

An Economic Model for Resource Exchange in Mobile Peer to Peer Networks¹

Ouri Wolfson Bo Xu A. Prasad Sistla

Department of Computer Science, University of Illinois at Chicago

Abstract

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 the available parking spaces around the current location. Or, the driver may be interested in the traffic conditions (e.g. average speed) one mile ahead.

In this paper we examine the dissemination of information about resources in mobile peer-to-peer networks, where vehicles and sensors communicate with each other via short-range wireless transmission. Each disseminated resource represents information about a spatial-temporal event, such as the availability of a parking slot at a particular time. We propose an opportunistic dissemination paradigm, in which a moving object transmits the resources it carries to encountered vehicles and obtains new resources in exchange. We develop and analyze a family of economic models that provide incentive for collaboration in data dissemination. 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 Potential drawbacks of this approach are (i) the responses to queries may be outdated, (ii) the response time may not meet the real-time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], Bluetooth [4], 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 [6, 32] and are being planned for deployment on a large-scale environment [31]. Similar communication capabilities are being planned between sensors in the infrastructure and moving vehicles [31]. 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 (“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

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

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, in order to match specific needs with expertise (e.g. burn victim and dermatologist). Thus, we use the term moving objects to refer to both, vehicles and pedestrians.

In this paper we propose to examine an *opportunistic* approach to resource (parking slot, taxi-cab customer, dermatologist, etc.) dissemination, in which an object propagates the resources it carries to encountered objects, and obtains new resources 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 resource propagation paradigm, which we call *opportunistic peer-to-peer*.

Like many other peer-to-peer systems, our opportunistic peer-to-peer paradigm heavily relies on cooperation among users. However, an opportunistic peer-to-peer network typically consists of moving objects owned by different authorities (different companies, different private owners) and each moving object has its own goal. In such an environment, the owner of a moving object may decide not to cooperate in resource exchange.

In this paper we develop various incentive mechanisms to stimulate cooperation in an opportunistic peer-to-peer environment. The incentive mechanisms are based upon virtual currency [7, 24]. Each mobile node carries virtual currency in the form of a counter that is protected from illegitimate manipulation by a trusted and tamper resistant hardware module [12]. The incentive mechanisms differ in who pays, who charges, and how much is paid or charged. We classify resources into two basic types depending on whether the producer or the consumer pays for the propagation of the resource. One resource type is *producer-paid resource*. The producer of a producer-paid resource pays an "advertisement" fee by attaching a certain amount of virtual currency in the announcement message. Each mobile node that transmits the resource withdraws a "commission" from the advertisement fee. A gas station that advertises to neighboring vehicles is an example of producer-paid communication. An issue naturally arising from the producer-paid model is how much advertisement fee the producer should put in order for the advertisement to reach a given percentage of moving objects within a certain size of area. Furthermore, how to achieve the best tradeoff between the advertising cost and the advertising effect, i.e.

how to use minimum advertising fee to reach maximum number of objects? We will address these issues by simulations.

Another resource type is *consumer-paid resource*, namely the consumer of a resource, rather than the producer, pays the fee. The amount of the fee depends on the relevance of the resource at the time of trading. Parking slot advertisements are examples of such resources. We will evaluate how much a consumer gains when searching for a resource using the resource information it bought, compared to not using the resource information. Specifically, we evaluate how much time is saved when a consumer uses the resource information to search a resource. This saving justifies why a consumer should and why she is willing to pay for the received resource information.

In summary, this paper makes the following contributions. First, we devise two incentive models to stimulate cooperation in a novel resource discovery environment: an opportunistic peer-to-peer environment. These two incentive models differ in who pays, who charges, and how much is paid or charged. Second, we experimentally show that paying for the propagation of resource is well justified by the benefit gained by the payer, which in turn justifies the incentive models. Finally, we quantify by simulations the best tradeoff between the payment and the benefit.

The rest of this paper is organized as follows. In section 2 we describe the system model. In section 3 we introduce two economic models, namely the producer-paid model and the consumer-paid model. In section 4 we experimentally analyze the producer-paid model and the consumer-paid model. In section 5 we discuss relevant work. In section 6 we conclude the paper and discuss future work.

