8

Vehicular Networks: Applications, Protocols, and Testbeds

Mario Gerla and Marco Gruteser

Abstract

Vehicular networks are expected to be one of the major new application areas for wireless and Internet services. There are more than 600 million vehicles worldwide and these will be networked to achieve improvements to safety, traffic management, navigation, and user convenience. Vehicular networks (VANETs) have several elements in common with ad hoc mesh networks, but also have unique new requirements including high mobility, rapidly changing topology, multiple usage modes (vehicle-to-infrastructure [V2I] and vehicle-to-vehicle [V2V]), and the central importance of geo-location.

In the first part of this chapter, emerging VANETs are shown to be unique in the broad family of MANETs (Mobile Ad Hoc Networks). VANET services are reviewed and classified. A location-aware content distribution ("car-torrent") is then presented. Next, vehicle urban sensing is showcased for applications that range from traffic congestion/pollution measurements to distributed civilian surveillance. MobEyes, an urban surveillance application that supports forensic investigations, is then described and contrasted to other urban sensing projects.

In the second part of the chapter, the enabling VANET protocols are reviewed. First, physical and MAC layer standards for vehicular communications (DSRC, WAVE, and IEEE 802.11p) are reviewed. Then, new VANET network level protocol requirements are identified and solutions are discussed. Geo-location-based protocol architectures are introduced and briefly touch on complementary techniques such as geo-based handoff and geo-based beam adaptation for smart antennas. Security and privacy issues are addresses, with particular attention to location privacy. These protocols are illustrated with urban sensing applications.

The third part describes the role of the infrastructure in VANETs, and introduces the notion of MobiMESH, the wireless mesh architecture consisting of

https://doi.org/10.1017/CBO9780511921117.009 Published online by Cambridge University Press

roadside (Access Points) APs. Functions such as Mobility Management (e.g., Geo Location Service) are supported by MobiMESH.

In the fourth part, experimentation of VANET protocols and applications is discussed, and two emerging VANET vehicular testbeds – C-VeT and ORBIT – are reviewed.

8.1 Introduction

Vehicular communications have been receiving increasing attention over the last ten years as a viable means of augmenting road safety and travel efficiency. The field has consequently attracted consistent investments from auto manufacturers and public transport authorities, further stimulating academic research. We have reached now a situation where the essential building blocks of vehicular networks (On Board Radios, Road Side APs, Reserved 5.9 Ghz spectrum, and dedicated communication standards [Standard Specification for Telecommunications and Information Exchange between Roadside and Vehicle Systems 2003]) are (almost) available, thus opening up interesting opportunities for a wealth of car-to-car applications.

On the one side, security-oriented applications are still the top priority for auto industry and transport authorities, and recent testbed experiments have proven the effectiveness of vehicular communications in preventing intersection crashes (ElBatt et al. 2006). On the other side, the availability of the technology is stimulating interesting debates on new and challenging applications to be supported by vehicular communication systems, and visionaries are looking beyond safety applications. Automatic and efficient traffic control services (using "Intelligent Transport" techniques) can greatly benefit from vehicular communications by reducing traffic congestion, possibly keeping under control the associated chemical pollution. Imagine a comprehensive urban traffic planning system that receives inputs from vehicles (e.g., route plans, destinations, sensor readings, positions, driver's preferences, etc.), processes such information to generate an "urban routing" plan, and implements the plan through the careful control of traffic lights. The control may be extended to actual vehicle routes, possibly rerouting the vehicle to alternate, less congested routes with the assistance of "navigator" companies.

The aforementioned traffic planning system also can be equipped with entertainment-oriented functionalities providing information on locally available resources (e.g., restaurants, movie theaters, museums, etc.) and supporting content distribution, sharing, and file streaming through peer-to-peer systems (e.g., Car-torrent [Nandan et al. 2005]) and e-commerce applications, as well as mobile Internet gaming. Moreover, a new paradigm of applications arises from the observation that vehicles can actually behave as collectors (i.e., "sensors")

of information from the surrounding environment. Indeed, vehicles can be easily equipped with several sensing devices monitoring specific physical processes/phenomena (cameras, microphones, pollution sensors, humidity, temperature, etc). Such sensing devices can be used to build up a distributed and enriched awareness of the vehicular environment, which, in turn, can boost the creation of "environment-aware" applications. As an example, vehicular surveillance systems can be built to support crime investigation, homeland protection, and suspicious activities monitoring. Further, massive distributed databases can be created and maintained storing commercial, entertainment, and cultural information.

From a network architecture point of view, we argue that to support all the aforementioned applications/services, vehicle-to-vehicle communications need to be supported and integrated into *roadside infrastructure*, which in turn must provide Internet connectivity and communication resiliency. As an example, crash prevention and intelligent transport applications would not be feasible or effective if they relied only on pure car-to-car communications under sparse vehicle distributions. Similarly, content distribution (via CarTorrent, say) services most likely must retrieve the original content in the Internet, thus calling for a fixed infrastructure to bridge the vehicles to the Internet. Thus, roadside infrastructure must be ubiquitous and instantly available to support all the above functions.

Roadside APs providing the contact point between the vehicular realm and the infrastructure are to be placed in special locations, to best serve the fast-moving vehicles, as opposed to the APs designed to support pedestrians, which are generally placed in shopping malls, popular bars, restaurants, bus/train stations, and other public places. To this extent, ideal places to install the roadside APs are traffic lights and more generally light poles, overpasses, and other public structures. Traffic lights in particular are perfectly positioned to act as traffic routers: They are ubiquitously distributed throughout urban centers in precisely the locations where traffic management is most required; they are equipped with power and directly maintained by local municipalities; and they have the best "view" of approaching vehicles and crossing pedestrians. Traffic lights and other roadside access points form neighborhood *wireless meshes* that are interconnected with each other via the infrastructure. Not all the roadside AP's have wired access to the Internet, due to cost and physical limitations. The wireless mesh will provide this interconnection in a simple and cost-effective way.

Vehicular protocols and applications can be adequately evaluated and validated only in an experimental setting. Various vehicular testbeds have recently been announced, many of them offering open access to experimenters. Given the difficulty to create test environments that capture the scale of an urban grid

with millions of vehicles, there must be provision for powerful emulation platforms and rigorous validation tools that help bridge the gap between small-scale testbeds and large-scale simulators.

This chapter will introduce VANET architectures using a top-down approach. Requirements and applications are introduced first, followed by enabling protocols, supporting infrastructure functions, and testebds. The chapter is organized as follows. First, in Section 8.2, the VANET is compared and contrasted to closely related MANETs, and VANET unique properties are highlighted. Next, emerging VANET applications are described, including content delivery (CarTorrent/CodeTorrent) and "urban sensing." Section 8.3 follows, with the protocols that make such applications possible. The main focus is safety messaging/broadcast standards; mobility models/generators; routing, including emerging geolocation-based protocol architectures; DTN routing; and vehicular security and privacy. Section 8.4 identifies the role of the infrastructure and introduces the notion of a wireless mesh network and its role in support of mobility management. Section 8.5 will cover the emerging VANET testbeds (UCLA C-VeT; Rutgers Vehicle Testbed + ORBIT).

8.2 Vehicular Network and Application

8.2.1 VANET vs. MANET: What Is the Difference?

The first MANET (Mobile Ad Hoc Network) was borne about forty years ago, on the wake of the ARPANET successful debut. It was called Packet Radio Network (Kahn 1977) and was mainly viewed as a portable (at the light weight of 40 lb) radio for packet radio communications among soldiers in the battlefield. In the past forty years, the MANET has received enormous attention by wireless network researchers in academia as well as in the aerospace and military industry. Supported by steady funding from government and defense agencies, it has evolved to be an extremely sophisticated system both in radio and protocol designs. The most important application is tactical networking, followed by emergency and civilian protection scenarios. Excluding a few sensor networks (which are fixed anyway), commercial MANET applications are still in their infancy. The VANET is the prominent example of emerging MANET. In fact, it is the researchers' dream because it enables a number of exciting and compelling applications that have commercial potential. However, if researchers expect to extend mature MANET protocols to the VANET, they are going to be quickly disappointed. That is because the VANET is anything but an ordinary MANET.

To start, the conventional MANET is instantly deployable and reconfigurable in areas without infrastructure. Figure 8.1 contrasts the multi-hop, instantly deployable MANET with the wireless infrastructure network. The urban VANET

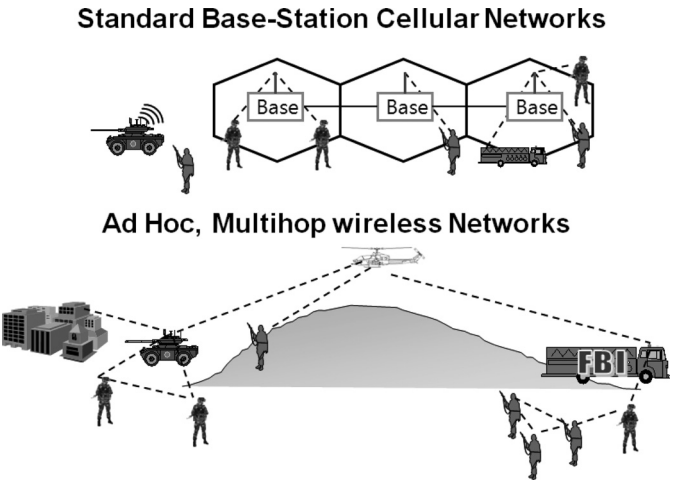

Figure 8.1. Ad hoc multihop instantly deployable MANET versus wireless infrastructure network.

is also dynamically reconfigurable; however, in normal operating conditions, it can tap one or more different "infrastructures (3G cellular, WiFi or IEEE 802.11p, WiMAX). MANETs are typically deployed to satisfy a "temporary" need (e.g., battlefield, emergency, etc) – as qualified by the term "ad hoc." VANETs run welldefined, permanent applications (like safe navigation, crash prevention, road congestion monitoring, etc.).

Mobility is the key attribute of a MANET and is characterized by a motion pattern. In battle and emergency scenarios, the motion pattern is generally not well known in advance; routing architectures are compared under various semiarbitrary assumptions, like random way-point, group motion, coordinated motion (say, follow the leader, gather/scatter, etc.) depending on the specific applications. In the VANET, there is a much better understanding of the motion pattern (say, commuting traffic to/from work; business traffic to/from train station/ airport/convention center; shopping expeditions to malls, etc.). In fact, most of the VANET architecture evaluations are based on traces or on traffic patterns that have been validated by traces. Mobility in MANETs also implies battery constraints – like in a scouting team equipped with portable radios and exploring the forest for a few days. Thus, low energy protocols are a must. In a VANET, battery power can be assumed infinite for the purpose of communications and computing.

Mobile-to-mobile multi-hop routing on dynamically changing paths has been the trademark of MANETs. In fact, in MANETs, the lead application so far has been reliable data delivery (uni or multicast) to remote destinations. Most of the challenges in MANETs design stem from designing stable routing protocols and

https://doi.org/10.1017/CBO9780511921117.009 Published online by Cambridge University Press

data transfer sessions (UDP and TCP) over such multi-hop paths. In VANETs, the delivery of data to remote destinations is not the lead application. Besides, the Internet infrastructure takes care of that. At most, data will travel a few vehicle-to-vehicle hops until a roadside AP is reached. Typical VANET applications require neighbor interactions like broadcasting alarms, P2P sharing of content, exchanging sensor information, and so on. So VANET routing is "proximity" driven rather than multi-hop to far destinations.

Multi-hop routing in VANETs still plays a role to get to a roadside AP a few hops away. More important, efficient V2V routing is required in special situations – for example, when the entire infrastructure has failed because of a disaster (e.g., Hurricane Katrina scenario), or when the infrastructure cannot be used for covert operations (e.g., homeland defense or peacekeeping operations in an unfriendly city). When V2V multi-hop routing is required, the preferred routing scheme is geographic routing, considering that virtually all vehicles will soon be equipped with GPS, and there are efficient techniques to fill in the gap in tunnels of urban canyons where the GPS signal is weak. Moreover GPS jamming is not as critical in VANET applications as it is in tactical MANET scenarios.

