Opportunistic Resource Exchange in Inter-vehicle Ad-hoc Networks*

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

Abstract

In this paper we examine resource discovery in inter-vehicle ad-hoc networks in an urban area, where moving vehicles communicate with each other via short-range wireless transmission. Our focus is on real-time location-specific information. We explore an opportunistic *approach to resource discovery, in which a vehicle obtains information about resources from encountered vehicles. The vehicle uses a spatio-temporal relevance function to sort the resources, and save only the most relevant ones. Our theoretical and experimental analysis indicates that the opportunistic exchange algorithm automatically limits the distribution of a resource to a bounded spatial area and to the duration for which the resource is of interest.*

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 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 Potential drawbacks of this approach are (i) the responses to queries may be outdated, (ii) the response time may not meet the real-time requirements, and (iii) access to infrastructure communication service is costly. In this paper we explore a new paradigm that is based on inter-vehicle communications.

We assume that two vehicles can communicate with each other when their distance is smaller than a threshold. This communication can be enabled by a local area wireless protocol such as IEEE 802.11 [13], Ultra Wide Band (UWB) [16], Bluetooth [14], or CALM [15]. These protocols provide broadband (typically tens of Mbps) but short-range (typically 50-100 meters) peer-to-peer communication. With inter-vehicle communication, a mobile user discovers the desired information from the vehicles it encounters, or from distant vehicles by multi-hop transmission relayed by intermediate moving vehicles. Thus, resource discovery is performed in an *inter-vehicle ad hoc network*. Compared to the traditional fixed-station based information query, this paradigm has the following advantages:

1. 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.
2. It is free of charge, assuming that vehicles are willing to relay messages for free (in exchange for their messages being relayed).

In this paper we propose an *opportunistic* approach to resource discovery, in which a vehicle either senses the resources or obtains new resources from its exchanges with encountered vehicles. For example, a

* Research supported by NSF Grants ITR-0086144, NSF-0209190, ITR 0326284, NSF-0330342, NSF-0320956, CCR-0070738, and EIA-0000516.

vehicle V either senses the location of the parking space V vacated, or it finds out about available parking spaces from other vehicles. In turn, either these spaces have been vacated by these encountered vehicles, or these vehicles have obtained this information from other, previously encountered, vehicles. Thus the parking space information transitively spreads out across vehicles. This simple flooding procedure raises an important issue: what constraints on the resource spreading behavior are caused by limited memory space in each vehicle. With a limited memory space, a vehicle can only keep the most "relevant" resources it obtains. This in turn raises the question on what should be an appropriate relevance function in this environment? What parameters must be taken in consideration in this function? In the parking example, the duration since the parking slot became available and the distance of this slot from the vehicle ought to play a significant role in determining relevance. As both duration and distance increase, the likelihood is very low that the parking slot will remain available. The relevance function must therefore reflect these intuitive considerations. In this paper, we will introduce a spatio-temporal relevance function to compute the relevance of resources on the basis of both time and distance.

Our theoretical and experimental analysis reveals some interesting and useful properties of the opportunistic exchange approach based on the relevance function. First, a resource only spreads within a limited geographic area beyond the home of the resource (e.g., the location of a parking space). Second, within this limited area, the replication of a resource increases as a vehicle approaches the home of the resource. Furthermore, the boundary of this limited area varies with time. It initially expands until a time threshold beyond which there is no copy of the resource in the system. These properties show that our opportunistic exchange algorithm automatically limits the distribution of a resource to a bounded spatial area and to the duration for which the resource is of interest. Thus, for example, a vehicle in Chicago will never hear about a free parking slot in New York.

The rest of the paper is organized as follows. In section 2 we introduce the resource model. In section 3 we describe the opportunistic exchange algorithm and analyze it theoretically. In section 4 we analyze it by simulations. In section 5 we discuss relevant work. Finally in section 6 we conclude the paper and discuss future work.

2. Resource Model

In this section we introduce several notions regarding resources and the relevance of resources.

2.1. Resources and Their Organization

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. Information about a free parking space is a spatio-temporal resource.

