Benefit and Pricing of Spatio-temporal Information in Mobile Peer-to-Peer Networks

Bo Xu and Ouri Wolfson
University of Illinois at Chicago
{boxu,wolfson}@cs.uic.edu

Naphtali Rishe
Florida International University
rishen@cs.fiu.edu

Abstract

In this paper we examine the dissemination of reports about resources in mobile peer-to-peer networks, where moving objects communicate with each other via short-range wireless transmission. Each disseminated report represents information about a spatial-temporal resource, such as the availability of a parking slot at a particular time and location. We introduce an architecture and a data model for dissemination of such reports. We develop an analytical model to quantify the benefit of report dissemination, where the benefit is measured in terms of time-saving. We further propose an incentive mechanism for participation in resource dissemination, and a method of pricing the resource information.[1]

1. Introduction

A mobile peer-to-peer network is a set of moving objects that communicate via short-range wireless technologies such as IEEE 802.11 and Bluetooth. With such communication mechanisms, a moving object receives information from its neighbors, or from remote objects by multi-hop transmission relayed by intermediate moving objects. A killer application of mobile peer-to-peer networks is resource discovery in transportation. For example, the mobile peer-to-peer approach can be used to disseminate the information of available parking slots, which enables a vehicle to continuously display on a map to the driver, at any time, the available parking spaces around the current location of the vehicle. Or, the driver may use this approach to get the traffic conditions (e.g. average speed) one mile ahead. Similarly, a cab driver may use this approach to find a cab customer, or vice versa.

A mobile peer-to-peer network 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, a passenger who arrives at an airport may use this approach to find another passenger for cab-sharing from the airport to downtown, so as to split the cost of the cab. Furthermore, 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.

This approach can also be used for endangered species animal assistance. For example, sensors can be installed on wild animals. Each sensor monitors its carrier's health condition, and it disseminates a report when an emergency symptom is detected. Thus we use the term moving objects to refer to all, vehicles, pedestrians, and animals.

In this paper we quantify the value/benefit introduced by the dissemination of resource reports in mobile peer-to-peer networks. Each report represents the availability of a resource (parking slot, taxi-cab customer, cab-sharing passenger, etc.). Since a user uses the mobile peer-to-peer network to discover resources, we measure the value/benefit of the network in terms of the amount of search-time saved in discovering and taking possession of the resource when using the network.

Observe that in all the applications given above, each resource is "transient", in the sense that it is valid only for a certain period of time. It also pertains to a particular location. In other words, it is spatio-temporal. For many of these resources, resource information benefits a user only if the user reaches the resource while it is valid. In other words, "resource

[1] This research is supported in part by NSF Grants 0326284, 0330342, ITR-0086144, 0513736, 0209190, EIA-0320956, EIA-0220562, and HRD-0317692.

discovery″ in these applications means discovering and taking possession of the resource, not only becoming aware of it. For example, a report that points to an available parking slot generates benefit only if the user reaches the slot before it is occupied; a report about a customer waiting for a cab generates benefit only if the cab driver reaches the customer before the customer leaves with another cab.

We consider an opportunistic peer-to-peer (OP2P) approach for information dissemination in mobile peer-to-peer networks. In this approach, resource-reports are announced to the neighborhood using short range wireless communication. The problem is that a local broadcast often cannot reach all the moving objects for which the information may be useful[2]. To address this problem, an object propagates the reports it carries to encountered objects (i.e. objects that come within transmission range) 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 moving objects by gossiping. Similarly, information about an accident or a taxi cab customer is propagated transitively.

In this paper we develop an analytical model for computing the benefit that is expected to be generated by a resource report in the mobile peer-to-peer environment. The benefit is measured by the amount of time saved when a consumer uses resource reports to capture a resource, compared to not using resource information.

Like many other peer-to-peer systems or mobile ad-hoc networks, the ultimate success of our OP2P paradigm heavily relies on cooperation among users. However, a peer-to-peer network such as a vehicular network may consist of moving objects owned by different authorities and each moving object has its own goal. In such an environment, the owner of a moving object may decide not to cooperate in supplying or relaying information. So the second issue we address in this paper is to develop an economic model to stimulate cooperation.

We consider incentive mechanisms that are based on virtual currency [5]. Each mobile node carries virtual currency in the form of a coin counter that is protected from illegitimate manipulation by a trusted and tamper resistant hardware module [6]. A consumer pays the sender a fee for each report it receives.

Assume now that the monetary value of a coin is fixed arbitrarily (e.g. 1 cent = 1 coin). A natural question arising from our economic model is how to ″price″ the reports, namely how many coins a resource report is worth? The production of the report is virtually free, thus we propose to price the report based on its benefit to the consumer (and hence the connection between the two subjects of this paper, namely the benefit of reports and the economic model). In this paper we define the requirements of an economic model for trading spatial-temporal resource-reports, and propose an economic model that satisfies these requirements.

