

A Cyber Physical Systems Perspective on the Real-time and Reliable Dissemination of Information in Intelligent Transportation Systems

Aniruddha Gokhale and Mark P. McDonald

Vanderbilt University

Nashville, TN 37235, USA

Tel: +1-615-322-8754 E-mail: {a.gokhale,mark.p.mcdonald}@vanderbilt.edu

Steven Drager and William McKeever

Air Force Research Laboratory

Rome, NY 13441, USA

Tel: +1-315-330-3045 E-mail: {Steven.Drager,William.McKeever}@rl.af.mil

Received: September 30, 2010 Accepted: October 11, 2010 DOI: 10.5296/npa.v2i3.480

Abstract

Timely and reliable dissemination of traffic-related information to drivers is a key property that intelligent transportation systems (ITS) should support. Numerous impediments stemming due to (a) physical factors, such as mobility and speed of vehicles, density of vehicles, characteristics of the wireless radio channel, and power and bit rate of radio transceivers, and (b) cyber issues, such as MAC layer access point associations and address resolutions (ARP), network layer addressing, routing and handoffs, and transport layer retransmissions lead to unpredictability in the timely and reliable dissemination of information to drivers. This paper presents compelling arguments in favor of new research directions in this area that are based on a cyber-physical systems (CPS) perspective. In particular, this paper makes three contributions. First, it considers a vehicle-centric perspective to survey and study the physics-and cyber-imposed impediments to the timely and reliable dissemination of information. Second, it presents a promising CPS solution to overcome a subset of the impediments discovered. Third, it outlines lessons learned indicating the need for more focused research and realistic testbeds. The evaluations

presented in the paper are based on simulations conducted in the OMNeT++/INETMANET simulator framework for IEEE 802.11 networks. Lack of open ITS testbeds motivated us to choose simulation as an alternative to conduct our study.

Keywords: Cyber physical systems, design and runtime techniques, real-time and reliable information dissemination, intelligent transportation systems.

1. Introduction

Car drivers often experience getting stuck in a traffic jam on their way to work despite having checked the traffic reports prior to leaving their homes. These unanticipated events can have an adverse societal impact, *e.g.*, students are marked as tardy in school and business professionals may miss important meetings or flights. Unanticipated events can also adversely impact mission-critical operations, *e.g.*, medical doctors who may be scheduled to perform a surgery may be delayed or first responders may find it hard to reach the scene of an accident.

Intelligent Transportation Systems (ITS) [19] are envisioned to address the numerous challenges faced by the transportation sector. One category of solutions envisioned in ITS pertains to the real-time and reliable delivery of traffic-related information to drivers both for safety-critical applications (such as blind spot warnings during lane changing) and for applications that improve driving experience and help the environment (such as notification of congestion and rerouting advice that can help to alleviate traffic congestion and lost productivity).

Supporting these applications requires a thorough understanding of the ITS problem space. An ITS comprises three types of communication networks:

- (1) a wireless network formed among the vehicles for vehicle-to-vehicle (V2V) communication,
- (2) a wireless network that involves the vehicles communicating with the road-side infrastructure (V2I), and
- (3) a predominantly wireline network that connects the multiple infrastructure elements.

Real-timeliness and reliability of information dissemination via V2V and V2I communication is a hard problem due to multiple challenges. Some challenges are imposed by the physics of the system including the wireless radio transceiver power, shared nature of the wireless channel, mobility of the vehicles, and density of the vehicles. Other challenges arise from the vagaries of the cyber infrastructure including behavior of protocols like 802.11 media access layer (MAC), address resolution protocol (ARP), IP addressing and routing, and TCP retransmission and congestion control.

Cyber Physical Systems (CPS) [18] provide a natural framework to address these challenges since they integrate the physical, cyber, and human factors in a single framework. This paper makes the following contributions to CPS-based solutions for real-time and reliable dissemination of ITS information:

- It surveys the literature to highlight key insights gained by prior efforts that either pinpoint the impact on timeliness and reliability made by the behavior of the physical

medium or illustrate software-driven solutions that attempt to maintain the quality of service (QoS).

- It highlights an integrated perspective of the physics-and cyber-imposed impediments to realizing the ITS property of real-time and reliable information dissemination based on additional evaluations we conducted.
- It presents a CPS-based solution to overcome some of these impediments. The evaluations presented in the paper are conducted using the OMNeT++ simulator (www.omnetpp.org) and the INETMANET framework developed for OMNeT++.

The rest of the paper is organized as follows: Section 2 discusses related researches that provide key insights into the problem space; Section 3 discusses the impediments in realizing the ITS vision; Section 4 demonstrates how our CPS solution approach overcomes these impediments; and finally Section 5 presents concluding remarks.

2. Related Work

Prior research closely related to the ideas presented in this paper can be classified along two dimensions: (1) those that evaluate the performance of the wireless channel for vehicular communications and (2) those that develop higher-level capabilities (*e.g.*, in middleware) to overcome the challenges imposed by mobility. These prior efforts are important since they provide crucial insights into the behavior and performance of the physical medium, or provide key software-based algorithms that can tune performance.