A *spatio-temporal resource,* or a *resource* for short, is a piece of information about a spatio-temporal event e.g., the availability of a parking space at a certain location at a certain time, or the vehicle speed at a particular time and location. The *home* of the resource is the point location of the event. For example, the home of an available parking space is the location of the parking space. The *age* of a resource is the length of time since the resource has been created. For example, consider the resource indicating that a parking slot has become available at 2pm. The age of the resource at time 2:02pm is 2 minutes. Similarly, consider the resource indicating that the speed of a vehicle at time 2pm at location x is v. The age of the resource at time 2:05 is 5 minutes.

We assume that a moving object has a fixed amount of memory allocated to each application (e.g. the user allocates 10 entries for relevant parking slots. In other words, the user wants only 10 parking slots to be saved and displayed[1]). In this paper, we will only investigate the behavior of the resource propagation in the case of one application. The number of entries allocated for the application is referred to as the *memory allocation*.

2.2. Relevance of Resources

The relevance of spatial resources decreases as distance increases, and similarly, the relevance of temporal resources decreases with age. The relevance of spatio-temporal resources decreases as distance or time increase. For the parking slot example, 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

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

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 increases, and the distance from its home increases. In this paper we use the following function to compute the relevance of resource *R*:

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

t is the age of *R* and *d* is the distance from the home of *R*. α and β are non-negative constants that represent the decay factors of time and distance respectively. The bigger the ratio α/β, the more the relevance is sensitive to time than to distance; conversely, the relevance is more sensitive to distance than to time.

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

3. Opportunistic Exchange

In subsection 3.1 below we describe the opportunistic exchange procedure and follow it up in subsection 3.2 with a theoretical analysis.

3.1. Procedure Description

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, *A* and *B* first exchange their resources. Upon receiving new resources, vehicle *A* computes the relevance for each received resource and re-evaluates the relevance of its own resources. If all the resources do not fit in *A*'s memory space, the least relevant ones are purged. If two moving objects travel within the transmission range for a period of time, after the initial exchange only newly arrived resources are exchanged.

We assume that 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.

3.2. Spatial and Temporal Boundaries of Resource Distribution

In this subsection, we theoretically analyze the opportunistic exchange procedure, and show that with the above procedure, a resource is always propagated within a bounded area, and there exists an age threshold beyond which the resource disappears from the system. In the following analysis we assume that the wireless transmission range *r* is negligible and the time consumed by each resource exchange is negligible.

Denote by *M* the memory allocation, and by *v* the maximum speed a vehicle can travel with. We say that a resource is *rejected* by *O* at time *t* if *O* receives but does not save the resource at *t*. A resource is *new* for a time interval $[t_1,t_2]$ if the resource is created during this time interval.

Theorem 1: If $r = 0$ and the time consumed by each resource exchange is 0, and each object receives or generates at least *K* new resources per time unit, then for each resource *R*,

(1) If $\alpha > 0$, then there is no copy of *R* in the system after $\frac{M}{K} \cdot (1 + \frac{\beta}{\alpha} \cdot v)$ time units since the creation of *R*.

(2) At any point in time, there is no copy of *R* at any location that is more than $v \cdot \frac{M}{K}$ distance units away from the home of *R*.

Theorem 1 indicates that if there are sufficient interactions between vehicles in the system, and if new resources are generated at a certain pace, then the spread of each resource is limited to a neighborhood around its home, and the resource disappears from the system after a bounded amount of time.

Example: If the memory allocation is 10, the maximum speed is 60 miles/hour, $\beta/\alpha = 0.01$, and each object receives or generates at least 100 new resources per hour, then the age threshold is 9.6 minutes and the spatial boundary is 6 miles.

Before proving Theorem 1, let us introduce the following three lemmas.

Lemma 1: Let *R* be a resource created at time t_0. If *R* is rejected or purged by an object *O* at time *t* ($t \geq t_0$), then at any time point *t'* after *t*, the relevance of the least relevant resource in *O*'s memory is higher than or equal to $-(\alpha + \beta \cdot v) \cdot (t' - t_0)$.