2. System 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 (e.g. the availability of a parking slot, or a car accident, the speed of a vehicle at a particular time-point, a taxi-cab request at a particular location, 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 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*. The start time of the valid duration is referred to as the *create time* of the

resource. For example, the valid duration of the available-parking-slot resource is the period of time since the slot becomes available until it is occupied. 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 an accident starts when it occurs, and lasts until it is cleaned up.

Let us comment further about spatial resources, such as gas stations, ATM machines, etc. Opportunistic dissemination of such resources is an alternative paradigm to geographic web searching (see e.g. [11]). 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; but she does not need to do so in our scheme.

2.2 Station to Vehicle Transmission

The system consists of fixed stations and moving objects. Each station senses resources. The station announces the resource to the moving objects within the transmission range of the station by broadcast or unicast. Each announcement message contains the home and the create-time of the resource. 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, and when it became available. 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). The station announces the location of the intersection and the time at which the button is pressed. In the car accident example, the event may be announced by the sensor that deploys the air-bag.

A station announces a resource only within the valid duration of the resource. So for example, the parking slot sensor announces the slot only when it is available; occupation of the slot is sensed and the broadcast ceases. The cab-call station stops announcing the request once the customer takes a cab.

2.3 Virtual Currency and the Security Module

The system circulates a virtual currency called *coins*. The coins owned by each moving object is represented by a coin counter that is physically stored in that object. The coin counter is decreased when the object pays out for buying

resources and increased when the object earns in for selling. Each object has a trusted and tamper resistant hardware module called *security module* [12]. The coin counter is stored in the security module and thus is protected from illegitimate manipulation. A resource producer (i.e. a station) may also own coins. A producer is trusted and tamper resistant, and their coins are also protected from illegitimate manipulation. Each coin is bought for a certain amount of real money and it can be cashed for the same amount of real money. The resources saved by a moving object are stored in the security module.

2.4 Vehicle to Vehicle Exchange

The user specifies to the security module what types of resources she is interested in buying and selling. For example, “buy the information about parking slots that are closer than 100 meters from my current location”. Denote by r the wireless transmission range. We say that two vehicles encounter each other when their distance is smaller than r . When two vehicles A and B encounter each other, if both A and B have their security module open, then A and B start a secure session to exchange resources². During each encounter A buys resources A is interested in and B has, and sells resources A has and B is interested in. Similarly for B. If A encounters two or more vehicles simultaneously, the exchanges occur sequentially. In other words, we assume that there is a mechanism to resolve interference and conflicts.

3. Resource Types and Economic Models

Each resource belongs to a resource type, and for each resource type there is an economic model that provides the incentive for brokering and trading in resources of this type. The types differ depending on whether the producer or the consumer pays for the resource. The economic model associated with the producer-paid resource types consists of a policy obeyed by the participants in the exchanges involving a producer-paid resource. The policy is implemented in the security module. Similarly for the consumer paid resource type.

For each resource type T , the owner/user of a moving object may decide not to participate in the exchange of resources of type T . However, if it participates in the game, then it behaves according to the economic model associated with T . The two models are described in the next two subsections respectively, and in the last subsection we discuss hybrid resources.

² The secure session is established based on some public key infrastructure that is omitted in this paper due to space limitations.

3.1 Producer-Paid Resources

Resources which the owner is interested in advertising are producer-paid. A gas station that advertises its location to neighboring vehicles is an example of the producer-paid policy. A car that reports its accident or breakdown and requests assistance from other drivers (e.g. jump-start cables) is another example.

At a high level, the producer-paid model works as follows. When a resource R is announced by its producer, the producer loads with R a certain number of coins, C , called the *initial budget* of R . When R is received by a moving object, it carries a certain budget C_0 . A moving object earns a flat commission fee f each time it transmits the resource to another moving object. The remaining budget of the resource is divided between the sender and receiver (in order for both to keep propagating the resource).

Now let us describe the detailed procedure. Let A be a moving object that carries a resource R . Let C be the current budget of R . Suppose that A encounters B . If both A and B have R , then A and B will not update their coin counters. Otherwise, A increases its coin counter by f , and sets the

budget of R to be $\frac{C-f}{2}$. B sets the budget of its copy of R

to be $\frac{C-f}{2}$. If $\frac{C-f}{2} < f$, namely the remaining budget

is less than the brokerage fee, then both A and B will stop transmitting R in the future.