2.1 Research on Performance Evaluation of Wireless Networks in the Presence of Mobility

The IEEE 802.11 technologies form the basis of most wireless networks in wireless LANs and Wi-Fi, which have become affordable and commoditized. The vehicular networks are also designed to leverage these dominant technologies. In particular, the IEEE 802.11p [12, 11] is an emerging standard for dedicated short range communication (DSRC) used in vehicular networks.

Bai [1] describes real-world experiments involving three vehicles. The authors focus primarily on the reliability aspect of 802.11p DSRC for vehicular safety communications needed in applications, such as forward collision warning or blind spot advising. The authors measured the packet delivery ratio (*i.e.*, the percentage of packets transmitted that were correctly received) and distribution of consecutive packet drops (*i.e.*, the probability distribution of consecutive packet drops). The authors concluded that DSRC does not incur packet loss in bursts and hence the reliability of DSRC is adequate.

Despite their critical findings, the insights in [1] are limited due to measurements with only three vehicles using only V2V communications. Vehicular networks simultaneously must support both the V2V and V2I communication networks, and that too with variation in density of vehicles, speeds, and direction of traffic flow.

Prior studies [2, 8, 16, 17, 20] have evaluated the performance of 802.11p using simulations and analytical models. Eichler [8] conducted OMNeT++-based simulations to measure the collision probability, throughput, and delay for 802.11p. The study determined that 802.11p was unable to handle many high priority messages in a dense network of

vehicles. Similarly, lower priority messages suffer from exponential delays as the density of the network increases.

Bilstrup et. al. [2] study 802.11p MAC layer for supporting real-time applications. In contrast to the findings in [1], they conclude that the channel access (which is a property controlled by the MAC layer) cannot be assured in a manner that supports predictable behavior. As a result, 802.11p is not suitable for low-latency communications.

Similarly, Wang et. al [17] conclude that certain parameter settings in 802.11p MAC can lead to undesirable throughput. For example, the backoff window sizes (which are needed after collisions) are not adaptive to the dynamics of the number of vehicles involved in the communication. This problem stems from the density of vehicles and the percentage of them trying to communicate.

All these findings are crucial since solutions to overcome the inherent limitations are needed. In [16], the authors modify the 802.11p MAC layer to avoid wastage of bandwidth resulting from switching of channels between 802.11p's control channel and service channel. This results in better observed performance at the TCP/UDP level. Similarly, in [4], the authors modify the MAC layer to alleviate the problems of collisions and priority inversions thereby supporting real-time communications.

Although the solutions in our paper also overcome similar challenges, we focus on integrating individual solutions at the both the physical and cyber level.

2.2 Research on Application Layer Software Capabilities

The JECho event infrastructure research [7] argues that a static topology of brokers is insufficient to assure reliable delivery of information when end hosts are mobile. To overcome this problem, the authors present an approach called *opportunistic event channels*, in which brokers are dynamically created or deleted depending on the mobility of the producers and consumers. An algorithm to optimize the event delivery path is presented that considers the network-level handoffs to kickstart an application-level handoff of the producer or consumer from its broker to a new broker that can provide higher assurances of data delivery.

The RT-STEAM [3] project describes a space-elastic model to provide real-time information dissemination. The model is elastic because the range of dissemination can be adapted to continue to meet real-time requirements. The approach is spatial because it relies on geographical proximity of interested subscribers from the producers to adapt the region of dissemination. This approach is geared towards ad hoc communications between vehicles.

The ComPoScan [13] project trades off the available window for data communications with positioning accuracy by switching between active scanning of signal-to-noise ratio versus monitor sniffing, which is supported on majority of 802.11 cards. This work is important since it can form a building block for a more holistic ITS capability.

Despite significant prior work, however, these efforts tend to focus on individual problems addressing them in isolation. Our focus is to address the challenges holistically as a CPS. Our CPS solution relies on sensing physical parameters, such as vehicle speed and signal-to-noise ratio, to trigger the runtime adaptation logic supported by the cyber layer, which in turn controls physical parameters, such as wireless radio power.

3. Impediments to Real-time and Reliable Information Dissemination

Section 2 outlined how prior work has studied individual issues in isolation, such as the challenges stemming from the MAC layer of 802.11p. This section describes experiments conducted to better understand known challenges and uncover hitherto unknown challenges from an integrated perspective of the physical and cyber artifacts of ITS for the following reasons:

1. *Physics of the system:* Since vehicles are mobile, the vehicular networks they create among themselves and/or with the infrastructure elements are *ad hoc*. The temporal (*i.e.*, lifespan of the network) and spatial (*i.e.*, size and topology) properties of such network are determined by multiple factors including the speed, direction, and density of the vehicles, as well as the wireless radio transceiver power, frequency, and signal-to-noise ratio. These factors impact timeliness and reliability in information dissemination.
2. *Vagaries of the cyber infrastructure:* Protocols at the different layers of the network stack (*e.g.*, the five layer TCP/IP stack) significantly impact the timeliness and reliability of information dissemination. For example, the MAC layer requires address resolutions using ARP; the network layer requires addressing and routing to work despite handoffs. Likewise, transport layer protocols (such as TCP) involve sophisticated mechanisms for retransmissions, flow control, and congestion control.