Proof: Consider the relevance of *R* to *O* at time *t*. Since *R* is created at t_0 and the transmission range is 0, the distance between the location of *O* at *t* and the home of *R* cannot exceed $v \cdot (t - t_0)$. Thus the relevance of *R* for *O* at *t* cannot be lower than

$-\alpha \cdot (t - t_0) - \beta \cdot v \cdot (t - t_0)$. Since R is rejected or purged by O at t, the relevance of the least relevant resource in O's memory at t cannot be lower than $-\alpha \cdot (t - t_0) - \beta \cdot v \cdot (t - t_0)$. From time t to t', the maximum distance O can move is $v \cdot (t' - t)$. Thus the decrease of the relevance of the least relevant resource in O's memory from t to t' cannot exceed $\alpha \cdot (t' - t) + \beta \cdot v \cdot (t' - t)$. Therefore, the relevance of the least relevant resource in O's memory at t cannot be lower than $-\alpha \cdot (t - t_0) - \beta \cdot v \cdot (t - t_0) - (\alpha \cdot (t' - t) + \beta \cdot v \cdot (t' - t))$ $= -(\alpha + \beta \cdot v) \cdot (t' - t_0)$.

Lemma 2: Let R be a resource created at time t_0. If R is received by an object O at time t ($t \geq t_0$), then at any time point t' after t, the relevance of R for O at t' is not lower than $-(\alpha + \beta \cdot v) \cdot (t' - t_0)$.

Proof: Consider the relevance of R to O at time t. Since R is created at t_0 and the transmission range is 0, the distance between the location of O at t and the home of R cannot exceed $v \cdot (t - t_0)$. Thus the relevance of R for O at t cannot be lower than $-\alpha \cdot (t - t_0) - \beta \cdot v \cdot (t - t_0)$. From time t to t', the maximum distance O can move is $v \cdot (t' - t)$. Thus the decrease of the relevance of R from t to t' cannot exceed $\alpha \cdot (t' - t) + \beta \cdot v \cdot (t' - t)$. Therefore, the relevance of R for O at t cannot be lower than $-\alpha \cdot (t - t_0) - \beta \cdot v \cdot (t - t_0) - (\alpha \cdot (t' - t) + \beta \cdot v \cdot (t' - t))$ $= -(\alpha + \beta \cdot v) \cdot (t' - t_0)$.

Lemma 3: At any point in time, the relevance of any resource in any object's memory is higher than or equal to $-\frac{M}{K} \cdot (\alpha + \beta \cdot v)$.

Proof: Let t be an arbitrary time point. Let P be the set of the new resources that an object O has received or generated during the time interval $[t - \frac{M}{K}, t]$ (namely the last $\frac{M}{K}$ time units before t). Since O receives or generates at least K new resources per time unit, the size of P is at least M. Let Q be the set of resources in O's memory at t. There are two cases.

1. $Q \subseteq P$. Let R be a resource in Q. Suppose that R was created at t_0. According to Lemma 2, the relevance of R for O at t is no less than $-(\alpha + \beta \cdot v) \cdot (t - t_0)$ which is no less than $-\frac{M}{K} \cdot (\alpha + \beta \cdot v)$.
2. There is at least one resource in Q that is not in P. Consider any resource R that is in P but not in Q. R must have been rejected or purged by O at sometime t' between $t - \frac{M}{K}$ and t. Suppose that R was created at t_0. According to Lemma 1, the relevance of the least relevant resource in O's memory at t is higher than or equal to $-(\alpha + \beta \cdot v) \cdot (t - t_0) \geq -\frac{M}{K} \cdot (\alpha + \beta \cdot v)$.

In either of the above two cases, the relevance of the least relevant resource in O's memory at t is higher than or equal to $-\frac{M}{K} \cdot (\alpha + \beta \cdot v)$. In other words, at any point in time, the relevance of any resource in any object's memory cannot be lower than $-\frac{M}{K} \cdot (\alpha + \beta \cdot v)$.

Now we prove Theorem 1.

Proof of Theorem 1: First let us prove that there is no copy of R in the system after $\frac{M}{K} \cdot (1 + \frac{\beta}{\alpha} \cdot v)$ time units since the creation of R. When the age of a resource is greater than $\frac{M}{K} \cdot (1 + \frac{\beta}{\alpha} \cdot v)$, its relevance is lower than $-\frac{M}{K} \cdot (\alpha + \beta \cdot v)$. According to Lemma 3, this resource cannot exist in any object's memory.