Observe that there is yet another way of distributing the remaining budget. Namely, instead of splitting it between the sender and the receiver, the sender always keeps all the remaining budget. However, in this way, only the moving objects that receive the resource directly from the producer have chance to replicate the resource. Thus replications occur only along the routes of these objects. Intuitively, this leads to a less uniform spreading of the resource in comparison with “budget splitting”, and therefore it generates poorer advertisement effect. This observation has been confirmed by experiments.

3.2 Consumer-paid Resources

Resources which the consumer is interested in receiving are consumer-paid. Available parking slots is an example of a resource type for which the consumer pays.

For each consumer-paid resource R , the security module of a moving object operates in one of the following modes:

1. *Consumer*. In consumer mode, the security module displays to the user the resource R . The consumer pays for the resource an amount that is dependent on the relevance of the resource; this amount is discussed in subsection 3.2.1.

2. *Broker*. In broker mode the security module simply stores R and forwards it to other moving objects, but R is not displayed to the user, thus the user in broker mode cannot directly benefit from R . A broker does not pay for R , and it is paid only in case it sells R to a consumer. Thus the broker participates in the game speculatively, and it is paid for R only if R is bought by a consumer. The price constitutes the incentive of the broker to participate in the game. A broker may sell R to multiple consumers.

At any point in time a moving object may operate in broker mode or in consumer mode for a resource type T . Which mode the security module operates in is controlled by the user, per the user’s interest. For example, a traveler may operate in broker mode for parking slot resources while she is far from her destination, and then switch to the consumer mode when she prepares to park. For another example, a cab driver may operate in the consumer mode while he is looking for a customer and switch to the broker mode after the cab is hired. When switching to consumer mode, a user allocates a budget for the purchase of resources. When the user switches back to broker or drops out of the game for this resource type, any remaining funds in the budget go back to the user.

Resources acquired in consumer mode are consumer resources, and resources acquired in broker mode are broker resources. At a particular moving object a resource cannot switch from broker to consumer mode, or vice versa.

3.2.1 Price of a Resource

The consumer price of R is a function of its relevance when traded. In this paper we will use the relevance function introduced in [28]:

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

t is the time length since the creation of R and d is the distance from the home location of R to the consumer. α and β are constants that represent the decay factors of time and distance respectively.

For example, the relevance (i.e. usefulness) of a parking slot depends on both, the period of time since the slot became available, and on the distance of the consumer from the slot when receiving the resource (i.e. the slot-availability information). As the information about a parking slot grows stale, it becomes less and less relevant as the likelihood of its availability decreases. Its relevance should therefore be less than that of a more recent one. This comparison must however be tempered by the distance factor. A parking slot that is closer to a vehicle is certainly more relevant than one that is farther away. In general, the relevance of a resource decays as its age (i.e. the length of the period of time since its creation-time) increases, and the distance from its home increases.

How should α and β be determined? One possibility is to fix them such that the relevance is $-PT$ where PT is the length of the period of time starting when the slot became available, and ending when the consumer can reach the slot. The larger PT grows, the less relevant the slot is. Assuming that the speed of the consumer vehicle is v and it moves to the slot in a straight line, $PT = t + d/v$. In this case, $\alpha=1$ and $\beta=1/v$, and the relevance of a parking slot is:

$$F(R) = -t - \frac{d}{v}$$

where v is the travel speed of a consumer.

Based on the relevance function, the price is formulated as follows.

$$P(R) = \begin{cases} E - (\alpha \cdot t + \beta \cdot d) & \alpha \cdot t + \beta \cdot d < E \\ 0 & \alpha \cdot t + \beta \cdot d \geq E \end{cases}$$

where E is the value of the resource to a consumer that is at the parking slot when it becomes available. Clearly this value may differ from consumer to consumer, however, our model is fixed-price (as in a store) rather than auction based.

E , α , and β may differ for different resources of the same type. For example, these values may be different for different parking slots, and even for the same parking slot they may be different at different times of day. These values are included in the resource announcement message, and propagated by the broker.