It is important to obtain an integrated perspective of the physical and cyber-imposed challenges so that practical solutions that will realize the vision of ITS can be conceived. To obtain such a perspective we conducted a number of simulations emulating a traffic system. Lack of open ITS testbeds motivated us to choose simulation as an alternative to conduct our study.

3.1 Simulation Framework

The experiments in this paper use OMNeT++ (www.omnetpp.org), which is an extensible, modular, component-based, C++ simulation library and framework, with an Eclipse-based IDE and a graphical as well as command-line runtime environment. Domain-specific functionality (*e.g.*, support for simulation of communication networks, queuing networks, performance evaluation) is provided by model frameworks, developed as independent projects.

OMNeT++ provides component architecture for models. Basic components (called simple modules) are declared in a high level language called the network definition (NED) and its logic is programmed in C++. Individual components can be assembled into larger components and models by declaring them as compound modules in the NED file. Modules can communicate with each other via messages.

OMNeT++ has extensive GUI support, and due to its modular architecture, the simulation kernel (and models) can be embedded easily into applications. Due to its generic architecture OMNeT++ can be used in various problem domains, such as modeling of wired and wireless communication networks, protocol modeling, modeling of multiprocessors and other

distributed hardware systems, validating of hardware architectures, evaluating performance aspects of complex software systems, and any other domain where a discrete event simulation approach is suitable.

For the evaluations in this paper, we leverage the INETMANET framework also available from the OMNeT++ web portal. The INETMANET framework contains support for both wired and wireless networks including IPv4, IPv6, TCP, UDP protocol implementations, and several application models. The framework includes an MPLS model with RSVP-TE and LDP signaling. Link-layer support includes PPP, Ethernet and various flavors of 802.11. Static routing can be set up using network autoconfigurators, or one can use routing protocol implementations.

Mobile ad hoc networking protocols, such as AODV, are also available in addition to many different kinds of mobility models including linear, random, mass and circular mobility.

3.2 Experiment Setup

To simulate a scenario in OMNeT++, developers must describe the modules in a NED file. Configurations that never change can be specified inside the NED files. Any configuration parameters whose values can vary must be specified at initialization-time of the simulation through initialization files.

For experiments that involve mobility, the INETMANET framework's *ChannelControl* module requires the size of the region (called playground) to be specified. We used a playground of size 5,000 meters X 1,000 meters for our experiments. Since vehicles do not move in random directions but follow the road, we limited our simulations to scenarios comprising straight roads, such as those found on highways. This was obtained using the *LinearMobility* module from INETMANET.

The INETMANET framework does not provide an implementation of the IEEE 802.11p protocol since the latter is still not a standard. To mimic 802.11p as closely as possible, we chose the standard 802.11a implementation with a 5.9 GHz frequency range and a channel bit rate of 6 Mbps. The radio transceiver's sensitivity was set to -85dBm while the thermal noise threshold was kept at -110dBm. These parameters help the simulator to distinguish a valid packet reception from noise. The signal to noise ratio threshold was set to 4 dB.

In INETMANET the computation of the distance for wireless signal propagation and power of received signal is based on the Friis transmission equation [9]. Since this technique works in idealized conditions, the simulation results may deviate somewhat from reality. Since we are interested in understanding the general trends and challenges, however, this limitation is acceptable for our work.

We narrowed the scope of our simulations to V2I communications where infrastructure elements, such as road side units (RSUs), mediate communication between vehicles or vehicles to internet-based services. Our motivation to restrict the study stems from the fact that the ITS problem space is very large, which argues for narrowing the scope of the integrated cyber-and physics-imposed challenges and studying these challenges one at a time.

To understand the combined impact of both the cyber-and physics-imposed challenges, we developed a series of simulation scenarios that help to elicit these challenges. All our simulations assume straight roads without any obstacles, such as buildings, hills and trees.

Naturally, more studies are needed without making these assumptions, which is the focus of our future work.

3.3 Impediments to real-time and reliable data dissemination due to physics of the system

In 802.11 networks a node acquires network connectivity only when it gets associated with an access point (AP). In ITS, vehicles involved in V2I communications must be associated with a RSU, which serves as the AP. Association with an AP in 802.11 networks involves complex activities. At the physical layer, the radio transceiver on the mobile node (in our case the vehicle), which operates in the 802.11 STA mode (*i.e.*, station mode), can choose to actively scan the radio channel to discover an AP. At the same time, any 802.11 device acting in the AP mode (*i.e.*, access point) periodically broadcasts a beacon frame which can be received by any node in the radio range that is tuned to the same channel.

Whenever either the mobile node or the AP discovers the other entity, complex MAC layer signaling takes place that first involves an authentication request (AUTH) by the node to the AP followed by an acknowledgment from the AP, only if the security credentials are acceptable. This is followed by an association request (ASSOC) by the node and a grant message by the AP. Once the node is associated with the AP, the node stores the AP's MAC address for any subsequent V2I communications.

