Suitability of Distributed Mobile Wireless Networking
For Urban Traffic Congestion Mitigation

by

Shailesh Niranjan Humbad

B.S. in Computer Science and Mathematics
University of Michigan
Ann Arbor, Michigan (1999)

Submitted to the Department of Urban Studies and Planning
in partial fulfillment of the requirements for the degree of

MASTER IN CITY PLANNING
AT THE
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2001

© 2001 Shailesh Niranjan Humbad
All rights reserved

The author hereby grants to MIT permission to reproduce and to distribute publicly
paper and electronic copies of this thesis document in whole or in part.

Signature of Author_____
Department of Urban Studies and Planning
May 17, 2001

Certified by_____
Joseph Ferreira, Jr.
Professor of Urban Planning and Operations Research
Thesis Supervisor

Accepted by_____
Professor Dennis Frenchman
Chair, MCP Committee
Department of Urban Studies and Planning

ACKNOWLEDGMENTS

Suitability of Distributed Mobile Wireless Networking
For Urban Traffic Congestion Mitigation

by

Shailesh Niranjan Humbad

Submitted to the Department of Urban Studies and Planning
on May 17, 2001 in Partial Fulfillment of the
Requirements for the Degree of Master In City Planning

ABSTRACT

A suitability study is performed into the use of distributed mobile wireless networking for the purposes of urban traffic congestion mitigation. The technologies of global positioning system (GPS), wireless networking, and mobile ad-hoc networking (MANET) protocols are surveyed for potential usability and applicability in a peer-to-peer highway vehicle network. Analysis of traffic statistics for the Boston, MA metropolitan area reveal the parameters required to build an initial network. The estimated parameters are a two percent level of penetration (50,000 vehicles), two Megabit per second usable data bandwidth, one half mile average transmission range, two hundred dollars cost per device, and fifteen million dollar total system cost for five years of operation. Using a hop-count routing algorithm, the network would support collection of area-wide vehicle positions for automated highway traffic sampling and fleet tracking on congested roadways.

Following this first stage system are presented two more application scenarios according to increasing levels of penetration and increased reliability of the network. The medium-term application is the provision of mobile Internet access to allow consumer and business services. The long-term application is the ability to perform automated transactions. Envisioned in this long-term scenario is the ability to do area-wide road pricing to reduce congestion levels and influence land-use decisions.

Technology options and design choices for privacy protection are discussed including voluntary participation, incentivized participation, blackout zones, aggregation of data, non-identifiable data, and anonymous routing protocols. Centralized toll tables and transactions are shown to reduce privacy but increase convenience as opposed to distributed toll tables and in-vehicle transactions. Insitutional implemention through Federal ITS funding of a State-run public-private partnership is suggested to maximize mutual benefit. Given these options for handling the issues, the staging presented, and the flexibility, coverage, and application benefits of the system, the conclusion is that such a network would be suitable for mitigation of urban traffic congestion.

Thesis Supervisor: Joseph Ferreira, Jr.
Title: Professor of Urban Planning and Operations Research

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1: INTRODUCTION

1.1 Purpose and Scope of Study

Charging for use of public roadways on an as-you-go basis most equitably distributes the cost of roadways and can tailor demand to maximize efficiency. The latter possibility, either charging tolls for use of congested roads or giving credits for using underused roads, is called 'congestion pricing' and could eliminate the major problem of traffic congestion faced by urban areas. Despite its merits, congestion pricing has been unpopular in the U.S. due to its difficulty of implementation. This thesis presents a possible method for implementing a congestion pricing system from the perspective of a metropolitan transportation organization. By combining several recent technological innovations – the Internet, fast, short-range radio modems, and inexpensive global position system (GPS) access – densely spaced vehicles can communicate information between each other and to the Internet. With the ability to communicate comes the ability to collect traffic data and enact congestion pricing, as well as the provision of many other commercial services. An implementable scenario would first provide traffic data collection. Building on this initial proof-of-concept system would lead to Internet access for vehicles, which then would allow for the automated transactions necessary for congestion pricing.

1.2 Theoretical Benefits of Road Pricing

Automobile congestion is the excessive fullness of a road. Normally, cars will flow through a road at fixed speed and adding more cars to the road will increase the total volume of traffic that the road is carrying. However, after a certain point, adding more cars reduces the speed of traffic, which in turn reduces the volume of traffic that the road can carry. In the worst case, the road turns into a parking lot and the volume of traffic carried by the road goes down to zero. At this point, everyone on the road is wasting their time and wasting fuel idling.

Congestion increases pollution levels and also the risk of accidents. When such congestion happens on a daily basis due to overuse, it is termed peak level congestion. Another form of congestion is incident congestion, which is caused by non-recurring events such as accidents and construction. The problem at hand is concerned with peak level congestion, for which a solution seems feasible and prudent.

By imposing charges on road users, road pricing internalizes the external costs of congestion such as time delays, accidents, and pollution. For the purposes here, *road* pricing and *congestion* pricing are used interchangeably. Arthur Pigou, the Cambridge economist, argued for such pricing more than 75 years ago (Button, 1998, 'Introduction'). His explanation was simple, and is illustrated by Figure 1.1. On a hypothetical one-lane road, identical vehicles have a willingness to pay for mobility indicated by the downward sloping line D, which is equivalent to Marginal Private Benefit (MPB) and Marginal Social Benefits (MSB). This demand curve reflects the fact that as price goes up, less mobility is demanded. Because each vehicle causes congestion whose burden is felt by all road users, the Marginal Social Cost (MSC) is greater than the Marginal Private Cost (MPC) or Average Social Cost (ASC). Without road pricing, the system equilibrates at the intersection of D and MPC, and a mobility of N^0 demanded. At this level of mobility, congestion exacts an actual cost on society, represented by the MSC at N^0, which is much greater than the cost to each individual user, the MPC. Congestion causes time delays felt by all users. To remedy the situation, a toll or use price of r^* is charged, bringing the equilibrium mobility to N^* and the MSC in line with the MPB. Assuming redistribution of the collected revenue, social welfare improves by the amount represented by the shaded area (Button, 1998, 'Introduction').
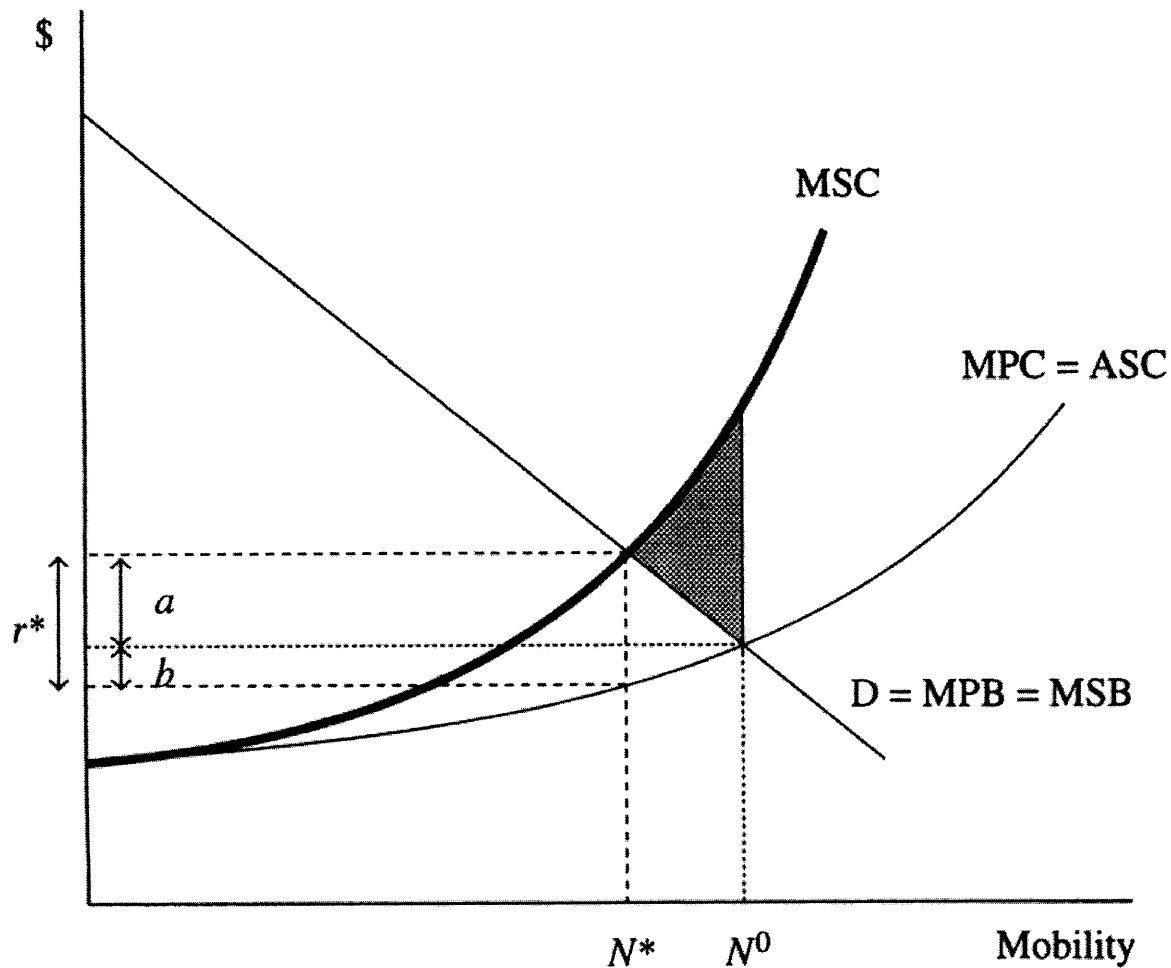
Figure 1.1: Efficiency Gains from Congestion Pricing (Button, 1998, 'Introduction')

If the revenue is not redistributed, then the people who no longer use the road because of

the toll – the 'tolled off' – are definitely worse off, and the people who remain users of the road –

the "tolled on" – are either just as well off or worse off. The benefactor is the government that

collects the revenue (Hau, 1998, p. 48). Therefore, there can be no political impetus to enact

road tolls unless the revenue is used for a publicly supported cause (such as road building), or is

redistributed to the public directly. Indeed, congestion pricing has only succeeded in places

where the public felt a dire need for it. Normally, people would like to travel on the roads for

free, but because society pays all costs of roadways, it is not surprising that the French playwright, Jean Anouilh, once said, "What you get for free, costs too much."

The late Nobel Prize winning economist William Vickrey (1969) argued that as investment in roadways and other transport facilities proceeds over time, a larger and larger proportion of it is devoted primarily to relieve congestion and expand capacity. To explain the rationale for this, Vickrey created the dynamic bottleneck model of congestion. Under his simple bottleneck scenario, with no tolls, all vehicles will arrive at a bottleneck at the same time and cause congestion visible due to the queuing of the vehicles. Under a suitable toll structure, the queue is eliminated, and to the extent that the toll revenues reduce preexisting flat taxes, motorists as a whole are better off.

Vickrey (1969) states, "In the absence of congestion charges, a decision to expand facilities may have to be taken on an all or nothing basis." This is because construction of additional capacity is not continuous, but piecemeal. Combined with the difficulty of projecting demand, the situation commonly leads road builders to substantially overinvest or underinvest. The canonical example is of expanding capacity on a congested road only to later find that the additional traffic attracted has caused just as much congestion as before. On the other hand, pricing can be adjusted often and in small increments, upward or downward depending on needs. Pricing doesn't mean highway investment will cease; it provides a mechanism to know when it is justified.

If all routes were appropriately tolled, it would be relatively easy to see where the highest prices are needed to limit congestion, and these places would be the most likely candidates for expansion. Inherent in price is information, a basic concept of economics. In congestion pricing, there is inherent information about traffic demand. Thus, any sophisticated congestion

pricing system necessarily comes with a sophisticated traffic data collection system. This information is valuable by itself, and can eliminate the often arbitrary rules of thumb used by existing planning methods.

1.3 Traffic Data Collection Needs

"Data collection forms the basis for all studies on urban traffic management (OECD, 1979)." Collected data and its analyses are used during the evaluation stage of setting urban policy objectives. The general problem areas to which it can be applied are accessibility, mobility, safety, environment, and energy. In each of these problem areas, various strategies are applied to influence tripmaking, improve throughput, influence modal choice, improve traffic safety, alleviate environmental problems, and reduce energy consumption. These strategies represent the traffic engineering perspective and more strategies exist. As the need for urban traffic management increases, so does the need for traffic data collection of motorized vehicles. In summary, traffic data collection is important for public agencies for improving the management of the transportation infrastructure.

Increasingly, private entities are demanding traffic data in the form of vehicle location information. Private entities can use vehicle location information, combined with other traffic data, to provide many kinds of services. These generally fall under the heading of Advanced Traveler Information Systems (ATIS) and include services such as pre-trip information about traffic, in-trip driver information about traffic, personal information services, and route guidance and navigation (Ciccarelli, 2001, p. 9). At a more commercial level, the private sector uses fall under the categories of fleet tracking, mayday systems, and dynamic route guidance and information. Apogee Associates in May 1997 projected that these services would reach a market

size of $420 billion over the next twenty years (1997). As the quality and timeliness of data improves, the services become more valuable.

When location information from numerous vehicles is collected and processed, one can obtain detailed information about regional traffic. This traffic information consists of presence, volume (vehicles per hour), speed (kilometers per hour), and density (vehicles per hour). In order to do this, there must be enough 'probes', vehicles whose location can be tracked, to provide a representative sample. One viable option has emerged recently to tackle this option, that of tracking the location of wireless cell phones in standby mode. The ability to track the phones comes as a result of a Federal Communications Commission (FCC) regulation in 1996 mandating that mobile phone companies be able to ascertain the location of callers in emergency situations. Work is currently underway in public-private partnerships to implement these systems to allow private companies to offer location-based services and public sector institutions to collect traffic data. While the public sector uses the traffic information to decide policy, the private sector uses the location information to provide services to vehicles.

1.4 Emergence and Convergence of Technologies

Tracking of cell phones represents one technological means of obtaining vehicle position information, but is by no means the only one. Cell phones are based on radio communication, which has recently become more and more sophisticated. As part of the digital revolution, it has become possible to transmit data, in binary form, over the electromagnetic medium of radio waves reliably at high rates of speed. The transmission works best when the transmitter of the signal and receiver are close to each other, with few obstacles, such as walls, buildings, people, and trees in between them. Cell phones thus require many cellular towers that act as local carriers of the signal within a cell. The drawback of this system is the huge capital costs of

building, maintaining, and upgrading hundreds of thousands of towers over wide geographic areas. The fact that wireless radio communication works best over short distances, requiring many towers, begs the question of whether or not a cellular system is the best means of deploying wireless networking services.

If not cellular, then what else? Before answering that question, it is important to notice another technology that has recently blossomed – global positioning system (GPS). With a GPS receiver, one can pinpoint one's location in latitude and longitude anywhere on the Earth. In addition, one can obtain the precise time. GPS was once limited to military use only, and accuracy was purposely limited for civilian use. However, since May 2000, the induced error has been eliminated and civilian GPS receivers have become much more accurate. Due to technological advancements, GPS receivers have come down drastically in cost, size, and power consumption, to the point where handheld versions are readily available. At the same time, their accuracy has improved due to better error correction using advanced methodology, as well as from the government lifting restrictions. Knowledge of location and the ability to transmit data short distances converge in the technology of the mobile ad-hoc network (MANET).

In a MANET, vehicles communicate with each other over short distances using radio modems. For communication between parties not in radio range, information 'hops' from one vehicle to the next, with each intermediate node passing along information for other nodes. The nature of this network implies several things. First of all, there must be sufficient density of nodes in order for information to be successfully relayed. Second, due to the ad-hoc configuration of the network, there is no need for costly, fixed base stations or towers. Finally, a routing protocol is necessary to making this work. A routing protocol is a set of rules for relaying data that ensures the integrity and success of senders' and receivers' data transmissions

over multiple hops. When combined with GPS location knowledge, the routing protocol becomes robust, efficient, and practical for large numbers of nodes.

The nodes do not necessarily have to be mobile or vehicular; they could be people carrying handhelds, or even rooftop transmitters. However, vehicles have some characteristics that make them ideal for large-scale networks of this type. First of all, vehicles do not have the power constraints faced by battery operated electronic devices. Whereas a battery might be limited to 0.25 watts, a vehicle can easily provide 100 watts or more on a continuous basis, enough to run a small computer. Second, the movement patterns of vehicles are predictable and their routes are known; they simply follow the road networks. Third, as in-vehicle telematics and electronics become ever more popular and commonplace, the process of adding another device is well known and understood. In fact, it appears that such devices will end up in vehicles regardless of any government action. Finally, people are spending ever more time in their vehicles, often in congested routes. By implementing a MANET that exploits the inherent characteristics of vehicle congestion, one can implement solutions to combat congestion.

1.5 Structure of Thesis

This first chapter covers the basic elements and points of view of this thesis to motivate the remaining chapters. What follows in Chapter Two is a discussion in further detail of the limitations of current congestion pricing methods, traffic data collection methods, and highway finance. Congestion pricing systems in Singapore and in LA (FASTRAK) will be used as examples. Chapter Three will cover the basics of networking technology: wireless networking, GPS, and multi-hop protocols for mobile networks. Chapter Four will present possible scenarios for implementing an actual system, including an analysis of feasibility and suitability. A phased build-out will be discussed, showing the applications enabled at each level of implementation.

Chapter Five will conclude with a discussion of general issues surrounding congestion pricing, the prospects and hurdles of actual implementation, and a summary of results.

CHAPTER 2: SURVEY OF CURRENT METHODS

2.1 U.S. Highway Finance System

The Federal-Aid Highway Act of 1956 launched the construction of *National System of Interstate and Defense Highways* (Weiner, 1992). Under this legislation, the federal government pays states a matching grant of 90% for interstate highway related projects. The Highway Revenue Act of 1956 provided the funding for this massive construction project by increasing the gasoline tax and several other motor related excises. In passing the law, Congress broke its tradition of generalizing tax receipts and earmarked the revenues into a *highway trust fund* whose funds could only be spent for highway purposes (Davis and Cunningham, 1994). The momentum of this policy has carried forward to this day, but has softened somewhat with the passage of the Intermodal Surface Transportation Efficiency Act (ISTEA) of 1991. States have more flexibility in deciding how funds are used, and the federal matching grant will decline to 75%. However the incentive to focus on highway projects persists because, "states simply will not forego the 75 percent financial support available for a commitment of 25 percent of their own resources (Davis and Cunningham, 1994, p. 135)." This huge incentive to construct highways has been a major factor in the U.S. becoming an auto-centric nation with over two hundred million vehicles.

The federal government raises more than 80% of its road funding from a national 18 cent per gallon gasoline tax (Michigan DOT Act 51 Study Committee, 2000). The problem with this system of funding highways is that the taxes do not directly apply to a driver's consumption of the space on a roadway. In a small town with ample room, an excise tax on motor fuels is simple and effective, but in a large metropolis, where roadway space is constrained, it leads to congestion. There are too many people fighting over a too little space because fees that are

collected are not directly related to the actual usage of the roads. For example, in Massachusetts

most roadway improvements are financed through a state gasoline tax and the proceeds of

periodic bond issues. This money is used through federal aid programs that contribute up to 90%

financing for major improvements to state and local roadways and bridges (Boston Metropolitan

Planning Organization, 1989). Funding level is determined by the classification of the road, with

the most funding going for interstate highways. For example, the federal government is

contributing about 58% of the cost of the Big Dig in downtown Boston, or about seven billion

out of twelve billion dollars (Massachusetts Turnpike Authority, 2000). In this funding scheme,

there is no way to curtail overconsumption or in other words, control congestion.

2.2 Cost of Congestion

How much does congestion actually cost people in terms of dollars? In the Urban

Roadway Congestion Annual Report--1982 to 1994, the Texas Transportation Institute (TTI)

measured areawide congestion levels in 50 urban areas (1992). The data is based upon the

Federal Highway Administration's Highway Performance Monitoring System, which organizes

data collected by state and local agencies. The areawide congestion levels are measured by

calculating the density of traffic on roadways. The density formula they use incorporates the

volume of cars (vehicle kilometers of travel) on the roads divided by the amount of roads (lane-

kilometers of roadway). (A similar method will be used in Chapter 4.) A variable weighting

factor is used so that results are comparable across urban areas. The formula gives a *roadway*

congestion index (RCI) where a value greater than 1 indicates areawide congestion, and a value

less than 1 indicates an uncongested average mobility level. Note that even though areawide it

may be uncongested, there could still be places where there is heavy congestion (Texas

Transportation Institute, 1992).

