

The producer reports the resource only within the valid duration of the resource. So for example, the parking slot sensor reports the slot only when it is available; occupation of the slot is sensed and the report transmission ceases. The cab-call station stops reporting the request once the customer takes a cab.

2.3. Relevance Model

For each application, a moving object maintains a database that records at any point in time the top M relevant availability reports the object knows at that time. This database is referred to as the *availability reports database* or *reports database* for short. M is referred to as the *interest threshold*. For example, a user who is looking for a parking slot has the reports database that keeps top 10 relevant parking slot availability reports. In other words, the user wants only 10 most relevant parking slot reports to be saved and displayed¹. In this paper we use the relevance function introduced in [7]:

$$F = -\alpha \cdot t - \beta \cdot d \quad (\alpha, \beta \geq 0) \quad (1)$$

¹ This parking slot information is only relevant to the driver when he is close to his destination. However, the driver should not be bothered indicating this and the default is always showing all the slots saved in the memory.

t is the time length since the transmission-time (i.e. the time when the report is transmitted by the producer), and d is the distance from the home location of the reported resource. α and β are constants that represent the decay factors of time and distance respectively. α and β may vary per resource type and per moving object.

In this paper we set $\alpha=1$ and $\beta=1/v$, where v is the motion speed of the consumer. Thus the relevance of the report is: $F = -t - d/v$. The absolute value of the relevance represents the length of the period of time starting when the report is transmitted by the producer, and ending when the consumer can reach the reported resource. The larger this length grows, the less relevant the report is.

2.4. Opportunistic Report Dissemination (ORD)

When a moving object o encounters another object o' , namely when the distance between o and o' is smaller than the transmission range, o and o' exchange their availability reports. Upon receiving new reports, o computes the relevance of the received reports and re-evaluates the relevance of its own reports. If all the reports do not fit in the reports database of o , the least relevant ones are purged such that only the top M reports are kept in the database.

During report exchange, if o and o' move within the transmission range for a period of time, after the initial exchange only newly arrived reports are exchanged. In this way, a new report is flooded to all the moving objects that do not have the report but are connected to o or o' by single or multi-hop links. If o encounters two or more objects simultaneously, the exchanges occur sequentially. In other words, we assume that the underlying communication module provides a mechanism to resolve interference and conflicts.

3. Data Accuracy of ORD

Section 3.1 defines several notions related to accuracy. Section 3.2 describes the simulation method, and section 3.3 presents the simulation results.

3.1. Definitions

An availability report Q is *valid* at time t if the resource reported by Q is available at t . Otherwise Q is *invalid* at t . We say that the *relevance position* of Q at object o at t is i if Q is the i -th relevant report at o at t . The *valid percentage of position* i , denoted by CP_i , is the percentage of time when the i -th relevant report in o 's database is valid. The report at a certain relevance position i may change over time. However, the valid

percentage of position i represents the likelihood that at an arbitrary time the report at i is correct.

3.2. Simulation Method

We synthetically generated and moved objects within a 5mile×5mile grid network. The distance between two neighboring grid points is 0.1 mile (approximately the length of 1 street block). The objects move on the network in a random way-point model [6]. The motion speed of an object is randomly chosen from the interval $[v, v+10]$ where v is a parameter. We use the *route-distance* as the distance metric. The route-distance between two points on the grid network is the length in miles of the shortest path between them.

There are six parameters for each simulation run, namely the interest threshold M , the transmission range r , the lower bound speed v , the object density g , the resource density s , and the mean of the length of resource valid duration e . Each simulation run is executed as follows. At the beginning of the simulation run, $5 \times 5 \times g$ objects are generated and they start to move at the same time (time 0). $5 \times 5 \times s$ resources are generated such that the location of each resource is at a random position on a random link of the grid network. The length of the valid duration of each resource follows an exponential distribution with mean e . Whenever a resource becomes unavailable, it is removed from the system, and a new resource is generated in the same random fashion. The reports of the available resources are acquired by the objects that pass by. When the distance between two objects is smaller than r , they exchange their reports, re-evaluate the relevance, and keep only the M most relevant reports. Both the acquisition and the exchange are finished instantaneously. The length of each run is 10 simulated hours. The system parameters and their values are listed in Table 1 below.

Table 1: Simulation parameters and their values

Parameter	Symbol	Unit	Value
Interest threshold	M		1-30
Mean of valid duration length	e	second	30-120 with increment 5
Resource density	s	resources/mile ²	50-400 with increment 10
Transmission range	r	meter	50-200 with increment 5
Motion speed	v	miles/hour	10-60 with increment 5
Object density	g	objects/mile ²	100-1000 with increment 50

3.3. Simulation Results

Table 2 shows the valid percentages of each relevance position i for different values of the interest threshold M , under the same combination of other parameters. First let us examine the valid percentages for a given M value. For a given M value (i.e. at a single row in Table 2), the valid percentage decreases as the position i increases. In other words, the valid percentage decreases as the report relevance decreases. For example, when $M=10$, CP_1 decreases from 93% to 80% when i increases from 1 to 10. This indicates a direct correlation between the relevance of a report and its likelihood of being valid. Intuitively, when the relevance of a report Q is high, it means that the age of Q is small (i.e. Q was announced very recently) and the distance to the home of the reported resource is small. Since the age is small, it is likely that Q is still valid.