Since the node is mobile, with passage of time the received power at the node's radio transceiver will eventually fall below the sensitivity value at which point the RSU is considered out-of-range. Our goal is to identify the impact of vehicular speed, density of vehicles, and the power of the wireless radio – all physical factors – on the time it takes to discover the RSU, associate with it, and ultimately lose contact with it. Determining these times are important since they determine the window of opportunity for the middleware and application layer – all cyber artifacts – to schedule data communication operations, which are key to realizing the different kinds of ITS applications.

3.3.1 Scenario 1: Impact of vehicle speed on RSU association

The first experiment comprises a single vehicle whose speed can be configured per iteration of the experiment. This special case of a single vehicle is used to understand the impact of speed without the impact of collisions arising from multiple vehicles contending for the shared channel.

Figure 1 shows our simulation setup. A RSU (*i.e.*, AP) is placed on the playground at coordinates (500, 100). A single vehicle moves east (*i.e.*, from left to right) starting from coordinates (200, 105). The speed of the vehicle is varied between 15 MPH (miles per hour) and 75 MPH. The higher speeds represent typical speed limits on interstates in USA. The lower speeds represent speeds during rush hour traffic or speeds on entry and exit ramps.

Two iterations of the experiment were conducted for each speed value: once using active scanning by the vehicle for a RSU and once relying solely on beaconing from the RSU. The wireless radio transceiver power for both the RSU and vehicle is set to 5mW.




Figure 1: Experimental setup for discovering the road-side unit

Figure 2 shows the time taken for the vehicle from its starting point to discover the RSU, authenticate with it, associate with it, and ultimately lose contact with it for the active scanning mechanisms. This figure confirms our intuition that higher speeds will result in faster discovery of the RSU but also a rapid disconnection from the RSU. The beacon-only approach reveals negligible difference in performance compared to the active scan approach. The latter results are not shown due to lack of space. Note also that the RSU discovery takes time in the order of a few seconds which is detrimental to response times for applications.




Figure 2: RSU relationship for single vehicle with varying speed and active scanning

We argue that with increasing number of vehicles on the road, the bandwidth of the channel will have to be shared between both the scan requests (which each vehicle transmits) and beacons (which are broadcasted by the RSU). Hence, in the rest of the experiments we use only the beaconing approach. Note that the 802.11p standard proposes to drop the association process altogether but instead relies on wildcard AP addresses. This approach will require active scanning by the node but no beaconing by the AP.

3.3.2 Scenario 2: Impact of radio power on RSU association

The second experiment maintains a constant speed of 60 MPH for the vehicle but varies the power level of the wireless radio transceivers on the vehicle and the RSU. Our aim is to determine the impact of power levels on the time window available for data communication via the RSU. The power level used in individual iteration of the experiment was varied between 5mW and 50mW. The starting coordinates of the vehicle was the same as before, however, the RSU is now placed at coordinates (1500, 100) to ensure that the vehicle is out of range of the RSU at the start of the simulation even for the higher power levels.

Figure 3 depicts the impact of different power levels on the time it takes to discover the RSU, authenticate with it, associate with it, and ultimately lose connection with it. The results are in tune with our intuition that the higher the power level, the earlier is the RSU discovered and later is the connection lost.

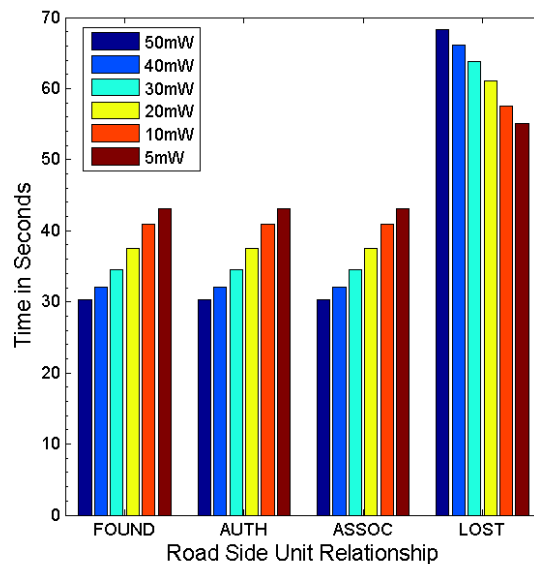


Figure 3: RSU relationship for single vehicle with varying radio power

These results are important since they provide insights into how higher layer cyber solutions can make runtime adaptations to the radio power to improve reliability and timeliness of information dissemination.

3.3.3 Scenario 3: Impact of vehicular density on RSU association

In the third experiment we created a scenario representing two-way traffic. The highway comprises two lanes each in both directions. The RSU is assumed to be placed on the median. We varied the number of vehicles heading east-to-west and west-to-east so that platoons of vehicles are formed in both directions. The platoon size was varied from one vehicle in each platoon to 16 vehicles in each platoon for a total of 32 vehicles.