The city with the worst congestion was Los Angeles, with an RCI of 1.52 in 1994. By comparison, Detroit was roughly in the middle of the top ten most congested with an RCI of 1.24. All of the cities in the top ten had increasing levels of congestion from the period of 1988 to 1994. The study also calculated an average speed ratio, which was the ratio between the observed speeds and the freeflow speeds. The top ten cities had an ASR of about 0.7, which means that traffic was flowing on average 70% of what it could be flowing if there were no congestion. From these measures, the study calculates the cost of congestion, in terms of wasted time and fuel. For each driver, the cost was roughly $900 per year in 1994 in the worst ten cities. This works out to approximately $3.50 per workday per driver, or about $75 dollars a month. For Boston, the estimated total congestion cost was about 1.65 billion dollars per year in 1994. For comparison, this represents approximately 10% of the entire Massachusetts State budget that year of about 17 billion dollars. In another study by the Environmental Defense Fund (EDF), Michael Cameron (1994) calculates the cost of congestion (including pollution) for Southern California, to be about 28% of all transportation costs in the state, or about 11 billion out of 38 billion dollars. Clearly, congestion is a serious and expensive issue facing urban areas in the United States.

2.3 Methods For Relieving Congestion

Due to the enormous scale of the problem, many people have proposed solutions. The first, most basic solution is to simply build more roads. The Texas Transportation Institute (TTI) calculated how much roadway would be required to maintain congestion levels, and shows that there is a shortfall every year in the amount of roads that are built and the amount of roads that should be built. The shortfall arises due to the fact that the number of vehicles and vehicle miles traveled are increasing much faster than the size of the physical infrastructure, which is growing

slowly if at all. Detroit was the worst offender in 1994 when it built only 50% of the roads that it technically should have to maintain congestion levels (TTI, 1992).

In addition to TTI's analysis and the theoretical justifications provided by Vickrey in the 1960's, anecdotal evidence shows that building more highways just attracts more traffic. For example, the Detroit News reports on the condition of Interstate 275 in suburban Detroit, a fairly new highway built in the 1970's. "Now, more than 100,000 cars travel parts of I-275 daily near Canton Township, the second fastest growing township in southeast Michigan. Cars often back up two miles on I-275 near the Ford Road exit during rush hour. Crowded Ford and Haggerty just off the exit are contributing to the interstate backups (Puls, 1999)." In another article, a Michigan Department of Transportation spokesman says, "We cannot build our way out of congestion (Albert, 2000)." This pattern repeats itself in cities all across the country.

Another set of solutions falls under the heading of Travel Demand Management (TDM). The primary purpose of TDM is to reduce the number of vehicles using the road system while providing a wide variety of mobility options to those who wish to travel. As the name implies, the end result of these changes in behavior is a reduction in travel demand. However, inducing these changes in behavior is not easy, and often requires monetary incentive. Therefore, these methods have had limited success and work best only with wide cooperation. In addition, these are only short-term solutions to congestion in the case of increasing population and fixed infrastructure. These methods include the following:

- carpools and vanpools

- public and private transit, including buspools and shuttles

- non-motorized travel, including bicycling and walking

- compressed work weeks

- flexible work schedules

- service improvements to transit service that provide savings in costs and travel time

- provision of preferential lanes

Intelligent Transportation Systems (ITS) are a recently promoted solution for congestion and improvement of transportation infrastructure. "ITS encompasses a broad range of innovative technologies, systems, and information management strategies that when linked together can greatly improve the safety and efficiency of travel and mobility in urban and rural areas, on transit systems, and on Interstate highways (Federal Highway Administration, 1998a)." These methods are useful and important, but only have limited effects. Their salient feature is that they make the current system incrementally more efficient. Thus, they do not tackle the issue of congestion directly (except for electronic toll collection in places where tolling did not previously exist). ITS includes the following:

- Traffic Signal Control

- Electronic Toll Collection

- Traveler Information Services

- Fleet Scheduling and Management

- Vehicle Identification and Record-Keeping Systems

Land-use planning can also be used to control congestion. Ostensibly, planners can zone land and then try to estimate and control the impact development will have on traffic levels. This is notoriously difficult, because first of all planners can not completely control development. Second, the predictions are based on estimates and assumptions about future traffic levels. Third, there is a lack of coordination between the many governmental units that impact land use and transportation. Fourth, planners exert local power over land use decisions, making regional

cooperation difficult. Fifth, many land-use policies with good local intentions, such as minimum

lot sizes and minimum parking provisions, actually encourage regional over-reliance on autos.

And finally, even absent the political obstacles, the simulation and prediction problem is

extremely difficult, mainly because gathering the data and running the simulations requires

extensive resources and expertise. Though this latter obstacle will likely diminish as computers

become more powerful and prolific, it will still present a problem in cities where resources are

severely constrained. The underlying problem is a lack of pricing information for roads that

works against policy makers. *Land-use planning is a frail and ad-hoc exercise when road prices

aren't known.* The following example from the perspective of a developer will help make this

clear.

Currently, developers have a silo mentality (understandably). They develop a parcel in

the manner of a computer scientist's greedy-algorithm, i.e. they optimize revenue for the

development locally, with little regard for the fabric of the city in which they are embedded. For

example, suppose a developer owns a land parcel downtown that is equidistant from an airport

and a convention center. In the absence of any other information, the developer may consider

the development to be worthwhile. Now, consider what would happen if the roads were priced.

Suppose the round-trip car journey from the parcel to the convention center cost $10 while the

route to the airport cost $5 in congestion-based tolls. In this case, the developer could calculate

the cost to his customers of the trips to the convention center and the airport. He or she may

decide it is not economical to build a hotel on his parcel because of the high cost of the trips to

the convention center and airport. Suppose there was a residential building next to the

convention center. The owner of that property, upon seeing the road prices, would suddenly

realize the value of having a hotel there and convert it, because the owner knows of the high

costs of travel to the convention center from other areas. Alternatively, the hotel developer may develop the hotel and provide regular bus service to the convention center.

In any case, the existence of road prices leads naturally and automatically to a more efficient layout of the city with reduced traffic loads. While it may be true that the developer knows that the route to the convention center is normally heavily congested, this may not be significant enough to be realistically considered by the developer. Multiply that oversight by hundreds of developers, and one can see easily how regional congestion can arise. The developer is forced into the silo mentality by the lack of pricing information for the roads. Having the roadway prices will make the developers job of locally maximizing revenue easier and it will make his or her decision consistent with the greater requirements of the city's infrastructure. Nothing enters into the mind more clearly than a dollar value price. Businesses and households can't make correct decisions about where to locate because the costs of using the roads are hidden. So, whenever a new road is built, everyone decides to move near it to take advantage of the lower congestion, but since there is no price signal to tell everyone that the road is becoming too congested, people don't know to stop development until the highway becomes congested.

Because road pricing is such a tantalizing option for controlling congestion, many studies have been done to ascertain its feasibility. Evaluating Congestion Pricing Alternatives for the Puget Sound Regional Council summarizes these options well (Puget Sound Regional Council, 1994). This document lists three functional points in the transportation decision process where pricing can be applied. These are Roadway Use, Vehicle Use, and Vehicle Ownership. Examples of Vehicle Ownership charges include purchase taxes, registration fees, and flat license fees. Of the three categories of pricing, these are least directly related to the decision to

make a particular trip on a particular road. Examples of Vehicle Use charges are vehicle miles traveled (VMT) charges, parking charges, and fuel taxes. These affect the decision to use a vehicle, but not the decision to use a particular roadway. Examples of Roadway Use charges are tollgates, automatic vehicle identification (AVI) systems, and area licensing. The prices set for these methods directly affect the decision to use a particular roadway's capacity. Thus, Roadway Use charges are theoretically a *best* solution because they directly reduce congestion. However, they are least used in practice due to significant implementation difficulties. The two real world examples of road pricing that follow will provide design guidelines and illustrate how implementation difficulties have been overcome.

2.4 Singapore Road Pricing System

Road pricing methods are most notably used in Singapore. The country is a densely populated island city-state in Asia that has severe constraints on available land. The country is also wealthy and generally well managed with strict laws and law-abiding citizens. In 1995, management and administration of the transport sector was combined into a single, amalgamated agency called the Land Transport Agency (LTA). Singapore, like all industrializing countries of Asia, has been faced with increasing numbers of vehicles, vehicle-miles traveled, and cars per person. Its problem has been more acute due to its severe land scarcity. Despite implementing high import duties, registration fees, gasoline taxes, and parking charges, all measures of vehicle usage have increased steadily and rapidly over the past 50 years. (In 1990, Singapore implemented a Vehicle Quota System (VQS) to directly control the number of vehicles allowed in the country. The VQS limits the increase in the number of vehicles each year by auctioning off a limited number of Certificates of Entitlement (COE), without which vehicles may not be purchased. The bidding process resulted in COE costs of approximately US$ 40,000 for mid-

sized cars in 1997 (Foo, 1998). However, since this is a complementary measure to congestion pricing and only indirectly affects peak-level traffic, it is not salient to this discussion.)

To combat increased congestion, Singapore implemented an Area Licensing Scheme (ALS) in 1975 as a crude and simple way of road pricing. As described by Tuan Seik Foo in 1997, under the ALS, vehicles were required to purchase a permit in order to obtain permission to enter designated restricted zones in the city during peak traffic periods. The ALS boundaries encompassed approximately 7.25 square kilometers of Singapore's Central Business District (CBD) and had 33 entry points demarcated by prominent, lighted signs. In 1997, the charges ranged from US$ 0.67 to US$ $4.02 for whole day passes (that allowed entrance during any designated peak period). (One Singaporean dollar is taken to be worth US$ 0.67 throughout this text, as it was in 1997. As of May 2001, it is worth about US$ 0.60.) Part-day and monthly passes were also available at entrance points and by prepurchase. Passes were sold through a wide variety of outlets including roadside booths, post offices, various retail stores, and gas stations. Enforcement was done at the entry points by trained traffic wardens through visual identification. The number of violators averaged below 1%, and vehicles were booked by mail and charged a fine of US$ 33.50. Easy escape routes were provided just before entry points so motorists would not unwittingly enter without a permit.

The ALS successfully reduced peak congestion levels, leading to fewer vehicles in the city center and higher average speeds. One study reported average travel speeds increasing from 23 km/hr to 30 km/hr (Seventh Parliament of Singapore, 1990). Traffic accident and air pollution levels were also shown to have decreased. Motorists rescheduled trips to the CBD during off-peak hours, changed routes to avoid the CBD altogether, or switched modes to public transportation. There was no major effect on the business sector in the CBD due to little change

in customer volume. Due to low enforcement and maintenance costs, the system generated significant amounts of revenue, which was used in the integrated transportation policy. The success of ALS was helped significantly by a number of characteristics of Singapore that are listed below.

- The severe land-scarcity in Singapore.

- Single-level government and no outside traffic.

- An integrated approach to transportation policy.

- Availability and promotion of alternative transport. (buses, rail, park-and-ride)

- Mostly middle-class vehicle owners. (less likely to have political influence to protest than the upper-class)

- Literate, well-informed populace with low corruption.

Criticism of the ALS has been largely unfounded. For example, businesses did not move out of the CBD *en masse*, because the reduced traffic congestion improved accessibility and lowered pollution, resulting in benefits outweighing the cost of the ALS. Another criticism is that ALS fees are regressive, affecting lower income groups more than higher income groups. However, the government channels funds directly to lower income groups by funding affordable mass transit that is extensively available. This funding includes subsidized capital outlays for transit, controlled fares, and ALS fee exemptions for mass transit vehicles. Due to various reasons, these funds are shown to come from general public funds, not directly from the ALS revenue stream, but the effect is the same – a transfer of wealth. It is worth repeating that the major assumption is the availability of mass transit as an alternative to driving. Because of this and the characteristics listed above, the ALS was able to overcome its largest obstacle, significant public opposition.

Other limitations of ALS are more pronounced. For example, congestion has increased on peripheral roads unabated. The government has begun to implement peak period tolls on those roads as well. Another limitation is the inflexibility of setting prices for the permits, which typically must remain static for several years before changing, even though traffic conditions vary at a much finer grain. In short, the ALS lacks flexibility of pricing to match the conditions of the roads. In order to overcome these limitations, Singapore transportation officials studied new methods of state-of-the-art electronic toll collection systems during the mid to late 1980s. These studies culminated in the LTA proposing an Electronic Road Pricing (ERP) scheme in its 1996 White Paper.

Tuan Seik Foo describes the ERP scheme in his year 2000 paper on demand management in Singapore. The ERP scheme supplanted the ALS in 1998 with the installation of overhead gantries at each restricted zone entrance. All vehicles wishing to enter the zone must be installed with an in-vehicle unit (IU). IUs were originally installed for free, but now must be purchased or leased by the vehicle owner. The IU is a small, passive, low-power, electronic device that allows automatic debiting of charges from a cash card upon passing under one of the overhead gantries. In such a way, vehicles are now electronically charged a toll whenever they enter into the designated areas. Cash cards store money and are sold throughout Singapore and can be refilled at ATMs. The overhead gantries record a photograph of any vehicle that does not pay due to insufficient funds, no cash card, no IU, or a malfunctioning IU. A bill for the fine is mailed to the violator, and the violator may appeal if there was no actual violation, i.e. there was a mechanical system error. The toll charges are weighted according to size of the vehicle, e.g. motorcycles are charged half the cost of passenger vehicles while large buses and very heavy goods vehicles are charged double. Charges are also varied by area, with different rates for the

CBD and the various expressways that are also equipped with gantries. See Table 2.1 below for rates in the CBD restricted zone. Charges vary at most every thirty minutes according to a predetermined schedule. These rates are about half the rate of charges under the ALS for full-day passes. Also, the charges are reviewed quarterly and adjusted if they are found to be too high or too low at a particular time, with the criteria for adjustment primarily being a pre-determined optimum travel speed range.

Table 2.1: ERP Rates in Restricted Zone as of July 1, 1999	
Time	Rate for Cars (In Singaporean Dollars)
7:30 – 8:00 AM	1.00
8:00 – 8:30 AM	2.00
8:30 – 9:00 AM	2.50
9:00 – 9:30 AM	2.00
9:30 AM – 5:30 PM	1.00
5:30 – 6:00 PM	1.50
6:00 – 6:30 PM	2.00
6:30 – 7:00 PM	1.00
(Source: Foo, 2000) (S$ 1 ~ US$ 0.60)	

The ERP system construction contract was awarded to the Philips consortium for US$ 132 million in 1995. By the time the system was made operational in July 1998, 96% of Singapore's 680,000 vehicles had been fitted with the devices at IU installation centers such as vehicle inspection centers, auto-repair shops, and mobile units. The ERP system collected 23% less revenue (US$ 53 million versus US$ 69 million) than the ALS scheme, but cost 47% less to maintain (US$ 6 million versus US$ 11.4 million) per year. However, *congestion mitigation is the goal of the system and not revenue collection.* The LTA announced optimum travel speed ranges of 45-65 km/hr on expressways and 20-30 km/hr on arterials in the restricted zone, and has successfully adjusted charges each quarter until those criteria are met. The system has been

more equitable than the ALS because roads are exactly pay-per-use. Despite the high initial investment costs, the ERP system has lower maintenance costs and much higher flexibility.

The ERP system has been very reliable too. The reliability is mostly due to a well-engineered, sturdy, and reliable IU and Gantry system that underwent significant testing before being deployed. The ERP system also requires no human enforcement personnel at the gantries, besides enforcement to prevent queuing or speeding before price changes. Originally, Singapore had considered a passive ERP system, whereby individuals' car movements would be tracked by a central computer, and the owner sent a monthly bill. Due to privacy concerns that the government would trace the movement of individuals, this version proved to be politically infeasible. By using a smart cash card, the ERP system eliminates the need for tracking vehicles. Transaction records stored in LTA's central computer system are deleted as soon as the payments are settled, usually within one working day. "There are stringent internal security and control measures at the LTA's central computer room to ensure that there is no abuse of the system by the staff (Foo, 2000)."

Although the ERP system has a high initial investment cost, its lower maintenance cost and increased flexibility has allowed it to largely eliminate congestion in downtown Singapore. The cash card system is appealing to customers because it is easy to use, convenient, and protects their privacy. Channeling revenues to mass transit and other low-income transportation modes can offset regressivity of charges. The ERP system, however, is mechanical and subject to breakdown, requiring teams of technicians to be ready in case of emergency. Another shortcoming is that congestion on untolled peripheral roadways has not been controlled. Expansion of the system would require the expense of building gantries at every toll point. Some reconfiguration of the roads would also be required to enable people to divert from the toll

road if they should choose. Overall, the ERP system in Singapore serves as a model for implementing large-scale congestion pricing in cities.

2.5 California Road Pricing System

The FASTRAK project in Orange County, California shows another method of implementing congestion pricing. As described by Richardson and Bae (1998), in this system tolls are applied only on *new capacity*, a 10-mile privately financed road paralleling an existing free expressway (Fwy 91). Vehicles lease a transponder affixed to their windshield that allows automatic tolling when entering the expressway. The toll rate automatically varies to 5 preset levels ranging from $0.25 to $2.75 depending on the level of congestion detected. Vehicles with more than 3 persons (High Occupancy Vehicles, HOV) are given access to a toll-free lane. Toll violators are caught and fined either manually or by an automatic video surveillance system. Before the three entrances to the freeways, there is a one-and-a-half mile segment of road where drivers can evaluate the congestion levels and prices and decide which route to take. The major benefit of this system is that it overcomes the equity issue by allowing choice. The major drawback is that its scope is limited only to areas where access can be controlled and a choice of routes is available.

These two examples highlight how the practical, technical, and political hurdles have been overcome in existing congestion pricing systems. However, both of these systems are still limited in the geographic area that they cover, and are expensive to extend. To maximize equity and minimize traffic spillover effects, the ideal system would be ubiquitous, a system where all urban roads are covered. The main reasons for not building such a system are the administrative and operating costs (Puget Sound Regional Council, 1994). To build gantries throughout an entire metropolitan area, physically reconfigure entrances, and adjust entrances as needs changed

would be an enormously expensive task, even if automatic vehicle identification were used to collect tolls. As we will see, new technologies present the opportunity for drastically lowering or eliminating these costs, possibly making a ubiquitous toll collection system feasible.

2.6 Traffic Data Collection Limitations

Chapter 1 mentions another issue intertwined with congestion pricing; traffic data collection. This is because in order to charge someone per use of a facility, one by necessity must know the usage level of that facility. The Clean Air Act Amendments of 1990 (CAAA 90') and ISTEA 91' mandate the collection of more and new types of traffic related data. Each Transportation Management Area (TMA), a metropolitan area with more than 200,000 people, must create and utilize a Congestion Management System (CMS). These laws are intended to encourage greater cooperation amongst agencies for the purpose of collecting traffic data to address congestion. In practice, the laws have created the push for ITS and the more accurate collection of Vehicle Miles Traveled (VMT) data. However, despite the push from the federal government, there are many deficiencies in current traffic data collection systems implemented by state and local agencies. These results are covered in the Overview of Traffic Monitoring Programs in Large Urban Areas (Mergel, 1997) sponsored by the Federal Highway Administration (FHWA).

The study concludes that there is a general lack of knowledge of traffic data collection systems, lack of coordination among agencies, and no central source of information. In addition, funding and staffing cuts have limited the ability to collect data and implement new technologies. Finally, "The quality of urban area traffic data collection efforts, and presumably of the resulting data, varies widely. While many programs would appear to meet currently accepted standards, many others would not, and in many cases there is no program (Mergel,

1997, p. xvi)." Regarding newer technologies, the study states, "Other technologies such as AVI, GPS, video, aerial photography, and GIS that could have been considered are not yet widely used in practice and are still somewhat experimental, and are more appropriate to the collection of travel time/speed/delay/congestion type of traffic data rather than volume. While this type of traffic data will assume increased importance in the future, it is not currently a significant part of most agencies data collection programs. GIS, while holding significant promise as a data management, data display, and analysis tool, also appears to be in an early developmental stage, at least as far as traffic data collection agencies are concerned (Mergel, 1997, p. 84)." In summary, it is an increased emphasis on technology that is forecast to overcome congestion. Therefore, the next chapter will discuss some of these potential technologies in detail.

CHAPTER 3: TECHNOLOGY

3.1 Introduction

A combination of several technologies provides the basis of the Mobile Ad-Hoc Network (MANET). This chapter provides a brief introduction to them in conceptual sequence. It will begin with a brief introduction to Global Positioning System (GPS). What follows is a general framework for understanding networking devised by the International Standards Organization (ISO). Within this framework, the basics of wireless networking will be discussed. Descriptions and comparisons of various MANET routing protocols complete the building of the MANET technology for wireless communication between mobile nodes. Following this chapter will be the discussion of how this technology solution can be applied to the problem of congestion

3.2 Global Positioning System (GPS)

A GPS receiver allows the user to pinpoint his or her location from anywhere on the Earth by reading signals from satellites in orbit above the Earth. GPS technology is described by Trimble Navigation's web-based GPS tutorial (2001). A constellation of 24 satellites continuously beam signals down to Earth such that nearly every place on Earth has a view of 3 or more satellites at any one time. By comparing and processing the signals from the satellites, (in much the same way cellular phone companies triangulate the location of a 911 caller) a GPS receiver can calculate its position on the Earth in latitude and longitude. A lesser known fact is that a GPS receiver also calculates the precise time from the satellite signals, a requisite for calculating position. However, this can be regarded as a side benefit of knowing one's location.