The P2P method provides a search engine for local resources that may be only temporarily available. The same search engine can be provided by a central-site that is updated wirelessly and broadcasts the resource information. The advantage of the mobile P2P approach is that it avoids the costs of maintaining and communicating to and from a central site. Such communication would involve the regulated wireless spectrum. In contrast, the mobile P2P communication occurs over the unregulated spectrum.

In summary, this paper makes the following contributions. We propose a method of quantifying the benefit of spatio-temporal resource-information in mobile peer-to-peer networks. Intuitively the benefit is given by the time reduction in resource capturing. To address the incentive problem, we offer a solution in the sense of an economic model for this environment. An additional problem and solution offered in this context is a method of pricing the resource information.

The rest of the paper is organized as follows. Section 2 describes the architecture. Section 3 quantifies the benefit of a resource report. Section 4 presents the economic model. Section 5 discusses security and atomicity issues. Section 6 compares the paper to relevant work. Section 7 concludes the paper.

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, and so on.

[2] In heavily built downtown areas the effective distance of such wireless communication technologies may be 10-50 meters.

Formally in our model there are N *resource types* T_1, T_2, ..., T_N . At any point in time there are M *resources* R_1, R_2, ..., 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. 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. 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. [7]). 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.

2.2. Peers and Validity Reports

The system consists of two types of peers, namely fixed hotspots and moving objects. Each peer m 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 m that produces validity reports, called the *report producer* for R. A peer may be the report producer for multiple resources. Each report $a(R)$ contains at least the following information, namely *resource-id*, *create-time*, and *home-location*. Resource-id is the identification of R that is unique among all the resources of the same type in the system; create-time is the time when report $a(R)$ is created (it is also the time when R is sensed valid); home-location is the home of R.

In the parking slots example, a sensor in the parking slot monitors it, and when the slot becomes free, it produces a validity report. In the car accident example, the report is produced by the sensor that deploys the air-bag.

We say that $a(R)$ is a type T_i report if R is a type T_i resource. Let $a(R)$ be a type T_i report. At any point in time, a peer m is either a *consumer* or a *broker* of $a(R)$. m is a consumer of $a(R)$, and $a(R)$ is a *consumer report* to m, if m is attempting to discover or find a type T_i resource. m is a broker of $a(R)$ and $a(R)$ is a *broker report* to m, if m is not attempting to discover/find T_i but is brokering $a(R)$, i.e. the only purpose of m storing $a(R)$ is to relay it to other peers.

2.3. Reports Relations

There are two relations in the reports database of a peer m. One is the *consumer relation*, which stores all the reports that m knows about and for which m is a consumer. Another is the *broker relation*, which stores all the reports that m knows about and for which m is a broker. The two relations have a common object-relational schema. The schema contains three columns: (i) resource-type which indicates the type of the reported resource; (ii) resource-id; (iii) report-description, which is an abstract data type that encapsulates all the attributes of a report. All the report description data types inherit from a single data type called *AbstractReport*. AbstractReport contains two attributes, namely create-time and home-location. Thus every report description data type has these two attributes.

2.4. Peer-to-Peer Report Exchange

We assume that 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.

The user of a peer specifies to the communication module what types of validity reports she is interested in consuming. There are several strategies to determine which report types to broker, e.g. MADM (Multiple Attribute Decision Making) [2] or machine learning [1]. But we do not discuss this subject since it is not relevant to the main topic of this paper. 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, A acquires reports that A is interested in and B has in B's broker relation. In turn A provides reports that A has in A's broker relation and B is interested in.

3. Benefit of a Report

In this section we develop an analytical model that computes how much a user gains when capturing a resource using validity reports, compared to the case where no resource information is used. The model is based on a simplified system environment, which is described in subsection 3.1. In subsection 3.2 we provide the analytical model and the main theorem. In subsection 3.3 we prove the theorem and comment on its experimental validation.

3.1. Environment

Let T be a resource type. We assume that there is only one consumer M that has T in its consumer relation, namely that is searching for a resource of type T (The rationale for this assumption will be given shortly). The other peers are brokers of type-T reports. M starts to search for the resource at time 0. It moves along a simple closed curve[3] C with length l at a constant speed v. For example, a vehicle circles within a geographic area to find a parking slot. Resources are generated by a Poisson process in time with intensity λ and uniformly along the curve C. The length of the valid duration of each resource follows an exponential distribution with mean u. We use this distribution to model the competition for the resource generated by other consumers. In other words, instead of analyzing multiple consumers competing for resources, we analyze a single consumer; but each resource is valid for only u time units approximately. This accommodates the scenario of endanger species monitoring and that of traffic resources in the sense that in these there is no competition.