Figure 4 shows our experimental setup with one vehicle each on either side. The speeds of all vehicles are maintained at 60 MPH. Although it is unrealistic to expect all vehicles in a

platoon to travel at exactly the same speeds, we chose this scenario since it reflects a worst case scenario where due to the same speed, each vehicle will demonstrate the same pattern for association with the RSU. We surmise that this situation will result in higher number of collisions in the authentication and association messages (and data messages) leading to unpredictability in the association time (and application-level data throughput).




Figure 4: Experimental setup for two way traffic with platoons

Two vehicles in neighboring lanes are separated by 5 meters. Thus, the lead vehicle moving east in the inner lane starts at coordinates (200, 105) while the one in the outer lane starts at (200, 110). Each following vehicle is separated from the vehicles in the front also by 5 meters. An exactly symmetric arrangement exists for the westward moving traffic.

Figure 5 shows the variation in the time it takes for the lead vehicle moving east to discover the RSU, authenticate with it, associate with it, and then disconnect from it as the size of the platoon increases from 2 to 32 (1 to 16 in each direction). The jitter in the times shown is the result of contention for the shared channel and resulting collisions.




Figure 5: RSU relationship for platoons of vehicles with varying platoon sizes

3.4 Impediments to real-time and reliable data dissemination due to vagaries of higher layer cyber space protocols

To realize any useful application in ITS, simply discovering a RSU is not enough. In other

words, understanding only the impact of physics on physical layer protocols is not sufficient. Cyber-level algorithms often control physical elements to provide useful capabilities of societal importance. An important requirement is application-specific data that must be communicated between vehicles and other network-based services. For example, real-time traffic updates need to flow from servers based in the wired Internet to the vehicles. On the other hand, applications, such as periodic transmission of health status of a vehicle, must send data from a vehicle to a network-based server. All such communication is mediated through the RSUs.

The ITS applications and their QoS will require leveraging traditional cyber infrastructure including middleware; network protocols, such as TCP and UDP for transport-level end-to-end properties; IP for addressing and routing; and 802.11 MAC for transmitting the physical frames over the wireless channel.

The physics of the system (*e.g.*, speed, density, radio power) impact the performance of traditional cyber technologies, such as TCP, IP and 802.11 MAC. For example, a disconnection from a RSU due to mobility creates multiple problems at every higher layer protocol of the cyber infrastructure. TCP will experience packet loss and resort to retransmissions. It is imperative to study the effects of such retransmissions on application-level throughput and message latencies, particularly, when the density of vehicles in a region increases.

At the IP level, addressing and routing are key issues that must be handled when a vehicle is handed over from one RSU to the other. Handoffs may impact the timeliness and reliability of information dissemination. It is necessary to understand the impact of handoffs and determine how promising techniques, such as routing in delay tolerant networks [10], may need to be enhanced for ITS.

3.4.1 Scenario 1: Impact on application-level communication due to interplay between speed-imposed communication window and transport-layer connection establishment

In this experiment we measured the request-response latencies between a vehicle and a server. Figure 6 shows our experimental setup. A wired network-based server is shown directly attached to the RSU's LAN-based interface. The RSU simply relays packets between the wireless and wired network. The server is directly attached to the RSU to eliminate additional impact in our measurements due to routing and switching.




Figure 6: Experimental setup for data communication mediated by a RSU

In the experiment, the vehicle opens up a TCP connection with the server. Every request is 1 Kbytes while the response is 5 Kbytes. A request will fit in a single MAC-layer (whose maximum transfer unit is approximately 1,500 bytes) frame while multiple frames will be needed for the response. Trying to open up a connection to the server even before associating with a RSU is useless. Hence, using the timings we measured in our earlier experiment (see Section 3.3.1) we ensured that the application opens a connection to the server right after the RSU is associated.

We observed that the TCP SYN connection establishment segment triggers an ARP request as shown in Figure 7. Even with a UDP communication, an ARP request must be made. ARP is needed by the vehicle because it needs to determine the destination MAC address to which it must relay the message in order to reach the network-based server. Naturally, the ARP overhead further reduces the time window of communication with RSU available to a vehicle.




Figure 7: Address resolutions initiated by TCP packets

Figure 8 shows the number of packets received successfully at the application level by the vehicle for different speeds. Since it is already known that higher speeds will have a smaller window of time for communication, the number of packets received will also be smaller. Note that these measurements include TCP-level effects, such as congestion control and retransmissions if at all any of these were experienced during the experiments.




Figure 8: TCP-level packets delivered for different speeds

These results are important since higher level cyber artifacts, such as middleware and application level protocols, as well as tools for real-time scheduling can determine how best to schedule data transmissions and what message sizes are appropriate to ensure retrieval of information in the available time window.

3.4.2 Scenario 2: Impact on application-level communications due to interplay between packet collisions caused by dense vehicular traffic and transport-level retransmissions

We extended the experimental setup shown in Figure 6 by including the two-way traffic and platoon of vehicles described in Section 3.3.3. The goal is to estimate the impact of collisions on application-level data throughput. For this experiment we kept the speed of all vehicles constant at 60mph.