Now we prove that at any point in time, there is no copy of R at any location that is more than $v \cdot \frac{M}{K}$ distance units away from the home of R. Consider the relevance of R for an object O that is more than $v \cdot \frac{M}{K}$ away from the home of R. If O has R, the age of R is greater than $\frac{M}{K}$. Thus the relevance of R is lower than $-\alpha \cdot \frac{M}{K} - \beta \cdot v \cdot \frac{M}{K} = -\frac{M}{K} \cdot (\alpha + \beta \cdot v)$. This contradicts Lemma 3.

4. Simulation

In this section we first describe the simulation method, and then present the simulation results.

4.1. Simulation Method

We synthetically generated and moved objects within a 50mile×50mile square area. For each object *i*, we randomly chose two points within the square area, and assigned them as the start point and the destination point of *i* respectively. The path of *i* is the straight-line segment between the start point and the destination point. *i* moves along its path from the start point to the destination point at a constant speed. All the objects use the same constant speed. A resource is generated and carried by *i* at the time point when *i* starts to move, representing that for example, a parking slot is available because of the leaving of *i*. The home of the resource is the start point of *i*. The memory allocation is the same for all the objects.

There are three parameters for each simulation run, namely the memory allocation *M*, the transmission range *r*, the constant speed *v*, and the traffic density *d* (i.e. the number of objects per square mile). Each simulation run is executed as follows. At the beginning of the simulation run, 50×50×*d* objects are generated and they start to move at the same time (time 0). When the distance between two objects is smaller than *r*, they exchange their resources, re-evaluate the relevance, and purge less relevant resources if needed. Each exchange is finished instantaneously. When an object reaches its destination point, it is removed from the system, and a new object is generated and started (again with the start and destination points randomly chosen within the square area). A resource is also generated for the start of the new object. The length of each simulation run is 10 simulated hours.

During a simulation run, we trace the distribution of each resource *R* at each time unit during *R*'s lifetime (*R*'s lifetime is the time period from the time when *R* is generated up to the time when it disappears from the system). For the purpose of this, we generate 30 rings centered at the home of *R* and with the width of 0.1 mile, such that the outer radius of the *i*-th (1≤i≤30) ring is *i*×0.1 miles. Let *t* be the time of the *k*-th time unit of *R*'s lifetime. The distribution of *R* at *t* is described and recorded as follows. We compute the density of *R* within each ring *i*, by dividing the number of copies of *R* within *i* at *t* by the size of *i*. Thus we obtain a vector $(d_1^{Rk}, d_2^{Rk}, ..., d_{30}^{Rk})$, where d_i^{Rk} is the density of *R* within ring *i* at *t*. We call this vector the *density histogram* of *R* for the *k*-th time unit.

At the end of the simulation run, we average the density histograms of each *R* for the *k*-th time unit, and thus get a vector $(d_1^k, d_2^k, ..., d_{30}^k)$, where $d_i^k = \frac{\sum_R d_i^{Rk}}{N}$, and *N* is the total number of resources that have been generated. This vector is the *average density histogram* for the *k*-th time unit. From the average density histogram other two measures are derived. One is the number of copies at the *k*-th time unit, which equals to $\sum_{i=1}^{30}(d_i^k \cdot A_i)$ where A_i is the size of ring *i*. The other is the 95%-boundary, which gives the smallest distance such that 95% of the copies are within this distance from the home. The system parameters and their values are listed in Table 1 below:

Table 1: Simulation parameters and their values

Parameter	Symbol	Unit	Value
Decay factor of time	α		1
Decay factor of distance	β		1
Side length of the geographic area		mile	50
Memory allocation (given in the number of memory entries allocated)	*M*		25, 50, 75, 100, 125
Transmission range	*r*	meter	50, 100, 150, 200, 250
Traffic speed	*v*	miles/ hour	20, 40, 60, 80
Traffic density	*d*	objects/mile^2	25, 50, 100
Ring width in density histogram		mile	0.1
Duration of a simulation run		hour	10