To continue the above example, if $\alpha=1$ and $\beta=1/v$, then E can be set to be a value that is bigger than the average length of the valid duration of a parking slot. It means that a parking slot resource has a non-zero price if (statistically) the consumer can reach the slot before it is occupied by some other moving object.

3.2.2 The Consumer-paid Policy – Summary

When a moving object switches to consumer mode for a resource type, it acquires new resources of that type in consumer-mode. For each such resource, the consumer pays $P(R)$. This price goes to the broker or the resource-advertiser from which the user acquired the resource. A broker does not pay anything for the resource (but, remember, cannot view the resource). A user cannot sell or exchange any resource it holds in consumer-mode (so that it does not sell the parking-slot information after it has occupied the slot). However, a broker can sell the same resource to multiple consumers, and it obviously can transfer (for free) the same resource to multiple brokers.

Clearly, there are other possible variants of this economic model. For example, we experimented with a broker-to-broker commission fee as a percentage of the value of the resource. But a commission fee of 0 turned out

to optimize broker incentive. The consumer-paid policy and the producer-paid policy are summarized in the table below.

producer-paid	sender (A)	
receiver (B)	A increases its coin counter by f and sets the budget of R to be $(C-f)/2$. B sets the budget of R to be $(C-f)/2$.	
	consumer-paid	sender (A)
receiver (B)	consumer	B pays A $P(R)$
	broker	no payment

3.3 Producer/Consumer-Paid Resources

Some resources may be both, consumer-paid and producer-paid. For example, consider taxi-cab requests. In principle, both the cab driver and the customer are interested in the match. Thus there are two variants of the cab-request type, consumer-paid and producer-paid. A customer may advertise its cab request as a consumer-paid resource. However, a free cab that does not have coins, or is not willing to pay will not be able to view this request. Thus, for maximum exposure, the customer may advertise the request in both variants.

4. Experimental Analysis

In this section we analyze by simulations the two incentive policies described above.

4.1 Analysis of the Producer-paid Policy

4.1.1 Advertising Strategies

For a producer-paid resource, the producer wishes the resource to be advertised to all the moving objects within a certain area (called the *coverage area*), and it wants to do so at minimum cost. This is the purely spatial variant of the problem, i.e. the relevance does not decay with time. This is the case, e.g., for a gas station or a restaurant.

Complete coverage is not always achievable, for the following reason. Even when the price is disregarded, due to the opportunistic nature of the mobile peer-to-peer network, an object is not guaranteed to receive the advertised resource when it travels through the coverage area. It may enter and exit the area without meeting another moving object, and without passing by the provider. So probabilistically speaking, the coverage can never be 100 percent.

Cost minimization is also elusive. When an object O that has the resource leaves the coverage area, the remaining

budget associated with the resource is also carried out of the coverage area. Hopefully the remaining budget will eventually be used out as O encounters more objects. However, some of these later encountered objects may not travel to the coverage area, in which case some of the remaining budget is wasted. This suggests that it is essentially impossible for the producer to precisely control the advertising cost spent on each object reached.

For a producer-paid resource, there are two decisions the producer has to make. First, which objects it should announce the resource to. Should it announce to every object that passes by, or should it announce for every k objects, and for the latter choice should k be a constant or variable? Second, how much initial budget should the producer include in each announcement. In this paper we consider a simple strategy that we call the *simple advertising strategy* (SAS). In this strategy, the producer announces the resource to each object that passes by, and it includes in each announcement the same amount of initial budget.

In this subsection we will analyze SAS. We will study how much initial budget is needed in order to (probabilistically) reach a given percentage of objects in the coverage area. We will also study the tradeoff between the cost of advertisement and the effect of the advertisement (i.e. the percentage of objects reached by the advertisement).

Finally, we will compare SAS with a variant of the producer-paid policy in which the cost efficiency is optimized, namely the cost per each reached object is exactly the commission fee f . In this variant, the producer does not include any initial budget in an announcement, and an object does not withdraw any virtual currency from resource transmission. Instead, there is a centralized credit clearance system. Each producer has a credit account in the system. The producer attaches its account number to each announcement. The announced resource is transmitted only within the coverage area. When an object that has the resource leaves the coverage area, it removes the resource from its memory. Each object that has ever transmitted the announcement remembers this account number and the number of transmissions it has done. When the object connects to the centralized credit clearance system (possibly through a wired network or a cellular wireless network), it submits the record. The clearance system charges the producer based on the number of transmissions and pays to the object accordingly. We refer to this variant as the *offline producer-paid* policy, or OPP. Clearly, in OPP the cost per each reached object is exactly the commission fee f .