From the preceding discussion emerges the picture of a VANET that is quite different from the conventional MANET. In fact, vehicles will connect in most cases single-hop to the infrastructure like in a WLAN. V2V ad hoc networking will occur only it is necessary because of lack of nearby APs or applications latency constraints – say crash prevention, or emergencies and covert operations. We may describe the VANET as an Opportunistic Ad Hoc Network, where direct access to Internet (via WiFi, WiMAX, or 3G) is readily available but is opportunistically "bypassed" using the "ad hoc" if too costly or inadequate. For example, V2V is preferred for the exchange of navigation safety beacons and alarms among cars as shown in Figure 8.2. This drastic difference between the VANET and the conventional MANET is in part a loss, in the sense that it precludes the use of much of the classic MANET research generated over the last forty years. It does, however, open a tremendous opportunity of new research on this very novel environment in many areas including:

Physical and MAC layers:

- Radios (MIMO, multichannel, cognitive, SDR)
- Positioning in GPS deprived areas

Network Layer & Routing:

- Mobility models
- Network Coding

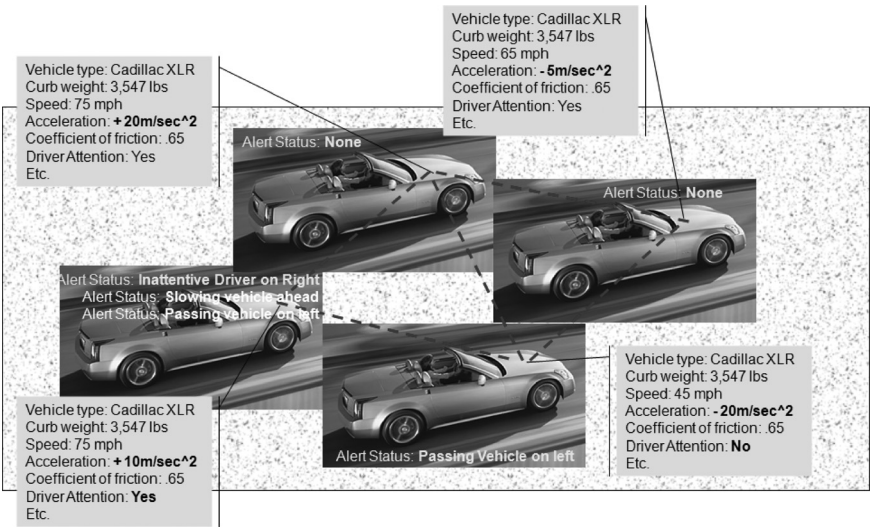

Figure 8.2. V2V navigation safety beacons and alarms exchanged among cars.

Routing:

- Geo routing, Content based routing, DTNs

Security and privacy
New Applications:

- Content distribution, mobile sensing, safety, etc.

8.2.2 Emerging Vehicular Applications

8.2.2.1 Classification and Requirements

As mentioned earlier, VANETs offer the opportunity to deploy, in addition to traditional MANET type applications, a broad range of innovative, peer-to-peer content sharing and dissemination applications. Although P2P sharing has been so far confined to the wired Internet (e.g., BitTorrent), the much increased storage and processing capacity of VANETs with respect to personal or sensor-based ad hoc networks make such applications now feasible in the mobile domain as well. Moreover, the fact that car passengers are a captive audience provides incentive for content distribution and sharing at a scale that would be unsuitable to other ad hoc network contexts. We describe a representative set of VANET P2P applications and classify them by the vehicle's role in managing data: as a data source, data consumer, source and consumer, or intermediary.

https://doi.org/10.1017/CBO9780511921117.009 Published online by Cambridge University Press

First, the vehicle is a unique *source* of data. It provides an ideal platform for mobile data gathering, especially in the context of monitoring urban environments (i.e., vehicular sensor networks) (Eriksson et al. 2008; Hull et al. 2006; Lee et al. 2006; Lee et al. 2008a; Lee et al. 2008b). Each vehicle can sense events (e.g., images from streets or the presence of toxic chemicals), process sensed data (e.g., recognizing license plates), and route messages to other vehicles (e.g., forwarding notifications to other drivers or police officers). As the vehicle removes processing power and storage space constraints, these sensors can generate and handle data at a rate not imaginable for traditional sensor networks. Vehicular sensor applications require persistent and reliable storage of data for later retrieval. Namely, they require networking protocols (including sophisticated query processing) to efficiently locate/retrieve data of interests (e.g., finding all the vehicles at a certain time and location).

Second, vehicles can be significant *consumers* of content. The on-board equipment is capable of supporting high-fidelity data retrieval and playback. For the duration of each trip, drivers and passengers make up a captive audience for large quantities of data. Examples include locality-aware information (map-based directions) and content for entertainment (streaming movies, music, and ads) (Nandan et al. 2005; Lee et al. 2006b; Nandan et al. 2006; Caliskan 2006). These applications require high network data rates and fast access to stored data.

In a third class of compelling applications, vehicles are both the *producers and consumers* of content. Examples include services that report on road conditions and accidents, traffic congestion monitoring, and emergency neighbor alerts – for example, my brakes are malfunctioning (Dikaiakos et al. 2005; Guo et al. 2005; Lee et al. 2006a; Nadeem et al. 2003; Park et al. 2006). Also, interactive applications (e.g., voice-over-V2V and online gaming) belong to this category. These applications require location-aware data gathering/dissemination and retrieval. In particular, interactive applications require real-time communication among vehicles.

Finally, all of the previously mentioned applications will need to rely on vehicles in an *intermediary role*. Individual vehicles in a mobile group setting must cooperate to improve the quality of the applicant experience for the entire network. Specifically, vehicles will provide temporary storage (caching) for others, as well as forwarding of both data and queries. In this capacity, they require reliable storage as well as efficient location of and routing to data sources and consumers.

The demands of these applications give us a list of requirements and challenges for vehicular applications.

> **Time sensitivity** – Time-sensitive data must be retrieved or disseminated to the desired location within a given time window. Failure to do so renders the data useless. This mirrors the needs of multimedia

streaming across traditional networks, and one can leverage relevant research results from the related areas.

Location awareness – Both data gathered from vehicles and data consumed by vehicles are highly location-dependent. This property has direct implications on the design of data management and security components. Data caching and indexing should focus on location as a first-order property, whereas data dissemination must be location-aware in order to maintain privacy and prevent tampering.

Most applications require methods of storing/retrieving such location/time sensitive information. As in MANETs, we can use structured approaches such as geographic hashing (Ratnasamy et al. 2002) and DHT (Caesar et al. 2006), or structureless approaches such as epidemic dissemination (Vahdat and Becker 2000). However, it is nontrivial to maintain structure in VANETs due to the high mobility, nonuniform distribution of vehicles and intermittent connectivity. Thus, most application protocols rely on variants of epidemic data dissemination such that the produced information is disseminated to nodes in an area where the information is produced (Caliskan et al. 2006; Dikaiakos et al. 2005; Lee et al. 2006a; Nadeem et al. 2003; Zhou et al. 2005).

8.2.2.2 Vehicles as Data Consumers: Content Distribution

Content distribution to vehicles ranges from multimedia files to road condition data and to updates/patches of software installed in the vehicle. Nandan et al. (2005) proposed SPAWN, a BitTorrent-like file swarming protocol in a VANET. In SPAWN, a file is divided into pieces and is uploaded into an Internet server. Each file has a unique ID (e.g., hash value of the file content), and each piece has a unique sequence number. Users passing by the APs download parts of the file. Once out of the range of APs, they cooperatively exchange missing pieces.

SPAWN is composed of the following components: peer/content discovery and peer/content selection. Due to intermittent presence of APs, SPAWN cannot use a centralized server as in BitTorrent that keeps track of all the peers. Instead, SPAWN uses a decentralized "gossiping" mechanism for peer/content discovery that leverages the broadcast medium of the wireless networks. A gossip message of a node contains a file ID, a list of pieces that the node has, a hop-count, and so on. For efficient gossiping, SPAWN uses gossiping methods, namely probabilistic spawn and rate-limited spawn. In the probabilistic spawn, nodes forward gossip messages with a certain probability, whereas in rate-limited spawn, nodes forward gossip messages in their buffer with a certain rate; for example, forwarding a random gossip message in the buffer every two seconds. The hop-count of a gossip message is incremented whenever a gossip message

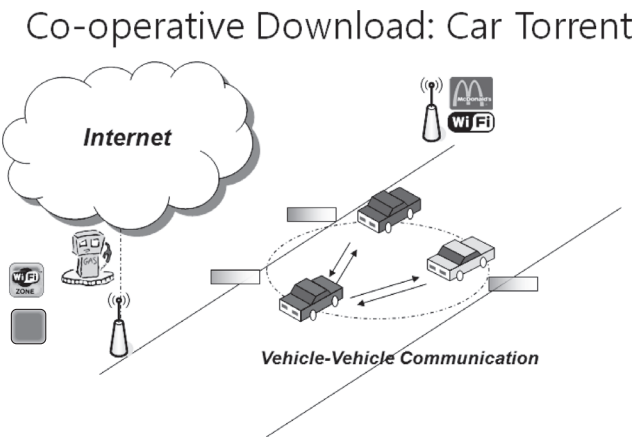

Figure 8.3. Cooperative file downloading in a VANET.

is forwarded. For a given file, there are three types of users in the network: those who are interested in downloading the files, those who are uninterested in downloading the files, and those who do not understand the SPAWN protocol. These roles are considered in the gossiping. For instance, interested users may have a higher probability of packet forwarding than uninterested users.

After the peer/content discovery, a node has to select a peer to download a piece. Given that TCP connections spanning fewer hops perform better in multi-hop wireless networks, SPAWN uses proximity-driven piece selection strategies where the proximity is estimated by the hop-count in the gossip messages: (1) Rarest-Closest First chooses the rarest piece among all the peers in one's peer list, and breaks the tie based on proximity; (2) Closest-Rarest First selects the rarest piece among all the closest peers. Recall that BitTorrent uses a rarest piece first-selection strategy where the rarest piece among all the peers in its list is selected. After peer selection, the node finally downloads pieces by setting up a TCP connection. Any routing protocols such as AODV and DSR can be used for this purpose.

By simplifying SPAWN, Lee et al. (2007) proposed CarTorrent (Figure 8.3). Given that proximity is the key factor of peer selection, CarTorrent uses k-hop limited probabilistic gossiping, and Closest-Rarest First is used for peer selection. CarTorrent uses a cross-layer approach in that route discovery of underlying on-demand protocols is utilized for gossiping. Lee et al. (2006a) proposed CodeTorrent, a network coding-based content distribution protocol. Recall that BitTorrent-like protocols suffer from a coupon collection problem – that is, as a node collects more pieces, it will take progressively longer time to collect a new piece. It is known that network coding can mitigate this problem (Gkantsidis and Rodriguez 2005; Chiu et al. 2006). Figure 8.4 shows that CodeTorrent improves

https://doi.org/10.1017/CBO9780511921117.009 Published online by Cambridge University Press

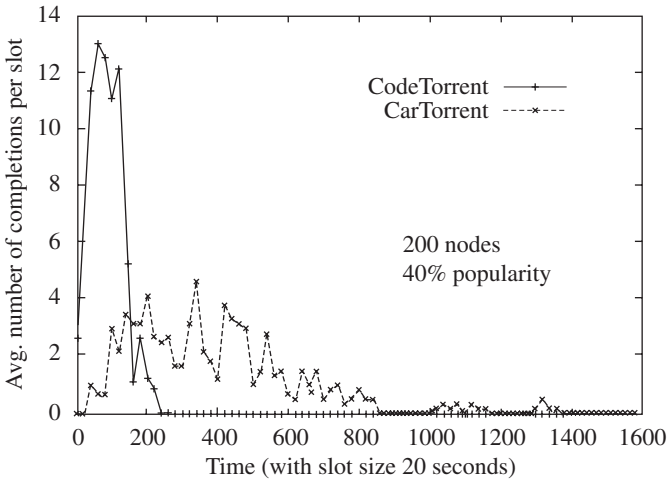

Figure 8.4. CodeTorrent improves download completion time versus CarTorrent.

download completion time versus CarTorrent by almost tenfold in a simulation experiment with 200 vehicles because it eliminates the "coupon collection" problem.