The accuracy of GPS can now be enhanced by two technologies called Differential GPS (DGPS) and. Carrier-phase GPS. Differential GPS involves the cooperation of two receivers, a mobile one and a stationary one. By comparing signals from both receivers and removing timing

distortions caused by the atmosphere, accuracies of several meters can be obtained for mobile applications. Many public beacon stations transmit the corrections (on a separate band, usually in the 300kHz range) allowing GPS receivers within range to correct their readings. The corrections can also be stored, and GPS readings post-processed for error correction. The aviation industry is also developing Augmented GPS, transmitting corrections using satellites to allow for continent-wide DGPS. Carrier-phase GPS relies on the GPS signal's higher frequency carrier signal to improve timing measurements and accuracy. By combining these two enhancements, accuracies of within centimeters can be achieved at reasonable cost.

One limitation of GPS is that the satellite signals are blocked and bounced by tall buildings, causing inaccurate readings in downtown areas. In vehicles, these distortions can be corrected by using extra sensors in the vehicle measuring the vehicle's speed and direction. This process, called Dead Reckoning (DR), corrects the GPS readings, which tend to occasionally jump to inaccurate positions randomly as the vehicle moves. Such distortions can also be corrected by mathematically disallowing jumps of greater than a certain magnitude, those that would imply an acceleration higher than possible for the vehicle. This requires no additional sensors, but the GPS readings would be only interpolated estimates for those disallowed readings. If the vehicle had a stored and current map of the roadway system, these errant readings could be made even more precise. In any case, this issue will not be so relevant because the system described later is targeted toward large, congested highways and arterials--roadways that the majority of the time will not have many large buildings nearby.

GPS technology is already in use by many people. The applications range from surveying to monitoring vehicle fleets like ambulances. Mobile GPS units are commercially available and mountable in vehicles. The applications of mobile GPS go under the heading of

Automatic Vehicle Location (AVL) services. Examples of such services include safety, security, navigation, dynamic routing, timing, and driver monitoring. The demand for these services continues to drive improvements and cost-reductions in GPS technology. In order for some of these applications to work, some way of communicating with the driver must be arranged. Currently, cellular phones, satellite messaging, or proprietary radio networks are used for this messaging (see Appendix C). An upcoming technology for mobile communications is that of the MANET, whose description begins with an outline of the generic paradigm for networks.

3.3 Framework For Networking: OSI Network Model

A framework for understanding networking functionality is provided by the Open Systems Interconnect (OSI) network model devised by the International Standards Organization (ISO). The network is divided into seven 'layers' each of which provides a specific functionality to the network's operation. Data flows through the layers, being handled at each step in sequence. Typically, the first four layers involve the actual communication of data from host to host, while the last three layers involve non-transmission related processing. An actual implementation of a network model need not have all layers or features, may combine features of adjacent layers, and may split layers apart. For example, Transmission Control Protocol/Internet Protocol (TCP/IP), the lingua franca of the Internet, is designed around a four-layer scheme (The Trustees of Indiana University, 1996). The layers and their descriptions follow (Griffioen, 2000 and Hansen-Childers, 2001):

1. Physical Layer

   Responsible for transmitting raw bits on the physical medium, e.g. fiber optic cable, electromagnetic frequency, and copper Ethernet cable.

2. Data Link Layer

Responsible for the detection and notification of physical errors and establishment and termination of logical links. Responsible for creation of frames (groups of bits) and their transmission.

3. Network Layer

Responsible for flow control, defining the basic unit of data transfer (packet), and the routing of (unordered) data.

4. Transport Layer

Responsible for transmission control, the ordering of received data and acknowledgement of correctly received data. Handles messages, process addressing, and end-to-end reliability.

5. Session Layer

Responsible for controlling sessions and state information between hosts.

6. Presentation Layer

The viewing, encryption, and conversion of data on the network.

7. Application Layer

The end-user applications which use the network.

3.4 General Description of Wireless Networking

At the physical layer, a sender emits a radio signal using an antenna. The electromagnetic signal travels in an omni-directional, multi-directional, or uni-directional fashion depending on the kind of antenna used. The strength of the signal declines over distance, and is roughly proportional to d^{-2} where d is distance. Information can be transmitted on this signal using various encoding schemes, most familiar of which are the amplitude modulation or frequency modulation of AM and FM radio, respectively. Information can also be encoded in binary, as 0s

and 1s, making it digital and understandable to computers. A receiver can pick up this signal using its own antenna and use the information encoded in it. The signal is picked up as an alternating current on the antenna and its strength is a function of the voltage of the signal divided by the background noise. If the signal is above a threshold value, then it is picked up successfully. The details of the encoding scheme, error correction methods, and signal quality assurance are not relevant to this discussion except that they ensure the delivery of data wirelessly.

If there are multiple senders and receivers who wish to communicate on the same frequency, they must make provision for access conflicts (collisions) so that no two senders emit signals to a receiver at the same time. The data link layer logically provides this functionality. If the senders did not cooperate, then the receiver would pick up only the signal of the strongest sender, or be simply unable to comprehend the signal. If the senders can communicate with each other, then there is no problem because each can simply take turns, pass a token, or request permission before sending. However, since we are dealing with wireless nodes that can be in any location, the two senders might be out of range of each other, but within range of a receiver. This is called the "hidden terminal" problem and is explained below.

In Figure 3.1, each letter represents a transmitter, which can send a signal within the range of communication represented by the ellipse surrounding it. Each transmitter can receive or detect a signal if it is within range of the sending transmitter. Two transmitters who are in each other's ellipse can communicate with each other. An example of the hidden terminal problem is that if both A and C unknowingly emitted a signal at the same time to B, there would be a collision at B resulting in data loss.

Figure 3.1: Hidden and Exposed Terminals in Wireless Networks

Another less severe problem is called the "exposed terminal" problem. For example, if B

wished to transmit to A while C wished to transmit to D, there is technically no problem if they

transmit at the same time because A and D will receive only one signal each. The problem

occurs if B senses the signal of C and waits until the medium is free before transmitting. While

the "hidden terminal" problem results in data loss, the "exposed terminal" problem simply

reduces efficiency (Livani, 1999).

3.5 MAC Protocols

A Medium Access Control Protocol (MAC Protocol) is a set of rules transmitters follow

to transmit data over a particular non-sharable medium, in this case an electromagnetic frequency

band, that may or may not solve the problems above. This protocol provides the functionality of

the network layer. The simplest rule is that every transmitter can begin sending data at any time,

and the transmission is simply retried if there is a collision. This rule comprises what is called

the ALOHA protocol, a highly inefficient protocol with a theoretical upper bound of 18.4 % on

the rate of successful transmissions.

"The first step toward high-performance broadcast communication protocols is the

Carrier Sense Multiple Access (CSMA) protocol (Livani, 1999)." The rule in CSMA is that a

transmitter sends data only if it detects a free medium by having sensed a baseline signal level.

A variant of CSMA is CSMA/CA (CA: Collision Avoidance), and is part of the 802.11 wireless

networking standard. In this protocol, the rule is that the transmitter sends data only if it senses the signal level as being free for a fixed length of time (DIFS: Distributed Inter-Frame Space) agreed upon by all the transmitters. If the signal is not free, the transmitter waits for the DIFS plus a random multiple of a fixed time slot. Both of these protocols reduce the probability of collisions and do not solve the hidden terminal problem (Livani, 1999).

The hidden and exposed terminal problems are solved by the Multiple Access Collision Avoidance (MACA) protocol. In this protocol, transmitters can broadcast special data packets called request-to-send (RTS) and clear-to-send (CTS). These inform all other transmitters within range that data transmission is about to take place. In the Figure above, if both A and C wished to communicate with B, they would both send an RTS signal. B would respond with a CTS signal heard by both A and C, but with an identifier saying only C can send. Once C finished transmitting, then B would send a CTS signaling A to send its message. If the RTS signals had collided at B, both A and C would not have received the CTS signal, and would have retried after a random amount of time. This protocol solves the problems efficiently for unicast communication, but behaves like CSMA for multicast communication. As such, the 802.11b standard uses this RTS-CTS handshake for unicast communication, and CSMA/CA for multicast communication. This choice of protocols means that under congested situations, the likelihood of failure is higher for broadcast communication than for unicast communication.

Other MAC protocols solve the problems of medium access in different ways and with different tradeoffs. There are CDMA, TDMA, NAK, TTP, BTMA, FAMA, and many others. The numerous MAC protocols are continually developed, improved, and published in technical journals. The point is that the physical layer constraints of the frequency can be managed by these protocols to allow for efficient communication between mobile nodes, in either unicast

(point-to-point) or multicast (broadcast) fashion. The choice of MAC protocol depends upon the specific requirements of the system being designed. MAC protocols for wireless data transmission exist and are continually being tweaked for better efficiency and reliability. The design choice of the MAC protocol is a technical decision which can be regarded as a black box when considering higher level issues because the solutions are well developed and understood, if not the best they can be. Once this issue is resolved, the MANET requires solution to a higher level (transport layer) issue, how can data be efficiently routed in the face of significant network partitioning?

3.6 Mobile Ad-Hoc Networks

In the game of telephone, one person whispers a sentence into the ear of a neighboring person. This continues along a chain of people until what comes out at the other end is amusingly different from the original sentence. The concept of passing along a message from one person to another is the concept underpinning *mobile ad-hoc networks*. If we break down the term, we get the following definitions: network is a means for communication, *ad-hoc* means there is no reliance on a pre-established hierarchy, and *mobile* implies non-stationary elements. Thus, a mobile ad-hoc network is a system for communicating between moving nodes whose configuration may change at any time.

Mobile ad-hoc networking (MANET) is somewhat synonymous with Mobile Packet Radio Networking, Mobile Mesh Networking, and Mobile, Multihop, Wireless Networking, and has its origins in military research during the 70's and 80's (Corson, 1999). (The bulk of this information was taken from the Internet Engineering Task Force (IETF) Request for Comments (RFC) number 2501. This informational RFC is entitled Mobile Ad Hoc Networking: Routing Protocol Performance Issues and Evaluation Considerations (Corson, 1999).) MANET is related

to "mobile IP" technologies, but not the same. "Mobile IP" technologies are systems that allow

nomadic roaming of hosts across different physical access point to the Internet. The

technologies involve address management and protocol interoperability, and any hop-by-hop

routing exists within the framework of existing fixed networks. By contrast, in a MANET, the

nodes themselves self-configure to form a routing infrastructure, without external support.

However, as we will see later, a working implementation of a MANET does not necessarily have

to be completely independent of support from fixed nodes.

Each node of a MANET consists of an electronic device capable of two things: bi-

directional communication with nodes within range and routing of data. GPS capability may

also be included. The wireless communications operates according to a MAC protocol. The

routing of data, or the passing of messages between nodes, operates according to a routing

protocol. A routing protocol consists of a set of rules by which a node directs the sending of

received data packets. In the game of telephone, there is one simple rule; just pass the message,

slightly changed, along to the next person in line. In an actual protocol, the rules may be

complex and numerous. The routing protocol is essential because nodes must rely on each other

to forward each other's messages. A number of routing protocols are discussed below.

3.7 Mobile Ad-Hoc Network Routing Protocols

Several generic issues are faced by each of the protocols. Each must complete a startup

and initialization phase during which routes or neighboring nodes are discovered. During the

operation phase, one baseline for performance is that the protocol must allow every node to find

and send data to every other reachable node. (A *reachable* node is a node to which a data packet

may be sent by hopping through intermediate nodes.) In addition to finding another node, the

route of a node must be as short as possible, but not necessarily the shortest path. The protocol

must also allow for the propagation of the fact that a link has been broken and the node, unreachable. Link breakage detection or route maintenance can be automatic or on-demand. Automatic requires constant sensing or message passing between nodes while on-demand only requires messages or sensing when a packet needs to be sent. Automatic will tend to require more overhead due to the periodic beaconing/broadcasting necessary to continuously monitor connection status. Another issue is that the protocol needs to inhibit or prohibit looping of data packets, whereby a packet is repeatedly routed between two or more nodes. Protocols can also perform load balancing by nature or by design if some segments of the network become overburdened with traffic. They also have a choice between hierarchical routing, in which the network is divided into areas or clusters, and flat (non-hierarchical) networks. Another design consideration is whether or not and how much to cache network and/or configuration data.

Broch *et al.* have conducted a performance comparison of four recent multi-hop wireless ad-hoc network routing protocols for MANETs (1998). They used the network simulation software *ns-2* to simulate the performance of these network protocols. The simulation methodology intended to measure performance on a theoretical basis in order to evaluate the protocols under a range of conditions. Thus, the simulation parameters are not derived from any real-world application. The simulation parameters were 50 wireless nodes randomly distributed over a rectangle of 1500 m by 300 m moving according to a "random waypoint" model. Nine hundred seconds of activity were simulated. The random waypoint model means that nodes move to a random destination at a speed distributed uniformly between 0 and a maximum speed. Once at the destination, the nodes pause for a fixed amount of time before performing another movement to another waypoint. Numerous movement pattern scenarios were tested and the

performance results were averaged over those runs. To obtain an impression of what these routing protocols involve, a very general description for each follows (Broch *et al.*, 1998):

DSDV – Destination-Sequenced Distance Vector

In this protocol, every node is assigned a monotonically increasing, even sequence number. Routes are maintained at each node and broadcast to neighboring nodes (every fifteen seconds in the simulation runs performed by Broch *et al.*). The routing information consists of the sequence number of the next hop in the chain, the number of hops to the destination, and the destination's address. Packets are routed using the sequence numbers in such a way that looping is prevented. Basically, the packet travels along the highest sequence numbered path with number of hops being equal and minimized.

TORA – Temporally-Ordered Routing Algorithm

This protocol attaches a 'height' value that decreases linearly over each hop to a destination; the height is basically the number of hops. Packets travel 'down' like water towards the destination. Searchers find routes by broadcasting a QUERY packet to all neighbors. Each neighbor relays this packet to all of its neighbors until the destination, or a node knowing the height to the destination, is found. The final node (or nodes) then broadcasts an UPDATE packet that contains the height, and this packet propagates throughout the network, with each node updating its height accordingly. A CLEAR packet is used to reset routing data when a node becomes unreachable. TORA relies on the Internet MANET Encapsulation Protocol (IMEP) for ensuring message delivery to all neighbors and notification of broken links. IMEP uses periodic BEACON and HELLO packets to sense link status and to maintain a list of a node's neighbors.

DSR – Dynamic Source Routing

DSR nodes contain a cache of full route paths between node pairs. To discover a route, a ROUTE REQUEST packet is broadcast and flooded to all nodes. The destination node or a node knowing a route to the destination replies with a ROUTE REPLY packet that is also propagated throughout the network. Each node hearing a ROUTE REPLY message stores the entire path to the destination in its cache. Cached paths are used to aggressively minimize the frequency and propagation of the route packets. ROUTE ERROR packets are generated and broadcast when a link is broken. The protocol is thus divided into two mechanisms, Route Discovery and Route Maintenance. Data packets are transmitted with the entire route stored in the packet itself. The advantage of this protocol is that nodes do not have to maintain the topology of the network, because packets contain their desired path. Also, because Route Discovery and Route Maintenance are performed on an as-needed basis, there is no need for periodic route advertisements or neighbor detection packets.

AODV – Ad Hoc On-Demand Distance Vector

AODV works by combining features of DSR and DSDV. As in DSR, routes are discovered by controlled propagation of ROUTE REQUEST and ROUTE REPLY packets. Each node only keeps track of the next hop towards a particular destination and utilizes sequence numbers to avoid looping. Periodic HELLO messages are broadcast once per second to maintain link status. Three failed messages indicate a link is down. Alternatively, physical layer or link layer methods may be used to keep track of neighbors. UNSOLICITED ROUTE REPLY messages are generated when links go down to reset routing information.

Under the simulations run by Broch *et al.*, each of the protocols performed well in some cases and poorly in other cases. However, testing with 600 nodes by Li *et al.* (2000) of MIT showed significantly diminishing performance as the number of nodes increased. The group at MIT have developed a system based on Geographic Forwarding and GPSR, the Greedy Perimeter Stateless Routing protocol (Karp, 2000). In this protocol, sending nodes are assumed to know the identity and geographic location of the receiver, and this information is embedded in the packet header. Each node maintains only a list of neighbors currently within range, and in this sense, the network is 'stateless'. When a node receives a packet, it forwards it to its neighbor who is closest *geographically* to the receiver. GPSR uses a network theory based routing system to recover from dead ends and avoid loops. Basically, packets travel along the reachable perimeter of the unreachable destination node until they come back to the same point, signifying that the destination node is unreachable. By combining this protocol with a location service, called Grid Location Service (GLS), full protocol functionality is achieved. The GLS allows a node to find out the geographic position of any other reachable node. To accomplish this, it uses a system similar to DSR except that instead of propagating queries to all nodes, only a certain subset of nodes is queried. Only the nodes in this subset are programmed to know the location of the particular vehicle, meaning that all other nodes must query these nodes to find the location. This affords some anonymity to vehicles. The Grid system, or GLS combined with GPSR, was able to sustain what appears to be scalable performance even under 600 nodes moving under the random waypoint model on a field 2900 meters square. These simulation runs took approximately 20 hours to complete (Li, 2001).

Besides these protocols, research on MANETs is ongoing by the military, universities, and the IETF. Even more recent protocols that will not be covered in any detail include TBRPF,

FTSP, STAR, OSPF, and a multitude of others (Bellur, Ogier, and Templin, 2001). A comparison study of recent protocols for *multicast* communications is done by Lee *et al.* of the University of California at Los Angeles (2000). A list of currently evolving protocols can be found at the IETF MANET current Internet drafts website http://www2.ietf.org/ids.by.wg/manet.html. These protocols provide various improvements to the above protocols, and are generally more complex as a result.

The funding for MANETs originally came from military sources for the purposes of developing easily deployable communication systems for applications such as situation awareness on battlefields. As such, research momentum in the area focuses on a number of things. First of all, the protocols must minimize power consumption (by reducing the amount of overhead data transmissions) in order to be suitable for hand-held or power-constrained devices. Secondly, the protocols tend to be suited for small-scale (100's of nodes) networks. Third, the protocols generally do not use GPS or network topology information like road layouts. Fourth, the protocols are usually designed for bi-directional and fully reliable communications. Finally, the protocols are completely independent of external assistance, for both initialization and operation. Because of these characteristics, the protocols under research and development are not necessarily best suited for a metropolitan scale, vehicle congestion mitigation scheme. In this latter case, power is not a major issue, there may be millions of nodes, communications may not need to be 100% reliable, and a fixed supporting infrastructure is not out of the question. However, the work on the protocols generates ideas and allows for testing of the ideas. This is important because a new application, such as vehicle congestion mitigation, can use the ideas, methodologies, and results to create a new, better-suited protocol. Such a scenario is presented in the next chapter.

CHAPTER 4: POSSIBLE SCENARIOS

4.1 Introduction

After introducing the issue, discussing congestion pricing, and covering the basics of the technology, the argument of this chapter is that a congestion pricing system could be suitably implemented using existing technology. The required technology can be divided into three distinct logical components: Short Range Wireless Data Packet Transmission, Geographic Positioning System (GPS), and Data Routing Protocols. Using these technologies, what is presented is a hypothetical scenario for implementing a basic system in the Boston Metropolitan area based on a Mobile Ad-Hoc Network.

We first calculate vehicular densities for Boston, MA from statistical data to get a sense of the highway system. Next we present a basic hopping protocol for collecting traffic data. Following this, we describe three different applications and how they can be implemented using a distributed mobile network. For each application, hardware and technical requirements will be discussed, issues covered, and conclusions drawn about their rationale and suitability.

The first application is a traffic data collection system that can retrieve vehicle positions for a sample of vehicles on congested highways. The necessary physical infrastructure of this system broadly separates into two pieces: gateway devices and mobile devices. The traffic statistics for Boston and estimated prices for the components lead to an estimation of costs. The second application is the provision of consumer and business information services for which people would be willing to pay. The third and final application is location-based transaction processing, including congestion pricing. We will see that each application has an increasingly demanding set of requirements.