Figure 9 shows the minimum, average, and maximum number of packets received by any vehicle in the platoon. With increasing density of vehicles, the number of packets received starts decreasing. After about 16 vehicles overall in the platoon (8 on each side), some vehicles do not even receive any packet.




Figure 9: Packets delivered with increasing platoon size

The behavior in Figure 9 can be attributed to increasing number of MAC layer collisions and resulting retransmissions by TCP, as shown in Figure 10. The figure shows that the data throughput is acceptable until about 14 vehicles in the platoon since the number of collisions and jitter remains low until that point.




Figure 10: MAC-level collisions with increasing platoon size

3.4.3 Scenario 3: Impact on application-level data communication due to interplay between RSU-handoffs and network-layer routing

In our final experiment we introduced multiple RSUs so that the vehicle can be handed over from one RSU to other as the vehicle moves out of range of one RSU and into that of the other. Figure 11 shows the experimental setup. In this experiment we moved the server deep

inside the network so that it can be reached to the vehicle only through a router, an Ethernet switch and a RSU. This arrangement enables us to determine the impact of routing in addition to all the earlier studied impediments.




Figure 11: Experimental setup for RSU handoff

In addition to all the other impediments uncovered earlier, in this experiment we faced an additional challenge that adversely impacts the timeliness and reliability of information dissemination. The router that routes packets from the server to the vehicle consults its routing table entries. When the vehicle is associated with RSU A (shown as *ap A*), the router forwards the packet on the interface towards the Ethernet switch 1. When the vehicle moves into the range of *ap C*, the routing table entry is no longer valid. Similar issues are faced at the Ethernet switches, whose tables also become stale as the vehicle moves out of range.

Traditional approaches based on Dijkstra or Bellman-Ford algorithms used in the Internet to update routing tables at routers are based on the assumption that routing updates are infrequent. In ITS vehicles move at high speeds which makes it infeasible to rely on traditional algorithms. Solutions relying on an RSU assigning an IP address using DHCP will incur network address translations at the RSU imposing additional overhead. Moreover, it does not address the routing table problem since the router must still route a packet to a different RSU when handoffs occur.

Solutions relying on Mobile IP [14], though promising, are not directly applicable since Mobile IP relies on a home agent and a foreign agent. Every vehicle will need to maintain a home agent with some RSU, which may be hard and impose an extra overhead. Moreover, due to high speeds, the foreign agents will change frequently requiring significant overhead of maintaining the mapping from home agent to foreign agent, and redirecting requests.

4. A Proposed CPS solution for reliable and timely V2I communications

Sections 3.3 and 3.4 described a range of impediments to real-time and reliable data dissemination caused due to physical- and cyber-level issues. This section describes a proposed solution to address the impediments to real-time and reliable information dissemination that focuses on ITS V2I networks. Our CPS solution includes a tight integration of cyber artifacts (such as middleware) with the physical elements of ITS. In particular the cyber artifacts provide runtime adaptations that are effected by sensing of physical parameters (such as speed, radio power, and signal-to-noise ratio) and their projected impact on the timeliness of data. These projections are obtained from the experimental results

of Sections 3.3 and 3.4. The details of our CPS solution are discussed in the remainder of this section.

4.1 Overcoming RSU association overhead

Our solution to overcome the RSU association overhead is based on Service Oriented Architecture (SOA) that comprises a web service that provides a map of RSUs that a vehicle is expected to encounter along the route it is going to travel. The client for such a service could very well be an add-on to an existing navigation system built in the vehicle or a portable GPS system. In either case, such a navigation system not only computes the effective route but also retrieves a map of RSUs. RSUs can be categorized as points-of-interest (POIs) inside the navigation systems.

For an 802.11x based network, the retrieved map of RSUs comprise their MAC addresses and GPS locations. The MAC addresses are necessary so that the vehicle can eliminate the process of RSU association. For improving reliability, alternate RSUs can be simultaneously used through a technique known as beacon stuffing [6] but our solution currently does not utilize this method. The GPS coordinates provide the RSU location, which can be used by a vehicle to determine when to proactively switch from one RSU to another as the vehicle moves out of range of one RSU.

In our simulation-based solution, such a map is provided to a vehicle as part of the configuration file used at simulation initialization-time.

4.2 Overcoming speed-imposed impediments

Having an RSU map does not mean that a vehicle will automatically be in range with one of the RSUs. The speed of the vehicle is one factor that determines the time window available for communication. Scheduling data communications only during this window is necessary since otherwise mechanisms for flow control and congestion control may significantly impact performance.

To enable a vehicle to determine when to schedule data communications, cyber artifacts (such as middleware) must be aware of when the vehicle gets in the range of the RSU and how much time window is available for the current vehicle speed. When the signal-to-noise ratio increases beyond a threshold, the wireless radio in a vehicle gets in range of a RSU. Data from our experiments for signal-to-noise ratio for different speeds but the same radio power is shown in Figure 12. The figure reveals that a ratio above 25 made the vehicle get in range.