4.2. Simulation Results

4.2.1. Propagation of a resource.

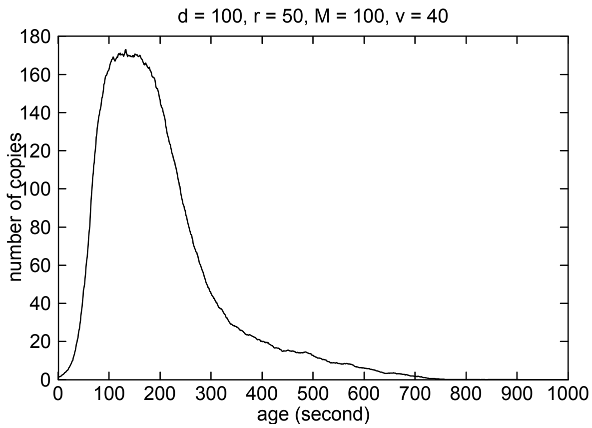

Figure 1: Number of copies as a function of age

First let us examine how the number of copies of a resource evolves during its lifetime. Figure 1 shows the number of copies as a function of age. As can be seen from Figure 1, the number of copies increases rapidly at beginning, until a maximum value is reached. Then the number decreases rapidly and then slowly until it becomes zero. The number increases at beginning because of the propagation of the resource caused by opportunistic exchanges. However, as time proceeds, the relevance decreases, causing two effects: (i) more objects purge the resource out; and (ii) fewer objects take it. These two effects make the number of copies start to decrease. After some time, the relevance becomes so low such that all the objects that have carried it have purged it out and no objects take it upon exchange. The resource thus disappears from the system.

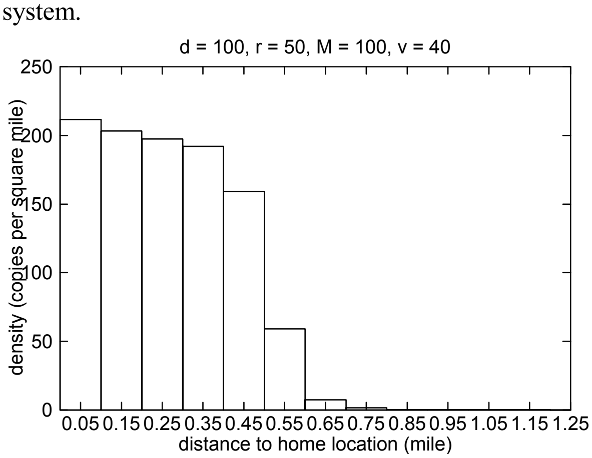

Figure 2: The average density histogram when the number of copies reaches the maximum value

Now we study how a resource is spatially distributed. We choose the time point at which the number of copies reaches the maximum value. Figure 2 shows the average density histogram at this time point. From Figure 2 it can be seen that the density decreases as the distance to the home increases. This indicates that the copies converge to the home. Observe that there is a boundary radius (in this case about 0.8 miles) such that there is no copy outside the area defined by the boundary radius. Further observe that this radius is very small compared to the size of the entire geographic area (50mile×50mile). This indicates that the opportunistic exchange algorithm automatically obtains a balance between the availability of resources and the cost of exchanging and storing them.

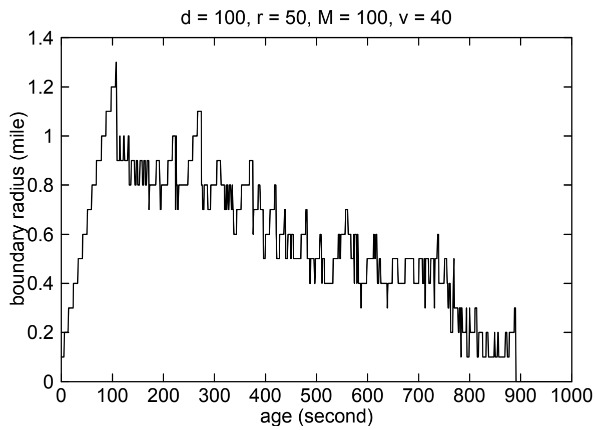

Figure 3: Boundary radius as a function of age

Figure 3 shows how the boundary radius changes as function of age. From Figure 3 we can see that the coverage of the resource (the area where the density is not zero) first expands, until a maximum value is reached. Then it starts to shrink until finally the resource disappears from the system. Observe the synchronization between Figure 3 and Figure 1. The coverage of the resource expands as the number of copies increases, and it shrinks as the number decreases.