M receives a report $a(R)$ at time 0 at location G, and it does not receive any other report afterwards. With the initial moving direction of M, the travel distance from G to the home location of R is d. d is greater than $l/2$. $a(R)$ was transmitted by its producer at time $-t_0$ ($t_0>0$), and we call t_0 the *age* of $a(R)$. We use the age to model the delay of OP2P dissemination.

In section 4.4 we discuss the extension of the environment to more applicable cases, via a simulation system.

[3] A simple closed curve is a curve that is closed and does not intersect itself, also called Jordan curve.

3.2. Analytical Model

We analyze the benefit of $a(R)$ in terms of the amount of time that is saved for M to capture a resource when M uses $a(R)$, compared to if it does not use $a(R)$. Therefore we consider two cases depending on whether M uses the report or not. In the first case M keeps its original direction as if it did not receive $a(R)$. In the second case M changes direction and goes to R. Denote by P_1 the length of the time period starting from 0 until M captures a resource in the first case (i.e. not using the report), P_2 the length of the time period starting from 0 until M captures a resource in the second case (i.e. using the report). Let us define function $G(x)$ which is used by the theorem we are going to present.

$$G(x)=\frac{l\cdot(1-(1+u\cdot\lambda\cdot x/l)\cdot e^{-u\cdot\lambda\cdot x/l})}{u\cdot\lambda\cdot v}+$$

$$e^{-u\cdot\lambda\cdot x\cdot l^{-1}}\cdot e^{-\frac{t_0+x/v}{u}}\cdot\frac{x}{v}+$$

$$(1-F_1(\frac{x}{v}))\cdot\frac{(l+u\cdot\lambda\cdot x)+(l+u\cdot\lambda\cdot l)\cdot e^{-u\cdot\lambda\cdot(l-x)/l}}{u\cdot\lambda\cdot v}+$$

$$(1-F_2(\frac{l}{v}))\cdot\frac{l}{u\cdot\lambda\cdot v\cdot(1-e^{-l/(u\cdot v)})}$$

where

$$F_1(\frac{x}{v})=1-e^{-u\cdot\lambda\cdot x/l}\cdot(1-e^{-\frac{t_0+x/v}{u}})$$

$$F_2(\frac{l}{v})=F_1(\frac{x}{v})+(1-F_1(\frac{x}{v}))\cdot(1-e^{-u\cdot\lambda\cdot(l-x)/l})$$

x can be any real number but we will use only two values d and l–d. As will be shown in section 3.4, $G(d)$ is the expected value of P_1 and $G(l\text{-}d)$ is the expected value of P_2.

Theorem 1: P_1-P_2 is a random variable. The expected value of P_1-P_2 is $G(d)-G(l-d)$ □

We will derive $G(x)$ and prove the theorem in subsection 3.3.

3.3. Proof of Theorem 1

Due to space limitations we only list the lemmas that lead to Theorem 1 without proving them. Recall that P_1 denotes the length of the time period starting from 0 until M captures a resource without using the report; P_2 denotes the length of the time period starting from 0 until M captures a resource using the report. Define random variable P to be the length of the time period starting from 0 until M captures a resource, under the condition that R was never generated and M did not receive $a(R)$. Let $\lambda' = \lambda / l$.

Lemma 1: The distribution function of P is $D(p)$=

$$\begin{cases} 1-e^{-u\cdot\lambda'\cdot v\cdot p} & p \le l/v \\ D(\frac{l}{v})+(1-D(\frac{l}{v}))\cdot(1-e^{u\cdot\lambda'\cdot v\cdot e^{-l/(u\cdot v)}\cdot(p-l/v)}) & p > l/v \end{cases}$$

□

Lemma 2: The distribution function of P_1 is $D_1(p)$=

$$\begin{cases} 1-e^{-u\cdot\lambda'\cdot v\cdot p} & p < d/v \\ D(\frac{d}{v})+(1-D(\frac{d}{v}))\cdot e^{-\frac{t_0+d/v}{u}} & p = d/v \\ D_2(\frac{d}{v})+(1-D_2(\frac{d}{v}))\cdot(1-e^{-u\cdot\lambda'\cdot v\cdot(p-d/v)}) & \frac{d}{v} < p \le \frac{l}{v} \\ D_2(\frac{l}{v})+(1-D_2(\frac{l}{v}))\cdot(1-e^{u\cdot\lambda'\cdot v\cdot e^{-l/(u\cdot v)}\cdot(p-l/v)}) & p > l/v \end{cases}$$