We focus mainly on the first system, the data collection system, because it would be difficult to satisfy the requirements of the other systems immediately. It would be nearly impossibly and prohibitively expensive to install perfectly working devices in all vehicles all at once. We also do not know all of the possible useful applications. Also, the later systems require more reliable data transmission, something that would not be achievable in the short term. We discuss the later systems as possible paths for the system to grow into, but these are more prospective and therefore not as fully fleshed-out. The conclusion is that the MANET, once developed, will be suitable for implementing measures for congestion mitigation.

4.2 The Boston Metropolitan Area

In order to determine the feasibility of a short-range wireless network on the highway system, one must calculate the approximate density of vehicles on the roads at each hour. These calculations will result in a crude, but meaningful approximation of the distance between vehicles and the number of vehicles per mile of highway during each hour of a typical day. Caveats to the assumptions underlying these calculations will be discussed at the end of this section. From the density estimates, we calculate the required number of probe vehicles as a percentage of total vehicles, the sample rate. A probe vehicle represents a node participating in the MANET. Data from Table HM-71 and Table HM-72 of the Federal Highway Statistics in 1997 provides traffic and demographic statistics for the Boston, MA Federal Aid Urbanized Area as defined by the Census Bureau (Federal Highway Administration, 1998b). (Boston is the 8[th] largest in the US by population.) A summary of relevant information from these tables is presented in Table 4.1 and Table 4.2 below for later reference. Calculations are done only for major highways including interstates, freeways and other expressways. These are generally the

largest, most heavily used roads and hence would be the first target of a system reliant on density of vehicles.

Table 4.1: Boston, MA Roadway and Travel Distribution, 1997		
Classification	**Miles of Roadway**	**Daily VMT**
Interstate	147	17,373,000
Other freeways and expressways	67	4,424,000
Other principal arterials	764	16,110,000
Minor arterials	1,248	9,144,000
Collectors	1,178	3,445,000
Local	6,668	7,790,000
Total	10,072	58,286,000
Source: *Table HM-71, Highway Statistics 1997*		

Table 4.2: Boston, MA Demographics and Traffic Summary, 1997	
Name	**Number**
Estimated Population	2,890,000
Estimated Vehicle Population	2,5000,000 (1,800,000 cars; 700,000 trucks)
Net land area	1138 sq. miles
Total freeway DVMT (Daily Vehicle Miles Traveled)	21,797,000
Total DVMT per capita	20.2
Percent of total DVMT served by freeways	37.4%
Average daily traffic volume per freeway mile	101,855
Total miles of freeway	214
Total estimated freeway lane miles	1304
Average daily traffic per freeway lane	16,711
Source: *Table HM-72, Highway Statistics 1997*	

The density is calculated using two different methods in the tables in Appendix A. Both methods are based on a proxy for hourly traffic volume taken from a traffic calming study done on a commuter collector road in Nebraska (Office of Transportation, 1996). This study included a graph of hourly traffic volumes over a 24-hour period. These counts are used to calculate the percent of vehicles that are on the road during each hour, and these percentages are used as weighting factors for our traffic density calculations. Another assumption we make is an

estimated truck length of 50 feet and car length of 15 feet, giving a weighted average length for the Boston vehicle fleet of 24.8 feet (Federal Highway Administration, 1998b, Table HM-71). An average highway speed of 55 miles per hour is assumed. The number of vehicles was actually 750 at 5:00 PM, but was reduced to 500 to make more sense for highways, which should typically have a more even distribution of volume. Further explanation of this adjustment is given later.

The first density estimate is derived from the total freeway daily vehicles miles traveled of 21,797,000. This number is multiplied by each percentage to obtain the number of vehicle miles traveled each hour during a typical commuting day. Next, these numbers are divided by the total miles of freeway to obtain the hourly traffic volume in each mile of freeway. For example, if vehicles traveled 100,000 miles on freeways during the 9:00 AM to 10:00 AM interval, and there are 100 miles of freeways, then 1000 vehicles traveled on each mile of freeway during that hour. Note that this assumes each highway mile is traveled equivalently throughout the metropolis. As mentioned earlier, this density measure is similar to the one used by the Texas Transportation Institute to calculate area-wide congestion levels (1992). We also calculate the hourly traffic volume in each lane-mile of freeway, by dividing by total freeway lane miles.

To calculate density, we assume that the entire volume of vehicles passes through the mile at a constant speed and with an equal distance between each vehicle. At the beginning of the hour, the first vehicle is at the end of the mile-long segment. At the end of the hour, the last vehicle has just entered the mile-long segment. This way, the segment has vehicles on it throughout the whole hour. Thus, in order to calculate the number of vehicles on any mile-long stretch of highway at any given point in time during an hour, we divide the hourly traffic volume

during that hour by the average vehicle speed plus one. For example, if there are 1000 vehicles traveling over a mile of highway between 9:00 AM and 10:00 AM at 55 miles per hour, the number of vehicles on that mile-long stretch of highway at any point in time is on average about 18 (1000/56). This gives us the vehicle density in terms of the number of vehicles per mile of highway.

Next, average distance in feet between each vehicle in each lane is calculated assuming the even distribution of volume as before and with each vehicle taking up 24.8 feet. Here is the formula:

Average distance (ft.) between vehicle in lane = (((milesperhour+1)*feetpermile)-
(hoursvolume*vehiclelength))/(hoursvolume-1)

Or using a train analogy:

Distance between each train car = (length of train – total length of train cars)
/ number of spaces between train cars

The second density estimate is derived from the estimated number of vehicles registered in the Boston urbanized area. The major assumption is that all vehicles are used during the day and that each vehicle is driven for one hour on the highways (discussed later). The number of registered vehicles is multiplied by the proxy percentage to obtain the estimated number of vehicles on the highways each hour. These numbers are divided by the total highway miles and the total highway *lane* miles to obtain the estimated number of vehicles per highway mile and per highway lane-mile, respectively. From these numbers and using the assumptions and formulas as before, we arrive at the estimated distance in feet between cars in each lane.

Finally, the required sample rates (number of probe vehicles out of all vehicles) for various levels of probe density are calculated. For example, if we would like to have on average 2 probe vehicles per mile, and there are 20 vehicles per mile during that hour, then we need a

10% sample of the vehicles (2/20 = 10%). We calculate the sample rates required for having

one, five, and ten probe vehicles per mile. In cases where the required sample rate exceeds

100%, there are less than the required number of probe vehicles in that mile during that hour. In

other words, if there are only five vehicles on the highway segment, we cannot know the

locations of ten vehicles even with a 100% sample rate.

These results are based on very rough estimates of actual vehicular density. They do not take

into account the stochastic nature of vehicles' locations. In other words, even if on average there

is one probe vehicle per mile, on any given stretch of several miles, there may be no probe

vehicles or numerous probe vehicles. In addition, the assumptions of even and equal distribution

of all vehicles or all vehicle miles traveled over the highways will not reflect reality. For the first

calculations, the fact is that some sections of highway are traveled more than others. For the

second calculations, all registered vehicles are not used during the day and some are used more

than once. Nevertheless, one gains a rough sense of vehicular densities on the freeways, and we

can use our first-hand experience of the highways to fine-tune our interpretation of the results.

As we can see in Appendix A, there is more density when numbers are derived from the total

number of vehicles than when they are derived from vehicle miles traveled (VMT). The

following distance between vehicles is 100-300 feet during peak periods for the VMT-based

calculations while it is only 10-30 feet for the vehicle-count-based calculations. This is probably

due to the fact that not all registered vehicles are used during the day on the highways, contrary

to the assumption. In addition, each vehicle is probably not used for a whole hour. From this

point of view, the calculations based on VMT may be more accurate, because they reflect actual

usage of the highways exclusively. However, the more dense results of the second calculations

give us results for heavily congested segments of highway, which truly exist. The adjustment

from 750 to 500 vehicles mentioned earlier is due to the fact that at 750, the space taken up by the vehicles estimated to be on the highway would have exceeded one mile, the available road space.

From these results, we can estimate the bare minimum requirements to implement a basic system based on peer-to-peer relay between vehicular nodes. For now, we assume a random sample of vehicles are selected to be participating nodes. Also, we use the results of the VMT calculations, because these are both more conservative and likely to be more accurate. Here is a summary of results for various sample rates:

- To obtain 1 probe vehicle per mile on freeways during only peak congestion periods, approximately 0.5% of vehicles must be nodes, or 12,500 probe vehicles in the Boston urbanized area.

- To obtain 1 probe vehicle per mile on freeways during all heavy use hours, approximately 1% of vehicles must be nodes, or 25,000 probe vehicles in the Boston urbanized area.

- To obtain 2 probe vehicles per mile on freeways during all heavy use hours, approximately 2% of vehicles must be nodes, or 50,000 probe vehicles in the Boston urbanized area.

4.3 A Simplified MANET Routing Protocol For Traffic Data Collection

The above calculations provide estimates of sample rates and distances between vehicles. What follows is a description of a simplified MANET routing protocol that could take advantage of the low sample rates and high distances previously calculated. In order to demonstrate a reasonable routing protocol designed especially for an actual implementation system in highway vehicles, this thesis introduces a very simple, smallest hop-count algorithm. This implementation focuses simply on collecting periodic position information from each probe

vehicle. The protocol is called the hopping algorithm and is similar to GPSR, which was discussed in the last chapter.

The hopping algorithm is a simple, *uni-directional* protocol for transferring data from a vehicle to an Internet IP address. It requires a device consisting of a wireless transmitter, GPS device, and low-end processing system to be installed in each vehicular node. The device acts as a router for data packets and as a source of data, namely, the location data for the vehicle. Location data is comprised of *latitude* and *longitude*. Each device has an expected *range* of communication in meters. It has no need for uniquely identifiable properties in the simplest case, and hence the protocol would be anonymous. The entire system is composed of two basic elements, routers (installed in mobile vehicles) and sinks (installed at fixed sites). We will refer to the routers interchangeably as probe vehicles, mobile locations, or nodes, and sinks interchangeably as gateways, fixed locations, or nodes.

The sinks take data off the wireless network and either process the data themselves, or transmit it over the conventional Internet to a processing location of choice. Routers are made aware of the presence of sinks by hearing the sink's presence broadcast message given at a periodic interval, *beacon_period*. Other sinks are simply notified of presence through the Internet. In such a way, all elements connected to the network become aware of the location of the sink. The sinks are the destinations for all data transmitted from cars. Therefore, their locations must be known by each router so that packets can be efficiently routed over the network. The sinks can also keep track of other sinks, but this information is not necessary.

Each router also sends a broadcast message to other routers to make them aware of their presence. This broadcast message is not relayed. Upon hearing such a broadcast, a router simply updates its list of neighboring nodes, which contains all routers within range. This is

similar to the IMEP protocol described last chapter. The broadcast is done every neighbor_refresh seconds, e.g. 500 milliseconds. The receiver of the broadcast updates the neighbor's current location, previous location, number of hops to the closest sink, and distance (optional & calculated). Each router also maintains some information about itself, including its present location, previous location, next expected location (optional & calculated), and the number of hops to the closest sink. If a router has no neighbors, the number of hops it assigns to itself is infinity. The number of hops is set to 0 if the router's neighbor list contains one or more sinks. Otherwise, the router scans its neighbor table for the neighbor with the lowest hop count and sets its hop count to 1 plus the hop count of that neighbors. As a further simplification, if routers and sinks are treated equally as neighbors, and the number of hops of a sink is always set to -1, then the number of hops is calculated as simply the minimum number of hops from the list of (all) neighbors plus one. Because of the importance of knowing the location of a sink, the change in hop count from infinity to a real number could trigger a beacon immediately rather than waiting for the next scheduled beacon. This would not increase overhead if the next scheduled beacon was skipped, but propagation of sink knowledge will be much faster.

Suppose a vehicle wishes to transmit its location to the sink every second. Every second, it creates a data packet containing a source location, time stamp, time-to-live (say 5 seconds), and a maximum number of hops (say 100). (The data packet could also contain a uniquely identifying serial number, or a serial number that is randomly generated per trip, or some variation in between. Each method has some tradeoffs regarding privacy, and trip trackability.) The entire routing protocol can be described by one rule: the data packet it transmitted to the neighbor with the minimum number of hops to a sink. This works because all nodes keep track of the presence of other nodes within communications range. For example, if the router were

within range of a sink, it would simply transmit its packet to the sink in one 'hop'. If the router were not in range of a sink, it would transmit the packet to a neighbor closest to a sink. If the router had to relay a packet, it would simply follow the same rule.

In order to increase reliability, each router keeps a buffer of packets to be routed (*buffer_packets*). This buffer fills up when a node has no neighbors or when all neighbors do not know the location of a sink. To prevent looping and overflowing of packets, the time-to-live and maximum number of hops are used to expire packets. Also, packets should be expired when the buffer is full. In this way, packets would eventually reach a sink if a sink were reachable. Upon reaching the sink, the location data are processed and sent over the Internet to a central collection server. For example, the sink could buffer the incoming location data packets in a file. Once every minute, it could transmit the file via a secure connection to a web server that would insert the data into a database. The entire process can be completely automated.

We use a hopping method instead of a straight-line-distance-to-sink method. Therefore, the routing protocol does not need any of the GPS information. Technically, by knowing the location coordinates of every router and sink, a node could simply route to the node closest to the destination geographically. However, the transportation grid is not evenly distributed across the geography of the Earth. If it were true that vehicles with routers were evenly distributed across an evenly distributed transportation network, then a straight-line method would work better. However, the transportation network is like spaghetti and vehicles are not evenly distributed. Therefore, sending a packet the closest distance to an already known sink might often end up sending the packet into a dead end under some configurations of the roads. The GPSR protocol overcomes this issue, but does so by adding complexity to the protocol. Also, the hop-count

method by design seems to always route along the shortest available path, while GPSR would sometimes takes less than optimal routes (this statement requires formal proof).

The parameters of this system could be empirically determined through modeling, simulation, and experimentation. These calculations and simulations are well beyond the scope of this thesis. We can, however, determine the goals of such analysis. It would need to incorporate the stochastic nature of vehicle position, which would create disconnects in the network. Such effort would determine the optimal parameters for the range of the device, the size of the packet buffer, the beacon interval, the time-to-live, and the maximum number of hops. The choice of these parameters would simultaneously influence the desired number and sample choice of probe vehicles and the desired number and placement of sinks. All of this would begin with consideration of the particular transportation grid, target locations, and applications under consideration.

4.4 Initial Application: Traffic Data Collection System Description

The first goal is to obtain a sample of peak traffic flows during rush hour on metropolitan Boston freeways. This goal would be the simplest for the first major system implementation. (The first few tests would be local test cases in controlled environments in order to determine the parameters.) This particular application can take advantage of several natural characteristics of MANETs. The brevity of the data packets containing position information has benefits because there is less possibility of collisions when vehicles move into proximity of existing conversations. Because vehicles tend to move in unison, node mobility into and out of range would be less under forward to backward communications. Another advantage is that the information being collected is non-essential. Finally and most importantly, this application does not need every vehicle to be outfitted with a device.

There are a number of technical modifications that improve efficiency. Since the periodic beacons by routers also carry their location information, there is actually no need for a separate data packet to transmit location information. If the beacon length were 500 ms and we wanted to collect location information every 1 s, then we can simply attach a relay flag to the beacon that is set to 'true' every other beacon. Receivers would know to relay a beacon packet according to the routing protocol if its relay flag were set to true. This improvement would reduce overhead. The availability of more channels or alternate means of neighbor detection would significantly increase bandwidth and improve reliability by allowing non-conflicting and more frequent neighbor updates. However, we are assuming that only one channel or band is available for both beaconing and data transmission, as is true in the 802.11b protocol.

Another technical modification would be to adjust the beacon period depending on the speed of the vehicle, making it more frequent at high speeds and less frequent at low speeds. The beacon period must take into account the fact that nodes are moving at up to 80 miles per hour, or 117 feet per second. If the range is typically 5280 feet and we want at most 2% of it to be traversed in one second, then we should have a beacon period of $0.02*5280/117 = 0.9$ seconds. The GPS data could also be combined with a regional roadway map to predict possible next locations and improve routing efficiency. However, this would require changes to the protocol that were dependent on a fixed road layout. The roadway map would need to be updated as infrastructure changes were made. Such a protocol improvement is possible, but further research would be required. Another more advanced improvement would be to use the precise time and location coordinates from the GPS receiver to create a new multiple access protocol for handling MAC issues. Theoretically, a complex combination of time, frequency,

and position could be used in the protocol. However interesting, these improvements are not

necessary to implementing a working system.

By using a stepped power level, the radio transmitters can reduce interference. For

example, the transmitter could begin by using its lowest power level. If no neighbors were

found, then it could increase its power level, repeating until it reached its highest power level.

The importance of this modification comes to light in heavily congested areas, where hundreds

or thousands of nodes may be within the maximum radio range. By reducing the power level of

the transmitter when more than a certain number of neighbors are detected, interference can be

reduced. For example, the lowest power level might work only from ten to fifteen feet, while the

highest power level may work from two to ten miles. A *max_neighbors* value can be empirically

determined depending on the performance of the multiple access protocol under congestion.

This allows scalability and predictability of the network from dense configurations to very sparse

configurations.

By modifying the number and location of sinks, we can make sure there is enough

bandwidth to transmit all the data. For example, suppose there are 50,000 vehicles, each

generating one data packet each second containing latitude, longitude, ID, and timestamp, say a

16 byte (128 bit) packet. The total data rate of all vehicles, if they are on the network all at once

is $50,000 * 128 = 6,400,000$ Bits Per second, or 6.4 Megabits. This is within the maximum

capacity of 10 Megabits for wireless networking. However, in practice, there will be some

constraints on this maximum bandwidth. Namely, we have to subtract from the maximum

capacity the overhead of the networking protocol and make another subtraction for interference.

For argument's sake, assume we have a two Megabit data rate after accounting for these things.

In this case, if there was only a single sink, then there could be a point failure if all the data

funneled down to a single line of vehicles. In other words, demand would be 6.4 Megabits through that path, but capacity would only be 2 Megabits. To avoid this situation, we can add an additional gateway and locate it far away from the first one. In this way, packets will hop to a closer gateway and we effectively reduce the demand on any single link by a factor of two. Adding more gateways further reduces the load. Thus the hopping algorithm by design allows for easy load balancing. The best geographic distribution of sinks should be empirically determined, not theoretically. In any case, to allow for 6.4 Megabits, we need at least 4 sinks, assuming each sink is geographically located to split the observed demand in an even distribution. Appendix B describes the range and bandwidth of currently available networking protocols. Note that the increased outdoor bandwidths require modified antennae because the FCC regulates power output on the bandwidths used by those protocols.

Depending on which vehicles are chosen to participate and how many vehicles participate, we can determine the range of communications required by each node. For example, the 2 % sample rate (50,000 vehicles) would provide on average 2 vehicles per mile with an average spacing of ½ mile, but this does not imply a minimum device range of ½ mile. In reality, if a true random sample was used on a very large metropolitan area, the range may need to be several miles in order to work on the congested routes, or the packet buffer and timeouts fairly large. However, it would still work reliably in the most congested and heavily used areas. The actual required range and sample rates will be less in these areas. However, the greater the range, the more geographic area is accessible and, equivalently, the less likely nodes will be in disconnected partitions of the network. We can also reduce the required range by taking a sample of vehicles whose travel characteristics are known. It would be prudent to do this partially anyway because installing devices in unused vehicles would be wasteful, but not too

wasteful if the devices became cheap enough. For example, if we select daily commuters to Boston from a single suburb, then the traffic data we can expect to collect will be data along the highways from Boston to that suburb during the morning and afternoon rush hours. We would naturally locate a sink along one of these highways.

To increase range further, we can mount antennas on top of the vehicle. This should usually allow for line of sight communications between probe vehicles on the highway. The only obstructions would be temporary ones, from curves in the road and from trucks. For front-to-back communication along a straight road, there would be no trees or other water-containing matter to block the low-power microwave signal. Possibly, the antennas could be made to project their signal directionally in a cone to the front and the back of the vehicle. This would reduce crosstalk across parallel highways or intersecting highways and increase the communications range significantly. Since highway curves are typically long and nodes can buffer packets if they are temporarily out of range, the drawbacks of this modification would be limited. Extremely hilly or curvy local roads would present a larger problem for directional antennas, but conveniently, we usually do not need to collect traffic data from these roads. Other edge cases to consider are an isolated sharp turn, two parallel congested freeways, an isolated T-intersection, a cloverleaf highway interchange, a complicated multiple freeway and roadway intersection, light traffic condensing suddenly into congested traffic, and vice versa. The actual best radio footprint can be determined by modeling the city and testing various configurations. Table 4.3 shows the range of an off-the-shelf wireless Ethernet adapter, which does not use a directional antenna. With slight modification through increased power or a directional antenna, we can see that inexpensive and readily available hardware supplies enough range and bandwidth for our purposes.