4.1.2 Simulation Setup

We built a synthetic road network as a grid. The distance between two neighboring grid points is 0.1 mile (approximately the length of one street block). The home of the resource is located at the center of the grid. The

coverage area with diameter h is defined to be the maximum area such that all the location points inside the area are within $h/2$ distance units away from the home of the resource. We use the *block-distance* as the distance metric. The block-distance between a point and the home of the resource is the length of the shortest path between them. Under this definition, the coverage area is the diamond area centered at the home of the resource, and the diameter of the diamond area is h . Figure 1 shows a coverage diamond area with diameter 6 blocks.




Figure 1: A coverage diamond area with diameter 6 blocks

There are 4 parameters to each simulation run, namely the traffic speed v , the transmission range r , the traffic density g , and the diameter of the coverage area d . v ranges from 10 to 50 miles/hour with increment 10. r has values from 10, 50, 100, 150, 200 meters. g has values 100 and 500 objects per square mile. d ranges from 0.4 to 2 miles with increment 0.4. For SAS there is an extra parameter, namely the initial budget. The initial budget ranges from 0 to 1000 with increment 50. The commission fee is fixed to be 1 for SAS.

At the beginning of each simulation run, $g \times d^2 / 2$ objects are generated as follows. Denote by S the set of the points on the border of the road network. The start location of each object O is randomly picked up from S . O travels through the network along a straight line at a constant speed that is randomly picked up between v and $v+10$ miles/hour. When O reaches the destination, it is removed from the system. At the same time a new object is generated in the same random fashion.

During each simulation run, we collect two values. The first one is how many objects have ever traveled the coverage area during the simulation run. The second one is how many of these objects have the resource some time during they travel the coverage area. The percentage of the second value over the first value is referred to as the *percentage of reached*. For SAS we collect another value, which is the total number of coins that are spent by the producer.

4.1.3 Simulation results

4.1.3.1 Simple Advertising Strategy (SAS)

Percentage of reached versus initial budget. Figure 2 shows the percentage of reached as a function of initial budget for different values of traffic density. The figure indicates that the percentage of reached initially increases rapidly as the initial budget increases. At this stage, the increased initial budget is used to support more transmissions within the coverage area. However, after a certain point, it becomes stable. Further increase of initial budget generates only very marginal increase to the percentage of reached. For example, when the traffic density $g=500$ objects per square mile, this point is around $b=50$ (b is the initial budget), and the percentage of reached is about 90% at this point. Increasing the initial budget from 50 to 1000 introduces only 2 percentage points of increase to the percentage of reached. The percentage of reached is maximized when the initial budget is infinite or when the commission fee is zero. The maximum percentage of reached when $g=500$ is 96% (not displayed in Figure 2).

These observations suggest that at $b=50$, the coverage area is almost saturated already, meaning that the initial budget is big enough to support all the transmissions of the resource that could have occurred if the initial budget is infinite. Further increase to the initial budget are mostly wasted in the interactions outside the coverage area. Therefore, 50 is the value of the initial budget that leads to the best trade-off between the effect and the cost of advertisement. In other words, the optimal initial budget is 50 for the traffic environment with the given parameters.

Similarly, the optimal initial budget is 200 when the traffic density $g=100$. By simulations, we are able to identify the optimal initial budget values for other traffic environments. And this is the significance of Figure 2 from the decision-making perspective.




Figure 2: Percentage of reached as a function of initial budget for different values of traffic density ($r=50$ meters, $v=40$ miles/hour, $d=2$ miles)

Cost per reached object. The cost per reached object is computed by dividing the total number of objects that have the resource some time during they travel the coverage area by the total number of coins that are spent by the producer. Figure 3 shows cost per reached object as a function of initial budget for different values of traffic density. The figure indicates that the cost per reached object increases linearly as the traffic density increases. The minimum cost per reached object is 1 and it is achieved when the initial budget is 0. However, the initial budget of 0 leads to the worst advertisement coverage (see Figure 2), because with 0 initial budget only the objects that pass by the producer can get the resource, and there is no information propagation at all. When the traffic density $g=500$, the cost per reached object is about 3 for the optimal initial budget (50).