□

Lemma 3: The distribution function of P_2 is $D_2(p)$=

$$\begin{cases} 1-e^{-u\cdot\lambda'\cdot v\cdot p} & p < \frac{l-d}{v} \\ D(\frac{l-d}{v})+(1-D(\frac{l-d}{v}))\cdot e^{-\frac{t_0+(l-d)/v}{u}} & p = \frac{l-d}{v} \\ D_1(\frac{l-d}{v})+(1-D_1(\frac{l-d}{v}))\cdot(1-e^{-u\cdot\lambda'\cdot v\cdot(p-(l-d)/v)}) & \frac{l-d}{v} < p \le \frac{l}{v} \\ D_1(\frac{l}{v})+(1-D_1(\frac{l}{v}))\cdot(1-e^{u\cdot\lambda'\cdot v\cdot e^{-l/(u\cdot v)}\cdot(p-l/v)}) & p > l/v \end{cases}$$

□

Lemma 4: The expected value of P_1 is $G(d)$.

Lemma 5: The expected value of P_2 is $G(l\text{-}d)$.

where $G(x)$ is defined in subsection 3.2.

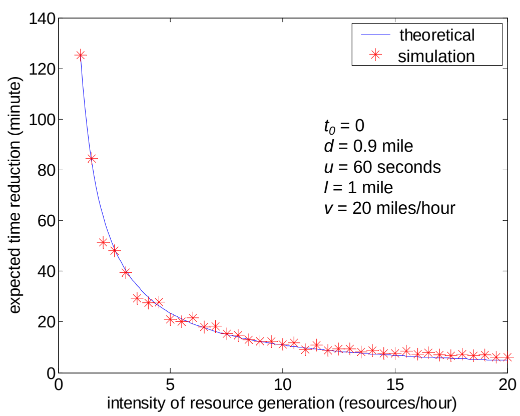

Figure 1: Comparison between analytical model and simulation results

We also conducted simulations to obtain the time reduction caused by a report. The simulation setup was the same as used in theoretical analysis. Figure 1 shows two curves of expected time reduction versus intensity of resource generation, generated by the analytical model and simulations respectively. The theoretical analysis formalizes the simulation results we obtained because the analytical and simulation curves are identical.

4. The Economic Model

In this section we introduce an economic model that stimulates peers to participate in report dissemination even if they are not interested in using a resource. The economic model needs to satisfy the following requirements:

It should handle two categories of reports, depending on whether the producer or the consumer pays for the reports. Reports that the owner is interested in advertising are *producer-paid*. Reports that the consumer is interested in knowing are *consumer-paid*. A resource may have both producer-paid and consumer-paid reports, if both the producer and the consumer are willing to pay for the reports. For example, reports that include the location of a gas station may be producer-paid because the gas station wishes to advertise them to neighboring vehicles. They may also be consumer-paid because a consumer may be willing to pay for a gas station report if he really needs one. Similarly for taxi-cab requests and reports of available parking slots. In this paper we concentrate on consumer-paid reports.

It should consider peers that may be producers, consumers, and brokers. For consumer paid reports, both producers and brokers should be incentivized. For producer paid reports, brokers should be incentivized.

It should allow any peer to turn-off the spatio-temporal information module. But if it turns on the spatio-temporal information module, then the module behaves according to the economic model.

It should protect from the following attacks:

1. A peer creates and sells fictitious validity reports.
2. A propagator modifies a report.
3. A consumer-paid report is overheard by an intruding-consumer that that does not pay;
4. A peer illegally increases its virtual currency counter.
5. A consumer buys a report, captures the resource, then sells the useless report.

Now we present our solution that satisfies the above requirements. Section 4.1 introduces two fundamental components of our economic model, namely virtual currency and the security module. Section 4.2 describes the consumer-paid trading policy. Section 4.3 discusses how a consumer decides what reports to buy and act on. Section 4.4 discusses pricing of validity reports.

4.1. Virtual Currency and Security Module

The system circulates a virtual currency called *coins*. The coins owned by each peer is represented by a coin counter that is physically stored in that peer. The coin counter is decreased when the peer pays out for buying validity reports and increased when the peer earns in for selling. Each peer has a trusted and tamper resistant hardware module called the *security module*. The module exports to the environment a fixed set of input/output primitives. A common example of a low-cost security module is smart card with an embedded one-chip computer [6]. The coin counter is stored in the security module and thus is protected from illegitimate manipulation. Each coin is bought for a certain amount of real money but it cannot be cashed for real money, and therefore the motivation for breaking into the security module is significantly reduced. The validity reports database, including the consumer relation and the broker relation, are stored in the security module. A standard electronic cash encryption model can be used in addition to the security module.