Table 4.3: Operating Range Of Linksys Instant Wireless PCI Adapter WDT11 Cost $49.95 at Outpost.com	
Indoor:	Outdoor:
50M (164 ft.) @ 11 Mbps	250M (820 ft.) @ 11 Mbps
80M (262 ft.) @ 5.5 Mbps	350M (1148 ft.) @ 5.5 Mbps
120M (393 ft.) @ 2 Mbps	400M (1312 ft.) @ 2 Mbps
150M (492 ft.) @ 1 Mbps	500M (1640 ft.) @ 1 Mbps
Source: *Product Information*, [Online] Available: http://shop3.outpost.com/product/61576/ [Apr. 23, 2001].	

4.5 Initial Application: Public and Private Institutional Uses

Two probe vehicles per mile should be enough to do sophisticated travel demand statistics collection. By collecting second-by-second position updates from vehicles, we can calculate that vehicle's speed on any particular stretch of highway. By knowing the speeds of several vehicles on the highway, we can guess the prevailing conditions. For example, at 9:30 AM, three vehicles report their position data on a particular half-mile long segment of highway, and we calculate their speed to be 30 miles per hour. At 9:35, three other vehicles report data from that segment and show a speed of 25 miles per hour. From this information, we can reasonably conclude that the segment of highway is congested, and its speed is reduced to 25-30 miles per hour. Over time, we can make graphs of typical speeds on various segments of highway. Thus, we can pinpoint trouble spots on the highways, or places where addition investment or mitigation measures are needed. If our sample is random and large enough, we can even estimate traffic volumes on the highways during all heavy use hours. Indeed, the calculations in Appendix A show that with a 2% sample rate, we can collect data from 7:00 AM to 6:00 PM. (The required sample rates for the 1 probe vehicle per mile set are roughly below 1 % during these times in the DVMT table.) Once this data is collected, we can correlate it with other streams of data such as air pollution, weather, major sporting events, holidays, etc.

Such detailed position information is the ideal traffic information for highway administration agencies. The data can be filtered, packaged, and analyzed and made available to state, local, and the federal level through the Internet. The data could replace costly origin-destination surveys, and allow for more efficient allocation of highway development resources. Types of data that can be produced include traffic volume, vehicle classification, travel time, speed, vehicle miles traveled, vehicle occupancy, origin-destination, and truck weight. Some of these would require additional pieces of information from the vehicle, such as attributes, status, and identification information. As described in Chapter 2, this data would fulfill much of the requirements mandated by ISTEA 1991 and the CAAA of 1990, which require transportation management areas (TMAs) to develop congestion management systems (CMS) through data collection efforts. By going through the process of building the system, TMAs would require collaboration at a more metropolitan level and increased intergovernmental coordination. These process improvements are also mandated by the federal government, so that regions can deal with congestion and air quality issues.

Besides the government, this information can also be used by the private sector. In Advanced Transportation Management Systems (ATMS), the public and private uses are intertwined. For example, the traffic data collected by the government agency could be distributed to the public by a private agency for government purposes; viz. to improve traffic flows. However, the primary benefit for commercial users would be for metropolitan-area vehicular fleet tracking. Private companies can be given access to detailed vehicle position data for a fee. This might be highly valuable for certain types of vehicles. Commercial fleets in the shipping industry such FedEx, UPS, and long distance shippers; delivery companies such as Dominoes Pizza, Kozmo.com, and FTD; rental companies such as Avis and Hertz; and company

and institutional fleets would be potential customers of the data. Note that long-distance (interstate) tracking is not available unless the system is coupled with another vehicle tracking system. Some of these uses would benefit from additional information being collected from the vehicle. For example, a taxicab could relay information about whether it is free or occupied. Government fleets such as police cars, fire engines, ambulances, and other government vehicles can also use the fleet tracking ability. A summary of potential uses is listed in Table 4.4.

Table 4.4: Potential Applications of a Metropolitan Vehicle Tracking Service	
Customer	Application
Taxi Cabs	A central location can have a graphical display of all vehicles and their locations. When customers call in to request a ride, the closest free cab can be dispatched to the customer's location.
Package Delivery	A central location can have a graphical display of all vehicles and their locations. When customers call in to request a package pick-up, the administrator can update the route for the delivery van to include the customer pick-up. This allows for same-day pick-up service.
New Delivery applications	A company could have a fleet of vehicles whose sole purpose is to handle local deliveries, like food, groceries, etc. When requests are placed with the system, the vehicle's route can be updated, e.g. by the driver phoning in or being paged. Vehicle tracking would increase the efficiency of this service and improve the precision of delivery time estimates.
Business Vehicle Fleets	The business vehicle fleets of local public or private organizations can be tracked to ensure legitimate use of the vehicles. Asset tracking can be done for commercial and public transportation carriers.
Ambulance, Fire, Police Dispatch	Dispatchers can route the closest vehicle to an emergency situation.

4.6 Initial Application: Individual Driver Uses

From the perspective of an individual driver, the benefit of this system is not so clear. Individuals would first have to be convinced that their privacy was secure, the single most important factor that could easily prevent widespread acceptance. Privacy issues will be discussed in the next chapter, but there are a few smaller benefits an individual may obtain by allowing the tracking device to be installed in their vehicle. First of all, there is a weak security benefit; if the vehicle is stolen and thief is unaware, the vehicle's location can be tracked. The device and its antenna can be securely fastened inside the vehicle to decrease the likelihood of

tampering. This benefit is weak due to the possibility of tampering and also due to jamming technologies currently under development that would render the device useless. For example, jammers being developed to quiet cell phones in theatres, restaurants, lecture halls, etc. could easily be used against this device. Furthermore, there is no benefit anytime or place when the device is out of range of the MANET, like at night. Another weak benefit may be personal gratification and perhaps bragging rights for participation in a transportation research study. This would appeal to certain types of people even if there were no other tangible benefits.

The benefits are weak if we assume that people do not want to be tracked and the device is part of a government program. However, tracking ability may soon be demanded by people who desire such applications as route guidance, emergency roadside assistance, automatic crash notification, other customized traffic information, and concierge services. Also, people may wish to access the location of their family member's vehicle for various reasons. In the end, a government tracking program may need to provide some monetary incentive for installing a pure tracking device with no other benefits for the individual. This could be in the form of rebates on vehicle registration fees. The revenue collected from commercial users could be used to offset subsidies paid to individuals, if no other benefits are provided to the individual. Also, the government could guarantee privacy and anonymity, as it does through the IRS and the Census bureau. The device's identifier can be unlinked from the individual's information. It is important to note that the device requires no maintenance on the part of the driver because it automatically initializes onto the network and begins transmitting data as soon as the vehicle is turned on, and stops once the vehicle is turned off.

Another way to allay privacy concerns is to have optional blackout zones from which position information is not relayed. These blackout zones would cover local and private roads

and would be pre-programmed into the device using roadway data. Conveniently, position information from these roads is not altogether necessary for the purposes of congestion management. For example, a subdivision with one entrance might be blacked out, but the arterial it empties onto might be included. Thus, any vehicle leaving or entering would be known to have a destination or origin in that subdivision, but no other information could be gleaned. Most places, however, have more than one access point; thus, it would be nearly impossible to regenerate an individual's movements without knowing some other information about that individual, like when and where he/she worked. But in this latter case, significant privacy has already been lost. Once local areas are blacked out, vehicles entering and leaving highways and major arterials would be anonymous. The blackout areas can be advertised to the individuals to reassure them of their privacy. This would be akin to what the Census bureau does by aggregating individual data to a point where personal information is hidden, but meaningful statistics are still available. By having a small random sample of paid users who do not black out their local data, one can recover origin-destination survey information. The smaller the sample of users, the more blackout zones are needed to guarantee anonymity. In the beginning, with a very low sample rate, perhaps only major highways would be included.

4.7 Initial Application: Implementation and Recap

Both public and private entities have use of such a system, and therefore the system would be built ideally in a public-private partnership. The extreme sensitivity of data and results of data analyses necessitates that a private venture be regulated or a public venture be closely watched. The first step to implementing the system would be to design the devices and simulate the network to obtain appropriate parameters. Second, one would create a statistically valid random sample of target vehicles, or an alternative sample based on known driver characteristics.

Third, the agency would need to install the devices into vehicles and locate and install sink devices. Once this is completed, data collection can begin. The system must be reliable enough so that the data can be automatically filtered. One should be able to determine the difference between a vehicle stopping at a place, turning off, then moving on, and the device going out of range or malfunctioning, then coming back on. The performance of the system can be measured by how much and how reliably data is collected, and by maintenance checks of devices.

Because this is a metropolitan-level system and the U.S. does not have explicit, elected metropolitan-level governing bodies, it would most likely be implemented at the State level. The State could authorize the project to be organized through a Metropolitan Planning Organization (MPO), which typically consists of local and county leaders. The system could be part of a metropolitan Congestion Management System (CMS). Being implemented at the State level does not rule out incentives or funding from the Federal level; such inducements would probably be very likely through existing ITS initiatives. Implementing the system at a local, single-city level would be inadvisable because many benefits would be lost leading to a higher cost to benefits ratio. If it is implemented at the State level with Federal funding, then the existing State/Metropolitan highway or transportation agency can handle the traffic statistics collection and analysis. This would be beneficial because this task is already their core competency. Generally, state transportation agencies are heavily weighted in terms of institutional momentum and workforce composition towards traffic and civil engineering.

Building, testing, and deploying the networking infrastructure would most naturally fall into the hands of a private-sector telecommunications company. They would have the expertise needed to understand the technical issues related to sophisticated electronics. They would also have a natural advantage in the ability to market and sell a complex technology to consumers.

Furthermore, a telecom firm would be more nimble and flexible to the rapid changes in technology. Having the telecom firm as the public interface may also be advantageous because these firms generally have a relatively better repertoire with the public in terms of customer service, the ability to provide information, and image.

Similar to Singapore, this technical aspect of the system could be contracted to a large telecom company in exchange for a lump sum payment. In return for providing all this, the telecom would receive revenues from selling fleet tracking information, fast track approval of wireless spectrum, and lowering of any other regulatory hurdles. The State would receive full rights to the use and dissemination of information. It would also retain control and oversight provisions to prevent abuse and encourage standardization in the best interest of the public. In any case, the details of this partnership would be highly dependent upon the political situation and institutional framework of the particular State, and there is considerable room for flexibility. Significant bargaining and trading amongst parties will be required to achieve the most equitable and mutually beneficial scheme.

Having the Federal government's involvement would be necessary to prevent market segmentation caused by competing technologies. States would not automatically have an incentive to standardize, causing inconvenience to consumers who would have to install multiple, redundant devices in their vehicles. An alternative way to gain the benefits of competition would be to allow different metropolitan regions to implement different *initial* systems and compare results. Once these first systems are played out, then the Federal government could step in to establish national standards for MANET networking protocols and, as we will see later, congestion pricing systems. Regions could implement those standards when they upgrade their equipment. In such a way, drivers and truckers could drive across the country

without worrying about how their position information is collected or tolls deducted in any one particular State.

To briefly recap, with a sufficiently large sample, which is estimated to be about 2 % of all vehicles in a metropolitan area, one can achieve the goals of traffic data collection on sufficiently dense and congested roads. With higher sample rates, congestion prediction can be done by observing when a large number of vehicles are headed towards a bottleneck area. This congestion news and traffic data can be provided to individuals and companies in exchange for having a tracking device in their vehicle. The system provides ubiquitous and instant availability through sampling of all traffic parameters including vehicle presence, volume, speed, and density; a complete, metropolitan-scale picture of traffic conditions. In addition, technical steps such as black-out zones and unidentifiable vehicles can protect individual's privacy. The next section discusses how much such a system would cost and its comparison with other systems.

4.8 Initial Application: Cost Calculations

In order to ascertain the cost of an initial MANET traffic data collection system, we must analyze the cost of the various components. We use today's prices in our estimate, but expect the cost of electronic components to fall dramatically over time. The cost reduction is historically true and there is no reason to expect that it will not continue. Combine that with the fact that transportation issues are extremely long range, and we have virtual certainty that the cost of electronic components will one day become very affordable. The cost components fall under the following four categories:

– Mobile Device (variable location), *mobile node*

  – Wireless Data Transmission: Antenna, DSP, Amplifier, Processor, etc.

- Routing and Control: CPU, memory, software, supporting logic.

- GPS: antenna(e), DSP, etc.

- Gateway Device (fixed location), *fixed node*

  - Same as Mobile Device, except GPS is optional.

  - Internet Connection

  - Power Source

- Data Processing and Collection Center, *data center*

  - Database System

  - Internet Connection

  - Additional Software for Analysis and Dissemination

- Administrative, Management, and other Soft Costs

The traffic data collection scenario presented requires thousands of mobile nodes, one to one hundred fixed nodes, and a single data center. The number of government staff would be comparable to a typical Congestion Management System (CMS) initiative, consisting of a moderate amount of executive, administrative, and technical staff. One can see immediately that the cost of the mobile nodes comprises the bulk of the initial deployment cost. However, the maintenance cost of these nodes is nominal once they are deployed. The other capital costs are relatively minor in comparison. The hardware and maintenance cost of the data center will be the comprised of the cost of the high speed internet connection, the cost of the servers, the cost of the people to run it, the building rent, and electricity. For an initial deployment, this could simply be a room in a government transportation building near offices for administrators of the project.

The cost of the fixed nodes depends upon where they are installed. The beauty of this system is that these nodes can be installed *anywhere* within range of a major road. The hardware size and power consumption will be comparable to that of a desktop computer. Since there is so much flexibility in locating this node, one major assumption we make is that the location cost of these nodes can be minimized drastically. For example, the node could be placed in a state-owned building and connected to that building's existing Internet connection and power supply. Alternatively, the node could be placed in the space of a networking cabinet rented from a telecommunications company, using the existing power supply and benefiting from the direct connection to the Internet backbone. These locations would tend to be in densely populated areas, which is convenient because the system thrives on density and proximity to congested roadways.

The cost of the fixed nodes and the data center are orders of magnitude less than the cost of the thousands of mobile nodes, which comprise the bulk of the costs. The cost for custom-made mobile devices might be $500 to $1000 for small-scale (1 to 1000 node) runs. CarNet, an MIT experiment in mobile networks, uses off the shelf components: a pocket computer running the free Linux operating system for routing and control, a GPS card, and an 802.11b Ethernet card for data transmission. The hardware cost for this setup is approximately $800 per node, breaking down to around $500 for the PDA, $100 for the wireless card, and $100-200 for the GPS card (Li, 2001). For a larger scale contract (1,000 to 10,000) from a major manufacturer, at today's prices, the cost might be from $400 to $600 per node. In five years, through mass-production and efficiency improvements, the cost might decline to $100 to $300 per node. In ten years, the cost might be below $100. In any case, the cost of these nodes comprises the major portion of the initial outlay and capital expense for this system.

Also, the devices will need to be replaced and upgraded as the technology matures, representing a major, ongoing maintenance expense. The hardware and protocol standards, once established, would require expensive realignments, possibly hindering extensive rollouts initially. That is why the traffic data collection system only has an estimated five-year lifespan and a relatively small sample. This is similar to the way DSL and cable modems are replaced and upgraded as better technologies are rolled out. Various technical and design improvements to the device, such as Flash ROMs, upgradeable software, and component-based hardware upgrades, can allow for cheaper upgrade paths. Such methods would be necessary because the need to regularly replace thousands or millions of nodes would be impractical. The design would avoid duplication and replacement of working functionality (like GPS) for the many possible applications. The mobile nodes require no user interface or display, reducing their cost further. These features could be added to give users additional features, but the user can be expected to pay for the cost of these add-ons.

If each device hypothetically costs $200 in design, fabrication, and deployment, then the cost for a 2 % sample rate will be $10,000,000 for 50,000 vehicles on the Boston metropolitan region's interstate freeway system. Assuming the system lasts for five years and maintenance and other costs amount to an additional $5,000,000, the cost for five years of metropolitan-wide traffic data sampling is $15,000,000, or $3,000,000 per year. Per probe vehicle, the cost would be $60 per year, or about $5 dollars per month. This compares favorably to other commercially available mobile wireless services described in Appendix C, which can cost anywhere from $30 to $80 per month, and may not have enough capacity to handle such a volume of data. The U.S. Department of Transportation (DOT) in conjunction with a State DOT could deploy such a system as part of an ITS initiative. If participants are chosen carefully, i.e. daily commuters on

major congested routes, a complete sample can result with dense enough coverage to measure traffic volumes on congested roads. By comparison, the Singapore ERP system was built in 1995 by the Philips consortium for $116 million. This project fitted 96% of 680,000 vehicles with in-vehicle units within a year, resulting in an initial investment cost of about $178 per vehicle (Foo, 2000). These calculations assume no revenues from paying customers, no payments or subsidies to individual participants, and only the limited functionality of one-way vehicle position data collection.

4.9 Medium-Term Application: Full-Duplex Network Services

Once the viability of this system has been established, and implementation experience gained, the next step is to increase the number of participating vehicles and enable two-way data communications. As participation increases, reliability and available bandwidth will also increase, leading to the ability to provide Consumer and Business Services. These services are not necessarily related to congestion mitigation, but are a necessary intermediate step before superior congestion mitigation measures like congestion pricing can be implemented. The major difference between this system and the previous system is the enabling of two-way data transfers, from vehicle to fixed location, fixed location to vehicle, and vehicle to vehicle.

There are two ways to send data to a vehicle: by broadcasting data through conventional systems or by routing data over the MANET. A broadcast system would be similar to some first generation satellite-based Internet access services like Hughes Direct Network and StarBand. In this method, a user connects to the service by modem for uploading data and sending requests, but downloads are received on a dish from a broadcasted satellite signal. The other method is to deliver the data by routing it through participating vehicles in the MANET. Routing over the MANET will be much like routing over an IP network, except that in IP, the routing algorithm is

symmetric for upstream and downstream connections, while over the MANET, the routing

algorithm must be different for upstream and downstream. For the purpose of explaining how

the MANET downstream connection works, we will assume that data to vehicles is transmitted

only through the MANET and not by any other means. Another assumption needed is that the

data is being sent because the vehicle specifically requested it, or in other words, this is not a

broadcast or 'pushed' data stream. For example, the user requesting a specific webpage would

be an example of a 'request' while a radio station would be a 'push'.

In order to perform data transmission to a vehicle, the basic requirement is that one must

be able to find the location of the vehicle. This implies two things: the vehicle must be uniquely

identifiable, and the vehicle must make its location known to a directory service. These in turn

imply the loss of anonymity. However, the loss does not have to be complete. For example,

MIT's Grid Location Service described in the last chapter satisfies these requirements, but

reduces the likelihood that a random node can track a vehicle. In any case, to allow a vehicle to

receive data, some privacy must be lost. For example, when a cell phone is turned on to be able

to receive calls, the cell towers must keep track of the location of the phone, thereby

compromising the privacy of the individual. However, this loss of privacy is voluntary, optional,

and done only in exchange for a service. That is why a quasi-mandatory traffic data collection

system must take all steps to preserve privacy, while in this networking services system,

individuals willingly give up some degree of their privacy in exchange for the services.

Vehicles can be uniquely identified by several methods. They can have a global unique

identifier (GUID) like the MAC hardware Ethernet addresses that are unique to each Ethernet

network interface card in the world. They can also be assigned a unique identifier by a registrar

that keeps track of all assigned identifiers. Lastly, they can assign themselves a unique identifier,

in which case there would be a problem if the ID were not unique. We will assume the each vehicle has a GUID. Before discussing how data is routed to a vehicle, some characteristics of the MANET are introduced.

The MANET is typically thought of as a 'stub' network, which means that it does not permit exogenous traffic to "transit" through it, but only allows traffic originating from and/or destined for internal nodes. It is like a neighborhood street that has no outlet. Any external data destined for an internal node in the MANET must pass through one of the bridges between the MANET and the external world. The bridge or gateway will be assumed to be in a fixed location, either combined with a sink or separate from it. In a more advanced and complicated scenario, a vehicle that has some other high speed wireless service, such as an uplink to a satellite, could share that bandwidth for other vehicles and act as a gateway for the MANET to the external world.

Technically, a MANET could be implemented as a "transit" network, by allowing externally generated data to pass through it to other external sites. This would be technically challenging, but not impossible. Work is underway by IETF working groups to ensure interoperability with existing network infrastructures should such a system be implemented. With sufficient bandwidth, it would be possible to sell Internet access to private companies located along the roadways. This would be most valuable in dense cities where laying wires is difficult and other line-of-site technologies are inadequate. It would not be effective in rural areas due to the sparsity of vehicle traffic. Given current premiums for network access, the revenue hypothetically could be enough to pay for the system's administrative costs.