Figure 3 has a lot of "noise": the boundary radius suddenly jumps up and goes back to "normal". One reason for this is as follows. If an object happens to travel a relatively long distance without interacting with any other object, then the boundary radius for a resource it carries will have a jump. In order to more precisely describe the distribution of a resource, we use 95%-boundary. Figure 4 shows the 95%-boundary as a function of age. It has a similar shape as Figure 3, but is much smoother.

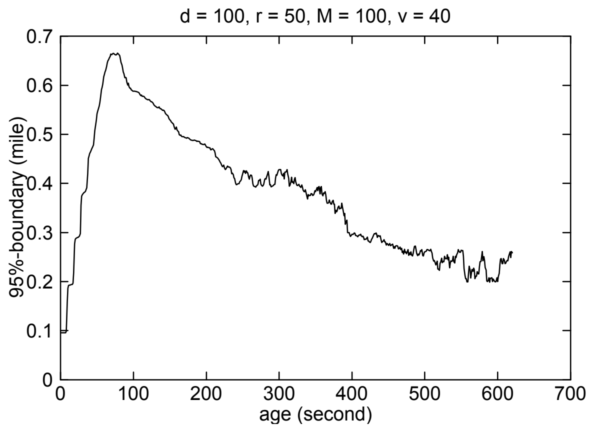

Figure 4: 95%-boundary as a function of age

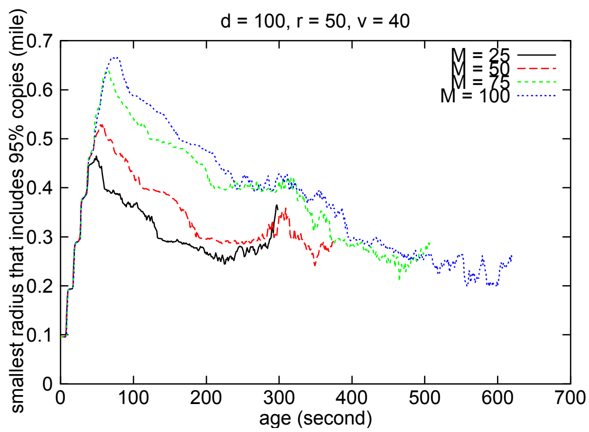

Figure 5: 95%-boundary as a function of age for different values of memory allocation

4.2.2. Impact of the memory allocation. Figure 5 shows the 95%-boundary curves for different values of memory allocation. It can be seen from Figure 5 that, (i) with higher memory allocation, the coverage of a resource expands to a higher maximum value and it expands to that value later; (ii) as the memory allocation increases, the length of the lifetime increases. These are because higher memory allocation accommodates more resources and thus delays the purge of resources.

4.2.3. Impact of the transmission range. Figure 6 shows the 95%-boundary curves for different sizes of transmission range. It can be seen from Figure 6 that, (i) with a bigger transmission range, the coverage of a resource expands to the maximum value sooner; (ii) as the size of transmission range increases, the length of lifetime decreases. This can be explained as follows. An increased transmission range results in a vehicle receiving a greater number of resources in a given time interval, thus causing greater contention for the memory, and causing entries to be bumped out of the memory sooner.

4.2.4. Impact of the traffic speed

Figure 7 shows the 95%-boundary curves for different values of traffic speed. It can be seen from Figure 7 that, as the traffic speed increases, the coverage of a resource expands to the maximum value sooner and the length of lifetime is shorter. Intuitively, a higher speed causes more frequent exchanges and thus speeds up the convergence.