When two moving objects m_1 and m_2 encounter each other, if both m_1 and m_2 have their security module open, then m_1 and m_2 start a secure session to trade validity reports [4]. The trading policy is implemented in the security module. For each resource type T, the owner/user of a moving object may decide not to participate in the exchange of type T reports. The owner/user may also turn off the security module. However, if it participates in the game, then security module behaves according to the economic model.

4.2. The Consumer-paid Policy

Each report $a(R)$ is acquired by the security module of a peer in one of the following two modes:

1. *Consumer*. In consumer mode, the report $a(R)$ is saved in the consumer relation. The consumer relation is accessible by the user so the user can read $a(R)$. The consumer buys new reports according to some strategy (Report buying/usage strategies are discussed in section 4.3) but cannot sell them. The price of a report is a function of the relevance of the report.

2. *Broker*. In broker mode the report $a(R)$ is saved in the broker relation. The security module simply stores $a(R)$ and forwards it to other peers. A broker pays a percentage of the price of the report. It is paid the same percentage when selling the report to another broker, and it is paid the full price when selling the report to a consumer. How to setup the percentage to maximize the incentive is a subject of our future work. The received payment constitutes the incentive of the broker to participate in the game. A broker may sell $a(R)$ to multiple consumers or brokers. A producer always operates in broker mode for the reports it transmits. To prevent a consumer from purchasing a report at broker price, the broker relation is not accessible by the user. Thus the user cannot benefit from a broker-report as a consumer by viewing the information in the report.

Validity reports acquired in consumer mode are consumer reports, and reports acquired in broker mode are broker reports. A report cannot switch between broker and consumer relations.

4.3. Report Buying/Acting Strategies

When multiple consumers hear about the same competitive resource (such as a parking slot or a cab customer), they may all head to that resource, leading to contention. In order to address this phenomenon of ″herding″, a consumer needs to be selective when buying and acting on reports. We propose the following strategy for a consumer to decide what reports to buy and act on. We assume that for each resource type T there is a probability distribution $Prob(T)$ that describes the length of the valid time of type T. This distribution can be obtained empirically from historical data. In sections 2 and 3 we assumed the exponential distribution but it does not necessarily have to be so.

Now assume that each consumer M knows this distribution. When M decides whether to buy a report $a(R)$ of T, M uses $Prob(T)$ to compute the probability that R will remain valid when M reaches it, based on how much time has passed since $a(R)$ has been generated and how long it will take for M to reach R. This probability is referred to as the *capture probability* of R (In the proof of Theorem 1 we showed how to compute this probability for exponential distribution). If the capture probability of R is higher than a certain threshold (e.g. 0.5), then M buys $a(R)$ and goes to R. If on the way to R consumer M receives another report $a(R')$ and the capture probability of R' is higher than that of R, then M buys $a(R')$ and goes to R'. In other words, if the capture probability is lower than the threshold then the consumer assumes that it is not worthwhile to purchase the report, and does not pursue the respective resource. The capture probability threshold prevents all, but the most likely consumers to capture

[4] The secure session is established based on the public key infrastructure described in 5.1. The details about the establishment of the secure session are omitted in this paper due to space limitations.

it, from pursuing the resource. In other words, the concept of capture probability threshold addresses the herding phenomenon.

Observe that we implicitly made two assumptions in order to determine how much time has passed since $a(R)$ has been generated and how long it will take for M to reach R. The first assumption is that the consumer knows its location when receiving $a(R)$, so the travel time needed to reach R can be computed. The second assumption is that the clock between the report producer and the consumer is synchronized, so the age of the report can be accurately computed. Both assumptions can be satisfied if each peer is equipped with a GPS that reports both location and time. If/when the user moves indoors, the latest outdoors GPS reading can be used for (albeit imprecise) location determination and for clock synchronization.

4.4. Pricing of Availability Reports

In traditional commerce, a good or service is priced based on the cost to the seller of producing the good or service, plus a profit. In our model, the production cost is zero. Thus we propose that the price of a report is based on its value to the consumer buying it, namely the amount of time it is expected to save consumer for resource discovery. The reduction of discovery time can be converted to a certain amount of real money. For example, if on average a consumer is willing to spend 2 cents in order to save one minute, then the average value of a minute is 2 cents. Let D be the dollar worth of a coin, Q be the dollar worth of a unit of time, and E be the amount of time $a(R)$ is expected to save. The price of $a(R)$ in coins is $\frac{Q \cdot E}{D}$. For example, if each coin is bought for 1 cent, on average a consumer is willing to pay 10 cents to save one minute, and $a(R)$ is expected to save 12 seconds (0.2 minute), then the price of R is $10\times0.2/1=2$ coins.