Eriksson et al. (2008a) proposed techniques to improve data delivery throughput. Quick – a streamlined WiFi client – reduces the end-to-end link establishment delay to a WiFi AP, and Cabernet Transport Protocol (CTP) improves the data throughput by differentiating congestion in wired links and packet loss in wireless links. Recently, Yoon et al. (2008) proposed Mobile Opportunistic Video-on-demand (MOVi), a mobile peer-to-peer (P2P) video-on-demand application. Since switching WiFi modes (between infrastructure and ad hoc modes) takes time, MOVi exploits the opportunistic mixed usage of roadside WiFi APs and direct P2P communications using Direct Link Service (DLS) in 802.11 standards that enables direction communications between nodes within a single BSS.

8.2.2.3 Vehicles as Data Sources: Vehicular Sensor Platforms

Vehicular networks are emerging as a new network paradigm of primary relevance, for example for proactive urban monitoring using sensors and for sharing and disseminating data of common interest. In particular, we are interested in urban sensing for effective monitoring of environmental conditions and social activities in urban areas using vehicular sensor networks (VSNs). Differently from traditional wireless sensor nodes, vehicles can easily be equipped with powerful processing units, wireless communication devices, GPS, and sensing

https://doi.org/10.1017/CBO9780511921117.009 Published online by Cambridge University Press

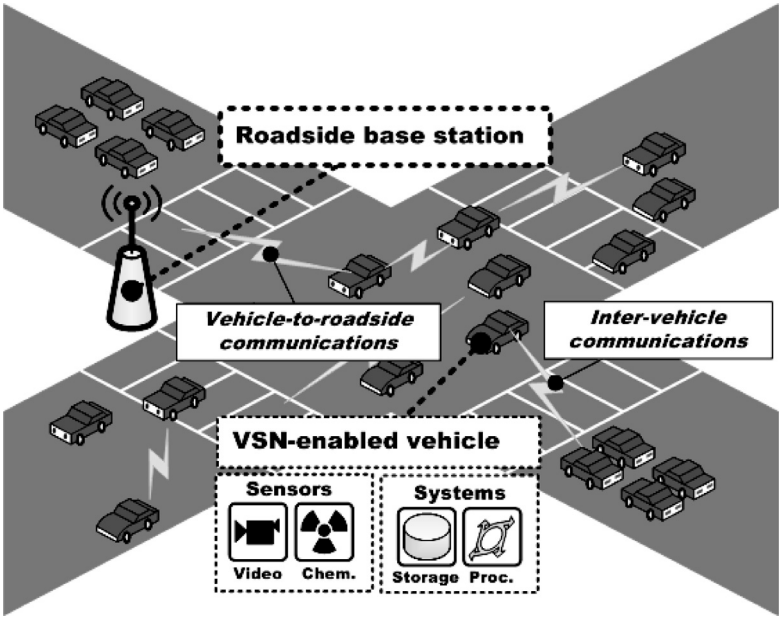

Figure 8.5. Vehicular Sensor Network (VSN).

devices such as chemical detectors, still/video cameras, and vibration/acoustic sensors. Figure 8.5 shows an application scenario.

MobEyes: Proactive Urban Monitoring Services

MobEyes aims to provide proactive urban monitoring services where vehicles continuously monitor events from urban streets, maintain sensed data in their local storage, process them (e.g., recognizing license plate numbers), and route messages to vehicles in their vicinity to achieve a common goal (e.g., to allow police agents to pursue the movements of specific cars). However, this requires the collection, storage, and retrieval of massive amounts of sensed data. In conventional sensor networks, sensed data is dispatched to "sinks" and is processed for further use (e.g., Direct Diffusion [Intanagonwiwat et al. 2000]), but that is not practical in VSNs due to the sheer size of generated data. Moreover, it is impossible to filter data a priori because it is usually unknown which data will be of use for future investigations. Thus, the challenge is to find a completely decentralized VSN solution, with low interference to other services, good scalability, and tolerance to disruption caused by mobility and attacks.

MobEyes is a novel middleware that supports VSN-based proactive urban monitoring applications (Lee et al. 2006a; Lee et al. 2008a; Lee et al. 2008b). Each sensor node performs event sensing, processing/classification of sensed data, and periodically generates data summaries with extracted features and

https://doi.org/10.1017/CBO9780511921117.009 Published online by Cambridge University Press

context information such as timestamps and positioning coordinates. Summaries are then disseminated to other regular vehicles, making it possible for patrol cars to move to the scene of an accident, say, and opportunistically harvest from neighbor vehicles the summaries relative to that accident.

Summary Diffusion: Any regular node periodically advertises a packet with newly generated summaries to its current neighbors. Each packet is uniquely identified (generator ID + locally unique sequence number). This advertisement to neighbors provides more opportunities to the agents to harvest the summaries and thus reduces the delay to collect the desired data. The advertise period is set to optimize the tradeoff between harvesting latency and data channel load.

Neighbors receiving a packet store it in their local summary database. Therefore, depending on node mobility and encounters, packets are opportunistically diffused into the network. MobEyes is usually configured to perform "passive" diffusion: Only the packet source can advertise its packets. Two different types of passive diffusion are implemented in MobEyes: single-hop passive diffusion (packet advertisements only to single-hop neighbors) and k-hop passive diffusion (advertisements travel up to k-hop as they are forwarded by j-hop neighbors with $j < k$). MobEyes can also adopt other diffusion strategies – single-hop active diffusion, for instance – where any node periodically advertises all packets (generated and received) in its local database at the expense of a greater traffic overhead. As detailed in the following section, in a usual urban VANET, it is sufficient for MobEyes to exploit the lightweight k-hop passive diffusion strategy, with very small k values, to achieve an efficient level of harvesting.

Figure 8.6 depicts the case of a VSN node C1 encountering other VSN nodes while moving (for the sake of readability, only C2 is explicitly represented).

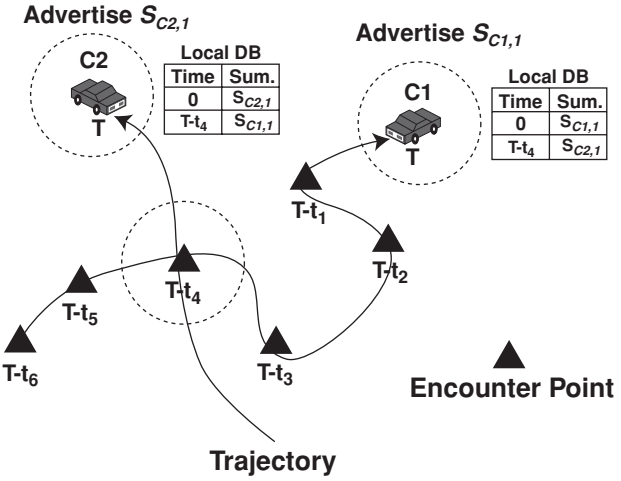

Figure 8.6. MobEyes single-hop passive diffusion.

https://doi.org/10.1017/CBO9780511921117.009 Published online by Cambridge University Press

Encounters occur when two nodes exchange summaries, that is, when they are within their radio ranges and have a new summary packet to advertise. In the figure, dotted circles and timestamped triangles represent, respectively, radio ranges and C1 encounters. In particular, the figure shows that C1 (while advertising $S_{C1,1}$) encounters C2 (advertising $S_{C2,1}$) at time T-t_4. As a result, after T-t_4, C1 includes $S_{C2,1}$ in its storage, and C2 includes $S_{C1,1}$.

Summary Harvesting: A MobEyes police agent harvests diffused summaries that meet particular criteria (typically, time and place where the data was collected) by proactively querying its neighbors. The ultimate goal is to collect all the relevant summaries generated in a given area. To focus only on missing summaries, a MobEyes agent compares its already collected set with the summary list at each neighbor (set difference problem) by exploiting a space-efficient data structure for membership checking, namely a Bloom filter. A Bloom filter for representing a set of n elements consists of m bits, initially set to 0. The filter applies k independent random hash functions h_1, . . . , h_k to MobEyes summary identifiers and records the presence of each element into the m bits by setting k corresponding bits. To check the membership of the element x, it is sufficient to verify whether all h_i(x) are set.

In summary, the MobEyes harvesting procedure consists of the following steps:

- The police agent broadcasts a "harvest" request with its Bloom filter.
- Each neighbor prepares a list of "missing" summaries from the received Bloom filter.
- One of the neighbors returns missing summaries to the agent.
- The agent sends back an acknowledgment with a piggybacked list of just-received summaries. Upon listening or overhearing this, neighbors update their missing summary lists for the agent.
- Steps 3 and 4 are repeated until there are no missing summaries.

Note that Bloom filter membership checking is probabilistic. In particular, false positives may occur and induce MobEyes regular nodes not to send summaries still missing to the agent. The probability of a false positive depends on m and n (Fan et al. 1998). Nevertheless, in MobEyes, the agent can obtain a missing summary with high probability, because it is highly probable that other nodes have the summaries as time passes, and the harvesting procedure is repeated as the agent moves. For example, in usual VSN deployment scenarios (e.g., with ten neighbors on average), we can show that the probability of missing one summary due to false positives after repeating the procedure multiple times is very low.

Data Retrieval: Harvesting leads to summaries that only contain metadata. From the metadata the agent gets the ID of the actual vehicle that owns the data. The

data must then be obtained from that vehicle. This entails finding the current vehicle location (via a *Location Server*) and routing a request (via *Geographic Routing*) to said vehicle to upload the data at its earliest convenience to the Internet. Both Geographic Routing and Location Server implementation are described in later sections of this chapter.

Related Urban Mobile Sensor Platform Projects

Recently, there have been several projects that addressed the sensing of urban data (traffic, pollution, road conditions, etc.) using mobile platforms (cell phones or vehicles). In CarTel (Hull et al. 2006), users submit their queries about sensed data on a portal hosted on the wired Internet. Then, an intermittently connected database is in charge of dispatching queries to vehicles and of receiving replies when vehicles move in the proximity of open access points to the Internet. Eriksson et al. (2008) proposed a system called Pothole Patrol that uses mobility of vehicles, opportunistically gathering data from vibration and GPS sensors, and processing the data to access road surface conditions. Yoon et al. (2007) proposed a method of identifying traffic conditions on surface streets using the GPS location traces collected from vehicles. Eisenman et al. (2006) proposed a three-tier architecture called MetroSense: Servers in the wired Internet are in charge of storing/processing data sensed by cell phones; Internet-connected stationary Sensor Access Points (SAP) act as gateways between servers and cell phone users viewed as mobile sensors (MS); MS move in the field opportunistically delegating tasks to each other and "muling" (Shah et al. 2003) data to SAP. MetroSense requires infrastructure support, including Internet-connected servers and remotely deployed SAP. Wang et al. (2006) proposed data delivery schemes in Delay/Fault-Tolerant Mobile Sensor Network (DFT-MSN) for cell phone-based pervasive information gathering. The trade-off between data delivery ratio/delay and replication overhead is mainly investigated in terms of buffer and energy resource constraints. CENS Urban Sensing project (Burke et al. 2006) addresses "participatory" sensing where cell phone-equipped agents of the same interest participate in an urban monitoring campaign. The data is uploaded to Internet servers via WiFi access points or the 3G network.

As it may have been noted, the previously mentioned urban sensing applications upload the data to the Internet at the earliest opportunity. This is because the data is of immediate need (e.g., traffic information in CarTel or pollution measurements in the CENS Participatory Sensing project). MobEyes differs from the previously described applications in that it collects indiscriminately a massive amount of information (like a security video camera in a shopping mall). Only a fraction of this information will be needed for forensic investigations in case of an accident. Thus MobEyes uploads no data or metadata to the Internet. Rather it opportunistically "disseminates" the data in the urban grid, making it easier for future investigator to search.

8.3 Enabling Protocols

Whereas vehicular networks enable a broad spectrum of applications, automotive safety is still a key motivating application that drives most of the protocol development activities. Initial safety applications (Robinson et al. 2007) such as lane change assistance (LCA), emergency electronic brake lights (EEBL) (Zang et al. 2008), and cooperative collision warning (CCW) (ElBatt et al. 2006) will steer the driver's attention through indicator light, as well as auditory and haptic signals, to warn of potentially dangerous situations. In the longer term, applications may also actively intervene – for example, by conducting automatic collision avoidance maneuvers (Ferrara and Paderno 2006) or through vehicle crash preparation that can reduce injuries for vehicle occupants.