Now let us consider the valid percentages of a given position i . For a given position i (i.e. at a single column in Table 2), the valid percentage is approximately the same for different M values. For example, consider the valid percentage of position 1. As M changes from 1 to 30, CP_1 varies within a pretty small range (between 91% and 94%), and there is not a trend, either increasing or decreasing. This is a very interesting property. It suggests that when a small interest threshold (M) is used, the valid percentages of the top positions are the same as if a bigger interest threshold is used. In other words, saving more reports at each exchange does not help any on lifting the information accuracy.

Table 2: Valid percentages for different M values ($e = 60, s = 100, r = 100, v = 40, g = 500$)

	i=1	i=2	i=3	i=4	i=5	i=6	i=7	i=8	i=9	i=10
M=1	94									
M=2	94	91								
M=3	93	91	89							
M=4	93	91	89	87						
M=5	92	90	89	87	86					
M=6	93	90	89	88	87	85				
M=7	91	89	87	86	85	84	82			
M=8	92	89	89	87	85	84	83	82		
M=9	91	88	87	85	84	83	83	82	81	
M=10	93	90	88	87	86	84	84	83	82	80
M=20	92	90	88	86	84	83	82	81	81	80
M=30	92	88	86	85	84	83	82	81	81	81

Simulations with other parameter combinations demonstrate a similar pattern as shown in Table 2.

4. Determination of the Optimal Interest Threshold

For the purpose of decision-making, a moving object o needs to have a certain confidence level on the accuracy of the reports it uses. For example, o may want that the valid percentages of all the reports used for decision-making should be higher than 90%. The analysis on data accuracy in section 3.3 revealed that the valid percentage consistently decreases as the relevance position number increases. This suggests that o should use only the top- k reports where k is the maximum relevance position number for which the valid percentage is higher than or equal to 90%. In the mean time, section 3.3 tells us that keeping more reports does not increase the data accuracy. This suggests that o does not need to keep the reports that are below position k . Therefore the optimal interest threshold for the 90% confidence level in this case is k . k is optimal in the sense that (i) a higher interest threshold is not necessary because it does not provide higher data accuracy; (ii) a lower interest threshold is not enough because it does not give all the reports that reach the desired the confidence level.

In this section, we identify the optimal interest thresholds for different system environments. We derive the optimal interest thresholds from the simulation results generated in section 3. Specifically, for each combination of the object density g , the transmission range r , the motion speed v , the resource density s , and the mean of valid duration length e , we look at the valid percentages for $M=30$. We find the maximum relevance position i such that the valid percentage of i is higher than or equal to a given confidence level. This i is the optimal interest threshold for the given confidence level and for the given system environment. Now we present the results. Due to space limitations, we only show the optimal interest threshold as a function of the object density (see Figure 1). The figure gives two curves, for confidence levels 85% and 90% respectively.

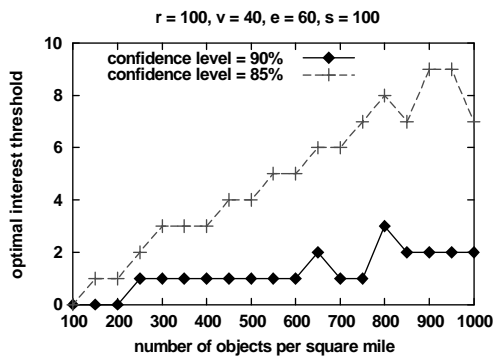


Figure 1: Optimal interest threshold versus object density

Figure 1 shows that the optimal interest threshold increases as the object density increases. Intuitively, as the object density increases, the interactions between moving objects and producers, and the interactions among moving objects become more intensive. Thus the reports get propagated more quickly and are able to reach more objects before they become “old” or invalid. These reports purge out older reports, which are more likely to be invalid, and therefore the valid percentage is increased.

5. Conclusion

In this paper we devised a 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 moving objects. We introduced a resource dissemination algorithm with which a moving object builds a local database of availability reports. We evaluated by simulations the accuracy of the information recorded in the reports database. The results show that in most situations, more than 85% of time the most relevant report at a moving object correctly reflects the availability of the reported resource. We also developed an approach to determining the optimal size of the reports database.

References

- [1] http://www.usatoday.com/tech/news/2003-12-01-ms-cars_x.htm
- [2] IEEE Computer Society. *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*. 1997.
- [3] <http://www.ubisense.net/technology/uwb.html>
- [4] <http://www.its.dot.gov/speeches/madridvii2003.ppt>
- [5] <http://www.etsi.org/frameset/home.htm?/technic/activ/ITS/ITS.htm>
- [6] G. Resta, P. Santi. An Analysis of the Node Spatial Distribution of the Random Waypoint Mobility Model for Ad Hoc Networks. *Proc. 2th ACM Int. Workshop on Principles of Mobile Computing*, Toulouse, France, 2002, pp. 44-50.
- [7] B. Xu, A. Ouksel, O. Wolfson. Opportunistic Resource Exchange in Inter-vehicle Ad Hoc Networks. *Proc. MDM 2004*. Jan. 2004.
- [8] <http://www.cartalk2000.net/>
- [9] http://firechief.com/ar/firefighting_roborescuers_increase_disaster/