D and Q are inputs to the system. D represents the rate between the real money and virtual currency, and can be an arbitrary value (e.g. 1 dollar = 1000 coins). Q represents the dollar worth of a unit of time averaged among the consumers. So both P and Q are constants. Note that the dollar worth of a unit of time may be different for different consumers. However, Q is the average across all the consumers. For example, Q could be determined as a percentage of the average salary per time unit, or it could be generated from a survey.

As revealed by the analysis in section 3, the benefit E varies depending on many parameters including u the mean of valid duration of resources, λ the intensity with which resources are generated, and v the motion speed of the consumer. Thus the price of R also depends on these parameters. The price of R varies in the same sense that the dollar worth of an airline mile varies depending on ticket prices, seat availability, season, and many other parameters. However, our method enables to price a report similarly to the way United Airlines has priced a frequent flyer mile at 3 cents [20].

Clearly, section 3 analyzes a simple case that is limited in its applicability (motion in a circle). But it demonstrates the issues involved in an analytical model, and provides the motivation for building a simulation system to determine the benefit of a report. We have built such a system, and verified the analytical results of sec. 3. This gives the assurance necessary to use the simulation system in more realistic environments (e.g. search for a taxi-cab customer in downtown Chicago).

In real time, the price of the resource is established by the security module. The module knows its time, location, and type of resource involved in the exchange. Based on the time and location it determines the age (t_0) and the distance (d) of a report. Based on the time, location, and resource type the security module also determines the environmental parameters, by a table-lookup. An entry in the table gives the mean of invalidation duration and the intensity of resource generation as functions of the area where the trade occurs, time of day, type of resource, etc. The table is embedded in the security module by the vendor of the OP2P system. The vendor populates the table, i.e. determines the environmental parameters, using a simulation system. Observe that the buyer and seller of a report have the same table, thus their price for the report traded should be identical, assuming that the buyer's location is used for the computation by both the buyer and the seller.

5. Discussion

5.1. Security Issues

In this subsection we discuss important security issues. For example, how the announced resources are protected from illegitimate hearing; how the resource exchange can be protected from various types of attacks.

We build the security mechanisms based on a public key infrastructure. The security module of each producer or each moving object that is willing to participate in the game has a public key and a private key. The public key is certified by a certificate authority that is trusted by all these parties. The private key, the public key certificate[5] and the public key of the certificate authority are stored in the tamper-resistant (as the car odometer) and read-resistant security module.

When a producer announces a resource, it signs the announcement using its private key and attaches to the announcement its public key certificate. The signed announcement and the public key certificate are sent as a single data item when the announcement is transmitted or propagated. When a producer and a moving object become neighbors, or when two moving objects become neighbors, their security modules run the STS protocol [12] to establish a symmetric session key. The STS protocol uses the aforementioned public key infrastructure to ascertain that the two communication parties are who they claim to be. The session key is used to protect the communication messages from disclosure or modification. With these mechanisms, the following attacks are protected against.

1. A moving object creates and sells fictitious validity reports. Prevented by the fact that the object does not know the private key of the producer security module, thus cannot ″sign″ the resource announcement.
2. A propagator modifies a report. Again prevented by announcement signature.
3. A consumer-paid report is overheard by an intruding-consumer that that does not pay. Prevented by the session key.
4. A moving object user illegitimately creates coins, i.e. increases its coin counter. Prevented by the fact that the coin counter is stored in the security module.

5.2. Transactional/Atomicity Issues

The transaction between two vehicles consists of a handshake initiation that includes the types of resources each one is interested in consuming/brokering, followed by the resource exchange and coin charge/credit for each resource. Observe that these operations must be executed as a distributed atomic transaction. For example, the credit of one account should be committed only if the debit of the other account is committed; and in turn, this should occur if and only if the corresponding report was received properly. Therefore, the transaction must be followed by a commit protocol. The problem is that, due to the high mobility at which the transaction occurs, the commit protocol between two moving objects may not begin or may not complete.

We propose to resolve this problem by a Mobile Peer-to-Peer Transaction (MOPT) mechanism which is a combination of an audit trail (or log) maintained online in the security module, and a central bank to which the audit trails of all moving objects are transmitted periodically, e.g. once a day. Our proposed MOPT mechanism has an online component that executes at the security module for each transaction, and an offline component.

The online component of MOPT at a security module S performs the following functions. It keeps a log of the reports that have been exchanged and the credit/debit charged for each one. The records of this log correspond to the log records in database transaction recovery. When a transaction completes unsuccessfully, then the user of S is still charged and can use the reports it received, and gets credit for the reports it (thinks it) sold. So if a broker B sent a report to a consumer C, but didn't receive the commit message, it still gets (temporary) credit.