The key protocol challenges in enabling safety applications are:

Reliability and Timeliness in Sparse and Dense Networks

Protocols must deliver messages with high reliability to other nearby vehicles over a broad range of different network scenarios. It must provide reliable message delivery under very sparse traffic conditions – say, two vehicles on a rural highway, with shadowing effects from roadside structures, and in time for vehicles to allow avoiding accidents. Consider the case where a vehicle blocking a highway after a curve sends warning messages to approaching vehicles. Stopping distances at highway speed under wet conditions can exceed 250 m. Messages must also be reliably delivered under extremely dense conditions, when the network is primarily interference limited. In dense urban areas or around major highway intersections during rush hour, hundreds of vehicles may be within the nominal communication range and can potentially interfere with a transmission.

Authenticity and Anonymity

The communication system must be able to validate that messages were generated by a trusted agent. One step toward this goal is authentication of the transmitter. If vandals can spoof safety-critical messages such as a collision warning message, the warning system itself could create enormous psychological stress on drivers and occupants or even lead to rear-end collisions due to sudden breaking. At the same time, the communication system should protect driver's and vehicle's anonymity (except perhaps under well-defined circumstances for law enforcement). Periodically emitting a radio signal with a unique identifier would enable more efficient surveillance technologies that can monitor which vehicles arrive at certain sensitive locations (e.g., hospitals, political meetings, etc.).

To address these challenges, protocol research and development is carried out both in industry standard bodies and in academic venues. This section will first review physical and MAC layer standards, then discuss protocol design for safety applications, and finally review emerging geographic protocols.

8.3.1 Regulations and Standards

There exist a large number of standard activities that cover different aspects of vehicular network communications. This section will focus primarily on the more established spectrum regulations and physical and MAC layer standards.

Spectrum for automotive dedicated short-range communications (DSRC) has been allocated in several countries around the world. In the United States, for example, the Federal Communications Commission (FCC) has reserved spectrum in the 5.9 GHz band and regulates permissible transmission powers. Even though the FCC allows an Equivalent Isotropically Radiated Power (EIRP) as high as 44.8 dBm for public safety applications, regular vehicles are limited to an EIRP of 2 W using omni-directional antennas. Still EIRPs and the 800 mW maximum antenna input power is much higher than the maximum allowed by 802.11a to enable communications ranges up to about 1,000 m under line-of-sight conditions. Considering that the frequency band of 5.9 GHz is significantly affected by shadow fading, the effective range is often much less. Figure 8.7 shows the delivery probability for different distances simulated using ns-2 Rayleigh fading channel with parameters tuned from outdoor vehicle experiments. This

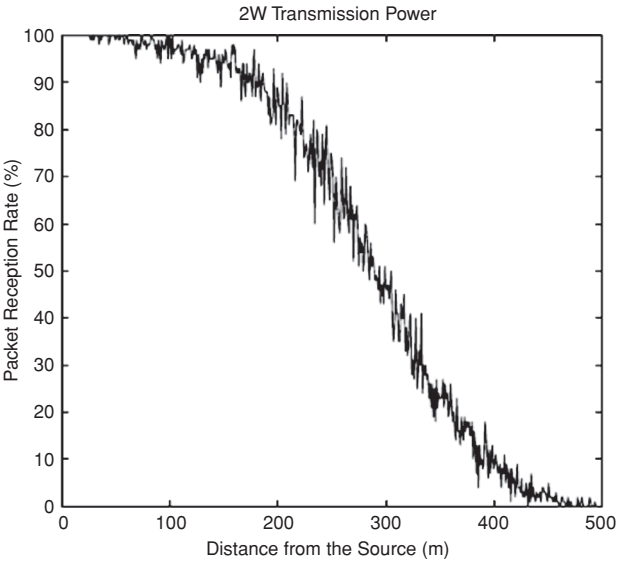

Figure 8.7. Simulated delivery probability versus distance in a Rayleigh fading channel using a 5.9 GHz 802.11p setup at 2W transmission power.

https://doi.org/10.1017/CBO9780511921117.009 Published online by Cambridge University Press

simulation shows that the effective range even at highest transmit powers can be expected to be a few hundred meters, with the range further reduced if more severe obstructions are present.

MAC and PHY Layer. Vehicular ad hoc communications are generally assumed to operate over an OFDM physical layer and CSMA/CA MAC. Such protocols are defined in the IEEE 802.11p working group (IEE 2006), which has adapted IEEE 802.11 protocols for vehicular characteristics. The physical layer remains very similar to an 802.11a OFDM PHY, except for the following changes. The 802.11p works in the band allocated for ITS applications, 5.850-5.925 GHz, allowing a total of eight channels of 10 MHz bandwidth. This differs from the 20 MHz channels in 802.11a, but some channels (i.e., channels 174,176 and 180,182) may be optionally combined to yield 20 MHz channels. The reduced channel bandwidth reduces bitrates to a maximum of 27 Mbps, but using the same number of subcarriers as in 802.11a makes 802.11p more robust to frequency selectivity of wideband channels. A higher OFDM Guard Interval of 1.6 μsec also makes 802.11p more robust to intersymbol interference caused by the high Root-Mean-Square (RMS) delay spreads that are encountered in vehicular environments 10 ns to 40 ns for vehicle separations of 10 m to 30 m and LOS, up to 400 ns in NLOS scenarios (Zang et al. 2005) as compared to 50 ns indoors. In addition, 802.11p has a longer preamble than 802.11a, allowing for better channel estimation.

The multiple-access mechanisms remain largely unchanged from 802.11a. The association mechanism has, however, been redesigned to account for the more dynamic nature of vehicular networks. To reduce the need for active scanning, 802.11p designates channel 178 as a control channel. The exact protocols are still under consideration, but it is expected that each station must periodically listen to this channel so that stations can negotiate the use of the other service channel. Frames on the control channel are always transmitted at a rate of 6 Mbps. Any prospective WAVE BSS user starts listening to the control channel for "WAVE Announcement action frames" that contain all the information required to join the BSS.

Wireless Access in Vehicular Environments (WAVE) Standards. The IEEE P1609 standard family defines an architecture and key services for vehicular networks. These include resource management, security, and multichannel operations. Finally, standards such as SAE J2735 contain a message dictionary that defines message formats for the exchange of vehicle and road information. It is typically

assumed that vehicles know their own location, for example, through Global Positioning System (GPS) receivers. Using a message defined in this standard, vehicles can then transmit their current position and past trajectory to other nearby vehicles. Messages for advertising road topology and infrastructure also exist. Whereas the standards cover many aspects of the communication protocols, other aspects of the system, particularly antennas and applications, are left to car manufacturers. On new vehicles, the communication system may be connected to the vehicle bus to in-vehicle sensors including brake, traction control sensors, and the radar and lidar sensors deployed in some new vehicle models for adaptive cruise control. New vehicles with built-in systems carry the antenna on the center rear part of the roof. Other vehicles may mount the system and antenna near the rear-view mirror on the inside windshield, similar to current electronic toll tags. This position would enable quick deployment on legacy vehicles. Some systems may interpret the message simply to provide driver warnings, whereas other vehicles may use the information to configure vehicle systems such as the braking system.

8.3.2 Broadcast Protocols for Safety Applications

To support safety applications, each vehicle must have knowledge of the surrounding vehicle constellation, which is the position, speed, acceleration, and yaw rate of other nearby vehicles (typically vehicles within a 300 m radius). Thus, a key communication primitive for vehicular safety applications is a periodic broadcast from each vehicle to disseminate vehicle movement information. By receiving these position announcements, each vehicle can then combine all received reports to create a view of the surrounding vehicle constellation. The exact use of this information then depends on the specific safety application; it may be used to issue warnings to drivers or to take precautionary actions. Current U.S. standard deliberations are considering a messaging rate of 10 Hz for each vehicle.

8.3.2.1 Scalability and Density

Eventually, the vehicular network must scale to include all motor vehicles in the country. Although it may take many years of deployment efforts to reach this goal, it is worthwhile to consider how the technology could scale to such large and dense network scenarios early on to avoid costly recalls at a later time. According to the 2004 Highway Statistics (Federal Highway Administration 2004), there are about 240 million registered motor vehicles in the United States. The number of vehicles is similar to the number of phones supported by cell

Figure 8.8. Example regions with potentially high vehicle densities: (a) Junction of Freeway 110 and 105 (b) Intersection in New York City.

phone systems, but the network challenges are fundamentally different in that vehicles must dynamically organize themselves into local networks and allocate spectrum resources rather than relying on a carefully planned base station setup for coordination.

When deployed to a large number of vehicles, the system must meet its reliability requirements even in very-high-node density environments, which can be expected in rush-hour traffic on highways or in urban centers. Example regions that may encounter high densities are depicted in Figure 8.8.

The exact specifications are still under deliberation, but the Wireless Access in Vehicular Environment (WAVE) (IEE) and SAE standard groups are currently defining wherein each vehicle disseminates its position and vehicular dynamics information via periodic broadcasts. The broadcast rate is application-specific, but it is generally assumed to be one message every 100 ms per vehicle. The message size is typically less than 100 bytes but can reach larger sizes due to a large authentication header. In some cases, total message sizes can reach 500 bytes if optional payload information like path histories is included. Without the security overhead and optional information but considering MAC protocol overhead, this yields a typical data rate requirement of about 5 kbit/s/vehicle. With security overhead, it becomes approximately 10 kbit/s/vehicle. At first glance, this appears to be a very modest data rate requirement, but vehicular networks are interference-limited because the node densities that can be expected in vehicular scenarios make meeting this requirement challenging. For example, consider a congested, slow-moving, two-way highway with four lanes each and one car every 10 m. This results in 480 cars within a 600 m interference range of a 300 m transmission, thus requiring approximately 5 Mbps capacity if transmissions can be perfectly scheduled but more than 10 Mbps with the currently envisioned less efficient CSMA-based protocols. Unfortunately, current state-of-the-art DSRC technology cannot provide this capacity for the required communication range

because radios have to operate at lower channel bandwidths of 10 MHz to reduce multipath delay spreads and Doppler effects in the vehicular environment.

These interference limitations motivate data aggregation approaches that can reduce bandwidth requirements while still achieving similar communication ranges.

8.3.3 Emerging Geo-Protocols

Vehicular networks are intricately linked to the physical world, and their applications require that each vehicle be able to monitor its position. These characteristics have also lead to a number of proposals that use geographic position information to improve network performance, rather than just dissemination vehicle positions over a location-agnostic network stack. One class of such protocols is geocasting – the delivery of messages to all nodes within a defined area. There exists a natural match to the typical requirement to disseminate vehicular movement information to all vehicles within a radius of say 300 m. Particularly when data aggregation and multi-hop forwarding are used, message propagation is no longer limited by a single vehicle's transmission range and requires other mechanisms to prevent flooding of the network. Establishing a geographic boundary for message propagation can fill this need.

In a multi-hop message forwarding scenario, geocast protocols could further increase network efficiency if safety applications can define smaller message delivery zones based on map information or recent vehicle trajectories. For example, consider the extended electronic brake light scenario illustrated in Figure 8.9. Here a brake message should be reliably delivered to all following

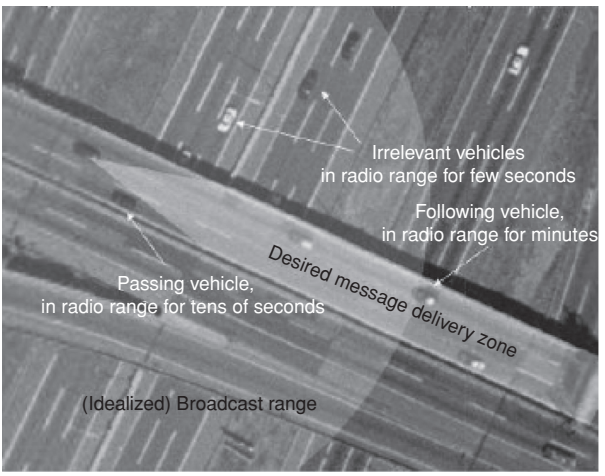

Figure 8.9. Use of persistent geocast in vehicular networks.

https://doi.org/10.1017/CBO9780511921117.009 Published online by Cambridge University Press

vehicles, where some vehicles might be out of communication range of the origin vehicle. Thus, some of the following vehicles must forward the message. The message is, however, only relevant to vehicles following on the same road; vehicles on the intersecting highway do not need to be notified. In this scenario, the intended message recipients are frequently changing, thus enumerating destination identifiers for each vehicle is cumbersome. A naive flooding approach with a time-to-live parameter might propagate in unintended directions, causing unnecessary network congestion. A geocast services provides a natural way to identify the destination vehicles through a geographic perimeter and can optimize message forwarding to only include nodes in the intended destination zone.