The transmission of data from a fixed location to a vehicle can be divided into three steps. The first step is the vehicle registering itself on the network and sending a request. The

second step is the decision of where to send the response. The last step is the actual routing of data to the vehicle. In this scenario, we assume a vehicle sends a request to a fixed location gateway device that forwards the request to the actual destination, which could be an Internet website. The destination, or Internet website, returns its response to the gateway whose job is now to route the data back to the vehicle. However, the vehicle must makes its presence known to the gateway first through a registration process.

Generically, during the registration process, a vehicle must make its position known to a directory service. In the Grid Location Service, the directory service is implemented by a subset of the vehicular nodes themselves. However, there is no reason why the directory service could not be at a fixed location. One nice thing about implementing a traffic data collection system first is that the same vehicle position data being transmitted for that purpose can be used to maintain a list of vehicles registered on the network. At each fixed location gateway, we have a computer with an Internet connection and a wireless transmitter. This computer can maintain a list of all recently collected position packets. Hence, it contains a list of vehicles within proximity and their positions. Typically, this computer will receive multiple position packets from each vehicle, and so it can keep track of where the vehicle is moving. The computer can also have a geographic information system (GIS) which stores the layout of the local roads.

By combining these two sources of information, the computer has a good idea of which way the vehicle might be heading. If the computer also knows of and can communicate with neighboring fixed location gateways, it can tell which fixed location might be the next place to 'pass on' the 'ownership' of a particular vehicle. Here, ownership signifies that the vehicle is in closest proximity of that fixed location. The computer can continually maintain and update this list with a modest amount of processing power available from an ordinary personal computer

(PC), assuming the PC is dedicated to the task and does not have to process any radio signals. This future position estimation of the vehicles enhances the reliability of return data transmission.

The gateway upon receiving the response to a request must transmit the response to the vehicle. Suppose a vehicle creates an Internet request packet and sends to the gateway through the MANET using the hopping protocol mentioned before. *When a vehicle desires to receive a request, the vehicle can broadcast its unique identifier to allow the return data, but at other times the vehicle could remain anonymous.* Once the request packet arrives, the gateway will then perform the Internet request subject to a typical timeout of 30 seconds. Assume for now that the response data is a short piece of data that fits into a single packet. If the request is successful, response data is repackaged by the gateway with some additional information for routing back to the vehicle. The additional data is as follows: the GUID of destination, the last known position of the destination, and optionally, the estimated position, and/or the specific route to take.

The gateway, using its prediction capability, may decide to forward the packet over the Internet to another gateway who 'owns' the vehicle at its predicted position. At the owning node, the data packet is sent to a single vehicle within range that is closest to the predicted location of the destination. In that vehicle, the list of neighbors is searched, and if the GUID of the destination is found, the packet is immediately relayed to that vehicle. If the destination has been reached, the vehicle unpackages and delivers the information to the requesting entity, such as a pager, location-based service request, or in-vehicle web device. If the destination is not in the list of neighbors and it is not the current host, then it is forwarded according to the GPSR routing protocol. As described earlier, this protocol forwards packets to the neighbor whose GPS

coordinates are closest to the predicted position of the destination vehicle. GPSR routes around blindspots in an efficient manner and can continue until the GUID in the packet header matches the GUID of the receiving vehicle.

A more complicated protocol might route the packet according to a specific route stored in the packet. This specific route might be a simple, ordered list of GPS coordinates that specifies the line vector path for the packet to take, or it could be a list of GUIDs of vehicles to route through. The path would be generated by the fixed location gateway computer from information about the layout of the roads. A vehicle could forward such a packet in the next direction from present along the path specified by the specific route. This second protocol would enhance reliability by following roads whose traffic levels are known. Every packet should have an originating timestamp, and all devices throughout the network should have the same TTL, say 30 seconds, for each packet, after which the packet is destroyed or assumed to be lost.

Some applications of two-way data communications are described below (Schaffnit, 1999):

- Vehicle diagnostics and owner/dealer/insurance notification

- Collision warning/avoidance: Intersection, railroad crossing

- Personal security: Identification

- Real-time Information: Traffic, News, Business Directories, Internet/Email

- Entertainment: Broadcast

- Automatic Transactions: Tolls, Purchases

- Communications: Broadcast and Two-way

- Services on Demand: Door lock, Lights on/off

Besides the consumer and business services, there is at least one possible congestion mitigation measure. There is the possibility that people could pre-program in their desired trip

and 'request' a central tracking computer to tell them the likelihood of delay given other peoples'

plans. The person could then modify his or her route dynamically during his drive at the

suggestion of the computer, which knows the likelihoods of congestion at various points.

However, this dynamic route changing is not recommended for consumers as it could add

distraction to the commute. The advice from the computer may be in the form of news

advisories like, "congestion predicted on route x," but not in the form of imperatives like,

"change to route y now." Mid-trip route planning and adjustment would increase the likelihood

of traffic accidents if users made sudden movements across highways and roadways to change

their direction. For example, the computer may reasonably say, "there is an accident up ahead

five miles blocking traffic." But it should not say, "exit this freeway now," when the exit is half

a mile away—that would be highly dangerous, especially if many vehicles on the freeway

received the same information at once. It might also be less efficient overall because routing all

the vehicles over secondary roads could be slower despite the accident delay on the freeway.

The computer could technically route vehicles to maximize the efficiency of the network, i.e. by

telling some drivers to stay on the road and some to exit based on some selection mechanism.

However, the computer will not always be right due to the unpredictable nature of some roadway

events. Therefore, such a system would not be advisable or very useful in the end. Congestion

itself is a useful message for people to change routes. As we will see in the next section, price is

an even better message.

4.10 Long-Term Application: Congestion Pricing

At some point, it is likely that almost all vehicles will be equipped with a GPS, computer,

and networking capability. This could be due to market forces, government action, or both. At

this stage, congestion pricing becomes feasible because one has the ability to do automated

transactions for *all* vehicles. When all vehicles are participating, there is no free-rider problem

or equity issue between regions. Conceptually, there are several services and key attributes

required of such a system. As before, the services include a positioning service such as GPS and

network connectivity that is uni-directional, bi-directional, and/or multicast. In terms of

attributes, to perform government mandated transactions such as tolls, the system should be close

to 100% reliable during times of operation. Again, it should be available to all vehicles to avoid

equity issues. Finally, it should be feasible in terms of cost and technical complexity. These

issues lead to two separate questions. How can one design a region-wide mobile wireless

network? And, given such a network, how can we implement congestion pricing?

Over other mobile wireless networks, the advantages of distributed wireless network for

congestion pricing are:

1) It fits better to the problem, because tolling and data collection capabilities are needed only

   where traffic is dense.

2) It also fits better because the system automatically extends to new congested areas.

3) There are smaller capital costs because there is no need for roadside gantries, satellites, cell-

   phone towers, or other major physical infrastructure.

4) We can use the system for many other purposes, such as fleet tracking, consumer services,

   and intercity data transmission.

5) The system is upgradeable simply by updating code in the device, or upgrading components

   of the device to soften the issue of large-scale upgrades.

6) Once standardized, the system can be relatively easily deployed in any city anywhere in the

   world, particularly in developing countries with limited resources and infrastructures.

Due to these factors, this kind of network is arguably better than the fixed, cellular systems for implementing ubiquitous wireless network connectivity for the purposes of congestion pricing. The comparisons with other mobile wireless networking technologies (Appendix C) suggests that this system would be cheaper to implement, especially in the long run as electronic component costs continue to decrease and physical infrastructure costs increase. Other mobile networking technologies typically rely on a cellular system, requiring hundreds or thousands of physical sites to be installed with expensive cellular relay towers. These tower installations require capital expenses, upgrading expenses, and space rental expenses. In addition, current technologies provide only meager bandwidth. Though higher bandwidths are promised, the capital expenses for these systems (3G) will generally be even higher due to smaller cell sizes. These cellular networks require expensive build-out to reach vast suburban areas and ever increasing numbers of users. By contrast, the peer-to-peer network using inexpensive and standardized in-vehicle devices extends naturally and automatically along dense roadways, improves in reliability as more participants join, and decreases in cost as component costs decrease. Having argued for the advantages of the mobile peer-to-peer network in terms of price and coverage alignment with congestion, the next step is to show how it can be used for congestion pricing.

There are a number of options and issues with congestion pricing. For all of these options, we assume that tolls are preset on a regular basis, perhaps monthly or quarterly, and the information about them is disseminated through conventional media outlets such as the newspapers and the Internet. The toll information would consist of the rates for various highways at various times, and the locations where the tolls will be charged – the toll checkpoints. The tolls would be set depending on the observed congestion patterns for the

previous month. Typically, since travel demand is a largely a function of the physical layout of the city, which does not change quickly, setting tolls in this fashion would be sufficient to impact congestion. This is the same method that is used successfully in Singapore, but is different from the L.A. FASTRAK system, which updates tolls continuously depending on prevailing conditions. The method chosen is appropriate because we assume that there is little or no modal choice for a large number of congested routes in the U.S.

To implement congestion pricing, we need to decide when, where, and how to perform the tolling transaction. There are several options here assuming GPS and network access to all vehicles. One option is to broadcast toll tables to vehicles at set intervals through a conventional broadcast medium such as radio or television. The in-vehicle device would then subtract tolls automatically when passing a toll collection point, either from a remote user account or from a cash-card inserted into the device. A variant of this option would be to transmit the toll tables to vehicles through the MANET from the roadside gateway devices. Finally, the vehicle could transmit its location data to central location continuously. The central location could then charge tolls to users by mail or by debit accounts based on internal toll tables. In the first option, the toll checkpoint locations are stored in the vehicle, while in the latter option, the toll checkpoints are stored at a central locations. This latter option has a benefit because there is no issue of updating toll tables in vehicles, or tampering of in-vehicle toll tables. However, this system would require personally identifiable position information of the vehicle to be relayed to the central government agency each time the user passed a toll checkpoint, but not necessarily at all other times.

In order to facilitate the second-tier, consumer and business, mobile applications services, individuals will require some mechanism for their vehicle to be tracked. This mechanism may

be implemented by private companies to do mobile transactions to user accounts for commercial services, or it may be put into place by the government. In either case, the individual can maintain the right to his or her privacy because personally identifiable tracking is only required when data needs to be transmitted back to the vehicle. This should allow enough flexibility to protect consumer privacy from both the government and industry. The government can take advantage of the transaction capability to solve the major public issue of congestion, and with careful design, minimize the tradeoff of privacy.

Continuous tracking by the government is very likely to be politically infeasible. Therefore, we present a system that protects privacy and does congestion pricing using the broadcast method of distributing toll tables. As mentioned before, in this mechanism, all toll checkpoints and prices are broadcast to the vehicle. When the device records that it has passed a checkpoint, it can simply generate a transaction and send that through the network to the toll collection agency. In this system, the location of the vehicle is only needed for routing of communications. The data delivered to the collection agency could only consist of the dollar amount of the toll, the time stamp, and the user's identification. In order to do enforcement, an officer would simply have to listen for the signature of the wireless transmission containing the transaction information at the toll checkpoint. If no signature is heard, then either the device has failed, the person has disabled the device, or the vehicle contains no device. In this case, either an automated photography system could identify the vehicle or enforcement could be done manually, as it is done for existing toll plazas. This system would provide a better balance between privacy and automation, but there may be an unacceptably high risk of fraud without enforcement. The collection agency would know only the time, the amount of the toll, and the identity of the user, but not the location.

In an even more private scenario, the system could be virtually the same as the Singaporean system, with transactions taking place through cash cards readers equipped in the device. In this case, toll tables must be broadcast to the vehicles, and some information must be broadcast from the vehicles to enable remote enforcement. The MANET network's utility would be reduced in this scenario. Though it could be used for broadcasting the toll tables to all the vehicles, there would be little use for the vehicle-to-fixed capabilities, besides the anonymous traffic data collection.

These options show that there is considerable flexibility in designing the system to protect privacy. There are also several technical issues with a congestion pricing system not related to privacy. For example, people may slow down or queue at a tolling checkpoint just before the toll is scheduled to decrease creating a possibly dangerous situation. This problem occurs in Singapore, and the same method they use to combat it can be used here. One could simply make it a traffic violation and use police enforcement to curb it, through video surveillance or police cruiser presence.

A trickier issue is the need for rules to prevent cutting through neighborhoods and side streets to avoid tolls. There are a finite number of routes a person can take to bypass a given toll. In addition, the number of routes is cut down by the fact that some routes will require a time and fuel cost much greater than the toll. While vehicle movements can be tracked to a greater or lesser extent (side streets may be out of range of the network), there is a question of whether or not the driver is legitimately passing through a neighborhood, or is simply on the way to work and cutting through a neighborhood to avoid a toll. To mitigate this problem, the first step is to make a list of all possible routes a vehicle may take to avoid the toll. The second step is to put an enforcement mechanism into place.

A basic mechanism would be for police to patrol common shortcuts and give tickets to repeat violators. However, this strategy would be rather costly, subject to dispute, and ineffective in some scenarios where there are many possible shortcuts. A technical solution would be to apply tolls region to region, thereby including all neighborhood-shortcut routes. In this case, there would be no incentive to take one route over another between the regions. The system could also passively or actively track vehicles and look for repeat offenders, whereupon violators can be sent letters requesting them to stop, or can be imposed a fine. This tracking could be passive by programming shortcut routes into the device and sending a signal when certain criteria have been met, whereupon it would be up to an officer to determine if a law has been broken. It may simply do to take a hands-off policy to this problem. If a neighborhood complains of a serious problem, then any one of the steps mentioned could be implemented.

More general benefits of the system are now discussed. If all vehicles are participating in this electronic road pricing system, then enacting tolls is trivial because there is no free rider problem. When the road pricing is limited to one area, that area tends to become less attractive to development, causing development to spread out. If an area becomes less attractive, there are obvious local political obstacles. However, from a metropolitan perspective, spreading out development may very well be a policy objective, even though at present, this uncontrolled sprawl is considered a serious problem.

With ubiquitous road pricing available comes the ability to control development through road pricing policy. Limitations on development can be enacted by relatively higher tolls and spurs to development can be created by relatively lower tolls. For example, if a downtown city location were extremely congested, enacting gradually increasing tolls would slowly force development to move out of the city, perhaps to city subcenters. Also, mass transit, whether

available or not, becomes a more attractive option. If the downtown city location is underused, while suburban areas are congested, relatively lower city tolls could attract development there.

One of the benefits of this system is the great deal of flexibility it affords. For example, it becomes feasible to offer credits for people who park and ride transit. One could also give credits to people who use a highway during off-peak hours. Enforcement of traffic violations could also be done automatically, though this would be more objectionable at present. However, transactions would only work under dense roadway conditions. The system will not work in very sparse traffic. This benefits the tolling for congestion pricing application, where one does not want to charge for usage of the roads during off-peak hours, but is an obstacle to providing credits at these times. In any case, flexibility of policy choices exists in dense urban areas.

With the GPS ability, one could also charge on some lanes of a roadway and not charge on other lanes. This brings up the ability to do congestion pricing where there is *less than full participation* or ubiquitous access. The incentive could be to offer diamond or fast lanes to participating drivers in exchange for installing the device and paying tolls. Such a graduated toll system may suit larger, dispersed metros where there are plenty of lanes to be divided. In this scenario, where some users pay and others do not, there is in some sense, a better allocation of equity because the people who cannot afford to pay tolls can travel on the subsidized, free highways. By providing free options to those who cannot pay, the system becomes somewhat more equitable. It is only somewhat more equitable because it will cause worse congestion for those priced out. However, this would lead to reduced traffic, the stated design goal, with its side benefits of increased traveler safety due to fewer accidents, lower pollution, and faster commutes.

In summary, congestion pricing can be implemented when this MANET-based network is nearly 100% reliable. The reliability is needed because when money is involved, there is no room for errors. The kinds of transactions that could take place can be divided into public sector and private sector, with some overlap. The public sector could collect tolls or collect fines and the private sector could offer real-time, position-based information or purchase transactions. With partial participation, one can implement some congestion pricing measures such as access to fast lanes and the ability to control congestion is limited. Privacy issues can be mitigated because identifiable location information need only be collected at toll points, or not at all when distributing toll tables to vehicles. With full participation and ubiquitous implementation, one has the best control over congestion on roadways, and by extension, development of the land-use in a city.

4.11 System Vulnerabilities

At some times, the MANET would break down and wireless connectivity would be lost to vehicles. The breakdown could be localized or regional, and would likely be temporary until the problem was solved. Examples of such breakdowns follow:

1. A failing device or unknown source emits erroneous or interfering radio signals.

2. A major solar storm or lightning storm disables or resets on-board computers.

3. Maintenance shutdowns temporarily halt data collection at fixed node gateways.

4. Hackers or vandals attack the system at one of its vulnerable points.

5. Construction activity or traffic accidents disconnect part of the network.

For each of these cases, the MANET could be eventually restored. To prevent and reduce the likelihood of outages, several steps can be taken. First of all, one can increase several MANET parameters to improve robustness, including signal range, packet buffer size, number of

fixed nodes, and even market penetration. Second, the system can be improved by careful design of the device, the security controls such as encryption, and the routing protocol. Finally, the system can be popularized by image-conscious marketing to ensure political support and funding.

At this point, the MANET system's fit with congestion pricing becomes even more apparent, because congestion pricing is not an essential, mission-critical, or life-threatening service. At most, a day, maybe a week, worth of toll revenue will be lost and congestion will return to pre-toll levels. On some days, when the haze and smog clear, the Santa Monica mountains can be seen from the city of Los Angeles, giving inhabitants a beautiful glimpse at their surrounding physiography. More importantly, it reminds them of their air pollution problem. Similarly, in the case of congestion pricing, breakdowns of the system would remind people of the congestion problem as it existed before the tolls were enacted.

CHAPTER 5: ISSUES & CONCLUSION

5.1 Summary

Automobile congestion is a serious and growing problem faced by urban areas around the world. The combined effects of space limitations and increased vehicle usage are causing congestion. Whenever there is a scarce resource to be divided amongst people, pricing is often the most effective way of equitably and efficiently distributing the benefits to those who value it most. Congestion pricing, charging drivers per use of a roadway, has been around for decades in the form of tollbooths. However, mechanical or electronic tolling has been highly unpopular in the U.S. due to high infrastructure costs and the political costs of tolling one region and not another. As such, many alternative means of controlling congestion are being implemented, including travel demand management and ITS. However, none of these methods price roadway usage directly, and therefore, they do not directly address the problem of congestion.

Singapore has had the most success with implementing congestion pricing due to its unique situation. The city-state is a prosperous island with educated citizens and serious land constraints. Nearly all vehicles in Singapore are equipped with devices from which tolls are automatically deducted when passing under toll checkpoints in the CBD and on various highways. The city-state can control congestion to an unprecedented degree, literally being able to specify the average vehicular speeds they would like to see on the roadways. The effects have been beneficial, leading to increased transit usage, cleaner air, and lower traffic levels. Because of the option of using transit, equity between tolled and untolled regions has not been seriously affected. Similarly, part of the FASTRAK system in California tolls one of two parallel highways allowing consumer choice and leading to less inequity.

The ideal system would have low infrastructure costs and be applicable to all congested areas. Such a system has become a possibility with the advent of three technologies: GPS, wireless communications, and MANET routing protocols. GPS device accuracies can reach the centimeter level at a reasonable cost. Many wireless communications standards are now available that provide bandwidth up to 10 Megabits per second and outdoor ranges from 2 to 20 miles. MANET routing protocols under development, such as GPSR and GLS, are available to route data across the peer-to-peer network. We estimate that a vehicle-mountable device with the combination of these technologies could be mass-produced today for an estimated $200 per device, with the cost decreasing significantly over time.

Such low-cost devices installed in a sample of vehicles could provide bi-directional communications capabilities with the Internet. In congested areas with sufficient numbers of equipped vehicles, the devices can relay information for each other to any desired destination. Initially with low numbers of participants, data most reliably travels one-way from vehicles to collecting servers allowing for fleet tracking and traffic data collection. In the medium-term, bi-directional communications can be implemented allowing for the provision of consumer and business services. Finally, when there is sufficient density of vehicles in order to allow nearly 100% reliable relaying of data between vehicles, one can allow financial transactions on the network. This paves the way for automatic congestion pricing based on a vehicle's position.

From traffic statistics of Boston, MA, the result of this thesis is that about two percent of vehicles in the metropolitan area must have devices in order for the initial data collection application to work for congested highways during daylight, heavy-use hours. This sample rate is estimated (not in detail) to be enough to provide both network connectivity and enough data points to algorithmically determine traffic measures. We presented a routing protocol that is

simple and effective at accomplishing the uni-directional communications. The protocol would

rely on a directional device range of several miles and a usable bandwidth of about 2 Megabits

per second. Supplementary devices in fixed locations could handle any persistent gaps in the

network. In terms of cost, we estimate approximately fifteen million dollars for the 50,000

participating vehicles, with two-thirds of the cost being for the devices themselves. Such a

system would allow collecting traffic statistics and tracking vehicle locations on congested

roadways throughout a metropolitan area. For the medium and long-term application, we

described the generic procedure for how bi-directional communications can be enabled to

provide the services such as in-vehicle Internet access. The technical issues described with these

systems become easier to manage as the number of participating vehicles increases to the point

of being able to handle financial transactions. However, the non-technical issues become more

important.