The offline component of MOPT, at the end of the day sends to a central bank the logs of the transactions that completed unsuccessfully during the day. After receiving all the logs from all the moving objects, the central bank does the following for the transactions that completed unsuccessfully at one or both participants (thus it ignores transactions that completed successfully at both participants). If the same transaction completed unsuccessfully at both participants, then the traces from the respective security modules are used to settle the credit/charge to both accounts. In the example above, if C didn't receive the report, B's credit will be reversed. If the transaction completed unsuccessfully at only one of the participants, i.e. the transaction is absent from the other security module trace, this fact indicates how the account at the unsuccessful participant should be settled. In this case, in the example above, B's credit will be made permanent.

Observe that our MOPT mechanism needs to remember only the logs of unsuccessfully completed transactions, but can forget successfully completed transactions. Considering that moving objects may execute thousands of transaction per day, this is an important property.

Observe that this offline banking mechanism violates to some extent our principle of a completely decentralized economy. However, this overhead is relatively minor, and can be easily automated. Moreover, this overhead occurs offline whereas the real-time data dissemination is still decentralized.

[5] The public key certificate is a digital document encrypted using the private key of the certificate authority. The document contains some information (e.g. the name) of the certificate holder and the holder's public key.

6. Relevant Work

Traditional peer-to-peer approaches. A traditional peer-to-peer approach like Gnutella [21] could be used to search spatio-temporal resources, the problem addressed in this paper. In Gnutella, a query for a resource type (expressed by key words) is flooded on the overlay network (within predefined hops), and replies are routed back to the querying node along the same path as the query message. In other words, resource information is *pulled* by the querying node from the resource producer. This generates two problems in our context. First, since resources are transient and consumers do not know when they are generated, a consumer will have to constantly flood its query in order to catch resource information. Second, this does not work if there is not a path between the querying node and the resource producer. In our approach, a resource report is *pushed* by the resource producer to consumers via opportunistic dissemination and the dissemination area is automatically bounded by information prioritization ([28]). Gridella [23] and DHT systems such as Chord [22] have similar problems as Gnutella in that they use a *pull* model. In addition, Gridella and DHT systems require that the complete identifier (or key) of the searched data item be provided in a query, whereas in our case a consumer does not know a priori the keys of the searched resources.

Mobile P2P data dissemination. A lot of work has been done on data dissemination in mobile peer-to-peer environments (e.g. [16, 17, 18, 25, 26]). Some use the gossiping/epidemic paradigm (e.g. [16, 17, 18]) which is similar to our OP2P approach. 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 expected benefit to express the utility of data (loosely speaking, to measure the consistency), and we provide incentive for moving objects to participate as information suppliers and intermediaries.

Opportunistic dissemination of spatio-temporal information has been studied in [19, 28]. This paper differs from these works in multiple aspects. The theoretical analysis of information benefit is new. The consumer-paid pricing scheme is new, and so is the proposed approach to security and atomicity issues.

Incentive mechanisms for P2P and MANET. Our economic model, including virtual currency, security module, and consumer-paid policy, is inspired by the work of Buttyan and Hubaux [5] on stimulating packet forwarding in MANET. In their work, a node receives one unit of virtual currency for forwarding a message of another node, and such virtual currency units are deducted from the sender (or the destination). In our model, however, the amount of virtual currency charged by an intermediary node (broker) for forwarding a report is proportional to the expected benefit of the report, the latter depending on the dynamic spatio-temporal properties of the report (age and distance) as well as various system environmental parameters.

To the best of our knowledge, our work is the first one that attempts to quantify the benefit of spatio-temporal information and to price based on the benefit of information to the consumer rather than the cost of forwarding it. This distinguishes our work from many other incentive mechanisms (see e.g. [9, 10, 13, 15, 27]) which concentrate on compensating forwarding cost in terms of battery power, memory, CPU cycles. In a vehicular network such a cost is negligible.

Another possibility to build incentive is to use a reputation system (see e.g. [8]). 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. Moreover, management of reputation in a mobile distributed environment is difficult to implement.

Finally, game theory models have been proposed for incentives in peer-to-peer networks, including strategy-proof computing [3] and DAMD (Distributed Algorithmic Mechanism Design) [11, 24]. These models are aimed at providing incentive compatible mechanisms (i.e. mechanisms that result in desirable system-wide outcome from selfish behavior by the system′s agents) at tractable computational and communication cost, without the need for a centralized control. However, many mechanisms based on these models assume perfect connectivity among nodes [14], which may rarely exist in our environment. Applying the strategy-proof principle to a highly mobile environment is an open problem.