Figure 12: Variation in signal-to-noise ratio with changing speeds

Note how the signal-to-noise ratio spectrum shifts in time for different speeds. Although for this paper we did not consider vehicle-to-vehicle (V2V) communications, it is possible to leverage the phase shifts to switch between V2V and V2I communications to improve reliability of communication. For example, a slow moving vehicle that may lose connectivity with a RSU may offload some communications to a faster moving vehicle that may have acquired connectivity with the next RSU. This dimension of work will be the focus of our future work.

Next, the vehicle must determine the time window it has to communicate. A mathematical relationship between the vehicle speed and the time window can be determined from experimental data shown in Figure 2 and/or mathematical analysis [15]. This relation can be used not only for discrete speed values but even when the vehicle is accelerating or decelerating.

Finally, if the vehicle is hosting multiple applications, the cyber artifacts must be aware of the total data that can be received within the time window, which can be obtained from Figure 8 so that the cyber artifacts can prioritize and control the flow of traffic from different applications [5].

Sensing of physical factors and relaying them to higher level middleware artifacts will require cross-layer communications and optimizations. Our simulation leverages a module called the *NotificationBoard* that serves as an event-channel to relay events of interest to subscribers.

4.3 Overcoming impediments due to radio power

Deployment of RSUs is optimized for the speed limit for that road. When vehicles move faster than their speed limits but maintain high radio power levels, there is every possibility that their radio range will overlap multiple RSUs. On the other hand, slow moving vehicles will encounter regions of no connectivity.

Figure 13 shows how the time window for communication with RSU can be adjusted

by varying the power level of the radio, which is an example of how cyber artifacts can control physical factors. The implementation techniques are similar to the ones used to overcome the impediments imposed by speed. This paper, however, does not integrate the control of speed and power simultaneously, which is the focus of our future work.




Figure 13: Variation in signal-to-noise ratio with changing radio power

To handle this problem, higher layer cyber artifacts within a vehicle can dynamically adjust their radio power levels so that connectivity with a RSU can be maximally maintained without overlaps.

4.4 Overcoming impediments imposed by higher layers

Section 3.4.3 revealed how routing tables must be updated when handoffs occur. Note that routes must be changed only at those switches and routers that are closest to the RSUs. Since the next RSU along the route is stored in the map at the vehicle, the cyber layer can send a message to the closest router or switch with the address of the next RSU so that the routing table entry can be updated.

It is conceivable that the system may get overloaded with such messages if every vehicle attempts to send route update messages. Although we do not address this problem in this paper, a potential solution is to leverage the platoon effect and elect one vehicle as a leader who alone uses V2I communications while other vehicles use V2V communication.

Our simulation-based implementation leverages the *NotificationBoard* module to sense the signal-to-noise ratio values. When the value is close to falling below the sensitivity threshold, an event is scheduled for the router to trigger a table update.

5. Conclusion

This paper presented the causes for impediments to real-time and reliable information dissemination in vehicle-to-infrastructure (V2I) communication networks for intelligent

transportation systems. The primary source of the impediments was a combination of physical (*e.g.*, vehicular speed, radio range) and cyber (*e.g.*, behavior of protocols such as TCP, IP and 802.11 MAC) factors. The paper then outlined a CPS-based solution to overcome the uncovered challenges. Simulation was used to conduct the study and implement the solution, which is available at www.dre.vanderbilt.edu/~gokhale/ITS.zip.

Below we mention some important lessons learned during our research.

- The work presented focuses only on the V2I aspect of ITS. A general ITS solution should comprise both V2V and V2I, which is the focus of our future work.
- The experiments performed in this paper made many assumptions on the terrain and road types. In future we will perform additional experiments that removes several of the assumptions made in this paper.
- Real-world vehicular traffic traces are needed to obtain accurate behavior, however, simulators like OMNeT++ are not geared for traffic engineering. Our current work is using the SUMO traffic simulator and using its TRaCI facility to create a bridge between the SUMO and OMNeT++ simulators. We are also exploring the use of an emerging simulator called NCTuNS.
- Real testbeds for ITS are needed for real data and test solutions but such testbeds are not readily available. We are setting up a miniature testbed using a combination of wireless routers serving as RSUs and Netbooks serving as the in-vehicle computers.

Acknowledgment

This research was made possible in part by grants from the US Air Force Research Laboratories Visiting Summer Faculty Research Program, US National Science Foundation RAPID Award # 1047792, and Vanderbilt University Discovery Grant. We would also like to thank the anonymous reviewers for their excellent feedback that helped to improve the contents of this paper.