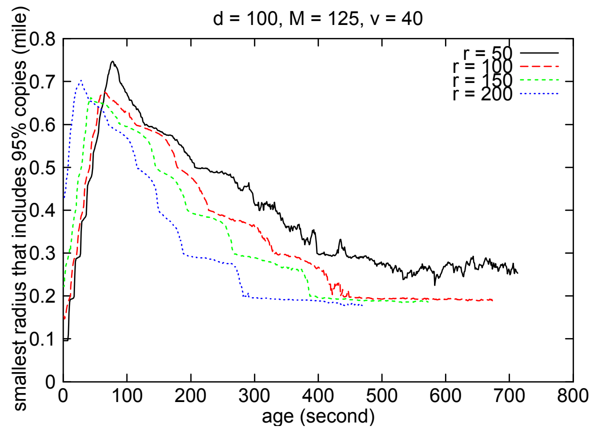

Figure 6: 95%-boundary as a function of age for different values of transmission range

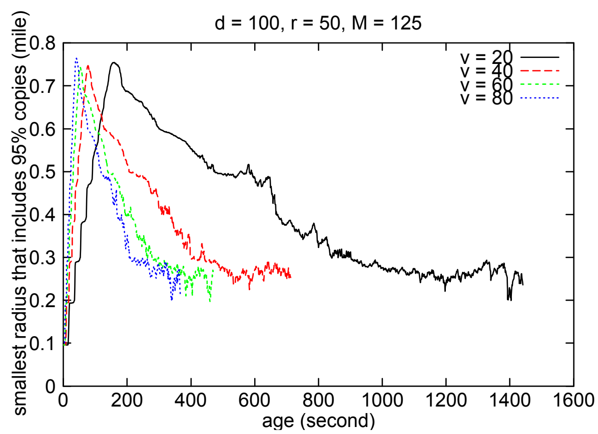

Figure 7: 95%-boundary as a function of age for different values of traffic speed

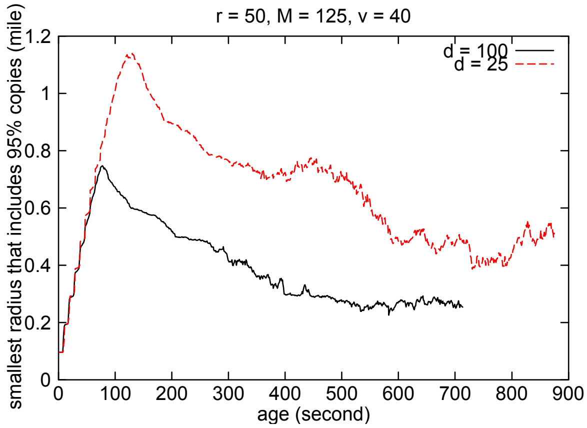

Figure 8: 95%-boundary as a function of age for different values of traffic density

4.2.5. Impact of the traffic density. Figure 8 shows the 95%-boundary curves for different values of traffic density. Observe that there is a significant difference between *d*=100 and *d*=25. With the density of 100, the coverage expands to a lower maximum value, but it expands to the value sooner. Intuitively, as the traffic density increases, a vehicle receives new resources more frequently, and therefore a resource is likely to be purged sooner.

5. Relevant Work

Different resource discovery architectures 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 allows resource providers to find a directory agent and to register with it, and a naming convention for resources. Examples are the Service Location Protocol (SLP) [1], Jini [2], Salutation [3], and UPnP [5]. In inter-vehicle ad-hoc networks, due to high variability of the network topology, we cannot rely on any one component to be always available. Therefore, it is important to develop approaches that rely more on opportunistic exchanges of resources than on a dedicated resource directory.

This paper reports on the preliminary steps in our on-going research to develop a comprehensive model that leads to a better understanding and prediction of the scale and the speed of resource spreading in a mobile network. An accurate model will provide insight into the viability of building a resource service infrastructure in a mobile network, and will aid in identifying the weaknesses in the resource spreading chain. For this purpose, we have been considering opportunistic resource spreading as a form of epidemics [4, 7, 9, 10, 11], where vehicles with resources are viewed as *infectious* individuals and vehicles accepting resources are similar to *susceptible* individuals. Starting from the classical epidemic Kermack-Mckendrick model [6], we investigate ways of adapting and expanding it to our more complex application.