Figure 3: Cost per reached as a function of initial budget ($r=50$ meters, $v=40$ miles/hour, $d=2$ miles)

4.1.3.2 Offline Producer-paid (OPP)




Figure 4: Percentage of reached as a function of the diameter of the coverage area for different values of traffic density ($r=50$ meters, $v=40$ miles/hour)

Impact of size of the coverage area. Figure 4 shows the percentage of reached as a function of the size of the coverage diamond for different values of traffic density. The figure indicates that the percentage of reached increases as

the size of the coverage diamond increases. This is because the bigger the diamond area, the more interactions an object tends to have, and thus the bigger chance it will receive the resource sometime during its travel in the diamond area.

Impact of transmission range. Figure 5 shows the percentage of reached as a function of the transmission range. The figure indicates that the percentage of reached increases as the transmission range increases. Intuitively, the bigger the transmission range, the more intensive interactions an object tends to have, and thus the bigger chance an object will receive the resource within the diamond area.




Figure 5: Percentage of reached as a function of transmission range for different values of traffic density ($v=40$ miles/hour, $d=2$ miles)

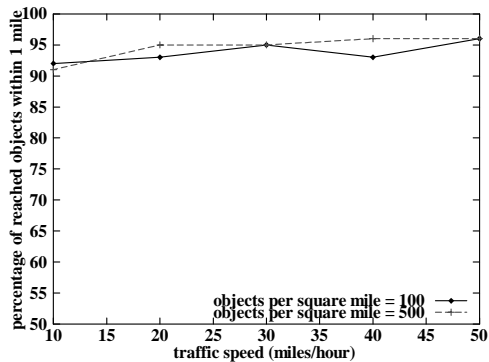


Figure 6: Percentage of reached as a function of traffic speed for different values of traffic density ($r=50$ meters, $d=2$ miles)

Impact of traffic speed. Figure 6 shows the percentage of reached as a function of traffic speed. The figure suggests that traffic speed does not have much impact to performance. Intuitively, increasing the traffic speed essentially creates a “fast-replay” of the lower speed scenario. This is similar to “fast-replaying” a video: everything proceeds faster, but all the occurrences (here interactions between objects) are truly

preserved. Note that, however, in the simulations we assumed that each resource exchange occurs instantaneously and correctly. With this assumption, the interaction is always successful regardless for how long two objects are within the transmission range of each other. This is not necessarily true in a realistic environment where two objects may leave out of the transmission range before they finish the resource exchange. Taking into consideration the effect of traffic speed on the successful interaction rate is part of our future work.

Impact of traffic density. Finally, observe that the percentage of reached in a more dense traffic environment is always higher than that in a less dense one, regardless of the other traffic parameters. This is because higher traffic density results in more intensive object interactions and thus increases the probability that an object receives the resource within the coverage area.

4.2 Analysis of the Consumer-Paid Policy

In this subsection we evaluate how much a consumer gains when searching for a resource using the resource information it bought, compared to not using the resource information. Specifically, we evaluate how much time is saved when a consumer uses the resource information to search a resource. This saving justifies why a consumer should and why she is willing to pay for the received resource information. We consider searching of resources that can be used by a single customer at a time (e.g. a parking slots). These resources are referred to as *competitive resources*. First we describe two strategies for a consumer to search competitive resources, one using resource information and another not. Then we compare these two strategies.

The first resource searching strategy is *blind search*, or BS. With this strategy, a consumer drives around the area where a resource of interest could possibly be located, and it occupies the first resource that is available at the time when the consumer reaches it. The second strategy is *information guided search*, or IGS. With this strategy, a consumer starts with a blind search, until either a resource is available at the time when the consumer reaches it, or some resource R that is located in the searching area is received. In the latter case, the consumer goes to R if the distance between R and the consumer’s current location is within a certain threshold h^3 . If another resource R’ is received during traveling to R, then the consumer goes to either R’ or R whichever is closer.