The concept of a geocast has first been proposed by Navas and Imielinsky of WINLAB in 1997 (Navas and Imielinsky 1997). For such highly mobile vehicular applications, this concept requires adaptation for ad hoc communication and persistence to notify new cars in the area. Whereas work in the MANET and sensor network field such as Mobicast (Mobile Just-in-time Multicasting) (Huang et al. 2002) has addressed multicasting in mobile ad hoc networks, protocols tailored specifically to vehicular networks have been proposed. For example, Maihöfer (2005) proposed *abiding geocast* that extends the earlier geocast models by including the notion of message validity duration. The abiding geocast protocol will not only deliver the message to all nodes present in the geocast region during the initial message transmission, but also continue to deliver the message to all vehicles that enter the geocast zone during the message validity duration. It also considers infrastructure-assisted geocast.

Another form of geographic protocols, *georouting*, can also find application in vehicular networks, for example, to transmit a message to a roadside unit that cannot be reached in a single hop. Rather than using the knowledge of logical link level associations to find a path from the source to the destination, as in typical topology-based routing schemes (Perkins and Royer 1999; Johnson and Maltz 1996), geographic routing finds the path by using location information of the destination and potential forwarding nodes. A node chooses as next-hop forwarder the neighboring node that is closest in geographic space to the destination node. The main advantage of location-aware routing is that it does not require route establishment and maintenance, which can be costly in highly dynamic vehicular networks. It does require, however, that the location of the destination node is known, which is easier to achieve for stationary roadside infrastructure nodes compared to mobile nodes.

Resiliency to mobility and channel variations in vehicular networks can be further improved through opportunistic protocol techniques. Standard routing protocols (including georouting) select a next-hop neighbor based on their routing metric and instruct the MAC layer to unicast a packet to this selected node.

This appears wasteful in dense wireless networks, because this particular destination may be unreachable due to fast or slow fading, whereas there are likely neighboring nodes that correctly receive the frame immediately but discard it due to the incorrect MAC address. A realization of an opportunistic protocol could use a soft destination, where any node close to these coordinates can forward the packet. This approach takes advantage of the additional resources in a dense network, and conceptually it realizes a form of cooperative diversity gain that recent works in the information theory community (e.g., Laneman and Wornell 2003; Laneman et al. 2004; Nostratinia et al. 2004) have shown to provide large gains in networks with idle nodes, or in a slow, fading environment. Ignoring all protocol overheads and assuming a Rayleigh channel, best-case gain estimates follow those for selection combining, which predicts a 12 dB diversity Signal-to-Noise Ratio (SNR) gain using two receivers (over one receiver) for an outage probability of 0.01 (Goldsmith 2005). Adding further receivers yields diminishing returns, but a third and fourth receiver still provides an additional 7 dB and 4 dB gain, respectively. By realizing similar diversity gains, cooperative protocols can operate at lower transmission power and thus increase spatial reuse. The key challenge in such cooperative protocols lies in low-overhead distributed forwarder selection algorithms. One approach is to allow multiple forwarders to contend through a backoff mechanism that skews the probability of channel access to forwarders in closer proximity of the destination (Kaul et al. 2008).

8.3.4 Security

Another key challenge in vehicular network protocol design is providing authenticated communication while maintaining anonymity or pseudonymity of the vehicles. Let us first consider the authentication mechanisms. To authenticate messages, the current standard for security services in DSRC/WAVE considers digital signature-based authentication primitives. Keys and certificates for vehicles could be issued during the vehicle registration process. Using elliptic curve cryptography, the overhead of these signatures on a packet amounts to a manageable few tens of bytes. The public keys for verifying messages could be distributed through certificates appended to the messages, which would increase message length noticeably, or could be periodically broadcast by each node.

Security of Aggregated Messages. Aggregated messages, however, require additional protection because a spoofed or faulty message can misrepresent sensor information from a large number of vehicles. Consider again the position-monitoring application that collects the

location of nearby vehicles. Let us assume that a regular record contains the current location and speed of the vehicle, a timestamp, and a signature with certificate to authenticate the message. An aggregator could then combine multiple records syntacticly: By listing the vehicles' positions in a single record authenticated with a single signature and certificate. It could also aggregate information semantically, for example, by describing a bounding box that contains all vehicle positions. Again, to reduce overhead, the aggregated message ideally would contain only a single signature and certificate. Because authentication only establishes the source but not the correctness of the content, a regular message may contain incorrect or spoofed positions information for a vehicle. Simultaneously inserting a large number of spoofed vehicles would require having the same number of valid keys available. An aggregated message, however, could contain an arbitrary number (subject only to packet size constraints) of position claims using only a single key.

One approach to address this issue is probabilistic validation. A receiver can probabilistically verify the correct aggregation by requesting the original record for a randomly chosen identifier contained in the aggregated message. Full validation means obtaining the complete set of original records, checking their signatures, and confirming that applying the aggregation function to this yields the aggregated message. To reduce the bandwidth and resource consumption of this validation process, the receiver can use the probabilistic method. The randomly chosen identifier acts as a random challenge, so that the sender does not know beforehand which original message will be verified. If the same vehicle repeatedly does not supply a valid original message for the requested record, the receiver can assume that this record is spoofed. Note that this validation method can only catch spoofing of additional vehicles, not omission of existing vehicles. Omission of vehicles can, however, be more easily addressed by the nodes surrounding the aggregator. If they overhear an aggregated message that undercounts vehicles in the area, they can send a corrected aggregated message to add the additional vehicles (assuming the majority of vehicles are trustworthy). Removing spoofed vehicles would require collaboration between multiple nearby nodes, because no single node can be sure that the additional vehicles do not exist outside its radio range. Note also that probabilistic validation is most effective if a method for recourse exists, which penalizes the sender of the message. For example, the receiver could notify authorities of the suspected tampering with the vehicle communication system, who can track down repeated offenders.

Temporary Keys. Signatures and certificates are essentially pseudo-identifiers because a receiver knows that two messages originate

https://doi.org/10.1017/CBO9780511921117.009 Published online by Cambridge University Press

from the same sender if both signatures can be verified with the same public key. Since DSRC communications are broadcast over the wireless medium other vehicles and any unauthorized parties can record such pseudo-identifiers from vehicles in the vicinity. By monitoring the identifiers at multiple locations, third parties could calculate a vehicles' average speed, or they could monitor the identifiers of vehicles visiting a sensitive location (e.g., medical clinic). Over time, these bits of information can create a profile that identifies the driver. To address these privacy concerns, inter-vehicle communication protocols should ideally be free of such static pseudo-identifiers but still provide basic authentication functions.

One approach to provide privacy while authenticating is switching among a large number of temporary keys. Because storage is relatively affordable, each vehicle could store a large number, say ten thousand, of certificates that can be used. Used keys could be replenished during vehicle maintenance or be remotely updated over a wide-area connection. It is critical that only one of these certificates is valid at any given time to prevent spoofing of other vehicles. The degree of privacy can be increased by using a higher switching frequency. More frequent key changes, however, make it more difficult to implement the safety applications that rely on tracking paths of nearby vehicles. It also creates a tension with the secure aggregation approach, because messages from a suspicious node cannot be filtered when it switches to a new key. A good solution must balance all these requirements.

8.4 The Role of the Infrastructure: MobiMESH and GLS

One of the unique features of the VANET is the omnipresence of the infrastructure. In fact, the wired infrastructure is accessed through a thin *wireless mesh* layer. It is thus important to understand the interaction and interdependence between vehicular networks, wireless mesh, and Internet. To start, VANET applications benefit from the support of Internet and wireless mesh network in many services:

Mobility: The infrastructure manages mobility. Mobility management (e.g., knowing where vehicle X is at time T) requires the registration with a location server (centralized or distributed) that is built in the infrastructure. The location server accepts registrations of participating vehicles as they roam the city and maintains the equivalent of a DNS mapping vehicle IDs to current estimated geolocation and AP to reach the mobile. Moreover, the infrastructure must facilitate AP to AP soft-session handoff of roaming mobiles. This is provided for both TCP sessions and stream and is critical for real-time applications (voice, videoconference, and interactive games).

Security and Authentication: An important overarching concern in VANETs is security and privacy. The infrastructure helps "authenticate" the alarms received from other vehicles – filtering bogus attacks, for example. It also helps preserve "location" privacy considering that vehicles potentially jeopardize such privacy by exchanging beacons, advertisements, and warnings with other drivers. Security and privacy guarantees require a certifying authority residing in the infrastructure.

Routing: The shortest path between two vehicles may go through the infrastructure. Using in part information supplied by the infrastructure (e.g., city map, vehicles density in various sectors, urban WiFi channel load, etc.), each vehicle can determine whether it is better to route a packet totally within the VANET or partly through the wired infrastructure (Gerla et al. 2006). Namely, the "Data Routing Advisory" is analogous of the Navigator Advisory for data packets.

Urban measurement repository: The infrastructure can serve as storage of various vehicle measurements ranging from traffic, pollution, and mobility pattern all the way to individual vehicle traces and bogus alarm attack reports. In particular, it keeps records of the data collected during VANET experiments.

Emergency operations: As drivers become progressively dependent on VANET services, such services should be maintained even when the infrastructure partially or totally fails. Critical VANET protocols (routing, capacity estimation, location service, resource allocation, and security management) must be carefully designed so as to allow a reliable transition from full Internet support to completely autonomous ad hoc operations in case of infrastructure facility failure or destruction. This is particularly important because in such situations, the VANET will be the only "infrastructure" available for emergency services such as vehicle evacuation and search-and-rescue team networking. It will offer an important backup to Public Emergency Networks like TETRA. The roadside wireless mesh will play an important role. APs powered by solar generators will use cognitive radio capabilities to reestablish a fixed, wireless emergency backbone throughout the affected urban area.

The services described here are best illustrated by describing the functionalities of the vehicular mesh, called MobiMESH (Capone et al. 2006), that is being installed at UCLA as part of the C-VeT testbed. The following properties make MobiMESH particularly suited for C-VeT support:

- *Broadband Backhauling* – the MobiMESH networks are able to build up and dynamically maintain a broadband wireless backbone that can be used to support/complement vehicle-to-vehicle communications;

https://doi.org/10.1017/CBO9780511921117.009 Published online by Cambridge University Press

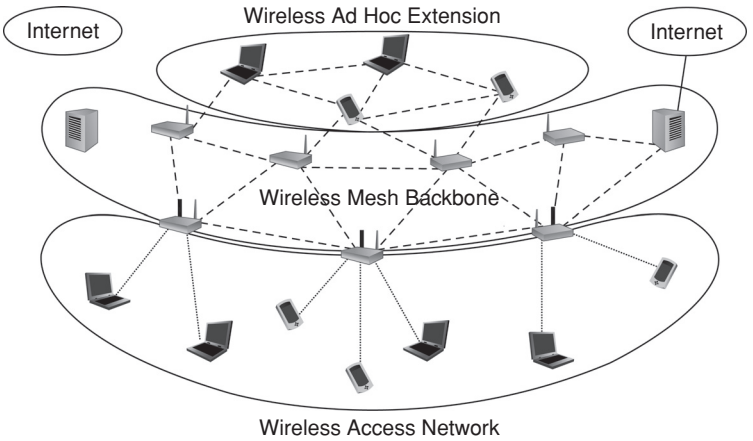

Figure 8.10. MobiMESH network architecture.

- *Mobility Support* – wireless devices are allowed to seamlessly roam within MobiMESH networks without losing active connections;
- *Flexibility* – the MobiMESH networks are self-configuring and self-managing.

A concise technical description of MobiMESH follows.

8.4.1 The MobiMESH Architecture

MobiMESH features a hybrid mesh network architecture. Indeed, the network consists of three main architectural building blocks shown in Figure 8.10:

- a mesh backbone composed of MobiMESH wireless mesh routers that provide the routing and mobility management infrastructure, and is further connected to gateways;
- an ad hoc extension responsible for extending MobiMESH functionalities to mobile nodes;
- an access network that can be used by standard WiFi clients to get connectivity.