5.2 Privacy

Arguably the largest non-technical issue is privacy. Most people would strongly object to

the thought of the government tracking their movements in their personal automobiles. This fear

of government tracking is not irrational. Once detailed vehicle position information is available,

there will be a powerful incentive for people to know various other people's movements. Some

of these incentives may be acceptable to society, while some may not. Fleet tracking for

business purposes seems to be an acceptable application of the technology. Perhaps a more

questionable example, at least to youths, might be parents who wish to keep track of their

children. Most objectionable may be allowing police to automatically enforce laws or conduct

unrestrained investigations using the data.

In one way, the technology already exists through widely available cell phones, whose positions are now mandated to be traceable. In this case, people voluntarily give up their right to privacy in order to be able to receive phone calls and be locatable in an emergency. In the same way, we expect people to give up some of their privacy in order to receive Internet access or other data streams in their vehicle. As we have seen, by necessity, this functionality requires someone, somewhere to keep track of the vehicle's position, so that the vehicle may be reached.

A trade of service gained in exchange for a loss of some privacy may be more palatable to people. Indeed, in the author's conversations with peers, they showed great aversion to such a technology, but when told of possible benefits they may receive, they were willing and eager to accept it. Namely, in exchange for being tolled and to a modest degree tracked, people expected to receive the benefits of collected revenues through road improvements and tax cuts, such as elimination or reduction of vehicle registration fees, gas taxes, and other automobile-related fees. This same pattern is true in Singapore, where people tolerate the system only with an expectation of benefits. Some U.S. opinion surveys showed congestion pricing approval ratings to be only 15 to 25% for 'time-based tolling.' However, approval ratings jumped to over 60% when people were asked about their willingness to accept 'congestion pricing to pay for improved roads' (Coughlin, 2001).

People value movement information because it can be abused, just like medical histories and financial records. Legislation exists to protect of these latter types of information, and similarly, strong regulations could prevent abuse of the movement data. One rule would be to establish that a person's movement history is under the sole ownership of that person and that no entity—parent, employer, or government—has rights to that data without explicit permission of

that person. The enactment of such enforceable legislation to protect privacy would be a necessary step for quelling people's legitimate fears about their privacy risks.

However, the most persuasive argument for the privacy issue is that there are technical options and design choices available to protect people's privacy. We have presented a number of these options for congestion pricing and vehicle tracking. For non-commercial vehicle tracking, the options include voluntary participation, paid participation, and blacked-out local zones with non-identifiable position data. Individuals may also gain some safety benefit or convenience of knowing the location of their vehicle if they desired. Traffic statistics could also be aggregated in the manner of the Census in order to further hide the identities of tracked vehicles.

For congestion pricing, the options are centralized toll tables and collection versus distributed toll tables and collection. The toll tables contain the time, location, and toll amount at each toll checkpoint, and must be referenced each time a vehicle is to be tolled. In general, the centralized system is more convenient, but sacrifices more privacy. In the centralized system, vehicles would need to transmit their identifiable location continuously to the central computer, which would then reference the toll tables to automatically apply tolls to user accounts. The location data would be unnecessary after toll deduction is completed, but people would likely be uncomfortable with this situation. One tradeoff is to centralize toll collection but distribute the toll tables. In this method, toll amounts but not location data are required by the central tolling agency. This would work for those who did not wish to participate in vehicle tracking, but wanted to use tolled roads in the system. Location information would only be required internally for the routing of data. As we have seen in the example of GLS, by distributing the location service across multiple nodes, nearly complete anonymity is preserved during routing. The toll tables could be updated by broadcasting over the network, or by alternative road-side or area-

wide broadcast mechanisms. The additional complexity and overhead would be the tradeoff for the extra privacy.

These methods show the possibilities for protecting privacy of location information. The particular design choice is up to the implementers. For example, in Singapore, the first system design included centralized billing, but was scrapped due to fears over privacy. Instead, a cash card system with somewhat more privacy and accompanying inconvenience was approved. We provide options instead of a fixed privacy policy and implementation scenario, and simply state that tracking technologies are inevitably coming along with their benefits and drawbacks. Even without government intervention, private companies may implement such a system and the government may simply become another consumer of the location information (Ciccarelli, 2001). Knowledge of the possible scenarios will allow implementers to decide the best solution for their current political situation. It is worth noting, however, that the privacy issues have been more formally addressed.

Ogden (2001) gives a thorough account of the privacy issues surrounding electronic toll collection (ETC) in his article in the Transportation Research Journal. To introduce the issue, he states, "The development of intelligent transport systems (ITS), and in particular ETC, will provide massive amounts of detailed, cumulative, personal, and potentially real-time location and identification data. This raises the prospect of real-time or retrospective surveillance of the movement of vehicles and/or people and will provide the holders of such information with substantial marketing and data-matching capabilities for use in surveillance and telemarketing applications." He summarizes the privacy issues surrounding electronic toll collection into eight categories:

1. The use of electronic tags which are dependent upon assigning unique identifiers.

2. The collection of positioning data and the compilation of records related to individual travel behavior and patterns of use.

3. Enforcement issues.

4. The use of tags for surveillance.

5. Access to information by third parties.

6. Integration of disparate databases.

7. Data security.

8. Privacy policy.

Ogden then describes the ETC privacy and standards and guidelines developed for Australian ETC systems by Standards Australia, an independent non-profit organization. These privacy principles are contained in six propositions, which describe policy guidelines for dealing with the eight issues listed above. The impression from his research is that the privacy issues surrounding electronic toll collection are well understood and tameable.

5.3 Equity

With congestion pricing also comes the issue of social equity, or fairness and justice according to natural law or right. This issue applies to any system of tolling on roadways, not just the one described here. The goal of social equity, according to Webster's definition of equity, is to provide a system free of bias or favoritism. Because transportation is an emotive subject with a huge influence on our ability to function in society, we expect everyone, by natural law, to have some access to it free from bias. Our society is built upon the notion of equal opportunity for all citizens.

In terms of equity, a congestion pricing system based on a MANET would be no different from an ordinary toll road or more conventional road pricing system. Therefore, we leave the

issue to be handled generically in a more suitable forum. However, we remark that the issue of equity is separable from the way a service is priced. For example, food is normally priced according to the market, by supply and demand. In order to achieve social equity goals, we have such programs as food stamps and welfare. Thus, there is a transfer of wealth from the well-to-do who are taxed, to the poor, allowing for a greater and more equal access to opportunity. Through the public process, people have agreed to this system of trading some of their income in exchange for achieving a publicly held value. This has been done without significantly skewing the market prices of the goods themselves. Such methodology can and should be a part of any congestion pricing scheme. Singapore's congestion pricing strategy takes this into account. Similarly, the system presented here shows a way to improve the pricing of roadways in order to eliminate market inefficiencies, but does not disallow reapportionment of funds to achieve social equity goals.

5.4 Institutional

The difficulty of congestion pricing lies in the implementation details, according to Dr. Joseph Coughlin, Professor of Transportation Studies at MIT (2001). He remarks that while there has been limited success in Singapore, Israel, Ireland, and Norway, in the U.S., the history of congestion pricing has been mostly negative. For example, ISTEA 1991 devoted $125 million of funds to be used for congestion pricing projects, but over five years only $15 million was used. Many cities refused the funds on the grounds that congestion pricing in their city would make them less competitive with other cities. Also, many cities refused due to difficult technical and infrastructure barriers. He suggests that any congestion pricing system will have to overcome business and public opinion, satisfy economic and program cost constraints, have an

acceptable infrastructure venue, handle concerns over equity and privacy, have a workable institutional (public/private) framework, and have reasonable enforcement mechanisms.

Three conditions seem to exist where congestion pricing succeeds and a solution is developed to satisfy the aforementioned requirements. One, it should be a politically charged issue in that the congestion is intolerable enough and no other options have worked. Two, the benefits must outweigh the costs. Three, there must be a robust transportation alternative for those who leave their cars. Chances of successful implementation and public approval also increase dramatically if the program funds are linked to some tangible and reasonable benefit for taxpayers. This means that in marketing a congestion pricing scheme, one must never say or imply that the purpose is to change driver behavior. The advertised purpose must be to obtain funding for pressing public needs, like improved roadways or transit. The latter angle is more palatable than the former. As in Singapore, the goal of the system would be to reduce congestion, not to collect revenue. In terms of venue of implementation, political opposition will be galvanized if the scheme is focused on a single concentrated area, e.g. Long Island bridges in New York. It will be easier to overcome if people are dispersed and it is harder for them to create identifiable groups or develop a group consciousness.

Besides the public implementation hurdles, political opposition can also come from businesses who are concerned about costs to their employees. Will toll reimbursement programs become another perk demanded by employees? Will businesses have to shift work schedules? Some of the cities that took ISTEA money actually gave it back after vocal media opposition from businesses. From the trucking & shipping industry viewpoint, their benefits might exceed their costs because of increased reliability, speed, and predictability of the roadway network. However, though politically powerful, truckers are not a popular group and are heavily

taxed already. They will prefer a simpler system only to reduce their chance of being victimized, e.g. by excessive tolls on trucks. A legal hurdle might be that currently, Interstate highways can not be tolled. However, this requirement has been relaxed recently. For example, there are exemptions for high-cost bridges and tunnels, and now for rebuilt Interstates through the Transportation Efficiency Act for the 21st Century (TEA-21) (Coughlin, 2001).

5.5 Conclusion

The designers of a practical system should use the previous suggestions for implementation to answer the questions left unanswered by this thesis.

- How should funds be distributed to prevent people from being singled out?

- What are the exact incentives given to people to sign up for the system?

- Which level of government—federal, state, metropolitan, or local—is best positioned to implement this?

- What partnerships with transit agencies or private agencies can be worked out?

- What are the implications for land use in a city where roads are priced?

- How should a particular road be priced and at what level?

- How should information about pricing be disseminated to users and how should they be billed?

- How can such a system achieve political acceptance?

We answer these questions partially with the initial implementation scenario of the traffic data collection system. This system would provide traffic statistics by sampling far superior to currently collected traffic data. The statistics would be more timely, available for all heavily congested routes, and completely automated in complete constrast to current traffic data collection efforts which are usually out-of-date, fragmented, and labor intensive. The statistics

97

would allow highway planners to better understand their system and allocate resources more efficiently at a metropolital level. The $15,000,000 cost of this system over an estimated five years could be completely offset by better allocations of roadway money. Another benefit would be that police, fire, and ambulance vehicle tracking information could be provided. In fact, the system could actually generate revenue by selling the location tracking services to certain commercial fleets, such as taxi-cabs, local delivery companies, rental car companies, and institutional fleets. This location tracking service would be cheaper to implement ($5 per month per vehicle) than currently available mobile wireless data technologies (Appendix C).

Participation of commercial fleets would be supplementary to the primary purpose of the system, to collect a sample of traffic data. We assume the 2% sample does not include the participating fleets. In the random sample, some individuals may want access to their location information for security or safety purposes, and agree to be tracked in detail in exchange for this service. However, for others, privacy would be an issue. This can be handled by limiting the collected information (through blackout zones), by aggregating it, by making it unidentifiable, or by providing a small monetary incentive to the participants. In any case, the cost of the system is small compared to the benefits that could be received.

This is one reason why the initial implementation scenario is simply for traffic data collection and fleet tracking. It allows for a manageable and relatively small-scale project to test the concept of the MANET. We have suggested that the funding could be provided by the Federal government through an ITS grant or as part of a congestion mitigation grant. The system could be implemented by a State government, which could have an MPO contract the work in-house or to a private company as part of a CMS. The proof of concept provided by this system would give policy-makers experience and confidence to make further improvements and upgrade

the system to support the medium-term and long-term applications. The Federal government's

funding and organizing role will be necessary to provide standards and long-range planning to

facilitate easy upgrade paths. The private sector's role will be to provide the technology and

help the public understand it.

We have presented a staged rollout scenario in Appendix D, describing the possibilities

and applications available at the various levels of participation. With higher levels of

participation, between five and ninety percent, a wide range of consumer and business

information applications can be implemented. These possibilities arise due to increased

reliability and the enabling of bi-directional communications. These scenarios will require

strong public-private partnerships because of the high degree of overlap of benefits and

responsibilities. Once participation reaches the saturation point and reliability is virtually

guaranteed, the system can support financial transactions such as location-based tolling, allowing

one to implement region-wide congestion pricing. The issue of people taking shortcuts to avoid

tolls can be handled through enforcement similar to today's mechanisms, or by combining

enforcement with technical provisions such as region-to-region tolling or tracking of those who

take shortcuts. System vulnerabilities from vandalism and breakdowns are less severe because

tolling is not vital to roadway function, the system has inherent failsafe provisions, and several

design steps can be taken to prevent such occurrences. The wide coverage of this system

satisfies regional equity considerations because one has the ability to price any road in the

metropolitan area. Privacy can be protected by distributing to toll tables and having the devices

themselves determine when a toll is to be deducted. The incredible road pricing flexibility this

system affords allows one to exert indirect control over metropolitan land use.

In summary, while not answering every question, we have argued that a metropolitan-scale traffic data collection and congestion pricing system can technically be implemented at a competetive cost using today's technology. These technologies include wireless networking, GPS, and mobile network routing protocols. By integrating these technologies into an inexpensive device installed in vehicles, one can create a valuable and useful communications network where vehicles rely on message passing over short-range links to transfer data. The system benefits are the low physical infrastructure needs and greater flexibility over existing mobile wireless networks. The privacy and technical issues can be solved by specific design choices. The conclusion is that distributed mobile wireless networking will be a suitable technology for urban traffic congestion mitigation.

BIBLIOGRAPHY

Act 51 Study Committee, Michigan Department of Transportation (2000) *Act 51 White Paper - The Future Of Federal Highway Funding* [Online], Available: http://www.mdot.state.mi.us/act51/act51study/financial/fedfunhwy.htm [2000, Nov. 16].

Albert, Tanya (2000) *Growth flooded freeways sooner than projected*, The Detroit News, 2 Jul., [Online], Available: http://detnews.com/2000/metro/0007/02/d01-84794.htm.

Apogee Research, Inc. (1997) *Intelligent Transportation Systems – National Investment and Market Analysi*, Prepared for ITS America and U.S. Department of Transportation.

Bellur, B., Ogier, R.G. and Templin, F.L. (2001) *Topology Broadcast Based on Reverse-Path Forwarding (TBRPF)*, March, [Online] Available: http://www2.ietf.org/internet-drafts/draft-ietf-manet-tbrpf-01.txt, (work in progress).

Boston Metropolitan Planning Organization (1989) *Transportation Programs - A Guide for Local Officials in the Metropolitan Boston Region*, [Online], Available: http://ntl.bts.gov/DOCS/mpo.html [2000, Nov. 16].

Broch, J. *et al.* (1998) *A Performance Comparison of Multi-Hop Wireless Ad Hoc Network Routing Protocols* in Proceedings of the Fourth Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom October '98), ACM: Dallas, TX.

Button, K.J. and Verhoef, E.T. (1998) *Introduction* in Road Pricing, Traffic Congestion and the Environment: Issues of Efficiency and Social Feasibility, Button, K.J. and Verhoef, E.T., Edward Elgar Publishing Limited, Norhampton, MA, pp. 3-13.

Cameron, Michael W. (1994) Efficiency and Fairness on the Road - Strategies for Unsnarling traffic in Southern California, Environmental Defense Fund: Oakland, CA.

Ciccarelli, Armand (2001) An Analysis of the Impact of Wireless Technology on Public vs. Private Traffic Data Collection, Dissemination, and Use, Massachusetts Institute of Technology, M.C.P. Thesis.

Corson, S. and Macker, J. (1999) *Mobile Ad hoc Networking (MANET): Routing Protocol Performance Issues and Evaluation Considerations (RFC 2501)*, Jan., [Online], Available: http://www.ietf.org/rfc/rfc2501.txt [2001, Feb. 25].

Coughlin, Joseph (2001) Lecture on Congestion Pricing in Transportation and Environmental Policy course at MIT, 16 March.

Davis, G. and Cunningham, W. (1994) A Primer on Highway Finance, University Press of America: New York.

Federal Highway Administration (1998a) *The National ITS Architecture: A Framework for Intelligent Transportation Systems*, Publication No. FHWA-SA-98-056, [Online], Available: http://www.nawgits.com/fhwa/sheet2.pdf.

Federal Highway Administration (1998b) *Highway Statistics 1997*, Report No. FHWA-PL-98-020, [Online], Available: http://www.fhwa.dot.gov/ohim.

Foo, Tuan Seik (1997) *An effective demand management instrument in urban tranport: the Area Licensing Scheme in Singapore*, Cities, Vol 14, No. 3, pp. 155-164.

Foo, Tuan Seik (1998) *An effective demand management instrument in urban transport: the Vehicle Quota System in Singapore*, Cities, Vol. 15, No. 1, pp. 27-39.

Foo, Tuan Seik (2000) *An advanced demand management instrument in urban tranport: Electronic road pricing in Singapore*, Cities, Vol 17, No. 1, pp. 33-45.

Griffioen, Jim (2000) *The OSI Network Layers*, 11 Jan., [Online], Available: http://www.dcs.uky.edu/~griff/classes/cs471/notes/ch14a/notes/node10.html [2001, Apr. 08].

Hansen-Childers, Judy (2001) *OSI Network Layers* [Online], Available: http://www2.thomas.edu/hansenj/osi.htm [2001, Apr. 08].

Hau, Timothy D. (1998) *Congestion pricing and road investment* in Road Pricing, Traffic Congestion and the Environment: Issues of Efficiency and Social Feasibility, Button, K.J. and Verhoef, E.T., Edward Elgar Publishing Limited, Norhampton, MA, pp. 39-78.

Karp, Brad and Kung, H. T. (2000) *GPSR: Greedy perimeter stateless routing for wireless networks* in Proc. ACM/IEEE MobiCom, August, [Online], Available: http://www.eecs.harvard.edu/~karp/gpsr-mobicom2000.ps [2001, May 2].

Lee, S.-J. *et al.* (2000) *A performance comparison study of ad hoc wireless multicast protocols* in Proceedings of IEEE INFOCOM 2000, vol. 2, pp. 565-574.

Li, Jinyang *et al.* (2000) *A Scalable Location Service for Geographic Ad Hoc Routing* in the Proceedings of the 6th ACM International Conference on Mobile Computing and Networking, (MobiCom August '00), Boston, Massachusetts, pages 120-130. [Online] Available: http://www.pdos.lcs.mit.edu/papers/grid:mobicom00/ [2001, May 2].

Li, Jinyang (2001) Doctoral Candidate in Computer Science at the Massachusetts Institute of Technology 2001, Conversation with the author, 12 Apr.

Livani, Mohammad Ali (1999) *A Wireless Multicast Protocol for Real-Time Communication Between On-Road Vehicles* in Intelligent transportation systems (ITS): research and applications, SP-1467, Society of Automotive Engineers Technical Papers, 1999-01-2918.

Massachusetts Turnpike Authority (2000) *Central Artery/Tunnel Project -- Financial Plan: Section 5. Funding Status*, [Online], Available: http://www.bigdig.com/thtml/finplan/dod-fp5.htm [2000, Nov. 16].

Mergel, Joseph (1997) An Overview of Traffic Monitoring Programs in Large Urban Areas, Federal Highway Administration, Office of Highway Information Management.

OECD-Organisation for Economic Cooperation and Development (1979) Road Research: Traffic Measurement Methods for Urban and Suburban Areas, OECD: Paris.

Ogden, K.W. (2001) *Privacy issues in electronic toll collection*, Transportation Research Part C: Emerging Technology 9, pg. 123-124.

Office of Transportation (1996) *SW Corbett - Phase 2, Boundary to Nebraska, Final Project Evaluation*, City of Portland, Oregon, [Online] Available: http://www.trans.ci.portland.or.us/Traffic_Management/TrafficCalming/PDXeval s/local/corbett.htm [2001, March 18].

Puget Sound Regional Council (1994) *Evaluating Congestion Pricing Alternatives for the Puget Sound Regional Council*, Technical Paper MTP-17a, [Online], Available: http://www.bts.gov/smart/cat/ecp.html [2000, Nov. 16].

Puls, Mark (1999) *Freeway traffic explodes as Metro Detroit prospers*, The Detroit News, 21 Jan., [Online], Available: http://detnews.com/1999/metro/9901/21/01210171.htm.