For several reasons, the form of resource spreading in inter-vehicle ad-hoc network cannot be modeled using the simple SIS (Susceptible, Infectious, Susceptible) or the SIR (Susceptible, Infectious, Removed) models [6]. The propagation in these models concerns a single disease in a population where individuals are initially all susceptible to the disease, then possibly infected for some time, and later immune. In our case, we need to investigate the propagation of multiple diseases and an individual is infected depending on the characteristic of each disease and those infecting him at a given point in time. This characteristic in our case is expressed by the relevance function. An individual may not be susceptible to an infection at one time and become infected at a later time. Because the relevance of a resource for a moving vehicle is dynamic, a vehicle may refuse a resource at one time and then later find it of interest. Furthermore, an individual may become immune to a disease at one point and later become susceptible. A vehicle may shed away a resource at one time just to determine it to be relevant at another.

6. Conclusion and Future Work

In this paper we devised a model for discovery of spatio-temporal resources in an infrastructure-less environment, in which the database is distributed among the moving objects. In this model, two vehicles exchange their local databases when their distance is smaller than the wireless transmission range. Least relevant resources are purged during the exchange. We analyzed the model theoretically and experimentally, using a random waypoint model. The analysis suggests that the opportunistic exchange algorithm automatically restricts the propagation of a resource to a limited spatial area and a limited temporal interval. Furthermore, these spatial and temporal limits vary when system parameters change.

In this paper we assumed that the user only queries the local database. However, another possibility which we plan to investigate in future research, involves query delivery to all the vehicles in a geographic area, and the collection of results. There has been a lot of relevant research in mobile ad hoc networks (see [8] for a survey). The problem is how to deliver a message from a source to a destination via multiple mobile hops. We expect that some of this research will be relevant, however, usually the mobility assumed in ad hoc networks is different than in vehicular networks.

7. References

[1] E. Guttman, C. Perkins, J. Veizades, M. Day, Service Location Protocol, Version 2. *RFC2608*, June 1999. http://www.ietf.org/rfc/rfc2608.txt

[2] W. Keith Edwards, Core JINI, Prentice Hall, 1999.

[3] http://www.salutation.org/

[4] W. R. Heinzelman, J. Kulik, and H. Balakrishnan, Adaptive Protocols for Information Dissemination in Wireless Sensor Networks, in *Proceedings of the Fifth Annual ACM/IEEE International Conference on Mobile Computing and Networking (MOBICOM)*, pages 174--185, Seattle, Washington, 1999.

[5] Universal Plug-and-Play (UPnP) Forum, *Microsoft Corporation*. http://www.upnp.org

[6] N. Bailey. *The Mathematical Theory of Infectious Diseases and its Applications* (second edition), Hafner Press, 1975

[7] A. Beaufour, M. Leopold, P. BonnetSmart-Tag Based Data Dissemination, in *Proceedings of the First ACM International Workshop on Wireless Sensor Networks and Applications (WSNA)*, Atlanta, Georgia, 2002.

[8] E. Royer and C. Toh, A Review of Current Routing Protocols for Ad Hoc Mobile Wireless Networks, *IEEE Personal Communications*, pages 46--55, Apr. 1999.

[9] A. Khelil, C. Becker, J. Tian, K. Rothermel. An Epidemic Model for Information Diffusion in MANETs, MSWiM'02, pages 54-60, Atlanta, Georgia, September, 2002.

[10] A. Demers, D. Greene, C. Hauser, et al. Epidemic Algorithms for Replicated Database Maintenance, in *Proceedings of the 6th Annual ACM Symp. On Principles of Distributed Computing*, pages 1-12, Vancouver, BC, August 1987.

[11] Vahdat and D. Becker. Epidemic Routing for Partially Connected Ad Hoc Networks, *Technical Report CS-200006*, Duke University, April 2000. http://citeseer.nj.nec.com/vahdat00epidemic.html

[12] R. Vaughan. Urban Spatial Traffic Patterns. Pion, 1987.

[13] IEEE Computer Society. Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications.1997.

[14] 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.

[15] Continuous air interface for long and medium distance, http://www.magcom.no/Tolling/chicago/KEv%20Finale%20paper%202154%20CALM%20-%20The%20new%20ITS%20Communication%20concept%20r21.pdf

[16] New public safety applications and broadband internet access among uses envisioned by FCC consideration of ultra-wideband technology. Federal Communications Commission announcement, May 10, 2000. http://www.fcc.gov/Bureaus/Engineering_Technology/News_Releases/2000/nret0006.html.