We simulated searching of parking slots. We considered a simplistic situation, in which there are only one consumer and two parking slots in the system. We assumed that resource information, once announced, propagates at a constant speed. The simulations results indicate that when

³ h depends on the contention and the density of the searched resource type.

the information is propagated 20 or more times faster than the traffic speed, the searching time of IGS is half of that of BS. The detailed simulation setup and the complete results are omitted due to space limitations.

5. Relevant Work

Different resource discovery architectures (SLP [14], Jini [15], Salutation [16], and UpnP [17]) 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 [21, 22], architecture [26], transactional and query processing issues [20, 23], economics [25] have all been studied in previous works. There are two major differences between these works and ours. 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 traffic density can vary in a big range from rush hours to midnight. The underlying communication network is thus subject to topology changes and disconnections. 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 [8], DHTs like [9]). Second, in our environment information sharing is opportunistic. Pre-defined data access structures such as search routing tables used in Gridella [10] are impractical in mobile ad hoc networks. They are replaced by opportunistic peer-to-peer interactions.

Our work is also relevant to self-organized mobile ad-hoc networks, which have been investigated extensively (see [27] for a survey). Most existing works in this area studied message forwarding and routing. The main issue is how to send a message from one node to another, possibly by multi-hops. The issue is treated independently of the semantic properties of applications. In opportunistic peer-to-peer networks, however, what is needed is data-centric routing, with group communications referenced by attribute-based names, rather than being node-centric with communications referenced by global node identifiers. For

example, a vehicle would query the average speed within a certain area, rather than the speed of a particular vehicle.

Incentive mechanisms have been studied in the context of mobile ad-hoc networks (see for example [7, 29]). The issue is how to stimulate a mobile node to cooperate in forwarding information to other nodes. There are two major differences between our research and these incentive mechanisms. First, they consider abstract messages without looking into the semantics of the data included in the message, whereas in our model we deal specifically with spatio-temporal resources. Second, here the network is subject to partitions, and the communication mode can be unicast (e.g. when returning an answer to a query), multicast (e.g. when delivering a query to a geographic area), or broadcast (e.g. when spreading an advertisement). The combination of these characteristics makes our problem different than the existing work in this area. However, we expect to be able to draw on existing ideas in this area.

Incentive mechanisms have also been studied for static peer-to-peer networks (see e.g. [25]). In this case the static nature of the problem is often relied upon heavily, for example, by ‘punishing’ a user that is found non-cooperative over time. Such a longer-time perspective is missing in our opportunistic paradigm, which may rarely involve the same pair of moving objects in an exchange.

A lot of work has been done on sensor networks in the database community (e.g. [19, 30]). However, the existing work considers only stationary sensors, whereas in our model there are both stationary sensors (stations) and moving sensors (moving objects).

Finally, in our prior work [28] we studied the way in which resource information is propagated in space and time in an opportunistic peer-to-peer environment, but did not address incentive issues.

6. Conclusion and Future Work

In this paper we discussed incentive mechanisms for stimulating moving objects to participate in resource exchange in an opportunistic peer-to-peer network. We distinguished between two types of resources, namely producer-paid resources and consumer-paid resources, and developed incentive mechanisms for these two types respectively. By experimental analysis, we showed that the payment made by producers or consumers is well justified by the benefit they gain, which in turn justifies the proposed incentive models.

We also studied important security issues, which are not discussed in the paper due to space limitations. For example, how to prevent an object from generating fictitious resources, and be paid for them? How to prevent a broker from modifying a resource? How to prevent a consumer-paid resource being overheard by an intruding-consumer

that does not pay? We resist again these attacks based on some public key infrastructure.

A lot remains to be done in the future. For example, privacy and security need further study. Atomicity of resource exchange transactions needs to be addressed. 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.