References

- [1] F. Bai and H. Krishnan. Reliability Analysis of DSRC Wireless Communication for Vehicle Safety Applications. In *IEEE Intelligent Transportation Systems Conference (ITSC'06)*, pages 355–362, Toronto, Canada, Sept. 2006.
- [2] K. Bilstrup, E. Uhlemann, E. G. Ström, and U. Bilstrup. Evaluation of the IEEE 802.11p MAC method for Vehicle-to-Vehicle Communication. In *IEEE 68th Vehicular Technology Conference (VTC '08-Fall)*, pages 1–5, Calgary, Alberta, Canada, Sept. 2008.
- [3] M. Bourouche, B. Hughes, and V. Cahill. Building Reliable Mobile Applications with Space-Elastic Adaptation. In *Proceedings of the 2006 International Symposium on World of Wireless, Mobile and Multimedia Networks (WOWMOM '06)*, pages 627–632, Niagara Falls, NY, USA, 2006. IEEE Computer Society.
- [4] B. D. Bui, R. Pellizzoni, M. Caccamo, C. F. Cheah, and A. Tzakis. Soft Real-Time Chains for Multi-Hop Wireless Ad-Hoc Networks. In *13th IEEE Real Time and Embedded Technology and Applications Symposium (RTAS'07)*, pages 69–80, Seattle, WA, 2007.

- [5] J. Burgess, B. Gallagher, D. Jensen, and B. N. Levine. MaxProp: Routing for Vehicle-Based Disruption-Tolerant Networks. In *Proceedings of the 25th IEEE International Conference on Computer Communications (Infocom '06)*, pages 1–11, Barcelona, Spain, 2006. IEEE Comm Soc.
- [6] R. Chandra, J. Padhye, L. Ravindranath, and A. Wolman. Beacon-Stuffing: Wi-Fi without Associations. In *Proceedings of the Eighth IEEE Workshop on Mobile Computing Systems and Applications (HOTMOBILE '07)*, pages 53–57, Tucson, AZ, USA, 2007. IEEE Computer Society.
- [7] Y. Chen, K. Schwan, and D. Zhou. Opportunistic Channels: Mobility-aware Event Delivery. In *Proceedings of the ACM/IFIP/USENIX 2003 International Conference on Middleware (Middleware '03)*, pages 182–201, Rio de Janeiro, Brazil, 2003. Springer-Verlag New York, Inc.
- [8] S. Eichler. Performance Evaluation of the IEEE 802.11 p WAVE Communication Standard. In *IEEE 66th Vehicular Technology Conference (VTC Fall '07)*, pages 2199–2203, Baltimore, MD, Oct. 2007.
- [9] H. T. Friis. A Note on a Simple Transmission Formula. *proc. IRE*, 34(5):254–256, 1946.
- [10] S. Jain, K. Fall, and R. Patra. Routing in a Delay Tolerant Network. In *Proceedings of the 2004 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM '04)*, pages 145–158, Portland, Oregon, USA, 2004. ACM.
- [11] D. Jiang and L. Delgrossi. IEEE 802.11p: Towards an International Standard for Wireless Access in Vehicular Environments. In *IEEE Vehicular Technology Conference (VTC '08)*, pages 2036–2040, May 2008.
- [12] D. Jiang, V. Taliwal, A. Meier, W. Holfelder, and R. Herrtwich. Design of 5.9 Ghz DSRC-based Vehicular Safety Communication. *IEEE Wireless Communications*, 13(5):36–43, Oct. 2006.
- [13] T. King and M. B. Kj'aergaard. ComPoScan: Adaptive Scanning for Efficient Concurrent Communications and Positioning with 802.11. In *6th International ACM Conference on Mobile Systems, applications, and services (MobiSys '08)*, pages 67–80, Breckenridge, CO, USA, 2008. ACM.
- [14] C. E. Perkins. Mobile IP. *IEEE Communications*, 35(5):84–99, 1997.
- [15] M. Rudack, M. Meincke, and M. Lott. On the Dyanamics of Ad Hoc Networks for Inter Vehicle Communications (IVC). In *Proceedings of the International Conference on Wireless Networks (ICWN '02)*, Las Vegas, NV, USA, 2002.
- [16] S.-Y. Wang, H.-L. Chao, K.-C. Liu, T.-W. He, C.-C. Lin, and C.-L. Chou. Evaluating and Improving the TCP/UDP Performances of IEEE 802.11(p)/1609 Networks. In *IEEE Symposium on Computers and Communications (ISCC '08)*, pages 163–168, Marrakech, Morocco, July 2008.
- [17] Y. Wang, A. Ahmed, B. Krishnamachari, and K. Psounis. IEEE 802.11p Performance Evaluation and Protocol Enhancement. In *IEEE International Conference on Vehicular Electronics and Safety (ICVES '08)*, pages 317–322, Columbus, OH, Sept. 2008.
- [18] W. Wayne. Cyber-Physical Systems. *IEEE Computer*, 42(3):88–89, 2009.
- [19] R. J. Weiland and L. B. Purser. Intelligent Transportation Systems. *Transportation*

Research, 1:40AM, 2009.

- [20] M. Wellens, B. Westphal, and P. Mahonen. Performance Evaluation of IEEE 802.11-based WLANs in Vehicular Scenarios. In *IEEE 65th Vehicular Technology Conference (VTC '07-Spring)*, pages 1167–1171, Dublin, Ireland, Apr. 2007.

Copyright Disclaimer

Copyright reserved by the author(s).

This article is an open-access article distributed under the terms and conditions of the Creative Commons Attribution license (<http://creativecommons.org/licenses/by/3.0/>).