The mesh backbone and the ad hoc extension are based on the ad hoc network paradigm, where all nodes and mesh routers collaborate to route traffic. Routing on the mesh is provided through a proactive ad hoc routing protocol based on OLSR (Clausen and Jacquet 2003) and properly modified to account for multiple radios at the mesh nodes, and for varying link quality metrics. The backbone network is also responsible for the integration with the wired network, through

https://doi.org/10.1017/CBO9780511921117.009 Published online by Cambridge University Press

gateways equipped with a wired interface that can route traffic to the Internet. The access network is rather flexible and operates in the infrastructure mode, so that standard clients perceive the network as a standard WLAN and behave accordingly; in this way, MobiMESH can also be accessed by standard WLAN clients (e.g., pedestrians) with no specific software installed.

The MobiMESH Mesh Routers represent the main building block of the MobiMESH network because they are responsible for creating the broadband backhaul, further offering access to wireless mobile clients. The MobiMESH Mesh Routers can be equipped with two to four radio interfaces that can be flexibly used either as backbone or access interfaces. Moreover, any interface can be tuned to any available channel in the two frequency bands 2.4 GHz, 5.7 GHz, and 5.9 GHz (via DSRC). Mesh Routers with an interface dedicated to WiFi access are called Access Routers.

An important overarching concern in C-VeT environment is security and privacy. The MobiMESH network provides security functions, so that it can be safely employed to deliver any kind of traffic and to extend preexisting secure networks. In a Wireless Mesh Network (WMN), it is very important that only authorized devices can join the network; MobiMESH Mesh Routers are in fact authenticated through the use of X.509 certificates, and the backbone traffic is encrypted through a time-changing key encryption algorithm. Moreover, centralized MAC filtering and captive portal functionalities are supported.

MobiMESH architecture implements a proprietary mobility support daemon that dynamically handles the MAC-IP address association as clients roam throughout the network. Experiments carried out on real deployments have shown that the handover latency for a wireless client changing Access Router is upper bounded by 20 ms in most of the cases. Consequently, the handover is not perceived during VoIP calls.

In the following section, as an example of Infrastructure Service, we describe the Geo-Location Service (GLS) targeted for implementation in C-VeT.

8.4.2 The Geo-Location Service (GLS)

The Geo-Location Service (GLS) is a distributed service that maps any car ID to its most recent geo location. Exploiting MobiMESH, we propose an Overlay Location Service (OLS) implementation. As shown in Figure 8.11, an overlay structure is established in MobiMESH. Periodically (say, every minute) each car registers to the nearest MobiMESH APs with its ID (license#, IP address(es), time, owner name, owner IP address billing address, etc.) and the current geo-location. In normal operating conditions, OLS spans both the MobiMESH and the wired Infrastructure. In case of infrastructure failure, OLS can be completely supported (with some loss in performance) by MobiMESH, assuming the latter

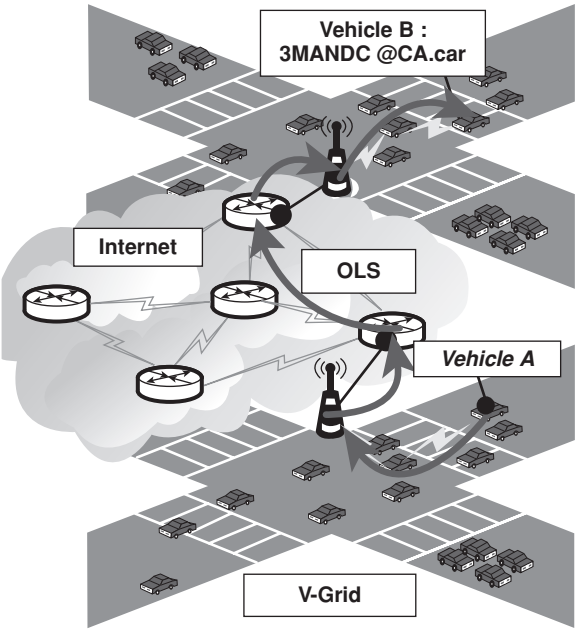

Figure 8.11. Location service and routing in the MobiMESH.

is fully connected by virtue of long-range Cognitive Radio links. OLS maintains an index of vehicle IDs. Each ID is mapped to the most recent geo coordinates (thus allowing motion prediction). The index is distributed across the overlay. It may be managed via DHT (Distributed Hash Table).

To illustrate the OLS operation, suppose that mobile host A wants to establish a TCP connection to mobile host B (see Figure 8.11). Host A injects in the nearest MobiMESH AP the query: 3MANDC@CA. It gets back the "most recent" set of time-tagged geo-locations of host B. From these, it can estimate vehicle speed and direction and thus infer the current location of B. A then selects the best AP to reach the destination. Host A encapsulates the message in an IPv6 network envelope with destination geo address in the extended header. The destination AP geo-routes the packet into the vehicular network to B using geo address, car ID, and MAC in the header. Upon successful delivery, car B responds with its own IP address and geo address. It directs its response (encapsulated in the overlay envelope) to the sender IP address.

8.5 Vehicular Testbeds

The primary goal of the vehicle testbed is to enable V2V and V2I experiments aimed at the evaluation of VANET protocols and applications in a realistic

https://doi.org/10.1017/CBO9780511921117.009 Published online by Cambridge University Press

setting. It must allow external users to define, execute, and monitor various experiments. It must allocate resources so that users can efficiently share the testbed. It must assist the experimenters with software tools such as traffic generators, measurement collection, and preprocessing facilities, and possible interface to emulators and simulators.

In addition, the vehicle testbed must interact and interwork with the infrastructure so that the applications being tested can benefit from the various services of the latter. In particular, it manages coexistence of car-to-car 802.11p channel with WiFi-based mesh infrastructure; it interfaces with the infrastructure for support in mobility management, routing, traffic control, and congestion control; it facilitates transparent interconnection of vehicles across the city via the wired Internet; it enables the VANET to operate with and without infrastructure support with smooth transition between the two modes and phasing out of noncritical applications.

In this section, we present two vehicular testbed implementations: the UCLA C-VeT testbed and the ORBIT based Rutgers testbed.

8.5.1 C-VeT Architecture

C-VeT is an open platform that supports vehicular network and urban sensing research and related applications. It is inspired to the pioneering work done by Larry Peterson and Tom Anderson with Planet Lab (Peterson et al. 2002). It features an always-on, fully virtualized, Internet-accessible, sensor-equipped testbed infrastructure. The UCLA campus, with its 10 acres of urban development, reproduces many of the scenarios, propagation, and communication challenges typical of a city, in a realistic manner but yet relatively small-scale. In particular, the C-VeT architecture provides:

- A fully virtualized platform that runs both Linux-based and Windows-based operating system with full insulation among the guest virtual machines, and enables the users to redesign low-level protocols such as, for instance, MAC protocols. This feature will be key for network centric experiments.
- A Campus Wide Mesh network developed using OPEN WRT and optimized for the integration and support of the vehicular network. It will help cope with network disruptions (quite common in small-scale testbeds) and enable opportunistic, interactive, and delay-tolerant experiments that exploit the infrastructure.
- 30 facility management vehicles equipped with the C-VeT hardware/software, providing an always-on platform to run experiments and collect traces and measurements. The facility management vehicles perform both routine maintenance trips and on-demand interventions in response to

emergencies resulting in a varied mobility pattern that well approximates real city traffic.

- 30 commuting vans, equipped with the C-VeT-Census platform that will survey the environment gathering traffic and air quality, and stereoscopic images. The aim is to build a large micropollution database that enables new models and also facilitates visual environment surveys (see Figure 8.1).
- A number of downloadable, preconfigured virtual appliances to allow users to develop the protocols to be tested at home with a compatible software configuration.
- A large-scale emulator that will allow users to debug their algorithms and protocols on the same hardware as the actual C-VeT nodes but with an emulated network component developed with the Qualnet hybrid simulation.
- A robust Internet interface that will manage the users and deploy the experiments in a streamlined fashion. The Web server will provide the front-end for a number of user-friendly services and tools enabling users to focus on research rather than testbed implementation. For example, services to set up the experiments and gather the data; APIs to low-lever interfaces for hardware component virtualization; virtual MadWiFi layer for the support of virtual machines.
- The ability to develop algorithms, applications, and protocols that directly operate at Layer 2 using a TUN/TAP mechanism for both Windows and Linux OS. Recent research showed that the TCP/IP suite may not be the most appropriate choice for vehicular networks and a ground-up protocol stack redesign is needed.
- An organized live database of mobility traces, sensed environmental data, road traffic information, Vehicle CanBus statistics, MAC layer statistics (through MAD WiFi) and physical layer statistics taken using a variety of radios (Cognitive Radios, MIMO, etc.). This data collection will be made available to the research community in collaboration with existing trace collection programs and archives such as CRAWDAD (Kotz and Henderson 2005).

The testbed was designed using a top-down approach; the whole system can be described through a number of relatively simple building blocks: the C-VeT mobile node, the C-VeT mesh node, the C-VeT-Census platform, the Web-based control center, and the emulation platform.

The C-VeT infrastructure is designed to provide an always-on facility for research in wireless vehicular network. To achieve this goal, we chose to install our equipment in the UCLA campus facility management and van pool vehicles.

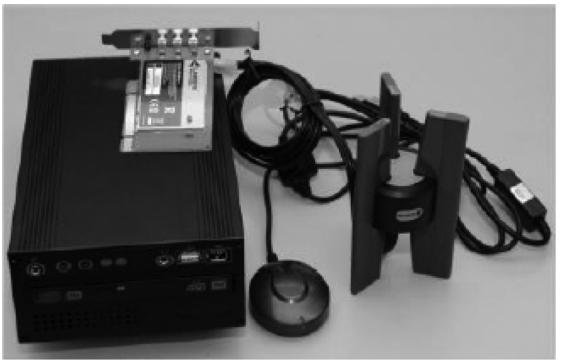

Figure 8.12. C-VeT mobile node.

Those cars and vans are driven everyday to fulfill the campus needs and perform both routine and nonroutine tasks. In addition to the permanent facility vehicles, there is a small pool of private vehicles equipped with C-VeT nodes that can be driven by the researchers themselves for customized, controlled experiments.

The C-VeT mobile node (Figure 8.12) is an industrial-strength Cappucino PC powered by an Intel Dual Core Duo processor at 2.5GhZ, 2GB of RAM, and 320GB of disk. Hard drive and internal parts are rugged to sustain physical stress (i.e., large temperature fluctuations, vibrations, etc.). The PC has three wireless interfaces: IEEE802.11a/b/g/n based on the Atheros AR9160 chipset; IEEE802.11p interface based on a Daimler-Benz customized chipset; and a standard Bluetooth interface mostly for internal communications.

Other radios can also be retrofitted in the mobile node platform. In particular, a few vehicles may be equipped with programmable Silvus SC2000 MIMO platforms (4x4 configuration) that provide full access to the physical layer and enable a new generation of experimental MAC layer research.

On Board Sensors: The C-VeT nodes are instrumented with a customized sensor platform designed to provide a flexible data collection. This includes Infrared-based CO_2 sensors; electrochemical CO sensors; SIRF III or Ublox-based GPS sensors; temperature, and humidity sensors; and a megapixel camera. Using the C-VeT cars as mobile air quality sensors will enable a new wave of atmospheric research aimed at the use of mobile sensing agents to study the air quality at the neighborhood level. Part of the fleet will feature high-performance exhaust particulate sensors DC2000CE by Echocem [ECO], thus being the first testbed able to support the currently leading research in microclimate air quality.

The C-VeT mesh node is based on MobiMESH hardware. C-VeT mesh nodes feature Open WRT OS and Atheros Chipset with MadWiFi support, thus easing up the integration with mobile nodes. The fixed infrastructure will be installed on

Figure 8.13. C-VeT infrastructure.

the roof tops of UCLA buildings aiming at full campus coverage and integration with the existing campus WiFi infrastructure. The mesh allows opportunistic Internet access from vehicles and also provides a control channel to the vehicles. The mesh network can be configured via the Web; e.g., customized routes can be set up by the network operator to perform particular experiments. This C-VeT integrated approach with infrastructure and vehicles broadens the experimental scenarios. In the initial phase, we will cover the south campus, and creating an initial backbone of six mesh points. The initial campus coverage map is shown in Figure 8.13.