Richardson, Harry W. and Bae, Chang-Hee Christine (1998) *The Equity Impacts of Road Congestion Pricing* in Road Pricing, Traffic Congestion and the Environment: Issues of Efficiency and Social Feasibility, Button, K. J. and Verhoef, E. T., Edward Elgar Publishing Limited: Norhampton, MA.

Schaffnit, Tom (1999) *Intelligent Transportation Systems – Vehicular Communications Options*, Society of Automotive Engineers, Technical Report 1999-01-2884.

Texas Transportation Institute (1992) Urban Roadway Congestion Annual Report-1982 to 1994, Texas A & M University System: College Station, Texas, Available: National Technical Information Service.

Trimble Navigation Limited (2001) *Trimble GPS Tutorial*, [Online], Available: http://www.trimble.com/gps/fsections/aa_f0.htm [2001, Apr. 23].

The Trustees of Indiana University (1996) *ISO/OSI Network Model*, 28 Jun., [Online], Available: http://www.uwsg.indiana.edu/usail/network/nfs/network_layers.html [2001, Apr. 08].

Seventh Parliament of Singapore, Select Committee on Land Transportation Policy (1990) *Report of the Select Committee on Land Transportation Policy*, Government of Singapore.

Vickrey, William S. (1969) Congestion Theory and Transport Investment, *American Economic Review*, 59:251-260.

Weiner, E. (1992) History of Urban Transportation Planning, Distributed in Cooperation with Technology Sharing Program, U.S. Department of Transportation, Washington, D.C., DOT-T-93-02, [Online], Available: http://www.bts.gov/tmip/papers/history/utp/toc.htm, p. 46-74.

Appendix A-1: Traffic Statistics in the Boston Urbanized Area

Statistic	Value	Units	Formula
Total freeway miles [1]	214	miles	totalmiles
Total estimated freeway lane miles	1304	miles	lanemiles
Total freeway DVMT [2]	21,797,000	miles/day	totalvolumeMA
Vehicles over each mile per day	101,855	vehicle/day	vehiclespermile =totalvolumeMA/totalmiles
Vehicles over each lane-mile per day	16,715	vehicle/day	vehiclesperlanemile =totalvolumeMA/lanemiles
Est.average vehicle length in feet	24.8	feet	vehiclelength=(n_autos*l_autos+n_trucks*l_trucks)/numbervehicles
Number of automobiles	1,800,000	autos	n_autos
Number of trucks	700,000	trucks	n_trucks
Total number of vehicles	2,500,000	vehicles	numbervehicles
Est. average auto length	15	feet	l_autos
Est. average truck length	50	feet	l_trucks
Est. population	2,890,000	people	population
Est. freeway speed	55	miles/hour	milesperhour
Est. average lanes per mile	6.09	lanes	lanes=lanemiles/totalmiles
All roads DVMT per capita	20.2	miles/day	-
Est. overall vehicle ownership rate	87%	percent	numbervehicles/population
Feet per mile	5280	feet/mile	feetpermile

[1] Includes Interstate and Other Freeways and Expressways

[2] DVMT = Daily Vehicle Miles Traveled

Source: Federal Highway Administration, Highway Statistics 1997, Table HM-71 and HM-72

Appendix A-2: Proxy for Hourly Traffic Volume

Time	Vehicles	% of total vehicles during day
12:00 AM	30	0.60%
1:00 AM	20	0.40%
2:00 AM	20	0.40%
3:00 AM	20	0.40%
4:00 AM	20	0.40%
5:00 AM	40	0.80%
6:00 AM	120	2.39%
7:00 AM	350	6.97%
8:00 AM	350	6.97%
9:00 AM	240	4.78%
10:00 AM	240	4.78%
11:00 AM	290	5.78%
12:00 PM	360	7.17%
1:00 PM	320	6.37%
2:00 PM	300	5.98%
3:00 PM	340	6.77%
4:00 PM	450	8.96%
5:00 PM	500	9.96%
6:00 PM	400	7.97%
7:00 PM	250	4.98%
8:00 PM	150	2.99%
9:00 PM	120	2.39%
10:00 PM	60	1.20%
11:00 PM	30	0.60%
Totals	5020	100.00%
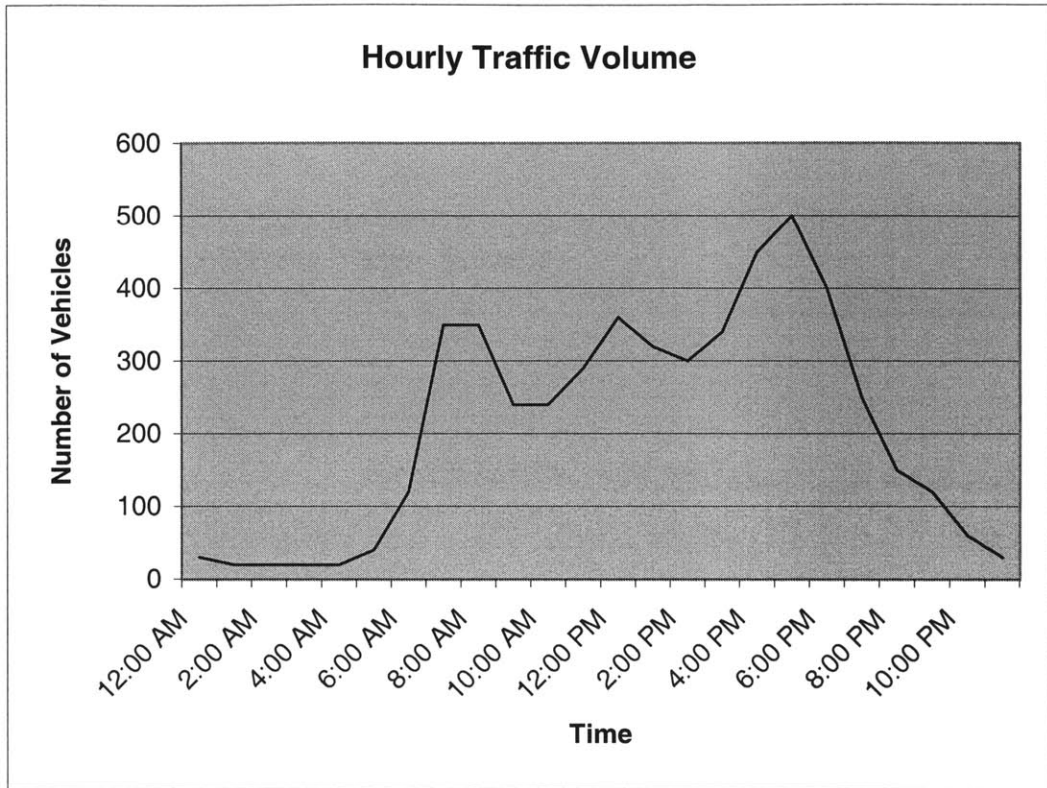
Source: Office of Transportation, 1996, Nebraska Project Evaluation

Appendix A-3: Density Calculations Based on Daily Vehicle Miles Traveled

Time	VMT/hour	Hourly traffic volume over each mile of freeway	Hourly traffic volume over each lane-mile	Vehicles per mile of freeway each hour	Vehicles per lane-mile of freeway each hour	Average feet between cars in each lane-mile	Sample rate needed to have X probe vehicles per mile of freeway			per lane-mile of freeway		
							1	5	10	1	5	10
12:00 AM	130,261	609	100	11	2	2,965	9.2%	46.0%	92.0%	56.1%	280.3%	560.6%
1:00 AM	86,841	406	67	7	1	4,482	13.8%	69.0%	138.0%	84.1%	420.4%	840.9%
2:00 AM	86,841	406	67	7	1	4,482	13.8%	69.0%	138.0%	84.1%	420.4%	840.9%
3:00 AM	86,841	406	67	7	1	4,482	13.8%	69.0%	138.0%	84.1%	420.4%	840.9%
4:00 AM	86,841	406	67	7	1	4,482	13.8%	69.0%	138.0%	84.1%	420.4%	840.9%
5:00 AM	173,681	812	133	14	2	2,212	6.9%	34.5%	69.0%	42.0%	210.2%	420.4%
6:00 AM	521,044	2,435	400	43	7	717	2.3%	11.5%	23.0%	14.0%	70.1%	140.1%
7:00 AM	1,519,711	7,101	1,165	127	21	229	0.8%	3.9%	7.9%	4.8%	24.0%	48.1%
8:00 AM	1,519,711	7,101	1,165	127	21	229	0.8%	3.9%	7.9%	4.8%	24.0%	48.1%
9:00 AM	1,042,088	4,870	799	87	14	346	1.1%	5.7%	11.5%	7.0%	35.0%	70.1%
10:00 AM	1,042,088	4,870	799	87	14	346	1.1%	5.7%	11.5%	7.0%	35.0%	70.1%
11:00 AM	1,259,189	5,884	966	105	17	282	1.0%	4.8%	9.5%	5.8%	29.0%	58.0%
12:00 PM	1,563,131	7,304	1,199	130	21	222	0.8%	3.8%	7.7%	4.7%	23.4%	46.7%
1:00 PM	1,389,450	6,493	1,066	116	19	253	0.9%	4.3%	8.6%	5.3%	26.3%	52.6%
2:00 PM	1,302,610	6,087	999	109	18	271	0.9%	4.6%	9.2%	5.6%	28.0%	56.1%
3:00 PM	1,476,291	6,899	1,132	123	20	237	0.8%	4.1%	8.1%	4.9%	24.7%	49.5%
4:00 PM	1,953,914	9,130	1,498	163	27	173	0.6%	3.1%	6.1%	3.7%	18.7%	37.4%
5:00 PM	2,171,016	10,145	1,665	181	30	153	0.6%	2.8%	5.5%	3.4%	16.8%	33.6%
6:00 PM	1,736,813	8,116	1,332	145	24	197	0.7%	3.4%	6.9%	4.2%	21.0%	42.0%
7:00 PM	1,085,508	5,072	832	91	15	331	1.1%	5.5%	11.0%	6.7%	33.6%	67.3%
8:00 PM	651,305	3,043	499	54	9	568	1.8%	9.2%	18.4%	11.2%	56.1%	112.1%
9:00 PM	521,044	2,435	400	43	7	717	2.3%	11.5%	23.0%	14.0%	70.1%	140.1%
10:00 PM	260,522	1,217	200	22	4	1,462	4.6%	23.0%	46.0%	28.0%	140.1%	280.3%
11:00 PM	130,261	609	100	11	2	2,965	9.2%	46.0%	92.0%	56.1%	280.3%	560.6%
Totals	21,797,000											

Appendix A-4: Density Calculations Based on Total Number of Vehicles

Time	Est. # of vehicles on freeways each hour	Est. vehicles per mile of freeway	Est. vehicles per lane-mile of freeway	Concatenated length (ft.) of all vehicles in each lane-mile	Remaining length in lane-mile (ft.)	Average feet between cars in each lane-mile	Sample rate needed to have X probe vehicles per mile of freeway			per lane-mile of freeway		
							1	5	10	1	5	10
12:00 AM	14,940	70	11	284	4996	436	1.4%	7.2%	14.3%	8.7%	43.6%	87.3%
1:00 AM	9,960	47	8	189	5091	666	2.1%	10.7%	21.5%	13.1%	65.5%	130.9%
2:00 AM	9,960	47	8	189	5091	666	2.1%	10.7%	21.5%	13.1%	65.5%	130.9%
3:00 AM	9,960	47	8	189	5091	666	2.1%	10.7%	21.5%	13.1%	65.5%	130.9%
4:00 AM	9,960	47	8	189	5091	666	2.1%	10.7%	21.5%	13.1%	65.5%	130.9%
5:00 AM	19,920	93	15	379	4901	321	1.1%	5.4%	10.7%	6.5%	32.7%	65.5%
6:00 AM	59,761	279	46	1137	4143	90	0.4%	1.8%	3.6%	2.2%	10.9%	21.8%
7:00 AM	174,303	814	134	3315	1965	15	0.1%	0.6%	1.2%	0.7%	3.7%	7.5%
8:00 AM	174,303	814	134	3315	1965	15	0.1%	0.6%	1.2%	0.7%	3.7%	7.5%
9:00 AM	119,522	559	92	2273	3007	33	0.2%	0.9%	1.8%	1.1%	5.5%	10.9%
10:00 AM	119,522	559	92	2273	3007	33	0.2%	0.9%	1.8%	1.1%	5.5%	10.9%
11:00 AM	144,422	675	111	2747	2533	23	0.1%	0.7%	1.5%	0.9%	4.5%	9.0%
12:00 PM	179,283	838	137	3410	1870	14	0.1%	0.6%	1.2%	0.7%	3.6%	7.3%
1:00 PM	159,363	745	122	3031	2249	18	0.1%	0.7%	1.3%	0.8%	4.1%	8.2%
2:00 PM	149,402	698	115	2841	2439	21	0.1%	0.7%	1.4%	0.9%	4.4%	8.7%
3:00 PM	169,323	791	130	3220	2060	16	0.1%	0.6%	1.3%	0.8%	3.9%	7.7%
4:00 PM	224,104	1047	172	4262	1018	6	0.1%	0.5%	1.0%	0.6%	2.9%	5.8%
5:00 PM	249,004	1164	191	4736	544	3	0.1%	0.4%	0.9%	0.5%	2.6%	5.2%
6:00 PM	199,203	931	153	3789	1491	10	0.1%	0.5%	1.1%	0.7%	3.3%	6.5%
7:00 PM	124,502	582	95	2368	2912	31	0.2%	0.9%	1.7%	1.0%	5.2%	10.5%
8:00 PM	74,701	349	57	1421	3859	67	0.3%	1.4%	2.9%	1.7%	8.7%	17.5%
9:00 PM	59,761	279	46	1137	4143	90	0.4%	1.8%	3.6%	2.2%	10.9%	21.8%
10:00 PM	29,880	140	23	568	4712	206	0.7%	3.6%	7.2%	4.4%	21.8%	43.6%
11:00 PM	14,940	70	11	284	4996	436	1.4%	7.2%	14.3%	8.7%	43.6%	87.3%
Totals	2,500,000											

Appendix B: Short Range Wireless Data Specifications

Examples for use in distributed wireless networking.

Name	Protocol	Sponsors	Range	Speed	Application	Technology	Availability
HomeRF 2.0[1]	Shared Wireless Access Protocol (SWAP)	HomeRF Working Group, Inc.	Up to 150 ft	10 Mbps	Home Networking	2.4 Ghz, up to 256 stations per network, seamless handoffs, up to 8 simultaneous active streams	Now (20 Mbps planned)
Bluetooth[2]	Bluetooth Protocol	Bluetooth SIG[3]	30 ft to 300 ft[4]	1 Mbps	Device Interoperability	2.4 Ghz.	Now
Wireless Ethernet (WiFi)[2]	IEEE 802.11b	Wireless Ethernet Compatibility Aliance (WECA)	Up to 25 mi. open air. Up to 150 ft indoor.	11 Mbps	Corporate & Home Networking	2.4 Ghz.	Now (75 Mbps planned)

References:
1. HomeRF – http://www.homerf.org, http://www.homerf.org/data/tech/techpres.pdf
2. Bluetooth Specifications Book – http://www.bluetooth.com/developer/specification/Bluetooth_11_Specifications_Book.pdf
3. Bluetooth SIG – http://www.bluetooth.com/sig/sig/sig.asp - The Bluetooth Special Interest Group (SIG) includes promoter companies 3Com, Ericsson, IBM, Intel, Lucent, Microsoft, Motorola, Nokia and Toshiba, and more than 2000 Adopter/Associate member companies.
4. Bluetooth Message Forum –
   http://www.bluetooth.com/developer/forum/forum_threads.asp?action=show&id=41&conference=4&forum=4&name=Unknown – "Bluetooth devices come in three possible power levels, two of which are commonly being implemented now. [Class 1] 100mW (+20dBm) range approx 100m, [Class 2] 1mw (0dBm) range approx 10m"

Appendix C: Mobile Wireless Data Technologies

Not requiring peer-to-peer relay.

Technology	Providers	Coverage	Speed	Consumer Cost	Technology Details	Service	Availability
Ricochet	Metricom[1]	Ricochet Cities only	28.8 Kbps – 128 Kbps	$80/month $100-$300 setup	packet-switched radio technology, cellular base stations 900 Mhz and 2.4 Ghz	internet access	Now.
Mobitex Network[3]	Cingular Wireless[2] Motient[6]	93% of US urban popln.	9600 baud[4]	$45/month $400 setup	packet-switched radio technology, cellular base stations 900 Mhz	messaging	Now.
Cellular Digital Packet Data	Bell Atlantic, Omnisky, GoAmerica	CDPD urban areas only	19.2 Kbps	$40-$60/month[5]	packet-switched radio technology, cellular base stations 900 Mhz	internet access	Now.
AMPS Cellular	Aeris.net	Nationwide	19.2 Kbps	geared to corporate customers	Data on control channel of existing AMPS cellular network, 800 Mhz[12], 1 to 20 km cells[11]	messaging	Now.
PCS Cellular	Sprint PCS and others	Limited Nationwide	14.4 Kbps	$9.95 – $179.95/month[14]	fixed network cellular system at 1.8/1.9 Ghz[12], .1 to 100 m cells[11], TDMA or CDMA	internet access	Now.
IMT-2000 ('3G' 3rd Generatiion Cellular)	International Telecommunications Union	Global	144 Kbps mobile, 2 Mbps fixed[10]	n/a	fixed network cellular system. around 2 Ghz.[9] CDMA2000	high speed internet	planned
High Altitude Stratospheric Platform (HAPS)	SkyStation (2002)[7]	Urban areas	384 Kbps – 2 Mbps for data[7]	n/a	47 Ghz. WCDMA and CDMA2000. solar powered floating platform	internet access	planned
Mobile Ad-Hoc Network (MANET)	n/a	Urban areas	up to 10 Mbps	n/a	GPS based Mobile Ad-Hoc Network (peer-to-peer relay)	internet access	n/a

References:
1. Get Ricochet – http://www.metricom.com/getricochet.htm#partners
2. Cingular Interactive – http://www.bellsouthwd.com
3. Ericsson Mobitex – http://www.ericsson.com/wireless/products/mobsys/mobitex/mobitex.shtml
4. Product Specifications – http://www.fleetcommunications.com/specs2.html
5. Bay Area Wireless Internet – http://www.cruzio.com/~jeffl/nooze/wireless.htm
6. Motient Mobile Internet – http://www.motient.com/
7. SkyStation Telecommunication Service Characteristics – http://www.skystation.com/service.html#2GHz
8. "WIRELESS" INTERNET: What the 3G Challenge Means for U.S. Competitiveness – http://www.ntia.doc.gov/ntiahome/threeg/3gintro.htm
9. ITU Frequently Asked Questions – http://www.itu.int/imt/7_faqs/index.htm
10. ITU 3rd Generation Mobile Services And Applications – http://www.itu.int/imt/what_is/3rdgen/index.html – "minimum 144 kbit/s in all radio environments and 2 Mbit/s in low-mobility and indoor environments."
11. Can PCS Dethrone Cellular and Become the Wireless King? – http://fiddle.visc.vt.edu/courses/ee4984/Projects1997/eickmeyer_davison.html
12. Cellular Networking Perspectives: AMPS Cellular – http://www.cnp-wireless.com/cellular.html
13. A Survey of Mobile Data Networks – http://www.comsoc.org/pubs/surveys/3q99issue/salkintzis.html
14. News: Sprint PCS "Wireless Internet" – http://www.computeruser.com/newstoday/99/09/21/news4.html

Notes:
'Setup' includes cost of access device and setup fees
'cost per month' is cost for an unlimited access service plan

Appendix D: Applications of the Vehicular MANET

	Traffic Data	Services	Ubiquitous
Expected Timeframe	3 – 10 years	5 – 15 years	10 – 20 years
Fraction of Probe Vehicles	1-5 % (sparse)	5 – 90% (increasingly dense)	90-100% (dense)
Reliability	90% in target areas	99% in densely populated regions	100% across metropolis
Applications	Collect sample of traffic information. Sell position information for fleet tracking. Use information to provide traveler information by conventional means. Use information to develop urban and roadway plans Role-based multicast	Internet access Car-phone Traveler information. Business route planning information. inter-vehicle information services	Tolls/Credits to curb congestion. Automated vehicular routing. Automated vehicular collision avoidance.
Technology Description	uni-directional data traffic from vehicle to internet	bi-directional data traffic, location query service	scalability, reliability, and flexibility to roadway demands
Range and Speed of internode communication	½ mile to 1 mile, 2 Mbits	100 ft to ½ mile, 11 Mbits	auto-sensing range setting, 20+ Mbits
Routing Protocol	GPSR, fixed node 'helpers', full location service not needed	GPSR with GLS	GPSR with GLS or more advanced protocol