7. Conclusion

In this paper we devised an architecture and a data model for dissemination of spatio-temporal resource-information in mobile peer-to-peer networks, in which the resource information database is distributed among the hotspots and moving objects. We developed an analytical model for computing the expected time reduction in capturing a resource, given the age and the distance of the reported resource. Then we used this model to analyze the impact of various factors on the benefit of a report. We devised an economic model to stimulate moving objects to provide or relay resource information in mobile peer-to-peer networks. We also proposed an approach to determining the virtual currency price of the resource information. The price is based on the amount of time reduction that the information is expected to generate.

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.

Acknowledgements: We thank Aris Ouksel for helpful discussions.

References

[1] T. Michael. *Machine Learning*. McGraw-Hill, 1997.

[2] K. Yoon and C. Hwang. *Multiple Attribute Decision Making: An Introduction. Sage Publications*, 1995.

[3] Chaki Ng, et al. Strategyproof Computing: Systems Infrastructures for Self-Interested Parties. *1st Workshop on the Economics of P2P systems*, 2003.

[4] http://firechief.com/ar/firefighting_roborescuers_increase_disaster/

[5] L. Buttyan and J.P. Hubaux. Stimulating Cooperation in Self-Organizing Mobile Ad Hoc Networks. *ACM/Kluwer Mobile Networks and Applications*, 8(5), Oct. 2003.

[6] A. Pfitzmann, B. Pfitzmann, and M. Waidner. Trusting Mobile User Devices and Security Modules. *IEEE Computer*, February 1997.

[7] A. Markowetz, et al. Exploiting the Internet As a Geospatial Database, *International Workshop on Next Generation Geospatial Information*, 2003.

[8] 1st Workshop on the Economics of P2P Systems. http://www.sims.berkeley.edu/research/conferences/p2pecon

[9] S. Zhong, et al. Sprite: A Simple, Cheat-Proof, Credit-Based System for Mobile Ad-Hoc Networks. In *INFOCOM* 2003.

[10] R. Krishnan, et al. The *economics of peer-to-peer networks*, Carnegie Mellon University, 2002.

[11] J. Shneidman and D. Parkes. Rationality and self-interest in peer-to-peer networks. *Second International Workshop on Peer-to-Peer Systems*, 2003.

[12] W. Diffie, P. Oorschot, M. Wiener. Authentication and Authenticated Key Exchanges. *Designs, Codes and Cryptography* 2 (1992), 107-125.

[13] J. Crowcroft, et al. Modelling incentives for collaboration in mobile ad hoc networks. *WiOpt*'03, France, Mar. 2003.

[14] J. Liu, V. Issarny, On Service Allocation in Selfish Mobile Ad Hoc Networks. *PIM 2004*.

[15] V. Srinivasan, et al. Cooperation in wireless ad hoc networks. *INFOCOM* 2003.

[16] K. Rothermel, C. Becker, and J. Hahner. Consistent Update Diffusion in Mobile Ad Hoc Networks. *TR 2002/04, CS Dept*, University of Stuttgart, 2002.

[17] Vahdat and D. Becker. Epidemic Routing for Partially Connected Ad Hoc Networks, *TR CS-200006*, Duke University, April 2000.

[18] H. Hayashi, T. Hara, and S. Nishio. Cache Invalidation for Updated Data in Ad Hoc Net-works. *CoopIS* 2003.

[19] B. Xu, O. Wolfson. Pricing in an Economic Model for Mobile Peer-to-Peer Networks. *Technical Report DBMC-2004-05*, CS Dept, UIC, 2004.

[20] Personal Miles. http://www.ualmiles.com/PersonalMiles.jsp

[21] Gnutella website. http://gnutella.wego.com

[22] I. Stoica, et al. Chord: A Scalable Peer-to-Peer Lookup Service for Internet Applications. In *SIGCOMM*, 2001.

[23] Karl Aberer, et al. Improving Data Access in P2P Systems, *IEEE Internet Computing*, 6(1), 2002.

[24] J. Feigenbaum, S. Shenker. Distributed Algorithmic Mechanism Design: Recent Results and Future Directions. *Proc. of DIALM* 2002.

[25] M. Papadopouli et al. Effects of Power Conservation, Wireless Coverage and Cooperation on Data Dissemination Among Mobile Devices. *MobiHoc* 2001.

[26] G. Karumanchi, S. Muralidharan, and R. Prakash. Information Dissemination in Partitionable Mobile Ad Hoc Networks. *SRDS'99*, pp.4-13, 1999.

[27] A. Mok, et al. FAIR: Fee Arbitrated Incentive Architecture in Wireless Ad Hoc Networks. *RTAS'04*.

[28] B. Xu, A. Ouksel, O. Wolfson. Opportunistic Resource Exchange in Inter-vehicle Ad-hoc Networks. *MDM*04.