To achieve seamless integration between the CVET-Mesh and the Vehicular network components, we will develop Layer 3 and Layer 2 routing and VLAN support. Level 3 network layer routing between moving vehicle and the fixed nodes will enable communications across campus and to the Internet. The Layer 2 routing will enable the experimenter to force mobiles to be in the same broadcast domain, ignoring the fact that there are several fixed nodes in between.

8.5.1.1 Testbed Deployment and Preliminary Results

Infrastructure Nodes Coverage

To find the best placement of the infrastructure nodes, we ran a campaign of coverage tests around the UCLA campus. The main focus is on the coverage of the roads. This represents a hard challenge because we experienced that the WiFi radio signal basically propagates only in Line of Sight (LOS). To assess the coverage of a single infrastructure node, we equipped a car with a laptop, a GPS receiver, and a IEEE802.11b/g wireless card. The car node would log every second its position and if it is in reach of the infrastructure node or not. Using this

https://doi.org/10.1017/CBO9780511921117.009 Published online by Cambridge University Press

Figure 8.14. Coverage experiment from the Ashe Center Building at UCLA.

information, we were able to plot the coverage map of each single infrastructure node. Figure 8.14 shows the coverage map for the infrastructure node placed on the top of the Ashe Center Building at UCLA. White dots represent the covered locations and red dots the unreachable ones. The results show that we were able to cover the whole area called Westwood Plaza that extends up to 700.

Video Streaming

As a preliminary experiment, we wanted to test the feasibility of a video transfer from a mobile node to an infrastructure node via the wireless mesh. The mesh consists of four nodes on the four corners of Engineer IV building at UCLA. In this configuration, each node could reach only the two nodes that are next to it. This means that to reach the farther node, two hops are required, as shown in Figure 8.15.We placed a webcam in the moving car and used VLC to stream

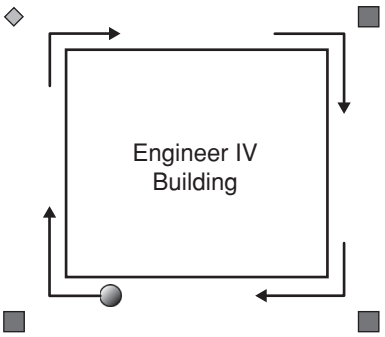

Figure 8.15. Video Streaming experiment: 1 moving video source (CSircle), 3 fixed nodes providing connectivity (Squares), and 1 fixed receiving node (Diamond) meters away from the infrastructure node. On the other hand, as soon as we lose the LOS, the connection breaks, as evidenced by traces on one of the crossing roads.

https://doi.org/10.1017/CBO9780511921117.009 Published online by Cambridge University Press

the video to one of the fixed nodes. With this setup, the car is always connected to the mesh and at most two hops away from the receiving node. To maintain connectivity and fresh routes, we used the OLSR (Clausen 2003) implementation provided by INRIA. The webcam was generating a video stream at resolution of 176 × 144 pixels at 15 frames per second. Thus the stream was generating an average of 128 Kbps (since the codec used was DIV3 the bitrate was not constant due to dynamic compression). The video was streamed using UDP, so the lost frames were not retransmitted. The VLC server was set with a cache of 200 ms.

On the receiving node, we were both saving and displaying the video. In the real-time video transfer, the missing frames were much more than 10 percent, but because we were saving the raw data received from the source, we were able to reconstruct and re-encode the video received. In Figure 8.16, we show the loss rate for the video after the reconstruction. As shown in Figure 8.16, the percentage of loss for both frames and blocks is approximately 10 percent. Such a loss still grants the possibility of actually displaying the video. For real-time delivery, the reconstruction buffer cannot be used. Forward error correction schemes and adaptive coding rate may be used in this case. Another important result of this experiment was the time when the frame losses occurred. In fact they occurred when the mobile node was swapping from one relay to another. This means that the refresh of the route is not fast enough to be transparent for the video stream. These experiments were useful to determine the impact of wireless mesh multihopping on real time traffic. Clearly, buffers and coding strategies must be properly matched to the topology and user requirements.

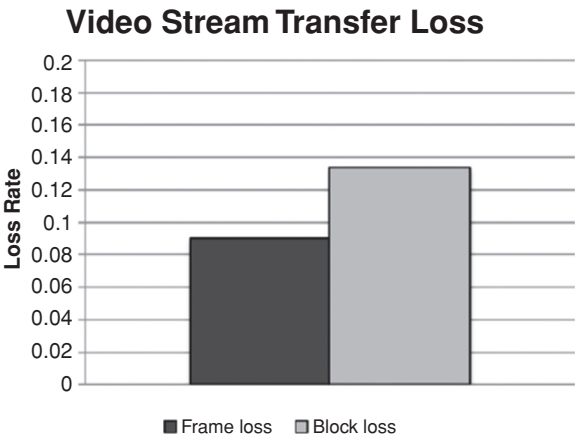

Figure 8.16. Loss rate for the video stream.

https://doi.org/10.1017/CBO9780511921117.009 Published online by Cambridge University Press

8.5.2 ORBIT Vehicular Testbed

ORBIT Indoor Testbed. The ORBIT laboratory testbed (Figure 8.17) comprises 800 IEEE 802.11a/b/g devices attached to 400 nodes in a 20-by-20 meter space that provides a controlled environment to generate reproducible results.

Mobile Outdoor Testbed. The ORBIT testbed also includes a vehicular outdoor field trial component. It comprises several building-mounted 802.11 base stations, vehicular nodes, programmable smart phones, and 3G data accounts for experimental purposes provided by a campus cellular network operator. Stationary nodes are deployed at five different locations, with ten nodes close to the ORBIT facility at the NJ Tech Center and three locations in Rutgers University Busch Campus. All outdoor nodes are connected through back-end Internet links using Ethernet tunnels to each other and the ORBIT control facility. The back-end interface can be used for experiment control, remote data collection, and to allow configuration of different network topologies.

As shown in Figure 8.17, the vehicular nodes use the same base node platform as used in the indoor ORBIT testbed to enable seamless moving of software by copying disk images between the testbeds. Every node is a custom-designed small form factor PC with 1GHz Via C3 CPU, 512 MB RAM, and 20 GB hard disk with remote management interface. The nodes include two IEEE 802.11 a/b/g interfaces whose PHY and MAC layers are similar to the ones defined in the DSRC/WAVE standards. For positioning, Garmin 18 5 Hz Global Positioning System receivers are used to obtain position updates at higher frequency (standard receivers provide only 1Hz samples, during which a vehicle can move

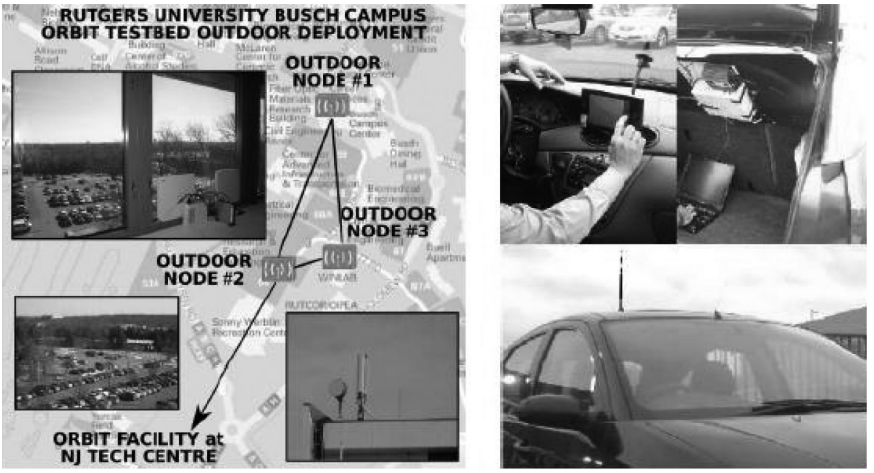

Figure 8.17. The ORBIT wireless research testbed: (a) Campus outdoor setup (b) Vehicular setup.

https://doi.org/10.1017/CBO9780511921117.009 Published online by Cambridge University Press

up to 30 m). The cars use magnetic mount omnidirectional external antennas for 2.4/5GHz. A 12-to-120V power inverter that serves as the power supply (via the car battery) and the setup includes optional keyboard and 7in LCD for experiment control.

Also available in the outdoor testbed are 10 Nokia N95 smart phones powered by an ARM11-based Texas Instruments OMAP2420 running at 330MHz. It is equipped with 64MB RAM, 160MB internal memory, and a flash memory that can be expanded up to 8GB. Short-range communication options include wireless LAN (802.11 b/g) and Bluetooth 2.0 EDR. The N95 also includes a built-in GPS receiver based on TI's GPS5300 NaviLinkô 4.0 single-chip solution for GPS and A-GPS. The Nokia N95 runs Symbian OS v9.2 and is programmable using C++, Java J2ME (MIDP 2.0, CLDC 1.1), and various scripting languages.

8.6 Conclusions

In this chapter, we have surveyed the emerging VANET applications, ranging from vehicular sensors to entertainment. We have contrasted VANET to traditional MANET design, identifying the unique VANET features and requirements. Given these unique features, we have proceeded to classify a representative set of VANET applications based on the vehicle's role in managing data: as source, consumer, source/consumer, or intermediary. We have then reported a vehicular sensing application – MobEyes – and a content distribution application – CarTorrent.

We have then introduced the protocol suite that makes such applications possible. The main focus was on routing and on emerging geolocation-based protocol architectures; on delay-tolerant routing; and on security and privacy.

We then identified the critical role of the infrastructure in the deployment of VANET applications; we introduced the notion a wireless mesh network and its role in support of mobility management.

Finally, we introduced the VANET testbeds that are being deployed at UCLA (C-VeT) and Rutgers (ORBIT-based Vehicular Testbed). We also reported preliminary experiments with live video uploads to an Internet client via a four-node mesh network.

The future of VANET research is bright. There are a number of compelling applications ready to be deployed, and users are eager to try them out. The protocols and the standards are nicely coming into place. The remaining roadblocks in VANET deployment and broad adoption are liability, privacy, and penetration. However, even these roadblocks will soon be removed. The liability is restricted to only a small class of applications (such as intersection crash prevention); moreover, rapid progress is being made in that area. Privacy issues have been practically resolved in two ways: by virtue of technology advances, and by the

fact that users are getting accustomed to give up privacy for other benefits. Full penetration (say, of DSRC radios) is no longer critical for the deployment of many applications (such as navigation and Intelligent Transport) that are increasingly relying on 3G, WiFi, and WiMAX.

References

Burgess, J., Gallagher, B., Jensen, D., and Levine, B. N. 2006. MaxProp: Routing for Vehicle-Based Disruption-Tolerant Networks. *IEEE INFOCOM*.

Burke, J., Estrin, D., Hansen, M., A. Parker, A., Ramanathan, N., Reddy, S., and Srivastava, M. B. 2006. Participatory Sensing. *ACM WSW*.

Caesar, M., Castro1, M., Nightingale, E. B., O'Shea, G., and Rowstron, A. 2006. Virtual Ring Routing: Network Routing Inspired by DHTs. *SIGCOMM'06*.

Caliskan, M., Graupner, D., and Mauve, M. 2006. Decentralized Discovery of Free Parking Places. *ACM VANET, Los Angeles*.

Capone, A., Napoli, S., and Pollastro, A. 2006. Mobimesh: An Experimental Platform for Wireless Mesh Networks with Mobility Support. *Proc. of ACM QShine 2006 Workshop on Wireless Mesh: Moving Towards Applications.*

Chiu, D. M., Yeung, R. W., Huang, J., and Fan, B. 2006. Can Network Coding Help in P2P Networks? *NetCod'06*.

Clausen, T., and Jacquet, P. 2003. Optimized Link State Routing Protocol (OLSR). *RFC3626*.

Dikaiakos, M. D., Iqbal, S., Nadeem, T., and Iftode, L. 2005. VITP: An Information Transfer Protocol for Vehicular Computing. *ACM VANET*.