7. 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] J. Haartsen, M. Naghshineh, J. Inouye, O. J. Joeresson, and W. Allen. Bluetooth: Vision, Goals, and Architecture. *ACM Mobile Computing and Communications Review*, 2(4):38-45, October 1998.
- [5] <http://www.etsi.org/frameset/home.htm?/technicalactivities/ITS/ITS.htm>
- [6] <http://www.oki.com/en/otr/html/nf/otr-158-2-4.html>
- [7] L. Buttyan and J. Hubaux. Nuglets: a Virtual Currency to Stimulate Cooperation in Self-Organized Mobile Ad Hoc Networks. *Technical Report DSC/2001/001*, Swiss Federal Institute of Technology – Lausanne, Dept. of Communication Systems, 2001.
- [8] Gnutella website. <http://gnutella.wego.com>
- [9] I. Stoica, et al. Chord: A Scalable Peer-to-Peer Lookup Service for Internet Applications. In *Procs. ACM SIGCOMM*, 2001.
- [10] Karl Aberer, Manfred Hauswirth, Magdalena Puceva, Roman Schmidt. Improving Data Access in P2P Systems, *IEEE Internet Computing*, 6(1), January/February 2002.
- [11] A. Markowetz, et al. Exploiting the Internet As a Geospatial Database, *International Workshop on Next Generation Geospatial Information*, 2003.
- [12] A. Pfitzmann, B. Pfitzmann, and M. Waidner. Trusting Mobile User Devices and Security Modules. *IEEE Computer*, February 1997.
- [13] W. Diffie, P.C. van Oorschot, and M.J. Wiener. Authentication and Authenticated Key Exchanges. *Designs, Codes and Cryptography* 2 (1992), 107-125.
- [14] E. Guttman, C. Perkins, J. Veizades, M. Day, Service Location Protocol, Version 2. *RFC2608*, June 1999. <http://www.ietf.org/rfc/rfc2608.txt>
- [15] W. Keith Edwards, *Core JINI*, Prentice Hall, 1999.
- [16] <http://www.salutation.org/>
- [17] Universal Plug-and-Play (UPnP) Forum, *Microsoft Corporation*. <http://www.upnp.org>
- [18] N. Daswani, H. Molina, B. Yang. Open Problems in Data-Sharing Peer-to-Peer Systems. *ICDT 2003*.
- [19] L. Qiao, D. Agrawal, A. Abbadi. Supporting Sliding Window Queries for Continuous Data Streams. *SSDBM 2003*, Cambridge, MA, July 2003.
- [20] A. Popovici, G. Alonso. Ad-Hoc Transactions for Mobile Services. Proc. Of the 3rd VLDB Workshop on Transactions and Electronic Services. 2002.
- [21] V. Kalogeraki, A. Delis, D. Gunopulos: Peer-to-Peer Architectures for Scalable, Efficient and Reliable Media Services. *IPDPS 2003*
- [22] 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.
- [23] P. Kalnis, et al. An adaptive peer-to-peer network for distributed caching of OLAP results. Proc. Of SIGMOD 2002, pp. 25-36.
- [24] D. Turner and K. Ross. The Lightweight Currency Protocol. *Internet Draft*. September 2003, <http://cis.poly.edu/~ross/papers/draft-turner-lcp-00.txt>.
- [25] 1st Workshop on the Economics of P2P Systems website, <http://www.sims.berkeley.edu/research/conferences/p2pecon>
- [26] M. Papazoglou, B. Krämer, J. Yang. Leveraging Web-Services and Peer-to-Peer Networks. In *Proceedings of CAiSE 2003*, June, 2003.
- [27] E. Royer and C. Toh, A Review of Current Routing Protocols for Ad Hoc Mobile Wireless Networks, *IEEE Personal Comm.*, pp. 46-55, Apr. 1999.
- [28] B. Xu, A. Ouksel, O. Wolfson. Opportunistic Resource Exchange in Inter-vehicle Ad Hoc Networks. *Proc. of MDM 2004*. Jan. 2004.
- [29] Sheng Zhong, Jiang Chen and Yang Richard Yang. Sprite: A Simple, Cheat-Proof, Credit-Based System for Mobile Ad-Hoc Networks. In *Proceedings of IEEE INFOCOM 2003*.
- [30] P. Bonnet, et al. Towards Sensor Database Systems. *Proceedings of the Second International Conference on Mobile Data Management*, Hong Kong, Jan. 2000.
- [31] <http://www.its.dot.gov/speeches/madridvii2003.ppt>
- [32] <http://www.cartalk2000.net>