Eisenman, S. B., Ahn, G.-S., Lane, N. D., Miluzzo, E., Peterson, R. A., and Campbell, A. T. 2006. MetroSense Project: People-Centric Sensing at Scale. *ACM WSW*.

ElBatt, T., Goel, S. K., Holland, G., Krishnan, H., and Parikh, J. 2006. Cooperative Collision Warning Using Dedicated Short Range Wireless Communications. *Proceedings of the 3rd International Workshop on Vehicular Ad hoc Networks*, pages 1–9.

Eriksson, J., Balakrishnan, H., and Madden, S. 2008a. Cabernet: A Content Delivery Network for Moving Vehicles. *Technical Report TR-2008-003, MIT-CSAIL.*

Eriksson, J., Girod, L., Hull, B., Newton, R., Balakrishnan, H., and Madden, S. 2008b. The Pothole Patrol: Using a Mobile Sensor Network for Road Surface Monitoring. *MobiSys'08*.

Fan, L., Cao, P., and Almeida, J. 1998. Summary Cache: A Scalable Wide-Area Web Cache Sharing Protocol. *ACM SIGCOMM*.

Federal Highway Administration, U.S. Department of Transportation. 2004. *Highway statistics 2004*. http://www.fhwa.dot.gov/policy/ohim/hs04/htm/mv1.htm

Ferrara, A., and Paderno, J. 2006. Application of Switching Control for Automatic Pre-Crash Collision Avoidance in Cars. *Nonlinear Dynamics*, 46, 307–321.

Gerla, M., Zhou, B., Lee, Y.-Z., Soldo, F., Lee, U., and Marfia, G. 2006. Vehicular Grid Communications: The Role of the Internet Infrastructure. *WICON'06*.

Gibbons, P. B., Karp, B., Ke, Y., Nath, S., and Seshan, S. 2003. IrisNet: An Architecture for a Worldwide Sensor Web. *IEEE Pervasive Computing*, 2(4): 22–33.

Gkantsidis, C., and Rodriguez, P. 2005. Network Coding for Large Scale Content Distribution. *INFOCOM'05*.

Goldsmith, A. 2005. *Wireless Communications*. Cambridge University Press.

Guo, M., Ammar, M. H., and Zegura, E. W. 2005. V3: A Vehicle-to-Vehicle Live Video Streaming Architecture. *PerCom'05*.

Huang, Q., Lu, C., and Roman, G. 2002. Mobicast: Just-in-Time Multicast for Sensor Networks under Spatiotemporal Constraints. *Proc. of the 2nd International Workshop on Information Processing in Sensor Networks*, pages 442–457.

Hull, B., Bychkovsky, V., Chen, K., Goraczko, M., Miu, A., Shih, E., Zhang, Y, Balakrishnan, H., and Madden, S. 2006. CarTel: A Distributed Mobile Sensor Computing System. *ACM SenSys*.

IEEE 1609 – Family of Standards for Wireless Access in Vehicular Environments (WAVE).

IEEE. 2006. "Draft Amendment to Standard for Information technology – Telecommunications and information exchange between systems – Local and metropolitan networks – specific requirements – Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications: Amendment 3: Wireless Access in Vehicular Environments (WAVE).

Intanagonwiwat, C., Govindan, R., and Estrin, D. 2000. Directed Diffusion: A Scalable and Robust Communication Paradigm for Sensor Networks. *ACM MOBICOM*.

Johnson, D. B., and Maltz, D. A. 1996. Dynamic Source Routing in Ad Hoc Wireless Networks. *Mobile Computing*, 153–181.

Juang, P., Oki, H., Wang, Y., Martonosi, M., Peh, L.-S., and Rubenstein, D. 2002. Energy-Efficient Computing for Wildlife Tracking: Design Tradeoffs and Early Experiences with ZebraNet. *ACM ASPLOS-X*.

Kahn, R. 1977. The Organization of Computer Resources into a Packet Radio Network. *IEEE Transactions on Communications*, 25(1): 169–178.

Kaul, S., Gruteser, M., Onishi, R., Vuyyuru, R., and T.I.T. Center. 2008. GeoMAC: Geo-Backoff Based Co-operative MAC for V2V networks. *IEEE International Conference on Vehicular Electronics and Safety*, pages 334–339.

Kotz, D., and T. Henderson T. 2005. "Crawdad: A Community Resource for Archiving Wireless Data at Dartmouth. *IEEE Pervasive Computing*, 4(4), 12–14.

Laneman, J., Tse, D., and Wornell, G. 2004. Cooperative Diversity in Wireless Networks: Efficient Protocols and Outage Behavior. *IEEE Transactions on Information Theory*, 50, 3062–3080.

Laneman, J., and Wornell, G. 2003. Distributed Space-Time-Coded Protocols for Exploiting Cooperative Diversity in Wireless Networks. *IEEE Transactions on Information Theory*, 49, 2415–2425.

Lee, K. C., Lee, S.-H., Cheung, R., Lee, U., and Gerla, M. 2007. First Experience with CarTorrent in a Real Vehicular Ad Hoc Network Testbed. *MOVE'07*.

Lee, U., Magistretti, E., Gerla, M., Bellavista, P., Lio, P., and K.-W. Lee, K.-W. 2008a. Bio-inspired Multi-Agent Data Harvesting in a Proactive Urban Monitoring Environmen. *Elsevier Ad Hoc Networks Journal, Special Issue on Bio-Inspired Computing and Communication in Wireless Ad Hoc and Sensor Networks*, 7(4), 725–741.

Lee, U., Magistretti, E., Zhou, B., Gerla, M., Bellavista, P., and Corradi, A. 2006a. MobEyes: Smart Mobs for Urban Monitoring with Vehicular Sensor Networks. *IEEE Wireless Communications*, 13(5): 51–57.

Lee, U., Magistretti, E., Zhou B., Gerla, M., Bellavista, P., and Corradi, A. 2008b. Dissemination and Harvesting of Urban Data using Vehicular Sensor Platforms. *IEEE Transaction on Vehicular Technology*.

Lee, U., Park, J.-S., Amir, E., and Gerla, M. 2006b. FleaNet: A Virtual Market Place on Vehicular Networks. *V2VCOM'06*.

Lee, U., Park, J.-S., Yeh, J., Pau, G., and Gerla, M. 2006c. CodeTorrent: Content Distribution Using Network Coding in VANETs. *MobiShare'06*..

Maihöfer, C., Leinmüller, T., and Schoch, E. 2005. Abiding Geocast: Time-Stable Geocast for Ad hoc Networks. *Proceedings of the 2nd ACM International Workshop on Vehicular Ad hoc Networks*, pages 20–29.

Nadeem, T., Dashtinezhad, S., Liao, C., and Iftode, L. 2003. TrafficView: Traffic Data Dissemination Using Car-to-Car Communication. *ACM Mobile Computing and Communications Review (MC2R)*, 8(3): 6–19.

Nandan, A., Das, S., Pau, G., Gerla, M., and M. Y. Sanadidi, M. Y. 2005. Co-operative Downloading in Vehicular Ad-Hoc Wireless Networks. *IEEE/IFIP WONS*.

Nandan, A., Tewari S., Das, S., Pau, G., Gerla, M., and L. Kleinrock. 2006. AdTorrent: Delivering Location Cognizant Advertisements to Car Networks. *IEEE/IFIP WONS, Les Menuires.*

Nath, S., Liu, J., and Zhao, F. 2006. Challenges in Building a Portal for Sensors World-Wide. *ACM WSW*.

Navas, J. C., and T. Imielinski, T. 1997. GeoCast: Geographic Addressing and Routing. *Proceedings of the 3rd Annual ACM/IEEE International Conference on Mobile Computing and Networking*, pages 66–76.

Nosratinia, A., Hunter, T., and Hedayat, A. 2004. Cooperative Communication in Wireless Networks. *Communications Magazine, IEEE*, 42, 74–80.

Ott, M., Seskar, I., Siracusa, R., and Singh, M. 2005. Orbit Testbed Software Architecture: Supporting Experiments as a Service. *Proceedings of IEEE Tridentcom*, pages 136–145.

Park, J.-S., Lee, U., Oh, S. Y., Gerla, M., and Lun, D. 2006. Emergency Related Video Streaming in VANETs Using Network Coding. *ACM VANET'06*.

Peterson, L., Anderson, T., Culler, D., and Roscoe, T. 2002. A Blueprint for Introducing Disruptive Technology into the Internet. *Proceedings of HotNets–I*.

Perkins, C., and Royer, E. 1999. Ad-hoc On-Demand Distance Vector Routing. *Proceedings of the 2nd IEEE Workshop on Mobile Computing Systems and Applications*, pages 90–100.

Ratnasamy, S., Karp, B., Yin, L., Yu, F., Estrin, F., Govindan, R., and Shenker, S. 2002. GHT: A Geographic Hash Table for Data-Centric Storage. *WSNA'02*.

Rheingold, H. 2003. *Smart Mobs: The Next Social Revolution*. Basic Books.

Riva, O., and Borcea, C. 2007. The Urbanet Revolution: Sensor Power to the People! *IEEE Pervasive Computing*, 6(2), 41–49.

Robinson, C., Caveney, D., Caminiti, L., Baliga, G., Laberteaux, K., and Kumar, P. 2007. Efficient Message Composition and Coding for Cooperative Vehicular Safety Applications. *IEEE Transactions on Vehicular Technology*, 56, 3244–3255.

Shah, R. C., Roy, S., Jain, S., and Brunette, W. 2003. Data MULEs: Modeling a Threetier Architecture for Sparse Sensor Networks. *Elsevier Ad Hoc Networks Journal*, 1(2–3): 215–233.

Small, T., and Haas, Z. J. 2003. The Shared Wireless Infostation Model – A New Ad Hoc Networking Paradigm (or Where There Is a Whale, There Is a Way). *ACM MOBIHOC*.

Sormani, D., Turconi, G., Costa, P., Frey, D., Migliavacca, M., and Mottola, L. 2006. Towards Lightweight Information Dissemination in Inter-vehicular Networks. *ACM VANET'06*.

Soroush, H., Banerjee, N., Balasubramanian, A., Corner, M., Levine, B., and Lynn, B. 2009. DOME: A Diverse Outdoor Mobile Testbed. *UMass Technical Report UM-CS-2009-23*.

Standard Specification for Telecommunications and Information Exchange Between Roadside and Vehicle Systems – 5 GHz Band Dedicated Short Range Communications (DSRC) Medium Access Control (MAC) and Physical Layer (PHY) Specifications, September 2003.

Vahdat, A., and Becker, D. 2000. Epidemic Routing for Partially-Connected Ad Hoc Networks. *Technical Report CS-200006, Duke University*.

Wang, Y., and Wu, H. 2006. DFT-MSN: The Delay/Fault-Tolerant Mobile Sensor Network for Pervasive Information Gathering. *INFOCOM'06*.

Wu, H., Fujimoto, R., Guensler, R., and Hunter, M. 2004. MDDV: A Mobility-Entric Data Dissemination Algorithm for Vehicular Networks. *ACM VANET*.

Yin, J., ElBatt, T., Yeung, G., Ryu, B., and Habermas, S. 2004. Performance Evaluation of Safety Applications over DSRC Vehicular Ad Hoc Networks. *ACM VANET'04*.

Yoon, H., Kim, J., Tan, F., and Hsieh, R. 2008. On-demand Video Streaming in Mobile Opportunistic Networks. *PerCom'08*.

Yoon, J., Noble, B., and Liu, M. 2007. Surface Street Traffic Estimation. *MobiSys'07*.

Zang, Y., Stibor, L., Orfanos, G., Guo, S., and Reumerman, H. 2005. An Error Model for Inter-vehicle Communications in Highway Scenarios at 5.9GHz. *Proceedings of the 2nd*

ACM International Workshop on Performance Evaluation of Wireless Ad hoc, Sensor, and Ubiquitous Networks, pages 49–56.

Zang, Y., Stibor, L., Reumerman, H., and Chen, H. 2008. Wireless Local Danger Warning Using Inter-vehicle Communications in Highway Scenarios. *14th European Wireless Conference*, pages 1–7.

Zhou, P., Nadeem, T., Kang, P., Borcea, C., and Iftod, L. 2005. EZCab: A Cab Booking Application Using Short-Range Wireless Communication. *IEEE PerCom'05*.